

Learning Player Strategies Using Weak Constraints

Author

ELLIOT GREENWOOD

Supervisor

DR. KRYSIA BRODA

Co-Supervisor

MARK LAW

Second Marker

PROF. ALESSANDRA RUSSO

IMPERIAL COLLEGE LONDON

DEPARTMENT OF COMPUTING

18th June 2018

ABSTRACT

Computer programs are now able to easily beat human world champions at complex games like *Chess* and *Go*, but the reasoning behind *how* these programs choose their moves remains relatively obscure. This is due to popular methods using neural networks, which are difficult to interpret without heavy analysis.

In this work we explore an alternative method of learning strategies from gameplay using Inductive Logic Programming (ILP), a logical learning framework that has been extended to the Answer Set Programming (ASP) paradigm. Specifically, we use *weak constraints* from ASP to learn preferences between board states based on human-driven examples. Learning weak constraints is a novel technique that has recently been introduced into ILASP, a learning system created at Imperial College London.

We discuss and show through experimentation ILASP's suitability to learning strategies through weak constraints. We provide a methodology for describing, abstractly, game mechanics using the Game Description Language, learning the rules of the game, strategies to use, and employing the learnt strategies with the aid of a planner. In game theory, search trees are used to employ minimax, a way of expressing optimal play by reasoning about the value of moves in the future. We take this notion and extend existing ILP frameworks to rank moves with respect to game trees.

ACKNOWLEDGEMENTS

I would like to thank my supervisor Dr. Krysia Broda for her guidance, advice and enthusiasm throughout the project. I wish to express my gratitude to Mark Law for the time and help he has dedicated towards this project and allowing me to use ILASP, he has also given me countless ideas during the last year.

I also wish to thank Lottie, for her incredible support and patience over the past few years.

Thank you to Jack and Dan, for all your help and making me laugh. For playing hours of games against me, I would like to say thank you to James, Michael, Florian, Jonathan.

Finally, thank you to my parents, my sister, and my friends for their encouragement during my time at Imperial.

CONTENTS

I	PRELIMINARIES	1
1	INTRODUCTION	3
1.1	Motivation	3
1.2	Objectives	4
1.3	Contributions	4
2	GAMES UNDER STUDY	7
2.1	Onitama	7
2.1.1	Rules	7
2.1.2	Three Card Variation	8
2.1.3	Example Game	9
2.2	Five Field Kono	14
2.2.1	Rules	14
2.3	Cross-Dot Game	15
2.3.1	Rules	15
3	BACKGROUND	17
3.1	Normal Logic Programs	17
3.1.1	Syntax	17
3.1.2	Herbrand Models	19
3.1.3	Stable Model Semantics	20
3.2	Answer Set Programming	22
3.2.1	Choice Rules	22
3.2.2	Weak Constraints	23
3.3	Inductive Logic Programming	24
3.4	Inductive Learning of Answer Set Programs	25
3.4.1	Learning from Answer Sets	25
3.4.2	Learning Weak Constraints	28
3.4.3	Context Dependent Examples	30
3.4.4	Learning from Noisy Examples	32
3.4.5	Constraining the Hypothesis Space with Bias Constraints	33
3.4.6	ILASP Meta-Level Representation	34
3.5	Game Theory	35
3.5.1	Game Types	35
3.5.2	Utility Functions	35
3.5.3	Minimax Theorem	35
4	RELATED WORK	39
4.1	Knowledge Representation	39
4.2	Representing Games in Formal Logics	39
4.3	Learning Answer Set Programs	40
4.4	Machine Learning and Games	40
4.5	Explainable AI	40

II	IMPLEMENTATION	43
5	GAME MODEL	45
5.1	Intuition	45
5.2	Game Description Language	46
5.2.1	Specification	46
5.2.2	Translation into asp	50
5.3	Simplifications	52
6	DIGITISED GAME & PLANNER	53
6.1	Program Flow	53
6.1.1	Example Collection	54
6.1.2	Minimax Planner	54
6.1.3	Assistive Movement	55
6.2	Extensibility	57
6.2.1	Adding new games	57
7	LEARNING PREFERENCES FROM GAME TREES	59
7.1	Motivation	59
7.2	Inductive Learning Programs with Deep Orderings	60
7.3	Implementation	65
7.3.1	ILASP with Meta-Program Injection	65
7.4	Translation to Context Dependent LOAS Task with Meta-Program Injection	65
7.5	Automatically Generating the Game Trees	68
8	CASE STUDY: CROSS-DOT GAME	73
8.1	The Game	73
8.1.1	Representation	73
8.2	Learning Strategies	74
8.2.1	Simple Strategies	74
8.2.2	Combined Strategies	75
8.2.3	Forward-Thinking Strategies	77
8.3	Comparison	78
III	EVALUATION	81
9	LEARNING THE GAME RULES	83
9.1	Process	83
9.2	Learning	83
10	LEARNING AND EXPRESSING STRATEGIES	87
10.1	Immediate Strategies	87
10.1.1	Winning	87
10.1.2	Capturing Piece	88
10.1.3	Space Advantage	90
10.2	Complex Strategies	92
10.3	Tournaments	94
10.4	Summary	94
11	CONCLUSION	99
11.1	Achievements	99
11.2	Future Work	100
11.2.1	Performing Quiescent Search with Weak Constraints	100
11.2.2	Identifying Examples with Strong Strategic Choices	101

IV	APPENDIX	103
A	LOGIC PROGRAMS	105
B	ILASP LEARNING EXAMPLES	111
B.1	Onitama	111
B.1.1	Experiment 9.1: Onitama Rules	111
B.1.2	Experiment 10.7: Defend Pawns	113
B.2	Five Field Kono	121
B.2.1	Experiment 9.2: Five Field Kono Rules	121
B.3	Cross-Dot	123
B.3.1	Experiment 9.3: Cross-Dot Rules	123
	BIBLIOGRAPHY	125

LIST OF FIGURES

Figure 2.1	Onitama setup configuration, with the cards OX, MONKEY, CRANE, RABBIT, COBRA.	7
Figure 2.2	Use of the rabbit card	8
Figure 2.3	Use of the rabbit card with the opposite player	9
Figure 2.4	Move Cards featured in the Game	10
Figure 2.5	Example game following the moves from Example 2.1.3	10
Figure 2.6	Example game following the moves from Example 2.1.3 (cont.)	11
Figure 2.7	Example game following the moves from Example 2.1.3 (cont.)	12
Figure 2.8	Example game following the moves from Example 2.1.3 (cont.)	13
Figure 2.9	Legal moves for the red player	14
Figure 2.10	The initial state of the game of Five Field Kono alongside some winning states	14
Figure 2.11	Some example states of the Cross-Dot Game ($m = 6, k = 2$)	15
Figure 3.1	Graphical representation of X in Example 3.1.5	21
Figure 3.2	Simplified graphical representation of the background facts of B in Example 3.4.2	27
Figure 3.3	Examples e_1, e_2, e_3, e_4 in their respective contexts with the preferred action shown by the highlighted move.	31
Figure 3.4	Some of Alan's moves from a particular game against Betty	32
Figure 3.4	Some of Alan's moves from a particular game against Betty (cont.)	32
Figure 3.5	Decision tree of Nim with backpropagation	36
Figure 3.6	α - β pruning tree, dashed lines denote the pruned branches.	37
Figure 3.7	Result of using an insufficient quiescent search	38
Figure 4.1	Performance vs. Explainability (adapted from Gunning (2016))	41
Figure 4.2	Visualisations of neurons from Layer 4c. Figure from Olah et al. (2018)	41
Figure 5.1	Game Description Language (GDL) Current State	47
Figure 5.2	Game Description Language (GDL) Initial State	47
Figure 5.3	Game Description Language (GDL) Next State	48
Figure 5.3	Game Description Language (GDL) Next State	48
Figure 5.4	Game Description Language (GDL) Legal and Chosen Moves	49
Figure 5.5	Game Description Language (GDL) Goal Relation	49
Figure 5.6	Translation of Tic-Tac-Toe program	51
Figure 6.1	User flow through the digital game	53
Figure 6.2	Current board state midway through a game	55
Figure 6.3	Possible leaf nodes of the game tree pruned to depth 2	56
Figure 6.3	Possible leaf nodes of the game tree pruned to depth 2 (cont.)	56
Figure 6.4	Example of a user correcting a move to a more defensive alternative	57
Figure 7.1	Possibility of an exchange	59
Figure 7.2	Play is considered to depth 0	60
Figure 7.4	Play is considered to depth 2	61
Figure 7.3	Play is considered to depth 1	61
Figure 7.5	Cross-Dot <i>defence</i> game tree generated from mid-game state	70

Figure 8.1	Cross-Dot <i>defence</i> game tree generated from mid-game state . . .	78
Figure 10.1	Illustrations of how valuing different areas of the board affects strategy	91
Figure 10.2	Illustrations of how controlling space on a board can give be advantageous. Grey areas represent spaces that the cards allow each player to move to. The strategies learnt in Experiment 10.5 and Experiment 10.7 both have this concept in their mode declarations (<i>valid_translation</i> ($\cdot, \cdot, \cdot, \cdot$))	91
Figure 10.3	Guaranteed win by following the moves shown	92

LIST OF TABLES

Table 3.1	Wrapper predicates to indicate structural elements of rules in the hypothesis space	33
Table 5.1	Game Description Language (GDL) Keywords	46
Table 5.2	Rules generated for normal/simplified rules with different parameters	52
Table 8.1	Categorised strategy rules from Zhang and Thielscher (2015) . . .	74
Table 10.1	Hypothesis tests between various strategies,★ represents a significant value	95
Table 10.2	Summary of experiments performed throughout the report	96

LIST OF LOGIC PROGRAMS

Logic Program A.1	Onitama Background Knowledge	105
Logic Program A.2	Five Field Kono	107
Logic Program A.3	Cross-Dot Game	108
Logic Program B.1	Examples of legal and illegal moves in Onitama	111
Logic Program B.2	Deep Ordering Translation for the Defend Task	114
Logic Program B.3	Examples of legal and illegal moves in Five Field Kono	121
Logic Program B.4	Examples of legal and illegal moves in Cross-Dot	123

LIST OF ALGORITHMS

Algorithm 7.1	Branch Generation	69
Algorithm 7.2	Children Generation	69

LIST OF DEFINITIONS

3.1.1	Definition (literal)	17
3.1.2	Definition (rule)	17
3.1.3	Definition (logic program)	18
3.1.4	Definition (variable safety)	19

3.1.5 Definition (Herbrand Base)	19
3.1.6 Definition (Herbrand Interpretation)	19
3.1.7 Definition (Herbrand Model)	19
3.1.8 Definition (Least Herbrand Model)	19
3.1.9 Definition (reduct)	20
3.2.1 Definition (choice rule)	22
3.2.2 Definition (weak constraint)	23
3.4.1 Definition (language bias)	25
3.4.2 Definition (partial interpretation)	26
3.4.3 Definition (Learning from Answer Sets task)	26
3.4.4 Definition (hypothesis length)	28
3.4.5 Definition (language bias)	29
3.4.6 Definition (ordering example)	29
3.4.7 Definition (Learning from Ordered Answer Sets task)	29
3.4.8 Definition (Context Dependent Partial Interpretation)	30
3.4.9 Definition (Context Dependent Ordering Example)	30
3.4.10 Definition (Context Dependent Learning from Ordered Answer Sets (LOAS) task)	30
3.4.11 Definition (reify)	34
3.4.12 Definition (append)	34
7.2.1 Definition (Minimax Phases)	61
7.2.2 Definition (Explanation Condition)	62
7.2.3 Definition (Deep Context Dependent Ordering Example)	64
7.2.4 Definition (Deep, Context Dependent LOAS task)	64
7.3.1 Definition (Meta-Program Injection)	65

ACRONYMS

- AI Artificial Intelligence
- ASP Answer Set Programming
- CDOE Context Dependent Ordering Example
- CDPI Context Dependent Partial Interpretation
- GDL Game Description Language
- GGP General Game Playing
- ILASP Inductive Learning of Answer Set Programs
- ILP Inductive Logic Programming
- KIF Knowledge Interchange Format
- LAS Learning from Answer Sets
- LOAS Learning from Ordered Answer Sets
- NAF Negation as Failure
- PGN Portable Game Notation

Part I

PRELIMINARIES

1 | INTRODUCTION

Humans have been playing and mastering games for thousands of years. The earliest games date back to the ancient Egyptian and Mesopotamian era, with games such as *Senet*¹ and the *Royal Game of Ur*². Nowadays, people all over the world compete in fierce championships to prove themselves to be the best *Chess* and *Go* players, but in recent years computers have begun to surpass human level skill in many games (Mnih et al., 2013; Silver et al., 2017b). The rate at which these programs are learning is incredible, with world champions describing the games they play as alien (Peter Heine Nielsen, 2017). However, the moves that the programs make appear counter-intuitive or simply inexplicable even by masters in the field, such as Michael Redmond³.

In this project we explore learning game strategies using a Inductive Logic Programming (ILP) system, ILASP (Law, Russo and Broda, 2014), built on top of a logical paradigm known as Answer Set Programming (ASP). We utilise *weak constraints* as the method of comparing board states and explore what strategies, if any, can be expressed using them. The advantages of using a logical framework is that they are explainable. Strategies that can be learnt in an ILP system can be conveyed in English.

1.1 MOTIVATION

Inductive Learning of Answer Set Programs (ILASP) has recently been extended to learn preferences through weak constraints (Law, Russo and Broda, 2015a). By specifying an ordering over particular concrete examples it learns generic rules that can be used to score future examples. Whilst it has been used in some small cases to learn very guided hypotheses (Law, Russo and Broda, 2015b), in this project we aim to use ILASP to learn about more open ended, practical problems.

Board games have long been used in the field of machine learning as a way of demonstrating the power of a particular system, e. g. Reinforcement Learning in *AlphaZero* to learn the game *Go*. Games are a convenient way of demonstrating machine learning for several reasons:

- they are easy to follow,
- it is easy to determine the success or failure of the agent,
- the environment is self-contained and therefore fairly simple.

Learning strategies in board games is a domain specific view of the more general problem of *preference learning*. When a player chooses a move to play they are expressing their preference for that move over other possible options. Preference learning is a problem that has been attempted in many fields for other purposes (e. g. PageRank Algorithms, Liu (2007)). Preference learning in ILASP tries to find a minimal hypothesis that *exactly*

¹ <https://www.boardgamegeek.com/boardgame/2399/senet>

² http://www.britishmuseum.org/research/collection_online/collection_object_details.aspx?objectId=8817&partId=1

³ When talking about move 37 of game 2, *AlphaGo* vs. Lee Sedol

describes the non-noisy examples, the hypotheses can contain variables, allowing them to generalise over other examples. We explore the effectiveness and suitability of using ILASP, developed by Law, Russo and Broda (2014) at Imperial College London, as a solution to this problem.

In this project we primarily use the strategy game *Onitama* to test and explore ILASP. *Onitama* was chosen as it has not been solved but is not as complex as *Chess*. The game also is a little different to traditional abstract strategy games. Specifically, the movement of the pieces changes every turn (in a predictable manner) through the use of a set of cards and so having a strategy that prefers a particular piece is not viable (cf. the Queen in *Chess*), instead ILASP will need to learn higher level movement features. Additionally, *Onitama* has a random set up and so there is no similar notion of standard opening theory that exists in *Chess*.

1.2 OBJECTIVES

The objectives for this project are:

- formulate a methodology of learning a game from rules to winning;
- learn the rules of various games from incomplete examples of legal and illegal moves, for completeness of the methodology;
- explore what strategies can be expressed using weak constraints and potential extensions to current ILASP tasks in order to cover a larger space;
- evaluate the effectiveness of the strategies learnt by ILASP against various hard-coded strategies.

1.3 CONTRIBUTIONS

There are four main contributions in this project. Firstly, we provide an exploration of ILASP's ability to learn immediate strategies, starting with choosing an appropriate representation for the game in logic (Chapter 5). Strategies such as capturing, evading capture, and controlling space are explored in Chapter 10. In order to complete the study of learning games, Chapter 9 shows ILASP's ability to learn the rules of the game after being shown some, but not all, legal and illegal moves in only a handful of board states.

We have implemented a planner that can run in two modes, *training-mode* and *tournament-mode*, which collect different types of examples to use when learning (Chapter 6). It interfaces directly with Clingo (Kaminski, 2014), an answer set solver, in an intelligent manner in order to batch evaluate many states.

We present experimentation into learning game strategies using two techniques in Chapter 10. The first is to observe the moves of the winner of a game in order to learn their strategy, the second involves having a player critique the AI whilst playing and suggesting better moves and learning solely from these counter examples.

Finally, we present an extension to the Learning from Answer Sets (LAS) framework, ILP_{LOAS}^{deep} , that is capable of learning strategies that reason into the future (Chapter 7 and Section 10.2), by not comparing examples of the same board state, but by incorporating the minimax theorem into the search of the game tree in order to find learning examples

that express the strategy further down the tree. An example of a strategy that would need this technique is learning to exchange pieces on the board. Along with this extension, we present a proof of the correspondence between the set of orderings chosen by the system and the minimax theorem. The creation of the additional examples can be achieved in a *partially* automated manner given the representation of the games described in Chapter 5.

2 | GAMES UNDER STUDY

2.1 ONITAMA

Onitama is a two-player, perfect information, abstract strategy game. Unlike traditional strategy games such as chess, *Onitama* has a random starting configuration. Additionally, the players' possible moves are determined by *move cards*, which vary from game to game.

Each player has two types of pawns: one *master* (♠) and four *student* (♣) pawns. The master starts on the respective players' *temple* squares.

2.1.1 RULES

OBJECTIVE

There are two winning conditions in the game, achieved either by capturing the opponent's master pawn, or by moving one's master pawn to the opponent's *temple* square.

SETUP

The game is played on a 5-by-5 grid, with each player starting with one master and four student pawns.

Five *move cards* are then drawn and each player is given two, with the last placed in the middle.

Remark. The centre card denotes who will start by means of a coloured icon on the card; in diagrams I use the side of the board to denote this. The right side of the board is blue, the left is red, i. e. the card is always on the players' right. In Figure 2.1 the blue player is starting.




Figure 2.1. Onitama setup configuration, with the cards OX, MONKEY, CRANE, RABBIT, COBRA.

MOVING

Move cards show how a piece can move; a player may move any of their pieces according to the cards. A pawn may jump over other pawns and a player may capture an opponent's pawn only by moving on to the space that it occupies. Once a *move card* has been used it is swapped with the *move card* in the centre.

A move card shows a 5×5 grid, a central dark grey square representing the pawn's starting location, lighter grey squares representing valid moves for that pawn (these are relative to the start location). Example 2.1.1 shows all valid moves from the RABBIT card.

Remark. The location of a pawn is denoted by the pawn's symbol followed by a tuple (Row, Column) where (1, 1) is the bottom-left, e. g. ♞ (2, 3) means the red master at row 2, column 3, and ♜ (5, 5) means a blue student at row 5, column 5 (i. e. the top-right corner).

Example 2.1.1. Figure 2.2 below shows the use of the rabbit card on the blue player's turn. It depicts several key points:

- The *move card* is being used by both ♞ and ♜ pawns
- Pawns may jump over pawns, e. g. ♞ (3, 1) jumping over ♞ (3, 2)
- Pawns cannot move to a position occupied by a pawn of the same team, e. g. ♞ (2, 2) cannot move to (2, 4)
- Pawns cannot move off the board, nor does it wrap, e. g. ♞ (3, 1) and ♞ (2, 4)
- Pawns can capture opponent pawns, e. g. ♞ (3, 1) captures ♞ (4, 2)



Figure 2.2. Use of the rabbit card

Example 2.1.2. Figure 2.3 shows the rabbit card being used by the red player. The only point of note here is that the card has been rotated 180° to face the red player.

2.1.2 THREE CARD VARIATION

A game of Onitama normally uses five cards, offering players a maximum of 8 moves per pawn to choose from, and allowing players to selectively withhold cards from their



Figure 2.3. Use of the rabbit card with the opposite player

opponent. This creates a much larger state space and invites highly complex strategies. To simplify the problem whilst learning, I used just three cards; one per player and a centre card. It follows that players can now deterministically calculate the cards they will receive and therefore have certainty over which moves will be available to them in the future.

2.1.3 EXAMPLE GAME

Games of Onitama can be expressed in a modified Portable Game Notation (PGN) (Edwards, 1994). A halfmove is denoted in a similar way to a pawn's location: the pawn's symbol, starting and ending location, and the card name used. A full move is denoted by the move number, followed by the halfmove of each player. A capture is denoted using a \times between the starting and ending locations.

Example 2.1.3. The following is a short example of a full game of Onitama, with the board states depicted in Figure 2.8. The cards used in this game can be found in Figure 2.4.

- | | |
|---------------------------------|------------------------------|
| 1. ♁ (1,2)(2,2) MOUSE | ♁ (5,4)(4,4) BOAR |
| 2. ♁ (1,3)(2,3) PANDA | ♁ (5,5)(4,5) MOUSE |
| 3. ♁ (1,4)(2,4) BOAR | ♁ (4,4)(3,3) PANDA |
| 4. ♁ (2,3) \times (3,3) MOUSE | ♁ (5,2)(4,2) BOAR |
| 5. ♁ (1,5)(2,5) PANDA | ♁ (4,5)(4,4) MOUSE |
| 6. ♁ (3,3)(3,4) BOAR | ♁ (4,4) \times (3,4) PANDA |





Figure 2.4. Move Cards featured in the Game



(a) Initial Board

Figure 2.5. Example game following the moves from Example 2.1.3




Figure 2.4. Move Cards featured in the Game (repeated from page 10)




Figure 2.6. Example game following the moves from Example 2.1.3 (cont.)




Figure 2.4. Move Cards featured in the Game (repeated from page 10)

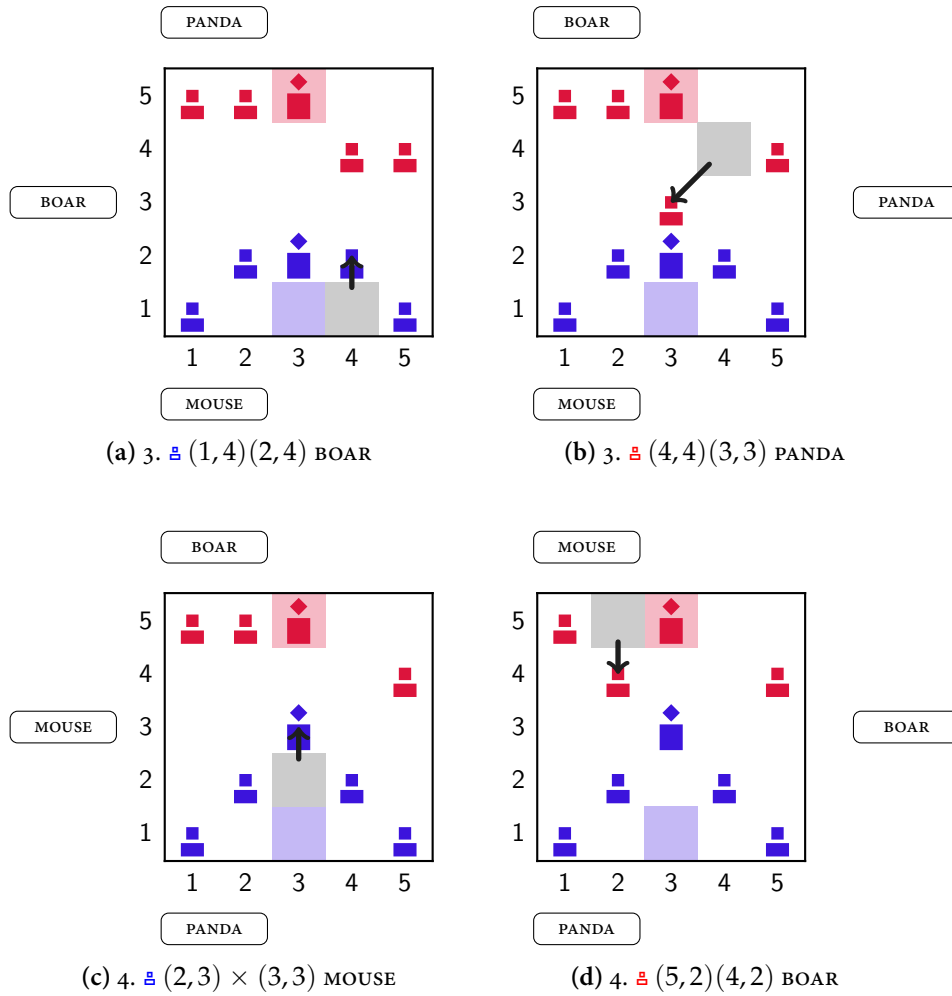


Figure 2.7. Example game following the moves from Example 2.1.3 (cont.)




Figure 2.4. Move Cards featured in the Game (repeated from page 10)

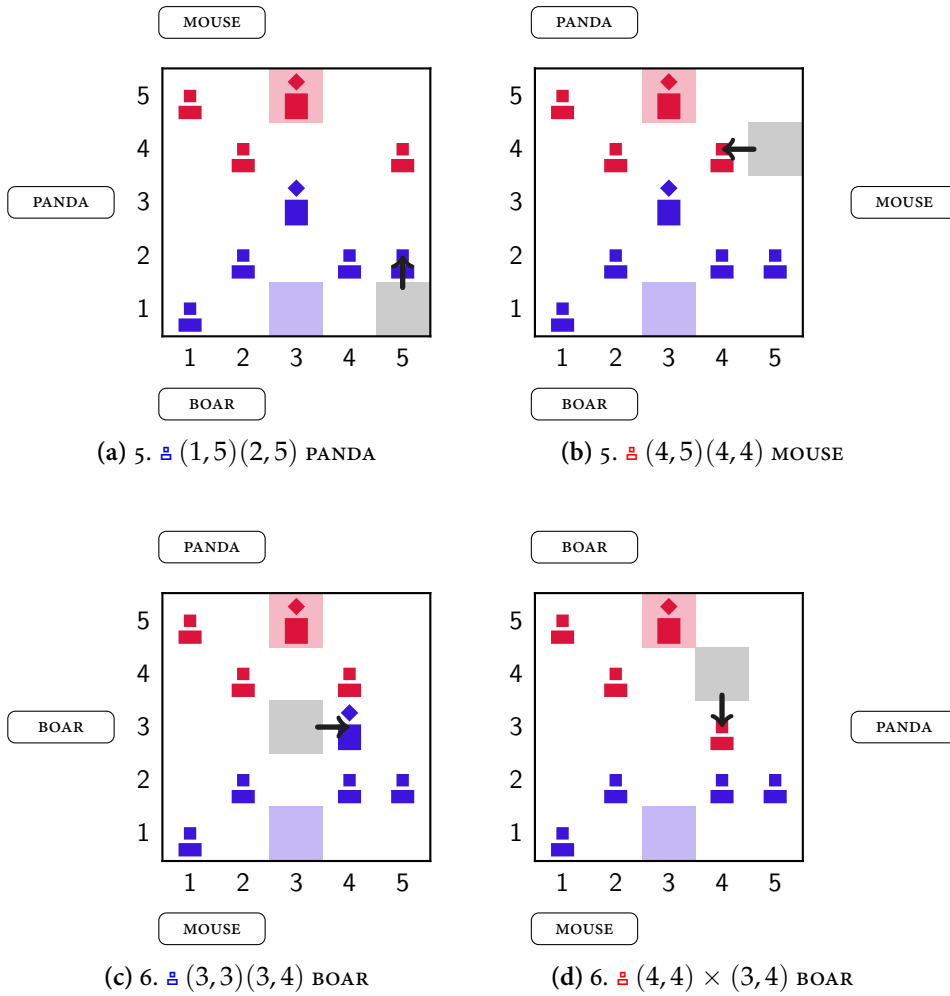


Figure 2.8. Example game following the moves from Example 2.1.3 (cont.)

2.2 FIVE FIELD KONO

Five Field Kono is a checkers-like strategy game originating in Korea. Two players try to block and out manoeuvre each other in order to occupy the starting spaces of their opponent. All counters move in the same fashion, diagonally, and there are no captures in the game.

In this report I use Five Field Kono to demonstrate the techniques used throughout in the setting of a different game.

2.2.1 RULES

MOVES

All counters can move one space along any diagonal, but cannot jump pieces and cannot occupy the same space as another counter, as illustrated in Figure 2.9.




Figure 2.9. Legal moves for the red player

OBJECTIVE

A player is trying to move from their starting position to their opponent's starting position before their opponent does the same. If a player leaves counters in their starting position they count in their opponent's favour, thus their opponent only has to occupy the vacated spaces. A player must have at least one of their counters in their opponent's starting positions in order to win. Figure 2.10 shows the starting and some winning conditions of the game.

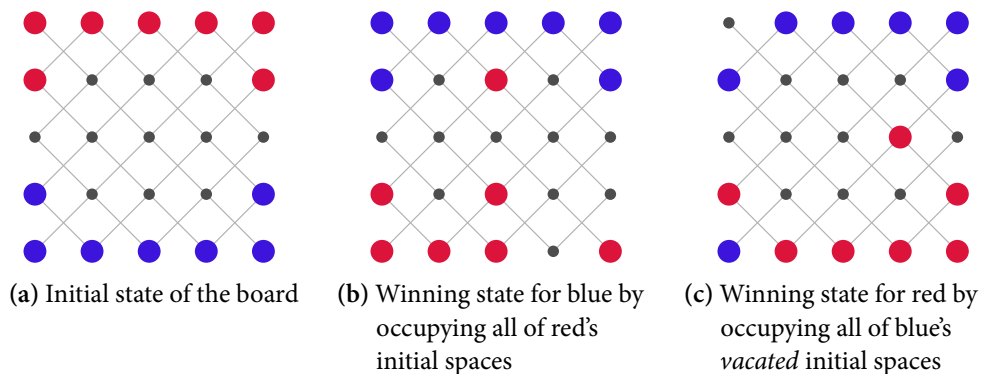


Figure 2.10. The initial state of the game of Five Field Kono alongside some winning states

2.3 CROSS-DOT GAME

The Cross-Dot game is a variation of Tic-Tac-Toe which I have taken from the literature (Zhang and Thielscher, 2015). The Cross-Dot game, another name for an m - k game, where players aim to get k boxes in a row marked with their player symbol (from a possible m boxes).

2.3.1 RULES

The rules of the game are very straight-forward, on your turn you may mark any empty box with your marker (\times or \cdot). The board is initially completely empty, and the game ends when one of the players achieves k contiguous boxes for the win, or when all of the boxes are filled. The \times player starts the game.





Figure 2.11. Some example states of the Cross-Dot Game ($m = 6, k = 2$)

3 | BACKGROUND

3.1 NORMAL LOGIC PROGRAMS

Normal logic programs extend definite logic programs to include the notion of Negation as Failure (NAF). Negation as Failure, denoted by the operator `not`, differs from classical negation (\neg). Informally, `not p` means that it cannot be proven that p holds. Below is an example of NAF to illustrate its semantics.

Example 3.1.1. Below is an example game state of *Onitama* in which the blue player has just won.



On the right is a set of predicates that we know to be facts about the board. Below is a rule about the *master* being captured.

$$master_captured(P) \leftarrow player(P), not\ master_on_board(P).$$

From this rule we can deduce that $master_captured(red)$, because it cannot be shown that $master_on_board(red)$ is true. However, $master_captured(blue)$ cannot be deduced because, by the set of facts above, the blue master is on the board.

3.1.1 SYNTAX

Logic programs are constructed from *terms*, which consist of variables, constants, and n-ary function symbols, *atoms*, which are n-ary predicates defined over terms, and *literals* and *clauses* (defined below).

Definition 3.1.1 (literal). A *literal* is an expression of the form A , a *positive* literal, or $not\ A$, a *negative* literal, where A is an atom.

Definition 3.1.2 (rule). *Clauses*, or *rules*, are formulæ of the form

$$h \leftarrow b_1, \dots, b_m, not\ c_1, \dots, not\ c_n \tag{r}$$

where h , each b_i , $i \in [1, m]$ and c_j , $j \in [1, n]$ are atoms, $m \geq 0$, $n \geq 0$.

$head(r) = h$, is the head of the rule, $body^+(r) = \{b_1, \dots, b_m\}$ are the positive body literals, and $body^-(r) = \{c_1, \dots, c_n\}$ are the negative body literals.

A *fact* is a rule, r , with no body literals, i. e. $body^+(r) \cup body^-(r) = \emptyset$. A *constraint*, or *hard constraint*, is a rule, r , with no head, i. e. $head(r) = \perp$.

A *ground* rule is one which contains no variables, i. e. all terms that appear in the clause are constants. The set of ground rules created by substituting all variables with all combinations of atoms from the language is known as the *ground instances*. For example, the ground instances¹ of $master_on_board(Player) \leftarrow location(pawn(master, Player), Cell)$ are:

$$\left\{ \begin{array}{l} master_on_board(red) \leftarrow player(red), location(pawn(master, red), cell(1, 1)) \\ master_on_board(blue) \leftarrow player(blue), location(pawn(master, blue), cell(1, 1)) \\ master_on_board(red) \leftarrow player(red), location(pawn(master, red), cell(1, 2)) \\ master_on_board(blue) \leftarrow player(blue), location(pawn(master, blue), cell(1, 2)) \\ \dots \\ master_on_board(red) \leftarrow player(red), location(pawn(master, red), cell(3, 3)) \\ master_on_board(blue) \leftarrow player(blue), location(pawn(master, blue), cell(3, 3)) \\ \dots \\ master_on_board(red) \leftarrow player(red), location(pawn(master, red), cell(5, 5)) \\ master_on_board(blue) \leftarrow player(blue), location(pawn(master, blue), cell(5, 5)) \end{array} \right.$$

Definition 3.1.3 (logic program). A *logic program* is a set of rules. The *grounding* of the logic program is the union of the sets of ground instances of each rule.

Example 3.1.2. Let the following *normal logic program*², $\Pi_{location}$, represent where pawns are on the board.

$$location(pawn(master, red), cell(2, 1)). \quad (3.1)$$

$$location(pawn(master, blue), cell(3, 2)). \quad (3.2)$$

$$next(location(pawn(Rank, Player), Cell)) \quad (3.3)$$

$$\begin{aligned} &\leftarrow location(pawn(Rank, Player), Cell), \\ &\quad \text{not } does(_, move(_, Cell, _)), \\ &\quad \text{not } does(Player, move(Cell, _, _)). \end{aligned}$$

$$next(location(pawn(Rank, Player), To)) \quad (3.4)$$

$$\begin{aligned} &\leftarrow location(pawn(Rank, Player), From), \\ &\quad does(Player, move(From, To, _)). \end{aligned}$$

$$\leftarrow location(P1, Cell, T), location(P2, Cell, T), P1 < P2. \quad (3.5)$$

Rules 3.1 and 3.2 are *facts*, they represent where the two master pawns are currently. Rules 3.3–3.4 are *normal rules*. They describe where a pawn is at the next time step. Rule 3.3 says a pawn will be in a cell if it is currently in that cell, nothing has been moved

¹ For ease of reading I have only included groundings that make semantic sense, i. e. not substituting the *Cell* variable with *red*, although they are a part of the set of ground instances

² $_$ is an anonymous variable, meaning we do not care about its value

there, and it has not been moved (i. e. a pawn is in the same location as before and it has not been captured or moved). When a pawn is moved rule 3.4 says that it will be in the new location at the next time step. There is a subtle difference between the use of the anonymous variables used in rule 3.3 and those in rule 3.4. In rule 3.3 they create a projection of the *does* predicate, and the meaning becomes “there is not a move that ends in this location”. If the anonymous variables were actual variables³ then even if there was a move that ended in *Cell*, you could satisfy the *NAF* literal with any other player, or card, for example. On the other hand, in rule 3.4 the anonymous variable can be substituted with a variable as the meaning simply means “for some card”. Finally, rule 3.5 is a *constraint* saying that only one piece can be in a square at any given time.

Note. Rule 3.5 uses $<$ for inequality instead of \neq . This is because using \neq means the constraint is symmetric, doubling its grounding. All terms are totally ordered and so these are equivalent.

Definition 3.1.4 (variable safety). A variable is *safe* within a rule iff it appears in a positive body literal. A rule is *safe* iff all of its variables are safe.

Example 3.1.3. The following rule are *not* safe.

$$\text{opponent}(\text{pawn}(\text{Rank1}, \text{red}), \text{pawn}(\text{Rank2}, \text{blue})). \quad (3.6)$$

$$\begin{aligned} \text{center_card}(\text{Card}) \leftarrow \text{not } \text{in_hand}(\text{red}, \text{Card}), \\ \text{not } \text{in_hand}(\text{blue}, \text{Card}). \end{aligned} \quad (3.7)$$

Rule 3.6 has variables (*Rank1*, *Rank2*) in the head of a rule with no body, therefore they appear in no positive body literal. In rule 3.7 the variable *Card* only appears in negative body literals.

3.1.2 HERBRAND MODELS

Definition 3.1.5 (Herbrand Base). The *Herbrand Base* of a program, Π , is the set of all ground atoms that can be constructed from predicates in Π and the ground terms that can be constructed using the terms and function symbols that occur in Π . It is denoted by $\text{atoms}(\Pi)$.

Definition 3.1.6 (Herbrand Interpretation). Let Π be a *definite* logic program (i. e. one without Negation as Failure), which is written in language \mathcal{L} . A *Herbrand Interpretation* is created by assigning every $\alpha \in \text{atoms}(\Pi)$ either *true* (\top) or *false* (\perp). The interpretation is usually written as the set of atoms that have been assigned to true, everything else is false.

Definition 3.1.7 (Herbrand Model). A *Herbrand Model* is a Herbrand Interpretation, I , which satisfies every rule in Π , that is to say for every rule $R \in \Pi$ if $\text{body}^+(R) \subseteq I$ and $\text{body}^-(R) \cap I = \emptyset$ then $\text{head}(R) \in I$.

Definition 3.1.8 (Least Herbrand Model). A *least Herbrand Model* is a Herbrand Model which has a \subseteq -minimal set of *true* ground atoms, which for definite logic programs is always unique. The least Herbrand Model of Π is denoted $M(\Pi)$.

³ You would also need to make the rule safe with types for the new variables

Example 3.1.4. Take the following logic program Π :

$$\text{opponent}(\text{red}, \text{blue}) \quad (3.8)$$

$$\text{opponent}(\text{blue}, \text{red}) \quad (3.9)$$

$$\text{control}(\text{red}) \quad (3.10)$$

$$\text{next}(\text{control}(\text{Player})) \leftarrow \text{control}(\text{Opp}), \text{opponent}(\text{Player}, \text{Opp}). \quad (3.11)$$

and the set $\text{atoms}(\Pi)$ ⁴:

$$\left\{ \begin{array}{l} \text{opponent}(\text{red}, \text{blue}), \text{opponent}(\text{blue}, \text{red}), \text{control}(\text{blue}), \text{control}(\text{red}), \\ \text{next}(\text{control}(\text{red})), \text{next}(\text{control}(\text{blue})) \end{array} \right.$$

One possible Herbrand Interpretation is:

$$\left\{ \begin{array}{l} \text{opponent}(\text{blue}, \text{red}), \text{control}(\text{red}), \\ \text{next}(\text{control}(\text{red})), \text{next}(\text{control}(\text{blue})) \end{array} \right.$$

However this is not a model as $\text{opponent}(\text{red}, \text{blue})$ is not true and neither is the rule 3.11. Another Herbrand Interpretation is:

$$\left\{ \begin{array}{l} \text{opponent}(\text{red}, \text{blue}), \text{opponent}(\text{blue}, \text{red}), \\ \text{control}(\text{red}), \text{next}(\text{control}(\text{blue})) \end{array} \right.$$

this interpretation is a Herbrand Model. Further, it is a Least Herbrand Model.

3.1.3 STABLE MODEL SEMANTICS

The stable model semantics shown below is based on that presented by Gelfond and Lifschitz (1988).

Definition 3.1.9 (reduct). Let Π be any ground normal logic program. Let $\text{atoms}(\Pi)$ be the Herbrand Base. Let $X \subseteq \text{atoms}(\Pi)$ be a set of atoms. The *reduct* Π^X is the set of clauses obtained from Π as follows:

- (i) delete any clause in Π that has a formula $\text{not } A$ such that $A \in X$
- (ii) delete all remaining formulæ of the form $\text{not } A$ in the bodies of the remaining clauses

or equivalently:

$$\Pi^X \triangleq \{ \text{head}(r) \leftarrow \text{body}^+(r) \mid r \in \Pi, \text{body}^-(r) \cap X = \emptyset \} \quad (3.12)$$

X is a *stable model*, or *answer set*, of Π iff $X = \text{M}(\Pi^X)$.

⁴ For convenience, we only show the subset of the Herbrand Base that needs to be constructed, other atoms such as $\text{next}(\text{control}(\text{control}(\text{blue})))$ would also be in the Herbrand Base

Example 3.1.5. Let Π be the grounding of a subset of the logic program from Example 3.1.2, with an extra fact denoting the \diamond master's move:

$$\Pi = \text{ground} \left(\begin{array}{l} \text{location}(\text{pawn}(\text{master}, \text{blue}), \text{cell}(3, 2)). \\ \\ \text{next}(\text{location}(\text{pawn}(\text{Rank}, \text{Player}), \text{Cell})) \\ \quad \leftarrow \text{location}(\text{pawn}(\text{Rank}, \text{Player}), \text{Cell}), \\ \quad \text{not } \text{does}(_, \text{move}(_, \text{Cell}, _)), \\ \quad \text{not } \text{does}(\text{Player}, \text{move}(\text{Cell}, _, _)). \\ \\ \text{next}(\text{location}(\text{pawn}(\text{Rank}, \text{Player}), \text{To})) \\ \quad \leftarrow \text{location}(\text{pawn}(\text{Rank}, \text{Player}), \text{From}), \\ \quad \text{does}(\text{Player}, \text{move}(\text{From}, \text{To}, _)). \\ \\ \text{does}(\text{blue}, \text{move}(\text{cell}(3, 2), \text{cell}(2, 1), \text{monkey})). \end{array} \right)$$

Let X be the following set of atoms:

$$\left\{ \begin{array}{l} \text{location}(\text{piece}(\text{master}, \text{blue}), \text{cell}(3, 2)) \\ \text{next}(\text{location}(\text{piece}(\text{master}, \text{blue}), \text{cell}(2, 1))) \\ \text{does}(\text{blue}, \text{move}(\text{cell}(3, 2), \text{cell}(2, 1), \text{monkey})) \end{array} \right.$$




Figure 3.1. Graphical representation of X in Example 3.1.5

Constructing Π^X gives us:

$$\left\{ \begin{array}{l} \text{location}(\text{piece}(\text{master}, \text{blue}), \text{cell}(3, 2)) \\ \text{does}(\text{blue}, \text{move}(\text{cell}(3, 2), \text{cell}(2, 1), \text{monkey})) \\ \text{next}(\text{location}(\text{pawn}(\text{master}, \text{blue}), \text{cell}(3, 2))) \\ \quad \leftarrow \text{location}(\text{pawn}(\text{master}, \text{blue}), \text{cell}(3, 2)), \\ \quad \quad \text{does}(\text{blue}, \text{move}(\text{cell}(3, 2), \text{cell}(3, 2), \text{monkey})). \\ \text{next}(\text{location}(\text{pawn}(\text{master}, \text{blue}), \text{cell}(2, 1))) \\ \quad \leftarrow \text{location}(\text{pawn}(\text{master}, \text{blue}), \text{cell}(3, 2)), \\ \quad \quad \text{does}(\text{blue}, \text{move}(\text{cell}(3, 2), \text{cell}(2, 1), \text{monkey})). \\ \text{next}(\text{location}(\text{pawn}(\text{master}, \text{blue}), \text{cell}(3, 2))) \\ \quad \leftarrow \text{location}(\text{pawn}(\text{master}, \text{blue}), \text{cell}(2, 1)), \\ \quad \quad \text{does}(\text{blue}, \text{move}(\text{cell}(2, 1), \text{cell}(3, 2), \text{monkey})). \\ \text{next}(\text{location}(\text{pawn}(\text{master}, \text{blue}), \text{cell}(2, 1))) \\ \quad \leftarrow \text{location}(\text{pawn}(\text{master}, \text{blue}), \text{cell}(2, 1)), \\ \quad \quad \text{does}(\text{blue}, \text{move}(\text{cell}(2, 1), \text{cell}(2, 1), \text{monkey})). \end{array} \right.$$

When we compute the least Herbrand model of Π^X we get:

$$\left\{ \begin{array}{l} \text{location}(\text{pawn}(\text{master}, \text{blue}), \text{cell}(3, 2)) \\ \text{next}(\text{location}(\text{pawn}(\text{master}, \text{blue}), \text{cell}(2, 1))) \\ \text{does}(\text{blue}, \text{move}(\text{cell}(3, 2), \text{cell}(2, 1), \text{monkey})) \end{array} \right.$$

which is X , and therefore X is an *answer set*.

3.2 ANSWER SET PROGRAMMING

Answer Set Programming (ASP) is a form of declarative programming oriented towards difficult search problems (Lifschitz, 2008). It introduces new concepts such as *choice rules* and *weak constraints* (Calimeri et al., 2013), described below. Throughout this report I will be using the answer set solver Clingo.

3.2.1 CHOICE RULES

Definition 3.2.1 (choice rule). A choice rule is a *rule* with an aggregated head. It is of the form

$$l \{L_1, \dots, L_h\} u \leftarrow B_1, \dots, B_m \quad (3.13)$$

where $l, u \in \mathbb{N}$ are the lower and upper bound, respectively, and L_i where $i \in [1, h]$ is a literal. This aggregated head creates between l and u new rules where the head $L \in \{L_1, \dots, L_h\}$ and the body is B_1, \dots, B_m . Due to the fact there can be many ways of satisfying these conditions choices rules provide a way of generating multiple answer sets.

Example 3.2.1. This example gives the background knowledge for the moves that can be made on a turn T .

$$0 \{ \text{does}(\text{Player}, \text{Action}) \} 1 \leftarrow \text{legal}(\text{Player}, \text{Action}), \text{not } \text{terminal}. \quad (3.14)$$

$$\leftarrow \text{does}(\text{Player}, \text{move}(\text{From1}, \text{To1}, \text{Card1})), \quad (3.15)$$

$$\text{does}(\text{Player}, \text{move}(\text{From2}, \text{To2}, \text{Card2})),$$

$$\text{neq}(\text{From1}, \text{To1}, \text{Card1}) < \text{neq}(\text{From2}, \text{To2}, \text{Card2}).$$

$$\leftarrow \text{legal}(_, _), \text{not } \text{does}(_, _), \text{not } \text{terminal}. \quad (3.16)$$

Rules 3.14–3.16 describe what is allowed to be a move within the game. Rule 3.14 is a *choice rule* with an lower bound of 0 and an upper bound of 1. It states that for each grounding of a valid move we can either choose to make a move or not. Rules 3.15 and 3.16 together enforce that there is at least one unique move at any given time, so long as there is a valid move that can be made. The case when there is no valid move is dealt with separately.

Remark. If Rule 3.14 was a normal rule the program would be unsatisfiable because it would force all valid moves to be chosen moves, violating the uniqueness constraint.

3.2.2 WEAK CONSTRAINTS

Weak constraints were originally introduced in Disjunctive Datalog (Buccafurri, Leone and Rullo, 1997) in order to specify integrity constraints that should be satisfied *if possible*. They were used to express optimisation ideas such as “ensure there are as few timetable clashes as possible”. In this report we are trying to optimise a player’s performance in a game. Thus, we can harness weak constraints to represent tactics in a game that build upon each other in order to describe a strategy.

Definition 3.2.2 (weak constraint). *Weak constraints* create an ordering over $\mathcal{AS}(\Pi)$, specifying which answer sets are “better” than others. Unlike hard constraints, weak constraints do not affect what is or is not in an answer set of a program Π (Law, Russo and Broda, 2015a).

A weak constraint is of the form

$$\leftarrow b_1, \dots, b_m, \text{not } c_1, \dots, \text{not } c_n. [w@l, t_1, \dots, t_o] \quad (3.17)$$

where $b_1, \dots, b_m, c_1, \dots, c_n$ are atoms from \mathcal{L}_H , w and l are terms specifying the *weight* and *level*, and t_1, \dots, t_o are terms. All variables in weak constraint 3.17 must be safe, i. e. all variables that appear in t_1, \dots, t_o must appear in the positive body literals.

When a weak constraint’s body is satisfied a tuple (w, l, t_1, \dots, t_o) is generated, where t_i are the ground terms in the weak constraint. For any ASP program Π and answer set $A \in \mathcal{AS}(\Pi)$ there is a set $\text{weak}(\Pi, A)$ of tuples (w, l, t_1, \dots, t_o) corresponding to weak constraints whose body is satisfied.

Example 3.2.2. The following weak constraints are one possible way of expressing that a player wishes to win first and foremost, and capture pieces whilst maintaining their own

if possible. The $\text{goal}(\cdot, \cdot)$ predicate⁵ is which denotes that a player has received some reward for meeting some game condition (e. g. winning).

$$\leftarrow \text{not } \text{goal}(\text{red}, 100). [5@2] \quad (3.18)$$

$$\leftarrow \text{location}(\text{pawn}(\text{student}, \text{blue}), \text{Cell}) [1@1, \text{Cell}] \quad (3.19)$$

$$\leftarrow \text{location}(\text{pawn}(\text{student}, \text{red}), \text{Cell}) [-1@1, \text{Cell}] \quad (3.20)$$

Weak Constraint 3.18 says that if we cannot show that *red* has won then we incur a penalty of 5. The priority level 2 means that this constraint will be minimised first. Weak constraint 3.19 says that if we find a blue student (♟) on the board then we incur a penalty of 1. The variable *Cell* means we incur this penalty for each unique grounding of these variables, i. e. each blue student on the board. Weak constraint 3.20 says we gain a reward for each red student (♞) on the board. Both weak constraint 3.19 and 3.20 are at priority level 1 and so the penalties are added (e. g. if there were $2 \times \text{♟}$ and $1 \times \text{♞}$ on the board, and it is not a winning state for red, the score would be 5 at level 2, and 1 at level 1).

The semantics for weak constraints presented below are based on those in Calimeri et al. (2013) and Law, Russo and Broda (2015a). For any level $l \in \mathbb{N}$, $A \in \mathcal{AS}(\Pi)$ let

$$P_l^A = \sum_{w \in W_A^l} w$$

denote the penalty of the answer set $A \in \mathcal{AS}(\Pi)$, at level l where $W_l^A = \{w \mid (w, l, t_1, \dots, t_o) \in \text{weak}(\Pi, A)\}$. A_1 dominates A_2 ($A_1 \succ_{\Pi} A_2$) iff $\exists l$ such that $P_l^{A_1} < P_l^{A_2}$ and $\forall l' > l \Rightarrow P_{l'}^{A_1} = P_{l'}^{A_2}$

Example 3.2.3. Given some program Π containing weak constraints 3.18–3.20 such that the *only* answer sets are $\mathcal{AS}(\Pi) = \{A_1, A_2\}$, let

$$\text{weak}(\Pi, A_1) = \{(5, 2), (1, 1, \text{cell}(3, 2)), (1, 1, \text{cell}(2, 4)), (-1, 1, \text{cell}(2, 3))\}$$

$$\text{weak}(\Pi, A_2) = \{(5, 2), (1, 1, \text{cell}(5, 4)), (-1, 1, \text{cell}(3, 4))\}$$

$$P_{A_1}^2 = 5$$

$$P_{A_1}^1 = 1 + 1 - 1 = 1$$

$$P_{A_2}^2 = 5$$

$$P_{A_2}^1 = 1 - 1 = 0$$

therefore $A_2 \succ_{\Pi} A_1$ as the penalties at level 2 are equal and $P_{A_2}^1 < P_{A_1}^1$. Additionally as no answer set dominates A_2 it is *optimal*.

Remark. Clingo uses the notation :- in place of \leftarrow , and :- in place of \leftarrow .

3.3 INDUCTIVE LOGIC PROGRAMMING

The field of Inductive Logic Programming (ILP) essentially combines Machine Learning with logical knowledge representation (Muggleton et al., 2011). The definition of an inductive learning problem is given in Muggleton (1991) and Muggleton and Raedt (1994). The general setting for an inductive logic program has three languages:

⁵ We will later see in Chapter 5 that this predicate is defined by a game description language

\mathcal{L}_E : The language of the examples

\mathcal{L}_B : The language of the background knowledge

\mathcal{L}_H : The language of the hypotheses

The general induction learning problem is then defined as follows: given a set of *ground, atomic* examples or observations $E^+ \subseteq \mathcal{L}_E$, a background knowledge $B \subseteq \mathcal{L}_B$, both consistent, find a hypothesis $H \subseteq \mathcal{L}_H$ such that $B \cup H \models E^+$, that is, the background knowledge together with the learnt hypothesis entail, or *cover*, the observations.

A set of negative examples $E^- \subseteq \mathcal{L}_E$ can also be given (Muggleton and Raedt, 1994), it is then required that the background knowledge together with the hypothesis is *consistent* with respect to the negative examples, $\forall e^- \in E^- . B \cup H \cup e^- \not\models \square$.

3.4 INDUCTIVE LEARNING OF ANSWER SET PROGRAMS

In Law, Russo and Broda (2014) the concept of *ILASP* was introduced. This expanded on previous work in the *ILP* field such as *Progol*, *Metagol* and *HAIL* (Muggleton, 1995; Muggleton et al., 2014, and; Ray, Broda and Russo, 2003). The previous work mainly focussed on definite and normal logic programs, whereas *ILASP* looks at a different class of programs; Answer Set Programs (Law, Russo and Broda, 2014). In this section I will look into the first version of *ILASP*. I also explore further advances, for example, learning from weak constraints (Law, Russo and Broda, 2015a) and learning from context dependent examples (Law, Russo and Broda, 2016).

The problem is structured in a similar manner to the general inductive problem. However the sets of examples, E^+ , E^- , are now no longer individual atoms, but instead *partial interpretations*, meaning a problem from *Progol*, for example, could be constructed with a single example.

In this section, it is assumed that background knowledge and hypotheses are ASP programs, as defined in Section 3.2.

3.4.1 LEARNING FROM ANSWER SETS

In this section, I define the *ILASP* paradigm based on Law, Russo and Broda (2014).

In an *ILP* task the hypothesis space is defined by a language bias, which specifies how \mathcal{L}_H is built.

Definition 3.4.1 (language bias). A Learning from Answer Sets (*LAS*) *language bias* is defined as a pair of mode declarations $M^{\text{LAS}} = \langle M_h, M_b \rangle$, where M_h and M_b are the *head* and *body mode declarations*, respectively. Both M_h and M_b are sets of literals with their arguments replaced with v_{type} and c_{type} denoting ‘variable’ and ‘constant’, respectively, where *type* is the type of the variable or constant.

Given a language bias M^{LAS} , a rule, r , of the form $L_h \leftarrow L_1, \dots, L_m$, not L_{m+1}, \dots , not L_n such that $n \geq m \geq 0$ is in the search space S_M^{LAS} iff

- (i) either
 - a) L_h is empty or
 - b) L_h is compatible with M_h or

c) L_h is an aggregate $l \{h_1, \dots, h_k\} u$ such that $0 \leq l \leq u \leq k$ and $\forall i \in [0, k]$
 h_i is compatible with M_h

(ii) $\forall i \in [1, n]$ L_i is compatible with M_b

(iii) All variables are safe

Finally, each rule $r \in S_M^{\text{LAS}}$ is given a unique identifier r_{id} .

Remark. Informally, a literal is *compatible* with a mode declaration m if there exists an instance of v (resp. c) that can be replaced by a variable (resp. constant).

Example 3.4.1. Let $M^{\text{LAS}} = \langle \{goal(v_{role}, c_{reward})\}, \{opponent(v_{role}, v_{role}), temple(v_{role}, v_{cell}), location(pawn(c_{rank}, v_{role}), v_{cell})\} \rangle$. Then the following rules are in S_M^{LAS} :

$$\begin{aligned} goal(Player, 100) &\leftarrow opponent(Player, Enemy), \\ &\quad temple(Enemy, Temple), \\ &\quad location(pawn(master, Player), Temple). \\ \\ &\leftarrow not\ opponent(Player, Enemy), \\ &\quad temple(Enemy, Cell), \\ &\quad location(pawn(master, Player), Cell). \end{aligned}$$

However the following is not:

$$\begin{aligned} goal(red, Reward) &\leftarrow opponent(red, Enemy), \\ &\quad temple(Enemy, Cell), \\ &\quad location(pawn(master, red), Cell). \end{aligned}$$

because there is no head mode declaration that allows the *goal* predicate with this combination of variables and constants, similarly for the body literal *opponent*, and because the second argument in *pawn* should be a variable according to the language bias.

Definition 3.4.2 (partial interpretation). A *partial interpretation* E is a pair $\langle E^{inc}, E^{exc} \rangle$ of sets of ground atoms from \mathcal{L}_E , called the *inclusions* and *exclusions*, respectively. An Answer Set A *extends* $\langle E^{inc}, E^{exc} \rangle$ iff $E^{inc} \subseteq A$ and $E^{exc} \cap A = \emptyset$.

A partial interpretation E is *bravely* entailed by a program Π iff there exists $A \in \mathcal{AS}(\Pi)$ such that A extends E . E is *cautiously* entailed by Π iff for every $A \in \mathcal{AS}(\Pi)$ A extends E .

Definition 3.4.3 (Learning from Answer Sets task). A *Learning from Answer Sets task* is a tuple $T = \langle B, S_M^{\text{LAS}}, E^+, E^- \rangle$ where B is the background knowledge, S_M^{LAS} is the search space defined by the language bias M^{LAS} , and E^+ and E^- are partial interpretations called, respectively, the positive and negative examples.

A hypothesis H is an inductive solution of T ($H \in \text{ILP}_{\text{LAS}}(T)$) iff:

- (i) $H \subseteq S_M^{\text{LAS}}$
- (ii) $\forall e^+ \in E^+ \exists A \in \mathcal{AS}(B \cup H)$ such that A extends e^+
- (iii) $\forall e^- \in E^- \neg \exists A \in \mathcal{AS}(B \cup H)$ such that A extends e^-

Example 3.4.2. In this example we try to learn who the start player of a game of onitama is. Recall from Section 2.1 that all the *move cards* have a icon depicting either the red or the blue player, here this is represented by $\text{icon}(\cdot, \cdot)$. $\text{does}(\cdot, \cdot)$ denotes an action taken by a player, $\text{control}(\cdot)$ denotes whose turn it is, other predicates have their intuitive meaning.

Let $T = \langle B, S_M^{\text{LAS}}, E^+, E^- \rangle$ be a learning task where B is the following logic program:^{6 7}

$$\text{in_play}(\text{panda}; \text{mouse}; \text{fox}). \quad (3.21)$$

$$\text{in_hand}(\text{red}, \text{panda}). \quad (3.22)$$

$$\text{in_hand}(\text{blue}, \text{mouse}). \quad (3.23)$$

$$\text{does}(\text{red}, \text{move}(\text{cell}(4, 2), \text{cell}(3, 2), \text{panda})). \quad (3.24)$$

$$\text{icon}(\text{panda}, \text{red}). \quad (3.25)$$

$$\text{icon}(\text{mouse}, \text{blue}). \quad (3.26)$$

$$\text{icon}(\text{fox}, \text{red}). \quad (3.27)$$

$$\text{icon}(\text{viper}, \text{red}). \quad (3.28)$$

$$\text{control}(\text{Player}) \leftarrow \text{init}(\text{control}(\text{Player})), \text{not control}(_). \quad (3.29)$$

$$\text{next}(\text{control}(\text{red})) \leftarrow \text{control}(\text{blue}). \quad (3.30)$$

$$\text{next}(\text{control}(\text{blue})) \leftarrow \text{control}(\text{red}). \quad (3.31)$$

$$\begin{aligned} \text{next}(\text{in_hand}(\text{Player}, \text{Card})) \leftarrow \text{in_hand}(\text{Player}, \text{Card}), \\ \text{not control}(\text{Player}). \end{aligned} \quad (3.32)$$

$$\text{next}(\text{in_hand}(\text{Player}, \text{Card})) \leftarrow \text{center_card}(\text{Card}), \text{control}(\text{Player}). \quad (3.33)$$

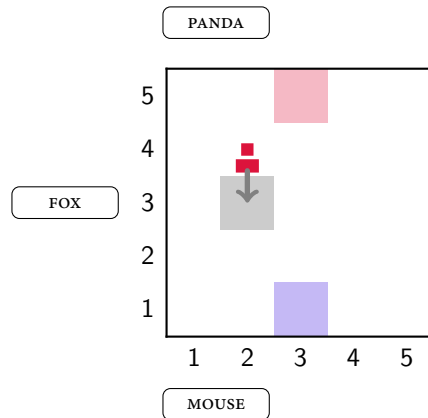


Figure 3.2. Simplified graphical representation of the background facts of B in Example 3.4.2

$$M^{\text{LAS}} = \left\langle \begin{aligned} & \{ \text{init}(\text{control}(\mathbf{v}_{\text{role}})), \text{center_card}(\mathbf{v}_{\text{card}}) \}, \\ & \{ \text{in_play}(\mathbf{v}_{\text{card}}), \text{in_hand}(\mathbf{c}_{\text{role}}, \mathbf{v}_{\text{card}}), \text{icon}(\mathbf{v}_{\text{card}}, \mathbf{v}_{\text{role}}), \text{center_card}(\mathbf{v}_{\text{card}}) \} \end{aligned} \right\rangle$$

⁶ Some types have been omitted for brevity

⁷ ; is used to abbreviate atoms e.g. $a(b;c). \equiv a(b).a(c)$.

Then the following set of positive examples say that on the second turn the red player must be in possession of the FOX card, and that neither MOUSE nor PANDA are the initial centre card.

$$E^+ = \{\langle\{next(in_hand(red, fox))\}, \{center_card(mouse), center_card(panda)\}\rangle\}$$

Additionally, the negative examples below say that in no answer set can both players start, and that the VIPER is never the centre card.

$$E^- = \{\langle\{init(control(red)), init(control(blue))\}, \emptyset\rangle, \langle\{center_card(viper)\}, \emptyset\rangle\}$$

The learnt hypothesis H is:

$$init(control(Player)) \leftarrow center_card(Card), icon(Card, Player). \quad (3.34)$$

$$center_card(Card) \leftarrow in_play(Card), \quad (3.35)$$

$$not\ in_hand(blue, Card),$$

$$not\ in_hand(red, Card).$$

Rule 3.34 says that the first player is given by the player icon on the centre card. Rule 3.35 says that the centre card is the card in play that is in neither player's hand.

It is possible that several hypotheses will cover the examples, ILASP will choose one of the shortest length, defined below.

Definition 3.4.4 (hypothesis length). Given a hypothesis H , the length of the hypothesis $|H|$ is the number of literals in H^D where H^D is obtained from H by replacing all the aggregates by their disjunctive normal form. For example, hypothesis H above (rules 3.34 and 3.35) is length 7.

3.4.2 LEARNING WEAK CONSTRAINTS

So far we have seen that answer sets can represent states of the board (i. e. Figure 3.2) and that multiple moves can generate several answer sets using choice rules (Definition 3.2.1), which can be used to represent small expanded sections of a game tree. When learning strategies we need some notion of preference over these answer sets, and by extension the moves. In commonly used machine learning methods for playing games, such as reinforcement learning, a *state-action value function* can be used to estimate the reward of taking an action in a certain state. Answer Set Programming provides us with weak constraints (Definition 3.2.2) which can assign penalties⁸ to answer sets, acting as a form of state-action value function. Law, Russo and Broda (2015a) introduces LOAS, a system that extends LAS to learn these weak constraints.

In previous ILP systems such as TILDE (Blockeel and De Raedt, 1998), preferential learning was approached as a classification problem (Dastani et al., 2005), with each labelled as 'good' or 'bad'. The drawback is that it offers no relative preference between two examples classified as 'good' (or 'bad').

In order to learn weak constraints we need to extend the notion of mode declaration to be able to generate a language bias that includes weak constraints. Therefore two new mode declarations are added: M_o which specifies what can appear in the body of a weak

⁸ or negative penalties, i. e. rewards

constraint, and M_w which specifies what can appear as a weight. We also give $l_{max} \in \mathbb{N}$ as the maximum level that can appear in H .

Weak constraints are the crux of learning the strategies of a player. They encode why a player perceives one move to be better than the other possible moves. We use the weight of the weak constraint as a penalty and the best move is the answer set that receives the lowest penalty.

Definitions 3.4.5, 3.4.6 and 3.4.7 are from Law, Russo and Broda (2015a), and extend definitions seen in previous sections.

Definition 3.4.5 (language bias). This definition extends Definition 3.4.1. In a LOAS task the language bias is defined by the mode declaration $M = \langle M_h, M_b, M_o, M_w, l_{max} \rangle$. The search space S_M is the set of rules such that $r \in S_M$ satisfies one of the following conditions:

- (i) $r \in S_M^{\text{LAS}}$, where $M^{\text{LAS}} = \langle M_h, M_b \rangle$
- (ii) r is a weak constraint $\Leftarrow b_1, \dots, b_m, \text{not } c_1, \dots, \text{not } c_n.[w@l, t_1, \dots, t_o]$ such that $w \in M_w, l \in [0, l_{max}], t_1, \dots, t_o$ is the set of terms in $body^+(r) \cup body^-(r)$ and each b_i, c_j is compatible with M_o where $i \in [0, m], j \in [0, n]$.

Remark. Because weak constraints do not alter what is contained in an answer set one could always remove the weak constraints from a hypothesis and it would be more optimal and still cover the examples. To solve this we need to introduce the notion of preferred examples to the learning task, the weak constraints are then used to cover this type of example.

Definition 3.4.6 (ordering example). An *ordering example* is a pair $o = \langle e_1, e_2 \rangle$ where e_1, e_2 are partial interpretations. An ASP program Π *bravely* respects o iff $\exists A_1, A_2 \in \mathcal{AS}(\Pi)$ such that A_1 extends e_1, A_2 extends e_2 and $A_1 \succ_{\Pi} A_2$. An ASP program *cautiously* respects o iff $\forall A_1, A_2 \in \mathcal{AS}(\Pi)$ such that A_1 extends e_1, A_2 extends e_2 , it is the case that $A_1 \succ_{\Pi} A_2$.

Definition 3.4.7 (Learning from Ordered Answer Sets task). A Learning from Ordered Answer Sets task is a tuple $T = \langle B, S_M, E^+, E^-, O^b, O^c \rangle$, where

- (i) B is the background knowledge,
- (ii) S_M is the search space defined by the mode declaration,
- (iii) $M = \langle M_h, M_b, M_o, M_w, l_{max} \rangle$,
- (iv) E^+, E^- are the positive and negative examples, respectively,
- (v) O^b, O^c are sets of ordering examples over E^+ called the brave and cautious orderings, respectively.

A hypothesis $H \subseteq S_M$ is in $ILP_{\text{LOAS}}(T)$, the inductive solutions of T , iff:

- (i) $H' \in ILP_{\text{LAS}}(\langle B, S_M^{\text{LAS}}, E^+, E^- \rangle)$, where H' is the subset of H with no weak constraints, $M^{\text{LAS}} = \langle M_h, M_b \rangle$.
- (ii) $\forall o \in O^b$, such that $B \cup H$ bravely respects o
- (iii) $\forall o \in O^c$, such that $B \cup H$ cautiously respects o

3.4.3 CONTEXT DEPENDENT EXAMPLES

When learning strategies and rules of a game it is very likely that one board state will not encapsulate the full rules or nuances of the strategy. For example, preferring to win than capture a student pawn in *Onitama* cannot be inferred from one state in which it is possible to win but not capture, or for that matter one in which it is not possible to win! Therefore, when providing an example we wish to couple it to the state of the board when the move was being made (the context of the move).

To take this into account I present extensions of Definitions 3.4.2, 3.4.6, and 3.4.7, in accordance with Law, Russo and Broda (2016), where the concept of coupling each example with its context was introduced.

Definition 3.4.8 (Context Dependent Partial Interpretation). This definition extends Definition 3.4.2. A Context Dependent Partial Interpretation (CDPI) is a tuple $\langle e, C \rangle$ where e is a partial interpretation and C is an ASP program without weak constraints, called the *context*.

Definition 3.4.9 (Context Dependent Ordering Example). This definition extends Definition 3.4.6. The only difference is that the answer sets include the contexts of their respective CDPI. A Context Dependent Ordering Example (CDOE) o is a tuple of CDPIs $\langle \langle e_1, C_1 \rangle, \langle e_2, C_2 \rangle \rangle$. An ASP program Π *bravely* respects o iff $\exists A_1 \in \mathcal{AS}(\Pi \cup C_1), A_2 \in \mathcal{AS}(\Pi \cup C_2)$ such that A_1 extends e_1 , A_2 extends e_2 and $A_1 \succ_{\Pi} A_2$. An ASP program *cautiously* respects o iff $\forall A_1 \in \mathcal{AS}(\Pi \cup C_1), A_2 \in \mathcal{AS}(\Pi \cup C_2)$ such that A_1 extends e_1 , A_2 extends e_2 , it is the case that $A_1 \succ_{\Pi} A_2$.

Definition 3.4.10 (Context Dependent LOAS task). This definition extends Definition 3.4.7. A Context Dependent Learning from Ordered Answer Sets task is a tuple $T = \langle B, S_M, E^+, E^-, O^b, O^c \rangle$ where the examples are now CDPIs and the orderings are CDOEs. A hypothesis H is an inductive solution of T , $H \in ILP_{LOAS}^{context}(T)$ iff:

- (i) $H \subseteq S_M$,
- (ii) $\forall \langle e, C \rangle \in E^+, \exists A \in \mathcal{AS}(B \cup C \cup H)$ such that A extends e ,
- (iii) $\forall \langle e, C \rangle \in E^-, \neg \exists A \in \mathcal{AS}(B \cup C \cup H)$ such that A extends e ,
- (iv) $\forall o \in O^b$, such that $B \cup H$ bravely respects o ,
- (v) $\forall o \in O^c$, such that $B \cup H$ cautiously respects o .

Example 3.4.3. Let $T = \langle B, S_M, E^+, E^-, O^b, O^c \rangle$ be a $ILP_{LOAS}^{context}$ task where B is the background knowledge from Logic Program A.1, together with the facts $\{in_play(mouse; boar; bear); control(red)\}$,

$$M = \langle \emptyset, \emptyset, \{goal(c_{role}, 100)\}, \{-1, 1\}, 2 \rangle,$$

$$C_1 = \begin{cases} location(pawn(student, blue), cell(1, 2)). \\ location(pawn(master, blue), cell(2, 3)). \\ location(pawn(student, red), cell(2, 4)). \\ location(pawn(master, red), cell(3, 4)). \\ in_hand(red, mouse). \\ in_hand(blue, boar). \end{cases}$$

$$C_2 = \begin{cases} location(pawn(master, blue), cell(1, 3)). \\ location(pawn(student, red), cell(2, 2)). \\ location(pawn(master, red), cell(5, 3)). \\ in_hand(red, bear). \\ in_hand(blue, mouse). \end{cases}$$

$$e_1 = \langle \{does(red, move(cell(2, 4), cell(2, 3), mouse))\}, \emptyset, C_1 \rangle$$

$$e_2 = \langle \{does(red, move(cell(2, 4), cell(1, 4), mouse))\}, \emptyset, C_1 \rangle$$

$$e_3 = \langle \{does(red, move(cell(2, 2), cell(1, 3), bear))\}, \emptyset, C_2 \rangle$$

$$e_4 = \langle \{does(red, move(cell(2, 2), cell(1, 2), bear))\}, \emptyset, C_2 \rangle$$

$$E^+ = \{e_1, e_2, e_3, e_4\},$$

$$E^- = \emptyset,$$

$$O^b = \{\langle e_1, e_2 \rangle, \langle e_3, e_4 \rangle\},$$

$$O^c = \emptyset.$$

The examples are shown in Figure 3.3.

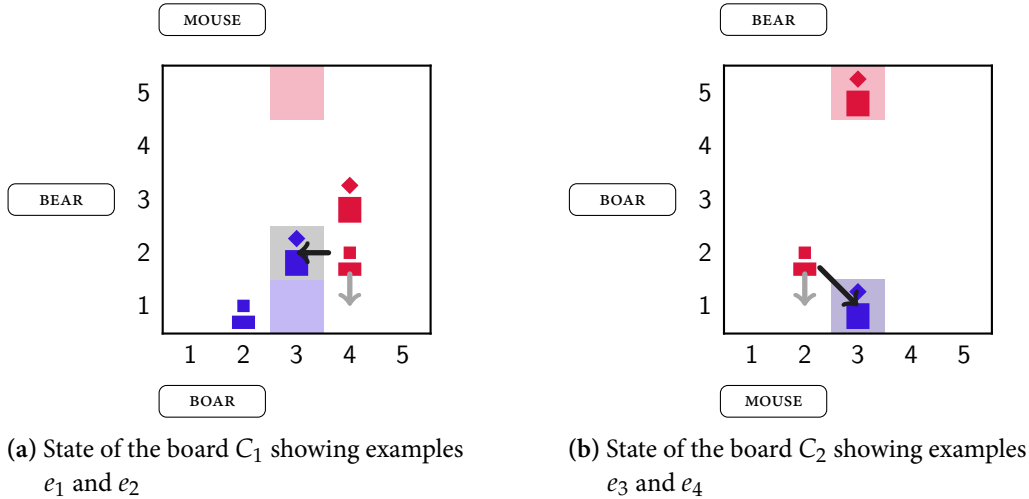


Figure 3.3. Examples e_1, e_2, e_3, e_4 in their respective contexts with the preferred action shown by the highlighted move.

The hypothesis $H = \{\leftarrow \text{not } goal(red, 100).[1@1]\}$ is the inductive hypothesis for T , alternatively you could have $H' = \{\leftarrow \text{not } goal(red, 100).[-1@1]\}$.

3.4.4 LEARNING FROM NOISY EXAMPLES

In complex games a player’s strategy will not always be strictly enforced; the player might deviate or make a mistake. When this happens the moves in the game will not all conform to a consistent strategy that can be learnt. To account for this we use a feature of ILASP called *noise*. Examples or ordering examples can be assigned a *noise weight*, $w \in \mathbb{N}$, which will be accounted for when learning a hypothesis.

Normally, the length of a hypothesis (Definition 3.4.4) is used to choose the optimal hypothesis for a task. When noise is involved the heuristic is $|H| + W$, where W is the sum of all noise weights of the *uncovered* examples.

Example 3.4.4. Let us assume Alan is using the strategy “win if possible, otherwise capture a piece, otherwise move randomly”. Some example moves may look like those in Figure 3.4a–3.4c. However in Figure 3.4d we can see a noisy example from earlier in the game.

Recall $\Pi_{location}$ from Example 3.1.2: B is created by removing the facts from $\Pi_{location}$; $M = \langle \emptyset, \emptyset, \{next(location(pawn(v_{rank}, c_{role}), v_{cell})), goal(c_{role}, 100)\}, \{-1, 1\}, 2 \rangle$; E^+ is the set of examples created with empty inclusions and exclusions, using the four states in Figure 3.4a–3.4d and all possible moves from each board position as the context; and O^b is the set of orderings created by comparing the moves shown in Figure 3.4a–3.4d to all other moves from that state.

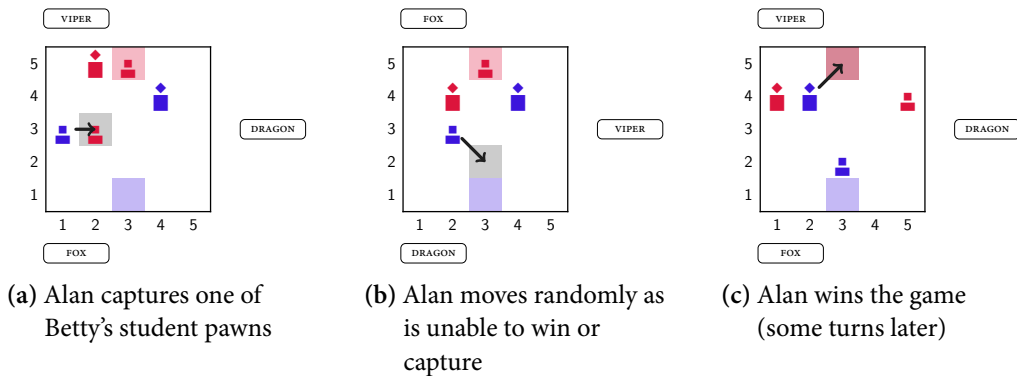


Figure 3.4. Some of Alan’s moves from a particular game against Betty




Figure 3.4. Some of Alan’s moves from a particular game against Betty (cont.)

ILASP is able to learn the following strategy, treating Figure 3.4d as noise:

$$\Leftarrow \text{goal}(\text{blue}, 100).[-1@1] \quad (3.36)$$

$$\Leftarrow \text{next}(\text{location}(\text{pawn}(\text{Rank}, \text{red}), \text{Cell}).[1@1, \text{Rank}, \text{Cell}] \quad (3.37)$$

Weak constraint 3.36 provides a reward for winning, and weak constraint 3.37 says minimise the number of red pawns on the board. This is thus akin to Alan's strategy.

Remark. As there is no example in which there is a choice between winning and capturing they have been learnt at the same level.

3.4.5 CONSTRAINING THE HYPOTHESIS SPACE WITH BIAS CONSTRAINTS

When the mode bias grows with numerous heads and bodies the number of potential rules in the hypothesis space becomes vast. This effect is magnified when there are a lot of common variables and the maximum number of variables is greater than 7-8.⁹ A large proportion of these rules can be considered nonsense semantically, or represent something you know is uninteresting (e. g. you may wish for a series of variables of the same type to be distinct). In order to filter these rules from the hypothesis space *bias constraints* are used. They are a set of meta rules that are applied to the hypothesis space, clause by clause (i. e. not globally), a series of predicates have been defined to aid in filtering. They can be found in Table 3.1.

LOGIC	DESCRIPTION
<code>head(·)</code>	The predicate is in the head of the rule
<code>body(·)</code>	The predicate is in the body of the rule
<code>constraint</code>	The rule is a (hard) constraint
<code>weak_constraint</code>	The rule is a weak constraint
<code>weight(·)</code>	The weight of the weak constraint, variables are made lower case
<code>term(·,·)</code>	The index and lower case version of a term in the weak constraint
<code>naf(·)</code>	The literal is a NAF literal

Table 3.1. Wrapper predicates to indicate structural elements of rules in the hypothesis space

Example 3.4.5. Take the following mode declarations:

$$M = \left\langle \left\{ \text{does}(\mathbf{v}_{\text{role}}, \text{move}((\mathbf{v}_{\text{index}}, \mathbf{v}_{\text{index}}), (\mathbf{v}_{\text{index}}, \mathbf{v}_{\text{index}}))) \right\}, \left\{ \text{cell}(\mathbf{v}_{\text{index}}, \mathbf{v}_{\text{index}}), \text{role}(\mathbf{v}_{\text{role}}) \right\} \right\rangle$$

⁹ Based on empirical evidence of running various ILASP tasks with -s

ILASP produces 1062 rules from this mode bias. This can be made much smaller using the following bias constraints.

$$\leftarrow \text{constraint}. \quad (3.38)$$

$$\leftarrow \text{body}(\text{naf}(\text{cell}(_, _))). \quad (3.39)$$

$$\leftarrow \text{body}(\text{naf}(\text{role}(_))). \quad (3.40)$$

$$\leftarrow \text{head}(\text{does}(_, \text{move}((V, V), (_, _)))). \quad (3.41)$$

$$\leftarrow \text{head}(\text{does}(_, \text{move}((_, _), (V, V)))). \quad (3.42)$$

$$\leftarrow \text{head}(\text{does}(_, \text{move}((V, _), (V, _)))). \quad (3.43)$$

$$\leftarrow \text{head}(\text{does}(_, \text{move}((_, V), (_, V)))). \quad (3.44)$$

$$\leftarrow \text{head}(\text{does}(_, \text{move}((V, V), (_, _)))). \quad (3.45)$$

$$\leftarrow \text{head}(\text{does}(_, \text{move}((V, V), (_, _)))). \quad (3.46)$$

$$\leftarrow \text{head}(\text{does}(_, \text{move}((V, _), (_, V)))). \quad (3.47)$$

$$\leftarrow \text{head}(\text{does}(_, \text{move}((_, V), (V, _)))). \quad (3.48)$$

Equation 3.38 means that there can be no constraints.¹⁰ Equation 3.39–3.40 express that we do not want the types to be NAF literals.¹¹ Equation 3.41–3.48 express that no two indices in the head should be forced to be identical.

With the addition of these bias constraints into the program the search space reduces to only 24 rules.

3.4.6 ILASP META-LEVEL REPRESENTATION

ILASP uses a meta-level representation of the learning task that is solved in Clingo. It does so by manipulating the predicates in the program and ‘tagging’ examples, there are notions of encoding the weights which translate into summing aggregates for each level, and checking for a violating reason can be done with the parity of the single-level penalty. The full encoding is not important for this project, however it is important to note its existence as the grounding of the task is not only dependent on the background knowledge but also the meta-level representation, and there can be repercussions if it is not carefully considered. The complete translation and the definitions of some functions that are used in this report can be found in Law, Russo and Broda (2015a).

META-LEVEL FUNCTIONS

The following two definitions are taken from Law, Russo and Broda (2015a). They are functions that allow the manipulation of predicates in a program, they are used in Chapters 6 and 7 to enable batching and encoding searches.

Definition 3.4.11 (reify). Given a program, Π , a predicate, $\text{pred}(\cdot)$, and a term, term , $\text{reify}(\Pi, \text{pred}, \text{term})$ is the program created by replacing all atoms $\alpha \in \text{atoms}(\Pi)$ with $\text{pred}(a, \text{term})$.

¹⁰ ILASP actually provides a command line flag `-nc` for this feature, I am expressing it here as a bias constraint for the purpose of example

¹¹ ILASP also has a flag for a mode declaration (`positive`), stating that a literal can only appear positively

Definition 3.4.12 (append). Given a program, Π , and an atom α , $\text{append}(\Pi, \alpha) \triangleq \left\{ \text{head}(r) \leftarrow \text{body}^+(r), \text{body}^-(r), \alpha \mid r \in \Pi \right\}$.

3.5 GAME THEORY

Within this report references will be made to occasional game theoretic concepts.

3.5.1 GAME TYPES

ZERO-SUM The reward of a player is the loss of another.

NON-COOPERATIVE Players are competing against each other.

EXTENSIVE-FORM Games are played with sequential turns according to some turn function, for example “trailing player goes first” (see the game *Glen More*¹²) or simply “clockwise around the table”.

Note. All of the games that are considered in this report are two-person, zero-sum, non-cooperative, extensive-form games.

3.5.2 UTILITY FUNCTIONS

A utility function $u^p : \mathcal{S} \mapsto \mathbb{N}$ is a function from states to natural numbers, with a higher utility indicating that player p prefers this state.

Remark. In zero-sum games there is a single utility function u which is used by all players. Any other player in the game receives the negation of the utility as a reward.

3.5.3 MINIMAX THEOREM

Extensive-form games can be visualised using decision trees, with each node representing a game state and each edge being a legal action from one state to another. The minimax algorithm is a method of selecting an action at a node by assuming what would happen if one’s opponent plays optimally, i. e. you choose the move that maximises the utility given the fact your opponent minimises it.

In simple games such as *Nim* and *Tic-Tac-Toe* the tree can be exhaustively searched. Backwards reasoning can be used in order to calculate the exact move to make. Take the example below which describes using the minimax theorem to compute optimal play for the game of *Nim*¹³.

Example 3.5.1. Below is a decision tree for the game of *Nim* a game where a player can split any number into two, unless the number is 1 or 2. A player loses if they cannot make a legal move, i. e. left with only 1s and 2s. Figure 3.5 starts the game at the value of 7, leaf nodes are scored with +1 for a win and −1 for a loss.

¹² <https://boardgamegeek.com/boardgame/66362/glen-more>

¹³ <https://wikipedia.com/nim>

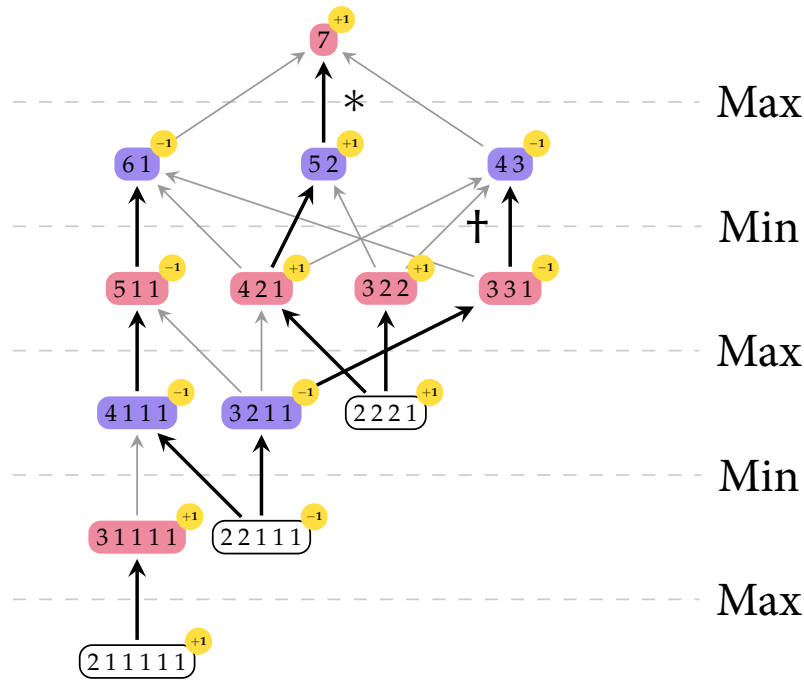


Figure 3.5. Decision tree of *Nim* with backpropagation. Bold arrows denote the backwards propagation of the values in the tree. Coloured nodes represent the player's choice from its children. White nodes are terminal states. Based on the highlighted arrows the red player should choose the middle (*) move as the reward will be +1 which is the greatest of the three options $\{+1, -1, -1\}$. (†) shows a min step, the opponent wishes to minimise the proponent's reward and so chooses the -1 action.

However in more complex games, such as *Chess*, *Go* or *Onitama*, the tree is too large and the search would be too computationally intensive. Therefore, the search is only considered up to a certain depth and the states are compared based on the player's utility function.

α - β PRUNING

α - β pruning is an optimisation technique that can be applied to minimax (Heineman, Pollice and Selkow, 2008). It yields exactly the same result, but examines fewer nodes. During the search the algorithm keeps track of two values α and β which denote the highest min score and the lowest max score, respectively. This way if you find something outside of these bounds you need not examine the branch further. This technique is used in the minimax planner in Section 6.1.2. α - β pruning is one of the optimisation techniques applied to Romstad et al. (*Stockfish*), one of the best *Chess* engines.

Take the following example depicted in Figure 3.6. The algorithm starts by setting $\alpha = -\infty$, $\beta = \infty$, and then visits state *a* and so $\alpha = 6$. Next state *b* is visited, $3 < 6$ so α stays the same, the maximum of these values, 6, is assigned to state *c* and $\beta = 6$. The algorithm then traverses the *f* branch ($\alpha = -\infty$, $\beta = 6$) by visiting *d*, updating $\alpha = 7$, as $\beta \leq \alpha$ the rest of the branch *e* is pruned. Using similar logic for the rest of the tree (*) is also pruned.




Figure 3.6. α - β pruning tree, dashed lines denote the pruned branches.

HORIZON EFFECT

One common issue with the minimax theorem is the *horizon effect*, named after the leaves of the tree which make up the horizon. Informally, it is the lack of knowing what would happen if you evaluated the tree to the next depth, and how that would affect the result. In practice it means that you may make a move that you thought left you in a strong position but five moves later you realise that if you had thought ahead another move it would have turned out to be a bad move. In this report I ignore this effect, though it is something that could be used in the planning phase (Section 6.1.2) in order to create a more formidable and robust Artificial Intelligence (AI). One method of overcoming the effect is using *quiescent search* is a selective search that looks into tactical play. “Programs with a poor or inadequate quiescence search suffer more from the horizon effect” (Marsland, 1986). In many *Chess* engines this is used applied after some minimax lookahead, commonly it is used to look for checks and captures (creating ‘capture trees’). In Chapter 10 we look at defending pieces and exchanges using minimax lookahead. The following chess position studied by Marsland (1986).




Figure 3.7. Using insufficient quiescent search could result in $1\dots b2!\pm$ Blocking the bishop's attack on the Queen. However, this just delays the Queen's capture — this is missed by an 8-ply search ($1\dots b2$ 2 $\text{♙}\times b2$ 3 $\text{♜}\times c3$ 4 $\text{♜}\times d4$ 5 $\text{♜}\times e5$). A better variation would be $1\dots f6$, leading to a draw. FEN: 5r1k/4Qpq1/4p3/1p1p2P1/2p2P2/1p2P3/3P4/BK6 b --

4 | RELATED WORK

Learning strategies in Inductive Logic Programming touches on many areas of computer science such as Logic, Machine Learning, and Game Theory. This diverse background leads to many interesting bodies of work being produced on the topic. Ranging from formal logics to experimental work. I will be looking into areas of Knowledge Representation, use of ASP in planning problems (specifically single-player games), formal logic for describing and verifying properties of strategies, and finally, explainable AIs and ‘black box’ models.

4.1 KNOWLEDGE REPRESENTATION

There are many considerations to make when choosing how to represent an environment in an ASP logic program. Opting for a more terse representation may mean that you lose some useful information. Conversely, a verbose representation can greatly increase the grounding of the program making the learning task infeasible.

There are several formalisations that can be used to represent games. There exist a series of Action Languages (Gelfond and Lifschitz, 1998) that represent transition systems that are very general (e.g. Fangzhen Lin’s suitcase). Some of the languages have translations into ASP (Gebser, Grote and Schaub, 2010; Lee, 2012). However these languages are too general for the simple games. Most games follow similar patterns and can be described using some common features, such as players and actions. These concepts can be expressed in the Action Languages, but there is nothing guiding the representation. This can lead to two games have vastly different representations and therefore varying results when trying to use them in learning tasks.

To avoid this we have chosen to use a language that has been designed for expressing games and is studied and used throughout the General Game Playing (GGP) community to express games. It is simple language that is specific to games, called the Game Description Language (GDL), the specification of which is given in Section 5.2. The basic version covers single player games, later iterations incorporating multi-player games (Love et al., 2006), and incomplete information (Thielscher, 2010). It includes concepts that are specific to games (e.g. `goal` and `role`) which make the logic program more intuitive to read. Translations from GDL to ASP exist (Cerexhe, Sabuncu and Thielscher, 2013) with an initial state, and with timestamped states (defined in terms of the current timestep and actions). A slightly modified version of this translation will be defined in Section 5.2.2 that does away with the time stamps as we are only interested in the current and next state.

4.2 REPRESENTING GAMES IN FORMAL LOGICS

Using the GDL as a basis for further work was completed by Zhang and Thielscher (2015). They formalised the language, allowing them to prove interesting properties of games and extending it to represent strategies of games. The language proposed by Zhang et al. is a modal logic encoding the GDL specification, which intuitively fits the nature of GDL.

The extension of introducing strategies as a subset of the legal moves given some logical “strategy rule” is particularly interesting. Strategy rules can be composed to express priorities using two new logical operators defined in the paper, a prioritised conjunction and a prioritised disjunction.

In their paper, Zhang et al. study various strategies of a simple game, building in complexity. Later, in Chapter 8, I look at learning the presented strategies, comparing ILASP’s representation with their own. I also look at applying some of the techniques I describe in Chapter 7 to learn strategies involving their modal operator, which in their paper they are unable to encode strategies for in ASP.

Kaiser (2012) presents a game learning system that learns first order sentences describing the rules (legal moves and outcomes) of several games, including *Connect4*, *Gomoku*¹ and *Tic-Tac-Toe*, by watching videos of the games being played. The system uses a similar structure to the ILP tasks used here, defining relational structures to represent the board state. Kaiser’s program successfully learns five games, with the longest learning task taking 906 seconds (excluding video processing). Despite the main objective of the study being to improve upon state-of-the-art visual learning, the use of first order logic and Inductive Logic Programming to describe the games is interesting.

4.3 LEARNING ANSWER SET PROGRAMS

ASP has had success in many constraint and planning problems, however little work has been done using ASP to learn rules, or preferences (as we do in this report). For example (Grasso, Leone and Ricca, 2013) outlines a real world example of a travel website that uses an ASP program to suggest places to go. This program could be improved by learning user preferences of locations and times of year based on previous examples.

4.4 MACHINE LEARNING AND GAMES

Early developments in machine learning systems playing board and video games include Temporal-Difference Backgammon programs (Tesauro, 1995), and Atari video games (Mnih et al., 2013). Artificial intelligence has been prominent in the media recently with *AlphaZero* (Silver et al., 2017a) defeating *Stockfish* in Chess, *Elmo* in Shogi, and its predecessor, *AlphaGo Zero*, (Silver et al., 2017b) in Go. As well as current work looking into playing StarCraft II, a 3-D real-time strategy game (Vinyals et al., 2017).

Many of the examples that have seen public success have been based on deep reinforcement learning, and other statistical methods. The downside of this approach is that they cannot explain the reasoning behind the move. In the case of *AlphaGo Lee*, the version of *AlphaGo* that beat master Go player Lee Sedol, some moves played by *AlphaGo Lee* were beyond the comprehension of expert human players (e. g. game 2 move 37²).

4.5 EXPLAINABLE AI

In recent decades we have begun to see an increase in the number of artificially intelligent systems that are equal to or better than human level at complex tasks, e. g. *AlphaZero*, *My-*

¹ <https://boardgamegeek.com/boardgame/11929/go-moku>

² <https://youtu.be/JNrXgpSEEIE>

cin (Shortliffe, 1977). These complex models are essentially black boxes as far as human interpretation is concerned.

Mycin (Shortliffe, 1977), a medical diagnosis system created in the 1970s by Edward Shortliffe. The program was not rolled out to hospitals due to legal and ethical debates, despite achieving higher accuracy than physicians. The main issue was accountability; no one knew how the algorithm came to its conclusions.

A study by DARPA (a research branch of US defense agency) compared different machine learning methods and how ‘explainable’ each method is. Figure 4.1 shows a summary of this. Clearly they believe we have some way to go before we can consider our machine learning methods to be fully explainable, having said this ILASP is not an approximation (without noise) of the examples and interpreting the predicates can be done by people with relative ease.

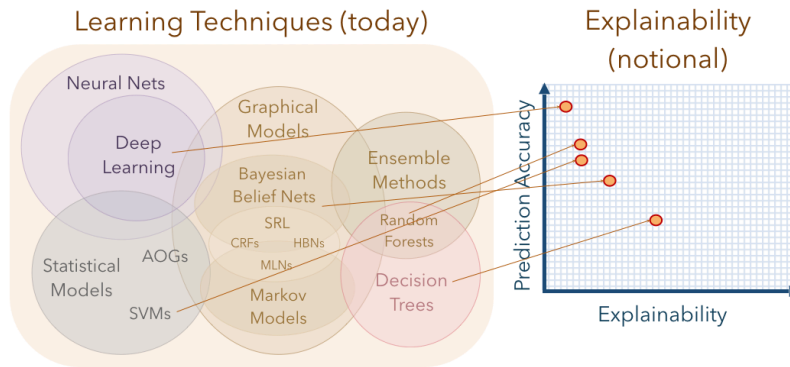


Figure 4.1. Performance vs. Explainability (adapted from Gunning (2016))

Google, in collaboration with Carnegie Mellon University, have looked into explaining how features are perceived by neural networks and visualising neurons to achieve some quite bizarre results (Olah et al., 2018); some examples of the visualisations can be seen in Figure 4.2. Their work also shows key areas of the image that are important in the classification process. For example, a picture of a cottage may be classified as such and



Figure 4.2. Visualisations of neurons from Layer 4c. Figure from Olah et al. (2018)

the reason is that the neuron for ‘house’ from Figure 4.2 was highly activated along with a ‘thatch’ neuron, thereby justifying its decision. Similar work has also been done by others, see Samek, Wiegand and Müller (2017).

Part II

IMPLEMENTATION

5 | GAME MODEL

All learning tasks require a set of background knowledge, which may be empty.¹ Each context dependent example will also contain some background knowledge specific to the example, i. e. state of the board, at that time. As mentioned in Section 4.1 there is a balance to be struck between preserving information and achieving concision. Section 4.1 also introduced a framework for games; this language is formally introduced here alongside a translation into ASP, the format used by ILASP.

Throughout this section *Onitama* is used as a running example. The other games studied throughout this report (Chapter 2) can be represented similarly, and their full representations can be found in Appendix A.

5.1 INTUITION

In order to represent the state of the board the game's components must be encoded as a series of predicates and function symbols. Outlined below is a list of steps I considered when encoding the game.

1. **CREATE TYPES** These are simple predicates to define the types of variables later in the program; these are mostly used to ensure that all variables are safe. They could be more rigorous to ensure the rules will make sense if a bad context is given, for example checking a card name is valid, but this is not necessary.
2. **CREATE COMPONENTS**
 - BOARD** The board is broken up into a 2-dimensional grid of spaces (or cells).
 - PAWNS** Pawns are represented as a tuple of their rank and associated player.
 - CARDS** Cards have a name, a set of moves, and a starting colour associated with them.
3. **GAME STATE** There are two common ways of representing boards in games: (1) describing what is in each space of the board, (2) describing where each pawn is on the board. In my model of *Onitama* I opted for the latter. This is because there are 25 spaces on the board but only a maximum of 10 pawns that can be on the board, therefore reducing the state space. This is because Clingo computes a subset of the grounding based on the facts in the program. All other spaces are assumed to be empty.

I also provide what cards players are holding in this state, and the card in play that is not held is the centre card.
4. **RULES** The predicate `legal(·)` is learnt using examples of legal and illegal moves in different board states. I provide the winning conditions in the background knowledge, though these could also be learnt.

¹ Though in all examples presented here this is not the case

5.2 GAME DESCRIPTION LANGUAGE

The Game Description Language (GDL) is a language used in GGP to specify game states and rules. The language is a variant of Datalog with function constants, negation and recursion (Love et al., 2006), though in practice the language Knowledge Interchange Format (KIF) is used, which can very easily be translated into ASP (Section 5.2.2). The advantages of writing the background and examples of the game in this manner are that (a) the system can easily be used in a GGP competition or using past GGP games, (b) as seen in Section 4.2, other work has been done on languages for games derived from GDL, and finally, (c) it is easy to generate the future possible states of the game from the `next` relation which is useful in Chapter 7. The other style that has been used when writing games in ASP is a time based model, allowing answer sets to be paths down the tree. In the end I decided not to use this approach for numerous reasons. The grounding becomes much larger when you increase the number of time steps into the future the game is simulated. The evaluation function evaluates a single state not paths from the current state.

5.2.1 SPECIFICATION

Here I outline the specification as given in Love et al. (2006). KIF statements are written in prefix notation, e. g. `(cell 1 3)` is a function `cell` applied to arguments 1 and 3. Variables start with question marks, e. g. `(pawn ?rank red)` matches any red pawn. Once again I will use Onitama (Section 2.1) as the running example.

GDL defines several keywords that are used as a common way to represent the states and actions within a game, listed in Table 5.1. The second section of the table lists keywords that are not part of the structure of the game.

KEYWORD	DESCRIPTION
<code>(role ?r)</code>	?r is a player
<code>(init ?p)</code>	?p is a predicate true in the initial state of the game
<code>(true ?p)</code>	?p is a predicate true in the current state of the game
<code>(next ?p)</code>	?p is a predicate true in the next state of the game, i. e. after an action has been made
<code>(does ?r ?a)</code>	?r does action ?a
<code>(legal ?r ?a)</code>	?r can make action ?a
<code>(goal ?r ?v)</code>	?r gets payoff ?v
<code>terminal</code>	This is a terminal state
<code>(distinct ?a ?b)</code>	?a and ?b are different
<code>(not ?p)</code>	The negation of ?p
<code>(<= head body1 body2 ... bodyN)</code>	Construction of a logical rule

Table 5.1. GDL Keywords

PLAYERS: role RELATION

The role relation specifies the players in the game, control is used to show whose turn it is. In Onitama the roles are (role red) and (role blue).

GAME STATE: true RELATION

The true relation describes what is true before an action is made, i. e. the state of the game. Figure 5.1 shows a game state and its corresponding true relation.




Figure 5.1. GDL Current State

INITIAL STATE: init RELATION

init is similar to true, but specific to the initial state of the game. Figure 5.2 shows an example for Onitama.




Figure 5.2. GDL Initial State

GAME STATE UPDATE: `next` RELATION

The `next` relation represents the updates to the game state after an action has been made. This is usually defined as a set of rules, as seen in Figure 5.3c. Figure 5.3 shows an example of this update, where Figure 5.3a shows the current position and the move that is to be made and Figure 5.3b shows the board after the update.

The rules in Figure 5.3c say that if `red` is in control now, `blue` will be in control in the next state and vice versa, i. e. play alternates between the players. The other rule states that a player's piece will be at a new location after being moved, (see `MOVES: does` RELATION on page 49). Several rules have been omitted here for simplicity, such as persistence rules and rules regarding the cards.




Figure 5.3. GDL Next State

```
(=<= (next (control blue)) (true (control red)))
(=<= (next (control red)) (true (control blue)))

(=<= (next (location (pawn ?rank ?role) ?to))
     (does ?role (move ?from ?to ?card))
     (true location (pawn ?rank ?role) ?from)
    )
```

(c) next Relation

Figure 5.3. GDL Next State

LEGAL MOVES: `legal` RELATION

The `legal` moves relation maps players onto moves that can be made in this state. It is quite common to use the atom `noop` to represent that no action has been made, especially in turn based games. Figure 5.4 shows some examples of the `legal` relation in Onitama.

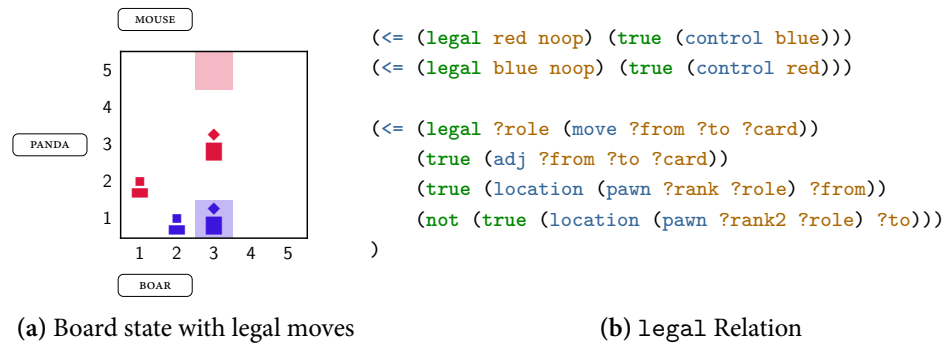


Figure 5.4. GDL Legal and Chosen Moves

MOVES: `does` RELATION

The `does` relation maps players to chosen actions, which will be a subset of the `legal` relation.

Note. Some games allow simultaneous moves by players and some games allow players to make multiple moves, therefore the `does` relation does not necessarily map one player to one action. However, without loss of generality, we can model a player making multiple actions as one combined action.

For example, from Figure 5.4 red could choose (`does red (move (cell 2 1) (cell 1 2) mouse)`), and blue is forced to choose (`does blue noop`).

GOAL STATES: `goal` RELATION

```

(<= (goal ?role 100) (win ?role))
(<= (goal ?role 0) (lose ?role))

(<= (lose red) (win blue))
(<= (lose blue) (win red))

; Winning conditions from Section 2.1.1
(<= (win red) (not (true (location (pawn master blue) ?cell))))
(<= (win blue) (not (true (location (pawn master red) ?cell))))
(<= (win red) (true (location (pawn ?rank red) (cell 5 3))))
(<= (win blue) (true (location (pawn ?rank blue) (cell 1 3))))

```

Figure 5.5. GDL Goal Relation

TERMINAL STATES: `terminal` RELATION

The `terminal` atom is a nullary predicate that indicates that a state has no actions from this point. This could be due to a winning condition being triggered, a round limit being reached, or lack of available options, for example.

In Onitama the game is over if a player wins, the rules do not state what happens if neither player can move, nor is there a condition where players draw if a certain number of moves have passed without capture (cf. chess). In GDL this would be (`<= terminal (win ?role)`).

5.2.2 TRANSLATION INTO ASP

Here I describe a translation into ASP that maintains the instantaneous nature of GDL (cf. time series models Cerexhe, Sabuncu and Thielscher (2013)). Given the set of GDL rules \mathcal{G} , $\llbracket \mathcal{G} \rrbracket = \Pi$, where Π is the equivalent logic program.

$$\llbracket \mathcal{G} \rrbracket = \bigcup_{r \in \mathcal{G}} \llbracket r \rrbracket \cup \text{generator} \quad (5.1)$$

$$\llbracket (\leq h \ b1 \ \dots \ bn) \rrbracket = \llbracket h \rrbracket \leftarrow \llbracket b1 \rrbracket \ \dots \ \llbracket bn \rrbracket \quad (5.2)$$

$$\llbracket (\text{distinct } t1 \ t2) \rrbracket = \llbracket t1 \rrbracket \neq \llbracket t2 \rrbracket \quad (5.3)$$

$$\llbracket (\text{not } a) \rrbracket = \text{not } \llbracket a \rrbracket \quad (5.4)$$

$$\llbracket (\text{true } (p \ t1 \ \dots \ tn)) \rrbracket = p(\llbracket t1 \rrbracket, \dots, \llbracket tn \rrbracket) \quad (5.5)$$

$$\llbracket (p \ t1 \ \dots \ tn) \rrbracket = p(\llbracket t1 \rrbracket, \dots, \llbracket tn \rrbracket) \quad (5.6)$$

$$\llbracket ?v \rrbracket = V \quad (5.7)$$

$$\llbracket a \rrbracket = a \quad (5.8)$$

where

$$\text{generator} = \left\{ \begin{array}{l} 0\{\text{does}(\text{Role}, A)\}1 \leftarrow \text{legal}(\text{Role}, A) \\ \leftarrow \text{role}(\text{Role}), \text{not } \text{does}(\text{Role}, _), \text{not } \text{terminal} \end{array} \right\} \quad (5.9)$$

The *generator* set encapsulates the choice of actions and ensures that an action is made at each non-terminal state.

Example 5.2.1. Figure 5.6 shows the translation of Tic-Tac-Toe from GDL into ASP².

² A few rules have been omitted for brevity


```

1  (role x)
2  (role o)
3
4  (init (cell 1 1 b))
5  (init (cell 1 2 b))
6  (init (cell 1 3 b))
7  (init (cell 2 1 b))
8  (init (cell 2 2 b))
9  (init (cell 2 3 b))
10 (init (cell 3 1 b))
11 (init (cell 3 2 b))
12 (init (cell 3 3 b))
13 (init (control x))
14
15 (<= (next (cell ?m ?n ?role))
16     (does ?role (mark ?m ?n))
17     (true (cell ?m ?n b)))
18 (<= (next (cell ?m ?n ?role))
19     (true (cell ?m ?n ?role))
20     (distinct ?role b))
21 (<= (next (cell ?m ?n b))
22     (does ?role (mark ?j ?k))
23     (true (cell ?m ?n b))
24     (distinct ?m ?j))
25 (<= (next (cell ?m ?n b))
26     (does ?role (mark ?j ?k))
27     (true (cell ?m ?n b))
28     (distinct ?n ?k))
29 (<= (next (control x)) (true (control o)))
30 (<= (next (control o)) (true (control x)))
31
32 (<= open (true (cell ?m ?n b)))
33
34 (<= (legal ?role (mark ?x ?y))
35     (true (cell ?x ?y b))
36     (true (control ?role)))
37 (<= (legal x noop) (true (control o)))
38 (<= (legal o noop) (true (control x)))
39
40 (<= (goal x 100) (line x))
41 (<= (goal x 0) (line o))
42 (<= (goal o 100) (line o))
43 (<= (goal o 0) (line x))
44 (<= (goal ?role 50)
45     (role ?role) (not open)
46     (not (line x)) (not (line o)))
47
48 (<= terminal (line ?role))
49 (<= terminal (not open))

```

(a) Tic-Tac-Toe in KIF adapted from Love et al. (2006)

```

1  role(x).
2  role(o).
3
4  init(cell(1, 1, b)).
5  init(cell(1, 2, b)).
6  init(cell(1, 3, b)).
7  init(cell(2, 1, b)).
8  init(cell(2, 2, b)).
9  init(cell(2, 3, b)).
10 init(cell(3, 1, b)).
11 init(cell(3, 2, b)).
12 init(cell(3, 3, b)).
13 init(control(x)).
14
15 next(cell(M, N, Role) :-
16     does(Role, mark(M, N)),
17     true(cell(M, N, b)).
18 next(cell(M, N, Role) :-
19     true(cell(M, N, Role)), Role != b.
20 next(cell(M, N, b) :-
21     does(Role, cell(J, K)),
22     true(cell(M, N, b)), M != J.
23 next(cell(M, N, b) :-
24     does(Role, cell(J, K)),
25     true(cell(M, N, b)), N != K.
26 next(control(x)) :- true(control(o)).
27 next(control(o)) :- true(control(x)).
28
29 open :- true(cell(M, N, b)).
30
31 0 { does(Role, A) } 1 :- legal(Role, A).
32 :- role(Role), not does(Role, _),
33     not terminal.
34
35 legal(Role, mark(X, Y)) :-
36     true(cell(X, Y, b)), true(control(Role)).
37 legal(x, noop) :- true(control(o)).
38 legal(o, noop) :- true(control(x)).
39
40 goal(Role, 100) :- line(Role).
41 goal(x, 0) :- line(o).
42 goal(o, 0) :- line(x).
43
44 goal(Role, 50) :-
45     role(Role), not open,
46     not line(x), not line(o).
47
48 terminal :- line(Role).
49 terminal :- not open.

```

(b) Translation of Tic-Tac-Toe into ASP

Figure 5.6. Translation of Tic-Tac-Toe program

5.3 SIMPLIFICATIONS

The representation described in the previous section can become quite large, with many untied³ variables and a large grounding because of some of the predicates used.

Untied variables are much easier to simplify, clingo and ILASP both have support for anonymous variables ($_$) which are projected away by the grounder (gringo). This means that, for example, if the following card predicate is used to check ownership of the card the translation would not be needed, allowing us to replace them with anonymous variables.

$$\text{card}(\text{Name}, \text{DR}, \text{DC}) \Rightarrow \text{card}(\text{Name}, _, _)$$

By the grounder this is then treated as $\text{card}(\text{Name})$ and the arity 3 predicate is now only 1.

Arithmetic expressions can be calculated at ground time by gringo, but when generating mode declarations For example take the following rule for adjacency:

$$\begin{aligned} \text{adj}((\text{Row1}, \text{Col1}), (\text{Row2}, \text{Col2})) \leftarrow & \text{cell}((\text{Row1}, \text{Col1})), & (5.10) \\ & \text{cell}((\text{Row2}, \text{Col2})), \\ & \text{Row2} == \text{Row1} + 1, \\ & \text{Col2} == \text{Col1} + 1. \end{aligned}$$

This means that two squares of the form \square are adjacent to each other. However, because the variables cannot be free we must bind the variables by ensuring they are valid cells. On a 5-by-5 board this generates 25 unique cells.

This can be simplified by substituting arithmetic expressions into the head of the equation (where applicable), however the types are still needed.

$$\begin{aligned} \text{adj}((\text{Row1}, \text{Col1}), (\text{Row1} + 1, \text{Col1} + 1)) \leftarrow & \text{cell}((\text{Row1}, \text{Col1})), & (5.11) \\ & \text{cell}((\text{Row1} + 1, \text{Col1} + 1)) \end{aligned}$$

Clingo is optimised so that these are equivalent when grounding. However when generating the mode declarations in ILASP using the second format is better providing you restrict the maximum length of the body and the number of variables needed. Table 5.2 shows the number of rules generated for the two formats and different parameters passed to ILASP, the (*) indicates that the correct rule does not appear in this set of rules.

Max Variables	Max Weak Constraint Length	Format for rule 5.10	Format for rule 5.11
2	2	1*	6
4	4	591	2238

Table 5.2. Rules generated for normal/simplified rules with different parameters

³ A word which here means that a variable occurs in the body of a rule only once

6 | DIGITISED GAME & PLANNER

In order to collect examples to learn from I created some digital versions of the games that can record different sorts of moves depending on what I was interested in. I created a multi-purpose digital version of each game I was studying: (1) they *collect examples* from recorded games, (2) they utilised learnt strategies to play games in conjunction with a *minimax planner*, (3) and finally they could be used to help the AI learn from it's mistakes with *assistive movement*. The digital games have two running modes and can be configured for human vs. human games, human vs. AI play, and AI vs. AI self-play. When playing a game involving an AI it is possible to run it in tournament-mode or training-mode. Tournament-mode is just like a normal game where the AIs will play the moves they choose based on their current strategy. In training-mode the AI will ask the human whether or not their chosen move is a good one and ask for a suggestion if the human disagrees.

I chose to write the planning library in Haskell as the games can be easily described succinctly using the language's type system. Each game needs a way of getting legal moves, validating if a move is legal, and an update function for moves; I utilised Haskell's Type-Classes in order to make the AI player agnostic to the game.

6.1 PROGRAM FLOW

The main flow of the program is presented in Figure 6.1, the red (●), green (●) and yellow (●) nodes refer to Sections 6.1.1, 6.1.2 and 6.1.3, respectively.




Figure 6.1. User flow through the digital game. Blue states represent intermediate nodes between the sections discussed in this chapter. The red state is the step in which examples are collected, this is at the end of the program after the game has been played. Green states represent the steps made by the minimax planner. Finally, the yellow state is the assistive move step where the user can provide a counter example to the AIs move if needed.

- 0 Select game
- 1 Choose player configurations (human or AI)
- 2 Start Game
- 3 Human plays move (if applicable)
- 4 AI generates the game tree to a given depth
- 5 Clingo rules are generated to represent the leaves of the tree, and example is given below (Example 6.1.1)
- 6 Clingo is run in batches of 20,000 to evaluate the scores of the leaves
- 7 The game tree is re-mapped so that leaves now contain scores instead of board states
- 8 The minimax algorithm is run over the tree in order to select the best move
- 9 If in tournament-mode the move is played and it is the next players turn
- 10 If in training-mode the move is presented to the user, if accepted the move is played otherwise the player's chosen move is played and the ordering is saved as a counter example
- 11 If the game is in a terminal state the game is over and (optionally) the winners orderings are saved as described in Section 6.1.1

6.1.1 EXAMPLE COLLECTION

Digital versions of the games can be played by two players and the winner's moves are chosen as examples. We make the assumption that a player has played as well as they could according to their chosen strategy, and that when facing a competent opponent some strategic choices will be made by the winner in order to beat their opposition. The moves they made are then turned into ordering examples in the following way:

1. Create the positive example for their chosen move, m_{pref}^i where i is the turn number, the context is the board state and the inclusions is the single fact $does(\cdot, \cdot)$ detailing the move.
2. From the given board state generate all possible alternative moves (i. e. ignoring the move they made).
3. Create the positive examples for alternative moves, $m_{alt_j}^i$ where i is the turn number and j is an index for the move.
4. Create the brave orderings $\langle n, m_{pref}^i, m_{alt_j}^i \rangle \forall i \forall j$ for some noise weight n .

6.1.2 MINIMAX PLANNER

When calculating which move to make the AI uses Clingo as a scoring mechanism and minimax (Section 3.5.3) as an adversarial search method for looking ahead in the tree of possibilities to better estimate the optimal move.

Example 6.1.1. Given we start in the state pictured in Figure 6.2, and we generate a game tree to depth 2 we could end up in any of the board states pictured in Figure 6.3. In

order to score these board states we need Clingo to evaluate the penalty given by the weak constraints that encode the current strategy. Whilst it is possible to do using a single state of the board and getting the score of the single answer set, this method becomes *very* slow when computing the scores of leaves from a depth 4 or more game tree. The bottleneck of this computation is the I/O that must be performed. Therefore, to combat this, the board states can be batched together in the following manner: let B be the background knowledge from Logic Program A.1 excluding the predicates $\text{next}(\cdot)$ and the action generator, Π_1, Π_2 be two programs representing two board states from Figure 6.3

$$\Pi_1 = \begin{cases} \text{location}(\text{pawn}(\text{master}, \text{red}), \text{cell}(5, 3)) \\ \text{location}(\text{pawn}(\text{student}, \text{red}), \text{cell}(4, 3)) \\ \text{location}(\text{pawn}(\text{student}, \text{red}), \text{cell}(4, 5)) \\ \text{location}(\text{pawn}(\text{master}, \text{blue}), \text{cell}(1, 3)) \\ \text{location}(\text{pawn}(\text{student}, \text{blue}), \text{cell}(3, 2)) \\ \text{location}(\text{pawn}(\text{student}, \text{blue}), \text{cell}(1, 2)) \\ \text{control}(\text{red}) \\ \text{in_hand}(\text{goose}, \text{red}) \\ \text{in_hand}(\text{monkey}, \text{blue}) \end{cases} \quad \Pi_2 = \begin{cases} \text{location}(\text{pawn}(\text{student}, \text{red}), \text{cell}(5, 4)) \\ \text{location}(\text{pawn}(\text{student}, \text{red}), \text{cell}(4, 5)) \\ \text{location}(\text{pawn}(\text{master}, \text{blue}), \text{cell}(1, 3)) \\ \text{location}(\text{pawn}(\text{student}, \text{blue}), \text{cell}(4, 2)) \\ \text{location}(\text{pawn}(\text{student}, \text{blue}), \text{cell}(3, 3)) \\ \text{location}(\text{pawn}(\text{student}, \text{blue}), \text{cell}(1, 2)) \\ \text{control}(\text{red}) \\ \text{in_hand}(\text{goose}, \text{red}) \\ \text{in_hand}(\text{monkey}, \text{blue}) \end{cases}$$

Recall the meta-function `append` from Section 3.4.6. We use it here to create the program $\Pi = B \cup \text{append}(\Pi_1, \text{option}(1)) \cup \text{append}(\Pi_2, \text{option}(2))$ which encodes each board state as a different option that Clingo must score.

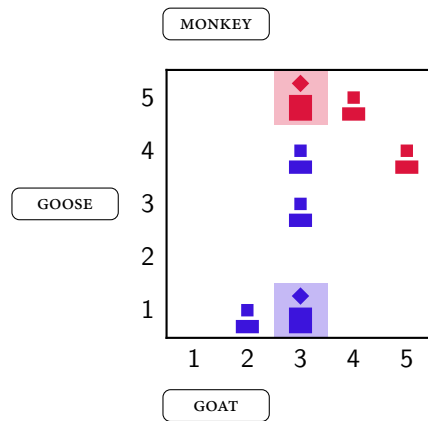




Figure 6.2. Current board state midway through a game

6.1.3 ASSISTIVE MOVEMENT

The assistive movement feature was implemented to replicate the style of learning that a parent might take on with their child, or an experienced player with a beginner. When a less experienced player makes a blunder the more experienced player may offer an alternative move and allow them to try to figure out why that move is better. This is what is offered in the training-mode, which when in use, diverts the flow of the program through the yellow (●) node in Figure 6.1. After the AI player has selected a move the user is




(a) Figure 6.2 after the moves
 $\text{red}(5,4) \times (4,3)$ MONKEY.
 $\text{blue}(3,3) (3,2)$ GOAT.




(b) Figure 6.2 after the moves
 $\text{red}(4,5) (3,4)$ MONKEY.
 $\text{blue}(1,2) (2,3)$ GOAT.

Figure 6.3. Possible leaf nodes of the game tree pruned to depth 2



(c) Figure 6.2 after the moves
 $\text{red}(5,3) (4,2)$ MONKEY.
 $\text{blue}(4,3) \times (4,2)$ GOAT.



(d) Figure 6.2 after the moves
 $\text{red}(5,4) \times (4,3)$ MONKEY.
 $\text{blue}(1,3) (2,4)$ GOAT.

Figure 6.3. Possible leaf nodes of the game tree pruned to depth 2 (cont.)

offered the opportunity to select a correction, Figure 6.4 shows an example of this, the brave ordering created in this example would be:

$$\begin{aligned}
 e_1 &= \langle \{ \text{does}(\text{blue}, \text{move}(\text{cell}(4,2), \text{cell}(3,1), \text{dog})) \}, \emptyset, C \rangle \\
 e_2 &= \langle \{ \text{does}(\text{blue}, \text{move}(\text{cell}(5,3), \text{cell}(5,2), \text{dog})) \}, \emptyset, C \rangle \\
 o &= \langle e_1, e_2 \rangle
 \end{aligned}$$

where C is the context of the board:

$$C = \begin{cases} \text{location}(\text{pawn}(\text{student}, \text{red}), \text{cell}(2,2)) & \text{location}(\text{pawn}(\text{student}, \text{blue}), \text{cell}(5,1)) \\ \text{location}(\text{pawn}(\text{student}, \text{red}), \text{cell}(2,3)) & \text{location}(\text{pawn}(\text{student}, \text{blue}), \text{cell}(4,2)) \\ \text{location}(\text{pawn}(\text{student}, \text{red}), \text{cell}(2,5)) & \text{location}(\text{pawn}(\text{student}, \text{blue}), \text{cell}(4,3)) \\ \text{location}(\text{pawn}(\text{student}, \text{red}), \text{cell}(1,5)) & \text{location}(\text{pawn}(\text{student}, \text{blue}), \text{cell}(4,4)) \\ \text{location}(\text{pawn}(\text{master}, \text{red}), \text{cell}(1,3)) & \text{location}(\text{pawn}(\text{master}, \text{blue}), \text{cell}(5,3)) \end{cases}$$




Figure 6.4. Example of a user correcting a move (which directly puts an undefended student in danger) to a more defensive alternative

6.2 EXTENSIBILITY

The digital games and minimax planner have been designed to allow new games to be added easily. The games are implemented as a package within the `Planner.Games.Data` package and a `Main` file responsible for running the game loop and asking for user moves. New AI tactics can be added under the `Planner.Games.AI` package. There are a few additional packages such as `Planner.Games.Data.Tree`, `Planner.Logic`, and `Planner.Minimax` responsible for game trees, translation into (gringo) logic programs, and minimax, respectively.

6.2.1 ADDING NEW GAMES

In addition to *Onitama* I have implemented *Five Field Kono*, and *Cross-Dot*. This can be done by simply describing the games using Haskell's type system and writing a data type that conforms to the `GameState s` type class that is parameterised on a game state, the class uses type families¹ in order to couple the game state type with the types of moves, players, and pieces. The type class is presented below, some methods are necessary for gameplay (e.g. `getBoardArray` and `makeMove`) and some are solely for example collection

¹ https://wiki.haskell.org/GHC/Type_families

(e.g. `lastMove` and `relevantMoves`). This type class is the used by the game trees, game loop and the AI players to be able to interface with the different game states.

```

1 class
2   (Eq (Player s), Eq (Piece s), Eq (Move s), Show s, Show (Move s), Translate s,
3     ↪ Translate (Move s,s), Translate (Player s, Move s), Translate (Player s),
4     ↪ Translate (Piece s))
5     => GameState s where
6
7   type Move s
8   type Player s
9   type Piece s
10  getBoardArray    :: s -> Board (Piece s)
11  getPossibleMoves :: s -> [Move s]
12  getResult        :: s -> Maybe (Result (Player s))
13  makeMove         :: Move s -> s -> s
14  whoseTurn        :: s -> (Player s)
15  legalmove        :: Move s -> s -> Bool
16  lastMove         :: s -> Maybe (Move s)
17  relevantMoves    :: s -> [(Move s, Move s)]

```

Haskell's expressive types mean that lots of parameterised data types could be created that many games can utilise. Below are the types that are used throughout all three games implemented so far.

```

1 type Location = (Int, Int)
2 data Position p = Empty | Filled p deriving (Eq, Ord, Show)
3 type Board p = Array Location (Position p)
4 data Result p = Draw | Win p
5   deriving (Eq, Show)

```

Most abstract strategy games have the notion of a board or grid on which the game is played (e.g. *Chess*, *Go*, *Onitama*, *Tak*² etc.). Here the board is represented by an array indexed by a `Location` tuple. Each location might contain a piece, which can be anything at all. In *Five Field Kono* and *Cross-Dot* a piece can be simply denoted by the player. More complex games such as *Chess* and *Onitama* must use a custom piece data type to represent the difference between pieces (e.g. Knights and Rooks or Masters and Students). A game such as *Tak* can even have stacks of different types of pieces, yet here we can still easily represent this, perhaps using `Data.Stack`³. The fact that a space *might* contain a piece suggests that the `Maybe`⁴ data type could be used. Instead we defined the data type `Position` with two constructors `Empty` and `Filled p`, which is in fact isomorphic to the `Maybe` data type but affording us more clarity when reading the type signatures.

Whilst this framework does not cover all games, it covers most games based around a grid that pieces move on — even the cards in *Onitama* can still be represented and used. This could be abstracted further by allowing the game to represent the indexing into the board array or the result, but the extra generality covers a set of games that are not yet being looked at by the learning tasks presented in this report.

2 <https://boardgamegeek.com/boardgame/197405/tak>

3 <http://hackage.haskell.org/package/Stack-0.3.2/docs/Data-Stack.html>

4 <http://hackage.haskell.org/package/base-4.11.1.0/docs/Data-Maybe.html>

7 | LEARNING PREFERENCES FROM GAME TREES

7.1 MOTIVATION

In many games a player can gain a strategic advantage by considering possible moves in the future, consider chess where Grand Masters are often calculating games 5–10 steps into the future in their head at any time. In some games it is very easy to think ahead due to, say, only a few moving pieces on the board. Whereas others can quickly branch and it is difficult to consider all possibilities, so you may choose just a few likely options to consider. Take, for example, this situation in Onitama:

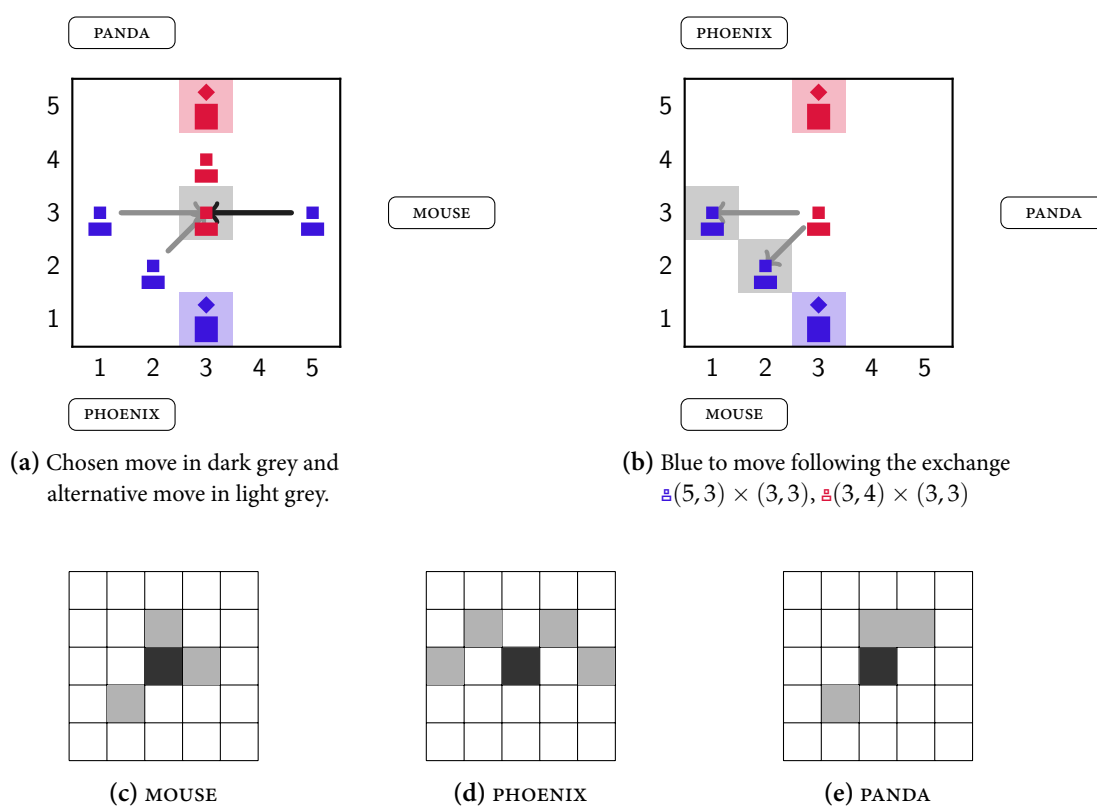


Figure 7.1. Possibility of an exchange

Example 7.1.1. Figure 7.1a depicts blue’s turn with the PHOENIX card. If blue’s strategy was to “capture a piece if possible, otherwise move randomly” then blue would pick one of the three moves shown. Suppose blue chooses $\blacktriangleleft(5,3) \times (3,3)$, one possible series of (unfortunate) events that could unfold for blue is shown in Figure 7.1b. From this position blue cannot capture red, and has few pawns, thus few options. Furthermore, no matter what move blue makes red is able to take a piece weakening blue’s position, this tactic is known as a *fork*. However, if we consider the move again but looking a few steps into the future we can see that even though that pawn is in danger, it is possible to pick a move that avoids this scenario and puts the blue player into a good position.

Note. When reasoning like this using minimax you can suffer from the *horizon effect* (Section 3.5.3), imagine that this was in fact the state we got to after look several moves into the future and was the best option. If we had looked *just one more level* further we would have seen a different result. Thus, looking further down the tree is not guaranteed to always give better results.

Looking into this effect is not in the scope of this project, but is an interesting area for further work (see Section 11.2.1).

7.2 INDUCTIVE LEARNING PROGRAMS WITH DEEP ORDERINGS

In this section we define a new ordering type that extends Context Dependent Ordering Example (CDOE) to take this tree structure of games into account when learning. In all the learning tasks we have seen up to this point all the examples have been independent from each other. In order to reason in the future it is important to consider possible outcomes at future states, and so we may wish to look for certain preferences within groups of examples that represent states in the future. In order to define these groups we must create a relation between a subset of the examples. The new orderings use minimax trees to work out the best strategy for the player. We can then say that a player prefers a move to another iff it is the best action to take given the minimax tree from this state. These trees will define the relationship between certain examples.

Note. We make the assumption that both players are rational and using the same utility.

A *deep* ordering takes two states and looks down the game tree to a certain depth. In order to illustrate this consider the following trees in Figures 7.2, 7.3 and 7.4. The root node is the game board in Figure 7.1. Dashed lines represent a choice, where only one node is needed as part of the requirement, solid lines mean that all nodes are considered.

Suppose Alan is playing a game and chooses the *left* action, assuming Alan is rational and playing optimally we can deduce that the action *left* was a better choice than action *right*. Figure 7.2 illustrates this choice, with the numbers in the nodes representing the utility Alan gains for choosing that action. Here we see that the utility of *left* is greater than that of *right*.

For Alan to have known what the utility of *left* was he would have had to considered possible moves his opponent, Betty, could have made. Alan can then reason about which action to take assuming Betty is trying to make him lose and she is also rational and playing optimally. For the *left* action to be the optimal choice it must be the case that no matter what Betty chooses to do it cannot be worse for him than one of the options she would have, had he chosen the *right* action. Figure 7.4 illustrates what this might look like. We see that the *left* action has expanded into two actions, *l-left* and *l-right*, from which Betty can choose between. Further, we can observe that *all* options reachable from *left* are better (i. e. have a higher utility) than the single option, *r-only*, reachable from *right*.




Figure 7.2. Play is considered to depth 0




Figure 7.4. Play is considered to depth 2

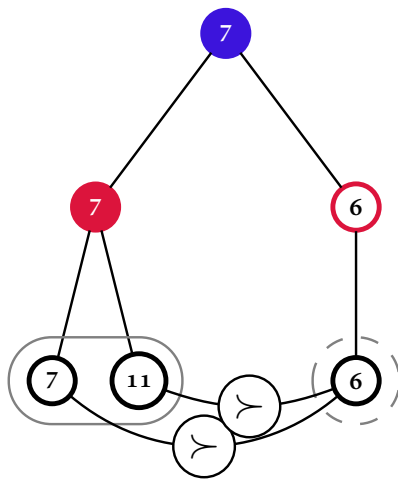


Figure 7.3. Play is considered to depth 1

However, notice that this can now be broken into two sub games, one for the game after Alan selects *left* and one after Alan selects *right*. Betty, in both cases, can do as Alan did and consider possible moves that Alan can respond with. This leads us onto Figure 7.4.

In Figure 7.4 we once again unfold the tree down one more level. This now shows the choices Alan has. Given Alan chose action *left*, and assuming Betty selected the *l-left*, Alan can choose between *l-l-left*, *l-l-middle*, and *l-l-right*. As Alan is rational we know that Alan will choose the option that is better than any of the options he

would be faced with given the action Betty could have chosen if he had originally selected *right*. In other words, one of *l-l-left*, *l-l-middle*, and *l-l-right*, must be better than both *r-o-left* and *r-o-right* (the actions highlighted with a solid line in Figure 7.4). However, this must also hold if Betty had chosen action *l-right*! Therefore, *l-r-only* must also be better than both *r-o-left* and *r-o-right*.

We can clearly from the diagram that this holds, and so the action *left* that Alan originally selected was the optimal move (considering up to a depth of two). It is also apparent how this recursive structure can be generalised to an arbitrary depth, and arbitrary branching factors.

Formally, we represent the game trees, \mathcal{T} , as a tuple (e, \triangleleft) where $e \in \mathcal{S}$ is the root of the tree, a state of the game board, and $\triangleleft : \mathcal{S} \times \mathcal{S}$ the child relation ($c \triangleleft p$ reads “ c is a child of p ”) which is a partial ordering over \mathcal{S} , where \mathcal{S} is the set of all states of the game board. Let $u : \mathcal{S} \mapsto \mathbb{N}$ be the utility function of the game.

Definition 7.2.1 (Minimax Phases). The minimax algorithm is broken up into a *max_phase* and a *min_phase*, where *max_phase*, *min_phase* are functions with signature $\mathcal{T} \mapsto \mathbb{N}$.

$$\text{max_phase}((e, \triangleleft), d) \triangleq \begin{cases} u(e) & \text{if } d = 0 \\ \max \{ \text{min_phase}(e', d-1) \mid e' \triangleleft e \} & \text{otherwise} \end{cases}$$

$$\text{min_phase}((e, \triangleleft), d) \triangleq \begin{cases} u(e) & \text{if } d = 0 \\ \min \{ \text{max_phase}(e', d-1) \mid e' \triangleleft e \} & \text{otherwise} \end{cases}$$

A variation of the minimax algorithm that selects between just two children of a node is defined as:

$$\text{minimax}: \mathcal{T} \times \mathcal{T} \times \mathbb{N} \mapsto \mathcal{T}$$

$$\text{minimax}((e_1, \triangleleft), (e_2, \triangleleft), d) \triangleq \arg \max \left\{ \begin{array}{l} \text{min_phase}((e_1, \triangleleft), d), \\ \text{min_phase}((e_2, \triangleleft), d) \end{array} \right\},$$

where $d \geq 0$.

Definition 7.2.2 (Explanation Condition). The following *explanation condition* is a predicate that takes a utility function, two trees and a depth and will be true if there are paths in the left tree that explain why the left branch is preferred to the right branch. This is done by alternating quantifiers at each depth, following the intuition given previously.

$$\text{explanation}: (\mathcal{S} \mapsto \mathbb{N}) \times \mathcal{T} \times \mathcal{T} \times \mathbb{N} \mapsto \text{Bool}$$

$$\text{explanation}(u, (e_1, \triangleleft), (e_2, \triangleleft), d)$$

$$\triangleq (\forall e_1^1 \triangleleft e_1)(\exists e_1^2 \triangleleft e_1^1)(\forall e_1^3 \triangleleft e_1^2) \cdots (\forall e_1^{d-1} \triangleleft e_1^{d-2})(\exists e_1^d \triangleleft e_1^{d-1}) \\ (\exists e_2^1 \triangleleft e_2)(\forall e_2^2 \triangleleft e_2^1)(\exists e_2^3 \triangleleft e_2^2) \cdots (\exists e_2^{d-1} \triangleleft e_2^{d-2})(\forall e_2^d \triangleleft e_2^{d-1}) \\ u(e_1^d) > u(e_2^d)$$

Proposition 1. Given two trees (e_1, \triangleleft) , (e_2, \triangleleft) , a depth, d , and a utility function, u ,

$$\text{min_phase}((e_1, \triangleleft), d) > \text{min_phase}((e_2, \triangleleft), d) \iff \text{explanation}(u, (e_1, \triangleleft), (e_2, \triangleleft), d)$$

Proof.

Case 1 (Base Case).

Case 1.1 ($d = 0$). To show:

$$\text{min_phase}((e_1, \triangleleft), 0) > \text{min_phase}((e_2, \triangleleft), 0) \iff \text{explanation}(u, (e_1, \triangleleft), (e_2, \triangleleft), 0)$$

From the definition of *min_phase* we get that $u(e_1) > u(e_2)$, which is the definition of *explanation* $(u, (e_1, \triangleleft), (e_2, \triangleleft), 0)$, as when $d = 0$ there are no quantifiers.

Case 1.2 ($d = 1$). To show:

$$\text{min_phase}((e_1, \triangleleft), 1) > \text{min_phase}((e_2, \triangleleft), 1)$$

$$\iff (\forall e_1^1 \triangleleft e_1)(\exists e_2^1 \triangleleft e_2)u(e_1^1) > u(e_2^1)$$

$$\min_phase((e_1, \triangleleft), 1) > \min_phase((e_2, \triangleleft), 1) \quad (7.1)$$

$$\iff \min_{e_1^1 \triangleleft e_1} \left(\max_phase(e_1^1, 0) \right) > \min_{e_2^1 \triangleleft e_2} \left(\max_phase(e_2^1, 0) \right)$$

$$\iff (\forall e_1^1 \triangleleft e_1)(\exists e_2^1 \triangleleft e_2) \max_phase(e_1^1, 0) > \max_phase(e_2^1, 0)$$

This follows from being able to select any child of e_1 as it is true for the minimum and the fact that one can select the child of e_2 that yields the minimum

$$\iff (\forall e_1^1 \triangleleft e_1)(\exists e_2^1 \triangleleft e_2) u(e_1^1) > u(e_2^1) \quad (\text{from definition } \max_phase) \quad (7.2)$$

Case 2 (Inductive Step, $d \geq 2$). For any trees (f_1, \triangleleft) and (f_2, \triangleleft) let the inductive hypothesis be:

$$IH \triangleq \min_phase((f_1, \triangleleft), d-2) > \min_phase((f_2, \triangleleft), d-2)$$

$$\iff (\forall f_1^1 \triangleleft f_1)(\exists f_1^2 \triangleleft f_1^1)(\forall f_1^3 \triangleleft f_1^2) \cdots (\forall f_1^{d-3} \triangleleft f_1^{d-4})(\exists f_1^{d-2} \triangleleft f_1^{d-3})$$

$$(\exists f_2^1 \triangleleft f_2)(\forall f_2^2 \triangleleft f_2^1)(\exists f_2^3 \triangleleft f_2^2) \cdots (\exists f_2^{d-3} \triangleleft f_2^{d-4})(\forall f_2^{d-2} \triangleleft f_2^{d-3})$$

$$u(f_1^{d-2}) > u(f_2^{d-2})$$

To show:

$$\min_phase((e_1, \triangleleft), d) > \min_phase((e_2, \triangleleft), d)$$

$$\iff (\forall e_1^1 \triangleleft e_1)(\exists e_1^2 \triangleleft e_1^1)(\forall e_1^3 \triangleleft e_1^2) \cdots (\forall e_1^{d-1} \triangleleft e_1^{d-2})(\exists e_1^d \triangleleft e_1^{d-1})$$

$$(\exists e_2^1 \triangleleft e_2)(\forall e_2^2 \triangleleft e_2^1)(\exists e_2^3 \triangleleft e_2^2) \cdots (\exists e_2^{d-1} \triangleleft e_2^{d-2})(\forall e_2^d \triangleleft e_2^{d-1})$$

$$u(e_1^d) > u(e_2^d)$$

$$\min_phase((e_1, \triangleleft), d) > \min_phase((e_2, \triangleleft), d) \quad (7.3)$$

$$\iff \min_{e_1^1 \triangleleft e_1} \left(\max_phase(e_1^1, d-1) \right) > \min_{e_2^1 \triangleleft e_2} \left(\max_phase(e_2^1, d-1) \right) \quad (7.4)$$

$$\iff (\forall e_1^1 \triangleleft e_1)(\exists e_2^1 \triangleleft e_2) \quad (7.5)$$

$$\max_phase(e_1^1, d-1) > \max_phase(e_2^1, d-1)$$

$$\iff (\forall e_1^1 \triangleleft e_1^1)(\exists e_2^1 \triangleleft e_2)$$

$$\max_{e_1^2 \triangleleft e_1^1} \left(\min_phase(e_1^1, d-2) \right) > \max_{e_2^2 \triangleleft e_2^1} \left(\min_phase(e_2^1, d-2) \right) \quad (7.6)$$

$$\begin{aligned} &\iff (\forall e_1^1 \triangleleft e_1)(\exists e_2^1 \triangleleft e_2)(\exists e_1^2 \triangleleft e_1^1)(\forall e_2^2 \triangleleft e_2^1) \\ &\quad \min_phase(e_1^2, d-2) > \min_phase(e_2^2, d-2) \end{aligned} \quad (7.7)$$

Using the inductive hypothesis, where $(f_1, \triangleleft) = (e_1^2, \triangleleft)$ and $(f_2, \triangleleft) = (e_2^2, \triangleleft)$, yields:

$$\begin{aligned} &\iff (\forall e_1^1 \triangleleft e_1) (\exists e_2^1 \triangleleft e_2) (\exists e_1^2 \triangleleft e_1^1) (\forall e_2^2 \triangleleft e_2^1) \\ &\quad (\forall e_1^3 \triangleleft e_1^2) (\exists e_1^4 \triangleleft e_1^3) (\forall e_1^5 \triangleleft e_1^4) \cdots (\forall e_1^{d-1} \triangleleft e_1^{d-2}) (\exists e_1^d \triangleleft e_1^{d-1}) \\ &\quad (\exists e_2^3 \triangleleft e_2^2) (\forall e_2^3 \triangleleft e_2^2) (\exists e_2^4 \triangleleft e_2^3) \cdots (\exists e_2^{d-1} \triangleleft e_2^{d-2}) (\forall e_2^d \triangleleft e_2^{d-1}) \\ &\quad u(e_1^d) > u(e_2^d) \end{aligned} \quad (7.8)$$

$$\begin{aligned} &\iff (\forall e_1^1 \triangleleft e_1) (\exists e_1^2 \triangleleft e_1^1) (\forall e_1^3 \triangleleft e_1^2) \cdots (\forall e_1^{d-1} \triangleleft e_1^{d-2}) (\exists e_1^d \triangleleft e_1^{d-1}) \\ &\quad (\exists e_2^1 \triangleleft e_2) (\forall e_2^2 \triangleleft e_2^1) (\exists e_2^3 \triangleleft e_2^2) \cdots (\exists e_2^{d-1} \triangleleft e_2^{d-2}) (\forall e_2^d \triangleleft e_2^{d-1}) \\ &\quad u(e_1^d) > u(e_2^d) \end{aligned} \quad (7.9)$$

Equation 7.9 follows from the fact you can always choose the child that produces the maximum/minimum and $(\forall a)(\forall b) P(a, b) \equiv (\forall b)(\forall a) P(a, b)$ and $(\exists a)(\exists b) P(a, b) \equiv (\exists b)(\exists a) P(a, b)$. □

Now, let the states be Context Dependent Partial Interpretations and we can define a new $ILLP_{LOAS}^{deep}$ task that uses the explanation condition to select a subset of the leaves of a game tree to bravely respect, these preferences *explain* why at the root of the tree there is a preference, according to the minimax theorem. In order to define this new task we must first define a new ordering type that incorporates *explanation*. To do so we must first alter the definition of *explanation* – in a manner that preserves Proposition 1 – by replacing the comparison of utilities $u(e_1^d) > u(e_2^d)$ with $B \cup H$ bravely respects $\langle e_1^d, e_2^d \rangle$, where $B \cup H$ are given. This can be done as there is a mapping from the penalties given by the weak constraints to the natural numbers.

Definition 7.2.3 (Deep Context Dependent Ordering Example). This definition extends Definition 3.4.9. A deep CDOE, o , is a tuple $\langle e_1, e_2, d, \triangleleft \rangle$ of two trees, (e_1, \triangleleft) and (e_2, \triangleleft) , where e_1, e_2 are CDPIS, \triangleleft is a child relation over the CDPIS, and a depth d . An ASP program Π deeply respects o iff *explanation* $(\Pi, (e_1, \triangleleft), (e_2, \triangleleft), d)$.

Definition 7.2.4 (Deep, Context Dependent LOAS task). This definition extends Definition 3.4.10. A deep, context dependent Learning from Ordered Answer Sets task is a tuple $T = \langle B, S_M, E, O, \triangleleft \rangle$ where B, S_M, E are as before in Definition 3.4.10. $O = \langle O^b, O^c, O^d \rangle$ where O^d is the set of deep orderings over $\mathcal{S} \subseteq E^+$, \mathcal{S} is the set of all board

states.¹ \triangleleft is a partial ordering over \mathcal{S} denoting the child relation. A hypothesis H is an inductive solution of T , $H \in ILP_{LOAS}^{deep}$ iff:

- (i) $H' \in ILP_{LAS}^{context}(\langle B, S_M^{LAS}, E^+, E^- \rangle)$, where H' is the subset of H with no weak constraints, $M^{LAS} = \langle M_h, M_b \rangle$.
- (ii) $\forall o \in O^b$, such that $B \cup H$ bravely respects o
- (iii) $\forall o \in O^c$, such that $B \cup H$ cautiously respects o
- (iv) $\forall o \in O^d$, such that $B \cup H$ deeply respects o .

7.3 IMPLEMENTATION

The implementation of the new ILP_{LOAS}^{deep} task is detailed below. Firstly we describe an experimental feature of ILASP 3.2 (meta-program injection), This feature allows certain examples to be activated, through the custom definition of the `example_active(·)` predicate. From here we present the injected meta-program we use to encode the explanation condition, and prove that this is equivalent to the ILP_{LOAS}^{deep} task. We also present a method of generating the child relation, and board states from the GDL specification.

7.3.1 ILASP WITH META-PROGRAM INJECTION

In order to be able to modify the meta-level representation we must use the meta-program injection feature of ILASP. The following definition of the feature is from (Law, Russo and Broda, 2018):

Definition 7.3.1 (Meta-Program Injection). Let $T = \langle B, S_M, \langle E^+, E^-, O^b, O^c \rangle \rangle$ be an $ILP_{LOAS}^{context}$ task and \mathcal{Q} be an ASP program. Given any interpretation I of \mathcal{Q} , we T_I denotes the task $\langle B, S_M, \langle E_I^+, E_I^-, O_I^b, O_I^c \rangle \rangle$, where:

$$\begin{aligned} E_I^+ &= \{e \in E^+ \mid \text{example_active}(e_{id}) \in I\} \\ E_I^- &= \{e \in E^- \mid \text{example_active}(e_{id}) \in I\} \\ O_I^b &= \{o \in O^b \mid \text{example_active}(e_{id}) \in I\} \\ O_I^c &= \{o \in O^c \mid \text{example_active}(e_{id}) \in I\} \end{aligned}$$

A hypothesis $H \subseteq S_M$ is said to be an *inductive solution* of T with respect to the injection of \mathcal{Q} iff $\exists I \in \mathcal{AS}(\mathcal{Q})$ such that $H \in ILP_{LOAS}^{context}(T_I)$.

7.4 TRANSLATION FROM DEEP, CONTEXT DEPENDENT LOAS TASK TO CONTEXT DEPENDENT LOAS TASK WITH META-PROGRAM INJECTION

For any ILP_{LOAS}^{deep} task, $T = \langle B, S_M, E, O, \triangleleft \rangle$, we translate T to $T' = \langle B, S_M, \langle E^+, E^-, O_{inject}^b, O^c \rangle \rangle$ the $ILP_{LOAS}^{context}$ with \mathcal{Q} as follows:

¹ When presenting examples only the board states that are used will be shown

E^+ , E^- and O^c are left unchanged. O_{leaves}^b is the set of all possible combinations of leaves at depth d on the left branch and leaves at depth d on the right branch

$$O_{leaves}^b = \bigcup_{\langle e_1, e_2, d, \triangleleft \rangle \in O^d} \left\{ \langle e_1^d, e_2^d \rangle \mid \forall e_1^d \triangleleft^d e_1, \forall e_2^d \triangleleft^d e_2 \right\}$$

$$O_{inject}^b = O^b \cup O_{leaves}^b$$

$$\begin{aligned} \mathcal{Q} = & meta(E^+) \cup meta(E^-) \cup meta(O^b) \cup meta(O_{leaves}^b) \\ & \cup meta(O^c) \cup meta(O^d) \cup meta(\triangleleft) \end{aligned}$$

\triangleleft^d is the d^{th} power of \triangleleft , defined inductively by

$$\triangleleft^0 = \{(e, e) \mid e \in \mathcal{S}\}$$

$$\triangleleft^1 = \triangleleft$$

$$\triangleleft^{i+1} = \triangleleft \circ \triangleleft^i \quad (\text{for } i > 0, \circ \text{ is functional composition})$$

and the *meta* function is as follows:

$$meta(E^+) = \{example_active(e_{id}) \mid e_{id} \in E^+\} \quad (7.10)$$

$$meta(E^-) = \{example_active(e_{id}) \mid e_{id} \in E^-\} \quad (7.11)$$

$$meta(O^b) = \{example_active(o_{id}) \mid o_{id} \in O^b\} \quad (7.12)$$

$$meta(O^c) = \{example_active(o_{id}) \mid o_{id} \in O^c\} \quad (7.13)$$

$$meta(O_{leaves}^b) = \left\{ ord(e_{id}^1, e_{id}^2, o_{id}) \mid o = \langle e^1, e^2 \rangle \in O_{leaves}^b \right\} \quad (7.14)$$

$$meta(O^d) = \left\{ \begin{array}{l} root(e_1, chosen) \\ root(e_2, other) \end{array} \mid o \in O^d \right\} \quad (7.15)$$

$$meta(\triangleleft) = \{child(a, b) \mid b \triangleleft a\} \quad (7.16)$$

Equations 7.10–7.13 add facts to the meta-program ensuring that the positive, negative, brave and cautious orderings must be covered. Equation 7.14 adds a set of facts denoting the possible orderings of the leaves at depth d . Equation 7.15 adds the root trees to the meta-program. Finally, we add a fixed program responsible for searching through all the possible brave orderings, O_{leaves}^d , in order to find the set

$$\bigcup_{\langle e_1, e_2, d, \triangleleft \rangle \in O^d} \left\{ \langle e_1^d, e_2^d \rangle \mid explanation(B \cup H, (e_1, \triangleleft), (e_2, \triangleleft), d) \right\}$$

for a hypothesis, H , that will be learnt by the task. This can be achieved using a choice rule for selection over the children of a branch such that $\exists e^{i+1} \triangleleft e^i$, and a normal rule for enforcing $\forall e^{i+1} \triangleleft e^i$.

$$explanation(EX_ID, forall) \leftarrow root(EX_ID, chosen) \quad (7.17)$$

$$explanation(EX_ID, exists) \leftarrow root(EX_ID, other) \quad (7.18)$$

$$1\{explanation(Child,forall) \mid child(Parent,Child)\}1 \quad (7.19)$$

$$\leftarrow explanation(Parent,exists),$$

$$child(Parent,_)$$

$$explanation(Child,exists) \leftarrow child(Parent,Child), \quad (7.20)$$

$$explanation(Parent,forall)$$

$$example_active(ORD_ID) \leftarrow ord(EX_ID_1,EX_ID_2,ORD_ID), \quad (7.21)$$

$$explanation(EX_ID_1,_),$$

$$explanation(EX_ID_2,_).$$

Proposition 2. Given the task $T = \langle B, S_M, E, O, \triangleleft \rangle$, the meta-program \mathcal{Q} from the translation of T into a meta-program injection task, and a hypothesis $H, \exists I \in \mathcal{AS}(\mathcal{Q})$ such that

$$\forall o \in O_{leaves}^b(I), B \cup H \text{ bravely respects } o \iff \forall o \in O^d, B \cup H \text{ deeply respects } o$$

$$\text{where } O_{leaves}^b(I) = \{o \in O_{leaves}^b \mid example_active(o_{id}) \in I\}$$

Proof. Let $T = \langle B, S_M, E, O, \triangleleft \rangle$, \mathcal{Q} be the translation of T into a meta-program injection task, and hypothesis H .

$$(\forall o \in O^d) B \cup H \text{ deeply respects } o \quad (7.22)$$

$$\iff (\forall o \in O^d)(\forall o' \in O_{leaves}^b(I, o)) B \cup H \text{ bravely respects } o \quad (7.23)$$

$$\iff I \in \mathcal{AS}(\mathcal{Q}) \forall example_active(o_{id}) \in I. H \text{ covers } o \quad (7.24)$$

$$\iff \forall o \in O_{leaves}^b(I) B \cup H \text{ bravely respects } o \quad (7.25)$$

Equation 7.24 comes from the fact that Rules 7.20 and 7.19 clearly generate all the children and exactly one child, respectively. All the terms needed to satisfy these predicates are given as facts in \mathcal{Q} (see $meta(O^d)$, $meta(O_{leaves}^b)$, $meta(\triangleleft)$). We know that $B \cup H$ bravely respects this ordering from the explanation condition. *Note.* The upper bound of 1 is not required, it just limits the number of examples that are generated, the important bound is the lower bound, stating that *there must be at least one*. \square

Theorem 1. Given the task $T = \langle B, S_M, E, O, \triangleleft \rangle$ and the task T' with \mathcal{Q} created from the translation of T into a meta-program injection task then for all $H \subseteq S_M$,

$$H \in ILP_{LOAS}^{deep}(T) \iff H \in ILP_{LOAS}^{context}(T') \text{ with } \mathcal{Q}$$

Proof. In order for $H \in ILP_{LOAS}^{deep}(T)$ it must satisfy the following (from Definition 7.2.4):

- (i) $H' \in ILP_{LAS}^{context}(\langle B, S_M^{LAS}, E^+, E^- \rangle)$, where H' is the subset of H with no weak constraints, $M^{LAS} = \langle M_h, M_b \rangle$.
- (ii) $\forall o \in O^b, B \cup H$ bravely respects o
- (iii) $\forall o \in O^c, B \cup H$ cautiously respects o

(iv) $\forall o \in O^d, B \cup H$ deeply respects o .

In order for $H \in ILP_{LOAS}^{context}(T')$ with \mathcal{Q} the following must hold: $\exists I \in \mathcal{AS}(\mathcal{Q})$ such that $H \in ILP_{LOAS}^{context}(T'_I)$, and for $H \in ILP_{LOAS}^{context}(T'_I)$ the following must be satisfied (from Definition 7.3.1, Definition 3.4.7):

(v) $H' \in ILP_{LAS}^{context}(\langle B, S_M^{LAS}, E_I^{+'}, E_I^{-'} \rangle)$, where H' is the subset of H with no weak constraints, $M^{LAS} = \langle M_h, M_b \rangle$.

(vi) $\forall o \in O_I^{b'}, B \cup H$ bravely respects o

(vii) $\forall o \in O_I^{c'}, B \cup H$ cautiously respects o

It remains to show that (i) \Leftrightarrow (v), (iii) \Leftrightarrow (vii), and (ii) and (iv) \Leftrightarrow (vi).

1. (i) \Leftrightarrow (v) B and S_M are left unchanged by the translation. From $meta(E^+)$, $\forall e \in E^+$ the fact $example_active(e_{id})$ is in \mathcal{Q} therefore it is also in I therefore $\forall e . e \in E^+ \Leftrightarrow e \in E_I^{+'}$. A similar argument can be made for E^- .
2. (iii) \Leftrightarrow (vii) A similar argument as above shows that $\forall o . o \in O^c \Leftrightarrow O_I^{c'}$.
3. (ii) and (iv) \Leftrightarrow (vi) $O_I^{b'}$ is split, by definition, into $O_I^b \cup O_{leaves}^b(I)$. Again, using a similar argument to 1. and 2., it is sufficient to show that

$$\forall o \in O_{leaves}^b(I), B \cup H \text{ bravely respects } o \iff \forall o \in O^d, B \cup H \text{ deeply respects } o$$

where $O_{leaves}^b(I) = \{o \in O_{leaves}^b \mid example_active(o_{id}) \in I\}$. This follows directly from Proposition 2. □

7.5 AUTOMATICALLY GENERATING THE GAME TREES

As the depth and number of deep orderings increases the number of examples and orderings that are needed grow rapidly. Fortunately, the structure of the GDL-based logic programs can be leveraged in order to generate the additional positive examples in the tree, and the child relation, \triangleleft , automatically. Each root example can be combined with its context and the background and using Clingo the children can be created using the methodology outlined below.

We define two mutually recursive functions (algorithm 7.1 and algorithm 7.2). One which for each answer set of of an example will generate the branches and collect the child relation and all leaf examples from the branches. The other will take an answer set an extract the next state (i. e. take the $next(\cdot)$ and make that the current state), and then aggregate the results of making the branches from each child.

The *Branch Generation* algorithm is the entry point of the generation, a root node is passed in along with the background knowledge and a depth to generate. If $d = 0$ then the entry node is the child, which is already in the examples, and so empty sets are returned. Otherwise, the *Child Generation* algorithm is called for each answer set of $B \cup C$, this is done by calling out to Clingo, and the results are then aggregated.

Algorithm 7.1: Branch Generation

Input: Background Knowledge: B **Input:** Positive Example: $\langle e, C \rangle \in \mathcal{S}$ **Input:** Depth: $d \geq 0$ **Result:** A set of $\langle e, C \rangle$ for each leaf at depth d , and the child relation, \triangleleft

```

begin
  if  $d = 0$  then
    return  $\emptyset, \emptyset$ 
  else
     $(decendents, E_{child}^+) \leftarrow$ 
       $unzip \{ChildGen(B, e_{id}, d - 1, A) \mid A \in \mathcal{AS}(B \cup C)\}$ 
    return  $\cup decendents, \cup E_{child}^+$ 
  end
end

```

Algorithm 7.2: Children Generation

Input: Background Knowledge: B **Input:** Positive Example ID: e_{id} **Input:** Depth: $d \geq 0$ **Input:** Answer Set: A **Result:** Tuple of the child relation, \triangleleft , and the child examples, E_{child}^+

```

begin
   $state \leftarrow \{f \mid next(f) \in A\}$ 
   $E_{child}^+ \leftarrow \{\{move\}, \emptyset, state\} \mid move \in legal(state)\}$ 
   $\triangleleft \leftarrow \{(e'_{id}, e_{id}) \mid e' \in E_{child}^+\}$ 
  if  $d = 0$  then
    return  $\triangleleft, E_{child}^+$ 
  else
     $(decendents, E_{leaves}^+) \leftarrow unzip \{BranchGen(B, e', d) \mid e' \in E_{child}^+\}$ 
    return  $\cup_{\triangleleft' \in decendents \cup \triangleleft} \triangleleft', \cup E_{leaves}^+$ 
  end
end

```

The *Child Generation* algorithm takes as input the background knowledge, the ID of the current example and an Answer Set of the example and the background knowledge. The depth, d , is also passed in, and is the value we reduce on in order to terminate. Firstly the next state is created from the answer set, and the child examples and child relation are generated from this state. Here we use the *legal* which simply returns the set of legal moves for the game, in practice this can be done by, again, calling out to Clingo.

Example 7.5.1. Take the following tree from Chapter 8, it depicts a game tree for the CrossDot game (Section 2.3). The left branch denotes the chosen path by the player, and the right branch denotes another possible move. The strategy here is that not letting your opponent win is a good idea. If, however, you were playing by the strategy “move into an isolated box or next to your own marker” (a dominant strategy for the first player) then the right branch would be preferred.




Figure 7.5. Cross-Dot *defence* game tree generated from $\langle (\cdot, [X][][][], a_2), (\cdot, [X][][.][], a_3), 2 \rangle$, the ordering example, with a depth of 2. Green and red represent a winning and losing state for \cdot , respectively.

The background knowledge, B , for this task can be found in Logic Program A.3. The mode declaration is $M = \langle \emptyset, \emptyset, \{goal(v_{role}, 100)\}, \{1, -1\}, 1 \rangle$. The only positive examples, E^+ , given are:

$$\begin{aligned}
 e_1 &= \langle \emptyset, \emptyset, C_1 \rangle \\
 e_2 &= \langle \emptyset, \emptyset, C_2 \rangle \\
 C_1 &= \begin{cases} box(1, x). & box(2, b). & box(3, b). & box(4, b). \\ control(o). \\ \leftarrow not\ does(o, mark(2)). \end{cases} \\
 C_2 &= \begin{cases} box(1, x). & box(2, b). & box(3, b). & box(4, b). \\ control(o). \\ \leftarrow not\ does(o, mark(3)). \end{cases}
 \end{aligned}$$

The only deep ordering example given is $\langle e_1, e_2, 2, \triangleleft \rangle$. No other examples are given. The full task is $T = \langle B, S_M, E, O, \triangleleft \rangle$.

Next we generate the additional children examples and the child relation \triangleleft .

$$\begin{aligned}
 e_{1,1,1} &= \langle \emptyset, \emptyset, C_3 \rangle & e_{1,2,1} &= \langle \emptyset, \emptyset, C_4 \rangle & e_{2,1,1} &= \langle \emptyset, \emptyset, C_5 \rangle & e_{2,2,1} &= \langle \emptyset, \emptyset, C_6 \rangle \\
 C_3 &= \begin{cases} box(1, x). & box(2, o). & box(3, b). & box(4, x). \\ control(o). \\ \leftarrow not\ does(o, mark(4)). \end{cases}
 \end{aligned}$$

$$C_4 = \begin{cases} \text{box}(1, x). \text{ box}(2, o). \text{ box}(3, x). \text{ box}(4, b). \\ \text{control}(o). \\ \leftarrow \text{not does}(o, \text{mark}(3)). \end{cases}$$

$$C_5 = \begin{cases} \text{box}(1, x). \text{ box}(2, x). \text{ box}(3, o). \text{ box}(4, b). \\ \text{control}(o). \\ \leftarrow \text{not does}(o, \text{noop}). \end{cases}$$

$$C_6 = \begin{cases} \text{box}(1, x). \text{ box}(2, b). \text{ box}(3, o). \text{ box}(4, x). \\ \text{control}(o). \\ \leftarrow \text{not does}(o, \text{mark}(2)). \end{cases}$$

$$\triangleleft = \begin{cases} (e_1, e_{1,1}), & (e_1, e_{1,2}), \\ (e_2, e_{2,1}), & (e_2, e_{2,2}), \\ (e_{1,1}, e_{1,1,1}), & (e_{1,2}, e_{1,2,1}), \\ (e_{2,1}, e_{2,1,1}), & (e_{2,2}, e_{2,2,1}), \end{cases}$$

All e are added to E^+ and \triangleleft is as defined above. Next we translate T into T' the meta-program injection task. $T' = \langle B, S_M, \langle E^+, E^-, O_{inject}^b, O^c \rangle \rangle$ with \mathcal{Q} . E^+ , E^- , and O^c are unchanged, as per the translation.

$$O_{inject}^b = O^b \cup \left\{ \begin{array}{l} \langle e_{1,1,1}, e_{2,1,1} \rangle, \quad \langle e_{1,1,1}, e_{2,2,1} \rangle \\ \langle e_{1,2,1}, e_{2,1,1} \rangle, \quad \langle e_{1,2,1}, e_{2,2,1} \rangle \end{array} \right\}$$

$$\mathcal{Q} = \begin{cases} \text{root}(e_1, \text{chosen}) & \text{root}(e_2, \text{other}) \\ \\ \text{ord}(e_{1,1,1}, e_{2,1,1}, o_1) & \text{ord}(e_{1,1,1}, e_{2,2,1}, o_2) \\ \text{ord}(e_{1,2,1}, e_{2,1,1}, o_3) & \text{ord}(e_{1,2,1}, e_{2,2,1}, o_4) \\ \\ \text{example_active}(e_1) & \text{example_active}(e_2) \\ \text{example_active}(e_{1,1,1}) & \text{example_active}(e_{1,2,1}) \\ \text{example_active}(e_{2,1,1}) & \text{example_active}(e_{2,2,1}) \\ \\ \text{explanation}(E, \text{forall}) \leftarrow \text{root}(E, \text{chosen}) \\ \text{explanation}(E, \text{exists}) \leftarrow \text{root}(E, \text{other}) \\ \\ 1 \{ \text{explanation}(C, \text{forall}) \mid \text{child}(P, C) \} 1 \leftarrow \text{child}(P, _), \text{explanation}(P, \text{forall}) \\ \text{explanation}(C, \text{exists}) \leftarrow \text{child}(P, C), \text{explanation}(P, \text{forall}) \\ \\ \text{example_active}(O) \leftarrow \text{ord}(E1, E2, O), \text{explanation}(E1, _), \text{explanation}(E2, _) \end{cases}$$

The solution of the task is H iff $\exists I \in \mathcal{AS}(\mathcal{Q})$ such that $H \in ILP_{\text{LOAS}}^{\text{context}}(T'_I)$. H is the hypothesis below:

$$\rightsquigarrow \text{goal}(P, 100), \text{control}(P). [-1@1, P]$$

and $I \in \mathcal{AS}(\mathcal{Q})$ contains exactly the following $\text{example_active}(o_{id})$ predicates (and other predicates, e. g. root etc.): $I_{\text{example_active}} = \{ \text{example_active}(o_1), \text{example_active}(o_2) \}$

8 | CASE STUDY: CROSS-DOT GAME

In this chapter I focus on an alternative game, which I have taken from the literature. The game I will look at has also been studied in Zhang and Thielscher (2015). The Cross-Dot Game, another name for a m - k game, where players aim to get k boxes in a row marked with their player symbol (from a possible m boxes), the full description and rules of the game can be found in Section 2.3. An example of a game state may look as follows: $\boxtimes \square \boxdot \square \square$.

As discussed in Section 4.2, Zhang and Thielscher (2015) present an extension to Game Description Language which includes a modal operator (\bigcirc), the ability to represent strategies in their language (not just the game rules), some preference operators for the strategies (Δ , ∇) and a translation of a subset of their language into ASP.

Zhang et al. use the following definition of a strategy. A strategy of player i is $S(\varphi) \subseteq W \times A^i$ where W is the set of game states, A^i is the set of actions player i can take and φ is a strategy rule that is true in all ω such that $(\omega, a) \in S(\varphi)$. A strategy S is only valid if $S(\varphi) \neq \emptyset$.

The notion of prioritised conjunction (Δ) and prioritised disjunction (∇) were introduced by Zhang and Thielscher in order to combine smaller strategy rules into more complex strategies. The semantics of $r_1 \Delta \dots \Delta r_n$ are: apply as many strategy rules from the left as possible whilst still maintaining a valid strategy for this state. The semantics of $r_1 \nabla \dots \nabla r_n$ are: try r_1 and if that does not work try r_2 and so on.

In their paper Zhang et al. present strategies written in their modal language, they then compare them and show interesting properties of each. The experiments in this chapter take examples of games that have been played using the strategies and learn an equivalent strategy in ASP. Experiments 8.1–8.5 use context-dependent LOAS tasks and Experiment 8.6 onwards use our deep, context-dependent LOAS tasks from Chapter 7.

8.1 THE GAME

The game state is denoted by a tuple of current player and current state $(p, \square \square \dots \square)$. For instance, $(\times, \boxtimes \boxdot \square \square)$ where \times is to move and each player has had one turn playing in the first two boxes. Moves, or state-action pairs, are represented as a tuple of state and action $(p, \square \square \dots \square, a_b)$ where a_b represents marking box b .

Example 8.1.1. Taking the example game state above and the action a_3 , we get the move $(\times, \boxtimes \boxdot \square \square, a_3)$. When this move is made the following state change occurs: $(\times, \boxtimes \boxdot \square \square) \xrightarrow{a_3} (\cdot, \boxtimes \boxdot \boxtimes \square)$. If we preferred this move to $(\times, \boxtimes \boxdot \square \square, a_4)$ we can express this as the CDOE (Definition 3.4.9): $\langle (\times, \boxtimes \boxdot \square \square, a_3), (\times, \boxtimes \boxdot \square \square, a_4) \rangle$. This notation is used throughout the chapter.

8.1.1 REPRESENTATION

To represent the game I have used a modified version of the ASP presented in the paper. This is to eliminate certain duplication, and due to some restrictions imposed by the ASP allowed in ILASP, e.g. conditions. The representation I have used follows from the GDL

specification in Section 5.2. `next` means the predicate is true after the move. `does` means a move is taken. `legal` means a move is possible given `true`. The representation used in the paper has multiple rules for some predicates that deal with boxes that can be on either side; which I overcome by introducing `adj(·, ·)` meaning two boxes are adjacent. Additionally, I introduce `longest_chain(·, ·)` to denote a player's longest chain to avoid creating multiple instances of the program for different ks . The full logic program used as background knowledge for the game can be found in Logic Program A.3.

8.2 LEARNING STRATEGIES

In order to demonstrate the capabilities of weak constraints I use `ILASP 3` to learn the strategies presented and compare the actions chosen with the actions selected using the strategies from Zhang and Thielscher (2015). The strategies in the paper can be divided into three categories: simple, prioritised and forward-thinking. These strategy rules are shown in Table 8.1. In all the examples in the following section partial interpretations are written in a simplified format using the aforementioned state-action representation.

Note. Throughout this chapter I use E_O^+ to denote the positive examples drawn from the ordering examples (i. e. O^b , O^c and O^d).

SIMPLE	PRIORITISED	FORWARD-THINKING
<i>fill_next</i>	<i>combined</i>	<i>defence</i>
<i>fill_isolated</i>	<i>fill_leftmost</i>	<i>cautious</i>
<i>fill_any</i>	<i>thoughtful</i>	<i>fill_o_next</i> <i>passive_defence</i>

Table 8.1. Categorised strategy rules from Zhang and Thielscher (2015)

8.2.1 SIMPLE STRATEGIES

Simple strategies can be learnt using only a single priority level for the weak constraints.

Experiment 8.1 (Fill Next). The *fill_next* strategy means if possible, mark a box next to one of your own. Given the Context Dependent LOAS task (Definition 3.4.10) $T_{next} = \langle B, S_M, E_O^+, \emptyset, O^b, \emptyset \rangle$ where

B is the background knowledge from Logic Program A.3,

$$M = \langle \emptyset, \emptyset, \{adj(v, v), does(c, mark(v)), box(v, c)\}, \{-1\}, 1 \rangle$$

$$O^b = \left\{ \begin{array}{l} \langle (\times, \boxtimes \square \square \square \square \square, a_2), (\times, \boxtimes \square \square \square \square \square, a_3) \rangle \\ \langle (\times, \square \square \boxtimes \square \square \square, a_2), (\times, \square \square \boxtimes \square \square \square, a_5) \rangle \\ \langle (\times, \square \square \boxtimes \square \square \square, a_2), (\times, \square \square \boxtimes \square \square \square, a_6) \rangle \\ \langle (\times, \square \square \boxtimes \square \square \square, a_4), (\times, \square \square \boxtimes \square \square \square, a_5) \rangle \\ \langle (\times, \square \square \square \square \boxtimes \square, a_6), (\times, \square \square \square \square \boxtimes \square, a_3) \rangle \end{array} \right.$$

The ordering examples shown here are taken from a few games against a player (\cdot) using a random strategy. Each example is from the \times player's point of view.

ILASP was able to learning the following hypothesis¹:

$$\Leftarrow \text{does}(x, \text{mark}(\text{Box1})), \text{adj}(\text{Box1}, \text{Box2}), \text{box}(\text{Box2}, x).[-1@1, \text{Box1}, \text{Box2}]$$

In English this hypothesis means you² gain a reward of 1 if you mark a box next to a box you have already marked.

Experiment 8.2 (Fill isolated). The *fill_next* strategy has a flaw, if you play in the first box your opponent can easily block you off. To remedy this you may wish to play in one of the centre cells as you are guaranteed to win on your next turn (assuming $k = 2$). The *fill_isolated* strategy means if possible, mark a box with no mark either side of the box. Given the learning task $T_{iso} = \langle B, S_M, E_O^+, \emptyset, O^b, \emptyset \rangle$ where

B, M are the same as in T_{next} ,

$$O^b = \begin{cases} \langle (\times, \square \square \square \square \square \square, a_2), (\times, \square \square \square \square \square \square, a_1) \rangle \\ \langle (\times, \square \square \square \square \square \square, a_2), (\times, \square \square \square \square \square \square, a_6) \rangle \\ \langle (\times, \square \square \square \square \square \square, a_3), (\times, \square \square \square \square \square \square, a_6) \rangle \\ \langle (\times, \boxtimes \square \square \square \square \square \square, a_4), (\times, \boxtimes \square \square \square \square \square \square, a_3) \rangle \end{cases}$$

The examples here show what you might play at the beginning of the game with this strategy, and what you might do later in the game if you made a mistake earlier on.

ILASP was able to learning the following hypothesis¹:

$$\Leftarrow \text{does}(x, \text{mark}(\text{Box1})), \text{adj}(\text{Box1}, \text{Box2}), \text{true}(\text{box}(\text{Box2}, b)).[-1@1, \text{Box1}, \text{Box2}]$$

In English this hypothesis means you² try to maximise the number of blank boxes either side of you.

Remark. *fill_any* expresses a lack of preference over the actions and so does not need anything to be learnt.

8.2.2 COMBINED STRATEGIES

Combined strategies use the prioritised operators Δ and ∇ to compose basic notions into something more complicated.

Experiment 8.3 (Combined). The *combined* is the composite strategy $\text{fill_next} \nabla \text{fill_isolated} \nabla \text{fill_any}$, meaning you first try to fill a box next to your own, then fill an isolated box and if there are no moves that conform to this strategy you fill any box you can. Given the learning task $T_{comb} = \langle B, S_M, E_O^+, \emptyset, O^b, \emptyset \rangle$ where

B is the same as in T_{next} ,

$$M = \langle \emptyset, \emptyset, \{ \text{adj}(v, v), \text{does}(c, \text{mark}(v)), \text{true}(\text{box}(v, c)) \}, \{-1, \text{box}\}, 2 \rangle$$

$$O^b = O_{next}^b \cup O_{iso}^b$$

¹ Variable names have been replaced for clarity

² Here we assume you are playing as \times

Note. Here we are able to use the union of the examples used for the simple strategies (O_{next}^b, O_{iso}^b) as they encode the preference of *fill_next* over *fill_isolated*. Specifically, the example $\langle (\times, \boxtimes \square \square \square \square \cdot, a_2), (\times, \boxtimes \square \square \square \square \square, a_3) \rangle$, and there are no counter examples to this preference. Though this is not always the case.

ILASP 3 was able to learning the following hypothesis¹:

$$\begin{aligned} &\rightsquigarrow \text{does}(x, \text{mark}(\text{Box1})), \text{adj}(\text{Box1}, \text{Box2}), \text{true}(\text{box}(\text{Box2}, x)).[-1@2, \text{Box1}, \text{Box2}] \\ &\rightsquigarrow \text{does}(x, \text{mark}(\text{Box1})), \text{adj}(\text{Box1}, \text{Box2}), \text{true}(\text{box}(\text{Box2}, b)).[-1@1, \text{Box1}, \text{Box2}] \end{aligned}$$

In English this hypothesis means you² try to go next to your own piece first, and if this is not possible minimise the number of dots next to you. I believe this is equivalent to the other strategy in this game.

Experiment 8.4 (Fill Leftmost). The *fill_leftmost* is the composite strategy $c_m \Delta \dots \Delta c_1$ where $c_i = \bigvee_{j=1}^i \text{does}(a_j)$ such that $1 \leq i \leq m$ (e.g. if $\text{does}(a_2)$ is an option at the current state ω and $\text{does}(a_1)$ is not then $(\omega, a_2) \in S(c_m \Delta \dots \Delta c_1)$ but $(\omega, a_{-2}) \notin S(c_m \Delta \dots \Delta c_1)$ where -2 is everything bar 2).

Remark. This representation of this strategy is rather unintuitive, we will see that in ILASP this is far simpler and more self explanatory.

Given the learning task $T_{left} = \langle B, S_M, E_O^+, \emptyset, O^b, \emptyset \rangle$ where

B, M are the same as in T_{comb} ,

$$O^b = \{ \langle (\times, \boxtimes \cdot \square \square \square \square \boxtimes, a_3), (\times, \boxtimes \square \square \square \square \square \boxtimes, a_4) \rangle \}$$

ILASP 3 was able to learning the following hypothesis¹:

$$\rightsquigarrow \text{does}(x, \text{mark}(\text{Box1})).[\text{Box1}@1, \text{Box1}]$$

In English, this hypothesis means you gain a penalty proportional to the position of the box, i.e. the further left the lower the penalty.

Experiment 8.5 (Thoughtful). The *thoughtful* strategy means the same as the combined strategy but in all cases pick the left most option when faced with a choice. Formally this is, *combined* Δ *fill_leftmost*. Given the learning task $T_{thoughtful} = \langle B, S_M, E_O^+, \emptyset, O^b, \emptyset \rangle$ where

$$\begin{aligned}
B &= B_{comb} \cup \{\text{leftmost}(B, 7 - B) \leftarrow \text{box}(B)\}, \\
M &= \langle M^h, M^b, M^o, \{-1, \text{box}, \text{score}\}, 2 \rangle \\
M^h &= \{\text{fill_next}(v_{\text{box}}, v_{\text{box}}), \text{fill_isolated}(v_{\text{box}}, v_{\text{box}})\} \\
M^b &= \{\text{adj}(v_{\text{box}}, v_{\text{box}}), \text{does}(c_{\text{role}}, \text{mark}(v_{\text{box}})), \text{true}(\text{box}(v_{\text{box}}, c_{\text{state}}))\} \\
M^o &= \{\text{fill_next}(v_{\text{box}}, v_{\text{box}}), \text{fill_isolated}(v_{\text{box}}, v_{\text{box}}), \text{leftmost}(v_{\text{box}}, v_{\text{box}})\} \\
O^b &= \left\{ \begin{array}{l}
\langle (\times, \square \square \square \square \square \square, a_2), (\times, \square \square \square \square \square \square, a_3) \rangle \\
\langle (\times, \square \square \square \square \square \square, a_2), (\times, \square \square \square \square \square \square, a_1) \rangle \\
\langle (\times, \square \boxtimes \cdot \square \square \square \square, a_1), (\times, \square \boxtimes \cdot \square \square \square \square, a_5) \rangle \\
\langle (\times, \square \boxtimes \square \square \square \square, a_3), (\times, \square \boxtimes \square \square \square \square, a_4) \rangle \\
\langle (\times, \square \boxtimes \square \cdot \square \square \square \square, a_1), (\times, \square \boxtimes \square \cdot \square \square \square \square, a_3) \rangle \\
\langle (\times, \square \square \boxtimes \cdot \square \square \square \square, a_2), (\times, \square \square \boxtimes \cdot \square \square \square \square, a_1) \rangle \\
\langle (\times, \square \square \boxtimes \cdot \square \square \square \square, a_2), (\times, \square \square \boxtimes \cdot \square \square \square \square, a_6) \rangle \\
\langle (\times, \square \cdot \boxtimes \square \square \square \square, a_4), (\times, \square \cdot \boxtimes \square \square \square \square, a_1) \rangle \\
\langle (\times, \square \cdot \boxtimes \square \square \square \square, a_4), (\times, \square \cdot \boxtimes \square \square \square \square, a_5) \rangle \\
\langle (\times, \square \cdot \square \boxtimes \square \square \square \square, a_2), (\times, \square \cdot \square \boxtimes \square \square \square \square, a_6) \rangle
\end{array} \right.
\end{aligned}$$

O^b contains two main groups of examples, the first set are taken from playing using the strategy from the beginning of the game. However just using these examples can cause ILASP to learn a hypothesis that learns features of the states not the moves. With all of the examples, ILASP was able to learning the following hypothesis¹:

$$\begin{aligned}
\text{fill_1}(\text{Box2}) &\leftarrow \text{adj}(\text{Box1}, \text{Box2}), \text{does}(_, \text{mark}(\text{Box1})). \\
\text{fill_2}(\text{Box2}) &\leftarrow \text{adj}(\text{Box1}, \text{Box2}), \text{box}(\text{Box1}, \text{Player}), \\
&\quad \text{does}(\text{Player}, \text{mark}(\text{Box2})). \\
&\Leftarrow \text{fill_1}(\text{Box}), \text{leftmost}(\text{Box}, \text{Score}).[-\text{Score}@1, \text{Box}, \text{Score}] \\
&\Leftarrow \text{fill_2}(\text{Box}), \text{leftmost}(\text{Box}, \text{Score}).[-\text{Score}@2, \text{Box}, \text{Score}]
\end{aligned}$$

ILASP has created two predicates, fill_1 which describes a box that is adjacent to a box that is being marked, and fill_2 which describes a player marking a box next to one of their own boxes. Putting this together with the weak constraint we get: “at the highest priority gain the reward of how far left a box is if there is already an adjacent box that you have marked, otherwise choose the leftmost box with the most neighbours”. Because of the complexity of the game this is equivalent to the correct strategy presented in (Zhang and Thielscher, 2015).

8.2.3 FORWARD-THINKING STRATEGIES

Forward-thinking strategies involve the modal operator \bigcirc which means “will be true in the next state”. Further, this operator can be nested giving rise to strategies that reason several steps into the future, such as the *defence* strategy rule below. In their paper Zhang et al. state they are unable to represent strategies involving \bigcirc in ASP. In this section, I

experiment with using Deep Context Dependent LOAS tasks in order to learn strategies using \circ .

Experiment 8.6 (Defence). If the cross player were to play in the first box the dot player would lose using the *thoughtful* strategy, assuming the cross player is rational. This is illustrated in Figure 8.1 below, where the right branch represents the *thoughtful* strategy. It would make more sense for the dot player to block the cross player in order to try and force a draw.

The *defence* strategy means that if there is a move that the opponent could play to win, make that move instead. The learning task described below uses the *deep orderings* from Chapter 7 (Definition 7.2.3). Given the learning task $T_{def} = \langle B, S_M, E_O^+, O \rangle$ where

B is the same as in $T_{thoughtful}$,

$M = \langle \emptyset, \emptyset, M^o, \{-1, 1\}, 1 \rangle$

$M^o = \{goal(v_{role}, c_{reward}), role(v_{role}), control(v_{role})\}$

$O^d = \{ \langle (\cdot, \boxtimes \square \square \square, a_2), (\cdot, \boxtimes \square \square \square, a_3), 2 \rangle \}$

The relevant examples are then expanded into a game tree (see Figure 8.1). Along with generated examples, and the injected meta program (described in Section 7.4) ILASP was able to learning the following hypotheses¹: and at depth 2 you get

$\Leftarrow \text{not } goal(V0, 0), control(V0).[-1@1, 1, V0]$

When applied to minimax at depth two this corresponds to the strategy “make sure you have not lost on your next turn”.




Figure 8.1. Cross-Dot *defence* game tree generated from $\langle (\cdot, \boxtimes \square \square \square, a_2), (\cdot, \boxtimes \square \square \square, a_3), 2 \rangle$, the ordering example, with a depth of 2. Green and red represent a winning and losing state for \cdot , respectively.

8.3 COMPARISON

In this chapter we have seen that, using weak constraints, ILASP is able to learn the strategies presented by Zhang and Thielscher. The weak constraints have the advantage of using integer weights in order to penalise or reward moves, whilst the logical language described in Zhang and Thielscher (2015) suffer from verbosity, for example in the *leftmost* strategy. Using these weights to create a ‘heatmap’ of locations is an idea that is revisited in Experiment 10.4.

Zhang and Thielscher translate a subset of strategies written in their modal logic into ASP. Specifically, all strategies using the modal operator \bigcirc . They use traditional constraints in order to restrict the set of legal moves down to a few that are considered good with respect to the strategy. This means that if a particular strategy does not satisfy some of the properties shown in the paper (e. g. completeness and determinism) then one could find themselves in a situation where no moves can be recommended, despite legal moves existing. This problem can be quite easily rectified by saying that any legal move is part of the strategy, if rest of the strategy failed. In their modal logic this is written as:

$$\begin{aligned} fill_any^i &= \bigvee_{a^i \in A^i} (legal(a^i) \wedge does(a^i)) \\ strat^i &= r_1 \nabla r_2 \nabla \dots \nabla fill_any^i \end{aligned}$$

Which is to be read as “try all strategies r_i in decreasing order of priority and if they all fail then simply select a legal action”. However, when translated into ASP this becomes:

```

1 strat(A, T) :- r_1(A, T).
2 strat(A, T) :- r_2(A, T), not r_1(_, T).
3 %...
4 strat(A, T) :- fill_any(A, T), not r_n(_, T).
5 fill_any(A, T) :- legal(A, T).

```

Whereas when using weak constraints this process is completely omitted as if no weak constraints are violated (i. e. no penalties are gained) then every legal action will by default score 0 and any one can be selected.

Part III
EVALUATION

9 | LEARNING THE GAME RULES

Learning the rules of a game is the first step in coming up with a strategy that will allow you to defeat your opponents. The rules introduce you to the board, cards, and pieces of the game. They show you how everything can interact and what actions you have available to you on your turn. Most importantly the rules let you know how to win the game, and by putting all of these components together you can begin to formulate a plan that will allow you to have a better chance of winning.

In this chapter, I will outline the methodologies used in order to learn the rules of *Onitama*, *Five Field Kono* and *Cross-Dot*. Each experiment below lists the examples and mode bias provided to ILASP, in Appendix B you can find the exact options passed to ILASP and what features were used.

9.1 PROCESS

For each game a representation was created in GDL and translated into ASP. In the games looked at here the legal actions available can be calculated directly from the current state of the board. Thus the `next(·)` and `does(·,·)` predicates can be removed from the representation, significantly reducing the grounding of the meta representation.

Each game was translated with as few helper predicates as possible, for example `rows(·)` and `columns(·)` were added to *Onitama* in order to extract rows and columns from a pair of cells. All games had types added to the background knowledge in order to overcome free variables and bound cell ranges etc.. The background knowledge for each game can be found in Appendix A.

Whilst you could achieve the same hypothesis from observing many games being played, examples have been provided by an ‘oracle’. This way mimics the way a parent may teach their child a game, additionally it means that examples can be chosen to specifically show a new rule or edge case. In order for ILASP to learn an accurate representation of the rules ‘nonsense’ moves must be explicitly ruled out as illegal moves. For example, you are unable to move from a empty square to another empty square.

9.2 LEARNING

Experiment 9.1 (Onitama Rules). The task takes the background knowledge from Logic Program A.1 without the action choice rule and the `next(·)` predicate. The mode bias has been built such that instead of creating one rule that defines the legal moves it learns sub predicates each with a restricted search space. This reduces the maximum number of variables that need to be used, and the number of literals that a rule contains.

The mode declaration of this task added some additional rules to the background knowledge as, at the time, ILASP did not support arithmetic expressions in the bias constraints.

$$\text{delta}(D) \leftarrow \text{card}(_, (D, _)).$$

$$\text{delta}(D) \leftarrow \text{card}(_, (_, D)).$$

$$\begin{aligned} \mathit{math}(B - C, B, C, 1) &\leftarrow \mathit{cell}(B - C, B), \mathit{delta}(C). \\ \mathit{math}(B + C, B, C, -1) &\leftarrow \mathit{cell}(B + C, B), \mathit{delta}(C). \end{aligned}$$

The mode declaration used is as follows, the bias constraints used can be found in Section B.1.1, along with the positive and negative examples:

$$M_h = \begin{cases} \mathit{legal}(\mathbf{v}_{role}, \mathit{move}(\mathbf{v}_{pair}, \mathbf{v}_{card})) \\ \mathit{pred}_1(\mathbf{v}_{role}, (\mathbf{v}_{cell}, \mathbf{v}_{cell})) \\ \mathit{pred}_2(\mathbf{v}_{card}, \mathbf{v}_{pair}, \mathbf{v}_{dir}) \\ \mathit{pred}_3(\mathbf{v}_{role}, \mathbf{v}_{cell}) \\ \mathit{pred}_4(\mathbf{v}_{role}, \mathbf{v}_{cell}) \\ \mathit{pred}_5((\mathit{cell}(\mathbf{v}_{index}, \mathbf{v}_{index}), \mathit{cell}(\mathbf{v}_{index}, \mathbf{v}_{index})), (\mathbf{v}_{index}, \mathbf{v}_{index})) \\ \mathit{pred}_6(\mathbf{v}_{delta}, \mathbf{v}_{pair}, \mathbf{v}_{dir}) \\ \mathit{pred}_7(\mathbf{v}_{delta}, \mathbf{v}_{pair}, \mathbf{v}_{dir}) \end{cases}$$

$$M_b = \begin{cases} \mathit{control}(\mathbf{v}_{role}) & \mathit{pred}_1(\mathbf{v}_{role}, \mathbf{v}_{pair}) \\ \mathit{location}(\mathit{pawn}(\mathbf{v}_{rank}, \mathbf{v}_{role}), \mathbf{v}_{cell}) & \mathit{pred}_2(\mathbf{v}_{card}, \mathbf{v}_{pair}, \mathbf{v}_{dir}) \\ \mathit{in_hand}(\mathbf{v}_{role}, \mathbf{v}_{card}) & \mathit{pred}_3(\mathbf{v}_{role}, \mathbf{v}_{cell}) \\ \mathit{card}(\mathbf{v}_{card}, (\mathbf{v}_{delta}, \mathbf{v}_{delta})) & \mathit{pred}_4(\mathbf{v}_{role}, \mathbf{v}_{cell}) \\ \mathit{dir}(\mathbf{v}_{role}, \mathbf{v}_{dir}) & \mathit{pred}_4(\mathbf{v}_{role}, \mathbf{v}_{cell}) \\ \mathit{cell}(\mathbf{v}_{index}, \mathbf{v}_{index}) & \mathit{pred}_5(\mathbf{v}_{pair}, (\mathbf{v}_{index}, \mathbf{v}_{index})) \\ \mathit{math}(\mathbf{v}_{index}, \mathbf{v}_{index}, \mathbf{v}_{delta}, \mathbf{v}_{dir}) & \mathit{pred}_6(\mathbf{v}_{delta}, \mathbf{v}_{pair}, \mathbf{v}_{dir}) \\ & \mathit{pred}_7(\mathbf{v}_{delta}, \mathbf{v}_{pair}, \mathbf{v}_{dir}) \end{cases}$$

Remark. The hypothesis that is learnt can be reduced to one rule which has a smaller grounding by replacing the occurrences of $\mathit{pred}_X(\cdot)$ by their bodies.

$$\mathit{pred}_7(DR, (From, To), Dir) \leftarrow \mathit{math}(Row2, Row1, DR, Dir), \quad (9.1)$$

$$\mathit{rows}(From, To, (Row1, Row2)).$$

$$\mathit{pred}_6(DC, (From, To), Dir) \leftarrow \mathit{math}(Col1, Col2, DC, Dir), \quad (9.2)$$

$$\mathit{columns}(From, To, (Col1, Col2)).$$

$$\mathit{pred}_2(Card, Coords, Dir) \leftarrow \mathit{card}(Card, (DC, DR)), \quad (9.3)$$

$$\mathit{pred}_6(DC, Coords, Dir),$$

$$\mathit{pred}_7(DR, Coords, Dir).$$

$$\mathit{location}(Role, Cell) \leftarrow \mathit{location}(\mathit{pawn}(_, Role), Cell). \quad (9.4)$$

$$\mathit{pred}_1(Role, (From, To)) \leftarrow \mathit{location}(\mathit{pawn}(_, Role), From), \quad (9.5)$$

$$\mathit{not location}(Role, To),$$

$$\mathit{cell}(To), \mathit{control}(Role).$$

$$\mathit{legal}(Role, \mathit{move}(Coords, Card)) \leftarrow \mathit{in_hand}(Role, Card), \mathit{dir}(Role, Dir), \quad (9.6)$$

$$\mathit{pred}_1(Role, Coords),$$

$$\mathit{pred}_2(Card, Coords, Dir).$$

The hypothesis learnt can be described in English in the following way:

- (9.1) $pred_7$ is true when there is a correct translation to the rows of the coordinates;
- (9.2) $pred_6$ is true when there is a correct translation to the columns of the coordinates;
- (9.3) $pred_2$ is true when there is a card that represents the translation of the pawn;
- (9.4) $location$ is the projection of the $location$ predicate, removing the pawn's rank;
- (9.5) $pred_1$ is true when the starting location contains *your* pawn and the ending location does not;
- (9.6) $legal$ is true when a card in your hand provides the translation of the pawn (given your direction) and that you have moved *your* pawn to a valid location.

This hypothesis can be compressed into the following form:

$$\begin{aligned}
 legal(Role, move((From, To), Card)) \leftarrow & in_hand(Role, Card), dir(Role, Dir), \\
 & location(pawn(_, Role), From), \\
 & not\ location(pawn(_, Role), To), \\
 & cell(To), control(Role), \\
 & card(Card, (DC, DR)), \\
 & math(Col1, Col2, DC, Dir), \\
 & columns(From, To, (Col1, Col2)), \\
 & math(Row2, Row1, DR, Dir), \\
 & rows(From, To, (Row1, Row2)).
 \end{aligned}$$

Experiment 9.2 (Five Field Kono Rules).

$$\begin{aligned}
 adj(cell(Row1, Col1), cell(Row2, Col2)) \leftarrow & \hspace{15em} (9.7) \\
 & cell(Row1, Col1), cell(Row2, Col2), \\
 & Row2 == Row1 + 1, Col2 == Col1 + 1.
 \end{aligned}$$

$$\begin{aligned}
 adj(cell(Row1, Col1), cell(Row2, Col2)) \leftarrow & \hspace{15em} (9.8) \\
 & cell(Row1, Col1), cell(Row2, Col2), \\
 & Col1 == Col2 + 1, Row1 == Row2 - 1.
 \end{aligned}$$

$$\begin{aligned}
 adj(cell(Row1, Col1), cell(Row2, Col2)) \leftarrow & \hspace{15em} (9.9) \\
 & cell(Row2, Col2), cell(Row1, Col1), \\
 & Col2 == Col1 + 1, Row2 == Row1 - 1.
 \end{aligned}$$

$$\begin{aligned}
 adj(cell(Row1, Col1), cell(Row2, Col2)) \leftarrow & \hspace{15em} (9.10) \\
 & cell(Row1, Col1), cell(Row2, Col2), \\
 & Col1 == Col2 + 1, Row2 == Row1 - 1.
 \end{aligned}$$

$$\begin{aligned}
 legal(Role, move(cell(Row1, Col1), cell(Row2, Col2))) \leftarrow & \hspace{15em} (9.11) \\
 & adj(cell(Row1, Col1), cell(Row2, Col2)), control(Role) \\
 & state(Row1, Col1, Role), state(Row2, Col2, e).
 \end{aligned}$$

The hypothesis learnt can be described in English in the following way:

- (9.7)–(9.10) *adj* represents that immediate diagonals are adjacent to the current location (this is more obvious after some arithmetic manipulation);
- (9.11) A move is *legal* when the destination is empty and adjacent to the starting location, which has one of *your* pieces.

This hypothesis can be compressed into the following form:

$$\begin{aligned}
 &adj(cell(Row, Col), cell(Row + 1, Col + 1)) \leftarrow cell(Row, Col), cell(Row + 1, Col + 1). \\
 &adj(cell(Row, Col), cell(Row + 1, Col - 1)) \leftarrow cell(Row, Col), cell(Row + 1, Col - 1). \\
 &adj(cell(Row, Col), cell(Row - 1, Col + 1)) \leftarrow cell(Row, Col), cell(Row - 1, Col + 1). \\
 &adj(cell(Row, Col), cell(Row - 1, Col - 1)) \leftarrow cell(Row, Col), cell(Row - 1, Col - 1). \\
 &legal(Role, move(cell(Row1, Col1), cell(Row2, Col2))) \leftarrow \\
 &\quad adj(cell(Row1, Col1), cell(Row2, Col2)), control(Role) \\
 &\quad state(Row1, Col1, Role), state(Row2, Col2, e).
 \end{aligned}$$

Experiment 9.3 (Cross-Dot Rules). As the search space of the Cross-Dot game is much smaller I also added the head for learning the *noop* action, i. e. does not move.

$$legal(Role, mark(Box)) \leftarrow box(Box, b), control(Role). \quad (9.12)$$

$$legal(Role, noop) \leftarrow not\ control(Role), role(Role). \quad (9.13)$$

10 | LEARNING AND EXPRESSING STRATEGIES

Capturing, advancing pawns and forking¹ are all tactics that a player may enforce within a game in order to try and gain the advantage over their opponent. These tactics are also calculable by only looking at the current board position, and they form the building blocks of winning strategies. A lot of powerful strategies in games come from composing smaller strategies that solve particular sub-goals within the game. For example, during the opening you may be looking to capture student pawns in order to restrict the opponents movement, and in the end game look at advancing your master pawn across the board to your opponent's temple square. Later these simple strategies are used in order to create something more complex. As well as being combined they can be reasoned about at different depths of the tree and, in a similar fashion to Chapter 8, we will explore applying deep orderings to player's strategies to learn stronger strategies still.

10.1 IMMEDIATE STRATEGIES

10.1.1 WINNING

Whilst it seems obvious, achieving the winning conditions must be encoded within the strategy. The idea of winning, losing or drawing is also the basis of algorithms such as Monte Carlo tree search, which when they arrive at a winning node propagate a positive value up the game tree, a negative value if it is a losing node or zero for a drawing position. This simple heuristic can be a good indicator of a board position if many games are simulated by randomly walking the tree.

In many games, including *Onitama*, there are multiple ways to win. Recall from the rules (Section 2.1) that in *Onitama* you can either capture your opponent's master, or navigate your own master to your opponent's temple. It may be of interest to learn which of these strategies is preferred and if one method is dominant over another. However, learning the dominant/preferred strategy will require a more complex task, as the best option can often depend on both the cards, and the strategy of your opponent.

Experiment 10.1 (Try to Win). For this experiment we used ILASP2i to perform the learning task as there is no noise. The task given to ILASP was $T = \langle B, S_M, E^+, \emptyset, O^b, \emptyset \rangle$ where,

B is the logic program from Logic Program A.1

$M = \langle \emptyset, \emptyset, \{goal(v_{role}, c_{score}), control(v_{role})\}, \{-1, 1\}, 1 \rangle$

E^+ is the set of examples, e_i , with the contexts, C_i , below and empty inclusions and exclusions

$O_b = \{ \langle e_1, e_2 \rangle, \langle e_1, e_3 \rangle \}$

Note. For this task we can take advantage of the built in GDL specified predicate *goal*.

¹ Threatening to capture two pawns with one

```

% Context: C1
location(pawn(student,red),cell(5,5)).
location(pawn(student,blue),cell(4,3)).
location(pawn(student,red),cell(4,5)).
location(pawn(master,red),cell(2,1)).
location(pawn(student,red),cell(2,2)).
location(pawn(student,red),cell(2,4)).
location(pawn(student,blue),cell(1,1)).
location(pawn(student,blue),cell(1,5)).
control(blue).
in_play(monkey).
in_play(ox).
in_play(sable).
in_hand(red,monkey).
in_hand(blue,ox).

% Context: C2
location(pawn(student,blue),cell(4,3)).
location(pawn(student,red),cell(4,4)).
location(pawn(student,red),cell(4,5)).
location(pawn(student,red),cell(3,5)).
location(pawn(master,red),cell(2,1)).
location(pawn(student,red),cell(2,2)).
location(pawn(master,blue),cell(2,4)).
location(pawn(student,blue),cell(1,1)).

location(pawn(student,blue),cell(1,5)).
control(blue).
in_play(monkey).
in_play(ox).
in_play(sable).
in_hand(red,monkey).
in_hand(blue,ox).

% Context: C3
location(pawn(student,red),cell(5,5)).
location(pawn(student,blue),cell(4,3)).
location(pawn(student,red),cell(3,4)).
location(pawn(student,red),cell(3,5)).
location(pawn(master,red),cell(2,1)).
location(pawn(student,red),cell(2,2)).
location(pawn(master,blue),cell(2,4)).
location(pawn(student,blue),cell(1,1)).
location(pawn(student,blue),cell(1,5)).
control(blue).
in_play(monkey).
in_play(ox).
in_play(sable).
in_hand(red,monkey).
in_hand(blue,ox).

```

The hypothesis learnt by ILASP is:

$$\Leftarrow \text{goal}(\text{Role}, 100), \text{not control}(\text{Role}).[-1@1, \text{Role}]$$

At first this result may look odd as it says “a reward is given when the player *not* in control reaches the goal of 100”, however when looking at the states it is because the player who has just moved is no longer having their turn.

10.1.2 CAPTURING PIECE

A common strategy among *Onitama* beginners is to capture a piece when they can, and whilst this strategy can sometimes be short-sighted (as seen in Chapter 7) it is important to be able to learn the strategy and represent it in the background knowledge.

Intuitively to humans a hypothesis of ‘capturing’ could be thought of as “moving into a square with containing opponent pawn”, This contains the notions of ‘moving’, ‘location’ and ‘opponent pawns’. Therefore predicates relating to these concepts would be useful to add to the mode declarations. With this in mind we present the following experiment.

Experiment 10.2 (Capture). Using the digital versions of the games many examples of capturing a pawn if at all possible (and randomly choosing in the case of multiple possibilities) were collected. For this experiment we used ILASP 2i in order to do the learning as there is no noise. The task given to ILASP was $T = \langle B, S_M, E^+, \emptyset, O^b, \emptyset \rangle$ where,

B is the logic program from Logic Program A.1

$$M = \left\langle \emptyset, \emptyset, \left\{ \text{location}(\text{pawn}(\mathbf{v}_{\text{rank}}, \mathbf{v}_{\text{role}}), \mathbf{v}_{\text{cell}}), \right. \right. \\ \left. \left. \text{control}(\mathbf{v}_{\text{role}}), \text{does}(\mathbf{v}_{\text{role}}, \text{move}(\mathbf{v}_{\text{cell}}, \mathbf{v}_{\text{cell}})) \right\}, \{-1, 1\}, 1 \right\rangle$$

E^+ is the set of examples, e_i , with the contexts, C_i , below and empty inclusions and exclusions

$$O_b = \{\langle e_1, e_2 \rangle, \langle e_3, e_4 \rangle\}$$

```

% Context: C1
location(pawn(student,red),cell(5,1)).
location(pawn(student,red),cell(5,3)).
location(pawn(master,red),cell(4,2)).
location(pawn(student,blue),cell(3,2)).
location(pawn(student,blue),cell(3,5)).
location(pawn(student,blue),cell(2,2)).
location(pawn(master,blue),cell(1,3)).
location(pawn(student,blue),cell(1,5)).
control(red).
in_play(sable).
in_play(ox).
in_play(monkey).
in_hand(red,sable).
in_hand(blue,ox).

% Context: C2
location(pawn(student,red),cell(5,1)).
location(pawn(student,red),cell(5,3)).
location(pawn(master,red),cell(4,2)).
location(pawn(student,blue),cell(3,2)).
location(pawn(student,blue),cell(3,3)).
location(pawn(student,red),cell(3,5)).
location(pawn(student,blue),cell(2,2)).
location(pawn(master,blue),cell(1,3)).
location(pawn(student,blue),cell(1,5)).
control(red).
in_play(sable).
in_play(ox).
in_play(monkey).
in_hand(red,sable).
in_hand(blue,ox).

% Context: C3
location(pawn(student,red),cell(5,1)).
location(pawn(student,red),cell(5,5)).
location(pawn(student,blue),cell(4,2)).
location(pawn(student,blue),cell(3,5)).
location(pawn(student,blue),cell(2,2)).
location(pawn(master,blue),cell(1,3)).
location(pawn(student,blue),cell(1,5)).
control(red).
in_play(monkey).
in_play(sable).
in_play(ox).
in_hand(red,monkey).
in_hand(blue,sable).

% Context: C4
location(pawn(student,red),cell(5,1)).
location(pawn(student,red),cell(5,5)).
location(pawn(master,red),cell(4,2)).
location(pawn(student,blue),cell(3,2)).
location(pawn(student,blue),cell(3,5)).
location(pawn(student,blue),cell(2,2)).
location(pawn(master,blue),cell(1,4)).
location(pawn(student,blue),cell(1,5)).
control(red).
in_play(monkey).
in_play(sable).
in_play(ox).
in_hand(red,monkey).
in_hand(blue,sable).

```

The hypothesis learnt by ILASP is:

$$\rightsquigarrow \text{location}(\text{pawn}(\text{Rank}, \text{Player}), \text{Cell}).[1@1, \text{Rank}, \text{Player}, \text{Cell}]$$

This hypothesis does not contain the notion of movement and has a penalty instead of a reward! On closer inspection it can be seen that this weak constraint is attempting to minimise the number of unique pawns on the board, as you are unable to capture your own pieces this translates to capturing the opponent.

This experiment is not particularly interesting, it has a single predicate in the body and a single priority. Also consider the possibility of having the choice to capture a master piece or a student piece, the hypothesis above would weight them the same. In order to express this preference we can change the mode declaration by replacing v_{rank} with c_{rank} , and adding the following ordering example: $O^{b'} = O^b \cup \{\langle e_5, e_6 \rangle\}$

```

% Context: C5
location(pawn(student,red),cell(3,4)).
location(pawn(student,red),cell(5,5)).
location(pawn(student,blue),cell(2,5)).
location(pawn(student,blue),cell(3,3)).

location(pawn(student,blue),cell(2,1)).
location(pawn(master,blue),cell(1,3)).
control(red).
in_play(boar).
in_play(sheep).

```

```

in_play(eel).
in_hand(red,boar).
in_hand(blue,eel).

% Context: C6
location(pawn(student,red),cell(5,5)).
location(pawn(master,red),cell(3,3)).
location(pawn(student,blue),cell(3,4)).
location(pawn(student,blue),cell(3,2)).

location(pawn(student,blue),cell(2,1)).
location(pawn(master,blue),cell(1,3)).
control(red).
in_play(boar).
in_play(sheep).
in_play(eel).
in_hand(red,boar).
in_hand(blue,eel).

```

Experiment 10.3 (Capture: Master or Student?). Given the new task $T = \langle B, S_{M'}, E^+, \emptyset, O^{b'}, \emptyset \rangle$, based on the task in Experiment 10.2 above, where

$$M' = \langle \emptyset, \emptyset, \{location(pawn(c_{rank}, v_{role}), v_{cell}), control(v_{role})\}, \{-1, 1, -2, 2\}, 1 \rangle$$

notice how extra weights have been added, alternatively we could have chosen to allow two priority levels. They have been added as one might expect at a high priority the strategy describes the master and at a lower priority, the students.

ILASP learns the hypothesis:

$$\begin{aligned} &\Leftarrow location(pawn(master, Player), Cell1), \\ &location(pawn(student, Player), Cell2).[1@1, Player, Cell1, Cell2] \end{aligned}$$

The immediate observation is that ILASP has been able to cover the ordering examples using a single rule. This hypothesis means the following: *for each player penalise each of their students when their master is also on the board.* In a scenario when you cannot capture the master, then the masters will still be on the board next turn and so this strategy is equivalent to before. However, when you can capture the opponent's master it will no longer be on the board and so only the winning player's students are counted.

10.1.3 SPACE ADVANTAGE

Experiment 10.4 (Controlling Regions of the Board). A common tactic in (certain) abstract strategy games is to control as much space a possible (or certain spaces of the board). In chess, controlling the centre is often touted as good practice. It can free up paths for your bishops, and allow knights to control more of the important squares. Figure 10.2 shows two illustrations, on the left is a board state with blue's control shaded, and on the right is the same for red. We can see that blue is making it very difficult for red to advance. Red is only controlling their half of the board, meanwhile blue is moving pawns up across the board, creating lots of space to move into.

This experiment used a large set of noisy examples collected from full games but was unable to learn a suitable strategy, due to the noise in the examples and mode declarations not providing enough information.

Experiment 10.5. In this experiment the examples are generating from a player correcting the decisions of an AI. These examples were collected using the *training-mode* of digital *Onitama*, see Section 6.1.1 for details. The task was run using ILASP21 because we assume there is no noise as the player is explicitly correcting a move.

The mode declaration for the task is:

$$M = \left\langle \emptyset, \emptyset, \left\{ \begin{array}{l} valid_translation(_, v_{cell}, _, v_{role}), \\ control(v_{role}), opponent(v_{role}, v_{role}), \\ location(pawn(c_{rank}, v_{role}), v_{cell}) \end{array} \right\}, \{-1, 1\}, 1 \right\rangle$$

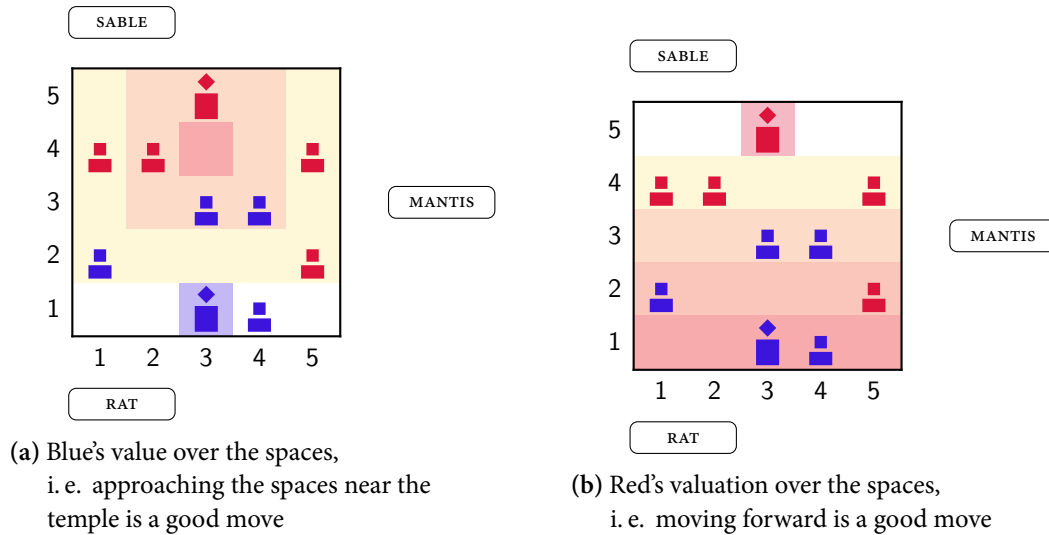


Figure 10.1. Illustrations of how valuing different areas of the board affects strategy. Red represents higher valuation and white is the lowest.




Figure 10.2. Illustrations of how controlling space on a board can give be advantageous. Grey areas represent spaces that the cards allow each player to move to. The strategies learnt in Experiment 10.5 and Experiment 10.7 both have this concept in their mode declarations (*valid_translation(·,·,·,·)*)

The contexts used to create the example are as follows

The learnt hypothesis learnt is:

- ↔ not *valid_translation*(_, To, _, Player), *location*(*pawn*(*master*, Player), To).[1@1, To, Player]
- ↔ not *valid_translation*(_, To, _, Player), *location*(*pawn*(*master*, Enemy), To), *control*(Enemy), *opponent*(Player, Enemy).[1@1, To, Player, Enemy]

This hypothesis says that you get penalised for not being able to defend your master pawn. The second rule says that if one is not able to threaten the opponent's master then a penalty is received. Note that the *control*(·) here is after the move has been made and so this refers to the opponent's master.

10.2 COMPLEX STRATEGIES

Experiment 10.6 (Planning a Path for the Master). Many chess engines (e.g. *Stockfish*) use endgame tables (e.g. Nalimov endgame tablebases²) towards the end of the game, normally when there are fewer than six pieces on the board, as these small portions of the sub-trees have been solved. A wrong move can quickly turn a winning position into a drawing position (or worse a losing position). In *Onitama* because of the fixed moves you can find yourself in a situation where providing you notice it early enough you can force a win by navigating the master correctly to the temple square, this occurs more often in the three card variant. Learning this from examples is something that could be done using enough forethought, and is a good test of how powerful the deep orderings are, and what they can achieve. Figure 10.3 demonstrates a series of moves for the blue master \square which red is unable to defend against. This experiment is an extension of the “try to win” strategy (Experiment 10.1) involving deep orderings.

While in theory reasoning 5 moves into the future is possible the $ILLP_{LOAS}^{context}$ with \mathcal{Q} that is generated contains 1.15×10^7 brave orderings for the leaves. This is due to the branching factor of the game. As the task has to be fully described before being solved by ILASP there currently is no way of partially evaluating the tree. In our future work (Section 11.2) we discuss one method of pruning the tree by evaluating ‘quiescent nodes’ (see Section 3.5.3). For this reason the task below uses a deep ordering based on Figure 10.3b which only has 24 649 brave orderings, meaning an average branching factor of 5.39 moves.

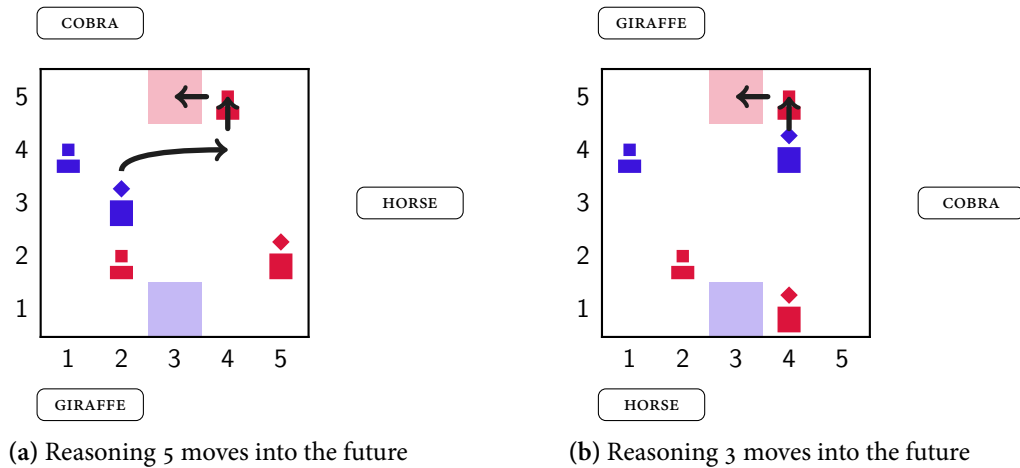


Figure 10.3. Guaranteed win by following the moves shown

We use ILASP 3 with injection in order to encode the following task: $T = \langle B, S_M, E, O, \triangleleft \rangle$ where,

$$M = \langle \emptyset, \emptyset, \{goal(v_{role}, v_{reward}), control(v_{role})\}, \{v_{reward}\}, 1 \rangle \quad (10.1)$$

$$E^+ = \{e_1, e_2\} \quad (10.2)$$

$$O^d = \{\langle e_1, e_2, 3, \triangleleft \rangle\} \quad (10.3)$$

² <http://www.k4it.de/index.php?topic=egtb&lang=en>

```

% Context: C1
location(pawn(student, red), cell(2,2)).
location(pawn(student, red), cell(5,4)).
location(pawn(master, red), cell(2,5)).
location(pawn(student, blue), cell(4,1)).
location(pawn(master, blue), cell(3,2)).
control(blue).
in_play(giraffe).
in_play(cobra).
in_play(horse).
in_hand(red, cobra).
in_hand(blue, giraffe).
:- not does(blue, move(cell(3,2), cell(4,4),
  ↪ giraffe)).

% Context: C2
location(pawn(student, red), cell(2,2)).
location(pawn(student, red), cell(5,4)).
location(pawn(master, red), cell(2,5)).
location(pawn(student, blue), cell(4,1)).
location(pawn(master, blue), cell(3,2)).
control(blue).
in_play(giraffe).
in_play(cobra).
in_play(horse).
in_hand(red, cobra).
in_hand(blue, giraffe).
:- not does(blue, move(cell(4,1), cell(3,1),
  ↪ giraffe)).

```

This lead to ILASP generating the hypothesis:

$$\leftarrow \text{goal}(\text{Player}, \text{Reward}).[\neg \text{Reward}@1, \text{Player}, \text{Reward}]$$

Despite this being a similar hypothesis to the one originally learnt in Experiment 10.1, given the two examples e_1 and e_2 a normal context-dependent LOAS task would have returned this as *unsatisfiable* because of the forward thinking. Next we describe a more complex use case for deep orderings.

Experiment 10.7 (Defending Pawns). In *Onitama* there is very little space to manoeuvre, however it can still be the case that a pawn is left stranded on the board. Take Figure 10.2 as an example, the red student ♙ on (2,5) is left without defence (note the lack of shading from red’s perspective).

Unlike chess where the pieces have fixed movement, in a game of *Onitama* the cards will rotate. Therefore to think about defending your pieces you need to calculate both where your pieces will be and what moves they will have. In order to achieve this as part of an ILASP learning task, you can encode the preference as a deep ordering looking at depth 2. As if the other player takes your pawn then you have a chance to capture back (giving a better valuation than just losing a pawn).

This task was created using *training-mode*, with one alteration: a brave to a deep ordering, of depth 1 i. e. considering the opponent’s next turn. The mode declaration of the tasks is the same as in Experiment 10.5, the main difference is in this experiment we are looking to defend our own pawns instead of avoid capture completely. For ease, M is restated below:

$$M = \left\langle \emptyset, \emptyset, \left\{ \begin{array}{l} \text{valid_translation}(_, \mathbf{v}_{\text{cell}}, _, \mathbf{v}_{\text{role}}), \\ \text{control}(\mathbf{v}_{\text{role}}), \text{opponent}(\mathbf{v}_{\text{role}}, \mathbf{v}_{\text{role}}), \\ \text{location}(\text{pawn}(\mathbf{c}_{\text{rank}}, \mathbf{v}_{\text{role}}), \mathbf{v}_{\text{cell}}) \end{array} \right\}, \{-1, 1\}, 1 \right\rangle$$

The contexts for the partial interpretations using the deep ordering are below, all other contexts used in this task can be found in Section B.1.2.

```

% Context: C7
location(pawn(master, red), cell(5,3)).
location(pawn(student, red), cell(5,4)).
location(pawn(student, red), cell(4,1)).
location(pawn(student, red), cell(4,4)).

location(pawn(student, red), cell(4,2)).
location(pawn(student, blue), cell(2,1)).
location(pawn(student, blue), cell(2,4)).
location(pawn(student, blue), cell(1,1)).
location(pawn(master, blue), cell(1,3)).

```

```

location(pawn(student,blue),cell(1,4)).
control(red).
in_play(rat).
in_play(crab).
in_play(mantis).
in_hand(red,mantis).
in_hand(blue,crab).
:- not does(red, move(cell(4,2), cell(3,1),
  ↪ mantis)).

% Context: C8
location(pawn(master,red),cell(5,3)).
location(pawn(student,red),cell(4,1)).
location(pawn(student,red),cell(4,2)).
location(pawn(student,red),cell(4,4)).

location(pawn(student,red),cell(5,4)).
location(pawn(student,blue),cell(2,1)).
location(pawn(student,blue),cell(2,4)).
location(pawn(student,blue),cell(1,1)).
location(pawn(master,blue),cell(1,3)).
location(pawn(student,blue),cell(1,4)).
control(red).
in_play(rat).
in_play(crab).
in_play(mantis).
in_hand(red,mantis).
in_hand(blue,crab).
:- not does(red, move(cell(5,4), cell(4,5),
  ↪ mantis)).

```

The brave ordering examples, O^b , are:

$$\langle e_1, e_2 \rangle$$

$$\langle e_3, e_4 \rangle$$

$$\langle e_5, e_6 \rangle$$

and the deep ordering example: $\langle e_7, e_8, 1, \triangleleft \rangle$, a translation of this task can be found in Logic Program B.2.

ILASP learnt the following hypothesis:

$$\Leftarrow \text{not } \textit{valid_translation}(_, \textit{To}, _, \textit{Player}), \textit{location}(\textit{pawn}(\textit{student}, \textit{Player}), \textit{To}).[1@1, \textit{Player}, \textit{To}]$$

This hypothesis translates into English as: “minimise the number of students one has that you cannot move to” which means that if your piece were to be captured, you could move to that square and capture back.

10.3 TOURNAMENTS

In order to test whether or not the hypotheses we have learnt in this chapter are usable in a game and able to perform well we pitted strategies against each other and against a random (no preference) and hand coded strategy. Each experiment involves 100 simulated games of two strategies, the random nature of the game means that each AI gets to start roughly evenly. We then perform a binomial, two-tailed hypothesis test to see whether or not the strategies are equal or if there is a bias towards one. Table 10.1 shows the p-values of the various tests, as you can see most of the experiments show that we should reject the hypothesis that the strategies are equal.

10.4 SUMMARY

In this section we give a high level overview of the learning tasks in the project and comment on what our results show. Table 10.2 shows the hypotheses learnt across the various chapters of this report.

STRATEGY 1 VS. STRATEGY 2	P-VALUE	RESULT
Win or Capture vs. Random	$1.9114 \times 10^{-15} \star$	Win or Capture is clearly a better strategy than random.
Win or Capture vs. Avoid Capture	0.67181	These strategy's are roughly equivalent, we should not reject the null hypothesis
Win or Capture vs. Defend Pieces	0.01203 \star	The Defend Pieces strategy is considered the better strategy

Table 10.1. Hypothesis tests between various strategies, \star represents a significant value

EXPERIMENT	HYPOTHESIS	NOTES
Experiment 10.1	\leftrightarrow <i>goal(</i> Player, 100).[-1@1, Player]	Useful later when combined with other strategies
Experiment 10.2	\leftrightarrow <i>location</i> (<i>pawn</i> (<i>Rank</i> , <i>Player</i>), <i>Cell</i>). [1@1, <i>Rank</i> , <i>Player</i> , <i>Cell</i>]	Instance of the penalty being used to minimise a board feature
Experiment 10.3	\leftrightarrow <i>location</i> (<i>pawn</i> (<i>master</i> , <i>Player</i>), <i>Cell</i> 1), <i>location</i> (<i>pawn</i> (<i>student</i> , <i>Player</i>), <i>Cell</i> 2).[1@1, <i>Player</i> , <i>Cell</i> 1, <i>Cell</i> 2]	
Experiment 10.4	—	This experiment was unsuccessful based on numerous example games, this could be due to noise
Experiment 10.5	\leftrightarrow <i>not valid_translation</i> (_, <i>To</i> , _ <i>Player</i>), <i>location</i> (<i>pawn</i> (<i>master</i> , <i>Player</i>), <i>To</i>).[1@1, <i>To</i> , <i>Player</i>] \leftrightarrow <i>not valid_translation</i> (_, <i>To</i> , _ <i>Player</i>), <i>location</i> (<i>pawn</i> (<i>master</i> , <i>Enemy</i>), <i>To</i>), <i>control</i> (<i>Enemy</i>), <i>opponent</i> (<i>Player</i> , <i>Enemy</i>).[1@1, <i>To</i> , <i>Player</i> , <i>Enemy</i>]	
Experiment 10.6	\leftrightarrow <i>goal</i> (<i>Player</i> , <i>Reward</i>).[- <i>Reward</i> @1, <i>Player</i> , <i>Reward</i>]	This was a similar task to Experiment 10.1, but using deep orderings
Experiment 10.7	\leftrightarrow <i>not valid_translation</i> (_, <i>To</i> , _ <i>Player</i>), <i>location</i> (<i>pawn</i> (<i>student</i> , <i>Player</i>), <i>To</i>).[1@1, <i>Player</i> , <i>To</i>]	
Experiment 8.1	\leftrightarrow <i>does</i> (<i>x</i> , <i>mark</i> (<i>Box</i> 1)), <i>adj</i> (<i>Box</i> 1, <i>Box</i> 2), <i>box</i> (<i>Box</i> 2, <i>x</i>).[-1@1, <i>Box</i> 1, <i>Box</i> 2]	
Experiment 8.2	\leftrightarrow <i>does</i> (<i>x</i> , <i>mark</i> (<i>Box</i> 1)), <i>adj</i> (<i>Box</i> 1, <i>Box</i> 2), <i>true</i> (<i>box</i> (<i>Box</i> 2, <i>b</i>)).[-1@1, <i>Box</i> 1, <i>Box</i> 2]	
Experiment 8.3	\leftrightarrow <i>does</i> (<i>x</i> , <i>mark</i> (<i>Box</i> 1)), <i>adj</i> (<i>Box</i> 1, <i>Box</i> 2), <i>true</i> (<i>box</i> (<i>Box</i> 2, <i>x</i>)).[-1@2, <i>Box</i> 1, <i>Box</i> 2] \leftrightarrow <i>does</i> (<i>x</i> , <i>mark</i> (<i>Box</i> 1)), <i>adj</i> (<i>Box</i> 1, <i>Box</i> 2), <i>true</i> (<i>box</i> (<i>Box</i> 2, <i>b</i>)).[-1@1, <i>Box</i> 1, <i>Box</i> 2]	
Experiment 8.4	\leftrightarrow <i>does</i> (<i>x</i> , <i>mark</i> (<i>Box</i> 1)).[<i>Box</i> 1@1, <i>Box</i> 1]	
Experiment 8.5	(see page 77)	
Experiment 8.6	\leftrightarrow <i>not goal</i> (<i>V</i> 0, 0), <i>control</i> (<i>V</i> 0).[-1@1, 1, <i>V</i> 0]	

Table 10.2. Summary of experiments performed throughout the report

In this chapter we have presented a series of experiments created from constructed situations in order to evaluate the feasibility of using weak constraints to learn strategies in games. We started with single predicate weak constraints to assess the winning conditions, afterwards we moved onto looking at basic tactics such as capturing, and combining the two strategies in order to create a more robust strategy. This strategy has shown to be effective in practice, but not as strong as some of the later positional and forward-thinking strategies, just as one might expect.

We then looked into strategies that assign value to positions on the board, utilising weights of weak constraints as a simple way of summing over the pawns on the board.

Finally, we look into applying the deep orderings in order to calculate positions from the given states, just as players do when playing games. We have shown that for smaller depths these tasks work well, and are superior in practice to the immediate strategies.

11 | CONCLUSION

Learning from weak constraints means we can express strategies using English once they have been learnt. This allows us to not only understand in more depth what has been learnt but also it allows us to learn more about a game from the strategies of better players. Weak constraints allow us to build upon existing strategies by assigning a priority level to each rule. Using *deep orderings* we are now able to learn about a state based one preferences within the game tree using our ILP_{LOAS}^{deep} tasks.

Attempting to learn from full games, i. e. comparing the move a player made to other possible moves per turn, *did not yield interesting or powerful strategies*. Further experimentation looking into reducing the grounding and constructing a hypothesis space that is better suited to the games could be an option.

Learning only from suggestions provided by a human player allowed us to create a smaller set of examples that we were more certain reflected the strategy of the player. The results that came from these experiments *did show that ILASP could learn strategies that are good in competition*. Having fewer examples means that the tasks can run quickly, this is important for future work in order to incorporate the learning into an *online* AI player.

Learning from *deep orderings* has taken advantage of being able to inject additional meta-programs into ILASP allowing us to create rules which activate only certain learning examples. In this report we have shown proof-of-concept tasks that highlight the added expressiveness of deep orderings. Though, there is further work that needs to be done in order to reduce the size of the tasks for larger depths.

In conclusion, the experimentation in this report has demonstrated that weak constraints are fully capable to learn strategies; what is more, it has shown the explainable nature of the hypotheses learnt.

11.1 ACHIEVEMENTS

1. We have described a method of applying the Game Description Language to learning strategies (Chapter 5). In particular it means that we can use its defined structure in order to construct tree-based examples from the answer sets of a turn. It also means that the framework presented is not fixed to any of the games presented here in this report, the background knowledge for *Onitama*, *Five Field Kono* and *Cross-Dot* all conform to this specification.
2. In Chapter 6, we created a program to simulate games using different types of players (e. g. human, Monte Carlo and the Clingo AIs) in both *tournament-mode* and *training-mode*. In the two modes we can record different types of learning examples that are given directly to ILASP. Again, new types of board games can be added to this program in a simple manner.
3. We present ILP_{LOAS}^{deep} , a new Inductive Logic Programming framework that introduces *deep orderings*. A new way of describing preferences over *possible* futures based on a single move based on the minimax theorem.

4. Following from a review of the literature (Chapter 4) we explored strategies for *Cross-Dot* (see (Zhang and Thielscher, 2015)).
5. To evaluate the learning methods we provide experimentation in Chapter 10 as well as running tournaments (Section 10.3) to compare the strategies that have been learnt.

11.2 FUTURE WORK

Whilst experimenting with ILASP and *weak constraints* we came across many points to explore, such as:

- Extending the hypothesis space to allow custom weights for rules, this would allow *features* of the games to be learnt and added as rules to background knowledge and used in the mode declarations with a length greater than one in order to not bias the mode declaration too much in their favour,
- Extending the ILP_{LOAS}^{deep} to perform a Monte Carlo style search, this would involve only considering a subset of the children that are considered to be ‘better’ using some heuristic — perhaps learnt or represented using weak constraints.
- Experiment with simple games involving non-determinism, for example the *Royal Game of Ur*, a simple race game using a 4-sided die where the strategy lies in choosing your pieces correctly to hinder your opponent’s movement.
- Looking into relating our explanation condition with formal strategy and game logic, such as those in Benthem (2011) and Zhang and Thielscher (2015).

However, there are two more significant pieces of work that could be explored whilst using *weak constraints* to learn strategies. The first is learning features of the game to perform a quiescent search, and the second is effectively identifying examples that encode the strategy.

11.2.1 PERFORMING QUIESCENT SEARCH WITH WEAK CONSTRAINTS

As the minimax tree can explode rapidly, *deep orderings* become infeasible after a certain depth, depending on the average branching factor of the game. In order to cut down on the branches that are expanded an heuristic could be applied to a state and any branches that are deemed to not be of use can be pruned from the tree. This heuristic can be a set of weak constraints that have been learned previously. Clingo has an option to only return answer sets below a certain penalty, this would potentially require using a wider set of weights when learning than done so in this report. In order to not bias future tasks, you could prune the branch with a certain probability, similar to the exploitation/exploration techniques used in Reinforcement Learning.

The weak constraints learnt for this purpose may be more ‘feature-based’, i. e. a set for captures, a set for checks etc., mimicking the capture-trees and endgame tables in Chess engines. Performing this selective search after a fixed depth minimax is the basis of a quiescent search, and not only prunes the tree but it is able to mitigate against the horizon effect.

11.2.2 IDENTIFYING EXAMPLES WITH STRONG STRATEGIC CHOICES

Based on the results from the experiments we have run in this report we see that having lots of noise and positions where no interesting tactical decisions have been made make learning the strategies more challenging. However, using a human player means that collecting examples is expensive and cumbersome. One future extension to the example collection mechanism could be to have a method of selecting which examples taken from a game are worth keeping and which are not. This could be achieved through a hybrid logic-neural solution using a trained neural network to select examples from a game that are key in different decision processes. The logic framework would offer the same accurate and explainable properties that are so desired in an AI that we have seen throughout this report.

Part IV
APPENDIX

A | LOGIC PROGRAMS

Logic Program A.1. Onitama Background Knowledge

```
1 % Game Constants
2
3 cell(5, 1). cell(5, 2). cell(5, 3). cell(5, 4). cell(5, 5).
4 cell(4, 1). cell(4, 2). cell(4, 3). cell(4, 4). cell(4, 5).
5 cell(3, 1). cell(3, 2). cell(3, 3). cell(3, 4). cell(3, 5).
6 cell(2, 1). cell(2, 2). cell(2, 3). cell(2, 4). cell(2, 5).
7 cell(1, 1). cell(1, 2). cell(1, 3). cell(1, 4). cell(1, 5).
8
9 cell(cell(1..5, 1..5)).
10
11 temple(red, cell(5,3)).
12 temple(blue, cell(1,3)).
13
14 opponent(red, blue).
15 opponent(blue, red).
16
17 role(red).
18 role(blue).
19
20 dir(red, -1).
21 dir(blue, 1).
22
23 captured_master(Player) :- role(Player), not true(location(pawn(master, Player),
  ↪ _)).
24
25 % Action Generation
26 0 { does(Player, Action) } 1 :- legal(Player, Action), not terminal.
27
28 %:- does(Player, A1), does(Player, A2), A1 != A2.
29
30 :- does(P, noop), does(P, move(_, _, _)).
31 % Try splitting this up
32 :- does(P, move(C1, _, _)), does(P, move(C2, _, _)), C1 < C2.
33 :- does(P, move(_, C1, _)), does(P, move(_, C2, _)), C1 < C2.
34
35 :- legal(Player, _), not terminal, not does(Player, _).
36
37
38 % Game State
39
40 % Pawn location
41 next(location(pawn(Rank, Player), Cell)) :-
42 true(location(pawn(Rank, Player), Cell)),
43 not does(_, move((_, Cell), _)), not does(Player, move((Cell, _), _)).
```

```

44 next(location(pawn(Rank, Player), To)) :- true(location(pawn(Rank, Player),
    ↪ From)), does(Player, move((From, To), _)).
45
46 % Card in center
47 true(center_card(Card)) :- in_play(Card), not true(in_hand(red, Card)), not
    ↪ true(in_hand(blue, Card)).
48
49 % Player's hand
50 next(in_hand(Player, Card)) :- true(in_hand(Player, Card)), not
    ↪ true(control(Player)).
51 next(in_hand(Player, Card)) :- true(center_card(Card)), true(control(Player)).
52 % Only needed for 5 card variant
53 next(in_hand(Player, Card)) :- true(in_hand(Player, Card)), Card != Card2,
    ↪ does(Player, move(_, Card2)).
54
55 % Player's turn
56 next(control(Player)) :- true(control(Opp)), opponent(Player, Opp).
57 next(control(blue)) :- true(control(red)).
58
59 goal(Player, 100) :- captured_master(Opp), opponent(Player, Opp).
60 goal(Player, 100) :- temple(Opp, Temple), true(location(pawn(master, Player),
    ↪ Temple)), opponent(Player, Opp).
61 goal(Player, 0) :- goal(Opp, 100), opponent(Player, Opp).
62
63 terminal :- goal(_, 100).
64
65 legal(Player, noop) :- role(Player), not true(control(Player)).
66 legal(Player, move(cell(FromR, FromC), cell(ToR, ToC), Card)) :-
67     true(control(Player)),
68     true(location(pawn(_, Player), cell(FromR, FromC))),
69     not true(location(pawn(_, Player), cell(ToR, ToC))),
70     dir(Player, D), card(Card, (DC, DR)),
71     true(in_hand(Player, Card)),
72     ToR == FromR-(DR*D), ToC == FromC-(DC*D).
73
74 % Cards
75 card(tiger,(0,-2)). card(tiger,(0,1)).
76 card(crab,(0,-1)). card(crab,(-2,0)). card(crab,(2,0)).
77 card(monkey,(1,-1)). card(monkey,(1,1)). card(monkey,(-1,1)).
    ↪ card(monkey,(-1,-1)).
78 card(crane,(0,-1)). card(crane,(-1,1)). card(crane,(1,1)).
79 card(mantis,(-1,-1)). card(mantis,(1,-1)). card(mantis,(0,1)).
80 card(boar,(0,-1)). card(boar,(-1,0)). card(boar,(1,0)).
81 card(dragon,(-1,1)). card(dragon,(1,1)). card(dragon,(2,-1)).
    ↪ card(dragon,(-2,-1)).
82 card(elephant,(-1,-1)). card(elephant,(-1,0)). card(elephant,(1,0)).
    ↪ card(elephant,(1,-1)).
83 card(eel,(-1,-1)). card(eel,(-1,1)). card(eel,(1,0)).
84 card(goose,(-1,-1)). card(goose,(1,1)). card(goose,(1,0)). card(goose,(-1,0)).
85 card(frog,(-1,-1)). card(frog,(1,1)). card(frog,(-2,0)).
86 card(horse,(0,-1)). card(horse,(-1,0)). card(horse,(0,1)).

```



```

87 card(rabbit,(1,-1)). card(rabbit,(-1,1)). card(rabbit,(2,0)).
88 card(ox,(0,-1)). card(ox,(1,0)). card(ox,(0,1)).
89 card(cobra,(1,-1)). card(cobra,(1,1)). card(cobra,(-1,0)).
90 card(rooster,(1,-1)). card(rooster,(-1,1)). card(rooster,(-1,0)).
   ⇒ card(rooster,(1,0)).
91 card(kirin,(-1,-2)). card(kirin,(1,-2)). card(kirin,(0,2)).
92 card(turtle,(-2,0)). card(turtle,(-1,1)). card(turtle,(1,1)).
   ⇒ card(turtle,(2,0)).
93 card(giraffe,(-2,-1)). card(giraffe,(2,-1)). card(giraffe,(0,1)).
94 card(phoenix,(-2,0)). card(phoenix,(-1,-1)). card(phoenix,(1,-1)).
   ⇒ card(phoenix,(2,0)).
95 card(otter,(-1,-1)). card(otter,(1,1)). card(otter,(2,0)).
96 card(viper,(-2,0)). card(viper,(0,-1)). card(viper,(1,1)).
97 card(iguana,(-2,1)). card(iguana,(0,-1)). card(iguana,(1,1)).
98 card(rat,(-1,0)). card(rat,(0,-1)). card(rat,(1,1)).
99 card(bear,(-1,-1)). card(bear,(0,-1)). card(bear,(1,1)).
100 card(dog,(-1,-1)). card(dog,(-1,0)). card(dog,(-1,1)).
101 card(sable,(1,-1)). card(sable,(-1,1)). card(sable,(-2,0)).
102 card(seasnake,(2,0)). card(seasnake,(0,-1)). card(seasnake,(-1,1)).
103 card(tanuki,(2,1)). card(tanuki,(0,-1)). card(tanuki,(-1,1)).
104 card(mouse,(1,0)). card(mouse,(0,-1)). card(mouse,(-1,1)).
105 card(panda,(1,-1)). card(panda,(0,-1)). card(panda,(-1,1)).
106 card(fox,(1,-1)). card(fox,(1,0)). card(fox,(1,1)).

```

Logic Program A.2. Five Field Kono

```

1  role(w).
2  role(b).
3
4  cell(5, 1). cell(5, 2). cell(5, 3). cell(5, 4). cell(5, 5).
5  cell(4, 1). cell(4, 2). cell(4, 3). cell(4, 4). cell(4, 5).
6  cell(3, 1). cell(3, 2). cell(3, 3). cell(3, 4). cell(3, 5).
7  cell(2, 1). cell(2, 2). cell(2, 3). cell(2, 4). cell(2, 5).
8  cell(1, 1). cell(1, 2). cell(1, 3). cell(1, 4). cell(1, 5).
9
10 init(state(1,1,w)).
11 init(state(1,2,w)).
12 init(state(1,3,w)).
13 init(state(1,4,w)).
14 init(state(1,5,w)).
15 init(state(2,1,w)).
16 init(state(2,5,w)).
17
18 init(state(5,1,b)).
19 init(state(5,2,b)).
20 init(state(5,3,b)).
21 init(state(5,4,b)).
22 init(state(5,5,b)).
23 init(state(4,1,b)).
24 init(state(4,5,b)).
25

```

```

26 state(A, B, e) :- not state(A, B, w), not state(A, B, b), cell(A, B).
27
28 not_in_starting_location(w) :- init(state(A, B, b)), not state(A, B, w).
29 not_in_starting_location(b) :- init(state(A, B, w)), not state(A, B, b).
30
31 goal(P, 100) :- not not_in_starting_location(P), role(P).
32 goal(P, 50) :- not legal(_, _), role(P).
33 goal(P, 0) :- not_in_starting_location(P).
34
35 0 { does(P, A) } 1 :- legal(P, A), not terminal.
36 :- role(P), not does(P, _), not terminal.
37 :- does(P, A1), does(P, A2), A1 < A2.
38
39 next(control(x)) :- control(o).
40 next(control(o)) :- control(x).
41
42 next(state(A, B, P)) :- role(P), state(A, B, P), not does(P, move(cell(A, B),
    ↪ _)).
43 next(state(A, B, P)) :- does(P, move(cell(A, B), _)).
44
45 next(box(M,P)) :- does(P,mark(M)), init(box(M,b)).
46 next(box(M,P)) :- init(box(M,P)), P!=b.
47 next(box(M1,b)) :- does(P,mark(M2)), init(box(M1,b)), M1!=M2.
48
49 legal(P, noop) :- role(P), not control(P).
50
51 adj(cell(V0, V1), cell(V2, V3)) :- cell(V0, V1), cell(V2, V3), V2 == V0+1, V3 ==
    ↪ V1+1.
52 adj(cell(V0, V1), cell(V2, V3)) :- cell(V0, V1), cell(V2, V3), V1 == V3+1, V0 ==
    ↪ V2-1.
53 adj(cell(V2, V3), cell(V0, V1)) :- cell(V0, V1), cell(V2, V3), V1 == V3+1, V0 ==
    ↪ V2-1.
54 adj(cell(V0, V1), cell(V2, V3)) :- cell(V0, V1), cell(V2, V3), V1 == V3+1, V2 ==
    ↪ V0-1.
55 legal(V4, move(cell(V0, V1), cell(V2, V3))) :- adj(cell(V0, V1), cell(V2, V3)),
    ↪ state(V0, V1, V4), state(V2, V3, e), control(V4).
56
57 terminal :- goal(_, 100).
58 terminal :- goal(_, 50).

```

Logic Program A.3. Cross-Dot Game

```

1 role(x).
2 role(o).
3 box(1).
4 box(2).
5 box(3).
6 box(4).
7 next(box(M,P)) :- does(P,mark(M)), box(M,b).
8 next(box(M,P)) :- box(M,P), P!=b.
9 next(box(M1,b)) :- does(_,mark(M2)), box(M1,b), M1!=M2.

```

```

10 next(control(x)) :- control(o).
11 next(control(o)) :- control(x).
12 legal(P,mark(M)) :- box(M,b), control(P).
13 legal(P,noop) :- not control(P), role(P).
14 legal(P,noop) :- terminal, role(P).
15 open :- box(_,b).
16 terminal :- not open.
17 terminal :- goal(_, 100).
18 chain_segments(1,M,P) :- role(P), box(M,P).
19 chain_segments(2,M+1,P) :- chain_segments(1, M, P), box(M+1,P).
20 longest_chain(M,P) :- chain_segments(M,_,P), not chain_segments(M+1,_,P).
21 longest_chain(0,P) :- role(P), not chain_segments(,_,P).
22 0 { does(P, mark(B)) } 1 :- legal(P, mark(B)), not terminal.
23 does(P, noop) :- legal(P, noop).
24 :- role(P), not does(P, _), not terminal.
25 :- does(P,mark(B1)), does(P,mark(B2)), B1<B2.
26 goal(P,100) :- longest_chain(2,P).
27 goal(o,0) :- goal(x, 100).
28 goal(x,0) :- goal(o, 100).
29 adj(M,M+1) :- box(M), box(M+1).
30 adj(M+1,M) :- box(M), box(M+1).

```


B | ILASP LEARNING EXAMPLES

B.1 ONITAMA

B.1.1 EXPERIMENT 9.1: ONITAMA RULES

Logic Program B.1. Examples of legal and illegal moves in Onitama

```
1 #pos({legal(blue,move((cell(1,2),cell(2,3)),rabbit))}, {}, {
2 in_play(seasnake).
3 in_play(dog).
4 in_play(rabbit).
5 control(blue).
6 in_hand(red,dog).
7 in_hand(blue,rabbit).
8 location(pawn(master,red),cell(5,3)).
9 location(pawn(student,red),cell(5,1)).
10 location(pawn(student,red),cell(5,2)).
11 location(pawn(student,red),cell(5,4)).
12 location(pawn(student,red),cell(5,5)).
13 location(pawn(master,blue),cell(1,3)).
14 location(pawn(student,blue),cell(1,1)).
15 location(pawn(student,blue),cell(1,2)).
16 location(pawn(student,blue),cell(1,4)).
17 location(pawn(student,blue),cell(1,5)).
18 card(seasnake,(2,0)). card(seasnake,(0,-1)). card(seasnake,(-1,1)).
19 card(dog,(-1,-1)). card(dog,(-1,0)). card(dog,(-1,1)).
20 card(rabbit,(1,-1)). card(rabbit,(-1,1)). card(rabbit,(2,0)).
21 }).
22
23 #pos({ legal(blue,move((cell(2,3),cell(3,3)),seasnake))
24 , legal(blue,move((cell(2,3),cell(1,2)),seasnake))
25 },
26 { legal(blue,move((cell(1,5),cell(1,7)),seasnake))
27 , legal(blue,move((cell(1,5),cell(2,4)),seasnake))
28 , legal(blue,move((cell(2,3),cell(1,4)),seasnake))
29 , legal(red,move((cell(4,1),cell(5,2)),rabbit))
30 }, {
31 in_play(dog).
32 in_play(rabbit).
33 in_play(seasnake).
34 control(blue).
35 in_hand(red,rabbit).
36 in_hand(blue,seasnake).
37 location(pawn(master,red),cell(3,4)).
38 location(pawn(student,red),cell(4,1)).
39 location(pawn(student,red),cell(4,5)).
40 location(pawn(master,blue),cell(1,3)).
```

```

41 location(pawn(student,blue),cell(2,1)).
42 location(pawn(student,blue),cell(2,3)).
43 location(pawn(student,blue),cell(1,5)).
44 card(seasnake,(2,0)). card(seasnake,(0,-1)). card(seasnake,(-1,1)).
45 card(dog,(-1,-1)). card(dog,(-1,0)). card(dog,(-1,1)).
46 card(rabbit,(1,-1)). card(rabbit,(-1,1)). card(rabbit,(2,0)).
47 }).
48 #pos(
49 { legal(red,move((cell(4,3),cell(2,2)),kirin))
50 },
51 { legal(red,move((cell(3,3),cell(5,3)),kirin))
52 , legal(red,move((cell(3,2),cell(1,1)),kirin))
53 , legal(red,move((cell(3,3),cell(2,3)),rat))
54 }, {
55 in_play(fox).
56 in_play(kirin).
57 in_play(rat).
58 control(red).
59 in_hand(red,kirin).
60 in_hand(blue,rat).
61 location(pawn(master,red),cell(5,3)).
62 location(pawn(student,red),cell(4,3)).
63 location(pawn(student,red),cell(3,3)).
64 location(pawn(master,blue),cell(3,5)).
65 location(pawn(student,blue),cell(4,4)).
66 location(pawn(student,blue),cell(2,2)).
67 card(fox,(1,-1)). card(fox,(1,0)). card(fox,(1,1)).
68 card(rat,(-1,0)). card(rat,(0,-1)). card(rat,(1,1)).
69 card(kirin,(-1,-2)). card(kirin,(1,-2)). card(kirin,(0,2)).
70 }).
71
72 #max_penalty(1000).
1 #bias(":- head(pred_1(V0, (V1, V1))).").
2 #bias(":- not head(pred_3(_, _), body(location(_, _))).").
3 #bias(":- not head(pred_4(_, _), body(naf(location(_, _))).").
4 #bias(":- not head(pred_3(_, _), not head(pred_4(_, _), body(control(_))).").
5
6 #bias(":- not head(pred_5(_, _), body(cell(_, _))).").
7 #bias(":- head(pred_5((cell(V0, V0), _, _))).").
8 #bias(":- head(pred_5( (_, cell(V0, V0)), _))).").
9 #bias(":- head(pred_5((cell(_, V0), cell(_, V0)), _))).").
10 #bias(":- head(pred_5((cell(_, V0), cell(V0, _)), _))).").
11 #bias(":- head(pred_5((cell(V0, _), cell(_, V0)), _))).").
12 #bias(":- head(pred_5((cell(V0, _), cell(V0, _)), _))).").
13 #bias(":- head(pred_5(_, (V0, V0))).").
14 #bias(":- body(pred_5(_, (V0, V0))).").
15
16 #bias(":- head(pred_5(_, _), body(cell(V0, _)), body(cell(_, V0))).").
17
18 #bias(":- not head(pred_2(_, _, _), body(card(_, _))).").
19 #bias(":- not head(pred_2(_, _, _), body(pred_6(_, _, _))).").

```

```

20 #bias(":- body(card(_, (V0, V0))).").
21
22 #bias(":- head(pred_3(_, _)), body(pred_1(_, _)).").
23 #bias(":- not head(pred_1(_, _)), body(pred_3(_, _)).").
24 #bias(":- not head(pred_1(_, _)), body(pred_4(_, _)).").
25 #bias(":- not head(pred_6(_, _, _)), body(math(_, _, _, _)).").
26 #bias(":- not head(pred_6(_, _, _)), body(pred_5(_, _)).").
27
28 #bias(":- head(legal(_, move(V1, V1, _))).").
29 #bias(":- body(math(V0, V0, _, _)).").

```

B.1.2 EXPERIMENT 10.7: DEFEND PAWNS

```

% Context: C1
location(pawn(student,red),cell(5,1)).
location(pawn(student,red),cell(5,2)).
location(pawn(master,red),cell(5,3)).
location(pawn(student,red),cell(5,4)).
location(pawn(student,red),cell(4,4)).
location(pawn(student,blue),cell(1,1)).
location(pawn(student,blue),cell(1,2)).
location(pawn(master,blue),cell(1,3)).
location(pawn(student,blue),cell(1,4)).
location(pawn(student,blue),cell(1,5)).
control(blue).
in_play(rat).
in_play(crab).
in_play(mantis).
in_hand(red,rat).
in_hand(blue,crab).

% Context: C2
location(pawn(student,red),cell(5,1)).
location(pawn(student,red),cell(5,2)).
location(pawn(master,red),cell(5,3)).
location(pawn(student,red),cell(5,5)).
location(pawn(student,red),cell(4,5)).
location(pawn(student,blue),cell(1,1)).
location(pawn(student,blue),cell(1,2)).
location(pawn(master,blue),cell(1,3)).
location(pawn(student,blue),cell(1,4)).
location(pawn(student,blue),cell(1,5)).
control(blue).
in_play(rat).
in_play(crab).
in_play(mantis).
in_hand(red,rat).
in_hand(blue,crab).

% Context: C3
location(pawn(student,red),cell(5,1)).
location(pawn(master,red),cell(5,3)).
location(pawn(student,red),cell(5,4)).
location(pawn(student,red),cell(4,2)).
location(pawn(student,red),cell(4,4)).
location(pawn(student,blue),cell(2,1)).
location(pawn(student,blue),cell(2,4)).
location(pawn(student,blue),cell(1,1)).
location(pawn(master,blue),cell(1,3)).
location(pawn(student,blue),cell(1,4)).
control(red).
in_play(mantis).
in_play(crab).

location(pawn(student,blue),cell(2,5)).
location(pawn(student,blue),cell(1,1)).
location(pawn(master,blue),cell(1,3)).
location(pawn(student,blue),cell(1,4)).
control(red).
in_play(crab).
in_play(rat).
in_play(mantis).
in_hand(red,crab).
in_hand(blue,rat).

% Context: C4
location(pawn(student,red),cell(5,1)).
location(pawn(master,red),cell(5,3)).
location(pawn(student,red),cell(5,4)).
location(pawn(student,red),cell(4,2)).
location(pawn(student,red),cell(4,4)).
location(pawn(master,blue),cell(2,2)).
location(pawn(student,blue),cell(2,5)).
location(pawn(student,blue),cell(1,1)).
location(pawn(student,blue),cell(1,2)).
location(pawn(student,blue),cell(1,4)).
control(red).
in_play(crab).
in_play(rat).
in_play(mantis).
in_hand(red,crab).
in_hand(blue,rat).

% Context: C5
location(pawn(master,red),cell(5,3)).
location(pawn(student,red),cell(5,4)).
location(pawn(student,red),cell(4,1)).
location(pawn(student,red),cell(4,2)).
location(pawn(student,red),cell(4,4)).
location(pawn(student,blue),cell(2,1)).
location(pawn(student,blue),cell(2,4)).
location(pawn(student,blue),cell(1,1)).
location(pawn(master,blue),cell(1,3)).
location(pawn(student,blue),cell(1,4)).
control(red).
in_play(mantis).
in_play(crab).

```

```

in_play(rat).
in_hand(red,mantis).
in_hand(blue,crab).

% Context: C6
location(pawn(master,red),cell(5,3)).
location(pawn(student,red),cell(5,4)).
location(pawn(student,red),cell(4,1)).
location(pawn(student,red),cell(4,2)).
location(pawn(student,red),cell(4,4)).
location(pawn(student,blue),cell(2,1)).

location(pawn(master,blue),cell(2,3)).
location(pawn(student,blue),cell(2,5)).
location(pawn(student,blue),cell(1,1)).
location(pawn(student,blue),cell(1,4)).
control(red).
in_play(mantis).
in_play(crab).
in_play(rat).
in_hand(red,mantis).
in_hand(blue,crab).

```

Logic Program B.2. Deep Ordering Translation for the Defend Task

```

272 location(pawn(master,red),cell(5,3)).
273 location(pawn(student,blue),cell(2,1)).
274 location(pawn(student,blue),cell(2,4)).
275 location(pawn(student,blue),cell(1,1)).
276 location(pawn(master,blue),cell(1,3)).
277 location(pawn(student,blue),cell(1,4)).
278 in_hand(blue,crab).
279 in_hand(red,crab).
280 control(blue).
281 location(pawn(student,red),cell(5,4)).
282 location(pawn(student,red),cell(4,1)).
283 location(pawn(student,red),cell(4,4)).
284 location(pawn(student,red),cell(3,1)).
285 does(blue,move(cell(2,4),cell(2,2),crab)).
286 does(red,noop).
287 }).
288 #pos(ex_a_0_4, {}, {}, {
289 location(pawn(master,red),cell(5,3)).
290 location(pawn(student,blue),cell(2,1)).
291 location(pawn(student,blue),cell(2,4)).
292 location(pawn(student,blue),cell(1,1)).
293 location(pawn(master,blue),cell(1,3)).
294 location(pawn(student,blue),cell(1,4)).
295 in_hand(blue,crab).
296 in_hand(red,crab).
297 control(blue).
298 location(pawn(student,red),cell(5,4)).
299 location(pawn(student,red),cell(4,1)).
300 location(pawn(student,red),cell(4,4)).
301 location(pawn(student,red),cell(3,1)).
302 does(blue,move(cell(2,4),cell(3,4),crab)).
303 does(red,noop).
304 }).
305 #pos(ex_a_0_5, {}, {}, {
306 location(pawn(master,red),cell(5,3)).
307 location(pawn(student,blue),cell(2,1)).
308 location(pawn(student,blue),cell(2,4)).
309 location(pawn(student,blue),cell(1,1)).
310 location(pawn(master,blue),cell(1,3)).

```



```

311 location(pawn(student,blue),cell(1,4)).
312 in_hand(blue,crab).
313 in_hand(red,rat).
314 control(blue).
315 location(pawn(student,red),cell(5,4)).
316 location(pawn(student,red),cell(4,1)).
317 location(pawn(student,red),cell(4,4)).
318 location(pawn(student,red),cell(3,1)).
319 does(blue,move(cell(1,3),cell(1,5),crab)).
320 does(red,noop).
321 }).
322 #pos(ex_a_0_6, {}, {}, {
323 location(pawn(master,red),cell(5,3)).
324 location(pawn(student,blue),cell(2,1)).
325 location(pawn(student,blue),cell(2,4)).
326 location(pawn(student,blue),cell(1,1)).
327 location(pawn(master,blue),cell(1,3)).
328 location(pawn(student,blue),cell(1,4)).
329 in_hand(blue,crab).
330 in_hand(red,rat).
331 control(blue).
332 location(pawn(student,red),cell(5,4)).
333 location(pawn(student,red),cell(4,1)).
334 location(pawn(student,red),cell(4,4)).
335 location(pawn(student,red),cell(3,1)).
336 does(blue,move(cell(2,1),cell(2,3),crab)).
337 does(red,noop).
338 }).
339 #pos(ex_a_0_7, {}, {}, {
340 location(pawn(master,red),cell(5,3)).
341 location(pawn(student,blue),cell(2,1)).
342 location(pawn(student,blue),cell(2,4)).
343 location(pawn(student,blue),cell(1,1)).
344 location(pawn(master,blue),cell(1,3)).
345 location(pawn(student,blue),cell(1,4)).
346 in_hand(blue,crab).
347 in_hand(red,rat).
348 control(blue).
349 location(pawn(student,red),cell(5,4)).
350 location(pawn(student,red),cell(4,1)).
351 location(pawn(student,red),cell(4,4)).
352 location(pawn(student,red),cell(3,1)).
353 does(blue,move(cell(1,3),cell(2,3),crab)).
354 does(red,noop).
355 }).
356 #pos(ex_a_1_1, {}, {}, {
357 location(pawn(master,red),cell(5,3)).
358 location(pawn(student,blue),cell(2,1)).
359 location(pawn(student,blue),cell(2,4)).
360 location(pawn(student,blue),cell(1,1)).
361 location(pawn(master,blue),cell(1,3)).

```

```

362 location(pawn(student,blue),cell(1,4)).
363 in_hand(blue,crab).
364 in_hand(red,rat).
365 control(blue).
366 location(pawn(student,red),cell(4,1)).
367 location(pawn(student,red),cell(4,2)).
368 location(pawn(student,red),cell(4,4)).
369 location(pawn(student,red),cell(4,5)).
370 does(blue,move(cell(1,4),cell(1,2),crab)).
371 does(red,noop).
372 }).
373 #pos(ex_a_1_2, {}, {}, {
374 location(pawn(master,red),cell(5,3)).
375 location(pawn(student,blue),cell(2,1)).
376 location(pawn(student,blue),cell(2,4)).
377 location(pawn(student,blue),cell(1,1)).
378 location(pawn(master,blue),cell(1,3)).
379 location(pawn(student,blue),cell(1,4)).
380 in_hand(blue,crab).
381 in_hand(red,rat).
382 control(blue).
383 location(pawn(student,red),cell(4,1)).
384 location(pawn(student,red),cell(4,2)).
385 location(pawn(student,red),cell(4,4)).
386 location(pawn(student,red),cell(4,5)).
387 does(blue,move(cell(2,1),cell(3,1),crab)).
388 does(red,noop).
389 }).
390 #pos(ex_a_1_3, {}, {}, {
391 location(pawn(master,red),cell(5,3)).
392 location(pawn(student,blue),cell(2,1)).
393 location(pawn(student,blue),cell(2,4)).
394 location(pawn(student,blue),cell(1,1)).
395 location(pawn(master,blue),cell(1,3)).
396 location(pawn(student,blue),cell(1,4)).
397 in_hand(blue,crab).
398 in_hand(red,rat).
399 control(blue).
400 location(pawn(student,red),cell(4,1)).
401 location(pawn(student,red),cell(4,2)).
402 location(pawn(student,red),cell(4,4)).
403 location(pawn(student,red),cell(4,5)).
404 does(blue,move(cell(2,4),cell(2,2),crab)).
405 does(red,noop).
406 }).
407 #pos(ex_a_1_4, {}, {}, {
408 location(pawn(master,red),cell(5,3)).
409 location(pawn(student,blue),cell(2,1)).
410 location(pawn(student,blue),cell(2,4)).
411 location(pawn(student,blue),cell(1,1)).
412 location(pawn(master,blue),cell(1,3)).

```

```

413 location(pawn(student,blue),cell(1,4)).
414 in_hand(blue,crab).
415 in_hand(red,rat).
416 control(blue).
417 location(pawn(student,red),cell(4,1)).
418 location(pawn(student,red),cell(4,2)).
419 location(pawn(student,red),cell(4,4)).
420 location(pawn(student,red),cell(4,5)).
421 does(blue,move(cell(2,4),cell(3,4),crab)).
422 does(red,noop).
423 }).
424 #pos(ex_a_1_5, {}, {}, {
425 location(pawn(master,red),cell(5,3)).
426 location(pawn(student,blue),cell(2,1)).
427 location(pawn(student,blue),cell(2,4)).
428 location(pawn(student,blue),cell(1,1)).
429 location(pawn(master,blue),cell(1,3)).
430 location(pawn(student,blue),cell(1,4)).
431 in_hand(blue,crab).
432 in_hand(red,rat).
433 control(blue).
434 location(pawn(student,red),cell(4,1)).
435 location(pawn(student,red),cell(4,2)).
436 location(pawn(student,red),cell(4,4)).
437 location(pawn(student,red),cell(4,5)).
438 does(blue,move(cell(1,3),cell(1,5),crab)).
439 does(red,noop).
440 }).
441 #pos(ex_a_1_6, {}, {}, {
442 location(pawn(master,red),cell(5,3)).
443 location(pawn(student,blue),cell(2,1)).
444 location(pawn(student,blue),cell(2,4)).
445 location(pawn(student,blue),cell(1,1)).
446 location(pawn(master,blue),cell(1,3)).
447 location(pawn(student,blue),cell(1,4)).
448 in_hand(blue,crab).
449 in_hand(red,rat).
450 control(blue).
451 location(pawn(student,red),cell(4,1)).
452 location(pawn(student,red),cell(4,2)).
453 location(pawn(student,red),cell(4,4)).
454 location(pawn(student,red),cell(4,5)).
455 does(blue,move(cell(2,1),cell(2,3),crab)).
456 does(red,noop).
457 }).
458 #pos(ex_a_1_7, {}, {}, {
459 location(pawn(master,red),cell(5,3)).
460 location(pawn(student,blue),cell(2,1)).
461 location(pawn(student,blue),cell(2,4)).
462 location(pawn(student,blue),cell(1,1)).
463 location(pawn(master,blue),cell(1,3)).

```

```

464 location(pawn(student,blue),cell(1,4)).
465 in_hand(blue,crab).
466 in_hand(red,rat).
467 control(blue).
468 location(pawn(student,red),cell(4,1)).
469 location(pawn(student,red),cell(4,2)).
470 location(pawn(student,red),cell(4,4)).
471 location(pawn(student,red),cell(4,5)).
472 does(blue,move(cell(1,3),cell(2,3),crab)).
473 does(red,noop).
474 }).
475 #brave_ordering(ord_a_1_ex_a_0_1_ex_a_1_1, ex_a_0_1, ex_a_1_1).
476 #brave_ordering(ord_a_1_ex_a_0_1_ex_a_1_2, ex_a_0_1, ex_a_1_2).
477 #brave_ordering(ord_a_1_ex_a_0_1_ex_a_1_3, ex_a_0_1, ex_a_1_3).
478 #brave_ordering(ord_a_1_ex_a_0_1_ex_a_1_4, ex_a_0_1, ex_a_1_4).
479 #brave_ordering(ord_a_1_ex_a_0_1_ex_a_1_5, ex_a_0_1, ex_a_1_5).
480 #brave_ordering(ord_a_1_ex_a_0_1_ex_a_1_6, ex_a_0_1, ex_a_1_6).
481 #brave_ordering(ord_a_1_ex_a_0_1_ex_a_1_7, ex_a_0_1, ex_a_1_7).
482 #brave_ordering(ord_a_1_ex_a_0_2_ex_a_1_1, ex_a_0_2, ex_a_1_1).
483 #brave_ordering(ord_a_1_ex_a_0_2_ex_a_1_2, ex_a_0_2, ex_a_1_2).
484 #brave_ordering(ord_a_1_ex_a_0_2_ex_a_1_3, ex_a_0_2, ex_a_1_3).
485 #brave_ordering(ord_a_1_ex_a_0_2_ex_a_1_4, ex_a_0_2, ex_a_1_4).
486 #brave_ordering(ord_a_1_ex_a_0_2_ex_a_1_5, ex_a_0_2, ex_a_1_5).
487 #brave_ordering(ord_a_1_ex_a_0_2_ex_a_1_6, ex_a_0_2, ex_a_1_6).
488 #brave_ordering(ord_a_1_ex_a_0_2_ex_a_1_7, ex_a_0_2, ex_a_1_7).
489 #brave_ordering(ord_a_1_ex_a_0_3_ex_a_1_1, ex_a_0_3, ex_a_1_1).
490 #brave_ordering(ord_a_1_ex_a_0_3_ex_a_1_2, ex_a_0_3, ex_a_1_2).
491 #brave_ordering(ord_a_1_ex_a_0_3_ex_a_1_3, ex_a_0_3, ex_a_1_3).
492 #brave_ordering(ord_a_1_ex_a_0_3_ex_a_1_4, ex_a_0_3, ex_a_1_4).
493 #brave_ordering(ord_a_1_ex_a_0_3_ex_a_1_5, ex_a_0_3, ex_a_1_5).
494 #brave_ordering(ord_a_1_ex_a_0_3_ex_a_1_6, ex_a_0_3, ex_a_1_6).
495 #brave_ordering(ord_a_1_ex_a_0_3_ex_a_1_7, ex_a_0_3, ex_a_1_7).
496 #brave_ordering(ord_a_1_ex_a_0_4_ex_a_1_1, ex_a_0_4, ex_a_1_1).
497 #brave_ordering(ord_a_1_ex_a_0_4_ex_a_1_2, ex_a_0_4, ex_a_1_2).
498 #brave_ordering(ord_a_1_ex_a_0_4_ex_a_1_3, ex_a_0_4, ex_a_1_3).
499 #brave_ordering(ord_a_1_ex_a_0_4_ex_a_1_4, ex_a_0_4, ex_a_1_4).
500 #brave_ordering(ord_a_1_ex_a_0_4_ex_a_1_5, ex_a_0_4, ex_a_1_5).
501 #brave_ordering(ord_a_1_ex_a_0_4_ex_a_1_6, ex_a_0_4, ex_a_1_6).
502 #brave_ordering(ord_a_1_ex_a_0_4_ex_a_1_7, ex_a_0_4, ex_a_1_7).
503 #brave_ordering(ord_a_1_ex_a_0_5_ex_a_1_1, ex_a_0_5, ex_a_1_1).
504 #brave_ordering(ord_a_1_ex_a_0_5_ex_a_1_2, ex_a_0_5, ex_a_1_2).
505 #brave_ordering(ord_a_1_ex_a_0_5_ex_a_1_3, ex_a_0_5, ex_a_1_3).
506 #brave_ordering(ord_a_1_ex_a_0_5_ex_a_1_4, ex_a_0_5, ex_a_1_4).
507 #brave_ordering(ord_a_1_ex_a_0_5_ex_a_1_5, ex_a_0_5, ex_a_1_5).
508 #brave_ordering(ord_a_1_ex_a_0_5_ex_a_1_6, ex_a_0_5, ex_a_1_6).
509 #brave_ordering(ord_a_1_ex_a_0_5_ex_a_1_7, ex_a_0_5, ex_a_1_7).
510 #brave_ordering(ord_a_1_ex_a_0_6_ex_a_1_1, ex_a_0_6, ex_a_1_1).
511 #brave_ordering(ord_a_1_ex_a_0_6_ex_a_1_2, ex_a_0_6, ex_a_1_2).
512 #brave_ordering(ord_a_1_ex_a_0_6_ex_a_1_3, ex_a_0_6, ex_a_1_3).
513 #brave_ordering(ord_a_1_ex_a_0_6_ex_a_1_4, ex_a_0_6, ex_a_1_4).
514 #brave_ordering(ord_a_1_ex_a_0_6_ex_a_1_5, ex_a_0_6, ex_a_1_5).

```

```

515 #brave_ordering(ord_a_1_ex_a_0_6_ex_a_1_6, ex_a_0_6, ex_a_1_6).
516 #brave_ordering(ord_a_1_ex_a_0_6_ex_a_1_7, ex_a_0_6, ex_a_1_7).
517 #brave_ordering(ord_a_1_ex_a_0_7_ex_a_1_1, ex_a_0_7, ex_a_1_1).
518 #brave_ordering(ord_a_1_ex_a_0_7_ex_a_1_2, ex_a_0_7, ex_a_1_2).
519 #brave_ordering(ord_a_1_ex_a_0_7_ex_a_1_3, ex_a_0_7, ex_a_1_3).
520 #brave_ordering(ord_a_1_ex_a_0_7_ex_a_1_4, ex_a_0_7, ex_a_1_4).
521 #brave_ordering(ord_a_1_ex_a_0_7_ex_a_1_5, ex_a_0_7, ex_a_1_5).
522 #brave_ordering(ord_a_1_ex_a_0_7_ex_a_1_6, ex_a_0_7, ex_a_1_6).
523 #brave_ordering(ord_a_1_ex_a_0_7_ex_a_1_7, ex_a_0_7, ex_a_1_7).
524 #inject("r_o_o_t(ex_a_0,chosen).
525 r_o_o_t(ex_a_1,other).
526 c_h_i_l_d(ex_a_0,ex_a_0_1).
527 c_h_i_l_d(ex_a_0,ex_a_0_2).
528 c_h_i_l_d(ex_a_0,ex_a_0_3).
529 c_h_i_l_d(ex_a_0,ex_a_0_4).
530 c_h_i_l_d(ex_a_0,ex_a_0_5).
531 c_h_i_l_d(ex_a_0,ex_a_0_6).
532 c_h_i_l_d(ex_a_0,ex_a_0_7).
533 c_h_i_l_d(ex_a_1,ex_a_1_1).
534 c_h_i_l_d(ex_a_1,ex_a_1_2).
535 c_h_i_l_d(ex_a_1,ex_a_1_3).
536 c_h_i_l_d(ex_a_1,ex_a_1_4).
537 c_h_i_l_d(ex_a_1,ex_a_1_5).
538 c_h_i_l_d(ex_a_1,ex_a_1_6).
539 c_h_i_l_d(ex_a_1,ex_a_1_7).
540 :- c_h_i_l_d(P1,C), c_h_i_l_d(P2,C), P1<P2.
541 example_active(EX_ID,forall) :- r_o_o_t(EX_ID,chosen).
542 example_active(EX_ID,exists) :- r_o_o_t(EX_ID,other).
543 1 {example_active(Child,forall) : c_h_i_l_d(Parent,Child)} 1 :-
    ⇨ example_active(Parent,exists), c_h_i_l_d(Parent,_).
544 example_active(Child,exists) :- c_h_i_l_d(Parent,Child),
    ⇨ example_active(Parent,forall).
545 example_active(ORD_ID) :- o_r_d(EX_ID_1,EX_ID_2,ORD_ID),
    ⇨ example_active(EX_ID_1,_), example_active(EX_ID_2,_).
546 o_r_d(ex_a_0_1,ex_a_1_1,ord_a_1_ex_a_0_1_ex_a_1_1).
547 o_r_d(ex_a_0_1,ex_a_1_2,ord_a_1_ex_a_0_1_ex_a_1_2).
548 o_r_d(ex_a_0_1,ex_a_1_3,ord_a_1_ex_a_0_1_ex_a_1_3).
549 o_r_d(ex_a_0_1,ex_a_1_4,ord_a_1_ex_a_0_1_ex_a_1_4).
550 o_r_d(ex_a_0_1,ex_a_1_5,ord_a_1_ex_a_0_1_ex_a_1_5).
551 o_r_d(ex_a_0_1,ex_a_1_6,ord_a_1_ex_a_0_1_ex_a_1_6).
552 o_r_d(ex_a_0_1,ex_a_1_7,ord_a_1_ex_a_0_1_ex_a_1_7).
553 o_r_d(ex_a_0_2,ex_a_1_1,ord_a_1_ex_a_0_2_ex_a_1_1).
554 o_r_d(ex_a_0_2,ex_a_1_2,ord_a_1_ex_a_0_2_ex_a_1_2).
555 o_r_d(ex_a_0_2,ex_a_1_3,ord_a_1_ex_a_0_2_ex_a_1_3).
556 o_r_d(ex_a_0_2,ex_a_1_4,ord_a_1_ex_a_0_2_ex_a_1_4).
557 o_r_d(ex_a_0_2,ex_a_1_5,ord_a_1_ex_a_0_2_ex_a_1_5).
558 o_r_d(ex_a_0_2,ex_a_1_6,ord_a_1_ex_a_0_2_ex_a_1_6).
559 o_r_d(ex_a_0_2,ex_a_1_7,ord_a_1_ex_a_0_2_ex_a_1_7).
560 o_r_d(ex_a_0_3,ex_a_1_1,ord_a_1_ex_a_0_3_ex_a_1_1).
561 o_r_d(ex_a_0_3,ex_a_1_2,ord_a_1_ex_a_0_3_ex_a_1_2).
562 o_r_d(ex_a_0_3,ex_a_1_3,ord_a_1_ex_a_0_3_ex_a_1_3).

```

```
563 o_r_d(ex_a_0_3,ex_a_1_4,ord_a_1_ex_a_0_3_ex_a_1_4).
564 o_r_d(ex_a_0_3,ex_a_1_5,ord_a_1_ex_a_0_3_ex_a_1_5).
565 o_r_d(ex_a_0_3,ex_a_1_6,ord_a_1_ex_a_0_3_ex_a_1_6).
566 o_r_d(ex_a_0_3,ex_a_1_7,ord_a_1_ex_a_0_3_ex_a_1_7).
567 o_r_d(ex_a_0_4,ex_a_1_1,ord_a_1_ex_a_0_4_ex_a_1_1).
568 o_r_d(ex_a_0_4,ex_a_1_2,ord_a_1_ex_a_0_4_ex_a_1_2).
569 o_r_d(ex_a_0_4,ex_a_1_3,ord_a_1_ex_a_0_4_ex_a_1_3).
570 o_r_d(ex_a_0_4,ex_a_1_4,ord_a_1_ex_a_0_4_ex_a_1_4).
571 o_r_d(ex_a_0_4,ex_a_1_5,ord_a_1_ex_a_0_4_ex_a_1_5).
572 o_r_d(ex_a_0_4,ex_a_1_6,ord_a_1_ex_a_0_4_ex_a_1_6).
573 o_r_d(ex_a_0_4,ex_a_1_7,ord_a_1_ex_a_0_4_ex_a_1_7).
574 o_r_d(ex_a_0_5,ex_a_1_1,ord_a_1_ex_a_0_5_ex_a_1_1).
575 o_r_d(ex_a_0_5,ex_a_1_2,ord_a_1_ex_a_0_5_ex_a_1_2).
576 o_r_d(ex_a_0_5,ex_a_1_3,ord_a_1_ex_a_0_5_ex_a_1_3).
577 o_r_d(ex_a_0_5,ex_a_1_4,ord_a_1_ex_a_0_5_ex_a_1_4).
578 o_r_d(ex_a_0_5,ex_a_1_5,ord_a_1_ex_a_0_5_ex_a_1_5).
579 o_r_d(ex_a_0_5,ex_a_1_6,ord_a_1_ex_a_0_5_ex_a_1_6).
580 o_r_d(ex_a_0_5,ex_a_1_7,ord_a_1_ex_a_0_5_ex_a_1_7).
581 o_r_d(ex_a_0_6,ex_a_1_1,ord_a_1_ex_a_0_6_ex_a_1_1).
582 o_r_d(ex_a_0_6,ex_a_1_2,ord_a_1_ex_a_0_6_ex_a_1_2).
583 o_r_d(ex_a_0_6,ex_a_1_3,ord_a_1_ex_a_0_6_ex_a_1_3).
584 o_r_d(ex_a_0_6,ex_a_1_4,ord_a_1_ex_a_0_6_ex_a_1_4).
585 o_r_d(ex_a_0_6,ex_a_1_5,ord_a_1_ex_a_0_6_ex_a_1_5).
586 o_r_d(ex_a_0_6,ex_a_1_6,ord_a_1_ex_a_0_6_ex_a_1_6).
587 o_r_d(ex_a_0_6,ex_a_1_7,ord_a_1_ex_a_0_6_ex_a_1_7).
588 o_r_d(ex_a_0_7,ex_a_1_1,ord_a_1_ex_a_0_7_ex_a_1_1).
589 o_r_d(ex_a_0_7,ex_a_1_2,ord_a_1_ex_a_0_7_ex_a_1_2).
590 o_r_d(ex_a_0_7,ex_a_1_3,ord_a_1_ex_a_0_7_ex_a_1_3).
591 o_r_d(ex_a_0_7,ex_a_1_4,ord_a_1_ex_a_0_7_ex_a_1_4).
592 o_r_d(ex_a_0_7,ex_a_1_5,ord_a_1_ex_a_0_7_ex_a_1_5).
593 o_r_d(ex_a_0_7,ex_a_1_6,ord_a_1_ex_a_0_7_ex_a_1_6).
594 o_r_d(ex_a_0_7,ex_a_1_7,ord_a_1_ex_a_0_7_ex_a_1_7).
595 example_active(ex_d_0).
596 example_active(ex_d_1).
597 example_active(ex_c_0).
598 example_active(ex_c_1).
599 example_active(ex_b_0).
600 example_active(ex_b_1).
601 example_active(ex_a_0).
602 example_active(ex_a_1).
603 example_active(ex_a_0_1).
604 example_active(ex_a_0_2).
605 example_active(ex_a_0_3).
606 example_active(ex_a_0_4).
607 example_active(ex_a_0_5).
608 example_active(ex_a_0_6).
609 example_active(ex_a_0_7).
610 example_active(ex_a_1_1).
611 example_active(ex_a_1_2).
612 example_active(ex_a_1_3).
613 example_active(ex_a_1_4).
```

```

614 example_active(ex_a_1_5).
615 example_active(ex_a_1_6).
616 example_active(ex_a_1_7).

```

B.2 FIVE FIELD KONO

B.2.1 EXPERIMENT 9.2: FIVE FIELD KONO RULES

Logic Program B.3. Examples of legal and illegal moves in Five Field Kono

```

1  #pos(init,
2  { legal(w, move(cell(1,2), cell(2,3)))
3  , legal(b, noop)
4  , legal(w, move(cell(2,5), cell(3,4)))
5  },
6  { legal(w, noop)
7  , legal(w, move(cell(1,2), cell(2,2)))
8  , legal(w, move(cell(1,2), cell(1,1)))
9  },
10 {
11 control(w).
12
13 state(1,1,w).
14 state(1,2,w).
15 state(1,3,w).
16 state(1,4,w).
17 state(1,5,w).
18 state(2,1,w).
19 state(2,5,w).
20
21 state(5,1,b).
22 state(5,2,b).
23 state(5,3,b).
24 state(5,4,b).
25 state(5,5,b).
26 state(4,1,b).
27 state(4,5,b).
28 }).
29
30 #pos(mid,
31 { legal(b, move(cell(5,1), cell(4,2)))
32 , legal(w, noop)
33 , legal(b, move(cell(4,3), cell(5,4)))
34 },
35 { legal(b, noop)
36 , legal(b, move(cell(4,3), cell(3,4)))
37 , legal(b, move(cell(4,2), cell(3,1)))
38 , legal(b, move(cell(4,2), cell(3,3)))
39 , legal(w, move(cell(1,2), cell(2,3)))
40 , legal(w, move(cell(2,3), cell(3,2)))
41 },
42 {
43 control(b).
44
45 state(1,1,w).
46 state(2,3,w).
47 state(1,3,w).

```

```

48 state(1,4,w).
49 state(1,5,w).
50 state(2,1,w).
51 state(3,4,w).
52
53 state(5,1,b).
54 state(5,2,b).
55 state(5,3,b).
56 state(4,3,b).
57 state(5,5,b).
58 state(4,1,b).
59 state(4,5,b).
60 }).

61 3 ~ pred(cell(V0, V1), cell(V3, V4), V2) :- state(V0, V1, V2), state(V3, V4, e).
62 3 ~ pred(cell(V0, V1), cell(V4, V3), V2) :- state(V0, V1, V2), state(V3, V4, e).
63 3 ~ pred(cell(V0, V3), cell(V1, V4), V2) :- state(V0, V1, V2), state(V3, V4, e).
64 3 ~ pred(cell(V0, V3), cell(V4, V1), V2) :- state(V0, V1, V2), state(V3, V4, e).
65 3 ~ pred(cell(V0, V4), cell(V1, V3), V2) :- state(V0, V1, V2), state(V3, V4, e).
66 3 ~ pred(cell(V0, V4), cell(V3, V1), V2) :- state(V0, V1, V2), state(V3, V4, e).
67 3 ~ pred(cell(V1, V0), cell(V3, V4), V2) :- state(V0, V1, V2), state(V3, V4, e).
68 3 ~ pred(cell(V1, V0), cell(V4, V3), V2) :- state(V0, V1, V2), state(V3, V4, e).
69 3 ~ pred(cell(V1, V3), cell(V0, V4), V2) :- state(V0, V1, V2), state(V3, V4, e).
70 3 ~ pred(cell(V1, V3), cell(V4, V0), V2) :- state(V0, V1, V2), state(V3, V4, e).
71 3 ~ pred(cell(V1, V4), cell(V0, V3), V2) :- state(V0, V1, V2), state(V3, V4, e).
72 3 ~ pred(cell(V1, V4), cell(V3, V0), V2) :- state(V0, V1, V2), state(V3, V4, e).
73 3 ~ pred(cell(V3, V0), cell(V1, V4), V2) :- state(V0, V1, V2), state(V3, V4, e).
74 3 ~ pred(cell(V3, V0), cell(V4, V1), V2) :- state(V0, V1, V2), state(V3, V4, e).
75 3 ~ pred(cell(V3, V1), cell(V0, V4), V2) :- state(V0, V1, V2), state(V3, V4, e).
76 3 ~ pred(cell(V3, V1), cell(V4, V0), V2) :- state(V0, V1, V2), state(V3, V4, e).
77 3 ~ pred(cell(V3, V4), cell(V0, V1), V2) :- state(V0, V1, V2), state(V3, V4, e).
78 3 ~ pred(cell(V3, V4), cell(V1, V0), V2) :- state(V0, V1, V2), state(V3, V4, e).
79 3 ~ pred(cell(V4, V0), cell(V1, V3), V2) :- state(V0, V1, V2), state(V3, V4, e).
80 3 ~ pred(cell(V4, V0), cell(V3, V1), V2) :- state(V0, V1, V2), state(V3, V4, e).
81 3 ~ pred(cell(V4, V1), cell(V0, V3), V2) :- state(V0, V1, V2), state(V3, V4, e).
82 3 ~ pred(cell(V4, V1), cell(V3, V0), V2) :- state(V0, V1, V2), state(V3, V4, e).
83 3 ~ pred(cell(V4, V3), cell(V0, V1), V2) :- state(V0, V1, V2), state(V3, V4, e).
84 3 ~ pred(cell(V4, V3), cell(V1, V0), V2) :- state(V0, V1, V2), state(V3, V4, e).
85 5 ~ adj(cell(V0, V1), cell(V2, V3)) :- cell(V0, V1), cell(V2, V3), V0 == V2+1, V1 == V3+1.
86 5 ~ adj(cell(V2, V3), cell(V0, V1)) :- cell(V0, V1), cell(V2, V3), V0 == V2+1, V1 == V3+1.
87 5 ~ adj(cell(V0, V1), cell(V2, V3)) :- cell(V0, V1), cell(V2, V3), V0 == V2+1, V3 == V1+1.
88 5 ~ adj(cell(V2, V3), cell(V0, V1)) :- cell(V0, V1), cell(V2, V3), V0 == V2+1, V3 == V1+1.
89 5 ~ adj(cell(V0, V1), cell(V2, V3)) :- cell(V0, V1), cell(V2, V3), V1 == V3+1, V2 == V0+1.
90 5 ~ adj(cell(V2, V3), cell(V0, V1)) :- cell(V0, V1), cell(V2, V3), V1 == V3+1, V2 == V0+1.
91 5 ~ adj(cell(V0, V1), cell(V2, V3)) :- cell(V0, V1), cell(V2, V3), V2 == V0+1, V3 == V1+1.
92 5 ~ adj(cell(V2, V3), cell(V0, V1)) :- cell(V0, V1), cell(V2, V3), V2 == V0+1, V3 == V1+1.
93 5 ~ adj(cell(V0, V1), cell(V2, V3)) :- cell(V0, V1), cell(V2, V3), V0 == V2+1, V1 == V3-1.
94 5 ~ adj(cell(V2, V3), cell(V0, V1)) :- cell(V0, V1), cell(V2, V3), V0 == V2+1, V1 == V3-1.
95 5 ~ adj(cell(V0, V1), cell(V2, V3)) :- cell(V0, V1), cell(V2, V3), V0 == V2+1, V3 == V1-1.
96 5 ~ adj(cell(V2, V3), cell(V0, V1)) :- cell(V0, V1), cell(V2, V3), V0 == V2+1, V3 == V1-1.
97 5 ~ adj(cell(V0, V1), cell(V2, V3)) :- cell(V0, V1), cell(V2, V3), V1 == V3+1, V0 == V2-1.
98 5 ~ adj(cell(V2, V3), cell(V0, V1)) :- cell(V0, V1), cell(V2, V3), V1 == V3+1, V0 == V2-1.
99 5 ~ adj(cell(V0, V1), cell(V2, V3)) :- cell(V0, V1), cell(V2, V3), V1 == V3+1, V2 == V0-1.
100 5 ~ adj(cell(V2, V3), cell(V0, V1)) :- cell(V0, V1), cell(V2, V3), V1 == V3+1, V2 == V0-1.
101 5 ~ adj(cell(V0, V1), cell(V2, V3)) :- cell(V0, V1), cell(V2, V3), V2 == V0+1, V1 == V3-1.
102 5 ~ adj(cell(V2, V3), cell(V0, V1)) :- cell(V0, V1), cell(V2, V3), V2 == V0+1, V1 == V3-1.
103 5 ~ adj(cell(V0, V1), cell(V2, V3)) :- cell(V0, V1), cell(V2, V3), V2 == V0+1, V3 == V1-1.
104 5 ~ adj(cell(V2, V3), cell(V0, V1)) :- cell(V0, V1), cell(V2, V3), V2 == V0+1, V3 == V1-1.
105 5 ~ adj(cell(V0, V1), cell(V2, V3)) :- cell(V0, V1), cell(V2, V3), V3 == V1+1, V0 == V2-1.

```



```

106 5 ~ adj(cell(V2, V3), cell(V0, V1)) :- cell(V0, V1), cell(V2, V3), V3 == V1+1, V0 == V2-1.
107 5 ~ adj(cell(V0, V1), cell(V2, V3)) :- cell(V0, V1), cell(V2, V3), V3 == V1+1, V2 == V0-1.
108 5 ~ adj(cell(V2, V3), cell(V0, V1)) :- cell(V0, V1), cell(V2, V3), V3 == V1+1, V2 == V0-1.
109 5 ~ adj(cell(V0, V1), cell(V2, V3)) :- cell(V0, V1), cell(V2, V3), V0 == V2-1, V1 == V3-1.
110 5 ~ adj(cell(V2, V3), cell(V0, V1)) :- cell(V0, V1), cell(V2, V3), V0 == V2-1, V1 == V3-1.
111 5 ~ adj(cell(V0, V1), cell(V2, V3)) :- cell(V0, V1), cell(V2, V3), V0 == V2-1, V3 == V1-1.
112 5 ~ adj(cell(V2, V3), cell(V0, V1)) :- cell(V0, V1), cell(V2, V3), V0 == V2-1, V3 == V1-1.
113 5 ~ adj(cell(V0, V1), cell(V2, V3)) :- cell(V0, V1), cell(V2, V3), V1 == V3-1, V2 == V0-1.
114 5 ~ adj(cell(V2, V3), cell(V0, V1)) :- cell(V0, V1), cell(V2, V3), V1 == V3-1, V2 == V0-1.
115 5 ~ adj(cell(V0, V1), cell(V2, V3)) :- cell(V0, V1), cell(V2, V3), V2 == V0-1, V3 == V1-1.
116 5 ~ adj(cell(V2, V3), cell(V0, V1)) :- cell(V0, V1), cell(V2, V3), V2 == V0-1, V3 == V1-1.
117 #max_penalty(50).

```

B.3 CROSS-DOT

B.3.1 EXPERIMENT 9.3: CROSS-DOT RULES

Logic Program B.4. Examples of legal and illegal moves in Cross-Dot

```

1 #pos({legal(x, mark(1)), legal(o, noop)}, {legal(x, noop), legal(o, mark(1))}, {
2   box(1, b).
3   box(2, b).
4   box(3, b).
5   box(4, b).
6   control(x).
7 }).
8 #pos({legal(x, noop), legal(o, mark(1))}, {legal(o, noop), legal(o, mark(2))}, {
9   box(1, b).
10  box(2, x).
11  box(3, b).
12  box(4, b).
13  control(o).
14 }).
15 #pos({legal(o, noop), legal(x, mark(3))}, {legal(x, mark(1)), legal(x, mark(2))},
16   → {
17   box(1, o).
18   box(2, x).
19   box(3, b).
20   box(4, b).
21   control(x).
22 }).
23 #modeh(legal(var(role), mark(var(box)))).
24 #modeh(legal(var(role), noop)).
25 #modeb(box(var(box), const(state))).
26 #modeb(control(var(role))).
27 #modeb(role(var(role))).
28 #constant(state, x).
29 #constant(state, o).
30 #constant(state, b).
31
32 #maxv(2).

```


BIBLIOGRAPHY

- Benthem, Johan van (2011). 'Logic Games: From Tools to Models of Interaction'. In: *Proof, Computation and Agency: Logic at the Crossroads*. Ed. by Johan van Benthem, Amitabha Gupta and Rohit Parikh. Dordrecht: Springer Netherlands, pp. 183–216. ISBN: 978-94-007-0080-2. DOI: 10.1007/978-94-007-0080-2_11. URL: https://doi.org/10.1007/978-94-007-0080-2_{_}11.
- Bloekel, Hendrik and Luc De Raedt (1998). 'Top-down induction of first-order logical decision trees'. In: *Artificial Intelligence* 101.1-2, pp. 285–297. ISSN: 00043702. DOI: 10.1016/S0004-3702(98)00034-4. URL: <http://linkinghub.elsevier.com/retrieve/pii/S0004370298000344>.
- Buccafurri, Francesco, Nicola Leone and Pasquale Rullo (1997). 'Strong and weak constraints in disjunctive datalog'. In: *Logic Programming And Nonmonotonic Reasoning*. Vol. 1265, pp. 2–17. ISBN: 3-540-63255-7. DOI: 10.1007/3-540-63255-7_2. URL: http://dx.doi.org/10.1007/3-540-63255-7_{_}2.
- Calimeri, Francesco et al. (2013). 'ASP-Core-2 Input Language Format'. In: URL: <https://www.mat.unical.it/aspcomp2013/files/ASP-CORE-2.03c.pdf>.
- Cerexhe, Timothy, Orkunt Sabuncu and Michael Thielscher (2013). 'Evaluating Answer Set Clause Learning for General Game Playing'. In: *Logic Programming and Nonmonotonic Reasoning*. Ed. by Pedro Cabalar and Tran Cao Son. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 219–232. ISBN: 978-3-642-40564-8.
- Čermák, Petr et al. (2014). 'MCMAS-SLK: A model checker for the verification of strategy logic specifications'. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol. 8559 LNCS, pp. 525–532. ISBN: 9783319088662. DOI: 10.1007/978-3-319-08867-9_34. arXiv: 1402.2948.
- Dastani, Mehdi et al. (2005). 'Modelling user preferences and mediating agents in electronic commerce'. In: *Knowledge-Based Systems* 18.7, pp. 335–352. ISSN: 09507051. DOI: 10.1016/j.knosys.2005.05.001.
- Edwards, Steven J. (1994). *Portable Game Notation Specification and Implementation Guide*. URL: <https://opensource.apple.com/source/Chess/Chess-110.0.6/Documentation/PGN-Standard.txt> (visited on 09/01/2018).
- Gebser, Martin, Torsten Grote and Torsten Schaub (2010). 'Coala: A Compiler from Action Languages to ASP'. In: *Logics in Artificial Intelligence*. Ed. by Tomi Janhunen and Ilkka Niemelä. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 360–364. ISBN: 978-3-642-15675-5.
- Gelfond, Michael and Vladimir Lifschitz (1988). *The stable model semantics for logic programming*. DOI: 10.1.1.24.6050.
- Gelfond, Michael and Vladimir Lifschitz (1998). 'Action languages'. In: *Electronic Transactions on AI* 3.16, pp. 1–23. URL: <http://ssdi.di.fct.unl.pt/rcr/geral/biblio/assets/gelfond98action.pdf>.
- Grasso, G, N Leone and F Ricca (2013). *Answer set programming: Language, applications and development tools*. DOI: 10.1007/978-3-642-39666-3_3. URL: https://www.scopus.com/inward/record.uri?eid=2-s2.0-84881127729_{_}&

- }doi=10.1007/978-3-642-39666-3_3&partnerID=40&md5=d476b39aec47180e307cc402ef5b94a1.
- Gunning, David (2016). *Explainable Artificial Intelligence (XAI)*. URL: <https://www.cc.gatech.edu/~alanwags/DLAI2016/978Gunning29IJCAI-16DLAIWS.pdf> (visited on 28/04/2018).
- Heineman, George T, Gary Pollice and Stanley Selkow (2008). 'Path Finding in AI'. In: *Algorithms in a Nutshell*, pp. 213–217. ISBN: 9780596516246. DOI: 10.1093/aje/kwq410. URL: <http://www.ncbi.nlm.nih.gov/pubmed/21047818>.
- Kaiser, Łukasz (2012). 'Learning Games from Videos Guided by Descriptive Complexity'. In: *Aaai*, pp. 963–969. URL: <http://www.aaai.org/ocs/index.php/AAAI/AAAI12/paper/viewPDFInterstitial/5091/5508>.
- Kaminski, Roland (2014). *Clingo*. URL: <https://github.com/potassco/clingo>.
- Law, Mark, Alessandra Russo and Krysia Broda (2014). 'Inductive Learning of Answer Set Programs'. In: *European Conference on Logics in Artificial Intelligence (JELIA) 2*. Ray 2009, pp. 311–325. ISSN: 16113349. arXiv: 1608.01946.
- Law, Mark, Alessandra Russo and Krysia Broda (2015a). 'Learning weak constraints in answer set programming'. In: *Theory and Practice of Logic Programming* 15.4-5, pp. 511–525. ISSN: 14753081. DOI: 10.1017/S1471068415000198. arXiv: 1507.06566.
- Law, Mark, Alessandra Russo and Krysia Broda (2015b). *The ILASP system for learning Answer Set Programs*. URL: <https://www.doc.ic.ac.uk/~m11909/ILASP> (visited on 22/12/2017).
- Law, Mark, Alessandra Russo and Krysia Broda (2016). 'Iterative Learning of Answer Set Programs from Context Dependent Examples'. In: *Theory and Practice of Logic Programming*. Vol. 16. 5-6, pp. 834–848. DOI: 10.1017/S1471068416000351. arXiv: 1608.01946.
- Law, Mark, Alessandra Russo and Krysia Broda (2018). *The Meta-Program Injection Feature in ILASP*. Tech. rep. Imperial College London, p. 2. URL: <https://www.doc.ic.ac.uk/~m11909/ILASP/inject.pdf>.
- Lee, Joohyung (2012). 'Reformulating action language C+ in answer set programming'. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 7265, pp. 405–421. ISSN: 03029743. DOI: 10.1007/978-3-642-30743-0_28.
- Lifschitz, Vladimir (2008). 'What Is Answer Set Programming?'. In: *Aaai 2008*, pp. 1594–1597.
- Liu, Tie-Yan (2007). 'Learning to Rank for Information Retrieval'. In: *Foundations and Trends® in Information Retrieval* 3.3, pp. 225–331. ISSN: 1554-0669. DOI: 10.1561/15000000016. arXiv: arXiv:1208.5535v1. URL: <http://www.nowpublishers.com/article/Details/INR-016>.
- Love, Nathaniel et al. (2006). 'General Game Playing: Game Description Language Specification'. In: *Science LG-2006-01*. URL: http://games.stanford.edu/language/spec/gdl_spec_2008_03.pdf.
- Marsland, T. A. (1986). 'A review of game-tree pruning'. In: *ICCA journal* 9.1, pp. 3–19. URL: <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.67.7446&rep=rep1&type=pdf>.
- Mnih, Volodymyr et al. (2013). 'Playing Atari with Deep Reinforcement Learning'. In: *arXiv preprint arXiv:1312.5602*.

- Muggleton, Stephen H. et al. (2014). 'Meta-interpretive learning: Application to grammatical inference'. In: *Machine Learning* 94.1, pp. 25–49. ISSN: 08856125. DOI: 10.1007/s10994-013-5358-3.
- Muggleton, Stephen (1991). 'Inductive logic programming'. In: *New Generation Computing* 8.4, pp. 295–318. ISSN: 02883635. DOI: 10.1007/BF03037089.
- Muggleton, Stephen (1995). 'Inverse entailment and prolog'. In: *New Generation Computing* 13.3-4, pp. 245–286. ISSN: 02883635. DOI: 10.1007/BF03037227.
- Muggleton, Stephen and Luc de Raedt (1994). 'Inductive Logic Programming: Theory and Methods'. In: *Journal of Logic Programming* 19.20, pp. 629–679. ISSN: 07431066. DOI: 10.1016/0743-1066(94)90035-3. URL: <http://homes.soic.indiana.edu/natarasr/Courses/I590/Papers/ilp.pdf>.
- Muggleton, Stephen et al. (2011). 'ILP turns 20'. In: *Machine Learning* 86.1, pp. 3–23. ISSN: 0885-6125. DOI: 10.1007/s10994-011-5259-2. URL: <http://www.springerlink.com/content/9463m43357074631/>.
- Olah, Chris et al. (2018). 'The Building Blocks of Interpretability'. In: *Distill*. DOI: 10.23915/distill.00010.
- Ray, Oliver, Krysia Broda and Alessandra Russo (2003). 'Hybrid abductive inductive learning: A generalisation of Prolog'. In: *Proceedings of the International Conference on Inductive Logic Programming* 2835, pp. 311–328. ISSN: 3-540-20144-0. DOI: 10.1007/978-3-540-39917-9_21. URL: <http://www.springerlink.com/index/LXQFK7E5V8K3D0AV.pdf>.
- Romstad, Tord et al. *Stockfish*. URL: <https://stockfishchess.org>.
- Samek, Wojciech, Thomas Wiegand and Klaus-Robert Müller (2017). 'Explainable Artificial Intelligence: Understanding, Visualizing and Interpreting Deep Learning Models'. In: arXiv: 1708.08296. URL: <https://arxiv.org/pdf/1708.08296.pdf> <http://arxiv.org/abs/1708.08296>.
- Shortliffe, Edward H (1977). 'Mycin: A Knowledge-Based Computer Program Applied to Infectious Diseases'. In: *Proceedings of the Annual Symposium on Computer Application in Medical Care*, pp. 66–69. ISBN: 0195-4210. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2464549/pdf/procascamc00015-0074.pdf> <http://www.ncbi.nlm.nih.gov.proxygw.wrlc.org/pmc/articles/PMC2464549/>.
- Silver, David et al. (2017a). 'Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm'. In: *arXiv preprint arXiv:1712.01815*.
- Silver, David et al. (2017b). 'Mastering the game of Go without human knowledge'. In: *Nature* 550.7676, pp. 354–359. ISSN: 14764687. DOI: 10.1038/nature24270. arXiv: 1610.00633.
- Takizawa, Makoto. *Elmo*. URL: <https://github.com/mk-takizawa/elmo-for-learn>.
- Tesauro, Gerald (1995). 'Temporal difference learning and TD-Gammon'. In: *Communications of the ACM* 38.3, pp. 58–68. ISSN: 00010782. DOI: 10.1145/203330.203343. URL: <http://portal.acm.org/citation.cfm?doid=203330.203343>.
- Thielscher, Michael (2010). 'A general game description language for incomplete information games'. In: *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI-10)*, pp. 994–999. URL: <https://www.aaai.org/ocs/index.php/AAAI/AAAI10/paper/view/1727>.

- Vinyals, Oriol et al. (2017). 'StarCraft II: A New Challenge for Reinforcement Learning'. In: *arXiv preprint arXiv:1708.04782*. DOI: <https://deepmind.com/documents/110/sc21e.pdf>. arXiv: 1708.04782. URL: <http://arxiv.org/abs/1708.04782>.
- Zhang, Dongmo and Michael Thielscher (2015). 'Representing and Reasoning about Game Strategies'. In: *Journal of Philosophical Logic* 44.2, pp. 203–236. ISSN: 15730433. DOI: 10.1007/s10992-014-9334-6. arXiv: 1407.5380.