Gatik

We own the middle mile.™

ML4AD Award Ceremony

Aleksandr Petiushko, Head of AI Research

AAAI 2025, ML4AD

March 4, 2025

Company Overview

Background

The Leader In Autonomous Short-Haul Logistics

- Founded in 2017 by veterans of the autonomous technology industry
- Customers: Walmart, Kroger, Tyson Foods, Georgia-Pacific and more
- Current locations include Texas, Arkansas & Ontario (Canada)
- Expanding to new markets throughout 2025
- Use case leverages point-to-point movement of goods to optimize safety and efficiency and meet customer needs

Class 8 vehicles for highway-only driving between hubs

Distribution Center

Class 3-7 cold chain capable vehicles for urban, semi-urban, and highway driving environments

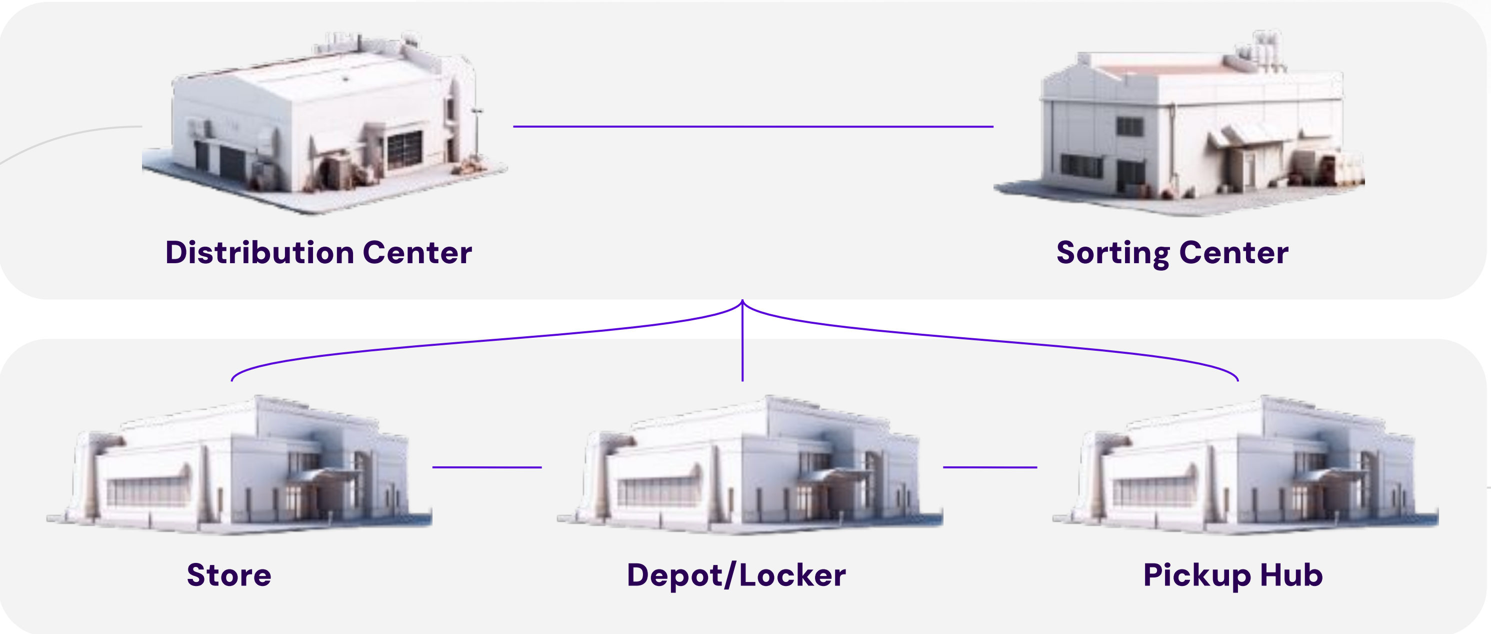

Slower moving vehicles with limited capacity

Home

Definition	Highway only; hub-to-hub; Class 8; >400 miles	Highway & Semi Urban; DC-hub-store; short-haul; Up to 400 miles	Urban; store-to-home; smaller robots; 1-5 miles
Technological Differentiation	Highway-only capabilities	Purpose built technology for fixed & repeatable routes/networks. Tailored for urban, semi-urban & highway driving	Geofenced use-cases Leading to countless route combinations

Gatik

Making the Supply Chain More Responsive and Efficient

Service	Savings	Safety	Sustainability
Increase Product Flow	Lower Costs	Reduce Accidents	Meet emission targets
Driver Shortage Hedge	Higher Utilization	Improve Visibility	Full EV solutions
Dedicated Capacity	Real-time Tracking	Increase Control	Simplify Maintenance
99%+ Delivery On-Time	**Savings 20%+**	**Exemplary Safety Record**	**EV/AV Solution**

Structured Autonomy

Customized solution for restricted route and roadway interactions **shorten validation time** and **optimizes for safe operations**

- **Hyper-Constrained** | Custom-fitting AV technology for known routes
- **Route Optimized for Safety** | Pre-defined and risk-mitigated

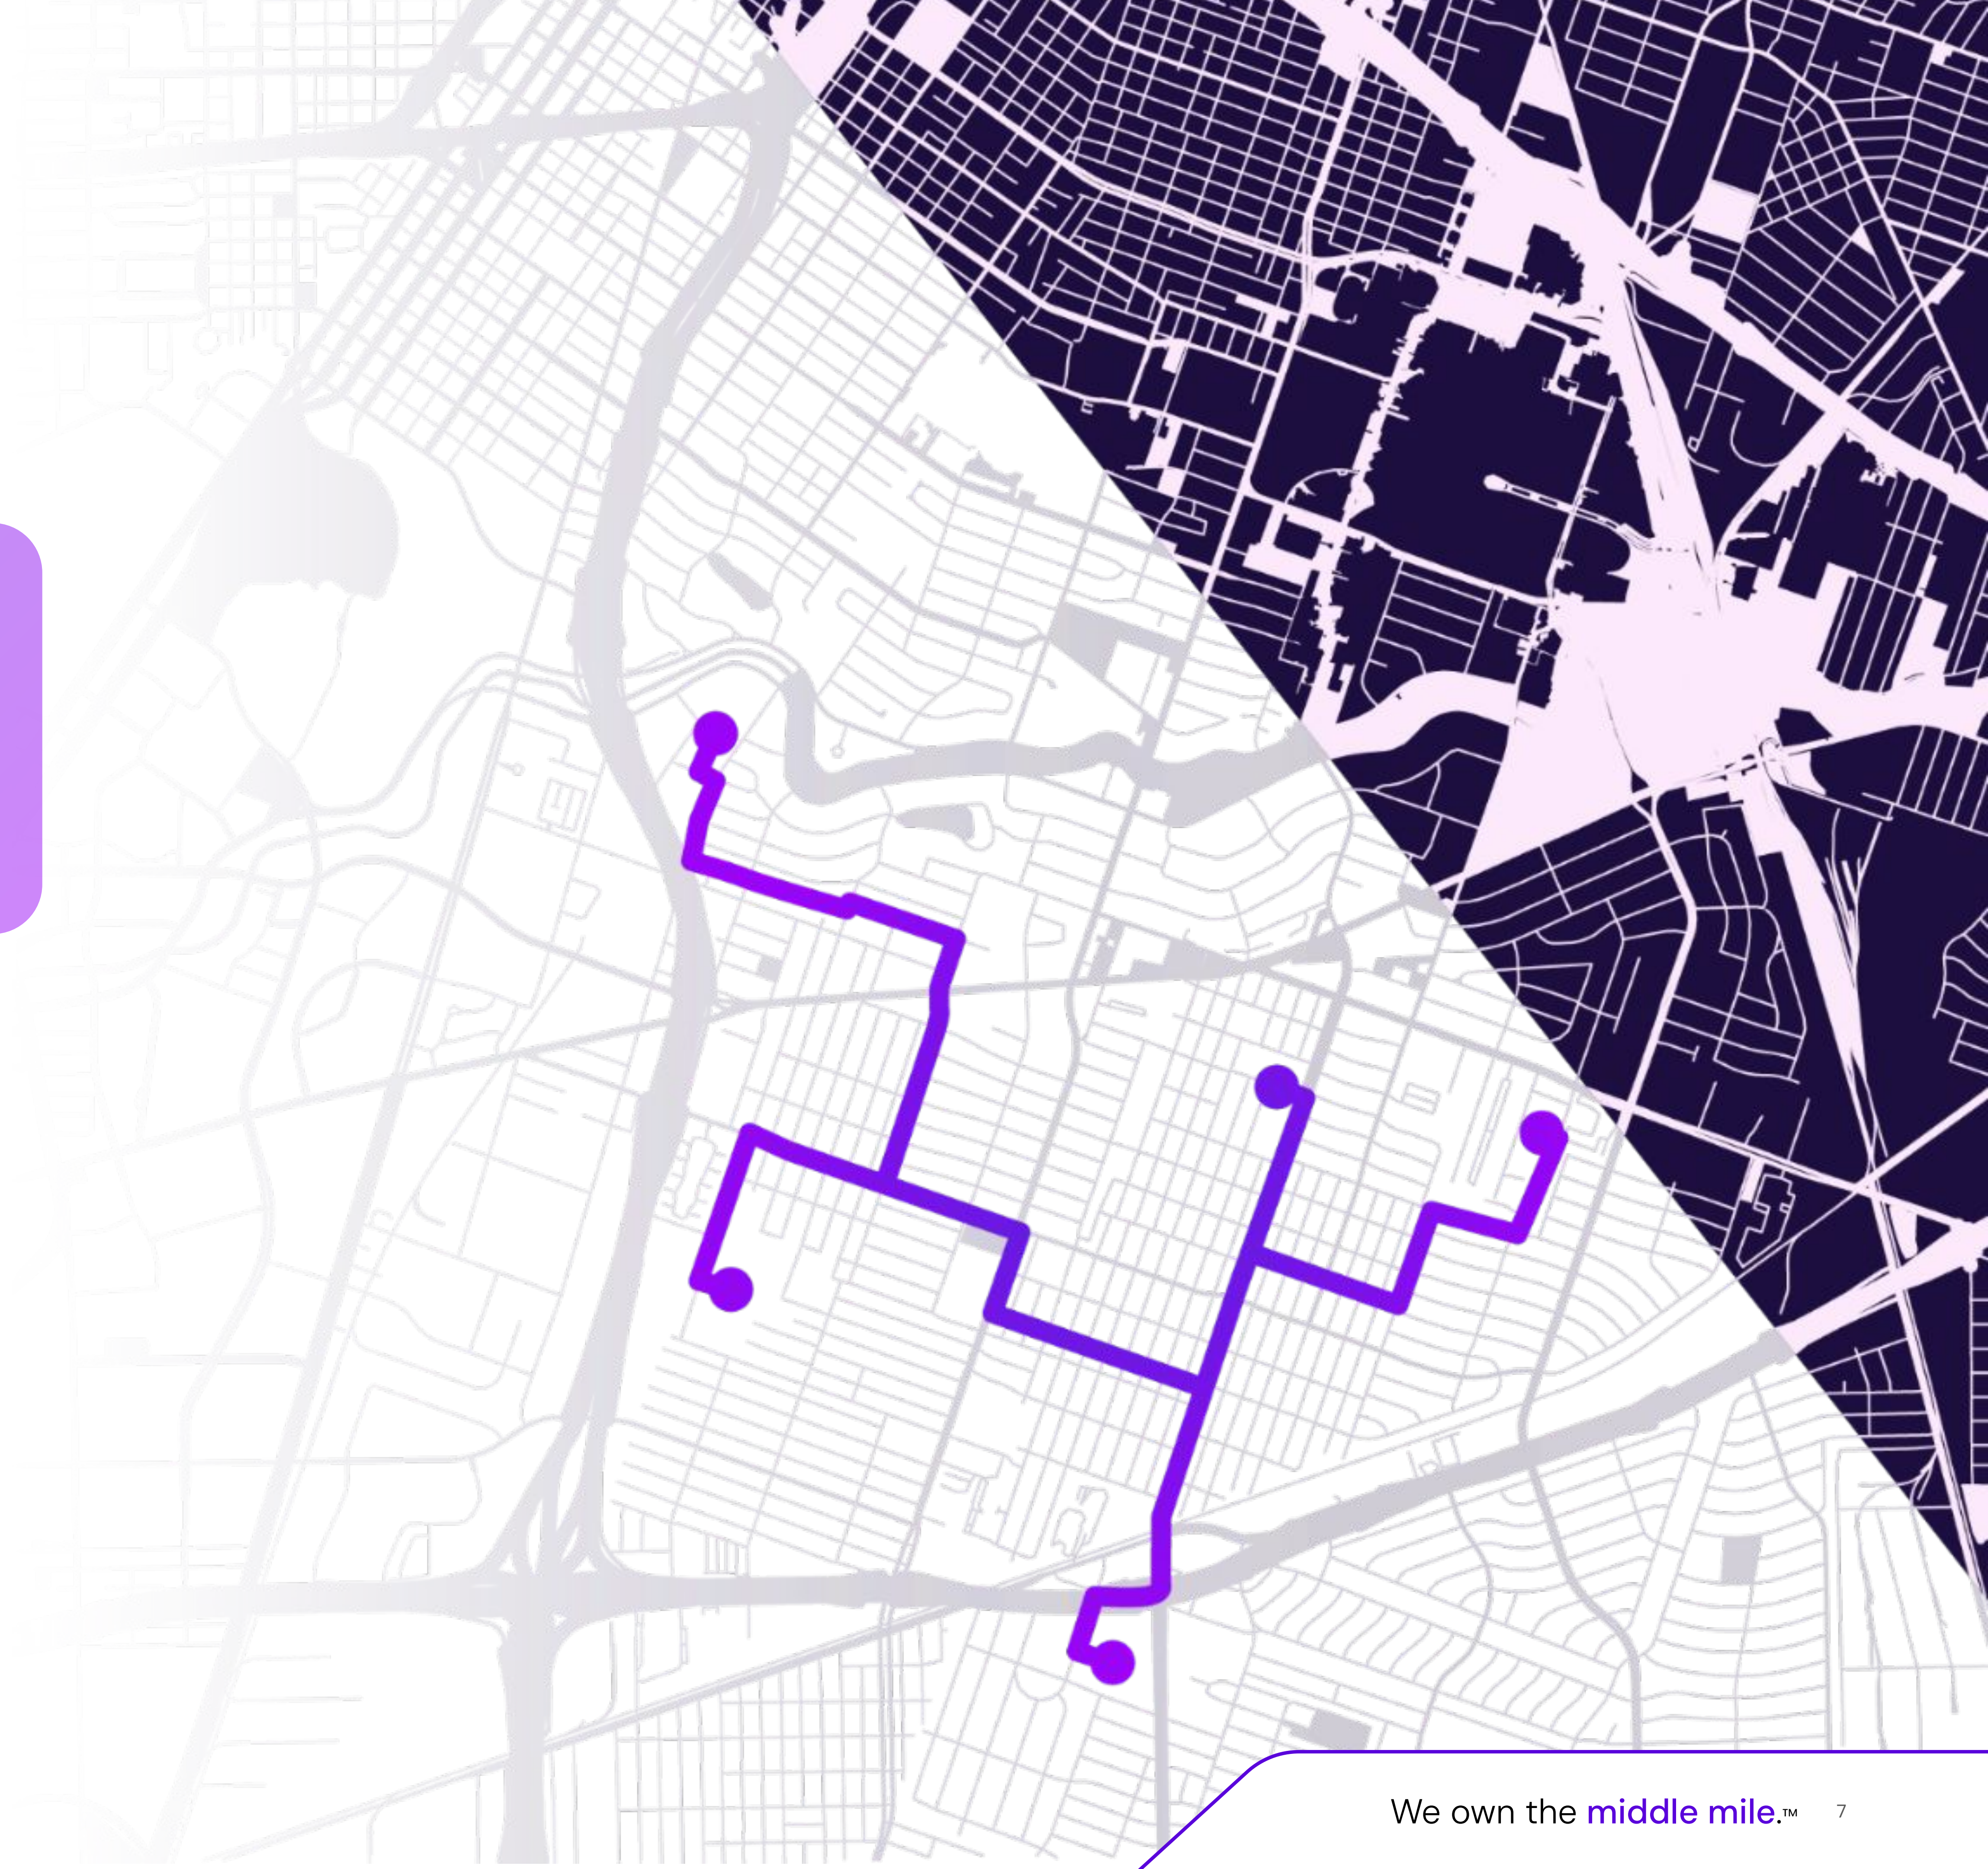

Allows for Incremental Expansion of Operational Design Domain

- Broader ODDs like Geofenced regions: Value proposition is to enable transport between many to many locations. Solving for a single route doesn't really provide any value. **A given route connection may not even see any customer demand during service**
- Before deployment, each of these route variations need to be validated

Each route variation, might or might not provide value - but to enable the service all nodes need to be considered, developed and validated

Approach needs to be generalized and resources are required to be spent to validate all routes

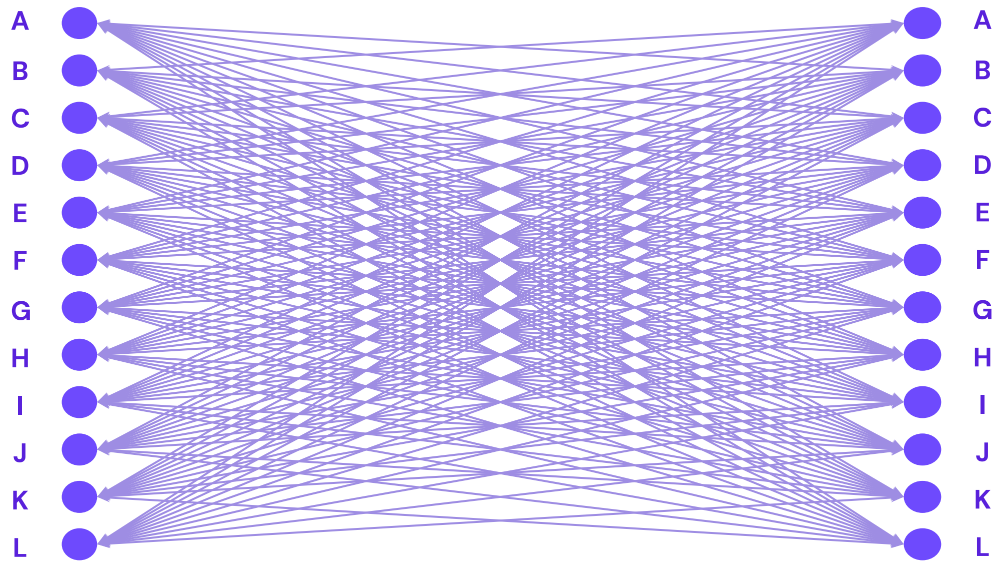

- Gatik's fixed route ODDs: Value proposition is to enable transport between one to one location. **Solving for a single route immediately provides value – Promise of trips – multiple times a day, 7 days a week**
- Before deployment, only specific route(s) needs to be validated.

Each route(s), provides guaranteed value for service.

Targeted use of resources for development & validation of each route – also very high confidence validations

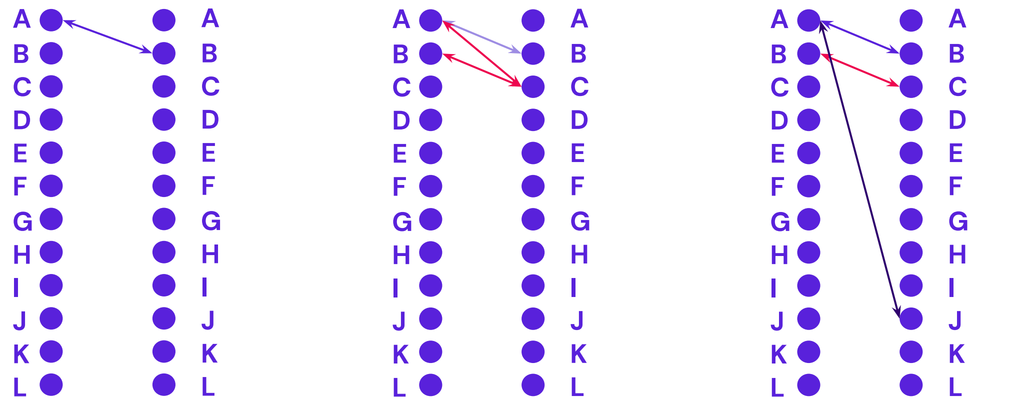

Completely New Routes
Only routes that add commercial value (A<->B) are validated instead of all the variations

Extending Existing Routes
Adding route from B to C, also extends route from A to C. Allows incremental extension

Adding New Destinations
Adding a new route from A to J, doesn't necessarily mean we need to solve for all routes to J

Research Opportunities

- **We are hiring!**
  - Research Scientists
  - ML Infra Engineers
  - Directions:
    - Mapping
    - Perception
    - Behavior (Prediction and Planning)
    - End-to-end Systems
    - Simulation
    - Safety and Uncertainty
  - Apply here: https://gatik.ai/careers/

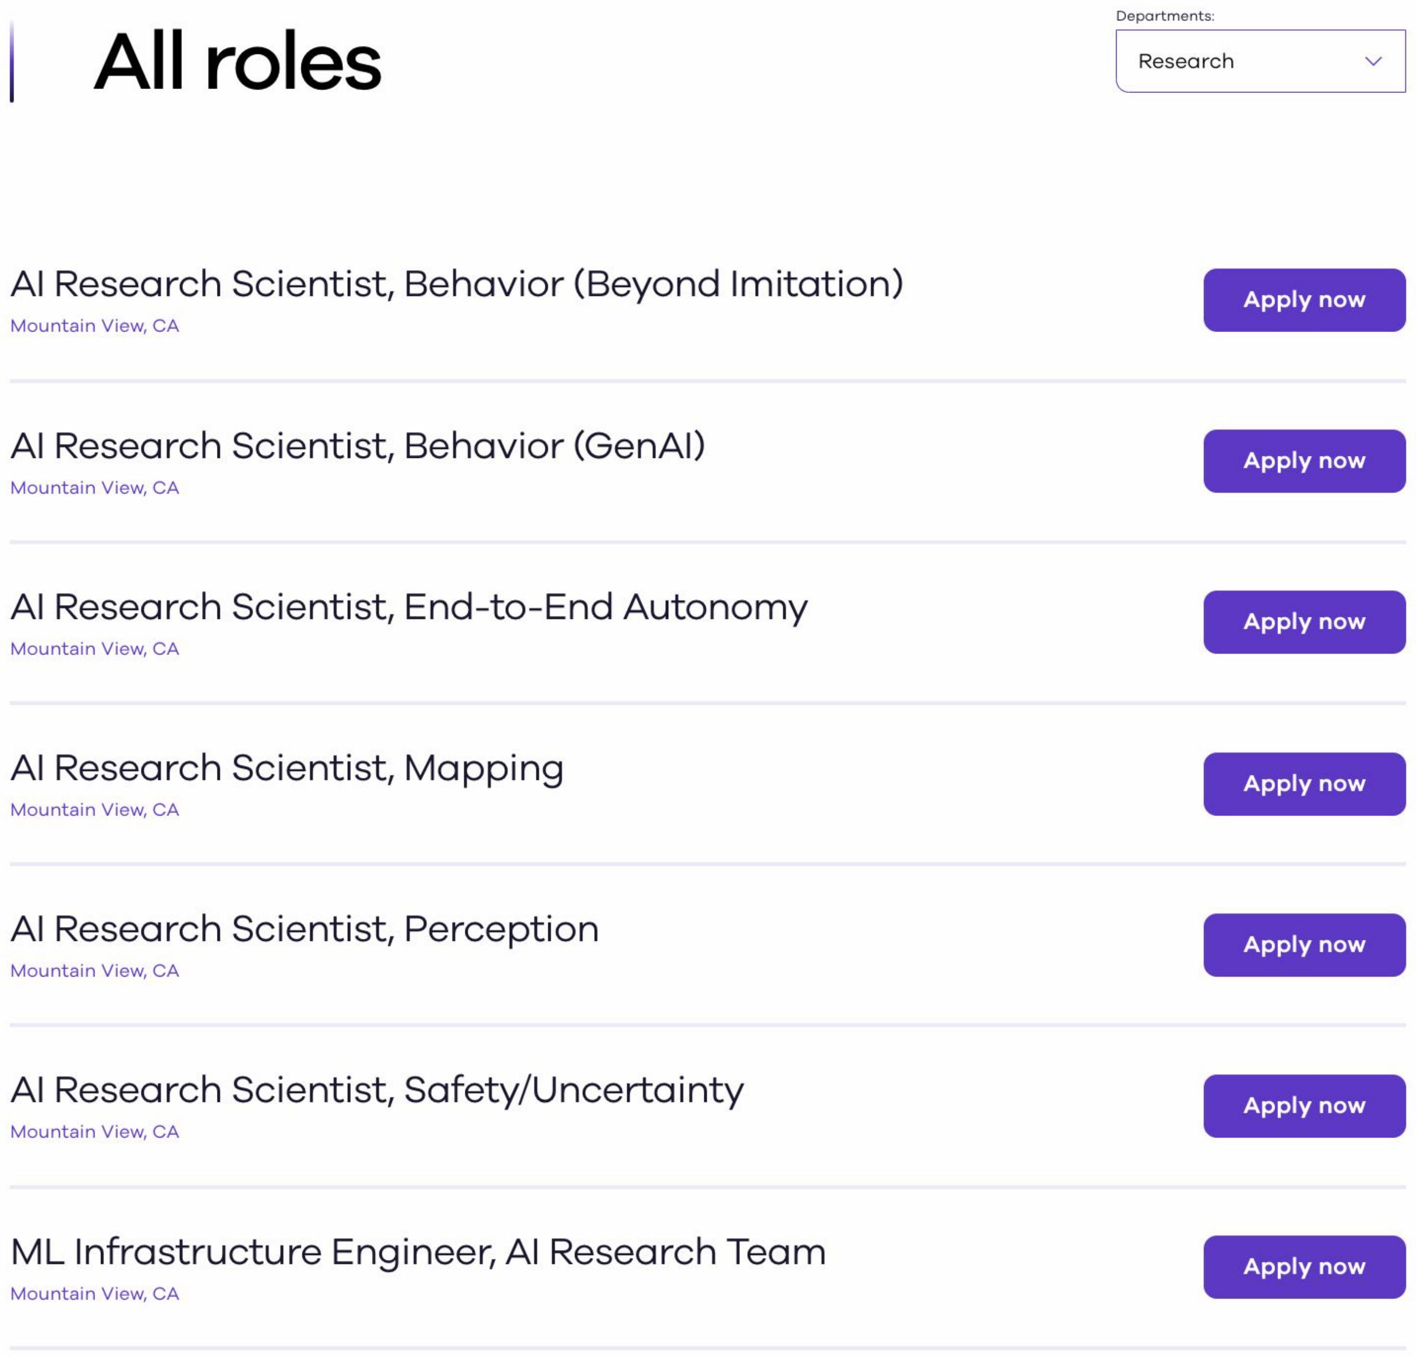

ML4AD Submitted Papers

Awarding

* Prepared together with **Amir Yazdani**, Sr. Research Scientist

Workshop Topics

- Prediction and Planning for AD with LLMs
- Foundation Models for AD
- Mapless Autonomous Driving
- Scaling Laws for AD
- Diffusion modeling for prediction, planning
- Closed loop training and evaluation
- Causal/counterfactual analysis of interactive multi-agent scenarios
- Real-time inference and prediction
- Data-driven AD simulation
- Human driver in the loop for interaction modeling
- Coordination with vehicles (V2V) or infrastructure (V2I)
- Uncertainty propagation through AD software pipelines
- Imitation learning, Reinforcement learning for AD
- Off-road autonomous driving
- Adaptive driving styles based on user preferences
- Metrics/benchmarks for AD

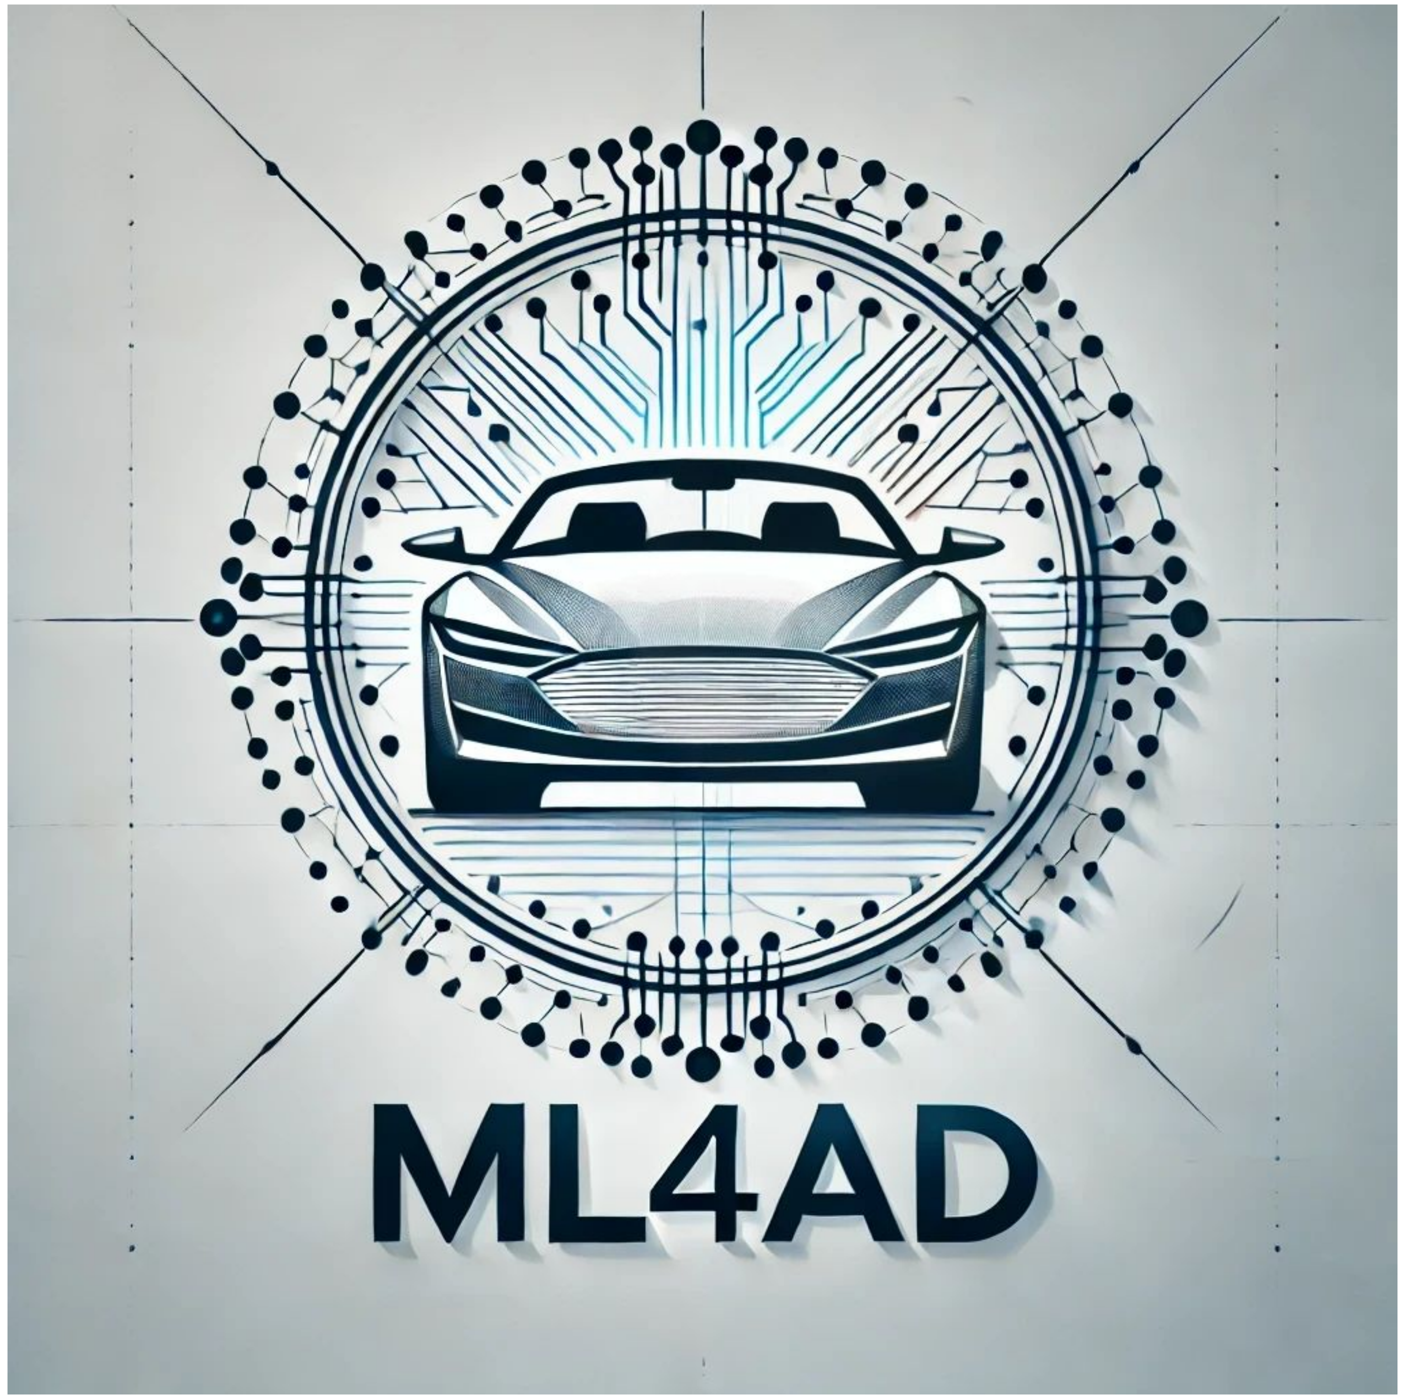

Paul*’s Design :)*

Papers Topics

- **Perception**:
  - 3D
    - Data Augmentation
    - Occupancy Prediction
    - Representation
- **Behavior**:
  - Simulation
    - Real2Sim w/ VLM
    - Diffusion Models
    - RL w/ adversarial agents
    - Tournament Elo Rating
  - Planning
    - Hierarchical approach: RL + Optimization
    - KD

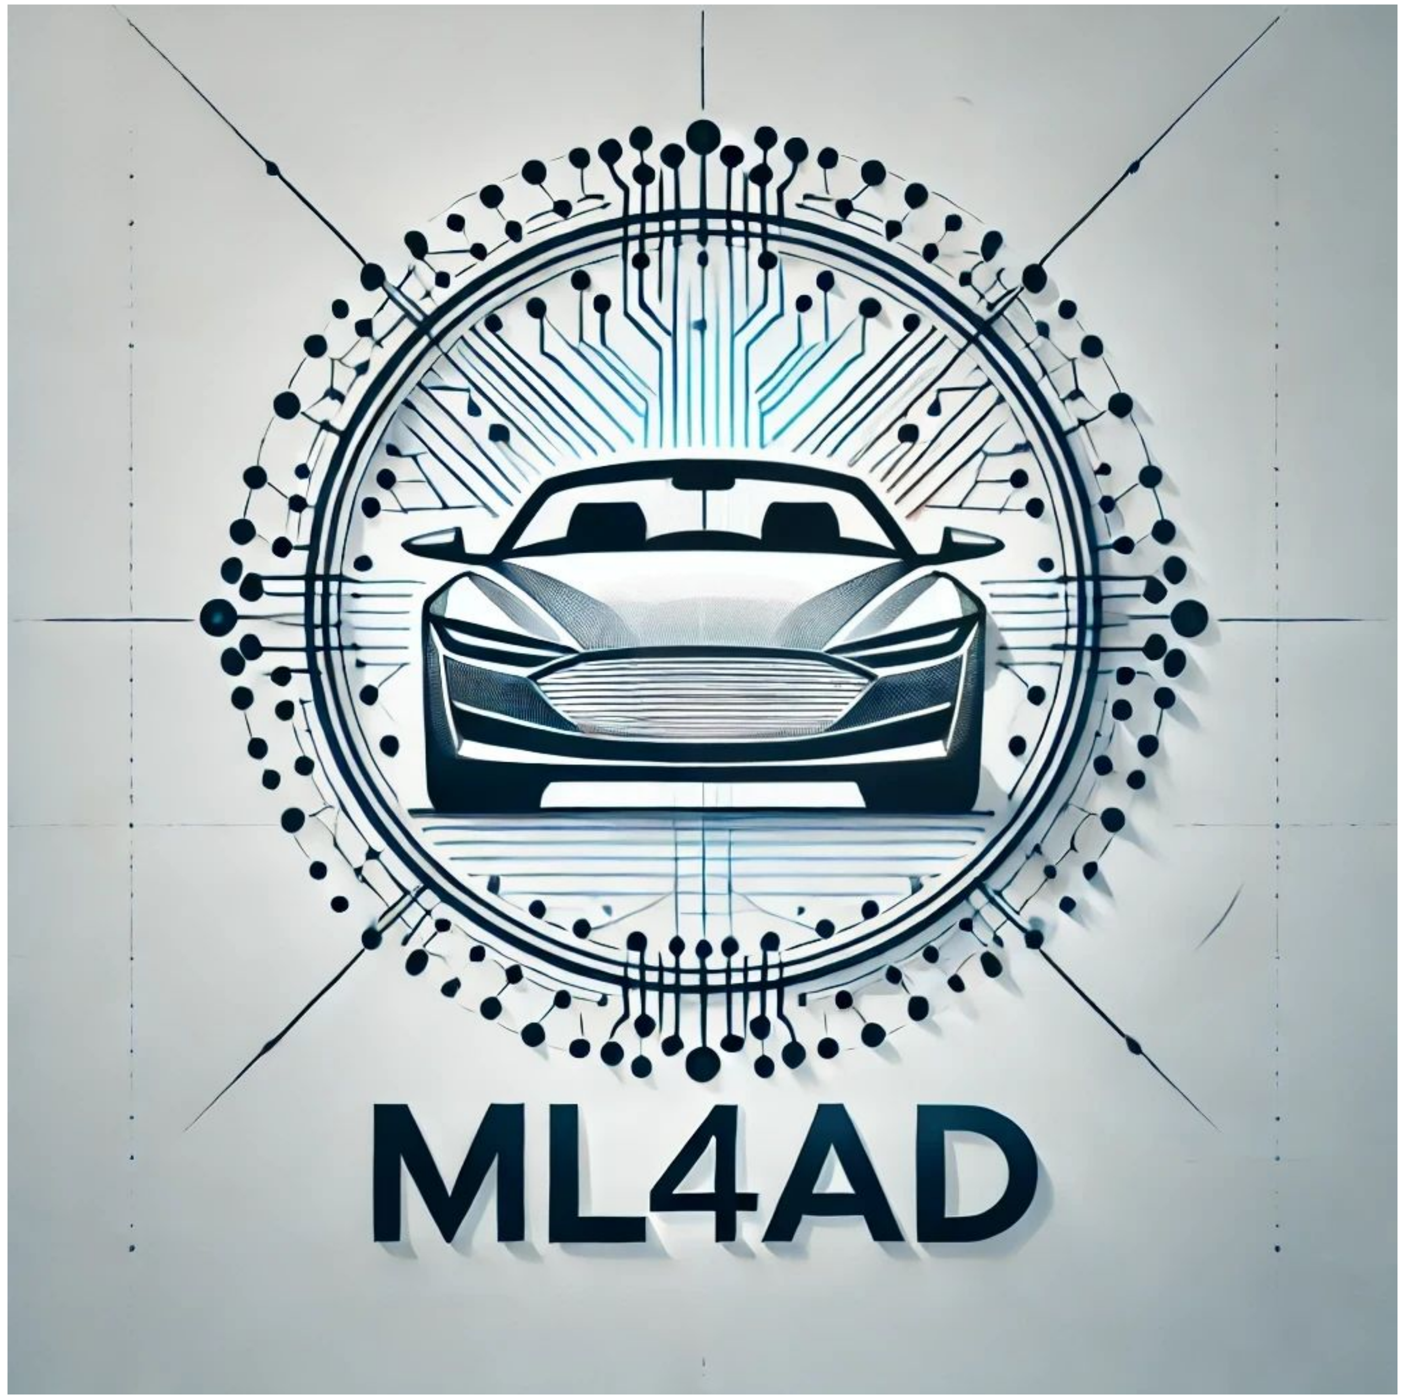

***Paul**'s Design :)*

Gatik

Papers Assessment

- **8 Axes**:
  - Significance
  - Novelty and Innovation
  - Technical Depth
  - Practical Applicability
  - Experimental Rigor
  - Clarity
  - Reproducibility
  - Relevance to Workshop

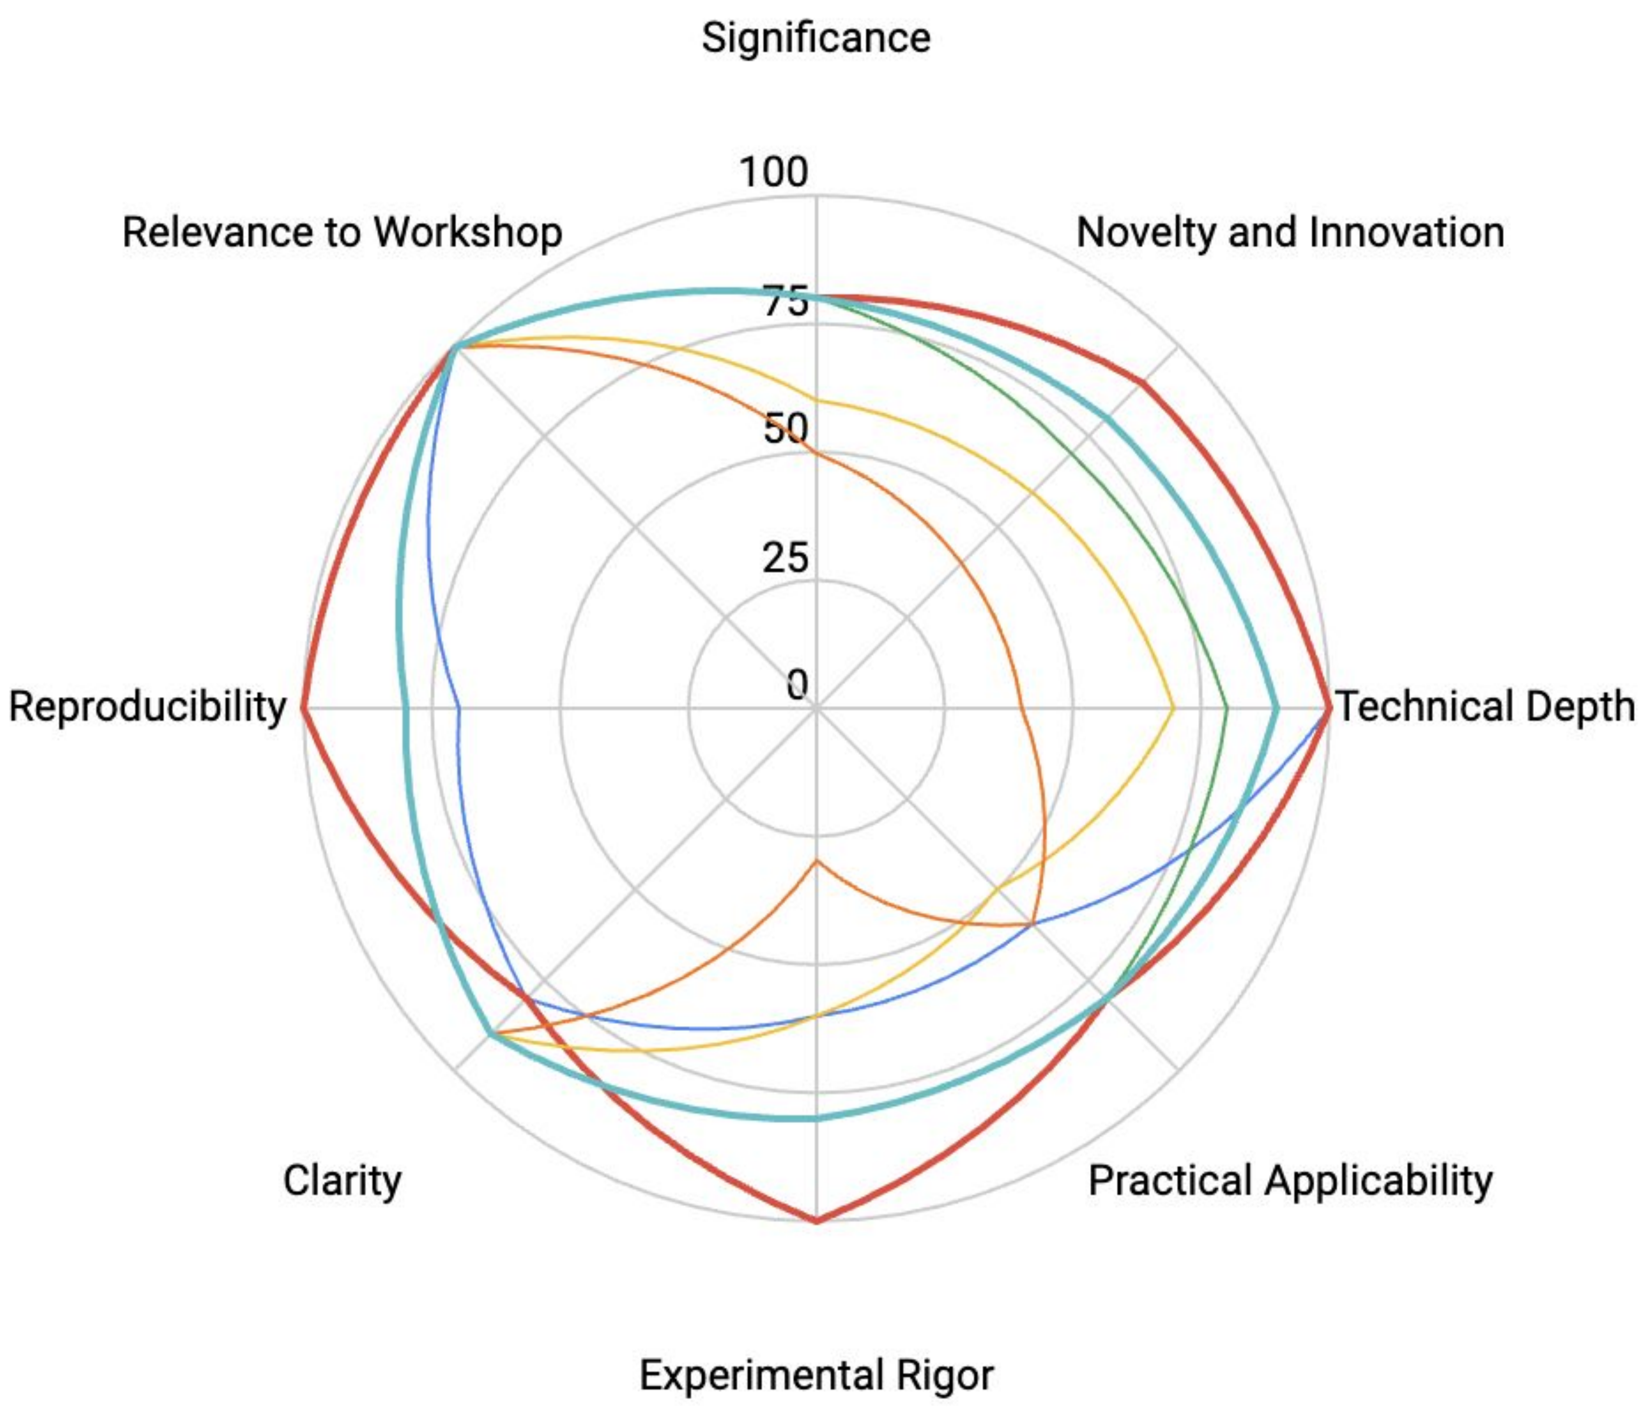

Best Papers

- **Best Perception**:
  - Title: *A Spatiotemporal Approach to Tri-Perspective Representation for 3D Semantic Occupancy Prediction*
  - Authors: *Sathira Silva, Savindu Wannigama, Gihan Jayatilaka, Muhammad Haris Khan, Roshan Ragel*

- **Best Behavior**:
  - Title: *Teacher-guided Off-road Autonomous Driving*
  - Authors: *Vedant Mundheda, Zhouchonghao Wu, Jeff Schneider*

Results Table

Title	Points (% out of max)	Award
Teacher-guided Off-road Autonomous Driving	91%	Best **Behavior**
A Spatiotemporal Approach to Tri-Perspective Representation for 3D Semantic Occupancy Prediction	85%	Best **Perception**
Paper 3	82%	-
...		
Paper N	61%	-

Gatik

Best Perception: Foundation

- **Baseline work: Tri-Perspective View (TPV[1])**
  - Voxel: Expensive
  - BEV: Not Expressive Enough (lost z)
  - TPV: A compromise!

- Moreover, let's avoid LiDAR

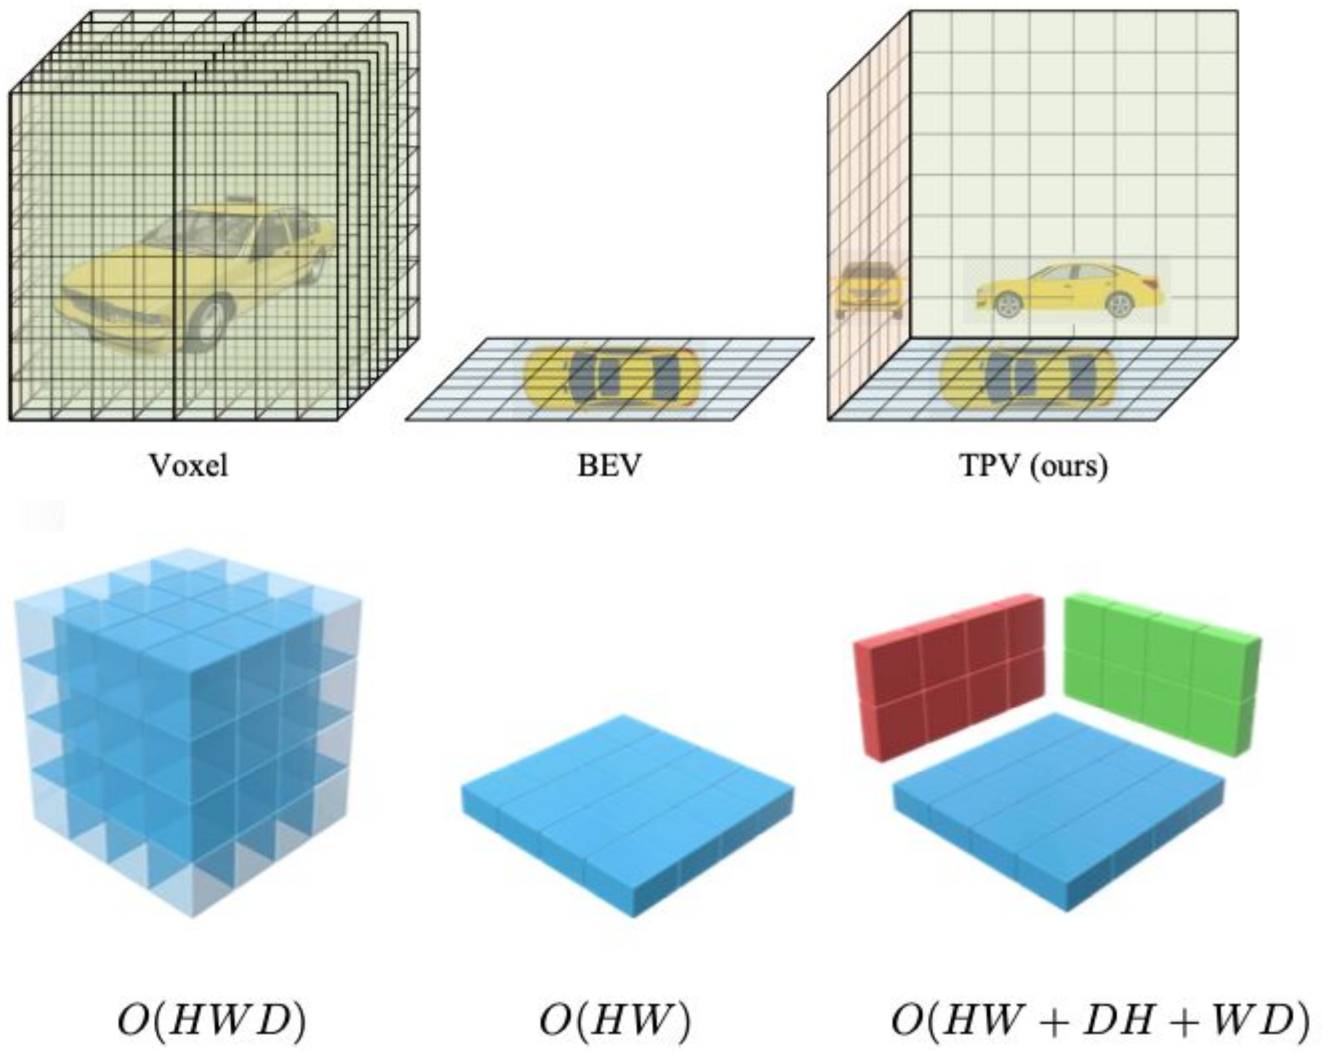

[1] Huang, Yuanhui, et al. "Tri-perspective view for vision-based 3d semantic occupancy prediction" 2023

Best Perception: Adding Time – TPV → S2TPV

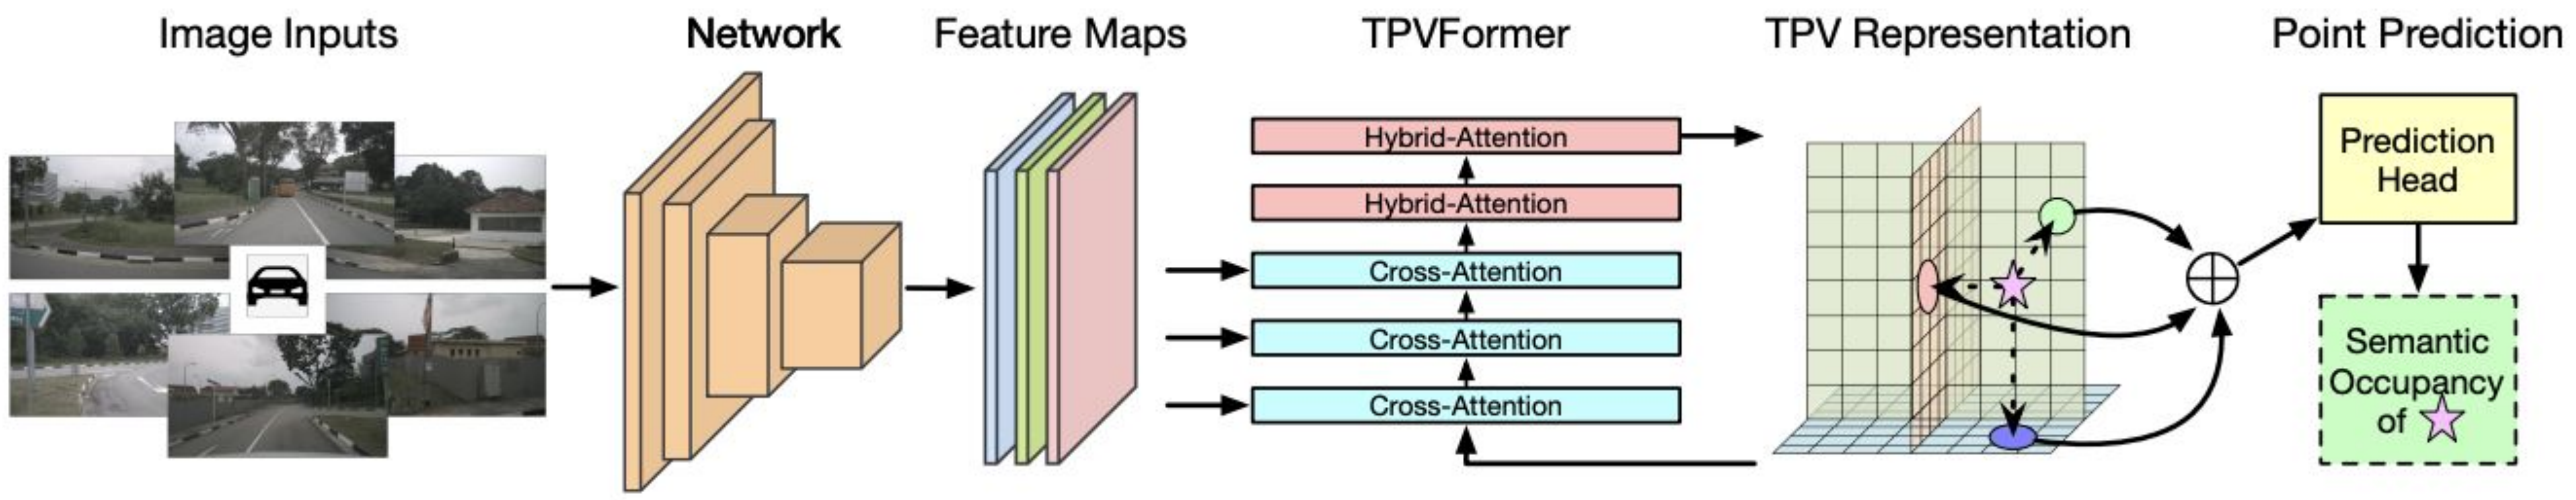

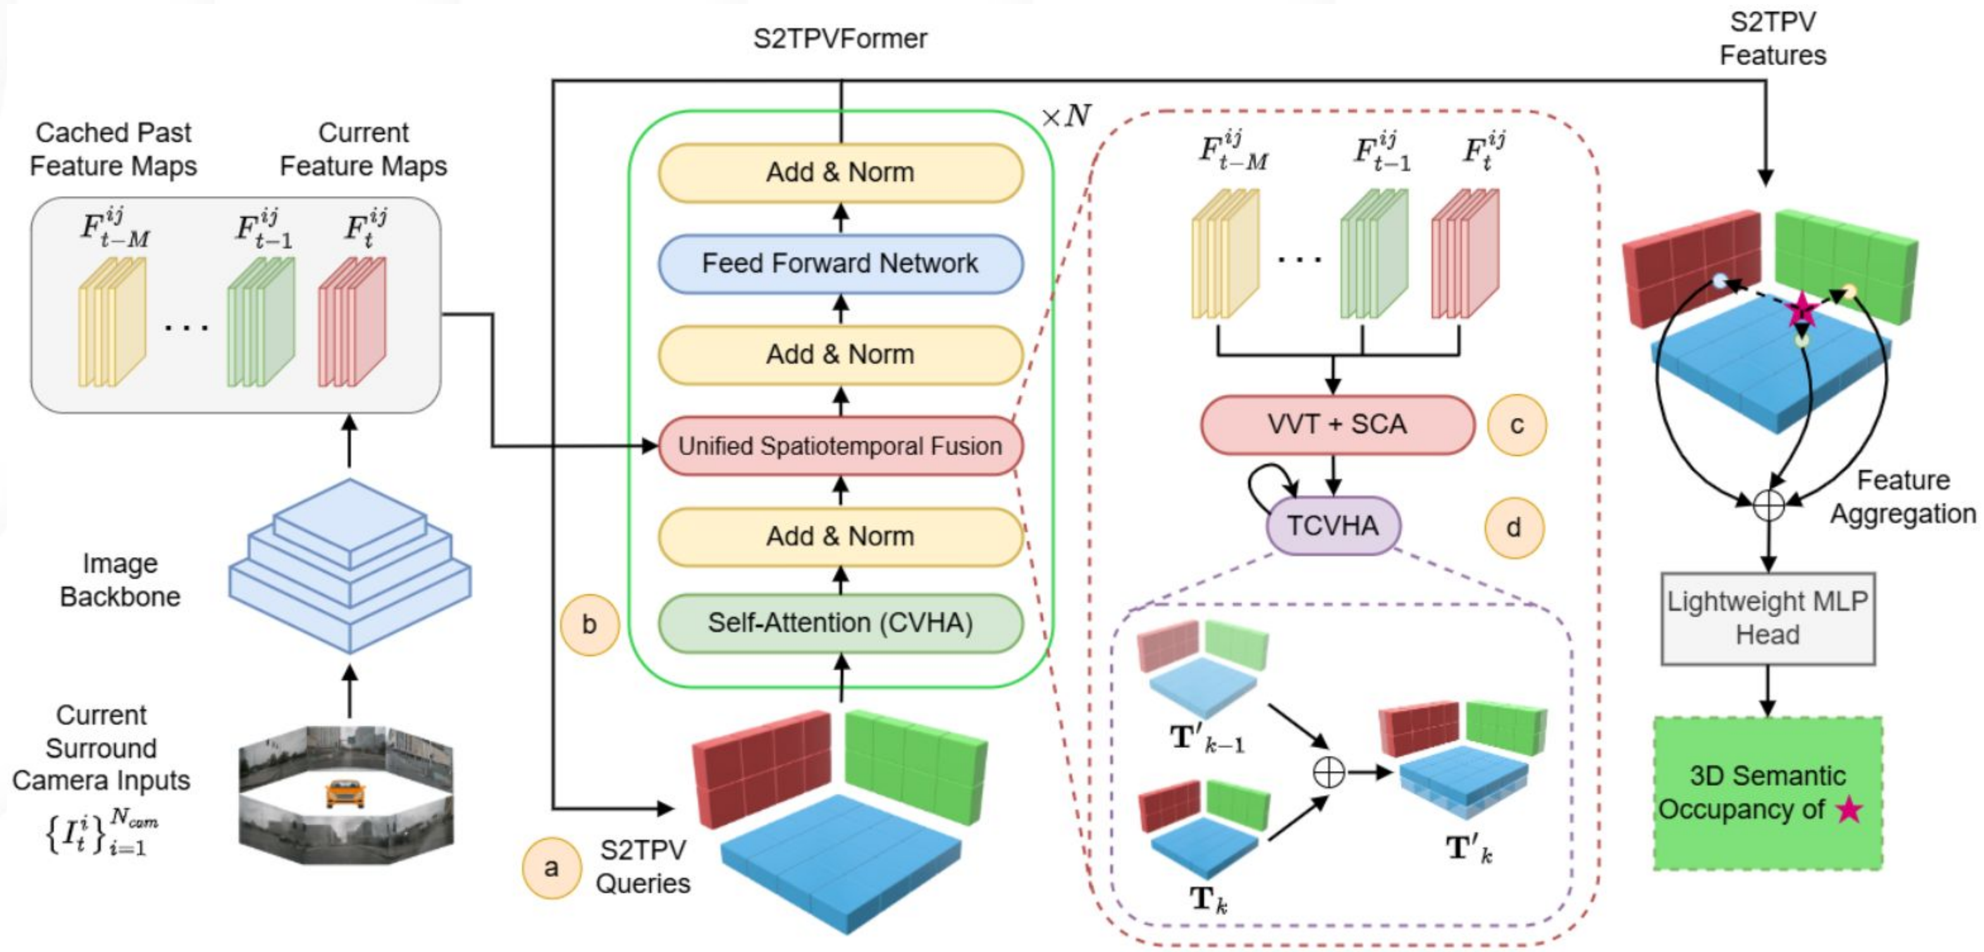

Best Perception: Opportunities

- **TPVFormer vs S2TPVFormer**
  - Not uniformly better
- **Latency**
  - No timings and its comparison to other TPV / BEV / Voxel solutions

Method	Input Modality	mIoU (%)
MINet	LiDAR	56.3
LidarMultiNet	LiDAR	81.4
UniVision	LiDAR	72.3
PanoOcc	LiDAR	71.4
OccFormer	LiDAR	70.8
TPVFormer-Small†	Camera	59.2
TPVFormer-Base†	Camera	69.4
S2TPVFormer (Base)	Camera	60.4

Method	Input Modality	mIoU
BEVFormer	Camera	56.2
TPVFormer-Base†	Camera	68.9
TPVFormer-Small†	Camera	59.3
S2TPVFormer (base)	Camera	61.6

Ablation	mIoU (%)
TPVFormer-Small*	**44.4**
S2TPVFormer (Small)	43.4
TPVFormer	52.0
S2TPVFormer (Base)	**55.0**

Best Behavior: Foundation

- **Baselines:**
  - Model Predictive Path Integral (MPPI[1])
    - Accurate, but computationally expensive
  - RL / Proximal Policy Optimization (PPO[2])
    - Exploration, fast
- **Approach**:
  - Hierarchy: low-frequency MPPI → RL
    - RL agent is distilled from the high-frequency MPPI teacher: Teacher Action Distillation with Policy Optimization (**TADPO**)

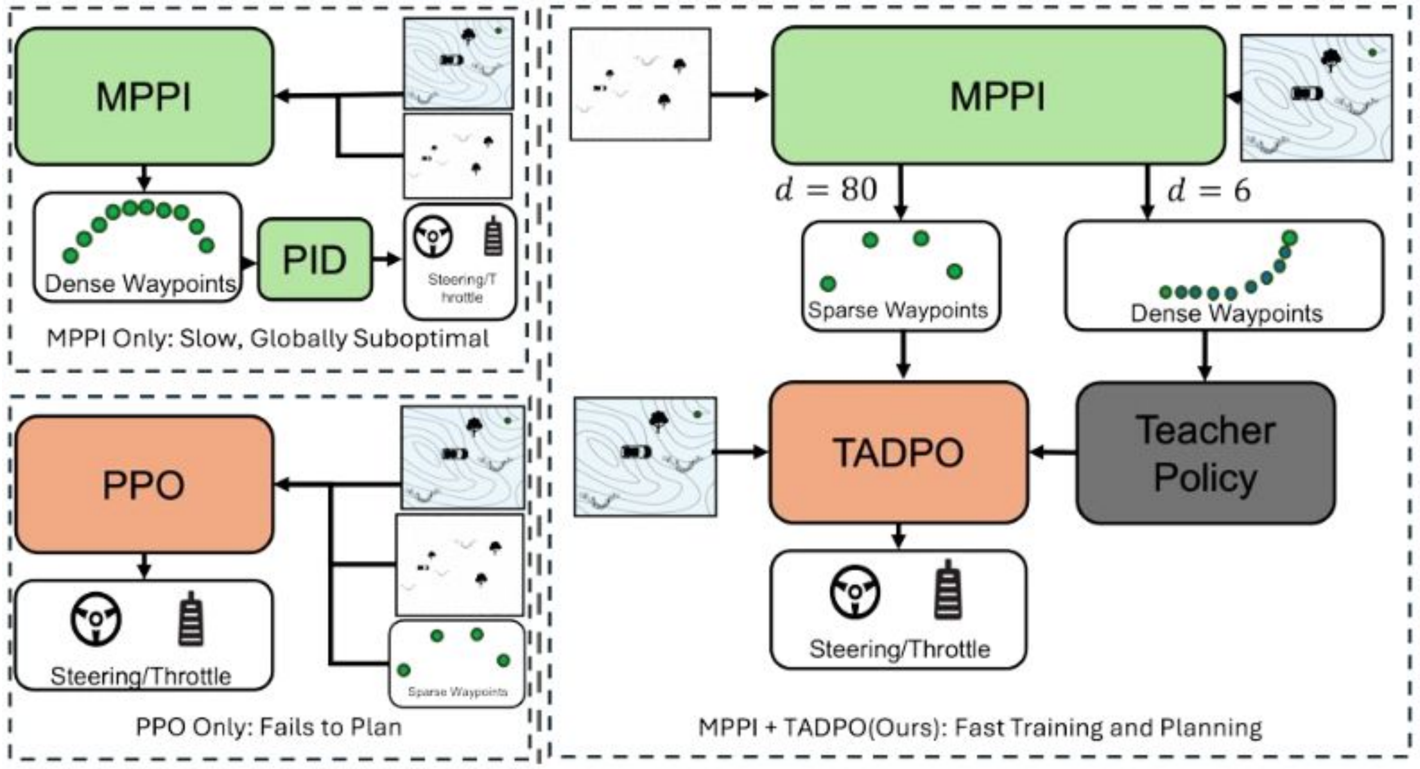

Gatik

[1] Williams, Grady, et al. "Model predictive path integral control using covariance variable importance sampling." 2015.
[2] Schulman, John, et al. "Proximal policy optimization algorithms." 2017.

Best Behavior: TADPO

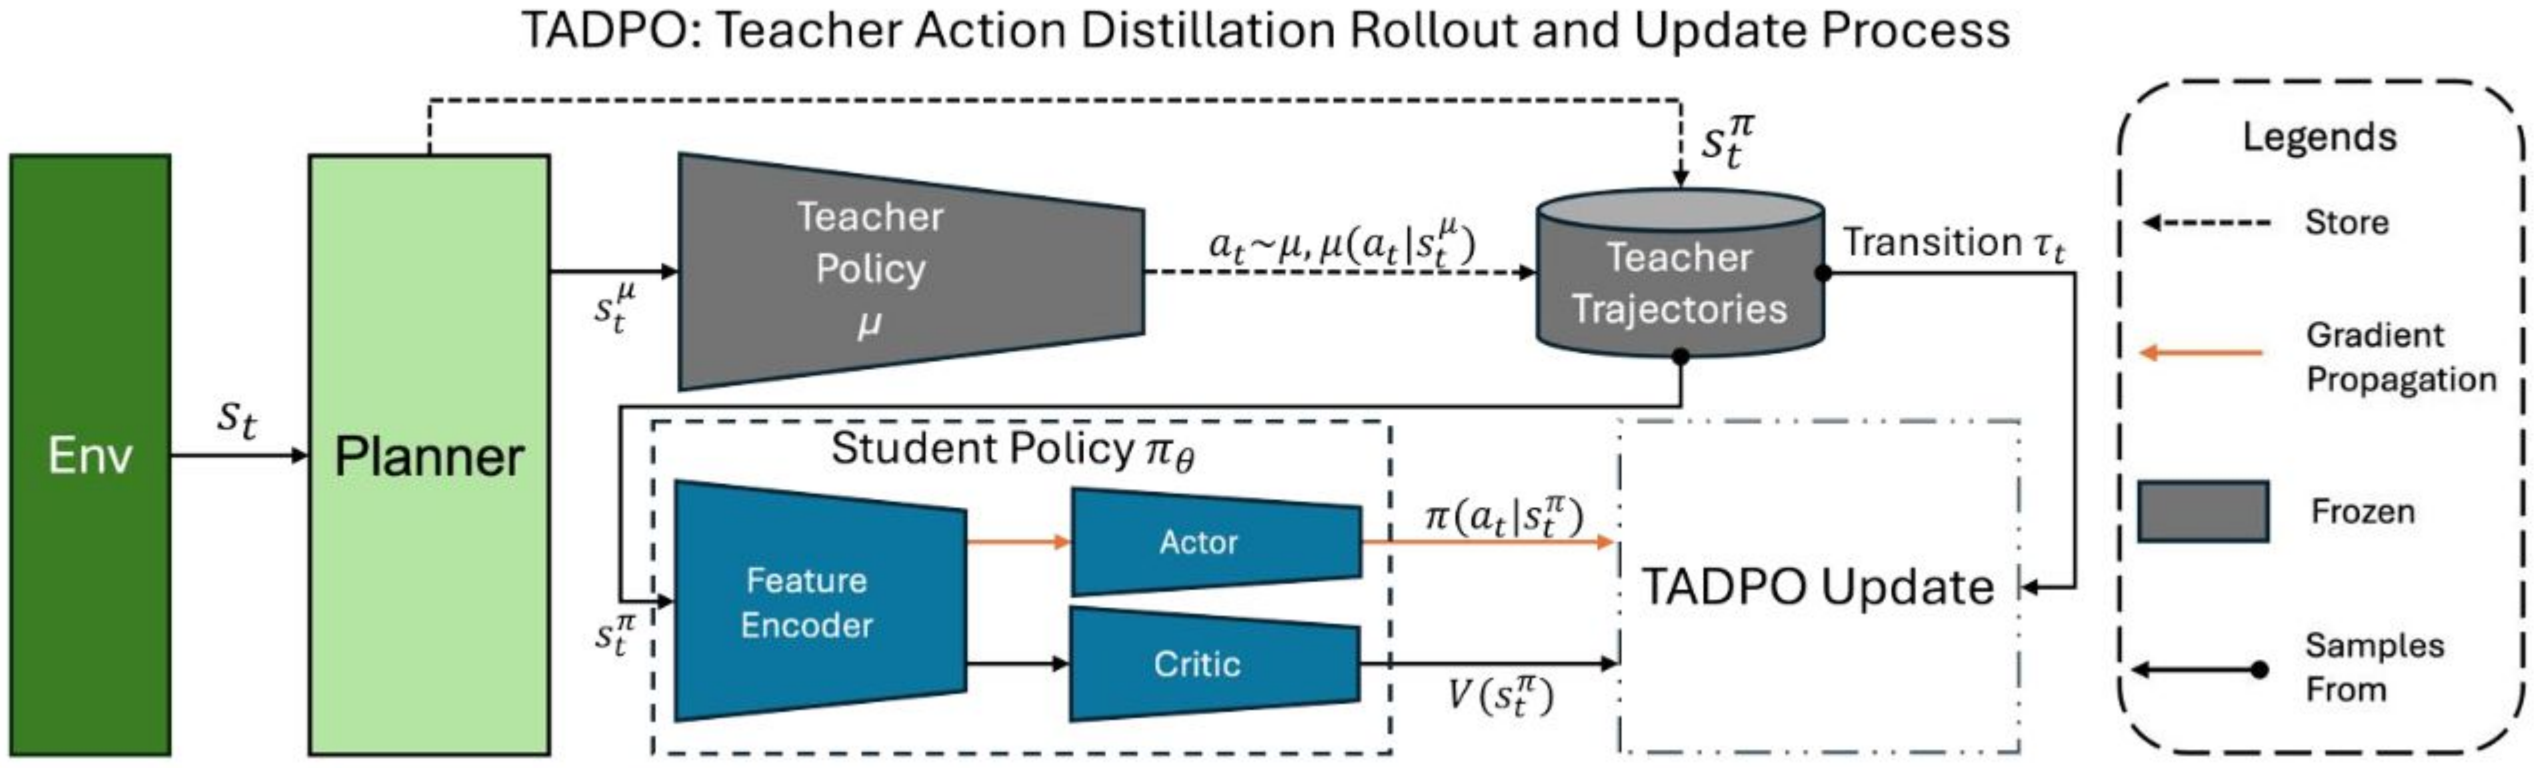

$$L^{\text{CLIP}}(\theta) = \mathbb{E}_t\left[\min\left(r_t(\theta)\hat{A}_t, \text{clip}(r_t(\theta), 1-\epsilon, 1+\epsilon)\hat{A}_t\right)\right]$$

$$L^{\text{VF}}(\theta) = \mathbb{E}_t\left[(V_{\pi_{\theta_{\text{old}}}}(s_t) - R_t)^2\right]$$

$$L^{\text{entropy}}(\theta) = \mathbb{E}_t\left[-H[\pi_\theta(\cdot|s_t)]\right]$$

$$L^{\text{PPO}}(\theta) = L^{\text{CLIP}}(\theta) - c_1 L^{\text{VF}}(\theta) + c_2 L^{\text{entropy}}(\theta)$$

PPO

→

$$L^{\text{TAD}}(\theta) = L^{\mu}(\theta) + c_2 L^{\text{entropy}}(\theta)$$

$$\rho_t(\theta) = \frac{\pi_\theta(a_t|s_t^\pi)}{\mu(a_t|s_t^\mu)}$$

$$\hat{\Delta}_t = R(a_t, s_t) - V_{\pi\theta_{\text{old}}}(s_t^\pi)$$

$$L^{\mu}(\theta) = \mathbb{E}_{a_t\sim\mu}\left[\max\left(0, \min(\rho_t(\theta), 1+\epsilon_\mu)\hat{\Delta}_t\right)\right]$$

TADPO

Best Behavior: Opportunities

- **Usecase**
  - Interesting to see TADPO usage for usual onroad driving (CARLA[1])
- **Comparison**
  - Hierarchical approach:
    - Instead of Optimization → RL, RL → IL / other generator (CIMRL[2])
  - Balance:
    - RL exploration vs IL speed / determinism (BC-SAC[3])

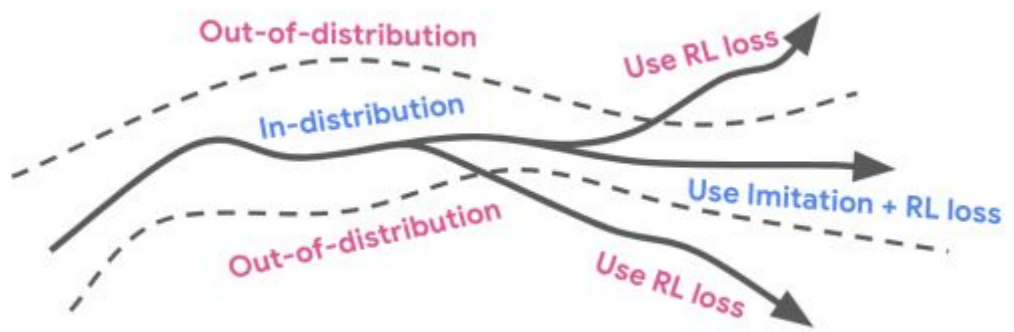

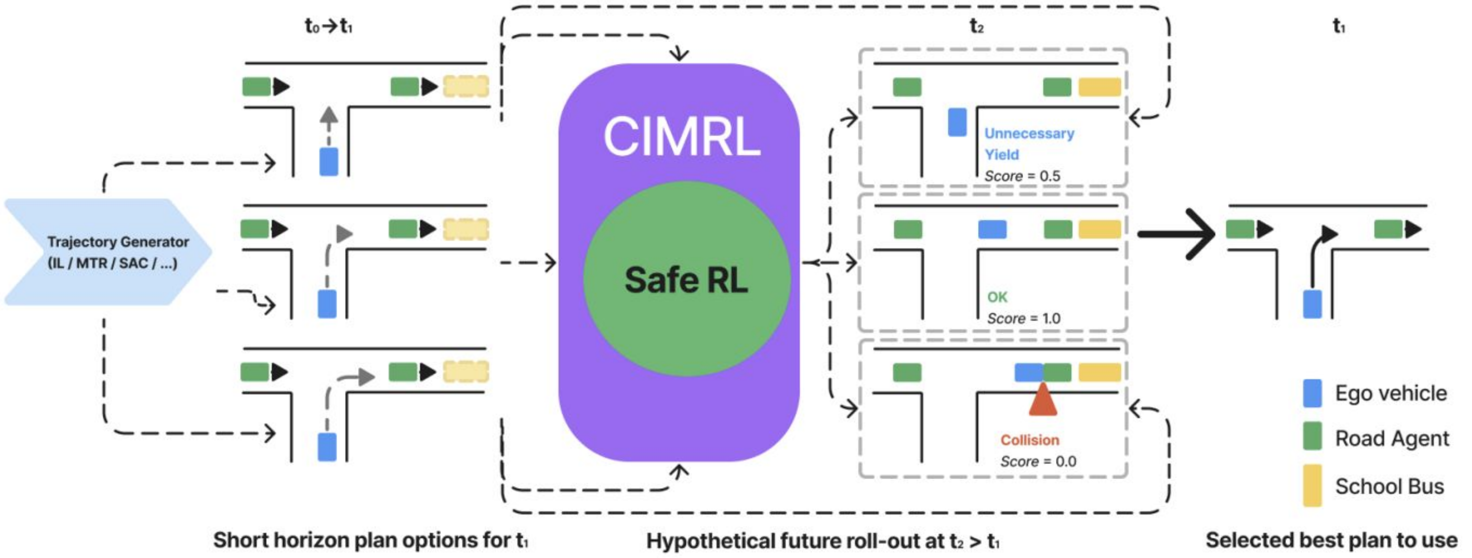

[1] Dosovitskiy, Alexey, et al. "CARLA: An open urban driving simulator." 2017.
[2] Booher, Jonathan, et al. "CIMRL: Combining IMitation and Reinforcement Learning for Safe Autonomous Driving." 2024.
[3] Lu, Yiren, et al. "Imitation is not enough: Robustifying imitation with reinforcement learning for challenging driving scenarios." 2023.

Conclusion

Submitted papers

are well-aligned with the **ML4AD** Workshop topics

Perception papers

concentrate on **3D**

Behavior papers

focus on **Simulation**

Q&A

Your Questions, Our Expertise

Gatik