LatRec: Recognizing Goals in Latent Space

Leonardo Amado[†], **João Paulo Aires**[†], **Ramon Fraga Pereira**[†],
Maurício C. Magnaguagno[†], **Roger Granada**[†], **Gabriel Paludo Licks**[†] and **Felipe Meneguzzi**[‡]
Pontifical Catholic University of Rio Grande do Sul (PUCRS), Brazil
Graduate Program in Computer Science, School of Technology
[†]{leonardo.amado, joao.aires.001, ramon.pereira, mauricio.magnaguagno, roger.granada, gabriel.paludo}@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. LATREC applies modern goal recognition algorithms directly to real-world data (images) by building planning domain knowledge using an unsupervised learning algorithm that generates domain theories from raw images. We demonstrate this approach in an online simulation of simple games, such as the n-puzzle game.

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 for goal and plan recognition rely on planning theories, which require a substantial amount of domain knowledge (Ramírez and Geffner 2009; Pereira, Oren, and Meneguzzi 2017). 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).We overcome this limitation in LATREC[1] 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. Thus, our approach ultimately applies modern goal recognition algorithms directly on real-world data, rather than using a domain expert to describe domain knowledge.

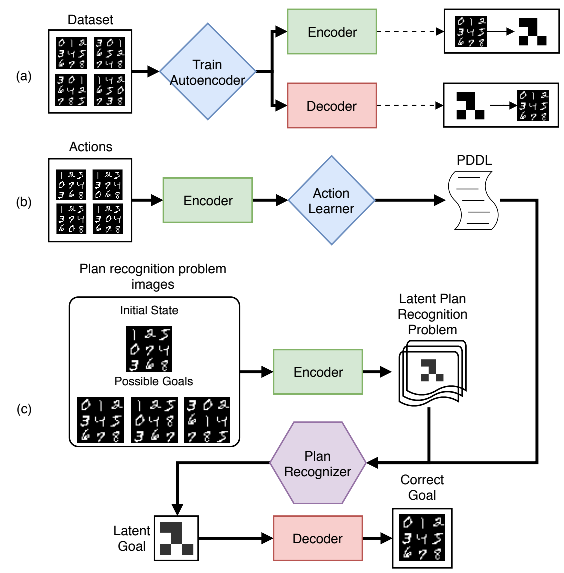

Figure 1: Pipeline of planning and goal recognition.

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 (Pereira, Oren, and Meneguzzi 2017). 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 autoencoder (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

[1]Demo video: https://goo.gl/CQoa9Y

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 bitwise 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 techniques[2], 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.

Application Overview

LATREC is a desktop application implementing our approach for recognizing goals in latent space. The user can start a "New Game", "Reset" the current game or "Solve" the puzzle, where all steps to solve the puzzle are displayed. The user should solve the puzzle by playing the board with blue tiles, where the navy tile containing "0" indicates the tile that the user is allowed to move. At the right side, the system outputs the three most likely goals that the user is trying to achieve. LATREC initially generates a randomly sorted sequence of blue tiles representing the initial state. Note that the moving tiles are blue to make it easier for the user to visualize the current board configuration, but in the background application it represents an image of the dataset. We display the set of possible goals the user can aim to achieve on the right side of the screen. With each move of the user, LATREC generates an encoded representation of the transition caused by the user action. After each movement, LATREC uses the goal recognizer to compute which are the most likely goals for the given action sequence. Finally, LATREC uses its decoder to convert the latent representation to images and display them on the right side of the screen. Figure 2 contains two screens showing the movement of the moving tile and the updating of their ranking.

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

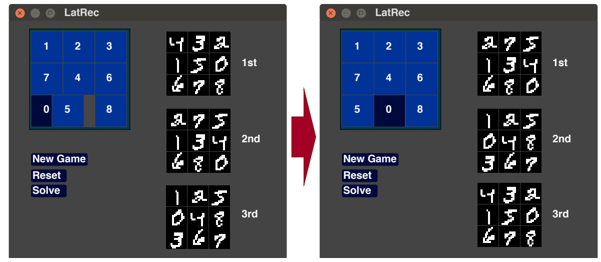

Figure 2: Application screen showing the transition of two states and the candidate goals for the current state.

Conclusion

We developed the demo for LATREC, a tool that performs goal recognition using real-world data (images). This demo showcases the effectiveness of our approach in different games, where the user can experiment and act as an agent in these domains. Currently, LATREC can recognize goals in three distinct games: lights-out; 8-puzzle and tower of hanoi. As future work, we aim to incorporate our advances in learned-model goal recognition, focusing in improving both the quality of our PDDL extraction and the performance of the auto-encoders we use.

Acknowledgements: Felipe acknowledges support from CNPq under project numbers 407058/2018-4 and 305969/2016-1, and FAPERGS process number 18/2551-0000500-2. We acknowledge support from HP Brasil using incentives of Brazilian Informatics Law n^{o} 8.2.48 of 1991. We acknowledge the support given by CAPES under Finance Code 001 and the Pro-Alertas project (88887.115590/2015-01).

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*.

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*.

Vincent, P.; Larochelle, H.; Bengio, Y.; and Manzagol, P.-A. 2008. Extracting and composing robust features with denoising autoencoders. In *ICML*.