Optimizing Doppler Velocity Estimation with Deep Learning

Project 4, Deep Learning F24

Xixi Wang, Vincenzo Genovese, Artur Habuda, Alexandru-Stefan Marin

Overview

- **Goal:** Estimate the radial velocity of a golf ball using 6-channel spectrograms with CNN, aiming to minimize parameters while maintaining low RMSE.
- **Key Innovation:** Introduced **skip connections** to enhance feature flow and reduce parameters. Two models were designed: one with residual connections and another with Squeeze-and-Excitation blocks, optimized with alternative activation functions (LeakyReLU, SiLU) and pooling strategies.
- **Result:** Achieved **1**% of the baseline model's parameters with slightly better RMSE.

Architectures Overview

- **Baseline Model (SpectrVelCNNRegr):**
  1. Serves as a benchmark for comparison.
  2. Uses basic ReLU activations and MaxPooling for downsampling.

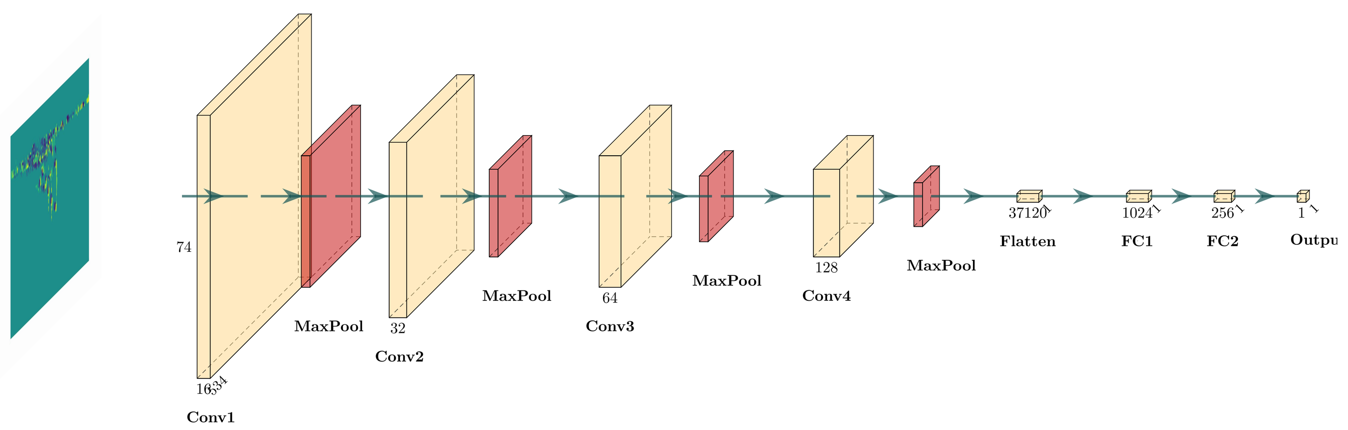

- **Residual-ReLU Regressor:**
  1. Trained with two activation functions: LeakyReLU and SiLU.
  2. Introduced residual connections to prevent vanishing gradients and improve feature flow substitute Pool with Convolutional Layer and GAP
  3. Added Dropout and BatchNorm for regularization to mitigate overfitting.
  4. Test hyper. para. influence on test rmse: Adam+41.89% - Kaiming+6.64% - SiLU+7.99%

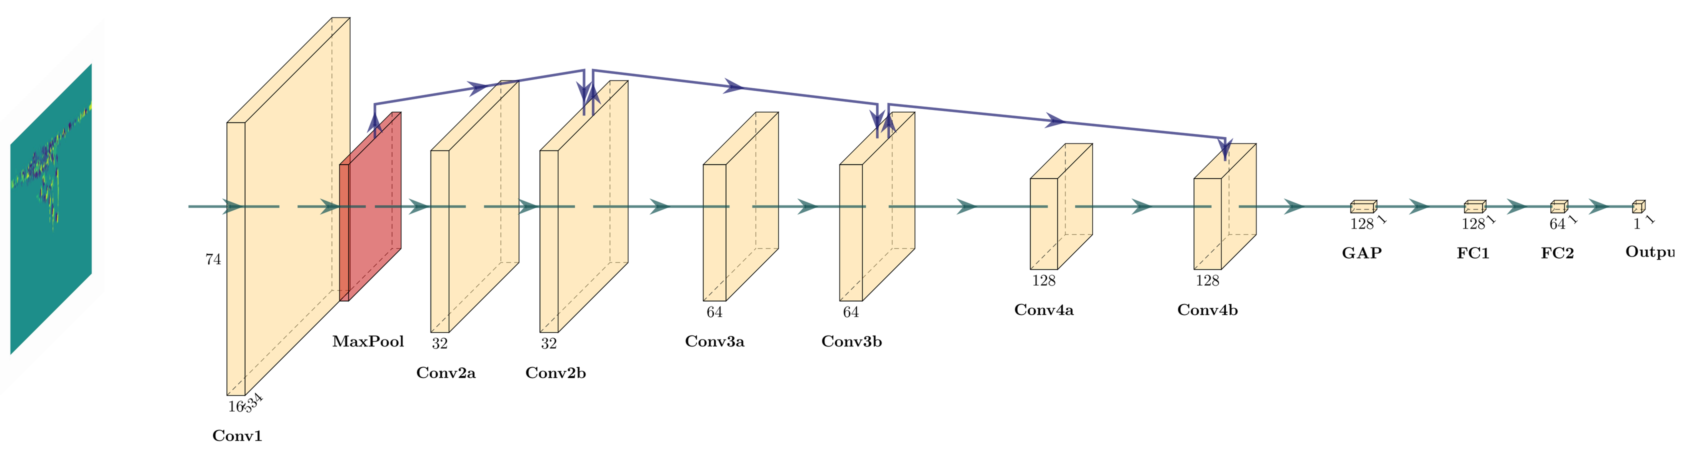

- **Squeeze-Excite Residual-SiLU Regressor :**
  1. Introduced Squeeze-and-Excitation (SE) blocks for channel-wise attention.
  2. Implemented versions with different activation functions: SiLU and LeakyReLU+1.54%.
  3. Utilized Global Average Pooling to efficiently reduce spatial dimensions.

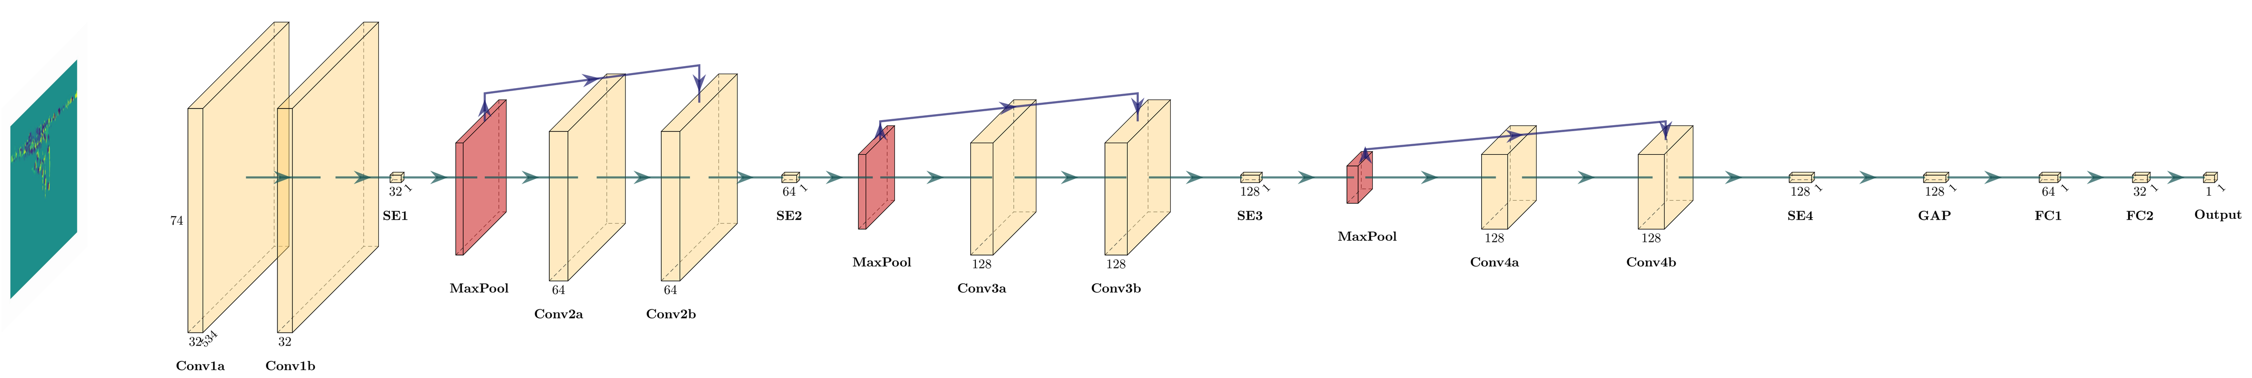

Hyperparameters

- **Learning rate:** 1e-5
- **Epoch:** 500
- **Batch size:** 10
- **Optimizer:**
  - **SGD:** better generalization and final convergence quality, despite slower initial progress.
- **Activations:**
  - **LeakyReLU:** Reduces dead neurons, improving training stability through negative slope.
  - **SiLU:** Smoother gradients lead to consistent and improved accuracy.

Test Results

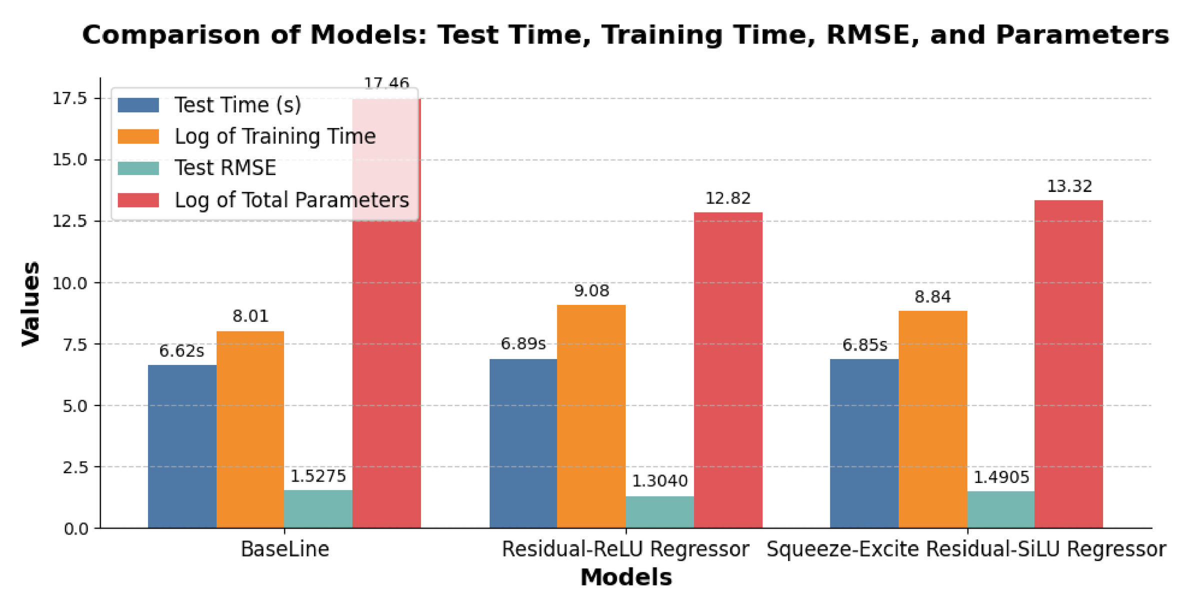

The percentual values below compare each model against the baseline :(+/-) indicates an increase or decrease in parameter value. green/red indicates improvement or deterioration. Residual-RELU Regressor(RU), Squeeze-Excite Residual-Silu(SU).

- **Test rmse**: RU:+15%, SU:+3%
- **Model Parametrs**:RU:-98%, SU:-99%
- **Inference Time**:RU:+4.1%, SU:+3.4%
- **Training Time**:RU:+190%, SU:+120%

Inference time results have been obtain by averaging out time between 5 tests.

Dataset

- **Shape of the data: [79, 918, 6]**
  4 power channels reveal scattering properties: **ball detection**.
  2 phase channels provide phase difference: **speed estimation**.
- **Data cropping, Normalization, Interpolation.**

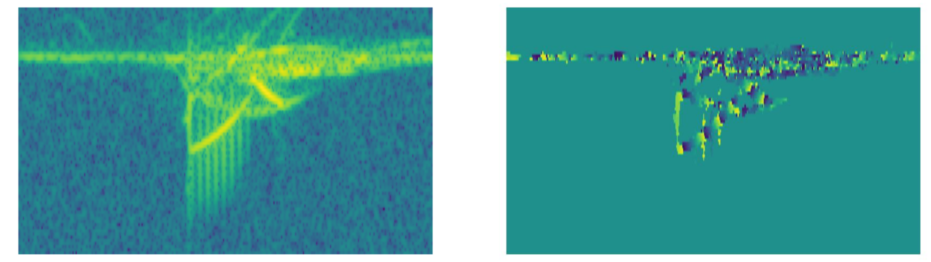

Power Spectrogram Phase Spectrogram

- **Data Split**

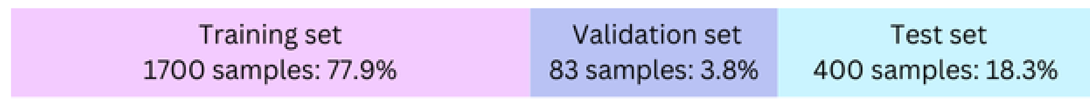

It was decided to switch validation and test sets, in order to have a higher sample count on test set. Achieving a more reliable final evaluation.

Convergence graphs