

LEARNING TO SCORE FINAL POSITIONS IN THE GAME OF GO

E.C.D. van der Werf, H.J. van den Herik, J.W.H.M. Uiterwijk
*Institute for Knowledge and Agent Technology, Department of Computer Science,
Universiteit Maastricht, P.O. Box 616, 6200 MD Maastricht, The Netherlands*
{e.vanderwerf,herik,uiterwijk}@cs.unimaas.nl, <http://www.cs.unimaas.nl/~vanderwerf/>

Abstract This paper presents a learning system for scoring final positions in the Game of Go. Our system learns to predict life and death from labelled game records. 98.9% of the positions are scored correctly and nearly all incorrectly scored positions are recognized. By providing reliable score information our system opens the large source of Go knowledge implicitly available in human game records, thus paving the way for a successful application of machine learning in Go.

Keywords: Go, learning, neural net, scoring, game records, life and death

1. Introduction

Evaluating Go positions is one of the hardest tasks in Artificial Intelligence (AI). In the last decades, stimulated by Ing's million-dollar price for the first computer program to defeat a professional Go player (which has expired unchallenged), Go has received significant attention from AI research (Bouzy and Cazenave, 2001; Müller, 2002). Yet, despite all efforts, the best computer Go programs are still no match even for human amateurs of only moderate skill. Partially this is due to the complexity of Go, which makes brute-force search techniques infeasible on the 19×19 board. However, on the 9×9 board, which has a complexity between Chess and Othello (Bouzy and Cazenave, 2001), the *current* Go programs perform nearly as bad. The main reason lies in the lack of good positional evaluation functions. Many (if not all) of the current top programs rely on (huge) static knowledge bases derived from the programmers' Go skills and Go knowledge. As a consequence the top programs are extremely complex and difficult to improve. In principle a learning system should be able to overcome this problem.

In the past decade several researchers have used machine-learning techniques in Go. After Tesauro's (1995) success story many researchers, including Dahl (2001), Enzenberger (1996) and Schraudolph et al. (1994), have applied Tem-

poral Difference (TD) learning for learning evaluation functions. Although TD-learning is a promising technique, which was underlined by NEUROGO's latest performance at the 21st Century Championship Cup (Myers, 2002), there has not been a major breakthrough, such as in Backgammon, and we believe that this will remain unlikely to happen in the near future as long as most learning is done from self-play or against weak opponents.

Over centuries humans have acquired extensive knowledge of Go. Since that knowledge is implicitly available in the games of human experts, it should be possible to apply machine-learning techniques to extract that knowledge from game records. So far game records have only been used successfully for move prediction (Enderton, 1991; Dahl, 2001; van der Werf et al., 2002). However, we are convinced that much more can be learned from these game records.

One of the best sources of game records on the Internet is the No Name Go Server game archive (NNGS, 2002). NNGS is a free on-line Go club where people from all over the world can meet and play Go. All games played on NNGS since 1995 are available on-line. Although NNGS game records contain a wealth of information, the automated extraction of knowledge from these games is a non-trivial task at least for the following three reasons.

Missing Information. Life-and-death status of blocks is not available. In scored games only a single numeric value representing the difference in points is available.

Unfinished Games. Not all games are scored. Human games often end by one side resigning or abandoning the game without finishing it, which often leaves the status of large parts of the board unclear.

Bad Moves. During the game mistakes are made which are hard to detect. Since mistakes break the chain of optimal moves it can be misleading (and incorrect from a game-theoretical point of view) to relate positions before the mistake to the final outcome of the game.

The first step toward making the knowledge in the game records accessible is to obtain reliable scores at the end of the game. Reliable scores are obtained by correct classification of life-and-death stones on the board. This paper focuses on determining life and death for final positions. By focusing on final positions we avoid the problem of unfinished games and bad moves during the game, which will have to be dealt with later.

It has been pointed out by Müller (1997) that proving the score of final positions is a hard task. For a set of typical human final positions, Müller showed that a combination of complex static analysis and search, still leaves around 75% of the board-points unproven. Heuristic classification of his program EXPLORER classified most blocks correctly, but still left some regions unsettled (and to be played out further). Although this may be appropriate for computer-computer games it can be annoying in human-computer games, especially under the Japanese rules which penalize playing more stones than necessary.

Since proving the score of most final positions is not (yet) an option, we focus on learning a heuristic classification. We believe that a learning algorithm for scoring final positions is important because: 1) it provides a more flexible framework than the traditional hand-coded static knowledge bases, and 2) it is a necessary first step toward learning to evaluate non-final positions. In general such an algorithm is good to have because: 1) large numbers of game records are hard to score manually, 2) publicly available programs still make too many mistakes scoring final positions, and 3) it can avoid unnecessarily long human-computer games.

The rest of this paper is organised as follows. Section 2 discusses the scoring method. Section 3 presents the learning task. Section 4 introduces the representation. Section 5 provides details about the dataset. Section 6 reports our experiments. Finally, section 7 presents our conclusions.

2. The Scoring Method

The two main scoring methods in Go are territory scoring and area scoring. Territory scoring, used by the Japanese rules, counts the surrounded territory plus the number of captured opponent stones. Area scoring, used by the Chinese rules, counts the surrounded territory plus the alive stones on the board. The result of the two methods is usually the same up to one point. The result may differ since one player placed more stones than the other, for three possible reasons; (1) because Black made the first and the last move, (2) because one side passed more often during the game, and (3) because of handicap stones. (Under Japanese rules the score may also differ because territory surrounded by alive stones in *seki* is not counted.). In this paper area scoring is used since it is the simplest scoring method to implement for computers.

Area scoring works as follows: First, the life-and-death status of blocks of connected stones is determined. Second, dead stones are removed from the board. Third, each empty point is marked Black, White, or neutral (the non-empty points are already marked by their colour). The empty points can be marked by flood filling or by distance. Flood filling recursively marks empty points to their adjacent colour. In the case that a flood fill for Black overlaps with a flood fill for White the overlapping region becomes neutral. (As a consequence all non-neutral empty regions must be completely enclosed by one colour.) Scoring by distance marks each point based on the distance toward the nearest remaining black or white stone(s). If the point is closer to a black stone it is marked black, if the point is closer to a white stone it is marked white, otherwise (if the distance is equal) the point does not affect the score and is marked neutral. Finally, the difference between black and white points, together with a possible *komi*, determines the outcome of the game.

In final positions scoring by flood filling and scoring by distance should give the same result. If the result is not the same, there are large open regions with unsettled interior points, which usually means that some stones should have been removed or some points could still be gained by playing further. Comparing flood filling with scoring by distance is therefore a useful check to detect whether the game is finished and scored correctly.

3. The Learning Task

The task of learning to score comes down to learning to determine which blocks of connected stones are dead and should be removed from the board. This is learned from a set of labelled final positions, for which the labels contain the colour controlling each point. A straightforward implementation would be to learn to classify all blocks based on the labelled points. However, for some blocks this not a good idea because their status can be irrelevant and forcing them to be classified just complicates the learning task.

The only blocks required for a correct score are either alive and at the border of their area, or dead in the opponent's area. This is illustrated by Figure 1. Here all marked stones must be classified. The stones marked by triangles must be classified alive. The stones marked by squares must be classified dead. The unmarked stones are irrelevant for scoring because they are not at the border of their area and their possible capturability does not affect the score. For example, the two black stones in the top-left corner kill the white block and are in Black's area. However, they can always be captured by White, so forcing them to be classified as alive or dead is misleading and even unnecessary. (The stones in the bottom left corner are alive in *seki* because neither side can capture. The two white stones in the upper right corner are adjacent to two neutral points and therefore also at the border of White's region.)




Figure 1. Blocks to classify.

3.1 Recursion

Usually blocks of stones are not alive on their own. Instead they form chains or groups which are only alive in combination with other blocks. Their status also may depend on the status of neighbouring blocks of the opponent, i.e., blocks can live by capturing the opponent. (Although one might be tempted to conclude that life and death should be dealt with at the level of groups this does not really help because the human notion of a group is not well defined, difficult to program, and may even require an underlying notion of life and death.)

Because life and death of blocks is strongly related to the life and death of other blocks the status of other (usually nearby) blocks has to be taken into account. Partially this can be done by including features for nearby blocks in the representation. In addition, it seems natural to consider a recursive framework for classification which employs the predictions for other blocks to improve performance iteratively. In our implementation this is done by training a cascade of classifiers which use previous predictions for other blocks as additional input features.

4. Representation

In this section we will present the representation of blocks for classification. Several representations are possible and used in the field. The most primitive representations typically employ the raw board directly. A straightforward implementation is to concatenate three bitboards into a feature vector, for which the first bitboard contains the block to be classified, the second bitboard contains other friendly blocks and the third bitboard contains the enemy blocks. Although this representation is complete, in the sense that all relevant information is preserved it is unlikely to be efficient because of the high dimensionality and lack of topological structure.

4.1 Features for Block Classification

A more efficient representation employs a set of features based on simple measurable geometric properties, some elementary Go knowledge and some hand-crafted specialised features. Several of these features are typically used in Go programs to evaluate positions (Chen and Chen, 1999; Fotland, 2002). The features are calculated for single friendly and opponent blocks, multiple blocks in chains, and colour-enclosed regions (CERs).

For each block our representation consists of the following features: (All features are single scalar values unless stated otherwise.)

- *Size* measured in occupied points.
- *Perimeter* measured in number of adjacent points, including points over the edge.
- *Opponents* are the occupied adjacent points.
- (*First order*) *liberties* are the free adjacent points.
- *Protected liberties* are the liberties which cannot be played by the opponent, because of suicide or being directly capturable.
- *Auto-atari liberties* are liberties which by playing them reduce the liberties of the block from 2 to 1, which means that the blocks would become directly capturable (such liberties are protected for an adjacent opponent block).

- *Second-order liberties* are the liberties of (first-order) liberties (excluding the first-order liberties).
- *Third-order liberties* are the liberties of second-order liberties (excluding first- and second-order liberties).
- *Adjacent opponent blocks*
- *Local majority* is the number of opponent stones minus the number of friendly stones within a Manhattan distance of 2 from the block.
- *Centre of mass* represented by the distance to the closest and second-closest edge.
- *Bounding box size* is the number of points in the smallest rectangular box that can contain the block.

Adjacent to each block are colour-enclosed regions. CERs consist of connected empty and occupied points, surrounded by stones of one colour or the edge. It is important to know whether an adjacent CER is fully accessible, because a fully accessible CER surrounded by safe blocks provides at least one sure liberty. To detect fully accessible regions we use so-called miai strategies as applied by Müller (1997). In contrast to Müller's original implementation we also add miai accessible interior empty points to the set of accessible liberties, and also use protected liberties for the chaining. For fully accessible CERs we include:

- *Number of regions*
- *Size*
- *Perimeter*
- *Split points* are crucial points for preserving connectedness in the local 3×3 window around the point. (The region could still be connected by a big loop outside the local 3×3 window.)

For partially accessible CERs we include:

- *Number of partially accessible regions*
- *Accessible size*
- *Accessible perimeter*
- *Size of the unaccessible interior.*
- *Perimeter of the unaccessible interior.*
- *Split points of the unaccessible interior.*

The size, perimeter and number of split points are summed for all regions. We do not address individual regions because the representation must have a fixed number of features, whereas the number of regions is not fixed.

Another way to analyse CERs is to look for possible eyespace. Points forming the eyespace should be empty or contain capturable opponent stones. Empty points directly adjacent to opponent stones are not part of the eyespace. Points on the edge with one or more diagonally adjacent alive opponent stones and

points with two or more diagonally adjacent alive opponent stones are false eyes. False eyes are not part of the eyespace (we ignore the unlikely case where a big loop upgrades false eyes to true eyes). Initially we assume all diagonally adjacent opponent stones to be alive. However, in the recursive framework (see below) the eyespace is updated based on the status of the diagonally adjacent opponent stones after each iteration. For directly adjacent eyespace of the block we include:

- *Size*
- *Perimeter*

Since we are dealing with final positions it is often possible to use the optimistic assumption that all blocks with shared liberties can form a chain (during the game this assumption is dangerous because the chain may be split). For this, so-called, optimistic chain we include:

- *Number of blocks*
- *Size*
- *Perimeter*
- *Split points*
- *Adjacent CERs*
- *Adjacent CERs with eyespace*
- *Adjacent CERs, fully accessible* from at least one block.
- *Size of adjacent eyespace*
- *Perimeter of adjacent eyespace*
- *External opponent liberties* are liberties of adjacent opponent blocks which are not accessible from the optimistic chain.

Adjacent to the block in question there may be opponent blocks. For the weakest (measured by the number of liberties) directly adjacent opponent block we include:

- *Perimeter*
- *Liberties*
- *Shared liberties*
- *Split points*
- *Perimeter of adjacent eyespace*

The same features are also included for the second-weakest directly adjacent opponent block and the weakest opponent block directly adjacent to or sharing liberties with the optimistic chain of the block in question.

By comparing a flood fill starting from Black with a flood fill starting from White we find unsettled empty regions which are disputed territory (assuming all blocks are alive). If the block is adjacent to disputed territory we include:

- *Direct liberties* in disputed territory.
- *Liberties of all friendly blocks* in disputed territory.
- *Liberties of all enemy blocks* in disputed territory.

4.2 Additional Features for Recursive Classification

For the recursive classification the following six additional features are used:

- *Predicted value* of the strongest friendly block with a shared liberty.
- *Predicted value* of the weakest adjacent opponent block.
- *Predicted value* of the second-weakest adjacent opponent block.
- *Average predicted value* of the weakest opponent block’s optimistic chain.
- *Adjacent eyespace size* of the weakest opponent block’s optimistic chain.
- *Adjacent eyespace perimeter* of the weakest opponent block’s optimistic chain.

Next to these additional features the predictions are also used to update the eyespace, i.e., dead blocks can become eyespace for the side that captures, alive blocks cannot provide eyespace, and diagonally adjacent dead opponent stones are not counted for detecting false eyes.

5. The Data Set

In the experiments we used game records obtained from the NNGS archive (NNGS, 2002). All games were played on the 9×9 board between 1995 and 2002. We only considered games which are played to the end and scored, thus ignoring unfinished or resigned games. Since the game records only contain a single numeric value for the score, we had to find a way to label all blocks.

5.1 Scoring the Data Set

For scoring the dataset we initially used a combination of GNUGO and manual labelling. Although GNUGO has the option to finish games and label blocks the program could not be used without human supervision. The reasons for this are bugs, the inherent complexity of the task, and the mistakes made by weak human players which ended the game in positions that are not final, or scored them incorrectly. Fortunately, nearly all mistakes are easily detected by comparing GNUGO’s scores and labelled boards with the numeric scores stored in the game records. As an extra check all boards containing open regions with unsettled interior points (where flood filling does not give the same result as distance-based scoring) were also inspected manually.

Since the scores did not match in many positions the labelling proved to be very time consuming. We therefore only used GNUGO to label the games played in 2002 and 1995. With the 2002 games a classifier was trained. When we tested the performance on the 1995 games it outperformed GNUGO’s labelling. So therefore our classifier replaced GNUGO for labelling all other games (1996-2001), retraining it each time a new year was labelled. Although this speeded up the process it still required a fair amount of human intervention mainly because of games that contained incorrect scores in their game record.

A few hundred games had to be thrown out completely because they were not finished, contained illegal moves, contained no moves at all (for at least one side), or both sides were played by the same player. In a small number of cases, where the last moves would have been trivial but not actually played, we made the last few moves manually.

Eventually we ended up with a dataset containing 18,222 final positions. Around 10% of these games were scored incorrectly (by the players) and were inspected manually. (Actually the number of games we inspected is significantly higher because of the games that were thrown out and because our initial classifiers and GNUGO made mistakes). On average the final positions contained 5.8 alive blocks, 1.9 dead blocks, and 2.7 irrelevant blocks. (In the case that one player gets the full board all his blocks were assumed irrelevant although at least one block should of course be classified as alive.)

Since the Go scores on the 9×9 board range from -81 to $+81$ the chances of an incorrect labelling leading to a correct score are low, nevertheless it could not be ruled out completely. On inspecting an additional 1% of the positions randomly we found none that were labelled incorrectly. Finally, when all games were labelled, we re-inspected all positions for which our best classifier seemed to predict an incorrect score. This final pass detected 42 positions (0.2%) which were labelled incorrectly, mostly because our initial classifiers had made the same mistakes as the players that scored the games.

5.2 Statistics

Since many game records contained incorrect scores we looked for reasons and gathered statistics. The first thing that came to mind is that weak players might not know how to score. Therefore in Figure 2 the percentage of incorrectly scored games related to the strength of the players is shown. (Although in each game only one side may have been responsible for the incorrect score, we always assigned blame to both sides.) The two marker types distinguish between rated and unrated players. Although unrated players have a value for their rating, it is an indication given by the player and not by the server. Only after playing sufficiently many games the server assigns players a rating.

Although a significant number of games are scored incorrectly this is usually not considered a problem when the winner is correct. (Players typically forget to remove some stones when they are far ahead.) Figure 3 shows how often incorrect scoring by rated players converts a win to a loss.

It should be noted that the percentages in Figures 2 and 3 were weighted over all games regardless of the player. Therefore they do not necessarily reflect the probabilities for individual players, i.e., the statistics can be dominated by a small group of players that played many games. This group at least contains some computer players which have a tendency to get robbed of their points




Figure 2. Incorrect scores.

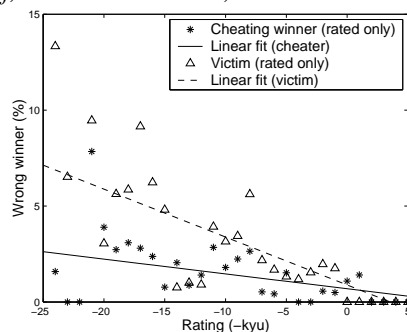


Figure 3. Incorrect winners.

in the scoring phase. We therefore also calculated some statistics that were normalised over individual players. For rated players the average probability of scoring a game incorrectly is 4.2%, the probability of cheating (the incorrect score converts loss to win) is 0.66 %, and the probability of getting cheated is 0.55%. For unrated players the average probability of scoring a game incorrectly is 11.2%, the probability of cheating is 2.1 %, and the probability of getting cheated is 1.1%. The fact that the probability of getting cheated is lower than the probability of cheating is the result of a small group of players (several of which are computer programs) that systematically lose points in the scoring phase, and a larger group of players that take advantage of them.

6. Experiments

In this section experimental results are presented for: (1) selecting a classifier, (2) performance of the representation, (3) recursive performance, (4) full board performance, and (5) performance on the 19×19 board. Unless stated otherwise the various training and validation sets, used in the experiments, were extracted from games played between 1996 and 2002. The test set was always the same, containing 7149 labelled blocks extracted from 919 games played in 1995.

6.1 Selecting a Classifier

An important choice is selecting a good classifier. In pattern recognition there is a wide range of classifiers to choose from (Jain et al., 2000). We tested a number of well-know classifiers for their performance on datasets of 100, 1000, and 10000 examples. The classifiers are: Nearest Mean Classifier (NMC), Linear Discriminant Classifier (LDC), Logistic Linear Classifier (LOGLC), Quadratic Discriminant Classifier (QDC), Nearest Neighbour Classifier (NNC), K-Nearest Neighbours Classifier (KNNC), BackPropagation Neural net Classifier with momentum and adaptive learning (BPNC), Levenberg-Marquardt Neural net Classifier (LMNC), and RProp Neural net Classifier (RPNC). Some preliminary experiments with a Support Vector Classifier, Decision Tree Clas-

sifiers, a Parzen classifier and a Radial Basis Neural net Classifier were not pursued further because of excessive training times and/or poor performance. All classifiers except the neural net classifiers, for which we directly used the standard matlab toolbox, were used as implemented in PRTools3 (Duin, 2000).

The results, shown in Table 1, indicate that performance first of all depends on the size of the training set. The linear classifiers perform better than the quadratic classifier and nearest neighbour classifiers. For large datasets training KNNC is very slow because it takes a long time to find an optimal value of the parameter k . The number of classifications per second of (K)NNC is also low because of the large number of distances that must be computed (all training examples are stored). Although the performance of the nearest neighbour classifiers might be improved by editing and condensing the dataset, we did not investigate them further.

Classifier	Training size	Training error (%)	Test error (%)	Training time (s)	Classif. speed (s^{-1})
NMC	100	2.8	3.9	0.0	4.9×10^4
	1,000	4.0	3.8	0.1	5.2×10^4
	10,000	3.8	3.6	0.5	5.3×10^4
LDC	100	0.7	3.0	0.0	5.1×10^4
	1,000	2.1	2.0	0.1	5.2×10^4
	10,000	2.2	1.9	0.9	5.3×10^4
LOGLC	100	0.0	9.3	0.2	5.2×10^4
	1,000	0.0	2.6	1.1	5.2×10^4
	10,000	1.0	1.2	5.6	5.1×10^4
QDC	100	0.0	13.7	0.1	3.1×10^4
	1,000	1.0	2.1	0.1	3.2×10^4
	10,000	1.9	2.1	1.1	3.2×10^4
NNC	100	0.0	18.8	0.0	4.7×10^3
	1,000	0.0	13.5	4.1	2.4×10^2
	10,000	0.0	10.2	4.1×10^3	2.4×10^0
KNNC	100	7.2	13.1	0.0	4.8×10^3
	1,000	4.2	4.4	1.0×10^1	2.4×10^2
	10,000	2.8	2.8	9.4×10^3	2.6×10^0
BPNC	100	0.5	3.6	2.9	1.8×10^4
	1,000	0.2	1.5	1.9×10^1	1.8×10^4
	10,000	0.5	1.0	1.9×10^2	1.9×10^4
LMNC	100	2.2	7.6	2.6×10^1	1.8×10^4
	1,000	0.7	2.8	3.2×10^2	1.8×10^4
	10,000	0.5	1.2	2.4×10^3	1.9×10^4
RPNC	100	1.5	4.1	1.4	1.8×10^4
	1,000	0.2	1.7	7.1	1.8×10^4
	10,000	0.4	1.1	7.1×10^1	1.9×10^4

Table 1. Performance of classifiers.

The best classifiers are the neural network classifiers. It should however be noted that their performance may be slightly overestimated with respect to the size of the training set, because we used an additional validation set to stop training (this was not possible for the other classifiers because they are not trained incrementally). The Logistic Linear Classifier performs nearly as good as the neural network classifiers, which is quite an achievement considering that it is just a linear classifier.

The results of Table 1 were obtained with neural networks that employed one hidden layer containing 15 neurons with hyperbolic tangent sigmoid transfer functions. Since our choice for 15 neurons was quite arbitrary a second experiment was performed in which we varied the number of neurons in the hidden layer. In Figure 4 results are shown for the RPNC. The classification errors marked with triangles represent results for training on 5,000 examples, the stars indicate results for training on 15,000 examples. The solid lines are measured on the independent test set, whereas the dash-dotted lines are obtained on the training set. The results show that even moderately sized networks easily overfit the data. Although the performance initially improves with the size of the network, it seems to level off for networks with over 50 hidden neurons (the standard deviation is around 0.1 %). Again clearly the key factor in improving performance is in increasing the training set.




Figure 4. Sizing the neural network for the RPNC.

6.2 Performance of the Representation

In section 4 we claimed that a raw board representation is inefficient for predicting life and death. To validate this claim we measured the performance of such a representation and compared it to our specialised representation.

The raw representation consists of three concatenated bitboards, for which the first bitboard contains the block to be classified, the second bitboard contains other friendly blocks and the third bitboard contains the enemy blocks. To remove symmetry the bitboards are rotated such that the centre of mass of the block to be classified is always in a single canonical region.

Since high-dimensional feature spaces tend to raise several problems which are not directly caused by the quality of the individual features we also tested two compressed representations. These compressed representations were generated by performing Principal Component Analysis (PCA) on the raw representation.

For the first PCA mapping the number of features was chosen identical to our specialised representation. For the second PCA mapping the number of features was set to preserve 90% of the total variance.

The results, shown in Table 2, are obtained for the RPNC with 15, 35, and 75 neurons in the hidden layer, for training sets with 100, 1,000 and 10,000 examples. All values are averages over 11 runs with different training sets, validation sets (same size as the training set), and random initialisations. The errors, measured on the test set, indicate that a raw representation alone requires too many training examples to be useful in practice. Even with 10,000 training examples the raw representation performs much weaker than our specialised representation with only 100 training examples. Simple feature-extraction methods such as Principal Component Analysis do not seem to improve performance, indicating that preserved variance of the raw representation is relatively insignificant for determining life and death.

Training Size	Extractor	Test error 15 neurons (%)	Test error 35 neurons (%)	Test error 75 neurons (%)
100	-	29.1	26.0	27.3
100	pca1	22.9	22.9	22.3
100	pca2	23.3	24.3	21.9
1000	-	13.7	13.5	13.4
1000	pca1	16.7	16.2	15.6
1000	pca2	14.2	14.5	14.4
10000	-	7.5	6.8	6.5
10000	pca1	9.9	9.3	9.1
10000	pca2	8.9	8.2	7.7

Table 2. Performance of the raw representation.

6.3 Recursive Performance

Our recursive framework for classification is implemented as a cascade of classifiers which use extra features, based on previous predictions as discussed in subsection 4.2, as additional input. The performance measured on an independent test set for the first 4 steps is shown for various sizes of the training set in Table 3. The results are averages of 5 runs with randomly initialised networks containing 50 neurons in the hidden layer (the standard deviation is around 0.1 %).

The results show that recursive predictions improve the performance. However, the only significant improvement comes from the first iteration. The improvements are by far not significant for the average 3- and 4-step errors. The reason for this is that sometimes the performance got stuck or even worsened after the first iteration. Preliminary experiments suggest that large networks

were more likely to get stuck after the first iteration than small networks, which might indicate some kind of overfitting. A possible solution to overcome this problem is to retrain the networks a number of times, and pick the best based on the performance on the validation set. If we do this the best networks, trained on 100,000 training examples, achieve a 4-step error of 0.25%.

Training Size	Direct error (%)	2-step error (%)	3-step error (%)	4-step error (%)
1,000	1.93	1.60	1.52	1.48
10,000	1.09	0.76	0.74	0.72
100,000	0.68	0.43	0.38	0.37

Table 3. Recursive performance.

6.4 Full Board Performance

So far we have concentrated on the percentage of blocks that are classified correctly. Although this is an important measure it does not directly tell us how often boards will be scored correctly (a board may contain multiple incorrectly classified blocks). Further we do not yet know what the effect is on the score in number of board points. Therefore we tested our classifier on the full-board test positions (which were not used for training or validation).

For our best 4-step classifier trained with 100,000 examples we found that 1.1% of the boards were scored incorrectly. For 0.5% of the boards the winner was not identified correctly. The average number of incorrectly scored board points (using distance-based scoring) was 0.15, however in case a board is scored incorrectly this usually affects around 14 board points (which counts double in the numeric score).

6.5 Performance on the 19×19 Board

The experiments presented above were all performed on the 9×9 board which, as was pointed out before, is a most challenging environment. Nevertheless, it is interesting to test whether the techniques scale up to the 19×19 board. So far we did not have the time to label large quantities of 19×19 games. So, training directly on the 19×19 board was not an option. Despite of this we tested our classifiers, which were trained from blocks observed on the 9×9 board, on the problem set *IGS_31_counted* from the Computer Go Test Collection. This set contains 31 labelled 19×19 games played by amateur dan players, and was used by Müller (1997). On the 31 final positions our 4-step classifier classified 5 blocks incorrectly (0.5% of all relevant blocks), and as a consequence 2 final positions were scored incorrectly. The average number of incorrectly scored board points was 2.1 (0.6%).

In his paper Müller stated that heuristic classification of his program EXPLORER classified most blocks correctly. Although we do not know the exact performance of EXPLORER we believe it is safe to say that our system, which scored 99.4% of all board points correctly, is performing at least at a comparable level. Furthermore, since our system was not even trained explicitly for 19×19 games there may still be significant room for improvement.

7. Conclusions

We have developed a system that learns to score final positions from labelled examples. On unseen game records our system scored around 98.9% of the positions correctly without any human intervention. Compared to the average rated player on NNGS (who for scored 9×9 games has a rating of 7 kyu) our system is more accurate at removing all dead blocks, and performs comparable on determining the correct winner.

By comparing numeric scores and counting unsettled interior points we can efficiently detect nearly all incorrectly scored final positions. Although some final positions are assessed incorrectly by our classifier, most are in fact scored incorrectly by the players. Detecting these games is important because most machine-learning methods require reliable training data for good performance.

7.1 Future Work

By providing reliable score information our system opens the large source of Go knowledge which is implicitly available in human game records. The next step will be to apply machine learning in non-final positions. We believe that the representation and techniques presented in this paper provide a solid basis for static predictions in non-final positions.

The good performance of our system was obtained without any search, indicating that static evaluation is sufficient for most human final positions. Nevertheless, we believe that some (selective) search can still improve the performance. Adding selective features that involve search and integrating our system into MAGOG, our 9×9 Go program, will be an important next step.

Although the performance of our system is already quite good for labelling game records, there are, at least in theory, still positions which may be scored incorrectly when our classifier makes the same mistake as the human players. Future work should determine how often this happens in practice.

Another point where our system can be improved is the representation. Although the current representation performs adequately, some features may be redundant or correlated. Feature extraction, feature selection, and possibly adding some new features may improve performance even further.

References

- Bouzy, B. and Cazenave, T. (2001). Computer Go: An AI oriented survey. *Artificial Intelligence*, 132(1):39–102.
- Chen, K. and Chen, Z. (1999). Static analysis of life and death in the game of Go. *Information Sciences*, 121:113–134.
- Dahl, F. (2001). Honte, a go-playing program using neural nets. In Fürnkranz, J. and Kubat, M., editors, *Machines that Learn to Play Games*, chapter 10, pages 205–223. Nova Science Publishers, Huntington, NY.
- Duin, R. (2000). *PRTools, a Matlab Toolbox for Pattern Recognition*.
- Enderton, H. (1991). The Golem Go program. Technical Report CMU-CS-92-101, School of Computer Science, Carnegie-Mellon University.
- Enzenberger, M. (1996). The integration of a priori knowledge into a Go playing neural network. Unpublished Manuscript.
- Fotland, D. (2002). Static eye analysis in ‘THE MANY FACES OF GO’. *ICGA Journal*, 25(4):201–210.
- Jain, A. K., Duin, R. P. W., and Mao, J. (2000). Statistical pattern recognition: A review. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(1):4–37.
- Müller, M. (1997). Playing it safe: Recognizing secure territories in computer Go by using static rules and search. In Matsubara, H., editor, *Proceedings of the Game Programming Workshop in Japan '97*, pages 80–86. Computer Shogi Association, Tokyo, Japan.
- Müller, M. (2002). Computer Go. *Artificial Intelligence*, 134(1-2):145–179.
- Myers, B. (2002). The 21st century championship cup 2002. *ICGA Journal*, 25(4):245.
- NNGS (2002). The no name go server game archive.
<http://nngs.cosmic.org/gamesearch.html>.
- Schraudolph, N., Dayan, P., and Sejnowski, T. (1994). Temporal difference learning of position evaluation in the game of Go. In Cowan, J. D., Tesauro, G., and Alspector, J., editors, *Advances in Neural Information Processing 6*, pages 817–824. Morgan Kaufmann, San Francisco, CA.
- Tesauro, G. (1995). Temporal difference learning and TD-Gammon. *Communications of the ACM*, 38(3):58–68.
- van der Werf, E. C. D., Uiterwijk, J. W. H. M., Postma, E. O., and van den Herik, H. J. (2002). Local move prediction in Go. In *Proceedings of the 3rd International Conference on Computers and Games (CG'02) (Edmonton, CA, July 25-27)*. To appear in LNCS, © Springer-Verlag.