

DEMO: Visualizations for an Explainable Planning Agent

Tathagata Chakraborti¹ and Kshitij P. Fadnis² and Kartik Talamadupula² and Mishal Dholakia²
Biplav Srivastava² and Jeffrey O. Kephart² and Rachel K. E. Bellamy²

¹ Computer Science Department, Arizona State University, Tempe, AZ 85281 USA

tchakra2 @ asu.edu

²IBM T. J. Watson Research Center, Yorktown Heights, NY 10598 USA

{ kpfadnis, krtalamad, mdholak, biplavs, kephart, rachel } @ us.ibm.com

Abstract

In this demonstration, we report on the visualization capabilities of an Explainable AI Planning (XAIP) agent that can support human in the loop decision making. Imposing transparency and explainability requirements on such agents is crucial for establishing human trust and common ground with an end-to-end automated planning system. Visualizing the agent’s internal decision making processes is a crucial step towards achieving this. This may include externalizing the “brain” of the agent: starting from its sensory inputs, to progressively higher order decisions made by it in order to drive its planning components. We demonstrate these functionalities in the context of a smart assistant in the Cognitive Environments Laboratory at IBM’s T.J. Watson Research Center.

Recently, there have been concerted efforts towards making the outputs of planning processes more palatable to human decision makers – e.g. eXplainable AI Planning (XAIP) (Fox, Long, and Magazzeni 2017; Langley et al. 2017). In this paradigm, emphasis is laid on the qualities of *trust*, *interaction*, and *transparency* of an AI system. One of the key features that an XAIP agent must support is visualization. The planning community has recently made a concerted effort to support the visualization of key components of the end-to-end planning process: from the modeling of domains (Bryce et al. 2017); to assisting with plan management (Izygon, Kortenkamp, and Molin 2008); and beyond (Sengupta et al. 2017; Benton et al. 2017). For an end-to-end planning system – which goes from lower level sensory data (e.g. vision, speech) to progressively higher level decision-making capabilities (planning, plan recognition) – this becomes even more challenging. This is due to the fact that the system’s state is determined by information at different levels of abstraction, all of which is being coalesced in the course of decision making. It is in this spirit that we present *Mr. Jones*, a set of visualization capabilities for an XAIP agent that assists with human-in-the-loop decision-making in an instrumented meeting space. These capabilities are essential toward the establishment of common ground in the collaborative planning process.

Introducing Mr. Jones – *Mr. Jones* (Chakraborti et al. 2017c), situated in the CEL – the Cognitive Environments Laboratory – at IBM’s T.J. Watson Research Center is de-

signed to embody the key properties of a proactive assistant while fulfilling the properties desired of an XAIP agent. Similar to (Manikonda et al. 2017), we divide the responsibilities of *Mr. Jones* into two processes (c.f. Figure 1) – **Engage**, where plan recognition techniques are used to identify the task in progress; and **Orchestrate**, which involves active participation in the decision-making process via real-time plan generation, visualization, and monitoring.

Mind of Mr. Jones – The externalization of the “mind” of *Mr. Jones* – i.e. the various processes that feed the different capabilities of the agent (c.f. Figure 2) – consists of five widgets. The largest widget on the top represents the probability distribution that indicates the confidence of *Mr. Jones* in identifying the task being collaborated on, along with a button that displays the provenance of each such belief. The information used as provenance is generated directly from the plans used internally by the recognition module (Ramirez and Geffner 2010) and justifies why, given the model of the underlying planning problems, these tasks look likely in terms of plans that achieve those tasks. The system is adept at handling uncertainty in its inputs (in coming up with an explanatory plan, it has announced likely assignments to unknown agents in its space).

Below this is a set of four widgets, each of which give users a peek into an internal component of *Mr. Jones*. The first of them (top left) presents a wordcloud representation of *Mr. Jones*’s belief in each of the tasks; the size of the word representing that task corresponds to the probability associated with that task. The second widget (top right) shows the agents that are recognized as being in the environment currently – this information is used by the system to determine what kind of task is more likely. This information is obtained from four independent camera feeds that give *Mr. Jones* an omniscpective view of the environment; this information is represented via snapshots (sampled at 10-20 Hz) in the third widget (bottom left). Finally, the fourth widget (bottom right) represents a wordcloud based summarization of the audio transcript of the environment. This transcript provides a succinct representation of the things that have been said in the environment in the recent past via the audio channels. The interface thus provides a (constantly updating) snapshot of the various sensory and cognitive organs

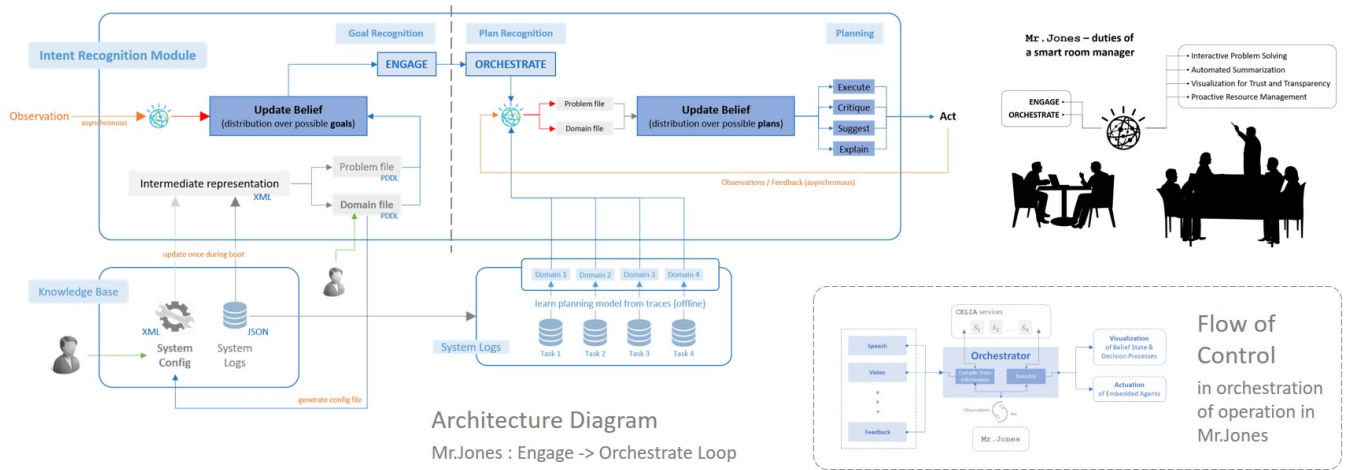


Figure 1: The building blocks of Mr. Jones – the two main components *Engage* and *Orchestrate* situate the agent proactively in a decision support setting with human decision makers in the loop. The top right inset shows the roles of Mr. Jones as a smart room orchestrator. The bottom right inset illustrates the flow of control – each service runs in parallel and asynchronously to maintain anytime response.

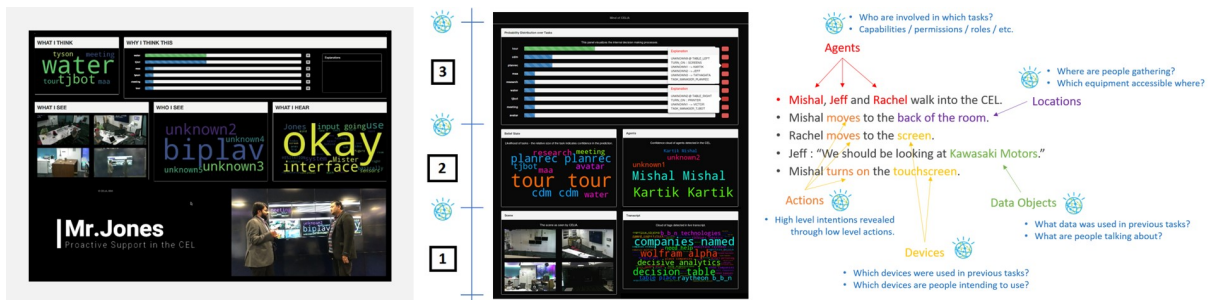


Figure 2: Snapshot of the mind of Mr. Jones. A video of the system in action can be viewed at <https://youtu.be/ZEHxCKodeGs>.

associated with Mr. Jones – the eyes, ears, and mind of the CEL. This is organized at increasing levels of abstraction –

- [1] *Raw Inputs* – These show the camera feeds and voice capture (speech to text outputs) as received by the system. These help in externalizing what information the system is working with at any point of time, and can be used in debugging at the input level if the system makes a mistake or in determining whether it is receiving enough information to make the right decisions. It is especially useful for an agent like Mr. Jones, which is not embodied in a single robot or interface but is part of the environment as a whole; in such cases, it is difficult to attribute specific events and outcomes to the agent.
- [2] *Lower level reasoning* – The next layer deals with the first stage of reasoning over the raw inputs – *What are the topics being talked about? Who are the agents in the room? Where are they situated?* This helps a user identify what knowledge is being extracted from the input layer and fed into the reasoning engines. It increases the situational awareness of agents by visually summarizing the contents of the scene at any point of time.
- [3] *Higher level reasoning* – Finally, the top layer uses information extracted at the lower levels to reason about ab-

stract tasks in the scene. It visualizes the outcome of the plan recognition process, along with the provenance of the information extracted from the lower levels (agents in the scene, their positions, speech intents, etc.). This layer puts into context the agent’s current understanding of the processes in the scene.

In addition to this, we support a plan visualization tool *Fresco* (Chakraborti et al. 2017a), that builds on recent work in top-K planning (Katz et al. 2018) and model-based plan explanations (Chakraborti et al. 2017b) to provide a concise visualization of a plan. Specifically, we cast the plan visualization problem as a plan explanation process and filter out irrelevant information from the underlying model given a specific planning problem.

A detailed description of the system can be accessed at <https://arxiv.org/abs/1709.04517>.

Acknowledgements. A significant part of this work was initiated and completed while Tathagata Chakraborti was an intern at IBM’s T. J. Watson Research Center during the summer of 2017. The continuation of his work at Arizona State University is supported by an IBM Ph.D. Fellowship.

References

- Benton, J.; Smith, D.; Kaneshige, J.; and Keely, L. 2017. CHAP-E: A plan execution assistant for pilots. In *Proceedings of the Workshop on User Interfaces and Scheduling and Planning*, UISP 2017, 1–7.
- Bryce, D.; Bonasso, P.; Adil, K.; Bell, S.; and Kortenkamp, D. 2017. In-situ domain modeling with fact routes. In *Proceedings of the Workshop on User Interfaces and Scheduling and Planning*, UISP 2017, 15–22.
- Chakraborti, T.; Fadnis, K. P.; Talamadupula, K.; Dholakia, M.; Srivastava, B.; Kephart, J. O.; and Bellamy, R. K. E. 2017a. Visualizations for an explainable planning agent. *CoRR* abs/1709.04517.
- Chakraborti, T.; Sreedharan, S.; Zhang, Y.; and Kambhampati, S. 2017b. Plan Explanations as Model Reconciliation: Moving Beyond Explanation as Soliloquy. In *IJCAI*.
- Chakraborti, T.; Talamadupula, K.; Dholakia, M.; Srivastava, B.; Kephart, J. O.; and Bellamy, R. K. 2017c. Mr. Jones – Towards a Proactive Smart Room Orchestrator. *AAAI Fall Symposium on Human-Agent Groups*.
- Fox, M.; Long, D.; and Magazzeni, D. 2017. Explainable Planning. In *First IJCAI Workshop on Explainable AI (XAI)*.
- Izygon, M.; Kortenkamp, D.; and Molin, A. 2008. A procedure integrated development environment for future spacecraft and habitats. In *Space Technology and Applications International Forum*.
- Katz, M.; Sohrabi, S.; Udrea, O.; and Winterer, D. 2018. A Novel Iterative Approach to Top-k Planning. In *International Conference on Automated Planning and Scheduling (ICAPS)*.
- Langley, P.; Meadows, B.; Sridharan, M.; and Choi, D. 2017. Explainable Agency for Intelligent Autonomous Systems. In *AAAI/IAAI*.
- Manikonda, L.; Chakraborti, T.; Talamadupula, K.; and Kambhampati, S. 2017. Herding the crowd: Using automated planning for better crowdsourced planning. *Journal of Human Computation*.
- Ramirez, M., and Geffner, H. 2010. Probabilistic plan recognition using off-the-shelf classical planners. In *AAAI*.
- Sengupta, S.; Chakraborti, T.; Sreedharan, S.; and Kambhampati, S. 2017. RADAR - A Proactive Decision Support System for Human-in-the-Loop Planning. In *AAAI Fall Symposium on Human-Agent Groups*.