

AI Training Platform

Validated with NVIDIA® Tesla®
P100/V100 accelerators



Key Features

- 8x PCIe x16 Gen3 GPGPU slots (300W/slot)
- Validated with NVIDIA® Tesla® P100/V100 accelerators
- Dual Intel® Xeon® Scalable processors
- 24x DDR4-2666 RDIMM
- 8x SATA 6Gb/s hot-swappable 2.5" drives

Headlines

- Microsoft acquires Lobe to help bring AI development capability to everyone
- New NVIDIA Data Center Inference Platform to Fuel Next Wave of AI-Powered Services
- Baidu Announces Apollo 1.5 and a 10 Billion yuan Autonomous Driving Fund
- Clearing the Confusion: AI vs Machine Learning vs Deep Learning Differences

In this Edition:

Headlines:

- COVER STORY: AI Training Platform
- ADLINK to Showcase Carrier-grade Network Security Platform and Latest AI Training Platform at Edge Computing Congress 2018
- Baidu Announces Apollo 1.5 and a 10 Billion yuan Autonomous Driving Fund
- New NVIDIA Data Center Inference Platform to Fuel Next Wave of AI-Powered Services
- Microsoft acquires Lobe to help bring AI development capability to everyone
- SK Telecom Deploys Xilinx FPGAs for AI Acceleration, Achieves 5X Performance/16X Performance-per-watt over GPUs
- Mercury Systems Announces High-Performance Compute Blade for Artificial Intelligence and Other Advanced On-Platform Processing Applications
- Clearing the Confusion: AI vs Machine Learning vs Deep Learning Differences



- Will Artificial Intelligence empower or hinder creativity?
- AI Accelerating Discovery

Daniel Dierickx
CEO & co-Founder
at e2mos
Acting Chief Editor



*Over 3 decades
Chips & Embedded Systems
Global Market Expertise*

Dear Reader,

Here is your free copy of **AI World**, one of our five e-magazines published by e2mos.

Those e-magazines are distributed Free of charge WORLDWIDE to our PREMIER Database.

Each e-magazines has its own Dedicated Website.

Our aim is to provide you with relevant information in relation with your business, we are listening to our READERS in order to adapt the content & the right mix.

In the mean time we have produced over 2.000 Editions.

FREE Subscription

Click on the logos

aiworld

IoT World

Telecom COTS World
Broadband Broadcast IoT AI Convergence

Embedded Systems World

ATCA World

Editor/Publisher: e2mos

WEB: www.e2mos.com

Contact: mgt@e2mos.com

About e2mos SERVICES

- Business Development
- Coaching Biz Discovery
- Publications & e-mailings

ADLINK to Showcase Carrier-grade Network Security Platform and Latest AI Training Platform at Edge Computing Congress 2018

As a premier supplier to the telecom and network markets, ADLINK delivers high performance platforms offering I/O density and scalability, enabling customers to build solutions meeting the demands of next-generation, high-end application scenarios

Mannheim, Germany | 05-Sep-2018

ADLINK Technology, Inc., a global provider of advanced Edge Computing products, will demonstrate its carrier-grade network security platform CSA-7400 and latest AI training platform ALPS-4800 at Edge Computing Congress from Sept. 18 - 20 in Berlin. In addition, ADLINK's Xavier Serra, telecom and networking business development manager and subject matter expert, will present on MEC Solutions for Different Vertical Markets: NFV/SDN, 5G, Smart Cities and IIOT on Sept. 19 at 3:20PM.

ADLINK's CSA-7400 is a next-generation, high performance, carrier grade COTS network security platform built on the Open Compute Carrier-grade Edge Reference Architecture (OCCERA) by ADLINK, integrating network interfaces, switches, and overall computing capacity. The CSA-7400's flexibility and configurability enables cross-business product deployment and easy integration to other high-end network security markets, such as next-generation firewalls, telecommunications, and Multi-access Edge Computing (MEC). Featuring high throughput capacity and I/O density, parallel computing and computing density, carrier-grade high availability, and support for standardized API management, the application-ready CSA-7400 meets the stringent requirements for network security systems, and enhances solution integrators' competitive positions by allowing them to focus their development efforts on differentiating end applications.



4U 19" Network Appliance
CSA-7400 [Click Here](#)



AI Training Platform
ALPS-4800 [Click Here](#)

ADLINK's ALPS-4800 is the company's latest AI training platform, validated with up to 8 NVIDIA® Tesla® P100/V100 accelerators in a 4U server design, providing more than just a hardware server system. Rather, the ALPS-4800 is an optimized, workload-conscious acceleration system for Machine Learning (ML), Deep Learning (DL) and High Performance Computing (HPC) applications for smart factory, smart city, mobile robotics, defense, medical, and digital security and surveillance (DSS). Integrating a server node powered by dual Intel® Xeon® Scalable processors and PCIe expansion box with PCIe switching, the ALPS-4800 platform is highly configurable and scalable to support the combination of NVIDIA GPUs required to meet the needs of a variety of applications, supporting different workloads and demands, multiple topologies and bandwidths between GPUs and CPUs with simple cable routing adjustments. Moreover, InfiniBand support allows the platform to be easily scaled up to multiple GPU clusters.

By leveraging more than 20 years of expertise in developing highly reliable and available embedded computing systems, ADLINK is a premier supplier of extensive, cost-effective COTS, as well as fast time-to-market ODM solutions to worldwide tier-one TEMs and network security integrators. ADLINK ensures best practices in product obsolescence and lifecycle management by leveraging its long-standing strategic partnerships with major processor and software vendors. ADLINK offers design services in every major geographic region, benefiting customers with increased responsiveness, short delivery lead-time and ease of doing business. ADLINK focuses on continued development to help customers effectively mitigate budget constraint, and smoothly and seamlessly take on technology migration and product integration.

Edge Computing Congress is a meeting point for global leading operators and vendors to discuss the most pressing topics of Edge Computing, covering 5G networks, cyber security, artificial intelligence, augmented and virtual reality, IoT and Industry 4.0. To schedule a meeting with the ADLINK team at the event, please contact Xavier Serra (xavier.serra@adlinktech.com).

Baidu Announces Apollo 1.5 and a 10 Billion yuan Autonomous Driving Fund



BEIJING, Sept. 21, 2017 (GLOBE NEWSWIRE) -- Baidu, Inc. (NASDAQ: BIDU), announced on Wednesday the release of Apollo 1.5, the latest iteration of the company's Apollo open-source autonomous driving platform that has gathered 70 strong partners so far. The company also announced a 10 billion RMB (\$1.5 billion) Apollo Fund to invest in 100 autonomous driving projects in the next three years.

Building on Apollo 1.0, Apollo 1.5 opens up five additional core capabilities which include obstacle perception, planning, cloud simulation, High-Definition (HD) maps and End-to-End deep learning, providing more comprehensive solutions to developers and ecosystem partners to accelerate the deployment of autonomous driving.



A car deployed with Apollo capabilities is on display at Baidu Developer Conference in Beijing in July.

The obstacle perception capability enables vehicles to accurately identify obstacles during both day and night. With the planning capability, vehicles can plan the correct driving path and make optimal driving decisions. The cloud simulation system on Apollo is the only platform that provides open, built-in HD maps tailored for autonomous driving. Based on large scale cloud computing capacity, Apollo has a vast amount of real traffic data in China, and has a virtual running capacity of one million kilometers per day.

Apollo is widely regarded as the "Android of the auto industry." First announced in April, Apollo is an open platform that provides a comprehensive, secure, and reliable all-in-one solution that supports all major features and functions of an autonomous vehicle. Its first iteration, Apollo 1.0, with capabilities enabling vehicles to do autonomous waypoint driving in enclosed venues, was announced in July at Baidu's inaugurate AI Developers Conference in Beijing.

Over the past 2 months, Apollo has seen dozens of code updates each week and the addition of more than 65,000 lines of new code. At the same time, Apollo has received a positive response from global developers. To date, more than 1,300 companies have downloaded Apollo source code and nearly 100 companies have applied for open data via the Apollo website.

Apollo has attracted [70 global and Chinese partners](#), including OEMs, Tier 1 suppliers, developer platforms and technology start-ups. Newly joined members include Hyundai Motor, ROS, esd electronics, Neousys Technology, and autonomous driving startups such as Momenta and iDriver+ Technologies.

Many of those partnerships have already begun to bear fruit. Baidu has signed over 50 cooperation agreements with Apollo partners on mass production or joint product development plans. King Long, a Xiamen-based commercial vehicle manufacturer, performed autonomous waypoint driving in enclosed venues using buses deployed with Apollo's 1.0 capabilities. Momenta, a Beijing-based autonomous driving startup, successfully conducted testing on designated lanes using Apollo 1.5 enabled cars, which were able to accurately recognize obstacles, passengers, and make optimal decisions even at night when visibility is very low.

At Wednesday's event, Baidu and Velodyne, the world's leading developer of LiDAR sensors, jointly announced that Velodyne will provide Apollo ecosystem members with priority access to product information, technical support as well as a shorter lead time to product purchase. Baidu will also partner with Udacity, a Silicon Valley based online education platform, to foster autonomous talent globally by introducing autonomous courses, certifications, special competitions and talent services.

Developers and Apollo ecosystem partners can access the Apollo website (apollo.auto) for more information.

About Baidu

Baidu, Inc. is the leading Chinese language Internet search provider. Baidu aims to make a complex world simpler through technology. Baidu's ADSs trade on the NASDAQ Global Select Market under the symbol "BIDU". Currently, ten ADSs represent one Class A ordinary share.

Media Contact Baidu International Communications intlcomm@baidu.com



New NVIDIA Data Center Inference Platform to Fuel Next Wave of AI-Powered Services

Tesla T4 GPU and New TensorRT Software Enable Intelligent Voice, Video, Image and Recommendation Services

Wednesday, September 12, 2018

GTC Japan -- Fueling the growth of AI services worldwide, NVIDIA today launched an [AI data center platform](#) that delivers the industry's most advanced inference acceleration for voice, video, image and recommendation services.

The NVIDIA TensorRT™ Hyperscale Inference Platform features [NVIDIA® Tesla® T4 GPUs](#) based on the company's breakthrough NVIDIA Turing™ architecture and a comprehensive set of new inference software.

Delivering the fastest performance with lower latency for end-to-end applications, the platform enables hyperscale data centers to offer new services, such as enhanced natural language interactions and direct answers to search queries rather than a list of possible results.

"Our customers are racing toward a future where every product and service will be touched and improved by AI," said Ian Buck, vice president and general manager of Accelerated Business at NVIDIA. "The NVIDIA TensorRT Hyperscale Platform has been built to bring this to reality — faster and more efficiently than had been previously thought possible."

Every day, massive data centers process billions of voice queries, translations, images, videos, recommendations and social media interactions. Each of these applications requires a different type of neural network residing on the server where the processing takes place.

To optimize the data center for maximum throughput and server utilization, the NVIDIA TensorRT Hyperscale Platform includes both real-time inference software and Tesla T4 GPUs, which process queries up to 40x faster than CPUs alone.

NVIDIA estimates that the AI inference industry is poised to grow in the next five years into a \$20 billion market.

Industry's Most Advanced AI Inference Platform

The NVIDIA TensorRT Hyperscale Platform includes a comprehensive set of hardware and software offerings optimized for powerful, highly efficient inference. Key elements include:

- NVIDIA Tesla T4 GPU – Featuring 320 Turing Tensor Cores and 2,560 CUDA® cores, this new GPU provides breakthrough performance with flexible, multi-precision capabilities, from FP32 to FP16 to INT8, as well as INT4. Packaged in an energy-efficient, 75-watt, small PCIe form factor that easily fits into most servers, it offers 65 teraflops of peak performance for FP16, 130 TOPS for INT8 and 260 TOPS for INT4.
 - [NVIDIA TensorRT 5](#) – An inference optimizer and runtime engine, NVIDIA TensorRT 5 supports Turing Tensor Cores and expands the set of neural network optimizations for multi-precision workloads.
- NVIDIA TensorRT inference server – This containerized microservice software enables applications to use AI models in data center production. Freely available from the [NVIDIA GPU Cloud](#) container registry, it maximizes data center throughput and GPU utilization, supports all popular AI models and frameworks, and integrates with Kubernetes and Docker.

Supported by Technology Leaders Worldwide

Support for NVIDIA's new inference platform comes from leading consumer and business technology companies around the world.

"We are working hard at [Microsoft](#) to deliver the most innovative AI-powered services to our customers," said Jordi Ribas, corporate vice president for Bing and AI Products at Microsoft. "Using NVIDIA GPUs in real-time inference workloads has improved Bing's advanced search offerings, enabling us to reduce object detection latency for images. We look forward to working with NVIDIA's next-generation inference hardware and software to expand the way people benefit from AI products and services."

Chris Kleban, product manager at [Google Cloud](#), said: "AI is becoming increasingly pervasive, and inference is a critical capability customers need to successfully deploy their AI models, so we're excited to support NVIDIA's Turing Tesla T4 GPUs on Google Cloud Platform soon."

More information, including details on how to request early access to T4 GPUs on Google Cloud Platform, is available [here](#).

... to next page

New NVIDIA Data Center Inference Platform to Fuel Next Wave of AI-Powered Services

... from previous page

Additional companies, including all major server manufacturers, voicing support for the NVIDIA TensorRT Hyperscale Platform include:

“Cisco's UCS portfolio delivers policy-driven, GPU-accelerated systems and solutions to power every phase of the AI lifecycle. With the NVIDIA Tesla T4 GPU based on the NVIDIA Turing architecture, Cisco customers will have access to the most efficient accelerator for AI inference workloads — gaining insights faster and accelerating time to action.”
— Kaustubh Das, vice president of product management, Data Center Group, Cisco

“Dell EMC is focused on helping customers transform their IT while benefiting from advancements such as artificial intelligence. As the world's leading provider of server systems, Dell EMC continues to enhance the PowerEdge server portfolio to help our customers ultimately achieve their goals. Our close collaboration with NVIDIA and historical adoption of the latest GPU accelerators available from their Tesla portfolio play a vital role in helping our customers stay ahead of the curve in AI training and inference.”
— Ravi Pendekanti, senior vice president of product management and marketing, Servers & Infrastructure Systems, Dell EMC

“Fujitsu plans to incorporate NVIDIA's Tesla T4 GPUs into our global Fujitsu Server PRIMERGY systems lineup. Leveraging this latest, high-efficiency GPU accelerator from NVIDIA, we will provide our customers around the world with servers highly optimized for their growing AI needs.”
— Hideaki Maeda, vice president of the Products Division, Data Center Platform Business Unit, Fujitsu Ltd.

“At HPE, we are committed to driving intelligence at the edge for faster insight and improved experiences. With the NVIDIA Tesla T4 GPU, based on the NVIDIA Turing architecture, we are continuing to modernize and accelerate the data center to enable inference at the edge.”
— Bill Mannel, vice president and general manager, HPC and AI Group, Hewlett Packard Enterprise

“IBM Cognitive Systems is able to deliver 4x faster deep learning training times as a result of a co-optimized hardware and software on a simplified AI platform with PowerAI, our deep learning training and inference software, and IBM Power Systems AC922 accelerated servers. We have a history of partnership and innovation with NVIDIA, and together we co-developed the industry's only CPU-to-GPU NVIDIA NVLink connection on IBM Power processors, and we are excited to explore the new NVIDIA T4 GPU accelerator to extend this state of the art leadership for inference workloads.”
— Steve Sibley, vice president of Power Systems Offering Management, IBM

“Kubeflow - We are excited to see NVIDIA bring GPU inference to Kubernetes with the NVIDIA TensorRT inference server, and look forward to integrating it with Kubeflow to provide users with a simple, portable and scalable way to deploy AI inference across diverse infrastructures.”
— David Aronchick, co-founder and product manager of Kubeflow

“Oracle Cloud Infrastructure - Open source cross-framework inference is vital to production deployments of machine learning models. We are excited to see how the NVIDIA TensorRT inference server, which brings a powerful solution for both GPU and CPU inference serving at scale, enables faster deployment of AI applications and improves infrastructure utilization.”
— Kash Iftikhar, vice president of product development, Oracle Cloud Infrastructure

“Supermicro is innovating to address the rapidly emerging high-throughput inference market driven by technologies such as 5G, Smart Cities and IOT devices, which are generating huge amounts of data and require real-time decision making. We see the combination of NVIDIA TensorRT and the new Turing architecture-based T4 GPU accelerator as the ideal combination for these new, demanding and latency-sensitive workloads and plan to aggressively leverage them in our GPU system product line.”
— Charles Liang, president and CEO, Supermicro

Keep Current on NVIDIA

Subscribe to the NVIDIA blog, follow us on Facebook, Google+, Twitter, LinkedIn and Instagram, and view NVIDIA videos on YouTube and images on Flickr.

About NVIDIA

NVIDIA's (NASDAQ: NVDA) invention of the GPU in 1999 sparked the growth of the PC gaming market, redefined modern computer graphics and revolutionized parallel computing. More recently, GPU deep learning ignited modern AI — the next era of computing — with the GPU acting as the brain of computers, robots and self-driving cars that can perceive and understand the world.

More information at <http://nvidianews.nvidia.com/>.

Microsoft acquires Lobe to help bring AI development capability to everyone



Sep 13, 2018 | [Kevin Scott - Executive Vice President and Chief Technology Officer, Microsoft](#)



Lobe's simple visual interface lets people easily create intelligent apps that can understand hand gestures, hear music, read handwriting, and more.

Technology has already transformed the world we live in. Computers are ubiquitous, from the machines on our desks to the devices in our pockets and in our homes. Now, breakthroughs in artificial intelligence (AI) and deep learning are helping scientists treat [cancer more effectively](#), farmers figure out how to [grow more food using fewer natural resources](#), and give people from different countries the ability to communicate [across language barriers](#).

In many ways though, we're only just beginning to tap into the full potential AI can provide. This in large part is because AI development and building deep learning models are slow and complex processes even for experienced data scientists and developers. To date, many people have been at a disadvantage when it comes to accessing AI, and we're committed to changing that.

Over the last few months, we've made [multiple investments](#) in companies to further this goal. The acquisition of [Semantic Machines](#) in May brought in a revolutionary new approach to conversational AI, and the acquisition of [Bonsai](#) in July will help reduce the barriers to AI development through the Bonsai team's unique work combining machine teaching, reinforcement learning and simulation. These are just two recent examples of investments we have made to help us accelerate the current state of AI development.

Lobe logo

Today, we're excited to announce the acquisition of [Lobe](#). Based in San Francisco, Lobe is working to make deep learning simple, understandable and accessible to everyone. Lobe's simple visual interface empowers anyone to develop and apply deep learning and AI models quickly, without writing code.



We look forward to continuing the great work by Lobe in putting AI development into the hands of non-engineers and non-experts. We're thrilled to have Lobe join Microsoft and are excited about our future together to simplify AI development for everyone.

SK Telecom Deploys Xilinx FPGAs for AI Acceleration,



Achieves 5X Performance/16X Performance-per-watt over GPUs

SAN JOSE, Calif. and SEOUL, South Korea, Aug. 15, 2018 /PRNewswire/ -- Xilinx, Inc. (NASDAQ: XLNX) and SK Telecom (SKT) today jointly announced that SKT has deployed Xilinx® FPGAs as their artificial intelligence (AI) accelerators in its data center. The Xilinx® Kintex® UltraScale™ FPGAs are now running SKT's automatic speech-recognition (ASR) application to accelerate NUGU, its voice-activated assistant. SKT achieved up to five times higher performance in ASR applications when compared to GPUs, and more importantly, 16 times better performance-per-watt. This is the first commercial adoption of FPGA accelerators in the AI domain for large-scale data centers in South Korea.



The FPGA-based accelerator lowers the total cost of ownership (TCO) of ASR application servers by populating existing CPU-only servers with efficient Xilinx FPGA add-in cards. The ASR servers easily and simply accelerate multiple voice service channels with Xilinx FPGA cards in their empty slots. One FPGA card provides more than five times the performance of a single server, resulting in substantial cost savings. This performance and TCO advantage benefits SK Telecom and its customers.

"Over many years we have seen the shape of the industry evolve, and we are proud to be at the forefront of developing AI accelerators. By designing our solution based on the Xilinx KCU1500 board and our own bitstream image, we have developed a cost-effective, high-performance application," said Kang-Won Lee, senior vice president, software research and development center at SKT.

The adaptive nature of Xilinx FPGAs enables fast deployment of custom hardware accelerators for the rapidly evolving field of AI and deep learning. Moreover, FPGAs provide higher performance and lower latency at lower power when compared to CPUs and GPUs. SK Telecom joins a fast-growing list of high-profile commercial data centers that have deployed FPGAs for compute acceleration.

"We are delighted to have the opportunity to deploy Xilinx FPGAs to SKT's AI data center, a first in South Korea," said Manish Muthal, vice president, marketing for data center at Xilinx. "Xilinx Kintex UltraScale KCU1500 FPGAs show that Xilinx has a major competitive edge in applications. Xilinx will continue to focus its technological capabilities and innovations on data center acceleration."

SKT plans to demonstrate its FPGA-based AI solution at the [Xilinx Developer Forum \(XDF\)](#) being held in Silicon Valley on October 1-2, and in Beijing on October 16.

About SK Telecom

Established in 1984, SK Telecom is the largest mobile operator in Korea by both revenue and number of subscribers. As of December 2017, the company holds around 50 percent of the market, with 30.2 million mobile subscribers including 22.87 million LTE subscribers. It has reached KRW 17.520 trillion in revenue in 2017. SK Telecom has led the advancement of mobile technologies ranging from 2G to 4G, and is currently setting important milestones in its journey to 5G. The company is not only leading innovation in the field of mobile network, but also providing IoT, media, home and platform services. SK Telecom is determined to play a significant role in the Fourth Industrial Revolution by achieving innovations and promoting shared growth with other players in the industry. For more information visit www.globalskt.com.

About Xilinx

Xilinx develops highly flexible and adaptive processing platforms that enable rapid innovation across a variety of technologies – from the endpoint to the edge to the cloud. Xilinx is the inventor of the FPGA, hardware programmable SoCs and the ACAP, designed to deliver the most dynamic processor technology in the industry and enable the adaptable, intelligent and connected world of the future. For more information, visit www.xilinx.com.



Mercury Systems Announces High-Performance Compute Blade for **Artificial Intelligence** and Other Advanced On-Platform Processing Applications

Scalable Xeon and UltraScale FPGA processing power, and extensive memory packaged in rugged, secure, small form-factor blade for deployment anywhere

ANDOVER, Mass., Sept. 13, 2018 (GLOBE NEWSWIRE) -- Mercury Systems, Inc. (NASDAQ: MRCY, www.mrcy.com) announced the EnsembleSeries™ LDS3517 processing blade for advanced on-platform processing, machine learning and artificial intelligence (AI) applications. Each blade combines an Intel Xeon® D server-class processor, a Xilinx® UltraScale® FPGA and a mezzanine site in a 3U OpenVPX™ form-factor. Wafer-stacking and system-in-package (SiP) miniaturization technologies enable the latest general and FPGA processing capabilities, their extensive supporting memory and a versatile mezzanine expansion site to fit into this compact form-factor. The LDS3517 blade is ideally suited to on-platform cognitive electronic warfare (EW), next-generation radar, machine learning and AI applications that require small, powerful and scalable processing engines. Optionally configured with embedded BuiltSECURE™ systems security engineering and packaged with modified-off-the-shelf-plus (MOTS+) technology, the LDS3517 compute blade can support military missions anywhere.

"Our customers are asking us for small, powerful processing blades they can run their most compute-intense applications on and these resources need to be rugged and secure for deployment at the tactical edge," said Joe Plunkett, Mercury's Senior Director and General Manager for Sensor Processing Solutions. "The EnsembleSeries LDS3517 meets these needs, making all manner of new missions possible from advanced on-platform electro-optical/infrared (EO/IR) and EW cognitive processing to platform autonomy and AI."

EnsembleSeries LDS3517 blades feature:


Embedded BuiltSECURE technology: Proven, built-in, system-wide security across software, firmware and hardware in the domains of system security engineering (SSE), trust and cyber-hardening.

Dense, versatile processing capability: The LDS3517 is a 3U OpenVPX Xeon D family-powered processing blade with powerful UltraScale FPGA support and a mezzanine site for customization and broad-streaming I/O pre-processing, off-boarding and deep packet inspection.

Versatile mezzanine site: An XMC mezzanine site supports I/O customization, making the LDS3517 exceptionally versatile for a wide variety of high-bandwidth signal processing applications. Each blade is interoperable with Mercury's other EnsembleSeries OpenVPX processing building blocks for rapid processing subsystem pre-integration.

Rugged packaging: The LDS3517 optionally includes MOTS+ technology that equips blades with the highly rugged embedded packaging. MOTS+ is an additional layer of environmental protection and ruggedness that requires the soldering of all board-level devices, regardless of their native packaging/terminations, to their respective substrates for reliability.

On-board AI processing: AI applications require higher levels of processing power and larger memories. The LDS3517's Xeon D family processor is supported with a powerful UltraScale FPGA and the system memory required to run machine learning and AI applications.



Mercury's EnsembleSeries™ processing solutions are designed, made, programmed and supported in the USA in DMEA-accredited facilities using devices from our trusted, managed supply chain. Customer orders are currently being accepted for LDS3517 blades in conduction and Air Flow Through (VITA 48.8) one-inch packages with shipments beginning in Q3, 2018.

For more information, visit www.mrcy.com/LDS3517 or contact Mercury at +1 (866) 627-6951 or info@mrcy.com.

Mercury Systems – Innovation That Matters™

Mercury Systems (NASDAQ:MRCY) is a leading commercial provider of secure sensor and safety-critical processing subsystems. Optimized for customer and mission success, Mercury's solutions power a wide variety of critical defense and intelligence programs. Headquartered in Andover, Mass., Mercury is pioneering a next-generation defense electronics business model specifically designed to meet the industry's current and emerging technology needs. To learn more, visit www.mrcy.com.

Clearing the Confusion: AI vs Machine Learning vs Deep Learning Differences

By: Dr. Michael J. Garbade, Founder & CEO of Education Ecosystem | 14-Sep-2018



Raise your hand if you've been caught in the confusion of differentiating artificial intelligence (AI) vs machine learning (ML) vs deep learning (DL)...

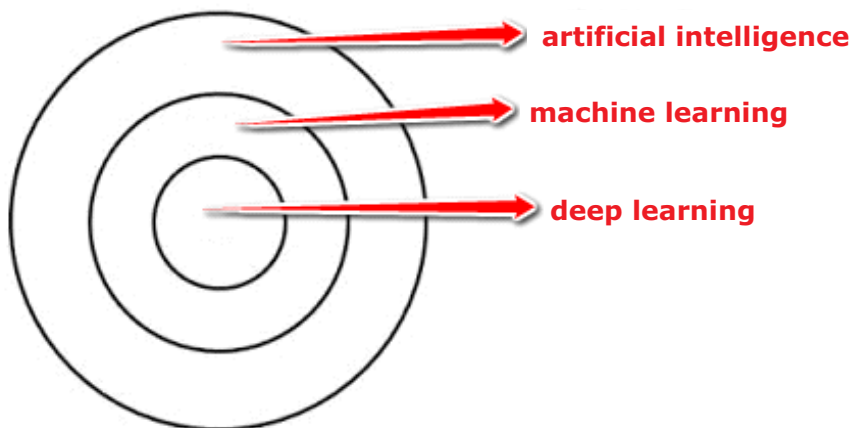
Bring down your hand, buddy, we can't see it!

Although the three terminologies are usually used interchangeably, they do not quite refer to the same things.

[Andrey Bulezyuk](#), who is a German-based computer expert and has more than five years of experience in teaching people how artificial intelligence systems work, says that "practitioners in this field can clearly articulate the differences between the three closely-related terms."

Therefore, is there a difference between artificial intelligence, machine learning, and deep learning?

Here is an image that attempts to visualize the distinction between them:



As you can see on the above image of three concentric circles, DL is a subset of ML, which is also a subset of AI.

Interesting?

So, AI is the all-encompassing concept that initially erupted, then followed by ML that thrived later, and lastly DL that is promising to escalate the advances of AI to another level.

Let's dig deeper so that you can understand which is better for your specific use case: artificial intelligence, machine learning, or deep learning.

... to next page

Will Artificial Intelligence empower or hinder creativity?

Published on September 26, 2017
Christian von Reventlow, Group Executive Products & Technologies at Telstra



Discussed this topic on DTs second Museum-Talk session on Digital Responsibility with 100 participants in Berlin Germany. Highlights:.

Professor Hartmut Fladt (Music) suggested that AI expands creative composing. However feels AI is lacking the emotions – specifically the typical human element of controversial emotions. Plus that music is strongly influenced by cultural upbringing. Both difficult to train an AI on.

Professor Ulrich Furbach (AI) pointed out, that AIs will learn to judge their own behaviours – as a human. This does not need to be adversarial - it might be supporting each other. AI might end up being a muse for an artist.

My view – a core element of creativity is to generate surprises for the observer. And AI's are able to find surprising solutions to – so they are creative.

Professor Sebastian Schnieder (experimentell Psycho-Physiology) does not see Hollywood's horror movies becoming true. Suggested to experiment and learn. His biggest concern: overpromising and subsequent disappointment.

Great discussion with the audience. And more to come:

October 5 on Cyber Security in Flensburg Germany's Phänomenta,

October 26 on Digitization of Design and Production Processes at the Vitra Design Museum.

Tags: #DigitalResponsibility #Museumstak #Innovation #AI #KI #Innovation

Source: [Click Here](#)

Clearing the Confusion: AI vs Machine Learning vs Deep Learning Differences

... from previous page

What is artificial intelligence?

As the name suggests, artificial intelligence can be loosely interpreted to mean incorporating human intelligence to machines.

Artificial intelligence is the broader concept that consists of everything from Good Old-Fashioned AI ([GOFAI](#)) all the way to futuristic technologies such as deep learning.

Whenever a machine completes tasks based on a set of stipulated rules that solve problems (algorithms), such an "intelligent" behavior is what is called artificial intelligence.

For example, such machines can move and manipulate objects, recognize whether someone has raised the hands, or solve other problems.

AI-powered machines are usually classified into two groups—general and narrow. The general artificial intelligence AI machines can intelligently solve problems, like the ones mentioned above.

The narrow intelligence AI machines can perform specific tasks very well, sometimes better than humans—though they are limited in scope.

The technology used for classifying images on Pinterest is an example of narrow AI.

What is machine learning?

As the name suggests, [machine learning](#) can be loosely interpreted to mean empowering computer systems with the ability to "learn".

The intention of ML is to enable machines to learn by themselves using the provided data and make accurate predictions.

ML is a subset of artificial intelligence; in fact, it's simply a technique for realizing AI.

It is a method of training algorithms such that they can learn how to make decisions.

Training in machine learning entails giving a lot of data to the algorithm and allowing it to learn more about the processed information.

For example, here is a table that identifies the type of fruit based on its characteristics:

Weight (grams)	Texture	Type of Fruit
155	Rough	Orange
180	Rough	Orange
135	Smooth	Apple
110	Smooth	Apple
120	Smooth	?

As you can see on the table above, the fruits are differentiated based on their weight and texture.

However, the last row gives only the weight and texture, without the type of fruit.

And, a machine learning algorithm can be developed to try to identify whether the fruit is an orange or an apple.

After the algorithm is fed with the training data, it will learn the differing characteristics between an orange and an apple.

Therefore, if provided with data of weight and texture, it can predict accurately the type of fruit with those characteristics.

... to next page

Clearing the Confusion: AI vs Machine Learning vs Deep Learning Differences

... from previous page

What is deep learning?

As earlier mentioned, deep learning is a subset of ML; in fact, it's simply a technique for realizing machine learning. In other words, DL is the next evolution of machine learning.

DL algorithms are roughly inspired by the information processing patterns found in the human brain.

Just like we use our brains to identify patterns and classify various types of information, deep learning algorithms can be taught to accomplish the same tasks for machines.

The brain usually tries to decipher the information it receives. It achieves this through labelling and assigning the items into various categories.

Whenever we receive a new information, the brain tries to compare it to a known item before making sense of it --which is the same concept deep learning algorithms employ.

For example, artificial neural networks (ANNs) are a type of algorithms that aim to imitate the way our brains make decisions.

Comparing deep learning vs machine learning can assist you to understand their subtle differences.

For example, while DL can automatically discover the features to be used for classification, ML requires these features to be provided manually.

Furthermore, in contrast to ML, DL needs high-end machines and considerably big amounts of training data to deliver accurate results.

Wrapping up

Do you now understand the difference between AI vs ML vs DL?

Then, raise your hands...

We promise to develop an [AI algorithm](#) that tells us whenever someone raises their hand.



Dr. Michael J. Garbade

Founder & CEO of Education Ecosystem.
Serial entrepreneur with experience from Amazon, GE & Rebate Networks

AI Accelerating Discovery

Published by Joanne Itow | Mon, 2018-04-30 16:25

In early April 2018, the Materials Research Society held their spring meeting and exhibit at the Phoenix, Arizona convention center. With over 110 symposium presentations, it was difficult to select which sessions to attend. But one forum caught my eye, "AI for Materials Development". These days AI seems to be everywhere.


As we all speculate about the impact of AI on autonomous driving and the next killer app, Carla Gomes, Professor of Computer Science and director of the Institute for Computational Sustainability at Cornell University, is focusing on large-scale constraint-based reasoning. She pointed out that AI still can't compete with good ol' human common sense. Human reasoning and inference planning are still lacking in most AI systems. One of the key fundamentals of AI is building a neural network that resembles the human brain. Even with the advancements of 7nm silicon technology, this is a daunting task, not to mention the complexities of software algorithms to mimic the human thought and decision process.

But in the world of materials development, AI excels. By integrating material experimentation and AI, the discovery of new materials and the application of materials in the real world is progressing at an accelerated pace. AI is capable of developing the hypotheses and—along with robotics—is following through with new scientific discovery.

Subbarao Kambhampati is Professor of Computer Science & Engineering at Arizona State University and President of AAI (Association for the Advancement of Artificial Intelligence). He pointed out that AI has developed in reverse of a child's mind. A child begins with visual stimulation, moves to vocal patterning and eventually reasoning. Machine learning started with the programming from human reasoning and moved to speech recognition and now image recognition.

Both speakers place a high value on the impact of AI on research science, but Kambhampati pointed out the importance of Polany's Paradox. Michael Polany claimed that all knowledge relies on personal judgement. AI is based on algorithms and rules, but human knowledge and capability relies on skills and experiences that often lie beneath our conscious mind and are transmitted to us via culture and tradition. The main takeaway is that 'we know more than we can tell'. Semico believes that AI is invaluable in research, but many people believe it will take a long time before AI can replace the human quality of common sense.

Semico Research recently released a Tech Brief, Artificial Intelligence: Powering the Next Generation of Processors. For more details about Semico's report, contact Rick Vogeley at rickv@semico.com.



Joanne Itow

Managing Director, Manufacturing for Semico Research
<https://semico.com/content/joanne-itow>