Antenna Company announced the availability of high performance antenna technologies for Wi-Fi, GPS and Cellular networks.
See page > 7

Clippard has announced a series of direct actuating valves that offer a fast response time for accurate dosing of minute volumes.
See page > 23

PRODUCT DESIGN & DEVELOPMENT®

COMPONENT, SYSTEM A… FOR DESIGN ENGINEERS

WWW.PDDNET.COM

SEPTEMBER 2016

71ST YEAR, ISSUE 6

Advantage Business Media

A.I. comes to consumer electronics products

MORE THAN **25** NEW PRODUCTS

Artificial Intelligence Invades Consumer Electronics

By Lee Goldberg, Megan Crouse, and Alianna Maren, Ph.D.

While even experts can't all agree on a precise definition for "artificial intelligence" (AI), the technology behind it continues to become more powerful, less expensive, and easier to implement. As a result, we're seeing AI capabilities beginning to appear in many types of consumer electronic products. In this article, we'll look at some of the ways designers are coaxing their designs across the fuzzy chasm between "smart" and "intelligent," along with a few notable examples of AI-enabled products.

Defining AI

The origins of modern AI theory date back to the early 1940's when Alan Turing and Alonzo Church postulated that digital computers might be capable of simulating any process of formal reasoning. Their pioneering work was in good part the inspiration for the group of researchers who gathered at Dartmouth University in the summer of 1956 to define and solve a set of basic problems (a.k.a. goals) which continue to define modern AI research. These include *reasoning, knowledge representation, planning, learning, natural language processing, perception, and social intelligence.*

Solving these problems has required the development of new concepts and tools to navigate the new problem space. This has given rise to symbolic reasoning, expert systems, neural networks, and fuzzy logic; fields that made possible the first practical applications of AI in the 1980's.

AI Challenges

At that time, however, expert systems suffered from serious deficiencies. In attempting to codify how human experts identified key distinctions (ranging from diagnosing medical symptoms to space shuttle communications), they had become too rigid. Human experts think in terms of patterns and weighted combinations of factors, more than in terms of rules.

Another big challenge that had emerged by the early 1990's was that many of the crucial problems required dealing with lots of little pieces of low-level data, rather than a few chunks of largely symbolic data. For example, image understanding (a field related to AI), requires processing huge volumes of pixel-level data, using various algorithms to extract the lines and regions in an image. The AI-based challenge is to connect these lines and regions to known and identifiable objects, such as people, landmarks, and activities or interactions.

Deep Learning Allows AI to Move Ahead

In order to overcome these limitations, those early AI systems had to develop the ability to deal with more ambiguous data sets and to perform some level of self-learning. These challenges made it difficult to field practical AI systems, until the recent advances made possible by the development of *deep learning*. Jesse Clayton, product manager, Autonomous Machines at NVIDIA, explains "Deep learning is a neural-network based approach to machine learning. What makes it fundamentally different than previous

The Moorebot voice-interactive robot assistant with a warm, funny digital personality. Its AI capabilities allow it to interpret and respond to human speech, as well as infer activity and emotion from the images its camera captures. Photo: Courtesy of Moorebot.

approaches is that it makes it practical to train systems using very large datasets. Training on large datasets results in better accuracy for tasks like image recognition, natural language processing, medical diagnosis, game playing, and many others."

Deep learning technologies have enabled a key breakthrough in AI technology - the synthesis of multiple ways of describing something. When applied to image processing, this means first identifying the line and region segments, and then connecting "like-with-like," using algorithms that are loosely based on the workings of the brain's visual cortex. Once provisional shapes are assembled, algorithms call upon higher-level symbolic knowledge to suggest possible shape interpretations. Matching the shapes against possible interpretations ultimately yields an interpreted image. This capability is essential, for example, to Google's self-driving car, as well as to robots that can move around houses and factory floors.

Thanks to these developments, by the late 1990s, AI was being commonly used for many applications such as data mining, financial analysis, logistics management, and medical diagnosis. At the time, it was expected that AI would eventually be applied to nearly any application which involved "big data".

AI Gets Personal

What we did not expect was what is occurring now – that AI is becoming both hidden and pervasive, popping up in unexpected places like our smartphones and fitness devices. This emergence of consumer-level, functionally-useful, and adaptive AIs has been made possible by the confluence of several key factors; (1) substantial evolutions in core AI algorithms, (2) powerful new insights in machine learning theory, (3) advances in low-cost, high-throughput processing chips, and (4) the prevalence of cloud-based computing resources.

Cloud-based computing, in particular, made many of the first consumer AI applications practical and affordable by eliminating the need for the host device to have a processor powerful enough to support all the compute-intensive AI functions itself. Smart Phones, for example, can now understand your natural way of speaking and react intelligently, thanks to digital assistants like Siri or Google Assistant, both of which are powered by deep learning algorithms running in the cloud.

Another more recent application of cloud-based AI is the Moorebot, a robot assistant that's designed to be cute as well as helpful. The affable one-eyed device can perform customer service and shop inventory in a business setting, read emails, or give reminders and play videos at home. Moorebot reaches into the cloud for voice recognition services and information that requires an internet connection, such as weather and news. However, despite the fact that Moorebot is based on AI technologies such as facial and voice recognition, the company's CEO Jun Ye said that the makers intentionally avoided using the term "artificial intelligence" in marketing material. "We think some regular consumers start to hate artificial intelligence," Ye said.

New Hardware Essential to Embedded AI

The June Intelligent Oven. Photo courtesy of June Life Inc.

Although many AI-enhanced consumer devices have much of their processing down on the cloud, it is not always the best solution - especially for real-time, or safety-critical applications. Self-driving cars, for example, need to have their deep learning functionality residing onboard in order to function safely even when network connections may be unreliable, or there's lots of

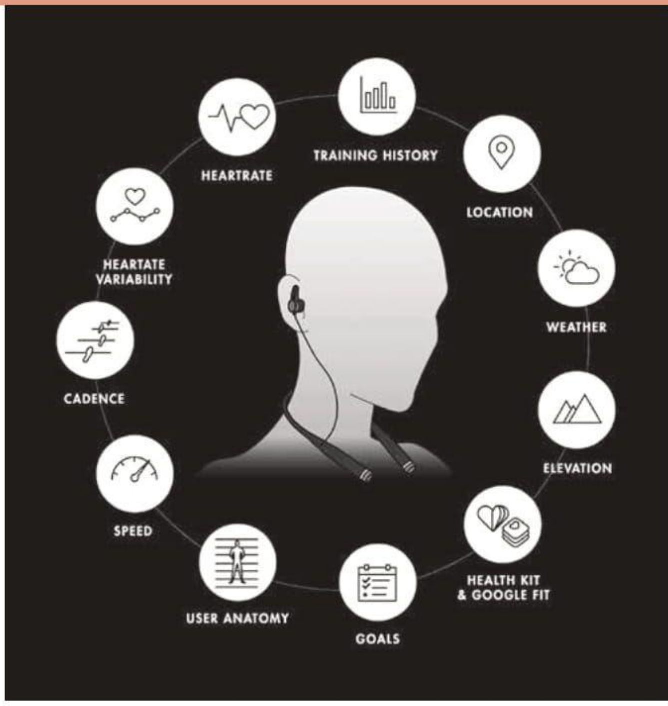

LifeBeam's VI intelligent fitness monitor. Image courtesy of LifeBeam

latency in the connection.

In these applications, specialized multi-core processors known as graphics processing units (GPUs) have been essential to efficiently supporting locally-hosted deep learning. NVIDIA's Clayton describes GPUs as "highly parallelized architectures that can run those algorithms fast, making it possible to train very large data sets in hours or days, instead of weeks or months." For example, NVIDIA's Jetson TX1 is a module designed specifically for deep learning for embedded systems. It incorporates a powerful GPU, and only uses 10W, and it's smaller than a credit card. The Jetson TX1, with its multi-GPU designed specifically for embedded applications, has been a key enabler for new devices.

One of the recently-released consumer devices powered by NVIDIA's Jetson TX1 chip is the Horus wearable assistant for the blind. It uses a small head-mounted camera to gather imagery and then translates what it is seeing to an audio description. For instance, it can enable a blind mother to read a written book to her child.

The TX1 chip also powers the June Intelligent Oven, which uses computer vision and deep learning to help you optimally prepare your food. It has a variety of sensors and it's been trained to determine exactly the right level of doneness for all sorts of recipes. It can even tell what kind of food you're putting in just by looking at it.

Partitioning between the Device and the Cloud

Although the Horus assistant and the June Oven keep their intelligence entirely locally-resident, it's not a cost-effective approach for many types of products where price is a primary design driver. As a result, most of the next generation of consumer-facing AI systems will likely partition various percentages of their large-scale algorithms and deep learning processes between on-board and cloud-based platforms, while keeping user-specific and identifying information device-local.

This approach is essential not only for personal assistant and household management and utility devices, but also for the emerging realm of intelligent wearable devices. One such device is the Vi, an AI-based personal trainer developed by LifeBeam which can respond to voice prompts and offer "coaching" advice based on its analysis of the user's performance and vital signs.

A Lifebeam's intelligence resides partially on the headset device itself and partially on the connected phone app. A company representative divides VI's artificial intelligence according to two different aspects. One is a set of rules based on pre-defined data inputs, through which the user's behavior is filtered. The second is a more adaptable data collection and pattern recognition capability, and this is the part that actually learns. Eventually, after it gets to know the user's behavior, it can suggest activities and audio or come to conclusions based on that behavior.

Regardless of how next-generation intelligent consumer devices will partition their processing between local and cloud-based platforms, the one aspect of the future is clear: intelligent devices will become more ubiquitous and will find roles in many aspects of our lives, both individually and at as a society.

PDD

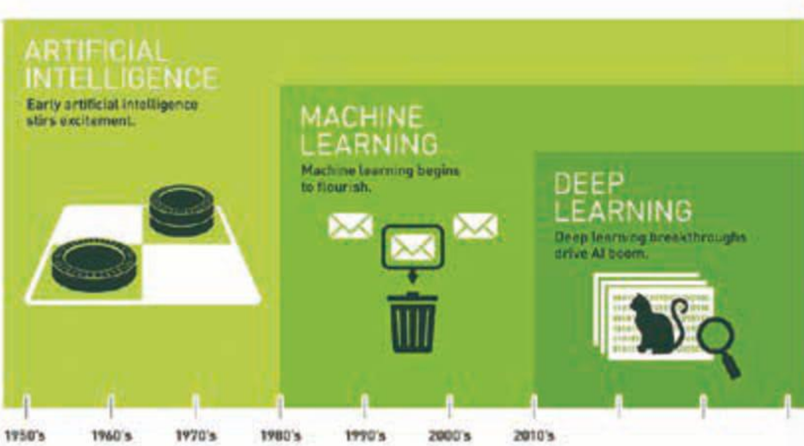

Since an early flush of optimism in the 1950s, smaller subsets of artificial intelligence – first machine learning, then deep learning, a subset of machine learning – have created ever larger disruptions.

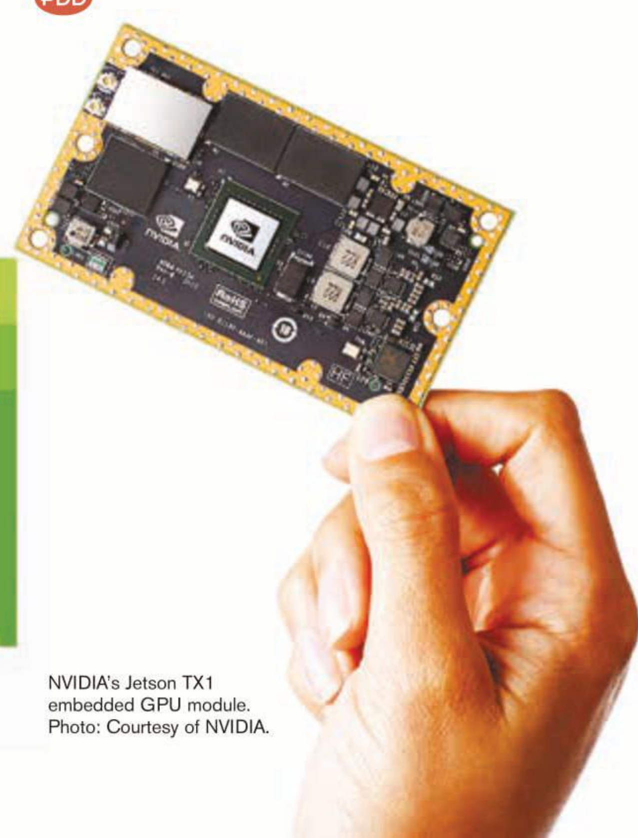
NVIDIA's Jetson TX1 embedded GPU module. Photo: Courtesy of NVIDIA.