DEEP LEARNING FOR IN-GAME NFL PREDICTIONS

Cameron Taylor | cntaylor@stanford.edu | Link to presentation

PROBLEM / MOTIVATION

Problem: In-game NFL play-calling crucial, pre-play predictions important input to strategizing.

Idea: Use pre-play data, including images, to predict outcome of play.

Model: Benchmark ML, CNN, Transfer learning

- Input = Pre-play situation + Image
- Output = (1) Yards on play
  (2) Offensive play call (pass or run)

Results:

- (1) is too difficult with data and models
- Higher accuracy on (2)
- Image data not much value

DATA / FEATURES

1. Kaggle play-by-play NFL data [1]
   - Provides labels: Y1 = yards gained, Y2 = 1{pass play}
   - Other pre-play info (X1); Ex: Score, time, etc.
2. Manually collected 1,049 play images (X2) on NFL Rewind

Merge 1. + 2. by play

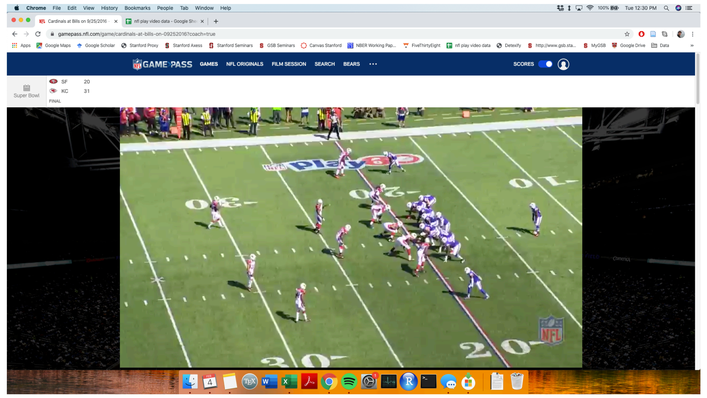

MODELS

1. Benchmark (1) = median yards (2) = guess most common play
2. Benchmark ML w/o image data (just X1) (LASSO and RF)
3. Shallow CNN w/ pre-play and image (X1 and X2)
4. Transfer Learning w/ VGG19 [2] trained on ImageNet (X1 and X2)

M = 1,049

Train / Dev / Test = 75/12.5/12.5

Tune focus

- SCNN: learn rate, # hidden units
- Transfer: learn rate, # layers from VGG19

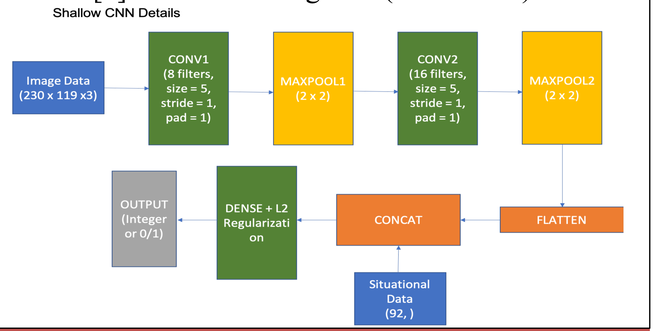

RESULTS

Model	Mean Absolute Yards
Benchmark: Guess Median of Training y	6.18
LASSO (HPs: $\alpha = 0.0176$)	6.18
Random Forest (HPs: number trees = 3, max depth trees = 3)	6.16
Shallow CNN (HPs: learning rate = 2.5×10^{-5}, epochs = 15, mini-batch = 32, # hidden units in dense layer =5, L2 reg = 0.001)	6.18
VGG19 Transfer Learning (HPs: learning rate = 2.5×10^{-6}, epochs = 15, mini-batch = 32, # hidden units in dense layer = 2, conv layers from VGG = 5)	6.21

Notes: HP = Hyperparameters.

Table 1: Results for Predicting Yardage Outcomes

in Train = 786
in test = 132

Training Absolute Yards
Approx 3

Approx 3

Model	Accuracy
Benchmark: Guess all plays are pass	0.546
LASSO (HPs: $\alpha = 0.0176$)	0.568
Random Forest (HPs: number trees = 2, max depth trees = 7)	0.614
Shallow CNN (HPs: learning rate = 1.0×10^{-3}, epochs = 15, mini-batch = 32, # hidden units in dense layer =4, L2 reg = 0.005)	0.606
VGG19 Transfer Learning (HPs: learning rate = 1.0×10^{-3}, epochs = 10, mini-batch = 32, # hidden units in dense layer = 4, conv layers from VGG = 5)	0.606

Hyperparameters

Table 2: Results for Predicting Play Call

Training Accuracy
Approx 0.9

Approx 0.9

DISCUSSION

1. Models did NOT do well on predicting yards
   - Why? Benchmark as good
   - Interpretation: not enough signal in the image or pre-play data for problem, too much overfitting

2. Models perform well on predicting play call
   - Why? Achieve better test accuracy
   - Interpretation: learn important times for certain play calls and how to predict based on player locations

3. Image data does not provide higher value
   - Why? Pre-play RF does better
   - Interpretation: pre-play data as valuable as player locations in problem

FUTURE WORK

1. Figure out how to extract more signal for predicting yards (more detailed video data)
2. Add more structure to image data (label players) and transfer learning on other models (YOLO) to improve performance
3. Evidence for interpretations

REFERENCES

[1] Max Horowitz, Ronald Yurko, and Sam Ventura. "Detailed NFL Play-by-Play Data 2009-2018." Kaggle.

[2] Karen Simonyan and Andrew Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. arXiv:1409.1556 [cs], April 2015