4. Learning Mechanisms

Overview of Learning

- Biology: synaptic plasticity

- Computation:
  - Self organized:
    - statistical regularities $\Rightarrow$ internal models
    - long time scale
  - Error-driven: getting the right answers
    - outcomes $\Rightarrow$ expectations
    - short time scale
 - Integration of two forms of learning

A. Biology of Synaptic Plasticity

Mechanisms of Synaptic Weight Change

- Amount of neurotransmitter released

- Nonspecific (extra-synaptic) neurotransmitter

- Change in number of post-synaptic receptors

- Rapid change in shape of dendritic spines

- Formation of new synapses

Synapses Change Strength
(in response to patterns of activity)

(slide < O'Reilly)

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

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

Hebb's Rule

- "When an axon of cell A is near enough to excite a cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased." — Donald Hebb

- I.e., neurons that fire together, wire together

- $\Delta W = xy$

  where x is sending activity and y is receiving activity

Hebb's Rule Learns Correlations

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

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

Causal Learning?

(slide < O'Reilly)

Let's Get Real…

(slide < O'Reilly)
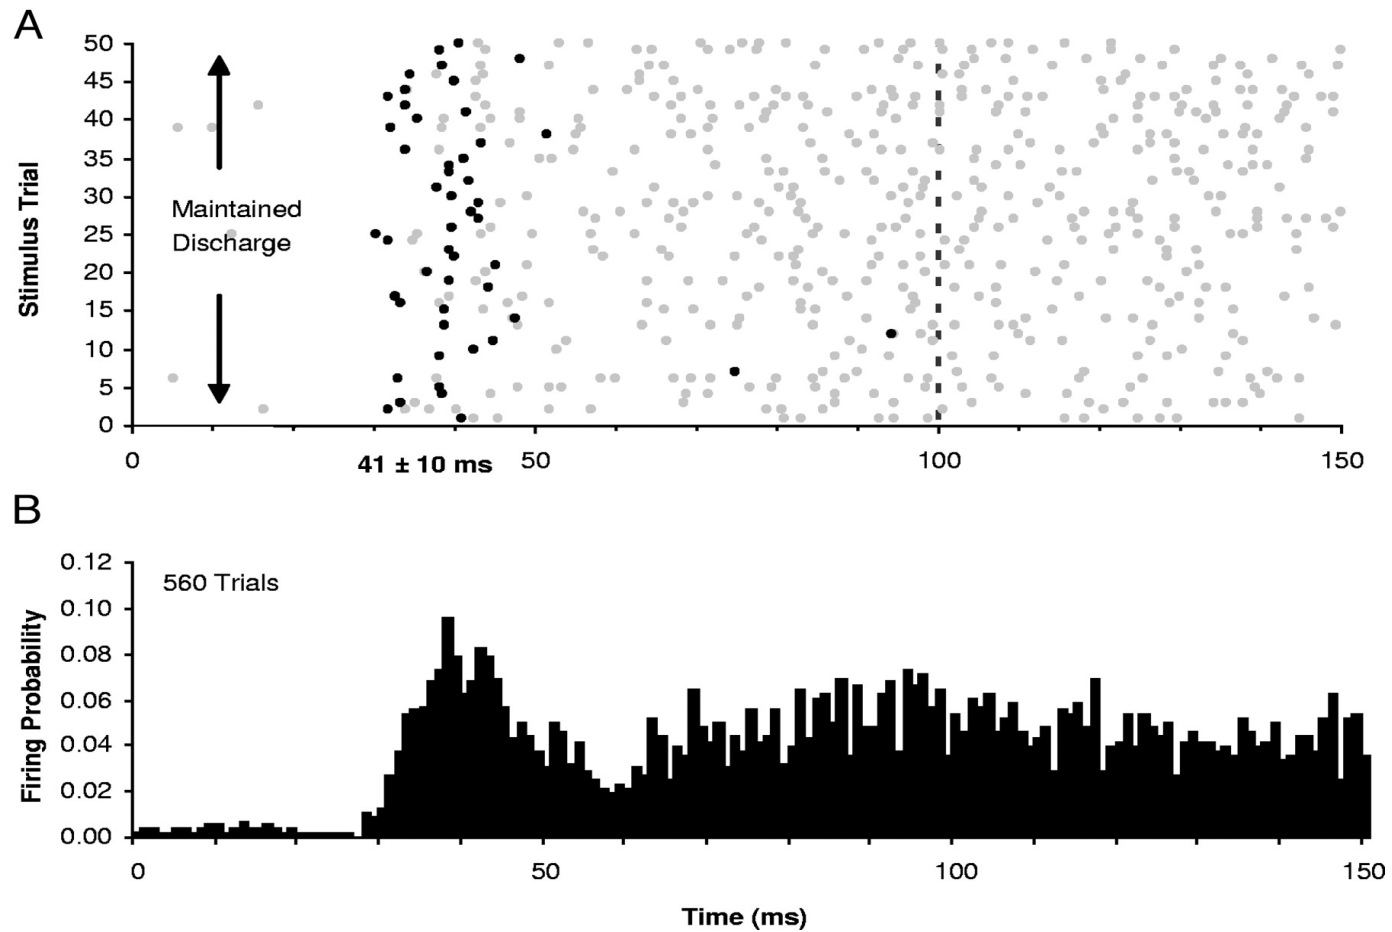
Urakubo et al. (2008) Model

- Highly detailed combination of 3 existing strongly-validated models:

(slide < O'Reilly)
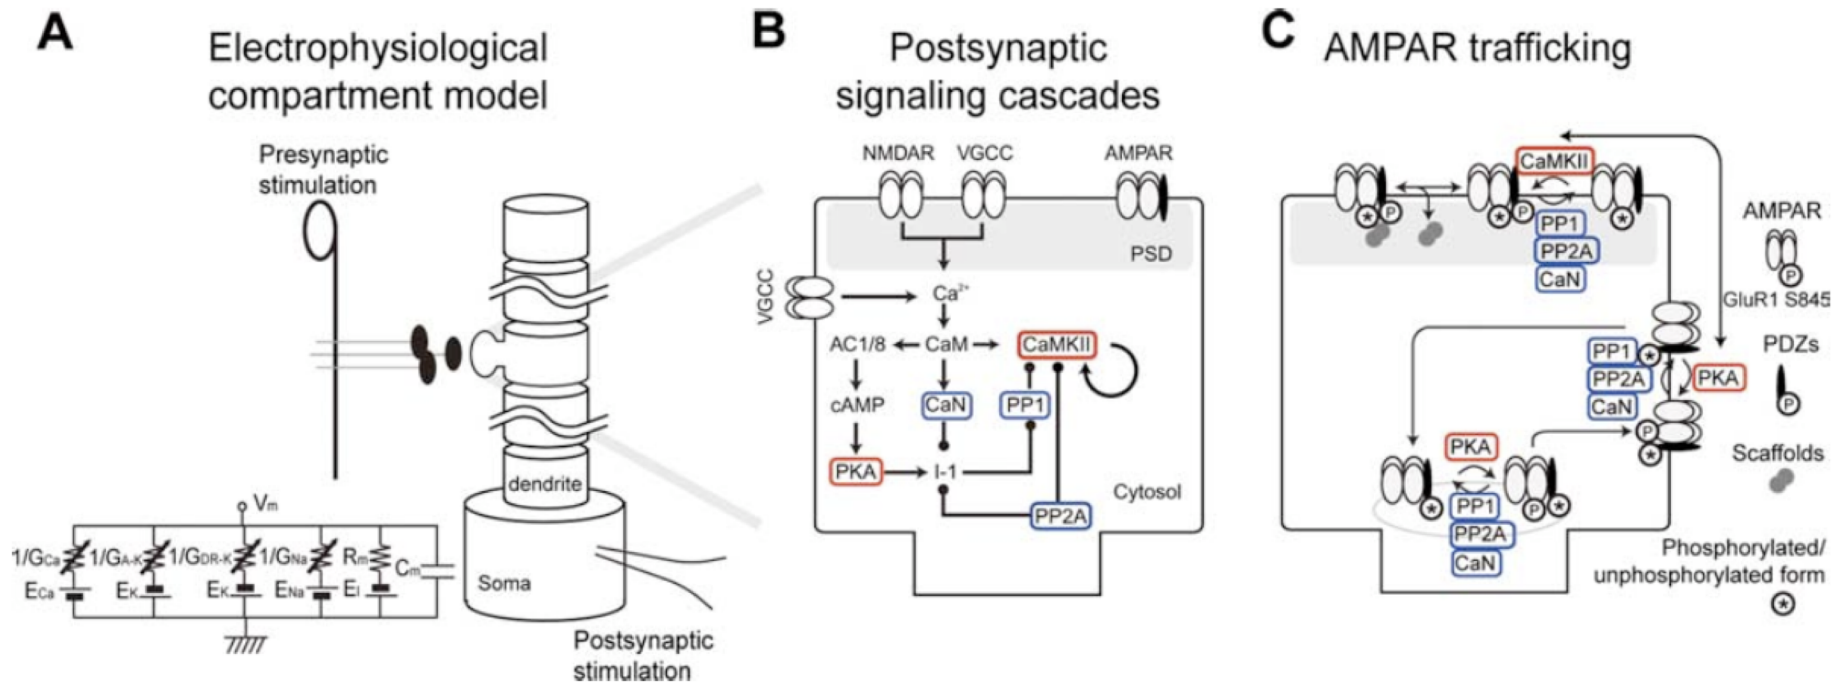
B. XCAL & BCM Models

XCAL = Linearized BCM

- Bienenstock, Cooper & Munro (1982) – BCM

- Adaptive threshold θ_M

  ✓ Lower when less active

  ✓ Higher when more

  ✓ Homeostatic

  ✓ Weights can decrease

(slide based on O'Reilly)

Evidence for Floating Threshold

COSC 421/521

(slide based on O'Reilly)
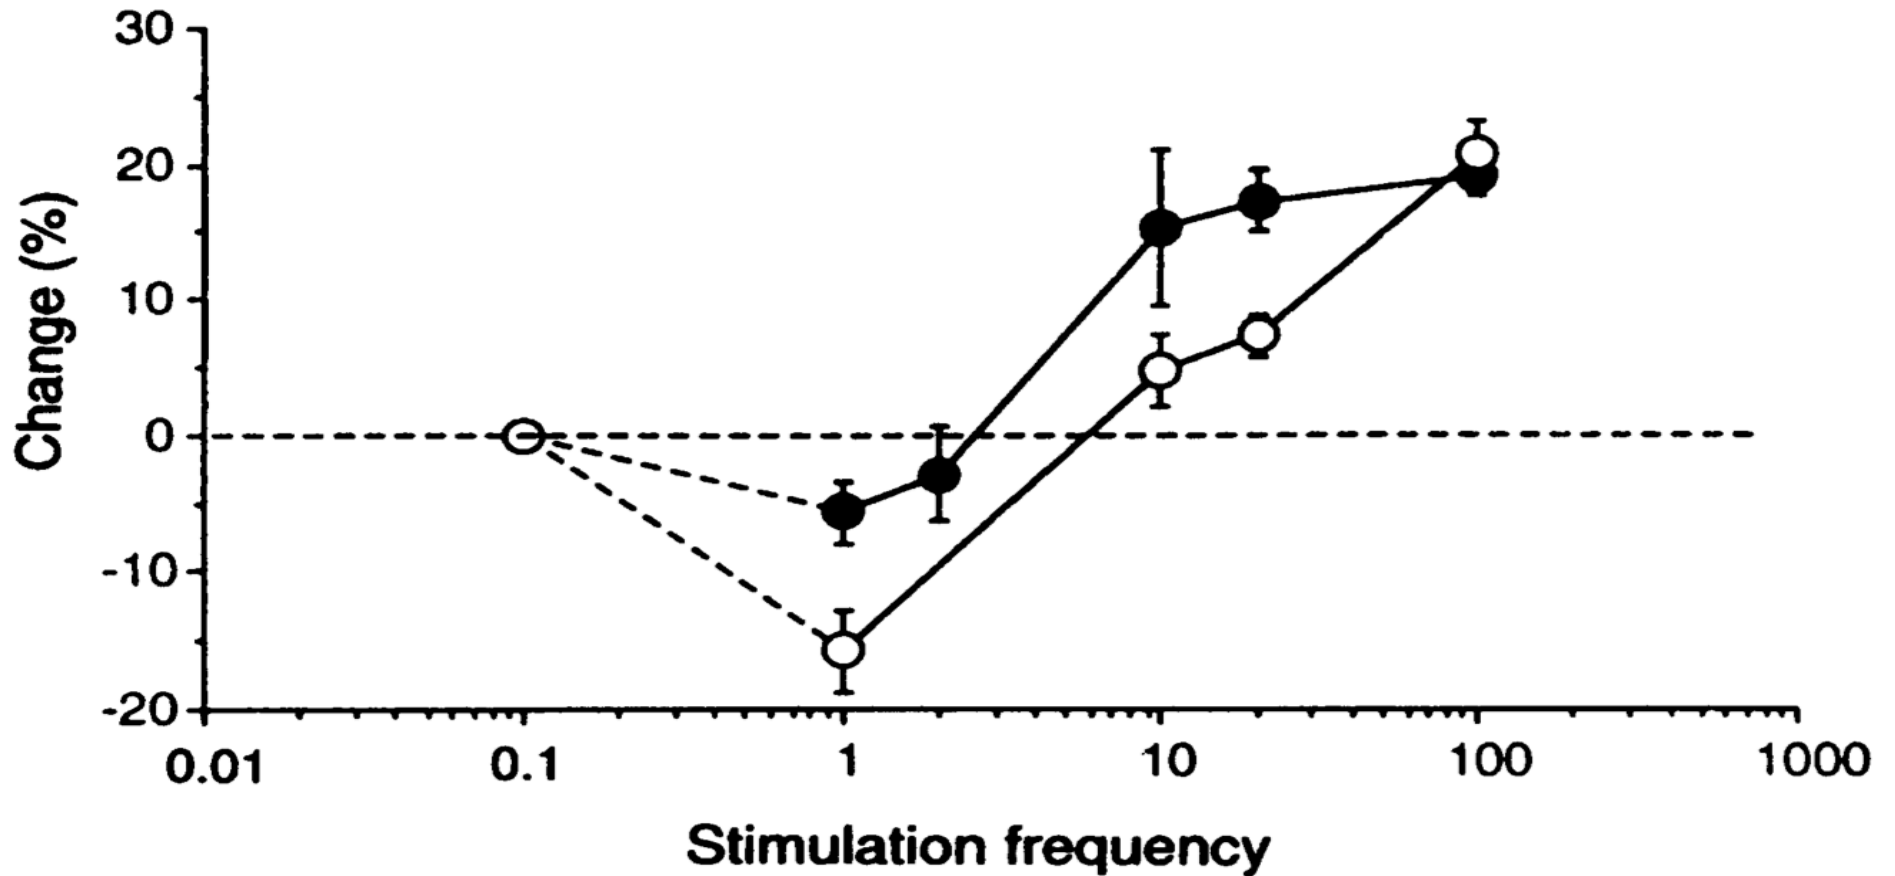
Homeostatic Behavior

Floating threshold adapts to long-term postsynaptic activity

Tends to equalize activity among neurons

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

LTP/LTD Approximation

Piecewise linear approximation to LTP/LTD

Typical $\theta_d = 0.1$

Floating threshold

$$\Delta W = \epsilon f_{\text{XCAL}}(\langle xy \rangle_s, \langle y \rangle_l)$$

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

Note: $\Delta W \approx \epsilon x_s(y_s - y_l)$

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

Simplified Computation
of Floating Threshold

$$\text{if } y > .2 \text{ then } y_l = y_l + \frac{1}{\tau_l} \left(\max - y_l \right)$$

$$\text{else } y_l = y_l + \frac{1}{\tau_l} \left(\min - y_l \right)$$

Integrating Over Experience

- Units that are activated (initially) by a wide range of stimuli end up representing an *average* of all these stimuli

    — The unit will learn what they have in common

- Units that are activated (initially) by a narrow range of stimuli end up representing very *specific* features

    — If a unit is only activated by a single stimulus, it will only represent that stimulus

- Thus the learning rule allows you to represent concepts at varying levels of abstraction, depending on the excitability of the unit

(slide based on Frank)

Multiple Units

- One detector can only represent one "thing" (pattern of correlated features)

- <u>Goal</u>: We want to have different units in the network learn to "specialize" for different things, such that each thing is represented by at least one unit

- Random initial weights and inhibitory competition are important for achieving this goal

(slide based on Frank)

Simple Competitive Learning

- Competitive learning network
  - — two layers, randomly initialized weights
  - — second is self-reinforcing, mutually inhibitory
  - — "winner takes all" dynamics
  - — implements nearest neighbor classification
- Learning
  - — winner moves toward last
  - — weight vectors move to centers of clusters
  - — implements k-means clustering

Competitive Network

W

COSC 421/521

Diagram of Competitive Learning

$$y_j = \sum_{k=1}^{n} W_{jk} x_k = \mathbf{W}_j \cdot \mathbf{x}$$

$$\mathbf{W}_j \cdot \mathbf{x} = \|\mathbf{W}_j\| \, \|\mathbf{x}\| \, \cos \varphi_j$$

If the $\mathbf{W}_j$ are normalized,

$$\mathbf{W}_j \cdot \mathbf{x} = \|\mathbf{x}\| \, \cos \varphi_j$$

$$y_j > \theta \text{ if } \|x\| \, \cos \varphi_j > \theta$$

Competitive Cluster Learning
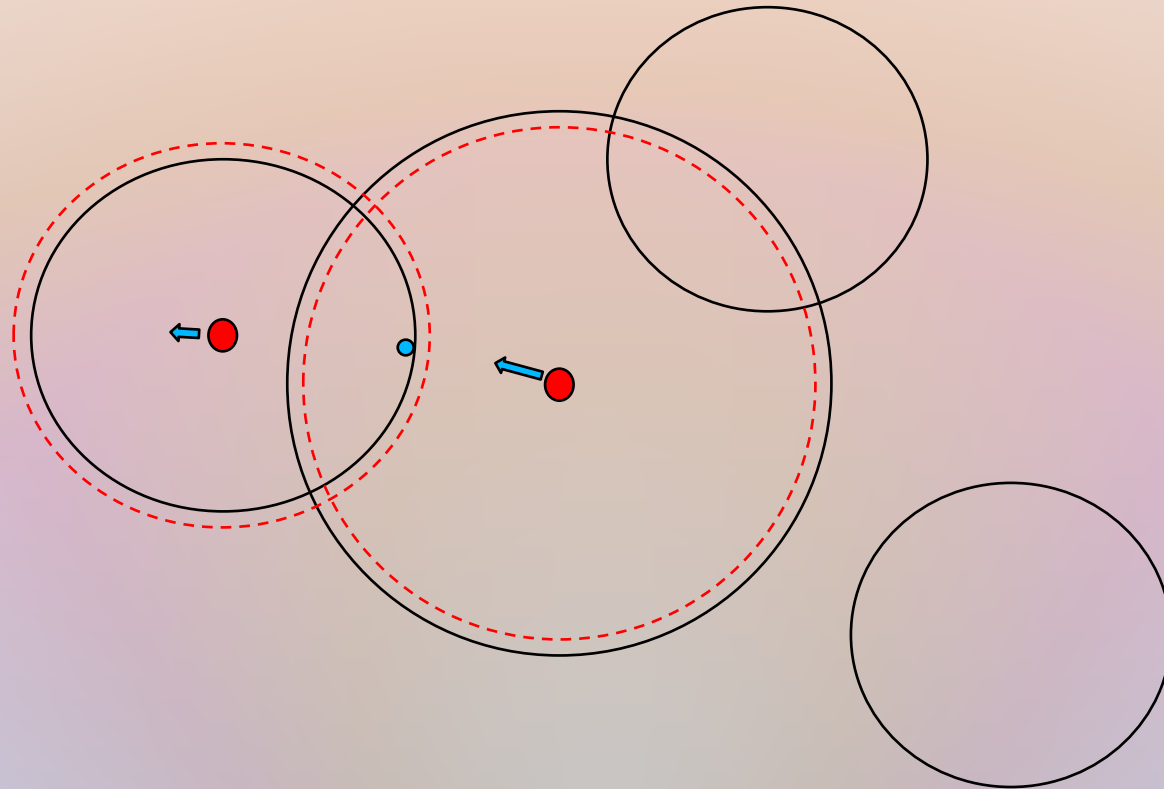
Self-Organized Learning

- 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

emergent demonstration: Self_Organizing

C. Error-driven Learning

Self-organized vs. Error-driven Learning

- Previously we have discussed self organized Hebbian learning

  — Leverage correlations to grow detectors that correspond to things in the world (cats, professors...)

- Next we will discuss error-driven (task) learning

  — Task = producing a specific output pattern in response to an input pattern

  — e.g., reading; giving the correct answer to 3 + 3

(slide based on Frank)

Task Learning

- Task learning encompasses:

  — Giving an appropriate response to a stimulus

  — Arriving at an accurate interpretation of a situation

  — Generating a correct expectation of what will happen next

- In all of the above cases, there is a correct answer...

(slide < Frank)

Floating Threshold = Medium Term Synaptic Activity (Error-Driven)

Error-Driven Learning

- For achieving intended outcomes

- Faster 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 backpropagation algorithm when combined with bidirectional connections

Backpropagation:
Mathematics of Error-driven Learning

(slide < O'Reilly)
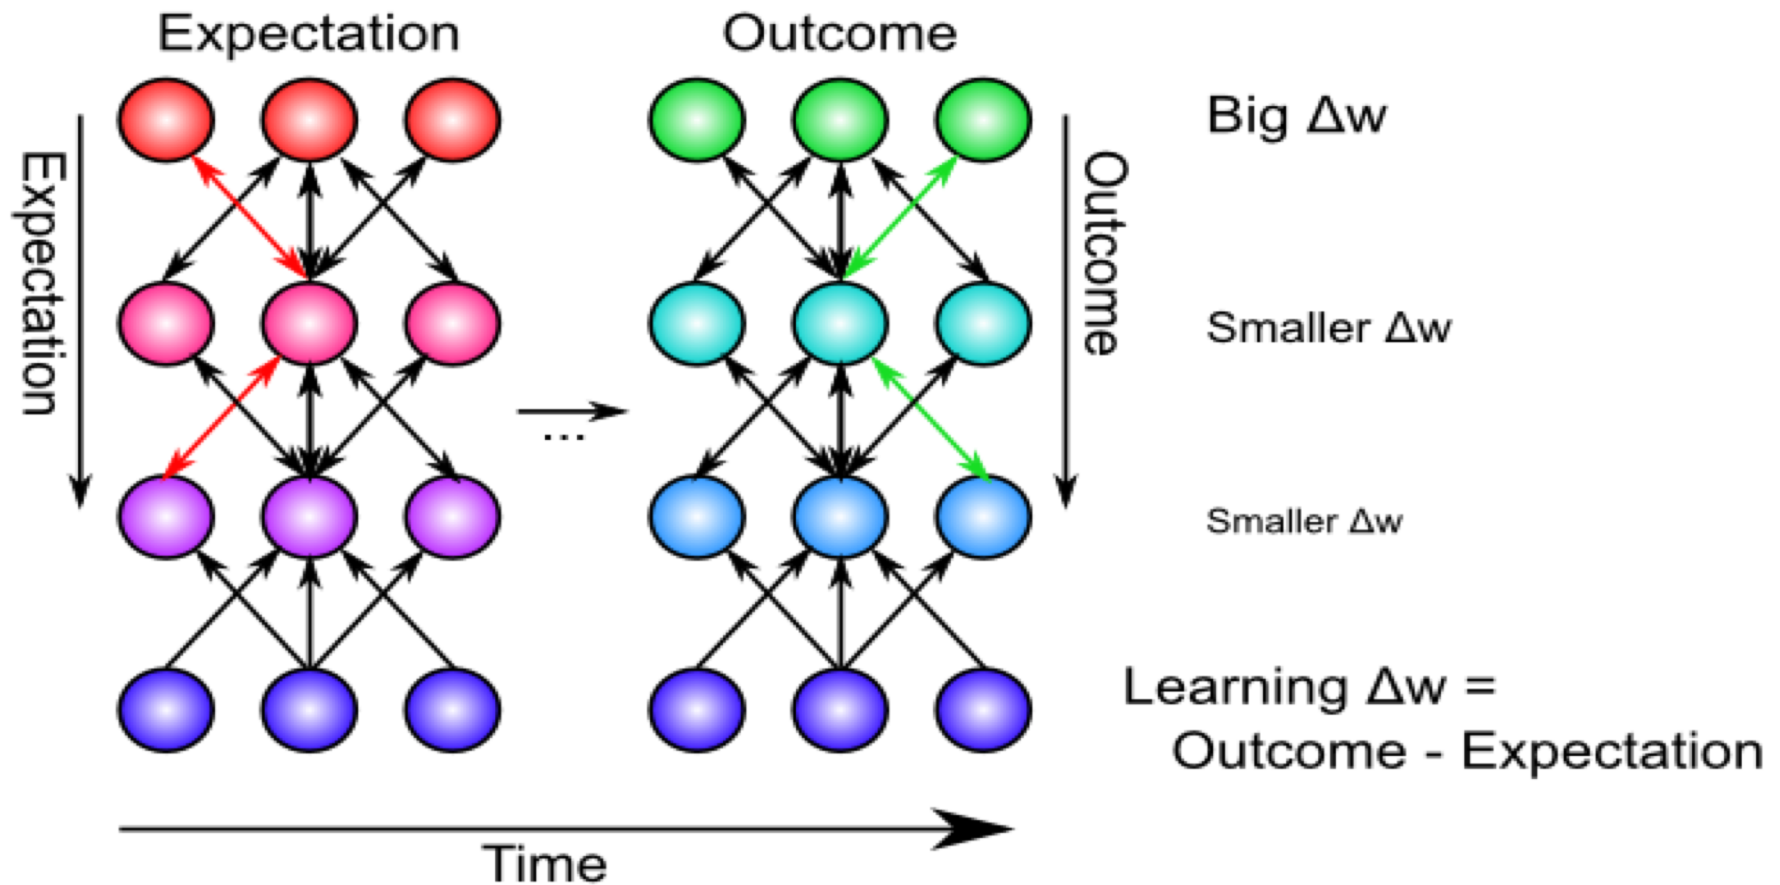
Credit/Blame Assignment via GeneRec

a) Expectation / Minus

$$z = f(\sum_j y_j w_j)$$

$$y = f(\sum_i x_i w_i + \sum_k z_k w_k)$$

b) Outcome / Plus

$$\delta = t - z = z^+ - z^-$$

$$\Delta w = y\delta = y(z^+ - z^-)$$

$$\delta = (\sum_k t_k w_k - \sum_k z_k w_k) y' = y^+ - y^-$$

$$\Delta w = x\delta = x(y^+ - y^-)$$

The Plus Phase helps identify bridging units that are well connected to both the input and the target output, and GeneRec adjusts weights to maximize the activity of these units. — M. Frank

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

More precisely…

Minus Phase:

Clamp x_i^- (inputs)

$$y_j^- = f\left(\sum_i x_i^- W_{ij} + \sum_k W_{jk} z_k^-\right)$$
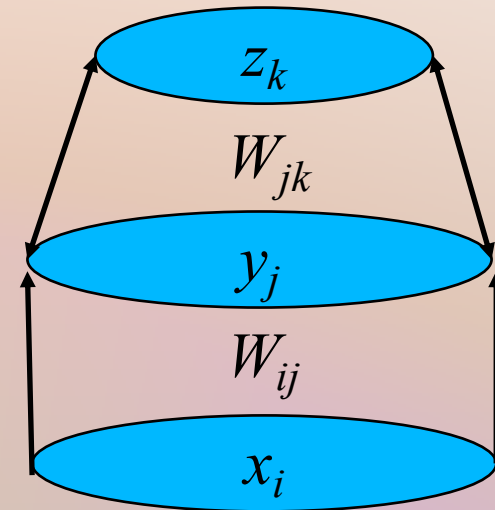
$$z_k^- = f\left(\sum_j y_j^- W_{jk}\right)$$

Plus Phase:

Clamp z_k^+ (target outputs)

$$\Delta W_{jk} = y_j^-(z_k^+ - z_k^-) = y_j^- \delta_k \text{ where } \delta_k = z_k^+ - z_k^-$$

$$y_j^+ = f\left(\sum_i x_i^- W_{ij} + \sum_k W_{jk} z_k^+\right)$$

$$\Delta W_{ij} = x_i^-(y_j^+ - y_j^-) = x_i^- \delta_j \text{ where } \delta_j = y_j^+ - y_j^-$$

Two Issues

- Two issues:
  - Need weights to be symmetric
  - Why should we use minus phase sending activity instead of plus phase?

- Solution: Average together plus and minus phase sending activation, and average together feedforward and feedback weight changes

$$\Delta W_{ij} = \epsilon \left[\frac{x_i^+ + x_i^-}{2} \left(y_j^+ - y_j^- \right) + \frac{y_j^+ + y_j^-}{2} \left(x_i^+ - x_i^- \right) \right]$$

$$= \frac{\epsilon}{2} \left[\left(x_i^+ + x_i^- \right) \left(y_j^+ - y_j^- \right) + \left(x_i^+ - x_i^- \right) \left(y_j^+ + y_j^- \right) \right]$$

$$= \epsilon \left(x_i^+ y_j^+ - x_i^- y_j^- \right)$$

(slide based on Frank)

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

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

$$\Delta W = \epsilon f_{\text{XCAL}}(\langle xy \rangle_{\text{s}}, \langle xy \rangle_{\text{m}})$$

$$\approx \epsilon f_{\text{XCAL}}(x_{\text{s}}y_{\text{s}}, x_{\text{m}}y_{\text{m}})$$

$$\approx \epsilon(x_{\text{s}}y_{\text{s}} - x_{\text{m}}y_{\text{m}})$$

Backpropagation Algorithm

- Adjusts weights in multilayer artificial neural net by "steepest descent" in order to minimize error

  — does this by estimating partial derivatives of error with respect to weights

- (Re)invention in 1986 rejuvenated artificial neural net research

  — had been neglected for 15 years

- Basis for current "deep learning" methods

- Uses biologically implausible backwards error propagation through synapses

- Nevertheless, the biologically plausible XCAL algorithm is a good approximation to BP

Relation to BP Algorithm (backward propagation of errors)

$$\delta_j = y_j^+ - y_j^-$$

$$= f\left(\sum_i x_i^- W_{ij} + \sum_k W_{jk} z_k^+\right) - f\left(\sum_i x_i^- W_{ij} + \sum_k W_{jk} z_k^-\right)$$

$$\approx f'\left(\sum_i x_i^- W_{ij} + \sum_k W_{jk} z_k^-\right)\left(\sum_k W_{jk} z_k^+ - \sum_k W_{jk} z_k^-\right)$$

$$= y_j' \left[\sum_k W_{jk}(z_k^+ - z_k^-)\right]$$

$$= y_j' \sum_k W_{jk} \delta_k$$

Summary: Error-driven Learning

- Adjusts expectations to better match outcomes (better prediction)

- Uses later, better information to train earlier expectations (later training earlier)

- Allows network to more quickly settle into attractor (quicker convergence)

- But at this time there is limited empirical support for medium timescale floating threshold

emergent demonstration: Pattern_Associator

emergent demonstration: Error_Driven_Hidden

Weight Bounding

- Biologically, synaptic weights are bounded

- Approximate computationally by exponential approach to limits 0 and 1:

```
if dwt > 0 then wt = wt + (1—wt)×dwt;

else wt = wt + wt×dwt;
```

- Gradual approach preserves signal better than clipping

Contrast Enhancement

- To enhance contrast in activities for midrange weights, distinguish:

  — internal weight *w* determined by learning

  — contrast-enhanced weight $\hat{w}$ determines synaptic efficacy

- Contrast enhancement function:

$$\hat{w} = \frac{1}{1 + \left(\frac{w}{\theta(1-w)}\right)^{-\gamma}}$$

COSC 421/521

D. What drives learning?

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: $\quad \Delta W = \eta f_{\text{XCAL}}\left(x_s y_s, x_m\left(\lambda y_1 + (1-\lambda) y_m\right)\right)$

- Probably not the whole story

Learning Situations

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

Leabra

Leabra Mechanisms

Deep (many layers)

Bidirectional Excitation

Learning

$$\Delta w = \langle xy \rangle_s - \langle xy \rangle_m \quad \text{(Error-driven)}$$
$$+ \langle xy \rangle_s - \langle y \rangle \quad \text{(BCM Hebbian)}$$

k-Winners Inhibitory Competition

(sparse distributed representations)

$$\frac{dV_m^t}{dt} = t \sum_c g_c^t \overline{g}_c (E_c - V_m^t)$$

Integrate & Fire Point Neurons

Leabra = Learning in an Error-driven and Associative, Biologically Realistic Algorithm

COSC 421/521

(fig. < O'Reilly)

emergent demonstration: Family_Trees

Supplementary

Oja's Rule (1)

- To avoid instability due to unlimited growth, suppose we keep the weight vectors normalized:

$$\mathbf{w}' = \frac{\mathbf{w} + \epsilon y \mathbf{x}}{\|\mathbf{w} + \epsilon y \mathbf{x}\|}$$

- Consider one component: $w_j' = \dfrac{w_j + \epsilon y x_j}{\left[\sum_k (w_k + \epsilon y x_k)^2\right]^{1/2}}$

- Call RHS $f(\epsilon)$ and expand in Taylor series:

$$f(\epsilon) = f(0) + \epsilon f'(0) + \epsilon^2 f''(0) + \cdots$$

Oja's Rule (2)

- You can show (algebra!):
$$w_j' \approx \frac{w_j}{\|\mathbf{w}\|} + \epsilon \left[\frac{yx_j}{\|\mathbf{w}\|} - \frac{w_j \sum_k yx_k w_k}{\|\mathbf{w}\|} \right] + O(\epsilon^2)$$

- If **w** is normalized,
$$w_j' = w_j + \epsilon \left(yx_j - w_j y \sum_k x_k w_k \right)$$
$$= w_j + \epsilon (yx_j - w_j y^2)$$

- In vector form, $\Delta \mathbf{w} = \epsilon y \mathbf{x} - \epsilon y^2 \mathbf{w}$

- This is Hebbian learning with decay

Stability and Convergence

- Under broad conditions, will stabilize: $0 = \Delta\mathbf{w} = \epsilon(\mathbf{x}y - y^2\mathbf{w})$

- Note $y = \sum_k x_k w_k = \mathbf{x}^\mathrm{T}\mathbf{w} = \mathbf{w}^\mathrm{T}\mathbf{x}$

- Therefore, $0 = \epsilon[\mathbf{x}(\mathbf{x}^\mathrm{T}\mathbf{w}) - (\mathbf{w}^\mathrm{T}\mathbf{x})(\mathbf{x}^\mathrm{T}\mathbf{w})\mathbf{w}]$
  $0 = \epsilon[(\mathbf{x}\mathbf{x}^\mathrm{T})\mathbf{w} - \mathbf{w}^\mathrm{T}(\mathbf{x}\mathbf{x}^\mathrm{T})\mathbf{w}\mathbf{w}]$
  $0 = (\mathbf{x}\mathbf{x}^\mathrm{T})\mathbf{w} - \mathbf{w}^\mathrm{T}(\mathbf{x}\mathbf{x}^\mathrm{T})\mathbf{w}\mathbf{w}$

- The parenthesis is the covariance matrix $\mathbf{R} = \langle\mathbf{x}\boldsymbol{x}^\mathrm{T}\rangle$

- Hence $\mathbf{R}\mathbf{w} = (\mathbf{w}^\mathrm{T}\mathbf{R}\mathbf{w})\mathbf{w}$

- Therefore $\mathbf{w}$ is an eigenvector of $\mathbf{R}$ with eigenvalue $\lambda = \mathbf{w}^\mathrm{T}\mathbf{R}\mathbf{w}$

- In fact, you can show λ is the largest eigenvalue, and therefore $\mathbf{w}$ is the *first principal component*, which accounts for most variance in $\mathbf{x}$

Summary: Hebb's and Oja's Rules

- A single Hebbian neuron with decay (Oja's Rule) extracts the first principal component of a stationary input

- The variance of this component is given by the corresponding (maximum) eigenvalue

- Oja's Rule can be extended into a *generalized Hebbian algorithm* that extracts all the principal components and their variances

Summary:
Learning in Brains and Computers

- Competitive learning

  — Competitive k-means clustering

- Error-driven learning

  — Back-propagation

- Hebbian learning with decay

  — Oja's rule, generalized Hebbian algorithm, PCA