Levels and Explanations

Jon Opie (jon.opie@adelaide.edu.au)

Discipline of Philosophy
School of Humanities
University of Adelaide
South Australia 5005

Abstract

It is a mainstay of the philosophy of science that reduction is a relationship between theories pitched at different levels of nature. But the relevant sense of “level” is notoriously difficult to pin down. A promising recent analysis links the notion of level to the compositional relations associated with mechanistic explanation. Such relations do not order objects by scale or physical type; one and the same kind of entity can occur at several levels in a single mechanism. I will sketch this approach to levels and consider some of its implications for our understanding of the relationship between cognitive psychology and neuroscience.

Keywords: reduction; explanation; levels of nature; mechanism; mechanistic explanation

Introduction

It is a mainstay of the philosophy of science that reduction is a relationship between theories pitched at different levels of nature. So, for example, thermodynamics, which deals with certain bulk properties of matter, is said to be reducible to quantum mechanics, a theory operating at the level of the molecular constituents of matter. Likewise, the reduction of cognitive psychology to neuroscience would involve replacing a behavioural-level theory with a theory at the neuronal level. But the relevant sense of “level” is notoriously difficult to pin down, and attempts to offer an analysis that is consistent with scientific practice are plagued with difficulties.

A promising new approach (Bechtel 2008; Craver 2007) links the concept of level to the compositional relations associated with mechanistic explanation. This approach has some unorthodox consequences, including that levels are *local* rather than global (in a sense to be explained), and that one and the same kind of system or entity can occur at several levels in a mechanistic hierarchy. According to the received view, by contrast, every material system belongs to one and only one level in a monolithic, global hierarchy.

In what follows I will unpack these two competing approaches to levels, offer some grounds for preferring the mechanistic approach, and sketch a few implications for our understanding of the relationship between cognitive psychology and neuroscience.

DN Explanation

Once upon a time, to explain a phenomenon was to derive it from statements describing laws and background conditions. This view, most clearly expounded in the work of Hempel (1966), is known as the Deductive Nomological (DN) model because it treats explanation as the *deduction* of explananda from *laws* (Greek: nomoi) and boundary conditions. For example, to explain why the pressure on the walls of a gas-filled piston roughly doubles when its volume is halved, we invoke Boyle’s law. This law states that the product of pressure and volume for an ideal gas at a fixed temperature is constant ($PV = c$, where P is pressure, V is volume, and c is a constant determined by the quantity and temperature of the gas).[1] We can mathematically derive the measured change in pressure using this law, with the change in volume and the other criteria operating as background or boundary conditions.[2]

What distinguishes one scientific discipline from another on the DN model (insofar as disciplines are regarded as primarily engaged in the business of explaining some range of phenomena), is the theoretical vocabulary, ontology, and proprietary laws or theories applicable to their respective domains. The various gas laws thus serve to pick out the discipline of thermodynamics, which (among other things) deals with gases conceived as fluids with a characteristic set of macroscopic properties. Kinetic theory, by contrast, deals with the behaviour of the microscopic constituents of matter, treated either as classical particles governed by the Newtonian laws of motion, or quantum systems governed by quantum mechanics.

DN Reduction

A question that naturally arises is how disciplines so conceived relate to one another. One possibility is that they are unrelated; this seems to be the right thing to say about cosmology and economics. But many disciplines operate in overlapping or ontologically related domains, and appear to be explanatorily connected. In particular, attempting to explain laws, as opposed to phenomena, usually requires one to cross discipline boundaries. To explain Boyle’s gas law we have to invoke mechanical and statistical principles which do not belong to classical thermodynamics. If we can

---

[1] Additional criteria are implicit in the reference to an “ideal gas”. An ideal gas is a model system that is exactly described by the equation $PV = nRT$, where n is the amount of gas (measured in moles), R is the gas constant, and T is its temperature. A real gas approximates the behaviour of an ideal gas, but only under conditions of high temperature and low density.

[2] $PV = c;\ V_1 = ½V_0 \therefore P_1 = c/V_1 = 2c/V_0 = 2P_0$

derive a law of discipline A from the laws of discipline B, then we have explained that law (in the DN sense) and achieved a *partial reduction* of A to B. A *complete reduction* of A to B is a derivation, without remainder, of the laws of A from those of B. A *pure reduction* is a derivation using only premises that belong to the domain of the reducing discipline B. Pure reductions are rare. Derivations typically rely on premises about boundary conditions which can't be stated in the vocabulary of the reducing discipline. For example, to explain fluid transport across a cell wall we require not only the laws of fluid dynamics and an account of the molecular-scale interactions involved, but also postulates about the structure of cell walls.[3] A pure, complete DN reduction permits one to *eliminate* the reduced discipline, since the explanations it trades in are redundant, but such derivations are few and far between.

DN Levels

With this view of explanation and reduction in mind, Oppenheim and Putnam (1958) developed an influential scheme for organizing natural objects. They divide nature into the following six levels, where each level comprises a set of entities with a characteristic scale, and subject to the laws of an attendant discipline: Elementary Particles (Particle Physics), Atoms (Atomic Physics), Molecules (Chemistry), Cells (Cell Biology), Organisms (Biology), Society (Social Science).

A principal motivation for this scheme is to provide a basis for the putative unity of science, a unity constituted by reductive relationships among levels. The idea is that social science will eventually reduce to biology, biology to cell biology, cell biology to chemistry, and so on. DN reduction is *transitive* in this scheme: if cell biology reduces to chemistry, and chemistry to atomic physics, then cell biology also reduces to atomic physics because the derivation of cell biology piggy-backs on the derivation of chemistry, allowing one to explain cellular phenomena in atomic terms. Were this sequence of reductions to be completed, all of science would be unified by its grounding in particle physics, which would provide a common starting point for the derivation of phenomena at all levels.

Oppenheim and Putnam's approach to levels has several characteristics that continue to shape our thinking about scientific explanation. First, they envisage a *single, global level hierarchy*. The order of levels, with elementary particles at the lowest level, society at the highest, is determined by the order of the derivations required to reduce one theory to another. By assumption, every natural system belongs to one and only one level, and no system falls outside the scheme. Second, this hierarchy corresponds to an ordering in terms of both *scale* (atoms are smaller than molecules, molecules are smaller than cells, etc.) and *part-whole relations* (atoms are components of molecules, which are components of cells, and so on). Third, disciplines are *bound to levels* and vice versa. That is, each level has an associated discipline, and disciplines do not deal with entities belonging to more than one level.

This scheme has come in for a deal of criticism, both in terms of its assumptions about the nature of reduction (e.g., Fodor 1974) and its simple picture of inter-disciplinary relationships (Wimsatt 1976). Craver (2007, pp.173-7) highlights the following defects in the scheme:

i. it has significant gaps—many things, such as cellular organelles, ecosystems, and galaxies, don't have any obvious home;
ii. it doesn't tell us how to relate various kinds of things at the same scale, for example, transistors and cells;
iii. it is inconsistent with the fact that distinct disciplines often target phenomena at the same scale, e.g., crystallography and biochemistry;
iv. it is inconsistent with the fact that disciplines such as neuroscience span levels of scale and composition.

Although it might be possible to finesse Oppenheim and Putnam's scheme (for example, by adding levels), many philosophers suspect that it's the underlying approach to explanation and reduction which is really at fault here.

Mechanistic Explanation

The hegemony of the DN model of explanation started to come unstuck in the 1970s with Salmon's resurrection and defense of the causal-mechanical account of explanation (Salmon 1978, 1989). Salmon argued that scientists generally seek to discover the *causes* of some range of phenomena, rather than subsume them under laws. To explain variations in the pressure on the walls of a piston one describes the system involved — a fixed quantity of gas made up of a huge number of fast-moving particles; a closed but expandable container with rigid walls — and the way the parts of the system produce the phenomenon — pressure variations are caused by changes in the mean rate of particle-wall impacts as a result of, say, changes in the size of the container. In other words, one describes the mechanism that produces the explanandum phenomenon.

Recent work in the mechanistic tradition has focused on more fully explicating the nature of mechanisms, and the way these feature in scientific explanations (Bechtel 2006; Bunge 1997; Craver 2001, 2007; Glennan 2002; Machamer, Darden & Craver 2000; Woodward 1989). Our everyday conceptions of mechanism, which are informed by experience with relatively simple devices such as clocks and corkscrews, are potentially at odds with scientific usage. Such artifacts certainly embody mechanisms, but they are poor models for the kinds of processes scientists typically invoke to explain natural phenomena. Natural mechanisms — such as thermal conduction, protein synthesis, bacterial conjugation, perspiration, sexual selection, colonization, gravitation, and so on — need not have rigid parts, nor must they be linear, denumerable, or easily understood. Indeed, many natural phenomena result from complex, non-linear processes that defy our best attempts at analysis.

---

[3] An example of a pure reduction is the derivation of optics from electromagnetic theory (Bunge 1977, R79).

In the broadest terms, a mechanism is *a process in a material system that produces (or prevents) some change, or brings something into being*.[4] Material systems themselves are sometimes referred to as "mechanisms" when the aim is to highlight stable structural elements with the capacity to produce some systemic behavior. Thus, a clock is a "mechanism" even if it doesn't keep time. However, it is more consistent with scientific usage to identify mechanisms with productive processes. Bechtel stipulates that "[a] mechanism is a structure performing a function in virtue of its component parts, its component operations, and their organization" (2006, p.26). Likewise, Craver regards a mechanism as "a set of entities and activities organized such that they exhibit the phenomenon to be explained" (2007, p.5). The emphasis here is on the way operations or activities are organized to produce a specific outcome. To specify a mechanism one must identify: i) the relevant parts of the system, ii) the activities of those parts, and iii) how the organization of those parts and their activities gives rise to the phenomenon of interest. For example, the pumping of the heart depends on certain of its parts (ventricles, atria, valves), their activities (contraction and relaxation, opening and closing), and their spatial, temporal and causal relations (valves connect chambers and vessels, atria and ventricles contract and valves open in a specific sequence). A given system may embody more than one mechanism if various combinations of its parts and their activities produce distinct phenomena. Hearts not only pump blood, but also produce hormones involved in governing blood pressure.[5]

What distinguishes mechanistic explanation from abstract *functional analysis* (Cummins 1975) is a focus on the material basis of systemic behaviours. A functional analysis decomposes an overall system capacity into a set of sub-capacities, but is usually silent on how those sub-capacities are realized. By contrast, a mechanistic explanation reveals how some phenomenon depends on the constitution and organization of particular material entities. Organization has spatial, temporal, causal and hierarchical dimensions. The parts of a mechanism have characteristic structures, positions and arrangements. Their activities occur with specific timings, rates and durations, and in sequences or cycles which may incorporate feedback and other kinds of orchestration. Organization also has a hierarchical dimension because mechanisms typically contribute to the behaviour of superordinate systems, and are composed of subordinate systems with structure of their own. For example, the heart is part of the circulatory and respiratory systems, and its parts (atria, ventricles, valves) are composed of cells of various types whose organized activity produces atrial contraction, distension of the valves, and so on. (Bechtel 2008, pp.10-17; Craver 2001)

Mechanistic Levels

According to the causal-mechanical account of explanation, to explain some function ϕ of a material system S is to identify and describe a process M in S that produces ϕ. If S is an *open system* (a system that exchanges matter and/or energy with its environment) M will typically be subject to outside influences, which means that fully specifying ϕ will involve specifying the environmental conditions that bear on M. For example, blood is part of the heart's environment, and the condition of the blood (e.g., its level of oxygenation) has an impact on the heart's activity. If we wish to explain the pumping of the heart, we must first establish how variation in the properties of the blood (pH, pO_2, viscosity, etc.) affects this behaviour. Thus, mechanistic explanation almost always encompasses at least two distinct *levels of organization*: i) the level of the system S and its environment, including any relevant containing systems[6]; and ii) the level of the active parts of S (Bechtel 2008, p.148). Since mechanistic explanation iterates, deeper explanations will reveal further levels of organization in the structure of the parts of S.

Bechtel (2008) and Craver (2007) argue that material organization is the basis of a level hierarchy suited to mechanistic explanation. As remarked, every mechanistic explanation involves at least two levels of organization, and it is rare not to refer to the organization of one or more subsystems of the primary system. What belong at each level are the active entities whose organization produces the explanandum phenomenon at the next level. Thus, atria, ventricles and valves are at one level, because their organized activity constitutes the pumping of the heart; the heart, blood and vessels a next higher level because they collectively act to circulate nutrients, hormones and gases around the body. The relation between levels is one of *composition*. A mechanism at one level is composed of active entities at the next lowest level of organization. Those entities are themselves composed of active entities at a yet lower level, and so on.[7] Composition implies spatial and temporal *containment* (that is, the parts and activities of a mechanism cannot exceed the size and duration of the whole process), but is more than this. A mechanism is not a mere bag of parts and activities, but a *physical gestalt* (Kohler 1920); a material whole that is constituted by the organization of its parts and their activities.

There are a number of significant respects in which this mechanistic approach to levels differs from Oppenheim and Putnam's scheme. First, mechanistic levels are always *local*

---

[4] This formulation is adapted from Bunge (1997). A *material system* is a bundle of real things that behave in some respects as a unit by virtue of their interactions or bonding. Atoms, crystals, synapses, transistors, cells, organisms, families, firms and galaxies are material systems, which are to be contrasted with *conceptual systems*, such as theories and classifications. (ibid, p.415)

[5] Excessive stretching of the atria and ventricles causes the release of two peptides that lower blood pressure by relaxing arterioles, inhibiting the secretion of renin and aldosterone, and inhibiting the reabsorption of sodium ions by the kidneys.

[6] The respiratory system, for example, in the case of the heart. The heart is both a *component* and a *target* of this system because it requires a supply of oxygen in order to function.

[7] The exception is entities such as electrons, which so far as we know have no internal structure.

(Craver 2007, pp.192-3). They do not comprise a single, global hierarchy, because such levels are defined relative to the current explanatory target. To explain how nutrients are transported around the body one unpacks the circulatory system into heart, blood, veins and arteries. These all belong at a single level for the purpose of explaining circulation. But the heart also has a role in controlling blood pressure, and finds itself in rather different company in that context (see footnote 5). Moreover, to explain the heart's endocrine function one must decompose it into a set of active parts distinct from those that explain its ability to pump blood. The heart is thus a nexus for numerous mechanisms which overlap in that system, but ascend and descend along very different composition-decomposition pathways. Instead of a single level hierarchy, there are as many hierarchies as there are distinct phenomena to which the heart contributes.

Second, the order associated with a mechanistic hierarchy is *partial*, not total (Bechtel 2008, p.147; Craver 2007, p.191). Some things are related by level membership, others are not. Even though the heart is composed of cells, osteocytes (bone cells) do not appear at a lower level than the heart in the mechanism of circulation, because they are not part of that mechanism. Likewise, there simply is no answer to questions about the ordering of transistors and cells, galaxies and ecosystems, atoms and economies, except where these systems figure in some common multi-level mechanism. It is also important to be aware that level relations only apply to entities that belong to one and the same mechanism; they do not apply to entity types. My pulmonary artery is part of a mechanism of circulation that includes my heart, and belongs at the same level as the heart in that mechanism, but it doesn't contribute to circulation in anyone else so there is no question about its relationship to hearts in general.

Third, entities of a given kind can appear at more than one level in the same mechanism. For example, free protons are part of an acid-sensing mechanism in sensory neurons. Such neurons contain proton-gated ion channels, which respond to proton binding by opening a pore in the cell membrane, thereby inducing a sensory signal (Waldmann et al, 1997). Protons are at the same level as ion channels in this context, because the two must interact for acid-sensing to occur. However, since protons are a basic building block of matter, protons are also constituents of ion channels themselves. Thus protons appear on at least two organizational levels in this mechanism.[8]

The mechanistic approach to levels is at odds with some very entrenched intuitions. It is by now second-nature to associate levels with well-defined physical types, each having a fixed place in a monolithic, global hierarchy. The simplicity of this picture is one of its chief attractions, but also its principal flaw, because it fails to reflect the complexities of scientific practice. Most disciplines pay little attention to boundaries defined by scale or type, seeking out whichever entities, activities and forms of organization illuminate their explanatory targets. Indeed, it is reasonable to ask why anyone would expect a global level hierarchy to make sense of the very different explanatory demands of, say, cell biology and plasma physics. Both disciplines investigate processes whose ultimate constituents are protons, neutrons and electrons, but organized in such vastly different ways that the two domains essentially have no structures in common at any scale.[9]

Mechanistic Reduction

There is both reduction and emergence here. Mechanisms innocently emerge from the organized activity of lower-level entities, and have effects that their constituents in isolation lack.[10] The existence of this kind of novelty is not miraculous, although it can be difficult to understand, especially when the organization involved is complex or non-linear. Mechanistic explanation is reductive in the sense that it reveals the connection between phenomena and the lower level entities that produce them. Yet it doesn't thereby eliminate high-level phenomena, which depend not only on the constituents of their mechanisms, but also on the way those constituents are organized. Nor does mechanistic explanation eliminate disciplines, in the sense of making them redundant. Disciplines span levels, and must work together to discover the multi-level mechanisms of complex phenomena (such as cognition). The unity of science, such as it is, does not depend on reductive relationships among disciplines, but on the shared ambition to reveal the hidden structure of nature. It is not an "imperial unity", in which one discipline dominates all others, but a "federal unity"; a polity founded on the principles of autonomy and cooperation (Auyang 1998).

The Fate of Cognitive Psychology

If we accept the DN account, the ultimate fate of cognitive psychology will be elimination by mature neuroscience. The latter, operating at the "neuronal level", will provide a pure, complete reduction of cognitive psychology, rendering the latter redundant. But the DN model of science is suspect in all kinds of ways, not least because its picture of levels is so at odds with scientific practice. In this final section I will consider what alternative fate awaits cognitive psychology on the mechanistic account of explanation.

Important critical reactions to the radical reductionism of the DN model arose prior to the re-emergence of the mechanistic perspective. Fodor (1974) argued that the special sciences (everything except fundamental physics) do not reduce to physics, because the properties subsumed by their laws are *multiply realizable*. The property of being a

---

[8] See Bechtel (2008, p.147) for a similar example.

[9] A plasma is a high-temperature gas of unbound protons, neutrons and electrons, which do not combine to form atoms because of their high kinetic energy.

[10] A property, process or entity is "innocently emergent" if it can be explained in terms of lower-level entities and their activities. A "radically emergent" phenomenon, by contrast, isn't determined by or explicable in terms of some underlying material substrate. Such phenomena are inherently mysterious. Both life and consciousness are sometimes regarded as radically emergent.

cell, for example, doesn't have a unique decomposition into the properties of elementary particles. From the perspective of microphysics, cells do not show any interesting commonalities. But cells, tectonic plates, airfoils, galaxies, and so on, are the very stuff of the special sciences. Even if such objects are nothing but agglomerations of fundamental particles, we still require the laws of the special sciences in order to frame useful generalizations about them.

Multiple-realizability thus provides grounds for rejecting the reduction of cognitive psychology to microphysics. However, many cognitive scientists go further, claiming that psychology and neuroscience address distinct levels of nature: the "cognitive" and "neural" levels, respectively. On this view, the role of cognitive psychology is to characterize cognition in a substrate-neutral way, as a set of formally specified input-output or information-processing relations. Neuroscience will eventually show us how such cognitive functions are realized in the case of neurozoans[11], but the presumed autonomy of the cognitive level is consistent with there being other realizations of the self-same functions in non-neural or even non-biological hardware. Indeed, the possibility of radically different realizations of cognitive functions suggests that cognitive explanation is a matter of causal or functional analysis[12] pure and simple.

We need not follow functionalists down this road. To begin with, multiple realizability is not ruled out by a focus on structure rather than function. Consider the humble *airfoil*. An airfoil generates lift when it moves through the air because its upper surface has greater curvature than its lower surface.[13] Airfoils are thus distinguished from other kinds of things not merely by their function (*what they do*), but also, and more fundamentally, by their shape (*what they are*). Many things can generate lift, e.g., rocket engines, but not all of them are airfoils. What unifies airfoils, as a kind, is a common *structural* property. And that property is multiply realizable. An airfoil can be built using sheet metal, fiberglass, canvas stretched across balsa, or any other materials rigid enough to maintain the correct shape.

Mechanistic explanation, unlike causal analysis, reveals both functional and structural forms of multiple realizability. A mechanistic explanation of flight, for example, not only describes how the various parts of an aircraft work together (the engine produces thrust which causes movement of air across the wings, the joystick changes the orientation of ailerons and elevators which in turn alter the attitude of the aircraft), but also how the structures of those parts are responsible for what they do (the differential curvature of each wing causes air pressure to be greater on its lower surface than its upper surface). Flight control is multiply realizable because there are numerous ways of establishing a causal relationship between pilot and control surfaces — using mechanical, hydraulic, or "fly-by-wire" systems, for example. But, as already remarked, the material and structural properties of components such as wings are also multiply realizable. My point is that if we want to argue from the existence of multiple realizability to the failure of DN-style reduction, we are not thereby forced to regard the explanation of macroscopic phenomena as mere causal analysis.

One might accept this general point about explanation, but nevertheless regard cognition as a special case. This view is typical of those who treat cognition as *digital computation*: the rule-governed manipulation of in-the-head symbols. Symbols are representations that bear an arbitrary physical relationship to the things they represent. This means that symbols don't carry their meaning intrinsically, but acquire meaning by virtue of how they are manipulated. Arithmetic is the familiar case. Numerals represent numbers not because of their physical form, but because we operate on them in accordance with rules that support a numerical interpretation. Likewise, what guarantees the semantic coherence of a digital computer is the set of physically-implemented rules that govern its operations.[14] All of this has been taken to suggest that computation, and by extension cognition, has no essential reference to the substrate in which it is implemented. So long as the transitions between symbols are well-behaved from a semantic perspective, their material properties are of little consequence. And this has led many theorists to accept the existence of a "cognitive level" that abstracts from structure in a very profound way. It is no wonder that neuroscience is so often treated like the handmaid of cognitive psychology — useful for discovering how cognitive functions are implemented in neural systems, but with nothing to say about their fundamental nature.

Although still influential, this approach to cognition is no longer the only game in town. In the 1980s connectionists began to rebuild the foundations of cognitive psychology on the assumption that the structural properties of the brain are not mere implementation details, but core determinants of human cognition (Rumelhart, McClelland et al 1987). More recently, accounts of cognitive representation better suited to this focus on material structure have started to emerge (see, e.g., Churchland 2001; Cummins 1996; O'Brien & Opie 2004). O'Brien and Opie (2006) argue that neural representation depends on the existence of a physical analogy between systems of neural states and their target domains. Two systems are analogous when there is a relation-preserving mapping between them. For example, a tree's growth rings are analogous to the climatic conditions in which the tree developed. Plentiful seasons produce wide growth rings, whereas drought years produce comparatively narrow rings. The relative thickness of growth rings therefore reflects the variations in the seasons, and permits us to treat the former as a record of the latter.

---

[11] A *neurozoan* is any organism that has a nervous system of some kind (not necessarily a brain).

[12] Functional and causal analysis aren't equivalent in general, except when it comes to explaining symbolic (digital) computation.

[13] The lift generated by an airfoil, such as a wing, doesn't just depend on its shape, but also on its surface area and angle of attack. Nonetheless, the term "airfoil" is usually reserved for objects with differential curvature of their two principal surfaces.

[14] See O'Brien & Opie 2009 for further discussion.

Connectionist studies suggest that neural networks acquire meaning in a similar way. In a face recognition network developed by Garrison Cottrell (reported in Churchland 1995, pp.35-55) the similarity relations between hidden-unit activation patterns were found to correspond to facial relationships. Plotting hidden unit activity as points in “activation space”, with one axis for each hidden unit, reveals distinct regions for male and female faces, and sub-regions associated with individual faces. Similar faces are represented by patterns that are nearby in activation space, whereas dissimilar faces are represented by patterns that are correspondingly further apart.[15] In other words, the hidden unit activation patterns generated by Cottrell’s network are physically analogous to (features of) the set of faces on which the network has been trained.

Assuming that this case generalizes, it appears that neural representations do not acquire their meanings by virtue of how they are manipulated, but because they support suitable physical analogies. Such analogies underpin the capacity of neural networks to engage in computation. Cottrell’s network can make sensible facial discriminations, such as correctly assigning gender and identity to its inputs, because of the analogy between its system of hidden-unit activation patterns and the faces they represent. Systems that operate in this way are called *analog computers*. Just like digital computers, analog computers are built from a variety of materials and are governed by diverse physical principles. However, unlike digital computers, analog computers can’t be adequately described in purely causal terms, that is, in terms that make no reference to their material and structural properties. Quite the reverse. It is the manner in which an analog computer is realized that determines the kinds of things it can represent, and sustains the semantic coherence of its internal processes. (O’Brien & Opie 2009) Thus, if human cognition has a significant analog dimension, the realization-independent cognitive level envisaged by many theorists is illusory.

The upshot of all this is that commitment to the existence of a cognitive level of analysis isn’t a consequence of the computational perspective per se, but of buying into the assumptions of the classical (digital) approach to cognition. If our best account of cognition turns out to be an analog theory, we need not abandon the mechanistic style of explanation that works so well in other sciences. But neither need we accept the radical reductionism of the DN model. Cognitive phenomena don’t cease to exist once we begin to understand their mechanisms. On the contrary, it is by revealing the mechanisms of cognition that we find out what cognition really *is*. We know that cognition in neurozoans is crucially dependent on structural properties at a number of scales, from the shapes of ion channels in axons, to the organization of inputs on dendritic trees, the laminar structure of the thalamus and its homologues, the connectivity patterns within the eye and other input systems, and the general layout of the mammalian brain. Moreover, there are strong causal dependencies between processes at these various scales. A change in the molecular structure of a neurotransmitter, for example, can have significant ramifications for the behaviour of the whole organism. There is no neat isolation of processes occurring at different scales, as suggested by the DN model, nor are there distinctive neural and cognitive levels.

What, then, of cognitive psychology? As I remarked above, DN levels were conceived with disciplinary relationships in mind. They are meant to bind disciplines to phenomena at a characteristic scale, and unify them (in the limit) by derivation from a common basic science. Mechanistic levels, by contrast, reflect the complex causal structure of the world. They are local and partial, crisscrossing groupings defined by scale or physical type, and they *are no respecters of discipline boundaries*. One might think this undermines the autonomy of cognitive psychology. If there is no cognitive level, what exactly is the subject matter of this discipline? But again, this view only betrays the powerful grip of the DN model on our thinking. There may be no cognitive level, but the world is overflowing with cognitive phenomena. At its most general, cognition comprises those information-sensitive processes whereby organisms track their environment and orchestrate their own survival and reproduction. Such processes come in a multitude of forms, arguably spanning everything from multi-organism societies to individual bacteria.[16] Cognitive psychology is and will remain a core discipline in pursuit of this quarry, which is no less significant for our growing appreciation of the diversity of cognitive phenomena, nor the long list of sciences (neuroscience, genetics, ecology, microbiology, and physics, to name a few) which today contribute to unpacking its mechanisms.

Cognition no doubt has many features that generalize across organisms, some of which may be reproducible in artificial systems. But it is a conceit to imagine that we can first determine the function, and then figure out its realization. Structure and function are deeply intertwined in biological systems. To understand cognition we must bring multiple disciplines to bear in a process Sunny Auyang calls “synthetic analysis” (1998), repeatedly taking cognitive systems apart and putting them together again (both conceptually and experimentally). By this means we will finesse our understanding of both cognition and its mechanisms. Cognitive psychology and neuroscience are partners in this enterprise, neither of which is dispensable either in practice or in principle.

References

Auyang, S. Y. (1998). *Foundations of Complex-System Theories*, Cambridge: Cambridge University Press

Bechtel, W. (2006). *Discovering Cell Mechanisms: The Creation of Modern Cell Biology*, New York: Cambridge University Press.

---

[15] See O’Brien & Opie 2001 for further details.

[16] For discussions of bacterial smarts see Shapiro 2007 and Stock et al 2002.

Bechtel, W. (2008). *Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience*, New York: Lawrence Erlbaum

Bunge, M. (1977). Levels and reduction. *American Journal of Physiology, 233(3),* R75-R82.

Bunge, M. (1997). Mechanism and explanation. *Philosophy of the Social Sciences, 27(4),* 410-65.

Churchland, P. M. (1995). *The Engine of Reason, the Seat of the Soul.* MIT Press.

Churchland, P. (2001). Neurosemantics: On the Mapping of Minds and the Portrayal of Worlds. In K. E. White (ed.), *The Emergence of Mind*. Milan, Fondazione Carlo Elba, pp. 117-47.

Craver, C. F. (2001). Roles functions, Mechanisms, and Hierarchy. *Philosophy of Science, 68,* 31-55.

Craver, C. F. (2007). *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience.* New York: Oxford University Press.

Cummins, R. (1975). Functional Analysis. *The Journal of Philosophy, 72(20), 741-65.*

Cummins, R. (1996). *Representations, Targets, and Attitudes*. MIT Press.

Fodor, J. A. (1974). Special sciences. *Synthese, 28,* 97-11.

Glennan, S. (2002). Rethinking Mechanistic Explanation. *Philosophy of Science, 69*: S342-53.

Hempel, C. (1966). Philosophy of Natural Science. Prentice Hall, pp. 47-67.

Köhler, W. (1920) Die physischen Gestalten in Ruhe und im stationären Zustand. Braunschweig: Vieweg und Sohn.

Machamer, P., Darden, L. and Craver, C. (2000). Thinking about Mechanisms. *Philosophy of Science, 67*, 1–25.

O'Brien, G., & Opie, J. (2001). Connectionist vehicles, structural resemblance, and the phenomenal mind. *Communication and Cognition 34*, 13–38.

O'Brien, G. & Opie, J. (2004). Notes toward a structuralist theory of mental representation. In H. Clapin, P. Staines and P. Slezak (eds.), *Representation in Mind: New Approaches to Mental Representation*. Elsevier.

O'Brien, G. & Opie, J. (2006). How do Connectionist Networks Compute? *Cognitive Processing 7*, 30-41.

O'Brien, G. & Opie, J. (2009). The Role of Representation in Computation. *Cognitive Processing, 10*, 53-62.

Oppenheim, P. and Putnam, H. (1958). Unity of Science as a Working Hypothesis. In H. Feigl, M. Scriven, and G. Maxwell (Eds.), *Minnesota Studies in the Philosophy of Science, 2*. Minneapolis: University of Minneapolis Press.

Rumelhart, D., McClelland, J., & the PDP Research Group (1987). *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, vol. 1: Foundations.* Cambridge, MA: MIT Press.

Salmon, W. (1978). Why ask "Why?"?, *Proceedings and Addresses of the American Philosophical Association, 51*, 683-705.

Salmon, W. (1989). Four Decades of Scientific Explanation. In P. Kitcher and W. Salmon (eds.), *Minnesota Studies in the Philosophy of Science, 13*. Minneapolis: University of Minneapolis Press.

Shapiro, J. A. (2007). Bacteria are small but not stupid: cognition, natural genetic engineering and socio-bacteriology. *Studies in History and Philosophy of Biological and Biomedical Sciences, 38,* 807-19.

Stock, J. B., Levit, M. N. & Wolanin, P. M. (2002). Information processing in bacterial chemotaxis. *Science's Signal Transduction Knowledge Environment* (STKE) http://www.stke.org/cgi/content/full/OC_sigtrans;2002/132/pe25.

Waldmann, R. et al (1997). A proton-gated cation channel involved in acid-sensing. *Nature, 386(6621),* 173-7.

Wimsatt, W. (1976). Reductionism, Levels of Organization, and the Mind-Body Problem. In G. Globus, I. Savodnik, and G. Maxwell (eds.), *Consciousness and the Brain*. New York: Plenum Press, pp. 199-267.

Woodward, J. (1989) The causal mechanical model of explanation. In P. Kitcher and W. Salmon (eds.), *Minnesota Studies in the Philosophy of Science, 13*. Minneapolis: University of Minneapolis Press.