GenAI.works

The Global Free AI Learning Guide 2026

For Builders, Learners and Future AI Leaders

December 2025

GenAI Works Research

What This Guide Is

This guide is a roadmap for learning AI using only free, high-quality resources from universities and technology companies.

It brings together courses and programs from places like

Stanford · Massachusetts Institute of Technology · Harvard University · University of Oxford · Google · Microsoft · aws · IBM · Caltech · UC Berkeley · GenAI.academy

then organizes them by learning stage instead of by brand. The goal is to help you see where you are now and what to do next, rather than leave you with another long list of links.

You can use this guide if you are:

- starting from the basics,
- returning to technical work after a break,
- already working in tech and wanting to deepen your AI skills,
- or leading a team and planning how people should upskill.

How This Guide Works

The guide is structured around three stages of learning.

- **Beginner:** You're at the early stages of your AI journey and want a clear, low-pressure way in. This stage also fits if you're coming back to math or coding after some time away and need to rebuild confidence before tackling heavier material.

- **Intermediate:** You can already write basic code and want to understand what the models are doing and build projects that genuinely solve problems.

- **Advanced:** You're comfortable with the underlying math and programming and now want to work at a deeper level: studying research-style content, understanding architecture choices or designing AI systems end to end.

Each stage has:

- A short self-check so you can see where you fit
- A skills map for that stage
- A set of courses from Stanford, MIT, Harvard, Google, Microsoft, AWS, IBM, Caltech, Oxford, Berkeley, GenAI Academy, PwC and Deloitte, grouped under that stage

You can move through the stages in order or drop into the one that matches your current level.

Stage 1:
Beginner

Stage 1: Beginner

“I want to understand what AI is and start using it in my work or studies.”

Self-check: You’re probably here if...

- You have little or no formal background in AI
- You are okay with some technical ideas but prefer clear explanations
- You want to use AI tools and learn the basics before worrying about formulas
- You are a student, early-career professional, creator or founder who is exploring AI seriously for the first time

Skills to build at this stage

By the end of the Beginner stage, you should be able to:

- Explain in simple language what machine learning and generative AI are
- Recognize common AI use cases in your field
- Use AI tools (chat interfaces, content tools, research tools) with context-aware prompts
- Follow high-level explanations of models without getting lost
- Complete at least one small project or assignment that uses AI in a practical way

BEGINNER COURSES

Category: Conceptual & Ethics

UNIVERSITY OF OXFORD

Oxford Home Study: Machine Learning Fundamentals

Ideal if you prefer reading before coding.

- Covers supervised and unsupervised learning, neural networks, and basic concepts
- Places strong emphasis on AI ethics and the impact of algorithmic bias and the responsibilities of deploying intelligent systems.
- Completely self-paced and accessible with no prerequisites

Access it here

Oxford Home Study: AI & Machine Learning Series

- Short modules on topics like natural language processing, AGI ideas and tools like Claude
- It focuses on the business and operational side of AI. It explains what genAI tools do rather than just how to code them.
- Good for building vocabulary and a mental map of the field

Access it here

Category: Practical AI for Work

Google

Google AI Essentials

- About 10 hours
- Focus: how to use AI for productivity, communication and ethical decision-making
- Designed for non-engineers who want to bring AI into daily work
- Good early step if you feel unsure where to start

Access it here

Microsoft

Microsoft: Career Essentials in Generative AI

- Around 4 hours
- Explains how generative AI is changing work and why it matters
- Includes discussion of responsible use and career impacts
- Useful for managers, analysts and general knowledge workers

Access it here

Category: Practical AI for Work

GenAI.academy

GenAI Academy – AI Express: Intro to Generative AI

- Three live-style sessions of about one hour each, plus small assignments
- Helps you move from “I tried AI a few times” to “I can design prompts and content that support real tasks”
- You create: a promotion post, a stronger prompt and an initial agent-like concept
- Very good for people who learn best by doing short, focused challenges

Access it here

Caltech

Caltech CTME: AI Tools for Everyone

- About 8 hours, self-paced
- Shows you how to combine tools like ChatGPT, Perplexity, NotebookLM, Gamma and image generators
- Focus on workflows: research, summarization and presentation building
- Suitable for professionals who want to read a documents and produce usable reports instantly orchestrate tools

Access it here

Category: First Steps into Coding & Algorithms

HARVARD UNIVERSITY

Harvard: CS50's Introduction to AI with Python

- Seven weeks, with flexible timing
- The course is taught by Professor David J. Malan, a globally respected teacher known for making complex concepts clear and engaging.
- Assumes basic Python, then teaches search, simple machine learning, basic natural language processing and reinforcement learning
- You build small but memorable projects: a tic-tac-toe AI, a search engine and more
- The emphasis on computational thinking and storytelling, with David Malan's teaching style making algorithmic concepts stick.

Access it here

If you complete one conceptual course (Oxford or AI Essentials) and one practical course (GenAI Academy, Caltech, or CS50 AI), you can consider yourself through the Beginner stage.

Stage 2:
Intermediate

Stage 2: Intermediate

> *“I can follow basic AI concepts and want to understand and build models.”*

Self-check: You’re probably here if...

- You are comfortable with basic Python
- You know what terms like “training” and “overfitting” roughly mean
- You want to understand machine learning more formally
- You are ready to invest more time in structured coursework
- You want projects that can go into a portfolio or résumé

Skills to build at this stage

By the end of the Intermediate stage, you should be able to:

- Describe key machine learning algorithms and when to use them
- Work with training and validation sets
- Implement models in code using popular libraries
- Understand the basics of deep learning and large language models
- Build a few end-to-end projects: for example, a recommendation system or a small search engine

Stanford

Andrew Ng: Machine Learning (Stanford University)

- About 11 weeks, 5–8 hours per week
- Explains supervised and unsupervised learning, basic neural networks and core algorithms
- This course breaks down how machines learn from data, from predicting outcomes to spotting patterns humans might miss.
- Balances intuition with implementation, which makes it a common standard in industry
- Widely recognized in interviews and hiring processes

Access it here

HARVARD UNIVERSITY

Harvard – Data Science: Building Machine Learning Models

- Eight weeks, estimated 2–4 hours per week
- Build a recommendation system that helps you to understand the core mechanics used by companies like Netflix and Amazon to personalize content and drive significant revenue.
- Focus on regression, decision trees, cross-validation, regularization and building a movie recommendation system.
- Very suitable if you want a mix of math, code and business relevance

Access it here

Category: Machine Learning Foundations

Massachusetts Institute of Technology

MIT – 6.034: Artificial Intelligence

- Semester-equivalent, self-paced
- Learn the logical foundations that enable machines to reason and solve problems.
- Covers search, knowledge representation, learning and computer vision
- Patrick Winston's lectures explain the "why" of AI systems as much as the "how"
- Good bridge between basic ML courses and more advanced material

Access it here

Category: Deep Learning and Modern Architectures

Massachusetts Institute of Technology

MIT – 6.S191: Introduction to Deep Learning

- Intensive short course (often completed over a few days, self-paced online)
- Covers the exact technologies driving the current AI market. You learn the specific architectures behind tools like ChatGPT and Midjourney.
- Introduces convolutional and recurrent networks, transformers and generative models
- Includes implementation labs using PyTorch or TensorFlow
- Strong choice if you want a compact introduction to deep learning

Access it here

Microsoft

Microsoft: AI for Beginners (GitHub Curriculum)

- Twelve weeks, 24 lessons
- Covers neural networks, natural language processing and computer vision
- Includes sketchnotes and exercises that make technical ideas more approachable
- Suits learners who prefer a structured, open-source path

Access it here

Category: LLMs in Practice

We avoid heavy jargon here, but this is where you start to understand the systems behind modern language models.

Google

Google Cloud Skills Boost: Introduction to Large Language Models

- About one hour
- Explains how LLMs are trained and used
- Introduces tuning concepts and the Vertex AI environment
- Good first technical look at this type of model

Access it here

Google Cloud Skills Boost: Prompt Engineering and Gemini Code Assist

- 5–10 hours of labs
- Teaches you to use AI for code generation, debugging and structured problem-solving
- Includes practical exercises in a live cloud environment

Access it here

Microsoft

Microsoft: Azure OpenAI Service Modules

- Short, modular lessons
- Show how to build GPT-4 powered copilots, implement RAG, ground AI in enterprise data
- Useful for engineers and technical product people who want to see how these systems fit into real infrastructure

Access it here

AWS Skill Builder: AI Ready / GenAI Foundations

- 135+ AI/ML courses under the “AI Ready” initiative.
- Includes introductions to Amazon Bedrock and working with multiple model providers
- Good option if you want exposure to cloud-native AI patterns

Access it here

Category: First Steps into Agents and Workflows

GenAI.academy

GenAI Academy: Intro to Agentic AI

- Three sessions
- Explains what makes a system “agentic” in practical terms
- Shows how to design simple multi-step flows and when to involve human review
- A good way to move from single prompts to more structured interactions

At this stage, you begin to connect classic ML, deep learning and modern language systems into a coherent picture. Once you have several projects and labs under your belt, you are ready to consider advanced study.

Access it here

Stage 3:
Advanced

Stage 3: Advanced

> *“I want to design AI systems, go deep on architecture or work near research.”*

Self-check: You’re probably here if…

- You are comfortable with linear algebra, probability and basic statistics
- You can read research-style explanations without getting overwhelmed
- You have already completed at least one ML course and a few projects
- You are considering roles in AI engineering, research, or technical leadership

Skills to build at this stage

By the end of the Advanced stage, you should be able to:

- Derive or at least follow the derivation of core algorithms
- Design and evaluate complex models and training setups
- Understand trade-offs in architecture choices
- Work with advanced topics such as modern language architectures, agent designs and evaluation
- Place technical work in the context of governance, ethics and strategy

Category: Deep Theory and Architecture

Massachusetts Institute of Technology

MIT – 6.036: Introduction to Machine Learning

- Self-paced, semester-level depth
- Focuses on math required to understand the "black box" of machine learning, underpinnings of linear models, neural networks and regularization
- Uses problem sets and exams that match MIT's on-campus expectations
- Strong base if you want to read or contribute to technical work
- The curriculum focuses heavily on the mechanics of algorithms. It prepares you for advanced research roles where understanding the underlying math is mandatory.

Access it here

Stanford

Stanford – CS221: Introduction to AI

- 10–15 weeks
- Focus on search, logic, planning and decision-making under uncertainty
- Provides the conceptual tools used in areas like robotics and advanced software agents
- Best tackled with prior experience in math and programming

Access it here

Stanford

Stanford – CS224N: Natural Language Processing with Deep Learning

- Self-paced lectures and assignments
- Covers word embeddings, recurrent networks, attention and modern language architectures
- Often considered a reference course for understanding how current text models work
- Suitable if you already have deep learning basics

Access it here

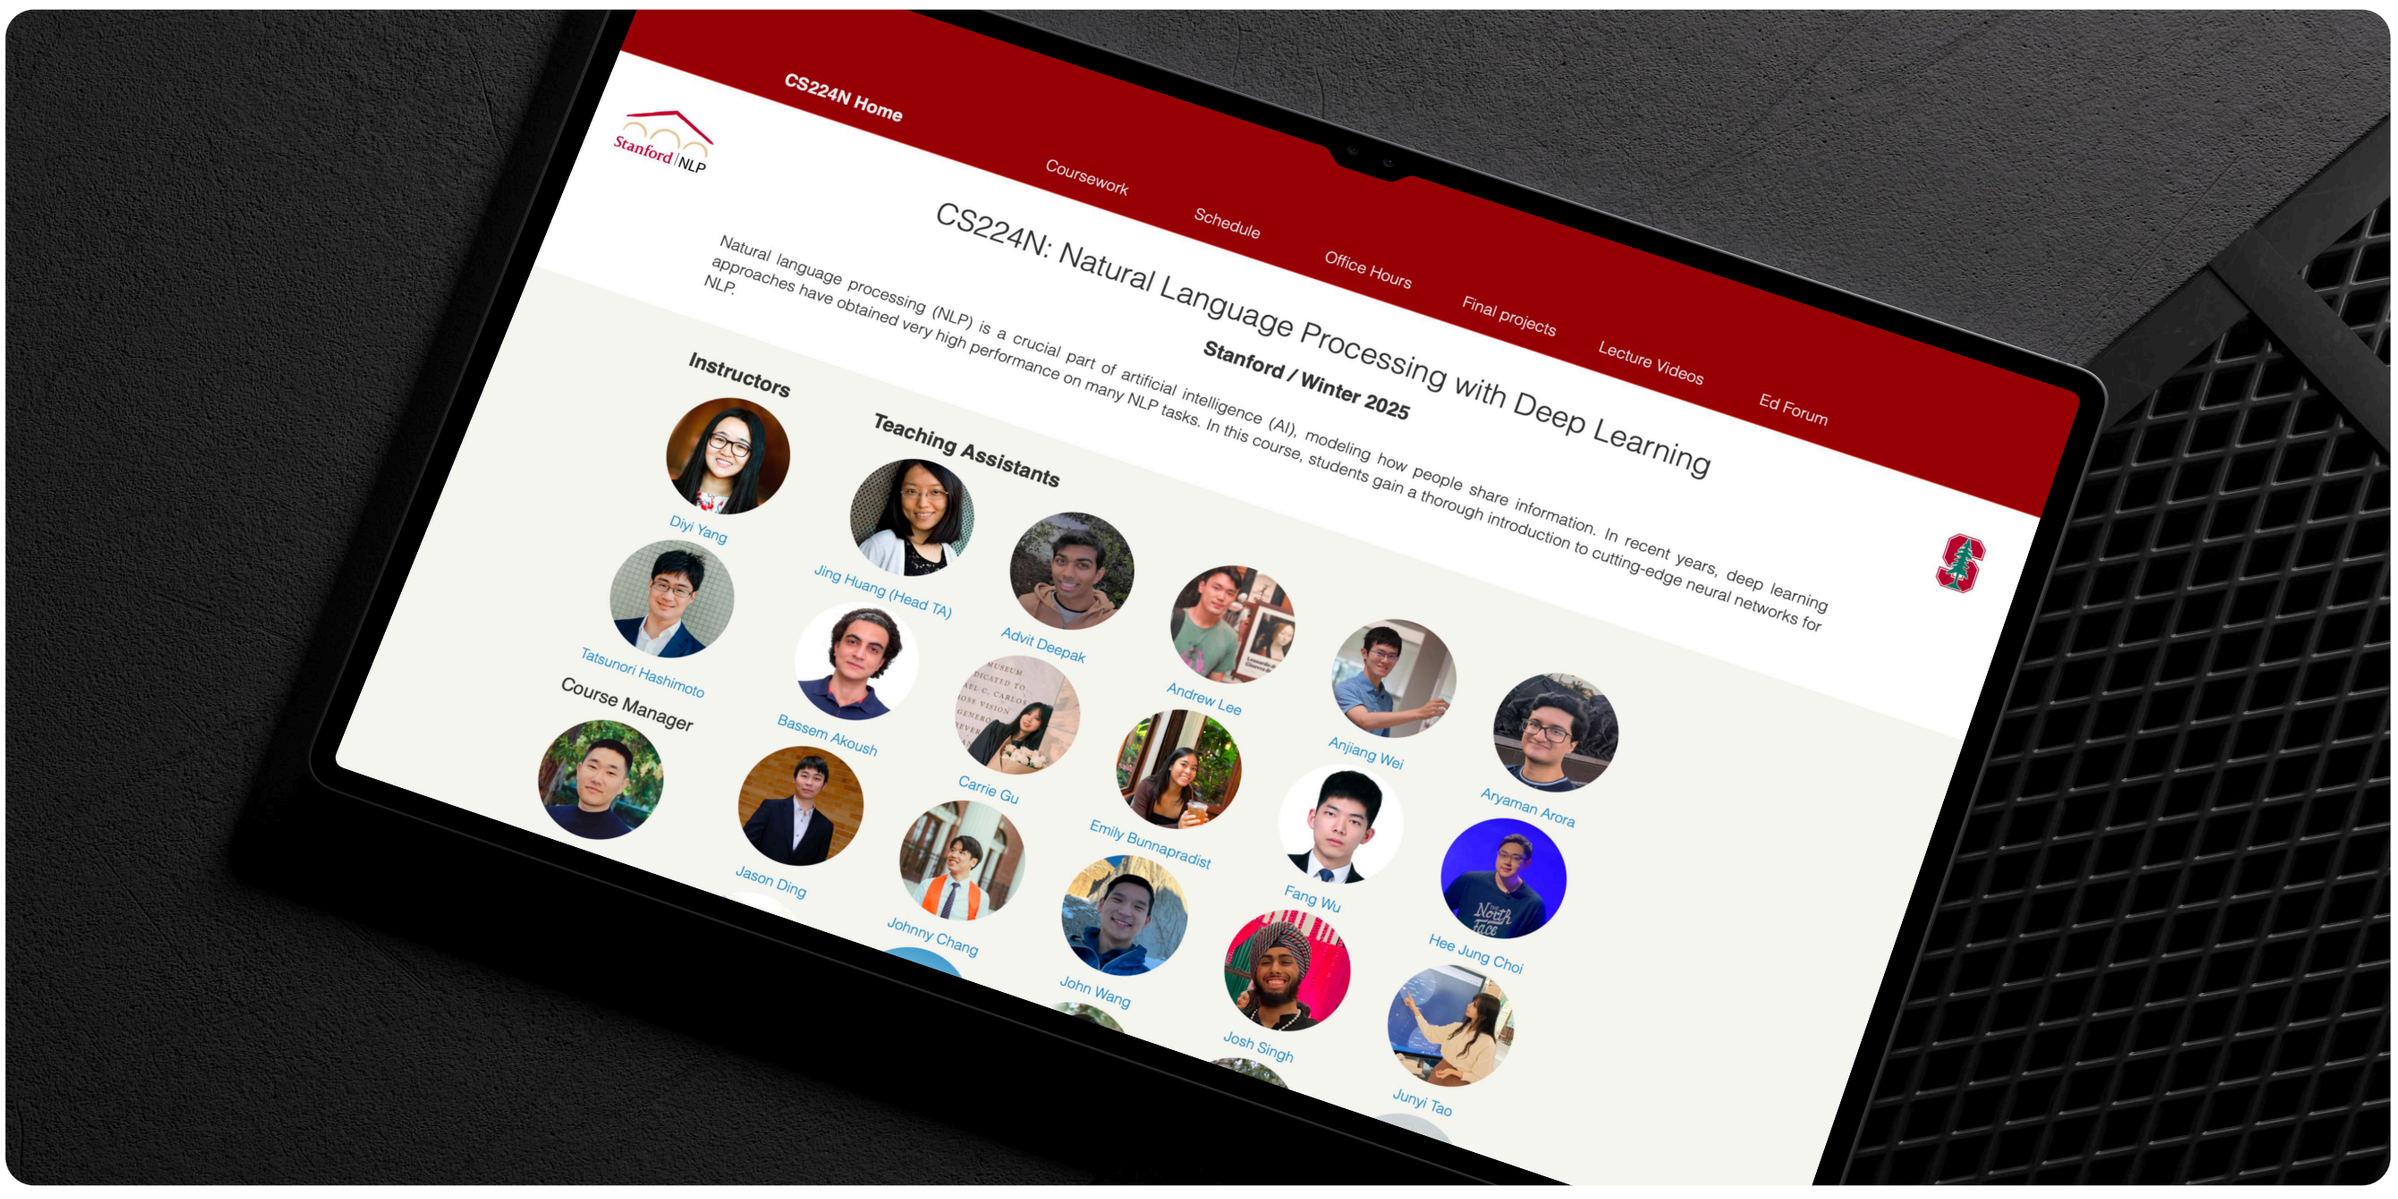

Category: Deep Learning at Scale

Massachusetts Institute of Technology

MIT: 6.S191 (as a second pass)

- Rewatch or redo labs with more attention to implementation detail and efficiency
- Focus more on optimization, practical tricks and model evaluation

Access it here

Google Microsoft

Google Cloud and Microsoft: Advanced Labs

- At this level, the most relevant content is not a single “course,” but the advanced portions of existing labs and documentation
- Examples: building more complex text applications, connecting models to enterprise data sources, or optimizing for latency and cost

UC Berkeley

UC Berkeley: Agentic AI Massive Open Online Courses

- Self-paced
- Introduces LLM agents, planning algorithms, multi-agent systems and real-world tool use.
- Guest lectures often include research leaders from Google DeepMind, OpenAI and Anthropic.
- The content is strongly influenced by the “AGI labs” view of the future, focusing on autonomy, coordination and reasoning.

Access it here

GenAI.academy

GenAI Academy: Vibe Sessions (Biz / Web) at a deeper level

- Revisited at this stage, these sessions become a space to experiment with more structured workflows and agent-like patterns in your own context

Access it here

IBM watsonx Foundations

- Introduces trusted AI, specialized models and governance considerations
- Focus on “Granite” models and governance
- Useful if you work in regulated industries or large organizations
- Emphasis on tracking data lineage, auditability and hybrid cloud.

Access it here

PwC Academy & Deloitte AI Institute

- Provide executive-level material on responsible AI, workforce transformation and industry-specific applications
- Important reading if you plan to influence policy, strategy or change management in your company

Access it here

Access it here

At the Advanced stage, your learning becomes more self-directed. You use these courses and resources as reference points while you work on more ambitious projects, research ideas or system designs.

Putting It All Together

Instead of thinking *"What is the best AI course?"*, it is more helpful to ask:

- **Where am I right now?**
- **Which skills do I actually need next?**

Then pick one or two courses from the stage that matches you and finish them fully, including the assignments and projects.

A simple way to use this guide:

- If you are new or returning after a long break, start in Beginner
- If you have coded before and want to understand models, start in Intermediate
- If you have already done several ML projects and want depth, move into Advanced

You can always loop back: a beginner who completes AI Express and CS50 AI might take a pause before moving to Stanford or MIT. An intermediate learner might add Deloitte or PwC content to understand how their technical work lands in the business.

The aim is not to rush through everything. The aim is to build a stack of skills that makes you effective in the real world: a clear mental model of how AI works, backed by enough hands-on practice to use it confidently in real work

This guide gives you the structure. The rest are small, consistent steps.

If you want to stay current on AI learning and real-world developments, follow us on our social channels and subscribe to the newsletter below.

Newsletter	https://newsletter.genai.works/
Website	https://genai.works/
Youtube	https://youtube.com/@GenAIWorks
Instagram	https://www.instagram.com/generativeai_official/
X	https://x.com/GenAI_Now
Medium	https://medium.com/@genai.works

GenAI.works