3. Neocortical Dynamics

Functions of Layers

- Input
  - layer 4
  - from sensation or other areas
- Hidden
  - layers 2 & 3
- Output
  - layers 5 & 6
  - to motor systems or other areas

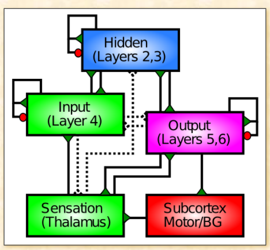

(fig. < O'Reilly, *Comp. Cog. Neurosci.*)

Connection Directions

- Feedforward
  - from Hidden in lower to Input in higher
- Feedback
  - from Hidden & Output in higher to Hidden & Output in lower
- Lateral
  - from Hidden and Output to all three layers in same area
- Bidirectionality
  - pervasive

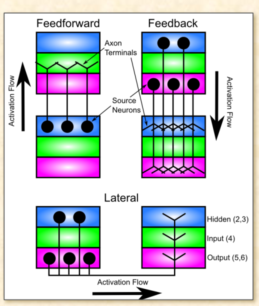

(fig. < O'Reilly, *Comp. Cog. Neurosci.*)

Hierarchical Categorical Representations

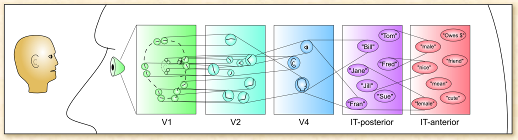

- Successive layers of neural detectors
- Progressively more abstract

4/4/17 (fig. < O'Reilly, *Comp. Cog. Neurosci.*) 4

Cell Responses in V4

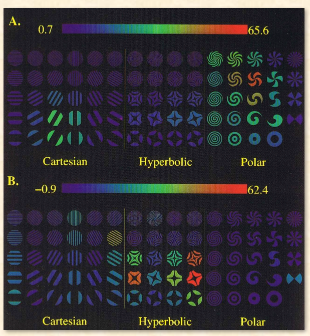

4/4/17 (fig. < Clark, *Being There*, 1997) 5

Sparse Distributed Representation

- Localist representation
  - "grandmother cells"
  - unlikely in brain
- K-out-of-N detectors
  - typically 15–25% of neurons active
- Approximate orthogonality

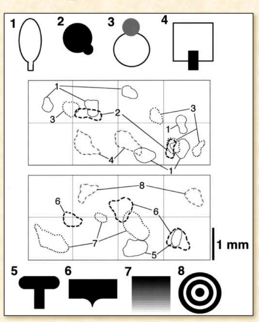

(monkey IT cortex)

4/4/17 (fig. < O'Reilly, *Comp. Cog. Neurosci.*, from Tanaka, 2003) 6

Sparse Distributed Representations

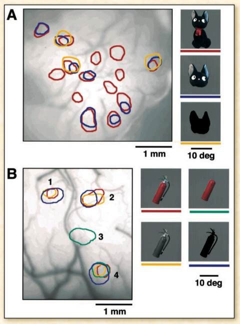

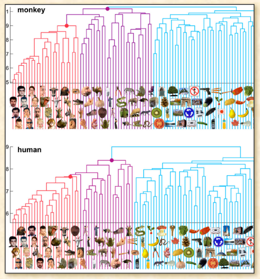

4/4/17 (figs. < O'Reilly, *Comp. Cog. Neurosci.*, from Tanaka 2003 and Kriegeskorte et al. 2008) 7

Coarse Coding

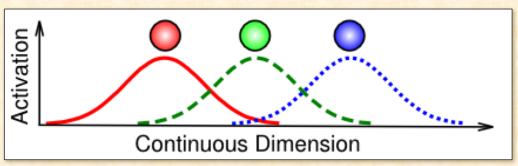

- Broadly-tuned receptive fields
- Population-coding of precise values
- Common throughout sensory and motor areas

4/4/17 (fig. < O'Reilly, *Comp. Cog. Neurosci.*) 8

Orientation Columns

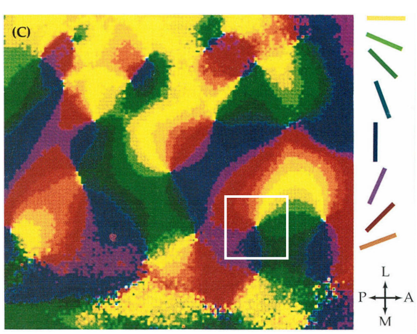

4/4/17 (fig. < Nicholls & al., *Neur. to Brain*) 9

Orientation Columns

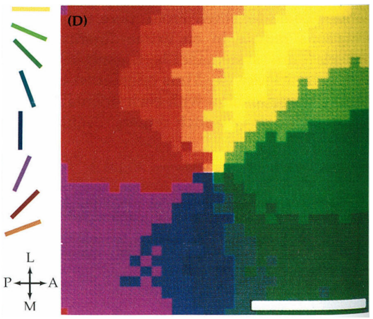

4/4/17 (fig. < Nicholls & al., *Neur. to Brain*) 10

Topographic Maps: Bat Auditory Cortex

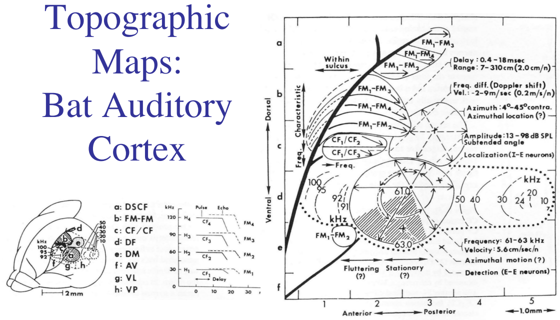

4/4/17 (figs. from Suga, 1985) 11

Bidirectional Excitation

- Functions
  - recognition
  - top-down imagery
  - ambiguity resolution
  - pattern completion
- Attractor dynamics
  - convergence on good representation
  - energy vs. harmony

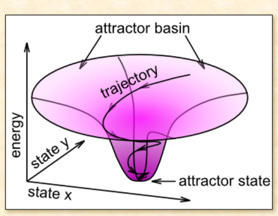

(fig. < O'Reilly, *Comp. Cog. Neurosci.*)

4/4/17 12

Ambiguity Resolution

4/4/17 13

(fig. < O'Reilly, *Comp. Cog. Neurosci.*)

Inhibitory Competition and Activity Regulation

- Activity regulation
- Selective attention
- Competition
  - K winners take all
  - can be implemented algorithmically
- Sparse distributed representation

4/4/17 14

Activity Regulation

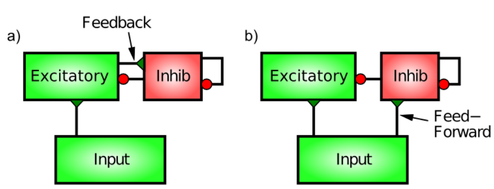

- Feedback: reactive, reflects actual level of activity, robust, responsive, may be unstable
- Feedforward: anticipatory, limits feedback oscillation, slow, brittle
- Work well together

4/4/17 15

(fig. < O'Reilly, *Comp. Cog. Neurosci.*)

4. Learning Mechanisms

4/4/17 16

Spike Timing Dependent Plasticity (STDP)

1. V_m elevated by backpropagating action potential
2. Repels Mg^+ opening NMDA channels
3. Presynaptic neuron fires, releasing glutamate
4. Glutamate binds unblocked NMDA channels, allowing Ca^{++} influx
5. Ca^{++} increases number & efficacy of AMPA receptors

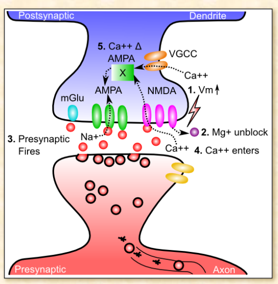

4/4/17 (fig. < O'Reilly, *Comp. Cog. Neurosci.*) 17

Long-term Potentiation (LTP) vs. Long-term Depression (LTD)

- LTP vs. LTD depends on Ca^{++} concentration over several 100 msec
- Records possible causal connection
- Actual situation is more complicated with multiple APs

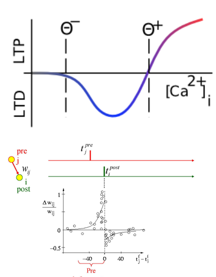

4/4/17 (figs. < O'Reilly, *Comp. Cog. Neurosci.*) 18

LTP/LTD Approximation

- Piecewise linear approximation to LTP/LTD
- Typical $\theta_d = 0.1$
- Floating threshold
- $\Delta W = \eta f_{\mathrm{XCAL}}\left(\langle xy \rangle_s, x\langle y \rangle_l\right)$

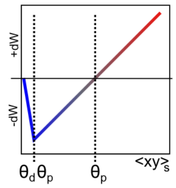

$$f_{\mathrm{XCAL}}(c, \theta_p) = \begin{cases} c - \theta_p & \text{if } c > \theta_p \theta_d \\ -c(1-\theta_p)/\theta_d & \text{otherwise} \end{cases}$$

4/4/17 (fig. < O’Reilly, *Comp. Cog. Neurosci.*) 19

Homeostatic Behavior

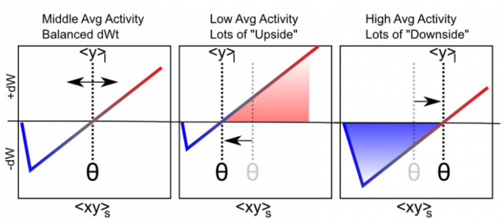

- Floating threshold adapts to long-term postsynaptic activity
- Tends to equalize activity among neurons

4/4/17 (fig. < O’Reilly, *Comp. Cog. Neurosci.*) 20

Competitive Learning

- Competitive learning network
  - two layers, randomly initialized weights
  - second is self-reinforcing, mutually inhibitory
  - “winner takes all” dynamics
- Learning
  - winner moves toward last
  - weight vectors move to centers of clusters

4/4/17 21

Competitive Network

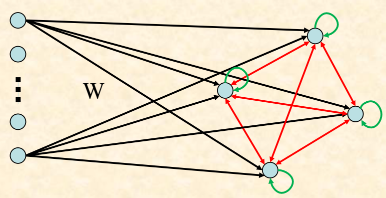

4/4/17 22

Diagram of Competitive Learning

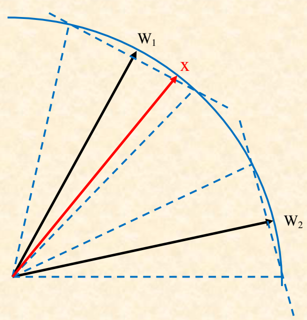

Self-Organizing Learning

- Inhibitory competition
  - ensures sparse representation
- Hebbian "rich get richer"
  - adjustment toward last pattern matched
- Slow threshold adaptation
  - adjusts receptive fields
  - equalizes cluster probabilities
- Homeostasis
  - distributes activity among neurons
  - more common patterns are more precisely represented
- Gradually develops statistical model of environment

4/4/17 24

Error-Driven Learning

- For achieving intended outcomes
- Fast threshold adaptation
- Short-term outcome – medium-term expectation
  - ✓ "plus phase" – "minus phase"
- Depends on bidirectional connections
  - ✓ communicates error signals back to earlier layers
- Contrastive Attractor Learning (CAL)
  - ✓ approximately equivalent to BP when combined with bidirectional connections

4/4/17 25

Contrastive Attractor Learning

- Network learns contrast between:
  - early phase/expectation (minus)
  - late phase/outcome (plus)
- Gets more quickly to late phase, which has integrated more constraints

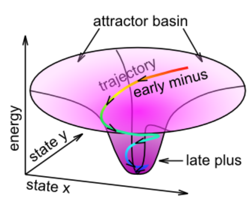

$$f_{\text{XCAL}}(c, \theta_{\text{p}}) = \begin{cases} c - \theta_{\text{p}} & \text{if } c > \theta_{\text{p}}\theta_{\text{d}} \\ -c(1-\theta_{\text{p}})/\theta_{\text{d}} & \text{otherwise} \end{cases}$$

$$\Delta W = \eta f_{\text{XCAL}}\left(\langle xy \rangle_{\text{s}}, x\langle y \rangle_{\text{m}}\right)$$

$$\approx \eta\left(x_{\text{s}} y_{\text{s}} - x_{\text{m}} y_{\text{m}}\right)$$

4/4/17 26

Learning Signals?

- What constitutes an "outcome"?
- Dopamine bursts arise from unexpected rewards or punishments (reinforcers)
  - violation of expectation
  - needs correction
- Dopamine modulates synaptic plasticity
  - controls λ: $\Delta W = \eta f_{\text{XCAL}}\left(x_{\text{s}} y_{\text{s}}, x_{\text{m}}\left(\lambda y_{\text{l}} + (1-\lambda) y_{\text{m}}\right)\right)$
- Probably not the whole story

4/4/17 27

Learning Situations

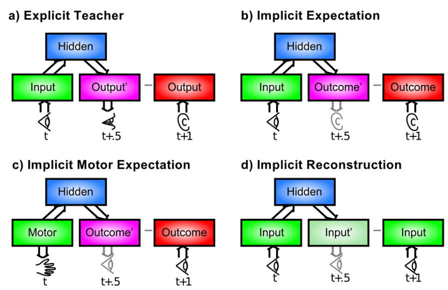

(fig. < O'Reilly, *Comp. Cog. Neurosci.*)

4/4/17 28