

Playing “Invisible Chess” with Information-Theoretic Advisors

A.E. Bud, D.W. Albrecht, A.E. Nicholson and I. Zukerman
{bud,dwa,annn,ingrid}@csse.monash.edu.au
School of Computer Science and Software Engineering, Monash University
Clayton, Victoria 3800, AUSTRALIA
phone: +61 3 9905-5225 fax: +61 3 9905-5146

From: AAAI Technical Report SS-01-03. Compilation copyright © 2001, AAAI (www.aaai.org). All rights reserved.

Abstract

Making decisions under uncertainty remains one of the central problems in AI research. Unfortunately, most uncertain real-world problems are so complex that any progress in them is extremely difficult. Games model some elements of the real world, and offer a more controlled environment for exploring methods for dealing with uncertainty. Chess and chess-like games have long been used as a strategically complex testbed for general AI research, and we extend that tradition by introducing an imperfect information variant of chess with some useful properties such as the ability to scale the amount of uncertainty in the game. We discuss the complexity of this game which we call *invisible chess*, and present results outlining the basic values of invisible pieces in this game. We motivate and describe the implementation and application of two information-theoretic advisors that assist a player of invisible chess to control the uncertainty in the game. We describe our decision-theoretic approach to combining these information-theoretic advisors with a basic strategic advisor. Finally we discuss promising preliminary results that we have obtained with these advisors.

1 Introduction

Making decisions under uncertainty remains one of the central problems in AI research. An agent in an uncertain world needs to select actions from the action search space — the set of all possible actions in that world. As the uncertainty increases, this task can become increasingly difficult. The number of possible actions may increase, the number of possible situations in which those actions may be applied may increase, or both. The effects of these growing search spaces are amplified as the agent tries to search further ahead. Each action on each possible world state requires more possible world states to be evaluated for future moves. This property of imperfect information domains makes tackling real-world problems extremely difficult.

Games and game theory model some of these properties of real-world situations in a more controlled environment and thereby allow analysis and empirical testing of decision-making strategies in these domains. In this capacity, games have long been used as a testbed for general artificial intelligence research and ideas. In particular, chess has a long

Copyright © 2001, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

history of use in the AI community because of its strategic complexity, and well-studied and understood properties.

Despite these advantages, chess has two significant drawbacks as a general AI testbed: the first is the success of computer chess players that use hard-coded, domain specific rules and strategies to play; and the second is the fact that standard chess is a perfect information domain.

A number of researchers have tackled the first of these drawbacks. For example (Berliner 1974) investigated generalised strategies used in chess play as a model for problem solving, and (Pell 1993) introduced *metagame*. Metagame is a system for generating new games arbitrarily generated from a set of games known as Symmetric Chess-Like games (SCL Games). Thus there is no point hand-training a computer program to play one particular game, as each game requires different strategic planning and position evaluation. A good computer metagame player has to somehow encapsulate a higher level of “strategic knowledge” than is possible in a single game such as chess.

In this paper, we address the second drawback of chess as a general AI testbed — that of perfect information. We describe a missing (or imperfect) information variant of standard chess which we call *invisible chess* (Bud *et al.* 1999).¹ Invisible chess involves a configurable number of invisible pieces, i.e., pieces that a player’s opponent cannot see. Invisible chess is thus a representative of the general class of strategically complex, imperfect information, two player, zero-sum games.

Many researchers have investigated games with missing information including poker ((Findler 1977), (Korb, Nicholson, & Jitnah 1999), (Koller & Pfeffer 1997)), bridge ((Gambäck, Rayner, & Pell 1991), (Ginsberg 1999), (Smith, Nau, & Throop 1996)), and multi-user domains (Albrecht *et al.* 1997). With the exception of Albrecht *et al.*, who use a large uncontrollable domain, all of these domains are strategically simple given perfect information.

On the other hand, invisible chess retains all of the strategic complexity of standard chess with the addition of a controllable element of missing information. Invisible chess

¹Invisible chess is in the set of Invisible SCL Games, an extension to the set of SCL Games introduced by Pell. We have chosen invisible chess as the first game to explore because of the availability of standard chess programs.

is related to kriegspiel ((Li 1994) and (Ciancarini, Dalla-Libera, & Maran 1997)), a chess variant in which all the opponent's pieces are invisible and a third party referee determines whether or not each move is valid.

In addition to the complexity of playing standard chess, a player of invisible chess must maintain the possible positions of the opponent's invisible pieces. These positions may be represented as a probability distribution over the possible squares on the chess board. Section 4 describes our design that enables a central module to maintain approximations to the invisible piece probability distributions for both players. Section 5 presents a brief analysis of the advantages of playing with various invisible pieces. Interestingly, we show that the relative value of the minor invisible pieces differs from the standard chess piece ranking. We report on the uncertainty caused by the different invisible pieces and the effect that those pieces have on the opponent's ability to play strategically. Section 6 discusses the relationship between uncertainty and the information-theoretic concept of entropy, and the application of entropy to our results. It also motivates the design of information-theoretic advisors, describes these advisors and presents results that demonstrate the efficacy of our approach. Section 7 contains conclusions and ideas for further work.

2 Related Work

Since the 1950s, when Shannon and Turing designed the first chess playing programs (Russell & Norvig 1995), computers have become better at playing certain games such as chess using large amounts of hand-coded domain specific information. Continuing the tradition of using games as a testbed for general AI investigation, (Berliner 1974) proposed a tactical analyser for chess which used strategies and tactics, but did not play as well as the existing hard-coded systems.

In answer to the success of hard-coded algorithms, (Pell 1993) introduced a class of games known as Symmetric Chess Like (SCL) Games, and a system called Metagamer that plays games arbitrarily generated from this class using a set of advisors representing strategies in the class of games. Pell did not consider imperfect information games. However, his class of SCL Games is easily extended to the class of Invisible SCL Games, where one or more of a player's pieces is hidden from their opponent.

(Koller & Pfeffer 1995) investigated simple imperfect information games with an initial goal of solving them. However, their approach does not scale up to more complex games such as invisible chess.

(Smith, Nau, & Throop 1996) wrote a bridge playing program using a modified form of game tree with enumerated strategies rather than actions, effecting a form of forward pruning of the game tree. Smith *et al.* stated that forward pruning works well for bridge, but not for chess. (Ginsberg 1996) introduced partition search to reduce the effective size of the game tree. He showed that this approach works well for bridge and other games with a high degree of symmetry. Partly because in chess pawns only move forward, and partly because of the strategic nature of the game, the same posit-

ions do not tend to occur in a significant number of nodes in the game tree,² and therefore partition search is not likely to be very effective on any variants of chess. More recently (Ginsberg 1999) included Monte Carlo methods in his bridge playing program to simulate many possible outcomes, choosing the action with the highest expected utility over the simulations. Ginsberg claims that trying to glean or hide information from an opponent is probably not useful for bridge. In contrast, we present results that show that both information hiding and gleaning can be useful in invisible chess.

(Frank & Basin 1998) and (Frank & Basin 2000) have performed a detailed investigation of search in imperfect information games. They have concentrated almost exclusively on bridge. Frank & Basin focus on a number of search techniques including attenuating the game tree which they have shown to be an NP-Complete problem, and Monte Carlo methods. Their investigation suggests that Monte Carlo methods are not appropriate for imperfect information games. This problem is exacerbated in invisible chess given the relative lack of symmetry and the size of the game tree.

In the closest work to that presented here, (Ciancarini, Dalla-Libera, & Maran 1997) considered king and pawn end games in the game of kriegspiel. Kriegspiel is an existing chess variant where neither player can see the opponent's pieces or moves. They used a game-theoretic approach involving substantive rationality and have shown promising results in this trivial version of kriegspiel. Thus far they have not applied their approach to a complete game of kriegspiel.

Kriegspiel differs from invisible chess in a number of important areas. All the opponent's pieces are invisible to a player of kriegspiel. Thus there is no way to reduce the uncertainty in the domain without reducing the strategic complexity. By contrast, in invisible chess, the number and types of invisible pieces is controllable. As all pieces are invisible, every move involves substantial increases in uncertainty. In invisible chess, a player may choose whether to move a visible or invisible piece. Finally, in kriegspiel, a player attempts moves until a legal move is performed. A third party arbitrator indicates whether or not each move is legal. In invisible chess, the player is given specific information when a move is impossible or illegal, and may miss a turn as a result of attempting an impossible move (see Section 3).

These differences make kriegspiel a substantially more complex and less controllable domain than invisible chess. Specifically, exploring the relationship between uncertainty and strategic play would be more difficult in kriegspiel. Ultimately, the pathological case of invisible chess where all pieces are invisible approaches the complexity of kriegspiel. Thus invisible chess provides a convenient stepping stone to a much more difficult problem.

To our knowledge there exist no computer kriegspiel players that are able to play a full game of kriegspiel with which to compare our system.

²In fact if the same position occurs three times in a game of chess, the game is declared a draw.

3 Domain: Invisible Chess

Invisible chess is based on standard chess with the following difference. In invisible chess we differentiate between visible and invisible pieces and define them as follows: a *visible piece* is one that both players can see; an *invisible piece* is one that a player's opponent cannot see. Thus, every time a visible piece is moved, the board is updated as per standard chess. However, when a player moves an invisible piece, their opponent is informed which invisible piece has moved, but not where it has moved from or to. This information enables each player to maintain a probability distribution of their opponent's invisible pieces across the squares on the board. In this paper, we frame invisible chess pieces (♞) to distinguish them from visible chess pieces (♞).

We define three terms for referring to moves in invisible chess: A **possible move** is any legal chess move given complete information about the board. That is, no invisible pieces are in the way of the move, and the piece can move to the desired position. An **impossible move** is an attempted move that violates the laws of chess because of an incorrect assumption as to the whereabouts of an invisible piece. For example, a player attempts to move a piece "through" an invisible piece, or attempts to use a pawn to capture an invisible piece that is not diagonally in front of it. See Figure 1. An **illegal move** is an impossible move that would not allow the game to continue. For example, a player is not allowed to move their king into check by an invisible piece. See Figure 3.

The rules of invisible chess are based on the rules of chess. The only modifications pertain directly to invisible pieces and their impact on the game. In general if a move is possible then it is accepted; if a move is impossible then it is rejected and the player's turn is forfeited; and if a move is illegal, some information regarding the reason the move is illegal is revealed to the player attempting the move, and the player must supply another move. Note that we do not allow invisible kings in this basic invisible chess, as a version of invisible chess with invisible kings would drastically modify the goal of the original game.

(Bud *et al.* 2000) gives details of all scenarios where new rules come into effect, but in summary, only the rules for the following scenarios differ from those of standard chess.

- 1: Impossible moves are disallowed, the position of the first invisible piece that caused the move to be impossible is revealed, and the player's turn is forfeited. Figure 1 shows an example of an impossible move.
- 2: If a player's king is threatened, the invisible piece causing the check is revealed, and the threatened player moves normally. Additionally, the player can infer that there are no invisible pieces between the threatening piece and their king (unless the king is being threatened by a knight). Figure 2 shows an example of this scenario. Black is told that the invisible white bishop is on g5.
- 3: If a player is in check, and attempts to move into check again or does not move out of check, any invisible pieces causing the check are revealed, and the player must supply another move. In Figure 3, white is in check and attempts to capture the invisible black knight on f1. Conse-

quently the invisible black bishop on c4 is revealed, and white must supply another move.

- 4: If a player is not in check, and attempts to move into check, the invisible piece causing the check is revealed, and the player's turn is forfeited. In standard chess, the move would be disallowed and the player would provide an alternate move. In Figure 4, white attempts to move their king from d1 to d2. However, black's invisible bishop on a5 is threatening d2; consequently the invisible black bishop on a5 is revealed, and white forfeits their turn.
- 5: If a player attempts to castle through check from an invisible piece or through an invisible piece itself, the invisible piece is revealed and the turn is forfeited.

3.1 Domain Complexity

The complexity of standard chess is well understood. Chess has an average branching factor of around 35. A typical game lasts approximately 50 moves per player so the entire game tree has approximately 35^{100} nodes ((Russell & Norvig 1995), p 123).

In the trivial case where the opponent has no invisible pieces, the game tree for invisible chess is exactly as large as that for standard chess. Once invisible pieces are introduced into the game, each node of the game tree must be expanded for each possible move for each possible combination of positions of the opponent's invisible pieces. In a game of invisible chess the player and opponent each have m invisible pieces. If each invisible piece has a positive probability of occupying n_i squares, then the branching factor is approximately the branching factor of chess multiplied by the combination of invisible squares that could be occupied as shown by the following formula:

$$\text{Branching Factor} = 35 \times n_1 \times n_2 \times \dots \times n_m \quad (1)$$

If each player has two invisible pieces, and each of those pieces has an average of only four squares with a positive probability of occupation, then the average approximate branching factor of that game of invisible chess is $35 \times 16 = 560$. For three invisible pieces each, the branching factor is around $35 \times 64 = 2240$. Assuming that the invisible pieces are on the board and moving for approximately half the game, then the complete expanded game tree for three invisible pieces each is in the order of 2240^{50} which is around 10^{150} nodes.

To make the domain even more complex, chess has virtually no axes of symmetry that allow the size of the game tree to be reduced as in highly symmetrical games such as tic-tac-toe, and Smith and Nau (Smith & Nau 1993) claim that forward pruning to reduce the branching factor of chess has been shown to be relatively ineffective.

In addition to the combinatorially expansive nature of the invisible chess game tree, in order to play invisible chess, a player must maintain beliefs about the positions of the opponent's invisible pieces.

A player using a simplistic belief updating scheme such as assuming the most likely destination will ignore many probable squares. Further, a player that uses any belief updating




Figure 1: An impossible move. The black bishop on b2 is invisible. White attempts an impossible move by trying to move the bishop on c1 to a3.

scheme that does not take into account every possible destination square when an invisible piece is moved, will lose important information about the flow of the game. Suppose a player incorrectly assumes that the opponent's invisible piece is on a certain square; if the opponent then moves a visible piece to that square, the player now has no way of backtracking, and has no information about the location of the invisible piece.

Assuming no strategic knowledge (i.e., each square that an invisible piece can move to is as likely to be visited as any other), and only one invisible piece, a player can easily maintain the precise distribution for that piece (see (Bud *et al.* 2000)). In addition, whenever any piece (other than a knight) moves, the probability distribution for each invisible piece may be updated to reflect the fact that other squares in the moving piece's path are vacant. Similarly, whenever a player's king is threatened by any piece (except a knight), visible or invisible, the probability distributions for all invisible pieces may be updated to reflect the known vacancies between the piece causing check, and the threatened king.

However, for multiple invisible pieces, the positions of invisible pieces are not conditionally independent with respect to each other, so the probability of one invisible piece occupying any particular square affects the probability of another invisible piece occupying that square. Maintaining the probability distributions of multiple invisible pieces involves combinatorial calculations in the number of invisible pieces or the storage of combinatorially large amounts of data in order to maintain the complete joint distributions of all invisible pieces over the entire chess board.

We address this problem by utilising a central module to obtain an approximation of the distribution that can be calculated quickly enough to use in real time (Section 4). The GCDMM calculates the distributions for both players, and has access to the true positions of all invisible pieces at all times. We use the true positions of other invisible pieces




Figure 2: An invisible piece is revealed to warn of the check. Black is in check because of the invisible white bishop on g5.

as an approximation to the joint distribution of all invisible pieces across all squares. Thus, the calculation of each invisible piece's move is reduced to the calculation to be performed when there is only one invisible piece. Because the resulting approximation is more accurate than the distribution that could be calculated by a real player of invisible chess, our results slightly underestimate the advantages of playing chess with invisible pieces. For the purposes of this research, using this method to maintain probability distributions allows us to focus our efforts on the effects of manipulating the amount of uncertainty inherent in these distributions.

4 Basic Design

In Section 3.1 we estimate the branching factor of the game tree for invisible chess with 2 invisible pieces to be around 560, and for 3 invisible pieces to be in the order of 2000, compared to the branching factor of standard chess which is about 35 ((Russell & Norvig 1995), p. 123). To cope with this combinatorial explosion, and the strategic complexity of invisible chess, we employ a "divide and conquer" approach. We split the problem of choosing the next move into a number of simpler sub-problems, and then use utility theory (Raiffa 1968) to recombine the calculations performed for these sub-problems into a move. We use information-theoretic ideas (Rényi 1984) to deal with the uncertainty in the domain, and standard chess reasoning to deal with the strategic elements.

This modular, hybrid approach is implemented with advisors or experts connected and controlled by a Game Controller and Distribution Maintenance Module (GCDMM). In the current implementation, we include three advisors: (1) the *strategic* advisor; (2) the *move hide* advisor; and (3) the *move seek* advisor. The GCDMM controls the game state, has knowledge of the positions of all invisible pieces and maintains distributions of invisible pieces on behalf of




Figure 3: An illegal move. White is in check from the black bishop on a5, and attempts to capture the invisible black knight on f1 where it would be in check from the invisible black bishop on c4.




Figure 4: Attempting to move into check. White attempts to move the king from d1 to d2. This is impossible because it results in the king being in check from the invisible black bishop on a5.




Figure 5: An invisible bishop each. White believes the invisible black bishop on b4 is on one of a3, b4, c5 or d6. Black believes the invisible white bishop on e2 has possible squares e2, d3, c4, b5 and a6.

Move	Strat ($W_1:1$) - Pos				EU	Hide $W_2:5$	Seek $W_3:5$	Total
	b4	a3	c5	d6				
d4c5	4	-32	378	-24	81.50	0	0.56	84.30
b2a3	4	399	-28	-24	87.75	0	0.56	90.55
a2a4	25	-32	376	61	107.50	0	0	107.50
a2a3	21	-32	376	64	107.25	0	0.77	111.10
e2h5	-24	386	374	21	189.25	0.77	0	193.10
c3e4	4	397	363	29	198.25	0	0	198.25
...								
e2f1	24	384	376	54	209.50	0.77	0	213.35
b2b3	30	395	374	55	213.50	0	0	213.50
a1b1	26	395	376	58	213.75	0	0	213.75
g1h3	33	397	374	52	214.00	0	0	214.00
e2c4	24	387	376	54	210.25	0.77	0	214.10
d1d3	31	395	377	62	216.25	-0.22	0	215.15
h2h3	31	395	376	63	216.25	0	0	216.25
e2a6	24	397	376	54	212.75	0.77	0	216.60
g1f3	34	394	376	64	217.00	0	0	217.00
e2d3	32	395	377	63	216.75	0.77	0	220.60

Figure 6: Possible Moves sorted by their relative total scores.

the two players. The GCDMM is responsible for deciding whether a move is legal, impossible or illegal and ensuring that the game progresses correctly. When it is a player's turn to move, the GCDMM requests the next move from the Maximiser for that player. The Maximiser, responsible for choosing the best move suggested by the available advisors, generates all possible boards and all possible moves, and requests utility values from each of the advisors for each move. Each advisor evaluates the possible moves, across as many boards as possible in the available time, according to an internal evaluation function. The *strategic advisor*, a modified version of GNU Chess which returns all possible moves

and their utilities (evaluated to a specified depth), gives the Expected Utility (EU) of each move from a strategic perspective.³ The EU for each move or action (A) is calculated by multiplying its utility value by the probability of the game state (X) for which the utility value was calculated,

³Note that GNU Chess has no knowledge of invisible pieces. The modifications to GNU Chess are to perform extended depth searching, and to return all moves and utilities rather than only the best move. In standard chess, a poor move may be discounted early in the search so that the search on the game tree can focus on more promising moves. However, in invisible chess, a move that is poor in one board position may be good in other positions.

and summing this across all possible game states (G) as per Equation 2.

$$EU(A) = \sum_{\forall X \in G} (Utility(A|X) \times Prob(X)) \quad (2)$$

The move seek advisor scores moves according to how much they reduce the player’s uncertainty. For example, a move that fails because it “collides” with an invisible piece removes the uncertainty as to the position of that piece. Similarly, the move hide advisor scores moves according to how much they increase the opponent’s uncertainty. For example, moving a previously revealed invisible piece will greatly increase the opponent’s uncertainty. The move hide and seek advisors are described in Section 6.

Each advisor has a weight associated with it that reflects the relative value of its advice. The Maximiser multiplies the value returned from each of the advisors by its weight (W_i) and sums these values. The move with the highest overall value is passed back to the GCDMM which implements that move and requests a move from the other player. If there are multiple moves with equal highest score, one of these moves is chosen at random by the Maximiser.

4.1 Example

Figure 5 shows a typical position in a game of invisible chess with each player playing with one invisible bishop; it is white’s turn to move. Black’s belief about the position of white’s invisible bishop is represented by the probability distribution: e2 0.2, d3 0.2, c4 0.2, b5 0.2, a6 0.2, and white believes that black’s invisible bishop has the probability distribution: a3 0.25, b4 0.25, c5 0.25, d6 0.25. White knows that the invisible black bishop cannot be on e7 as the black queen traversed e7 to get to its current position at h4. The GCDMM asks the white Maximiser for its next move. The white Maximiser passes each possible board position, i.e., the current board with black’s invisible black bishop in each of b4, a3, c5 and d6 to the strategic advisor. Figure 6 shows some possible moves across the four possible boards together with their utilities and the EU of each move from the strategic advisor’s perspective. Notice that the move g1f3 has the highest strategic EU of 217. However, moving the knight on g1 has no effect on either player’s information. If there were no other advisors present, the Maximiser would choose this move.

Some moves do assist in discovering the location of the invisible black bishop, e.g. d4c5, a2a3 and b2a3, however, none of these moves is strategically strong. For this reason, and because of the small number of squares in the invisible black bishop’s distribution, the move seek advisor has little effect at this point in the game. Later in the game, once the invisible black bishop has moved several times, many squares will have a probability of being occupied by the invisible black bishop, and the move seek advisor will have more effect.

The move hide advisor evaluates each move based on the increase in the opponent’s uncertainty. When the invisible white bishop moves, its destination from the opponent’s perspective, may be any square accessible from any square currently in the invisible white bishop’s distribution. Thus all

	1B	1N	1R	1Q	2B	2N	2R	N.I.
1B	X	61.0	65.0	35.0	17.2	41.5	43.6	89.8
1N	39.0	X	53.4	39.2	13.2	37.4	27.6	89.0
1R	35.0	46.6	X	38.2	25.4	34.4	27.4	93.0
1Q	65.0	60.8	61.8	X	43.8	41.4	47.6	93.0
2B	82.8	86.8	74.6	56.2	X	72.0	57.6	95.6
2N	58.5	62.6	65.6	58.6	28.0	X	45.8	96.2
2R	56.4	72.4	72.6	52.4	42.4	54.2	X	98.2
N.I.	10.2	11.0	7.0	7.0	4.4	3.8	1.8	X

Table 1: Win percentages for different combinations of invisible pieces.

possible destination squares need to be added to its distribution, and regardless of where the invisible white bishop moves to (unless it causes check), the opponent’s uncertainty is increased. On the other hand, when a visible piece traverses a square that has a positive probability of occupation by an invisible piece, the opponent’s uncertainty decreases as they can infer that the traversed square is actually vacant. For example, the move d1d3 tells the opponent that the square d3 is empty.

The Maximiser multiplies each move value from each advisor by the advisor’s weight (W_i) and sums over all the advisors, giving the “Total” column in Figure 6. Notice that the move with the highest overall value is now e2d3 (220.60). The extra added uncertainty has been enough in this case to make this move better than the strategic choice of g1f3.

5 The Strategic Advisor

In this section we present results for playing invisible chess with a single strategic advisor for each player. Each player played with a different configuration of invisible pieces, using only a strategic advisor to decide the next move. Each result is obtained from 500 games run with a particular configuration of invisible pieces that ended with a win.⁴ White has a slight advantage in standard chess, and the strategic advisor moves pieces differently depending on whether it is playing white or black. To remove colour biases from the results, each set of 500 games was broken into two runs of 250 games each, the invisible piece configurations were swapped between white and black for each run, and the results were averaged between the two runs.

Table 1 shows results for games played with major (non-pawn) invisible pieces against each other and against no invisible pieces (N.I.). Reading across a row, it shows the percentage of games won by the combination of invisible pieces in the row heading against the invisible piece combination of a particular column. For example, one invisible bishop (1B) won 65% of the time against one invisible rook (1R), but only 43.6% of the time against two invisible rooks (2R).

These results show that the values of several invisible piece combinations differ between invisible chess and standard chess. For example, in standard chess, a rook is considered more valuable than a bishop. However, in invisible chess,

⁴Drawn games are largely the result of repeating moves continuously. Consisting of less than 10% of games played, draws are not counted in these results.

one bishop beat one rook, and two bishops beat every other combination of invisible pieces considered including rooks. This apparent anomaly is due to the bishop's early involvement in the game, while rooks tend to come into play later. By causing uncertainty early in the game, a player with two invisible bishops has an early advantage against a player with two invisible rooks. Further analysis of these results is presented in (Bud *et al.* 2000).

6 Building Information Theoretic Advisors

A reasonable inference from the results shown above is that players with more information about the game tend to win more often; i.e., the closer a player's belief about a board position is to the true board position, the more likely the player is to play well strategically. In this section, we describe our use of information theory (Rényi 1984) to quantify a player's uncertainty about the positions of an opponent's invisible pieces (Section 6.1), present two information-theoretic advisors (Section 6.2), and discuss some preliminary results obtained with these advisors (Section 6.3).

6.1 Uncertainty and Entropy

Information theory provides a measure for quantifying information represented by a sequence of symbols. That is, using information theory we can determine the minimum number of bits required to transmit the sequence of symbols to someone else. This number represents a measure of the amount of information intrinsic to the sequence of symbols. The calculation of this number requires a distribution of the probability that any particular symbol to be transmitted will occur. Thus any probability distribution can be said to have an *entropy* or information measure associated with it. This entropy measure is bounded from below by zero, when there is only one possibility and therefore no uncertainty, and increases as the distribution spreads.

In invisible chess, given a probability distribution of each of the opponent's invisible pieces, it is possible to derive a probability distribution across all possible board positions.⁵ Thus we can calculate the entropy (H) of a set of board states together with their associated probabilities as follows:

$$H = - \sum_{\forall X \in G} Prob(X) \times \log_2(Prob(X)) \quad (3)$$

where X represents a single game state, from the set of possible game states (G), which is one possible combination of positions of the opponent's invisible pieces, and $Prob(X)$ represents the probability of that game state.

As invisible pieces move, the number of squares they may occupy increases. This leads to an increase in the number of possible board states, and therefore the entropy of the distribution of those board states, i.e., it increases the opponent's uncertainty.

⁵Assuming that the piece distributions are independent, this can be calculated by combinatorially cycling through the invisible piece positions. In our implementation, these distributions of invisible pieces are independent because of the way they are maintained (Bud *et al.* 2000).

Further examination of our results shows that the more a player moves their invisible pieces, the more impossible moves their opponent attempts (see (Bud *et al.* 2000) for details). This movement of invisible pieces and the corresponding opponent uncertainty can be summarised using entropy to quantify each player's uncertainty in a game. Figures 7 and 8 show the uncertainty of each player regarding the positions of the opponent's invisible pieces for games played with one invisible bishop against two invisible bishops.⁶ For each game, the entropy of each player's distribution of invisible pieces is summed across all moves. Thus the graphs show the total amount of uncertainty each player had to deal with over the course of each game. Figure 7 shows each player's uncertainty in the 190 games that were won by white; the solid line shows black's uncertainty, sorted by entropy, while the dashed line shows white's uncertainty in the corresponding game. Figure 8 shows each player's uncertainty in the 60 games that were won by black; the dashed line shows white's uncertainty, sorted by entropy, while the solid line shows black's uncertainty in the corresponding game. As can be seen from Figure 7, all the games won by white, except for one (game number 86) show greater uncertainty for black. The games won by black (Figure 8) are much closer in uncertainty and there are many games where white had greater uncertainty than black.⁷ This example corroborates our intuition that players with less uncertainty are more likely to win, and underpins our information-theoretic advisors.

6.2 Information-Theoretic Advisors

Following an analysis of the results in Section 5, we have implemented two information-theoretic advisors. These advisors are *move hide* and *move seek*.

Move Hide Advisor. Working on the premise that the more uncertain the opponent is, the worse they will play, the move hide advisor advises a player to perform moves that hide information from the opponent. That is, each move is scored according to its expected effect on the opponent's perceived uncertainty about the positions of the player's invisible pieces. This effect may be an increase, a decrease or no change in the opponent's uncertainty. Moves by invisible pieces that do not cause check result in an increase in the opponent's uncertainty. Moves that cause check or moves by visible pieces may cause a decrease in the opponent's uncertainty by revealing vacant squares or invisible pieces themselves. As discussed in Section 6.1, we use the entropy of the distribution of the positions of the opponent's invisible pieces as a measure of the opponent's uncertainty.

In order to calculate this entropy, the move hide advisor needs to model the opponent's distribution update strategy. Of course a player can use the real positions of each non-moving invisible piece and thereby avoid the cost of storing

⁶The many games with zero entropy for white are generally due to the invisible bishop (queen's bishop) capturing a piece on its first move and subsequently being captured without moving again.

⁷Note that the same relationship between entropy and win can be seen for games between more evenly matched invisible pieces (Bud *et al.* 2000).

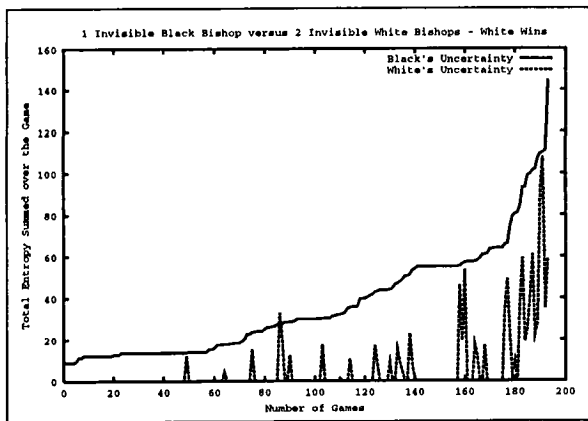


Figure 7: Uncertainty (as entropy) in games with one invisible black bishop against two invisible white bishops that were won by white.

or calculating the complete joint distribution of the positions of all invisible pieces. This method provides a model of the best distribution the opponent could have without taking strategic information into account. To incorporate strategic information into the model, a player could use the combined expected utility for each possible destination square that an invisible piece could move to.

For a proposed move, the move hide advisor uses Equation 3 to calculate the entropy of the opponent model of the distribution of the player's invisible pieces prior to and following the move. The exponential of the perceived change in entropy is returned as the move hide utility. The exponential is taken in order to allow a comparison between the log based utilities returned by the move hide and move seek advisors, and the utility returned by the strategic advisor which scores moves on a linear scale.

Move Seek Advisor. The move seek advisor suggests that a player perform moves that are more likely to discover information about the positions of the opponent's invisible pieces. That is, each move is scored according to the expected decrease in the entropy of the opponent's invisible pieces following the move. This expected decrease in entropy must be greater than or equal to zero as no move can make a player less certain about the position of the opponent's invisible pieces.

A move that covers a large number of squares with a positive probability of occupation by the opponent's invisible pieces will yield a certain amount of information whether it is successful or not. Clearly this type of move will yield more information if it is successful, as the player now knows that all of those squares are vacant. On failure, a player can only conclude that at least one of those squares is occupied. On the other hand, a move that traverses only one square with a positive probability of occupation by the opponent's invisible pieces will yield more information if unsuccessful. That is, the player now knows that an invisible piece is on that square. Thus, the move seek advisor multiplies the projected decrease in entropy for each outcome by the probability of that outcome to get an expected utility. The exponential




Figure 8: Uncertainty (as entropy) in games with one invisible black bishop against two invisible white bishops that were won by black.

Weight	Hide		Seek	
	Q vs Q	B vs B	Q vs Q	B vs B
0.5	57.2	77.2	47.2	51.0
5.0	43.6	54.6	45.0	71.6
50.0	13.6	30.3	47.6	67.0

Table 2: The effect of the move hide and move seek advisors.

of this expected change in entropy is returned as the move seek utility.

6.3 Advisor Results

This section presents and discusses some preliminary results that show the individual effects of the move hide and move seek advisors with varying weights. Each result was obtained by playing 500 games separated into runs of 250 games as before, with one player using the strategic advisor only, against an opponent using one of the information-theoretic advisors and the strategic advisor. The invisible piece configurations in this section were chosen because their results are typical of those obtained using our information-theoretic advisors.

The move hide column of Table 2 shows the results of adding the move hide advisor to an invisible queen each (Q vs Q) and an invisible white bishop versus an invisible black bishop (B vs B). The move seek column shows the results of adding the move seek advisor to an invisible queen each and one invisible white bishop versus one invisible black bishop (B vs B). The first column headed "Weight" shows the weight of the information-theoretic advisor relative to the strategic advisor. Thus, with a weight of 0.5, in games played with an invisible queen each, the player playing with the move hide advisor won 57.2% of the games.

Move Hide Results. With a weight of 0.5, the move hide advisor significantly assists, increasing the win rate to 57.2% and 72.2% for queen versus queen and bishop versus bishop respectively. However, as the move hide advisor weight increases past 0.5, the win rate decreases quite rapidly. This

behaviour is typical of all observed move hide runs, and results from the player taking less notice of the strategic advisor's advice. Although the opponent may be slightly more uncertain when the move hide advisor is weighted highly, the player is making enough strategically poor moves to counter that advantage. Nonetheless, the move hide advisor is de nitely helpful with small weight, and the results shown for a weight of 0.5 are statistically signi cantly different from the base case of 50%.

Figure 9 shows the difference in entropy between white and black, each playing with an invisible queen, averaged over the entire game for each of 250 games. To make the trends clearer, the data points are sorted by entropy. The middle curve represents games played between invisible queens with no information-theoretic advisors present, and shows that the entropy difference is fairly even. In approximately half the games it is negative, and in the other half it is positive, and it ranges between around -1 and +1. The bottom curve represents games where white played with the move hide advisor. This represents an entropy advantage to white. The difference between white and black's uncertainty ranges from under -2 to around +1. The rest of the curve is also shifted downwards indicating the relative increase in black's uncertainty. This increase in black's uncertainty leads to more wins for white while using the move hide advisor.

When an invisible piece moves, the uncertainty increases in the same manner regardless of the nal destination of the invisible piece. Thus the move hide advisor is most valuable when any strategically advantageous move by an invisible piece is possible.

Move Seek Results. Table 2 (column 5) shows the move seek advisor's effectiveness against an opponent's invisible bishop. As the weight increases to around 5.0, the percentage of wins increases up to 71.6%. This is largely a result of the move seek advisor assisting the player to nd and capture the opponent's invisible bishop early in the game. The player then has an information advantage equivalent to an invisible bishop versus no invisible pieces and is very likely to win. Figure 10 shows the entropy advantage of playing with the move seek advisor. The data points are sorted by entropy to clarify the trends. The top curve shows the entropy difference with no move seek advisor, and the bottom curve shows that the entropy is signi cantly lower when playing with a move seek advisor. In this example, the move seek advisor assists white to reduce uncertainty over the course of the game and therefore play strategically better than black.

Table 2 (column 4) shows the move seek advisor as applied to a game between two invisible queens. In this situation, the move seek advisor is much less effective and would appear to probably be detrimental in some cases. It seems likely that the large number of squares an invisible queen may have a positive probability of occupying, and the frequency of queen movement, make it dif cult for the opponent to nd an invisible queen. The top curve in Figure 9 represents games where white played with the move seek advisor. An examination of the entropy difference slightly favours black compared to the base case. This is almost certainly because white is spending moves trying to nd black's in-

visible queen rather than moving their own invisible queen. Only moves that traverse squares that have a positive probability of occupation by the opponent's invisible pieces are valued by the move seek advisor. These moves may or may not correspond to good strategic moves. The more a player listens to the move seek advisor, the fewer good strategic moves they are able to perform. The difference between the bishop-seeking behaviour and the queen-seeking behaviour depends on the dif culty of nding the opponent's invisible piece. An invisible bishop can have a positive probability of occupying at most half the available squares on the board, while an invisible queen can have a positive probability of occupying all the available squares. Thus, the advice provided by the move seek advisor often aids in the early capture of the invisible bishop, thereby removing the uncertainty from the game. In contrast, following this advice when the invisible piece is an invisible queen leads to wasted moves. Thus the move seek advisor assists when the uncertainty in the game is low, but may be useless or detrimental when the uncertainty is high.

7 Conclusion and Future Work

The results presented in this paper are preliminary and further exploration of the domain is required. There are a number of areas that require more in depth investigation. Specifically, more accurate prediction of the likely positioning of the opponent's invisible pieces is needed. This prediction could take the form of using strategic information about the likely destination of a moving invisible piece, or involve evaluating the complete search tree for more ply. Improving this distribution would reduce its entropy and therefore improve strategic performance. A side effect of this type of distribution improvement is the possibility of incorporating bluf ng into invisible chess. That is, moving an invisible piece to an unlikely position in order to confuse an opponent.

As indicated above, one way to improve the prediction of the positions of the opponent's invisible pieces would be to model the uncertainty regarding a player's pieces from the opponent's perspective for more ply. However, the problem of the combinatorial expansion in the search required as a result of this prediction needs to be resolved. The most obvious way to manage this explosion is to nd effective ways to prune the game tree.

As our system currently stands, a non-integrated player of invisible chess (whether human or machine) could not have the benefit of our GCDMM which maintains the approximation to the distribution of the positions of the opponent's invisible pieces by using the true positions of all other invisible pieces when updating the distribution. Such a player would need to nd other ways of approximating that distribution.

In summary, we have presented and discussed a complex, but controlled domain for exploring automated reasoning in an uncertain environment with a high degree of strategic complexity. We have motivated and introduced the use of information-theoretic advisors in the strategically complex imperfect information domain of invisible chess. We have shown that our distributed-advisor approach using a combi-

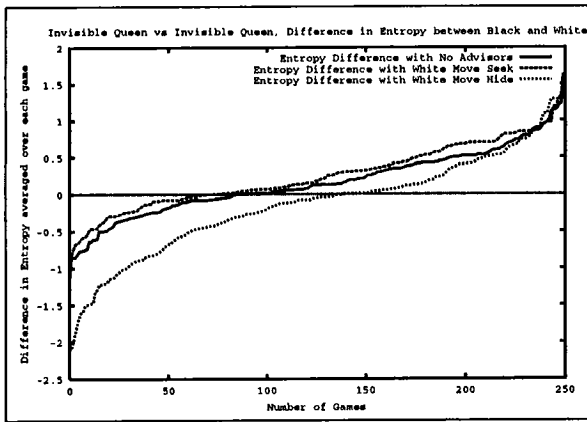


Figure 9: Entropy difference between white and black for a base case, move seek (weight 1.0) and move hide (weight 0.5), sorted by entropy.

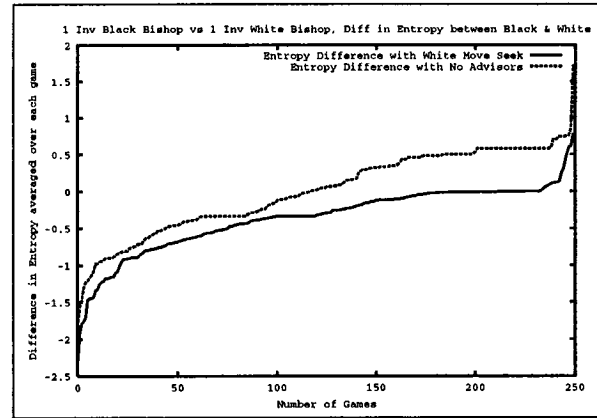


Figure 10: The difference in entropy between white and black for a base case, and with move seek, sorted by entropy.

nation of information-theoretic and strategic aspects of the domain lead to performance advantages compared to using strategic expertise alone. Given the simplicity and generality of this approach, our results point towards the potential applicability to a range of strategically complex imperfect information domains.

Acknowledgements

The authors would like to thank Chris Wallace for his insight, advice and patience.

References

- Albrecht, D. W.; Zukerman, I.; Nicholson, A. E.; and Bud, A. 1997. Towards a bayesian model for keyhole plan recognition in large domains. In *User Modelling. Proceedings of the Sixth International Conference UM97*, number 383 in CISM Courses and Lectures, 365–376. Springer Wien, New York NY USA.
- Berliner, H. J. 1974. *Chess as Problem Solving*. Ph.D. Dissertation, Carnegie Mellon University, Pittsburgh, PA.
- Bud, A.; Nicholson, A.; Zukerman, I.; and Albrecht, D. 1999. A hybrid architecture for strategically complex imperfect information games. In *Proceedings of the 3rd International Conference on Knowledge-Based Intelligent Information Engineering Systems*, 42–45. IEEE Press.
- Bud, A.; Zukerman, I.; Albrecht, D.; and Nicholson, A. 2000. Invisible chess: Rules, complexity and implementation. Technical Report forthcoming, Dept. of Computer Science, Monash University.
- Ciancarini, P.; DallaLibera, F.; and Maran, F. 1997. Decision Making under Uncertainty: A Rational Approach to Kriegspiel. In van den Herik, J., and Uiterwijk, J., eds., *Advances in Computer Chess 8*, 277–298. Univ. of Rulimburg.
- Findler, N. V. 1977. Studies in machine cognition using the game of poker. *Communications of the ACM* 20(4):230–245.
- Frank, I., and Basin, D. 1998. Search in games with incomplete information: A case study using bridge card play. *Artificial Intelligence* 100:87–123.
- Frank, I., and Basin, D. 2000. A theoretical and empirical investigation of search in imperfect information games. In *Theoretical Computer Science — To appear*.

Gambäck, B.; Rayner, M.; and Pell, B. 1991. An architecture for a sophisticated mechanical Bridge player. In Levy, D. N., and Beal, D. F., eds., *Heuristic Programming in Artificial Intelligence — The Second Computer Olympiad 2*, 87–107. Chichester, England: Ellis Horwood.

Ginsberg, M. 1996. Partition search. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, 228–233. AAAI Press / MIT Press.

Ginsberg, M. 1999. GIB: steps toward an expert-level bridge-playing program. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, 584–589. Morgan Kaufmann.

Koller, D., and Pfeffer, A. 1995. Generating and solving imperfect information games. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, 1185–1192. Morgan Kaufmann.

Koller, D., and Pfeffer, A. 1997. Representations and solutions for game-theoretic problems. *Artificial Intelligence* 94(1):167–215.

Korb, K. B.; Nicholson, A. E.; and Jitnah, N. 1999. Bayesian poker. In Laskey, K. B., and Prade, H., eds., *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*, 343–350. Stockholm, Sweden: Morgan Kaufmann.

Li, D. 1994. *Kriegspiel: Chess Under Uncertainty*. Premier Publishing.

Pell, B. 1993. *Strategy Generation and Evaluation for Meta-Game Playing*. Ph.D. Dissertation, University of Cambridge.

Raiffa, H. 1968. *Decision analysis: introductory lectures on choices under uncertainty*. Addison-Wesley.

Rényi, A. 1984. *A diary on information theory*. John Wiley and sons.

Russell, S., and Norvig, P. 1995. *Artificial Intelligence: A Modern Approach*. Prentice Hall.

Smith, S. J. J., and Nau, D. S. 1993. Strategic planning for imperfect-information games. In *Games: Planning and Learning, Papers from the 1993 Fall Symposium*, 84–91. AAAI Press.

Smith, S. J. J.; Nau, D. S.; and Throop, T. A. 1996. Total-order multi-agent task-network planning for contract bridge. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, 108–113. AAAI Press / MIT Press.