DFT Summer School: Lecture 5

John P. Spencer
Professor
School of Psychology
University of East Anglia
Norwich, UK

Where we've been thus far…

- How does a population of neurons form a stable pattern of activation to create a local decision about, say, the presence/absence of a stimulus?
  - Node dynamics; DST; attractors; relaxation time; bifurcations; instabilities; self-stabilized state; self-sustaining state
- How does a neural population also reflect the meaning of a stimulus or response—we need metric spaces
  - Dynamic fields; peaks; different forms of neural interaction; convolution; the impact of noise; inputs; tracking; peak drift; peak interactions
- What are the neural principles that link DFT and neurophysiology; put differently, how neural is DFT?
  - DPA; evidence for neural interactions; functional topography; receptive fields; adaptive nature of broad receptive fields; evidence for population coding
- What's the link between fields and motor behaviors? How do we get a field to actually behave in the world?
  - Robotics case study; turning peaks into behavioral dynamics; attractors and repellors along behavioral dimensions like heading direction

We've come quite far!

- Pat yourselves on the back!

University of East Anglia

The next arc: higher-level cognition

Although we've come quite far, still a long way to go…Thus far, we've only handled very simple situations

- Detecting the presence/absence of a stimulus and maintaining a working memory of that perceptual decision
- Detecting/forming a memory for a color or the location of a stimulus
- Generating behavioral responses (button presses, changes in heading direction) in response to a simple task

- At present, we can't even encode and form a memory for a colored, oriented bar…For that, need to move into higher dimensions.

University of East Anglia

Multidimensional fields

How do we represent multi-feature objects?

- Typical approach in cognition: multi-dimensional space → one dimension for each feature
  - Color
  - Orientation
  - Shape
  - Material
  - Texture
- If we want to sample each dimension with 100 neurons, the full representation of this 5D space neurally would require 100^5 (10 billion) neurons [vs. 500 for 5 1D fields]
  - there are approximately 100 billion neurons in a brain; so this multidimensional approach would quickly exhaust the computational power of the brain
- The moral: Multidimensional representations are costly!

University of East Anglia

Multidimensional fields

Are there multidimensional fields in the brain?

- Yes!
- Many cells in early vision are 3D…sensitive to variations in 2D spatial position and also tuned to a feature dimension (e.g., simple orientation cells)
- What does it mean for the representation to be multi-dimensional?
  - The neural system is tuned to all possible combinations of the dimensions
  - Offers powerful framework for flexibility and generalization
  - Consider two spatial dimensions as in visual cortex…

University of East Anglia

Multidimensional fields

- The mathematics behind multiD field should look familiar:

$$\tau \dot{u}(\mathbf{x}) = -u(\mathbf{x}) + h + s(\mathbf{x}) + \int_F k(\mathbf{x}-\mathbf{x}')g(u(\mathbf{x}'))d\mathbf{x}'$$

$$\tau \dot{u}(x,y) = -u(x,y) + h + s(x,y) + \iint k(x-x',y-y')g(u(x',y'))dx'dy'$$

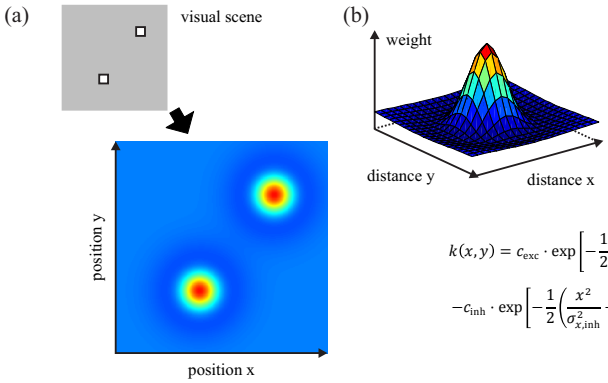

$$k(x,y) = c_{\text{exc}} \cdot \exp\left[-\frac{1}{2}\left(\frac{x^2}{\sigma^2_{x,\text{exc}}} + \frac{y^2}{\sigma^2_{y,\text{exc}}}\right)\right]$$

$$-c_{\text{inh}} \cdot \exp\left[-\frac{1}{2}\left(\frac{x^2}{\sigma^2_{x,\text{inh}}} + \frac{y^2}{\sigma^2_{y,\text{inh}}}\right)\right] - c_{\text{glob}}$$

University of East Anglia

What about different feature dimensions?

- Need to think about the metrics along each dimension
- Conceptually, the metrics specify the similarity relations along each dimension
  - Metrics along one space don't have to be similar across dimensions (two 'units' along a vertical dimension don't have an intrinsic relation to two 'units' along a color dimension)
- Mathematically, this just means that the kernels don't have to be symmetric
  - How do we set this? Could be based on behavioral data (e.g., model discrimination data for similar values along a dimension)

University of East Anglia

Combinatorial explosion: revisited

- Multidimensional fields exist in the brain
- Multidimensional fields are straightforward mathematically/dynamically
- But multidimensional fields are costly (in neural resources and computational resources)…so use them sparingly!

Reconsider the multidimensional picture from cognitive science

- We don't really need a full multidimensional space for all feature combinations…some regions of multidimensional space are actually quite sparse [purple polka dot elephants…]
- Instead of representing a big empty space, what if we had a way to effectively pull things into and out of a 'bound', high-dimensional representation…

University of East Anglia

Integration and selection in DFT

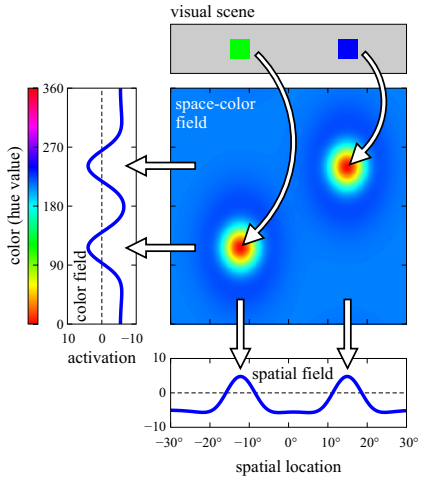

University of East Anglia

Ridge projections: 1D → 2D

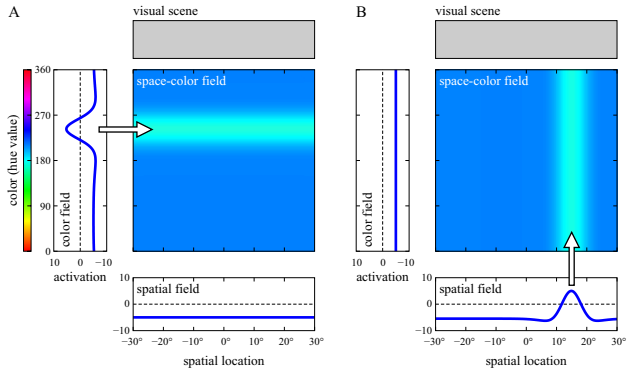

University of East Anglia

Spatial selective attention

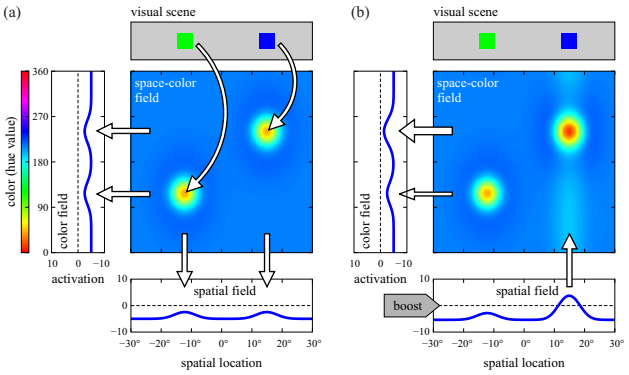

University of East Anglia

Feature attention

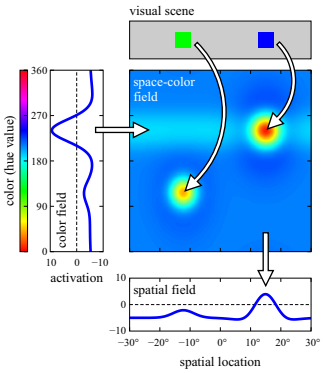

$$\tau \dot{u}_{\text{spt}}(x) = -u_{\text{spt}}(x) + h_{\text{spt}} + s_{\text{spt}}(x) + \int k_{\text{spt,spt}}(x-x')g\left(u_{\text{spt}}(x')\right)dx' + \int \left[k_{\text{spt,vis}}(x-x')\int g(u_{\text{vis}}(x',y))dy\right]dx'$$

$$\tau \dot{u}_{\text{col}}(y) = -u_{\text{col}}(y) + h_{\text{col}} + s_{\text{col}}(y) + \int k_{\text{col,col}}(y-y')g(u_{\text{col}}(y'))dy' + \int \left[k_{\text{col,vis}}(y-y')\int g(u_{\text{vis}}(x,y'))dx\right]dy'$$

$$\tau \dot{u}_{\text{vis}}(x,y) = -u_{\text{vis}}(x,y) + h_{\text{vis}} + s_{\text{vis}}(x,y) + \iint k_{\text{vis,vis}}(x-x',y-y')g(u_{\text{vis}}(x',y'))dx'dy' + \int k_{\text{vis,spt}}(x-x')g\left(u_{\text{spt}}(x)\right)dx' + \int k_{\text{vis,col}}(y-y')g(u_{\text{col}}(y))dy'$$

University of East Anglia

Multidimensional integration and selection

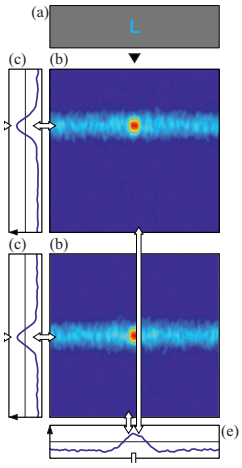

- Selective attention gives us a way to move into and out of higher-dimensional representation
- Consider a complex object: we can represent a two-feature ‘L’ (color, shape) using two, 2D representations
- The ‘bound’ object is the pattern of peaks coupled along the spatial dimension
- Could have a 5D representation of the object using 50,000 neurons instead of 10 billion

University of East Anglia

Multidimensional integration and selection

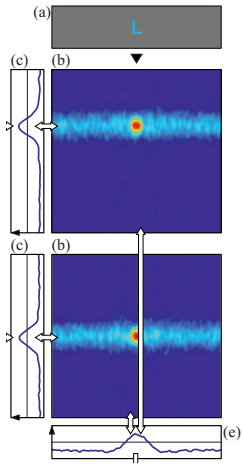

- Limitation: have to 'bind' one object at a time using selective attention
  - Can lead to errors under some conditions (illusory conjunctions)
- Also need a way to keep track of bindings: need to deal with binding in working memory (not just binding in perception)
  - Tackle one piece of the working memory story now…the rest will be in chapter 8
  - Then we'll use this bit of WM to help explain illusory conjunctions

University of East Anglia

Integration and selection with WM

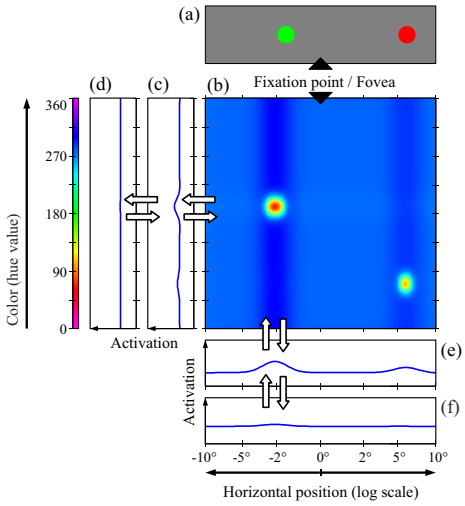

- Add WM layers
- Allows us to implement 'top-down' goals → WM can affect attentional selection

University of East Anglia

WM affects attentional selection

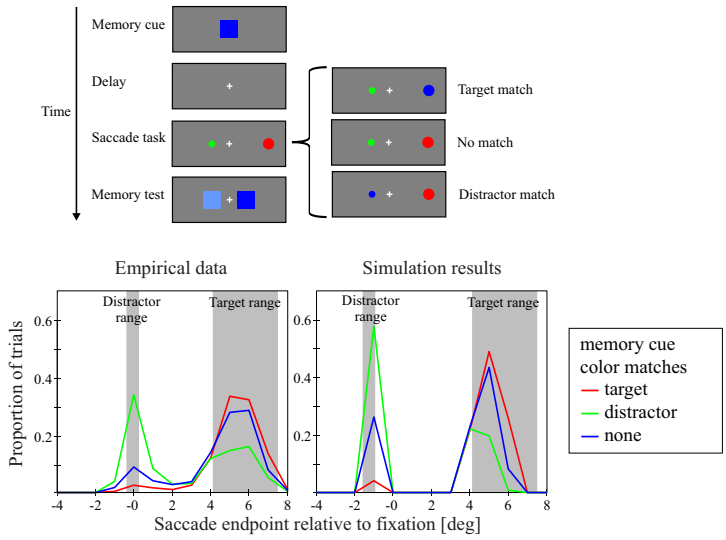

University of East Anglia

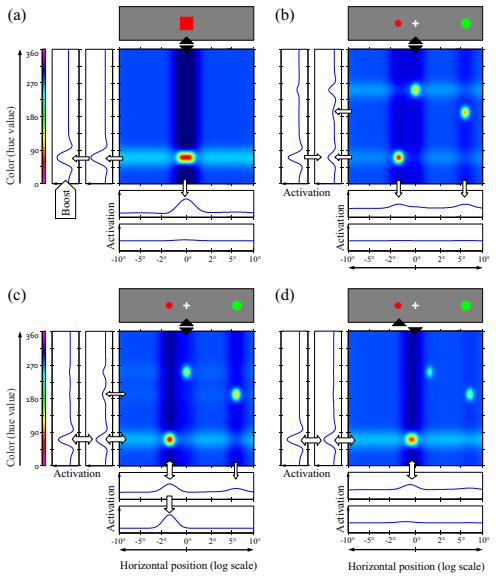

University of East Anglia

WM, Attention & Illusory Conjunctions

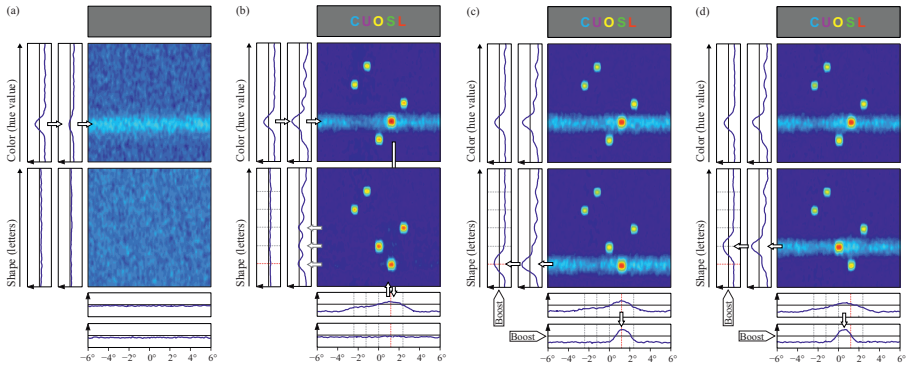

University of East Anglia

Conclusions

- Started with a sad statement: can't even represent a colored, oriented bar with DFT
- Ended with a framework that solves the combinatorial explosion problem, explains illusory conjunctions, and starts to give us insights into attentional selection and selection via 'top-down' goals in WM
- Higher cognition might actually be within reach…

University of East Anglia