Bidding Decision Model for the Card Game Tarok

DOMEN MARINČIČ, MATJAŽ GAMS, MITJA LUŠTREK
Department of Intelligent Systems
Jozef Stefan Institute
Jamova 39, SI-1000 Ljubljana
SLOVENIA
domen.marincic@ijs.si, matjaz.gams@ijs.si, mitja.lustrek@ijs.si

Abstract: In the paper we present a decision model for bidding in the card game four-player tarok. The decision model was built on the basis of Bayesian networks. We developed a program called Tarok7 for playing four-player tarok, which served as a testbed for evaluating the decision model. With the help of Tarok7 program, the model was compared to human experts. The results of the comparison show high matching of the program's bidding with the bidding of human experts.

Key-Words: Decision models, Bayesian networks, Card games, Bidding, Tarok

1 Introduction

Bidding is a part of various card games, for example bridge, whist, poker and tarok. Before the actual card play, players offer to play more and more difficult types of games with the one making the most ambitious offer choosing the type that will be played. The winner of bidding is the one who actually scores in the game. Since bidding requires the prediction of the final outcome of each type of a game, it is in a way more difficult than card play itself.

In bidding, players bid in a sequence divided into rounds. In each round every player makes one bid. Each bid must be higher than the previous one. Bidding is finished when all the players but one pass, i.e. do not continue with higher bids. The remaining player is called the declarer.

In four-player tarok the declarer can play against the other three players teamed together or choose one partner, depending on his last bid. The strategy of a bidder is to choose the most appropriate bid according to the strength of his hand and the estimated strength of the other players. Generally, the types of games associated with high bids require the bidder to have a higher advantage over the other players to win the game.

In general, there are two approaches to solving bidding problems: knowledge-based and simulation-based. The advantage of knowledge-based systems is their better explicability. As a result, the decision process can be controlled and tuned more easily than in simulation-based systems. Another advantage is the speed of the decision process. In highly complex non-perfect information games, simulation-based systems consume a lot of time. On the other hand, it is a lot easier to build a simulation-based system, because no knowledge acquisition and implementation is needed.

We decided to use the knowledge-based approach. We developed a decision model on the basis of Bayesian networks [3]. The model was implemented and tested in the program for playing four-player tarok called Tarok7 described in [7].

In the paper we first give a brief overview of related work. Then we describe the structure of the decision model. After that we present the evaluation tests of the model. At the end, a conclusion is presented and suggestions for future work are discussed.

2 Related Work

Bridge is probably the best-known game with bidding and GIB [2] the best-known bridge-playing program. GIB uses a database of bidding strategies, but its principal bidding mechanism is simulation. Another well-known program for playing bridge is Bridge Baron, described in [8], where authors mostly discuss the card play while bidding is not treated in detail.

Whist is also a card game that includes bidding. [4] describes a program for playing whist that uses game search for card play. Experiments to perform bidding with simulation have been performed. The author reports poor results, so bidding was not included in the program.

Betting in poker is also a form of bidding. Poki, a program for playing poker described in [1], bases its decisions on simulation aided by statistical modelling of opponents. In [5], another program for playing poker, Bayesian networks are used to represent (i) the relationships between the current hand type, (ii) the final hand type after the five cards have been dealt and (iii) the behaviour of the opponent. Thus the posterior probability of winning a game is obtained.

Another game that includes bidding is three-player tarok. Silicon tarokist, a program for three-player tarok described in [6], [9], also employs simulation for bidding.

3 Description of the Decision Model for Bidding

The Bayesian network that represents the bidding decision model is presented in Figure 1. The top-level nodes represent the state of the game at the moment a player has to make a bid. The nodes connected to the node "Strength of bidder" represent various features describing the bidder's cards. The nodes connected to the node "Strength of other players" represent last bids of other players and their overall quality of play as estimated by the bidder. The mid-level nodes semantically integrate the attributes in the top-level nodes. They are not strictly necessary, but they make the network more compact and easier to design. Each bottom-level node represents one of the possible bids and therefore the type of game associated with that bid. The random variables associated with the bottom-level nodes represent the possible final scores of the game and can have the following values: "high defeat", "low defeat", "low win" and "high win".

To determine the optimal bid, the prior probabilities of all top-level are set according to the current state of the game. In our case, the node "Bid of player A" can have three values: "pass", "low bid" and "high bid". If the player's last action was a low bid, then the probability of this value is set to 1 while the other two values are assigned probabilities equal to 0. Probabilities of the other top-level nodes are set in the similar way.

The posterior probabilities of the values in the bottom-level nodes are calculated according to the inference rules of Bayesian networks. The particular structure of the Bayesian network allows us to use an adapted version of general inference rules. Let v_b be a value of a bottom-level node B. Let $M = \{M_1, M_2, ..., M_k\}$ be the set of k parent nodes of the node B. Let $V_{m_i} = \{v^1_{M_i}, ..., v^{n_i}_{M_i}\}$ be the set of n_i values of the parent node M_i. $P(B = v_b)$ is then calculated by Equation (1) First, we calculate the probabilities of the mid-level nodes. Probabilities of the bottom-level nodes are calculated recursively with the same formula.

The values ("high defeat", "low defeat", "low win" and "high win") of the bottom-level nodes are assigned discrete numeric values -1, -1/3, 1/3 and 1 respectively. The result of the inference rules are the probabilities of the values of the bottom-level nodes. Then, the expectations of the probability distributions of the bottom-level nodes are calculated. The bid associated with the highest expectation is chosen. If all the expectations are negative, the bidder should pass.

4 Tuning the Decision Process

The model incorporates three essential factors of bidding decisions: (i) the strength of players (in the mid-level nodes, which summarise the bidder's cards and previous bids of the other players), (ii) the values of the types of the games associated with bids and (iii) the level of risk. The second and the third factor are incorporated in the probability distributions in the conditional probability tables of the bottom-level nodes.

Figure 2 illustrates how the model deals with

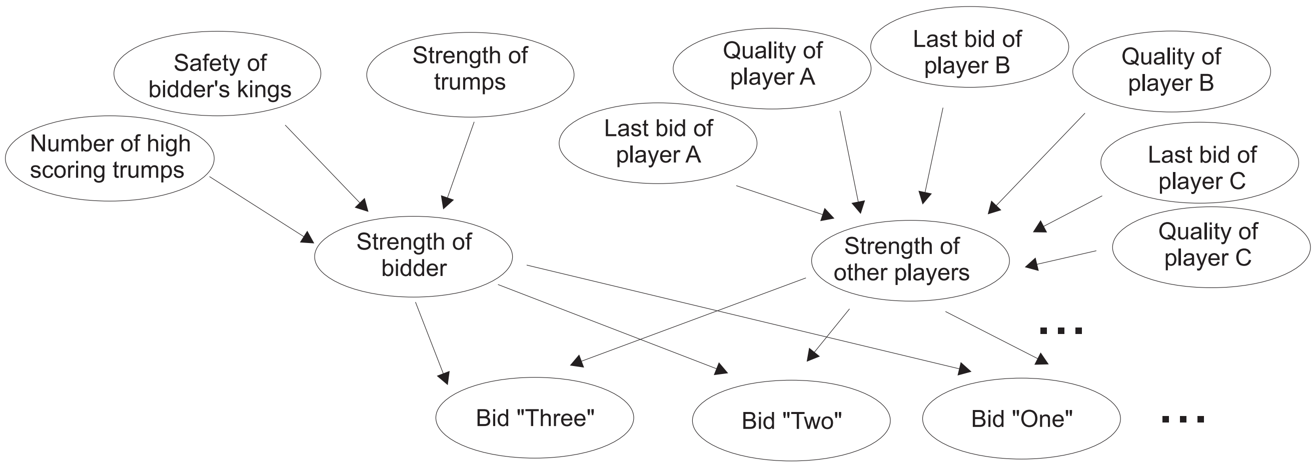

Figure 1: Decision model for bidding in four-player tarok

$$P(B = v_b) = \qquad (1)$$

$$\sum_{j_1=1\ldots n_1,\ldots,j_k=1\ldots n_k} P(v_b|M_1 = v_{M_1}^{j_1}, \ldots, M_k = v_{M_k}^{j_k})P(M_1 = v_{M_1}^{j_1}) \ldots P(M_k = v_{M_k}^{j_k})$$

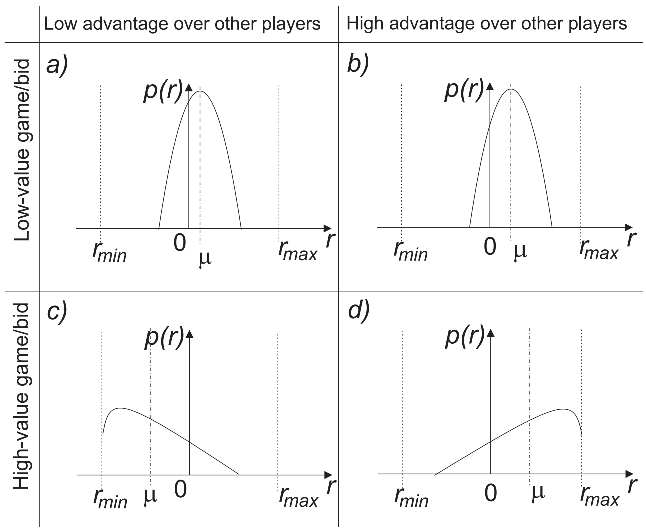

Figure 2: Four basic decision situations at bidding

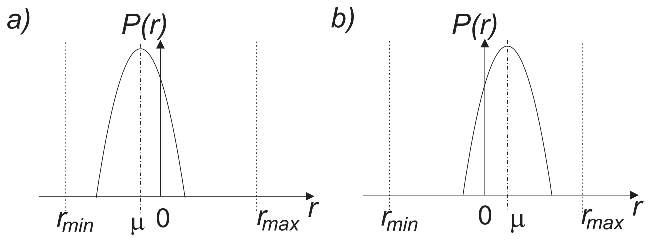

Figure 3: Adapting risk at bidding

the factors (i) and (ii) with four simple cases regarding the advantage of the bidder over the other players (low/high; i) and the game value (low/high; ii). Each of the probability density functions represents the discrete probability distribution in a bottom-level node, as calculated in the decision process. To make the example more informative we present it using continuous values, although discrete values are used in the real model.

On the horizontal axis r are the expected game scores. The lower and the upper score limits are denoted by r_{min} and r_{max}. The probability density functions associated with these scores are on the vertical axes and are denoted by $p(r)$. The expectations are denoted by μ. Note that this is only a schematic presentation of the probability distributions in the model.

Two decision situations are depicted in Figure 2. In the left column, the bidder and his partner are slightly stronger or at least not much weaker than the other players; in the right column the bidder estimates that together with the partner they possess much better cards than the other players. In both situations the bidder can choose a low-value bid where low negative or positive scores are expected, a high-value bid with high

Expert	A	B	C
Matching of the program with the experts	92%	82%	80%
Percentage of the bids with difference of more than one degree	1%	2%	2%
Percentage of the program's more aggressive bids when the program and the expert bid differently	35%	75%	72%

Table 1: Comparison of bidding of the Tarok7 program with human experts

scores expected or pass.

The bidding decision is performed in two steps: First, a particular situation, e.g. "Low advantage" is determined. Then, μ for the low-value bid (a) and μ for the high-value bid (b) are computed. The bidder will evidently choose the bid corresponding to the higher value of μ. If μ is negative, pass is a reasonable choice.

Over multiple games, e.g. in the course of a tournament, new information is obtained to change the bidder's strategy in the sense of greater or lower aggressiveness or risk. This can be easily modelled by modifying the probability distributions in the conditional probability tables of the bottom-level nodes, making bidding more or less aggressive. An example in Figure 3 shows two distributions, encouraging less (a) or more (b) risky bidding.

5 Evaluation of the Decision Model

In this test four computer players represented by Tarok7 program and an expert human player were bidding at the same time. When it was the fourth player's turn to bid, first the expert made a bid followed by the fourth computer player. In this way the human and the computer player were put in exactly the same position at bidding.

Table 1 summarises the results of the test. Bidding of Tarok7 is compared to three human experts: A, B and C. Expert A made 500 bids, while the other two made 100 bids each. The percentages in the second row denote the proportion of bids when the program and the humans chose the same action. The result 100% would mean a complete match. The third row represents the cases when the difference between the program's bid and the expert's bid was more than one degree. For the cases when the experts and the program bid differently, the fourth row shows the percentages of bids when the program bid higher than the human. The value 100% would mean that the program always bid higher than the expert when they bid differently.

Bidding of Tarok7 is more similar to expert A than to the other experts, which was expected since expert A designed the decision model for bidding. According to the results in the fourth row, expert A bid slightly more aggressively than the program, while the other two experts were less aggressive. Overall, there are very few cases when the experts and the program disagree strongly in their decisions.

6 Conclusion

In this paper we described a decision model for bidding in the four-player tarok card game. The model is based on a Bayesian network. The test that we performed shows high matching between the model's decisions and the decisions of human experts in playing tarok. Moreover, the decision model rarely chooses bids that are radically different than experts' bids.

We explained how it is possible to adapt the model relatively easily when changes of bidding strategy are required. We can do this by changing the probability distributions in the conditional probability tables of the bottom-level nodes of the model that represent the possible bids. One of the possibilities for future work is automatic learning of the values of the conditional probability tables. The final results of games played under different bids could serve as the feedback for the learning algorithm.

References:

[1] D. Billings, A. Davidson, J. Schaeffer, D. Szafron, "The challenge of poker," *Artificial Intelligenc2e Journal*, Vol. 134, No. 1-2 , pp. 201-240, 2002.

[2] M. L. Ginsberg, "GIB: Imperfect information in computationally challenging game," *Journal of Artificial Intelligence Research*, Vol. 14, pp. 303-358, 2001.

[3] F. V. Jensen, "Bayesian networks and decision graphs," *Springer Verlag: New York*, 2001.

[4] J. Fong, "Developing an Artificial Intelligence for Whist", University of California, Computer Science, Los Angeles, 2005.

[5] K. Korb, A. Nicholson, N. Jitnah, "Bayesian poker," in *Proceedings of the 15th Conf. on Uncertainty in Artificial Intelligence Stockholm, Sweden*, Vol. , pp. 343-350, 1999.

[6] M. Luštrek, M. Gams, I. Bratko, "A program for playing tarok," *ICGA Journal*, Vol. 23 , No. 3., pp. 190-197, 2003.

[7] D. Marinčič, "Kombinacija teorije odloanja in preiskovanja pri raunalnikem igranju iger," *Masters thesis, Faculty of computer and information engineering, University of Ljubljana*, 2004.

[8] S. J. J. Smith, D. Nau, T. Throop, "Computer bridge: a big win for AI planning," *AI Magazine*, Vol. 19, No. 2, pp. 93-105, 1998

[9] Silicon tarokist website. http://tarok.bocosoft.com.