CaT: Coaching a Teachable Student

Jimuyang Zhang, Zanming Huang, Eshed Ohn-Bar

Boston University

Paper Tag: TUE-PM-352

JUNE 18-22, 2023 CVPR VANCOUVER, CANADA

BOSTON UNIVERSITY

H2X Lab

How to ***Effectively Teach*** Sensorimotor Agents?

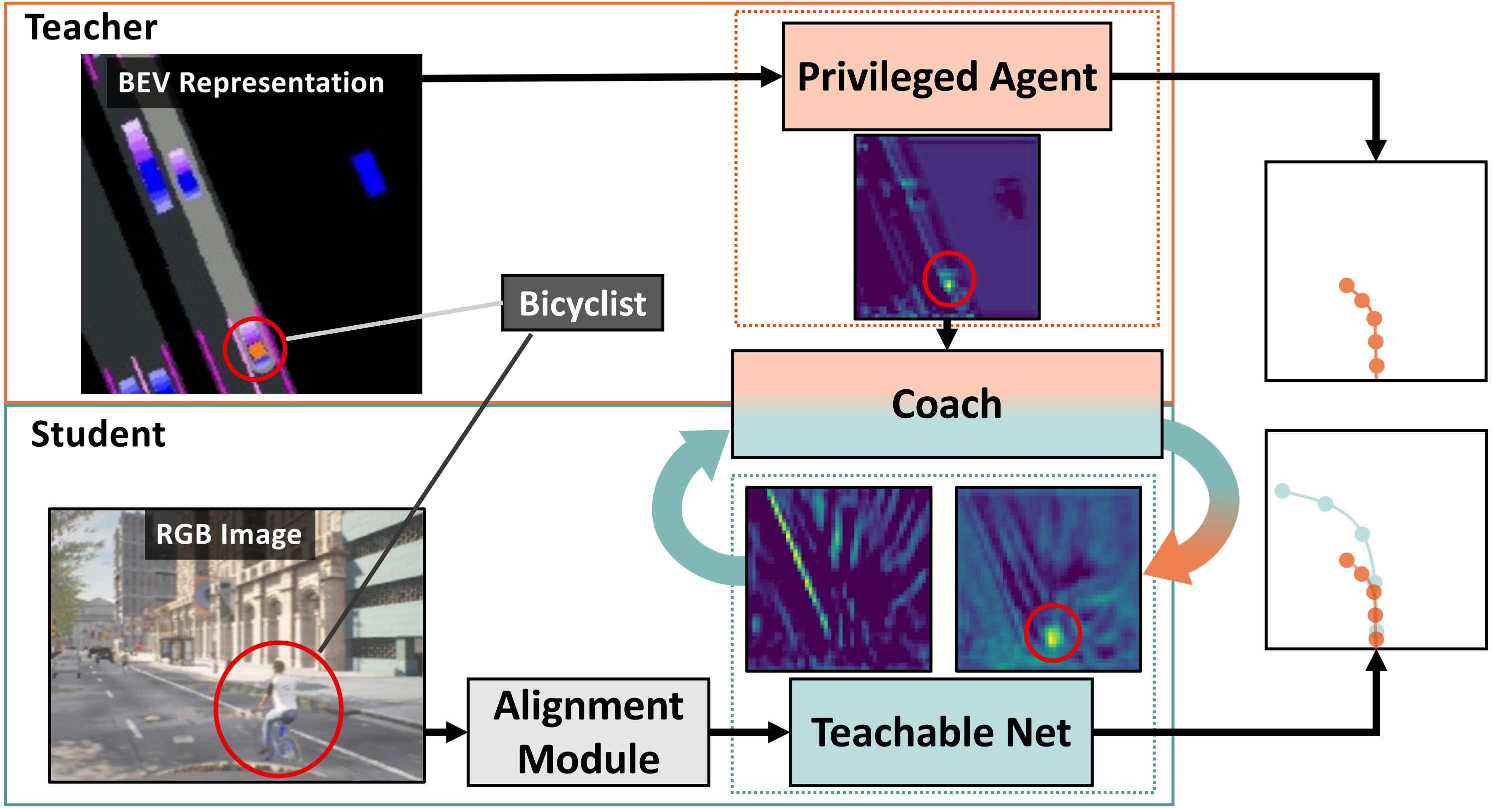

Step 1: Learn an ***Effective Teacher*** from a Privileged BEV with ***Safety Hints***

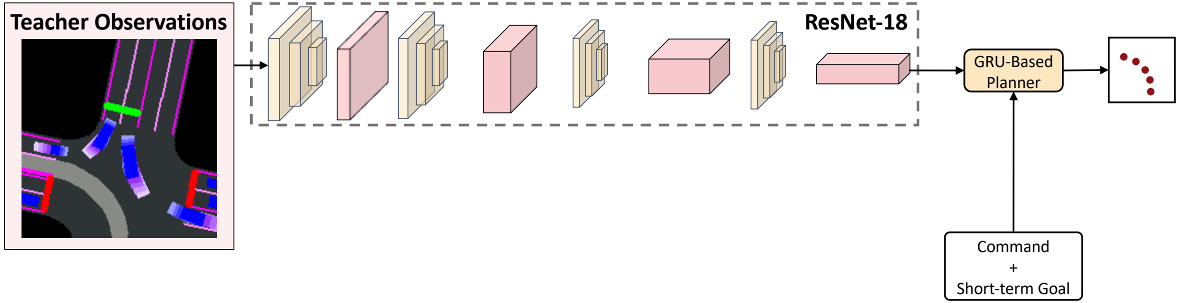

Step 2: Learn a Student Model with an Image-to-BEV ***Feature Alignment Module***

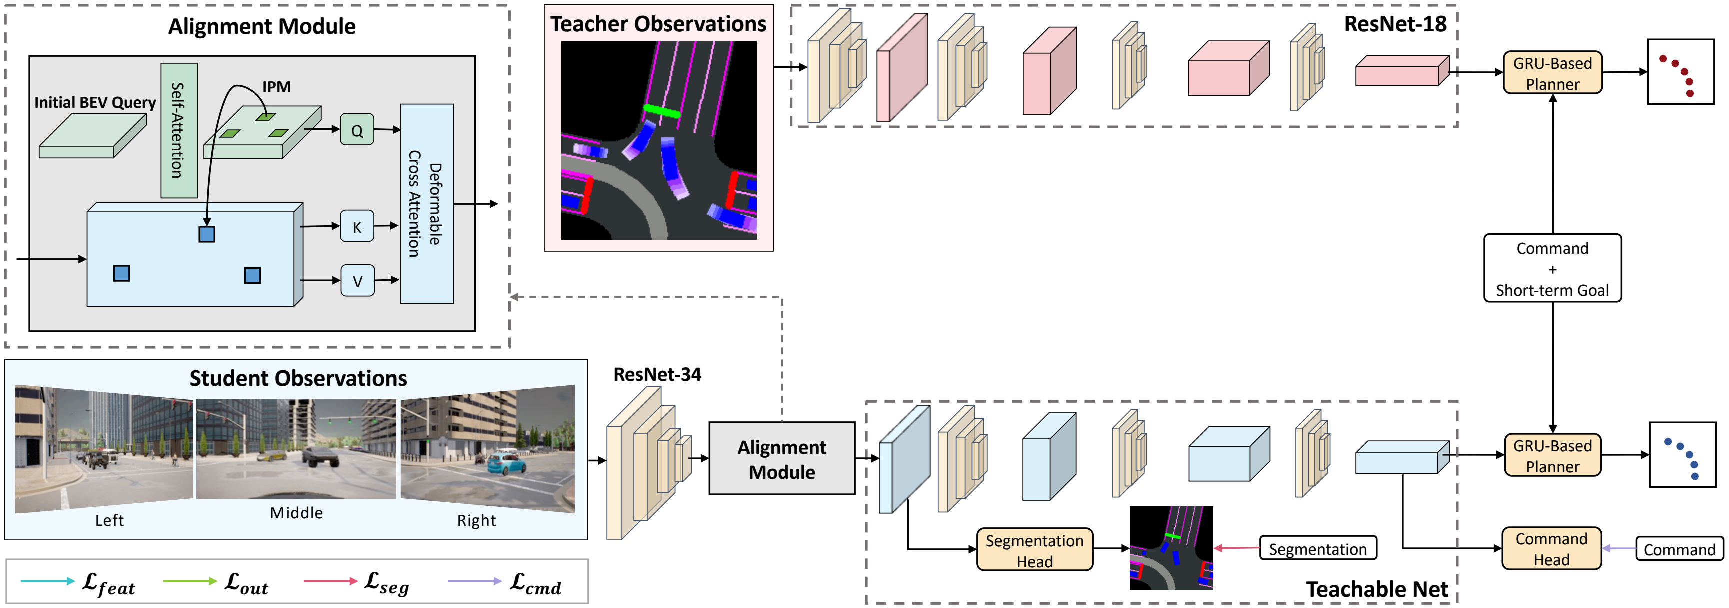

Step 2: Learn a Student Model with an Image-to-BEV ***Feature Alignment Module***

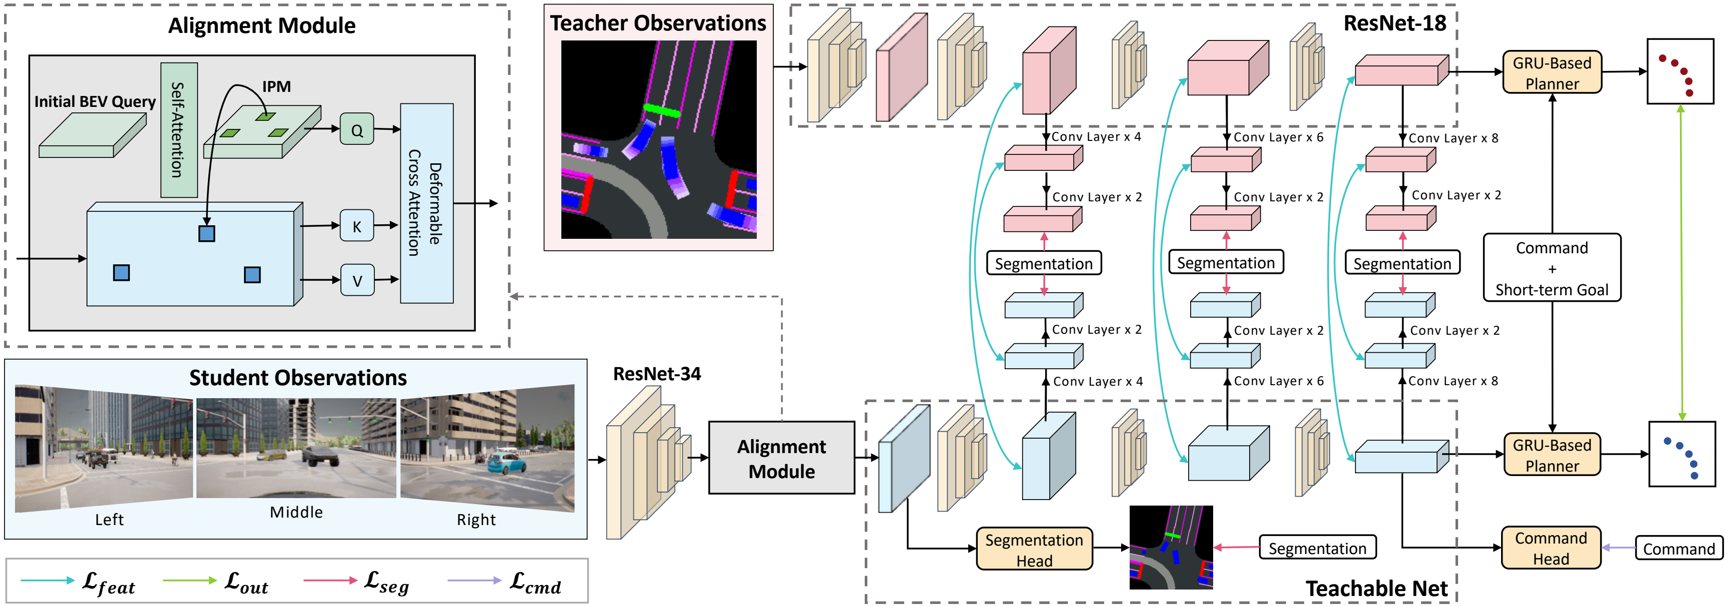

Step 3: ***Student-paced Coaching*** Scaffolds the Difficult Sensorimotor Learning Task

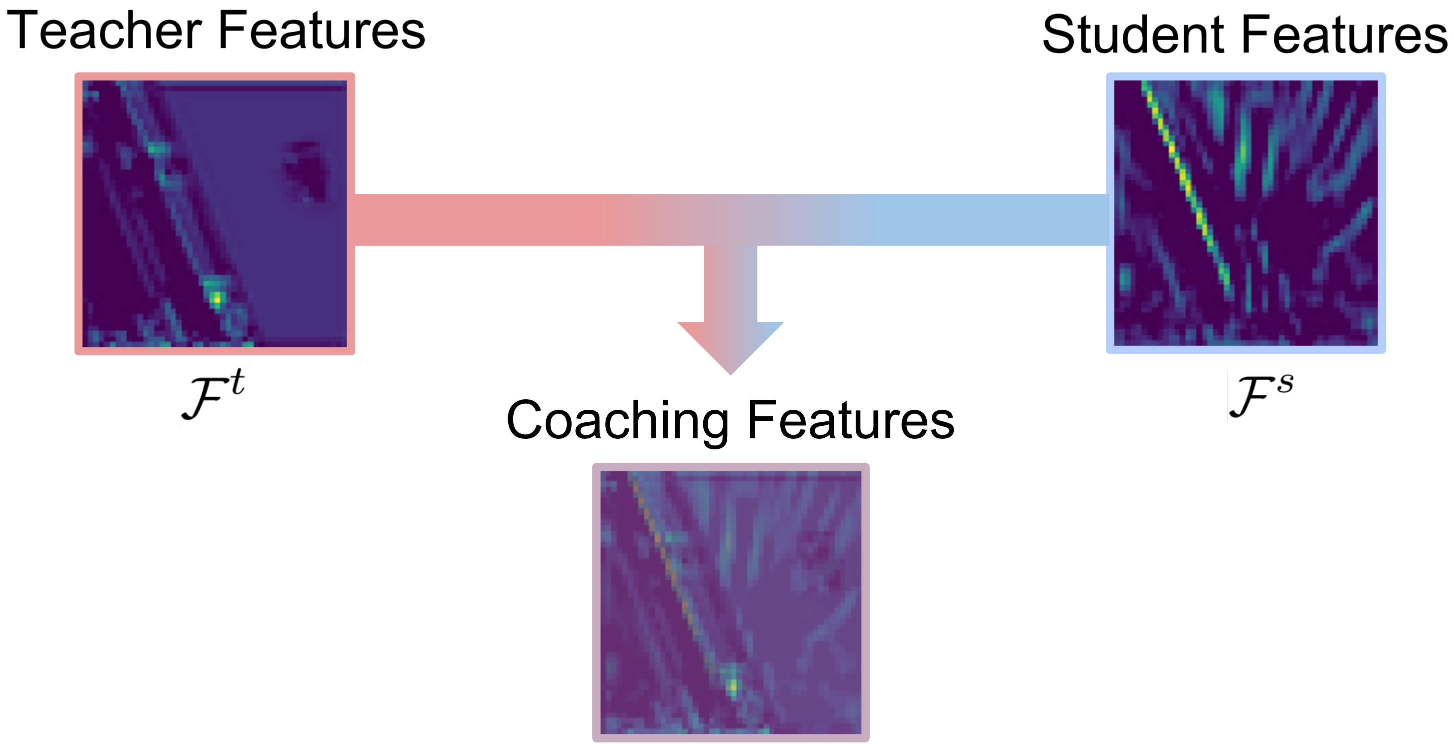

$$\mathcal{F}^t \leftarrow \lambda_i \mathcal{F}^s + (1-\lambda_i)\mathcal{F}^t, \text{ if } \mathcal{L}_{CaT} > \tau_i$$

Baseline

CaT

Learning from a Privileged Teacher

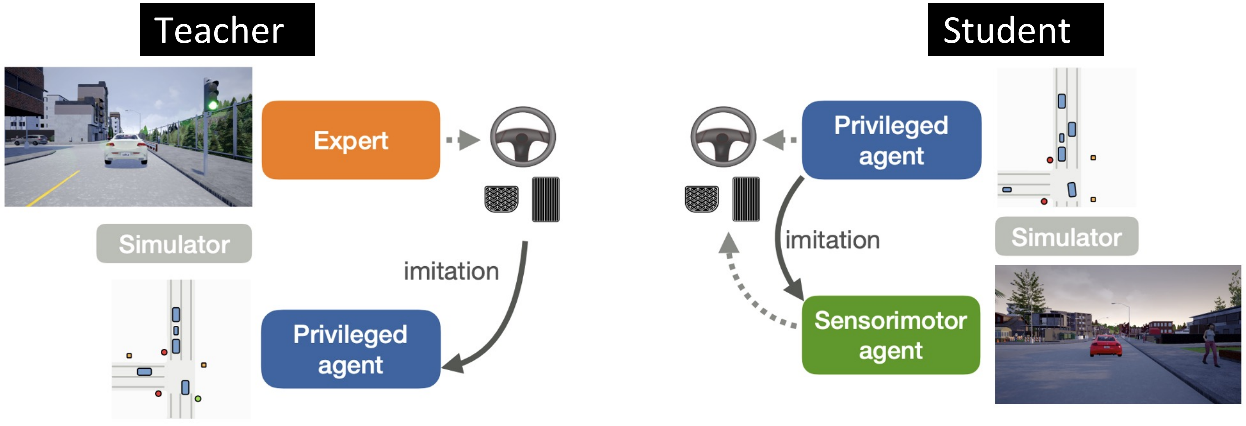

(a) Privileged agent imitates the expert (b) Sensorimotor agent imitates the privileged agent

Does not address:

- *Inherent differences between inputs*
- *Only output distillation – what about internal features?*
- ***Modeling capacity*** *of the student?*

Task	LBC	PV	AT
Empty	70 ± 0	100 ± 0	100 ± 0
Regular	62 ± 2	93 ± 2	99 ± 1
Dense	39 ± 8	45 ± 10	59 ± 6

Method	DS ↑	RC ↑	IS ↑
WOR [110]	20.53 ± 3.12	48.47 ± 3.86	$\mathbf{0.56} \pm 0.03$
Latent TransFuser (Ours)	37.31 ± 5.35	$\mathbf{95.18} \pm 0.45$	0.38 ± 0.05
LAV [46]	32.74 ± 1.45	70.36 ± 3.14	0.51 ± 0.02
Late Fusion (LF)	22.47 ± 3.71	83.30 ± 3.04	0.27 ± 0.04
Geometric Fusion (GF)	27.32 ± 0.80	91.13 ± 0.95	0.30 ± 0.01
TransFuser (Ours)	$\mathbf{47.30} \pm 5.72$	93.38 ± 1.20	0.50 ± 0.06
Expert	*76.91 ± 2.23*	*88.67 ± 0.56*	*0.86 ± 0.03*

Chen, et al. CoRL 2020, He, et al. NeurIPS 2013, Weihs, et al. NeurIPS, 2021, Chitta et al., PAMI 2022

CaT: Coaching a Teachable Student

We propose an effective ***deep knowledge distillation*** for sensorimotor students with:

(1) A strong teacher model

(2) ***Alignment module*** for transforming image features to BEV space

(3) A ***coaching optimization mechanism*** for scaffolding the difficult learning task

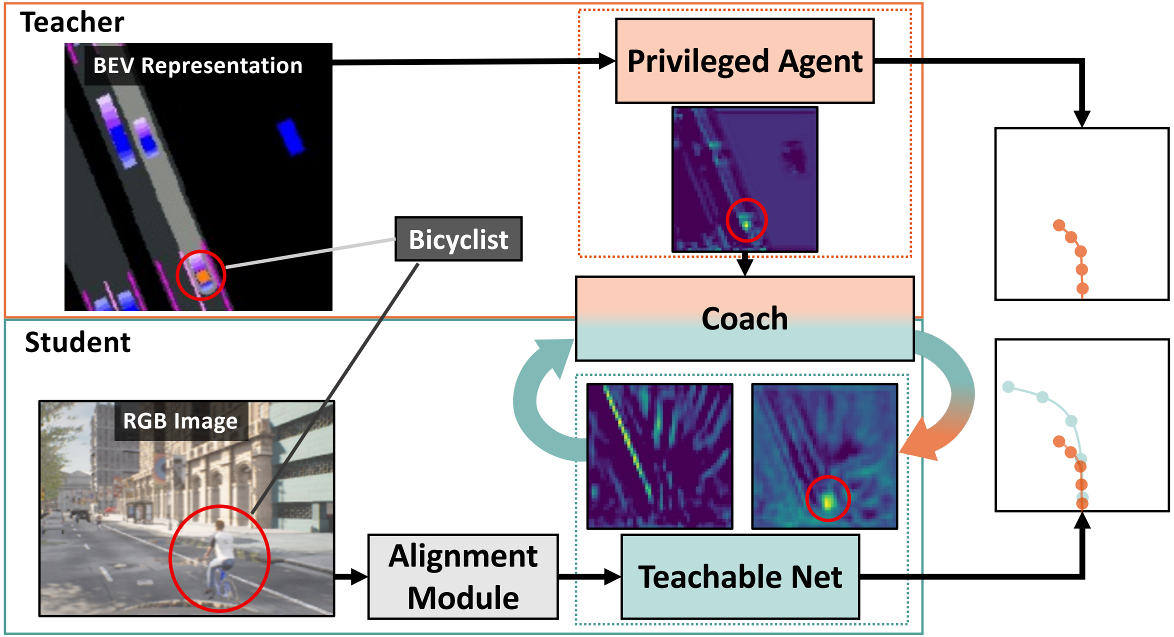

Problem Setup

Objective: Given a dataset $\mathcal{D}$ comprising sensory and privileged observations and a loss function $\mathcal{L}$, the student can be optimized from the teacher using

$$\boxed{\underset{\theta}{\operatorname{argmin}}\, \mathbb{E}_{(\mathbf{x}^s,\mathbf{x}^t)\sim\mathcal{D}}\left[\mathcal{L}(\mathcal{F}^s(\mathbf{x}^s;\theta),\mathcal{F}^t(\mathbf{x}^t;\psi))\right]}$$

Three RGB Camera Views:	$\mathbf{I} = [\mathbf{I_0}, \mathbf{I_1}, \mathbf{I_2}] \in \mathbb{R}^{W\times H\times 3}$
Privileged Bird's-Eye-View (BEV):	$\mathbf{B} \in \{0,1\}^{W_B\times H_B\times C_B}$
Student Observations:	$\mathbf{x}^s = (\mathbf{I}, \mathbf{g}, c) \in \mathcal{X}^s$
Teacher Observations:	$\mathbf{x}^t = (\mathbf{B}, \mathbf{g}, c) \in \mathcal{X}^t$
Student Agent:	$f_\theta^s\colon \mathcal{X}^s \to \mathcal{Y}$, $\theta \in \mathbb{R}^d$
Teacher Agent:	$f_\psi^t\colon \mathcal{X}^t \to \mathcal{Y}$, $\psi \in \mathbb{R}^d$
Student Network Feature Maps:	$\mathcal{F}^s(\cdot;\theta)$
Teacher Network Feature Maps:	$\mathcal{F}^t(\cdot;\psi)$
Categorical Navigational Command:	$c \in \{1,\ldots,6\}$

Learning an Effective Teacher

Incorporating explicit safety-driven (*Agent Forecast*, *Entity Attention*) cues to BEV results in a strong teacher

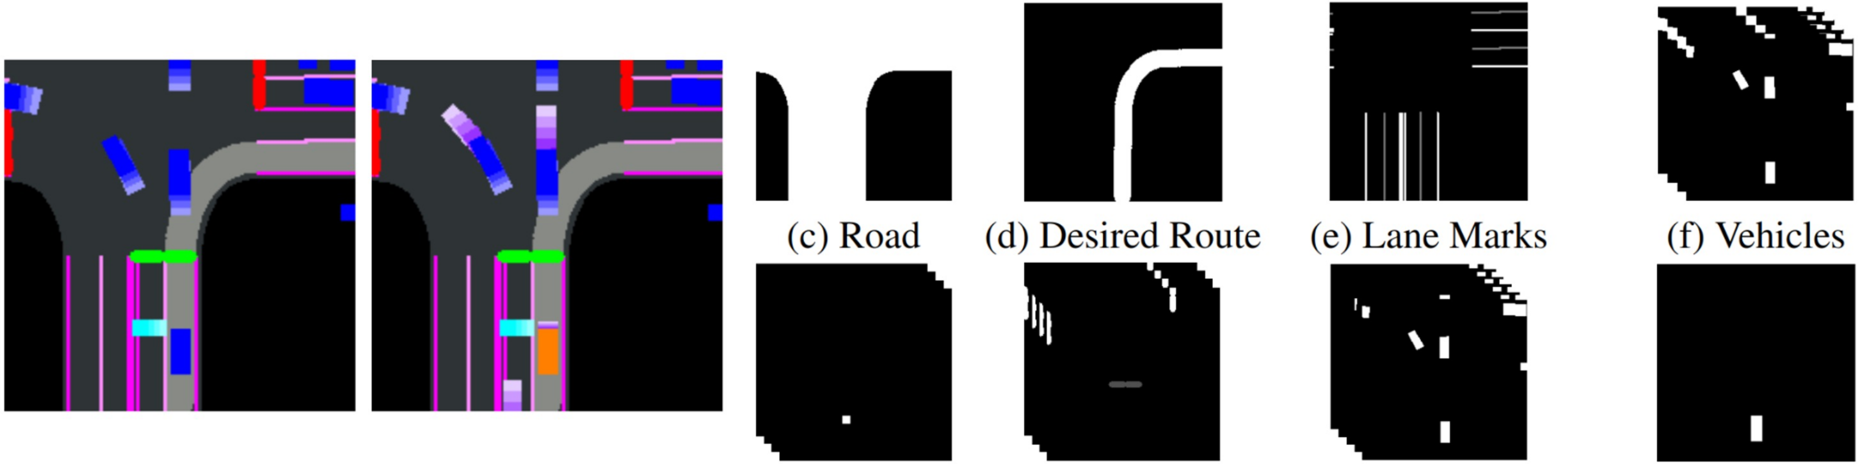

(a) Baseline [79] (b) With Safety Hints (c) Road (d) Desired Route (e) Lane Marks (f) Vehicles (g) Pedestrians (h) Traffic Lights (i) Agent Forecast (j) Entity Attention

Learning a Teachable Student

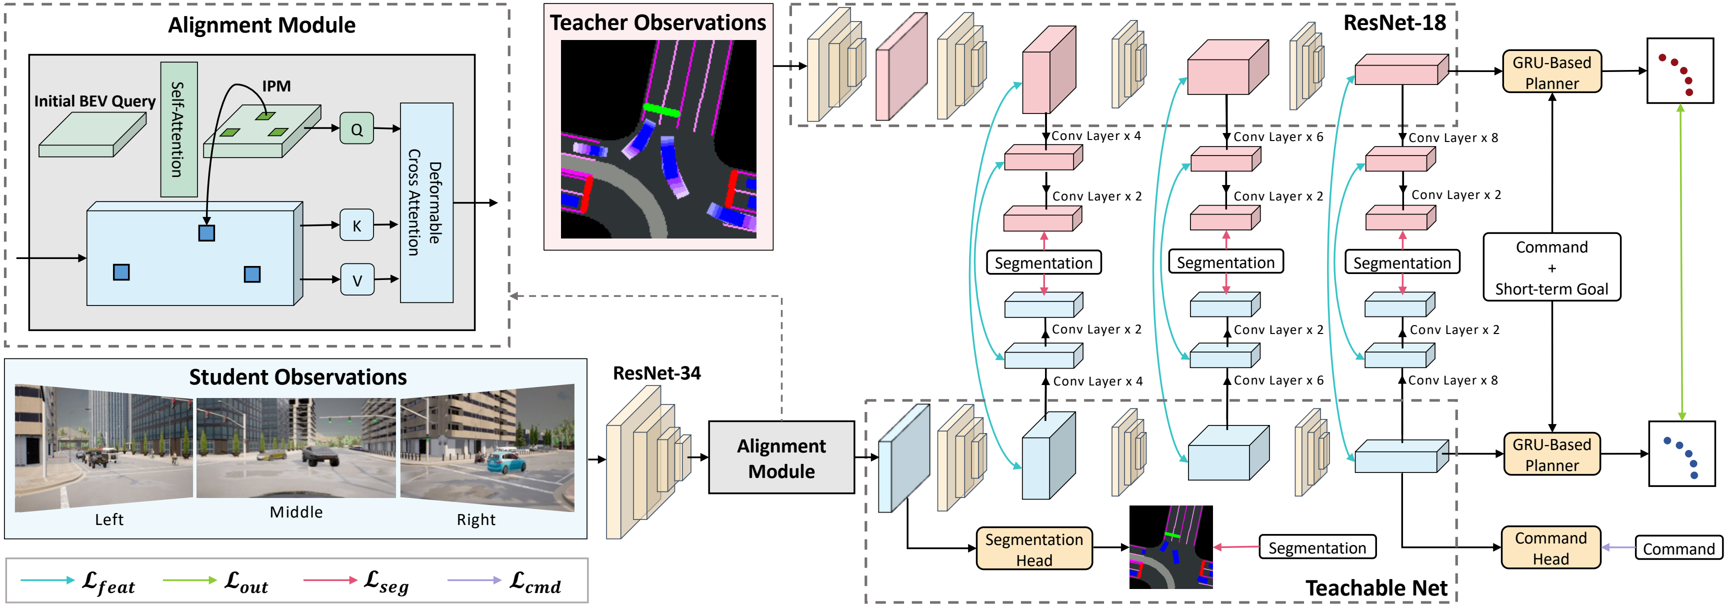

Loss Function: $\mathcal{L}_{CaT} = \mathcal{L}_{out} + \mathcal{L}_{feat} + \mathcal{L}_{seg} + \mathcal{L}_{cmd}$

Output Distillation: $\mathcal{L}_{out} = \sum_{c=1}^{C} \left\| f_\theta^s(\mathbf{x}^s, c) - f_\psi^t(\mathbf{x}^t, c) \right\|_1$

Feature Distillation: $\mathcal{L}_{feat} = \sum_{i=1}^{3} [\left\| \mathcal{F}_i^s(\mathbf{x}^s) - \mathcal{F}_i^t(\mathbf{x}^t) \right\|_2 + \left\| T_i^s(\mathcal{F}_i^s(\mathbf{x}^s)) - T_i^t\left(\mathcal{F}_i^t(\mathbf{x}^t)\right) \right\|_2 + \lambda_{CD} \left\| \mathcal{F}_i^s(\mathbf{x}^s) - \mathcal{F}_i^t(\mathbf{x}^t) \right\|_{CD}]$

Student-paced Coaching

Student-paced coaching adjusts the learning rate in a sample-selective manner, which aims to stabilize training by ***reducing the difficulty*** when the student is unable to perform the optimal action.

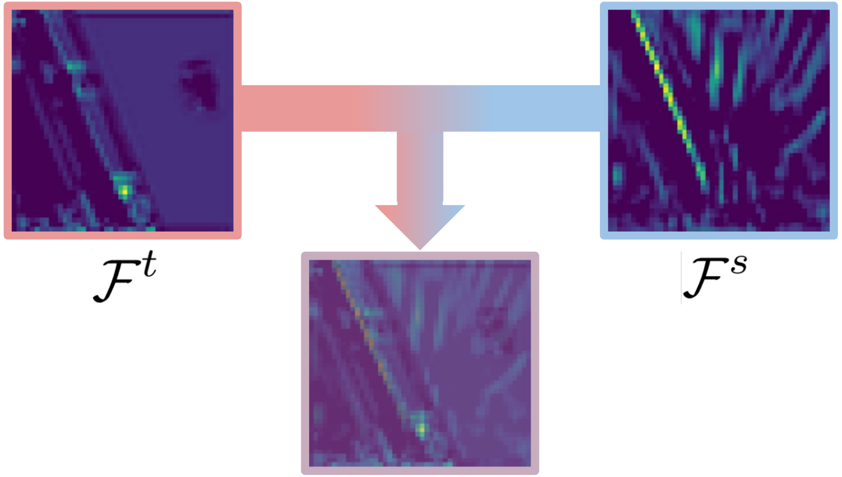

$$\mathcal{F}^t \leftarrow \lambda_i \mathcal{F}^s + (1-\lambda_i)\mathcal{F}^t, \text{ if } \mathcal{L}_{CaT} > \tau_i$$

Results

Method	RGB	LiDAR	DS ↑	RC ↑	IS ↑
LAV [11]	✓	✓	48.41±3.40	80.71±0.84	0.60±0.04
TransFuser [16]	✓	✓	46.20±2.57	**83.61±1.16**	0.57±0.00
WOR [10]	✓	✗	17.36±2.95	43.46±2.99	0.54±0.06
NEAT [15]	✓	✗	24.08±3.30	59.94±0.50	0.49±0.02
TCP* [71]	✓	✗	42.86±0.63	61.83±4.19	0.71±0.04
CaT (w/o Alignment, Coaching, FD)	✓	✗	39.48±0.67	60.96±1.65	0.68±0.01
CaT (w/o Alignment, Coaching)	✓	✗	40.64±0.98	62.45±0.46	0.67±0.01
CaT (w/o Coaching, FD, SH)	✓	✗	44.10±0.40	65.84±5.55	0.72±0.03
CaT (w/o Coaching, SH)	✓	✗	49.69±2.28	81.10±0.58	0.64±0.02
CaT (w/o Coaching)	✓	✗	55.55±1.41	81.97±2.34	0.68±0.01
CaT	✓	✗	**58.36±2.24**	78.79±1.50	**0.77±0.02**
Privileged Agents:					
RL Expert (Roach) [79]	-	-	60.14±2.40	85.83±0.60	0.69±0.03
Rule-based Expert	-	-	71.96±2.13	77.46±3.11	**0.91±0.00**
Basic BEV Agent [13]	-	-	24.08±2.83	73.36±1.08	0.31±0.06
+ History and Desired Path	-	-	52.81±1.79	79.34±3.65	0.71±0.06
+ Agent Forecast	-	-	65.73±0.93	83.50±1.18	0.79±0.02
+ Entity Attention	-	-	**73.30±1.07**	**87.44±0.28**	0.83±0.02

Results

Ablation Study on Feature Distillation Layers

Method	DS ↑	RC ↑	IS ↑
No Distillation	44.10	65.84	0.72
One Layer [71, 79]	45.23	69.33	0.68
Three Layers $\mathcal{L}_2$	49.31	66.92	0.78
Three Layers $\mathcal{L}_2 + \mathcal{L}_{CD}$	51.95	62.82	**0.87**
Three Layers $\mathcal{L}_{feat}$	**55.55**	**81.97**	0.68

Open-Loop Evaluation on nuScenes.

Method	ADE (m) ↓	FDE (m) ↓	Coll. (%) ↓
BEV Agent	0.33	0.52	0.49
CaT (w/o Coaching, FD, SH)	0.48	0.43	0.68
CaT	**0.41**	**0.36**	**0.27**

Qualitative Results

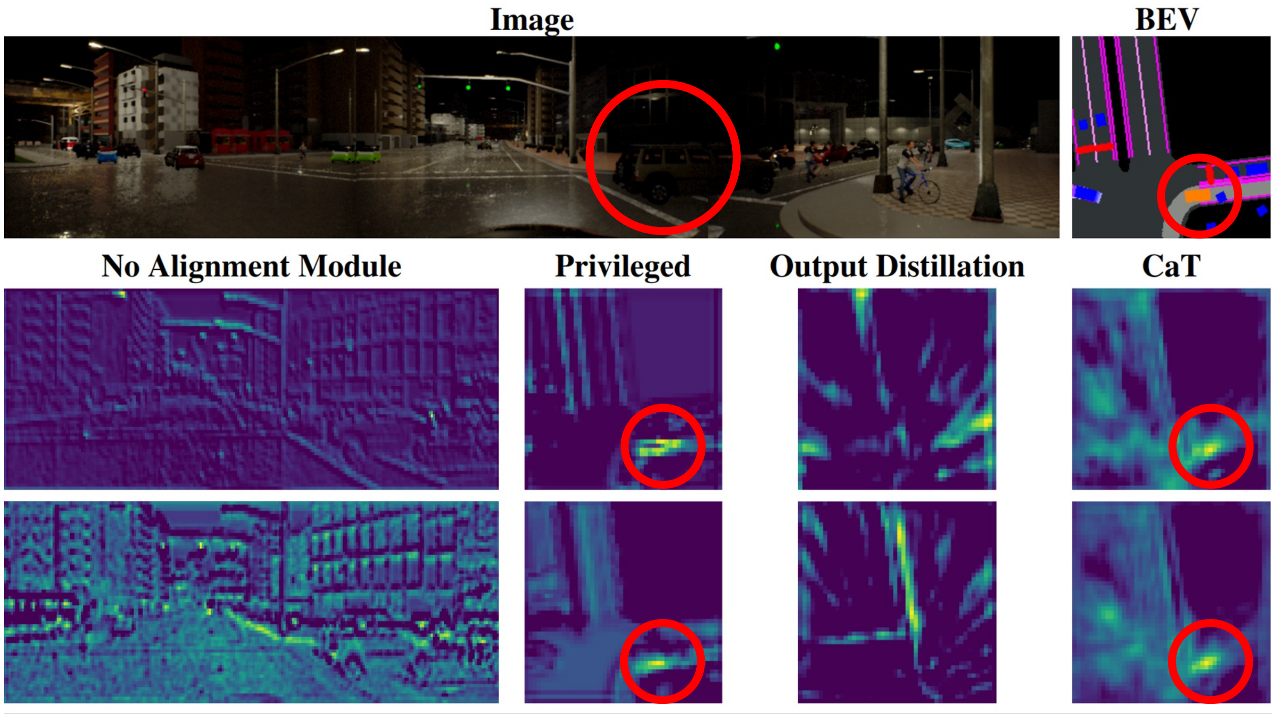

Scenario: Night-time driving with agent turning right at an intersection with a vehicle in the way.

Qualitative Results

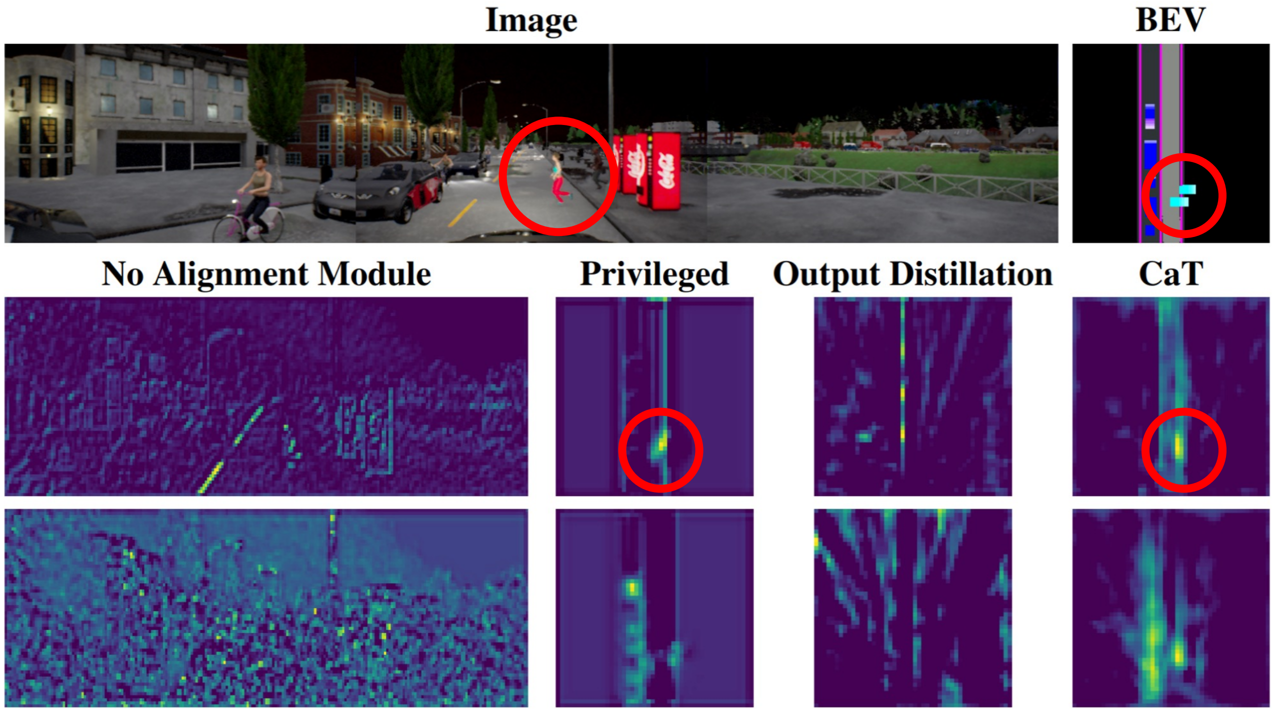

Scenario: Night-time driving with a pedestrian abruptly emerging from the right.

Baseline

CaT

Summary

- Develop an ***alignment module***, enabling extensive supervision from the privileged teacher over the intermediate feature learning
- Incorporate explicit ***safety-aware cues*** into the BEV space that facilitate an effective teacher agent
- Design ***student-paced coaching*** that scaffolds knowledge and leads to improved model optimization by considering the learning ability of the student

Thank You for Watching!

JUNE 18-22, 2023 CVPR VANCOUVER, CANADA

BOSTON UNIVERSITY

H2X Lab