

Neuroscience Foundations

Study Guide

by Richard J. Young, Ph.D.

richard@deepneuro.ai

Chapter 2 of "The NeuroAI Handbook"

What a single neuron actually computes, how it spikes, and how the brain learns without breaking.

One-sentence thesis. A single biological neuron is already a small network, and the brain's learning is stable only because three forces cooperate: a rule that changes synapses, a brake that keeps them bounded, and a signal that says *when* the change matters.

PART 1: The Neuron Is Not a Perceptron

Inside a Single Cell

The textbook cartoon is a bag that sums its inputs: dendrites collect signals, the soma adds them up, and the axon hillock fires a spike if the total clears threshold. That cartoon is the artificial neuron. The real cell is richer. Its **dendrites** are not passive wires; they contain active machinery that computes *before* the soma ever sees the result.

Two mechanisms do the heavy lifting. **NMDA spikes** are regenerative plateaus in a dendritic branch that fire only when inputs arrive together, an AND-like coincidence detector. **Calcium spikes** in the apical (top) dendrites let distant, top-down signals gate or amplify whatever the basal dendrites are receiving. Each branch behaves as a quasi-independent nonlinear subunit, and the soma pools their outputs.

The upshot. A neuron approximates a shallow multilayer network in one cell. That extra layer of computation is exactly what improves credit assignment and lets context (top-down input) gate feedforward drive.

ANALOGY

Think of the neuron as a committee, not a single voter. Each dendritic branch is a subcommittee that debates and votes first; the soma only tallies the subcommittees' verdicts. A perceptron is one voter. That is the difference.

A Zoo of Synapses

The connections are not all alike either, and the differences are computational. **AMPA** receptors give fast, roughly linear excitation. **NMDA** receptors are blocked by magnesium until the cell depolarizes, so they act as coincidence detectors and gate plasticity. **GABA** provides inhibition, fast shunting (GABA-A) or slow modulation (GABA-B). Layered on top, neuromodulators like dopamine and acetylcholine adjust gain and plasticity thresholds through slower metabotropic pathways.

Synapses also change strength on the fly. **Short-term plasticity** makes a synapse *facilitate* (grow stronger over a burst) or *depress* (weaken as vesicles deplete). Functionally this is temporal filtering: a facilitating synapse emphasizes sustained bursts, a depressing one emphasizes onsets, a built-in sense of what just changed.


PART 2: How a Spike Happens

The Hodgkin-Huxley Machine

Why is a spike all-or-nothing, with that same sharp shape every time? Hodgkin and Huxley answered this in 1952 (and won a Nobel for it) by treating the membrane as a **capacitor** in parallel with voltage-gated ion channels. Current either charges the membrane or leaks out through sodium, potassium, and leak channels:

$$C_m \frac{dV}{dt} = I_{ext} - I_{Na} - I_K - I_L$$

The genius is in the **gating variables** m , h , n , each the probability that a class of channel gate is open, each with its own speed. The spike is their choreography: depolarization opens fast m gates so sodium rushes in (positive feedback, the upstroke); slow h gates then close and shut sodium off; delayed n gates open and potassium flows out, dragging the voltage back down and briefly past rest. Fast excitation, delayed inhibition, one clean pulse.



The action potential: sodium drives the upstroke, potassium the downstroke and undershoot.

From Biophysics to a One-Line Neuron

You rarely need all four channels. Simpler models trade biophysical detail for speed, and the one AI uses most, the **leaky integrate-and-fire** neuron, keeps only four operations: integrate, leak toward rest, compare to threshold, reset. The ladder below runs from most faithful to most efficient; where you sit on it depends on whether you want mechanism or throughput.


Neurons Adapt

Give a real neuron a constant input and it does *not* fire at a constant rate. It fires fast, then slows, **spike-frequency adaptation**, driven by slow potassium and calcium-activated currents that accumulate with activity. Computationally this is a high-pass filter: the cell shouts about change and whispers about steady state. The same instinct shows up in AI as normalization layers, attention decay, and plain habituation.

Model	Equations	Key feature
Hodgkin-Huxley	4 ODEs	Full biophysics with explicit ion channels; the reference model
FitzHugh-Nagumo	2 ODEs	Phase-plane view of excitability; strips channels to two variables
Izhikevich	2 ODEs + reset	Reproduces most cortical firing patterns very cheaply
Exponential IF	1 ODE + reset	Adds a realistic spike-initiation nonlinearity
Leaky IF	1 ODE + reset	Simplest useful spiking unit; the workhorse in AI models

Cells That Fire Together

Donald Hebb’s slogan, “cells that fire together, wire together,” is the ancestor of every learning rule in this book. Modern experiments sharpened it into **spike-timing-dependent plasticity**: the exact millisecond order matters. If the presynaptic neuron fires *just before* the postsynaptic one, the synapse strengthens (LTP); reverse the order and it weakens (LTD). Timing encodes causality, cause-before-effect gets rewarded.



The STDP window. Pre-before-post (right) strengthens; post-before-pre (left) weakens.

Why Learning Needs a Brake

Left alone, STDP is unstable: strong synapses fire the cell more, which strengthens them further, a runaway loop that saturates every weight. The brain prevents this with **homeostasis**, synaptic scaling that renormalizes a neuron’s total input, and sliding plasticity thresholds (the BCM idea) that make potentiation harder when a cell has been too active. Without a brake, no learning rule stays useful for long. AI rediscovered the same necessity as weight normalization

and regularization.

The Third Factor

Timing tells you *who* fired with whom, but not whether the outcome was any good. The brain adds a third factor: a broadcast neuromodulatory signal, dopamine for reward, acetylcholine for surprise, that gates when a Hebbian change is actually written in. A useful shorthand for these **three-factor rules** is a product of pre-activity, post-activity, and a modulator:

$$\Delta w \propto x_{pre} \cdot x_{post} \cdot M$$

This is the biological answer to **credit assignment**. Backpropagation solves the same problem with a global error signal and symmetric backward weights, biologically implausible but effective. Local plasticity plus a neuromodulatory “now it counts” is the brain’s more distributed alternative.

ANALOGY

A student takes notes constantly (Hebbian pairing), but only burns something into memory when a surprise or a reward marks the moment as important (the dopamine third factor). Constant note-taking without that stamp is just a full notebook nobody rereads.

PART 4: A Map of the Brain for AI

The brain is not one learner but several, each specialized, and each maps onto a familiar AI idea. The **cortex** runs hierarchical predictive processing, the biological cousin of deep learning. The **hippocampus** is a memory index: the dentate gyrus separates similar patterns, area CA3 completes a whole memory from a partial cue, CA1 binds and times, and replay during rest consolidates and plans, the blueprint behind retrieval-augmented and world-model systems.

The **basal ganglia** are an actor-critic: a “Go” pathway (D1) and a “NoGo” pathway (D2) compete to select actions, and a dopamine reward-prediction error, the same quantity as the temporal-difference error, tunes the balance. The **cerebellum** is a supervised learner in the Marr-Albus-Ito mold, with the climbing fiber delivering an explicit teacher signal. The **thalamus** routes and gates traffic between cortical areas, the brain’s attention and control layer.

Brain system	What it computes	AI analog
Cortex	Hierarchical predictive processing	Deep networks
Hippocampus	Pattern separation, completion, replay	Memory nets, RAG, world models
Basal ganglia	Action selection via dopamine reward-prediction error	Actor-critic / TD learning
Cerebellum	Supervised error correction (climbing-fiber teacher)	Supervised learning
Thalamus	Routing, gating, attentional control	Attention and controller modules

Quick Self-Test

1. Why is “a neuron is a single perceptron” wrong? Name the dendritic mechanism that makes a cell act like a shallow network.
2. In the Hodgkin-Huxley spike, what do the sodium and potassium currents each contribute to the waveform?
3. STDP is unstable on its own. What two homeostatic mechanisms keep it in bounds, and why are they needed?
4. What is the “third factor” in a three-factor learning rule, and what problem does it solve that plain Hebbian pairing cannot?
5. Match three brain systems to their AI analogs (cortex, basal ganglia, cerebellum).

If You Remember One Thing

Learning in the brain is not one rule but a balance of three: a Hebbian change (STDP) that rewrites synapses, a homeostatic brake that keeps them bounded, and a neuromodulatory signal that says when the change is worth keeping. Remove any one and the system either fails to learn or tears itself apart, which is the real lesson for anyone building brain-inspired AI.

Richard J. Young, Ph.D. · richard@deepneuro.ai · Senior AI Research Scientist, UnitedHealth Group · UNLV Lee Business School · ORCID 0000-0002-1109-7552