Improving Movement Prediction of Traffic Actors using Off-road Loss and Bias Mitigation

Matthew Niedoba, Henggang Cui, Kevin Luo,
Darshan Hegde, Fang-Chieh Chou, Nemanja Djuric

Uber ATG

Background and motivation

RasterNet: trajectory prediction with rasterization and CNN

- RasterNet [1, 2] rasterizes the HD-map and dynamic object states into an object-centric raster input image for each object.
- A CNN is then applied to predict the future positions of the target object.
- The model predicts multiple trajectories for each object with per-waypoint Gaussian uncertainties.

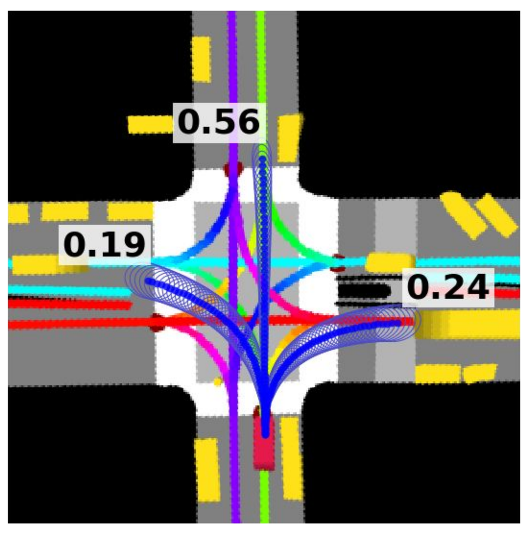

Fig 1. RasterNet input raster and output trajectories

Unrealistic off-road predictions

- Existing approaches use only the L2 distance to the ground-truth as the supervision for every waypoint.
- However, such loss might not have enough penalty on some bad unrealistic predictions (e.g., an off-road actor prediction).

Biased distribution of action categories

- Traffic datasets used for training and evaluating a trajectory prediction model are often dominated by straight-going actors.
- Models trained and evaluated on such datasets often have suboptimal performance on turning actors.

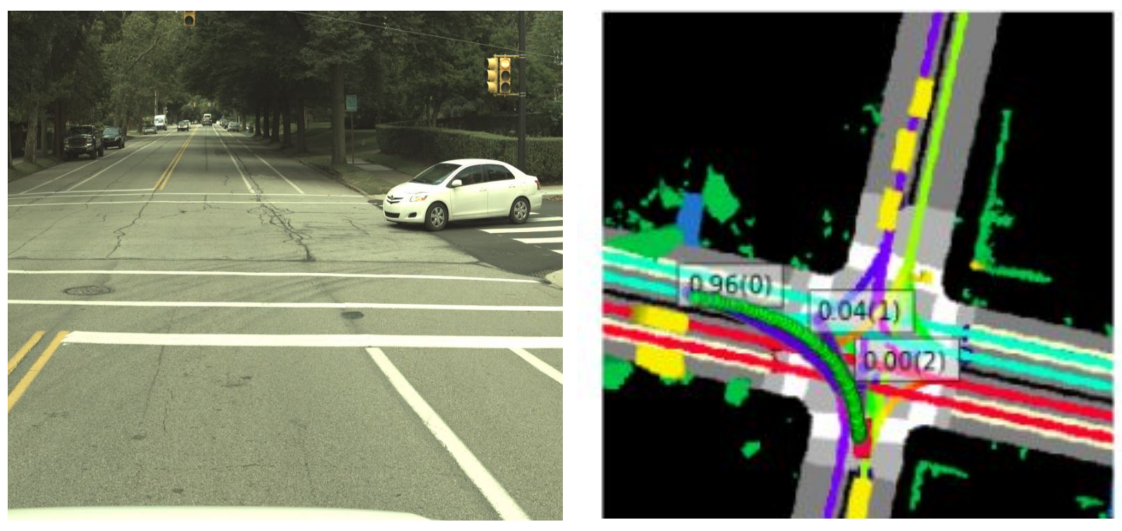

Fig 2. Incorrect off-road prediction for a turning actor caused SDV to hard-brake

Contributions

- In this work, we propose to improves the traffic actor predictions with two novel methods: *action category upweighting* and two *off-road losses*.

Mitigating biases with action category upweighting

- We determine the action category for each data sample based on the heading difference θ between the first and last ground-truth waypoints.
  - Straight if $|\theta| < 20°$
  - Left turn if $-135° < \theta < -20°$
  - Right turn if $20° < \theta < 135°$
  - Sharp turn if $|\theta| > 135°$
- We observe significant bias in three different autonomous driving datasets
- To mitigate this issue, we propose to upweight the left-, right-, and sharp-turn samples in the loss and slice the prediction error metrics based on the action category.

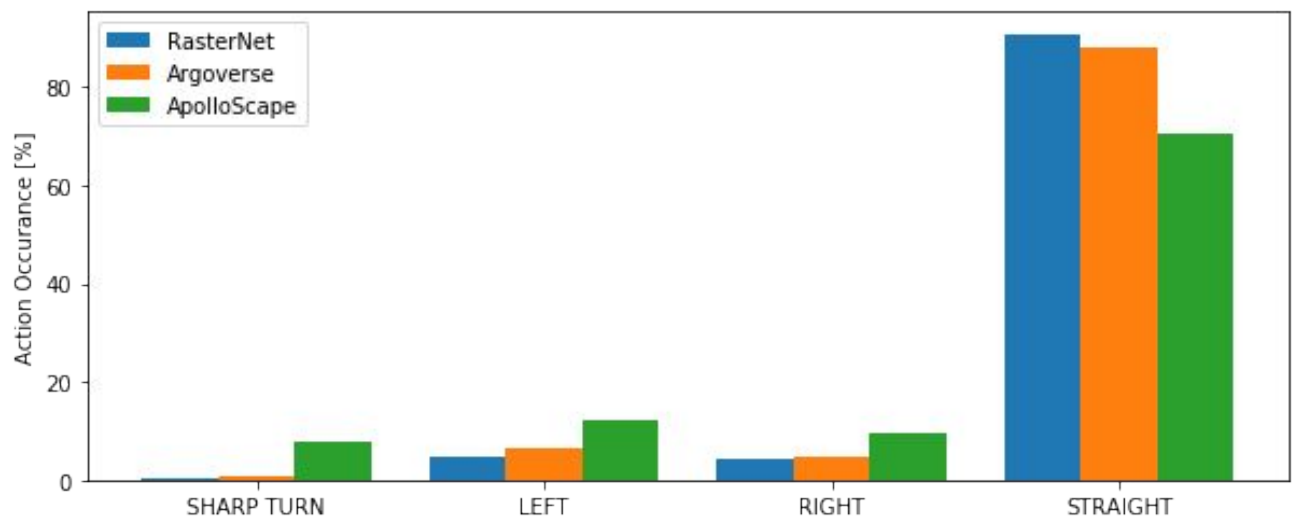

Fig 3. Action category distribution of three different datasets

Mitigating unrealistic off-road predictions with off-road losses

Determine drivable regions

- We first propose a method to identify the *drivable region* for each actor of interest.
- A most straightforward approach would be to identify all road surface as the drivable region. However, such naive approach would include lanes that the actor is unlikely to be driving on (e.g., opposing lane for a high-speed actor) in the drivable region.
- In this work, we perform lane path association and scoring for each actor, and mark the lanes with high occupancy scores for the actor as its drivable region.
- The drivable region information will be used by our two novel off-road losses presented below.

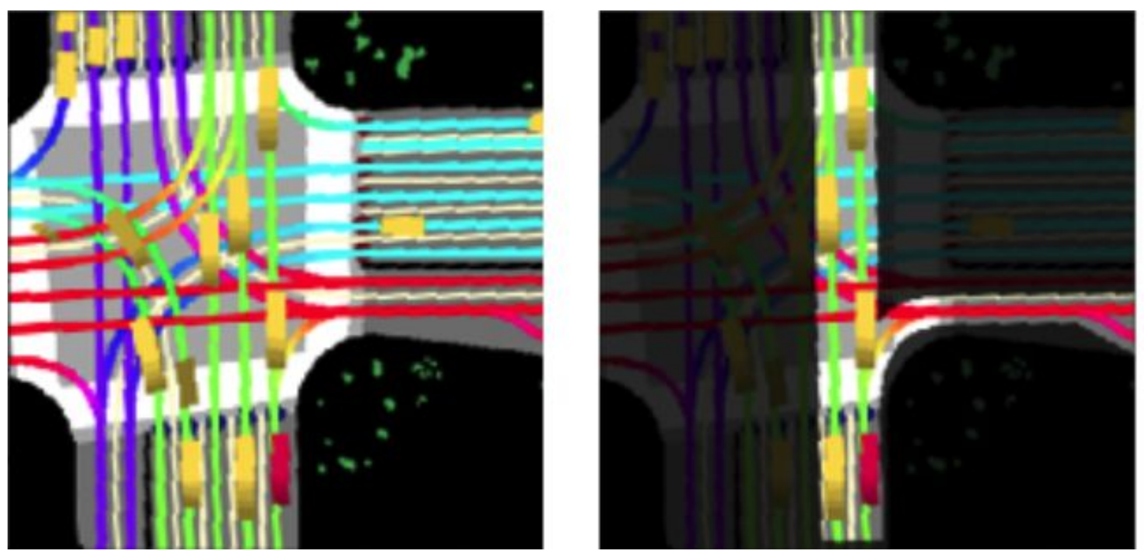

Fig 4. Map raster and the drivable region for the actor of interest

Off-road false positive loss

- We propose an *off-road false positive loss* that upweights the traditional L2 distance loss for the off-road *false positive* predictions.
- An off-road false positive prediction is defined as a prediction that is outside the drivable region while the corresponding ground-truth is inside the drivable region.
- This loss term helps the training to have more penalties on the off-road false positive predictions.

$$\mathcal{L}_{ofp} = \frac{\lambda_{ofp}}{T} \sum_{t=1}^{T} \mathbb{1}_{ofp;\mathcal{M}}(\hat{p}_t, g_t) \, \|\hat{p}_t, g_t\|_2$$

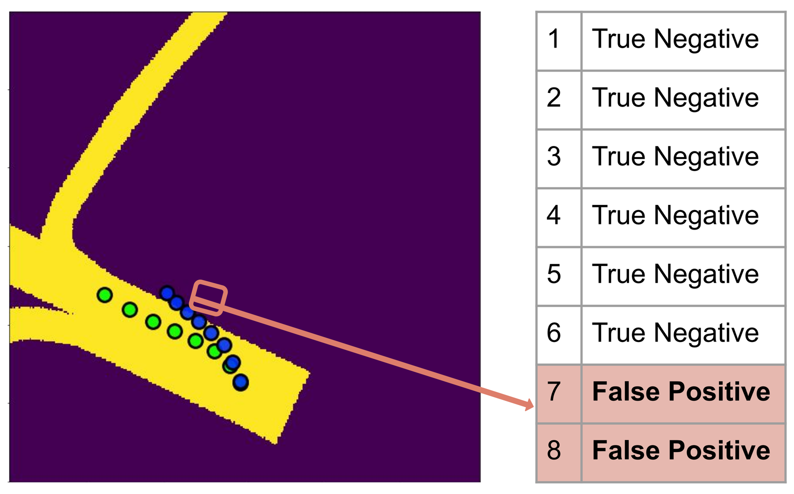

Fig 5. Offroad false positive

Off-road distance loss

- We propose another *off-road distance loss* which is computed as the L2 distance from the prediction p to its corresponding *nearest on-road point* $r(p)$.
- The mapping from each prediction point to its nearest on-road point is pre-computed and stored as labels for each sample.
- This loss term helps drag off-road predictions back to the road.

$$\mathcal{L}_{od} = \frac{\lambda_{od}}{T} \sum_{t=1}^{T} \|\hat{p}_t, r_{\mathcal{M}}(\hat{p}_t)\|_2$$

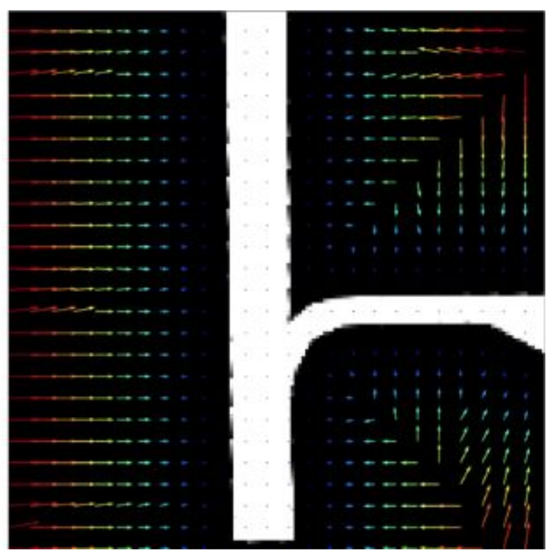

Fig 6. Nearest on-road point mapping

Evaluation results and conclusion

Experimental setups

- Models to compare
  - RasterNet
  - RasterNet + action_upweight (2x upweighting for turning samples)
  - RasterNet + action_upweight + offroad_fp_loss (λ_{ofp}= 5) + offroad_dist_loss (λ_{od}= 0.25)
- Metrics
  - L2 error, cross-track error (CT), along-track error (AT), and off-road distance (OD)

	Overall				Straight				Turning			
Model	ℓ_2	CT	AT	OD	ℓ_2	CT	AT	OD	ℓ_2	CT	AT	OD
RasterNet	**1.52**	**0.34**	**1.39**	0.05	**1.46**	**0.28**	**1.37**	0.04	2.15	0.85	1.69	0.10
+ action upweight	**1.52**	**0.34**	**1.39**	0.05	1.47	0.29	**1.37**	0.04	**2.06**	**0.82**	**1.64**	0.11
+ off-road losses	1.58	0.37	1.44	**0.04**	1.52	0.32	1.41	**0.03**	2.13	0.84	1.68	**0.08**

Tab 1. Prediction error metrics @3s sliced by action category

	Overall				Straight				Turning			
Model	ℓ_2	CT	AT	OD	ℓ_2	CT	AT	OD	ℓ_2	CT	AT	OD
RasterNet	**4.79**	**0.90**	4.42	**0.14**	**4.54**	**0.61**	**4.37**	**0.11**	7.50	3.12	5.64	0.39
+ action upweight	**4.79**	0.93	**4.40**	0.15	4.59	0.67	4.39	**0.11**	**6.97**	2.90	**5.29**	0.44
+ off-road losses	4.91	0.96	4.52	**0.14**	4.72	0.73	4.49	**0.11**	7.03	**2.75**	5.47	**0.35**

Tab 2. Prediction error metrics @6s sliced by action category

Conclusions

- Action upweighting improves the predictions for turning actors.
- Adding off-road losses reduce the off-road distance error.
- With the action upweighting and off-road loss methods combined, we achieve the best predictions for the turning actors.

Reference

[1] N. Djuric, V. Radosavljevic, H. Cui, T. Nguyen, F.-C. Chou, T.-H. Lin, and J. Schneider. Short-term motion prediction of traffic actors for autonomous driving using deep convolutional networks. *arXiv preprint arXiv:1808.05819, 2018*.

[2] H. Cui, V. Radosavljevic, F.-C. Chou, T.-H. Lin, T. Nguyen, T.-K. Huang, J. Schneider, N. Djuric. Multimodal trajectory predictions for autonomous driving using deep convolutional networks. ICRA, 2019.