PRECOG: PREdiction Conditioned On Goals in Visual Multi-Agent Settings

Nicholas Rhinehart[1], Rowan McAllister[2], Kris Kitani[1], Sergey Levine[2]

[1]Carnegie Mellon University; [2]University of California, Berkeley

https://arxiv.org/abs/1905.01296 • https://sites.google.com/view/precog

Contributions

- ✓ State-of-the-art multi-agent forecasting
- ✓ Goal-conditioned multi-agent forecasting
- ✓ Multi-agent imitative planning objective

Multi-agent Forecasting

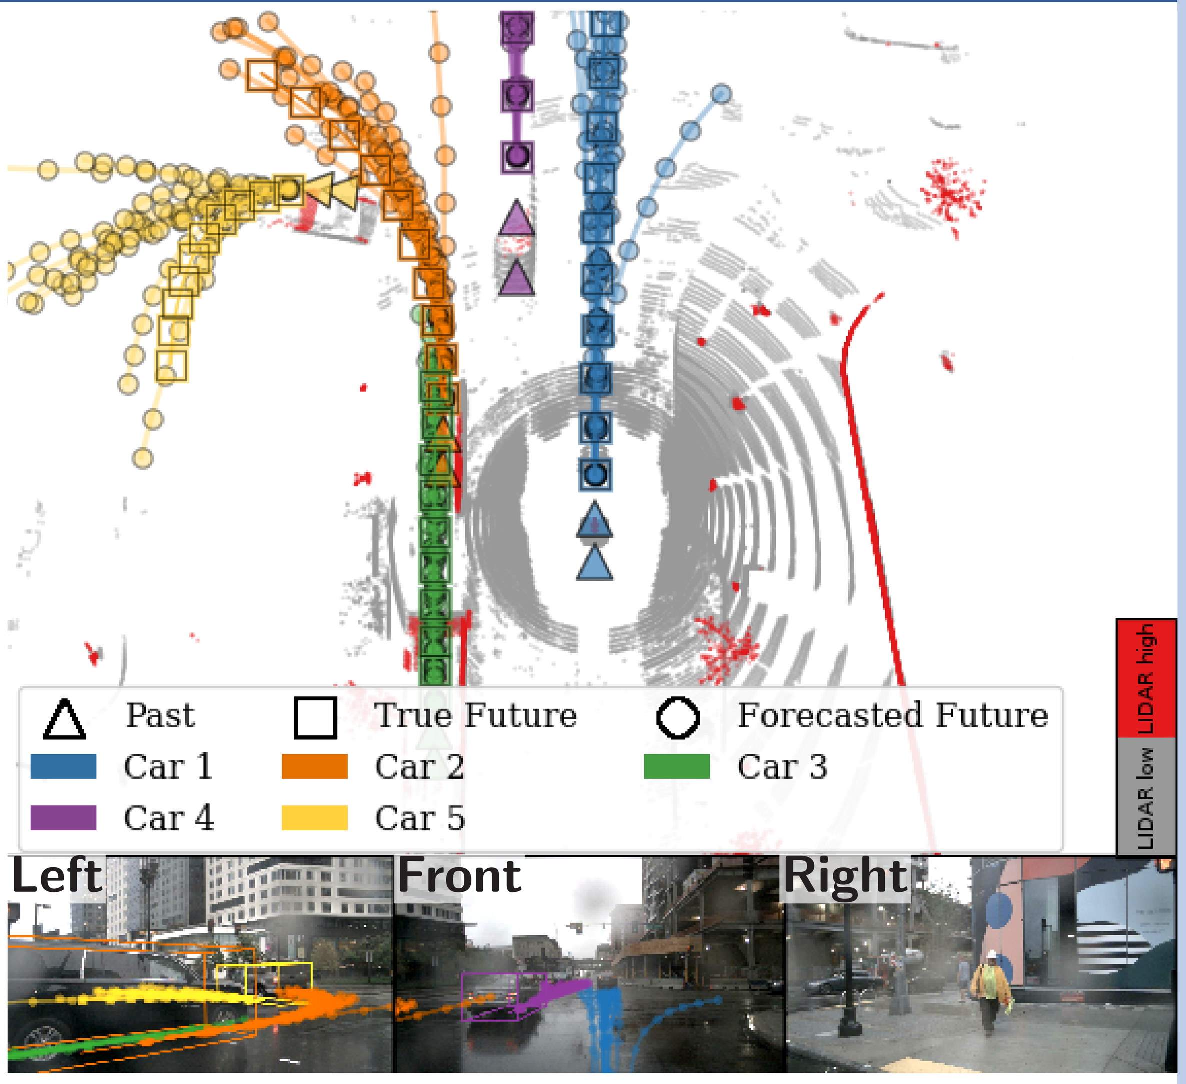

Goal-Conditioned Forecasting

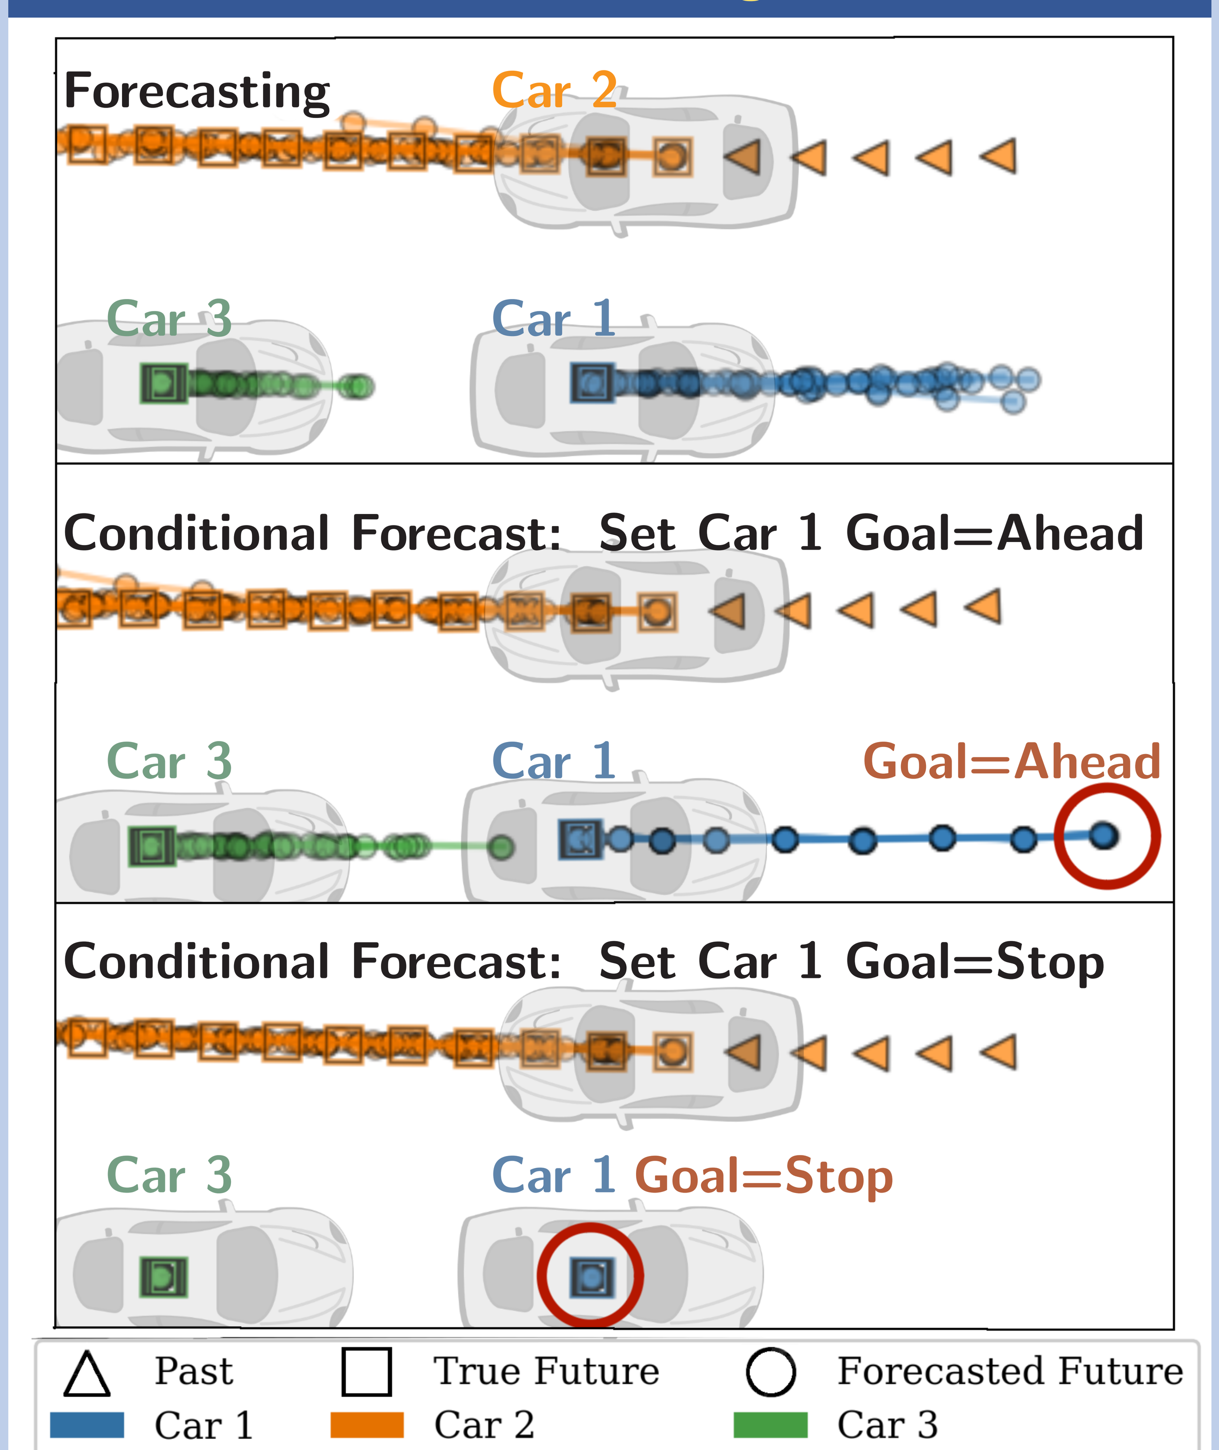

Prior Work: R2P2 (ReparameteRized Pushforward Policy) [3]

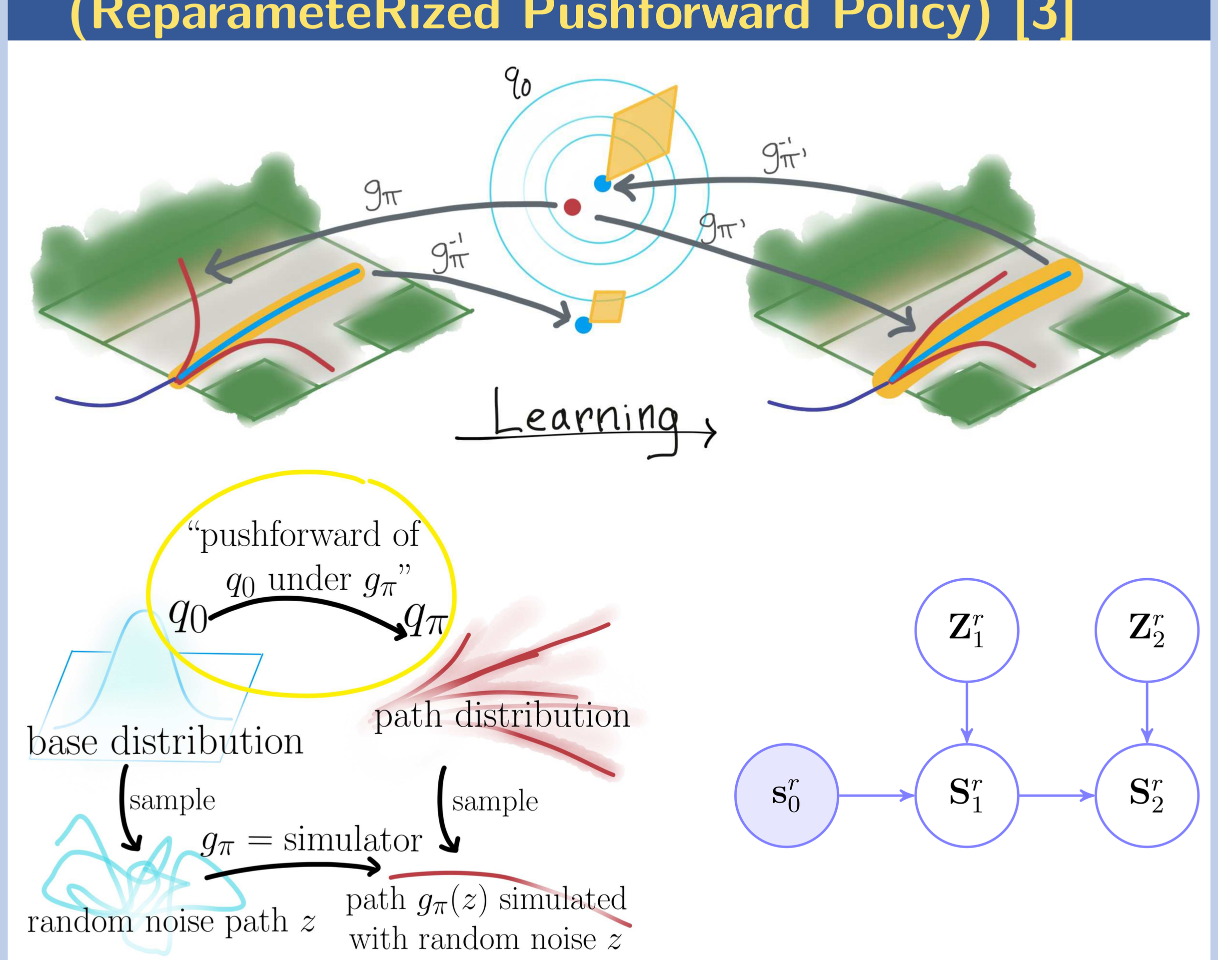

Prior Work: Imitative Models [4]

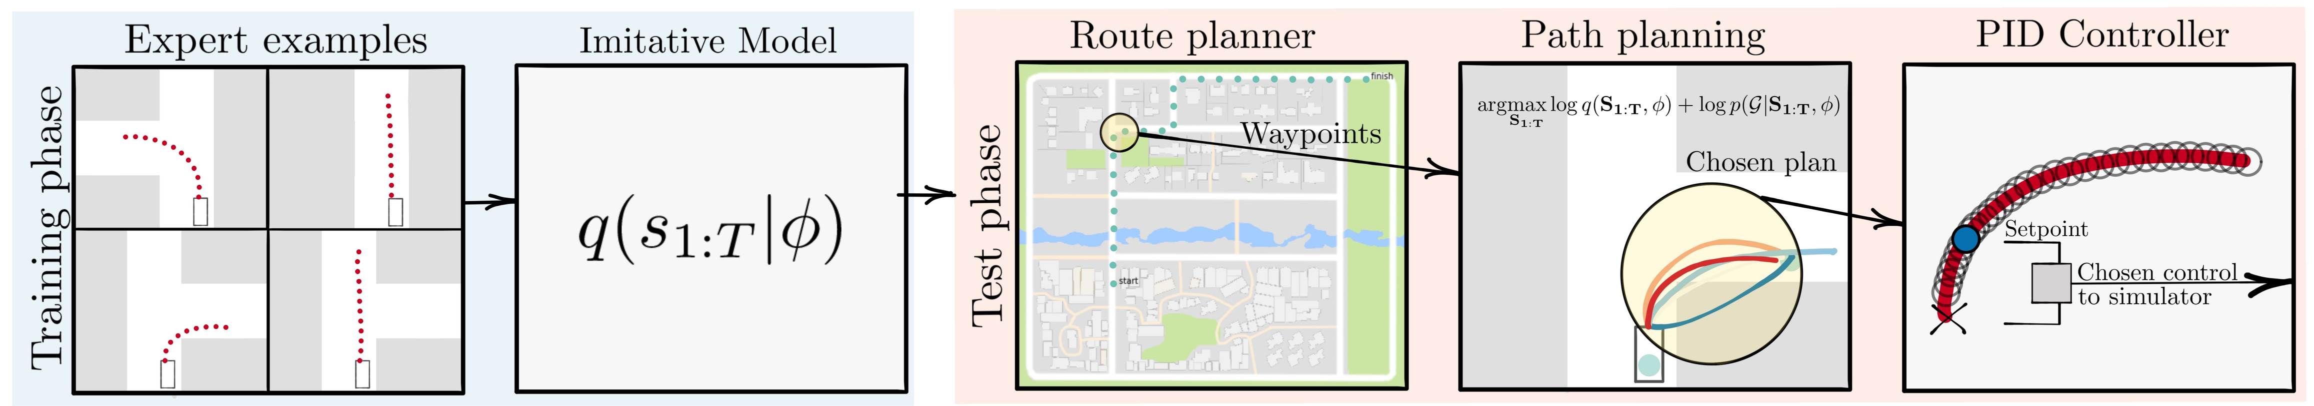

Plan what expert "would have likely done" to reach goal.

$$\max_{s_{1:T}} \log p(s_{1:T}|\mathcal{G},\phi) = \max_{s_{1:T}} \log q(s_{1:T}|\phi) + \log p(\mathcal{G}|s_{1:T},\phi) - \log p(\mathcal{G}|\phi)$$

$$= \max_{s_{1:T}} \log \underbrace{q(s_{1:T}|\phi)}_{\text{imitation prior}} + \log \underbrace{p(\mathcal{G}|s_{1:T},\phi)}_{\text{goal likelihood}}$$

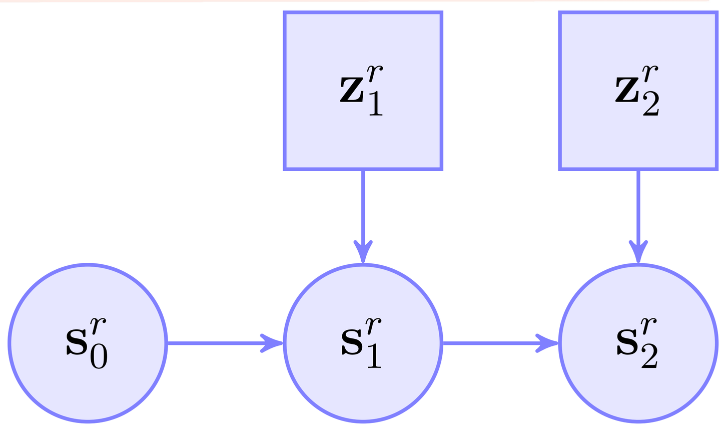

Model of Agent Interaction

ESP (Estimating Social-forecasting Probability)

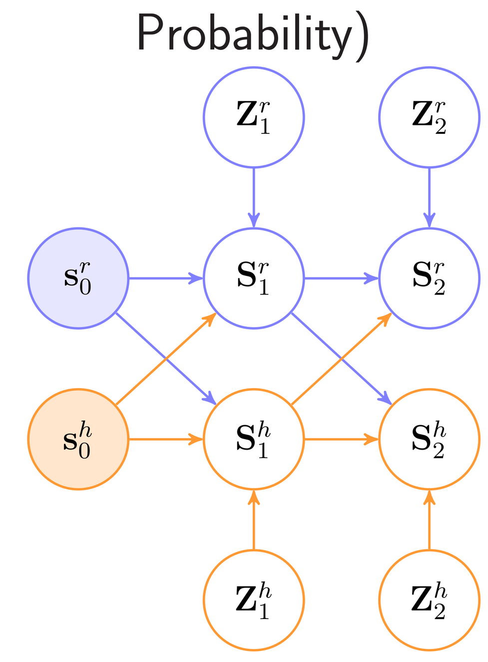

PRECOG (PREdiction Conditioned on Goals)

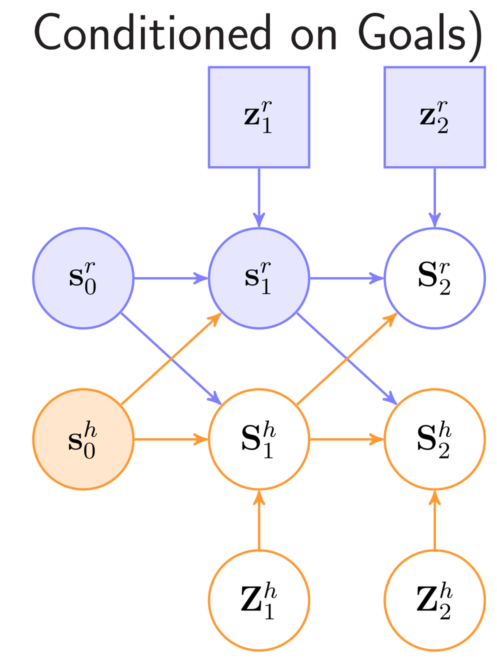

Goal-Conditioned Forecasting

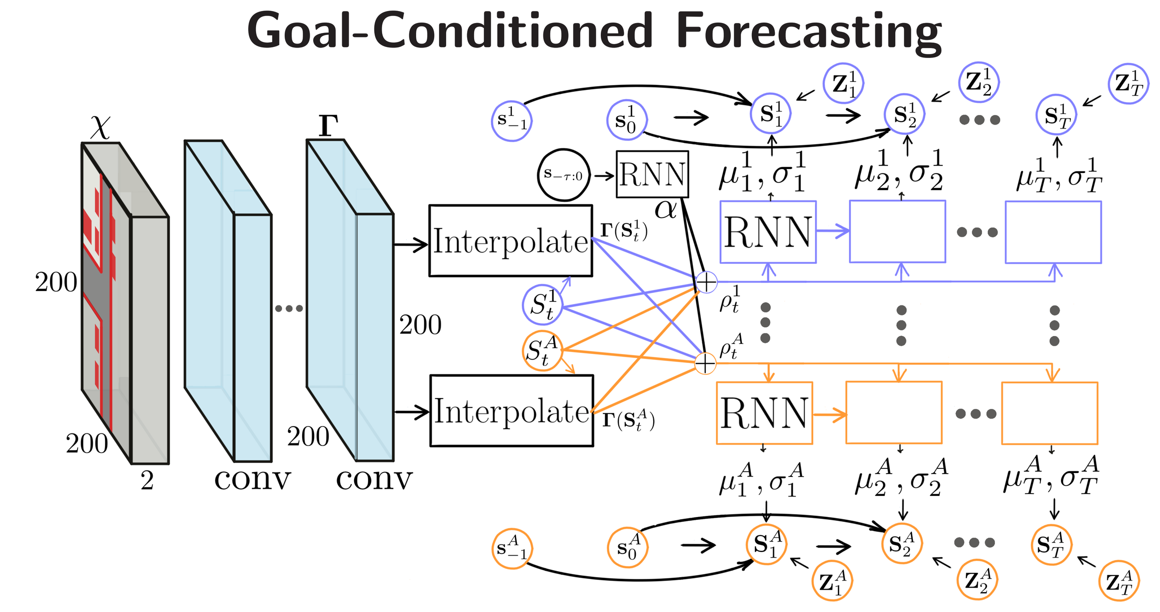

Multi-Agent Planning Objective

$$z^{robot*} = \operatorname*{argmax}_{z^{\text{robot}}} \mathbb{E}_{Z^{\text{human}}}[\log p(S|\mathcal{G},\phi)]$$

$$= \operatorname*{argmax}_{z^{\text{robot}}} \mathbb{E}_{Z^{\text{human}}}[\log(q(S|\phi)p(\mathcal{G}|S,\phi))] - \log p(\mathcal{G}|\phi)$$

$$= \operatorname*{argmax}_{z^{\text{robot}}} \mathbb{E}_{Z^{\text{human}}}[\log \underbrace{q(f(Z)|\phi)}_{\text{multi-agent prior}} + \log \underbrace{p(\mathcal{G}|f(Z),\phi)}_{\text{goal likelihood}}]$$

Forecasting Evaluation

Approach		Test $\hat{m}_{K=12}$	Test $\hat{e}$	Test $\hat{m}_{K=12}$	Test $\hat{e}$	Test $\hat{m}_{K=12}$	Test $\hat{e}$
CARLA Town02 Test	2 agents	3 agents		4 agents		5 agents	
DESIRE	[2]	1.587 ± 0.020	–	2.234 ± 0.023	–	2.422 ± 0.017	–
SocialGAN	[1]	0.812 ± 0.013	–	1.098 ± 0.014	–	1.141 ± 0.015	–
R2P2-MA	[3]	0.387 ± 0.008	0.645 ± 0.002	0.690 ± 0.009	0.621 ± 0.002	0.770 ± 0.008	$\mathbf{0.618 \pm 0.002}$
Ours: ESP, no LIDAR		0.719 ± 0.011	0.640 ± 0.002	0.919 ± 0.011	0.650 ± 0.002	1.102 ± 0.011	0.652 ± 0.002
Ours: ESP		$\mathbf{0.385 \pm 0.007}$	$\mathbf{0.585 \pm 0.002}$	$\mathbf{0.509 \pm 0.007}$	$\mathbf{0.599 \pm 0.002}$	$\mathbf{0.675 \pm 0.007}$	0.630 ± 0.001
nuScenes Test	2 agents	3 agents		4 agents		5 agents	
DESIRE	[2]	4.421 ± 0.130	–	5.957 ± 0.162	–	6.575 ± 0.198	–
SocialGAN	[1]	3.033 ± 0.110	–	3.484 ± 0.129	–	3.871 ± 0.148	–
R2P2-MA	[3]	2.055 ± 0.093	0.989 ± 0.008	2.695 ± 0.100	1.020 ± 0.011	3.311 ± 0.166	$\mathbf{1.050 \pm 0.012}$
Ours: ESP, no LIDAR		2.240 ± 0.084	$\mathbf{0.955 \pm 0.008}$	3.201 ± 0.113	1.033 ± 0.012	3.442 ± 0.139	1.107 ± 0.018
Ours: ESP		1.705 ± 0.089	1.018 ± 0.011	2.547 ± 0.095	1.053 ± 0.015	3.266 ± 0.155	1.082 ± 0.013
Ours: ESP+Road		$\mathbf{1.505 \pm 0.070}$	1.016 ± 0.011	$\mathbf{2.360 \pm 0.093}$	$\mathbf{1.013 \pm 0.012}$	$\mathbf{2.892 \pm 0.162}$	1.114 ± 0.024

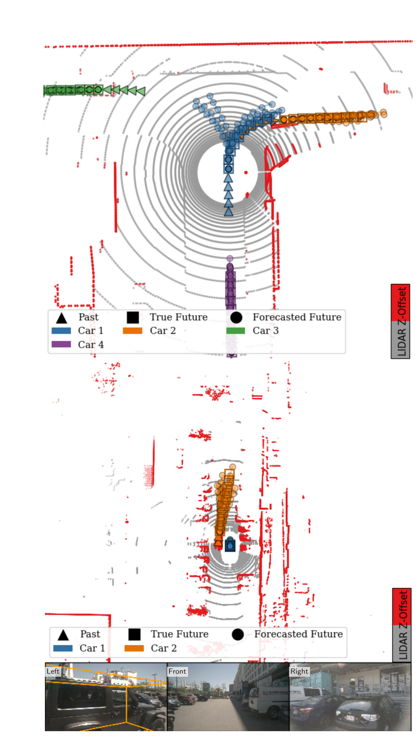

Goal-Conditioned Forecasting Evaluation

Data	Approach	Test $\hat{m}_{K=12}$	Test $\hat{m}^{a=1}_{K=12}$	Test $\hat{m}^{a=2}_{K=12}$	Test $\hat{m}^{a=3}_{K=12}$	Test $\hat{m}^{a=4}_{K=12}$	Test $\hat{m}^{a=5}_{K=12}$
CARLA 3	ESP	0.426 ± 0.013	0.204 ± 0.009	0.556 ± 0.027	0.519 ± 0.021	–	–
	PRECOG	$\mathbf{0.355 \pm 0.012}$	$\mathbf{0.052 \pm 0.003}$	$\mathbf{0.519 \pm 0.025}$	$\mathbf{0.493 \pm 0.020}$	–	–
CARLA 5	ESP	0.718 ± 0.012	0.340 ± 0.011	0.759 ± 0.024	0.809 ± 0.025	0.851 ± 0.023	0.828 ± 0.024
	PRECOG	$\mathbf{0.640 \pm 0.011}$	$\mathbf{0.066 \pm 0.003}$	$\mathbf{0.741 \pm 0.024}$	$\mathbf{0.790 \pm 0.024}$	$\mathbf{0.804 \pm 0.022}$	$\mathbf{0.801 \pm 0.024}$
nuScenes 3	ESP	1.511 ± 0.077	1.128 ± 0.061	1.543 ± 0.118	1.862 ± 0.147	–	–
	PRECOG	$\mathbf{1.016 \pm 0.062}$	$\mathbf{0.121 \pm 0.005}$	$\mathbf{1.320 \pm 0.105}$	$\mathbf{1.606 \pm 0.122}$	–	–
nuScenes 5	ESP	2.921 ± 0.175	1.861 ± 0.109	2.369 ± 0.188	2.812 ± 0.188	3.201 ± 0.254	4.363 ± 0.652
	PRECOG	$\mathbf{2.508 \pm 0.152}$	$\mathbf{0.149 \pm 0.021}$	$\mathbf{2.324 \pm 0.187}$	$\mathbf{2.654 \pm 0.190}$	$\mathbf{3.157 \pm 0.273}$	$\mathbf{4.254 \pm 0.586}$

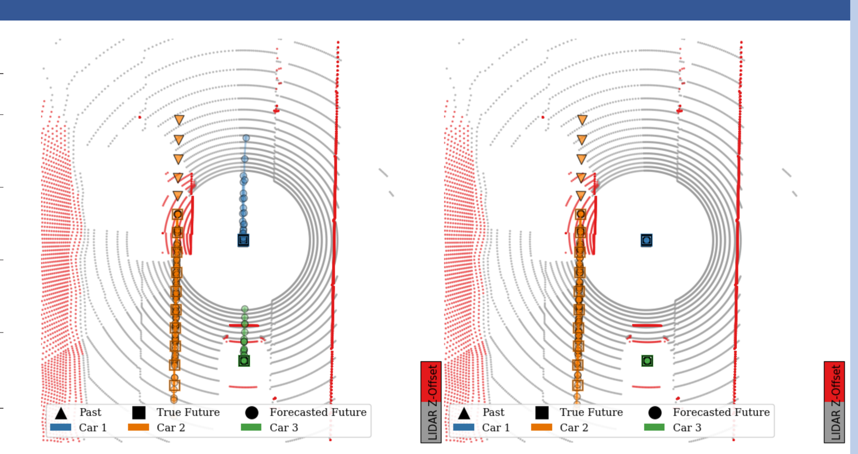

References

[1] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi. Social GAN: Socially acceptable trajectories with generative adversarial networks. In *Computer Vision and Pattern Recognition (CVPR)*, 2018.

[2] N. Lee, W. Choi, P. Vernaza, C. B. Choy, P. H. Torr, and M. Chandraker. DESIRE: Distant future prediction in dynamic scenes with interacting agents. In *Computer Vision and Pattern Recognition (CVPR)*, pages 336–345, 2017.

[3] N. Rhinehart, K. M. Kitani, and P. Vernaza. R2P2: A reparameterized pushforward policy for diverse, precise generative path forecasting. In *European Conference on Computer Vision (ECCV)*, September 2018.

[4] N. Rhinehart, R. McAllister, K. Kitani, and S. Levine. Deep imitative models for flexible inference, planning, and control. *arXiv preprint arXiv:1810.06544*, 2018.

{nrhineha,kkitani}@cs.cmu.edu {rmcallister,svlevine}@berkeley.edu