INTERO: A Model of Robotic Interoceptive Sensing

Amber Maimon[1,2,*,†], Iddo Yehoshua Wald[1,*,†], Mihai Pomarlan[1,*,†], Shiyao Zhang[1], Daniel Beßler[1], Robin Nolte[1], Dennis Küster[1], Robert Porzel[1], Tanja Schultz[1] and Rainer Malaka[1]

[1]*University of Bremen, Bibliothekstr. 5, 28359 Bremen, Germany*
[2]*University of Haifa, 199 Abba Khoushy Ave., 3103301 Haifa, Israel*

Abstract

Interoception is the ability of an organism to sense its internal state. It plays a central role beyond physiological self-regulation, underlying processes such as decision making, self-awareness, motivation and adaptability, as well as emotional experience and social interaction. We consider here an analogous capability for artificial agents and present the first steps towards a formal description of what it means for a signal to be an interoceptive signal for some artificial system. This formal description is in the form of an ontology of interoceptive sensing. We present some competency questions that such an ontology must address. Our presented questions focus on the distinction between inner and outer signals for a system, and the relevance of inner signals for self-regulation (also referred to as homeostasis). We have started building this ontology by combining the sensing and actuation ontologies SOSA and SSN into the robotics ontology SOMA, and intend to use the systemic view being developed by Mizoguchi and Borgo as a guide to our approach to the formalization of systems capable of interoception.

Keywords

Interoceptive signals, Interoception, Homeostatic control, Embodied cognition, Physiological signals

1. Introduction

In humans, interoception is the continuous monitoring of bodily signals, including one's heartbeat, breathing, body temperature, hunger, and other visceral sensations [1, 2]. It underpins homeostasis [3] and plays a central role in emotion, self-awareness, and decision-making [2, 4, 5, 3]. It is through a constant integration of intero- and exteroceptive sensations that our brains shape what we perceive and feel, what motivates us, and how we develop and experience a sense of self [6, 7, 8, 9].

Unlike humans, robotic agents lack mental states, yet they employ mechanisms reminiscent of internal states in the sense that they are insulated from direct interference from the outside world but which may affect behavior and actions. A robot's representation of physical states such as battery level, temperatures, power consumption, and CPU load, could be referred to as the robot's information about its own bodily states. These signals fluctuate dynamically during task execution and could offer insights into the "inner state" of robotic agents. Thus, while a robot cannot be aware of its bodily states in the sense of human interoception, we propose considering these types of information as the closest available robotic equivalents to biological signals such as heart rate, respiration, or metabolic energy

RobOntics'25: International Workshop on Ontologies for Autonomous Robotics, December 10, 2025
*Corresponding author.
†These authors contributed equally.
maimon.amber@gmail.com (A. Maimon); wald@uni-bremen.de (I. Y. Wald); pomarlan@uni-bremen.de (M. Pomarlan); shiyao@uni-bremen.de (S. Zhang); danielb@uni-bremen.de (D. Beßler); nolte@uni-bremen.de (R. Nolte); kuester@uni-bremen.de (D. Küster); porzel@uni-bremen.de (R. Porzel); tanja.schultz@uni-bremen.de (T. Schultz); malaka@tzi.de (R. Malaka)
https://www.ambermaimon.com/ (A. Maimon); https://iddowald.com (I. Y. Wald); https://www.uni-bremen.de/dmlab/team/shiyao-zhang-izzie (S. Zhang); https://www.uni-bremen.de/dmlab/team/robin-nolte (R. Nolte); https://www.uni-bremen.de/csl/institut/team/mitarbeitende/dr-dennis-kuester (D. Küster); https://www.uni-bremen.de/dmlab/team/dr-ing-robert-porzel (R. Porzel); https://www.uni-bremen.de/csl/institut/direktorin (T. Schultz); https://www.uni-bremen.de/dmlab/team/rainer-malaka (R. Malaka)
0000-0002-4592-0711 (A. Maimon); 0000-0001-7998-0133 (I. Y. Wald); 0000-0002-1304-581X (M. Pomarlan); 0000-0001-5965-0428 (S. Zhang); 0000-0002-9876-5743 (D. Beßler); 0009-0004-2975-6378 (R. Nolte); 0000-0001-8992-5648 (D. Küster); 0000-0002-7686-2921 (R. Porzel); 0000-0002-9809-7028 (T. Schultz); 0000-0001-6463-4828 (R. Malaka)

in living organisms. By doing so, they form part of the robot's embodied self-model, enabling it to reason not only about what it perceives and knows about the surrounding world but also about its own physical states.

The analog, for a robot, of human interoception, would enable it to connect its internal states with cognition, planning, and control. This would enable homeostatic regulation by adjusting movement, timing, or sensory precision in response to internal strain or limited energy, and support anticipatory control by predicting and mitigating such changes before they impair performance. This predictive regulation parallels homeostatic reinforcement learning, where actions are guided by minimizing internal error rather than by external reward.

By making these internal variables explicit in the reasoning architecture, artificial interoception provides a mechanism for adaptive autonomy. Robots can dynamically reorganize priorities, alter strategies, or delay non-essential tasks when internal states exceed thresholds. Over time, this could support the emergence of embodied self-models, where decision-making reflects both task goals and the system's physical and computational viability. Artificial interoception thus forms a foundation for stable, resource-aware, and resilient robotic cognition. Artificial interoception could be relevant for any artificial agents, and particularly to embodied ones such as robots. It adds a missing "physiological" dimension, allowing robots to link knowledge about their actions and goals (form an introspective model) with knowledge about their internal physical condition (interoceptive state).

In this work we seek to establish the representational groundwork for explicating these interoceptive signals and states via an ontological model of inner signals and their connection to subsystems of homeostatic control. Based on such a model robots can reason about their internal states and communicate them in various ways. The main contribution of this work, therefore, lies in connecting existing models of sensors and sensing to their potential significance for a robotic agent's internal state.

Our long-term objective is to define a formal structure that allows artificial agents to reason about their physical internal states as part of their (meta)cognitive processes, and in this paper we will present some initial steps in this direction. We will build on existing ontologies of robotics – SOMA [10] and its extension MOI [11] –, sensor ontologies SSN [12] and SOSA [13], and the work on formalizing systems done by Mizoguchi and Borgo [14]. By enhancing the MOI ontology with interoceptive categories, we lay the groundwork for robotic systems capable of self-regulation, adaptive decision-making, and resource-aware planning, similar to human embodied cognition.

2. State of the Art

2.1. Embodied Cognition and Interoception

There is a growing consensus that human cognition is embodied [8] and dependent upon a crucial balance between interoception (awareness of internal bodily signals) and exteroception (perception of external signals) [1]. Interoception has been argued to play a crucial role in cognitive systems that underlie human abilities such as self-awareness, decision-making [15], and emotional regulation [2, 4, 5, 3] alongside being important for homeostasis [3].

Embodied cognition emphasizes the interplay between bodily states and cognitive processes [6, 7, 8, 9]. It posits that cognition cannot be fully understood without considering the influence of the body and its interactions with the environment. This concept is well exemplified in foundational work, which demonstrated how embodied agents can exhibit adaptive and intelligent behavior through situated interactions with their environment [16]. Moreover, recent advances in active inference and predictive processing have refined these models. These theories propose that interoception is not merely a passive process but an active regulatory mechanism for bodily states [17]. This framework, known as interoceptive inference, suggests that the brain continuously generates predictions about internal bodily states to minimize physiological uncertainty.

2.2. Artificial Cognition and Homeostatic Regulation

We adopt the definition of Artificial Cognition (ACo), which has previously been suggested as a more general term for artificial agents with embodied cognition, and which aims to avoid the trap of mind-body dualism in robotics [18]. Some recent work has begun to explore the concept of interoception in artificial systems. For example, recent work has proposed mathematical models of interoception for artificial cognitive systems [19], while other work introduced the concept of Embodied Neural Homeostat—an agent that self-regulates based on internal energy levels, demonstrating adaptive behaviors [20]. Similarly, internal state variables for autonomous robots have been proposed to enable self-regulation in unpredictable environments [21]. Our project will build on these approaches to develop artificial systems equipped with interoceptive analogues, exploring their adaptability and resilience.

The lack of interoceptive mechanisms in most current artificial cognitive systems possibly results in rigid adaptability and poor uncertainty management [19, 21]. Most existing models typically rely solely on external sensory input, making them unable to self-regulate or adapt their behavior based on internal states. Integrating interoception into virtual agents could enable agents that set their own goals based on internal needs, much as animals do.

In reinforcement learning research, the idea of homeostatic reinforcement learning (HRL) has been introduced: instead of using a fixed external reward, the agent's reward signal is derived from internal state errors (deviations from a desired setpoint). This approach explicitly links classical drive-reduction theories from psychology (e.g. hunger reduction) with RL algorithms [22], allowing agents to learn behaviors that restore internal balance. Initial studies in simulated agents have shown that such systems can learn survival-enhancing behaviors (like seeking "food" when energy is low) purely from the imperative to maintain internal variables within optimal ranges [22]. These developments point to a novel frontier: simulating interoception in non-biological systems as a route to more lifelike, context-sensitive artificial cognitive systems.

3. The INTERO Ontology

We propose expanding the existing MOI ontology (Metacognition Ontology Infrastructure) [11] originally developed to capture self-monitoring and introspection in cognitive architectures by introducing the Interoception for Robots Ontology (INTERO) - a formal model of interoceptive signals and their impact in a system's operation.

It is first worth considering what we might want to regard as relevant "interoceptive signals" and their computational analogues in this context. In our view, an ontology of interoceptive signals for artificial agents would include things such as:

- **Physical signals**: Signals that represent the state of a physical system, measured by sensors. They reflect the system's physical condition or strain. E.g. temperature measurements, current draws, motor strain.
- **Computational signals**: Signals that reveal aspects of a computational process, yet are likewise captured through physical measurement. They are often expressed as temporal or performance indicators. E.g. inference duration, data access latency, decision time.

For the purposes of this ontology, interoceptive signals will be limited to analogs of physical/computational signals in artificial agents that are observable, even if with some degree of error, by some kind of sensor. This raises the question of what distinguishes interoceptive signals from other types of signals. Previous literature suggests three aspects that are crucial to making sense of interoception [19]:

- **Factorization**: it is possible to partition the variables characterizing an environment into those belonging to the "inside" of an agent and to the agent's "outside", with conditional independence between the variable sets when given information about a "boundary" between the agent and the environment. This aspect is fundamental to making interoception a meaningful notion, as without a clear separation between inside and outside there is no basis for speaking of interoception as a distinct kind of perception;

- **Homeostasis**: control loops inside the agent make use of interoceptive signals to gauge how far away the agent is from some desired state for its inner variables, such that these control loops will act towards restoring the desired state;
- **Reward**: in cases where the agent modifies its control loops (e.g. via learning), interoceptive signals offer a reward signal to guide the learning process.

We will focus on the first two aspects – factorization and homeostasis – in this paper, and leave reward for future work. We will first provide an overview of the approach we consider towards capturing the factorization and homeostasis aspects in our ontology, and then detail the two factors in turn in sections 4 and 5.

In brief, our approach is to treat an interoceptive signal as an ObservableProperty, as per the SSN ontology [12, 13], which originates inside a system defined as explicated by Mizoguchi and Borgo [14]. We intend to eventually use the Mizoguchi-Borgo concept of system, as opposed to the SSN concept of system, because of its more sophisticated account of function and structure[1]. However, as an axiomatization of the Mizoguchi-Borgo approach is not yet available, we will here refer to $SSN : System$ instead. An interoceptive signal may also originate inside a subsystem. In such a case, it will also be an interoceptive signal for the system as a whole. This approach allows us to associate signals to (sub)systems and/or their boundaries, where we impose the further constraint that a properly defined system is only one where the variables characterizing its boundary form a Markov blanket for the variables originating inside the system. I.e., that variables originating inside the system are conditionally independent from variables outside the system, as long as the boundary variables are given.

To account for homeostasis, we make again reference to the Mizoguchi-Borgo description of subsystems as playing functions inside a larger system [14] and assert in our ontology that interoceptive signals originating from a system or one of its subsystems will modulate some function of some of its subsystems[2].

Thus, we aim for our ontology to enable reasoning about boundary conditions and internal regulationwithin artificial agents. We assume and build on previously existing formalizations of behavior, function, and system [14, 23]. The following competency questions describe some of the modelling aspects that the ontology is intended to address.

1. *What is a system boundary? What does it mean for a variable to be "internal" or "external" to a system?* Formalizes how to represent a separation between a system and the rest of the world using the notion of Markov blankets: variables internal to a system only directly depend on each other and variables on the system's boundary, while the boundary variables are allowed to also directly depend on variables external to the system.
2. *Is a system boundary part of that system? Are variables describing the boundary state part of the variables describing the system state?* While it is convenient and intuitive to assert that the boundary, and its characterizing variables, belong to a system, it is also intuitive to keep a distinction between boundary and interior.
3. *When is an internal variable also an interoceptive one?* An internal variable is interoceptive if the system makes it available to internal monitoring and can use it to modulate its operation..

[1]SSN is by design minimal in its commitments as to what systems are and how they are structured.

[2]In the Mizoguchi-Borgo approach, the function of a system always makes reference to the system's exterior. The typical homeostatic loop closes inside a system, however, and is, in this understanding, better seen as the function of some "life-support" subsystem.

Together, these competency questions capture the ontology's capacity to reason about artificial interoception as a continuous, embodied process. They constrain the representation of mechanisms for adaptive regulation, predictive and explanatory inference, and the identification of critical boundaries and cross-domain dependencies between internal and external states.

Our ontology is intended both as a tool for human designers of interoceptive robots to describe and reason about the systems they design, as well as for the "online" use by such systems to organize their interoception-related reasoning. In this paper, when we will give examples, we will focus on the second use case and illustrate how concepts from the ontology could be used to define relevant queries to answer some of the competency questions above.

4. Internal-External Separation

A key challenge in modeling artificial interoception is determining which signals count as the "inside", or "outside", of a system. According to the free energy principle [24, 25], any system that maintains its organization over time must possess a statistical boundary, a *Markov blanket*, that separates internal from external states while allowing controlled information exchange between them. This boundary defines the minimal conditions for self-organization and provides a unified framework for describing the relationship between interoceptive (internal) and exteroceptive (external) processes [26].

4.1. Formalizing Agent-Environment Boundaries

In the subsequent discussion, we will be using signal and (random) variable interchangeably because to each signal that a system may measure as part of its interoception or perception, one can attach a random variable for the value of that signal. We begin by referring to a concept from the field of probabilistic inference.

Definition 4.1. Let there be a set of random variables $\mathscr{U}$. A *Markov blanket* for a set of variables $\mathscr{S} \subset \mathscr{U}$ is a set of variables $\mathscr{B} \subset \mathscr{U}$, disjoint from $\mathscr{S}$, such that $\mathscr{S}$ are conditionally independent of all other variables in $\mathscr{U}$ when given the blanket:

$$P(\mathscr{S}|\mathscr{U} - \mathscr{S}) = P(\mathscr{S}|\mathscr{B})$$

This notion is especially important in the domain of graphical probabilistic models, which are methods to describe, and reason with, relations of conditional dependence between random variables. In such models, vertices correspond to variables, and edges to (unmediated) dependency relations. There are many possible Markov blankets for a set of variables in a probabilistic graphical model, but a good intuition for the concept is given by thinking of the smallest blanket − the smallest set of vertices such that all paths between the variables of interest and other variables in the graph pass through the variables of the blanket.

The usefulness of this notion for our application is that it allows us to make precise what it means for a variable to be "inside" a system, and therefore for the signal it corresponds to, to be an internal one. Thus, in our view, a system is properly defined if it has a boundary such that the variables that characterize the boundary act as a Markov blanket for the variables that characterize the inner state of the system, separating them from the variables characterizing the environment.

We note that the concept of Markov blanket has a rich history of application in the literature about self-regulating systems. In an active inference context [24], it formalizes how internal and external states interact through sensory and active states. A Markov blanket separates the variables that belong to the agent from those that belong to the environment. Internal and external states influence one another only indirectly, through the blanket states.

Following [25], biological and cognitive systems can be viewed as hierarchically composed of such blankets, forming "Markov blankets of Markov blankets" that span multiple scales, from cells to organs to whole organisms and their sensorimotor loops. This recursive structure provides a principled way to describe hierarchies of subsystems.

Our proposal is to expand the notion of system defined by Mizoguchi and Borgo [14] to include the notion of variables characterizing a system state and its boundary, and a direct dependence relation (*directlyDependsOn*) that may exist between such variables, corresponding to the direct dependence relation from probabilistic graphical models. A constraint on a system is then that there would be no direct dependence relations between variables inside it and variables outside, with the sets of internal and external variables relative to a system being disjoint. Chains of dependence relations may exist however, as long as these pass through variables characterizing the boundary. This proposal also addresses the observation of [25] cited above, because our notion of subsystem would be similarly constrained.

Formally, our new axioms for this extension would include:

$$hasInnerVariable(X, V) \rightarrow \neg hasOuterVariable(X, V) \tag{1}$$

$$hasOuterVariable(X, V) \rightarrow \neg hasInnerVariable(X, V) \tag{2}$$

$$hasSubSystem(X, T) \wedge hasInnerVariable(T, V) \rightarrow hasInnerVariable(X, V) \tag{3}$$

$$\begin{aligned} SSN \;:\; & System(X) \wedge hasInnerVariable(X, V) \wedge \neg hasInnerVariable(X, W) \wedge \\ & \neg hasBoundaryVariable(X, W) \rightarrow \neg directlyDependsOn(V, W) \end{aligned} \tag{4}$$

Note that the absence of a direct dependence relation between variables does not mean an absence of dependence between them. Typically, variables characterizing the system and the environment will depend in complex ways on each other, and it is virtually always possible to infer something about a variable by knowing something about another one.

Example: Consider a robot's ambient and CPU temperatures as random variables. These variables would be separated by another variable for the temperature of the robot's outer shell, even if this variable is not actually measured. We would then say in our modelling that the ambient and CPU temperature variables do not directly depend on each other, but there would be a chain of dependence that links CPU temperature to shell temperature and shell temperature to ambient temperature. This is not a deterministic connection. E.g., there may be some other factor that controls the ambient temperature and which is unknown to the robot. However, even if the robot only measures its CPU temperature, it has some information to estimate the ambient temperature (it will likely be similar to the CPU, and possibly cooler).

5. Interoceptive State and Decision-Making

Interoceptive inference links self-regulation with decision-making by signaling how well the system maintains internal balance. Interoceptive sensors track signals of variables such as energy or temperature and indicate when they move away from preferred conditions. In such cases, regulatory systems are triggered to attempt to steer the interoceptive variable back to its desired range.

We propose to expand existing formalizations of functions [23] to account for functions that measure and modulate the variables characterizing a system. Following these formalizations, it is not a system's function as a whole to measure or regulate its internal variables, but rather these functions are attributed to some of its subsystems. We might say as shorthand that a robot "monitors and controls its temperature", but really these functions are more precisely understood when allocated to the robot subsystems.

Now the difference becomes clear between interoceptive and merely internal variables. A variable is interoceptive if it is internal to a system, monitored by the system, and (at least to some extent) controllable by the system. We do not make strong commitments about the system being able to always control the variable, or actively attempting to control it at all times. Rather, we propose that to be an interoceptive variable in a system means there exists a subsystem whose function it is to regulate this variable, independent of whether this subsystem is activated at some time or not.

Formally:

$$hasInteroceptiveVariable(X, V) \rightarrow hasInnerVariable(X, V) \tag{5}$$

$$hasInteroceptiveVariable(X, V) \rightarrow SSN : hasSubSystem(X, M) \wedge SSN : observes(M, V) \tag{6}$$

$$hasInteroceptiveVariable(X, V) \rightarrow SSN : hasSubSystem(X, R) \wedge SSN : actsOnProperty(R, V) \tag{7}$$

Another relevant aspect of interoceptive variables is that often – though not always – they are related to conditions under which a system may function. That is, some of these variables must remain within some given ranges, or else the continuation of the system's "life" may be compromised [26]. Using the SSN ontology [12] also allows us to represent assertions about interoceptive variables that capture their meaning in terms of the conditions in which a system should operate. This can be done by treating interoceptive variables as instances of the SSN concepts of operating and survival properties, each of which has ranges attached. Since these assertions are straightforward to make with the conceptual machinery of SSN, we do not outline them here.

Example: A navigation task requiring continuous sensing cannot proceed if CPU load or battery level are near critical thresholds. The system must first cool or recharge before continuing. In this way, interoceptive state functions as a decision variable linking physical regulation to motivational valuation. Note that this regulation is not actively maintained forever – we do not expect the robot to be constantly recharging. Likewise, regulation will not always be successful, as a robot might run out before reaching the charger. But we do expect that if something like battery level is to count as an interoceptive signal, then the robot should be able to notice it and have the machinery or behavioral routines to do something about it.

6. Conclusion and Future Work

In this paper, we have proposed the INTERO ontology, a descriptive framework for modeling robotic interoception grounded in the SSN and SOSA ontologies. The model formalizes how internal signals, system boundaries, and regulatory processes can be represented to capture the maintenance and interpretation of an agent's internal state. By doing so, it provides a structured basis for describing how internal and external variables interact through defined boundaries and how these relations support stability and adaptive control.

This formalization constitutes a step toward artificial agents that are capable of evaluating and regulating their internal states in ways that can be compared to human interoception. By integrating interoceptive reasoning into ontological representations of system behavior, INTERO connects regulation with adaptation and decision-making. Over time, such representations may help explain how consistent self-regulation can give rise to intrinsic motivation and other emergent forms of adaptive behavior.

Future extensions will address how interoceptive variables can inform learning, reward, and self-modification in order to enable agents to refine their behavior through feedback from their own internal processes. It will also be important to examine how embodiment and physical design influence interoceptive reasoning in practice. Further work should also connect INTERO to existing semantic representations used in robotics, to facilitate interoperability between models of internal regulation and conventional tasks or configuration ontologies.

In parallel, extending the framework to virtual and simulated agents would enable reasoning about interoceptive processes that emerge from computational or energetic constraints rather than physical ones. Finally, it will be valuable to explore how interoceptive representations interface with metacognitive models of self-monitoring and control, linking regulation to higher-order reflection. Ultimately, this work aims to advance modular ontologies that describe the foundational mechanisms underlying adaptive, self-regulating, and cognitively integrated artificial systems.

Declaration on Generative AI

The author(s) have not employed any Generative AI tools.

Acknowledgements

This work was supported by the German Research Foundation DFG, as part of Collaborative Research Center (Sonderforschungsbereich) 1320 Project-ID 329551904 “EASE - Everyday Activity Science and Engineering”, subproject “P01 – Embodied semantics for the language of action and change: Combining analysis, reasoning and simulation”.

References

[1] T. Engelen, M. Solcà, C. Tallon-Baudry, Interoceptive rhythms in the brain, Nature Neuroscience 26 (2023) 1670–1684. URL: https://www.nature.com/articles/s41593-023-01425-1. doi:10.1038/s41593-023-01425-1, publisher: Nature Publishing Group.

[2] A. D. Craig, How do you feel? Interoception: the sense of the physiological condition of the body, Nature Reviews Neuroscience 3 (2002) 655–666. URL: https://www.nature.com/articles/nrn894. doi:10.1038/nrn894, publisher: Nature Publishing Group.

[3] M. Tsakiris, H. Critchley, Interoception beyond homeostasis: affect, cognition and mental health, Philosophical Transactions of the Royal Society B: Biological Sciences 371 (2016) 20160002. URL: https://royalsocietypublishing.org/doi/10.1098/rstb.2016.0002. doi:10.1098/rstb.2016.0002.

[4] J. Gibson, Mindfulness, Interoception, and the Body: A Contemporary Perspective, Frontiers in Psychology 10 (2019). doi:10.3389/fpsyg.2019.02012.

[5] S. S. Khalsa, R. Adolphs, O. G. Cameron, H. D. Critchley, P. W. Davenport, J. S. Feinstein, J. D. Feusner, S. N. Garfinkel, R. D. Lane, W. E. Mehling, Interoception and mental health: a roadmap, Biological psychiatry: cognitive neuroscience and neuroimaging 3 (2018) 501–513. URL: https://www.sciencedirect.com/science/article/pii/S2451902217302343, publisher: Elsevier.

[6] F. Adams, K. Aizawa, Embodied cognition and the extended mind, in: The Routledge companion to philosophy of psychology, Routledge, 2019, pp. 193–213. URL: https://www.taylorfrancis.com/chapters/edit/10.4324/9780429244629-13/embodied-cognition-extended-mind-fred-adams-ken-aizawa.

[7] S. Gallagher, Embodied and enactive approaches to cognition, Cambridge University Press, 2023. URL: https://www.cambridge.org/core/elements/embodied-and-enactive-approaches-to-cognition/FFAB4E136E12255CD3139A87ECB2DCAD.

[8] L. Shapiro, Embodied cognition, Routledge, 2019. URL: https://www.taylorfrancis.com/books/mono/10.4324/9781315180380/embodied-cognition-lawrence-shapiro.

[9] M. Wilson, Six views of embodied cognition, Psychonomic Bulletin & Review 9 (2002) 625–636. URL: https://doi.org/10.3758/BF03196322. doi:10.3758/BF03196322.

[10] D. Beßler, R. Porzel, M. Pomarlan, A. Vyas, S. Höffner, M. Beetz, R. Malaka, J. Bateman, Foundations of the socio-physical model of activities (SOMA) for autonomous robotic agents, in: F. Neuhaus, B. Broderic (Eds.), Proceedings of the 12th International Conference on Formal Ontology in Information Systems (FOIS 2021), volume 344 of *Frontiers in Artificial Intelligence and Applications*

(FAIA), IOS Press, Bozen-Bolzano, Italy, 2021, pp. 159–174. doi:10.3233/faia210379, iSBN: 978-1-64368-248-8.
[11] R. Nolte, M. Pomarlan, D. Beßler, R. Porzel, R. Malaka, J. Bateman, Towards an ontology for robot introspection and metacognition, in: N. Aussenac-Gilles, T. Hahmann, A. Galton, M. M. Hedblom (Eds.), Proceedings of the 13th International Conference on Formal Ontology in Information Systems (FOIS 2023), volume 377 of *Frontiers in Artificial Intelligence and Applications (FAIA)*, IOS Press, Sherbrooke, Québec, Canada, 2023, pp. 318–333. doi:10.3233/FAIA231137, iSBN: 978-1-64368-469-7.
[12] M. Compton, P. Barnaghi, L. Bermudez, R. García-Castro, O. Corcho, S. Cox, J. Graybeal, M. Hauswirth, C. Henson, A. Herzog, V. Huang, K. Janowicz, W. D. Kelsey, D. Le Phuoc, L. Lefort, M. Leggieri, H. Neuhaus, A. Nikolov, K. Page, A. Passant, A. Sheth, K. Taylor, The ssn ontology of the w3c semantic sensor network incubator group, Web Semant. 17 (2012) 25–32.
[13] K. Janowicz, A. Haller, S. J. Cox, D. Le Phuoc, M. Lefrançois, Sosa: A lightweight ontology for sensors, observations, samples, and actuators, Web Semant. 56 (2019) 1–10. URL: https://doi.org/10.1016/j.websem.2018.06.003. doi:10.1016/j.websem.2018.06.003.
[14] R. Mizoguchi, S. Borgo, The role of the systemic view in foundational ontologies, in: proceedings of the 5th Workshop on Foundational Ontology, 2021.
[15] S. Yonekura, H. Kanazawa, A. Narita, Y. Kuniyoshi, A core driver of somatic markers: Interoception shapes the transition of decision making, 2025. URL: https://www.biorxiv.org/content/10.1101/2025.03.05.641768v1. doi:10.1101/2025.03.05.641768, pages: 2025.03.05.641768 Section: New Results.
[16] L. Steels, R. Brooks, The Artificial Life Route to Artificial Intelligence: Building Embodied, Situated Agents, Routledge, 2018. Google-Books-ID: 7b5aDwAAQBAJ.
[17] A. K. Seth, K. J. Friston, Active interoceptive inference and the emotional brain, Philosophical Transactions of the Royal Society B: Biological Sciences 371 (2016) 20160007. URL: https://royalsocietypublishing.org/doi/10.1098/rstb.2016.0007. doi:10.1098/rstb.2016.0007, publisher: Royal Society.
[18] G. Sandini, A. Sciutti, P. Morasso, Artificial cognition vs. artificial intelligence for next-generation autonomous robotic agents, Frontiers in Computational Neuroscience 18 (2024). URL: https://www.frontiersin.org/journals/computational-neuroscience/articles/10.3389/fncom.2024.1349408/full. doi:10.3389/fncom.2024.1349408, publisher: Frontiers.
[19] S. Lee, Y. Oh, H. An, H. Yoon, K. J. Friston, S. J. Hong, C.-W. Woo, Life-inspired Interoceptive Artificial Intelligence for Autonomous and Adaptive Agents, 2023. URL: http://arxiv.org/abs/2309.05999. doi:10.48550/arXiv.2309.05999, arXiv:2309.05999 [cs].
[20] N. Yoshida, H. Kanazawa, Y. Kuniyoshi, Synthesising integrated robot behaviour through reinforcement learning for homeostasis, 2024. URL: https://www.biorxiv.org/content/10.1101/2024.06.03.597087v2. doi:10.1101/2024.06.03.597087, pages: 2024.06.03.597087 Section: New Results.
[21] X. Zhou, C. C. Menassa, V. R. Kamat, Interoceptive Robots for Convergent Shared Control in Collaborative Construction Work, 2025. URL: http://arxiv.org/abs/2501.09290. doi:10.48550/arXiv.2501.09290, arXiv:2501.09290 [cs] version: 1.
[22] H. Laurençon, C.-R. Ségerie, J. Lussange, B. S. Gutkin, Continuous Homeostatic Reinforcement Learning for Self-Regulated Autonomous Agents, 2021. URL: http://arxiv.org/abs/2109.06580. doi:10.48550/arXiv.2109.06580, arXiv:2109.06580 [cs].
[23] R. Mizoguchi, Y. Kitamura, S. Borgo, A unifying definition for artifact and biological functions, Appl. Ontol. 11 (2016) 129–154. URL: https://doi.org/10.3233/AO-160165. doi:10.3233/AO-160165.
[24] K. Friston, Life as we know it, Journal of the Royal Society Interface 10 (2013) 20130475.
[25] M. Kirchhoff, T. Parr, E. Palacios, K. Friston, J. Kiverstein, The markov blankets of life: autonomy, active inference and the free energy principle, Journal of The royal society interface 15 (2018) 20170792.
[26] M. Solms, K. Friston, How and why consciousness arises: some considerations from physics and physiology, Journal of Consciousness Studies 25 (2018) 202–238.