Week 9

DAVID JENSEN

JULY 16, 2010

What I was doing

- Looking at how drivers' actual actions compare to the prediction of the Intelligent Driver Model
  - If a vehicle is behaving unsafely, they will NOT fit the model
- Able to detect 82% of aggressive drivers with 1.22% false positives

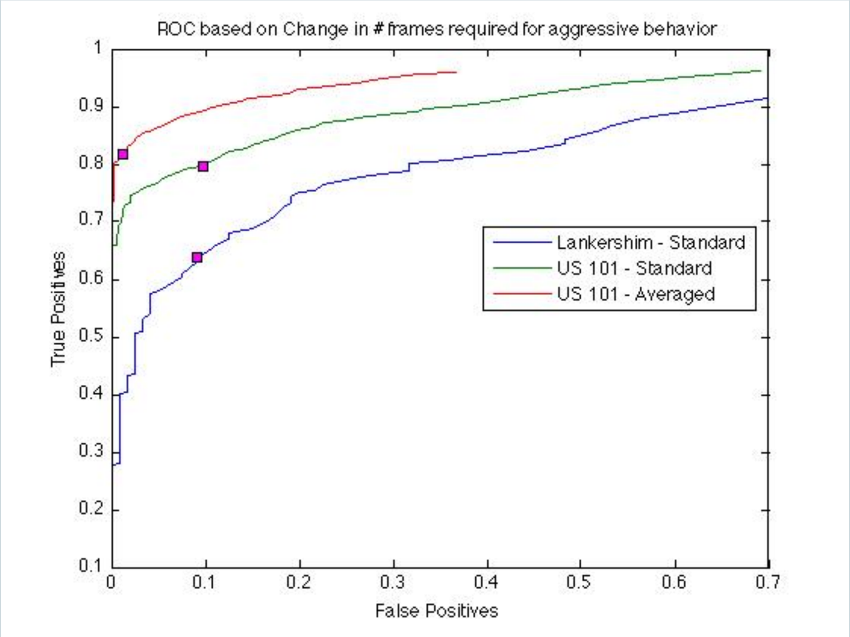

Weaknesses of the IDM

- The model can actually only detect vehicles that speed and/or tailgate
- The model cannot detect unsafe or excessive lane change behaviors

$$\dot{v}_\alpha = a^{(\alpha)}\left[1 - \left(\frac{v_\alpha}{v_0^{(\alpha)}}\right)^\delta - \left(\frac{s^*(v_\alpha, \Delta v_\alpha)}{s_\alpha}\right)^2\right].$$

$$s^*(v, \Delta v) = s_0^{(\alpha)} + s_1^{(\alpha)}\sqrt{\frac{v}{v_0^{(\alpha)}}} + T^\alpha v + \frac{v \Delta v}{2\sqrt{a^{(\alpha)} b^{(\alpha)}}}$$

- If a vehicle is identified by the IDM as aggressive, it is not obvious which behavior they are actually performing, thus making the model too general
  - Many vehicles also demonstrate multiple aggressive behaviors

Weaknesses of the IDM

- Using the model can provide fair results for general unsafe behavior: following a bit too closely or traveling just over the speed limit. If we would like to identify only vehicles that are following way too closely, the model gives too many false positives to make the results meaningful.
- The intelligent driver model is designed for microscopic simulations, but the parameters are very difficult to tune
- Using gradient descent to optimize the parameters yields impossible values:
  - T = -.8
  - a,b, imaginary

The New Approach

- We have decided to instead use a trained classifier to identify aggressive behaviors
- Using well-chosen features, we can use SVM to classify drivers into aggressive groups.
- Each vehicle has a feature vector based on their behavior over a slice of time (n frames)
  - If a vehicle is tracked for 100 frames with $n = 10$, that vehicle will have 10 feature vectors
  - This method accounts for the sparseness of aggressive behaviors
- Using the same feature vectors, a binary classifier is trained for each behavior we wish to identify, allowing for a single vehicle to have multiple behaviors

Features Used

- Currently, three features are used for each vehicle, although this will be expanded in time.
  - The average velocity of the vehicle over the slice (`mean(v(i:i+n))`) compared with the spatio-temporal mean of velocities in the area (`v_ave`)
    - `feature1 = mean(v(i:i+n-1)) – v_ave`
  - The mean of the squared error of the vehicle's acceleration over the slice and the IDM expected acceleration
    - `feature2 = mean(a(i:i+n-1)–IDM(i:i+n-1)^2)`
  - Average time behind the preceding vehicle
    - `feature3 = mean(T(i:i+n-1))`

Calculating the Ground Truth

- Currently, two behaviors are used: speeding and tailgating
- We are interested in extreme cases of aggressive behavior. The following are automatically determined based on vehicle trajectories.
  - Vehicles in excess of 10m/s of the average velocity of surrounding vehicles are considered speeding
  - Vehicles less than 1s behind the preceding vehicle
- Ideally, we would like to have several people go through the video by hand and identify those vehicles which they consider aggressive

Some Results

- The entire data set is divided into approx 119,000 slices
- LibSVM v 2.91 for Matlab is used to handle the training and testing
- Using the first 5,000 slices for both training and testing, the following results were obtained:
  - Speeding: 98.62% accuracy (4931/5000) classified correctly
    - 156 speeding slices
  - Tailgating: 96.46% accuracy (4823/5000) classified correctly
    - 358 tailgating slices