Towards mixed-initiative planning systems:
building upon automated planning and plan recognition systems to construct solutions collaboratively

Supervisors: Caroline CHANEL, Baptiste PELLETIER, Alexandre ALBORE

in collaboration with Felipe MENEGUZZI (PUCRS - Brazil, University of Aberdeen, King's College)

1 Context and Relevancy

In scenarios as natural resource preservation, surveillance, mapping or search and rescue, it is relevant to deploy efficiently heterogeneous teams of autonomous agents to cover large environments – see Fig. 1a. These teams can be composed by human and robotic agents with different characteristics and capabilities. For instance, for natural resource monitoring and preservation, it could be useful to coordinate a heterogeneous team of agents composed by ground and aerial robots and human beings, to survey areas and to detect–and keep away–undesirable actors. Another example could be the deployment of collaborative exploration between humans and robots to assess the damage caused by a natural disaster while searching for survivors. Note that in such missions some environment properties or parameters could be ill-defined, uncertain, or unknown before mission starts, e.g., the number of locations of interest to visit, the number of survivors, the locations of hazards or traps, or the number and the locations of (undesirable) actors.

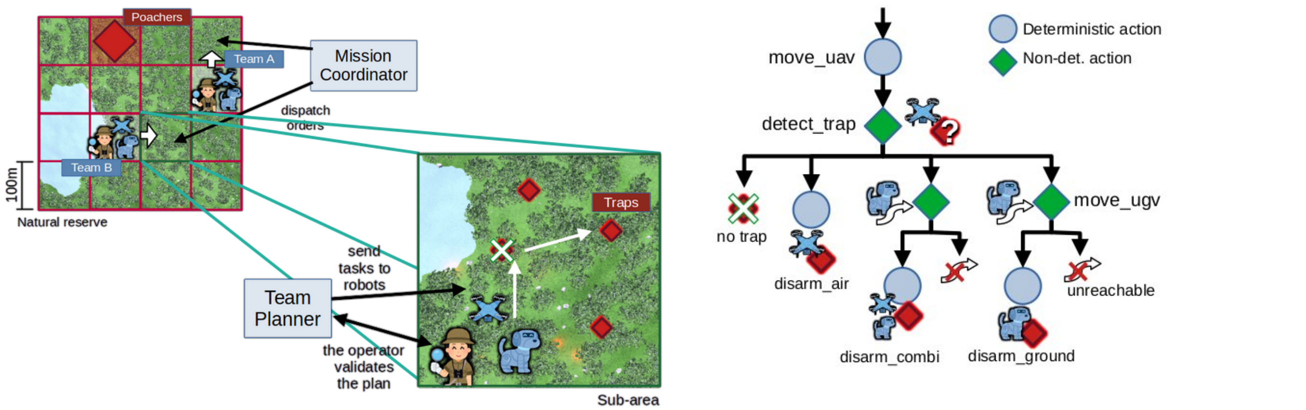

(a) A heterogeneous team.

(b) Conditional plan to check a given location.

Figure 1: Illustrative example of a natural preservation mission employing groups of heterogeneous teams.

Assuming complete knowledge of environment properties and dynamics, which enables setting up a precise planning problem, it is possible to leverage automated planning tools to coordinate the team actions in order to achieve mission goals. The resulting plan determines which action should be executed by each agent – see Fig. 1b. Formally, automated planning can be defined as the automated process of choosing and organizing actions while anticipating their effects (Ghallab et al., 2025). Automated planning is often used for long- or short-term sequential decision-making problems, for single- or multi- actors climbing on the category of combinatorial decision-making problems. Automated planning can be used for deciding which task to assign to each agent, or which sequence of actions should be executed by the team of agents in view of satisfying goals and/or optimizing one (or multiple) metric(s).

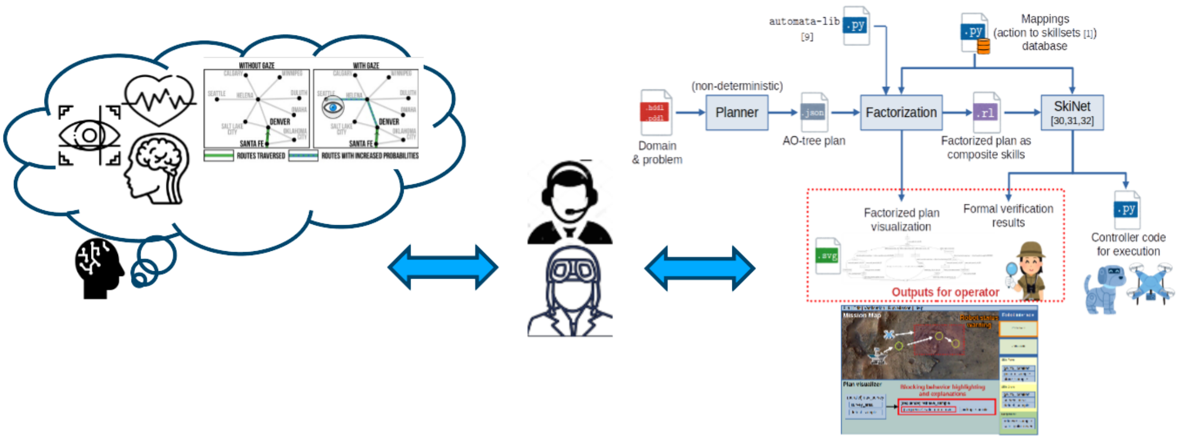

Figure 2: Illustrative example of mixed-initiative planning framework building upon automated planning tools and PGR systems.

Mixed-initiative planning (Cox and Zhang, 2005; Bresina et al., 2005) is a particular planning approach to design a planning engine system that attempts to collaborate with a human operator to construct planning solutions. This planning paradigm, allowing the implementation of decision support systems, has gained attention in particular when the planning problem, if solved only by the human, would require a huge cognitive effort, or if solved only by the machine, would suffer from lack of human expert knowledge. Examples of mixed-initiative planning can be cited, as for example, crew-computer collaboration for space mission scheduling (Zheng et al., 2023), manned-unmanned teaming operations planning (Maier et al., 2024; Schmitt and Schulte, 2015), mixed-initiative systems for planning and scheduling in industry (Prévôt et al., 2024), or robotic planning and control (Chrpa et al., 2015; Gianni et al., 2011).

Despite these recent works, literature is still scarce, and several challenges remain, as for instance: How to account with human preferences regarding parts of the solution plan without treating it as a constraint in the initial planning problem? How to express plan solutions and to explain the reasoning employed by the planning engine? How to implement initiatives from the planning system, *i.e.* plan suggestions, that could well fit to human preferences or eventually push the human to accept a suggestion not necessarily aligned with his initial thoughts?

To push towards flexible mixed-initiative planning systems, a first step would consist on enabling the system to recognize human intentions, goals, or plan preferences among possible options in some contexts. Interestingly, Plan, Activity, and Intent Recognition systems (Sukthankar et al., 2014; Singh et al., 2020) employ distinct inference techniques to recognize the goals of agents under observation. Following (Meneguzzi and Pereira, 2021; Amado et al., 2024), Plan-and Goal Recognition (PGR) are strongly related to automated planning because the system in charge of the recognition, employs abstract deductive reasoning to infer the actual plan being carried out by the observed agent, or to infer the most likely desired goal, from a sequence of observations.

In this context, this project aims to build upon automated planning tools and goal/plan recognition systems towards the design of flexible planning engine capable of collaborating with human operators to iteratively draw up solutions for uncertain or even ill-defined planning problems.

This original research extends previous work, such as (Sohrabi et al., 2019) that exploited automated planning and plan recognition to design planning advisors. The mixed-initiative planning application cases we target consists in planning missions where human operators are required to validate, or to take part on planning and on the execution of the plan – see Fig. 2. For instance we target natural resource monitoring and preservation, or collaborative manned-unmanned team operations, where the human operator is in charge of, or eventually accompanied by, ground or aerial robots capable of autonomous navigation behavior, information acquisition through embedded sensors or even physical interaction with the environment they evolve.

2 Research Problem

Automated planning is commonly used for deciding a sequence of actions while anticipating their effects in view of satisfying goals and/or optimizing one (or multiple) metric(s). It tackles long- or short-term sequential decision-making problems, for single- or multi- actors. Among the automated planning methods of the literature, Classical planning is usually based on an goal-satisfying approach, where the states variables describe properties of the agent(s) and the environment they evolve, and actions describe operations that change the properties of the state. From an initial state, the planning engine is in charge of finding the sequence of actions (operations) that satisfy goal requirements. Going behind the classical planning approach, some other planning systems can deal with non-deterministic action effects and/or partial state observability, such as (Bryce, 2006; Albore et al., 2009; Albore and Geffner, 2009; Komarnitsky and Shani, 2016; Pereira et al., 2022; Messa and Pereira, 2023). These planning paradigms and algorithms propose strong foundations for tackling real-life problems where actions can lead to different outcomes and where some state properties can not be observed (all the time).

For example, let us consider a team illustrated in Figure 1a, composed by a human operator, a ground robot, and a drone, which is asked to inspect a large area of interest. The drone can fly over the large area, and *possibly detect* locations of interest to be further inspected by the team. This `fly_over_area` action could define the number of locations of interest where supposed hazards or traps where detected. Then the human agent or the robots can be assigned to inspect one of the locations of interest to closely observe the supposed trap. In order to execute an inspection action, the ground agent first needs to move to the position where the supposed trap was localized, however the trap *could not be reachable* by this agent. The words *possibly detected* or *could not be reachable* characterize two types of non-deterministic effects of the `fly_over_area` or `move_ugv` actions, respectively, leading to different course of actions. Moreover, `detect_trap` attests on the presence of partial observability, because each time a given location is inspected, the information gathered only serves to this location, and does not inform about the other locations of interest.

Despite the usefulness of automated planning engines for solving complex and combinatorial problems, the challenge this previous example brings remains not truly addressed in literature where one aims to produce a plan for an ill-defined and/or uncertain problem, leading to limitations on planning problem definition or to consider a (too) complex planning problem. Some properties (number of locations of interest to be visited), or actions results can only be observed at the place. To deal with these challenges, some authors adopted the planning-and-acting paradigm (Chanel et al., 2019; t'Hooft et al., 2016; Maliah et al., 2022). Where the planning engine produces a first (partial) plan, then while executing the first action, it supposes its results to (re-)plan towards goal requirements. This planning-and-acting architecture builds on the so-called online planning paradigm (Ghallab et al., 2025).

From a human perspective, the inputs one may give to the planning engine do not naturally capture some problems aspects, leading the human user to simplify some parts of the problem, or to make use of tricks to get around some abstract description limitations. For instance, despite the huge use of the STRIPS-based[1] based description languages, as Planning Domain Description Language (PDDL)[2] and variants, to describe planning domains and associated problems, it is not straightforward to model the effects of a `fly_over_area` action regarding that the number of locations of interest is not known in advance.

Moreover, the solutions obtained may not account with human preferences, because they were not developed for it. As said, planners search to satisfying goals and/or optimize metrics while respecting some constraints. For instance, taking our example again, the human agent, which is in place with the robots, could prefer (but not constraint the system) to visit some locations in a given order, what could be cost-inefficient regarding a distance metric optimized by the planner. Maybe the human operator could be willing to accept some costly solution to better fit her preferences. As another example, the human, which has experience on the terrain, knows which locations are more likely to the reachable by the ground robots. But this experience-based knowledge is hard to be feed into actual planning engines without expressing it as a constraint in precedence of actions, needing the user to

[1] The Stanford Research Institute Problem Solver (STRIPS) is an automated planner developed by Richard Fikes and Nils Nilsson in 1971.

[2] see `https://fareskalaboud.github.io/LearnPDDL/` for a PDDL introduction

be an expert on planning tools to express such insights.

Closely to how a human being proceeds to solve complex problems, hierarchical planning is a paradigm of automated planning approaches in which the planning engine searches a decomposition of complex tasks into smaller sub-tasks using predefined methods (Ghallab et al., 2025). In particular, Hierarchical Task Networks (HTN) provide an intuitive methodology for specifying high-level instructions on how robots and agents should perform tasks, while also giving the planner enough flexibility to choose the lower-level steps and their ordering (Lallement et al., 2014). However, the hierarchy imposed by the decompositions - a given library of domain control knowledge - are handwritten, and possibly restricts which sub-tasks can be part of the solution regarding the initial state of the system, actions, and goals.

Despite that, the decomposition control over the feasible solutions, usually evaluated in a AND/OR tree, makes hierarchical planning more expressive than the classical goal-satisfying STRIPS-based approaches (Höller, 2023). Such an expressivity, that hierarchical planning brings, makes it a promising approach in cases where the planning system collaborates with humans, in particular if the human operator is asked to interact with the planning engine to achieve desired solutions (Alnazer et al., 2022; Lallement et al., 2014). Regarding Plan-and Goal Recognition (PGR) and Hierarchical models, previous works have proposed approaches linking recognition systems and plan tree grammars (Geib and Goldman, 2009). And, only recently (Höller et al., 2018) proposed to combine the expressive HTN representation with the off-the-shelf planning systems for PGR. Their work was able to handle large models with many possible goals achieving high recognition rates.

However, classical hierarchical planning does not accounts with non-deterministic actions, i.e., actions for which the effects can lead to different outcomes, or partial observability regarding the state of the system, i.e, unknown state variables status or values. Very few studies start to address non-deterministic effects in the HTN context (Yousefi and Bercher, 2024; Halacheva, 2022), but several challenges still remain relating to planning domains with cycles, with metrics to be optimized, and in particular accounting with non-deterministic action affects or partial observability.

Moreover, despite the HTN expressivity advantages from a user point of view, how to feed human preferences into HTN solvers remains an open question. Let's take again our example of the human, ground and aerial robots team which is asked to visit some large area of interest (see Figure 1b). Each `check_loc` high level action can be decomposed differently: depending on the chosen agent that does `move` action, depending if a trap is detect or not. If a trap is detect, `disarm_trap` can also be decomposed differently depending on the type of trap and on the agents needed for that action. In conclusion these actions can have non-deterministic effects and can furnish only partial knowledge about the true (hole) state. Thus, we believe there is still place for hierarchical-based planning extensions working with non-deterministic effects and partial observability.

In conclusion, within mixed-initiative planning paradigm, this project aims to investigate approaches that could be applied, and at least, partially addressing limitations, such as:

L1 Hierarchical planning seems to be a promising approach for mixed-initiative planning. However there is still place to develop hierarchical planners enabling to stop resolution and delivering a partial plan, or handling non-deterministic effects and/or partial observability,

L2 Regarding plan and goal recognition systems, former works proposed PRG systems based on hierarchical planners. Nevertheless, there is still place for coupling PRG and hierarchical planners in the scope of mixed-initiative planning.

L3 In mixed-initiative planning, there is place to develop flexible planning engines able to account with human inputs such as possible parts of solution plans, that could be identified through PGR systems, instead of treating human inputs as a constraint in the initial planning problem;

Building upon these limitations, we expect to lay the foundations towards flexible mixed-initiative planning systems. Such systems would fit solutions to human preferences, or eventually push the human operator to accept plan suggestions.

3 Objectives and Work Methodology

Objective 1 : State-of-the-art reviewing In objective 1, the student will review the literature works, to identify automated planning tools and plan-and goal recognition systems that could be applicable as it. Note, it could lead to planning problem simplifications. Additionally, it is expected from this reviewing work to identify promising mixed-initiative planning frameworks, deeply analyzing recent works such as (Alnazer et al., 2022; Prévôt et al., 2024; Maier et al., 2024). It would enable to define in which extent we could rely on them to combine automated planning and PRG systems.

Objective 2 : Plan and goal recognition (PGR) system Following Objective 1, the student will evaluate in Objective 2 the applicability of the state-of-the-art goal and plan recognition approaches (Singh et al., 2020), probably based in a simplified model (or non-deterministic or HTN model). The main goal is to a first pipeline based on PRG allowing to integrate human choices iteratively.

Objective 3 : Partial observable, non-deterministic, hierarchical planning (POND-HP) Following Objective 1, the student will evaluate on-the-shelf automated planning tools, probably based in a simplified planning problem, using or non-deterministic or HTN approaches in an online planning paradigm. We, expected to identify opportunities for further developments. In particular, we expect to automatize planning problem definition and to exploit insights from plan compression tools (Pelletier et al., 2025) for improve the solving process.

Objective 4 : Towards mixed-initiative planning systems In Objective 4, we expect the student to start to combine the approaches studied in Objectives 2 and 3 in order to sketch up a mixed-initiative planning engine. This activity will necessitate specific research on interaction modalities, visualization tools, and translation or encoding means to be used by the mixed-initiative planning engine. In this sense, we plan to at least initiate the software architecture developments during the present project. But we acknowledge that it is not realistic to expect final implementation results on this activity given the project's duration.

4 Candidate's profile

Technical Skills Candidates should showcase scientific aptitude, programming proficiency in languages like java, C#, Python, or C++, and familiarity with automated planning techniques. However, those keen to learn coding will also be considered.

Personal Attributes A passion for cutting-edge research, the capability to work independently, adaptability, and exemplary communication skills are essential.

Laboratory Location DCAS Department
ISAE-SUPAERO (Institut Supérieur de l'Aéronautique et de l'Espace)
10 Av. Marc Pélegrin, 31400 Toulouse

Compensation Interns will receive a stipend of 4.35€ net per hour, equating to approximately 609€ per month.

Duration Six months, commencing in February or March 2026, with potential opportunities for a subsequent PhD (financial support already acquired).

Application to be sent to supervisors: Prospective candidates should forward their CV and cover letter to:
Caroline Chanel at `caroline.chanel@isae-supaero.fr`
Alexandre Albore at `alexandre.albore@onera.fr`
Baptiste Pelletier at `baptiste.pelletier@isae-supaero.fr`

References

Albore, A. and Geffner, H. (2009). Acting in partially observable environments when achievement of the goal cannot be guaranteed. In *Proc. of ICAPS Workshop on Planning and Plan Execution for Real-World Systems*.

Albore, A., Palacios, H., and Geffner, H. (2009). A translation-based approach to contingent planning. In *IJCAI*, volume 9, pages 1623–1628.

Alnazer, E., Georgievski, I., Prakash, N., and Aiello, M. (2022). A role for htn planning in increasing trust in autonomous driving. In *2022 IEEE International Smart Cities Conference (ISC2)*, pages 1–7. IEEE.

Amado, L., Shainkopf, S. P., Pereira, R. F., Mirsky, R., and Meneguzzi, F. (2024). A survey on model-free goal recognition. In *IJCAI 2024: 33rd International Joint Conference on Artificial Intelligence*.

Bresina, J. L., Jónsson, A. K., Morris, P. H., and Rajan, K. (2005). Mixed-initiative planning in mapgen: Capabilities and shortcomings. In *Workshop on Mixed-Initiative Planning and Scheduling*.

Bryce, D. (2006). Pond: The partially-observable and non-deterministic planner. *Proceedings of the International Conference on Automated Planning and Scheduling*, page 58.

Chanel, C. P. C., Albore, A., T'hooft, J., Lesire, C., and Teichteil-Königsbuch, F. (2019). Ample: An anytime planning and execution framework for dynamic and uncertain problems in robotics. *Autonomous Robots*, 43:37–62.

Chrpa, L., Pinto, J., Ribeiro, M. A., Py, F., Sousa, J., and Rajan, K. (2015). On mixed-initiative planning and control for autonomous underwater vehicles. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1685–1690.

Cox, M. T. and Zhang, C. (2005). Planning as mixed-initiative goal manipulation. In *ICAPS*, pages 282–291.

Geib, C. W. and Goldman, R. P. (2009). A probabilistic plan recognition algorithm based on plan tree grammars. *Artificial Intelligence*, 173(11):1101–1132.

Ghallab, M., Nau, D., and Traverso, P. (2025). *Acting, Planning, and Learning*. Cambridge University Press.

Gianni, M., Papadakis, P., Pirri, F., and Pizzoli, M. (2011). Awareness in mixed initiative planning. In *AAAI Fall Symposium: Robot-Human Teamwork in Dynamic Adverse Environment*.

Halacheva, A.-M. (2022). Htn planning under uncertainty for sustainable buildings. B.S. thesis, Universität Stuttgart.

Höller, D. (2023). *Hierarchical Planning: Expressivity Analysis, Solving Techniques, and Problem Compilations*. PhD thesis, Universität Ulm.

Höller, D., Behnke, G., Bercher, P., and Biundo, S. (2018). Plan and goal recognition as htn planning. In *2018 IEEE 30th International Conference on Tools with Artificial Intelligence (ICTAI)*, pages 466–473. IEEE.

Komarnitsky, R. and Shani, G. (2016). Computing contingent plans using online replanning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30.

Lallement, R., De Silva, L., and Alami, R. (2014). Hatp: An htn planner for robotics. In *2nd ICAPS Workshop on Planning and Robotics, PlanRob 2014*.

Maier, S., Kiam, J. J., and Schulte, A. (2024). Adaptive mission planning: Evaluation of a hybrid cognitive mixed-initiative planning assistant in manned-unmanned teaming operations. In *2024 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 3492–3499.

Maliah, S., Komarnitski, R., and Shani, G. (2022). Computing contingent plan graphs using online planning. *ACM Transactions on Autonomous and Adaptive Systems (TAAS)*, 16(1):1–30.

Meneguzzi, F. R. and Pereira, R. F. (2021). A survey on goal recognition as planning. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI), 2021, Canadá.*

Messa, F. and Pereira, A. G. (2023). A best-first search algorithm for fond planning and heuristic functions to optimize decompressed solution size. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 33, pages 277–285.

Pelletier, B., Chanel, C. P., and Albore, A. (2025). A formal factorization approach of non-deterministic plans: Application to an anti-poaching robotic mission scenario. In *Workshop on Agents and Robots for reliable Engineered Autonomy*, pages 72–92. Springer.

Pereira, R. F., Pereira, A. G., Messa, F., and De Giacomo, G. (2022). Iterative depth-first search for fond planning. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 32, pages 90–99.

Prévôt, E., Ménard, M.-A., Gaudrcault, J., and Ali, M. B. (2024). Human-machine mixed-initiative system for integrated planning and scheduling in the manufacturing industry. In *IISE Annual Conference. Proceedings*, pages 1–6. Institute of Industrial and Systems Engineers (IISE).

Schmitt, F. and Schulte, A. (2015). Mixed-initiative mission planning using planning strategy models in military manned-unmanned teaming missions. In *2015 IEEE International Conference on Systems, Man, and Cybernetics*, pages 1391–1396.

Singh, R., Miller, T., Newn, J., Velloso, E., Vetere, F., and Sonenberg, L. (2020). Combining gaze and ai planning for online human intention recognition. *Artificial Intelligence*, 284:103275.

Sohrabi, S., Katz, M., Hassanzadeh, O., Udrea, O., Feblowitz, M. D., and Riabov, A. (2019). Ibm scenario planning advisor: Plan recognition as ai planning in practice. *Ai Communications*, 32(1):1–13.

Sukthankar, G., Geib, C., Bui, H. H., Pynadath, D., and Goldman, R. P. (2014). *Plan, activity, and intent recognition: Theory and practice*. Elsevier.

t'Hooft, J., Cabaniols, C. L., and Chanel, C. C. (2016). Proactive planning and execution strategies with multiple hypotheses. In *11th National Conference on "Software and Hardware Architectures for Robots Control" & Quatrièmes Journées Architectures Logicielles pour la Robotique Autonome, les Systèmes Cyber-Physiques et les Systèmes Auto-Adaptables (SHARC 2016).*

Yousefi, M. and Bercher, P. (2024). Laying the foundations for solving fond htn problems: grounding, search, heuristics (and benchmark problems). In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, pages 6796–6804.

Zheng, J., Shelat, S. M., and Marquez, J. J. (2023). Facilitating crew-computer collaboration during mixed-initiative space mission planning. In *SpaceCHI 3.0 A Conference for Human-Computer Interaction for Space Exploration*.