

Training Schedule

Module 1: Foundational Aspects of Generative AI

- Evolution of AI to Foundation Models and Generative AI
- Foundations of Generative AI: LLMs, Diffusion Models, Multimodal AI
- Large Language Models (LLMs)
 - Transformer architecture (conceptual understanding)
 - Tokenization, embeddings, attention mechanisms
 - Pre-training, fine-tuning, instruction tuning
 - Parameter-efficient adaptation (LoRA, QLoRA – overview)
- Prompt Engineering
 - Prompt patterns and best practices
 - Structured prompts for scientific and government workflows
- Retrieval-Augmented Generation (RAG)
 - Need for grounding and trust
 - Vector databases and semantic search
 - Knowledge-assisted GenAI systems

Hands-on: Prompt engineering and RAG-based assistant

Module 2: Generative AI Ops Pipelines & Products

- GenAI lifecycle
- Data pipelines for GenAI systems
- Model orchestration and workflow automation
- MLOps vs GenAIOps: key differences
- Model evaluation, logging, and observability tools
- Managing hallucination, drift, and performance decay
- Overview of GenAI-enabled products and platforms

Hands-on: Designing an end-to-end GenAI pipeline

Module 3: Applied Generative AI Use-Cases

Generative AI in Research & R&D

- Literature review automation
- Hypothesis and experiment design support
- Code and documentation generation
- Synthetic data generation

Generative AI in Healthcare

- Clinical text summarization
- Medical report generation
- Data privacy and compliance considerations

Generative AI in Smart Governance

- Citizen service chatbots
- Policy drafting and document summarization

- Decision-support for administrators
- Multilingual service delivery

Case Studies: Designing Sector Playbooks in GenAI

Module 4: Building & Deploying Generative AI Applications at Scale

- Architecture of GenAI applications
- API-based GenAI services
- Backend integration with existing IT systems
- Model serving and inference optimization
- Hardware considerations (CPU/GPU/Edge)
- Scaling, latency, and cost optimization
- Security, access control, and system hardening

Hands-on: Building and deploying a GenAI application

Module 5: Open-Source & Sovereign AI Systems

- Open-source LLM ecosystem overview
- Advantages of open vs proprietary models
- On-premise and air-gapped GenAI deployments
- Data localization and national AI sovereignty
- Cost-effective AI strategies for public institutions
- Customizing open models for local languages and domains

Discussion: Sovereign AI for developing countries

Module 6: Responsible, Safe & Ethical Generative AI

- Bias, fairness, and inclusivity
- Hallucinations and misinformation risks
- Privacy-preserving GenAI
- Security risks and misuse prevention
- AI governance frameworks and risk management
- Alignment with national AI strategies and UNESCO AI Ethics

Course Summary

This course blends foundational concepts with advanced techniques, offering participants a comprehensive understanding of Generative AI. Covering a range of topics from foundational aspects on Large Language Models (LLMs), Transformers, embeddings, prompt engineering, and Retrieval-Augmented Generation (RAG) using open-source tools like Hugging Face, PyTorch, LangChain, and FAISS. Week 2 focuses on use case of applied Generative AI on exploring research, healthcare, and smart governance , building and deploying AI applications at scale, Generative AI Ops pipelines, open-source and sovereign AI systems, and responsible, ethical AI practices and Hands on Sessions.