

Coordinate Transformations in Parietal Cortex

Computational Models of Neural Systems

Lecture 7.1

David S. Touretzky


November, 2015

Outline


- Anderson: parietal cells represent locations of visual stimuli.
- Zipser and Anderson: a backprop network trained to do parietal-like coordinate transformations produces neurons whose responses look like parietal cells.
- Pouget and Sejnowski: the brain must transform between multiple coordinate systems to generate reaching to a visual target.
- A model of this transformation can be used to reproduce the effects of parietal lesions (hemispatial neglect).

Parietal Cortex

- Four sections of IPL (inferior parietal lobule):
 - 7a: visual, eye position
 - 7b: somatosensory, reaching
 - MST: visual motion, smooth pursuit
 - medial superior temporal area
 - 19/37/39 boundary in humans
 - V5a in monkeys
 - LIP: visual & saccade-related
 - lateral intra-parietal area



Monkey and Human Parietal Cortex





Inferior Parietal Lobule

- Posterior half of the posterior parietal cortex.
- Area 7a contains both visual and eye-position neurons.
- Non-linear interaction between retinal position and eye position.
 - Model this as a function of eye position multiplied by the retinal receptive field.
- No eye-position-independent coding in this area.

Results from Recording in Area 7a (Anderson)


- Awake, unanesthetized monkeys shown points of light
- 15% eye position only
- 21% visual stimulus only
- 57% respond to a combination of eye position and stimulus
- Most cells have spatial gain fields; mostly planar
- Approx. 80% of eye-position gain fields are planar

Spatial Gain Fields




Spatial Gain Fields of 9 Neurons


- Cells b,e,f:
 - Evoked and background activity co-vary
- Cells a,c,d:
 - Background is constant
- Cells g,h,i:
 - Evoked and background activities are non-planar, but total activity is planar



Types of Gain Fields



Neural Network Simulation



Simulation Details

- Three layer backprop net with sigmoid activation function
- Inputs: pairs of retinal position + eye position
- Desired output: stimulus position in head-centered coords.
- 25 hidden units
- ~ 1000 training patterns
- Tried two different output formats:
 - 2D Gaussian output
 - Monotonic outputs with positive and negative slopes

Hidden Unit Receptive Fields

10



20

30





40

a





No units



b




c




Random weights; no training

Real and Simulated Spatial Gain Fields

Real



Simulated



Summary of Simulation Results


- Hidden unit receptive fields sort of look like the real data.
- All total-response gain fields were planar.
 - In the real data, 80% were planar
- With monotonic output, 67% of visual response fields planar
- With Gaussian output, 13% of visual response fields planar
- Real data: 55% of visual response fields planar
- Maybe monkeys use a combination of output functions?
- Pouget & Sejnowski: sampling a sigmoid function at 9 grid points can make it appear planar. Might be a sigmoid.

Discussion

- Note that the model is not topographically organized.
- The input and output encodings were not realistic, but the hidden layer does resemble the area 7a representation.
- Where does the model's output layer exist in the brain?
 - Probably in areas receiving projections from 7a.
 - Eye-position-independent (i.e., head-centered) coordinates will probably be hard to find, and may not exist at a single cell.
 - Cells might only be independent over a certain range.
- Prism experiments lead to rapid recalibration in adult humans, so the coordinate transformation should be plastic.

Pouget & Sejnowski: Synthesizing Coordinate Systems



- The brain requires multiple coordinate systems in order to reach to a visual target.
- Does it keep them all separate?
- These coordinate systems can all be synthesized from an appropriate set of basis functions.
- Maybe that's what the brain actually represents.




Basis Functions


- Any non-linear function can be approximated by a linear combination of basis functions.
- With an infinite number of basis functions you can synthesize any function.
- But often you only need a small number.
- Pouget & Sejnowski: use the product of gaussian and sigmoid functions as basis functions.
 - Retinotopic map encoded as a gaussian
 - Eye position encoded as a sigmoid

Gaussian-Sigmoid Basis Function



Coordinate Transformation Network






Can derive either head-centered or retinotopic representations from the same set of basis functions. The model used 121 basis functions.

Summary of the Model





- Not a backprop model.
 - Input-to-hidden layer is fixed set of nonlinear basis functions
 - Output units are linear; can train with Widrow-Hoff (LMS algorithm)
- Less training required than for Zipser & Anderson, but model uses more hidden nodes.
- Assume sigmoid coding of eye position, unlike Zipser & Anderson who use a linear (planar) encoding.
 - But sigmoidal units can look planar depending on how they're measured.

Evidence for Saturation (Non-Linearity)

- Cells **B** and **C** show saturation, supporting the use of sigmoid rather than linear activation functions for eye position.




Sigmoidal Units Can Still Appear Planar




Map Representations

- Alternative to spatial gain fields idea.
- Localized “receptive fields”, but in head-centered coordinates instead of retinal coordinates.
- Not common, but some evidence in VIP (ventral intraparietal area).




Vectorial Representations

- Unit's response is the projection of stimulus vector A along the units' preferred direction: dot product.
- Units are therefore linear in a_x and a_y ; response to θ is a cosine function.
- But 20% of real parietal neurons were non-linear.
- Motor cortex appears to use this vector representation to encode reaching direction.




Hemispatial Neglect

- Caused by posterior parietal lobe lesion (typically stroke).
- Can also be induced by TMS.
- Patient can't properly integrate body position information with visual input.




Copies of a clock and a daisy


Line Bisection Task



Artist's Rendition of Left Hemisphere Neglect (Depict Impaired Attention as Loss of Resolution)




Retinotopic Neglect Modulated By Egocentric Position



Stimulus-Centered Neglect


B



Note that target **x** is in same retinal position in C1 vs. C2. Only the distractors have moved.


Pouget & Sejnowski Model of Neglect

- Parietal cortex representations are biased toward the contralateral side.
- Similar model to previous paper, but...
- Neglect simulated by biasing the basis functions to favor right-side retinotopic and eye positions, simulating a right side parietal lesion (loss of left side representation).



Selection Mechanism


- Present the model with two simultaneous stimuli, causing two hills of activity in the output layers.
- Select the most active hill as the response. Zero the activities of those units to cause the model to move on. Allow them to slowly recover.



Simulation Results



- Right side stimuli are selected and activation set to zero.
- But stimuli eventually recover and are selected again.
- Left side stimuli have poor representations and are frozen out.

A




Simulation Results

A





B



dashed line: looking straight ahead
solid line: body turned to the left

Simulation Results



Discussion

- Neglect patients show a mixture of retinotopic, head-centered, trunk-centered, and object-centered effects.
- This argues for a representation that combines multiple types of information.
 - Damage to that area could explain the mixture of effects.
- The proposed parietal basis function representation encodes information in a way that allows any desired reference frame to be extracted by a simple linear output layer.
- Tradeoff: to encode more information, the basis functions must be more complex.
 - And you need more of them.
 - And decoding becomes more complex (even if linear).