

LatRec: Recognizing Goals in Latent Space (Student Abstract)

Leonardo Amado,[†] Felipe Meneguzzi[‡]

Pontifical Catholic University of Rio Grande do Sul (PUCRS), Brazil
Graduate Program in Computer Science, School of Technology

[†]leonardo.amado@acad.pucrs.br

[‡]felipe.meneguzzi@pucrs.br

Abstract

Recent approaches to goal recognition have progressively relaxed the requirements about the amount of domain knowledge and available observations, yielding accurate and efficient algorithms. These approaches, however, assume that there is a domain expert capable of building complete and correct domain knowledge to successfully recognize an agent's goal. This is too strong for most real-world applications. We overcome these limitations by combining goal recognition techniques from automated planning, and deep autoencoders to carry out unsupervised learning to generate domain theories from data streams and use the resulting domain theories to deal with incomplete and noisy observations. Moving forward, we aim to develop a new data-driven goal recognition technique that infers the domain model using the same set of observations used in recognition itself.

Introduction

Goal recognition is the task of identifying the desired goal of an agent by observing its behavior in an environment. Plan recognition approaches aim to identify the specific plan to which the observed agent has committed to perform to achieve a particular goal. Most approaches to goal and plan recognition require a substantial amount of domain knowledge (Sukthankar et al. 2014, Chapter 1). However, such approaches assume that a domain expert can provide a correct and complete domain knowledge for the algorithm to successfully recognize an agent's goal. The dependence on an expert limits the applicability of such algorithms in many real-world domains, so recent work has relaxed such requirements (Pereira, Pereira, and Meneguzzi 2019; Pereira et al. 2019; Zhuo 2019). We overcome this limitation in LATREC by building planning domain knowledge using an unsupervised learning algorithm to generate domain theories from raw images. LATREC uses the learned domain knowledge on traditional goal recognition techniques (Pereira, Oren, and Meneguzzi 2017) to recognize the correct intended goal from image data. Our domain knowledge is automatically generated using a transition function derived from a state representation learned by auto-encoder.

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Thus, our approach applies modern goal recognition algorithms directly on real-world data, rather than using a domain expert to describe domain knowledge.

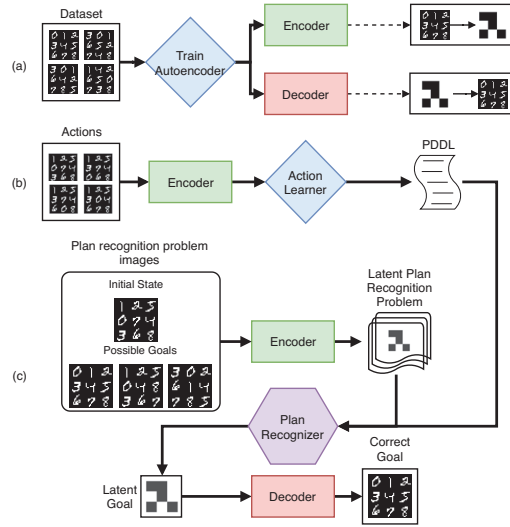


Figure 1: Pipeline for recognizing goals in latent space

Goal recognition in Latent Space

Planning algorithms are based on the *factored* transition function that represents states as discrete facts. This transition function is traditionally encoded manually by a domain expert, and virtually all existing plan recognition approaches require varying degrees of domain knowledge in order to recognize observations. Automatically generating such domain knowledge involves at least two processes: converting real-world data into a factored representation (*i.e.*, predicates for the planning process); and generating a transition function (*i.e.*, the set of possible actions in the planning domain) from traces of the factored representation. A recent approach by Asai and Fukunaga (Asai and Fukunaga 2018) uses an auto-encoder (Vincent et al. 2008) neural network to automatically generate domain models from images of simple games and problems. The neural network uses an encoder to

convert an input image into a discretized representation. Our encoder receives 42x42 black and white images and outputs a 6x6 latent representation activated by Gumbel-Softmax (Jang, Gu, and Poole 2016). The decoder reconstructs the input image from the discretized representation.

Images processed through the encoder become bit vectors that provide analogues for states comprising propositional attributes. Sequences of such propositional states implicitly represent sequences of action executions, which, in turn, allow us to infer a PDDL domain by comparing states before and after the execution of such implicit actions. We use bit-wise comparisons from the states before and after an action to infer preconditions and effects for a number of unique actions, and then try to group unique actions together to account for potential noise in the actions. Thus, we compute all the elements of a PDDL action and call this process Action Learner, as illustrated in Figure 1 (b).

To be able to plan using this domain we generate a planning problem by providing two images to the auto-encoder: one corresponding to the initial state and one corresponding to the desired state. This, in turn, allows us to generate a goal recognition problem as formalized by previous work (Ramírez and Geffner 2009) and apply off-the-shelf goal recognition techniques¹, such as (Ramírez and Geffner 2009; Pereira, Oren, and Meneguzzi 2017). The output of such techniques is a list of goals ordered by probability of being the correct one. We then decode the recognized goal, obtaining its image representation using the decoder. We illustrate this process in Figure 1(c), and detail the process in Amado et al. 2018.

Data-driven model

To compute a data-driven model capable of recognizing goals in latent space, we train an LSTM (Hochreiter and Schmidhuber 1997) that receives a sequence of encoded states and predicts an encoded goal. To perform a fair comparison to the state-of-the-art, we use as input encoded states generated by the encoder module from the autoencoder created by Asai and Fukunaga (2018). Thus, we convert each image-state into a latent representation (a 6x6 binary matrix). First, given a set of image-states representing a sequence of states and the goal of a certain plan, we use the encoder to generate the latent representation for each image. Second, we train the LSTM to predict the goal given the latent states. The output is a representation of this goal. Finally, we use the decoder from Asai and Fukunaga autoencoder to convert the produced representation into an image.

To train the LSTM network, we require data extracted from plans for each domain. We use plan traces generated by LatPlan(Asai and Fukunaga 2018), observing the states that were reached in each plan. Each trace generated a list of states, and then we included the goal of each trace as a class to the LSTM. To improve accuracy in low observability scenarios, we included partially observable traces (which means some states were removed from the plan trace), including 10%, 30%, 50%, 70% of observability. No candidate

¹In this work, we used the landmark-based heuristics of (Pereira, Oren, and Meneguzzi 2017).

goals are needed to recognize goals in this approach.

Conclusion

We developed an approach for goal recognition using image data as evidence, obviating the need for human engineering to create a task for goal recognition. We tested it with three state-of-the-art approaches of goal recognition to evaluate the domain we derived from image evidence and a data-driven model to recognize goals. We aim to improve pruning of redundant actions in the domain inference process, encode domains with incomplete information to account for the noise and inconsistencies generated by the autoencoder, and develop more complex data-driven models.

Finally, the current approach is limited in that it is incapable of recognizing goals not included in the training dataset, as it deals with the goal recognition task as a classification task. Moving forward, we aim to develop a data-driven model that computes a probability of each predicate being true, instead of fully reconstructing the recognized goal. This should allow us to compute the probability of each candidate goal by measuring the overall distance of each goal hypothesis using the probabilities of individual predicates.

References

- Amado, L.; Pereira, R. F.; Aires, J. P.; Magnaguagno, M.; Granada, R.; and Meneguzzi, F. 2018. Goal recognition in latent space. In *IJCNN*.
- Asai, M., and Fukunaga, A. 2018. Classical Planning in Deep Latent Space: Bridging the Subsymbolic-Symbolic Boundary. In *AAAI*.
- Hochreiter, S., and Schmidhuber, J. 1997. Long short-term memory. *Neural Comput.* 9(8):1735–1780.
- Jang, E.; Gu, S.; and Poole, B. 2016. Categorical reparameterization with gumbel-softmax. *CoRR* abs/1611.01144.
- Pereira, R. F.; Vered, M.; Meneguzzi, F.; and Ramírez, M. 2019. Online Probabilistic Goal Recognition over Nominal Models. In *IJCAI*.
- Pereira, R. F.; Oren, N.; and Meneguzzi, F. 2017. Landmark-Based Heuristics for Goal Recognition. In *AAAI*.
- Pereira, R. F.; Pereira, A. G.; and Meneguzzi, F. 2019. Landmark-Enhanced Heuristics for Goal Recognition in Incomplete Domain Models. In *ICAPS*.
- Ramírez, M., and Geffner, H. 2009. Plan Recognition as Planning. In *IJCAI*.
- Sukthankar, G.; Geib, C.; Bui, H. H.; Pynadath, D.; and Goldman, R. P. 2014. *Plan, Activity, and Intent Recognition: Theory and Practice*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1st edition.
- Vincent, P.; Larochelle, H.; Bengio, Y.; and Manzagol, P.-A. 2008. Extracting and composing robust features with denoising autoencoders. In *ICML*.
- Zhuo, H. H. 2019. Recognizing multi-agent plans when action models and team plans are both incomplete. *ACM Trans. Intell. Syst. Technol.* 10(3):30:1–30:24.