

Symposium in Memory of Jeff Elman: Language Learning, Prediction, and Temporal Dynamics

James L. McClelland (jlmcc@stanford.edu)

Department of Psychology, Stanford University
Stanford, CA 94305 USA

Arielle Borovsky (aborovsky@purdue.edu)

Speech, Language, & Hearing Sciences
Purdue University
West Lafayette, IN 47907 USA

James S. Magnuson

(james.magnuson@uconn.edu)

Department of Psychological Sciences, University of
Connecticut, Storrs, CT 06269-1020, USA

Ken McRae (kenm@uwo.ca)

Department of Psychology, University of Western
Ontario, London, ON N6A 5C2 Canada

Gina R. Kuperberg

(gkuperberg@mgh.harvard.edu)

Department of Psychology, Tufts University,
Medford, MA, 02155 USA

Felix Hill (felixhill@google.com)

DeepMind, 6 Pancras Square
London, N1C 4AG UK

Keywords: Jeff Elman, simple recurrent networks, prediction, TRACE, language development, event knowledge

Introduction

Jeffrey Locke Elman (1948-2018) devoted his career to studying human language. He investigated how people use language flexibly and productively, and how these abilities are learned from linguistic and other input. Jeff was a faculty member at the University of California, San Diego from 1977 until he passed away in 2018. His early research focused on phonetics and phonology. This work began his theoretical journey that resulted in the ideas for which he is best known: seemingly discrete combinatorial units of language, such as phonemes, may best be understood as emergent properties of underlying continuous multidimensional representations, such as phonetic input.

In the early 1980's, Jeff was a key part of transformative developments at UCSD in connectionist modeling, working with Jay McClelland, Dave Rumelhart, Geoff Hinton, and Liz Bates. With McClelland, he developed the TRACE model of speech perception. In TRACE, speech perception is seen as a constraint satisfaction process in which prior and subsequent context combine with incoming sensory evidence to determine how humans perceive speech.

Jeff then turned his attention to a central but often neglected aspect of cognition: time. Jeff's work on Simple Recurrent Networks, beginning with his classic 1990 article *Finding Structure in Time*, proposed that time-evolving continuous hidden-state representations are fundamental to language processing, and enable prediction-based learning of language. This work remains among the most influential in the history of Cognitive Science. Jeff's subsequent work explored the initial conditions under which a simple recurrent network would recover grammatical structure. He then led a collaborative project to rethink the nature of what must be built in as a foundation for language, and more generally for cognition (Elman et al., 1996). In later work,

he focused on the relationship between language and event knowledge. He argued that words do not have meanings, but instead provide clues that a listener uses to understand language. He also focused on event knowledge as a basis for prediction during language comprehension (Elman, 2009; Metusalem et al., 2012). Jeff's final major contribution was a model of how event knowledge is learned. He argued that knowledge of the components and temporal structure of events emerges as a consequence of prediction-based learning (Elman & McRae, 2019).

Jeff also played a major role in advancing Cognitive Science as a field. At UCSD, he and colleagues co-founded the interdisciplinary Center for Research in Language in 1985. In 1986, Jeff was a major part of the first Cognitive Science department, which he chaired from 1995 to 1998. Jeff also served as Dean of Social Sciences, and a founder of both the Kavli Institute for Mind and Brain and the Halicioğlu Data Sciences Institute. Finally, Jeff provided guidance for the field by serving as President of the Cognitive Science Society, and a highly respected Chair of the NIH Language and Communication study section.

This symposium honors Jeff's memory. The introduction and discussion will be led by the organizers (McClelland & McRae). In between, four speakers whose work reflects the legacy of Jeff's contributions will present research from the perspectives of cognitive neuroscience, cognition and perception, language development, computational modeling, and deep learning in simulated embodied agents.

Talks

Gina R. Kuperberg

Language prediction over time and space: Evidence from multimodal neuroimaging studies

In his seminal paper, *Finding structure in Time*, Elman argued that predictions are based not just on input from the world, but on the ever-changing state of the cognitive system. He emphasized the idea that these predictions are

non-deterministic, implicit, and inevitable. He also pointed out that prediction error not only provides feedback to the system (to learn and maximize its performance), but that it also provides valuable clues for the scientist: it can tell us about the structure of the input and the nature of cognition. These ideas have far-reaching implications for thinking about what neural measures can tell us about the architecture of language comprehension. I will discuss evidence from multimodal neuroimaging studies (ERP, fMRI and MEG) that, during comprehension, spatiotemporally distinct neural signatures reflect neural prediction error and updating at multiple time scales. I will argue that they point to a language comprehension system in which probabilistic predictions are generated and incrementally updated over time, at multiple levels and grains of representation, with the ultimate goal of inferring the latent cause that best explains the full set of inputs encountered — the message that the communicator intended to convey. Consistent with Elman's ideas, I also will argue that the neural responses evoked by prediction violations play a crucial role in triggering us to rapidly adapt to the statistical structure of our ever-changing communicative environments so that we can predict more efficiently in the future.

Arielle Borovsky

Prediction in a changing world

This talk connects with several of Elman's contributions, including his perspective on prediction, learning over time, and event knowledge in language learning and processing. Numerous language processing models emphasize the importance of listeners' ability to predict upcoming information for efficient language comprehension and learning. Much of the evidence for these models is derived from studies of comprehension in well-known or familiar (i.e. predictable) contexts. However, speakers are pressed to prioritize novel information, suggesting that everyday conversation does not typically rehash redundant events. In developmental and learning contexts, this problem may be compounded by the fact that listeners may still be learning about the language and the world. Therefore, they may not have sufficient knowledge to generate predictions. In all of these circumstances, prediction might be counter-productive for comprehension. I will discuss recent studies of how adults and children engage in prediction while learning about new events. The findings illustrate that while adult listeners can rapidly modify their predictions in the face of change, children develop this flexibility gradually over a protracted period. By incorporating developmental insights and learning paradigms into studies of linguistic prediction, we can develop richer models of how predictive mechanisms support everyday communication and learning.

James S. Magnuson

Elman's agenda for the cognitive science of language processing

I will review the remarkable breadth and depth of one of Elman's major contributions: the TRACE model of speech

perception and spoken word recognition (McClelland & Elman, 1986). I then will apply one of his other major contributions – the simple recurrent network (SRN; Elman, 1990) – to the same domain. Remarkably, SRNs have not been applied deeply to problems in spoken word recognition. Even more remarkably, despite seemingly large differences in architecture, TRACE and SRNs make extremely similar predictions, including item-specific predictions for large sets of items. I will conclude by considering how deeply Elman's ideas and work have shaped the cognitive science of language processing.

Felix Hill

Embodied neural network agents that learn language in a simulated world

I will describe a neural network 'agent' that is situated in a fully-navigable simulated 3D world as a model of early child word learning. The agent perceives its world via first-person continuous raw visual input and must learn to respond, with appropriate sequences of fine-grained motor actions, to symbolic language-like stimuli that describe simple goals. Recurrent components inspired by Jeff Elman's work play an important part in this architecture both for processing language word-by-word and for making sense of experience timestep-by-timestep. I explain how, under certain training conditions, the agent learns to reflect some known aspects of human word learning, including the emergence of semantic classes, vocabulary spurts, curriculum effects and word-learning biases. I further demonstrate how word learning can be sped up by incorporating an offline experience replay mechanism. Finally, I discuss the strengths and weaknesses of modelling early word learning with deep reinforcement learning agents in this way.

References

- Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, 14, 179-211.
- Elman, J. L. (2009). On the meaning of words and dinosaur bones: Lexical knowledge without a lexicon. *Cognitive Science*, 33, 547-582.
- Elman, J. L., Bates, E., Johnson, M. H., Karmiloff-Smith, A., Parisi, D. & Plunkett, K. (1996). *Rethinking Innateness*. Cambridge, MA: MIT Press.
- Elman, J. L., & McRae, K. (2019). A model of event knowledge. *Psychological Review*, 126, 252-291.
- McClelland, J. L., & Elman, J. L. (1986). The TRACE model of speech perception. *Cognitive Psychology*, 18, 1-86.
- Metusalem, R., Kutas, M., Urbach, T. P., Hare, M., McRae, K., & Elman, J. L. (2012). Generalized event knowledge activation during on-line language comprehension. *Journal of Memory & Language*, 66, 545-567.