

Language as a Tool. An Insight From Cognitive Science

Bartosz Brożek

Jagiellonian University; Copernicus Center for Interdisciplinary Studies, Krakow, Poland

e-mail: bbroze@yahoo.com

Mateusz Hohol

Pontifical University of John Paul II; Copernicus Center for Interdisciplinary Studies,
Krakow, Poland

e-mail: mateuszhohol@gmail.com

Abstract:

In this paper it has been argued that the theory of conceptual maps developed recently by Paul M. Churchland provides support for Wittgenstein's claim that language is a tool for acting in the world. The role of language is to coordinate and shape the conceptual maps of the members of the given language community, reducing the cross-individual cognitive idiosyncrasies and paving the way for joint cognitive enterprises. Moreover, Churchland's theory also explains our tendency to speak of language as consisting of concepts which correspond to things we encounter in the world. The puzzle of common sense reference is no longer a puzzle: while at the fundamental level language remains a tool for orchestrating conceptual maps, the fact that the maps encode some communally shared categorization of experience fuels our talk of concepts capturing the essences of things, natural kinds, prototypes, etc.

Keywords: cognitive science, conceptual maps, language processing, meaning, philosophy of language.

1. Prelude

In *Philosophical Investigations* Ludwig Wittgenstein famously argues that language is not a “mirror of reality”; rather, it is a tool for acting in the world. What this claim amounts to, is a revolutionary change in the understanding of the nature of linguistic phenomena. Wittgenstein underlines this point by contrasting his novel conception with the Augustinian picture of language, in which

the individual words in language name objects—sentences are combinations of such names.—In this picture of language we find the roots of the following idea: Every word has a meaning. This meaning is correlated with the word. It is the object for which the word stands [18, §1].

What does it mean to reject this view? It is commonly believed that in the *Investigations* Wittgenstein posits that the meaning of an expression is determined by the ways in which it can be used (let us mark this fact by referring to the Wittgenstein's understanding of meaning with an asterisk "*" and use no asterisk with the "received" understanding of the term). Arguably, it is whole sentences, rather than individual words, that have meaning*, since it is sentences that are tools for acting in the world (even if we sometimes use single words – for example, when shouting "Brick!" – what we intend to say by this is something like "Watch out, because there is a brick falling down on you from the roof of the building by which you are standing!"). If we say that the word "brick" has a meaning, then it is only derivative in character: the meaning of a word can be abstracted from the meanings* of the utterances in which the word occurs. Let us consider another example. The sentence "Jane likes mangoes" means* something, but *not* because the words "Jane," "like," and "mangoes" have particular meanings. It is sentences – or their functional equivalents – that serve as the means for acting in the world and it is sentences that have meaning*, with the meanings of the component words being just *ex post* theoretical reconstructions. We can "construe" the meanings of "Jane," "like," and "mangoes" because we know in what sentences these words can occur and in what contexts those sentences can be used. This constitutes a break in the long philosophical tradition which posited the primacy of a concept over a sentence: in Plato, Aristotle, Descartes, Kant, and even in Frege, Russell and the early Wittgenstein (though with certain reservations in the latter cases), it is the meaning of individual concepts that is primary, while the meaning of a sentence is composed of the meanings of the words contained in it. Meaning* is not meaning: the latter is primarily predicated of concepts, while the former – of sentences. The latter is connected to the idea that words in language somehow reflect the nature of the things in the world they name, why the former sees linguistic expressions as ways to act in that world.

However, this dramatic shift in the understanding of the nature of language leads to some serious theoretical difficulties. Let us consider two of them. Firstly, it seems plainly clear that at least some of the words in our languages express concepts which correspond to things in the world. The word "table" expresses a concept which refers to tables, the word "apple" expresses a concept which refers to apples, and the word "water" expresses a concept that refers to water. Of course, the exact nature of such concepts is a focus of fierce philosophical debates. For example, the classical theory of concepts holds that "a lexical concept C has definitional structure in that it is composed of simpler concepts that express necessary and sufficient conditions for falling under C" [12]. According to the prototype theory, on the other hand, "a lexical concept C doesn't have definitional structure but has probabilistic structure in that something falls under C just in case it satisfies a sufficient number of properties encoded by C's constituents" [12]. The point is, however, that on both competing accounts linguistic concepts correspond to some things, which seems incompatible with the view suggested by Wittgenstein that words are mere tools for acting in the world.

Secondly, it is also undeniable that the theories we construct and express in language often capture the objective features of reality. This is particularly visible in physics, where mathematical equations not only enable to precisely predict a course of physical events, but often "know more" than their creators. An instructive example is given by Wigner. When Heisenberg formulated his quantum mechanics based on matrix calculus, the theory was applicable only to a few idealized problems. Applied to the first real problem, of the hydrogen atom, it also proved successful:

This was (...) still understandable because Heisenberg's rules of calculation were abstracted from problems which included the old theory of the hydrogen atom. The miracle occurred only when matrix mechanics, or a mathematically equivalent

theory, was applied to problems for which Heisenberg's calculating rules were meaningless. Heisenberg's rules presupposed that the classical equations of motion had solutions with certain periodicity properties; and the equations of motion of the two electrons of the helium atom, or of the even greater number of electrons of heavier atoms, simply do not have these properties, so that Heisenberg's rules cannot be applied to these cases. Nevertheless, the calculation of the lowest energy level of helium (...) agrees with the experimental data within the accuracy of the observations, which is one part in ten million. Surely in this case we 'got something out' of the equations that we did not put in [17, p. 10].

This, according to Wigner, shows that “the mathematical language has more to commend it than being the only language which we can speak; it shows that it is, in a very real sense, the correct language” [17, p. 7]. But if so, it is difficult to conclude that the language of mathematics is just a tool for acting in the world (*pace* mathematical instrumentalism). Mathematical equations of the contemporary physics seem to correspond – in one way or another – to the structural aspects of reality.

Therefore, the question is, how to reconcile the persuasive picture of the nature of language as a tool provided by Wittgenstein with two different facts: that we are forced to claim that at least some of our linguistic concepts have referents in the world we encounter in our everyday experience; and that the theories developed in physics capture such aspects of the structure of the universe which lie beyond that experience. Let us deem those problems, respectively, the puzzle of the common sense reference and the puzzle of the scientific reference. In what follows, we would like to show how the first of those puzzles may be solved – or, at least, illuminated – by adopting the view of human cognition proposed by Paul M. Churchland. Further, we shall argue that although the second puzzle remains a mystery, its significance transcends the framework of the Wittgensteinian view of language as a tool.

2. The received views of cognition

In order to grasp the significance and the novelty of Paul M. Churchland's theory of cognition, it is reasonable to begin with an outline of the two traditional views of human cognitive abilities – of thinking as seeing and of thinking as conversation. The model of thinking as seeing is best encapsulated in the conceptions championed by Locke, Hume, Berkeley, and other early modern philosophers. It is a view that all we know comes from experience: when one perceives something, a mental picture (referred to as an idea or an impression) is produced, which is stored in one's memory, from which they can be retrieved even if the object which caused the experience is no longer present. Those pictures may also be processed in a number of ways. For instance, Locke believed that simple ideas may be combined into compounds, or brought together so as to establish a relation between them, or abstracted from to produce a general idea [11]. On this view, understanding something boils down to “seeing clearly” the ideas representing it, their compounds, relations between them, or some features they share. It further presupposes a special relationship between ideas and things in the external world: the former *resemble* the latter, just as pictures resemble the pieces of reality they portray. Still, Locke, Hume and Berkeley play down the role of language, which – on their account – becomes a system of purely conventional symbols used to designate ideas. That a table is called ‘a table’ in English, ‘der Tisch’ in German, or ‘stół’ in Polish, is a result of a convention; but the important thing is that all three words correspond to the same idea. Our understanding of reality is extra-linguistic: we know what tables are, what they look like, and what their functions are since we have the relevant ideas, and not because we use certain expressions.

This view of thinking faces two kinds of objections. On the one hand, and contrary to much of the current psychological and neurobiological theories, it encapsulates a particularly passive view

of cognition, in which judgment and action are logically depend on prior perception. For example, Descartes believes that “there are only two modes of thinking in us, viz., the perception of the understanding and the action of the will. (...) To perceive by the senses (*sentire*), to imagine, and to conceive things purely intelligible, are only different modes of perceiving (*percipiendi*); but to desire, to be averse from, to affirm, to deny, to doubt, are different modes of willing. (...) The will as well as the understanding is required for judging. I admit that the understanding is necessary for judging, there being no room to suppose that we can judge that which we in no way apprehend; but the will also is required in order for us to assent to what we have in any degree perceived” [5, XXXII, XXXIV].

On the other hand, the model of thinking as seeing totally disregards the role of language in cognition. The observational data gathered, and the theories developed since the times of Benjamin Lee Whorf, strongly suggest that the language one speaks influences one’s perception and thinking; and even if the strong reading of linguistic relativity may not be supported by facts, there is no escape from the conclusion that the language and culture into which we are born decisively shape the way we perceive and the way we act. The failure to *see* this is one of the grave errors of those who believe that thinking resembles seeing. This problem is best illustrated in relation to the question of how abstract concepts are generated. Let us recall that according to Locke one produces a general idea by abstracting from concrete ideas; moreover, Locke claimed that the mechanism of abstraction is an inborn ability, something we have at our disposal from the very beginning of our lives. But what does “seeing” a general idea consist in? For instance, what would seeing a general idea of a triangle be like? Such a general idea of a triangle can neither represent an acute, a right, nor an obtuse triangle. Immanuel Kant was quick to realize that fact – he claimed that all we can ‘see’ in intuition are particulars: if one ‘sees’ a triangle, it is a particular acute- or right- or obtuse-angled triangle; if one ‘sees’ a tree, it is a particular tree, with a particular height, shape, number of branches, type of leafs, etc. Of course, says Kant, knowledge does not consist of ‘pictures’ of particular instances of triangles or trees; knowledge is universal and conceptual. Hence, Kant claimed that the concepts we have are general and refer to classes of things. The concept of a triangle embraces all triangles, be they acute, right or obtuse; the concept of a tree refers to all trees; etc. Naturally, a concept (of a triangle or of a tree) is not something one sees; rather, every concept is associated with the so-called transcendental scheme, i.e. a set of instructions one has to follow in order to construct in the intuition a particular instantiation of a given concept. Thus, the concept of a triangle is associated with a transcendental scheme which tells us how to ‘draw’ an imaginary triangle in the intuition. The outcome of the procedure is always a particular triangle; but the procedure itself, the set of instructions which constitute the transcendental scheme, encapsulate the general properties of *any* triangle. Kant might have failed to realize the extent of the influence language has on perception and action, but he clearly identified the shortcomings of the “mental eye” view of cognition [10], [3].

Unfortunately, the alternative to the model of thinking as seeing – the view that thinking is like a conversation of the soul with itself – is equally lacking. This is a theory which, with some reservations, and in different forms, may be ascribed to numerous philosophers and psychologists such as the late 19th century German psychologist Max Müller, who emphasized that true thought cannot be carried on without words [13]; Lev Vygotsky, who claimed that the structure of language determines the deep structure of thought [16]; or Jerry Fodor who believes in the existence of a universal language of thought or *Mentalese* [8].

The claim that thinking cannot take place outside language – be it talking to oneself in the language our mothers teach us or the use of the mysterious *Mentalese* – is easily dismissed when we consider that according to such a view thinking would be an exclusively human affair, and hence constitute an evolutionary enigma. Moreover, when we limit thinking to the use of words, much of our conscious and meaningful experience would need not only a different label, but also a different explanation. Imagine yourself ‘playing in your mind’ Bach’s *Kunst der Fuge* or remembering someone’s face; it surely is non-verbal, but it is deserving of the label “thinking”. It may be argued

that the problem is purely definitional and is based on a misunderstanding. The proponents of thinking as conversation are simply not interested in non-verbal mental experiences – they do not deny their existence, but claim that they have different foundations than thinking in words and as such do not deserve the name. But in such a case, one is entitled to ask for an explanation what is the relationship between thinking (which requires language) and other forms of mental life, as well as what are the evolutionary reasons for such a dualism of thought.

3. Conceptual maps

Therefore, it seems that we need an alternative to the two received models of thinking. Such an alternative has recently been sketched by a number of scholars, most notably Peter Gärdenfors and Paul Churchland. In a nutshell, their views – even if different in motivations, scope, empirical basis, and other details – boil down to the thesis that our mental representations and, in consequence, the process of thinking which takes advantage of them, are neither (structured) collections of ‘mental pictures’, nor (structured) sets of propositions. For Gärdenfors, the best description of human individual knowledge is provided by what he calls *conceptual spaces* [6], [7], while Churchland prefers to speak of *conceptual maps* [4]. Let us consider in some detail the latter proposal.

When one’s brain receives some input from the environment, it is through the sensory neurons (e.g., those located in rods and cones in the retina, mechanoreceptors within the skin, etc.). However, those populations of sensory neurons do not produce the final representation of the perceived phenomena; rather, they only serve as the first rung in a complex, hierarchical neural structure, where information retained at one level is transformed – through synaptic connections – to a higher level. “Each rung of each of these ladders constitutes a unique cognitive canvas or representational space, a canvas or space with its own structured family of categories, its own set of similarity and difference relations, and its own peculiar take on some enduring aspects of the external world. What happens, as sensory information ascends such a ladder, is its progressive transformation into a succession of distinct representational formats, formats that embody that brain’s background ‘expectations’ concerning the possible ways in which the world can be” [4, p. 35]. In this way, face recognition, the categorization of colours, and any other cognitive activity is enabled by the existence of an appropriate neuronal ‘ladder’; it follows, *inter alia*, that the representation of some concept – say, of ‘red’ – is not an activation pattern at the level of sensory neurons, but rather ‘redness’ is a part of the conceptual system encapsulated in a hierarchy of neural networks. In such a hierarchy, the information from the lower-level network is transformed through the synaptic connections to the higher-level, and hence the ability to recognize colour red, and to distinguish it from other colours, is determined by the configuration of the connection weights across the transforming population of the synapses. Such a cognitive architecture cannot be properly captured by any system of logic; it is rather a vector space, and thus our representation of the world is geometrical. It is why Paul Churchland is justified in speaking of conceptual maps in our brains, and Peter Gärdenfors – of conceptual spaces.

Crucially, such vector spaces provide a framework for representing both perception and action. “The brain’s representations of the world’s enduring categorical and causal structure (its ‘factual’ knowledge), and the brain’s representations of its various acquired motor skills and abilities (its ‘practical’ knowledge), are both embodied in the carefully sculpted metrics of similarities and differences that provide lasting structure to each one of the brain’s many activation spaces” [4, p. 49]. Thus – contrary to many philosophical theories and the “received wisdom” of folk psychology – no unbridgeable gap exists between the representations of perception and action.

Another important aspect of the geometrical view of the mind is that the neural activation patterns are a result of the long process of individual learning. We are not born with the ability to discern between red and orange, or distinguish female from male faces. But our brains are so constructed as to enable the gradual formation of neural activation patterns, and this process leads firmly to the establishment of hierarchies of neural networks which make it possible for us to

categorize and conceptualize our experience and act in the world. If so, our mental representations are clearly embodied and enacted. They emerge in the interactions of our bodies with the environment, and they are shaped by what we do, and not – at least not only – by what we passively experience.

There is one further aspect of the above described cognitive mechanism that should be stressed here. The mechanism “displays the capacity to reach beyond the shifting vagaries of one’s sensory inputs so as to get a grip on the objective and enduring features of one’s perceptual environment” [4, p. 67]. The training of a given neural network, the painstaking orchestration of the numerous synaptic weights, takes place over a long period of time, fine-tuning the entire system to the environment, and hence capturing its unchanging, invariant features. Naturally, not all such features are represented in the human brain – only those which happen to be relevant for our actions. Moreover, the cognitive apparatus of other animals, from quite simple organisms to non-human primates, also capture objective aspects of reality. For example, bats must have representations of the environment which are quite distinct from human conceptual maps, even in relation to the same aspect of the world. “A blind bat, (...) may know what flamingos are, and be perfectly able to discriminate flamingos from other flying creatures, but it will do so by focusing (actively) on the distinctive acoustic signature of the sonar echo returned from such a bird, or (passively) on the equally distinctive acoustic signature (‘swoosh, swoosh’) given off when any flamingo takes flight” [4, p. 89]. But conceptual maps differ not only across but also within species. No two people have exactly the same neural representations of the given phenomenon: the relevant synaptic connections and their weights may differ to a greater or lesser degree (even though the fact that humans have similar genetic underpinning, as well as that they deal with the same type of environment, significantly reduces such differences).

Churchland claims further that in addition to the slow process of learning, which consists of the formation of the hierarchies of neural networks as described above, there exists another, faster, dynamic process which may also be called “learning”. It is conceptual redeployment and it takes place when – without altering any synaptic connections or adjusting the network’s weights – the sensory input receives a significantly novel conceptual interpretation. An example of such an occurrence may be the realization that dolphins – with all their fish-like features – are best categorized as mammals; or Darwin’s observation that the concepts of natural variation and selective reproduction are applicable to all species throughout the history of life on Earth and that they are capable of explaining the diversity of those species; or Newton’s insight that the Moon’s elliptical trajectory is only an instance of a flung stone’s trajectory [4, p. 188-192].

3. The nature of language

But the two kinds of learning Churchland identifies – the long process of the formation of conceptual maps, and a relatively quick one of conceptual redeployment – do not exhaust the list of mechanisms operating at the human cognitive *niveau*. What is typically human, and what constitutes one of the greatest discoveries in the evolutionary history of our species, is the ability to use language. When a child learns to speak, the effect on her conceptual maps is that they are formed in a way which is typical for the society in which she was born. Of course, some level of idiosyncrasy remains, but it is by far smaller than it would be were there no language, since “most obviously (...), a shared language and a shared theoretical vocabulary allows us to coordinate our individual conceptual maps of a given domain, both in our background understanding of that domain’s abstract structure and in our local sensory and instrumental indexings of that map in the business of interrogating the objective world” [4, p. 269]. Apart from amplifying, regulating, and coordinating conceptual maps, there is one more effect that language has on human cognition – it becomes collective:

It then involves a number of distinct brains – at least a handful, and perhaps many millions – engaged in a common endeavor. Their consensual understanding of the world’s general structure, of its local and present configuration, of its immediate past and expected future, is then shaped, not by the activities of and the sensory inputs to a single brain, but by the activities of and inputs to a large number of different brains, similarly but not identically situated [4, 252].

This fact resonates with a number of theories which stress that human cognition is informed by culture, and culture is our common enterprise. In this context, Michael Tomasello speaks of “the cultural ratchet” – a genetically rooted mechanisms for accumulating patterns of conduct and passing them on from generation to generation [15]. Culture – at its core – consists of such patterns, which are not limited to purely linguistic behaviour. Thus, the claim that language amplifies, regulates, and coordinates conceptual maps among individuals, is not precise; what plays those roles are any socially shared patterns of conduct, not only linguistic.

All this reinforces the Wittgensteinian claim that language is not a mirror of reality, but rather a tool humans use to coordinate their perception and action. The fundamental role of linguistic expressions is not to mimic the structure of the world, but to coordinate conceptual maps among the members of the given society. To come back to the already considered example, the sentences such as “Jane likes mangoes” help to form the abstract structure (i.e., the relevant set of conceptual maps) of the domain of experience pertaining to persons, their attitudes and feelings, etc., and a particular utterance of the sentence “Jane likes mangoes” serves to index some concrete part of the interlocutor’s conceptual map. At the same time, it is now easier to understand why we tend to think of words as associated with concepts: we believe that the word ‘mango’ denotes a concept of a certain fruit since our perceptual apparatus is so shaped – through the long process of the formation of conceptual maps – that our brains have a vector activation space, which enables us to distinguish mangoes from other fruits. Therefore, we can repeat that it is sentences that have meaning, since they are the functional units for acting in the world, while the meaning of their component words is something derivative, a theoretical construction – or better a reconstruction – of the “living language”. Meanwhile, this reconstruction is made possible by the fact that our conceptual apparatus is organized in such a way that it encodes – in its vector space – the relevant aspects of prototypical mangoes, dogs, tables, and other entities, as well as actions such as grasping, speaking, etc. We believe that in this way the puzzle of the common sense reference disappears: within the framework provided by Churchland’s theory of conceptual maps, it is possible to reconcile the Wittgensteinian picture of language as a tool with our strong conviction that at least some linguistic concepts refer to things in the world.

But is coordinating conceptual maps everything that language does? It may be argued that the answer to this question is negative, and that the role of language far exceeds the mere cross-individual orchestration of perception and action: it opens up a completely new cognitive niche. Let us compare the following *ex post* reconstructions of two scientific discoveries. Richard Hamming tries to imagine how Galileo arrived at the discovery of the law of falling bodies. He says:

Well, Galileo was a well-educated man and a master of scholastic arguments. He well knew how to argue the number of angels on the head of a pin, how to argue both sides of any question. He was trained in these arts far better than any of us these days. I picture him sitting one day with a light and a heavy ball, one in each hand, and tossing them gently. He says, hefting them, "It is obvious to anyone that heavy objects fall faster than light ones – and, anyway, Aristotle says so." "But suppose," he says to himself, having that kind of a mind, "that in falling the body broke into two pieces. Of course the two pieces would immediately slow down to their appropriate speeds. But suppose further that one piece happened to touch the other one. Would they now be one piece and both speed up? Suppose I tied the two pieces

together. How tightly must I do it to make them one piece? A light string? A rope? Glue? When are two pieces one?" The more he thought about it – and the more you think about it – the more unreasonable becomes the question of when two bodies are one. There is simply no reasonable answer to the question of how a body knows how heavy it is – if it is one piece, or two, or many. Since falling bodies do something, the only possible thing is that they all fall at the same speed – unless interfered with by other forces. There's nothing else they can do. He may have later made some experiments, but I strongly suspect that something like what I imagined actually happened [9, p. 87].

It is clear that Galileo's reasoning – as portrayed by Hamming – may be understood in terms of a sophisticated conceptual redeployment. Language might have played a heuristic role here, but a negligible one. Crucial was the insight that the conceptual map which associates the speed of a falling body with its weight is in fact a blend of two conceptual maps, and that these can be separated.

Now, let us consider Abraham Pais' reconstruction of Heisenberg's discovery of quantum mechanics:

What, in one dimension, is a classical orbit? It is described by one coordinate x that varies continuously as a function of the time t , an orbit is given symbolically as $x(t)$. Now Heisenberg seeks inspiration from his previous work with Kramers. There the issue had been to find amplitudes $A(\nu)$ for the scattering of light with frequency ν by an atom. $A(\nu)$ should depend on the transitions from atomic states n to states m , as indicated by the symbol (*not* used by Kramers-Heisenberg) $A_{mn}(\nu)$. Now Heisenberg reasoned (I think), let us try to do something similar for $x(t)$, represent it by the 'quantum symbol' $x_{mn}(t)$, where, to fix ideas, m and n refer to quantum states of a harmonic oscillator, the simplest example he discussed in his paper. There are two possibilities. Either m equals n : $x_{nn}(t)$, which shall represent the coordinate at time t insofar as the system is in state n . Or m does not equal n , when $x_{mn}(t)$ shall represent what one might call a coordinate in transition. Likewise the classical velocity $v(t)$ in the orbit shall be represented by $v_{mn}(t)$. All these quantities satisfy Heisenberg's criterion of being 'in principle observable'. Classically the continuous orbit $x(t)$ satisfies an equation of motion which tells us how the particle moves from one position and velocity to another. Heisenberg assumes that each of the quantities $x_{mn}(t)$ satisfies that same equation. Next he asks: what is the energy of the particle in the state n ? Again he takes over the classical expression for the energy, a function of $x(t)$ and $v(t)$. Classically one proceeds by finding solutions of the equations of motion, substituting those in the energy expression and so obtaining the corresponding energy values. Heisenberg proceeded likewise in trying to find the quantized energies E_n . But now he had to make a crucial decision. The energy of the oscillator (our example) depends on the squares of $x(t)$ and $v(t)$. If one represents $x(t)$ by $x_{mn}(t)$, then how should one represent $x^2(t)$? It seems that the simplest and most natural assumption would be:

$$x_{mn}^2(t) = \sum_k (x_{mk}(t) \times x_{kn}(t)),$$

where \sum_k means: sum over all possible values of k [14].

There is no need to continue with Pais' reconstruction. It should be clear by now that the discovery of quantum mechanics would be impossible without the essential, non-heuristic use of language – in this case, the language of mathematics as applied to the fundamental problems of physics. Of

course, Heisenberg's reasoning required quite complex conceptual redeployments – the connections between his conceptual maps must have changed to attain new insights pertaining to the behaviour of the hydrogen atom. However, it is quite unlikely that if someone else independently discovered quantum mechanics, her conceptual redeployments would resemble what happened in Heisenberg's brain. The reason is that mathematics and mathematical physics are quite abstract, and abstract language does not necessarily lead to cross-individual coordination of conceptual maps. For example, many studies suggest that the differences between brain representations of relatively simple mathematical concepts (such as natural numbers) are quite substantial [2].

This point may be generalized: arguably all abstract linguistic concepts in any domain, not only in mathematics, have no well-defined counterparts at the level of the conceptual maps in the brain. Let us consider, for example, the concept of justice. It is unlikely that one would be able to pinpoint any (aspect of) a particular conceptual map or a set of such maps that would correspond to 'justice' (while there may be representations of some classes of acts which usually count as just or unjust). This is not surprising: given the function of language there is no – and there cannot be any – one-to-one correspondence between linguistic concepts and the vector spaces which represent reality at the brain level. Abstract concepts feature in linguistic expressions which, when uttered by parents, teachers, and other people we talk to, do in fact shape our conceptual maps, making them more similar to the maps of the members of our society. But while one can identify corresponding aspects of the conceptual maps for linguistic concepts such as 'mango', 'table' or 'grasping', the same cannot be said of 'justice', 'punitive damages', 'noncommutative geometry' or 'muon'. Higher mathematics, mathematical physics, philosophies of law and of literature, etc., do not exist in our brains – they are publicly shared theories expressed in a suitable language. Of course, through the coordination of individual conceptual maps they do influence the way we perceive and act in the world. But they do more than that. As the example of Heisenberg's discovery and other theories of contemporary physics show, language can give us access to the world which by far exceeds the grasp of reality enabled by the mechanism of the formation of conceptual maps. This fact only reinforces the mystery behind the puzzle of scientific reference: language is not only a tool for acting in the world. While it serves us well in coordinating our conceptual maps, it does more than that – it enables an insight into the structure of reality which lies beyond any individual cognitive apparatus.

4. Coda

In this paper we have argued that the theory of conceptual maps developed recently by Paul M. Churchland provides additional support for Wittgenstein's claim that language is a tool for acting in the world. The role of language is to coordinate and shape the conceptual maps of the members of the given language community, reducing the cross-individual cognitive idiosyncrasies and paving the way for joint cognitive enterprises. Moreover, Churchland's theory also explains our tendency to speak of language as consisting of concepts which correspond to things we encounter in the world. The puzzle of common sense reference is no longer a puzzle: while at the fundamental level language remains a tool for orchestrating conceptual maps, the fact that the maps encode some communally shared categorization of experience fuels our talk of concepts capturing the essences of things, natural kinds, prototypes, etc. At the same time, Churchland's theory fails short of resolving the puzzle of scientific reference: that the mathematical equations of physics help us capture some enduring aspects of reality, which lie far beyond the limits of our sensual experience, remains a mystery [2, p. 203-236]. We believe, however, that this is not an argument against the Wittgensteinian view of language as a tool or Churchland's theory of conceptual maps. Rather, it is a phenomenon which requires metaphysical explanations [1, p. 151-187].

References

1. Brożek, B. *Rule-Following. From Imitation to the Normative Mind*, Kraków: Copernicus Center Press, 2013.
2. Brożek, B., Hohol, M. *Umysł matematyczny*, Kraków: Copernicus Center Press, 2014
3. Brożek, B., Olszewski, A. "The Mathematics of the Transcendental Ego", *Copernicus Center Reports* 2, 2011, 75-124.
4. Churchland, P.M. *Plato's Camera How the Physical Brain Captures a Landscape of Abstract Universals*, Cambridge, MA: The MIT Press 2012.
5. Descartes, R. *The Principles of Philosophy*, translated by J. Veitch, <http://www.gutenberg.org/>, 1901.
6. Gärdenfors, P. *Conceptual Spaces: The Geometry of Thought*, Cambridge, MA: The MIT Press, 2000.
7. Gärdenfors, P. *The Geometry of Meaning: Semantics Based on Conceptual Spaces*, Cambridge, MA: The MIT Press, 2014.
8. Fodor, J. *The Language of Thought*, Crowell Press: New York, 1975.
9. Hamming, R. "The Unreasonable Effectiveness of Mathematics", *The American Mathematical Monthly*, 87(2), 1980, p. 81-90.
10. Kant, I. *Critique of Pure Reason*, translated by P. Guyer and A.W. Wood, Cambridge University Press: Cambridge 1998.
11. Locke, L. *An Essay Concerning Human Understanding*, Hazleton: Pennsylvania University Press, 2009
12. Margolis, E., Laurence, S. "Concepts", *The Stanford Encyclopedia of Philosophy* (Spring 2014 Edition), E.N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2014/entries/concepts/>.
13. Müller, M. "On Thought and Language", *Monist* 1(4), 1891, p. 572-589.
14. Pais, A. *Niels Bohr's Times in Physics, Philosophy, and Polity*, Oxford – New York: Oxford University Press, 1991.
15. Tomasello, M. *The Cultural Origins of Human Cognition*, Cambridge, MA: Harvard University Press, 2001.
16. Vygotsky, L. *Thought and Language*, Cambridge, MA: Harvard University Press, 1986.
17. Winger, W. "The Unreasonable Effectiveness of Mathematics in the Natural Sciences", *Communications on Pure and Applied Mathematics* 13(1), 1960, p. 1-14.
18. Wittgenstein, L. *Philosophical Investigations*, translated by G.E.M. Anscombe, Oxford: Basil Blackwell, 1986.