Combining AI and HEAT for glassmaking

Emhart Glass Vision is the branch of Emhart which produces container glass inspection machines. Niki Estner* discusses the latest trends in AI research, and how they apply to industrial quality inspection.

AI Research moves in waves. A promising approach is discovered. Scientists around the world build on the approach, improve it, drive it to the edge of what is possible. At some point, making progress harder and slower, until someone finds a new promising approach - a new mother lode - and the cycle begins again.

The last big "wave" of AI research can be summarised as: Build large databases of training data. Label the training data manually. I.e., for every input, tell the system what the desired output is.

Train a deep neural network to reproduce the desired output directly. Use the network to replace manual labour. This is called supervised learning.

The supervised learning approach has been tremendously successful.

But the success depended on having larger databases – and that ultimately limited its potential.

Supervised learning systems also tend to learn things we don't want them to learn: For example, if you only train some type of defect, say a bird swing, on one type of glass colour, say amber glass, and never on other glass colours – then a supervised learning systems will quickly generalise that these defects only happen on amber glass – true for the training set, but not in general.

This also fuels the need for more and more accurately labeled samples.

The Present

In contrast, in the current "wave", gigantic databases of unlabeled data are used to train a deep neural network to build some form of understanding of the data. A much smaller labeled set of data is then used to train the desired output based on the pre-trained understanding of the data.

This approach was a true breakthrough. It brought us applications like Chat-GPT or AI image generators like Midjourney.

This is elegant because unlabeled data is much easier to come by. For example, production monitoring systems like the BEG HEAT system record images of millions of defects every week.

HEAT

The Bucher Emhart Glass (BEG) Hot End Advisory Tool (HEAT) shows glass plant operators a view of the current production quality, as it is seen by inspection machines: Unlike common line information systems, it shows defect classification information and images. With HEAT, operators can see what defects are currently produced, instead

Continued>>

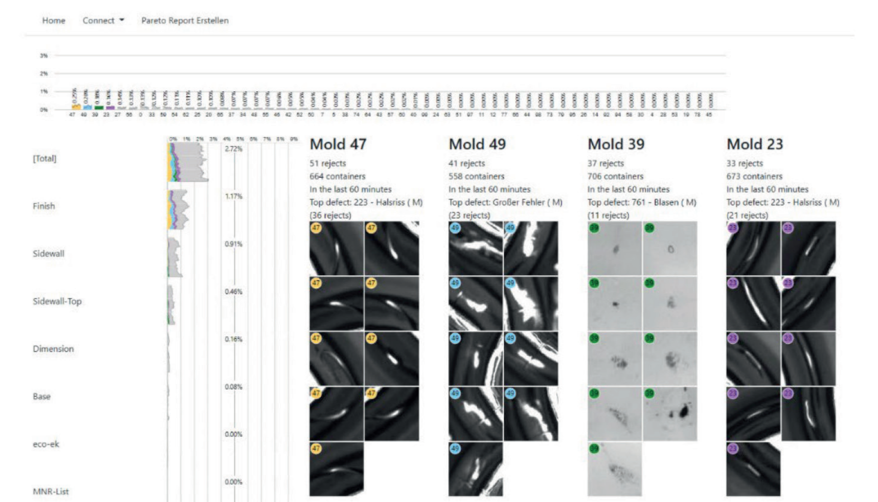

Fig 1. Screenshot of the HEAT overview screen, including defect classification.

of just how many. HEAT also includes a database of common glass defects with descriptions, causes and remedies.

Seeing classification data and images is useful for optimising live production quality. It is also invaluable for handling customer complaints: HEAT can store data for weeks or months (depending on available server storage). In case of a complaint, you can easily look at past production and defect images (***Fig 1***).

As a side effect, HEAT systems installed in over 80 production lines globally collect defect images for AI training.

Currently, Emhart Glass Vision uses a database of ~250, 000,000 images.

How

So how do you train a machine "understanding" of images? One common task is to train it to create new images that look like they're real. We all have seen pictures of the pope wearing gangsta-rap clothing.
A neural network that can produce images like this must have learned concepts like faces, eyes, body structure and many more (***Fig 2***).

So how does all this help visual quality inspection? One big problem with glass defects is that the same type of defect will look different on different glass containers – flint glass vs. amber, thin vs. thick glass change the appearance of defects.

With older machine learning approaches, that meant that we needed multiple bird swings on amber-, flint-, thin wall-, thick wall-containers, in neck area and in sidewall area – and all combinations, for training. Otherwise, a neural network would quickly learn things like "bird swings only appear on flint jars in the neck area". Which may be true for the training set, but is not true in general. And for all of these samples, a human had to manually train the system "this defect is a bird swing" for each sample – so the network can differentiate between critical and cosmetic defects.

This problem virtually goes away if you can generate new defect images. Train a lot of unlabeled samples, label a few of them with the correct defect type, and you get a system that can generate artificial images of any defect type, anywhere on any shape. These images are then used to train a much more robust defect detector.

Benefits

But were we not able to detect and classify defects years ago? True. Every container glass inspection machine produced in the last 20 years can be set up to detect all common defect types. But:

- It must be set up, and that requires a skilled operator who has time to set it up.
- Classification is often not accurate enough for process control.
- And there is always some level of false rejects. Some containers could be palletised, but are in fact rejected because inspections today cannot distinguish them from good containers.

The pre-trained Emhart AI system has a deep understanding of what acceptable and defective containers are.

Unlike in the past where you had to train inspection systems to learn through many containers about the acceptable limits in your plant, all you need to do now is to introduce a few new containers per article type to customise a system to analyse all of the containers in your plant

Classification is also much more accurate out of the box because it is no longer based on simple features like "brightness" or "width / height ratio" but instead on a deep image understanding learned from hundreds of millions of images. Better classification also leads to a lower false reject rate. In production trials, we have seen false reject reduction of up to 3% of the total production.

An average glass plant releases about 70k tons of CO_2 every year, so this effect alone could potentially save 2000 tons of CO_2 emission every year .

Summary

A typical container glass inspection machine produces hundreds of high-resolution images of every single produced container. By combining existing systems like HEAT and modern unsupervised AI methods, it is now possible to use this wealth of data directly for AI training. ■

***Manager of software development at Emhart Glass Vision, Bucher Emhart Glass, Cham, Switzerland, https://emhartglass.com/**

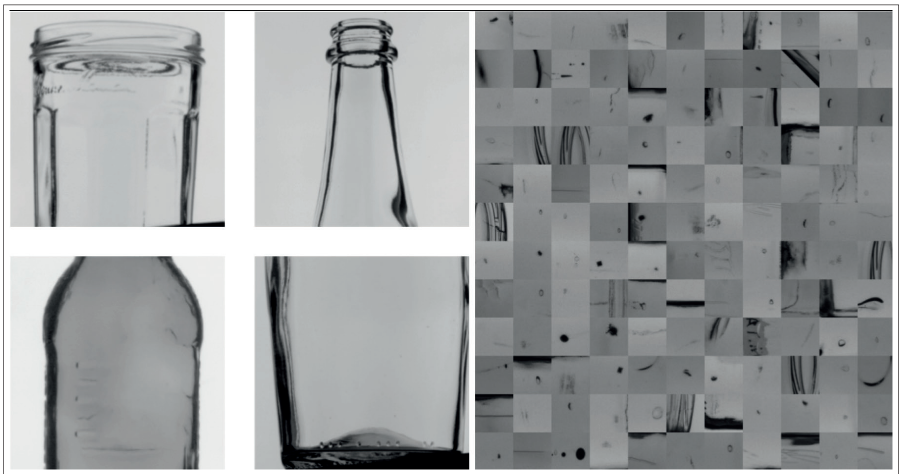

Fig 2. These are not images of real containers or real defects – they were created by a BEG AI system.