

AI RESOURCE GUIDE

Understand AI. Apply It. Think Ahead.

Curated recommendations for podcasts, books, and training programs, I organized by entry level and depth.

Not a link dump, but an opinionated selection: what was worth my time, why, and what you'll take away. Every recommendation comes with my personal take: what I got out of it, what stuck, and who it's best for.

Daniel Mühlbauer-Kerber · March 2026



Podcasts

Audio is the most efficient way to stay current. This selection covers the spectrum from accessible entry points to deep technical dives.

● GETTING STARTED

Hard Fork

Kevin Roose & Casey Newton · The New York Times

My take: The best entry-level podcast for people who want to understand AI in context, not in isolation. Roose and Newton place AI developments within the broader tech landscape without the hype. Their strength: making complex implications accessible in 40 minutes.

Free · Weekly · ~45 min

Practical AI

Chris Benson & Daniel Whitenack

My take: The name says it all. Not a research deep-dive, but: how are organizations actually deploying AI? Especially useful for HR and business professionals who want to understand what happens beyond the demos. The hosts keep technical topics accessible without oversimplifying.

Free · Weekly · ~50 min

AI For Humans

Gavin Purcell & Kevin Pereira

My take: Entertaining yet substantial. Great for anyone looking for a lighter entry point without sacrificing depth. The hosts break advanced concepts into everyday language, ideal for your commute or workout.

Free · Weekly · ~60 min

● GOING DEEPER

Latent Space

Alessio Fanelli & Swyx

My take: The podcast AI engineers actually listen to. Technically demanding but extremely current. If you want to understand where foundation models, agents, and AI infrastructure are heading, this is essential. Not for beginners, but indispensable if you need to understand the architecture layer.

Free · Weekly · ~60-90 min

Cognitive Revolution

Nathan Labenz

My take: Labenz interviews builders, researchers, and investors with a depth you usually only find at specialist conferences. Particularly valuable: his "AI Scouting Reports" that distill complex developments into clear takeaways. My go-to when I need a grounded assessment of a trend.

Free · Biweekly · ~90 min

Eye on AI

Craig S. Smith

My take: Smith comes from journalism (former NYT correspondent) and brings a perspective most tech podcasts lack: ethics, policy, regulation. For anyone who wants to understand AI not just as a tool, but as a societal phenomenon. Especially relevant for governance topics.

Free · Biweekly · ~30 min

No Priors

Sarah Guo & Elad Gil

My take: The investor perspective on AI. It shows where the money is flowing, which business models work, what's hype? Particularly useful for understanding how the market values AI products. Helps build stronger internal business cases for AI investment.

Free · Weekly · ~45 min

Books

Books remain the best way to build deep understanding of a topic. This selection provides both foundation and foresight.

● GETTING STARTED

Co-Intelligence: Living and Working with AI

Ethan Mollick · 2024

My take: The best entry book for anyone who wants to understand AI as a working tool and not an abstract technology. Mollick (Wharton professor) bridges practical application and strategic thinking. His concept of "co-intelligence", human and machine as a team, is the mental model I use in my own work. Required reading, but he tends to get a bit hyped about some topics.

Purchase · ~260 pages

AI Snake Oil: What AI Can Do, What It Can't, and How to Tell the Difference

Arvind Narayanan & Sayash Kapoor · 2024

My take: The antidote to AI hype. Narayanan and Kapoor (Princeton) systematically show where AI delivers real value and where the market massively overpromises. Particularly strong on "predictive AI" in domains like HR, criminal justice, and healthcare. Essential for contributing informed perspectives rather than recycling hype.

Purchase · ~300 pages

Nexus: A Brief History of Information Networks from the Stone Age to AI

Yuval Noah Harari · 2024

My take: Harari embeds AI in the grand history of information processing. Not technical, but philosophical and societal. Helps put the current AI wave into historical context. Worth reading for anyone who wants to understand why information networks shift power and what that means for organizations. I do not share some of his dystopian views.

Purchase · ~400 pages

● GOING DEEPER

Agentic Artificial Intelligence

Pascal Bornet, Jochen Wirtz, Tom Davenport et al. · 2025

My take: The most current reference on AI agents, the topic that will dominate the next 12-24 months. Bornet and team connect architecture concepts with real business applications. For anyone who wants to understand what "agentic AI" means beyond the buzzword and how it will reshape workflows.

Purchase · ~350 pages

Reshuffle: How AI Is Reshaping the Systems Behind Work and Life

Sangeet Paul Choudary · 2025

My take: Choudary takes a different path than most AI books: he discusses systems, not intelligence. How is AI changing the structures within which we work? Especially relevant for People & Organization, because the question isn't whether AI replaces jobs, but how AI restructures the systems in which jobs exist.

Purchase · ~280 pages

Empire of AI: A History of Artificial Intelligence

Karen Hao · 2025

My take: Investigative journalism on OpenAI and the power dynamics behind the AI industry. Hao (former MIT Technology Review) shows how a non-profit mission became a commercial powerhouse. Important for understanding the political and economic dimension of AI, not just the technical one.

Purchase · ~320 pages

Courses & Training

Structured learning for anyone who wants to build skills beyond podcasts and books. Good news: the world's best AI courses cost nothing. Stanford, Harvard, Google, Andrew Ng, Andrej Karpathy ... all free.

● GETTING STARTED

Elements of AI

University of Helsinki & MinnaLearn

My take: AI without math, without code. Over 2 million people have taken this course, for good reason. It breaks down fundamental AI concepts (machine learning, neural networks, Bayesian probability) into clear, interactive modules. If you want a solid foundation without any technical prerequisites, start here. Available in 26 languages.

Free · ~30 hours · Self-paced

→ elementsofai.com

AI for Everyone

Andrew Ng · Coursera / deeplearning.ai

My take: Andrew Ng is the best explainer in the AI world. This course is built for non-technical professionals and delivers a solid foundation in ~6 hours: what AI can do, what it can't, and how to think about it as an organization. No code, no math, pure understanding.

Free to audit · ~6 hours

→ [Course on Coursera](https://www.coursera.org/ai-for-everyone)

● GENERATIVE AI

Generative AI for Everyone

Andrew Ng · Coursera / deeplearning.ai

My take: The updated companion to "AI for Everyone," focused entirely on generative AI. Explains how LLMs work, what prompt engineering is, and where the real business applications lie. Especially useful for leaders who want to understand what's behind ChatGPT without needing to prompt themselves.

Free to audit · ~8 hours

→ [Course on Coursera](https://www.coursera.org/generative-ai-for-everyone)

Prompt Engineering for Developers

DeepLearning.AI & OpenAI

My take: The most practical short course on prompt engineering available. Built in collaboration with OpenAI, it teaches you how to write effective prompts not through guesswork, but through systematic principles. About one hour of focused content. Worth it for anyone who uses LLMs regularly and wants significantly better results.

Free · ~1 hour

→ [Course on DeepLearning.AI](#)

Intro to Large Language Models

Andrej Karpathy

My take: One video. One hour. And afterward, you genuinely understand what ChatGPT actually is, not the marketing version, the real version. Karpathy (former OpenAI, former Tesla AI lead) explains LLMs with a clarity that's hard to find elsewhere. Accessible enough for non-engineers, precise enough for technical people. The single best hour you can invest in understanding the technology.

Free · ~1 hour · YouTube

→ [Watch on YouTube](#)

● MACHINE LEARNING

Google Machine Learning Crash Course

Google

My take: 15 hours, completely free, now updated with LLM content. Google built this as internal training material and then made it public. Covers the fundamentals of ML with interactive exercises and real-world examples. The best option for anyone who wants a structured, hands-on ML foundation without committing to a multi-month course.

Free · ~15 hours

→ [Google ML Crash Course](#)

Machine Learning Specialization

Andrew Ng · Stanford / Coursera

My take: The gold standard. If you want to understand what's happening under the hood of AI systems, this is it. Three courses that explain ML from the ground up with hands-on exercises in Python. Demanding but worth it. Not necessary for everyone, but essential for anyone who wants to have technical conversations or build things themselves.

Free to audit · ~3 months · 9h/week

→ [Course on Coursera](#)

● ADVANCED**CS50's Introduction to AI with Python**

Harvard University (Brian Yu & David Malan)

My take: Harvard's legendary CS50, focused on AI. Covers search algorithms, knowledge representation, optimization, machine learning, and reinforcement learning all with hands-on Python projects. Rigorous but well-taught. For anyone who wants a proper computer science foundation in AI rather than just tool-level skills.

Free · ~7 weeks · edX

[→ CS50 AI](#)**An Introduction to Statistical Learning (ISLR)**

Stanford University (Hastie, Tibshirani et al.)

My take: The mathematical foundation behind machine learning. If you've ever wondered why a model makes the predictions it does, or how to properly evaluate model performance, this is where you build that understanding. The textbook is free online, and Stanford's companion course walks you through it. Not light reading, but the kind of depth that separates people who use AI from people who understand it. I did the course in 2015 and still profit from it.

Free · Textbook + video lectures

[→ statlearning.com](#)**Transformers & Large Language Models (Stanford CS25)**

Stanford University

My take: Under the hood of GPT. This Stanford course goes deep into transformer architecture the technology that powers every major LLM. Attention mechanisms, training procedures, fine-tuning, RLHF. Technically the most demanding course on this list, but if you want to truly understand how models like GPT-4 and Claude work at the architecture level, this is where you go.

Free (lectures online) · Graduate-level

[→ Stanford CS25](#)



Bonus: Stay Current

One newsletter that's worth more than most courses.

One Useful Thing (Substack)

Ethan Mollick

My take: Not a course in the traditional sense, but the single best ongoing learning channel for AI in a work context. Mollick constantly experiments with new AI tools and shares his findings as short, practical essays. Read this newsletter regularly and you'll be better informed than 90% of people who have "taken AI courses."

Free · Newsletter · 1-2x per week

→ oneusefulthing.org

Note: This guide reflects my state of knowledge as of March 2026. The AI field moves fast. If you've found something you think belongs here, let me know. No sales pitches, please. I did also not include my own course.

Curated by Daniel Mühlbauer-Kerber · AI & People Analytics Expert · March 2026