Deep Learning approaches to Goal Recognition

Doctoral Consortium

Lorenzo Serina
Università degli Studi di Brescia
Brescia, Italy
lorenzo.serina@unibs.it

ABSTRACT

The Goal Recognition task is a problem in the field of Artificial Intelligence, with significant applications across different domains. The current state-of-the-art approach involves transforming the Goal Recognition problem into one that can be solved using classical planning algorithms. This work aims to extend the boundaries of existing state-of-the-art systems by applying deep learning models to the Goal Recognition task.

KEYWORDS

Goal Recognition, Automated Planning, Deep Learning, LSTM, BERT

ACM Reference Format:
Lorenzo Serina. 2025. Deep Learning approaches to Goal Recognition: Doctoral Consortium. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 3 pages.

1 INTRODUCTION

Our research focusses on approaching the **Goal Recognition** task using **Deep Learning** techniques. Goal Recognition is the task of inferring an agent's goal, from a set of hypotheses, given a model of the environment dynamic, and a sequence of observations of such agent's behavior.

In the literature, several systems have been proposed to solve goal recognition problems [3]. The state-of-the-art approach is based on the transformation of a Goal Recognition problem into a problem solvable by classical planning algorithms [4, 5, 11]. This approach requires the domain knowledge, consisting of the model of all the actions that can be performed by the agent and a description of the initial state of the world in which the agent performs these actions.

Our objective is to extend the boundaries of current state-of-the-art systems through the application of deep learning models for the Goal Recognition task. In particular, we are trying to address the goal recognition problem as a multi-classification task where each goal fluent can be considered true or false, as developed in [1]. This approach can provide several advantages like higher accuracy, faster recognition times and it can operate with less information, since the domain and the model of the problems are not required.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

2 RESEARCH

Our main focus has been the development of systems based on Deep Learning architectures, in particular **Long Short Term Memory** (LSTM) and **BERT**, which can solve offline or online goal recognition instances.

2.1 Offline and Online Goal Recognition

Our current approach, as many others in the literature, belongs to the "goal recognition over a domain theory" field [4, 6, 9], into which the available knowledge consists of an underlying model of the agent's behavior and its environment.

An *instance of the GR problem* in a given domain is then specified by: an initial state I of the agent ($I \subseteq F$), where F consists of all the possible state fluents of the considered domain; a sequence $O = \langle o_1, .., o_n \rangle$ of observations ($n \geq 1$), where each o_i is an action in A performed by the agent, and a set $\mathcal{G} = \{G_1, .., G_m\}$ ($m \geq 1$) of possible goals of the agent, where each G_i is a set of fluents over F that represents a partial state. The observations form a trace of the full sequence π of actions performed by the agent to achieve a goal G^*.

We refer to *Offline Goal Recognition* when the observation trace consists of actions in π, ordered as they appear in π; note that the selected actions can be non-consecutive in π.

Whereas, we refer to *Online Goal Recognition* when the observation trace is revealed incrementally and consists of a prefix of π; therefore, starting from the first action in π, at each execution step the subsequent action is added to the observation sequence. For this task, we followed the assumptions stated in [10].

2.2 BERT for Offline GR

For the *Offline Goal Recognition* task, we trained a **BERT** model on different classical planning domains. This model can be used as a *foundation model* for the considered domain and this work is published in [8]. Once obtained a foundation model, we fine-tuned it to perform the goal recognition task.

2.2.1 Architecture. BERT [2] is a Transformer-based model designed to be pre-trained using unlabelled text. Once pre-trained, we can fine-tune it with just one additional output layer. To perform the goal recognition task, a *Selection Mechanism* is attached to the net output layer. The architecture is depicted in Figure 1.

2.2.2 Pre-Training. Our BERT model is trained using a slight variation of the typical masked language modelling task that we call *Planning Language Modelling* (PLM). We substitute a certain percentage of actions with the special token [MASK] and give this incomplete sequence as input. The BERT model has to predict the masked tokens from the overall context of the action sequence.

Figure 1: Architecture for goal recognition using a fine-tuned BERT model.

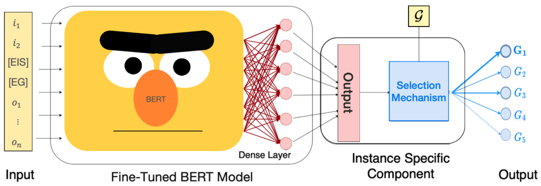

2.2.3 *Fine-Tuning.* For fine-tuning, we took a random subset of the training set and we derived the observation sequences by randomly selecting actions from the plans (preserving their relative order). The generated data set consists of pairs $(\langle I, O\rangle, G^*)$ where I is the initial state of the planning problem, O is a sequence of observed actions obtained by sampling a plan π, and G^* is the hidden goal corresponding to the goal of the planning problem solved by π. Given a goal recognition instance, the output of the i-th neuron $\bar{o}_i$ corresponds to the i-th fluent f_i and the activation value of $\bar{o}_i$ gives a rank for f_i being true in the goal. These ranks are fed to the *Selection Mechanism*, to select the predicted goal.

2.2.4 *Results.* The tests conducted in the Offline Goal Recognition Task show an excellent understanding of the domains and how they work. We tested it on different classical planning domains like DEPOTS, DRIVERLOG and ZENOTRAVEL, with mean improvement of 8%, 9% and 4% w.r.t. the state-of-the-art.

2.3 CLERNet for Online GR

For the *Online Goal Recognition* task, we developed the **Causal Link Enanched Recurrent Network** (CLERNet)[7]. We implemented a system based on a **Long Short-Term Memory Network** (LSTM), to process the sequence of actions observed, trained with a special labelling technique based on Causal Links.

2.3.1 *Dataset creation and Labelling Strategies.* We address Online Goal Recognition as a many-to-many sequential task, so we need a label for each observed action in input. In classical planning, different actions can contribute to different goal fluents, and this kind of information can be very important for our Neural Component to model the connections among actions and goals. In Automated Planning, it is provided by **Causal Links**, and we extracted and used them as labels.

2.3.2 *Architecture.* Our architecture is depicted in Figure 2. It consists of three main parts: the *Neural Component* (in blue), the *Aggregation Component* (in dotted green) and the *Selection Component* (in dashed red). The input of the first component is an instance of Online Goal Recognition, and its output is a prediction vector for each input observation $o_i \in O$. First, an LSTM layer processes the sequence of observations. The output of each cell is then passed to a time-distributed feed-forward layer, which has $N = |F_G|$ output neurons that provides a prediction for each observation we receive in input. Therefore, considering the prediction made after seeing the i-th observed action o_i, the output of the j-th neuron y_i^j corresponds to the score associated to the j-th fluent f_j (fluents are lexically ordered), and the activation value of y_i^j gives a score for f_j belonging to the agent's goal.

Figure 2: Architecture of our online goal recognition system.

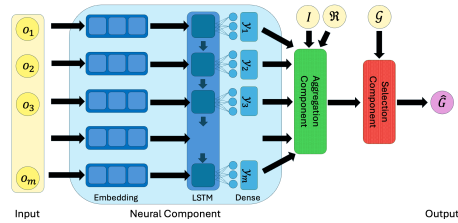

The Aggregation component takes as input the score vectors generated by the Neural Component, the initial state I of the agent and the information related to which fluents in F_G are mutually exclusive. Then, it calculates a unique aggregated score that summarises all this information.

Finally, the Selection Component receives in input the aggregated score and the set of candidate goals $\mathcal{G}$, and it outputs the predicted goal $\hat{G}$.

2.3.3 *Results.* The tests conducted in the Online Goal Recognition task show that this approach recognises the correct goal for a huge number of different problems earlier w.r.t. other approaches. We calculate two key metrics: Ranked First (RF) and Convergence (CV). RF represents the mean number of correct predictions per instance. In contrast, CV refers to the mean number of accurate predictions when models decide on a goal and maintain this choice throughout the entire recognition process. For instance, consider the sequence $<$ True, False, True, True $>$; here, RF equates to 3/4, and CV equals 2/4, because the model changes prediction on the second observation. We evaluated the system on different classical planning domains like DEPOTS, LOGISTICS and SATELLITE, achieving on average an improvement of 3% on RF and 5% on CV wrt. the state-of-the-art.

3 FUTURE WORK

In the future, we aim to study the robustness of goal recognition models, considering both model-based and data-driven, on cyber security attacks, like adversarial attacks and stealth attacks. In order to make it possible, we are studying deception techniques to move the acting agent in order to confuse and make the model predict the wrong goal. Then we will approach the problem of improving the models in order to be more robust to the studied attacks.

An alternative development might involve moving away from the PDDL framework. The systems showcased represent an initial step towards integrating Deep Learning models into Automated Planning tasks. The findings demonstrate these models' strong ability to comprehend tasks and planning domains via PDDL formalism. However, the strength of Deep Learning lies in its capability to discern patterns within data without relying on formal rule-based learning, by modeling probability distributions. Consequently, these models could be applied to various domains and data types, such as graphs or sequences of images illustrating an actor's movement in an environment.

REFERENCES

[1] Mattia Chiari, Alfonso Emilio Gerevini, Francesco Percassi, Luca Putelli, Ivan Serina, and Matteo Olivato. 2023. Goal recognition as a deep learning task: the GRNet approach. In *Proceedings of the International Conference on Automated Planning and Scheduling*, Vol. 33. 560–568.

[2] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).

[3] Felipe Meneguzzi and Ramon Fraga Pereira. 2021. A Survey on Goal Recognition as Planning. In *Proceedings of IJCAI 2021*.

[4] Ramon Fraga Pereira, Nir Oren, and Felipe Meneguzzi. 2020. Landmark-based approaches for goal recognition as planning. *Artif. Intell.* 279 (2020).

[5] Miquel Ramírez and Hector Geffner. 2009. Plan Recognition as Planning. In *Proceedings of IJCAI 2009*.

[6] Miquel Ramírez and Hector Geffner. 2010. Probabilistic Plan Recognition Using Off-the-Shelf Classical Planners. In *Proceedings of AAAI 2010*. AAAI Press.

[7] Lorenzo Serina, Mattia Chiari, Alfonso Emilio Gerevini, Luca Putelli, and Ivan Serina. 2025. Towards Efficient Online Goal Recognition through Deep Learning. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*. International Foundation for Autonomous Agents and Multiagent Systems / ACM.

[8] Lorenzo Serina, Mattia Chiari, Alfonso E Gerevini, Luca Putelli, Ivan Serina, et al. 2022. A Preliminary Study on BERT applied to Automated Planning. In *CEUR WORKSHOP PROCEEDINGS*, Vol. 3345. CEUR-WS.

[9] Franz A. Van-Horenbeke and Angelika Peer. 2021. Activity, Plan, and Goal Recognition: A Review. *Frontiers Robotics AI* 8 (2021).

[10] Mor Vered, Gal A. Kaminka, and Sivan Biham. 2016. Online goal recognition through mirroring: humans and agents. In *Fourth Annual Conference on Advances in Cognitive Systems (Advances in Cognitive Systems)*, Kenneth Forbus, Tom Hinrichs, and Carrie Ost (Eds.). Cognitive Systems Foundation.

[11] Mor Vered, Ramon Fraga Pereira, Mauricio Cecilio Magnaguagno, Gal A. Kaminka, and Felipe Meneguzzi. 2018. Towards Online Goal Recognition Combining Goal Mirroring and Landmarks. In *AAMAS*. International Foundation for Autonomous Agents and Multiagent Systems Richland, SC, USA / ACM, 2112–2114.