Researchers link robots to surveillance teams

April 11 2017, by Bill Steele

Seeing many possible views of the same area from fixed and mobile cameras could be confusing to a human, but a computer can combine it all, track people and objects and notice significant events. Credit: Engineering Communications

If you were monitoring a security camera and saw someone set down a backpack and walk away, you might pay special attention – especially if you had been alerted to watch that particular person. According to Cornell researchers, this might be a job robots could do better than humans, by communicating at the speed of light and sharing images.

The researchers are developing a system to enable teams of robots to share information as they move around and, if necessary, get help in interpreting what they see, enabling them to conduct surveillance as a single entity with many eyes. Beyond surveillance, the new technology might help when teams of robots relieve humans of dangerous jobs like disposing of landmines, cleaning up after a nuclear meltdown or surveying the damage after a flood or hurricane.

"Once you have robots that cooperate you can do all sorts of things," said Kilian Weinberger, associate professor of computer science, who is collaborating on the project with Silvia Ferrari, professor of mechanical and aerospace engineering, and Mark Campbell, the John A. Mellowes '60 Professor in Mechanical Engineering.

Their work, "Convolutional-Features Analysis and Control for Mobile Visual Scene Perception," is supported by a four-year, $1.7 million grant from the U.S. Office of Naval Research (ONR). The researchers will call on their extensive experience with computer vision to match and combine images of the same area from several cameras, identify objects and track objects and people from place to place. The work will require groundbreaking research, Weinberger said, because most prior work in the field has focused on analyzing images from just a single camera as it moves around. And often, Ferrari added, a camera that doesn't move at all. The new system will fuse information from fixed cameras, mobile observers and outside sources.

The mobile observers might include autonomous aircraft and ground

vehicles and perhaps humanoid robots wandering through a crowd. They will send their images to a central control unit, which might also have access to other cameras looking at the region of interest, as well as access to the internet for help in labeling what it sees. What make of car is that? How do you open this container? Identify this person.

The system might notice a significant face, then track that face through the crowd, Weinberger suggested. In earlier work, Ferrari developed what might be described as a robot game of Marco Polo, where a team of hunters track targets through a complex environment.

Knowing the context of a scene, robot observers may detect suspicious actors and activities that might otherwise go unnoticed. A person running may be a common occurrence on a college campus but may require further scrutiny in a secured area.

The core technology is a combination of "deep learning" that lets a computer interpret images, and "Bayesian modeling," which allows it to continuously update its model of the world as new data comes in. The programming will also include a "planning function" to figure out how to obtain additional data that might be needed to resolve an uncertainty. It will help its mobile agents avoid obstacles and, if necessary, direct them to locations where a closer look is needed.

While the Navy might deploy such systems with drone aircraft or other autonomous vehicles, the researchers plan early tests on the Cornell campus, using research robots to "surveil" crowded areas while drawing on an overview from existing webcams. This work might lead to incorporating the new technology into campus security, Ferrari suggested.

Provided by Cornell University