SEÇÃO: VARIA

Mental or Neural Representations: Justifying the terminology used in cognitive neuroscience[1]

Representações mentais ou neurais: justificando a terminologia utilizada na neurociência cognitiva

Representaciones mentales o neurales: una justificación de la terminología utilizada en neurociencia cognitiva

Fabiana Mesquita Carvalho[2]
orcid.org/0000-0001-8730-3060
bibimesq@gmail.com

Nara M. Figueiredo[3]
orcid.org/0000-0003-0270-7467
naramfigueiredo@gmail.com

Abstract: In this paper we investigate whether one of the most common uses of the concept of representation is justifiable by suggesting the conditions under which it can be accepted and how it can be related to mental states. We present mental states in terms of private experiences and public events. We argue that a representation is a relation involving three main elements as well as the user of the representation, and defend that the conditions in which we can conceive neural activity as representational are set by the context of observing a correlation between public events and patterns of neural activity. We aim at demonstrating that neural activity can be seen as both representational and non-representational - but rather constitutive - depending upon if we are considering public events under the perspective of the observer, or if we are considering private experiences under the subjective perspective.

Key words: neural representation, mental representation, cognitive neuroscience, Wittgenstein.

Resumo: Neste artigo, investigamos se um dos usos mais comuns do conceito de representação é justificável, sugerindo as condições sob as quais ele pode ser aceito e como pode ser relacionado aos estados mentais. Apresentamos estados mentais em termos de experiências privadas e eventos públicos. Argumentamos que uma representação é uma relação que envolve três elementos principais, bem como o usuário da representação, e defendemos que as condições nas quais podemos conceber a atividade neural como representativa são definidas pelo contexto da observação de uma correlação entre eventos públicos e padrões de atividade neural. Nosso objetivo é demonstrar que a atividade neural pode ser vista como representacional e não representacional - mas constitutiva - dependendo se estamos considerando eventos públicos sob a perspectiva do observador, ou se estamos considerando experiências privadas sob a perspectiva subjetiva.

Palavras-chave: representação neural, representação mental, neurociência cognitiva, Wittgenstein

Resumen: En este artículo investigamos si uno de los usos más comunes del concepto de representación es justificable al sugerir las condiciones bajo las cuales puede aceptarse y cómo puede relacionarse con los estados mentales. Caracterizamos los estados mentales en términos de experiencias privadas y eventos públicos. Argumentamos que una representación es una relación que involucra tres elementos principales además del usuario de la representación, y defendemos que las condiciones en las que podemos concebir la actividad neuronal como representativa se establecen contextualmente, a partir de la observación de una correlación entre eventos públicos y patrones de actividad neural. Nuestro objetivo es demostrar que la actividad neuronal puede considerarse tanto representativa como no representativa (más bien constitutiva), dependiendo de si estamos considerando eventos públicos bajo la perspectiva

---

del observador, o si estamos considerando experiencias privadas bajo la perspectiva subjetiva.

Palabras clave: representación neural, representación mental, neurociencia cognitiva, Wittgenstein.

1 Introduction

Cognitive neuroscience is surely developing at a fast pace. Eric Kandel has suggested that "cognitive neuroscience - with its concern about perception, action, memory, language and selective attention - will increasingly come to represent the central focus of all neurosciences in the 21st century." (Kandel & Squire, 2000). On that account, there is an urge for conceptual clarification. Some foundational concepts of cognitive neuroscience, such as the one of *mental representation*, have been brought from cognitive science, which for decades was entirely dependent upon the presupposition of symbolic internal representations and computations over these representations (Fodor, 1975). In fact, with very few exceptions, the commonly held theories in psychology and cognitive science still make use of the notion of representation, but there has been little concern to further explain what representation exactly is, and whether its use is appropriate within these contexts. The concept of representation in the cognitive sciences covers the description of a wide compass of phenomena, from mental states at the personal-level (e.g., beliefs, feelings, desires), to patterns of neuronal activity in a certain brain area at the sub-personal-level (e.g., firing rate of a neuronal population in the orbitofrontal cortex). On the other hand, the use of representational vocabulary in cognitive neuroscience is more specific and directed to characterizing neural processes and patterns of activation, being thus focused on sub-personal representations and their role in explaining behavior. Cognitive neuroscientists talk about sub-personal representations both in terms of active or on-line representations (e.g. a current perceptual or emotional state) and stored or off-line representations (e.g. memory). According to their causal description, when the neural representation performs "the right functional role in the rest of the neural architecture, then they can be the basis on which a whole person has mental properties like perceiving and desiring" (Shea, 2013a).

In this paper, we address the question of what neural representation is. We focus on clarifying whether neural activity can indeed be taken as representational, and how it could be related to mental states. First, we introduce the foundations of the concept of representation, and refer to its use in cognitive neuroscience (Section 2). Next, we will examine the conditions under which the conception of neural activity as representational can be accepted (Section 3). And, finally, we will consider Wittgenstein's *Philosophical Investigations* to clarify the conceptual misunderstandings that have led us to the problem of mental representation (Section 4).

These steps will allow us to defend that neural activity can both *represent* and *constitute* mental states. We will first argue that neural activity *represents* to the extent that there are data strongly associated, by correlation, with public events, which are expressions of mental states, and could indeed be interpreted as neural representation in cognitive neuroscience. Then, we will argue that neural activity *is constitutive of* mental states (e.g. perception, feeling or desire) when it is considered from the perspective of the subject. Finally we will suggest, based on wittgensteinian arguments, that the use of the concept of representation referring to any private mental state is awry.

2 The foundations of the concept of representation

We can certainly say that a map represents a city, that smoke represents fire, and many people defend that language represents thoughts and scenarios. Nevertheless, the means of representation and, therefore the reasons why these cases are called 'representation' are quite different. The map example displays an association based on rules of representation. This can be said to be the most fitting example of representation, for it displays clearly, the represented element, namely, the city, the very representation, namely, the map, and the conditions in which it represents (the associative rules – e.g. the ratio, the position of the map). The fire example, on the other hand, is

Fabiana Mesquita Carvalho · Nara M. Figueiredo
Mental or Neural Representations: Justifying the terminology used in cognitive neuroscience

3/9

understood as representational with regards to causality. There is a causal association between fire and smoke that allows us to take that as a paradigmatic relation. Although not every time that there is fire we also have smoke, and not every time that there is smoke there is also fire, we take smoke as representative of fire. This is due to the logically inductive aspect of a relation between natural events that we call causality.

Linguistic representations, alternatively, are uses of language that we can call descriptions. Language represents to the extent that a word stands for something else. Not all language is representational. Most of our use of language is, in fact, instrumental, not representational. The linguistic representation relation is established by means of learning processes in which we grasp the uses of words, and this learning requires training (Wittgenstein, 2009)[4].

Despite the fact that there are different kinds of representation, there are similarities amongst them that give us a picture of the use of the concept of representation. These similarities are (i) a represented element, (ii) the one that represents, and (iii) the means of representation (Table 1). In the map example, the map represents the city by means of *associative rules*. In the fire example, the smoke represents fire by means of a reference to a *paradigmatic causal relation*. In the language example, language represents the world by means of *associative training*. These relations are justifiably considered means of representation for the three of them can be presented as an answer to the question 'How do you know that A represents B?'.

In addition to that, cognitivist theories emphasize the *user of* the representation. The tripartite concept of representation states that a representation must be used as a 'stand-in' by someone to count as a representation (Slezak, 2002). Therefore, a basic definition of representation is: 'A' stands for 'B' by means of 'm' according to 'U'. Where 'A' is the represented element, 'B' is the one that represents, 'm' the means of representation and 'U' is the user of the representation.

In recent years the cognitive sciences have been facing substantial disagreement over whether the concept of representation can be used to refer to mental processes or states (Chemero, 2009, Clark, 2013). As aforementioned, the concept of representation in the cognitive sciences can be used in the context of describing both personal- and subpersonal-level phenomena. The former concerns conscious mental states, such as a perceptual state, and the latter applies to unconscious patterns of neuronal activity underpinning the perceptual state. In other words, the perceptual state of seeing a book is a personal-level phenomena whereas the firing patterns of neurons in the visual cortex as the agent sees the book is a subpersonal-level property of the agent (Shea, 2013). In the present article, we will adopt the notions of *mental representation* and *neural representation*, which are theoretical constructs of the cognitive sciences (also used in cognitive neuroscience) (Decharms & Zador, 2000), to refer to personal- and subpersonal-level phenomena, respectively.

As we suggested, given that (I) to represent is to stand for something else other than what is represented, (II) that we should consider the means of representation, and (III) the user of the representation must be taken into account, our question is 'In what sense do mental or neural representations represent?'

In the next section we argue that neural activity can be said to be representational to the extent that there is a correlation, observed by the researcher, between them and public events. Public events or objects are those capable of being publicly observed (as opposed to private states or experiences). It includes verbal and non-verbal expressions of private states of a subject (e.g., conscious and unconscious physiological signs) as well as any social, cultural and natural object or phenomena that are susceptible to interpretation. The observed correlation allows us to infer one from the other (i.e., the patterns of neural activity and the public events) (Figure 1). The question

---

4   The distinction between descriptive and instrumental uses of language (Wittgenstein, 2009) leads us to the idea that language does not represent thoughts and emotions. Language represents a scene, but it *expresses* thoughts and emotions. *Expressing* is not a relation as much as representing. (See more on language as expressing thoughts on neo-expressivist accounts of private experiences, e.g.: Bar On, D. 2004)

of how neural representations represent is, then, transformed into the question for the means of representation (which is, the observed correlation).

3 Conditions under which the conception of neural activity as representational can be accepted

Regardless if a theory takes perception, cognition and action as brain-secluded and passive processes or as a dynamic and unified process coupled with the environment, both views are capable of accepting perception, cognition and action as processes underpinned by neurally-distributed representations (Clark, 2015, 2016). Whether these models assume representations to be richly reconstructive or non-reconstructive, this matter is not important for the argument that will follow. What does it mean to refer to brain activity as neurally-distributed representations?

Neural activity recordings can be done using several different techniques, and the analysis and modeling of neural activity data can also be done using several different approaches (Quiroga & Panzeri, 2013). Despite that the different analysis and modeling approaches focus on different data encoding schemes and different statistical estimates, they are all based on some type of *correlational* measure (Figure 2) which will result on the so called neural population codes (or simply 'neural codes').

It is important to clarify that neural codes are not the same as the raw neural activity recorded during the whole extension of an experiment. Raw data is meaningless until patterns are identified through statistical modelling. Thus, for the neural data to be informative, the raw neural activity must undergo a data processing stage (i.e., statistical analysis phase) in which spatio-temporal markers are used as reference points to identity activity patterns within raw data. These spatio-temporal markers are dictated by the experimental protocol (i.e., stimuli presentation rules). Since these reference points are previously known, it is possible to search for patterns of neural activity, that is, to infer an association between stimulus presentation and certain patterns of neural activity. In other words, it is only because there is a data process-

ing algorithm able to identify activity patterns in raw data that neural codes are symbolic. Thus, neural codes can represent once the correlation (i.e., the means of representation) between the represented element (i.e., the stimulus) and the very representation is established.

Cognitive neuroscientists infer that, by means of neural codes, neural activity represents personal-level phenomena, such as perceiving, desiring, feeling and believing. In this sense, they refer to a representation as a pattern of neural activity that reliably co-occurs with the presentation of a particular stimulus. For instance, functional MRI studies have shown neural codes in the lower and higher visual areas that reliably represent, respectively, spatial orientation of the stimulus and object category, both in seen and imagined objects (Kamitani & Tong, 2005; Haxby et al., 2001; Horikawa & Kamitani, 2017).

For instance, a recent analysis technique, called representational similarity analysis, exploits the observation that the same population of neurons can represent several stimuli since it can undergo different patterns of activity (Figure 2b). Hence, there are many possible combinations of activity patterns within the population which provide a rich representational space – space which can be analyzed geometrically. Representational similarity analysis techniques assess the correlation between pairs of neural representations, and can be used to compare representational geometries between brain regions and between stimulus types within and between subjects, for instance (Cichy et al 2014; Charest & Kriegeskorte, 2015; Kriegeskorte & Kievit, 2013).

Given the descriptions above, we can consider the analogies presented in the previous session as a sound argument in favour of the idea that the use of representation in cognitive neuroscience is legitimate, for the representational structure shown by the analogies is identified on the relation between neural codes and stimuli. This structure (see Table 1) can be summarized as the represented element (i.e., the public event), the means of representation (i.e., the method or medium) and the representation itself (i.e., the symbol). In other words, a representation is a relation which involves the public event, the method, and the symbol. These three aspects

are intrinsically related, they need each other to be what they are. In this sense, what justifies the use of the term representation for the neural codes is the idea that they are only symbolic when they are in the representational relation. A symbol that is not part of relational structure does not represent (it is empty). As we mentioned before, the relation requires the interpreter or user of the representation. In other words, the obtaining of the representational relation depends on subjects with cognitive capacities.

Our interlocutor might argue that something (e.g., a certain symbol) that doesn't have content doesn't represent anything. We are arguing for the opposite: something can represent and yet not carry content. According to the foundations of the notion of representation that we explore in Section 2, the fact that something represents doesn't entail that it has representational content. We show, by analogical argument that something can represent by other means (see Table 1). This is possible in virtue of a reader. In other words, the "aboutness" is dependent on an observer that knows (or has learnt) how to read a certain symbol.

Therefore, we consider legitimate the use of the concept of representation by neuroscientists to express how neural activity is connected to public events including expressions of personal-level phenomena (or mental states). The neuroscientist is capable of showing that processed neural activity represents certain stimulus-induced (e.g. perception) or imagery-induced (e.g. memory, imagination) public events by means of correlation (given by the neural codes) (Figure 1). We defend that neuronal activity represents to the extent that it is possible to establish a relation based on the observation of neuron activity patterns and public events. Neural representations are wrongly believed to possess content as a result of their relations to the external world. To this extent, although content can be related to neural activity, it is not a property of it, i.e. neural activity does not *possess* content, but it can be content *related*. Thus, it only makes sense to refer to neuronal activity as neural encodings or neural representation from the perspective of an observer (e.g. the neuroscientist). It is the neuro-

scientist who is going to interpret the measured neuronal activity, according to her own analytical methods, as representational. Indeed, describing neuronal activity as representational is a helpful tool developed and used by neuroscientists to build theories of the brain functioning.

4 Mental Representation - why private mental states cannot be accepted as representational

What we call a private mental state is the personal level of experience and it is, therefore, essentially inner. Contrary to representationalist accounts in cognitive science which embrace the notion that mental states can be taken as mental representations, we argue that in the personal level there is no representation relation. Perceiving the world is not, in the personal level, representing the world. Perceptions, sensations, desires, thoughts and all the scope personal level of experience can reach are also simply mental states. These states are not content related as neural representations can be content related, because they cannot be in a representative relation as we described in section 2. This is justified by two arguments from Wittgenstein's *Philosophical Investigations (2009),* namely, the *private language argument* and the *beetle in a box* example (Figueiredo, 2019). According to these two arguments Wittgenstein dismisses the idea that private states are referential elements because (1) there can be no identity criteria for inner states, and (2) there can be no verification criteria for inner states. These two reasons are given when he considers relations between private states and between a private state and a name, respectively.

More specifically, the *private language argument* tells us why the private state cannot count as a representational reference to the meaning of intentional states words, and the *beetle in the box example* tells us that the inner state is also not a necessary condition to the meaning of the word, once the meaning is defined publicly.

The *private language argument* consists, in a nutshell, in denying that we can, or should, refer to private states to define them, as we can do to

define objects ostensively. This argument is based on the rule following example which shows that a private rule cannot count as a rule, for, thinking I am following my private rule is the same as following it. In other words, private states cannot be taken as standard for correctness for whatever we think they are, it is what they are (Wittgenstein, 2009; § 258). This lack of identity criteria can also be illustrated by memory cases. If I remember an event x and cannot count on any objective, or public, criteria to verify if the memory is, in fact, correct, I will continue to think I remember the event correctly, and, more importantly, this is what I remember. In other words, I cannot check my own memory to verify the correctness of it, because I simply remember.

The problem Wittgentein addresses with the beetle in the box example is our assumption that the relation between private states and language must have the private state as a reference in the same way that public objects are taken as a reference when we define them ostensively. He suggests that the supposed private 'mental object' - let us call it private state - is dropped out of consideration if we conceive a relation between private states and a name (Wittgenstein, 2009; §293). As we can see in the *beetle in the box* example: Every person has a box and has exclusive access to its content. Everyone calls the content of its own box 'beetle'. Although no-one has ever seen the content of the box of the other, they all speak of their private content and of each other's by means of the same word, namely, 'beetle'(Wittgenstein, 2009; §293). This example can be more illustrative if one takes an non-existent English word: take 'nafa'. Now imagine I'm referring to some inner state of mine and using this word. Wouldn't you be asking for more information? Would you know what I'm referring to? Wouldn't you be trying to determine whether I'm referring to some pain, or to some thought, to some image or memory? Wouldn't you, after that, try to determine it more specifically on the basis of the words and things you know by public reference? Let's say you get to discover that I'm imagining an object. Wouldn't you point to several things and ask, is it this? Is it that? Is it this color? Wittgenstein's analogy shows that the private experience is not relevant to the

meaning of the words that we use to talk about them. When we consider them as irrelevant, this does not imply that they do not exist or that we do not (or cannot) know them, but merely that they are not a reference for us to publicly use words and understand words that express them.

One can understand the beetle in the box as the subjective introspection. It is acknowledged in experimental psychology, since Wundt, that subjective introspection is inaccurate, and the search for objective investigation has been driving scientists since then. What Wittgenstein is saying with this argument is that one does not look into his own experience in order to know what one is seeing. One simply has the experience and this is seeing.

These two arguments don't imply that our private states are irrelevant for our understanding of them, but merely that they are not to be taken as a reference when we want to determine their meaning (Figueiredo, 2019). Our private states are a condition for our understanding (for example, that I see red is a condition for me to understand what the word 'red' refers to when it refers to a color, but my mental image is not a reference for the meaning of the word 'red' in the same way 'red' stands-for for the color red in a visual description). In other words, we need a common perceptual biological apparatus allowing us to perceive what can be taken as paradigms that serve as identification criteria, for instance, a public sample of red. Perceiving is a condition for us to be able to make distinctions. Even if we cannot know whether our private experience is the same as of the other person, it must be possible that I can distinguish in the same way that you can distinguish. That is, even if the perception is not the same in terms of private experiences, the paradigms can be standardized when it comes to public distinctions.

As we described in section 2 we should not consider representation as an object with semantic properties, such as content, reference and truth values. A representation is a relation that involves three aspects: what is represented, the means of representation and the representation itself – and in addition to that, it also involves the interpreter or user of the representation. Thus, private states cannot be part of a representational relation, for

Fabiana Mesquita Carvalho · Nara M. Figueiredo
Mental or Neural Representations: Justifying the terminology used in cognitive neuroscience

7/9

there is nothing to count as a standard for correctness. This means that they can be neither a represented element, nor a representation itself. Understanding that the mental private state cannot be representational is a great step towards an understanding of the mental scope. For it shows us that the status of the private states cannot be analogous to the status of public objects.

We should also say that physical properties, aspects or states that are intrinsic to the neural system (see Shea 2013b) do not have semantic properties. For, as we have mentioned before, physical properties, aspects or states can represent as long as there is a correlation between them and other observed public events. The relation by means of an observed correlation is a requirement for a neural activity pattern to be taken as representational. Our claim that neural representation is only conceivable if we consider an observer conforms to the theories in which the user of the representation must be taken into account (Slezak, 2002; Shapiro, 2011).

5 Concluding remarks

We have suggested that representation is a relation (section 2), that neural activity can represent public events (section 3) and that intentional states are not representational. The three sessions of this paper come together to build a picture that allows us to answer the question of how neural activity can both *represent* and *be constitutive of* mental states.

Public events, which include, but are not exhausted by verbal and non-verbal expressions of private states, can be assessed in scientific experiments. Given that, the neural activity underlying public events can be apprehended. They can be measured, encoded into patterns, correlated to the public events in question, and considered as neural representation. From the observer's perspective (e.g., researcher), expressed private states are *represented* in patterns of neural activity by means of correlation (Figure 1).

On the other hand, when the subject perceives, thinks, desires and so on, she is not representing, she *is* merely perceiving, thinking and desiring. For instance, a subject who perceives something (e.g., sees a cat) is neither perceiving her own represen-

tations of the cat nor representing her perception of the cat to herself. The neural activity going on in the subject's brain is not representational, but *is constitutive* of the subject's private state of perceiving the cat. The physical aspect of neural activity is part of the very nature of private states.

If we accept that neural activity is constitutive of private experiences, we can say that neural activity can be seen as both representational and non-representational (but rather constitutive) depending upon if we are assuming the position of the observer, which is tied to public events, or the position of the subject, which concerns private states. We believe that this perspective shift allows us to clarify the grounds on which the concept of representation can be properly used in cognitive neuroscience. This is fundamental for avoiding conceptual misunderstandings and misleading inferences on the search for the neural underpinnings of the mind.

7 Appendix

FIGURE 1
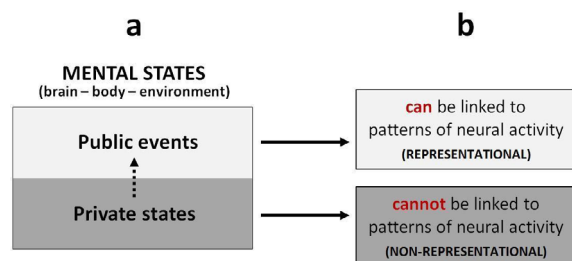
Figure 1 – Representational and non-representational aspects of mental states. a. Mental states can be seen as complex phenomena which emerge from nonlinear interactions of the physical properties of an agent's brain, body, and environment (Beer 2000; please note that our view on neural representation is committed neither with reductionist nor with non-reductionist accounts of the brain/mind relation). Based on Wittgenstein, we can say that mental states can be understood in terms of private states (dark grey) and their public expressions (dashed arrow), which become public events and can be observed (light grey). **b**. Only public events can be linked to patterns of neural activity and taken as representational, both for data averaged across subjects and for subject-unique data analysis. Any form of private state determination (including that of neural representation) is impossible since there is no reference or standard for it to be compared or related to.

FIGURE 2
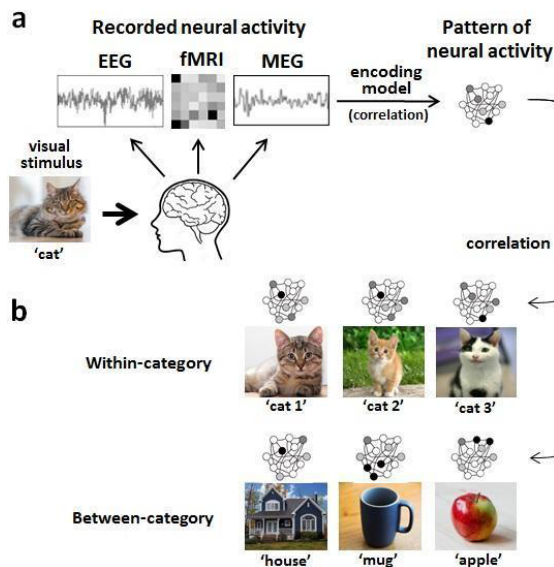
Figure 2 – Schematic illustration of the steps from sensory stimulus to neural representation in a simple experimental design. a. EEG, fMRI or MEG activity is recorded while subjects view natural images. The recorded neural activity elicited by the visual stimulus ('cat') is then encoded into patterns of activity by an encoding model. In the figure, each image elicits activity in the same neural network, with individual neurons displaying low to high firing rates (colour light grey to black), so each image is said to be represented by a unique spatiotemporal pattern of activity across the network. **b**. These encoded activity patterns, which are neural representations, could be determined for several other visual stimuli, and could be, for instance, compared (correlated) to the activity patterns representing within-category or between-category stimuli. These correlations assess the categorical distance between representations, and could also be compared to behavioural categorization judgements (Carlson *et al*., 2014). To this end, the correlations reveal to what extent neural representations reflect stimulus properties, and also, to what extent they are reflected in behaviour

TABLE 1

Representation itself OR The symbol (public event)	Means of representation OR Elo	Represented element OR Actual object or event (public event)
Map E.g. A city map	Associative rules	What is mapped E.g. The city
Smoke	Causal relation	Fire
Language E.g. 'Book'	Associative training/learning processes	What is meant/understood E.g. The object book
Neuronal activity (The neural codes [5])	Correlation (causal ou not).	Stimulus/public event

TABLE 1 – Scheme of the three elements of a representational relation. The first column on the left shows examples of what we call 'representation itself'. It refers to what can be considered as symbolic. The second column, in the middle, shows examples of the means that are necessary for a relation to obtain. The third column, on the right, shows the element of the relation that can be referred to when one has to explain what the representation means. We have opted for suggesting alternative names for each of the elements of the representational relation in the hopes of avoiding confusions. Thus, 'representation itself', 'means of representation' and 'represented element' can be named as 'symbol', 'elo' and 'public event' respectively.

---

[5] Neural codes are the processed neural activity, not the raw one. This is why they are symbolic. Because there is a data processing stage (this can be understood as statistical analysis algorithms) in which spatio-temporal markers are used as reference points to identity activity patterns in raw data (raw data is the measurement and recording of pure neuronal activity. This data is meaningless until patterns are identified through statistical modelling).
Neural codes can represent once the correlation (the means of representation) between the represented element and the very representation is known.

Fabiana Mesquita Carvalho · Nara M. Figueiredo
Mental or Neural Representations: Justifying the terminology used in cognitive neuroscience

9/9

References

BAR ON, D. *Speaking my Mind*. New York: Oxford University Press, 2004. https://doi.org/10.1093/0199276285.001.0001.

CARLSON T. A; RITCHIE J. B.; KRIEGESKORTE, N.; DURVASULA, S.; MA, J. Reaction time for object categorization is predicted by representational distance. *Journal of Cognitive Neuroscience*, Cambridge, Mass., v. 26, n. 1, 1, p.132–142, 2014. https://doi.org/10.1162/jocn_a_00476.

CHAREST I.; KRIEGESKORTE, N. The brain of the beholder: honouring individual representational idiosyncrasies, Language, Cognition and Neuroscience, [*s. l.*], v. 30, n. 4, 367-379, 2015. https://doi.org/10.1080/23273798.2014.1002505.

CICHY R. M.; PANTAZIS, D.; OLIVA, A. Resolving human object recognition in space and time. Nature Neuroscience, London, v. 17, p. 455-462, 2014. https://doi.org/10.1038/nn.3635.

CLARK A. Predicting peace: An End to the Representation Wars. *In*: METZINGER, T.; M. Windt, J. *Open Mind*. 7.th ed. Frankfurt am Main: MIND Group, 2015. p. 1-7.

CLARK A. Busting out: predictive brains, embodied minds, and the puzzle of the evidentiary veil. Noûs, 2016. Disponível em: doi: 10.1111/nous.12140. Acesso em: 02 mar. 2020. https://doi.org/10.1111/nous.12140.

DECHARMS, R. C.; & ZADOR, A. Neural representation and the cortical code. *Annual review of neuroscience*, Palo Alto, Calif., US, v. 23, n. 1, 613-647, 2000. https://doi.org/10.1146/annurev.neuro.23.1.613.

DENNETT, D. *Content and consciousness*. London: Routledge, 1969.

FIGUEIREDO, N. M. The Problem of the Unbridgeable Ontological Gap Between the Brain and Our Private Experiences in Wittgensteinian Terms. *Discurso*, São Paulo, v. 49, n. 1, p. 185-203, 2019. Disponível em: doi. org/10.11606/issn.2318-8863. discurso.2019.159310. Acesso em: 02 mar. 2020. https://doi.org/10.11606/issn.2318-8863.discurso.2019.159310.

FODOR, J. *The Language of Thought*. Cambridge: Harvard University Press, 1975.

HAXBY, J. V. *et al*. Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*, New York, v. 293, p. 2425–2430, 2001. https://doi.org/10.1126/science.1063736.

HORIKAWA, T.; KAMITANI, Y. Generic decoding of seen and imagined objects using hierarchical visual features. *Nature Communications*, v. 8, Article number: 15037, p. 1-15, 2017. Disponível em: doi: 10.1038/ncomms15037. Acesso em: 02 Mar. 2020. https://doi.org/10.1038/ncomms15037.

KAMITANI, Y.; TONG, F. Decoding the visual and subjective contents of the human brain. *Nature Neuroscience*, London, v. 8, p. 679-685, 2005. https://doi.org/10.1038/nn1444.

KANDEL E. R.; SQUIRE, L. R. Neuroscience: Breaking down scientific barriers to the study of brain and mind. *Science*, New York, v. 290, p. 1113–1120, 2000. https://doi.org/10.1126/science.290.5494.1113.

KRIEGESKORTE, N.; KIEVIT, R. A. Representational geometry: integrating cognition, computation, and the brain. Trends Cognitive Science, [*s. l.*], v. 17, p. 401–12, 2013. https://doi.org/10.1016/j.tics.2013.06.007.

QUIROGA, R. Q.; PANZERI, S. *Principles of neural coding*. [*s. l.*], CRC Press, 2013. https://doi.org/10.1201/b14756.

SHAPIRO, L. *Embodied cognition*. [*S. l.*]: Routledge, 2011. https://doi.org/10.4324/9780203850664.

SHEA, N. Neural mechanisms of decision-making and the personal level. *In*: FULFORD, K. W. M.; DAVIES, M.; GIPPS, R. G. T.; GRAHAM, G.; SADLER, J. Z. G.; THORNTON, T. (Ed). *The Oxford Handbook of Philosophy and Psychiatry*. Oxford: Oxford university press, 2013a. p. 1063-1082. https://doi.org/10.1093/oxfordhb/9780199579563.013.0062.

SHEA, N. Naturalising representational content. *Philosophy Compass*, [*S. l.*], v. 8, n. 5, 496–509, 2013b. https://doi.org/10.1111/phc3.12033.

SLEZAK, P. The tripartite model of representation. P*hilosophical Psychology*, Abingdon, v. 15, n. 3, p. 239-270, 2002. https://doi.org/10.1080/0951508021000006085.

WITTGENSTEIN, L. *Philosophical Investigations*. Trad. G. E. M. Anscombe, P. M. S. Hacker and Joaquim Shulte. Oxford,UK: Wiley Blackwell, 2009.

Fabiana Mesquita Carvalho

Post-Doctoral Researcher — University of Sao Paulo — FFLCH, USP.

Nara M. Figueiredo

Post-Doctoral Researcher - University of Campinas. UNICAMP. This is my current affiliation. This work was developed while Dr. Fabiana and I were both researchers in the philosophy department of the University of São Paulo - USP, under the supervision of Prof. Osvaldo Pessoa. We would like to thank the department and the fundamental support of Prof. Osvaldo.

Mailing address:

Dra. Fabiana Mesquita Carvalho

Universidade de São Paulo - USP

Av. Prof. Luciano Gualberto, 315, Sala 1007

Cidade Universitária 05508-010

São Paulo, SP, Brasil

Dra. Nara Miranda de Figueiredo

Universidade Estadual de Campinas - UNICAMP

Rua Sérgio Buarque de Holanda, 251

Cidade Universitária "Zeferino Vaz"

Barão Geraldo 13083-859

Campinas, SP, Brasil