LabelFormer: Object Trajectory Refinement for Offboard Perception from LiDAR Point Clouds

Anqi Joyce Yang, Sergio Casas, Nikita Dvornik, Sean Segal,
Yuwen Xiong, Jordan Sir Kwang Hu, Carter Fang, Raquel Urtasun

waabi · University of Toronto

What is Offboard Perception?

- **Motivation:** Modern self-driving systems require a large set of annotated data, but human labelling is slow and costly
- **Task:** Automatically label object trajectories from LiDAR data
- **Setting:** Access to a limited set of human annotations, access to future observations, no real-time constraints
- **Goal:** Accurate bounding boxes + computationally cheap
- **Application:** Generate large-scale auto-labelled dataset for training downstream onboard perception models

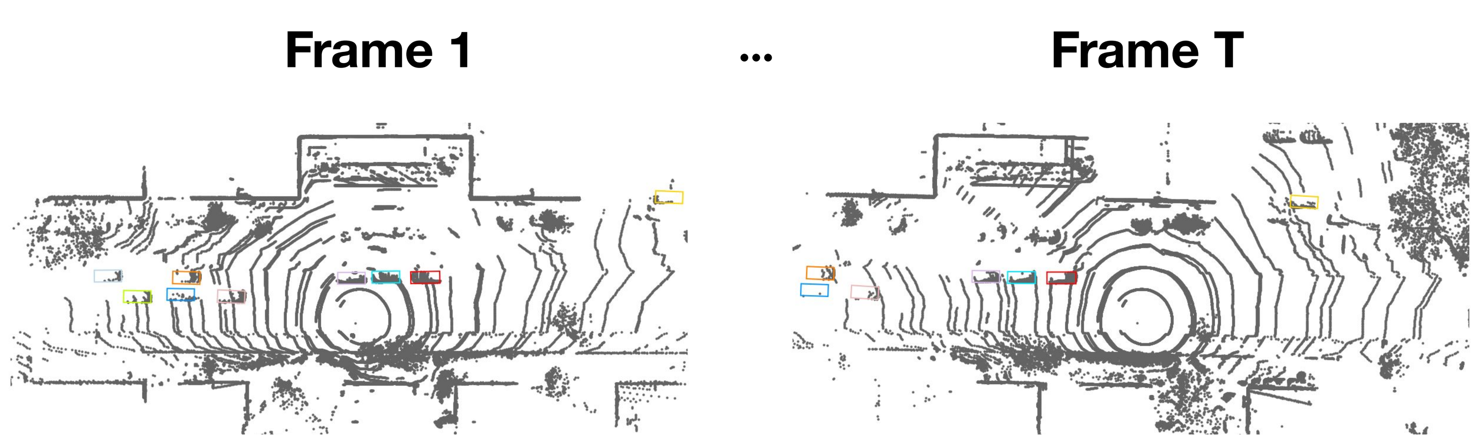

Intuition behind LabelFormer

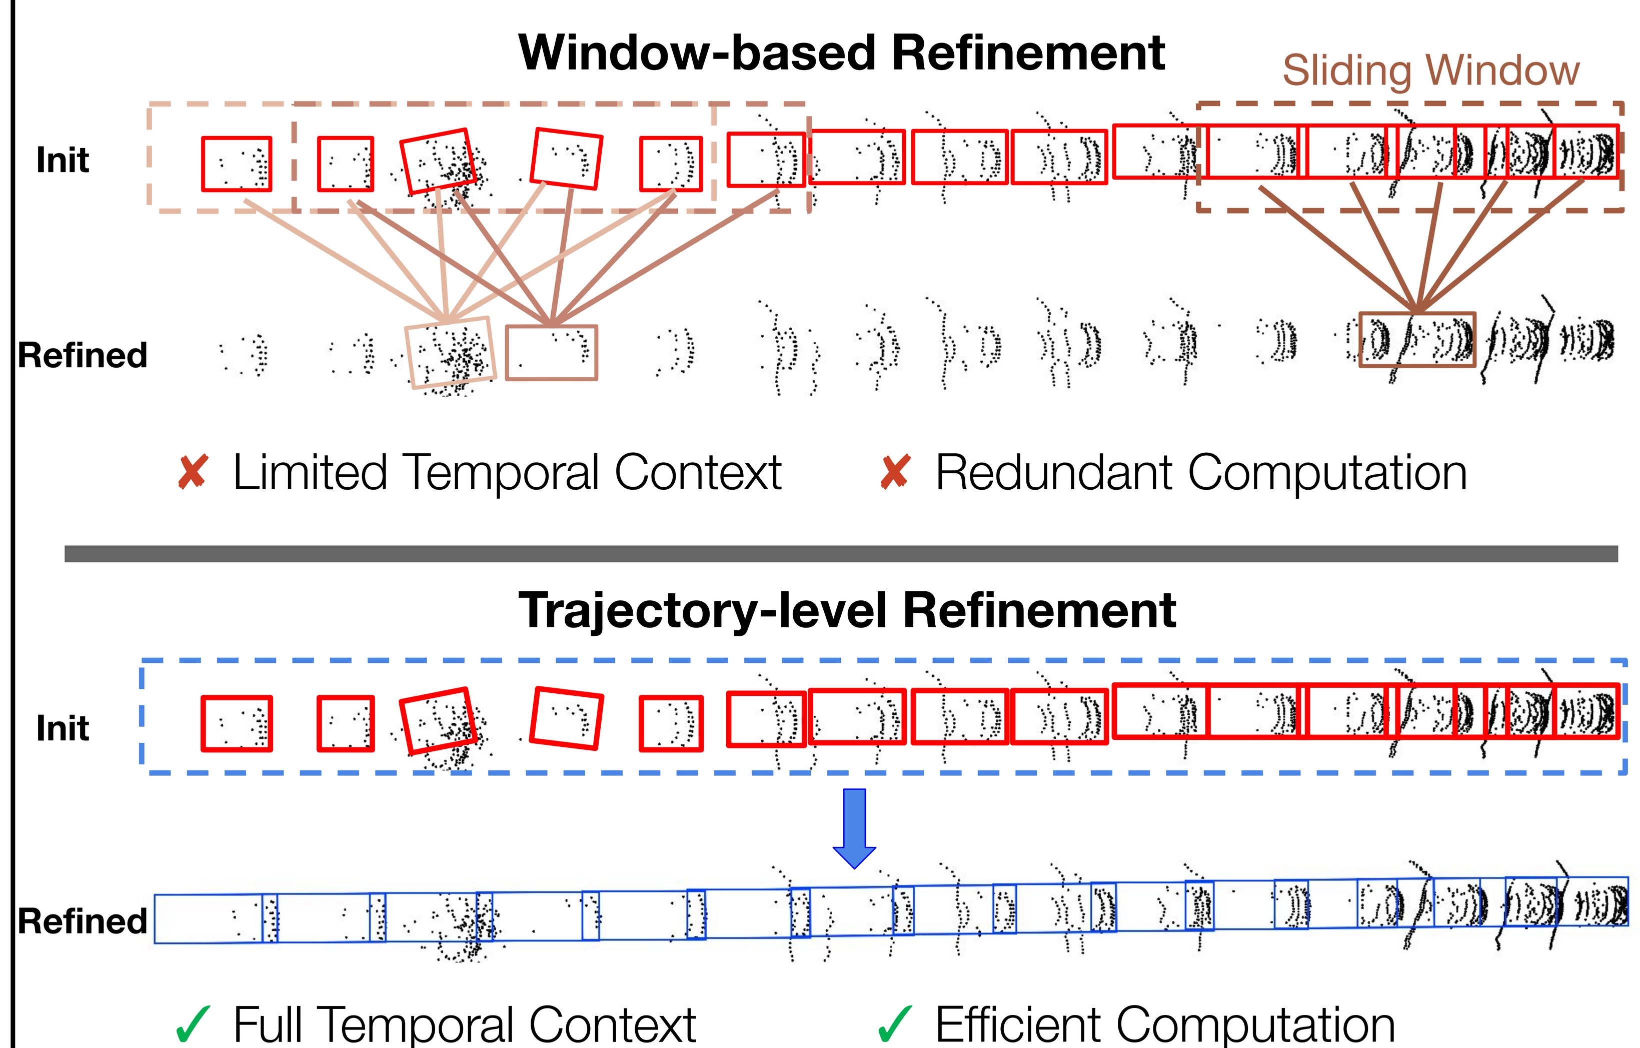

Qualitative Results

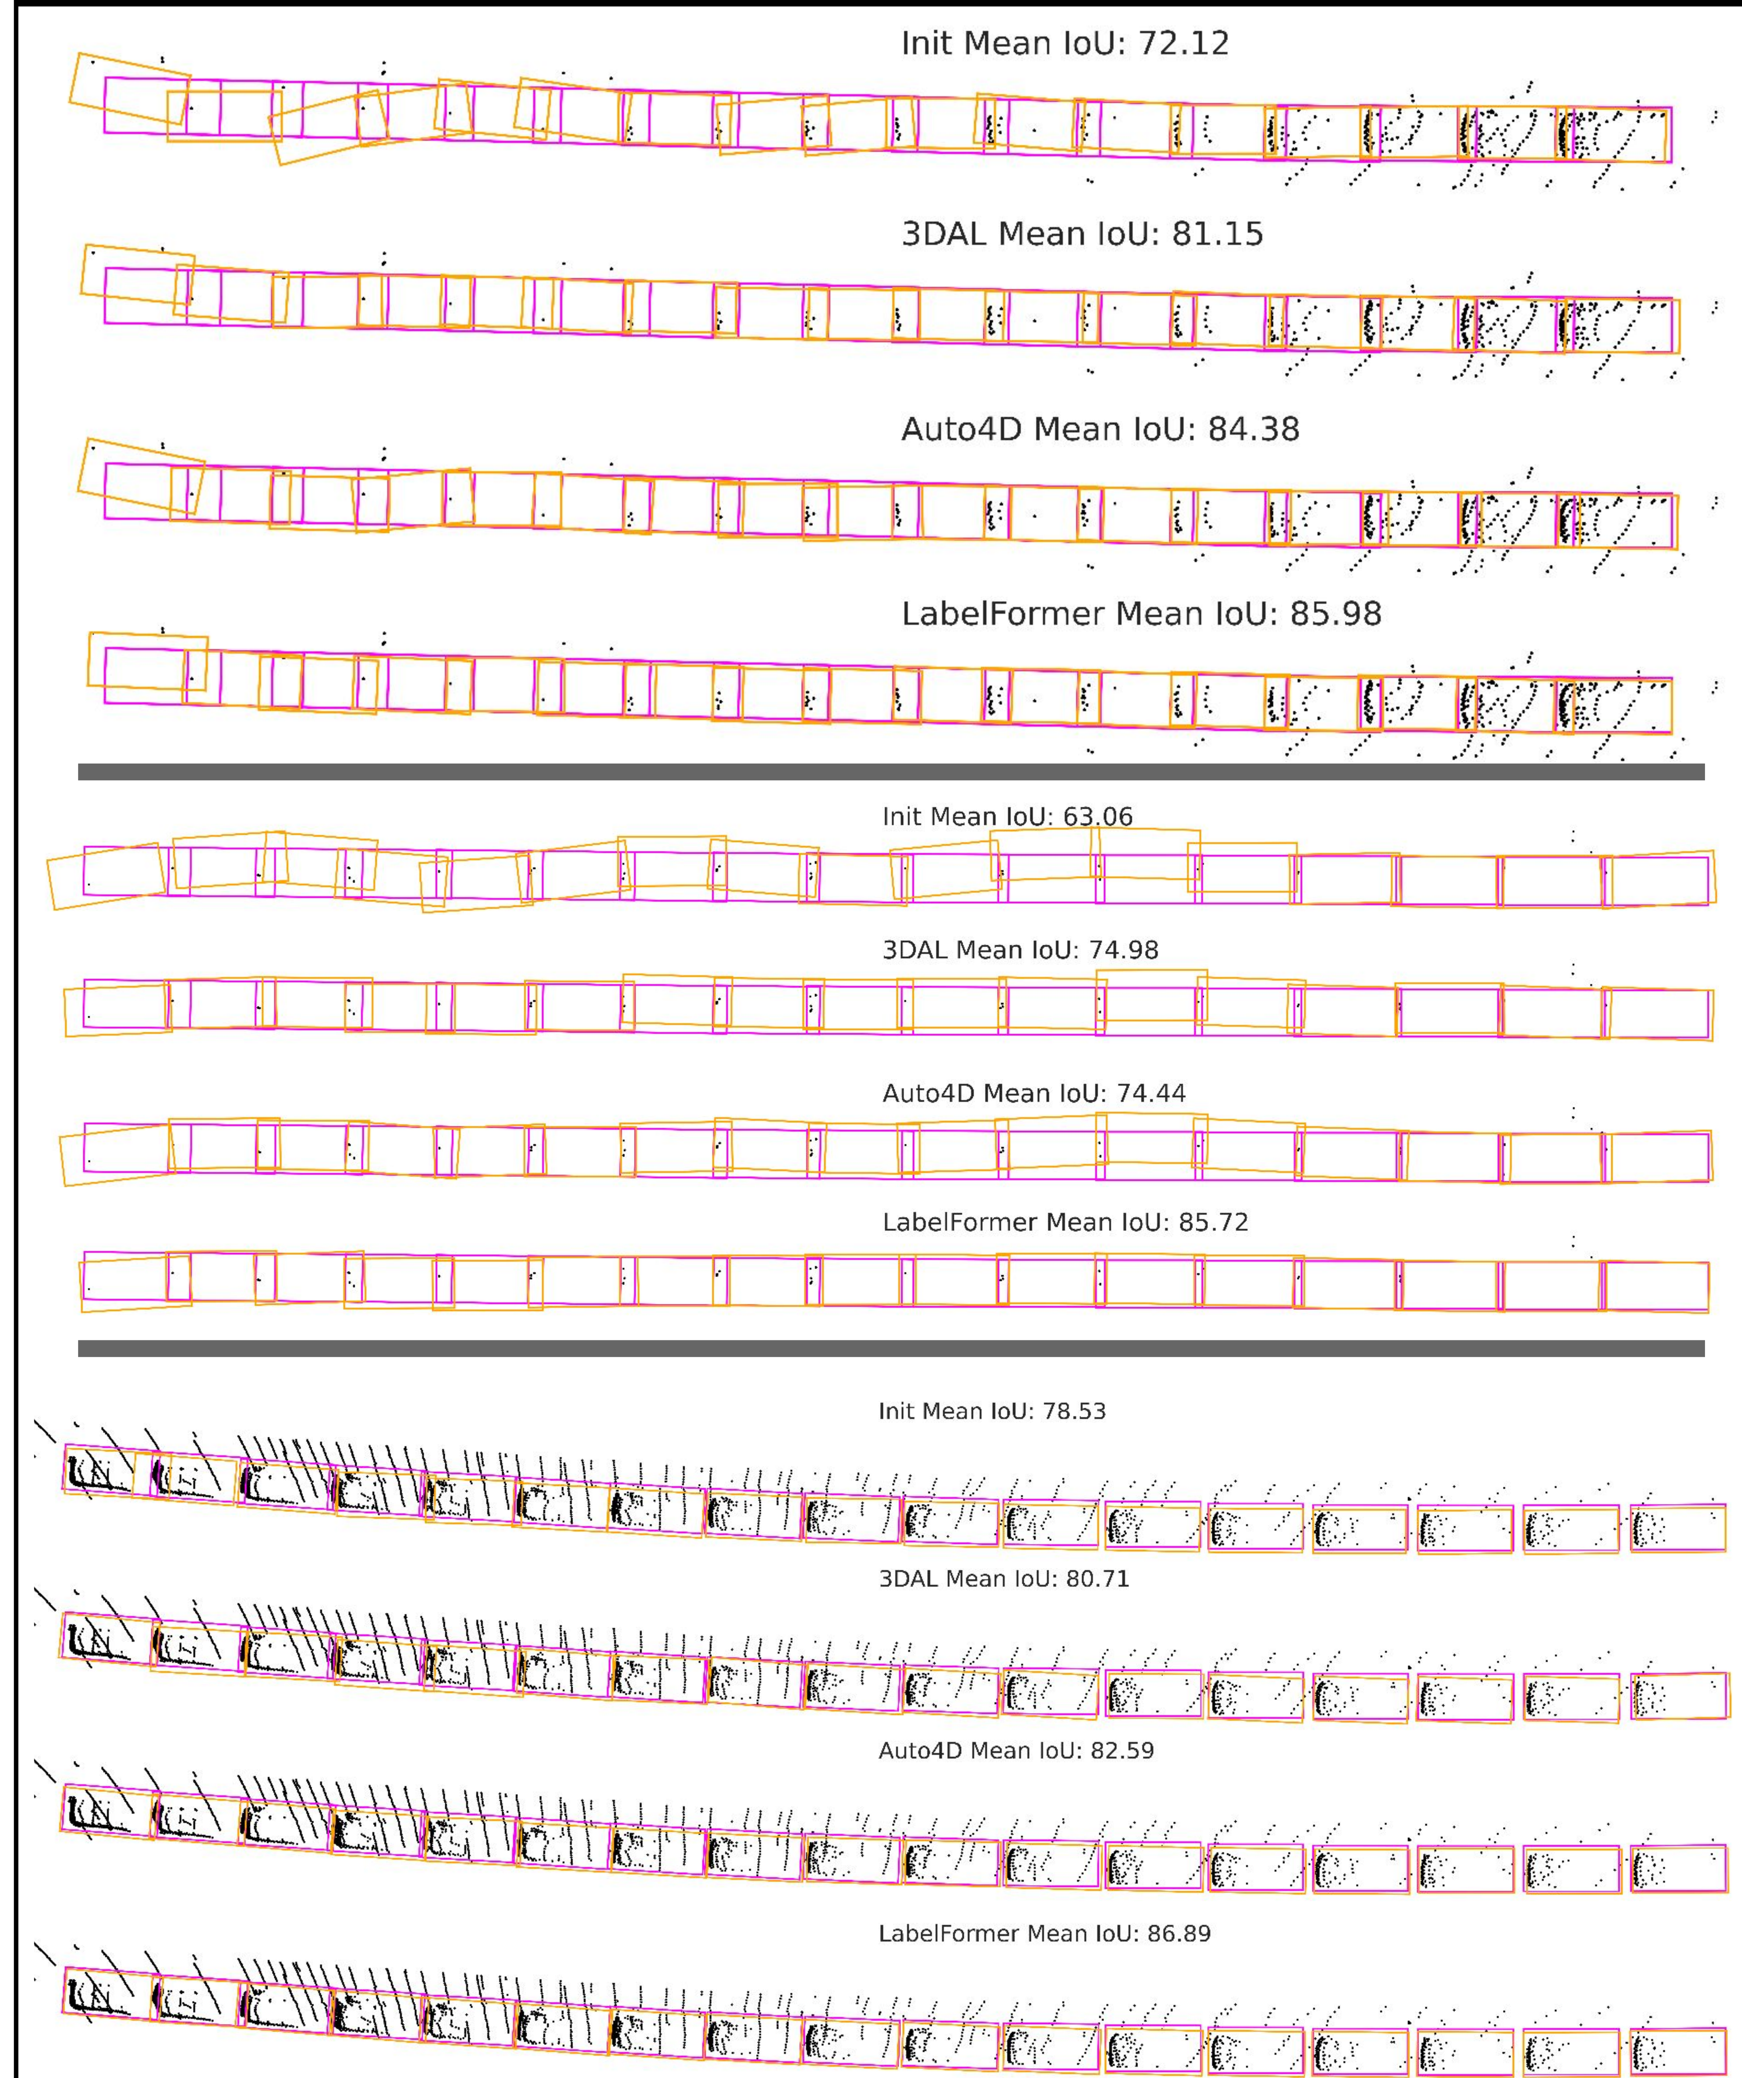

Two-Stage Auto-labelling Paradigm

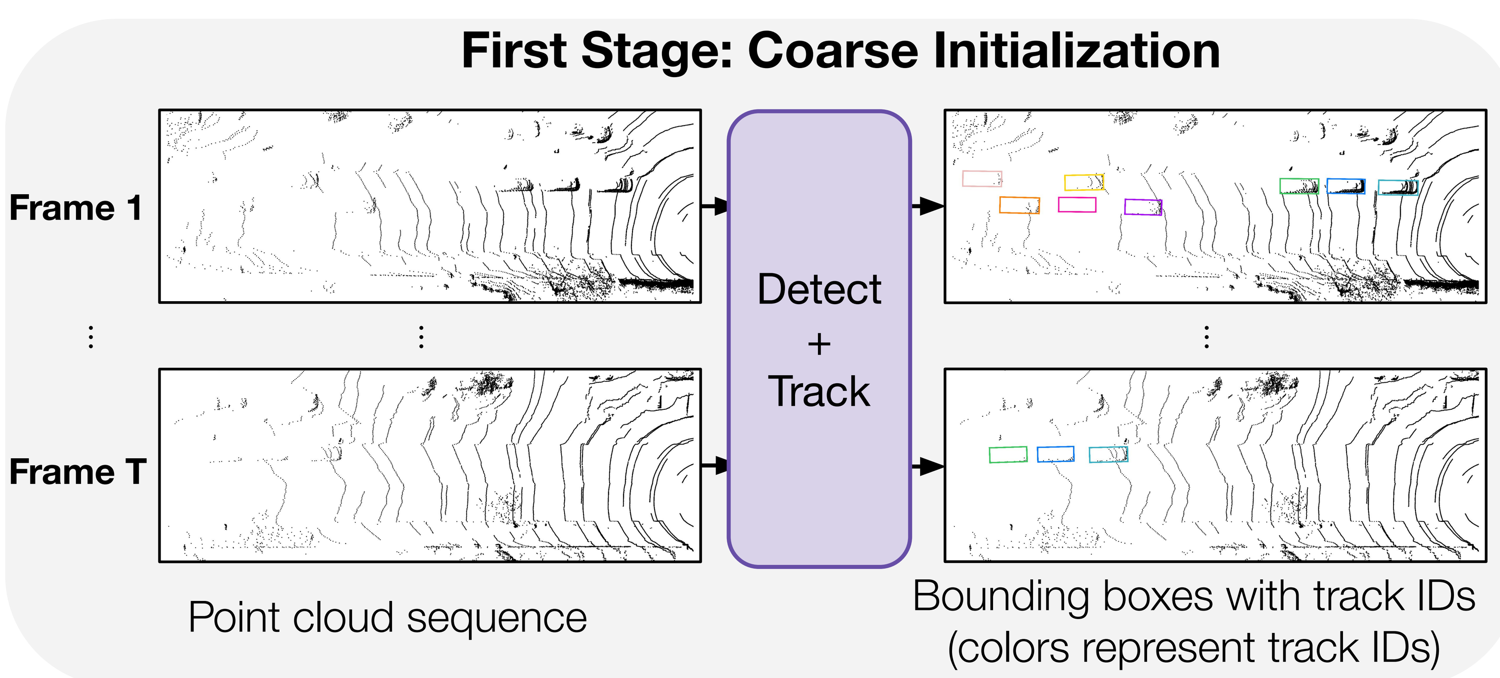

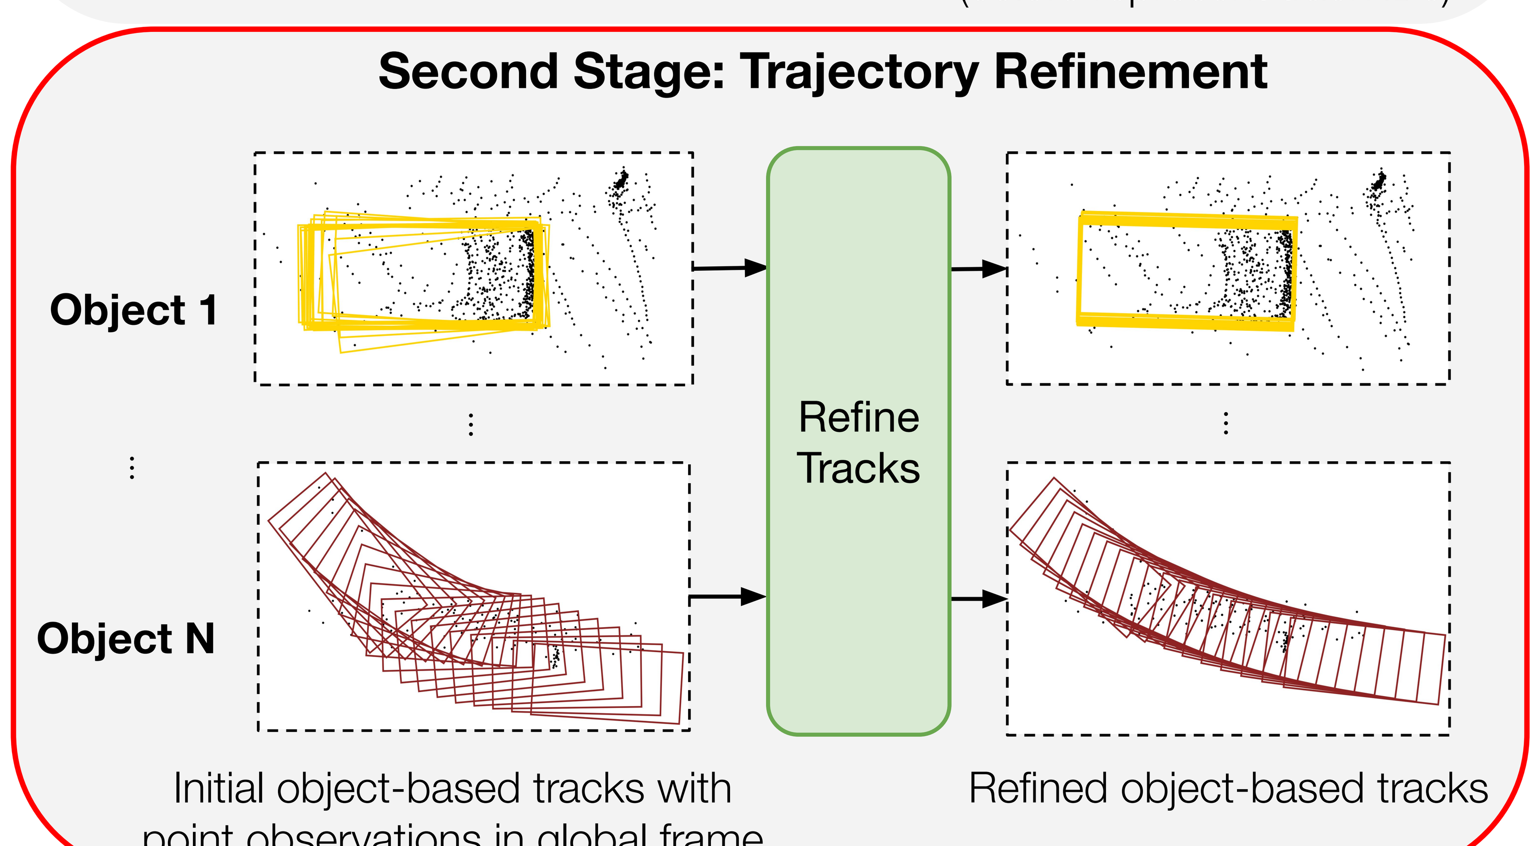

In this work we tackle the second-stage object trajectory refinement.

LabelFormer Architecture

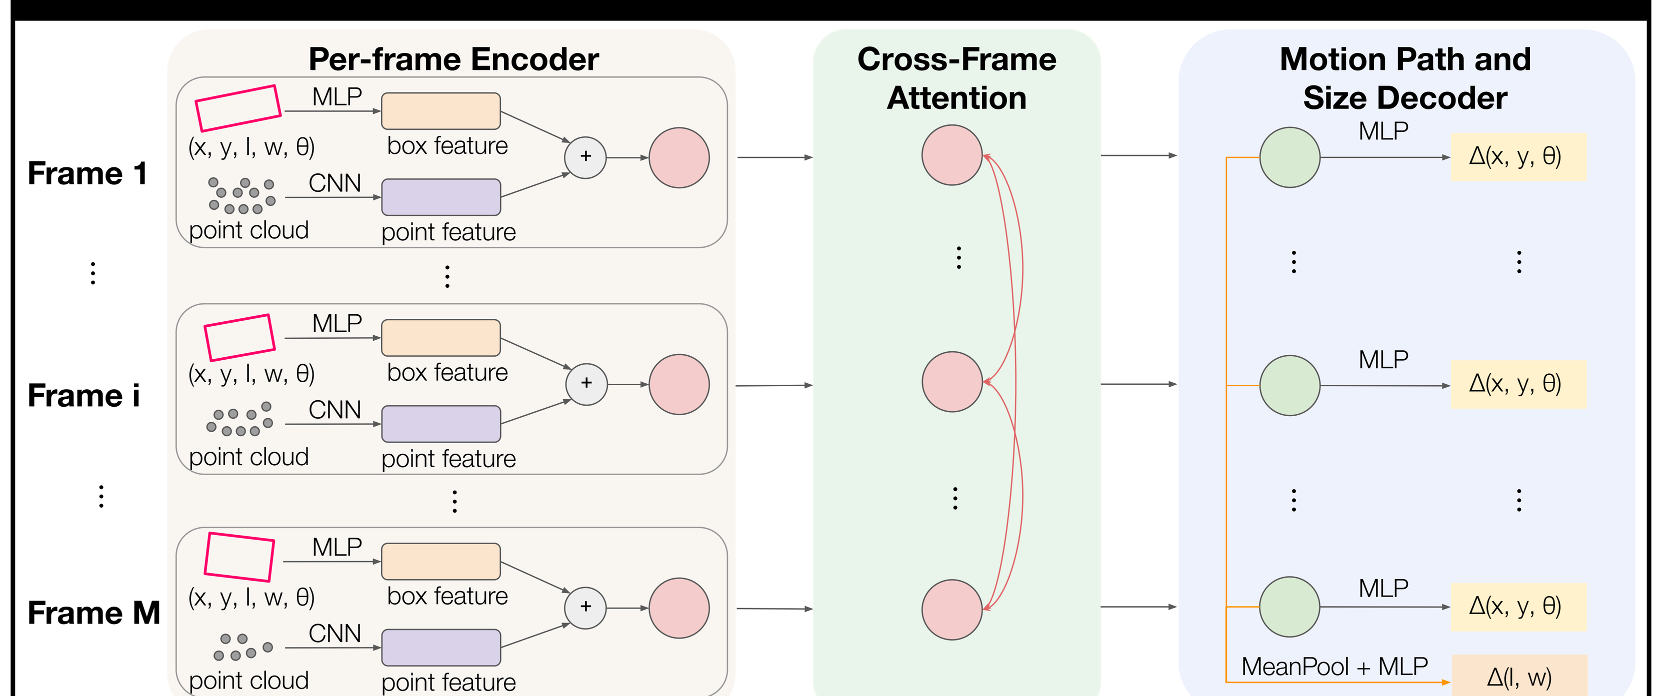

- **Simple:** single network
- **Effective:** temporal reasoning with full temporal context
- **Efficient:** no redundant computation

Comparison with SoTA

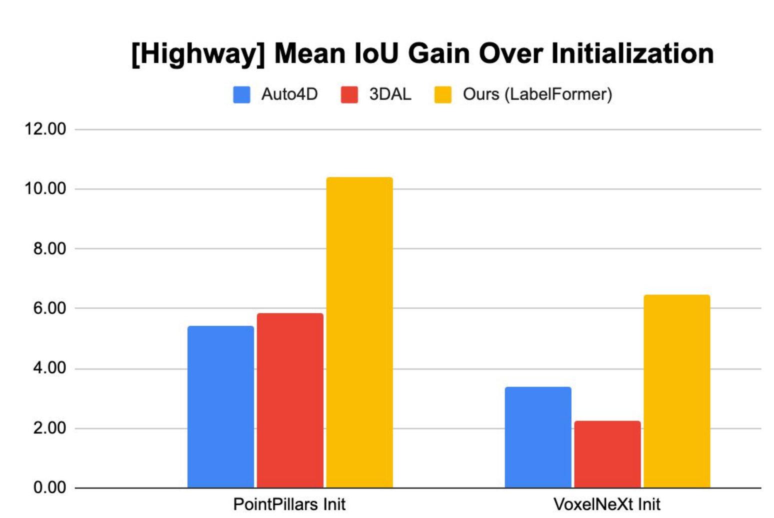

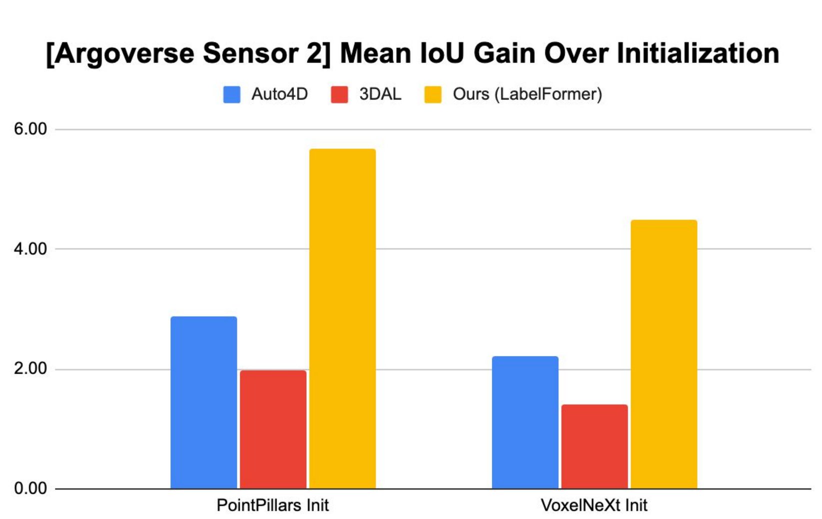

Downstream Evaluation

We apply different auto-labellers to augment the Highway dataset. With a larger annotated dataset, we train a downstream onboard object detector.

Auto-Label	Mean AP	AP@0.5	AP@0.7	AP@0.8
N/A	82.98	91.62	79.17	55.97
Init	83.63	92.67	79.51	55.30
Auto4D	83.42	92.71	79.32	55.07
3DAL	83.64	92.66	79.76	55.79
Ours	**84.81**	**92.91**	**80.91**	**59.00**