

Seeing is Believing in Embedded Machine Vision

[Electronic Design](#)

[Stephen Turnbull](#)

Thu, 2016-07-28 14:53



Talk to your friends and family about “machine vision” and you might get strange looks, possibly followed by a discussion about a movie they saw where robots become self-aware and dangerous. Fortunately, the reality isn’t nearly as ominous. Maybe the industry should consider a description of the technology segment that’s more user-friendly, because this cutting-edge field has dramatic and positive potential for embedded applications.

At its core, machine vision is simply leveraging the information available in an image to make a decision about what to do next with the object in the image. A simple pass/fail examination of a product on the assembly line or before shipping is one of the more simple examples.

Printed-circuit-board (PCB) inspection is a common use case, where an image of a master, correctly populated board can be quickly and easily compared to production PCBs as they move from an automated pick-and-place system onto the next stage. This is an invaluable step of quality assurance and scrap reduction that the human eye and brain could never consistently repeat hundreds or even thousands of times per day.

Related

[Conquer the Challenge of Integrating Efficient Embedded Vision](#)

[Intelligent Imaging Yields Smarter Machine Vision Solutions](#)

[Q&A: Embedded Vision Alliance Founder Jeff Bier on Visual Intelligence’s Future](#)

Processing Needs of Machine Vision

As the resolution of image-capturing systems increases, so too does the potential for machine vision. That’s because the detail available for evaluation increases at a corresponding rate. Smaller and smaller subsets of visual information can be evaluated against a master template, intensifying the burden of the system processor to churn through the data and quickly deliver a decision on next steps (pass/fail, hold, return to start, etc.).

Vegetable grading is a case where simple size and pass/fail for product quality is not optimal since product standards are different from country to country and product quality varies over the course of a season. To be able to minimize scrap for the producer and still maintain the right quality for the customer, more optimal algorithms are needed for quality grading—a nearly impossible task for the human eye and brain.

One company addressing this application is [Qtechnology](#) of Denmark. The company delivers smart cameras for vegetable grading of production volumes up to 25 tons per hour, which requires analyzing more than 250,000 products from around 500,000 images. At 6.2 MB for each image, this particular case requires analyzing more than 2.5 terabytes of image data an hour per machine, a colossal amount of information to process. This amount

lata would take more than six hours of transfer time on a single Gigabit Ethernet connection.

To solve this with simpler algorithms would require multiple stages and cameras, lighting in the machine, more real estate in the factories, etc. The alternative is to apply extensive processing power, either as a centralized processing unit over high-bandwidth connections or distributed processing with smart cameras. Data would be processed in real time directly in the camera, with only the results per product delivered to the final mechanical grading system.

To address different image-capture technologies, Qtechnology relies on exchangeable heads with different sensor arrays to go with the smart-camera systems. Its hyperspectral imaging head, for example, allows for non-destructive detection of food quality and safety.

In standard vision systems, food quality and safety are usually defined by external physical attributes like texture and color. Hyperspectral imaging is giving the food industry the opportunity to include new attributes into quality and safety assessment, like chemical and biological attributes for determining sugar, fat, moisture, and bacterial count in the products.

In hyperspectral imaging, three-dimensional image cubes of spatial and spectral information are obtained from each pixel. More spectral characteristics give better discrimination of attributes and enable more attributes to be qualified. The image cubes include the intensity (reflected or transmitted light) of each pixel for all acquired wavelengths of light, which results in each image cube containing a mass of information. This amount of data represents an exponential increase in the computational challenge to extract qualitative and quantitative results for product grading in real time.

Applying Heterogeneous Compute

Supporting such processing demands today and into the future requires high-performance, scalable processing.

Qtechnology uses an accelerated processing unit (APU) in its smart-camera platforms that combine the GPU and CPU on the same die, enabling the system to offload intensive pixel data processing in the vision applications to the GPU without high-latency bus transactions between processing components. This lets the CPU serve other interrupts with lower latency, helping improve the real-time performance of the entire system and addressing the rising processing demands of modern vision systems.

The pairing of a different processing engine on a single die or in a system to apply the right processing power to the problem is at the heart of heterogeneous computing. The [Heterogeneous System Architecture \(HSA\) Foundation](#) was founded in 2012 specifically to help the industry define open specifications for processors and systems that leverage all available processing elements to improve processing efficiency.

The GPU is a massively parallel engine that can apply the same instructions across large data sets (in this case, pixels) at the same time; that's exactly what's needed to deliver a 3D game on your favorite gaming console or PC. Coincidentally, this is also exactly what is needed for machine vision.

Performance can further be increased by pairing the APU with an external, discrete GPU in a Mobile PCI Express Module (MXM) form factor. When doing so, companies are able to add further GPU processing resources to support even more intensive vision tasks when needed.

Software is a critical part of the equation. With HSA, the whole processing platform can be governed by a standard Linux kernel, which requires only modest development support with each new kernel release. The Yocto Project, an open-source collaboration project, provides templates, tools, and methods to help users create custom Linux-based systems for embedded products.

A huge ecosystem support for x86 allows companies to tap open-source and third-party image-processing libraries such as OpenCV, [Mathworks'](#) MATLAB, and HALCON. Debugging tools, latency analyzers, and profilers (perf, ftrace) are also widely available.

Machine vision is the latest example of silicon's processing power being applied to help reduce costs, speed manufacturing, increase quality, and provide a wealth of benefits to the world in myriad different businesses and applications. This positive economic, cultural, and personal impact is becoming widely available thanks to the innovation of the industry and insightful thinking of embedded engineers.



Looking for parts? Go to sourceesb.com

Source URL: <http://electronicdesign.com/embedded/seeing-believing-embedded-machine-vision>