

AI BEZ STRACHU - 2026

The Visual Encyclopaedia of AI for the Over-40s

80 focused chapters. 10 leadership topics. AI for
professionals who ship.

FREE 10-CHAPTER SAMPLE - AGE 40+

Inside this sample

The first ten chapters are reproduced from the complete edition. Download, read and share this sample; purchase links for the complete PDF and physical print-on-demand edition are available at aibezstrachu.matsiems.com/books.

- 01 What AI Actually Is
- 02 The 2023 Inflection Point
- 03 Three Waves of AI in Business
- 04 What AI Cannot Do
- 05 The AI Competitive Landscape
- 06 ROI Realism
- 07 Vocabulary You Need
- 08 The Boardroom Question
- 09 Build vs Buy vs Integrate
- 10 The AI Maturity Model

CHAPTER 01

What AI Actually Is

Strip away the hype: AI is statistical pattern-matching at massive scale. It predicts the next token, the next move, the next diagnosis. Understanding this one fact changes how you evaluate every AI claim.

Strip away the hype: AI is statistical pattern-matching at massive scale. It predicts the next token, the next move, the next diagnosis. Understanding this one fact changes how you evaluate every AI claim.

Key ideas

- GPT-4 has ~1.76 trillion parameters — each a tunable number.
- The model never 'understands' — it finds the most statistically likely continuation.
- Every AI product you use today is built on this same foundation.

CHAPTER 02

The 2023 Inflection Point

ChatGPT reached 100M users in 60 days — the fastest product adoption in history. The shift wasn't a gradual evolution; it was a step-change. Leaders who grasped this early moved fast.

ChatGPT reached 100M users in 60 days — the fastest product adoption in history. The shift wasn't a gradual evolution; it was a step-change. Leaders who grasped this early moved fast.

Key ideas

- ChatGPT hit 1M users in 5 days (Netflix took 3.5 years).
- Within 6 months, every major tech firm had a competing product.
- The inflection was driven by scale: models became useful at ~10B parameters.

CHAPTER 03

Three Waves of AI in Business

Wave 1 (2016–2020): narrow AI in specific tasks. Wave 2 (2021–2023): large language models. Wave 3 (2024–): agentic AI that plans and acts autonomously. You are managing Wave 3 now.

Wave 1 (2016–2020): narrow AI in specific tasks. Wave 2 (2021–2023): large language models. Wave 3 (2024–): agentic AI that plans and acts autonomously. You are managing Wave 3 now.

Key ideas

- Wave 1 AI required labelled training data per task.
- Wave 2 enabled zero-shot generalisation across domains.
- Wave 3 agents can book meetings, write code, and iterate without human prompts.

CHAPTER 04

What AI Cannot Do

AI hallucinates with confidence. It has no persistent memory by default, no real-time data access, no genuine reasoning. Leaders who know the limits make better decisions than those dazzled by the demos.

AI hallucinates with confidence. It has no persistent memory by default, no real-time data access, no genuine reasoning. Leaders who know the limits make better decisions than those dazzled by the demos.

Key ideas

- Hallucination rates vary 3–20% depending on task and model.
- Without retrieval augmentation, models know nothing after their training cut-off.
- AI cannot reliably perform multi-step logical deduction without chain-of-thought scaffolding.

CHAPTER 05

The AI Competitive Landscape

OpenAI leads but Anthropic, Google DeepMind, Meta, and Mistral are all formidable. Open-source models (Llama 3, Mistral) close the gap monthly. Your vendor strategy must account for rapid change.

OpenAI leads but Anthropic, Google DeepMind, Meta, and Mistral are all formidable. Open-source models (Llama 3, Mistral) close the gap monthly. Your vendor strategy must account for rapid change.

Key ideas

- Anthropic Claude, Google Gemini, and OpenAI GPT-4o are the enterprise frontrunners.
- Meta's Llama 3 70B runs on a single A100 — no API needed.
- Model performance benchmarks are obsolete within 3–6 months.

CHAPTER 06

ROI Realism

Most AI pilots fail to scale because the business case was vague. Real ROI comes from automating specific, high-frequency, low-stakes decisions — not from 'digital transformation'.

Most AI pilots fail to scale because the business case was vague. Real ROI comes from automating specific, high-frequency, low-stakes decisions — not from 'digital transformation'.

Key ideas

- McKinsey: AI adds \$4.4T in annual economic value — most in knowledge work.
- Average enterprise AI pilot: 6–18 months, 40% abandoned before production.
- Quick wins: document summarisation, code review, first-draft generation.

CHAPTER 07

Vocabulary You Need

LLM, prompt, token, fine-tuning, RAG, agent, embedding — these terms are used loosely. Knowing precise definitions prevents vendor confusion and misaligned expectations.

LLM, prompt, token, fine-tuning, RAG, agent, embedding — these terms are used loosely. Knowing precise definitions prevents vendor confusion and misaligned expectations.

Key ideas

- A token is ~ 0.75 English words — pricing and context windows are measured in tokens.
- Fine-tuning adjusts model weights; prompting does not.
- RAG (Retrieval-Augmented Generation) injects external knowledge at inference time.

CHAPTER 08

The Boardroom Question

Every executive is being asked: how does AI affect our business model? The right answer is a matrix: which tasks get faster, which jobs transform, which competitors are already ahead.

Every executive is being asked: how does AI affect our business model? The right answer is a matrix: which tasks get faster, which jobs transform, which competitors are already ahead.

Key ideas

- PwC: 30% of UK jobs could be automatable with current AI by 2030.
- Automation displaces tasks, not roles — most jobs will be reshaped, not eliminated.
- Early movers in AI adoption show 15–20% productivity gains in knowledge work.

CHAPTER 09

Build vs Buy vs Integrate

Building a custom model is expensive and rarely necessary. Buying an AI product is fast but shallow. The real value is in integrating AI into your existing workflows via API. Know which strategy fits each use case.

Building a custom model is expensive and rarely necessary. Buying an AI product is fast but shallow. The real value is in integrating AI into your existing workflows via API. Know which strategy fits each use case.

Key ideas

- Training a frontier model costs \$10M-\$100M+.
- Off-the-shelf AI tools pay back in weeks for admin tasks.
- API integration typically costs £5-50k and delivers in 2-4 months.

CHAPTER 10

The AI Maturity Model

Level 1: experimental (individual tools). Level 2: functional (department automation). Level 3: strategic (AI in decision-making). Level 4: transformational (AI-native business model). Where is your organisation?

Level 1: experimental (individual tools). Level 2: functional (department automation). Level 3: strategic (AI in decision-making). Level 4: transformational (AI-native business model). Where is your organisation?

Key ideas

- 70% of large enterprises are at Level 1-2.
- Level 3+ requires data infrastructure, governance, and cultural change.
- AI-native companies (Airbnb, Stripe, Uber) bake AI into their core product loops.