Finding limits of deep learning with bodybuilder image comparison

Luke Corwin
David Osterman
Andrei Prokhorov

Motivation

We would like to create a tool to help judges to
evaluate athletes at the bodybuilding competitions
based on their photos
using convolutional neural networks.

Data collection

- Photos from the National Physique Committee website: https://contests.npcnewsonline.com/contests/2024
- Only the year 2024
- Men's Physique division
- Selected 2 photos (one front, one back) from each competitor
- Data was paired photos (front only so far)
  - The winner in a pair was the competitor with lower rank
- Divided data into training, validation and test sets
  - Approx. 60/20/20 ratio
  - No competition or competitor appears in more than one set.

Model description

TwoInputNet

- Wrote ourselves based on PyTorch book example
- Two image feature recognition components, each consisting of:
  - 3 Conv2D layers
  - ReLU non-linear elements
  - MaxPool2d
- Image comparison component:
  - Linear layer
  - ReLU
  - Linear layer with output size 1
  - Sigmoid activation
- Output rounded to 0 or 1

ResNet

- Replaces the image recognition layers in TwoInputNet with ResNet50
- Otherwise identical

Results: TwoInputNet

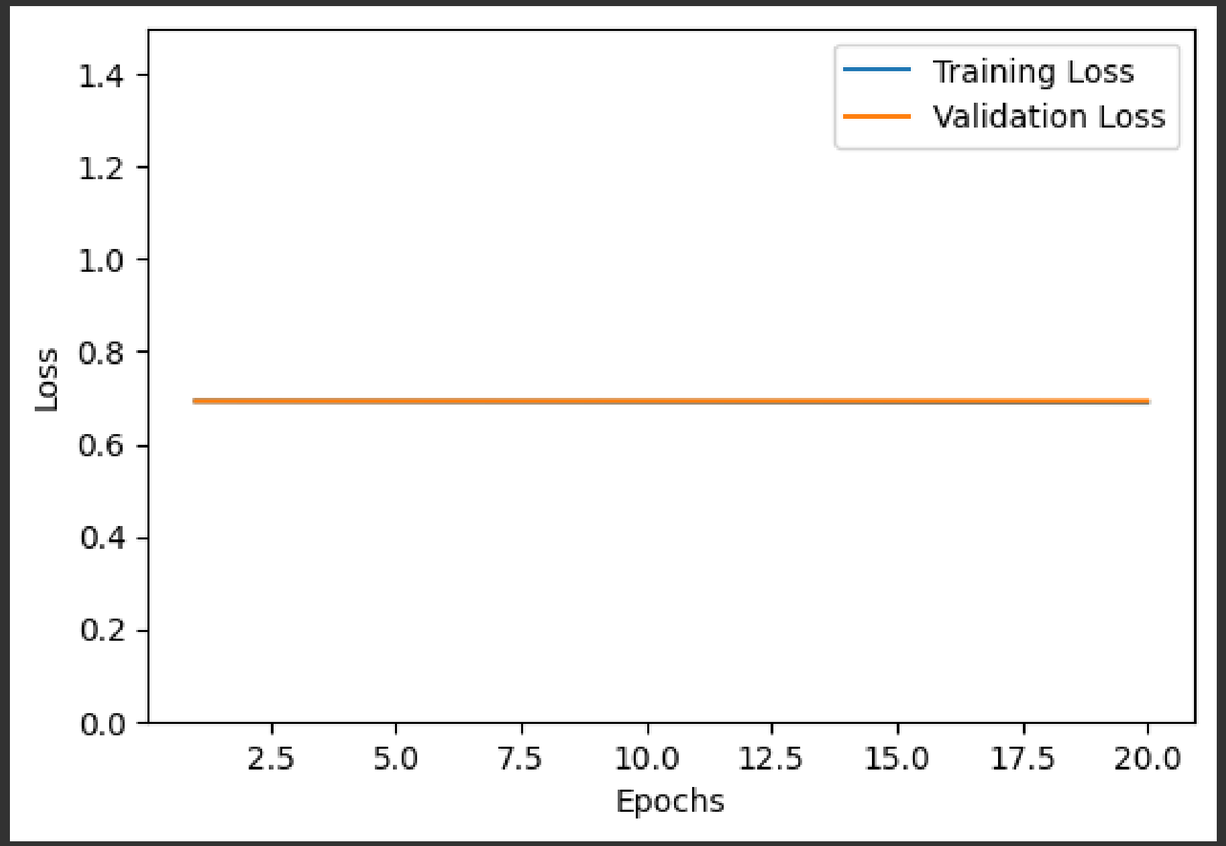

Loss vs. epochs

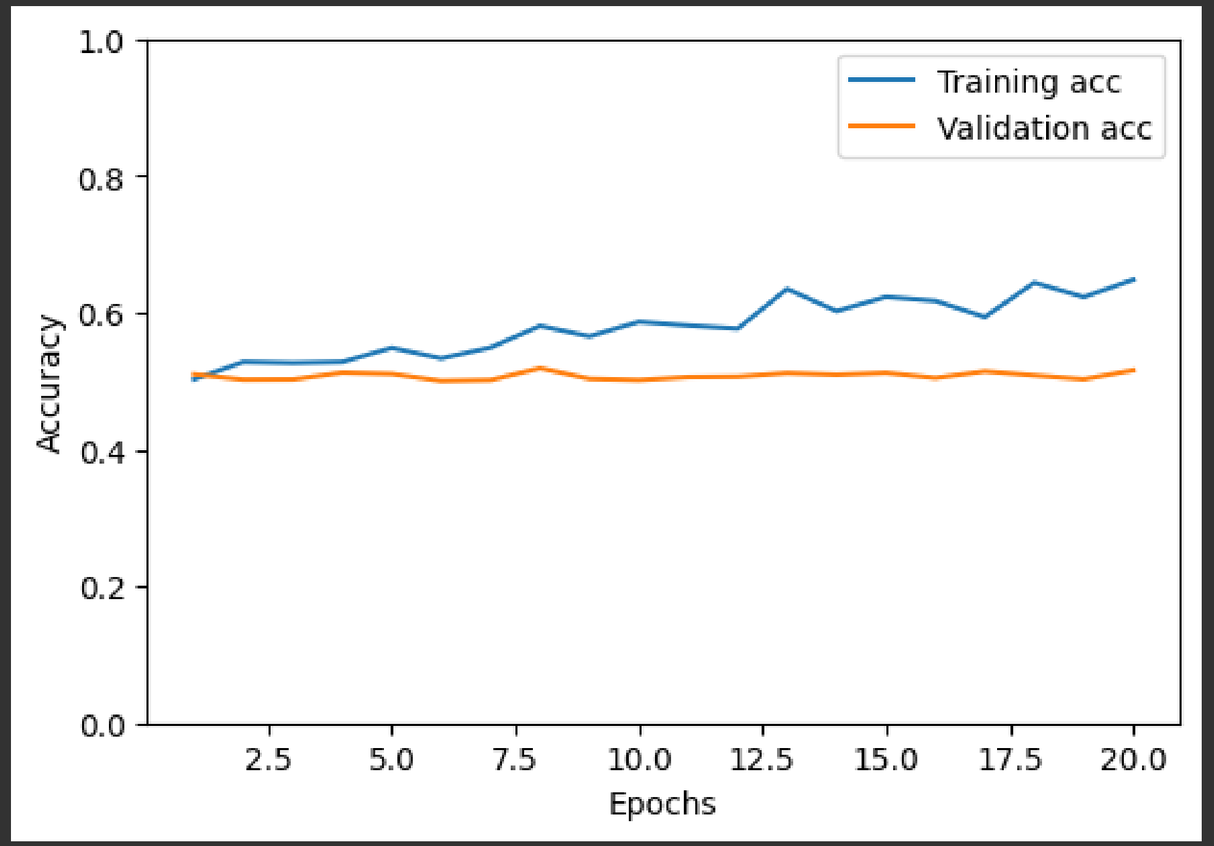

Accuracy vs. epochs

Model does not appear to be learning

Results: ResNet50

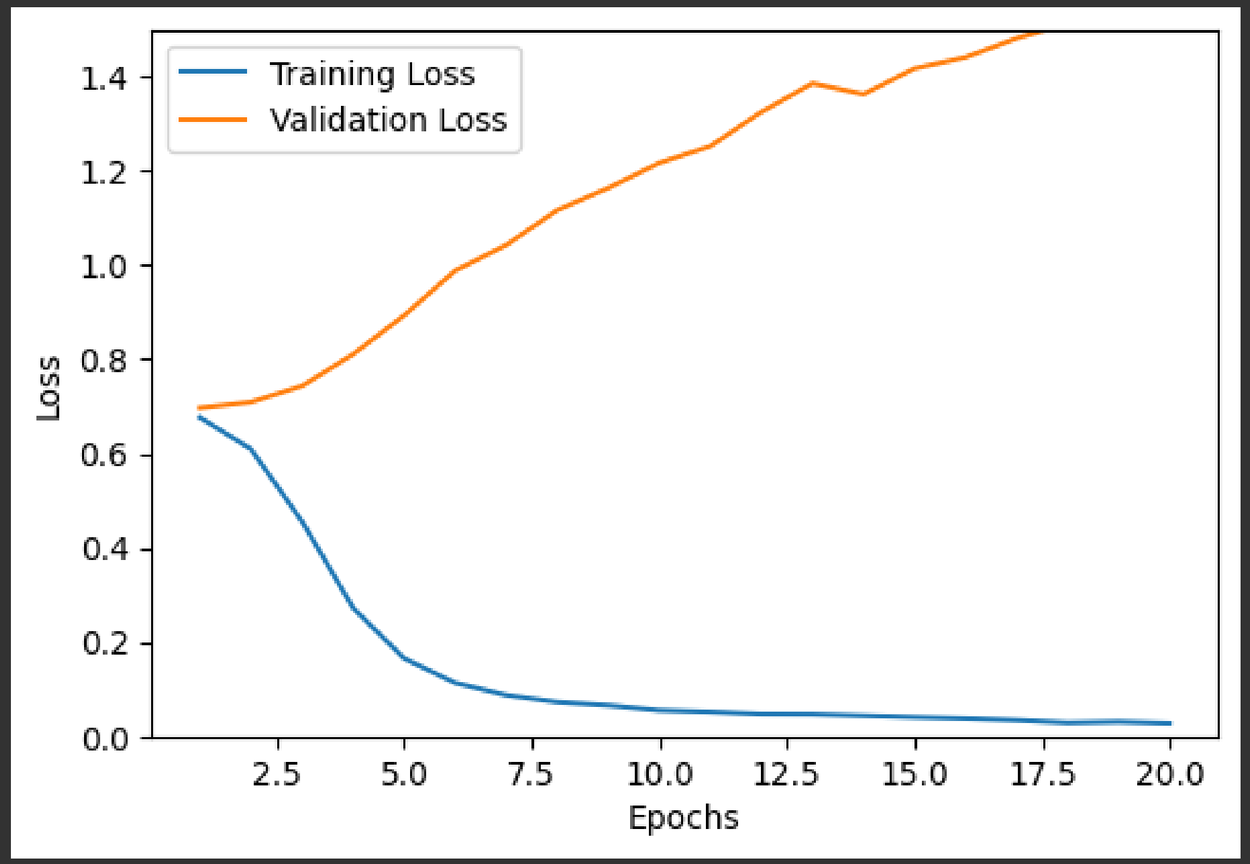

Loss vs. epochs

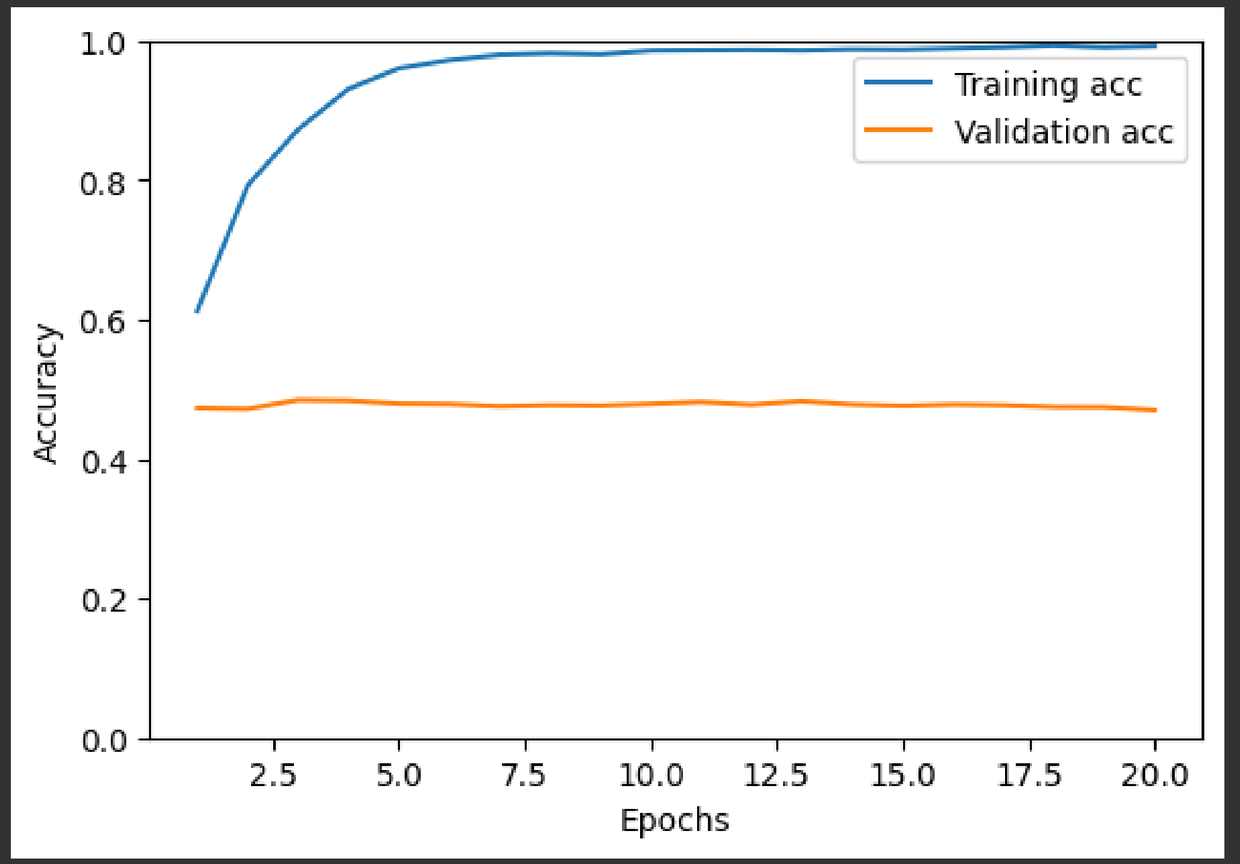

Accuracy vs. epochs

Model appears to be overtraining

Possible Future Improvements

- Train the network which would facilitate the photo selection process.
- Create larger dataset including 10 previous years using it and train the model for it. We expect better performance of the model on the validation set.

Thank you for your support!