

## **Teaching robots to identify human activities**

July 19 2011, By Bill Steele



To simplify computation, a human image is reduced to a skeleton lof lines and angles. To infer what a person is doing, a computer breaks down activities into a sequence of short actions and compares them with previous observations.

(PhysOrg.com) -- If we someday live in "smart houses" or have personal robots to help around the home and office, they will need to be aware of what humans are doing. You don't remind grandpa to take his arthritis pills if you already saw him taking them -- and robots need the same insight.

So Cornell researchers are programming robots to identify human activities by observation. Their most recent work will be described at the 25th Conference on <u>Artificial Intelligence</u> in San Francisco, in an Aug. 7 workshop on "plan, activity and intent recognition." Ashutosh Saxena, assistant professor of computer science, and his research team report that they have trained a robot to recognize 12 different human activities, including brushing teeth, drinking water, relaxing on a couch and working on a computer. The work is part of Saxena's overall research on



## personal robotics.

Others have tried to teach robots to identify human activities, the researchers note, using video cameras. The Cornell team used a 3-D camera that, they said, greatly improves reliability because it helps separate the human image from background clutter. They used an inexpensive Microsoft Kinect camera, designed to control video games with body movements. The camera combines a video image with infrared ranging to create a "point cloud" with 3-D coordinates of every point in the image. To simplify computation, images of people are reduced to skeletons.

The computer breaks activities into a series of steps. Brushing teeth, for example, can be broken down into squeezing toothpaste, bringing hand to mouth, moving hand up and down and so on. The computer is trained by watching a person perform the activity several times; each time it breaks down what it sees into a chain of sub-activities and stores the result, ending with an average of all the observations.

When it's time to recognize what a person is doing, the computer again breaks down the activity it observes into a chain of sub-activities, then compares that with the various options in its memory. Of course no human will produce the exact same movements every time, so the computer calculates the probability of a match for each stored chain and chooses the most likely one.

In experiments with four different people in five environments, including a kitchen, living room and office, the computer correctly identified one of the 12 specified activities 84 percent of the time when it was observing a person it had trained with, and 64 percent of the time when working with a person it had not seen before. It also was successful at ignoring random activities that didn't fit any of the known patterns.



The researchers note that some people may regard robot monitoring of their activities as an invasion of privacy. One answer, they suggest, is to tell the robot that it can't go into rooms where the door is closed.

Computer code for converting and processing Kinect data is publicly available at <u>pr.cs.cornell.edu/humanactivities/</u>.

Provided by Cornell University

Citation: Teaching robots to identify human activities (2011, July 19) retrieved 10 May 2024 from <u>https://phys.org/news/2011-07-robots-human.html</u>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.