

GUIDEXIA LEARNING

AI ENGINEER

COMPLETE ROADMAP

Week-by-Week | 8-9 Months | Zero to AI Engineer | All Free Resources

36

Weeks

12

Phases

300+

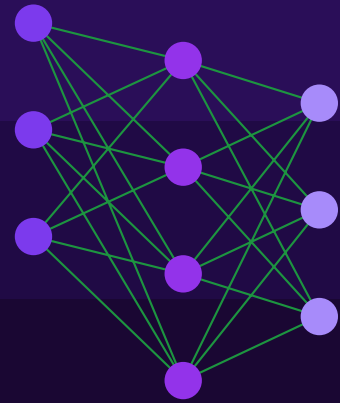
Skills

8-9

Months

100%

Free



Follow us: guidexia.com

This is your complete, week-by-week AI Engineer roadmap spanning 8 to 9 months — 12 structured phases and 36 weeks of deliberate learning. It covers everything from Python fundamentals to fine-tuning LLMs, building RAG systems, deploying AI to production, and landing your first AI Engineer role. Every topic is practical, every week has a clear deliverable, and every resource listed is 100% free.

MONTH-BY-MONTH OVERVIEW — 9 MONTHS TO AI ENGINEER

Month	Focus Area	Phases Covered	Key Deliverable
Month 1	Python for AI + Mathematics	Phase 1-2	Jupyter notebooks with NumPy/Pandas projects
Month 2	Classical Machine Learning	Phase 3	Scikit-learn pipeline on Kaggle dataset
Month 3	Deep Learning with PyTorch	Phase 4	Trained CNN on custom image dataset
Month 4	NLP + Computer Vision	Phase 5-6	Fine-tuned BERT + YOLO object detector
Month 5	LLMs + Prompt Eng + RAG	Phase 7	Production RAG chatbot with Gradio UI
Month 6	Fine-tuning + MLOps	Phase 8-9	Fine-tuned Llama model deployed to cloud API
Month 7	Specialisation + AI Apps	Phase 10	Streaming AI app with Next.js + FastAPI
Month 8	Research Skills + Paper Impl.	Phase 11	Reproduced a transformer paper from scratch
Month 9	Capstone + Portfolio + Job Readiness	Phase 12	Live AI product + polished portfolio + offer

PRE-REQUISITES

A laptop, stable internet, and 2-3 hours/day. Basic Python helps but is not required. All tools, platforms and courses referenced in this roadmap are completely free.

1

Python for AI Engineers

Weeks 1–3 (Month 1)

PHASE 01

Python is the primary language of AI. Unlike general Python courses, this phase focuses on the Python patterns AI engineers use daily — vectorised code, functional pipelines, and data structures built for scale.

WEEK Weeks 1-2 — Python Core + Scientific Stack

Python AI Essentials	NumPy Deep Dive
<ul style="list-style-type: none"> - Data types, lists, dicts - Functions & lambda - Comprehensions & generators - Itertools & functools 	<ul style="list-style-type: none"> - ndarray & dtypes - Vectorised operations - Broadcasting rules - Array slicing & masking

WEEK Week 3 — Pandas, Plotting & Environment

Pandas for Data	Visualisation & Env
<ul style="list-style-type: none"> - DataFrame CRUD - GroupBy & aggregation - Merge & pivot tables - Missing value strategies 	<ul style="list-style-type: none"> - Matplotlib & Seaborn - Jupyter notebooks - Conda environments - Google Colab & GPU runtimes

SKILLS YOU WILL GAIN IN THIS PHASE

Python	NumPy	Pandas	Matplotlib	Seaborn
Jupyter	Colab	Conda	Vectorisation	Broadcasting
Generators	Functional Python			



Phase Outcome:

You can write efficient, vectorised Python code for data manipulation and set up a full AI development environment.

You don't need a PhD, but you do need the maths. This targeted phase covers exactly the linear algebra, calculus and statistics that AI engineers encounter every single day.

WEEK Week 4 — Linear Algebra — The Language of Data

Vectors & Matrices	Decompositions
<ul style="list-style-type: none"> - Dot products & norms - Matrix multiply & transpose - Inverse & determinants - Geometric intuition 	<ul style="list-style-type: none"> - Eigenvalues & eigenvectors - SVD — what it means in ML - PCA preview - Covariance matrices

WEEK Week 5 — Calculus & Optimisation

Calculus for AI	Optimisation
<ul style="list-style-type: none"> - Derivatives & chain rule - Partial derivatives - Gradient vectors - Jacobians & Hessians 	<ul style="list-style-type: none"> - Gradient descent intuitively - Learning rate effects - Convex vs non-convex - Loss surface visualisation

WEEK Week 6 — Probability & Statistics

Probability	Statistics for AI
<ul style="list-style-type: none"> - Random variables & distributions - Bayes theorem — deep intuition - MLE & MAP estimation - Entropy & information 	<ul style="list-style-type: none"> - Central limit theorem - Hypothesis testing - A/B testing basics - Statistical significance

SKILLS YOU WILL GAIN IN THIS PHASE

Linear Algebra	Dot Products	SVD	Eigenvalues	PCA
Derivatives	Gradient Descent	Chain Rule	Probability	Bayes
MLE	Entropy	A/B Testing		



Phase Outcome:

You can read ML papers, understand algorithm derivations, and explain why models work at a mathematical level.

Classical Machine Learning

Weeks 7–10 (Month 2)

PHASE 03

Before you prompt GPT, you must understand the models that came before it. Classical ML is the foundation every AI engineer is expected to know cold in interviews and on the job.

WEEK Week 7 — Supervised Learning I

Regression Models	Classification Models
<ul style="list-style-type: none"> - Linear & polynomial regression - Regularisation: Ridge, Lasso - Feature scaling: StandardScaler - Metrics: MSE, MAE, R-squared 	<ul style="list-style-type: none"> - Logistic regression - Decision trees & pruning - Naive Bayes classifier - KNN algorithm

WEEK Week 8 — Supervised Learning II — Ensembles

Tree Ensembles	SVMs & Kernels
<ul style="list-style-type: none"> - Random Forests — intuition - Gradient Boosting — how it works - XGBoost deep dive - LightGBM for speed 	<ul style="list-style-type: none"> - Margin maximisation - Kernel trick (RBF, poly) - SVM for classification - When to choose SVM

WEEK Week 9 — Unsupervised & Semi-supervised


Clustering	Dimensionality Reduction
<ul style="list-style-type: none"> - K-Means & elbow method - DBSCAN for noisy data - Hierarchical clustering - Gaussian Mixture Models 	<ul style="list-style-type: none"> - PCA from scratch - t-SNE for visualisation - UMAP — faster alternative - Autoencoders preview

WEEK Week 10 — Model Evaluation & Pipelines

Evaluation Mastery	Scikit-learn Pipelines
<ul style="list-style-type: none"> - Train/val/test split - Cross-validation (k-fold, stratified) - Precision, recall, F1, AUC-ROC - Confusion matrix analysis 	<ul style="list-style-type: none"> - Pipeline object - ColumnTransformer - Feature engineering - GridSearchCV & Optuna

SKILLS YOU WILL GAIN IN THIS PHASE

Linear Regression	Logistic Regression	Decision Trees	Random Forest	XGBoost
LightGBM	SVM	K-Means	DBSCAN	PCA
t-SNE	UMAP	Cross-Validation	AUC-ROC	Scikit-learn
Optuna				

 **Phase Outcome:** You can train, evaluate and explain any classical ML model, build Scikit-learn pipelines, and ace ML interview questions.

Deep Learning Fundamentals

Weeks 11–14 (Month 3)

PHASE 04

Neural networks power everything from image recognition to large language models. Understanding them from the ground up — not just as black boxes — is what separates AI engineers from prompt engineers.

WEEK Week 11 — Neural Networks from Scratch

Perceptrons & Activation	Backpropagation
<ul style="list-style-type: none"> - Perceptron model - Forward pass - Activation functions (ReLU, sigmoid, tanh, GELU) - Universal approximation theorem 	<ul style="list-style-type: none"> - Chain rule in networks - Computational graphs - Vanishing gradient problem - Weight initialisation (He, Xavier)

WEEK Week 12 — PyTorch Core

PyTorch Basics	Data Loading
<ul style="list-style-type: none"> - Tensors & autograd - Building nn.Module - Optimisers (Adam, SGD, AdamW) - Training loop from scratch 	<ul style="list-style-type: none"> - Dataset & DataLoader classes - Custom datasets - Data augmentation - GPU training (.cuda())

WEEK Week 13 — CNN — Computer Vision Intro


CNN Architecture	Transfer Learning
<ul style="list-style-type: none"> - Conv layers, filters, stride - Pooling & feature maps - Receptive field concept - BatchNorm & Dropout 	<ul style="list-style-type: none"> - Pretrained models (ResNet, VGG) - Fine-tuning strategies - Feature extraction - timm library

WEEK Week 14 — RNN & Sequential Models

RNNs & LSTMs	Training Deep Nets
<ul style="list-style-type: none"> - Recurrent cell mechanics - LSTM gates (forget, input, output) - GRU — lighter alternative - Sequence-to-sequence 	<ul style="list-style-type: none"> - Learning rate schedules - Gradient clipping - Early stopping - Experiment tracking (W&B;)

SKILLS YOU WILL GAIN IN THIS PHASE

PyTorch	Neural Networks	Backpropagation	Autograd	CNNs
Convolutions	BatchNorm	Dropout	Transfer Learning	ResNet
timm	RNNs	LSTMs	GRU	Adam
W&B;	GPU Training			

 **Phase Outcome:** You can build, train and debug neural networks in PyTorch from scratch and apply transfer learning to real datasets.

5

Natural Language Processing

Weeks 15–18 (Month 4)

PHASE 05

Language is the primary interface of modern AI. This phase takes you from bag-of-words all the way to transformer fine-tuning — the most in-demand AI skill in 2024-2025.

WEEK Week 15 — Classical NLP

Text Preprocessing	Feature Extraction
<ul style="list-style-type: none"> - Tokenisation (word, subword, BPE) - Stemming & lemmatisation - Stop words & punctuation - Regex for text cleaning 	<ul style="list-style-type: none"> - Bag-of-Words & TF-IDF - N-grams model - Word2Vec & GloVe - FastText embeddings

WEEK Week 16 — Transformers — Architecture Deep Dive

Attention Mechanism	Transformer Architecture
<ul style="list-style-type: none"> - Self-attention from scratch - Multi-head attention - Positional encodings - Scaled dot-product attention 	<ul style="list-style-type: none"> - Encoder-decoder structure - Encoder-only (BERT-style) - Decoder-only (GPT-style) - Layer norm & residuals

WEEK Week 17 — Hugging Face Ecosystem

HF Transformers	Fine-tuning BERT
<ul style="list-style-type: none"> - from_pretrained API - Tokenizer deep dive - Pipeline API for quick inference - AutoModel & AutoTokenizer 	<ul style="list-style-type: none"> - Text classification task - Trainer API - Custom training loop - Evaluation & metrics (seqeval)

WEEK Week 18 — NLP Projects & Applications

NLP Tasks	Deployment
<ul style="list-style-type: none"> - Sentiment analysis - Named Entity Recognition (NER) - Question answering - Text summarisation (T5/BART) 	<ul style="list-style-type: none"> - FastAPI NLP endpoint - HF Spaces deployment - Gradio demo UI - Model quantisation (INT8)

SKILLS YOU WILL GAIN IN THIS PHASE

Tokenisation	BPE	TF-IDF	Word2Vec	GloVe
Transformers	Self-Attention	Multi-Head Attention	BERT	GPT
Hugging Face	Fine-tuning	Trainer API	NER	Summarisation
Gradio	Quantisation			



Phase Outcome:

You can fine-tune transformer models for any NLP task and deploy them as production-ready APIs with a demo UI.

Vision is the second major modality of AI. Modern AI engineers work with multimodal systems that combine images and text — from CLIP to GPT-4V. This phase gives you that capability.

WEEK Week 19 — Computer Vision Deep Dive

Object Detection	Image Segmentation
<ul style="list-style-type: none"> - YOLO architecture (v8) - Anchor boxes & IoU - mAP metric explained - Roboflow datasets 	<ul style="list-style-type: none"> - Semantic vs instance segmentation - SAM (Segment Anything) - DeepLab & FCN - Masks & bounding boxes

WEEK Week 20 — Vision Transformers & CLIP

Vision Transformers	CLIP & Contrastive Learning
<ul style="list-style-type: none"> - ViT — image as patches - DeiT & Swin Transformer - Image classification with ViT - Attention maps visualisation 	<ul style="list-style-type: none"> - CLIP architecture & training - Zero-shot image classification - Contrastive loss - ALIGN & SigLIP

WEEK Week 21 — Generative Vision Models

Diffusion Models	Multimodal AI
<ul style="list-style-type: none"> - DDPM from intuition - Stable Diffusion architecture - DALL-E 3 & Midjourney - ControlNet & LoRA 	<ul style="list-style-type: none"> - LLaVA & GPT-4V APIs - Image + text pipelines - Video understanding basics - Practical multimodal projects

SKILLS YOU WILL GAIN IN THIS PHASE

YOLO	Object Detection	SAM	Segmentation	ViT
CLIP	Contrastive Learning	Stable Diffusion	DDPM	ControlNet
LoRA	LLaVA	GPT-4V	Multimodal	Roboflow

LLMs & Prompt Engineering

Weeks 22–25 (Month 5)

PHASE 07

Large Language Models are the defining technology of this decade. This phase teaches you to work with them deeply — not just call an API, but understand, optimise, and build systems with them.

WEEK Week 22 — LLM Foundations

How LLMs Work

- Pretraining & next-token prediction
- RLHF & instruction tuning
- Scaling laws (Chinchilla)
- Context window & attention

LLM Landscape

- GPT-4, Claude, Gemini, Llama
- Open vs closed models
- Choosing the right model
- API rate limits & costs

WEEK Week 23 — Prompt Engineering Mastery

Core Techniques

- Zero-shot, one-shot, few-shot
- Chain-of-thought (CoT)
- Tree-of-thought (ToT)
- ReAct prompting

Advanced Prompting

- System prompts & personas
- Prompt injection & defence
- Structured output (JSON mode)
- DSPy framework

WEEK Week 24 — RAG — Retrieval Augmented Generation

RAG Architecture

- Chunking strategies
- Embedding models (OpenAI, E5)
- Vector databases (Pinecone, Chroma, Weaviate)
- Retrieval: BM25 + dense hybrid

RAG Optimisation


- Re-ranking (Cohere, cross-encoder)
- HyDE — hypothetical document
- Multi-query retrieval
- RAG evaluation (RAGAS)

WEEK Week 25 — LLM Agents & Tool Use

AI Agents	Frameworks
<ul style="list-style-type: none"> - ReAct agent architecture - Tool calling & function calling - Planning & memory - Multi-agent systems 	<ul style="list-style-type: none"> - LangChain chains & agents - LlamaIndex for data - CrewAI multi-agent - AutoGen & custom agents

SKILLS YOU WILL GAIN IN THIS PHASE

LLMs	GPT-4	Claude	Llama	Prompt Engineering
CoT	Few-Shot	RAG	Vector DBs	Pinecone
Chroma	Embeddings	Re-ranking	RAGAS	LangChain
LlamaIndex	AI Agents	Function Calling	CrewAI	DSPy





8

Fine-tuning & Model Customisation

Weeks 26–28 (Month 6)

PHASE 08

Calling an API is level 1. Fine-tuning is level 2. This phase teaches you to customise foundation models for specific tasks, domains, and behaviours — the premium AI engineering skill.

WEEK Week 26 — Fine-tuning Fundamentals

Full Fine-tuning	Efficient Fine-tuning
<ul style="list-style-type: none"> - When to fine-tune vs RAG - Dataset preparation (alpaca format) for fine-tuning - Instruction tuning datasets 	<ul style="list-style-type: none"> - LoRA — intuition & math - QLoRA (4-bit quantisation) - PEFT library - Rank & alpha selection

WEEK Week 27 — Advanced Fine-tuning & Alignment

Training Techniques	Alignment Methods
<ul style="list-style-type: none"> - Unsloth for fast training - TRL library - SFT Trainer - Multi-GPU training with DeepSpeed 	<ul style="list-style-type: none"> - RLHF overview - DPO — Direct Preference Optimisation - Constitutional AI - Reward models

WEEK Week 28 — Evaluation & Model Management

LLM Evaluation	Model Hub
<ul style="list-style-type: none"> - Automated: BLEU, ROUGE, BERTScore - LLM-as-a-judge (GPT-4 eval) - Benchmark suites (MMLU, HellaSwag) - Human eval design 	<ul style="list-style-type: none"> - Hugging Face Hub publish - Model cards best practice - GGUF & llama.cpp - Ollama for local models

SKILLS YOU WILL GAIN IN THIS PHASE

Fine-tuning	LoRA	QLoRA	PEFT	Unsloth
TRL	DPO	RLHF	Constitutional AI	Reward Models
BERTScore	LLM-as-judge	MMLU	Ollama	llama.cpp
GGUF	Model Cards	DeepSpeed		



Phase Outcome:

You can fine-tune open-source LLMs efficiently, align them to desired behaviour, and evaluate them rigorously.

Building a model that works in a notebook is not engineering. MLOps bridges the gap between AI research and production — this is what makes you a real AI engineer, not just an AI hobbyist.

WEEK Week 29 — ML Pipeline Engineering

Experiment Tracking	Feature Stores & Pipelines
<ul style="list-style-type: none"> - MLflow — runs, metrics, artifacts - Weights & Biases sweeps - DVC for data versioning - Reproducible experiments 	<ul style="list-style-type: none"> - Feature stores (Feast, Tecton) - Airflow / Prefect pipelines - Data validation (Great Expectations) - Schema drift detection

WEEK Week 30 — Model Serving & APIs

API Design for AI	Serving Frameworks
<ul style="list-style-type: none"> - FastAPI async endpoints - Streaming responses (SSE) - Batch inference endpoints - OpenAI-compatible API format 	<ul style="list-style-type: none"> - vLLM for LLM serving - TorchServe - Triton Inference Server - BentoML packaging

WEEK Week 31 — Cloud Deployment & Monitoring

Cloud AI Platforms	AI Monitoring
<ul style="list-style-type: none"> - AWS SageMaker end-to-end - GCP Vertex AI - Azure ML - Modal & Replicate (serverless GPU) 	<ul style="list-style-type: none"> - Data drift (Evidently AI) - Model performance degradation - Prompt logging & tracing - LangSmith for LLM ops

SKILLS YOU WILL GAIN IN THIS PHASE

MLflow	W&B;	DVC	Airflow	Prefect
FastAPI	Streaming	vLLM	TorchServe	Triton
BentoML	SageMaker	Vertex AI	Modal	Replicate
Evidently AI	LangSmith	LLM Ops		



Phase Outcome:

You can deploy AI models to production, serve them at scale with proper monitoring, and maintain them over time.



10

Specialisation — Choose Your AI Path

Weeks 32–33 (Month 7)

PHASE 10

AI engineering has multiple career tracks. In these two weeks you go deep on your chosen specialisation while also exploring the frontier of what is being built today in industry.

WEEK Week 32 — AI Application Development

Full-Stack AI Apps	Voice & Audio AI
<ul style="list-style-type: none"> - Next.js + AI SDK (Vercel) - Streaming chat UI - Supabase + pgvector - Auth for AI apps 	<ul style="list-style-type: none"> - Whisper ASR (speech-to-text) - TTS: ElevenLabs & Coqui - Real-time voice pipelines - Audio preprocessing

WEEK Week 33 — AI at Scale & Frontier Topics

Scaling AI Systems	Emerging AI
<ul style="list-style-type: none"> - Mixture of Experts (MoE) - Speculative decoding - KV-cache optimisation - Quantisation: GPTQ, AWQ 	<ul style="list-style-type: none"> - AI coding agents (Devin-like) - AI for robotics (RT-2) - Multimodal agents - AI safety & alignment basics

SKILLS YOU WILL GAIN IN THIS PHASE

Next.js AI SDK	Streaming Chat	pgvector	Supabase	Whisper
TTS	ElevenLabs	MoE	Speculative Decoding	GPTQ
AWQ	KV-Cache	AI Agents	Robotics AI	AI Safety



Phase Outcome:

You have a deep specialisation in your chosen AI track and understand the frontier of what is being built in industry.



11

Research Skills & Reading Papers

Weeks 34–35 (Month 8)

PHASE 11



The best AI engineers can read papers and implement ideas from scratch. This phase teaches you to consume the research frontier and stay current in a field that moves faster than any textbook.

WEEK Week 34 — Reading & Implementing Papers

Paper Reading Strategy	Implementation Practice
<ul style="list-style-type: none"> - Anatomy of an ML paper - How to read in 3 passes - ArXiv navigation & alerts - Papers with Code resource 	<ul style="list-style-type: none"> - Implement: Attention is All You Need - Implement: LoRA paper - Git-based paper replication - Ablation studies

WEEK Week 35 — Contributing to AI Research

Research Tools	Open Source AI
<ul style="list-style-type: none"> - LaTeX basics for AI - wandb Reports for sharing - Writing experiment summaries - Reproducing SOTA results 	<ul style="list-style-type: none"> - Contributing to HuggingFace - Open LLM Leaderboard - Submitting to AI conferences - Building research reputation

SKILLS YOU WILL GAIN IN THIS PHASE

ArXiv	Paper Reading	Attention Paper	LoRA Implementation	Papers with Code
LaTeX	W&B; Reports	Reproducing SOTA	Open Source AI	HuggingFace
Research Writing				Contribution



Phase Outcome:

You can read any AI research paper, implement it from scratch, and contribute meaningfully to the open-source AI community.



12

Portfolio, Capstone & Job Readiness

Week 36 (Month 9)

PHASE 12



Everything you have built across 35 weeks now becomes your story. This final phase packages your expertise into a portfolio, resume and interview preparation that lands AI engineering roles.

WEEK Week 36 Part A — Capstone AI Project & Portfolio

Capstone Project	Portfolio Assets
<ul style="list-style-type: none"> - Full end-to-end AI application - Fine-tuned model + RAG + Agent - Deployed API + Gradio/Next.js UI - HuggingFace Space + GitHub 	<ul style="list-style-type: none"> - AI portfolio website - 3 case-study writeups (Medium) - YouTube demo walkthrough - GitHub with clean READMEs

WEEK Week 36 Part B — Resume, Interviews & Job Search

Resume & LinkedIn	Interview Mastery
<ul style="list-style-type: none"> - AI-specific resume format - Quantify: 'Reduced latency by 60%' - LinkedIn: AI headline + posts - 5 technical recommendations 	<ul style="list-style-type: none"> - ML fundamentals questions - LLM & RAG system design - Live coding (NeetCode AI list) - Behavioural: STAR + impact stories

SKILLS YOU WILL GAIN IN THIS PHASE

Capstone	End-to-End AI	HuggingFace Spaces	Portfolio	Case Studies
Technical Writing	Resume	LinkedIn	System Design	ML Interview
LLM Interview	RAG Interview	Job Search	Negotiation	



Phase Outcome:

You are a job-ready AI Engineer with a live capstone project, published portfolio, and full interview preparation to land top AI roles.

SUCCESS STRATEGIES & FREE RESOURCE DIRECTORY

Build Every Week Without Exception	Every week must produce a commit, a notebook, or a deployed demo. Consumption without creation is the #1 failure mode.
Understand, Don't Just Copy	When you use a library, look at the source. When you run a script, understand each line. Depth compounds.
Read One Paper Per Week from Month 3	Start with 'Attention Is All You Need'. Then BERT, then GPT-2. Papers build intuition no tutorial can give you.
Join the AI Community Publicly	Post your projects on X/Twitter, LinkedIn, HuggingFace Spaces. The AI community is generous — build in public.
Kaggle Competitions from Month 2	Enter competitions even if you finish last. The discussions and top notebooks are worth more than any course.
Track Experiments from Day One	Use MLflow or W&B; from your very first ML project. The habit of reproducibility will save you hundreds of hours.
Review & Refactor Old Projects	Every month, pick one old project and make it 30% better. That is how senior-level code intuition develops.

FREE RESOURCES REFERENCED IN THIS ROADMAP

Resource	URL / Platform	Phase / Best For
fast.ai	fast.ai	Phases 3-4 — Best practical DL course
Hugging Face Learn	huggingface.co/learn	Phases 5-8 — Transformers & fine-tuning
DeepLearning.AI	deeplearning.ai	All phases — Andrew Ng's short courses
CS231n Stanford	cs231n.stanford.edu	Phase 4, 6 — Best CV course ever
Andrej Karpathy	youtube: Andrej Karpathy	Phase 4, 7 — Build GPT from scratch
Papers with Code	paperswithcode.com	Phase 11 — Find & implement papers
ArXiv	arxiv.org	Phase 11 — AI research frontier
LangChain Docs	python.langchain.com	Phase 7 — Agents & RAG
LlamaIndex Docs	docs.llamaindex.ai	Phase 7 — Advanced RAG
Kaggle	kaggle.com	Phases 2-6 — Competitions & datasets
Modal	modal.com	Phase 9 — Serverless GPU deployment
Guidexia	guidexia.com	All phases — Structured learning + mentorship

In 9 months, you will be an AI Engineer. The only variable is whether you start today.

[guidexia.com](https://www.guidexia.com)