54

PSYCHOLOGY

Richard Samuels

Introduction

The philosophy of psychology is concerned with issues that span work in the philosophy of science, philosophy of mind, and empirical psychology. Psychology is not a unified field but a diverse confederation of subfields and research programs, any of which could form a focal point for philosophical attention; and indeed many have, including psychoanalysis, social psychology, and abnormal psychology. But it is *cognitive psychology* – and the field of cognitive science, of which it is a central part – that has dominated research in the philosophy of psychology; and it is this research that I focus on here.

Though cognitive scientists disagree on many issues, one widespread commitment is that the mind is a *mechanism* of some sort: roughly, a physical device decomposable into functionally specifiable subparts. On this assumption, a central task for psychology is to characterize the nature of that mechanism: its basic operations, component parts, and development. Much philosophy of psychology is concerned with the project; and in the following sections I aim to provide a flavor of the research by considering three prominent issues:

- Is the mind is a *computer* of some sort?
- To what extent are minds *modular* in organization?
- To what extent is our mental structure *innately* specified?

Each issue combines in complex ways empirical and philosophical considerations; and collectively they identify many of the major faultlines that divide central positions in the philosophy of psychology and cognitive science.

Computationalism

If the mind is a machine, then what sort of machine might it be? One very influential answer is that the mind is a *computer*. According to this view, psychological processes such as perceiving, reasoning and remembering are – or, at any rate, depend on – computational processes. Although this general idea has dominated much philosophy of psychology and cognitive science, it has been elaborated in different ways; and among the most important and widely discussed distinctions is that between so-called *classical* and *connectionist* (or parallel distributed processing) versions.

Classical computationalism

Classicism is a view with deep historical roots, though it is perhaps the research of twentieth-century logicians, such as Alan Turing, that have exerted greatest influence on its conception of computation and, hence, of psychological processes. According to this view, the mind is a *symbol manipulation* device: an information-processing mechanism that operates on internally encoded bodies of information, called "data structures" or "symbols." Slightly more precisely, according to classicism:

(a) *Psychological processes employ mental symbols*. Such representations are language-like in that they possess both *semantic* properties – such as reference and meaning – and formal, or *syntactic*, properties: they are composed from constituents combined according to grammatical rules. For this reason classicists are sometimes said to advocate a "language of thought hypothesis" (Fodor 1975).

(b) *Psychological processes are sensitive to the syntactic structure of symbols*. Though symbols have semantic properties, cognitive processes are sensitive only to their *syntactic* or formal properties.

(c) *Psychological processes are algorithmic*. Roughly put, they can be characterized by sets of basic operations that are guaranteed to produce a determinate outcome in some finite number of steps. Those basic operations are sometimes said to be merely mechanical in the sense that no insight or ingenuity is required either to perform them or to determine what step to perform next.

Together these claims yield a general conception of psychological processes as algorithmically specifiable ones defined over the syntactic properties of mental symbols. For almost fifty years this proposal has been central to much work in cognitive science, where researchers have sought to specify the representations and algorithms on which such cognitive capacities as language, vision and reasoning depend. Moreover, it has played double duty as a metaphysics of mind. Minds, it is claimed, just are classical computers of the right sort; and having a thought (belief, desire, etc.) just is to bear an appropriate computational relation to some symbolic mental representation.

Virtues

Advocates of classical computationalism typically defend their view on explanatory grounds; for not only has it underwritten much productive empirical research, but it also helps explain some pervasive and fundamental aspects of cognition. Two are especially worthy of mention.

Rational causation

Many mental processes – most obviously reasoning – involve relations between mental states that are both causal and inferential (or rational). If I believe, for example, that all men are mortal and that I am a man, I may thereby come to believe I am mortal as well. In such a case, the earlier beliefs not only cause the latter, their meanings are also related in such a way as to provide premisses from which to infer the latter. Historically, this phenomenon posed a serious explanatory challenge: a version of the notorious *homunculus regress*. To explain such rational-cum-causal relations, it seems that meanings themselves must be causally efficacious, which in turn appears to require some inner interpreter – an intelligent subsystem, or homunculus – for which thoughts have meanings. But then the same problem of coordinating semantic and causal relations recurs for the homunculus, resulting in a regress of interpreters.

Classicists seek to address the problem by rejecting the assumption that rational causation is explicable only if meanings are causally efficacious. Instead they invoke an idea familiar to logicians, that inferences can be characterized proof-theoretically in terms of formal rules. (*Modus ponens* is a simple example.) When applied to the task of understanding cognition, the idea is that mental processes are inferential not because of any unexplained sensitivity to meanings, but because they depend on formal rules which, though defined over the syntax of representations, are like logical rules in that they preserve semantic relations. Moreover, since by assumption cognitive processes are algorithmic, they are ultimately decomposable into combinations of operations the execution of which requires no intelligence at all. The threat of regress is thus blocked and the homunculi expelled.

Productivity and systematicity

A second, widely cited, virtue of classicism is that it explains the productivity and systematicity of thought (Fodor and Pylyshyn 1988). Human thought seems *productive* in at least the sense that at any particular time we are capable of entertaining a great many thoughts, many of which are novel to us. Further, human thought seems *systematic* in roughly the sense that if someone is capable of entertaining some thoughts, he or she is thereby capable of thinking others as well. So far as we know, for example, no one is capable of entertaining the thought that John loves Mary, yet incapable of entertaining the thought that Mary loves John.

Classicists purport to explain those phenomena by assuming that thought depends on a combinatorial system of representations. On this view, thought is productive because relatively simple representations – if you like, words in the language of thought – can be combined according to syntactic rules to produce more complex expressions, which can in turn be combined according to the very same rules to produce still more complex representations, and so on *ad infinitum*. Similarly, thought is systematic because given some set of mental representations – "MARY," "LOVES," and "JOHN," for example – the very same rules, being defined over the syntax of the representations, permit the generation of multiple complex expressions – in the present case, both "JOHN LOVES MARY" and "MARY LOVES JOHN." Classicism's ability to provide elegant explanations of systematicity and productivity is widely regarded as among its main virtues.

Objections

For all its apparent virtues classicism has been subject to a bewildering array of objections. Some are relatively *a priori* in character. In his notorious *Chinese Room argument*, for example, John Searle purports to show that performing the right computations is insufficient for understanding. The argument proceeds from a thought-experiment:

> A native English speaker who knows no Chinese [is] locked in a room full of boxes of Chinese symbols (a database) together with a book of instructions for manipulating the symbols (the program). Imagine that people outside the room send in other Chinese symbols which, unknown to the person in the room, are questions in Chinese (the input). And imagine that by following the instructions in the program the man in the room is able to pass out Chinese symbols which are correct answers to the questions (the output). (Searle 1999: 115)

From outside it seems that the system understands Chinese. But, according to Searle, no matter what program the man executes, he won't know what the symbols *mean*. Thus mastery of syntactic operations – of the program – is insufficient for semantics; and since understanding a sentence requires a grasp of what the sentence *means*, running a program is insufficient for understanding as well.

The critical discussion surrounding Searle's argument is too large to consider in detail here (see Preston and Bishop 2002). But one common response is that, as an objection to classicism, it misses the mark. Classicists do not claim that executing the right program is, by itself, sufficient for thought. That would require the acceptance of a claim which classicists routinely deny: that computational role – the way the program uses a representation – determines its meaning. Rather, what classicists maintain is that thinking is a computational process operating on semantically evaluable representations, while *leaving open* – indeed frequently *endorsing* – the option that semantic properties are determined by something other than computational role, such as causal relations to the environment. Thus, according to the objection, the conclusion of Searle's argument is wholly compatible with the truth of classicism.

Another, more empirically oriented, kind of objection to classicism seeks to draw conclusions from explanatory failures of cognitive science. One major class of difficulties, often subsumed under the heading of the "frame problem," concern the explanatory challenge posed by our ability to determine the information that is *relevant* to the tasks we perform (Ford and Pylyshyn 1996). In particular, when making plans or revising our beliefs, we somehow manage to identify the information that is relevant to the task at hand and ignore the rest. How is this *relevance sensitivity* to be explained in classical terms? It is implausible that we survey *all* our beliefs, since such a strategy would require more time and computational power than we possess. Some more computationally feasible process is required. Yet many doubt that such a process can be specified in classical terms. It has been suggested, for example, that relevance is unlikely to be explicable in classical terms because it is a *holistic* property of thought, in roughly the sense that the relevance of a given thought depends on a broad array of *surrounding conditions*, such as one's background beliefs and intentions.

Connectionism

Whether classicists can address this and other problems remains a point of considerable dispute. But many have taken such challenges as grounds for exploring alternative accounts of cognition, of which the most influential is *connectionism*. Though connectionist proposals vary considerably in detail, they share a basic, neurally inspired, conception of our cognitive machinery. Cognitive systems are, on this view, *multilayer* networks of *nodes* attached to one another by *weighted connections*. In prototypical networks, activation spreads from an input layer of nodes to an output layer – typically via *hidden* layers of units – and the weights of connecting nodes are adjusted by some sort of *learning algorithm*, such as back-propagation, so that the system can "learn" to perform various tasks. This general conception of cognitive systems has proven to be of considerable utility to psychologists and has been used with varying degrees of success to model many psychological processes and capacities, including vision, language acquisition, concept-learning, and motor control.

On some conceptions of connnectionism, there is no conflict with classicism. For example, one common view, known as "implementational connectionism," seeks not to replace classicism but merely to explain how classical systems are implemented or realized in the brain. But even those who seek to displace classical accounts – so-called "eliminative connectionists" – typically acknowledge many important commonalities. Specifically, they often share with classicists the assumptions that cognition is both representational and computational. It is representational because the nodes in a connectionist network – especially input and output nodes – are widely assumed to represent properties and objects; and they are computational both because learning rules are algorithmic and because the spread of activation from input to output nodes can be interpreted as computing a function.

Where, then, *do* the main differences reside? Perhaps the most widely cited difference is that connectionist representations are typically not *syntactically* structured. As a consequence, connectionists typically reject both the classical conception of mental representation and the attendant account of cognitive processes as defined over the syntactic properties of representations.

Virtues

Connectionist systems are often said to possess characteristics that make them apt for modeling cognition, including:

- *Speed*: because networks process information in parallel they can be fast.
- *"Graceful degradation"*: in contrast to classical computers, the performance of a neural network remains relatively unaffected by degradation in the input signal or by damage to the system.
- *Neural-realism*: networks are more brain-like than are classical computers.

- *Learning*: connectionist networks show an impressive ability to "learn from experience."
- *Multiple constraint satisfaction*: neural networks easily address problems that require the resolution of many conflicting constraints in parallel.

Critics respond that some of those virtues (e.g., speed and "graceful degradation") are not reasons for rejecting classicism, but at most reasons for adopting implementational connectionism. Other putative virtues, they claim, have been over-sold. For example, it has been argued that the resemblance to real brains is a very loose one; and that classical systems also learn and solve problems involving multiple constraints. An assessment of those claims remains a topic of ongoing debate.

Objections

It has also been argued that eliminative connectionism exhibits some serious deficiencies. Perhaps the most common complaint is that it fails to explain core aspects of our representational capacities. For instance, Fodor and Pylyshyn (1988) argue that connectionists lack a satisfactory explanation of the systematicity and productivity of thought. More recently, Gary Marcus (2001) has argued that connectionist networks of the normal sort fail to accommodate the fact that humans not only represent categories (such as the category of "cats") but also individuals (e.g., Tiddles and Tom).

Another concern is that connectionism has done little to address the deepest problems encountered by classical approaches. For example, the frame problem arises most clearly in relation to flexible, knowledge-intensive, processes such as reasoning and planning. But connectionism has made relatively little progress in understanding such processes, let alone in providing any systematic account of how we successfully identify relevant information when engaged in reasoning or planning.

Hybrid views and radical alternatives

In recent years, theorists have become less inclined to view the *classicism–connectionism* debate as a dispute between two mutually exclusive versions of computationalism. One common proposal is that we need to posit *hybrid* models that combine both classical and connectionist components. It has been suggested, for example, that "higher" cognitive processes, such as planning and deliberative reasoning, depend on a classical architecture, while more associative processes, such as implicit learning, depend on connectionist mechanisms (Sloman 1996).

Another, more radical, development is the claim that both classicists and most connectionists are wrong to assume that the mind is a computer of *any* sort. Instead, it is claimed that we should think of the brain's neural networks and the connectionist systems used to model them as *dynamical systems* best described by the sorts of differential equations found in physics (Port and van Gelder 1995). Assessing this dynamical systems theory and other alternatives remains a central project for the philosophy of psychology.

Modularity

The classicism–connectionism debate is concerned largely with the mind's *micro-architecture*: the basic elements and operations on which mental activity depends. But there is widespread agreement that minds are also organized into larger *macro-architectural* units. Historically, these were called "faculties," though contemporary theorists tend to speak of "cognitive systems"; and in recent years much discussion of the nature of those systems has occurred in the context of debate over *modularity*.

To a first approximation, debates over modularity concern the extent to which minds are composed from autonomous systems dedicated to restricted information-processing tasks. Systems that are restricted in those ways tend to be referred to as "modules"; and those relatively free from such constraints are said to be "non-modular." At one extreme, for example, is the sort of radically *non-modular* view of minds as comprised of one (or perhaps a few) general-purpose computers that can process many different kinds of information, and thereby perform many different tasks. At the other extreme, is the sort of radical modularity on which minds are composed from thousands of highly specialized and entirely autonomous devices, each dedicated to a very specific task and capable of only processing a highly restricted range of information. In reality, neither position is taken seriously. Instead, the debate is concerned largely with articulating and assessing a range of intermediate positions.

Fodorian modularity

One well-known modularity hypothesis defended by Fodor (1983) and others is that the modular structure of the mind is restricted to *input* systems (those responsible for perception, including language perception) and *output* systems (those responsible for producing behavior). On this view, central systems – those responsible for reasoning and decision-making – are non-modular. Thus minds are modular only at the periphery.

Fodor's defense of this proposal goes hand-in-hand with an attempt to articulate an appropriate notion of modularity. Fodorian modules are characterized by a cluster of features that they tend to exhibit to some interesting degree. Specifically, modules are prototypically:

- *domain-specific*: they operate on a limited range of inputs, defined by some task domain like vision or language-processing;
- *informationally encapsulated*: they have limited access to information in other systems;
- *inaccessible*: other mental systems have only limited access to a module's computations;
- *shallow*: their outputs are not conceptually elaborated;
- *mandatory*: they respond automatically to input;
- *fast*: their operation is relatively fast;
- *neurally localized*;
- subject to *characteristic* and *specific breakdowns*; and
- their development exhibits a *characteristic pace* and *sequence*.

Not all of these characteristics are of equal theoretical importance. Domain specificity and informational encapsulation are the most central, while the others, in large measure, are empirical consequences of those more central characteristics. Fodor argues, on the basis of evidence from the study of vision and speech comprehension that input systems are modular in the above sense. In contrast, he maintains, central systems are likely to be both domain-general and informationally unencapsulated. They are likely to be domain-general because the processes responsible for reasoning and decision-making function to combine inputs from different perceptual domains. And they are likely to be unencapsulated because there are few constraints on the sorts of information we can use in determining what to believe or what to do. For example, Fodor maintains that almost any of a person's beliefs can be relevant to the sort of reasoning characteristic of science – what is sometimes called "abductive" reasoning, or "inference to the best explanation."

Massive modularity

Though Fodor's view has been challenged from many directions, one of the most recent and intriguing responses comes from those who advocate a *massive modularity hypothesis* (MM). Advocates of MM accept that input and output systems are modular. But, *pace* Fodor, they maintain that central systems are largely or entirely modular as well. So, for example, it has been suggested that there are modules for such central processes as social reasoning, biological categorization, and probabilistic inference.

What should we make of that proposal? As one would expect, it will depend in large measure on an assessment of evidence for and against the existence of particular modules – evidence which at this time is inconclusive. But advocates of MM also defend their views on the basis of quite general considerations about the nature of cognition. Consider the following example:

> *Task Specificity Argument:* There are a great many cognitive tasks whose solutions impose quite different demands. So, for example, the demands on vision are distinct from those of speech recognition, probabilistic judgment, grammar induction, and so on. Moreover, since it is very hard to believe there could be a single general inference mechanism for all of them, for each such task we should postulate the existence of a distinct mechanism, whose internal processes are computationally specialized for processing different sorts of information in the way required to solve the task. (Carruthers 2006)

This argument is not intended as a deductive proof of MM, but only to render it plausible. Nonetheless, I doubt it shows even this much. If the only alternative to MM were a mind comprised of a single *general-purpose* mechanism treating all problems in the same way, then MM would be the more plausible option. But these are manifestly *not* the only options. First, denying MM is wholly compatible with the existence of many specialized mechanisms for perception and motor control. But even if we focus on central systems, positing multiple dedicated modules is not the only way of explaining our capacity to perform many different reasoning tasks. A familiar alternative is that relatively unspecialized inference mechanisms use different bodies of specialized information in solving different problems. A major difficulty with the present argument is that it fails to adjudicate between MM and this familiar alternative.

So, it is far from clear that the standard arguments for MM are satisfactory. It is also worth noting that MM, at least in radical form, struggles to accommodate some central aspects of human cognition. For example:

- *Conceptual integration*: we are capable of freely combining concepts across different subject matters or content domains. Not only can I have thoughts about colors, about numbers, about shapes, and so on, but I can have thoughts that concern *all* these things – for example, that I had two, roughly round, red steaks for lunch.
- *Generality of thought*: not only can we freely combine concepts, we can also deploy the resulting thoughts in our theoretical and practical deliberations – to assess their truth or plausibility, but also to assess their relevance to our plans and projects.
- *Inferential holism*: given surrounding conditions – especially background beliefs – the relevance of a representation to the theoretical or practical tasks in which one engages can change dramatically. Indeed, it would seem that given appropriate background assumptions, almost any belief can be relevant to the task in which one engages.

Although some maintain that those features can be accommodated within a wholly modular account of cognition, a more plausible approach is to posit some genuinely non-modular central systems. This does not require that all central systems be modular in the way Fodor appears to suppose. Another possibility is that central processes are subserved by both modular and non-modular systems (Stanovich 2004). According to its advocates, this *dual systems* account possesses the virtues of MM while better accommodating a host of phenomena, including those outlined above.

Nativism

Thus far I have discussed two general issues about the structure of the mind. A related issue concerns the *acquisition* of mental structure: To what extent is the mind's structure *innately specified*? Discussion of the question is often couched as a dispute between *nativism* and *non-nativism* of which empiricism is a central sort. In brief, nativists claim that the mind contains *lots* of innate structures: concepts, bodies of information, psychological mechanisms, and modules. In contrast, non-nativists maintain that the mind contains relatively little innate structure. For example, empiricists typically suggest that the mind comes equipped with little more than perceptual mechanisms and a few systems for domain-general learning, such as associative learning mechanisms (e.g., Pavlovian conditioning) and general-purpose, inductive, learning mechanisms.

Linguistic nativism

Disputes over innateness have emerged in connection with a broad array of psychological phenomena, including our intuitive understanding of the physical world, arithmetic, and concept acquisition. But it is in connection with language that the issues have been most extensively explored. Here, largely under Noam Chomsky's influence, nativist proposals have dominated research for almost half a century.

Researchers working on language tend to suppose that when acquiring a language one comes to possess an internal grammar – or an internal representation of a grammar– for that language. (This helps explain, among other things, the systematicity and productivity of language.) Clearly, it is implausible that the grammar possessed by a competent speaker – for instance, a grammar for English as opposed to French or Hindi – is innately specified since the grammar that one acquires depends on the linguistic environment that one inhabits. Nonetheless, in contrast to other organisms, all humans everywhere – save those suffering extreme pathology or environmental deprivation – reliably acquire competence in some natural language within the first few years of life. That suggests, with only a hint of idealization, that humans share some set of innate resources – some *initial state* – that permits the acquisition of a grammar for the language they speak. A central problem for any account of language acquisition is thus to characterize the initial state: those innate resources which reliably enable a grammar to be acquired on the basis of the available environmental information.

What are the options? One major distinction is that between linguistic *empiricism* and linguistic *nativism*. Empiricists claim that language acquisition depends on the same domain-general mechanisms that are responsible for cognitive development in other domains. In contrast, linguistic nativists claim that at least some of the innate resources on which language acquisition depends are specific to the domain of language. But even if one endorses some version of linguistic nativism, there is still plenty of room for disagreement over the nature and extent of our innate language-specific resources. For instance, Chomskians claim that we possess an innate *universal grammar*: a rich body of innately specified knowledge that specifies the properties shared by all natural languages (Chomsky 1980). But one may be a linguistic nativist without being a Chomskian. For example, one might think there is an innately specified, language-specific, learning mechanism or module, while denying that there is an innate universal grammar.

Arguments

The debate over linguistic nativism is a largely empirical one; and like other empirical debates, different proposals are assessed in terms of their overall ability to accommodate evidence in a simple, powerful, and conservative manner. Here, there are many sorts of evidence that are relevant, including: evidence for linguistic universals; evidence concerning the relative ease of language acquisition; evidence concerning the specific patterns of error that occur during language acquisition; evidence of selective impairment and genetic disorders; and evidence from computational modeling. But perhaps the most influential argument for linguistic nativism – and the one that has received most attention from philosophers – has come to be known as the *poverty of the stimulus argument* (PoSA).

The PoSA has been formulated in a number of different ways. But the rough idea is that some version of linguistic nativism must be correct because the information that children receive from the environment is too impoverished to permit an *empiricist learner* – one lacking any innate language-specific knowledge, mechanisms, or biases – to acquire the grammar for their language.

Though the PoSA has been widely accepted by linguists, it has also been subjected to sustained criticism. One major challenge concerns the issue of *what* environmentally derived information is available in the course of language acquisition. Nativists have tended, for example, to suppose that children are seldom provided with negative data – roughly, information about when an utterance is *not* grammatical. But recently that assumption has come under scrutiny; and researchers have argued that such data are both available to and used by children in the course of language development (Chouinard and Clark 2003).

Another major challenge concerns the nature of empiricist learners. Almost everyone agrees that traditional empiricist accounts of language-learning, such as those that have emerged from the behaviorist tradition, are inadequate. But in recent years there has been an explosion of research on statistical learning (Pereira 2000); and some have suggested that this research may form the basis for a satisfactory empiricist account of language acquisition. Though a systematic assessment of the methods is beyond the scope of the present chapter, it is far from clear that they undermine the PoSA for linguistic nativism. Recall: What the PoSA purports to show is merely that language acquisition requires some set of innate language-specific structures or biases. But the current state of research on statistical learning seems wholly compatible with this claim. Specifically, our most successful computational models of language-learning invariably assume language-specific constraints. For example, they assume some model (or representational scheme) relevant to the domain of language; and they presuppose constraints on the inputs that the learning system receives (e.g., sentences in the target language as opposed to the myriad other kinds of inputs that a learning device may receive). Though there is much more to say on the matter, it is far from clear that without an account of how such constraints are acquired by empiricist learning, those models vindicate empiricism as opposed to suggesting a variant on linguistic nativism: one which posits an innate language-specific, statistical, learning mechanism or module.

What is innateness?

Much debate over *innateness* in cognitive science proceeds under the assumption that the notion is clear enough to permit the framing of substantive empirical issues. But there are, in fact, considerable difficulties in understanding what innateness is; and some prominent theorists have even suggested that very concept is "fundamentally confused" (Griffiths 2002). If such a claim could be sustained, it would appear to have important implications for research in psychology. For not only would it undermine nativism in its various forms, but it would also threaten the main empiricist alternatives, since they too presuppose the coherence of the innateness concept.

One standard reason for claiming that innateness is a confused concept is that it is said to confound several properties under a single term: properties that are neither co-extensive nor, by themselves, adequate to characterize what we mean by "innate." For instance, it is sometimes claimed that innate traits are those that are *present at birth*, even though presence at birth is neither necessary nor sufficient for innateness. It is not sufficient, because prenatal learning is possible; and is it not necessary, because, as Descartes observed long ago, innate characteristics can be acquired quite late in development. (Illustration: pubic hair is plausibly innate but clearly not present at birth.) Similarly, it is sometimes said that innate traits are solely the products of internal (including genetic) causes, even though this is clearly not necessary for innateness, since, like all contemporary theorists, nativists wholeheartedly accept the banal thesis that cognitive traits are caused jointly by internal *and* environmental factors.

In view of the problems with standard claims about innateness, philosophers of psychology have responded in a variety of ways. One response is to conclude that innateness is a confused concept and map out the implications of this for future psychological research. Another response is to try to make systematic sense of the notion of innateness that figures in psychology and allied sciences. Though this is not the place to pursue the matter in detail, at least two proposals merit further consideration. The first is that innate traits are those that are environmentally *canalized*. Roughly put: a trait is innate on this view when it is relatively insensitive to the range of environmental conditions under which it emerges (Ariew 1999). The second suggestion is that an innate psychological trait is one that is *psychologically primitive*. Roughly: it is acquired in the normal course of events, though not by psychological processes, such as learning or perception (Cowie 1999; Samuels 2002). Like so many other issues in the philosophy of psychology, deciding which (if any) of these options to adopt remains a topic for active and ongoing debate.

See also Observation; Cognitive science.

References

Ariew, André (1999) "Innateness Is Canalization: A Defense of a Developmental Account of Innateness," in V. Hardcastle (ed.) *When Biology Meets Psychology*, Cambridge, MA: MIT Press.

Carruthers, Peter (2006) *The Architecture of the Mind: Massive Modularity and the Flexibility of Thought*, Oxford: Oxford University Press.

Chomsky, Noam (1980) *Rules and Representations*, New York: Columbia University Press.

Chouinard, M. M. and Clark, E. V. (2003) "Adult Reformulations of Child Errors as Negative Evidence," *Journal of Child Language* 30: 637–69.

Cowie, Fiona (1999) *What's Within? Nativism Reconsidered*, New York: Oxford University Press.

Fodor, J. A. (1975) *The Language of Thought*, New York: Thomas Crowell.

—— (1983) *The Modularity of Mind*, Cambridge, MA: MIT Press.

Fodor, J. and Pylyshyn, Z. (1988) "Connectionism and Cognitive Architecture: A Critical Analysis," *Cognition* 28: 3–71.

Ford, K. M. and Pylyshyn, Z. W. (eds) (1996) *The Robot's Dilemma Revisited: The Frame Problem in Artificial Intelligence*, Norwood, NJ: Ablex.

Griffiths, Paul (2002) "What Is Innateness?" *Monist* 85: 70–85.

Marcus, Gary F. (2001) *The Algebraic Mind*, Cambridge, MA: MIT Press.

Pereira, Fernando (2000) "Formal Grammar and Information Theory: Together Again?" in *Philosophical Transactions of the Royal Society* A358: 1239–53.

Port, R. and van Gelder, T. J. (1995) *Mind as Motion: Explorations in the Dynamics of Cognition*, Cambridge, MA: MIT Press.

Preston, John and Bishop, Michael (eds) (2002) *Views into the Chinese Room: New Essays on Searle and Artificial Intelligence*, New York: Oxford University Press.

Samuels, Richard (2002) "Nativism in Cognitive Science," *Mind and Language* 17: 233–65.

Searle, John (1999) "The Chinese Room," in R. A. Wilson and F. Keil (eds) *The MIT Encyclopedia of the Cognitive Sciences*, Cambridge, MA: MIT Press.

Sloman, S. A. (1996) "The Empirical Case for Two Systems of Reasoning," *Psychological Bulletin* 119: 3–22.

Stanovich, K. E. (2004) *The Robot's Rebellion: Finding Meaning in the Age of Darwin*, Chicago: University of Chicago Press.

Further reading

There are a number of good anthologies and introductory texts in the philosophy of psychology: José Luis Bermúdez, *Philosophy of Psychology: A Contemporary Introduction* (New York: Routledge, 2005); Andy Clark, *Mindware: An Introduction to the Philosophy of Cognitive Science* (Oxford: Oxford University Press, 2001); and George Botterill and Peter Carruthers, *The Philosophy of Psychology* (Cambridge University Press, 1999) all provide good, though quite different, introductions to the field. Denise Delarosa Cummins and Robert Cummins (eds) *Minds, Brains and Computers* (Oxford: Blackwell, 2000) contains many influential papers, especially on computational approaches to cognition; and the relative merits of classicism and connectionism are discussed at length in Cynthia Macdonald and Graham Macdonald (eds) *Connectionism: Debates on Psychological Explanation* (Oxford: Blackwell, 1995). For very different assessments of connectionist theory see: William Bechtel and Adele Abrahamson, *Connectionism and the Mind*, 2nd edn (Malden, MA: Blackwell, 2002); and GaryMarcus, *The Algebraic Mind* (Cambridge, MA: MIT Press, 2001). For discussion of various facets of debate over nativism, see Cowie (1999); and for differing treatments of modularity, see Steven Pinker, *How the Mind Works* (New York: W. W. Norton, 1997); Jerry Fodor, *The Mind Doesn't Work That Way* (Cambridge, MA: MIT Press, 2000). For an impressive range of papers on innateness and modularity, see Peter Carruthers, Stephen Laurence, and Stephen Stich's 3-volume *The Innate Mind* (New York: Oxford University Press, 2005, 2006, 2007). Finally, Rob Stainton (ed.) *Contemporary Debates in Cognitive Science* (Malden, MA: Blackwell, 2006) contains state-of-the-art discussions of many central topics in the philosophy of psychology.