



Ashesh Jain's work with Shikhar Sharma (IIT summer intern), Thorsten Joachims and Ashutosh Saxena on robot checkout clerk featured in media

[HOME](#) < [ABOUT](#) < [NEWS](#)

DATE POSTED: 12/03/2013

Ashesh Jain's work with Shikhar Sharma (IIT summer intern), Thorsten Joachims and Ashutosh Saxena on robot checkout clerk was on FOX News , NBC, French Tribune, Daily Mail (UK), Kurweil AI and several other venues. His algorithm teaches robots to learn how to plan trajectories through incremental online feedback. His YouTube video already has more than 100k views in the last few days.



[NetID Login \(internal use only\)](#)



© 2025 Cornell University

402 Gates Hall, Cornell University, Ithaca, NY 14853



If you have a disability and are having trouble accessing information on this website or need materials in an alternate format, contact web-accessibility@cornell.edu for assistance.

The first StrictlyVC of 2026 hits SF on April 30. Tickets are going fast. [Register now.](#)

Latest Security
Startups AI
Venture Apps
Apple



Events
Podcasts
Newsletters

Advertisement



Cornell Researchers Help Robot Unlearn Stabby Motions With A Human Trainer

John Biggs

6:25 AM PST · November 5, 2013



Advertisement

We humans enjoy not having knives inside of us. [Robots](#) don't know this, three laws be damned. Therefore it's important for humans to explain this information to robots using careful training. Thankfully, the good dudes at Cornell are on the case.

Ashutosh Saxena, assistant professor of computer science, and his team have created a system for fixing robotic motions. In their demo they show the robot lifting a knife from a counter and nearly stabbing a guy. The trainer explains that stabbing is not OK and the robot begins to learn through a process of trial and error that, in the end, ensures minimal stabbage. The system uses trajectory mapping (the robot decides on three potentially un-stabby motions) and the human selects the best one and moves the robot in order to ensure minimal stabbage. From the paper:

Advertisement

[OBJ] Then humans can give corrective feedback. As the robot executes its movements, the operator can intervene, manually guiding the arms to fine-tune the trajectory. The robot has what the researchers call a “zero-G” mode, where the robot’s arms hold their position against gravity but allow the operator to move them. The first correction may not be the best one, but it may be slightly better. The learning algorithm the researchers provided allows the robot to learn incrementally, refining its trajectory a little more each time the human operator makes adjustments or selects a trajectory on the touch screen. Even with weak but incrementally correct feedback from the user, the robot arrives at

Advertisement

way to move a cereal box, but that wouldn't work with a carton of eggs. Also, since eggs are fragile, the robot is taught that they shouldn't be lifted far above the counter. Likewise, the robot learns that sharp objects shouldn't be moved in a wide swing; they are held in close, away from people.

You can read their [entire paper here](#) or simply watch the amazing, non-stabbing robot below. Sadly, when the robots [become TIDWRTWHUFOO](#) we may not be so lucky – or unstabbed.

[youtube=http://youtu.be/uLktpkd7ojA]

Topics: [Hardware](#) [robots](#) [stabbage](#) [stabbing](#) [stabby](#) [TIDWRTWHUFOO](#)

When you purchase through links in our articles, we may earn a small commission. This doesn't affect our editorial independence.



John Biggs

Contributor

`



John Biggs is a writer, consultant, programmer, former East Coast Editor and current contributing writer for TechCrunch. He writes mainly about technology, cryptocurrency,...

[View Bio](#) >

Advertisement

Advertisement

April 30 | San Francisco, CA

StrictlyVC kicks off the year in SF. Register now for unfiltered fireside chats and VC insights with leaders from Uber, Replit, Eclipse, and more. Plus, high-value connections that actually move the needle. Tickets are limited.

REGISTER NOW >

Most Popular



- OpenAI releases GPT-5.5, bringing company one step closer to an AI 'super app'
- Microsoft offers buyout for up to 7% of US employees
- Duolingo is now giving users access to advanced learning content
- Unauthorized group has gained access to Anthropic's exclusive cyber tool Mythos, report claims
- Tim Cook stepping down as Apple CEO, John Ternus taking over
- Blue Origin's New Glenn put a customer satellite in the wrong orbit during its third launch
- Palantir posts mini-manifesto denouncing inclusivity and 'regressive' cultures

Loading the next article



[TechCrunch](#)

[Staff](#)

[Contact Us](#)

[Advertise](#)

[Crunchboard Jobs](#)

[Site Map](#)

[Terms of Service](#)

[Privacy Policy](#)

[RSS Terms of Use](#)

[Tim Cook](#)

[John Ternus](#)

[New Glenn](#)

[WhatsApp](#)

[Mythos](#)

[Tech Layoffs](#)

[ChatGPT](#)

© 2026 TechCrunch Media LLC.

Advertisement



Human touch makes robots smarter: On Learning context-driven user preferences



ashesh jain
126 subscribers

Subscribe

277



Share

Download

Save



141,721 views Nov 1, 2013

More details can be found at <http://pr.cs.cornell.edu/coactive>


This research was done at the Cornell University Personal Robotics Lab by PhD student Ashesh Jain and Prof. Ashutosh Saxena.

In this video, the robot does not know how to move a knife safely and learns through Interactive zero-G and tablet touch user interaction on Baxter for grocery checkout application. With high degree-of-freedom manipulators, one needs to find a trajectory that is not only valid from a geometric point-of-view (i.e., feasible and obstacle-free, the criterion that most existing planners focus on), but also satisfy the user's preferences. In this work, we propose an algorithm to learn such preferences via eliciting online feedback from the user, which does not need to be an optimal demonstrations.

Our Facebook presence: [facebook.com/cornell.robotics](#)

Key moments

View all



0:08

Heatmap without learning Preferable...



0:14

Without Learning User Preferences




0:25

Iteration 1



0:36

Robot Improves



0:40

Top 3 trajectories are shown in simulator



0:55

Iteration 3

Transcript

Follow along using the transcript.

Show transcript



ashesh jain
126 subscribers

Videos

About

Show less



Add a comment...

**@asheshjain** 12 years ago

"Just like any other machine, Robots can be used to improve humans living conditions or they can be misused and abused. It's not a technical choice we are faced with, it's a social one." -- quoting a famous roboticist.

4 Reply

**@GlutesEnjoyer** 12 years ago

we're one step closer to the 3 rules of robotics becoming a reality

Reply

**@mysteryprize** 12 years ago

"Our research laboratory has recently acquired this giant robot arm. After the first few days of seeing how accurately it can throw office stationery at students from the 4th storey lab windows, we decided to try something more practical: teaching it how to hold a knife in a grocery store."

Reply

**@AbundantJD** 12 years ago

This is pretty damn raw if you ask me.

Reply

**@harpodjangorose9696** 12 years ago

Happy birthday, Skynet! :O

Reply

**@Dirtfire** 12 years ago

Now this is awesome. It'll be very useful in coming years.

Reply

**@jackmoore846** 12 years ago

IT'S LEARNING!!

Reply

**@pinochska** 12 years ago

he is just trolling man.. awesome video and what language did you use for the software interface? looks quite nice

Reply

**@asheshjain** 12 years ago

It's a mix two kinds of feedback. One in which the user re-rank trajectories and in the other user corrects the waypoints, we call this feedback zero-G.

Reply

**@grethamaes8712** 12 years ago

Holland got talent

Reply

@osc11771 12 years ago

Interesting...

Reply

@asheshjain 12 years ago

We teach robots how to safely manipulate common objects like knife in presence of humans. To us it is very obvious that knife can hurt humans, but to robots it is not.

Reply

@Lacombe57 12 years ago

User Preference #4.: Have humans working as cashiers; 1. Establishing user preferences not required. 2. Humans retain employment.

Reply

@asheshjain 12 years ago

Good point. Refer to the link in the video description for many more examples and details.

Reply

@harpodjangorose9696 12 years ago

Good, robot. Here's a penny.

Reply



@optimistic12 12 years ago



Sign in

1 Reply

@crimsonxtide50 12 years ago
Is this the beginning of Skynet?

Reply

@nakada1996 12 years ago
They are taking our jobsss xD

Reply

@jarochito73 12 years ago
NEERRDDSSSS!!!

Reply

@JayLofa51 12 years ago
Nice sandals

Reply

Robots can learn to hold knives – and not stab humans



— A technician works with Baxter, an adaptive manufacturing robot created by Rethink Robotics at The Rodon Group manufacturing facility, Tuesday, March 12, 2013, in Hatfield, Pa. There are three jobs open at Rodon Group, a plastic parts manufacturer near Philadelphia. But despite the reports of a shortage of skilled workers nationwide, CEO Michael Araten isn't sweating it. (AP Photo/Matt Rourke)

Matt Rourke / AP file

Nov. 5, 2013, 11:37 PM GMT+5:30

By Nidhi Subbaraman

Humans, most of us anyway, can move a coffee cup without spilling the liquid inside. But such delicate tasks – how to balance a plateful of food, or how to place an egg crate in the fridge

without smashing the contents – are, bizarrely, hard lessons to teach a robot.



Subscribe to read this story ad-free

Get unlimited access to ad-free articles and exclusive content.



A new training technique being perfected at Cornell University will let regular humans easily train household robots to perform delicate or dangerous tasks by simply guiding their arms. Recently, the researchers trained a robot as a checkout clerk, and taught it not to harm a person standing nearby when it rang up a pointy object, a knife.

"Many of the tasks that look simple to us are hard for robots," [Ashuthosh Saxena](#), associate professor of computer science at Cornell University, told NBC News. "Not being clumsy, which most robots are, is a hard problem [to solve]" and a hard thing to teach, he said.

Today, most robots learn their moves from experts. They can be commanded by pre-written programming – which may not be universal for all actions – or trained to copy movements after watching precise expert demonstrations.

HUMAN TOUCH MAKES ROBOTS SMARTER. ON LEARNING CONTEXT-DRIVE
ashesh jain



Watch on

But Saxena's learning technique will let future folks at home train their household helper robot without perfect instructions. When a person moves its arm with the knife shaft pointed away from the person standing in front of it, the bot gets the message. Then robot then


displays movement options on a screen, so the user can continue to refine the command by sight.

"When it arrives at the house, people will be able to give feedback to say: 'I like this, or I don't like this,'" Saxena said. Hey robot! Don't walk in front of the TV when it's on. Don't overturn the plate that's holding the pizza slices. (But it's OK to hold an empty plate upside down!) Safety lessons about how to put away knives, as the research team demonstrated in a recent YouTube video, go on the list as well. The work is due to be presented at the Neural Information Processing Systems conference in Lake Tahoe this December.

"Don't harm the human." Asimov's first law may be simple enough for us to understand, but a hard one to teach robots. Saxena's work could help to change that. Guess it's time to [bring on the salad-chopping robots?](#)

Nidhi Subbaraman writes about technology and science. You can follow her on [Facebook](#), [Twitter](#) and [Google+](#). 

 Nidhi Subbaraman



HUFFPOST

Knife-Wielding Supermarket Checkout Robot Taught Not To Stab Customers

Would You Trust A Knife-Wielding Checkout Robot?

Huffington Post UK

07/11/2013 05:10am GMT



In what could prove to be one of the most important moves ever made by humanity, scientists have taught our soon-to-be-robot-overlords not to stab us.

[Cornell University engineers](#) have modified an assembly line machine and taught it to [scan items at a check out.](#)

But rather than have it flailing around with knives and frozen baguettes they have used a clever for of algorithmic learning to

teach it the best trajectory for [scanning a potentially dangerous item](#) in the presence of people.



An early miss

Baxter, as the checkout assistant of the future is called, is a multi-jointed machine that can move more flexibly than a human.

By giving Baxter feedback about its chosen moves when scanning something it can learn how best to perform the task without spilling an innocent shopper's guts over the conveyer belt.

[**LIKE US ON FACEBOOK**](#) | [**FOLLOW US ON
TWITTER**](#)

Ashutosh Saxena, assistant professor of computer science, said: "We give the robot a lot of flexibility in learning.

"We build on our previous work in teaching robots to plan their actions, then the user can give corrective feedback."

In tests with users who were not part of the research team, most users were able to train the robot successfully on a particular task with just five corrective feedbacks.

The robots also were able to generalize what they learned, adjusting when the object, the environment or both were changed.

The research was supported by the U.S. Army Research Office, a Microsoft Faculty Fellowship and the National Science Foundation.

RELATED

[UK TECH](#)[SUPERMARKET CHECKOUT ROBOT](#)[ROBOTS](#)[KNIFE WIELDING ROBOT](#)[CORNELL UNIVERSITY](#)

HUFFPOST

[ABOUT US](#)[WORK FOR US](#)[STANDARDS AND CORRECTIONS](#)[USER AGREEMENT](#)[PRIVACY POLICY](#)[PRIVACY SETTINGS](#)[CONTACT US](#)[ADVERTISE WITH US](#)[RSS](#)[COMMENT POLICY](#)[CONSENT PREFERENCES](#)

Part of HuffPost Tech. ©2026 BuzzFeed, Inc. All rights reserved.

The Huffington Post

the Kurzweil Library Tracking breakthroughs in tech, science, and world progress. resources

- about + contact
- newsletter
- *the*
Kurzweil
Library
- books by Ray Kurzweil
- writings by Ray Kurzweil
- Ray Kurzweil lifetime *of* innovation
- Ray Kurzweil press + appearances

Attention, Kmart shoppers: there's a robot checker open in aisle six

November 5, 2013

Wow, it's happening already. Cornell robotics researchers are now teaching robots complex tasks such as performing as grocery-store checkout clerks.

Can robot butlers (as in Robot & Frank) be far behind?

This is a big step beyond what Cornell researchers in Prof. Ashutosh Saxena's lab were teaching their smart robots the last time we (virtually) visited the lab in May.

We saw then how a robot could predict a human's action and likely trajectory (a "heat map"), such as opening a refrigerator in anticipation of a human's need for help: putting a dish into the fridge.

How to train a robot as a grocery checkout clerk

Now the robot's tasks have become more complex and nuanced. The Cornell researchers set out to program complex mobile manipulator robots, like Willow Garage PR2 and Rethink Robotics Baxter. These robots have arms with high degrees of freedom, enabling them to perform complex household chores or assembly-line tasks.

But that high level of complexity also creates a gnarly challenge: how do we enable a general user to train the robot's multi-arm movements for different environments, tasks, and user preferences, and




Bad robot! Dodging a knife, heroic Cornell robotics researcher teaches Baxter how to be a checkout clerk (without stabbing the paying customers) by moving the knife to a safe distance.

with movements that are valid from a geometric standpoint (feasible and obstacle-free)?


The Cornell researches developed a novel solution: “The human user does not need to demonstrate optimal trajectories as training data, but merely needs to iteratively provide trajectories that slightly improve over the trajectory currently proposed by the system.

“We argue that this co-active preference feedback can be more easily elicited from the user than demonstrations of optimal trajectories, which are often challenging and non-intuitive to provide on high degrees of freedom manipulators.”


Some examples of their new robots during training:



Baxter checkout robot dazzles with fancy egg-carton moves



Oops! — Baxter holds a knife too close to the flowers
("Cleanup, aisle six")



OK, who gets the seven cups of Quad Grande, Non
Fat, Extra Hot Caramel Macchiato Upside Down?

Human touch makes robots smarter. on Learn

ashesh jain



Watch on

The researchers will present a paper [1] on their impressive research at NIPS 2013.

OK, fine. Now when can I get one of those robot butlers?

Previous virtual visits by *KurzweilAI* to the Cornell Robot Learning Lab

Teaching a *robot* to anticipate human actions

A *robot* that learns how to tidy up after you

Getting ‘hallucinating’ robots to arrange your room for you

More Cornell robot videos

(Image credits: Cornell Robot Learning Lab)

the Kurzweil Library



- [about](#)
- [newsletter](#)
- [privacy](#)
- [Ray Kurzweil Facebook Profile](#)

- [letters + requests for Ray Kurzweil](#)
- [mail @ KurzweilTech.com](#)
- [interviews @ KurzweilTech.com](#)
- [speaking @ KurzweilTech.com](#)
- [writing @ KurzweilTech.com](#)

[sapphi.re](#)

© 2026

The human touch makes robots defter

By Bill Steele

November 6, 2013



(/linkedin)



(/facebook)




(/twitter)



(/email)



(<https://www.addtoany.com/share?url=https%3A%2F%2Fnews.cornell.edu%2Fstories%2F2013%2F11%2Fhuman-touch-makes-robots-defter&title=The%20human%20touch%20makes%20robots%20defter>)




Provided

Graduate students Ian Lenz, left, and Ashesh Jain with the many-jointed Baxter model robot from Rethink Robotics. In Ashutosh Saxena's Personal Robotics Lab the robot is called "Yogi" because all their robots are named for bears, in honor of Cornell's mascot.



Provided

With multiple joints, a Baxter robot can move more flexibly than a human, but it would be hard for a human to decide how best to use those arms, so the robot is programmed to plan its own movements, then allow humans to make corrections.



Provided

As a first step, the robot computes three possible trajectories for moving an object and displays them on a touch screen. After an operator selects one, the robot goes through the motions and the operator can make refinements.

Cornell engineers are helping humans and robots work together to find the best way to do a job, an approach called “coactive learning.”

“We give the robot a lot of flexibility in learning,” said Ashutosh Saxena, assistant professor of computer science. “We build on our previous work in teaching robots to plan their actions, then the user can give corrective feedback.”

Saxena’s research team will report their work at the Neural Information Processing Systems conference in Lake Tahoe, Calif., Dec. 5-8.

Modern industrial robots, like those on automobile assembly lines, have no brains, just memory. An operator programs the robot to move through the desired action; the robot can then repeat the exact same action every time a car goes by.

But off the assembly line, things get complicated: A personal robot working in a home has to handle tomatoes more gently than canned goods. If it needs to pick up and use a sharp kitchen knife, it should be smart enough to keep the blade away from humans.

Saxena’s team, led by Ph.D. student Ashesh Jain, set out to teach a robot to work on a supermarket checkout line, modifying a Baxter robot from Rethink Robotics in Boston, designed for assembly line work. It can be programmed by moving its arms through an action, but also offers a mode where a human can make adjustments while an action is in progress.

The Baxter’s arms have two elbows and a rotating wrist, so it’s not always obvious to a human operator how best to move the arms to accomplish a particular task. So the researchers, drawing on previous work, added programming that lets the robot plan its own motions. It displays three possible trajectories on a touch screen where the operator can select the one that looks best.

Then humans can give corrective feedback. As the robot executes its movements, the operator can intervene, guiding the arms to fine-tune the trajectory. The robot has what the researchers call a “zero-G” mode, where the robot’s arms hold their position against gravity but allow the operator to move them. The first correction may not be the best one, but it may be slightly better. The learning algorithm the researchers provided allows the robot to learn incrementally, refining its trajectory a little more each time the human operator makes adjustments. Even with weak but incrementally correct feedback from the user, the robot arrives at an optimal movement.

The robot learns to associate a particular trajectory with each type of object. A quick flip over might be the fastest way to move a cereal box, but that wouldn’t work with a carton of eggs. Also, since eggs are fragile, the robot is taught that they shouldn’t be lifted far above the counter. Likewise, the robot learns that sharp objects shouldn’t be moved in a wide swing; they are held in close, away from people.

In tests with users who were not part of the research team, most users were able to train the robot successfully on a particular task with just five corrective feedbacks. The robots also were able to generalize what they learned, adjusting when the object, the environment or both were changed.

The research was supported by the U.S. Army Research Office, a Microsoft Faculty Fellowship and the National Science Foundation.

MEDIA CONTACT

Syl Kacapyr

vpk6@cornell.edu (mailto:vpk6@cornell.edu)

📞 607-255-7701 (tel:607-255-7701)

RELATED STORIES

Robots learn to recognize objects (March 2013)

(<http://www.news.cornell.edu/stories/2013/03/teachable-moments-robots-learn-our-humanistic-ways>)

Get Cornell news delivered right to your inbox.

SUBSCRIBE (/ABOUT/SUBSCRIBE-NEWSLETTER)

Researchers Are Finally Teaching Robots To Be Less Stabby

BY ANDREW LISZEWSKI

PUBLISHED NOVEMBER 5, 2013, 1:40 PM ET

READING TIME 1 MINUTE



READ LATER



COMMENTS (0)

Most of the news coming out of robotics research has us really worried about mankind's future, but Cornell University finally brings

us a glimmer of hope. Researchers there are working on developing an algorithm through physical feedback that will teach robots to be more careful with certain objects—like say when handling a sharp knife around highly stabbable humans.

The robot is free to first perform movements it assumes are ok, but a human operator can intervene at any point with the push of a button and then adjust the bot's trajectory and movements to something a little safer. So, for example, if it grabs a knife and is unaware the long blade can cause significant damage, the human operator can move the robot's arm to a safer position, and over time the algorithm will learn to be more careful when handling sharp objects.

The research isn't just about protecting humans, though. A robot could be taught to handle fragile items like eggs or LCD screens with greater care, making them more effective in a factory setting. But on the flip side to all the good this research could do, is it really such a great idea to reinforce the idea for robots that knives and humans don't mix? [[Cornell University](#) via [TechCrunch](#)]

EXPLORE MORE ON THESE TOPICS

ROBOTS

SHARE THIS STORY



Sign up for our newsletters

Subscribe and interact with our community, get up to date with our customised Newsletters and much more.

GIZMODO

IO9

EARTHER

THE NEXT INTERFACE

EMAIL ADDRESS

SUBSCRIBE NOW

Related Articles