The Fundamentals of AI

From First Principles to the Perfect Prompt

The Fundamentals of AI

From First Principles to the Perfect Prompt

A Practical, Step-by-Step Reference for Understanding and Mastering AI Prompting

Remco Dekkinga

Evergreen IT

Title: The Fundamentals of AI
Subtitle: From First Principles to the Perfect Prompt
Author: Remco Dekkinga
Publisher: Bookmundo

Code snippets with this icon can be downloaded from our website: https://evergreen-it.nl/books

ISBN: 978-94-0389-574-1
First print, May 2026

Table of Contents

Preface

AI is already in the tools you use every day. It can be genuinely useful, and it can also sound confident while being wrong.

This book gives you a method. We treat prompting like engineering: define the job, control the inputs, set clear constraints, and check the result.

My promise is practical: when an output is wrong or vague, you will know what to change. You will add context, tighten constraints, give a better example, or change the requested format, then test again.

This is for people who need AI to help in real work, not just in demos. Software professionals, architects, developers, technical leaders, educators, and low-code or no-code builders will find techniques you can reuse immediately.

If you are done guessing, start here. Let's get to work.

Remco Dekkinga

1 How This Book Is Structured

Each chapter introduces one core idea and then puts it to work right away. You will see an evolving system prompt written in Markdown, plus short exercises that help you improve it step-by-step. Read cover to cover and it will feel like a guided course. Dip in and out and each chapter will still hold up as a practical reference.

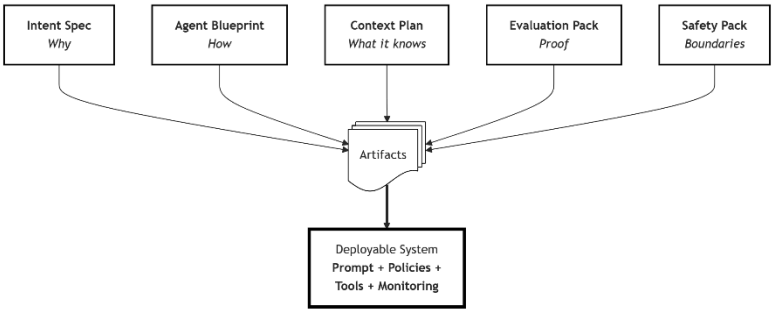

Along the way, you will build five "design artifacts". Think of them as the simple paperwork that turns a clever prompt into something you can trust in real use. Each artifact answers a different question: **Why are we building this? How should it behave? What is it allowed to know? How will we prove it works?** And **when should it stop?**

Fundamentals (Chapter 2): Before we design anything, we need a shared mental model of what language models do well, what they don't do at all, and why they sometimes sound confident while being wrong. This chapter explains tokens, context windows, and grounding. The goal is not to write better prompts yet. The goal is to understand the constraints of the machine, so the rest of the book feels predictable rather than magical.

Intent Spec (Chapter 3): Before we touch tools or architecture, we get clear on the why. This is the part many teams rush, and it is