Improving Inexact Measurements: Active and Passive Learning Compared

Martin Vikström Morin

June 12, 2017

When dealing with complex systems, the task of accurately measuring and identifying all of the involved parameters and states becomes difficult. One solution is to simply add more sensors of better quality but this isn't always resource effective or even possible. The research here covers a different approach where, with the help of computer algorithms, all available information is melded and fused together, producing a combined result that is better than any one single measurement could give. This passive approach is also combined and compared with a method that tries to actively learn about the true system by steering it so to give more informative measurements. However, this active learning greatly increases the complexity of the problem so it was a pleasant result that, for the case examined, the passive approach produced results of equal quality.

An often recurring problem in the modern world is the problem of control, i.e. given some process, what steering input should be chosen in order to achieve some desired goal? The processes can be almost anything, from large industrial chemical plants and mechanical systems such as anti-lock brakes in cars, to completely digital systems such as large computer networks.

How each of these systems are steered all differ but to achieve good control it's important to know the state of the processes. In an anti-lock breaking system these states might be the velocity of the car and the tires together with the amount of break pressure applied. For a completely digital system it might be something completely different such as the amount of traffic each router in a computer networks has to deal with. Regardless of what type of process that is controlled, the accurate identification of these states are key.

The motivation for the particular research performed here is a counter intuitive, voluntary introduction of bad measurements. The less precise measurements are easier to encode, reducing the amount of information that needs to be sent from the sensor to the controller. This is desirable for processes where this data is sent over some large network. These networks might have long transmission times and be shared with other services, risking traffic congestions. Reducing the amount of data makes the controller less sensitive to both of these problems.

To still get the necessary accuracy from the inexact measurement a method known as a particle filter was used. It's an increasingly popular approach that has been made possible by rapid development of computational power of the last couple of decades. The particle filter passively learns how to best combine all available information to improve the accuracy but it was also compared with a couple of active approaches.

How active learning works can be understood by imagining driving blindfolded on a road. There are rumble strips on either side so as long as no vibrations are felt you could be anywhere in the lane but as soon vibrations are felt, you know that you are at one of the edges. By steering right until a vibration is felt you would know that you are at the right side. This way you have actively learned and improved your knowledge of your position based on your steering.

When performing a case study it was found that the passive learning of the particle filter showed great performance improvements while the active approaches did not improve it further. In the blindfolded driving analogy it means that the swerving of actively steering towards the rumble strip is worse than simply waiting for the car to slowly drift towards one side. In both the driving example and case study are these result good, waiting and seeing is a very easy strategy to perform correctly.