

---

## X Goes First: Teaching Simple Games through Multimodal Interaction

---

**Thomas R. Hinrichs**

**Kenneth D. Forbus**

EECS, Northwestern University, Evanston, IL 60208 USA

T-HINRICHS@NORTHWESTERN.EDU

FORBUS@NORTHWESTERN.EDU

### Abstract

What would it take to teach a computer to play a game entirely through language and sketching? In this paper, we present an implemented program through which an instructor can teach the rules of simple board games using natural language and sketching. We describe the architecture, information flow, and vocabulary of instructional events and walk through an annotated example. In our approach, the instructional and communication events guide abductive reasoning for language interpretation and help to integrate information from sketching and language. Having a general target representation enables the learning process to be viewed more as translation and problem solving than as induction. Lastly, learning by demonstration complements and extends instruction, resulting in concrete, operational rules.

### 1. Introduction

A long term promise of cognitive systems is that they should be able to learn simple tasks by instruction, rather than requiring detailed programming. Yet formidable obstacles remain. For tasks that are structured, conditional, iterative, and context sensitive, a learner must construct fairly sophisticated representations. Resolving the imprecision and ambiguities of natural interaction depends critically on plausible inference and the expectations of the learner. What kinds of expectations can be brought to bear when learning from such interaction?

We have been exploring these issues through the multimodal instruction of simple games. The teacher and learner communicate via natural language textual dialog and sketching. From this interaction, the learner acquires the definitional rules of a game sufficient to play it. The benefit of this is that the simplicity of the learning task allows us to focus on instructional interaction patterns. It is clear when the rules have been learned correctly, incorrectly, or incompletely, whereas higher-level strategies are more subjective and harder to evaluate. Board games, in particular, are ideal because they are simple, highly spatial, and have explicit concrete rules.

This paper presents our initial analysis and experiments in teaching Tic Tac Toe. Although an extremely simple game, the instruction runs headlong into many of the difficulties to which we have alluded. In fact, the title of this paper suggests some of the issues: “X goes first.” By itself, this has no clear meaning. X is a metonymic reference to a player based on the form of their mark. “Goes” in this case, means takes a turn, and “first” means that turn is first in an implicit sequence of turns. The input sentence and the intended meaning are both simple, but the interpretation process is not. In our model, expectations are built up progressively by recognizing implicit instructional events in the dialog. This contextual bootstrapping is complemented by multimodal constraints that ground reference and guides generalization. We aim to identify

patterns of interaction by which depiction and verbal instruction mutually constrain each other to help build a coherent task representation.

We make four claims:

1) *A small set of instructional events can suffice to constrain interpretation.* These need not be strictly sequenced or form any kind of high-level grammar, but they serve as strong contextual hints for resolving ambiguity, in combination with temporal and spatial proximity clues.

2) *Natural interaction is facilitated by multiple levels of representation.* The content of natural instructional dialog primarily concerns concepts and examples. To produce an executable program, this content must be translated into operational rules. By translating incrementally through different levels of representation, appropriate expectations, such as instructional events, can help disambiguate and interpret input.

3) *Spatial depiction grounds verbal abstractions.* Because a sketch is intrinsically propositional, it can resolve ambiguous abstractions, while verbal descriptions help lift and generalize over-specific examples.

4) *Learning by demonstration complements instruction.* Verbal instruction can most easily express abstract taxonomic classifications, associative relations, and parameters, while demonstration can often present procedures and configurations more concisely. Even a simple game like Tic Tac Toe benefits from integrating these two learning strategies.

The next section describes the concrete task of teaching Tic Tac Toe, the target representation, and the general cognitive architecture. Sections 3 and 4 present the language and sketch understanding components respectively. Section 5 describes the multimodal fusion, learning, and game-playing processes. Section 6 presents an annotated transcript of learning Tic Tac Toe. Section 7 analyses the interaction patterns and their generality and briefly describes experiments in learning the piece-moving game Hexapawn. Lastly, we describe related and future work and summarize results.

## 2. Teaching Tic Tac Toe

Most of us have some experience teaching Tic Tac Toe to a child or remember learning it ourselves. The extreme simplicity of the game is deceptive - learning it assumes familiarity with a number of concepts such as making physical marks, taking turns, spatial relations between marks, learning from an instructor, rule-defined behavior, winning and losing, and facility with language. Just as a child is not a blank slate, our learning program also builds on prior knowledge.

Our approach treats learning as a kind of problem solving, rather than inducing rules over data sets. If you know what you're trying to learn, i.e., the form of the target representation, then learning becomes a matter of looking to fill gaps in a model. Moreover, if that model is operational – it can be employed in a task – then instruction can be augmented by demonstration, practice, and feedback. We therefore define our task as introducing concepts and rules of the game through natural language and sketch gestures, followed by interactively playing the learned game in the same interface with the same sketch.

### 2.1 Target Representation

The target representation to be learned must be concrete enough to support interactive play, yet general enough to cover a wide range of possible games. Primarily for this reason, we use an

implementation of the Game Definition Language (GDL) (Love, Hinrichs & Genesereth, 2006). GDL represents games as a set of Horn clauses that define a finite state machine (FSM). The predicates *initial*, *legal*, *next*, and *terminal* support the definition of a FSM, augmented by game-specific predicates *role* and *goal*. By fixing the form of the target representation in this way, we are able to support a wide variety of games with an interpreter that can translate between the logical state of a game and the sketch-based presentation.

## 2.2 Architecture

Our game learner is implemented within the Companions Cognitive Architecture (Forbus, Klenk & Hinrichs, 2009), as diagrammed in Figure 1. The learner consists of plans and rules that coordinate the interactions of three primary agents: a language interaction agent that processes natural language, a sketch agent that supports sketch interaction and gameplay, and a domain reasoning agent that assimilates events from the other two agents and incrementally builds the domain rules of the game. Each agent is endowed with a knowledge base containing the ResearchCyc ontology<sup>1</sup>. This provides a starting vocabulary of many thousands of interrelated concepts and predicates, denotations and semantic frames for mapping from English lexical tokens to symbolic concepts, and a mechanism for partitioning knowledge into *microtheories*, an inheritance lattice of collections of mutually consistent assertions. These agents use the FIRE reasoning engine (Forbus et al. 2010) which supports deductive retrieval and backward chaining, and/or solving, and hierarchical task network planning. Agents iteratively plan and execute actions from an agenda of tasks. This built-in agenda allows game rules to define a control strategy without requiring a separate, game-specific control loop.

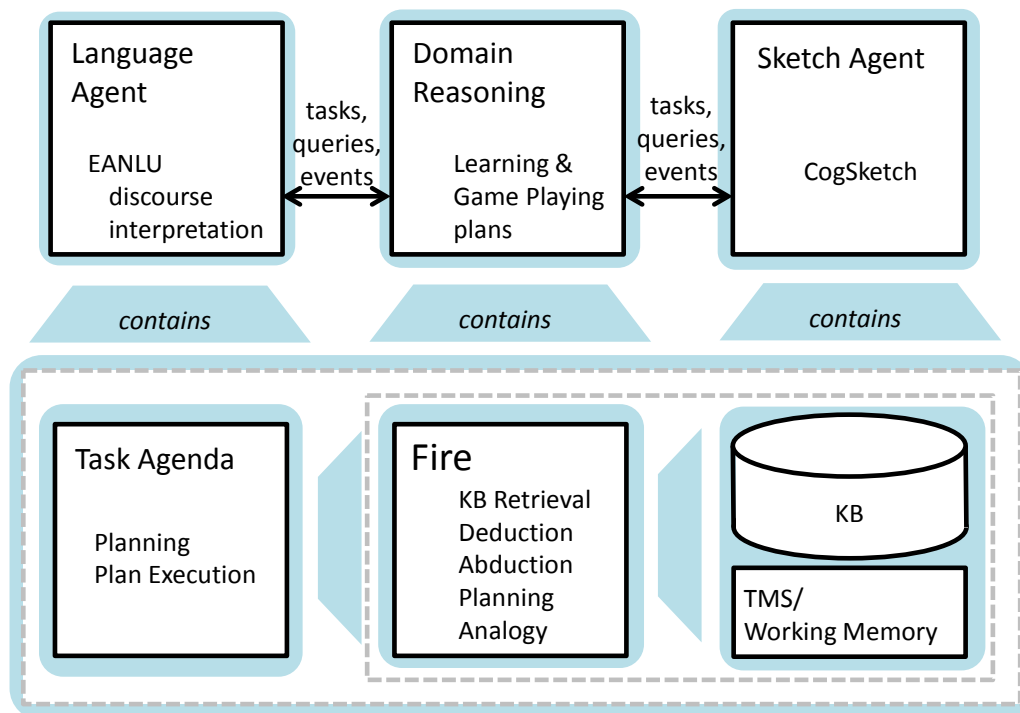


Figure 1: Companions Agent Architecture

### 3. Language Interaction

The interaction manager is built on top of the Explanation Agent natural language understanding (EA NLU) system (Tomai, 2009), which integrates several off-the-shelf components and resources to produce, given a sentence and discourse context, representations of the possible meanings of that sentence, expressed in the ResearchCyc ontology. The system preserves syntactic and semantic ambiguities in explicitly represented choice sets. Ambiguous choices are resolved using controlled abductive inference (Hobbs et al., 1990), constrained by high-level recognition rules that seek contextually meaningful interpretations.

As we instruct the learner, it applies these general rules to look for instructional events that could explain an utterance. For example, could a given utterance be interpreted as introducing or naming an entity in the game? When such a query can be satisfied by collapsing an ambiguous choice to a particular semantic interpretation, it will do so. This relies on inference bottoming out in propositional abductive rules that are compiled out by the parser for each sentence. By querying for individual conjuncts of an interpretation, it identifies the assumptions that could make it true and rules in the reified semantic choice. Collectively, these semantic choices are composed into a coherent representation of the sentence. Figure 2 illustrates an example sentence, a high-level interpretation rule for recognizing goal expressions, a portion of a low-level abductive rule, and the resulting interpretation.

The resulting sentence representation is a predicate calculus description that may contain

*"The first player to mark three squares in a row wins"*

```
(<== (definesGoal ?sent ?role ?goal) (and
  (winning ?sent ?role ?win) (isa win5927 Winning)
  (winningAction ?sent ?win ?role ?act) (performedBy win5927 player5762)
  (winningState ?sent ?role ?act ?goal) (isa player5762 GameRole)
  (goalStateConfiguration ?sent ?output)) (nthInSeries player5762 series5758 1)
  (doneBy mark5769 player5762)
  (isa mark5769 MarkingOnASurface)
  (objectMarked
    mark5769 group-of-square5801)
  (isa group-of-square5801
    Set-Mathematical)
  (selectedChoice
    (cardinality group-of-square5801 3)
    (and (isa win5927 Winning)
      (elementOf square5801
        group-of-square5801)
      (isa square5801 Square)
      (in-UnderspecifiedContainer
        square5801 row5880)
      (isa row5880 RowOfObjects)))
  ...
  (isa row5880 RowOfObjects)))
```

Figure 2. Example sentence, interpretive rule skeleton and resulting FOPC interpretation

tokens representing discourse variables (essentially skolems) and embedded microtheories. This representation, although formal, is not an operational specification of a game rule. Consequently, this logical form undergoes at least two additional levels of expectation-based translation. It is first lifted to a more language-neutral form that replaces embedded microtheories and discourse variables with 2nd order predicates. Based on whether the utterance is a declarative, interrogative, or imperative, it is then passed to the domain reasoning agent as the arguments of a *communication event*. The domain reasoning agent assimilates communication events and further translates the content into explicit game rules. This process of progressive translation is illustrated in Figure 3. The vertical arrows show the direction of translation from input to the target game representation.

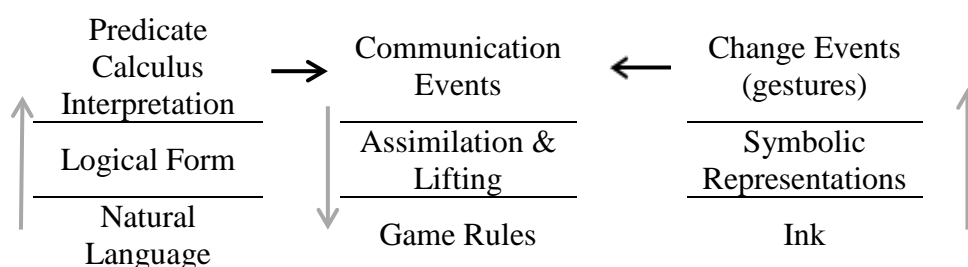


Figure 3. Representational Levels used in Language, Domain Reasoning, and Sketch Agents

#### 4. Sketch Interaction

Sketch interaction is supported through a Companions agent wrapper around the CogSketch sketch understanding system (Forbus et al, 2011). One of the major functions of CogSketch here is to convert between a visual depiction and a structured, symbolic representation of the ink. It also captures and represents mouse and/or pen gestures in a way that can be communicated to a reasoning system and supports action primitives for modifying, duplicating, and deleting elements in a sketch. Supporting such programmatic control allows the sketch interface to serve as the user interface for playing games.

Structurally, a sketch may contain a number of sub-sketches that can represent alternative views or different states in time of a depicted scene or object. The sub-sketches in turn consist of aligned layers (like tracing paper). Layers contain entities called *glyphs* which consist of ink segments. Glyphs also have semantic content, either an entity in the depicted world or they express a relationship or elaboration between such entities.

When we teach or play a board game, we draw glyphs to represent the pieces or marks and the board. As we introduce and identify pieces, they are added to a subsketch that we can think of as a catalog. The catalog holds prototypes of pieces or marks, but doesn't preserve any particular spatial arrangement between them. As objects are named, it associates the prototype name with the glyph. Figure 4 shows the state of a sketch when the teacher is about to introduce the second player, O.

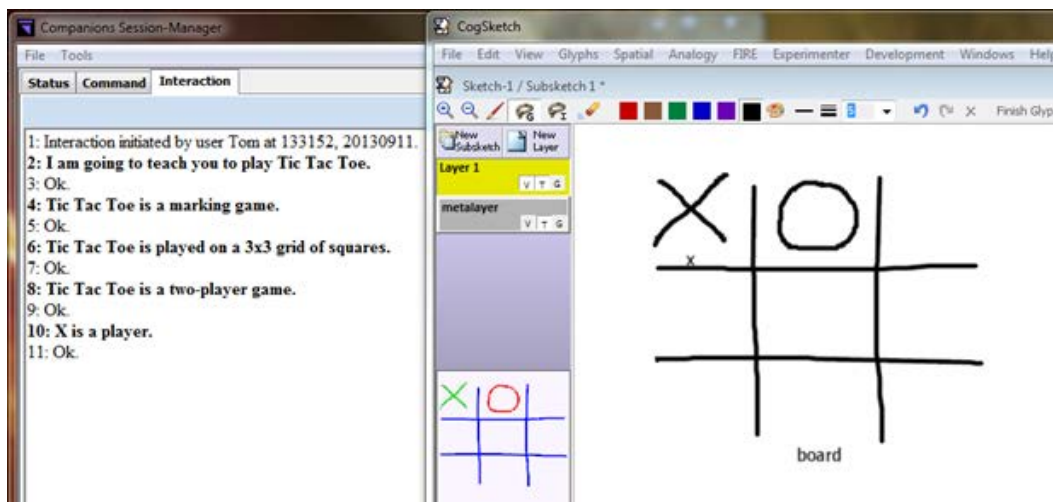


Figure 4. Prior to introducing ‘O’

Later, when we start to play a learned game, the system generates another sub-sketch to represent the current state of the game. It marks or populates the board by instantiating a prototype glyph from the catalog into a legal location on the game board.

## 5. Learning and Generalizing

User interactions are passed from the language and sketch agents to the reasoning agent in the form of reified *communication events*. This does two things: It preserves the relative ordering of utterances and gestures, and it supports flexible control for interactive learning and game play. The learner can respond to a request or declarative statement, or the creation, destruction, movement, or selection of a sketch entity. Moreover, it can wait for such a typed event, and keep track of the last entity created, moved, etc. This enables flexibility in presentation ordering and it allows implicit turn taking. There is no need for the human opponent to say “your turn”, because the agent is waiting for an entity creation event from which to update the state and start its move.

The reasoning agent is responsible for translating (possibly incomplete) predicate calculus statements into operational game rules. This is a translation from one formal representation to another, so it is much more straightforward than language interpretation.

### 5.1 Extending game rules through demonstration

Not all game rules can be conveniently expressed through language. In our example transcript, the goal is expressed as three marks in a row. This leads to a single rule that recognizes a horizontal row. When we play the game, the instructor can introduce new rules on the fly after they or the system wins through a vertical or diagonal line. “I win” or “you win” signifies that a new rule should be produced by lifting the game configuration for the winning role. It consults the existing goal rule and determines that the state is a spatial configuration of the winner’s marks. Given that, it lifts the current winning state in a similar manner. In the case where there

are extraneous marks that do not contribute to the goal, the instructor may explicitly select the relevant glyphs in the sketch, to avoid over-fitting. Of course, many games have goals that cannot be defined so simply, checkmate being an obvious example. In such cases, a more sophisticated mixed initiative dialog will be required.

## 6. Annotated Example

Tic Tac Toe can be learned in a very small number of interchanges. In the following example, the instructor starts by introducing the topic:

1) *I'm going to teach you to play Tic Tac Toe.*

The purpose of this is to bootstrap the interpretation context. Because TicTacToe is already a concept in the ontology, the learner knows that it is a game, so it enables the game-learning interpretation rules and instructional events. Since there is no natural language generation, the system responds with "Ok." to signal that it understood. We omit these responses from here on.

Next, the instructor creates a new blank sketch in which to draw the elements such as marks and the board. He then classifies the game:

2) *Tic Tac Toe is a marking game.*

This distinguishes it from a piece-moving game or a construction game and imports some general rules about marking, such as the property that marks accumulate monotonically.

*Instructor draws a hash for the game board.*

The system does not recognize the ink or have any expectations about what it could be, so it merely records the fact that an entity was created.

3) *TicTacToe is played on a 3x3 grid of squares.*

Now the meaning of the hash mark in the sketch becomes clear. The utterance describes a spatial configuration which is inferred to be a game board. The most recent un-interpreted sketch entity is assumed to be that board and is implicitly divided into a 3x3 Cartesian coordinate system.

4) *Tic Tac Toe is a two-player game.*

This classifies the game again and imports basic rules for turn-taking.

5) *X is a player.*

Because this is a game domain, player is understood to mean a game role, rather than an athlete. This is another example of the abductive mechanism at work. Each domain has a set of abductive preferences, specified in a microtheory associated with that domain. Thus the preferences are expressed declaratively, making them potentially easier to learn.

*Instructor draws an X.*

The most recently introduced role name is assigned to the new sketch entity. Tic Tac Toe is a slightly special case in that the player's roles and their marks are the same, whereas in a game like Chess, entities would correspond to pieces.

*Instructor draws an O.*

The O can be drawn either before or after naming it.

6) *O is a player.*

Again, it defines the role and associates the mark entity with it.

7) *X and O take turns marking empty squares.*

Now that the roles have been defined, the instructor has elaborated the marking action precondition to require empty squares and turn taking.

8) *The first player to mark three squares in a row wins.*

This describes one of the winning conditions – a horizontal row of marks. Rather than over-generalize the concept of “row”, we leave the remaining winning conditions to be learned by demonstration.

9) *X goes first.*

X is now understood to be one of the players, and although the sequence of turns is not explicit in the language, the expectation of turn taking primes the interpretation.

10) *Start a game.*


The imperative utterance becomes a request to play a game. Enough information has been conveyed to play a legal game.

11) *You win.*

The learner plays with one-step lookahead, but beyond that, its actions are almost entirely random legal moves. By letting it win with a vertical column, the instructor conveys a new winning condition from which it must lift a new rule. Figure 5 shows the state of the board game at this point of the game, plus the rule extracted as a consequence of being told that this state is an example of a win condition.

## 7. Analysis

Learning by instruction presents some challenges for evaluation, because it doesn’t happen incrementally over many trials, so there are no learning curves to show effectiveness. This is known as *one-shot* learning. Instead, we will analyze the contributions of different parts of our



```
(<== (gdl-goal ?player 100)
      (isa ?player GameRole)
      (entityLabel ?player ?mark)
      (gdl-true (cell 3 ?y ?mark))
      (gdl-true (cell 2 ?y ?mark))
      (gdl-true (cell 1 ?y ?mark)))
```

Figure 5. Learning to win with a vertical column plus induced rule representation.



Table 1. Instructional Events for teaching spatial games

Instructional Event	Description	Utterance	Ambiguity Reduction
Topic Introduction	Set up context for future interpretation	1	20
Game Classification	Extract type of game, enable partial	2,4	1
Role Definition	Associate names with players	5,6	2
Entity Introduction	Associate names with marks or pieces	-	-
Action Elaboration	Describe an action type	7	432
Configuration - Spatial	Define or reference a spatial configuration	3	1584
Configuration - Temporal	Define or reference a temporal configuration	9	3
State Description	Define or reference a named state	11	1
Win Conditions	Define condition for winning the game	8	384
Directive	Request an action	10	7

model and catalog the interaction patterns employed.

## 7.1 Instructional Events

Most of the expectations about games are manifested as *instructional events* in the language interpretation. These are similar to dialog acts, except that they don't drive the learner from one explicit high-level state to another, but instead, a side effect of recognizing them is that they make abductive choices that collapse ambiguity. Table 1 presents our catalog of instructional events for game learning indexed by the example utterance that triggers them. Note that Entity Introduction is not used in this example because Tic Tac Toe conflates roles with marks. Ambiguity reduction is a measure of the number of possible semantic interpretations of the input sentence. This can vary widely for different sentences and grammatical constructions, but shows the effectiveness of the abduction mechanism. How will this change as we address additional games and types of games? The set of recognition rules will certainly need to grow, to support new phrasings and constructions. However, we do not expect the set of instructional events to grow significantly. We only see two additions, based on preliminary analyses: (1) describing strategies, which requires combining multiple existing event types and (2) describing tradeoffs.

In addition to instructional events, a small set of *communication events* integrate the multiple modalities. Our set of communication events includes UserStatement, UserRequest, UserQuery, EntityAdded, EntityDeleted, EntityMoved, and EntitySelected. Although others are possible (e.g. EntityRotated), we do not expect to need them for board game domains.

## 7.2 Compression Ratio

One way to measure the effectiveness of an instructional system is to measure the compression ratio. How concise can the instruction be and how much information can be conveyed by a single




Figure 6. Compression Profile by Utterance

sentence? When the instructor says "The first player to mark three squares in a row wins", this is a very concise way of conveying a lot of information, as illustrated in Figure 2.

In this case, the intermediate representation is quite verbose. We've translated from a single sentence, to a conjunction of 13 expressions, and on to a Horn clause with six sub-expressions. On the other hand, a sentence like "TicTacToe is a marking game" produces an interpretation with one expression, (isa TicTacToe MarkingGame), but that imports 13 expressions into the game definition (in the form of 2 assertions and 5 rules). Figure 6 presents the compression profile for the sentences in the example transcript.

As a qualitative measure, this tells us something about the instruction process. In some cases, background knowledge is being applied in the translation and operationalization of the interpretation, and in other cases, it is the composition of general, but operational knowledge that is applied. We take this as support for the incremental, progressive interpretation and translation process.

Another way to look at Figure 6 is that it shows the relative contributions of prior knowledge from Research Cyc to learned knowledge. The intermediate representation consists of Cyc concepts and predicates, while the target representations are the expressions constituting GDL rules and properties. The final learned definition of Tic Tac Toe contains 12 rules, 10 initialization statements, 2 numeric parameters and 3 entity names.

### 7.3 Learning Hexapawn

In order to better understand the generality of our approach, we applied our system to learning the game of Hexapawn. This is a game that is similar in size and complexity to Tic Tac Toe, but involves moving pawns on a 3x3 board.

Rather than require the instructor to draw pawns, we instead created a sketch with an initial catalog of glyphs by importing vector graphics (See Figure 7). This resulted in fewer deictic references and virtually no gestures during the instructional phase. This would need to change if we were to learn spatial relations by example, or strategic concepts, or to explain how a knight moves in Chess

The main difference between Tic Tac Toe and Hexapawn is that the latter is a piece-moving game. This implies some background knowledge about the persistence of pieces, e.g. that movement removes a piece from its original location and places it at the destination. We can

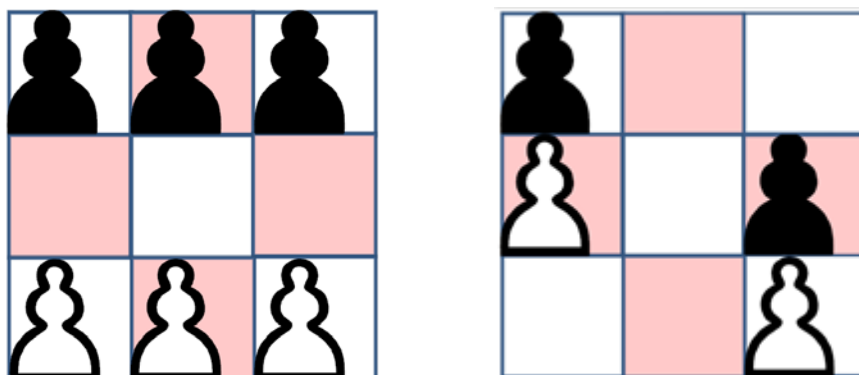


Figure 7. Hexapawn initial state and win configuration

assume by default that a location only holds a single piece, though this can be easily overridden through instruction. Another difference is that while Tic Tac Toe conflates the player with his mark, there is a definite distinction in a piece moving game between a player and his or her pieces. The representation must support that level of indirection, and allow for the casual use of a piece to stand in for a player and vice versa. A third difference is that a piece moving game doesn't necessarily start out with an empty board. It was necessary to recognize references to the initial state and translate that into rules for populating the board.

In terms of the process of instruction, the two games are very similar. We are able to teach the rules of Hexapawn in 16 sentences. We introduce players, pieces, and game configurations in much the same way, but assume prior knowledge of a few more spatial relations than were necessary for Tic Tac Toe. For example, one instruction is: "White may move wp diagonally up to capture bp." *Up*, *down*, *diagonal*, and *across* are all adjectives (and sometimes adverbs) that we operationalize directly in terms of grid coordinates. Clearly, these are concepts that could be induced through examples, but our purpose is not to start with a tabula rasa, but to find an appropriate level of discourse that makes instruction feasible and effective.

The ambiguity reduction in the Hexapawn transcript is similar to that of Tic Tac Toe. The number of possible interpretations of the input sentences typically ranges from 1 to 20, with outliers at 198 and 2160 for the sentence "The first player to move a piece across the board wins."

## 8. Related Work

We are taking a collaborative problem solving approach to instruction (Allen et al, 2002). Our task and approach have a different emphasis than PLOW (Allen et al, 2007), which also uses this approach. Both do one-shot learning, but PLOW uses language more as a running commentary on a demonstration than as instruction. In our case, demonstration is treated as a repair to instruction. Our architecture is less stratified than Allen's because there is no grammar of higher-level states that the learner passes through. Our reified communication events to assimilate statements, glyphs, etc., correlate with the communicative acts in his account.

Several prior systems have been built that learn Tic Tac Toe from purely visual inputs, either generated by a person manipulating a physical game board (Kaiser, 2012) or from a robot manipulating a specially engineered game board (Barbu, Narayanaswamy & Siskind, 2010).

Both of these systems require substantially more input data than ours to start playing the game. This is the advantage of using natural language to concisely and quickly communicate basic concepts about the game. On the other hand, communicating the full range of winning conditions is something that can conveniently be done via demonstration.

The approach that is most similar and complementary to ours is being pursued by the Soar group at the University of Michigan (Kirk & Laird, 2013). They are also teaching games, including Tic Tac Toe among others, through language, demonstration and spatial interaction, in this case using a physical robot. Our sketch interface allows us to avoid many of the difficult perceptual problems and focuses our efforts in different areas. The Soar project has emphasized spatial preposition learning early on, while we have deferred that in favor of starting with some basic prepositions and seeing what can be expressed with that vocabulary. We have also not pursued natural language generation so there is less mixed-initiative dialog.

## 9. Limitations and Future Work

Our current implementation has a number of limitations which can be broadly classified into language, learning, and scaling issues. Although it sounds natural, the language interaction is not especially robust. Seemingly small changes in phrasing can cause parsing or interpretation to fail. While this is an active area of research for us, it is not the focus of this project in particular. One positive effect of the abductive interpretation process is that it is mostly additive. That is, additional knowledge does not break existing interpretation capabilities.

Our learner does not currently induce new intermediate predicates. This has not been an issue with the simple games we have addressed so far, but we will probably need to borrow some techniques from Inductive Logic Programming (Muggleton & De Raedt, 1994) before attempting to learn more complex games.

We do not yet have a facility for repairing incorrectly learned rules through language, though there doesn't seem to be a deep theoretical impediment to this. We address the problem of overfitting by allowing the instructor to select entities on the board to designate the relevant part of a winning configuration. This is only necessary when it is ambiguous, and the technique might need to be extended for more sophisticated games.

Scaling is one reason we haven't yet attempted to learn the rules of Chess. The sheer number of legal moves makes it unpleasantly slow to play through the backchaining reasoner, and the sketch understanding system automatically computes qualitative spatial relations that slows significantly when there are 32 pieces on the board. Neither of these are intrinsic impediments. More control over spatial inference and compiling out rules to a more efficient form for execution would make Chess a reasonable domain to learn. In the meantime, we intend to pursue sub-problems, such as learning how different pieces move, and learning endgame strategies.

In addition to teaching other games and game strategies, we are also looking at teaching physical reasoning tasks through multimodal interaction. Our main research goal is to identify instructional techniques that are amenable to multimodal interaction. A prime candidate is analogy. Analogy is one of the core capabilities of the Companions architecture, but so far, this project has made no use of analogy at all. Nevertheless, we see a large role for it in bringing across concepts and action models from one game to another. While our current implementation uses inheritance to import fragments from abstract game types such as `MarkingGame`, transfer from one concrete game to another could greatly simplify instruction.

Abduction has proven to be an extremely powerful tool. So much so that we are currently re-implementing the facility to embed it more broadly in our reasoning engine and to support more general weighted abduction. We expect this to improve the flexibility of the language interpretation system and to be useful in other kinds of problem solving as well, such as plan recognition.

## 10. Conclusions

We have demonstrated one way to produce an executable program through informal language and sketching. We have found sketching to be a good medium for teaching board games. Our main purpose is to identify kinds of knowledge and expectations that help enable instruction, while being general enough to not presuppose too much knowledge of the domain.

The implemented system supports our earlier claims in the following ways. We presented ten instructional events that constrain interpretation in the domain of learning simple board games. Translation from English into a conceptual representation and then on to more procedural rules allowed contextual expectations to be at the level of concepts and examples, rather than concrete operational rules. Propositional sketch representations ground verbal abstractions, especially during the process of lifting winning configurations to goal rules. This might be termed *synergy of modality*. Finally, instruction is complemented by demonstration when we provide feedback to the learner during gameplay. This might be termed *synergy of instruction*. The result is a dialog that we believe to be a natural kind of instruction, in which we are able to convey twelve operational rules in only ten natural language utterances.

One reason for focusing on learning the game specification rather than strategies for playing well is that once you have the basic rules for playing, everything else is optimization. It may well be the case that optimization (construed broadly) is a domain that itself is amenable to instruction. This leads to the obvious open question: Can the effectiveness of instruction reach the break-even point such that more and more background knowledge can be learned with less and less effort? We expect it will, but that is an empirical question.

## Acknowledgements

This research was supported by the Defense Advanced Research Projects Agency under contract D12AP00026. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressly or implied, of the DARPA or the U.S. Government.

## References

- Allen, J., Blaylock, N., & Ferguson, G. (2002). A Problem Solving Model for Collaborative Agents. *1st Int'l. Joint Conf. on Autonomous Agents and Multi-Agent Systems (AAMAS-02)*, Bologna, Italy, ACM Press
- Allen, J., Chambers, N., Ferguson, G., Galescu, L., Jung, H. Swift, M. & Taysom, W. (2007). PLOW: A collaborative task learning agent. *National Conference on Artificial Intelligence (AAAI)*. Vancouver, BC.

- Barbu, A., Narayanaswamy, S., & Siskind, J. (2010). Learning physically-instantiated game play through visual observation. *Proc. Of ICRA '10*, pp. 1879-1886.
- Forbus, K., Hinrichs, T., de Kleer, J. & Usher, J. (2010) FIRE: Infrastructure for experience-based systems with common sense. *AAAI Fall Symposium on Commonsense Knowledge*. Arlington, VA.
- Forbus, K., Klenk, M., & Hinrichs, T. (2009). Companion Cognitive Systems: Design Goals and Lessons Learned So Far. *IEEE Intelligent Systems*, vol. 24, no. 4, pp. 36-46, July/August.
- Forbus, K., Usher, J., Lovett, A., & Wetzel, J. (2011). CogSketch: Sketch understanding for Cognitive Science Research and for Education. *Topics in Cognitive Science*. pp 1-19.
- Hobbs, J., Stickel, M., Appelt, D., & Martin, P. (1990). Interpretation as Abduction. *Artificial Intelligence*, 63, 69-142.
- Kaiser, L. (2012) Learning games from videos guided by descriptive complexity. *Proceedings of AAAI-2012*, pp. 963-969.
- Kirk, J. & Laird, J. (2013) Learning Task Formulations through Situated Interactive Instruction. Submitted to ACS 13.
- Love, N., Hinrichs, T., & Genesereth, M. (2006). General Game Playing: Game Description Language Specification. Stanford Logic Group Technical Report LG-2006-01.
- Muggleton, S. & De Raedt, L. (1994) Inductive Logic Programming: Theory and Methods. *Journal of Logic and Algebraic Programming*, vol. 19/20, pp. 629-679.
- Tomai, E. (2009). A Pragmatic Approach to Computational Narrative Understanding (Tech. Rep. No. NWU-EECS-09-17). Doctoral dissertation, Northwestern University, Department of Electrical Engineering and Computer Science, Evanston, Illinois.