# INSIGHT: The quest for causal explanations

Pablo Bustos[1*], Pilar Bachiller[1†], Ismael García-Varea[2*], Jesus Martínez-Gómez[2†], Antonio Bandera[3*], Rebeca Marfil[3†]

[1*]RoboLab, Universidad de Extremadura, Spain.
[2]SIMD, Universidad de Castilla-La Mancha, Spain.
[3], Universidad de Málaga, Spain.

*Corresponding author(s). E-mail(s): pbustos@unex.es; ismael.garcia@uclm.es; ajbandera@uma.es;
Contributing authors: pilarb@unex.es; jesus.martinez@uclm.es; rebeca@uma.es;
†These authors contributed equally to this work.

**Abstract**

Autonomous machines currently lack a complete understanding of their surroundings, missions, and the unexpected events that may affect their performance. It is evident that robots require a deeper level of understanding to take actions that make sense. The low levels of awareness in robots are a significant barrier in various robotic application domains. Machines lack the same level of understanding as their engineers and the people around them, which limits their performance, dependability, and trustworthiness. There is a need for a more profound core architecture that enables robots to comprehend their world and mission better. Deep learning alone is not enough to solve this problem. The ability to create explanations about a robot's own experiences, especially the capacity to causally explain unexpected events, is crucial in building socially aware robots. Our paper outlines the key proposition of the INSIGHT project, which is to equip robots with the ability to navigate their surroundings, comprehend them by verifying the coherence of their internal models, and develop new causal explanations in response to detected inconsistencies. The cognitive architecture used in the project will be expanded by integrating new memories (both semantic and episodic) and internal processes that control the flow of information between them. The new key element added is an internal simulator, which the robot will use to test alternative hypotheses until an effect is causally linked to its potential causes. Once established, the cause will be internalized as new knowledge, which will be used to accomplish future missions. A key output of this process will be a set of human-understandable explanations that will enable the robot to engage in a clarifying dialogue with its human supervisor.

**Keywords:** robotics, causal inference, cognitive, reasoning

# 1 Introduction

In the scope of mobile robotics, a profound need exists for them to possess a comprehensive understanding of their open environments. This understanding goes beyond mere identification of interactions with surrounding objects, such as recognising charging opportunities when encountering a battery station or utilising environmental data for traditional location and navigation tasks [1]. In today's landscape, the challenges are increasingly ambitious, striving to equip robots with the capacity to articulate their own experiences. Similar to humans, robots may encounter anomalies or unforeseen situations during their journeys that defy explanation based on prior knowledge. The capability to apprehend and elucidate these anomalies by establishing causal relationships among surrounding objects offers robots a priceless source of data for enhancing their behaviour and, consequently, refining future experiences [2]. This method of knowledge discovery can be denoted as the 'coordination of claims and evidence [3]' and assumes some pre-existing awareness of surrounding objects, typically achieved through the use of models and ontologies. This paper presents the INSIGHT project, a collaborative effort across three different universities aimed at advancing the development of self-understandable and self-extendable robots. Our principal objective within the INSIGHT project is to analyse unexpected experiences to construct new knowledge, including explanation capabilities.

Causal reasoning is one of the cornerstones of INSIGHT, even though it remains relatively underexplored in robotics [4]. Causality [5] will be investigated to generate and compare new hypotheses that could account for anomalous experiences. Furthermore, causal relations serve as a valuable tool for providing additional explanations, allowing newly discovered knowledge to be expounded through natural human-robot interaction processes.

The INSIGHT project draws upon the expertise of three research groups in the fields of robotics and artificial intelligence. Furthermore, it will also leverage certain pre-existing components that will streamline the developmental phases. These components include a mobile robot with a cognitive architecture featuring a working memory (referred to as DSR: Deep State Representation). This memory comprises a multi-labelled directed graph containing both symbolic and geometrical information. In addition to the physical robot, there is an embedded physics simulator (EPS) capable of integrating data from both the real world and simulations. The simulator's role will be pivotal in evaluating causal hypotheses through scenario assessments.

The design and development stages involved in the research project are situated within a use case, which is employed to illustrate the challenge faced when a robot encounters unexpected experiences based on its prior knowledge. In these scenarios, the robot may be incapable of providing a valid rationale for the events that transpired, necessitating an adaptation of its knowledge and, consequently, the acquisition of new insights from the experience. This learning process is facilitated through a self-understanding procedure, as detailed in Section 4, which endeavours to identify causal explanations for unexpected events. Within this process, the semantic memory is examined to propose and validate new causal relationships between objects, with the assistance of the internal simulator, that can account for the preceding experience. These newly established relationships should also be suitable for presenting a coherent explanation to a human through a phase of natural human-robot interaction.

## 2 Related work

In recent years, the incorporation of causal reasoning in robotics, a relatively underexplored area, has been increasingly noted as pivotal in fostering the development of truly smart robots [4]. Current research efforts are leaning towards equipping robots with the ability to establish causal rules through various methods such as imitation learning, where robots discern the rules of intuitive physics by observing actions and their subsequent effects [6], or learning from demonstrations to enhance spatial, temporal, and causal knowledge for task execution [7, 8].

A critical area of focus in the current landscape lies in the integration of deep learning techniques into robotic systems to promote enhanced awareness and response mechanisms. Research conducted by LeCun, Bengio, and Hinton in [9] underscores deep learning as a potent tool for enabling machines to learn from a substantial volume of data. However, it also underscores the necessity of complementing it with other methodologies to attain a higher level of comprehension and interaction. In [10], the authors advocate a shift in paradigm that places common sense at the core of the discourse. They identify functionality, physics, intent, causality, and utility (FPICU) as the five fundamental domains of cognitive AI imbued with humanlike common sense. From this vantage point, addressing the 'how' and 'why' becomes the more pressing inquiry, necessitating reasoning about unobservable factors beyond visible pixels.

Moreover, the concept of robots possessing the capability to generate explanations comprehensible to humans has garnered scholarly interest. This concept is in harmony with the objectives pursued within the INSIGHT project. Research by Hayes and Shah [11] in the realm of human-robot interaction has illuminated the significance of robots being adept at participating in elucidatory dialogues with humans. This proficiency requires a degree of causal reasoning to elucidate their experiences and actions effectively.

One significant advancement in this field is the development of an inner simulator that facilitates the identification and exploration of potential causes behind specific effects. This simulator aids in evaluating error situations, permitting a robot to discern possible reasons behind an unsuccessful action and adjust its approach accordingly, thus stepping beyond the limitations noted by Wyatt et al., where robots were restricted to predefined notions of possible causes [12]. In a broader scope, the idea of embedding an internal simulator as a core element of the cognitive architecture is gaining momentum. From the findings of mirror neurons [13] to research in consciousness [14, 15] and cognitive science [16], the playground has been prepared for cognitive robots. Internal simulation is being explored from different angles, whether using hyperrealistic simulators [17], probabilistic models that can be inverted to yield explanations [18], or as tools to reason under occlusion [19].

The INSIGHT project introduces a novel approach in this domain, distinguishing itself from other similar initiatives like the one presented by Brawer et al. [20]. Unlike the latter which relies on physical exploration, INSIGHT utilizes a combination of observation and inner simulation, enabling robots to probe various potential causes by manipulating influencing factors within simulated environments. Furthermore, this approach potentially enhances data efficiency and facilitates the generation of self-conducted experiments, transcending the observational learning constraints delineated in works by Han et al., where causal variables were mandated to be observable [21].

The progressive trend in robotics research is steering towards the integration of causal reasoning and learning, fostering a generation of robots capable of intelligent, autonomous decision-making through simulation and analysis of a broad spectrum of potential causal factors. This approach, pioneered by initiatives like INSIGHT, promises to revolutionize the robotics domain by engendering robots that are not only receptive to causal influences but also capable of innovatively responding to them.

## 3 The core experiment: What tipped over the medicine bottle?

The core of the INSIGHT project focuses on one detailed experiment that should be carried out and that conveys all the theoretical problems and hypotheses raised in the project. The experiment has been designed to keep the essential elements of the problem while simplifying the context in which the robot evolves. The scenario consists of a robot named Shadow (see Fig. 1)that follows its human supervisor, Ana, on a drug delivery mission at a healthcare facility. Before starting, Ana leaves a medicine box on the robot's tray, making it aware of it. As she proceeds towards her destination, Shadow encounters an unforeseen obstacle, causing the medicine to spill and scatter on the ground. Despite the robot's IMU sensing the disruption, it remains unaware of the missing medicine until it scans the area with its head camera, in accordance with its pre-programmed mission, a few moments later. The robot halts its movements and contacts Ana, recognizing that the medicine is a crucial prerequisite for the plan to proceed. Ana promptly responds to the notification, realizing the absence of the medicine, and initiates a conversation with Shadow.

ANA: What happened?
SHADOW: I hit a bump, and the medicine fell to the ground.
ANA: Where is it now?
SHADOW: It should be 2 metres behind, on the floor.
ANA: Do you think the medicine might be broken?
SHADOW: No, I don't think so.
ANA: OK, let me pick it up and put it back on the tray (The medicine is put back on the tray. The robot looks at it and updates its state).
SHADOW: Let's keep going! We are running late.

**Fig. 1**: The robot Shadow

In this conversation, Shadow successfully communicates the circumstances to Ana and explains that a bump was the cause of the accident. To meet the deadline, it is necessary to complete most of the explanation in the second line before Ana asks any questions. As a necessary effect of building this explanation, the new BUMP concept will be grounded by synthesizing a BUMP detector and making it available to the control system. Also, the future behaviour of Shadow will take into account this new functionality to detect, anticipate and avoid potential bumps on its way. Before we proceed with the use case, it is important to establish a set of conditions that must be met prior to searching for an explanation.

a) Shadow is running an instance of the CORTEX architecture (see Figure 2) that includes several memories, including working, semantic and episodic memories, and other modules that connect them [22][23].
b) The robot's sensors include a camera, an IMU, and a microphone.
c) The robot's IMU senses the inclination of the robot but, as it is not categorised as a known event, it is quietly recorded in episodic memory and ignored. The same situation applies to the camera.
d) One of these elements is an embedded Physics simulator, bottom-right in Figure 2, that can run synchronised with the robot's activity. The speed of the robot and the objects detected by the perception system are fed into the simulator to update its state.
e) The simulator is initialised with a nominal plan, namely Follow[Ana], and with the environment's elements perceived by the robot: the person, the floor, and the structural elements around it.

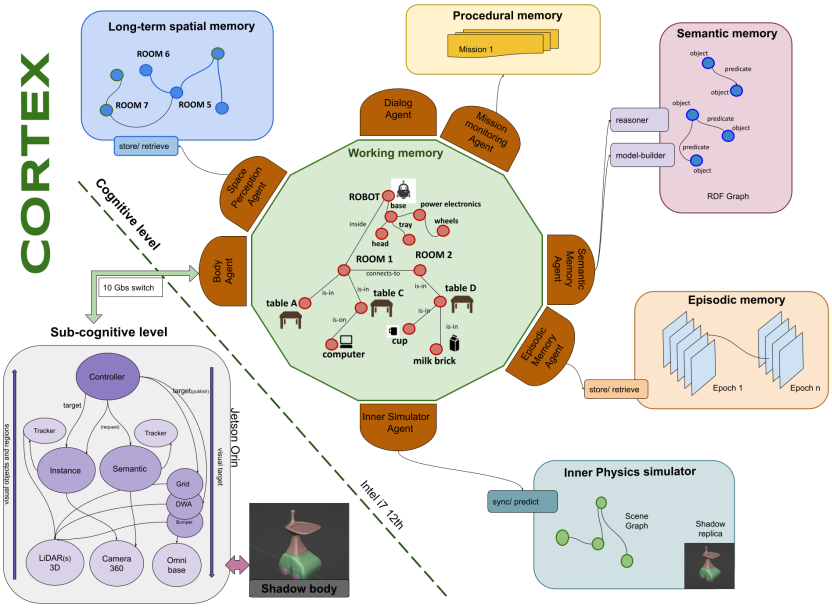


**Fig. 2**: The CORTEX cognitive architecture

f) The output of the simulated robot's sensors is injected back into the working memory in real time as dual, synthetic versions of the robot's real sensors. These "efferent" copies can be used to detect mismatches between the state predicted by the model inside the simulator, and the state read by the sensors.
g) After the initial interaction with the human, the robot starts a mission that consists of following Ana while checking that the medicine remains safely on the tray.

## 4 Outline of a solution

Below we present a sequence of steps that, according to the project's hypothesis, would lead to the construction of a causal explanation and the grounding of the BUMPER concept. These internal actions on the architecture are a form of learning that changes it structurally and functionally.

1. As the robot moves forward, its camera registers the bump on the floor, but it is silently ignored since there is no detector for it. The presence of a bump is not transmitted to the simulator.
2. When the robot hits the bump, a difference is measured between the real data read from the IMU and the synthetic data generated within the simulator. The sensorial anomaly (event) is registered but without further consequences. Nothing is triggered by such an event.
3. Because of the bump, the box tips over and falls on the floor. The robot notices that the medicine is missing and proceeds to stop the mission, halting and calling

the human. A new sensorial event is registered, i.e., the absence of the box, but this time it triggers an action since its presence was mandatory in the mission.

4. The self-understanding process starts here. It is inside the robot's control architecture to trigger the search for a causal explanation whenever an unexpected event, such as FALLEN_BOX, occurs.
5. The embedded simulator is taken out of sync with the external world. Then, it is restarted but this time as a contrastive hypothesis evaluator. The goal now is to search for the simplest scenario that generates a set of synthetic perceptive data approximating the timing and readings from the real sensors. In doing so, different runs will have to be performed, each one with a different initial situation. The loop will continue until a valid scenario is found that reproduces both the IMU perturbation and the box-dropping event, at approximate times.
6. To achieve this, the episodic memory is queried to retrieve the initial scenario occurring just before the beginning of the FALLEN _BOX event. This scenario will be extended with different new elements, extracted from the semantic memory, that could be a plausible cause of the experienced events.
7. The semantic memory is queried at this point for cause-effect statements containing the target known effect: FALLEN_BOX. Any other contextual symbols that might have been elicited by the robot's perceptual system, such as TILTED(robot), or MOVING(robot) will be added to the query. In the following, we try to maintain the generality of the case but assume a minimum set of conditions:

   (a) There is a symbol BUMP in the semantic memory's knowledge base. This is easy to assume since external sources of generic knowledge, i.e. public robotics ontologies knowledge graphs or large language models, can be accessed online.
   (b) There are also rules linking BUMP with the effects it causes on moving things, i.e. tilt, lean over, bouncing, roll over, etc. such as $BUMP \wedge MOVING(x) \implies TILT(x)$

   We now show a possible answer set after an informed search over the space of causal relations is completed:

   (a) $MOVING(x) \wedge BUMP \implies TILT(x)$ (base knowledge)
   (b) $MOVING(x) \wedge BROKEN\_WHEEL(x) \wedge CARRYING(x, y) \implies TILT(x) \wedge FALLING(y)$ (promising)
   (c) $MOVING(x) \wedge BUMP \wedge CARRYING(x, y) \implies TILT(x) \wedge FALLING(y)$ (optimal)
   (d) $BURGLAR(h) \wedge MOVING(x) \wedge CARRYING(x, y) \wedge VALUABLE(y) \implies STEALING(h, y) \wedge MISSING(y)$ (too elaborated)

   The tags at the end of the rules are only to show different types of possible candidates obtained from the available sources. No previous evaluations of their suitability are supposed.
8. The next step is to transform each accepted entry from the answer set into an initial scenario for the simulator. All initial scenarios are derived from the same situation recovered from the episodic memory and then extended with some additional variations derived from the rule being tested. The extension of the scenario is a crucial part that involves addressing the *grounding* problem since we need to transform a

group of ungrounded, externally obtained symbols, i.e. BUMP, into functional elements of the simulation. In the case of the concept BUMP, we will assume that one of its defining properties is a 3D mesh representing a prototypical protrusion. The simulator's initial scenario is a textual description of the elements in the scene and their kinematics properties. To extend it with the bump, we need to add a new object defined by the mesh file and its position and size in the scene. This gives us three free parameters, the 2D position and the size, that have to be evaluated along a reasonable range of values. The complete process has thus to iterate over two loops, the list of candidate hypotheses and their free parameters, until a scenario is found whose synthetic sensorial experiences match the real experience. The specific scenario that best matches the IMU and medicine tilting event, in time and in the shape of the response, will be selected as the optimal one and will be considered the causal explanation of the event.

9. The end of the search triggers for crucial new processes:
   (a) First, the selected rule that links the bump with the tilting and falling of the medicine is marked in semantic memory as a *grounded* rule, meaning that it has passed a contrastive hypothesis procedure triggered by a perceptual anomaly.
   (b) This is followed by the construction of the human-understandable explanation demanded by Ana, and that will be verbalised in the ongoing dialogue with her.
   (c) The additional details requested by Ana involve the location of the fallen box and the possibility of it being broken. This is a reasoning process that needs to access the results of the simulation. The distance to the medicine can be obtained from the final spatial coordinates of the bump. The integrity of the medicine could also be similarly obtained if it had the capability to simulate the breaking of falling crystal bottles. If not, a pre-existent formula or a table relating falling height and bottle material characteristics with its breaking point would be necessary. This data could be obtained from an external oracle, such as a large language model (LLM).
   (d) To advance in the grounding process of the BUMP concept, the robot should be able to recognize a bump if it hits it, but also in advance of this hit. The first situation would involve a classifier on the IMU data, and the second on the image and/or lidar data. In doing so, it could avoid future accidents and by knowing now why it did that, it would be able to explain it to the humans. But building a classifier on the fly is not an easy task, even a simple one. A possible approach would be to collect data from the two classes, bump/no-bump, in both modalities and train a neural network as a binary classifier. The problem is that there is only one example of the bump class recorded in the episodic memory. The only alternative is to synthesize more examples with the internal simulator. Since we do have a virtual bump, it is a matter of running variations of the robot going over it, changing its size and the robot's speed. Note that this is a third different use of the same simulator in the CORTEX architecture. Assuming that we have the code to operate the simulator as a data collector, the rest of the process implies the selection of a neural network architecture and the training until convergence. Once cross-validated, the model can be used to generate a new component that would be added to the architecture. Once installed, these

components, IMU and RGB, will detect and instantiate the BUMP concept every time it is experienced by the sensors.

(e) A final change in the existing architecture should define how the behaviour of the robot changes upon the perception of a new bump. How can the robot activate, modify or create an avoiding behaviour that was not there before and that avoids a bump on the floor? The concept BUMP will not be fully grounded until it is causally connected with the robot and its environment. If the robot already knows how to avoid obstacles, creating for instance a coastal map used by MPC controller, a simple solution would be to add the bump to the list of image regions that count as obstacles for the grid. To do this autonomously implies that the list is not hard coded and that the integration of the new rule implies this action, by interpreting that the box falling to the floor is something that must be avoided.

# 5 Reconstruction of a virtual experiment

Assuming that all the steps presented in Section 4 were correctly performed, we can run a virtual experiment in which the robot meets a new bump during a fresh mission. To build this experiment, we only need to focus on the working memory that occupies the central part of Figure 2, where the robot and objects close to it are represented, and on the semantic memory. Agents, in brown, connect memories, pumping information among them.

Starting from a similar initial situation with the robot Shadow following Ana on a new delivery mission, the robot's extended perception system visually detects the bump ahead before running over it. The corresponding region of space is added to the coastal map as an obstacle and the navigation algorithm executes a detour to avoid it. As a consequence of the triggering of the new BUMP detector, a corresponding concept node is injected into the working memory, with an arc connecting it to the node representing the floor (or the room, hallway, etc.). According to CORTEX's dynamics, the semantic memory will immediately detect the new node and will retrieve some potentially relevant rules that include the concept BUMP. Assuming that the one just learnt is found, it will be injected into the working memory edges labelled *TILT* and *FALLING*, connecting it with the node representing the medicine box. This new information internally added from the semantic memory will be used to build a verbal explanation if the human notices the deviation taken by the robot and asks for the cause of it. A final effect of completing this new mission will be the acquisition of new images of a bump and their storage for later retraining of the recently created classifier.

An initial analysis of this problem and the path to a solution outlined here underlines three steps of major difficulty. First, the process that goes from the candidate rules obtained from the semantic memory to the initialization of the internal simulator with the corresponding variations; second, the synthesis of a classifier on the fly using only data from one real experience and the augmented samples that could be derived from it with the simulator; and third, the use of the new concepts or relations introduced from previously ungrounded knowledge, such as BUMP or FALLING, by the hard-coded controllers, so they end up *making sense* inside the architecture. This is an instance of the open-world problem that could be partially overcome if all new elements have known super types previously known to the controllers. These open problems and many others that will arise are the subject of this research effort that will last for the next few years.

# 6 Conclusions

In this paper, we have introduced the new INSIGHT research project, which aims to examine the current challenges in robot learning and offer solutions grounded in contemporary causal approaches. The inception of this project stemmed from the imperative to enhance robots' comprehension of their environment and enable them to construct explanations based on their own experiences. This objective aligns with two prominent themes in the field of artificial intelligence: causality and explainability.

To concretize the ambitious objectives of the project, it incorporates a use case that represents a plausible scenario a robot might encounter, involving events that defy explanation based on its pre-existing knowledge. This use case serves as a common foundation for project members to develop and integrate novel proposals and advancements. As a result of the INSIGHT research project, we anticipate yielding compelling outcomes across various domains, including cognitive architectures and human-robot interaction.

# 7 Acknowledgments

This work has been partially funded by the Spanish Ministry of Science and Innovation TED2021-131739-C22, supported by MCIN/AEI/10.13039/501100011033 and the European Union "NextGenerationEU"/PRTR, by the Spanish Ministry of Science and Innovationan PDC2022-133597-C41 and by FEDER Project 0124 EUROAGE MAS 4 E (2021-2027 POCTEP Program). This work has been funded by the project SBPLY/21/180225/000062 funded by the Government of Castilla-La Mancha and "ERDF A way of making Europe". This work has also been supported by Universidad de Castilla-La Mancha and "ERDF A way of making Europe" under project 2022-GRIN-34437. It is also partially funded by MCIN/AEI/10.13039/501100011033 and "ESF Investing your future" through the projects PID2019–106758GB–C33, PID2022-137344OB-C31, PID2022-137344OB-C32 and PID2022-137344OB-C33.

# References

[1] Stocking, K.C., Gopnik, A., Tomlin, C.: From robot learning to robot understanding: Leveraging causal graphical models for robotics. In: Conference on Robot Learning, pp. 1776–1781 (2022). PMLR

[2] Pelivani, E., Cico, B.: Toward self-aware machines: Insights of causal reasoning in artificial intelligence. In: 2021 International Conference on Information Technologies (InfoTech), pp. 1–4 (2021). IEEE

[3] Sandoval, W.A.: Epistemic goals. Encyclopedia of science education, 393–398 (2015)

[4] Hellström, T.: The relevance of causation in robotics: A review, categorization, and analysis. Paladyn, Journal of Behavioral Robotics **12**(1), 238–255 (2021)

[5] Pearl, J.: Causality, 2nd edn. Cambridge University Press, Cambridge, UK (2009). https://doi.org/10.1017/CBO9780511803161

[6] Agrawal, P., Nair, A., Abbeel, P., Malik, J., Levine, S.: Learning to poke by poking: Experiential learning of intuitive physics. In: Proceedings of the 30th International Conference on Neural Information Processing Systems. NIPS'16, pp. 5092–5100. Curran Associates Inc., Red Hook, NY, USA (2016)

[7] Xiong, C., Shukla, N., Xiong, W., Zhu, S.-C.: Robot learning with a spatial, temporal, and causal and-or graph. 2016 IEEE International Conference on Robotics and Automation (ICRA), 2144–2151 (2016)

[8] Angelov, D., Hristov, Y., Ramamoorthy, S.: Using causal analysis to learn specifications from task demonstrations. In: Proc. 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), pp. 1341–1349. International Foundation for Autonomous Agents and Multiagent Systems, ??? (2019). International Conference on Autonomous Agents and Multi-Agent Systems 2019, AAMAS 2019 ; Conference date: 13-05-2019 Through 17-05-2019. http://aamas2019.encs.concordia.ca/,http://aamas2019.encs.concordia.ca/

[9] LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. Nature **521**(7553), 436–444 (2015) https://doi.org/10.1038/nature14539

[10] Zhu, Y., Gao, T., Fan, L., Huang, S., Edmonds, M., Liu, H., Gao, F., Zhang, C., Qi, S., Wu, Y.N., Tenenbaum, J.B., Zhu, S.C.: Dark, Beyond Deep: A Paradigm Shift to Cognitive AI with Humanlike Common Sense. Engineering **6**(3), 310–345 (2020) https://doi.org/10.1016/j.eng.2020.01.011 arXiv:2004.09044

[11] Hayes, B., Shah, J.A.: Improving robot controller transparency through autonomous policy explanation. In: Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction. HRI '17, pp. 303–312. Association for Computing Machinery, New York, NY, USA (2017). https://doi.org/10.1145/2909824.3020233 . https://doi.org/10.1145/2909824.3020233

[12] Wyatt, J., Aydemir, A., Brenner, M., Hanheide, M., Hawes, N., Jensfelt, P., Kristan, M., Kruijff, G.-J., Lison, P., Pronobis, A., Sjöö, K., Vrecko, A., Zender, H., Zillich, M., Skočaj, D.: Self-understanding and self-extension: A systems and representational approach. IEEE T. Autonomous Mental Development **2**, 282–303 (2010)

[13] Gallese, V.: Embodied Simulation and Its Role in Cognition. The Embodied Self.

Dimensions, Coherence and Disorders **7**(13), 77–91 (2018)

[14] Hesslow, G.: The current status of the simulation theory of cognition. Brain Research **1428**, 71–79 (2012) https://doi.org/10.1016/j.brainres.2011.06.026

[15] Holland, O., Goodman, R.: Robots with internal models. Journal of Consciousness Studies **10**(4), 1–45 (2003)

[16] Bass, I., Smith, K., Bonawitz, E., Ullman, T.D.: Partial Mental Simulation Explains Fallacies in Physical Reasoning. PsyArXiv (2021). https://doi.org/10.31234/osf.io/y4a8x . psyarxiv.com/y4a8x

[17] Mania, P., Kenfack, F.K., Neumann, M., Beetz, M.: Imagination-enabled Robot Perception. IEEE International Conference on Intelligent Robots and Systems (June), 936–943 (2021) https://doi.org/10.1109/IROS51168.2021.9636359 arXiv:2011.11397

[18] Kenghagho, F.K., Neumann, M., Mania, P., Tan, T., Siddiky, F.A., Weller, R., Zachmann, G., Beetz, M.: NaivPhys4RP - Towards Human-like Robot Perception 'Physical Reasoning based on Embodied Probabilistic Simulation'. IEEE-RAS International Conference on Humanoid Robots **2022-Novem**(December), 815–822 (2022) https://doi.org/10.1109/Humanoids53995.2022.10000153

[19] Trinidad Barnech, G., Tejera, G., Valle-Lisboa, J., Núñez Trujillo, P.M., Bachiller Burgos, P., Castro, P.: Initial results with a simulation capable in robotics cognitive architecture. In: Proceedings of ROBOT2022: Fifth Iberian Robotics Conferenc, Zaragoza, España (2023)

[20] Brawer, J., Qin, M., Scassellati, B.: A causal approach to tool affordance learning. 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 8394–8399 (2020)

[21] Han, Z., Yanco, H.: Communicating missing causal information to explain a robot's past behavior. J. Hum.-Robot Interact. **12**(1) (2023) https://doi.org/10.1145/3568024

[22] Bustos García, P., García, J.C., Cintas Peña, R., Martinena Guerrero, E., Bachiller Burgos, P., Núñez Trujillo, P., Bandera, A.: Dsrd: A proposal for a low-latency, distributed working memory for cortex. In: Bergasa, L.M., Ocaña, M., Barea, R., López-Guillén, E., Revenga, P. (eds.) Advances in Physical Agents II, pp. 109–122. Springer, Cham (2021)

[23] Bustos García, P., Manso Argüelles, L., Bandera, A.J., Bandera, J.P., García-Varea, I., Martínez-Gómez, J.: The cortex cognitive robotics architecture: Use cases. Cognitive Systems Research **55**, 107–123 (2019) https://doi.org/10.1016/j.cogsys.2019.01.003