



Aurora™ Deep Learning



Zebra
Aurora™
Deep Learning

The Aurora™ Vision for OEM Portfolio

Intuitive Software for Industrial Image Analysis

Zebra is now a leading provider of user-friendly machine vision software for industrial image analysis. Our comprehensive Aurora™ Vision for OEM portfolio of graphical software helps users easily create custom machine vision applications.

Enhanced and optimized by machine vision experts for nearly 15 years, these world-class software products – Aurora™ Vision Studio, Aurora™ Vision Library, Aurora™ Deep Learning – offer state-of-the-art, industrial reliability, quality, and speed. They have been used by machine builders, vision system integrators, robotic designers, and industrial end users worldwide to enable rapid development of vision applications in verticals ranging from food production and retail to agriculture and healthcare.

Ready-made tools for dataflow and comprehensive image analysis filters are all hardware-agnostic – to suit your customers' specific needs. They enable your engineers to quickly and easily construct powerful, customized machine vision applications to augment your operations.

Introducing Aurora™ Deep Learning

Deep Learning to Further Enhance Solution Quality

This add-on product within our comprehensive Aurora™ for OEM software portfolio offers a complete set of industrial-quality deep learning tools which can be used to solve problems that are far too complex for traditional machine vision algorithms and further enhance the outputs of Aurora™ Vision Studio and Aurora™ Vision Library software.

Features & anomaly detection, advanced optical character recognition (OCR) technology, and deep learning capabilities combine to recognize the way real-world images are used – and evolve accordingly.


Internally, Aurora™ Deep Learning uses large neural networks designed and optimized by our research team for use in industrial inspection systems.

Together with Aurora™ Vision Studio, Aurora™ Deep Learning constitutes a complete solution for training and deploying modern machine vision systems.



Learns from just a few samples

Typical applications require between 20 and 50 images for training. More sample images are better but the Aurora™ Deep Learning software learns key characteristics from a limited training set and then generates thousands of new artificial samples for effective training.



Works on GPU and CPU

A modern GPU is required for effective training. At production, you can use either GPU or CPU. Using GPU will typically be 3-10 times faster (with the exception of Object Classification, which is equally fast on CPU).



High performance guaranteed

Typical training time on a GPU takes usually 5-15 minutes. The inference time varies, depending on the tool and hardware – from between 5ms and 100ms per image. The highest performance is guaranteed by using Aurora™ Weaver, our industrial inference engine.



All-In-One Software Package

The Aurora™ Vision for OEM portfolio offers the most comprehensive approach to the development of custom machine vision applications:



2D & 3D algorithms



C++ and .NET libraries



Rapid development environment



HMI Designer



Technical support and know-how



Deep Learning

Aurora™ Deep Learning vs Traditional Machine Vision

Aurora™ Deep Learning is a new reliable solution for machine vision problems that could not have been solved before. There are, however, applications that can still only be realized with traditional methods. How do you know which approach is better? Here is a quick guide:

Deep Learning

Typical applications:

- Surface inspection (cracks, scratches)
- Food, plant, wood inspection
- Plastics injection molding
- Textile inspection
- Medical imaging

Traditional machine vision

Typical applications:


- Dimensional measurements
- Code reading
- Presence or absence checking
- Location of fiducials on PCB
- Print inspection

Typical characteristics:

- Deformable objects
- Variable orientation
- Customer provides vague specifications with examples of good and bad parts
- Reliability: 99%

Typical characteristics:

- Rigid objects
- Fixed orientation
- Customer provides formal specifications with tolerances
- Reliability: 100%



Hardware Requirements

Aurora™ Deep Learning can work on a standard industrial PC. But for better performance, we recommend using modern GPU boards from the NVIDIA® GeForce® and Tesla series with compute capability 3.5 or higher.



Training Interface for end users

- Aurora™ Deep Learning allows end users who are non-vision experts to retrain a Deep Learning model on a factory floor.
- Users of Aurora™ Vision Library can create their own training interface for end users using the C++ API.

Training Procedure

1

Collect and normalize images

- Acquire between 20 and 50 images, both good and bad, representing all possible object variations, and save them to your hard drive.
- Make sure the object scale, orientation and lighting are as consistent as possible.

2

Training

- Open Aurora™ Vision Studio and add one of the Deep Learning tools.
- Open an editor associated with the tool and load your training images.

3


Execute

- Run the program and see the results.
- Go back to Step 1 or 2 until results are fully satisfactory.

Training and Validation Sets

In Aurora™ Deep Learning, as in all fields of machine learning, it is very important to follow the correct methodology. The most important rule is to separate the Training set from the Validation set.

The Training set is a set of samples used for creating a model. We cannot use it to measure the model's performance, as this often generates overly optimistic results. Thus, we use separate data – the Validation set – to evaluate the model. Our Deep Learning tool automatically creates both sets from the samples provided by the user.



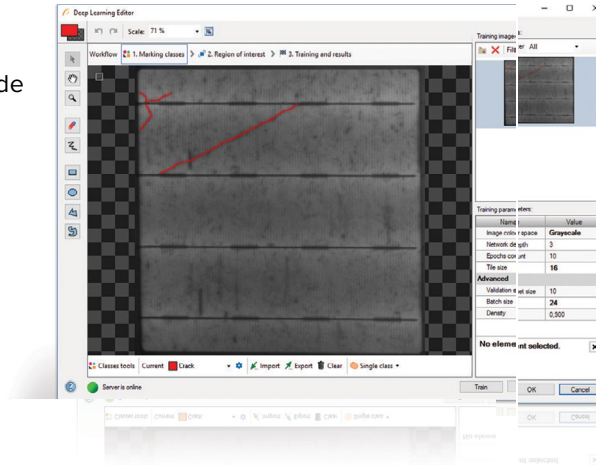
Application Examples

Feature Detection (Supervised)

In the supervised detection mode, the user needs to carefully label pixels corresponding to defects in the training images. Aurora™ Deep Learning then learns to distinguish good and bad features by looking for their key characteristics.

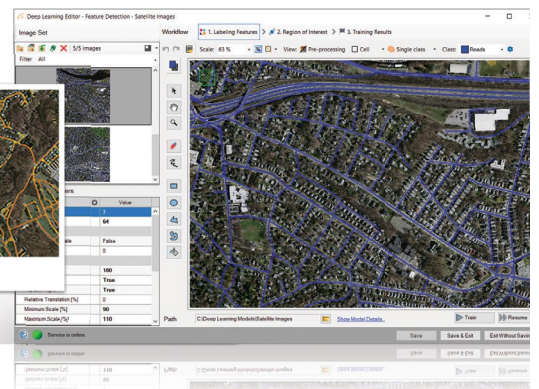


Photovoltaic Panels Inspection

In this application, cracks and scratches must be detected on a surface that includes complicated features. Using traditional methods, this requires complicated algorithms with dozens of parameters which must then be adjusted for each type of solar panel. Using Aurora™ Deep Learning, it is enough to train the system in the supervised mode with just one tool.




Satellite Image Segmentation

Satellite images are difficult to analyze as they include a huge variety of features. Nevertheless, Aurora™ Deep Learning can be trained to detect roads and buildings with very high reliability.



Other Examples


Marble cracks



Wood knots



Fabric defects



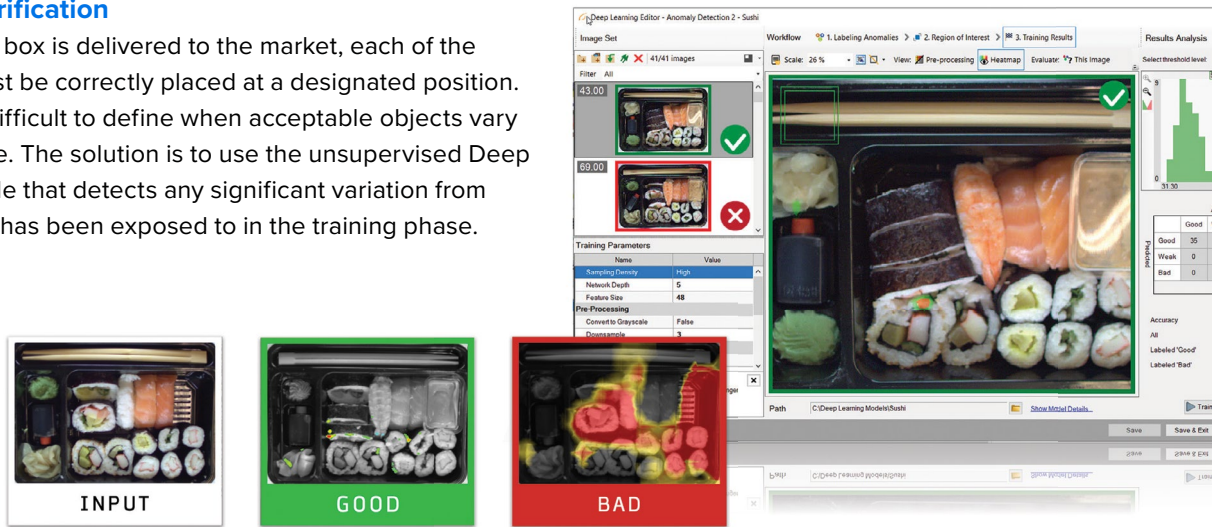
Anomaly Detection (Semi-Unsupervised)

In the semi-supervised mode, training is simpler. A defect is not strictly defined – the tool is trained with good samples and then looks for deviations of any kind.

Aurora™ Deep Learning provides two variants of the Anomaly Detection tool. They are both designed for detecting anomalies, but in a different way. The first uses image reconstruction techniques, while the second performs one-class classification of every part of the input image. When highly precise defect heatmaps are needed – even at the expense of higher computational time – the first variant is recommended.

Package Verification

When a sushi box is delivered to the market, each of the elements must be correctly placed at a designated position. Defects are difficult to define when acceptable objects vary in appearance. The solution is to use the unsupervised Deep Learning mode that detects any significant variation from what the tool has been exposed to in the training phase.



Plastics Injection Molding

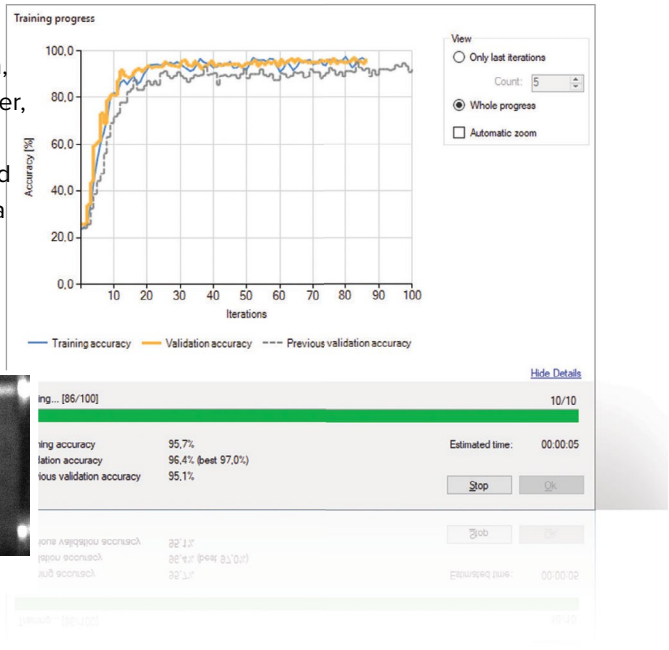


Injection molding is a complex process, with many production problems that might occur. Plastic objects may also include folding or other kinds of shape deviations that are acceptable for the customer. Aurora™ Deep Learning can learn all the acceptable deviations from the provided samples and then detect anomalies of any type when running on a production line.



Object Classification


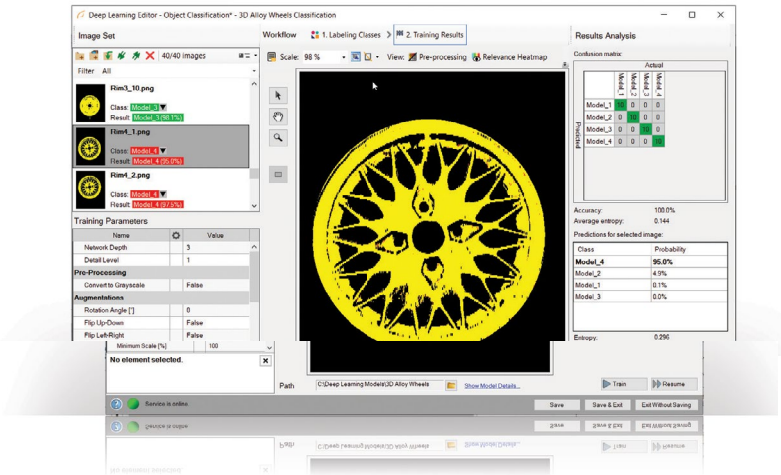
Caps: Front or Back

Plastic caps may sometimes accidentally flip in the production machine. If the customer would like to detect such a situation, the task can be completed using traditional methods. However, it requires an expert to design an algorithm specific to this application. Alternatively, Aurora™ Deep Learning can be used to automatically learn to recognize the front and the back of a cap from a set of training pictures.



3D Alloy Wheel Identification

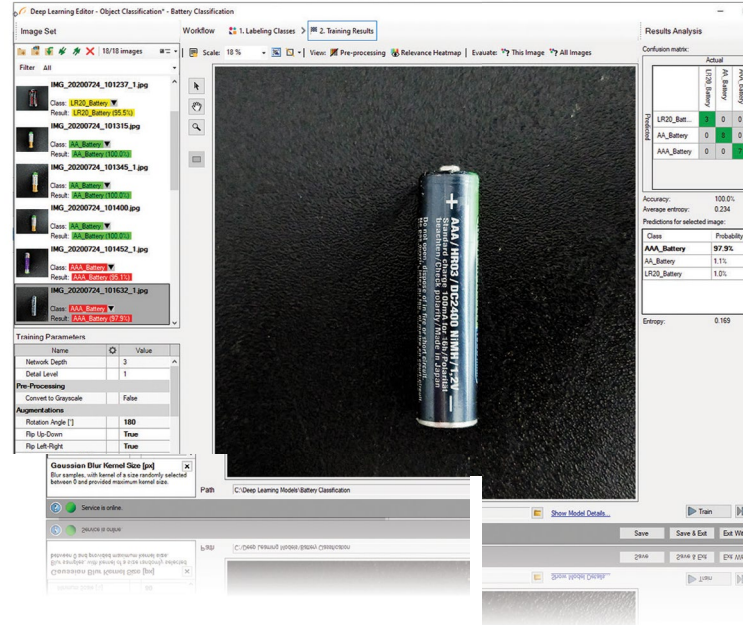
There may be hundreds of different alloy wheel types being manufactured at a single plant. The identification of the particular model among such variety is virtually impossible using traditional methods. Template matching would require a huge amount of time, trying to match hundreds of models, while handcrafting of bespoke models would simply require too much development and maintenance. Aurora™ Deep Learning is an ideal solution, allowing the program to learn directly from sample pictures and come up with reliable results.



Batteries Classification

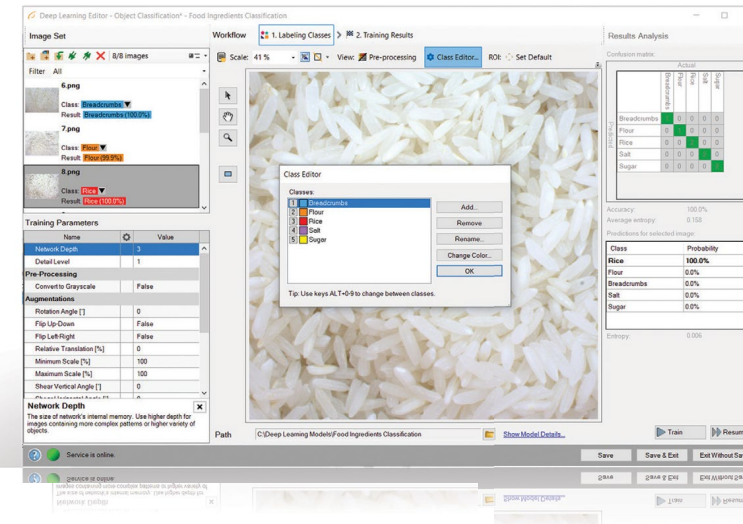
Batteries can be found in every room in the house nowadays. Unfortunately, most of them end up in trash cans and then are taken to landfill sites, where they begin to rot away and may leak dangerous chemicals into the ground, causing soil and water pollution.

Aurora™ Deep Learning tools make classification of used batteries easy. Simply teach your program what selected types of batteries look like and it will classify them automatically. The range of this application is extremely wide – from sorting batteries in big recycling plants to small automatic battery collection containers in the streets.



Food Ingredients Classification

Although it may appear easy at first, especially for a human brain, it is very difficult for a traditional machine vision system to distinguish between sugar and flour when it is being transported at a speed of a few meters per second. In food ingredients packaging systems, customers use Aurora Deep Learning to ensure that the correct material is loaded.



Breadcrumbs



Flour



Rice



Salt



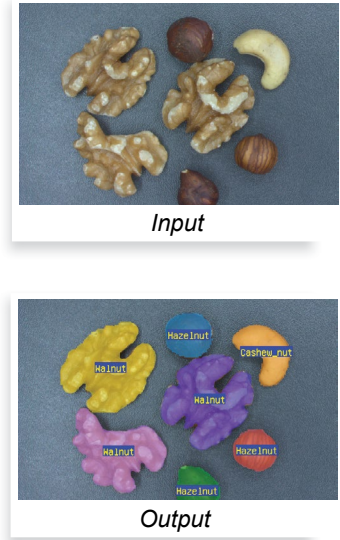
Sugar

Instance Segmentation

The instance segmentation technique is used to locate, segment, and classify single or multiple objects within an image. Unlike the feature detection technique, this technique detects individual objects and may be able to separate them even if they touch or overlap.

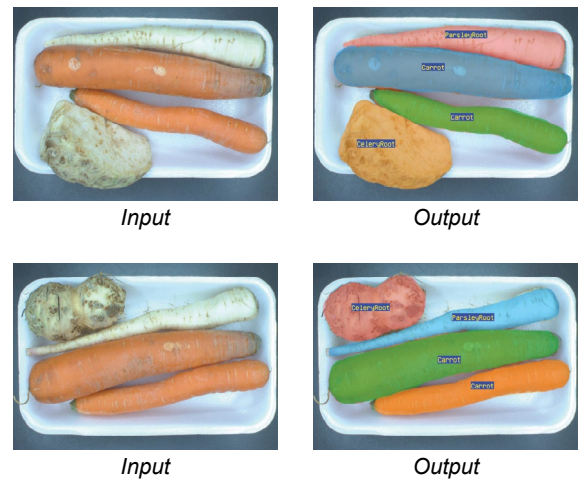
Nuts Segmentation

Mixed nuts are a very popular snack food, consisting of various types of nuts. As the percentage composition of nuts in a package should be in accordance with the list of ingredients printed on the label, the customer wants to ensure that the correct number of each type of nut is going to be packaged. The instance segmentation tool is an ideal solution in such application, as it returns marks corresponding to the segmented objects.



Package Verification

A typical set of soup greens used in Europe is packaged on a white plastic plate in a random position. Production line workers may sometimes accidentally forget to put one of the vegetables on the plate. Although there is a system that weighs the plates, the customer often wants to verify the completeness of the product just before the sealing process. As there are no two vegetables that look the same, the ideal solution is to use deep learning-based segmentation. In the training phase, the customer just has to mark the regions corresponding to selected vegetables.



Point Location

The Point Location tool looks for specific shapes, features, or marks that can be identified as points on an input image. It may be compared to traditional template matching, but here the tool is trained with multiple samples and becomes robust against huge variability in the objects of interest.

Tracing Bees

The task that seems impossible to achieve with traditional methods of image processing can be done using Aurora™ Deep Learning tools. In this case we use them to detect bees. When it is done, we can check whether they are infected by varroosis – a disease caused by parasitic mites attacking the honey bees. The parasite attaches itself to their bodies and we can use a characteristic red inflammation spot to classify the bees according to their health condition. Not only does this example show that it is an easy solution for a complex task, but also that Aurora™ software can be used in many different branches of industry, e.g., agriculture.



Healthy bee




Bee with varroosis

Pick and Place

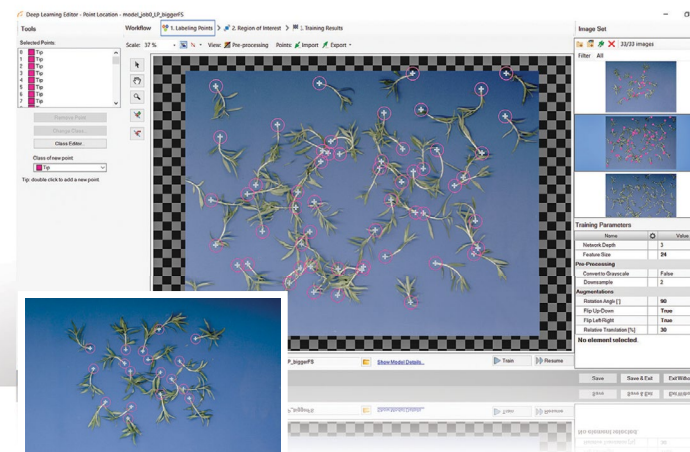
In these applications, we need to guide a robotic arm to pick up items, most typically from a conveyor belt or from a container. A good example of such application is picking small stem cuttings and then placing them vertically in pots. Any inaccuracies in detection may cause them to be planted too deep or upside down, which will result in cuttings not forming roots. Aurora™ Deep Learning tools make it possible to quickly locate the desired parts of the plants and provide accurate results required for this operation.



Location accuracy



Located points



Optical Character Recognition

The optical character recognition (OCR) capabilities within Aurora™ Deep Learning are the answer for challenging character recognition projects, in which complex non-uniform backgrounds, blurred, damaged, distorted, or obscured characters, or reflective metal surfaces make it impossible to use traditional OCR techniques.

The tool comes with a ready-to-use neural network that is pre-trained using thousands of different image samples. It can achieve up to ~97% accuracy straight out of the box, even when dealing with very difficult cases, and enables the user to create a robust OCR application in just a few simple steps – without the need for machine vision expertise.

Key features

- Ready-to-use, comes with a pre-trained neural network.
- Can deal with difficult OCR cases, impossible to achieve using traditional methods.
- Very high accuracy straight out of the box.
- Easy to use, no need for machine vision expertise.
- Works on both NVIDIA GPU and CPU.



To find out more, visit zebra.com/aurora-for-oem



NA and Corporate Headquarters
+1 800 423 0442
inquiry4@zebra.com

Asia-Pacific Headquarters
+65 6858 0722
contact.apac@zebra.com

EMEA Headquarters
zebra.com/locations
contact.emea@zebra.com

Latin America Headquarters
zebra.com/locations
la.contactme@zebra.com