LAHORE UNIVERSITY OF MANAGEMENT SCIENCES
Department of Electrical Engineering

EE514/CS535 Machine Learning
Spring 2025

Project Description

Total Marks: 50
Contribution to Final Assessment: 20%

1 Objectives

You can choose one of the following course projects:

Project 1: Robust Real-Time Age and Gender on Edge Computing devices

Project 2: Earth System Digital Twin for rainfall forecasting over Punjab, Pakistan

The primary objectives of this course project are:

- Develop hands-on experience in applying machine learning techniques to real-world problems.
- Understand the role of feature extraction and preprocessing in improving model performance.
- Explore state-of-the-art machine learning methods for classification, regression, and forecasting tasks.
- Learn to evaluate and compare models based on standard performance metrics.
- Gain experience in model deployment and optimization for real-time applications.
- Enhance documentation, reporting, and presentation skills related to machine learning solutions.

Students will work in teams of up to three members. Each student in the group will receive the same score.

2 Project 1 - Robust Real-Time Age and Gender on Edge Computing devices

2.1 Objective

This project aims to develop a machine learning model capable of predicting age and gender from facial images and deploy it on web-based platforms using TensorFlow.js for real-time inference on edge devices (mobile phones, laptops). The project will capture images from the user's device (webcam or mobile camera) via a frontend interface and process them using a machine learning model. The project will explore the trade-offs between model size, accuracy, and computational efficiency by progressively implementing different machine-learning techniques.

The learning outcomes of this project include:

- Develop a solid understanding of convolutional neural networks (CNNs) and their variants.
- Learn to deploy machine learning models using WebML.
- Develop a working understanding of probabilistic machine learning methods.

2.2 Data Set Description

The UTKFace dataset is a widely used benchmark dataset in computer vision and machine learning, specifically for age, gender, and ethnicity classification tasks. It contains over 20,000 images of human faces, each labeled with age, gender, and ethnicity information. The dataset spans a diverse range of ages from 0 to 116 years, making it one of the most comprehensive facial datasets for age estimation. The dataset has been extensively used in tasks such as:

- Age Estimation - Predicting the age of a person based on facial features. Since age is a continuous numerical value, age estimation is treated as a regression problem. The model predicts a real-valued output representing the estimated age.
- Gender Classification – Distinguishing between male and female faces. Since gender typically has discrete labels (e.g., male or female), gender classification is treated as a classification problem. The model assigns each face to a specific class.

You can download the filtered dataset from this link (Google Drive).

2.3 Phases of the Project

You need to build the model from scratch but can leverage the pretrained weights from ImageNet to enhance training efficiency and performance. The dataset must be divided into train (70%), validation (15%), and test (15%) sets.

2.3.1 Phase 1: Initial Deployment with Lightweight Models

- Implement small models such as ResNet-34.
- Deploy the models on a web platform using TensorFlow.js.

  - TensorFlow.js $\longrightarrow$ To load and run the trained machine learning model in the browser.
  - WebRTC or Web APIs $\longrightarrow$ For accessing the webcam and capturing images in real-time.
  - React.js / Vue.js $\longrightarrow$ To create a dynamic and interactive UI.
- Evaluate real-time performance on different edge devices (e.g., mobile phones, laptops).
- Analyze performance in terms of metrics, latency, and computational efficiency.

2.3.2 Phase 2: Deployment of State-of-the-Art Models

- Train and deploy state-of-the-art (ResNet-152), more accurate deep learning models.
- Compare performance improvements with the lightweight models.
- Evaluate the feasibility of running large models on WebML while ensuring low latency.

2.3.3 Phase 3: Uncertainty Quantification

- Explore probabilistic machine learning methods to improve robustness. In probabilistic machine learning, we aim to learn the probability distribution of the output layer weights rather than deterministic values, enabling the model to quantify uncertainty in its predictions.
- Predict age or gender with certainty.

2.4 Deliverables

- Conduct a thorough review of the relevant literature on age and gender prediction using machine learning. Incorporate innovative ideas into your model design and deployment strategy to enhance performance and robustness.
- In the project report, comprehensively evaluate model performance by presenting key performance metrics on the test set. Include results from real-time testing on different devices and provide their specifications. Additionally, include visualizations of real-time testing outcomes by displaying captured images with corresponding predictions.
- To further analyze model behavior, discuss the training and validation loss curves to assess potential overfitting or underfitting. You may also analyze latency vs. accuracy trade-offs, and quantify uncertainty estimates (if applicable) to ensure robustness. You can report your results on different devices and provide their specifications.
- If deploying on edge devices, discuss optimization techniques (in the report) such as model quantization, pruning, or knowledge distillation to enhance efficiency without compromising accuracy.

3 Project 2 - Earth System Digital Twin for rainfall forecasting over Punjab, Pakistan

3.1 Motivation

Understanding rainfall patterns is crucial for climate resilience, disaster preparedness, and sustainable water management. However, global climate models often lack the fine spatial resolution needed for accurate local decision-making. This project aims to bridge that gap by leveraging machine learning, specifically a Long Short-Term Memory (LSTM)-based model, to downscale ERA5 precipitation data from a coarse 150 km resolution to a fine 25 km resolution.

This work is inspired by the research of Bittner et al. (A.5), where they successfully applied an LSTM-based downscaling framework to improve precipitation projections in Australia. Building on their methodology, we will develop a first-ever regionalized model for Punjab, Pakistan, tailored to its unique climatic conditions and monsoon-driven rainfall patterns.

This project provides an excellent opportunity for students interested in climate science, machine learning, and geospatial analytics to gain practical experience in applying AI to environmental challenges, including flood forecasting, water resource management, climate modeling, and sustainability-focused industries.

The learning outcome of this project include:

- Develop expertise in processing and analyzing climate datasets i.e. ERA5 reanalysis data.
- Gain hands-on experience in designing, training, and evaluating LSTM-based deep learning models over same spatio-temporal datasets.
- Learn to generate and visualize high-resolution precipitation maps, supporting applications in flood forecasting and climate adaptation.

3.2 Dataset

The ERA5 reanalysis dataset, produced by ECMWF under the Copernicus Climate Change Service (C3S), provides global climate and weather data with spatial (25 km) and temporal (hourly) resolution from 1940 to the present. It integrates historical observations from satellites, weather stations, and other sources with advanced numerical models to offer detailed information on atmospheric, land, and oceanic conditions, including temperature, wind, humidity, and precipitation. The dataset is freely accessible through the Copernicus Climate Data Store (CDS) via a web interface, Python API, and Google Earth Engine. For this project you can download the dataset from this https://cds.climate.copernicus.eu/datasets/reanalysis-era5-single-levels?tab=download

3.3 Implementation Details

You will be using a set of predictor variables (input) in coarse resolution (150km) to predict (output) precipitation values at fine resolution (25km). For training of the model, the coarse resolution dataset would be generated by upscaling the same fine resolution dataset, which would also serve as the ground truth.

3.3.1 Dataset Processing

- Download the climate predictors A.1 from "ERA5 hourly data on single levels from 1940 to present" from the Climate Data Store (link provided earlier) in GRIB format of the year 1999-2011 for the Punjab region A.4.
- Compute the daily means (note: aggregates for precipitation and not the mean value) for downloaded parameters, as well as calculate the A.2 predictors. Save the data in tiff format for daily averages, aggregates (precipitation), and min/max values (temperature).
- Using the tiff files, calculate the spatial encoding and assign temporal encodings to daily values i.e. A.3
- Upscale the daily fine-scale data (25km) to a coarse resolution of 150km.
- Divide the data into train and test for the years 1999-2009 and 2010-2011, respectively.

3.3.2 Model Formulation and Training

- Create LSTM model(s) with model parameters from the title paper.
- We can assign an LSTM model for each coarse-resolution pixel. For this project we will be focusing only on the central 04 pixel of the coarse resolution Punjab map. i.e. 04 LSTM models
- Model Input (for each LSTM):
  - Coarse resolution climate variables (A.1, A.2) from self and all adjacent pixels i.e. 09 parameters from each 06 pixel values.
  - Spatial encoding for each fine-resolution spatial grid (36, 3) of the central pixel.
  - Coarse-resolution central pixel, temporal encoding (1,2)
  - window size for each LSTM : 56 days.
- Model Output (for each LSTM):
  - Precipitation value for the next day for each of the fine-resolution pixels of the reference coarse-resolution pixel (1,36).
- Model Training:
  - We would be training each LSTM separately.
  - Train the model on historical ERA5 data (1999-2009).
  - Test the model on recent years (2010-2011).

3.3.3 Generation of high-resolution precipitation maps for Punjab

- Convert the predicted values into daily spatial maps (25km).
- Generate a side-by-side (coarse, predicted, fine res) daily variations (preferable as a video/gif) of precipitation for the season monsoon season of 2011.

4 Expectations and Scope of Work

4.1 Scope of Work (Project 1 and Project 2)

For both projects, students are expected to follow these key steps:

1. **Problem Formulation**: Define the problem statement, research questions, and objectives.
2. **Literature Review**: Conduct a comprehensive study of related work to identify best practices and state-of-the-art approaches.
3. **Dataset Processing and/or Feature Engineering**: Extract and preprocess relevant features from the dataset to improve model performance.
4. **Dimensionality Reduction (if applicable)**: Analyze the impact of dimensionality reduction techniques on model performance.
5. **Model Implementation**: Train and evaluate multiple machine learning models for the given problem.
6. **Performance Evaluation**: Compare different models based on key performance metrics and discuss trade-offs.
7. **Final Report and Presentation**: Summarize findings, insights, and conclusions in a well-structured report and presentation.

4.2 Assessment

The project will be assessed based on the following components:

- **Project Report (20 marks)**: Detailed documentation covering problem formulation, methodology, implementation, and results.
- **Project Code and Documentation (15 marks)**: Well-structured and documented codebase.
- **Video Presentation (10 marks)**: A 3-minute summary of the project findings, highlighting key contributions.
- **Timely Submission (5 marks)**: Adherence to deadlines for deliverables.

We encourage you to use a template from your favorite machine learning conference (e.g., NIPS or ICML) for report. We encourage Masters and PhD students to use the LaTeX template for their reports. You can seek guidance from TAs regarding the sections that are required to be included in the report.

5 Timeline of Deliverables

We want you to adhere to the following timelines.

- **Deliverable 1:** Project selection and group formation.
  - Due: Week 9, 21st March, Friday 23:55
  - The spreadsheet for groups can be found here.
- **Deliverable 2:** Mid-Term Report and Initial Implementation
  - **Due Date**: Week 12, 11th April (Friday, 23:55)
  - Submission must include draft report sections covering:
    * Abstract
    * Introduction
    * Detailed Project Overview
    * Problem Formulation
  - Initial implementation:
    * For Project 1: A working prototype for feature extraction from facial images and preliminary classification model(s) implementation.
    * For Project 2: Preliminary data processing results and initial design of the LSTM-based downscaling model for rainfall forecasting.
- **Deliverable 3:** Video presentation (3 minutes)
  - Due: Week 14: 22nd April, Tuesday, in class. Any group member may be asked to present.
- **Deliverable 4:** Submit code (documented) and final report
  - Due: 04th May, Sunday 23:55

A Appendices

A.1 Climate Predictors - ERA5 Reanalysis

We will utilize the following climate variables as predictors for precipitation downscaling:

- Temperature (°K) – 2m temperature
- Geopotential (m^2 s^{-2})
- 10m U Wind Component (m s^{-2})
- 10m V Wind Component (m s^{-2})
- Mean Sea Level Pressure (Pa)
- Total Precipitation (mm)
- Total cloud cover

A.2 Climate Predictors - Generated

We can generate the following predictors from daily max and min values of temperature:

- Maximum Temperature at 2m (°K)
- Minimum Temperature at 2m (°K)

A.3 Spatial and Temporal Encodings

You may use the reference from the refernce paper for these encodings:

- Spatial Position (Latitude, Longitude)
- Time Features (Day, Month, Season Encoding)

A.4 Punjab Region Bounds

- North: 34.05
- South: 27.65
- East: 75.5
- West: 69.25

A.5 Reference Paper

Bittner, Matthias, et al. "An LSTM-based downscaling framework for Australian precipitation projections." NeurIPS 2023 Workshop on Tackling Climate Change with Machine Learning. 2023.