

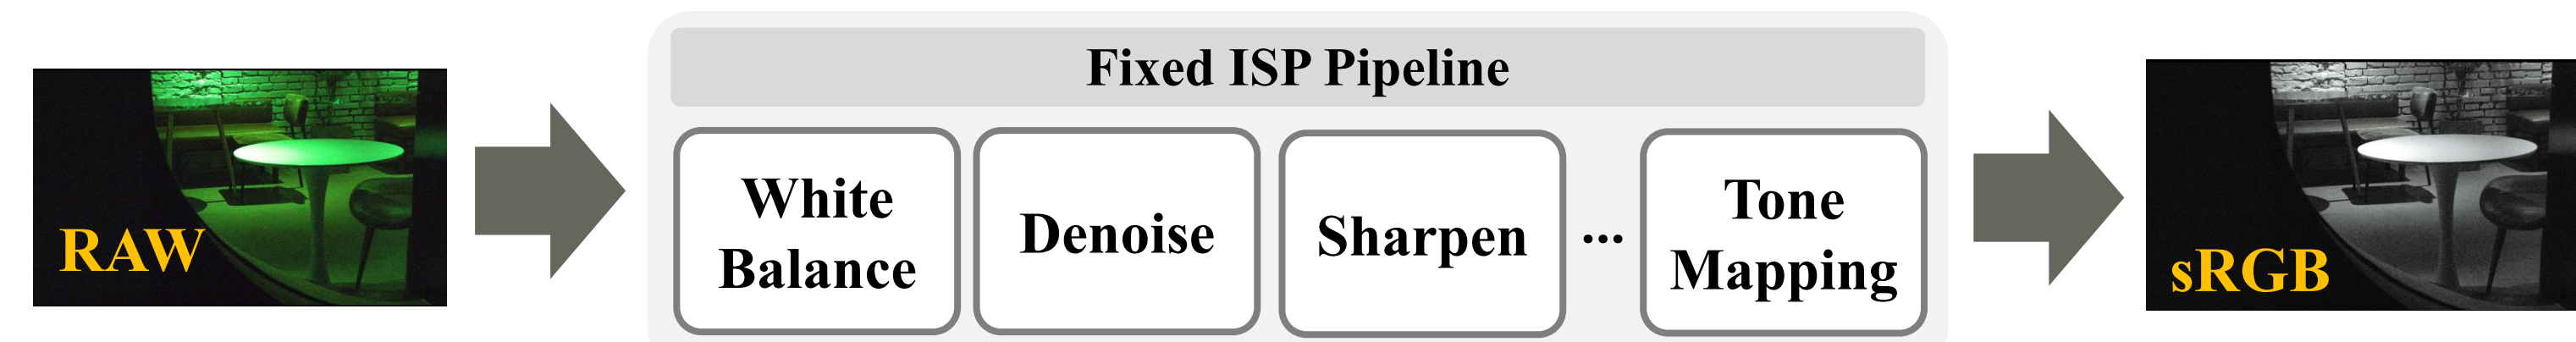


TL;DR: Sequence-level RL makes task-aware ISP more stable, faster, and lighter while improving downstream performance.

Camera ISP

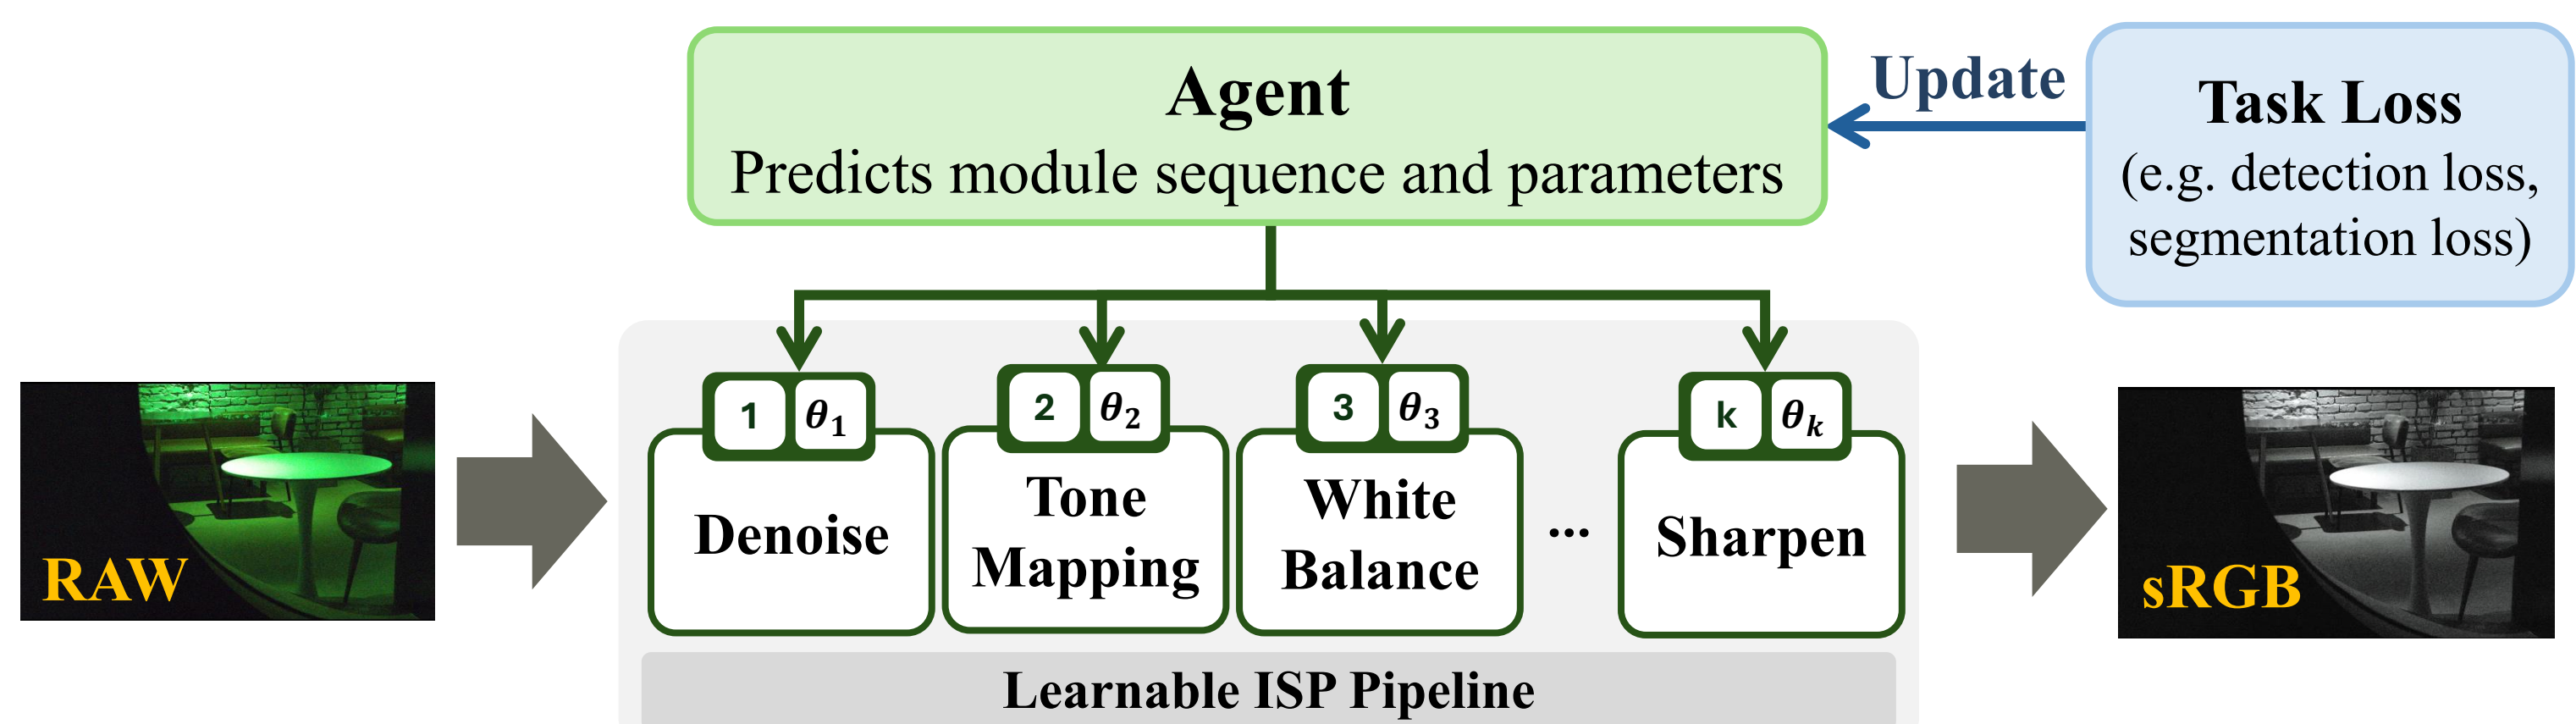
Conventional Camera ISP (Image Signal Processor)

- Convert a RAW image into a visually pleasing sRGB image
- Uses a fixed, human-designed sequence of ISP modules



Task-aware Camera ISP

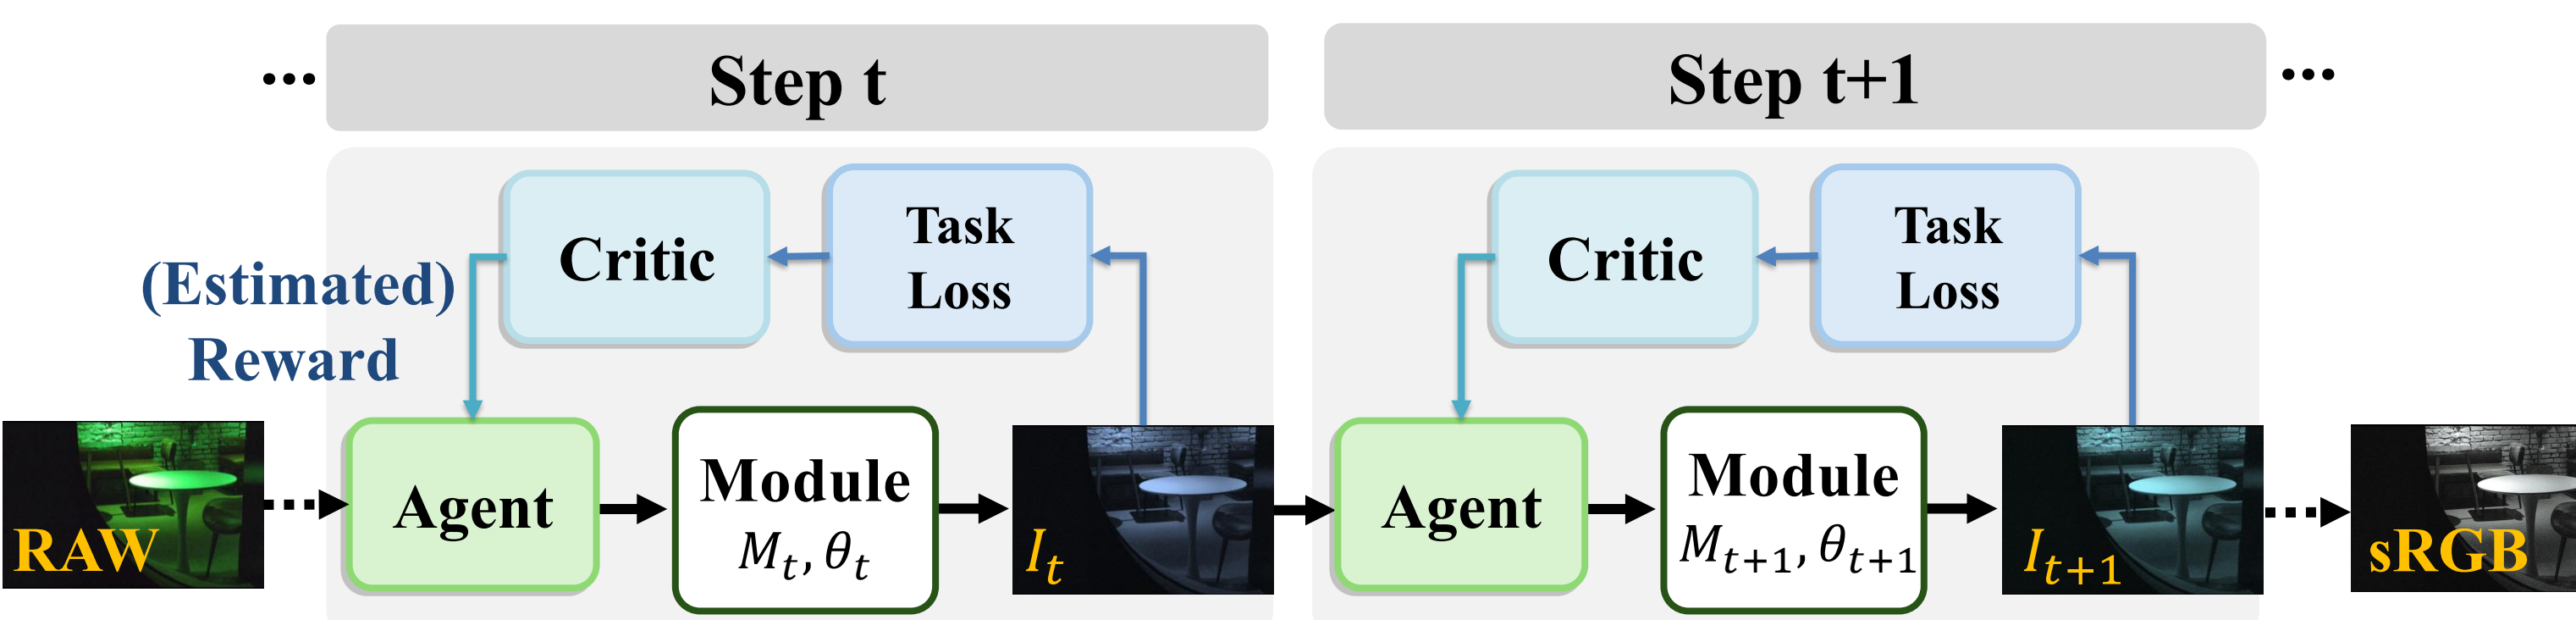
- Optimize the ISP pipeline for downstream tasks (e.g., object detection, instance segmentation)
- Learns both module order and parameters using task loss



Limitations of Previous Methods

Stepwise Reinforcement Learning

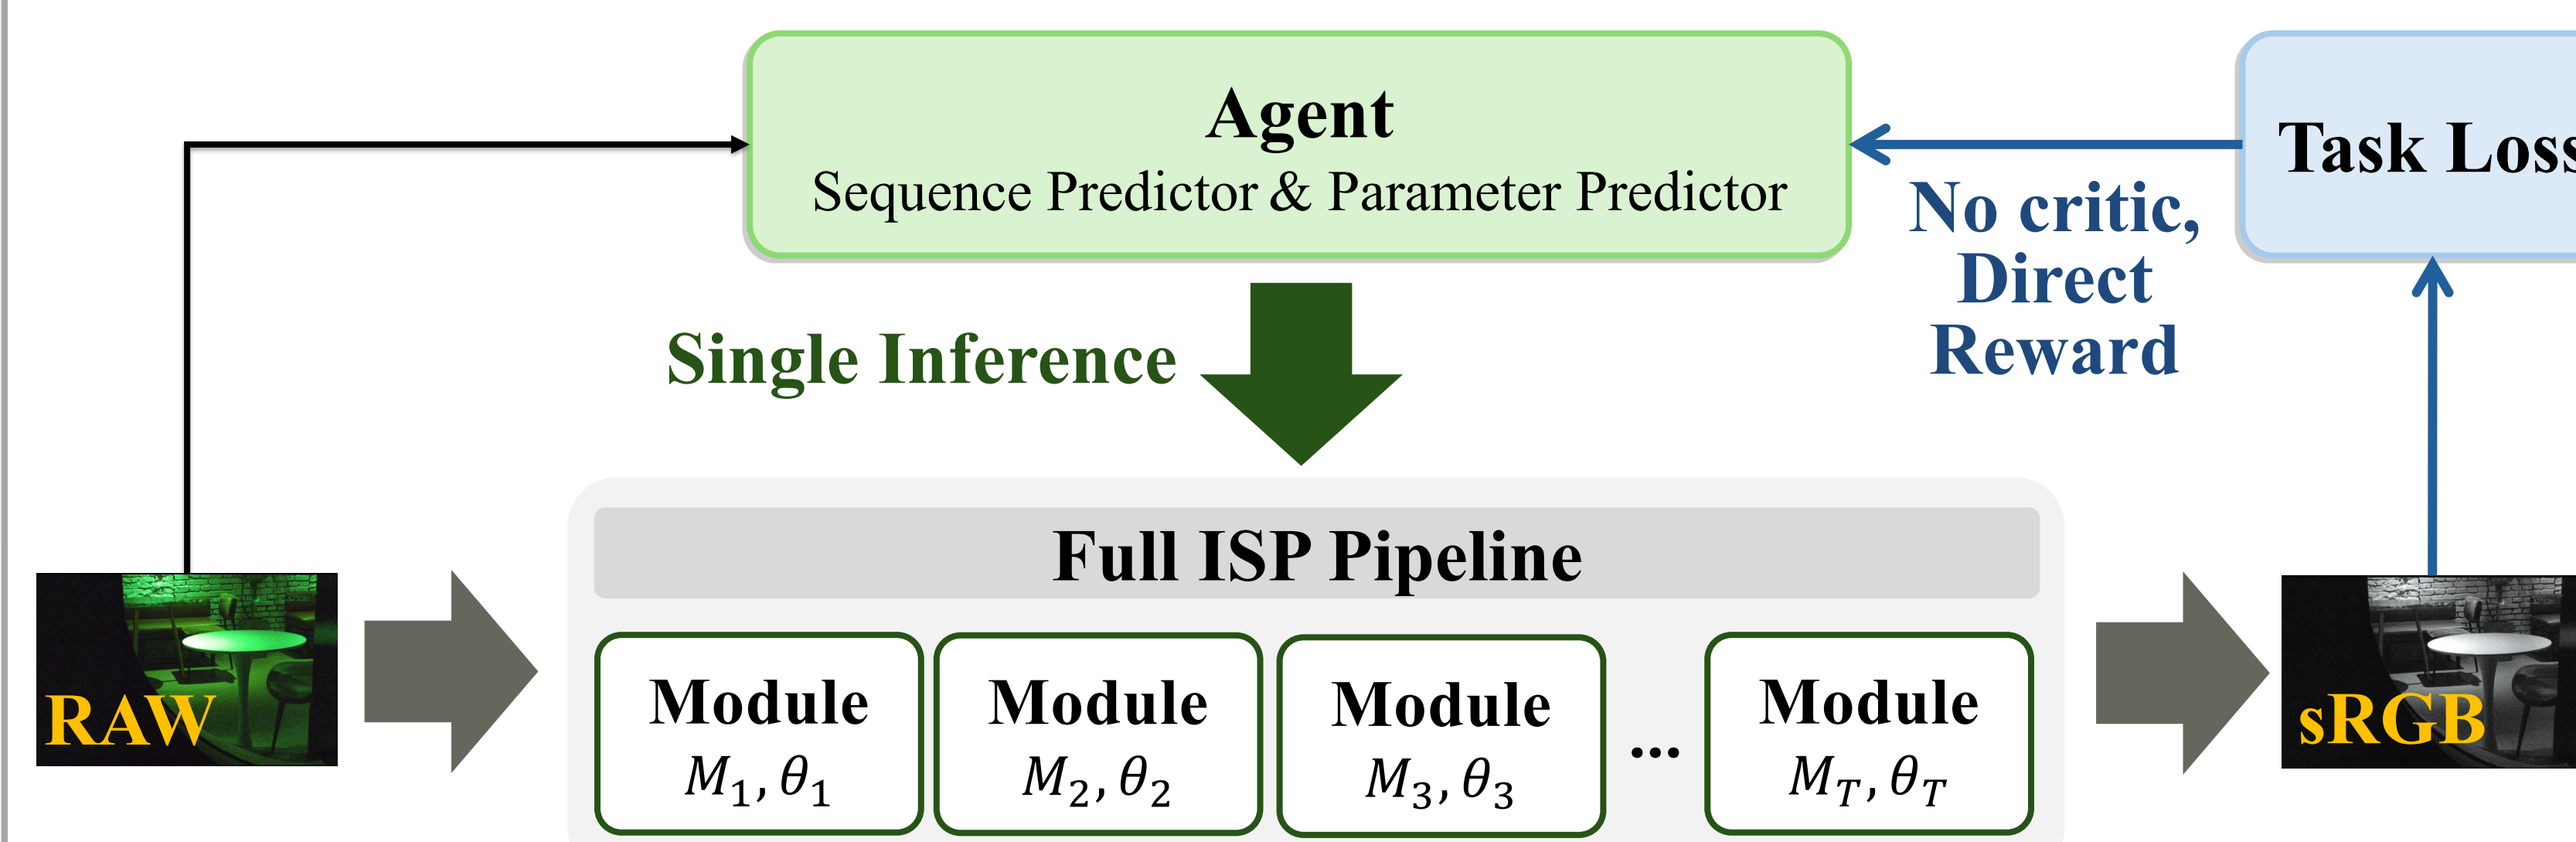
- Step-wise pipeline optimization
 - ⇒ Repeated agent decisions, increasing computational cost
- Intermediate reward estimation
 - ⇒ Relies on critic-estimated rewards, causing unstable learning



Our Approach

Key Idea: Direct task reward learning at the sequence level

- Predict the entire ISP pipeline at once
 - ⇒ Improves computational efficiency
- Final task reward optimization
 - ⇒ Eliminates unstable reward estimation, leading to stable training



Quantitative Comparison

+0.6 mAP

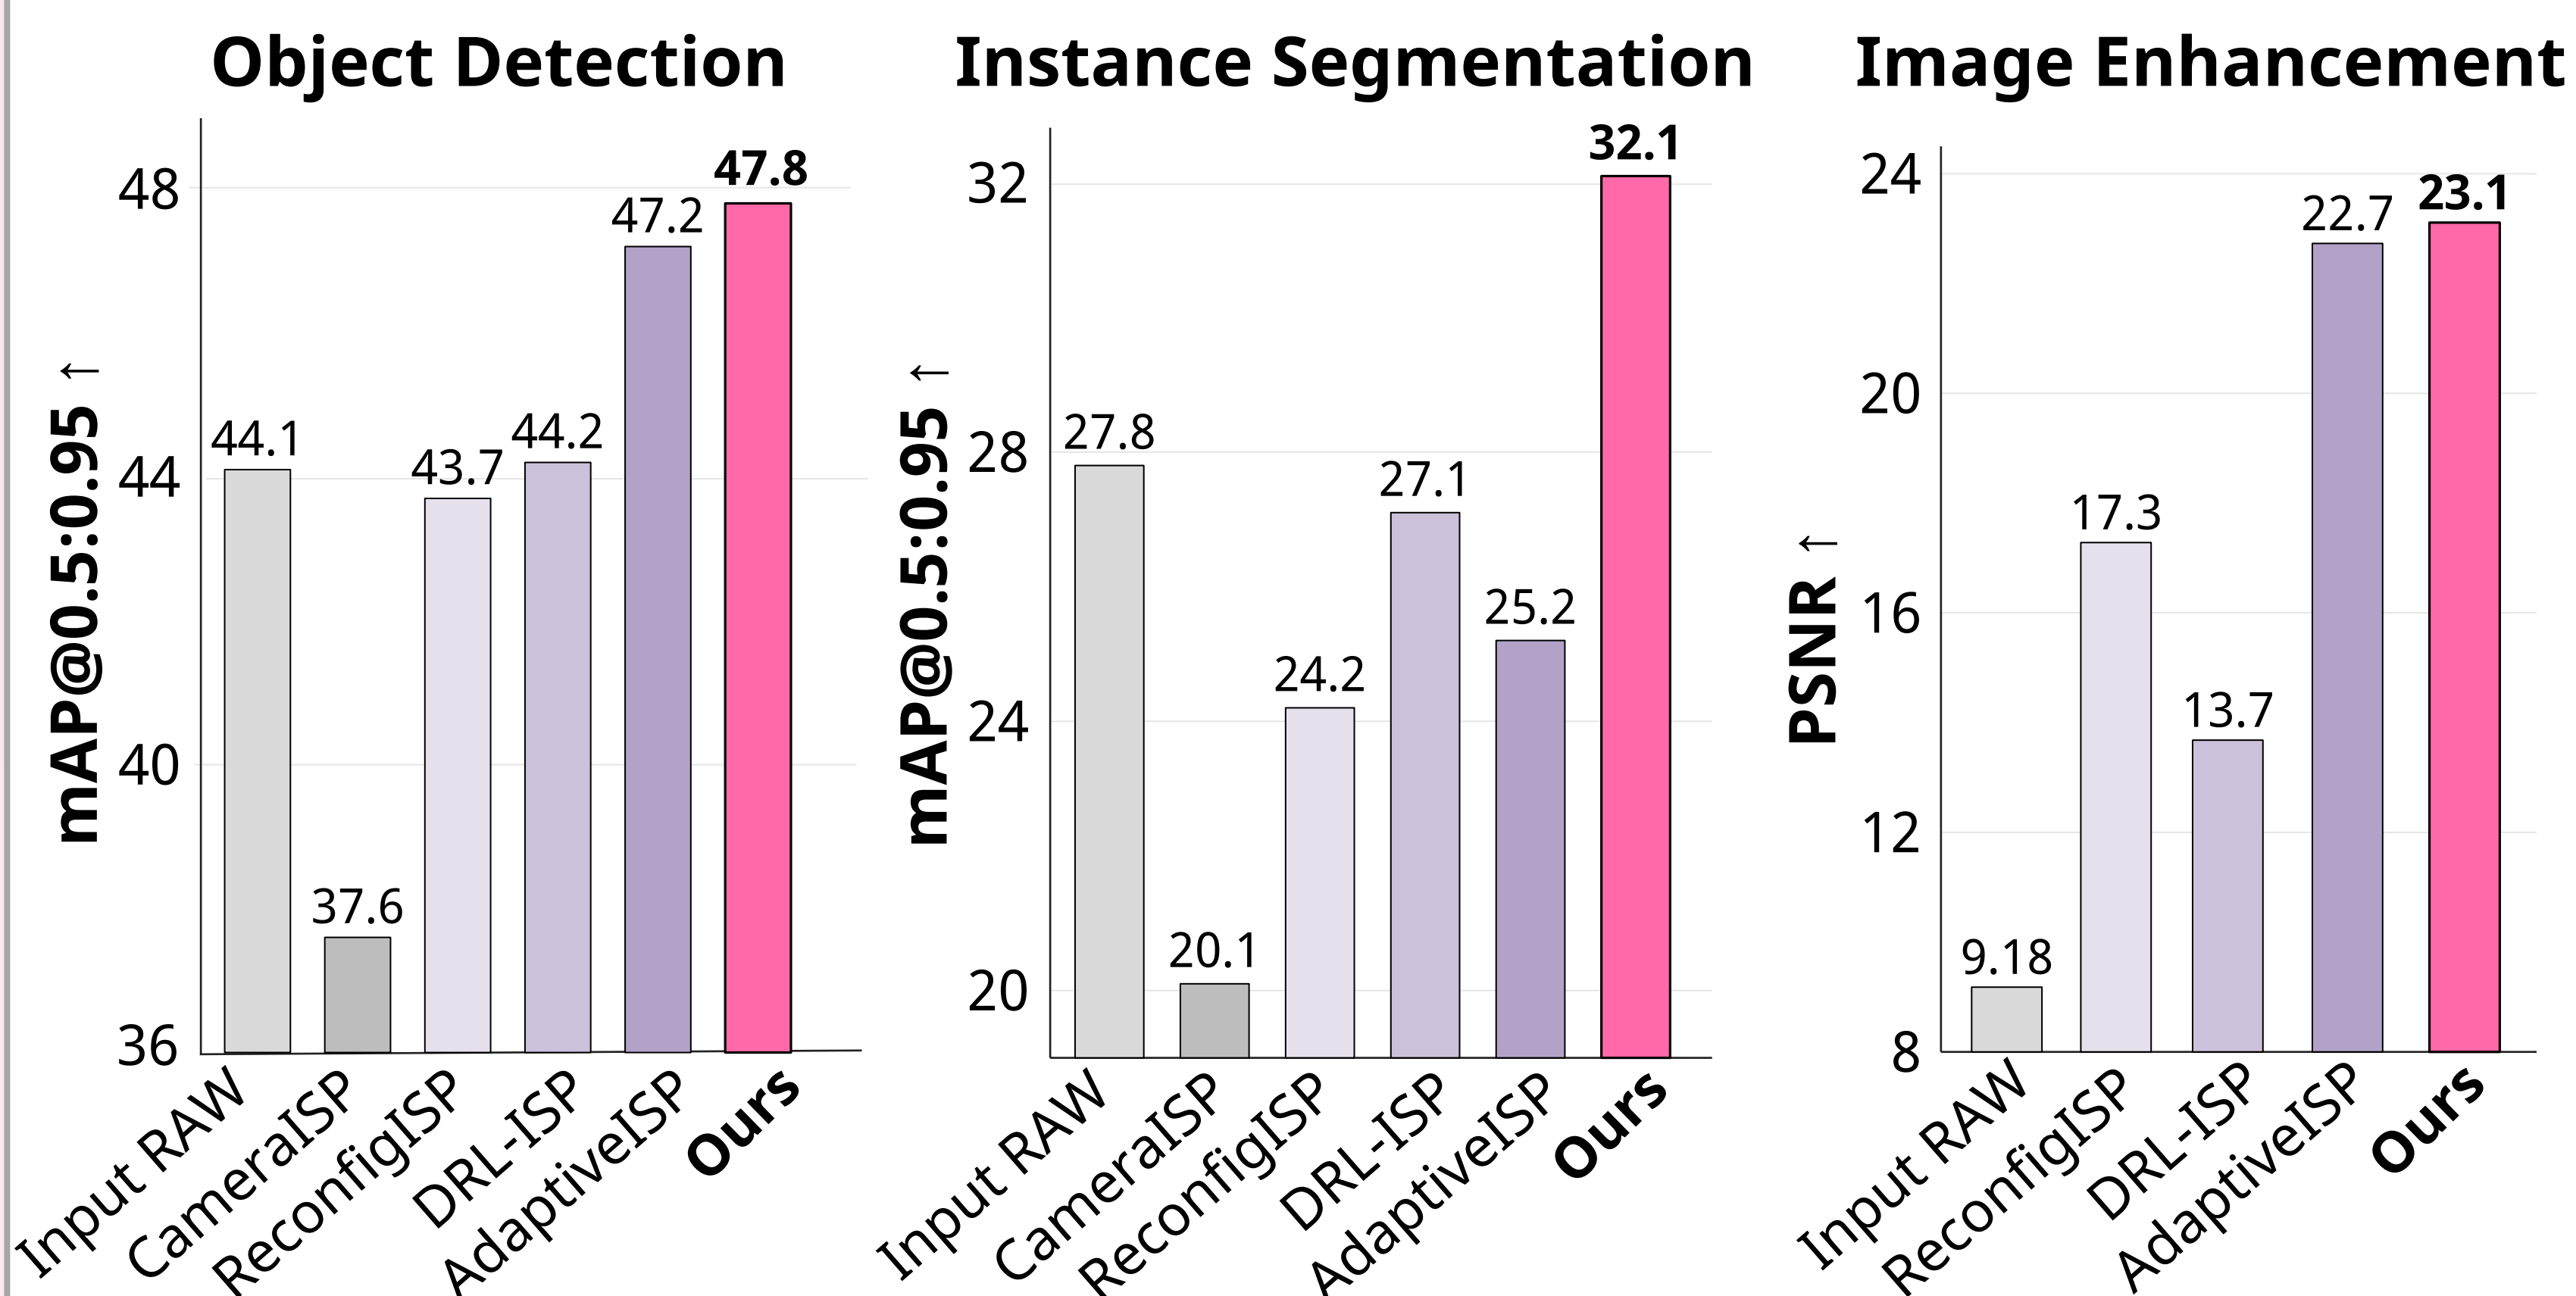
over previous best on detection

+5.0 mAP

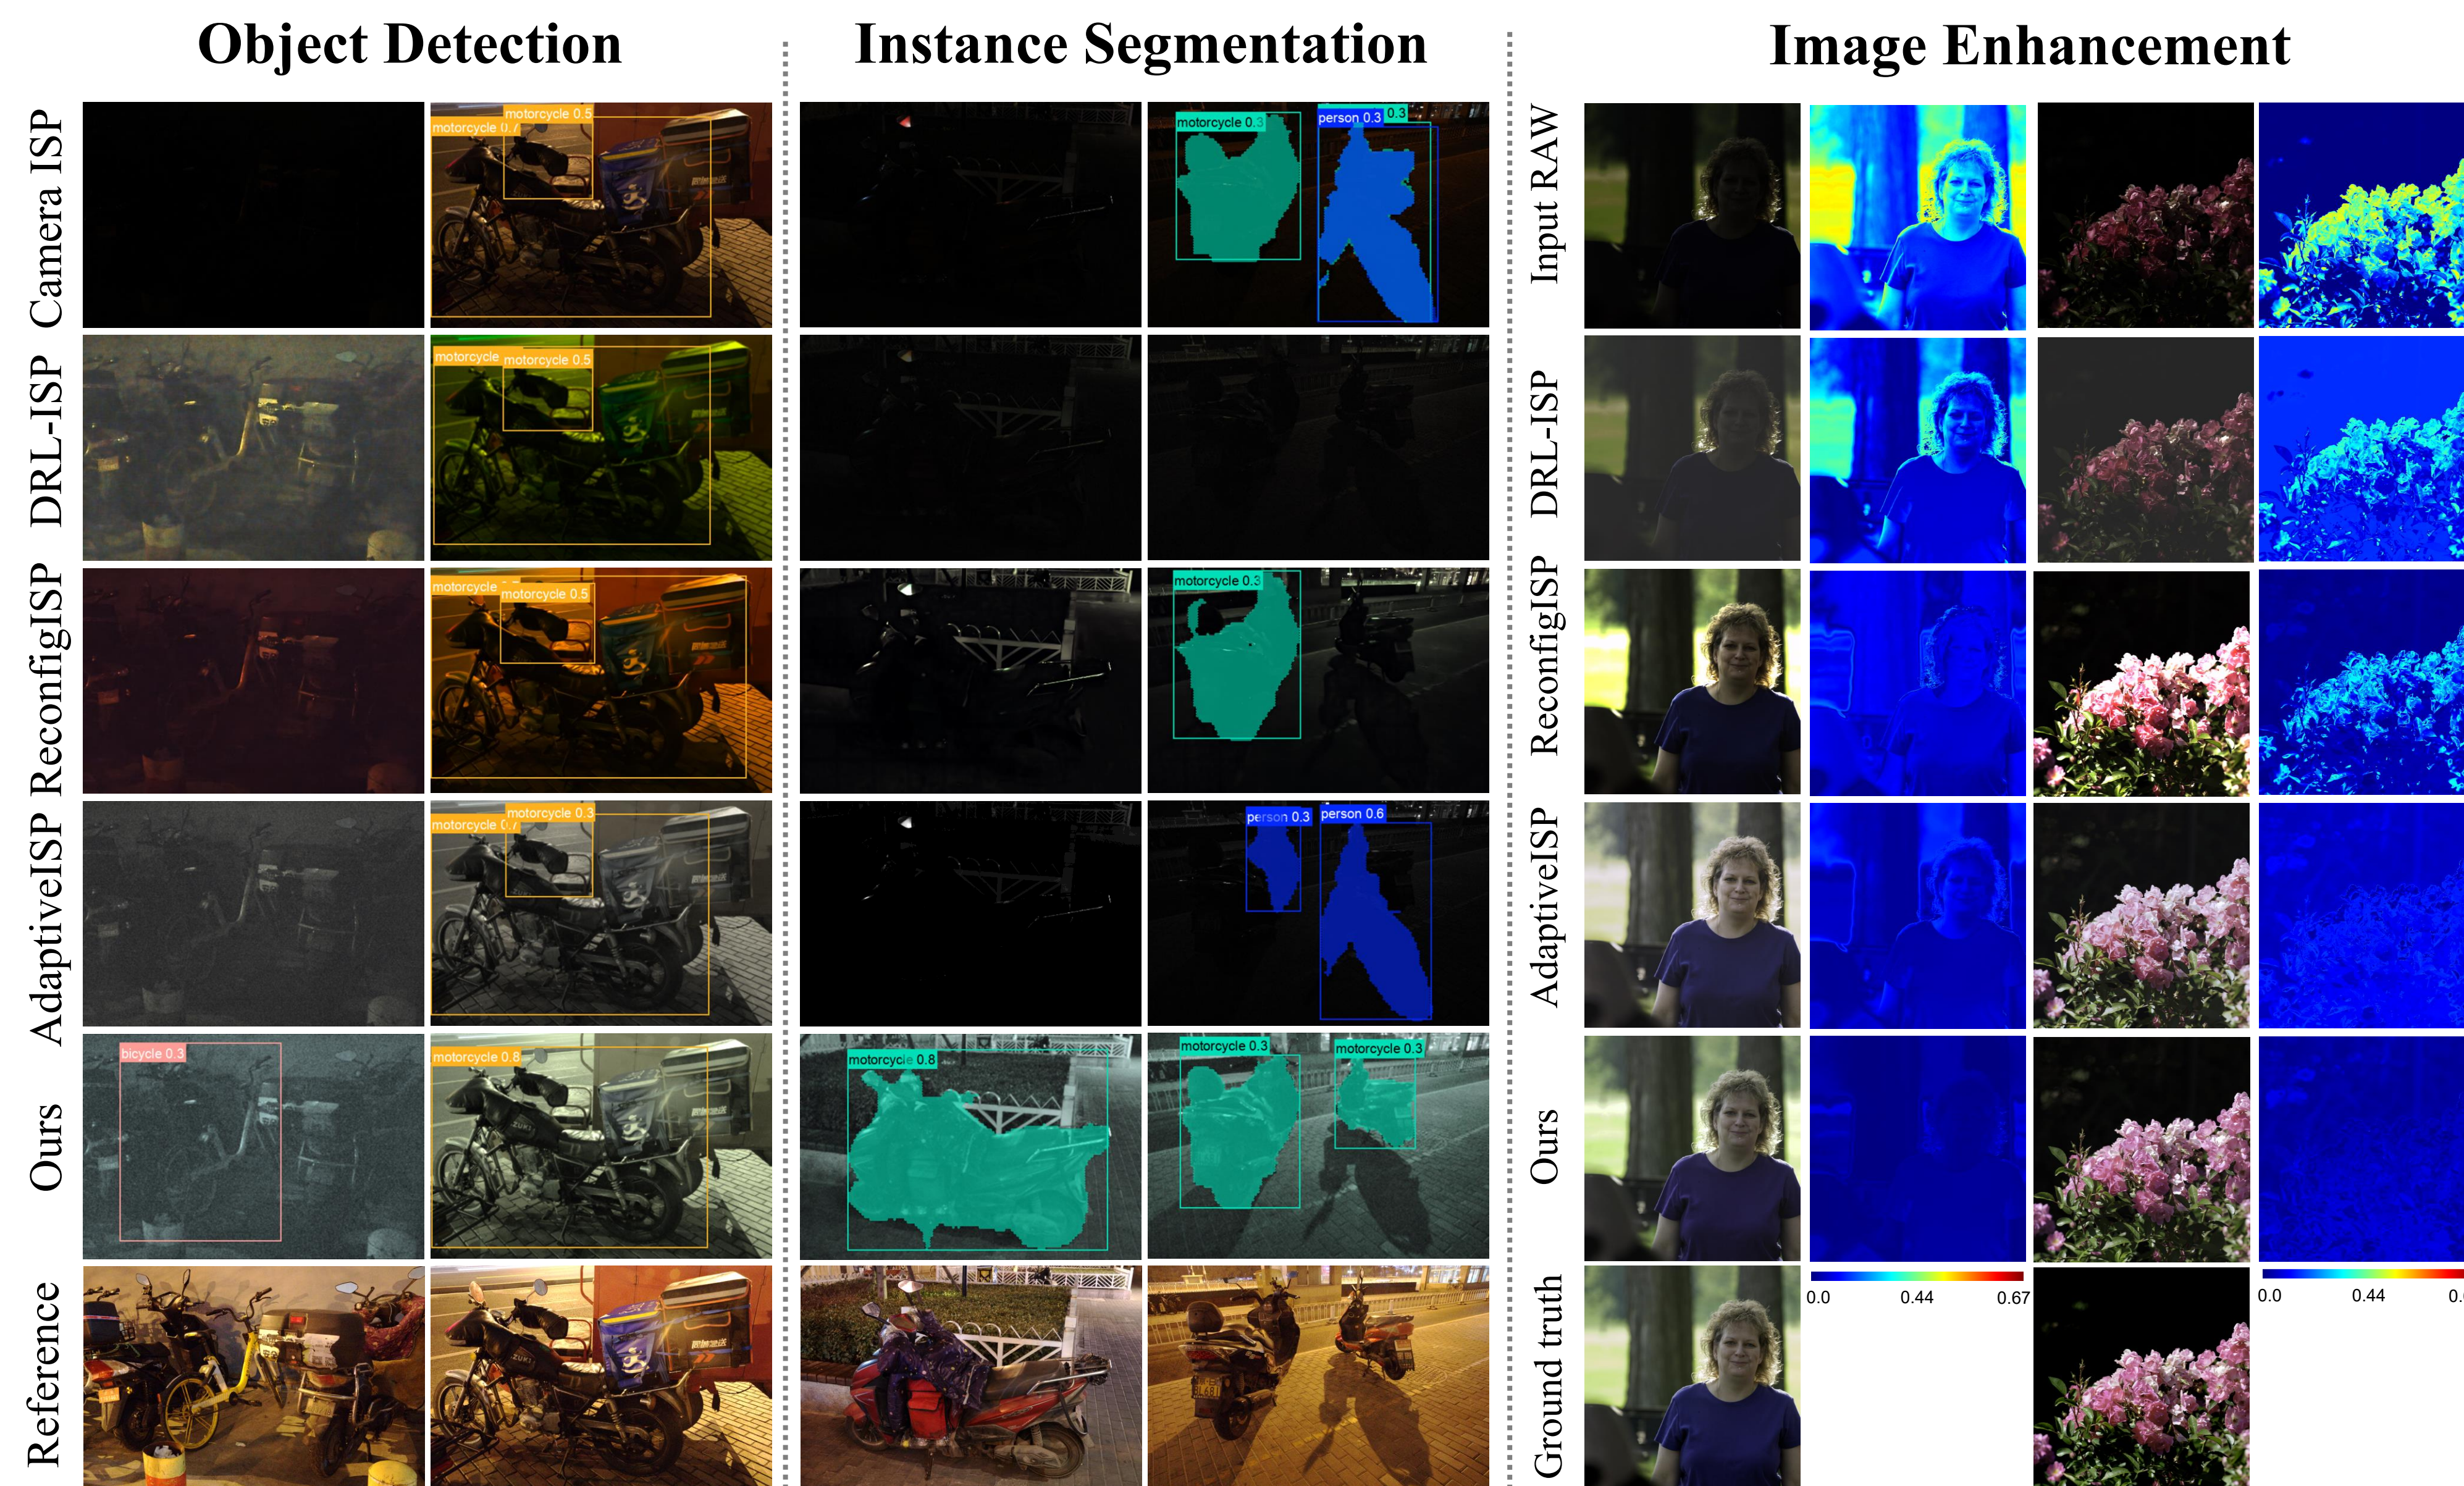
over previous best on segmentation

+0.38 dB

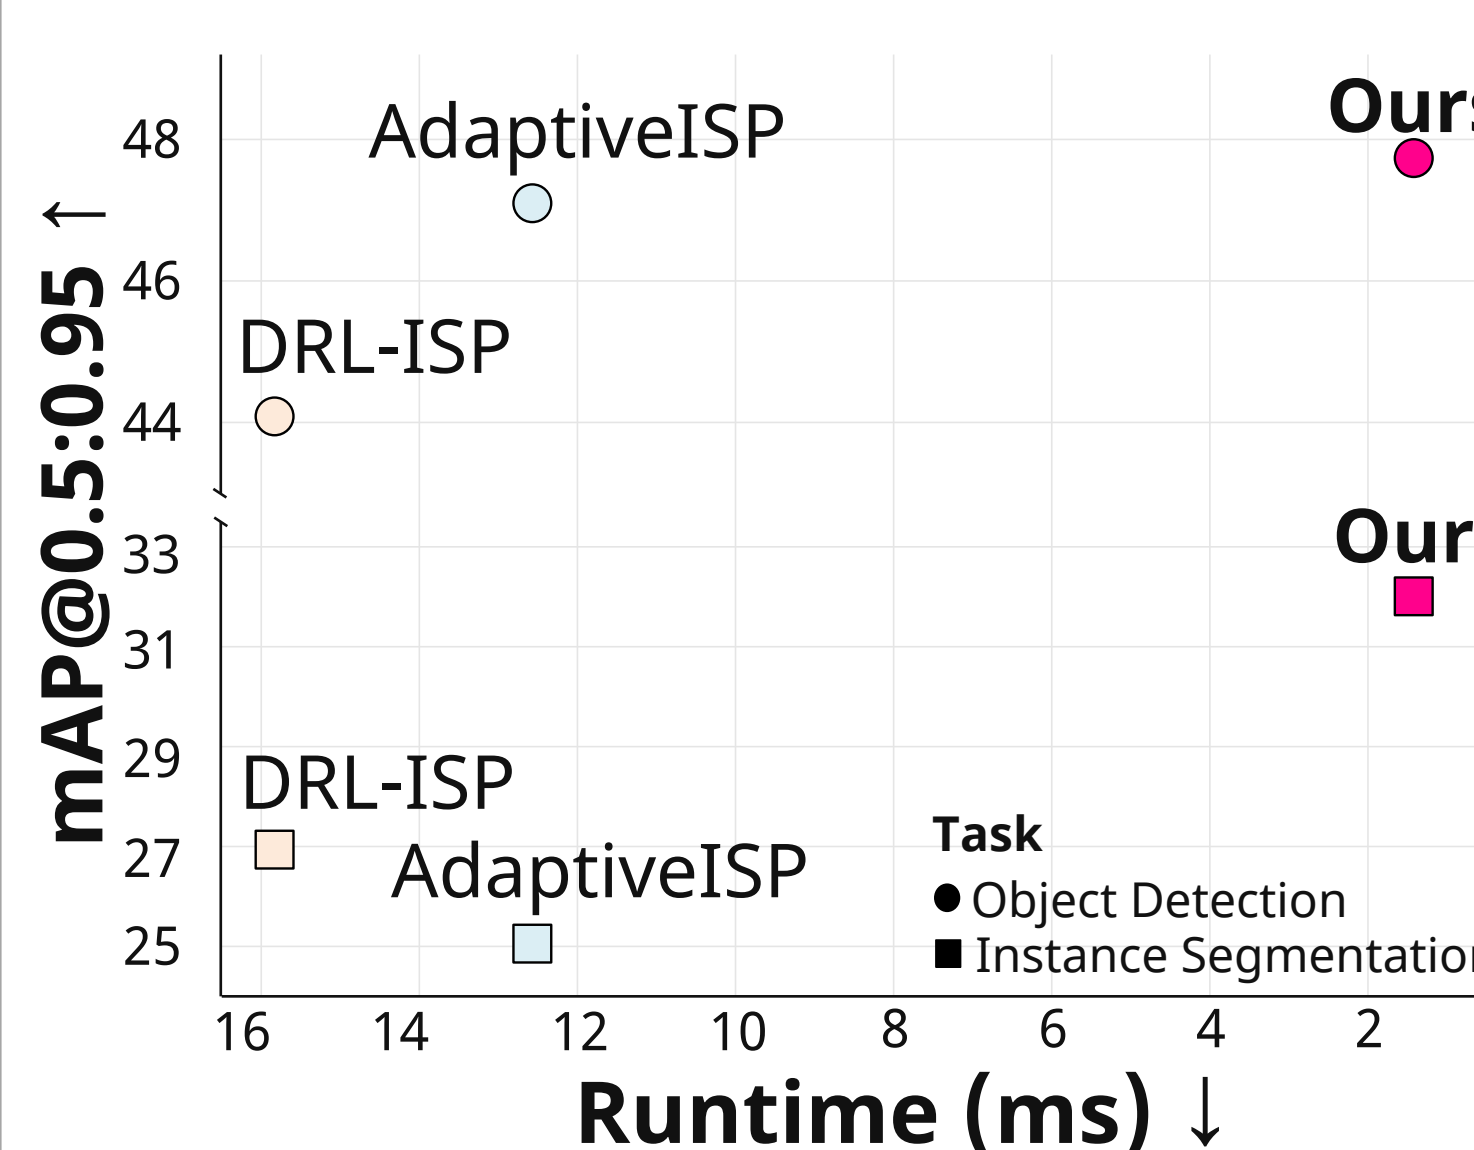
over previous best on enhancement



Qualitative Comparison



Computational Efficiency



8.2x speedup
over previous best on runtime

12.4x fewer params
over previous best on parameters

Method	Params (M)	MACs (M)	Peak GPU Memory (MB)	Runtime (ms)
DRL-ISP [1]	6.57	155.3	1013.9	15.71
AdaptiveISP [3]	7.18	70.2	39.6	12.72
Ours	0.53	15.1	14.4	1.55

Reference

- Yu et al. "ReconfigISP: Reconfigurable camera image processing pipeline." ICCV. 2021.
- Shin et al. "DRL-ISP: Multi-objective camera ISP with deep reinforcement learning." IROS. 2022.
- Wang et al. "AdaptiveISP: Learning an adaptive image signal processor for object detection." NeurIPS. 2024.