A Neural Network Model of Event Representations: Sensorimotor Sequencing, Place Coding, Self-organization, and Bayesian Inference

Martin Takac, Alistair Knott
Comenius University, Slovakia

Abstract:
In our talk we will present a model of how events are stored in working memory (WM).

Firstly, the model assumes a particular account of how episodes are experienced in the sensorimotor (SM) system. The key idea is that the process of experiencing an episode has a canonical sequential structure: the observer attends first to the agent, then to the patient (if there is one), and then classifies the action.

Assuming sequential attention to the agent and patient allows us to posit a WM medium that holds features of 'the currently attended individual', in particular its location and its identity (as a type or a token). This 'WM individual' medium is shown on the left of the figure below, along with a medium whose units hold short-term associations between location and type. The WM individual medium is occupied first by the agent, and then by the patient.

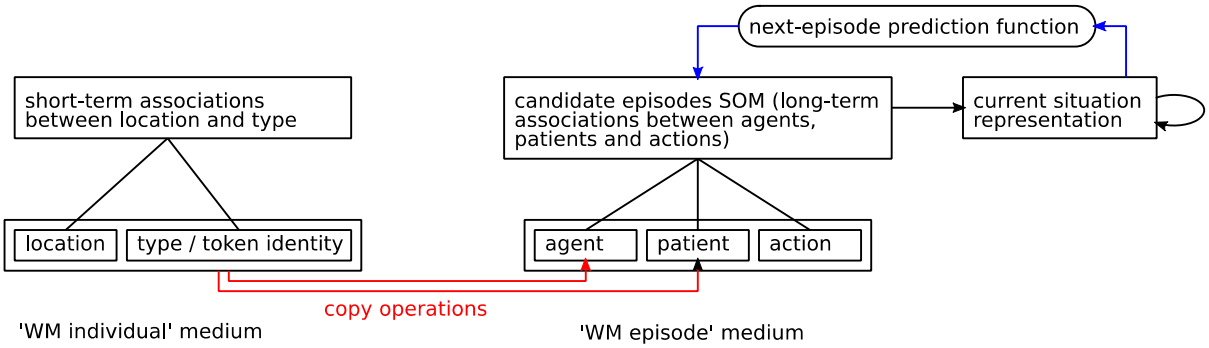

This sequential structure also allows an interesting method for encoding episodes. We assume a 'WM episode' medium, with separate fields for agent, patient, and action. Crucially, the agent and patient fields hold copies of representations in the WM individual medium - in particular, representations of type and token identity. Models where representations of the agent and patient are coded 'by place', in separate media, are normally prone to obvious problems: there is nothing in common between the representation of 'John as agent' and 'John as patient'. But in our case, where the place-coded representations are just pointers into a single WM individual medium, these problems do not arise.

The place-coded representations in the WM episode medium have several advantages. In particular, they allow the learning of localist representations of whole episodes, in units encoding associations between 'agent', 'patient', and 'action' representations. These are held in a self-organizing map (SOM; Kohonen, 1982), which learns to represent the most commonly occurring episodes. Since its inputs represent participants as both object tokens and object types, SOM units can come to represent a mixture of token episodes and episode types - another interesting advantage.

Episodes arrive sequentially in the WM system. The boundaries between events have recently been the subject of much interesting research (see Radvansky and Zacks, 2014 for a review). Our model features another SOM – this time a recurrent one - that is updated by each episode as it arrives. Units in this SOM come to represent the situations that are most frequently encountered, that is, the most commonly experienced sequences of episodes. A final interesting feature of our architecture is that we can train a function to predict the next episode in the candidate episodes' buffer. Since this function learns to predict localist episode representations, after training it predicts a distribution over possible episodes – a very rich and useful structure.