CellPress REVIEWS

coding is ill suited for flexibility, but is fundamental for adaptive fast reactions.

[1]Center for Neural Science, New York University, New York, NY, USA
[2]MySpace Lab, Department of Clinical Neuroscience, University Hospital of Vaud (CHUV), Lausanne, Switzerland

*Correspondence: andrea.serino@unil.ch (A. Serino).

https://doi.org/10.1016/j.tics.2019.01.001

References

1. Bufacchi, R. and Iannetti, G. (2018) An action field theory of peripersonal space. *Trends Cogn. Sci.* 22, 1076–1090
2. Serino, A. *et al.* (2015) Body part-centered and full body-centered peripersonal space representations. *Sci. Rep.* 5, 18603
3. Cisek, P. and Kalaska, J.F. (2010) Neural mechanisms for interacting with a world full of action choices. *Annu. Rev. Neurosci.* 33, 269–298
4. Matelli, M. and Luppino, G. (2001) Parietofrontal circuits for action and space perception in the macaque monkey. *Neuroimage* 14, S27–S32
5. Pellencin, E. *et al.* (2017) Social perception of others shapes one's own multisensory peripersonal space. *Cortex* 104, 163–179
6. de Haan, A.M. *et al.* (2016) Approaching threat modulates visuotactile interactions in peripersonal space. *Exp. Brain Res.* 234, 1875–1884
7. Taffou, M. and Viaud-Delmon, I. (2014) Cynophobic fear adaptively extends peripersonal space. *Front. Psychiatr.* 5, 122
8. Duhamel, J.R. and Goldberg, M.E. (1998) Ventral intraparietal area of the macaque: congruent visual and somatic response properties. *J. Neurophysiol.* 79, 126–136
9. Graziano, M.S. *et al.* (1997) Visuospatial properties of ventral premotor cortex. *J. Neurophysiol.* 77, 2268–2292
10. Noel, J.-P. *et al.* (2018) From multisensory integration in peripersonal space to bodily self consciousness: from statistical regularities to statistical inference. *Ann. N. Y. Acad. Sci.* 1426, 146–165
11. Cléry, J. *et al.* (2015) Impact prediction by looming visual stimuli enhances tactile detection. *J. Neurosci.* 35, 4179–4189

Letter

The Value of Actions, in Time and Space

Rory J. Bufacchi[1,2] and
Gian Domenico Iannetti[1,2,*]

Noel and Serino raise several interesting points about our recent functional interpretation of peripersonal space (PPS) [1]. We are grateful because their commentary gives us an opportunity to clarify aspects of our reasoning. In our critique of the vague notions surrounding PPS we may, ironically, have been too vague ourselves. As a result, Noel and Serino appear to repropose more or less exactly what we intended to originally propose.

Noel and Serino eloquently show how an action-value framework can be fruitful to discussions about PPS – something with which we can only strongly agree. We briefly note, for the sake of clarity, that the word 'value' they use (a reinforcement learning concept, Box 1) corresponds to the term 'action relevance' in our opinion article [2].

The main objection that Noel and Serino raise to our perspective seems to arise from a slight misunderstanding of what 'value' signifies in reinforcement learning (and therefore of what we meant by 'action relevance'). They state that 'given enough time, values . . . will exist for all of space and time', and on this basis they conclude that our argument is incomplete because PPS neurons do not directly encode all possible values. However, this objection is inaccurate because we do not claim that values will *exist* (in the brain) for all of space and time, but instead that they *can be calculated* for all of space and time. Values are intrinsically instantaneous in that they are only calculated when a particular state is visited. Specifically, action values are the *discounted cumulative expectation of future rewards due to a particular action in a particular state, given a particular assumption on what future states will be visited* (Box 1). These values are therefore conditional on an assumption of how the future will unfold, and of what states, actions, and rewards will be available to an agent. However, these assumptions and future values are not stored explicitly but are instead encapsulated in a function which takes the current state as input and returns the current value of actions as output.

This value-output function can be a neural network, in which case the assumptions about the future are stored in the precise network configuration. The values that such a network outputs, or at least the intermediate steps necessary for calculating the final values, are the 'action relevances' we mention in our original paper (in the case of the brain, the inputs to such a value-calculating network should be state estimators, which likely include activity coming from the ventral stream, frontal areas, and limbic regions [3]). Our claim was thus that PPS-related measures reflect the *instantaneous* value of particular types of actions, and not that PPS measures explicitly reflect the value of any possible action at any given time (i.e., for any possible state): PPS measures reflect the instantaneous *output of* a function

Box 1. Reinforcement Learning in Brief

The goal of reinforcement learning is to allow an agent to optimally interact with its environment. There are four core concepts in reinforcement learning: states, actions, rewards, and values. States describe the configuration of the agent–environment system. Actions are the options available to the agent, and cause it to enter another state in the next time-step (this next state can be identical to the previous state). Each action–state pair is associated with a reward, which can be positive, negative, or zero. The objective of the agent is to maximize rewards by estimating at each time-step the value of each possible action. To estimate these action values, the agent makes assumptions about what future states can be visited (i.e., physical laws) and will be visited (i.e., the policy of the agent). Based on those (learned) assumptions, expected future rewards are summed and weighted, typically in a manner inversely proportional to time, such that predicted rewards in the near future weigh more heavily than those in the far future. Thus, when the agent finds itself in a particular state, each of its available actions is associated with one value: the discounted cumulative expectation of future rewards.

CellPress REVIEWS

rather than the *infinite array of values that the output of this function could take*. We might have contributed to this misunderstanding when claiming that a field is 'a quantity that has a magnitude for each point in space and time'. We should have clarified that the magnitude of a PPS measure can be seen as a specific sample from a field in the here and now rather than as a database containing all possible field values.

There is one further clarification we would like to make. Although all PPS measures reflect action value (at least under the perspective we propose), not all action values are reflected in PPS measures. The opinion of Noel and Serino about this issue is unclear because their title states that 'high action values occur near the body', implying that, for any type of action, action values *can only* be high when an object is near the body. However, they later specifically refer to contact creation/avoidance actions, implying that their title holds true only for this type of action. To be explicitly clear: our claim was that PPS measures reflect the value of only those actions which create or avoid contact with the body, and therefore are *in part* dependent on proximity to the body. There certainly are, however, action values which do not depend on body proximity. After all, it is undeniable that non-contact actions can be valuable, and that their value does not necessarily have anything to do with proximity: merely imagine tracking a distant cloud with your head to gather information about future storms.

Acknowledgments

We thank Richard Somervail and Marina Kilintari for their valuable input to this response. We also acknowledge the support of The Wellcome Trust (COLL JLARAXR) and the European Research Council.

[1]Neuroscience and Behaviour Laboratory, Istituto Italiano di Tecnologia (IIT), Rome, Italy

[2]Department of Neuroscience, Physiology, and Pharmacology, University College London (UCL), London, UK

*Correspondence: g.iannetti@ucl.ac.uk (G.D. Iannetti).

https://doi.org/10.1016/j.tics.2019.01.011

References

1. Noel, J.-P. and Serino, A. (2019) High action values occur near our body. *Trends Cogn. Sci.* 23, 269–270
2. Bufacchi, R.J. and Iannetti, G.D. (2018) An action field theory of peripersonal space. *Trends Cogn. Sci.* 22, 1076–1090
3. Cisek, P. and Kalaska, J.F. (2010) Neural mechanisms for interacting with a world full of action choices. *Annu. Rev. Neurosci.* 33, 269–298

Forum

A Semantic Network Cartography of the Creative Mind

Yoed N. Kenett[1,*] and Miriam Faust[2,3]

The role of semantic memory in creativity is theoretically assumed, but far from understood. In recent years, computational network science tools have been applied to investigate this role. These studies shed unique quantitative insights on the role of semantic memory structure in creativity, via measures of connectivity, distance, and structure.

What do we need to know to have creative ideas? Embedded in theories on creativity is the notion that knowledge plays a role in one's ability to generate creative ideas. The main theory relating creative thinking to semantic memory – the memory system that stores concepts and facts – is the associative theory of creativity [1]. According to this theory, creativity involves the connection of weakly related, remote concepts into novel and applicable concepts. The farther apart the concepts are, the more creative the new combination will be. For this new combination to be applicable – to make sense – a broad enough body of knowledge is required. Thus, the structure of semantic memory plays an important role in the creative process. Furthermore, this theory argues that low and high creative individuals differ in their structure of semantic memory, with high creative individuals having a structure that facilitates such a process [1]. However, this theory has been challenging to investigate due to the complexity of modeling and representing semantic memory, which would allow examination of this theory. Recently, computational methods to study knowledge and memory structure in creativity are paving the way to uniquely examine their role in the creative process [2–4] and examine the associative theory of creativity [1]. Here, we outline one such approach, based on the application of network science methodologies [5].

Network science is based on mathematical graph theory, providing quantitative methods to investigate complex systems as networks [5,6]. A network is comprised of nodes that represent the basic units of a system (semantic memory) and edges that signify the relations between them (semantic similarity). While the application of network science methodologies has become a popular approach to study brain structure and function [7], it has been used to study cognitive phenomena to a lesser extent. This is despite classic cognitive theory in language and memory being highly related to a network perspective [5,6,8]. By structuring memory as a network [5], network science can directly and quantitatively examine classic cognitive theory and the operations of cognitive processes such as those taking place during memory retrieval and associative thought [8]. Such an approach provides powerful quantitative methods to examine the structure and dynamics of