

THE SCIENCE OF EVERYTHING

COSMOS

14 June 2021/
Mark Pesce

The surprisingly complicated technology that goes into picking winners

Choosing and gathering an apple is simple for humans, but it's a task that has frustrated robotics researchers for decades. Now, thanks to a *Matrix* simulation-style training regime, it seems it's finally within reach.



Credit: Ambi Robotics.

Can robots be trained to pick up after us? A recent breakthrough in machine learning could point the way toward a much more tactile future for these smartest of machines.

In April the ABC [carried a report](#) about a new machine put to work in the apple orchards of Tasmania, helping pickers – of whom there are far fewer, because of pandemic-closed borders – to shoulder the burden. Instead of bearing a heavy sack of freshly picked fruit, the machinery offers a conveyor belt from picker’s hand to storage bin. That seems a small thing, until you consider what it means to be freed of 10 or 15 kilos of dead weight every time they reach up to pick another apple. Even that little bit of help makes the job easier.

And while that bit of kit marks a great improvement, why haven’t we seen a piece of farm equipment capable of picking apples directly from the tree? The answer is simple – picking isn’t as easy as it looks.

Touch is perhaps our most frequently overlooked sense (because it so infrequently fails us) and it makes picking possible. You can hold a cricket ball, or an egg, or an apple, and know – instinctively – just the right amount of pressure to apply to it, to keep it steady in your hand. A broad set of tactile and [proprioceptive](#) senses (a ‘sixth sense’ that allows us to

sense our muscles and joint positions) feed that awareness, and if you should lose even a part of them, you would find yourself dropping the ball, crushing the egg, or bruising that apple. That back-and-forth between touch and muscle contraction gives us a grip that's as strong or delicate as each situation requires, even without prior experience: you can pick up something you've never held before, and not make a complete muddle of it.

All that “embedded” cognition has always been far beyond the capacities of any robots. For decades they've been equipped with claws – or tools like spanners and drills – but have always operated on strictly programmatic lines: move the tool to precisely such-and-such a position, operate the drill for exactly this period of time, then withdraw, wait a few seconds, and do it all over again. That works quite well on an automotive assembly line, but in less predictable situations, such as a full-grown and fully laden apple tree, that programmatic capacity doesn't offer anything useful.

Welcome to the ‘picking problem’, one of the most interesting – and most difficult – in the century-long story of robotics. It's a problem that has long obsessed [Dr Ken Goldberg](#), chair of the robotics program at the University of California, Berkeley. [Full disclosure – I've been a friend of Ken's for 25 years.] Goldberg has trained an entire generation of roboticists (and [co-written a children's book](#) with his daughter) and along the way he's given them a touch of his own obsession for picky robots. As a result, we're at a watershed moment in the history of the field, a moment when robots finally get ‘good enough’ at picking to be useful.

Just as with humans, for a robot, picking presents a “sensory integration” problem. Try picking up a few unknown objects in the dark and you'll quickly realise that sight plays as important a part in picking as does touch or proprioception. You have to be able to see the apple in order to reach out and pick it off the tree. That means robots need good cameras but, more importantly, they need good image recognition systems, so they can “understand” what their cameras look at: flower, branch, or fruit.

Over the last decade, the confluence of very powerful computer hardware and sophisticated, deep-learning algorithms have led to a revolution in

image recognition: show a computer a few million images of an object, from a few thousand different angles, and it will probably be able to recognise it in the wild. (Human beings will only need to see the same object presented once or twice to be able to recognise it again. Such are the gulfs that remain between designed and evolved capabilities.)

Once the robot can recognise an object, it has to have a go at reaching out and picking it up. A robot learns how to do that just as we do as toddlers – by trying and failing. The robot will fail and fail and fail and fail thousands of times, before it learns enough from its failures to do it right.

You might reckon that sort of extended learning cycle would be sufficient to give a robot the nous to be able to denude an orchard of its apples, but the organic nature of the tree itself presents more complications than any robot can handle. Can the robot work out a path through the branches to reach the apple? How hard does the robot need to pull to release the apple? Does it need to add a gentle twist to snap it off? And how much pressure should it apply? These complexities, Goldberg argues, place a robot apple picker toward the pinnacle of solutions of the picking problem, rather than a low-hanging fruit.

Goldberg and his graduate students recently revealed a solution that greatly speeds the training of picking robots: put the robot inside a simulation. Like some weird inversion of *The Matrix*, Goldberg's research team sent their machines into the artificial world of [Dex-Net 4.0](#), and, within that environment, where they controlled every element, they could run the robot through simulations far faster than would be possible in the real world. *Ten thousand* times faster. Mistakes that take hours or even days to learn to avoid in the real world could be overcome in minutes or seconds. And since every single object a robot has to pick has to be separately trained for, that big difference in training times quickly adds up.

All of this went into the design of a new robot created by a startup, [Ambi Robotics](#), co-founded by Goldberg and some of his students. Designed to automate the difficult “pick and place” tasks required to fill boxes for shipment, or necessary to sort through a container of random objects, the firm's robots promise a “picking rate” far faster than existing robots, although not yet quite the equal of humans' evolved capabilities.

Ambi, along with competitors such as Boston Dynamics (makers of the creepy “digidog” and [new “Stretch” robot](#)), have their eyes firmly fixed on a big prize – the mushrooming potential of online order fulfilment. Anyone who’s seen *Nomadland* has had an insight into the life of Amazon’s human pick-and-place army, gathering up all the pieces for a customer’s parcel. Solve that problem, worth billions of dollars, and then, with a robot that can grip and learn, you can go on to grasp the world.

Half a decade ago, [Goldberg pointed](#) to the utter failure of domestic robotics as a sign of how much work the field had left ahead of it. “You can’t even get a robot to clear the table after dinner,” he assessed. “Something we don’t even need to think about, but it’s still well beyond the capacity of any robot – at any price.” Once they leave the warehouse, this new generation of picky robots could soon be picking up after us.

<https://cosmosmagazine.com/technology/robotics/the-surprisingly-complicated-technology-that-goes-into-picking-winners/>