CogniSight
Miniature Trainable
Vision Sensor

Control Panel

Version 2.0

A product of General Vision Inc.

GENERAL
VISION

Table of Contents

License Agreement

This License Agreement is a legal agreement between you (LICENSEE) and General Vision Inc. (General Vision) for the enclosed computer software, and any associated media, printed materials, and "online" or electronic documentation ("SOFTWARE PRODUCT").

License, rights and limitations
General Vision grants to the LICENSEE a revocable, non-transferable, and non-exclusive license to use the SOFTWARE PRODUCT. This license may not be shared or used concurrently on different computers. You may install one copy of the SOFTWARE PRODUCT, or any prior version, on a single computer for use by a single user. LICENSEE may make one copy of SOFTWARE PRODUCT for backup purposes. No other copies shall be made without General Vision's prior written consent.

Ownership and Copyright
The SOFTWARE PRODUCT, including all supporting documentation, is a proprietary product of General Vision. General Vision retains all title and copyrights in and to the SOFTWARE PRODUCT, the accompanying printed materials or electronic documentation and any complete or partial copies of the SOFTWARE PRODUCT.

Limited Warranty
General Vision warrants that the SOFTWARE PRODUCT will perform substantially in accordance with the accompanying written materials or electronic documentation and General Vision will make reasonable efforts to solve any problem issues arising out of the intended use of the SOFTWARE PRODUCT. To the extent allowed by applicable law, implied warranties on the SOFTWARE PRODUCT, if any, are limited to ninety (90) days. General Vision disclaim all other warranties and conditions, either express or implied, including, but not limited to, implied warranties of merchantability, fitness for a particular purpose, with regards to the SOFTWARE PRODUCT.

Limitation of Liability
To the maximum extent permitted by law, in no event shall General Vision be liable for any special, incidental, indirect, or consequential damages whatsoever (including, without limitation, damages for loss of business profits, business interruption, loss of business information, or any other pecuniary loss) arising out of the use of or inability to use the SOFTWARE PRODUCT or the provision of or failure to provide support services, even if General Vision has been advised of the possibility of such damages. In any case, General Vision's entire liability under any provision of the license agreement shall be limited to the amount actually paid by you for the SOFTWARE PRODUCT.

Terms and termination
This license agreement is effective from the date of receipt and shall remain in full force until terminated. Without prejudice to any other right, General Vision may terminate this LICENSE AGREEMENT if you fail to comply with any of the terms and conditions of this license agreement. In such event, you must destroy all complete and partial copies of the SOFTWARE PRODUCT in your possession.

Trademarks
CogniSight is a trademark of General Vision Inc.
Product names listed are trademarks of their respective manufacturers.

Support
support@general-vision.com
Tel 707-765-6150

Introduction

The CogniSight Control Panel is a very simple interface to configure a CogniSight MTVS to recognize images whether for part inspection, video surveillance or else.

CogniSight MTVS is a miniature trainable vision sensor which can learn and recognize in real-time a region of interest within a video frame. It features a "fully hardwired" image recognition engine and neural network, both sinking less than 15 milliamps...Recognition is performed at 60 frames per second without a single line of software.

At power up, CogniSight MTVS starts digitizing video frames and recognizing their contents provided that it is loaded with a knowledge. This knowledge is actually the contents of the neurons (describing the patterns to recognize) and the definition of a region of interest (describing where to look for in the video images). In addition to handling the recognition, the neurons are also capable of adaptive learning. This means that they will build the knowledge themselves and the only task left to the user is to find good examples and teach them. The enclosed Control Panel provides a very easy to use interface to teach different categories of objects, save the knowledge to the Flash memory of CogniSight MTVS or to a file for backup purposes.

Installation

- Connect the CogniSight MTVS to the USB adaptor
- Connect the USB adaptor to a PC USB2 port.
- Windows will detect the new device and prompt for the location of its driver. The later is supplied on the CD delivered with the CogniSight Control Panel.
- Install the CogniSight Control Panel software
- Launch the application from the Start\Program Files menu\ General Vision menu.
- The live video image will appear instantaneously on screen.

Hardware specifications

- Recognition at up to 60 frames per second
- Neural network of 18 neurons
- Single region of inspection (ROI) from 16x16 pixels up to full frame
- Category output to I/O lines at frame rate
- Resume inspection at power up
- Flash memory to save application settings
- I2C communication
- LVDS video transmission
- Delivered with USB adaptor

Control Panel software

- Sensor settings (trigger, gain and shutter speed)
- Select the region of interest graphically
- Teach by examples within the region
- Up to 4 categories of objects
- Request the transmission of all or specific categories
- Record images at up to 15 fps
- Load and save project files (i.e. image knowledge) to/from the board

Getting started

At start-up, CogniSight Control Panel shows the live video in a window. A region of interest (ROI) is outlined in the image. It represents the area to learn and/or recognize and its coordinates are reported in the status bar at the bottom of the image. The status bar also reports the ratio between the number of neurons presently loaded with a knowledge and the total number of neurons available. Both the region of interest and the contents of the neurons are restored from the Flash memory of CogniSight MTVS when the latter is powered up. If the Flash memory has never been programmed the region of interest is centered in the video frame and occupies half its surface, and no neuron has a knowledge.
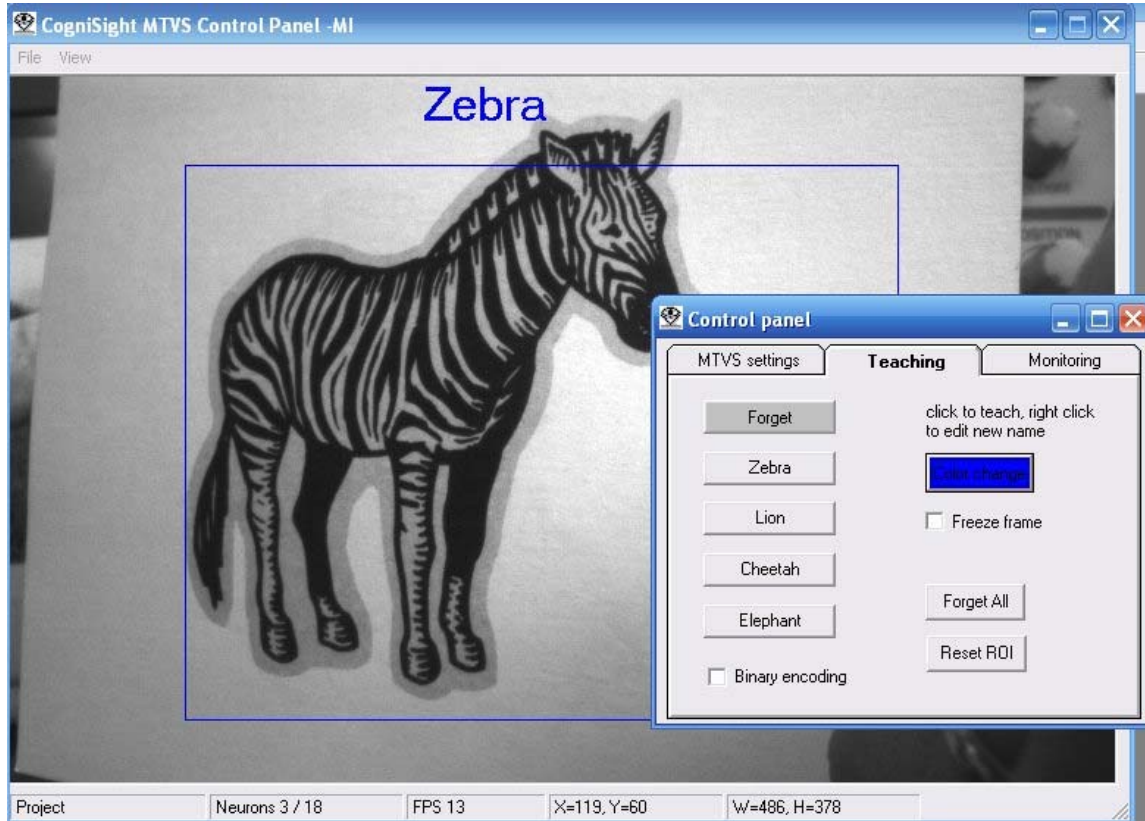
The Control Panel, accessible from the View menu, allows to edit the hardware settings of the device, teach what to recognize, monitor if the video images are classified properly and enable selective transmission . When the accuracy of a recognition engine is validated, the knowledge can be saved to a project file and also into the Flash memory of CogniSight MTVS. This last operation will ensure that the device can resume recognition autonomously at power-up without any need for a computer!

Sensor Settings

The first tab of the Control Panel gives access to the settings of the vision sensor. Changes can be observed immediately in the image on screen.
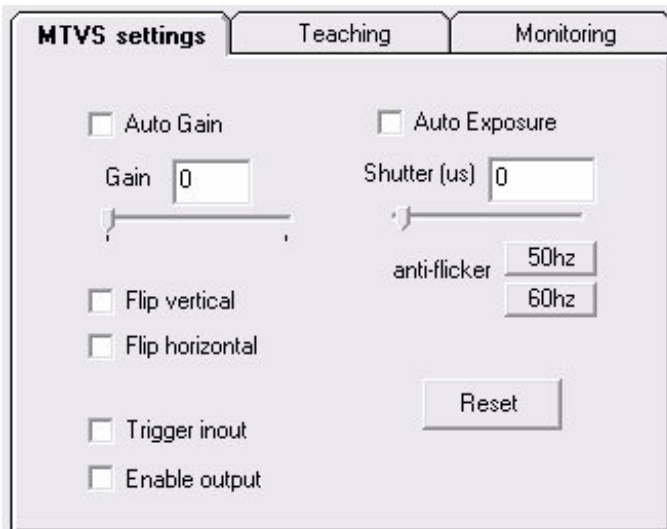
Auto-Gain and **Auto Exposure** are set by default. These options can be disabled in order to adjust the Gain and Shutter values manually.

The 50Hz and 60Hz **anti-flicker** buttons set the shutter speed to a value limiting the flickering effects of a fluorescent light powered with a 50 or 60 hertz power supply.

Flip vertical and **Flip horizontal** change the readout order of the lines and columns of the CMOS sensor. These options are convenient to retrieve images showing objects always in the proper orientation independently from the orientation of the CogniSight MTVS.

The capture of video images can be triggered by an external signal connected to the Trigger input line of CogniSight MTVS. Selecting the **Trigger** option. automatically turns off the use of the Auto-Gain and Auto-Exposure so the Gain and Shutter speed can be adjusted manually.

The **Enable Output** checkbox enables the output of the recognized category onto the output lines of CogniSight MTVS. Refer to the hardware manual of CogniSight MTVS for information on how to connect these output lines to external devices such as ejectors, PLCs, recorders, etc.

Project Settings

Starting a new application requires the definition of a region of interest and the description of the categories to discriminate.

Region of Interest

The size and position of the ROI is easily edited on screen using the mouse cursor (left mouse click moves the region, right mouse click resizes the region).


The region of interest should be defined as the smallest rectangle including the objects to recognize. Note that it does not need to necessarily include the entire object, but can be limited to the portion showing its discriminating features. The only constraint for a given project is to keep the same aspect ratio for the ROI at all time.

In the example to the left, the region of interest contains the entire bottle cap. However, if the purpose of the inspection is to detect that a cap is present and the filling level is correct, a thinner and longer vertical rectangle would be appropriate.

Categories of objects

Four categories of objects can be taught from the Teaching tab of the Control Panel. Their default names are Category 1, Category 2, Category 3 and Category 4 and appear on the four teaching buttons. The names can be changed at any time by double-clicking at the corresponding button. An Edit box is prompted and a new name can be edited.
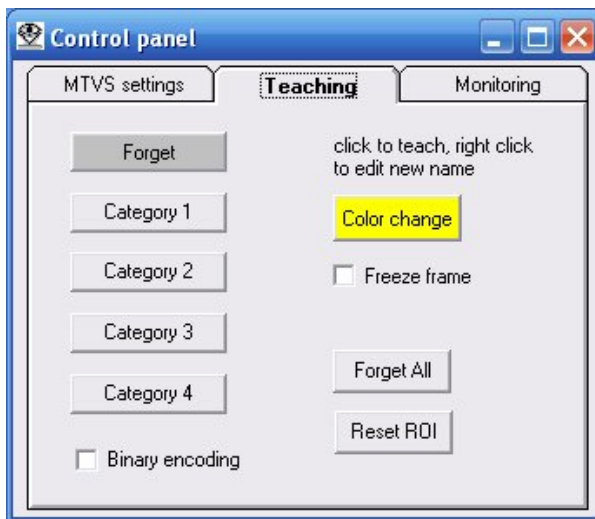
Examples of categories for a part inspection application:
- Good (category 1)

- Acceptable (category 1)
- Recyclable (category 2)

- Grade A (category 1)
- Grade B (category 2)
- Grade C (category 3)
- Grade D (category 4)

Examples of categories for a traffic monitoring application:
- Vehicle (category 1),
- Pedestrian (category 2)

- Car (category 1)
- Truck (category 2)
- Van (category 3)

Category Binary Coding option

The category names are associated to the category values 1, 2, 3 and 4 which are the ones transmitted on the output data bus of MTVS. If an application requires connection of this data bus to some electronic components, it is possible to encode the category values in binary format 2, 4, 8 and 16 so that each category activates a separate line of the data bus as shown in the table below.

	Default	Binary coding	Binary output value
Unknown	0	1	00000001
Cat 1	1	2	00000010
Cat 2	2	4	00000100
Cat 3	3	8	00001000
Cat 4	4	16	00010000

Teaching

To teach an example, move the ROI over a relevant location and click at a corresponding category button. Examples can be picked up at different locations in different images.
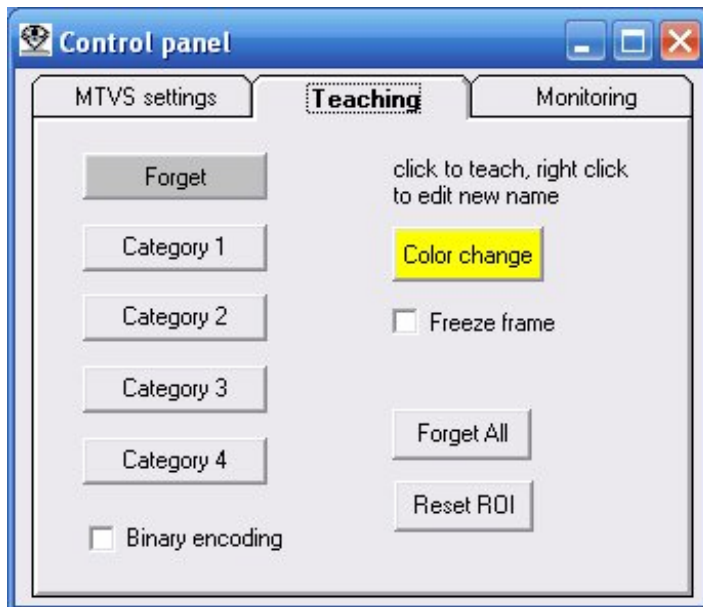
As soon as one example is taught, CogniSight MTVS starts recognizing the region of interest in real-time. The recognized category is reported in the image window.

At first, the engine will classify all images as belonging to the one and only known category. If the image content looks really different from the learned example, the response can be "Unknown". This will continue until a second and more categories are taught. The more teaching, the better the discrimination between different categories of objects.

Click **Forget** when the recognition engine makes mistakes recognizing a region which does not belong to any of the defined categories. This action does not expand the existing knowledge, but helps moderate its tendency to over-generalize.

Click **Forget All** to clear the entire knowledge and start a learning session from scratch.

Remember that teaching an incorrect example can corrupt an entire knowledge, so back up can be important. The knowledge can be saved at any time during a learning session using the **File/Save project** menu. It is a good idea to do so regularly naming the project file with an incremental index as the number of neurons increases. Another alternative to avoid the mis-happening of involuntary knowledge corruption consists of teaching with the Image Knowledge Builder application.

Questions and Answers

Q&A: I correct the erroneous recognition of Category A as Category B, the engine reports Category B but no neuron is added to the knowledge

This behavior is typical when the region of interest is recognized with uncertainty. Several neurons fire but do not agree on the recognized category. The category A seen on screen is the first category of the neurons with the highest confidence level. However, a neuron with category B also recognizes the region. By teaching category B again, the firing neurons with a category A will reduce their confidence level for the profit of the neuron with the category B. No neuron is therefore created but the existing neurons adapt their influence field to respect the teacher's instruction.

Q&A: I correct the erroneous recognition of Category A as Category B, but the engine does not take my request into account

This behavior is typical when the region of interest is recognized with uncertainty. Several neurons fire with a same confidence level but different category values. The category A seen on screen is the first one of this list, but a neuron with category B also recognizes the region. Therefore, teaching category B again does not create any new neuron. A method to confirm this behavior is to check the option "Conservative Recognition" on the Monitoring tab. In this case, the display should report "Uncertain".

Q&A: I want to reinforce an uncertain recognition of Category A, but the engine does not take my request into account
This correction requires that the firing neurons reporting a category different from A correct their response (and reduce the range of their similarity domain). Keep trying and Category A should become the only recognized category after a few learning. The "Uncertain" lable will dissapear.

Q&A: It seems that I have to learn in every frame to get a correct result. Should training be so tedious?

One common mistake which can require a tedious training is to select a region of interest which is too big and include irrelevant information, especially from the background, but also from a portion of an object which does not include any discriminant feature. For example, to detect the presence of a person, one can decide to learn the profile of an entire face with or without shoulders, or a pair of eyes with or without the nose. The requested working conditions of an application will determine which approach seems the most suitable. The teaching utilities of Image Knowledge Builder are also very easy and quick to test multiple approaches. One way to understand why a knowledge does not produce accurate results is to review the contents of the neurons (under View/Models menu). If the contents of a neuron is not relevant to what should be recognized, a new learning session should be started.

Q&A: I want to train on a large number of images to validate thoroughly the accuracy of the recognition

Using the Image Knowledge Builder application to teach a recognition engine is a more serious method than teaching on live images through the MTVS Control Panel. With IKB, multiple trial and validation scenarios can be executed on a variety of images. When the accuracy of the recognition is satisfying, an image knowledge file (*.ikf) can be saved for export to CogniSight MTVS. In To proceed with this recommended methodology, images need to be collected through the Control Panel. If applicable, it can be very useful to arrange the recorded images into folders describing their categories or the environmental conditions under which they were acquired (i.e. good lighting, poor lighting, uniform background, moving background, close up, etc).

Image Knowledge Builder (IKB) features a very simple user interface to build a knowledge by learning automatically series of images pre-sorted into folders per category. The accuracy of this knowledge can be tested by comparing the classification returned by the engine with the location of the file in a category folder. These batch processes are very quick and practical to help identify difficult or complex classifications, pinpoint possible erroneous or inconsistent teaching, evaluate compromises between throughput and accuracy, etc.

View models used by the neural network

The View Model menu allows to review the contents of the neurons which represent the knowledge. This utility can be helpful to understand why some erroneous or uncertain classifications occur. It can also point to examples which are irrelevant because showing too much background.

The content of a neuron comprises a reference pattern, a category and a similarity range. The reference pattern is displayed in a format proportional to the region of interest.
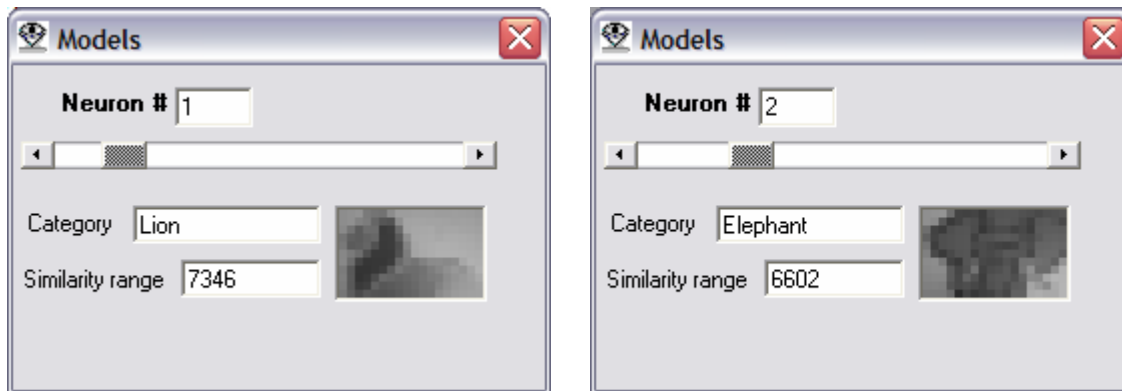

In the illustrations above, the first neuron describes the Lion. If the signature of a region of interest (ROI) looks similar to this "iconic" representation of the lion and the distance between the two patterns is less than the similarity range of 7346, Neuron #1 fires and reports the category of the ROI as being Lion. Neuron #2 will perform the same comparison and report the category "Elephant" if the difference between its reference patterns and the ROI is less than 6602. In the event that both Neuron #1 and #2 fires at the same time, the classification is uncertain. Looking at the reference patterns stored in the neurons can reveal the lack of examples showing discriminant features, or the abundance of examples which are too similar.

Image Monitoring

The recognition engine of CogniSight MTVS continuously reports the recognized category in the image window. This category is the one of the first neuron with a "best match" distance. More explicitly, it is the smallest category value of all the neurons firing with a same smallest distance value.

If additional categories are identified by other neurons (more or less confident than the first neuron), the label "Uncertain" is displayed.
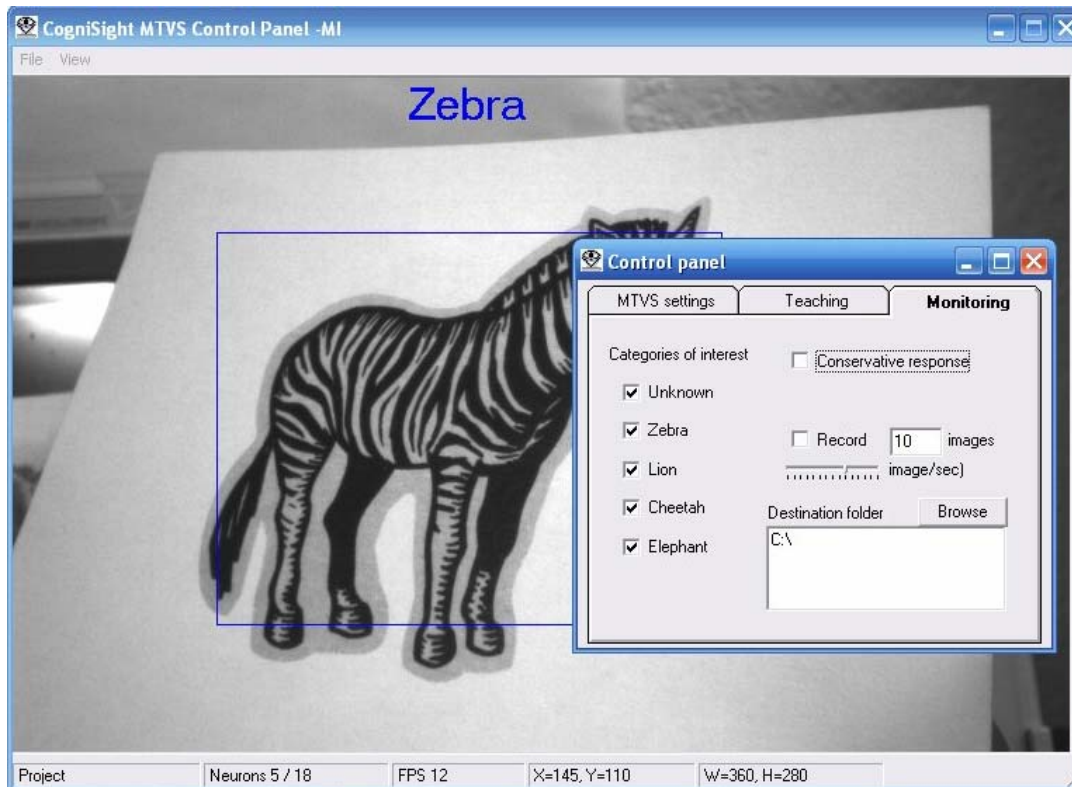
A first method to verify the accuracy of the recognition is through selective image transmission. A second and more robust method consists of using Image Knowledge Builder application (see paragraph below)

The Monitoring tab allows to select a specific category (or categories) of images to display on screen. If a single category (Grade A for example) is checked, only this type of images should appear on screen. The appearance of an image which clearly does not belong to the requested category indicates that the engine is still making confusion and requires more training. If it is difficult to isolate images classified incorrectly on the screen, they can be saved to disk by checking the "Record" option.

Image Collection

Image collection can serve two important purposes:

(1) Build a data base of images for training and validation with Image Knowledge Builder or other applications running a CogniSight recognition engine
(2) Verify the proper recognition of the video images by recording them selectively based on their content.

Prior to collecting images, a destination folder must be selected. Clicking the Browse button prompts a Save dialog box to select, create and delete folders.

Collection for off-line training

A CogniSight recognition engine has very good generalization capabilities and is also very adaptive. It is best taught by showing as many examples as possible, representative of the variety of products and environmental conditions of the final application. Among all the examples, the CogniSight model generator will automatically learn the ones which are significant and discard the ones which do not add information to the existing knowledge.

When starting a new application, if applicable, it is a good idea to collect images per type of objects to recognize and store them into corresponding folders. Examples can also be organized per type of environmental conditions they represent such as lighting conditions, cases of known and/or unknown background, ranges of scale, ranges of angular positions, etc.

Images collected for an application should be arranged into of folders and subfolders representative of the possible case studies. Such organization is not mandatory, but can help expedite the training efforts and the identification of the easy and difficult cases it encompasses. Indeed, using the Image Knowledge Builder application, training can be done in a few mouse clicks using the "Learn by Folder" utility.

Following is the description of how to proceed with a new application using the example of the animal cards supplied with the CogniSight MTVS.

Step 1:   Click at the Destination folder and create a "Cheetah" folder. Close the dialog box.
Step 2:   Edit the number of images to collect (50 for example)
Step 3:   Place the card of the Cheetah in front of the CogniSight sensor
Step 4:   Click the Collect checkbox and wait until it gets unchecked meaning that the collection is done. Note that a variety of "slightly different" examples can be saved by moving the card closer and farther from the sensor, or by rotating it (all this within reasonable amplitudes).
Step 5:   If a recognition engine must cope with varying lighting conditions, translations, rotations, and more, add subfolders under the Cheetah folder and give them names describing the conditions they represent.

Repeat: Proceed in a similar way with all the animal cards

Step 6:   Using IKB, apply the "Learn by Folder" to the images collected in Step 4 for each category of animals.
Step 7:   Once the recognition of these images is correct, save the project to an *.ikf file
Step 8:   Open the Control Panel and load the project. Place a card in front of the sensor and verify that it is recognized.

Repeat: Proceed in a similar way with all the animal cards

Collection for verification of live recognition

Step 9:   Click at the Destination folder and create a "Recognized_Cheetah" folder. Close the dialog box.

Step 10:   Uncheck all the category checkbox of the Monitoring panel except for the Cheetah.

Step 11:   Edit the number of images to collect (100 for example)

Step 12:   Check the Collect box and pass the different cards in front of the sensor. Continue until the Collect checkbox is unmarked.

Step 13:   Using IKB or Windows Explorer, open the "Recognized_Cheetah" folder and verify that the images only show images of the Cheetah. If, not, more training has to be done!

Project File Management

Once a recognition engine delivers the required accuracy, the knowledge can be saved to a project file (*.ikf) or in the flash memory of the CogniSight MTVS.

Image knowledge files (ikf)

This project file contains all the information necessary to resume the current recognition. This information includes:
- the sensor settings,
- the definition of the region of interest
- the names of the categories and
- the contents of the neurons

The file is fully compatible with the Image Knowledge Builder and vice-versa. This means that both applications can exchange *.ikf files.

Remark about project saved from IKB

A project file (*.ikf) originally saved in Image Knowledge Builder contains camera settings which are all set to "0". When loading such file the camera settings have to be adjusted manually under the Control Panel and produce images with a contrast and brightness similar to the ones used to build the knowledge under Image Knowledge Builder.

A project file (*.ikf) originally saved in the MTVS Control Panel contains the camera settings. Loading this file in Image Knowledge builder and adding more knowledge to it does not erase the camera settings.

Save to Flash

The Save to Flash command saves to the Flash memory of the CogniSight MTVS all the information necessary to resume the recognition at the power up of the device. This information includes:
- the definition of the region of interest
- the contents of the neurons
- Status of the recognition status register indicating if the live recognition is enable or not, and if the outputs are enabled or not.

This information excludes the following data:
- the sensor settings,
- the names of the categories

This means that if recognition does not work properly, the quality of the images must be verified and look similar to the one seen at the time of the training. Also the teaching buttons will still be labeled with their default names.

Index