

Neural networks for contract bridge bidding

B YEGNANARAYANA, DEEPAK KHEMANI and
MANISH SARKAR

Department of Computer Science & Engineering, Indian Institute of Technology, Madras 600 036, India
email: yegna@iitm.ernet.in

Abstract. The objective of this study is to explore the possibility of capturing the reasoning process used in bidding a hand in a bridge game by an artificial neural network. We show that a multilayer feedforward neural network can be trained to learn to make an opening bid with a new hand. The game of bridge, like many other games used in artificial intelligence, can easily be represented in a machine. But, unlike most games used in artificial intelligence, bridge uses subtle reasoning over and above the agreed conventional system, to make a bid from the pattern of a given hand. Although it is difficult for a player to spell out the precise reasoning process he uses, we find that a neural network can indeed capture it. We demonstrate the results for the case of one-level opening bids, and discuss the need for a hierarchical architecture to deal with bids at all levels.

Keywords. Artificial neural networks; backpropagation; games; contract bridge bidding; knowledge; artificial intelligence.

1. Introduction

Artificial Neural Networks (ANN) can be trained to capture the implicit associations between an input pattern and the corresponding output response of complex systems (Haykin 1994). Often such associations can be used to perform tasks like pattern classification and pattern mapping. But associations can be very complex, depending upon the richness of the domain. We consider one such domain, namely, bidding in the game of Contract Bridge.

In contract bridge, a player makes a bid to convey information about the pattern of the thirteen cards in his hand. If he is the first to make a bid, it is called “opening bid”, which he makes based only on the pattern of the cards he is holding. That is, he has no *a priori* knowledge of the rest of the cards in other players’ hands. But during the rest of the auction he makes a bid based on not only his own cards, but also on the bids that have been made till then. The problem of a machine making even an opening bid in this game is challenging. There are many hand patterns and situations which are not covered precisely by the “rules” in the convention system used in the game. Thus a straightforward

application of a rule-based system cannot logically explain the reasoning process used by a player in making a bid. The objective of this paper is to explore the possibility of capturing the reasoning process of a player in making an opening bid based on the pattern of the cards presented to him. It may be noted that although there are rules describing the conventions (such as Standard American) of the bidding system used by a player, the player uses these conventions only as a guideline for making a bid with a given hand. It is generally difficult to precisely state consistently the reasoning process used in making a particular bid. Sometimes for a given hand, the same player may make a different bid at a different time. It is this variability in the reasoning of the player that we intend to capture, if possible, by an artificial neural network.

Games have been of interest to Artificial Intelligence (AI) researchers because they provide well defined domains which are nevertheless rich in computational complexity. Games like bridge and chess have been engaging the attention of AI researchers for a long time. The intrinsic symbolic nature of these games implies that one does not have to compromise in representation. In contrast, many compromises have to be made in representation while dealing with real world problems, such as in speech recognition and image understanding.

Traditionally, in the context of AI, game theoretic methods are applicable to two-person complete information games like chess. Bridge, however is different (Khemani & Ramakrishna 1984). Making a bid, in particular, is an incomplete information activity. A knowledge-based approach seems to be better suited. But we observe that subtle interactions between many simple rules make the traditional rule-based approach difficult to implement. It is here that we are hoping that ANNs might prove to be more useful.

In § 2 we describe the nature of the bidding problem and also the scope of the present work. In § 3 we describe the issues involved in exploring ANNs for the bidding problem. The basic network architecture and experiments with different number of nodes are described in §§ 4 and 5, respectively. A hierarchical architecture is suggested in § 5 to overcome some of the difficulties of the basic network.

2. Bidding problem: Scope of the present study

The goal of bidding in bridge is to cooperatively estimate the playing strength of the two hands held by a partnership, and arrive at an optimal contract. Some details of the game of bridge are given in appendix A. Since each player can see only his/her hand, cooperation with the partner takes place in the form of exchange of information encoded in the bids. The formal interpretation of staking a claim through a bid is described below.

A bid consists of two parts.

- (a) The denomination or the level which determines the number of tricks the bidding side is claiming to make.
- (b) The suit which the side selects as “trump”. For example, a bid of “3 spades”(3S) means that with spades as trump the side will make nine (3 + 6) tricks. Apart from the four suits (spade(S), heart(H), diamond(D) and club(C)), one can also choose to bid “no trumps”(N), meaning that no suit is the trump, and all are equal.

The different bids that can be made have an order imposed by the rules of the game, as follows: 1C, 1D, 1H, 1S, 1N, 2C, 2D, 2H, 2S, 2N, 3C, 3D... upto 7N. Bidding can only proceed left to right in the order. That is, a player can only choose, other than a "Pass"(P) or no bid option, a bid on the right of the last bid he heard. In addition a "double" can be used (bid) over an opponent's bid, and a "redouble" over an opponent's double. Thus the total number of bids to exchange information is limited.

A further constraint is imposed by the fact that bidding higher levels implies a greater commitment in terms of the number of tricks, which may not be backed up in play by the strength of the cards. Thus making a bid at a higher level has to be a judicious choice between the need to convey more information and the risk of overshooting the estimate in the number of tricks that can be made. Many bidding systems have evolved in an effort to make an optimum use of the limited bidding space for conveying maximum information.

Bidding thus is a distributed cooperative process whereby a team of two players endeavors to find an optimum contract. During the course of a bidding sequence a player may aim to convey the suits held by him, support for the suits held by the partner, the strength of his hand in terms of high cards (such as Ace, King, Queen and Jack, which are usually encoded in terms of numerical values), and other kinds of strength (such as single card and void in a suit). The estimated strength of the hand changes dynamically, depending on the information received from the partner. Information exchange may be intercepted by the bids opponents may make. Or additional information may be available from the bids opponents make.

2.1 Information in the hands

The fifty two cards of a pack can be dealt in about 5.36×10^{28} ways. Thirteen cards can be picked from fifty two in 6.35×10^{11} ways. There are 39 possible hand patterns as shown in table 1. These range from the most balanced 4-3-3-3 shape (4 cards of one suit, and 3 each of the rest) to the most unbalanced 13-0-0-0. These patterns and the percent probability with which they appear are shown in table 1. The percentage probabilities listed under TOTAL do not distinguish between the suits. The column named SPECIFIC lists the percentage probabilities when the suit is also specified along with the length. For example, the percentage probability of getting a 8-2-2-1 hand (pattern number 24) is 0.1924 percent. However, the percentage probability of being dealt with 8 clubs, 1 heart, and 2 each of diamonds and spades is one twelfth of 0.1924, i.e., 0.016 percent.

The aim of a bidding system is to convey the maximum amount of information for cases which are most frequent. Obviously all possible hands cannot be described because of "low bandwidth" of the bidding channel. Usually, for instance, players may open the bidding from a choice of 21 bids beginning from the pass level upto the level 4. It would be wasteful to reserve one of these bids to describe a hand, say of 7-6-0-0 shape, which may occur once only in twenty thousand times.

Since four-card and five-card suits are most likely, most bidding systems are generally designed taking this *a priori* knowledge into account. Thus, a simple bidding rule might be "IF you have at least 4 cards in Spades THEN bid Spades". A bidding system is usually a collection of such rules, and also it contains rules which can take into account the strength of the high cards. For example, a rule may say "IF you have 13 points THEN you can make

Table 1. Distribution of hand patterns.

No.	Pattern	Total	Specific	No.	Pattern	Total	Specific
1	4-4-3-2	21.5512	1.796	24	8-2-2-1	0.1924	0.016
2	4-3-3-3	10.5361	2.634	25	8-3-1-1	0.1176	0.010
3	4-4-4-1	2.9932	0.748	26	8-3-2-0	0.1085	0.005
	—	—	—	27	8-4-1-0	0.0452	0.002
4	5-3-3-2	15.5165	1.293	28	8-5-0-0	0.0031	0.0003
5	5-4-3-1	12.9307	0.539		—	—	—
6	5-4-2-2	10.5797	0.882	29	9-2-1-1	0.0178	0.001
7	5-5-2-1	3.1739	0.264	30	9-3-1-0	0.0100	0.0004
8	5-4-4-0	1.2433	0.104	31	9-2-2-0	0.0082	0.0007
9	5-5-3-0	0.8952	0.075	32	9-4-0-0	0.0010	0.00008
10	6-3-2-2	5.6429	0.470	33	10-2-1-0	0.0011	0.00004
11	6-4-2-1	4.7021	0.196	34	10-1-1-1	0.0004	0.0001
12	6-3-3-1	3.4482	0.287	35	10-3-0-0	0.00015	0.00001
13	6-4-3-0	1.3262	0.055		—	—	—
14	6-5-1-1	0.7053	0.059	36	11-1-1-0	0.00002	0.000002
15	6-5-2-0	0.6511	0.027	37	11-2-0-0	0.00001	0.000001
16	6-6-1-0	0.0723	0.006		—	—	—
17	7-3-2-1	1.8808	0.078	38	12-1-0-0	0.000003	0.0000003
18	7-2-2-2	0.5129	0.128		—	—	—
19	7-4-1-1	0.3918	0.033	39	13-0-0-0	0.09×10^{-8}	0.02×10^{-7}
20	7-4-2-0	0.3617	0.015		—	—	—
21	7-3-3-0	0.2652	0.022		—	—	—
22	7-5-1-0	0.1065	0.005		—	—	—
23	7-6-0-0	0.0056	0.0005		—	—	—

Values listed under specific are for named suits having specified length. Numbers under total sum up the values for all possible ways of choosing suits for the given pattern or shape.

an opening bid". Here 13 is a numerical encoding of the card strength, which counts an Ace as 4, a King as 3, a Queen as 2 and a Jack as 1. A combination of the above two sets of rules may thus allow a player to bid 1 spade (1S) with the following hand,

Q 8 6 3 2, K J 5, A 9 7, Q J,

where the suits from left to right are: Spades, Hearts, Diamonds, and Clubs.

2.2 Why not rule-based system?

One could thus think of building a rule based system which helps in making a bid. The difficulty with such an approach is that players normally use these rules only as a guideline, and often they make bids the reasoning for which they cannot articulate in terms of the given rules. For example, for a hand containing 4 spades and 4 diamonds, a rule may suggest opening 1S, or possibly 1D. But for the two hands given below, which are only

slightly different, a player may choose different bids as shown:

A J 9 3, K 8 4, K 7 4 3, A 9 – bid : 1S

A J 9 3, K 8, K 7 4 3, A 9 4 – bid : 1C

This change comes because of the subtle reasoning process the player uses, since he is also concerned about his *next* bid. There are other patterns too in the hand for such a reasoning. For example “K84” is a “support” for a possible bid of 2H by partner, while “K8” is not. “A94” in clubs, on the other hand is an openable suit if the hand has no five carder. One could possibly list all such possibilities as rules, but the number of rules may be too many. One of the aims of our experiment is to test whether an ANN can capture this reasoning based on input patterns which expert players seem to be using. A neural network could capture the implicit knowledge because ANN learns from examples. Thus in this case an expert need not articulate all the rules, but only adopt them appropriately while making a bid.

Another interesting feature of bidding is that sometimes the bid for a given hand need not be unique. Two or more alternate bids for the same hand may seem reasonable to the human players. The choice is also likely to be influenced by contextual factors in the game like the past history of the match and the disposition of the player. Thus in general, a bid is judged subjectively with respect to the available knowledge. But often this knowledge is a dynamic process as an *a posteriori* analysis causes new associations to be learnt.

2.3 Scope – opening bids

In this paper we consider only the opening bids, when the player is first to bid, and no other information is available or relevant. This is the simplest case where the player has to make a bid based only on the pattern in the 13 cards of his hand. The bidding is not influenced by other factors such as bids by other players, the position of the game, etc. However the bid is influenced by the bidding convention that players adopt. The convention system (bidding system) helps a player to look for pattern features to evaluate the strength of a given hand.

The reason for considering the opening bid is that it is the only bid that is made solely on the basis of the pattern in the given hand. Any bid made after the opening bid utilizes the information conveyed by the preceding bids also, besides the given hand pattern. For example, if your partner has opened with a bid of 1S, then you may bid 2S with the following hand,

Q 9 8 5, A 3, J 8 5 4, 9 6 5.

Clearly, the bid of 2S is based on the information that the partner has a good hand with a spade suit. This information has to be *combined* with the patterns of the given hand to produce the bid of 2S. Also, in some sense the information contained in the given hand is in a raw form, while the information received from the bid of 1S is in an abstracted form. Combining such information is a more complex process and beyond the scope of this paper.

2.4 Bidding convention

The bidding convention assumed in the current exercise is the *Standard American bidding* system. This was chosen because it is less artificial than other systems, for example *Precision bidding* system. All systems use some bids for which the semantics is not straightforward. Artificial systems use more of such bids, where the meanings are very well defined, but may not have a direct relation with the obvious pattern in the hand. For example, a 2D bid in the *Precision* system denies a diamond suit, rather than promising the suit. Natural systems, on the other hand, have more direct mappings. In the process they also allow for subtle changes in the bidding rules, often called “expert judgment”, a phenomenon we hope to capture.

There are about 10^{11} different hands a player can start with. A look at table 1 shows that hand patterns containing seven or more cards in a suit constitute less than 5% of this number. A random hand generator will find it difficult to produce representative samples of all possible patterns. Seven card or longer suits may be normally opened only at third level or above. Long training cycles would become necessary if the training set were large enough to include all these less frequently occurring patterns. In this paper we have limited our study only to low (upto 2) level bids. This decision was arrived at by a process of experimentation described later in the paper.

3. Why artificial neural networks?

Artificial neural networks are attractive because the networks learn from examples. If in the process they can generalize (Hertz *et al* 1991), then they can also provide a useful interpolation capability. Due to generalization capability, it is possible that the system may capture the subtle reasoning process used in making a bid, which may be difficult to incorporate explicitly in the rule-based formulation.

Table 2. Variations in players' bids.

Set no.	No. of people who bid	No. of hands for which their bids differed	%
1	2	14	17.50
2	3	17	21.25
3	2	26	32.50
4	2	24	30.00
5	2	14	17.50
6	2	13	16.25
7	2	17	21.25
8	2	29	36.25

Sets of 80 hands were bid by human players. It can be observed that for every set the players differed on a significant number of hands, suggesting that more than one correct bid may exist.

The bidding problem, as we see it, is an exercise in high level perception. It involves mapping complex patterns in a hand onto a single output corresponding to the bid for the hand. ANN's do precisely this. Given a set of input patterns and the corresponding target patterns, they try to capture the implicit relation between the two. Once the network has been trained to generalize, then it can respond meaningfully to a new pattern.

Realizing a network even for the opening bid for a given hand is not a trivial task. In fact the relation between the input pattern and the corresponding output bid is not unique even amongst expert players as shown in table 2 from the data collected at a bridge tournament.

We have decided to explore the possibility of training a multilayer feedforward neural network with a backpropagation algorithm (Rumelhart *et al* 1986) to capture the implicit reasoning from several examples of input(pattern)-output(bid) pairs of data. In order to perform this study several other issues need to be considered. Some of them are described below.

3.1 Representation

The input can be represented as raw data as shown in figure 1a, where the input selects 13 of the 52 input nodes, representing the 13 cards that have been dealt. Thus the input layer here consists of 52 nodes, each of which may have a value either 1 or 0.

The input can also be represented in the form of feature patterns as shown in figure 1b. These patterns are based on the evaluation of the strength of the hand by a bidding system. In this representation there are 16 nodes in the input layer, which can take values between -1 and $+4$. Thirteen nodes are used to represent the cards, while three are used as markers (-1) between suits. In this representation an attempt was made to feed some feature information in the form of relative weights given to various cards. These weights were close to the points given to high cards in most bidding systems. Bidding systems use numerical values for Aces, Kings, Queens and Jacks, as described earlier. In our representation we gave appropriate values to the other cards as well. Our initial experimentation showed that the first representation was preferable, as the feature representation is somewhat subjective. During training we found that the network converged with the first representation, whereas the feature based representation failed to converge in some cases. This is interesting because the information in the second case is in an interpreted, or an abstracted form. It appears that abstraction from raw data, if not done properly, may not be useful for obtaining generalization by a network. Experiments described in this paper therefore use the first representation.

It is also interesting to note that from the representation of the raw data as shown in figure 1a the network is able to extract the necessary pattern information to capture the reasoning process without explicitly describing the pattern. This is to be contrasted with pattern recognition problems in speech and vision, where features from the input data have to be extracted carefully and fed as input to the neural network. In fact the performance critically depends on the feature extraction stage in such cases.

The second important issue is the representation of the bids. A straightforward representation of all possible bids each by a separate output node leads to problems in training. Since the number of sample hands for higher level bids will be fewer relative to hands for the lower level bids, the network may not converge for a given set of limited data.

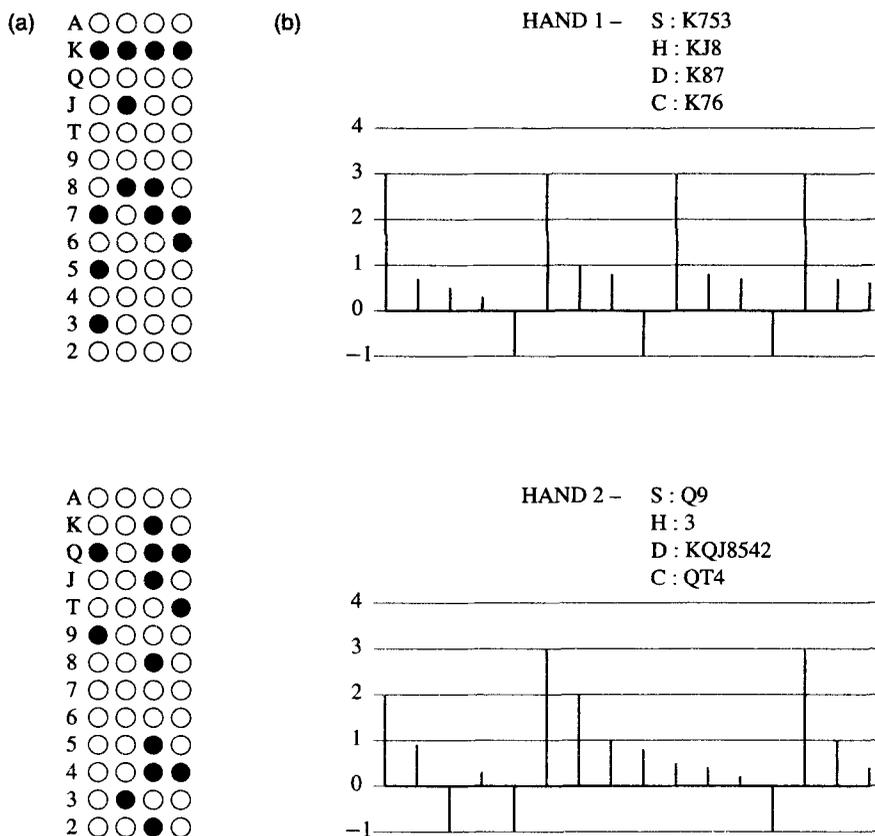


Figure 1. Illustration of input layer patterns for two hands. In the first representation (a) the input is in raw form, while in the second (b) some features (high card points) have been extracted. The networks perform better with the raw information.

To study the convergence behaviour of the network, we have conducted experiments with networks limited to different number of output nodes, corresponding to different levels of bidding. These experiments suggest that a hierarchical network approach may be needed for implementing the opening bid problem.

A third issue is the interpretation of the output results. It is quite possible that during testing the network produces a bid which is different from a player's bid. But then the player should also find out whether the network bid is also reasonable for the given hand.

We have provided one node in the output layer for each possible bid. Starting with a network with output nodes for all possible bids, upto seven levels we have reduced the size of the network upto two level bids at which it can be trained within a reasonable time.

There are two ways in which the output can be interpreted. One way is to accept the bid represented by the node giving highest activation value for the output. But one has to be careful here, as two likely bids may have nearly equal output activation values. The second way is to have a threshold, so that a node with an activation value above the threshold is considered to be a plausible bid. This has the drawback that no such node may exist for a

particular example, and the output will have to be reexamined. We have adopted the first strategy of picking the highest node with the activation value as the output bid.

4. Backpropagation network for opening bid problem

A multilayer feedforward network was used with the backpropagation (Rumelhart *et al* 1986) algorithm for training. There are 52 nodes in the input layer. The number of nodes in the hidden and output layers were varied as described in the following experiments.

4.1 Data generation

The hands used for training of the network were generated by a program which simulates shuffling of cards. The distribution of the hand patterns generated by the program matches the distribution given in table 1. A representative set of 19 hands are given in table 3. Note that the hands which contain suits of maximum length 5 constitute about 80% of all the hands. On the other hand, bids of 2H and 2S, in the configuration named 2-NT network, require the following features: a six-card suit, with no singleton or void in the hand, about 8 to 10 high card points, with most of the high cards in the bid suit. To successfully train the network for these bids, it is necessary to have a large number of these samples in the training set. This in turn would mean a correspondingly large training set and hence a large training period. We have used a generating program to produce hands according to a given set of constraints, for example, length of heart suit to be at least 6, number of points to be at least 6. In this way, we can produce more hands for which we want the system to learn patterns. But this method of generating large number of sample hands does not reflect the pattern environment of the bidding problem. A pattern environment is described by the patterns together with the probabilities of occurrence of these patterns.

4.2 Representation of input patterns

The input layer has 52 nodes, one for each card. Each node has a value of 1, if the card is present, or 0, if the card is not present in the given hand. Thus the input vector marks out a subset 13 out of the 52 cards. Each training data pair consists of 13 activated nodes in the input layer and one desired activation node in the output layer. For example, the first hand in table 3 and the corresponding input pattern vector are given by

Hand: K753-KJ8-K87-K76
Input pattern: 0100000101010010100100000001000011000000100000110000

The desired output is activation of the node corresponding to the bid of 1C in the 2-NT network. The 2-NT network has output nodes corresponding to all the five one level bids, all the five two level bids, a "pass"(P) bid and an "unknown"(U) category bid. The "unknown" category includes all hands corresponding to the levels above the 2-NT level.

Table 3. Sample hands generated by a program.

No.	Hand (S)-(H)-(D)-(C)	Points	Desired bid (by authors)
1	K753-KJ8-K87-K76	13	1C
2	Q9-3-KQJ8542-QT4	10	3D
3	863-KQJT954-K9-5	9	3H
4	T853-KT83-A9-AQ7	13	1S
5	J62-AJ5-K652-942	9	P
6	7-AQT976-874-K85	9	2H
7	Q87-KQJT63-92-95	8	2H
8	A98-K2-AQJ62-AT6	18	1N
9	QT75-A73-J3-AKJ8	15	1C
10	T75-K52-AK86-QJ7	13	1D
11	Q94-QJ832-A95-AK	16	1H
12	AKT7642-Q53-8-T5	9	3S
13	J986-72-AQ543-J8	8	P
14	A7-K94-KT65-AK74	17	1N
15	A92-4-AK82-AJ832	16	1C
16	QT-76-AQJ752-AJ3	14	1D
17	A-KQJ9873-J98-T6	11	1H
18	AK842-AKT93-6-82	14	1S
19	AQJT-QT4-AJ-QT92	16	1N

Bids in the table were made by the authors, to illustrate the preparation of the training set. In this training set some less frequent hands are present which are generated specially to ease learning.

5. Network training

This section describes the development of the network architecture. The training algorithm used was the backpropagation algorithm. The network architecture was evolved by a trial and error process. Initially we have started with a network aimed at capturing all the bids. These larger networks failed to converge. Therefore we have pruned the network size by reducing the number of output nodes first. Thus we have designed smaller networks with fewer output nodes. The process of pruning suggests a method to evolve a set of modular networks organised in a hierarchical fashion, where each module specialises on a subset of bids. In the following, we describe our trial experiments for evolving a suitable architecture for the bidding problem.

5.1 Experiment 1

An approach using 13 output nodes to capture all the bids was explored. Seven nodes were assigned to the 7 levels of bids, five were assigned for the suit and one node for “pass”

Table 4. Performance of the 1-level network.

Hidden nodes	tss	No. of epochs	No. of hands
40	22.20	5000	1600
45	23.80	5000	1600
50	23.01	5000	1600
55	21.40	5000	1600
60	19.12	5000	1600

Here 52 input nodes and 7 output nodes were used.

bid. Thus, for each input, except for a “pass” hand, two output nodes are expected to be activated, one for the level of the bid, and the other for the suit. The training set consists of 900 hands. Different network architectures were examined. Some of them are:

- (a) 30 and 20 nodes in the 2 hidden layers,
- (b) 30 and 6 nodes in the 2 hidden layers.

It was found that the network did not converge. This was probably because there were a large number of bids (some 2-level bids and all 3- or higher level bids) for which very few training patterns were available. Since opening bids are rarely made at levels higher than 3, quite a few of the output nodes in this representation were not used. This is motivated by the fact that the two components of the bid, viz. level and suit, may not be independent of each other. For subsequent experiments, we have decided to use a simpler format for output nodes with one output node for each bid.

5.2 Experiment 2 – The 1-level network

To resolve the convergence problem, we reduced the number of output nodes to seven, by restricting the bids to 1 level only, including “pass” and the “unknown” category bids. Then we observed that the network converged. Actually, our motivation behind this experiment was to verify the feasibility of learning input patterns.

The resulting network consists of 52 input nodes, 7 output nodes and one hidden layer. The number of nodes in the hidden layer was varied to study its effect on the performance. Table 4 gives the total sum of squared error(tss) for 40, 45, 50, 55 and 60 hidden nodes. The network with 60 hidden nodes gave the highest accuracy of 92% correct bids on the test set, when compared with the bids made by an expert player. The test data consisted of 200 randomly generated hands. Some test results are given in appendix B and in table 5. It should be noted that while evaluating the performance, if the output of the network was also acceptable by the expert player as a possible bid, then it was taken as correct output. Errors in the 1-level network are likely to occur because there are many borderline hands which may fall into the categories of either 1-level or 2-level. Fewer such hands are likely at the borderline of higher level bids, such as between 2-level and 3-level.

Table 5. Bids made by the 1-level network.

No.	(S)-(H)-(D)-(C)	Points	Expert bids	1-level network
1	KJ6-Q62-K94-A864	13	1C	1C
2	-4-AT942-AKJ8974	12	U	1C
3	J8-KQT643-KQ92-T	11	P	1H
4	K64-AQ874-AT953	13	1C	1C
5	AKQT9-J765432-9-	10	U	P
6	AT63-KQJT6-AT73	14	1D	1D
7	AQT87-K4-AKT4-Q7	18	1S	1S
8	AKJ9542-QJ852-6	11	U	U
9	62-AT3-J942-AKT7	12	1C	1C
10	Q653-AKJ8-AJ-974	15	1H	1S

Strong imbalanced hands were labelled “unknown” (U) for training purposes. However sometimes (e.g. example 2) the system did better by opening 1C. Also, in example 3, the network’s bid seems to be better! The discrepancy in the last example is also typical of human players.

5.3 Experiment 3 – The 2-NT network

We consider a network to include 2-level bids. Here, the number of output nodes are 12, one each for P, 1C, 1D ... up to 1N, 2C, 2D ... up to 2N and U.

Initially we attempted to train the network with a training set conforming to the theoretical distribution of hand patterns. But the network could not be trained. The reason is that the network was unable to learn the patterns for 2C, 2D, 2H, 2S, 2N bids since they are very rare. Getting suitable samples of such hands requires large numbers of training sets. Instead we have decided to selectively insert the patterns, which are rare, into the training set. Nearly 250 hands from these 2-level bids were added along with nearly 1350 hands from the 1-level bids. As a result we have obtained 1600 hands to train the network.

This network was trained using five different architectures having 40, 45, 50, 55, and 60 hidden nodes. The performance of the network for this training set is shown in table 6.

Results produced by the network are given in appendix B and in table 7. The network has bid correctly for about 92% of the test hands, which were not part of the training set. On the training set hands, the performance of the networks varied from 95% to 100%.

Another point is that initially we planned to give only positive samples of hands for the bids which we wanted the system to make. But we found that for the system to perform well, we also had to give a large number of hands for which we did not want the system to make a bid. So, we have introduced all those hands under the bid “unknown”. The performance of the resulting network is shown in the table 7.

Table 6. Performance of the 2-NT network.

Hidden nodes	tss	No. of epochs	No. of hands
40	11.01	5000	1600
45	12.78	5000	1600
50	10.52	5000	1600
55	11.35	5000	1600
60	8.14	5000	1600

Here 52 input nodes and 12 output nodes were used.

5.4 A proposed architecture

It can be clearly observed that smaller networks are easier to train, and consequently they also perform better. On the other hand, looking at the task environment, one can see that all the bids made at higher levels are specialized. In addition, they deal with hands that are less frequent. This is consistent if one associates a cost with each bid proportional to the level at which it is made. Hands with four-card and five-card suits are most common (80%) and the bidding systems are designed to use the cheaper (low level) bids for these hands. To design a complete ANN system would require sufficient training samples for hands of all patterns. It appears reasonable to consider the high level specialized bids as exceptions, and train different networks to deal with them. Thus one would have a modular structure of the network, each module catering to a specialized situation. These

Table 7. Bids made by the 2-NT network.

No.	(S)-(H)-(D)-(C)	Points	Expert bids	2-NT network
1	KJT4-QT9762-A-T3	10	P	1H
2	A8-QT2-KQJ97532-	12	1D	1D
3	7-QJ98-QJ8752-T3	6	P	P
4	Q-KQJ852-73-K983	11	P	2H
5	K94-A5-AJT743-74	12	1D	1D
6	AK8742-J4-K852-9	11	1S	1S
7	AT9872-5-K54-643	7	2S	2S
8	A-K82-QJ9642-Q97	12	P	1D
9	J6-JT-8-AKQ98654	11	3N	1C
10	84-KJT9542-Q8-A3	10	2H	2H

Many experts would open hand 4 with 2H, because the key feature – long solid suit is present. In hand 8 the system has in fact done better by opening 1D. In hand 9 it possibly had to choose between “unknown” and 1C, since it does not know the 3N bid, which is very specialised.

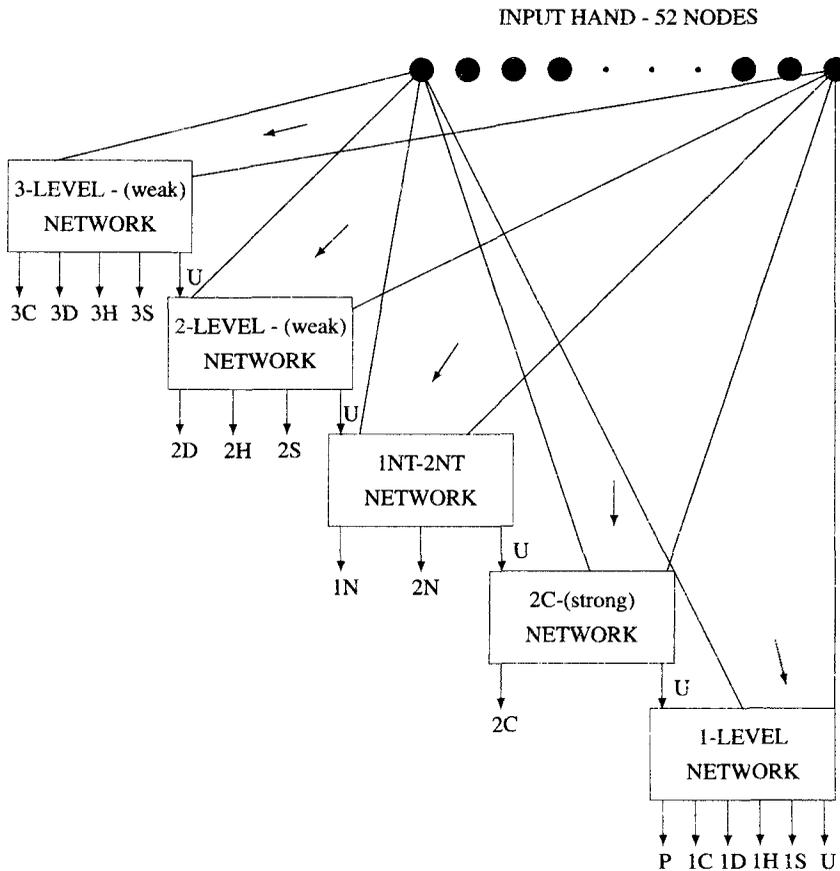


Figure 2. A hierarchically structured neural network. Specialized modules, which can be trained individually, look for specific patterns before passing on control to the next module in the hierarchy. The input hand is fed into each module.

can be organized hierarchically so that control flows from the exceptions to the general hands.

A possible hierarchical architecture is shown in figure 2. Each network in the hierarchy looks for specific patterns. For example the network for the 3-level weak bids looks for a pattern of seven-cards in a suit with no other four-card suit and about 7 to 10 high card points with most of the points being in the bid suit. It either gives one of 3C, 3D, 3H, 3S bids or classifies the hand as unknown. If the “unknown” output node is activated, it in turn activates the following network. In this way control flows in a cascaded form from the specialized networks to the more general one. Preliminary investigations suggest that the specialized networks can be trained easily with much smaller sets of hands.

But it is not clear whether a hierarchical architecture could be developed based only on these specialized networks. There are several issues in training a network made of many modules. One is the flow of control and the interaction between different modules. The different modules may not be independent of each other. There may be overlap among

the output classes. It is not clear how the training set for the different modules should be organized.

6. Conclusions

The studies reported in this paper clearly demonstrate that a neural network can be trained to capture the implicit reasoning used for bidding a hand in bridge. From the above experiments, the following points are worth noting:

- (a) A large network is difficult to train.
- (b) A 1-level network performs well.
- (c) Even a 2-NT network does not learn well because of lack of data for 2-level bids.
- (d) But if specific data are added for the 2-level bids, then the network may perform well.
- (e) This leads to the idea of a modular hierarchical network. It may be possible to optimize the network design taking into consideration probabilities of input–output patterns.
- (f) Finally, it requires a more sophisticated network architecture to construct a system to bid hands through a complete auction. This is because of the necessity of combining raw information (the hand) with the processed information (the bids).

The present study clearly brings out several interesting research issues to be explored using neural network architectures. The first issue is the representation of the input data. In situations like card games, representation in raw form appears preferable, as any feature representation is likely to be subjective and may result in loss of information. In contrast, in problems dealing with speech and image data, it is essential to represent the data in a manner that reflects the auditory and visual sensory processing, respectively. Errors in the feature representations are usually responsible for poor generalization performance in the pattern recognition tasks involving speech and image. The second major issue is training a network with patterns occurring with widely different probabilities. This is a difficult issue in many practical problems, including for example speech, where different speech sounds occur with widely different probabilities.

References

- Haykin S 1994 *Neural networks – A comprehensive foundation* (New York: Macmillan)
- Hertz J, Krogh A, Palmer R G 1991 *Introduction to the theory of neural computation* (Reading, MA: Addison-Wesley)
- Khemani D, Ramakrishna R S 1989 Bridge: A benchmark for knowledge based planning. *J. Artif. Intell., Cogn. Sci. Appl. Epistemol.* 6(2/3): 137–151
- Rumelhart D E, Hinton G E, Williams R J 1986 Learning internal representations by error propagation. In *Parallel and distributed processing* (eds) D E Rumelhart, J L McClelland (Cambridge, MA: MIT Press)

Appendix A. A note on contract bridge: the problems

Contract bridge is played with a regular pack of 52 cards dealt randomly and equally among 4 players. Let us call them North, South, East and West, according to their position at the table. North and South are partners, as are East and West. The cards are ranked in the order Ace, King, Queen, Jack, 10, 9, ...2 in each suit. Each player plays a card, in clockwise order and the highest ranking card wins the trick. Thirteen such tricks are played, each time the winner of the previous trick starting play. This constitutes one *deal* or one *hand*.

There are two stages of play in each deal, viz. bidding, followed by the play of cards. The goal in a deal is to maximize *points*. The points essentially depend upon bidding. Bids are made for the number of tricks the side promises to make, given the stated “trump” suit. Eventually the highest bid is accepted in each deal. This is known as the *contract*. Generally, the higher a side bids, the more points it is likely to win, *provided* it can fulfil the contract. That is, if the side can make the number of tricks it has bid for, it wins some points. Let us call them success-points. If it loses, then the opponents get some points instead, which we can call penalty-points.

The straightforward goal in bidding is to bid the *highest* number of tricks one thinks the side *can* make. That is, to maximize success-points won. The means used in this process are the following,

- (a) Evaluation of own hand.
- (b) Communication with partner.
- (c) Projection of play.

Of these, the first two are simpler and can possibly be handled by heuristic methods. The third is more difficult, as it would involve constructing plausible distributions (based on the bids heard, and on probability) and then projecting the play. Using neural nets we hope to implicitly capture all three components for the opening bid situation.

A more complex goal is to make a *sacrifice bid*. This essentially means intentional overbidding, over an opponent bid, with the hope that the penalty-points loss will be lesser than the opponents' expected success-points gain, thus being an overall gain.

Even more complex goals are to sabotage the opponents' communication. This may mean consuming the bidding space (jamming the communications channel), or even making “false” bids to confuse opponents. In the process, an enterprising planner may make an “advance sacrifice” to “push” the opponents higher than they can manage, or to escape with a lighter penalty.

Considering that all these processes happen when the planner can see only one hand, one observes that bidding is probably the more difficult part of the game.

Once bidding is over, the goal for the play stage has been defined. One side has the contract, and is required to make the bid number of tricks. At this stage one player of the contracting side (called the dummy) exposes the cards to everybody, while the other (called the declarer) plans and executes the play. The opposing side (called defenders) are said to *defend* the contract, trying in fact to defeat the accomplishment of the contract by the declarer.

One can see that the situation at this stage is not symmetric. The declarer knows the entire strength of his side, and is in total control of the play of the cards. He is also aware of the

entire assets of the defence, in terms of material strength, since they have the remaining 26 cards. Each defender knows only his own hand, and cannot see his partner's hand. Therefore the two defenders have to combine their efforts to try and achieve the goal. This necessarily involves (formal) communication between the two. Both can see the dummy also.

Since the cards of all the players cannot be seen, one cannot project moves into the future. Methods like minimax search are therefore ruled out immediately. Instead, the success of a strategy can only be estimated based on the probabilistic distribution of cards, and any information gleaned from the communication taking place. The strategies themselves are derived from knowledge about the various known methods of tackling different card combinations.

The straightforward goal in the play of the hand is to make the number of tricks as stated in the contract. The emphasis is on maximizing the probability of success. If success is assured, then the goal can be revised to increase the number of tricks won, as some more points can then be gained. If success seems unlikely, then a planner may even choose to minimize losses, i.e., the penalty-points won by the opponents. Like in bidding, the planner may attempt to do better than par, by exploiting the incomplete information that the opponents have. This may introduce complex "meta-level" goals of protecting information, or sending out misleading signals.

Thus, we see that unlike games like chess, where a clear-cut strategy of aiming for the minimax value (saddle) points is meaningful, in bridge one has to largely grapple with incomplete information. In the face of such uncertainty, planning in bridge can only be a complex knowledge intensive activity.

Appendix B. Sample outputs from bidding networks

The bid made by the 1-level and 2-NT network for some hands are shown in table B1. Also included are bids made by two players earlier. It can be seen that for most of the hands the networks perform well.

Table B1. Expert bids and the bids made by the 1-level and 2-NT networks for the same hand.

(S)-(H)-(D)-(C)	Points	Expert bids	1-level network	2-NT network
AT5-J983-K5-AJT5	13	1C, 1H	1C	1H
AQ6-A752-AT2-KQ3	19	1C, 1H	U	1C
A95-AQ9-543-AQ93	16	1C, 1N	1N	1C
K85-K94-KQJ94-96	12	P, 1D	1C**	1D
T-AK6-J8642-QJ63	11	P, 1D	P	1D
KJ3-T-QT9643-AQ9	12	P, 1D	1D	1D
AQ84-AJ984-J92-9	12	P, 1H	1S	P
AQ8-AQJ6-63-AJ94	18	1H, 1N	1C	1C
974-A732-AK6-QJ6	14	1C, 1H	1C	1H
8753-4-AT5-AKT98	11	P, 1C	P	P
98-AK6-K854-AK43	17	1C, 1N	1C	1C
QJ72-K743-K9-A54	13	1S, 1H	1S	1S
Q72-A98-Q64-KJ92	12	P, 1C	P	P
JT4-J8-AKQ6-AQ64	17	1C, 1N	1N	1C
AKQ97-52-T3-9853	9	P, 2S	U	2S
K843-A64-A52-AT3	12	P, 1S	1S	1S
K832-AKJ3-QT5-K6	16	1N, 1H	1S	1H
-Q98762-KQ753-Q3	9	P, 1H	P	P
8-QJ5-A74-KJ9875	11	P, 1C	P	P
K8-AT94-JT7-AT65	12	P, 1H	1H	P
97-5-AKQ8754-AK2	16	2C, 1D	1D	1D
96542-AQJ2-AKQ7	16	1S, 1D	1S	1C
AJ75-4-KJT6-AKT8	16	1D, 1C	1C	1D
A6-QJ86-J94-KQJ9	14	1H, 1C	1C	1C
T8-AT5-K976-A874	11	1N, P	1C	1C
KQ-QJ7642-QJ92-T	11	P, 1H	1H	1H
K87-J8-AKQ8-AKQ2	22	1C, 2N	U	2N
KT84-KQT7-K63-T2	11	P, 1S	P	P
AJ-Q64-J42-AKJ85	15	1N, 1C	1C	1C
A632-3-96-KQJ863	10	P, 1C	1C	P
965-AT8762-J7-A9	9	1H, P	P	P
Q7-A74-A2-KQJ972	16	1D, 1C	1C	1C
3-T6-KT74-AQJT74	10	P, 1C	P	P
AKJ3-982-J85-KQ5	14	1S, 1C	1S	1C
92-KQ6-A93-AQJ73	16	1N, 1C	1C	1C
94-KQ3-AQT2-J632	12	P, 1D	1D	1D
-A874-AJT52-AQT5	15	1C, 1D	1D	1D
2-JT86-AQT9864-Q	9	P, 3D	U	P

(Continued)

Table B1. *Continued.*

(S)-(H)-(D)-(C)	Points	Expert bids	1-level network	2-NT network
K95-94-A852-AT83	11	P, 1D	P	P
Q-A743-AJ5-A8732	15	1H, 1C	1C	1H
-A642-AKT83-J943	12	1H, 1D	1H	P
AJ76-AQ73-75-J87	12	1S, 1H	1S	1S
AQJ73-KT85-42-72	10	1S, P	P	1S
KQ98754-83-Q7-A2	11	1S, 2S	1S	2S
QJ6-Q-KT864-A854	12	1D, P	1D	1D
743-Q65-AKT7-KQ9	14	1C, 1D	1N	P
J4-AJ654-J53-KJ9	11	P, 1H	P	P
K8653-3-7-AKJ753	11	3C, 1C	1C	1S
AT5-AK-Q9532-KQ3	18	1N, 1D	U	1N
A3-AT2-KT64-AKQJ	21	2C, 1C	U	1N**
75-A3-AQT632-AK5	17	1N, 1D	1D	1D
K53-QJ3-QJ8753-3	9	3D, P	U	2D
Q-KQ65-AKQ74-Q43	18	1N, 1D	1H	1D
A7-K8643-K863-A8	14	1S, 1H	1H	1H
K-AJ8-AK965-QT87	17	1N, 1D	U	1N
A943-Q432-Q52-A2	12	1C, P	P	1S
K87-AT9-J87-A972	12	1C, P	P	1C
JT86-AK43-76-AQ7	14	1S, 1H	1S	1S
982-Q652-AKQT-K5	14	1H, 1D	1D	1H
KQJT85-A96-743-8	10	3S, 2S	P	2S
AT982-QT2-A3-J65	11	1S, P	1S	1S
AJ32-AQ93-5-AQJ2	18	2C, 1C	U	1N**
76-AKQJ653-7-KQ9	15	1H, 4H	1H	1H
5-AJT8753-K4-K93	11	3H, 1H	1H	1H

** These bids are incorrect.