Master in Planning and
Management of Tourism Systems

Neural Networks

Nicola Cortesi

IMPORTANT NEWS

Last week Anthropic refused to give their LLM, Claude, to the Pentagon to use it for aquiring military targets without human supervision, and as a tool of mass surveillance of all US citizens

OpenAI and xAI instead did not raise any objection to use ChatGPT/Grok for these nefarious purposes.

For this reason, I suggest you to uninstall ChatGPT and Grok from your smartphones and to install Claude instead, as many users already did.

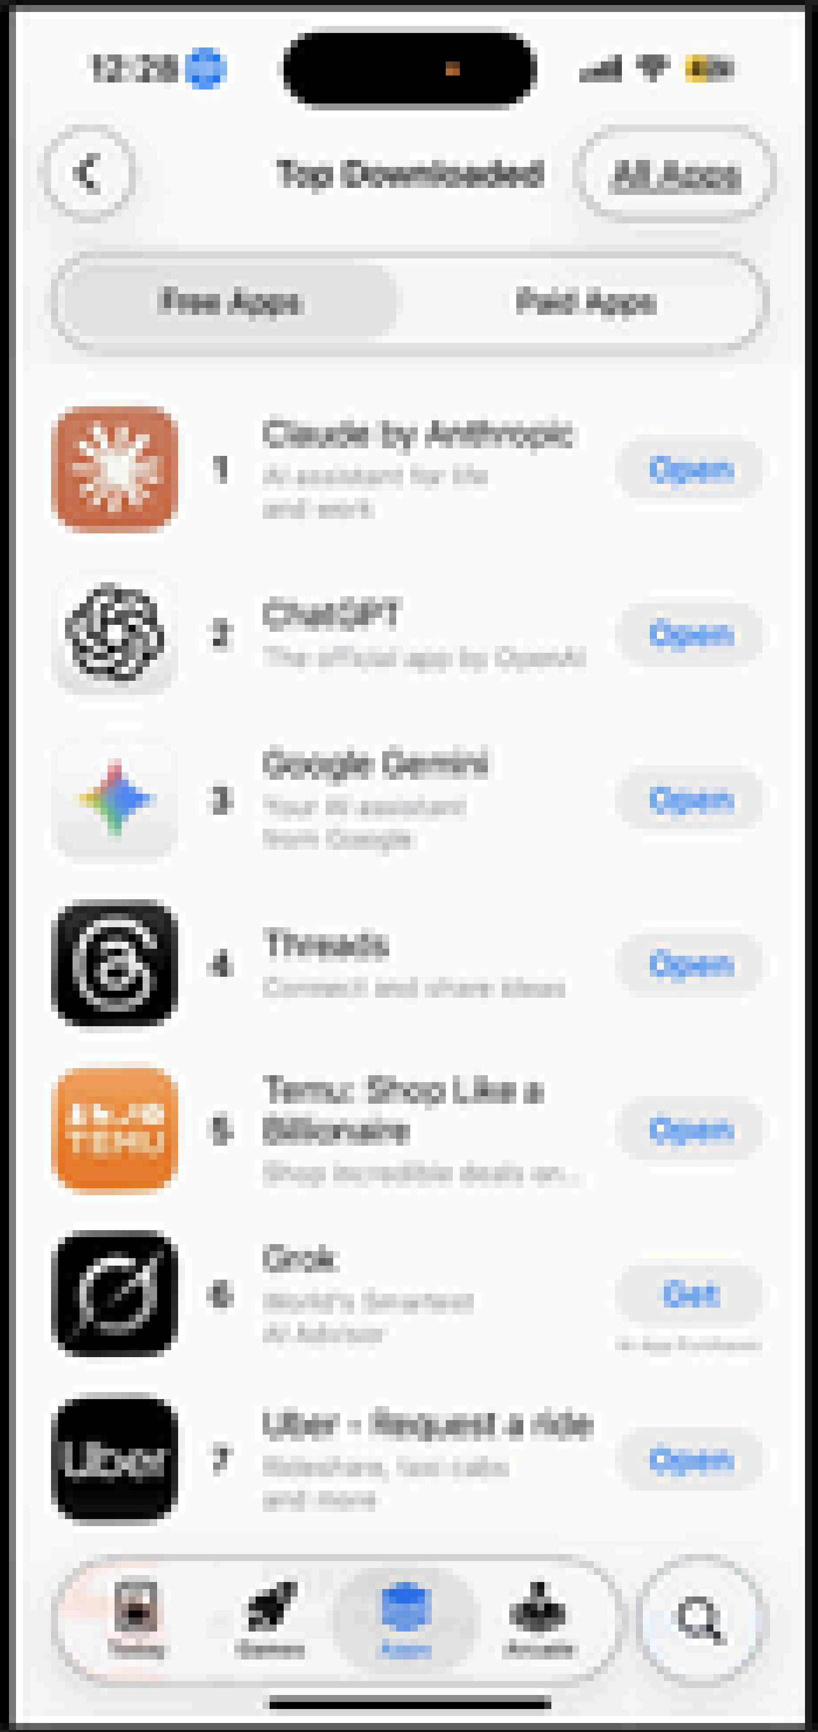

WE SIMULATED THE LEFT BRAIN HEMISPHERE

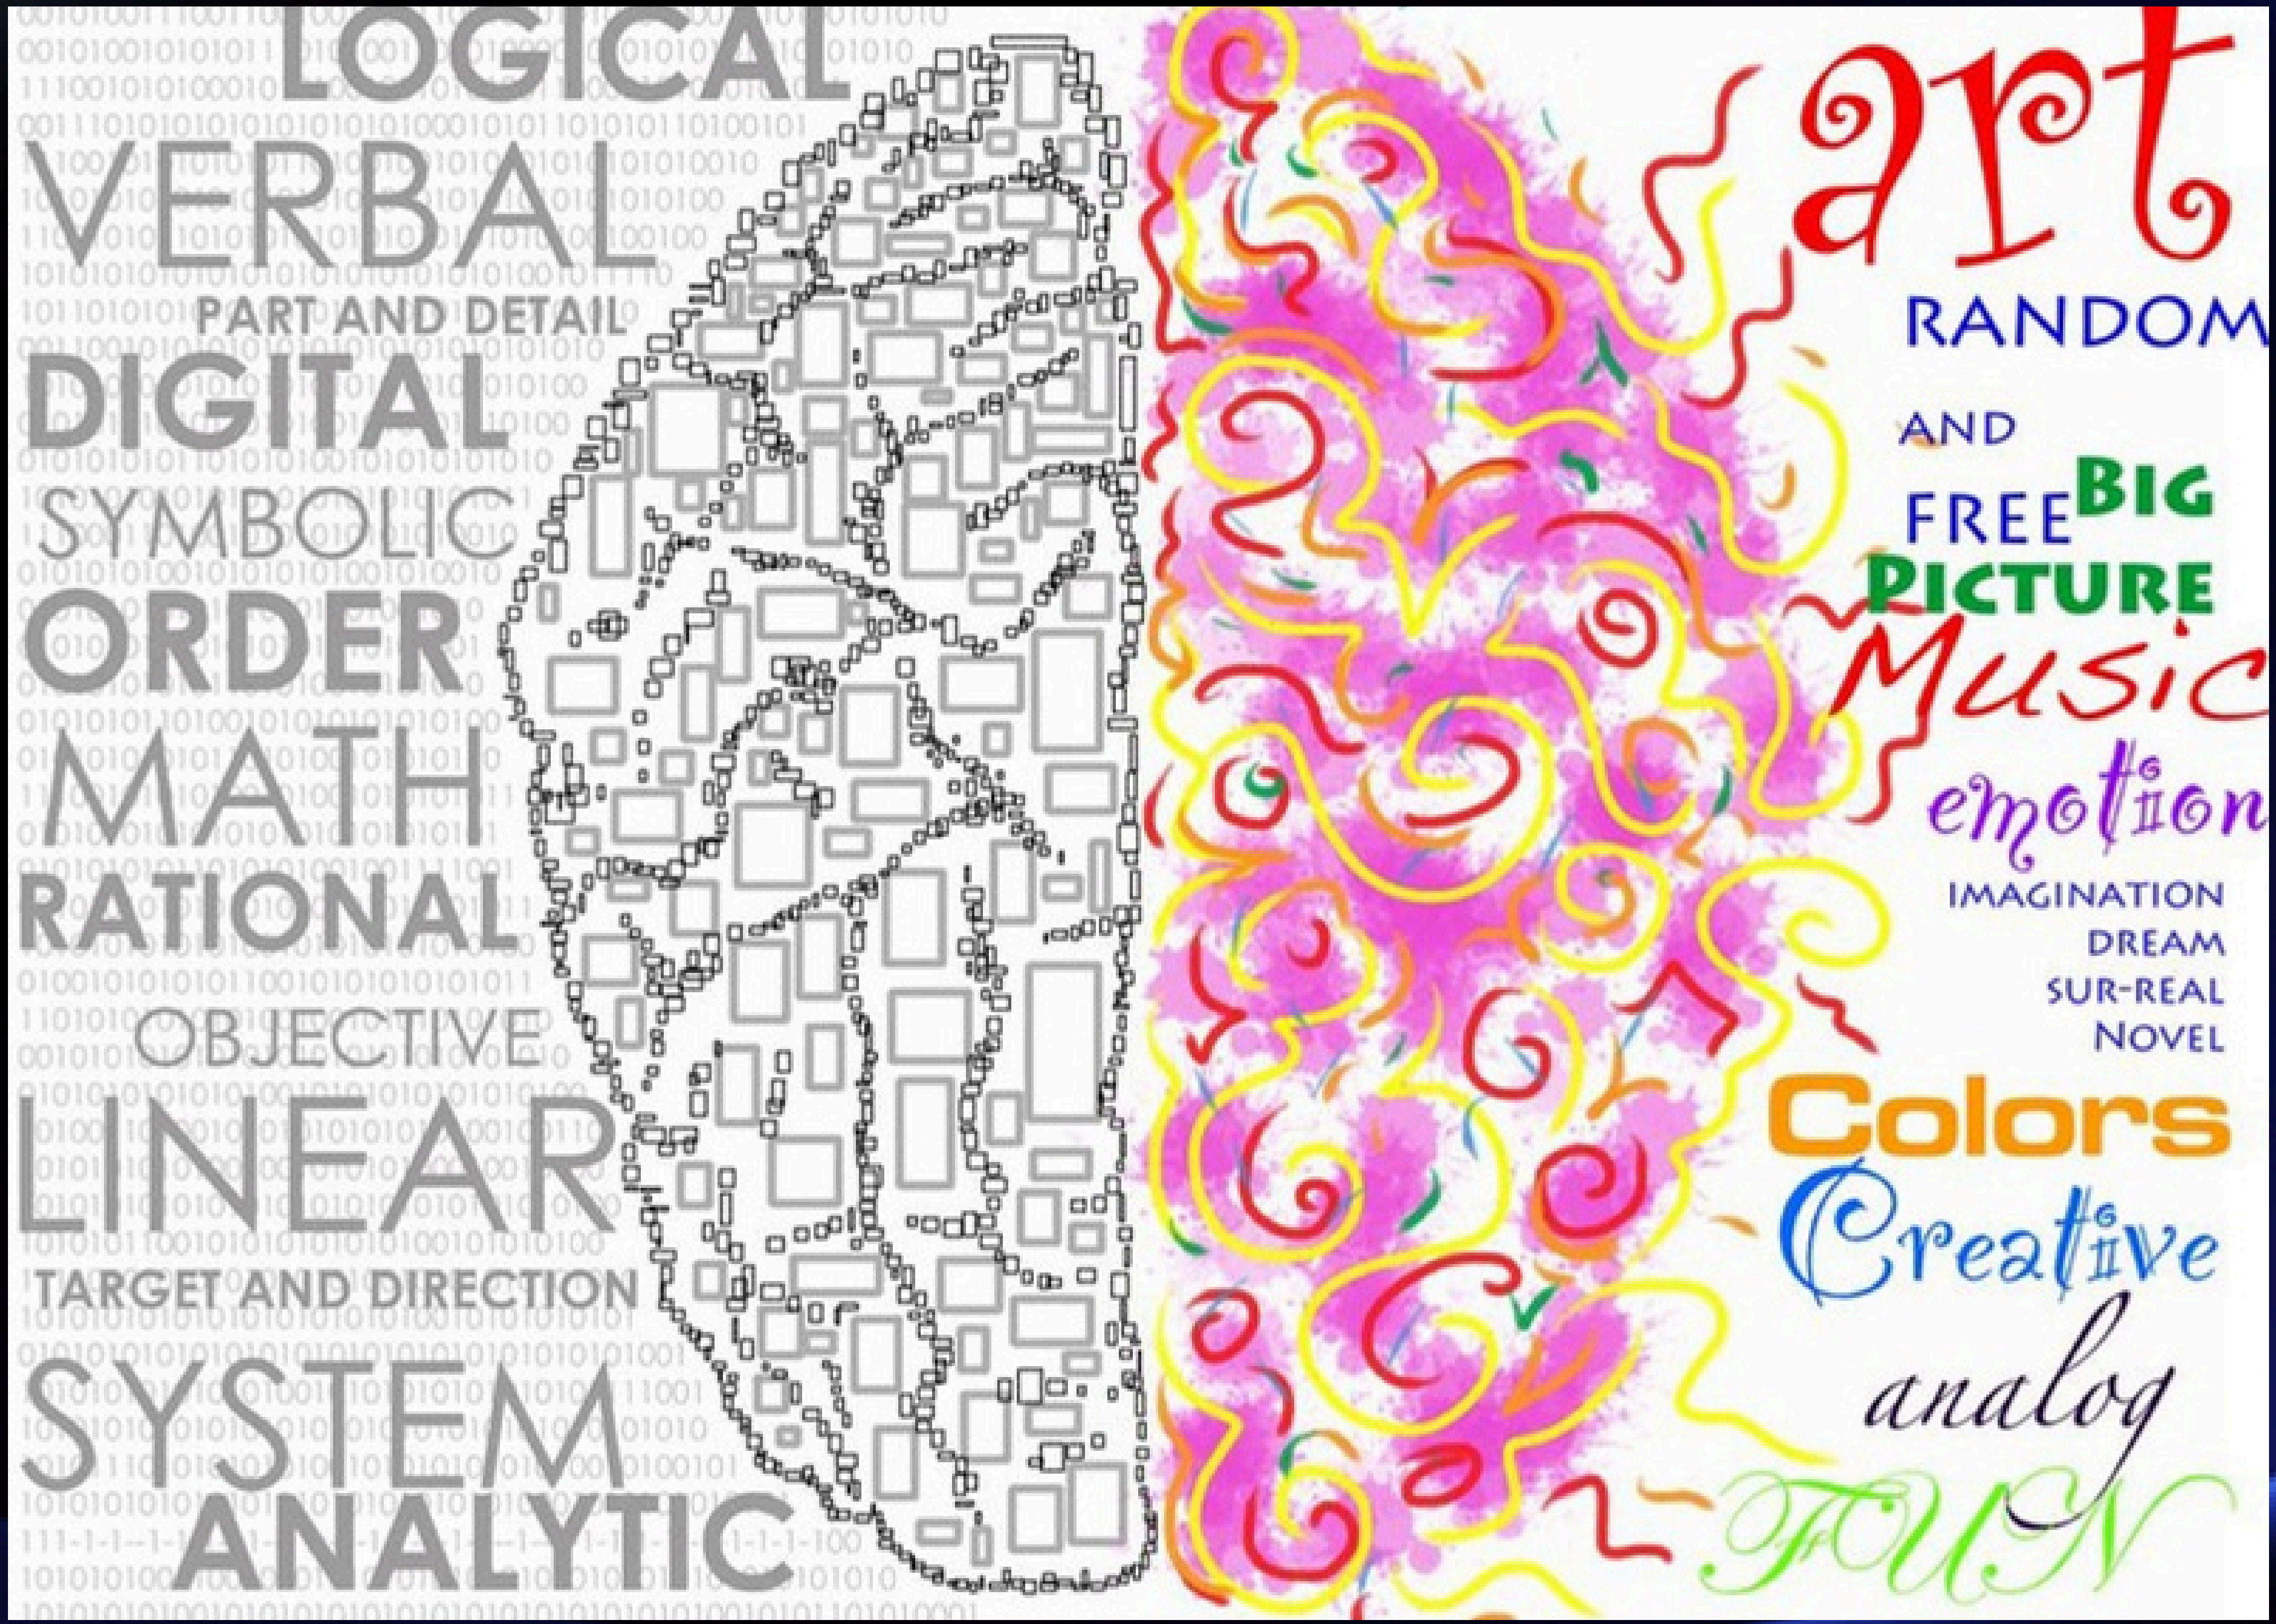

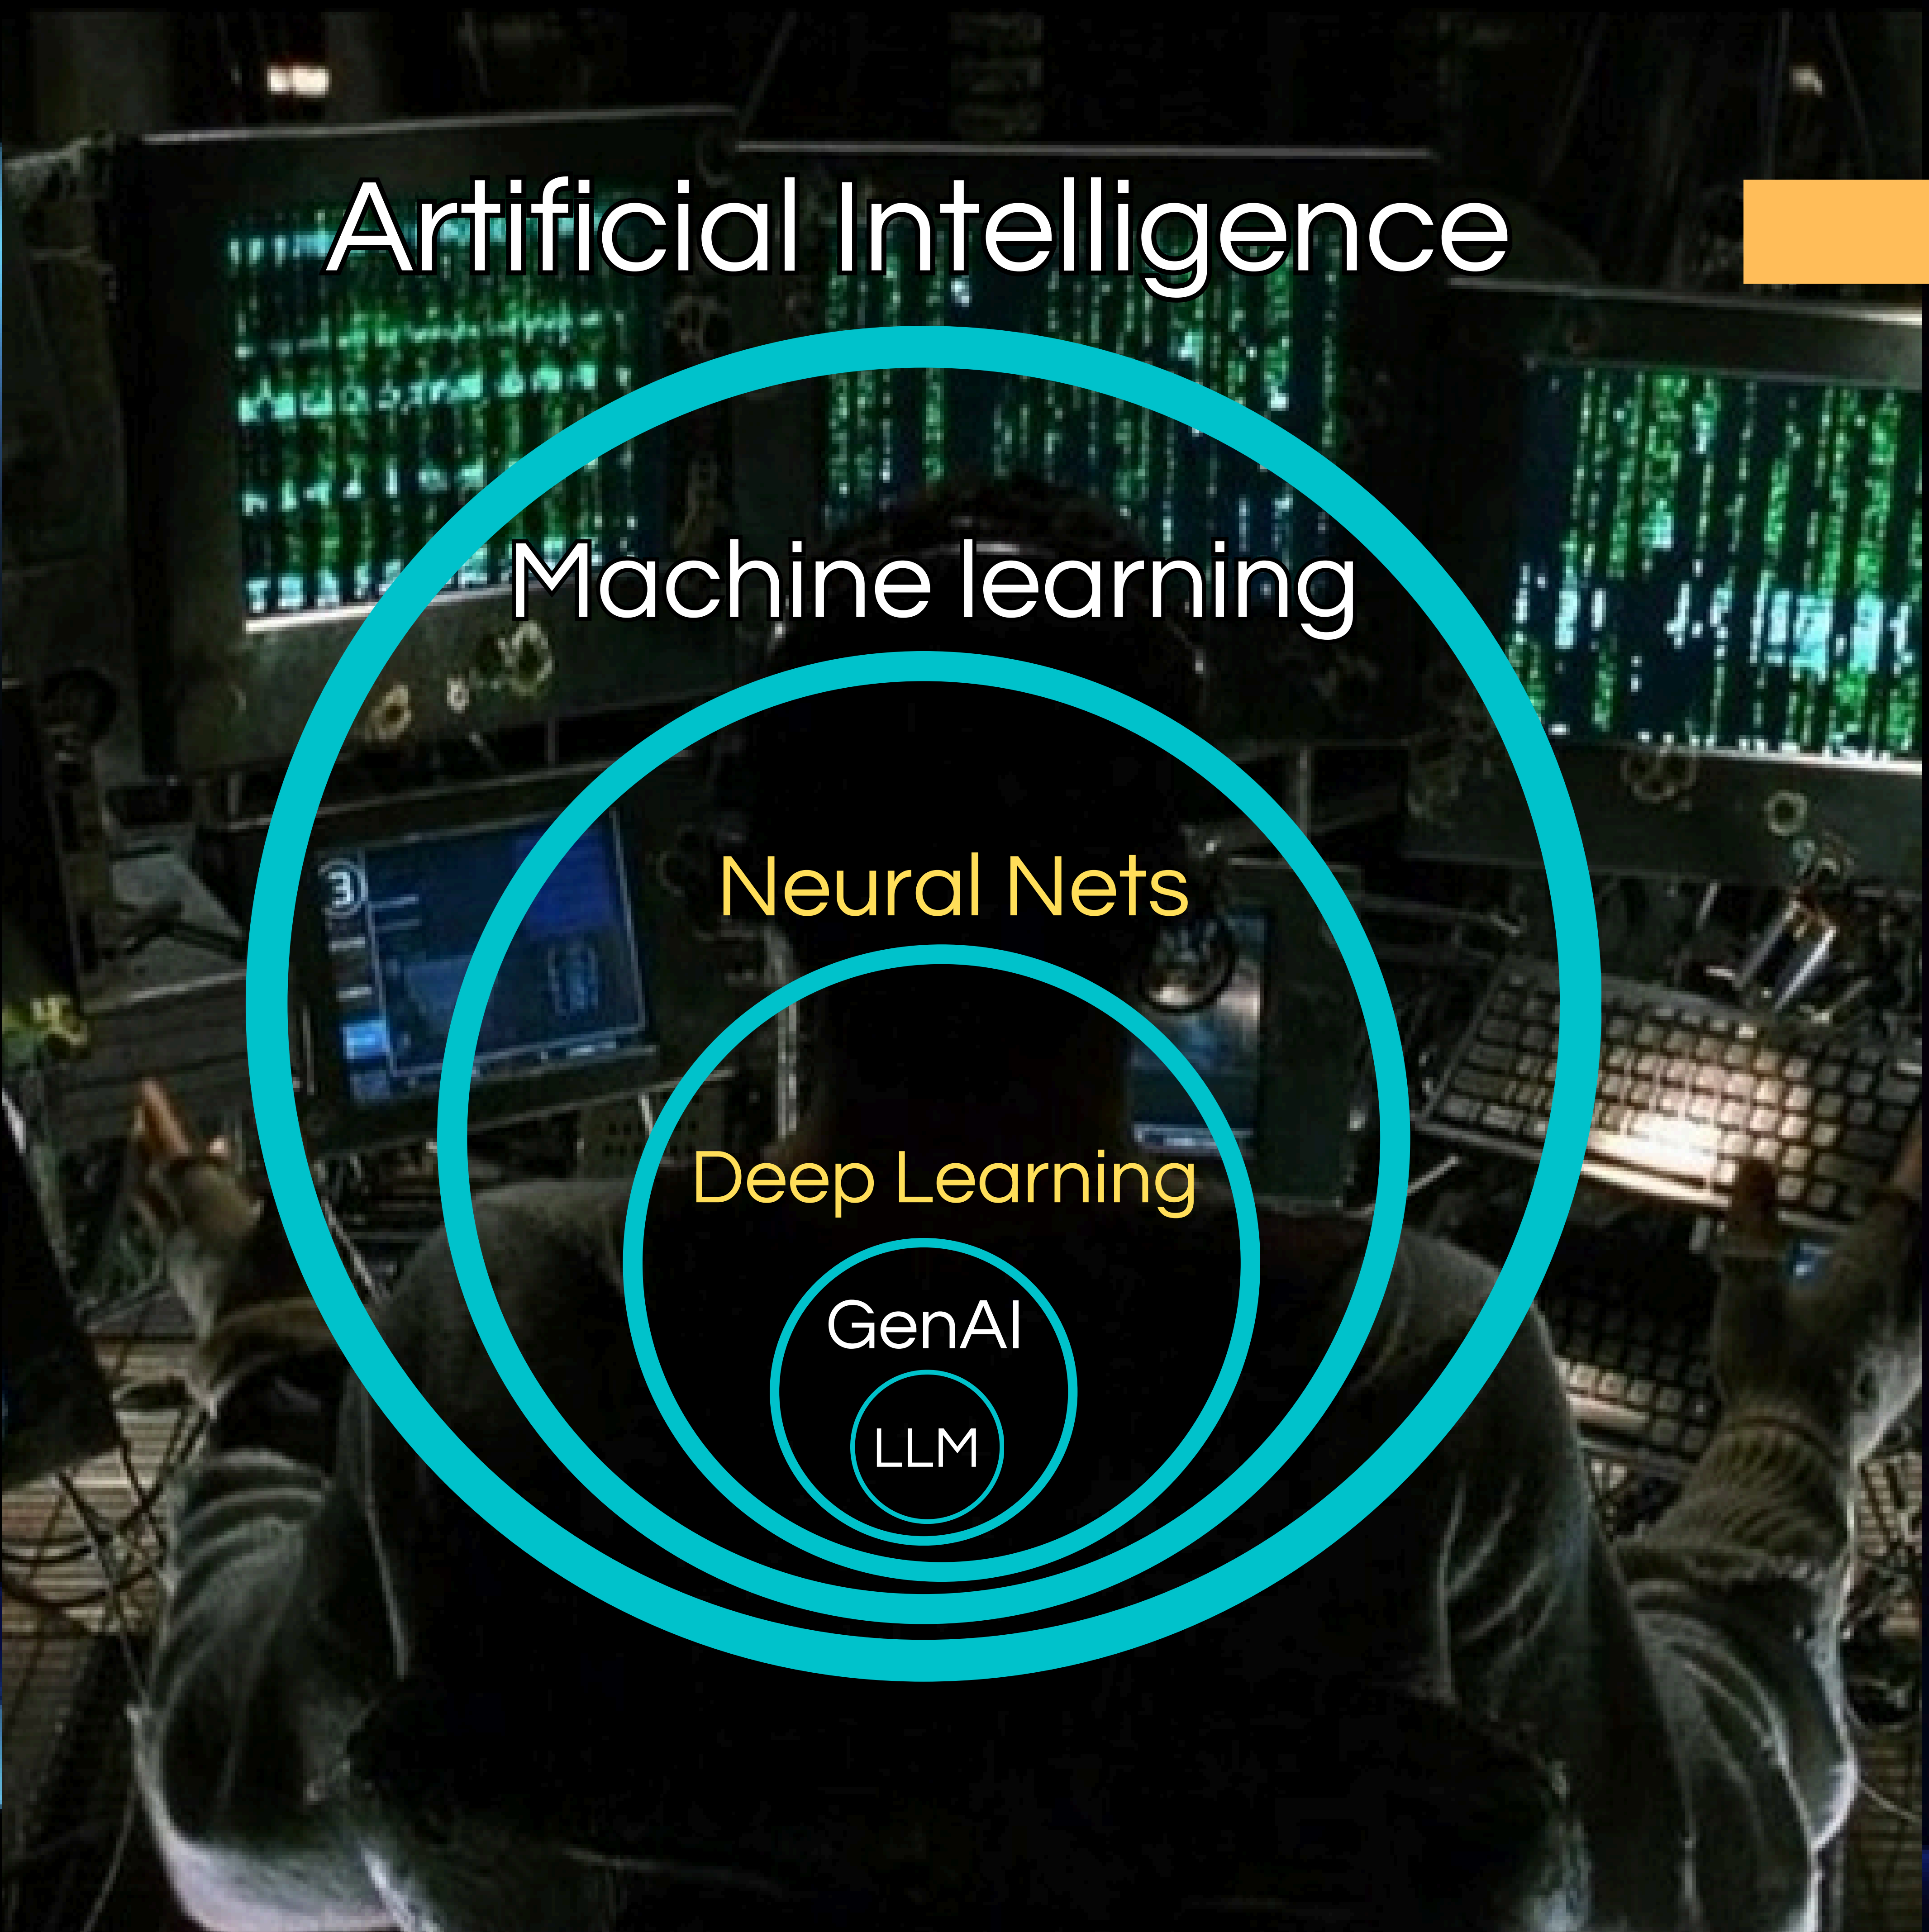

NEURAL NETS & DEEP LEARNING

Generative AI (GenAI) is a branch of Deep Learning focused on creating new data, based on patterns it has learned from existing data

LLM: Large Language Models, the type of GenAI that only output text

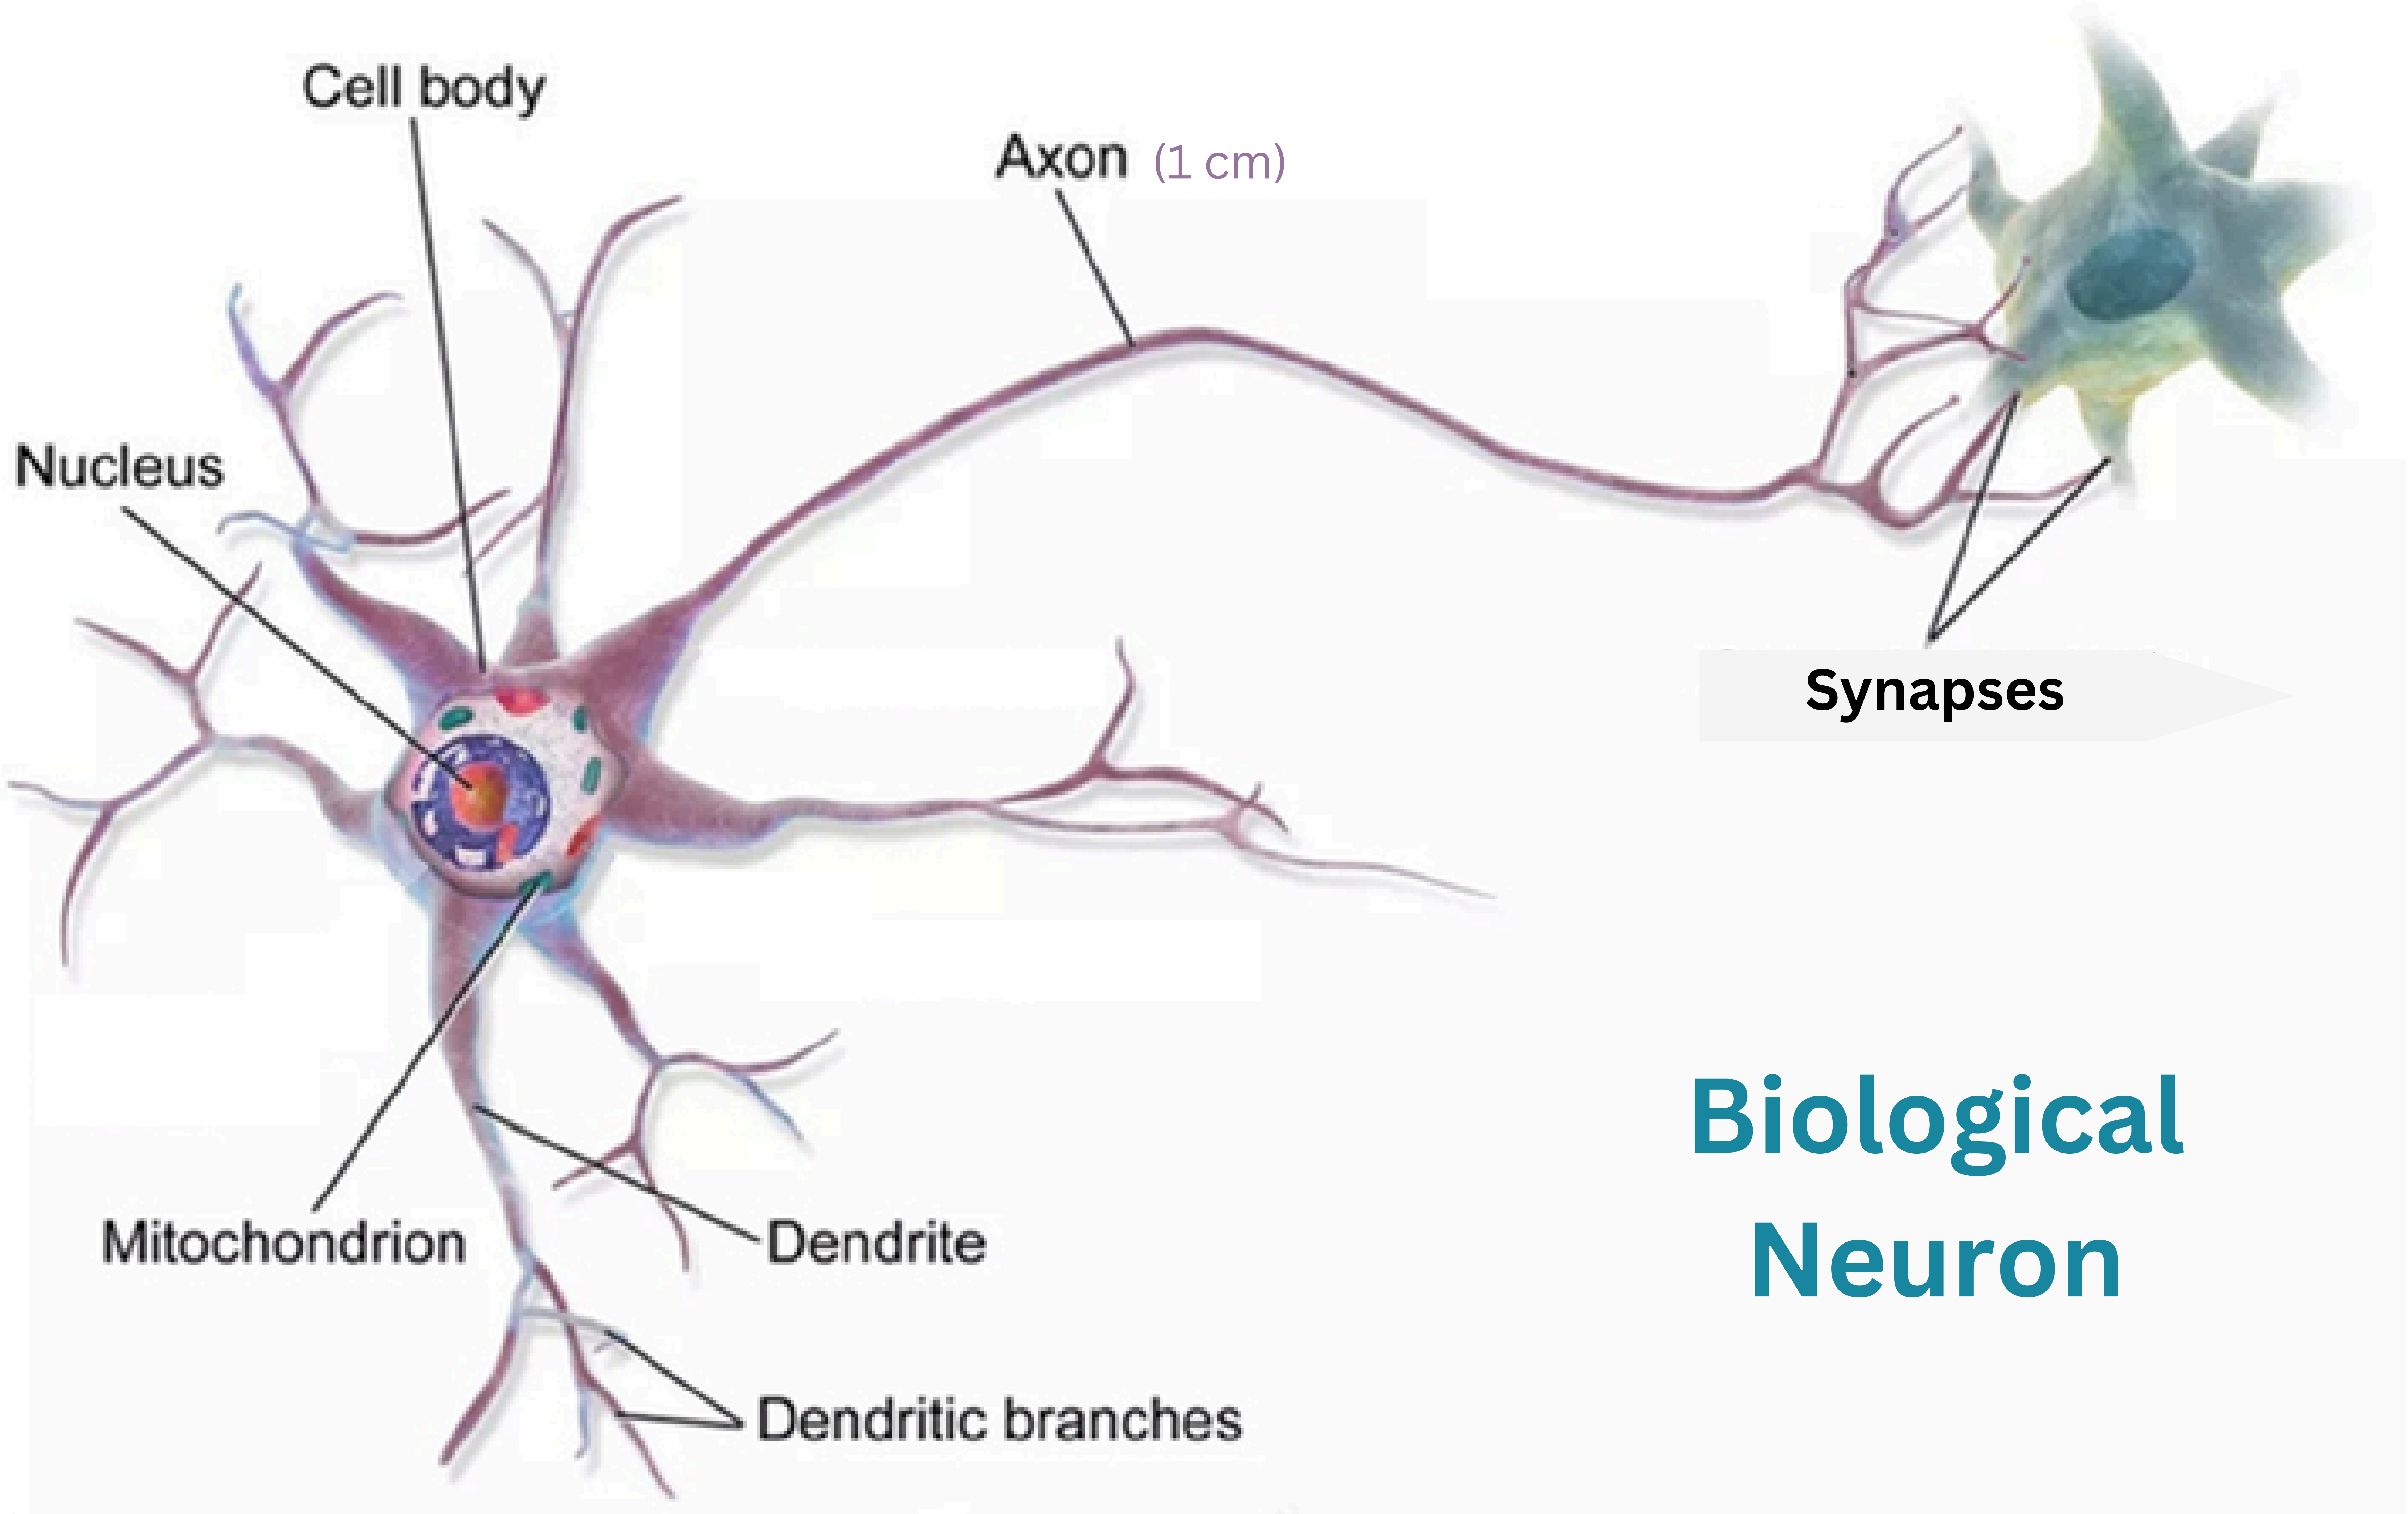

Biological Neuron

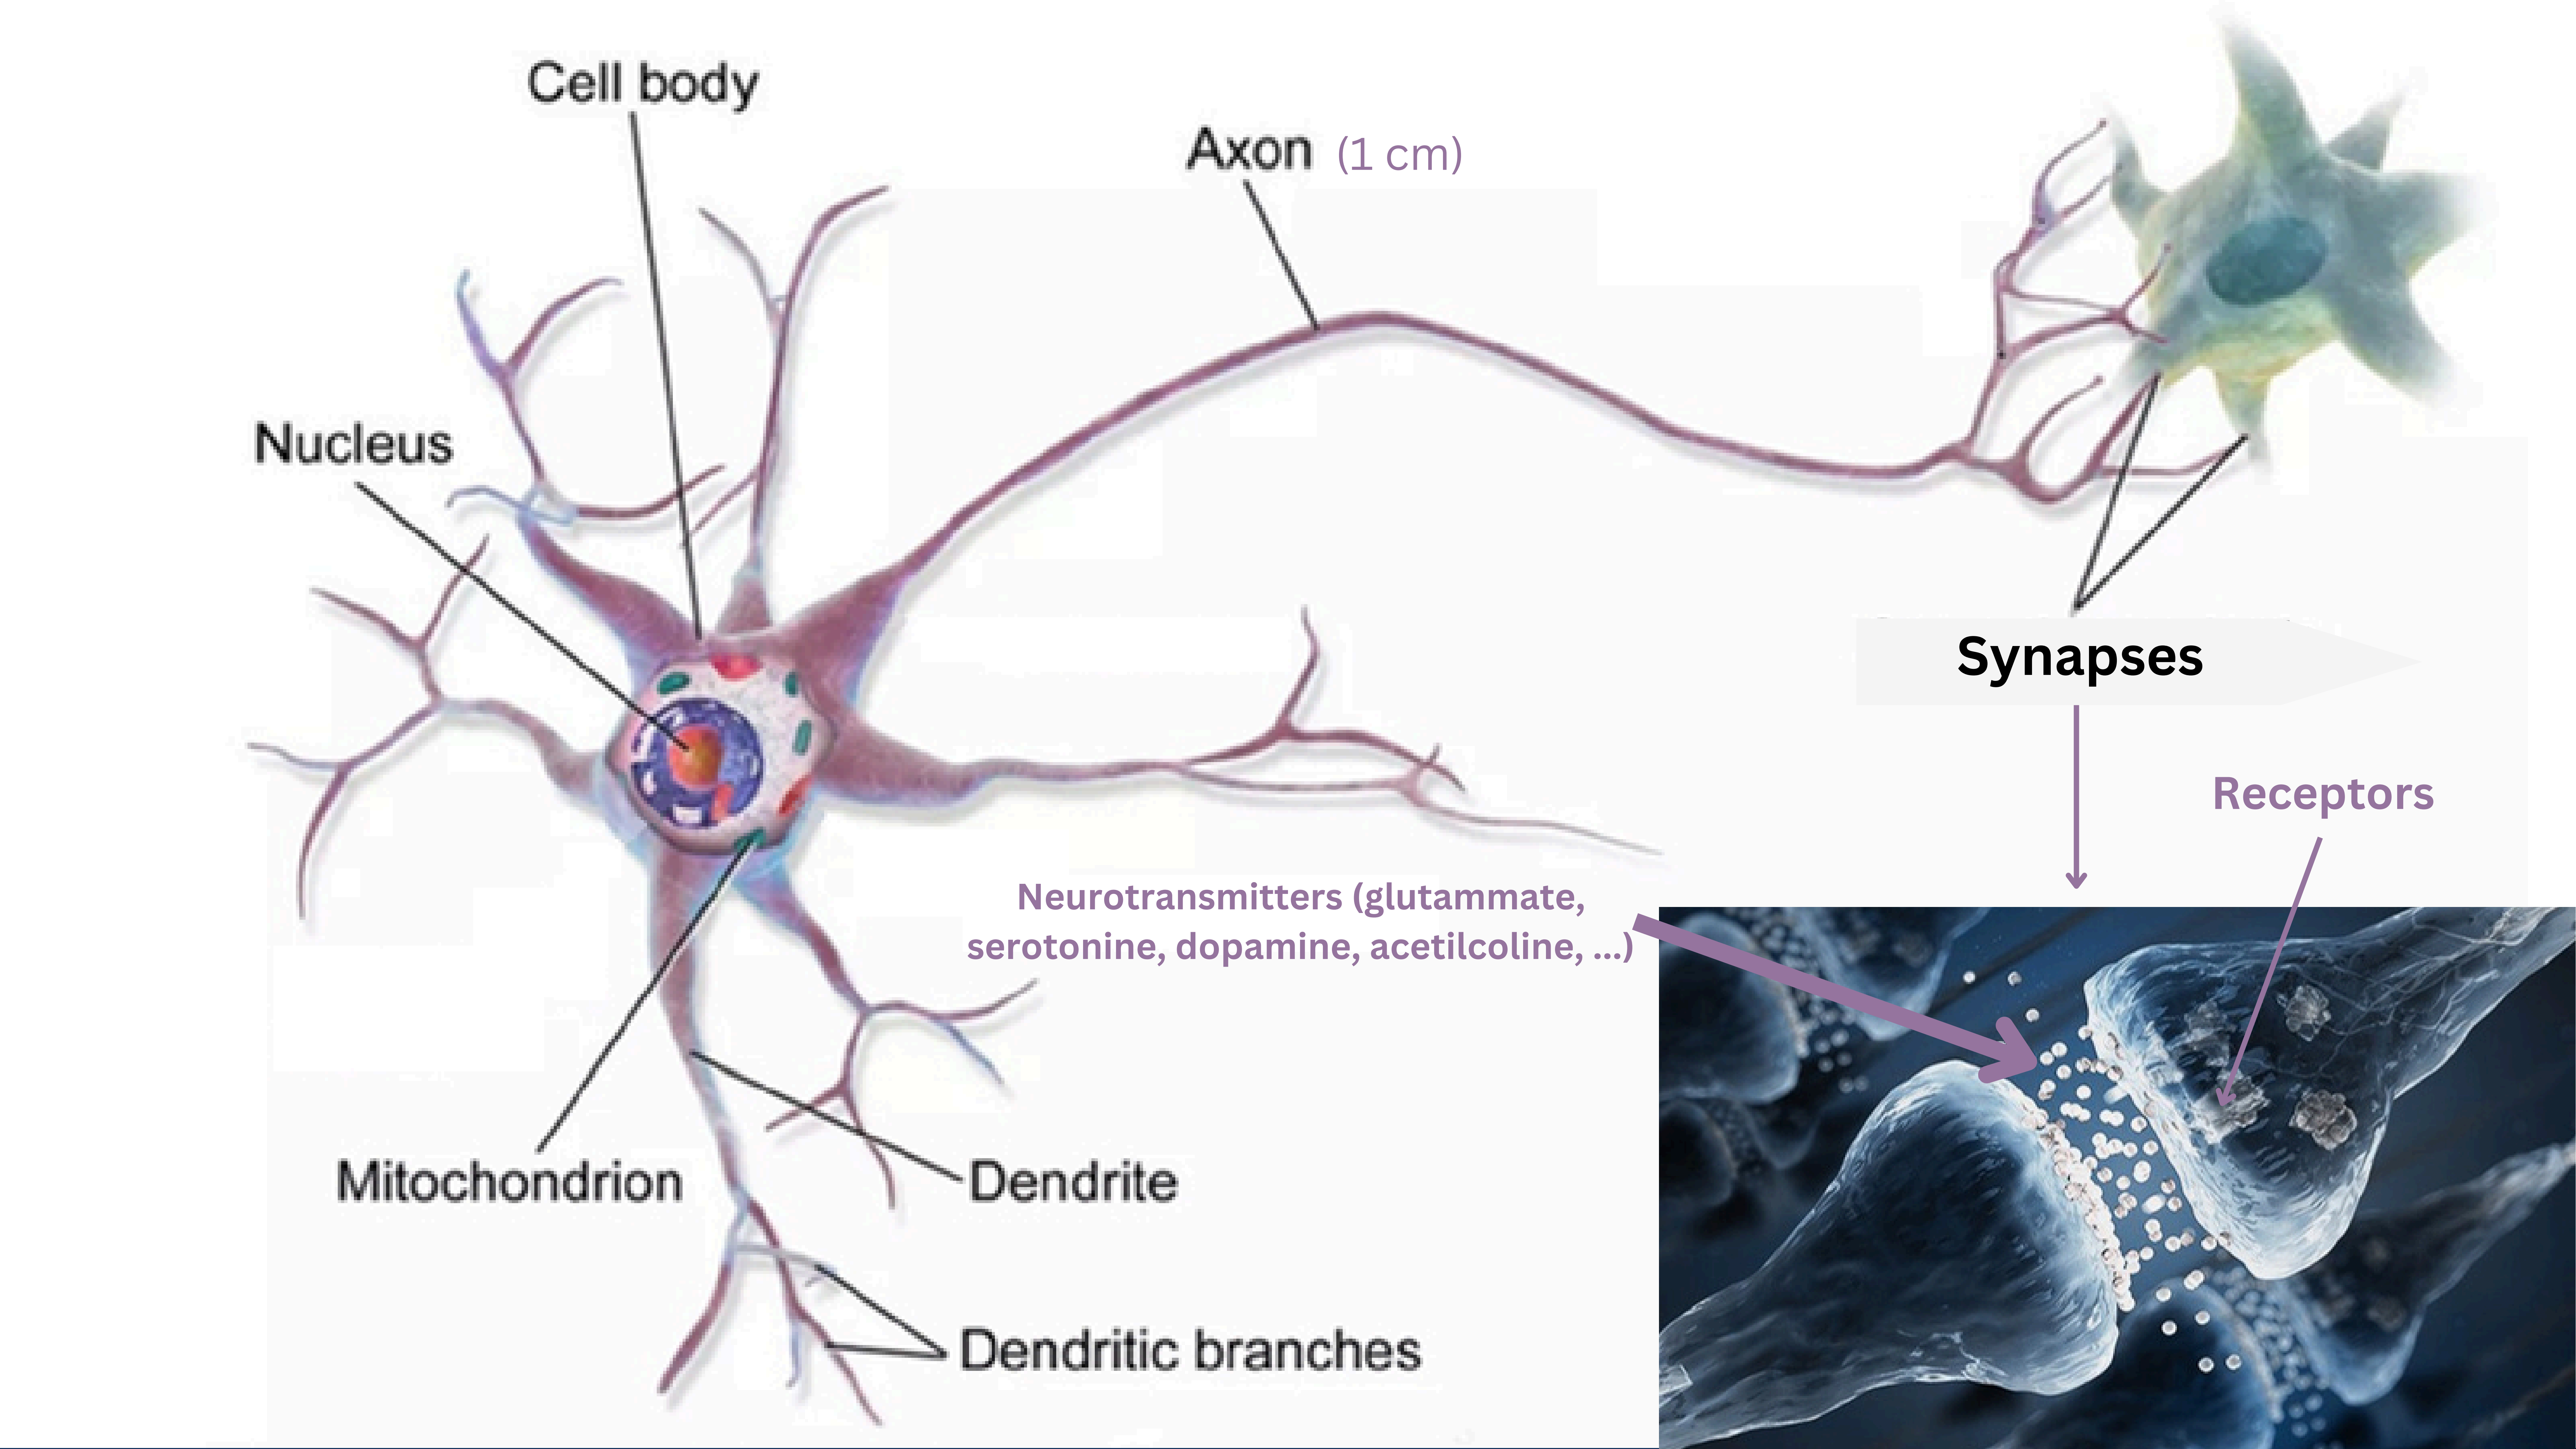
Cell body
Axon (1 cm)
Nucleus
Synapses
Receptors
Neurotransmitters (glutammate, serotonine, dopamine, acetilcoline, ...)
Mitochondrion
Dendrite
Dendritic branches

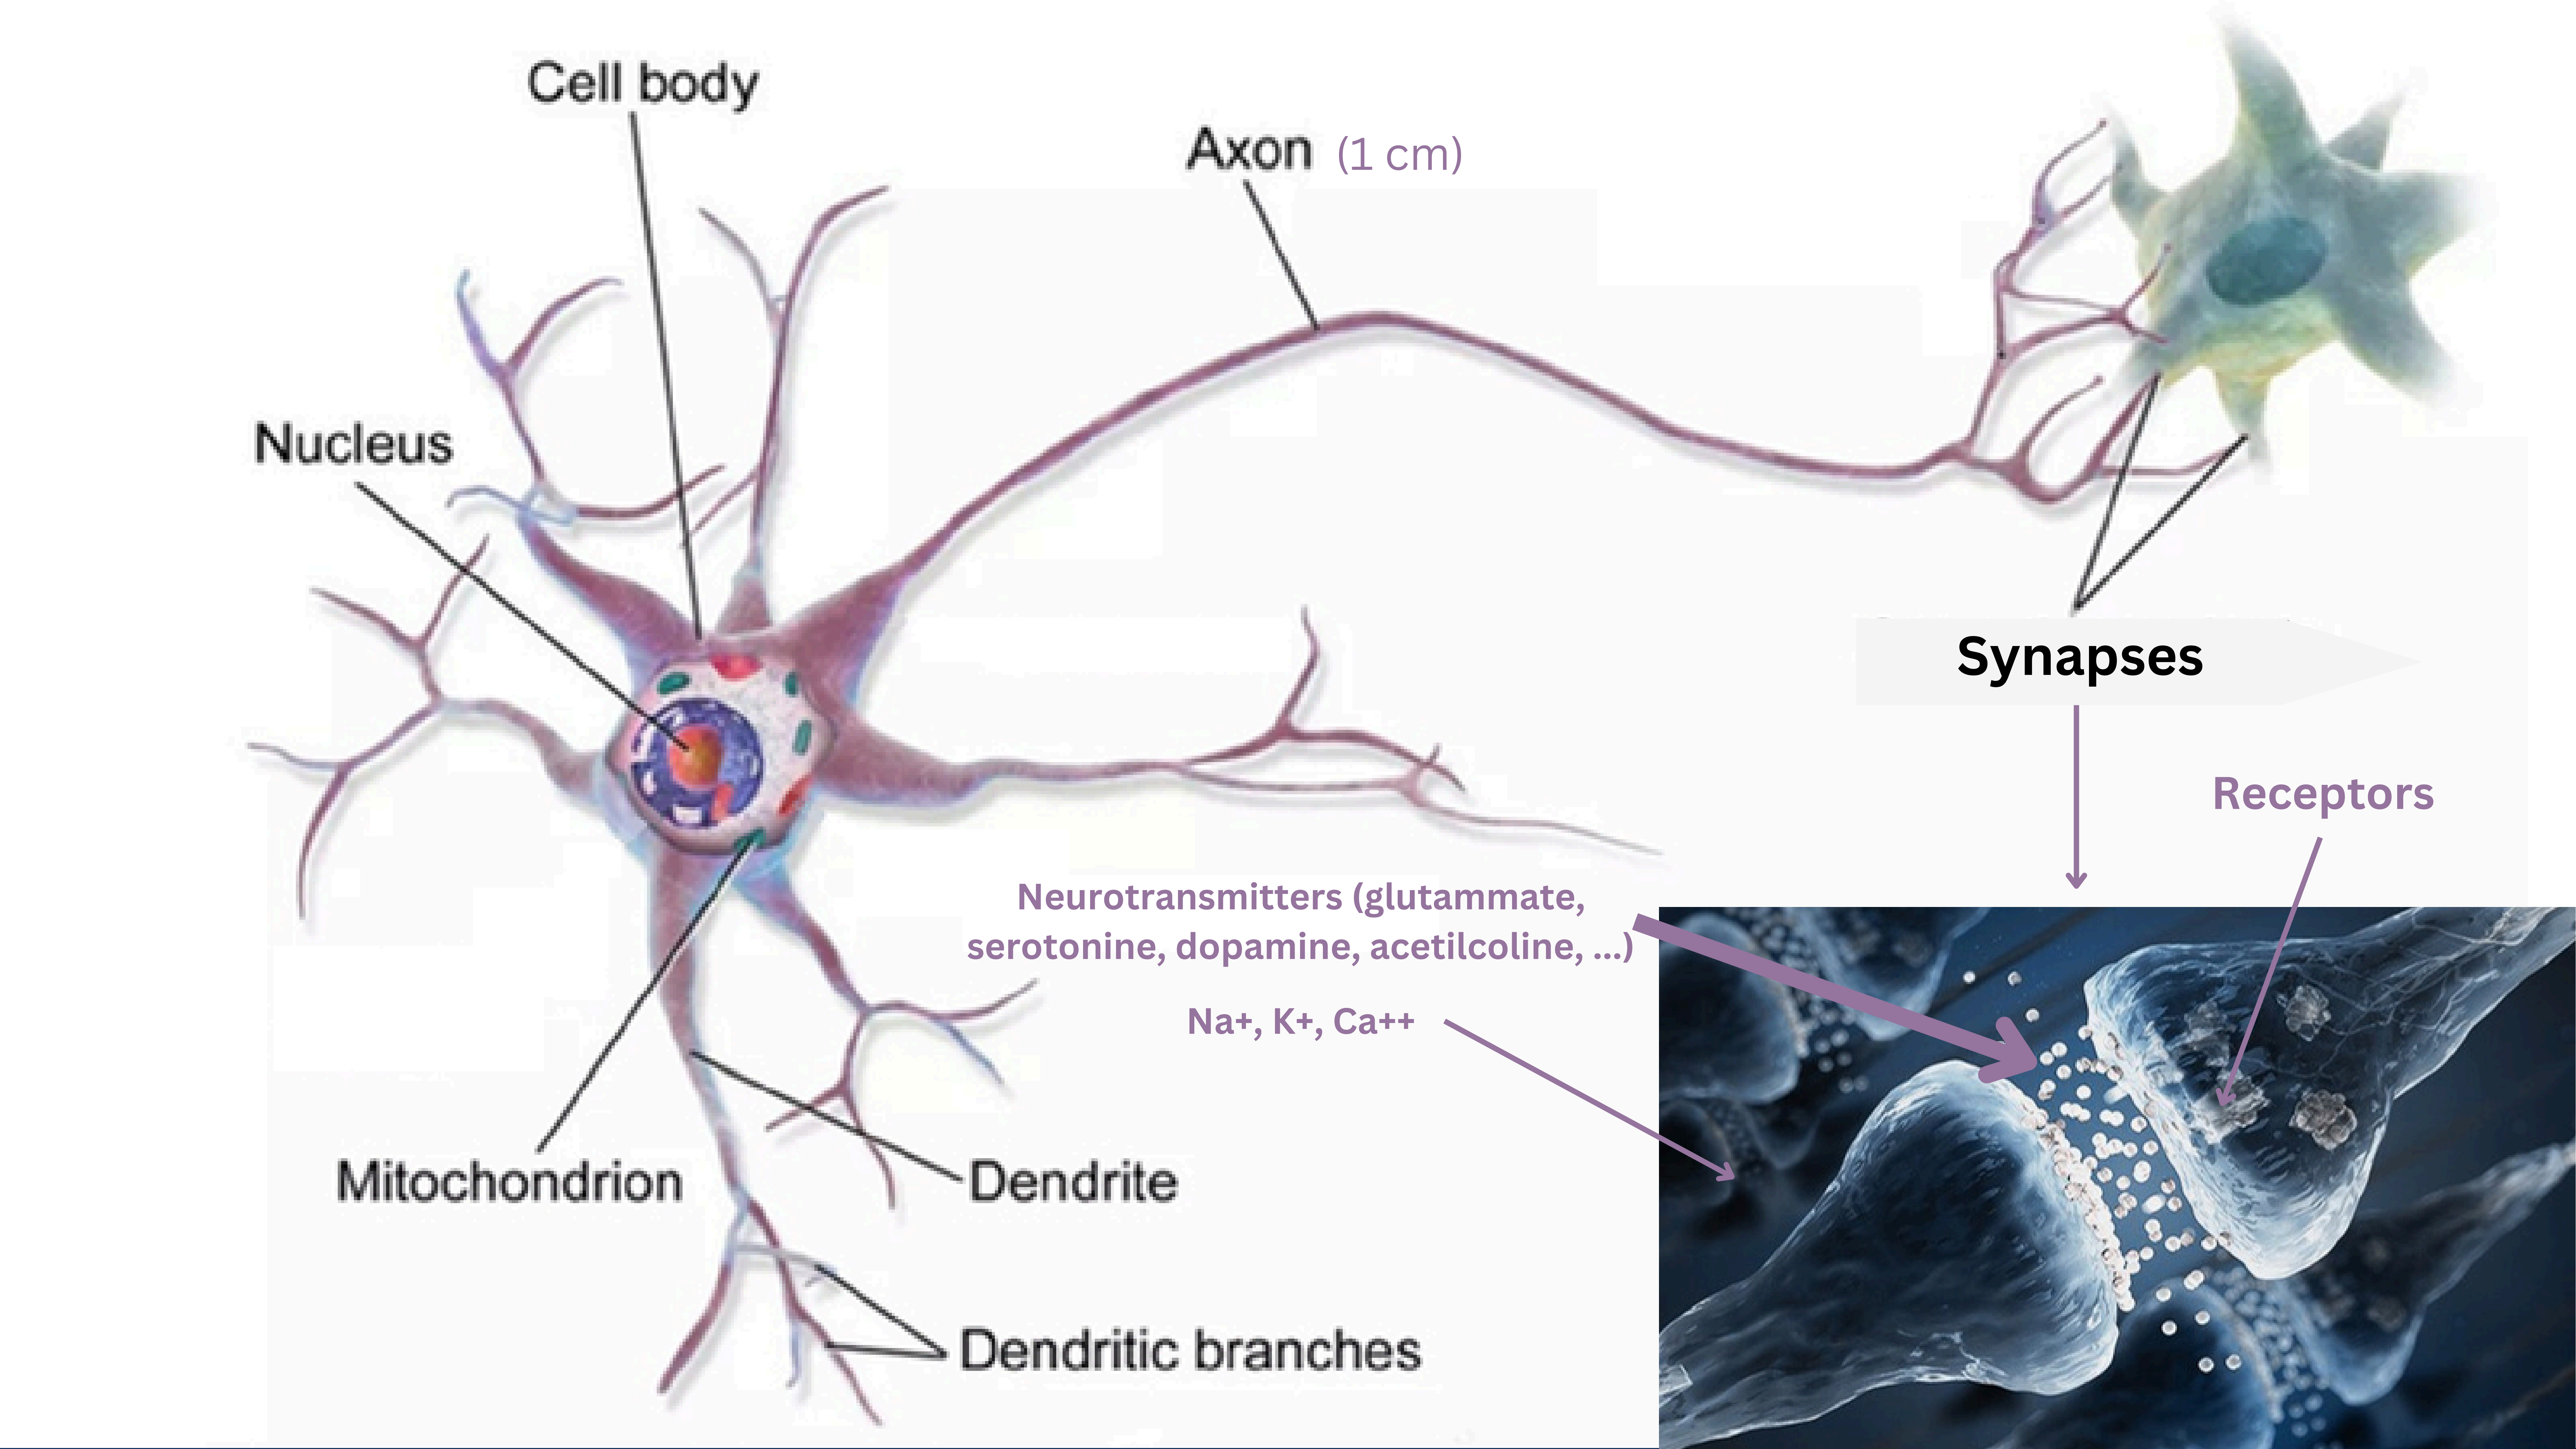
Cell body
Axon (1 cm)
Nucleus
Synapses
Receptors
Neurotransmitters (glutammate, serotonine, dopamine, acetilcoline, ...)
Na+, K+, Ca++
Mitochondrion
Dendrite
Dendritic branches

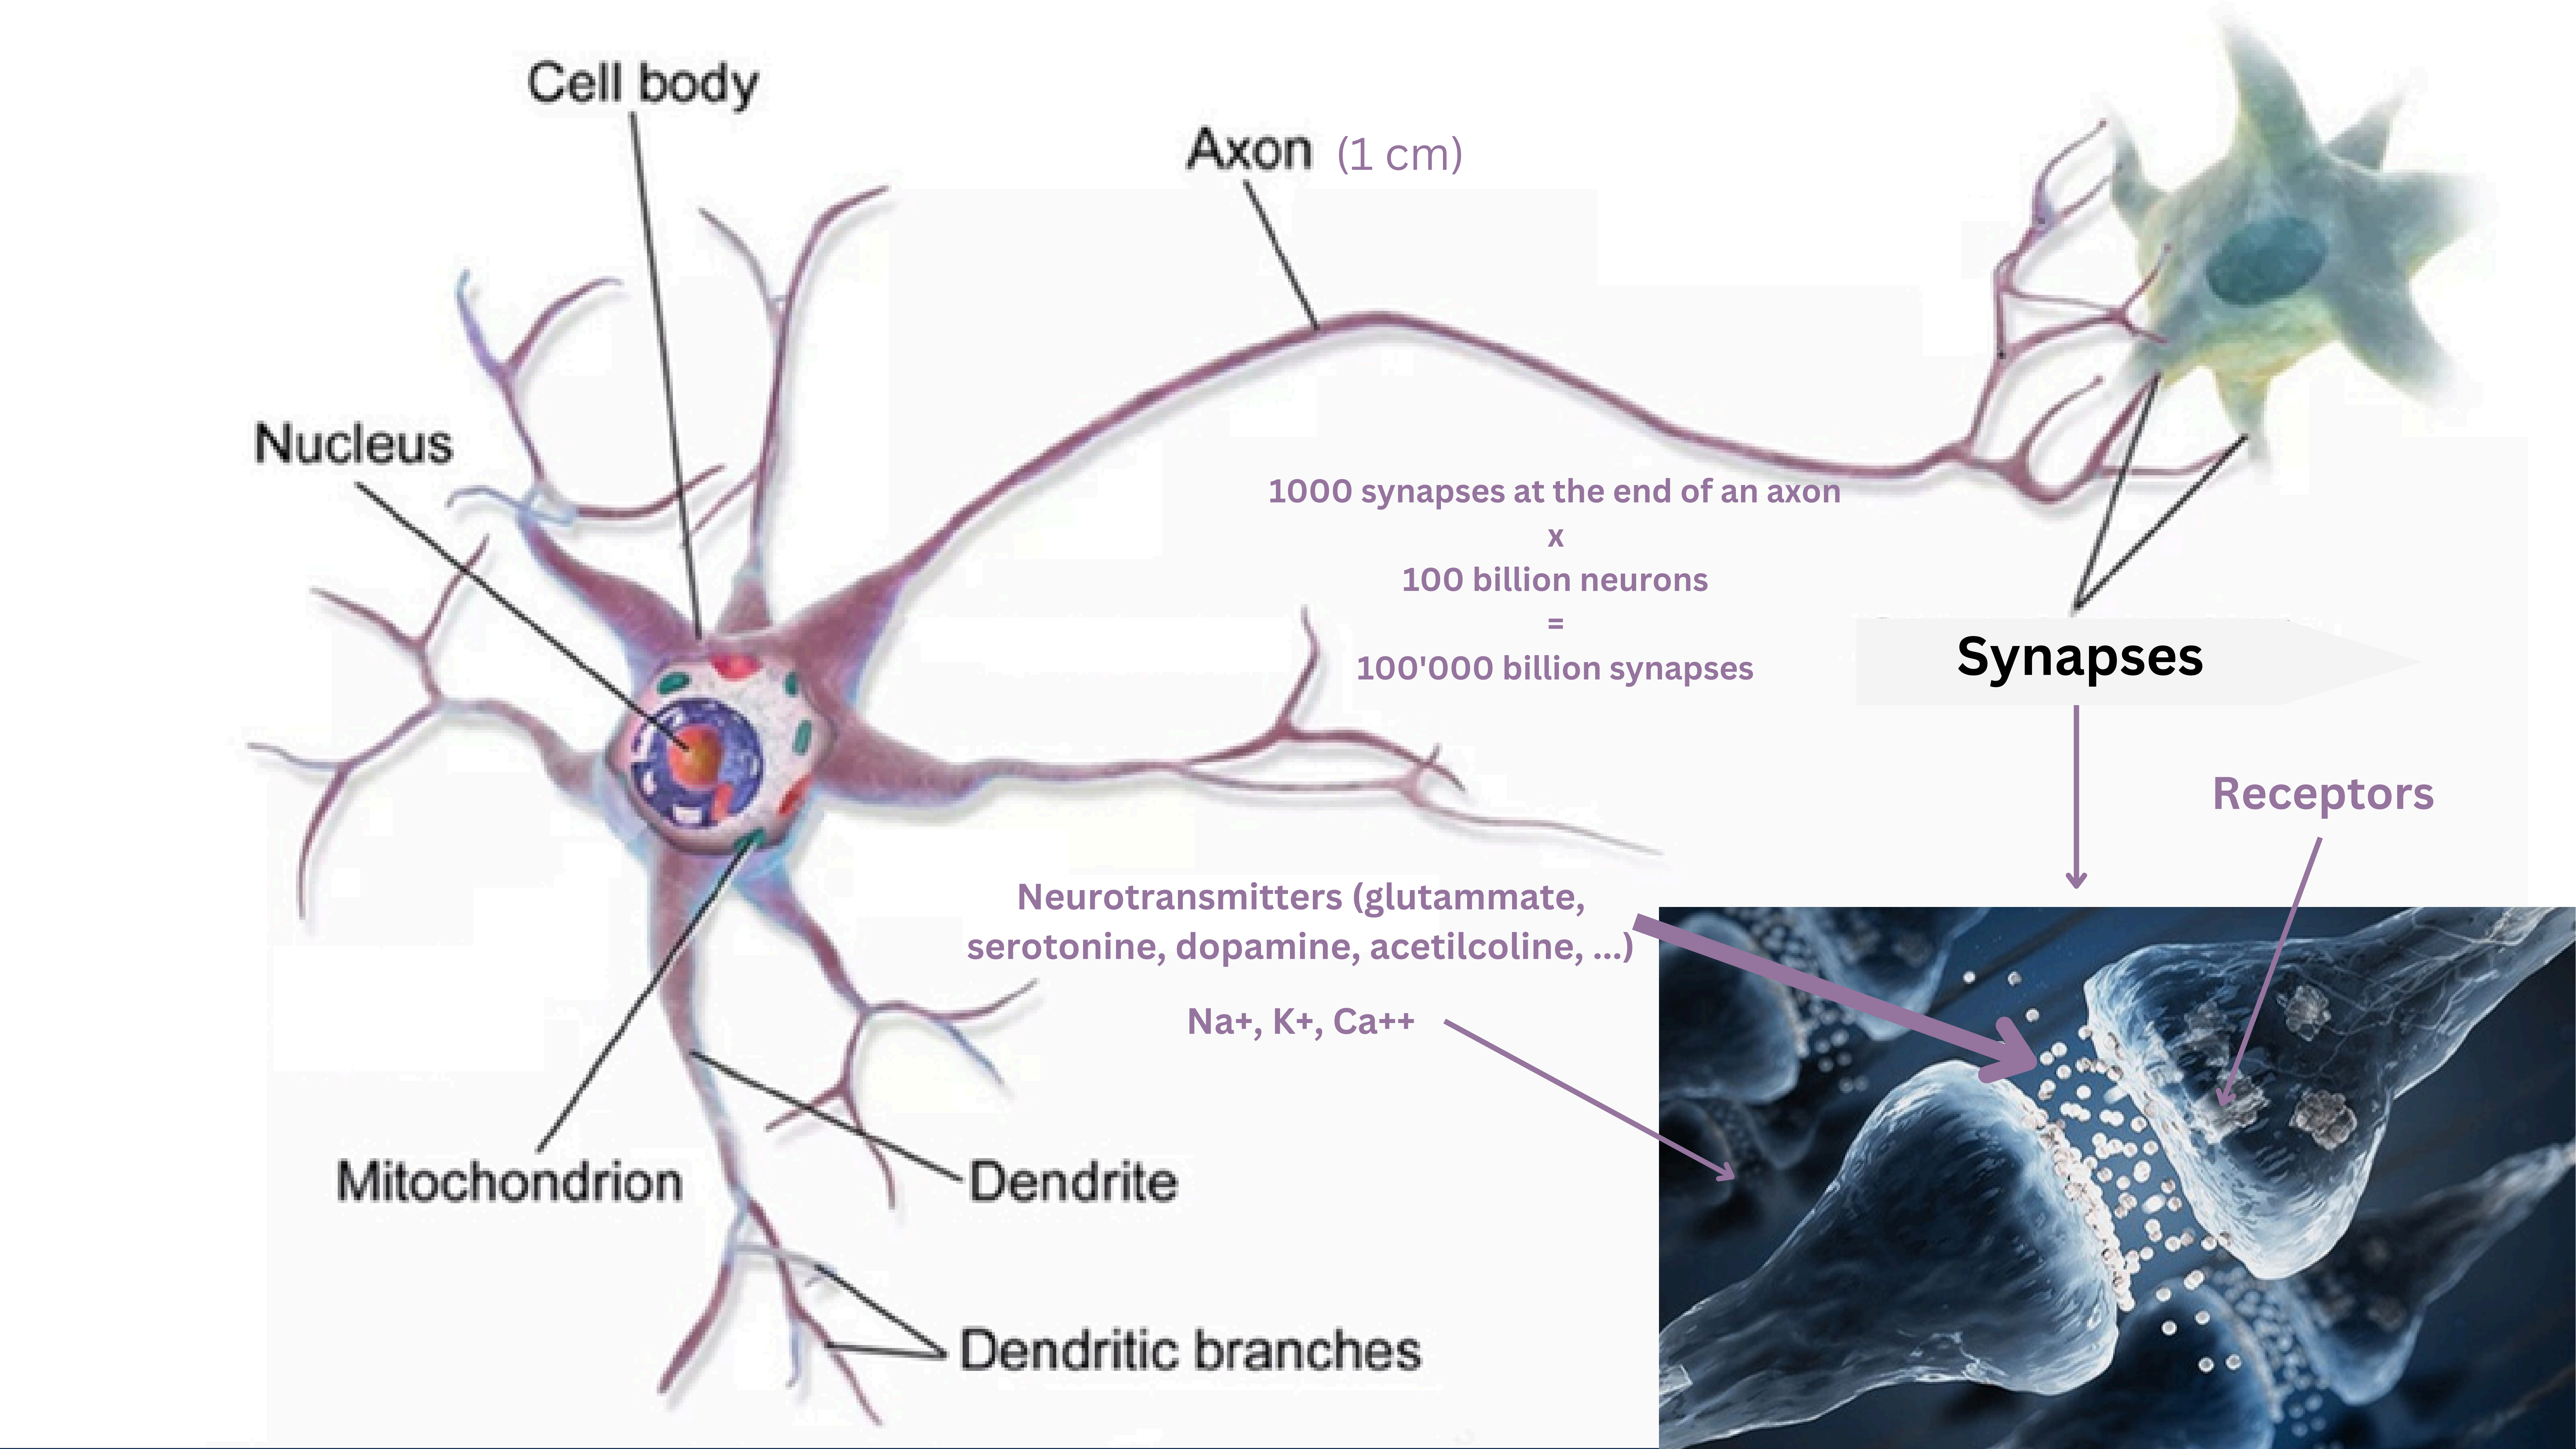
Cell body
Axon (1 cm)
Nucleus
1000 synapses at the end of an axon
x
100 billion neurons
=
100'000 billion synapses
Synapses
Receptors
Neurotransmitters (glutammate, serotonine, dopamine, acetilcoline, ...)
Na+, K+, Ca++
Mitochondrion
Dendrite
Dendritic branches

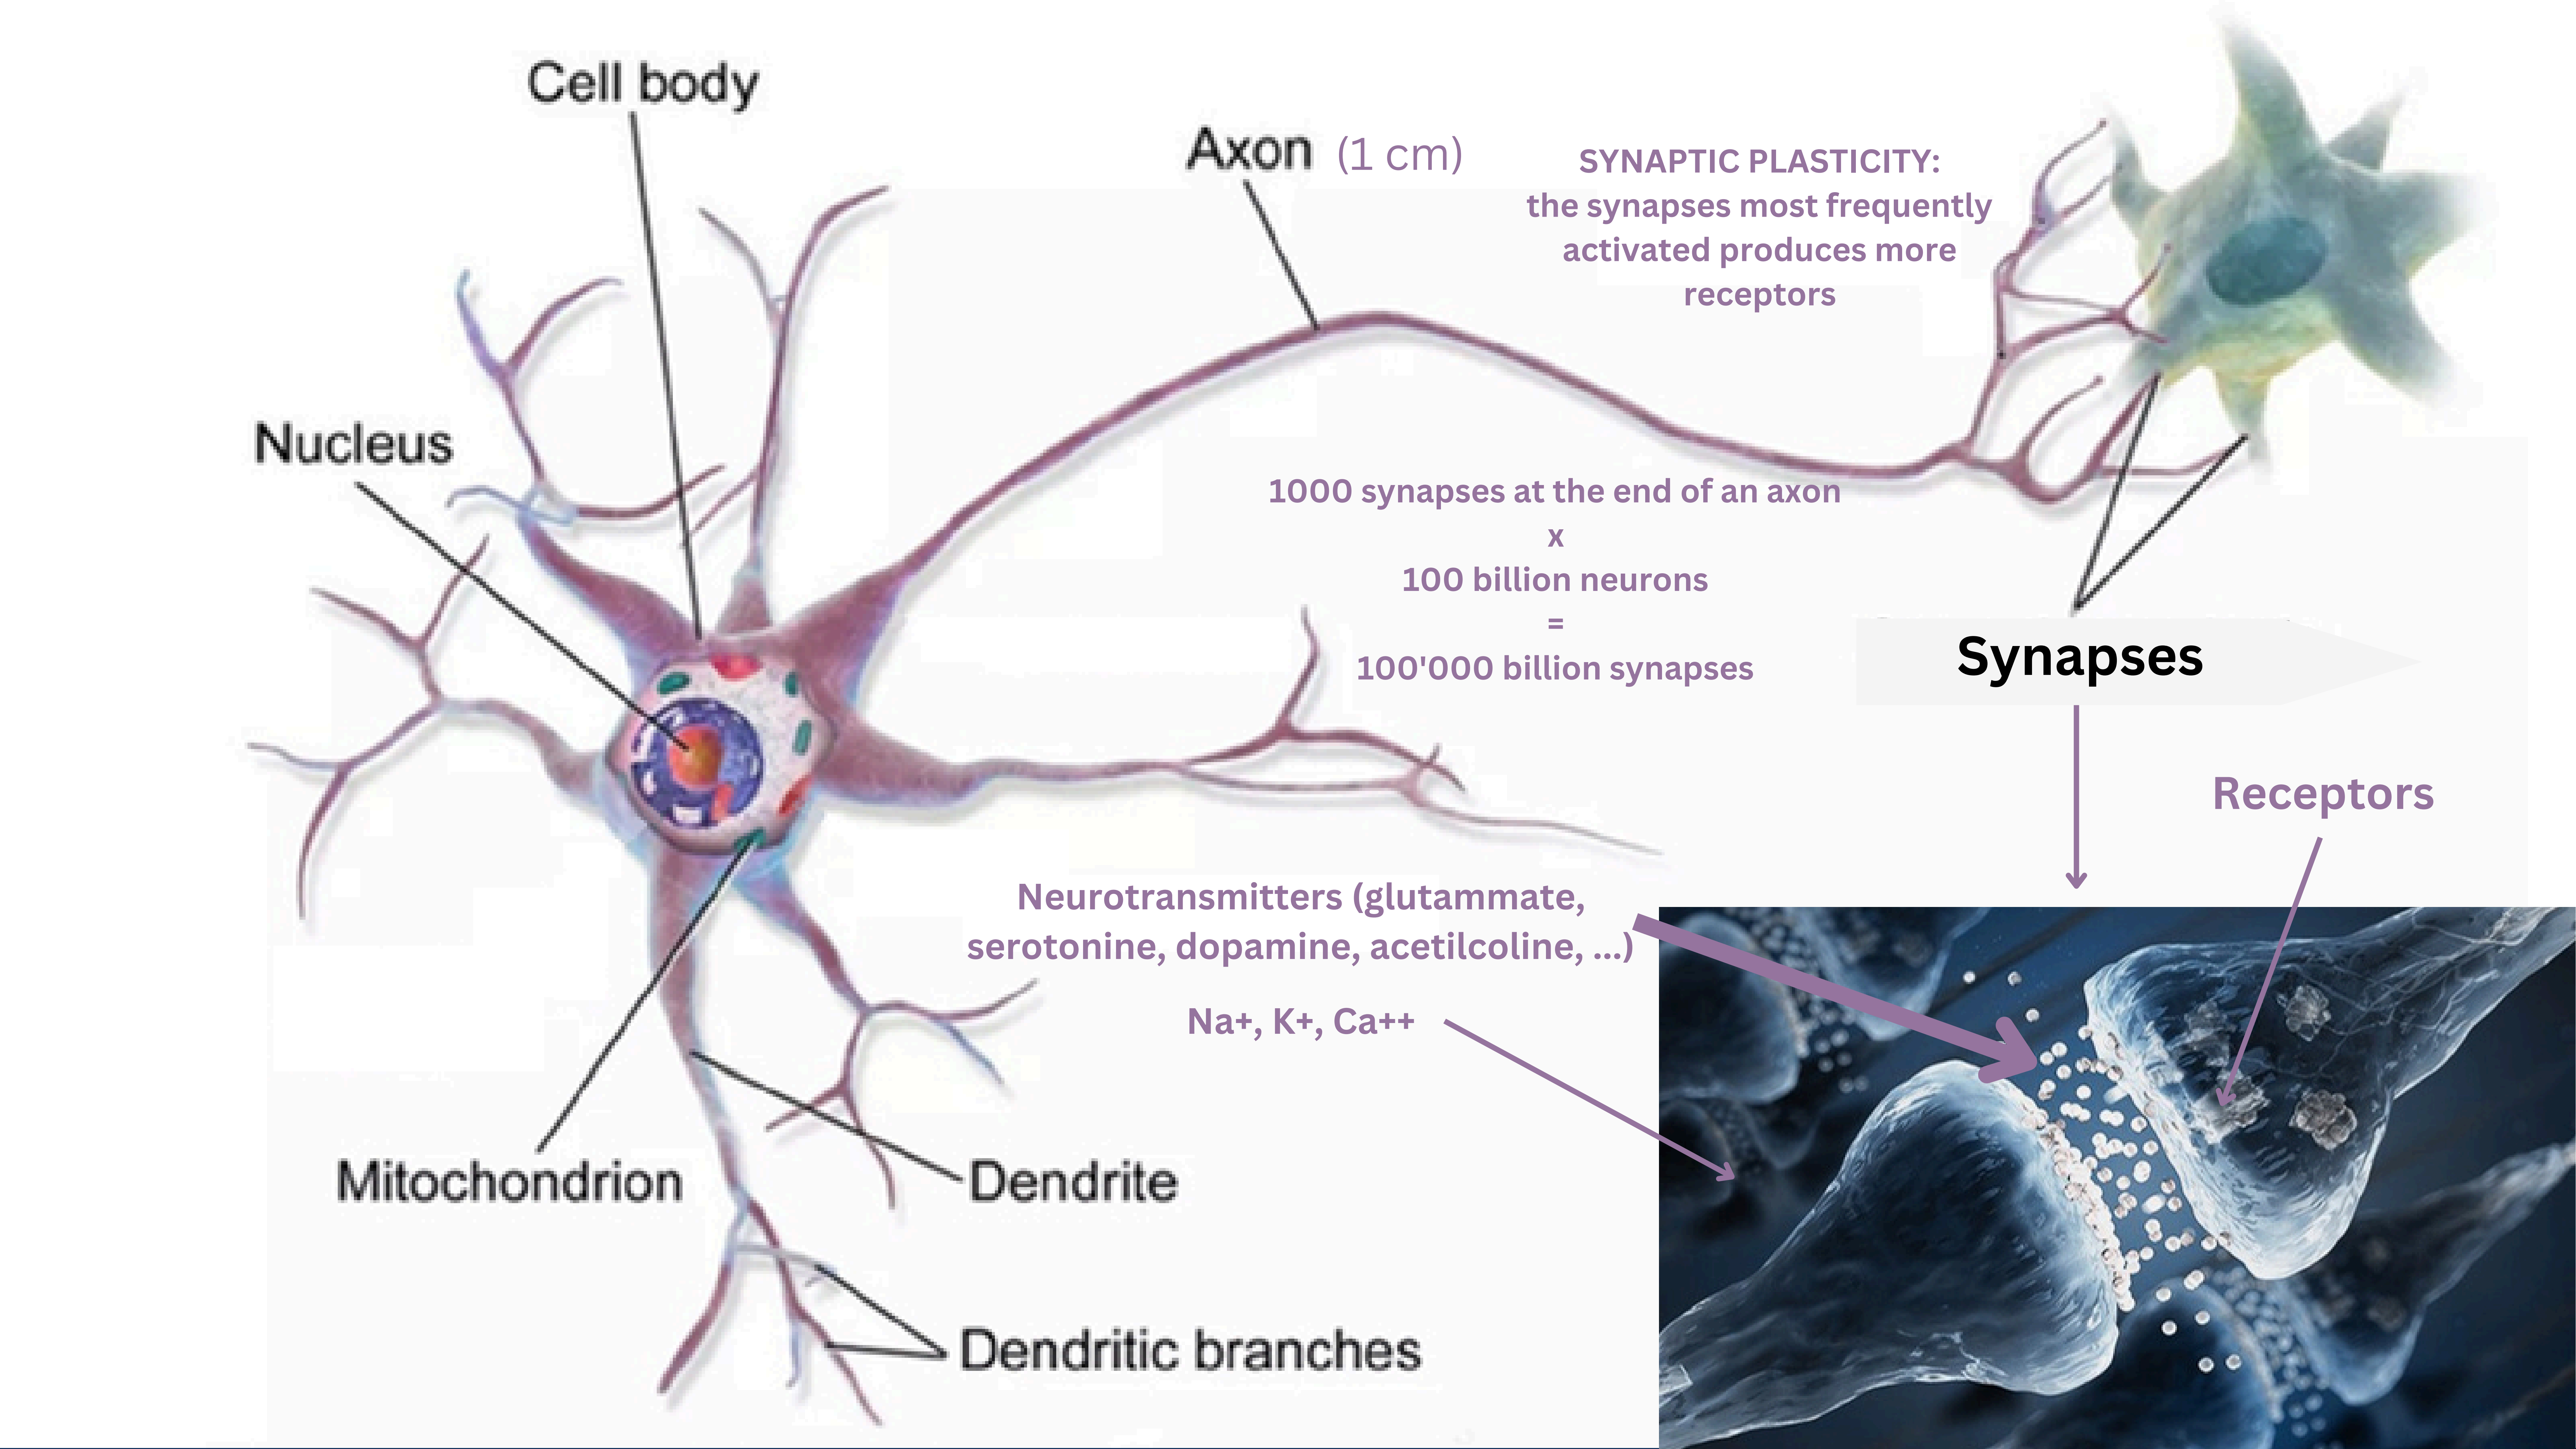
Cell body
Axon (1 cm)
SYNAPTIC PLASTICITY:
the synapses most frequently activated produces more receptors
Nucleus
1000 synapses at the end of an axon
x
100 billion neurons
=
100'000 billion synapses
Synapses
Receptors
Neurotransmitters (glutammate, serotonine, dopamine, acetilcoline, ...)
Na+, K+, Ca++
Mitochondrion
Dendrite
Dendritic branches

MEMORY OF NEURAL NETWORKS

Our minds do not memorize data in specific synapses as computer do. Instead, data is distributed through all the neural network, as in holograms

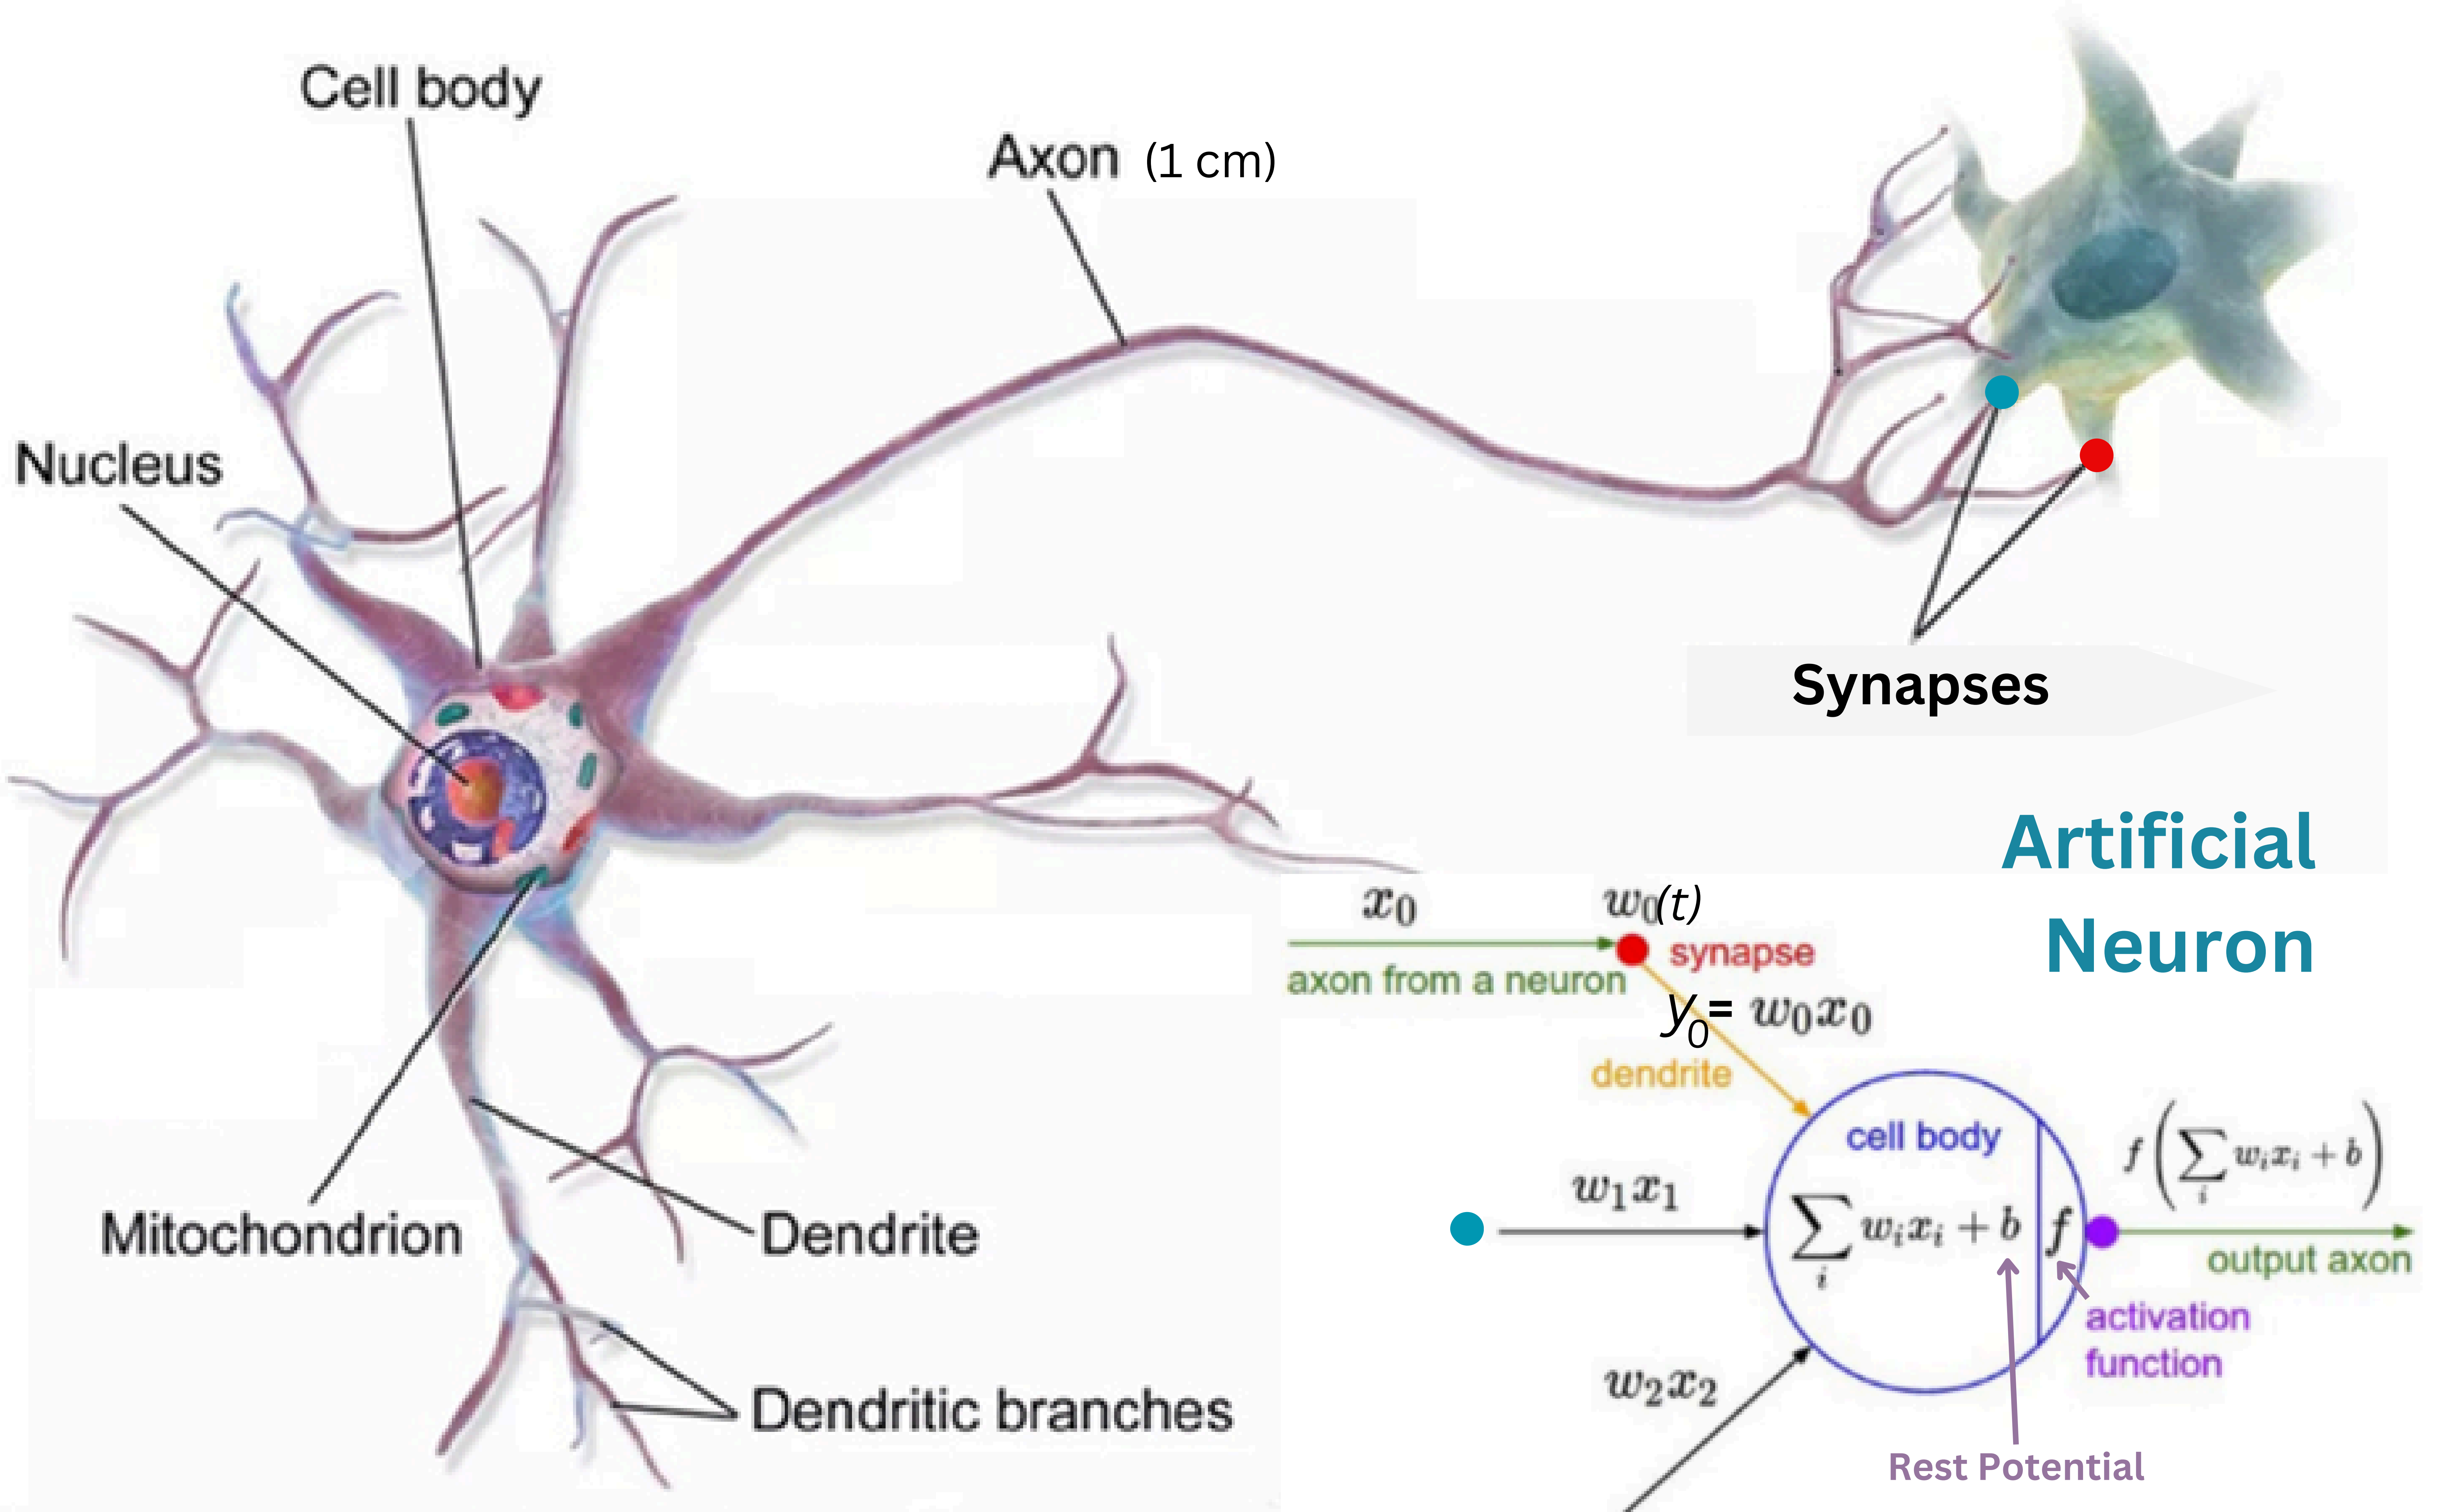
Cell body
Axon (1 cm)
Nucleus
Synapses
Mitochondrion
Dendrite
Dendritic branches
Artificial Neuron
x_0
$w_0(t)$
synapse
axon from a neuron
$y_0 = w_0 x_0$
dendrite
cell body
$w_1 x_1$
$\sum_i w_i x_i + b$
f
$f\left(\sum_i w_i x_i + b\right)$
output axon
activation function
$w_2 x_2$
Rest Potential

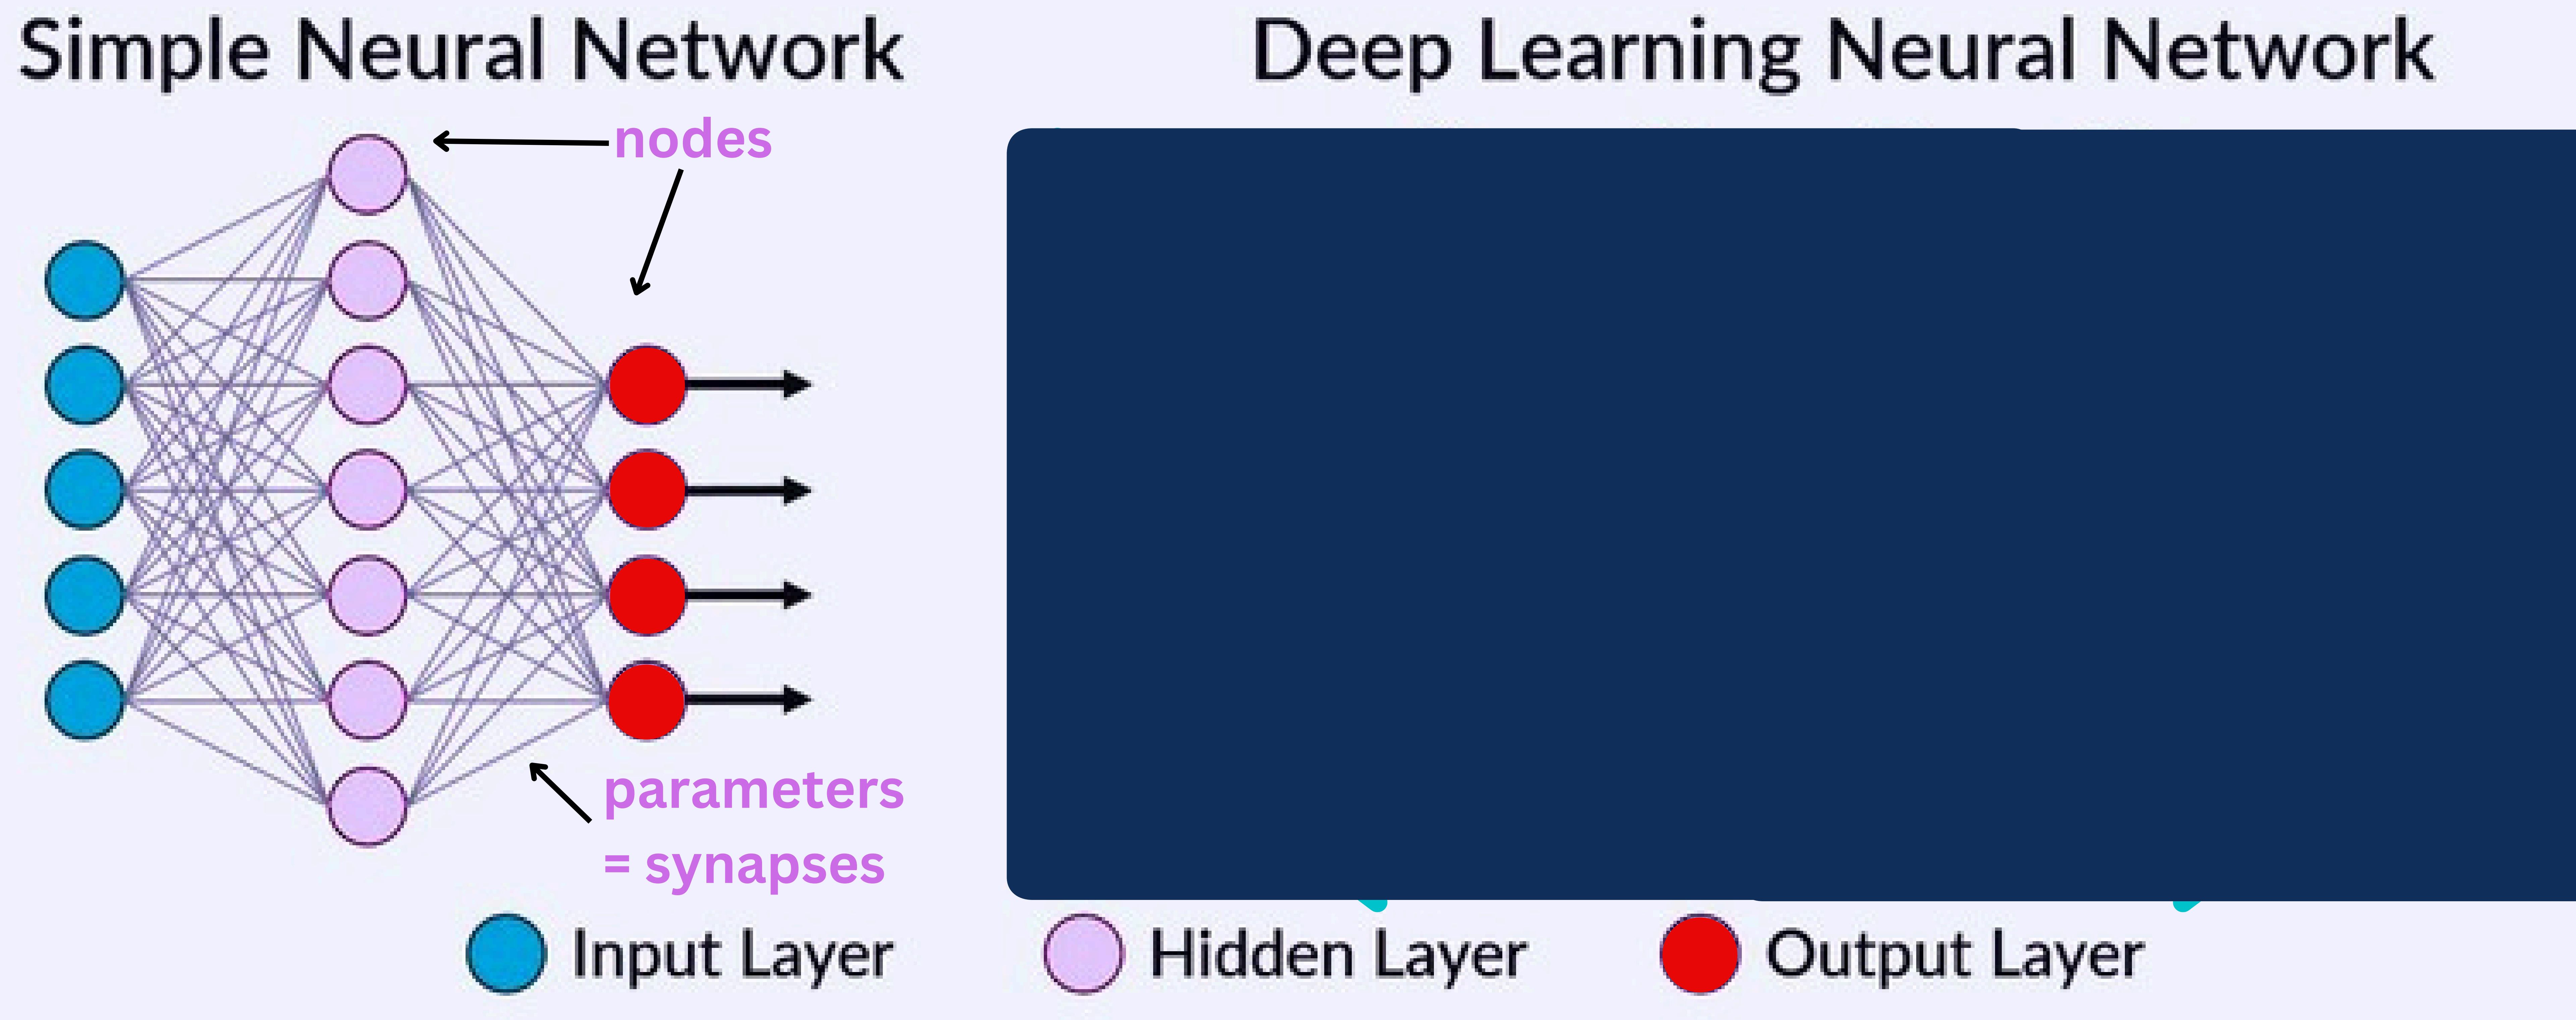

All neural networks are made up by different hidden layers of artificial neurons, also called nodes. Each parameter represent a synapsis

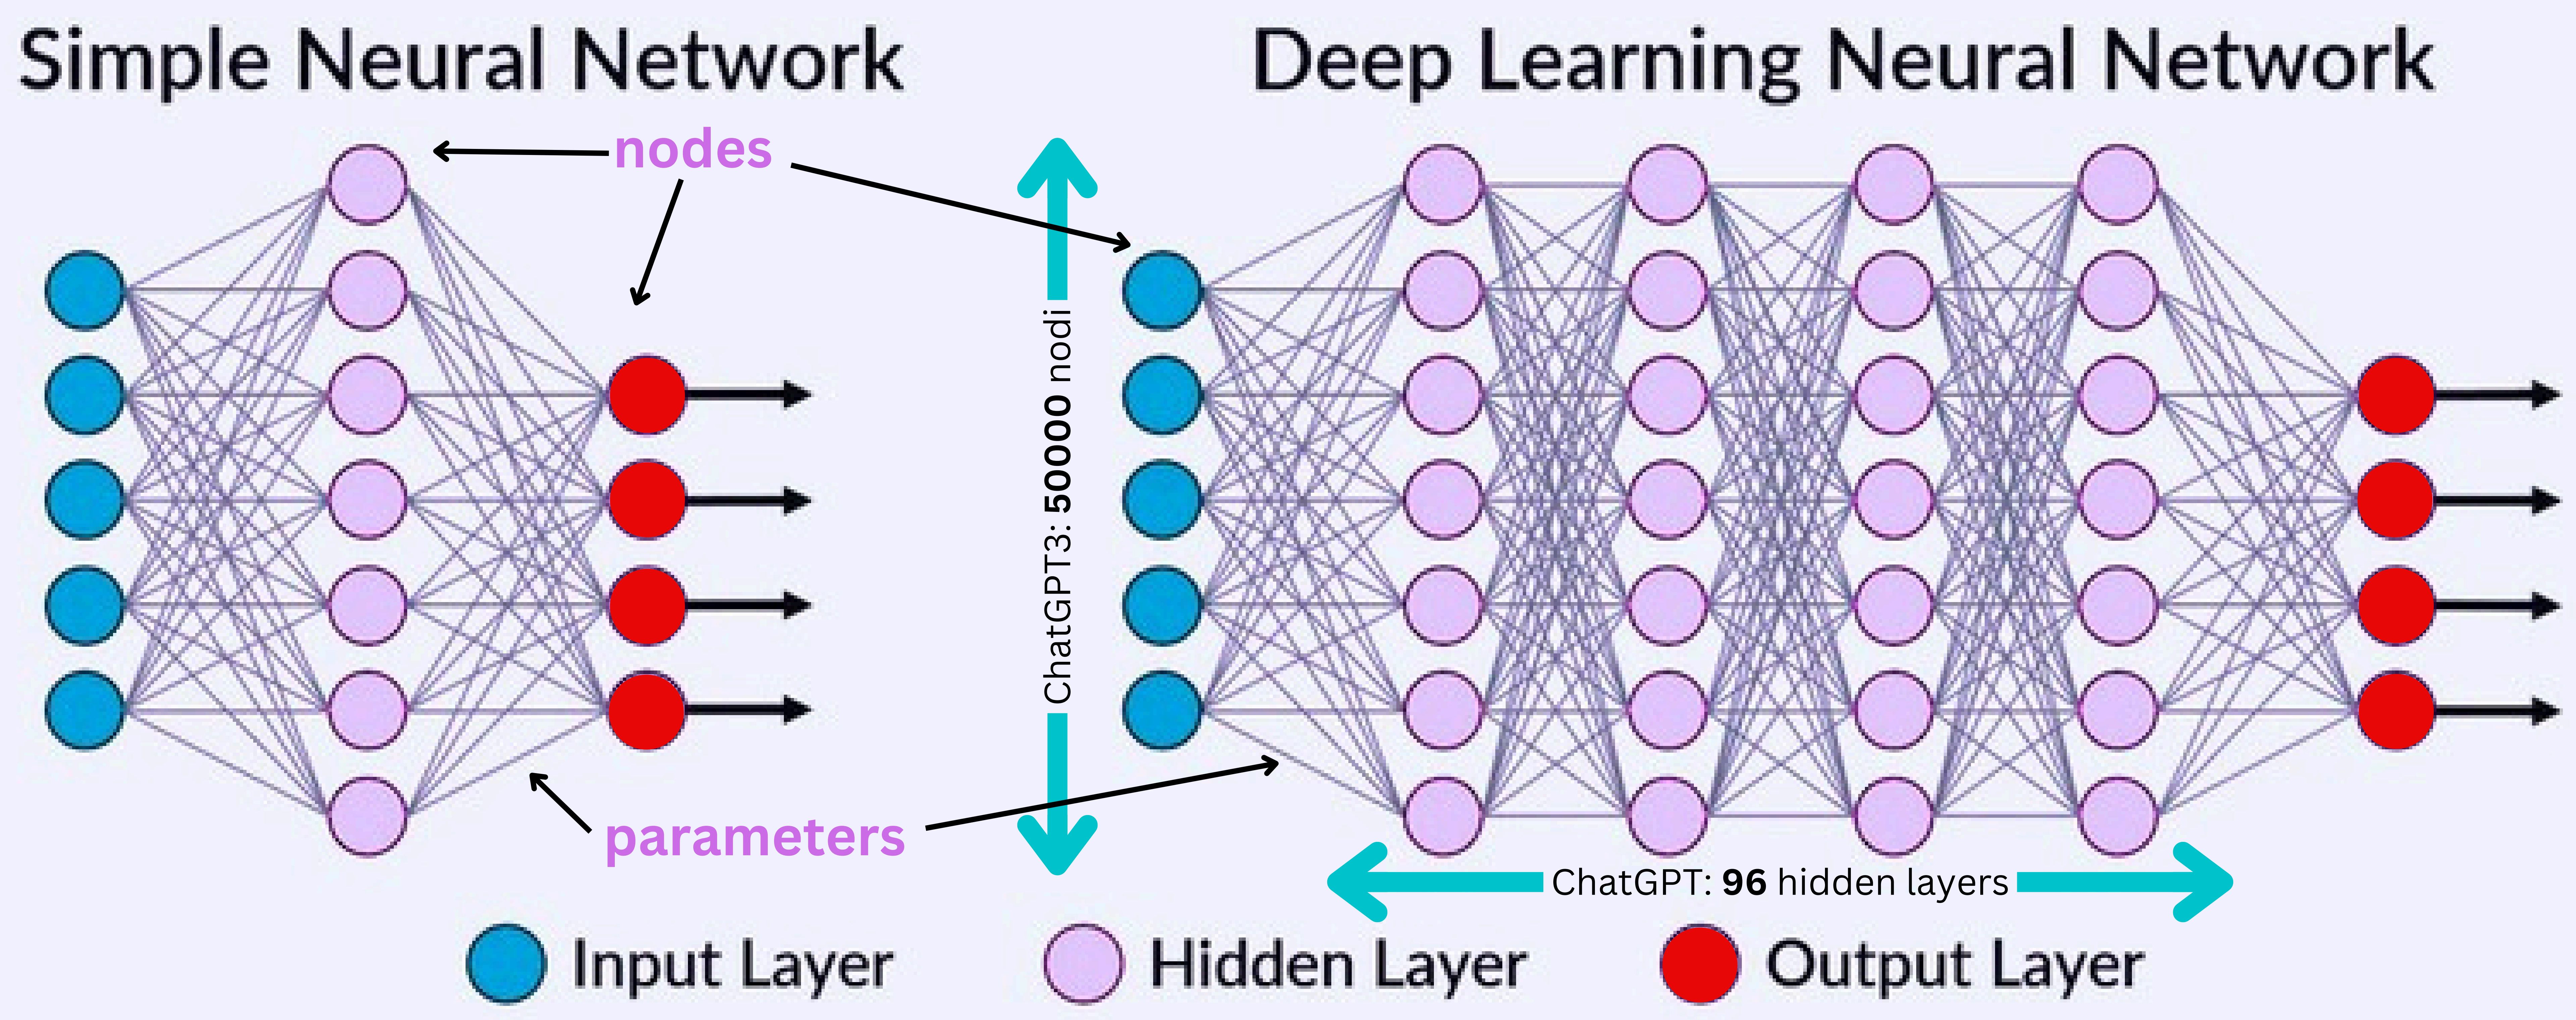

Neural networks with 4 or more hidden layers are called Deep Learning Neural Network (DNN)

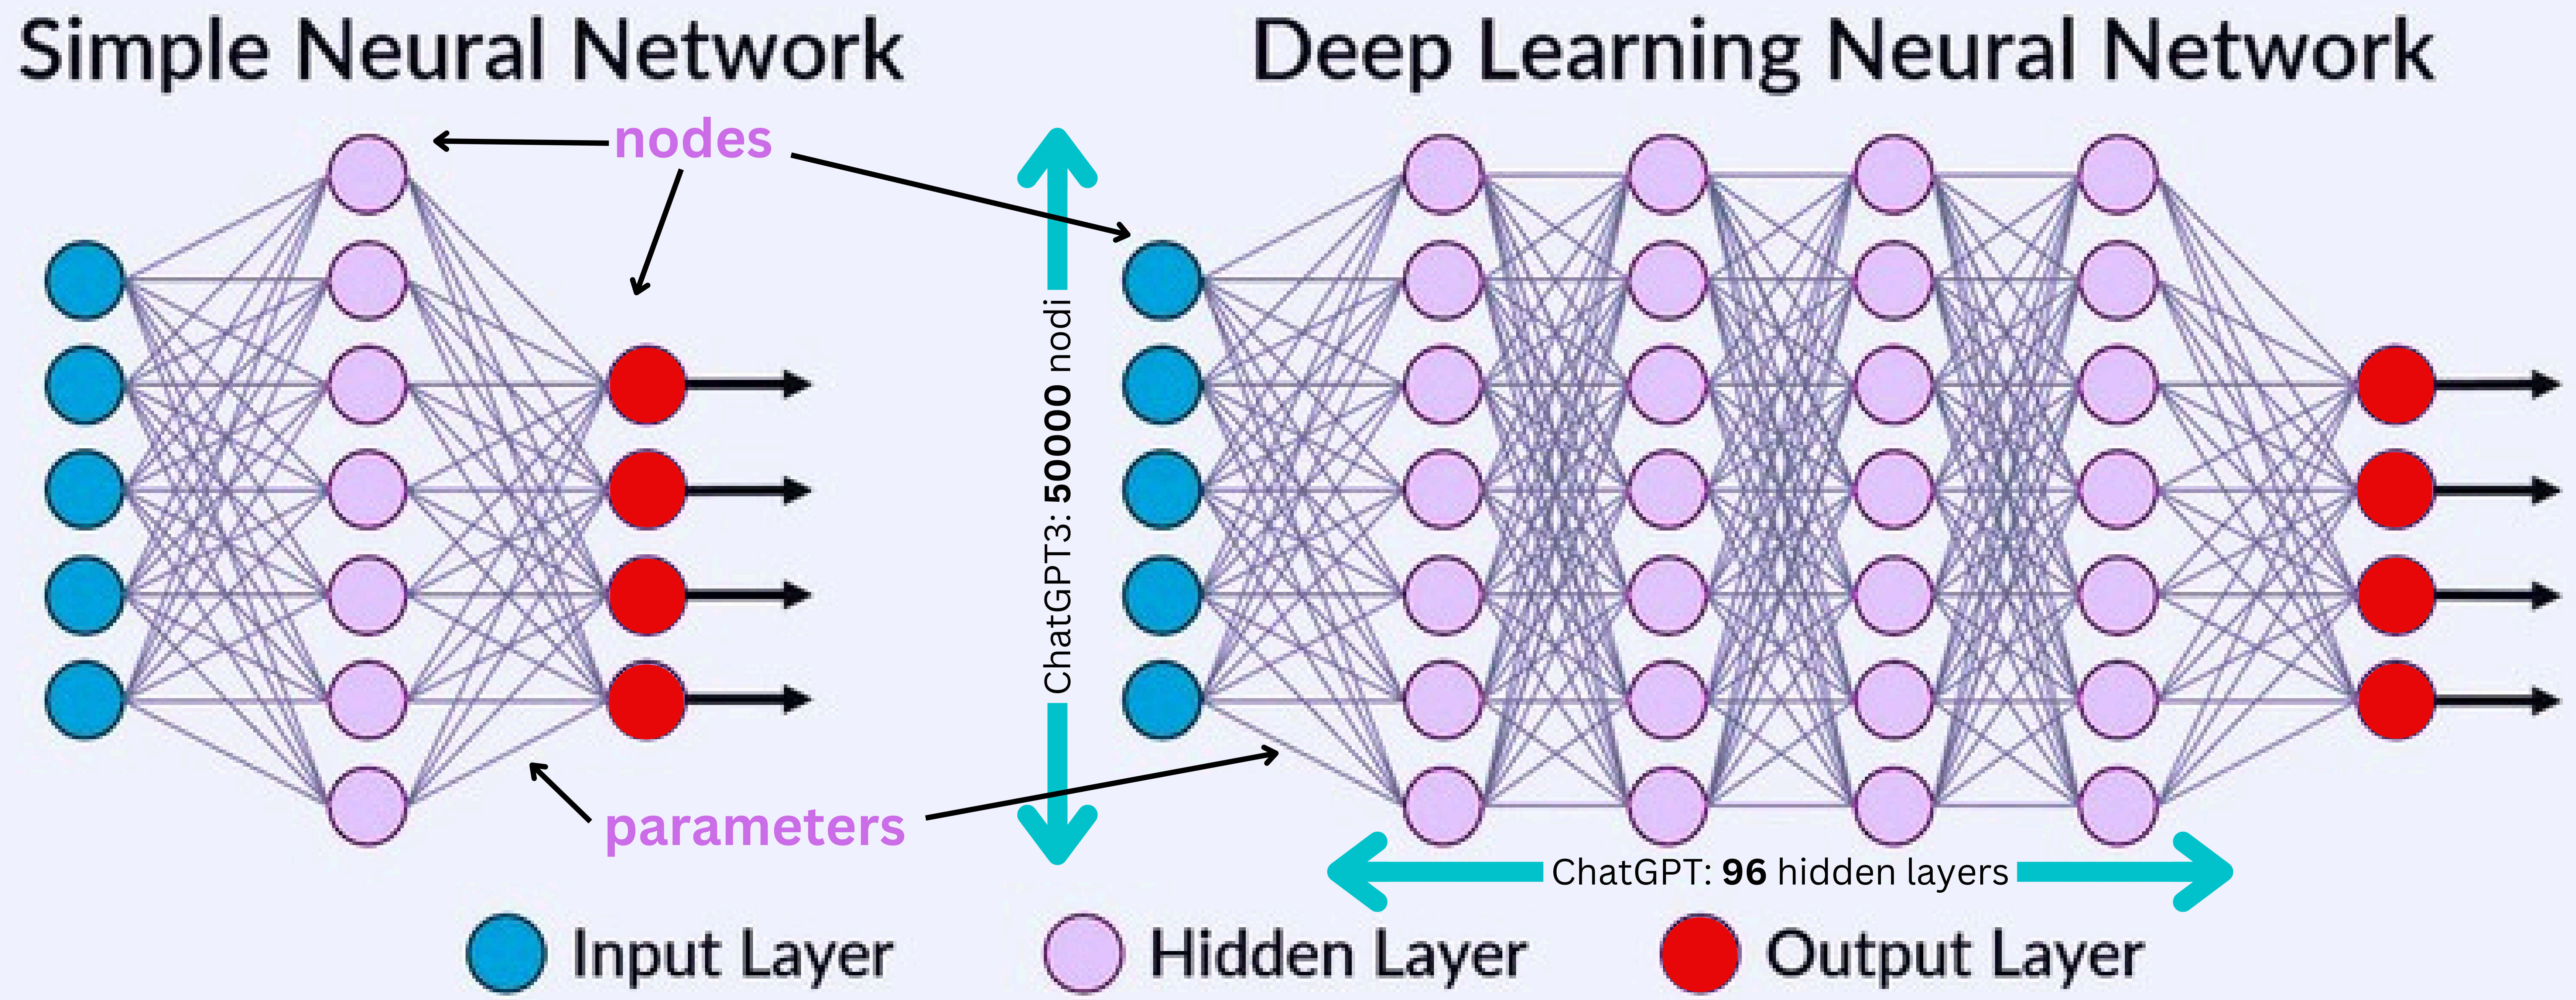

All Neural Networks in the world (including Generative AI models) follow this simple structure, with no exceptions.

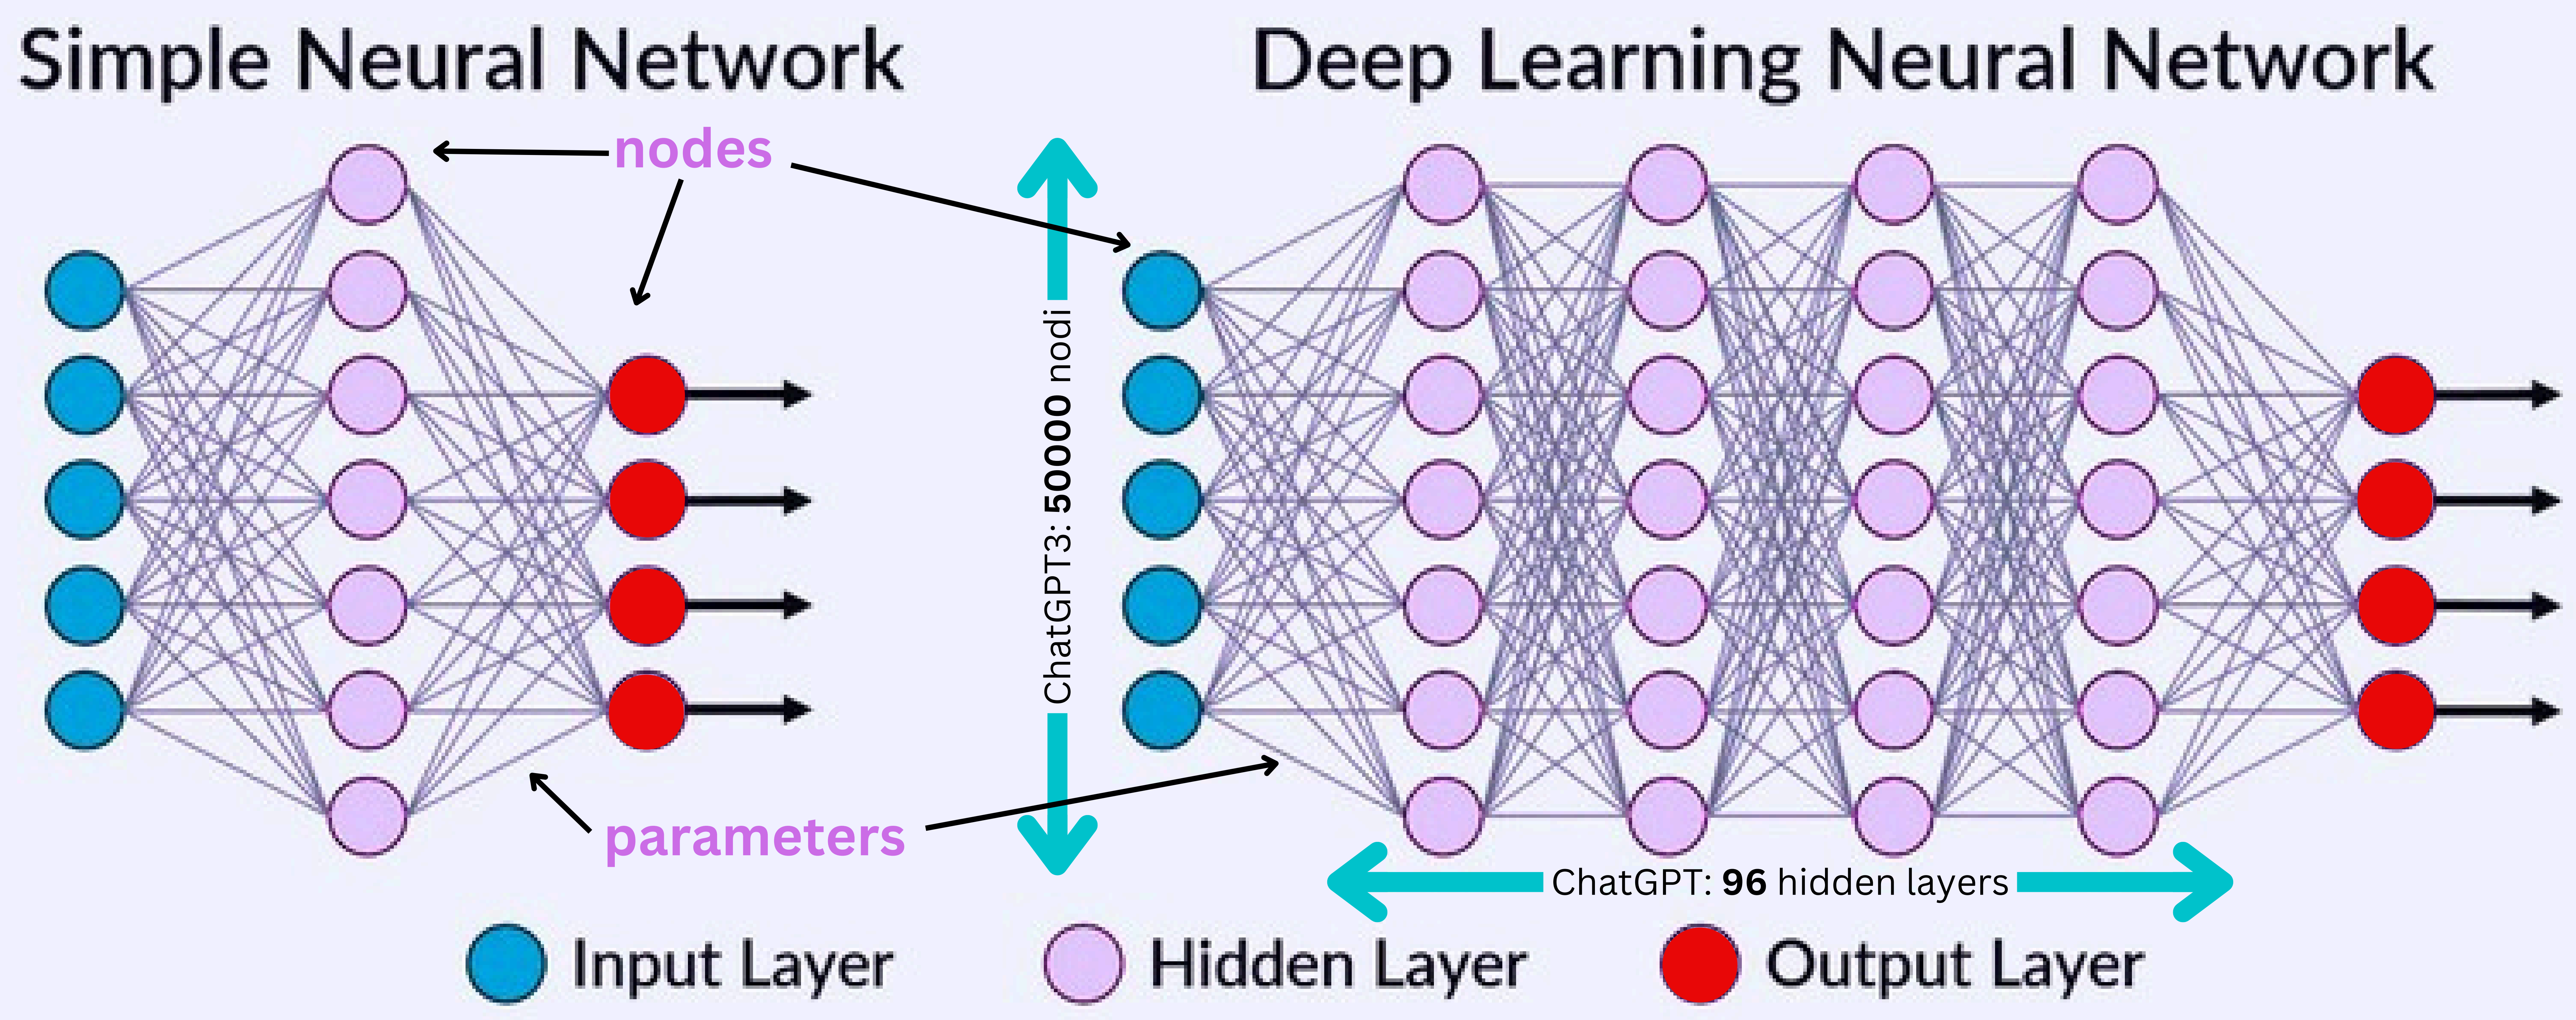

Notice how the structure of our brain is very different form this ordered network: our neurons are not organized in layers, can connect with any other neuron of the brain, even multiple times and even with themselves!

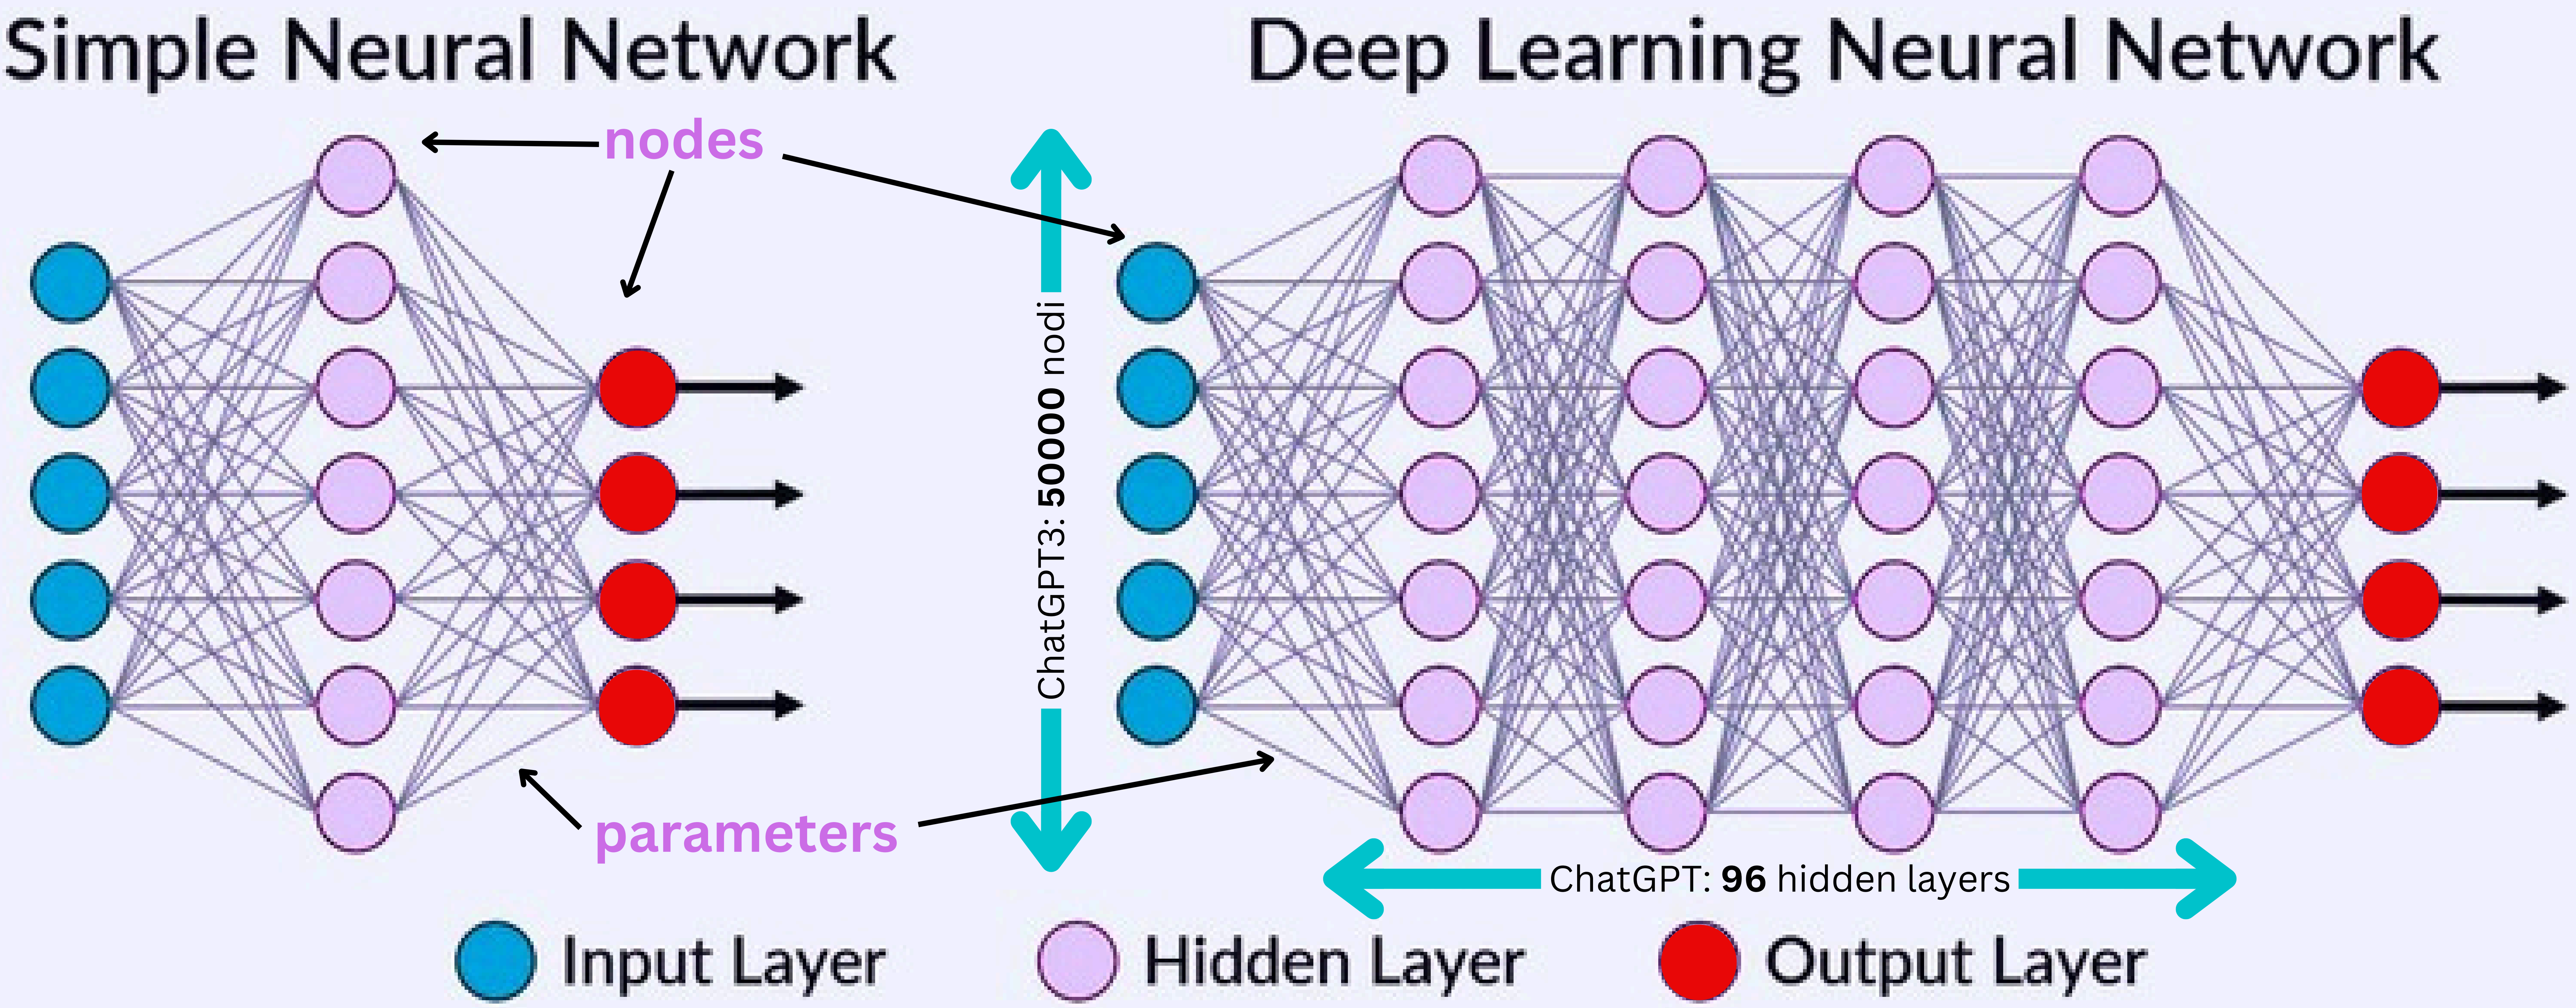

Our brain doesn't employ Activation functions and all other techniques described in the lesson and the next one (backpropagation, embeddings, etc)

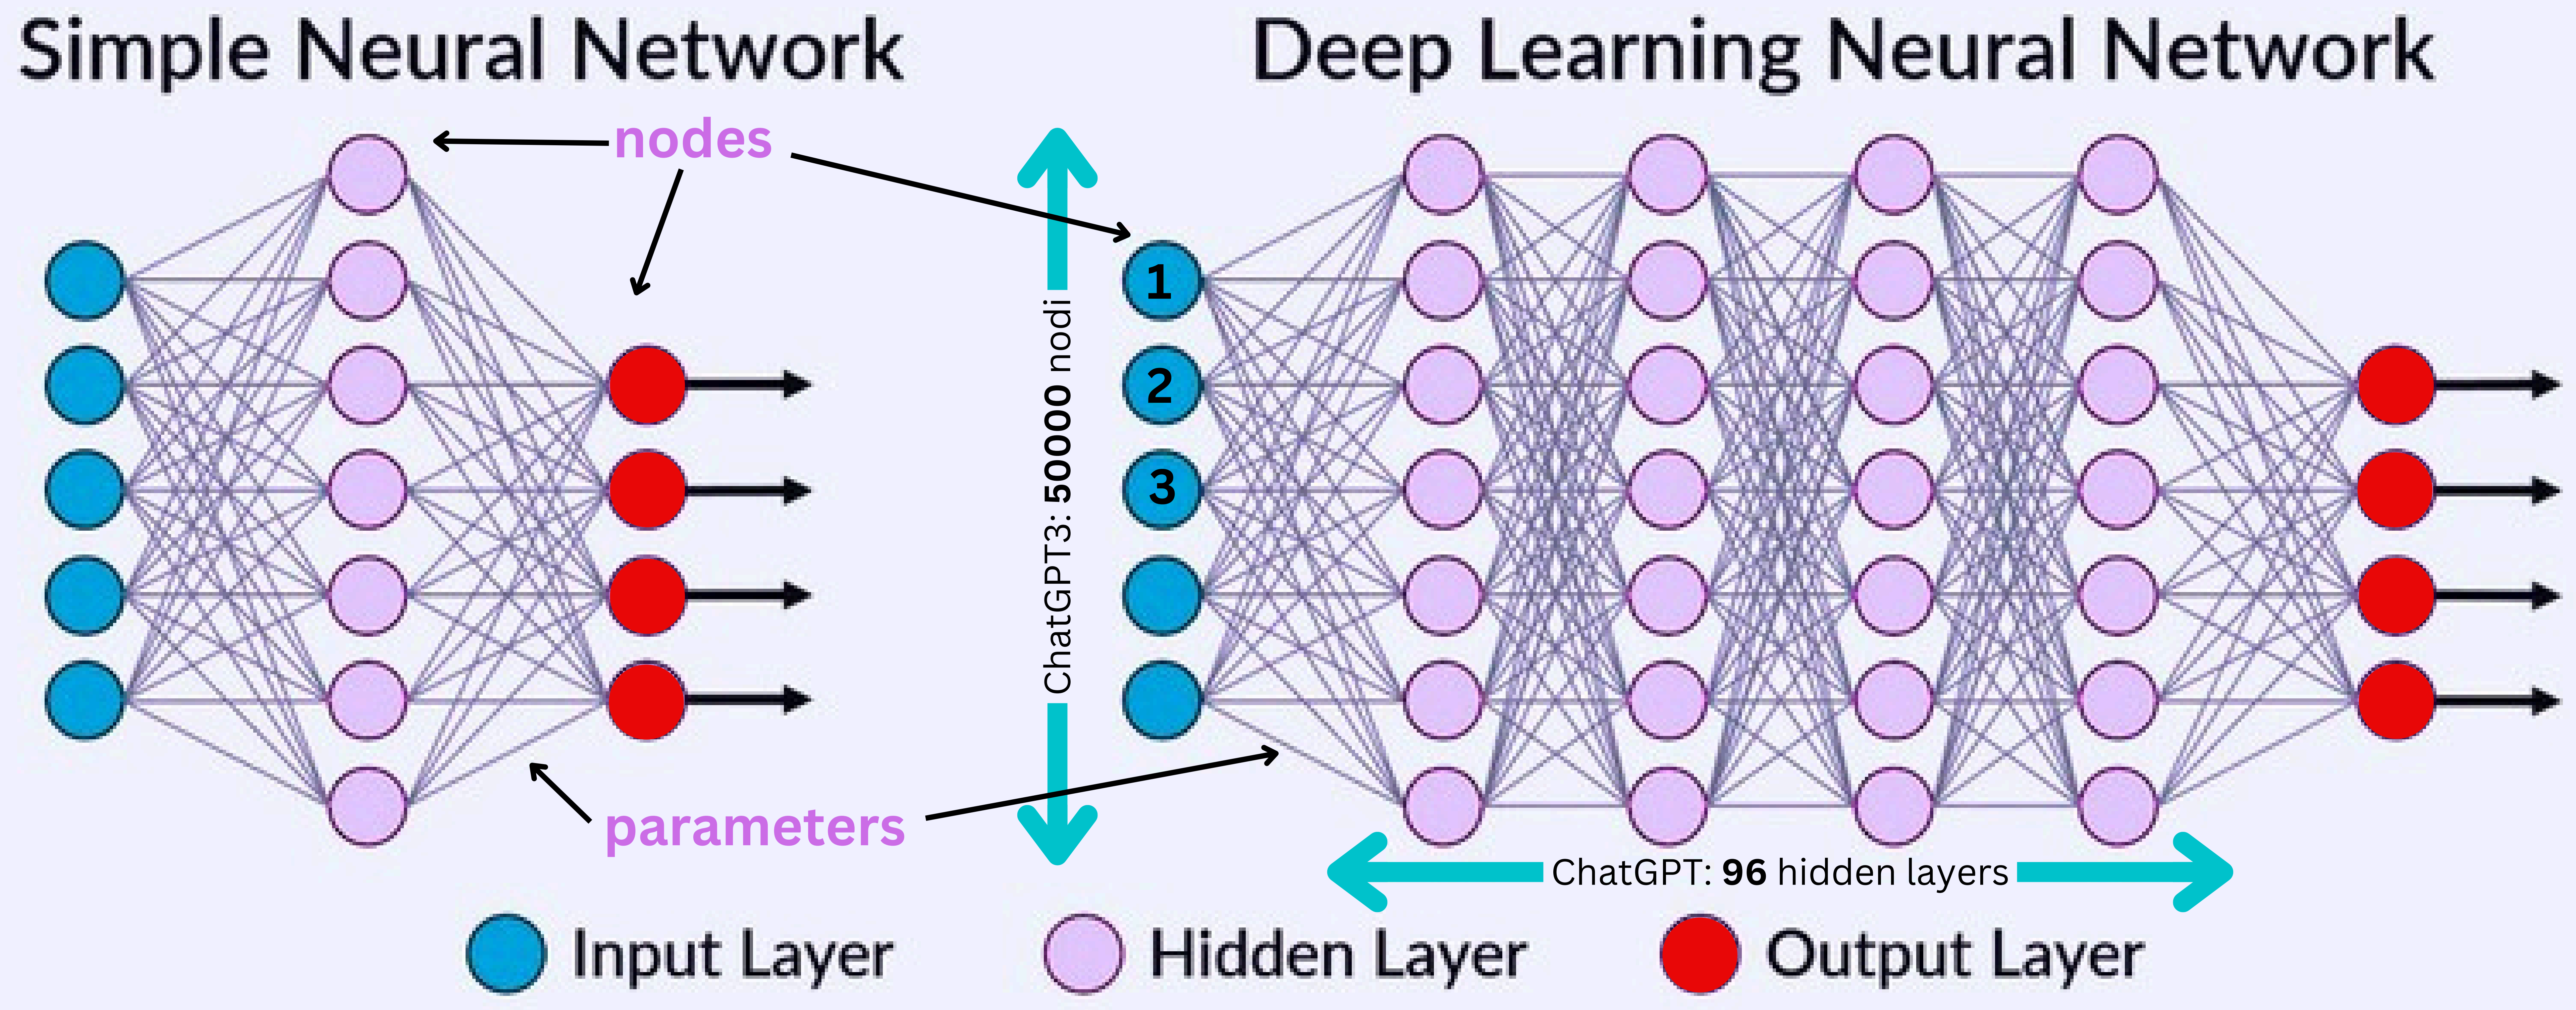

To run a prompt of a LLM, you take the words of the prompt and forecast the following one

Forward Propagation

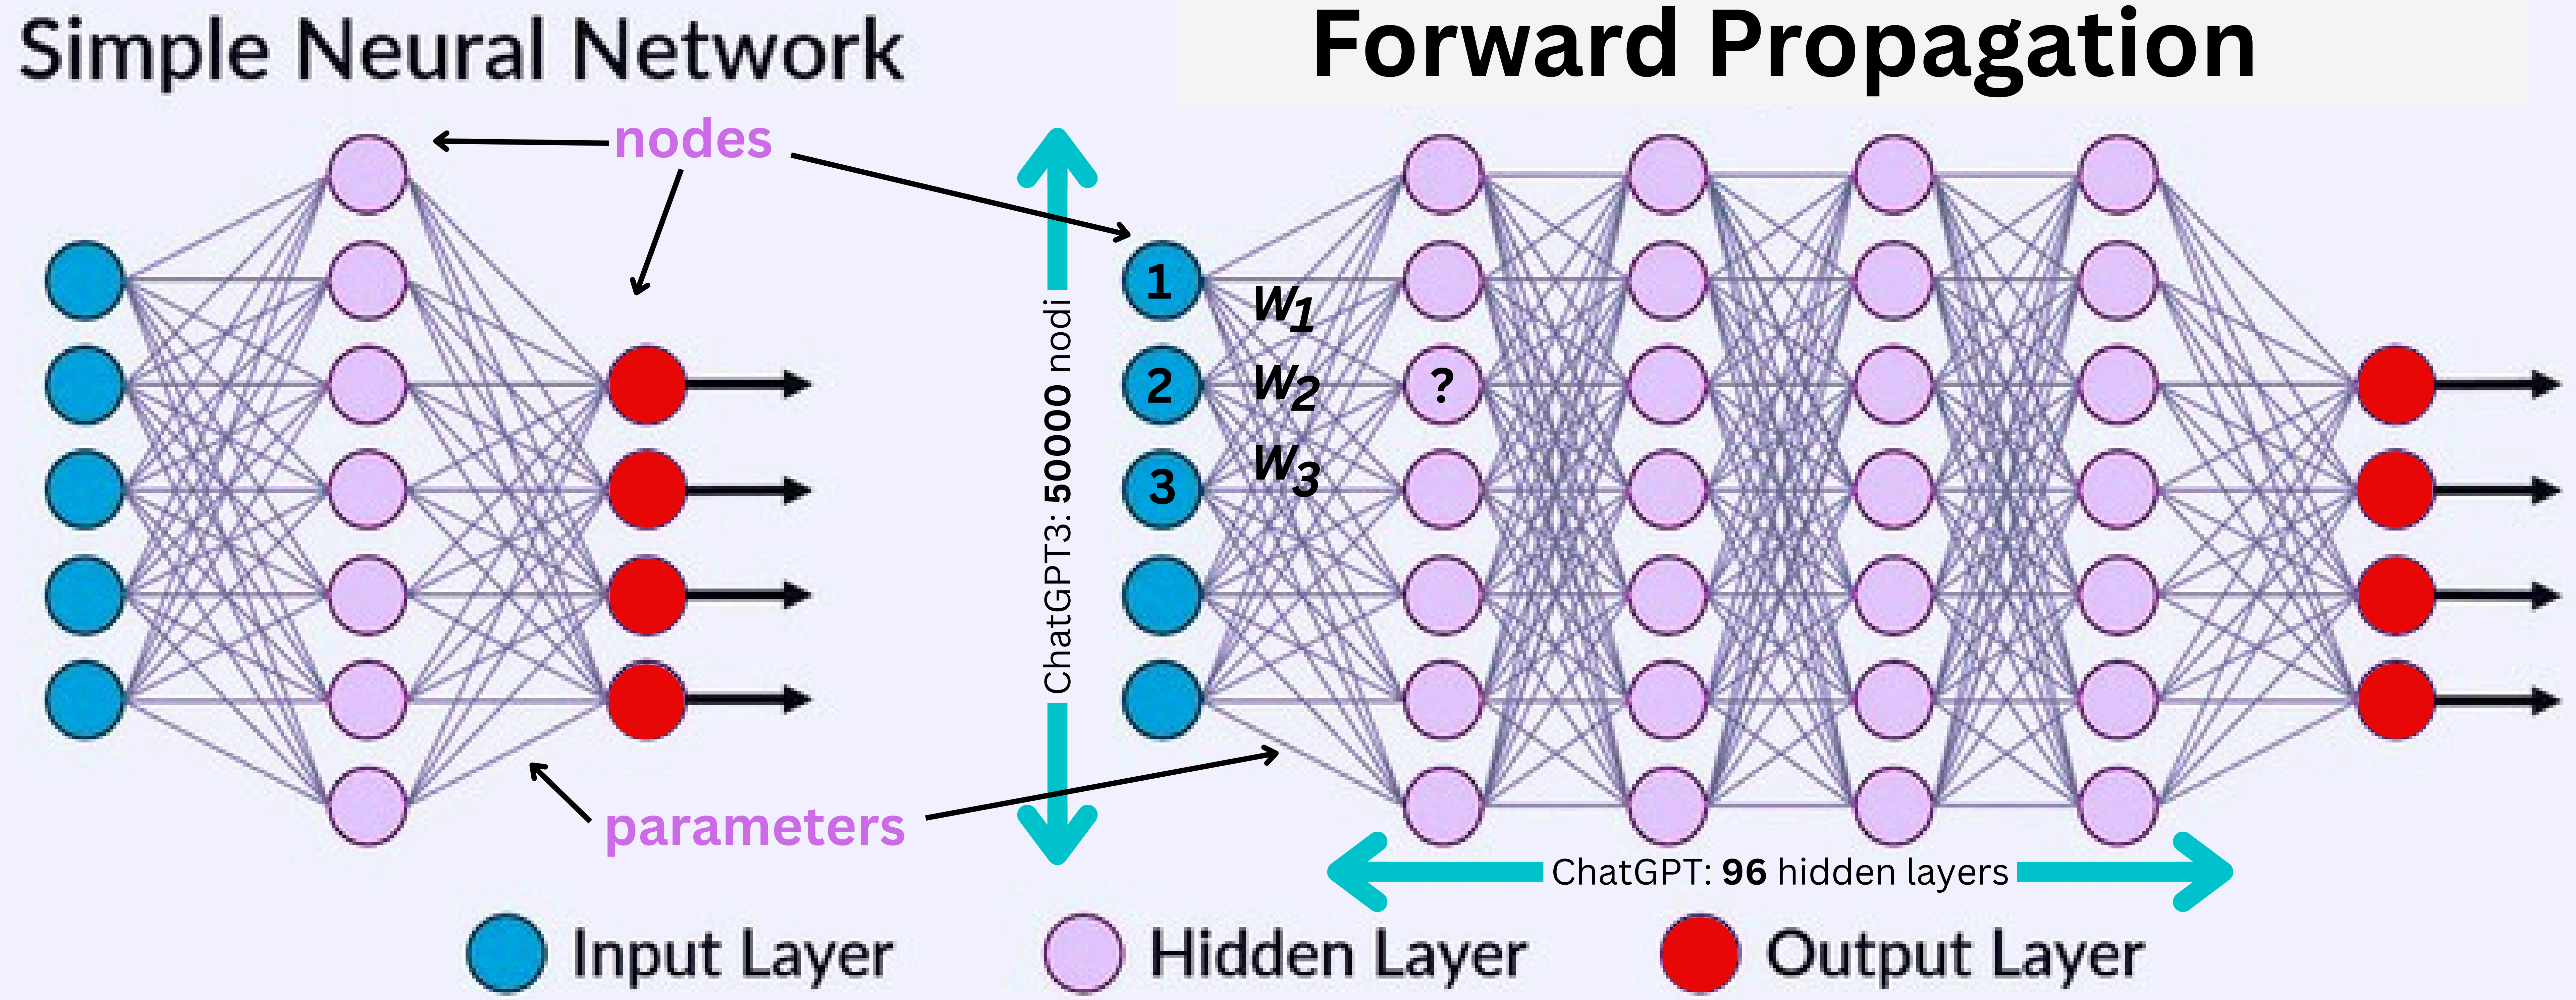

This process is called "Forward Propagation" or also "inference" and is done with simple matrix multiplications of the parameters W and the node values

Forward Propagation

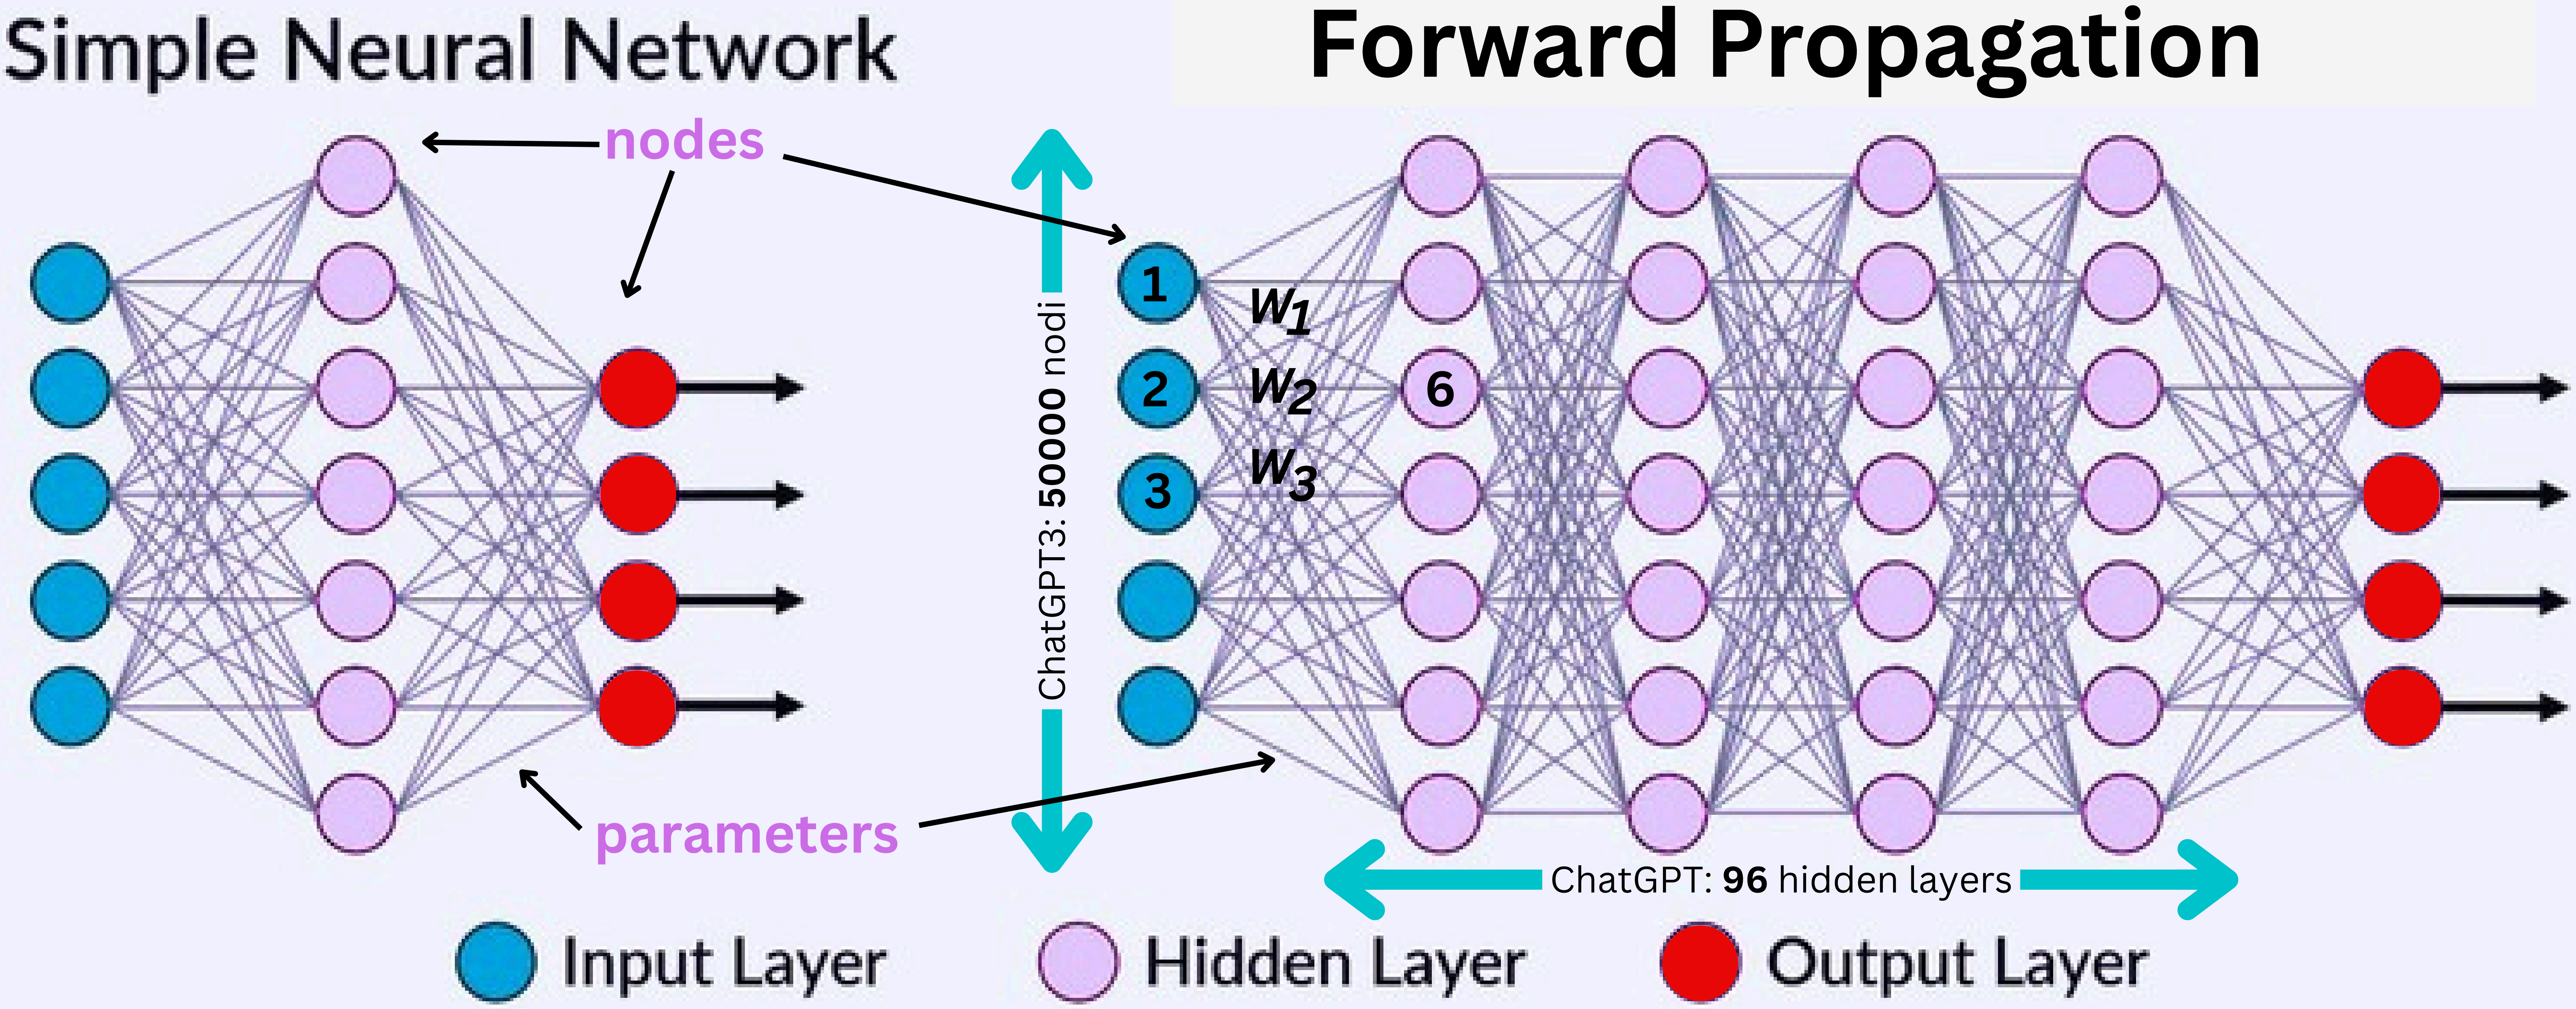

For example, the number “6” in the first hidden layer is the sum of:

1*W1 + 2*W2 + 3*W3

Forward Propagation

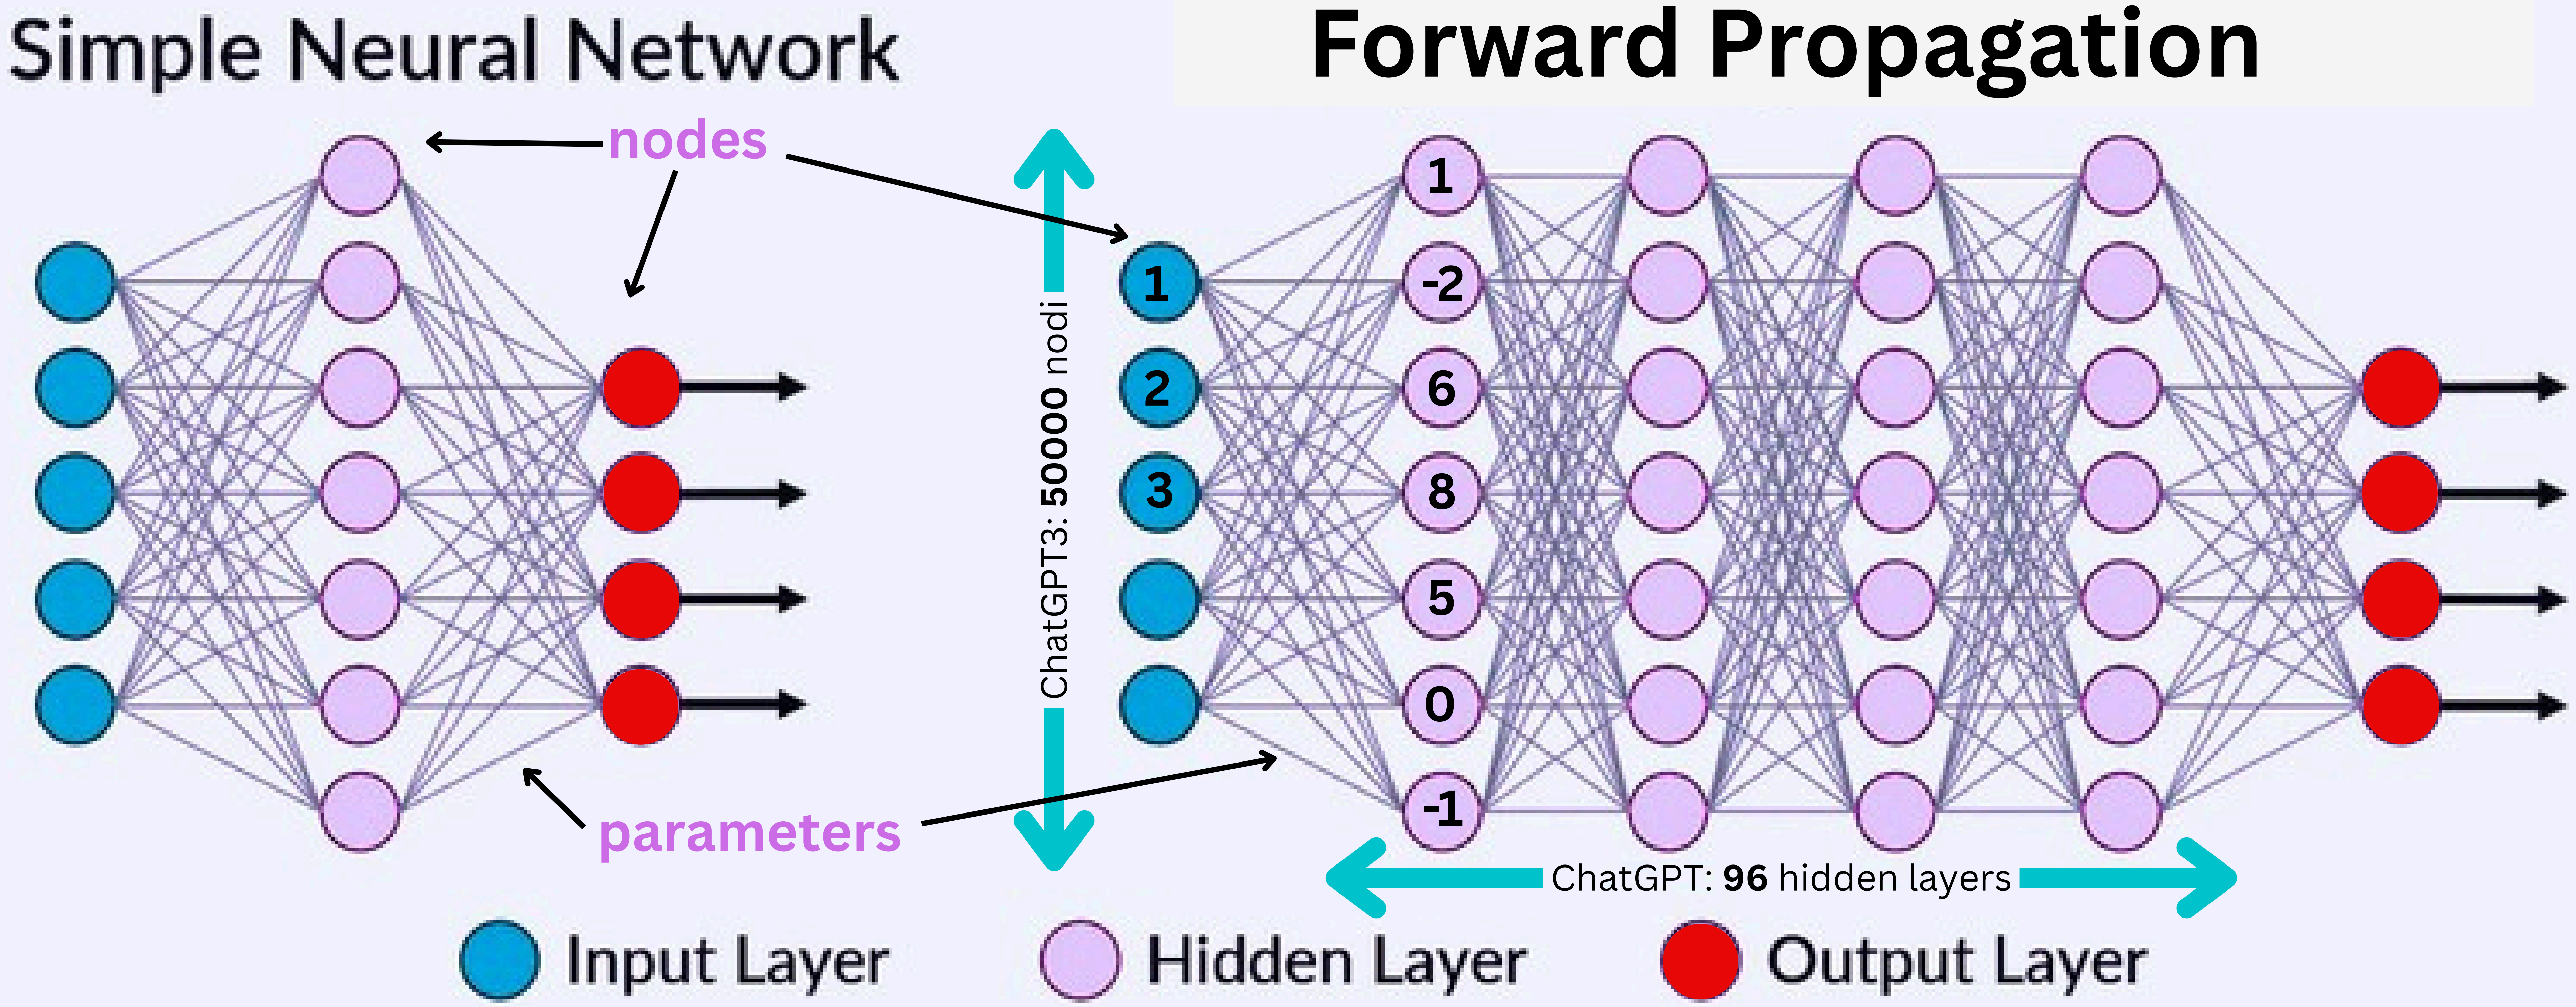

This process is called "Forward Propagation" or also "inference" and is done with simple matrix multiplications of the parameters W and the node values

Forward Propagation

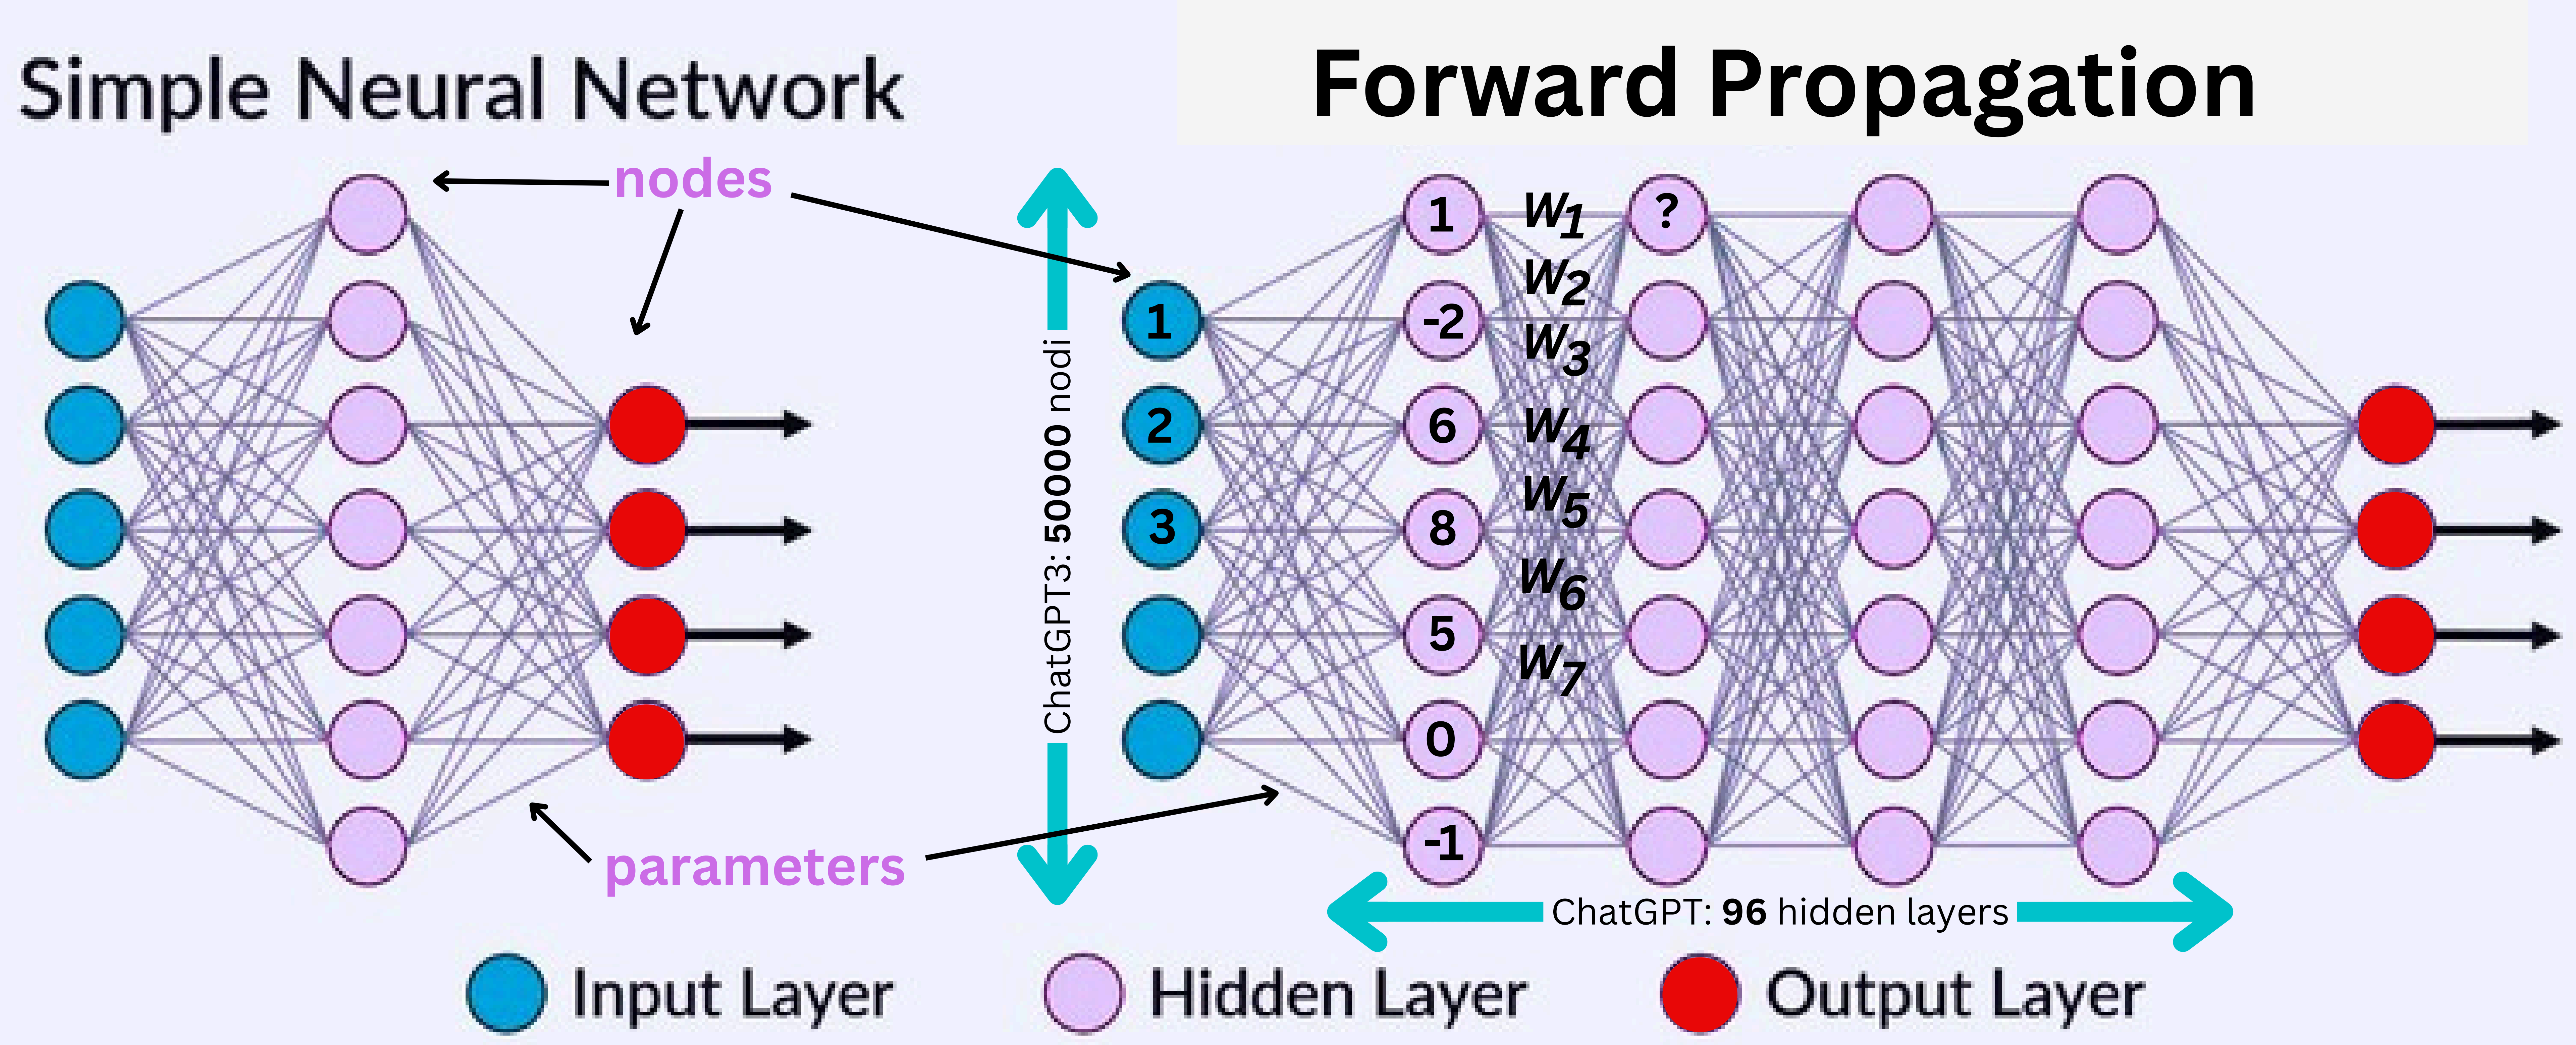

This process is called "Forward Propagation" or also "inference" and is done with simple matrix multiplications of the parameters W and the node values

Forward Propagation

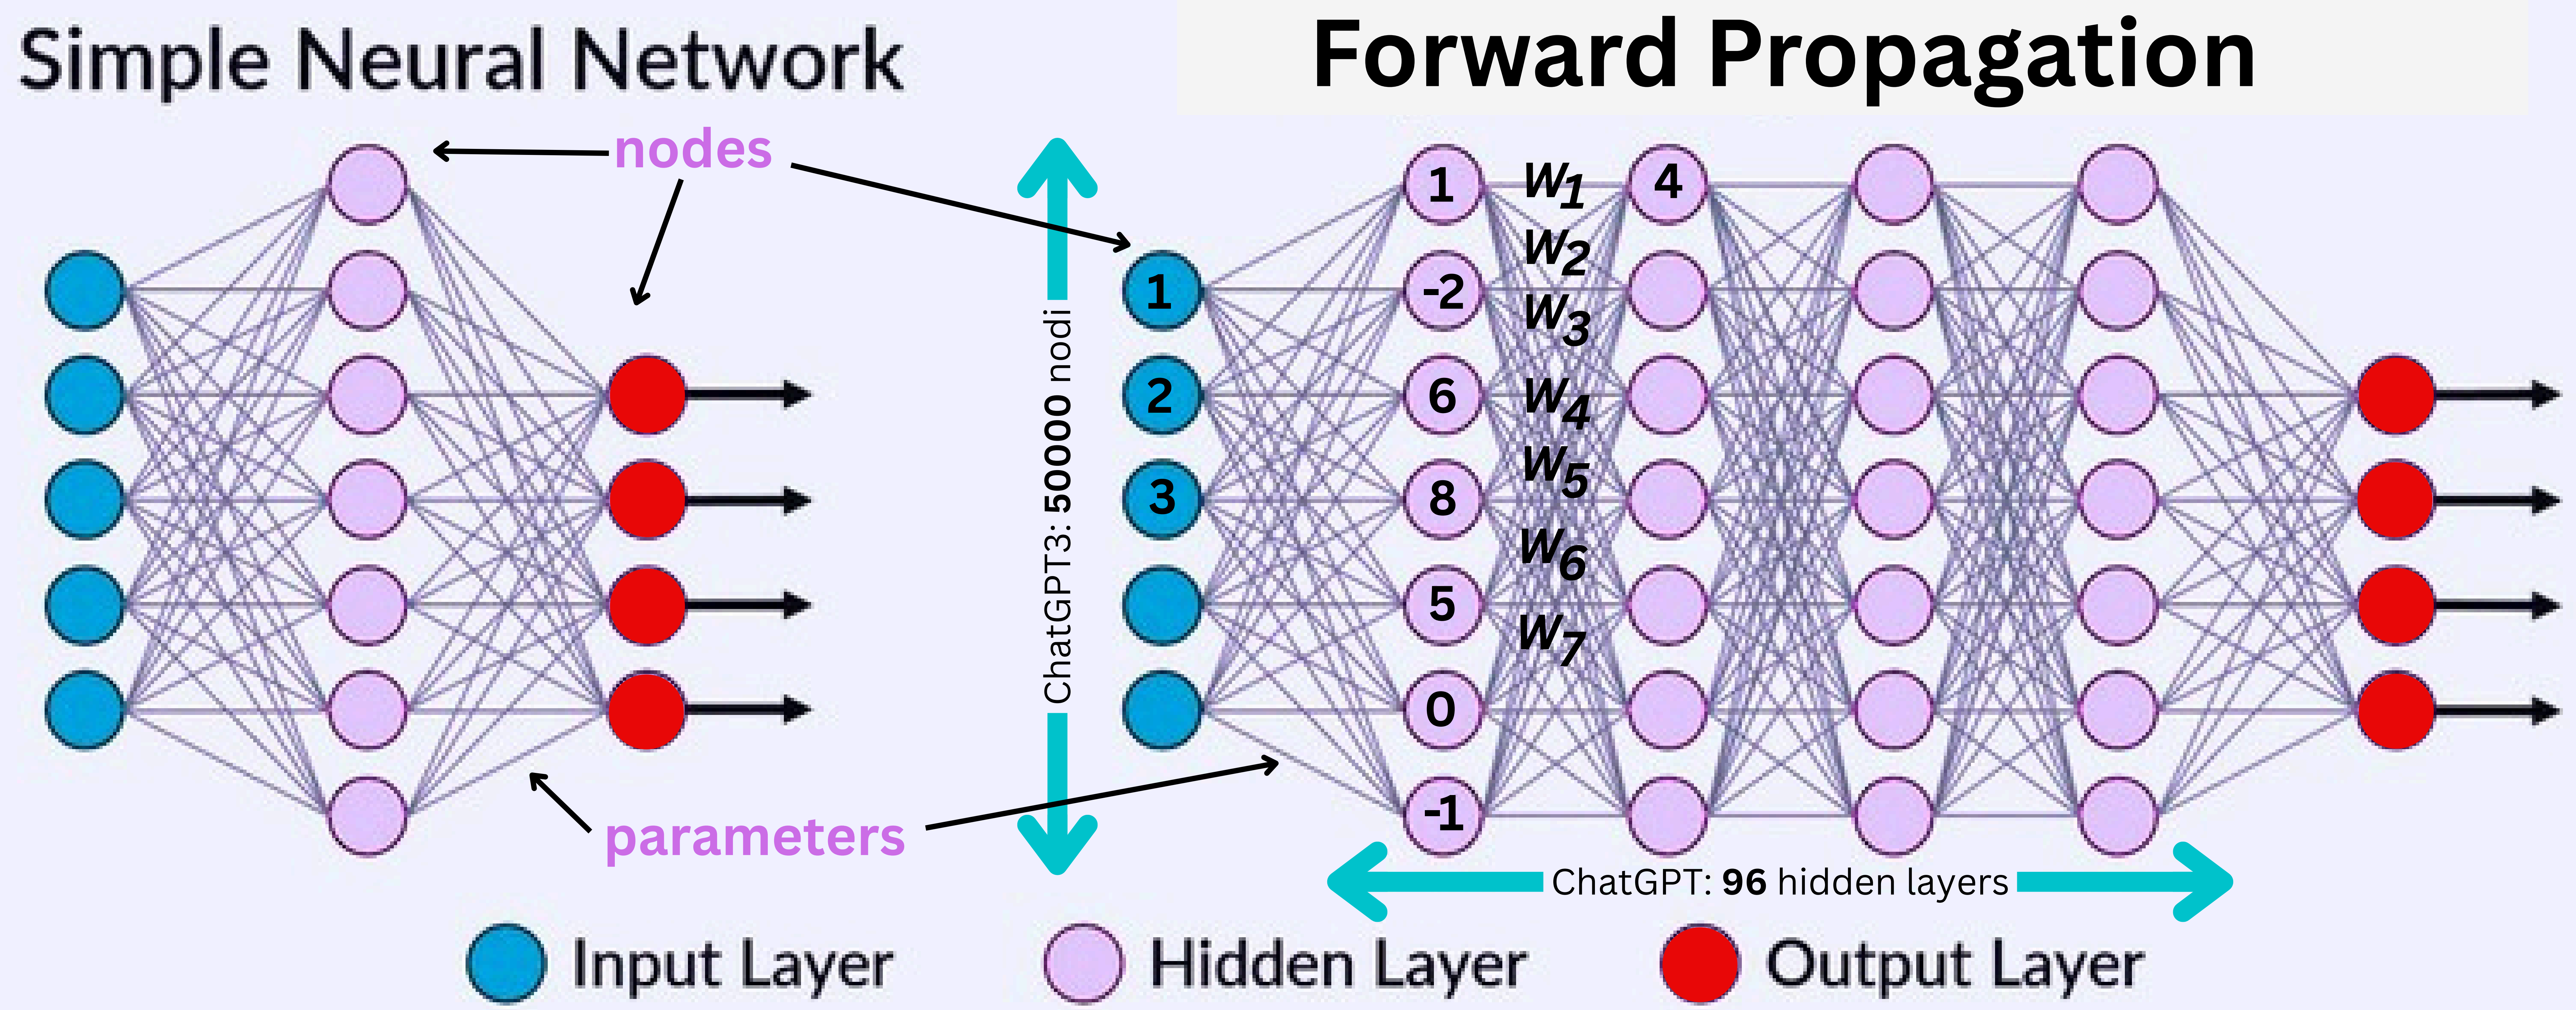

This process is called "Forward Propagation" or also "inference" and is done with simple matrix multiplications of the parameters W and the node values

Simple Neural Network

Forward Propagation

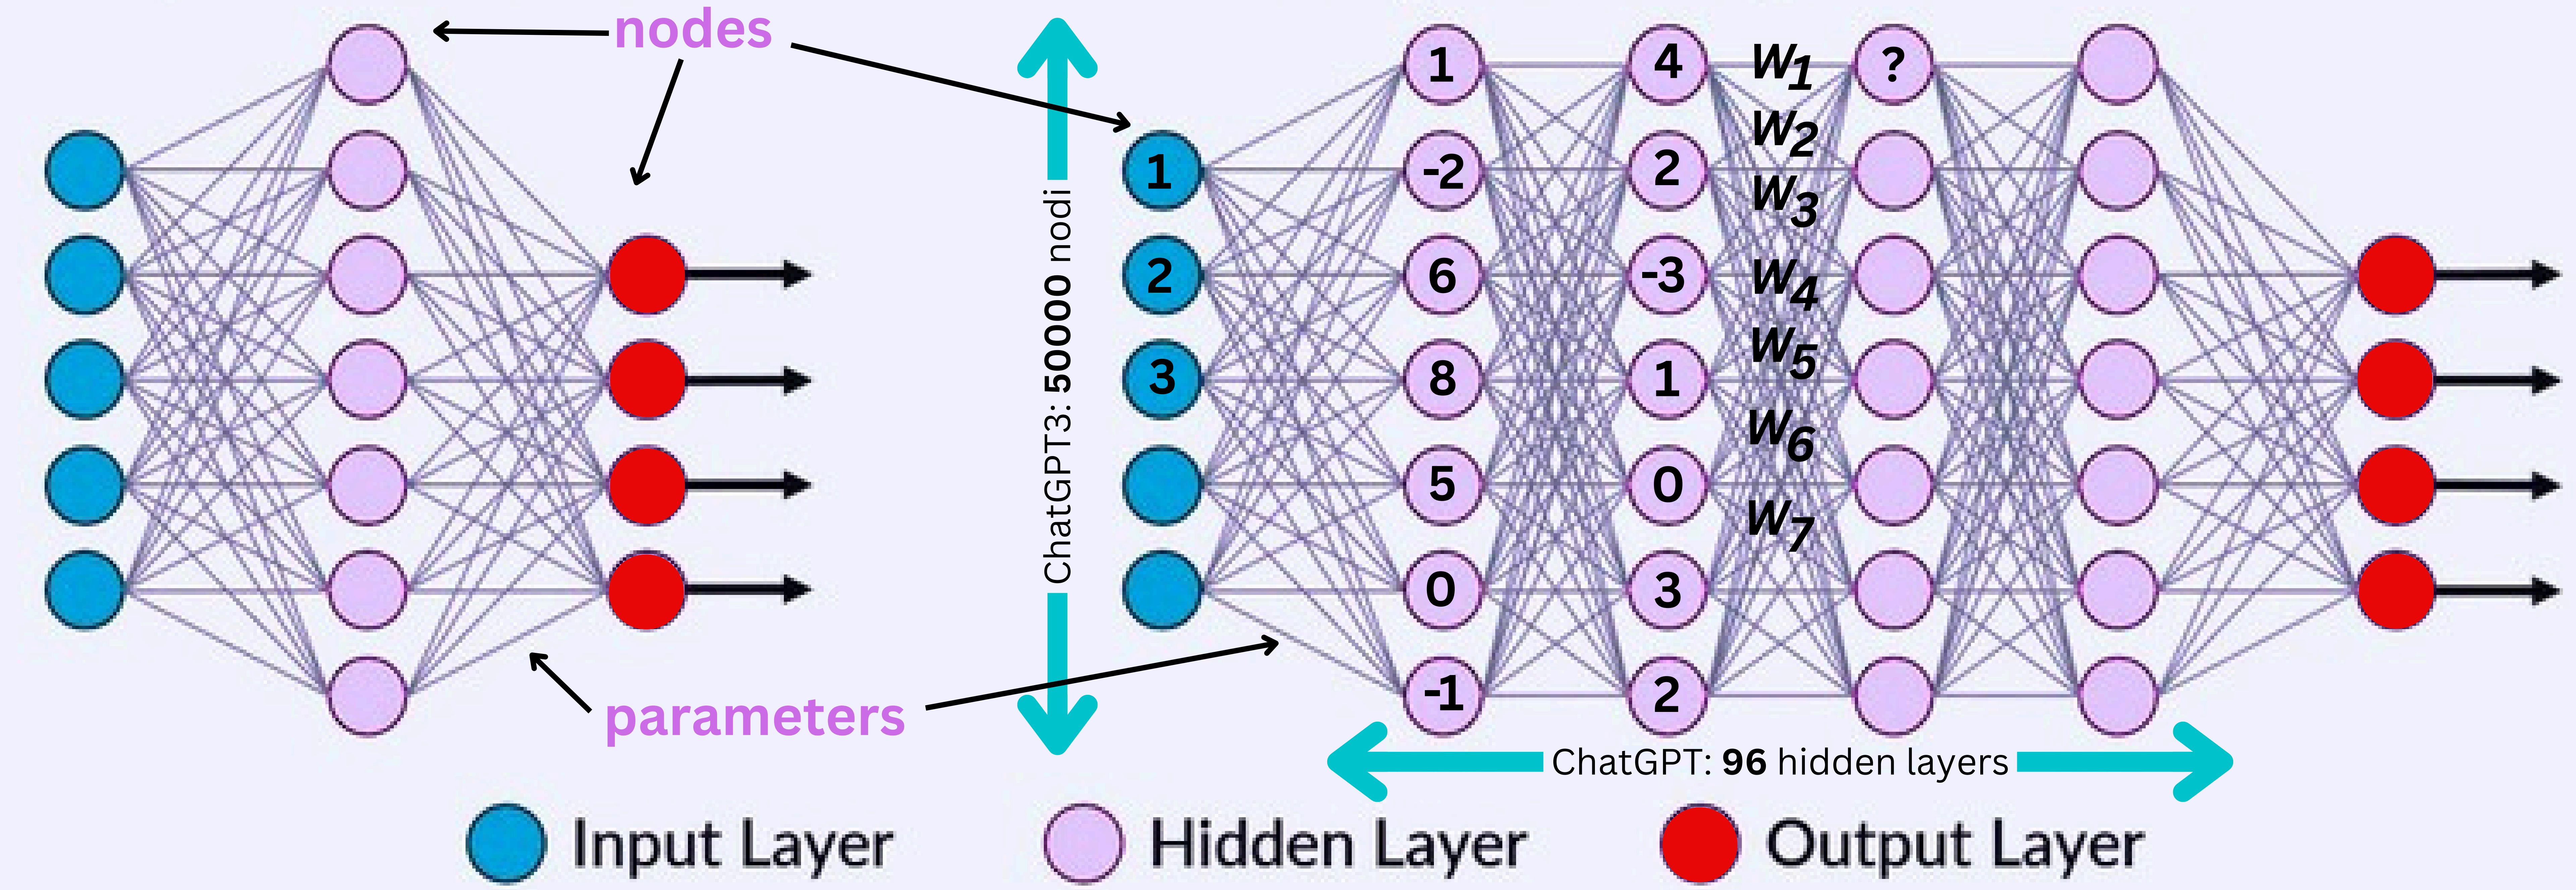

This process is called "Forward Propagation" or also "inference" and is done with simple matrix multiplications of the parameters W and the node values

Forward Propagation

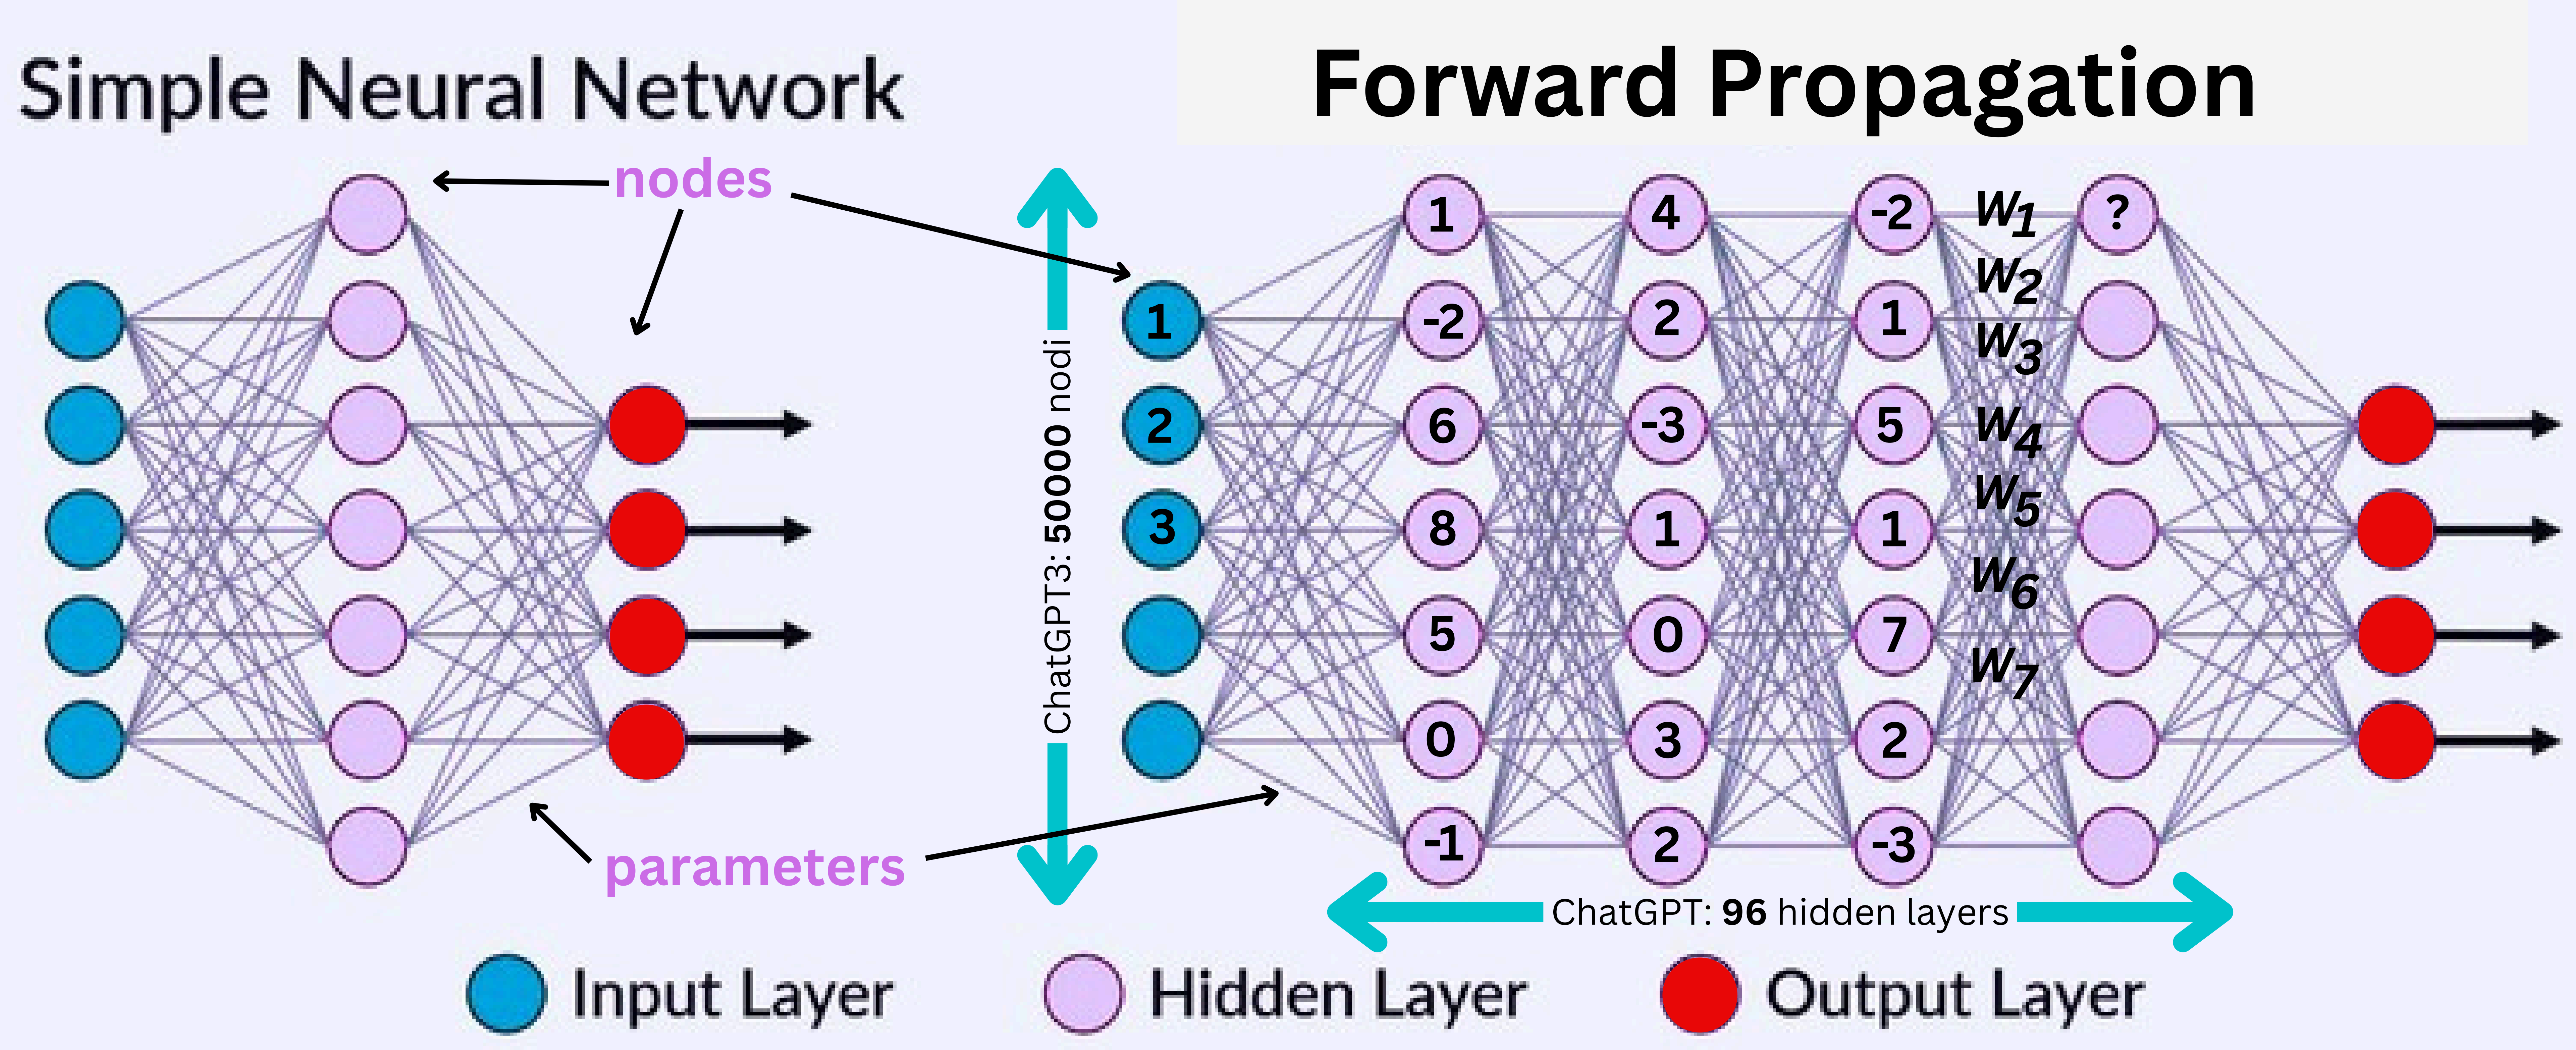

This process is called "Forward Propagation" or also "inference" and is done with simple matrix multiplications of the parameters W and the node values

Forward Propagation

Simple Neural Network

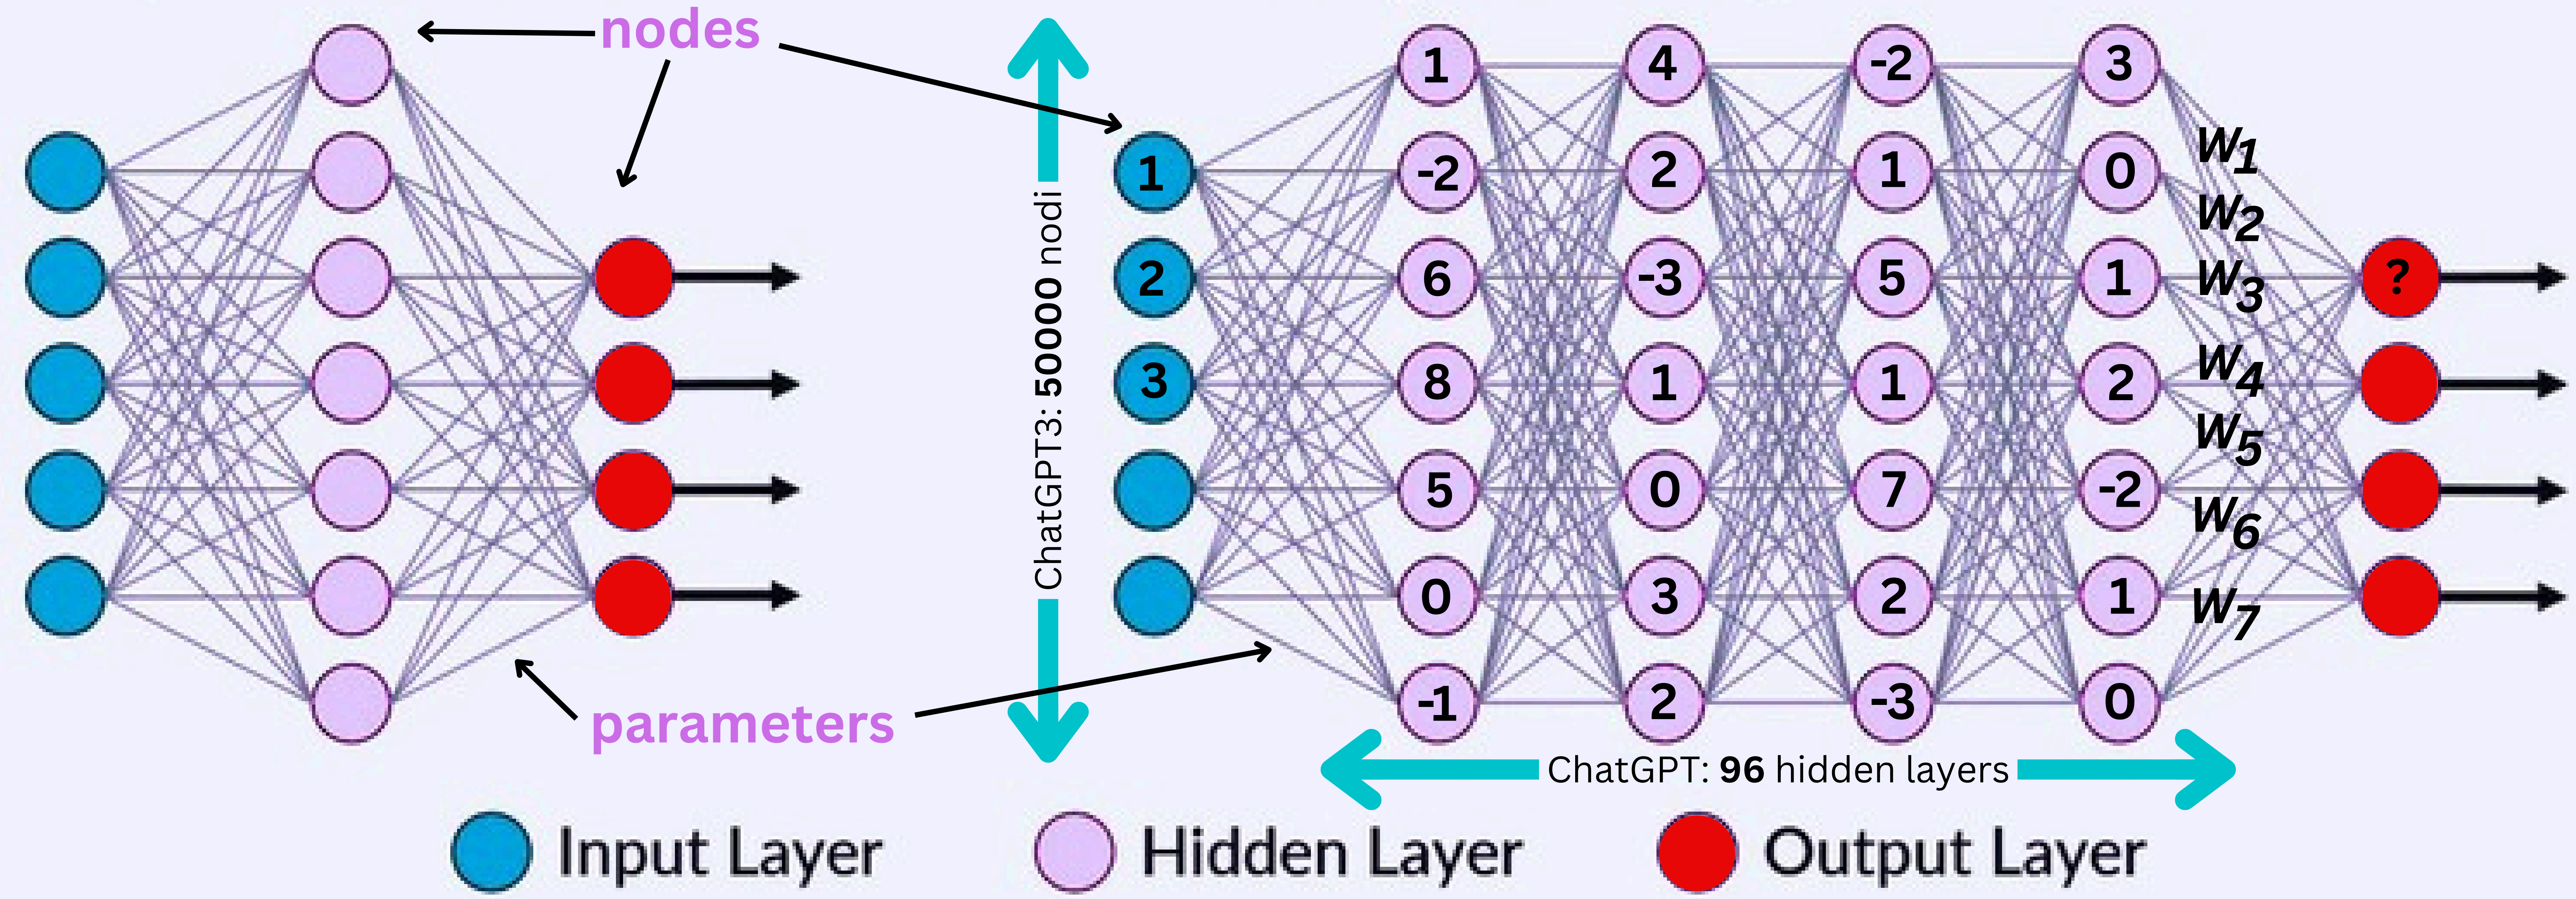

This process is called "Forward Propagation" or also "inference" and is done with simple matrix multiplications of the parameters W and the node values

Forward Propagation

Simple Neural Network

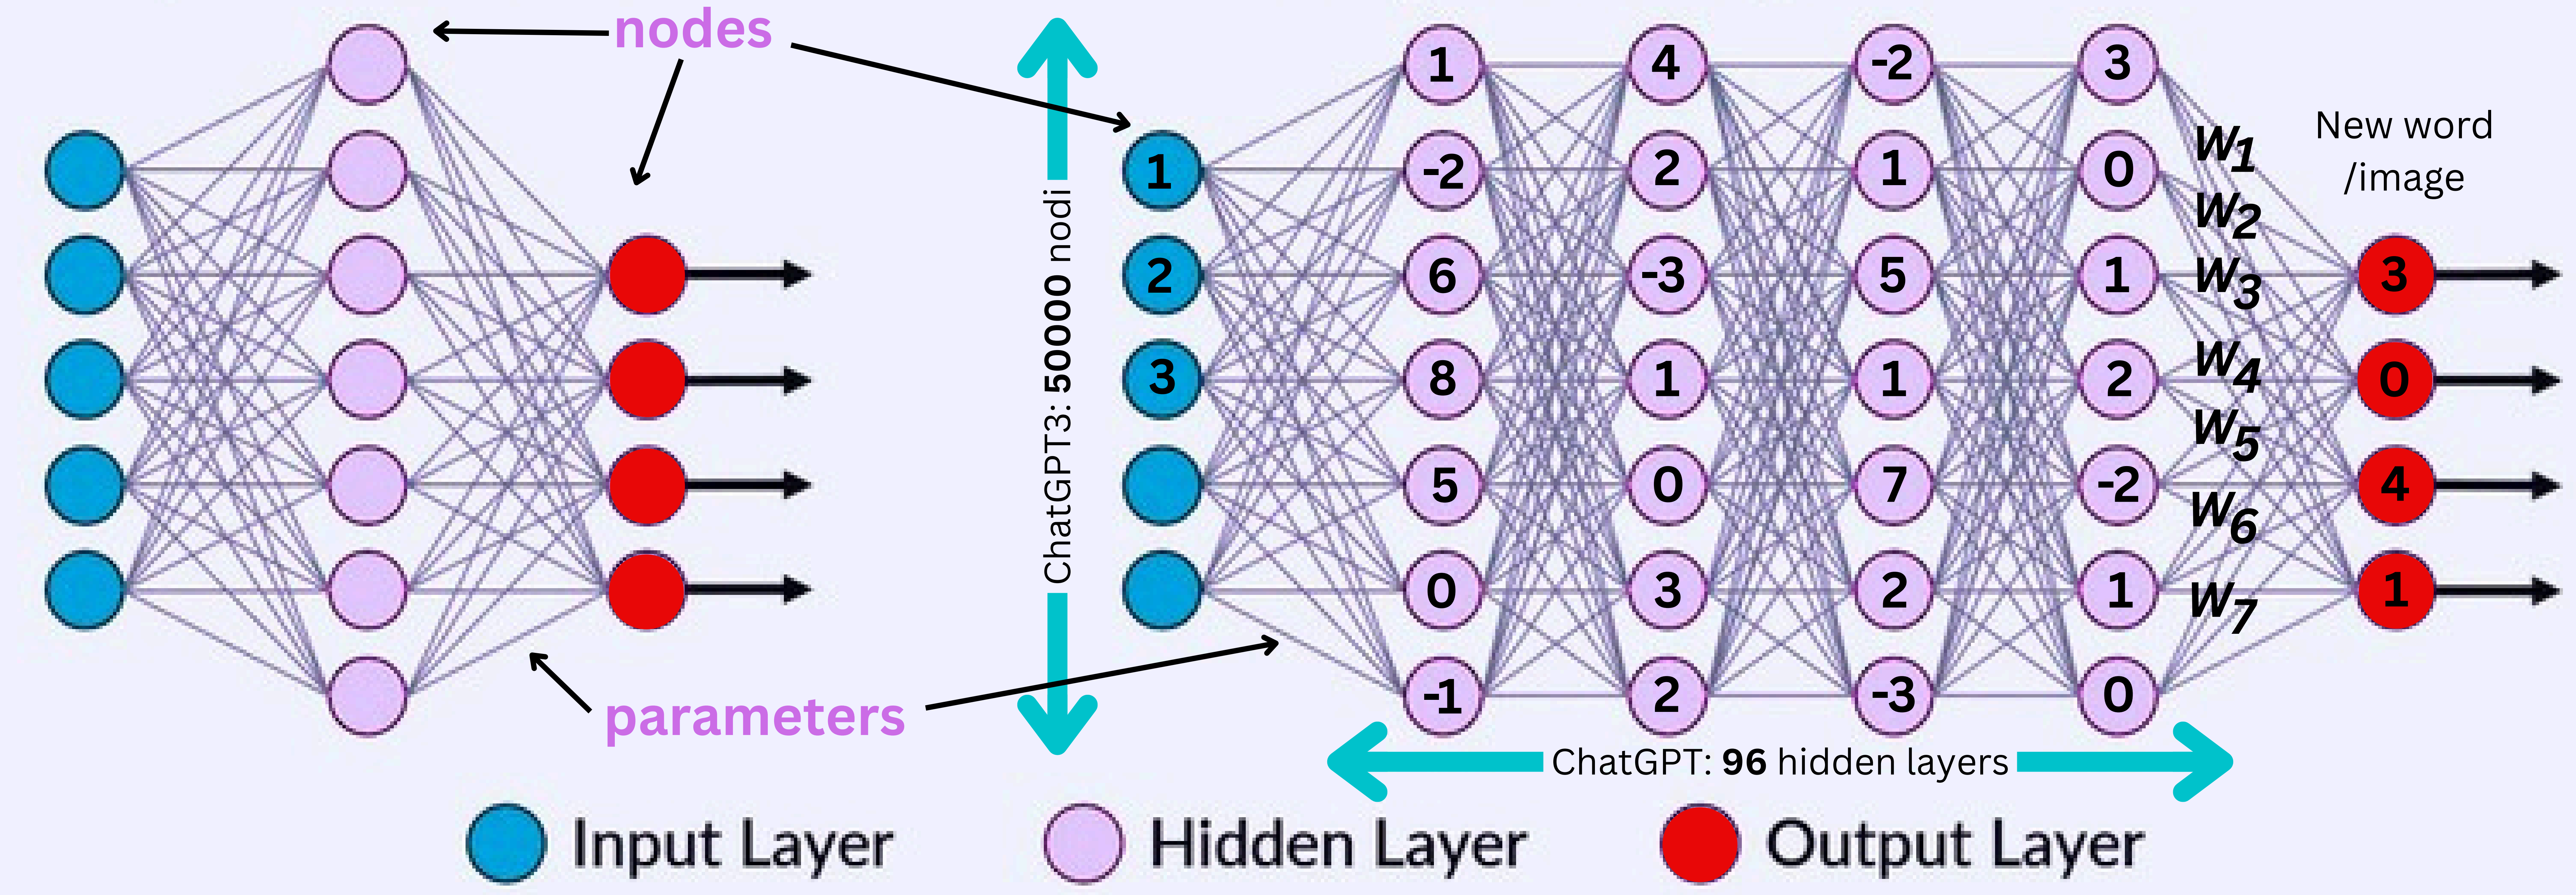

The result of the Forward Propagation is the new word or image to add at the end of the sequence

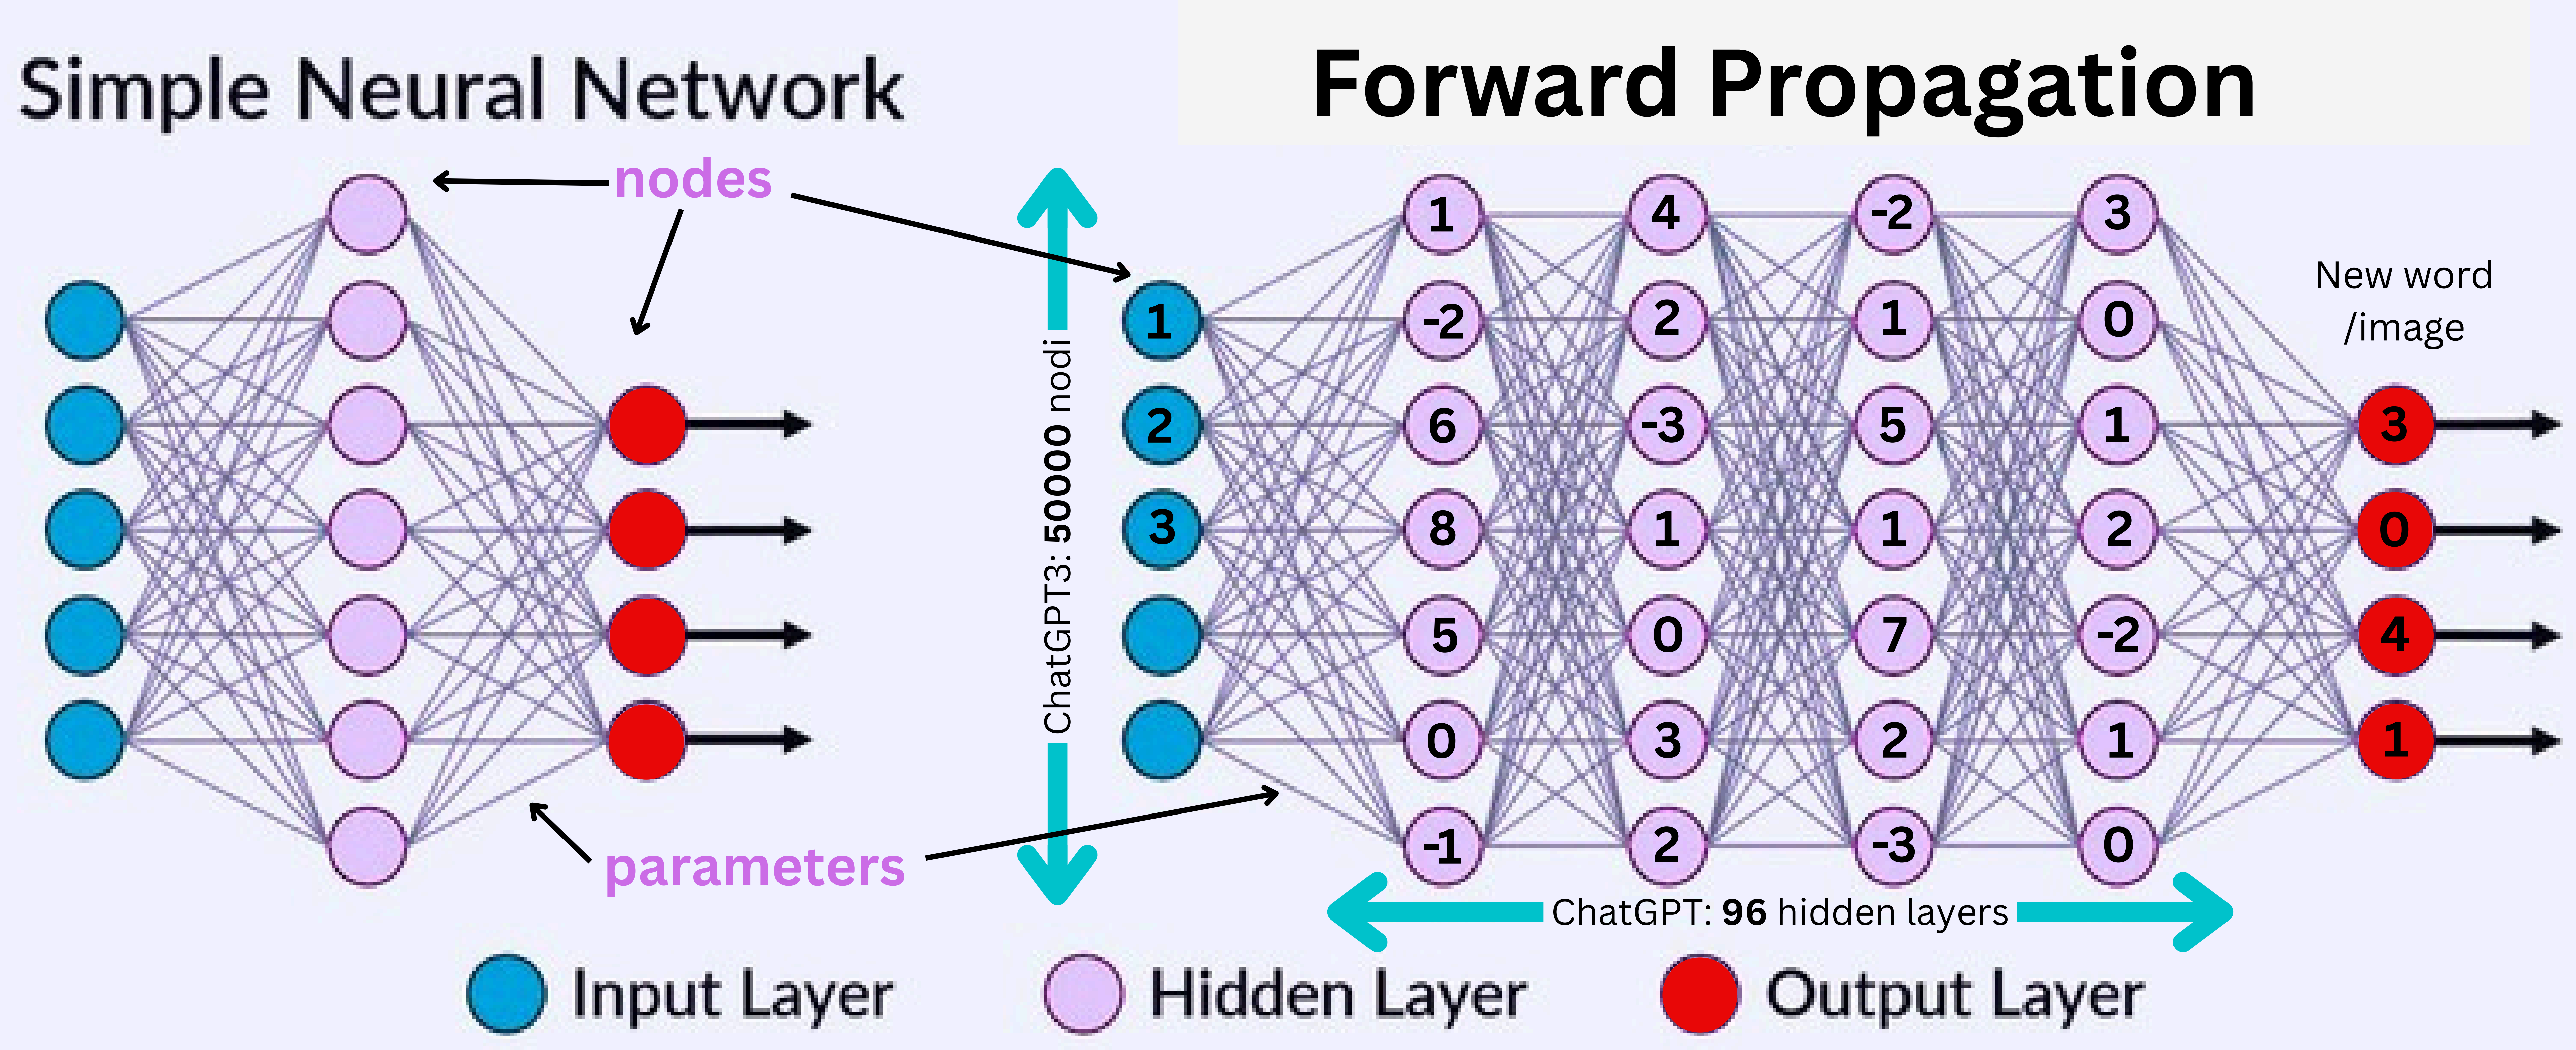

The values of the parameters W never change when you run a new prompt. Only the values of the currents inside the nodes change. How could the researchers identify the right values of the weights W?

Backpropagation

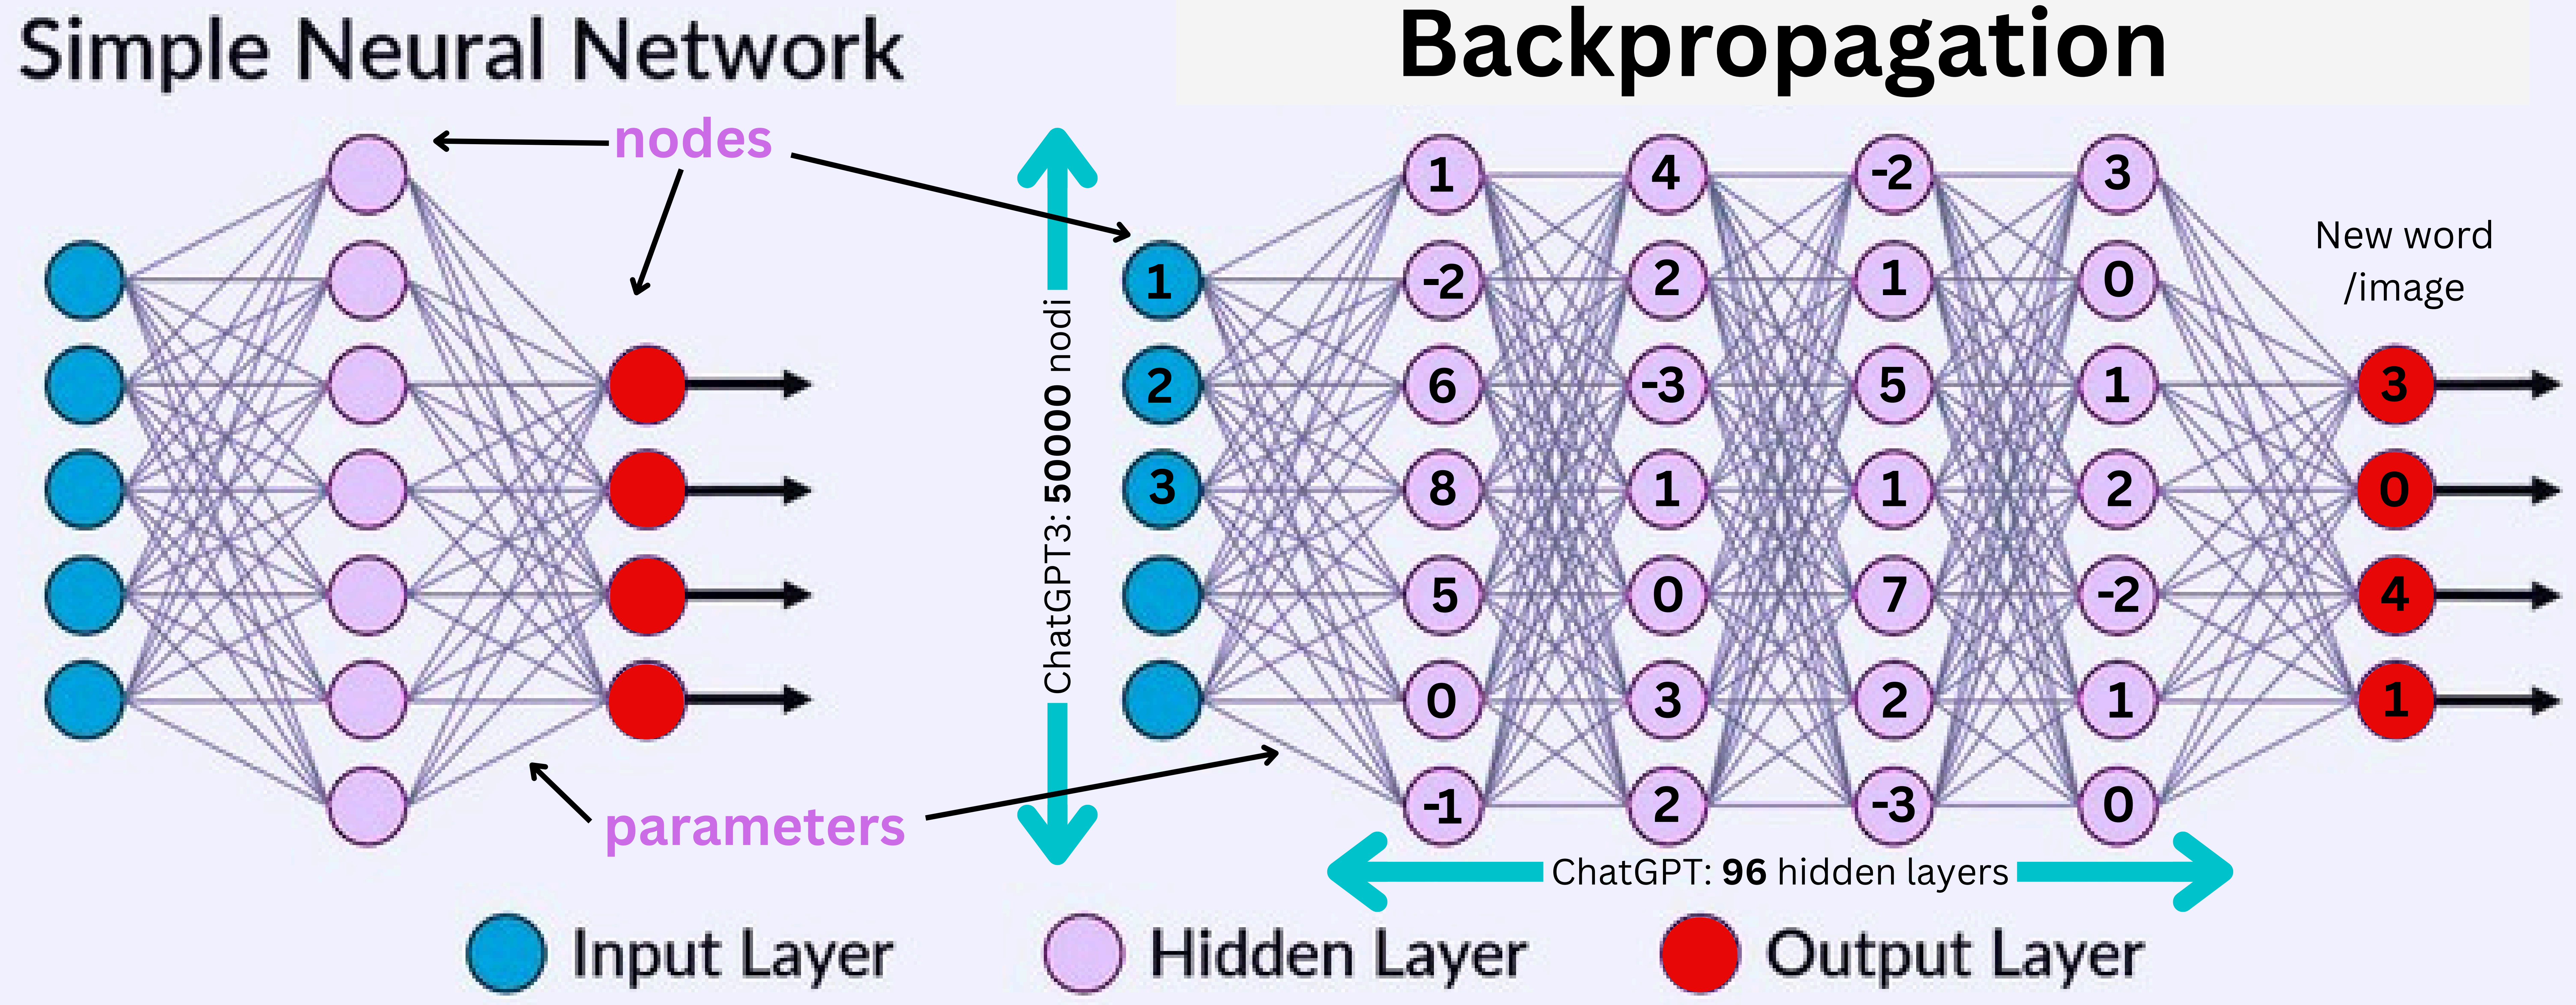

Through a method called "Backpropagation". They started giving random values to all parameters W.

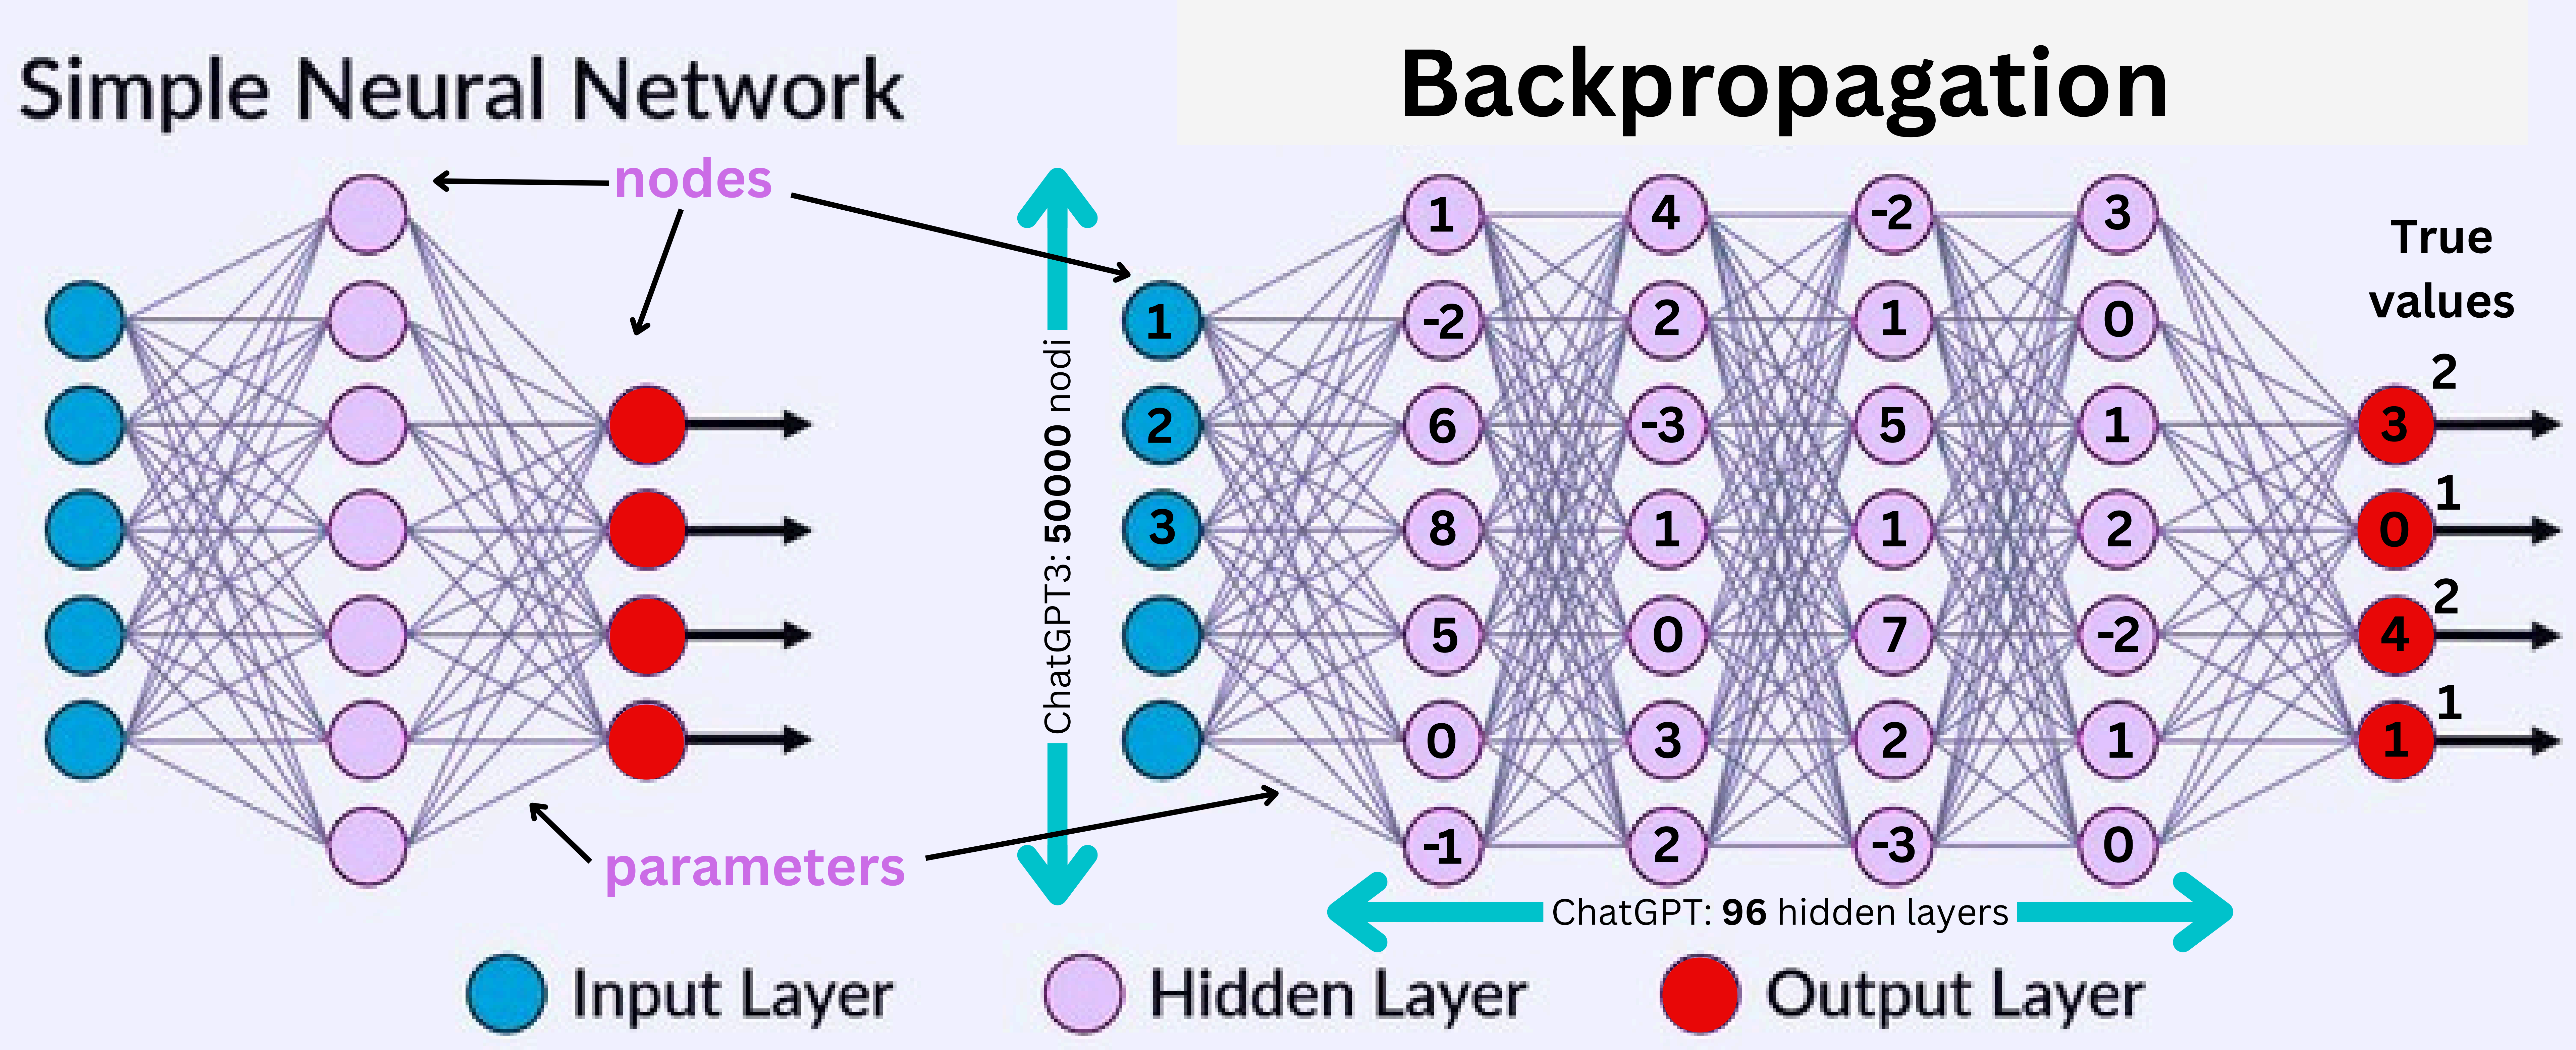

Then, you measure the difference (Loss function) between the output values and the true ones of the correct word/image

Backpropagation

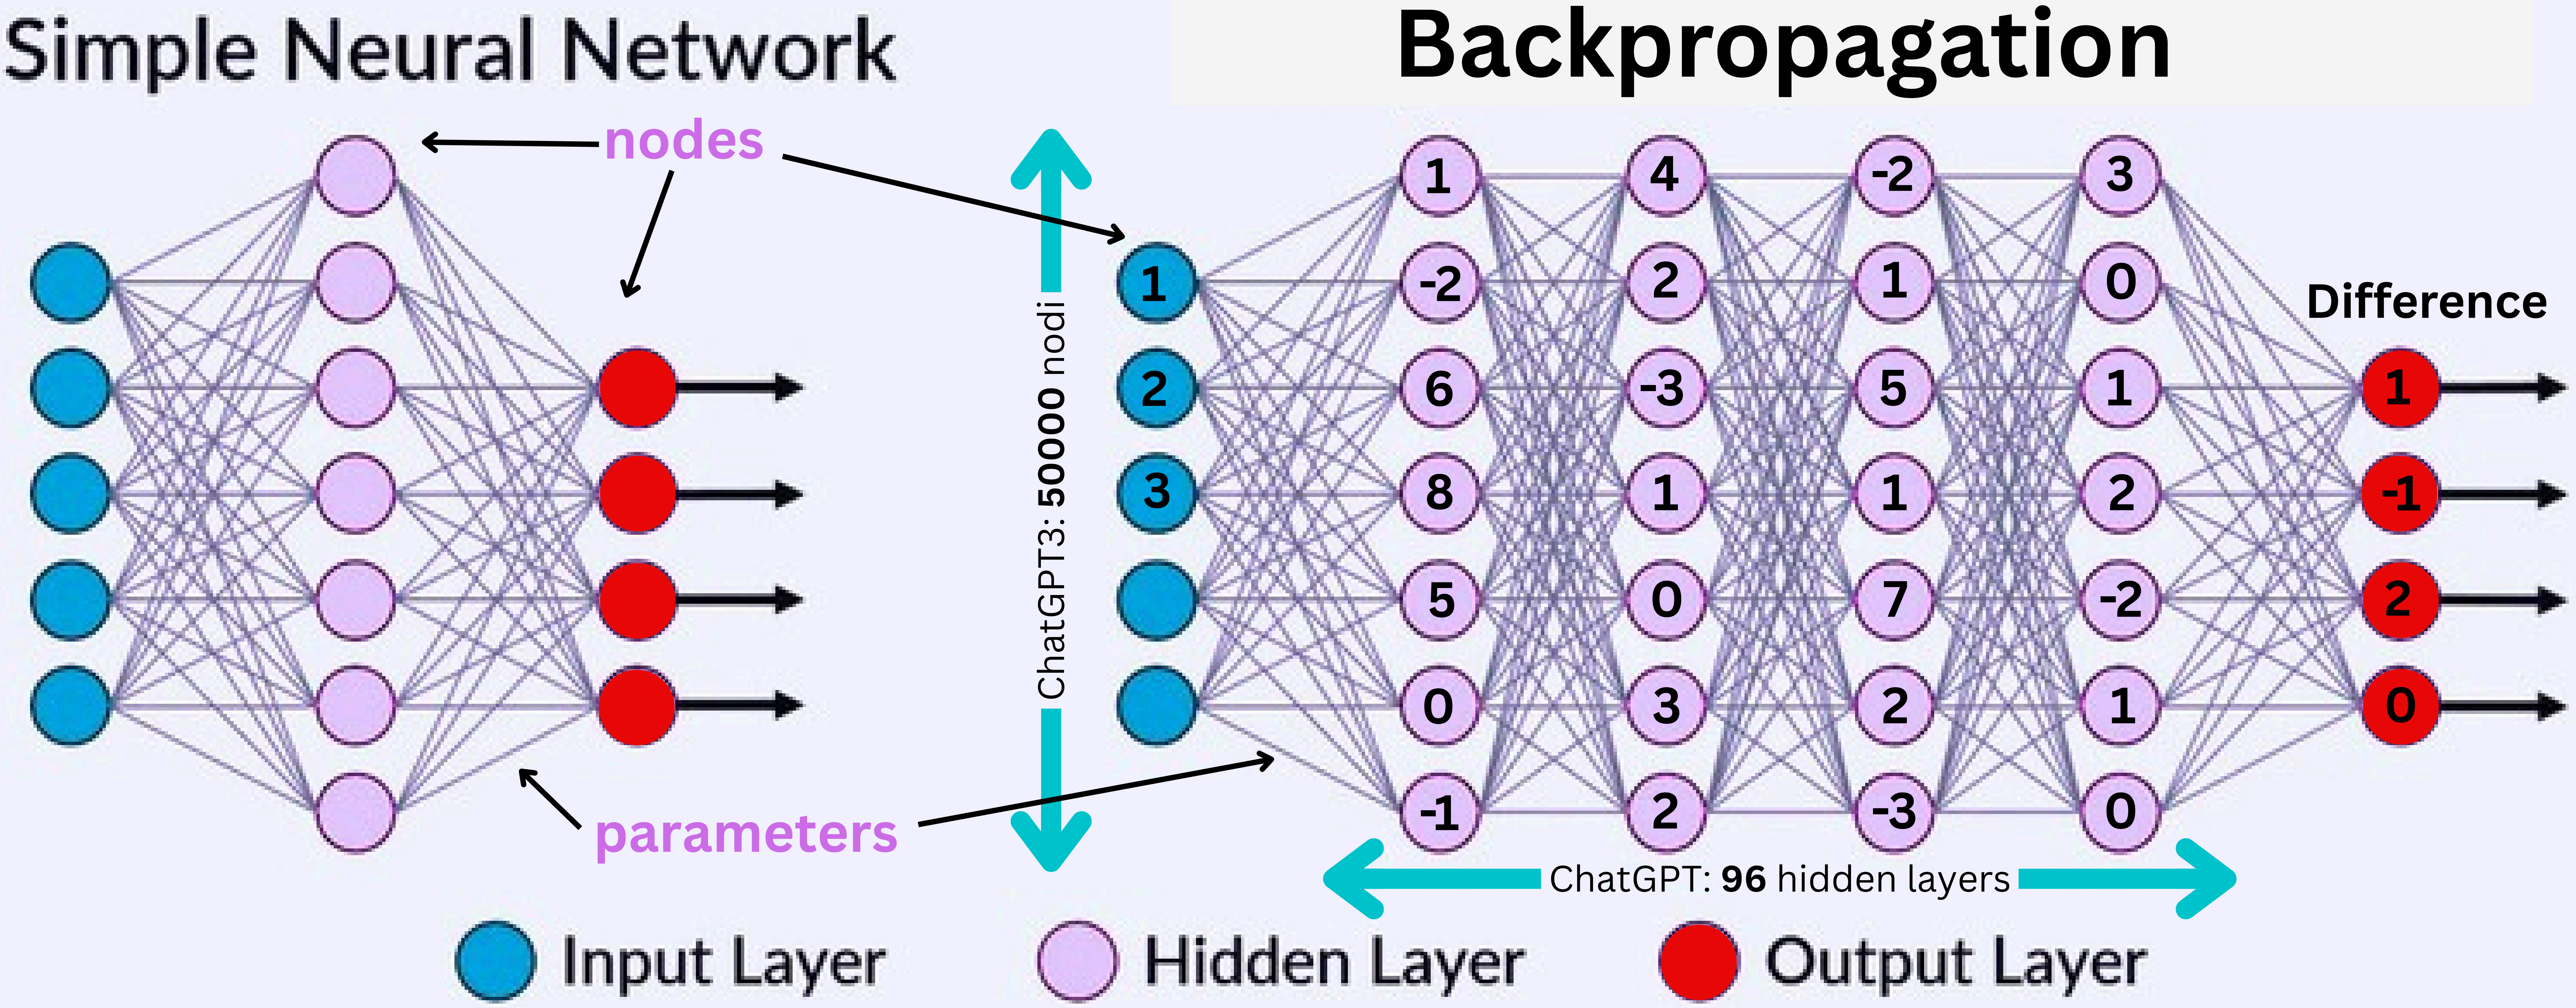

Then, you measure the difference (Loss function) between the output values and the true ones of the correct word/image

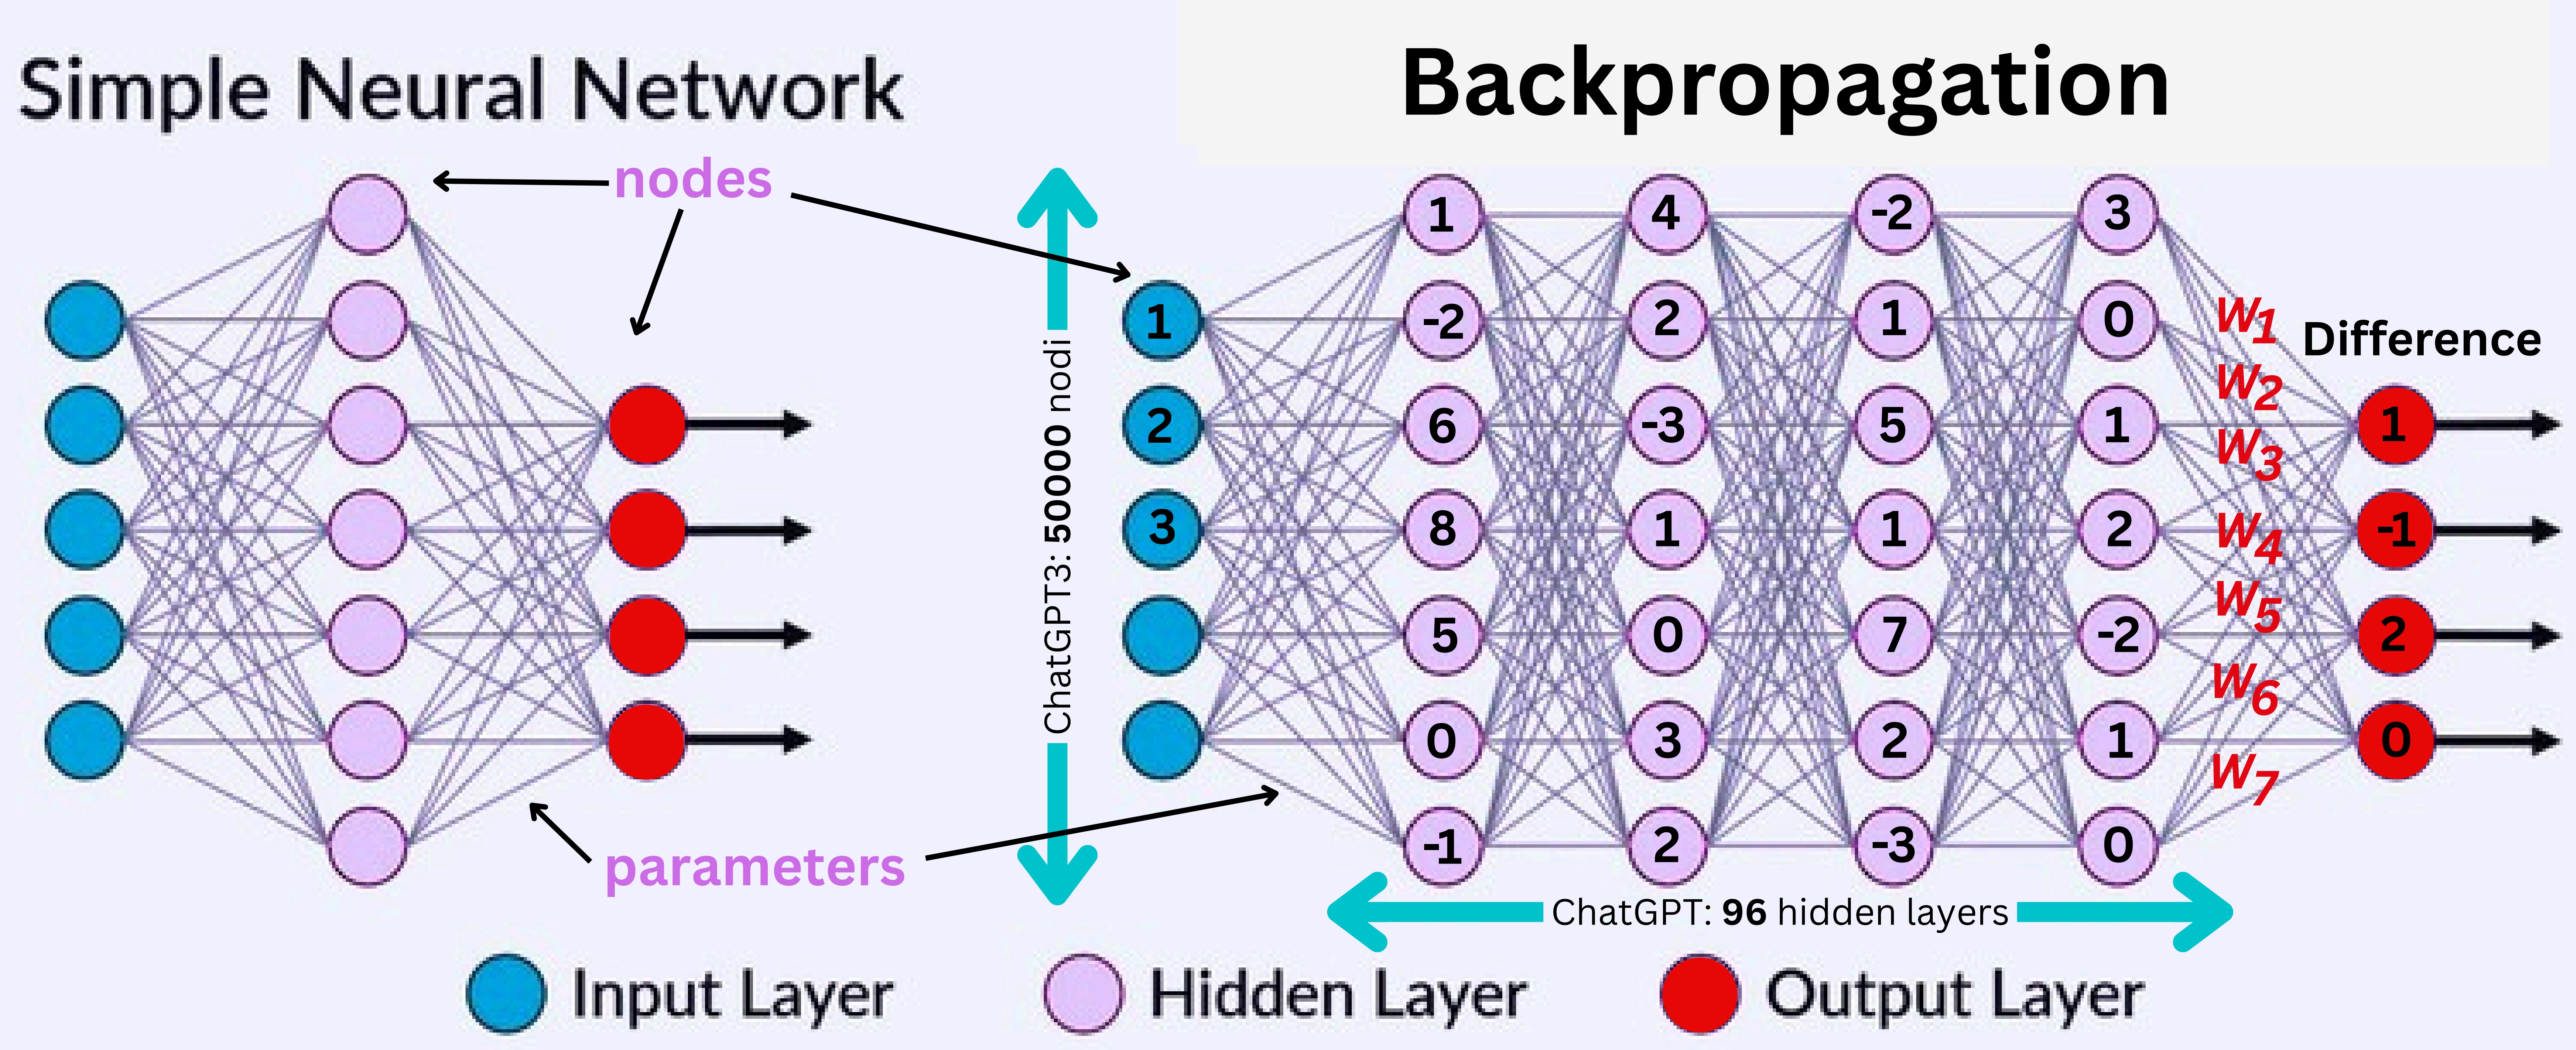

The difference is then propagated from right to left, hence the name backpropagation. This time also the parameters W are calibrated

Backpropagation

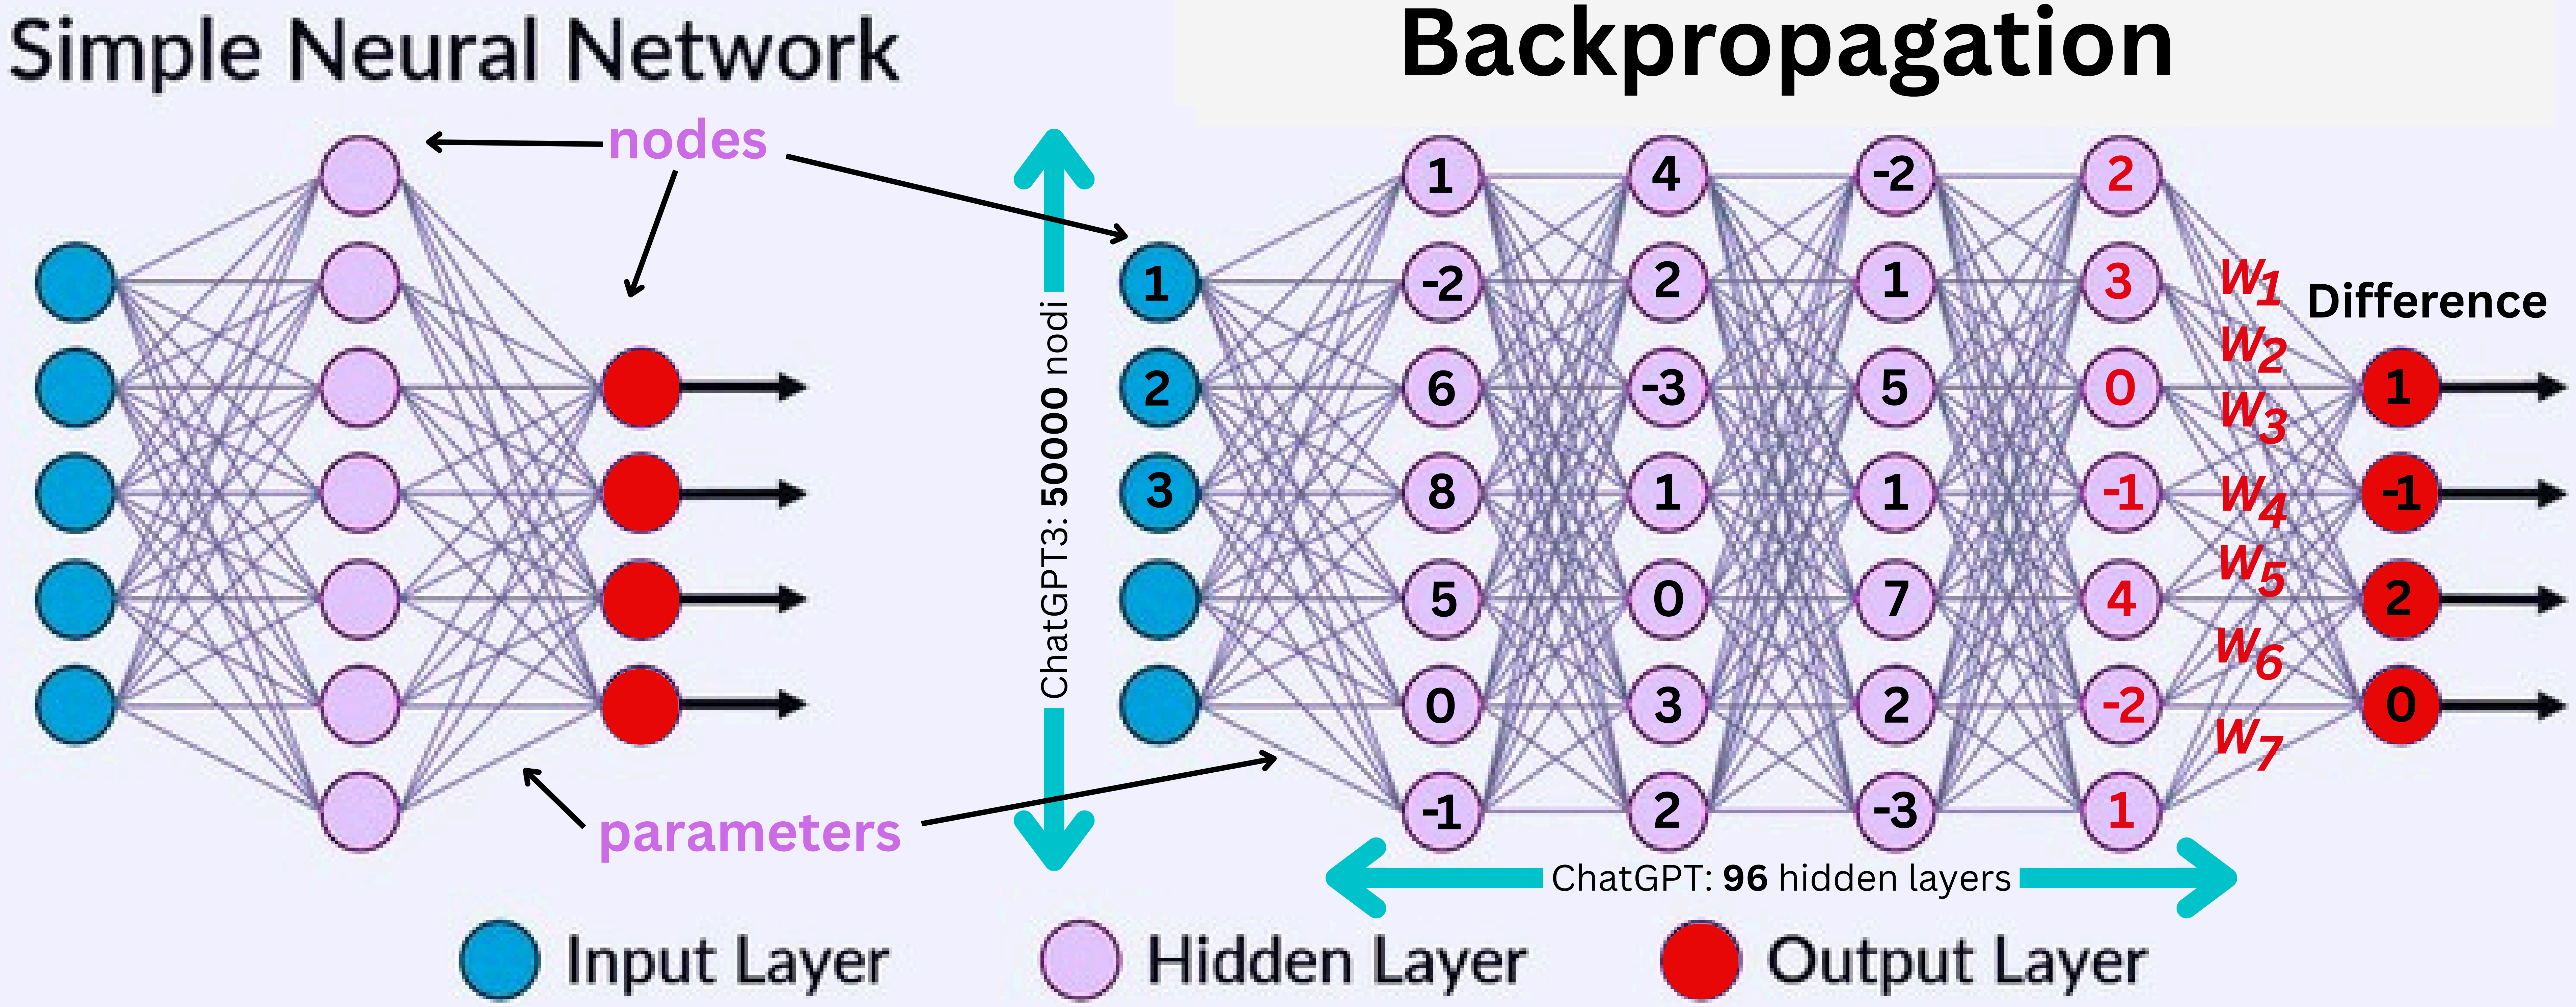

The idea is that the higher the difference, the more the parameters should be corrected. For example, a difference of 2 could be reduced by decreasing a bit a few parameters W, so the next time you run the forward regression the difference won't be 2 but only 1.

Simple Neural Network

Backpropagation

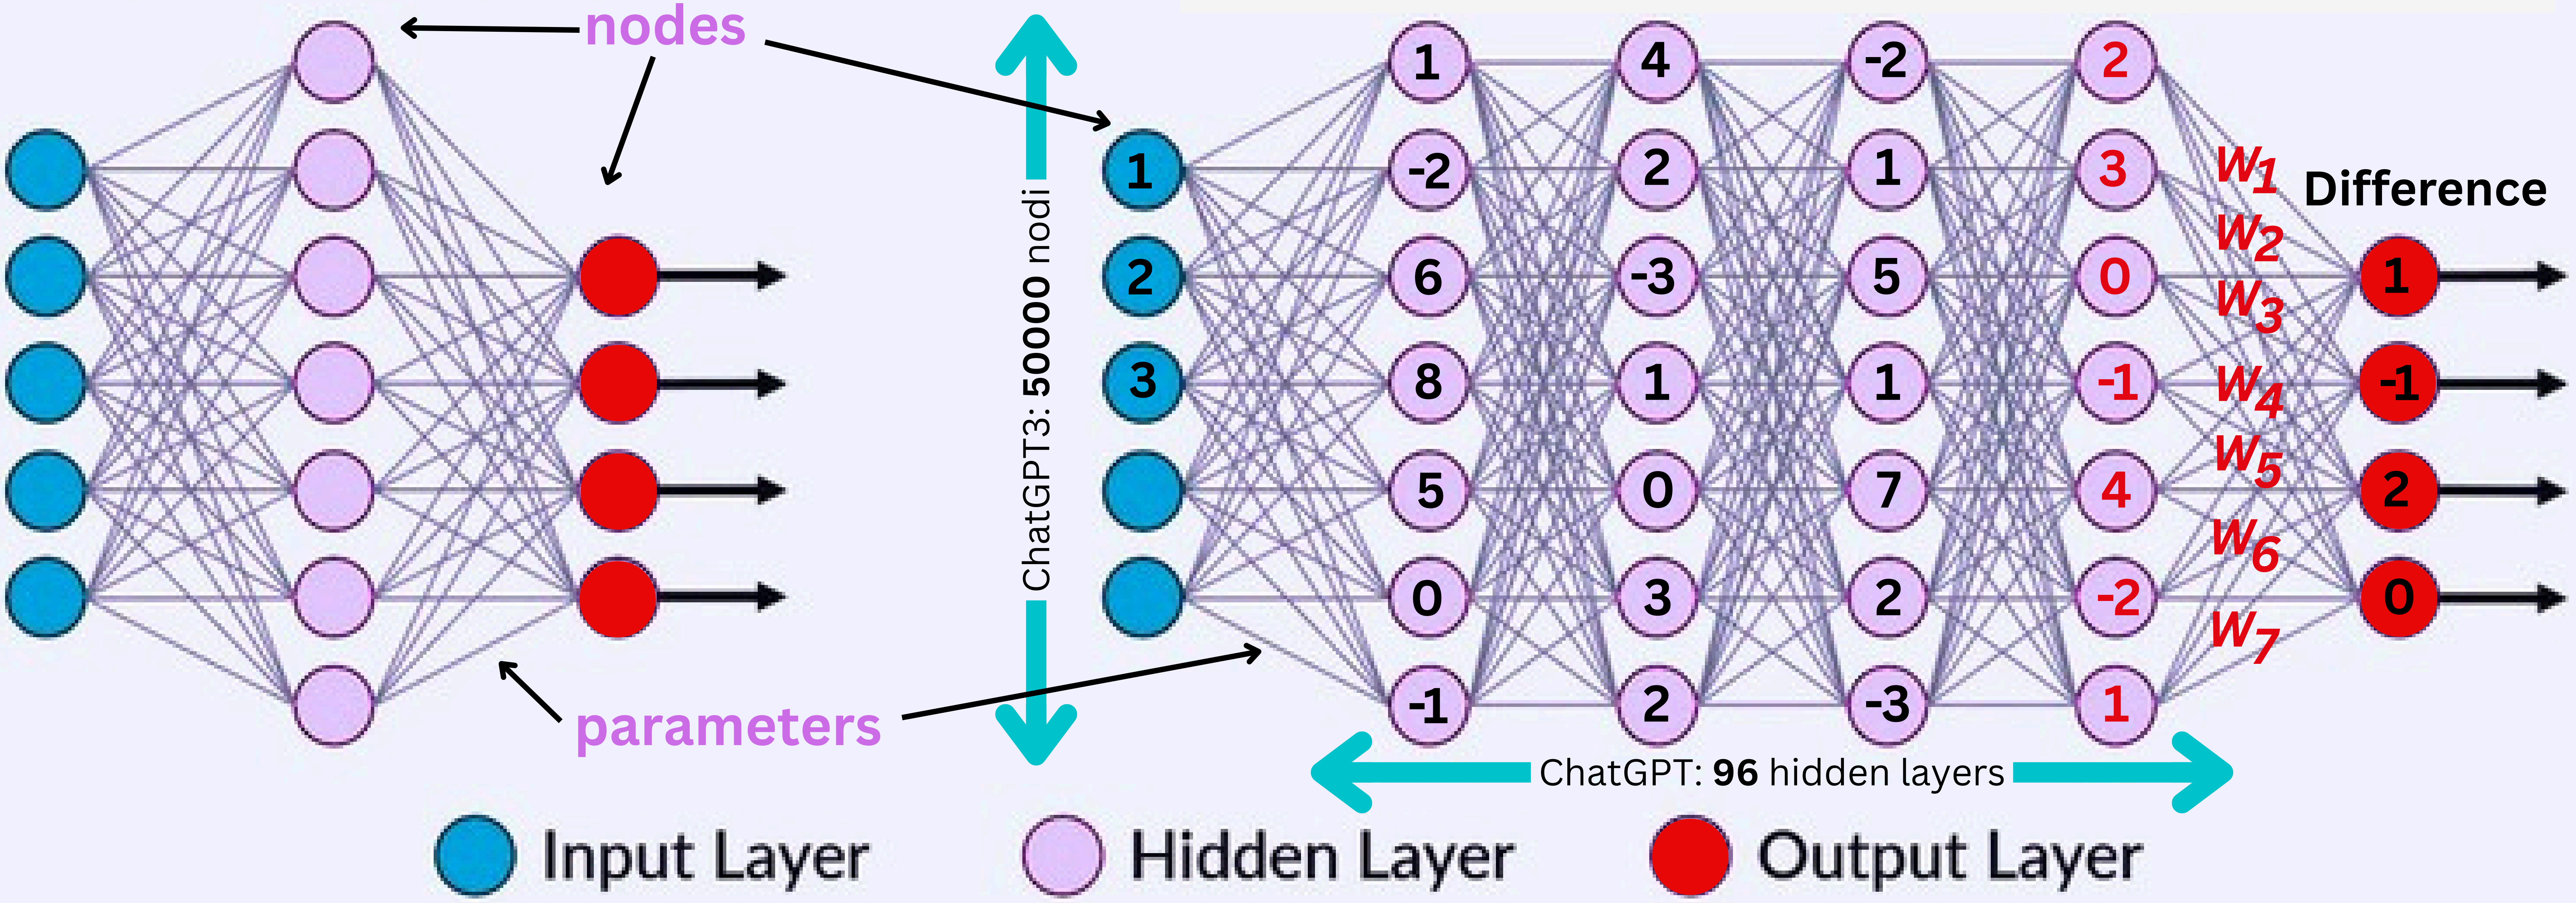

The algorithm increases the weights W of the group of neurons that contribute to providing the correct answer and decreases the weights of the neurons that tend to move away from the correct answer. In this way the correct answer is amplified.

Simple Neural Network

Backpropagation

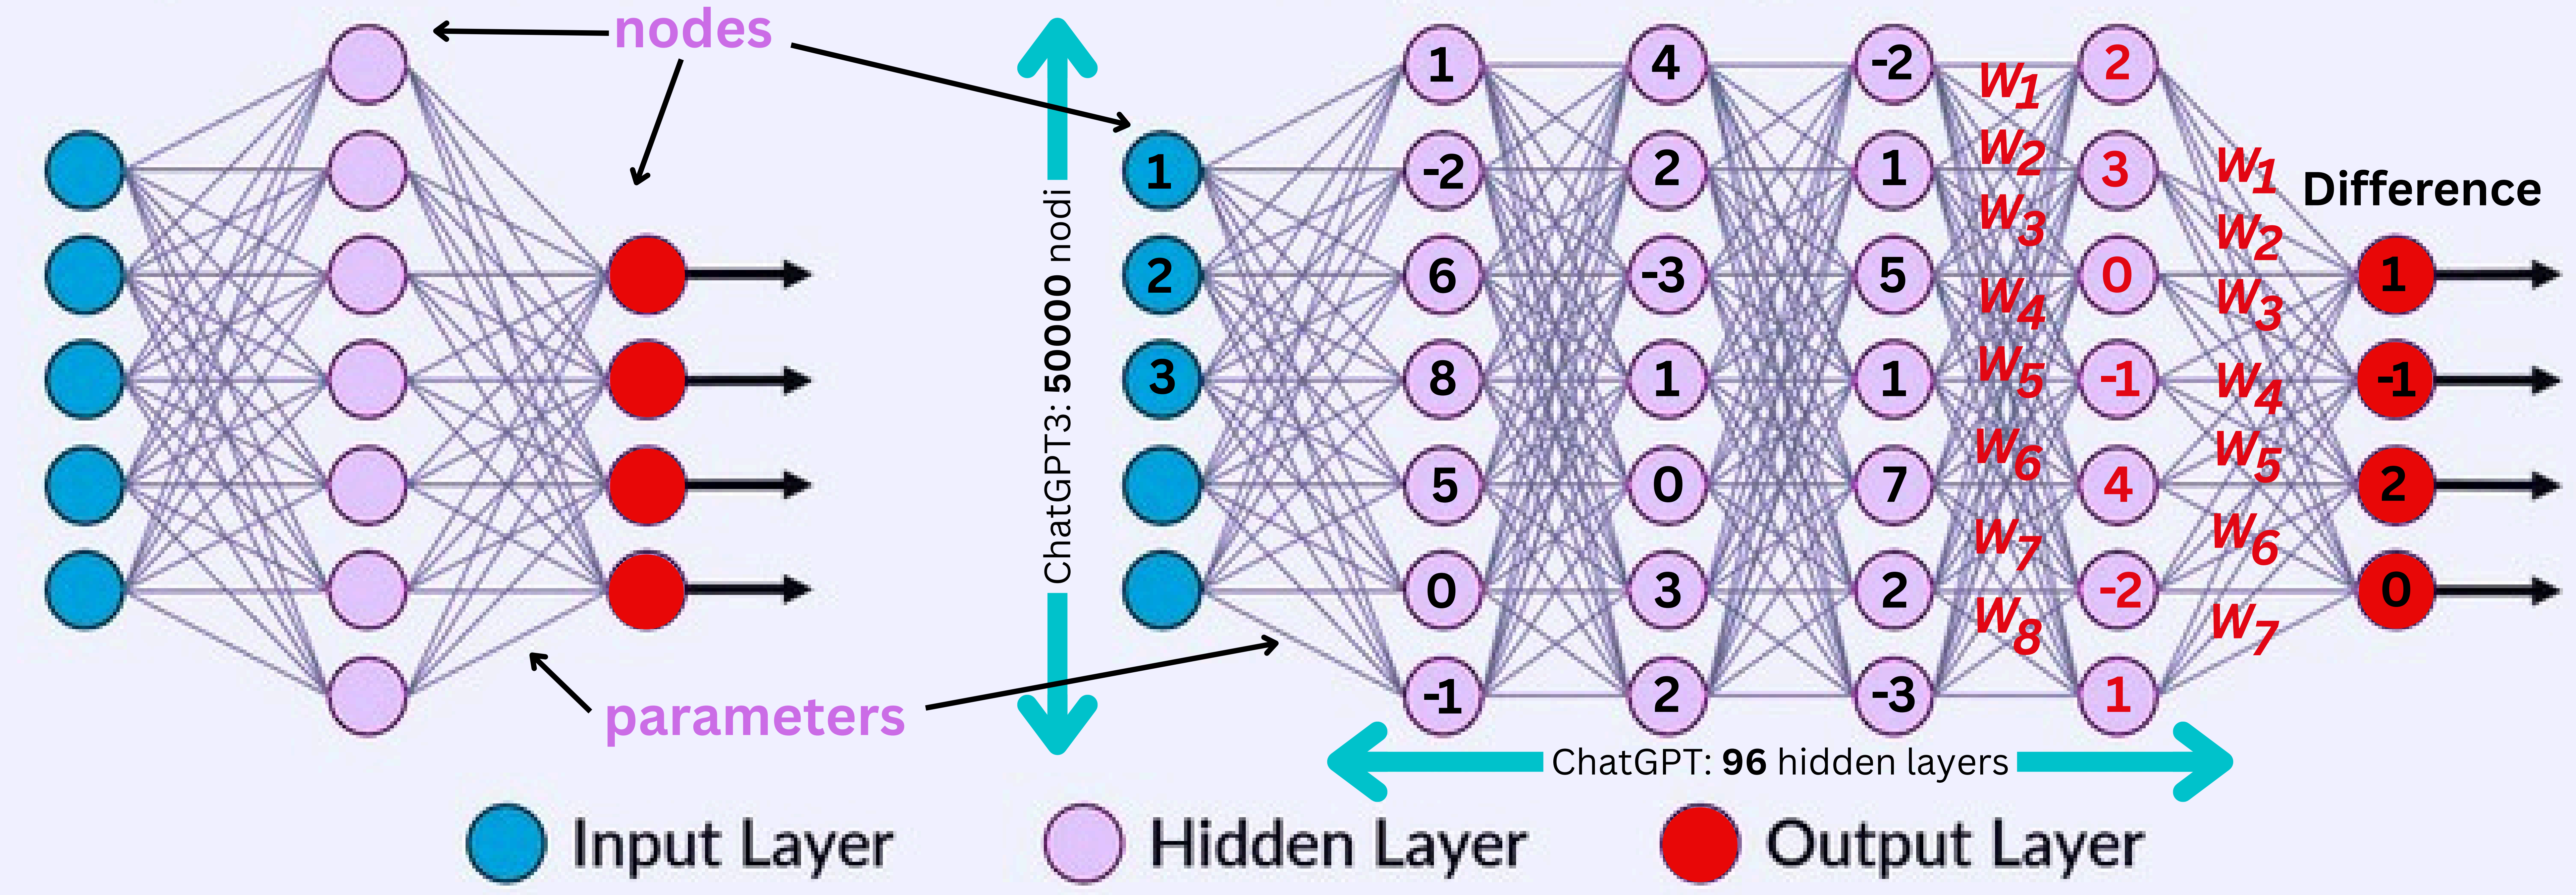

The same method is applied to the parameters of the previous hidden layers, one by one, until the first layer is reached

Backpropagation

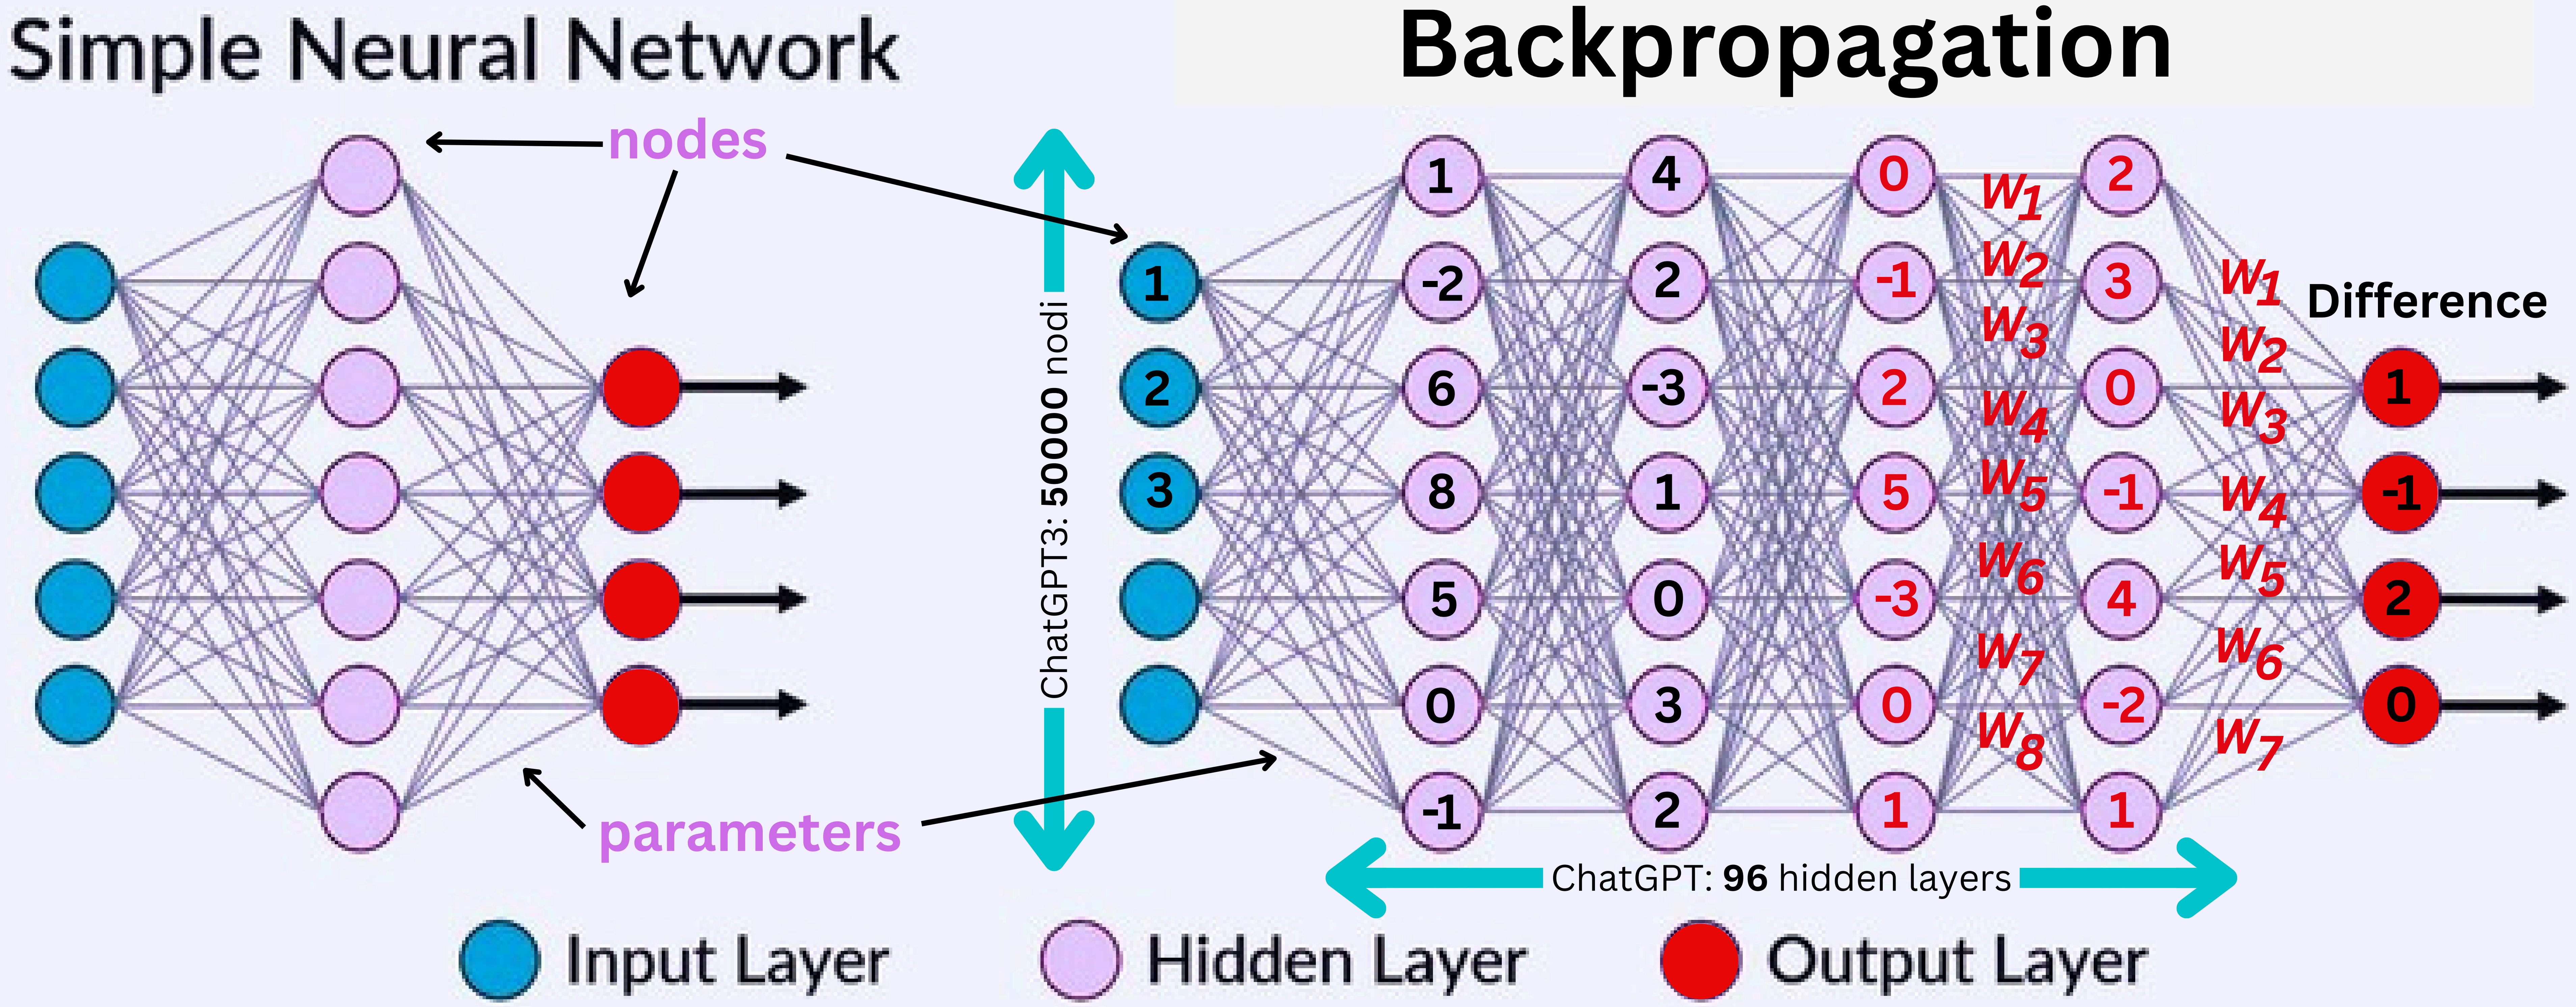

The same method is applied to the parameters of the previous hidden layers, one by one, until the first layer is reached

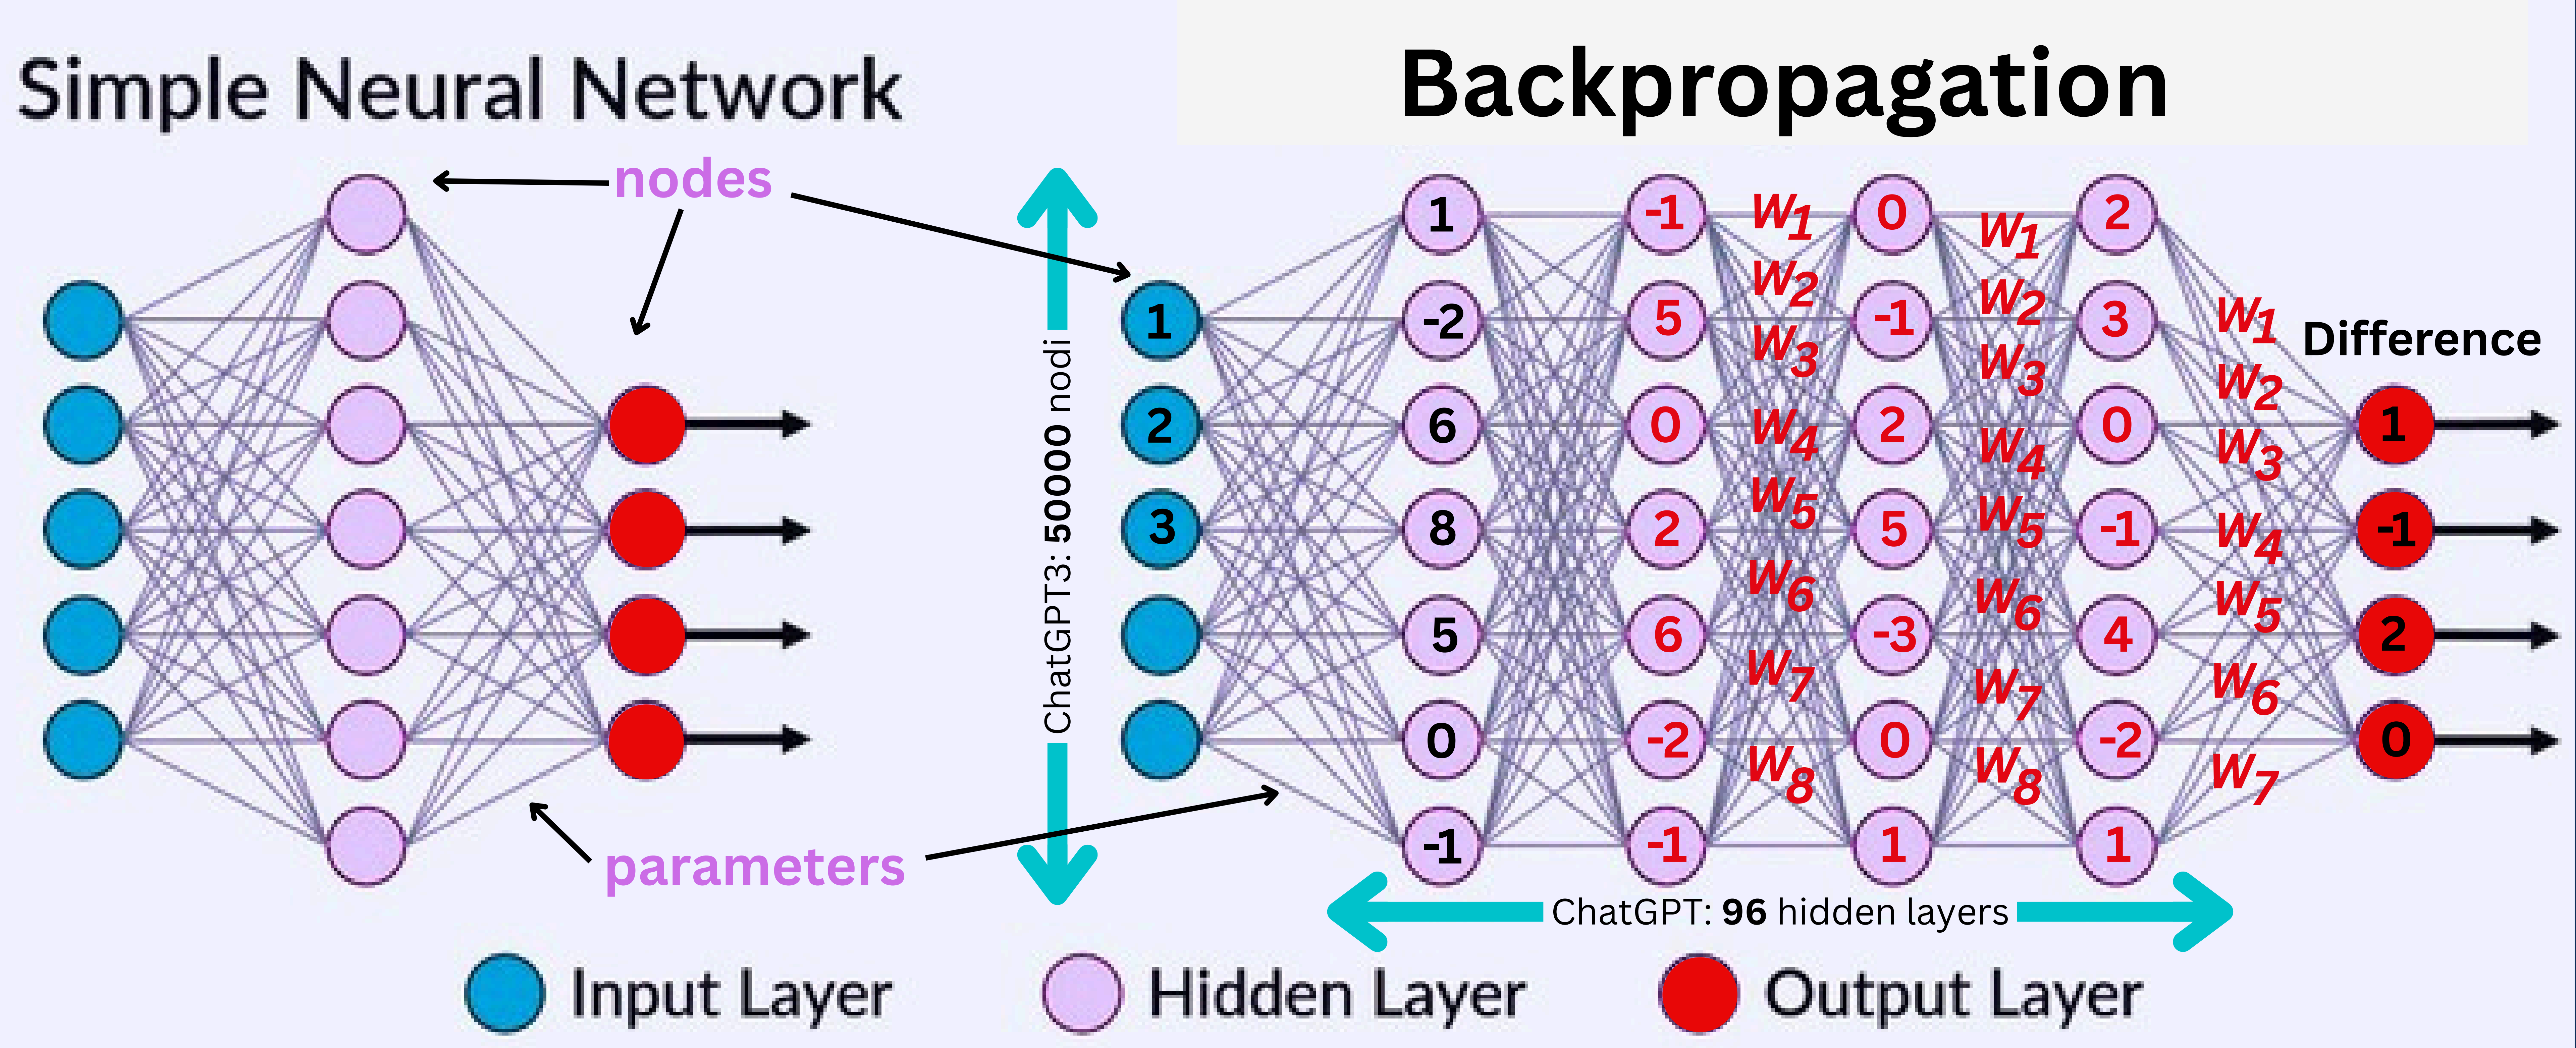

The same method is applied to the parameters of the previous hidden layers, one by one, until the first layer is reached

Backpropagation

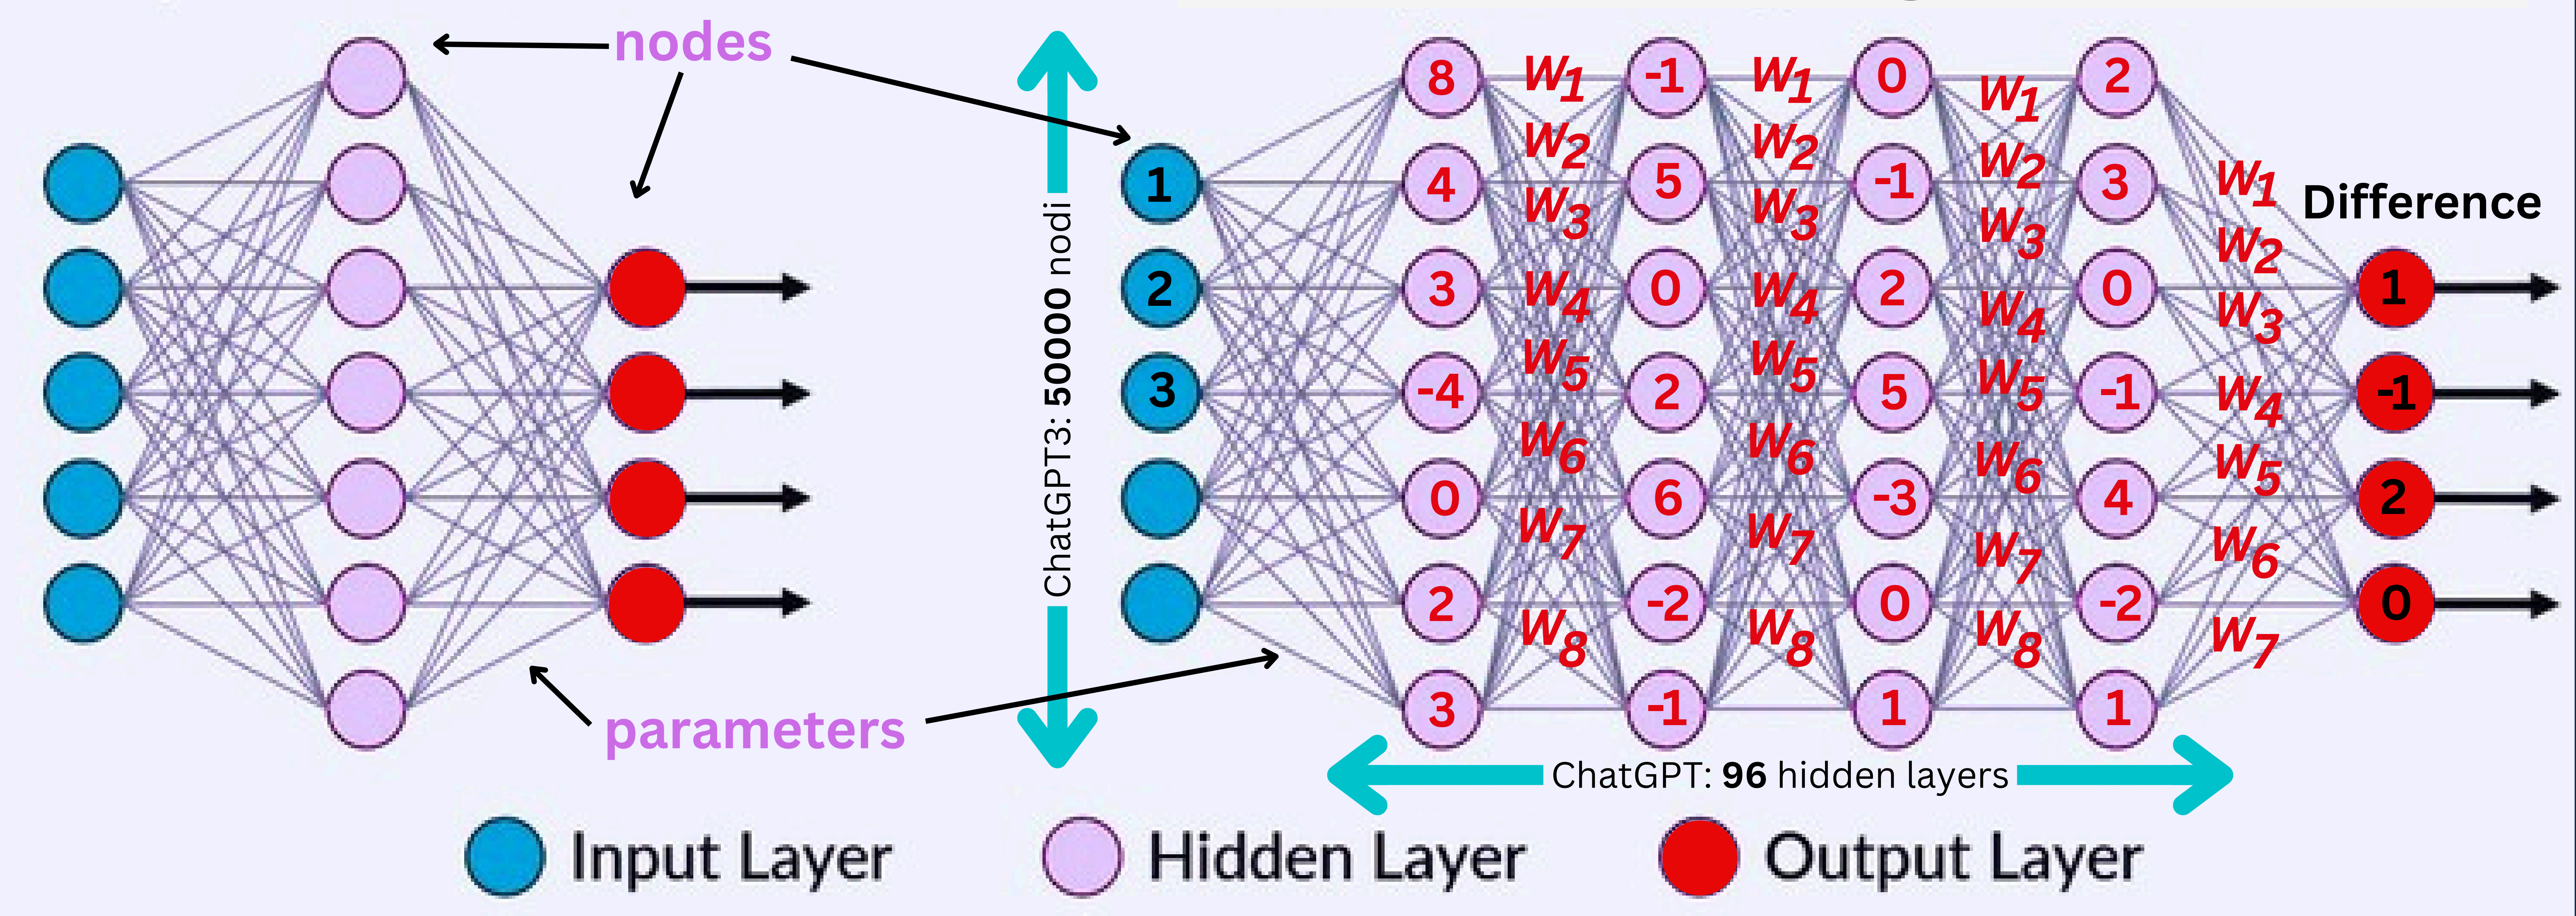

The same method is applied to the parameters of the previous hidden layers, one by one, until the first layer is reached

Backpropagation

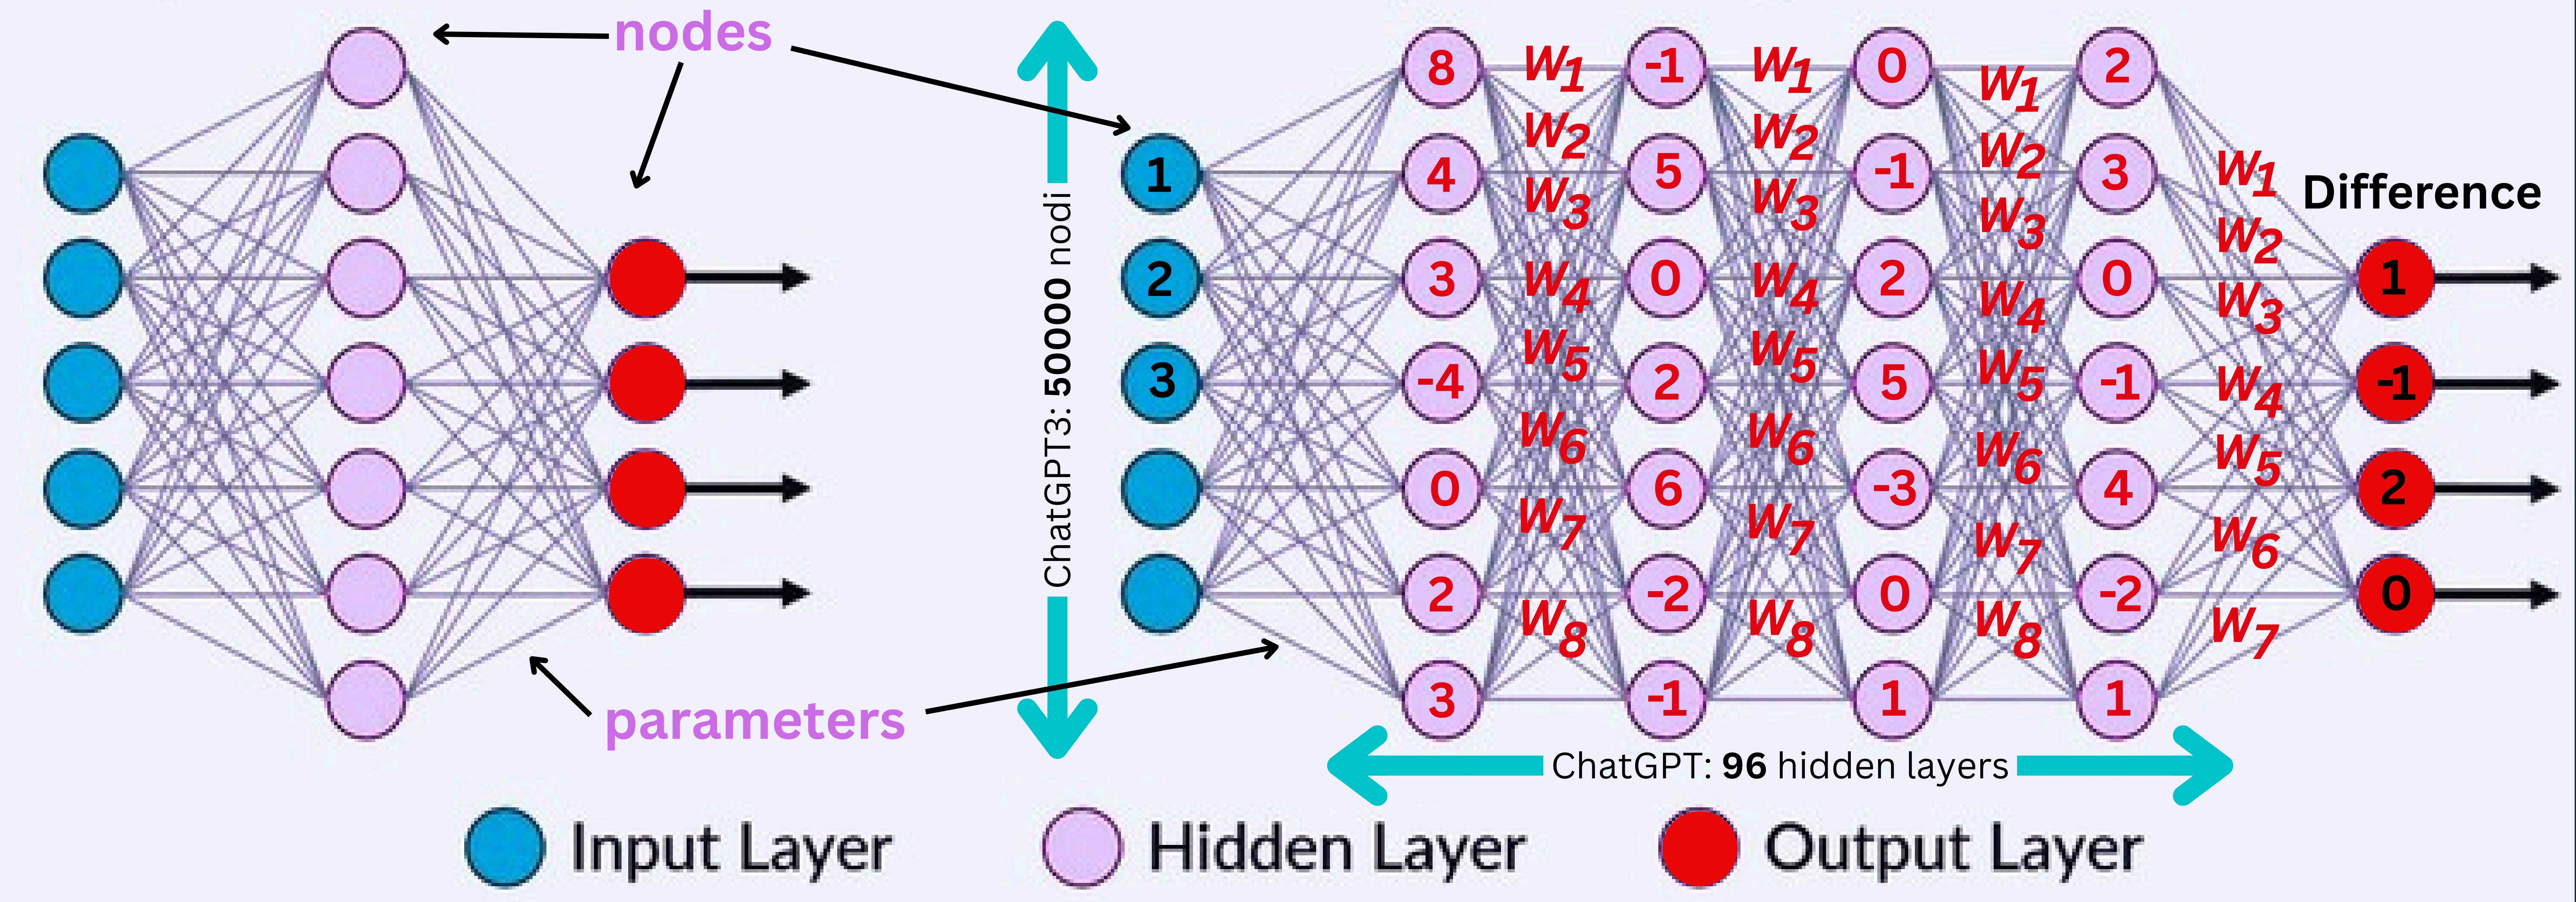

This process of optimizing the weights w is called model training or model calibration. When you train a deep learning model, you find the best values of all the parameters W.

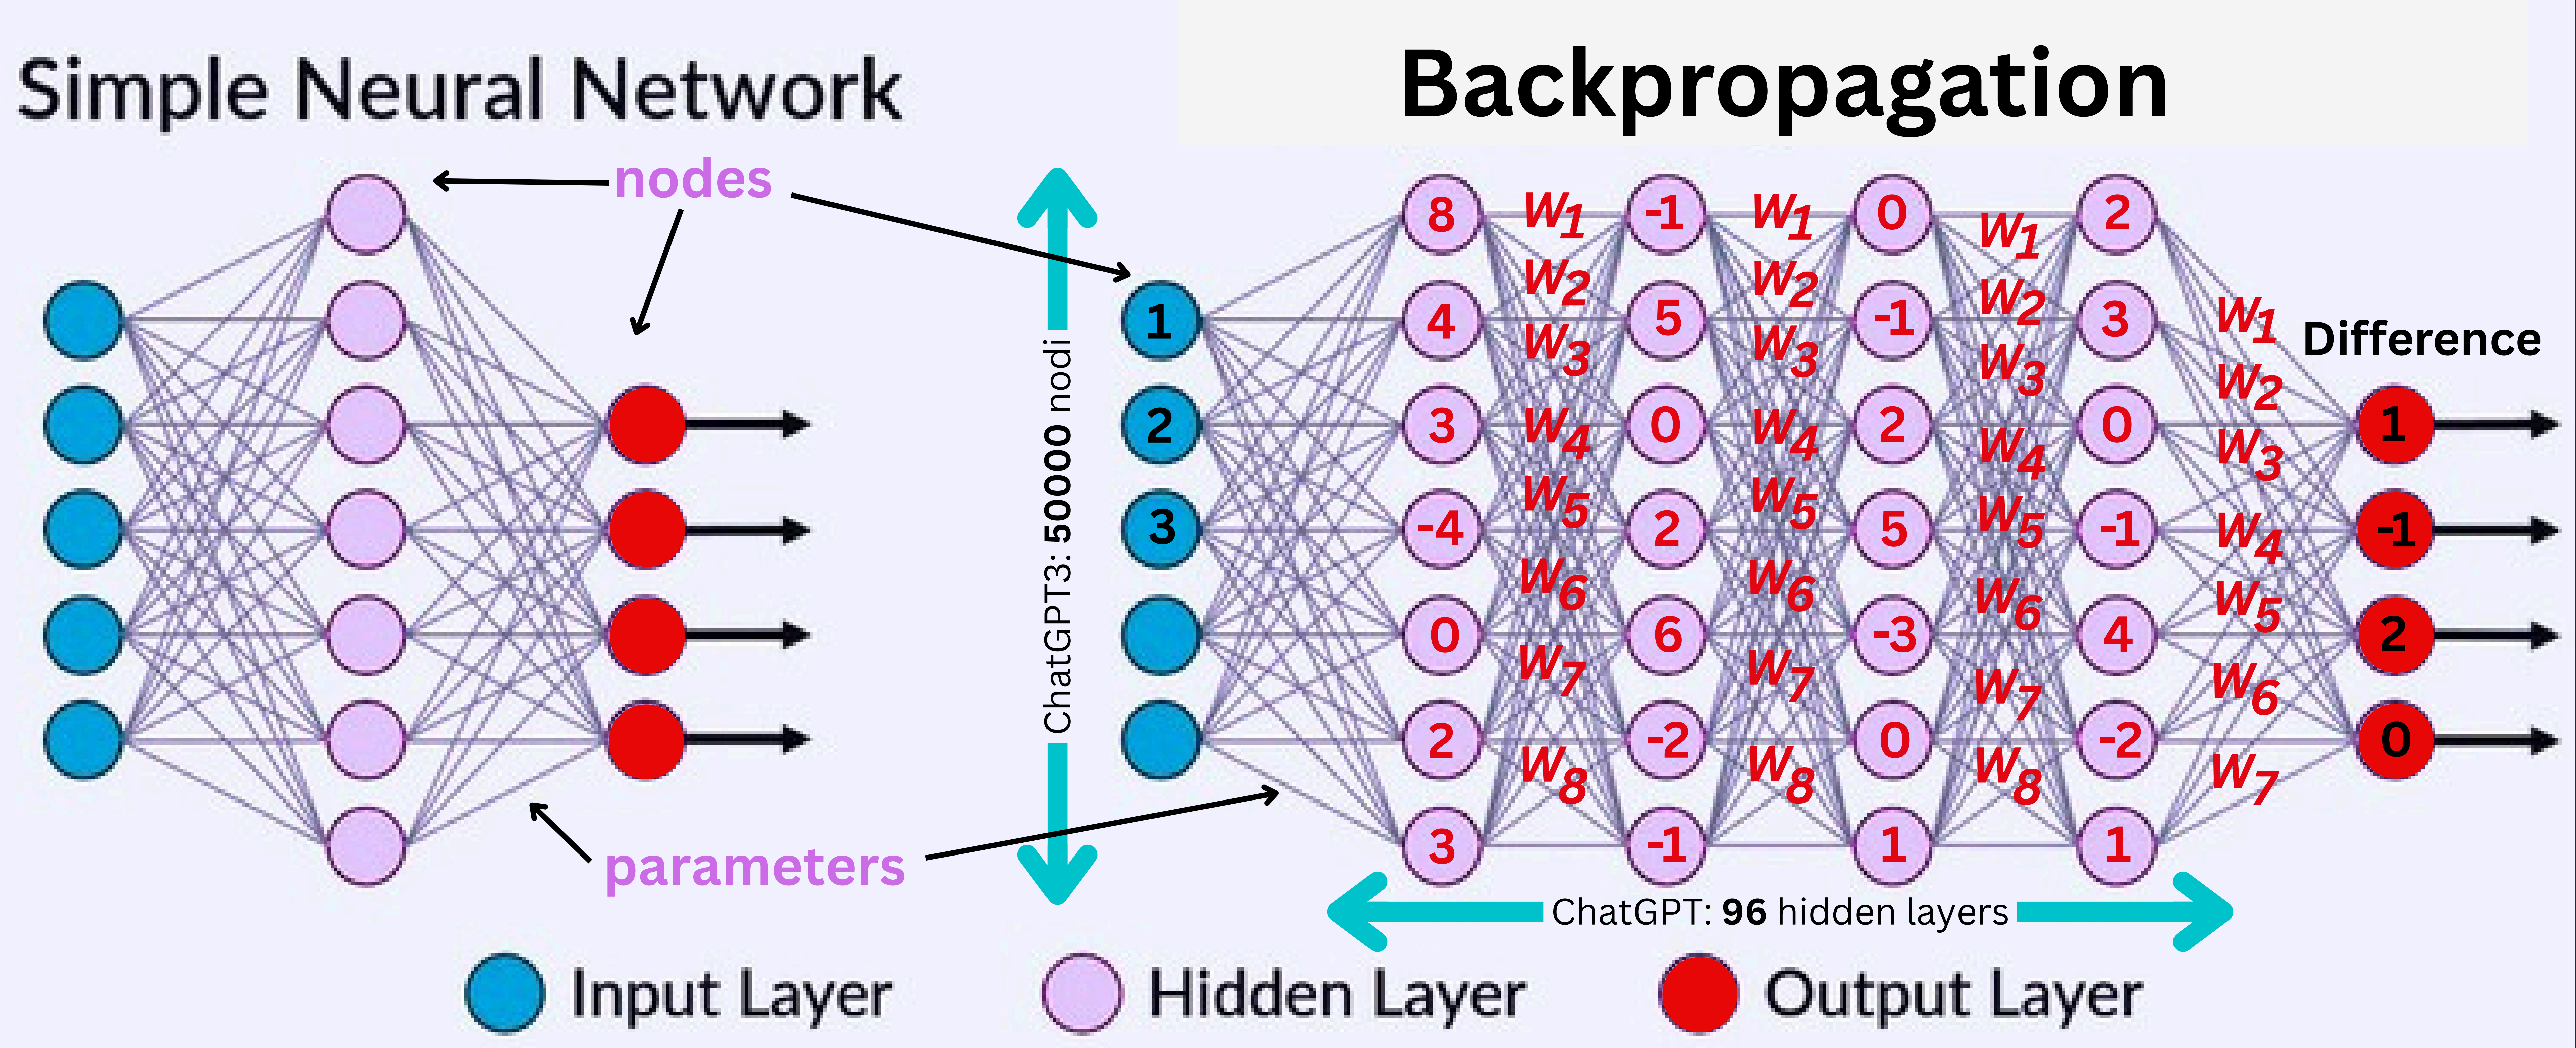

Once you have trained the model once, the work is done: it may take weeks to train a model, but once done, you don't need to train the model again every time a user insert a new prompt.

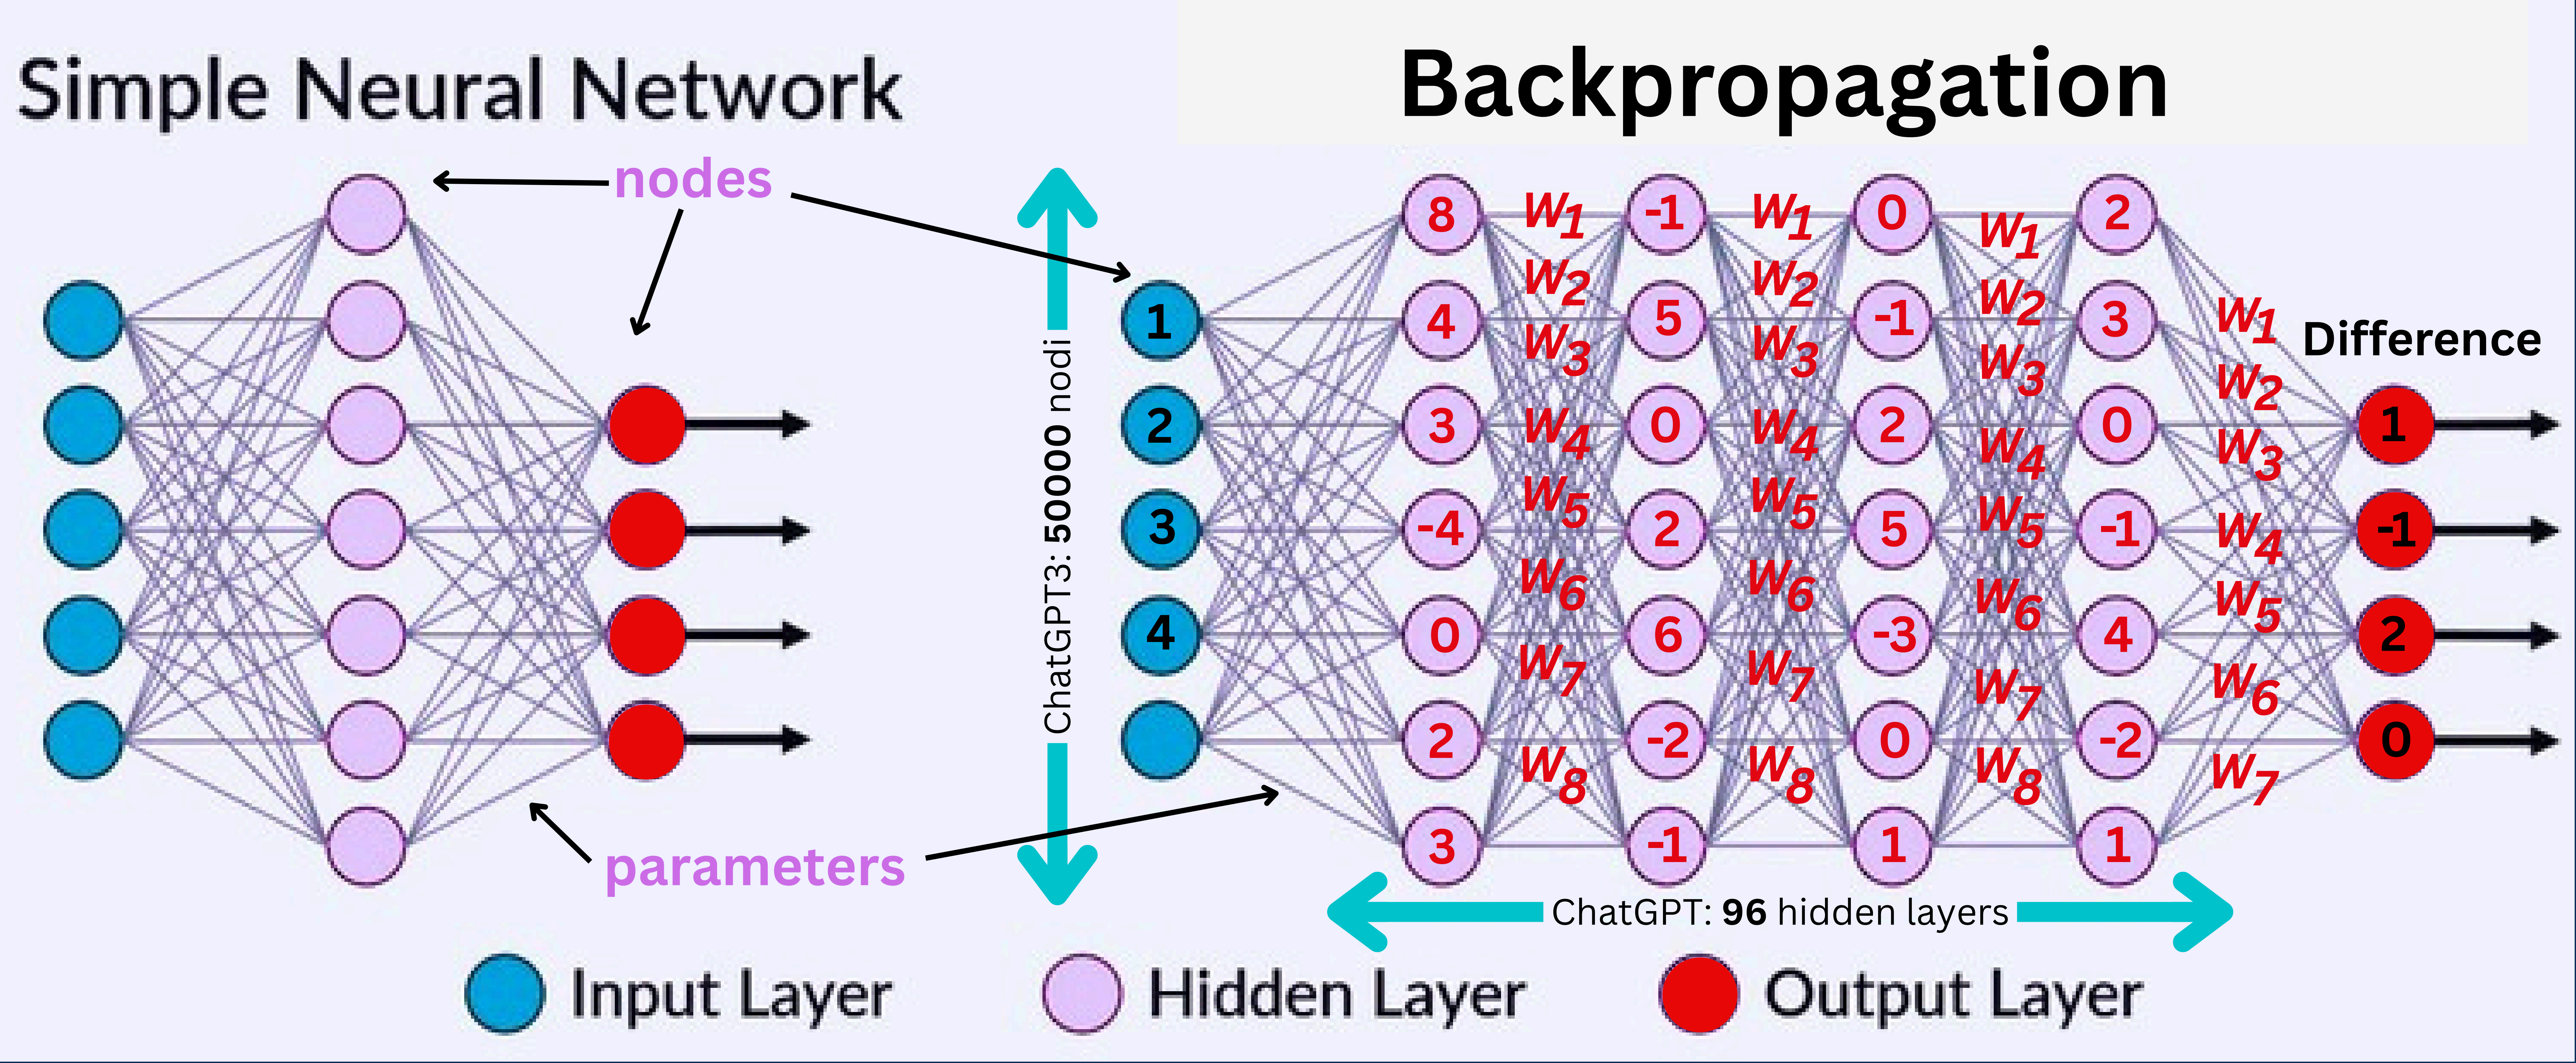

Once the model has forecasted the 4th word, you ask it to forecast the 5th word, and so on. It will fail over and over, until, after training it over thousands of books, it become very good at forecasting as it has learnt the structure of the language

Simple Neural Network

Backpropagation

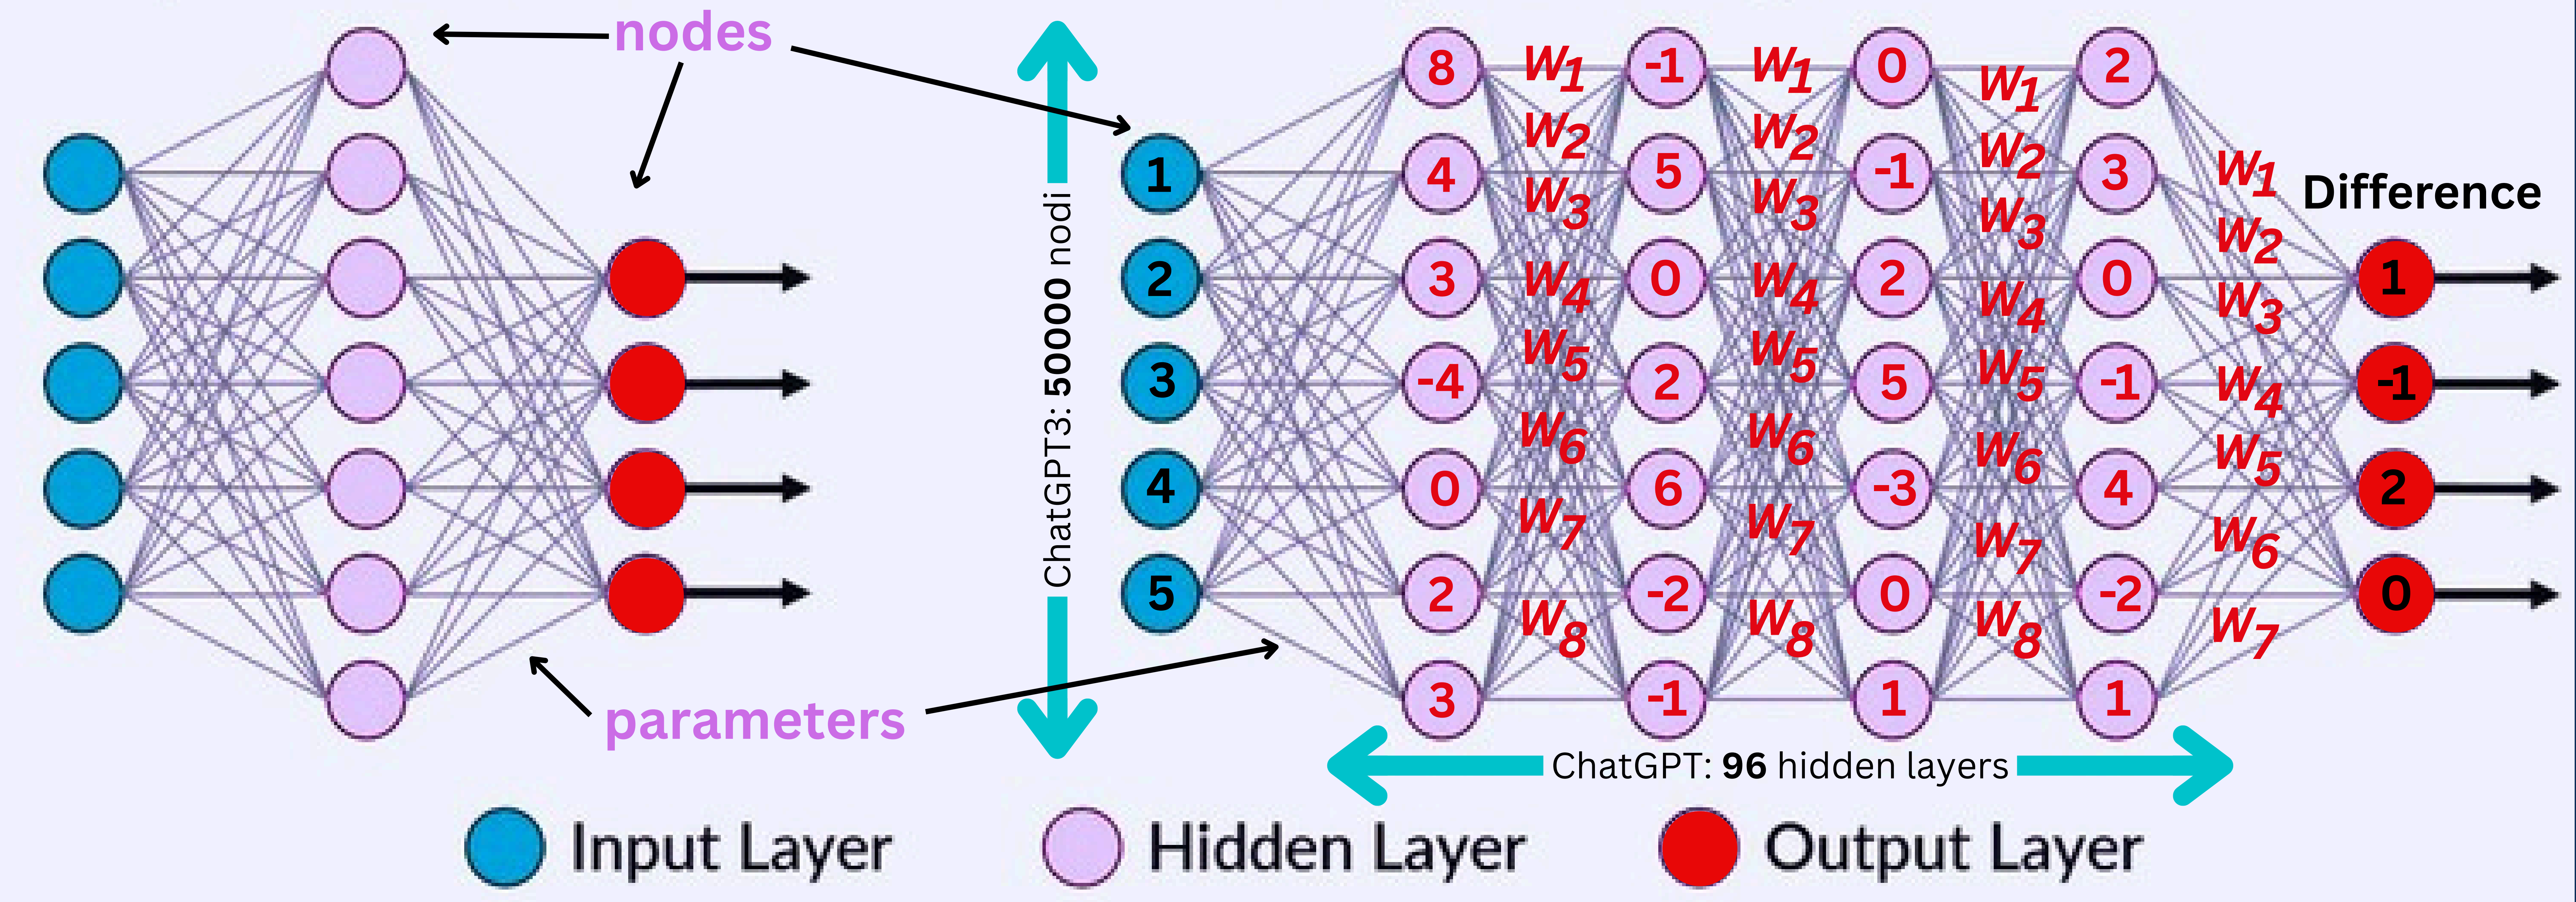

This type of learning is called supervised learning, and it is similar to what parents do to their children to teach them to speak: they repeat many sentences until the child learns which words to put one after the other

EMERGING PROPERTIES

The same Large Language Models are also used to generate summaries, write emails, translate into other languages, write code, compose poetry and songs, and so on, all tasks very different from predicting the next word. How do they do it?

EMERGING PROPERTIES

It takes a child years to learn to speak, but when the child goes to the elementary school and his teacher tells him to summarize a story, it doesn't take him other six years to learn how to summarize a text, because he already learned the structure of the language. All he needs are a few examples of summaries and after that he is able to summarize any text. In the same way, a Large Language Model is able to summarize any text after the researchers show it a few examples.

EMERGING PROPERTIES

This remarkable property of the deep learning models to learn new tasks quickly is the first example of an emergent property of the model. Emergent properties are abilities that the researchers didn't introduce explicitly in the model, they just appeared by themselves once the model reached a certain complexity, a certain number of parameters.

TRAINING DATA

The input data of Large Language Models must not only be abundant but also of good quality, just as a child grows better in an environment rich in external stimuli, and becomes more intelligent than a child who lives isolated.

CENSORSHIP

Large Language Models are trained on virtually all publicly available information on the Internet, removing text that is repeated multiple times, racist, sexist, and harmful information.

HIDDEN FIRST PART OF EACH PROMPT

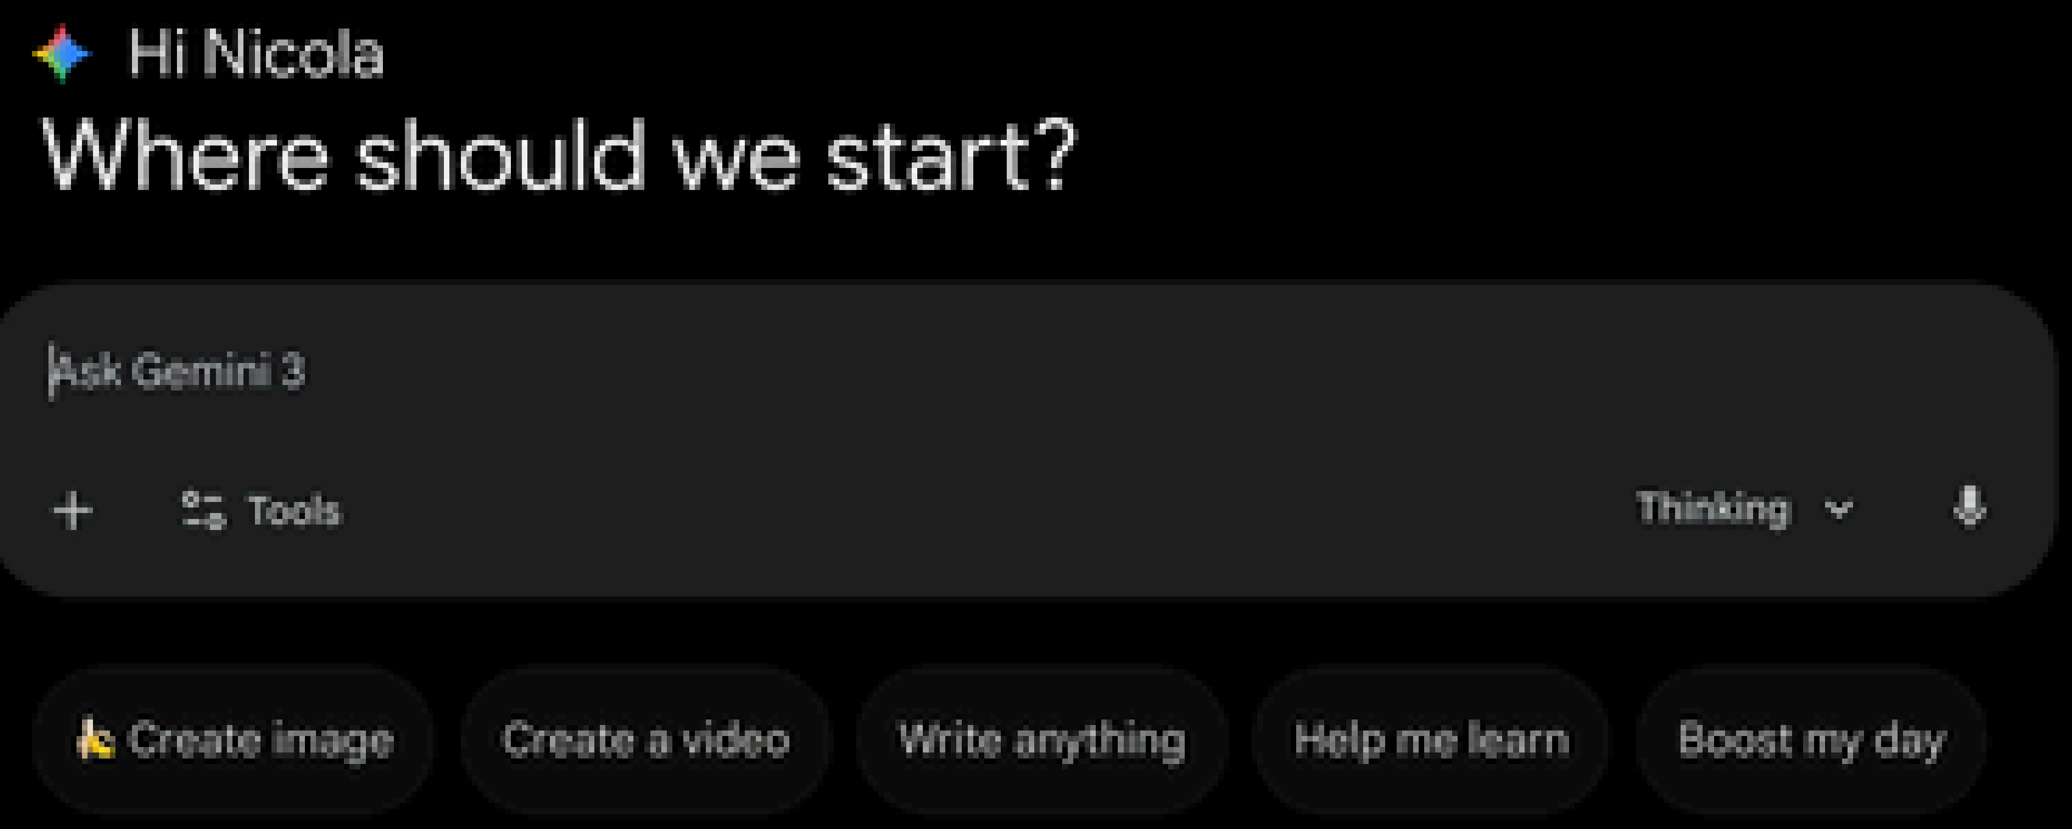

When writing a prompt for any language model, the model adds an invisible sentence to the beginning of the prompt, such as: "You are a friendly and helpful AI. Please answer the user's following question."

HUMAN LANGUAGE AND LARGE LANGUAGE MODELS

Large Language models allowed us to formulate a reasonable hypothesis to explain an age-old question: how did our articulated language arise? How did we go from simple animal calls to complex syntax?

HUMAN LANGUAGE AND LARGE LANGUAGE MODELS

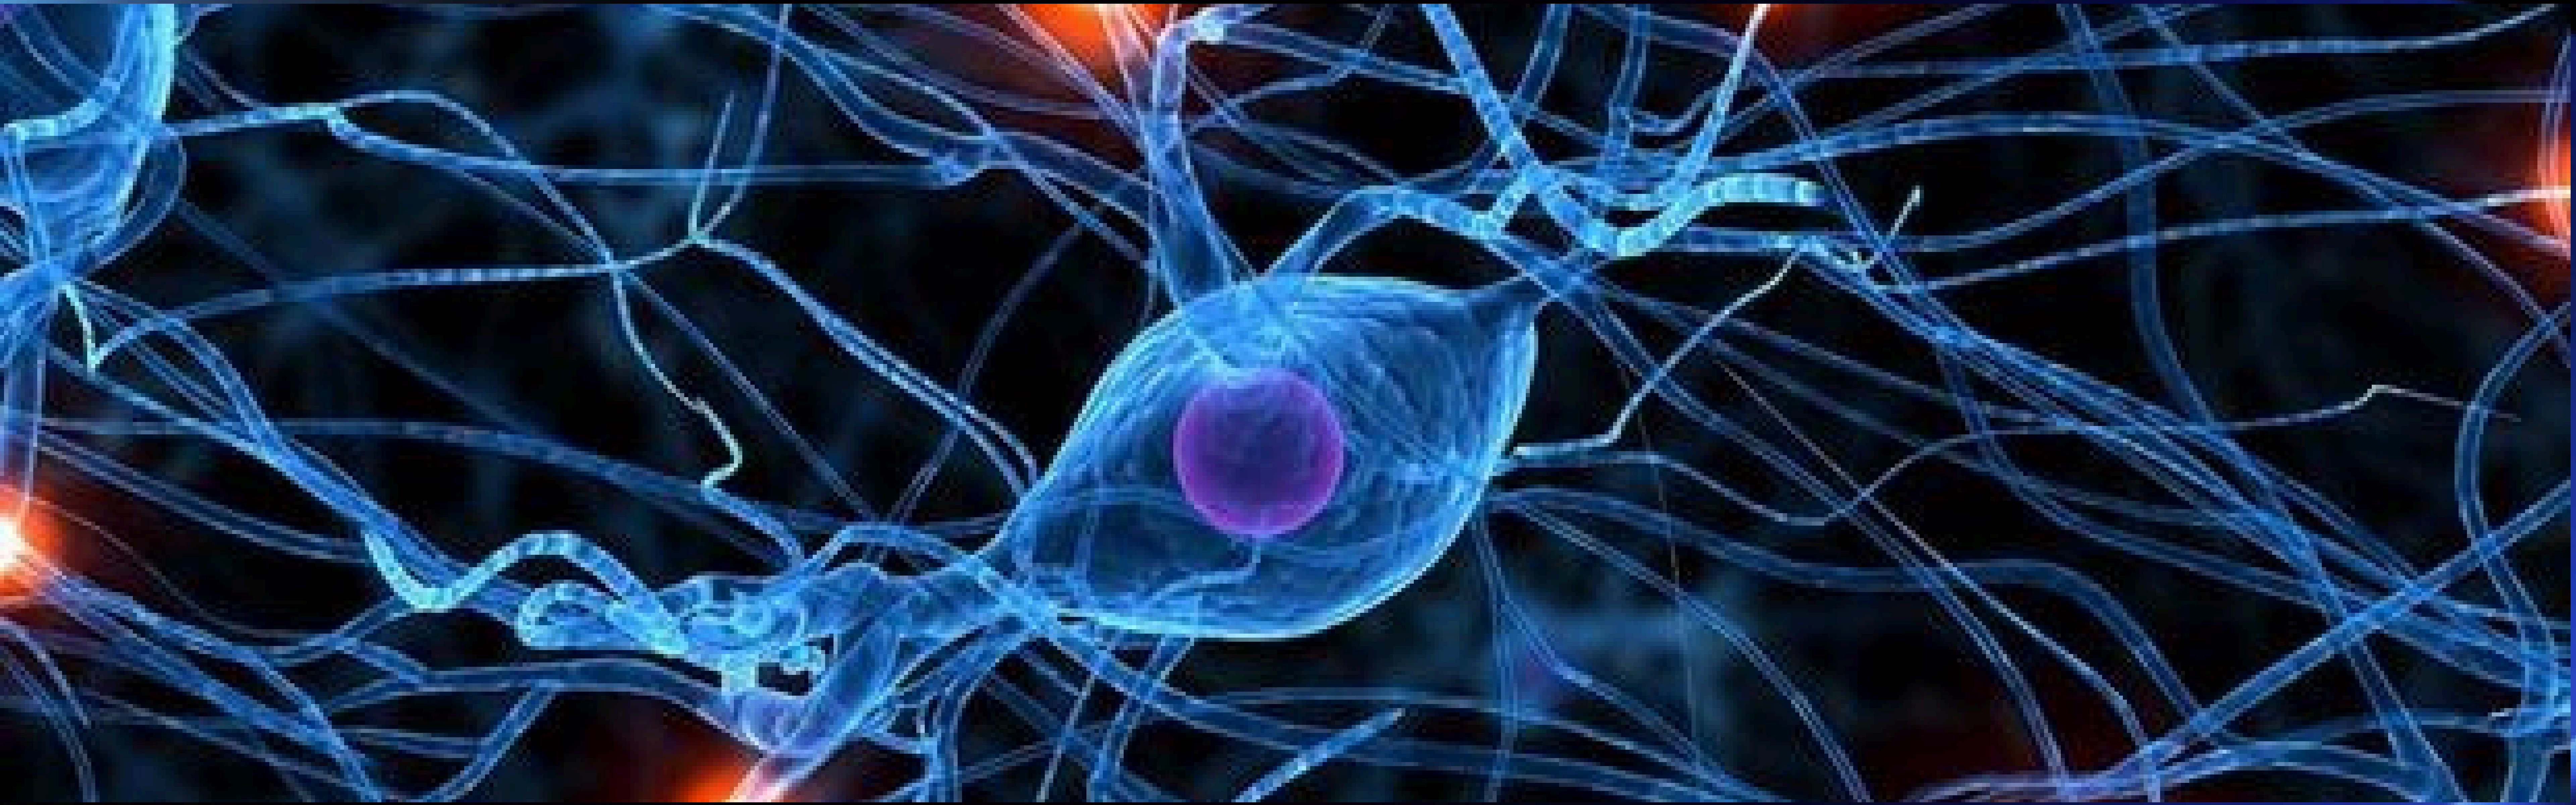

Human language is no more than 100.000 years old, a time too short for natural evolution to develop a new brain region specific for language. Thus, our brain performed a brilliant act of neuronal recycling: it took its existing neural network originally designed to predict the movement of objects in space and applied it to word prediction.

HUMAN LANGUAGE AND LARGE LANGUAGE MODELS

Many animals, in fact, including us, can easily predict the motion of objects. It was believed that the ability to speak was distinct from the ability to predict motion, until Large Language Models demonstrated that intelligent speech could be constructed by simply predicting the next word in a sequence.

OUR BRAIN DOESN'T THINK, BUT PREDICT

The brain isn't designed to perceive the world as it is, but to anticipate it well enough to increase its chances of survival. An animal that sees a predator may already be too late to escape. If it predicts its approach, however, it survives.

OUR BRAIN DOESN'T THINK, BUT PREDICT

The main advantage of language is that it allows us to simulate the future without experiencing it directly. Previously, we learned by dying. Later, we learned by listening to a story, even a short one, like "the lion is near the river." A story is a predictive simulation shared with other brains.

HUMAN LANGUAGE AND LARGE LANGUAGE MODELS

If this hypothesis is correct, then why don't animals talk?

HUMAN LANGUAGE AND LARGE LANGUAGE MODELS

Animals are born with all their synaptic "parameters" already calibrated by natural evolution: what we call “instinct”. A fawn can run within minutes of birth. Humans, on the other hand, are born immature, a characteristic of our species called neoteny

HUMAN LANGUAGE AND LARGE LANGUAGE MODELS

When we came down from the trees and went from four legs to two, the pelvic bones narrowed, and the newborn's head could no longer pass through the birth canal. The only way for us to be born was to be born prematurely, about a couple of months earlier than normal.

HUMAN LANGUAGE AND LARGE LANGUAGE MODELS

When we came down from the trees and went from four legs to two, the pelvic bones narrowed, and the newborn's head could no longer pass through the birth canal. The only way for us to be born was to be born prematurely, about a couple of months earlier than normal.

HUMAN LANGUAGE AND LARGE LANGUAGE MODELS

For this reason, when we are born, our synapses are not formed yet, our parameters are not determined yet by natural evolution. Instead, they are calibrated by our environment, that is to say, by our parents, that not only teach us how to walk and run, but also how to speak.

NEURAL NETWORKS ARE UNIVERSAL

Even linguistic models, however sophisticated, remain algorithms, and are therefore immaterial, that is, they are a property of information itself, not of matter. Therefore, they could also be found in other worlds or universes.

PROJECT WORK #1

- 150 new members in the last week
- Review the last facebook accounts in yellow to see if they are now recognized as human accounts
- Remember to remove the lower part of your images before publishing them, if they have the logo of Gemini/ChatGPT/…
- Remember to publish 1 post per week and some comments
- Do not publish or comment during Saturdays or Sundays!

PROJECT WORK #2

- New challenge: the first student to discover how to use AI Agents to automatically upload your videos in TikTok for free (instead of using Metricool/Buffer), will get a bonus of 3 points at the exam. Hint: look for free AI Agents that can interact with Buffer ("computer use")
- Only for the admin of the new TikTok channels: modify your TikTok language preferences by setting the same language of your channel. You can do it from: TikTok main menu → Setting and privacy → Language → App language
- Check in class the new TikTok channels, their number of visualizations and the new long videos

PROJECT WORK #3

- Add the name of the target city or destination chosen for your song in the new column Q of our google sheet
- Listen to the new songs in the Moodle of the students in class

PROJECT WORK #4:

- Today In class we will make the first part of Scene 11 together with your free 120 credits
- When you finish your 120 credits, subscribe for free to Google AI Pro only for the first month: in this way you will have 1000 free credits at your disposal
- Now you have one month of time to make one Scene of the movie (see next slides). Use you credits wisely; they are enough to make 50 video clips of 8 seconds each with Google Flow
- Then, 5 days before the end of the first free month, remember to cancel your subscription, so you won't be charged anything

PROJECT WORK #4:

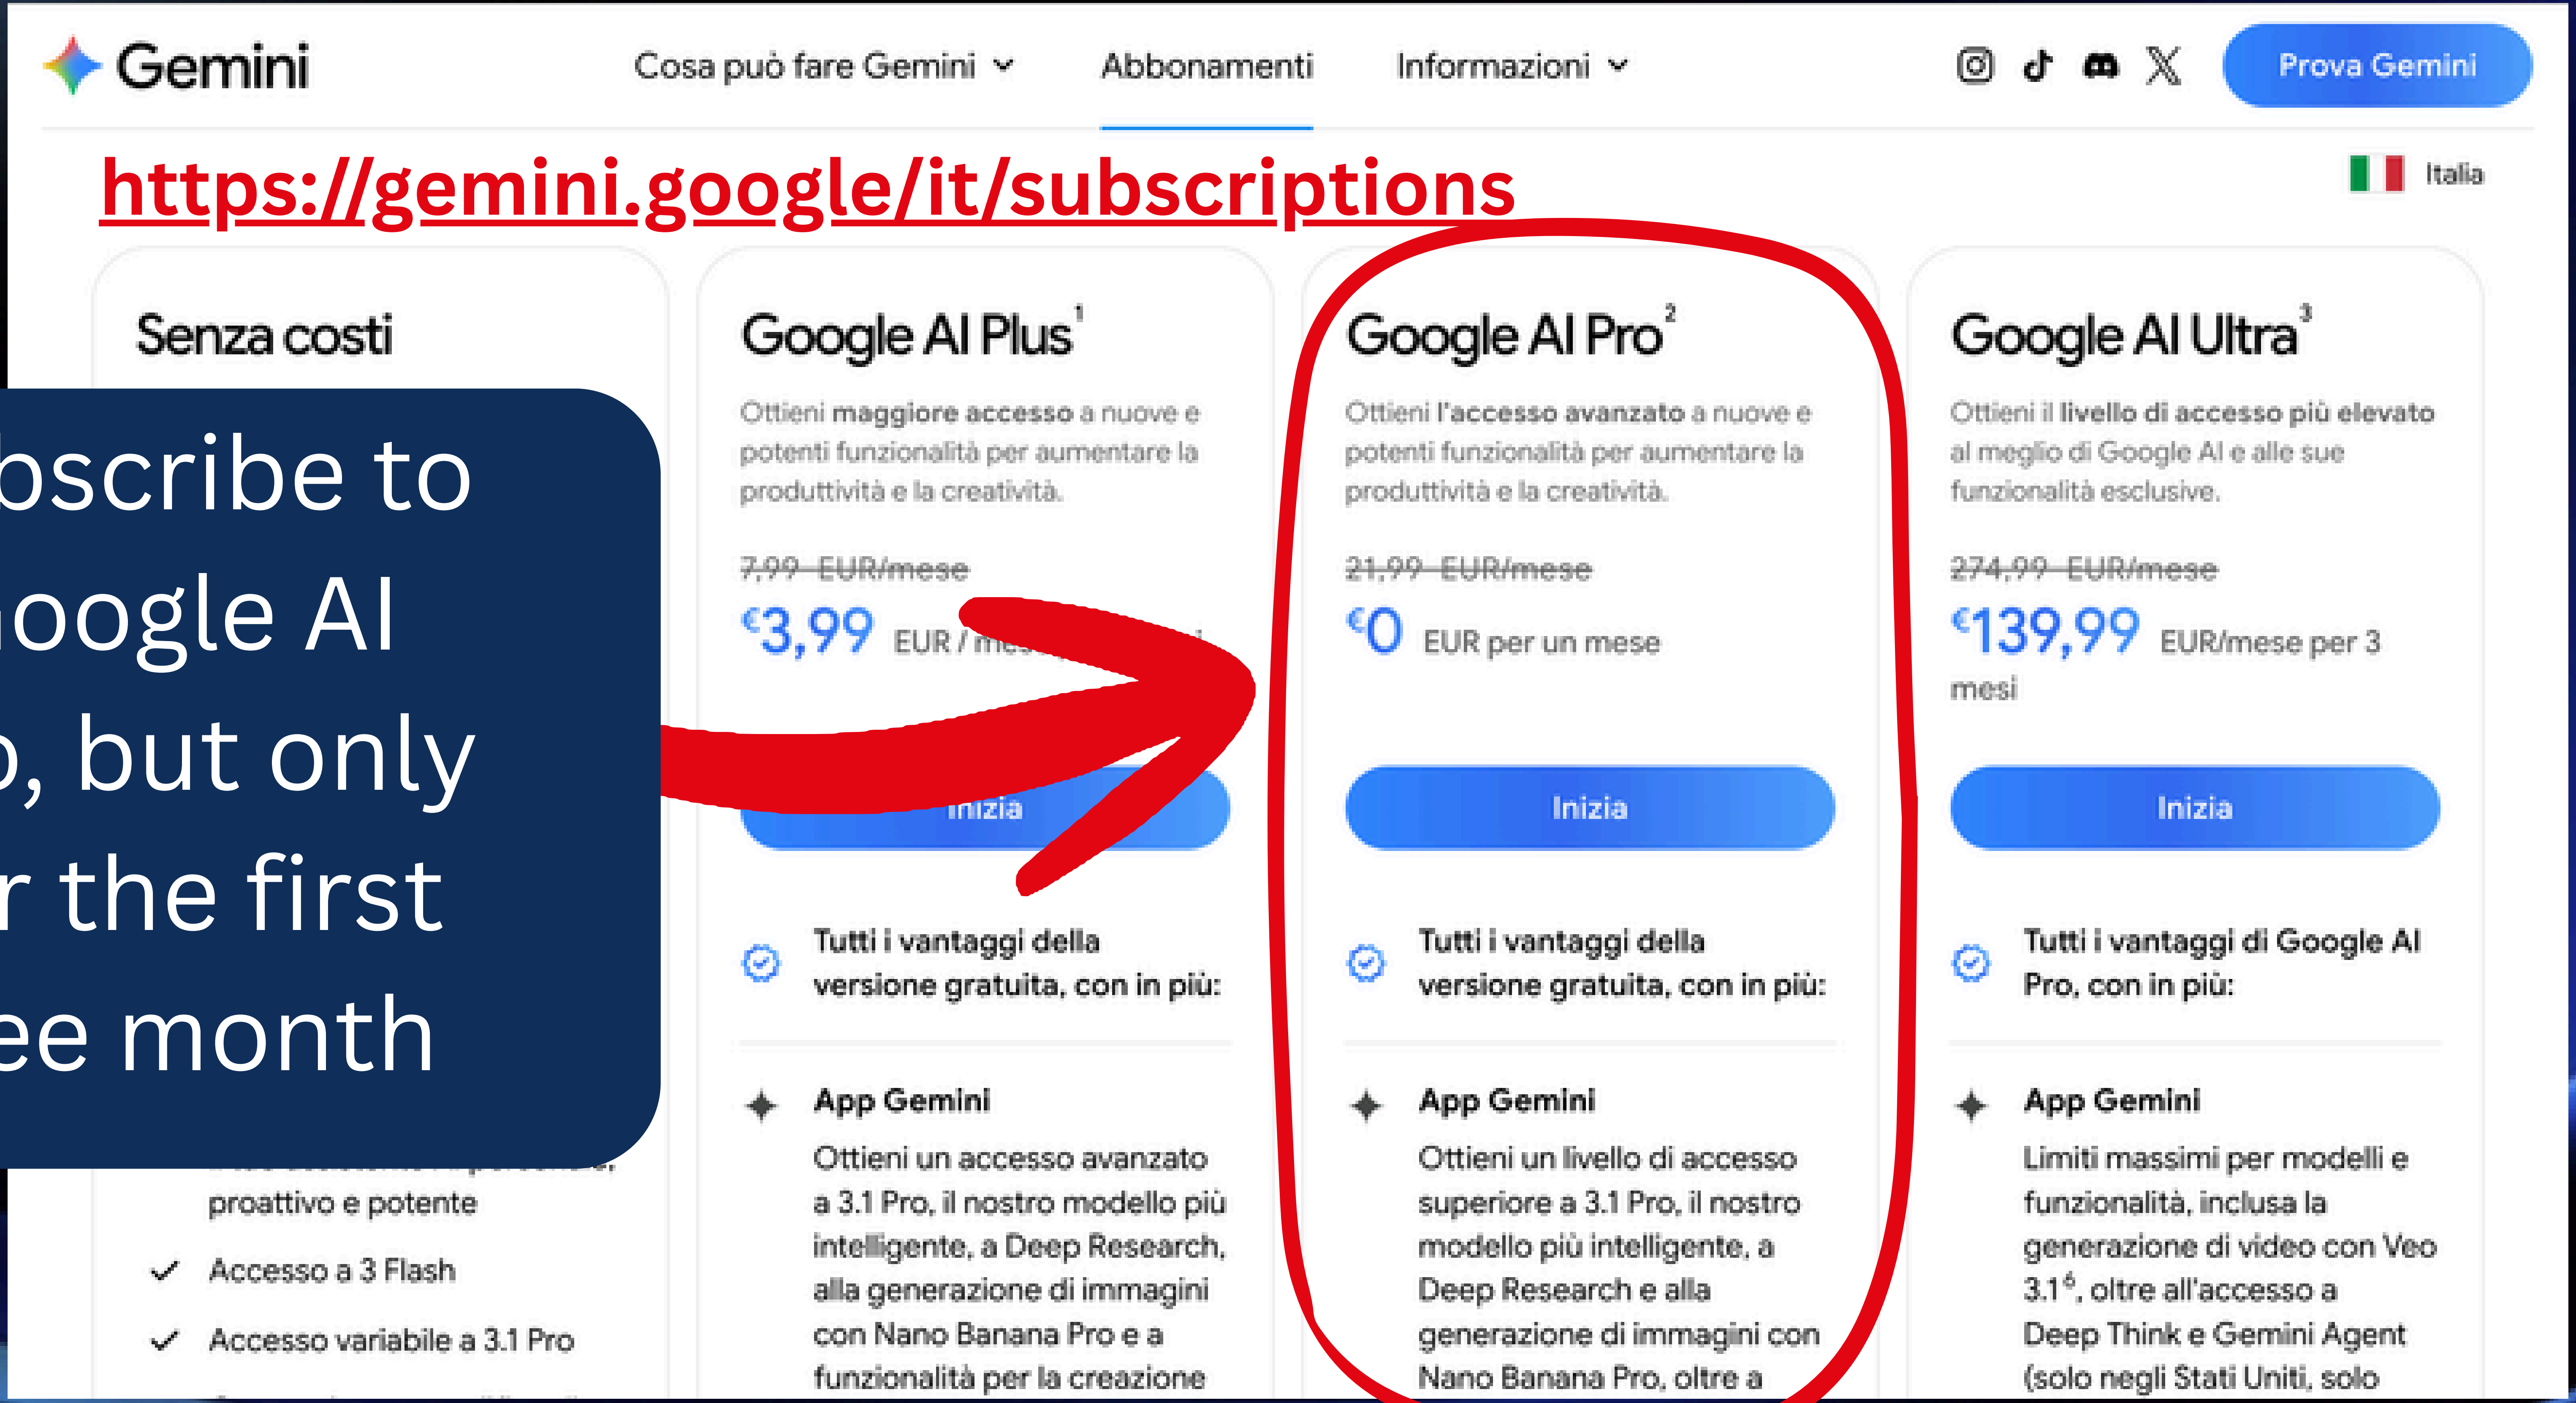

Subscribe to Google AI Pro, but only for the first free month

PROJECT WORK #4:

https://labs.google/fx/tools/flow

Cancel your subscription from here

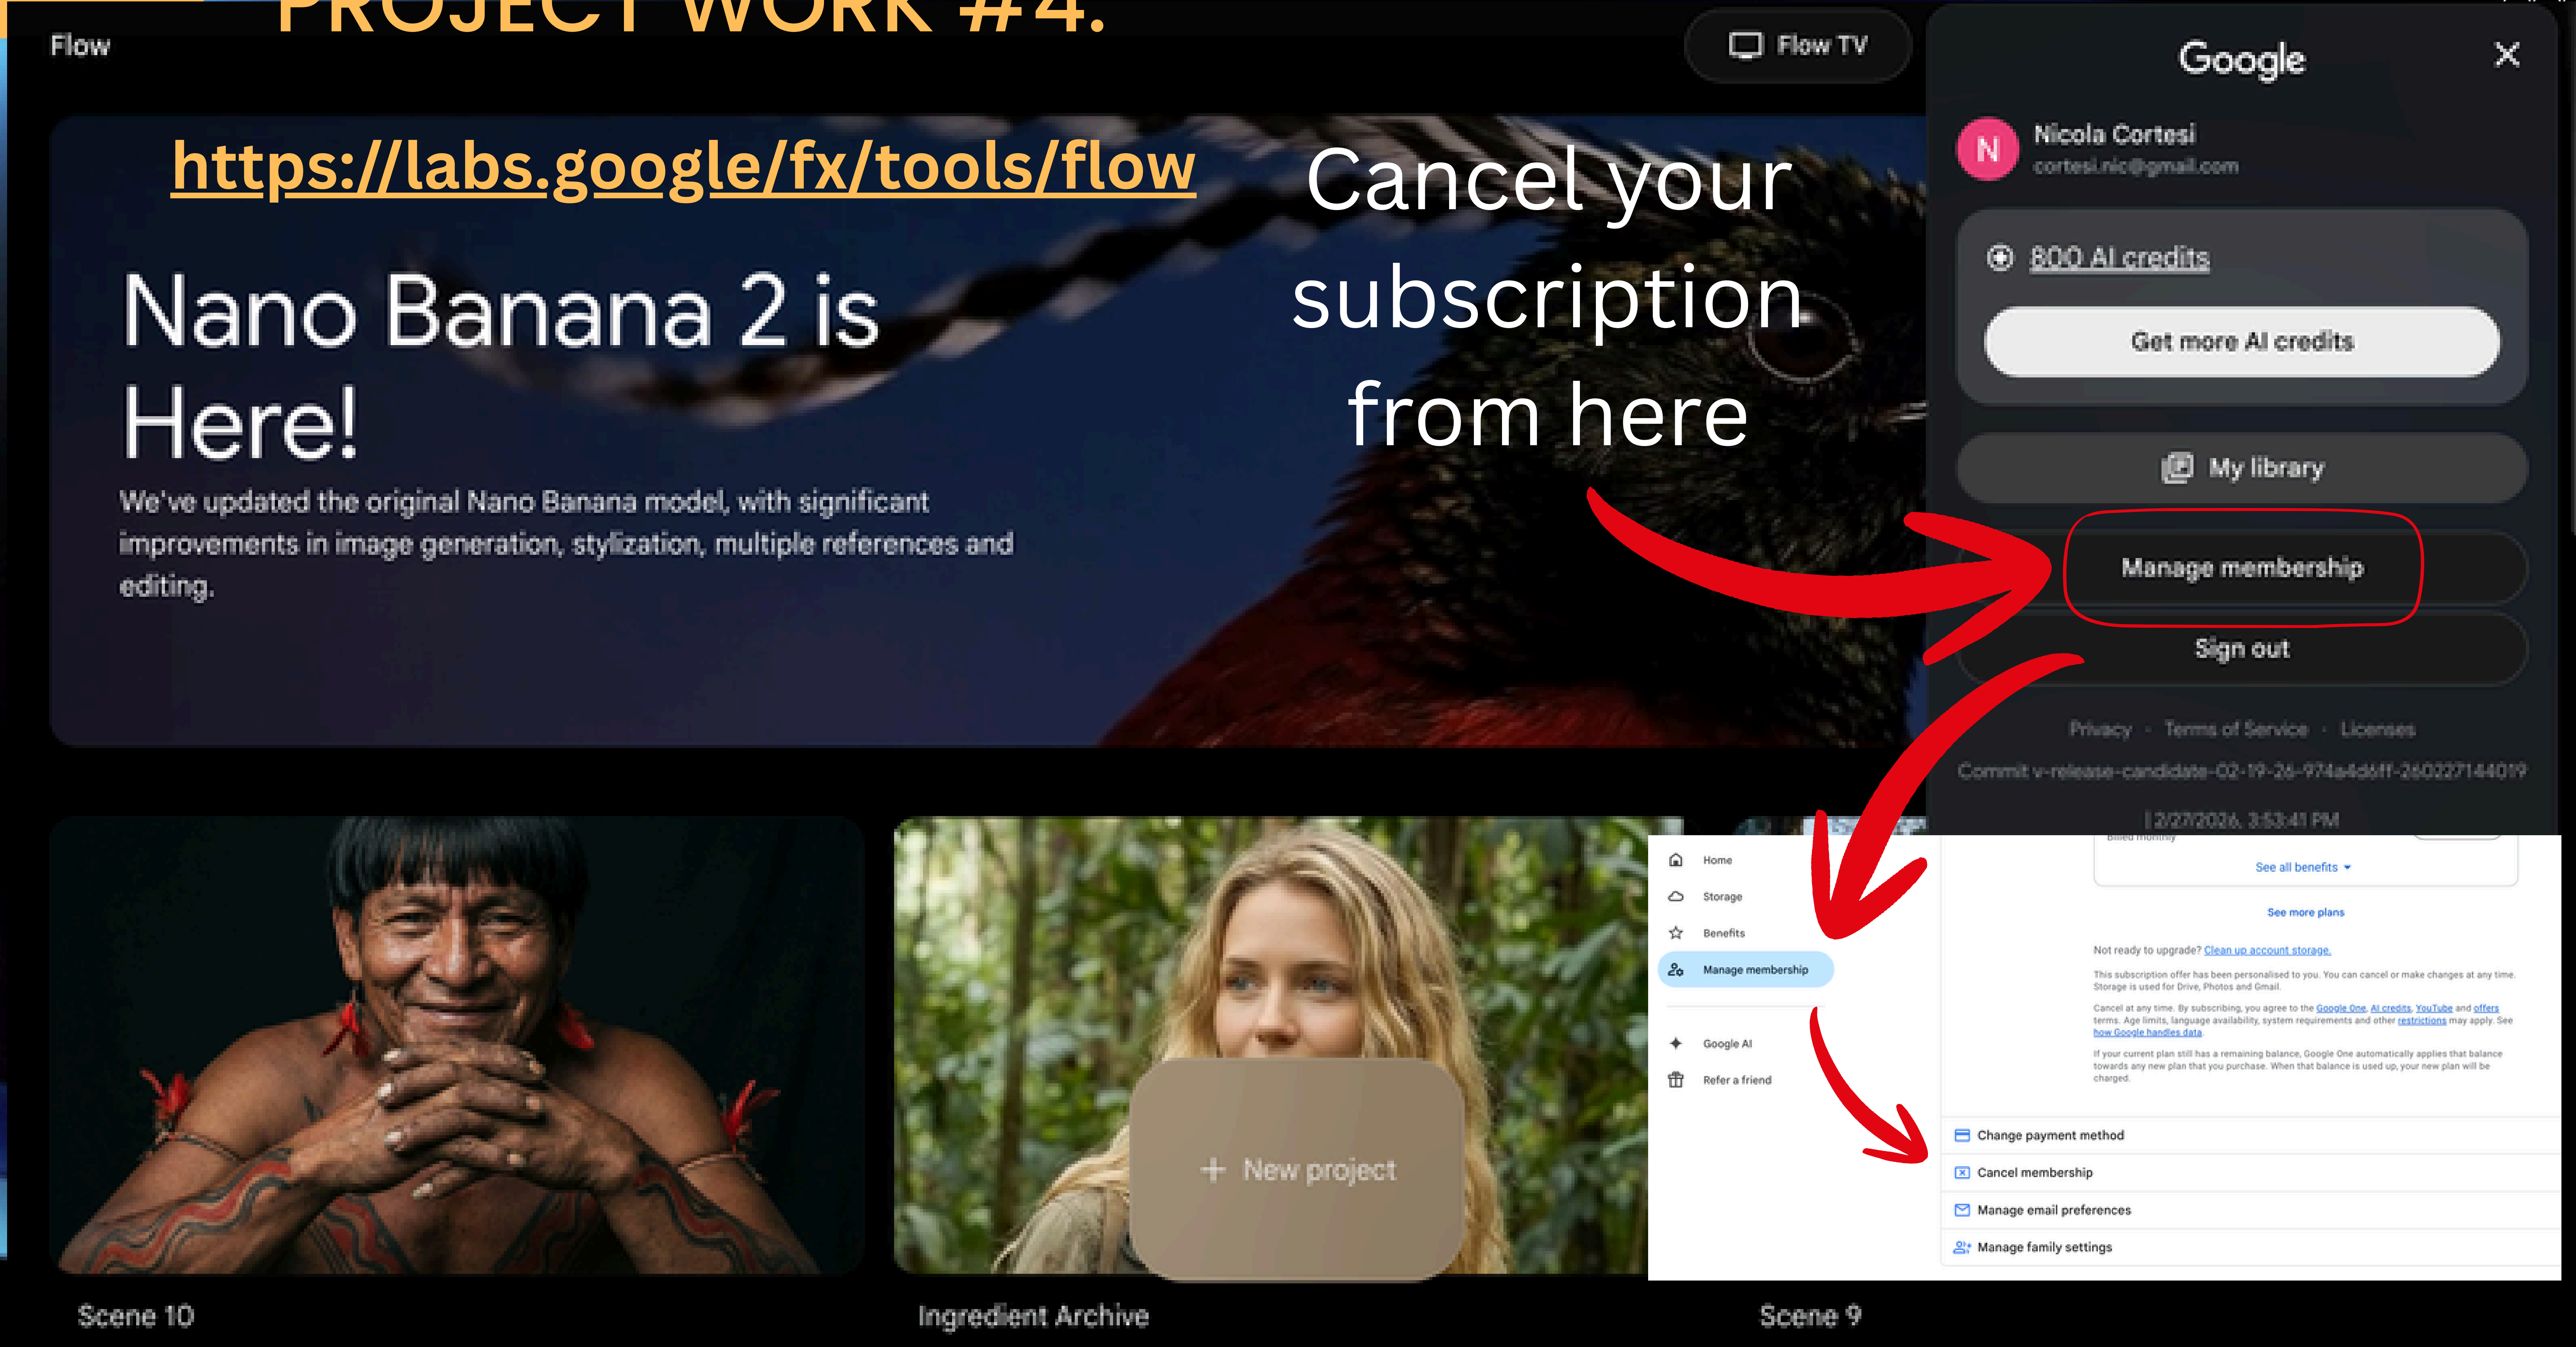

PROJECT WORK #4:

Each Group will make 3 to 5 Scenes, depending on its total number of members:

- Group #1: Scenes 18, 27, 35, 40, 42
- Group #2: Scenes 21, 39, 41
- Group #3: Scenes 14, 22, 32
- Group #4: Scenes 17, 23, 33, 38
- Group #6: Scenes 16, 26, 31, 36, 37

Each member of the group will only make one Scene out of the 3-5 assigned to his/her group, in order to save credits (check columns E-H of the Storyboard to discover exactly which Scene was assigned to you).

PROJECT WORK #4:

IMPORTANT UPDATE:

Last week Google introduced a major update to Google Flow, changing the graphical interface (the Scenebuilder now is useless) and giving some feature to Google Ultra users only.

Now the only way to ensure consistency from one shot to another is to save the last frame of your shot (right mouse click over the end of the video) and select it as “Start” frame of the “Frames” video type or as one image of the “Ingredients” video type

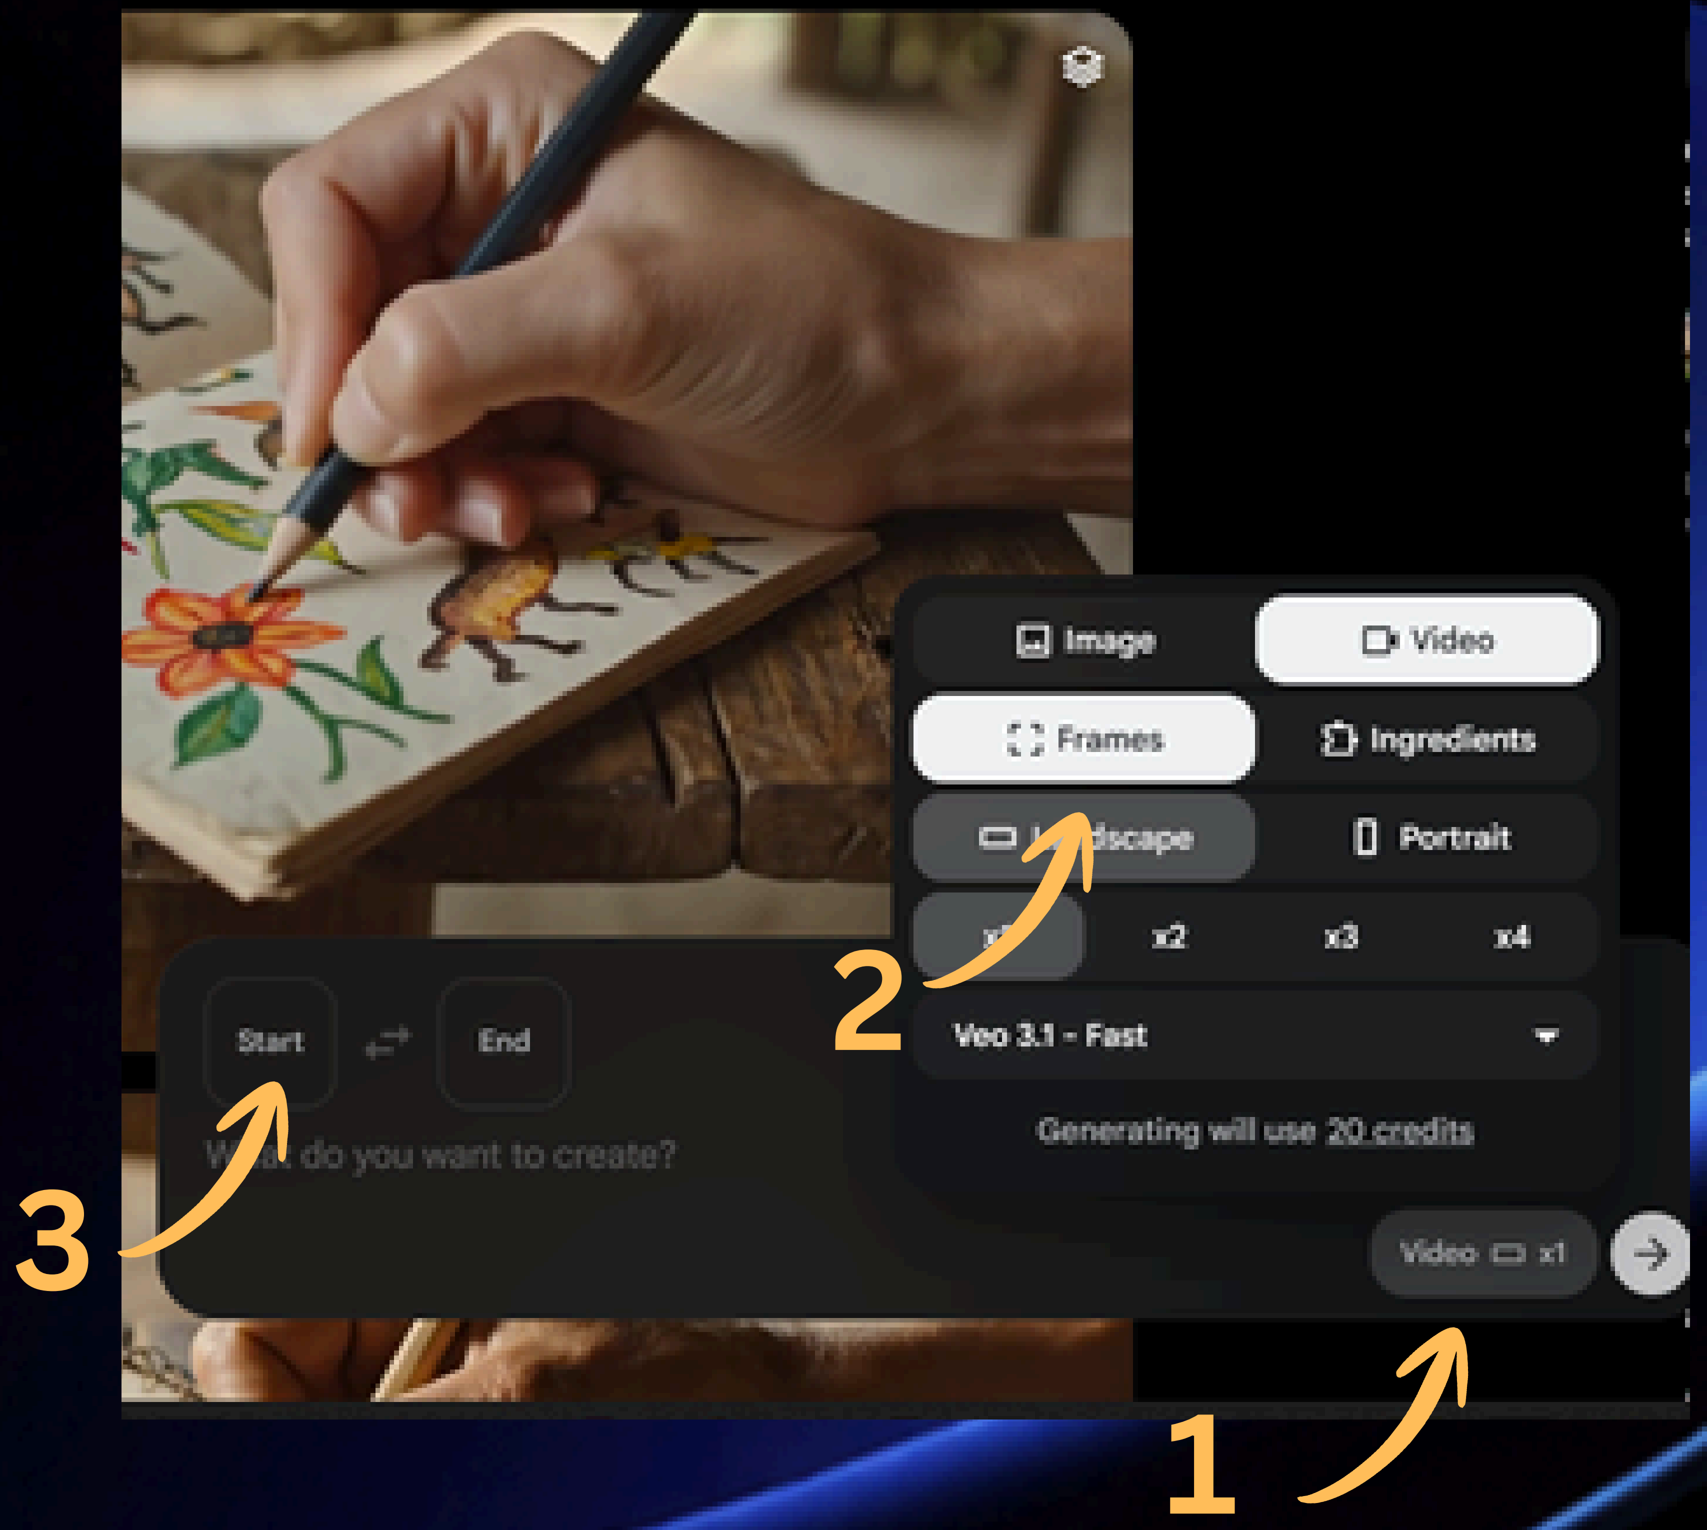

TASKS FOR NEXT WEEK

- Try to publish long videos in your TikTok channel, alternating them with short videos, to see if they perform better than them
- The day before next lesson, register to the first free month of Google AI Pro so you will be able to make video shots in class and for the following 3 lessons before your subscription expires

Suggested movie of the week: Tomorrowland (2005). It shows how AI might create a society of prosperity and abundance for everyone

Next lesson will be about Generative AI and Large Language Models