Euro Jnl Phil Sci (2012) 2:453–480
DOI 10.1007/s13194-012-0053-y

ORIGINAL PAPER IN PHILOSOPHY OF SCIENCE

The relationship between psychological capacities and neurobiological activities

Gregory Johnson

Received: 17 May 2011 / Accepted: 10 April 2012 / Published online: 9 May 2012

Abstract This paper addresses the relationship between psychological capacities, as they are understood within cognitive psychology, and neurobiological activities. First, Lycan's (1987) account of this relationship is examined and certain problems with his account are explained. According to Lycan, psychological capacities occupy a higher level than neurobiological activities in a hierarchy of levels of nature, and psychological entities can be decomposed into neurobiological entities. After discussing some problems with Lycan's account, a similar, more recent account built around levels of mechanisms is examined (Craver 2007). In the second half of this paper, an alternative is laid out. This new account uses levels of organization and levels of explanation to create a two-dimensional model. Psychological capacities occupy a high level of explanation relative to the cellular and molecular levels of organization. As a result, according to this model, psychological capacities are a particular way of describing the activities that occur at the cellular and molecular levels of organization.

Keywords Psychological capacity · Levels of organization · Levels of explanation

1 Introduction

This paper proposes a model for understanding the relationship between psychological capacities and neurobiological activities. Here *psychological*

G. Johnson (✉)
Department of English and Philosophy, Drexel University, Philadelphia, PA, USA
e-mail: gregory.s.johnson@drexel.edu

capacities, or just *the psychological*, will refer to the capacities that are defined and explained within the domain of cognitive psychology. Typically, *psychological capacity* refers to a certain class of abilities that can be addressed from a number of perspectives besides cognitive psychology, for example, from within neurobiology or cognitive science. In this paper, however, *psychological capacity* will only refer to capacities understood in psychological terms (i.e., described in psychological language). And *neurobiological activities* will refer to the processes—and the entities that participate in those processes—that occur within the brain.

Herein I am not addressing whether psychology is autonomous from neurobiology or whether psychology can be reduced to neurobiology. These are issues that concern the ultimate status of the science, psychology. The topic in this paper is the relationship between psychological things and neurobiological things, and I begin by assuming that they are related. The question then is, how can each be understood such that it is explicitly consistent with the other? The answer to this question will indicate how they are related.

The standard way of addressing this problem is to say that the relationship between psychological capacities and neurobiological activities is similar to the relationship that exists between the entities described in chemistry and the entities that belong to basic physics. In the latter case, to put it simply, physics is more basic than chemistry, and the entities that belong to the domain of chemistry are composed of the entities found in physics.

William Lycan has developed a detailed account of the relationship between psychological capacities and neurobiological acitivites that utilizes this chemistry-physics model and the hierarchical organization on which it depends (1987, see also 1981, 1991). According to his account, each psychological capacity can be decomposed into component parts. Each of these components can then be similarly decomposed, and if the process is continued, the decomposition will eventually yield neurobiological and neurochemical entities—and still lower level entities if it is further continued.[1]

Lycan is correct that the hierarchical organization found in nature has a central role in explaining the relationship between psychological capacities and neurobiological activities. But his account contains certain problems that are a result of locating psychological capacities directly on this hierarchy. As an alternative, I outline a model that utilizes a similar hierarchy, but does not locate psychological capacities (as they are described within cognitive psychology) directly on that hierarchy. An additional set of levels, a hierarchy of levels of explanation, is needed to correctly situate the psychological with respect to the neurobiological.

[1]This account is similar in some respects to one developed by Dennett (1978), to whom Lycan attributes the basic idea.

2 Lycan: the continuity of levels of nature

2.1 Three central features

Lycan's account has three central features. The first is the claim that psychological capacities should be understood in teleological terms. Thus, for Lycan, the function of a psychological capacity is the purpose of that capacity with the understanding that this purpose has been established by evolutionary processes.[2]

The second feature is what Lycan calls *homuncular functionalism*. This is the idea that psychological capacities can be decomposed into simpler components, each of which can be thought of as a homunculus that carries out a specific task. Together, this collection of homunculi performs the psychological capacity. Each one of these homunculi, which are basically sub-capacities, can then be explained by another set of homunculi (i.e., sub-sub-capacities), and so on. As Lycan says,

> To characterize the psychologists' quest in the way I have is to see them as first noting some intentionally or otherwise psychologically characterized abilities of the human subject at the level of data or phenomena, and positing—as theoretical entities—the homunculi or sub-personal agencies that are needed to explain the subject's having those abilities. Then the psychologists posit further, smaller homunculi in order to explain the previously posited molar behavior of the original homunculi, etc., etc. (1987, p. 40)

The motivation for analyzing psychological capacities in this way is explanatory. Each of these steps to "smaller" homunculi introduces new sub-capacities that are, individually, contributing less to the performance of the psychological capacity itself. The collection of simpler components explains how the higher level function is performed.

The third feature of Lycan's account is a hierarchy of levels of nature. The basic idea here is familiar. Compositionally, the entities found in nature can be put into a hierarchy. For example, organisms are composed of cells, cells are

[2]The reference to evolutionary processes is not defended by Lycan, but it does position his view within a certain context. He says initially, "I hope, and am inclined to believe, that the teleological characterizations that Homunctionalism [i.e., homuncular functionalism] requires can be independently explained in evolutionary terms" (1987, p. 43). And later, "If teleological characterizations are themselves explicated in evolutionary terms, then our capacities for mental states themselves become more readily explicable by final cause; it is more obvious why we have pains, beliefs, desires, and so on" (1987, p. 45). To illustrate, he offers, "Why does pain hurt? Why could we not have a damage-signaling and repair-instigating system that was not uncomfortable? The answer is simple. Suppose I had just such a system, like the red warning light on my auto engine. Just as I habitually though irrationally ignore the warning light and vaguely hope it will go away, I would ignore a personal warning light if it did not intrinsically provide me with an urgent motive to do something about it" (1987, p. 138, n. 17).

composed of molecules, molecules are composed of atoms, and so on.[3] In this hierarchy, as Lycan points out, the constituents of any particular level can be explained functionally, by referring to their purpose, or structurally, by referring to their parts, which occupy a lower level in the hierarchy. The entities at this lower level can also be characterized functionally or, by dropping down another level, structurally. Thus, "function" and "structure" occur throughout the levels of nature; the labels are simply relative to a place on the hierarchy.

2.2 Lycan's account, part 1

The three features just discussed are used by Lycan to generate an account of the relationship between psychological capacities and neurobiology activities. The first step is determining how teleology is related to homuncular functionalism. To accomplish this, Lycan suggests that "teleologicalness of characterizations is a matter of degree" (1987, p. 43). A psychological capacity that is at the top of a hierarchy created by homuncular functionalism is characterized in very robust teleological terms. As the decomposition proceeds downward, however, the characterizations, while still functional, become progressively less teleological. As a result, the jobs that the smallest homunculi perform are understood in mechanical terms, not in teleological terms (1987, p. 44).[4]

The next step is combining homuncular functionalism and the hierarchy of levels of nature. Since they both have a levels structure, they can, in theory, be combined. And Lycan claims that "for single organisms, degrees of teleologicalness of characterization correspond rather nicely to levels of nature" (1987, p. 45). By this he means that the functions of the entities found at higher levels of nature are the result of evolutionary processes, while the functions of the entities that occur at lower levels are less likely to be thought of in terms of purpose or design. Of course, the hierarchy created by homunucular functionalism also has this feature: higher level homunculi are more teleological than lower level ones. Thus, Lycan is able to claim that homuncular functionalism, which was originally a resource dedicated to understanding psychological capacities, is actually a part of the hierarchy of levels of nature.

[3] Lycan does not provide a very thorough explanation of the significant features of a level of nature. His definition is "levels [of nature] are nexus of interesting lawlike generalizations, and are individuated according to the types of generalizations involved" (1987, p. 38). Referring to "nexus of interesting lawlike generalizations" does not appear to be (and may not be intended as) a precise criterion. In any case, this definition presumably allows Lycan to draw on anything that falls within the purview of science.

[4] Lycan uses an example to illustrate this notion of degrees of teleologicalness that is discussed in the next section of this paper (see Fig. 2). In that example, as we move down this hierarchy created by homuncular functionalism from a *face recognizer* to an *analyzer* to a *scanner* to a *light meter* and to *photosensitive chemicals*, each characterization is, Lycan suggests, less teleological than the previous one.

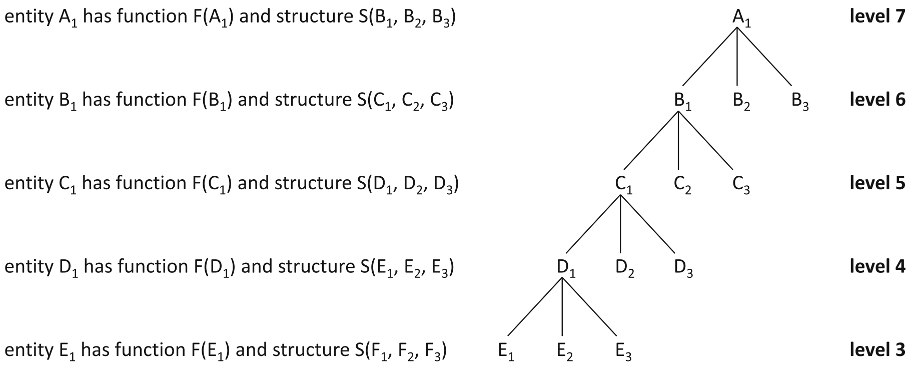

Fig. 1 On the *left* are the functional and structural characterizations of the entities at each of several levels of nature. See text for further explanation. The same information is represented on the *right* as a hierarchy

We now have an account of the relationship between psychological capacities and neurobiological activities. The psychological, which are "highly teleological characterizations" at the top of the homuncular functionalism hierarchy, occupy a relatively high level of nature. As characterizations are offered that are less teleological—as psychological capacities are decomposed via homuncular functionalism—the move is made to lower levels of nature, eventually reaching a level occupied by neurobiological entities. In outline, this account is represented in Fig. 1.

Because each entity occurs at a level of nature, each can be characterized functionally or structurally, and how this is done is relative to the level of nature that is being investigated. For instance, at level 5 in the figure, entity C_1 has some particular function: function $F(C_1)$, and C_1's structure is made up of the entities that it is composed of: D_1, D_2, and D_3, which are found at the next level down. D_1, D_2, and D_3 are characterized functionally at this lower level of nature, and their structures are explained by the entities that occupy level 3.

With respect to the psychological and the neurobiological, psychological capacities are found at a higher level of nature, say level 7 in Fig. 1, while neurobiological activities occupy one of the lower levels.[5] Thus, $F(A_1)$ might be the ability, described in functional terms, to comprehend language. And, according to this model, the structure that has this ability is found at level 6. Each of these entities at level 6, B_1, B_2 and B_3, has its own function, which is a sub-capacity of language comprehension. Down at level 3 meanwhile, entity E_1 is, let's say, a particular protein found in neurons. Thus, according to Lycan,

[5]Lycan locates psychological capacities just below the level of the organism itself (the organism is an "institution" containing the psychological capacities (1987, 40)), but several levels above the neurobiological. To locate the psychological directly at the neuroanatomical level is, he says, "implausible" (1987, p. 59).

the relationship between psychological entities and neurobiological ones is this composition relationship that can be observed by looking across levels of nature.

2.3 A face recognizer

The example that Lycan uses to illustrate his model is the capacity for face recognition (1987, pp. 43–44). And although the details of the example are fictitious, it does show how, on his account, a psychological capacity is related to the relevant neurobiological activities.

One way that face recognition might be carried out, Lycan suggests, is by implementing the following process. A particular program is engaged only when the input is a face viewed from either the right, left, or straight-forward profile. A *viewpoint locator* identifies which profile it is, and an *analyzer* then codes the relevant features based on what the viewpoint locator has told it. This information is then passed on to a *librarian* who compares it with information that is stored in memory. If a match is found, the librarian will be able to look at the identification tag attached to the information stored in the memory. The librarian will then pass the identification tag on to the *public relations officer* who will be able to give instructions to the appropriate *motor subroutine*, which will then verbally enunciate a name (Fig. 2).

This description of these entities and their activities is one level below the teleologically defined face recognizer: a device that's purpose is to recognize faces. If information about any of these particular sub-capacities is called for, then each can itself be decomposed. For instance, the analyzer might be a *projector* that projects a grid onto the profile and a *scanner* that encodes each square of the grid in a binary code to be passed on to some sub-capacity of the

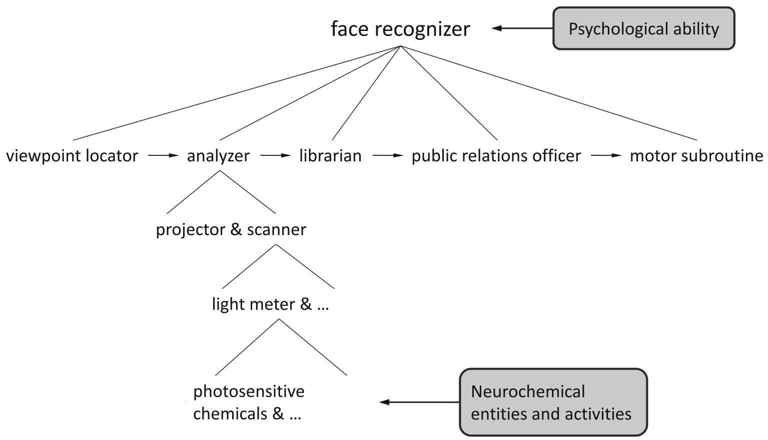

Fig. 2 Lycan's example of the decomposition of a face recognizer. These levels are, hypothetically, part of the hierarchy of levels of nature

librarian. If more information about the scanner is required, the scanner can be decomposed into a *light meter* and some way of reporting "0" or "1" based on the degree of darkness. The light meter can then be explained by invoking photosensitive chemicals, and so on. Thus, in several steps the face recognizer has been decomposed into relatively non-teleological chemical entities.

While this example is a good illustration of homuncular functionalism, it is also supposed to be consistent with the hierarchy of levels of nature. Each of the "homunculi"—e.g., the librarian, the projector, the scanner, and the lightmeter—are located at some level of nature. The face recognizer occupies the highest of these levels. One level below the face recognizer are the viewpoint locator, analyzer, librarian, public relations officer, and motor subroutine, which together make up the structure of the face recognizer. The entities at this lower level of nature have functional characterizations, and they are characterized structurally in terms of the entities found at the next lowest level. The function of the analyzer, for instance, is to encode the information provided to it, and the analyzer's structure is a projector and scanner, which are found one more level down.

Now Lycan is not actually suggesting that these are entities that occupy levels of nature or that the process he has described is how this psychological capacity, face recognition, is carried out. He does, however, believe that legitimate descriptions of the decomposition of psychological capacities into neurobiological (and lower level) entities can be produced using this format.

2.4 Lycan's account, part 2

Although this account has many attractive features, there are reasons to be skeptical about the way that it construes the relationship between psychological capacities and neurobiological activities. The source of the problem is that Lycan employs a single type of hierarchy, his levels of nature. Because he takes it that the hierarchy created by homuncular functionalism is part of the hierarchy of levels of nature, everything that is included in his account has to have a place on some level of nature. Entities occur at different places on the hierarchy, but, according to Lycan, all are found in nature and all can be characterized both functionally and structurally. But this claim is worth resisting. In this section, I will argue that the hierarchy created by homuncular functionalism cannot be mapped onto the hierarchy of levels of nature. If this is correct, then Lycan's account fails.

We can begin by noting that psychological capacities and the sub-capacities that subserve them are identified and individuated by their functional properties, not by any physical characteristics. This is hardly controversial, and it is explicitly part of Lycan's homuncular functionalism. He says at one point:

> What is functionalist about HF [i.e., homuncular functionalism] is obviously that it identifies a mental item in terms of that item's functional relations to other mental items (and to sensory inputs and motor outputs); nothing new there. (1991, p. 264)

But one result of Lycan's full account is that every entity in the hierarchy created by the decomposition of a psychological capacity has to have functional (i.e., relational) as well as intrinsic (i.e., non-relational) properties. (Again, this is because in his account the levels created by employing homuncular functionalism coincide with levels of nature.) There must be, on the one hand, *the component's function*, but also *the physical material that it is made of*.[6]

Now Lycan can use the face recognizer as an example of a psychological capacity because this "entity" can be defined in terms of its function—and likewise for the analyzer, librarian, public relations officer and so on. When Lycan introduces his example, the description only consists of functionally defined components that perform the face recognition task: first, the face recognizer itself and then the sub-capacities that collectively perform the face recognition task, the viewpoint locator, analyzer, librarian, and public relations officer.

The same can be seen by looking at a more serious example. The psychologists Craig Smith and Leslie Kirby describe part of the emotion process this way:

> A central feature of this model is the existence of what we call "appraisal detectors." These detectors continuously monitor for, and are responsive to, appraisal information from multiple sources. The appraisal information they detect determines the person's emotional state. ... [T]hey detect the appraisal information that is generated from different modes of processing. This detected information is then combined into an integrated appraisal that initiates processes to generate the various components of the emotional response, including an organized pattern of physiological activity, the action tendency, and the subjective feeling state. (2000, pp. 92–93)

The *appraisal detector* is a component (i.e., a sub-capacity) that, according to Smith and Kirby, has a role in the process that brings about an emotion response. But while this component has functional properties, Smith and Kirby do not identify—they could very well be completely unconcerned about—the physical structure or structures that perform the functions that they outline. This is not to say that "appraisal detecting" is not carried out by some physical means. It is just to say that physical properties are not attributed to a psychological capacity that is understood in purely functional terms.

[6]Ultimately, the distinction I am after is between *purely functional*, or nearly purely functional, and *physical*. This distinction is complicated, however, by the assumption that everything is physical. Thus, to start out, I will try to be more precise by using *relational* and *non-relational* (or *intrinsic*). For example, my belief that the sky is gray has only relational properties—so far as we understand the belief anyway. A neuron in my brain has some relational properties, but it also has intrinsic properties, for instance, being a neuron, as well as having length, width, and mass. And these intrinsic properties are also physical properties. An entity that is only known to have relational properties will not have any (known) physical properties.

The problem with Lycan's account, then, is that he claims that psychological capacities and their sub-capacities and sub-sub-capacities occupy levels of nature—levels that are above the levels where neurobiological entities reside. But to reside at a level of nature, an entity or collection of entities must have intrinsic properties that are different than the intrinsic properties of lower level entities. (Or otherwise stated, the psychological capacities, sub-capacities, and so on must be made out of stuff that is different—in some way—than the stuff that lower level entities are made of.) But insofar as these capacities are defined in purely functional terms, they do not have physical (or intrinsic) properties besides those of the relevant neurobiological entities. Hence, they cannot occupy a level of nature.

To recapitulate where we now stand, when Lycan combines homuncular functionalism with his levels of nature, he is making a very strong claim. And one for which there is very little support. A hierarchy created using humuncular functionalism contains entities that are purely functional. The hierarchy of levels of nature, meanwhile, contains entities that have both functional properties and physical (i.e., intrinsic) properties. Since the two hierarchies do not share the same properties, there is no reason to think that they can be combined. And the result is that Lycan's account fails.

It is useful to note that an alternative is to understand psychological capacities in purely functional terms and to treat the functional decomposition as just that—a decomposition that identifies all components by their functions and does not identify, or even make any implicit reference to, physical entities. With such a functional decomposition, the sub-capacities and sub-sub-capacities have an explanatory role, but they are not intended to occupy different levels of nature. This is, essentially, what we find in Cummins' (1975, 1983, 2000) account of functional analysis, although Cummins does add that the functional decomposition should ground out with some specific set of physical entities.

> Ultimately, of course, a complete theory for a capacity must exhibit the details of the target capacity's realization in the system (or system type) that has it. Functional analysis of a capacity must eventually terminate in dispositions whose realizations are explicable via analysis of the target system. Failing this, we have no reason to suppose we have analyzed the capacity as it is realized in that system. (2000, p. 126)

But until the functional decomposition does reach this final stage, an account like Cummins' does not provide any grounds for mapping any of the levels from the functional (i.e., homuncular) decomposition onto the hierarchy of levels of nature.[7]

[7] A different interpretation of Cummins' account of functional analysis is offered by Piccinini and Craver (2011). They suggest that a functional decomposition is not explanatorily useful unless it does, at least implicitly, refer to a physical system.

This leads to a further observation. If a psychological capacity and the functionally defined components that carry out that capacity do not have explicitly physical characterizations, and it seems clear that they do not, then the decomposition of the psychological to the neurobiological is impossible. As Lycan has demonstrated, a functional capacity can be decomposed into sub-capacities via homuncular functionalism. And a physical system can be decomposed into its component parts. But the move from the functional decomposition to the physical decomposition is not part of either one of these decompositions. Rather, the move from the purely functional to the physical is an identification—the identification of the physical parts that carry out a particular function or set of functions.

It might be helpful at this point to look, in a more general way, at the nature of the problem that Lycan's model encounters. Consider two sets of things. Set (1) is a cell, molecules, and atoms. Set (2) is a carrier of hereditary information, a gene, and a segment of deoxyribonucleic acid. In (1)—leaving our reductionist tendencies aside—the three can be understood as different physical things, the former composed of the latter. In (2), the relationship is different. Here, there is no way to understand these three as different physical entities. The only explicitly physical stuff is the segment of DNA; the others, the carrier of hereditary information and the gene, are just different ways of describing this physical material.

The difference between these two sets demonstrates where Lycan goes wrong. The entities in (1) clearly exhibit the hierarchical organization that Lycan relies on, and Lycan's account is capable of describing a relationship like the one between cells and atoms. Those in set (2), meanwhile, cannot be put into the same sort of hierarchy. This is because a *carrier of hereditary information* is a functional description that does not have any intrinsic properties besides those of the DNA. And this is exactly the type of observation that Lycan's account glosses over.

Now consider this set, set (3):

(i) the face recognizer;
(ii) the viewpoint locator, analyzer, librarian, and public relations officer;
(iii) the scanner and projector;
(iv) a collection of molecules somewhere in the brain.

Lycan claims that set (3) has all of the features of a decomposition that occurs across levels of nature (just like set [1]). And while it looks like it does, it does not. The components in (i), (ii), and (iii) have no physical (i.e., intrinsic) properties besides what are found at (iv). That is to say, just as the carrier of hereditary information and the gene are different functional descriptions of the same physical material, so are (i), (ii), and (iii). These three are all just different functional descriptions for the same collection of molecules (or whatever).

Lycan errs by not recognizing—or at least not owning up to—this difference between the entities in set (1) and those in sets (2) and (3). Because there isn't any "stuff" that a psychological capacity is, besides a collection of neurobiological entities and their activities, it is a mistake to think that psychological capacities and neurobiological entities occupy different levels in a hierarchy of levels of nature. As a result, the relationship between the psychological and the neurobiological cannot be as Lycan describes it.

2.5 Levels of mechanisms

Having examined Lycan's account in some detail, another account worth considering is the mechanistic model developed by Craver (2002, 2007). His *hierarchy of mechanistic levels* is a similar attempt to illuminate the relationship between psychological capacities and neurobiological activities. This account, however, encounters very much the same problem as Lycan's, underscoring the difficulty of characterizing this relationship.

Craver begins with mechanisms, which he defines as "collections of entities and activities organized in the production of regular changes from start or set up conditions to finish or termination conditions" (2002, p. 84). Then, as is illustrated in Fig. 3, taking an entity that participates in a mechanism and decomposing it into a separate mechanism produces two mechanisms that occupy different levels in a hierarchy of levels of mechanisms. Further decomposition of the entities in the lower level mechanism creates mores levels.

To illustrate this model, Craver uses a hierarchy consisting of four mechanistic levels that are each, in their own way, in the service of spatial memory. At the top of the hierarchy is the *level of spatial memory*. At this level, this type of memory is described: the contexts in which it is employed, how it stores and retrieves information, its limitations, and so on. Below the level of spatial memory is the *level of spatial map formation*. At this level are the neurons in the hippocampus that are active when rats navigate a maze, and which seem to collectively function as a "spatial map" of the rat's environment.

The formation of these spatial maps is, then, explained by the mechanisms at the *cellular-electrophysiological level*. This is "through LTP [long term

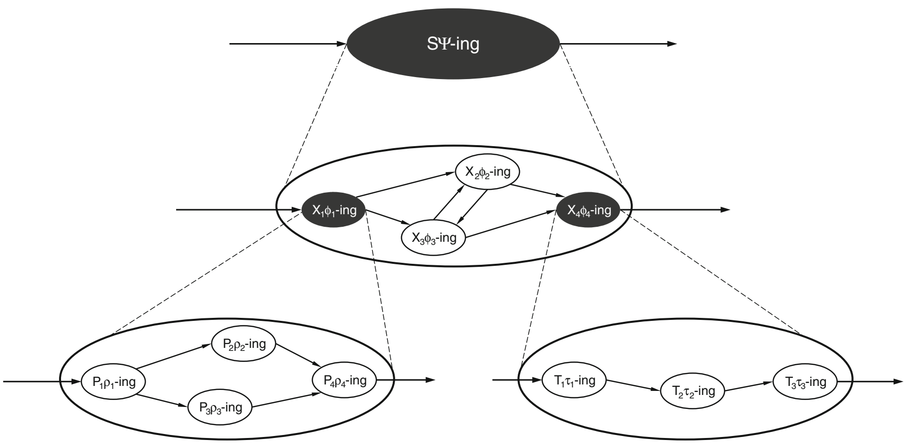

Fig. 3 Multiple levels in a hierarchy of levels of mechanisms. Adapted from Craver (2007, p. 194)

potentiation] in hippocampal synapses," a process by which the effectiveness of synaptic transmission is enhanced (2007, p. 167). And finally, the molecular mechanism that gives rise to long term potentiation is located at the *molecular level*.

According to Craver, these four levels of mechanisms constitute the relationship between this psychological capacity—spatial memory—and the relevant neurobiological activities. He adds,

> My decision to break this explanation into four levels is surely an oversimplification. There might be more levels. One might choose to identify networks of cells in the hippocampus, or cascades of molecules beneath a properly electrophysiological level. The hierarchy could also be expanded upward and downward. Upward, one can consider memory systems in the context of other cognitive and physiological mechanisms (such as emotion and sleep) or in the context of social groups and cultures. Downward, one can consider the protein folding mechanisms that give NMDA receptors their characteristic shapes and activities (2007, pp. 169–170).

Nonetheless, even if there are more levels than the four just described, the example illustrates Craver's account of how psychological capacities are related to neurobiological activities.

The problem with Craver's account is, perhaps, even more visible than it was for Lycan's. A mechanism, by definition, is a set of physical entities and the activities of those entities when all are organized in the appropriate way. As Craver stresses, the functions (his term is *activities*) are one aspect of any mechanism.

> The activities are the various doings in which these entities engage: neurons *fire*, neurotransmitters *bind* to receptors, brain regions *process*, and mice *navigate* mazes. Activities are the things that entities do; they are the productive components of a mechanism, and they constitute the stages of mechanisms. (2002, p. S84)

But just as relevant are the physical objects that that carry out those activities. About this, Craver says,

> Entities are the components or parts in mechanisms. They have properties that allow them to engage in a variety of activities. They typically have locations, sizes, structures, and orientations. They are the kinds of things that have masses, carry charges, and transmit momentum. (2007, pp. 5–6)

The problem here—not so different than it was for Lycan's model—is that psychological capacities (described using the resources of cognitive psychology) are not mechanisms. As has already been discussed at some length, psychological capacities do not have any of the properties of entities just listed by Craver, and from that it follows that the the decomposition of a psychological capacity cannot be a mechanistic decomposition (see Fig. 3). Moreover, although he does have the level of spatial memory in his hierarchy, Craver does not really even try to construct a "psychological" mechanism. About this level, he says the following:

> The topping-off point in this hierarchy is the spatial memory phenomenon. Call this the *level of spatial memory*. ... Spatial memory is tested in radial arm mazes, sunburst mazes, three-table problems, and the Morris water maze. The last of these is a circular pool filled with an opaque liquid covering a hidden platform. ... Researchers monitor the time that it takes the rat to find the platform and, in some cases, the trajectory of the rat through the pool. Experiments of this sort are used to define the phenomenon of spatial memory. (2007, pp. 165–167)

On the one hand, a description of the phenomenon is needed—it would be difficult to make sense of what is going on in the hippocampus without having some understanding of this psychological ability. But on the other hand, it is equally clear that the description of the phenomenon is not itself a mechanism, nor is it an entity that participates in a larger mechanism. Although Craver includes this description among his levels of mechanisms it fails to satisfy his definition of a mechanism.

And so again, we see that psychological capacities cannot be assigned the same type of level—that is a level in the same hierarchy—as neurobiological entities.

3 Levels of organization

A more satisfactory account can be developed if multiple hierarchies are used. One hierarchy is needed to organize the things that are found in nature, but a separate set of levels is needed to track the different descriptions of the things found in nature. To this end, the account developed here uses *levels of organization* to order the things that are found in nature and *levels of explanation* for the descriptions of these things. This section will examine levels of organization and explain how they should be understood. In the next section, the same will be done for levels of explanation.

Although levels of organization are similar to Lycan's levels of nature and Craver's levels of mechanisms, it is worth looking at a couple of the specific features of levels of organization in order to understand the exact role that they have in this new account. The first important feature is the composition relation. The entities at one level are composed of the entities found at lower levels, and so composition orders the levels in the hierarchy. But composition alone cannot be used to establish a hierarchy of levels of organization. If it were, then a new level would be created every time two entities were combined, and this would create far too many levels. The resulting hierarchy would not be helpful for thinking about how nature is organized. And so, in addition to composition, another feature needs to be invoked. The two usual candidates are either *structure* or *interaction*.

When structure is a feature of levels of organization, levels are specified in terms of the significant structures that appear at different scales. This is, for instance, how Churchland and Sejnowski (1992) delineate levels of organization (Fig. 4). As they employ this idea, empirical research determines which structures are the significant ones, and not merely aggregates of lower level components. Hence, on their view, when a scientific consensus determines that a particular natural structure is important, that agreement indicates a level of organization.

The second option, interaction, is a way of characterizing levels of organization that has been developed by William Wimsatt (1976, 2007). When interaction is used as a feature of levels of organization, levels are identified by the regular and predictable interactions that occur among certain entities. A collection of entities interacting with each other in regular and predictable ways—and in many cases depending on these interactions—constitutes a level of organization (1976, pp. 239–42). So, for example, organisms like ourselves have relatively regular and predictable causal interactions with each other and other animals. These interactions indicate that there is a level of organization for organisms. At a smaller scale, the same holds for molecules. They interact

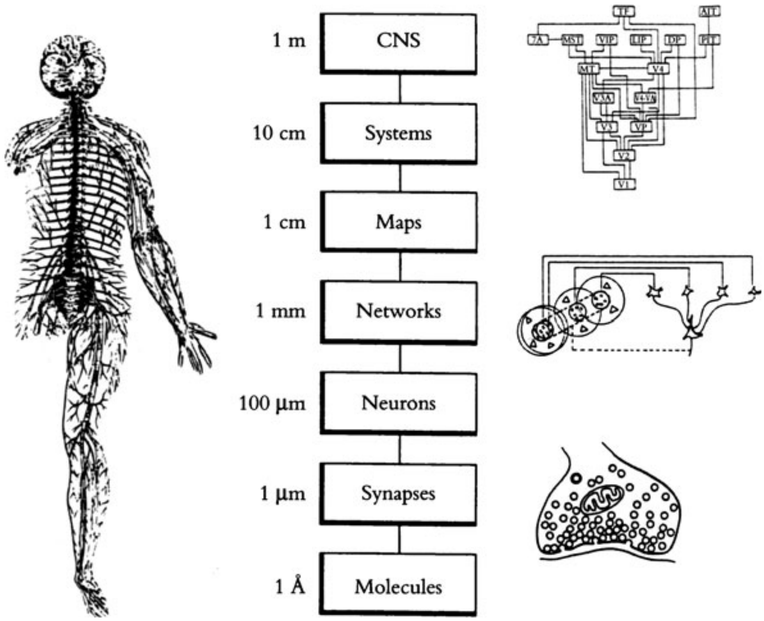

Fig. 4 In the *center* is the hierarchy of levels of organization that Churchland and Sejnowski suggest is found within the scope of the brain. They propose the following levels: systems, topographic maps, local networks, neurons, synapses, and molecules. A map of the visual system (*top*), a network for processing visual information about bars of light (*middle*), and a chemical synapse (*bottom*) are on the *right*. On the *left* is Vesalius's drawing of the human brain, spinal column, and peripheral nerves. From Churchland and Sejnowski (1988, p. 742; see also 1992, pp. 29–48). Reprinted with permission from AAAS

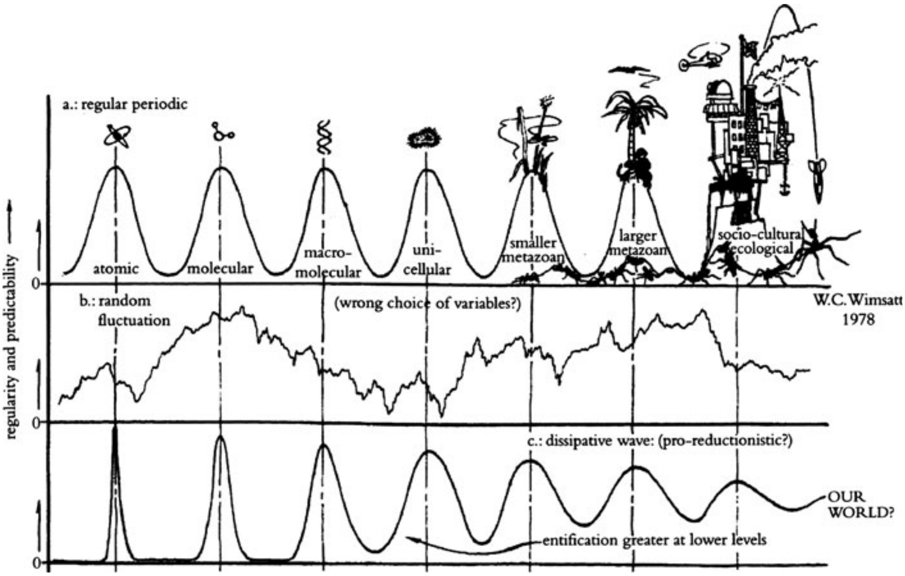

Fig. 5 Wimsatt's diagram of different possible plots for size versus the regularity and predictability of interactions. The *top plot* (**a**) suggests that the regularity and predictability of interactions (on the y-axis) is highest for the sizes of existing entities: atoms, molecules, cells, etc. The slightly less regular plot, **c**, ("Our World?") suggests that the regularity and predictability of interactions are quite high for smaller sizes and becomes progressively less regular as size increases, although they still exhibit an above average degree of regularity for the standard classes of entities. From Wimsatt (1994, p. 230). Reprinted with permission from University of Calgary Press

with other molecules, and these interactions signify a level of organization—and likewise for sub-atomic particles, atoms, and cells, to name a few more. Regular and predictable interactions indicate a sub-atomic level, an atomic level, and a cellular level of organization (Fig. 5).

Structure and interaction are each features that are useful for certain purposes. Here, because the goal is understanding the relationship between psychological capacities and neurobiological *activities*, levels of organization should identify activities, not just structures.[8] Therefore, the levels of organization have to be based on regular and predictable interactions among entities. Using these criteria, we get a level of organization for organisms, a level for cells, one for molecules, and a level for atoms (Fig. 6).

[8]Stepping back for a moment, the reason we are interested in neurobiological activities at all, rather than just neurobiological entities, is because psychological capacities are processes. They are temporally extended, and they typically transform an input into an output. Thus, in order to establish the relationship between psychological capacities and something neurobiological, we have to focus on neurobiological activities, not neurobiological structures.

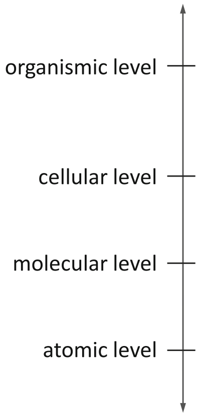

Fig. 6 Four levels of organization, including three that fall within the scope of the brain

There could be other levels of organization besides the cellular, molecular, and atomic that fall within the scope of our neurobiological interests, but a strict application of Wimsatt's criteria appears to suggest that there are only these three. Consider, for instance, the significant brain structures—the brain hemispheres, brain lobes, and functional brain areas. They are not included in this hierarchy because they do not, at least on the face of it, seem to participate in causal interactions. Their parts, neurons, interact with each other, but these aggregates do not themselves interact, and so there is not a level of organization dedicated to any of them (for a more detailed discussion of this issue, see Johnson 2009). But identifying every level of organization that falls within the scope of the brain is not especially important. Regular and predictable interactions among entities establishes several levels of organization, and that is all that is needed here.[9]

Now that we have a clear idea of what levels of organization are and how the hierarchy of levels of organization is constructed, we can see that psychological capacities do not have a place on this hierarchy, at least not qua psychological capacities. In order to integrate psychological capacities and the neurobiological activities that are found on the hierarchy of levels of organization, a second resource is needed, levels of explanation. But before turning to levels of explanation, we can take a moment to look at why psychological capacities, as they are described within cognitive psychology, do not have a place directly on the hierarchy of levels of organization.

Generally speaking, the identification and description of psychological capacities is accomplished by *methodological functionalism*. For cognitive psychology, methodological functionalism amounts to observing the inputs

[9] Craver's levels of mechanisms and Wimsatt's levels of organization have some similarities, most notably, an emphasis on causal interactions. I won't review the differences between the two positions here, but see Craver (2007, chapter 5) for his criticisms of Wimsatt.

that individuals receive and the outputs that they produce and then suggesting internal components that can explain how the inputs are turned into outputs.[10] Our understanding of psychological capacities is—at least until recently, and still for the most part—based entirely on this process. For example, a psychological capacity such as language comprehension is only understood in terms of its function. When internal components are proposed in order to explain language comprehension, no physical material is referred to. The functionally defined components are just elements that help us understand how the capacity is able to do the job that it does.

Psychological capacities are, no doubt, carried out by some physical means. But there is nothing about psychological capacities themselves that makes it appropriate to place them at a level of organization because neither they nor these components refer to (or pick out) any physical material. Thus, in order to understand how psychological capacities are related to the activities that occur in nature, a way of representing the relationship between these functional descriptions and the activities at the levels of organization is needed. For this we turn to levels of explanation.

4 Levels of explanation

In a hierarchy of levels of explanation, each level is a different type of description of the same process. Marr's *Vision* contains the clearest articulation of what levels of explanation are and how they are used (1982). In *Vision* and in the work that preceded it, Marr lays out his account of the visual process using three levels of explanation. Because these three types of descriptions are the standard ways of describing a mental process, Marr's basic account will be adopted here.[11]

Marr's three levels are, from the highest level to the lowest: the level of the computational theory, the level of the representation and algorithm, and the level of hardware implementation. The *level of the computational theory* is the level at which the "information processing tasks" that humans perform are described in a non-causal way; the description at this level is a description of what the task is, not how it is carried out. As Marr says, it is at this level that "the underlying nature of a particular computation is characterized .... One

[10] For more information on methodological functionalism, see Polger (2004, 2009).

[11] Marr does not make a distinction between the terms *explanation* and *description*. I will refer to *descriptions* that occupy *levels of explanation*. Also, although it will not have any bearing on what we do here, it is useful to note that there are other types of descriptions besides the three that Marr uses. (Statistical descriptions are one example.) Consequently, there can be other types of levels of explanation in addition to Marr's. And levels of explanation do not only apply to mental processes, but are used any time a higher level description is provided for a process.

It is also useful to note that while Marr speaks of descriptions of information processing tasks, this type of task is essentially the same as a capacity. A capacity is just the ability to perform a certain task.

can think of this part as an abstract formulation of *what* is being computed and *why*, and I shall refer to it as the 'theory' of a computation" (1977, p. 37).[12]

Consider, for instance, these tasks, which humans are able to perform: adding quantities, understanding a natural language, and holding information temporarily in short-term memory. The descriptions offered at the level of the computational theory specify what the tasks are that humans are able to perform so successfully. Sometimes providing this description is straightforward, as it is for a task such as adding quantities. Other times it takes some work simply to identify the task. That humans have short-term memory had to be discovered by Ebbinghaus in the 1880s (Squire and Kandel 1999). And for tasks such as understanding a natural language, although it is clear that humans are able to perform the task, specifying the exact nature of the task is a difficult problem.

Below the level of the computational theory is the *level of the representation and algorithm*. Here the procedures or operations that carry out—or at least operations that are sufficient to carry out—the task are described. The degree to which these operations describe the actual operations that are performed depends on the access that is available to the system. But, however accurate the proposed operation is, the description offered at this level does not reference the physical system that performs the task. It just describes the operation in terms of (a) an algorithm and (b) representations of the inputs for the algorithm to operate on. It is at the lowest level of explanation, the *level of hardware implementation*, that the physical mechanism that carries out the process is described.

Having reviewed the basics of Marr's account, we can now look at how the descriptions offered by the various mind sciences correspond to these levels of explanation. Descriptions of psychological capacities that are offered within the domain of cognitive psychology are descriptions at Marr's highest level. As was just explained, this level contains the theory of the capacity: a description of what the capacity is. Providing this description can be a substantial undertaking, and there are often competing descriptions offered at this level that correspond to differing ideas about the exact nature of a psychological capacity. Consider two descriptions of the ability to generate an appropriate emotion response. In a brief synopsis of his position, Scherer describes the ability this way:

> the nature of an emotional reaction is based on the individual's subjective appraisal or evaluation of an antecedent situation or event. The evaluation is generally considered to rely on cognitive processing of environmental or proprioceptive stimuli. (1997, p. 114)

[12] It has been pointed out that *computational theory* is a confusing label for this level because it is not where a computational operation is described (Bechtel 1994; Bechtel et al. 1998). As Bechtel et al. say, "[Marr] called his highest level *computational theory* (a label that many have found misleading; it is somewhat akin to Chomsky's notion of competence and might best be called task analysis)" (1998, p. 65).

Robinson, meanwhile, suggests that this same capacity should be understood rather differently.

> Affective appraisals respond automatically to events in the environment (either internal or external) and set off physiological changes that register the event in a bodily way and get the agent ready to respond appropriately. An emotional response is a response set off by a non-cognitive affective appraisal. I speculated that there are probably a limited number of basic emotion systems each identified by a specific non-cognitive appraisal and the particular suite of behaviour it prompts. (2005, p. 89)

Scherer and Robinson agree that generating an emotion response when presented with certain kinds of stimuli is a task that humans perform. They just disagree about the exact nature of this task. Scherer believes that it involves a cognitive evaluation of a stimulus, whereas Robinson is suggesting that it involves a simpler non-cognitive reaction to the stimulus.[13] And while it might be easy to imagine how each of their descriptions could be transformed into a description that belong at Marr's middle level—the level of representation and algorithm—their descriptions are, properly speaking, at the computational level. Scherer and Robinson may disagree about the contents of the middle and lowest levels of explanation as well, but the disagreement illustrated in these two quotations is only a disagreement about what the capacity is, not about how this capacity is carried out.

Describing how the capacity is carried out is the domain of Marr's two lower levels. The description provided at Marr's middle level does this in a more abstract way and must contain a description of the procedure that carries out the capacity, as well as representations of the inputs for the procedure. Typically, the descriptions that are developed within cognitive science belong at this level of the representation and algorithm. These descriptions can be generated with a symbolic framework, a connectionist one, or some other—as long as the description is explicit about the process that transforms representations of the inputs into outputs. The lowest level of explanation provides a description of how the capacity is carried out in a more concrete way: in terms of the relevant biological material. Thus, these are neurobiological descriptions, and they are developed within any of the neuroscientific disciplines.

[13]The descriptions that are in the quotations from Scherer and Robinson are brief; a complete description at this level can be quite substantial if the capacity is characterized in detail. See Scherer (2001) and Robinson (2004, 2005) for complete descriptions of their theories of emotion.

Using Scherer's and Robinson's descriptions is complicated by the fact that, while Scherer is a cognitive psychologist, Robinson is a philosopher. Nonetheless, a little blurring between fields is acceptable. And these two quotations are a good and concise illustration of Marr's highest level.

5 Psychological capacities and neurobiological activities

5.1 The two-dimensional model

Using levels of organization and levels of explanation separates (1) the levels that are used to order those things that are found in nature, and (2) the different descriptions of those things that are found in nature. By keeping (1) and (2) separate, but showing how they are related, we can see how psychological capacities are related to neurobiological activities.

To create a single framework with these two types of levels, we need only notice that the lowest level of explanation and some level of organization identify the same class of things: the interactions of entities found in nature. Regular and predictable interactions are, of course, one of the central features of levels of organization. And the description supplied at the lowest level of explanation is a description of a biological mechanism, which is a specific series of interactions. Thus, while the description at the lowest level of explanation focuses on the particular series of activities that carry out a capacity, this circumscribed series of activities is set within all of the activity at a level of organization, which includes more interactions than just those that are concerned with a single psychological capacity (Fig. 7).

The next step is determining where the hierarchy of levels of organization and this hierarchy of levels of explanation intersect. This requires finding the

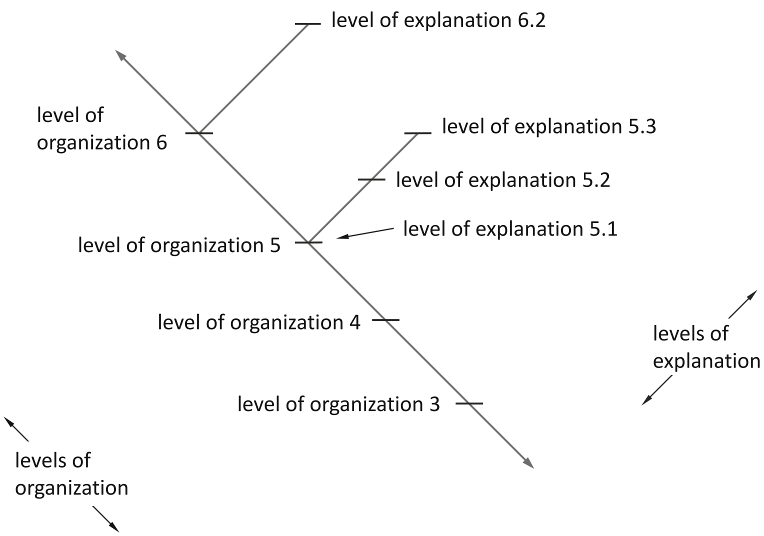

Fig. 7 The two-dimensional model created by the intersection of levels of organization and levels of explanation. Levels of organization are on the leftmost axis. Levels of explanation are the hierarchies on the right. Levels 5.3, 5.2, and 5.1 can be thought of as Marr's three levels. Other levels of explanation are also possible, however. For example, level 6.2 is a description of some of the activities occurring at level of organization 6

correct level of organization for the lowest level of explanation. The levels of organization that fall within the scope of the brain have already been identified: the cellular level, the molecular level, and the atomic level. Of these levels, the cellular level of organization seems to be an appropriate intersection point. The entities at this level, neurons, interact over the appropriate spatial and temporal scales to carry out psychological capacities, and much of the work within neuroscience confirms the claim that the activities at this level carry out psychological processes.[14] If this is the appropriate level of organization, then some of the activities that occur at the cellular level of organization—the activities and interactions of neurons—figure in the descriptions that occur at the lowest level of explanation. This intersection of levels of organization and levels of explanation creates a two dimensional model, a summary of which is shown in Fig. 8.

This model illustrates how psychological capacities are related to neurobiological activities. Psychological capacities are just a particular type of description of the activities that occur at the cellular level of organization. Psychological descriptions draw on a particular set of resources and, hence generate a particular kind of description. But they are, nonetheless, just descriptions of those activities that occur at the cellular level of organization. The psychological capacities that are described in cognitive psychology may seem distinct from the activities found at the cellular level of organization, but they are not, and moreover, they cannot be. For it to be otherwise, psychological capacities would either have to occupy their own level of organization or they would have to be unconnected from the activities that occur in nature.

A caveat is now in order. While Fig. 8 illustrates the significant features of this account, a more complete model should indicate that some aspects of psychological capacities are carried out by activities at the molecular level of organization. Molecular activity gives rise to the plasticity that occurs at the cellular level, and this does bear on the execution of psychological capacities. In the more detailed model, the intersection of levels of explanation and levels of organization is split between the cellular and the molecular levels of organization (Fig. 9).

[14]This is a substantial topic. Nevertheless, there are plenty of examples of neuroscientific investigations into how psychological processes are carried out at the cellular level. See, for example, Graziano's work on microstimulation of the neurons in the motor cortex (Graziano et al. 2002a, b); Rolls' investigation of the responses of neurons in the orbitofrontal cortex to certain kinds of visual information (Rolls et al. 2005, 2006); or the progress that is being made on the control of prosthetic devices by cellular activity (Carmena et al. 2003; Tillery and Taylor 2004; Velliste et al. 2008).

It is also interesting to note that the descriptions that Marr offered for the lowest level of explanation all involved activities at the cellular level of organization. The visual processes that he investigated are carried out, he suggested, by the activities of the neurons in the retina, lateral geniculate nucleus, and primary visual cortex (Marr 1982; Marr and Ullman 1981).

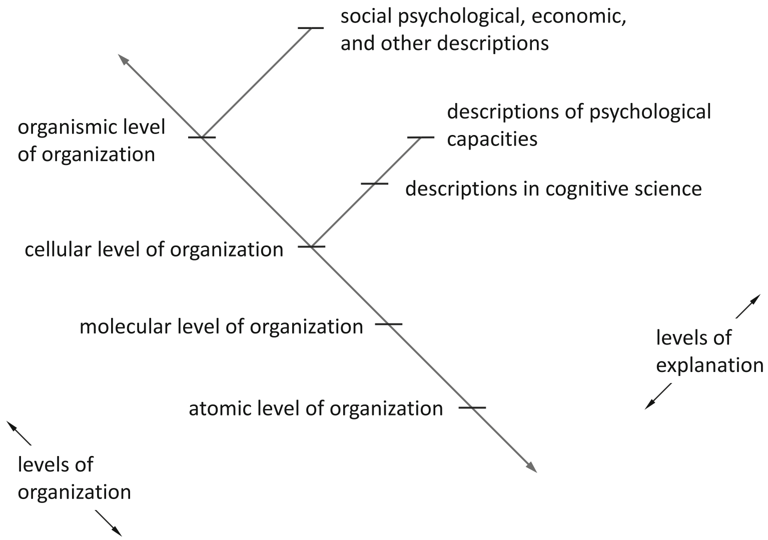

Fig. 8 A two-dimensional model of levels. Levels of organization are on the leftmost axis. Levels of explanation are the hierarchies on the right. The lower of the two hierarchies of levels of explanation is the one adopted from Marr. Descriptions of psychological capacities are at the *top*, descriptions offered in cognitive science are located at the *middle level*, and the *lowest level* of explanation coincides with the cellular level of organization. The other hierarchy of levels of explanation, which has social psychological and economic descriptions at the highest level, is included to demonstrate that levels of explanation are utilized any time descriptions are offered of the activities found at a specific level of organization

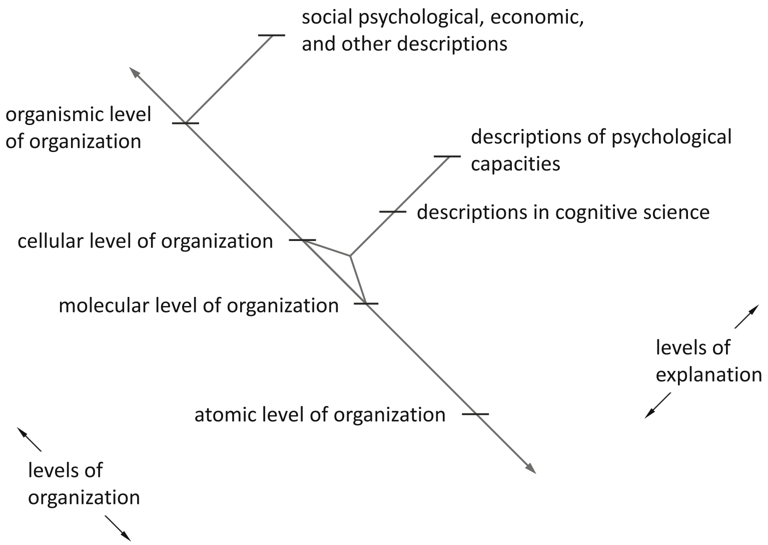

Fig. 9 In this diagram, the lowest level of explanation for psychological capacities intersects with both the cellular and the molecular levels of organization. See text for further explanation

5.2 Spatial memory again

Recall from section two that Craver uses spatial memory to illustrate his levels of mechanisms. Now this psychological ability can be used to illustrate the two-dimensional model. This will be brief, however. It will only include the highest and lowest levels of explanation, and it will only sketch out those descriptions. The middle level of explanation is not really needed in order to illustrate the relationship between the description of a psychological capacity and the relevant neurobiological activities. But keeping with the full model shown in Fig. 9, the lowest level of explanation in this example will focus first on the activities of the neurons in the hippocampus, and then, because spatial memory is especially dependent on neural plasticity, on a description of the activity at the molecular level of organization.

We begin with the lowest level of explanation; the one that occurs at the cellular level of organization. Spatial memory is—at least in part—carried out by neurons in the hippocampus. Three groups of excitatory neurons are arranged in a distinctive C-shaped pattern: granule cells in the dentate gyrus, and pyramidal cells in the CA3 and CA1 regions. As is shown in Fig. 10, these cells primarily make contact with each other and with neurons in the entorhinal cortex. But while Fig. 10 outlines the general ways in which these neurons are connected, there are many other interactions occurring here that belong

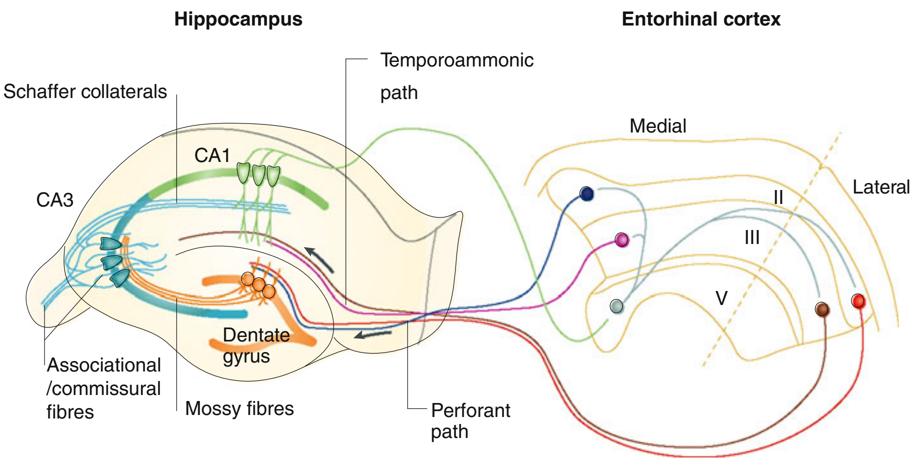

Fig. 10 The majority of the inputs to neurons in the hippocampus are from pyramidal cells in the entorhinal cortex (whose axons make up the perforant path). These axons terminate on granule cells in the dentate gyrus. From there, the basic trajectory is as follows. The axons of the granule cells (the mossy fibers) innervate CA3 pyramidal cells. The CA3 pyramidal cells contact pyramidal cells in the CA1 region ipsilaterally (by the Schaffer collaterals) and pyramidal cells in the contralateral CA3 and CA1 regions (via commissural fibers). The CA1 pyramidal cells then project back to neurons in the entorhinal cortex as well as to neurons in other cortical areas. From Neves et al. (2008, p. 66). Reprinted with permission from Macmillan Publishers Ltd: *Nature Reviews Neuroscience*

in the description of the activities that carry out spatial memory. For instance, CA3 pyramidal cells not only project to other areas, but also have many local connections with other pyramidal cells in the CA3 region. And although neurons in the entorhinal cortex mainly provide inputs to the granule cells in the dentate gyrus, they also innervate pyramidal cells in both the CA3 and CA1 regions. Plus, besides these excitatory granule and pyramidal cells, there are also many inhibitory neurons in the hippocampus that modulate the activities of the granule and pyramidal cells (Neves et al. 2008).

But besides outlining the diverse connections among these neurons, even a brief description of the activity at this level must point out that at least some of the pyramidal and granule cells in the hippocampus are *place cells*. These are cells that—based on what has been found in mice and rats—respond (i.e., generate action potentials) only when the animal is in a specific location. As Neves et al. report,

> Single-unit recordings from neurons in the hippocampus of freely moving rodents reveal that pyramidal and granule cells show a preference for firing in a particular location of an explored environment, regardless of the direction from which the animal enters the location. Hundreds of such 'place cells' fire in concert as a rat reaches a particular location, and place cells fire in sequence as the animal moves through a series of locations in a given environment. (2008, pp. 66–67)

In addition to the relevant activities at the cellular level of organization, the lowest level of explanation must also focuses on some of the activities at the molecular level of organization. The molecular activities that are relevant to a description of spatial memory are, in particular, those that give rise to long-term potentiation (LTP), a process that "strengthens" the connections between neurons.[15] LTP occurs when multiple pre-synaptic neurons simultaneously produce high frequency stimulation at different locations on a single post-synaptic neuron, causing the post-synaptic neuron to generate an action potential. In a relatively short amount of time, changes to the post-synaptic neuron—some temporary and some longer lasting—allow it to become active when receiving lower levels of stimulation from the pre-synaptic neurons.

[15]There are other processes that produce plasticity and probably occur in the neurons in the hippocampus. Long-term depression, spike-timing-dependent plasticity, and depotentiation are a few (Neves et al. 2008, p. 66). These other means of plasticity won't be discussed here, but they could be included in the descriptions that reside at the molecular level of organization (i.e., the lowest level of explanation that occurs at the molecular level of organization). But notice that the descriptions of molecular activities need only be those that are relevant to explaining spatial memory. Of course, "explanatorily relevant" can be widely interpreted, but, on the face of it, there does not seem to be a need to provide a molecular level explanation of every aspect of the activity that occurs at the cellular level in order accurately describe spatial memory.

The following is a very brief description of the activity at the molecular level during the later, longer lasting phase of LTP (see also Fig. 11). When neurotransmitter is released by a pre-synaptic neuron, AMPA receptors on the post-synaptic neuron open and positively charged sodium ions are drawn into the cell. This increase in intracellular Na^+ activates nearby NMDA receptors, and through these receptors positively charged calcium ions enter. The influx of Ca^{2+} sets in motion a cascade of molecular activity that results in gene transcription and the production of new proteins. These proteins are transported back to the synapse where they produce structural changes—probably the production of new AMPA receptors and the enlargement of the dendritic spine—which will allow a weaker pre-synaptic stimulation to produce a response in the post-synaptic cell (Kandel 2001; Bourne and Harris 2011. See Bickle 2003 for a discussion set within a philosophical context).

Meanwhile, at the highest level of explanation, the description focuses on what this spatial memory ability is that humans and other animals have. One suggestion, developed by the psychologists Amy Shelton and Timothy McNamara, is summarized this way:

> Shelton and McNamara (2001) proposed that learning and remembering the spatial structure of the surrounding environment involves interpret-

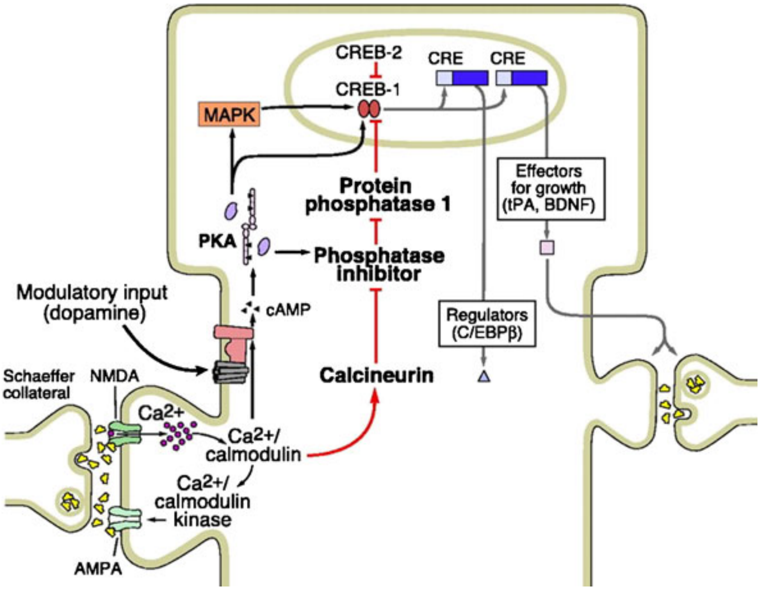

Fig. 11 Molecular activity in the post-synaptic neuron during late phase LTP. This activity begins at the synapses, proceeds to the nucleus in the cell body where transcription and the synthesis of new proteins occurs, and then returns to the synapses. The structural changes that occur are probably the placement of new AMPA receptors and the enlargement of dendritic spines (the synaptic portion of the post-synaptic neuron). From Kandel (2001, p. 1036). Reprinted with permission from AAAS

> ing the layout in terms of a spatial reference system. They suggested that this process is analogous to determining the "top" of a figure (e.g., Rock 1973); in effect, conceptual "north" is assigned to the layout, creating privileged directions in the environment. The frame of reference for this interpretation is selected using cues. The dominant cue, according to Shelton and McNamara (2001), is egocentric experience, but other cues can be used as well, including the structure of the environment itself. They also proposed that egocentric perspectives that are aligned with salient directions, axes, or planes in the environment are preferred to those that are not. (Mou and McNamara 2002, p. 162)

Notice that this psychological description does not describe how this ability is carried out, only what it is. How it is carried out is the province of the lower levels of explanation.

Furthermore, contra Lycan, there is no need to decompose this description so that it can smoothly align with the neurobiological description. Potentially, that could be accomplished, but to no obvious purpose. If this description accurately characterizes the psychological capacity in question, nothing more needs to be done besides recognizing that it is a higher level description of a certain set of neurobiological activities.[16]

6 Conclusion

The goal in this paper has been to outline an alternative model for understanding the relationship between psychological capacities and neurobiological activities. Several aspects of this model have been explained rather quickly. A more complete account should include a fuller discussion of the nature of psychological capacities, as well as a further defense of the claim that it is the activities at the cellular level of organization that primarily carry out psychological capacities. Nevertheless, what has been put forward here is a useful alternative to an account—like Lycan's—that locates psychological capacities and neurobiological activities on a single hierarchy of levels.

Based on this new model, it is clear that the psychological is not a part of nature in the way that Lycan suggests: the psychological is not one of the many types of things that occupy the various levels of organization. This model also indicates that the relationship between the psychological and the neurobiological is not analogous to the relationship between chemical processes and those of basic physics. Instead of having a composition relationship, psychological capacities are only a certain type of description of neurobiological processes.

[16]Some changes have been made to Craver's highest level (his "level of spatial memory"), and that description has become the highest level of explanation. Also, one of Craver's mechanistic levels has been eliminated. Besides a level for the cellular activities, he includes a level for LTP and then a separate level for the molecular process that carries out LTP. Dropping one of the latter two descriptions seems warranted.

This is a type of description that is useful. But because it is only a certain type of description, the psychological is not separable from the neurobiological in the same way as chemical processes are separable from those of basic physics.

Acknowledgements I am grateful for the help that I received from Thomas Polger and John Bickle on the many drafts of this paper. I am also indebted to two anonymous reviewers for this journal who provided very helpful comments on the paper's final versions.

References

Bechtel, W. (1994). Levels of description and explanation in cognitive science. *Minds and Machines*, *4*, 1–25.

Bechtel, W., Abrahamsen, A., & Graham, G. (1998). The life of cognitive science. In W. Bechtel, & G. Graham (Eds.), *A companion to cognitive science* (pp. 2–104). Malden, MA: Blackwell.

Bickle, J. (2003). *Philosophy and neuroscience: A ruthlessly reductive account*. Boston: Kluwer.

Bourne, J. N., & Harris, K. M. (2011). Coordination of size and number of excitatory and inhibitory synapses results in a balanced structural plasticity along mature hippocampal CA1 dendrites during LTP. *Hippocampus*, *21*, 354–373.

Carmena, J. M., Lebedev, M. A., Crist, R. E., O'Doherty, J. E., Santucci, D. M., Dimitrov, D. F., et al. (2003). Learning to control a brain–machine interface for reaching and grasping by primates. *PLoS Biology*, *1*, 193–208.

Churchland, P. S., & Sejnowski, T. J. (1988). Perspectives on cognitive neuroscience. *Science*, *242*, 741–745.

Churchland, P. S., & Sejnowski, T. J. (1992). *The computational brain*. Cambridge, MA: MIT Press.

Craver, C. F. (2002). Interlevel experiments and multilevel mechanisms in the neuroscience of memory. *Philosophy of Science*, *69*(Proceedings), S83–S97.

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*, 741–765.

Cummins, R. (1983). *The nature of psychological explanation*. Cambridge, MA: MIT Press.

Cummins, R. (2000). "How does it work?" versus "what are the laws?": Two conceptions of psychological explanation. In F. C. Keil, & R. A. Wilson (Eds.), *Explanation and cognition* (pp. 117–144). Cambridge, MA: MIT Press.

Dennett, D. C. (1978). *Brainstorms: Philosophical essays on mind and psychology*. Montgomery, VT: Bradford.

Graziano, M. S. A., Taylor, C. S. R., & Moore, T. (2002a). Complex movements evoked by microstimulation of precentral cortex. *Neuron, 34*, 841–851.

Graziano, M. S. A., Taylor, C. S. R., Moore, T., & Cooke, D. F. (2002b). The cortical control of movement revisited. *Neuron, 36*, 349–362.

Johnson, G. (2009). Mechanisms and functional brain areas. *Minds and Machines, 19*, 255–271.

Kandel, E. R. (2001). The molecular biology of memory storage: A dialogue between genes and synapses. *Science, 294*, 1030–1038.

Lycan, W. G. (1981). Form, function, and feel. *The Journal of Philosophy, 78*, 24–50.

Lycan, W. G. (1987). *Consciousness*. Cambridge, MA: MIT Press.

Lycan, W. G. (1991). Homuncular functionalism meets PDP. In W. Ramsey, S. P. Stich, & D. E. Rumelhart (Eds.), *Philosophy and connectionist theory* (pp. 259–286). Hillsdale, NJ: Lawrence Erlbaum.

Marr, D. (1977). Artificial intelligence–A personal view. *Artificial Intelligence, 9*, 37–48.

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco: W.H. Freeman.

Marr, D., & Ullman, S. (1981). Directional selectivity and its use in early visual processing. *Proceedings of the Royal Society of London B, 211*, 151–180.

Mou, W., & McNamara, T. P. (2002). Intrinsic frames of reference in spatial memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28*, 162–170.

Neves, G., Cooke, S. F., & Bliss, T. V. P. (2008). Synaptic plasticity, memory and the hippocampus: A neural network approach to causality. *Nature Reviews Neuroscience, 9*, 65–75.

Piccinini, G., & Craver, C. (2011). Integrating psychology and neuroscience: Functional analyses as mechanism sketches. *Synthese, 183*, 283–311.

Polger, T. W. (2004). *Natural minds*. Cambridge, MA: MIT Press.

Polger, T. W. (2009). Computational functionalism. In J. Symons, & P. Calvo (Eds.), *Routledge companion to the philosophy of psychology* (pp. 148–163). New York: Routledge.

Robinson, J. (2004). Emotion: Biological fact or social construction? In Robert C. Solomon (Ed.), *Thinking about feeling: Contemporary philosophers on emotion* (pp. 28–43). New York: Oxford University Press.

Robinson, J. (2005). *Deeper than reason: Emotion and its role in literature, music, and art*. Oxford: Oxford University Press.

Rock, I. (1973). *Orientation and form*. New York: Academic Press.

Rolls, E. T., Browning, A. S., Inoue, K., & Hernadi, I. (2005). Novel visual stimuli activate a population of neurons in the primate orbitofrontal cortex. *Neurobiology of Learning and Memory, 84*, 111–123.

Rolls, E. T., Critchley, H. D., Browning, A. S., & Inoue, K. (2006). Face-selective and auditory neurons in the primate orbitofrontal cortex. *Experimental Brain Research, 170*, 74–87.

Scherer, K. R. (1997). Profiles of emotion-antecedent appraisal: Testing theoretical predictions across cultures. *Cognition & Emotion, 11*, 113–150.

Scherer, K. R. (2001). Appraisal considered as a process of multilevel sequential checking. In K. R. Scherer, A. Schorr, & T. Johnstone (Eds.), *Appraisal processes in emotion: Theory, methods, research* (pp. 92–120). New York: Oxford University Press.

Shelton, A. L., & McNamara, T. P. (2001). Systems of spatial reference in human memory. *Cognitive Psychology, 43*, 274–310.

Smith, C. A., & Kirby, L. D. (2000). Consequences require antecedents: Toward a process model of emotion elicitation. In J. P. Forgas (Ed.), *Feeling and thinking: The role of affect in social cogition* (pp. 83–106). Cambridge, U.K.: Cambridge University Press.

Squire, L. R., & Kandel, E. R. (1999). *Memory: From mind to molecules*. New York: Scientific American Library.

Tillery, S. I. H., & Taylor, D. M. (2004). Signal acquisition and analysis for cortical control of neuroprosthetics. *Current Opinion in Neurobiology, 14*, 758–762.

Velliste, M., Perel, S., Spalding, M. C., Whitford, A. S., & Schwartz, A. B. (2008). Cortical control of a prosthetic arm for self-feeding. *Nature, 453*, 1098–1101.

Wimsatt, W. C. (1976). Reductionism, levels of organization, and the mind-body problem. In G. G. Globus, M. Grover, & I. Savodnik (Eds.), *Consciousness and the brain: A scientific and philosophical inquiry* (pp. 205–267). New York: Plenum Press.

Wimsatt, W. C. (1994). The ontology of complex systems: Levels of organization, perspectives, and causal thickets. In M. Matthen & R. Ware (Eds.), *Biology and society: Reflections on methodology* (Supplementary vol. 20 of the Canadian Journal of Philosophy, pp. 207–274). Calgary: University of Calgary Press.

Wimsatt, W. C. (2007). *Re-engineering philosophy for limited beings: Piecewise approximations to reality*. Cambridge, MA: Harvard University Press.