


Tips for Designing a Robust Computer Vision System for Self-driving Cars



Victor Cheng



The robustness and reliability of a self-driving car's computer vision system has received a lot of news coverage. As a vision software engineer at TI helping customers implement advanced driver assistance systems (ADAS) on our TDAx platform, I know how hard it is to design a robust vision system that performs well in any environmental condition.

When you think about it, engineers have been trying to mimic the visual system of human beings. As Leonardo DaVinci said, "Human subtlety will never devise an invention more beautiful, more simple or more direct than does nature, because in her inventions nothing is lacking and nothing is superfluous."

Indeed, I was recently painfully reminded of this fact. On March 2, 2018, I caught an eye infection which was eventually diagnosed as a severe adenovirus. After a month, my vision is almost back to normal. Throughout my health ordeal, I learned a few things about the human visual system that are applicable to our modern-day challenge of making self-driving cars.

The Importance of Sensor Redundancy

The virus affected my right eye first, and the vision in that eye became very blurry. When I had both eyes open, however, my vision was still relatively good, as though my brain was selecting the image produced only by my left eye but using the blurry image produced by my right eye to assess distance. From this, I can infer that stereo vision doesn't need to operate on full resolution images, although of course that is optimal. A downsampled image and even a reduced frame rate should be OK.

When my right eye became so painful I could no longer open it, I had to rely on my left eye only. Although my overall vision was still OK, I had difficulty assessing the distance of objects. During my recovery, my right eye started healing first, and my brain did the same thing once more: it relied on the eye that was improving.

From these observations, I can draw these conclusions about autonomous driving: for each position around the car where vision is used for object detection, there should be multiple cameras (at least two) pointing to the line of sight. This setup should be in place even when the vision algorithms only need monovision data.

Sensor multiplicity allows for failure detection of the primary camera by comparing images with auxiliary cameras. The primary camera feeds its data to the vision algorithm. If the system detects a failure of the primary camera, it should be able to reroute one of the auxiliary cameras' data to the vision algorithm.

With multiple cameras available, the vision algorithm should also take advantage of stereo vision. Collecting depth data at a lower resolution and lower frame rate will conserve the processing power. Even when processing is monocular by nature, depth information can speed up object classification by reducing the number of scales that need processing based on the minimum and maximum distances of objects in the scene.

TI has planned for such requirements by equipping its TDAx line of automotive processors with the necessary technology to handle at least eight camera inputs, and to perform state-of-the-art stereo-vision processing through a vision accelerator pack.

The Importance of Low-light Vision, Reliance on an Offline Map and Sensor Fusion

After the virus affected both of my eyes, I became so sensitive to light I had to close all the window blinds in my home and live in nearly total darkness. I managed to move around despite little light and poor vision because I could distinguish object shapes and remember their location.

From this experience, I believe low-light vision processing requires a mode of processing different than the one used in daylight, as images captured in low-light conditions have a low signal-to-noise ratio and structured elements, such as edges, are buried beneath the noise. In low-light conditions, I think vision algorithms should rely more on blobs or shapes rather than edges. Since histogram of oriented gradients (HOG)-based object classifications rely mostly on edges, I expect they would perform poorly in low-light conditions.

If the system detects low-light conditions, the vision algorithm should switch to a low-light mode. This mode could be implemented as a deep learning network that is trained using low-light images only. Low-light mode should also rely on data from an offline map or an offline world view. A low-light vision algorithm can provide cues in order to find the correct location on a map and reconstruct a scene from an offline world view, which should be enough for navigation in static environments. In dynamic environments, however, with moving or new objects that were not previously recorded, fusion with other sensors (LIDAR, radar, thermal cameras, etc.) will be necessary in order to ensure optimum performance.

TI's [TDA2P](#) and [TDA3x](#) processors have a hardware image signal processor supporting wide-dynamic-range sensors for low-light image processing. The TI Deep Learning (TIDL) library implemented using the vision accelerator pack can take complex deep learning networks designed with the Caffe or Tensor flow frameworks and execute them in real time within a 2.5W power envelope. Semantic segmentation and single-shot detector are among the networks successfully demonstrated on TDA2x processors.

To complement our vision technology, TI has been ramping efforts to develop radar technology tailored for ADAS and the autonomous driving market. The results include:

- Automotive radar millimeter-wave (mmWave) sensors such as the AWR1xx for mid- and long-range radar.
- A software development kit running on TDAx processors that implements the radar signal processing chain, enabling the processing of as many as four radar signals.

The Importance of Faulty Sensor Detection and a Fail-safe Mechanism

When my symptoms were at their worst, even closing my eyes wouldn't relieve the pain. I was also seeing strobes of light and colored patterns that kept me from sleeping. My brain was acting as if my eyes were open, and was constantly trying to process the incoming noisy images. I wished it could detect that my eyes were not functioning properly and cut the signal! I guess I found a flaw in nature's design.

In the world of autonomous driving, a faulty sensor or even dirt can have life-threatening consequences, since a noisy image can fool the vision algorithm and lead to incorrect classifications. I think there will be a greater focus applied to developing algorithms that can detect invalid scenes produced by a faulty sensor. The system can implement fail-safe mechanisms, such as activating the emergency lights or gradually bringing the car to a halt.

If you are involved in developing self-driving technology, I hope my experience will inspire you to make your computer-vision systems more robust.

Additional Resources

- Read these white papers:
 - [“Stereo vision – Facing the challenges and seeing the opportunities for ADAS applications.”](#)
 - [“AWR1243 sensor: Highly integrated 76-81GHz radar front end for emerging ADAS applications.”](#)
 - [“Embedded low-power deep learning with TIDL.”](#)
 - Order the [processor software development kit for TDAx ADAS systems on chip with Linux® and TI real-time operating system \(RTOS\) support](#).
 - [Learn more about TI’s ADAS solutions](#)

IMPORTANT NOTICE AND DISCLAIMER

TI PROVIDES TECHNICAL AND RELIABILITY DATA (INCLUDING DATA SHEETS), DESIGN RESOURCES (INCLUDING REFERENCE DESIGNS), APPLICATION OR OTHER DESIGN ADVICE, WEB TOOLS, SAFETY INFORMATION, AND OTHER RESOURCES "AS IS" AND WITH ALL FAULTS, AND DISCLAIMS ALL WARRANTIES, EXPRESS AND IMPLIED, INCLUDING WITHOUT LIMITATION ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT OF THIRD PARTY INTELLECTUAL PROPERTY RIGHTS.

These resources are intended for skilled developers designing with TI products. You are solely responsible for (1) selecting the appropriate TI products for your application, (2) designing, validating and testing your application, and (3) ensuring your application meets applicable standards, and any other safety, security, regulatory or other requirements.

These resources are subject to change without notice. TI grants you permission to use these resources only for development of an application that uses the TI products described in the resource. Other reproduction and display of these resources is prohibited. No license is granted to any other TI intellectual property right or to any third party intellectual property right. TI disclaims responsibility for, and you will fully indemnify TI and its representatives against, any claims, damages, costs, losses, and liabilities arising out of your use of these resources.

TI's products are provided subject to [TI's Terms of Sale](#) or other applicable terms available either on [ti.com](https://www.ti.com) or provided in conjunction with such TI products. TI's provision of these resources does not expand or otherwise alter TI's applicable warranties or warranty disclaimers for TI products.

TI objects to and rejects any additional or different terms you may have proposed.

Mailing Address: Texas Instruments, Post Office Box 655303, Dallas, Texas 75265
Copyright © 2023, Texas Instruments Incorporated