ABSTRACTS

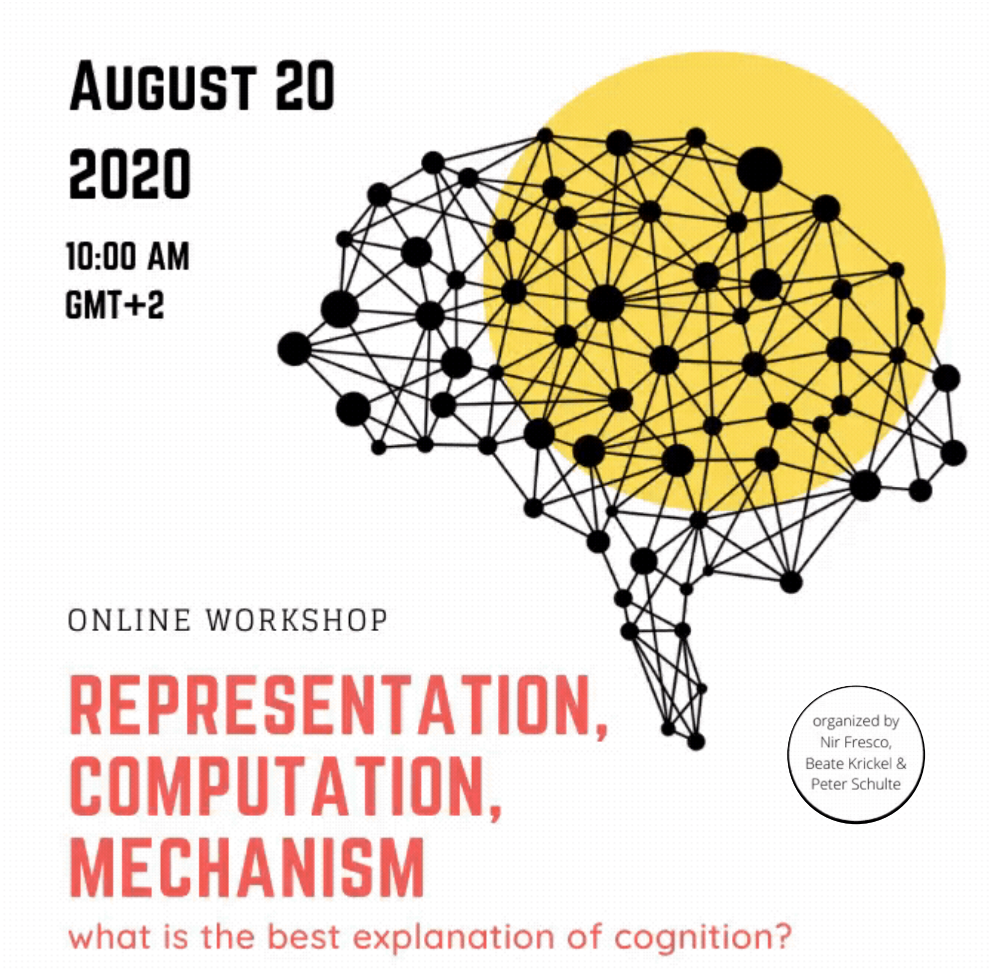

Lotem Elber (HUJI, Jerusalem)

How can we know what the brain computes when computations cannot be discovered?

It is common practice in cognitive neuroscience to identify neural representations and to relate them to hypothesized neural computations. However, it may seem that cognitive neuroscientists are facing an impossible task: they intend to identify neuronal representations and draw conclusions about what the brain computes, yet neuroscientific methods cannot tell us what the brain computes or represents. These methods allow the discovery of correlations and causal relations, but these relations, as has been extensively argued, are not sufficient to identify a specific computation or representation.

How can this impasse be overcome? Some may argue that neuroscientific talk of representations should be avoided. Others may argue that some specific correlations and causal relations are strong enough evidence to imply the existence of specific neural representations. I suggest a different route. Many agree that representation requires both correlation and function. I suggest that in cognitive neuroscience many general cognitive functions are simply posited as part of the scientific question – object recognition, motor control, planning, etc. Given this assumption, it is possible to derive the sub-functions of some brain activities. When this sub-function is to track some environmental feature, we can identify these brain areas as representing said feature. In some cases, this is easier than in other, and this may account for why some neuroscientific explanations are more convincing than other.

Marie Fellner (RU Bochum)

Cognitive neuroscience of memory: a shift from event related activity to representational formats

Cognitive Neuroscience and specifically research on human memory has shifted from a focus on neural correlates of cognitive processes to an emphasize of representations. Recent research has shifted from analyzing task specific activity changes to tracking stimulus specific activation patterns and representational structures. In this talk I will outline this shift focusing on research in episodic memory and specifically on voluntary directed forgetting. I will argue that understanding the format of mental representations is the core research question in cognitive neuroscience and fundamental to understand neural information processing.

Nir Fresco (BGU, Beer Sheva)

How context can determine the identity of physical computation

Computational explanations in the cognitive sciences span multiple levels of analysis, from a detailed biophysical model of neural activity, through an abstract algorithmic model of some cognitive phenomenon, to a mathematical specification of the function computed by some brain structure. The indeterminacy of computation complicates the endeavour of answering the question 'What does a particular neural—or physical—system do?' in computational terms. For a single physical process, taking place in whatever physical system, may often be described equally well as computing several different mathematical functions—none of which is explanatorily privileged.

At which level should the computational identity of a physical system S be determined? Dewhurst, for example, argues that the computational nature of S is wholly exhausted by S' basic physical structure (2018). Coelho-Mollo, on the other hand, argues that the computational identity of S is determined at a functional, rather than physical, level (2017). On both accounts, no contextual aspects determine the computational identity of S. Other mechanists, however, argue that contextual factors do play a role in determining S' computational identity, yet they diverge on what this role is (e.g., Piccinini 2015, Miłkowski 2017). Harbecke and Shagrir (2019) argue that contextual factors essentially determine the identity of S. I will survey some of these positions and claim that the environment can play a role in fixing the computational identity of S.

Matej Kohár (TU Berlin)

The scaling-up problem from a mechanistic point of view

The scaling-up problem for non-representational theories of cognition concerns extending an account sufficient for low-level interaction with the environment to more complex representation-hungry domains such as offline cognition and language. Those representational theories primarily aimed at explaining lower-level neural processes in perception and motor control face an analogous challenge to account for more complicated cognitive phenomena.

In this contribution, I compare the strategies for overcoming the scaling-up problem available to representationalists and to traditional non-representationalist views such as ecological psychology, dynamicism and enactivism. I show that a typical representationalist strategy involves compositionality and decouplability. Contents of basic representational vehicles are derived from teleosemantics, indicator semantics or structural correspondence. Cognitive processes then operate on these vehicles to compose and recombine them into more complex vehicles, whose content is not directly dependent on function, indication or structural correspondence. Typical non-representational theories cannot use this strategy, because they do not posit discrete internal states of the system which could be composed in the same way as representational vehicles in representationalist theories. For example, most dynamical models of cognitive phenomena posit that the cognitive system is divided into a large number of type-identical components, whose non-linear interactions are only described in aggregate by the state equations of the system. The compositional strategy cannot be employed here, because it draws much of its power from the fact that the components are different in kind and possess different causal powers. Thus, combining various components yields different results. This is not the case in dynamical models.

I will argue, however, that this worry does not apply to mechanistic non-representational models of cognitive phenomena. Mechanistic components are differentiated and organised in such a way that combining components in different ways, as well as composing mechanisms leads to predictably novel behaviours. Simple mechanisms subserving the organism's differential sensitivity to low-level perceptual stimuli can therefore be combined and organised into more complex mechanisms subserving sensitivity to higher-level properties. Likewise, mechanisms subserving a set of basic responses can be composed and organised into more complex mechanisms subserving conditional or stochastic response selection based on multiple cues and the current state of the system. Sensitivity to higher order properties and the ability to decouple responses from stimuli are hallmarks of the types of complex cognitive tasks usually discussed in connection with the scaling-up problem. I argue that this need not imply that the mechanism differentially sensitive to a stimulus is itself a representation, because such mechanism need not be decouplable and therefore fails to fulfil the functional role of a representational vehicle.

I conclude that the ability of both representational and non-representational theories of cognition to deal with the scaling-up problem depends on the degree in which they utilise the mechanistic approach to model-construction. The supposition that non-representational theories are less well equipped to deal with the scaling-up problem in virtue of lacking representational contents is false. Traditional non-representational theories are disadvantaged against representationalism because they lack mechanisms, not because they lack contents.

Beate Krickel (TU Berlin)

How to defend representation realism (or how to reject it)

Representation realism is the view that representations are real, i.e., that there are representational vehicles in the world that carry content. Different arguments have been put forward to defend representation realism (for recent attempts see Thomson & Piccinini (2018) and (2019)). At the core of all of these arguments is what I will call the Argument from Explanation: Representations provide the best explanation for cognitive phenomena. Hence, representations exist.

In this talk, I will analyze the conditions under which this inference is justified. I will show that in order for the inference from "best explanation" to "real" is justified only if a) representations add epistemic value to an explanation that cannot be had otherwise, and b) if the explanation picks out "features of the causal pattern of the world" (Craver 2007). Based on these considerations, I will argue that representational explanations need to be explanatory texts (Kohár & Krickel 2020) referring to constitutive mechanisms. In order for the Argument from Explanation to be convincing representational explanations have to provide the best texts of this format compared to non-representational explanations.

One way in which representations may enter constitutive mechanistic explanatory texts is by being constitutively relevant for the cognitive phenomenon. Kohár (dissertation) shows that this is not the case. However, representations may enter constitutive mechanistic explanatory texts by providing the best descriptions of the mechanistic components, and thereby add epistemic value that is missing in non-representational explanations. Following the literature on inference to the best explanation (Schurz 2016, Lipton 2003), I will discuss and analyze different criteria and thereby provide a to-do list for defenders of representational explanations and defenders of non-representational explanations alike.

Peter Schulte (Zurich)

Representational Explanations Defended (On All Counts)

In recent years, critics of representationalism have often questioned whether representational characterizations of cognitive processes have genuine explanatory force, i.e. whether these characterizations can be said to pick out explanatorily relevant properties. Starting from a teleosemantic account of representation, I will argue that we can secure an explanatory role for representational properties in a rather straightforward way, once we consider the right explanandum – namely, 'embedded' (world-involving) capacities of organisms. This, I contend, suffices to counter the objection of explanatory irrelevance against representationalism. However, it leaves another question wide open: do purely representational models, i.e. models that abstract away from neurophysiological details, enjoy some kind of 'explanatory autonomy'? Or are they (from an explanatory standpoint) unequivocally inferior to complete models that fill in these details? I will suggest, based on considerations of proportionality, that a version of the autonomy thesis can be successfully defended.