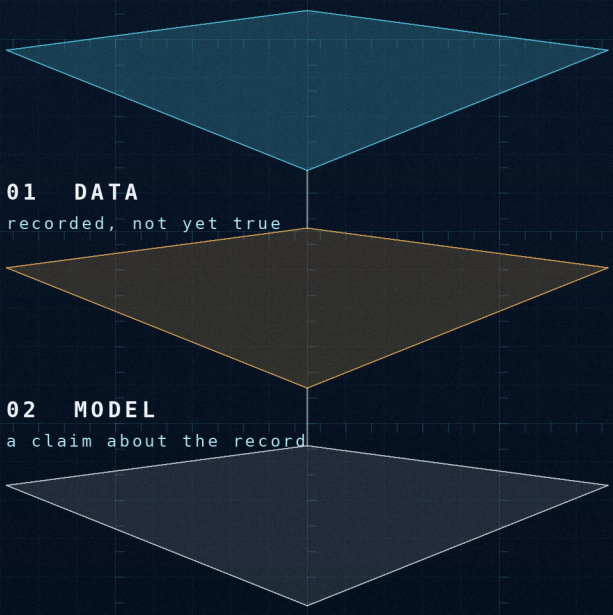
JACARANDA LABS PRESS · HARARE

Intelligence, Data and Judgement

An applied introduction to artificial intelligence, with a production case



COURAGE GARIKAI

FIRST EDITION · 2026

FIG. 1 — RAW RECORD, MODEL, AND THE PERSON WHO SIGNS

Intelligence, Data and Judgement

An applied introduction to artificial intelligence, with a production case

by Courage Garikai

Published by Coldcode SI Labs
First edition, 2026

Copyright © 2026 Courage Garikai. All rights reserved.

No part of this publication may be reproduced, distributed, or transmitted in any form or by any means without the prior written permission of the publisher, except for brief quotations in reviews.

Coldcode SI Labs Library · couragegarikai@atomicmail.io

CONTENTS

About this textbook 4

Who should read it 6

How to read it 7

How to use it in teaching 8

A note on evidence and honesty 9

Part One — Foundations 10

1 What intelligence is, and why the question is old 11

2 Four definitions, one discipline 18

3 A short history that explains today's defaults 23

4 What AI can do, what it cannot, and when to use it 28

5 Intelligence in the small firm 33

Part Two — The Material 37

6 Data: the raw material 38

7 Preparing data: the lifecycle 44

8 Features: the representation that decides the outcome 50

9 Seeing data: descriptive statistics and visualisation 56

10 Learning from data 61

11 Measuring a model honestly 68

12 The arithmetic of being wrong 74

Part Three — The Machine 80

13 Models in production 81

14 Algorithms, procedures and pseudocode 87

15 Languages, APIs and the cloud 92

Part Four — The Case 98

16 Designing under constraint: the rail 99

17 The gate: finding prospects without lying to yourself . . 106

18 What broke 112

Part Five — Judgement 118

19 Ethics, consent and the things a system should refuse . 119

20 Law, data protection and the state 126

21 Teaching it: curriculum, sequence and assessment 133

22 Conclusion: the discipline of the second question 139

Appendix A · Glossary 143

Appendix B · Eight labs 148

Appendix C · Mapping to the three source curricula 150

Appendix D · References and sources 154

Colophon 159

About this textbook

This book teaches the fundamentals of artificial intelligence to people who will have to use it — engineers, operators, teachers, and the owners of small businesses in places where the technology arrives later than the marketing. It is written to be taught from: every chapter opens with learning outcomes, carries a worked example, and closes with exercises and further reading.

The book has a companion that most introductions do not: a running case. **Tapstall** is a merchant commerce rail that I built and operate in Zimbabwe, South Africa, Kenya and Nigeria. It runs on free infrastructure, holds no customer money, and uses AI in precisely the places where a small firm needs it — finding prospects, writing copy, reconciling orders. Every abstract idea in this book is attached to something Tapstall does, or to something Tapstall got wrong. Chapter Sixteen onwards handles the case directly; earlier chapters refer to it in passing, so that the theory never floats.

Three existing curricula form the disciplinary spine: the COMCEC introduction to AI and its fundamental concepts (Kurt, 2025), prepared under the Organisation of Islamic Cooperation's Standing Committee for Economic and Commercial Cooperation; Microsoft's Fundamentals of Artificial Intelligence, Module 1 (Microsoft, 2020); and Let's Learn Artificial Intelligence, the Base Module published by NITI Aayog's Atal Innovation Mission with NASSCOM (NITI Aayog, 2020). These are cited throughout as [COMCEC], [Microsoft] and [NITI Aayog], and Appendix D maps every chapter of this book onto the sections of all three, so

that an instructor teaching from any of them can slot this material in.

Who should read it

Undergraduate and postgraduate students taking a first or second course in AI who want the standard material connected to an actual deployment rather than a toy dataset.

Practitioners who can call an API but cannot explain why the model's answer should be trusted, and what to do when it is not.

Non-specialist professionals — accountants, teachers, operators, public servants — who need enough vocabulary and judgement to evaluate what a vendor is selling them.

Instructors who want a one-semester course with labs, exercises, a case, and a mapping to the three widely used curricula above.

How to read it

The book is in five parts. **Part One** asks what intelligence is, why the definitions differ, and how those differences survive into systems. **Part Two** is the material: data, features, statistics, learning, and the arithmetic of being wrong. **Part Three** is the machine: algorithms, languages, APIs, cloud economics, and the lifecycle of a model in production. **Part Four** is the case: three chapters of Tapstall, ending with a forensic account of what fails. **Part Five** is judgement: ethics, law, and pedagogy.

Part One and Part Two are the core of any course. Part Three can be compressed for a non-technical cohort. Part Four should not be skipped by anyone who intends to deploy anything, because it is where the book stops describing AI and starts describing what happens when AI meets a Tuesday.

How to use it in teaching

A twelve-week semester at one two-hour session per week fits chapters one to twenty with two weeks on Part Four. A shorter six-week module should take chapters one, two, four, six, eight, nine, ten, nineteen. Appendix B supplies eight labs — four of them runnable with no code at all — and Appendix C collects the exercises by chapter. Each chapter’s exercises are graded by difficulty: recall questions test vocabulary, applied questions require a calculation or a design, and critical questions have no single right answer and should be argued in class.

A note on evidence and honesty

Where this book states something about the three source curricula, it is quoting or paraphrasing their content, attributed as [COMCEC], [Microsoft] or [NITI Aayog]. Where it states something about Tapstall, the figure comes from the running system or its public history — the release chain, the test suite, the production endpoints — and I say when a figure is a design ceiling rather than an observation. Where the book argues a position rather than reporting a fact, it says so. This distinction is the book's own first lesson: the difference between a claim about the world and a claim about the author is the difference between data and opinion, and most failures of applied AI begin with a blurring of the two.

One further admission. Tapstall has no paying customers at the time of writing. It is a working system with a verified test surface and a forensic record, not a validated commercial success. The case is used in this book for its architecture and its failure modes, both of which are real and inspectable. Any instructor who wants a case about market success will need a different case; any instructor who wants a case about what actually goes wrong will find this one useful.

PART ONE

Foundations

CHAPTER 1

What intelligence is, and why the question is old

Learning outcomes. By the end of this chapter you should be able to: state three common criteria by which human intelligence is described; explain why a single definition of artificial intelligence is unavailable; distinguish narrow from general intelligence; and describe the Turing test and the Chinese Room experiment and the different questions each raises.

1.1 The question a student asks first

The first question in every AI course is the one with no settled answer: what is intelligence? NITI Aayog's Base Module opens Unit 1 with it directly, and its method is instructive — rather than offering a definition and defending it, the module offers a checklist. When are we called intelligent? When we answer questions correctly; when we solve a difficult puzzle; when we decide the right action; when we learn quickly; when we recognise things and remember their names; when we use clues to guess. All of these, the module concludes, are examples of natural human intelligence, and a common set of abilities runs through the definitions: the ability to observe, recognise and understand; the ability to make smart decisions and solve problems; and the ability to learn and increase knowledge. An agent that can do those things is, in the module's phrase, an intelligent agent [NITI Aayog].

That checklist is doing serious work, and a student should notice what kind of work it is. It converts a metaphysical