
Contract Bridge: Multi-agent Adversarial Planning in an Uncertain Environment

Deepak Khemani
Shikha Singh

KHEMANI@IITM.AC.IN
CS16D008@SMAIL.IITM.AC.IN

Department of Computer Science and Engineering, IIT Madras, Chennai 600036, India

Abstract

First there was Samuel’s Checker program. Then came IBM’s Deep Blue, and more recently AlphaGo program from Deep Mind, and now the Poker program Libratus from CMU. One by one over the last twenty odd years computer programs have become better than humans at many games. And yet the citadel of Contract Bridge remains. We argue that Bridge is significantly and qualitatively different. So far, in the games where AI has done well, it was one on one, man versus machine. But Bridge is a team game, and with incomplete information to boot. Partners in bridge need to communicate, and where there is communication there is interception, and where there is interception there is the possibility of deception. A bridge player has to be an epistemic agent that needs to reason with knowledge, belief, plans, and probabilities. We survey the cognitive landscape, and chart out strategies for human like performance. We hope that this will also give us insights into building man-machine collaborative systems.

1. Introduction: Why Bridge?

Games, like Checkers, Chess, Backgammon, Scrabble, Go, Poker and Bridge, that have fascinated humankind have long been considered as test beds for Artificial Intelligence programs. This is not surprising. On the one hand these games have always been associated with intelligence, and pose problems that are inherently complex. On the other hand, games provide platforms where the interface with the external world is intrinsically digital and poses no challenges of perception and action, and they are domains where success or failure is simple to evaluate.

Of the games mentioned here, all but Contract Bridge have seen machines outperform humans. The last major citadel to fall was the game of Go, when the world champion Lee Sedol was beaten in 2016 by the program AlphaGo from DeepMind. In this paper we look at the game of Contract Bridge. We argue that the game of bridge is different from the games where success has been achieved so far. Unlike board games described below, bridge is a multi-faceted game requiring a multi-pronged approach, calling upon the entire cognitive faculties of *Memory*, *Reason*, and *Imagination* as described in the tree of Diderot and d’Alembert ¹.

Checkers, Chess and Go are two-player complete-information zero-sum games, as is the smaller game of Tic-tac-toe. Conceptually they are all simple and similar, and can be represented by a

1. https://en.wikipedia.org/wiki/Figurative_system_of_human_knowledge

game tree with alternate layers representing the choices for the two players. The difficulty, or the complexity arises from the size of the game tree. Chess has an estimated average game tree with 10^{120} leaf nodes, and it is not surprising that the minimax value of the game is not known, as it is for Tic-tac-toe (the game ends in a draw) and even Checkers (Schaeffer et al., 2007). Go has a much larger game tree estimated to be upwards of 10^{360} (Allis, 1994). Even though Chess and Go have not been solved, they have computer programs that beat the best amongst humans. IBM’s Deep Blue beat Garry Kasparov in 1997 (Campbell et al., 2002) and AlphaGo beat Lee Sedol in 2016 (Silver et al., 2016). These programs were able to do this, not by searching the entire game tree, but, by searching it selectively aided in the process by an evaluation function. The evaluation function itself is a product of human expert knowledge, encoded by humans and fine tuned by machine learning algorithms. Then in 2017 the program AlphaGo Zero (Silver et al., 2017b) was able to learn Chess and Go entirely from scratch without the aid of human knowledge, simply by repeatedly playing against itself, and the self learnt versions are much stronger than their predecessors (Silver et al., 2017a).

Scrabble and Backgammon are also board games, but they have an element of chance. There is uncertainty about the options available for the opponent. In Scrabble each player draws tiles from a fixed set of alphabet tiles, and the tiles drawn are hidden from the opponent. In Backgammon a pair of dice determines what moves a player can make. And yet computer algorithms have been devised that are much better than the best humans (Sheppard, 2002; Tesauro, 1995). Tesauro’s Backgammon player benefited from a self training neural network that employed temporal difference reinforcement learning to learn the evaluation function.

In 2017, the program Libratus from CMU (Brown & Sandholm, 2017) playing a two player limited version of Poker beat four leading players. Poker is an incomplete information game, and among other things one is required to bluff effectively, and also call opponent’s bluffs. Libratus employs a notion of card abstraction that is cognitively similar to the way bridge players think about cards.

In contrast to the above games, on the table Contract Bridge is a four-person game with opposing teams of two each. Each game begins with a pack of cards dealt out to the four players. Every deal played has a different, almost unique, starting position, unlike board games which have only one. Moreover, when an individual player is dealt a hand of thirteen cards, the holdings of the other three players are unknown to her. The battle that ensues involves many forms of reasoning. In one sense the task undertaken by each player is to cut through the fog of uncertainty so that the one knows the other players’ cards as much as is needed to succeed. In another sense the task is to plan for all possible worlds that remain in contention, so that one succeeds in as many of them as possible. In the midst of all this activity and opponent may try to obfuscate matters by *creating* the possibility of imaginary worlds to lead you astray. Above all bridge is a game of imagination, as much as it involves the traditional AI tools of deploying search, knowledge and memory.

In the next section we describe the game of bridge and its two phases. In the following section we do a quick review of the literature on Computer Bridge, and observe that most current programs rely on Monte Carlo methods. We then revisit the different tasks a bridge player is confronted with, and outline our approach to playing the game, describing the representations a cognitive system

creates, the knowledge retrieved from memory to reason with those representations, and the role of imagination in combating intelligent opponents in an incomplete information scenario.

2. Contract Bridge

We treat Contract Bridge as a concrete instance of a multi-agent game scenario where the players attempt to maximize their payoffs in resource bounded epistemic situations in which players have incomplete information about the resources and moves available to other agents they are cooperating or competing with. A key feature of such situations is the need for reasoning about the actions of other agents, and careful communication between collaborators using a public channel. Such games can also be models for protracted negotiations in different scenarios, like financial contracts, and for adversarial scenarios like troop movements during a war. The game captures the complex reasoning required in such situations, and at the same time provides a small manageable domain.

Contract Bridge is a bounded horizon game played by multiple players with limited private resources, drawn randomly from a fixed total pool. A team in a bridge game is composed either of two players, in a *pairs* event, or four, in a *team-of-four-duplicate* event. In a tournament many teams participate, and the goal for each team is to win the tournament. In all cases the final payoff for the players depends upon a series of *matches*, and the score for each match depends upon the score of each deal played on the table. On each deal on the table the general strategy is generally to maximize the payoff. This is not as straightforward as winning a chess game, because while there is the notion of a *par* score for each deal, the name of the game is to try and consistently beat the par. This is possible because of the incomplete information of the game, and also because the payoff is determined not just what happens on that table, but also on what happens on other table (or tables in a Pairs format). Furthermore, since the final payoff is a non-linear function of all matches (except when it is knock out format), and the match score a non-linear function of the score on each deal, players have to often formulate deal level goals in a context dependent manner, based on their estimate of how the tournament has gone so far. This makes tournament play complex enough. Playing a single deal has its own unique complexities as we observe below.

On the table, the one that is in our scope, a team of two players, which we will call *South* (S) and *North* (N), competes against the opposing team whose players we call *West* (W) and *East* (E). The game proceeds in a sequential manner in the clockwise direction. The game is played with a standard pack of 52 cards, which the *dealer* deals clockwise to the 4 players, 13 cards each. The 4 players become dealer turn by turn. The pack is made up of 4 *suits*, each of which is an ordered set. The *Ace* is the highest card with rank 1, followed by the *King*, *Queen*, *Jack*, and the rest in decreasing numeric value, the 2 being the lowest.

The *play* constitutes of 13 rounds, in each of which the players play a card on their turn. The first player can play any card, and the rest three have to *follow suit*, playing a card of the same suit, unless they do not have one in which case they can play any card. Each round is called a *trick*. The card with the highest rank wins the trick, and the player who played it gets to start the next round.

The general objective is for each team/side to win as many tricks as possible, since the payoff is proportional. What makes the game complex and interesting is the notion of a *contract*. The contract is a *declaration* by one side to make a certain number of tricks. It is decided by an *auction*

described below. The contract may notify a suit as the *trump suit*, in which case even the smallest card of the trump suit becomes higher than any card of any other suit. Further complexity is added by tailoring the payoff to the final contract bid. Certain landmark contracts – a *game*, a *small slam*, and a *grand slam*, with increasing number of tricks contracted – have additional bonus payoffs. This makes them attractive goals to bid for. But the payoffs are received *only* when the contract is fulfilled. If the contract fails, the opposing side, called *defenders*, get a (different) payoff determined by the number of tricks the contract fails by.

The contract itself is decided in an auction or *bidding* phase of the game that precedes the play phase. Thus the game is a two-stage game. In the first stage, the auction, each player makes a bid in turn, and bidding ends when the following three players have nothing further to say and pass consecutively. This is followed by the play phase, in which one of the declaring side players, the *dummy*, exposes his cards, *thus making more information available* to each of the other three players to reason with.

The auction begins with the *dealer* making the first bid. The minimum bid that one can make is for 7 tricks and the maximum is for 13. Given the 5 denominations made up of the 4 suits and a *No Trump* bid, there are 35 different possible contracts to bid for. This number is doubled to 70 by having each bid challenged by a bid called *Double*, and further doubled to 140 by a counter challenge called a *Redouble*. In addition a player can make a *Pass* bid. Since 2 Pass bids may precede each of the 140 bids, the length of the auction can be about 420 bids in the longest case (and 4 bids in the shortest case). This *space* allows considerable opportunity for conveying useful information, as we shall describe below.

This brief description of the general structure of the contract bridge game can give us some hints on the different kinds of reasoning that may be involved.

3. Related Work

There has been sporadic interest in programming for bridge, but much of it has been focused on converting the play problem into a complete information problem, known as a double dummy problem, and then employing game tree search based techniques. To account for the hidden hands, Monte Carlo sampling techniques are employed. We must emphasize that our aim is not to do so and instead explore the human approach of deploying memory, reason and cognition. In literature the closest work we found is on planning over multi-agent epistemic states (Muisse et al., 2015), where the authors state “*In the absence of prescribed coordination, it is often necessary for individual agents to synthesize their own plans, taking into account not only their own capabilities and beliefs about the world but also their beliefs about other agents, including what each of the agents will come to believe as the consequence of the actions of others.*” The only difference is that they adopt a classical planning approach on a much simpler problem, while bridge requires a multipronged approach for planning.

Most bidding programs use a rule-based approach (Carley, 1962; Lindelof, 1983; Wasserman, 1970) in which a certain bidding decision is made if the condition of the corresponding bidding rule match with auction patterns and constraints on the possible hands of the other players. A combined approach of rule-based system and lookahead search was used (Gambäck et al., 1993; Ginsberg,

2001) in order to make decision in states that matched either two or more rules or no rules at all. One of the earliest attempts at declarer play employed knowledge of specialized play techniques, called *Thematic Acts*, which were stored in the form of rules and a planner attempted to retrieve relevant ‘plays’ and string them together (Khemani, 1989, 1994). Another program, Bridge Baron by Smith et. al (Smith et al., 1996, 1998a) used Hierarchical Task Network (HTN) planning techniques to approach declarer play by generating and evaluating game trees whose branches represent the number of tactical schemes that an agent can use, and not the number of moves that agent can make as in conventional game trees, in the current state of play.

Bridge Baron augmented with Smith et. al’s HTN techniques won the 1997 World Computer Bridge Championships (Smith et al., 1998b), nonetheless, its performance was worse than that of an amateur bridge player. Ginsberg’s Intelligent Bridge player (GIB) (Ginsberg, 1999) used search based techniques- Monte-Carlo simulations in card playing and Borel simulations (Monte-Carlo like simulations) in bidding. Multiple instances of bridge’s perfect information variant, Double Dummy, consistent with bidding and previous card plays are analysed to suggest next play (first introduced by Levy (1989)). Search reduction techniques such as alpha-beta pruning, transposition tables and move ordering heuristic brought down the branching factor of search trees from 4 to 1.3, and partition search (Ginsberg, 1996) further reduced the search space to approx 18,000 nodes per deal. GIB also used Iterative broadening (Ginsberg & Harvey, 1992) to return a low width answer if a high width search fails to terminate in time.

However, Monte-Carlo approaches, as shown by Frank and Basin (Frank & Basin, 1998), do not encourage information gathering actions such as *discovery play* and tend to defer decisions until next round of play. To solve the problem of *strategy-fusion* and *non-locality* in sampling algorithms like MC simulations, (Frank et al., 1998) formalised *vector minimaxing* and *payoff-reduction minimaxing algorithms*. Frank demonstrated that *prm* algorithm outperforms the other two on finding optimal strategies. A slightly different approach for decision making and cooperation in bridge bidding was proposed (Amit & Markovitch, 2006) which employed *model-based Monte Carlo sampling* method to model partner(s)/opponents(s) information states and presented a learning framework that allows co-training of partners on a training set of various classes of states with conflicting actions. (DeLooze & Downey, 2007) used a combination of two self-organizing maps (SOMs) to find an optimal bidding strategy for no trump bridge hands. All the above work exploits human designed features for human bidding system but a recent attempt (Yeh & Lin, 2016) to automatically learn the bidding system directly from raw data was made using deep reinforcement learning, but it is yet to prove its mettle.

The currently successful bridge playing programs are written using a customization of the Monte Carlo and double-dummy techniques proposed by Ginsberg (1999) and most of them are regular entrants in the WCBCs. Since 1996 the American Contract Bridge League has been organising the official World Computer Bridge Championships annually, held along with international bridge events like ACBL NABCs, World Bridge Federation World Championships, and the first European Bridge Federation Open Championship. The 2017 event, the 21st World Computer-Bridge Championship, was held, August 19th - 24th, at the 43rd World Team Championship, Lyon, France and witnessed

the participation of Wbridge5 ², Shark Bridge ³, Micro Bridge ⁴, Q-Plus Bridge ⁵, Bridge Baron ⁶, RoboBridge ⁷ and Synrey Bridge ⁸.

4. The Possible Worlds

In board games like Chess and Go there is no uncertainty and there is only one possible world, visible to all. In Poker against one player the opponent could be holding any 2 or 4 cards from the remaining pack, which can be done in ${}^{45}C_2$ or ${}^{43}C_4$ ways. At various stages 3 or 4 or 5 cards are made visible to the players. In Backgammon the uncertainty arises due to the throw of the dice, and in Scrabble due to the draw of letter tiles. In both there is a board visible to all.

Unlike Chess and Go, in card games the initial position is different for each game when the cards are dealt. In Bridge, 13 cards are dealt initially to each player. We construct possible world models, or Kripke structures, in the style of the 3-card example used in Dynamic Epistemic Logic (van Ditmarsch et al., 2007) and specify in each possible world exactly which 13 cards are held by each of the four players. Then, the number of possible initial position, or deals is ${}^{52}C_{13} \times {}^{49}C_{13} \times {}^{26}C_{13} \times {}^{13}C_{13}$. This number, approximately 53×10^{27} or 53 octillion, is obviously too large to be handled with model checking.

However, our interest is to confront the possible worlds from the perspective of *one player*, who can see her own cards. The number of possible worlds for this player is ${}^{49}C_{13} \times {}^{26}C_{13} \times {}^{13}C_{13}$ in the bidding phase, still a dauntingly large number. During play, each player can also see the 13 cards of the dummy, and this number reduces to ${}^{26}C_{13} \times {}^{13}C_{13}$ which is about 10^7 . The current approaches employing Monte Carlo methods sample these possible worlds to create double dummy problems, in which there is only one possible world, and then solve each sample with game tree searching methods. The sampling can be aided by additional constraints derived from information inferred during play, but this approach is still in its infancy. While does work, it does so only in some of the deals. It is also not cognitively appealing, and there is no possibility of generating explanations and justifying the programs actions.

Human bridge players employ a mix of possible worlds and probabilities, as we will show below. But this reasoning is only to evaluate a plan that is constructed. Very often one has to do *teleological reasoning*, and make certain assumptions about the possible worlds before embarking upon planning, because only in those possible worlds would a plan succeed. An extreme example of such reasoning is in the following four card end position.

North: ♠ A, 4, 3, 2
South: ♠ Q, 10, 9, 8

2. wbridge5.com
3. sharkbridge.info
4. www.osk.3web.ne.jp/ mcbridge/
5. www.q-plus.com
6. www.bridgebaron.com
7. www.robobridge.com
8. www.xinruibridge.com

Needing four tricks in the given end position, the *only possible world* in which a plan (play the Q from South hand) would work is when the East player has the singleton J and West has the remaining spades.

Our current work deals with the play of the hand, but we also propose how to address the bidding phase. The basic idea is create a representation in which what is *known* about the different hands is explicit. To deal with complexity we often represent insignificant cards, and what is unknown, by abstraction of variables, where a ‘x’ is a logic like variable that stands for a small and/or insignificant card, as is done in bridge literature; or by creating a *few* clusters of possible worlds pertaining to significant cards, to reason with. The clusters are determined in a teleological manner during the planning process, and the task is to cater to either all distinct clusters, or as many as them. We will present another example of teleological reasoning combined with probabilistic and epistemic reasoning later in the paper.

Building a human like bridge playing program will demand a complex problem solving architecture where different kinds of reasoning will need to be integrated. Some of the processes needed are - planning and plan recognition, encoding and decoding of information, planning for communication, interception of communication along with deliberate disinformation and deception, reasoning with probabilities, counter planning, abduction and reasoning about intentions, epistemic reasoning about knowledge, opportunistic planning, and recognizing Pareto like optimal scenarios - all in the face of incomplete information.

The best possible contract is the one that maximizes the team’s payoff (usually measured in isolation for that deal). The actual payoff is determined by (a) the score obtained by a pair on the table and (b) the results from the other table(s). The score of a deal is a function only of (a) the contract reached and (b) the number of tricks made during play. Of this the first is decided in the bidding phase, while the second is decided during play. Both stages of play happen with incomplete information.

5. Bidding

The epistemic battle begins with bidding. The goal of each side in bidding is to try and discover the *playing strength* of the given deal, and try and arrive at the best contract from each side’s point of view. The best contract is dependent of the combined strength of the two hands, which both players try to estimate. The essence of bidding is thus communication.

Given the rules of bidding, the level of the contract becomes higher and higher as bidding proceeds. One side has to eventually let the other win the final contract, because bidding higher would entail losses greater than the payoff the other side is likely to get from their bid. A side may get a positive score either by bidding and making a contract, or by defeating the opponents’ contract wherein they extract a penalty. Often a player has to decide which one, the bonus or penalty is better, and decide to bid on or double the opponents’ bid. The latter decision is taken *opportunistically*, since in the first place the other side has to bid. Contrariwise the side with the weaker hands has to decide whether to let the opponents play in what they have bid, or to make a sacrifice bid (bidding higher even when likely to fail) yielding a penalty lower than what the opponents would otherwise get.

5.1 Syntax and Semantics

The vocabulary of bids is as follows:

- *Pass*. The player has nothing to say.
- nY where n is an integer in the range [1..7] and Y is an element from the ordered list (C, D, H, S, N) which stand for (Clubs, Diamonds, Hearts, Spades, No-trump). The *contractual meaning* of the bid nY is that the bidding side promises to win $n + 6$ of the 13 tricks in the deal with Y as trumps. The *semantics* of a suit being trump becomes operational during the play phase. In a no-trump contract no suit is the trump suit. Each suit has a value associated V with it. Clubs and diamonds are called *minors* while hearts and spades are called *majors*. V is 20 for minors, and 30 for the remaining three. The *value* of the contract is nV for suit contracts and $nV+10$ for no-trump contracts.
- *Double* is a bid that can be made after an opponent has made a bid nY . Its contractual meaning is that it doubles the value of the contract. Very often it is also used to convey some specific information. A doubled contract is denoted by nY^* .
- *Redouble* is a bid that can be made by the side that bid nY *in response* to a Double. Contractually it doubles the already doubled value of the deal. A redoubled contract is denoted by nY^{**} .

The scoring system defines the payoff. For a side successfully bidding and making a contract there are increasing bonuses for *part scores* ($value < 100$), *games* ($value \geq 100$), *small slams* (12 tricks) and *grand slams* (13 tricks). Each has an associated *added* bonus. If a side fails to fulfill the contract the other side gets a score that is proportional to the deficit. In addition all these scores are modulated by bids like *double* and *redouble*, and by the context defined by the state being *vulnerable* or *non-vulnerable*.

5.2 Bidding Systems

The rules of the game define the *contractual meaning* of a bid. To estimate the best contract each player communicates some information of her hand to the partner. This communication is done via an *encoded meaning* of the bid. A *bidding system* encodes information into bids, and various bidding systems have been devised. It must be noted that this encoding is not private between the two players (that would be cheating) but it must be explained on demand. This is also done via a *convention card* that each pair has to fill.

One must remember that the contractual obligation is *only* for the final bid contract. Bids made early in the auction are generally used to convey specific features of the players holding. The two main features of interest are lengths are suits and the number of high cards. Longer suits, when a side has eight or more combined length, are suitable for being designated trumps. High cards can win tricks, and the most common method is to assign high card points (HCPs) to the honour cards, ace 4, king 3, queen 2, and jack 1. Thus a pack has 40 high card points, and the side with more than 20 is *generally* stronger. Other features include the quality of suits (length and strength), singleton suits and voids (limiting the losers in a trump contract), specific counts of aces and kings, the shape of the hand (for example two suiter hands, that have lengths in two suits) and so on. The bidding activity involves exchange of such feature information. But it goes beyond incremental description.

As we see below, a bidding system can also assign a specific function to a bid, for example *asking* for a specific feature.

A bidding system is a language in which the allowed bids are the words (or sentences) and the features are the meaning. Different bidding systems are akin to different natural languages used by human beings to talk about the same things. Natural bidding systems have bids that “mean what they say”, while in an artificial bidding system the meaning may have nothing to do with the suit mentioned in the bid. Further, in all systems, the encoded meaning of a bid is *context sensitive*, that is the meaning depends on the sequence of bids made earlier. Different systems emphasize on different features, since the goal is to convey as much as possible in the limited bidding space. Designing informative bidding systems has also received considerable attention. The quest is for one that would allow the players to meaningfully communicate in a wide variety of situations.

There are 35 possible bids starting from 1C and the last being 7N. There are many bidding systems used by bridge players around the world. A commonly used system is the Standard American Yellow Card (SAYC)⁹. For example the bidding sequence 1N_S, P_W, 2C_N, P_E, 2D_S, P_W, 3N_N, P_E, P_S, P_W has the following exchange of information in SAYC. South_{1N}: *I have a balanced hand with 15-17 high cards points.* North_{2C}: *I have a supporting hand. Do you have 4 cards in Hearts or in Spades?* South_{2D}: *No, I don't have 4 cards in either.* North_{3N}: *In that case let us bid for the 3N game contract.*

We can categorize the contracts into the bands defined above - *part-score*, *game*, *small-slam*, *grand-slam* - based on the bonus. The primary goal of the bidders is to first identify the band in which the deal falls, and then the actual contract (suit and denomination) that is the goal during play. The basic process or algorithm is concerned with finding the highest band and the suitable contract that is likely to make. Each bidder must,

- create a picture of partner's hand by using the information conveyed by partner in the encoded bids, and also any inferences that one can make from all bids made by everyone so far.
- create a picture of their combined hands. This follows as a corollary.
- create a picture of opponent hands. This is similar to creating pictures of partner's hands.
- create a picture of her own hand as conveyed by her bids. This is crucial because it determines what else needs to be communicated.

The decision of what the final contract is based on the picture one created of the combined hands. Observe that the picture created by the two partners is different, because each knows her own hand completely. Very often the partner who has greater strength - high cards and/or playing tricks given a trump choice - takes *control* of the bidding, seeking specific information from her partner.

Bridge players categorize bids by the functional role the meaning plays during communication. A *descriptive* bid conveys information about the hand. A player should compare her hand with the hand her bids have shown, and try to encode the difference into the next bid. Very often careful planning is required where future bids (or the rebid) are already planned. A *sign-off* bid says that the player thinks that it is the final contract. A *preemptive* bid attempts to consume bidding space so that the opponents cannot exchange information. It is usually made when a player thinks that the

9. <http://www.acbl.org/tournaments\textunderscorepage/general-information/convention-cards/>

opponents have better cards. An *invitational* bid is made by a responder suggesting (and inviting) the first bidder to bid ahead if partner has a better hand. A *forcing* bid is a bid that tells partner that she should bid further. It may or may not convey information. An *asking* bid is a special kind of forcing bid that asks for some specific information for example the number of aces held. Other forcing bids include a *transfer* bid asking partner to bid a specific suit, a *take out* double, and a *cue* bid which explores a slam by showing *controls* (aces or kings). A *lead directing* bid (could be a double) is a bid indicating to partner what suit to lead during play.

Thus the bidding process involves information exchange in different forms. One can describe ones hand, ask about the partner's hand, respond to questions, make a bid to help decision making later in play, consume bidding space, decide when to hide information and when to use disinformation, to encourage or discourage further communication. A bidding program has to embody a high level strategy to choose amongst the above bid making goals, use the strategy to make appropriate bids by interpreting a given bidding system (and it should be able to use different bidding systems), and actively create a picture of the other three hands which in term will feed into the decisions needed for future bids during bidding.

The high level algorithm for making a bid must take the following information into account and then make a bid that will either convey some encoded meaning or decide the contract, or both. The bidding must end only when the players feel that the desired contract has been bid. Many a player has been left red faced on the table when the bidding ends before that!

- What is known about the two hands?
- Do we know the range (game/slam/part score)?
- Are we within the bidding space?
- Has partner made a forcing bid?
- Has partner responded to the forcing bid?
- Has opponent interfered (giving partner another chance to bid)?
- Has opponent offered a penalty opportunity?

One must also keep in mind that opponents are listening, even when they are not competing (because they do not have high cards). Any information revealed is also available to them, and they may profit from it by make better choices during play. They also sometimes have the option of throwing a spanner into your communication process by “eating up” space, specially early in the bidding phase. Such bids, called preempts, can be tactically very effective if backed up by a long suit, which would limit the penalty in case one is allowed to play there.

Some of the work in bidding is aimed at encoding a bidding system into a program (Amit & Markovitch, 2006; DeLooze & Downey, 2007; Ginsberg, 1999), while others try to learn a bidding system from labeled examples (Yeh & Lin, 2016). We advocate a more cognitive approach with plans to exchange information about the two hands and at the same time try and identify the best contract, all the time keeping the bids within a safe zone since eventually a bid will become a contract, and a positive payoff is desirable. The different bidding systems are like different languages whose semantics is defined in terms of the content the player wants to communicate. In our view the program for bidding should be concerned only with the semantics or content generation. The actual bids made, which depend on the bidding system, are the akin to surface level realization in natural language generation. Ideally a program should be able to read a convention card and plan the

bidding as allowed by the system. An added advantage of this approach would be that the program would be able to understand the bids when an opponent uses a different bidding system.

It is commonly accepted amongst expert bridge players that the bidding phase is by the harder of the two in game.

6. Knowledge Based Planning for Declarer Play

The basic unit of play is to play an eligible card. This is the action that is used by Monte Carlo methods after creating a double dummy sample. While such programs have been doing well, they provide no insights to an observer.

A bridge player, on the other hand, thinks in terms planning with larger patterns. Once the opening lead is made, and the dummy comes down, she takes stock of the situation, in the manner a general would before a battle. The high level planning is done as follows..

1. Count the top tricks, and compute the deficit. This can be done by a simple search.
2. Evaluate the leeway in terms of the tempo available to develop the needed tricks. This can be done by consulting a memory of suit combination plays.
3. Choose a set of suit combination plays that could generate the tricks. The choice is made based on the probability of the required lie of opponents cards.
4. Validate the complete plan. If not feasible, go back to step 3.

6.1 Thematic Acts

As described above, a declarer thinks in terms of combinations of cards. Bridge players have evolved a vocabulary to talk about patterns and associated sequence of cards to be played. For example *cashing* a trick means playing a winner card from one hand, and (usually) an insignificant card from the other hand. *Ducking* refers to playing small cards from both hands with the intention of letting opponents win the trick (often a battle has to be lost for a war to be won). A simple and common pattern used in planning the play of a hand is a *tenace*, illustrated in Figure 1(a). The figure illustrates the layout in a suit, say Spades. North has the rank 1 card (Ace) and the rank 3 card (Queen). The rank 2 card (King) is either with East or with West.

North: A Q 4 3 2	<u>TA-finesse</u> <u>Pattern:</u> A tenace <u>Play:</u> Play small card from the South hand IF West plays the rank 2 card play the rank 1 card from North hand ELSE play the rank 3 card from the North hand <u>Outcome:</u> Generates an extra trick with the rank 3 card. <u>Conditions:</u> Succeeds if West has the rank 2 card.
West: ? East: ?	
South: 8 7 6 5	
(a)	
	(b)

Figure 1. (a) A tenace is a pattern made up of rank 1 card and a rank 3 card, with the rank 2 card hidden with one of the opponents. (b) A finesse is a card combination play associated with a tenace

Associated with a tenace is a *Thematic Act* (Khemani, 1989) called a *finesse*, described in Figure 1(b). The attempt is to win a trick with the rank 3 card before the opponent can win a trick with the

rank 2 card, which is said to be *finessed*. In the given example it works as follows. South plays a small card (say 5) and waits to see what West plays. If West plays the King then North plays the Ace, else North plays the Queen. We have *finessed* the King! That is, won a trick with the Queen, a card lower in rank. In bridge literature one would refer to the play in short as “finesse the queen”.

A finesse can be played whenever a tenace exists. The play succeeds *only* if West has the king. One can compute the probability of success of this play, and use the value to compare with competing plans. With no other information the probability that either player has the king is 50%. But as one gets more information this probability can be revised *aposteriori*. For example if West has indicated that she has many HCPs then it is more that she holds the king. Likewise if East is known to have more spades than West, then the king is more likely to be one of those cards. Computing this a *posteriori* probability on the fly is an interesting problem.

If the high level planner requires 5 tricks from the given card combination in Figure 1(a), then the finesse would be the right play. But in addition the suit must *break 2-2*. This requirement will reduce the probability of success. If on the other hand only 4 tricks are required then a *safety play* is best, in which first the ace must be *cached*, and then one must play *small to the queen*. This succeeds when the king is with West, and it *also* works when East has the king *singleton* and West has J 10 9. Planning in bridge is replete with such context dependent choices, balancing risk and reward.

A bridge player has many such thematic plays in her repository, and the more expert a player is the larger the repository. A specialized class of plays called *squeezes* has been implemented in a logic program called *Python* (Sterling & Nygate, 1990). The high level algorithm is to evaluate the situation; investigate different thematic plays; construct a set of feasible plans with associated probabilities of success; select the best plan and execute it; keep monitoring the plan to spot an opportunistic situation or signs of failure. The play phase in contract bridge is essentially a battle in which each side tries to get their tricks before the other one does. There are only 13 tricks in a hand and only one side will succeed in their goal. Outmaneuvering the opponents is the prime motive, and all is fair as long as it is not illegal.

Assuming that the top tricks are not sufficient to fulfill the contract, or defeat it from the opponents’ perspective, then both sides try and develop tricks to overcome the deficit, before the other one can. This involves complex reasoning involving Control and Tempo, because the side which develops tricks faster, and can regain control, will succeed. The repertoire of thematic plays is a large one. Some of the names of such plays found in literature are – finesse, simple squeeze, double squeeze, criss-cross squeeze, dummy reversal, cross-ruff, end-play, and so on. A TA is like a high level action or a macro action that needs a pattern of cards, and some preconditions. The preconditions are about the hidden hands. They have probabilities associated with them. However these probabilities vary dynamically as more information becomes available.

6.2 Explainable Plans

If bridge programs are to play with humans, and not just against them, then there needs to be the possibility of analyzing and discussing the play. This has to be in terms of the abstract vocabulary that bridge players’ use, along with the mutual knowledge of high level plays with card combinations. We have taken the first steps, in implementing such a knowledge based program, written in the forward chaining rule-based language OPS5 (Khemani, 1989). After a hiatus of twenty odd

years, and with newer developments in epistemic reasoning and other areas of AI, we intend to resume the task. We illustrate below the use of knowledge chunks used, and how the triggering of relevant rules can lead to explanation generation from the trace of the program.

The following rule, in OPS5, is used to recognize an honour card lead, and also makes inferences about holdings of the player in that suit.

```
(p honour-sequence-broken
  (card-rank ^card « a k q j » ^played lead ^player <west> ^suit <suit> ^rank <rank>)
  (one ^plus <rank> ^is <second>)
  (card-rank ^player ew ^rank <second> ^suit <suit> ^card <jack>)
  (one ^plus <second> ^is <third>)
  (card-rank ^player « north south » ^rank <third> ^suit <suit> ^card <ten>)
  (one ^plus <third> ^is <fourth>)
  (card-rank ^player ew ^rank <fourth> ^suit <suit> ^card <nine>)
-- >
  (make type-of ^lead honour)
  (write record ..an honour lead.. <west> should have the <jack> and <nine> as well..)
  (make inference ^player <west> ^has <nine> ^suit <suit> ^rank <third>)
  (make inference ^player <west> ^has <jack> ^suit <suit> ^rank <second>))
```

The rule says that if the opponent has led one of ace, king, queen, and jack, then (a) recognize it as an honour card lead and (b) infer that the player has the lower and the third lower cards as well.

The Bath Coup is a well known tactic when a dangerous opponent has led from (say) a suite headed by K-Q-10. This play is applicable if the declaring side has the ace and the jack. The following rule recognizing this situation. Observe that the rule uses pcard patterns like “solidity” and length of suit. It recommends a play called “duck” and also recommends that a (deceptive) encouraging signal should be given by the declarer.

```
(p try-bath-coup
  (card-rank ^played lead ^player <west> ^suit <suit> ^rank 2)
  -(lead-trick ^of <suit>)
  (solidity ^suit <suit> ^rankclass 1 ^cards 1)
  (solidity ^suit <suit> ^rankclass 4 ^cards 1)
  (sequence-is ^plays <west> ^after <south>)
  (card-rank ^suit <suit> ^player <south> ^rank 4 ^played no)
  (holding ^suit <suit> ^hrc 1 ^length > 2)
  (holding ^player <south> ^suit <suit> ^length > 2)
  (tricks ^leeway >= 1)
-- >
  (write record ..let me try a bath coup..)
  (make encourage <suit>)
  (make lead-trick ^of <suit>))
```

(make duck $\hat{}$ rounds 1 $\hat{}$ of <suit>))

In the above rules there is an action called “write record” which essentially writes the trace of reasoning into a text file. The following excerpt is an illustration of the trace when the program planned a bath coup.

“the contract is for 9 tricks in no trumps..... ..lets look at the handthere is a major *tenace* in hearts *over east* ... k - 10 of hearts ..a *tenace over west* ...there is a *two way tenace* in hearts ... k - j of hearts ..a *tenace over west* ..can afford to *lose 1 more tricks*.. ..can stand 2 - 5 *break* in spadesno thats wrong.... spades are *critically dangerous*... ..can stand 3 - 5 *break* in clubsno thats wrong.... clubs are *critically dangerous*... ..okay lets *look at the lead*.. ... west has *led* the k of spadesan *honour lead*.. west should have the q and 10 as well.. ..let me try a *bath coup*.. ..lets see what the lead has to offer... ..*will the 8 encourage west ?* ...”

In the above trace we have italicized the key words, the features and the tactical plays mentioned in the rules. As one can see, the trace is easily readable by humans, and presents the reasoning process employed by the planner.

We now discuss the manner in which the possible world semantics of Kripke structures is meaningful both for reasoning with probabilities, and in deceptive plays.

6.3 Probabilities: Reasoning with Possible Worlds

A bridge player constructs different plans and selects one based on the likelihood of success. The likelihood of success depends of the assumptions about the lie of cards. Bridge players use a combination of possible worlds and their associated probabilities to choose between different plans.

Consider a situation in which the declarer needs to develop one trick to succeed in her contract. Let there be two options. Let Option-A be a finesse in hearts. The a *priori* probability of this succeeding is 50%. Let Option-B be to play ace, king and queen and hope for a 3-3 break in clubs. The a *priori* probability of this succeeding is 36%¹⁰. Of course these probabilities may change when more information about the opponents’ hand is available from bidding. But even given these probabilities the option of a better plan is to first try Option-B and if it fails revert to Option-A. This is because testing Option-B does not entail loss of control.

We look at a deal, to be discussed again later in the context of deception, in which the declarer has reason explicitly with possible worlds. Consider the situation in which the declarer has to develop 4 ticks from this diamond suit combination, under the condition that *at no time should East be allowed to win a trick*.

North: \diamond K J 10 7 6 3

South: \diamond 9 8 2

As one can see, the ace and the queen are missing. The first observation is that if the ace is with East then the above condition will be violated. So the declarer makes a *teleological assumption*

10. https://en.wikipedia.org/wiki/Contract_bridge_probabilities

that the ace is with West, and plans accordingly. The remaining possibilities that the declarer is faced with are depicted in Figure 2. The figure is a Kripke structure with only the diamond suit depicted. The *actual world* is labeled A, but the declarer does not know that. The edge between the possible worlds depict that the declarer cannot distinguish between these possible worlds. In addition we have used the letter ‘x’ to depict insignificant cards. As a result the possible worlds A and B have in fact two distinct instances, while the rest have only one. There are, thus, eight possible worlds, which reduce to six when we ignore small cards. Each of these six, or eight, possible worlds are themselves abstractions of many, because in each the other suit cards may have different distributions. But this level of representation is adequate for the selecting a plan for playing the diamond suit.

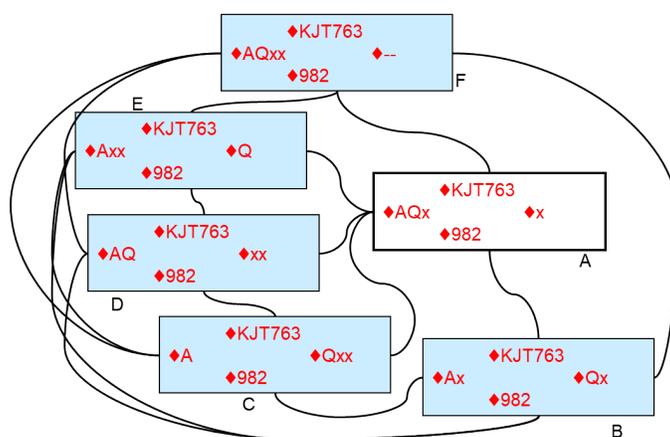


Figure 2. The remaining possible worlds in the diamond suit.

The declarer has to choose between two plans. Plan-A is to play a small card from the South hand and finesse the queen, by playing the jack. Plan-B is to play a small card to the king, and when this wins (given the assumption that the ace is with the West), play another diamond. As one can see Plan-A works in possible worlds A, D, and F. This is 4 out of 8 possible worlds. Plan-B works in the possible worlds A, B, D, E and F. That is 7 out of 8 possible worlds. Neither plan works in the possible world C. Clearly Plan-B is better, and the logical choice, and was in fact adopted by the declarer in the story described in the next section.

7. Epistemic Reasoning and Deception

Communication is an essential component of bridge. The bidding phase is all about communication. In *defense*, the opponents trying to defeat the contract also largely rely on communication. This is because *both the defenders* are active, and neither knows the combined strength of their two hands. The defenders employ signals, like high-card-encouraging, high-low, conventional leads like the fourth-best, and suit-preference signals. We have not yet addressed the defensive play in this paper, but suffice it to say the defenders have the means of communicating their own *plans* and *intentions*

to partner, and this information is also available for the declarer. So while each of the three active players cannot see two hidden hands, there is plenty of information inferred from the communication that is happening. Combined with the fact that bridge players are (generally) conversant with the set of TAs being deployed, there is also the possibility of recognizing the opponent's plan, and aim to scuttle it. One of the weapons one uses to scuttle opponent's plan is deception; and there is no better arena for skullduggery than the bridge table. An imaginative defender can advertise a fake plan to lead the declarer up the garden path. We illustrate this with a real life example from, aptly, a game during the World War I.

Maurice Gray (1889 - 1918) was a dispatch rider with the British Army and a keen bridge player. He was sitting West on the following hand (Truscott & Truscott, 2004) which we analyse from the declarer's perspective sitting South. The declarer was in a 3NT contract (to make 9 tricks) and could see the hand as depicted in Figure 3. South could count 5 top tricks (Spade A, Heart A K and Q, Club A) and needed to develop 4 more.

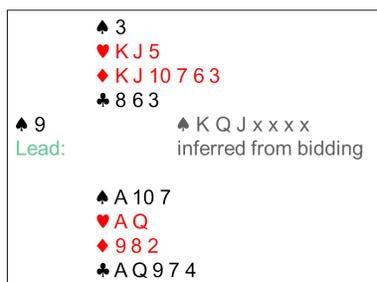


Figure 3. A 3 no trump contract after East has made a bid showing long spades and West has dutifully led the spade 9

The best option, as discussed above, was to develop the required tricks from diamonds. Since South had only one Control in spades, the Tempo was 0, since after he took the Spade Ace East would have many spade winners. East was the dangerous opponent and had to be kept out. The standard plan in this situation would be to duck 2 rounds of spades to make sure West does not have any more left, and hope that West has the Ace of Diamonds. The play would be to “play a small diamond to the king” It would work in the possible worlds A, B, D, E and F in Figure 3. We would not mind West winning 2 diamond tricks, as long East does not get any. In the layouts B and E East does have the Queen but does not get a trick for it, because it falls below the Ace or the King.

Consequently South embarked upon this plan. Figure 4 depicts the possible worlds as seen from the perspective of West, Maurice Gray. The possible worlds as seen by West are shown in dashed lines. Observe that the possible worlds for the two players have the actual world as common to both. Now West can see that given his own diamond holding, if declarer tries for extra tricks in diamonds he will succeed, whether he attempts Plan-A or Plan-B discussed above. West also knows that since one of the goals of declarer being to keep East at bay, he will most likely choose Plan-B. West knows his own cards in diamonds, but he also *knows that the declarer does not know* them, If he could convince South that the situation were the possible world B, then perhaps he could be

led astray. So, Gray played *as if the situation was the possible world B*, and made what would be a brilliant defense in possible world B.

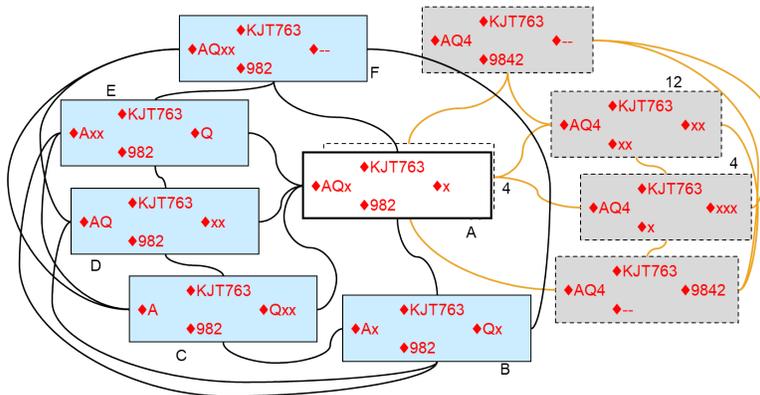


Figure 4. The possible worlds for West in the diamond suit.

The play went as follows. On trick 1 South ducked and allowed East to win the Jack of Spades. On trick 2 likewise East was allowed to win the King of Spade. South won trick 3 with the Ace of Spades and was surprised to see West jettisoning the Ace of Diamonds!

South now made the following epistemic inferences - “West has understood South’s plan, and is trying to counteract it by *creating an entry* for East. He does this by *jettisoning* the Ace of Diamonds so that in the layout B South cannot *develop* diamond tricks without letting East win a trick with the queen, who would then *cash* all his spades”. Consequently he abandons the plan for developing diamond tricks, discards the Diamond 6, and attempts to try the alternate plan of developing club tricks.

When the smoke cleared he discovered that the layout was actually A! Sitting West Maurice Gray had spun a web of deception leading South astray from the original plan that *would have worked*. Gray imagined the layout B and played as if that was the true layout, and succeeded in deceiving South.

8. Concluding Remarks

When a bridge playing program can execute such a play we would surely have to concede that it is intelligent. The key to such plays is the ability to imagine. The complexity arises because an agent has to imagine another agent thinking about possible worlds, and reason with imaginary worlds.

The complexity in contract bridge arises mainly because of incomplete information combined with the fact that the number of starting positions in the game is very large. As opposed to Monte Carlo simulations, human players adopt a knowledge based planning approach in which they retrieve partial plans from a repository of standard plays and string them into a whole. We envisage the following challenges for implementing bridge playing programs on similar lines.

- Representation of cards. Initially the cards are ranked from the ace to the 2 in each suit, with ace being rank 1 and the 2 being rank 13. However as the game proceeds and cards go out of

play, the ranks must be maintained dynamically, and the program must reason with these *fluent* ranks.

- Abstract representation of cards. One approach like human reasoning is to abstract away insignificant cards into a variable called ‘x’ as is done in the analysis above. But this has to be done carefully, because the ranks of cards change dynamically. If the A, K, Q and J have been played on trick 1, then the 10 has become the new ‘ace’! Consequently smaller cards also gain in stature.
- However, sometimes small cards may play an important role too, and must maintain their identity. One such situation is when one is trying to execute an end-play called a *throw-in*. To make it work the declarer may have to carefully preserve small cards so that an opponent is compelled to win when the suit is played.
- Representation of card patterns. Since a lot of reasoning is done with card patterns, like for example the tenace described in this paper, one needs an appropriate representation of such patterns. Moreover, as discussed above, old patterns vanish and new ones arise as play proceeds and cards go out of play.
- As a corollary of the above, one must develop algorithms to spot patterns dynamically as play proceeds. Perhaps a structure like a multi-threaded the Rete Net could be used here, to maintain the set of active patterns as play progresses.
- The play program must monitor the cards and the patterns for any opportunistic planning situations, or when the assumptions turn out to be false and replanning is exigent.
- Representation of the combined picture of the hands during bidding. This would involve making inferences from bids and creating fuzzy representations consistent with the 13 card ceiling for each hand. For example, one must be able to represent the fact that the partner, or an opponent, has a six or seven card spade suit (as was the case in the hand described in Section 6).
- Implementation of an algorithm that will estimate the combined playing strength of the two hands during bidding, and the band that the hands fall in.
- Implementation of decision making procedures that will determine the information to be conveyed next during bidding.
- Implementation of a “bidding language understanding” programs that will associate the features based semantics with bidding systems. This will be needed both for making bids (generation) and understanding opponents’ bids.
- Implementation of a system to interpret signals during play, and also generate signals during defense.
- Implementing a system for fuzzy epistemic representation that will incorporate the above inferences into a minimal set of possible worlds after abstracting away from irrelevant information.
- Implementation of a logic like reasoning system over epistemic states, as opposed to the model checking approaches currently in vogue. Studying the tradeoff between tractable reasoning against incompleteness of such a system.

The Epistemic Logic community has started investigating lying and deception, and contract bridge should be an ideal testing ground for doing so. An added challenge is to somehow circumvent the complexities of model checking with Kripke structures, and alternate ways of representations have to be found. Thus considerable work needs to be done in knowledge representation, but unless

we do so, there is little hope of reasoning with models of the agent’s environment in an informed manner. But when we do so, we should have a system in which a man-machine combination can meaningfully play the game of bridge. It will also hopefully give us some insights into building epistemic reasoning systems in adversarial incomplete information environments, which are not uncommon in the real world.

Acknowledgements

Our understanding of the Monte Carlo methods and current approaches to bidding greatly benefitted from discussions with Arun Bahulkar and his team at TRDDC Pune. We also thank R Ramanujam and Hans Ditmarsch of IMSc Chennai for insights into Dynamic Epistemic Logic and deception. We also thank the reviewers of ACS for the many insightful and helpful comments.

References

- Allis, L. (1994). *Searching for solutions in games and artificial intelligence*. Doctoral dissertation, Maastricht University, Limburg, Netherlands.
- Amit, A., & Markovitch, S. (2006). Learning to bid in bridge. *Machine Learning*, 63, 287–327.
- Brown, N., & Sandholm, T. (2017). Superhuman AI for heads-up no-limit poker: Libratus beats top professionals. *Science*, 359, 418–424.
- Campbell, M., Hoane Jr, A. J., & Hsu, F.-H. (2002). Deep Blue. *Artificial Intelligence*, 134, 57–83.
- Carley, G. L. (1962). *A program to play contract bridge*. Doctoral dissertation, Department of Electrical Engineering, Massachusetts Institute of Technology, Cambridge, MA.
- DeLooze, L. L., & Downey, J. (2007). Bridge bidding with imperfect information. *2007 IEEE Symposium on Computational Intelligence and Games* (pp. 368–373). Honolulu, HI: IEEE.
- van Ditmarsch, H., van Der Hoek, W., & Kooi, B. (2007). *Dynamic epistemic logic*. Springer.
- Frank, I., & Basin, D. (1998). Search in games with incomplete information: A case study using bridge card play. *Artificial Intelligence*, 100, 87–123.
- Frank, I., Basin, D. A., & Matsubara, H. (1998). Finding optimal strategies for imperfect information games. *Proceedings of the Fifteenth National Conference on Artificial Intelligence* (pp. 500–507). Madison, WI: AAAI Press.
- Gambäck, B., Rayner, M., Yard, M., & Pell, B. (1993). *Pragmatic reasoning in bridge*. Technical report, SRI International, Cambridge Computer Science Research Center, Cambridge.
- Ginsberg, M. L. (1996). Partition search. *Proceedings of the Thirteenth National Conference on Artificial Intelligence* (pp. 228–233). Portland, OR: AAAI Press.
- Ginsberg, M. L. (1999). GIB: Steps toward an expert-level bridge-playing program. *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence* (pp. 584–593). Stockholm, Sweden: Morgan Kaufmann.
- Ginsberg, M. L. (2001). GIB: Imperfect information in a computationally challenging game. *Journal of Artificial Intelligence Research*, 14, 303–358.

- Ginsberg, M. L., & Harvey, W. D. (1992). Iterative broadening. *Artificial Intelligence*, 55, 367–383.
- Khemani, D. (1989). *Theme based planning in an uncertain environment*. Doctoral dissertation, Department of Computer Science and Engineering, IIT Bombay, Mumbai, India.
- Khemani, D. (1994). Planning with thematic actions. *Proceedings of the Second International Conference on Artificial Intelligence Planning Systems* (pp. 287–292). Chicago: AAAI Press.
- Levy, D. N. (1989). The million pound bridge program. In D. Beal (Ed.), *Heuristic programming in artificial intelligence*, 95–103. Ellis Horwood.
- Lindelof, T. (1983). *COBRA: The computer-designed bidding system*. London: Victor Gollancz.
- Muise, C. J., Belle, V., Felli, P., McIlraith, S. A., Miller, T., Pearce, A. R., & Sonenberg, L. (2015). Planning over multi-agent epistemic states: A classical planning approach. *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence* (pp. 3327–3334). Austin, TX: AAAI Press.
- Schaeffer, J., Burch, N., Björnsson, Y., Kishimoto, A., Müller, M., Lake, R., Lu, P., & Sutphen, S. (2007). Checkers is solved. *Science*, 317, 1518–1522.
- Sheppard, B. (2002). World-championship-caliber Scrabble. *Artificial Intelligence*, 134, 241–275.
- Silver, D., et al. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, 529, 484–489.
- Silver, D., et al. (2017a). Mastering chess and shogi by self-play with a general reinforcement learning algorithm. *arXiv preprint arXiv:1712.01815*.
- Silver, D., et al. (2017b). Mastering the game of Go without human knowledge. *Nature*, 550, 354.
- Smith, S. J., Nau, D., & Throop, T. (1998a). Computer bridge: A big win for AI planning. *AI Magazine*, 19, 93–105.
- Smith, S. J., Nau, D. S., & Throop, T. A. (1996). A planning approach to declarer play in contract bridge. *Computational Intelligence*, 12, 106–130.
- Smith, S. J., Nau, D. S., & Throop, T. A. (1998b). Success in spades: Using AI planning techniques to win the world championship of computer bridge. *Proceedings of the Fifteenth National Conference on Artificial Intelligence* (pp. 1079–1086). Madison, WI: AAAI Press.
- Sterling, L., & Nygate, Y. (1990). Python: An expert squeezer. *The Journal of Logic Programming*, 8, 21–39.
- Tesauro, G. (1995). Temporal difference learning and TD-gammon. *Communications of the ACM*, 38, 58–68.
- Truscott, A., & Truscott, D. (2004). *The New York Times bridge book: An anecdotal history of the development, personalities, and strategies of the world’s most popular card game*. Macmillan.
- Wasserman, A. I. (1970). Realization of a skillful bridge bidding program. *Proceedings of the 1970 Fall Joint Computer Conference* (pp. 433–444). ACM.
- Yeh, C.-K., & Lin, H.-T. (2016). Automatic bridge bidding using deep reinforcement learning. *arXiv preprint arXiv:1607.03290*.