Teach Xplore

Teaching robots to hunt down prey like an animal

July 7 2016, by Bob Yirka

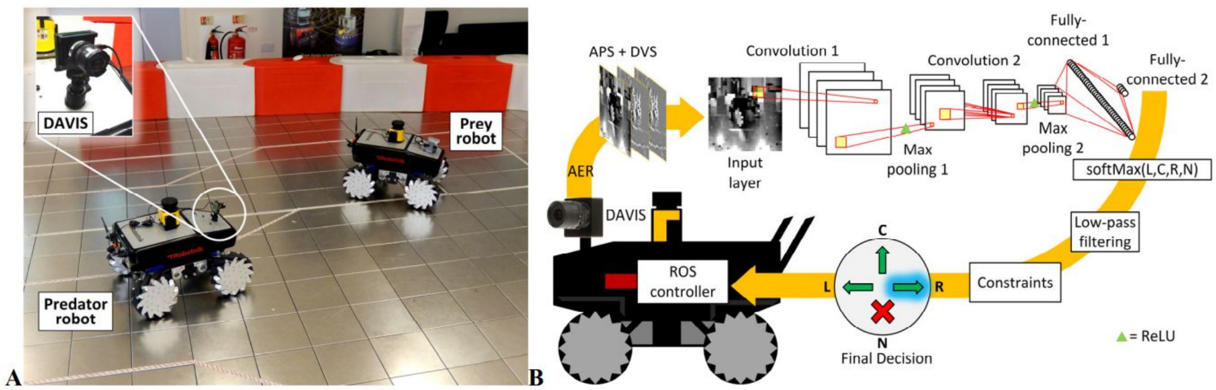

A: Summit XL predator (left) chasing the prey (right) in the robot arena of the University of Ulster, Londonderry. B: overall closed-loop system: the DAVIS sensor generates APS and DVS data which is alternately fed to the 4C5-R-2S-4C5-R-2S-40F-R-4F convolutional neural network. The results are filtered and the final decision is used in conjunction with the laser scanner output to control the Summit XL behavior in the ROS controller. Credit: arXiv:1606.09433 [cs.RO]

(TechXplore)—A team of researchers with the Institute of Neuroinformatics at the University of Zurich in Switzerland has been making news after uploading a paper to the preprint server *arXiv*, describing their experiments teaching robots to track and 'hunt' other robots. The idea is not new of course, it has been used in video games for decades—what is new is the technological approach that makes it a viable option for machines operating in a three dimensional world.

In order for a predator such as a lion to track and take down another animal, it must process an incredible amount of data very quickly; most of that information is visual of course, but some is also intuitive—guessing which way an antelope calf will turn, for example, based on prior experience. To give a robot such an ability, the researchers found, required the combination of two major elements: a different kind of camera and a specialized type of neural network (convolutional) that has been programmed to work in ways similar to the animal visual cortex.

The use of the network was an obvious choice, the software allows for learning, which is what animals must do in order to survive, in this case, learning about how to interpret movement. Giving one to a robot could help it learn to predict target movements. What was truly novel in this new work was the use of a silicon based camera, which was not really a normal camera at all. Modern digital cameras capture all of the information that passes through a lens at a given time, allowing for the production of a picture, or a video. But animal eyes do not really work that way—in addition to collecting information about color, brightness and the other information needed to form a picture, they also have parts that are sensitive to movement and change. To give a robot such an ability, the team used what they describe as a silicon retina—it senses changes the way animals do (except it uses pixels instead of neurons), which is critical for pursuing prey. Adding both to a tractor shaped robot gave it an ability to track and follow the movements of another similar robot, in real time.

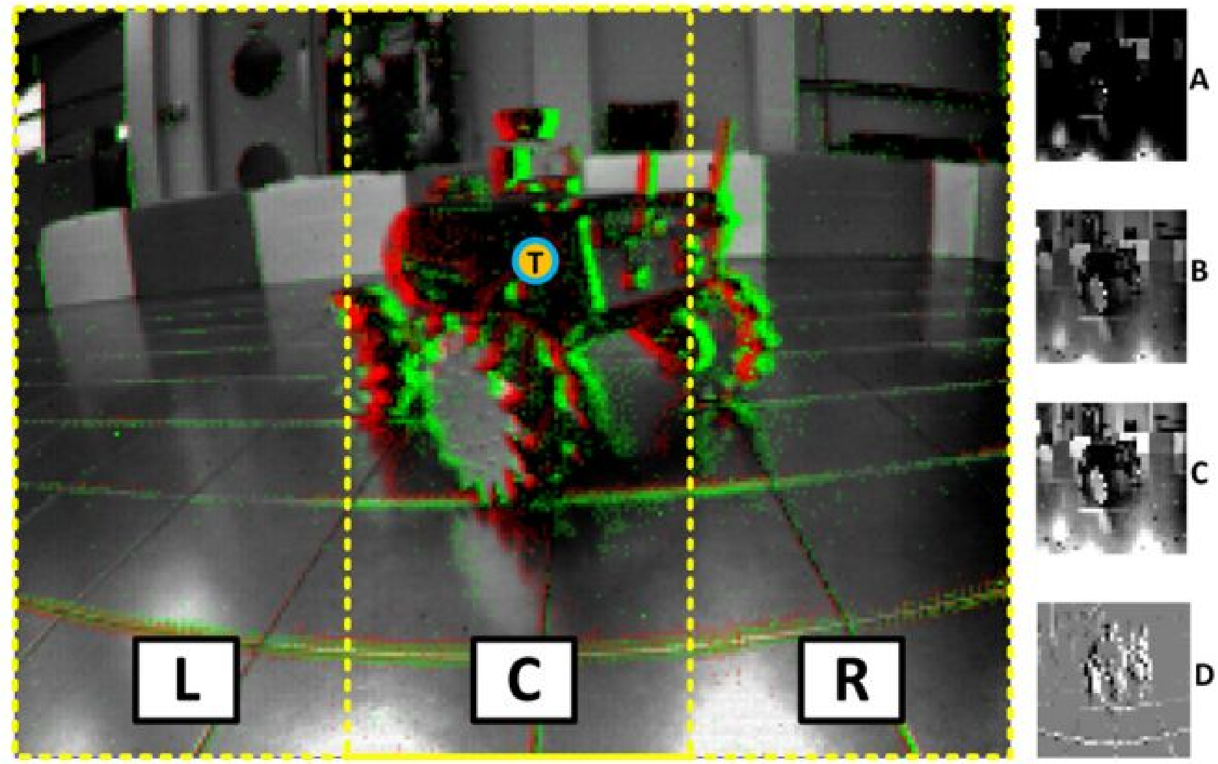

Example of raw recorded data with overlaying APS gray value data and DVS data (red are OFF events, green are ON events). The field of view is divided into the three regions and the target is labelled. A, B and C are the extracted 36 x 36 APS frames with falsely created exposures. D is a subsampled DVS histogram. Credit: arXiv:1606.09433 [cs.RO]

To ward off fears of robots tracking people and killing them, the researchers noted that tracking and following is used in many non-threatening ways, such as children following their parents. They note also that a robot of the future with such abilities might translate to piece of luggage or a grocery cart able to follow a person around so they do not have to lug or steer it.

More information: Steering a Predator Robot using a Mixed Frame/Event-Driven Convolutional Neural Network arXiv:1606.09433 [cs.RO] arxiv.org/abs/1606.09433

Abstract
This paper describes the application of a Convolutional Neural Network (CNN) in the context of a predator/prey scenario. The CNN is trained and run on data from a Dynamic and Active Pixel Sensor (DAVIS) mounted on a Summit XL robot (the predator), which follows another one (the prey). The CNN is driven by both conventional image frames and dynamic vision sensor "frames" that consist of a constant number of DAVIS ON and OFF events. The network is thus "data driven" at a sample rate proportional to the scene activity, so the effective sample rate varies from 15 Hz to 240 Hz depending on the robot speeds. The network generates four outputs: steer right, left, center and non-visible. After off-line training on labeled data, the network is imported on the on-board Summit XL robot which runs jAER and receives steering directions in real time. Successful results on closed-loop trials, with accuracies up to 87% or 92% (depending on evaluation criteria) are reported. Although the proposed approach discards the precise DAVIS event timing, it offers the significant advantage of compatibility with conventional deep learning technology without giving up the advantage of data-driven computing.