

Interlevel Experiments and Multilevel Mechanisms in the Neuroscience of Memory

Carl F. Craver^{†‡}

Washington University

The dominant neuroscientific theory of spatial memory is, like many theories in neuroscience, a multilevel description of a mechanism. The theory links the activities of molecules, cells, brain regions, and whole organisms into an integrated sketch of an explanation for the ability of organisms to navigate novel environments. Here I develop a taxonomy of interlevel experimental strategies for integrating the levels in such multilevel mechanisms. These experimental strategies include activation strategies, interference strategies, and additive strategies. These strategies are mutually reinforcing, providing a kind of interlevel and intratheoretic robustness that has not previously been recognized.

1. Introduction. Many theories in contemporary neuroscience are multilevel descriptions of mechanisms. One aim of experimentation is to integrate the different levels in such theories. In this paper I analyze the concepts of “mechanism” and “level” and deploy this analysis to describe three mutually reinforcing kinds of interlevel experiments used to integrate the levels in multilevel theories. My discussion is constructed by reference to recent attempts to integrate Long-Term Potentiation (LTP), a form of synaptic plasticity, into a multilevel mechanism of memory. The image of neuroscientific theory construction developed through this case diverges from traditional reductionistic perspectives (as advocated in, e.g., Schaff-

[†]Send requests for reprints to Carl F. Craver, Department of Philosophy, Busch Hall, One Brookings Drive, Washington University, St. Louis, St. Louis, MO 63130; e-mail: carl@twinearth.wustl.edu.

[‡]Special thanks to William Bechtel and Lindley Darden for comments and suggestions. Figures, which also appeared in Craver and Darden (2001), are reprinted with permission from the University of Pittsburgh Press.

Philosophy of Science, 69 (September 2002) pp. S83–S97. 0031-8248/2002/69supp-0008\$10.00
Copyright 2002 by the Philosophy of Science Association. All rights reserved.

ner 1993 and more recently by Bickle 1998), and does so more fully than (although somewhat consistently with) recent multilevel perspectives on neuroscientific practice (e.g., Bechtel and Richardson 1993; Schouten and Looren de Jong 1999). What is novel in this paper is the use of a nonformal account of theory structure to ground a taxonomy of specifically interlevel experiments and the use of this experimental taxonomy to explore the process of integrating levels in a multilevel neural mechanism.

2. Mechanisms and Their Organization. Mechanisms, as they are understood in contemporary neuroscience, are collections of entities and activities organized in the production of regular changes from start or setup conditions to finish or termination conditions (Machamer, Darden, and Craver 2000). The entities in neuroscience include things like neurons, neurotransmitters, brain regions, and mice. 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. When neuroscientists speak generally about activities, they use a variety of terms; activities are often called “processes,” “functions,” and “interactions.” When they speak specifically about activities, they use verbs and verb forms; they speak of attracting and repelling, phosphorylating and hydrolyzing, binding and breaking, and firing and releasing.

The entities and activities composing mechanisms are *organized*; they are organized such that they *do* something, *carry out* some task or process, *exercise* some faculty, *perform* some function or *produce* some end product. I will refer to this activity or behavior of the mechanism as a whole as the role to be explained by the description of the mechanism. The role is the activity at the top of Figure 1. Below it are the entities and activities composing the mechanism for that role.

The entities and activities composing mechanisms have a spatial and temporal organization that is crucial to their productivity. (By “crucial” I mean necessary in the circumstances (cf. Nagel 1977); this might usefully be fleshed out with the help of Mackie’s (1974) notion of an INUS condition (for **insufficient nonredundant** part of an **unnecessary but sufficient** condition)). Spatially, the entities composing the mechanism must be appropriately located, connected, structured and oriented with respect to one another if the mechanism is to work. The activities composing mechanisms also have crucial temporal orders, rates, and durations. Uncovering these temporal and spatial aspects of a mechanism’s organization is a major step in the construction of neuroscientific theories and so is a major focus of neuroscientific practice (see Craver and Darden, forthcoming).

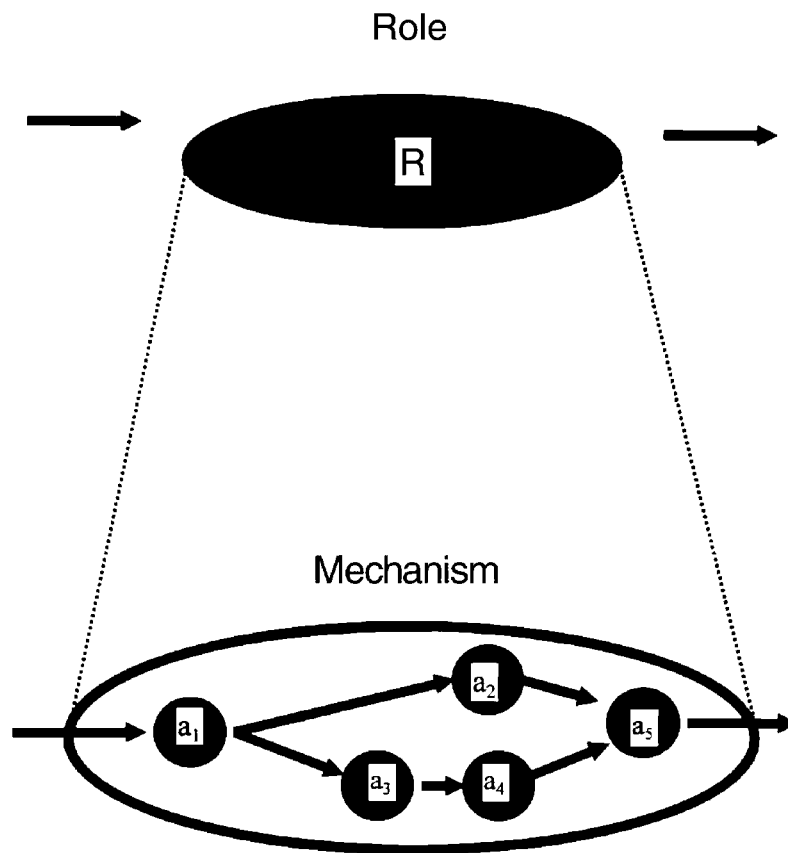


Figure 1

3. Example: The Mechanism of Long Term Potentiation. LTP is a means of strengthening synapses in the central nervous system. Many think that LTP is a crucial activity in the mechanisms of memory. LTP is a form of neural plasticity reminiscent of a memory mechanism proposed by D. O. Hebb in his 1949 *Organization of Behavior*. Hebb's idea was that memories might be formed by strengthening synapses when both the presynaptic and the postsynaptic neurons are simultaneously active. This hypothesis has had considerable staying power in contemporary neuroscience; in fact, it recently contributed to Eric Kandel's share of the Nobel Prize for medicine.

LTP is typically studied in the hippocampus, a brain structure long known to exhibit this form of synaptic modification. The hippocampus is another of the crucial entities in the mechanism of memory. Surgical re-

removal of the hippocampus produces profound memory deficits in human and nonhuman animals. A cross section of the hippocampus with some of its major anatomical regions and synaptic connections is shown in Figure 2. LTP can be induced at each of the three major excitatory synapses in this diagram.

There is no consensus about the mechanisms that produce LTP. One researcher has complained that the LTC (Long Term Controversy) over LTP is threatening to become an LTTP (a "Long Term Tar Pit") for neurobiologists (Malinow 1998, 1226). Nonetheless, an example of one plausible, if incomplete, sketch of the mechanism for LTP nicely illustrates several aspects of the mechanisms described in neuroscientific theories.

The hippocampal synapses that exhibit LTP use the neurotransmitter glutamate. Glutamate is released from the presynaptic cell with each action potential, and binds to receptors on the postsynaptic cell. LTP can be thought of as an increase in the effect of a single presynaptic action potential on the postsynaptic electrical response. This increase in the strength of the synapse could be due, for example, to the release of more glutamate from the presynaptic cell, or to the changing receptive properties of the postsynaptic cell, or perhaps to both.

One type of postsynaptic glutamate receptor in the hippocampus is called the NMDA receptor (for N-Methyl D-Aspartate, a chemical agonist that has a high affinity for this receptor). When glutamate binds to NMDA receptors on the postsynaptic cell, the NMDA receptors change their shape, exposing a pore in the cell membrane. If the postsynaptic cell is inactive, the channel remains blocked by large Mg^{2+} ions. But if the

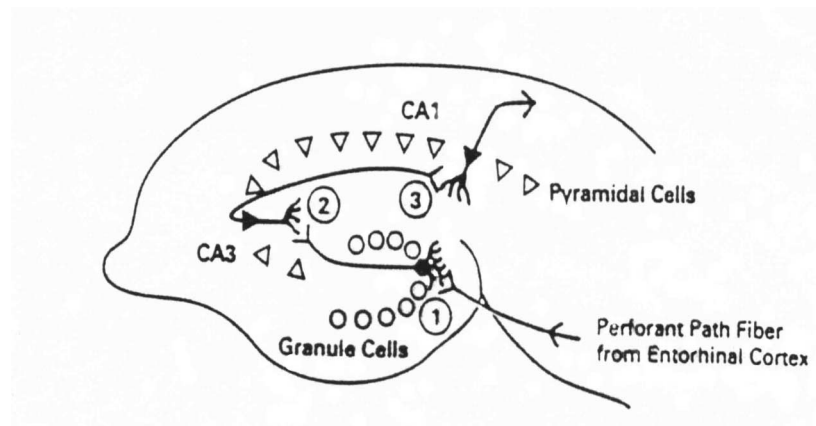


Figure 2

postsynaptic cell is depolarized, these Mg^{2+} ions float out of the channel, allowing Ca^{2+} to diffuse into the cell. The rising intracellular Ca^{2+} concentrations set in motion a long biochemical cascade terminating in the question marks of Figure 3.

The remaining details of this mechanism are more speculative, but three things are thought to happen. In the short term, it is thought that this cascade leads to an increase in the number or sensitivity of so-called AMPA receptors (perhaps by phosphorylation). These changes account for the rapid induction of LTP. In the long term, the cascade leads to the production of proteins in the postsynaptic cell body. These proteins are thought to alter the structure of the dendritic spines at that synapse (see, e.g., Engert and Bonhoeffer 1999; Maletic-Savetic, Malinow, and Svoboda 1999). Some suspect that there is also a presynaptic component of the LTP mechanism whereby, for example, the presynaptic cell releases more glutamate.

The entities in this mechanism are glutamate molecules, NMDA receptors, Ca^{2+} ions and the like. The activities include binding, diffusing, phosphorylating, and changing conformation. The working of the mechanism depends crucially upon its organization. It depends upon the order of the activities and on their relative rates and durations. It also depends crucially upon the structures, shapes, sizes, orientations, and locations of the com-




Figure 3

ponent entities. LTP is a representative example of mechanisms in contemporary neuroscience. But no analysis of mechanisms in neuroscience is complete without an analysis of their characteristic *multilevel* organization.

4. Three Kinds of Levels. Talk of “levels” is nearly ubiquitous in neuroscience and its philosophy. Both philosophers of mind and science are increasingly recognizing the difficulties and ambiguities attending such talk (see e.g., Heil 1999 and Kim 1993). But it is possible to take some initial steps toward greater clarity by disambiguating three different kinds of “levels”: levels of mere aggregates, functional levels, and mechanistic levels. Each of these kinds is individuated by a different asymmetrical decomposition relation. The multilevel theories of contemporary neuroscience exhibit multiple mechanistic levels, but the contrast with the other kinds is revealing.

Begin with **levels of mere aggregates** and the corresponding notion of an aggregative decomposition. An aggregative decomposition involves dividing some chunk of matter—some entity (it is always an entity that is aggregatively decomposed)—into smaller chunks of matter. The ball of wax and the hippocampus can each be sliced, diced, cubed, or spiral cut into parts; and these smaller parts could then, at least in principle, be put back together to fill the same volume of space occupied before the decomposition (see Haugeland 1998, chaps. 1 and 9). The intended sense of aggregativity is that developed by Wimsatt (1986): The properties of wholes are simple sums of the properties of parts (e.g., volume and mass); the wholes are stable under disaggregation and reaggregation of parts; and the parts do not significantly interact with one another. Talk of aggregate levels highlights relations of spatial inclusion and size among entities to the total neglect of activities and their organization.

Where aggregate levels are relationships among entities, **functional levels** are relations among abstract roles. Functional decomposition of one level into another involves taking a task, a routine, or a faculty and breaking it into sub-tasks, sub-routines, or sub-faculties. Functional decompositions are often treated by neuroscientists as if they were, at best, necessary oversimplifications in the generation of testable sketches or, at worst, pie in the sky speculations that are replaced or obviated as the details of a mechanism become available. This is because decomposition by functional role alone does not adequately embody those roles in the entities and activities that the ontic store of contemporary neuroscience has to offer. For the neuroscientist, purely functional decompositions are disembodied “how-possibly” descriptions of a mechanism; they are sometimes denigrated as “boxology.”

What neuroscientists are after is neither an aggregative decomposition

nor a purely functional decomposition, but rather a mechanistic decomposition into **mechanistic levels**—a decomposition into entities and activities organized in the performance of a higher level role. The activities and properties of the entities in the lower level mechanism may themselves be subject to mechanistic decomposition. In such cases, each mechanistic decomposition adds another level to what may become a multilevel mechanism. It is typically possible to distinguish levels by the different entities and activities that populate them and, as we will see, by the different techniques that are used to investigate those entities and activities. But how many levels there are and what kinds of entities are found at each level are empirical questions to be answered within a given research program.

Consider a sketch of the mechanisms of spatial memory. This sketch has roughly four distinct mechanistic levels, although we should expect the number of levels and the descriptions at different levels to change over time. At the top is a *behavioral-organismic* level, having to do with, for example, the various types of learning and memory, the conditions under which different memories may be stored or retrieved, and the conditions under which storage or retrieval are likely to improve or fail. Techniques for investigating phenomena at the behavioral-organismic level typically involve behavioral tasks, such as navigation, recognition of objects, and tests of avoidance, aversion, and preference.

Beneath this behavioral-organismic level is a *computational-hippocampal* level, having to do roughly with the role of the hippocampus in the mechanisms of memory, its cytological, anatomical, and structural features, its pathology, its connectivity with other brain regions, and the computational or processing stages it is thought to perform. Techniques for investigating phenomena at this level include ablation, pathological anatomy, multicellular recording, EEG, PET and MRI, as well as various computational approaches. Claims that the cells of the hippocampus (or some part of the hippocampus) may function as a “spatial map” (O’Keefe and Dostrovsky 1971), as an organ of “declarative memory” (Zola-Morgan and Squire 1993), as a relater of “an item and its context” (Schachter and Wagner 1999), or as “self localization and route replay” (Redish and Touretzky 1998) are hypotheses about this level.

The contribution of the hippocampus to the phenomena of memory is thought to involve LTP and various synaptic components. This *electrical-synaptic* level includes such entities as neurons, synapses, and dendritic spines and such activities as vesicular release and the generation and propagation of action potentials. Phenomena at this level are typically investigated with pharmacological and electrophysiological techniques.

Bottoming out this hierarchy are entities and activities at a *molecular-kinetic level*. At this level entities like the NMDA and AMPA receptors, glutamate, Ca^{2+} ions, and Mg^{2+} ions engage in activities like attracting




Figure 4

and repelling, binding and breaking, phosphorylating and hydrolyzing. These components are investigated with a host of biochemical, and increasingly, molecular biological techniques.

To summarize, the mechanism sketch for memory is multilevel; its current description includes mice learning and remembering, hippocampi generating spatial maps, synapses inducing LTP, and macromolecules binding and changing conformation. These levels are mechanistic levels in that they describe parts and wholes related as components to mechanisms or, more appropriately, as the activities of parts to the activities of the mechanism as a whole. It is a “sketch” because certain of its levels are poorly understood and because gaps exist in even the most well understood levels. This is often the case in the process of mechanism discovery.

The elaboration and refinement of such multilevel descriptions typically proceeds piecemeal with the goal of *integrating* the entities and activities at different levels (Craver and Darden 2001; Craver 2001). Integrating a component of a mechanism into such a hierarchy involves, first, contextualizing the item within the mechanism for the role to be explained. This involves “looking up” a level and identifying a mechanism that has the item as a component. Integrating involves, second, “looking down” a level and showing that the properties or activities of an entity can be explicated in terms of a lower level mechanism. For example, LTP might be integrated into a memory mechanism by looking up to see it as a component in a computational-hippocampal mechanism and by looking down to explain it in terms of its molecular level mechanisms. This understanding of mechanisms, levels, and integration yields a tidy taxonomy of interlevel experiments in neuroscience.

5. Interlevel Experimental Strategies. Interlevel experiments are tools for integrating the levels in hierarchical descriptions of mechanisms. Interlevel experiments tell us what the relevant entities and activities are, how they are nested in component/sub-component relations, and how the activities of the component entities fit into their mechanistic context.

As a first pass, experiments for testing mechanisms have three basic elements: (i) an experimental model (e.g., a strain of mouse), (ii) an intervention technique (e.g., electrical stimulation), and (iii) a detection technique (e.g., whole-cell recording). These elements are depicted in the abstract experimental protocol in Figure 5, which shows an experiment for a single mechanistic level. The connected circles and arrows represent a hypothesized mechanism putatively instantiated in an experimental model. On the left hand side of the figure are arrows standing for an intervention technique (I). (The intended sense of intervention might be explicated along the lines of Woodward 2000). The perturbation that is produced by I in the experimental preparation has “downstream” results which are detected or amplified using a detection technique (D).

It is easy to extend this view of experiments to interlevel experiments. Interlevel experiments are experiments in which the techniques for inter-

vening and detecting are targeted at different levels in the mechanistic hierarchy. For simplicity, start with experiments spanning two levels. The left hand side of Figure 6 exhibits a case of intervening to perturb a component in the lower level mechanism and detecting the consequences for a higher level role; these are bottom-up experiments. The right hand side of Figure 6 shows the opposite: intervention to perturb the higher level role and detection of the activities or properties of components in the lower level mechanism. These can be thought of as top-down experiments. I now want to consider three prevalent interlevel experimental strategies in contemporary neuroscience: activation strategies, interference strategies, and additive strategies.

5.1. *Activation Strategies.* Experiments exemplifying activation strate-




Figure 5




Figure 6

gies have a top-down structure; one activates, engages, triggers, or stimulates the role of interest and then detects the properties or activities of one or more putative components of the mechanism instantiating that role. (Activation strategies have been discussed in the context of functional PET and MRI techniques by Bechtel and Stufflebeam 1996 and Bogen 2001).

In the early 1970's, O'Keefe and Dostrovsky (1971) recorded the electrical activity of neurons in the rat hippocampus while the rats navigated a maze. The intervention in this case involves activating the spatial memory system by putting the rat in a maze. The detection technique is the electrical recording from hippocampal cells. They found that certain of these neurons generate bursts of action potentials whenever the rat enters a particular location while facing in a particular direction. These neurons have come to be called "place cells," and the region of space occupied by the rat when the place cell increases its activity is known as the cell's "place field." These place cells have slightly overlapping place fields that cover the animal's immediate spatial environment, and many believe for this reason that the hippocampus (or portions of it) could serve as a spatial map. These findings have recently been confirmed with multiunit electrodes that allow one to record from 70–150 pyramidal cells at once. Astonishingly, it is possible to *predict* the path taken by the rat on the basis of these recordings (Wilson and McNaughton 1993).

As compelling as these results are, they are, taken alone, far from establishing the mechanistic relevance of the hippocampus to navigation. For example, the electrical activity of the hippocampus and normal navigation may each be effects of a common cause. In that case, the activity of hippocampal cells would not be crucial for the mechanism and would in fact be incidental to the mechanism. Objections of this sort can be partly redressed by the second form of inter-level experimental strategy.

5.2 Interference Strategies. Interference experiments are bottom-up experiments in which one intervenes to diminish, retard, eliminate, disable, or destroy some component entity or activity in a lower level mechanism and then detects the results of this intervention for some higher level role. (The epistemic status of interference strategies has been discussed by Glymour 1994 and Bubb 1994; see also Pearl's 2000 discussion of "surgery" on a causal graph). Consider an example from recent gene knockout experiments on LTP and memory.

In late 1996, researchers at MIT, Columbia, and Cal Tech published a series of papers describing the effects of highly specific genetic deletions on the other levels in the memory hierarchy (McHugh et al. 1996; Rotenberg et al. 1996; Tsien et al. 1996a; Tsien et al. 1996b). The researchers invented a molecular scalpel for deleting the *NMDAR1* gene and for deleting it only in the CA1 region of the hippocampus. They then performed

detection techniques at each of the four levels in the multilevel theory. Knockout mice had difficulty learning the location of a submerged platform in a Morris water maze, a common behavioral-organismic experimental technique. Knockout mice swim randomly around the pool; controls quickly learn to swim directly to the platform. Multiunit recordings from knockout hippocampi revealed significant impairments in spatial map formation; the place fields were much larger and much less sharply defined. These deficits in spatial map formation are arguably the result of the absence of LTP since knockout synapses did not exhibit LTP under normal conditions.

This experiment is a bottom-up interference experiment with detection at multiple levels. The intervention technique interferes with the activities of the NMDA receptor by deleting the *NMDAR1* gene. The detection techniques register the results of this intervention on LTP, spatial map formation, and spatial memory.

Like activation strategies, interference strategies have their characteristic weaknesses. In the case at hand, for example, the intervention technique is, in effect, a form of cellular damage which may have unintended implications for cell function independent of the effect on LTP and, further, this cell damage may have unintended consequences for the organization of the hippocampus and the brain. One would be more certain of a role for NMDA-dependent hippocampal LTP in the mechanisms of memory if one could enhance memory by changing NMDA receptors and thereby enhancing hippocampal LTP. This is an additive strategy; it is the last I will discuss.

5.3. Additive Strategies. Additive experimental strategies are bottom-up strategies that involve intervening to stimulate, augment, hasten, intensify, or multiply some component in a mechanism. Additive strategies have been discussed by Bechtel and Richardson (1993).

Consider a recent additive experiment. In this case, researchers altered the ratio among types of the *NMDAR2* sub-unit in the mouse hippocampus (Tang et al. 1999). One type of *NMDAR2* subunit predominates in young mice; the other predominates in older mice. By influencing the patterns of gene expression, the researchers were able to get a transgenic strain of adult mice to express more NMDA receptors with the “young” subunit. This is the perturbation in the mechanism produced by the genetic intervention.

The detection techniques are scattered across a number of different levels. Using single unit recording techniques the researchers found that the transgenic NMDA current was greater than that in controls. Using electrophysiological techniques, they established that the LTP in these transgenic mice was different from that in adult controls and, in fact, that

it resembled the LTP typically seen in young mice. Finally, the researchers used a battery of behavioral-organismic tests (for, e.g., object recognition, fear learning, and spatial memory) to establish that transgenic mice learn faster and retain that learning longer than controls.

The ability to predictably alter memory by perturbing genes for the NMDA receptor addresses challenges to this multilevel description that cannot be addressed with activation and interference strategies alone. The capacity to control phenomena predictably by stimulating the mechanism helps to rule out both common causes and unintended effects as explanations for the experimental results. These experiments are by no means immune to criticism, but they represent a third source of convergent evidence for the contribution of the NMDA receptor (and LTP) to the mechanisms of learning.

6. Interlevel Experiments, Multilevel Integration, and Robustness. Working together, these three forms of interlevel experimental strategies provide an interestingly intratheoretic kind of robustness. Typically, “robustness” is defined in terms of manipulation or detection of a phenomenon via multiple theoretically independent routes (see e.g. Culp 1994, Wimsatt 1981). Interlevel experiments are interesting because they provide different independent paths of access to a phenomenon that are nonetheless part of the same multilevel theory. The independence of these paths is evidenced by the fact that each of the experimental strategies is prone to distinctive weaknesses that can be remedied, or corrected, by using the other experimental strategies. One lesson from this is that these experimental strategies cannot be evaluated adequately in isolation. A second lesson is that paths of access to a phenomenon can be theoretically independent but nonetheless intratheoretic.

Arguments about the existence and importance of LTP often turn on establishing such interlevel robustness. Integrating LTP into both lower and higher mechanistic levels is one of the means by which neuroscientists argue that LTP is real and of real significance for understanding the brain. Failure to find a lower level mechanism can often (though not always) be decisive in the fate of such a putative component, and the failure to find a role for that item, to situate it within a higher level mechanism, leaves LTP as no more than a curious laboratory phenomenon with no significance for the theories of neuroscience. Multilevel integration through interlevel experiments is thus a way of establishing the robustness of a phenomenon and thus securing its place in the ontic store of neuroscience. The mutually reinforcing, “back-watching,” fit of these experimental strategies makes them together a powerful combined attack in the pursuit of mechanisms.

These three experimental strategies (activation, interference, and ad-

dition) are increasingly common in the contemporary biological sciences generally and in neuroscience in particular. An adequate understanding of the findings of such experiments requires a clear explication of how these strategies work and what they show. In this paper, I have attempted to develop an adequate description of these experimental strategies by explicating the mechanistic structures that they are designed to elucidate. The task that remains is to use that explication to articulate the normative criteria for evaluating these different strategies. The ultimate status of LTP in the ontology of neuroscience is still to be negotiated; but its negotiation will benefit from careful thinking about the mechanisms into which LTP is to be integrated and the arguments by which that integration is to be achieved.

REFERENCES

- Bechtel, William and Robert C. Richardson (1993), *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*. Princeton: Princeton University Press.
- Bickle, John (1998), *Psychoneural Reduction*. Cambridge, Mass.: MIT Press.
- Bliss, Tim V. P. (1999), "Young Receptors Make Smart Mice", *Nature* 401: 25–27.
- Bogen, James (2001), "Functional Image Evidence: Some Epistemic Hot Spots", in P. Machamer, R. Grush, and P. McLaughlin (eds.), *Theory and Method in Neuroscience*. Pittsburgh, PA: University of Pittsburgh Press, 173–199.
- Bub, Jeffrey (1994), "Testing Models of Cognition Through the Analysis of Brain-Damaged Performance", *British Journal for the Philosophy of Science* 45: 837–855.
- Craver, Carl F. (2001), "Role Functions, Mechanisms and Hierarchy", *Philosophy of Science* 68: 53–74.
- (1998), *Neural Mechanisms: On the Structure, Function, and Development of Theories in Neurobiology*. PhD. dissertation. Pittsburgh, Pa.: University of Pittsburgh.
- Craver, Carl F. and Lindley Darden (2001), "Discovering Mechanisms in Neurobiology: The Case of Spatial Memory", in P. Machamer, R. Grush, and P. McLaughlin (eds.), *Theory and Method in Neuroscience*. Pittsburgh, Pa: University of Pittsburgh Press, 112–137.
- Culp, Sylvia (1994), "Defending Robustness", *Proceedings of the Biennial Meetings of the Philosophy of Science Association*. 46–57.
- Darden, Lindley (1991), *Theory Change in Science: Strategies from Mendelian Genetics*. New York: Oxford University Press.
- Darden, Lindley and Carl F. Craver (in press), "Discovering the Mechanism of Protein Synthesis", *Studies in the History and Philosophy of the Biological and Biomedical Sciences*.
- Engert, F. and T. Bonhoeffer (1999), "Dendritic Spine Changes Associated with Hippocampal Long-Term Synaptic Plasticity", *Nature* 399: 66–70.
- Frey, U. and R. G. M. Morris (1998), "Synaptic Tagging: Implications for Late Maintenance of Hippocampal Long-Term Potentiation", *TINS* 21: 181–188.
- Glymour, Clark (1994), "On the Methods of Cognitive Neursopsychology", *British Journal for the Philosophy of Science* 45: 815–845.
- Hacking, Ian (1988), *Representing and Intervening*. Cambridge: Cambridge University Press.
- Haugeland, John (1998), *Having Thought*. Cambridge, Mass.: Harvard University Press.
- Hebb, Donald O. (1949), *The Organization of Behavior*. New York: Wiley.
- Heil, John (1949), "Multiple Realizability", *American Philosophical Quarterly* 36: 189–204.
- Kim, Jaegwon (1993), *Supervenience and the Mind: Selected Philosophical Essays*. Cambridge: Cambridge University Press.
- Lomo, Terje (1966), "Frequency Potentiation of Excitatory Synaptic Activity in the Dentate

- Area of the Hippocampal Formation”, *Acta Physiologica Scandanavica* 68 (suppl. 277) 128.
- Machamer, Peter K, Lindley Darden, and Carl F. Craver (2000) “Thinking about Mechanisms,” *Philosophy of Science* 67: 1–25.
- Mackie, John L. (1974), *The Cement of the Universe: A Study of Causation*. Oxford: Oxford University Press.
- Malinow, Robert (1998), “Silencing the Controversy in LTP?”, *Neuron* 21: 1226–1227.
- Maletic-Savetic, M., R. Malinow, and K. Svoboda (1999), “Rapid Dendritic Morphogenesis in CA1 Hippocampal Dendrites Induced by Synaptic Activity”, *Science* 283:1923–1926.
- McHugh, T. J. et al. (1996), “Impaired Hippocampal Representation of Space in CA1-Specific NMDAR1 Knockout Mice”, *Cell* 87: 1339–1349.
- Nagel, Ernest (1977), “Functional Explanation in Biology”, *Journal of Philosophy*. 74: 280–301.
- Okeefe, J. and J. Dostrovsky (1971), “The Hippocampus as a Spatial Map. Preliminary Evidence from Unit Activity in the Freely Moving Rat”, *Brain Research* 34: 171–175.
- Pearl, Judea (2000), *Causality*. Cambridge: Cambridge University Press.
- Redish, A.D. and D. S. Touretzky (1998), “The Role of the Hippocampus in Solving the Morris Water Maze”, *Neural Computation* 10(1): 73–111.
- Rheinberger, Hans-Jörg (1997), *Toward a History of Epistemic Things*. Stanford: Stanford University Press.
- Rotenberg, Alexander, Mark Mayford, Robert D. Hawkins, Eric R. Kandel, and Robert U. Muller (1996), “Mice Expressing Activated CaMKII Low Frequency LTP and Do Not Form Stable Place Cells in the CA1 Region of the Hippocampus”, *Cell* 87: 1351–1361.
- Schachter, Daniel L. and A.D. Wagner (1999), “Remembrance of Things Past”, *Science* 285: 1503–1504.
- Schaffner, Kenneth (1993), *Discovery and Explanation in Biology and Medicine*. Chicago: University of Chicago Press.
- Schouten, Maurice K. D. and H. Looren de Jong (1999), “Reduction, Elimination and Levels: the Case of the LTP-Learning Link”, *Philosophical Psychology* 12: 237–262.
- Stufflebeam, R. S. and William Bechtel (1996), “PET: Exploring the Myth and the Method”, *Philosophy of Science* 64 (Supplement): S95-S106.
- Tang, Y., E. Shimizu, G. R. Dube, C. Rampon, G. A. Kerchner, M. Zhuo, G. Liu, and J. Z. Tsien (1999), “Genetic Enhancement of Learning and Memory in Mice”, *Nature*, 401: 63–69.
- Tsien, J. Z., D. F. Chen, et al. (1996a), “Subregion- and Cell Type-Restricted Gene Knockout in Mouse Brain”, *Cell* 87: 1317–1326.
- Tsien, J. Z., P. T. Huerta, and S. Tonegawa (1996b), “The Essential Role of Hippocampal CA1 NMDA Receptor-Dependent Synaptic Plasticity in Spatial Memory”, *Cell* 87: 1327–1338.
- Wilson, M. A. and Bruce McNaughton (1993), “Dynamics of the Hippocampal Ensemble Code for Space”, *Science* 261: 1055–1058.
- Wimsatt, William (1986), “Forms of Aggregativity”, in A. Donagan, A. N. Perovich, and M. V. Wedin (eds.), *Human Nature and Natural Knowledge*. Boston : D. Reidel, 259–291.
- (1994), “Robustness, Reliability, and Overdetermination”, in M. Brewer and B. Collins (eds.), *Scientific Inquiry and the Social Sciences*. San Francisco: Jossey-Bass.
- Woodward, James (2000), “Explanation and Invariance in the Special Sciences”, *British Journal for the Philosophy of Science* 52: 197–254.
- Zola-Morgan, S. and Larry Squire (1993), “Neuroanatomy of Memory”, *Annual Review of Neuroscience* 16: 547–563.