Neural circuits for cognition

Unsupervised learning: biology and computation

MIT Course 9.49/9.490

Instructor: Professor Ila Fiete

Announcement

- New homework posted. Due: 11/19 – TUESDAY! (nearly 2 weeks, but not a FRIDAY).

Unsupervised learning

- No task contingency.
- Learning simply from watching without feedback: e.g. watching TV, watching world from a café, listening to music, etc.
- Learn associations between inputs.
- Plasticity rules: based on coincident inputs, coincident input-output.
- Models: Network self-organizes to reflect correlations/statistical regularities in the input.

The dream of unsupervised learning

- Most learning in the real world is unsupervised.
- Supervision/supervised data is sparse/rare/precious
- Hope: After learning of regularities of world from unsupervised learning, will then require only a touch of supervision to solve tasks. (e.g. find clusters in the set of data; get label for one example, “a cat”; this label then applies to whole cluster, so you know many cats).

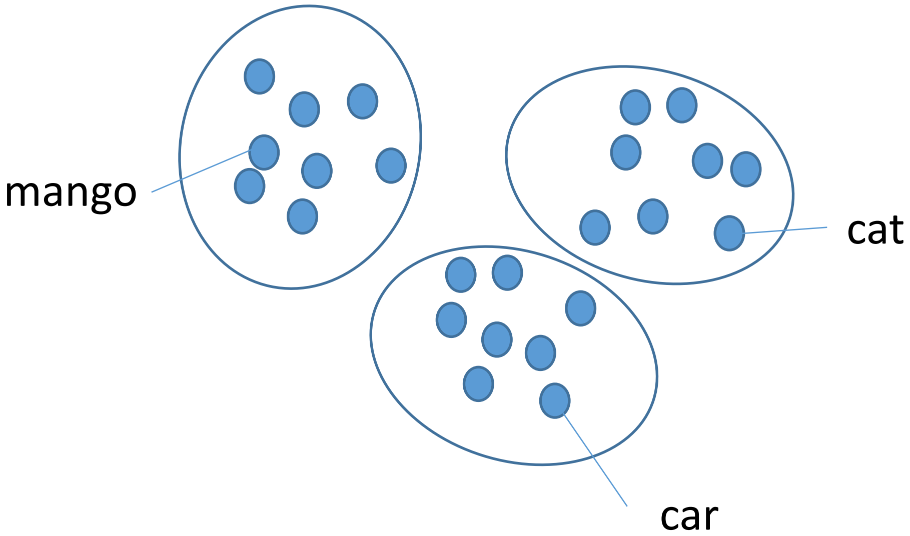

What is unsupervised learning

- “Unsupervised” is not categorical, it is part of a continuum.
- Also, can mean different things: *No explicit error function* - versus – *Error functions that do not require a teacher*.
- E.g.: learning to predict next state does have a specific task/loss, but also considered unsupervised because no supervisor. In ML, this type of unsupervised learning involves the use of supervised learning rules to propagate errors.
- Here, we will take the definition of “no explicit error function” in learning.

Synaptic plasticity rules (models) for unsupervised learning

- Temporally asymmetric rules: STDP. Already saw one example of this at work (sequence formation).
- Temporally symmetric: Hebbian learning, Oja's rule.

Hebb’s rule

- Neurons that fire together wire together (neurons that fire out of sync fail to link). Hypothesized by D. O. Hebb and others before him.
- Learn about co-occurence of inputs: if neurons A, B driven by co-occurring inputs, then they will fire together, predicting the correlation.
- Simple neuron/synapse-level explanation or mechanism for the associative learning seen in psychology (Pavlovian conditioning).

A neural substrate for associative (Hebbian) learning

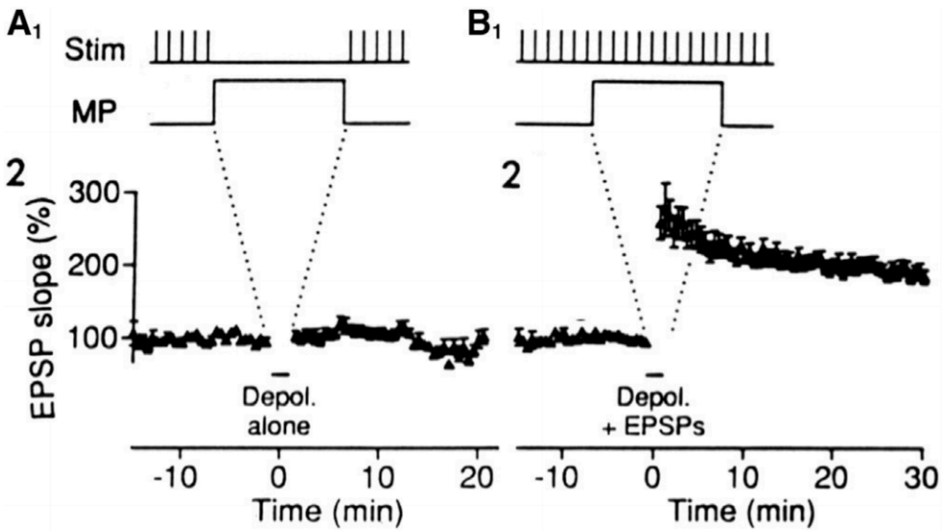

Figure 3. Minimum Requirements for Induction of LTP
The top series of diagrams (A1 and B1) illustrates schematically, at an expanded timescale, the stimulation of the excitatory synapses (Stim) and the control of the membrane potential (MP). The graphs below (A2 and B2) plot the maximal initial slope of the EPSP. In (A), synaptic stimulation was stopped and the cells were depolarized to 0 mV for 2 min. In (B), synaptic stimulation continued throughout the experiment and the cells were depolarized to 0 mV for 2 min. The recording electrode contained cesium to allow depolarization of the membrane. Each graph averages the results from 8–12 cells. Each slope measurement in an individual experiment was normalized to the average value of all points on the baseline (at least 10 min prior to each manipulation) for that experiment. Experiments were then divided into 20 s bins, each of which was averaged. Data are shown as mean ± SE (modified from Malenka et al., 1989).

Associative:
Need pre, post pairing.

Presynaptic spikes,
Postsynaptic activation.

Hypothesized mechanism for need for pre, post pairing (associativity)

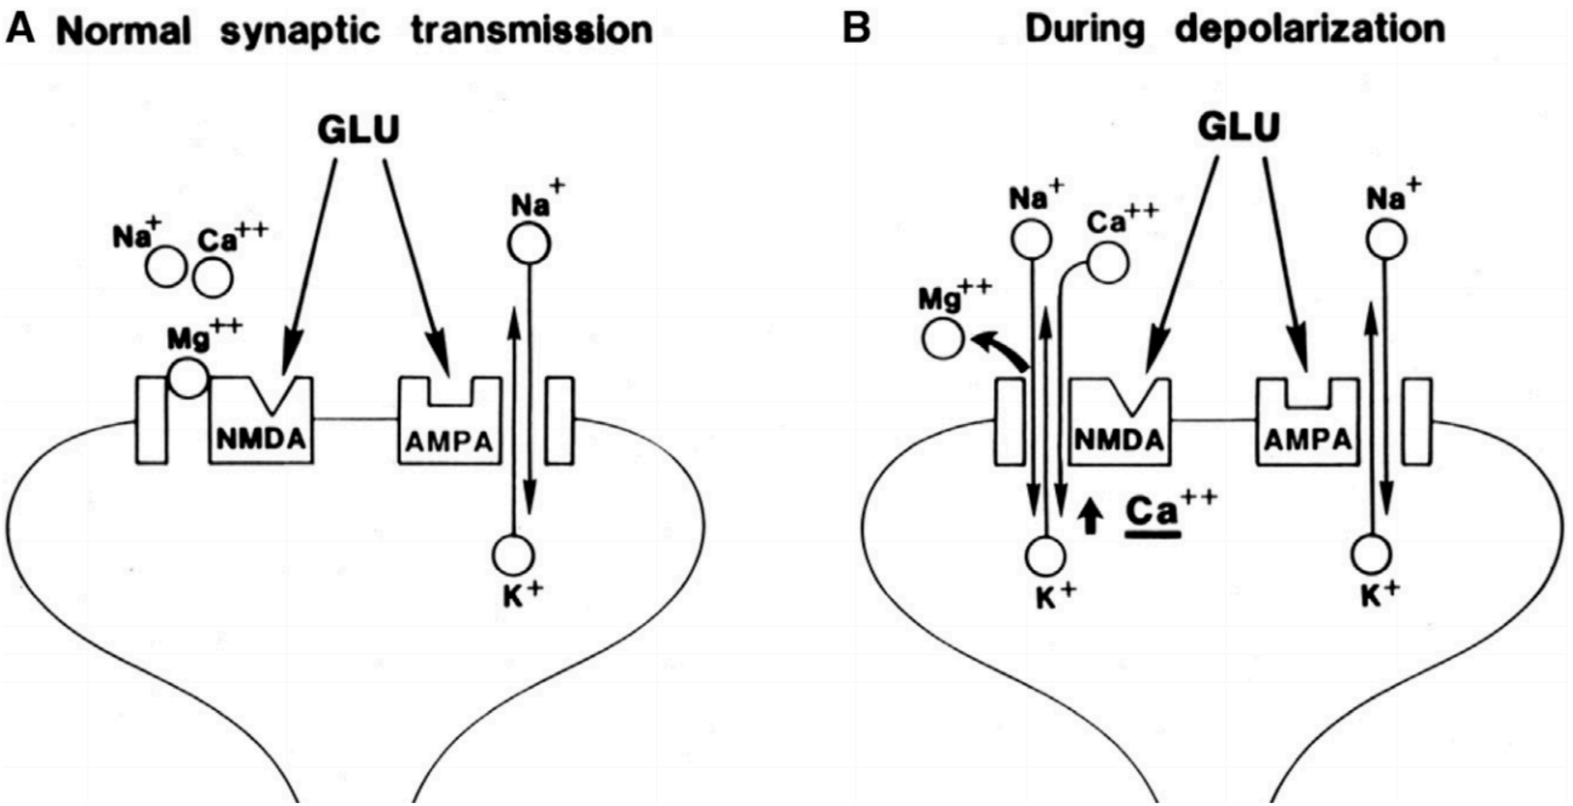

Figure 4. Model for the Induction of LTP in the CA1 Region of the Hippocampus

(A) The events occurring during low-frequency synaptic transmission. Glutamate is released from the presynaptic terminal and binds to both NMDA and AMPA receptors. Na^+ and K^+ flow through the AMPA receptor channel but not through the NMDA receptor channel, due to Mg^{2+} block of this channel. (B) The events occurring when the postsynaptic membrane is depolarized, as would occur during a high-frequency tetanus. The depolarization relieves the Mg^{2+} block of the NMDA receptor channel, allowing Na^+, K^+, and most importantly Ca^{2+} to flow through the channel. The primary target of the rise in Ca^{2+} in dendritic spines is the calcium-calmodulin-dependent kinase II (CaMKII) (modified from Nicoll et al., 1988).

Test: glutamate uncaging coupled with postsynaptic depolarization. Gives robust LTP.

Harvey and Svoboda, 2007; Lee et al., 2009; Matsuzaki et al., 2004; Tønnesen et al., 2014

Other forms of plasticity in the brain

- [Short-Term Synaptic Plasticity]
  - Short-term depression/facilitation
  - Dynamics may change on a long-term basis via LTP/LTD
- [Changes to intrinsic excitability of cell]
  - Density and distribution of various channels (ionic conductances)
  - Currently active research area
- [Growth and morphological changes in dendrites]
  - Currently active research area
- [Addition of new neurons?]
  - Hot topic of research in recent years…

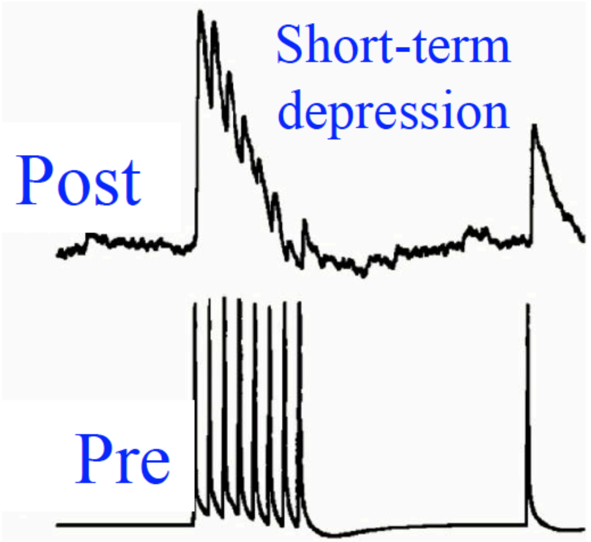

Back to Hebb's rule

- Consider the simplest network: single neuron, linear response, many inputs:

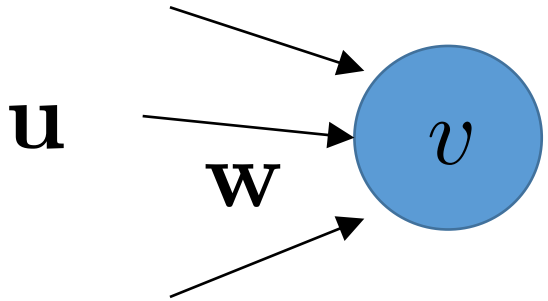

$$\text{inputs} \quad \mathbf{u}$$

$$\text{output} \quad v = \mathbf{w}^T \mathbf{u} = \mathbf{u}^T \mathbf{w}$$

On-board derivation of learning of input statistics

On-board derivation of learning:

- Hebbian rule and learning the top (principal) eigenvector.
- Simple Hebbian rule is unstable.
- Hebbian rule with simple weight decay: also unstable.
- Oja’s rule: nonlinear weight decay. Preserves norm of weights and learns principal eigenvector.