

LE SALON NOIR

AI COLLECTION · PHILOSOPHY

The mind behind artificial intelligence



Thinking, understanding, being conscious: a philosophical inquiry.

8 CHAPTERS · THOUGHT EXPERIMENTS

A voice replies. Who are we hearing?

You type a question. Seconds later an answer takes shape: orderly, nuanced, sometimes surprisingly considerate. You can challenge it, ask for an example, change your tone. The conversation continues. It becomes tempting to look behind the sentences for something resembling a mind.

This book starts at that precise point: between what a system does and what we imagine it is. It neither celebrates an artificial consciousness already demonstrated nor closes the question before examining it. It separates the problems. Could a performance be intelligent without being experienced? What does understanding a sentence involve? What evidence would justify attributing experience to another being?

The title also points to the human minds behind AI: the people who design it, those whose traces help shape its uses, and those who read intention into its responses. Thinking about machines brings us back to ourselves. We readily recognise an answer that satisfies us; we find it harder to explain why we think it is right.

Eight chapters move from these distinctions to thought experiments and ordinary situations. The scenes are fictional. They are neither observations about a particular model nor findings from a scientific experiment. The exercises invite you to develop an objection, compare criteria and specify what could change your mind.

You do not need an account with an AI provider to read the book or complete its exercises. A notebook is enough. If you choose to talk to a system, avoid confidential documents: invent an equivalent situation. The aim is not to secure a final answer, but to notice more carefully what each answer assumes.

About this edition

Our approach

An original introduction to philosophical inquiry. The positions discussed are arguments or hypotheses, not facts establishing the presence or absence of consciousness in a particular system. A convincing conversation alone is not a consciousness test.

Published by Le Salon Noir. First edition, September 2026. Text, translation and illustration created with artificial intelligence assistance and checked during this production. No affiliation with a research laboratory or model provider. The primary sources consulted are listed at the end of the book.

This edition is supplied for personal use. It includes neither resale rights nor access to an AI model or subscription. Referenced works and research remain attributed to their authors.

Your journey

What Do We Mean by a Mind?	5
Turing: Recognising Intelligence Through What It Does	8
The Chinese Room: Where Is Understanding Located?	11
Consciousness: Explaining an Experience, Not Just a Result	14
A Body, a World, an Experience	17
The Voice That Answers Us	20
What Must We Answer For?	23
Thinking with AI, Keeping Your Judgement	26
Words for thinking	29
Your judgement notebook	31
Sources and further reading	32

01 / A CORRECT ANSWER, AN IDEA UNDERSTOOD AND AN EXPERIENCE LIVED ARE NOT THREE NAMES FOR THE SAME THING.

What Do We Mean by a Mind?

The missed appointment

Imagine this scene. Aïcha asks a digital tool to move an appointment. It finds an available time, informs the participants and prepares a route. Everything works. That evening, a friend asks why she did not attend the meeting as originally planned. Aïcha replies, “I did not want to go.” The morning’s problem concerned a calendar; the evening’s concerned a decision whose motive she had hesitated to admit. A single action thus opens several questions: was the scheduling effective? Was the motive understood? Was anything experienced?

This fictional example establishes no permanent limit on machines. Instead, it teaches us not to let one success answer every question at once. A system may solve a particular problem admirably without that achievement, by itself, settling the nature of its relationship with the world. Philosophy begins here: by slowing down enough to establish what we are asking.

Intelligence: succeeding at what, under which conditions?

Let us adopt a working definition without presenting it as the only possible one: we will call intelligence the capacity to use information, learn or reason in order to respond to situations. This formulation remains deliberately open. Responding quickly, discovering a strategy, explaining an error and changing approach are different qualities. Excellence in one does not ensure excellence in another. A single overall score can conceal the very thing we wanted to examine.

In our scene, better planning also requires knowing what matters. Should the aim be to shorten the journey, accommodate a tired person or preserve a difficult conversation? An optimal solution depends on an objective. Choosing that objective is already a different question from pursuing it. We should therefore specify the task, constraints and criteria for success before declaring a behaviour intelligent. This does not diminish the achievement; it gives the claim a meaning that can be examined.

Understanding is more than repetition

To explore understanding, consider what we expect of someone who claims to have understood an instruction. They should be able to restate it, recognise when it applies and explain why an exception matters. None of these signs, taken alone, is an infallible certificate. Together, however, they suggest a richer requirement than reproducing a sentence. In this book, understanding will refer to making connections that render information usable and explainable within a context.

Meaning introduces another dimension. “It is late” can state the time, suggest leaving or request silence. The words remain the same; what they do in the exchange changes. We should therefore ask separately what a sentence refers to, how it fits a situation and what it enables its participants to do. Treating meaning as a single object hidden inside words would cause us to lose sight of this variety.

Thinking: one word, several kinds of activity

In ordinary speech, thinking can mean calculating, remembering, imagining a possibility or hesitating over a decision. These activities do not all have the same structure. Imagining a city without cars is different from adding two numbers; wondering whether to forgive someone cannot simply be replaced by predicting a consequence. We will therefore qualify the word whenever greater precision is needed.

This precaution avoids two opposite shortcuts: reserving thought for whatever resembles us, or calling every process that produces a result thought. Between these extremes, we can examine operations, explanations and architectures without deciding the answer in advance. The point is not to award an honorary title. It is to discover which similarities illuminate the phenomenon and which differences actually change our interpretation.

Experiencing something

Phenomenal consciousness means here that there is a lived experience: hearing a voice, feeling discomfort, seeing a colour. Thomas Nagel famously explored experience’s subjective character through the example of a bat, whose point of view is difficult to imagine from our own. His argument invites us to distinguish an external description from experiencing. It does not provide a universal consciousness test. [PHO4]

This distinction does not turn experience into magic. It simply identifies an additional question. Describing the conditions of a conversation and living through its embarrassment are not interchangeable descriptions. Conversely, uncertainty about an experience does not make a performance disappear. Let us keep the questions open alongside one another: what works, what means something and what is lived. Confusing them produces quick slogans; distinguishing them prepares a fruitful inquiry.

Three questions beneath one sentence

Take the claim: “This system understands my problem perfectly.” Rewrite it as three distinct propositions: one about what it accomplishes, one about the relationships it appears to grasp and one about any experience you may be attributing to it. For each, identify a relevant observation and what that observation would not establish. Aim for precision rather than a predetermined verdict.

Takeaway. Defining the questions does not delay inquiry into AI. It prevents a spectacular success from answering a different question.

Reading references: PH04