AV22 | 12th Autonomous Vehicles DETROIT

ODD BY USE CASE | GO-TO-MARKET STRATEGIES | ROAD TO AV > GO FROM L2 TO L5

EXCLUSIVE INTERVIEW WITH

RAMIN NABATI

AUTONOMOUS VEHICLE SENSOR FUSION ENGINEER

FORD MOTOR COMPANY

Automotive iQ

This September, Automotive IQ hosts the 12th Autonomous Vehicles conference in Detroit. The event will focus on the current manufacturing and solutions, as well as the big issues that Detroit is tackling to bring AVs into production and to the market. The event will showcase use cases and approaches on how to take technology, develop partnerships that are necessary to finish a product, develop a complete AV stack from top to bottom, bring it to market and gain acceptance from consumers or from enterprises.

We know that Sensor Fusion plays a huge part in the success of AVs, something that Ramin Nabati, Autonomous Vehicle Sensor Fusion Engineer at Ford Motor Company, focuses on in his role. Ramin will be speaking at the AVDT event this September, so we spoke to him to learn more.

Q: So Ramin, can you tell us a little bit about your role at Ford Motor Company?

Ramin:

I'm an AV Sensor Fusion Engineer with the AV Future Tech team, focusing on research, design, implementation, testing, and evaluation of sensor fusion algorithms for object detection and tracking in autonomous vehicles. This involves working with a variety of sensors such as camera, radar, LIDAR and IMU, from data acquisition to object detection, tracking and fusion.

Q: Well, we imagine you've got your hands full, especially as it's clear the industry is facing real challenges bringing AVs to production and market. In your personal opinion, what is the number 1 reason why OEMs are struggling to productionize AVs?

Ramin:

There are many reasons why we don't see AVs being productionized yet, ranging from a lack of regulations and standards to social acceptability. But being an engineer, I like to focus more on the technical aspects. In my opinion, the number 1 roadblock for productionizing AVs is cost, which might not sound like an engineering problem, but it soon becomes one when you need to design the system on a budget.

Over the past several years there has been tremendous progress in software and hardware technologies for autonomous driving, from accurate deep-learning-based object detection algorithms to solid-state LIDARs and 4D imaging radars. These technologies help solve many of the challenges associated with autonomous driving, but they are still very costly for production-level vehicles. State-of-the-art detection, tracking and fusion algorithms heavily rely on GPU processing for real-time performance, and many of them require dense point cloud and high-resolution images as inputs. Having all the required sensors to achieve an acceptable level of perception accuracy, and a powerful compute system capable of processing the generated data in real-time is prohibitively expensive for a production vehicle.

Q: In your personal opinion, how has sensor technology for AVs developed over the last 5 years, and what difference can we see in another 5 years to come?

Ramin:

As I mentioned before, the sensor technology for AVs has had a tremendous progress in the past few years. For LIDARS we've seen a shift towards using the solid-state technology, reducing mechanical components which results in lower cost and better integration. Automotive radar technology has also been improved upon significantly, with 4D imaging and software defined radars bringing a whole new dimension to radar detections and generating almost LIDAR-like point clouds. We've also seen improvements in power efficiency of some of these sensors, making it easier to design the system to meet the available power budget on a vehicle. I think we will continue to see such improvements, which hopefully will result in reduced cost for these sensors.

Q: What do you think are the right KPIs to ensure high levels of accuracy with Sensor Fusion?

Ramin:

It really depends on the fusion algorithm used, the functionality it's helping with and the stage it's applied at. For example, sensor fusion could be used to increase object detection accuracy, improve object tracking or even for obstacle trajectory prediction. Each one of these tasks would require a different set of metrics for evaluation. In object detection, the precision and recall values could be used as metrics to evaluate fusion-based detection results and compare them to detection results, using only a single sensing modality.

For object tracking you would use other metrics, such as number of identity switches or average multi object tracking accuracy (AMOTA). So, I don't think there is really a one-size-fits-all kind of KPI to ensure accurate sensor fusion results. What's more important than the specific metric you use, is to make sure you consider the sensors' capabilities and limitations when evaluating the fusion algorithm. For instance, how much a particular LIDAR-camera fusion algorithm can help with improving the object detection accuracy could be very different in daylight vs night-time.

Q: How important is robust design?

Ramin:

Having a robust fusion algorithm is crucial for autonomous driving applications. Robustness in sensor fusion could be defined in different ways. The system should be able to handle input streams that are distorted due to adverse weather conditions, such as snow or fog. These are scenarios that are common in driving scenarios but might not have been represented adequately in the datasets used for training and evaluation.

Depending on the functional safety requirements, the fusion algorithm may also need to be robust to temporary loss of input from one or multiple sensing modalities. Proper design of the fusion algorithm and the stage at which it is applied (early, middle, or late fusion) are important factors to mitigate these problems.

Q: We know that Sensor Fusion improves the overall performance capability of an Autonomous Vehicle, but in your personal opinion, what multiple fusion techniques are being used or are available? How do you determine which one to use (depending on the feature's Operation Design Domain)?

Ramin:

Depending on the stage at which fusion is applied to the sensing modalities, fusion algorithms could be categorized into 3 different classes: early, middle and late fusion. In early fusion, the raw data from different sensing modalities are combined with minimal or no pre-processing. The system is then able to learn a joint representation of different sensing modalities which could significantly improve accuracy.

In late fusion, on the other hand, data streams from different sensing modalities are processed in separate pipelines, and only the final results of each pipeline are fused together. This approach, also called decision-level fusion, is the most commonly used approach in autonomous driving applications due to its simplicity and robustness to loss of input.

Finally, middle fusion is the middle ground between early and late fusion, where fusion happens at a feature level. In this approach data streams from different modalities are processed first to extract useful features, and then are combined to generate decisions. Choosing the right fusion approach depends on many different factors, including the Operation Design Domain, functional safety requirements, sensor types and the available processing power.

Q: In your personal opinion, how important are partnerships across the AV ecosystem?

Ramin:

Partnerships at different levels could really help accelerate productionizing AVs. Tier 1 suppliers, for example, have valuable experience in mass production for OEMs, which could significantly help companies developing AV sensing and compute equipment to lower their production cost. This also applies to companies providing AV software solutions, where forming partnerships and sharing the data could greatly reduce the data collection costs. These partnerships could also help with shaping and developing the required standards for autonomous driving, which is an important factor in productionizing AVs.

Q: You've had an interesting career that has led you to become a sensor fusion specialist, can you please give us some highlights of your career, focusing on any mind-blowing discoveries with the development of sensor fusion, as well as any challenges that you've faced?

Ramin:

I started my journey in autonomous driving when I joined the University of Tennessee's EcoCAR3 team as the ADAS Lead in 2017. EcoCAR3 was the latest US Department of Energy (DOE) Advanced Vehicle Technology Competition (AVTC), focusing on improving energy efficiency and reducing environmental impact of a 2016 Chevrolet Camaro using advanced propulsion systems and ADAS technologies. My team were tasked with the design and implementation of a Forward Collision Warning (FCW) system using a radar and camera, providing haptic feedbacks to the driver.

Our FCW system used a very simple late fusion approach using target association in the camera frame, but performed extremely well in the competition, winning us the first place in ADAS demonstration among 16 teams from the US and Canada. Being hands-on in every step of this project, from sensor selection, simulation, placement, and calibration to data acquisition, sensor fusion and feedback generation really set the foundation for my PhD dissertation on radar-camera sensor fusion, as well as my career.

And the most important lesson I learned from that project is best captured by the Occam's razor principle: The simplest solution is almost always the best.

A big thank you to Ramin for speaking with us. We're looking forward to his presentation at the 12th Autonomous Vehicles Detroit 2022, Detroit, Michigan, September 20 - 22, 2022. If you'd like to attend, please register for the event.

Steven Wicks
Online Content Manager
Automotive IQ