Solution Brief

Manufacturing
Computer Vision (Worker Safety)

Flutura Cerebra VISION Intelligence Enhances Worker Safety

Intel® Optimizations for TensorFlow and Intel Distribution of OpenVINO™ toolkit create a faster, more efficient vision computing pipeline

Intel® AI Builders Member

Flutura
Actions! Not just Insights

Worker safety at construction sites and factories is among the most critical aspects of any organization's operations. Industrial injuries, such as falls, contact with harmful objects and chemicals, and others, result in lost workdays, lost revenues, and lost lives. Many injuries are preventable simply by ensuring compliance with workplace safety requirements. These include use of Personal Protective Equipment (PPE), such as hard hats, necessary for employee and visitor safety.

According to the U.S. Bureau of Labor Statistics (BLS), 84 percent of all workers who suffered head injuries were not wearing head protection.[1] Many companies use CCTV throughout their facilities and can forensically determine what led to the cause of an injury. However, real-time monitoring and analysis of the industrial environment help alert personnel to anomalies, like the absence of hard hats or safety vests in a job site or industrial arena. Early detection of violations of safety equipment can help personnel respond proactively and reduce injuries and fatalities.

Detecting Safety Violations

In work environments, monitoring personnel safety compliance through manual detection is difficult, error-prone, and time consuming. Reactive monitoring leads to inefficiencies, while manual monitoring discourages personnel, reducing morale. Video monitoring in large plants requires thousands of cameras and hundreds of people watching video monitors.

The Flutura Cerebra VISION Intelligence platform can help detect and alert personnel to potential real-time safety violations in an industrial environment. Cerebra VISION Intelligence analyzes video data from thermal cameras, industrial cameras, CCTVs, drones, and other sources. The primary objective is to address quality, worker safety, and production challenges across industries.

Cerebra vision focuses on object detection, event detection, and triangulation of sensors and video anomalies. Plant applications include production sequence monitoring, surface quality defect detection, PPE compliance in EHS, and others in the domains of safety, surveillance, and warehouse. A key feature is hard hat violation detection (see Figure 1). Cerebra VISION is ideal for manufacturing and engineering companies, or any industry with large-scale operations.

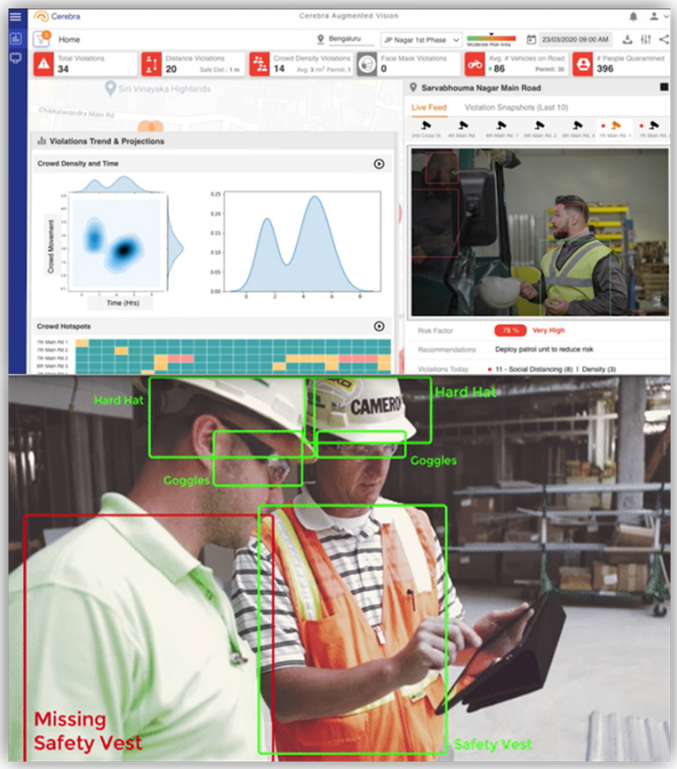

Figure 1. Cerebra VISION Intelligence platform inferences video imagery and identifies safety equipment non-compliance events.

Flutura developed their machine learning-based vision solution using the TensorFlow framework on a RetinaNet network with ResNet50 as the backbone. Their datasets are very large, which means they cannot run efficiently on discrete accelerators. Thus, the entire solution uses Intel® processors with large memory capacities for both training and inferencing.

Flutura is a member of the Intel AI Builders program. As a program member, developers had access to many tools and expertise provided by Intel to help them optimize their algorithms and training on Intel processors. Optimizing the Cerebra VISION Intelligence platform improved training and inference performance and thus offers assistance in the early detection of PPE violations on-site.

Intel Optimizations Make Cerebra VISION Run 2.7X Faster[2]

As an Intel AI Builder member, Flutura used Intel's development environment with access to 2nd Gen Intel Xeon® Scalable processor-based compute clusters and the latest Intel-optimized frameworks and libraries. These resources gave Flutura engineers tools for developing high-performance machine learning solutions to train their Deep Neural Networks (DNN).

Intel AI technologies include integrations into processor architecture that accelerate machine learning training and inferencing, such as Intel Deep Learning Boost (Intel DL Boost) and Intel Advanced Vector Extensions 512 for Deep Neural Networks.

Flutura developers engaged with Intel experts to leverage the Intel AI stack to tune their algorithms for Intel architecture. Flutura integrated the following Intel technologies in their solution (Figure 2):

- Intel Optimizations for TensorFlow, which includes the Intel MKL-DNN
- Intel Distribution of OpenVINO™ toolkit
- Intel Distribution for Python
- Optimized libraries for Keras, Scikit, and NumPy

The Intel Distribution of OpenVINO toolkit allows developers to create applications and solutions that emulate human vision. Based on convolutional neural networks (CNN), the toolkit extends computer vision workloads, such as the Cerebra VISION Intelligence platform, across Intel hardware (including accelerators) and maximizes performance.

Flutura engineers used Intel's development environment for optimization of model performance and accuracy, and to benchmark performance on Intel processor-based infrastructure. The Intel engineering team helped Flutura developers select the right Intel technologies and Intel hardware configurations and provided many recommendations, including:

- Number of cores, sockets, and processors used for model training and processes to run.
- How to configure the cluster to run multiple deep learning models on various CPUs for optimized processing.
- Knowledge on various clusters to benchmark a chosen infrastructure for different use cases and deployments.

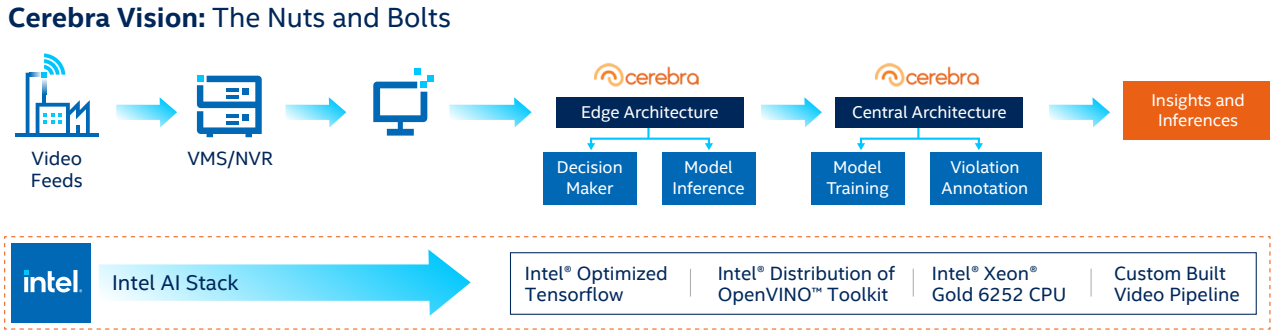

Figure 2. Cerebra VISION Intelligence on Intel AI technologies.

The collaboration helped Flutura to optimize their algorithm and define the best hardware specifications for various training and inferencing needs. The result achieved a model that ran 2.7x faster[2] inferencing with approximately 90 percent core utilization with Intel AI technologies and processors compared to a non-optimized model (Figure 3).

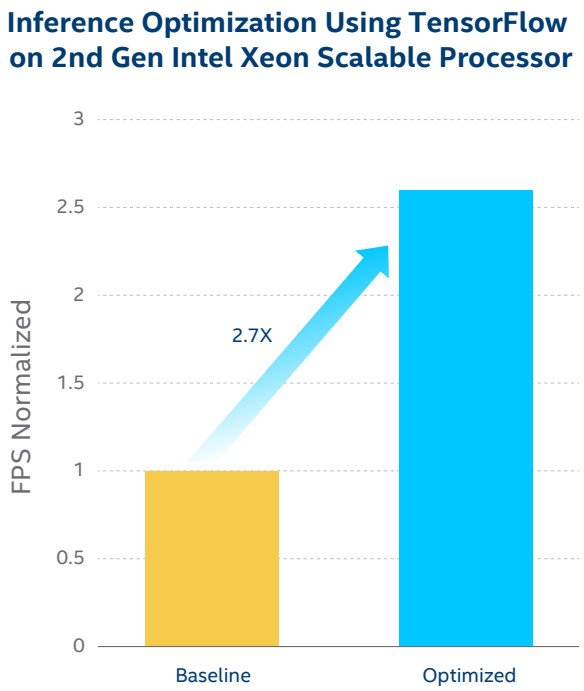

Figure 3. Optimized Flutura algorithms run 2.7x faster[2] with Intel Optimizations for TensorFlow on 2nd Gen Intel Xeon Scalable processors.

The faster platform was achieved with little change to the original software simply by using Intel optimizations and libraries, plus the Intel Distribution of OpenVINO toolkit. One of the major roadblocks when it comes to designing a very deep neural network is setting up the environment to train the model. With the help of Intel optimized toolkits and frameworks, this task was done in less time, which allowed Flutura developers to focus on important functionality of the neural network instead of hardware and software adaptations. The Intel tools and Intel DevCloud proved to be useful resources for Flutura developers.

Customer Case Study:

VISION Intelligence Reduces Workplace Accidents

Flutura worked with one of the largest adhesive manufacturing companies to deploy Cerebra VISION Intelligence to monitor regulatory safety compliance. They started with smaller use cases, such as PPE detection. Then, they evaluated the Cerebra VISION Intelligence solutions for mobile usage, crowd gathering, and loitering in restricted zones.

The model was trained and tuned using the Intel AI stack. The results met all the customer's Service Level Agreement (SLA) requirements with very high accuracy and low latency. This helped the manufacturing conglomerate reduce the compliance violations and legal costs due to workplace accidents.

The pilot was initiated in one plant. It is being scaled across all the facilities, globally. Flutura Cerebra VISION Intelligence with Intel optimizations made the workplace a safer and more productive environment.

Conclusion

Working with Intel, Flutura developers were able optimize their algorithms and configure the right hardware for various client needs—whether the deployment would be on-cloud or on-premise. The right hardware choices have improved the efficiency and computational speed. Using Intel Distribution for OpenVINO toolkit helped Flutura scale up their solution implementation and reduce model inferencing time without compromising accuracy. The more efficient video pipeline can be deployed on Intel processor-based edge computing systems for model inferencing using 2nd Gen Intel Xeon Scalable processors in multi-socket and single-socket platforms.

With optimizations complete, Flutura developers can further innovate by focusing on:

- Diverse Use Cases—Train the model to accommodate more use cases, especially on production monitoring, quality inspection, and security risk management.
- Scaling Up—In order to broaden their market and reach a larger audience.
- Improving the Self Learning and Feedback Loop—Further enhance the AI self-learning model with user feedback input and automatic model training.
- Unification—Move towards unified inference platforms to meet multiple performance needs for different use-cases and customers.

For more information about Cerebra VISION Intelligence platform, visit **www.flutura.com/video-analytics**

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

Intel® AI Builders Member

Flutura is an AI Platform Company focused on improving two core business objectives—Asset Uptime and Operational Efficiency. We do this with Cerebra, our AI Platform tuned for IIoT in Oil & Gas, Specialty Chemicals, and Heavy machinery manufacturing industries, powering connected asset and connected operations use cases.

---

[1] https://www.hexarmor.com/posts/the-hard-truth-about-safety-helmet-injuries-and-statistics

[2] RetinaNet (ResNet50) model inference on TensorFlow* Throughput Performance on Intel® Xeon® Gold 6252 Processor:

NEW: Tested by Flutura as of 5/29/2020. 2 socket Intel® Xeon® Gold 6252 CPU @ 2.10GHz, 24 cores HT On Turbo ON Total Memory 192 GB (12 slots / 16GB / 2666 MTs / DDR4 DIMM), BIOS: SE5C620.86B.02.01.0011.032620200659 (ucode: 0x400002c), Ubuntu 18.04.4 LTS, 5.3.0-53-generic, Deep Learning Framework: Regular TensorFlow* 1.15.2, Eigen, RetinaNet (ResNet50): Multi Layer, BS=5, customer data, 1 instance/2 socket, Datatype: FP32

BASELINE: Tested by Flutura as of 5/29/2020. 2 socket Intel® Xeon® Gold 6252 CPU @ 2.10GHz, 24 cores HT On Turbo ON Total Memory 192 GB (12 slots / 16GB / 2666 MTs / DDR4 DIMM), BIOS: SE5C620.86B.02.01.0011.032620200659 (ucode: 0x400002c), Ubuntu 18.04.4 LTS, 5.3.0-53-generic, Deep Learning Framework: Regular TensorFlow* 1.15.2, Eigen, RetinaNet (ResNet50): Multi Layer, BS=5, customer data, 1 instance/2 socket, Datatype: FP32

Intel is committed to respecting human rights and avoiding complicity in human rights abuses. See Intel's Global Human Rights Principles information at https://www.intel.com/content/www/us/en/policy/policy-human-rights.html. Intels products and software are intended only to be used in applications that do not cause or contribute to a violation of an internationally recognized human right.

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

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.

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

Printed in USA 1120/AU/HBD/PDF ♻ Please Recycle 344586-001US