Chess, Imagination, and Perceptual Understanding

PAUL COATES

1. Introduction

Chess is sometimes referred to as a 'mind-sport'. Yet, in obvious ways, chess is very unlike physical sports such as tennis and soccer; it doesn't require the levels of fitness and athleticism necessary for such sports. Nor does it involve the sensory-governed, skilled behaviour required in activities such as juggling or snooker. Nevertheless, I suggest, chess is closer than it may at first seem to some of these sporting activities. In particular, there are interesting connections between the way that we use our perceptual imagination in sports, and also in chess. The same distinction between calculation and natural instinct applies in chess as it does in many physical sports.

Meanwhile, a proper understanding of the way perceptual imagination works in chess points to something fundamental about the way we understand the physical nature of the world. A good chess player's appreciation of the tactical and strategical aspects of a chess position involves an imaginative grasp of spatial possibilities. This understanding is mirrored in a straightforward sense in physical sports when a player anticipates a return pass in football, or places a shot in tennis.

But there are connections at a deeper level as well, between the way we see chess positions, and the way we see ordinary objects such as apples and chairs. Our understanding of the physical world has much more of a formal nature than we unreflectively appreciate. Seeing the difference between a ball and a cricket bat is analogous in many ways to understanding the difference between a bishop and a knight in chess.

2. Calculation and the Combinatorial Explosion

In order to fully understand the way good players see chess positions, we need first of all to clarify the difference between chess tactics and strategy.

doi:10.1017/S1358246113000258 ©The Royal Institute of Philosophy and the contributors 2013

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

Chess involves the calculation of tactical possibilities latent in a position, so as to work out what advantages may be achieved in a given situation – such as winning a piece, or forcing checkmate. If the position is otherwise equal, winning even one pawn is usually sufficient for a good club player to force victory.

But chess involves more than mere calculation. For in nearly all chess positions, the number of possible sequences multiplies at an exponential rate, the further one looks ahead. In an average middle game position, the player has around thirty single moves to choose from. For each of these, the opponent will, likewise, have around thirty replies. So in most given positions there will be getting on for a thousand possible full, one-move, sequences (one move by each player). To assess *every* possibility, even three full moves ahead, would require examining many millions of possible outcomes. Fully calculating the possibilities latent in a position is clearly not feasible.

Human chess players therefore have to be selective. They need to think strategically. At every stage they have to massively restrict the number of moves that they examine. Accordingly, they focus upon only those moves that might impact on the few select aspects of the position that appear to be relevant. To do this, chess players need to use *intuition*. Intuition enables them to assess which side has the advantage in a position, and what particular moves and general plans are worth exploring, and are likely to promote an advantage.

Chess therefore involves two main skills. The first is the ability to calculate sequences of moves accurately. This requires *tactical expertise*, or what cognitive scientists call the ability to '*search*'.

The second skill involves being able to assess positions and select the moves most likely to lead to positions that are obviously better. This demands *strategical proficiency*. Good players become experts in pattern recognition, developed through experience, and rely upon *intuition*, which, in this context, involves an ability to recognise common patterns of pieces as '*chunks*' or '*templates*'.[1] These units they associate with likely outcomes. Players will see a group of pawns as mobile or fixed, for example, or a king position as weak or well defended, or the possibility of a back rank mate, and so on.

Today, computers have reached levels at which they are capable of defeating world champions at chess. But computers play in a different style from humans. In short-term tactics they play perfectly. However, it proves very hard to programme computers to 'assess positions' strategically. They rely mainly upon what is called 'brute

[1] See, for example, Chase and Simon (1973) and Lane and Gobet (2011).

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

force' – the ability to calculate literally billions of different positions very rapidly, so that they can compute in depth to around 8–10 full moves ahead, and with complete accuracy (the use of a large memory, for chess openings and endgame positions, also plays a role). However, their assessment of the resulting positions can sometimes be quite mistaken, as I shall illustrate in due course.

3. Chess and the Imagination

In both tactics and strategy, the good chess player makes use of his or her *imagination*, and I want first to distinguish various different senses of this term.

(1) *Basic imagination*:

In introductory books on chess, the nature of each type of piece is explained by diagrams indicating the legal moves open to it. Thus the bishop is characterised by the fact that it is only allowed to move along diagonal lines; the knight, by how it can jump to the opposite corner of a rectangle of 2 by 3 squares; and so on for the other pieces:

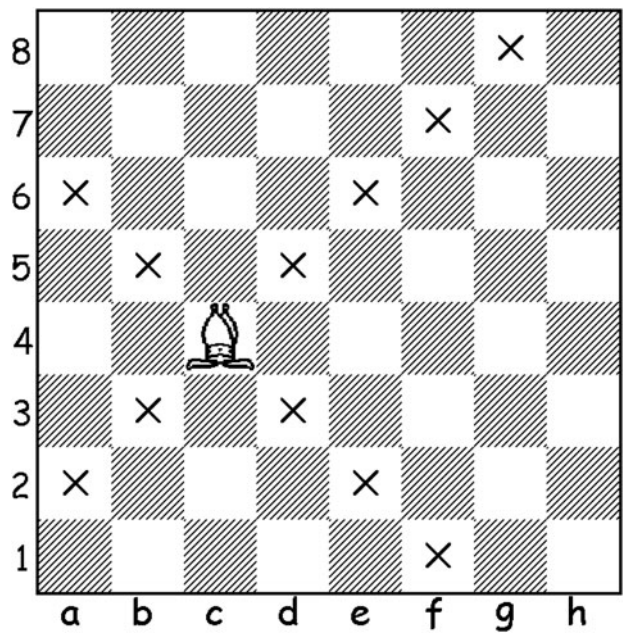


Diagram 1. Possible bishop moves indicated by the crossed squares.

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

This already gives us a first sense of 'imagination', which I shall call the *basic* sense. To see a given chess piece as belonging to a certain type, such as a bishop, is to be aware of what squares on the board it can move to. This sense of imagination is closely related to Kant's important notion of the *Productive Imagination*.

As Kant argued, when I grasp the nature of a common-place object that I directly and spontaneously recognise in perception, I exercise a low-level classificatory concept. I cannot see, or otherwise perceive, an object – in a manner which involves some awareness of what I see – without having an implicit grasp of its different possible appearances. To see a black object on a chair *as* a cat necessitates being aware of how that object would appear from different points of view. The productive imagination imposes a unity upon different *spatial* perspectives of one and the same object.[2]

In a related sense, to grasp what type a chess piece belongs to – whether it is a rook, or king, and so on – requires an understanding of the legal moves permitted to it in the context of a game. In grasping *this* particular piece as a rook, I understand where it could move to on its next move in *this* particular position. The imagination imposes a unity on different *temporal* perspectives, so that two successive appearances of a rook on different squares are taken to be the play of a single legal move by one piece. The different overall positions are seen as belonging to one and the same game of chess, related sequentially.

When a competent player sees his opponent's rook on a certain square, he is thus prepared for it to move in certain specific ways. This basic imaginative grasp of the possibilities open to a piece is *constitutive* of understanding the type of piece it is. Moreover, the understanding of its legal moves is *all* that is essential to understanding the nature of a chess piece. The assignment of standard shapes to each type of chess piece is a merely contingent matter, and was regularised in the nineteenth century by the adoption of 'Staunton pattern' figures.

Precisely what is involved in this kind of imaginative seeing turns out to be a complex matter, and we will examine later on further aspects of seeing a chess position that are connected with this point.

(2) We are more used to thinking of the imagination in chess in two further senses, which need to be distinguished from what I have termed the *basic* sense. These have to do the use of the imagination in clever tactical sequences, and in more profound strategical plans.

[2] A modern development of these ideas is set out in Sellars (1978).

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

The beginner will have learnt what a knight is in terms of the way that it can move and carry out forks – by attacking two pieces at once, as in Diagram 2:

1 Nf7ch Kg7
2 Nxg5

When we speak of an *imaginative tactical sequence* in chess, we have in mind the way that such simple ideas can enter into play, in more surprising ways. The proficient chess player, for whom the legal moves become second nature, instinctively guards against moves that straightforwardly lead to the capture and loss of an unprotected piece. The awareness of simple short sequences that result straightforwardly in loss guides the selection of sequences worth exploring further. Good players become familiar with forks, pins, discovered checks and other simple tactical devices, knowing what to avoid. But there is a critical balance to be achieved. While avoiding moves that lead to loss, the player must also develop a sense of the kind of situation that needs to be examined more closely. A move that appears initially unattractive may turn out to lead to greater gain.

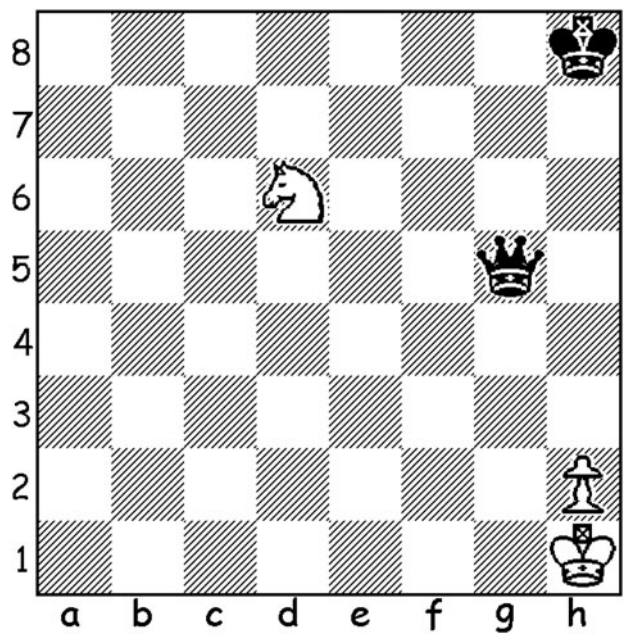


Diagram 2. The simple knight fork of king and queen.

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

This second use of the chess imagination enables good players to discover effective moves that lesser players might overlook.

Here is a relatively simple example, from the end of the 10th Match Game between Petrosian and Spassky, first match, 1966 (Diagram 3):

Petrosian finished off the game nicely by playing:

1 Bxf7ch Rxf7
2 Qh8ch!

At which point Spassky resigned, because of the forced sequence leading to a clearly lost position:

2 ….. Kxh8
3 Nxf7ch Kg7
4 Nxg5

In his famous game against Donald Byrne, played when he was only 13 years old, Bobby Fischer found a wonderfully imaginative tactical combination. Starting with a surprise queen sacrifice, he eventually won material after a forcing sequence of nearly ten moves (Diagram 4):

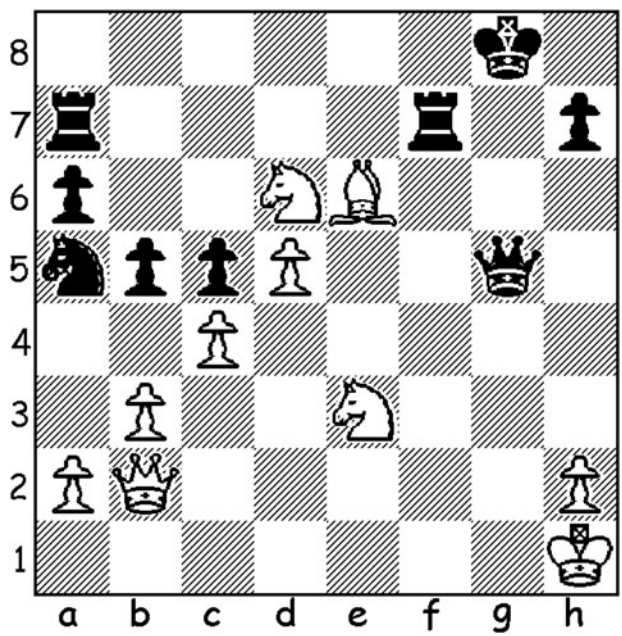


Diagram 3. Petrosian v. Spassky, First Match, end of Game 10.

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

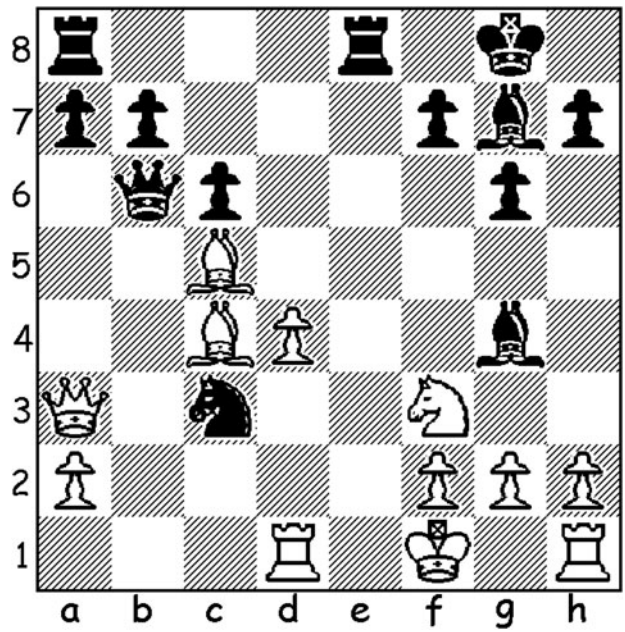


Diagram 4. D. Byrne v. Fischer, Rosenwald Trophy Tournament, US, 1956.

Fischer, playing black, achieved a won position by a spectacular combination as follows:

17….	Be6!!
18 Bxb6	Bxc4ch
19 Kg1	Ne2ch
20 Kf1	Nxd4ch
21 Kg1	Nxe2ch
22 Kf1	Nc3ch
23 Kg1	axb6
24 Qb4	Ra4
25 Qxb6	Nxd1

and with a rook and two bishops for a queen, Fischer had no problem in successfully winning the game.

In the Byrne-Fischer game, the actual combination itself is not too difficult for the average club player to follow. What was more remarkable was Fischer's imaginative ability in recognising the potential for the sacrifice in advance, and exploring its tactical possibilities. Imagination, in this *second* sense, involves examining

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

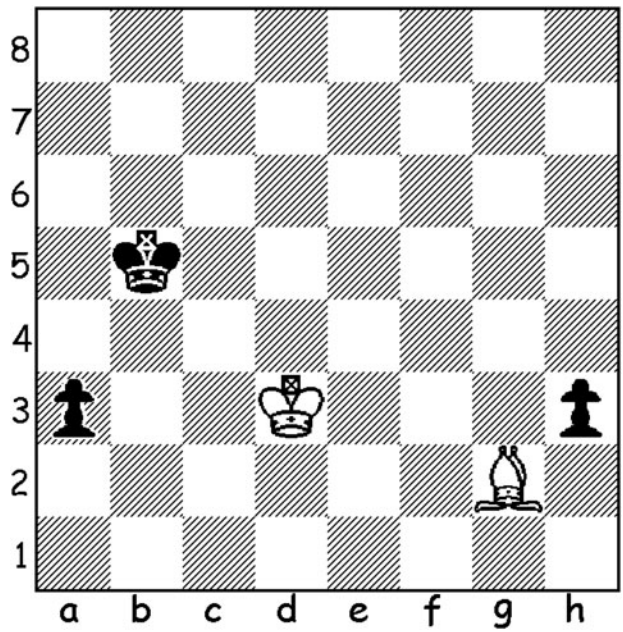


Diagram 5. Gratzer Endgame – queening threats: White to play and draw.

an unusual move or position, and seeing a tactical possibility that does not at first strike the player as obviously leading to an advantage.[3]

Here is another example of this kind of exercise of imagination, the end of a lovely endgame by Gratzer, in which the task is for White to find a way of forcing a draw (Diagram 5).

Although it is White to move, it looks initially as if one of the black pawns must go through to become a queen, in which case Black will win. It appears that White must first capture Black's pawn on h3, which then leaves the White king too far away to stop the Black pawn on a3 from queening. The sequence would be:

1 Bxh3 a2
2 Kc2 a1 = queen, and Black wins

[3] In his pioneering research, de Groot (1965) distinguished between a player's initial exploration of the chess position, the selection of moves worth investigating in a given position, and a subsequent investigation or analysis of the actual tactical possibilities resulting from the move.

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

However, White has a remarkable imaginative resource, a move that at first glance seems hardly worth considering, but which achieves a draw. The main variation is:

1 Bf1!! a2
2 Kc2ch

the discovered check saves the day for White.

2 … Kb4
3 Kb2 h2
4 Bg2

and with both Black pawns prevented from queening in safety, White achieves the draw.

(3) A third kind of imaginative ability is exercised when players consider the strategical possibilities in a position.

As we have already noted, the basic rules of the game stipulate which moves are legal. They are constitutive of chess. In a different sense, there are heuristic rules, rough principles that help players to select moves that will improve their position. Centuries of play were needed before players became able to articulate such heuristics, strategical rules such as:

Capture with pawns towards the centre;
Place rooks on open files, or on the seventh rank;
Knights are better placed in the middle of the board than on the side;
Aim not to move the same piece twice in the opening;
Overprotect weak points;
Don't let pawns become isolated;
and many more.

These rules are not legal requirements. None of them applies in all situations. But on the whole, following them promotes success. A good player will have internalised scores of such rules, and will be guided by them in strategical play.

Diagram 6 shows is a famous example, from the tournament game Emanuel Lasker v. Capablanca, St Petersburg 1914, final tournament:

Here Lasker, as White to move, successfully carried out the strong strategical plan of establishing his knight in the heart of his opponent's position, by playing: Ne6. The knight on the 6th rank was, for Black, 'like a bone in his throat' in Kasparov's words.[4]

But there is room for a more creative use of the imagination at a strategical level, where mere heuristic rule-following becomes

[4] Kasparov (2003), 212.

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

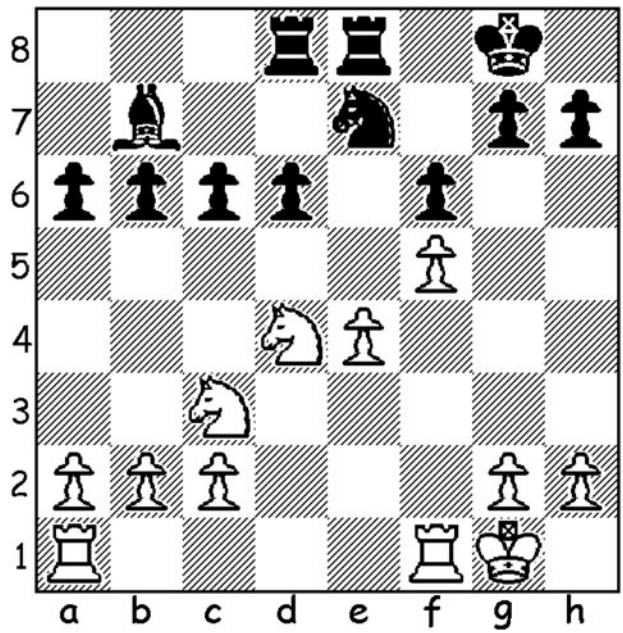


Diagram 6. Lasker v. Capablanca – White to play.

subservient to deeper aims. A nice example is given in another Fischer game, played against Matulovic in an international tournament (Diagram 7).

Here Matulovic had just played: 13 Bxe7, capturing Fischer's bishop. What might be thought of as the 'normal' move, the one that most averagely good players would automatically select, is a re-capture of the bishop by the queen:

13.... Qxe7

On the grounds that this would enable Black to castle his king into safety.

Fischer saw deeper into the position and realised that his king would be safe in the centre, and also that his queen would be needed on the queenside. He therefore played instead:

13 .... Kxe7!

and a dozen moves later his choice was vindicated when they reached a position where his forces dominated the board (Diagram 8).[5]

[5] After the moves: 14 Qd2 Nf6; 15 Bg2 Bb7; 16 Qd3 Qb6; 17 0–0 a5; 18 R(f)d1Ba6; 19 Qd2 R(h)c8; 20 h3 h5; 21 b3 Bxe2; 22 Qxe2 Rc3; 23 Rd3 R(b)c8; 24 Rxc3 Rxc3; 25 Kh2 Qc5; 26 Ra2....

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

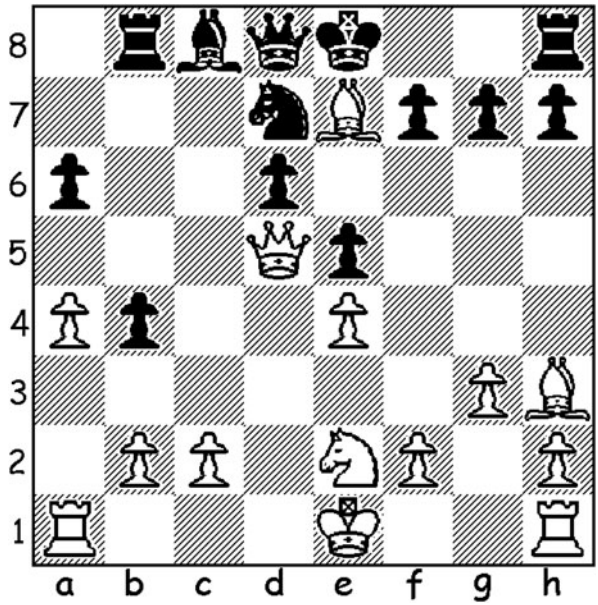


Diagram 7. Matulovic v. Fischer, Vinkovci, Yugoslavia, 1968.

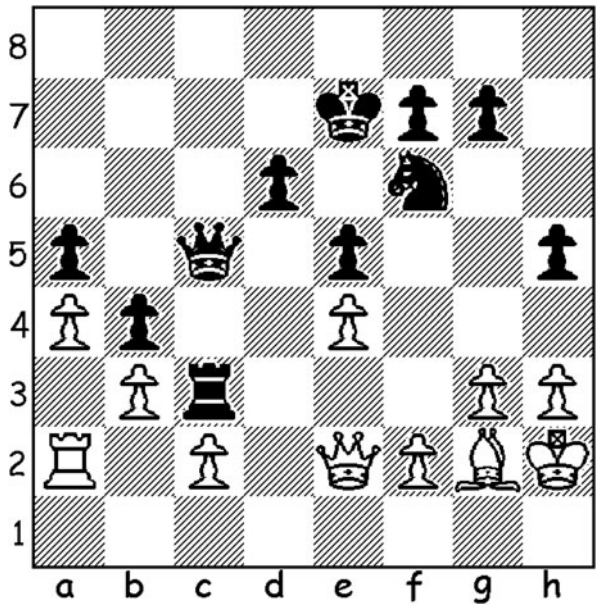


Diagram 8. Matulovic v. Fischer, position after White's 26th move.

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

Fischer won the game after a further dozen moves.

Here the imagination is being used in a different, creative way, looking more deeply into a position, and selecting an unusual move that proves to have beneficial long-term consequences.

Attractive as these examples of tactics and strategy are, the imaginative abilities they illustrate are complementary to an understanding of the rules of chess, rather than constitutive of it. Yet there is still a common idea involved. All three uses of imagination – in the basic sense of understanding how a piece moves, in discovering a tactical sequence, and in grasping the strategical aspects latent in a position – involve the subject considering future possibilities. Knowing that such possibilities exist enables a player to assess the potential of a position.

4. How Chess Players Think

Serious psychological investigation of the thought processes of chess players began with the work of Alfred Binet at the end of the nineteenth century. In 1946, the Dutch International Chess Master and psychologist, Adriaan de Groot, carried out a major study comparing the thought processes of top level Chess Grandmasters, who included world champions Alekhine and Euwe, with those of good club players.

De Groot's central finding was that the use of pattern recognition in soundly assessing positions was of greater importance than the ability to search accurately and to calculate positions in depth.[6] Of course, both abilities have to be well developed in a good chess player. At nearly all stages in a game a player will need to spend some time in careful calculation. But it is correct to say that chess players have to rely upon their intuitive skill in recognising patterns to guide their selection of the moves worth examining. It is this ability to select – usually in a matter of a few seconds – the few appropriate moves that merit further calculation that really distinguishes top level masters from averagely good players.

Current work on thinking in chess has upheld de Groot's findings about the central importance of intuition, or pattern recognition. Fernand Gobet summarises the findings of recent research as follows:

- Grandmasters are very good at the initial selection of moves worth considering, through pattern recognition, and they rely

[6] de Groot (1965).

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

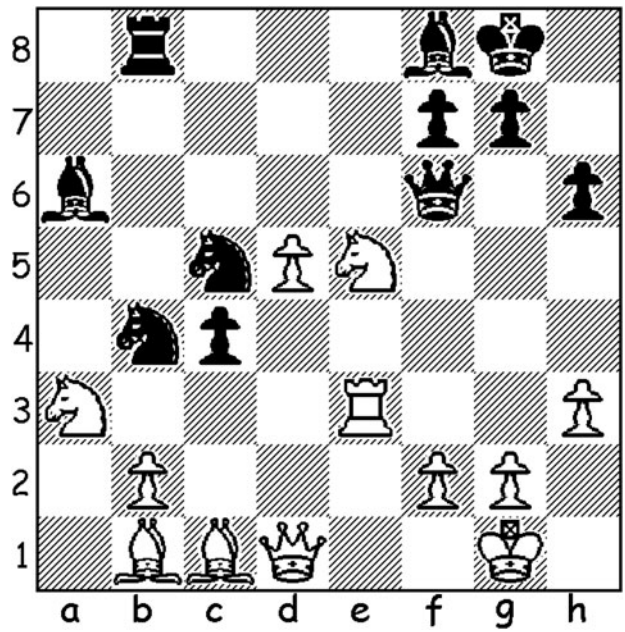


Diagram 9. Kasparov v. Karpov, Black to play.

heavily upon their use of intuition – that is by assessing and categorising a move or position as strong or weak, without explicit conscious calculation. They are also good at what is termed 'chunking': memorising typical groupings of pieces (empirical tests on memory confirms this finding);

- The knowledge of typical position types, and the use of intuition and pattern recognition, are more important than the ability to search and calculate deeply in a position;
- Intuition characterises the thinking of experts all the way through the decision processes.[7]

Yet experts also rely upon their ability to calculate more deeply, and slightly more widely, than average players. Gobet rightly criticises the work of Dreyfus in this respect. Drefus over-emphasises the reliance upon intuition by strong masters. He has claimed that the expert chess player proceeds intuitively for most of the time, rarely explicitly calculating sequences of possible future moves. This simply does not match up with the experiences of experts. Although they may intuitively, rapidly, and for the most part

[7] Gobet (2012).

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

accurately, *select* the right moves to analyse, they still need to go through a *justification* process, in which the likely responses are quickly and carefully – and consciously – calculated. And at times a grandmaster may spend an hour or more calculating different variations in a given position. An example occurred in the 16th game of the third championship match between Kasparov v. Karpov in 1986.

In Diagram 9, with only one pair of rooks and a few pawns exchanged, there is a myriad of relevant possible variations that need to be examined. Karpov had a long think, for around a full hour, mainly devoted to deciding which knight to play to d3. In his book on the match, Kasparov devotes several pages to the ramifications of this extremely complex position, concluding that in eventually deciding to play: 25 … Nbd3, Karpov selected the wrong knight to move to the d3 square.[8]

5. Seeing Positions and the Exercise of Concepts

In one respect, however, Gobet is misleading. He claims that, for the good chess player, the use of intuition precedes any conceptualisation. Now, it is true that chess players talk about having a good '*sight*' of the board, in swiftly and directly seeing the potential of a position. In this respect it has been claimed that at his peak, Capablanca was supreme. Reti complained that when he would show Capablanca a difficult endgame study that he had composed, the latter would instantly '*see*' the solution, almost before he had finished setting up the pieces. But it is a mistake to think of this kind of seeing as preceding the exercise of concepts by the player. Already, in seeing what the position on the board *is*, a chess player is employing classificatory concepts.

As we noted earlier, understanding the way a chess piece can legally move is internally related to, or necessary, for a grasp of which type of piece it is. It follows that what chess players call 'intuition' already involves the conceptual classification of specific pieces, and also of types of position. The exercise of this capacity is linked to an implicit grasp of the potential moves available to a piece. To see a piece as a knight in chess is to understand what it can do, to grasp its function.

This point is well illustrated by the occasional striking blunders perpetrated by even top grandmasters.

One tragic example occurred in a game between Petrosian – before he became world champion – and Bronstein (see Diagram 10):

[8] Kasparov (1987). Kasparov's next move in reply also turned out to be an error.

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

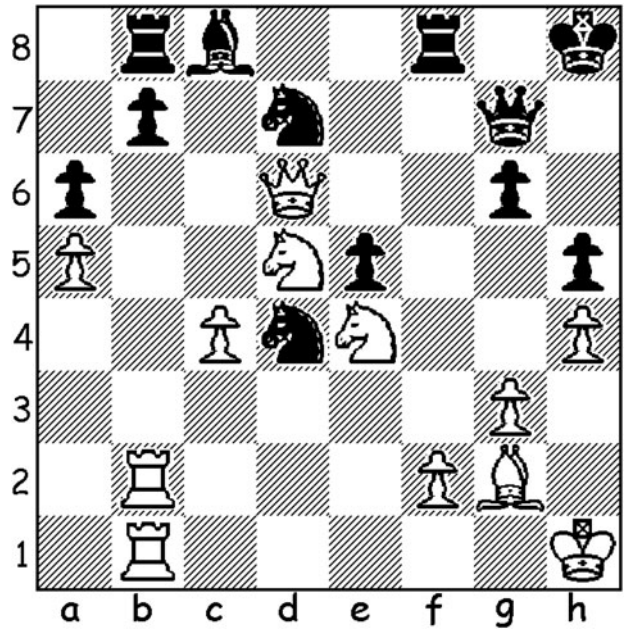


Diagram 10. Petrosian v. Bronstein, Amsterdam, 1956 – Black to play.

Petrosian had achieved a completely dominating position, and it seemed only a matter of time before he would win the game. Bronstein had been reduced to shuffling his knight around the squares: c6, d4 and f5.

In the diagram position, Petrosian had just played his queen to d6. Bronstein responded by making another knight move, 35 .... Nf5. Petrosian thought for a short while, and then blundered by:

> 36 Ng5??
>
> and after the obvious reply 36 ... NxQ he resigned immediately.

Petrosian must in some sense have seen the black knight, but only as a black shape. It was, after all, there in front of his eyes for several minutes as he looked at the board, while he pondered his next move. Yet Petrosian did not see the knight piece *as* a knight; he did not see its potential moves, that it was attacking his queen.

Another very interesting blunder occurred in a game between Fischer and Geller (Diagram 11).

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

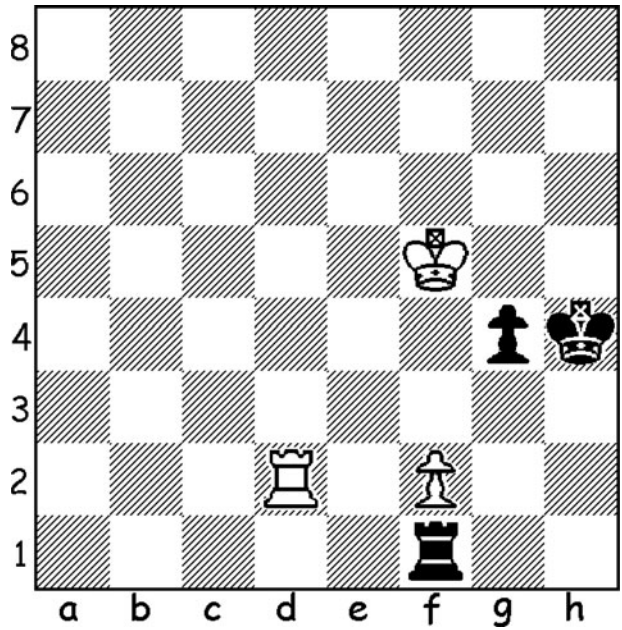


Diagram 11. Geller v. Fischer, Interzonal tournament, Palma de Mallorca 1970.

Geller was in a poor position. Being extremely tired at this point, he was relying upon the following variation, assuming that it led to an advantage for White:

67.... g3
68 fxg3 ch Kxg3
69 Kxf1

The problem is that this variation contains *two* illegal White moves:

** 68: pawn (f) takes pawn (g3), check,

and

** 69: king (f5) takes rook (f1)!

Geller was conceptualising his king piece as a rook, a piece that, of course, moves in a quite different way. Geller was no longer seeing his king piece *as* a king.

These examples illustrate one essential feature of the way chess players see positions. For a player to conceptualise a rook piece as a rook, it is necessary that he sees the potential moves open to it – for

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

example, that it might be moved along a certain file and capture a certain pawn. Thus conceptualising a piece involves a grasp of spatial possibilities – imagining, in the basic sense, the range of squares that it can legally move to. As we noted earlier, there is a sense in which the chess player is implicitly aware of the possible future appearances of the chess pieces on the board. Up to a point, this is correct. However, there are further issues that complicate the analysis of what is involved in playing a game, which require careful explication.

6. The Formal Nature of Chess

All perceptual experiences have two components. Firstly, they involve some kind of intentional component directed at an object, in virtue of which the subject focuses conceptually upon the nature of what is taken to be perceived. Experiences also contain a further, non-conceptual, phenomenal aspect. This second aspect is interpreted in different ways, either as the sensory *mode* of the intentional state, or as a distinct sensory *component*. In either case, it means that in addition to the intentional content, experiences involve a feature that is distinctive of perception, the awareness of phenomenal qualities.[9] These qualities are connected in some way to the external objects perceived.

In seeing a chess position on the board in an ordinary game, the player will entertain 'perceptual takings' about the game – perceptual thoughts, and intentions, involving the exercise of concepts focusing upon the pieces, their positions, and their possible future moves. Such thoughts will be guided by a visual awareness of phenomenal qualities, determined by the actual physical arrangement of pieces on the board.

If we consider any particular game of chess – for example, the sixth game between Tal and Botvinnik, in their 1960 world championship match – we need to distinguish three realms. First, there are the various concrete representations of the game, including the physical arrangement of pieces on the board in front of the players at the time of play; second, there are the players' visualization, thoughts, intentions and other mental states; but in addition to these two realms, there is a third realm, which includes the particular game itself, the

[9] This sensory component is, in turn, interpreted in contrasting ways by the Causal Theorist and the Direct Realist.

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

focus, or object, of the players' thoughts, an abstract entity that is distinct from the various mental and concrete representations of it.

There are several (interconnected) reasons for distinguishing the objects belonging in this third realm from those in the first two. Firstly, for the experienced player, in thinking about the game they are playing, the precise physical details of the appearance of the pieces drop away as unimportant – one might, for example, play a chess game with pieces taken from two sets, where the pawns are of a slightly different height and shape.

Secondly, what essentially determines a piece's identity are the legal moves that are permitted to it. Awareness of the type of piece occupying a square consequently has a normative dimension. A piece is conceived as a knight, or a rook, and so on, in virtue of the fact that the player is implicitly aware of the different new positions it could take up; how it might appear on the board a move or more hence.

Thirdly, if the players are both playing blindfold, there may be no board, or a board whose only purpose is to act as a record of the moves and position reached. In playing postal chess, or on the internet, different players and observers will be looking at different boards, all of which represent the same position. When Spassky was playing a candidates' match against Korchnoi in 1977, he alternated between contemplating the position on the board in front of him, and studying (the same) position on the large display board situated at the back of the playing arena. If a game is adjourned, it may be renewed with a different set of pieces, and so on. If an actual set of pieces is used, these might get accidentally moved or damaged during a game. (Thus a player is permitted to say 'j'adoube', and to adjust the placing of a piece on a square.) The true position in the game will be reconstructed from the previous intended moves; it is not determined by where the pieces actually sit on the board.

For all these reasons, we need to distinguish between the identity of a particular chess game that is being played between two players, and the concrete physical structure that may represent it for them during any period of time. The 'real' chess game that they are considering is something abstract, a structure or model conceived in formal terms. This model has an abstract spatial structure, one that reflects what is in common to all the possible representations of the position. The structure represented is not to be understood as equivalent to any concrete region of physical space. It is not determined solely by the physical arrangement of pieces on the board, but by the joint intentions and actions of the players, as these are integrated into the wider intentions of the chess-playing community that decides the rules of the game.

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

Hence it is important to realise that, when the chess player focuses upon those aspects of the position that determine the actual chess game being played, their thoughts about the game are not, in a strict sense, directly *about* the physical pieces which visually prompt and guide them. Strictly speaking, the real focus of the player's thoughts is *not* the actual, static, arrangement of physical pieces on the board in front of him, but an abstraction from them. Their perceptual thoughts do not refer to the concrete position, but to something akin to a formal model. When chess players look at positions, understanding what they see, they *project* the more immediate awareness of the colours, shapes and position of the pieces on the board onto the abstract position that the physical pieces represent.

The actual physical pieces on the board are proxies that are used to illustrate the abstract structure contemplated. They can be used by the players to represent the position arrived at after each move. In seeing a position, the good player is guided by a non-conceptual *awareness* of the actual pieces' shape, yet *perceives* each piece in terms of its spatial role or function. He or she does not pay much attention to the phenomenal appearance of the piece. Blindfold chess players report that when they imagine positions in their mind's eye, they do *not*, for the most part, imagine the actual Staunton pattern shapes of pieces, but instead, conceive of them more abstractly, as radiating lines of force, corresponding to the way that they are legally permitted to move.[10]

In effect, chess is a branch of mathematics. Finding a forcing continuation in a position is analogous to finding a geometrical proof. The proof that the angles of a Euclidean triangle add up to 180 degrees is not about an actual diagram of a triangle that happens to be used to illustrate it, but refers to an abstract model (or set of models). In the same way, a thought about what is achievable in a given chess position refers to the abstract model that reflects what is in common to all equivalent chess positions.

It is also important to note that the abstract structure focused upon in a chess game is conceived in entirely relational terms. What is represented is a spatial arrangement, with pieces whose nature is exhausted by formally specified relational possibilities. The moral here is that when we conceive of a chess piece as such, the *intrinsic* character is not relevant. A chess piece can be made of wood, plastic, metal or old bits of kneaded dough. Its shape can vary,

[10] See, for example, the classic work by Binet (1966). More recent studies confirm the finding that blindfold players imagine the chess pieces in terms of the 'lines of force' they exert: see Chabris and Hearst (2003).

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

likewise. An imaginative grasp of the future possibilities open to a piece is essential to our recognition of a piece for what it is. But the possibilities that in this way determine the nature of a piece are purely spatial. Only the potential relations of the piece to other squares and other pieces on the board matter. This idea, that the relational role or function of a piece is what is essential to our understanding of its nature, is one important feature that throws light upon the nature of our perception of everyday objects.

7. Phenomenology and Chess

The fact that our understanding of what is essential to a chess position refers to its formal, or abstract, nature contrasts with a paradoxical feature of chess to which I want to draw attention. Despite their abstract nature, from a phenomenological point of view, chess positions take on a vivid reality. From the good player's perspective, the pieces on the board come alive in a game. They are not treated merely as representations in some abstract formal system. Positions and pieces on the board can take on a meaningful presence, and have a strong emotional resonance. In part, the fact that positions take on significance has to do with practical goals, when a player strives to win the game. But in addition, the combinations that arise in practical play – and also in composed studies – can be assessed disinterestedly, in aesthetic terms, for beauty, elegance, wit and other related higher-level properties.

In his famous book, *The Luzhin Defence* – arguably the greatest novel about the game – Vladimir Nabokov captures some of these aspects of chess through his portrayal of the central figure of Luzhin, the doomed chess master whose thought processes, when playing blindfold chess, are described as follows:

> [Luzhin] saw then neither the knight's carved mane nor the glossy heads of the pawns – but he felt quite clearly that this or that imaginary square was occupied by a definite concentrated force, so he envisioned the movement of a piece as a discharge, a shock, a stroke of lightning – and the whole chess field quivered with tension ….[11]

Nabokov's description ties in nicely with the empirical studies of blindfold chess players of the way they imagine pieces in their mind.

[11] Nabokov (1964) pp. 91–2.

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

In commenting on the final game of his return match with Euwe, Alekhine wrote of the capture of his opponent's piece as 'removing the *most hated knight* of the match', and such a comment is typical of the passionate way that players engage with the game. The chess style of Petrosian was characterised by the way that he would exert a 'python-like' grip on his opponent, depriving him of all counter-play. Nimzowitsch spoke of a passed pawn as 'a criminal, who should be kept under lock and key'; the passed pawn has a 'lust to expand'. In these, and in multiple other ways, the manner in which chess players talk about the game reflects the fact that chess positions are full of meaning for them, figurative meaning that goes far beyond a grasp of formal structure. This feature of the way that we see chess positions has implications for theories about perception and the nature of physical world.

8. Structures and Objects

There is one further aspect of seeing, or perceptual understanding, in chess that also throws light on the nature of perception in our everyday perceptual dealings with the world, outside chess.

As noted, in order for human chess players to avoid having to calculate excessively large numbers of variations, they have to restrict the possible moves they examine in depth. In order to do this, they have to bracket off groupings of pieces on the board that are less central to their main focus. The player needs to treat some groupings of pieces on the board as semi-permanent units, like complex objects, which have no direct bearing on the task in hand. If the immediate task is, for example, to find a direct sacrifice to expose the opponent's king to a possible mate, then the player may ignore pieces on the other side of the board, at least in a temporary way, while exploring the initial sacrificial possibilities.

An example here will illustrate what I mean, and clearly shows the difference between the way computers 'think' or compute in chess, and the way that humans explore the tactical possibilities in positions.

This is from a very interesting endgame study by Mihai Neghina, first posted on the internet in 2009. The position was initially published without a solution, accompanied by the claim that computers were unable to evaluate the position correctly (Diagram 12).

This position is a forced win for White. However, even after some introductory moves, the computer wrongly calculates that Black has a winning advantage:

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

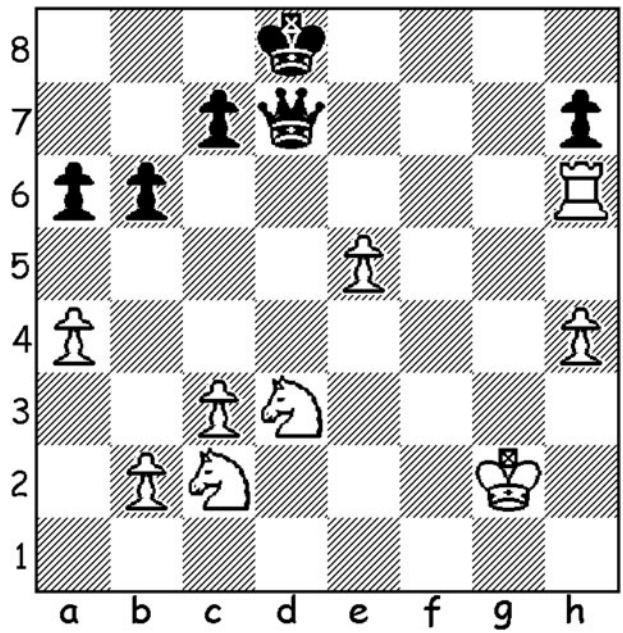


Diagram 12. Neghina Endgame Study position (version). White to play and win.

1 Nd4 Qg7ch
2 Kh3 Qxh6
3 Nf4 Kc8

(Because of the threat of a knight fork on e6, if the knight on f4 is captured)

4 Nde6 Kb7
5 h5

Black has a queen for two knights and a pawn, normally a winning advantage for the side with the queen. The computer's verdict on this position is that Black should be able to force a win. But what the computer cannot do is assess the position in the way that a human can, as involving a piece-complex in which the Black queen will be trapped, and out of play – a key idea that the human player can trade on – and hence work out that it is White who has the advantage.

Diagram 13 gives the key trapping position, reduced to its bare elements.

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

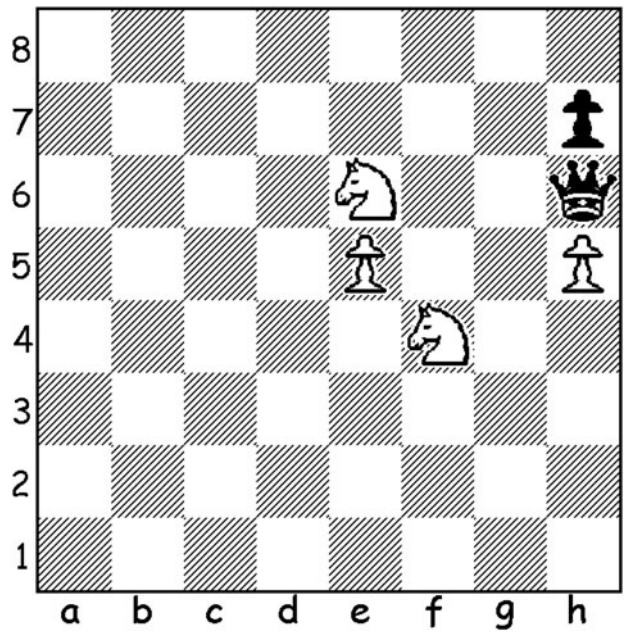


Diagram 13. Neghina Endgame Study: Trapping position.

The Black queen can make 8 moves in all, but they will all lead to her capture. It is very difficult to programme a computer to disregard this complex – at each move the computer will calculate these 8 possibilities, and their further ramifications, including over twenty possible replies by the White knights and the king. The possible variations multiply up to such an extent that they exceed even the computer's capacity for calculating billions of different positions.

The human player thinks differently. He, or she, can conceptualise the array on the right-hand side of the board as a semi-permanent object. The trapping piece-complex is treated as a stable unit. It can be imagined as a separate object, one that does not affect the play.

If the complex is treated in this way, it then becomes easier to ignore it when calculating the main lines, and analysing the best play for both sides. This results in White winning a crucial pawn on the queen side, so as to reach a won ending. The central variation continues: 5…c5 6 Kg4 c4 7 Kf5 Kc6 8 Ke4 b5 9 axb5ch axb5 10 Kd4 Kb6 11 Kd5 Ka5 12 Kc5 Ka6 13 Kc6 Ka5 14 Kb7 b4 15 Kc6 bxc3 16 bxc3 Ka4 17 Kc5 Kb3 Kd4 Kc2 19 Kxc4 and wins. It is only near the

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

final stages, after around 10 or so moves, that the computer begins to assess the position correctly, as a win for White.[12]

There are two morals to be drawn from this example. Firstly, with respect to the differences between the way that humans and computers think, the fact that humans are able set aside piece-complexes, treating them as irrelevant to the immediate focus of play, shows something important about their imaginative resources.[13] Here the imagination works in a negative manner, by restricting the assessment of future possibilities, and by treating the trapping position as a fixed unit, removed from the immediate scene of action. The future possible moves within the piece-complex are, as it were, 'suspended' and removed from consideration by the player, at least on a temporary basis.

Secondly, it is a phenomenologically important fact that the structure which is treated by the player as a relatively stable object is a complex spatial structure. This structure can be interpreted as a high-level property of the chess game, and illustrated by a concrete arrangement of pieces on a board, but its essence is nevertheless something abstract and formal.

9. Implications and Applications

To summarise what I have argued so far:

1) The imagination is exercised in different ways in chess. At the basic level, each piece is understood in terms of its spatial role. Players also exercise the imagination in further ways, in thinking through tactical sequences and assessing strategical possibilities, and in recognising larger structures, piece groupings that have different degrees of permanence and complexity.
2) A player's grasp of a chess position is based upon an implicit or explicit awareness of the potential moves open to each piece: players envisage the different transformations of the position that will result from possible legal moves, generating different spatial arrangements of pieces.
3) In thinking about a given position the player is *guided* by the representation of the position, either through seeing the

[12] For discussion of this endgame study, see Josten (2010).

[13] I am ignoring other differences between humans and machines that relate to the content and interpretation of chess moves, and the like. It is, of course, arguable that this constitutes a further dimension in which humans and computers differ.

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

qualitative aspects of the pieces on the board, or in imagining the positions of pieces when playing blindfold.

4) In a stricter sense, what the chess player grasps is an abstract model, a purely formal structure devoid of any intrinsic qualitative aspects, which is distinct from any concrete representation of the position reached in the game. In playing chess the player is dealing with what is at one level something purely mathematical: an abstract spatial structure governed by formally specifiable rules. Such structures are represented, for greater convenience, by arbitrarily chosen shapes on a board with a certain spatial configuration.

5) Yet the manner in which a good player responds to a chess position extends far beyond a grasp of its abstract spatial character. The arrangements of pieces on the board have phenomenological significance, and carry emotional charge. Although its essential aspects are purely formal, a chess position is seen as being richly meaningful.

The peculiar essence of chess as a game consists in this striking integration of two features of consciousness that are often considered to be starkly opposed: an *understanding* of formal or abstract structures, and an *awareness* of rich meaning, a meaning that transcends the practical goal of achieving a winning position.

These considerations about the way we grasp chess positions have important implications for epistemology. It is not only in playing chess (and similar games) that these two features of consciousness are combined. The structure of visual consciousness that is manifested in the case of seeing chess positions tells us something important about the general nature of perception. How we should analyse what goes on when a person plays chess has relevance to current philosophical disputes concerning the proper way to analyse the perceptual relation. The game of chess provides a model of our perceptual engagement with physical reality. In the concluding section of this paper I shall explain why this is so.

Much recent philosophy of perception has involved increasingly convoluted attempts to defend forms of Direct Realism (otherwise known as Disjunctivism).[14] The central tenet of this view is that experience is transparent, and that, in veridical perception, nothing mediates a subject's experience of the world. The subject is immediately aware of intrinsic properties belonging to the external physical

[14] Disjunctivism in its modern form stems principally from papers by Snowdon (1981) and McDowell (1982).

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

object perceived. An *irreducible experiential relation* connects the external object perceived with the subject's perceptual consciousness.[15]

Direct Realism has a superficial plausibility. It can appear to fit in with some of our common-sense convictions about our perceptual contact with the world. However, it faces two formidable objections. Firstly, there is the causal-scientific argument, developed by Valberg and Robinson. This turns on the point that both hallucinatory experiences, and also veridical experiences, are generated by exactly the same kinds of proximal brain states. Such states are sufficient, as the hallucinatory case indicates, to produce an experience *as of* an external world. The argument offers a good reason for holding that veridical experiences are *inner* states, which, like hallucinatory experiences, supervene solely upon the proximal brain state that is causally linked to the distal object perceived. It therefore follows that in seeing an object, I am immediately aware of the qualitative component of my inner experience – a phenomenal state of the same ontological kind that occurs in hallucination.[16]

The second objection involves a deeper metaphysical worry as to whether the Direct Realist theory is even coherent. What is an 'irreducibly experiential relation'? What kind of relation connects the subject's conscious experience to a mind-independent object? How does it fit into the natural world? It is difficult to spell out the Direct Realist theory in any clear way, so as to explicate the direct non-causal relation that is supposed to hold between the subject and the specific object he, or she, perceives. Without any positive account, there are no grounds to which the Direct Realist can appeal in order to specify which *particular* object is perceptually linked to the subject's experience, except grounds involving facts about causally related events. So Direct Realism becomes parasitic upon the causal account.[17]

The alternative to this view, the Causal Theory of Perception, holds that Direct Realism is misconceived. There is no irreducible *real* relation connecting a subject's conscious experience with any mind-independent physical item, which can make that very item phenomenally present in experience. The phenomenal redness that I am immediately aware of, when for example I stare at a juggling ball, supervenes on my brain state alone, and is an item logically distinct from that ball, ontologically in the same category as the green

[15] The phrase is from Hobson (*forthcoming*).

[16] See Valberg (1992) and Robinson (1994), and for criticism, Martin (2006) and Fish (2009).

[17] This argument is spelled out more fully in Coates (2007), Ch. 4.

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

after-image I will experience a moment later when I look away from the red ball onto a white surface. But to note this point is not to deny the fact that perception is also, in an important sense, direct. As Kant emphasised, there is a further component in experience, which comprises the exercise of low-level classificatory concepts. These concepts arise, without inference, in perceptual takings normally focused *directly* upon external objects and their features (surface or otherwise).[18] An *intentional* relation links the subject to the perceived object.[19] The awareness of phenomenal qualities causally prompts and guides our perceptual takings, which are about the external, mind-independent objects we perceive. The reddish shape I am non-conceptually aware of prompts me to see the object in front of me directly *as* a juggling ball.

The Causal Theory continues to meet with resistance. One objection to it interweaves semantical considerations about the content of experience with phenomenological considerations about the meaningfulness of our experience of objects. The charge is that the Causal Theory is unable to give a satisfactory account of how we come to have any kind of awareness of properties belonging to mind-independent objects, an awareness that forms the basis of our understanding of a rich and complex objective reality.[20] The Causal Theory thus appears to have the consequence that we are cut off from any knowledge of the intrinsic nature of physical things.

Part of the answer to this objection involves spelling out the structural realist conception of physical objects that, from a scientific perspective, is naturally allied with the Causal Theory. According to this conception, the intrinsic properties of objects are ultimately all primary qualities, qualities whose nature is essentially spatial. According to the form of structural realism originally outlined by Russell, the monadic properties that belong to a macroscopic object at a higher level of explanation are ontologically reducible. They

[18] It is of central importance in perception theory to distinguish the *directness* of manner in which the conceptual aspect of experience focuses upon the object, with the *immediacy* of non-conceptual, phenomenal, awareness of qualities that accompanies the conceptual aspect. These need not necessarily relate to the same entity.

[19] Note that to say this is not to endorse an intentionalist view of perceptual experience, since the intentional state is only one component of the overall perceptual experience.

[20] This objection appears to underlie McDowell's frequently expressed worries that a (Cartesian) epistemic divide between the subjectivity of experience and the external world has the consequence 'of putting subjectivity's very possession of an objective environment in question' (1986), 147.

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

can be accounted for in terms of spatial and causal relations between the micro-objects that constitute the higher-level structures.[21] Thus the nature of a certain kind, such as an apple, is explained in terms of its organic constitution, which in turn supervenes upon a certain chemical structure, and so on down to yet lower-level structures. What we understand, when we grasp the objective nature of an object, is something that can, in principle, be accounted for by reference to complexes of spatial structures connecting micro-objects, and the way that they causally interconnect.[22] However, we do not require full knowledge of these lower-level structures in order to have some understanding of the kind to which a given object belongs, so long as we have some knowledge of its causal nature. What we take to be the sensible qualities of objects in the world – the familiar colours, sounds, smells, tastes and feel of objects present to conscious experience – turn out to be states of our own minds, which we project onto the structures that we perceive.

By interpreting our perceptual grasp of the world along these structural realist lines, we can defend the Causal Theory against the charge that it is unable to explain how experiences have empirical content relating to an objective world.[23] Our formal grasp of spatial concepts allows us to form meaningful beliefs about the physical objects that transcend our experiences. Nevertheless, this appeal to our grasp of the essential spatial structure of the physical world does not exhaust the issues raised by the objection. There is a further dimension to the problem. The worry is that, on the structural realist account, we cannot do adequate justice to the phenomenology of our perceptual experience.

This deeper worry arises because, according to the Causal Theory, external physical objects can only be conceived of in austere formal terms, stripped of their intrinsic properties. The objection is that we are left with something that is only, at best, a grasp of a bare structure. Allegedly, this understanding fails to match up to our everyday grasp of the kinds of objects we encounter in perception. On the structural realist view, the notion of a physical object ends up as

[21] See Russell (1927).

[22] It may be that, at the lowest level, there are unknown intrinsic properties, as Langton (1998) and Lewis (2009) argue; alternatively, it has been argued that we do not need to posit monadic properties at the lowest level – structure goes all the way up, and down: see Ladyman (1998) and Esfeld (2003).

[23] See the defence of Russell's view in Maxwell (1970) and (1972), and also Lowe's vindication of Locke's views in his (1995) Ch. 3.

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

thin and attenuated. Yet the world that we engage with in our ordinary activities seems to be the very opposite of this. It is a world that is rich and complex, full of value and meaning. We encounter perceptually a multiplicity of complex natures, higher-level monadic properties, powers and norms.

However, from what has emerged about the way that we see chess positions, it should be clear that there is a good reply available to the Causal Theorist. As the examples from chess indicate, seeing (or otherwise perceiving) anything is always a complex affair. The way that we see chess positions is, to a large extent, indicative of how we see things in general. One component of visual experience – the sensory aspect of perceptual consciousness – involves something qualitative being immediately presented to the subject: an array of shapes and colours.[24] But the critical aspect of experience is the way that concepts enter into it. Experiences have a further component, which has intentional representational content. The patterns of pieces and squares presented in experience, when the player looks at the pieces on the board, need to be distinguished from the objects that are the focus of the chess player's higher-level cognitive attention, in which the pieces are taken to be related to each other according to formal rules that determine their potential moves. As we have noted, what the chess player sees, in the more inclusive sense, are abstract objects. In a strict sense, such objects are grasped in spatial terms, by virtue of the potential relations that the individual units bear to each other.

However, as the example of chess also shows, an essentially spatial grasp of an entity is compatible with that entity having a further level of meaningfulness for us. As has been argued above, chess structures can have a meaningfulness which goes beyond the merely formal. We perceive chess positions as containing relatively stable complexes of pieces, which can be considered, at least temporarily, as larger-scale unified objects. A chess player sees part of the board as a well-defended king's position, or a strong centre, a blocked pawn chain, and so on. We also noted that, despite the abstract formal nature of chess, pieces and piece-complexes are often conceived in value-laden terms, and they can take on a rich emotional significance. Parallel considerations apply to our perception of physical objects in general. In experience we intentionally represent objects around us as having shape, size, position and kind. Such features can be accounted for in terms of structure, resulting in the 'thin' conception

[24] The representational contents of this component can be accounted for in terms of Peacocke's scenario content.

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

of objective content. But in addition, the objective structures we interact with take on a much richer meaning for us. The strict, thin, account of content is compatible with a view of perception according to which experiences of the world have additional layers of significance. As in the example of chess, we directly perceive the physical world as comprising a wide range of different kinds of objects, instantiating rich and varied high-level properties, and permeated by value. The phenomenology of perception is compatible with an implicit understanding, in a stricter sense, of the relational structure of the physical world.

There is one obvious difference between our conception of chess positions and of physical objects, but it does not undermine the analogy I have been drawing. The abstract position that is the focus of my chess thoughts when I play a game of chess is conceived of as self-contained, unrelated to egocentric or allocentric space, and as causally inert. In contrast, the physical objects I see in my surroundings are potentially items I can causally interact with. This does not conflict with the structural view of objects. It means that my understanding of the nature of an object also comprehends an appreciation of the fact that my experiences of that object can undergo change as a result of changes in the object itself, and also as a result of my movements relative to it. An object of any given kind, such as an apple, will have a typical causal profile. In addition to understanding its essentially spatial nature, I have expectations of the different ways that an object will appear in different conditions, and how it will affect me.[25] I know that being hit by a hard cricket ball will cause me pain; and that biting into a ripe apple will cause me to have sensations of sweetness, and so on. This reinforces the conclusion that an essentially structural conception of objects is compatible with the rich experience of the world that is a feature of normal perception.

The conclusion suggested by these brief reflections is this. An analysis of the phenomenon of playing chess reveals something important about the general nature of perceptual experience, and of our relation to the physical world. Our perceptually based understanding of reality is of a world that has more of a structural character than we unreflectively appreciate. Our implicit grasp of physical objects has much in common with the kind of grasp we have of the spatial structures represented in chess positions. Seeing a chess position is in many ways like seeing in the ordinary sense; while conversely, our perceptually based understanding of the physical world has

[25] This point, again, goes back to the Kant's claims about the role of productive imagination in perception.

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

much in common with our grasp of the essential spatial character of chess positions.[26]

University of Hertfordshire
p.coates@herts.ac.uk

References

Binet, A. (1966) 'Mnemonic virtuosity: A study of chess players', *Genetic Psychology Monographs* **74**, 127–162. Translated from the *Revue des Deux Mondes* (1893), 117, 826–859.

Chase, W. & Simon, H. (1973) 'Perception in Chess' *Cognitive Psychology* **4**, 55–81.

Chabris, C. & Hearst, E. (2003) 'Visualization, Pattern Recognition, and Forward Search: Effects of Playing Speed and Sight of Position on Grandmaster Chess Errors' *Cognitive Science*, **27**, 637–648.

Coates, P. (2007) *The Metaphysics of Perception: Wilfrid Sellars, Perceptual Consciousness and Critical Realism*, London: Routledge.

Esfeld, M. (2003) 'Do relations require underlying intrinsic properties? A physical argument for a metaphysics of relations', *Metaphysica. International Journal for Ontology & Metaphysics*, 4, 5–25.

Fish, W. (2009) *Perception, Hallucination and Illusion*, Oxford: Oxford University Press.

Gobet, F. (2012) 'Concepts Without Intuition Lose the Game: Commentary on Montero and Evans 2011' *Phenomenology and the Cognitive Sciences*, 11: 2, 237–250.

Groot de, A. D. (1965) *Thought and Choice in Chess* (1st Ed.), The Hague: Mouton.

Hobson, K. (forthcoming) 'In Defense of Relational Direct Realism', *European Journal of Philosophy*.

Josten, G. (2010) *A Study Apiece*, Homburg: SAAR- Schach-Agentur.

Kasparov, G. (1987) *London-Leningrad Championship Games*, Oxford: Pergamon Press.

Kasparov, G. (2003) *On My Great Predecessors, Part I*, London: Everyman Chess.

[26] For helpful comments on an earlier draft I would like to thank Sam Coleman, Luciano Floridi, Peter Lane, and John Rose.

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press

Lane, P. & Gobet, F. (2011) 'Perception in Chess and Beyond: Commentary on Linhares and Freitas', *New Ideas in Psychology* 29, 156–161.
Ladyman, J. (1998) 'What is structural realism?' *Studies in History and Philosophy of Modern Science*, **29**, 409–424.
Langton, R. (1998) *Kantian Humility*, Oxford: Clarendon Press.
Lewis, D. (2009) 'Ramseyan Humility', in David Braddon-Mitchell & Robert Nola (eds) *Conceptual Analysis and Philosophical Naturalism*, Cambridge, Mass: MIT Press.
Lowe, J. (1995) *Locke on Human Understanding*, London: Routledge.
Martin, M. (2006) 'On Being Alienated' in Gendler, T. and Hawthorne, J. (eds) *Perceptual Experience*, Oxford: Oxford University Press.
Maxwell, G. (1970) 'Theories, Perception and Structural Realism' in R. Colodny (ed.) *The Nature and Function of Scientific Theories*, Pittsburgh: University of Pittsburgh Press.
Maxwell, G. (1972) 'Russell on Perception' in D Pears (ed.) *Bertrand Russell: A Collection of Critical Essays*, New York: Anchor Books.
McDowell, J. H. (1982) 'Criteria, Defeasibility, and Knowledge', *Proceedings of the British Academy*.
McDowell, J. H. (1986) 'Singular Thought and the Extent of Inner Space', in Pettit, P. and McDowell, J. H. (eds) *Subject, Thought, and Context*, Clarendon Press: Oxford.
Nabokov, V. (1964) *The Luzhin Defence*, London: Weidenfeld and Nicolson (With thanks to the Nabakov Estate).
Robinson, H. (1994) *Perception*, London: Routledge.
Russell, B. (1927) *The Analysis of Matter*, New York: Harcourt, Brace.
Sellars, W. (1978) 'The Role of Imagination in Kant's Theory of Experience' in Johnstone, H. W. (ed.) *Categories: A Colloquium*, Pennsylvania: Pennsylvania State University.
Snowdon, P. (1981) 'Perception, Vision and Causation', *Proceedings of the Aristotelian Society* **81**, 175–92.
Valberg, J. J. (1992) *The Puzzle of Experience*, Oxford: Clarendon Press.

https://doi.org/10.1017/S1358246113000258 Published online by Cambridge University Press