

AI v.s. Bridge Masters


Alfonso Ruiz (Escuela Bourbaki)

Program

1. AI's weapons
2. Chess
3. Guess Who?
4. Bridge
5. State of the art

AI's weapons

Computational power



Mathematicians





Mathematics: Counting

$$C(52, 13) = \frac{52!}{13!39!} = \frac{52 \times 51 \times 50 \times 49 \times 46 \times 45 \times 44 \times 43 \times 42 \times 41 \times 40}{13 \times 12 \times 11 \times 10 \times 9 \times 8 \times 7 \times 6 \times 5 \times 4 \times 3 \times 2 \times 1} = 635,013,559,600$$

Mathematics: Counting

$$\sim 10^{12}$$

$$10^{16}$$

$$10^{45}$$

Sagan's number



Mathematics: Probability Theory




$$P = \frac{C(4, 4) \times C(22, 9)}{C(26, 13)} = \frac{1 \times 497420}{10400600} = 4.8\%$$


Chess




Shannon, IBM, etc...



Legal moves ~ 30



Number of moves ~ 40



Possible games $\sim 10^{120}$

Evaluating functions

$$\begin{aligned} f(P) = & 200(K-K') + 9(Q-Q') + 5(R-R') + 3(B-B'+N-N') + (P-P') - \\ & 0.5(D-D'+S-S'+I-I') + \\ & 0.1(M-M') + \dots \end{aligned}$$

Honors counting



A-4

K-3

Q-2

J-1

This hand:

19

Minmax strategy

$$\max_{\forall p' = \text{next}(p)} \left\{ \min_{\forall p'' = \text{next}(p')} f(p'') \right\}$$


No learning?


Pruning


Guess who?




Reinforcement learning







Evaluating (reward) functions



$$f(P, Q, a)$$

$$f(P, Q, b)$$

$$f(P, R, c)$$

$$f(Q, R, d)$$

Markov Hypothesis

Bellman equations + Monte Carlo

Bridge

Difficulties

1. Randomness
2. Play: incompleteness
3. Biding: intelligence

Randomness

Play:

Best action v.s. best reward?

Finessing

	N	A Q X	
W			E
	S	X X X X	

Biding

Increasing auction

State of the art

GIB: Steps Toward an Expert-Level Bridge-Playing Program

Matthew L. Ginsberg

CIRL

1269 University of Oregon

Eugene, OR 97403-1269

ginsberg@cirl.uoregon.edu

Abstract

This paper describes GIB, the first bridge-playing program to approach the level of a human expert. (GIB finished twelfth in a hand-picked field of thirty-four experts at an invitational event at the 1998 World Bridge Championships.) We give a basic overview of the algorithms used, describe their strengths and weaknesses, and present the results of experiments comparing GIB to both human opponents and other programs.

1 Introduction

Of all the classic games of mental skill, only card games and Go have yet to see the appearance of serious computer challengers. In Go, this appears to be because the game is fundamentally one of pattern recognition as opposed to search; the brute-force techniques that have been so successful in the development of chess-playing programs have failed almost utterly to deal with Go's huge branching factor. Indeed, the arguably strongest Go program in the world (Handtalk) was beaten by 1-dan Janice Kim (winner of the 1984 Fuji Women's Championship) in the 1997 Hall of Champions after Kim had given the program a monumental 25 stone handicap.

reducing the poker "problem" to a large linear optimization problem which is then solved to generate a strategy that is nearly optimal in a game theoretic sense. Schaeffer, author of the world champion checkers program CHINOOK [Schaeffer, 1997], is also reporting significant success in this domain [Billings *et al.*, 1998].

The situation in bridge has been bleaker. In addition, because the American Contract Bridge League (ACBL) does not rank the bulk of its players in meaningful ways, it is difficult to compare the strengths of competing programs or players.

In general, performance at bridge is measured by playing the same deal twice or more, with the cards held by one pair of players being given to another pair during the replay and the results then being compared.² A "team" in a bridge match thus typically consists of two pairs, with one pair playing the North/South (N/S) cards at one table and the other pair playing the E/W cards at the other table. The results obtained by the two pairs are added; if the sum is positive, the team wins this particular deal and if negative, they lose it.

In general, the numeric sum of the results obtained by the two pairs is converted to International Match Points, or IMPs. The purpose of the conversion is to diminish the impact of single deals on the total, lest an abnormal result on one particular deal have an unduly large impact on the result of an entire match.

Learning to bid in bridge

Asaf Amit · Shaul Markovitch

Received: 8 February 2005 / Revised: 17 October 2005 / Accepted: 15 November 2005 /
Published online: 9 March 2006
Springer Science + Business Media, LLC 2006

Abstract Bridge bidding is considered to be one of the most difficult problems for game-playing programs. It involves four agents rather than two, including a cooperative agent. In addition, the partial observability of the game makes it impossible to predict the outcome of each action. In this paper we present a new decision-making algorithm that is capable of overcoming these problems. The algorithm allows models to be used for both opponent agents and partners, while utilizing a novel model-based Monte Carlo sampling method to overcome the problem of hidden information. The paper also presents a learning framework that uses the above decision-making algorithm for co-training of partners. The agents refine their selection strategies during training and continuously exchange their refined strategies. The refinement is based on inductive learning applied to examples accumulated for classes of states with conflicting actions. The algorithm was empirically evaluated on a set of bridge deals. The pair of agents that co-trained significantly improved their bidding performance to a level surpassing that of the current state-of-the-art bidding algorithm.

Keywords Bridge · Partner modeling · Opponent modeling · Learning to cooperate

Boosting a Bridge Artificial Intelligence

Véronique Ventos^{*,1}, Yves Costel², Olivier Teytaud^{†,1}, and Solène Thépaut Ventos³

¹LRI, INRIA, CNRS UMR 8623, University Paris-Saclay, France

²WBRIDGE5, Lyon, France

³LMO, CNRS UMR 8628, University Paris-Saclay, France

Preprint submitted and accepted to International Conference on Tools with Artificial Intelligence (ICTAI 2017). The final version will be included at the ICTAI proceedings

Abstract

Bridge is an incomplete information game which is complex both for humans and for computer bridge programs. The purpose of this paper is to present our work related to the adaptation to bridge of a recent methodology used for boosting game Artificial Intelligence (AI) by seeking a random seed, or a probability distribution on random seeds, better than the others on a particular game.

The bridge AI Wbridge5 developed by Yves Costel has been boosted with the best seed found on the outcome of these experiments and has won the World Computer-Bridge Championship in September 2016.

1 Introduction

Games have always been an excellent field of experimentation for the nascent techniques in computer science and in different areas of Artificial Intelligence (AI) including machine learning (ML). Despite their complexity, games problems are much easier to understand and to model than real life problems. Indeed, most games have a limited number of simple rules and have been subject to in-depth human analysis over time. Systems initially designed for games are then used in the context of real applications. Therefore, next-generation Watsons, the IBM system that has beaten two champions of Jeopardy!¹ in 2011, are used as consultant machines in fields as varied as medicine, cybersecurity, and business analytics.

Automatic Bridge Bidding Using Deep Reinforcement Learning

Chih-Kuan Yeh¹ and Hsuan-Tien Lin²

Abstract.

Bridge is among the zero-sum games for which artificial intelligence has not yet outperformed expert human players. The main difficulty lies in the bidding phase of bridge, which requires cooperative decision making under partial information. Existing artificial intelligence systems for bridge bidding rely on and are thus restricted by human-designed bidding systems or features. In this work, we propose a pioneering bridge bidding system without the aid of human domain knowledge. The system is based on a novel deep reinforcement learning model, which extracts sophisticated features and learns to bid automatically based on raw card data. The model includes an upper-confidence-bound algorithm and additional techniques to achieve a balance between exploration and exploitation. Our experiments validate the promising performance of our proposed model. In particular, the model advances from having no knowledge about bidding to achieving superior performance when compared with a champion-winning computer bridge program that implements a human-designed bidding system.

1 Introduction

Games have always been a challenging testbed for artificial intelligence (AI). Even for games with simple and well-defined rulesets, AI often needs to follow highly complex strategies to gain victory. One set of works on game AI focuses on full information games including chess, go, and Othello [18], whereas the other set studies incomplete information games such as poker and bridge [9, 17, 22]. In both cases, traditional works usually excel by embedding the knowledge of the best human players as computable strategies; however, researchers have recently shifted their focus to machine learning, allowing AI players to develop effective strategies automatically from data [9, 18, 22].

affects the score that the declarer’s team can get in the playing phase. The auction proceeds around the table in a clockwise manner, where each player chooses from one of the following actions: PASS, increasing the current value of the bid with respect to an ordered set of calls $\{1\clubsuit, 1\diamondsuit, 1\heartsuit, 1\spadesuit, 1NT, 2\clubsuit, \dots, 7NT\}$, DOUBLING and REDOUBLING. The first two actions are general ones for deciding the contract, while the latter two are special, less-used actions that modify the scoring function for the playing phase. The bidding sequence ends when three consecutive PASSES are placed, and the last bid becomes the final contract. The number in the final contract (such as 4 in 4♠) plus 6 represents the number of rounds that the team aims to win in the playing phase to achieve the contract (commonly referred to as “make”), and the symbol (such as ♠) reflects the trump suit in the playing phase.

In the playing phase of the bridge game, there are 13 rounds where each player shows one card from her/his hand and compares the values of the cards based on some rulesets with the trump suit having some priority. The player with the highest-valued card among the four is the winner of the round. After the 13 rounds, the score of the declarer’s team is calculated by a lookup table based on the final contract and the number of winning rounds of the declarer’s team, where making the contract leads to a positive score for the declarer’s team, and not making (failing) the contract results in a positive score for the the opponent’s team.

Bidding is an understandably a difficult task because of the incomplete-information setting. Given that each player can only see 13 out of 52 cards, it is impossible for a single player to infer the best contract for her/his team. Thus, each bid in the bidding phase needs to serve as a suggestion towards an optimal contract, information-exchanging between team members, or both. That is, a good bidding strategy should balance between exploration (exchanging information) and exploitation (deciding an optimal contract). Nevertheless, because the bid value needs to be monotonically increasing during

Escuela de matemáticas Bourbaki




Contact

alfonso@escuela-bourbaki.com

Catalog of ludic courses

- Escher
- Music
- Tortues
- Chaos Theory
- Geometry



Dibujo de Sky and Water I, Escher

ESCHERMATICS (d'après Sir R. Penrose)

Ludic courses



Contact

alfonso@escuela-bourbaki.com
