Predicting Global Wheat Crop Yields Using a Combination of CNN and LSTM Blocks

Sungjoon Park

Apr. 30th, 2024

Boston University Faculty of Computing & Data Sciences

BOSTON UNIVERSITY

Objective and Data

- Objective: Making a model that can accurately predict global wheat crop yields using meteorological image datasets
- Data
  - Input: Global near-surface temperature and precipitation heatmaps (Monthly)
    - Source: Climate Data Store by Copernicus
    - Samples (Left: near-surface temperature, Right: precipitation)

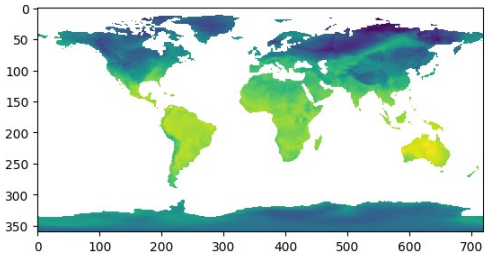

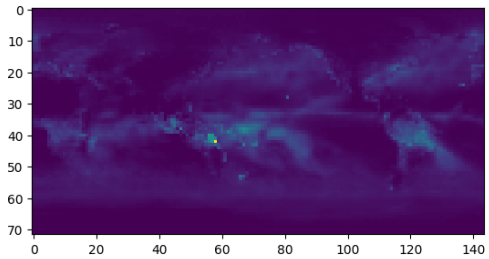

  - Target: Global yearly wheat crop yield per hectare (unit: ton)
    - Source: Our World in Data
  - Periods: Training (1980-2015), Validation (2016-2018), Test (2019)

BOSTON UNIVERSITY

Model

- CNN block

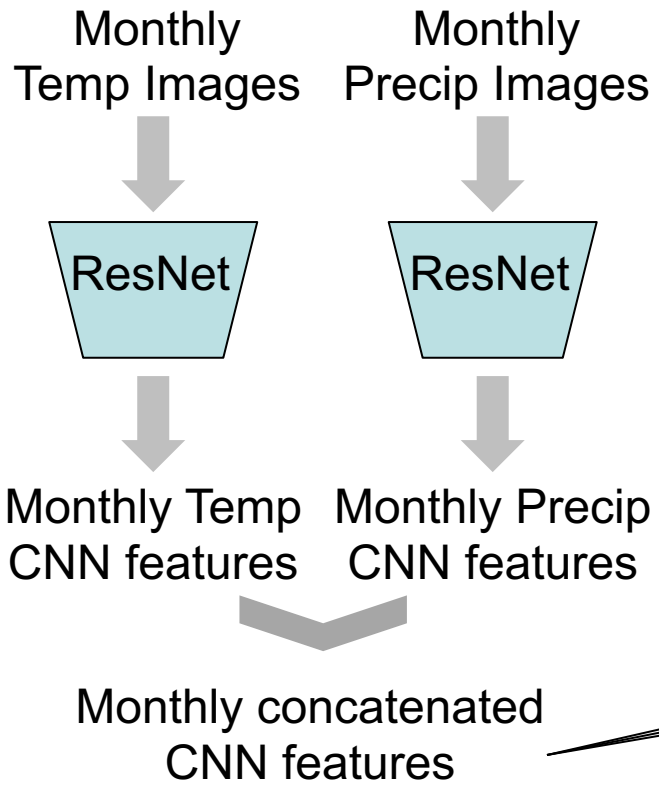

- LSTM block

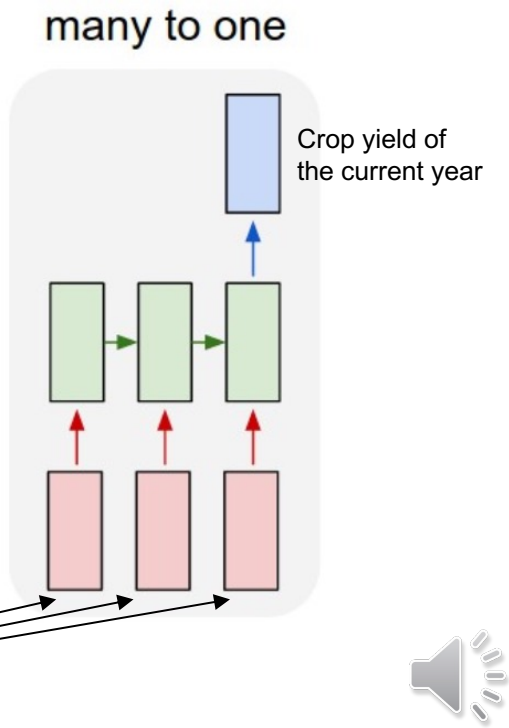

BOSTON UNIVERSITY

Two Tricks

- Auto-regressive prediction

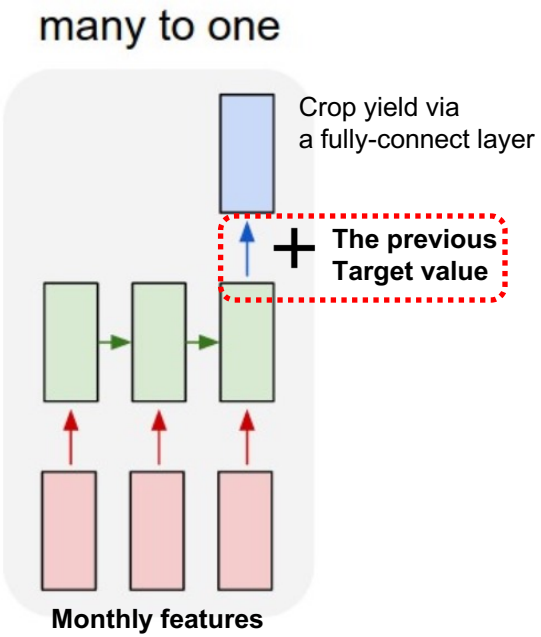

- Positional Encoding

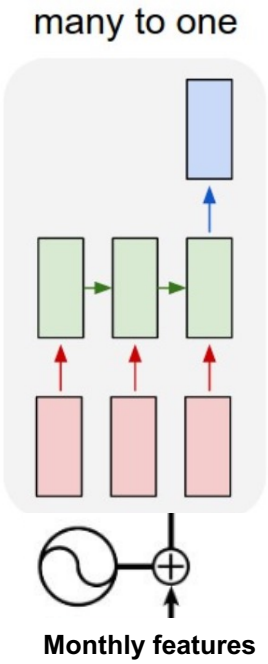

Positional Encoding Example

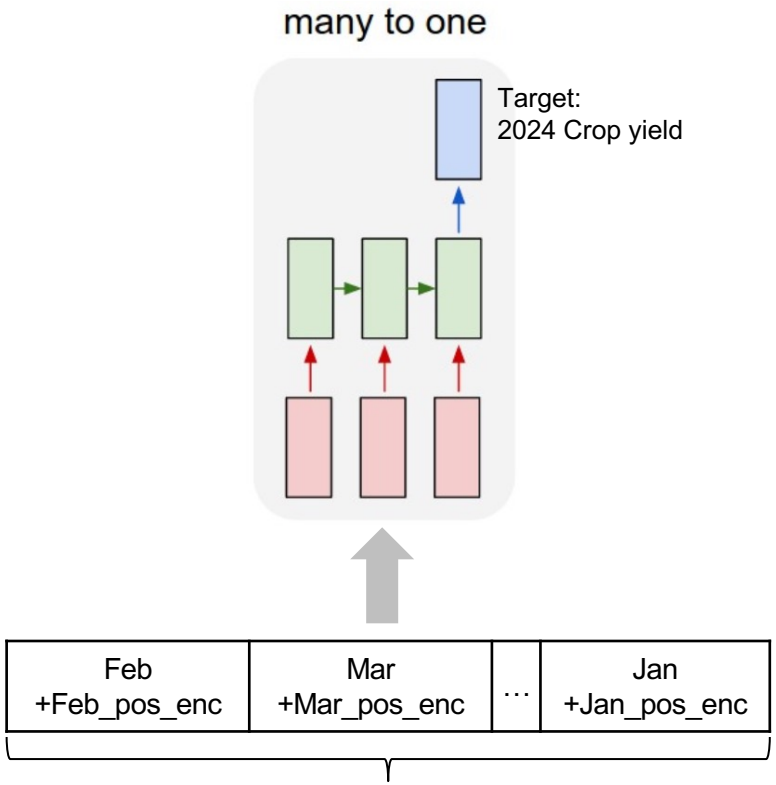

Monthly features for each layer
(Feb 2023 ~ Jan 2024)

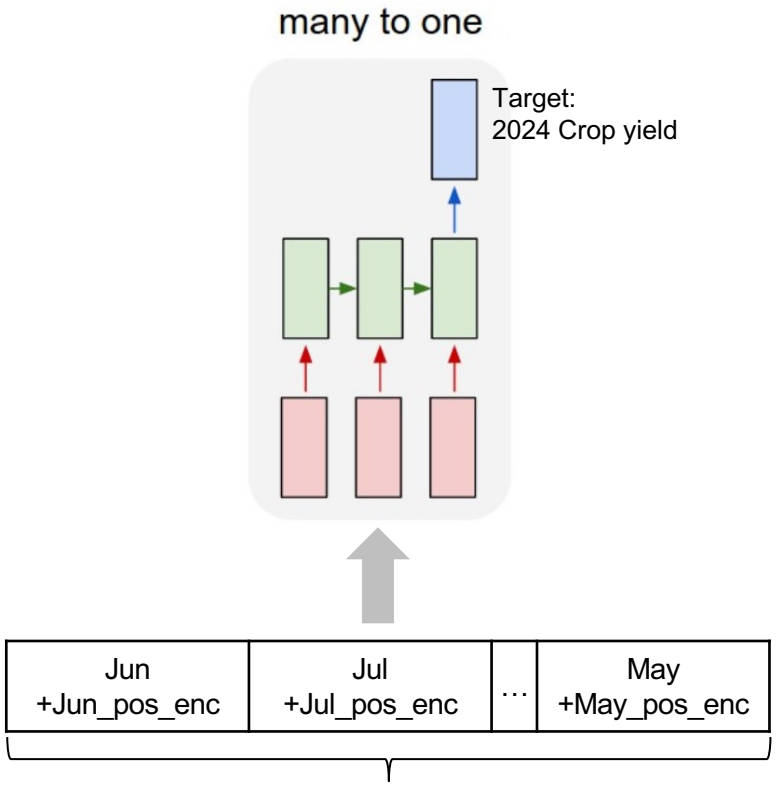

Monthly features for each layer
(Jun 2023 ~ May 2024)

BOSTON UNIVERSITY

Model Training and Performance

- Configurations
  - Framework: PyTorch
  - CNN Block Pre-trained Model: ResNet15
  - Loss function: MSE
  - Optimizer: Adam
  - Learning rate: 1e-3
- Hyperparameter-tuning: hidden size of the LSTM
- Results

Hidden Dim	64	128	256	**512**
Epoch when best	100	100	100	**29**
Test MSE	0.0057	0.005	0.0014	**0.0006**

BOSTON UNIVERSITY

Limitations and Future Work

- Incorporate additional inputs
  - Humidity and soil type heatmaps
- Hyperparameter-tuning: hidden size of the LSTM
  - Comparison study
- Results
  - More diverse hyperparameter-tuning
    - CNN Block pre-trained model
    - Position Encoding Strategy
    - Loss functions (e.g. L1 Loss)
    - Optimizers
    - Learning rates