Lappin S (1996) The interpretation of ellipsis. In: Lappin S (ed.) *The Handbook of Contemporary Semantic Theory*, pp. 145–175. Oxford, UK: Blackwell.
Lappin S and Benmamoun EA (eds) (1999) *Fragments: Studies in Ellipsis and Gapping*. Oxford, UK: Oxford University Press.
Lobeck A (1995) *Ellipsis: Functional Heads, Licensing, and Identification*. New York, NY: Oxford University Press.
Merchant J (2001) *The Syntax of Silence – Sluicing, Islands, and the Theory of Ellipsis*. Oxford, UK: Oxford University Press.
Rooth M (1992) Ellipsis redundancy and reduction redundancy. In: Berman S and Hestvik A (eds) *Proceedings of the Stuttgart Ellipsis Workshop*. Heidelberg: Arbeitspapiere des Sonderforschungsbereichs 340, Bericht Nr. 29, IBM Germany.
Schwabe K and Zhang N (2000) *Ellipsis in Conjunction*. Tübingen, Germany: Max Niemeyer Verlag.
Steedman M (1990) Gapping and constituent coordination. *Linguistics and Philosophy* **13**: 207–264.
Zoerner CE (1995) *Coordination: The Syntax of &P*. PhD dissertation. University of California, Irvine, CA.

Embodiment

Intermediate article

Ronald Chrisley, University of Birmingham, Birmingham, UK
Tom Ziemke, University of Skövde, Skövde, Sweden

CONTENTS

An understanding of how cognition is realized or instantiated in a physical system, especially a body, may require or be required by an account of a system's embedding in its environment, its dynamical properties, its (especially phylogenetic) history and (especially biological) function, and its nonrepresentational or noncomputational properties.

INTRODUCTION

In recent years a number of researchers in cognitive science and artificial intelligence (AI) have criticized many traditional approaches to modeling, building and understanding cognitive systems as not placing sufficient emphasis on the body or physical realization of such systems. Non-embodied approaches to cognitive science typically involve some or all of the following features, to a greater or lesser extent:

- The belief that cognition is computation, and thus can be understood in an implementation-independent way, allowing cognitive science to proceed independently of biology and neuroscience.
- A search for general-purpose cognitive abilities, not relativized to any particular (biological, sensorimotor, physical) context or need.
- A method of analysis, modeling and design that for the most part ignores temporal aspects of cognition, in that it focuses on behaviors (e.g. chess playing) that are evaluated in terms of 'getting the right answer' rather than exhibiting a particular dynamic profile, and sees cognition as a module that mediates between the deliverances of a causally prior perceptual module and the inputs to an autonomous action system.

In contrast, embodied approaches to cognition typically involve some or all of the following features, again to varying extents:

- Acknowledgment of the role that the body and its sensorimotor processes can and do play in cognition. Some aspects of the system that would, on the traditional view, be considered mere matters of implementation, are instead taken to be crucial components.
- Understanding of cognition in the context of its (especially evolutionary) biological function: to support the activities of the body.
- A view of cognition as a real-time, situated activity, typically inseparable from and often fully interwoven with perception and action.

'Embodied' cognitive science or artificial intelligence, then, refers to a range of loosely affiliated philosophies, explanatory frameworks and design methodologies that strive to redress a perceived neglect of the body in cognitive science.

Since the mid-1980s there has been a rapid increase in interest in embodied cognition (and use of the term 'embodiment'), but there are many aspects of cognitive science and artificial intelligence research conducted in the 1960s and 1970s that take embodiment into account.

ISSUES CONCERNING EMBODIMENT

The issue of embodiment is closely related to, though distinct from, several other issues of recent interest in cognitive science.

Embeddedness

Recognizing the role of the body in cognition facilitates an approach that sees cognition as partly constituted by, or in terms of its relationship to, the environment (Clark, 1997). Thus, embodied cognitive science is naturally related to investigations of externalism (the belief that mental states are partially constituted by states external to the cognizer), situatedness (the importance, in cognitive activity, of a cognizer's location in and relations to the environment), offloading (using aspects of the environment, such as numerals when doing long division, to reduce cognitive load), scaffolding (the assistance a developing infant gains from, e.g. interacting with adults who already have the ability being acquired), and interactivity (cognitive phenomena, such as turn-taking, which depend crucially on the dynamics of interaction between a cognizer and an object or other cognizer).

Noncomputationalism and Nonrepresentationalism

Once one acknowledges the presence of the body, it is possible to use bodily properties, dynamics and configurations to explain some abilities, behaviors and phenomena that previously were thought to require explanation in computational or representational terms (Varela *et al.*, 1991). However, it is still unclear whether such explanations which advert to bodily states are themselves noncomputational and nonrepresentational, or whether they instead invoke a new form of computation and/or representation. It certainly seems that an embodied approach need not be anticomputationalist or antirepresentationalist (Clark, 1997).

Dynamics

Many researchers who have taken an embodied approach have found it useful to turn away from discontinuous, nontemporal, logic-based formalisms and instead use the continuous mathematics of change offered by dynamical systems theory as a way to characterize and explain cognitive phenomena (Port and van Gelder, 1995). Again, while there may be natural affinities between embodied cognitive science and these tools, it is certainly possible to have one without the other.

Biology

Perhaps the most obvious connection between the mind and the body is the brain: surely an important part of understanding the mind is to understand the neurophysiology underlying cognition (Churchland, 1986). But there are other connections with biology as well. For example, some researchers have maintained that one can best understand natural cognitive systems in terms of the biological function and purposes that cognitive faculties served in the phylogenetic development of their bodies (Millikan, 1984). This requires understanding not only bodies, past and present, but the evolutionary niches of such bodies as well. An important constraint, then, on models of human cognition will be whether the proposed architecture is an evolvable one: whether it is the kind of architecture that could have been reached through a process of natural selection, given the conditions known to have been in place in our natural history (Sloman, 2001).

VARIETIES OF EMBODIMENT

There are several dimensions of variation in the views and theories of embodiment currently being considered in cognitive science.

Criteria for Embodiment

Perhaps the most important dimension of variation concerns what criteria must be met for something to be an embodied cognizer: what is to count as a body? One can, partly following Ziemke (2001), distinguish a range of views on this question, from the least to the most constraining.

Physical realization

According to one view, to be embodied is just to be realized in some physical substrate. All work in cognitive science is about embodied systems in this sense: even a virtual web agent must be realized in some physical facts at any given time. Only purely mental entities or spirits would lack this kind of embodiment. Thus, even traditional cognitive science is not as disembodied as some have claimed. For example, the influential notion of a physical symbol system (Newell and Simon, 1976) explicitly acknowledges the requirement that a cognitive system be embodied in this (weakest) sense.

Physical embodiment

Physical embodiment requires that the realizing physical system be a coherent, integral system, that to some degree persists over time. This would rule out virtual web agents, the physical realization of which can be radically distributed over the entire planet, but could still include any conventional robot. A tension arises here: if, as some theorists have argued, human cognition extends into the tools and other physical environmental states we exploit, then the localized, biological body is less relevant to cognitive science than the extended, constantly changing, distributed physical system that at any given time includes parts of our environment as well as parts of our bodies. Thus, active externalism (Clark and Chalmers, 1999) may be incompatible with strong notions of embodiment.

Organismoid embodiment

According to another view, the localized physical realization of the system must share some (possibly superficial) characteristics with the bodies of natural organisms, in terms of form or sensorimotor capacities, but need not be alive in any sense. The most prominent, and perhaps the most complex, examples of organismoid embodiment are humanoid robots such as Cog (Brooks and Stein, 1994) and Kismet (Breazeal and Scassellati, 2000). Work with these robots is based on the view that research in AI and cognitive robotics, in order to be able to investigate human-level cognition, has to deal with systems that have bodies which, although artificial and possibly non-living, have at least some humanlike characteristics. For example, Kismet is a humanoid robot that learns how to visually track objects. To do this, a human trainer must move objects of an appropriate size at an appropriate speed at an appropriate distance in front of Kismet's eyes. If the human trainer moves the tracked object too close to Kismet, Kismet responds by raising its eyebrows in a manner which in humans indicates a startle response. This naturally causes the human trainer to startle in return, which prompts a change in the training parameters of speed and distance. Thus, Kismet's organismoid embodiment, in the form of eyebrows and facial expressions, is an integral part of the human–robot training dynamic which should tend towards homeostasis.

Organismal embodiment

The strongest criterion of embodiment is that the body must not only be organism-like, but actually organismal and alive (e.g. Sharkey and Ziemke, 1998). Of course, this raises the question of what is required for something to be alive. Various answers to this question have been proposed, including the ability to metabolize, reproduce, regenerate, or grow; autonomy; and autopoiesis (e.g. Maturana and Varela, 1980; von Uexküll, 1928).

Other Dimensions of Variation

Another dimension of variation is the extent of the domain of cognition that requires an embodied approach. For example, one can ask whether reference to the body is required only for giving an account of low-level, sensorimotor aspects of human cognition, or if it is required for all forms of human cognition, including reasoning, mathematical thought, and language use (cf. the distinction between 'material' and 'full' embodiment in Nuñez (1999)).

Theorists might also disagree on how radical the effect of taking an embodied approach will be on the concepts, theories and methods of cognitive science. Clark (1999) distinguishes between simple and radical embodiment. In simple embodiment, the framework of traditional cognitive science is retained, and facts about embodiment are treated as mere constraints on theories of, for example, 'inner organization and processing'. Radical embodiment is more ambitious, and 'treats such facts as *profoundly altering the subject matter and theoretical framework of cognitive science*' (p. 348, emphasis in original). Radical embodiment is advocated by, for example, Lakoff and Johnson (1999). These authors claim that a central finding of cognitive science is that the mind is inherently embodied, and that this, together with the other two central findings (that thought is largely unconscious, and that abstract concepts are largely metaphorical), force us to reject not just the Western philosophy of mind, but most or all of Western philosophy, including especially Anglo-American analytic philosophy but also including postmodern philosophy. This throws everything into question: the nature of truth, meaning, time, space, language, rationality, and especially the self.

Finally, one can distinguish between epistemological and metaphysical approaches to an embodied cognitive science. An epistemological approach maintains that concepts concerning the body will be required to understand and explain cognition (even if the cognitive system itself is disembodied); a metaphysical approach maintains either that cognitive processes must be realized in a body or that AI should proceed by making embodied robots, but makes no claim as to whether embodiment must be adverted to in explanations of cognitive activity. For example, research involving the robot Shakey (Nilsson, 1984) was metaphysically embodied, but since its design and explanation focused primarily on Shakey's functional, computational aspects (namely, deliberation and planning) and not on Shakey's embodiment, this research was not epistemologically embodied. Conversely, computer simulations are by their very nature not (metaphysically) embodied, in most senses of 'embodiment', yet many researchers use (epistemologically) embodied simulations to model the crucial role a cognizer's body plays in its activity.

PHILOSOPHICAL CONCEPTIONS OF EMBODIMENT

Before the twentieth century, the most influential view of mind in Western thought was dualistic: the mind was regarded as composed of a separate, extensionless, nonphysical substance. This view led to many insoluble problems, both philosophical and empirical. For example, how do the mental and physical realms interact? How can we scientifically investigate something that is not in the physical world?

Behaviorism rejected dualism, and thereby opened up the possibility of scientific enquiry into an embodied mind, but it left little room for an understanding of the processes underlying much of mentality: it addressed little of what would be called 'cognition' today. Also, it actively avoided explaining or mentioning experience or consciousness (as did many later cognitivists). More promising steps towards an embodied understanding of mind were taken in the first part of the twentieth century: the relevance of von Uexküll's notion of the body has already been mentioned, but there were several other notable thinkers.

For example, in the 1920s Heidegger (1962) developed a phenomenology that understood human activity not as the result of the manipulation of context-free representations of objects, but as the contextualized experience of the body–environment system. Explicit, decontextualized representation of a hammer as an independent object occurs only when there is some kind of breakdown in the system (e.g. the hammer is too heavy). Dreyfus (1992) elaborated this and other aspects of Heidegger's philosophy into a critique of early, disembodied AI work, calling, for example, for systems that react to the particularities of the current perceptual/action situation rather than ones that attempt to create general-purpose, long-term plans. Although heartened by some aspects of neural or connectionist approaches to understanding the mind, Dreyfus provisionally concludes that, like their symbolic predecessors, connectionist models of cognition suffer from their lack of embodiment, in two ways. Firstly, by not being embedded in a real world with actual bodily concerns, most sophisticated connectionist learning must rely on the intervention of a human teacher, which prohibits the connectionist system from developing its own, genuine, meaningful attitude to the world. Secondly, he speculates that connectionist systems will never generalize in a way that we can recognize as being intelligent and meaningful until they have a form of life sufficiently similar to ours – which requires at least a body of some sort, perhaps even a humanoid one.

In the 1930s, Vygotsky (1978) claimed that language is an inherently socially situated activity, and that one can only understand a child's acquisition of language by recognizing this social context. It follows that inasmuch as social activity is embodied, the development and deployment of linguistic faculties will have to be understood as embodied as well.

While placing less emphasis on social embedding, Piaget (1954) was more explicit about the role of the body in the development of cognitive abilities. For example, his accounts typically made essential reference to the notion of a circular reaction, which in its primary form is the repetition of an activity in which the body starts in one configuration, goes through a series of intermediate stages, and eventually arrives at the starting configuration again. Thus, the kinds of abilities that an organism may acquire depend critically on what circular reactions are possible given that organism's body.

In the 1940s, Merleau-Ponty (1962) made the body central to his phenomenology of mind. For example, he claimed that we have intentions that we do not choose to have, by virtue of our bodies being the way they are. Furthermore, the way we perceive an object is determined by the modes of interaction that our bodies, given the nature of the object, allow. (This idea was a precursor of the notion of affordances (Gibson, 1979).) Even more strongly, Merleau-Ponty saw the body as the necessary medium for our having a world at all, with the nature of the activities of the body determining the nature of what could be experienced in our world. Further, the body could be augmented with tools to further develop the elements of our lived world.

Despite the emphasis that these thinkers placed on the body for understanding the mind and behavior, and partly because of their context outside the Anglo-American tradition, it is only recently that the notion of embodiment has had a significant influence on mainstream cognitive science and AI. Instead, these fields were, from their inception in the mid-1950s, dominated by the computer metaphor for mind (not surprising, perhaps, since the notion of computation had for centuries been a concept of the (human) mental activity of symbol and number manipulation). In particular, in the absence of any suggestion as to how they could constrain empirical investigation and modeling, philosophies of embodiment seemed metaphorical or even unintelligible to many cognitive and AI researchers at that time.

COGNITIVE SCIENCE AND EMBODIMENT

The fields of cognitive science and artificial intelligence have played a central role in the development of an embodied concept of mind and cognition. It was empirical work in cognitive science and artificial intelligence that allowed the development of a more robust and precise notion of embodied cognition to develop.

In particular, work involving mobile robots, at MIT Artifical Intelligence Laboratory and elsewhere, helped to establish the principles and concepts of embodiment and situatedness as the basis of a new approach to artificial intelligence and (later) an embodied cognitive science. For example, Brooks (1991) and his colleagues were able to get robots to perform tasks in the real world in real time that previously could only be done slowly and inflexibly, if at all. They did this by building robot bodies and robot controllers based on a design called 'subsumption architecture'. Rather than trying to graft a domain- and body-independent planning system onto a perceiving and acting robot, the subsumption architecture approach starts with an initial layer of simple perception–action mediation that implements some low-level behavior (e.g. obstacle avoidance). New behaviors (e.g. exploration) are added by adding further layers, which also mediate between perception and action in a simple way, but inhibit ('subsume') the lower layers when necessary to achieve the desired behavior. In such an architecture, there is no central locus of control, no separate planner, and no central model of the world that all processes must write to and read from in order to act appropriately. Communications between processes are not complex symbolic structures, but numerical values. What computation there is in the architecture is distributed, asynchronous, and non-hierarchical.

From an orthodox computationalist perspective, these design features have their disadvantages, but the designs of Brooks and his colleagues exploit the physical properties of the robots to overcome, or bypass, these limitations. For example, although internal communication between processes is limited, the world itself is often used as a medium for communication between the different layers and mechanisms. Much of what traditional thinking would say is required to perform a task is shown to be unnecessary if one takes advantage of regularities and information provided by the body–environment interaction.

The resulting empirical advances in robot engineering served as a springboard for a development and refinement of the notions and philosophies of embodiment. However, the failure of these approaches to quickly scale up to 'higher-level' aspects of cognition has led many to question the ability of the embodied approach to account for conceptual, abstract reasoning and representation. Kirsh (1991) correctly realizes that the issue is not one of representation per se; Brooks concedes that representation is required for some aspects of cognition. The question instead concerns having concepts: 'the ability to find an invariance across a range of concepts', as Kirsh puts it. As we move up a scale of accounting for ever more sophisticated cognitive activities, at what point must we stop limiting ourselves to designs that are tied to the particularities of the body, and begin to use designs that deploy concepts? Brooks (according to Kirsh) says 'almost never'; Kirsh disagrees, saying that much of not only reasoning and abstract thought but even perception and action must be understood in conceptual terms. Perhaps on a strict, rarefied notion of what concepts are, Kirsh is right: the explanation of concept-involving cognition can or must often go beyond what is provided by the body. But perhaps our very concept of concept needs revising (cf. Lakoff and Johnson, 1999); if we can understand how even full concept possession can be the result of being embodied in a particular way, then perhaps embodied robotics can serve as the model for far more of cognition than mere insect-like behavior.

The joint influence of embodied artificial intelligence research and philosophies and concepts of embodiment has prompted researchers to look for and formulate new forms of explanation for natural cognitive phenomena. For example, Thelen and Smith (1994) give an embodied explanation of the development of walking in infants. Rather than attempt to explain changes in gait as the result of changes in plans, rules or representations, Thelen and Smith give an elegant account that emphasizes changes in bodily factors such as the mass of the infant's leg. Such factors are then related to one another in a dynamical-systems framework, the phase transitions of which are used to explain the stage-like developments in infant walking behavior.

While the relevance of embodiment to research in mobile robotics, as discussed earlier, is obvious, some have claimed that concepts of embodiment are required for us to understand non-robotic artificial computational systems as well. Smith (1996) has argued that we will only be able truly to understand what is going on in ordinary desktop computers when we understand how the embodiment (in the sense of being located in space and time, having mass, and so on) of a computer enables it to achieve, for example, various forms of self-reference and even abstract mathematical reference. For example, it is the physically embodied 'two-ness' of a list `L=(a, b)` stored in computer memory that makes it possible for the computer to evaluate expressions such as `length(L)`.

Concepts of embodiment may also be necessary for us to theorize about the representational states of cognitive systems. The traditional means of specifying the content of, say, a belief state, is to provide a 'that' clause: a natural-language sentence that carries the same content as the belief that is being specified, for example, 'the child believes that the toy is within reach'. There are good reasons to believe that many representational contents, such as those of animals, infants, and sub-personal states, cannot be expressed in the conceptual framework of public language. To specify such contents for the purpose of a cognitive-scientific explanation, then, one may have to make essential reference to the body, and in particular the sensorimotor capabilities, of the system being explained (Cussins, 1990; Chrisley, 1995).

References

Breazeal C and Scassellati B (2000) Infant-like social interactions between a robot and a human caretaker. *Adaptive Behavior* **8**(1): 49–74.

Brooks RA (1991) Intelligence without representation. *Artificial Intelligence* **47**: 139–159.

Brooks RA and Stein LA (1994) Building brains for bodies. *Autonomous Robots* **1**: 7–25.

Chrisley R (1995) Taking embodiment seriously: non-conceptual content and robotics. In: Ford K, Glymour C and Hayes P (eds) *Android Epistemology*, pp. 141–166. Cambridge, MA: AAAI/MIT Press.

Churchland PS (1986) *Neurophilosophy: Toward a Unified Science of the Mind–Brain*. Cambridge, MA: MIT Press.

Clark A (1997) *Being There: Putting Brain, Body and World Together Again*. Cambridge, MA: MIT Press.

Clark A (1999) An embodied cognitive science? *Trends in Cognitive Science* **3**(9): 345–351.

Clark A and Chalmers D (1998) The extended mind. *Analysis* **58**: 10–23.

Cussins A (1990) The conectionist construction of concepts. In: Boden M (ed.) *The Philosophy of Artificial Intelligence*, pp. 368–440. Oxford, UK: Oxford University Press.

Dreyfus H (1992) *What Computers Still Can't Do*. Cambridge, MA: MIT Press.

Gibson JJ (1979) *The Ecological Approach to Visual Perception*. Boston, MA: Houghton Mifflin.

Heidegger M (1962) *Being and Time*. New York, NY: Harper and Row. [First published in 1927 as *Sein und Zeit*. Tübingen, Germany.]

Kirsh D (1991) Today the earwig, tomorrow man? *Artificial Intelligence* **47**: 161–184.

Lakoff G and Johnson M (1999) *Philosophy in the Flesh: The Embodied Mind and its Challenge to Western Thought*. New York, NY: Basic Books.

Maturana HR and Varela FJ (1980) *Autopoiesis and Cognition: The Realization of the Living*. Dordrecht, Netherlands: Reidel.

Merleau-Ponty M (1962) *Phenomenology of Perception*. London, UK: Routledge and Kegan Paul. [First published in 1945 as *Phénoménologie de la Perception*. Paris, France: Gallimard.]

Millikan R (1984) *Language, Thought and Other Biological Categories*. Cambridge, MA: MIT Press.

Newell A and Simon H (1976) Computer science as empirical enquiry: symbols and search. *Communications of the Association for Computing Machinery* **19**: 105–132.

Nilsson N (1984) *Shakey the Robot*. SRI Technical Note 323. Menlo Park, CA: SRI International.

Nuñez R (1999) Could the future taste purple? *Journal of Consciousness Studies* **6**(11–12): 41–60.

Piaget J (1954) *The Construction of Reality in the Child*. New York, NY: Basic Books. [First published in 1937 as *La Construction du Réel Chez l'Enfant*. Neuchâtel, Switzerland: Delachaux et Niestlé.]

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

Sharkey NE and Ziemke T (1998) A consideration of the biological and psychological foundations of autonomous robotics. *Connection Science* **10**(3–4): 361–391.

Sloman A (2001) Evolvable biologically plausible visual architectures. In: Cootes T and Taylor C (eds) *The Proceedings of the British Machine Vision Conference, Manchester, September 2001*, pp. 313–322. Manchester, UK: BMVC Press.

Smith B (1996) *On the Origin of Objects*. Cambridge, MA: MIT Press.
Thelen E and Smith L (1994) A dynamic systems approach to the development of cognition and action. Cambridge, MA: MIT Press.
von Uexküll J (1928) *Theoretische Biologie*. Berlin: Springer.
Varela FJ, Thompson E and Rosch E (1991) *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, MA: MIT Press.
Vygotsky LS (1978) *Mind in Society: The Development of Higher Psychological Processes*. Cambridge, MA: Harvard University Press. [First published 1934 in Russian.]
Ziemke T (2001) Are robots embodied? In: Balkenius C, Zlatev J, Kozima H, Dautenhahn K and Breazeal C (eds) *Proceedings of the First International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, pp. 75–83. Lund, Sweden: Lund University Cognitive Studies.

Further Reading

Brooks RA (1990) Elephants don't play chess. *Robotics and Autonomous Systems* **6**(1–2): 1–16.
Chrisley R (2000) *Artificial Intelligence: Critical Concepts*. London, UK: Routledge.
Hutchins E (1995) *Cognition in the Wild*. Cambridge, MA: MIT Press.
Pfeifer R and Scheier C (1999) *Understanding Intelligence*. Cambridge, MA: MIT Press.
Sheets-Johnstone M (1999) *The Primacy of Movement*. Amsterdam: John Benjamins.
von Uexküll J (1982) The theory of meaning. *Semiotica* **42**(1): 25–82.
Ziemke T (2001) The construction of 'reality' in the robot: constructivist perspectives on situated artificial intelligence and adaptive robotics. *Foundations of Science* **6**(1): 163–233.

Emergence

Intermediate article

Achim Stephan, University of Osnabrück, Osnabrück, Germany

CONTENTS

In ordinary language, to 'emerge' means to 'appear' or 'come into view'; but the technical use of the term is associated with features such as novelty, irreducibility, and unpredictability. The basic idea of emergence is that as systems become increasingly complex during evolution, some of them may exhibit novel properties that are neither predictable nor explainable on the basis of the laws governing the behavior of the systems' parts. Thus, complex wholes can come to have properties that are not reducible to the properties and relations of their constituents.

INTRODUCTION

During the 1990s, the term 'emergence' became widely used in such different fields as the philosophy of mind, self-organization, creativity, artificial life, dynamical systems, and connectionism. The term, however, is not used in a uniform way. It can imply novelty, unpredictability, irreducibility, and the unintended arising of systemic properties, particularly in artificial systems. Thus, it is rather controversial what the criteria are by which 'genuine' emergent phenomena should be distinguished from non-emergent phenomena. Some of the suggested criteria are very demanding, so that few, if any, properties would count as emergent. Others are inflationary, in that they count many, if not all, system properties as emergent. First of all, therefore, one should be clear about the various types of emergence. (*See* **Philosophy of Mind**; **Self-organizing Systems**; **Creativity**; **Artificial Life**; **Dynamical Systems, Philosophical Issues about**; **Connectionism**)

VARIETIES OF EMERGENTISM

Three theories among the different varieties of emergentism deserve particular interest: synchronic emergentism, diachronic emergentism, and a form of weak emergentism. In synchronic emergentism, the relationship between a system's