

Machine Learning reaches into the unknown

Originally published June 27, 2017 at <https://www.aiaa.net/artificial-intelligence/articles/machine-learning-reaches-into-the-unknown>



Teaching a machine to pick up objects it knows nothing about is quite a challenge in artificial intelligence research

The act of engaging with physical objects in a safe and non-damaging manner is a critical study in the field of Machine Learning. But the functional leap from robots with prior knowledge of their environment to AI-driven systems — capable of mapping and adapting to new objects and situations — represents a genuine evolutionary step away from the 'Henry Ford' model of production-line automation; a step towards authentic machine agency.

Yet the field of robot grasping is still haunted by templates and the need for 'fixed' conditions. Amazon Robotics is [investing significantly](#) in AI systems capable of negotiating cluttered warehouses and identifying objects which conform to pre-set 3D meshes held by the robots in a database. Meantime Google's Sergey Levine has led [elaborate research](#) to teach industrial robots how to cope with unexpected or cluttered assembly environments.

There's a lot at stake, not least the challenge of overcoming a fairly natural fear that robot hands may have incorrect understanding of the pressure and attitude-control needed to move different types of objects from one place to another. Unsurprisingly it's a fear that has often manifested in popular media.




For major manufacturers and for goods companies such as Amazon, the development of intelligent automated grasping represents an industrial lodestone on the road to genuine AI-driven services - not to mention a significant opportunity to downsize in terms of human contributors to the process.

Mobile dexterity

There's a symbiotic relationship between the well-funded industry concerns currently engaged with the problem, and medical and cybernetic research, which is seeking to utilise Machine Learning to create new AI-resourced dexterity solutions for people with impaired mobility.

At the College Of Computer and Information Science in Boston's Northeastern University, a research team — in collaboration with UMass Lowell — is developing a self-powered mobility system which uses lasers and neural networks to help identify, grasp and move 'unknown' objects for disabled users.



The project, entitled Open World Assistive Grasping Using Laser Selection, outfits a [Golden Avenger mobility scooter](#) with a versatile robotic arm taken from a Rethink Robotics [Baxter robot](#).

The team has developed two interfaces for the laser-guided grasping system; the first is aimed at users with limb tremors, which are dampened by the use of a foam ball to aim the laser, indicating to the arm's AI the area containing the object to be grasped.



The second interface (above, right) addresses users with far more severe upper body impairments, and uses a servo arrangement which could potentially be controlled by [sip-and-puff](#) (or 'single-switch') scanning. The versatility of the second interface permits other forms of control, including joystick or mouse (or, in theory, more ambitious leaps of mobility, such as [thought](#)).

"One of the major objectives of our research is to be able to deal with 'novel objects'," says project team member Marcus Gualtieri "which means we do not store a 3D mesh of the objects the robot will grasp. This is important in household and other human environments where there are a lot of objects."

"It would be tedious," he continues "to create a mesh model for each object. We hope the system can be more plug-and-play, where the non-expert operator does not need to do a lot of work up-front creating meshes of objects."

The prototype mobility unit uses a Machine Learning system which provides input to the robot arm from a separate but untethered laptop. The laptop can operate for about 90 minutes on its own battery, or can draw power from the Golden Avenger's own 24v batteries via an off-the-shelf 1000w pure-sine AC inverter. The longest self-powered run the unit has made was five hours, though this used up less than half of the battery capacity for the total rig.

Getting a grip

Despite the project's stricture regarding prior knowledge, basic methodology and procedures for learning had to be established, with the team using the [BigBIRD](#) (Big Berkeley Instance Recognition Dataset) large-scale 3D database of object instances, which contains sensor data on more than 125 objects.

"It is important for us," says Gualtieri, whose contribution to the project is mainly around motion planning and sensing components, "that the dataset be generated from sensors like the ones we will actually use, not synthetic sensor data. Of course, having a larger dataset and a larger network could help improve grasp detection accuracy at the cost of additional computational time required for training."

As with most Machine Learning hardware-based projects, OWAGULS has to balance accuracy against latency. Power consumption is also a practical consideration for an independent mobility unit which has to burn through compute cycles. But ultimately the effective independence of the Baxter arm grasps must remain the priority.

Unknown quantities

Anyone who has ever miscalculated the weight of a suitcase or bag (either wrenching arms or losing one's balance because it's empty) has experienced at least one of the many challenges an AI system faces in handling an object about which it has limited information.

Is it heavy? Is it *too* heavy? Can it be handled loosely (a sealed bottle) or should its orientation be maintained (a cup of hot liquid)? Is it, as it transpires, attached to anything? And if it is, should the 'keep pulling until maximum lift' approach be cancelled because of the knock-on consequences?

"Our system," says Gualtieri. "does not estimate an object's center of mass or other inertial properties. Knowing these properties would be helpful for reliable grasping of heavy objects. The challenge is perceptual."



"From a camera or depth image, the robot would have to estimate the inertial characteristics of an object it is only seeing for the first time. And visual sensing may not be enough; the robot wouldn't be able to tell how full the shampoo bottle is, for example, if the bottle is opaque. So the robot may need to adjust its estimate of the object's inertial properties based on feedback from force or tactile sensors."

Networked knowledge

The project's aim of robotic independence from excessive training is ambitious. Could a proposed system of Baxter-equipped Golden Avengers not at least form its own private learning network, benefitting from the same parameters for learned data? According to Gualtieri, it's a balancing act between automation and autonomy:

"Having a lot of prior information doesn't necessarily reduce the system's versatility," he says "as long as the system is not 'overfit' to the prior information."

"For example, if the system has access to a huge database of shampoo bottles, it can still be versatile if it is able to generalize to a new shampoo bottle, which is likely to be similar to but not exactly the same as several instances in its database.

"So the key is how the prior information is used. Deep neural networks seem to be the technology of choice today for converting a large database into a function that can generalize to some extent."

Gualtieri says that the project does have an interest in Reinforcement Learning (RL). "If the robot is to automatically adapt to its environment dynamically, to improve itself over time, RL is probably the right framework under which to consider this problem."

Counting the cost

Cost is another critical consideration in the viability of the system. The Baxter robot has become a standard initial purchase in AI-driven robotics projects, and technically one could obtain two serviceable arms from a single purchase (around \$25,000). However, Rethink Robotics requires a second control computer to be purchased to obtain individual arm control, with some software adjustments necessary to disassociate a two-armed system into a single limb.

Gualtieri admits that despite Baxter's ease-of-use, bespoke or more compact hardware would be a help in achieving a more lightweight robotic appendage. "Baxter is a great step in this direction, but the arm seems kind of large on the front of the scooter, especially when driving it through doorways. It would be nice if the arm were small and still had good reach. It would also be nice if the gripper had a wider opening and closing range."



The system's visual unit employs [OpenRAVE](#) for collision detection and Oxford University's [Infinitam](#) open source framework for depth tracking, enabled by a Kinect-style RGB-D stereoscopic camera. This part of the system too could be optimised, says Gualtieri, and not only in terms of cost.

"Currently Kinect-like sensors don't see black and transparent objects well," he observes. "Also, the field of view is limited for tasks like obstacle avoidance."



Gualtieri says that it would currently be 'risky' to test the system's grasping mechanism on ceramic and other breakable objects, observing that "it may not be safe to say it is a solved problem."

Steps towards the market

At the time of writing the Baxter-assisted unit is reportedly scheduled for tests with the actual target user group soon. "But I am not sure of the exact time. To get there, we need to make failures - like the arm colliding with the shelf and the arm failing to find a motion plan - to make them relatively rare. To have better collision avoidance, we need to add more sensors. We have already sped up the grasping process quite a bit."

The outlook for release to market is a challenge, even with the other problems overcome. The project represents a system capable of adjusting to unplanned environments, where its impossible that all the variables will have been tested in advance.

"I suspect that the scooter system will at first be limited to confined environments," Gualtieri predicts, "such as a gymnasium at an assisted living residence. Once the technology matures and people become more comfortable with it, then it can move to parks, pedestrian walkways, homes, and eventually grocery stores."

"The technology will most likely branch out to other application areas, such as manipulating rocks on Mars or assisting firefighters or other hazardous environments, where it can safely mature and gain people's confidence."

Paper: [Open World Assistive Grasping Using Laser Selection](#) (PDF)

IMAGES: [Arxiv](#) / YouTube