


Helping robots remember: Hyperdimensional computing theory could change the way AI works

May 16 2019, by Rebecca Copeland



The researchers' lab, as seen by the dynamic vision sensor. Credit: Perception and Robotics Group, University of Maryland.

The Houston Astros' José Altuve steps up to the plate on a 3-2 count, studies the pitcher and the situation, gets the go-ahead from third base, tracks the ball's release, swings ... and gets a single up the middle. Just another trip to the plate for the three-time American League batting champion.

Could a robot get a hit in the same situation? Not likely.

Altuve has honed natural reflexes, years of experience, knowledge of the pitcher's tendencies, and an understanding of the trajectories of various pitches. What he sees, hears, and feels seamlessly combines with his brain and [muscle memory](#) to time the swing that produces the hit. The robot, on the other hand, needs to use a linkage system to slowly coordinate data from its sensors with its motor capabilities. And it can't remember a thing. Strike three!

But there may be hope for the robot. A paper by University of Maryland researchers just published in the journal *Science Robotics* introduces a new way of combining perception and motor commands using the so-called hyperdimensional computing theory, which could fundamentally alter and improve the basic artificial intelligence (AI) task of sensorimotor representation—how agents like robots translate what they sense into what they do.

"Learning Sensorimotor Control with Neuromorphic Sensors: Toward Hyperdimensional Active Perception" was written by computer science Ph.D. students Anton Mitrokhin and Peter Sutor, Jr.; Cornelia Fermüller, an associate research scientist with the University of Maryland Institute for Advanced Computer Studies; and Computer Science Professor Yiannis Aloimonos. Mitrokhin and Sutor are advised by Aloimonos.


Integration is the most important challenge facing the robotics field. A robot's sensors and the actuators that move it are separate systems,

linked together by a central learning mechanism that infers a needed action given [sensor data](#), or vice versa.

The cumbersome three-part AI system—each part speaking its own language—is a slow way to get robots to accomplish sensorimotor tasks. The next step in robotics will be to integrate a robot's perceptions with its motor capabilities. This fusion, known as "active perception," would provide a more efficient and faster way for the robot to complete tasks.

In the authors' new computing theory, a robot's operating system would be based on hyperdimensional binary vectors (HBVs), which exist in a sparse and extremely high-dimensional space. HBVs can represent disparate discrete things—for example, a single image, a concept, a sound or an instruction; sequences made up of discrete things; and groupings of discrete things and sequences. They can account for all these types of information in a meaningfully constructed way, binding each modality together in long vectors of 1s and 0s with equal dimension. In this system, action possibilities, sensory input and other information occupy the same space, are in the same language, and are fused, creating a kind of memory for the robot.

The *Science Robotics* paper marks the first time that perception and action have been integrated.



Hyperdimensional pipeline. From the event data (b) recorded on the DVS during drone flight (a), “event images” (c) and 3D motion vectors (d) are computed, and both are encoded as binary vectors and combined in memory via special vector operations (e). Given a new event image (f), the associated 3D motion can be recalled from memory. Credit: Perception and Robotics Group, University of Maryland.

A hyperdimensional framework can turn any sequence of "instants" into new HBVs, and group existing HBVs together, all in the same vector length. This is a natural way to create semantically significant and informed "memories." The encoding of more and more information in turn leads to "history" vectors and the ability to remember. Signals become vectors, indexing translates to memory, and learning happens through clustering.

The robot's memories of what it has sensed and done in the past could lead it to expect future perception and influence its future actions. This active perception would enable the robot to become more autonomous and better able to complete tasks.

"An active perceiver knows why it wishes to sense, then chooses what to perceive, and determines how, when and where to achieve the perception," says Aloimonos. "It selects and fixates on scenes, moments in time, and episodes. Then it aligns its mechanisms, sensors, and other components to act on what it wants to see, and selects viewpoints from which to best capture what it intends."

"Our hyperdimensional framework can address each of these goals."

Applications of the Maryland research could extend far beyond robotics. The ultimate goal is to be able to do AI itself in a fundamentally different way: from concepts to signals to language. Hyperdimensional computing could provide a faster and more efficient alternative model to the iterative neural net and deep learning AI methods currently used in computing applications such as data mining, visual recognition and translating images to text.

"Neural network-based AI methods are big and slow, because they are not able to remember," says Mitrokhin. "Our hyperdimensional theory method can create memories, which will require a lot less computation, and should make such tasks much faster and more efficient."

Better motion sensing is one of the most important improvements needed to integrate a [robot](#)'s sensing with its actions. Using a dynamic vision sensor (DVS) instead of conventional cameras for this task has been a key component of testing the hyperdimensional computing theory.

Digital cameras and computer vision techniques capture scenes based on pixels and intensities in frames that only exist "in the moment." They do not represent motion well because motion is a continuous entity.

A DVS operates differently. It does not "take pictures" in the usual sense, but shows a different construction of reality that is suited to the purposes of robots that need to address motion. It captures the idea of seeing motion, particularly the edges of objects as they move. Also known as a "silicon retina," this sensor inspired by mammalian vision asynchronously records the changes of lighting occurring at every DVS pixel. The sensor accommodates a large range of lighting conditions, from dark to bright, and can resolve very fast motion at low latency—ideal properties for real-time applications in robotics, such as autonomous navigation. The data it accumulates are much better suited to the integrated environment of the hyperdimensional computing theory.

A DVS records a continuous stream of events, where an event is generated when an individual pixel detects a certain predefined change in the logarithm of the light intensity. This is accomplished by analog circuitry that is integrated on each pixel, and every event is reported with its pixel location and microsecond accuracy timestamp.

"The data from this sensor, the event clouds, are much sparser than sequences of images," says Cornelia Fermüller, one of the authors of the Science Robotics paper. "Furthermore, the event clouds contain the essential information for encoding space and motion, conceptually the contours in the scene and their movement."

Slices of event clouds are encoded as binary vectors. This makes the DVS a good tool for implementing the theory of hyperdimensional computing for fusing perception with motor abilities.

A DVS sees sparse events in time, providing dense information about changes in a scene, and allowing for accurate, fast and sparse perception of the dynamic aspects of the world. It is an asynchronous differential sensor where each pixel acts as a completely independent circuit that tracks the intensity changes of light. When detecting motion is really the kind of vision that is needed, the DVS is the tool of choice.

More information: A. Mitrokhin et al, Learning sensorimotor control with neuromorphic sensors: Toward hyperdimensional active perception, *Science Robotics* (2019). [DOI: 10.1126/scirobotics.aaw6736](https://doi.org/10.1126/scirobotics.aaw6736)

Provided by University of Maryland

Citation: Helping robots remember: Hyperdimensional computing theory could change the way AI works (2019, May 16) retrieved 17 April 2024 from <https://techxplore.com/news/2019-05-robots-hyperdimensional-theory-ai.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.