Combining LSTMs and Symbolic Approaches for Robust Plan Recognition

Extended Abstract

Leonardo Rosa Amado
PUCRS
Porto Alegre, Brasil
leorosaamado@gmail.com

Ramon Fraga Pereira
Sapienza University of Rome
Rome, Italy
pereira@diag.uniroma1.it

Felipe Meneguzzi
PUCRS
Porto Alegre, Brasil
felipe.meneguzzi@pucrs.br

ACM Reference Format:
Leonardo Rosa Amado, Ramon Fraga Pereira, and Felipe Meneguzzi. 2021. Combining LSTMs and Symbolic Approaches for Robust Plan Recognition: Extended Abstract. In *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), Online, May 3–7, 2021*, IFAAMAS, 3 pages.

1 INTRODUCTION

Plan recognition is the task of inferring the actual plan an observed agent is performing to achieve a goal, given domain theory and a partial, and possibly noisy, sequence of observations [3, 14, 22]. Applications include natural language processing [6], elder-care [5], multi-agent systems [4, 19], collaborative problem-solving[10, 11], epistemic problems [17] and more [7, 18]. Real-world plan recognition problems impose limitations on the quality and quantity of the observations, which may be missing or faulty from silent errors in the sensors [22]. While recent approaches to goal and plan recognition have substantially improved performance under partial observability and noisy conditions [12, 13, 20, 21, 23], dealing with these problems remains a challenge.

Recent work on goal and plan recognition use machine learning to assist planning-based approaches [23] in modeling domains. Such techniques [9, 16] yield robust models capable of accurate predictions with missing or noisy data. Thus inspired, we develop a novel approach to solve both goal and plan recognition tasks simultaneously by combining planning and machine learning techniques to mitigate problems of low and faulty observability. On the machine learning side, we use a set of plans to train a predictive statistical model of the most likely next states given a set of state observations. We combine such predictive models with landmark heuristics from state-of-the-art goal recognition techniques [13] to predict the states relevant towards a goal hypothesis, given a set of observations. We use such predictive models to address the two most common problems in observations in goal and plan recognition: missing and faulty (noisy) observations. First, the predictive model allows us to fill in missing observations and rebuild the complete sequence of states from a plan. Second, while completing missing observations, we can detect faulty (noisy) observations and build state sequences that do not necessarily comply with all the observations, if some are not consistent with the planning model.

We evaluate our approach in standard, handcrafted, classical planning domains, as well as in automatically generated domains in latent space [1], showing its effectiveness at recognizing both plans and goals. We compare our approach to classical plan recognition approaches [14, 15] in the literature, measuring the optimality of the computed plans and the accuracy of the predicted goals in scenarios with missing and faulty observations. Standard plan (goal) recognition approaches struggle to achieve high precision in latent space domains, as their spread (i.e., returned goals) in such domains is very high. Our approach achieves high precision in most domains, excelling in latent space with a precision increase of up to 60%, including in problems with noisy observations. Finally, we show that our approach can compute complete optimal plans is most problems, resulting in a reliable way to perform plan recognition.

2 PREDICTIVE PLAN RECOGNITION (PPR)

We improve plan recognition performance in low and noisy observability problems by combining both machine learning techniques and domain theory techniques into an approach called Predictive Plan Recognition (PPR). PPR infers plans even with very low observability and noisy observations using machine learning and classical planning techniques combined, achieving robust plan recognition. Our approach computes a sequence of intermediary states achieved by a plan π given a plan recognition problem $\Pi^{\Omega}_{\pi} = \langle \Xi, \mathcal{I}, \mathcal{G}, \Omega \rangle$, where $\Xi = \langle \mathcal{F}, \mathcal{A} \rangle$ is planning domain model ($\mathcal{F}$ is a set of fluents and $\mathcal{A}$ is a set of actions), $\mathcal{I}$ represents the initial state, $\mathcal{G}$ is the set of possible goals (including a correct goal $G^* \in \mathcal{G}$, unknown to the observer), and a sequence of observations Ω.

To compute such sequence of states for each possible goal, our approach rebuilds the sequence of states induced by a plan by iterating through the sequence of observations Ω and filling in any gaps due to partial observability. Using a domain model Ξ, we check if a sequence of state observations in the observations Ω_s is valid by evaluating whether a transition between each pair of consecutive state observations is valid, starting from the initial state $\mathcal{I}$ and the first observation. If the sequence of states is impossible, we conclude that an observation is missing at the point of the invalid transition of the observation trace, thus, we must predict the next state that should have been observed at this point. Figure 1 provides an overview of our architecture to solve goal and plan recognition problems using PPR. The approach PPR returns not only the agent's intended goal, but also a sequence of states (i.e., plan) that possibly achieves the goal. The green box labelled ComputeSequence($\mathcal{I}, \mathcal{A}, \Omega, G, \lambda$) represents the process of computing a plan for a given goal.

To compute the sequence of states (plan) of an agent by filling the missing observations, we must devise a way to predict the most

Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), U. Endriss, A. Nowé, F. Dignum, A. Lomuscio (eds.), May 3–7, 2021, Online.

likely next state. Hence, filling in the missing observations consists of a two step process that outputs a state s' as the next most likely state. First, we use a trained machine learning model $\mathcal{M}$ to compute the probability of each known stat being the most likely next state s' given a sequence of states S_π. This model computes a probability distribution of each state being the next state. Then, we select the k most likely next states from the probability distribution given by the machine learning model that are achievable after the current sequence of states S_π towards a goal G. From the k most likely states, we select the state with the highest landmark achievability as the next state s', which steers our sequence of states S_π closer to the goal condition G. We use k to weight the trade off between the probability distribution of the model $\mathcal{M}$ and the closeness to the goal G towards which we are recognizing the plan.

Once we are capable of predicting the most likely single missing state in a sequence of observations, we can use this prediction to fill in any number of missing observations. Thus, our approach fills in all missing observations for a sequence of observations, computing the sequence of states induced by the plan towards a single goal in the function $\textsc{ComputeSequence}(\mathcal{I}, \mathcal{A}, \Omega, G, \lambda)$ of Figure 1. We develop an approach to compute a plan for a possible goal G using a machine learning predictive model and a set of landmarks. To generate a sequence of states S that a plan π achieves for the goal condition G, we use the initial state $\mathcal{I}$ as the starting point. Then, we extract landmarks for the goal condition G. To extract the landmarks we use the algorithm proposed by Hoffmann et al.. We concatenate the list of state observations Ω with the possible goal G creating a new list $\hat{\Omega}$. We iterate through this list of observations $\hat{\Omega}$, checking if o (an observation from $\hat{\Omega}$) is the result of a valid transition from the last known valid state of $\mathcal{S}$ (denoted as $\mathcal{S}_{|\mathcal{S}|}$). If exists a valid transition between these two states, we add the observation to the sequence of states and move to the next observation of $\hat{\Omega}$. Otherwise, we use the method described before to compute the most likely next state, combining the machine learning model and the computed landmarks to generate the next state s'. We keep predicting states and adding them to the sequence of states S, until the predicted state s' can achieve the next observation in $\hat{\Omega}$. We repeat this process for every observation, including the desired goal G, until we compute a valid sequence of states in the model that can achieve G. The algorithm stops when it achieves the desired goal G during the prediction phase, returning the current sequence of states $\mathcal{S}$. If a particular goal hypothesis is very unlikely, trying to fill in missing observations from the last known valid transition necessarily induces a substantially suboptimal plan. In practice, incorrect goal hypotheses will induce plans much longer than the one for the correct hypothesis in domains where the state space is connected, or infinite plans that never reach it. Thus, to prevent the algorithm from generating such plans, we stop trying to complete the plan for a goal hypothesis G if during the prediction process we predict λ (a threshold) states consecutively that are unable to achieve the current observation o, returning the current sequence of states. This threshold can be any heuristic value, estimating the maximum length of a plan. Here, we use the length of the longest plan in the training dataset as λ.

We solve the problem of plan recognition by applying this process of generating a sequence of states to all goals $G \in \mathcal{G}$. In Figure 1, we can see in the white box that we have computed a sequence of states S_G for each goal hypothesis G. To decide which goal is the correct one, we compare the predicted sequence of states S_G for each goal. First, we discard any goal G that is not in the last state in the predicted sequence S_G. Then, we rank the remaining S_G based on their compliance with the set of observations Ω, selecting the sequence of states S_G that complies with most of the observations Ω. If there is a tie, we select the sequence of states S_G that has the best cumulative score during the predictions, normalized by the number of predictions, as the most probable sequence of states S_G, and its goal G as the most likely goal. Thus, we predict a single sequence of states (from which we can derive a plan), for a single goal as the most likely goal and plan the agent is pursuing, solving the problem of both goal and plan recognition.

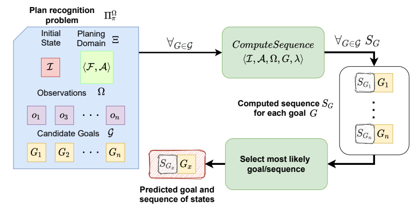

Figure 1: PPR Overview.

To deal explicitly with noisy observations, we develop a variation of the function $\textsc{ComputeSequence}(\mathcal{I}, \mathcal{A}, \Omega, G, \lambda)$. To do so, we detect when an observation constitutes noise, and ignore it in the predicted plan. We assume an observation o_i to be noisy if we can predict a state induced by subsequent observations (i.e. o_j, s.t. $j > i$) that can be reached by valid transitions between the last valid inferred state before o_i and o_j. Thus, we can build a valid plan skipping the observation o_i, now assumed to be noisy, by adding such state o_j to the current sequence of states S and continue iterating from observation o_{j+1}.

3 DISCUSSION AND FUTURE WORK

We developed an approach for goal and plan recognition that combines machine learning statistical prediction with domain knowledge within classical planning techniques. Our resulting approach achieves very high precision both in handcrafted and automatically generated plan recognition domains. We empirically show that our approach is capable of computing plans with very low observability (up to 90% missing observations) and noisy observations (up to 20% noise). While our approach does rely on data for its machine learning component, experiments show the amount of data to be much smaller than that needed to generate the learned domains [1, 2], which already necessitates data. Our machine learning model is simple enough that the same network architecture works for all domains, and thus, tuning the machine learning model is not really required for our approach to work. Indeed, we could use an entirely symbolic substitute for the neural network and predict the next states looking ahead one step in a heuristic search algorithm.

Acknowledgements: Felipe acknowledges support from CNPq with projects 407058/2018-4 (Universal) and 302773/2019-3 (PQ). Ramon acknowledges support from ERC Advanced Grant WhiteMech (No. 834228) and EU ICT-48 2020 TAILOR (No. 952215).

REFERENCES

[1] Leonardo Amado, Ramon Fraga Pereira, João Paulo Aires, Maurício Magnaguagno, Roger Granada, and Felipe Meneguzzi. 2018. Goal Recognition in Latent Space. In Proceedings of the International Joint Conference on Neural Networks (IJCNN).

[2] M Asai and A Fukunaga. 2017. Classical Planning in Deep Latent Space: From Unlabeled Images to PDDL (and Back). In AAAI Workshop on Knowledge Engineering for Planning and Scheduling.

[3] Sandra Carberry. 2001. Techniques for Plan Recognition. User Modeling and User-Adapted Interaction 11 (03 2001), 31–48.

[4] Stephen Cranefield, Felipe Meneguzzi, Nir Oren, and Bastin T. R. Savarimuthu. 2016. A Bayesian approach to norm identification. In Proceedings of the Twenty Second European Conference on Artificial Intelligence. 622 – 629. https://doi.org/10.3233/978-1-61499-672-9-622

[5] Christopher W. Geib. 2002. Problems with intent recognition for elder care. In Proceedings of the AAAI Conference on Artificial Intelligence.

[6] Christopher W Geib and Mark Steedman. 2007. On natural language processing and plan recognition. In Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI).

[7] Roger Granada, Ramon Fraga Pereira, Juarez Monteiro, Rodrigo Barros, Duncan Ruiz, and Felipe Meneguzzi. 2017. Hybrid Activity and Plan Recognition for Video Streams. In The AAAI 2017 Workshop on Plan, Activity, and Intent Recognition.

[8] J Hoffmann, J Porteous, and L Sebastia. 2004. Ordered Landmarks in Planning. Journal of Artificial Intelligence Research 22, 1 (April 2004), 215–278.

[9] S. Kim, H. Zhang, R. Wu, and L. Gong. 2011. Dealing with noise in defect prediction. In Proceedings of the International Conference on Software Engineering (ICSE).

[10] Jean Oh, Felipe Meneguzzi, Katia Sycara, and Timothy J. Norman. 2013. Prognostic normative reasoning. Engineering Applications of Artificial Intelligence 26, 2 (2013), 863 – 872. https://doi.org/10.1016/j.engappai.2012.12.006

[11] Jean Oh, Felipe Meneguzzi, Katia P Sycara, and Timothy J Norman. 2011. An Agent Architecture for Prognostic Reasoning Assistance. In IJCAI. 2513–2518.

[12] Ramon Fraga Pereira, Nir Oren, and Felipe Meneguzzi. 2017. Landmark-Based Heuristics for Goal Recognition. In Proceedings of the AAAI Conference on Artificial Intelligence.

[13] Ramon Fraga Pereira, Nir Oren, and Felipe Meneguzzi. 2020. Landmark-Based Approaches for Goal Recognition as Planning. Artificial Intelligence 279 (2020), 103217.

[14] Miquel Ramírez and Hector Geffner. 2009. Plan Recognition as Planning. In Proceedings of International Joint Conference on Artifical Intelligence (IJCAI).

[15] Miquel Ramírez and Hector Geffner. 2010. Probabilistic Plan Recognition Using Off-the-Shelf Classical Planners. In Proceedings of the AAAI Conference on Artificial Intelligence.

[16] Jeffrey C. Schlimmer and Richard H. Granger. 1986. Incremental Learning from Noisy Data. Machine Learning 1, 3 (March 1986), 317–354.

[17] Maayan Shvo, Toryn Q. Klassen, Shirin Sohrabi, and Sheila A. McIlraith. 2020. Epistemic Plan Recognition. In Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS).

[18] Maayan Shvo and Sheila A. McIlraith. 2020. Active Goal Recognition. In The AAAI Conference on Artificial Intelligence.

[19] Maayan Shvo, Shirin Sohrabi, and Sheila A. McIlraith. 2018. An AI Planning-Based Approach to the Multi-Agent Plan Recognition Problem. In Canadian Conference on Artificial Intelligence.

[20] S Sohrabi, A Riabov, and O Udrea. 2017. Planning-based Scenario Generation for Enterprise Risk Management. In AAAI Workshop on Knowledge Engineering for Planning and Scheduling.

[21] Shirin Sohrabi, Anton V. Riabov, and Octavian Udrea. 2016. Plan Recognition as Planning Revisited. In Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI).

[22] Gita Sukthankar, Robert P Goldman, Christopher Geib, David V Pynadath, and Hung Hai Bui. 2014. Plan, Activity, and Intent Recognition: Theory and Practice. Elsevier.

[23] Hankz Hankui Zhuo, Yantian Zha, Subbarao Kambhampati, and Xin Tian. 2020. Discovering Underlying Plans Based on Shallow Models. ACM Transactions on Intelligent Systems and Technology 11, 2 (2020).