Run, Model, Run!
Running Speed Prediction Using Lumo Run Wearable Device

Adam Gotlin[1], Michael Hittle[2], Salvatore Calvo[3]
{agotlin,mhittle,salcalvo}@stanford.edu

Our team would like to extend our thanks to: Jessica Selinger and Rachel Jackson of the Stanford Mobilize Lab, Jessica Torres of the Stanford Ashley Lab, the Lumo Run team, and the CS 230 teaching team.

Background

Problem Statement: Many fitness trackers use inertial measurement units (IMUs) to measure activity metrics like step cadence and bounce. This data is augmented with speed and distance metrics, but only when the tracker can access a GPS signal. We leverage deep learning and a subset of Lumo Run runner data to **predict running speed without a GPS signal.**

Model Inputs/Outputs: Runner anthropometric data is combined with time-series IMU measurements to predict average speed over a specified window of time.

Approach: Minimize mean absolute error (MAE) of speed prediction by leveraging CNN and FCN networks.

Results: The highest performing model is able to predict speed with an MAE of 0.115 m/s.

Models

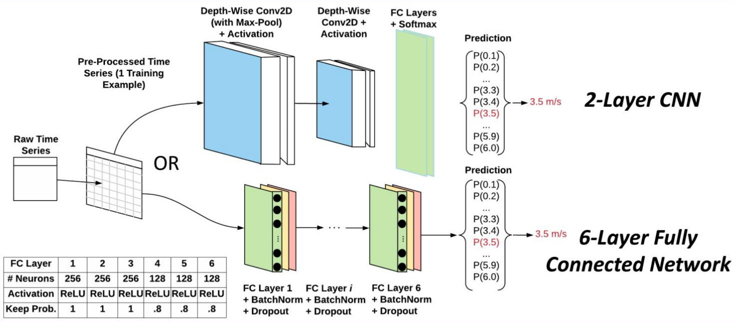

FC Layer	1	2	3	4	5	6
# Neurons	256	256	256	128	128	128
Activation	ReLU	ReLU	ReLU	ReLU	ReLU	ReLU
Keep Prob.	1	1	1	.8	.8	.8

We obtain our training instances by concatenating consecutive timestamp measurements and the associated user demographic data. The label, y_{true}, is the average speed over a specified input window bucketed to the nearest 0.1 m/s. Both models perform best with a **categorical cross-entropy loss function.**

Future Work

Our results are extremely promising and could be enhanced by:

- Leveraging **grid-search** for systematic hyperparameter tuning
- Training on a wider dataset (which is readily available, but requires greater compute resources)
- Consider a multi-model approach for distinct populations of runners (i.e. walkers vs. runners)
- Build a model to predict **current running speed** using trailing kinematic data only, which could inform real-time speed prediction on-device

Lumo Run Dataset

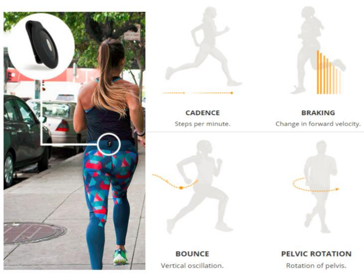

<u>Sampling Rate</u>: ~3 seconds
<u>Ground Truth</u>: GPS-derived speed
<u>Data Volume</u>: ~200,000 samples

RunnerID	Gender	Age	Height	Weight
1234	M	26	180	160
2345	M	35	186	165
3456	F	45	170	125

RunID	RunnerID	Time	Cadence	Braking	Bounce	Contact	Pelvic Rotation	Pelvic Drop	Pelvic Tilt	Speed
1234	5678	0.00	175	42	85	2	10	5	7	3.54
1234	5678	3.00	172	40	87	1	12	4	10	3.31
1234	5678	6.00	174	42	81	2	11	8	6	3.1

Results & Discussion

Best-performing CNN:

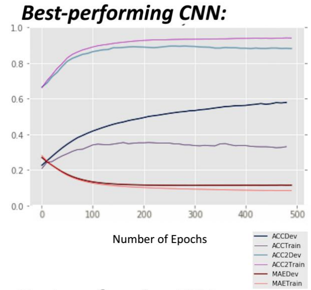

Best-performing FCN:

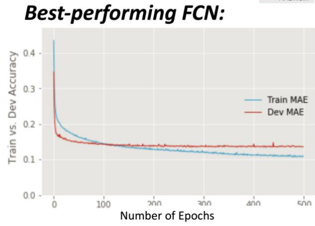

Hyperparameter Tuning and Performance:

Architecture	Conv. Layers	FC Layers	Learning Rate	Activations	Batch size	Input Window	Labeling Window	Classif. Accuracy	Buffered Classif. Accuracy*	MAE (m/s)
CNN	2	10	0.00001	ReLU > sigmoid	100	36	25	13.5 %	57.0 %	0.316
CNN	2	5	0.0001	ReLU > tanh	30	36	30	22.0 %	75.3 %	0.222
CNN	**2**	**50**	**0.0001**	**ReLU > tanh**	**50**	**36**	**30**	**33.1 %**	**88.2 %**	**0.115**
FCN	0	10	0.0001	ReLU	128	11	5	24.2 %	80.6 %	0.154
FCN	0	6	0.0001	ReLU	128	11	5	31.6 %	86.1%	0.136

*"Buffered" classification accuracy measures the percentage of instances where the predicted speed bucket falls at most 2 buckets from the true speed, or 0.2 m/s.

Error Analysis:

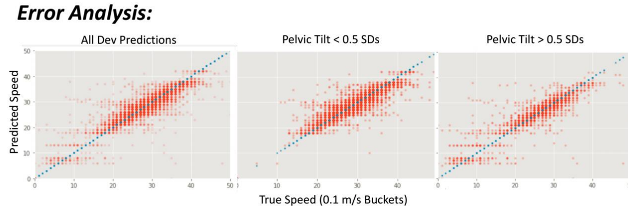

- The best model predicted speed with an **MAE of 0.115 m/s.**
- Though the goal is to predict speed (a scalar value), training with **classification outperformed regression.**
- The CNN model yields a lower dev set MAE when the time window for prediction is large.
- Confusion matrices show the FCN model tends to **perform poorly** during high pelvic tilt, which could indicate **elevation changes and/or speed-changes.**

Author Affiliations:
[1]Dept. of Mechanical Engineering
[2]Program in Human Biology, School of Humanities & Sciences; Division of Epidemiology, Dept. of Health Research and Policy
[3]Dept. of Management Science & Engineering