

Cellular neurosemiotics: Outlines of an interpretive framework

ANDREAS ROEPSTORFF

Introduction: neuroscience as a paradigmatic discipline for the 21st Century

There are at least two reasons to believe that neuroscience may become a paradigmatic discipline at the beginning of the 21st Century, like particle physics was it at the beginning of the 20th Century (Roepstorff, 2003). The first reason is internal to the discipline itself: in neuroscience questions about networks, agency, emergence, contextuality and historicity seem to pop up almost automatically. These are all central issues across a wide range of contemporary scientific practices and disciplines from the most hard core sciences to the softest humanities. In that respect, the place of neuroscience within the contemporary intellectual landscape resembles the role of particle physics 100 years ago, at that time with respect to questions of observability, relativity and causality.

None of these issues—networks, agency, emergence, and contextuality—were invented in neuroscience, but they appear to come into very interesting constellations in this field, probably because the subject matter in itself seems to provide structural affordance for these concepts. The most basic and fundamental question of neuroscience is, “how come that the brain, a vast network of relatively simple units, display this ability both for representing the world and for acting upon it”? It is very difficult to think about these matters without at some point having to consider issues like those mentioned above.

The second reason belongs to the sociology and anthropology of science rather than to the internal dynamics of neuroscience. It is an empirical fact that the last decade has seen the proliferation of a novel set of composite disciplines with or without hyphen—neuroaesthetics, neurophilosophy, neurohermeneutics, neuropedagogics, neurolinguistics, etc.—that complements old neuropsychology and neurophysiology. It is as if neurons and brains are becoming metaphors of grounding, a way to find a solid foothold for various disciplines from philosophy and the humanities to psychology and linguistics (Roepstorff, 1999b). This putative naturalisation is, however, not unidirectional, it is paralleled by a countercurrent attempting to culturalise nature, to demonstrate how facts about nature are not simply undisputed and eternal but reflect particular circumstances and historical trajectories (Roepstorff & Bubandt, 2003). Contemporary neurosciences are caught in the middle of this cosmological process, it is a hybrid issue located right between nature and culture. Also in that respect, it is paradigmatic for contemporary science (Latour, 1993).

Perhaps it is this element of double indeterminacy, of epistemological and ontological unfinishedness that makes “neuroscience” such a fascinating subject these days. It is interesting for a scientific insider because there are so many novel things to discover, so many new tool and technologies to play around with. But it is equally interesting from the perspective of an outside observer, who is fascinated by the exotic doings of the scientists (Roepstorff, 2002b), attracted to the strange facts and findings produced, and concerned about how they may reconfigure the grand narratives about man and woman, nature and culture and the relationship between them.

Assessing interdisciplinarity

Contemporary neuroscience has a bit of a frontier attitude to it. There are still new grounds to discover, new places to name and claim, new stories to tell. Why should one be interested in entering into this minefield under yet another heading, that of neurosemiotics¹? The most obvious reason is pretty banal; it is all about identity. Now that the phenomenologists², and the philosophers (Churchland, 1986) have taken their share of the brain why shouldn't the semioticians move in as well? We may call this line of argumentation a *symbolic* interrelation of semiotics and neuroscience: it would claim only weak and arbitrary connections between the respective subject matters. However, it would stabilise these by establishing a number of conventions which render them pragmatic, that is, useful and meaningful in novel discourses. Potentially, this would grant some credibility to semiotics via a connection to that well of discursive significance and credibility, which the neuroscience currently seems to be. This understanding of neurosemiotics is the one implied by the dictionary definition which defines neurosemiotics "as the branch of the neurosciences that investigate the neurophysiological basis of semiotic behaviour" (Bouissac, 1998). The problem with this delineation is, however, that although it may be attractive for semioticians to be able to claim to their territory "most of the field covered by the cognitive neurosciences" (op. cit.), it doesn't really make much difference to anyone other than semioticians. Consequently, "neurosemiotics today designates a rather virtual domain within the neurosciences" (op. cit.), and this appears to be as true today as when Bouissac published his definition.

However, once one has moved inside a semiotic vocabulary, one is not obliged to rely on symbolic connections only. Another possibility would be to examine putative *iconic* relations between semiotics and neuroscience, which could provide a rationale for a neurosemiotic approach. Are there similarities between the two fields, which render plausible the call for a neurosemiotic framework?

In the following, I will attempt such an examination. The first step is to apply a semiotic conceptual framework and vocabulary to phenomena within neuroscience. However, merely renaming well-known issues in a novel vocabulary does not suffice to demonstrate that neurosemiotics provides a meaningful heading. It only becomes relevant if the redescription makes a difference. In a recent paper, Veronica Box Mansilla and Howard Gardner (2003) have argued that the inherent quality of interdisciplinary work can be validated on three fundamental grounds: 1) the consistency with multiple separate disciplinary antecedents, 2) the balance in weaving together perspectives and 3) the effectiveness in advancing understanding. I propose to use these points as indicators—metaphorical *litmus tests*—for the potentials of a neurosemiotic approach.

Delineating cellular neurosemiotics

Contemporary neuroscience covers a vast number of subdisciplines at all levels genomics and proteomics via individual neurons to networks, brain regions and whole organisms, but most of the novel neuro-disciplines, neurohermeneutics, neurophilosophy,

-
1. For a fine overview of current neurosemiotic approaches see (Matrov, 2003).
 2. This framing has arguably been effectively initiated by Francisco Varela (e.g. 1996) and Shaun Gallagher (e.g. 2000), but the field is now coming into maturity through regular conferences and journals like *Phenomenology and Cognitive Science*.

neurolinguistics etc, only make sense at the highest level of description, the organism. This is also the case for neurophenomenology where the starting point for phenomenology, *what it is like to be xxx (in pain, introspecting, a bat etc)* (Varela, 2003) obviously is relevant for cognitive neuroscience (for an extended discussion see Jack & Roepstorff, 2003). However, the question appears increasingly irrelevant as one moves down the organisational hierarchy to the levels of brain regions, networks, and individual neurons.

In contrast, one of the potentials of neurosemiotics is that it may be applicable over a wide range of these levels since the mechanisms conceptualized by semiotics—communication, representation, exchange etc.—arguably are relevant at most of them. As noted by Donald Favareau (2002, note 2), “representation” occurs at a neuronal and neuroanatomical level in the theories of Gerald Edelman (Tononi & Edelman, 2000) and Antonio Damasio (e.g. 2000) two prominent contemporary neuroscientists, while questions of context, representations, anticipation and prediction are central in Karl Friston’s predictive modelling theory of brain function (e.g. 2002b; 2003). In the following, I will begin my examination of neurosemiotics at the levels of a single neuron and move on to simple interactions between neurons as they can be studied in the hippocampal slice preparation. Finally, I will discuss whether findings from “cellular neurosemiotics” may generalise on to higher levels of organisation such as brain regions, and organisms.

What is a neuron?

According to the standard text-book version, a neuron can be seen as a biological instantiation of an information processing device (Figure 1). At one end, the synapses on the bushy dendrites are the *receivers* of a neuronal signal in the form of *transmitter molecules* that act on receptors in the post-synaptic membrane. As a result, small currents propagate towards the cell body. These currents may be excitatory, this leads to a depolarization of the membrane potential of the neuron, or they may be inhibitory and prevent a depolarization of the neuron. If the neuron receives a sufficient amount of *signal* within a particular time window, it fires. This means that an electric pulse known as the action potential propagates down the axon at the other end of the neuron. The axon serves as a *conducting fibre* that allows the signal to jump from one node of Ranvier to the next till it reaches the axon terminals where presynaptic boutons are located. Here the action potential triggers the release of *neurotransmitter* into the synapse between the presynaptic bouton and the postsynaptic bouton. This, may, then, affect *receptors* on the postsynaptic boutons, and then the process continues. Merely looking at the concepts employed—“signal”, “transmission”, “receiver”, “transmitter” etc.—suggests that the standard understanding of neurons are grounded in a metaphor of information processing. In this version, the neuron becomes an advanced switch board that relatively passively transfers information from one end to another. This metaphor of communication and exchange allows for a first *rapprochement* between semiotics and neuroscience: similarities in the vocabulary allow for establishing Mansilla & Gardner’s first point: “consistency with multiple separate disciplinary antecedents”. The initial move from a neuroscience to a semiotic perspective can simply be done by lifting the usual words into a novel context.




Figure 1 A textbook version of a neuron, (reproduced with permission from www.enchantedlearning.com).

Neuronal activity is sign activity

In a seminal article in neurosemiotics, Donald Favareau (2002) performs this transformation. He suggests that moving the usual neurobiological vocabulary into a semiotic framing allows for an interpretation of the neuron where *code duality*, that is, transformation between digital and analogue representations (Hoffmeyer & Emmeche, 1991), becomes the key issue.

[T]he environmental surround that each neuron is situated in (and with which it interacts with most directly at the site of the synaptic cleft) is a Heraclitian world of ever-changing chemical and molecular interaction and constitution, whose analog representation (what neuroscientists call its “synaptic potential”) is constituted by whatever unique configurational state that environment is in at the moment of synaptic (which is presumed to be quantal) release. Conversely, the electric current generated within the neuron and which travels down the axon (referred to, semiotically enough, as an “action potential”) as a result of this analog release possesses all the attributes of a purely digital code: it is either wholly present or wholly absent, its amplitude is not variable, it does not decay over time or distance. Most critically: analog synaptic potentials generate digital action potentials which generate analog synaptic potentials which generate digital action potentials. This ongoing process of semiosis wherein the interactive, consequential interplay between digital and analog cell activity constitutes new signs and new information at every nodal (synaptic) point is, I believe, the starting point upon which the establishment of a discipline of neurosemiotic must be built (Favareau, 2002, 66).

Analysing the workings of the neuron in terms of *code duality* allows Favareau to get out of the textbook interpretation of neuroscience “the long-held neural conduit met-

aphor, — wherein 'information' flows through the circuitry of neurons in much the same way as electricity flows through a computer motherboard" (op. cit. p. 68). This switch-board model, where information is *in* the neuronal activity, is replaced with a semiotic understanding where "neuronal activity *is* sign-activity" (op. cit. p. 69). Favareau's vision thus becomes "an ongoing, dynamic process of sign-exchanging cells embedded in sign-exchanging brains embedded in sign-exchanging bodies embedded in sign-exchanging worlds".

In the following I will like to pursue the consequences of this idea through an examination of electrophysiological examinations in the hippocampal slice preparation.

Neurotransmission: interpretation and context

One of the best studied models of neurotransmission is the hippocampal slice preparation from rodent brains. Briefly, a rat is decapitated and the brain removed from the skull. The brain region known as the hippocampus is dissected out and cut into thin (400 μm) slices. These are placed on a thin nylon-mesh grid on the interface between a solution of nutrients and oxygen. In this way, the cells may be kept alive for many hours, and this makes it possible to study in great detail elements of synaptic transmission. Probably the best studied part of the hippocampus is the CA1 region where the large pyramidal cells are neatly organised with cell bodies on a visible line, the *stratum pyramidale*. These pyramidal cells receive excitatory inputs mainly mediated by the neurotransmitter glutamate through the Schaffer collaterals and commissural fibres. Activating these pathways brings the neuron closer to the firing threshold by depolarising the membrane potential. As importantly, inhibitory inputs bring the neuron further away from the firing threshold. These are mainly caused by inhibitory interneurons using the neurotransmitter GABA and located among the dendrites of the pyramidal cells in the *stratum radiatum*.

It is possible to insert a sharp glass microelectrode into the *stratum pyramidale* or the *stratum radiatum* and record the extracellular field potentials associated with the joint activity of a number of neighbouring neurons. Alternatively, one may penetrate into the body of an individual pyramidal cell and either passively record changes in the membrane potential or actively measure the currents needed to maintain the neuron at a particular state of depolarisation.

By stimulating the interneurons with an electric current, while applying various drugs to the solution and to the intracellular electrode, it is possible to isolate one component of the synaptic transmission: the fast GABA_A mediated synaptic inhibition (Nathan & Lambert, 1991). The end result is a preparation, where one may record the currents running into a single neuron as a result of the synchronous firing of a number of inhibitory neurons that are located just one synapse away.

Even at this very basic level, the system—understood as the postsynaptic neuron, where the activity is recorded and the presynaptic neurons that are stimulated—shows use-dependent plasticity. If two stimuli follow each other with an interval of 50 to 3000 ms, the size of the second response is significantly reduced both in amplitude and in decay time with a maximal reduction occurring at about 150 ms interstimulus interval (Figure 2).

It is generally thought that this short-term plasticity appears to be mediated by presynaptic *autoreceptors* located on the presynaptic boutons of the interneurons. This means that when GABA molecules are released, they bind to receptors on the neurons




Figure 2 Paired pulse inhibition of inhibitory postsynaptic currents (IPSC) recorded from pyramidal cells in CA1 of the rat hippocampus. A) in paired pulse stimulation with an interval of 150 ms, the second IPSC is notably smaller and decays faster than the first IPSC. D) Amplitude (squares) and half decay times (triangles) of second IPSC relative to the first IPSC as a function of interpulse interval (modified after Roepstorff & Lambert, 1994).

they are released from. As a result, the release of the neurotransmitter GABA to the second stimulation is inhibited by the release of neurotransmitter to the first pulse.

Does it make sense to describe this as a semiotic process? Let's take as our starting point Favareau's dictum: "neuronal activity *is* sign-activity" and rephrase the experimental situation along this line. This entails that not only the interaction between neurons, but also the interaction between the experimenter and the hippocampal slices

should be seen as an exchange of signs. When the experimenter stimulates the hippocampal slice, he provides a sign in the form of an electric pulse. The inhibitory interneurons may respond to this sign by initiating a cascade of further signs in the form of transmitter molecules that move towards neighbouring neurons, where they will induce inhibitory currents that will bring the cell further away from the cell-firing threshold. Some of the effect of these signs on a particular neuron is reflected in the intracellular recording shown in Figure 2. In this vocabulary the interneurons "interpret" the electric stimulation by releasing neurotransmitters and the pyramidal neuron, whose signal is being recorded, "interprets" this sign by a hyperpolarisation of the membrane potential, which, in turn, is being interpreted by the electrode inserted into the intercellular space³.

At this stage, a semiotic vocabulary offers nothing new. But what happens when the experimenter exposes the slice to a double pulse stimulation at an interpulse interval of 150 ms? It is quite obvious (Figure 2A) that the response measured to the second pulse is smaller than the response measured to the first pulse. In semiotic terms we may say that the interpretation of the electric stimulation is context dependent, as there is less inhibitory activity in the neuron when the stimulation is presented in the context of a previous stimulation within a particular time-window (Figure 2B).

This has important functional consequences. If we follow Favareau's idea that "code duality" is the key to neuronal functioning, the crucial step is whether the neuron will fire at a particular time, that is, whether it will send along a sign to the neighbouring neurons. This depends on a balance between the excitatory and the inhibitory inputs. The implication is that if the inhibitory input is reduced, the chances of an action potential following a particular excitatory input are increased.

This is seen in the phenomena in the hippocampal slice preparation known as "paired pulse facilitation". If one allows both inhibitory and excitatory inputs to affect the pyramidal neurons, these are more likely to fire in response to a stimulation if this has been preceded by a stimulation presented 50 ms -1000 ms earlier, (Figure 3). The mechanistic explanation for this phenomenon appears to be a reduction in the GABAergic inhibition similar to the one shown in Figure 2 (Nathan & Lambert, 1991). This allows, then, to phrase the finding depicted in Figure 2 in semiotic terms. The neuronal network connected to pyramidal cells in CA1 of the rat hippocampus is configured in such a way that the probability of the generation of a sign in response to a particular sign configuration in the afferent neurons is context dependent since it depends on the short term history of the network.

-
3. It may seem far-fetched to talk of "interpretation" in this situation, since in the standard "neural conduit" model, the neuron is usually considered to be more or less passively conducting the signals imposed on them. After 50 years of research, many of the mechanisms that link presynaptic depolarisation to the release of transmitters are known at a molecular level, but the link between excitation and amount of transmitter release is described statistically rather than mechanically. The standard assumption is that there is a constant stochastic release of vesicles from the nerve terminals, and that the probability of release is changed by the action potentials (Bennett & Kearns, 2000). This leads to a "quantal" model where neurotransmitters are released in discrete units. This allows potentially for sufficient semiotic freedom to warrant the notion of "interpretation", since the neuron, is influenced by the temporal and spatial context when actually doing something with the incoming stimulus/ the signs presented. This relatively simple scenario has, furthermore, been questioned recently (Burgoyne & Barclay, 2002; Vautrin & Barker, 2003).




Figure 3 Paired pulse facilitation in pyramidal cells measured as extracellular field potentials in stratum pyramidalis from rat hippocampus. A) the second response is facilitated at a stimulus interval of 150 ms. B second response superimposed on the first response (modified from Nathan & Lambert 1991).

In the following, I will attempt to rephrase this in terms of standard semiotic triangle (Nöth, 1990, 89).

The neuronal representamen

A representation is that character of a thing by virtue of which, for the production of a certain mental effect, it may stand in place of another thing. The thing having this character I term a representamen, the mental effect, or thought, its interpretant, the thing for which it stands, its object.

(CS Peirce, A Fragment, Collected Papers, 1.564, c. 1899)


The experiment examining the effect of paired pulse stimulation on the isolated GABAergic transmission may be described as a very simple semiotic system where the object or referent presented to the neurons consists of the individual electric pulses to the GABAergic interneurons. Figure 2 shows that every stimulation, each sign, opens up a time-window where a subsequent sign will evoke a smaller inhibitory potential. From the point of view of the neuron, this means that in the presence of both inhibitory and excitatory inputs, the sensitivity of the neuron to a second stimulation is increased within this time window (Figure 3). It seems thus that the one important aspect of the “sense”-made of a single stimulation to the inhibitory interneurons—the interpretant to use a peircean terminology—is a short term increase in sensitivity. This may be expressed in a “neuronal narrative”: *If the interneurons present another inhibitory signal within the time window, it is going to have less effect on the interpretation of the excitatory signs. They are therefore more likely to survive the code-duality transformation and create a new sign to be propagated from the pyramidal cell.*

Given that the mechanistic explanation of this effect is correct, we also have a description of the sign vehicle or representamen that mediates this effect: “be more sensitive to a second stimulation” is carried by a reduction in the amount of transmitter released into the synaptic cleft caused by an activation of the GABA_Aautoreceptor. This allows to render the experiment in a semiotic triangle (Figure 4B). In the context of a paired-pulse stimulation experiment, each stimulus is the *referent* of a semiotic process

in that neuron, which the experimenter records from. The *interpretant* evoked in the neuron by each interneuron stimulation is a short-term increase in sensitivity. Finally, the *representamen* can be given a physical interpretation in the form of a GABAautoreceptor that allows the system to react to this type of stimulation as a sign.

This representation of the neuronal network is of course done with reference to that very particular type of sign that the human experimenter presents to the neurons. The size of the stimulation, which enforces a number of closely located interneurons to fire in parallel, is much larger than would have occurred in normal neuronal communication. Therefore, we cannot be certain that neurons exchanging signs in a "natural" situation will show similar context dependent effects. However, the experiments suggest that the neuronal system has the capacity to interpret signs in a context dependent manner. Short-term plasticity as described for GABA synapses in the hippocampus is hence a very basic example of context-sensitivity with respect to time where "each response carries information about the temporal structure of the preceding pre-synaptic input spike train" (Fuhrman, Segev, Markram, & Tsodyks, 2002, quoted in Friston, 2003).


A Sense (interpretant)



Sign vehicle
(representamen)

Referent
(object)


B Short term increase in sensitivity



Decreased transmitter release

Stimulation of interneurons


C Actualised strengthening of specific connection



Increased transmitter release

Short train of stimuli

D Realised strengthening of specific connection



Increased transmitter release, increase in receptors

Long train of stimuli

Figure 4 Semiotic representations of neuronal processes in the hippocampus based on the semiotic triangle. See text for details

There are very good reasons to believe that in the brain, various forms of temporal context-sensitivity represent the rule rather than the exception.

Let us at this stage try to evaluate Mansilla & Gardner's three points for interdisciplinarity. The first point, "the consistency with multiple separate disciplinary antecedents" seems to be fulfilled. The narrative is able to move back and forth between the semiotic vocabulary and the neuroscience framework almost seamlessly. We have not introduced any major concepts novel to the neuroscience tradition, but peircean re-description of the experiment, which follows logically from Favareau's dictum: neuronal activity *is* sign activity, allows for a dynamic understanding of these notions. The key point is that the conduit metaphor of simple transmission is replaced by an understanding where the neurons are seen as active agents that react in context-sensitive albeit stereotypical ways. This is process of ascribing or granting agency to some entities is in semiotic terms, an actantialisation. The term "actant" is usually preferred over "agent" as it involves a softer version of agency, given the right context, even objects, or background figures may effectively be actants that come to play a role in the unfolding of events. This allows us to obtain a certain "balance in weaving together perspectives" (Mansilla & Gardner's second point) where the phenomena are given a more comprehensive interpretation by being embedded in two perspectives. The functional identification of GABA autoreceptors becomes, thus, at the same time the successful outcome of a classical neurophysiological analysis *and* the basic semiotic building block for the temporal interpretation of signs. Does this result in an "effectiveness in advancing understanding" (Mansilla & Gardner's third point)? This is of course the most crucial point. If the two descriptions are homologous, the outcome of the troublesome re-description is really meager. But lets retain a bit of optimism for the things to come. Perhaps shifting perspective is, at first, a matter of allowing shifts in the *style* of research (Fleck, 1979; Hacking, 1992; Roepstorff, 2001b), which affects the kinds of questions that one asks, the directions that one is pursuing. In the following, I will attempt to examine this idea further through an analysis the phenomenon known as long term potentiation.

The semiotics of long-term potentiation

Affect: to act upon or influence

Effect: Something produced by a cause or an agent

The Collins Paperback English Dictionary

One of the most intriguing neuronal mechanisms in the hippocampus is the phenomenon of long-term potentiation or LTP. First described by Bliss and Lomo thirty years ago (Bliss & Lomo, 1973), LTP is an increase in synaptic efficacy in the hippocampal slice preparation that may persist for several hours following a so-called tetanisation, a high frequency stimulation to the Schaffer collaterals. It is therefore an example of a use-dependent long-term change in neuronal connectivity, and it is currently the main model of memory formation in the mammalian brain (Malenka, 2003). Phrased in semiotic terms, the high-frequency stimulation—for the time of the experiment if not permanently—changes the context for the interpretation of subsequent stimulations.

During the three decades since the first discovery of LTP, numerous studies have looked into the neuronal and molecular mechanisms generating and sustaining LTP (extensively reviewed in Bliss, Collingridge, & Morris, 2003). It is generally considered

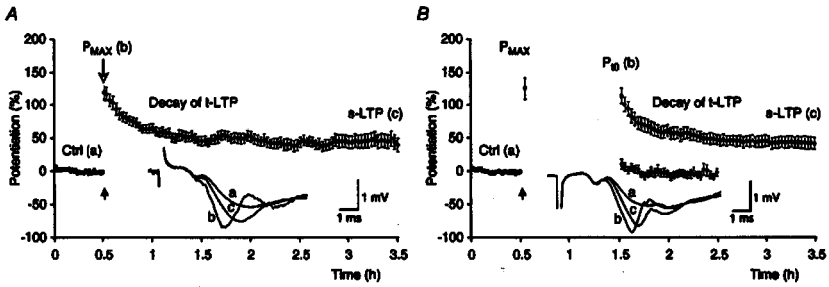


Figure 5 Transient and sustained phases of LTP in hippocampal slices. At the time indicated by the arrow, the slice receives a tetanic stimulation that induces potentiation. Representative field potentials at control (a) P_{MAX} (b) and at s-LTP (c) are shown in insert. In the presence of low-frequency stimulations (A), this potentiation decreases monoexponentially as a function of time. In the absence of stimulations (B), the potentiation remains the same for at least one hour, but decays as a function of time when low-frequency stimulation begins. This demonstrates that the decay of t-LTP is a function of activity rather than a function of time. Small insert above the field potentials show that simply leaving the slice alone for an hour without prior tetanisation does not induce potentiation (modified from Volianskis & Jensen, 2003).

that potentiation in CA1 of the hippocampal slice may be divided into three different phases, a post-tetanic potentiation (PTP) that decays in the time range of a few seconds, a short-term potentiation (STP) that decays in the time range of 30-60 minutes, and LTP lasting for more than 60 minutes. This understanding has, however, recently been modified by a very elegant set of experiments. The neurophysiologists Arturas Volianskis and Morten Skovgaard Jensen (2003) have demonstrated that if the hippocampal slice does not receive any new stimulation after tetanisation, there is no decay in the STP. However, as soon as new stimulations are presented, even after an hour of inactivity, the usual pattern of decay of STP is seen (Figure 5). This suggests that the decay of the phenomenon previously described as short-term potentiation is use-dependent rather than time-dependent.

This finding allows Volianskis and Jensen to discriminate between two types of LTP: sustained LTP (s-LTP) the "orthodox" understanding of LTP seen as the lasting potentiation in Figure 5 and transient LTP (t-LTP) formerly known as short-term potentiation (STP), which is seen as the decaying phase in Figure 5. Based on a very comprehensive set of experiments, they furthermore construct a dynamic model for synaptic plasticity consisting of (at least) four different stages (Figure 6). I think this model may be an instance of a small change in the understanding of synaptic dynamics. It suggests that the functional coupling and uncoupling of neurons rely not only on temporal parameters, but that the neurons may following particular exchanges, may remain in states of "tension". I think this is one example of a shift from a neuronal conduit metaphor to an understanding compatible with a semiotic perspective. In the following I will use a vocabulary imported from an analysis of semiotic existence (Greimas & Fontanille, 1991) to pursue this idea. I shall propose that the four stages of the synapse in the model can be paralleled with the notions of *virtualised*, *actualised*, *realised*, and *potentialised*

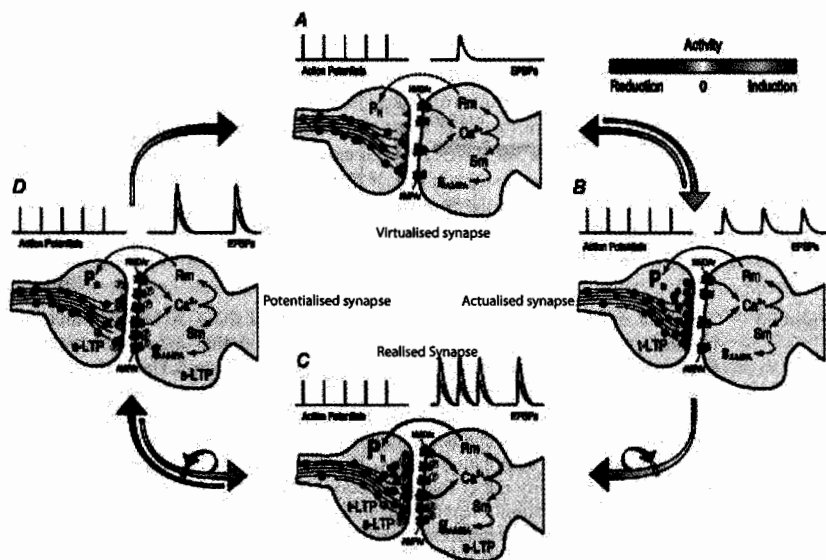


Figure 6 Modes of synaptic existence, modified from (Volianskis & Jensen, 2003), see text for details

synaptic stages as various positions in a semiotic square (op- cit. 25-7, 87-9).⁴ This semiotic redescription dissociates neuronal coupling into two orthogonal parameters the *affect*, which describes the extent to which activity in the presynaptic cell generates activity in the postsynaptic cell and the *effect*, which describes the consequences of a functional linkage in terms of the size of the excitatory post synaptic potential (EPSP).

The choice of the concepts *affect* and *effect* pays a humble homage to the fact that Greimas and Fontanille developed the four types of semiotic existence with reference to a semiotics of passion. Indeed, as we shall see, it suggests rephrasing the formation of strong synaptic connections between two neurons in terms of a love-story between the pre- and the postsynaptic cell.

4. The four types of semiotic existence are one of the most complicated ideas in an already very dense theory. Rather than attempting a detailed analysis of the concepts, I will, thus, proceed by way of more concrete examples. The argument for the translation of the four stages (virtualisation, actualisation, realisation and potentialisation) into vernacular (absence, wanting, presence, waiting) is based on (Fontanille, 1996; for discussion see Komoch, Roepstorff, & Vedel, 1996).

The synaptic “modes of existence”

In Volianskis & Jensen’s model there is a relatively low probability of transmitter release following a presynaptic spike in the normal state of a synapse, that is, prior to experimental tetanisation, and the resulting EPSPs are small. This is figuratively indicated by the low ratio between EPSPs and action potentials ($EPSP/AP = 1/5$) and the small EPSPs in Figure 6A. Rephrased, the affection between the two neurons are low, as is the effect of a functional coupling. In semiotic terms, this may be called the *virtualised* stage of the synapse. If the synapse receives a short train of presynaptic pulses, t-LTP may be induced (Figure 6B). This leads to an increase in the probability of transmitter release ($EPSP/AP=3/5$) without changing the effect of EPSPs. This becomes, then, the *actualised* stage of the synapse. The size of the t-LTP is a function of the stimulation frequency, but it is not related to the absolute number of spikes. In the absence of stimulation, this state of the synapse may persist for a long time, and this is the main finding of (Volianskis & Jensen, 2003). Although the actualised synapse appears stable with respect to time, it is highly affected by patterns of activity. If an actualised synapse receives a lot of stimulation (further spike trains), s-LTP may be induced along with the t-LTP (Figure 6C). This leads to a very active synapse with a high probability of transmitter release (the reliability of the synapse $EPSP/AP = 4/5$), while an increase in the postsynaptic AMPA receptors leads to a high efficacy (a high input/output relationship depicted as the red EPSPs in Figure 6C). There is a correlation between the number of spikes and the size of the s-LTP, but no correlation with spike frequency. This is then *realised* stage of the synapse. If, however, the actualised synapse only experience low frequency stimulations, it may again revert to its virtualised stage (Figure 6A).

If the realised synapse only receives sparse inputs, t-LTP will decay and, as a result, the reliability of the synapse decrease ($EPSP/AP = 2/5$, Figure 6D) but the efficacy may still remain high. We shall call this a *potentialised* synapse. Given a high activity, this synapse may again be realised (Figure 6B). Alternatively, a prolonged, low frequency stimulation may revert the potentialised neuron revert to its original virtualised stage (Figure 6A).

I propose that these synaptic dynamics may be described in semiotic terms, as an unfolding of the relationship between two neurons. If Volianskis and Jensen are right; synapses seem to code for the following dynamic relationship (Figure 7). In a normal “virtualised” synapse, it is difficult for the presynaptic neuron A to engage in a relationship with the neuron B as the affect, the synaptic reliability indicated by the AP/EPSP relationship, is low. It is as if the two neurons don’t take an interest in each other—in Greimasian terms: they are a non-conjoined and the mode of existence may be described as “absence” (Fontanille, 1996). However, a short spike train—in Volianskis and Jensen’s hands five spikes suffice—establishes a tension between neuron A and B. The likelihood that A will affect B is increased. The affect is high—but the effect of the affection, the synaptic efficacy, is still low. We may say that the synapse is now in an actualised state, characterised by a wanting (Fontanille, 1996) and in Greimasian terms the relation has become *disjoined*.

This sets up an affective relation between the two neurons. The faster the neuronal burst, the stronger the tension, the more likely A is to affect B. The proposed mechanism for t-LTP is highly interesting in semiotic terms: the post-synaptic neuron appears to send out a retrograde messenger—like the helper in a narrative—that allows the presynaptic neuron to increase the synaptic efficacy (in Figure 6B indicated by an increase in

The synaptic modes of existence

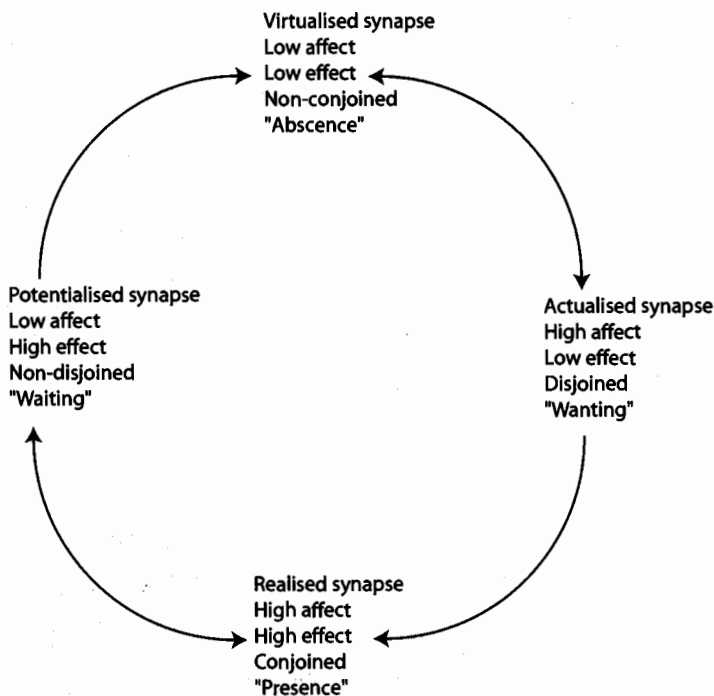


Figure 7 The synaptic modes of existence following Volianskis and Jensen's model of synaptic interaction. See text for details

the number of docked vesicles). In other words, actualising the synapse by increasing the affection between neuron A and neuron B requires — as is usually the case for human actants — the activity of both the neuronal actants.

The actualisation of the synapse opens a functional window where the activity of neuron A is likely to affect neuron B. It is the great finding of Volianskis and Jensen that the width of this window is governed by activity rather than by time. This changes the understanding of neuronal tension from a somewhat mechanical model that simply decays as a function of time to something that may be likened to falling in love at first sight in a coincidental encounter on the street. Traces of such event may persist in the memory for long, but it is unstable position that will be affected by further interaction. It seems that neurons follow a somewhat similar pattern, if neuron A utilises this "window of further engagement", that is, if it fires rapidly, this may lead to the generation of s-LTP which will stabilise the relation both by increasing the likelihood that A affects B (synaptic reliability) and the effect of this effect (the synaptic efficacy). The synapse has now been realised, and the two neurons are tightly coupled in that activity in the former is very likely to activate the latter. In a Greimasian vocabulary the two agents have become conjoined and the synapse is characterised by "presence".

This step appears to be irreversible: once the synapse has been realised, there is no turning back to an actualised "state of wanting". However, if the relationship is not maintained, if neuron A does not continue to engage passionately with neuron B, the synapse becomes potentialised with only s-LTP left. In Greimasian terms, the synapse has become non-disjointed like a relationship between two former lovers, and it is characterised by "waiting". The synaptic reliability, the affect, is low, but if A succeeds in activating B, the effect is still large. The transformation into potentialisation is reversible, the synapse may again be realised, that is, be effective and affective with a strong functional coupling.

Phrased in these terms, the synaptic dynamics, as proposed by Volianskis and Jensen, unfold a narrative about the establishing and breaking a relationship between two actants. The story begins when the A sends out a spike train which is picked up by B. Like "love at first glance" this establishes a tension in an actualised synapse, that in the absence of activity may remain for a long time. The synapse is in a tense state of "wanting" in the two senses of the word, at the same time a desire and a lack is installed. When A and B continue to interact following the first incidence, the trajectory depends on the level of activity. Without proper courting, the synapse declines to its virtual "low affect, low effect" state. If there is a strong level of activity, however, the relationship enters its realised stage, like an idealised marriage, where the activity of one is tightly coupled to but still independent of the activity of the other. But even the realised stage depends on constant activity. In the absence of activity s-LTP dies out, the synaptic affection decreases and the two neurons appear, again, to live separate lives, although the effect of an interaction is still large. The synapse is now potentialised, low effect, high affect, that may either be reinvoked to a relationship characterised by "presence" or deteriorate to the virtualised stage of "absence".

Evaluating the neurosemiotic description

Volianskis and Jensen's model of synaptic dynamics is very attractive. It is so new that it still remains to be seen whether it will survive the tough battle for recognition within the LTP-field (Bliss et al., 2003; Lisman, 2003), but it has already made it to the cover of one of the best physiological journals. It could be argued that it does not explicitly discuss some of the classic properties of LTP (Malenka, 2003) such as associativity and cooperativity, notions which reflect that firing in neighbouring synapses influences site-specific induction of LTP, however, this can be incorporated by expanding the model (both in the neuroscientific and in the semiotic formulation) to take into account the activity of neighbouring neurons. In semiotic terms, this means acknowledging the existence of a "neuronal sociality" where the general activity of many agents—closely located neurons—forms a significant context for the specific interaction among two neuronal actants.

In the context of this paper, the important issue is becomes how does this application of a neurosemiotic framework play itself out vis-à-vis Mansilla & Gardner's criteria for interdisciplinarity? At the first approximation, it allows for the same kind of redescription of the LTP phenomena that we previously saw for the paired-pulse stimulations. The particular stimuli presented to the neurons, short bursts and long bursts, are referents for semiotic processes that generate in the neurons the two interpretants "actualised strengthening" and "realised strengthening" (Figure 4 C-D). As with the paired-pulse stimulation, the representamens are mechanistic processes, but there is no agreement

within the LTP-community about what the substrates are, and not even whether the effects are post- or presynaptic (Lisman, 2003).

Phrasing LTP in terms of a Peircean triangle therefore allows for the first two criteria "the consistency with multiple separate disciplinary antecedents" and "balance in weaving together perspectives" to be met. However, the application of a more structural effective semiotic model, exemplified in (Figure 7), seems to open the third criterion for critical examination. Does an "effectiveness in advancing understanding" arise from this model? At this stage, this can not be settled, but there seems to be a number of prominent avenues. The application of a semiotic square highlights Favareau's dictum—*neuronal activity is sign activity*—by emphasising that the standard neuroscience notion of "synaptic plasticity" can be redescribed as particular processes whereby neuronal actants modify their internal relations along two dimensions: as a matter of increasing neuronal affectivity, which leads to a temporal coupling between neurons (activity in one neuron is linked on to activity in a connected neuron), and as a matter of neuronal effectivity (when neuron A affects neuron B it is likely to have consequences for whether B will fire an action potential. This extends on the passive conduit metaphor where neurons are either conducting or not conducting. These two states of absence/virtualisation and presence/realisation, are supplemented by two dynamic states of wanting (actualisation) and waiting (realisation). In semiotic terms (Brandt 1994) these are *modal* stages whose properties can only be understood on the basis of their trajectory. In all, this allows for a consistent understanding of how neurons exchange signs, where short term (actualisation) and long term (potentialisation) historicity plays a role. Volianskis and Jensen state in their paper and in conversations about the present paper, that the model as presented is indeed a simplification of the underlying mechanisms that they find. It is not likely that the mechanisms here described as *effect* and *affect* are in actual neuronal exchange discretised in two states high and low. Rather, a number of different molecular processes both post- and presynaptic seem to come together under these very gross headings. The semiotic redescription is, in other words, not a mechanistic understanding of the actual neuronal processes. It is a matter of applying a different model. But models matter. In hindsight, it may seem incredible that it took 25 years of research on LTP for someone to come up with the "forgotten experiment" where one simply left the slide alone following tetanisation. Whether it is an obvious question to ask what happens in the absence of stimulation depends on the general conceptualisation of neurons and neuronal processes. If neurons are passively conducting entities, the question seems obsolete. Who would ask what goes on in an electric cable when the current is off? However, if neurons are seen as semiotic entities, the question of the missing sign, of the expected sign, the anticipated sign and the novel sign are all equally relevant.

It is quite obvious that somewhere in the functioning of neuronal systems, dynamics of anticipation, expectancy, and tension must emerge. It seems attractive, at least as a hypothesis, if they occur already at the levels of interacting neurons. Examining the synaptic model in terms of a semiotic square may allow for new questions to emerge. Is it a co-incidence that the transformation from the subcontraries (actualised, potentialised) to the contraries (virtualised, realised) appears to be irreversible, or does it reflect more general semiotic processes where modal operators appear to be important "motors" in establishing emergence of meaning (Brandt, 1994)? Are there alternative trajectories through the model, and how can they be made understandable in terms of other examples of interacting actants?

Furthermore, some of the problems that seem to impose themselves on researchers within a traditional framework, are rendered more or less obsolete. One example is the discussion between Jeff Lichtman and Joshua Sanes on one hand and John Lisman on the other about whether LTP is a "social construction" (Lisman, Lichtman, & Sanes, 2003). This is phrased in terms of a standard conflict (Hacking, 1999; Roepstorff, 1999a) between social constructionists (Sanes and Lichtman) who believe that LTP may simply be the result of a particular human ordering of continuous universe and the naturalist (Lisman) who insists that LTP represents a natural kind. The neurosemiotic reformulation allows to transcend this dichotomy by insisting that both LTP experiments and the real-life neuronal processes should be characterised in terms of sign-exchanges, and it is not important whether a particular sign is "natural" or "cultural". It is obvious that when experimenters present an electric stimulus to the slice, this is very different from the type of signs that neurons exchange among themselves. However, experiments like those conducted by Volianskis and Jensen examine and indeed demonstrate some of the capacities of the neurons to react to particular types of signs. Here the critical analysis is how signs are processed and which types of signs may be interpreted differentially.

The neurosemiotic reformulation allows, in other words, for a consistent articulation of problems and for a creative generation of novel hypotheses. Potentially, this may lead to a fulfilment of Mansilla & Gardner's three criteria.

A question of levels: cells, modules or organisms?

The first is that whose being is simply in itself, not referring to anything nor lying behind anything. The second is that which is what it is by force of something to which it is second. The third is that which is what it is owing to things between which it mediates and which it brings into relation to each other.

(CS Peirce, 'A Guess at the Riddle', CP 1.356, c. 1890)

When I gave the talk behind this article (Roepstorff, 2002a), my main hypothesis was that the notion of neurosemiotics did not make much sense at the level of neurons, only carried allegoric relevance at the level of brain regions, but that it could be sensibly applied at the level of organisms, particularly humans. However, as it should be clear from this article, working through the dynamics of neuronal exchange has made me change that view. I have basically tried to argue that it *does* make sense to apply a semiotic vocabulary even at a neuronal level. Perhaps the phrasing could be even stronger by accepting Donald Favareau's dictum: neuronal activity *is* sign activity. This is, however, only the first step in the examination: as my renderings of processes at a neuronal level suggests (Figure 4) the semiotic mechanisms are very different from the type of processes that semioticians usually study in that all three elements, the referent, the interpretant and the representamen can be described by their physical and mechanical characteristics. If we should excel in neurophenomenology at this level, we may say that the life-world of individual neurons consists only of other neurons and of molecules, there is no reference to an "outside". This suggests that one should designate processes at a neuronal level as "protosemiotic" rather than as full-blown semiosis⁵. There are good

5. I owe this idea to Per Aage Brandt in a discussion of an earlier version of this paper.

reasons to take this position: it would render the cellular level an obvious pragmatic starting point for semiotic studies, and it would avoid that eternal semiotic regress, which biosemiotics potentially may lead to (if genes and proteins are semiotic, then why other molecules and, in the end, all matter (Hoffmeyer, 1996)). I am not saying that it is impossible to apply a semiotic vocabulary beyond the level of the individual cell or neuron, but it seems to me that the neuron is an obvious starting point, also because at this level a notion of actantiality clearly makes sense. In my understanding, actantiality is a key term in any semiotic consideration⁶.

I have not had the possibility to discuss in this paper how the neurosemiotic approach would play itself out at higher levels such as brain regions and individual organisms, and at this stage I can only draw the very rough contours. It seems to me that a neurosemiotic framework is highly compatible with one of the most convincing hypothesis of mammalian brain functioning, Karl Friston's predictive modelling (Friston, 2002a, 2002b, 2003). Starting out from a discussion of functional segregation and integration between brain modules, Friston argues that the brain can be likened to a hierarchical, Bayesian structure. Higher levels constantly estimate causes in the environment and the resulting incoming inputs to the senses, and these predictions are via backward connections in a top-down fashion sent to lower levels, where they are compared with the bottom-up representations driven by the actual environmental input to the senses. In this conceptualisation, representation, context, anticipation and prediction are central organisational principles at all levels of brain functioning from sensory regions to cognitive functions. The evaluation between expectancy and occurrence becomes one of the central 'currencies' of the brain, and a strong discord between the anticipated event and the actual sensory input leads to an updating and/or reconfiguration of the predictive model.

I think this model of brain functioning lends itself to a semiotic reformulation. One of the main points in the approach in this paper is that "contextuality" and "interpretation" do not emerge *ex nihilo* somewhere between the level of neurons and the level of the organism. Contextuality, historicity and interpretation are always there, also at the level of individual neurons communicating in an isolated hippocampal slice. The main problem becomes, rather, how these many contextualised and interpreted neuronal events lead to relatively stable representations of the environment and relatively stable prescriptions for actions. How come, that although only other neurons and chemical substances exist in the *Umwelt* of the individual neuron, they come to code for events and objects both in the outer world of things and agents, and in the inner world of imagined mental processes?

The key to this understanding lies, Friston claims, in an interplay between functional segregation, which anatomically and functionally modularises perception and action, and functional integration that ties the elements together by modulating the activities of the modules. It is functional segregation that allows for the many highly successful attempts at functional brain imaging, where particular characters such as "perception of faces" (O'Craven & Kanwisher, 2000) and "perception of speed and colour" (Chawla, Rees, & Friston, 1999) are processed in particular brain regions, and it is functional integration that allows the activities of these areas to be modified by various top-down processes like imagination and attention.

6. Taking the cellular level as the (pragmatic) starting point for a neurosemiotic analysis would align the approach with Francisco Varela's analysis of biological autonomy (Varela, 1987).

A neurosemiotic interpretation of these principles would suggest that at the level of individual brain regions, we have already moved beyond protosemiosis because the sign-processes can no longer be understood only in terms of the physical exchanges between neurons. Although the individual neurons in the various brain regions live out the normal "what it is like to be a neuron" life, the types of events, to which they react, are already a highly specific subset of the inner and/or outer world. Consequently, the scientific examination is best seen as an exchange of signs between the experimenter and the brains of the experimental subjects. Functional brain imaging is thus an attempt to identify those regions of the brain that are capable of responding differentially to the particular signs that the subjects are presented with and to examine how this interpretation may be affected by contextual and by mental factors. This is clearly a semiotic process.

Very importantly, the model seems also to allow for a straight-forward interpretation of how "communication" and other normal human semiotic competences and practices may play themselves out with respect to the brain. Tony Jack and I have argued elsewhere that cognitive brain imaging experiments can only be understood properly, when it is acknowledged that an exchange of a script for perception and/or action is a major part of cognitive experiments (Jack & Roepstorff, 2002; Roepstorff, 2001a). Furthermore, Chris Frith and I have argued that such processes, which by extension of the top-down metaphor may be described as top-top interactions (Roepstorff & Frith, in press), is a very important property of human cognition. We did not describe this in a semiotic vocabulary, but the translation is almost trivial, and it lends itself nicely to an implementation in a Friston-like understanding of the brain. One outcome of top-top interactions and script exchanges is a high-level specification of context, anticipation and prediction. Through these processes, the symbolic capability and semiotic openness of human cognition allows for an exchange of partial specifications of the most "top" levels of the human mind. As they are—somehow—translated into those types of signs that neurons understand, they are fed down through the hierarchical layers as specifications of context and configurations of expectancy, against which the bottom-up perceptual processes are being compared. One of the key elements in human cognition may thereby be top-top interactions that, by sharing the external context for the interpretation of events, allow for an internal synchronisation of neuronal context between individuals.

This section opened with a quote from CS Peirce that stated differences between three key peircian notions *firstness*, *secondness* and *thirdness*. Perhaps this allows for a semiotic interpretation of the difference between processes at the level of neurons, at the level of brain regions and at the level of the human mind. At the level of neurons, the signs exchanged do not point to anything outside of the neuronal code, in peircean terms only *firstness* "not referring to anything nor lying behind anything" is represented. At the level of brain regions, each individual neuron still only represents firstness, but as a consequence of the particular patterns of connectivity the activity carries reference outside of the neuronal activity itself, it is *secondness* "that which is what it is by force of something to which it is second". Finally, at the level of the organism, the scripts—the "highest" semiotic processes—specify the most general level of contextuality and pattern for activity. It is hence *thirdness* "that which is what it is owing to things between which it mediates and which it brings into relation to each other".

This understanding of the human mind seems both compatible with semiotic analysis and with the present knowledge of brain organisation, and it seems to allow both for hypothesis generation and for trying out a conceptual framework. It appears thus to be worthy of testing against the criteria identified by Mansilla & Gardner.

Conclusion: symbolic, iconic or indexical relations between semiotics and neuroscience?

I have in this paper claimed that an alliance between semiotics and neuroscience holds a number of potentials. Perhaps the most obvious attraction is symbolic, each discipline thus comes to adorn itself with a novel set of feathers. However, I have tried to argue that there are also more fundamental attractions: there seems to be so much conceptual overlap between the two disciplines that it warrants a closer examination. It is possible that the similarities are merely iconic, as my description of the synaptic exchanges in terms of a love-story between two actants. In my understanding, even iconic similarities would be important, and it could provide a framework that may prove fruitful to both traditions. A redescription in a different vocabulary may appear as being a matter of models, but models matter: they open up particular venues for research and render some questions obvious and others obsolete.

However, I am inclined to believe that the similarities may in certain areas run deeper than merely iconic similarity, that there are also structural similarities and indexical relations between the two disciplines. At their heart, they both have an examination of what happens, when interacting agents are caught up in complicated networks that seem to transcend the apparent boundary between the real and the imaginary. Paradoxically, it may be one of the most obscure branches of semiotics, the Paris school of structural semiotics that holds some of the important tools for this examination.

Acknowledgements

This article was supported by a grant from the Danish Research Fund to Centre for Functionally Integrative Neuroscience. The paper has benefited from constructive inputs from Per Aage Brandt, Chris Frith, Jakob Hohwy, Morten Skovgaard Jensen, Martin Skov, Arturas Volianskis, Svend Østergaard, Mikkel Wallentin and Don Favareau. However, this is very much an explorative paper, and I fear that it has not been able to solve all the problems they identified, answer all the questions they raised and dismantle all their rightful scepticism. It would never have been written without Joachim Schult's very competent workshop at TheorieLabor, Jena in 2002 and his efficient and patient role as editor of this special volume.

References

- Bennett, M. R. & Kearns, J. L. (2000): Statistics of transmitter release at nerve terminals. *Prog Neurobiol*, **60**(6), 545-606.
- Bliss, T., Collingridge, G. & Morris, R. (eds.) (2003): Special issue on LTP. *Philos Trans R Soc Lond B Biol Sci*, **358**(1432).
- Bliss, T. & Lomo, T. (1973): Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *Journal of Physiology*, **232**, 331-356.
- Bouissac, P. (1998): Neurosemiotics. In P. Bouissac (ed.). *Encyclopedia of Semiotics* (pp. 446-447). Oxford University Press: Oxford.
- Brandt, P. A. (1994): *Dynamique du sens. Études de sémiotique modale*. Aarhus University Press: Aarhus.

- Burgoyne, R. D. & Barclay, J. W. (2002): Splitting the quantum: regulation of quantal release during vesicle fusion. *Trends Neurosci.* **25**(4), 176-178.
- Chawla, D., Rees, G. & Friston, K. J. (1999). The physiological basis of attentional modulation in extrastriate visual areas. *Nat Neurosci.* **2**(7), 671-676.
- Churchland, P. S. (1986): *Neurophilosophy. Toward a Unified Science of the Mind-Brain*. MIT Press: Cambridge, Massachusetts.
- Damasio, A. R. (2000): *The Feeling of What Happens: Body and Emotion in the Making of Consciousness*. Harcourt: New York.
- Favareau, D. (2002): Beyond self and other: On the neurosemiotic emergence of intersubjectivity. *Sign System Studies*, **30**(1), 57-100.
- Fleck, L. (1979): *Genesis and Development of a Scientific fact*. Chicago University Press: Chicago.
- Fontanille, J. (1996): Le trope visuel, entre presence et absence. In G. i (Ed.), *Protee, 24-1 La Topologie visuelle*. Université de Quebec a Chicoutimi: Chicoutimi.
- Friston, K. (2002a): Beyond phrenology: what can neuroimaging tell us about distributed circuitry? *Annu Rev Neurosci.* **25**, 221-250.
- Friston, K. (2002b). Functional integration and inference in the brain. *Prog Neurobiol.* **68**(2), 113-143.
- Friston, K. (2003): Learning and inference in the brain. *Neural Networks*, **16**, 1325-1352.
- Fuhrman, G., Segev, I., Markram, H., & Tsodyks, M. (2002): Coding of temporal information by activity dependent synapses. *Journal of Neurophysiology*, **87**(140-148).
- Gallagher, S. (2000): Philosophical conceptions of the self: implications for cognitive science. *Trends in Cognitive Sciences*, **4**(1), 14-21.
- Greimas, A. J., & Fontanille, J. (1991): *The semiotics of passions. From states of affairs to states of feelings*. University of Minnesota Press: Minneapolis.
- Hacking, I. (1992): "Style" for historians and philosophers. *Stud. Hist.. Phil. Sci.*, **23**(1), 1-20.
- Hacking, I. (1999): *The social construction of what?* Harvard University Press: Cambridge, MA.
- Hoffmeyer, J. (1996): *Signs of meaning in the universe*. Indiana University Press: Bloomington.
- Hoffmeyer, J., & Emmeche, C. (1991): Code duality and the semiotics of nature. In M. Anderson & F. Merrell (Eds.), *On Semiotic Modeling* (pp. 166-177). Mouton de Gruyter: New York.
- Jack, A. I. & Roepstorff, A. (2002): Introspection and Cognitive Brain Mapping: from Stimulus-Response to Script-Report. *Trends in Cognitive Sciences*, **6**(8), 333-339.
- Jack, A. I., & Roepstorff, A. (eds.). (2003): *Trusting the Subject? The use of introspective evidence in cognitive science volume 1*. Exeter: Imprint Academics (also published as *Journal of Consciousness Studies*, 10(9-10)).
- Komoch, A., Roepstorff, A., & Vedel, K. (1996): Metaforer i Urbino. *Almen Semiotik*, **11/12**, 191-200.
- Latour, B. (1993): *We have never been modern*. Harvester Wheatsheaf: New York.
- Lisman, J. (2003): Long-term potentiation: outstanding questions and attempted synthesis. *Philos Trans R Soc Lond B Biol Sci*, **358**(1432), 829-842.
- Lisman, J., Lichtman, J. W. & Sanes, J. R. (2003): LTP: perils and progress. *Nature Reviews Neuroscience*, **4**, 926-929.
- Malenka, R. C. (2003): The long-term potential of LTP. *Nature Reviews Neuroscience*, **4**, 923-926.
- Mansilla, V. B. & Gardner, H. (2003): *Assessing interdisciplinary work at the frontier. An empirical exploration of "symptoms of quality"*. Available: <http://www.interdisciplines.org/interdisciplinarity/papers/6> [2003, 03/12].

- Matrov, D. (2003): *Neurosemiotic approaches and interpretations*. Unpublished Bachelor's thesis, Department of Semiotics, Tartu University, Tartu, Estonia.
- Nathan, T. & Lambert, J. (1991): Depression of the fast IPSP underlies paired-pulse facilitation in area CA1 of the rat hippocampus. *J Neurophysiol*, **66**, 1704-1715.
- Nöth, W. (1990): *Handbook of Semiotics*. Indiana University Press: Bloomington, IN.
- O'Craven, K. M. & Kanwisher, N. (2000): Mental imagery of faces and places activates corresponding stimulus-specific brain regions. *J Cogn Neurosci*, **12**(6), 1013-1023.
- Peirce, C. S. (1931-58): *Collected Papers of Charles Sanders Peirce*. Harvard University Press: Cambridge, Mass.
- Roepstorff, A. (1999a): Deconstructing Social Constructionism. *FOLK, Journal of the Danish Ethnographic Society*, **41**, 139-154.
- Roepstorff, A. (1999b): Det Neurale Menneske - et antropologisk perspektiv. In O. Høiris & H. J. Madsen & T. Madsen & J. Vellev (Eds.), *Menneskelivets mangfoldighed*. Moesgaard: Aarhus Universitetsforlag & Moesgaard Museum.
- Roepstorff, A. (2001a): Brains in Scanners: An Umwelt of cognitive neuroscience. *Semiotica*, **134**(1/4), 747-765.
- Roepstorff, A. (2001b): *Facts, Styles and Traditions. Studies in the ethnography of knowledge*. Unpublished ph.d. dissertation, Aarhus University, Aarhus Universitet, Afdeling for Etnografi og Social Antropologi.
- Roepstorff, A. (2002a): *Cognitive Brain Mapping—a Biosemiotic Perspective*. Paper presented at the Workshop Practical Applications and Consequences in Biosemiotics, TheorieLabor, Jena.
- Roepstorff, A. (2002b): Transforming subjects into objectivity. An ethnography of knowledge in a brain imaging laboratory. *FOLK, Journal of the Danish Ethnographic Society*, **44**, 145-170.
- Roepstorff, A. (2003): Mapping Brain Mappers: an ethnographic coda. In R. S. Frackowiak et al. (eds.), *Human Brain Function*, 2. ed. (pp. in press): Elsevier.
- Roepstorff, A. & Bubandt, N. (2003). Introduction: The critique of culture and the plurality of nature. In A. Roepstorff & N. Bubandt & K. Kull (Eds.), *Imagining Nature. Practices of cosmology and identity* (pp. 9-30). Aarhus University Press. Aarhus.
- Roepstorff, A. & Frith, C. D. (in press): What's at the top in the top-down control of action? *Psychological Research*.
- Roepstorff, A. & Lambert, J. D. C. (1994): Factors contributing to the decay of the stimulus-evoked IPSC in hippocampal CA1 neurons. *Journal of Neurophysiology*, **72**(6), 2911-2926.
- Tononi, G. & Edelman, G. (2000): *Consciousness: How matter becomes imagination*. Penguin Books: London.
- Varela, F. (1987): L'individualité: l'autonomie du vivant, *Sur l'individu* (pp. 88-94). Edition du Seuil: Paris.
- Varela, F. J. (1996): Neurophenomenology. A methodological Remedy for the Hard Problem. *Journal of Consciousness Studies*, **3**(4), 330-349.
- Varela, F. J. (2003): The point of view of the researcher: the cognitive sciences. In N. Depraz & F. J. Varela & P. Vermersch (Eds.), *On becoming aware* (pp. 115-128). John Benjamin Publishers: Amsterdam.
- Vautrin, J. & Barker, J. L. (2003): Presynaptic quantal plasticity: Katz's original hypothesis revisited. *Synapse*, **47**(3), 184-199.
- Volianskis, A. & Jensen, M. S. (2003): Transient and sustained types of long-term potentiation in the CA1 area of the rat hippocampus. *J Physiol*, **550**(Pt 2), 459-492.