

Neurostimulation

Study Guide

by Richard J. Young, Ph.D.

richard@deepneuro.ai

Chapter 6 of "The NeuroAI Handbook"

How the brain rewires itself, how neuromodulators gate learning, and how both echo the training of AI.

One-sentence thesis. The brain learns by rewiring itself, Hebbian strengthening held in check by homeostatic control and gated by neuromodulators, and both the rules (STDP, dopamine as reward-prediction error) and the interventions (tDCS, TMS, DBS) map onto how we train and fine-tune artificial networks.

PART 1: How Synapses Learn

Cells That Fire Together, Wire Together

Learning has to change something physical. In the brain, what changes is the *strength* of the synapse between two neurons. Donald Hebb's 1949 postulate stated the core idea: when neuron A repeatedly helps fire neuron B, the connection from A to B grows stronger. The slogan is "cells that fire together, wire together," and it remains the foundation of plasticity theory.

In symbols, the weight change is just the product of pre- and postsynaptic activity, scaled by a learning rate:

$$\Delta w_{ij} = \eta x_i y_j$$

The instability. Pure Hebbian learning is a positive-feedback loop: strong synapses keep strengthening without limit, driving runaway excitation. Something must push back, which is why the brain pairs Hebbian change with homeostatic control.

Strengthen and Weaken: LTP and LTD

Hebb's rule is implemented at the cellular level by two opposing processes. **Long-Term Potentiation (LTP)** strengthens synapses: high-frequency stimulation (around 100 Hz) opens **NMDA receptors**, which act as *coincidence detectors* because they require both presynaptic glutamate and postsynaptic depolarization at once. Calcium floods in, activates CaMKII, and more AMPA receptors are inserted (first shown in hippocampal slices by Bliss and Lomo, 1973). **Long-Term Depression (LTD)** does the reverse: low-frequency

stimulation (around 1 Hz) lets in only moderate calcium, activating phosphatases that remove AMPA receptors. These changes last hours to months, the cellular basis of memory.

Timing Is Everything: STDP


Hebb said nothing about *order*. **Spike-Timing-Dependent Plasticity** adds it: if the presynaptic neuron fires shortly *before* the postsynaptic one (within about 20 ms), the synapse strengthens; if it fires *after*, the synapse weakens; and the effect shrinks exponentially as the gap widens (Bi and Poo, 1998). This lets a network learn cause from effect, since the input that reliably precedes an output gets the credit.

Keeping It Stable: Homeostatic Plasticity

If Hebbian learning is the accelerator, **homeostatic plasticity** is the thermostat. **Synaptic scaling** multiplies all of a neuron's synapses toward a target activity level, preserving their relative sizes while capping the total (slow, over days). **Metaplasticity** (the BCM sliding threshold) raises the bar for potentiation after a burst of activity. **Intrinsic plasticity** tunes a neuron's own excitability. Together they prevent both silence and seizure while leaving room to learn.

ANALOGY

Hebbian and homeostatic plasticity are the accelerator and the cruise control. Hebbian rules speed up whatever synapse is being used; homeostatic scaling keeps the whole network from redlining. Learning needs both, or it either flat-lines or explodes.



$$\Delta t = t_{\text{post}} - t_{\text{pre}} \text{ (ms)}$$

The STDP window. Fire just before the target and the synapse strengthens; fire just after and it weakens. The magnitude decays exponentially with the timing gap.

The Brain's Global Settings

Local plasticity decides *which* synapses change. But something has to decide *when* learning should happen and *how strongly*, and that is the job of **neuromodulators**, chemical signals broadcast widely rather than sent point to point. They are the brain's global control panel for learning, and each one has a clean computational counterpart in AI.

Dopamine, the reward signal. Released from the VTA and substantia nigra, phasic dopamine encodes a *reward prediction error*: the gap between the reward you got and the reward you expected. An unexpected reward gives a burst; an omitted one gives a dip. This is the biological twin of the temporal-difference error in reinforcement learning.

Acetylcholine, the attention knob. From the basal forebrain, it raises the signal-to-noise ratio, sharpen-

ing processing of relevant stimuli and enabling memory encoding, much like an attention mechanism or dropout in a neural network.

Norepinephrine, the alertness dial. From the locus coeruleus, it sets arousal and flags novelty, nudging the system to explore. In AI terms it tunes the exploration-exploitation balance and the effective learning rate.

Serotonin, patience. From the raphe nuclei, it shapes mood, impulse control, and how heavily the future is discounted, the biological echo of the discount factor in RL.

They interact. These systems are not independent dials; dopamine-acetylcholine and norepinephrine-serotonin balances jointly set the brain's learning state.

Neuromodulator	Source	Brain role	AI analog
Dopamine	VTA, substantia nigra	Reward prediction error	TD error , actor-critic
Acetylcholine	Basal forebrain	Attention, signal-to-noise	Attention, dropout
Norepinephrine	Locus coeruleus	Arousal, novelty	Exploration, learning rate
Serotonin	Raphe nuclei	Mood, patience, impulse control	Temporal discounting

Nudging the Circuit

If plasticity is the brain's learning dial, neurostimulation reaches in and turns it from outside. The techniques span a spectrum from cheap and non-invasive to surgical and precise, and the trade-off is always the same: reach and focality versus safety and accessibility.

tDCS (transcranial direct current stimulation) passes a weak current (1 to 2 mA) between scalp electrodes. *Anodal* stimulation gently depolarizes neurons and raises excitability; *cathodal* does the reverse. It does not fire neurons directly, but its LTP-like after-effects can outlast the session. Cheap and portable, but coarse and highly variable between people.

TMS (transcranial magnetic stimulation) uses a brief 1 to 2 tesla magnetic pulse to induce real currents that *can* trigger action potentials. Repetitive TMS at low

frequency (at or below 1 Hz) tends to inhibit; high frequency (5 Hz and up) tends to excite. It reaches about 1 to 2 cm with good timing precision and is FDA-approved for treatment-resistant depression.

DBS (deep brain stimulation) implants electrodes that deliver high-frequency pulses (130 to 180 Hz) directly into deep nuclei. It is dramatically effective for Parkinson's disease, essential tremor, and dystonia, and is approved for OCD, but it is a neurosurgical procedure with real risk.

Also worth knowing. tACS entrains brain oscillations at chosen frequencies; focused ultrasound (tFUS) reaches deep targets non-invasively with millimeter precision; VNS is approved for epilepsy and depression; and optogenetics uses light-sensitive opsins for millisecond, cell-type-specific control, still a research tool in animals.

Technique	How it works	Invasiveness	Focality	Status
tDCS	Weak DC 1–2 mA, anodal/-cathodal	Non-invasive	Coarse (cm)	Research; variable
TMS	Magnetic 1–2 T pulses	Non-invasive	1–2 cm	FDA depression
DBS	Implanted electrodes, 130–180 Hz	Surgical	Precise, deep	Parkinson's, OCD

Dopamine Is Temporal-Difference Learning

The crown jewel of NeuroAI is a near-perfect match: the phasic firing of dopamine neurons looks exactly like the *temporal-difference error* that drives reinforcement learning. The TD error compares what actually happened to what was predicted:

$$\delta = r + \gamma V(s') - V(s)$$

The correspondence runs deep. An **actor-critic** architecture maps onto the basal ganglia: a *critic* (ventral striatum) learns state values and computes the error, while an *actor* (dorsal striatum) picks actions, and dopamine broadcasts δ to both. Extend Hebb's rule with this signal and you get a **three-factor learning rule**, $\Delta w = \eta x_i y_j \delta$, where an *eligibility trace* tags recently active synapses so a delayed reward can still find them, solving the credit-assignment problem.

Learning Without Backprop, and Fine-Tuning


STDP is a *local* rule: it updates a weight using only signals available at that synapse, no global error to

propagate backward. That makes it attractive for neuromorphic hardware such as SpiNNaker and BrainScaleS, and homeostatic scaling offers a handle on catastrophic forgetting. Neuromodulation itself reads like meta-learning: cyclical and decaying learning-rate schedules, and adaptive optimizers like Adam, are engineered stand-ins for the brain's shifting learning states.

The tightest practical parallel is **fine-tuning**. A pre-trained model is like a developed brain; adapting it to a new task by nudging a few parameters, through adapters or **LoRA** (low-rank adaptation), is the AI version of neurostimulation enhancing an existing circuit rather than growing one from scratch.

ANALOGY

Fine-tuning a large model with LoRA is renovating a house, not rebuilding it. The structure (pre-trained weights) stays; you make small, targeted changes to a few rooms. Neurostimulation does the same to the brain, turning up plasticity in one circuit while leaving the rest intact.



The three-factor rule: Hebbian coincidence (pre times post) is gated by a global neuromodulatory signal (dopamine as reward-prediction error) before any weight changes.

Quick Self-Test

1. Why is pure Hebbian learning unstable, and what kind of mechanism counteracts it?
2. How does STDP let a network distinguish cause from effect?
3. What does phasic dopamine encode, and which reinforcement-learning quantity does it match?
4. Rank tDCS, TMS, and DBS by invasiveness and by spatial precision, and give one clinical use of each where applicable.
5. In what sense is fine-tuning a pre-trained model with LoRA analogous to neurostimulation?

If You Remember One Thing

The brain's learning is a balance of forces: Hebbian strengthening, homeostatic stabilizing, and neuromodulators gating when it all counts. The single deepest bridge to AI is that dopamine behaves like a temporal-difference

error, which means reinforcement learning is not just inspired by the brain, it may be describing what the brain is actually doing.

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