


# 'Hallucinating' robots arrange objects for human use

18 June 2012, By Bill Steele



A robot populates a room with imaginary human stick figures in order to decide where objects should go to suit the needs of humans.

(Phys.org) -- If you hire a robot to help you move into your new apartment, you won't have to send out for pizza. But you will have to give the robot a system for figuring out where things go. The best approach, according to Cornell researchers, is to ask "How will humans use this?"

Researchers in the [Personal Robotics](#) Lab of Ashutosh Saxena, assistant professor of [computer science](#), have already taught robots to identify common objects, pick them up and place them stably in appropriate locations. Now they've added the human element by teaching robots to "hallucinate" where and how humans might stand, sit or work in a room, and place objects in their usual relationship to those imaginary people.

Their work will be reported at the International Symposium on Experimental Robotics, June 21 in Quebec, and the International Conference of Machine Learning, June 29 in Edinburgh, Scotland.

Previous work on robotic placement, the researchers note, has relied on modeling relationships between objects. A keyboard goes in front of a monitor, and a mouse goes next to the

keyboard. But that doesn't help if the robot puts the monitor, [keyboard](#) and mouse at the back of the desk, facing the wall.



Above left, random placing of objects in a scene puts food on the floor, shoes on the desk and a laptop teetering on the top of the fridge. Considering the relationships between objects (upper right) is better, but the laptop is facing away from a potential user and the food higher than most humans would like. Adding human context (lower left) makes things more accessible. Lower right: how an actual robot carried it out. (Personal Robotics Lab)

Relating objects to humans not only avoids such mistakes but also makes computation easier, the researchers said, because each object is described in terms of its relationship to a small set of human poses, rather than to the long list of other objects in a scene. A computer learns these relationships by observing 3-D images of rooms with objects in

them, in which it imagines human figures, placing them in practical relationships with objects and furniture. You don't put a sitting person where there is no chair. You can put a sitting person on top of a bookcase, but there are no objects there for the person to use, so that's ignored. The computer calculates the distance of objects from various parts of the imagined human figures, and notes the orientation of the objects.

Eventually it learns commonalities: There are lots of imaginary people sitting on the sofa facing the TV, and the TV is always facing them. The remote is usually near a human's reaching arm, seldom near a standing person's feet. "It is more important for a robot to figure out how an object is to be used by humans, rather than what the object is. One key achievement in this work is using unlabeled data to figure out how humans use a space," Saxena said.

In a new situation the a robot places human figures in a 3-D image of a room, locating them in relation to objects and furniture already there. "It puts a sample of human poses in the environment, then figures out which ones are relevant and ignores the others," Saxena explained. It decides where new objects should be placed in relation to the human figures, and carries out the action.

The researchers tested their method using images of living rooms, kitchens and offices from the Google 3-D Warehouse, and later, images of local offices and apartments. Finally, they programmed a [robot](#) to carry out the predicted placements in local settings. Volunteers who were not associated with the project rated the placement of each object for correctness on a scale of 1 to 5.

Comparing various algorithms, the researchers found that placements based on human context were more accurate than those based solely in relationships between objects, but the best results of all came from combining human context with object-to-object relationships, with an average score of 4.3. Some tests were done in rooms with furniture and some objects, others in rooms where only a major piece of furniture was present. The object-only method performed significantly worse in the latter case because there was no context to use. "The difference between previous works and

our [human to [object](#)] method was significantly higher in the case of empty rooms," Saxena reported.

The research was supported by a Microsoft Faculty Fellowship and a gift from Google. Marcus Lin, M.Eng. '12, received an Academic Excellence Award from the Department of Computer Science in part for his work on this project.

Provided by Cornell University

APA citation: 'Hallucinating' robots arrange objects for human use (2012, June 18) retrieved 17 April 2021 from <https://phys.org/news/2012-06-hallucinating-robots-human.html>

*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.*