16

Expert Against Oracle

A. J. Roycroft*
The Turing Institute,
Glasgow, UK

Abstract

A computer-generated combinatorial data base that plays optimally an almost undocumented and very difficult five-man chess endgame (i.e. the data base can be considered as an oracle) was matched against a domain specialist who had prepared for the contest with minimal prior access to the data base. His preparation and strategy are described and the results of the contest itself briefly summarized. The paper closes with an illustrated discussion of the selected endgame in Master practice.

1. PURPOSE OF PROJECT

The chess endgame specialist, in contrast to the tournament player, assesses his skill against the criteria of perfection, whether in adjudicating a position or in selecting a move. Automatic construction of complete look-up tables (data bases) by computer makes it possible to apply these exacting criteria in practice. By use of such an 'oracle' Kopec and Niblett (1980) were able empirically to verify the present author's claim to have acquired high-level mastery (in the actual test the play was move-perfect) of the play of won-for-White positions of king and rook against king and knight, an endgame of which full knowledge was lacking prior to the creation of a data base. The purpose of the new project is to investigate the problems and nature of skill-acquisition in an endgame selected as being so complex as to lie beyond the power of the unaided human endgame specialist to master thoroughly. For this purpose the author selected the ending king–bishop–bishop–king–knight (BBN). BBN was historically assumed to be a draw from a general position until at my suggestion Mr Ken Thompson computed an exhaustive BBN data base. The existence of the data base, and samples of its output, were published in the international quarterly magazine *EG* (Thompson and Roycroft, 1983).

* Present address: 17 New Way Road, London NW9 6PL, UK.

The Thompson data base contains complete information on how the two-bishops side can win from all but a few elementary or bizarre positions in a total space of some 250,000,000 positions. It thus exhibits 'skill' at a demonstrably unsurpassable level in a domain that is of formidable difficulty and complexity for human experts.

This characteristic makes possible in principle the absolute measurement of human performance in a difficult domain, and offers the opportunity to explore how 'inert but absolute' knowledge can be accessed and adapted for teaching, for learning, and for the development and validation of expert systems that aim to perform as well as the data base—but without it.

We also for the first time have the possibility to explore thoroughly a significant endgame data base at 'super-expert' level. Future combinatorial data bases in chess and in other domains may be more massive still. Early experience of handling them will be of value.

The present paper reports measurement of the performance of a domain specialist before (Part 1) and after (Part 2) he had been allowed unlimited access to the oracle data base, from which, for any legal position the user can retrieve the following:

(i) all optimal single-ply (that is, at a depth of one white or black move) continuations; and

(ii) the length of the remaining optimal path.

The author is a lifetime student of the chess endgame. He is the author of *The Chess Endgame Study* (1972 and 1981) and edits and publishes the international magazine *EG*. He is a strong, but not master, chessplayer. He is acknowledged worldwide as an endgame specialist.

2. THE TASK OF THE DOMAIN SPECIALIST (IN PART 1)

In October 1984 the author undertook:

1. To study the five-man pawnless chess endgame of (white) king and two bishops (one on light squares, one on dark squares) against (black) king and knight, using any available aid except the data base itself, with the exception of a set of 12 variations already excavated and provided in 1983 by Mr Thompson from the data base illustrating optimal play from one of the 32 worst-case-for-White (two bishops) starting positions. Time limit imposed on this period: none.
2. To record as fully and faithfully as possible all thought processes (the dated record to include time taken, chess positions and moves, sources of information used, discoveries, errors, corrections, trains of thought, going over previously trodden ground, etc.).
3. To announce when ready to face the data base, i.e. when no greater mastery of the material seemed achievable by private study.
4. To take the white (two bishops) side against the data base without

prior preparation of the particular positions to played and to play under strict tournament conditions (moves to be timed, and no moves taken back), with two exceptions: analysis on a separate board to be allowed (the domain specialist is not a practising chess master used to tournament conditions but an endgame scholar accustomed to analysing with board and men, similar to the situation that obtains in correspondence chess); and the so-called '50-move rule' to be disregarded.

The confrontation of domain specialist and oracle data base concludes Part 1. If, as is expected, the human performance is sub-optimal, Part 2 follows, in which the specialist is now allowed access to the data base. A second confrontation or test, with different positions but the same number of them, concludes Part 2.

3. BACKGROUND

The earliest known reference in chess literature to the pawnless endgame king and two bishops against king and knight is in the middle of the nineteenth century (Kling and Horwitz, 1851, pp. 62–5). R1, the first of two positions given by the authors, is the more important as it is largely independent of the positioning of the white force. It is given by them without supporting analysis but with the statement that the bishops 'cannot win if the weaker side can obtain a position similar to the above, but they win in most cases'. The second position, a win, is then given with a solution and a number of supporting variations extending to 14 moves. One or other of both positions is repeated in the subsequent literature up to 1983 (e.g. Pachman, 1983, pp. 19–20), with no modification to the verdict.

The author (Roycroft, 1972, p. 207) raised a doubt about the correctness of the claim that R1 (and positions like it) cannot be won. This doubt was confirmed in 1983 by output from the Ken Thompson data base.

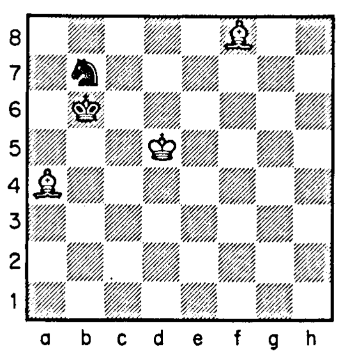

R1. Kling and Horwitz (1851). Either side to move.

The data base was generated by a method already known in principle (Ströhlein, 1970). First, all possible positions of checkmate (with the given force), and all positions where the knight is safely captured (without subsequent stalemate), are automatically generated. These comprise the finally won positions that the side with the bishops aims for, and at the same time they are the positions that the side with the knight wishes to postpone as long as possible. From this starting 'position set' the first 'derived set' of positions can be generated, the set of White to Move (WTM) positions that are 'Won in 1'. By an essentially similar, but logically more complex, process the antecedent Black to Move (BTM) position sets are generated and marked where and when appropriate 'Lost in 1'. The basis of an iterative 'maximin' or 'backing-up' procedure has now been established, whereby the solved depth increases in principle by one ply (one white move or one black) per pass. This iteration is initiated and relentlessly pursued until no more positions can be classified. At this stage all won positions will be marked with the solution depth. For a more detailed description of the process see, for example, Roycroft and Niblett (1979) and Thompson (1986). Residual positions still unmarked will be drawn, illegal, or, in a microscopic number of BTM instances, won for the knight's side. In Ken Thompson's solution only WTM positions are stored, the BTM positions being generated when required by program: for convenience we refer simply to the 'data base' whether WTM positions only or both WTM and BTM positions are physically stored.

The results have been widely reprinted in the world's other chess magazines. However, guidance in the domain literature as to how this endgame should be played remains (August, 1985) restricted to paraphrases of the sentence of Kling and Horwitz quoted above, that is that the defending side should always aim for a position like R1, because it is the only safe draw. (At the end of this paper we give an example of the influence of this advice on practical master play.) As a result of the present research it is likely that future advice to the superior side will be to steer towards the Kling and Horwitz position, since the winning method from that position is (or rather will be) well charted. (For a list of the principal authorities on the chess endgame see the entries within parentheses in the section References. However, Averbakh, the major modern authority, does not include the two bishops against knight endgame because endings with two pieces on one side are in principle excluded from its scope.)

4. THE FIVE PHASES OF THE PAWNLESS ENDGAME TWO BISHOPS AGAINST KNIGHT

The division of a maximum length solution to this endgame into five phases has been described by the author (Thompson and Roycroft, 1983). The following is an updated version. The quoted passages are taken from the article in *EG*. Where a number of moves is mentioned this refers to consecutive optimal moves by White, the side with the bishops.

4.1. Phase number

1. In a maximum depth solution position the white force will initially be under some constraint from the black force: either the white king or a bishop will be immobilized. It may take from six to 12 moves to lift this blockade, depending on its nature.

2. In the next phase Black retreats slowly and in good order and 'seeks refuge in the Kling and Horwitz position. This may be in any corner'. This phase takes us up to move 20, approximately.

3. White's task in phase 3 is to manoeuvre in order to set up any of a small number (probably only four) of 'exit' positions, that is, exits from a Kling and Horwitz position. Black is then forced out into the open. Typically this phase lasts six or seven moves.

4. 'The next stage is complex, fluid, lengthy and difficult. Black strives for maximum freedom, and frequently seems on the verge of achieving it. It takes White some 23 moves, not to be found in any book and characterized at times by excruciating slowness and mystery, before' Black, 'having failed time and again to repeat the Kling and Horwitz position, ends up' with his king 'on the board's edge, near a corner and accompanied by the black knight.'

5. 'The remaining dozen or so moves show the knight being lost, whether he stays close to the black king or runs away.'

The longest solutions have 66 white moves. There are 32 distinct positions that have this depth, though they group into 'families' of positions.

5. THE STORY OF PART 1

5.1. The 'private study' phase

The private study phase began in October 1984. The material available for study comprised:

(i) published books (in English, German, and Russian) on the chess endgame in general. In contrast to the thriving literature on individual chess openings there is very little published on specific endgames. The books do not cover the endgame two bishops against knight in any useful depth. (See References);

(ii) ten full-length (66 moves) solutions and associated list of one-ply-deep equi-optimal moves provided in August 1983 by Ken Thompson to the author in the latter's capacity as editor of *EG* magazine;

(iii) two further full-length solutions, also from Ken Thompson in 1983, with no alternative moves;

(iv) three 66-move and 67-move full-length solutions from an independent researcher. The 67-move solutions were later shown to be faulty (Comay and Roycroft, 1984), and the correctness of the 66-move maximum optimal depth thereby corroborated;

(v) the 32 distinct positions at the maximum optimal solution length, also provided by Ken Thompson to the author in 1983, but without chess moves;

(vi) the frequency table of the numbers of WTM positions at every solution length from 66 to 1, also provided by Ken Thompson to the author.

On 18th January, 1985 the domain specialist intimated in writing his readiness to confront the data base in the test to end Part 1.

5.2. The protocol

Separate publication of the protocol record of the domain specialist's thought processes is intended. It runs to over 200 pages and will be supplemented with appendices.

5.3. The Part 1 test and summarized results

The test began on Friday, 29 March 1985. Two test sessions were aborted due to program failure, and there was a two weeks' interruption for holiday. The 10th and final test position was played on Tuesday, 30 April 1985.

Two sessions were abandoned by the domain specialist, after 70 and 69 moves respectively. The remaining eight positions of the test were won by the domain specialist, giving a 'tournament' performance of 80%. The only other measurement of his performance that is available at present to the author is the ratio of the total of the optimal solution depths of the original positions to the number of moves actually taken by the author: 38%.

Both measurements are crude. If the sessions abandoned by the specialist after 70 and 69 moves had been abandoned at the outset without any winning attempt at all, the 38% figure would increase, thereby putting a premium on early abandonment. This is, however, not the case with the refinement (Doran–Michie 'path efficiency') used by Michie (1986) in his review of these same experiments, which rates abandonment at any stage as equivalent to taking infinitely many moves. It is not clear what measurement would be least unsatisfactory. Two other measurements will almost certainly give different figures and should at least be calculated:

1. The ratio of optimal moves to sub-optimal moves in all the moves played by the domain specialist.

2. A measurement that takes into account the domain specialist's division of the endgame into five phases. This would log a minimal penalty against a move that is sub-optimal but which kept the solution within the same phase; it would log a heavier penalty against a move that set the solution back a phase; an even heavier penalty would be imposed on an error that set back the solution two phases, and so on; the heaviest penalty would be for a blunder that gave away the win. If a penalty were measured in numbers of question marks (i.e. '?'), with the lowest penalty rated at a single question mark, then a session could be aborted by prior agreement if the total of accumulated penalties (i.e. the total of question marks deserved) in a session passed a certain threshold. This content-related measurement of performance was proposed by the author but not adopted, one argument against being that the division of solution into phases is at present subjective.

The reason not all measurements are available to the author is that work is proceeding and it is considered that even such ancillary information relating to an earlier test could be of indirect assistance to the domain specialist, who, at the time of writing has not had the test to conclude Part 2. Since a major object of the combined Parts 1 and 2 is to determine how and to what extent a human specialist can be aided in his comprehension of a complex domain, such additional information might, however slightly, distort the performance and measurement. Full statistics will be reported in the planned monograph on the total experiment covering Parts 1 and 2 (Further experiments with the present oracle, and experiments with other, even more massive data bases are envisaged.)

6. THE DOMAIN SPECIALIST'S STRATEGIES

Implemented strategies are necessarily domain specific: they have to be described in chess terms. But some general remarks for non-chess players may be helpful.

6.1. For non-chessplayers

Here is no place to debate what, if anything, chessplayers have and non-chessplayers lack. But being an amateur problem-solver as well as a chessplayer the author recognizes that all problem-solvers have some common skills and motivations, whatever their specialist knowledge or favoured domain. The awe in which non-chessplayers commonly hold chessplayers of even less than master strength is based partly, if not mainly, on myth. In the interests of better understanding the present section of this paper aims to demolish two specific myths.

The first myth: enormous numbers

A frequent argument to boost the myth of the arcane genius chessplayer invokes 'enormous numbers'. The number of possible chess positions exceeds the highest astronomical numbers; the number of possible chess games exceeds the number of possible chess positions, also astronomically. These are incontrovertible facts.

But a chessplayer does not have to remember or recognize all these positions and games, any more than any of us need remember or recognize all possible breakfasts in order to eat breakfast, or need remember or recognize all possible books in order to read a book.

Humans tame large numbers by ignoring them. Instead they seek and manipulate patterns, even if the patterns are initially tentative, approximate or unsound. It is a common-sense conjecture of everyday experience that a good pattern will have inferior patterns in its ancestry. If we persevere and are willing to learn, later patterns should be superior to earlier.

In the universal child's game of noughts and crosses (in American it is 'tic-tac-toe', in Russian 'Krestiki-noliki') how many different positions are there? Interestingly, the question is posed, and answers provided, in the literature of artificial intelligence (examples: Nilsson, 1971, pp. 137–8; Shirai and Tsuji, 1982, p. 10; Rich, 1983, p. 7; Alty and Coombs, 1984, p. 80). The usual answer given relies on the implied logic that there are nine initially empty cells to be filled, so we start with nine possibilities for the first play, leaving eight for the reply, seven for the third play, and so on. Factorial 9, or 9! is the (for a game like noughts and crosses, large) number: 362,880—which reduces through laws of legality and symmetry to 'three hundred or so distinct positions with which Nought (by convention the opening player) can be confronted' (Michie, 1961).

A contrasting answer results from arguing that there are three possible states for each cell: a nought, a cross, or emptiness. This gives us a ceiling of three to the ninth power (Rich, 1983), or: 19,683—and this is *before* eliminating symmetries.

Neither of these calculations impresses the *player* of noughts and crosses.

When tackling the identical question he will rather reason like this:

1. There are three rules or conventions that govern the game, and we can look on them as constraints:
—nought starts (constraint no. 1);
—play alternates between the placing of nought and cross (constraint no. 2);
—a completed row (in any direction, including diagonally) ends the game (constraint no. 3).

2. There is then a constraint of a different kind:
—elimination of all symmetries (constraint no. 4).

3. Finally, there is a constraint of a different kind again, one of experience or demonstration, namely:
—a well played game is inevitably drawn (constraint no. 5)

Constraints 1, 2, and 3 are part of the definition of the game. Constraint no. 5 amounts to the constraint of playing the game well. Constraint no. 4 is not essential but it is convenient to all parties.

The player then puts four questions based on the five constraints.

Q1: Can noughts occupy ALL FOUR corner cells?

The answer is 'no', because by constraint no. 5 there is no vacant cell for the fifth nought, needed by constraint no. 2.

Q2: Can noughts occupy just THREE corner cells?

The answer is 'yes', but in only one way.

Q3: Can noughts occupy just TWO corner cells?

The answer here has two parts: if the corner cells are diagonally opposite one another, then the remaining pair of (corner) cells must hold crosses, and any play thereafter into the centre cell will infringe constraint no. 5 again; on the other hand with a pair of cells in adjacent corners it is easy to show (by applying one or more of the five constraints) that there are just two possible distinct configurations.

Q4: Can there be a nought in one corner only or in no corner?

With a nought in just one corner cell, or in none, there is no way to satisfy all five constraints.

The player's answer to the question 'How many positions?' is therefore

1 + 2, i.e. 'three', (See R2.)

We have seen the following answers:

```
362,880
 19,683
      3
```

```
O X O        O X O   O X O
X X O        O X O   O O X
O O X        X O X   X O X
```

Noughts in three corners. Noughts in two corners.

R2. Noughts-and-crosses/tic-tac-toe/krestik i nulik. The only 'all constraints satisfied' configurations.

Which of them corresponds most closely to the reader's experience of the game?

Chess is more complex (i.e. it holds inordinately more patterns) than noughts and crosses, but humans cannot play chess well without forming, holding and manipulating patterns any more than they can play good noughts and crosses patternlessly. We shall return to this point after demolition of the second myth.

The second myth: how with all those mobile and differently moving pieces, can chessplayers possibly plan?

This argument, often in paraphrased form invoking large search trees and high branching factors, requires a different counter-demonstration.

Consider the puzzle of the 'solid pentominoes'. R3 shows the 12 'flat' pentomino shapes, namely all variations on five edge-contiguous identical squares. The solid pentominoes are made out of small cubes instead of squares, but the shapes are otherwise as shown. The puzzle we shall consider is the packing of all 12 shapes into a $3 \times 4 \times 5$ unit dimensioned box. (If preferred, the target shape can be thought of as a $3 \times 4 \times 5$ 'brick' to be assembled.) As there are 60 cubelets and 60 spaces to be filled there must be no empty space and no protruding cubelet.

Let us now describe from scratch how this quite tough puzzle might be tackled.

To begin with there is no strategy. We 'play' with the pieces, trying to solve randomly, but as we play we observe ourselves, looking for a pattern. Any discerned pattern is likely to be a pattern of failure, not a pattern of success, but this does not mean that it is not a pattern, nor

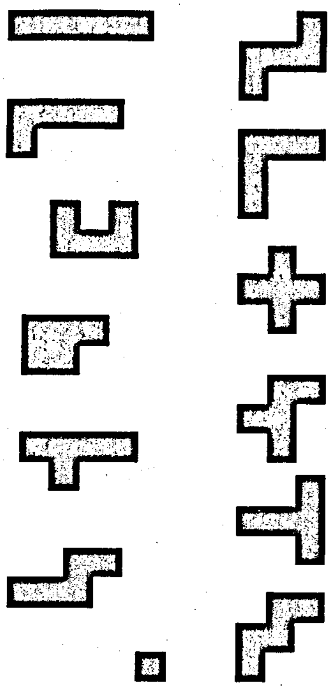

R3. the 12 different plane pentominoes that can be formed from juxtaposing five unit squares.

that it will not serve. The pattern which we observe, perhaps, is that after a failed attempt one or more of the shapes in the right-hand column of our figure tend to be unused, and that these are unused more frequently than the shapes in the left-hand column. (In passing, we can ask what the left-over 'awkward' pentominoes have in common with one another? Well, a 3×3 space (like noughts and crosses—a pattern!) is necessary and sufficient to hold any one of them.) Now from a pattern to a strategy is, for a human, no long journey, albeit not always a conscious or speedy one. In the case of our observed pattern, the derived plan or strategy might be:

> First fit 'several' of the 'awkward' pentominoes together into an *ad hoc* sub-assembly that will 'mutually absorb and minimize' the 'awkwardnesses' of individual pentominoes, and then arrange the remaining 'less awkward' pentominoes around the sub-assembly.

There are two important points about this strategy: it is not precise; and it is readily grasped by a human and implemented by a human, but a strategy it unquestionably is. It lends itself to objective evaluation in an experiment to compare the performance (time taken to solve, success ratios) of two groups of students, one group given the presumed benefit of the 'strategy' and the other group told nothing, but both groups tackling the identical constructional task. A complete set of the 3940 solutions to this puzzle has been computed and published (Bouwkamp, 1967).

That is typical of how the puzzle-solving mind works, whether in chess or pentominoes. Patterns are observed which lead to a strategy. A strategy will not find a solution by itself, but it serves the purpose of enabling the solver to be aware of what he is doing, to have a general aim which he can work with and refine. Such a general aim seems all the more manageable for being imprecise: 'several', 'awkward', 'mutually absorb and minimize', 'less awkward' are fuzzy terms and may be implemented differently (i.e. they may relate to different subsets of the pentominoes or to different subfeatures) by different people, though validly so: there is a distinction to be made between 'relevant fuzziness', which may even be essential in the initial communication of ideas or strategies, and ambiguity, which is to be avoided.

It is the same in chess. If you already have a strategy you can carry it out in your own way, adapting it as you go or discarding it for a better. If you don't, you can't.

If the foregoing account is accepted as valid it follows that for research purposes articulacy in the domain specialist is more important than expertise. A valuable corollary will be that a test of articulacy (to select good human subjects) can be general and standard, though designing such a test is not *a priori* the province of AI researchers.

6.2. For chessplayers

The problem is strange, even to an experienced chessplayer. Classic concepts, such as the importance of the centre of the board, mobility, sacrificial combinations and positional considerations, turn out to have either no, or limited, application. A long period was spent attempting to apply long-built-in concepts of the classic type, especially those that ought to be applicable to other endgames, but eventually they were largely replaced, or modified, as a result of hard experience.

The detailed story will be told in the 200-page protocol. Here only the principal new chess concepts that were found useful will be described. However, one extant concept that proved fundamental and fruitful was the Kling and Horwitz position, though even this tried and tested 133-year-old idea was inadequate in its basic formulation: elaboration was necessary.

6.3. New chess concepts

1. *'Knight's distance' from Kling and Horwitz position.* If we consider the square b7, then ONE knight's move's distance means the four specific squares a5, c5, d6, d8. (See R4.) TWO knight's moves' distance means the squares a4, a6, c4, c6, c8, e6, e8, plus the 'less frequent' squares b3, d3, e4, f5, f7 (R5 and R6). The latter five squares are less frequent in the sense that bN is less likely to occupy squares that are in the centre of the board (or in sub-regions controlled by the bishops) because it is the centre of the board that White must and can control in Phase 2 of the contest, the part that consists in driving Black out of the centre. The concept of knight's distance from a Kling and Horwitz position is particularly valuable in the most difficult part of the solution, namely Phase 4. It enables us with confidence (not with certainty) to estimate how far the solution has progressed and to consider particular special strategies and tactics appropriate to that stage, rejecting (i.e. not considering) irrelevant strategies and tactics. Now the four base squares

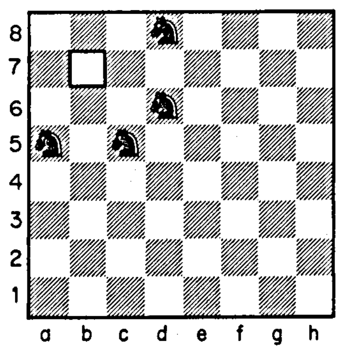

R4. Knight's distance ONE from b7.

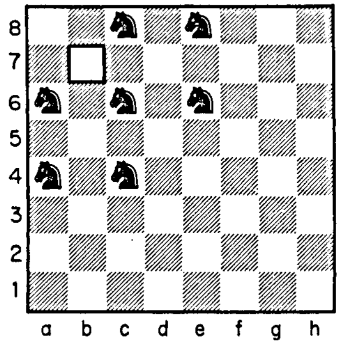

R5. Knight's distance TWO from b7 ('frequent' squares).

for Kling and Horwitz position mean different sets of squares in each case. But common squares begin to appear with greater frequency the greater the 'distance number'. Thus the square f4 is ONE from g2, TWO from b2 and g7, and THREE from b7—it is a 'good' square for Black's knight. (See R7.) In practice, since the black king moves slowly, and since the white force will in Phase 4 dominate the centre and some other sub-regions of the board, certain paths will be taboo to the knight: only those Kling and Horwitz positions which are accessible to the black king need be considered in applying the concept of knight's distance from a Kling and Horwitz position.

2. *Black 'king's distance' from a Kling and Horwitz position.* This concept is more static, that is, it is less liable to change from move to move, than the previous concept. The Kling and Horwitz position by definition requires Black's king; Black's king can move only to adjacent squares; it leaps to the eye when Black's king is occupying a Kling and Horwitz square; and distance simply means counting the king moves needed to reach such a square. On f4, for example, the black king is ONE from g3 (for a g2 position) and TWO from g6 (for a g7 position): he is

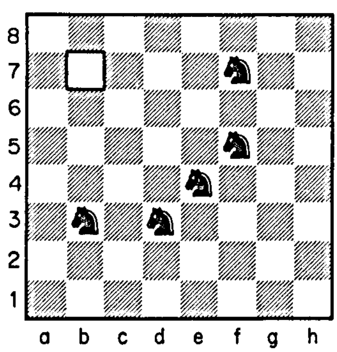

R6. Knight's distance TWO from b7 ('less frequent' squares).

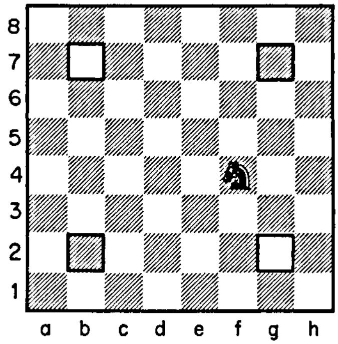

R7. Distance ONE from g2, distance TWO from g7 or b2, distance THREE from b7.

THREE from c2, but this can nearly always be ignored (in Phase 4). (See R8.)

3. Sum of 'distances'. This is simply the sum of the black knight's distance from a Kling and Horwitz position and the black king's distance from the same position. (The squares are, of course, different squares: a 'g7' Kling and Horwitz position implies the squares g7 for the black knight and g6 or f7 for the black king.) The summed distances are a rough-and-ready guide to progress in Phase 4. (See R9.)

4. A 'pseudo-fortress'. In chess endgame parlance the concept of a 'fortress' is familiar. The implication is that the materially inferior side sets up a position which, due to the geometry of the chessboard is impregnable to the particular attacking potential of the superior side. The edge of the board and especially the corners are suited to fortress positions. A 'pseudo-fortress' arises when Black is evicted from a Kling and Horwitz position (which has been called a fortress, but, we now know, in error) and adopts a posture in which king and knight are alongside each other between two Kling and Horwitz positions (such as

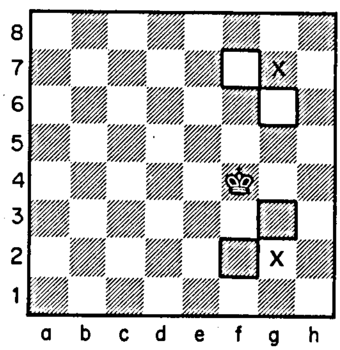

R8. King's distance ONE from g3 (a 'g2' square), King's distance TWO from g6 (a 'g7' square).

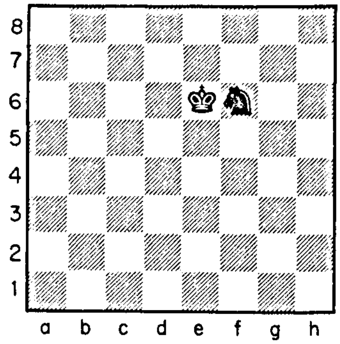

R9. Knight's distance TWO from g7, King's distance ONE from f7. Summed distance: THREE.

between b7 and g7) and at 'summed distances' of THREE or FOUR from each of them. Thus with bKd2 and bNe2 the 'summed distance' is FOUR from b2 and FOUR from g2. (See R10.) We may note, simply for the contrast, that the summed distance is EIGHT to b7 and EIGHT to g7. Such a position is strong for Black because he can adopt the strategy of oscillating to and fro with knight and king without heading for b2 or g2, which White presumably can prevent. Moreover, in his choice of an oscillating move Black, if he cannot check or usefully gain a tempo by attacking a bishop, will tend to choose a move that does not increase the summed distances. There are many manifestations of the pseudo-fortress.

5. The 'box'. This is the White 'counter-concept' useful in overcoming the 'pseudo-fortress'. It is a simple idea but its power is best explained by a comparison of chess diagrams and simultaneous consideration of the pseudo-fortress concept. A 'box' is a 2 × 2 array of squares controlled by the pair of bishops. A frequent tactic (in side-variations) to win the knight is a 'pin-crucifix' or a 'checking crucifix', which are simply special cases of the box. (See R11 and R12). Now if Black has a pseudo-fortress,

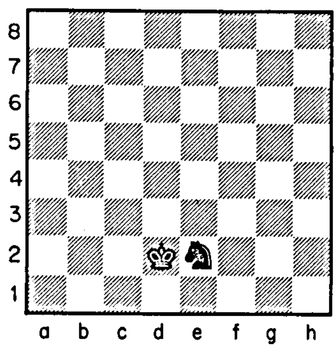

R10. A typical 'pseudo-fortress' with summed distances FOUR from 'b2' or 'g2' positions of the Kling and Horwitz type.

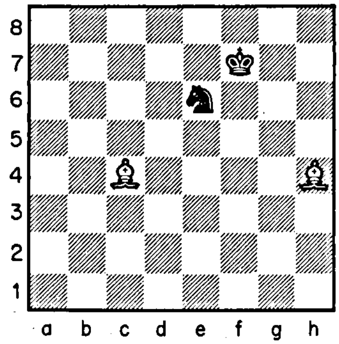

R11. A 'pin-crucifix' using a 'box' e6–f6,e7–f7. The black knight is lost without compensation.

and is happily moving king and/or knight to and fro, making sure to keep the white king at bay by checking when appropriate and then returning to the pseudo-fortress home square, how is he to be evicted? Apart from the king White has only his pair of bishops. They can be used in two obvious ways: on adjacent parallel diagonals from a distance, a classic technique, or by cross-fire to create a box. If we choose a 'box-building' strategy it is not difficult to decide what box is necessary and where the bishops must stand in order to set it up. With the black king on d2 and black knight on e2, the required box will be specified as d2–e2,d3–e3, if we assume that the white king prevents escape to c3. The result of such a box will be to drive the black king towards the cramping edge or a corner without the chance to set up a Kling and Horwitz position. (A box is not a universal panacea. It is a concept to be used with care and cunning.)

6. 'Advancing' the box. Place the black king on d6, with the knight alongside on e6—a pseudo-fortress with summed distance THREE to b7 and THREE to g7. Place the white king on c4, and the white bishops on b2

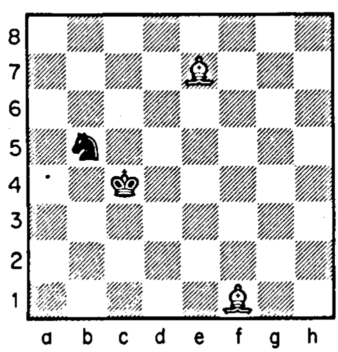

R12. A 'checking crucifix' using a 'box' b4–c4,b5–c5. The black knight is lost without compensation.

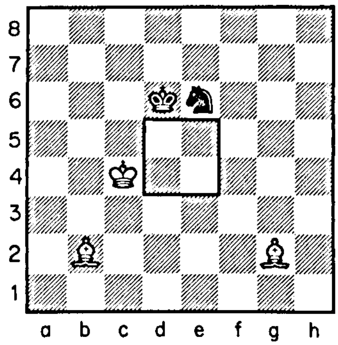

R13. Black has a pseudo-fortress. The box d4–e4,d5–e5 is ineffective. The whole box must be shifted one rank up the board.

and g2. (See R13.) The existing box (d4–e4,d5–e5) is a no-man's-land (Black cannot advance further towards the board's centre) but White's task is to drive Black out of his pseudo-fortress. If the chosen method is the box method, then the d4–e4,d5–e5 box must somehow be transformed into a d5–e5,d6–e6 box. To achieve this new box with the bishops that are able to travel only on either white squares or black squares it soon becomes evident that both bishops must switch sides of the board: the dark bishop on b2 must reach g3 (a square on the h2–b8 diagonal) and the light bishop on g2 must reach b3 (a square on the a2–g8 diagonal). (See R14.) This explains otherwise mysterious bishop moves away from the scene of action: they are unavoidable stepping-stones to where the bishops are needed to (threaten to) set up a new box.

7. *'Squinting' bishops.* Place the light bishop on b3 and the dark bishop on g3. Imagine (or place) the black king and knight on d7 and e7 respectively. (See R15.) The white king is centralized, but on no particular square, so it is omitted from the diagram. The bishops are well

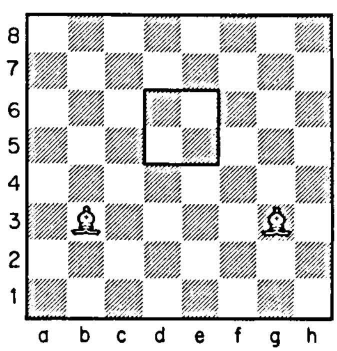

R14. A box d5–e5,d6–e6, showing the bishops on opposite sides of the board compared with the d4–e4,d5–e5 box.

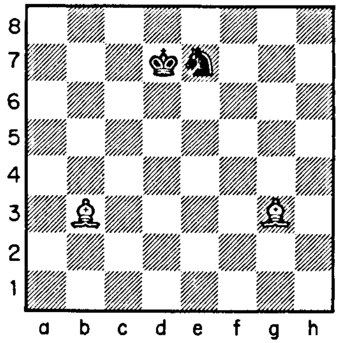

R15. The power of 'squinting' bishops (see text).

placed in at least five important respects: firstly, by controlling c7 and f7 they deter the black king from approaching both a b7 Kling and Horwitz and a g7 Kling and Horwitz; secondly, they are immune from tempo-gaining attack by the black king; thirdly, although not immune from attack by the knight they are relatively so, especially as the centralized white king will in Phase 4 control several squares that the knight would need to pass through to execute a bishop-harassing manoeuvre that is frequently a serious threat in Phase 2; fourthly, they create a d5–e5,d6–e6 box; and fifthly they are poised for a quick switch of sides of the board (with the probable, though not necessarily unique, purpose of advancing the box) in at most two moves (each) while still remaining relatively immune from attack (for instance the dark bishop can play to e1 and thence to b4). On the squares b3 and g3 the bishops are not 'glaring straight down' the long diagonals a1–h8 and h1–a8 but are just off-centre in this sense: hence 'squinting'. The term 'cross-eyed' graphically describes the pair of squinting bishops. The set-up is powerful and economical and it occurs frequently, with only occasionally a better square existing: for instance, the dark bishop on g3 might be vulnerable to attack or to a checking fork by the knight playing to f5, with g7 or b7 as possible defensive havens in consequence, and in this case the very remote square h2 might be superior to g3 for the bishop, but only in the short term.

7. THE STORY OF PART 2

This is the only section of the current paper to be written after the conclusion of the Part 2 test. Part 2 has no protocol corresponding to Part 1. Instead, dated files were created recording interactive sessions with the data base. Most of the time communing with the data base was spent trying out ideas, alternatives, 'what happens if' hypotheses, derived either from the specialist's own manual record of the Part 1 test, or from attempts at a methodical approach to a problem of a particular phase of play. Considerable time was spent in examination of the output from the interactive sessions. At best such examination would answer a question; at worst it would raise further questions. Close study was made of positions around a solution-depth of 20 moves, with the hope that an increased ability to recognize such long conclusions would ease the understanding of Phase 4, where the really deep play takes place.

The following are some of the major new patterns to emerge.

1. The black knight is on f3 (or its symmetrical equivalent) with the black king alongside. The knight has two paths to choose from in heading for a g2 Kling and Horwitz. It is extremely unusual for both paths to be blocked. (See R16.)

2. The pattern of white king making an outflanking move that avoids checks and covers a presumably important square to free a bishop for more important work. (See R17.)

3. Here there is a box that is prevented by the current placing of the black knight, but it is possible to attack the knight by one of the bishops. The box from Black's standpoint includes a potential outlet for his king, a 'valve', but the same square is where the knight may be lost by a pin-crucifix. This pattern might be dubbed the 'box-valve'. (See R18, R19.)

4. In the course of improving the position of the white king it 'follows' his opposite number across the board, keeping to the same or, at worst, adjacent orthogonal, especially during phase 4. In choosing the moment for such a move White must pay special regard to his king's vulnerability to checks from the knight. (See R20.)

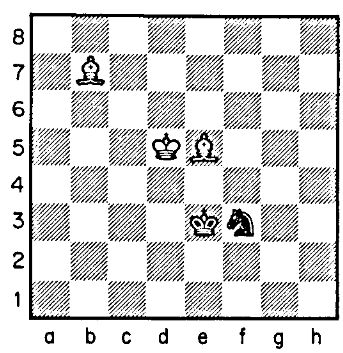

R16. White to move—depth 49. Black has two retreat paths from the square f3 to the Kling and Horwitz square g2: via h4 and via e1. If White prevents this by Bg3, then this bishop has lost its mobility and Black can threaten to set up a b2 Kling and Horwitz position. When Black can obtain a position like this it is a strong indication that the solution depth is near 50.

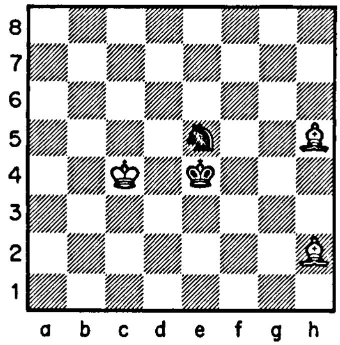

R17. White to move—depth 56. With the sequence Kc5,Nd3+;Kd6, White is safe from immediate checks and prepares to drive the black king towards the edge (any edge) starting with a bishop check on g6.

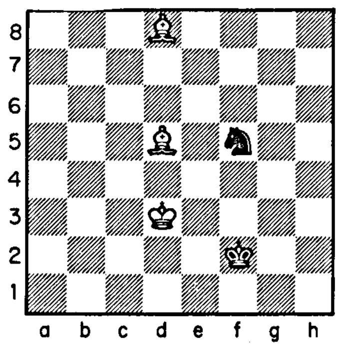

R18. White to move—depth 13. There is a latent box f3–g3,f2–g2. However, the black knight prevents the box move Bh4+. White plays Be4 and after the knight moves Bh4+ can be played.

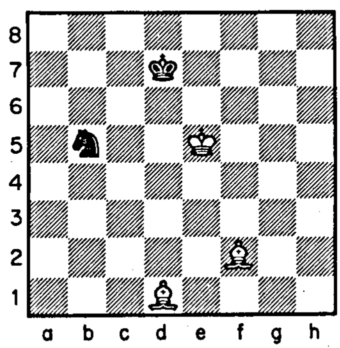

R19. Black to move—depth 19. Black's optimal move is Kc7: which allows White's reply Ke6. The explanation for not choosing Nc3; is that after Bb3,Kc6;Bd4, a box-valve position is created, with the cramping Ba4+ to come.

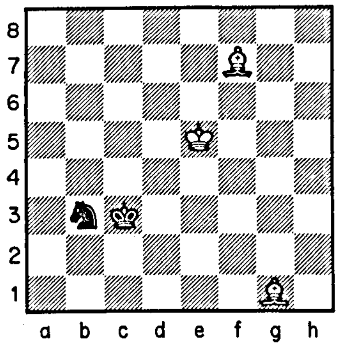

R20. White to move—depth 26. With the move Kd5, the white king is 'following' his opposite number across the board, but occupying central squares in contrast to the black king's more peripheral situation. The move also relieves a bishop of control of an important square, c4 in this case, illustrating the economy of square control by a less mobile piece when the more mobile piece is required for more active work.

5. Black has the occasional strong defence of forcing White to repeat moves as the only alternative to allowing a Kling and Horwitz. Thus a position that has the major feature of phase 4 (Black can be prevented from setting up a Kling and Horwitz) is nevertheless not a win unless progress can be made, and progress, it turns out, is only via phase 3 (a Kling and Horwitz). (See R21.)

6. The black king is forced to block a potential check by his own knight, thus allowing the white king to attack the knight with advantage. (See R22.)

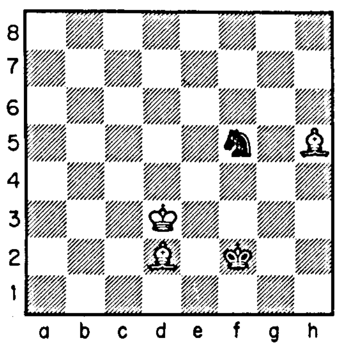

R21. White to move—depth 43. In terms of the Kling and Horwitz prevention heuristic White should control the square h4 by playing Bg5. But Black then renews the threat with Kg3. Then the only way to prevent Nh4; is to play Bd2, as Nh4;Be1+ is good for White. However, in reply to Bd2, Black plays Kf2;, repeating the diagram. White will not make progress by repetition.

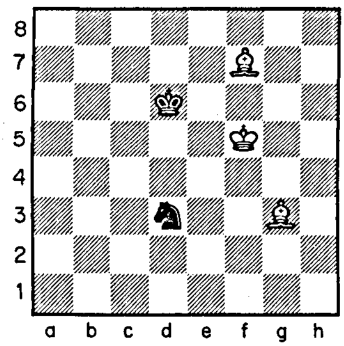

R22. Black to move—depth 49. Black is in check and plays his king to c5. This is optimal, despite the fact that it blocks this square for his own knight. It gives White the opportunity to advance his king with a useful gain of tempo by playing Ke4. However, and this illustrates the profundities of this endgame, the move Ke4 is not quite optimal. The sole optimal move is the mysterious Bh2.

7.1. Depthcharts

A novel technique for expressing in a concentrated visual form the content of the data base was invented. A diagram was produced showing only four of the five chessmen explicitly. The fifth man was added in the form of a number on each of the squares which the 'missing' man could legally occupy, having regard to which side was presumed to have the move. Such diagrams have been provisionally dubbed 'depthcharts'. A depthchart may identify localities where the depth is (for a white depthchart) significantly low, in which case one may confidently conjecture that if the missing white man is not in that area he ought to head towards it. The technique seems promising, but exploiting it methodically was not possible during the preparation period.

7.2. The part 2 test

The second test comprised, like the first, 10 positions.

Number of positions won: seven

Number of positions abandoned: two

One position is unaccounted for: in this, after 105 moves, when the depth of solution was low (nine) a data base move was erroneously executed on the board following the verbal notification over the internal telephone. The consequence was the data base (black) move 'king takes bishop' when the expert's board showed the knight on the square occupied, according to the system, by the black king. This was the only confusion of this kind in all the play of both tests. Performance measured by number of moves played divided by total of optimal lengths: 51.4%.

8. THE ENDGAME IN MASTER PRACTICE

This endgame has occurred several times in the tournament and match practice of the last 25 years but to the best of the author's knowledge the only serious and prolonged attempt to win a deep-solution position took place in the game between International Master Jozsef Pinter (Hungary) and International Grandmaster David Bronstein (USSR) at the international tournament at Budapest in 1978. (See R23; Benko, 1984.)

A comparison of the moves with the optimal moves (obtained by consulting the data base at every step of the game score) is of interest to chessplayers and to non-chessplayers.

Many observations and conjectures are possible arising from the comparison. However, firm conclusions are another matter. We draw none.

All moves played by the two masters from the moment this endgame appeared on the board are given below. The data base assumes that White has the bishops, in accordance with the normal convention of chess endgame theory. Out of respect for the valiant players we keep to the original game colours.

1. The numbering 68–117 corresponds to the serial move numbers of the actual game.
2. Where there is no move in parentheses the played move is optimal.
3. A move in parentheses is optimal, with the implication that the immediately preceding move played in the game is sub-optimal. Note that where there is more than one optimal move, only one is given. In many cases there is only one optimal move, and the occurrence of 12 equi-optimal moves for Black's 95th is unusual.
4. The two-digit number in parentheses after each black move is the optimal depth *after* that move has been played.

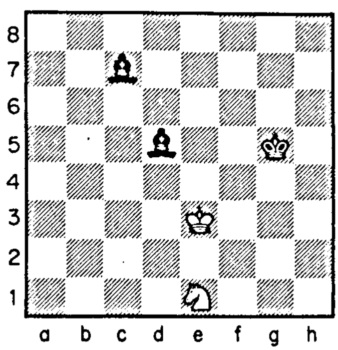

R23. Pinter (Hungary) vs. Bronstein (USSR) Budapest, 1978. White to play—position after Black's 67th move.

It follows that optimal play by both sides is identified by consecutive moves without any move in parentheses, when the depth will decrease by one at every move by Black.

The initial WTM position is lost for White in 54 optimal moves.

White	(optimal)	Black	depth	(optimal)
68.Kd4	(Nf3+)	Bf7	(52)	
69.Nd3	(Ke4)	Kf5	(50)	
70.Kc3	(Nc1)	Ke4	(43)	
71.Nb2		Be5+	(42)	
72.Kc2		Bg6	(42)	(Kd4)
73.Kb3		Kd5	(44)	(Kd4)
74.Na4	(Nc4)	Kc6	(49)	(Ke4)
75.Nb2	(Kc4)	Kb5	(38)	
76.Nd1		Bf7+	(37)	
77.Kc2		Kb4	(36)	
78.Kd3		Bg6+	(35)	
79.Ke3		Kc5	(45)	(Kc4)
80.Nf2	(Kf3)	Kd5	(41)	
81.Nh3		Bd4+	(40)	
82.Kf4		Be8	(48)	(Bg7)
83.Ng5		Bb6	(50)	(Bg7)
84.Nf3		Bc7+	(50)	(Bc6)
85.Ke3		Bg3	(50)	(Bc6)
86.Ng5		Bh5	(49)	
87.Nf3		Bc7	(50)	(Bd6)
88.Kf2	(Ne1)	Ke4	(41)	
89.Nh4		Kf4	(40)	
90.Ng2+		Ke4	(41)	(Kg5)
91.Nh4		Be8	(42)	(Kf4)
92.Ng2		Bc6	(43)	(Bh5)
93.Ne1		Bb5	(44)	(Bb7)
94.Ng2		Bd6	(45)	(Bc6)
95.Nh4		Kf4	(46)	(Bc7, and 11 other equally optimal moves)
96.Ng2+		Kf5	(45)	
97.Ne3+		Kg5	(48)	(Ke4)
98.Ng2	(Nd5)	Bc6	(43)	
99.Ne1		Kg4	(52)	(Kf5)
100.Ke3		Bc5+	(51)	
101.Kd3		Bf2	(50)	
102.Nc2		Kf4	(49)	
103.Na3	(Nd4)	Bh4	(49)	(Be4+)

White	(optimal)	Black	depth	(optimal)
104.Kd4		Be8	(48)	
105.Nc4		Bf7	(51)	(Bf2+)
106.Kc5	(Nd6)	Bf2+	(33)	
107.Kb4		Bd4	(32)	
108.Kb5		Bh5	(47)	(Be8+)
109.Kc6		Bf3+	(50)	(Ke4)
110.Kc7	(Kd6)	Kf5	(46)	
111.Na5	(Kd6)	Ke6	(44)	
112.Nb7	(Nc6)	Ke7	(45)	(Be5+)
113.Na5		Bf2	(44)	
114.Nc6+		Ke6	(43)	
115.Na5	(Nd8+)	Bg3+	(41)	(Ke4)
116.Kb6		Bf2+	(40)	
117.Kc7				

At this point the game was declared drawn by the '50-move rule', as no pawn had been moved and no capture had taken place for 50 consecutive moves.

We permit outselves five factual observations and one psychological comment.

Firstly, it may be seen that Pinter set up the Kling and Horwitz position no fewer than three times, each time in a different corner, namely after his moves 71 (in the b2 area), 90 (in the g2 area) and 112 (in the b7 area).

Secondly, Pinter's knowledge of the Kling and Horwitz position leads him consistently to head for it with unnecessary haste, this accounting for a number of his sub-optimal choices.

Thirdly, the quality of an individual sub-optimal move by either side can be crudely measured (if the opponent's previous and subsequent moves are optimal) by comparing the successive depth numbers. Thus it can be seen that Bronstein's 79...Kc5; increased the optimal depth from 35 to 45, a cruel penalty for not playing the optimal and so similar Kc4; while Pinter's 88.Kf2, reduced the optimal depth from 50 to 41. The great difficulty of this endgame is evident when one tries to give well-grounded reasons for the played moves being inferior to the optimal moves. (See R24.)

Fourthly, an optimal win within the confines of the traditional 50-move rule became possible after White's 70th move, but was lost, never to recur, with Black's 79th.

Fifthly, 36 of White's 50 moves were optimal, while only 26 of Black's

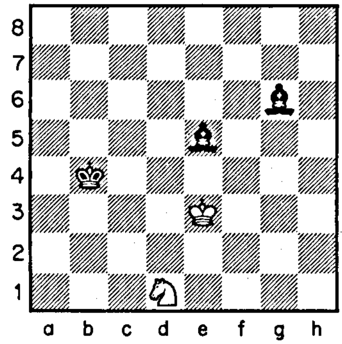

R24. Pinter vs. Bronstein. Position after White's 79th move. Is 79...Kc4; or 79...Kc5; the better move, and why?

49 were optimal. A conjecture is that this is evidence for the endgame being more difficult to play for the side with the bishops.

The psychological comment is that a mistake (as distinct from a crude blunder or oversight) of the kind of Bronstein's 79...Kc5; or Pinter's 88.Kf2, although it leaps to the eye when scanning the depth parentheses to the above game score, is recognized by the player, if at all, only several moves later. The player's general strategy will in all probability have been correct at the highest level, but his ability to calculate in order to reject moves that appear to meet the strategic objectives equally well (which nevertheless fail when countered by an optimal continuation) will be insufficient: one or more vital concepts are missing. Before the creation of the oracle data base no one could have described the feeling for position and depth of calculation needed to play this endgame really well. Now it begins to be possible. The missing concepts are waiting to be formed from data in the data base, and to be verified by reference to the same data base. The idea of *automatic* derivation of concepts or patterns meaningful to human domain specialists is a challenging, a tantalizing, possibility—is its realisation just round the corner or is it remote?

Acknowledgements

The 'BBN' data base was generated and generously made available by Mr Ken Thompson of Bell Laboratories, New York, USA. Dr Alen Shapiro of Intelligent Terminals Limited set up the data base in the Turing Institute, added an interactive interface, and was responsible for the admirably smooth conduct of the computing side of the contest. Professor Donald Michie directed the project. I am grateful to The Royal Society and to the Science and Engineering Research Council of Great Britain for the award of an Industrial Fellowship, and to my employer, IBM United Kingdom Limited, for their contribution and associated secondment.

REFERENCES

Entries entirely within brackets are recognized authorities on chess endgame theory not explicitly referred to in this paper.

Alty, J. I. and Coombs, M. J. (1984) *Expert systems—concepts and examples.* NCC, Manchester.

[Averbakh, Yu. I. (ed.) (1980–1984) *Shakhmatnye okonchaniya.* In five volumes. Fizkul'tura i Sport, Moscow.]

Benko, P. (July 1984) Electronic arts. In *Chess life,* pp. 44–5, 58. United States Chess Federation, New Windsor, NY.

Bouwkamp, C. J. (1967) Catalogue of solutions of the rectangular 3 × 4 × 5 solid pentominoes problem. Technological University Eindhoven and Philips Research Laboratories, Eindhoven.

[Chéron, A. (1960–1970) *Lehr- und Handbuch der Endspiele.* In four volumes. Siegfried Engelhardt, Berlin.]

Comay, O. and Roycroft, A. J. (April 1984) Two bishops against knight (continued from *EG*74). *EG*, no. 75. 249–52.

[Fine, R. (1941) *Basic chess endings.* David McKay, Philadelphia.]

[Hooper, D. V. (1970) *A pocket guide to chess endgames.* Bell, London.]

Kling, J. and Horwitz, D. (1851) *Chess studies, or endings of games.* Skeet, London.

Kopec, D. and Niblett, T. B. (1980) How hard is the play of the king–rook–king–knight ending? In *Advances in Computer Chess 2* (ed. M. R. B. Clarke), pp. 57–81. Edinburgh University Press, Edinburgh.

Michie, D. (1961) Trial and error. *Science Survey,* **2,** 129–45. Reprinted in *On Machine Intelligence* (2nd edn). Ellis Horwood, Chichester (1986).

Michie, D. (1986) The superarticulacy phenomenon in the context of software manufacture. *Proc. R. Soc. Lond.* A **405,** 185–212.

Nilsson, Nils, J. (1971). *Problem-solving methods in artificial intelligence.* McGraw-Hill, New York.

Pachman, L. (1983) *Chess endings for the practical player.* Routledge & Kegan Paul, London.

Rich, E. (1983) *Artificial intelligence.* McGraw-Hill, New York.

Roycroft, A. J. (1972) *Test tube chess.* Faber & Faber, London.

Roycroft, A. J. (1981) *The chess endgame study.* Dover, New York.

Roycroft, A. J. and Niblett, T. (June 1979) How the GBR class 0103 data base was created. *EG* no. 56, 145–6.

Shirai, Y. and Tsuji, J. (1982) *Artificial intelligence.* John Wiley, Chichester.

Ströhlein, T. (1970) *Untersuchungen uber Kombinatorische Spiele.* Technische Hochschule, Munich.

Thompson, K. and Roycroft, A. J. (November 1983) A prophecy fulfilled. *EG* no. 74, 217–20.

Thompson, K. (May 1986) *C* the programs that generate endgame data bases. *EG* no. 83, 2.