The 3 R's of Learning for Autonomy

Adrien Gaidon

Head of Machine Learning Research
Toyota Research Institute (TRI), CA, USA

adriengaidon.com

Machine Learning Research @ TRI

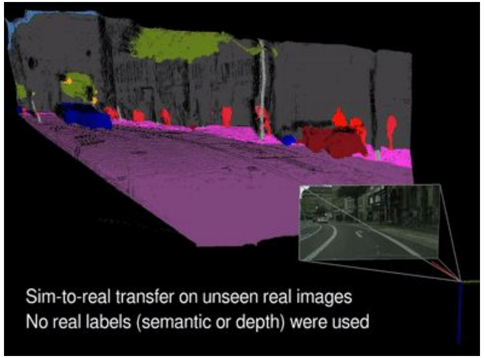

Scene Understanding

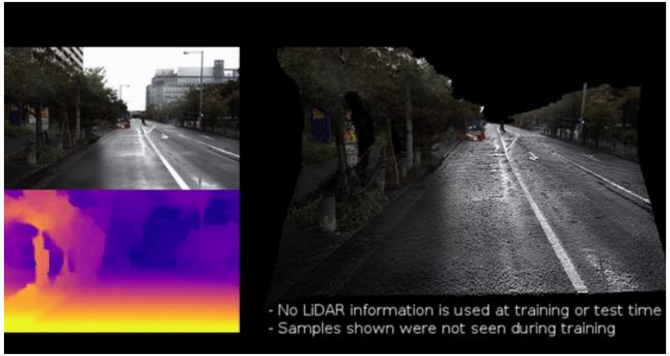

Self-Supervised Learning

Simulation & Auto-labeling

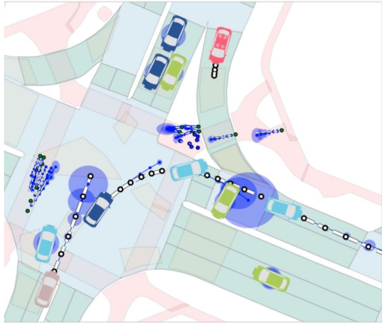

Behavior Modeling & Prediction

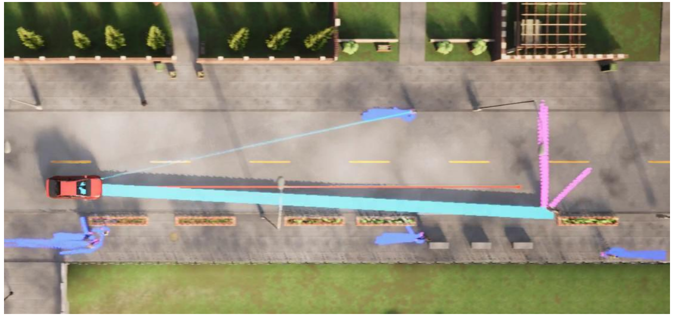

Learning for Planning & Control

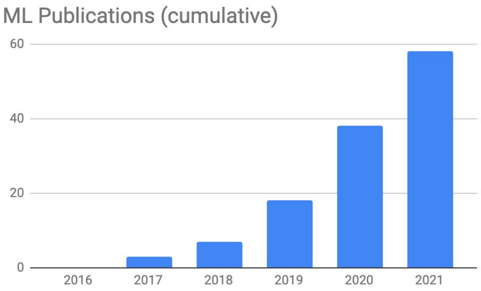

Cutting Edge Research & Transfer

Anatomy of Autonomy

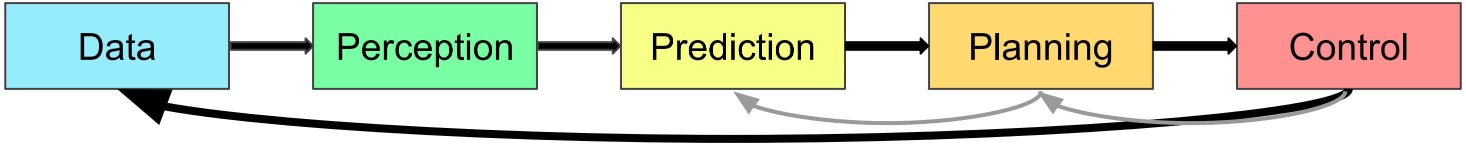

Conway's Law (paraphrased)

A system's design is isomorphic to the structure of the field.

Leaky Abstraction Law

All non-trivial abstractions, to some degree, are leaky.

Corollary: **The Arrows are Performance Bottlenecks.**

Do we even need to bother with modularity?

Demonstrations (good or bad): *trillions* of km/year!

Imitation Learning: simple, scalable, end-to-end

Exploring the Limitations of Behavior Cloning for Autonomous Driving,
F. Codevilla, E. Santana, AM. Lopez, A. Gaidon, ICCV'19 (oral)

Bigger models, pretraining, more data helps... but

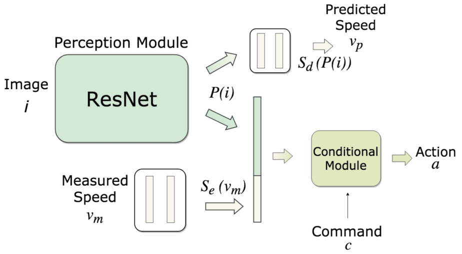

	Task	Variance
CILRS	Empty	23%
	Regular	26%
	Dense	42%
CILRS (ImageNet)	Empty	4%
	Regular	12%
	Dense	38%

Dataset Biases & Variance issues

+ Causal Confusion (de Haan el, NeurIPS'19)

No shortcuts: need intermediate representations

Does computer vision matter for action? Zhou *et al*, Science Robotics 2019

Driving Through Ghosts: Behavioral Cloning with False Positives, Bühler *et al*, IROS'20

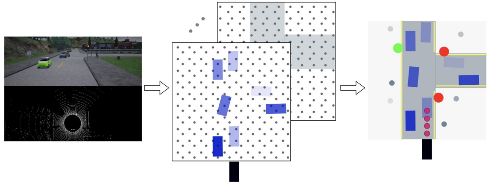

3 R & 3 P of Autonomy

Robustness in Perception
Randomness in Prediction
Risk-awareness in Planning

3 R & 3 P of Autonomy

Robustness in Perception

Randomness in Prediction

Risk-awareness in Planning

Robustness in Perception: Data

Assumption: no expectation for a *robot* to *recognize* something *completely new*.

Domain Coverage: World-scale Fleet Learning

Problem: cannot be supervised (too much data)

TOYOTA RESEARCH INSTITUTE

Supervised Learning

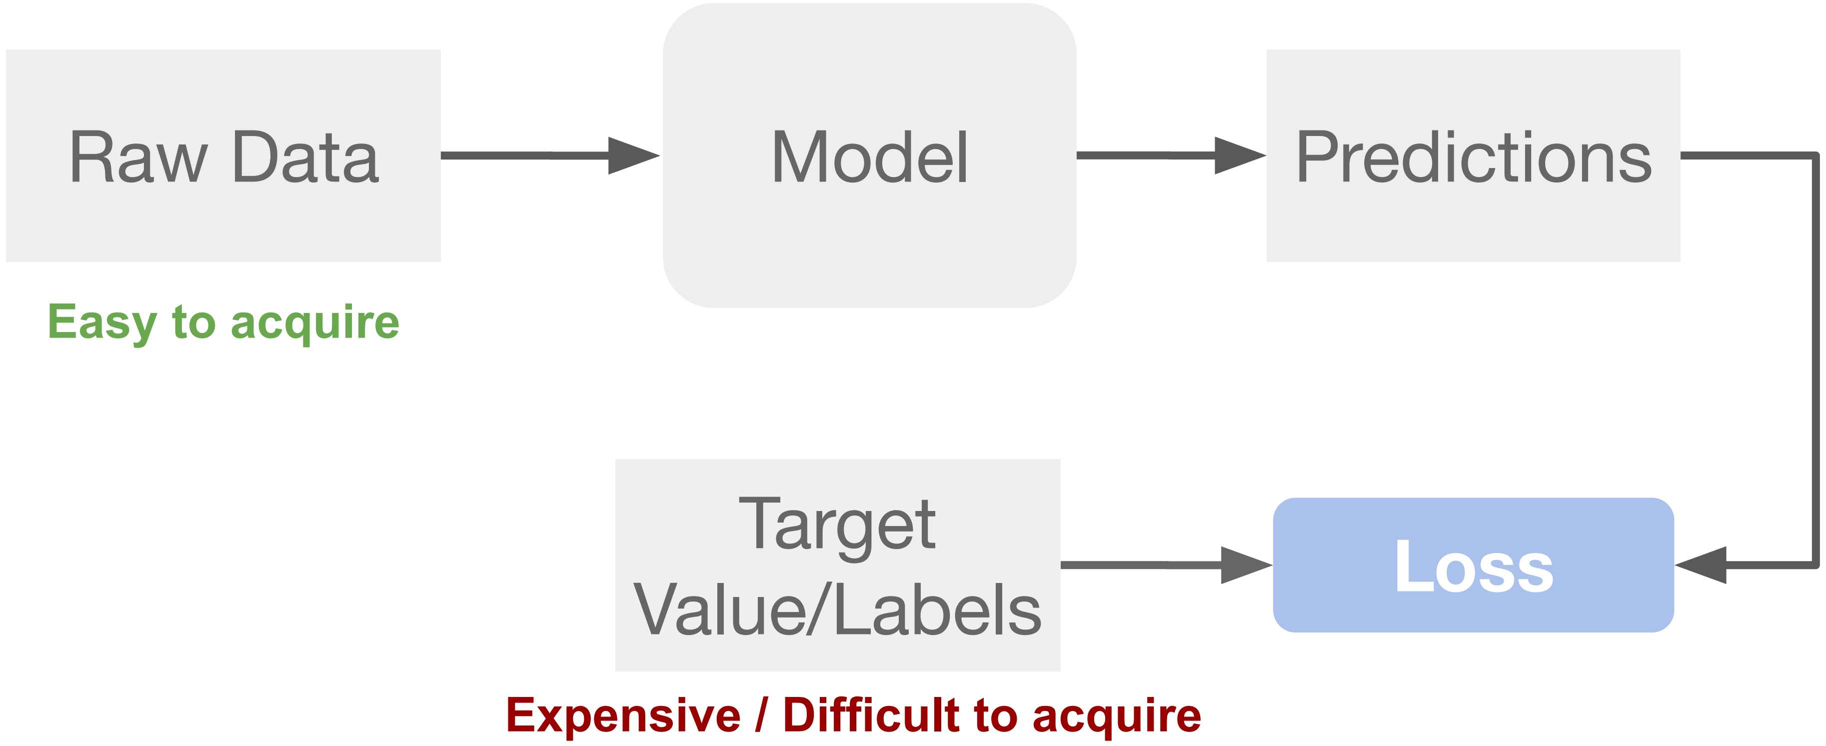

Self-Supervised Learning

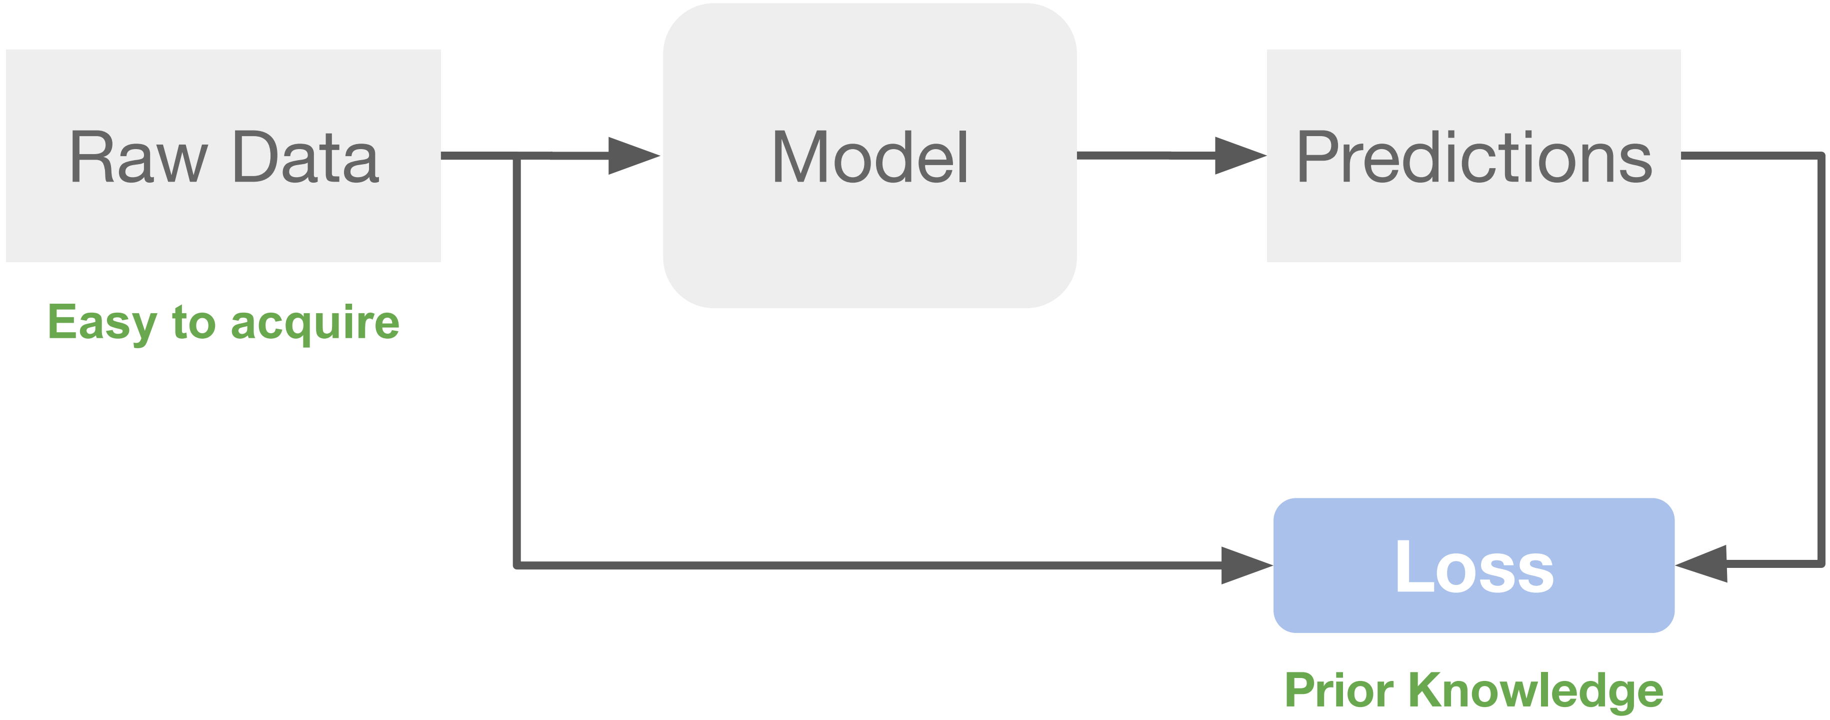

Pseudo-LiDAR / Monocular Depth Estimation

Single RGB Image

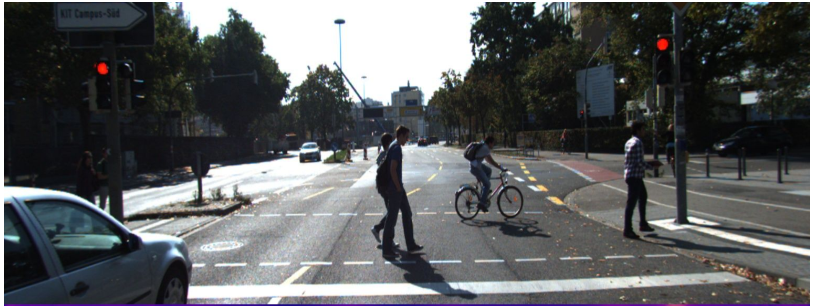

Predicted Depth Image

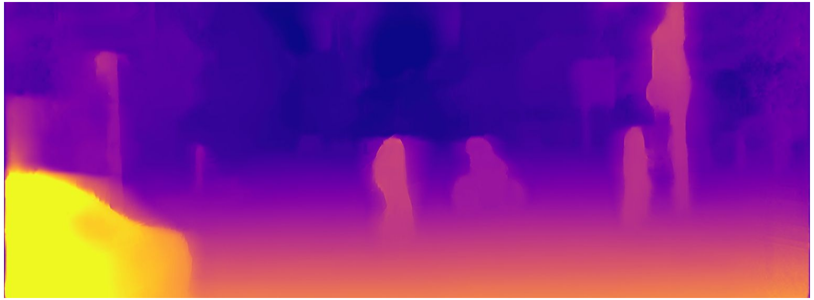

MonoDepth Network

Self-Supervised Structure-from-Motion (SfM)

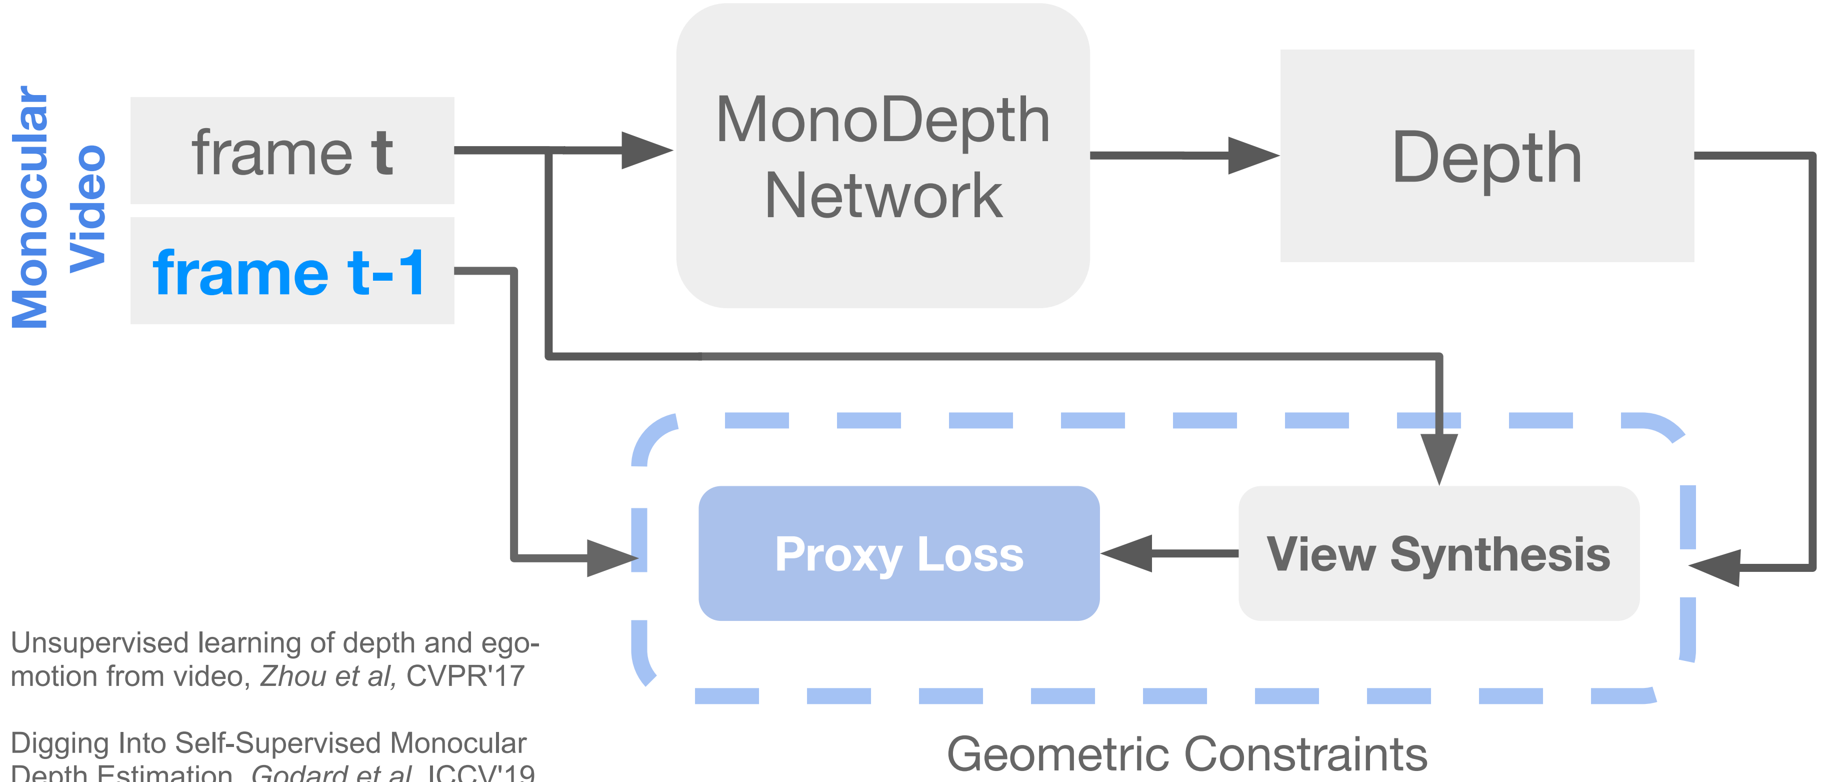

Unsupervised learning of depth and ego-motion from video, *Zhou et al,* CVPR'17

Digging Into Self-Supervised Monocular Depth Estimation, *Godard et al*, ICCV'19

Why Monocular Depth? (Driving, Robotics)

1 camera = cheapest sensor suite, most common

Complex sensor suites: wide baseline, redundancy

Monocular Depth = main bottleneck for 3D detection

Robust 3D Vision thanks to

Large-Scale Self-supervision?

Robustness in Perception: Data

Self-Supervised Learning: Geometry as Supervision

SuperDepth @ **ICRA'19**, *Robust Semi-Sup. monodepth* @ **CoRL'19**, *Two-Stream Networks for Self-Sup. Ego-Motion* @ **CoRL'19**, *Semantically-Guided monodepth* @ **ICLR'20**, ***Packnet-SfM*** @ **CVPR'20 (oral)**, *Neural Ray Surfaces* @ **3DV'20 (oral)**, Monodepth for Soft Visuotactile Sensors @ **Robosoft'21**, Packnet-SAN @ **CVPR'21**, Geometric Unsup. Domain Adaptation @ **ICCV'21**, pre-training for 3D detection @ **ICCV'21**, ...

$$\hat{I}_t(p_t) = I_s(p_s) \qquad p_s \sim K\hat{T}_{t\to s}\hat{D}_t(p_t)K^{-1}p_t$$

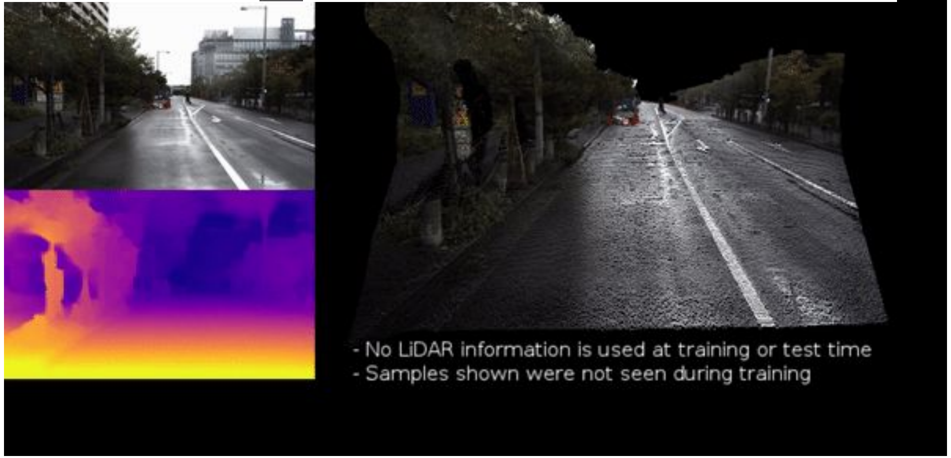

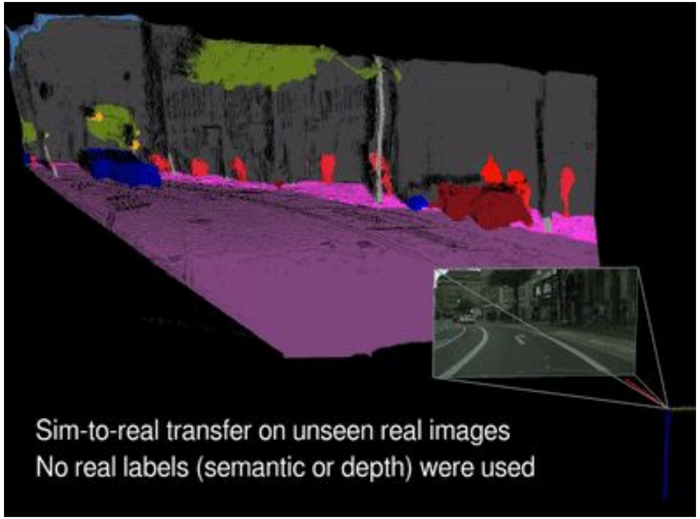

https://github.com/TRI-ML/packnet-sfm

Robustness in Perception: Data

Analysis by Synthesis: Vision as Inverse Graphics via Differentiable Rendering

ROI-10D @ **CVPR'19**, ***SDFLabel*** @ **CVPR'20 (oral)**, *MonoDR* @ **ECCV'20**, DR Survey @ arxiv'20,
Single Shot Scene Reconstruction @ **CoRL'21**

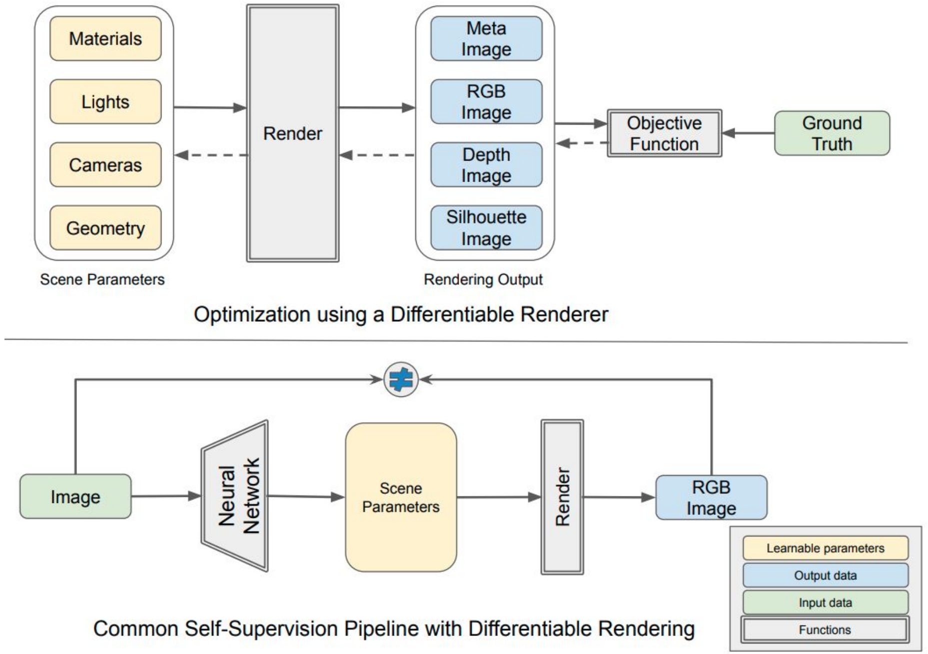

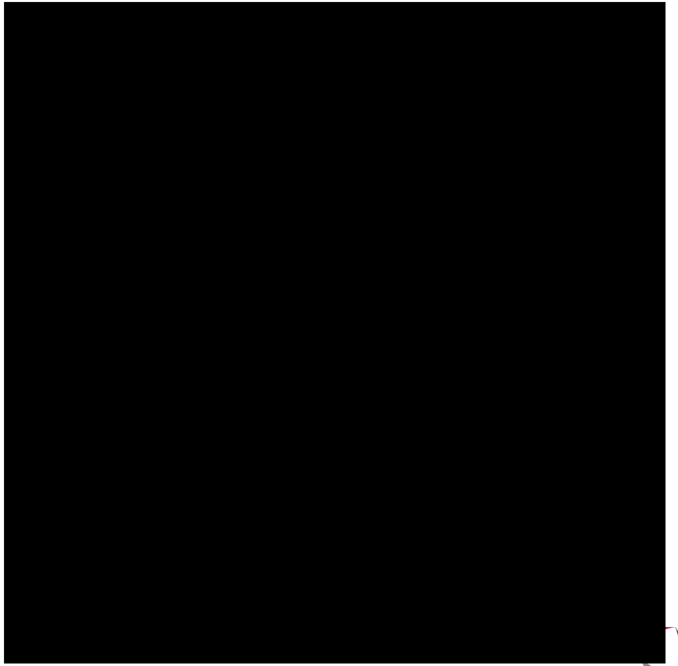

Single-Shot Scene Reconstruction

Sergey Zakharov, Rares Andrei Ambrus, Dennis Park, Vitor Campagnolo Guizilini, Wadim Kehl, Fredo Durand, Joshua B. Tenenbaum, Vincent Sitzmann, Jiajun Wu, Adrien Gaidon, **CORL'21** (paper)

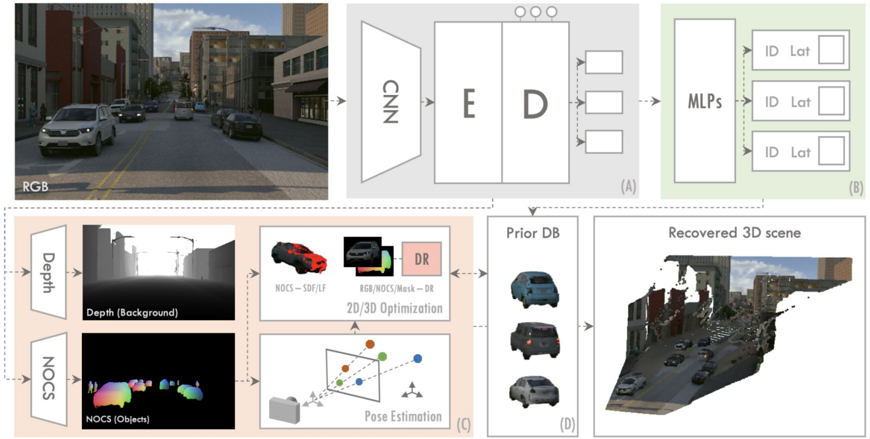

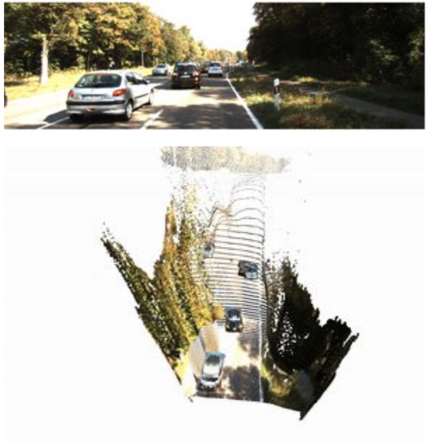

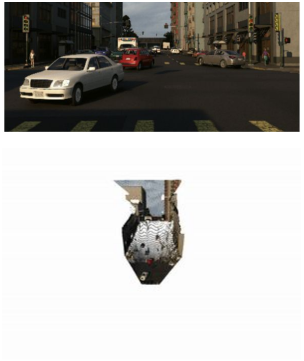

Fully editable and ***re-renderable*** *model* of a 3D scene from a ***single image***

Decompose the scene into ***object instances*** and ***background***

2D/3D optimization based on ***differentiable rendering***

Robustness in Perception: Data

Beware of Bias: Adaptive Regularization of the Long Tail

Learning Imbalanced Datasets with Label-Distribution-Aware Margin Loss, K. Cao *et al*, **NeurIPS'19**
Heteroskedastic and Imbalanced Deep Learning with Adaptive Regularization, K. Cao *et al*, **ICLR'21**

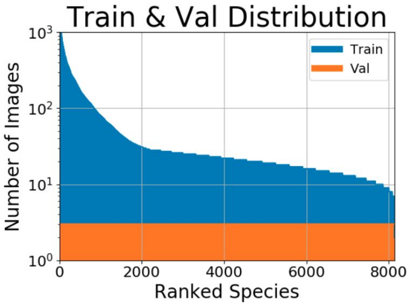

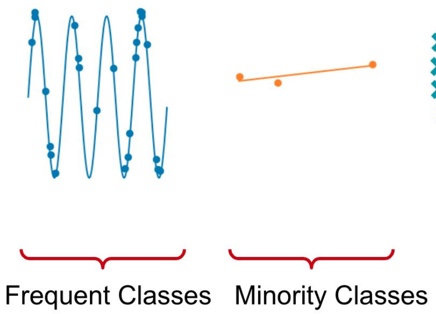

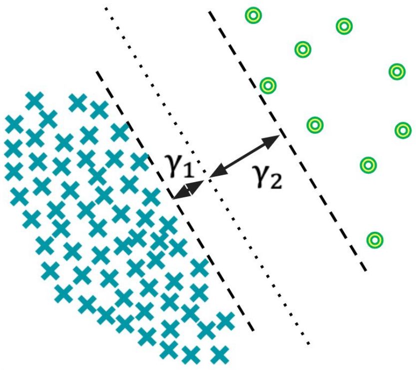

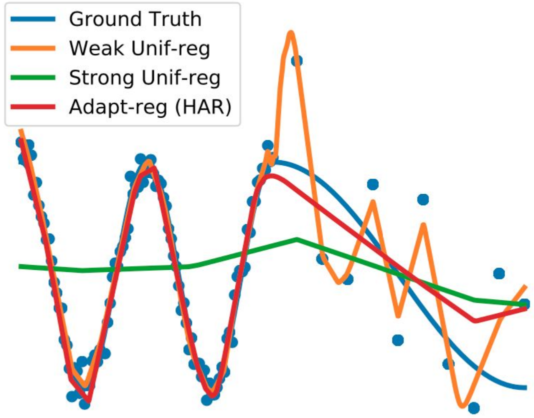

Robustness in Perception: Data

[Provable Guarantees for Self-Supervised Deep Learning with Spectral Contrastive Loss](), Jeff Z. HaoChen, Colin Wei, Adrien Gaidon, Tengyu Ma - **NeurIPS'21 (oral)**

[Self-supervised Learning is More Robust to Dataset Imbalance](), Hong Liu, Jeff Z. HaoChen, Adrien Gaidon, Tengyu Ma, arxiv'21

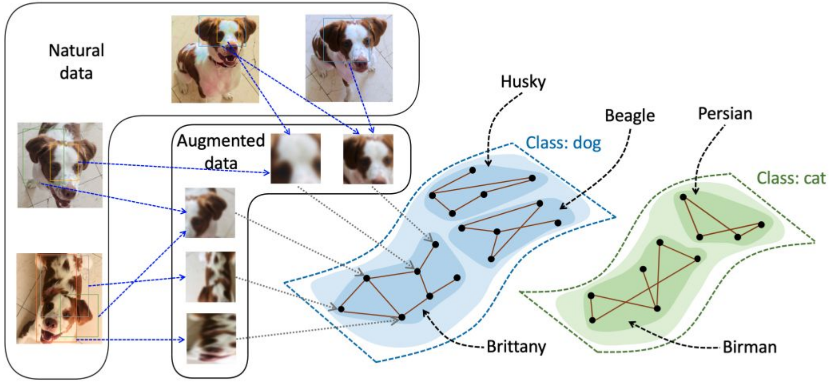

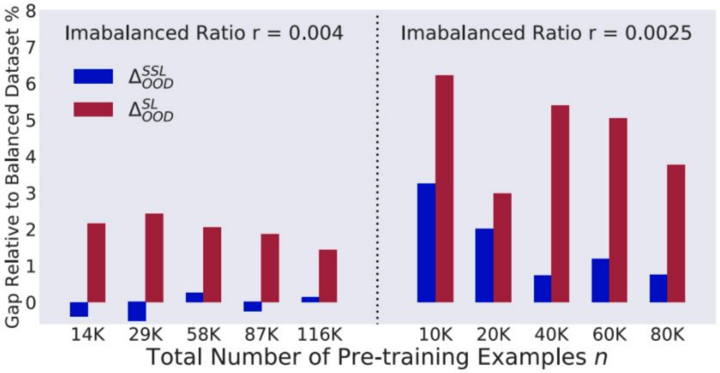

Robustness in Perception: Redundancy

Redundancy: sensors lie (Byzantine Generals)

→ Combine Radar + Lidar + Cameras + IMU + ...

Bottleneck: 3D Detection with Cameras (< Lidar)

Approach:

- Cross-sensor Auto-Labeling (2D-3D)
- Self/Semi-Supervised "pseudo-lidar"

  Robust Semi-Sup. monodepth @ **CoRL'19**, ***SDFLabel*** @ **CVPR'20 (oral)**, Packnet-SAN @ **CVPR'21**, *Full Surround Monodepth* @ **arxiv'21**

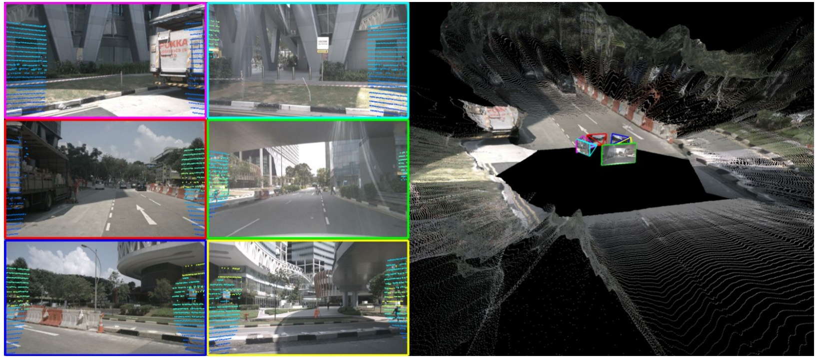

TOYOTA RESEARCH INSTITUTE

PackNet-SAN

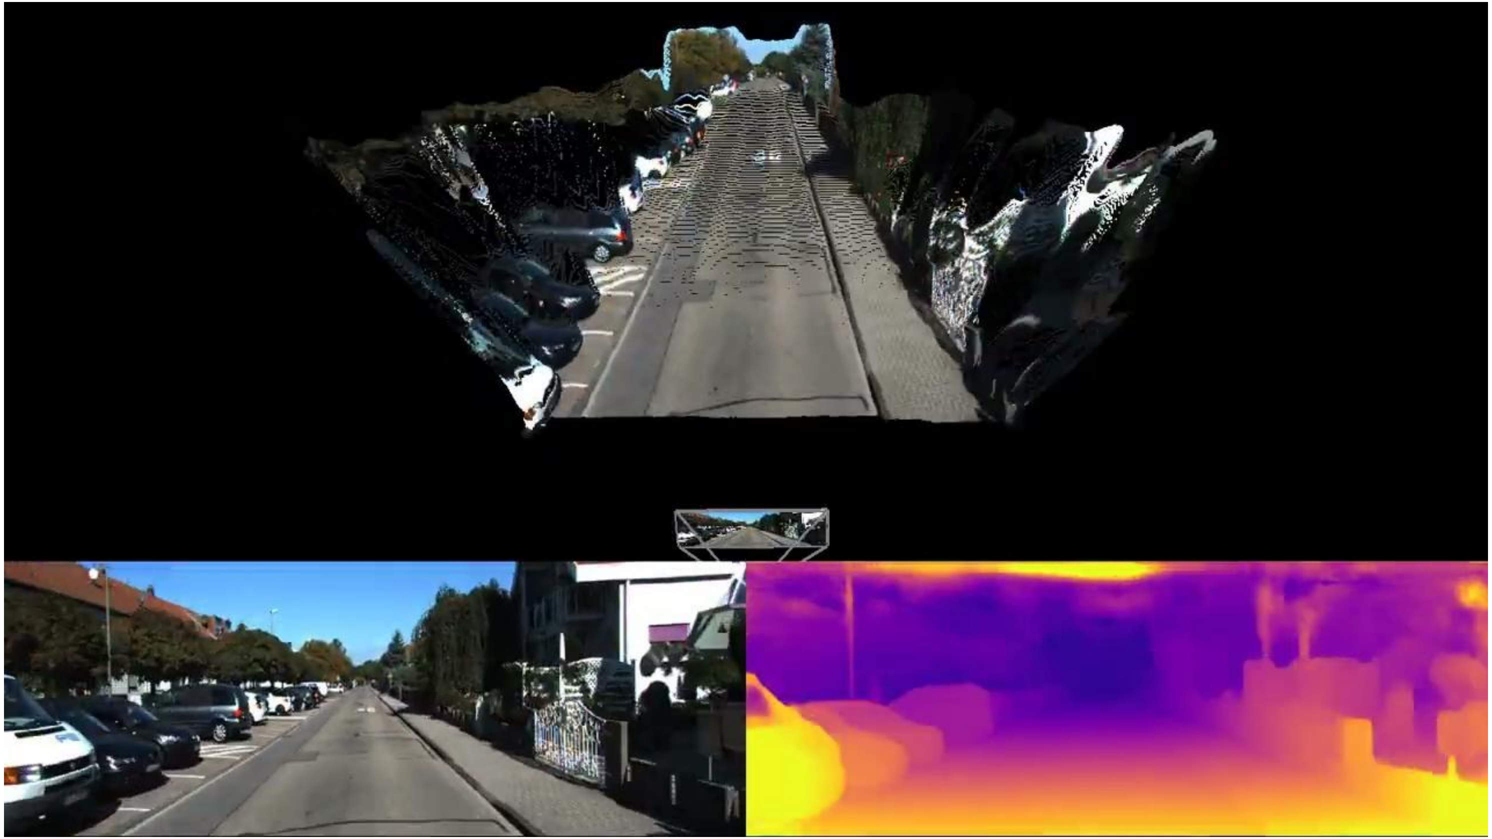

Robustness in Perception: Efficiency

Many sensors + complex online fusion

→ Deployment challenges (runtime, energy, stability...)

Efficiency: Core to Robustness!

- Hardware optimization

- Sharing Computations

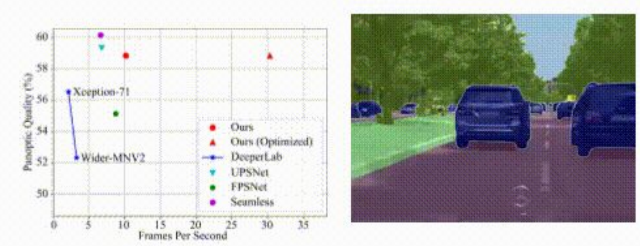

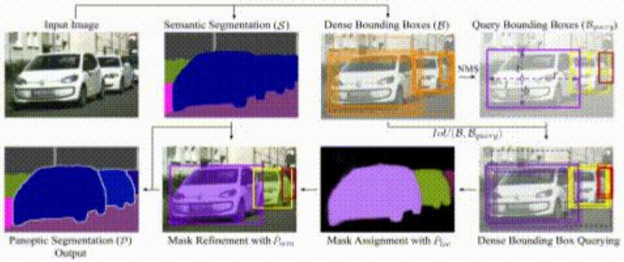

[Real-Time Panoptic Segmentation from Dense Detections](#), R. Hou, J. Li *et al*, **CVPR'20 (oral)**, *[Hierarchical Lovász Embeddings for Proposal-free Panoptic Segmentation](#)*, Kerola *et al*, **CVPR'21**

TOYOTA RESEARCH INSTITUTE

3 R & 3 P of Autonomy

Robustness in Perception

Data + Redundancy + Efficiency

Randomness in Prediction

Risk-awareness in Planning

3 R & 3 P of Autonomy

Robustness in Perception

Randomness in Prediction

Risk-awareness in Planning

TOYOTA RESEARCH INSTITUTE

Randomness in Prediction: Intent

Human Intent: *latent* yet governs Action

Approach: infer from Interactions and Context

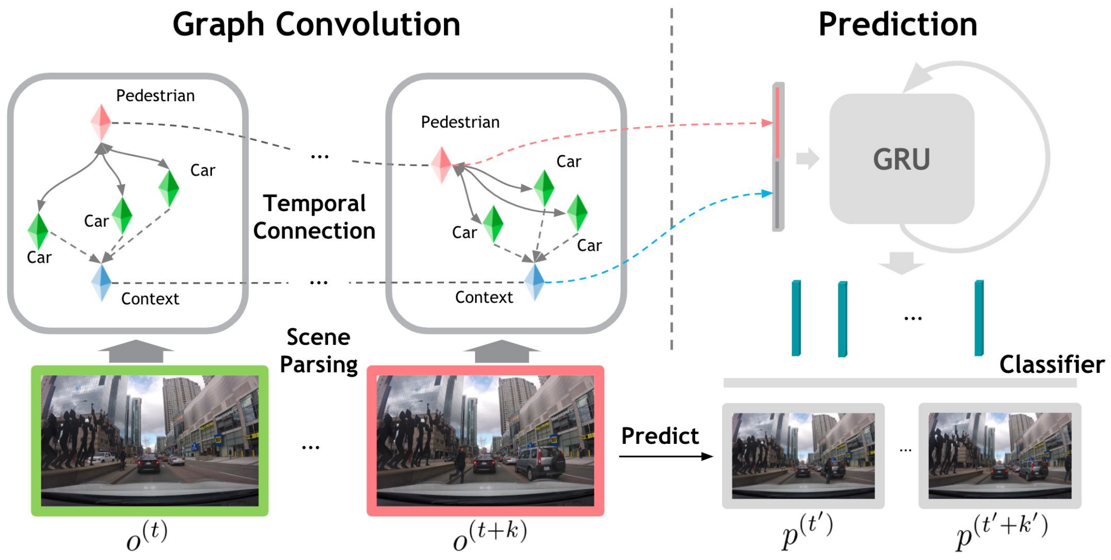

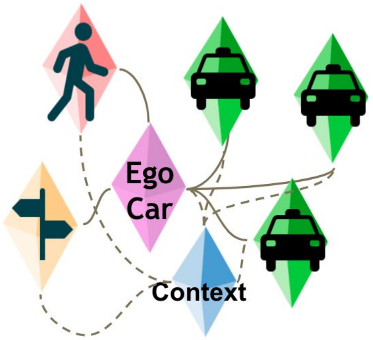

Spatiotemporal Relationship Reasoning for Pedestrian Intent Prediction, **RA-L & ICRA'20**

STIP: Stanford-TRI Intent Prediction
http://stip.stanford.edu/

Randomness in Prediction: Multi-Modality

Contingency: Predict *Distribution* over *Plausible* Futures

Structure Prediction like Human Decision-Making Process

It Is Not the Journey but the Destination: Endpoint Conditioned Trajectory Prediction,
K. Mangalam, H. Girase, S. Agarwal, K-H. Lee, E. Adeli, J. Malik, A. Gaidon, **ECCV 2020 (oral)**

Learn distribution of plausible destinations (intents)

Condition trajectory forecasting on sampled endpoints

Plan following social norms

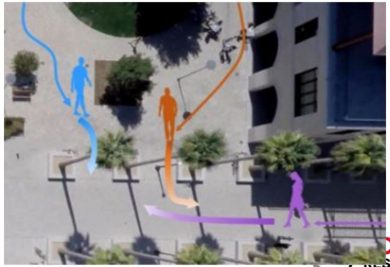

Randomness in Prediction: Uncertainty

Probabilistic Robotics: Every State is a Distribution!

Modular System: Uncertainty Must Propagate *Throughout*

Predicting Uncertainty: typically the main focus of research

Open challenge: *using* Uncertainty Upstream *and* Downstream

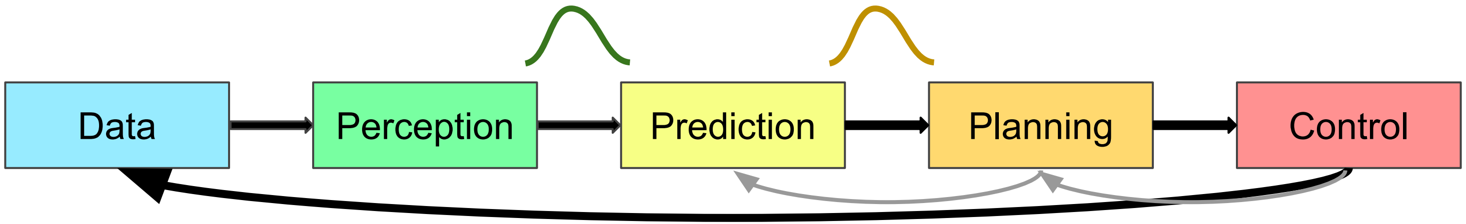

Using Upstream (Perceptual) Uncertainty

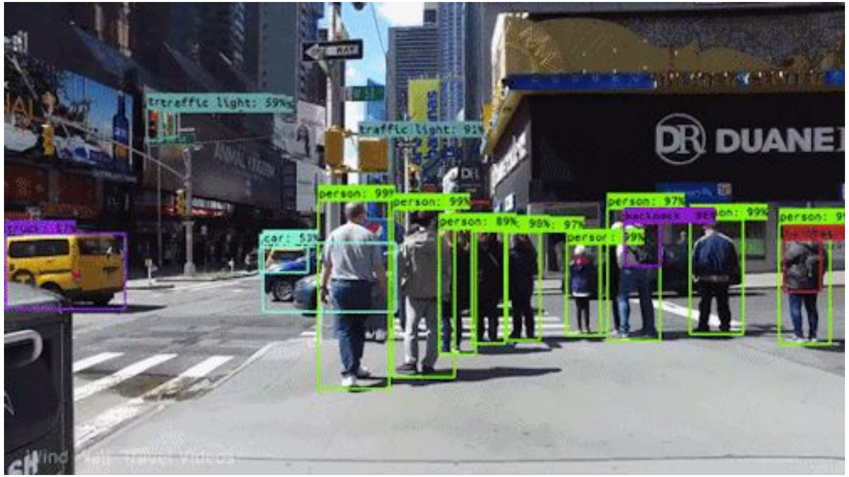

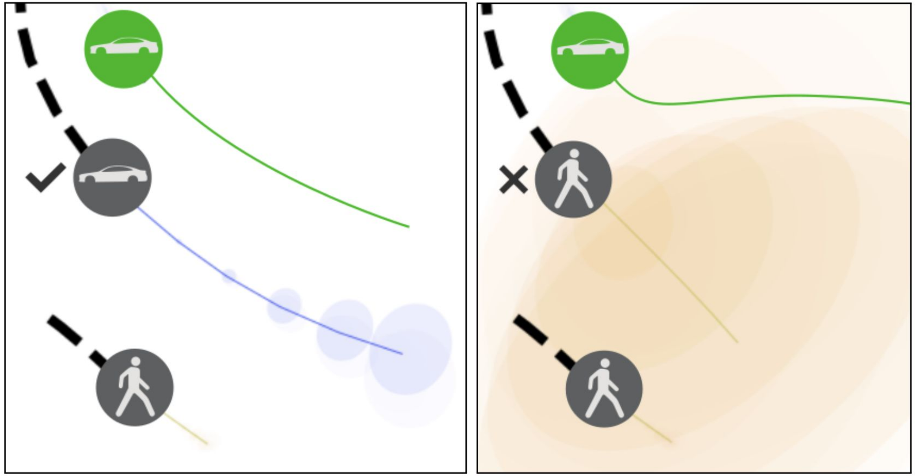

Heterogeneous-Agent Trajectory Forecasting Incorporating Class Uncertainty (**HAICU**), Ivanovic *et al*, arxiv:2104.12446

Using Upstream (Perceptual) Uncertainty

PUP: new dataset of driving logs with Perceptual Uncertainty (in challenging scenes for tracking) for Prediction

Class	PUP (Ours) Num. (%)	PUP (Ours) S_{probs}	Lyft Level 5 Num. (%)	Lyft Level 5 S_{probs}
bicycle	1.2k (0.8)	1.60	0.1M (0.4)	0.09
car	117k (81.9)	1.10	5.0M (24.5)	0.00
largevehicle	18k (12.4)	1.30	—	—
motorcycle	0.5k (0.3)	1.57	—	—
pedestrian	6.3k (4.4)	1.44	0.7M (3.3)	0.01
unknown	0.2k (0.1)	0.05	14.6M (71.8)	0.00

$$S_{\text{probs}} = -\sum_k P(C_i = k) \log P(C_i = k)$$

For reference, the maximum possible entropy for the PUP dataset is $\ln(11) \approx 2.40$ (uniform class probabilities)

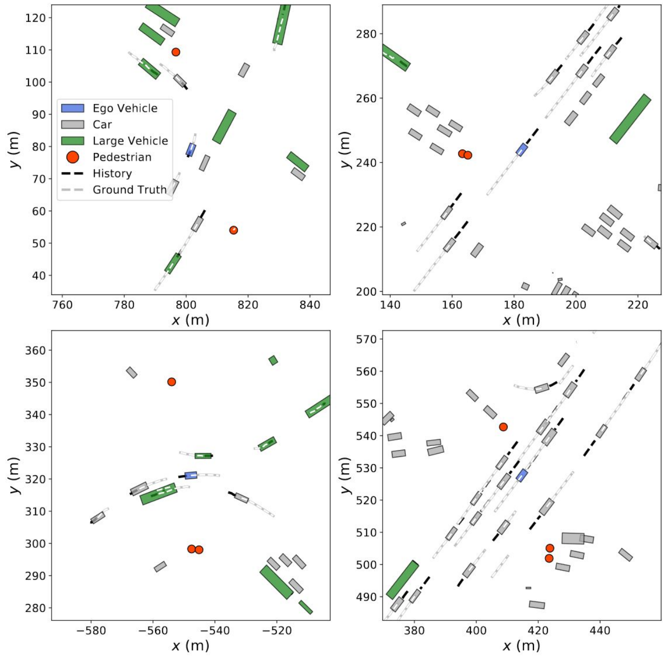

TOYOTA RESEARCH INSTITUTE

Using Upstream (Perceptual) Uncertainty

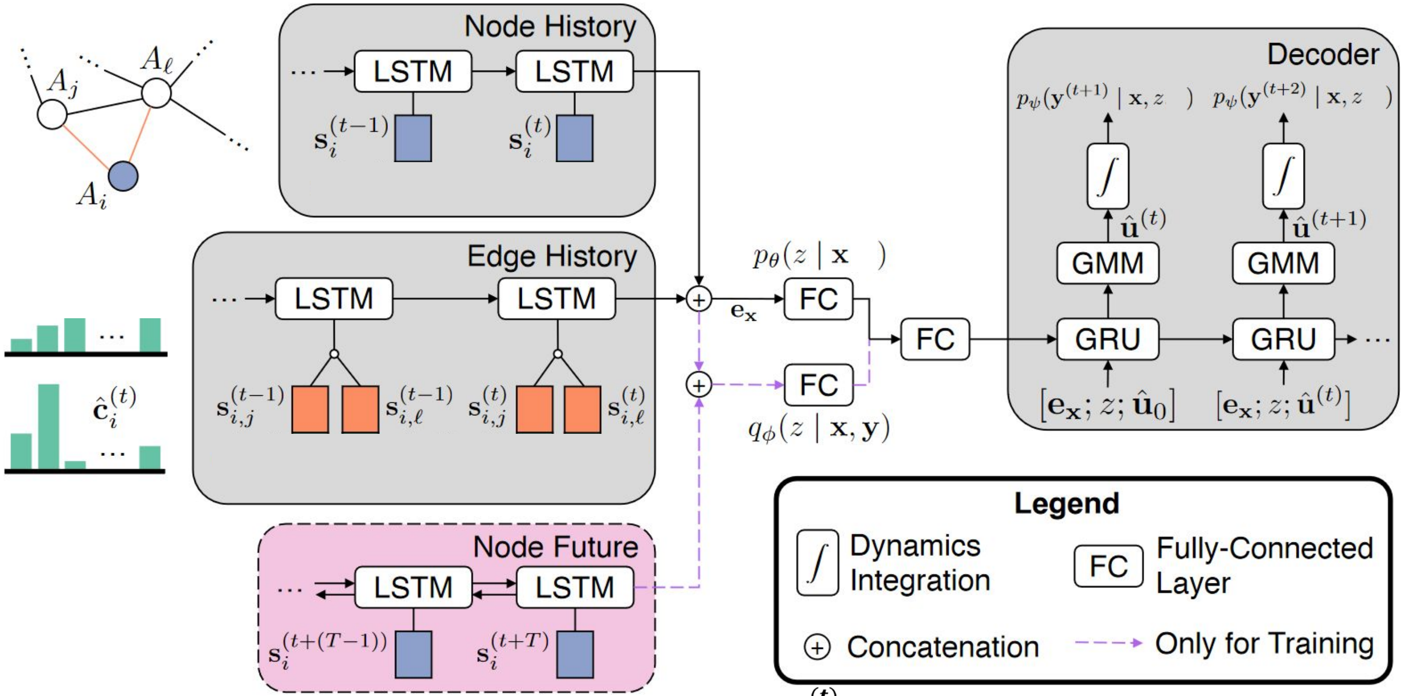

$\boldsymbol{s}_{\boldsymbol{i}}^{(\boldsymbol{t})}$ = state of agent i at time t

Using Upstream (Perceptual) Uncertainty

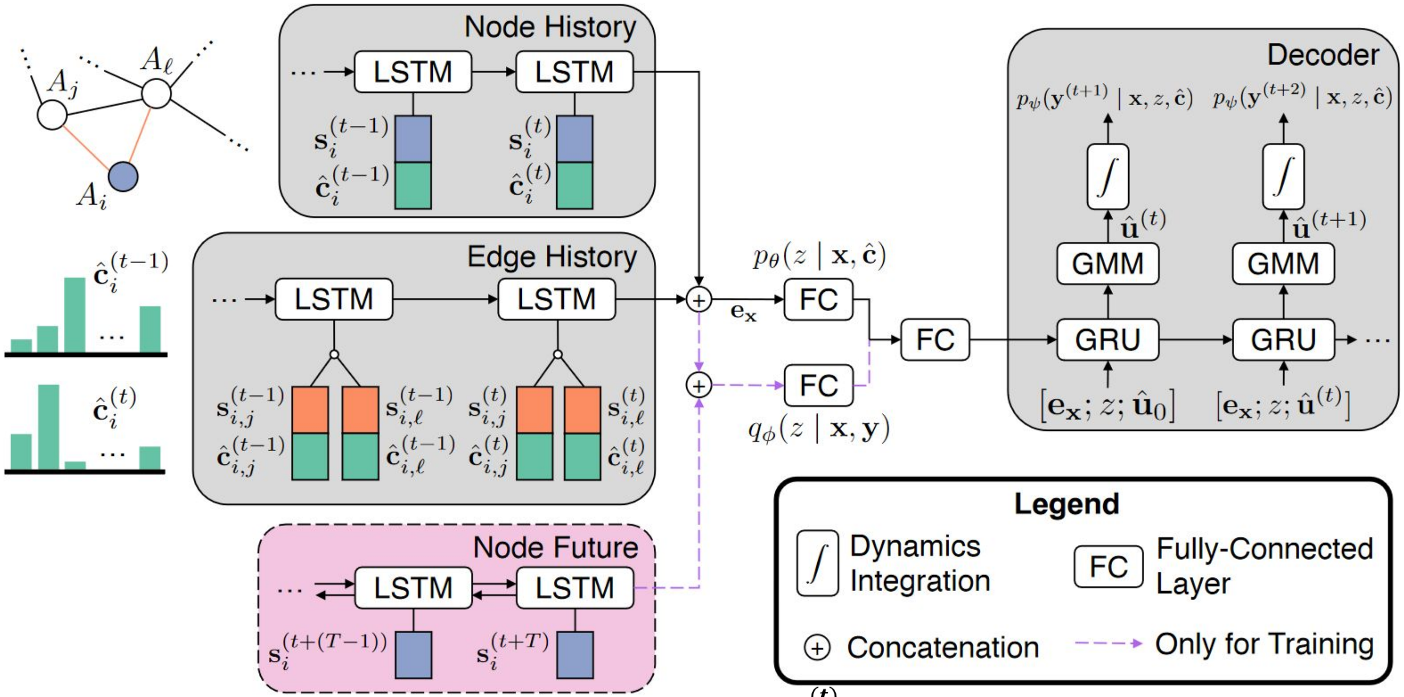

$s_i^{(t)}$ = state of agent i at time t

$\hat{c}_i^{(t)}$ = class probabilities of agent i at time t

Using Upstream (Perceptual) Uncertainty

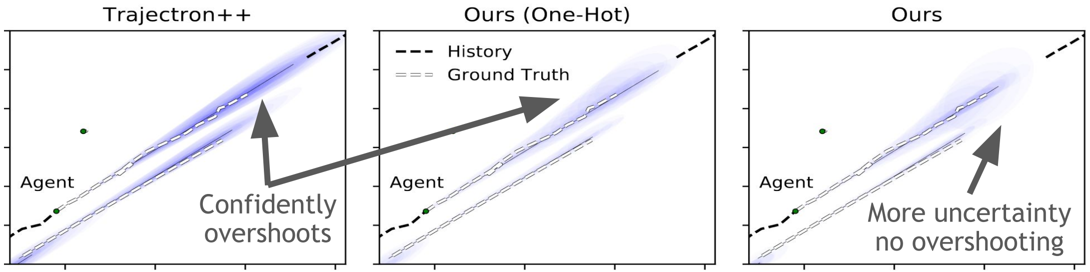

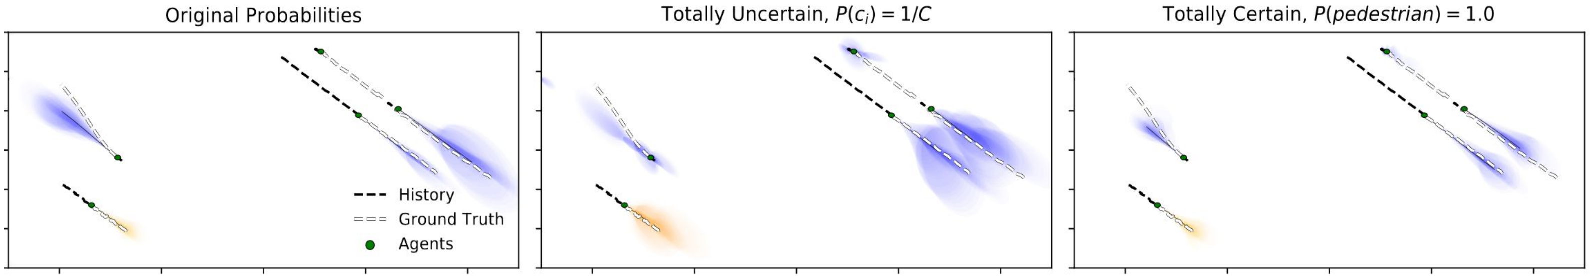

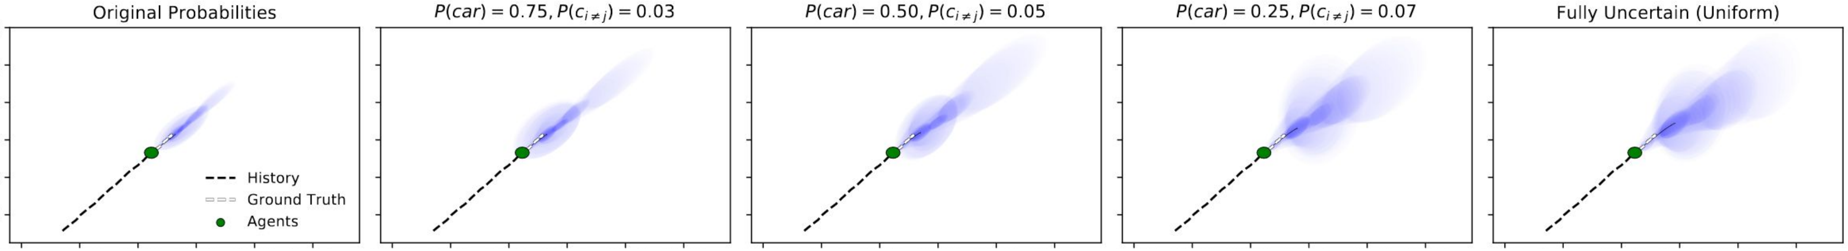

Uncertainty Representation for Control

MATS: An Interpretable Trajectory Forecasting Representation for Planning and Control, B. Ivanovic, A. Elhafsi, G. Rosman, A. Gaidon, M. Pavone, **CoRL'20**

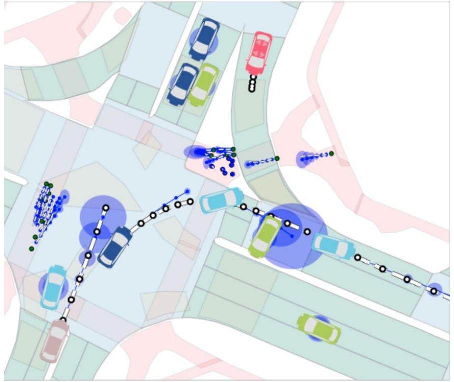

$$\mathbf{s}_z^{(t+1)} = \underbrace{A_z^{(t)}\mathbf{s}_z^{(t)}}_{\text{scene dynamics}} + \underbrace{B_z^{(t)}\mathbf{u}_R^{(t)}}_{\text{ego control effect}} + \underbrace{\mathbf{c}_z^{(t)}}_{\text{linearization}} + \underbrace{Q_z^{(t)}\mathbf{w}^{(t)}}_{\text{uncertainty}}$$

Blue = Dynamics, Green = Learned

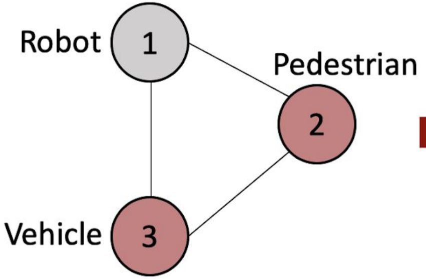

$$\begin{bmatrix} A_{11} & 0 & 0 \\ A_{21} & A_{22} & A_{23} \\ A_{31} & A_{32} & A_{33} \end{bmatrix} \begin{bmatrix} x_R \\ x_2 \\ x_3 \end{bmatrix} + \begin{bmatrix} B_1 \\ B_2 \\ B_3 \end{bmatrix} \begin{bmatrix} u_R \end{bmatrix} + \begin{bmatrix} c_1 \\ c_2 \\ c_3 \end{bmatrix} + \begin{bmatrix} Q_1 \\ Q_2 \\ Q_3 \end{bmatrix} \begin{bmatrix} w \end{bmatrix}$$

Control-Aware Prediction Objectives (CAPO)

Objective Mismatch: forecasting metrics are unaware of their use

CAPO: weight prediction metrics by their effect on control

- Cross-attention weighting [Mercat et al, ICRA'20]
- Counterfactual weighting: Action Discrepancy based on resampling

Improves forecasting **where it matters most** (potential collisions)

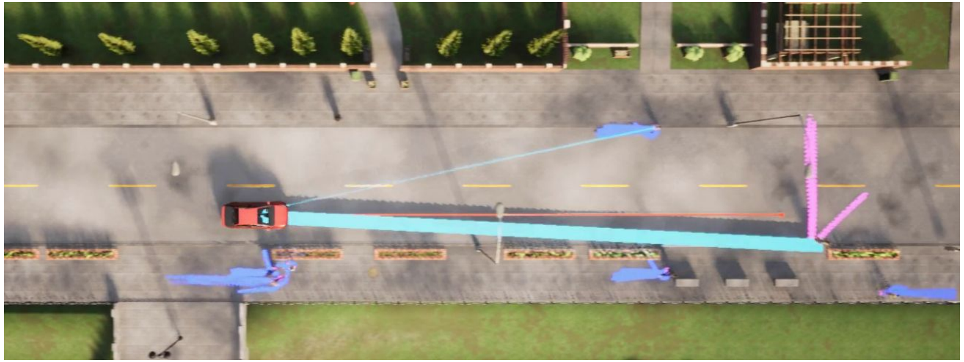

Control-Aware Prediction Objectives for Autonomous Driving

Rowan McAllister, Blake Wulfe, Jean Mercat, Logan Ellis, Sergey Levine, Adrien Gaidon (soon on arxiv)

3 R & 3 P of Autonomy

Robustness in Perception

Randomness in Prediction

Intent + Multi-modality + Uncertainty

Risk-awareness in Planning

3 R & 3 P of Autonomy

Robustness in Perception

Randomness in Prediction

Risk-awareness in Planning

TOYOTA RESEARCH INSTITUTE

Risk-awareness in Planning

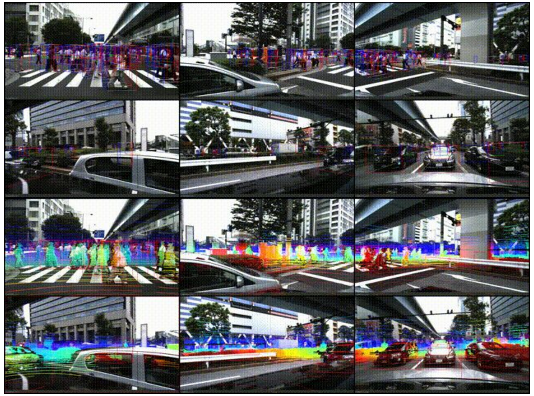
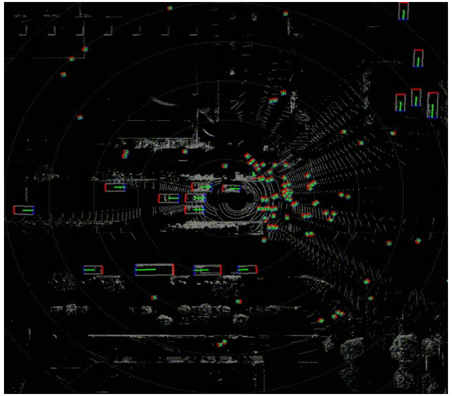

Perception / Prediction: always stochastic

Basis for *Safety-critical* Decisions in *Real-time*?

Safe Autonomy requires *Risk-Awareness*

Risk-awareness in Planning: Safety

Risk-Sensitive* *Sequential Action Control with* *Multi-Modal Human Trajectory Forecasting* *for Safe Crowd-Robot Interaction, H. Nishimura, B. Ivanovic, A. Gaidon, M. Pavone, M. Schwager, **IROS'20**

Environment

Probabilistic Human Motion Prediction

Model-based Planning / Control

Control

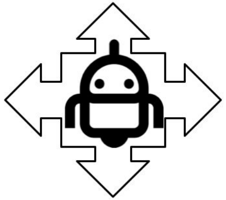

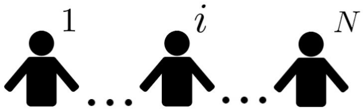

$$p^i_{k+1} = p^i_k + y^i_k \quad \{y^{1:N}_{1:T}\} \sim \mathcal{D}$$

- **Discrete-Time**
- **Stochastic**
- Arbitrary Distribution

Entropic Risk

$$R_{\mathcal{D},\sigma}(J) \triangleq \frac{1}{\sigma} \log\left(\mathbb{E}_{\mathcal{D}}[e^{\sigma J}]\right)$$

$$R_{\mathcal{D},\sigma}(J) \approx \mathbb{E}_{\mathcal{D}}[J] + \frac{\sigma}{2}\mathrm{Var}_{\mathcal{D}}(J)$$

Risk-neutral $\quad$ Risk-sensitive

$\sigma = 0 \quad \sigma > 0$

$$\dot{x}(t) = f\left(x(t)\right) + H\left(x(t)\right)u(t)$$

- **Continuous-Time**
- Deterministic
- Control-Affine

Risk-awareness in Planning: Safety

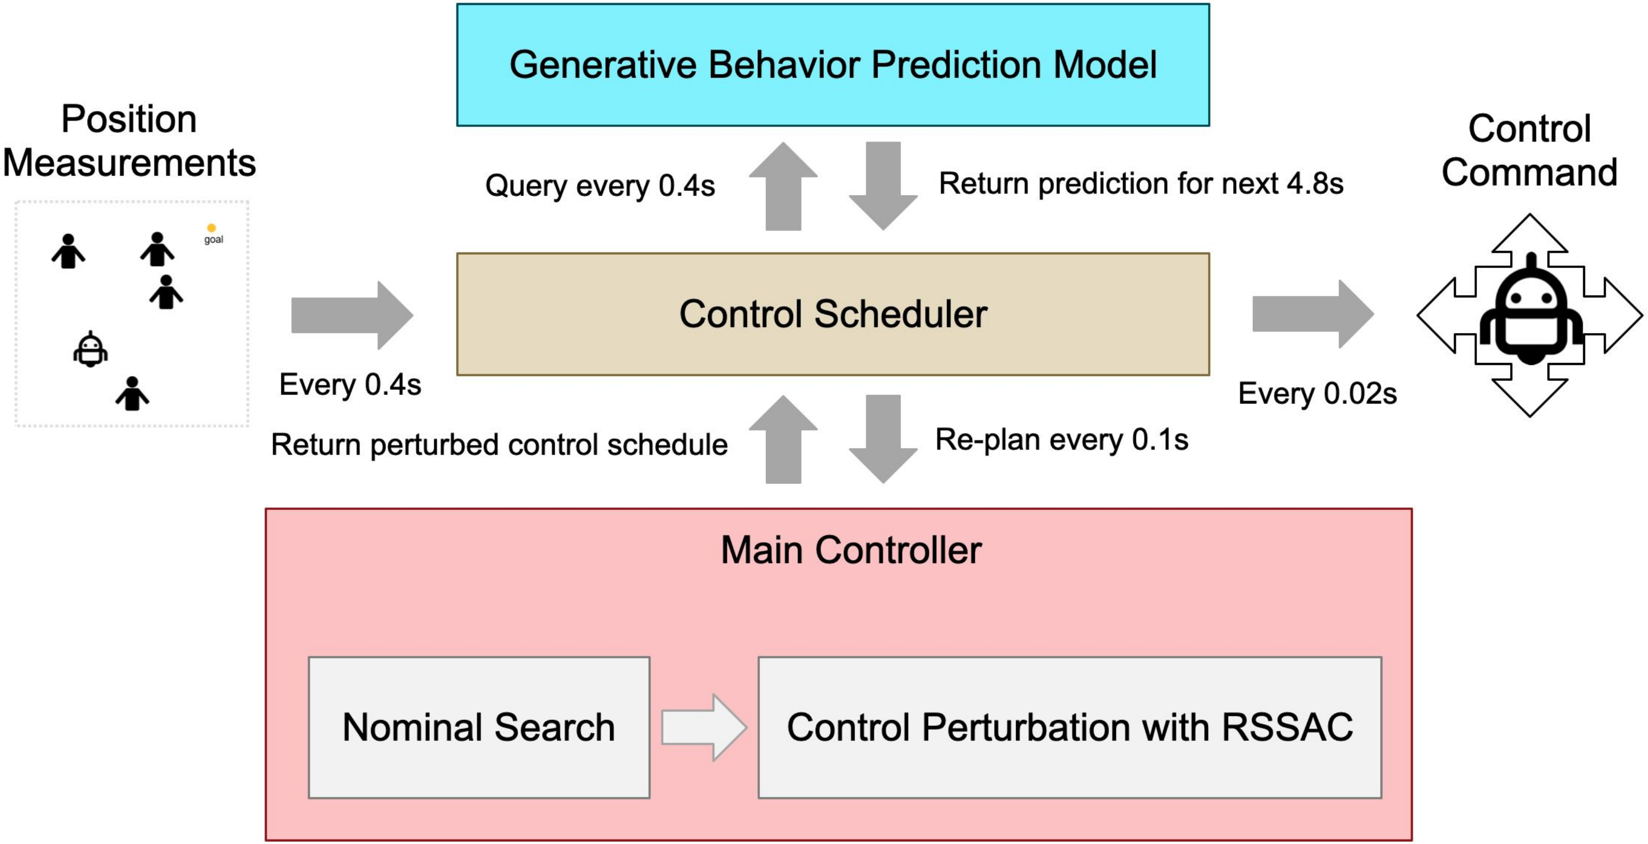

Risk-awareness in Planning: Safety

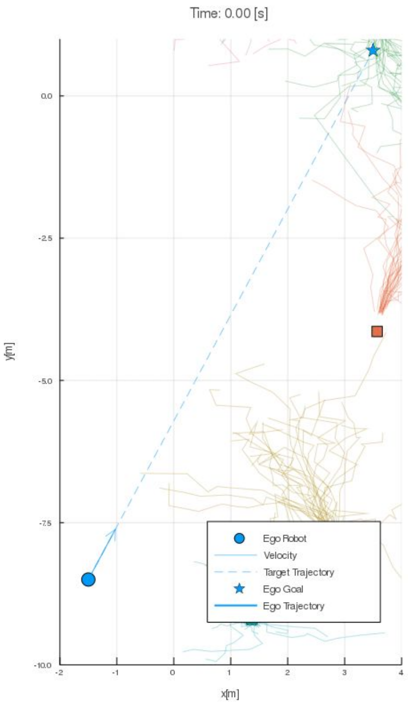

σ = 0.0 (Risk-Neutral)

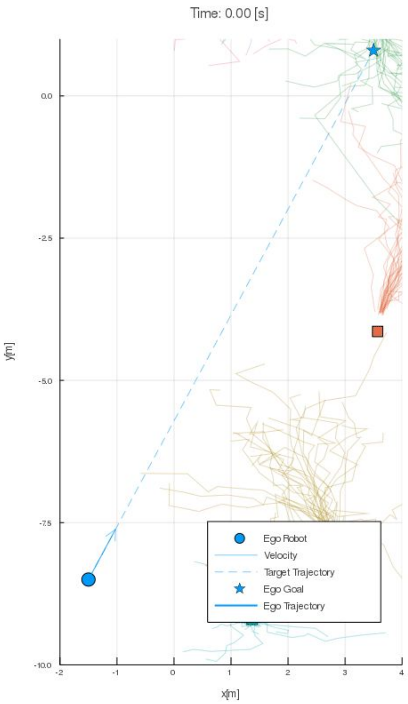

σ = 1.0 (Risk-Sensitive)

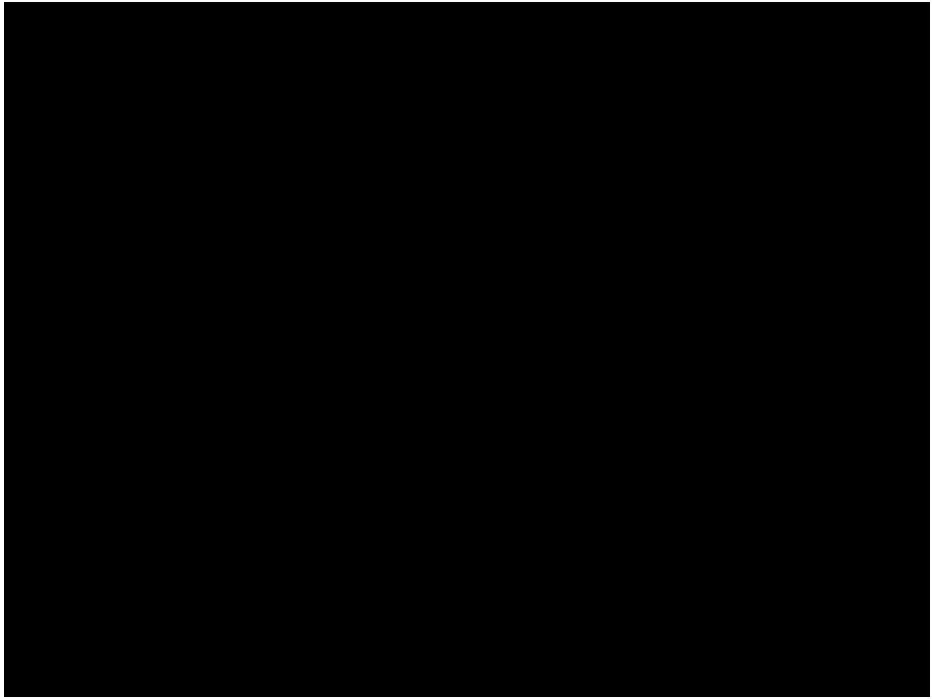

Game-Theoretic Planning for Risk-Aware Interactive Agents, M. Wang, N. Mehr, A. Gaidon, M. Schwager **IROS'20**

RAT iLQR: A Risk Auto-Tuning Controller to Optimally Account for Stochastic Model Mismatch, Nishimura *et al*, **RA-L/ICRA'21**

Risk-awareness in Planning: Causality

Decision-making: beyond pure data → Causal Inference

Counterfactuals in sim: find planner bugs and fixes

Behaviorally Diverse Traffic Simulation via Reinforcement Learning, S. Maruyama *et al*, **IROS'20**

Discovering Avoidable Planner Failures of Autonomous Vehicles using Counterfactual Analysis in Behaviorally Diverse Simulation, D. Nishiyama *et al,* **ITSC'20**

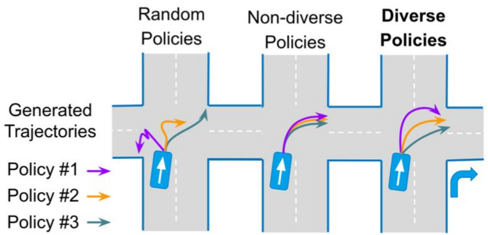

Diversity	High	Low	**High**
Driving Skills	Low	High	**High**

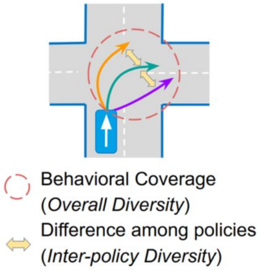

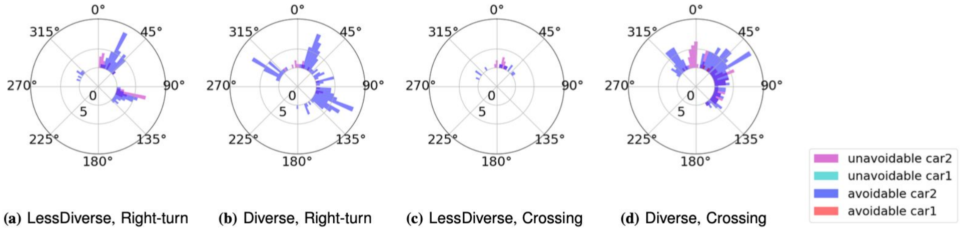

(a) LessDiverse, Right-turn (b) Diverse, Right-turn (c) LessDiverse, Crossing (d) Diverse, Crossing

Risk-awareness in Planning: Near-Accidents

Imitation in Near-Accidents? Phase Transitions

Reinforcement Learning based Control of Imitative Policies for Near-Accident Driving,
Z. Cao, E. Biyik, W. Z. Wang, A. Raventos, A. Gaidon, G. Rosman, D. Sadigh, RSS'20

RL to switch between basic IL policies

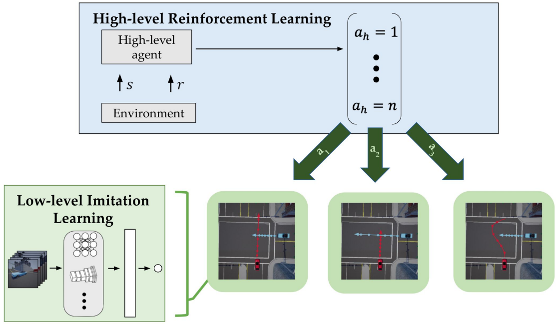

TOYOTA RESEARCH INSTITUTE

Risk-awareness in Planning: Near-Accidents

Imitation in Near-Accidents? Phase Transitions

Reinforcement Learning based Control of Imitative Policies for Near-Accident Driving, Z. Cao, E. Biyik, W. Z. Wang, A. Raventos, A. Gaidon, G. Rosman, D. Sadigh, RSS'20

RL to switch between basic IL policies

Less collisions + human-like

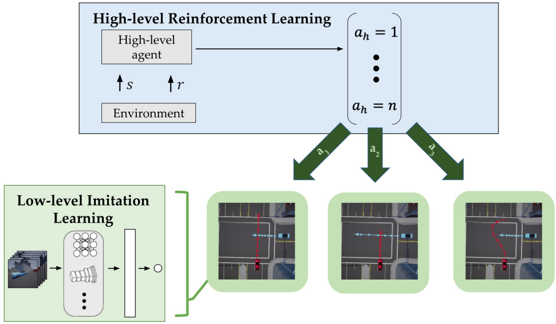

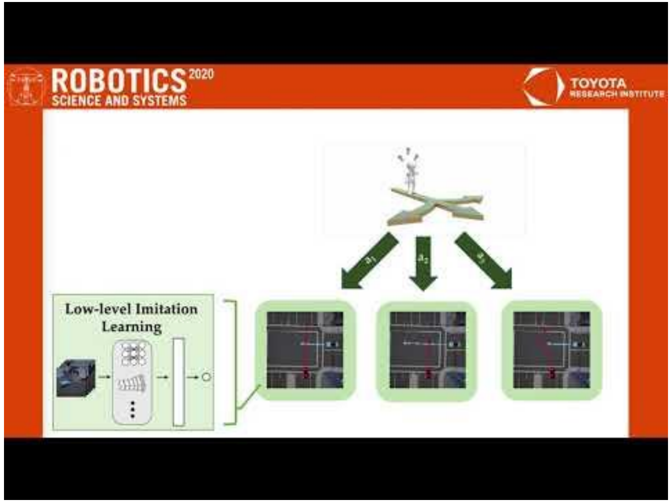

TOYOTA RESEARCH INSTITUTE

3 R & 3 P
of
Autonomy

Robustness in Perception

Randomness in Prediction

Risk-awareness in Planning

Uncertainty + Safety + Causality

Robot = Complex Sensorimotor Loop

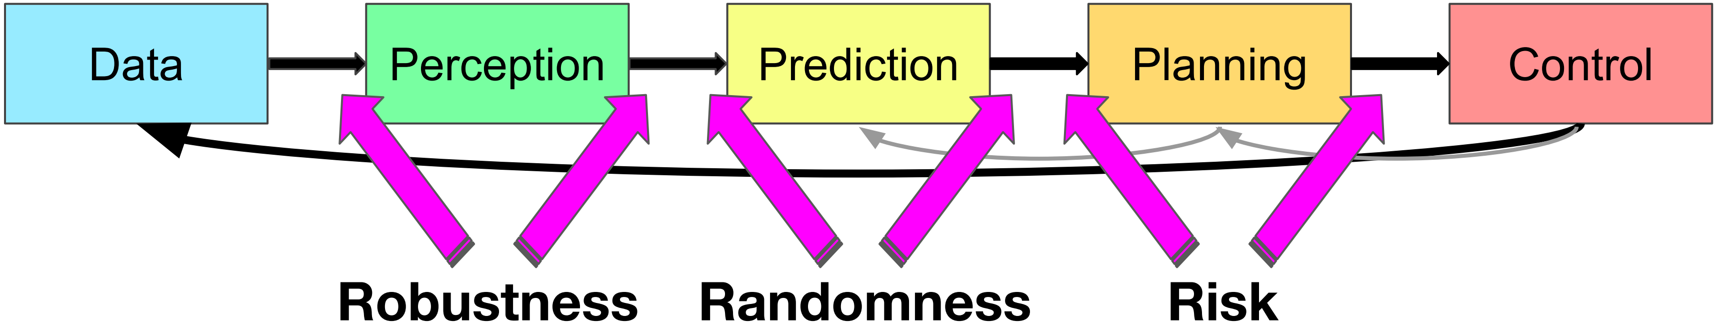

Law of Arrows: In a modular sensorimotor system, the performance bottlenecks are at the *interface* between modules.

Work on the Arrows!

Code & Data: https://github.com/TRI-ML
Blog posts: https://medium.com/toyotaresearch/
Twitter: https://twitter.com/ToyotaResearch

tri.global/careers