Predicting Human Trajectories in Multi-Class Settings

Karthik Raju (kraju444@stanford.edu), Max Chang (mchang4@stanford.edu)

Motivation

We investigate novel methods for predicting trajectories of different classes (pedestrians, bicyclists, skateboarders, carts) in crowded spaces. We use GRUs with pooled trajectories from neighboring occupants to predict the next time step at different time intervals in the future. We believe this work has applications in the navigation of autonomous agents through crowded spaces.

Approach

Naive GRU Model:

- Assume Independent trajectories (Basic Sequence Generation Problem)
- GRU: learning rate = 0.003 (with annealment), hidden state dim. = 128, 2 GRU layers, embedding with ReLU nonlinearity. Using Mean Squared Error & no BPTT truncation.

Pooling GRU Model:

1. Define 20 x 20 x 128 tensor around subject H
2. H(m,n) = pooled hidden states of neighbors at t-1 in square (m,n)
3. Embed into vector and feed to GRU

$$H_t^i(m,n) = \sum_{j \in N} 1_{m,n}[(x,y)^j] h_{t-1}^j$$

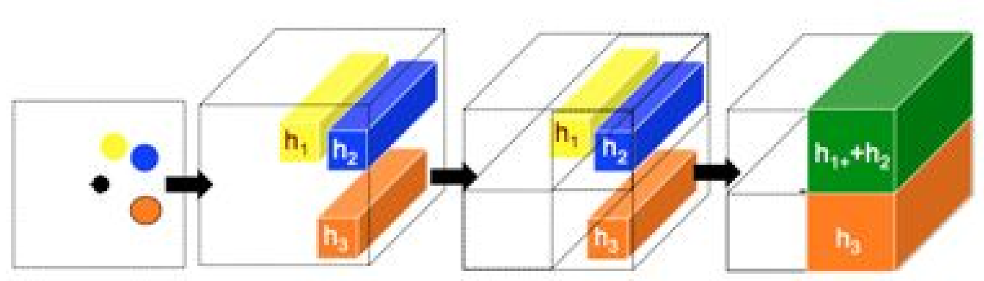

A. Alahi*, K. Goel*, V. Ramanathan, A. Robicquet, L. Fei-Fei, S. Savarese, Social LSTM: Human Trajectory Prediction in Crowded Spaces, in CVPR16

Multi-Class:

- 4 classes: Biker, Pedestrian, Skater, Cart
- Weights/parameters are shared between instances of same class

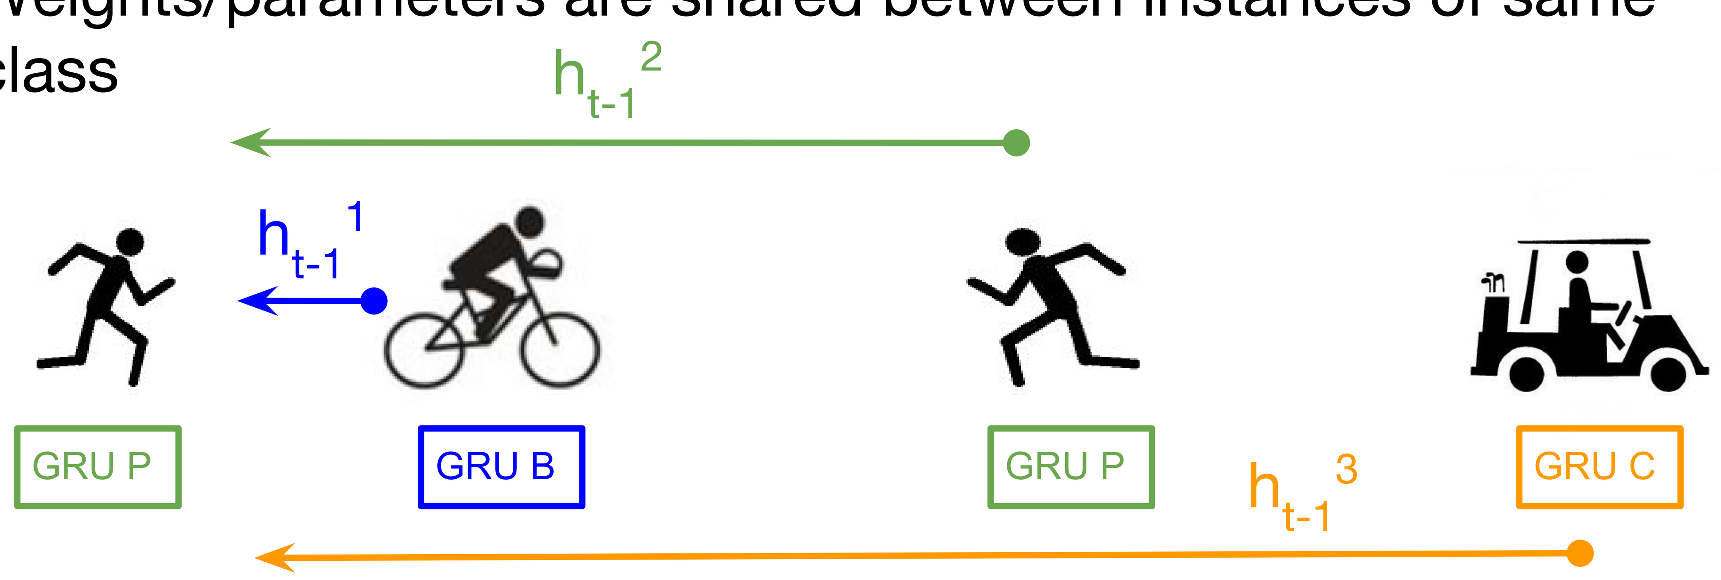

Data & Experimental Results (Path/Final)

	Biker	Pedestrian	Skater	Cart
Naive GRU	60.6 / 82.7	82.4 / 87.9	110.4 / 179.8	170.2 / 253.6
GRU + MC Pooling	50.5 / 66.3	52.5 / 65.9	123.2 / 184.8	144.6 / 245.7

Path Error = Distance at each predicted step between prediction and ground truth.
Final Displacement = Distance between the final points

- We used 5 pre-annotated Stanford Drone Dataset video-streams of the “The Circle of Death”
- total of 22 minutes of footage (with 1548 Bicyclists, 917 Pedestrians, 107 Skaters, and 132 carts)
  - Holdout: Training on first 3 scenes (17 min). Tested on 4th scene (5 min)
  - Training Set = 1307 Bikes, 831 Pedestrians, 91 Skaters, 111 Carts

Qualitative Analysis/Discussion of Results

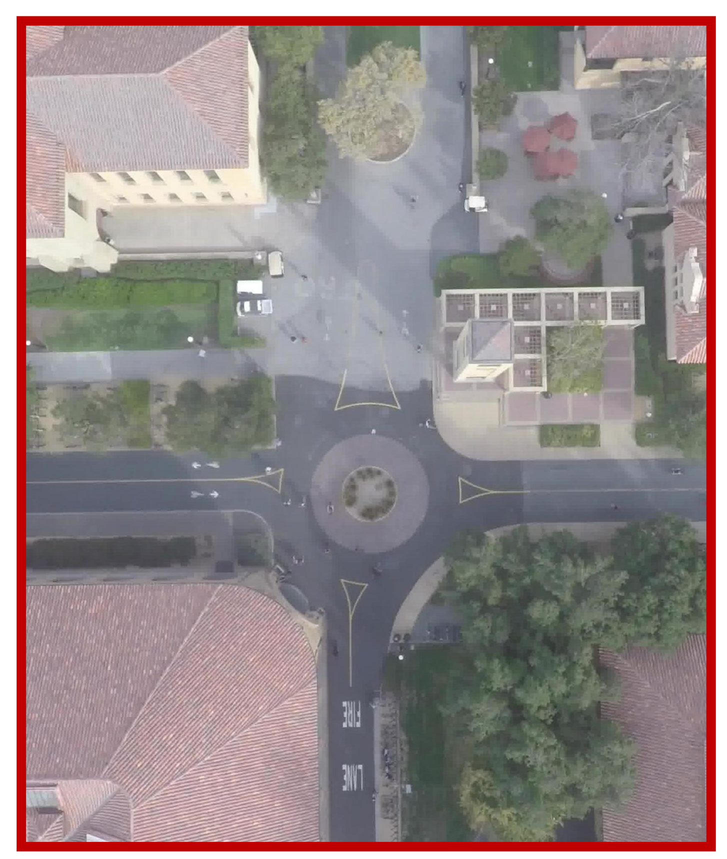

Circle of Death

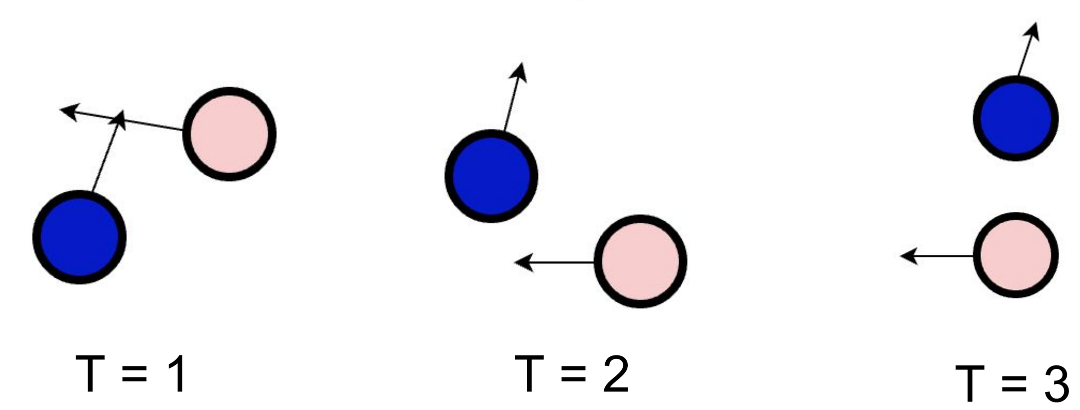

Visualization of Pedestrian Stopping for Bicycle.
Blue = Bicycle; Red = Pedestrian.

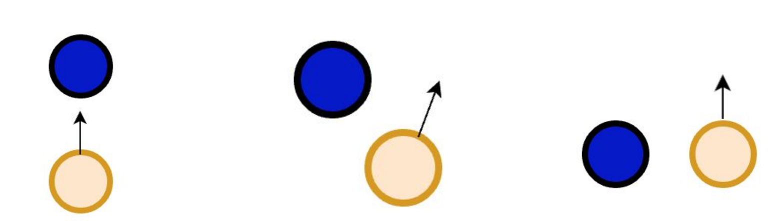

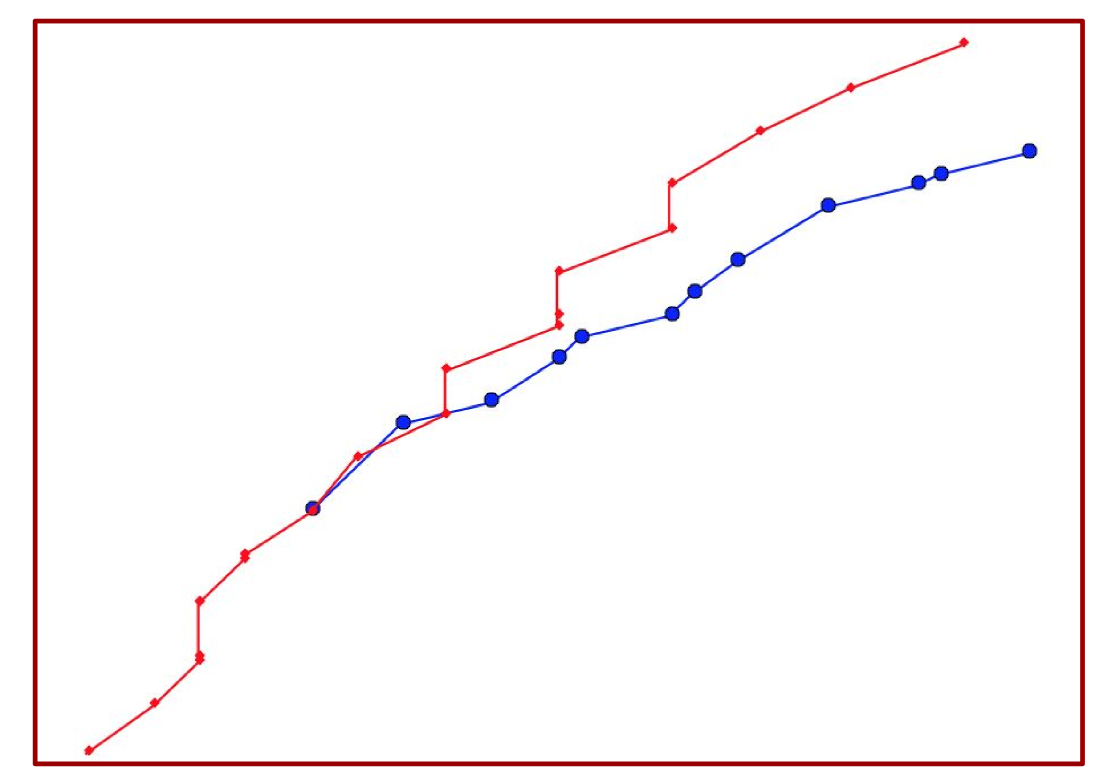

Sample Trajectory Prediction: Blue = Prediction

We were able to see the model map out non-linear paths and respond to other nearby agents. We saw good results on bikes and pedestrians but lacking results for the skaters and carts.

Future Work

We will first look to further substantiate our results through collecting a larger dataset, especially with respect to carts and skaters. We would also like to do more qualitative analysis since the errors don’t exactly capture how well objects avoid each other.