

Adaptive AI-BDI Architecture for Robotic Autonomy in Human-Centric Dynamic Environments



CIDAI-POC-2025-03

Some rights reserved. This work is available under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International ([CC BY-NC-SA 4.0](https://creativecommons.org/licenses/by-nc-sa/4.0/)).

Under the terms of this licence, you may copy, redistribute and adapt the work for non-commercial purposes, provided the work is appropriately cited, as indicated below.

In any use of this work, there should be no suggestion that CIDAI endorses any specific organization, products or services. The use of the CIDAI logo is not permitted. If you adapt the work, then you must license your work under the same or equivalent Creative Commons licence.

If you create a translation of this work, you should add the following disclaimer along with the suggested citation: “This translation was not created by the Centre of Innovation for Data tech and Artificial Intelligence (CIDAI). CIDAI is not responsible for the content or accuracy of this translation. The original English edition shall be the binding and authentic edition”.

Any mediation relating to disputes arising under the licence shall be conducted in accordance with the mediation rules of the [World Intellectual Property Organization](https://www.wipo.int/mediation/).

Suggested citation. CIDAI-POC-2025-03 // Adaptive AI-BDI Architecture for Robotic Autonomy in Human-Centric Dynamic Environments by Devis Dal Moro, Silvia Izquierdo Badiola, Adriana Maria Basbous and Magí Dalmau Moreno - Eurecat, CIDAI, 2025. Licence: [CC BY-NC-SA 4.0](https://creativecommons.org/licenses/by-nc-sa/4.0/)



1. ABSTRACT	4
2. DESCRIPCIÓ DEL PROBLEMA QUE RESOL	6
3. IMPLEMENTACIÓ DE LA SOLUCIÓ	8
4. IMPACT	17
5. BIBLIOGRAFIA	18
6. ACKNOWLEDGEMENT PROJECTS	19



1. Abstract

1.1. Objective

The primary objective of this work is to present a novel architecture for autonomous agents operating in human-centric, dynamic environments. This architecture aims to empower robots with the ability to reason dynamically, physically interact with the environment, and engage in natural, multimodal interactions. The overarching goal is to shift the paradigm from traditional robot programming to a more intuitive framework where humans specify abstract goals and robots independently determine the required steps, skills, and behaviors to achieve those goals. By leveraging continuous interactions and an evolving belief set, which accounts for objects' attributes, relationships, and human preferences, the architecture enables robots to deliver personalized and adaptive performance at the task level.

1.2. Vision

Our Cognition line team at Eurecat Robotics envisions a future where human-robot interaction is as seamless and intuitive as communication between humans. Robots will no longer require detailed programming but will instead act as collaborative agents capable of understanding abstract instructions, seeking clarifications, and autonomously planning and executing tasks. These agents will adapt to dynamic and uncertain environments while continuously improving their skills through human input and experiential learning. By incorporating human preferences and leveraging multimodal sensing, robots will deliver safe, efficient, and personalized assistance in diverse applications, ranging from industrial tasks to household chores.

1.3. Problem summary

Current robotic systems face significant limitations in dynamic and human-centric scenarios. They struggle to operate effectively under uncertainty, often requiring predefined and rigidly structured instructions. The gap between high-level human goals and low-level robotic execution poses a substantial challenge, making robot programming a labor-intensive and domain-specific task. Furthermore, the lack of adaptive mechanisms restricts robots' ability to generalize across different tasks and environments. Existing systems often fail to accommodate evolving human preferences or to capitalize on the wealth of information provided by continuous interactions. These deficiencies hinder the deployment of robots as autonomous agents capable of dynamic reasoning, interaction with the environment, and natural human interaction in real-world applications. This work seeks to address these challenges through an innovative architecture that prioritizes adaptability, autonomy, and seamless human-robot collaboration.

1.4. Solution

The proposed architecture introduces an AI-BDI (Belief-Desire-Intention) reasoning model, integrating perception, belief management, planning, and execution components to enable adaptive and autonomous behavior. It is designed and built on top of the advancements of ROS2-BDI¹[1]: a reference robotics framework to connect reasoning and planning with an agent-oriented approach but still did not allow sufficient flexibility and adaptability due to intrinsic constraints of classical planning approaches. This modern and AI-powered approach further pushes its capabilities. The architecture perceives the environment through multimodal inputs, including textual, visual, and low-level sensory data, which are processed to update the belief set. A belief manager maintains dynamic representations of objects' attributes, relationships, human preferences, and past interactions. High-level reasoning is facilitated by a Multimodal Large Language Model (MLLM) for goal distillation and desire set management, allowing the agent to interpret abstract instructions and clarify goals if needed. The grounded multi-objective problem is tackled using an online scheduler, continual planner, and action controller, ensuring real-time decision-making and adaptive revisioning of task scheduling under uncertainty. The architecture focuses on refining task-level planning and online decision-making through evolving belief sets, tailored to integrate individual preferences emerging from continuous human-robot interactions.

¹ Dal Moro, D., Robol, M., Roveri, M., Giorgini, P. (2022). Developing BDI-Based Robotic Systems with ROS2. In: Dignum, F., Mathieu, P., Corchado, J.M., De La Prieta, F. (eds) *Advances in Practical Applications of Agents, Multi-Agent Systems, and Complex Systems Simulation. The PAAMS Collection. PAAMS 2022. Lecture Notes in Computer Science()*, vol 13616. Springer, Cham. https://doi.org/10.1007/978-3-031-18192-4_9

2. Descripció del problema que resol

2.1. Background

To operate in open-world, dynamic environments, robots must overcome the limitations of traditional decision-making approaches. When generating behavior, the question arises: how can a robot decide what to do? Traditional methods such as handcrafted Behavior Trees (BTs) offer robust and explainable task execution, but they lack autonomy and adaptability. AI planning techniques, such as classical solvers, can efficiently solve well-defined tasks within a known domain but rely on strong assumptions about the world model: the initial state is fully known, the effects of actions are deterministic, and no sensing is required. Such assumptions fail in dynamic, open-world settings where knowledge is incomplete and continuously evolving.

Markov Decision Processes (MDPs) and Reinforcement Learning (RL) provide another approach, where models learn task policies from data. While MDPs are well-suited for structured, predictable tasks, they struggle to generalize in environments with imperfect knowledge, uncontrolled agents, or evolving human preferences. Additionally, the actual deployment of a system which is strongly based on black-box models might pose difficulties with respect to the formal verification and evaluation of safety measures and guarantees. This creates significant challenges in real-world applications, particularly in assistive robotics scenarios. For example, consider the simple task of bringing tea to a user. If the tea is expired (an open-world concept), or a grandchild begins interacting with the robot (an uncontrolled agent), the system must prioritize and adapt its plan. Similarly, it must handle uncertainties like tea temperature or conflicting goals (e.g., spoiling a surprise for another agent).

2.2. Why Traditional and Modern Approaches Fail

Given the provided overview, traditional methodologies struggle to address critical challenges in open-world, human-centric tasks, such as:

- Highly dynamic processes with imperfect or incomplete knowledge
- Uncontrolled agents and evolving human goals, desires, and preferences
- The need to prioritize tasks while managing tightly coupled task and motion reasoning

At the same time, directly employing modern approaches such as Large Language Models (LLMs) or Multimodal LLMs (MLLMs) for robot task planning introduces significant limitations. Auto-regressive LLMs, while powerful as general knowledge approximators, cannot inherently perform logical reasoning, structured planning, or self-verification¹ [2]. Studies reveal their poor robustness in arithmetic, logic, and abstract reasoning tasks, as well as their tendency to hallucinate outputs that lack semantic validity. Moreover, LLMs lack grounding in real-world environments, domain-specific information, and robot capabilities, making them ill-suited for adaptive, autonomous decision-making in physical settings.

¹ Valmeekam, K., Marquez, M., Sreedharan, S. & Kambhampati, S. On the planning abilities of large language models-a critical investigation. Advances in Neural Information Processing Systems 36 (2024).



2.3. Towards a Hybrid Solution

The solution lies in combining sound, traditional methodologies with the capabilities of LLMs and MLLMs. In this approach, LLMs act as universal approximators for open-world knowledge and assist in goal distillation² [3] and high-level reasoning³ [4]. These models are relevant for task planning as they can integrate real-world knowledge grounded in the robot's belief set. This belief set evolves continuously, capturing objects, their properties, relationships, and human preferences, ensuring coherent and adaptive task-level decision-making.

The architecture further incorporates Vision Language Models (VLMs) for open-world perception and enables plan adaptation through continuous interactions with humans. By enabling robots to identify what they don't know, actively seek clarifications, and update their semantic knowledge, this hybrid approach achieves robust, efficient, and semantically valid autonomous task planning.

MLLMs can therefore help in parsing and giving structure to unstructured and heterogeneous high and low-level information, enabling formal and sound reasoning techniques to integrate it into the continual planning and (re)scheduling process.

² Izquierdo-Badiola, S., Canal, G., Rizzo, C., & Alenyà, G. (2024). PlanCollabNL: Leveraging Large Language Models for Adaptive Plan Generation in Human-Robot Collaboration. In 2024 IEEE International Conference on Robotics and Automation (ICRA) (pp. 17344-17350). (Proceedings - IEEE International Conference on Robotics and Automation). IEEE. <https://doi.org/10.1109/ICRA57147.2024.10610055>

³ Kambhampati, S. et al. LLMs can't plan, but can help planning in llm-modulo frameworks. arXiv preprint arXiv:2402.01817 (2024).


3. Implementació de la solució

3.1. Arquitectura, tecnologia i dades utilitzades

The proposed architecture for autonomous agents combines a reasoning model with adaptive planning, perception, and execution capabilities. It facilitates dynamic decision-making, seamless human interaction, and continuous task-level adaptation to achieve robust and personalized performance in open-world environments.

The architecture is divided into three main subsystems: Perception, Reasoning, and Task Planning and Execution. Each subsystem plays a critical role in the flow of information and decision-making.

Figure 1 – AI-BDI framework showing the integration of Perception, Reasoning, and Planning subsystems for adaptive decision-making, dynamic belief management, and robust task execution in evolving environments.



Perception (AI-BDI Perception)

The perception block processes multimodal sensory inputs to extract meaningful representations of the environment. These inputs provide a foundation for updating the belief set, which forms the agent's understanding of the world:

- **Human Interaction:** Processes natural language inputs from humans. This allows the agent to interpret goals, clarifications, and instructions effectively.
- **Visual Perception:** Processes camera data and object recognition results, enabling the agent to identify and track objects, their relationships, and spatial attributes in a dynamic environment.
- **Low-Level Perception:** Integrates data from sensors such as tactile feedback, actuator feedback, pressure, or laser systems, providing real-time updates about the environment's physical state which might be relevant and feed as well to lower action controller modules.

These components collectively provide input to the Belief Manager, ensuring that the agent's understanding of its surroundings is grounded and continuously updated.

High-level reasoning and Knowledge Management (AI-BDI Reasoning Module)

At the heart of the architecture lies the AI-BDI (Artificial Intelligence - Belief-Desire-Intention) reasoning model, which manages high-level decision-making and dynamic goal interpretation. It includes the following components.

The central knowledge representation that evolves dynamically through continuous updates is referred as **Belief Set**. It keeps logs of the history of previous actions and outcomes, contributing to experience-based reasoning, while evolving in parallel the organization of data in a knowledge graph which counts at least for the following sets of data classes:

- **Objects' Attributes:** Properties such as object type, state (e.g., expired tea), and affordances.
- **Objects' Relationships:** Contextual and spatial relationships between objects (e.g., proximity, ownership).
- **Human States:** Captures human states such as contextual awareness information (e.g., what they're looking at, what they're holding), spatial relationships (e.g., their distance to targets), or emotional states (e.g., distracted)
- **Human Preferences:** Captures user-specific goals, desires, and behaviors through past interactions.

The **Belief Manager** is the module responsible for maintaining and updating the belief set with new perceptual data, interactions, and learned

insights. It takes leverage of MLLM(s) capabilities to parse, filter and organize such a variety of information without fully specified structure in the knowledge graph.

The **Reasoning Module** incorporates three essential components that set our system apart from traditional frameworks, enabling it to effectively address abstract goals, dynamic parameters, and unforeseen events. By leveraging the MLLM as a robust reasoning engine, guided through the designated agent role, the system engages in flexible reasoning about the environment and human inputs. These components are:

- **Goal Distillation:** The robot is provided with a description of its operational domain (e.g., “You are TIAGo, an assistive robot working in a residency for elderly individuals, responsible for tasks such as delivering the correct medications to patients and ensuring...”). Within this domain, the robot receives high-level goals in natural language (e.g., “Distribute assistive devices to the residents.”). This module’s role is to process the given goal alongside contextual environmental cues and translate it into a Desire Set. This set contains specific subgoals that can be forwarded downstream for conversion into actionable planning goals (e.g., “walking_cane provided_to resident1,” “glasses provided_to resident2”).
- **Action Refinement:** Our planning definition, unlike classical planning problems, allows for some object instances (e.g., walking_cane) to be defined without being grounded to specific instances in the environments (e.g., there are several walking canes in the scene). This is because in this way, a plan can be

computed even with partial information regarding the specific object, which can be refined during plan execution. This is what this module achieves. Before an action execution, if the argument is ambiguous, the robot will reason through the available information to ground the arguments to specific instances in the environment. This information may include feedback from both the belief set or from a triggered interaction with the user during task execution, and can include object attributes, spatial information or human preferences. In this way, the action can be refined and executed.

- **Issue Detection and Recovery:** Unexpected issues can arise during plan execution, regarding new ambiguities among the present objects, or unfeasibilities due to environment constraints (e.g., the target is not reachable) or human constraints (e.g., during a handover, the person is not looking). This module integrates an LLM with a set of tools to acquire pertinent information based on dynamically generated preconditions (e.g., detect_human_gaze(person1)). The module is provided with Actions are refined through feedback from the belief set and task execution, taking therefore in account previous data, spatial information and human preferences.

Task Execution (AI-BDI Planning and Execution Model)

The Multi-Objective Grounded Problem block bridges reasoning and action, enabling real-time planning, scheduling, and execution. This subsystem ensures that the agent's high-level goals are translated into executable actions that adapt to real-world dynamics. Key components include:

- **Low-Level Knowledge Base:** Contains pre-defined and selected domain knowledge which are obtained through handcrafted, mostly human provided pipelines for post processing and filtering low level data. It withholds low-level task skills related data, feedback and result of execution status for action primitives as well as environmental constraints and configuration settings for them. This ensures the agent can reliably execute basic operations (e.g., grasping, moving) and abstracts high-level reasoning and decision-making process from low-level data that are relevant exclusively at task execution.
- **Planner:** Using the grounded problem, dynamically generates and updates action sequences to fulfil the distilled goals. The planner incorporates real-time information from the belief set to adapt to changing conditions: revising constantly the plan(s) for forecasting and tackling beforehand possible plan failures, as well as improving the planned solution through unforeseen circumstances that occurs during plan execution¹ [5].

- **Online Scheduler:** Prioritizes and sequences multiple actions in relations to the plan, therefore the goals they aim to fulfil, managing task interdependencies and constraints. This component ensures timely execution and priorities compliance while pursuing in parallel multiple goals and possibly balancing conflicting objectives.
- **Action Controller:** Executes the refined actions in the environment while continuously monitoring feedback and success. It interacts directly with the physical components of the robot (e.g., actuators) and adapts execution strategies as needed, orchestrating adaptively the completion of the action through one (or many) handcrafted, modular and reusable Behavior Tree(s), essentially deploying the actuation of the real-world capabilities of the robot.


¹ A. Zanetti, D. D. Moro, R. Vreto, M. Robol, M. Roveri and P. Giorgini, "Implementing BDI Continual Temporal Planning for Robotic Agents," 2023 IEEE/WIC International Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT), Venice, Italy, 2023, pp. 378-382, doi: 10.1109/WI-IAT59888.2023.00062. keywords: {Aerospace electronics; Autonomous agents; Planning; Behavioral sciences; Intelligent agents; Robots; Continual planning; temporal planning; bdi; autonomous agents; robotics; ROS2},

3.2. Reptes resolts i resultats obtinguts

The implementation of this system was demonstrated in an assistive robotics scenario, with the objective of highlighting the unique capabilities that distinguish the proposed architecture from traditional frameworks. Specifically, these capabilities include [goal distillation](#), [action refinement](#), and [issue detection](#) and [recovery](#).

The diagram in Figure 2 illustrates the high-level operational flow of the proposed agent architecture, highlighting the links between goal distillation, planning, and execution with reasoning interleaved. The process begins with a high-level natural language goal provided to the agent (e.g., “Deliver the correct medicines to patients”). Through the [Goal Distillation](#) module, the agent interprets the input and generates a [Desire Set](#), a structured collection of subgoals or desires that the agent must fulfill to achieve the high-level objective. Using the [Desire Set](#), the agent extracts concrete [Planning Goals](#) and develops a plan that may include ambiguous actions or arguments (e.g., “pick medicine for patient1”), with the ambiguity addressed during the execution phase. During the execution, reasoning is highly involved to ensure robust task execution. For each action, the agent begins by checking for ambiguities and refines the action if necessary, such as clarifying which specific item or target is required, before attempting execution. If issues are detected, the agent takes corrective actions, such as prompting for clarifications or adjusting its strategy to overcome the obstacle. Once the parameters are fully refined and preconditions are satisfied, the agent executes the planned action and moves seamlessly to the next task. This dynamic feedback loop allows the agent to adapt its behavior in real-time, ensuring the system remains effective and reliable in uncertain and evolving environments.

Figure 2 – Flow of the agent’s execution pipeline, showing goal distillation, planning with ambiguous arguments, and interleaved reasoning during execution through action refinement, issue detection, recovery, and execution.




Demonstrated scenario and results.

The scenario involved the assistive robot, TIAGo, tasked with delivering the correct medications to two patients (Figure 3). After extracting goals from the human’s request, a plan is generated to execute the task (Figure 4).

Figure 3 – Tasking the robot to deliver the medicine to the patients.




Figure 4 – Goal distillation and plan generation from the natural language request.




The robot began by navigating to a medicine counter and identified a lack of specific information regarding which medicine to pick for the first patient. The system engaged in natural language interaction with a nearby human to clarify the required information. Upon obtaining the input, the robot successfully continued its plan, picked the correct medicine, and approached the patients in the room for handover (Figure 5).

Figure 5 – Action refinement. There are two medicine boxes on the medicine counter. More information is requested to pick the correct medicine for the first patient.



After approaching the patients, the robot detects an issue, as it could not identify which of the two patients is the recipient of the medicine, having no prior knowledge of their identities. The system resolves this by asking for the missing information and assigning it to the patient, allowing the robot to complete handing over the medicine (Figure 6).

Figure 6 – Issue detection and recovery. There are two patients whose names are not identified. More information is requested to resolve the issue and hand-over the medicine to the correct patient.



This process was repeated for the second patient, showcasing the robot's ability to dynamically resolve missing information and adapt its actions during execution.

The architecture's innovative use of placeholders for undefined object instances allowed the robot to compute plans despite incomplete initial information. These placeholders were resolved during execution through reasoning and interaction, bypassing significant limitations of classical planning. Additionally, advanced functionalities, such as pre-emptive detection and resolution of action feasibility issues (e.g., the patient is not identified or not looking at the robot during a hand-over), further demonstrated the robustness of the system.

The system achieved the following specific results:

1. **Goal Distillation:** The system transformed high-level goals (e.g., “Deliver the correct medicines to patients”) into grounded, specific desires that could be translated into planning goals (e.g., “handed_over adrianas_medicine adriana,” “handed_over johns_medicine john”).
2. **Action Refinement:** Ambiguous actions, such as “pick adrianas_medicine,” were refined during execution. The system sought clarifications (e.g., confirming the medicine type as paracetamol or specifying it’s the box near the bottle) and updated the action parameters before execution.
3. **Issue Detection and Recovery:** During execution, an LLM-driven agent dynamically reasons about preconditions for task completion, calling external tools to gather necessary information to identify and resolve issues. For example, before delivering a medicine, the robot ensured the human recipient was recognized, was looking at the robot, and was within reach. If any of these conditions were unmet, the system took corrective actions, such as prompting the humans to identify themselves or establish eye contact, or scanning the environment again to locate the intended recipient. The process iteratively continues until all issues are resolved.

Overall, the results demonstrated the architecture’s ability to:

- Adapt to dynamic and human-centric scenarios.
- Address challenges of incomplete or ambiguous information.
- Engage in natural, multimodal interactions to resolve execution barriers.

These achievements validate the architecture’s potential to enable robust, adaptive, and interactive task execution in complex real-world settings.

3.3. Limitacions actuals de la tecnologia

The proposed AI-BDI architecture for dynamic reasoning, task planning, and execution offers a solution to overcome the limitations associated with using either (model-based) classical planning or learning-based models in robotic applications. By integrating the strengths of both approaches into a cohesive hybrid architecture, this solution enables greater adaptivity and resilience in the agent’s performance compared to other architectures. This integration empowers the agent to effectively interpret abstract human instructions, define concrete goals, plan in the presence of uncertainties, and resolve ambiguities during execution, ultimately ensuring the successful achievement of desired outcomes.

While this hybrid solution provides significant potential for intelligent and complex robotic applications, it is still subject to the limitations of its underlying technologies. In particular, its reliance on MLLMs presents challenges, as their performance is heavily dependent on the quality and breadth of the training data. Additionally, LLMs are prone to generating hallucinations—incorrect or nonsensical outputs (e.g., using objects that do not exist in the scene or domain while generating subgoals). Furthermore, the results of MLLMs are highly sensitive to the structure of the prompts used, which affects the overall consistency and reliability of their output.

Moreover, the action refinement, issue detection, and recovery modules rely on VLMs, while the visual perception module depends on vision models. Both are limited by the quality of the training data and various environmental factors in interpreting images and videos. For example, if certain object classes were not includ-



ed in the training dataset, the model may fail to recognize or correctly classify them during inference. Variations in lighting, camera angles, occlusions, or background noise can further degrade the ability of the models to accurately interpret visual input in real-world environments. These models also face challenges with contextual understanding, handling complex scenes, and generalization across diverse scenarios—for example, they may struggle with recognizing the state of objects (e.g., open or closed), their attributes (e.g., blue, round), or the relationships between objects (e.g., above, below, left, right). Additionally, they impose a significant computational load, which can create bottlenecks in the pipeline and affect its overall performance, particularly in real-time and resource-constrained settings.

The planning domain currently relies on pre-defined domain knowledge, enabling the agent to reliably execute basic operations (e.g., grasping, moving), even in the presence of uncertainty in the world model which the agent resolves during task execution. However, this approach introduces limitations in terms of scalability to more complex environments and tasks. As the environment evolves—whether by introducing new actions, new types of objects, or adjusting action preconditions based on updated domain considerations—the domain becomes difficult to extend or adapt soundly without manual intervention. This limitation extends to the action controller, which relies on modular, reusable, handcrafted behavior trees. This limits the system's ability to easily scale to include new tasks or environments without significant re-engineering.

4. Impact

The implementation of the proposed architecture presents significant potential impacts, particularly in improving human-robot collaboration, enhancing adaptability in dynamic environments, and enabling personalized interactions in real-world human-centric applications.

Improved Human-Robot Collaboration

By enabling robots to interpret abstract instructions and adjust their actions based on real-time feedback, this solution can make interactions with robots feel more natural and intuitive. Robots equipped with such capabilities can become invaluable partners in tasks requiring close coordination, such as healthcare, elderly assistance, and collaborative industrial operations.

Adaptability in Dynamic Scenarios

The solution tackles the critical challenge of uncertainty in human-centric, dynamic scenarios. By maintaining an evolving belief set that integrates human preferences, object relationships, and environmental data, the architecture ensures robots can adapt their plans and actions to unforeseen changes related to these aspects.

Improved Personalization

Through continuous interaction, robots can tailor their behavior to the specific needs and preferences of individual users. This personalized approach enhances user satisfaction and efficiency, as the robots optimize their task execution over time. This makes the solution more effective in roles like caregiving or home assistance, where the ability to align with user routines and preferences creates a significant improvement in service quality and acceptance.



Reliable and Trustable Systems

The architecture's focus on constant feedback and structured decision-making helps reduce risks often associated with "black box" decision-making in AI systems. This ensures robots operate transparently and reliably, which is essential for tasks where safety and trust are critical, such as healthcare or public services.



5. Bibliografia

- [1] Dal Moro, D., Robol, M., Roveri, M., Giorgini, P. (2022). Developing BDI-Based Robotic Systems with ROS2. In: Dignum, F., Mathieu, P., Corchado, J.M., De La Prieta, F. (eds) Advances in Practical Applications of Agents, Multi-Agent Systems, and Complex Systems Simulation. The PAAMS Collection. PAAMS 2022. Lecture Notes in Computer Science(), vol 13616. Springer, Cham. https://doi.org/10.1007/978-3-031-18192-4_9
- [2] Valmeekam, K., Marquez, M., Sreedharan, S. & Kambhampati, S. On the planning abilities of large language models-a critical investigation. Advances in Neural Information Processing Systems 36 (2024).
- [3] Izquierdo-Badiola, S., Canal, G., Rizzo, C., & Alenyà, G. (2024). Plan-CollabNL: Leveraging Large Language Models for Adaptive Plan Generation in Human-Robot Collaboration. In 2024 IEEE International Conference on Robotics and Automation (ICRA) (pp. 17344-17350). (Proceedings - IEEE International Conference on Robotics and Automation). IEEE. <https://doi.org/10.1109/ICRA57147.2024.10610055>
- [4] Kambhampati, S. et al. Llms can't plan, but can help planning in llm-modulo frameworks. arXiv preprint arXiv:2402.01817 (2024).
- [5] A. Zanetti, D. D. Moro, R. Vreto, M. Robol, M. Roveri and P. Giorgini, "Implementing BDI Continual Temporal Planning for Robotic Agents," 2023 IEEE/WIC International Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT), Venice, Italy, 2023, pp. 378-382, doi: 10.1109/WI-IAT59888.2023.00062.



6. Acknowledgement projects

Pla de centres (Acció):

- Flagship Ronous2
- Traça IA Generativa

Projecte Adapta:

- Proyecto financiado por el programa Proyectos de I+D+i en líneas estratégicas – TransMisiones 2023 del Ministerio de Ciencia, Innovación y Universidades. Referencia del proyecto MIG-20232046.

CIDAI

Centre of Innovation
for Data tech
and Artificial Intelligence

C/Bilbao, 72 Edif. A. 08005 - Barcelona
Tel. +34 93 7419 100
info@cidai-catalonia-ai.eu
www.cidai.eu

Socis



Membres

