1

The Cognitive Sciences: One or Many?

1.0 Chapter Overview

When experimental psychology arose in the nineteenth century, it was a unified discipline. However, as the experimental method began to be applied to a larger and larger range of psychological phenomena, this new discipline fragmented, causing what became known in the 1920s as the "crisis in psychology," a crisis that has persisted to the present day.

Cognitive science arose in the 1950s when it became apparent that a number of different disciplines, including psychology, computer science, linguistics and philosophy, were fragmenting. Some researchers responded to this situation by viewing cognition as a form of information processing. In the 1950s, the only plausible notion of information processing was the kind that was performed by a recent invention, the digital computer. This singular notion of information processing permitted cognitive science to emerge as a highly unified discipline.

A half century of research in cognitive science, though, has been informed by alternative conceptions of both information processing and cognition. As a result, the possibility has emerged that cognitive science itself is fragmenting. The purpose of this first chapter is to note the existence of three main approaches within the discipline: classical cognitive science, connectionist cognitive science, and embodied cognitive science. The existence of these different approaches leads to obvious questions: What are the core assumptions of these three different schools of thought? What are the relationships between these different sets of core assumptions? Is there only one cognitive science, or are there many different cognitive sciences? Chapter 1 sets the stage for asking such questions; the remainder of the book explores possible answers to them.

1.1 A Fragmented Psychology

Modern experimental psychology is rooted in two seminal publications from the second half of the nineteenth century (Schultz & Schultz, 2008), Fechner's (1966) *Elements of Psychophysics,* originally published in 1860, and Wundt's *Principles of Physiological Psychology,* originally published in 1873 (Wundt & Titchener, 1904). Of these two authors, it is Wundt who is viewed as the founder of psychology, because he established the first experimental psychology laboratory—his Institute of Experimental Psychology—in Leipzig in 1879, as well as the first journal devoted to experimental psychology, *Philosophical Studies*, in 1881 (Leahey, 1987).

Fechner's and Wundt's use of experimental methods to study psychological phenomena produced a broad, unified science.

> This general significance of the experimental method is being more and more widely recognized in current psychological investigation; and the definition of experimental psychology has been correspondingly extended beyond its original limits. We now understand by 'experimental psychology' not simply those portions of psychology which are directly accessible to experimentation, but the whole of individual psychology. (Wundt & Titchner, 1904, p. 8)

However, not long after its birth, modern psychology began to fragment into competing schools of thought. The Würzberg school of psychology, founded in 1896 by Oswald Külpe, a former student of Wundt's, challenged Wundt's views on the scope of psychology (Schultz & Schultz, 2008). The writings of the functionalist school being established in North America were critical of Wundt's structuralism (James, 1890a, 1890b). Soon, behaviourism arose as a reaction against both structuralism and functionalism (Watson, 1913).

Psychology's fragmentation soon began to be discussed in the literature, starting with Bühler's 1927 "crisis in psychology" (Stam, 2004), and continuing to the present day (Bower, 1993; Driver-Linn, 2003; Gilbert, 2002; Koch, 1959, 1969, 1976, 1981, 1993; Lee, 1994; Stam, 2004; Valsiner, 2006; Walsh-Bowers, 2009). For one prominent critic of psychology's claim to scientific status,

> psychology is misconceived when seen as a coherent science or as any kind of coherent discipline devoted to the empirical study of human beings. Psychology, in my view, is not a single discipline but a collection of studies of varied cast, some few of which may qualify as science, whereas most do not. (Koch, 1993, p. 902)

The fragmentation of psychology is only made more apparent by repeated attempts to find new approaches to unify the field, or by rebuttals against claims of disunity (Drob, 2003; Goertzen, 2008; Henriques, 2004; Katzko, 2002; Richardson, 2000; Smythe & McKenzie, 2010; Teo, 2010; Valsiner, 2006; Walsh-Bowers, 2009; Watanabe, 2010; Zittoun, Gillespie, & Cornish, 2009).

The breadth of topics being studied by any single psychology department is staggering; psychology correspondingly uses an incredible diversity of methodologies. It is not surprising that Leahey (1987, p. 3) called psychology a "large, sprawling, confusing human undertaking." Because of its diversity, it is likely that psychology is fated to be enormously fragmented, at best existing as a pluralistic discipline (Teo, 2010; Watanabe, 2010).

If this is true of psychology, then what can be expected of a more recent discipline, cognitive science? Cognitive science would seem likely to be even more fragmented than psychology, because it involves not only psychology but also many other disciplines. For instance, the website of the Cognitive Science Society states that the Society,

> brings together researchers from many fields that hold a common goal: understanding the nature of the human mind. The Society promotes scientific interchange among researchers in disciplines comprising the field of Cognitive Science, including Artificial Intelligence, Linguistics, Anthropology, Psychology, Neuroscience, Philosophy, and Education. (Cognitive Science Society, 2013)

The names of all of these disciplines are proudly placed around the perimeter of the Society's logo.

When cognitive science appeared in the late 1950s, it seemed to be far more unified than psychology. Given that cognitive science draws from so many different disciplines, how is this possible?

1.2 A Unified Cognitive Science

When psychology originated, the promise of a new, unified science was fuelled by the view that a coherent object of enquiry (conscious experience) could be studied using a cohesive paradigm (the experimental method). Wundt defined psychological inquiry as "the investigation of conscious processes in the modes of connexion peculiar to them" (Wundt & Titchner, 1904, p. 2). His belief was that using the experimental method would "accomplish a reform in psychological investigation comparable with the revolution brought about in the natural sciences." As experimental psychology evolved the content areas that it studied became markedly differentiated, leading to a proliferation of methodologies. The fragmentation of psychology was a natural consequence.

Cognitive science arose as a discipline in the mid-twentieth century (Boden, 2006; Gardner, 1984; Miller, 2003), and at the outset seemed more unified than psychology. In spite of the diversity of talks presented at the "Special Interest Group in Information Theory" at MIT in 1956, cognitive psychologist George Miller,

> left the symposium with a conviction, more intuitive than rational, that experimental psychology, theoretical linguistics, and the computer simulation of cognitive processes were all pieces from a larger whole and that the future would see a progressive elaboration and coordination of their shared concerns. (Miller, 2003, p. 143)

The cohesiveness of cognitive science was, perhaps, a natural consequence of its intellectual antecedents. A key inspiration to cognitive science was the digital computer; we see in Chapter 2 that the invention of the computer was the result of the unification of ideas from the diverse fields of philosophy, mathematics, and electrical engineering.

Similarly, the immediate parent of cognitive science was the field known as cybernetics (Ashby, 1956; de Latil, 1956; Wiener, 1948). Cybernetics aimed to study adaptive behaviour of intelligent agents by employing the notions of feedback and information theory. Its pioneers were polymaths. Not only did cyberneticist William Grey Walter pioneer the use of EEG in neurology (Cooper, 1977), he also invented the world's first autonomous robots (Bladin, 2006; Hayward, 2001; Holland, 2003a; Sharkey & Sharkey, 2009). Cybernetics creator Norbert Wiener organized the Macy Conferences (Conway & Siegelman, 2005), which were gatherings of mathematicians, computer scientists, psychologists, psychiatrists, anthropologists, and neuroscientists, who together aimed to determine the general workings of the human mind. The Macy Conferences were the forerunners of the interdisciplinary symposia that inspired cognitive scientists such as George Miller.

What possible glue could unite the diversity of individuals involved first in cybernetics, and later in cognitive science? One answer is that cognitive scientists are united in sharing a key foundational assumption that cognition is information processing (Dawson, 1998). As a result, a critical feature of cognition involves representation or symbolism (Craik, 1943). The early cognitive scientists,

> realized that the integration of parts of several disciplines was possible and desirable, because each of these disciplines had research problems that could be addressed by designing 'symbolisms.' Cognitive science is the result of striving towards this integration. (Dawson, 1998, p. 5)

Assuming that cognition is information processing provides a unifying principle, but also demands methodological pluralism. Cognitive science accounts for human cognition by invoking an information processing explanation. However, information processors themselves require explanatory accounts framed at very different levels of analysis (Marr, 1982; Pylyshyn, 1984). Each level of analysis involves asking qualitatively different kinds of questions, and also involves using dramatically different methodologies to answer them.

Marr (1982) proposed that information processors require explanations at the computational, algorithmic, and implementational levels. At the computational level, formal proofs are used to determine what information processing problem is being solved. At the algorithmic level, experimental observations and computer simulations are used to determine the particular information processing steps that are being used to solve the information processing problem. At the implementational level, biological or physical methods are used to determine the mechanistic principles that actually instantiate the information processing steps. In addition, a complete explanation of an information processor requires establishing links between these different levels of analysis.

An approach like Marr's is a mandatory consequence of assuming that cognition is information processing (Dawson, 1998). It also makes cognitive science particularly alluring. This is because cognitive scientists are aware not only that a variety of methodologies are required to explain information processing, but also that researchers from a diversity of areas can be united by the goal of seeking such an explanation.

As a result, definitions of cognitive science usually emphasize co-operation across disciplines (Simon, 1980). Cognitive science is "a recognition of a fundamental set of common concerns shared by the disciplines of psychology, computer science, linguistics, economics, epistemology, and the social sciences generally" (Simon, 1980, p. 33). Interviews with eminent cognitive scientists reinforce this theme of interdisciplinary harmony and unity (Baumgartner & Payr, 1995). Indeed, it would appear that cognitive scientists deem it essential to acquire methodologies from more than one discipline.

For instance, philosopher Patricia Churchland learned about neuroscience at the University of Manitoba Medical School by "doing experiments and dissections and observing human patients with brain damage in neurology rounds" (Baumgartner & Payr, 1995, p. 22). Philosopher Daniel Dennett improved his computer literacy by participating in a year-long working group that included two philosophers and four AI researchers. AI researcher Terry Winograd studied linguistics in London before he went to MIT to study computer science. Psychologist David Rumelhart observed that cognitive science has "a collection of methods that have been developed, some uniquely in cognitive science, but some in related disciplines. . . . It is clear that we have to learn to appreciate one another's approaches and understand where our own are weak" (Baumgartner & Payr, 1995, p. 196).

At the same time, as it has matured since its birth in the late 1950s, concerns about cognitive science's unity have also arisen. Philosopher John Searle stated, "I am not sure whether there is such a thing as cognitive science" (Baumgartner & Payr, 1995,

p. 203). Philosopher John Haugeland claimed that "philosophy belongs in cognitive science only because the 'cognitive sciences' have not got their act together yet" (p. 103). AI pioneer Herbert Simon described *cognitive science* as a label "for the fact that there is a lot of conversation across disciplines" (p. 234). For Simon, "cognitive science is the place where they meet. It does not matter whether it is a discipline. It is not really a discipline—yet."

In modern cognitive science there exist intense disagreements about what the assumption "cognition is information processing" really means. From one perspective, modern cognitive science is fragmenting into different schools of thought—classical, connectionist, embodied—that have dramatically different views about what the term *information processing* means. Classical cognitive science interprets this term as meaning rule-governed symbol manipulations of the same type performed by a digital computer. The putative fragmentation of cognitive science begins when this assumption is challenged. John Searle declared, "I think that cognitive science suffers from its obsession with the computer metaphor" (Baumgartner & Payr, 1995, p. 204). Philosopher Paul Churchland declared, "we need to get away from the idea that we are going to achieve Artificial Intelligence by writing clever programs" (p. 37).

Different interpretations of *information processing* produce variations of cognitive science that give the strong sense of being mutually incompatible. One purpose of this book is to explore the notion of information processing at the foundation of each of these varieties. A second is to examine whether these notions can be unified.

1.3 Cognitive Science or the Cognitive Sciences?

One reason that Wilhelm Wundt is seen as the founder of psychology is because he established its first academic foothold at the University of Leipzig. Wundt created the first experimental psychology laboratory there in 1879. Psychology was officially part of the university calendar by 1885. Today, hundreds of psychology departments exist at universities around the world.

Psychology is clearly healthy as an academic discipline. However, its status as a science is less clear. Sigmund Koch, a noted critic of psychology (Koch, 1959, 1969, 1976, 1981, 1993), argued in favor of replacing the term *psychology* with *the psychological studies* because of his view that it was impossible for psychology to exist as a coherent discipline.

Although it is much younger than psychology, cognitive science has certainly matured into a viable *academic* discipline. In the fall of 2010, the website for the Cognitive Science Society listed 77 universities around the world that offered cognitive science as a program of study. Recent developments in cognitive science, though, have raised questions about its *scientific* coherence. To parallel Koch, should we examine "cognitive science," or is it more appropriate to inquire about "the cognitive sciences"? Investigating this issue is one theme of the current book.

According to psychologist George Miller (2003), cognitive science was born on September 11, 1956. At this early stage, the unity of cognitive science was not really an issue. Digital computers were a relatively recent invention (Goldstine, 1993; Lavington, 1980; Williams, 1997; Zuse, 1993). At the time, they presented a unified notion of information processing to be adopted by cognitive science. Digital computers were automatic symbol manipulators (Haugeland, 1985): they were machines that manipulated symbolic representations by applying well-defined rules; they brought symbolic logic to mechanized life. Even though some researchers had already noted that the brain may not work exactly like a computer, the brain was still assumed to be digital, because the all-or-none generation of an action potential was interpreted as being equivalent to assigning a truth value in a Boolean logic (McCulloch & Pitts, 1943; von Neumann, 1958).

Classical cognitive science, which is the topic of Chapter 3, was the first school of thought in cognitive science and continues to dominate the field to this day. It exploited the technology of the day by interpreting "information processing" as meaning "rule-governed manipulation of symbol" (Feigenbaum & Feldman, 1995). This version of the information processing hypothesis bore early fruit, producing major advances in the understanding of language (Chomsky, 1957, 1959b, 1965) and of human problem solving (Newell, Shaw, & Simon, 1958; Newell & Simon, 1961, 1972). Later successes with this approach led to the proliferation of "thinking artifacts": computer programs called expert systems (Feigenbaum & McCorduck, 1983; Kurzweil, 1990). Some researchers have claimed that the classical approach is capable of providing a unified theory of thought (Anderson, 1983; Anderson et al., 2004; Newell, 1990).

The successes of the classical approach were in the realm of well-posed problems, such problems being those with unambiguously defined states of knowledge and goal states, not to mention explicitly defined operations for converting one state of knowledge into another. If a problem is well posed, then its solution can be described as a search through a problem space, and a computer can be programmed to perform this search (Newell & Simon, 1972). However, this emphasis led to growing criticisms of the classical approach. One general issue was whether human cognition went far beyond what could be captured just in terms of solving well-posed problems (Dreyfus, 1992; Searle, 1980; Weizenbaum, 1976).

Indeed, the classical approach was adept at producing computer simulations of game playing and problem solving, but was not achieving tremendous success in such fields as speech recognition, language translation, or computer vision. "An overall pattern had begun to take shape. . . . an early, dramatic success based on the easy performance of simple tasks, or low-quality work on complex tasks, and then diminishing returns, disenchantment, and, in some cases, pessimism" (Dreyfus, 1992, p. 99).

Many abilities that humans are expert at without training, such as speaking, seeing, and walking, seemed to be beyond the grasp of classical cognitive science. These abilities involve dealing with ill-posed problems. An ill-posed problem is deeply ambiguous, has poorly defined knowledge states and goal states, and involves poorly defined operations for manipulating knowledge. As a result, it is not well suited to classical analysis, because a problem space cannot be defined for an ill-posed problem. This suggests that the digital computer provides a poor definition of the kind of information processing performed by humans. "In our view people are smarter than today's computers because the brain employs a basic computational architecture that is more suited to deal with a central aspect of the natural information processing tasks that people are so good at" (Rumelhart & McClelland, 1986c, p. 3).

Connectionist cognitive science reacted against classical cognitive science by proposing a cognitive architecture that is qualitatively different from that inspired by the digital computer metaphor (Bechtel & Abrahamsen, 2002; Churchland, Koch, & Sejnowski, 1990; Churchland & Sejnowski, 1992; Clark, 1989, 1993; Horgan & Tienson, 1996; Quinlan, 1991). Connectionists argued that the problem with the classical notion of information processing was that it ignored the fundamental properties of the brain. Connectionism cast itself as a neuronally inspired, biologically plausible alternative to classical cognitive science (Bechtel & Abrahamsen, 2002; McClelland & Rumelhart, 1986; Rumelhart & McClelland, 1986c). "No serious study of mind (including philosophical ones) can, I believe, be conducted in the kind of biological vacuum to which [classical] cognitive scientists have become accustomed" (Clark, 1989, p. 61).

The architecture proposed by connectionism was the artificial neural network (Caudill & Butler, 1992a, 1992b; Dawson, 2004, 2005; De Wilde, 1997; Muller & Reinhardt, 1990; Rojas, 1996). An artificial neural network is a system of simple processors, analogous to neurons, which operate in parallel and send signals to one another via weighted connections that are analogous to synapses. Signals detected by input processors are converted into a response that is represented as activity in a set of output processors. Connection weights determine the input-output relationship mediated by a network, but they are not programmed. Instead, a learning rule is used to modify the weights. Artificial neural networks learn from example.

Artificial neural networks negate many of the fundamental properties of the digital computer (von Neumann, 1958). Gone was the notion that the brain was a digital symbol manipulator governed by a serial central controller. In its place, the processes of the brain were described as subsymbolic and parallel (Smolensky, 1988); control of these processes was decentralized. Gone was the classical distinction between structure and process, in which a distinct set of explicit rules manipulated discrete symbols stored in a separate memory. In its place, the brain was viewed as a distributed system in which problem solutions emerged from the parallel activity of a large number of simple processors: a network was both structure and process, and networks both stored and modified information at the same time (Hillis, 1985). Gone was the assumption that information processing was akin to doing logic (Oaksford & Chater, 1991). In its place, connectionists viewed the brain as a dynamic, statistical pattern recognizer (Churchland & Sejnowski, 1989; Grossberg, 1980; Smolensky, 1988).

With all such changes, though, connectionism still concerned itself with cognition as information processing—but of a different kind: "These dissimilarities do not imply that brains are not computers, but only that *brains are not serial digital computers*" (Churchland, Koch, & Sejnowski, 1990, p. 48, italics original).

Connectionist models of cognition have had as long a history as have classical simulations (Dawson, 2004; Medler, 1998). McCulloch and Pitts described powerful neural network models in the 1940s (McCulloch, 1988a), and Rosenblatt's (1958, 1962) perceptrons were simple artificial neural networks that were not programmed, but instead learned from example. Such research waned in the late 1960s as the result of proofs about the limitations of simple artificial neural networks (Minsky & Papert, 1988; Papert, 1988).

However, the limitations of early networks were overcome in the mid-1980s, by which time new techniques had been discovered that permitted much more powerful networks to learn from examples (Ackley, Hinton, & Sejnowski, 1985; Rumelhart, Hinton, & Williams, 1986b). Because of these new techniques, modern connectionism has achieved nearly equal status to classical cognitive science. Artificial neural networks have been used to model a wide range of ill-posed problems, have generated many expert systems, and have successfully simulated domains once thought to be exclusive to the classical approach (Bechtel & Abrahamsen, 2002; Carpenter & Grossberg, 1992; Enquist & Ghirlanda, 2005; Gallant, 1993; Gluck & Myers, 2001; Grossberg, 1988; Kasabov, 1996; Pao, 1989; Ripley, 1996; Schmajuk, 1997; Wechsler, 1992).

In a review of a book on neural networks, Hanson and Olson (1991, p. 332) claimed that "the neural network revolution has happened. We are living in the aftermath." This revolution, as is the case with most, has been messy and acrimonious, markedly departing from the sense of unity that cognitive science conveyed at the time of its birth. A serious and angry debate about the merits of classical versus connectionist cognitive science rages in the literature.

On the one hand, classical cognitive scientists view the rise of connectionism as being a rebirth of the associationist and behaviourist psychologies that cognitivism had successfully replaced. Because connectionism eschewed rules and symbols, classicists argued that it was not powerful enough to account for the regularities of thought and language (Fodor & McLaughlin, 1990; Fodor & Pylyshyn, 1988; Pinker, 2002; Pinker & Prince, 1988). "The problem with connectionist models is that all the reasons for thinking that they might be true are reasons for thinking that they couldn't be *psychology*" (Fodor & Pylyshyn, 1988, p. 66). A *Scientific American* news story on a connectionist expert system included Pylyshyn's comparison of connectionism to voodoo: "'People are fascinated by the prospect of getting intelligence by mysterious Frankenstein-like means—by voodoo! And there have been few attempts to do this as successful as neural nets" (Stix, 1994, p. 44). The difficulty with interpreting the internal structure of connectionist networks has been used to argue against their ability to provide models, theories, or even demonstrations to cognitive science (McCloskey, 1991).

On the other hand, and not surprisingly, connectionist researchers have replied in kind. Some of these responses have been arguments about problems that are intrinsic to the classical architecture (e.g., slow, brittle models) combined with claims that the connectionist architecture offers solutions to these problems (Feldman & Ballard, 1982; Rumelhart & McClelland, 1986c). Others have argued that classical models have failed to provide an adequate account of experimental studies of human cognition (Oaksford, Chater, & Stenning, 1990). Connectionist practitioners have gone as far as to claim that they have provided a paradigm shift for cognitive science (Schneider, 1987).

Accompanying claims for a paradigm shift is the view that connectionist cognitive science is in a position to replace an old, tired, and failed classical approach. Searle (1992, p. 247), in a defense of connectionism, has described traditional cognitivist models as being "obviously false or incoherent." Some would claim that classical cognitive science doesn't study the right phenomena. "The idea that human activity is determined by rules is not very plausible when one considers that most of what we do is not naturally thought of as problem solving" (Horgan & Tienson, 1996, p. 31). Paul Churchland noted that "good old-fashioned artificial intelligence was a failure. The contribution of standard architectures and standard programming artificial intelligence was a disappointment" (Baumgartner & Payr, 1995, p. 36). Churchland went on to argue that this disappointment will be reversed with the adoption of more brain-like architectures.

Clearly, the rise of connectionism represents a fragmentation of cognitive science. This fragmentation is heightened by the fact that connectionists themselves freely admit that there are different notions about information processing that fall under the connectionist umbrella (Horgan & Tienson, 1996; Rumelhart & McClelland, 1986c). "It is not clear that anything has appeared that could be called a, let alone the, connectionist conception of cognition" (Horgan & Tienson, 1996, p. 3).

If the only division within cognitive science was between classical and connectionist schools of thought, then the possibility of a unified cognitive science still exists.

Some researchers have attempted to show that these two approaches can be related (Dawson, 1998; Smolensky & Legendre, 2006), in spite of the differences that have been alluded to in the preceding paragraphs. However, the hope for a unified cognitive science is further challenged by the realization that a third school of thought has emerged that represents a reaction to both classical and connectionist cognitive science.

This third school of thought is embodied cognitive science (Chemero, 2009; Clancey, 1997; Clark, 1997; Dawson, Dupuis, & Wilson, 2010; Robbins & Aydede, 2009; Shapiro, 2011). Connectionist cognitive science arose because it felt that classical cognitive science did not pay sufficient attention to a particular part of the body, the brain. Embodied cognitive science critiques both classical and connectionist approaches because both ignore the whole body and its interaction with the world. Radical versions of embodied cognitive science aim to dispense with mental representations completely, and argue that the mind extends outside the brain, into the body and the world (Agre, 1997; Chemero, 2009; Clancey, 1997; Clark, 2008; Clark & Chalmers, 1998; Noë, 2009; Varela, Thompson, & Rosch, 1991; Wilson, 2004).

A key characteristic of embodied cognitive science is that it abandons methodological solipsism (Wilson, 2004). According to methodological solipsism (Fodor, 1980), representational states are individuated only in terms of their relations to other representational states. Relations of the states to the external world—the agent's environment—are not considered. "Methodological solipsism in psychology is the view that psychological states should be construed without reference to anything beyond the boundary of the individual who has those states" (Wilson, 2004, p. 77).

Methodological solipsism is reflected in the sense-think-act cycle that characterizes both classical and connectionist cognitive science (Pfeifer & Scheier, 1999). The sense-think-act cycle defines what is also known as the classical sandwich (Hurley, 2001), in which there is no direct contact between sensing and acting. Instead, thinking—or representations—is the "filling" of the sandwich, with the primary task of planning action on the basis of sensed data. Both classical and connectionist cognitive science adopt the sense-think-act cycle because both have representations standing between perceptual inputs and behavioural outputs. "Representation is an activity that individuals perform in extracting and deploying information that is used in their further actions" (Wilson, 2004, p. 183).

Embodied cognitive science replaces the sense-think-act cycle with sense-act processing (Brooks, 1991, 1999; Clark, 1997, 1999, 2003; Hutchins, 1995; Pfeifer & Scheier, 1999). According to this alternative view, there are direct links between sensing and acting. The purpose of the mind is not to plan action, but is instead to coordinate sense-act relations. "Models of the world simply get in the way. It turns out to be better to use the world as its own model" (Brooks, 1991, p. 139). Embodied cognitive science views the brain as a controller, not as a planner. "The realization was that the so-called central systems of intelligence—or core AI as it has been referred to more recently—was perhaps an unnecessary illusion, and that all the power of intelligence arose from the coupling of perception and actuation systems" (Brooks, 1999, p. viii).

In replacing the sense-think-act cycle with the sense-act cycle, embodied cognitive science distances itself from classical and connectionist cognitive science. This is because sense-act processing abandons planning in particular and the use of representations in general. Brooks (1999, p. 170) wrote: "In particular I have advocated situatedness, embodiment, and highly reactive architectures with no reasoning systems, no manipulable representations, no symbols, and totally decentralized computation." Other theorists make stronger versions of this claim: "I hereby define radical embodied cognitive science as the scientific study of perception, cognition, and action as necessarily embodied phenomena, using explanatory tools that do not posit mental representations" (Chemero, 2009, p. 29).

The focus on sense-act processing leads directly to the importance of embodiment. Embodied cognitive science borrows a key idea from cybernetics: that agents are adaptively linked to their environment (Ashby, 1956; Wiener, 1948). This adaptive link is a source of feedback: an animal's actions on the world can change the world, which in turn will affect later actions. Embodied cognitive science also leans heavily on Gibson's (1966, 1979) theory of direct perception. In particular, the adaptive link between an animal and its world is affected by the physical form of the animal—its embodiment. "It is often neglected that the words *animal* and *environment* make an inseparable pair" (Gibson, 1979, p. 8). Gibson proposed that sensing agents "picked up" properties that indicated potential actions that could be taken on the world. Again, the definition of such affordances requires taking the agent's form into account.

Embodied cognitive science also distances itself from both classical and connectionist cognitive science by proposing the extended mind hypothesis (Clark, 1997, 1999, 2003, 2008; Wilson, 2004, 2005). According to the extended mind hypothesis, the mind is not separated from the world by the skull. Instead, the boundary between the mind and the world is blurred, or has disappeared. A consequence of the extended mind is cognitive scaffolding, where the abilities of "classical" cognition are enhanced by using the external world as support. A simple example of this is extending memory by using external aids, such as notepads. However, full-blown information processing can be placed into the world if appropriate artifacts are used. Hutchins (1995) provided many examples of navigational tools that externalize computation. "It seems that much of the computation was done by the tool, or by its designer. The person somehow could succeed by doing less because the tool did more" (p. 151).

Embodied cognitive science provides another fault line in a fragmenting cognitive science. With notions like the extended mind, the emphasis on action, and the abandonment of representation, it is not clear at first glance whether embodied cognitive science is redefining the notion of information processing or abandoning it altogether. "By failing to understand the source of the computational power in our interactions with simple 'unintelligent' physical devices, we position ourselves well to squander opportunities with so-called intelligent computers" (Hutchins, 1995, p. 171).

Further fragmentation is found within the embodied cognition camp (Robbins & Aydede, 2009; Shapiro, 2011). Embodied cognitive scientists have strong disagreements amongst themselves about the degree to which each of their radical views is to be accepted. For instance, Clark (1997) believed there is room for representation in embodied cognitive science, while Chemero (2009) did not.

In summary, early developments in computer science led to a unitary notion of information processing. When information processing was adopted as a hypothesis about cognition in the 1950s, the result was a unified cognitive science. However, a half century of developments in cognitive science has led to a growing fragmentation of the field. Disagreements about the nature of representations, and even about their necessity, have spawned three strong camps within cognitive science: classical, connectionist, and embodied. Fragmentation within each of these camps can easily be found. Given this situation, it might seem foolish to ask whether there exist any central ideas that can be used to unify cognitive science. However, the asking of that question is an important thread that runs through the current book.

1.4 Cognitive Science: Pre-paradigmatic?

In the short story *The Library of Babel*, Jorge Luis Borges (1962) envisioned the universe as the Library, an infinite set of hexagonal rooms linked together by a spiral staircase. Each room held exactly the same number of books, each book being exactly 410 pages long, all printed in an identical format. The librarians hypothesize that the Library holds all possible books, that is, all possible arrangements of a finite set of orthographic symbols. They believe that "the Library is total and that its shelves register . . . all that is given to express, in all languages" (p. 54).

Borges' librarians spend their lives sorting through mostly unintelligible volumes, seeking those books that explain "humanity's basic mysteries" (Borges, 1962, p. 55). Central to this search is the faith that there exists a language in which to express these answers. "It is verisimilar that these grave mysteries could be explained in words: if the language of philosophers is not sufficient, the multiform Library will have produced the unprecedented language required, with its vocabularies and grammars" (p. 55).

The fictional quest of Borges' librarians mirrors an actual search for ancient texts. Scholasticism was dedicated to reviving ancient wisdom. It was spawned in the tenth century when Greek texts preserved and translated by Islamic scholars made their way to Europe and led to the creation of European universities. It reached its peak in the thirteenth century with Albertus Magnus' and Thomas Aquinas' works on Aristotelian philosophy. A second wave of scholasticism in the fifteenth century was fuelled by new discoveries of ancient texts (Debus, 1978). "The search for new classical texts was intense in the fifteenth century, and each new discovery was hailed as a major achievement" (Debus, 1978, p. 4). These discoveries included Ptolemy's *Geography* and the only copy of Lucretius' *De rerum natura*, which later revived interest in atomism.

Borges' (1962) emphasis on language is also mirrored in the scholastic search for the wisdom of the ancients. The continued discovery of ancient texts led to the Greek revival in the fifteenth century (Debus, 1978), which enabled this treasure trove of texts to be translated into Latin. In the development of modern science, Borges' "unprecedented language" was first Greek and then Latin.

The departure from Latin as the language of science was a turbulent development during the scientific revolution. Paracelsus was attacked by the medical establishment for presenting medical lectures in his native Swiss German in 1527 (Debus, 1978). Galileo published his 1612 *Discourse on Bodies in Water* in Italian, an act that enraged his fellow philosophers of the Florentine Academy (Sobel, 1999). For a long period, scholars who wrote in their vernacular tongue had to preface their writings with apologies and explanations of why this did not represent a challenge to the universities of the day (Debus, 1978).

Galileo wrote in Italian because "I must have everyone able to read it" (Sobel, 1999, p. 47). However, from some perspectives, writing in the vernacular actually produced a communication breakdown, because Galileo was not disseminating knowledge in the scholarly *lingua franca*, Latin. Galileo's writings were examined as part of his trial. It was concluded that "he writes in Italian, certainly not to extend the hand to foreigners or other learned men" (Sobel, 1999, p. 256).

A different sort of communication breakdown is a common theme in modern philosophy of science. It has been argued that some scientific theories are incommensurable with others (Feyerabend, 1975; Kuhn, 1970). Incommensurable scientific theories are theories that are impossible to compare because there is no logical or meaningful relation between some or all of the theories' terms. Kuhn argued that this situation would occur if, within a science, different researchers operated under different paradigms. "Within the new paradigm, old terms, concepts, and experiments fall into new relationships one with the other. The inevitable result is what we must call, though the term is not quite right, a misunderstanding between the two schools" (Kuhn, 1970, p. 149). Kuhn saw holders of different paradigms as being members of different language communities—even if they wrote in the same vernacular tongue! Differences in paradigms caused communication breakdowns.

The modern fragmentation of cognitive science might be an example of communication breakdowns produced by the existence of incommensurable theories. For instance, it is not uncommon to see connectionist cognitive science described as a Kuhnian paradigm shift away from classical cognitive science (Horgan & Tienson, 1996; Schneider, 1987). When embodied cognitive science is discussed in Chapter 5, we see that it too might be described as a new paradigm.

To view the fragmentation of cognitive science as resulting from competing, incommensurable paradigms is also to assume that cognitive science is paradigmatic. Given that cognitive science as a discipline is less than sixty years old (Boden, 2006; Gardner, 1984; Miller, 2003), it is not impossible that it is actually pre-paradigmatic. Indeed, one discipline to which cognitive science is frequently compared—experimental psychology—may also be pre-paradigmatic (Buss, 1978; Leahey, 1992).

Pre-paradigmatic sciences exist in a state of disarray and fragmentation because data are collected and interpreted in the absence of a unifying body of belief. "In the early stages of the development of any science different men confronting the same range of phenomena, but not usually all the same particular phenomena, describe and interpret them in different ways" (Kuhn, 1970, p. 17). My suspicion is that cognitive science has achieved some general agreement about the kinds of phenomena that it believes it should be explaining. However, it is pre-paradigmatic with respect to the kinds of technical details that it believes are necessary to provide the desired explanations.

In an earlier book, I argued that the assumption that cognition is information processing provided a framework for a "language" of cognitive science that made interdisciplinary conversations possible (Dawson, 1998). I demonstrated that when this framework was applied, there were more similarities than differences between classical and connectionist cognitive science. The source of these similarities was the fact that both classical and connectionist cognitive science adopted the information processing hypothesis. As a result, both schools of thought can be examined and compared using Marr's (1982) different levels of analysis. It can be shown that classical and connectionist cognitive sciences are highly related at the computational and algorithmic levels of analysis (Dawson, 1998, 2009).

In my view, the differences between classical and cognitive science concern the nature of the architecture, the primitive set of abilities or processes that are available for information processing (Dawson, 2009). The notion of an architecture is detailed in Chapter 2. One of the themes of the current book is that debates between different schools of thought in cognitive science are pre-paradigmatic discussions about the possible nature of the cognitive architecture.

These debates are enlivened by the modern rise of embodied cognitive science. One reason that classical and connectionist cognitive science can be easily compared is that both are representational (Clark, 1997; Dawson, 1998, 2004). However, some schools of thought in embodied cognitive science are explicitly anti-representational (Brooks, 1999; Chemero, 2009; Noë, 2004). As a result, it is not clear that the information processing hypothesis is applicable to embodied cognitive science. One of the goals of the current book is to examine embodied cognitive science from an information processing perspective, in order to use some of its key departures from both classical and connectionist cognitive science to inform the debate about the architecture.

The search for truth in the Library of Babel had dire consequences. Its librarians "disputed in the narrow corridors, proffered dark curses, strangled each other on the divine stairways, flung the deceptive books into the air shafts, met their death cast down in a similar fashion by the inhabitants of remote regions. Others went mad" (Borges, 1962, p. 55). The optimistic view of the current book is that a careful examination of the three different schools of cognitive science can provide a fruitful, unifying position on the nature of the cognitive architecture.

1.5 A Plan of Action

A popular title for surveys of cognitive science is *What is cognitive science?* (Lepore & Pylyshyn, 1999; von Eckardt, 1995). Because this one is taken, a different title is used for the current book. But steering the reader towards an answer to this excellent question is the primary purpose of the current manuscript.

Answering the question *What is cognitive science?* resulted in the current book being organized around two central themes. One is to introduce key ideas at the foundations of three different schools of thought: classical cognitive science, connectionist cognitive science, and embodied cognitive science. A second is to examine these ideas to see whether these three "flavours" of cognitive science can be unified. As a result, this book is presented in two main parts.

The purpose of Part I is to examine the foundations of the three schools of cognitive science. It begins in Chapter 2, with an overview of the need to investigate cognitive agents at multiple levels. These levels are used to provide a framework for considering potential relationships between schools of cognitive science. Each of these schools is also introduced in Part I. I discuss classical cognitive science in Chapter 3, connectionist cognitive science in Chapter 4, and embodied cognitive science in Chapter 5.

With the foundations of the three different versions of cognitive science laid out in Part I, in Part II, I turn to a discussion of a variety of topics within cognitive science. The purpose of these discussions is to seek points of either contention or convergence amongst the different schools of thought.

The theme of Part II is that the key area of disagreement amongst classical, connectionist, and embodied cognitive science is the nature of the cognitive architecture. However, this provides an opportunity to reflect on the technical details of the architecture as the potential for a unified cognitive science. This is because the properties of the architecture—regardless of the school of thought—are at best vaguely defined. For instance, Searle (1992, p. 15) has observed that "'intelligence,' 'intelligent behavior,' 'cognition' and 'information processing,' for example are not precisely defined notions. Even more amazingly, a lot of very technically sounding notions are poorly defined—notions such as 'computer,' 'computation,' 'program,' and 'symbol'" (Searle, 1992, p. 15).

In Part II, I also present a wide range of topics that permit the different schools of cognitive science to make contact. It is hoped that my treatment of these topics will show how the competing visions of the different schools of thought can be coordinated in a research program that attempts to specify an architecture of cognition inspired by all three schools.