

Trainable Vision Sensors Better than Smart Cameras WHITE PAPER

The trend of the past years has been to childishly label cameras as “smart” because of a high-speed network connection or a powerful processor running complex algorithms. Still, cameras remain pretty hopeless in view of the amount of data they generate and do not respond to real needs for flexible and accurate image recognition with realistic speed performance.

The market has been ready for cameras which are smarter, smaller, faster and cheaper. This can be translated as follow:

- The sensor must have the ability to learn and recognize patterns autonomously.
- The sensor can be taught directly by the user by showing examples without any programming.
- The sensor can transmit results and/or images only when information of interest is detected.
- The sensor must be small for portability and cost reduction.

The trainable vision sensors developed by Cognimem branded as CogniSight Sensors respond to the above criteria and make image recognition ubiquitous because it is both easy and affordable.

CogniSight Sensors Merge Vision and Decision in a Miniature Apparatus

A CogniSight sensor features a video sensor and a brain which can learn and make instantaneous decisions and selective transmissions. This brain is trained and controlled by CogniSight, an image recognition technology developed by CogniMem Technologies and based on multiple neural networks.

The first time you install a CogniSight sensor, you will have to go through the few necessary steps common to all camera installations: mounting at the correct distance from the objects, lighting for good contrast, connection to a photocell or triggering device, if applicable, and finally interface to a PC for minimum configuration. This is where the roads between a CogniSight sensor and a smart camera split. The CogniSight sensor is equivalent to an image recognition server. It operates autonomously and just needs to receive the initial instruction from a client application to know what to recognize and which actions to take accordingly. This instruction can be reduced to the transmission of a knowledge file to the sensor which is equivalent to loading a knowledge into its associative memories. The knowledge is a simple data structure defined within the CogniSight technology. It is built by learning examples, can be extended and fine tuned at any time, saved to file for portability and read by any sensor powered by a CogniSight recognition engine. The client application connected to the CogniSight sensor can be designed to let a user perform immediate interactive training and/or to simply load existing knowledge files.

CogniSight Sensors Learn by Examples and Are Quick at Picking Up Patterns

Without programming any line of code, nor reading image processing manuals, you can teach a CogniSight sensor what to do just like you would teach a new operator. Simply outline your objects of interest in live or still images and label them with a category. On a factory floor, you can classify parts as Good or Bad; Acceptable, Defective or Recyclable; Type I to Type N. Inside a video surveillance station, you can classify and count passing objects as “car”, “truck”, “van”, “bicycle”, etc. As soon as you teach one example, the sensor commits a neuron and starts classifying incoming objects with this start-up knowledge. At first, it will probably classify all the objects as belonging to the one and only known category. If the features seen in an image are really different from the one learned by the sensor, the response can be “Unknown”. This will continue until you teach something new. As soon as you teach more categories, additional neurons get committed and they all start delimiting their similarity domain. The more you teach more neurons then the less confusion between different categories of objects.

Trainable Vision Sensors Better than Smart Cameras

WHITE PAPER

Immediate Recognition and Decision

Since the CogniSight sensor can recognize what is present in an image, you can let it take the appropriate action based on what is recognized in the image. On a factory floor, the sensor can directly control an ejector to reject or divert parts, transmit periodic statistical reports about the production and, if needed, record the images of parts with a certain type of defect. On a surveillance station, the sensor can send a warning signal, record the date and time a special event occurs, and transmit each image showing a particular object.

Once a CogniSight sensor is trained, it can save its knowledge for later use. This means the second time you power the sensor, it automatically and autonomously resumes inspection with no need for a PC or network connection. You can always connect a PC if you wish to monitor production statistics, start new training, or add more training to an existing engine. These interruptions are made over TCP/IP and do not interfere with the real-time recognition running between 15 and 20 frames per second.

So Simple Yet So Flexible and Adaptive

The learning mechanism of a CogniSight sensor is very intuitive since you only have to show and label examples. This “trainability” makes it possible to transfer to the sensor a human expertise built on years of practice, with no need for intermediaries such as interpreters, consultants or programmers. Still, there is no unique way to train a sensor to recognize a family of objects and you may have to teach and evaluate several engines before retaining the one which will be used for the real application. Let’s take the example of defect detection. Because it is impossible to model all the possible representations of defective parts, you can try two teaching methods: (a) be very conservative at teaching good parts and consider everything slightly different as defective; (b) be moderate at teaching defects so that more or less similar defects can be picked up. The selection of a teaching method can be influenced by the frequency of occurrence of the events or objects to recognize. If your production has 0.02% defects, it will be difficult to capture these defects and a conservative teaching might be wiser.

A Validation Driven by the Cost of a Mistake

Prior to letting the sensor take actions based on the contents of images; a thorough validation of the recognition engine is critical. This is simply done by monitoring the response of the engine over a significant number of examples. In a production environment, this verification may span over several days to ensure that the sensor has seen and classified properly a representative number and diversity of products. If corrections are required, they can be taught immediately by the end-user. Unlike conventional imaging systems, correcting the engine does not require fixing bugs in a program, but rather verifying that the taught examples were correct and detailed. The CogniSight technology has been used and field-tested for several years and its classifier is endorsed in many theses and applications. The key to the accuracy of a recognition engine is to have good teachers.

Depending on the cost of an error in an application, it might be necessary to publish strict guidelines for the teachers and even require that multiple experts annotate the same objects to make sure that there is no possible error of classification. If two human experts cannot agree to the classification of an object, you cannot expect wonders from a CogniSight sensor which acts as a simple parrot.

To help test a recognition engine, the CogniSight sensor can transmit images at high speed and selectively if requested. The CogniSight client application can then save these images and even dispatch them into folders per category. On a manufacturing line, you can collect images of products during different time and crew shifts. In an outdoor surveillance station, you can collect examples of objects under a variety of lighting and weather conditions, etc. The ability to request that the sensor transmits only certain categories of images can greatly speed up the validation of a recognition engine. For example, you can quickly see if the images stored in the folder “Category A” really belong to this category or not. If so, you can assume that the engine is well trained to identify category A. On the other hand, you will have to verify that images of category A cannot be found in the other

Trainable Vision Sensors Better than Smart Cameras

WHITE PAPER

folders. Now, if you notice that images stored in “Category A” often belong to category C, you can immediately teach more examples of objects belonging to category C. These observations help you identify the categories of objects which require more teaching due to the complexity or diversity of their features. You can also easily locate the categories which cause confusion and need to be taught with a conservative approach. These observations can also lead you to define intermediary categories or to merge categories together.

Building Image Knowledge Bases

On many occasions, it is more practical to teach and validate a recognition engine off-line, that is, without the sensor. You can take all the necessary time to teach and screen the recognition results, understand where confusion occurs, try new teaching methods and qualify recognition performance in terms of throughput and accuracy. Also, sometimes remote training might be necessary to consult an expert, or to compare the teaching instructions of multiple operators. Again, training is the key to the accuracy of an engine. The images collected with the CogniSight sensor can be used to build several knowledge files. The Image Knowledge Builder application from CogniMem is intended for this purpose. It can build a knowledge from a series of annotations made by an operator and test it on many images. Diagnostics reports guide you in the process of compromising between throughput, accuracy and speed. For example, an engine cannot discard what it has already learned. Image Knowledge Builder can pinpoint, however, when a knowledge starts to degenerate due to inconsistencies in the set of training examples. When satisfied with a recognition engine, you can simply save it to an Image Knowledge (IK) file and load it on the sensor.

Adapting a Knowledge to Changes and Variations

A CogniSight sensor can learn many models of objects and solve recognition problems with high non-linearity. For example, if it is easy to model a good screw from a bad screw, it is more tedious to describe a good fish from a bad fish, especially when the fish is not one among a dozen swimming in an aquarium, but rather one herring caught with a million of other herrings. A common challenge in industrial inspection is adaptability to changes in production, whether it is the introduction of a new part, the change of one component in a part, or other. With a CogniSight sensor, this problem can be solved in a very practical way: add more training to an existing recognition engine, or start a new recognition engine. For example, as awkward as it may sound, a CogniSight sensor trained on a fishing boat in December to sort herring might perform poorly in July. The explanation is very pragmatic: fish lose their scales in different seasons and this directly alters their reflectivity. As a result, the images seen by the camera in July are different from the ones seen and taught in December. The customer has the option to build a “universal” engine recognizing the scales of herring throughout the year, or to build several smaller engines for each season. Both options request a new image collection, some training and validation, and all these operations can be performed by the end-user.

Depending on the features of your objects, you might very well be able to teach a “universal” engine capable of classifying your entire production. This is especially true if you simply need to sort your parts into Pass and Fail. You can truly benefit from the non-linear classifier of the CogniSight sensor which will have no problem storing the same knowledge examples of good low-fat strawberry drinkable yogurt, pineapple flavored cottage cheese and vanilla yogurt. Your only risk trying to merge all these models in a same knowledge is to end up with a large knowledge base (which is not necessarily a problem) and, possibly, less accuracy. This can be tested quickly with tools like the Image Knowledge Builder so you can make decisions based on comparative results.

Another reason to start a new recognition engine can be un-related to the appearance of the objects to inspect but rather related to a level of moderation of the engine. For example, during hot summers, the quality of a label on a beer bottle might be less critical than during winter time. As a consequence a brewery might use two different engines, one being more tolerant to certain types of defects than the other, therefore increasing the throughput of the production.

Trainable Vision Sensors Better than Smart Cameras WHITE PAPER

Smaller Size and Price Tag

The evolution of the CogniSight sensors is going toward smaller, cheaper and faster appliances.

Smaller and cheaper make it possible to multiply inspection stations along a production line. Instead of building one expensive system checking all the acceptance criteria of a part, multiple CogniSight sensors can be distributed along a factory line and perform simple inspection tasks with standard equipment (lens, lighting, ejectors, etc). For example, in the bottling industry it makes sense to eject a bottle as soon as it is known that its filling level is below acceptable, or it is not properly sealed, or its label is improperly affixed. Taken individually, each of these controls are quite simple and prevent incoherence such as wasting supplies of caps and labels on empty bottles. Another advantage of early and local problem detection is the avoidance of damage to the machinery.