Across the board: Human-human game learning dialogues

Mattias Appelgren, Erik Lagerstedt, Asad Sayeed and Christine Howes
Department of Philosophy, Linguistics and Theory of Science, University of Gothenburg
firstname.surname@gu.se

Abstract

Learning novel tasks through dialogue and interaction is a common human activity. The field of Interactive Task Learning looks to give this capability to AI systems. However, the dialogue strategies used in those systems are not based on observed human strategies. This paper presents a data collection of instructive dialogue in the form of filmed interactions where a teacher teaches a board game to a student.

1 Introduction

The field of Interactive Task Learning (ITL) (Laird et al., 2017) looks to create Artificial Intelligence (AI) systems capable of learning interactively from human teachers.

The systems rely on dialogue strategies which are largely fixed and based on the capabilities of the systems or the decisions of the designers, not on any observed or theoretical dialogue strategy. In this paper we present a data collection effort which has the aim of studying instructive dialogue. We believe the effort will be valuable for the purpose of ITL, but that there are many possible phenomena that can be studied.

The domain we have chosen is board game instruction. Anyone who has played a board games will be familiar with this scenario: you come over to a friends house they suggest a game and they bring out a card board box that you have not seen before. Inside are different components and a rulebook which give meaning to the pieces of wood and cardboard. The friend then goes on to explain the rules.

Now, in many ways this setting is a good example of our human capability to learn a completely new task simply from interaction, and highlights many of the problems that an ITL system would face. For example, in a board game, different game components will be named different things. For example, a black cube might represent coal, a wooden figure referred to as a meeple. This is a type of grounding (Harnad, 1990) which cannot be solved by supervised learning beforehand, because the words and figures may be unique to the game being played. This type of interactive grounding is a common task in ITL related work (Matuszek, 2018; Krause et al., 2014; Scheutz et al., 2017; Lindes et al., 2017).

In a board game players generally perform particular actions during their turns which have specific conditions under which they are allowed or particular costs, with corresponding effects on the game state. These actions will be particular to the game and so a game player will have to learn what these actions are and when they are allowed. This is equivalent to learning actions in planning scenarios (Chai, 2018; She et al., 2014; Scheutz et al., 2017).

Finally, the game player will have to learn how to win the game: how to score or in what state the game ends. This can be likened to some kind of goal learning in planning (e.g. Appelgren and Lascarides (2019)) or inverse reinforcement learning, i.e. learning the reward function, tackled in e.g. Hadfield-Menell et al. (2016); Abbeel and Ng (2004).

Combined, board game instruction creates an interesting domain for ITL. However, it can also be interesting for research in dialogue, pragmatics, cognitive science, and games research. Due to the nature of board games, the context is fairly limited and controlled while remaining comparatively naturalistic and dynamic even when studied in a laboratory setting. In order to understand how people teach and learn games, we have decided to collect a dataset of board game instructing where an experienced and a novice player plays a game together.

Proceedings of the 29th Workshop on the Semantics and Pragmatics of Dialogue, September 3–5, 2025, Bielefeld, Germany.

2 Data Collection

We have elected to collect dialogues between a teacher and student learning the game Carcassonne. We chose Carcassonne for a number of different reasons. First, we want a game that is simple enough to learn and play quickly but not so simple that it would be trivial to learn the rules. Second, we wanted a game that is well known enough that there would be enough potential teachers that could be recruited, but not too popular so that the majority of potential learners would know of it, even if they may not have played it, and thus not have too much of a preconception (as might be the case for chess or monopoly). Third, we wanted a game that did not rely too much on hidden information or special rules on cards, which would be difficult to capture on camera or make learning the rules reliant on reading cards rather than understanding general rules. Carcassonne seemed to hit that balance well.

For data collection we set up the game on a table and set up three camera angles to capture the action. We had one camera above the table (Figure 1) facing down to capture the game play and hand movements. The other two cameras are pointed at each participant (Figure 2). Additionally, we fit each participant with a microphone and record their voice on two separate channels.

We recruit two participants for each recording session. One participant acts as a teacher and the other a student. Teachers have played Carcassonne before and feel confident enough to teach the game. For the students we recruit people who have not played Carcassonne before.

Prior to the actual recording we ask the participants to fill out a questionnaire where we ask them about what languages they speak, gender, familiarity with teaching, and familiarity with games.

At the time of writing we have had two pairs of participants recorded. Our goal is to recruit at least 10 pairs.

3 Hypotheses

In machine learning there is generally a supervised learning stage followed by a testing phase where the models are supposed to have learned everything they need.

A similar thing happens in the ITL dialogues where the teaching dialogue essentially enumerates all the possible rules, actions, states, etc., and once all the information has been conveyed the teaching is finished and the system is assumed to know the

Figure 1: Birds-eye view of game board.

Figure 2: Camera angle viewing the participant learning (left) and teaching (right) to play the game.

task.

However, people learn continuously. Fitts and Posner (1967) divide learning into three phases, where the learner first must understand the task conceptually, then a phase of trial and error where mistakes occur and can be corrected by a teacher, and finally an autonomous phase. In board games this would be equivalent to rule explanation, followed by supervised play, followed by autonomous play. We expect to find this pattern in the data.

We will transcribe the recordings and annotate the data according to when and how rules are introduced or referenced (e.g. as part of a correction). We will look at when students make mistakes and how they recover from them.

By investigating the data we collect we expect to find some of the ways in which these situations arise, gain a better understanding on the dynamics of humans teaching/learning these kinds of activities, and inform the design of future ITL systems, for the board game domain and beyond.

Despite the data being collected in a controlled environment with cameras and microphones recording the interaction, we argue that the ecological validity is unusually high in this study. The board game context lends itself to an integrative view of actions – both those that are performed linguistically and those that are non verbal but nevertheless an integral part of the dialogue.

Acknowledgements

This research was supported by ERC Starting Grant DivCon: Divergence and convergence in dialogue: The dynamic management of mismatches (101077927) and by the Swedish Research Council grant (VR project 2014-39) for the establishment of the Centre for Linguistic Theory and Studies in Probability (CLASP) at the University of Gothenburg.

References

Pieter Abbeel and Andrew Y. Ng. 2004. Apprenticeship learning via inverse reinforcement learning. In *ICML*.

Mattias Appelgren and Alex Lascarides. 2019. Learning plans by acquiring grounded linguistic meanings from corrections. In *In Proc. of the 18th Inter- national Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019, IFAAMAS*, page 9 pages.

Joyce Yue Chai. 2018. Language to action: Towards interactive task learning with physical agents. In *AAMAS*.

Paul M Fitts and Michael I Posner. 1967. *Human performance.* Brooks/Cole.

Dylan Hadfield-Menell, Anca D. Dragan, Pieter Abbeel, and Stuart J. Russell. 2016. Cooperative inverse reinforcement learning. In *NIPS*.

Stevan Harnad. 1990. The symbol grounding problem. *CoRR*, cs.AI/9906002.

Evan A. Krause, Michael Zillich, Thomas Emrys Williams, and Matthias Scheutz. 2014. Learning to recognize novel objects in one shot through human-robot interactions in natural language dialogues. In *AAAI*.

John E. Laird, Kevin A. Gluck, John R. Anderson, Kenneth D. Forbus, Odest Chadwicke Jenkins, Christian Lebiere, Dario D. Salvucci, Matthias Scheutz, Andrea Lockerd Thomaz, J. Gregory Trafton, Robert E. Wray, Shiwali Mohan, and James R. Kirk. 2017. Interactive task learning. *IEEE Intelligent Systems*, 32:6–21.

Peter Lindes, Aaron Mininger, James R. Kirk, and John E. Laird. 2017. Grounding language for interactive task learning. In *RoboNLP@ACL*.

Cynthia Matuszek. 2018. Grounded language learning: Where robotics and nlp meet. In *IJCAI*.

Matthias Scheutz, Evan A. Krause, Bradley Oosterveld, Tyler M. Frasca, and Robert Platt. 2017. Spoken instruction-based one-shot object and action learning in a cognitive robotic architecture. In *AAMAS*.

Lanbo She, Shaohua Yang, Yu Cheng, Yunyi Jia, Joyce Yue Chai, and Ning Xi. 2014. Back to the blocks world: Learning new actions through situated human-robot dialogue. In *SIGDIAL Conference*.