CAMBRIDGE
Cambridge University Press
978-1-107-08459-9 — The Neuroscience of Expertise
Merim Bilalić
Excerpt
More Information

CHAPTER ONE

Introduction to Research on Expertise

Learning Objectives

- What is expertise, what are expertise domains, and who are experts?
- How do experts accomplish seemingly impossible feats? What are cognitive mechanisms in expertise? Why is memory crucial for experts' outstanding performance and how is it connected to other cognitive processes such as attention and perception?
- How does the brain accommodate expertise?
- What is similar and what is different in the cognitive mechanisms of expertise and their neural implementations in perceptual, cognitive, and motor domains?
- How can expertise be used to investigate the human mind (and brain)?

1.1 Introduction

Imagine yourself on a tennis court. On the other side of the net is the five-time Wimbledon champion, Serena Williams, preparing to serve – and you are supposed to return her serve. With her serve regularly reaching 120 mph, you face a rather daunting task. The speed with which the tennis ball reaches you simply does not allow you enough time to perceive and react to its trajectory. In other domains that are seemingly based more on brainpower than on speed and physicality, the situation may be no less daunting. In the game of chess, not only are there numerous individual objects on the chessboard, but they are all connected with each other. The game of chess is so complex that some argue there are more possible combinations of moves in chess than there are atoms in the universe (Shannon, 1950). Yet you are expected to find the best solution in an environment in which even the most powerful computer would need an eternity to go through all the possibilities. You may be forgiven if you feel as though you are lost in a jungle, as many beginners do when they start playing chess. But losing a match in tennis or a game in chess is a small worry in comparison with the daily pressure that radiologists have to endure. Studying complex radiological images, they need to find suspicious tissue that is almost impossible for an untrained eye to spot. Missing even a tiny lump in a thorax X-ray may result in deadly complications for the unfortunate patient.

When one considers the complexity of the environment, it is no wonder that the feats performed by experts defy logic. The best tennis players not only return serves on a regular basis, but also launch their counter-attacks at the same time. The best chess players, known as grandmasters, need only a few glances to spot a promising solution, and experienced radiologists require a mere split second to spot an abnormality in an X-ray image. Research on expertise investigates exactly how such seemingly impossible feats come together. On the one hand, it studies how **cognitive processes**, such as perception, attention, and memory, enable experts' outstanding performance and how expertise has been implemented in the brain. On the other, it focuses on individuals and identifies the characteristics and activities necessary for the highest level of performance. I will tackle the widespread assumption that experts possess special abilities not found in mere mortals in the final fifth chapter. Here, in the introductory chapter, I will give an overview of the cognitive processes behind experts' outstanding performance and will illustrate the way experts' brains accommodate this performance.

1.2 Definition of Expertise and Its Domains

It may seem almost trivial to ask for a definition of **expertise**. After all, surely we know an expert when we see one. This might be the case with the best chess and tennis players, as well as with radiologists who save lives regularly. Their performances speak for themselves. However, there are also a number of domains where experts have been designated by general consensus and not on the grounds of their performance. We can assume that the politicians who are elected in local government or national parliaments are seen as experts. After all, they have been chosen by majority vote to tackle important problems in their society. Similarly, people who entrust Wall Street brokers with their money for investment presumably consider them experts in their business. Yet, on more than one occasion, you have probably been stunned by the decisions taken by your chosen representative, and these days it is clear that Wall Street brokers cannot reliably predict the movement of the financial market.

The main difference between radiology, chess, and tennis on the one hand, and politics and the financial world on the other, is the nature of the environments. Pathological elements in radiology rarely change; the rules in tennis and chess are constant. This consistency of environment enables practitioners consciously or unconsciously to acquire knowledge of regularities that will then be used in dealing with new situations. In contrast, politics and the financial market are regulated by numerous unknown factors, which make reliable prediction difficult if not impossible. Practitioners simply cannot acquire relevant knowledge as situations constantly change. Previously acquired knowledge is often of little use. Politicians and stockbrokers may be elected experts by people who trust them, or even by their peers, but their performances are not consistently outstanding enough for them to be considered experts. **Experts** are people who produce clearly above average (outstanding) performances on a regular basis (Ericsson, 2006). An expert performance is not a one-off. It is not something that comes and goes. If you were to wake skilled chess players in the middle of the night and show them a difficult chess puzzle, they would find the solution without much difficulty, just as skilled radiologists would find lesions in radiological images in the same situation. Politicians and stockbrokers would probably need all day and night, and a lot of luck, to get anywhere near that level of performance (for more information about differences between expertise and other domains, see Shanteau, 1992).

Classical expertise domains are stable environments. Changes do happen, such as new diseases, new makes of tennis balls, or new strings on rackets, but they are usually small enough that they do not change the environment radically and render previous knowledge irrelevant. Every expertise domain provides a wealth of consistent information to its practitioners. The co-occurring stable environmental constellations can be acquired and, as we will see later in this chapter, experts find ingenious ways of circumventing their cognitive limitations. Nonetheless, as anyone who has tried his or her hand at sports or games can testify, expertise domains are extremely complex, and mastering them takes years of dedicated practice. There are a lot of things to learn in any expertise domain. It is exactly this knowledge of the particular features of the ever-repeating constellations in a domain that enables experts to see the problems with different eyes from novices. As we will see in the course of this chapter, the reason why experts' strategies are more efficient is not that they execute the individual parts of the strategies more quickly than novices. Their performance is actually based on completely different kinds of strategies, which have been enabled by experts' knowledge of the domain. Novices lack this knowledge and consequently have to rely on rudimentary cognitive strategies.

Some other skills, or at least their components, take much less of our time to acquire. Take, for example, the relatively simple task of quickly rotating your foot, something we will consider later in the chapter. The skill necessary for this task is quickly acquired. The rest of the time is spent on refining the individual steps necessary to produce quicker and quicker performance. In the end, the performance becomes more and more efficient as the execution of the individual components becomes automatic. The simple tasks that enable participants to quickly acquire the skill are typical of the **skill acquisition** approach. The skill acquisition approach is similar to expertise in that it ultimately examines the same thing – skill. The skills, however, are rather simple, since they are designed for acquisition in a reasonable amount of time, unlike the classical expertise domains for which decades of intensive training are often necessary. Despite their differences, skill acquisition and expertise are complementary research streams. Skill acquisition provides insight into the very beginning of the road to excellence, whereas the expertise approach offers an understanding of the processes at the end of the same road. However, there are also marked differences. The strategy in skill acquisition tasks is the same in both skilled and unskilled practitioners. The simplicity of the tasks, or the short duration of practice, prevents participants from coming up with qualitatively different strategies. The experts 'merely' execute the same strategies more quickly. One of the hallmarks of experts' outstanding performance is the use of qualitatively different cognitive strategies based on domain-specific knowledge. These differences between skill acquisition and expertise have also been evident in their neural implementation, as we will see later in the chapter.

Now that we have cleared up the difference between skill acquisition and expertise, let us consider some distinctive expertise domains. In the opening paragraph we introduced some of the typical expertise domains. Tennis, chess, and radiology were not chosen at random: they are all representative of the three expertise domains, which we will examine in the following chapters. Expertise in radiology requires the visual intake of the information needed for the actual task of spotting lesions within radiological images. As such, it will be used as a typical task of **perceptual expertise**, relating to domains that predominantly rely on information from our senses. It is evident that experienced radiologists also need to engage their memory, as without it they would hardly be able to spot and categorize lesions. The task itself, however, is a purely visual search task that does not require the mental permutations we find in chess. Chess players rely on the visual information from the chessboard, but for their outstanding performance, they need to go beyond the available visual information. They have to retrieve previously stored chess constellations that may help them to understand the problem at hand, and then, in one of the main aspects of their expertise, to imagine how the game could proceed. Chess is an example of **cognitive expertise**, where information from our senses plays a secondary role compared to the subsequent engagement of memory and mental simulation. No chess game has been won by just perceiving the situation on the board. Both radiology and chess eventually require motoric responses, either indicating the lesion within a radiological image, or executing a chess move on the board. The motor component in these activities, however, is of no real significance. The essence of sports such as tennis, on the other hand, is exactly the motor component in the performance. Tennis will therefore serve as a prime example of **motor expertise**, relating to domains that are predominantly shaped by motoric responses.

This book deals with all three domains, and one chapter has been devoted to each of the three primary expertise domains. Just as in everyday life, where we perceive the world, make a mental image of it, and then act on that image, the structure of the book corresponds to this fundamental process. After the introductory chapter, which you are reading now, we will deal with radiology and other perceptual expertise domains in Chapter 2. Chapter 3 is dedicated to cognitive expertise, and in it we will see how the brain accommodates the highest levels of expertise in chess and other skills based on memory, such as mental calculations. The next part of the book, Chapter 4, will examine the cognitive and neural mechanisms behind tennis and other motor skills, which depend heavily on the motor component. In the final chapter, I will summarize the recurring themes running through the previous chapters, highlight the importance of expertise for neuroscience in general and discuss what is necessary to become an expert. The division into perceptual, cognitive and motor expertise is rather arbitrary, since all expertise domains, despite their differences, rely on similar if not identical **cognitive mechanisms**; that is, interaction between basic cognitive processes. We will briefly describe them in the next section before we turn to their neural implementation.

1.3 Cognitive Mechanisms in Expertise

How do experts achieve these incredible coups? To understand the way experts' minds have been wired, it helps to take a step back and look at everyday life. Believe it or not, you are an expert too, an everyday expert. It might seem banal from your current perspective, but just remember how many things a small child needs to learn. Unlike you, they cannot enter an unfamiliar room and immediately realize that it is an office, a bedroom, or a living room. You will have no problem in finding a light switch, should the lights suddenly turn off, but a small child would need to learn the position of the light switch first. You have encountered numerous versions of such rooms, you know what kinds of objects one would expect in such rooms and how those objects relate to each other, and you will certainly not look for the light switch on the floor or the ceiling. Children need to develop their ‘room expertise’ through years of exposure to rooms with all their contents and different variations. They will store things that occur together in their memory, even if they do not necessarily realize that they are picking up on such regularities in their environment. With lots of exposure, they will eventually reach your level of ‘room expertise’!

It is not much different with experts. Through years of exposure, experts have acquired knowledge about consistencies in their domain (Chase & Simon, 1973a; Gobet et al., 2001; Gobet & Simon, 1996d). Complex domains, such as radiology, chess, playing an instrument, or sport, obviously take more time to master than our everyday example of rooms. However, all these domains feature ‘rules’ that are stable and situations that arise again and again in one form or another. This knowledge is stored in **long-term memory (LTM)**, the process of material retention that we usually refer to when we talk about memory in everyday life. The name comes from the notion that the information stored here will remain available for retrieval for weeks, months, or even decades. This is in contrast to **short-term memory (STM)** where the content can stay only for several seconds. Once experts encounter a seemingly new situation in their domain, they will automatically activate the domain-specific knowledge stored in LTM for a long time (Richman, Staszewski, & Simon, 1995). The new situation will then be compared with previously encountered situations stored in LTM (Feigenbaum & Simon, 1984). The consequence of this automatic matching of patterns in the outside world and the brain is that experts quickly grasp the essence of the new situation. Their LTM has stored not only similar combinations of details to the one at hand, but also ways of dealing with such situations (Chase & Simon, 1973b). These methods are automatically retrieved and help to focus on the important aspects and ignore the irrelevant ones. Experts, then, do not need extraordinary abilities to comprehend the complex situations they face. Their knowledge enables them to look for the ‘light switch’ in the right place.

1.3.1 Perceptual and Cognitive Expertise

If this example of the light switch seems too abstract, consider the following example. Box 1.1, Figure 2 presents a chest X-ray that displays a potentially deadly disease – pneumonia. Spotting it is not very easy, but experienced radiologists manage to identify such lesions with remarkable success even if the image is presented for a fifth of a second, only slightly longer than the blink of an eye (Kundel & Nodine, 1975). In contrast, medical students who have encountered only a handful of chest X-rays are basically guessing when they perform this task. The task illustrates how a rich knowledge base of visual patterns enables experienced radiologists to quickly figure out what is going on in a problem presented to them. Once experts grasp the gist of the situation, they can immediately focus on the important aspects and ignore the irrelevant ones. Take a look at the image presented in Figure 1.1. It is again an X-ray image containing a lesion, circled in the figure, but this time experienced radiologists and medical students have more than just an eye blink to find the lesion. An eye tracker, a device for recording the direction that the eye is looking in, provides insight into their search strategies. We can see how radiologists do not waste much time and almost immediately focus on the lesion, leaving a large part of the image unexamined. Medical students in contrast, cannot afford to leave unturned any part of the X-ray if they are to spot the lesion. Their eyes cover the whole extent of the image.

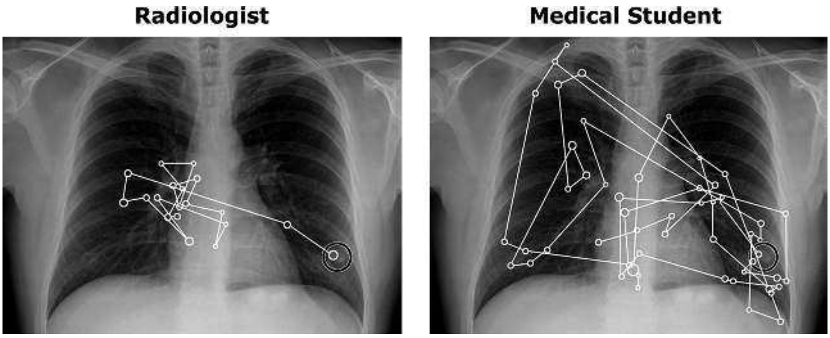

Figure 1.1 Radiological expertise. Experienced radiologists need only a few glances to figure out what is going on in an X-ray (left panel). Consequently, they fix their eyes on the lesion almost immediately, unlike less experienced radiologists, who need to investigate the whole image (right panel). The black circle is the location of the nodule, white circles represent where the eyes fixated, and lines represent the eye movements.

We find the same situation in the seemingly more cognitive domain of chess. Chess positions are made of numerous pieces and pawns (as the chess objects are called) spread across the board. These objects may not make much sense to you, but they form a meaningful unit to experienced chess players. Like the experienced radiologists, they need just a brief glance to figure out what is going on. When chess experts are asked to locate certain kinds of pieces (e.g., knights and bishops) among the other pieces and pawns, they focus on the objects of interest almost immediately, without having to examine the rest of the board. In contrast, novices need to examine the whole board to make sure they have located all the pieces of the specified kind (Bilalić, Langner, Erb, & Grodd, 2010). Chapter 3 on cognitive expertise will deal in great detail with this study (see Figure 3.13 for remarkably similar eye movement patterns to those of radiologists and medical students in Figure 1.1).

The search strategies of expert chess players and experienced radiologists are not only extremely efficient, but also surprisingly similar given how much radiology and chess differ at first sight. The similarity comes from the fact that in both domains LTM enables the fast intake of incoming information by matching it with its content. The matching between incoming sensory information and stored information in LTM is called **pattern recognition**. This pattern recognition process automatically draws information about many other aspects, including possible locations of certain objects connected to the recognized situation. The consequence of this represents the essence of expertise: attention is automatically drawn to important aspects of the situation. In this way an expert can reduce the complexity of the environment and deal with it successfully, despite limited cognitive resources. They are faster and more efficient, but not because they can examine all the aspects of the problem more quickly than novices. They focus their limited resources on the important aspects of the environment, disregarding other less informative elements. Their knowledge enables them to employ qualitatively different strategies from those used by novices. Novices may not have inherently weaker cognitive abilities than experts, but they lack specific knowledge that guides perception, and feel overwhelmed by the complexity of the situation. Their strategies are rudimentary and reflect the lack of domain-specific knowledge.

It is one thing to find a certain piece on the chessboard, and completely another to find a good solution to the problem that all the pieces and pawns on the board together pose. After all, chess players' task is to find good moves, not identify objects! How, then, can experts find the right path in the jungle of possibilities that chess constellations create?

A popular explanation is that they can calculate and foresee numerous moves in advance. Without this extraordinary ability it would be difficult to produce the performances that they do. How could they know if the situation in 10 moves from now will favor them, if they do not mentally simulate those situations in their head, in what is usually called the **mind's eye**? The Dutch psychologist Adrian de Groot (1978/1946) set out to investigate, among other things, this particular question. He devised a task that captures the core of chess expertise – finding the best solutions. Instead of letting players go ahead and play numerous moves, spending several hours on a single game, as they usually do, de Groot devised a laboratory task that is simple enough to be conducted in 15 minutes and yet truly mimics the behavior of chess players during the actual game. He presented chess players with a situation from an unknown tournament game depicted in Figure 1.2, and asked them to find the best move. He also asked them to verbalize their thoughts by the **think aloud technique** while they were looking for the best solution. It was not surprising that some of the world's best chess players, grandmasters, came up with better solutions than their weaker colleagues, whom I will call ordinary experts as they were indeed skilled chess players, but not at the highest level. The real surprise was the structure of the search, which did not differ between the groups: the best experts anticipated scarcely any more moves, as measured by the number of **half-moves**, or **plies**, they considered in advance, than ordinary experts. Both the best and the ordinary chess experts would first categorize the position as belonging to a certain

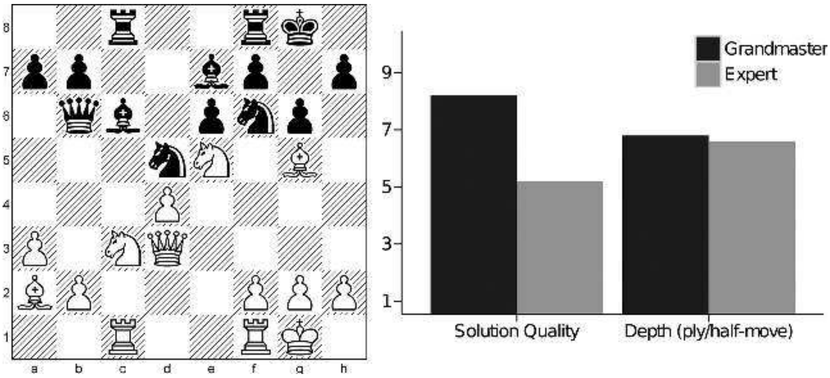

Figure 1.2 Chess expertise. Chess players were given an unknown chess position (left panel) and asked to think aloud while they looked for the best move (for aspiring chess players among readers, 1. Bxd5 wins). The best players (grandmasters) found better solutions, but they did not search more deeply than their weaker colleagues (right panel).

type, and would then on this basis retrieve common plans and possible solutions. The search after the initial phase did not differ, but the solution quality indicates that the initial phase did. The best players could grasp the essence of the position much better than their weaker colleagues. They could focus their analytic search efforts immediately on promising solutions, while weaker players were left investigating irrelevant paths. The pattern of results is reminiscent of the strategies found in the previously described visual searches in radiology and chess. Instead of examining irrelevant aspects of the environment like novices, experts could immediately focus on the informative elements.

One of the main reasons for experts' perceptual advantage is that they process the environment differently. Instead of perceiving individual objects, such as pawns, for example, experts form meaningful units of individual objects, also called **chunks**. In the case of chess, a king who has moved into a corner, as in Figure 1.2, would make a chunk together with the neighboring rook and the pawns. These chunks have been acquired through experiencing common occurrences of the objects in the environment, and have been stored in experts' LTM. They present the content of experts' memory, also called **knowledge structures**, which become more elaborate as experts gain more experience. The best experts have such sophisticated knowledge structures that they can grasp the essence of a complex situation within seconds. Chapter 3 will expand on the nature of the perceptual advantage of experts in the initial phase (see Figures 3.12 and 3.13).

The short historical account of the research on expertise demonstrates that domain-specific knowledge provides the core of the outstanding performance of experts. The acquired knowledge structures in LTM not only enable experts to orient themselves quickly in a new situation through clever guidance of attention to important aspects, but they also automatically provide good ways of dealing with the new situation. This also means that experts will always have a preconceived way of dealing with almost any situation relating to their expertise domain. Could this inseparable link between memory, attention, and problem solving make experts inflexible and blind to new alternatives? The study presented in Box 1.1 uncovers the cognitive mechanisms behind such a phenomenon.

Box 1.1 The Curse of Expertise, or Why Do Good Thoughts Block Better Ones?

We have seen that good ideas come easily to experts, almost immediately upon seeing the problem. What happens if the first idea that comes to experts is not the best one? Can experts get rid of their initial thought and