Functional Trajectory Forecasting and Consistency Guarantees for Self-Driving Cars in Social Settings

Scott Sussex*, Chiyu Dong**, John M. Dolan**

Harvard University*, Carnegie Mellon University**

Motivation

Self-driving cars will be required to navigate urban scenarios such as intersections, ramp merges, and lane changes. These are all highly social environments that require accurate estimates of the intentions of other cars. This requires models that can account for the interactions between multiple vehicles. Existing methods pose two limitations:

- Existing models represent a forecast as a series of points the car will reach, or simply as a categorical indicator of intention. Instead we desire a trajectory to be represented as a smooth function, where spatial coordinates change as a function of time.
- As the environment changes over time, a forecast will also change, potentially erratically. Existing methods don't provide guarantees for the consistency of a forecast over time.

Notation

$x \triangleq$	Training data features of n trajectories
$b \triangleq$	The ground truth label for training trajectories. Contains d points at times along each of the n trajectories.
$K \triangleq$	Kernel function between feature inputs. When no inputs given, assume the matrix $K(x, x)$
$k \triangleq$	Kernel function between two times.
$\eta_i \triangleq$	The model's estimated trajectory for the ith training input. This is a function of time.
$D \triangleq$	A derivative operator for k.
$\alpha \triangleq$	Learned weights.
$\lambda \triangleq$	Hyper-parameter to weight MSE against complexity

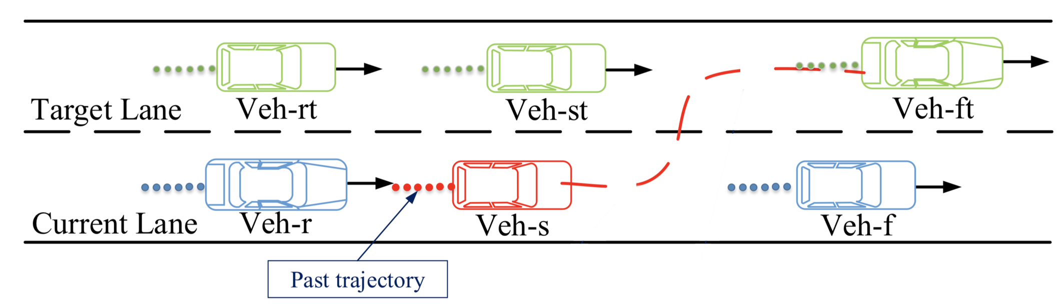

Method

Our approach minimizes a cost function, i.e

$$min \frac{1}{nd} \sum_{i,j} (b_i(t_j) - \eta_i(t_j))^2 + \lambda R$$

where R is a regularization term and $\eta_i(t) = k(t,).\alpha.K(x_i, x)$.

Possible values of R, where $f(x') = \alpha.K(x', x)$:

$\|\|f\|\|_H^2$	Naive application of typical RKHS regularizer. Not intuitive for a trajectory output. Representer theorem holds
$\Sigma_i \|\|D\eta_i\|\|_E^2$	Designed for a functional output. Intuitive definition of complexity. Representer theorem does not hold.

In either case we show that one can derive a consistency bound on what η will be if forecasted time T into the future by assuming Lipschitz continuity on the model input.

$$||X_1, X_2|| \leq CT$$

where C is a constant.

Experimental Results

We use lane-change scenarios extracted from the NGSIM dataset to evaluate our method with the use of different regularizers. For input x we use the previous trajectories of the host and all surrounding vehicles.

Result Di refers to the result of using our novel regularizer with D representing the ith derivative. "Base" represents the traditional regularizer. Direction y represents motion parallel to the lane dividers.

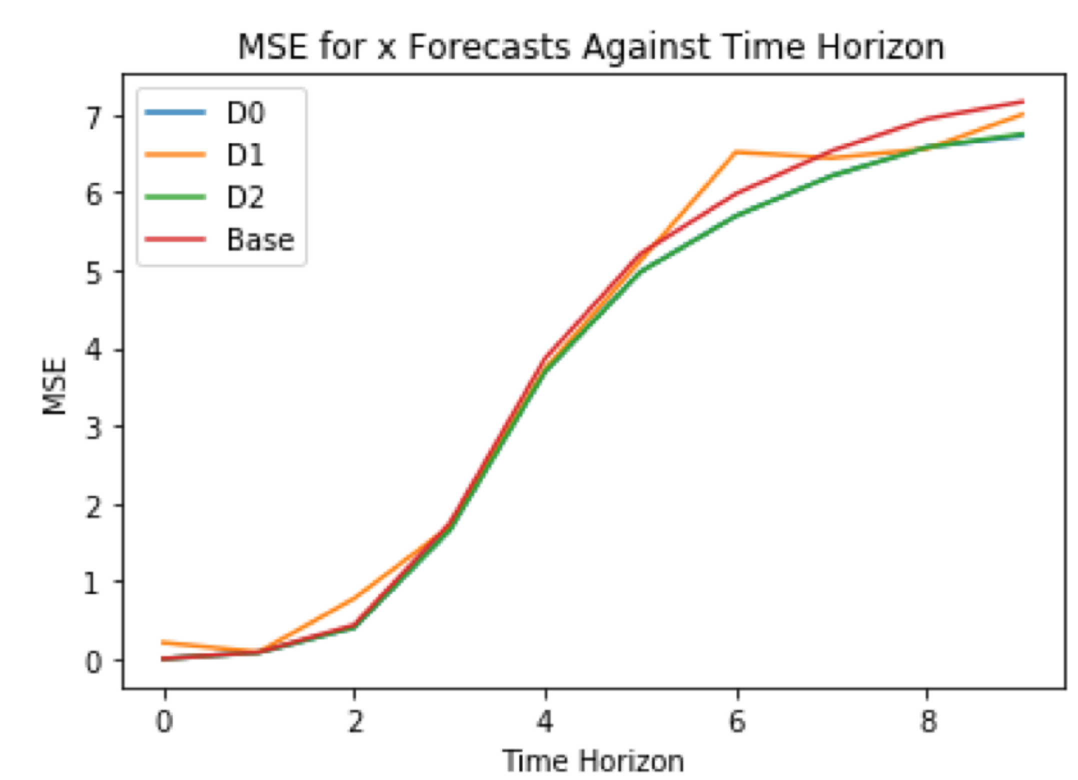

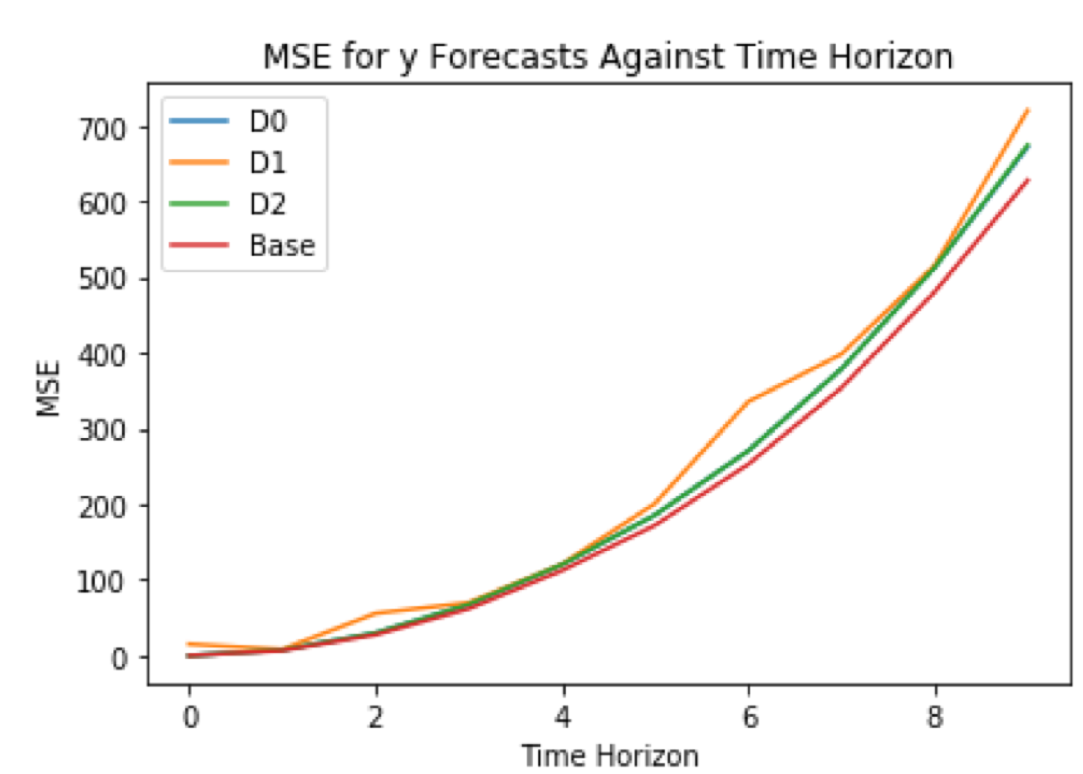

In general, D0, D2 and "Base" all perform competitively.

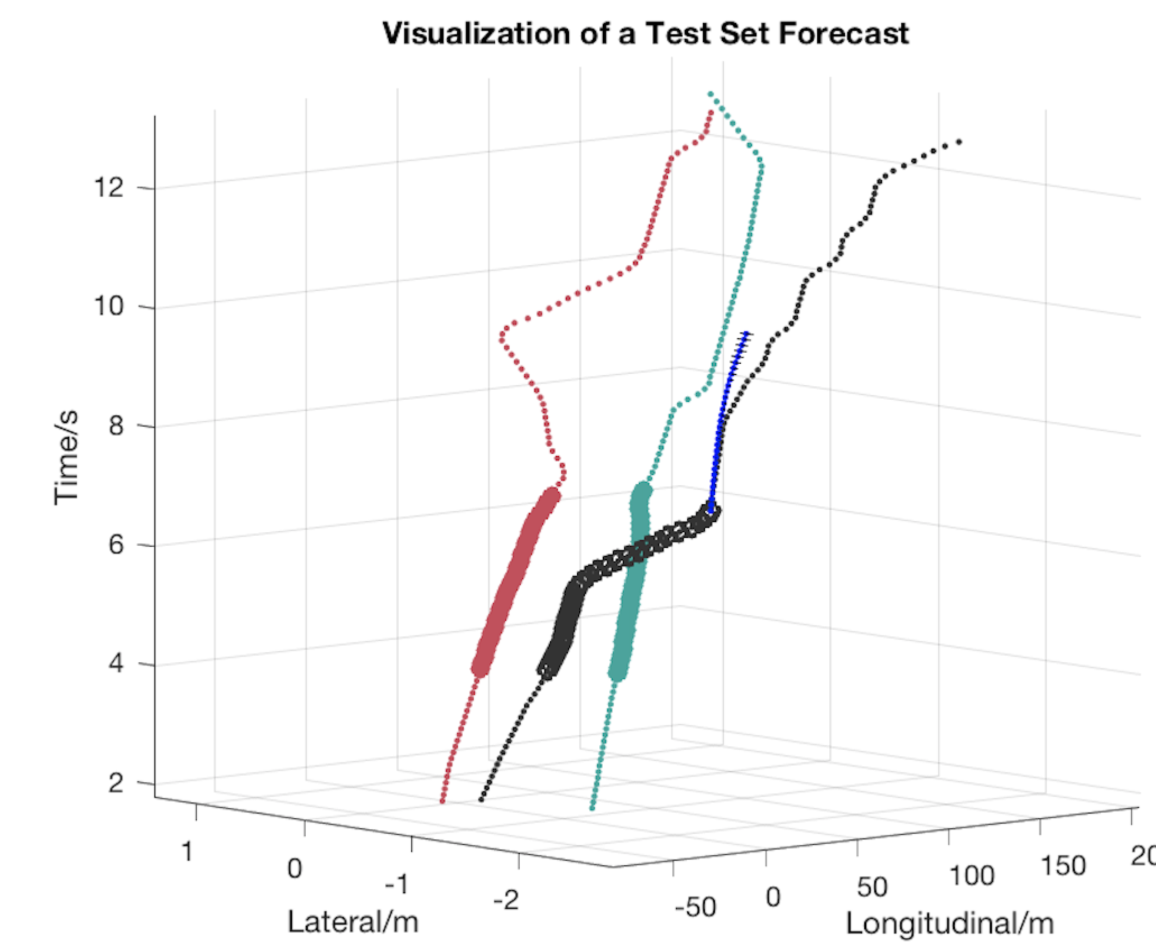

Thick regions are the host car. red and green are relevant cars. Blue is the forecast with a 3 second time horizon. Black bars are the bound on a trajectory forecast 0.5 seconds into the future.

Conclusion

We demonstrate how a non-parametric method can be used to output continuous trajectories for forecasting another vehicles motion, considering the interactions with other vehicles. We compare different regularizers for this setting and present a novel method for giving consistency bounds. We demonstrate the feasibility of the method and consistency bound in experiments.

Acknowledgements

The authors would like to thank Rachel Burcin and Ziqi Guo for their work on the RISS program, which is the excellent summer research program at CMU where this work was completed.

THE ROBOTICS INSTITUTE