Self-Supervised Learning of 4D Automotive Radar Scene Flow Estimation

Abstract

Scene flow allows autonomous vehicles to reason about the arbitrary motion of multiple independent objects which is the key to long-term mobile autonomy. While estimating the scene flow from LiDAR has progressed recently, it remains largely unknown how to estimate the scene flow from a 4D radar - an increasingly popular automotive sensor for its robustness against adverse weather and lighting conditions. Compared with the LiDAR point clouds, radar data are drastically sparser, noisier, and in much lower resolution. Annotated datasets for radar scene flow are also absent and costly to acquire in the real world. These factors jointly pose the radar scene flow estimation as a challenging problem.

This work aims to address the above challenges and estimate scene flow from 4D radar point clouds by leveraging self-supervised learning. To this end, we propose a novel scene flow estimation framework with a *Radar-Oriented Flow Estimation* (ROFE) module to derive a basic scene flow and a *Static Flow Refinement* (SFR) module to refine predictions of identified static points, as seen in Fig. 1. To enable robust self-supervised learning, we design three task-specific losses to cope with sparse and noisy radar point clouds. We collected an in-house multi-modal dataset by driving a vehicle in the wild for the evaluation of our method. Real-world experiment results validate that our method can robustly estimate the radar scene flow in the wild and effectively supports the downstream task of motion segmentation. Our *RaFlow* method is our first attempt at using 4D radars to enable robust long-term mobile autonomy. We plan to complement our work by exploiting the cross-modal supervision signals from other co-located sensors (e.g. RGB camera or IMU) to further improve the radar scene flow estimation. Applying radar scene flow to enable other downstream tasks, such as point cloud accumulation and multi-object tracking, is also to be investigated.

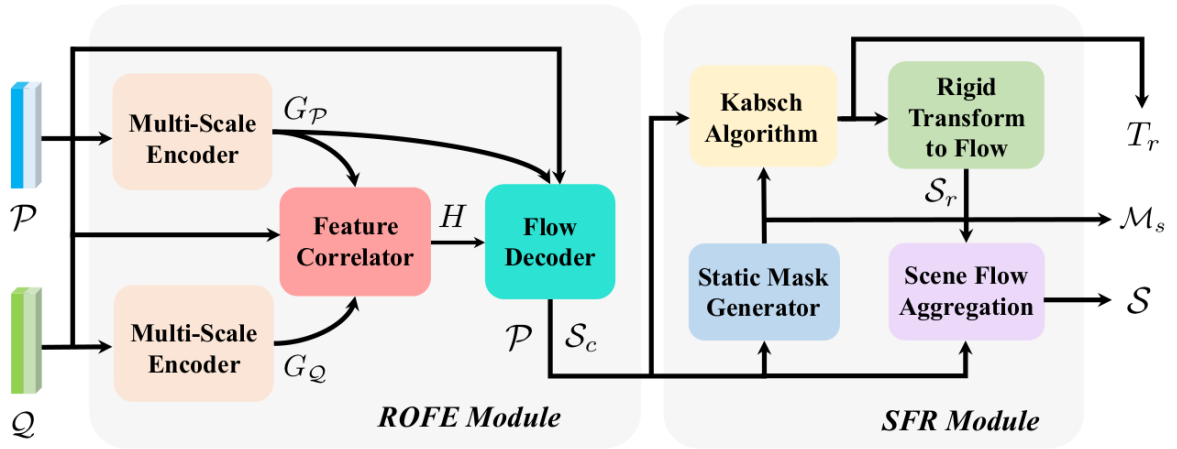

Figure 1: Overview of our radar scene flow estimation pipeline.