How the Brain Learns

From cells to memory and understanding

The question is not just “Which study technique works?” It is also “What changes inside me when something that was incomprehensible begins to make sense?” We will build the answer from the beginning, with no prior knowledge of biology required.

From the idea behind Euclid’s *Elements*, we keep the order: first name the objects, then explain their relationships. There is, however, an essential difference. Neuroscience is not derived from axioms. Its explanations are tested through observations and experiments; they can be extended or corrected.

Three words that do not mean the same thing

Learning: change resulting from experience that can affect future knowledge or behaviour. **Memory**: retaining, and being able to use, some effect of that experience. **Understanding**: here, we will test it through the ability to explain relationships, make predictions, and apply knowledge to a changed example. This is a criterion for assessment, not a complete philosophical definition.

The route through this book

2–5	What a neuron is made of, how it produces signals, and how it connects to others.
6–10	What can change with learning, what experiments show, and how we retrieve memories.
11–14	What makes an explanation difficult and what it means to acquire a usable model.
15–18	Which practices have experimental support, when caution is needed, and what remains unknown.

These numbers are also the pages of the main text. References [1], [2], and so on link to the sources at the end. The simple examples are our own teaching examples, not descriptions of experiments. The selection of studies is representative, not a systematic review of the entire field.

PART 2

The materials: cell, ion, membrane

A **cell** is a small living unit. Its interior is separated from its surroundings by a thin **membrane**. A **neuron** is a cell specialised in receiving, processing, and transmitting signals. It is not a little conscious entity that understands words.

In a common type of neuron, **dendrites** are branches that receive many inputs. The **cell body** contains essential machinery for maintaining the cell. The **axon** is an extension along which fast electrical signals travel. This is a useful starting picture, not an unbreakable blueprint for every neuron.



Where does the electricity come from?

Inside and outside the cell is fluid containing **ions**: small particles with an electrical charge. Sodium, Na^+, and potassium, K^+, are two positively charged examples. The movement of charged particles constitutes an electrical current. No metal wire is required.

The **membrane voltage** is the difference in electrical potential between the inside and outside. Voltage expresses a difference in electrical energy per unit of charge. If we write -70 mV, we mean that the inside is 70 thousandths of a volt more negative than the outside. This is an illustrative value, not a constant for every neuron.

Why does everything not equalise immediately?

The membrane does not let all ions pass freely. It contains **channels**: protein passageways selective for particular ions. Proteins are molecules whose structures enable them to perform particular functions. A channel can open or close.

An ion's movement depends both on the difference in its concentration across the membrane and on electrical force. These influences can reinforce or oppose each other. **Pumps** use energy to maintain unequal concentrations. The sodium–potassium pump is coupled to the consumption of ATP, a molecule that powers many cellular processes. The pump does not generate each individual fast electrical pulse.[2][3]

The step to keep. A neuron's electrical behaviour arises from physical things: charged particles, a membrane, channels, and differences in concentration.

PART 3

What does “firing” mean?

Firing means producing an **action potential**, or **spike**: a brief, rapid change in membrane voltage that can propagate along the axon. It does not mean “a thought has been created” or “a memory has been stored.”

A typical spike, in four steps

1. Some input changes the voltage. When the voltage becomes less negative, we call this **depolarisation**. A small depolarisation does not necessarily produce a spike.
2. If conditions reach a **threshold**, sodium entering through voltage-dependent channels causes further depolarisation. “Voltage-dependent” means that the channel’s state depends on the voltage. More channels open: the process temporarily reinforces itself.
3. Sodium channels inactivate, and potassium leaving the cell helps the voltage return towards negative values. The spike does not keep growing indefinitely.
4. For a short time, producing another spike is impossible or more difficult. This is called the **refractory period**. The ability to respond then recovers.

This mechanistic description was supported by measurements of ionic currents and a quantitative model, initially in the squid’s large axon. The threshold is not one universal number for all neurons.[1][2]

How does it travel, and where does it start?

Current from an activated section affects neighbouring sections of the membrane, where the spike can be regenerated. The same packet of ions does not travel from one end of the axon to the other.

In the brain, the trigger is often signals from other neurons; we will see next how these are transmitted. Sensory pathways convert external changes, such as light, into neural signals. But there is also internally generated activity: the brain does not passively wait for an external stimulus before it functions.

For a given cell under similar conditions, a stronger stimulus does not simply translate into a “taller spike.” It can change the number and timing of spikes. The number per second is called the **firing rate**.

The step to keep. A spike is an electrical event with an explainable mechanism. Its functional meaning depends on the cell, its connections, and its timing.

PART 4

How one neuron influences another

A **synapse** is a site of functional communication between cells. We will examine a chemical synapse between two neurons. The sending neuron is **presynaptic**, and the receiving neuron is **postsynaptic**. You do not need to memorise the names before seeing the process.

1. The spike reaches an ending of the first neuron's axon.
2. Channels open and calcium ions, Ca^{2+}, enter. Their entry helps release a chemical substance from small vesicles: membrane-bound "sacs."
3. The substance, the **neurotransmitter**, crosses the small gap between the cells and binds to **receptors**: proteins that respond to particular molecules.
4. Some receptors are themselves channels; others activate more indirect processes. This changes the behaviour of the second cell.

Calcium's importance for release, and the receiving cell's electrical response, were investigated experimentally in systems including the nerve–muscle junction. This simple experimental system allowed measurements that were harder to make inside the brain.[4][5]

Why does an incoming signal not always produce a spike?

Excitation tends to favour firing, whereas **inhibition** tends to make it harder. A neurotransmitter's effect depends on the receptors and the cell's conditions; its name alone is not enough. Inhibition does not always mean "the voltage becomes more negative": it can also limit the effect of other inputs.

A neuron combines inputs from many synapses. Where they arrive and how close together they occur in time affect the outcome. Imagine two weak excitatory inputs: if they arrive close enough together, they may jointly contribute to firing; if they are separated, the first effect may already have faded. This is a qualitative example, not a universal rule that "two inputs = one spike."

The step to keep. A neuron does not simply copy the previous neuron's message. Its output depends on many inputs, their timing relative to one another, and its own state.

PART 5

From signals to representations

A **neural circuit** is a set of connected neurons that influence one another. “Circuit” does not mean that the brain is constructed like a digital computer. Here it refers to the specific pathways of interaction.

When you see a spoon, different aspects of your experience include its shape, colour, name, and possible uses. These functions need not all be performed by the same cell. Nor is there a little observer inside your head looking at a second picture of the spoon.

What do we mean by “representation”?

A **neural representation** is a state or pattern of activity that is systematically related to something and can be used by the nervous system. “Pattern” might mean which neurons fire, how much, and when. It does not mean that words are written inside the cells.

The relationship can be distributed across many cells while some neurons are highly selective. **Selective** means that a neuron responds more under certain conditions than under others. It does not imply that this neuron alone contains the entire concept.

A human experiment

Ison and colleagues recorded individual neurons in patients with intracranial electrodes implanted for clinical reasons. While the patients learned new associations between pictures of people and places, some neural responses changed quickly: a neuron could now also respond to the picture associated with its originally preferred picture.[6]

This links a specific human learning experience to changes in recorded activity. It does not establish that “one person = one neuron” or that the recording reveals the person’s entire thought.

The step to keep. Information is not an extra substance inside a spike. It lies in the systematic relationship between activity, experience, and how the circuit uses that activity.

PART 6

What changes when we learn?

Activity is what a circuit is doing now. **Plasticity** is its capacity to change how it will respond later. A temporary signal and a lasting change are not the same thing. This is the first essential bridge from firing to memory.

A connection can change its effectiveness

Let us call two neurons A and B. If, under comparable conditions, the same activation of A produces a larger synaptic response in B after a learning process, we say that the connection has **strengthened**. This does not necessarily mean that a new neuron appeared or that A's spike became larger.

Long-term potentiation, LTP, is a prolonged increase in the effectiveness of synaptic transmission. **Long-term depression**, LTD, is a prolonged decrease. "Long-term" here contrasts with the momentary signal; it does not guarantee that a change lasts a lifetime.

Bliss and Lømo produced prolonged enhancement of responses in a region of rabbits' hippocampus following a particular electrical stimulation. The **hippocampus** is a brain structure deep within the temporal lobes; we will examine its role in memory on page 8.[7]

What can physically change?

One possibility is a change in the number or functioning of receptors at the synapse. **AMPA** and **NMDA** name two types of receptor for the neurotransmitter glutamate. At some excitatory synapses, the NMDA channel is usually blocked by a magnesium ion when the voltage is sufficiently negative. Depolarisation helps relieve this block. This combines the presence of transmitter from A with the electrical state of B.[28]

The calcium that can then enter activates proteins inside the cell. Depending on the conditions, these reactions can increase the presence of AMPA receptors at the synapse. Their movement has been observed experimentally: the same future input can now produce a larger response.[8] This explains how combined activity can selectively affect a connection; it does not mean that every instance of simultaneous firing strengthens it.

This is an **example of a mechanism**, not a recipe for every kind of memory. We do not learn only by "strengthening everything." Selective changes are needed, and decreases in responsiveness can also be functionally important.

The step to keep. The past can affect the future because experience changes a circuit's properties. "The synapses strengthened" always needs clarification: which ones, by what mechanism, and with what evidence?

PART 7

How do we know a change concerns memory?

When two events occur together, that is a **correlation**. To test whether one causally contributes to the other, we intervene in the first and examine what changes in the second. We also need suitable comparison conditions, so that the difference cannot be attributed to something unrelated.

Intervening in synaptic strength

Nabavi and colleagues paired artificial activation of a neural pathway with an unpleasant stimulus in **rats**. Subsequently, inducing LTD in that pathway reduced the expression of the learned response, while inducing LTP restored it. The study provides causal support for the involvement of synaptic changes in this particular learned behaviour. It did not demonstrate erasing and rewriting a complex human memory.[9]

Intervening in the cells of an experience

A **memory trace**, or **engram**, refers to the physical changes and cell ensembles involved in retaining and re-expressing a memory. It is a research concept, not a little file in one unique location.

Optogenetics allows selected cells to be influenced with light, using genetically introduced light-sensitive proteins. Liu and colleagues labelled neurons active during fear learning in **mice**. Activating these neurons later could produce freezing in a different environment, a behavioural indication of fear-memory retrieval.[10]

The finding and its limits

Producing a behaviour by activating cells shows that these cells can, under the experimental conditions, contribute to its expression. It does not automatically show that they are the only cells involved or that we know the animal's subjective experience.

A different question is **necessity**: if we silence an ensemble, does retrieval stop? Even then, we must rule out explanations such as a general inability to move. Observing "it did not respond" is not enough to conclude "it does not remember."

The step to keep. Experiments causally connect cells and synapses with particular memories. The leap from these findings to a complete explanation of human understanding remains very large.

PART 8

Memory is not one single thing

Consider three different statements: “I remember yesterday’s conversation,” “I know what a triangle means,” and “I know how to perform a practised movement.” All three concern effects of previous experience, but they do not require exactly the same functions.

Name	What it distinguishes
Episodic memory	Remembering an event as an experience in a context: what happened, where, and when.
Semantic memory	Knowledge of concepts and facts, without having to recall the particular moment when you learned them.
Procedural learning	Acquiring skills and ways of performing actions. This is not the same as being able to describe every step verbally.
Working memory	Temporarily retaining and using information for what you are doing now; we will examine it on page 11.

These are useful functional distinctions, not four independent “drawers.” A task can require more than one.

What have brain injuries taught us?

The classic study by Scoville and Milner described severe difficulty forming new lasting memories after bilateral operations in the medial temporal lobes. The well-known patient H.M. did not simply lose every mental ability or his entire past. The operations also did not affect the hippocampus alone. These observations help distinguish new learning from other functions.[11]

A useful way to think about the hippocampus’s contribution is connecting elements of a new experience: the person, the place, the temporal relationship. The human study on page 5 examines one such new connection.[6]

It would be wrong to imagine that the hippocampus “understands” while the rest of the brain merely stores. Nor does every everyday difficulty retrieving something while studying indicate a problem in this structure.

The step to keep. “I remember” can describe different abilities. Before looking for a mechanism, clarify exactly which ability you are examining.

PART 9

What happens after studying?

Consolidation refers to processes through which a memory stabilises or reorganises over time. There is not necessarily a moment when a temporary file is transferred intact to a permanent drive.

At the synaptic level, we are interested in maintaining and stabilising changes. At the level of brain systems, we are interested in how the contributions of different regions change. This distinction describes levels of investigation, not two completely separate processes.

Stabilisation can require the production of new proteins: in experiments on hippocampal tissue, inhibiting protein synthesis limited late phases of LTP.[29] It can also involve structural changes. **Dendritic spines** are small protrusions on dendrites where many excitatory synapses are located. In mice, learning was associated with the formation and selective retention of such protrusions.[30] We do not conclude that one spine is an entire memory or that every molecule in it remains the same forever.

Why are we interested in sleep?

Wilson and McNaughton recorded hippocampal neurons in rats. Relationships of co-activation that had appeared during exploration of a space reappeared during subsequent sleep. This is evidence of **reactivation**: a previous pattern of activity returns. The observation alone does not establish that reactivation caused better memory.[12]

In a human experiment by Rasch and colleagues, an odour was associated with learning object locations. Re-exposure to that odour during slow-wave sleep improved later retention on this task. This sleep stage is characterised by slow changes in recorded electrical activity. The intervention provides stronger evidence of a causal contribution than merely observing co-activation.[13]

What we should not conclude

These findings do not mean that sleep teaches you an unstudied subject from scratch. They do not imply that you should use odours or sounds on your own, or that every sleep stage has the same function for every kind of learning.

“The brain stores things when you sleep” is too crude. Learning and retention involve changes during both wakefulness and sleep. The precise relationship between reactivation, stability, and changes in memories remains an active research topic.

The step to keep. Biologically, a learning experience does not end the moment you close the PDF. This does not make sleep a substitute for initial understanding and practice.

PART 10

Retrieval: returning, not replaying a file

Retrieval is drawing on previous learning when you need it. A **retrieval cue** is something that helps activate the relevant knowledge: a question, a word, or part of a problem.

In **recognition**, a possible answer is in front of you: "I have seen this." In free recall, you must produce it without being shown it. Success under one condition does not guarantee success under the other.

Why does a hint help?

Suppose you cannot retrieve a concept, but remember it as soon as you see the diagram you studied it with. This does not prove that the information was entirely absent before the hint. It shows that the conditions of access matter. Research on encoding specificity examines precisely the relationship between how we learn something and the cues that later allow us to retrieve it. **Encoding** here means the processes of initial learning, not translating a sentence into some known digital code.[27]

A failure can reflect insufficient initial learning, difficulty accessing what was learned, or competition from other knowledge. **Interference** is when other memories make the desired retrieval harder. A single instance of "I got stuck" does not let us identify the responsible mechanism with certainty.

Why can a memory be convincing but wrong?

Roediger and McDermott showed that, after hearing related words, people could confidently report another related word that had not been presented. Semantic connections help organisation but can also lead to errors. This does not mean that all memories are unreliable.[14]

For studying, this gives you a simple reason to check your answer against the source. Sounding logical and feeling familiar are not sufficient tests of accuracy.

The step to keep. "I cannot bring it to mind now," "I never learned it," and "I remember it incorrectly" are different situations. They need different responses.

PART 11

Why unfamiliar terms block you

Attention concerns the selective priority given to particular information or actions. **Working memory** concerns the information you temporarily retain and use. They are not synonyms, although they work together.

When you read “A updates B before C replaces D,” you need to keep track of what each thing is and which relationship is being discussed. If none of the objects has been defined, you are not just missing a word: you are missing the very means of following what is happening.

The limit is not a magic number

In experiments on visual working memory, Luck and Vogel found a limited capacity to retain objects over very short periods. The result also depended on how features were grouped into objects. It does not imply a universal rule that “the brain holds exactly four ideas.” The tasks, materials, and prior knowledge matter.[15]

Knowledge changes what counts as one piece

A **chunk** is a set of elements you treat as one familiar unit. If you already know what a chemical synapse is, the phrase “synaptic transmission” can bring to mind a whole sequence. On your first reading, however, you must separately process the cell, the ending, the transmitter, and the receptor.

Chase and Simon studied the reconstruction of chess positions by players with different levels of experience. Domain-specific knowledge helped especially when positions had chess-like structure, and less when pieces were placed randomly. This small classic study does not establish that experts have universally superior memory.[16]

That is why a short slide can be clear to its author and opaque to a beginner. The author already has the connections they leave out. The reader is being asked both to invent those connections and to understand the new argument.

The step to keep. Brevity is not the same as simplicity. A longer explanation that defines the roles can require less mental effort than three compressed, ambiguous lines.

PART 12

What does “now it makes sense” mean?

A **mental model** is an internal, simplified representation of how the parts of a situation relate. It is not necessarily a picture. It can include rules such as “if this step is missing, that outcome will not follow.”

A **schema** is an organised structure of prior knowledge that helps you interpret new cases. A model of a particular situation can draw on such schemas. These terms describe functions; they do not point to one unique location in the brain.

First, I need to know what we are talking about

Read our deliberately vague sentence:

> “If it is there, we use it; otherwise we bring it from elsewhere, and perhaps remove something that was there before.”

Every word is simple. Yet the object, the two locations, and the objective are missing. You do not necessarily have a memory or concentration problem: the explanation itself is inadequate. On the next page, we will name everything and see whether we can predict the system’s behaviour.

In experiments by Bransford and Johnson, suitable context provided *before* a difficult-to-understand text improved comprehension and recall relative to conditions without that useful initial context. Providing the context later did not simply produce the same result.[17]

This offers a plausible interpretation of the feeling that a text is “hiding something”: perhaps it relies on a context the author never gave you. It is not a diagnosis of your difficulty or an experiment on Foundations.

The feeling of understanding needs checking

“It sounds clear” is useful, but not sufficient. Better evidence is whether you can name the parts, explain their relationships, and predict what will change if one condition changes. You can also have a coherent but incorrect model; that is why predictions must be checked against reality or the source.

The step to keep. An explanation should reveal the objects and their relationships, not force you to guess which unnamed thing it means.

PART 13

A complete small example

Let us build a simple library system. No knowledge of computers is required. We have a **storeroom** that initially contains all the books and a **table** near the librarian that holds only two. We want to serve requests quickly. Accessing the storeroom takes longer.

First, define the rules precisely

1. If the requested book is on the table, we use it there.
2. If it is missing and there is space, we bring it from the storeroom, put it on the table, and use it.
3. If it is missing and the table is full, first we return to the storeroom the book that has gone longest without being used. Then we bring, place, and use the requested book. Initially, the table is empty.

This is an analogy for a **cache**: a smaller, faster space that holds selected data from a larger space. Finding the requested item there is a **hit**; its absence is a **miss**. The third rule is a particular replacement rule, not a necessary feature of every cache.

Follow the rules rather than memorising the result

A, B, and C name three different books.

Request	Table afterwards	Why
A	A	Missing: we fetch it.
B	A, B	Missing: there is a free place.
A	A, B	Already present. B has now gone longer without use.
C	A, C	Missing, and there is no space. We remove B.

Where does understanding show up?

If B is requested now, a trip to the storeroom will be needed. If the table held three books, B would not have been removed in the last step. If we used a different replacement rule, the final state would not necessarily be the same. And if every request is for a different book, the table does not guarantee many hits.

The step to keep. You can now generate answers because you know the objects, the objective, and the rules. This example illustrates the structure of good teaching; it does not claim that human memory operates as a cache.

PART 14

How an explanation becomes your knowledge

A **worked example** shows not only the result but the route to it. For a beginner, it can supply an initial structure, so that the problem, the rules, and the solution strategy do not all have to be discovered at once.

Sweller and Cooper studied worked examples in learning algebraic manipulation. Their results supported the usefulness of examples during the initial acquisition of these skills. They do not imply that continually reading solutions replaces independent problem solving.[18]

Active explanation is not just repeating words

Self-explanation means explaining to yourself why a step follows, under which rule, and under which conditions. In our example:

“B is removed” states a result. “B is removed because the rule considers time since the last *use*, and A was just used” connects the result to the rule.

In a study by Chi and colleagues, students prompted to produce explanations while reading about the circulatory system improved more than a group that read the text twice. Complex comprehension questions were also examined. It was a small study with specific material, not a guarantee for every student and every text.[19]

A gradual reduction in support

The following sequence is a teaching proposal for our work, not a uniquely optimal protocol established by experiment:

1. You see a complete example with explicitly named objects.
2. You explain one crucial step without copying the sentence.
3. You complete a similar example with one step missing.
4. You solve a new example without having the solution in front of you.
5. You check the result and specifically correct the mistaken rule.

If the second step is impossible because you do not know what a term means, we return to the definition. We do not relabel confusion as “desirable difficulty.”

The step to keep. The aim is not for you to need the author’s explanation forever. It is to acquire the structure that lets you reconstruct the steps yourself.

PART 15

Why attempting retrieval helps

We have already distinguished retrieval from recognition. Now we add an experimental finding: attempting retrieval does not merely measure what you learned. It can also affect what you will remember later.

The comparison that matters

Roediger and Karpicke compared repeated study of texts with study followed by retrieval attempts. After a very short delay, additional study could perform better. On tests days later, retrieval performed better, even though rereading produced greater confidence.[20]

So “how easy it feels now” is not the same as “how well I will use it in a week.” This does not mean reading is useless: without sufficient initial access to meaning, you have nothing to retrieve.

It is not just about words and dates

Karpicke and Blunt found a benefit of retrieval practice with science texts, even on questions requiring inferences, compared with the particular concept-mapping activity they investigated. This does not condemn all concept maps or equate understanding with rote memorisation.[21]

For our small example, close page 13 and answer: “Which three cases do we distinguish when a book is requested?” Then check a different sequence of requests. The first task tests recall of the rules; the second tests whether you can use them.

What is the precise mechanism?

Possible explanations include improved future access to knowledge and reorganisation of its relationships. The behavioural result does not, by itself, identify which synapse changed or prove one unique mechanism. “Retrieval works because it causes LTP” would be a stronger claim than these educational experiments measure.

If you give a wrong answer, you need to compare it with a correct model. Propose a correction that identifies the mistake: “I confused the last insertion with the last use.” That is more useful information than the general statement “I don’t get it.”

The step to keep. Use retrieval both for learning and for checking learning. But check accuracy, not just how easily you produce an answer.

PART 16

Intervals and different examples

Spacing means separating opportunities to study over time rather than concentrating them in one session. **Interleaving** means mixing different types of problems. These are different variables: you can apply one without the other.

How far apart should reviews be?

Cepeda and colleagues varied both the interval before a review and the time until the final test. They showed that effective spacing depends on how much later we want to remember. Their study does not yield one universal schedule of review days.[22]

A practical starting proposal, then, is to return to the knowledge in later sessions and adjust the spacing to your performance. If every return requires learning almost from scratch, you may need a shorter interval or better initial teaching. This is a way to regulate studying, not a diagnosis of particular changes in the brain.

Why mix problems?

If the heading says “apply rule X” above twenty exercises, you practise executing the rule, but do less practice selecting it. In an exam, you often first have to decide which method fits.

Rohrer and Taylor separately investigated spacing practice and mixing types of mathematics problems. In their tasks, the respective alternatives performed better on a test one week later than massed or blocked practice. This does not mean you should mix unfamiliar topics before understanding even one example of each.[23]

Useful difficulty has prerequisites

Difficulty can help when it makes you retrieve knowledge or distinguish relevant cases. It is not useful simply because you are struggling. If a definition is missing, a question is ambiguous, or you receive no correction, the effort may not lead you to the intended rule.

The step to keep. Separate two goals: first learn how a method works; then recognise when it is needed and remember it later.

PART 17

Dopamine, surprise, and explanatory limits

So far, we have examined signal transmission and synaptic changes. Some chemicals can also alter how circuits respond and when changes are favoured. This is called **neuromodulation**. **Dopamine** participates in such functions; it is not simply "the pleasure chemical" or "the learning chemical."

Expected and actual outcomes

A **reward prediction error** is the difference between an expected and an observed reward. In a simple model, if you expected little and received more, the difference is positive. If an expected reward is absent, it is negative. This is a specific research term, not every mistake you make in an exercise.

Steinberg and colleagues showed in rats that selectively timed activation of dopamine neurons could drive learning of reward associations in a way consistent with a prediction-error signal. They did not show that increasing dopamine generally improves learning in every subject.[24]

For our studying, "predict, observe the outcome, explain the difference" is a reasonable teaching practice. We need not claim that we know the particular dopamine signal it produces in you.

The field did not stop at "stronger synapses"

Intrinsic excitability is how readily the cell itself generates spikes when it receives input. A 2026 study in mice examined temporary changes in this property in cells associated with fear memory, and their contribution to its formation.[25]

Glial cells are not neurons; they include **astrocytes**. In 2025 experiments in mice, particular ensembles of astrocytes participated in stabilising fear-related memory after re-exposure. They are not merely inert supporting material.[26]

The step to keep. We have specific mechanisms, but no single substance or switch that explains all learning. New findings do not automatically become new recipes for studying.

PART 18

What we can now say precisely

What causes firing? Ionic currents and membrane properties, under the influence of inputs and the cell's state. We explained this on pages 2–4 without requiring the neuron to "know" what it is doing.

What makes memory possible? Experience can leave lasting changes in synapses, cells, and circuits, so that a later input has a different effect. We have causal experiments for some forms of memory, not a complete decoding of every human memory.

What makes a concept understandable? At the functional level, acquiring relationships that let you explain and predict. At the neural level, this requires the operation and modification of many systems. We do not have a complete mechanistic map explaining, from every molecule to your experience, the moment when you understand a particular proof.

What this means for our work

The following is a design proposal, not an experimentally validated package or a guarantee of a particular grade:

1. **Name and object.** What are we talking about? What acts on what? Every new term should refer to something defined.
2. **Problem and mechanism.** Why do we need the concept? Which steps lead from the initial situation to the outcome?
3. **One complete example.** No hidden intermediate steps. Then change a condition and examine a case where the rule does not apply.
4. **Independent production.** Close the text, explain, and solve a new case. We check both the result and the reasoning.
5. **A later check.** Return after an interval. Record what you can do, not just how long you studied.

If a text already fails at the first step, the solution is not automatically more repetitions of the same unclear text. A better explanation is needed. If, instead, you explain correctly but forget later, the focus shifts to retrieval and retention.

The step to keep. Neuroscience explains parts of the biological mechanism. Human learning experiments test educational choices. We need both, without presenting either as complete proof of the other.

Sources and limits of the evidence

The following sources are primary research publications. Links lead to the publication or a copy of the original article. The age of a classic experiment does not, by itself, invalidate it; nor is a recent paper the final answer. Technical definitions and teaching examples are expressed in our own words, not as translated excerpts from the sources. References are grouped by theme; the final page adds more specific sources on the stabilisation of synaptic changes.

[1] Hodgkin, A. L. & Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve.

The Journal of Physiology 117, 500–544 · DOI: 10.1113/jphysiol.1952.sp004764

Measurements and a model in the squid axon. Supports the explanation of the action potential, not a theory of how mathematics is understood.

[2] Hodgkin, A. L. & Katz, B. (1949). The effect of sodium ions on the electrical activity of the giant axon of the squid.

The Journal of Physiology 108, 37–77 · DOI: 10.1113/jphysiol.1949.sp004310

Changes in external sodium and electrical activity. An experimental test of the signal's ionic basis in a particular biological preparation.

[3] Skou, J. C. (1957). The influence of some cations on an adenosine triphosphatase from peripheral nerves.

Biochimica et Biophysica Acta 23, 394–401 · DOI: 10.1016/0006-3002(57)90343-8

Biochemical experiments on crustacean nerve tissue; a foundation for investigating the sodium–potassium pump. Does not measure learning or memory.

[4] Fatt, P. & Katz, B. (1951). An analysis of the end-plate potential recorded with an intra-cellular electrode.

The Journal of Physiology 115, 320–370 · DOI: 10.1113/jphysiol.1951.sp004675

Intracellular recording at the neuromuscular junction. Useful for distinguishing a synaptic response from an action potential; not a recording inside a human brain.

[5] Katz, B. & Miledi, R. (1967). Ionic requirements of synaptic transmitter release.

Nature 215, 651 · DOI: 10.1038/215651a0

Experimental investigation of the ionic requirements for transmitter release, especially calcium. Does not, on its own, explain every type of synapse.

[6] Ison, M. J., Quian Quiroga, R. & Fried, I. (2015). Rapid Encoding of New Memories by Individual Neurons in the Human Brain.

Neuron 87, 220–230 · DOI: 10.1016/j.neuron.2015.06.016

Human intracranial recording during association learning. Selectivity of recorded cells does not establish that a concept is stored uniquely in one cell.

Sources · plasticity and memory

[7] Bliss, T. V. P. & Lømo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path.

The Journal of Physiology 232, 331–356 · DOI: 10.1113/jphysiol.1973.sp010273

Rabbits, electrical stimulation, and prolonged enhancement of responses. Observing LTP alone does not establish storage of a particular memory.

[8] Shi, S.-H. et al. (1999). Rapid Spine Delivery and Redistribution of AMPA Receptors After Synaptic NMDA Receptor Activation.

Science 284, 1811–1816 · DOI: 10.1126/science.284.5421.1811

Cellular imaging of receptor movement in a hippocampal preparation. An example of a physical mechanism of synaptic change, not a universal theory of learning.

[9] Nabavi, S. et al. (2014). Engineering a memory with LTD and LTP.

Nature 511, 348–352 · DOI: 10.1038/nature13294

Rats, a selected neural pathway, and a conditioned response. A causal intervention in synaptic strength, not erasure of human autobiographical memory.

[10] Liu, X. et al. (2012). Optogenetic stimulation of a hippocampal engram activates fear memory recall.

Nature 484, 381–385 · DOI: 10.1038/nature11028

Mice, cell labelling, and optogenetic activation. The behavioural measure is freezing, not a verbal report of subjective experience.

[11] Scoville, W. B. & Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions.

Journal of Neurology, Neurosurgery & Psychiatry 20, 11–21 · DOI: 10.1136/jnnp.20.1.11

Human clinical cases. The operations involved broader regions; the study does not isolate each structure as a selective animal experiment might.

[12] Wilson, M. A. & McNaughton, B. L. (1994). Reactivation of hippocampal ensemble memories during sleep.

Science 265, 676–679 · DOI: 10.1126/science.8036517

Recordings in rats. Reappearance of activity correlations during sleep; an observation requiring additional interventions for a strong causal conclusion.

[13] Rasch, B., Büchel, C., Gais, S. & Born, J. (2007). Odor cues during slow-wave sleep prompt declarative memory consolidation.

Science 315, 1426–1429 · DOI: 10.1126/science.1138581

A human reactivation experiment using an odour. The benefit concerned a particular already-learned task; it does not establish learning a new subject during sleep.

Sources · retrieval and understanding

[14] Roediger, H. L. & McDermott, K. B. (1995). Creating false memories: Remembering words not presented in lists.

Journal of Experimental Psychology: Learning, Memory, and Cognition 21, 803–814 · DOI: 10.1037/0278-7393.21.4.803

Human word-list experiments. Demonstrate particular recall errors, not that every confidently held memory is fabricated.

[15] Luck, S. J. & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions.

Nature 390, 279–281 · DOI: 10.1038/36846

Brief visual retention of features and objects. The limits in this task do not transfer unchanged to every concept or study activity.

[16] Chase, W. G. & Simon, H. A. (1973). Perception in chess.

Cognitive Psychology 4, 55–81 · DOI: 10.1016/0010-0285(73)90004-2

A small study of players with different experience. Supports the importance of structure and domain knowledge, not universally superior memory in all experts.

[17] Bransford, J. D. & Johnson, M. K. (1972). Contextual Prerequisites for Understanding: Some Investigations of Comprehension and Recall.

Journal of Verbal Learning and Verbal Behavior 11, 717–726 · Full original article

Experimental variation of context available before or after a text. Relevant to unclear references, but not an evaluation of our own PDFs.

[18] Sweller, J. & Cooper, G. A. (1985). The Use of Worked Examples as a Substitute for Problem Solving in Learning Algebra.

Cognition and Instruction 2, 59–89 · DOI: 10.1207/s1532690xci0201_3

A series of algebra-learning experiments. Initial support from worked examples is not an argument against later independent practice.

[19] Chi, M. T. H., de Leeuw, N., Chiu, M.-H. & LaVancher, C. (1994). Eliciting Self-Explanations Improves Understanding.

Cognitive Science 18, 439–477 · DOI: 10.1207/s15516709cog1803_3

14 students prompted to self-explain and 10 in the comparison group; a text on circulation. Understanding was also measured, but the small sample limits generalisation.

[20] Roediger, H. L. & Karpicke, J. D. (2006). Test-Enhanced Learning: Taking Memory Tests Improves Long-Term Retention.

Psychological Science 17, 249–255 · DOI: 10.1111/j.1467-9280.2006.01693.x

Human text-retrieval experiments. Different outcomes on immediate and delayed tests; no guarantee of transfer to every examination task.

Sources · practices and open questions

[21] Karpicke, J. D. & Blunt, J. R. (2011). Retrieval practice produces more learning than elaborative studying with concept mapping.

Science 331, 772–775 · DOI: 10.1126/science.1199327

A human experiment using science texts and inference questions. The comparison concerns the particular implementations of the two activities.

[22] Cepeda, N. J., Vul, E., Rohrer, D., Wixted, J. T. & Pashler, H. (2008). Spacing effects in learning: A temporal ridgeline of optimal retention.

Psychological Science 19, 1095–1102 · DOI: 10.1111/j.1467-9280.2008.02209.x

Fact-learning experiments with varying time intervals. They do not provide one common optimal review schedule for every subject.

[23] Rohrer, D. & Taylor, K. (2007). The shuffling of mathematics problems improves learning.

Instructional Science 35, 481–498 · DOI: 10.1007/s11251-007-9015-8

Two mathematics experiments: distributing practice over time and mixing problem types. The outcome depends on the material and phase of learning.

[24] Steinberg, E. E. et al. (2013). A causal link between prediction errors, dopamine neurons and learning.

Nature Neuroscience 16, 966–973 · DOI: 10.1038/nn.3413

Rats and optogenetic intervention in reward learning. Not a study of reading or evidence of benefits from “more dopamine.”

[25] Hadzibegovic, S. et al. (2026). Early intrinsic excitability plasticity of neocortical engram neurons defines memory formation and precision.

Nature Communications 17, article 291 · DOI: 10.1038/s41467-025-66975-3

Mice, fear memory, and transient changes in cellular excitability. Recent evidence of an additional mechanism, not a complete theory of human memory.

[26] The astrocytic ensemble acts as a multiday trace to stabilize memory (2025).

Nature 648, 146–156 · DOI: 10.1038/s41586-025-09619-2

Mice, astrocytic recordings, and interventions in fear memory. Extends the picture beyond neurons, without establishing a new study technique.

[27] Tulving, E. & Thomson, D. M. (1973). Encoding specificity and retrieval processes in episodic memory.

Psychological Review 80, 352–373 · DOI: 10.1037/h0020071

Human experiments and theoretical analysis of the encoding–retrieval relationship. Does not mean you must reproduce the same room or every detail of studying.

Sources · from a signal to a lasting change

[28] Nowak, L., Bregestovski, P., Ascher, P., Herbet, A. & Prochiantz, A. (1984). Magnesium gates glutamate-activated channels in mouse central neurones.

Nature 307, 462–465 · DOI: 10.1038/307462a0

Electrophysiological experiments in mouse neurons. The voltage dependence of the magnesium block helps explain why a receptor can combine a chemical signal with the cell's electrical state. This does not mean the experiment investigated human understanding, or that two requirements exhaust all the receptor's molecular regulation.

[29] Frey, U., Krug, M., Reymann, K. G. & Matthies, H. (1988). Anisomycin, an inhibitor of protein synthesis, blocks late phases of LTP phenomena in the hippocampal CA1 region in vitro.

Brain Research 452, 57–65 · DOI: 10.1016/0006-8993(88)90008-X

A pharmacological intervention in a hippocampal preparation. Inhibiting protein synthesis affected late phases of potentiation. Supports a distinction between early and later maintenance mechanisms; does not prove that every human memory requires the same proteins or the same timescale.

[30] Yang, G., Pan, F. & Gan, W.-B. (2009). Stably maintained dendritic spines are associated with lifelong memories.

Nature 462, 920–924 · DOI: 10.1038/nature08577

Tracking dendritic spines in mouse cortex after learning and sensory experience. Remodelling and selective retention were associated with behaviour. This is a structural correlation and support for a mechanism, not a mapping of one spine to one complete memory.

How to read the evidence

References support the particular findings described, not a certification that every broader interpretation has been proved. Animal experiments and isolated-tissue studies are identified as such. Proposed study practices are distinguished from descriptions of experiments. For learning practices, more weight is given to the direct human comparisons in sources [17]–[23] than to an analogy with a cellular experiment.