



Julia Computing and Intel: Accelerating the Artificial Intelligence revolution

Julia Computing and Intel are powering the next big wave in Artificial Intelligence with high performance, advanced Machine Learning and Deep Learning based on Intel® architecture

“With Julia, we believe that we have the architecture that has closed the gap between hardware and software. Intel recognized this benefit of Julia and since then, has been closely working with us to make sure that Julia runs the best on latest Intel® processors.”



Viral B. Shah
Co-Founder and CEO,
Julia Computing



Executive summary

Artificial Intelligence (“AI”) is at the heart of today’s technological innovation. From automatic speech recognition to driverless cars, AI is advancing rapidly and is undoubtedly, the new way forward. Artificial Intelligence is the next big wave in computing—and is poised to usher in a better world, in the order of major transformations before it like the agricultural revolution, industrial revolution, and the information age.

Julia Computing is among the frontrunners of the AI revolution with “Julia”, a high performance advanced Machine Learning language. Julia’s speed and the performance of Intel® technologies together are enabling to redefine the AI landscape. Intel and Julia Computing is driving state-of-the-art collaborative AI and enabling to maximize the potential of AI.

Julia, which can be easily integrated across a wide spectrum of domains, is helping in driving AI solutions for sectors such as finance, healthcare, 3D printing, astronomy and pharmacokinetics, among others.

AI – The next wave of computing

The future is set to witness an insatiable demand for AI computing. Even though the AI revolution started decades ago—as early as the 1960s—it is only in the recent past that the true potential of AI has emerged. Some of the key factors that are bringing AI to life today include:

Compute breakthrough: Driven by Moore’s Law, compute capability and architectural innovation have progressed to the point where the threshold required to support the intense demands of machine intelligence has been crossed.

Data deluge: The world of smart and connected devices has unleashed a data deluge. The Internet of Things (IoT) will include a projected 200 billion smart-and-connected devices by 2020¹, and the data produced is expected to double every two years to total 40 zettabytes by 2020².

Innovation surge: The road to AI is also being driven by a surge of innovation in myriad of fields. Each new AI algorithmic innovation and use case opens more eyes to the power of AI, leading more innovators to join the community and stimulating an ever-increasing demand for the technology.

Realizing the potential of AI with Machine Learning and Deep Learning

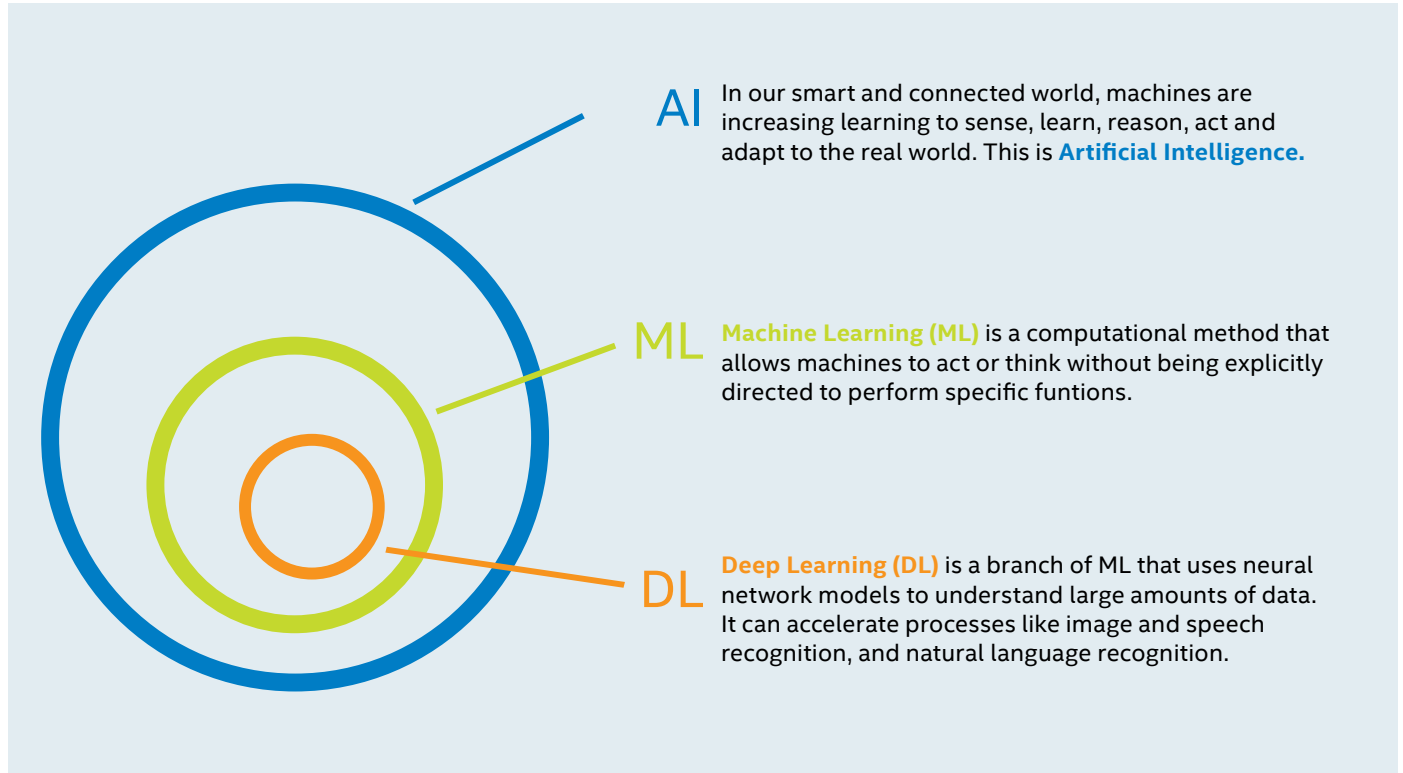


Fig. 1: AI solutions basically work on two concepts - Machine Learning and Deep Learning

Machine Learning and Deep Learning are some of the most capable AI tools for solving complex problems, and will continue to be developed and leveraged in the future. Machine Learning involves using a variety of algorithms that learn from analyzing data and improve performance based on real-world experience. Deep Learning is a rapidly emerging branch of Machine Learning, which relies on large data sets to iteratively train many-layered neural networks inspired by the human brain.

"AI has become increasingly interesting for everyone. Part of the allure for AI comes from this concept called Deep Learning, which basically tries to simulate the human brain. Neural work the way neurons work in the human brain. By simulating these neural networks, very complex problems that we believed to be difficult to be addressed by computers are now possible, whether it's language translation, speech recognition or image recognition," says Viral Shah, Co-Founder and Chief Executive Officer, Julia Computing.

Challenges of traditional AI solution development platforms

A large number of programming languages and development platforms for AI exist today. However, most of these languages or platforms are not optimized for maximizing the potential of AI. Mathematical or domain specialists prefer writing programs or algorithms on languages that are easy to use like Python* or R* or MATLAB* and then re-write the programs on C*, FORTRAN* or C#* for performance reasons.

Another key challenge for developing AI solutions is the requirement for extremely powerful, robust and scalable

computing platforms. "The challenge is that it needs a lot of computational horse power and that's where the world is moving towards and that's where Deep Learning needs a lot of data, it's using a lot of computational power on the cloud, lots of hardware and software coming around it," says Viral Shah.

Julia - A fresh approach to AI solution development

Viral Shah and his colleagues came up with the idea of creating an easy to use, scalable and high-performance Machine Learning language for AI. Julia Computing was founded in 2015 by the creators of the Julia programming language ("Julia") with a mission to develop products that make Julia easy to use, easy to deploy and easy to scale.

Julia provides the ideal programming environment for developing AI solutions by combining the functionality of advanced, quantitative environments such as R and Python with the speed of programming languages like Java* and C++*.

Not only does Julia solve the problem of using multiple languages but also provides high compatibility and interoperability with other AI languages for a seamless and optimal programming experience. Providing more insights into these advantages, Viral Shah adds, "Julia is ideal for AI algorithms and has interoperability with Python, R, C and FORTRAN as well. What Julia is able to do in the field of AI and Machine Learning is that one can write their own AI libraries in Julia from scratch. It gets the best performance from the hardware one may be running it on. Also, it's able to leverage existing Python or C based libraries such as TensorFlow* and MXNet*."

Why Julia for Machine Learning and Deep Learning?

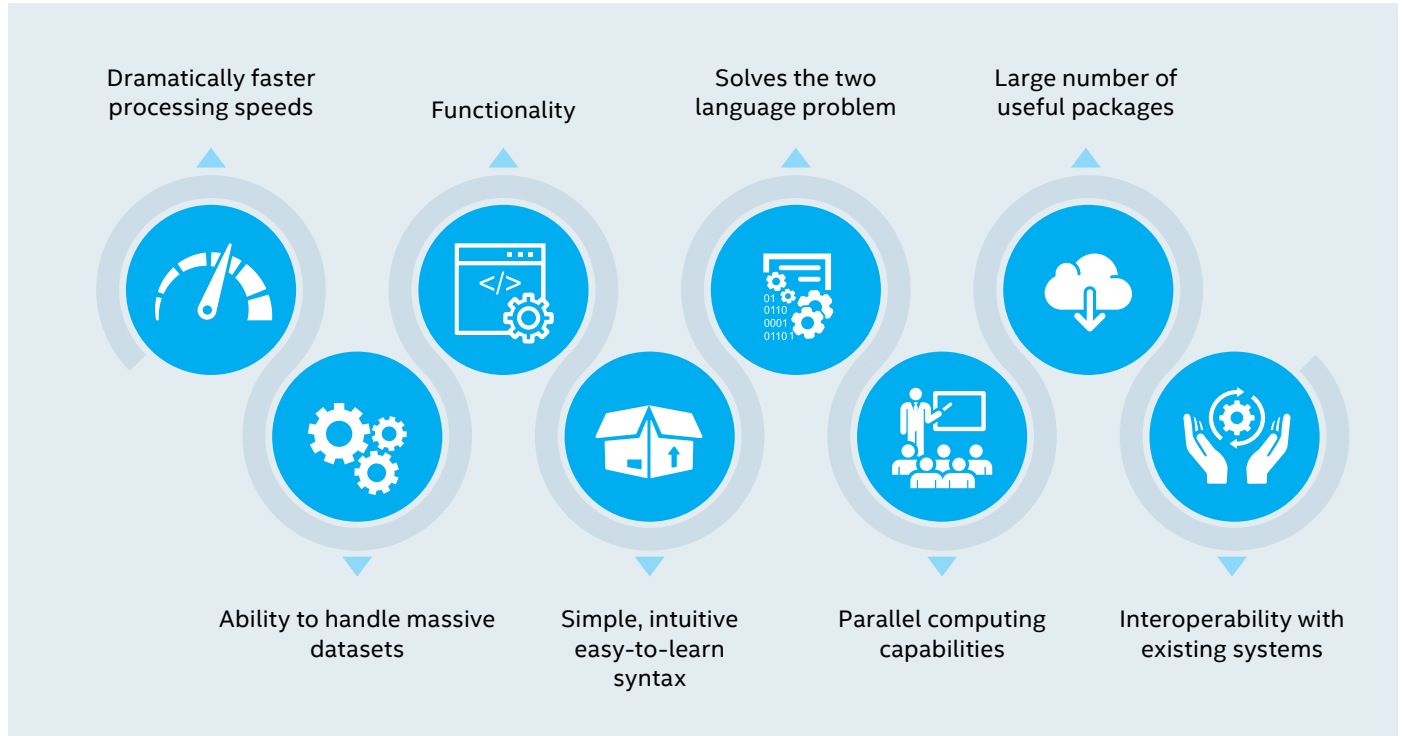


Fig. 2: Benefits of Julia for Machine Learning and Deep Learning

Julia Computing offers a complete portfolio of Julia based products that are built for the ease of programmers and program development. They are JuliaPro, JuliaBox, JuliaFin and JuliaRun. “With JuliaPro you can start developing without having to worry about open source. JuliaPro includes over 150 widely used Julia packages, including the best Machine Learning and AI libraries such as TensorFlow*, MXNet*, Mocha*, etc. JuliaRun helps you to deploy your Julia application to a large cluster cloud or a public cloud very easily. It helps you scale it to even thousands of nodes and manage your deployments. JuliaBox enables users to login free of cost, program and save it on the cloud. Finally, JuliaFin is built for finance professionals for asset management, risk management, capital allocation, trading and back testing among other functions,” enlightens Deepak Vinchhi, Co-Founder and Chief Operating Officer, Julia Computing.

Intel is unleashing the full compute potential for AI

Machine Learning entails parallel processing of vast amounts of mostly unstructured data—like video streams or sensor data feeds—in real time. Because of the specialized nature of this task, much of the preliminary effort has been focused on specialized graphics processors rather than industry-standard computing platforms.

Intel offers a comprehensive, flexible and performance optimized portfolio of hardware products for Machine Learning and Deep Learning, including Intel® Xeon® processors and Intel® Xeon Phi™ processors for general purpose infrastructures, as well as Field-Programmable Gate Arrays (FPGAs) and the Nervana™ platform for workload-optimized environments.

“With Julia, we believe that we have the architecture that has closed the gap between hardware and software. Intel recognized this benefit of Julia and since then has been closely working with us to make sure that Julia runs the best on latest Intel® processors,” says Viral Shah, explaining how Intel technology is key to advancing AI solutions.

As demonstrated above, Intel has the best-in-class hardware portfolio for Machine Learning and Deep Learning, but software is also a key element to unleash the full compute

“Machine Learning platforms like Julia require enormous computational ability that drives down the time to train. Intel has the best-in-class hardware portfolio and offers an optimized software stack to deliver game-changing AI applications.”



Margaret Burgraff
Vice President, Software and Services Group,
General Manager, Global Geographies & Business
Development, Intel

potential. Machine Learning platforms like Julia require both enormous computational ability that drives down the time to train and software frameworks that open up the possibilities of Machine Learning to software developers in every industry.

Intel offers an optimized software stack in order to deliver game-changing AI applications. At the basic level, Intel optimized primitive functions are used across a wide array of Machine Learning and Deep Learning frameworks and solutions. At the framework level, Intel is committed to optimizing the most popular analytics, Machine Learning and Deep Learning frameworks. Intel offers the Deep Learning SDK to accelerate Deep Learning training and deployment.

Strong collaborations to enable the future of AI

Intel and Julia Computing share a history of collaboration that goes back a long way. Emphasizing this, Viral Shah says, “We have been working closely with Intel from the very early days of Julia Computing, actually even before Julia Computing existed. We worked very close to leverage the best instruction sets that Intel provides in Julia programs so that the users of Julia can sort of leverage the power of Intel® processors.”

Intel was instrumental in driving the success of Julia, as we can gather from what Viral Shah says about the collaboration, “Very early on when Julia Computing was formed, Intel actually reached out to us. We had an agreement where we would get access to early release of hardware and get Julia running on their latest chip. Engineers from Intel would fix bugs as they show up, and at the same time we at Julia Computing would fix Julia related issues so that scientists using Julia on the latest Intel’s hardware get the performance benefits and ease of use that they are looking for. So it’s really a multi-party collaboration that has been possible because of our partnership with Intel.”

Driving innovative usages for AI

Julia Computing has been involved in groundbreaking efforts towards advancing and redefining the AI landscape and ultimately changing the lives of billions through innovative AI solutions. Key projects of Julia for redefining the future of AI include:

Drishti Eyecare*

Julia Computing is working on building a Deep Learning neural network software for Drishti Eyecare in India that will identify patients in rural India that have treatable abnormalities and can be brought to urban centers before the loss of vision.

Celeste

Intel, Lawrence Berkeley National Laboratory (Berkeley Lab)*, UC Berkeley*, MIT University* and Julia Computing are together working on this project that aims to create an astronomical catalog of the visible universe through scalable Bayesian inference.

Aviva*

A number of large investment banks and insurers are now using Julia for risk analytics, portfolio management, real-time trading and time series analysis. Aviva*, a global insurance company, uses Julia for all its regulatory compliance in the UK and Europe. The company replaced an old IBM* algorithmic system with Julia which is faster, takes lesser space and is cost-effective.

Autonomous Cars

The Autonomous Car Group at the University of California, Berkeley has produced a model car that runs Julia for doing all of its path planning. The model car can drive around avoiding obstacles and even drift park on its own. The same software supported by Julia has been put on a full sized car and is in the research and development stage.



Going ahead

Julia Computing has already seen great success with a large number of organizations making use of Julia. Julia Computing looks forward to a significant scale-up of engagement with Intel owing to Intel's investment in Low Level Virtual Machine (LLVM) and adoption of these technologies on the cloud. With this collaboration, Julia Computing plans to conceptualize many more breakthrough projects similar to Celeste. (see 'Driving Innovative Usages for AI' box for more details on Celeste).


"We expect to significantly scale our engagement with Intel by making Julia available on new hardware as soon as they come out. We see many more users and applications coming on board with Intel and Julia Computing especially supercomputing centers," says Viral Shah.

Intel, which has a vision to accelerate the AI revolution, aims to be at the forefront of AI with Julia Computing. From the business point of view, Julia Computing is strategizing to leverage Intel's vast sales force and connections to address a larger audience in the future ahead.

"We see a lot of demand in the areas of AI and Machine Learning. Even in finance, insurance or bioscience, people are taking our help to develop algorithms using Deep Learning and Machine Learning for solving problems in these verticals."



Deepak Vinchi
Co-Founder and COO,
Julia Computing



Julia is an Open Source Language.

All the views presented are independent views of the individuals or organizations featured in the white paper, and do not in any way represent the views of Intel.

¹IDC, Intel, United Nations - <http://www.intel.com/content/www/us/en/internet-of-things/infographics/guide-to-iot.html>

²IDC Study: Digital Universe in 2020 - <https://www.emc.com/collateral/analyst-reports/idc-the-digital-universe-in-2020.pdf>

© 2017 Intel Corporation. Intel, the Intel logo, Intel Xeon, and Intel Xeon Phi are trademarks of Intel Corporation in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.