SPAUN

Adam Marblestone
Cognitive Integration Class

Outline

- Neural Engineering Framework (NEF)
  - Compiling down functions onto connectivities among spiking neuron populations
    - Via an optimization procedure applied to each functional block in the system, as opposed to "learning"
- Semantic Pointer Architecture (SPA)
  - Vector symbolic architectures
- Overall SPAUN cognitive architecture

Neural Engineering Framework

spiking neural populations
--> represent vectors, x, of a specific length
e.g., 500-dimensional vector, several thousand neurons represent it

basis vectors, e, randomly assigned to neurons

nonlinear encoding of vector onto firing rates:

$$J = \alpha \boldsymbol{e} \cdot \boldsymbol{x} + J_{bias}$$

Neural Engineering Framework

spiking neural populations
--> represent vectors, x, of a specific length
e.g., 500-dimensional vector, several thousand neurons represent it

basis vectors, e, randomly assigned to neurons

nonlinear encoding of vector onto firing rates:

(input *current* to neuron) $$J = \alpha\, \boldsymbol{e} \cdot \boldsymbol{x} + J_{bias}$$

→ neuron model

→ spikes

Neural Engineering Framework

spiking neural populations
--> represent vectors, x, of a specific length
e.g., 500-dimensional vector, several thousand neurons represent it

basis vectors, e, randomly assigned to neurons

linear decoding by a readout population, via connection weights:

We can also perform the opposite operation: given a sequence of spikes we can estimate the original vector. As shown elsewhere (Eliasmith & Anderson, 2003), this can be done by deriving the decoding vectors $\boldsymbol{d}$ as per Equation 2, where a_i is the average firing rate for neuron i with a given vector $\boldsymbol{x}$, and the integration is over all values of $\boldsymbol{x}$.

$$\mathbf{d}=\Gamma^{-1}\Upsilon \quad \Gamma_{ij}=\int a_i a_j dx \quad \Upsilon_j=\int a_j \mathbf{x}\, dx \quad (2)$$

Symbolic Reasoning in Spiking Neurons:
A Model of the Cortex/Basal Ganglia/Thalamus Loop

Terrence C. Stewart (tcstewar@uwaterloo.ca)
Xuan Choo (fchoo@uwaterloo.ca)
Chris Eliasmith (celiasmith@uwaterloo.ca)
Centre for Theoretical Neuroscience, University of Waterloo
Waterloo, ON, N2L 3G1

Neural Engineering Framework

spiking neural populations

--> represent vectors, x, of a specific length

e.g., 500-dimensional vector, several thousand neurons represent it

basis vectors, e, randomly assigned to neurons

linear decoding by a readout population, via connection weights:

$$\hat{\boldsymbol{x}}(t)=\sum_{i,n}\delta(t-t_{i,n})*h_i(t)\boldsymbol{d}_i=\sum_{i,n}h(t-t_{i,n})\boldsymbol{d}_i \quad (3)$$

Symbolic Reasoning in Spiking Neurons: A Model of the Cortex/Basal Ganglia/Thalamus Loop

Terrence C. Stewart (tcstewar@uwaterloo.ca)
Xuan Choo (fchoo@uwaterloo.ca)
Chris Eliasmith (celiasmith@uwaterloo.ca)
Centre for Theoretical Neuroscience, University of Waterloo
Waterloo, ON, N2L 3G1

Neural Engineering Framework

spiking neural populations

--> represent vectors, x, of a specific length

e.g., 500-dimensional vector, several thousand neurons represent it

you can also decode using connection weights *with a transformation*

that transformation can approximate any function

optimize the weights to minimize error in that approximation

reminiscent of kernel methods:

nonlinearly map into high-dim space, then apply linear operators

Symbolic Reasoning in Spiking Neurons:
A Model of the Cortex/Basal Ganglia/Thalamus Loop

Terrence C. Stewart (tcstewar@uwaterloo.ca)
Xuan Choo (fchoo@uwaterloo.ca)
Chris Eliasmith (celiasmith@uwaterloo.ca)
Centre for Theoretical Neuroscience, University of Waterloo
Waterloo, ON, N2L 3G1

Neural Engineering Framework

So...

1) describe your system in terms of high-dimensional vectors
2) describe system functions as transformation on those vectors
3) compile onto neurons by randomly choosing input-current basis vectors, then optimizing connection weights

Symbolic Reasoning in Spiking Neurons:
A Model of the Cortex/Basal Ganglia/Thalamus Loop

Terrence C. Stewart (tcstewar@uwaterloo.ca)
Xuan Choo (fchoo@uwaterloo.ca)
Chris Eliasmith (celiasmith@uwaterloo.ca)
Centre for Theoretical Neuroscience, University of Waterloo
Waterloo, ON, N2L 3G1

Neural Engineering Framework

So...

1) describe your system in terms of high-dimensional vectors
2) describe system functions as transformation on those vectors
3) compile onto neurons by randomly choosing input-current basis vectors, then optimizing connection weights

Symbolic Reasoning in Spiking Neurons:
A Model of the Cortex/Basal Ganglia/Thalamus Loop

Terrence C. Stewart (tcstewar@uwaterloo.ca)
Xuan Choo (fchoo@uwaterloo.ca)
Chris Eliasmith (celiasmith@uwaterloo.ca)
Centre for Theoretical Neuroscience, University of Waterloo
Waterloo, ON, N2L 3G1

Neural Engineering Framework

symbol binding / compositionality:

circular convolution operation on vectors

blue⊗circle + red⊗square

Symbolic Reasoning in Spiking Neurons:
A Model of the Cortex/Basal Ganglia/Thalamus Loop

Terrence C. Stewart (tcstewar@uwaterloo.ca)
Xuan Choo (fchoo@uwaterloo.ca)
Chris Eliasmith (celiasmith@uwaterloo.ca)
Centre for Theoretical Neuroscience, University of Waterloo
Waterloo, ON, N2L 3G1

Neural Engineering Framework

symbol binding / compositionality:

circular convolution operation on vectors

$$(\mathbf{blue} \otimes \mathbf{circle} + \mathbf{red} \otimes \mathbf{square}) \otimes \mathbf{red}^{*}$$

$$= \mathbf{blue} \otimes \mathbf{circle} \otimes \mathbf{red}^{*} + \mathbf{red} \otimes \mathbf{square} \otimes \mathbf{red}^{*}$$

$$\approx \mathbf{blue} \otimes \mathbf{circle} \otimes \mathbf{red}^{*} + \mathbf{square}$$

Symbolic Reasoning in Spiking Neurons:
A Model of the Cortex/Basal Ganglia/Thalamus Loop

Terrence C. Stewart (tcstewar@uwaterloo.ca)
Xuan Choo (fchoo@uwaterloo.ca)
Chris Eliasmith (celiasmith@uwaterloo.ca)
Centre for Theoretical Neuroscience, University of Waterloo
Waterloo, ON, N2L 3G1

Neural Engineering Framework

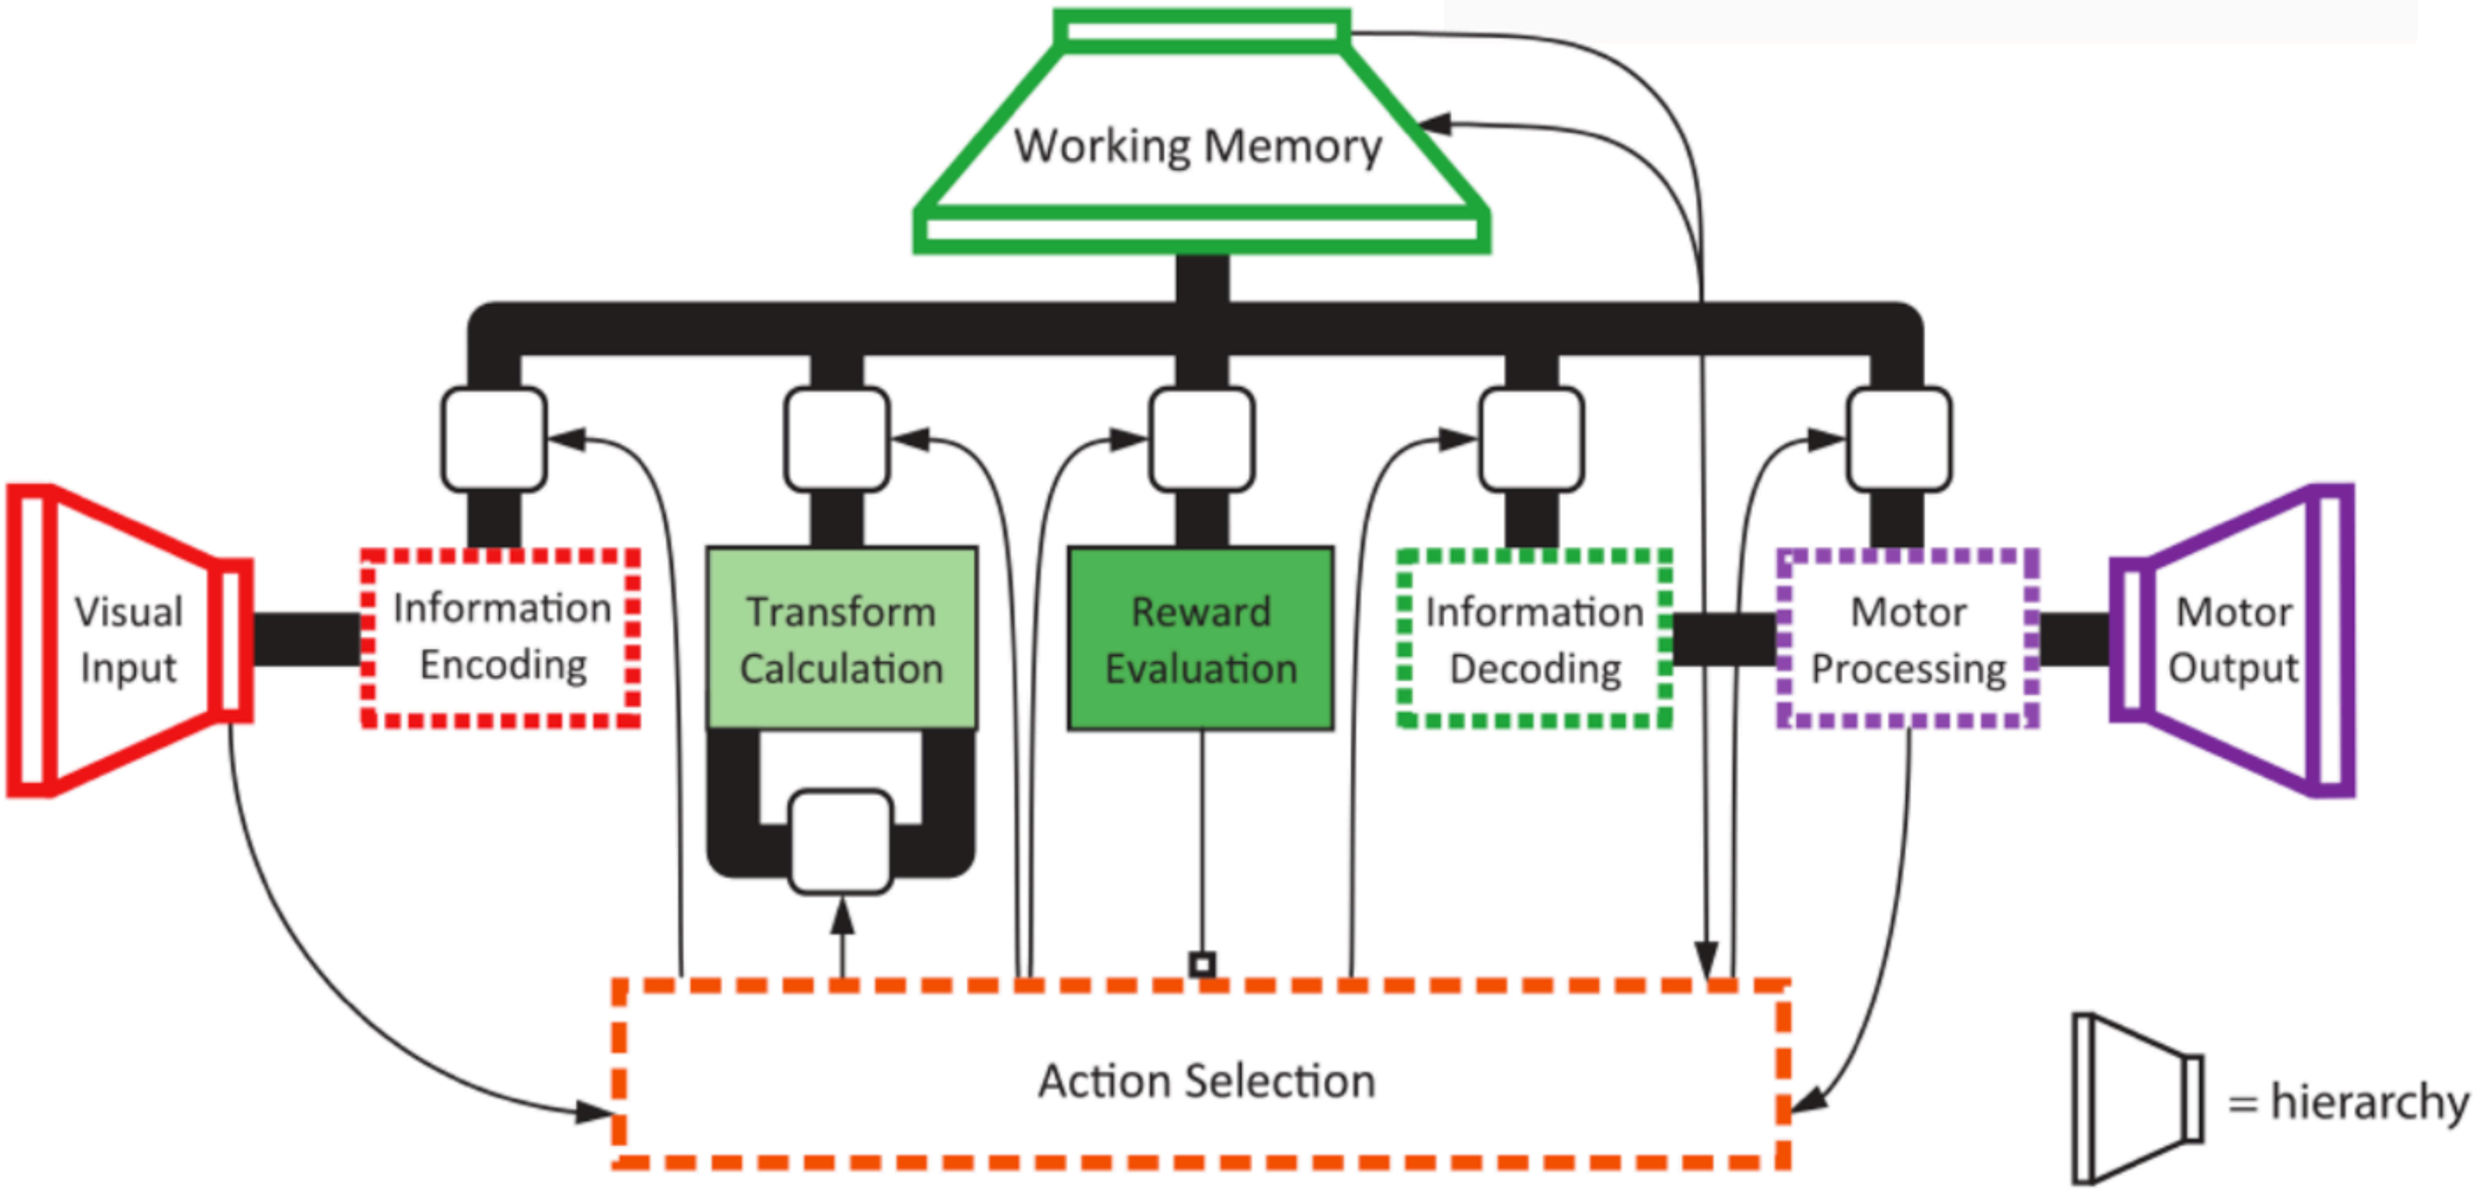

A Large-Scale Model of the Functioning Brain

Chris Eliasmith*, Terrence C. Stewart, Xuan Choo, Trevor Bekolay, Travis DeWolf, Yichuan Tang, Daniel Rasmussen

+ Author Affiliations

*To whom correspondence should be addressed. E-mail: celiasmith@uwaterloo.ca

Science 30 Nov 2012:
Vol. 338, Issue 6111, pp. 1202-1205
DOI: 10.1126/science.1225266

Neural Engineering Framework

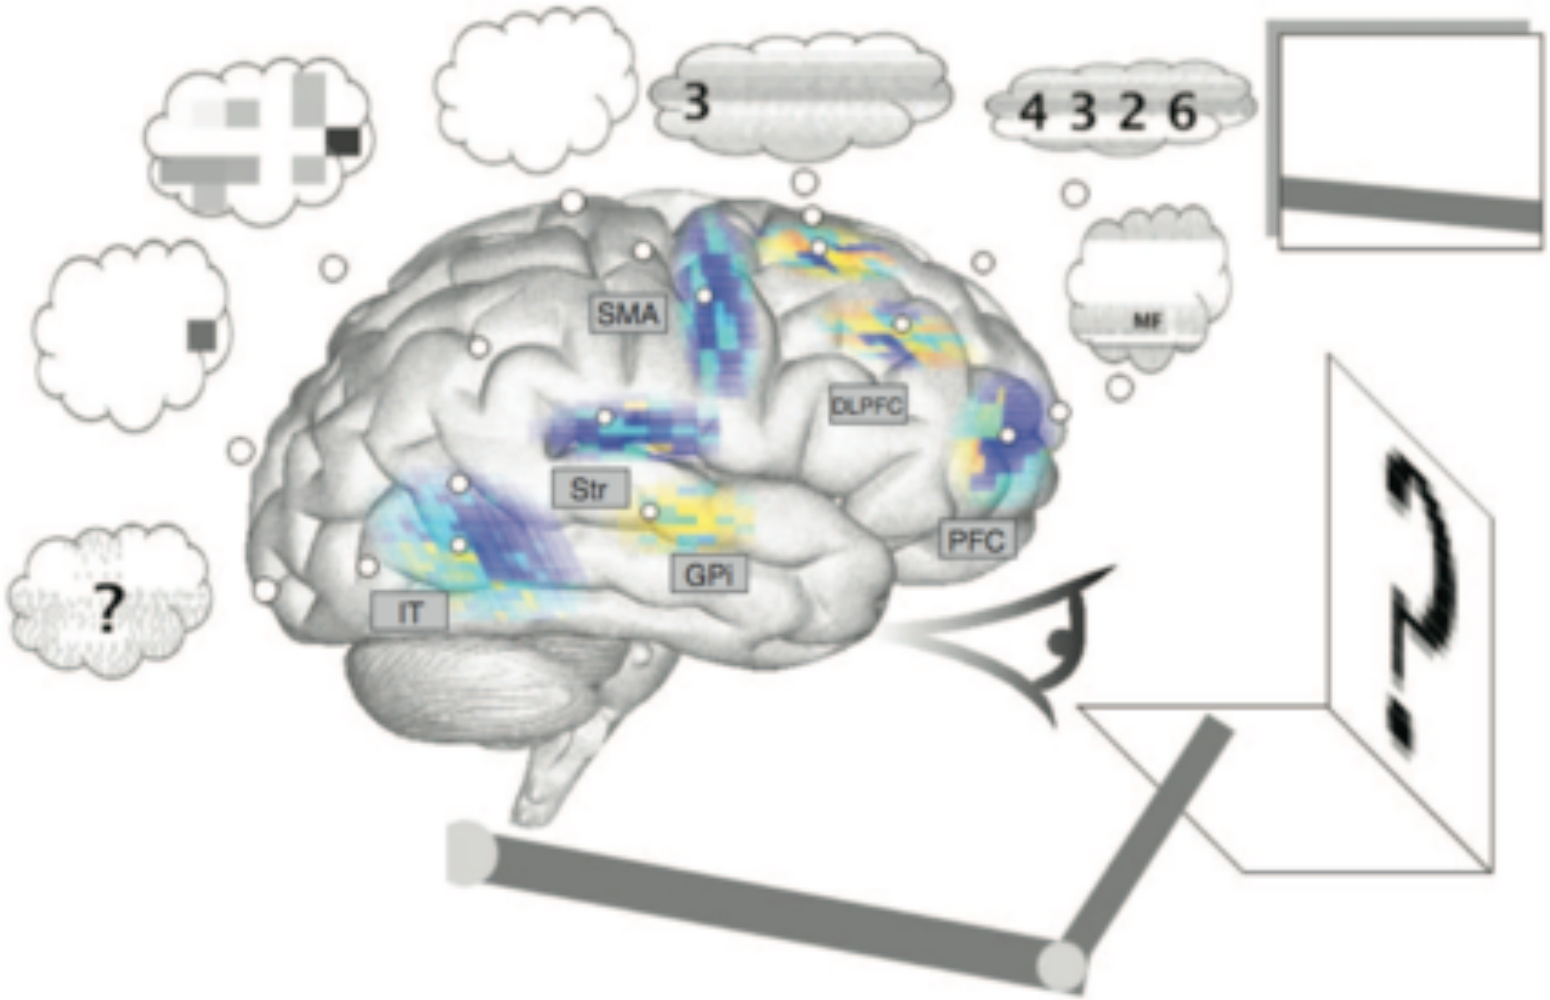

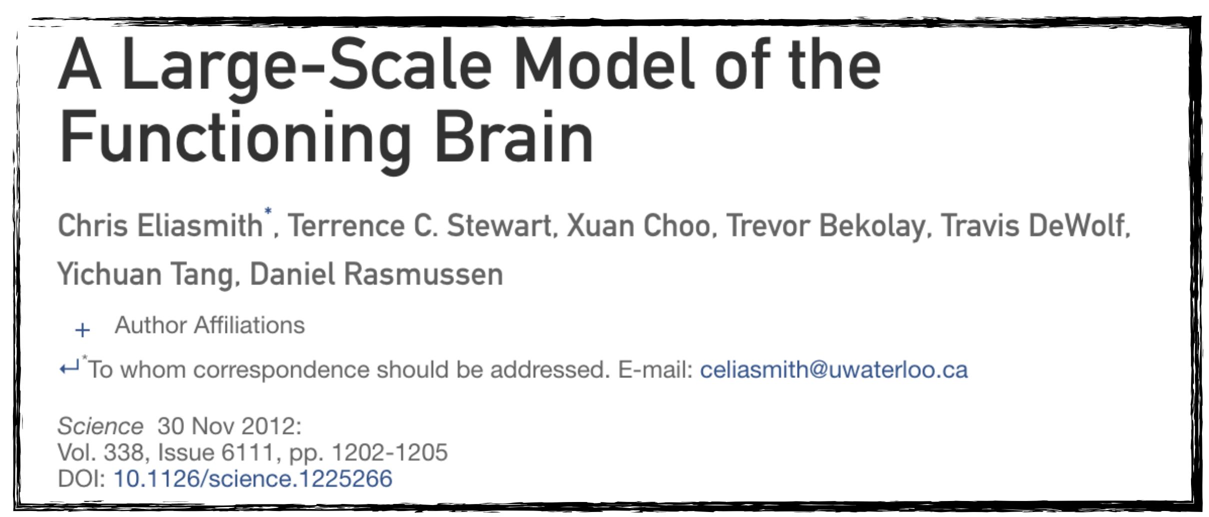
A Large-Scale Model of the Functioning Brain

Chris Eliasmith*, Terrence C. Stewart, Xuan Choo, Trevor Bekolay, Travis DeWolf, Yichuan Tang, Daniel Rasmussen

+ Author Affiliations

*To whom correspondence should be addressed. E-mail: celiasmith@uwaterloo.ca

Science 30 Nov 2012:
Vol. 338, Issue 6111, pp. 1202-1205
DOI: 10.1126/science.1225266

Neural Engineering Framework

Picking representations and block decomposition and locally optimizing connections

vs.

End-to-end optimization / learning

Good picture of a potential end-state of learning?

A Large-Scale Model of the Functioning Brain

Chris Eliasmith*, Terrence C. Stewart, Xuan Choo, Trevor Bekolay, Travis DeWolf, Yichuan Tang, Daniel Rasmussen

+ Author Affiliations

*To whom correspondence should be addressed. E-mail: celiasmith@uwaterloo.ca

Science 30 Nov 2012:
Vol. 338, Issue 6111, pp. 1202-1205
DOI: 10.1126/science.1225266