



Aurora Deep Learning™



Zebra
Aurora
Deep Learning™

The Aurora Vision 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™ 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 Vision 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.




All-In-One software package

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




2D & 3D algorithms



C++, NET libraries,
Python



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

Typical characteristics:

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

Traditional machine vision

Typical applications:

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

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 and System-on-Modules with appropriate compute capability (consult the Aurora Vision website for the latest requirements).



Training interface for end users

- Aurora Deep Learning allows end users who are nonvision 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++ or Python APIs.

Training procedure



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.

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.


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-supervised)

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.

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 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

Object location

This technique is used to locate and classify one or multiple objects within an image. The result is a list of the rectangles bounding the predicted objects with the corresponding class predictions and confidence scores. The model can deal with both single class as well as multiple object classes.

The tool returns rectangle regions containing predicted objects and shows their approximate location and orientation. It doesn't return the precise position of the key points of the object or the segmented region because they can be determined by the Locate Points tool instead.

Assembly inspection

Ensuring proper assembly is critical in manufacturing. Our Locate Object tool can be used to validate the orientation of parts, verify the presence of essential components as well as perform completeness checks on complex assemblies like automotive parts. These advanced solutions ensure accuracy, reduce errors, and keep production lines running smoothly.



INPUT




OUTPUT

Object orientation

Determining the orientation of objects is essential in many industries, and Zebra Aurora Vision Studio and Zebra Aurora Deep Learning excel in this task. In fish sorting facilities, for instance, these tools can be used to determine fish orientation for proper processing. Combined with other tools, they can also perform object identification, enabling tasks such as species classification and alignment for automated sorting. This ensures efficiency and accuracy in high-speed operations, helping streamline processes in the fishing industry.



INPUT



OUTPUT

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) and text location capabilities within Zebra 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. It also includes a deep learning-based text location tool, able to precisely pinpoint textual elements within an image for further OCR or processing.

Key features

- Ready-to-use, comes with a pre-trained neural network and different OCR models, including OCR-A and SEMI fonts.
- Provides quality metrics for Optical Character Verification (OCV) cases and character hint input to further improve reading accuracy.
- Can deal with difficult OCR cases, impossible to achieve using traditional methods.
- Includes a deep learning based LocateText tool, able to precisely locate textual elements within an image.
- 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