GIKI Full-Stack AI Bootcamp 2025: ML to LLMs

Overview:

GIK Institute is organizing the AI Spectrum Bootcamp 2025: From ML to LLMs & Beyond, a comprehensive four-week program designed to equip 50 STEM graduates with cutting-edge knowledge and hands-on expertise in artificial intelligence. Running from 28th July through 29th August 2025, the Bootcamp offers 240 hours of immersive training covering the entire AI spectrum—from foundational machine learning and classical algorithms to advanced deep learning, computer vision, transformers, MLOps, diffusion models, and the latest generative AI technologies. The program aims to enable participants to build internationally competitive AI skills through intensive practical sessions and theory classes led by experts in the field. Details of the Bootcamp structure, curriculum, and schedule are outlined in the following sections.

Program Learning Outcomes

- Apply foundational **machine learning** techniques including **linear regression**, **logistic regression**, **supervised** and **unsupervised learning**, and classification tasks.
- Develop, train, and optimize **neural networks** and ensemble models (**random forests**, **XGBoost**) using **TensorFlow**, incorporating regularization and bias-variance analysis.
- Design and implement advanced **deep learning** architectures such as **fully connected**, **convolutional**, **recurrent networks**, and **vision transformers** with best practices including **batch normalization**, **dropout**, and hyperparameter tuning.
- Employ state-of-the-art **computer vision** methods including **transfer learning**, **object detection** (**RCNNs**, **YOLO**), and **image segmentation** (**FCNs**, **U-Net**, **DeepLab**).
- Implement and fine-tune **natural language processing** models using **Word2Vec**, **transformers** (**BERT**, **GPT**), and **LLM** finetuning for tasks such as **NER** and question answering.
- Apply **MLOps** principles with tools like **Git**, **MLflow**, **Docker**, and **CI/CD** pipelines to enable reproducible and scalable ML workflows.
- Develop and deploy **generative models** including **GANs**, **diffusion models**, **GPT**, **WaveGAN**, **Magenta**, **DALL-E**, **Whisper**, and **Gemini** for multimodal generation tasks.
- Utilize **vision-language models** (**CLIP**, **BLIP**, **Flamingo**) for multimodal applications such as **image captioning**, **image-text retrieval**, and **visual question answering** (**VQA**).
- Critically evaluate ethical, fairness, transparency, and safety issues in AI, promoting responsible AI development.

Week	Theme	Summary of Content
1	Python Foundations, AI Introduction, & Classical ML	This week introduces foundational AI concepts using a clear, visual approach. You will explore key ideas, see how algorithms work in code, and gain insight into the essential mathematics behind them.. • **Day 1 & 2:** **(Ms. Safia Baloch)** Get hands-on training in Python programming and AI-powered tools for data analysis, automation, and visualization. Learn to build scripts that interact with large language models, automate real-world tasks, and analyze your own data. • **Day 3:** **(Dr. Ali Imran)** On Day 3, you'll immediately begin building practical machine learning models in Python using NumPy and scikit-learn. You'll learn the fundamental distinctions between supervised and unsupervised learning, and between regression and classification tasks. The core of the day will involve building your first linear regression model, understanding the critical concepts of cost functions, and implementing gradient descent to train your model effectively. • **Day 4 & 5:** **(Dr. Ali Imran)** You'll build and train neural networks in TensorFlow for multi-class classification, implement decision trees/ensembles (random forests, XGBoost), and master techniques to ensure real-world model generalization. Key Skills You Will Gain: – Build binary/multi-class classifiers (e.g., handwritten digit recognition). – Diagnose bias/variance using learning curves and apply regularization/data augmentation. – Construct decision trees, random forests, and boosted trees (XGBoost). – Choose between neural networks and tree ensembles for real-world tasks. • **Day 6:** Open ended projects on the topics covered during the week.

Week	Theme	Summary of Content
2	Deep Neural Networks	This week, you will explore the major technological trends fueling deep learning, learn to build and train various neural network architectures—including fully connected, convolutional, and recurrent networks—and apply them to real-world tasks in vision and language. You will practice best methods for training, regularization, and optimization, implement neural networks in TensorFlow, and use advanced techniques such as batch normalization, dropout, and hyperparameter tuning. By the end, you will be able to apply deep learning models to diverse applications, from image recognition to natural language processing using modern tools like HuggingFace transformers. • **Day 7 & 8:(Dr. Ali Imran)** The content covers setting up machine learning problems with a neural network mindset, using vectorization for efficiency, and building neural networks with one or more hidden layers through forward and backpropagation. Additionally, you'll gain best practices for training, developing test sets, and analyzing bias and variance, as well as techniques such as initialization, L2 and dropout regularization, hyperparameter tuning, batch normalization, and gradient checking. The week also introduces optimization algorithms like mini-batch gradient descent, Momentum, RMSprop, and Adam, and guides you in implementing neural networks using TensorFlow. • **Day 8&9:(Dr. Ali Imran)** You will learn about convolutional neural networks, including recent variations such as residual networks; apply your knowledge of CNNs to computer vision: object detection and semantic segmentation using self-driving car datasets.; and use neural style transfer to generate art and apply these algorithms to a variety of image, video, and other 2D or 3D data. • **Day 10 & 11: (Dr. Ahmar Rashid)** You will build and train RNNs, GRUs and LSTMs; apply RNNs to Character-level Language Modeling; gain experience with natural language processing and Word Embeddings; and use HuggingFace tokenizers and transformer models to solve different NLP tasks such as NER and Question Answering. • **Day 12:** Open ended projects on the topics covered during the week.

Week	Theme	Summary of Content
3	Computer Vision & Introduction to Natural Language Processing	This segment of the course offers a comprehensive exploration of advanced computer vision and natural language processing techniques. Starting with transfer learning and object detection frameworks, it progresses through image segmentation methods and culminates in foundational and cutting-edge NLP concepts, including transformers and large language model fine-tuning. Practical labs complement theoretical insights, ensuring hands-on mastery of these state-of-the-art technologies. • **Day 13:**(Dr. Ahmar / Dr. Khurram Jadoon) Introduction to Computer Vision Image classification Architectures. In-depth study of transfer learning techniques in computer vision, emphasizing the application and fine-tuning of pre-trained convolutional neural networks (CNNs) for tasks such as image classification. • **Day 14:**(Dr. Ahmar / Dr. Khurram Jadoon) Participants will examine both two-stage and one-stage object detection frameworks, including RCNNs, Faster R-CNN, Mask R-CNN, and YOLO. Anchor free detectors will also be covered in detail. The content integrates theoretical concepts with practical experience, enabling learners to implement and evaluate state-of-the-art transfer learning and object detection algorithms in real-world scenarios. • **Day 15:**(Dr. Ahmar / Dr. Khurram Jadoon) Covers advanced topics in computer vision, beginning with image segmentation techniques such as Fully Convolutional Networks (FCNs) and the U-Net architecture, with applications in semantic and medical image segmentation. Concepts about deep lab, YOLOseg, and sequence models will also be covered. Practical labs provide hands-on experience with these state-of-the-art techniques. • **Day 16:**(Dr. Ahmar / Dr. Khurram Jadoon) Introduction to natural language processing (NLP), Preprocessing Word2vec, and language modelling using sequence models. • **Day 17:**(Dr. Khurram Jadoon) Detailed coverage of Attention, Self Attention, Transformers, BERT GPT, and LLM finetuning. • **Day 18:** Open ended projects on the topics covered during the week.

Week	Theme	Summary of Content
4	Advanced Topics in Vision Langugae Models	• **Day 19:(Dr.Ahmar/Dr. Khurram Jadoon)** GPT, and in-context learning, introduction to reasoning (Chain of thought), and Parameter Efficient Fine Tuning and Retrieval Augmented Generation (RAG) • **Day 20-21:(Dr. Khurram Jadoon)** Vision transformers, Introduction to Vision-Language Models, Motivation, and Real-world Use cases: Image captioning, Image–Text Retrieval, and VQA, etc. How VLMs Work: Core Components (Vision encoders, Language encoders, Join Embedding (Contrastive vs Generative). Architecture I (CLIP, BLIP, and Flamingo, etc.). Architecture II, Zero Shot Classification. Hans-on-Lab on CLIP or BLIP for Vision-Language Tasks (VQA). • **Day 22-23:(Dr. Khurram Jadoon)** Introduction to Qwen Architecture and distinguished features such as grouped Query Attention, Rotary Positional Embeddings, RMSNorm, and SwiGLU Activation function. Introduction to the DeepSeek Model R1/VL with Distinguished Features such as Mixture-of-Experts (MoE) transformer architecture. • **Day 24:(Dr. Khurram Jadoon)** Deep seek Contd : Multi-Token Prediction and Memory Optimization, Multi-Head Latent Attention (MLA), High performance at reduced computational cost, and Fine-tuning and reinforcement learning. Concept of Multi-Modal CoT.
5	Adv. Topics in Generative AI & Introduction to MLOps	• **Day 25:(Dr. Shahab Ansari)** Generative AI for Text and Speech,Text Generation with GPT Models ,Speech Generation with WaveGAN, Music Generation with Magenta, Applications of Generative AI in Text, Speech, and Music, Implementing Generative Models for Text, Speech, and Music. • **Day 26:(Dr. Shahab Ansari)** State-of-the-Art Generative Models, GPT-4. The Next Generation Language Model ,Midjourney. Continual Learning Framework, DALL-E. Creating Images from Text Descriptions, Whisper and Gemini. Innovations in Generative AI. Hand-on-Lab on exploring State-of-the-Art Generative Models. • **Day 27-28:(Mr. Sajid Ali)** A practical overview of the machine learning operations lifecycle, covering essential concepts such as version control, experiment tracking, and containerization. Participants will learn about key tools like Git, MLflow, and Docker, and explore strategies for deploying and monitoring machine learning models in production. The course also introduces basic CI/CD automation practices and offers hands-on exposure to building simple deployment pipelines, equipping attendees with fundamental skills to streamline and manage the end-to-end ML workflow. • **Day 28-29:** Project

Schedule of Classes and Labs:

- Program Timeline: July 28, 2025 - August 29, 2025 (including Saturdays)
- Theory Classes: 8:30 am – 1:00 pm (4 Hours)
- Tea Break: 10:30 – 11:00 am
- Lunch: 1:00 pm – 2:00 pm
- Labs: 2:00 pm – 5:00 pm (3 Hours)