

Image Knowledge Builder

Interface with the V1KU board

Version 5.1
Revised 07/22/2011



This is a legal agreement between you (either as an individual or as an authorized representative of your employer) and CogniMem Technologies Inc ("CTI") concerning your rights to use the Software described in this manual and any accompanying written materials (the "Software"). In consideration for CTI allowing you to access the Software, you are agreeing to be bound by the terms of this Agreement. If you do not agree to all of the terms of this Agreement, do not download the Software. If you change your mind later, stop using the Software and delete all copies of the Software in your possession or control. Any copies of the Software that you have already distributed, where permitted, and do not destroy will continue to be governed by this Agreement. Your prior use will also continue to be governed by this Agreement.

LICENSE GRANT. Exclusively in conjunction with Licensee's development and sale of a product containing a CogniMem unit (e.g., an ASIC chip, FPGA chip, or IP core) supplied directly or indirectly from CTI and listed in the CTI product catalog, CTI grants to you, free of charge, the non-exclusive, non-transferable right to use the Software.

COPYRIGHT. The Software is licensed to you, not sold. CTI owns the Software, and United States copyright laws and international treaty provisions protect the Software.

SUPPORT. CTI is NOT obligated to provide any support, upgrades or new releases of the Software. If you wish, you may contact CTI and report problems and provide suggestions regarding the Software. CTI has no obligation whatsoever to respond in any way to such a problem report or suggestion. CTI may make changes to the Software at any time, without any obligation to notify or provide updated versions of the Software to you.

NO WARRANTY. TO THE MAXIMUM EXTENT PERMITTED BY LAW, CTI EXPRESSLY DISCLAIMS ANY WARRANTY FOR THE SOFTWARE. THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. YOU ASSUME THE ENTIRE RISK ARISING OUT OF THE USE OR PERFORMANCE OF THE SOFTWARE, OR ANY SYSTEMS YOU DESIGN USING THE SOFTWARE (IF ANY). NOTHING IN THIS AGREEMENT MAY BE CONSTRUED AS A WARRANTY OR REPRESENTATION BY CTI THAT THE SOFTWARE OR ANY DERIVATIVE WORK DEVELOPED WITH OR INCORPORATING THE SOFTWARE WILL BE FREE FROM INFRINGEMENT OF THE INTELLECTUAL PROPERTY RIGHTS OF THIRD PARTIES.

LIMITATION OF LIABILITY. IN NO EVENT WILL CTI BE LIABLE, WHETHER IN CONTRACT, TORT, OR OTHERWISE, FOR ANY INCIDENTAL, SPECIAL, INDIRECT, CONSEQUENTIAL OR PUNITIVE DAMAGES, INCLUDING, BUT NOT LIMITED TO, DAMAGES FOR ANY LOSS OF USE, LOSS OF TIME, INCONVENIENCE, COMMERCIAL LOSS, OR LOST PROFITS, SAVINGS, OR REVENUES TO THE FULL EXTENT SUCH MAY BE DISCLAIMED BY LAW.

HIGH RISK ACTIVITIES. You acknowledge that the Software is not fault tolerant and is not designed, manufactured or intended by CTI for incorporation into products intended for use or resale in equipment in hazardous, dangerous to life or potentially life-threatening environments, such as in the operation of nuclear facilities, aircraft navigation or communication systems or direct life support machines, in which the failure of products could lead directly to death, personal injury or severe physical or environmental damage ("High Risk Activities"). The inclusion of the Software as critical component in High-Risk Activities implies that the manufacturer assumes all risk of such use and in doing so indemnifies CTI against all charges.

This Agreement constitutes the entire agreement between you and CTI regarding the subject matter of this Agreement, and supersedes all prior communications, negotiations, understandings, agreements or representations, either written or oral, if any. This Agreement may only be amended in written form, executed by you and CTI.

Table of content

TABLE OF CONTENT	3
INTRODUCTION	4
WHAT CAN I DO WITH IMAGE KNOWLEDGE BUILDER?	4
INSTALL THE IKB_VIKU SOFTWARE.....	4
GETTING STARTED	4
CONFIGURE THE CAMERA	5
ADJUST FOCUS IF NECESSARY	5
ADJUST CONTRAST AND BRIGHTNESS	5
SETTING THE I/O CONTROLS	6
CLEANING THE LENS	6
DEFINE THE REGION OF INTEREST	7
SIZING THE REGION OF INTEREST	7
CHANGE POSITION.....	7
TEACH EXAMPLES	8
NAME YOUR CATEGORIES	8
TEACH EXAMPLES	9
ADVANCED OPTIONS	10
<i>Region Settings</i>	10
<i>Learn adjacent neighbors</i>	11
<i>Learn surrounding background</i>	11
<i>Similarity tolerance</i>	12
SURVEY THE RECOGNITION OF A REGION OF SEARCH	12
VIEW THE KNOWLEDGE CONTENT	13
IMAGE COLLECTION	13
INSTRUCTIONS	14
USAGE	15
<i>Collecting images for the training in the absence of external trigger</i>	15
<i>Collecting images for validation</i>	16
PROJECT MANAGEMENT	17
PROJECT FILE	17
SETTINGS FILE.....	17
DISTRIBUTING PROJECTS.....	17
<i>Missing settings file</i>	17
<i>Inter-platform distribution</i>	17

Introduction

IKB_V1KU software is intended for use with the V1KU module when configured with the default CogniSight engine for object recognition. You can adjust the video acquisition settings, define a region of interest and teach the neurons to recognize and classify this region. Once satisfied with the accuracy of the recognition, you can activate the use of the I/O lines to interface directly with actuators and other applications. The knowledge built by the neurons during the teaching phase can be saved to file for later use and even distribution to other V1KU boards.

What can I do with Image Knowledge Builder?

- View the live video images to adjust the positioning of the camera
- Adjust the sensor settings to obtain the best image quality
- Define the region of interest to monitor
- Teach up to 1024 categories of objects
- Monitor a region of interest in live images
- Record images to disk based on the content of the image
- Activate the output lines of the camera
- Save the knowledge you have built by teaching examples

Install the IKB_V1KU software

Start by installing the V1KU driver on your system (refer to the installation instruction of the USB driver).

Double click at the setup.exe file of the Install_IKB_V1KU folder.

Upon successful installation, click at the Windows Start menu and select Start\All_Programs\CogniMem\IKB_V1KU. Verify that you have a live video image on screen.

You are now ready to use the V1KU miniature vision sensor!

Getting started

A new project can be defined in 6 easy steps:

- Configure the camera (gain and shutter)
- Define the size of the objects to recognize (a part, a component within a part, a window, or else)
- Define the categories of objects to recognize (i.e. good, bad; or grade1, grade2, grade3; or A,B,C,D,E; etc)

- Teach examples for each category of objects
- View the knowledge generated by the neurons
- Monitor results on screen
- Collect images and save them to disk for further training and validation
- Save images based on the result of the recognition (i.e. selective recording)

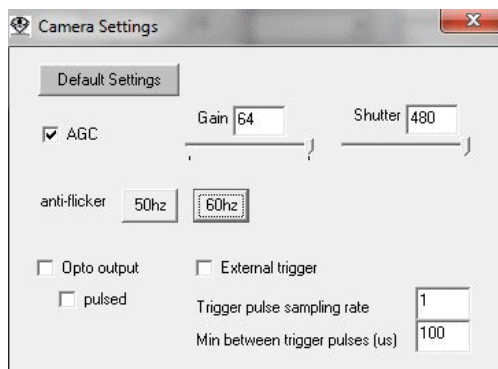
Configure the camera

Adjust focus if necessary

Insert the tip of a toothpick or tweezers in one of the groves surrounding the lens and rotate it as needed, clockwise to increase the focal length and counterclockwise otherwise.

Adjust contrast and brightness

This menu allows to adjust the video acquisition and output settings.



The **Auto-Gain** and **Auto-Exposure** options adjust the gain and shutter automatically in order to maintain the best intensity dynamic and contrast when the ambient light varies. Auto Gain and Auto Exposure are turned on by default.

Gain is the sensor analog gain. It can be changed when AutoGain is turned off and its value ranges between 16 and 64.

Shutter is the total integration time of the sensor in number of rows. It can be changed when AutoExposure is turned off and its value ranges between 0 and 480. This value multiplied by 30 gives the exposure time of the sensor in micro seconds.

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.

The **external trigger** is supported by the V1KU firmware revision 5.1 or later. When selected, two parameters can be adjusted:

- **Minimum lapse time** between valid trigger pulses. This delay is expressed in micro-seconds. It allows filtering invalid trigger pulses due to a bouncing input signal.
- **Sampling rate** or number of pulses to count prior to acquiring a new video frame. The default is one which means that every valid pulse triggers an image acquisition. If the value is 3 for example, an image is acquired only every three valid trigger pulses.

Setting the I/O controls

External Trigger sets the sensor to trigger mode. A new frame is captured only upon receipt of an external trigger pulse. This pulse must be received by the pins 6 and 10 of the J6 connector, corresponding to the anode and cathode of the input opto_coupler of the camera.

Opto outputs enables the transmission of the last 2 bits of the recognized category to the two opto-coupled output lines of the camera. Bit 0 can be read between pin 3 and 4 and Bit 1 can be read between pin 5 and 6.

Pulsed sets the output of the opto-coupled lines to a pulse format instead of the default sustained output. In this case the pulse lasts from the end of the recognition till the beginning of the next frame.

Cleaning the lens

If you observe fixed blurry spots in the field of view, you may have to clean the lens.

- To clean the lens, you will need a small screw driver, a Qtips and some rubbing alcohol
- Remove the 2 small screws located under the lens holder and extract the V1KU from the enclosure.
- Remove the two screws behind the holder
- Gently pull the lens holder straight (be careful the sensor is somewhat "clipped" inside)
- Rub the sensor with the wet QTips
- Clip back the lens holder over the sensor and connect to the Diagnostics or IKB software to view the image and verify that no dust spot appears in the field of view.
- When satisfied, screw the lens holder back


Define the region of Interest

The Region Of Interest (ROI) is the smallest rectangular area including your object. It should be limited to the smallest discriminant portion of your object and exclude any unnecessary background area.

For example, if a mechanical part can be classified as good or defective based on the presence or absence of a screw at a given position in the part, the region of interest should be not be defined as the whole part but rather limited to the rectangle where the screw is expected.

Sizing the region of interest

To size the region of interest, click at its lower right corner and drag the mouse cursor while keeping the right button of the mouse depressed.



Remark: Once you start teaching examples, the ROI can be resized but it has to keep its aspect ratio to remain compatible with the knowledge under construction. As a consequence you may notice that the ROI snaps to an invisible grid.

Change position


To change the position of the region of interest (ROI), click at its upper left corner and drag the mouse cursor while keeping the left button of the mouse depressed.

The ROI can also be moved one pixel at a time using the Arrow keys.

Teach examples

Name your categories

The Teach button opens the following panel. The first step consists of defining the categories of objects you want to discriminate.



The names of the Category buttons can be changed by clicking at the corresponding button with the right mouse button. An Edit box is prompted and a new name can be edited.

Examples of categories for a part inspection

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

Or

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

Examples of categories for a traffic monitoring application:

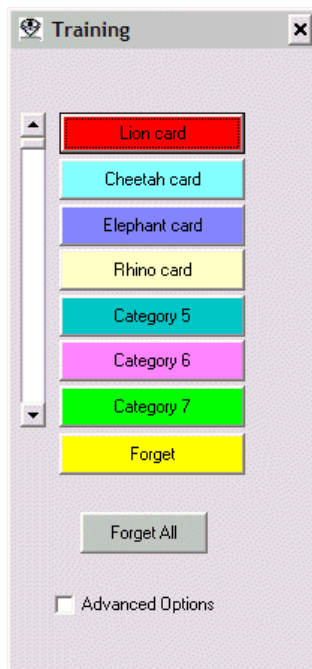
- Car (category 1)
- Truck (category 2)

- Van (category 3)


The color associated to a category can also be changed as follow: Click at the category button with the right mouse button while pressing down the Shift key. A color dialog appears so you can assign a new color to the category button.

Teach examples

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, the camera starts recognizing the region of interest in real-time. The recognized category is reported in the image window along with a distance value (Dist). This value quantifies the difference between the ROI and the original example taught to the neurons. A distance of 0 means that an exact match is found.



At first, the ROI might be recognized as the one and only known category in all the images. This overgeneralization made by the neurons will be corrected when an example of a different category is taught or when a background example is taught.

A background example is an example of ROI which should not belong to any category. It is taught by clicking the **Forget** button and has the effect to reduce the influence field of existing neurons so they do not over-generalize what they know.

The more examples are taught, the better the discrimination between different categories of objects.

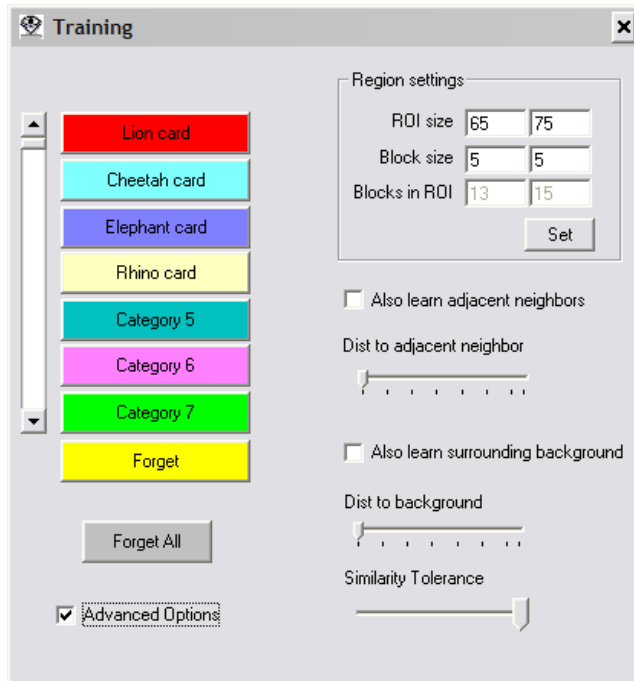
Click **Forget All** to clear the entire knowledge and start a learning session from scratch.

Remark:

Teaching an incorrect example can corrupt an entire knowledge, so remember to save the knowledge built by the neurons from time to time by calling the **File/Save project** menu. It is also a good idea to rename the project file with a different suffix documenting the number of neurons in use at the time of the Save.

Advanced options

The advanced options must be used with caution. Please read the descriptions below carefully and refer to the CogniMem technology reference guide and CM1K manual for details information.




Region Settings

Changing these settings should be done carefully.

The signature extracted from the ROI is called a sub-sample and is function of the width and height of the ROI, but also two additional attributes called Block Width and Block Height. This block size describes the internal geometry of the ROI and divides it into up to 256 sub-regions.


The pixels of the i^{th} block are averaged to produce the i^{th} component of the signature vector. The signature vector is the pattern learned or recognized by the neurons.



- If the teaching has not started yet, you need to ensure that the ROI size fits a maximum of 256 blocks. Note that the blocks do not have to be square.
- If the teaching has started and neurons are committed, you must ensure that the ROI size fits a maximum of 256 blocks but also that the distribution of blocks remain the same within the new ROI. Failure to do so will prevent the recognition of the new ROI with the existing neurons.

Learn adjacent neighbors


When this option is checked, clicking a category button teaches the current ROI, but also its two adjacent neighbors in such manner that it creates three conservative and close models of the example. This option creates neurons with small influence fields.



The adjacent neighbors are the upper left and lower right neighbor of the ROI taken at a distance specified by the “distance to neighbor” slider. This distance is expressed in pixel and cannot exceed the full width of the ROI.

Learn surrounding background

When this option is checked, clicking a category button teaches the current ROI, but also teaches its two adjacent neighbors as background examples. This option prevents the neurons from over-generalizing, but it must be used when the adjacent neighbors are indeed background examples. Note that this operation can be exercised cautiously and manually by moving ROI in the image and clicking the Forget button.




The adjacent neighbors are the upper left and lower right neighbor of the ROI taken at a distance specified by the “distance to neighbor” slider. This distance is expressed in pixel and cannot exceed the full width of the ROI.

Similarity tolerance

Reducing the Similarity tolerance helps prevent the over-generalization of the neurons using the same constraints for all of them. The slider defines the value of the maximum influence field authorized when a new neuron gets committed. The smaller this value the more conservative the recognition engine, but also the more teaching will be required.

Survey the recognition of a region of search

The Survey button and AutoProcess checkbox both launch the scanning of a Region Of Search (ROS) shown in blue and the reporting of all the locations where the ROI has been recognized. The color of the spots is the color of the recognized category if different from 0.



This function is especially useful to decide if more examples are required to improve the accuracy of the recognition or if background examples should be taught to prevent the false hits.

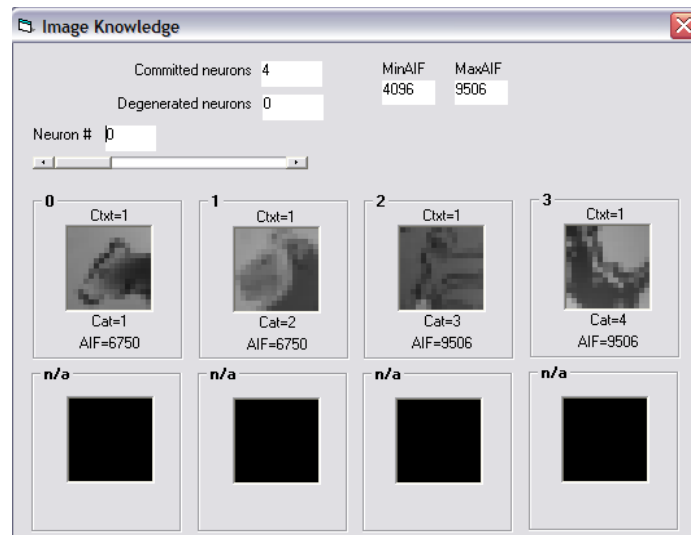
The survey of an image can take a few seconds depending on the size of the Region Of Search (ROS) and the value of the scanning step. A scanning step of (1,1) means that the ROI is recognized at each pixel inside the ROS. A scanning step of (2,4) means that the ROI is recognized every 2 columns and every 4 lines within the ROS.

The Region Of Search (ROS) can be changed with the mouse cursor when the AutoProcess checkbox is selected.

Remark: if the result of a survey does not show any recognized location, verify that the scanning step is not skipping too many positions within the ROS, and thus missing some important detection.

View the knowledge content

The View Knowledge menu allows reviewing the contents of the neurons, the pattern they hold, their category (Cat), and active influence field (AIF). This utility can be helpful to understand why some erroneous or uncertain classifications occur.



The AIF is a dynamic attribute entirely controlled by the neuron itself. From the time the neuron is created, its influence field can only reduce as new examples are taught and the knowledge gets more accurate. The smaller the influence field the more conservative the neuron.

Image Collection

Teaching and validation are the essence of building a knowledge which can recognize your objects accurately. For this purpose, it is recommended to teach and test the recognition engine against many images showing the variety of the objects and/or their surrounding and environmental conditions.

The Collect function lets you record live images and store them in folders for later use, either to teach or to verify proper recognition.

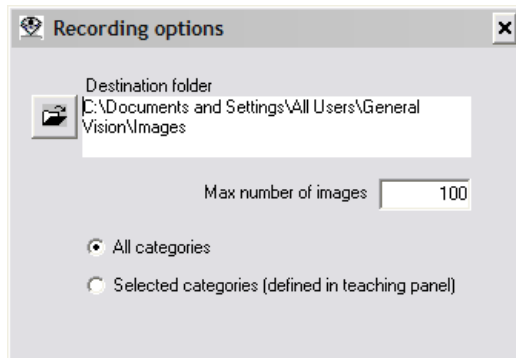
Examples of variability factors of an application:

- Day/Night
- Ambient lighting and shades
- Distances to the object
- Angular positions
- Changing background

Instructions

Recording for the first time since the application was started

The Recording options button gives you access to the following panel in which you can define which images must be saved, how many and where.

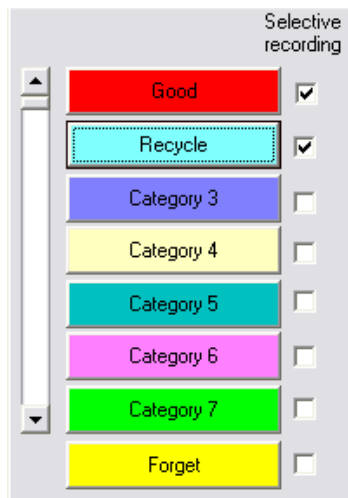


You can change the current destination folder by clicking the Open folder button.

Two settings can be adjusted:

(2) The maximum number of images to record per session. A session is initiated when you check the Record box in the main panel, but it can be stopped manually by checking the Record box again. The number of recorded images is updated in the status bar displayed below the image window.


(3) The type of images to record which can be (a) all images regardless of the category detected, or (b) only the images where specific categories are detected.



The selection of the categories to record is made in the teaching panel as shown to the left. Note that the default configuration of a selective recording excludes the images in which nothing is recognized. If you wish to include these images, you must check the box next to the Forget button.

Under selective recording the images are saved in separate sub-folders labeled with the category name.

Example:



Recording


The recording of the images is initiated by checking the Record box in the main panel. The number of recorded images is updated in the status bar displayed below the image window. The image recording will end automatically when the number of images requested in the Record Settings panel is reached, unless you have interrupted it before by checking the Record box again.

Collecting images for the training in the absence of external trigger

Continuous images collection can produce large amounts of images in a few minutes (at the rate of 4 to 6 images per second).

For industrial applications where the objects are passing on a conveyor belt, it can be easy to compensate for the lack of an external trigger (pulsed by a photocell) by training the camera to recognize when the edge of the object or a significant part of the object passes in a given region. Please note that this training is temporary for the duration of the image collection and has nothing to do with the final training which will define if the object is suitable or not based on some specific features.

Step-byStep instruction to record only interesting images

- Start a new project by selecting the menu Project/New
- Go to the Recording option menu and select a new destination folder ("Today_Collection" for example) and check the "Selected Categories" button.
- Return to the main menu and click the Teaching panel
- Change the label Category1 as GoodToRecord
- Click at the Snapshot button () to freeze an image in which the object to inspect is present as a whole or partially as long as its feature of interest can be seen.
- In the image window, adjust the yellow rectangle to include the edge of the object or another easily detectable feature of the object
- Go back to the Live mode by checking the Live box and observe when the category 'GoodToRecord' appears on the screen
- Without moving the yellow rectangle teach more examples of GoodToRecord if needed
- Without moving the yellow rectangle click the Forget button if the label GoodToRecord in the image appears when the object or its feature of interest is not seen
- When satisfied with recognition reported on the screen, check the Record box in the main panel

Note that the accuracy of this teaching session can be somehow "loose" since it is just intended to avoid recording images which will be useless for the teaching of the real inspection.

Example:

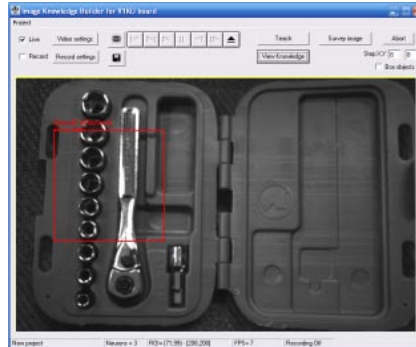
Final inspection = Verification that all the pieces are present in the toolbox. The final application will inspect a large area of the left flap.

Collection of images for the training = Record images showing the left flap.

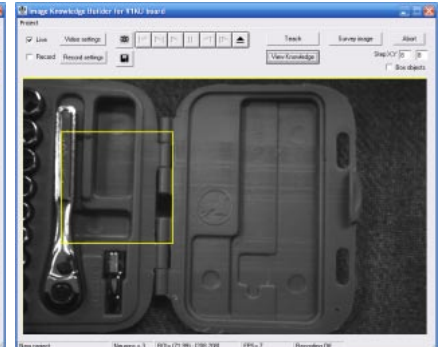
Region taught as
GoodToRecord



Region taught as
GoodToRecord



Region taught as Forget



When you have collected enough images, you can start a new project and use the images in the "GoodToRecord" folder to teach how you want the objects recognized and classified in the real inspection.

Collecting images for validation

When you are satisfied with the accuracy of the recognition seen on screen, you may want to validate results in a more thorough manner through the selective recording of the categories of objects over a significant period of time.

Step-by-Step instruction to record only interesting images

- Load your project by selecting the menu Project/Load
- Go to the Recording option menu and select a new destination folder called Today_Verification for example and check the "Selected Categories" button.
- Return to the main menu and click the Teaching panel
- If applicable, uncheck some categories of objects which do not interest you
- If applicable*, check the box next to the Forget button
- Check the Record box in the main panel

*If the acquisition of the images is made using an external trigger, the detection of an image showing an unknown or null category probably means that the part is missing or the knowledge is lacking some additional teaching. Recording these images can be useful.

Project Management

Projects defined under Image_Knowledge_Builder can be saved to disk in a project file or to the Flash memory of the camera.

A project is composed of two files associated to the same rootname:

- A project file (*.prj) containing the description of the recognition engine
- A settings file (*_settings.csv) with the description of the GUI preferences

Project file

- The *.prj file describes the recognition engine. It contains all the information necessary to resume the recognition of the images on an embedded system including the knowledge, the region of interest, the region of search and sensor settings
-

Settings file

- *_settings.csv describes the settings of the graphical user interface including the names and colors associated to the category values, as well as the preferred learning and recording settings.

Distributing projects

Missing settings file

When distributing project files for usage on another system, it is recommended to distribute both the project and the settings files together. If the settings file is missing the recognition engine will operate normally except that the program will report the recognized categories as numeric values. Also, if you decide to expand or tune the knowledge base by teaching new examples, the Teaching panel will show the default settings.

Inter-platform distribution

A project file generated initially when the IKB runs in simulation mode will not include the settings of the sensor. Usage of such file when IKB is connected to a V1KU board might require that you adjust the sensor settings for a proper recognition. These settings should produce the same type of contrast and brightness than in the image files used to teach the initial knowledge.

A project file generated initially when the IKB is interfaced to a V1KU board will include the settings of the sensor. These settings will be kept if you later update the project file while running in simulation mode.