Integrating an online HTN planner with external legal and ethical checkers*

Hisashi Hayashi[1] , **Yousef Taheri**[2] , **Kanae Tsushima**[3] ,
Gauvain Bourgne[2] , **Jean-Gabriel Ganascia**[2] , **Ken Satoh**[3]

[1]Advanced Institute of Industrial Technology, Tokyo, Japan
[2]Sorbonne University, Paris, France
[3]Research Organization of Information and Systems, Tokyo, Japan

Abstract

We present a planning agent with an online HTN planner integrated with external legal and ethical checkers. Three interaction modes are introduced and compared to enable efficient, adaptive compliance during online planning. The framework is evaluated using international data transfer scenarios.

1 Introduction

Online data sharing and processing raise numerous legal and ethical concerns, including privacy, safety, and algorithmic bias. These challenges are especially critical in international contexts, where systems must comply with diverse legal frameworks and ethical norms. Addressing this complexity requires not only normative awareness but also computational mechanisms that can adapt in real time to jurisdiction-specific constraints.

We present a modular framework in which legal and ethical checkers function as autonomous components, possibly deployed on separate servers. This enables scalable, context-aware specialization aligned with the localized nature of normative reasoning.

To enhance the efficiency of compliant and adaptive decision-making during online planning, we define and compare three interaction modes between the planner and the checkers. These modes differ in communication frequency and synchronization methods, and are evaluated in terms of their impact on computational and communication efficiency. We assess the proposed framework through scenarios involving international data transfer and processing.

2 Overall Architecture

Figure 1 illustrates the overall architecture, which consists of a planning agent, a legal checker, an ethical checker, and an action executor. The planning agent employs an online HTN planner to generate plans based on current beliefs and adjusts them dynamically during execution. It issues action commands to the executor and updates its beliefs and plans using feedback. The legal checker assesses each planned action for legal compliance, while the ethical checker compares alternative plans to select the one that best satisfies ethical norms. The action executor carries out the actions and returns outcomes to the planning agent. It also monitors the environment for unexpected changes, such as server activity, safety, or occupancy levels, and notifies the planning agent accordingly.

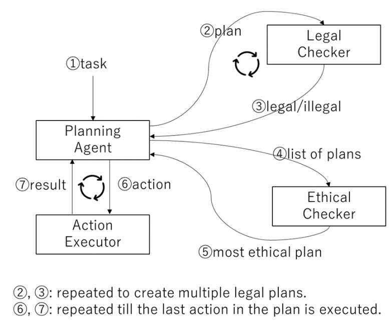

Figure 1: The overall architecture.

3 Interaction Modes

This section presents three interaction modes that define how the planning agent interacts with the legal and ethical checkers. To improve system efficiency, it is important to reduce the number of interactions and computation time, while ensuring that plans are based on the most recent information. The three modes are: *default*, *subscription*, and *all-subscription* mode. Each of these modes is detailed in the following subsections.

3.1 Default mode

The default mode offers the simplest interaction pattern and serves as the baseline for comparison. Figure 2a illustrates this configuration. In this mode, fluents representing the dynamic state of the world are stored as beliefs in the planning agent's database. When evaluating a plan, the legal and ethical checkers query the planning agent to retrieve the truth values of relevant fluents.

Whenever an action is executed or a fluent changes, the planning agent replans and updates multiple candidate plans. The legal checker assesses the legality of each updated plan, and the ethical checker selects the most ethical plan among those deemed legally valid.

*This extended abstract summarizes our paper (Hayashi et al. 2025) that was presented at ICAART 2025.

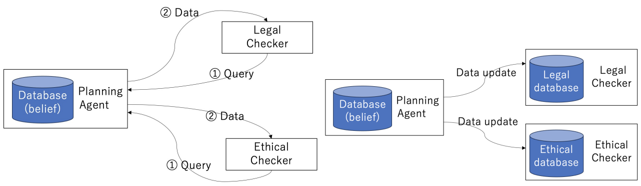

(a) Default mode. (b) Subscription mode.

Figure 2: Interaction modes between modules.

3.2 Subscription mode

The *subscription* mode improves efficiency by minimizing unnecessary communication between the planning agent and the legal and ethical checkers, while preserving compliance accuracy. In this mode, each checker specifies a set of *subscribed fluents*—fluents whose truth values are directly relevant to its normative evaluations. As shown in Figure 2b, checkers maintain synchronized local copies of these fluents. By maintaining local copies, checkers avoid frequent queries to the planning agent, significantly reducing interaction overhead. Furthermore, the planning agent can skip legal and ethical evaluations when no subscribed fluents are affected by the current action, thereby lowering computational cost.

3.3 All-subscription mode

The *all-subscription* mode is a special case of the subscription mode in which the legal and ethical checkers subscribe to all fluents. As a result, there is no need to explicitly specify subscribed fluents.

4 Use case

To demonstrate the characteristics and efficiency of our proposed approach, we apply it to eight scenarios involving data transfer and processing across a network of distributed nodes. The network consists of multiple interconnected nodes, each representing a segment of an organization that is either inside or outside the EU.

The system completes a task by locating the required data, transferring it to a processing node, executing the processing, and delivering the result to the requesting node. Within our architecture, a planner module generates multiple candidate plans—each corresponding to a different possible route for data transfer and processing within the network.

Table 1 summarizes these scenarios, each characterized by different types of dynamic changes that influence planning decisions via the legal and ethical checkers. These changes affect planning decisions through interactions with the legal and ethical checkers as depicted in Figure 1.

5 Evaluation

The results shown in Table 2 clearly demonstrate that the *subscription* mode significantly reduces both CPU time and interaction frequency compared to the default mode. Moreover, even when compared with the *all-subscription* mode—which monitors all fluents without discrimination—the targeted *subscription* mode yields notable improvements in computational and communication efficiency, while retaining architectural flexibility.

Table 1: Scenario summary and causes of replanning.

Scenario	Cause of replanning
basecase	No unexpected change
precond-replan-1	The route became unavailable
precond-replan-2	The route's availability changes twice
cost-ethical-replan	A cost parameter changes
ethical-replan	Ethical priorities are updated
legal-replan-1	Permission changes
legal-replan-2	Permission changes twice
legal-ethical-replan	Two replans: one legal, one ethical

Table 2: CPU time (top) and number of interactions (bottom).

CPU time (s)			
Scenario	Default	All-subscription	Subscription
basecase	1.38	1.37	**0.46**
precond-replan-1	2.38	2.30	**1.32**
precond-replan-2	2.83	2.75	**1.60**
cost-ethical-replan	2.31	2.28	**1.15**
ethical-replan	2.71	2.68	**1.32**
legal-replan-1	3.40	3.65	**0.83**
legal-replan-2	86.02	86.04	**49.70**
legal-ethical-replan	19.35	19.30	**11.01**

Number of interactions			
Scenario	Default	All-subscription	Subscription
basecase	16916	84	**25**
precond-replan-1	46038	121	**35**
precond-replan-2	50357	140	**51**
cost-ethical-replan	24617	104	**35**
ethical-replan	32760	121	**43**
legal-replan-1	41069	157	**47**
legal-replan-2	92655	337	**73**
legal-ethical-replan	60612	216	**107**

6 Conclusion

This paper presented an implementation of a planning agent that integrates an online planner with modular legal and ethical checkers. We examined three interaction modes and demonstrated that the fluent subscription technique significantly reduces both computational and communication overhead. These results highlight the effectiveness of the proposed architecture for efficient and scalable deployment in dynamic, norm-sensitive environments.

Acknowledgments

This work was funded by JST AIP Trilateral AI Research (JPMJCR20G4); JST Mirai Program (JPMJMI23B1); JSPS KAKENHI (22H00543 and 21K12144); Strategic Research Projects grant from ROIS, Research Organization of Information and Systems (2025-SRP-25); Agence Nationale de la Recherche (ANR, French Research Agency) project RECOMP (ANR-20-IADJ-0004).

References

Hayashi, H.; Taheri, Y.; Tsushima, K.; Bourgne, G.; Ganascia, J.-G.; and Satoh, K. 2025. A new planning agent architecture that efficiently integrates an online planner with external legal and ethical checkers. In *Proceedings of the 17th International Conference on Agents and Artificial Intelligence (ICAART 2025)*, volume 1, 263–272.