

Motivational System for the Discovery of Missions in Autonomous Robots

Diana Cerdas-Vargas, Emanuel Fallas-Hernández, Alejandro Romero, Fabián Álvarez-Juarez, and Richard J. Duro

Grupo Integrado de Ingeniería, Centro de Investigación CITIC, Universidade da Coruña, 15071 A Coruña, Spain

Correspondence: diana.cerdas@udc.es

DOI: <https://doi.org/10.17979/spu.23.c54>

Abstract:

This work introduces a motivational system for autonomous robots that allows the generation and learning of missions (and their associated drives) aligned with human purposes. The system is integrated into the e-MDB cognitive architecture, employing Large Language Models (LLMs) to interpret natural language instructions. The system employs three distinct LLM instances, each specialized for a specific function: semantic alignment, mission generation, and modeling of motivational drives. Experiments conducted within the Gazebo simulation environment demonstrated consistent alignment between robot behavior and human purposes, as well as high reliability in mission and drive validity.

1 Introduction

Autonomous robotics seeks to develop systems capable of learning and adapting to changing environments throughout their life cycle (Lifelong Open-ended Learning Autonomy) (Romero et al., 2023), while minimizing human intervention. This capability is essential in applications where direct supervision is not always possible, such as space exploration, elderly care, or high-level industrial automation. Over the years, various cognitive architectures have been proposed that incorporate motivational systems and/or decision-making mechanisms in order to meet this challenge. Some examples are the architectures H-GRAIL (Romero et al., 2025) or e-MDB (Becerra et al., 2021), among others (Kotseruba and Tsotsos, 2020).

In this context, the e-MDB (epistemic Multilevel Darwinist Brain) cognitive architecture provides robots with robust reasoning and decision-making mechanisms, which promote open and continuous adaptation throughout their operational life (Fallas Hernández et al., 2025). Recently, this cognitive architecture has been used for the PILLAR-Robots European project¹, which aims to develop a new generation of autonomous robots capable of setting their own goals, creating strategies, and learning from previous experience to support human needs in real-world scenarios. The project seeks to overcome the limitations of current robots, which are often restricted to predefined tasks, by enabling more flexible, adaptive, and independent systems that reduce reliance on human supervision. However, a major limitation of this architecture is that the missions and drives that guide the robot's behavior must be defined manually. This requirement not only reduces the level of autonomy that can be achieved, but also hinders the alignment between the purposes established by humans and the actions executed autonomously by the robot.

¹ <https://pillar-robots.eu/>

This work proposes a system that facilitates the transmission of human purposes to robots by receiving a purpose expressed in high-level language, semantically aligning it with human intent, and translating it into missions and drives that form the robot's motivational system. In order to better understand this work, the Section 2 outlines the challenges of human-robot interaction in purpose alignment. Then, Section 3 describes the key elements of the motivational system within the e-MBD cognitive architecture. Once all the theoretical concepts are clear, Section 4 presents the proposed interaction system and how it is coupled to the cognitive architecture. Finally, the experimental validation and results of this initial approach to purpose alignment in autonomous robots are reviewed in Sections 5 and 6.

2 Alignment problem in autonomous robots

Autonomous learning that is not explicitly aligned with human purposes poses significant risks, since the actions performed by the robot may be irrelevant or even harmful to humans. This challenge assumes particular significance in the context of Lifelong Open-ended Learning Autonomy (LOLA), wherein robots acquire novel competencies across diverse operational domains. A domain refers to the environment in which the robot operates as well as the robot itself, along with the set of tasks, objects, and interaction rules that govern its operation.

Although such flexibility is imperative for autonomy, it gives rise to a fundamental question: *How can we ensure that the skills acquired by a robot remain beneficial and aligned with human purposes?* To answer this question, (Baldassarre et al., 2025) proposed a purpose-based framework in which human users specify the objectives for the robot, either directly or indirectly, without constraining the means by which the robot accomplishes these goals. This framework highlights the translation of human purpose—user intentions in natural or symbolic language—into robot purpose, the internal needs, drives, and goals of the cognitive architecture.

The fundamental challenge, therefore, lies in bridging the semantic and representational gap between human purpose and robot purpose. Achieving a reliable translation between these two systems is crucial to ensure that open-ended learning capabilities remain beneficial to humans.

The subsequent section describes the e-MDB architecture, emphasizing the motivational system and its key components: needs, drives, and missions.

3 Motivational system: needs, missions, and drives

The e-MDB architecture has five fundamental components: The Long-Term Memory (LTM) system, interaction processes, learning system, motivational system, and off-line memory consolidation, imagination, and dreaming. Each of these parts has a fundamental role in how the robot is able to learn autonomously (Becerra et al., 2021). The proposed system will operate in the motivational system.

The motivational system regulates the robot's behavior through motivation management. It is composed of domain-independent needs and drives that gradually align with the specific goals identified by the robot in different domains (Fallas Hernández et al., 2025). Hence, the system enables the robot to discover new goals and determine which of them are active at a given time. This work focuses on the higher level of the motivational system (needs and drives); therefore, these concepts will be explained.

3.1 Needs and drives

Needs (n_j) represent a desired state in the robot's motivational state space (Romero et al., 2020). The motivational state space is the group of all the needs and drives the robot has, it is defined at design time and specifies dimensions that may correspond to internal parameters (e.g., battery level or temperature) or to externally defined human purposes. The satisfaction of a need is evaluated through its associated drive, which quantifies the distance between the current state

x_t and the target state n_j . In general, a drive can be expressed as in the Equation 45.1, where D_j decreases as the robot approaches the satisfaction of the need n_j (Romero et al., 2020). Drives therefore, act as continuous evaluators that determine the urgency of addressing each need.

$$D_j = f_j(x_t, n_j) \tag{45.1}$$

3.2 Types of needs and drives

In the context of the e-MDB cognitive architecture, two categories of needs and drives can be distinguished: cognitive and operational. Cognitive needs and drives promote exploration and knowledge acquisition, guiding how the robot expands its repertoire of skills. On the other hand, operational needs and drives correspond to human purposes (Romero et al., 2020). In this work, operational needs are referred to as missions, as they represent externally defined objectives that the robot must pursue. Missions thus serve as the link between human purposes and the robot’s adaptive behavior across different domains.

By integrating needs, drives, and missions, the motivational system provides the basis for balancing innate requirements with externally defined purposes. This framework supports the foundation for the solution described in the next section.

4 System description

This work presents a first approach to solve the alignment problem and also how to link the purpose translation with the motivational system of the cognitive architecture. The solution consists of three instances of large language models (LLMs) that work together as shown in Figure 1. The first instance uses **phi4:14b**, and it is responsible for aligning human purpose and ensuring that it is clearly specified. The second, also using **phi4:14b**, translates that aligned purpose into specific missions. Finally, a third instance, implemented with **GPT-4.1**, generates the drives associated with the missions and human purpose. The prompts, inputs and outputs of each LLM model can be found on: <https://github.com/DianaCerdasV/Interactive-System-for-e-MDB>.




Figure 1: Diagram of the interaction system with LLMs.

To adapt the interaction system into the e-MDB cognitive architecture, an *Alignment need* was created during the design time. This need activates the *Alignment policy*—a policy is a reactive decision structure that provides the best action to apply at a given moment—therefore, the *Alignment policy* executes the interaction system described above and, once the missions and drive functions have been generated, creates the cognitive nodes that represent them in the cognitive architecture space. In this way, the robot can learn and execute tasks that autonomously adapt to the purposes defined by humans.

5 Experiments

To validate the solution proposed in this paper, two general experiments were carried out. The following sections describe each experiment and its objective.

5.1 Alignment and translation system evaluation

This first experiment aims to evaluate the capacity of the designed system to align human purpose and generate valid missions and drive functions. Here, only the designed system is evaluated and, to simplify the evaluation, it is proposed to divide it into two tests: purpose alignment system evaluation, mission and drive function generators evaluation.

Purpose alignment system evaluation

The goal of this test is to validate the system's ability to interact with humans and transform a general purpose into a robot purpose that describes in detail the task to be performed by the robot. In this case, three different general purposes were used. For each of these tasks, a completely detailed purpose was written, which would be the "ideal" case that the system should achieve. Also, this is the one that will be used to compare the final results. The general purposes with their respective detailed purposes are shown below.

Then, 15 different ways of expressing these purposes were proposed² with a low level of detail, so that the system would have variability. To acquire the results of this test, each purpose was entered one by one into a new, trained instance of the system. The system was then interacted with until the final detailed description was collected. Once all the final descriptions from the alignment system had been registered, the semantic similarity between the initial purpose and the final description was calculated. To do this, the embedding vector for each of the phrases in the LLM space was obtained; for this, the **all-MiniLM-L6-v2 model** by UKPLab was used. This vector maps the meaning of a text into a numerical representation. Then, the cosine semantic similarity for the input and output purpose is calculated using Equation 45.2 Wang and Dong (2020). This was repeated for the 15 poorly detailed purposes and the final descriptions obtained from the three tasks.

$$\text{sim}(A, B) = \frac{A \cdot B}{\|A\| \|B\|} \quad (45.2)$$

Mission and drive functions generators evaluation

The aim of this test is to evaluate the system's ability to generate valid missions and drive functions from a series of fully defined human purposes. In this case, the detailed descriptions generated by the previously evaluated model and the robot's perceptions were taken as inputs, and the outputs were the missions and drive functions for each purpose. Then, to analyze the results, the missions and drive functions were classified as valid or invalid based on predefined criteria.

5.2 Simulated evaluation

This experiment was designed to evaluate the motivational system integrated with the e-MDB cognitive architecture and applied to a robot. The simulation was done in Gazebo, using the OSCAR robot as shown in the Figure 2. The experiment consists of the robot picking up the red cylinder and placing it inside the blue basket. In this case, it is assumed that the robot already knows how to perform policies such as: picking up the red cylinder, placing the cylinder inside the basket, and pressing the green button when the cylinder is out of reach. This experiment was repeated 10 times, each run had 2000 iterations and in each iteration, the robot could execute one policy.

² <https://github.com/DianaCerdasV/Interactive-System-for-e-MDB/>




Figure 2: Experiment configuration for OSCAR simulation in Gazebo.

6 Results

Figure 3 shows the cosine semantic similarity calculated between the inputs and outputs obtained for each task in the purpose alignment test. The test results show that the initial semantic similarity values between the ideal purposes and the generated descriptions were low (0.4 – 0.5), with high variability (around 0.1). However, after the alignment process, the average values increased to approximately 0.88 and, in addition, a significant reduction in dispersion was noted. This shows that the system managed to improve the accuracy of the correspondence between the ideal human purposes and processed descriptions.

Consequently, the system demonstrates its ability to bridge the semantic gap between human language and robotic representation, ensuring that ambiguous purposes are transformed into clearer and more operational instructions. This validates the system’s ability to effectively align purposes, regardless of initial variability.

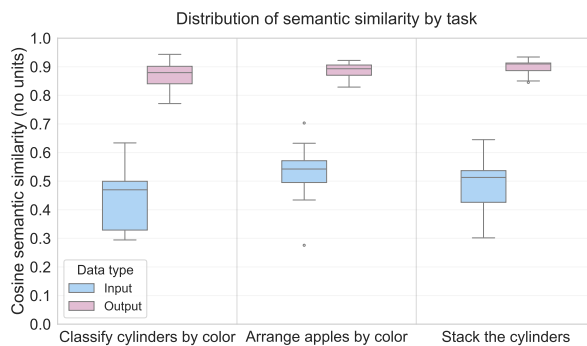


Figure 3: Box plot with the semantic similarity values of the input purposes and the output descriptions generated by the LLM system.

Figure 4 shows the percentage of valid missions and drives obtained after entering the fully defined purposes for each task in the motivational system. The results indicate that, across all three purposes, the number of valid missions and drive functions generated consistently exceeded the number of invalid ones. While the system demonstrated strong overall performance, variations were observed between tasks. The number of missions produced by the LLM also differed across tasks, a result associated with both task complexity and the formulation of the human purpose.

In general, it can be noted that the mission generator had an average performance of 89.4%, while the drive generator’s average performance was 84.1%. Hence, even though the system is not perfect, the results are satisfactory and meet the expected performance for a first approach

to the alignment problem.




Figure 4: Percentages of valid and invalid missions and drive functions per purpose.

Next, in the simulated experiment in Gazebo, the number of iterations it took to achieve the human purpose was counted. Then, based on these results, the graph shown in 5 was generated, which shows the number of iterations it took the robot to complete the purpose throughout the 10 runs. The blue curve represents a smoothed average of the attempts in the 10 runs, while the light blue shading indicates the variability of the results with a 95% confidence interval, and the purple shading indicates the ideal range of iterations needed to achieve the goal.





Figure 5: Variation in the number of iterations required to achieve the goal across trials. The smoothed average and its 95% confidence interval are shown in blue, while the purple area represents the ideal range of 2 to 3 iterations.

Also, to make sure the drives were created correctly and their evaluation was correct, the graph shown in the Figure 6 was created. This shows the drives evaluation for the first 25 iterations of the run 9, as well as the goals that were created and that rewarded the robot throughout these iterations.



(a) Drives evaluation over iterations (iterations (0-25)). (b) Rewarded goals over iterations (iterations 0-25).

Figure 6: Drive functions evaluation (above) and the goals that produced reward together with the purpose satisfaction (below).

At the beginning of the experiment, the cognitive architecture needed more iterations to fulfill the purpose (placing the red cylinder in the blue basket), because the robot did not know the correct sequence of actions and policies were executed randomly. As it received partial rewards for getting closer to the goal, the robot reduced the number of attempts until it stabilized between 2 and 4, managing to learn the task in approximately 10 repetitions. This behavior demonstrates rapid learning, which was mainly related to the simplicity of the goal, which requires only a few actions to complete.

In terms of performance, the results show a success rate of $\geq 99\%$. The failures are mainly explained by the fact that some objects left the camera's field of view, remaining out of the robot's reach. However, these cases were minimal and did not affect the validity of the experiment. In addition, the high success rate reflects that both the creation of missions and drives were adequate, as well as the distribution of weights. Finally, the ten repetitions performed showed consistent behavior and converged towards the ideal case, confirming the correct functionality of the system.

7 Conclusions

A motivational system was designed based on a series of LLMs to solve the problem of aligning and translating human purpose into robot purpose. The system was able to interpret and align different general purposes and translate them into valid missions and drive functions. Furthermore, through testing with the OSCAR robot in the Gazebo simulator, it was verified that the missions and drives proposed by the motivational system allow the robot to learn and satisfy human purpose without any direct user intervention in the architecture code. This represents a significant advance in solving the alignment problem to equip robots with Lifelong Open-ended Learning Autonomy (LOLA).

Future work consists of improving the robustness of the designed motivational system so that more complex tests can be performed and also tested on a physical robot.

Bibliography

- G. Baldassarre, R. J. Duro, E. Cartoni, M. Khamassi, A. Romero, and V. G. Santucci. A formalisation of the purpose framework: The autonomy-alignment problem in open-ended learning robots. *arXiv*, cs.RO(arXiv:2403.02514v2), Apr. 2025. URL <https://arxiv.org/abs/2403.02514v2>.
- J. A. Becerra, A. Romero, F. Bellas, and R. J. Duro. Motivational engine and long-term memory coupling within a cognitive architecture for lifelong open-ended learning. *Neurocomputing*, 452:341–354, 2021.
- E. Fallas Hernández, S. Martínez Alonso, A. Romero, J. A. Becerra Permy, and R. J. Duro. Autonomous generation of sub-goals for lifelong learning in robots. *arXiv preprint arXiv:2503.18914*, 2025. Preprint.
- I. Kotseruba and J. K. Tsotsos. 40 years of cognitive architectures: core cognitive abilities and practical applications. *Artificial Intelligence Review*, 53(1):17–94, 2020. doi: 10.1007/s10462-018-9646-y.
- A. Romero, F. Bellas, J. A. Becerra, and R. J. Duro. Motivation as a tool for designing lifelong learning robots. *Integrated Computer-Aided Engineering*, 27(3):353–372, 2020. doi: 10.3233/ICA-200633.
- A. Romero, F. Bellas, and R. J. Duro. A perspective on lifelong open-ended learning autonomy for robotics through cognitive architectures. *Sensors*, 23(3):1611, 2023.

- A. Romero, G. Baldassarre, R. J. Duro, and V. G. Santucci. H-grail: A robotic motivational architecture to tackle open-ended learning challenges. *IEEE Transactions on Cognitive and Developmental Systems*, pages 1–16, 2025. doi: 10.1109/TCDS.2025.3569352.
- J. Wang and Y. Dong. Measurement of text similarity: A survey. *Information*, 11(9):421, 2020. doi: 10.3390/info11090421. URL <https://www.mdpi.com/2078-2489/11/9/421>.