SOLUTION BRIEF

Artificial Intelligence
Data Science Model Building

SparkCognition Accelerates Automated Model Building with Intel® AI Technologies

2nd generation Intel Xeon® Scalable processors and Intel MPI library help boost Darwin model building application by up to 44.9X[1]

INTEL® AI BUILDERS MEMBER

Darwin™ by SparkCognition

Modern businesses are data driven. Many are turning to artificial intelligence (AI) to gain insights from their business and industry data. AI is based on a logical process—an algorithm—derived using various networks: Convolutional Neural Networks (CNNs), Deep Stacking Networks (DSNs), and others. Which network to use depends largely on the problem being solved and type of data that eventually will be inferenced.

The field of data science includes analyzing large amounts of data, defining the problem, realizing the best approach and network, and creating a model for inferencing data. All these tasks are complex and time-consuming, and they require a unique combination of analytical and technical skills. With AI becoming an important business driver, data scientists are in demand. A recent study expects the field to grow 28 percent by 2020.[2]

At the heart of developing AI applications is model building, a time-consuming, multi-stage process, which can involve significant computing capability. Once the model is deployed and inferencing data, the process of further training the model from results of the deployed model means going back into the process with new data results. Automated model builders accelerate this process.

Automate and Accelerate Model Building with Darwin

Intel AI Builder partner SparkCognition has an automated machine learning (AutoML) solution, Darwin, that accelerates data science at scale. Darwin uses a patented neuroevolutionary approach to custom-build models that perfectly fit the data, giving data scientists and developers the right tools to quickly prototype use cases and achieve results faster than traditional data science methods.

Darwin enables data professionals to harness the power of machine learning regardless of their level of expertise, by augmenting the skills of data science and analytics teams to prototype use cases faster, extract insights, and operationalize models at scale. Darwin's capabilities include:

Ability to Expedite Data Preparation Tasks—Darwin can proactively uncover problems within the dataset and interactively provides advice on how to address them, ensuring the data is ready for the automated model-building process.

Creation and Management of Thousands of Models— Darwin enables data science and analytics teams to build, deploy, and maintain all their models at scale from a single productive environment that fully integrates with an existing toolchain through SDK and API options.

Control of the Whole Process from Data to Model—Allows complete governance over methods used to prepare the data and selection of optimization functions, helping data scientists to influence the results they want while preserving the accuracy and performance of models.

Unparalleled Accuracy Through Deep Learning—Darwin's neuroevolutionary process specializes in the search and auto-tuning of neural architectures based on the intricacies of the data, defining and exposing every aspect of the network and resulting in more accurate, sophisticated models.

Handling of Complex Temporal Relationships—Darwin uses long short-term memory (LSTM) and temporal convolutional network (TCN) architectures to capture complex relationships over time and exploit them to make more accurate predictions.

True Generalization to Address the Unknown—Darwin's neuroevolutionary process is based on a fitness function. It quickly adapts to changes in upcoming data to achieve maximum accuracy under dynamic circumstances, evolving to maintain complexity, efficacy, and efficiency.

Darwin is delivered via the Google Cloud Platform (GPC),[3] providing storage and compute capabilities to accommodate large and complex data sets that demand scalable computing. It enables data science and analytics teams to build, deploy, and maintain all of their models from a single productive environment that fully integrates with an existing toolchain through SDK and API options.

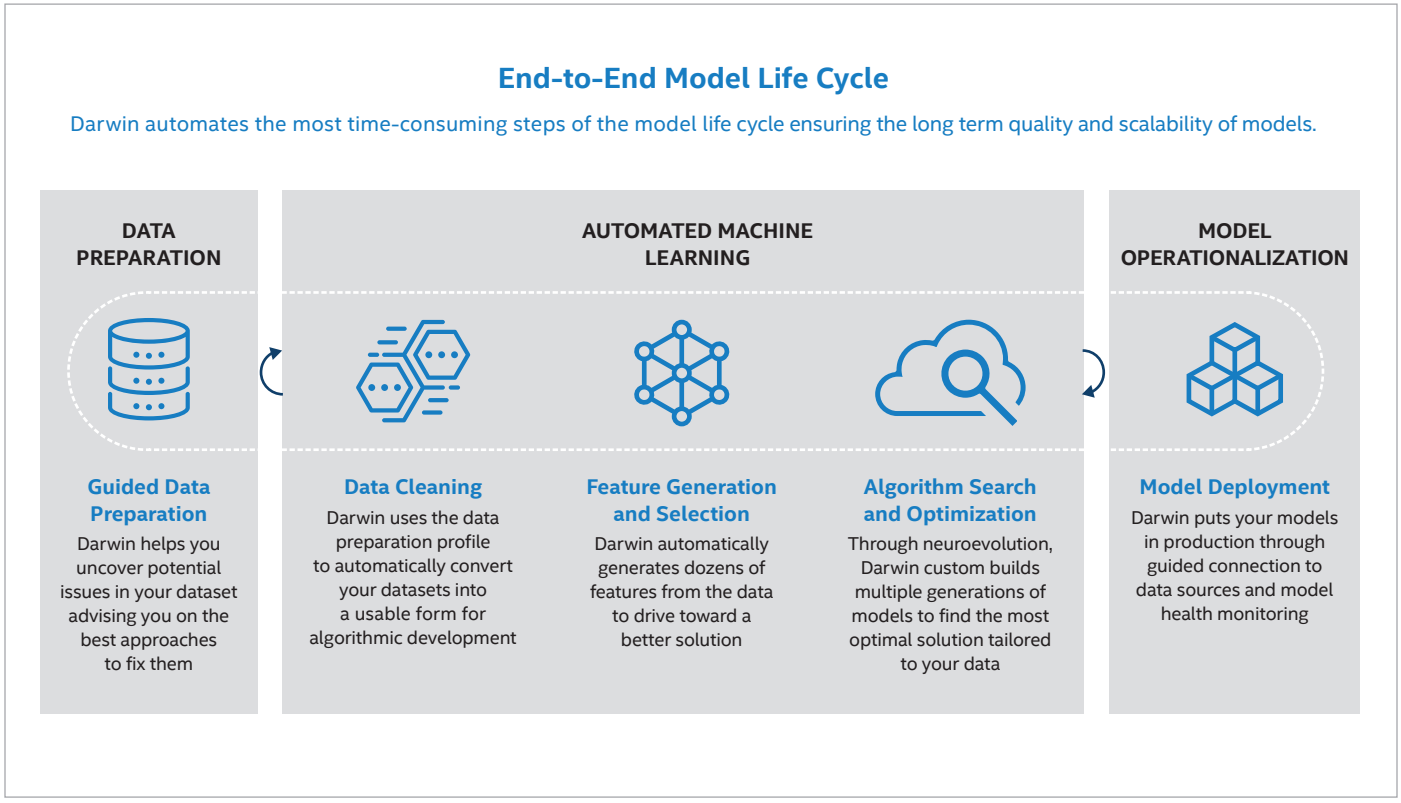

Figure 1. Darwin automates much of the model building process, accelerating data science at scale.

Accelerating Darwin Up to 44.9X[1] with Intel MPI Library and 2nd Gen Intel Xeon Scalable Processors

Darwin has a complex neural architecture search process that performs architecture and hyper-parameter optimization on top of model training, resulting in a very compute-intensive process. Most of the time in model building is spent on training various individual models, wherein the search and optimization processes are very time consuming.

Darwin was developed in Python using the PyTorch deep learning library. While the code was written to take advantage of multiple processors and GPUs, SparkCognition was able to further improve performance using Intel technologies and by simultaneously running multiple processes in parallel through distributed computing using the Message Passing Interface (MPI).

PyTorch's multiprocessing routines are not optimized for distributed processing on Intel architecture. The Intel® Distribution of Python and the Intel Message Passing Interface library (Intel MPI library) include optimizations that allow applications to take advantage of Intel processor features and run efficiently in a distributed environment.

To accelerate performance, SparkCognition developers added a mere five lines of code to Darwin for the application to spawn a new distributed process (called an MPI rank) for each model being built and optimized. This approach made Darwin extremely scalable to accommodate as many models as are available in Darwin's model portfolio.

When run on 2nd Generation Intel Xeon® Scalable processors with these optimizations incorporated, model training time was reduced by 97.7 percent, a speedup of 44.9X.[1]

Darwin could process models faster with the optimizations on Intel architecture. Parallelizing the entire training process across model types allowed Darwin to more effectively leverage the hardware and achieve this level of performance.

Darwin now uses Intel Processors to supplement GPUs in model training on the GCP for their automated model building services.

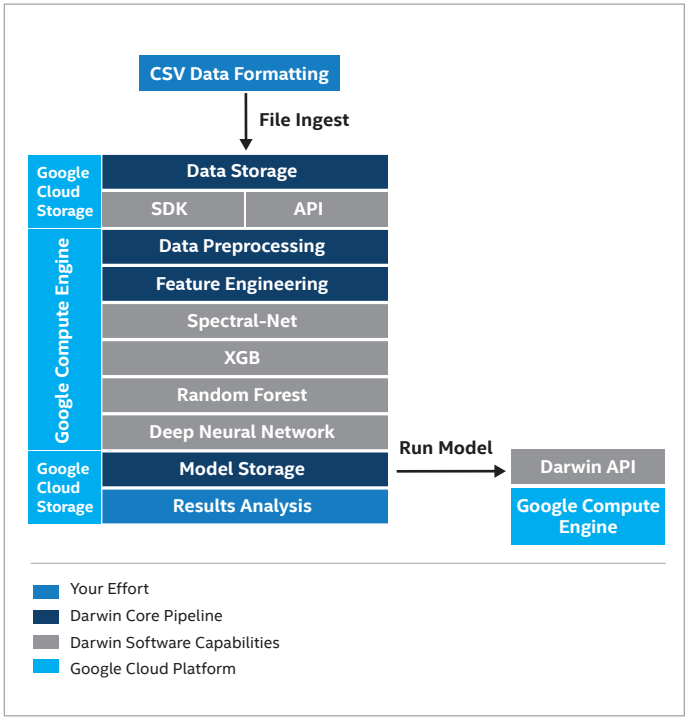

Figure 2. Darwin's automated model building service is delivered via the Google Cloud Platform (GCP).

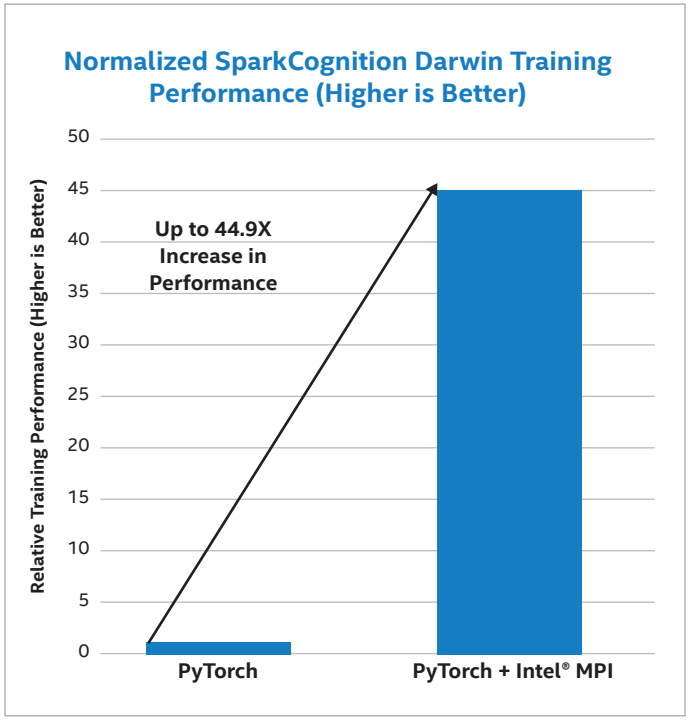

Figure 3. Model building and training with Intel technologies improved performance 44.9X.

Conclusion

Darwin augments the skills of data scientists with tools and technologies that accelerate model building to help democratize the use of AI across more industries. Working with the Intel AI Builders program to optimize their code for Intel architecture, SparkCognition improved Darwin performance by 44.9X. With this level of speedup from Intel optimizations, Darwin makes model building incredibly fast for original model runs and ongoing maintenance.

Faster model building accelerates downstream tasks, such as problem solving and tuning, helping to increase overall efficiency. Quicker model building can mean rapid deployments and faster insights from inferenced data. When run in a cloud environment, shorter run times compared to non-optimized automated machine learning systems, can help reduce cloud computing costs, or allow more models to be built in the same amount of time/costs.

For data-driven business leveraging the benefits of AI, Darwin helps companies attain insights faster and potentially at less cost with accelerated, automated model building.

For more information about the Darwin, visit https://www.sparkcognition.com/product/darwin/.

Learn more about the Intel AI Builders program at https://builders.intel.com/ai/.

To learn more about how SparkCognition collaborated with the Intel AI Builders program to accelerate AI, watch https://builders.intel.com/ai/social-hub/video/sparkcognition-accelerates-ai-innovation-with-intel.

SparkCognition is a member of the **Intel AI Builders Program**, an ecosystem of industry-leading independent software vendors (ISVs), system integrators (SIs), original equipment manufacturers (OEMs), and enterprise end users, which have a shared mission to accelerate the adoption of artificial intelligence across Intel platforms.

1. SparkCognition Darwin Throughput Performance on Intel® Xeon® Gold 6248 Processors:

   NEW: Tested by Intel as of 08/13/2019. Dual Intel® Xeon® Gold 6248 Processor (2 socket), 20 cores HT On Turbo ON Total Memory 384 GB (12 slots/32GB/ 2933 MHz), BIOS: SE5C620.86B.00.01.0015.110720180833 (ucode: 0x200004d), Ubuntu 18.04.2 LTS, Deep Learning Framework: PyTorch 1.1.0, Libraries: Intel® MKL-DNN v1.0, Intel® MPI (Message Passing Interface) 2019 Update 3 Build 20190214, BS=1000, N Models=40, N Processes=40, custom generated dataset

   BASELINE: Tested by Intel as of 08/13/2019. Dual Intel® Xeon® Gold 6248 Processor (2 socket), 20 cores HT On Turbo ON Total Memory 384 GB (12 slots/ 32GB/ 2933 MHz), BIOS: SE5C620.86B.00.01.0015.110720180833 (ucode: 0x200004d), Ubuntu 18.04.2 LTS, Deep Learning Framework: PyTorch 1.1.0, Libraries: Intel® MKL-DNN v1.0, BS=1000, N Models=40, N Processes=40, custom generated dataset

2. https://www.forbes.com/sites/louiscolumbus/2017/05/13/ibm-predicts-demand-for-data-scientists-will-soar-28-by-2020/#4c075d9c7e3b

3. On-premise deployment options are also available.

Performance results are based on testing as of dates shown in configuration and may not reflect all publicly available security updates. No product or component can be absolutely secure. See configuration disclosure for details.

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice. Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.

Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit: http://www.intel.com/performance

0120/AU/HBD/PDF ♻ Please Recycle 342166-001US