# An Empirical Analysis of Range for 3D Object Detection

Neehar Peri, Mengtian Li, Benjamin Wilson, Yu-Xiong Wang, James Hays, Deva Ramanan

## (1) Long-Range Detection

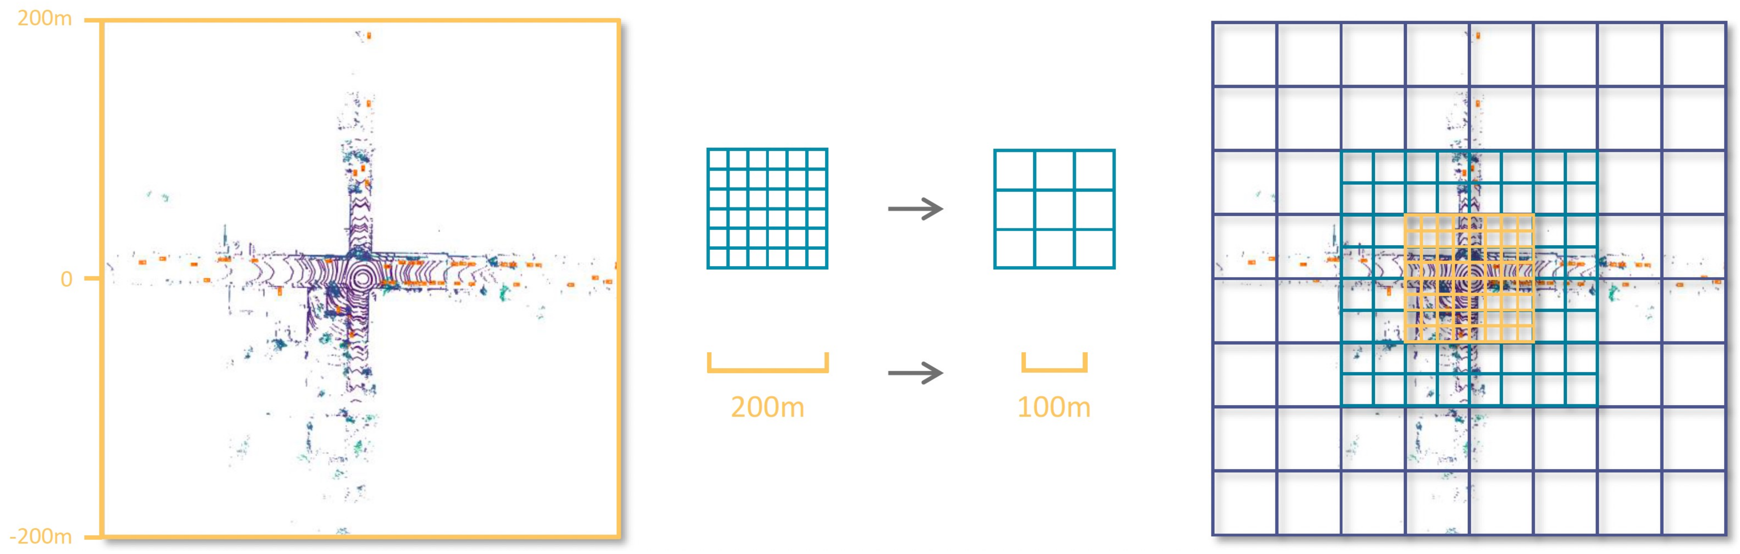


- Existing 3D LiDAR detectors struggle to detect far away objects (e.g. 300m) due to time and compute constraints.
- To manage compute, we can adopt a coarser grid or limit the processing range.

## (3) Near-Far Range Ensemble

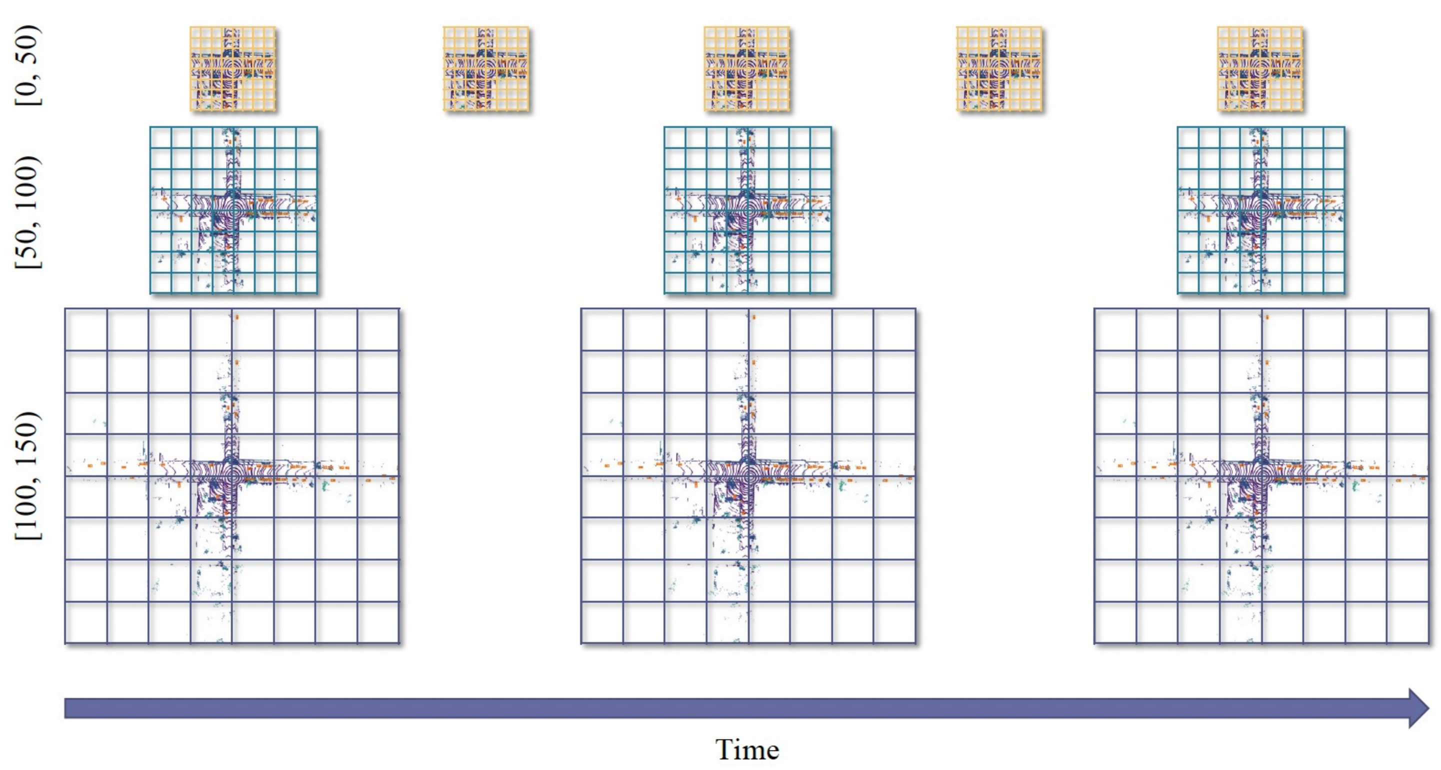


```
for time, lidar_sweep in enumerate(data):
    # Run near-field range expert
    near_dets = near_expert(lidar_sweep)

    if time % freq == 0:
        # Run far-field range expert
        cropped_sweep = donut_crop(lidar_sweep)
        far_dets = far_expert(cropped_sweep)
    else:
        # Forecast prev. detections
        far_dets = forecast(dets[time - 1])

    dets[time] = {near_dets, far_dets}
```

- We can trivially speed up multi-range ensembles via range-specific asynchronous processing.
- Inspired by hierarchical "slow-fast" planners that run a low-frequency planner with a high frequency reactive controller, we can run near-range experts at high frequency (to avoid immediate collision) and run far-range experts at lower frequency (for long-term planning).

## (4) Accelerating Inference

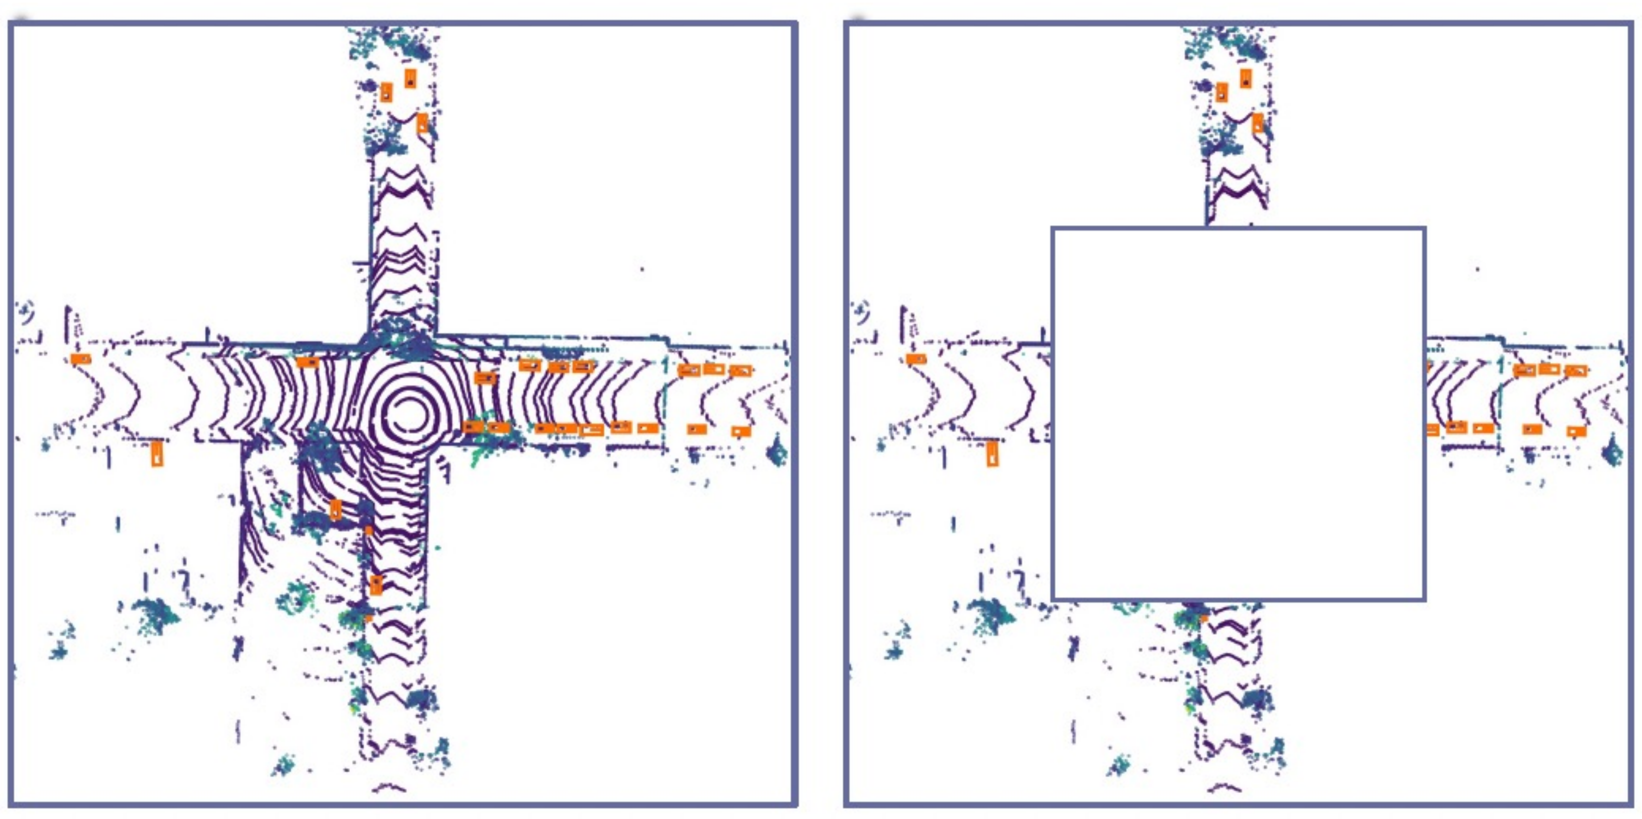

- Since each range-expert in our ensemble only contributes within a range interval, we can speed up inference by embracing sparse convolutional processing. We can simply mask out all other points outside of the processing range.

Full paper available at arxiv.org/abs/2308.04054

## (2) Accuracy vs. Latency Trade-Off

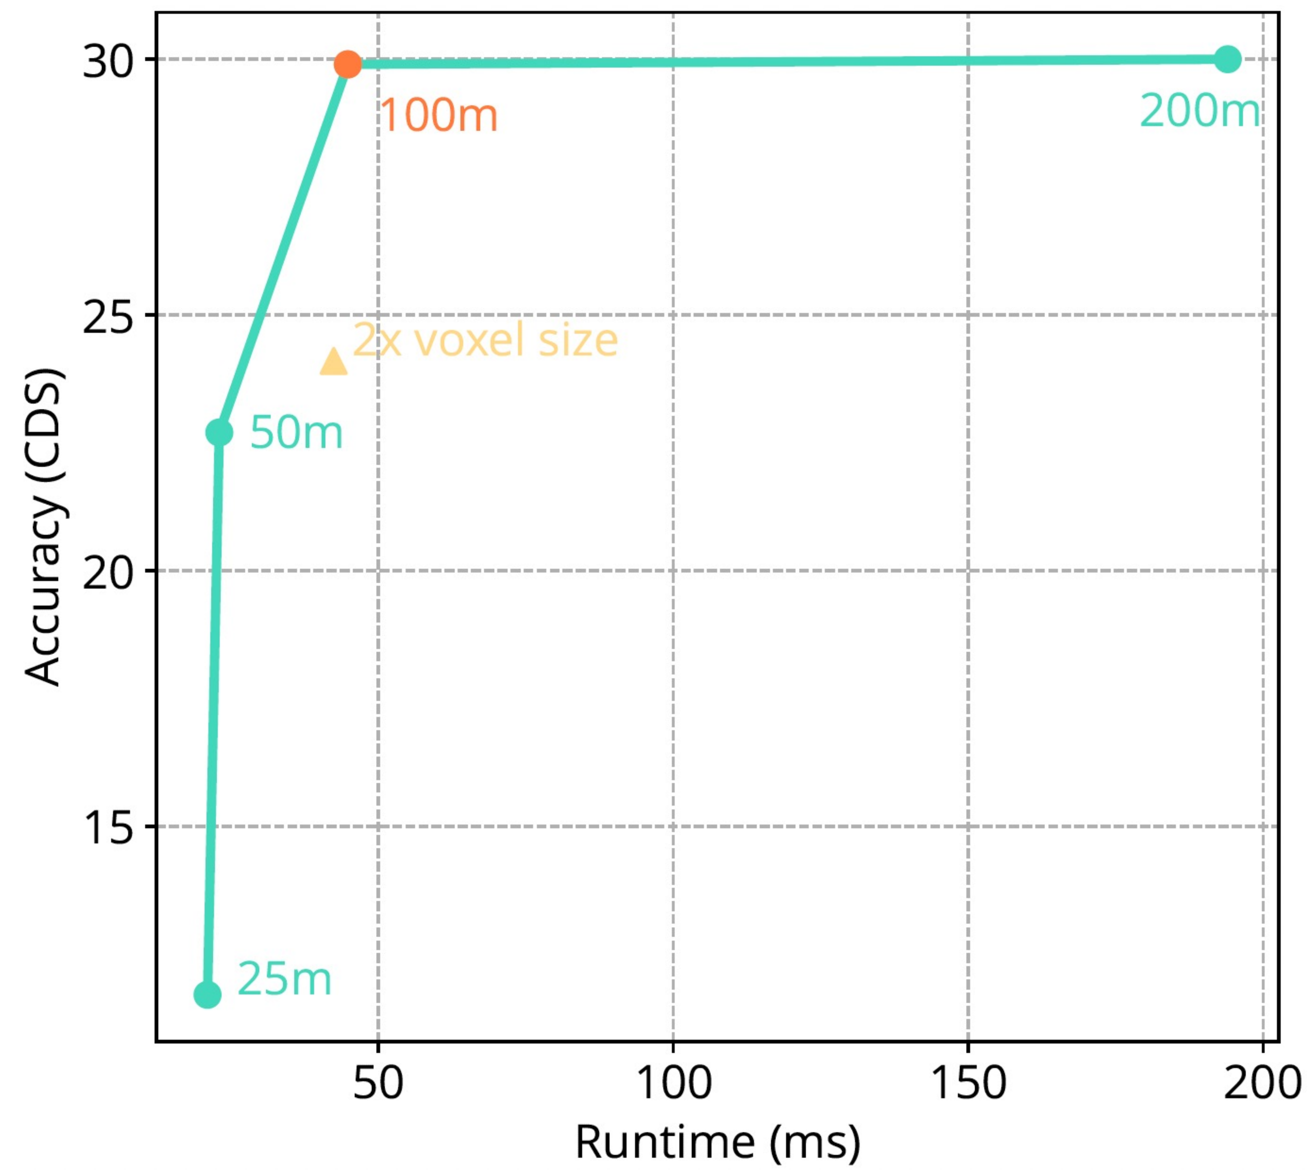


- **Range experts can generalize.** The 100m range-expert (orange point) generalizes to different ranges via fully-convolutional processing.
- **Range is the most effective "knob" for trading off accuracy-vs-latency.** Using 2x larger voxels (yellow triangle) improves latency but reduces performance.
- **It (apparently) pays to "give up" on long range.** Running the 100m range expert at 200m does not improve performance but increases latency. Can we do better? Use range ensembles!

## (5) Across-Range Generalization

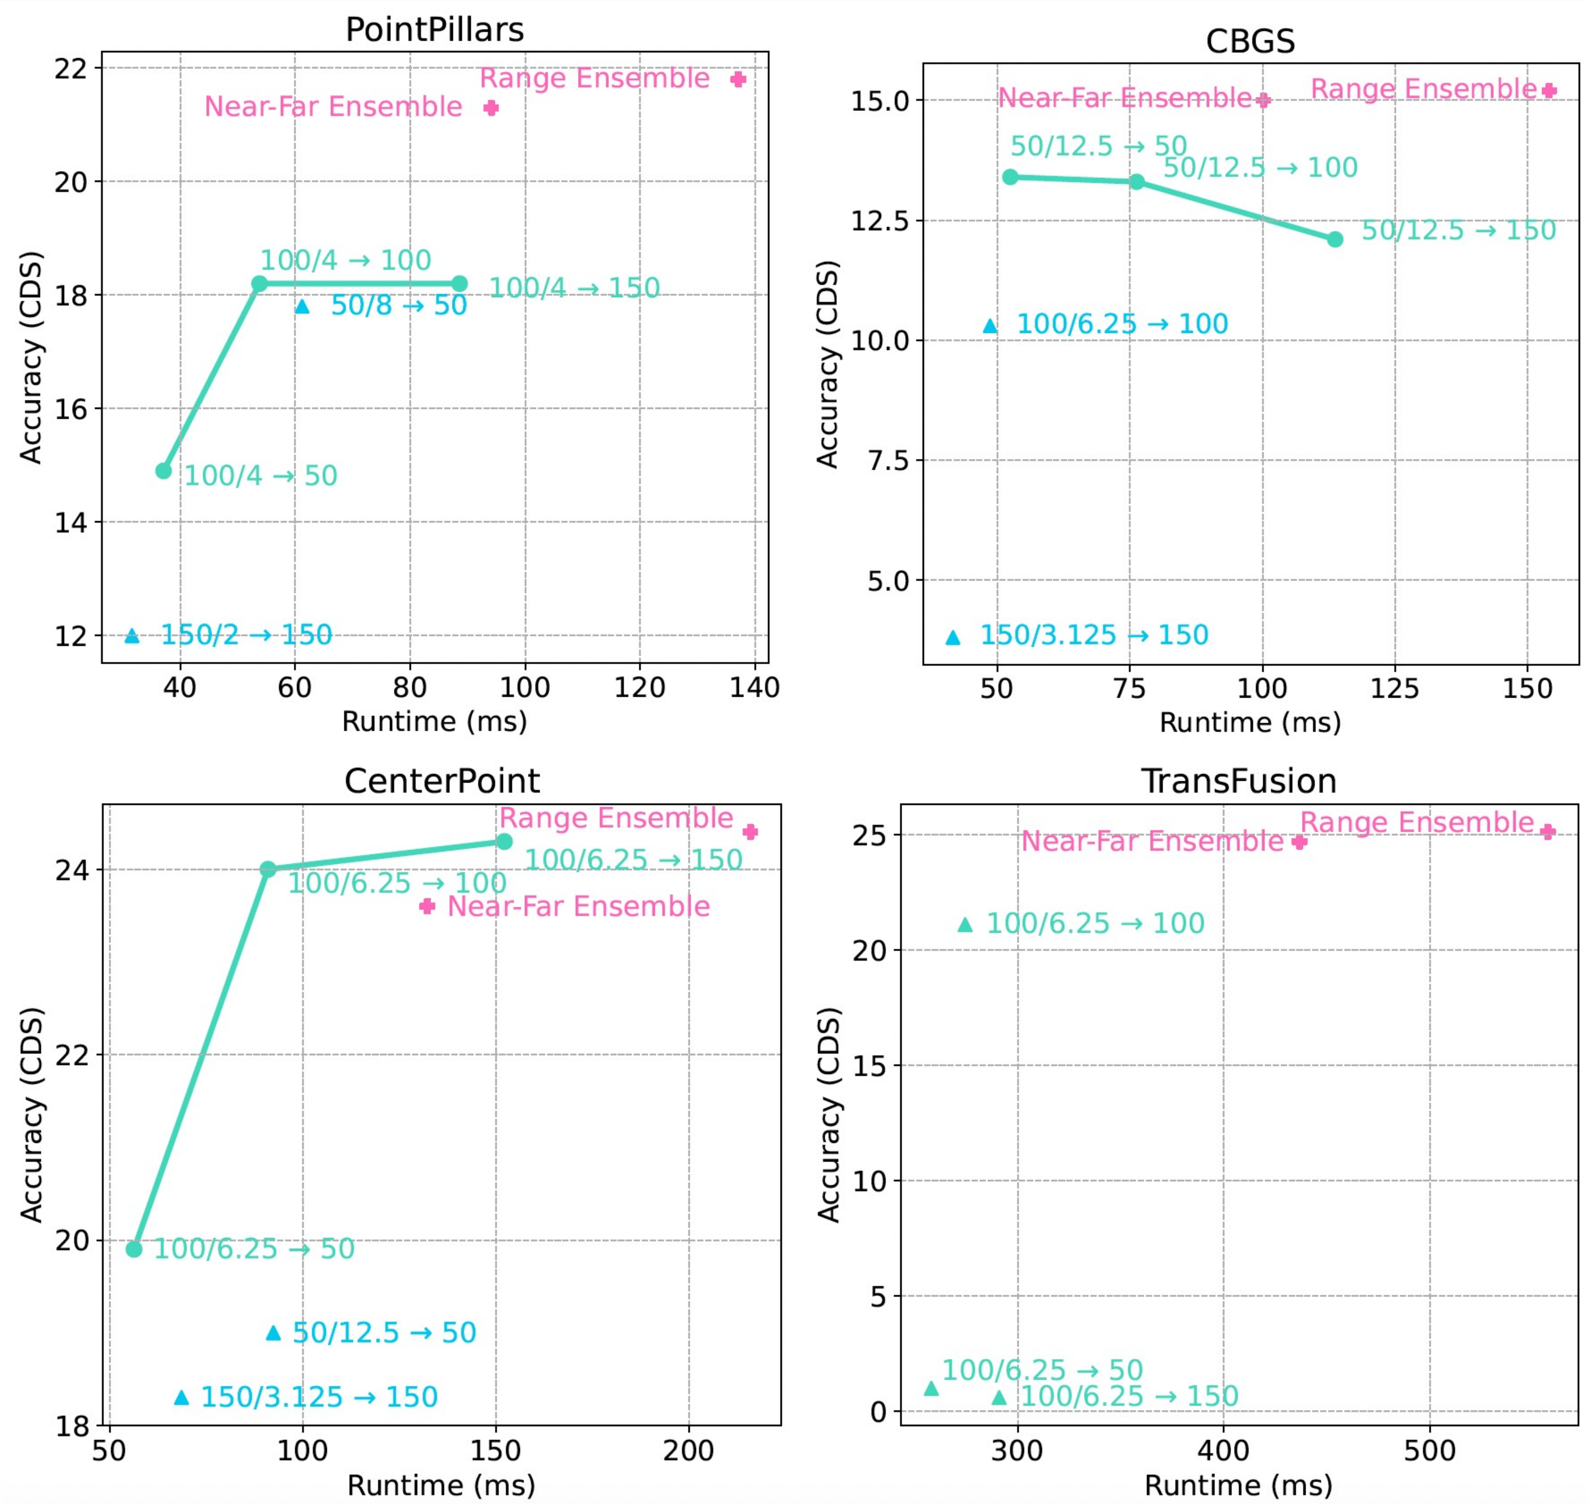


- Detectors have different across-range generalization characteristics with fully-convolutional processing
- **PointPillars.** The 100m range-expert outperforms both the 50m and 150m range experts, suggesting that we should "give up" on far-field detection.
- **CenterPoint.** The 100m range-expert evaluated at 150m nearly matches the performance of the range-ensemble, suggesting that model ensembles may not always be necessary
- **CBGS.** The 50m range-expert outperforms the 100m and 150m experts. However, running the 50m model at far range degrades performance, suggesting poor far-field generalization.
- **TransFusion.** We posit that the use of relative positional encoding rather than metric encoding leads to catastrophically poor across-range generalization