Commentary on Friedemann Pulvermüller's

The Neuroscience of Language. On Brain Circuits of Words and Serial Order

Giorgio Marchetti

This book offers a paradigmatic view of how it is possible to ground linguistic research on, and constrain linguistic theory by, neuroscientific evidence. The main aim of Pulvermüller's book is to present the putative neurobiological basis of language, that is, how language is organized in the human brain: using the author's own words, "to spell out language in the language of neurons" (Pulvermüller, 2005, p. 1[1]). In his opinion, once a neuronal language is developed which connects linguistic structures and processes to brain structures and processes, "it would be possible to explore the space of possibilities that is restricted on the one side by current neuroscientific knowledge and on the other side by linguistic phenomena" (Pulvermüller, 2005, p. 272). Neuroscientific data could then constrain linguistic theory and vice versa.

One of the fundamental reasons that leads Pulvermüller to set up his program is the belief that "The brain machinery is not just one arbitrary way of implementing the process it realizes, as, for example, any hardware computer configuration can realize almost any computer program or piece of software. The claim is that, instead, the hardware reveals aspects of the program. *Neuronal structure is information*" (Pulvermüller, 2005, p. 9). The structural and functional properties of the cortex must then be taken into account.

Neuroanatomical and neurophysiological evidence indicates that the cerebral cortex is a network of neurons characterized by ordered input and output connections in modality-specific areas, by heavy information mixing through short- and long-distance connections, and by correlation learning. Indeed, the cortex exhibits the following properties:

1. Ordered input and output connections in modality-specific areas: afferent projections (transmitting information from the sensory organs to the cortex) and efferent projections

[1] The pages of the quotations refer to the reprinted version of 2005.

(through which the cortex controls muscle activity) are ordered. They reach, or insert from, primary areas; sensory and motor projections are organized topographically;

2. Complex mappings of information patterns between modalities: intracortical connections permit mixing of afferent and efferent information. Adjacent neurons are heavily connected and form local clusters; primary areas tend to be linked through relay areas; adjacent areas are connected with high probability, whereas areas farther apart have a lower but still good chance to be connected; homotopic areas of the two hemispheres tend to be connected; connections between areas tend to be reciprocal;
3. Correlational learning: synaptic connections between neurons are modified depending on their activity. Neurons that fire together strengthen their mutual connections, whereas neurons that fire independently of each other weaken their connections (Pulvermüller, 2005, p. 21);
4. Laterality of language: language processes are lateralized to the dominant, left hemisphere, although, as Pulvermüller observes, "the nondominant hemisphere contributes to, is sufficient for, and is also necessary for the optimal processing of language" (Pulvermüller, 2005, p. 43).

These structural properties lead to conceive of the cortex as a device that can serve the following functions:

i) The architecture of the cortex is ideal for mixing and merging information immanent to sensory- and action-related activity patterns. Indeed, the cortex, which is supplied with information ordered according to modality, provides multiple indirect links between sensory- and action-related neurons, which are characterized both by great divergence (each neuron reaches thousands of other neurons) and by great convergence (each neuron receives input from multiple other neurons);

ii) The cortex can store complex relationships between input and output patterns. This is suggested by the fact that the mapping between primary areas (that is, the areas from which afferent and efferent fibers originate) is indirect, through relay neurons and areas;

iii) Webs of neurons can form that are distributed over various cortical areas. This is suggested by the fact that frequently co-occurring patterns of activity can be stored by means of strengthening of synaptic links between the participating neurons, and that such synaptic strengthening can take place not only between closely adjacent cortical neurons but also between neurons in distant areas (Pulvermüller, 2005, pp. 21-22).

Pulvermüller deduces that a device characterized by such structural and functional features allows for the formation of functionally coupled but distributed webs of neurons: what he calls *functional webs*. The development of functional webs would be driven by sensory-motor or sensory-sensory activation, and would be determined by the cortical projections indirectly connecting the coactivated neurons in primary areas to each other. A functional web will be a set of neurons:

a) that are strongly connected to each other;
b) that are distributed over a specific set of cortical areas;
c) that work together as a functional unit;
d) whose major parts are functionally dependent on each other so that each of them is necessary for the optimal functioning of the web (Pulvermüller, 2005, p. 24).

Each neuron member of the web would therefore contribute to holding the web together thereby playing an essential role in its functioning. In Pulvermüller's view, "a web of neuronal links strongly connecting all neurons involved in the specific processes triggered by an object in the input may become the cortical representation of this object. In this case, binding of object features would be established by mutual links within a *distributed functional web*" (Pulvermüller, 2005, p. 24). Therefore, a "concept" (for example, the concept of "cat") would be realized and neuronally represented as a distributed web of neurons forming a functional unit.

Evidence of the existence of functional webs comes from the fact that two important predictions entailed by the very idea of functional web are met. As we have seen, the idea of functional webs implies the existence of a large set of strongly linked neurons distributed over distant cortical areas. Two important predictions can then be made about functional webs: (a) if neurons in the functional webs are strongly linked, they should show similar response properties in neurophysiological experiments; (b) if the neurons of the functional web are necessary for the optimal processing of the represented entity, lesion of a significant portion of the network neurons must impair the processing of this entity. Both predictions were examined and observed in macaque monkeys using a memory paradigm in which the animal had to keep in mind the shape or colour of a stimulus and perform a concordant matching response after a delay of several seconds (delayed matching to sample task). Throughout the memory period, in which the animal had to keep in mind the colour or shape of the stimulus, neurons belonging to different cortical lobes fired at an enhanced and specific level (specific in the sense that they did not respond, or responded less, when a stimulus of a different colour was shown). Prediction (a) was then met. Further investigating neurons in frontal and inferior temporal areas that showed similar stimulus- and action-related activity, evidenced that by temporary cooling of the neurons in frontal areas led to a loss of stimulus specificity of the neurons in the temporal areas, and vice-versa. Thus temporary lesion of stimulus-specific neurons in one area led to functional impairment of the neurons in the respective other area, confirming prediction (b) (Fuster, 1997). These results are reminiscent of the studies on aphasias, which showed that the two areas most crucial for language processing in the cortex, the inferior frontal area of Broca and the superior temporal area of Wernicke, appear to be functionally interdependent.

In addition to this kind of empirical consideration about the existence of functional webs, Pulvermüller (2005, p. 23) also proposes two theoretical answers to the question why cognitive processes should be performed by numerous neurons cooperating in functional units rather than by single neurons. The first one is that only neuron ensembles working together in functional units can overcome the problem of the unreliability and noisiness of single neurons. The second argument in favour of functional webs comes from an estimate of the number of neurons necessary for carrying out the cognitive tasks our cortex is usually engaged in. According to an approximate estimation, one needs several 100,000 (about 1 million) engrams or representations to be able to perform one's usual cognitive tasks, including the possibility of speaking a language. If this estimation is correct, one million individual neurons might be sufficient for representing the various percepts, motor programs, words, etc. Therefore, considering that the cortex includes 10^{10} to 10^{11} neurons, the question can be asked why there are 100,000 to 1 million times as many neurons in the cortex as would be necessary? Pulvermüller's answer is that the cortex includes so many neurons because individual engrams are realized as populations of neurons of 10^5 to 10^6 neurons.

As to the *functional dynamics* of functional webs, that is, the activity states they can assume, Pulvermüller offers some tentative descriptions on the basis both of empirical data, theoretical considerations and the results of simulations performed with artificial models of networks of neurons. The main activity states a functional web can assume are the following:

i) Stimulation of a fraction of the functional web can lead to a full activation of the entire population of the neuron members of the web. This process is known as *ignition* (Pulvermüller, 2005, p. 29). If the functional web is considered as a memory representation of an object, the full ignition of a functional web can be considered as the neuronal correlate of the activation of the stored object representation. It has to be noticed that a functional web can be stimulated either by sensory input, which gives rise to the perception of an object present in the environment, or by other cortical neurons outside the functional web itself, which gives rise to the memory of the object. The ignition process likely takes place within a short period of time after stimulation: an educated guess indicates that the ignition occurs within 100-200 msec after sufficient information is present in the input (Pulvermüller, 2005, p. 54). Ignition is reflected in stimulus-evoked cortical potentials.

ii) Ignition is followed by a *reverberatory* state in which the assembly of neurons retains activity, although the level of activity may fall off exponentially with time (Pulvermüller, 2005, p. 30). This reduction of the activity level of functional webs likely occurs because of various reasons: fatigue effects, the refractory periods of neurons, and, not least, the working of a putative feedback regulation mechanism designed to keep the cortical level of activity within certain bounds and prevent the catastrophic overactivation of the network (Pulvermüller, 2005, pp. 78-80)[2]. The memory interval of reverberatory activity can last for tens of seconds, allowing the neuron set to retain its activity. Pulvermüller observes that this is a putative neurobiological basis of short-term memory: "The distributed cortical functional web itself would therefore be the organic side of a long-term (or *passive*) memory trace, and the sustained activity of the same web would realize the short-term (or *active*) memory" (Pulvermüller, 2005, p. 30). According to Pulvermüller, the putative mechanism that allows functional webs to reverberate and retain their activity for some time is that of "reverberatory synfire chains"[3], that is, reverberatory chains of neurons with many re-entrant loops through which activity waves can travel repeatedly (Pulvermüller, 2005, pp156-157). These reverberatory loops produce well-timed spatiotemporal patterns of activity within the functional web, and are assumed to be the basis of fast oscillation activity (as we will see later, dynamics of high-frequency cortical responses distinguish words from meaningless pseudo-words).

The main differences between ignition and reverberation can be so summarized. Ignition is a brief event; it involves all the neurons of the functional web; it does not imply a fixed spatiotemporal order of neuronal activity; it is the result of the overall strong connections within the functional web. On the contrary, reverberation is a continuous process lasting for several seconds or longer; it can be maintained by small neuron subgroups of the functional web; it is

[2] The problem of the overactivation of networks derives from the fact that in a kind of network such as functional webs, connected neurons that frequently fire together increase the strength of their wiring. In a network with many links between neurons, this fact leads to an increase in the connection strength in the entire network. After much learning and strengthening of connections, so many neuronal links may have become so effective that any stimulation of some of its neurons activates the entire network, thus causing catastrophic overactivation. According to Pulvermüller, a possible brain system realizing the feedback regulation mechanism designed to minimize the catastrophic overactivation, could be represented: "by the loop formed by projections from the cortex to the neostriatum (*Putanem* and *Nucleus candatus*), from there to the paleostriatum (or Pallidum), and finally to the thalamus, from where projects run back to the cortex" (Pulvermüller, 2005, p. 80). Inhibition between cortical neuronal assemblies would be provided indirectly by the inhibitory connections in the neostriatum.

[3] A synfire chain consists of subgroups of neurons connected in sequence. In such a kind of neuronal circuit, a given subgroup can be activated only by the synchronous activity of the preceding subgroup. An important feature of a synfire model is that different synfire chains can share the same neurons: a cortical neuron can be part of several synfire chains and it can therefore be a member of different spatiotemporal firing patterns (Pulvermüller, 2005, p. 149).

characterized by a fixed sequence of neuron activation; it is made possible by its strongest internal links envisaged to provide a preferred highway for spreading neuronal activity (Pulvermüller, 2005, p. 169).

iii) In addition to these two activity states – ignition and priming -, there is the possibility that a functional web is in an inactive state, that is, that it stays at its *resting* level.

iv) A fourth activity state, *priming*, is characteristic only of a specific kind of functional web, what Pulvermüller calls "neuronal set". Neuronal sets are functional webs that, through efferent connections, can influence other functional webs; more precisely, it is the excitatory processes of ignition and reverberation in a given ensemble that likely influence other webs connected to it. "Through connections between webs, which are assumed to be much weaker, on average, than their internal connections, one set of neurons can have an activating effect on the other" (Pulvermüller, 2005, pp. 169-170). Because ignition is a substantial excitatory process, it is assumed to prime connected neighbour sets regardless of how strong connections are between the sets; in contrast, the less powerful process of reverberation, to which less neurons contribute at any point in time, is assumed to prime neighbour sets only if the connection between the sets is strong.

How do these neuroscientific principles relate to language? Is there anything new that they can reveal about the representation and processing of words in the brain? Pulvermüller's hypothesis is: a) that functional webs represent meaningful language units: "words are cortically represented and processed by distributed functional webs of neurons" (Pulvermüller, 2005, p. 4); b) and that "there is a cell ensemble or functional web for each and every word" (Pulvermüller, 2005, p. 74).

How is it possible to prove the existence of functional webs relevant for the processing of words? As we have seen, one of the activity states characteristic of functional webs, namely reverberation, is assumed to be determined by reverberatory circuits realized through chains of neurons with many re-entrant loops, which in turn are supposed to be responsible for the production of timed high-frequency rhythms. If words are processed by functional webs, the prediction can then be made that words activate the corresponding functional web, including their relevant reverberatory circuits, thereby eliciting strong high-frequency rhythms; in contrast, phonologically and orthographically regular pseudo-words that are not part of the language would fail to activate a corresponding functional web, and no or very low high-frequency activity should be elicited. Experiments performed using MEG (magnetoencephalography) confirmed the prediction: about one-half second after the onset of spoken one-syllable words, high-frequency brain responses were significantly stronger compared to the same interval following pseudo-words, in particular over the left hemisphere (Pulvermüller, 2005, p. 53).

Moreover, also other kinds of neurophysiological techniques revealed physiological differences between words and pseudo-words: for example, differences in ERP (Event-related potential) of the brain have been found 100-200 ms after onset of visually presented stimuli. The fact that word-pseudoword differences in high-frequency activity tend to occur with longer latencies than the word-pseudoword differences in ERP, may indicate according Pulvermüller that different brain processes are occurring: "Early ERP differences may reflect the initial full activation, ignition of memory traces for words, a putative correlate of word recognition, whereas differences in high frequency responses may reflect continuous reverberatory activity of word-related functional networks, a putative state of active memory" (Pulvermüller, 2005, p. 55; but see also pp. 63-64).

Finally, functional webs would prove to explain empirical data better than alternative approaches, such as those based on distributed, not-discrete networks of neurons or single neurons (Pulvermüller, 2005, p. 64). Indeed, models based on distributed networks of neurons, in which no discrete representations exist and which process all stimulus types alike, could not easily explain the specific changes observed between words and pseudowords. Likewise, if words were represented by single neurons, the corresponding specific brain activity states could not be distinguished with large scale neuroimaging techniques, such as MEG, as on the contrary they actually are.

In summary, physiological studies seem to provide support for the hypothesis that meaningful language units are represented by distributed coupled neuronal assemblies, that is, functional webs. But *where* are the functional webs representing words and language localized? Over which cortical areas are they distributed?

Classical neurological language theories postulated the existence of two independent core language areas located in the *perisylvian region*: a) Broca's area, in the inferior frontal gyrus (Broadmann areas 44 and 45), which was supposed to store the representations of the articulatory movements performed to pronounce a word; b) Wernicke's area, in the superior temporal lobe (Broadmann area 22), which was supposed to store the images of the sound sequences of words. Contrary to classical neurological language theories, the advent of modern imaging research has made apparent that the language centers of Broca and Wernicke are mutually functionally dependent; moreover, it has also shown that although these core language areas are certainly important for language processing, they are not the only cortical areas contributing to and necessary for language processing: other areas as well become active when specific language stimuli are being processed. What is then the cortical topography of the functional webs representing words?

Pulvermüller argues that the correlation learning principle and the cortex's long-range connections between motor and sensory systems seem to imply associations between neurons in the classical cortical core language areas and in areas processing information about the words' referents. These associations would be brought about by word use in the context of objects and actions. "Functional webs could therefore provide the basis for the associations, in the psychological sense, between the name of an animal and the visual image it relates to, or between an action verb and the action it normally express" (Pulvermüller, 2005, p. 56). Pulvermüller names such associations linking phonological information and information about the actions and perceptions referred to by words *word webs*: it is word webs which are responsible for word-meaning processing and representation.

If word webs include areas processing aspects of words' typical referents, it may be argued that the cortical topography of word webs representing words primarily characterized by visual associations differs from the cortical topography of word webs representing words with strong action associations: "If the referent is an object usually perceived through the visual modality, neurons in temporal-occipital areas should be included in the web. If a word refers to actions or objects that are being manipulated frequently, neurons in fronto-central action-related areas are assumed to be wired into the cortical representations" (Pulvermüller, 2005, p. 59). Indeed, the differential cortical activation by action- and visually-related concepts and words, which had been strongly evidenced by neuropsychological studies on patients whose capacity to produce and comprehend nouns and verbs or animal and tool names was differentially affected by disease of the brain, was confirmed by metabolic imaging studies of category-specific processes using PET and fMRI, and by neurophysiological investigations of ERP and high-frequency cortical responses.

The postulate that words with different referential meaning may have functional webs characterized by different topographies implies even more fine-grained predictions. It is well known that the motor cortex is organized somatotopically: adjacent body muscles are represented in neighbouring areas within the motor cortex. Neurons controlling face movements are located in the inferior precentral gyrus, those controlling hand and arm movements are located in its middle part, and those controlling leg movements are located in its dorsomedial portion. "The correlation learning principle therefore suggests differential topographies for cell assemblies organizing leg-, arm-, and face-related words" (Pulvermüller, 2005, p. 62). Neurophysiological studies confirmed such prediction: action verbs referring to different types of actions are processed by functional webs located in different cortical areas. For example, in an EEG study, Pulvermüller compared face- and leg-related action verbs ('walking' vs. 'talking'). Current source density maps revealed differential activation along the motor strip. Words of the 'walking' type evoked stronger ingoing currents at dorsal site, over the cortical leg area, whereas those of the 'talking' type elicited the stronger currents at inferior sites, next to the motor representation of the face and articulators" (Pulvermüller, 2005, p. 62). Moreover, additional experimental data confirmed that information about the word form and the body parts with which the word-related actions are being carried out, are woven into the same word-related cortical networks and are activated near-simultaneously.

Another example of how the proposal that words are cortically represented and processed by distributed functional webs of neurons can help us understand the mechanisms of language, is represented by the way it explains deficits in language processing. Let us consider for example the double dissociation between agrammatism and anomia. Agrammatism and anomia are two different but complementary types of aphasia. Agrammatic patients have difficulty producing words primarily characterized by their grammatical function, such as articles, pronouns, auxiliary verbs, prepositions, and also inflexional affixes; for some patients, abstract nouns and verbs whose meaning is difficult to imagine are also difficult to produce. In a few words, for agrammatic patients low-imageability words are more difficult than comparable high-imageability words. In contrast, anomic aphasics do not have difficulty using the abstract "grammatical" or function words, but cannot find some well-imageable content words from the lexical categories of nouns, adjectives, and verbs.

How can this double dissociation be explained within a functional web framework? Pulvermüller suggests that:

> a highly imageable word would be represented by a perisylvian cell assembly with strong connections to neurons in other brain areas organizing referential aspects of the word, that is, the motor programs and perceptual patterns to which it can refer. (…) In contrast, a function word lacking any concrete associations that cannot be used to refer to concrete objects or actions would be represented by a functional web without strong links to neurons in action- or perception-related areas. (…) A further implication would be that the function words' networks are more strongly lateralized than are the networks representing highly imageable content words (Pulvermüller, 2005, p. 116).

Therefore, grammatical function words and also highly abstract and not imageable content words would be realized in the brain by more focal, strongly lateralized functional webs restricted to the perisylvian cortex, whereas concrete nouns, verbs, and adjectives would be realized by cell assemblies distributed over various areas of both hemispheres. Considering that the likelihood of a word to be affected by a brain lesion depends on the degree to which its corresponding functional web has been damaged, it can be deduced that if a lesion is restricted to the perisylvian region, only the grammatical function words' webs will be involved, thus giving rise to agrammatism; on the contrary, if a lesion affects areas primarily outside the perisylvian region, only the concrete content words' webs will be involved, thereby giving rise to anomia.

So far, we have seen that the notion of functional webs can help understand how single words and meaningful language units are represented and processed in the brain. But can they also help understand something about syntax? According to Pulvermüller, they can: indeed, also "grammar mechanisms in the brain can be thought of in terms of neuronal assemblies whose activity specifically relates to the serial activation of pairs of other neuron ensembles. These assemblies are called *sequence sets*" (Pulvermüller, 2005, p. 2).

Sequence sets are "neuronal sets" (that is, functional webs that, through efferent connections, can influence other functional webs) that "respond specifically to the ordered sequence of activations of other sets" (Pulvermüller, 2005, p. 177). The basic idea he puts forward is that "word webs and the sequence sets connected to them can respond in a specific manner to a sequence of words and fail to do so in response to the same words presented in a different order. The mechanism this is grounded in is the response of the sequence set to sequences it is prepared to detect" (Pulvermüller, 2005, p. 181). Consequently, a network capable of sequence detection exhibits a certain activity pattern when a string in the input is consistent with the network structure: it activates when it verifies that the sequence of the neuronal elements it represents are present in the input. In linguistic terms, this may be similar to the judgment that the string in the input is grammatical or well formed.

Sequence sets have a shortcoming. To make it possible for a network to decide if a certain sequence occurs, say *AB*, it is obvious to assume neuronal elements that specialize in the sequence detection of that specific sequence (*AB*). One may therefore postulate specific sequence sets for each sequence of words or morphemes. However, if each possible word sequence were to be represented by a specific sequence set, the number of sequence sets would be astronomically high. An obvious solution to overcome this problem is to categorize words into lexical categories and base serial-order algorithms on these grammatically defined word categories. The number of the necessary sequence detectors would thus be greatly reduced, because sequence sets could connect to a few category representations instead of an extremely large number of pairs of input units.

But how can the process of categorizing words into lexical categories be achieved? In Pulvermüller's view, "assigning a word *A* to a lexical category *a* is closely tied to specifying which kinds of words B_1, B_2, B_3, …. are required to occur – and therefore usually occur – together with word *A*" (Pulvermüller, 2005, p. 188). In other words, the problem of lexical categorization can be solved in terms of sequence regularities: the category of a nominative noun, for example, would be represented by two sequence detectors, one detecting that the element in question followed an article and the other examining whether it is followed by a verb; the representation of the category of a transitive particle verb may include sets sensitive to the word followed by a verb suffix, an accusative noun, and a particle; an so on. As one can see, some of the postulated sequence sets connected to a particular word representation would be sensitive to particular inputs preceding the word, whereas others would be sensitive to the elements following the word. A lexical category representation can then be defined: “as the union of several sequence detectors, a set of sequence sets” (Pulvermüller, 2005, p. 190).

The issue of lexical categorization is further complicated by the fact that a word can be a member of different lexical categories: the word “beat”, for example, can be used either as a noun or as a verb. This requires a mechanism deciding between the neuronal sets representing alternative lexical categories to which a given word can be assigned. The mechanism can be realized by implementing, for example, “winner-takes-all” dynamics allowing for most active lexical category representation to become fully active and the competitor category representations to be suppressed. Mutual inhibition would thus occur between the alternative, competing lexical category representations (Pulvermüller, 2005, pp. 194-196).

In addition to the solution to the problem of how lexical categorization can be realized in a neuronal network, Pulvermüller also offers plausible and elegant neurobiological solutions to other important grammatical and syntactical problems, such as the neural processing of: a) center-embedded strings (Pulvermüller, 2005, pp. 245-247); b) multiple occurrence of the same word form or lexical category in a sentence (Pulvermüller, 2005, pp. 238-245); c) long-distance dependency; d) distributed words; e) subject-verb agreement; f) the distinction between a constituent’s obligatory complements and its optional adjuncts (Pulvermüller, 2005, pp. 236-238); g) homophones (Pulvermüller, 2005, pp. 83-87); h) synonyms (Pulvermüller, 2005, pp. 87-88).

Despite all the promising perspectives for the studies on language opened up by the neuroscientific approach, some very important linguistic issues remain unexplained or only partly explained, as Pulvermüller himself admits. Without doubt, one of the most relevant ones concerns the processing of the meanings of words, that is, semantics. With reference to this issue, Pulvermüller highlights that there is still no common consensus among neuroscientists about which cortical areas (or neural processes) are involved in the processing of, and phenomenally constituting, the meanings of words. He states that: “The question concerning the cortical locus of the processing of word meaning has been addressed in many imaging studies, and a number of conclusions have been proposed” (Pulvermüller, 2005, p. 46). For example, on the basis of both lesion and positron emission tomography (PET) data, Tranel and Damasio (1999) emphasize the role of the left inferior and middle temporal gyri, while based on PET and electroencephalographic (EEG) data, Posner and Di Girolamo (1999) argue for semantic processes in left inferior frontal areas. At the same time, based on magnetoencephalographic (MEG) evidence, Salmelin, Helenius, and Kuukka (1999) claim that the left superior temporal lobe is relevant for word semantic; and

Skrandies (1999) reports EEG studies highlighting the importance of the occipital lobes in distinguishing word meanings, while Pulvermüller (Pulvermüller, 1999, 2005), as we have seen, stresses the importance of primary motor, premotor and prefrontal areas for word-meaning processing.

One of the possible reasons for this shortcoming is certainly of methodological nature. For example, the difficulty of the experiment and the task context in which words must be processed can certainly account for the differences in brain activation between tasks: it is one thing to focus one's attention on features of the stimulus, quite another to engage in language-related tasks, such as keeping items in memory, comparing items with each other, rejecting items, naming, lexical-decision, and so on (Pulvermüller, 2005, pp. 46, 53 and 65). Moreover, the modality of the stimulation plays an additional role in determining the set of active brain states: it is one thing to listen to a word, quite another to produce one; it is one thing to listen to a word, quite another to read one; it is one thing to listen to a frequently recurring or used word, quite another to listen to a rare word.

However, in my opinion, the shortcoming of cognitive neuroscience concerning semantics is mostly due to a lack of an adequate theoretical framework. Let us see briefly why an adequate theoretical framework should help overcome this shortcoming of cognitive neuroscience.

In general, cognitive neuroscience is interested in discovering how the brain makes possible and realizes the cognitive processes that allow us to think, speak languages, remember, pay attention to something, be conscious of something, etc.: in a word, to have a mental life. In order to achieve this aim, cognitive neuroscience usually tries to find out where in the brain, that is, in which brain areas, a particular cognitive process is located; what kind of neuron circuit implements the process; at which point in time the process takes place compared to other processes. Cognitive neuroscience tries then to give an account in *physical* terms of the cognitive processes constituting our *mental* life.

It is quite easy to realize that such an account cannot be given without having formulated a theoretical model that not only describes such cognitive processes - that is, which ones and how many they are - but also explains how they produce phenomena such as thought, language, consciousness, and meanings, what these processes consist of, and, more in general, how our mind works. Indeed, in order to identify the brain structures and mechanisms responsible for the production of these phenomena, one needs a criterion by means of which one understands where, what, how and when to look for and observe. The physical field can be subdivided in so many different levels (the micro-level of the atoms and sub-atomic particles; the medium-level of cells and neurons; the macro-level of the assemblies and systems of cells and neurons, etc.), and observed from so many different angles (as an isolated unit, as a composite system, as a dynamic structure, etc.) that without a criterion or guide for deciding where, what, how and when to observe, one cannot even start one's research: Where should one address it? What should be the more proper level of observation of the physical phenomena: the level of the atoms composing neurons, the level of neurons, or the level of the assemblies of neurons? How, and on what basis, could one explain the relationships between the various elements composing each physical level? What criterion should one adopt to analyze the relationships between the different physical levels?

Therefore, in order to be able to give an answer to these questions, one needs a theoretical framework that pilots the neuroscientific research, showing it what, where, when and how to look for in the brain: that is, a theoretical framework that provides cognitive neuroscience with the right level of description and explanation necessary to carry out its investigations. When the meanings of words are concerned, such a theoretical framework should specify their mental origin, how meanings take form in our mind, what mental mechanisms and operations produce meanings, how, given certain constraints (that is, what is for instance physically, psychologically and neurophysiological known), these mechanism produce meanings, how these very mechanism make us consciously experience meanings, what the meanings are for each and every word, and so on.

In this sense, a promising theoretical proposal is represented by those semantics such as Operational Semantics (Benedetti, 2004, 2006), Operative Linguistics (Ceccato, 1969, Marchetti, 1993), and Attentional Semantics (Marchetti, 2003, 2005, 2006), which, taking into proper account both our direct, conscious experience of meanings and what is known through the other sciences such as psychology and neurophysiology, aim at systematically describing the mental origin and formation of meanings. Such kinds of semantics, giving an *operational* account of the meanings of words, that is, identifying the elemental mental operations that make up the meaning of each and every word, and analyzing and representing our (conscious and unconscious) mental life as a function or set of functions performed by the working of some physical organs (the brain as a whole, or its parts), make it possible to assign every single mental property or function to some physical organ. In this way, they open the road to the systematic and detailed research of the physical bases of mental life: by subsequent and finer and finer manipulations of the physical substratum, it is then possible to empirically determine and isolate the organ that is responsible for the production of a specific mental property.

The natural and empirical counterpart to such kinds of semantics is represented by those neurophysiological approaches such as Operational Architectonics (Fingelkurts & Fingelkurts, 2001, 2005, 2006) that are specifically based on the notion of operation. According to Operational Architectonics, conscious phenomena are originated by the joint operations of functional transient neuronal assemblies, or Operational Modules (OMs). OMs, which are realized by the temporal synchronization of different brain operations executed by different local neuronal assemblies simultaneously (operational synchrony), can in turn be operationally synchronized between each other (on a new time scale), thus forming more abstract and more complex OM which constitute new and more integrated phenomenal experience.

As one can see, theoretical labour plays an important role in neuroscientific research, both for its capacity to direct and pilot research and for its explanatory and predictive power. A purely physical, neurocientific approach, centred uniquely on *brain*, that is not supported by an appropriate theoretical, operational model of *mind* is unavoidably doomed to continue to give unsatisfactory answers to the problem of how the brain realizes the cognitive processes that allow us to have a mental life.

Pulvermüller seems to be generally well aware of the importance of *theory* for neuroscientific research:

> What is necessary, then, are ideas about how to connect the level of language description to that of the description of neuron. Piling up more neurophysiological and imaging data may not help much in this enterprise. Empirical facts do not by themselves form a theory about the generation of sunlight or language. Theoretical work is required in the first place. The theoretical efforts can lead to the generation of predictions that can be addressed in crucial experiments. Lack of empirical data is never a very good excuse for postponing the necessary theoretical labour (Pulvermüller, 2002, p. 271);

but not to be fully aware of the importance of a *theory of mind*, given the primacy he recognizes of a *theory of brain*, as the following passage seems to attest:

> The instruments for monitoring brain activity do not by themselves tell the researcher what to look for when investigating linguistic representations and processes. There are infinite possibilities for describing and analyzing a shot time series obtained, for instance, using a multichannel electro- or magnetoencephalograph. What should the language and brain scientist begin with when searching for the pattern that, in the clear case, correlates with the occurrence of a wh- sentence? Answers to this question can be provided only by a theory about the brain mechanisms of language (Pulvermüller, 2002, p. 274);

Considering not only the necessity, which I mentioned earlier, for any empirical research to be performed within a precise theoretical framework, but also the fact that Pulvermüller himself resorts several times to theoretical models of language development, processing production and comprehension, such as for example Freud's or Lichtheim's, in order whether to adopt, criticize or simply refer to, them, this lack of awareness of the role played by a theory of mind seems quite unmotivated.

Moreover, when tackling the specific problem of the empirical investigation of word meaning processing, Pulvermüller seems to acknowledge and highlight the importance of theoretical labour not so much for its capacity to direct research, as for its *ex post* explanatory power:

In spite of the undeniable progress made possible by empirical results obtained with newly introduced techniques, it is likely that theoretical advances are necessary as well. This may help to explain *why* the processing of meaning – and other language and cognitive processes – should activate certain cortical areas but not others, and may, in the best case, help us understand the diversity of results reported in some areas of the cognitive neuroscience of language (Pulvermüller, 2002, p. 47).

Most probably, it is the partial knowledge of the advances of semantic studies that makes him underestimate the power they can have in directing and piloting empirical research, a partial knowledge that leads him to wrongly state that function words, such as "the" and the plural suffix "-s", add no semantic information to the sentence in which they are included, but play only a syntactic role: "Their inclusion in a sentence usually does not add semantic information to it, but may make the word string either acceptable or unacceptable. The neuronal representations of function words and affixes could therefore exist without a semantic part" (Pulvermüller, 2005, p. 117)[4]. Contrary to what Pulvermüller states, it has been extensively shown (Ceccato and Zonta, 1980, Glasersfeld, 1989, Marchetti, 1993) that function words and affixes do have a semantic role. As anyone can see comparing sentences (1), (2) and (3), the presence or the absence of an article actually makes the difference in a sentence, changing definitely its meaning:

(1) The man appeared
(2) A man appeared
(3) Man appeared

As to the plural suffix "-s", its use not only modifies the meaning of the sentence in which it is included, but also adds what Pulvermüller calls "referential meaning", in the sense that it influences and pilots the way in which we perceive and conceive things. Consider for example the following picture:

[4] Even though later on he seems to partly modify and adjust his opinion on function words and affixes: "At least, they can be conceptualized without a part including information about referential meaning that would be stored primarily outside the perisylvian region" (Pulvermüller, 2005, p. 117), thus very seemingly implicitly admitting that they can have some other kind of meaning that differs from a purely referential one. See also the following statement: "The assumption that no semantic information is being added to a sentence by the inclusion of function items does not generally hold true for all members of this category. The regular past suffix '-ed' and the auxiliary verb for 'was', for example, include information about time, and it is clear that a cortical representation of these lexical items must include this information as well and bind it to the phonological information about the respective form" (Pulvermüller, 2005, p. 117).

You can describe the picture in various ways: among them, there are, from the point of view of the grammatical number, at least two opposite alternatives, that is, “trees” and “thicket”, which – as anyone can realize - make you consciously experience the content of the picture in two completely different ways.

The importance of a theoretical model for neuroscientific research becomes evident also when dealing with more general issues concerning the conscious experiences elicited by the meanings of words. For example, theoretical considerations suggested by Attentional Semantics (Marchetti, 2003) suggest that, when considering word-meaning understanding during comprehension, at least two different and distinct phases must be distinguished. The first phase consists in having, so to say, a “pure” conscious experience of the meaning of the word. In this phase, independently of whether the word we hear is a function one or a content one, we do not consciously experience any “qualitative”, perceptible imagine: we simply understand the meaning of the word, without consciously representing it by means of images or other more concrete sensory modalities. The second phase, on the contrary, consists in having perceptive conscious experiences of the meaning of the word. This second phase takes place usually with content words, and only if we have enough time at our disposal to physically imagine, think extensively, or recall the events that are related to their meanings (nevertheless, sometimes it is possible for us to experience this second phase also with function words, as, for example, when we try to imagine what a word such as “but” or “or” reminds us of). Therefore, when empirically investigating word-meaning processing, one should carefully distinguish the two phases. As we have seen, Pulvermüller (Pulvermüller, 2005, p. 118) proposes that function or grammatical words are processed by strongly lateralized neuron ensembles restricted to the perisylvian cortex, whereas content words have less lateralized corresponding neuron ensembles distributed over various areas of both hemispheres. In view of the theoretical considerations suggested by Attentional Semantics, Pulvermüller’s hypothesis should be slightly refined. Indeed, while it seems highly plausible that, during the second phase of word-meaning processing, the extra-perisylvian nervous structures representing motor programs and perceptual patterns are involved, above all if content words are concerned, the absence of any qualitative, perceptive conscious experience during the first phase of word-meaning processing makes Pulvermüller's hypothesis quite problematic and in need of review.

On the whole, I think that Pulvermüller's work is highly valuable for his continuous and systematic attempt at constraining, empirically ground, and verifying linguistic data and theories. His effort to bridge the gap between brain and language is carried out in a serious and methodical way: however, in my opinion, it can be fully and successfully attained only when it is grounded on a well developed and viable theory of mind.

References

Benedetti, G. (2004). *Semantica operativa*. URL: www.mind-consciousness-language.com

Benedetti, G. (2006). Operational Noology as a new methodology for the study of thought and language: theoretical aspects and possible practical applications". *Cognitive Processing, 7*, 217-243.

Ceccato, S. (1969). (Ed.). *Corso di linguistica operativa*. Longanesi. Milano.

Ceccato, S., & Zonta, B. (1980). *Linguaggio consapevolezza pensiero*. Feltrinelli. Milano.

Fingelkurts, An. A., & Fingelkurts, Al. A. (2001). Operational architectonics of the human brain biopotential field: Towards solving the mind-brain problem. *Brain and Mind, 2*, 261-296.

Fingelkurts, An. A., & Fingelkurts, Al. A. (2005). Mapping of Brain Operational Architectonics. In F. J. Chen (Ed.), *Focus on Brain Mapping Research* (pp. 59-98). Nova Science Publishers, Inc.

Fingelkurts, An. A., & Fingelkurts, Al. A. (2006). Timing in cognition and EEG brain dynamics: discreteness versus continuity. *Cognitive Processing,* 7, 135-162.

Fuster, J. M. (1997). "Network memory". *Trends in Neurosciences, 20*: 451-459.

Glasersfeld, von E. (1989). *Linguaggio e comunicazione nel costruttivismo radicale*. Clup. Milano.

Marchetti, G. (1993). *The Mechanics of the Mind*. Espansione. Roma.

Marchetti, G. (2003). Foundations of Attentional Semantics. URL: www.mind-consciousness-language.com

Marchetti, G. (2005). The importance of non-attentional operations for Attentional Semantics. URL: www.mind-consciousness-language.com

Marchetti, G. (2006). A presentation of Attentional Semantics. *Cognitive Processing*, 7, 163-194.

Posner, M. I., & Di Girolamo, G. J. (1999). Flexible neural circuitry in word processing. *Behavioral and Brain Sciences, 22*, 299-300.

Pulvermüller, F. (1999). Words in the brain's language. *Behavioral and Brain Sciences, 22*, 253–336.

Pulvermüller, F. (2002). *The Neuroscience of Language. On Brain Circuits of Words and Serial Order*. Cambridge University Press. Cambridge.

Salmelin, R., Helenius, P., & Kuukka, K. (1999). Only time can tell – words in context. *Behavioral and Brain Sciences, 22*, 300.

Skrandies, W. (1999). Early effects of semantic meaning on electrical brain activity. *Behavioral and Brain Sciences, 22*, 301.

Tranel, D., & Damasio, A. R. (1999). The neurobiology of knowledge retrieval. *Behavioral and Brain Sciences, 22*, 303.