


Open eVision



Terms of Use

EURESYS s.a. shall retain all property rights, title and interest of the documentation of the hardware and the software, and of the trademarks of EURESYS s.a.

All the names of companies and products mentioned in the documentation may be the trademarks of their respective owners.

The licensing, use, leasing, loaning, translation, reproduction, copying or modification of the hardware or the software, brands or documentation of EURESYS s.a. contained in this book, is not allowed without prior notice.

EURESYS s.a. may modify the product specification or change the information given in this documentation at any time, at its discretion, and without prior notice.

EURESYS s.a. shall not be liable for any loss of or damage to revenues, profits, goodwill, data, information systems or other special, incidental, indirect, consequential or punitive damages of any kind arising in connection with the use of the hardware or the software of EURESYS s.a. or resulting of omissions or errors in this documentation.

This documentation is provided with Open eVision 2.6.1 (doc build 1110).
© 2018 EURESYS s.a.

Contents

PART I : INSTALLING OPEN EVISION	6
1. Installing Open eVision	7
PART II : MANIPULATING PIXELS CONTAINERS AND FILES	17
1. Manipulating Pixels Containers and Files	18
1.1. Pixel Container File Save	18
1.2. Pixel Container File Load	20
1.3. Memory Allocation	20
1.4. Image and Depth Map Buffer	21
1.5. Image Drawing and Overlay	24
1.6. 3D Rendering of 2D Images	25
1.7. Vector Types and Main Properties	26
1.8. ROI Main Properties	30
1.9. Flexible Masks	31
1.10. Profile	35
PART III : IMAGE PRE-PROCESSING LIBRARIES	37
1. EasyImage - Pre-Processing Grayscale Image	38
1.1. Intensity Transformation	38
1.2. Thresholding	41
1.3. Arithmetic and Logic	43
1.4. Non-Linear Filtering	45
1.5. Geometric Transforms	51
1.6. Noise Reduction and Estimation	53
1.7. Scalar Gradient	56
1.8. Vector Operations	56
1.9. Canny Edge Detector	58
1.10. Harris Corner Detector	60
1.11. Overlay	61
1.12. Operations on Interlaced Video Frames	62
1.13. Flexible Masks in EasyImage	62
2. EasyColor - Pre-processing Color Images	64
2.1. Bayer Transform	68
2.2. LUT for Gain/Offset (Color)	69
2.3. LUT for Color Calibration	70
2.4. LUT for Color Balance	70
3. EasyImage - Computing Image Statistics	73
4. EasyDeepLearning - Classifying Images	78
4.1. What EasyDeepLearning can do	78

4.2. Workflow	78
4.3. EasyDeepLearning Studio	80
4.4. Managing the Dataset	82
4.5. Using the Classifier	84
4.6. Accelerating the Computation with a NVIDIA GPU	87
PART IV : INSPECTION TOOLS	89
1. EasyObject - Analysing Blobs	90
1.1. Image Segmenters	93
1.2. Image Encoder	97
1.3. Holes Construction	99
1.4. Normal vs. Continuous Mode	101
1.5. Selecting and Sorting Blobs	104
1.6. Advanced Features	105
Computable Features	105
Draw Coded Elements	110
Flexible Masks in EasyObject	111
2. EasyGauge - Measuring down to Sub-Pixel	114
Workflow	114
Gauge definitions	115
Find transition points using peak analysis	118
Find shapes using geometric models	123
Gauge Manipulation: Draw, Drag, Plot, Group	124
Calibration and Transformation	125
Calibration using EWorldShape	127
Advanced Features	129
3. EasyMatch - Matching Area Patterns	134
Workflow	134
Learning Process	135
Matching Process	137
Advanced Features	137
4. EasyFind - Matching Geometric Patterns	139
4.1. Workflow	139
4.2. Learning Process	141
4.3. Finding Process	144
4.4. Advanced Features	144
5. Golden Template Validation (EChecker)	147
5.1. Image Comparison	152
PART V : TEXT IDENTIFICATION TOOLS	155
1. EasyOCR - Reading Texts	156
Workflow	156
Learning Process	157
Segmenting	157
Recognition	158
2. EasyOCR2 - Reading Texts (Improved)	161

3. EasyOCV - Validating Texts	169
3.1. Learning Passes	173
Inspect and compare image with model	173
3.2. Degrees of Freedom	173
3.3. Quality Indicators	176
3.4. Advanced Features	177
3.5. Programming with EasyOCV	178
PART VI : CODE IDENTIFICATION TOOLS	184
1. EasyBarCode - Reading Bar Codes	185
1.1. Reading Bar Codes	185
1.2. Reading Mail Bar Codes	189
2. EasyMatrixCode - Reading Matrix Codes	193
2.1. EasyMatrixCode vs EasyMatrixCode2	193
2.2. Specifications	194
2.3. Reading a Matrix Code	195
2.4. Learning a Matrix Code	197
3. EasyQRCode - Reading QR Codes	199
Workflow	199
QR code definition	200
Read a QR code	203
Advanced features	203
PART VII : 3D TOOLS	205
1. Understanding 3D Concepts	206
1.1. Basic Concepts	206
1.2. Laser Triangulation	209
1.3. The Laser Line 3D Acquisition Pipeline	211
2. Easy3D - Using 3D Toolset	213
2.1. Laser Line Extraction	213
2.2. Calibration	215
2.3. Point Cloud	220
Coordinates Transformations	220
Reducing a Point Cloud	221
Managing Planes	222
Aligning	224
2.4. Mesh	228
2.5. ZMap	229
Generating a ZMap	229
Managing the Coordinates	231
Static Methods	232
2.6. 3D Viewer	236

PART I
INSTALLING OPEN EVISION

1. Installing Open eVision

Installer Package

Open eVision comes as a single installer package "**Open_eVision_Installer_2.1.0.msi**". It contains everything needed to run or develop applications using **Open eVision**

Installation Types

The Open eVision Installer provides the following installation types:

- **Complete:** Everything needed for running or developing applications is installed on the system.
- **Typical:** Same as Complete, with the exception of Legacy components and VC++ 6.0 specific components.
- **Runtime:** Installs all binaries needed to run applications using Open eVision on the system.
- **Custom:** Allows to select exactly what components will be installed on the system.

Older Versions

Open eVision will not replace other Open eVision major versions, but install alongside. If the major version is identical, minor versions releases as well as maintenance releases will update automatically

Command-Line Interface

To install Open eVision with the command line, use:

```
msiexec /i "Open_eVision_Installer_2.1.0.msi" /qn INSTALLTYPE=[install_type]"
```

Where *[install_type]* can be **Complete**, **Typical** or **Runtime**. By default, installation type is 'Typical'.

For the command prompt to wait for the end of the installation add '*start /wait*' at the start of the command:

```
start /wait msiexec /i "Open_eVision_Installer_2.1.0.msi" /qn INSTALLTYPE=[install_type]"
```

License Activation

Open eVision licenses are activated from the Open eVision License Manager. The License Manager can be launched at the end of the installation, or from the Windows start menu.

Note: On Windows XP, the license Manager requires .NET 2.0.

Supported platforms and requirements

WES 2009

Windows Embedded Standard 2009 can install drivers and applications after FBA completes.

It is recommend to install Open eVision as follows:

- Add mandatory components to the Run-Time Image using Target Designer**
 - **.Net framework 3.0 setup** component: For Open eVision license manager
 - **Windows Installer Service** component: For Open eVision to install C/C++ run-time libraries.
 - Open eVision license manager needs either the internet support components for online activation, or USB Flash for an offline activation:.
 - **Sysprep** (Windows System preparation) **Component** or **System Cloning Tool Component**: To reseal the "**Master Target**" before deploying the image to multiple devices.
Note: To prevent **System Cloning Tool** from executing **FBRESEAL** automatically when FBA finishes:
Set System Cloning Tool Settings **Reseal Phase** to `Manual`, or change **System Cloning Tool Advanced Settings** `cmiResealPhase` from **12000** to **0**.
- Install Open eVision 2.1 and newer on the "Master Target"**
 - Once the Pre FBA OS Image has been built, boot the "Master Target" and allow FBA to complete.
 - When the "Master Target" has been booted for the second time, install the Open eVision libraries using the standard installer provided.
 - Optionally, install your own final application based on Open eVision, include all the run-time libraries needed by your application.
- Reseal the master package ready for mass deployment**
 - Run **FBRESEAL** or **Sysprep**. Once the computer shuts down, this image is the master.
 - Each time you deploy the "Master" in a new device, the Open eVision libraries need to be activated as described in the license manager documentation.

Open eVision in C++

Include the header (`Open_eVSION_2_0.h`) located in the installation folder > Include subfolder. No linker settings are required.

Microsoft Visual Studio C++ environments automatically adds the Open eVision Include folder at installation time. This must be done manually for Borland/CodeGear C++ environments.

Open eVision in .NET

Add a reference to the `Open_eVSION_NetApi_2_1.dll` in the development environment. No other DLL must be copied.

Using Open eVision in ActiveX

Add a reference to the `Open_eVSION_ActiveXApi_2_1.dll` component.

Visual Studio 6.0

- If you are using the regular API (new style API with exceptions and namespaces)
 1. Open your project settings, and add the following preprocessor macro definition:
`DO_NOT_USE_INLINE_OPEN_EVISION_2_1`
 2. Add the **Open_evVSION_2_0_VC6_Release.lib** and **Open_evVSION_2_0_VC6_Debug** in the corresponding configuration linker settings.
These files are in the Open eVision installation folder.
- If you are using the legacy support API (compatible with eVision 6.7.1):
 1. Open your project settings, and add the following preprocessor macro definition:
`DO_NOT_USE_INLINE_LEGACY_OPEN_EVISION_2_1`
 2. Add the **Legacy_Open_evVSION_2_0_VC6_Release.lib** and **Legacy_Open_evVSION_2_0_VC6_Debug** in the corresponding configuration linker settings.
These files are in the Open eVision installation folder.
- If you are using both the regular API and the legacy support API, you must perform all steps (and thus all the relevant libraries to your solution).

In order to use these libraries, your program must use the Multithreaded DLL (/MD) or Multithreaded Debug DLL (/MDd) code generation flags.

Visual Basic 6.0

Add a reference to the `Open_eVision_ActiveXApi_2_1.dll` (menu: "Project > Add Reference"). All objects are then directly usable in Visual Basic.

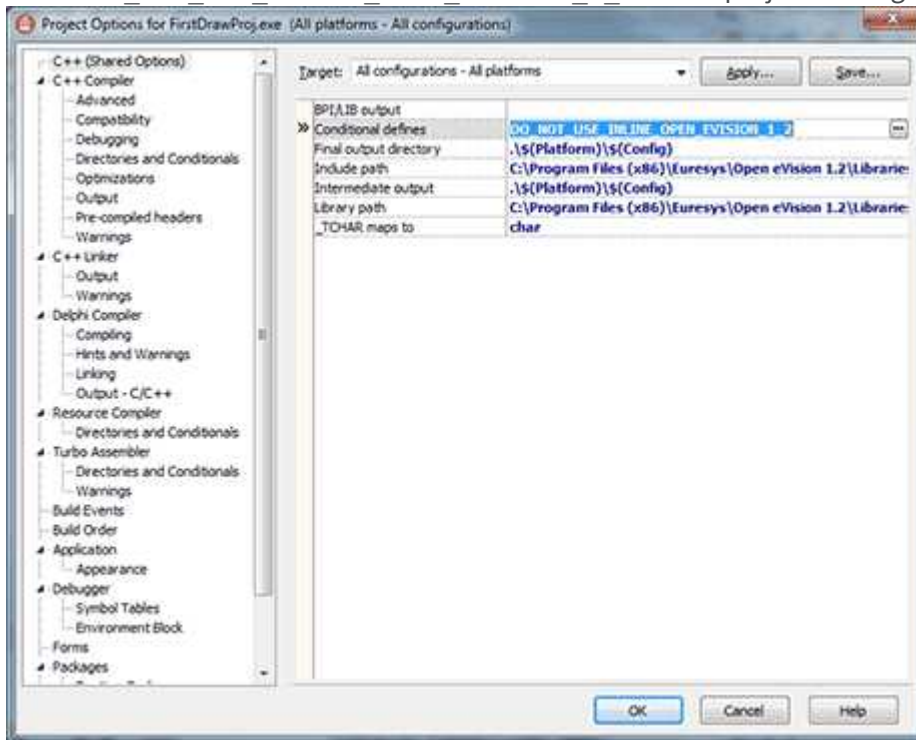
Embarcadero RAD Studio XE4/XE5

To configure projects in Embarcadero RAD Studio XE4 and XE5 to use Open eVision:

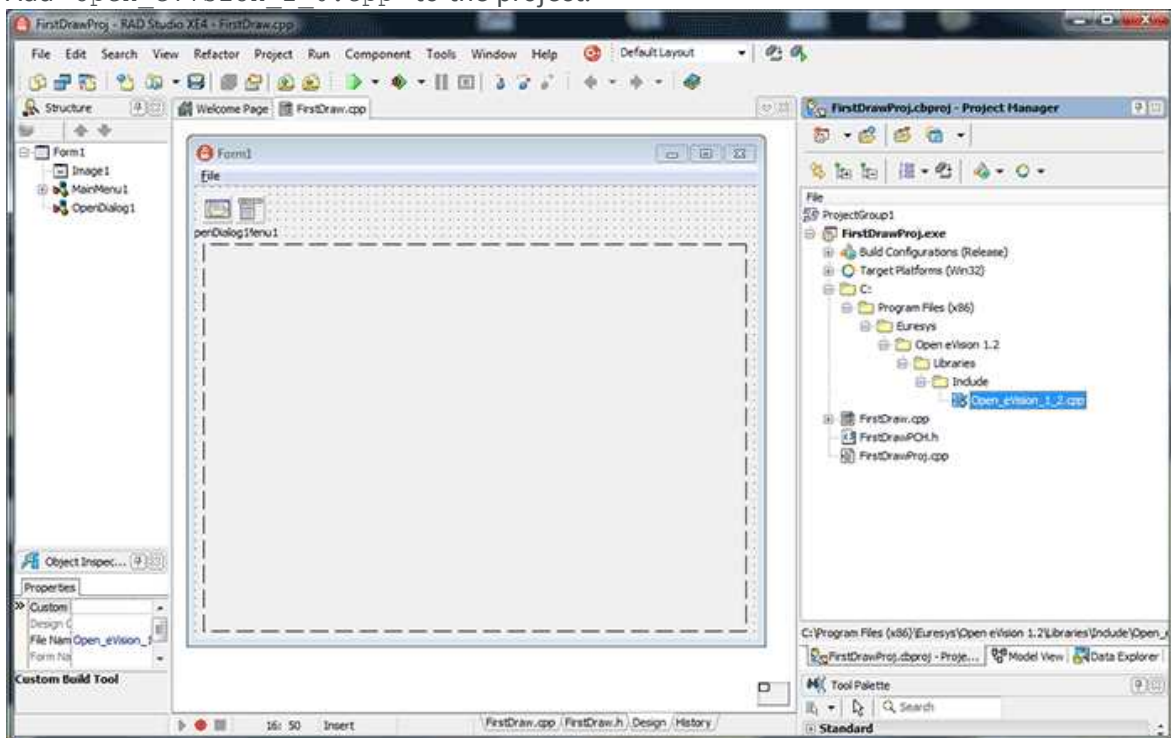
C++

1. Create a new C++ project.
2. Add the Open eVision include path to the project dependencies.

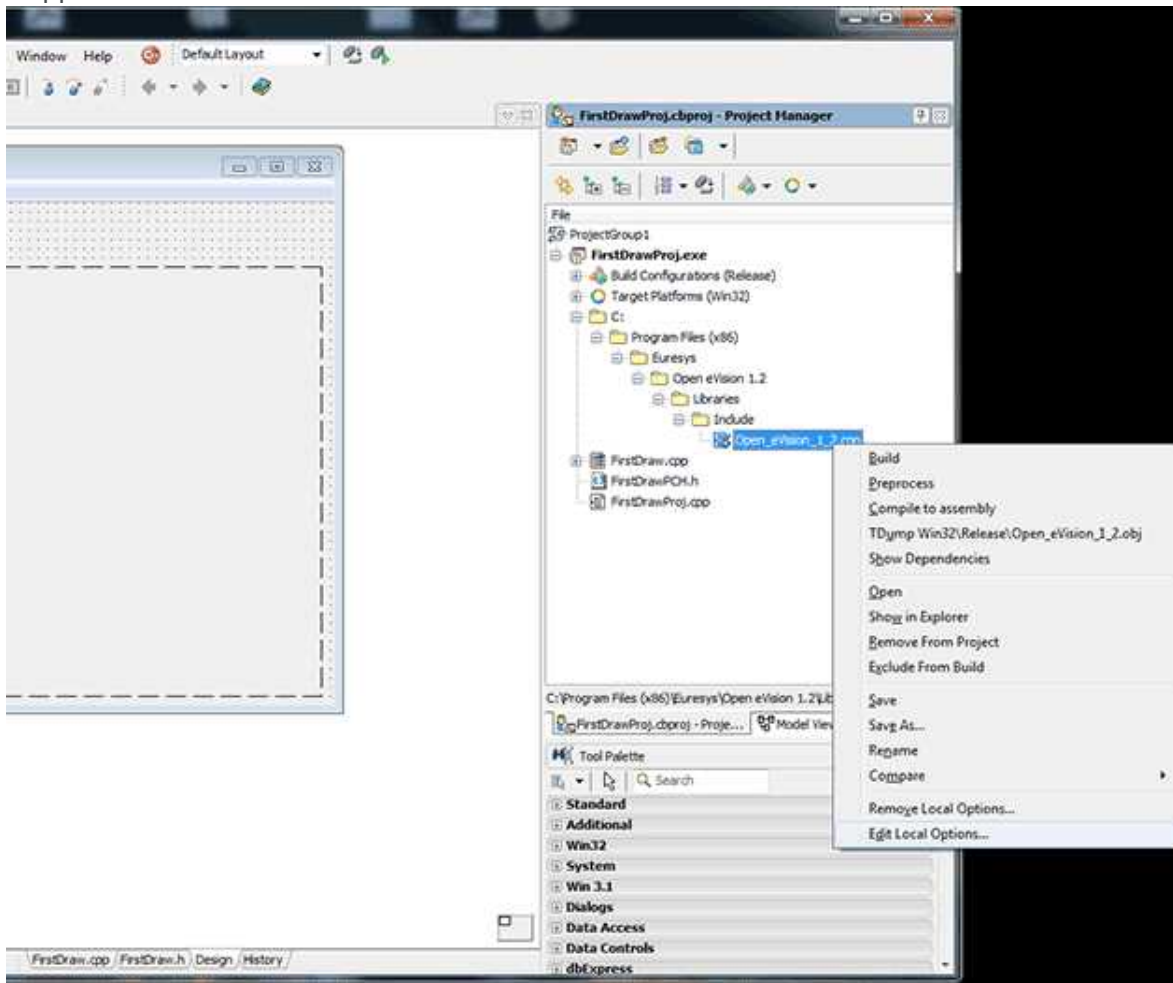
- Define `DO_NOT_USE_INLINE_OPEN_EVISION_2_1` in the project settings.



- Include "Open_evVSION_2_0.h" in the .cpp files where you want to use Open eVision.
- Add "Open_evVSION_2_0.cpp" to the project.

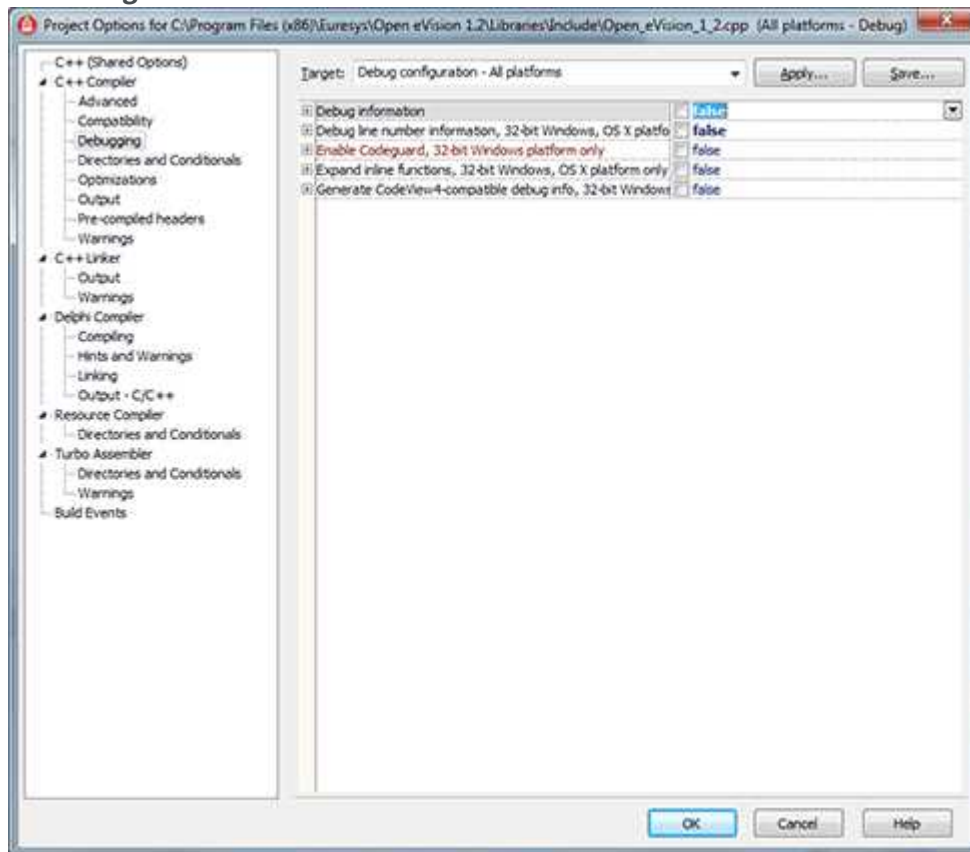


6. In both Debug and Release, modify the local C++ compiler options of the "Open_eVision_2_0.cpp" file:

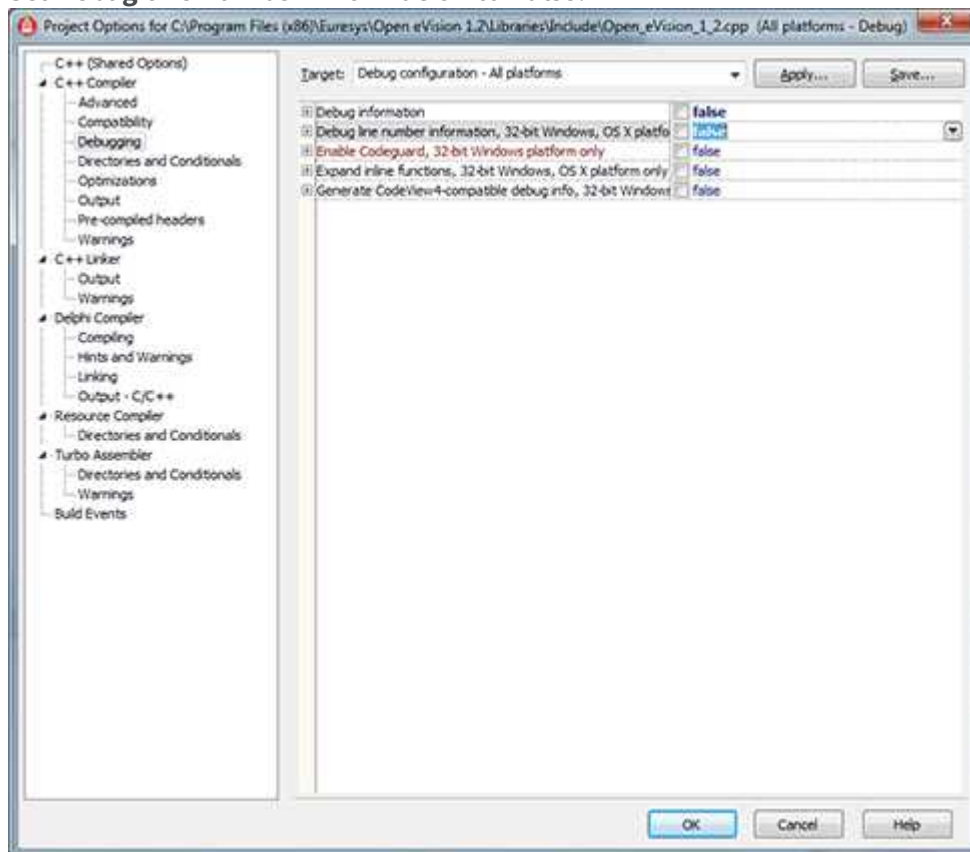


7. In **Debugging:**

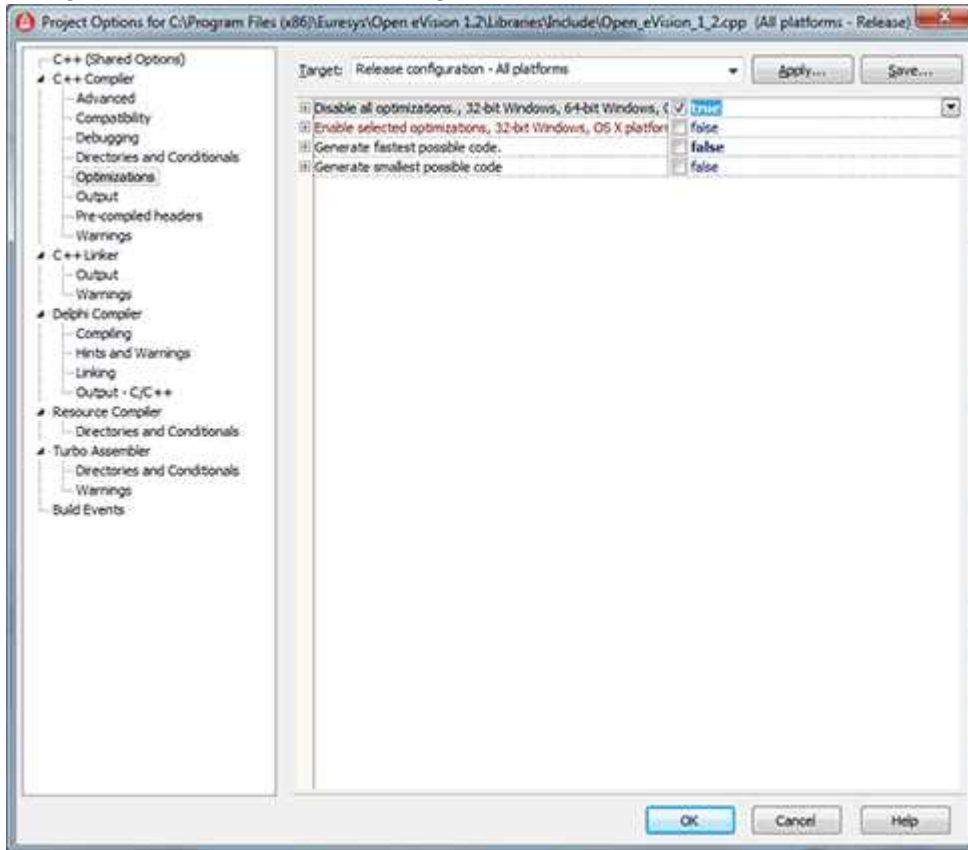
- a. Set **Debug information** to **False**.



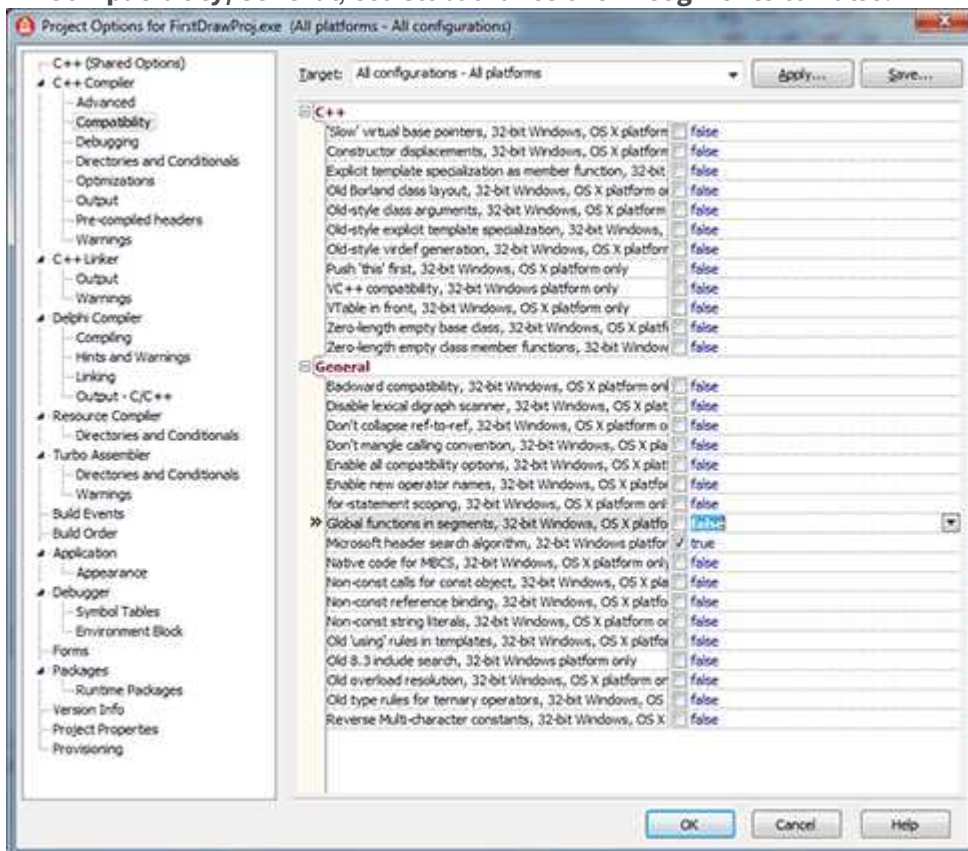
- b. Set **Debug line number information** to **False**.



8. In **Optimization**, set **Disable all optimizations** to **True**.



9. In **Compatibility/General**, set **Global functions in segments** to **False**.

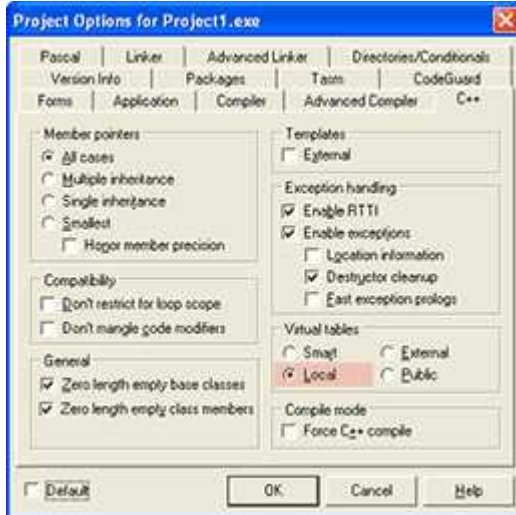


Borland C++

An error may occur due to a limitation in the number of functions (and virtual functions) in a single translation unit:

BCB6 Error E2491: Maximum VIRDEF count exceeded; check for recursion

If this problem occurs, change the **Virtual tables C++** option to **Local**:



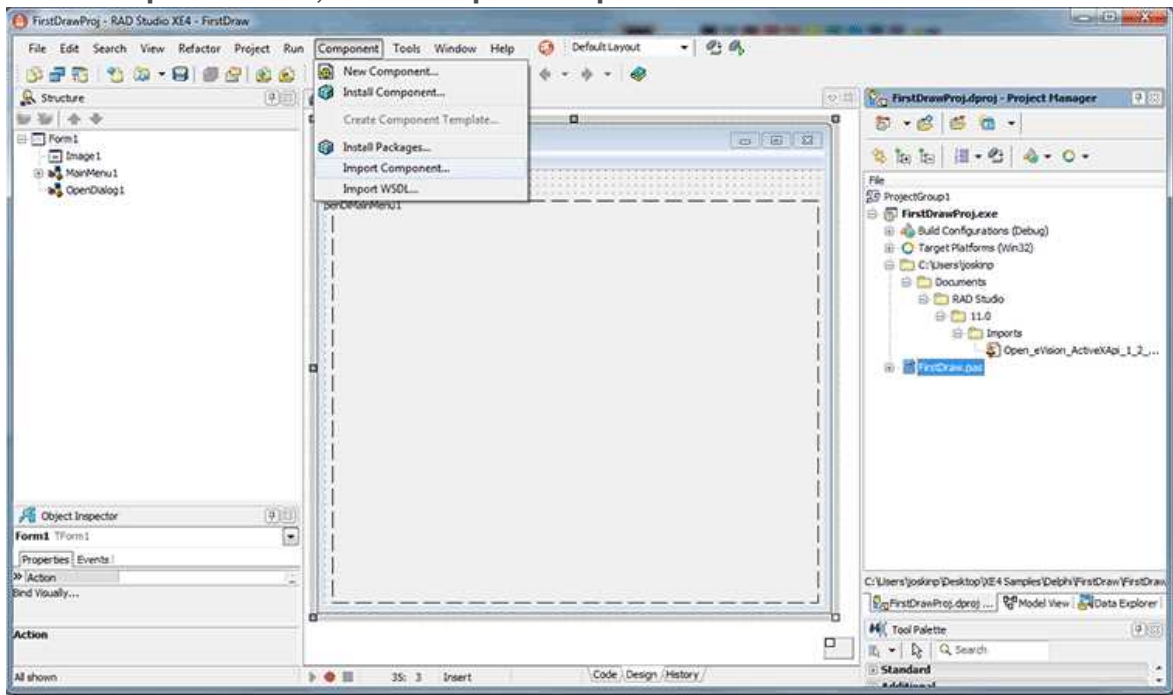
When using Open eVision objects as members of Borland GUI classes, like `TDialog` for instance, an unhandled exception can occur when the application is closed. To avoid this issue, create Open eVision objects dynamically using `new` and delete them in the destructor of the parent class.

Delphi

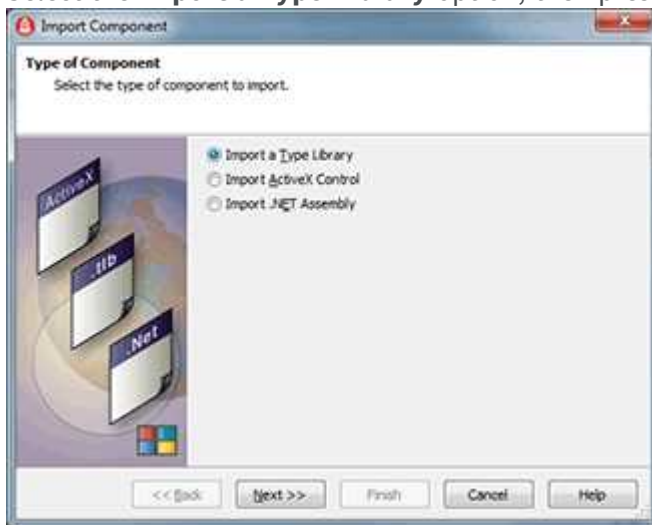
Note: Open eVision must be previously installed with the **Legacy ActiveX** option.

You must use the ActiveX DLL ("Component > Import Component... > Import .NET Assembly").

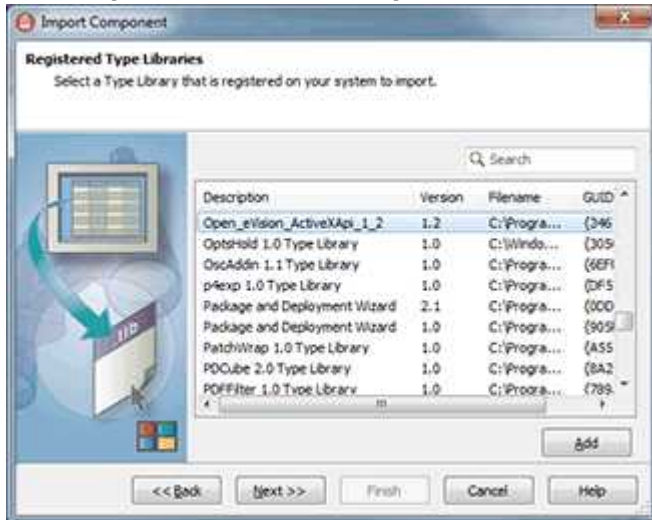
1. Create a new Delphi project.
2. In the **Component** menu, click on **Import Component**.



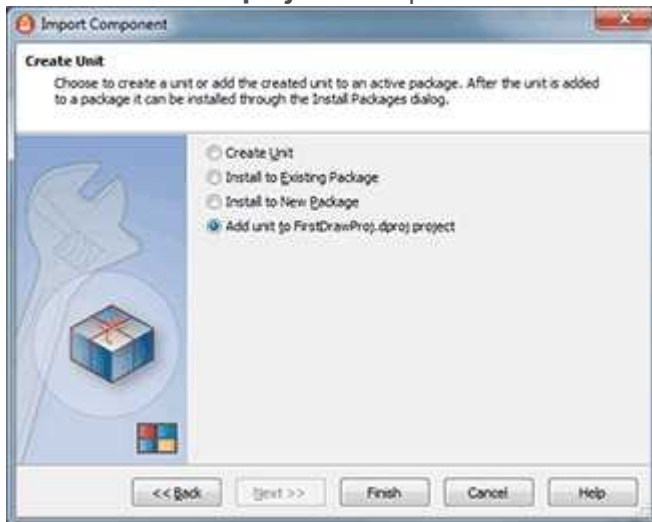
3. Select the **Import a Type Library** option, then press **Next**.



4. Select **Open_eVision_ActiveXApi_2_1** in the list, then press **Next**.



5. Don't change anything on the next form and press **Next**.
6. Select **Add unit to project** then press **Finish**.



PART II

*MANIPULATING PIXELS CONTAINERS
AND FILES*

1. Manipulating Pixels Containers and Files

1.1. Pixel Container File Save

Images and Depth Maps

The `Save` method of an image or the `SaveImage` method of a depth map or a ZMap saves the image data of an image or of a depth map or a ZMap object into a file using two arguments:

- Path: path, filename, and file name extension.
- Image File Type. If omitted, the file name extension is used.

Images bigger than 65,536 (either width or height) must be saved in Open eVision proprietary format.

Save throws an exception when:

- The requested image file format is incompatible with the image pixel types
- The Auto file type selection method and the file name extension is not supported

When saving a 16-bit depth map, the fixed point precision is lost and the pixels are considered as 16-bit integers.

image file type arguments

Argument	Image File Type
<code>EImageFileType_Auto(*)</code>	Automatically determined by the filename extension. See below.
<code>EImageFileType_Euresys</code>	Open eVision Serialization.
<code>EImageFileType_Bmp</code>	Windows bitmap - BMP
<code>EImageFileType_Jpeg</code>	JPEG File Interchange Format - JFIF
<code>EImageFileType_Jpeg2000</code>	JPEG 2000 File format/Code Stream - JPEG2000
<code>EImageFileType_Png</code>	Portable Network Graphics - PNG
<code>EImageFileType_Tiff</code>	Tagged Image File Format - TIFF

(*) Default value.

Assigned image file type if argument is `ImageFileType_Auto` or missing

File name extension(*)	Automatically assigned image file type
BMP	Windows Bitmap Format
JPEG, JPG	JPEG File Interchange Format - JFIF
JP2	JPEG 2000 file format
J2K, J2C	JPEG 2000 Code Stream
PNG	Portable Network Graphics
TIFF, TIF	Tagged Image File Format

(*) Case-insensitive.

Saving JPEG and JPEG2000 lossy compressions

`SaveJpeg` and `SaveJpeg2K` specify the compression quality when `saving` compressed images. They have two arguments:

- Path: a string of characters including the path, filename, and file name extension.
- Compression quality of the image file, an integer value in range [0: 100].
`SaveJpeg` saves image data using JPEG File Interchange Format – JFIF.
`SaveJpeg2K` saves image data using JPEG 2000 File format.

JPEG compression values

JPEG compression	Description
JPEG_DEFAULT_QUALITY (-1)	Default quality (*)
100	Superb image quality, lowest compression factor
75	Good image quality (*)
50	Normal image quality
25	Average image quality
10	Bad Image quality

(*) The default quality corresponds to the good image quality (75).

Representative JPEG 2000 compression quality values

JPEG 2000 compression	Description
-1	Default quality (*)
1	Highest image quality, lowest compression factor

JPEG 2000 compression	Description
16	Good Image Quality (*) (16:1 rate)
512	Lowest image quality, highest compression factor

(*) The default quality corresponds to the good image quality (16:1 rate).

Point Clouds

- Use the `Save` method to save the point cloud in Open eVision proprietary file format.
- Use the `SavePCD` method to save the point cloud in a file compatible with other software such as PCL (Point Cloud Library).

1.2. Pixel Container File Load

Images and Depth Maps

- Use the `Load` method to load image data into an image object:
 - It has one argument: the **path**: path, filename, and file name extension.
 - File type is determined by the file format.
 - The destination image is automatically resized according to the size of the image on disk.
- The `Load` method throws an exception when:
 - File type identification fails
 - File type is incompatible with pixel type of the image object

Serialized image files of Open eVision 1.1 and newer are incompatible with serialized image files of previous Open eVision versions.

When loading a BW16 image (with integer values) in a depth map, the fixed point precision set in the depth map (0 by default) is left unchanged and used.

Point Clouds

- Use the `Load` method to save the point cloud in Open eVision proprietary file format.
- Use the `LoadPCD` method to save the point cloud in a file compatible with other software such as PCL (Point Cloud Library).

1.3. Memory Allocation

An image can be constructed with an internal or external memory allocation.

Internal Memory Allocation

The image object dynamically allocates and unallocates a buffer. Memory management is transparent.

When the image size changes, re-allocation occurs.

When an image object is destroyed, the buffer is unallocated.

To declare an image with internal memory allocation:

1. Construct an image object, for instance `EImageBW8`, either with width and height arguments, OR using the `SetSize` function.
2. Access a given pixel. There are several functions that do this. `GetImagePtr` returns a pointer to the first byte of the pixel at given coordinates.

External Memory Allocation

The user controls `buffer` allocation, or [links a third-party image](#) in the memory buffer to an Open eVision image.

Image size and buffer address must be specified.

When an image object is destroyed, the buffer is unaffected.

To declare an image with external memory allocation:

1. Declare an image object, for instance `EImageBW8`.
2. Create a suitably sized and aligned buffer (see [Image Buffer](#)).
3. Set the image size with the `SetSize` function.
4. Access the buffer with `GetImagePtr`. See also [Retrieving Pixel Values](#).

1.4. Image and Depth Map Buffer

Image and depth map pixels are stored contiguously, from top row to bottom, from left to right, in Windows bitmap format (top-down [DIB¹](#)) into an associated buffer.

The buffer address is a pointer to the start address of the buffer, which contains the top left pixel of the image.

¹device-independent bitmap




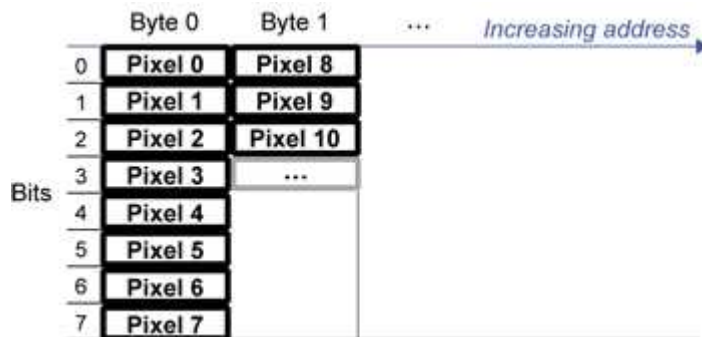
Image Buffer pitch

- Alignment must be a multiple of 4 bytes.
- Open eVision 1.2 onwards default pitch is 32 bytes for performance reasons (Open eVision 1.1.5 was 8 bytes).

Memory Layout


- `EImageBW1` stores 8 pixels in one byte.

Example memory layout of the first 2 pixels of a BW1 image buffer:




- `EImageBW8` and `EDepthMap8` store each pixel in one byte.

Example memory layout of the first pixels of a BW8 image buffer:




- `EImageBW16` stores each pixel in a 16-bit word (two bytes).

Example memory layout of the first pixels of a BW16 image buffer:




- [EImageC15](#) stores each pixel in 2 bytes. Each color component is coded with 5-bits. The 16th bit is left unused.

Example memory layout of the first pixels of a C15 image buffer:




- [EImageC16](#) stores each pixel in 2 bytes. The first and third color components are coded with 5-bits. The second color component is coded with 6-bits.

Example memory layout of the first pixels of a C16 image buffer:




- [EDepthMap16](#) store each pixel in 2 bytes using a fixed point format.
- [EImageC24](#) stores each pixel in 3 bytes. Each color component is coded with 8-bits.

Example memory layout of the first pixels of a C24 image buffer:



- `EImageC24A` stores each pixel in 4 bytes. Each color component is coded with 8-bits. The alpha channel is also coded with 8-bits.

Example memory layout of the first pixels of a C24A image buffer:



- `EDepthMap32f` store each pixel in 4 bytes using a float format.

1.5. Image Drawing and Overlay

- Drawing uses Windows `GDI`¹ system calls. `MFC`² applications normally use `OnDraw` event handler to draw, where a pointer to a device context is available. Borland/CodeGear's OWL or VCL use a **Paint** event handler.
- The color palette in 256-color display mode gives optimal rendering. Gray-level images can be improved using `LUT`³s (using histogram stretching techniques or pseudo-coloring).
- The zoom can be different horizontally and vertically.
- `DrawFrameWithCurrentPen` method draws a frame.
- **Non-destructive overlaying** drawing operations do not alter the image contents, such as `MoveTo/LineTo`.

¹Graphics Device Interface

²Microsoft Foundation Class

³LookUp Tables


- **Destructive overlaying** drawing operations alter the image contents by drawing inside the image such as `Easy::OpenImageGraphicContext`. Gray-level [color] images can only receive a gray-level [color] overlay.

1.6. 3D Rendering of 2D Images

These images are viewed by rotating them around the X-axis, then the Y-axis.

Gray 3D Rendering

`Easy::Render3D` prepares a 3-dimensional rendering where gray-level values are altitudes. Magnification factors in the three directions (X = width, Y = height and Z = depth) can be given. The rendered image appears as independent dots whose size can be adjusted to make the surface more or less opaque.



3D rendering


Color Histogram 3D Rendering

`Easy::RenderColorHistogram` prepares a 3-dimensional rendering of a color image histogram.

The pixels are drawn in the RGB space (not XY-plane) to show clustering and dispersion of RGB values.

This function can process pixels in other color systems (using `EasyColor` to convert), but the raw RGB image is required to display the pixels in their usual colors.

Magnification factors in all three directions (X = red, Y = green and Z = blue) can be given.




Color histogram rendering

1.7. Vector Types and Main Properties

A vector is a one-dimensional array of pixels (taken from an image [profile](#) or contour).

[EVector](#) is the base class for all vectors. It contains all non-type-specific methods, mainly for counting elements and serialization.



Profile in a C24 image

RGB values plot along profile

RGB values array
([EC24Vector](#))


A vector manages an array of elements. Memory allocation is transparent, so vectors can be resized dynamically. Whenever a function uses a vector, the vector type, size and structure are automatically adjusted to suit the function needs.

The use of vectors is quite straightforward:

1. **Create a vector of the appropriate type**, using its constructor and pre-allocate elements if required.


Vector types

- [EBW8Vector](#): a sequence of gray-level pixel values, often extracted from an image profile (used by [EasyImage::Lut](#), [EasyImage::SetupEqualize](#), [EasyImage::ImageToLineSegment](#), [EasyImage::LineSegmentToImage](#), [EasyImage::ProfileDerivative](#), ...).




Graphical representation of an [EBW8Vector](#) (see [Draw method](#))

- [EBW16Vector](#): a sequence of gray-level pixel values, using an extended range (16 bits), mainly for intermediate computations.




Graphical representation of an [EBW16Vector](#)

- [EBW32Vector](#): a sequence of gray-level pixel values, using an extended range (32 bits), mainly for intermediate computations (used in [EasyImage::ProjectOnARow](#), [EasyImage::ProjectOnAColumn](#), ...).



Graphical representation of an [EBW32Vector](#)

- [EC24Vector](#): a sequence of color pixel values, often extracted from an image profile (used by [EasyImage::ImageToLineSegment](#), [EasyImage::LineSegmentToImage](#), [EasyImage::ProfileDerivative](#), ...).




Graphical representation of an [EC24Vector](#)

- [EBW8PathVector](#): a sequence of gray-level pixel values, extracted from an image profile or contour, with corresponding pixel coordinates (used by [EasyImage::ImageToPath](#), [EasyImage::PathToImage](#), ...).




Graphical representation of an `EBW8PathVector` (see `Draw` method)

- `EBW16PathVector`: a sequence of gray-level pixel values, extracted from an image profile or contour, with corresponding pixel coordinates (used by `EasyImage::ImageToPath`, `EasyImage::PathToImage`, ...).




Graphical representation of an `EBW16PathVector` (see `Draw` method)

- `EC24PathVector`: a sequence of color pixel values, extracted from an image profile or contour, with corresponding pixel coordinates (used by `EasyImage::ImageToPath`, `EasyImage::PathToImage`, ...).




Graphical representation of an `EC24PathVector` (see `Draw` method)

- `EBWHistogramVector`: a sequence of frequency counts of pixels in a BW8 or BW16 image (used by `EasyImage::IsodataThreshold`, `EasyImage::Histogram`, `EasyImage::AnalyseHistogram`, `EasyImage::SetupEqualize`, ...).



Graphical representation of an `EBWHistogramVector` (see `Draw` method)

- `EPathVector`: a sequence of pixel coordinates. The corresponding pixels need not be contiguous (used by `EasyImage::PathToImage` and `EasyImage::Contour`).



Graphical representation of an `EPathVector` (see `Draw` method)

- `EPeakVector`: peaks found in an image profile (used by `EasyImage::GetProfilePeaks`).
 - `EColorVector`: a description of colors (used by `EasyColor::ClassAverages` and `EasyColor::ClassVariances`).
2. **Fill a vector with values.** First empty it, using the `EVector::Empty` member, then add elements one at a time by calling the `EC24Vector::AddElement` member. You can access any element by means of indexing.
 3. **Access a vector element**, either for reading or writing. Use the brackets operator, for instance, `EC24Vector::operator[]`.
 4. **Determine the current number of elements**, use member `EVector::NumElements`.
 5. **Draw the vector.**

A pixel vector is a plot of the element values as a function of the element index, so its graphical appearance depends on its type. You can draw a vector in a window. For legibility, the drawing should appear on a neutral background.

Drawing is done in the device context associated to the desired window. By default, curves are drawn in blue, annotations are drawn in black. The following parameters can be defined: `graphicContext`, `width`, `height`, `origin`, `color0`, `color1`, `color2`.

The `EC24Vector` has three curves drawn instead of one, each corresponding to a color component. By default, red, blue and green pens are used.

1.8. ROI Main Properties

ROIs are defined by a [width](#), a [height](#), and **origin x and y coordinates**.

The origins are specified with respect to the top left corner in the parent image or ROI.

The ROI must be wholly contained in its parent image.

The processing/analysis time of a BW1 ROI is faster if [OrgX](#) and [Width](#) are multiples of 8.

Save and load

You can [save](#) or [load](#) an ROI as a separate image, to be used as if it was a full image. The ROIs perform **no memory allocation** at all and never duplicate parts of their parent image, the parent image provides them with access to its image data.

The image size of the new file must match the size of the ROI being loaded into it. The image around the ROI remains unchanged.

ROI Classes

An Open eVision ROI inherits parameters from the abstract class [EBaseROI](#).

There are several ROI types, according to their pixel type. They have the same characteristics as the corresponding [image types](#).

- [EROIBW1](#)
- [EROIBW8](#)
- [EROIBW16](#)
- [EROIBW32](#)
- [EROIC15](#)
- [EROIC16](#)
- [EROIC24](#)
- [EROIC24A](#)

Attachment


An ROI must be [attached](#) to a parent (image/ROI) with parameters that set the parent, position and size, and these links are updated transparently, avoiding dangling pointers.

A normal image cannot be attached to another image or ROI.

Nesting

Set and Get functions change or query the width, height and position of the origin of an ROI, with respect to its immediate or topmost parent image.

An image may accommodate an arbitrary number of ROIs, which can be nested in a hierarchical way. Moving the ROI also moves the embedded ROIs accordingly. The image/ROI classes provide several methods to traverse the hierarchy of ROIs associated with an image.



Nested ROIs: Two sub-ROIs attached to an ROI, itself attached to the parent image

Cropping

`CropToImage` crops an ROI which is partially out of its image. The resized ROI never grows. An exception is thrown if a function attempts to use an ROI that has limits that extend outside of the parents.

Note: (In Open eVision 1.0.1 and earlier, an ROI was silently resized or repositioned when placed out of its image and sometimes grew. If ROI limits extended outside parents, they were silently resized to remain within parent limits.)

Resizing and moving

- ROIs can easily be resized and positioned by two functions and dragging handles:
 - `EBaseROI::Drag` adjusts the ROI coordinates while the cursor moves.
 - `EBaseROI::HitTest` informs if the cursor is placed over a dragging handle. Once the handle is known, the cursor shape can be changed by an `OnSetCursor` MFC event handler. `HitTest` is unpredictable if called while dragging is in progress. `HitTest` can be used in an `OnSetCursor` MFC event handler to change the cursor shape, or before a dragging operation like `OnLButtonDown`, (or `EvSetCursor` and `EvLButtonDown` in Borland/CodeGear's OWL) (or `FormMouseMove` and `FormMouseDown` in Borland/CodeGear's VCL). In VB6, `MouseDown`, `MouseMove`, `MouseUp` events return the current cursor position in twips rather than pixels, so conversion is mandatory.

1.9. Flexible Masks

ROIs vs flexible masks


ROIs and masks restrict processing to part of an image:

- "ROI Main Properties" on the previous page apply to all Open eVision functions. Using Regions of Interest accelerates processing by reducing the number of pixels. Open eVision supports hierarchically nested rectangular ROIs.
- Flexible Masks are recommended to process disconnected ROIs or non-rectangular shapes. They are supported by some `EasyObject` and `EasyImage` library functions.

Flexible Masks

A flexible mask is a BW8 image with the same height and width as the source image. It contains shapes of areas that must be processed and ignored areas (that will not be considered during processing):


- All pixels of the flexible mask having a value of 0 define the ignored areas.
- All pixels of the flexible mask having any other value than 0 define the areas to be processed.



Source image




Associated mask



Processed masked image

A flexible mask can be generated by any application that outputs BW8 images and by some [EasyObject](#) and [EasyImage](#) functions.


Flexible Masks in EasyImage



Source image (left) and mask variable (right)

Simple steps to use flexible masks in Easyimage

1. **Call the functions from EasyImage that take an input mask as an argument.** For instance, one can evaluate the average value of the pixels in the white layer and after in the black layer.
2. **Display the results.**



Resulting image

EasyImage Functions that support flexible masks

- `EImageEncoder::Encode` has a flexible mask argument for BW1, BW8, BW16, and C24 source images.

- [AutoThreshold](#).
- [Histogram](#) (function [HistogramThreshold](#) has no overload with mask argument).
- [RmsNoise](#), [SignalNoiseRatio](#).
- [Overlay](#) (no overload with mask argument for BW8 source images).
- [ProjectOnAColumn](#), [ProjectOnARow](#) (Vector projection).
- [ImageToLineSegment](#), [ImageToPath](#) (Vector profile).

Flexible Masks in EasyObject

A flexible mask can be generated by any application that outputs BW8 images or uses the Open eVision image processing functions.


EasyObject can use flexible masks to restrict blob analysis to complex or disconnected shaped regions of the image.

If an object of interest has the same gray level as other regions of the image, you can define "keep" and "ignore" areas using flexible masks and [Encode](#) functions.

A flexible mask is a BW8 image with the same height and width as the source image.

- A pixel value of 0 in the flexible mask masks the corresponding source image pixel so it doesn't appear in the encoded image.
- Any other pixel value in the flexible mask causes the pixel to be encoded.


EasyObject functions that create flexible masks



Source image

1) [ECodedImage2::RenderMask](#): from a layer of an encoded image


1. To encode and extract a flexible mask, first construct a coded image from the source image.
2. Choose a segmentation method (for the image above the default method `GrayscaleSingleThreshold` is suitable).
3. Select the layer(s) of the coded image that should be encoded (i.e. white and black layers using minimum residue thresholding).
4. Make the mask image the desired size using `mask.SetSize(sourceImage.GetWidth(), sourceImage.GetHeight())`.
5. Exploit the flexible mask as an argument to [ECodedImage2::RenderMask](#).



BW8 resulting image that can be used as a flexible mask

2) `ECodedElement::RenderMask`: from a blob or hole


1. Select the coded elements of interest.
2. Create a loop extracting a mask from selected coded elements of the coded image using `ECodedElement::RenderMask`.
3. Optionally, compute the feature value over each of these selected coded elements.



BW8 resulting image that can be used as a flexible mask


3) `EObjectSelection::RenderMask`: from a selection of blobs

`EObjectSelection::RenderMask` can, for example, discard small objects resulting from noise.




BW8 resulting image that can be used as a flexible mask

Example: Restrict the areas encoded by EasyObject



Find four circles (left) Flexible mask can isolate the central chip (right)

1. Declare a new `ECodedImage2` object.
2. Setup variables: first declare source image and flexible mask, then load them.
3. Declare an `EImageEncoder` object and, if applicable, select the appropriate segmenter. Setup the segmenter and choose the appropriate layer(s) to encode.
4. Encode the source image. Encoding a layer with just the area in the flexible mask is then pretty straightforward.
We see that the circles are correctly segmented in the black layer with the [grayscale single threshold segmenter](#):



5. Select all objects of the coded image.
6. Select objects of interest by filtering out objects that are too small.
7. Display the blob feature by iterating over the selected objects to display the chosen feature.

1.10. Profile

Profile Sampling

A **profile** is a series of pixel values sampled along a line/path/contour in an image.

- `EasyImage::ImageToLineSegment` copies the pixel values along a given line segment (arbitrarily oriented and wholly contained within the image) to a vector. The vector length is adjusted automatically. This function supports flexible mask.
- A **path** is a series of [pixel coordinates](#) stored in a vector.
`EasyImage::ImageToPath` copies the corresponding pixel values to the vector. This function supports flexible mask.

- A **contour** is a closed or not (connected) path, forming the boundary of an object. `EasyImage::Contour` follows the contour of an object, and stores its constituent pixels values inside a profile vector.

Profile Analysis

The profile can be processed to find peaks or transitions:

- A transition corresponds to an object edge (black to white or white to black). It can be detected by taking the first **derivative** of the signal (which transforms transitions (edges) into peaks) and looking for peaks in it.

`EasyImage::ProfileDerivative` computes the first derivative of a profile extracted from a gray-level image.

The `EBW8` data type only handles unsigned values, so the derivative is shifted up by 128.

Values under [above] 128 correspond to negative [positive] derivative (decreasing [increasing] slope).

- A **peak** is the portion of the signal that is above [or below] a given threshold - the maximum or minimum of the signal. This may correspond to the crossing of a white or black line or thin feature. It is defined by its:
 - **Amplitude**: difference between the threshold value and the max [or min] signal value.
 - **Area**: surface between the signal curve and the horizontal line at the given threshold.

`EasyImage::GetProfilePeaks` detects max and min peaks in a gray-level profile. To eliminate false peaks due to noise, two selection criteria are used. The result is stored in a **peaks vector**.

Profile Insertion Into an Image

`EasyImage::LineSegmentToImage` copies the pixel values from a vector or constant to the pixels of a given line segment (arbitrarily oriented and wholly contained within the image).

`EasyImage::PathToImage` copies the pixel values from a vector or a constant to the pixels of a given path.

PART III
IMAGE PRE-PROCESSING LIBRARIES

1. EasyImage - Pre-Processing Grayscale Image

EasyImage operations prepare images so that further processing gets better results by:

- isolating defects using thresholding or intensity transformations
- compensating perspective effects (for non-flat surfaces such as a bottle label)
- processing complex or disconnected shapes using flexible masks

The main functions are:

- **Intensity Transformations** change the gray-level of each pixel to clarify objects (histogram stretching).
- **Thresholding** transforms a binary image into a bi- or tri-level grayscale image by classifying the pixel values.
- **Arithmetic and logic** functions manipulate pixels in two images, or one image and a constant.
- **Non-Linear Filtering** functions use non-linear combinations of neighboring pixels (using a kernel) to highlight a shape, or to remove noise.
- **Geometric transforms** move selected pixels to realign, resize, rotate and warp.
- **Noise Reduction and Estimation** functions ensure that noise is not unacceptably enhanced by other operations (thresholding, high-pass filtering).
- **Gradient Scalar** generates a gradient direction or gradient magnitude map from a gray-level image.
- **Vector operations** extract 1-dimensional data from an image into a vector, for example to detect scratches or outlines, or to clarify images.
- **Harris corner detector** returns a vector of points of interest in a BW8 image.
- **Canny edge detector** returns a BW8 image of the edges found in a BW8 image.
- **Overlay** overlays an image on top of a color image.
- **Operations on Interlaced Video Frames** eliminate interlaced image artifacts by rebuilding or re-aligning fields.
- **Flexible Masks** help process irregular shapes in EasyImage.

1.1. Intensity Transformation


These EasyImage functions change the gray-levels of pixels to increase contrast.

Gain offset

`Gain Offset` changes each pixel to [old gray value * Gain coefficient + Offset].


- **gain** adjusts **contrast**. It should remain close to 1.
- **offset** adjusts **intensity** (brightness). It can be positive or negative.
- The resulting values are always saturated to range [0..255].

In this example, the resulting image has better contrast and is brighter than the source image.



Source and result images (with gain = 1.2 and offset = +12)

Color images have three separate gain and offset values, one per color component (red, green, blue).




Example of gain/offset applied on a color image

Normalization

`Normalize` makes images of the same scene comparable, even with different lighting.

It compares the average gray level (brightness) and standard deviation (contrast) of the source image and a reference image. Then, it normalizes the source image with gain and offset coefficients such that the output image has the same brightness and contrast as the reference image. This operation assumes that the camera response is reasonably linear and the image does not saturate.




The reference image (from which the average and standard deviation are computed), the source image (too bright), and the normalized image (contrast and brightness are the same as the reference image)

Uniformization

`Uniformize` compensates for non-uniform illumination and/or camera sensitivity based on one or two reference images. The reference image should not contain saturated pixel values and have minimum noise.


- When one reference image is used, the transformation is similar to an adaptive (space-variant) gain; each pixel in the reference image encodes the gain for the corresponding pixel in the source image.
- When two reference images are used, the transformation is similar to an adaptive gain and offset; each pixel in the reference images encodes either the gain or the offset for the corresponding pixel in the source image.



Example of an image uniformized with two reference images


Lookup tables

`Lut` uses a lookup table of new pixel values to replace the current ones - efficient for BW8 and BW16 images. If the transform function never changes, it is best to use a lookup table.



Example of a transform

1.2. Thresholding



Thresholding transforms an image by classifying the pixel values using these methods:

- "Automatic thresholding" on the next page (BW8 and BW16 images only)
- "AutoThreshold" on page 43 (BW8 and BW16 images only)
- "Manual thresholding" on the next page using one or two threshold values
- "Histogram based" on the next page (computed before using the thresholding function)

These functions also return the average gray levels of each pixel below and above the threshold.

Keys to successful thresholding

- Object and background areas should be of uniform color and illumination. Image uniformization may be required prior to thresholding.
- The gray level range of the object and background must be sufficiently different (all background pixels should be darker than the darkest object pixel).

- You must decide if the threshold value should be:
 - constant: **absolute** threshold
 - adapted to ambient light intensity: **relative** or **automatic** threshold


Automatic thresholding

The threshold is calculated automatically if you use one of these arguments with the `EasyImage::Threshold` function.

Min Residue: Minimizes the quadratic difference between the source and resulting image (default if Threshold function is invoked without a argument).



Max Entropy: Maximizes the entropy (i.e. the amount of information) between object and background of the resulting image.



Isodata: Calculates a threshold value that is an average of the gray levels: halfway between the average gray level of pixels below the threshold, and the average gray level of pixels above the threshold.

Manual thresholding

Manual thresholds require that the user supplies one or two threshold values:

- **one** value to the `Threshold` function to classify source image pixels (BW8/BW16/C24) into two classes and create a bi-level image. This can be:
 - `relativeThreshold` is the percentage of pixels below the threshold. The Threshold function then computes the appropriate threshold value, or
 - `absoluteThreshold`. This value must be within the range of pixel values in the source image.
- **two** values to the `DoubleThreshold` function to classify source image pixels (BW8/BW16) into three classes and create a tri-level image.
 - `LowThreshold` is the lower limit of the threshold
 - `HighThreshold` is the upper limit of the threshold

Histogram based

When a histogram of the source image is available, you can speed up the automatic thresholding operation by computing the threshold value from the histogram (using `HistogramThreshold` or `HistogramThresholdBW16`) and using that value in a manual thresholding operation.

These functions also return the average gray levels of each pixel below and above the threshold.

AutoThreshold

When no source image histogram is available, `AutoThreshold` can still calculate a threshold value using these [threshold modes](#): `EThresholdMode_Relative`, `_MinResidue`, `_MaxEntropy` and `_Isodata`.

This function supports flexible masks.

1.3. Arithmetic and Logic

Reasons you may use arithmetic and logic are:

- **to emphasize differences** between images by subtracting the pixels (a conformity check).
- **to compensate for non-uniform lighting** by dividing the target image by the image of the background alone.
- **to remove unwanted areas of an image** by preparing an appropriate mask, and clearing all the pixels that belong to the mask by using logical combinations of pixels.
- **to create a combined image** by combining the pixels of two source images to generate a resulting image.

Arithmetic operations are handled by the `Oper` function, `EArithmeticLogicOperation` enum lists all supported operators.

These operations can be applied to images and constants, they have one or two source arguments (image or integer constants) and one destination argument. If the source operands are a color and a gray-level image, each color component combines with the gray-level component to give a color image. [Histogram equalization](#) can improve your results.

arithmetic and logic combinations

Allowed combinations

	General	Copy	Invert	Shift	Logical	Overlay	Set
Const BW8 -> Image BW8		x					
Const C24 -> Image C24		x					
Image BW8 -> Image BW8		x	x				
Image BW8 -> Image C24		x	x			x	
Image C24 -> Image C24		x	x				
Const BW8, Image BW8 -> Image BW8	x						
Image BW8, Const BW8 -> Image BW8	x			x			x

	General	Copy	Invert	Shift	Logical	Overlay	Set
Image BW8, Image BW8 -> Image BW8	x				x		x
Image BW8, Image BW8 -> Image C24	x					x	
Const C24, Image C24 -> Image C24	x						
Image C24, Const C24 -> Image C24	x			x			
Image C24, Image C24 -> Image C24	x				x		
Image C24, Image BW8 -> Image C24	x				x	x	
Image BW8, Image C24 -> Image C24	x				x		x

Note: For logical operators, a pixel with value 0 is assumed **FALSE**, otherwise **TRUE**. The result of a logical operation is 0 when **FALSE** and 255 otherwise.

The classification of operations in the above table are:

General

- Compare (abs. value of the difference)
- Saturated sum
- Saturated difference
- Saturated product
- Saturated quotient
- Modulo
- Overflow-free sum
- Overflow-free difference
- Overflow-free product
- Overflow-free quotient
- Bitwise AND
- Bitwise OR
- Bitwise XOR
- Minimum
- Maximum
- Equal
- Not equal
- Greater or equal
- Lesser or equal

- Greater
- Lesser

Copy

- Sheer Copy

Invert

- Invert (negative)

Shift

- Left Shift
- Right Shift

Logical

- Logical AND
- Logical OR
- Logical XOR

Overlay

- Add an overlay

Set

Operators Copy if mask = 0 and Copy if mask \neq 0 are very handy to perform masking: the first image argument serves as a mask that allows or disallows changing the pixel values in the destination image.

- Copy if mask = 0
- Copy if mask \neq 0

1.4. Non-Linear Filtering

These functions use non-linear combinations of neighboring pixels to highlight a shape, or to remove noise.


Most can be destructive (except top-hat and median filters) i.e. the source image is overwritten by the destination image. Destructive operations are faster.

All have a gray image and a bilevel equivalent, for example [ErodeBox](#) and [BiLevelErodeBox](#).

1. They define the required shape by a "Kernel" on the next page (usually in a 3x3 matrix).
2. They slide this Kernel over the image to determine the value of the destination pixel when a match is found:

- **Erosion, Dilation:** shrinks / grows image regions.
- **Opening, Closing:** removes / fills image region boundary pixels.
- **Thinning, Thickening:** erodes / dilates using image pattern matching.
- **Top-Hat filters:** retains all the tiny image details while removing everything else.
- **Morphological distance:** indicates how many erosions are required to make a pixel black.
- **Morphological gradient:** indicates the outer and inner edges of the erosion and dilation processes.
- **Median filter:** removes impulsive noise.
- **Hit-and-Miss transform:** detects patterns of foreground /background pixels, can create skeletons.

Kernel



Rectangular kernel of half width = 3 and half height = 2 (left) Circular kernel of half width = 2 (right)

The morphological operators combine the pixel values in a neighborhood of given shape (square, rectangular or circular) and replace the central pixel of the neighborhood by the result. The combining function is non-linear, and in most cases is a rank filter: which considers the N values in the given neighborhood, sorts them increasingly and selects the K-th largest. Three special cases are most often used **erosion**, **dilation** and **median filter** where : K can be 1 (minimum of the set), N (maximum) or N/2 (median).


Erosion, Dilation, Opening, Closing, Top-Hat and **Morphological Gradient** operations all use rectangular or circular kernels of odd size. Kernel size has an important impact on the result.

examples

HalfWidth/HalfHeight	Actual width/height
0	1
1	3
2	5
3	7

Erosion, Dilation

Erosion reduces white objects and enlarges black objects, Dilation does the opposite.




Erosion thins white objects by removing a layer of pixels along the objects edges: [ErodeBox](#), [ErodeDisk](#). As the kernel size increases, white objects disappear and black ones get fatter.

Dilation thickens white objects by adding a layer of pixels along the objects edges: [DilateBox](#), [DilateDisk](#). As the kernel size increases, white objects get fatter and black ones disappear.

Opening, Closing

Opening removes tiny white objects / dust. Closing removes tiny black holes / dust.



An **Opening** is an erosion followed by a dilation using [OpenBox](#), [OpenDisk](#).

The global effect is to preserve the overall shape of objects, while removing white details that are smaller than the kernel size.

A **Closing** is a dilation followed by an erosion using [CloseBox](#), [CloseDisk](#).

The global effect is to preserve the overall shape of objects, while removing the black details that are smaller than the kernel size.

Thinning, Thickening

These functions use a 3x3 kernel to grow (**Thick**) or remove (**Thin**) pixels:

- Thinning: can help edge detectors by reducing lines to single pixel thickness.
- Thickening: can help determine approximate shape, or skeleton.

When a match is found between the kernel coefficients and the neighborhood of a pixel, the pixel value is set to 255 if thickening, or 0 if thinning. The kernel coefficients are:

- 0: matching black pixel, value 0
- 1: matching non black pixel, value > 0
- -1: don't care

Top-Hat filters

Top-hat filters are excellent for improving non-uniform illumination.



White top-hat filter: source and destination images

They take the difference between an image and its opening (or closure). Thus, they keep the features that an opening (or closing) would erase. The result is a perfectly flat background where only black or white features smaller than the kernel size appear.

- **White top-hat filter** enhances thin white features: [WhiteTopHatBox](#) ,[WhiteTopHatDisk](#).
- **Black top-hat filter** enhances thin black features:[BlackTopHatBox](#)[BlackTopHatDisk](#).

Morphological distance


[Distance](#) computes the morphological distance (number of erosion passes to set a pixel to black) of a binary image (0 for black, non 0 for white) and creates a destination image, where each pixel contains the morphological distance of the corresponding pixel in the source image.

Morphological gradient

The morphological gradient performs edge detection - it removes everything in the image but the edges.

The morphological gradient is the difference between the dilation and the erosion of the image, using the same structuring element.


[MorphoGradientBox](#), [MorphoGradientDisk](#).



Dilation - Erosion = Gradient

Median


The [Median](#) filter removes impulse noise, whilst preserving edges and image sharpness. It replaces every pixel by the median (central value) of its neighbors in a 3x3 square kernel, thus, outer pixels are discarded.



Median filter: source and destination images

Hit-and-Miss transform

Hit-and-miss transform operates on BW8, BW16 or C24 images or ROIs to detect a particular pattern of foreground and background pixels in an image.



Hit-and-miss transform

The `HitAndMiss` function has three arguments:

- A pointer to the source image of type `EROIBW8`, `EROIBW16`, or `EROIC24`
- A pointer to the destination image of type corresponding to the type of the source image. The sizes of the source and destination images must be identical.
- A kernel of type `EHitAndMissKernel` Two constructors are available for the kernel object:
 - `EHitAndMissKernel(int startX, int startY, int endX, int endY)` where:
 - `startX, startY` are coordinates of the top left of the kernel, must be less than or equal to zero.
 - `endX, endY` are coordinates of the bottom right of the kernel, must be greater than or equal to zero.
 The constructed kernel has no explicit restrictions on its size, and the following characteristics:


$$\text{kernel width} = (\text{endX} - \text{startX} + 1), \text{kernel height} = (\text{endY} - \text{startY} + 1)$$
 - `EHitAndMissKernel(unsigned int halfSizeX, unsigned int halfSizeY)` where:
 - `halfSizeX` is half of the kernel width - 1, must be greater than zero.
 - `halfSizeY` is half of the kernel height - 1, must be greater than zero.
 The constructed kernel has the following characteristics:

$$\text{kernel width} = ((2 \times \text{halfSizeX}) + 1), \text{kernel height} = ((2 \times \text{halfSizeY}) + 1)$$

$$\text{kernel StartX} = - \text{halfSizeX}, \text{kernel StartY} = - \text{halfSizeY}$$

example: detecting corners in a binary image.

The `hit-and-miss transform` can be used to locate corners.



Binary source image


1. Define the kernel by detecting the left corner. The left corner pixel has black pixels on its immediate left, top and bottom; and it has white pixels on its right. The following hit-and-miss kernel will detect the left corner:

```

- +
- + +
- +

```

2. Apply the filter on the source image. Note that the resulting image should be properly sized.



Resulting image, highlighted pixel is located on left corner of rhombus

3. Locate the three remaining corners in the same way: Declare three kernels that are the rotation of the filter above and apply them.

4. Detect the right, top and bottom corners.



1.5. Geometric Transforms

Geometric transformation moves selected pixels in an image, which is useful if a shape in an image is too large / small / distorted, to make point-to-point comparisons possible.

The selected area may be any shape, but the resulting image is always rectangular. Pixels in the destination image that have corresponding pixels that are outside of the selected area are considered not relevant and are left black.

When the source coordinates for a destination pixel are not integer, an interpolation technique is required.

The nearest neighborhood method is the quickest - it uses the closest source pixel.

The bi-linear interpolation method is more accurate but slower - it uses a weighted average of the four neighboring source pixels.

Possible geometric transformations are:

ReAlignment

The simplest way to realign two misaligned images is to accurately locate features in both images (landmarks or pivots, using pattern matching, point measurement or other) and realign one of the images so that these features are superimposed.

You can [register](#) an image by realigning one, two or three pivot points to reference positions. For best accuracy, the pivot points should be as far apart as possible.

- A **single pivot point** transform is a simple translation. If interpolation bits are used, sub-pixel translation is achieved.
- **Two pivot points** use a combination of translation, rotation and optionally scaling. If scaling is not allowed, the second pivot point may not be matched exactly. Scaling should not normally be used unless it corresponds to a change of lens magnification or viewing distance.
- **Three pivot points** use a combination of translation, rotation, shear correction and optionally scaling. A shear effect can arise when acquiring images with a misaligned line-scan camera.

Mirroring

This destructive feature modifies the source image to create a mirror image:

- [horizontally](#) (the columns are swapped) or
- [vertically](#) (the rows are swapped).


Translation, Scaling and Rotation

If the position or size of an object of interest changes, you can measure the change in position or size and generate a corrected image using the ScaleRotate and Shrink functions.

`EasyImage::ScaleRotate` performs:

- Image translation: you provide the position coordinates of a pivot-point in the source image and a corresponding pivot point in the destination image.
- Image scaling: you provide scaling factor values for X- and Y-axis.
- Image rotation: you provide a rotation angle value.

For resampling, the nearest neighbor rule or bilinear interpolation with 4 or 8 bits of accuracy is used. The size of the destination image is arbitrary.



Scale and rotate example

Shrink

`EasyImage::Shrink`: resizes an image to be smaller, applying pre-filtering to avoid aliasing.

LUT-based unwarping

If the feature of interest is distorted due to its shape (anamorphosized), you can unwarp a circular ring-wedge shape (such as text on CD labels) into a straight rectangle. A ring-wedge is delimited by two concentric circles and two straight lines passing through the center.

`EasyImage::SetCircleWarp` prepares warp images for use with function `EasyImage::Warp` which moves each pixel to locations specified in the "warp" images which are used as lookup tables.

1.6. Noise Reduction and Estimation

Noise can degrade the visual quality of images, and certain processing operations (thresholding, high-pass filtering) will enhance noise in a non-acceptable way. Acquired images are always noisy (this is best observed on live images where the pixel values fluctuate around the true intensity). When acquired with 8 bits of accuracy, the noise level typically amounts to about 3 to 5 gray-level values. One distinguishes several forms of noise:

- **additive**: noise amplitude is not related to image contents
- **multiplicative**: noise amplitude is proportional to local intensity
- **uniform**: noise amplitude follows a smooth distribution centered around zero
- **impulse**: noise amplitude is infinite.

Impulse noise produces a "salt and pepper" effect, while uniform noise blends.

Spatial noise reduction (if you only have 1 image):


Reduces uniform and impulse noise but blurs edges.

Cannot distinguish noise from actual signal changes, so always spoils part of the signal.

Uses the correlation between neighboring pixel values to perform convolution or median


filtering:

- **Convolution** replaces the value at each pixel by a combination of its neighbors, leading to a localized averaging. Linear filtering is recommended to reduce uniform noise. Beware that it tends to blur edges.



Uniform noise reduction by low-pass filtering

- **Median filtering** replaces each pixel by the median value in the pixel neighborhood (5-th largest value in a 3x3 neighborhood). This reduces impulse noise and keeps sharpness.



Impulse noise reduction by median filtering


- `EasyImage::Median`
- `EasyImage::BiLevelMedian`

Temporal noise reduction (for several images, e.g. moving objects):

Temporal noise reduction is achieved by combining the successive values of individual pixels across time. EasyImage implements recursive averaging and moving averaging.


EasyImage provides three ways to minimize noise by means of several images:

- **Temporal average:** just accumulates N images and average them; using standard arithmetic operations, as illustrated below. Creates de-noised image after N acquisitions using average values. Noise varies from frame to frame while the signal remains unchanged, so if several images of the same (still) scene are available, noise can be separated from the signal. The disadvantage of producing one de-noised image after N acquisitions only, is that fast display refresh is not possible.




Simple average

- **Temporal moving average:** accumulates the last N images and updates the de-noised image each time a new one is acquired, in such a way that the computation time does not depend on N. The whole process is handled by `EMovingAverage`. The disadvantage of this method is that it combines noisy images together.



Moving average

- **Temporal recursive average:** combines a noisy image with the previously de-noised image using `EasyImage::RecursiveAverage`.



Recursive average

Recursive averaging

This is a well known process for noise reduction by temporal integration. The principle is to continuously update a noise-free image by blending it, using a linear combination, with the raw, noisy, live image stream. Algorithmically, this amounts to the following:

$$DST_N = a * Src + (1-a) * Dst_{N-1}$$

where a is a mixture coefficient. The value of this coefficient can be adjusted so that a prescribed noise reduction ratio is achieved.

This procedure is effective when applied to still images, but generates a trailing effect on moving objects. The larger the noise reduction ratio, the heavier the trailing effect is. To work around this, a non-linearity can be introduced in the blending process: small gray-level value variations between successive images are usually noise, while large variations correspond to changes in the image.

`EasyImage::RecursiveAverage` uses this observation and applies stronger noise reduction to small variations and conversely. This reduces noise in still areas and trailing in moving areas.

For optimal performance, the non-linearity must be pre-computed once for all using function `EasyImage::SetRecursiveAverageLUT`.

Note: Before the first call to the `EasyImage::RecursiveAverage` method, the 16-bit work image must be cleared (all pixel values set to zero).

Noise estimation (of image compared to reference image):

To estimate the amount of noise, two or more successive images are required. In the simplest

mode, two noisy images are compared. (Other modes are available: if a noise-free image is available, it is compared to a noisy one; a noise-free image can also be built by temporal averaging.) Calculates the root-mean-square amplitude and signal-to-noise ratio.

- `EasyImage::RmsNoise` computes the root-mean-square amplitude of noise, by comparing a given image to a reference image. This function supports flexible mask and an input mask argument. BW8, BW16 and C24 source images are supported.
The reference image can be noiseless (obtained by suppressing the source of noise), or affected by a noise of the same distribution as the given image.
- `EasyImage::SignalNoiseRatio` computes the signal to noise ratio, in dB, by comparing a given image to a reference image. This function supports flexible mask and an input mask argument. BW8, BW16 and C24 source images are supported.
The reference image can be noiseless (obtained by suppressing the source of noise) or be affected by a noise of the same distribution as the given image.

Signal amplitude is the sum of the squared pixel gray-level values.

Noise amplitude is the sum of the squared difference between the pixel gray-level values of the given image and the reference.

1.7. Scalar Gradient

`EasyImage::GradientScalar` computes the (scalar) gradient image derived from a given source image.

The scalar value derived from the gradient depends on the preset lookup-table image.

The gradient of a grayscale image corresponds to a vector, the components of which are the partial derivatives of the gray-level signal in the horizontal and vertical direction. A vector can be characterized by a direction and a length, corresponding to the gradient orientation, and the gradient magnitude.

This function generates a gradient direction or gradient magnitude map (gray-level image) from a given gray-level image.

For efficiency, a pre-computed lookup-table is used to define the desired transformation.


This lookup-table is stored as a standard `EImageBW8/EImageBW16`.

Use `EasyImage::ArgumentImage` or `EasyImage::ModulusImage` once before calling `GradientScalar`.

1.8. Vector Operations

Extracting 1-dimensional data from an image generates linear sets of data that are handled as vectors. Subsequent operations are fast because of the reduced amount of data. The methods are either:


Projection



Projects the sum or average of all gray color-level values in a given direction, into various vector types (levels are added when projecting into an `EBW32Vector` and averaged when projecting into an `EBW8Vector`, `EBW16Vector` or `EC24Vector`). These functions support **flexible mask**.

- `EasyImage::ProjectOnAColumn` projects an image horizontally onto a column.
- `EasyImage::ProjectOnARow` projects an image vertically onto a row.

Profile



Samples a series of pixel values along a given segment, path or contour, then analyze and modify their Peaks and Transitions to make images clearer:

1. Obtain the profile of a line segment / path / contour.

`EasyImage::ImageToLineSegment` copies the pixel values along a given line segment (arbitrarily oriented) to a vector. The line segment must be entirely contained within the image. The vector length is adjusted automatically. This function supports flexible mask.

`EasyImage::ImageToPath` copies the corresponding pixel values to the vector. The function supports flexible mask. A **path** is a series of **pixel coordinates** stored in a vector.

`EasyImage::Contour` follows the contour of an object, and stores its constituent pixels values inside a profile vector. A **contour** is a closed or not (connected) path, forming the boundary of an object.

2. Analyse the profile to find peaks or transitions.

`EasyImage::GetProfilePeaks` detects max and min peaks in a gray-level profile. To eliminate false peaks due to noise, two selection criteria are used. The result is stored in a **peaks vector**.

A **peak** is a maximum or minimum of the signal which may correspond to the crossing of a white or black line or thin feature. It is defined by its:

- **Amplitude**: difference between the threshold value and the max [or min] signal value.
- **Area**: surface between the signal curve and the horizontal line at the given threshold.

A **transition** corresponds to an object edge (black to white, or white to black). It can be detected by taking the first **derivative** of the signal and looking for peaks in it.

`EasyImage::ProfileDerivative` computes the first derivative of a profile extracted from a gray-level image. This derivative transforms transitions (edges) into peaks.

`EBW8` data type only handles unsigned values, so the derivative is shifted up by 128. Values under 128 correspond to negative derivative (decreasing slope), values above 128 correspond to positive derivative (increasing slope).

3. Insert the profile into an image.


`EasyImage::LineSegmentToImage` copies the pixel values from a vector or a constant to the pixels of a given line segment (arbitrarily oriented). The line segment must be wholly contained within the image.

`EasyImage::PathToImage` copies pixel values from a vector or a constant to the pixels of a given path.

1.9. Canny Edge Detector


The Canny edge detector facilitates:

- Good detection: finds all edges
- Good localization: the found edges are as close as possible to the "real" edges in the image
- Minimal response: one edge response is accepted for each position, i.e. avoiding multiple close or intersecting edge responses



Source image and the result after a Canny edge detection

The EasyImage Canny edge detector operates on a grayscale BW8 image and delivers a black-and-white BW8 image where pixels have only 2 possible values: **0** and **255**. Pixels corresponding to edges in the source image are set to **255**; all others are set to **0**. It can adjust the scale analysis, it doesn't allow sub-pixel interpolation and it delivers a binary image after thresholding.



Canny edge detector example

The Canny edge detector requires only two parameters:

- **Characteristic scale of the features of interest:** the standard deviation of the Gaussian filter used to smooth the source image.
- **Gradient threshold with hysteresis:** maximum magnitude of the gradient of the source image expressed as a fraction ranging from 0 to 1 (two values).


The API of the Canny edge detector is a single class, [ECannyEdgeDetector](#), with the following methods:

- [Apply](#): applies the Canny edge detector on an image/ROI.
- [GetHighThreshold](#): returns the high hysteresis threshold for a pixel to be considered as an edge.
- [GetLowThreshold](#): returns the low hysteresis threshold for a pixel to be considered as an edge.
- [GetSmoothingScale](#): returns the scale of the features of interest.
- [GetThresholdingMode](#): returns the mode of the hysteresis thresholding.
- [ResetSmoothingScale](#): prevents the smoothing of the source image by a Gaussian filter.
- [SetHighThreshold](#): sets the high hysteresis threshold for a pixel to be considered as an edge.
- [SetLowThreshold](#): sets the low hysteresis threshold for a pixel to be considered as an edge.
- [SetSmoothingScale](#): sets the scale of the features of interest.
- [SetThresholdingMode](#): sets the mode of the hysteresis thresholding.

The **result image** must have the same dimensions as the input image.

1.10. Harris Corner Detector

The Harris corner detector is invariant to rotation, illumination variation and image noise. It operates on a grayscale BW8 image and delivers a vector of points of interest.



Harris corner detector example

The EasyImage Harris corner detector requires three parameters:

- The integration scale σ_i : the standard deviation of the Gaussian Filter used for scale analysis. $\sigma_d = 0,7 \times \sigma_i$, where σ_d is the differentiation scale: the standard deviation of the Gaussian Filter used for noise reduction during computation of the gradient.
- A corner threshold: a fraction ranging from 0 to 1 of the maximum value of the cornerness of the source image.
- A Boolean that toggles sub-pixel detection.

The following characteristics are available for every point of interest:

- Corner position (pixel coordinates with sub-pixel accuracy if enabled).
- Cornerness measurement.
- Gradient magnitude with regards to the differentiation scale σ_d .
- Gradient value along the X-axis with regards to the differentiation scale σ_d .
- Gradient value along the Y-axis with regards to the differentiation scale σ_d .

The API of the Harris corner detector is a single class named `EHarrisCornerDetector` and these methods:

- `Apply`: applies the Harris corner detector on an image/ROI.
- `EHarrisCornerDetector`: constructs a `EHarrisCornerDetector` object initialized to its default values.
- `GetDerivationScale`: returns the current derivation scale.
- `GetScale`: returns the integration scale.
- `GetThreshold`: returns the current threshold.
- `GetThresholdingMode`: returns the current thresholding mode for the cornerness measure.

- `IsGradientNormalizationEnabled`: returns whether the gradient is normalized before the computation of the cornerness measure.
- `IsSubpixelPrecisionEnabled`: returns whether the sub-pixel interpolation is enabled.
- `SetDerivationScale`: sets the derivation scale.
- `SetGradientNormalizationEnabled`: sets whether the gradient is normalized before the computation of the cornerness measure.
- `SetScale`: sets the integration scale.
- `SetSubpixelPrecisionEnabled`: sets whether the sub-pixel interpolation is enabled.
- `SetThreshold`: sets the threshold on the cornerness measure for a pixel to be considered as a corner.
- `SetThresholdingMode`: sets the thresholding mode for the cornerness measure.

Basic usage of Harris Corner Detector

An object of the `EHarrisCornerDetector` class can be reused across Harris detector applications, in order to reduce the setup time.

1. **Create an instance of the detector** and set the appropriate method, for instance, the integration scale, `SetScale`, with the structures of interest that could have a spatial extent of 2 pixels.
2. **Apply the detector** with two arguments to the new image : the input image and the interest points in the input image `EHarrisInterestPoints`.
3. Access the individual elements of the output vector.

1.11. Overlay

`EasyImage::Overlay` overlays an image on the top of a color image, at a given position.

If a color image is provided as the source image, all the pixels of this image are copied to the destination image, except the ones that equal the reference color. When a C24 image is used as overlay source image, the color of the overlay in destination image is the same as the one in the overlay source image, thus allowing multicolored overlays.

If a BW8 image is provided as the source image, all the overlay image pixels are copied to the destination image, apart from those that are the reference color which are replaced by the source images.

This function supports flexible mask and an input mask argument. C24, C15 and C16 source images are supported.

1.12. Operations on Interlaced Video Frames

When an image is interlaced, the two frames (even and odd lines) are not recorded at the same time. If there is movement in the scene, a visible artifact can result (the edges of objects exhibit a "comb" effect).

`EasyImage::RealignFrame` cures this problem if the movement is uniform and horizontal (objects on a conveyor belt), by shifting one of the frames horizontally. The amplitude of the shift can be estimated automatically.

`EasyImage::GetFrame` extracts the frame of given parity from an image while `EasyImage::SetFrame` replaces the frame of given parity in an image.

`EasyImage::MatchFrames` determines the optimal shift amplitude by comparing two successive lines of the image. These lines should be chosen such that they cross some edges or non-uniform areas.

`EasyImage::RebuildFrame` rebuilds one frame of the image by interpolation between the lines of the other frame.


`EasyImage::SwapFrames`: interchanges the even and odd rows of an image. This is helpful when acquisition of an interlaced image has confused even and odd frames.

The same image should be used as source and destination because only the shifted rows are copied. To use a different destination image, the source image must be copied first in the destination image object.

The size of the destination image is determined as follows:

```
dstImage_Width = srcImage_Width
dstImage_Height = (srcImage_Height + 1 - odd) / 2
```


1.13. Flexible Masks in EasyImage



Source image (left) and mask variable (right)

Simple steps to use flexible masks in Easyimage

1. **Call the functions from EasyImage that take an input mask as an argument.** For instance, one can evaluate the average value of the pixels in the white layer and after in the black layer.
2. **Display the results.**



Resulting image

EasyImage Functions that support flexible masks

- `EImageEncoder::Encode` has a flexible mask argument for BW1, BW8, BW16, and C24 source images.
- `AutoThreshold`.
- `Histogram` (function `HistogramThreshold` has no overload with mask argument).
- `RmsNoise`, `SignalNoiseRatio`.
- `Overlay` (no overload with mask argument for BW8 source images).
- `ProjectOnAColumn`, `ProjectOnARow` (Vector projection).
- `ImageToLineSegment`, `ImageToPath` (Vector profile).

2. EasyColor - Pre-processing Color Images


EasyColor makes color image processing as efficient as possible by detecting, classifying and analyzing objects. Several conversion functions mean that any color system can be processed.

Color definition and supported systems

What Is Color?

The human eye is sensitive to light:

- **Intensity**, or **achromatic** sensation, captured by **grayscale** images.
- **Wavelength**, or **chromatic** sensation, described in red, green and blue primary **colors**. True color digital images (24 bits per pixel; 8 bits per RGB channel) represent as many colors as the eye can distinguish.



Visible color gamut in the XYZ color space

There are three color systems:

- **Mixture** systems (RGB/XYZ) give the proportions of the three primaries to be combined.
- **YUV Luma/chroma** systems (XYZ/YUV) separate the achromatic (Y) and chromatic sensations (U & V). Used when a black and white image is required as well (television).
- **Intensity/saturation/hue** systems (RGB/XYZ/YUV) separate achromatic (black and white Intensity) from enhanced chromatic (color Saturation and Hue) sensations. Used to eliminate lighting effects, or to convert RGB images to another color system. More saturated colors are more vivid, less saturated ones are grayer.

In general:

- RGB is used by monitors, cameras and other display devices.
- YUV is used for efficient transmission of color images by compressing the chrominance information.
- XYZ is used for device-independent color representation.

All image processing operations can use **quantized** coordinates: discrete values in the [0..255] interval, which use a byte representation to store images in a frame buffer.

Color system conversion operations can also use simpler **unquantized** coordinates: continuous values, often normalized to the [0..1] interval.

Color Image Processing

A color image is a vector field with three components per pixel. All three RGB components reflected by an object have amplitude proportional to the intensity of the light source. By considering the ratio of two color components, one obtains an illumination-independent image. With a clever combination of three pieces of information per pixel, one can extract better features.

There are 3 ways to process a color image:

- **Component extraction:** you can extract the most relevant feature from the triple color information, to reduce the amount of data. For instance, objects may be distinguished by their hue, a pre-processing step could transform the image to a gray-level image containing only hue values.
- **De-coupled transformation:** you can perform operations separately on each color component. For instance, adding two images together adds the red, green and blue components and stores the result, component by component, in a resulting color image.
- **Coupled transformation:** you can combine all three color components to produce three derived components. For example, converting YIQ to RGB.

Supported color systems

EasyColor supports color systems RGB, XYZ, $L^*a^*b^*$, $L^*u^*v^*$, YUV, YIQ, LCH, ISH/LSH, VSH and YSH.

RGB is the preferred internal representation as it is compatible with 24-bit Windows Bitmaps.

	RGB-based	XYZ-based	YUV-based
Mixture	RGB	XYZ	—
Luma/Chroma	—	$L^*a^*b^*$ $L^*u^*v^*$	YUV YIQ
Intensity/Saturation/Hue	ISH LSH VSH	LCH	YSH

Transform using LUTs (LookUp Tables)

EasyColors Lookup tables provide an array of values that define what output corresponds to a given input, so an image can be changed by a user-defined transformation.

A color pixel can take 16,777,216 (2^{24}) values, a full color LUT with these entries would occupy 50 MB of memory and transforms would be prohibitively time-consuming. Pre-computed LUTs make color transforms feasible.

To transform a color image, you initialize a color LUT using one of the following functions:

"LUT for Gain/Offset (Color)" on page 69: `EasyImage::GainOffset`,

"LUT for Color Calibration" on page 70: `Calibrate`,

"LUT for Color Balance" on page 70: `WhiteBalance`,

`ConvertFromRGB`, `ConvertToRGB`.

This color LUT is then used in a transform operation such as `EasyColor::Transform` or you can create a custom transform using `EColorLookup` which takes unquantized values (continuous, normalized to [0..1] intervals), and specifies the source and destination color systems. Some operations use the LUT on-the-fly thus avoid storing the transformed image, for example to alter the U (of YUV) component while the image is in RGB format.

The optimum combination of **accuracy and speed** is determined by the choice of `IndexBits` and `Interpolation` - the accuracy of the transformed values roughly corresponds to the number of index bits.

- Fewer table entries mean smaller storage requirements, but less accuracy.
- No interpolation gives quicker running time, but less accuracy. Interpolation can recover 8 bits of accuracy per component. When the involved transform is linear (such as YUV to RGB), interpolation always gives exact results, regardless of the number of table entries.

Index Bits	Number of entries	Table size (bytes)
4	$2^{(3*4)} = 4,096$	14,739
5	$2^{(3*5)} = 32,768$	107,811
6	$2^{(3*6)} = 262,144$	823,875


Discrete Quantized vs. Continuous Unquantized

Color coordinates in the classical systems are normally **continuous** values, often normalized to the **[0..1]** interval. Computations on such values, termed **unquantized**, are simpler.

However, storage of images in a frame buffer imposes a byte representation, corresponding to **discrete** values, in the **[0..255]** interval. Such values are termed **quantized**.

All image processing operations apply to quantized values, but conversion operations can also be specified using unquantized coordinates.

Bayer transform



Bayer pattern encoded image

A Bayer encoded image is not compatible with a true color image (EC24), but white balance and gamma correction can be applied to it using `EColorLookup` parameter in `EasyColor::BayerToC24`.

A Bayer image is three times smaller, so processes much faster.

Easyobject can use the Bayer pattern to create a color image.

Transform YUV444 / YUV422

YUV images can be minimized without degrading visual quality using function `Format444To422` to convert from 4:4:4 to 4:2:2 format (or you can convert `Format 422 To 444`).

- 4:4:4 uses 3 bytes of information per pixel.
- 4:2:2 uses 2 bytes of information per pixel.

It stores the **even** pixels of U and V chroma with the **even and odd** pixels of Y luma as follows:

$$Y_{[even]} \quad U_{[even]} \quad Y_{[odd]} \quad V_{[even]}$$

Merge, extract and color

A color image contains three color planes of continuous tone images.

A gray-level image can be a component of a color system.

Merge and extract components

EasyColor can change or extract one plane at a time, or all three together. See `Compose`, `Decompose`, `GetComponent`, `SetComponent`.


These operations can use a color LUT to transform on the fly, they could build an RGB image from lightness, saturation and hue planes.

Note: *EasyColor functions perform the necessary interleaving / un-interleaving operations to support Windows bitmap format of interleaved color planes (blue, green and red pixels follow each other).*

Pseudo-color to transform gray-level images to color

The trick is to define a regular gamut of 256 colors and each color will be assigned to pixels with a corresponding gray-level value.


To define pseudo-color shades, you specify a trajectory in the color space of an arbitrary system. You can then pseudo-color using the drawing functions color palette (see Image and Vector Drawing) then save and/or transform it like any other color image.



Gray-level and pseudo-colored image

Separate color objects

This EasyColor process takes a set of distinct colors and associates each pixel with the closest color, using a layer index that can then be used in [EasyObject](#) with the labeled image segmenter to improve blob creation.



Raw image and segmented image (3 colors)

2.1. Bayer Transform

The Bayer pattern is a color image encoding format for capturing color information from a single sensor.

A color filter with a specific layout is placed in front of the sensor so that some of the pixels receive red light only, while others receive green or blue only.

An image encoded by the Bayer pattern has the same format as a gray-level image and conveys three times less information. The true horizontal and vertical resolutions are smaller than those of a true color image.


G	B	G	B	G	B	RGB	RGB	RGB	RGB	RGB	RGB
R	G	R	G	R	G	RGB	RGB	RGB	RGB	RGB	RGB
G	B	G	B	G	B	RGB	RGB	RGB	RGB	RGB	RGB
R	G	R	G	R	G	RGB	RGB	RGB	RGB	RGB	RGB
G	B	G	B	G	B	RGB	RGB	RGB	RGB	RGB	RGB
R	G	R	G	R	G	RGB	RGB	RGB	RGB	RGB	RGB

Bayer vs. true color format

Note: The Bayer pattern normally starts with a GB/RG block in the upper left corner. If the image is cropped, this parity rule can be lost, but parity adjustment is unnecessary when working on a Open eVision ROI.

The Bayer conversion method `EasyColor::BayerToC24` transforms an image captured using the Bayer pattern and stored as a gray-level image, into a true color image. There are three ways to reconstruct the missing pixels. The more complex the interpolation, the slower the conversion. However, it is highly recommended to use interpolation.

- **Non-interpolated mode:** duplicates the nearest pixel to above and/or to the left of the current pixel.
- **Standard interpolated mode:** averages relevant neighboring pixels.
- **Improved interpolated mode (recommended):** interpolates the unknown component values. This mode reduces visible artifacts along object edges.




Converted images with no (top), standard (left) and improved interpolation method (right)

2.2. LUT for Gain/Offset (Color)

Separate gains and offsets can be applied to each of the three components of an image (contrast enhancement transform). The RGB image must be transformed to the targeted color space, gains and offsets applied, then transformed back to RGB.

- When applied to a mixture representation, all three gains and offset should vary in a similar way.

- When applied to luma/chroma representations, the gain and offset of the chromatic components should vary in a similar way.
- When applied to intensity/saturation/hue representation, it makes no sense to apply gain and offset to the hue component.



Enhanced saturation / Uniform lightness

Note: *The contrast enhancement function can be used to uniformize a given component: setting the gain to 0 for some component has the same result as setting all pixels to the value of the offset for this component.*

2.3. LUT for Color Calibration

Color distortions introduced by the image acquisition chain can be corrected by comparing sample colors from the image with their true values. A calibrated color chart, such as the IT8, is required.

- Sample colors are the average color in a suitable ROI using [PixelAverage](#).
- True color values are specified in the XYZ color system. Even though the reference colors are described by their XYZ coordinates, the image to be calibrated must contain RGB information.

The calibration transform can be based on one, three or four reference colors. In the first case, calibration is a gain adjustment for the three color components. In the second and third case, a linear or affine transform is used.

2.4. LUT for Color Balance

A color image can be improved by changing gamma correction and white balance.

These effects can be corrected efficiently by setting up a lookup table using [WhiteBalance](#) and applying it on a series of images by means of [Transform](#). The LUT need only prepared once (it implements a de-coupled color transformation).

Gamma Pre-Compensation

Many color cameras use a gamma pre-compensation process that deals with the non-linear response of the display device (such as a TV monitor).

Gamma pre-compensation should be used after processing because using it before would change the result because of the non-linearity introduced.

The pre-compensation process applies the inverse transform to the signal, so that the image renders correctly on the display. Three pre-defined gamma values are available, depending on the video standard at hand:

Video standard	Gamma value	EasyColor property
NTSC	1/2.2	CompensateNtscGamma
PAL	1/2.8	CompensatePalGamma
SMPTE	0.45	CompensateSmpteGamma

Note: Pre-compensation cancellation and pure pre-compensation correspond to exponents that are inverse of each other.

Gamma Pre-Compensation Cancellation

Many color cameras have a built-in gamma pre-compensation feature that can be turned off. If this feature cannot be turned off and is not desired, its effect can be canceled by applying the direct gamma transform. The following pre-defined gamma values are available for this purpose:

Video standard	Gamma value	EasyColor property
NTSC	2.2	NtscGamma
PAL	2.8	PalGamma
SMPTE	1/0.45	SmpteGamma

White Balance

A camera may exhibit color imbalance, i.e. the three color channels having mismatched gains, or the illuminant (the light sources) not being perfectly white. When this occurs, the white areas appear as an unsaturated color. The white balance correction automatically adjusts three independent gains so that the components of a white pixel become equal. This means that a white balance calibration step is required, during which a white surface must be shown to the camera and the corresponding color component are measured. [PixelAverage](#) can be used for this purpose.



Raw image, and image with white balance and gamma pre-compensation

3. EasyImage - Computing Image Statistics

EasyObject statistics are related to the objects in an image.

EasyImage statistics are related to whole images (global illumination / contrast, saturation, presence or absence of an object).


Sliding window (creates new image of avg or std deviation of gray-level values)

The average and standard deviation of gray-level values can be computed in a sliding window, i.e., computed for every position of a rectangular window centered on every pixel. The window size is arbitrary.

Note: *The computing time of these functions does not depend on the window size.*


The result of the operation is another image.

The **local average**, `EasyImage::LocalAverage`, corresponds to a strong low-pass filtering.



Sliding window average


The **local standard deviation**, `EasyImage::LocalDeviation` enhances the regions with a high frequency contents, such as noisy or textured areas.



Sliding window standard deviation

Histogram computation and analysis(and LUT creation)

A histogram is a statistical summary of an image: it shows the number of occurrences of every gray-level value in an image, and its shape reveals characteristics of the image. For instance, peaks in the histogram curve correspond to dominant colors in the image. If the histogram is bi-modal, a large peak for the dark values corresponding to the background, and smaller peaks in the light values.



Typical image histogram

Histogram Computation

`EasyImage::Histogram` computes the histogram of an image. It has an input mask argument.

It supports flexible mask.

BW8, BW16 and BW32 source images are supported.

You can compute the cumulative histogram of an image, i.e. the count of pixels below a given threshold value, by calling `EasyImage::CumulateHistogram` after

`EasyImage::Histogram`.

Histogram Analysis

`EasyImage::AnalyseHistogram` and


`EasyImage::AnalyseHistogramBW16` provide statistics and thresholding values:

- Total number of pixels.
- Smallest and largest pixel value (gray-level range).
- Average and standard deviation of the pixel values.
- Value and frequency of the most frequent pixel.
- Value and frequency of the least frequent pixel.

Histogram equalization

`EasyImage::Equalize` re-maps the gray levels so that the histogram fills in the whole dynamic range as uniformly as possible.


This may be useful to maximize image contrast, or reveal a lot of image details in dark areas.



Equalized image and histogram

Setup a lookup table

`EasyImage::SetupEqualize` creates a LUT so you can work explicitly with the histogram and LUT vectors. It can be more efficient to keep the image histogram for other purposes (i.e statistics) and keep the equalization LUT to apply to other images.



Equalization lookup table

Image focus

Sharp focusing can be achieved if the `EasyImage::Focusing` quantity is maximum for a given image. This function must be called multiple times with multiple images with a different focus for the basis of an "auto-focus" system.

`EasyImage::Focusing` computes the total gradient energy of the image. You can then use this gradient as a measure of the focusing of an image.

The gradients of the image show the edges of the structures present in the image, with strong values if the image is well-focused and weaker values otherwise.

To compute the total gradient energy of the image, Open eVision:

- a. Squares the pixel values of the horizontal and vertical gradient images.
- b. Averages the squared pixel values over both images.
- c. Sums the averages.
- d. Takes the square root of the resulting value.

The resulting value is maximum if the image is well-focused.



A well-focused image, with its (absolute-valued) horizontal and vertical gradients. The gradients show the edges of the structures with strong values. The total gradient energy for this image is 17.9.



A badly focused image, with its (absolute-valued) horizontal and vertical gradients. The gradients show the edges of the structures with weak values. The total gradient energy for this image is 7.9.

EasyImage statistics functions

Area (number of pixels with values above/on/between thresholds)

- `EasyImage::Area` counts pixels with values above (or on) a threshold.
- `EasyImage::AreaDoubleThreshold` counts pixels whose values are comprised between (or on) two thresholds.

Binary and weighted moments (object position and extent)

- `EasyImage::BinaryMoments` computes the 0th, 1st or 2nd order moments on a binarized image, i.e. with a unit weight for those pixels with a value above or equal to the threshold, and zero otherwise. It provides information such as object position and extent.
- `EasyImage::WeightedMoments` computes the 0th, 1st, 2nd, 3rd or 4th order weighted moments on a gray-level image. The weight of a pixel is its gray-level value. It provides information such as object position and extent.

Gravity center (average pixel coordinates above/on threshold)

- `EasyImage::GravityCenter` computes the coordinates of the gravity center of an image, i.e. the average coordinates of the pixels above (or on) the threshold.

Pixel count (between 2 thresholds)

- `EasyImage::PixelCount` counts the pixels in the three value classes separated by two thresholds.

Minimum, maximum and average gray-level value

- `EasyImage::PixelMax` computes the maximum gray-level value in an image.
- `EasyImage::PixelMin` computes the minimum gray-level value in an image.
- `EasyImage::PixelAverage` computes the average pixel value in a gray-level or color image. For a color image, it computes the means of the three pixel color components, the variances of the components and the covariances between pairs of components.

Average, variance and standard deviation

- `EasyImage::PixelStat` computes min, max and average gray-level values.
- `EasyImage::PixelVariance` computes average and variance of pixel values.
- `EasyImage::PixelStdDev` computes average and standard deviation of pixel values. For a color image, it computes the standard deviations and correlation coefficients (covariance over the product of standard deviations) of the pairs of pixel component values.

Number of different pixels by comparing 2 images

- `EasyImage::PixelCompare` counts the number of different pixels between two images.

4. EasyDeepLearning - Classifying Images

4.1. What EasyDeepLearning can do

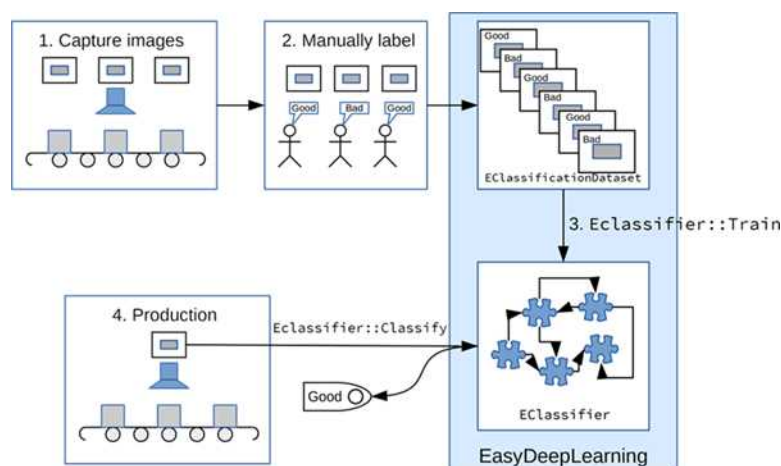
EasyDeepLearning is a library that classifies images using deep convolutional neural networks (CNNs). You can use it, for example, to identify a product in an image or to detect if the product is good or defective.

By opposition to traditional machine vision techniques, EasyDeepLearning does not require an explicit model of what to recognize inside an image. Instead, it learns this model from a set of example images. Thus it is able to solve machine vision problems where an explicit model is too complex to build.

To accelerate computations, we strongly recommend running the EasyDeepLearning library on a recent NVIDIA GPU. Refer to the section "Accelerating the Computation with a NVIDIA GPU" on page 87 for installing the required NVIDIA CUDA and deep learning library.

Note: The EasyDeepLearning library runs only on x64 platforms.

4.2. Workflow



To create an application based on EasyDeepLearning:

1. Capture a dataset of images representative of the problem you want to solve.
 - The capture conditions must be as close as possible of the production conditions.
 - Preferably, all images should have the same resolution.
 - The number of images needed to obtain a good classification performance depends on the complexity of the classification task.
However, EasyDeepLearning can solve some classification task steps with as few as 100 images per class.
2. Manually label the images in the dataset with the different categories you want to recognize.
3. Train a deep learning classifier on the labeled dataset.
 - a. Create an *EClassificationDataset* object.
 - b. Use the *EClassificationDataset::AddImage* method to add the manually labeled images to the object.
 - c. Use the *EClassificationDataset::SplitDataset* method to split the dataset into 2 parts called the training and the validation datasets,
Or manually create a new *EClassificationDataset* object for the validation dataset.
 - d. Create an *EClassifier* object.
 - e. Call the *EClassifier::Train* method with the *EClassificationDataset* objects as argument.
 - f. Use the *EClassifier::GetValidationMetrics* method to analyze the performance of the trained classifier by retrieving performance metrics computed on the validation dataset.
4. Use the *EClassifier::Classify* method to use the trained *EClassifier* object in production and classify new images.

The method returns an *EClassificationResult* object that contains a probability for each. Use the *EClassificationResult::GetBestLabel* method to get the label with the highest probability.

Overtfitting

It is important to use 2 separate datasets of images:

- A training dataset to train the classifier.
- A validation dataset to evaluate how well the classifier performs on images it was not trained on.

Deep learning techniques can suffer from overfitting; this means that the trained classifier is too focused on the specific images present in the training dataset and it is not able to learn a general model for the different labels. Such a classifier performs poorly in production.

Overfitting can happen when the training dataset is not representative enough of the problem you want to solve.

Solutions to prevent overfitting can be to capture additional images or, in some cases, to use data augmentation.

Data augmentation generates random transformations of the images in the training dataset to make the classifier robust to rotations, reflections or scaling that are not present in the original training dataset.

4.3. EasyDeepLearning Studio

EasyDeepLearning Studio

The EasyDeepLearning Studio is a graphical user interface for:

- Creating datasets and labeling images,
- Configuring and visualizing the data augmentation transformations,
- Training a classifier,
- Analyzing the performance of the classifier.

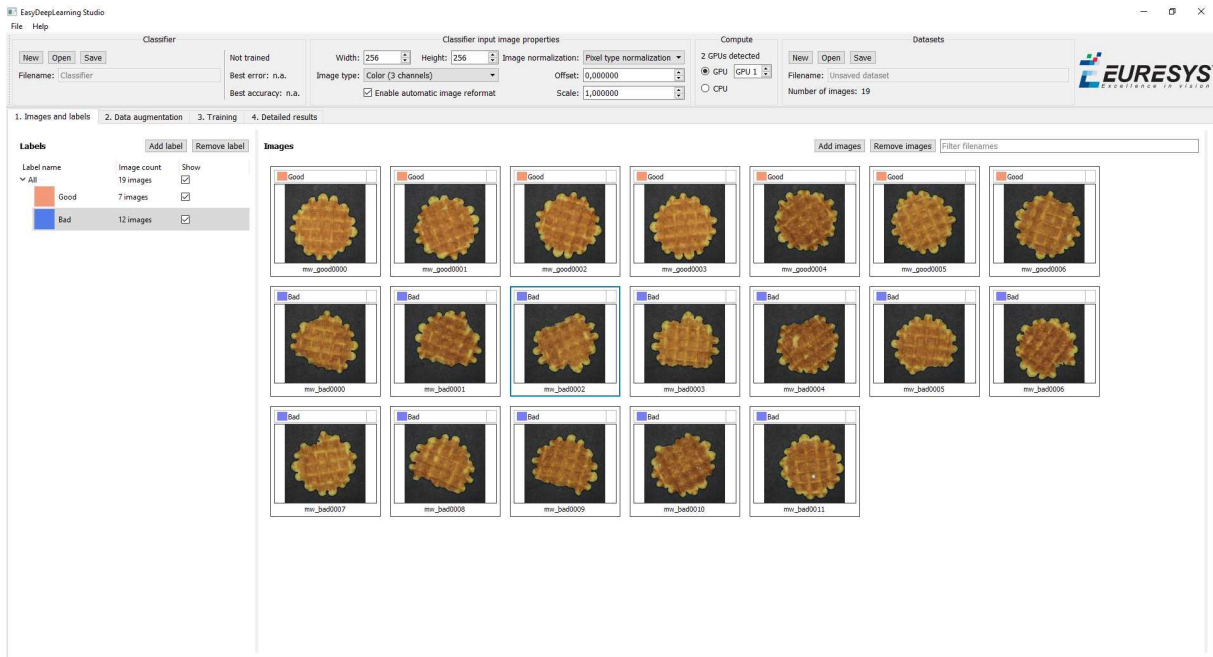
The EasyDeepLearning Studio is available in the installation folder of Open eVision.

Resources and code snippets

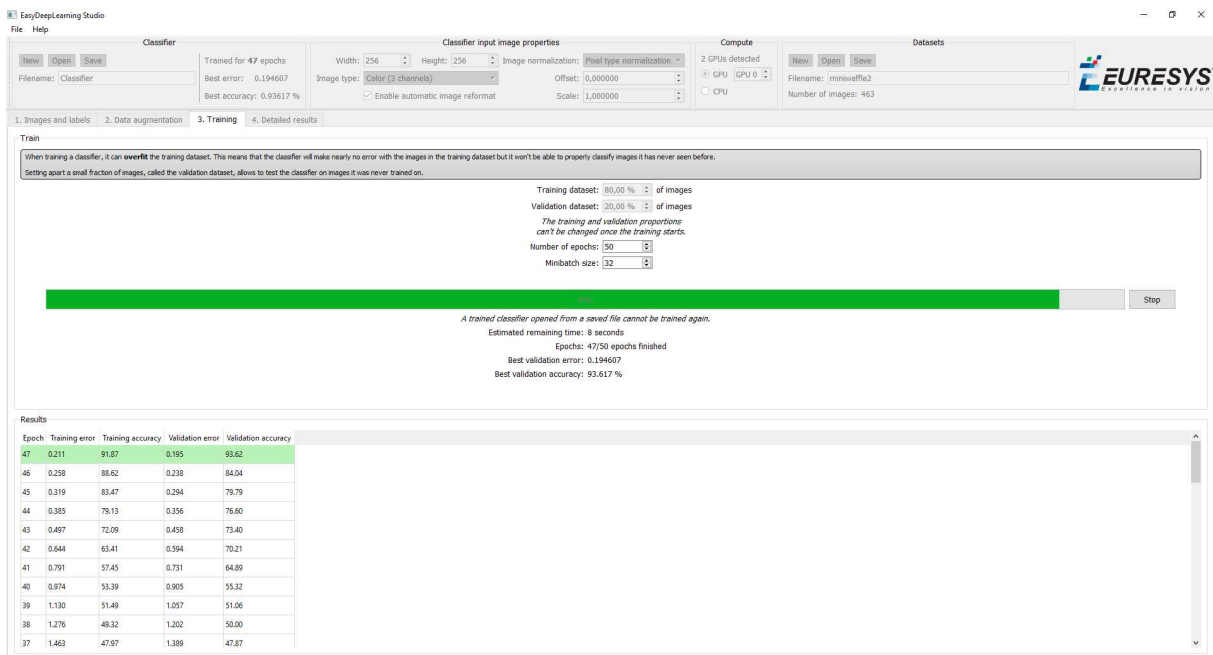
- The sample dataset called “MiniWaffle” illustrated below is available in the folder `Sample Images/EasyDeepLearning/MiniWaffle` of the installation folder. This dataset contains images of good and bad mini waffles.
- Some [code snippets](#) are also provided for illustration and reference.

Workflow illustration with EasyDeepLearning Studio

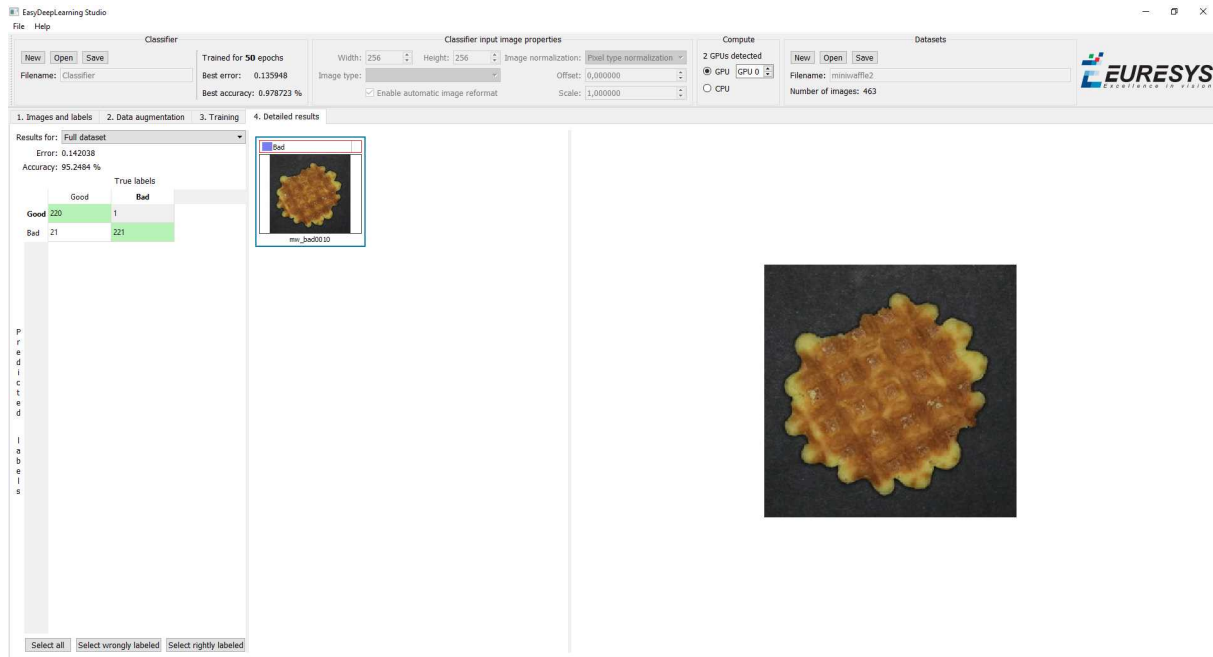
You can use the EasyDeepLearning Studio to perform steps 2 and 3 of the process described in section ["Workflow"](#) on page 78.



Manual labeling of images (step 2) and creating the dataset (steps 3a and 3b)



Splitting the dataset (step 3c) and starting the training (steps 3d and 3e)



Analyzing the performance (step 3f)

4.4. Managing the Dataset

Adding images

The supported image types are *ElImageType_BW8*, *ElImageType_BW16*, and *ElImageType_C24*.

- To add a single image and its label to a *EClassificationDataset*, use one of these methods:
 - *EClassificationDataset::AddImage(path, label)* for an image file,
 - *EClassificationDataset::AddImage(img, label)* for an Open eVision image object.
- To add several image files that share the same label, use the method *EClassificationDataset::AddImages(filter, label)*

where *filter* is a glob pattern with the wildcard characters:

- * means "zero or more character",
- ? means "a single character".

For example, *EClassificationDataset::AddImages("*good*.png", "good")* will add all PNG image files that contains "good" in their filename.

EasyDeepLearning automatically generates the set of labels from the labels of the images that you add to the dataset.

Note: The labels are case sensitive.

Setting label weights

Label weights represent the relative importance that the classifier gives to each class during the training.

By default, all labels are given the same weight of 1.

Set the weight of a label with `EClassificationDataset::SetLabelWeight(labelId, weight)`.

Using data augmentation

Data augmentation performs random transformations on images given to a `EClassifier` object during the training.

Use `EClassificationDataset::SetEnableDataAugmentation(true/false)` to enable or disable those transformations.

The possible transformations are:

- Horizontal and vertical flips (enabled with `EClassificationDataset::SetEnableHorizontalFlip` and `EClassificationDataset::SetEnableVerticalFlip`)
- Scaling (between a minimum and maximum value defined with `EClassificationDataset::SetMinScale` and `EClassificationDataset::SetMaxScale`)
- Horizontal and vertical shifts (between `-maxValue` and `maxValue` defined with `EClassificationDataset::SetMaxHorizontalShift(maxValue)` and `EClassificationDataset::SetMaxVerticalShift(maxValue)`)
- Rotations (between 0 and a maximum value defined with `EClassificationDataset::SetMaxRotationAngle`)
- Horizontal and vertical shear (between `-maxValue` and `maxValue` defined with `EClassificationDataset::SetMaxHorizontalShear` and `EClassificationDataset::SetMaxVerticalShear`)
- Brightness offset (between `-maxValue` and `maxValue` defined with `EClassificationDataset::SetMaxBrightnessOffset`)
- Contrast gain (between a minimum and maximum value defined with `EClassificationDataset::SetMinContrastGain` and `EClassificationDataset::SetMaxContrastGain`)
- Gamma corrections (between a minimum and maximum value defined with `EClassificationDataset::SetMinGamma` and `EClassificationDataset::SetMaxGamma`)
- Hue offset (between `-maxValue` and `maxValue` defined with `EClassificationDataset::SetMaxHueOffset`)
- Saturation gain (between a minimum and maximum value defined with `EClassificationDataset::SetMinSaturationGain` and `EClassificationDataset::SetMaxSaturationGain`)

Splitting the dataset

To evaluate how well the classifier performs, set some images apart and do not use them during the training.

To do this, either:

- Directly create 2 *EClassificationDataset* objects containing 2 different sets of images,
- Or randomly split an *EClassificationDataset* dataset into 2 parts with the method *EClassificationDataset::SplitDataset(trainingDataset, validationDataset, trainingProportion)*.

4.5. Using the Classifier

Input image format and normalization

Note: *The input image must be larger than a minimum size of 128 x 128.*

To ensure numerical stability, all the images given to an *EClassifier* object should have the same format (width, height and number of channels) and should be normalized.

- By default, a classifier uses the image format of the first image inserted in the training dataset:
 - All other images are automatically reformatted (anisotropic rescaling and conversion between color and grayscale).
 - If *EClassifier::SetEnableAutomaticImageReformat(false)* is called, the classifier throws an exception when attempting to train or classify an image that does not have the correct image format.
- You can also set manually the input image format with the methods *SetWidth*, *SetHeight* and *SetChannels* (1 channel for grayscale images and 3 channels for color images).

We recommend that you reduce the resolution of the original images using these methods to increase the processing speed of the classifier. Note that the “object of interest” must still be recognizable at the final resolution.

Histogram equalization

The classifier can also apply an histogram equalization to every input image. Use *EClassifier::SetEnableHistogramEqualization(true)* to activate it.

Training the classifier

To train a classifier, call the method `EClassifier::Train(trainingDataset, validationDataset, numberOfIterations)`.

- Iteration:
 - An iteration corresponds to going through all the images in the training dataset once.
 - The training process requires a large number of iterations to obtain good results.
 - The default number of iterations is 50.
 - The larger the number of iterations, the longer the training is and the better the results you obtain.

Calling the `EClassifier::Train()` method several times with the same training and validation dataset is equivalent to calling `EClassifier::Train()` once but with a larger number of iterations. The total number of iterations used to train the classifier is accessible through `EClassifier::GetNumTrainedIterations()`.

- The training process is asynchronous:
 - A call to `EClassifier::Train` launches a new thread that does the training in background.
 - The method `EClassifier::WaitForTrainingCompletion` suspends the program until the whole training is completed.
 - The method `EClassifier::WaitForIterationCompletion` suspends the program until the current iteration is completed.
 - During the training, use `EClassifier::GetTrainingProgression()` to follow the progression of the training.

- Mini batch size:

The mini batch size corresponds to the number of images that are processed together.

- The training is influenced by the mini batch size.
- A large mini batch size increases the processing speed of a single iteration but requires more memory.
- The training process is not able to learn a good model with too small or too large mini batch sizes.
- The default size of the mini batch is 32 (we recommend that you do not change this value unless the training fails because you don't have enough memory).
- It is common to choose powers of 2 as mini batch size for performance reasons.
- Use `EClassifier::SetMiniBatchSize` to change the size of your mini batch.

Validating the results

- After the completion of each iteration, EasyDeepLearning automatically computes several performance metrics about the training and validation dataset:
 - Call the methods `EClassifier::GetTrainingMetrics(iteration)` and `EClassifier::GetValidationMetrics(iteration)` to read these metrics.
 - The iterations are indexed between 0 and `EClassifier::GetNumTrainedIterations()-1`.
 - Call `EClassifier::GetBestIteration()` to retrieve the iteration that produced the best performance.
 - After the training, the classifier is back in the state corresponding to this best iteration.
- The metrics are represented by an `EClassificationMetrics` object that contains the following performance metrics:
 - The classification error (`EClassificationMetrics::GetError()`): the quantity that is minimized during the training.
 - The classification accuracy (`EClassificationMetrics::GetAccuracy()`): the number of images correctly classified divided by the total number of images in the dataset.
 - The confusion matrix (`EClassificationMetrics::GetConfusion(groundtruthLabel, predictedLabel)`): the number of images labeled as `groundtruthLabel` that are classified as `predictedLabel`.

Call `EClassifier::Evaluate` to evaluate a dataset independently of the training.

Classifying new images

Once the classifier is trained, call `EClassifier::Classify` to classify an Open eVision image.

This method returns a `EClassificationResult` object:

- `EClassificationResult::GetBestLabel()` returns the most probable label for the image.
- `EClassificationResult::GetBestProbability()` returns the probability associated with the most probable label.
- `EClassificationResult::GetProbability(label)` returns the probability associated with the given label.
- `EClassificationResult::GetRanking(label)` returns the ranking of the given label. The ranking goes from 1 (most probable) to `EClassifier::GetNumLabels()` (least probable).

You can also directly classify a vector of Open eVision images:

- Images are processed together in groups determined by the mini batch size.
- It is usually much faster to classify a group of images than a single image, especially on a GPU.

4.6. Accelerating the Computation with a NVIDIA GPU

CPU versus GPU

- Deep learning algorithms perform a lot of computations and can be very slow to train on a CPU.

For example, on a high-end Intel Core i9-7900X CPU with a single thread, with no data augmentation:

- The training can process up to 0.5 MegaPixels/seconds.
- The validation and classification can process up to 1.5 MegaPixels/seconds.

- Using a recent NVIDIA GPU greatly accelerates the processing speeds.

On a NVIDIA GeForce 1080Ti, with no data augmentation:

- The training can process up to 50 MegaPixels/seconds.
- The validation can process up to 160 MegaPixels/seconds.
- The classification of a single image can process up to 55 MegaPixels/seconds (equivalent to more than 800 256 x 256 grayscale images/second).

Please be aware that the actual speed varies with the input image format, the data augmentation, the batch size and the GPU model.

Using an NVIDIA CUDA® GPU

1. To use an NVIDIA GPU with the EasyDeepLearning library, install the following NVIDIA libraries on the computer:

- NVIDIA CUDA® Toolkit version v9.0 (<https://developer.nvidia.com/cuda-toolkit>)
- NVIDIA CUDA® Deep Neural Network library (cuDNN) v7.0.5 (<https://developer.nvidia.com/cudnn>)

2. According to the installation location:


- If you install the NVIDIA CUDA® Toolkit in its default location (C:\Program Files\NVIDIA GPU Computing Toolkit\CUDA\v9.0), the EasyDeepLearning library automatically finds what it needs.
- Otherwise, the DLLs `cusolver64_90.dll`, `curand64_90.dll`, `cuffft64_90.dll` and `cublas64_90.dll` should be copied in the Open eVision DLL folder (its default location is C:\Program Files (x86)\Euresys\Open eVision X.X\Bin64\).

3. Install the NVIDIA CUDA[®] Deep Neural Network library (cuDNN) that comes as a zip archive:
 - a. Unzip the files.
 - b. Copy the unzipped files to the NVIDIA CUDA[®] Toolkit installation directory as indicated in <https://docs.nvidia.com/deeplearning/sdk/cudnn-install/index.html#installwindows>.
 - c. If the NVIDIA CUDA[®] Toolkit is not installed in its default location, copy the DLL file `cuda64_7.dll` in the Open eVision DLL folder (its default location is `C:\Program Files (x86)\Euresys\Open eVision X.X\Bin64\`).
4. Use the method `EClassifier::SetEnableGPU` to tell EasyDeepLearning to use the GPU.
5. If there are more than one GPU installed on the computer, set the index of the GPU to use with the method `EClassifier::SetGPUIndex`.

PART IV
INSPECTION TOOLS


1. EasyObject - Analysing Blobs

The **EasyObject** library picks out features in an image by creating and processing blobs (objects or holes that have the same gray level range).



This library can be used for BW1, BW8, BW16 and C24 source images and is accessible from the `ECodedImage2` class which has improved execution time, especially for large images with many objects.

Workflow



Blob Definition

A blob is a grouping of neighboring pixels of the same gray level range. Blobs may be objects or holes in objects. EasyObject functions analyze both objects and holes.

When blobs are built, the inclusion relationship between holes and objects is computed.

Even though holes may be the actual objects of interest, it is easier to find an object of interest, then detect its holes (with EasyObject) and measure their characteristics (with EasyGauge or EasyObject).

Blobs are handled as independent entities:

- They can be selected by means of the layer they belong to, their position, a rectangular ROI or their computed features. The selection criteria can be combined (select the small objects; among these, select those close to the right edge...).
- They can be listed and sorted by their geometric characteristics: such as area, width, or ellipse of inertia.

Blob analysis can be restricted to rectangular and nested ROIs, and to complex or disconnected-shape regions using flexible masks.

Build Blobs

EasyObject chooses objects of interest and constructs blobs/holes in two steps:

1. **Segment**: classifies the source image pixels, creates layers, and constructs the runs (a run is a sequence of adjacent pixels in a row, that share the same property).
2. **Encode**: assembles runs, to build blobs for each layer.
You select which objects or holes are kept.

`EImageEncoder::Encode` analyzes the blobs and stores the result into a coded image which has a set of superimposed, mutually exclusive image layers, where the pixels of each layer have properties in common, such as being above a threshold.

Flexible masks can restrict encoding to an arbitrary shaped area.

There is no need to build **holes**, they are constructed on-the-fly when required.

Functions

- Segmentation `GetSegmentationMethod` and `SetSegmentationMethod`
- Grayscale single threshold `EGrayscaleSingleThresholdSegmenter`
- Grayscale double threshold `EGrayscaleDoubleThresholdSegmenter`
- Color single threshold `EColorSingleThresholdSegmenter`
- Color range threshold `EColorRangeThresholdSegmenter`
- Reference image `EReferenceImageSegmenter`
- Image range `EImageRangeSegmenter`
- Labeled image `ELabeledImageSegmenter`
- Binary images `EBinaryImageSegmenter`


Pixel aggregation (encoder)

- Layer selection
- Object construction: run aggregation into objects
- **Hole construction**: run aggregation into holes

Extract objects (using geometric parameters)

Once an image has been encoded, the coded elements (objects or holes) are accessible through the abstract class `ECodedElement` which provides a large set of methods applicable to a particular coded element:

Num	Area	Gravity Center X	Gravity Center Y	
59	2221	17.67	95.55	
37	387	161.69	53.46	
47	344	166.43	90.24	
32	327	226.23	72.07	
40	251	111.45	62.04	
50	239	120.41	91.51	
4	144	220.98	2.44	
68	142	72.05	123.10	



Features computation and display

The objects, holes and their features can be efficiently accessed randomly (in an index-based fashion).

1.1. Image Segmenters

There are several ways to segment pixels. The method is chosen with `GetSegmentationMethod` and `SetSegmentationMethod`.

1. Grayscale Single Threshold (default)

`EGrayscaleSingleThresholdSegmenter` is applicable to BW8 and BW16 grayscale images and produces coded images with two layers:

- The **black layer** (usually Layer 0) contains unmasked pixels with a gray value below the Threshold value.
- The **white layer** (usually Layer 1) contains the remaining unmasked pixels, i.e. unmasked pixels having a gray value greater or equal to the Threshold value.

EasyObject provides 5 thresholding methods:

- **Absolute** (integer value): represents the first gray value of the white layer. Set with `SetAbsoluteThreshold` method and got with `GetAbsoluteThreshold` method.
- **Relative** (%): represents the fraction of image pixels that belong to the Black layer, it is a user-defined float value in range 0 to 1. Set with `SetRelativeThreshold` method and got with `GetRelativeThreshold` method.
- **Minimum Residue** (default): The threshold is an automatically computed value such that the quadratic difference between the source and thresholded image is minimized.
- **Maximum Entropy**: automatically computed value such that the entropy (i.e. the amount of information) of the resulting thresholded image is maximized.

- **IsoData**: automatically computed value that lies halfway between the average dark gray value (gray levels below the threshold) and average light gray values (gray levels above the threshold).

Grayscale Single Threshold with a minimum residue thresholding method is the default. Only objects whose pixels have a value that is above this threshold are encoded.

2. Grayscale Double Threshold

`EGrayscaleDoubleThresholdSegmenter` is applicable to BW8 and BW16 grayscale images and produces coded images with three layers:

- The **black layer** (usually Layer 0) contains unmasked pixels having a gray value below the Low Threshold value.
- The **white layer** (usually Layer 2) contains unmasked pixels having a gray value above or equal the High Threshold value.
- The **neutral layer** (usually Layer 1) contains the remaining unmasked pixels.

The **Low Threshold** and **High Threshold** are user-defined integer values, set with `SetLowThreshold` and `SetHighThreshold` methods, and got with `GetLowThreshold` and `GetHighThreshold` methods.

3. Color Single Threshold

`EColorSingleThresholdSegmenter` is applicable to C24 color images; it produces coded images with two layers:

- The **white layer** (usually Layer 1) contains unmasked pixels that belong to the cube of the color space defined by the threshold point and the white point (255,255,255).
- The **black layer** (usually Layer 0) contains the remaining unmasked pixels.

The **Color Threshold** is a set of three **user-defined** integer values designating a color in the color space, set with `SetThreshold` method and got with `GetThreshold` method.

4. Color Range Threshold

`EColorRangeThresholdSegmenter` is applicable to C24 color images; it produces coded images with two layers:

- The **white layer** (usually Layer 1) contains unmasked pixels that belong to the cube of the color space defined by the Low Threshold point and the High Threshold point.
- The **black layer** (usually Layer 0) contains the remaining unmasked pixels.

The Low Threshold and High Threshold are each a set of three user-defined integer values designating a color in the color space, set with `SetLowThreshold` and `SetHighThreshold` methods and got with `GetLowThreshold` and `GetHighThreshold` methods.

5. Image Range

The following cases need a segmentation using **pixel-by-pixel thresholding** which gives an allowed range of values for each pixel:


- when the background is not uniform enough,
- when the illumination is not uniform across the image,
- when only differences between the image and a reference image (ideal) are to be enhanced,

The allowed range for each pixel is specified using two images: a low reference image with the minimum values allowed for each pixel, a high reference image with the maximum values. The reference images are thus the source image minus (or plus) a fixed value all over the image (assuming noise distribution is uniform and additive).

The difficulty is preparing suitable high and low reference images.


Preparing high and low reference images

You can start from an image of the scene without defects and add security margins before comparison.




Source image

Gray-level tolerance must be provided for noise and illumination variations.




Gray-level tolerance margins

The image may have a slight shift in some direction which can be corrected by enlarging the light and dark areas using dilate and erode morphological operations. This geometric tolerance margin is roughly as large as the morphological filter size.



Geometric tolerance margins

Combining both kinds of tolerance margins gives the best results.



Combined margins

Image Segmenter

[EImageRangeSegmenter](#) and [EReferenceImageSegmenter](#) are applicable to BW8, BW16, and C24 images; and produce coded images with two layers.

The low threshold and the high threshold are defined for each pixel individually by means of two reference images of the same type as the source image: the Low Image and the High Image. The Reference Image defines the reference threshold of each pixel is individually.

- For **grayscale** images, the **white layer** (usually Layer 1) contains unmasked pixels having a gray value in a range defined by the gray value of the corresponding unmasked pixels in the Low, High or Reference Image.
- For **color** images, the **white layer** (usually Layer 1) contains unmasked pixels having a color inside the cube of the color space defined by the colors of the corresponding unmasked pixels in the Low, High or Reference Image.
- The **black layer** (usually Layer 0) contains the remaining unmasked pixels.

Pointers to the **Low Image** can be set or got using the functions associated with the type of the source image:

- BW8: [SetLowImageBW8](#) [GetLowImageBW8](#)
- BW16: [SetLowImageBW16](#)[GetLowImageBW16](#)
- C24: [SetLowImageC24](#)[GetLowImageC24](#)

Pointers to the **High Image** can be set or got using the functions associated with the type of the source image:

- BW8: [SetHighImageBW8](#)[GetHighImageBW8](#)
- BW16 [SetHighImageBW16](#)[GetHighImageBW16](#)
- C24 [SetHighImageC24](#)[GetHighImageC24](#)

Pointers to the **Reference Image** can be set or got using the functions associated with the type of the source image:

- BW8: [SetReferenceImageBW8](#), [GetReferenceImageBW8](#)
- BW16: [SetReferenceImageBW16](#), [GetReferenceImageBW16](#)
- C24: [SetReferenceImageC24](#) , [GetReferenceImageC24](#)

6. Labeled Image

[ELabeledImageSegmenter](#) is applicable to is applicable to BW8 and BW16 grayscale images; it produces coded images with a number of layers equal to the maximum number of gray values: 256 for BW8 images or 65536 for BW16 images. The layer n contains all the unmasked pixels

having a gray value equal to n .

By default, all layers are encoded. However, it is possible to restrict the encoding to a single range of layers with `SetMinLayer` and `SetMaxLayer` functions which return the lowest and the highest values of the index range respectively.


7. Binary Image

`EBinaryImageSegmenter` is applicable to BW1 binary images; it produces coded images with two layers:

- Black layer (usually index 0) contains unmasked pixels with a binary value equal to zero.
- White layer (usually index 1) contains the remaining unmasked pixels, i.e. unmasked pixels with a binary value equal to one.

1.2. Image Encoder

The class representing the objects (`EObject`) derives from an abstract class `ECodedElement`.



Object building

Selecting the Layers to Encode

The segmentation methods (see [Image Segmenters](#)) determine which layer(s) to encode by default, and do not encode pixels from the other layers.

Function `GetMaxLayerIndex` returns the highest Layer Index value . It is available for all segmenters.

Enabling/disabling layer encoding for each layer individually

The following tables list, for each layer, the Set/Get function and the default enable/disable value.

Two-layer segmenters

Layer	Set LayerEncoded function	Get LayerEncoded function	Default value
Black layer	<code>SetBlackLayerEncoded</code>	<code>IsBlackLayerEncoded</code>	FALSE
White layer	<code>SetWhiteLayerEncoded</code>	<code>IsWhiteLayerEncoded</code>	TRUE

Three-layer segmenters

Layer	Set LayerEncoded function name	Get LayerEncoded function name	Default value
Black layer	SetBlackLayerEncoded	IsBlackLayerEncoded	FALSE
White layer	SetWhiteLayerEncoded	IsWhiteLayerEncoded	FALSE
Neutral layer	SetNeutralLayerEncoded	IsNeutralLayerEncoded	TRUE

Manually Assigning a Layer Index to Each Layer Individually

The following tables list, for each layer, the Set/Get function and the default value.

Two-layer segmenters


Layer	Set LayerEncoded function name	Get LayerEncoded function name	Default value
Black layer	SetBlackLayerIndex	IsBlackLayerIndex	0
White layer	SetWhiteLayerIndex	IsWhiteLayerIndex	1

Three-layer segmenters

Layer	Set LayerEncoded function name	Get LayerEncoded function name	Default value
Black layer	SetBlackLayerIndex	IsBlackLayerIndex	0
Neutral layer	SetNeutralLayerIndex	IsNeutralLayerIndex	1
White layer	SetWhiteLayerIndex	IsWhiteLayerIndex	2

Runs


For the sake of computational efficiency, the objects are described as lists of runs. A run is a sequence of adjacent pixels that share homogeneous properties (such as being above a given threshold). These runs are merged in objects by the image encoder.



A single object with five enhanced runs

EasyObject can work at object level, and at run level which allows faster processing in critical cases. This is useful to compute custom features on objects then list all runs belonging to a given object as shown in this example of working at run level, with colored runs in the output image.

1. **Declare a new `ECodedImage2` object.**
2. **Declare an `EImageEncoder`** and, if applicable, select the appropriate segmenter. Setup the segmenter and choose appropriate layer(s) to encode.
3. **Setup an output image.**
4. **Encode the image.**
5. **Color the runs in the output image.** Iterate over the objects of a specific layer by constructing a loop and then a `RunsIterator` object. This iterator allows to browse runs of the considered object. Once the iterator has finished a run of the considered object, the inner loop processes the pixels spanned by this run in the output image.
6. **Select a specific layer.**



Source image (left) with the white layer rendered (right)

Connexity

Pixels can touch each other along an edge or by a corner. In Four Connexity only pixels touching by an edge are considered neighbors. In Eight Connexity (the default) pixels touching by a corner are also considered neighbors.



4-connexity



8-connexity


Multiple images can be encoded in [continuous mode](#).

1.3. Holes Construction


A hole is a set of connected pixels that are entirely surrounded by a parent object (4 or 8 pixels depending on the connexity mode).

A hole has no child. Objects inside a hole are considered as separate objects.

`EObject` and `EHole` classes both derive from `ECodedElement`, so objects and holes are managed in the same way and share the same functions.



Encoding the white layer (3 objects and 3 holes)




Encoding the black layer (4 objects and 3 holes)

How to Color holes

1. **Declare a new** `ECodedImage2` object.
2. **Declare an** `EImageEncoder` and, if applicable, select and setup the appropriate segmenter, and choose the appropriate layer(s) to encode.
3. **Setup an output image.**
4. **Encode the image.**
5. **Declare a helper function to draw the runs.** A helper function (see also section Object Construction/Working at the Run Level) draws the runs in an output image, using, for example, a given color. This function can be shared for objects and holes.

6. **Draw the objects and their holes in the output image.** It is necessary to iterate over the objects of the chosen layer.
 - a. The helper function draws the runs of each object (`DrawRuns`) using a specific color.
 - b. The holes are iterated over the current object, and their runs are drawn.
 - c. Each hole of an object is drawn with a different color computed in the global function (`GetFadedColor`) which returns a color that depends upon the hole index, for example a gradation of red to green colors.



Raw image (left) Building of objects and all holes (right)

1.4. Normal vs. Continuous Mode

Normal Mode (default)

In normal mode, the image encoder does not track blobs across several successive images. EasyObject works with one image, without keeping blobs in memory. All the blobs are returned as objects.


Continuous Mode

In continuous mode EasyObject can process an image whose height is unknown or infinite (e.g. coming from a line-scan camera). The image is split into a several chunks that are fed into an image encoder. Objects that straddle several successive image chunks can be detected.

The image encoder encodes only the objects that contain no run touching the last row of the source image. Objects that touch the inferior border of the image are not written in the coded image because they are expected to continue in subsequent image chunks, but they are kept in memory and are processed when subsequent images are analyzed.

A large image is assumed to be divided in several chunks that are stored in the array `EImageBW8 chunk[x]`.

In this example, we generate a sequence of color images that exhibit objects encoded over successive chunks




Original image




Three chunks of the image


1. **Draw the objects encoded in a layer of a coded image.** This code is essentially the same as in "Browsing Runs" code snippet. The only difference is that an offset can be applied along the Y-axis.
2. **Define a function to draw the objects of a layer.** If a coded image contains objects that were started in a previous image: the runs of this object from the previous image are assigned with a negative Y-coordinate.
The zero Y-coordinate is the first row of the most recently encoded image. The convention is to assign the lowest Y-coordinate to the oldest run in the encoded objects.
The method `EImageEncoder::GetStartY` obtains the Y-coordinate of this oldest run. It is necessary to define a function that displays the content of a layer of a coded image. Each object can be displayed with a different color(computed by `GetFadedColor`). This function closely follows the function `DrawRuns`, but is adapted to continuous mode by taking **GetStartY** into account.
3. **Enable continuous mode** in property `EImageEncoder::SetContinuousModeEnabled`. Additional variables can be declared, for example to store the successive encoded image, or to hold the output images.
4. **Analyze the successive chunks.** To encode successive chunks use `Encode(chunk[count], codedImage)` and then `DrawLayer`. **Note:** The variable count spans integers 0, 1 and 2. When an object from a chunk is not complete it is kept in the internal memory of the image encoder.



Content of layerImage when count equals 0, after the application of DrawLayer.
 Chunk of the large image that is under consideration.
 Note that two objects in the lower-left of the image chunk are not encoded,
 because they touch the border of the chunk.




When count reaches 1, one of these two objects becomes completed,
 which leads to the encoding of the following image.
 Two other objects are not encoded yet at this time.
 Here is the result of the encoding of the last chunk (count = 2).



Three objects from the previous chunks have been closed, and have thus been encoded.

Flushing Continuous Mode

After encoding the three image chunks, there remains one object to be completed (in the bottom-right corner of the large image). However, as there are no more chunks, it is necessary to explicitly close this object and encode the remaining object using the [flushing of the image encoder](#). The internal memory of the image encoder is then empty.



Result of the flush

1.5. Selecting and Sorting Blobs

The object segmentation process considers any blob as an object, including noise pixels which appear as tiny objects. You can select which blobs to keep using the `EObjectSelection` class.

Create / Modify Selection

You can use the `EObjectSelection` `Add` and `Remove` methods to:

- Add or remove a single object, a hole or a whole layer to/from a selection.
- Add or remove objects or holes based on some specified **feature** (see the feature list in [Computing the Coded Element Features](#)).
- Add or remove objects or holes based on their specific **position**, or whether they lie within a specified ROI rectangle.

These actions can be cascaded and combined at will in a single selection.

Clear selection


You can clear a previous selection using `EObjectSelection::Clear`.

Sort selection

You can sort the elements of a selection according to any of their features.

Example

In this example, we select objects in the middle band of an image, with a surface >100 pixels.



Source image, and selection of objects

1. **Declare a new** `ECodedImage2` object.
2. **Declare an** `EImageEncoder` object and, if applicable, select and setup the appropriate segmenter and choose the appropriate layer(s) to encode.
3. **Encode the source image.**
4. **Create a selection of objects.** Create an instance of the `EObjectSelection` class and add objects to this selection, for instance through `EObjectSelection::AddObjects`.
5. **Remove objects based on the value of one feature at a time.** The objects in a selection can be unselected by calling one of the `EObjectSelection::Remove` methods.
6. **Remove the objects based on their position** using `EObjectSelection::RemoveUsingFloatFeature`. For details, see also "Working at the Run Level".
7. **Sort the selected objects** using `EObjectSelection::Sort`.
8. **Access the selected objects.**

1.6. Advanced Features

Computable Features

Methods prefixed with **Get** indicate a lazy evaluation: the result is computed on the first invocation of the method and cached.

Methods prefixed with **Compute** indicate that the function is reevaluated at every invocation and the result is never cached.

Position

Limit (top, bottom, left, right)	<code>ECodedElement::GetTopLimit</code>
	<code>ECodedElement::GetBottomLimit</code>
	<code>ECodedElement::GetLeftLimit</code>
	<code>ECodedElement::GetRightLimit</code>
Gravity center (X and Y)	<code>ECodedElement::GetGravityCenter</code>
	<code>ECodedElement::GetGravityCenterX</code>
	<code>ECodedElement::GetGravityCenterY</code>
Weight gravity center (X and Y)	<code>ECodedElement::ComputeWeightedGravityCenter</code>

Gravity center and weight gravity center



The **gravity center** returns the abscissa of the gravity center of the coded element.

The **weight gravity center** computes the gravity center of a given image over a coded element.

Extents

Area (pixel count)	<code>ECodedElement::Area</code>
Feret box (center X and Y, height, width with distinct orientation angles at 22, 45, 68 degrees)	<code>ECodedElement::ComputeFeretBox</code>
	<code>ECodedElement::GetFeretBox22Box</code>
	<code>ECodedElement::GetFeretBox22Center</code>
	<code>ECodedElement::GetFeretBox22CenterX</code>
	<code>ECodedElement::GetFeretBox22CenterY</code>
	<code>ECodedElement::GetFeretBox22Height</code>
	<code>ECodedElement::GetFeretBox22Width</code>
	<code>ECodedElement::GetFeretBox45Box</code>
	<code>ECodedElement::GetFeretBox45Center</code>
	<code>ECodedElement::GetFeretBox45CenterX</code>
	<code>ECodedElement::GetFeretBox45CenterY</code>
	<code>ECodedElement::GetFeretBox45Height</code>
	<code>ECodedElement::GetFeretBox45Width</code>
	<code>ECodedElement::GetFeretBox68Box</code>
	<code>ECodedElement::GetFeretBox68Center</code>
	<code>ECodedElement::GetFeretBox68CenterX</code>
<code>ECodedElement::GetFeretBox68CenterY</code>	
<code>ECodedElement::GetFeretBox68Height</code>	
<code>ECodedElement::GetFeretBox68Width</code>	
Bounding box (center X and Y, height, width)	<code>ECodedElement::GetBoundingBox</code>
	<code>ECodedElement::GetBoundingBoxCenter</code>
	<code>ECodedElement::GetBoundingBoxCenterX</code>
	<code>ECodedElement::GetBoundingBoxCenterY</code>

	ECodedElement::GetBoundingBoxHeight ECodedElement::GetBoundingBoxWidth
Min. enclosing rectangle (angle, center X and Y, height, width)	ECodedElement::MinimumEnclosingRectangle ECodedElement::MinimumEnclosingRectangleAngle ECodedElement::MinimumEnclosingRectangleCenter ECodedElement::MinimumEnclosingRectangleCenterX ECodedElement::MinimumEnclosingRectangleCenterY ECodedElement::MinimumEnclosingRectangleHeight ECodedElement::MinimumEnclosingRectangleWidth

Feret box

A feret box is a rectangle with the minimum surface rotated at a specified angle that contains all the pixels center points of an object.

- **Bounding box** is the Feret box at 0°.
- **Minimum enclosing rectangle** is the Feret box with the minimum surface across all the possible angles.
- **Width** of a **FeretBox rectangle** is the length of the rectangle side that exhibits the smallest angle with the X-axis. This is NOT necessarily the smallest side!
- The **height** of a Feret box rectangle is the length of the other side of the rectangle.

Miscellaneous

Starting point of the object contour (X and Y)	ECodedElement::GetContour ECodedElement::GetContourX ECodedElement::GetContourY
Largest run	ECodedElement::GetLargestRun
Run count	ECodedElement::GetRunCount
Object number (index)	ECodedElement::GetLayerIndex ECodedElement::GetElementIndex
Pixel gray-level value (average, deviation, variance)	ECodedElement::ComputePixelGrayAverage ECodedElement::ComputePixelGrayDeviation ECodedElement::ComputePixelGrayVariance
Pixel gray-level value (min and max)	ECodedElement::ComputePixelMax ECodedElement::ComputePixelMin

Ellipse of inertia



Eccentricity of the ellipse of inertia	<code>ECodedElement::Eccentricity</code>
Moment	<code>ECodedElement::GetCentralMoment</code>
	<code>ECodedElement::GetMoment</code>
	<code>ECodedElement::GetNormalizedCentralMoment</code>
Ellipse (angle, height, width)	<code>ECodedElement::GetEllipseAngle</code>
	<code>ECodedElement::GetEllipseHeight</code>
	<code>ECodedElement::GetEllipseWidth</code>
Second order geometric moments (Sigma: X, XX, XY, Y, YY)	<code>ECodedElement::GetSigmaX</code>
	<code>ECodedElement::GetSigmaXX</code>
	<code>ECodedElement::GetSigmaXY</code>
	<code>ECodedElement::GetSigmaY</code>
	<code>ECodedElement::GetSigmaYY</code>

Note: The object perimeter can be measured indirectly by tracing the object contour with contouring methods and counting the pixels.

From the standard geometric features, others can be derived. For instance, object elongation is computed as the ratio of large to short ellipse axis or max height over max width. Object circularity is defined as the ratio of the squared perimeter divided by four times pi multiplied by the object area.

Note: Note. Formulas (N = area):

$$\sigma_x = I_x = \frac{1}{N} \sum (y_i - \bar{y})^2$$

$$\sigma_y = I_y = \frac{1}{N} \sum (x_i - \bar{x})^2$$

$$\sigma_{xx} = I_{xx} = \frac{I_x + I_y}{2} + \sqrt{\left(\frac{I_x - I_y}{2}\right)^2 + I_x I_y + I_{xy}^2}$$

$$\sigma_{xy} = I_{xy} = \frac{1}{N} \sum (x_i - \bar{x})(y_i - \bar{y})$$

$$\sigma_{yy} = I_{yy} = \frac{I_x + I_y}{2} - \sqrt{\left(\frac{I_x - I_y}{2}\right)^2 + I_x I_y + I_{xy}^2}$$

$$\text{WIDTH} = 4 \sqrt{I_x}$$

$$\text{HEIGHT} = 4\sqrt{I_n}$$

$$\text{ANGLE} = \text{arccotan}\left(\frac{I_x - I_n}{I_{xy}}\right)$$

Convex Hull



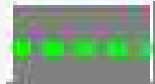




The convex hull of a shape is the convex polygon of minimum area that completely surrounds an object. The convex hull can be used to characterize the object footprint, as well as to observe concavities. Given that the number of vertices of the convex hull is variable, they are stored in a `EPathVector` container.

The corresponding function is `ECodedElement::ComputeConvexHull`.



Graphic Representation

The objects can be drawn onto the source image by means of `ECodedImage2::Draw`. The following features also have a graphical representation that can be drawn by the means of `ECodedImage2::DrawFeature`.

Objects	Graphic
Bounding box	
Convex hull	
Ellipse	
Feret box	
Feret box with an angle of 22°	
Feret box with an angle of 45°	
Feret box with an angle of 68°	

Gravity center	
Minimum enclosing rectangle	
Weighted gravity center	

Coordinate System and Conventions

Coordinate system

EasyObject uses a pixel coordinate system where the origin is conventionally at the top left corner of the top left pixel of an image. Consequently, the fractional part of the coordinates of the center of a pixel is ".5". This convention is best suited for the representation of sub-pixel coordinates.

Angles

Accordingly to the mathematical conventions, the angles are now counted inversely: A positive angle brings the X axis on the Y axis.

Evaluating the features

There is one property per feature, removing the need to access the feature through an **enum**.

Draw Coded Elements

Once an image has been encoded, the coded elements (object or hole) are accessible through the abstract class `ECodedElement` and a large set of methods:

To draw coded elements

1. **Declare a new** `ECodedImage2` object.
2. **Declare an** `EImageEncoder` object and, if applicable, select and setup the appropriate segmenter and choose the appropriate layer(s) to encode.
3. **Create an output image:** copy, pixel by pixel, the (grayscale) source image into a (color) output image if the drawing of the resulting features has to be colored.
4. **Encode the source image.**
5. **Draw the features for each object in a layer.**
6. Read the result, which can be rounded down. A specific drawing can be created to mark the feature (for example, draw a target for a gravity center).

To render flexible masks use `ECodedElement::RenderMask`.

The objects, holes and their features can be efficiently accessed randomly (in an index-based fashion).

Flexible Masks in EasyObject

A flexible mask can be generated by any application that outputs BW8 images or uses the Open eVision image processing functions.


EasyObject can use flexible masks to restrict blob analysis to complex or disconnected shaped regions of the image.

If an object of interest has the same gray level as other regions of the image, you can define "keep" and "ignore" areas using flexible masks and [Encode](#) functions.

A flexible mask is a BW8 image with the same height and width as the source image.

- A pixel value of 0 in the flexible mask masks the corresponding source image pixel so it doesn't appear in the encoded image.
- Any other pixel value in the flexible mask causes the pixel to be encoded.


EasyObject functions that create flexible masks



Source image

1) [ECodedImage2::RenderMask](#): from a layer of an encoded image


1. To encode and extract a flexible mask, first construct a coded image from the source image.
2. Choose a segmentation method (for the image above the default method `GrayscaleSingleThreshold` is suitable).
3. Select the layer(s) of the coded image that should be encoded (i.e. white and black layers using minimum residue thresholding).
4. Make the mask image the desired size using `mask.SetSize(sourceImage.GetWidth(), sourceImage.GetHeight())`.
5. Exploit the flexible mask as an argument to [ECodedImage2::RenderMask](#).



BW8 resulting image that can be used as a flexible mask

2) ECodedElement::RenderMask: from a blob or hole


1. Select the coded elements of interest.
2. Create a loop extracting a mask from selected coded elements of the coded image using `ECodedElement::RenderMask`.
3. Optionally, compute the feature value over each of these selected coded elements.



BW8 resulting image that can be used as a flexible mask


3) EObjectSelection::RenderMask: from a selection of blobs

`EObjectSelection::RenderMask` can, for example, discard small objects resulting from noise.



BW8 resulting image that can be used as a flexible mask


Example: Restrict the areas encoded by EasyObject



Find four circles (left) Flexible mask can isolate the central chip (right)

1. Declare a new `ECodedImage2` object.
2. Setup variables: first declare source image and flexible mask, then load them.
3. Declare an `EImageEncoder` object and, if applicable, select the appropriate segmenter. Setup the segmenter and choose the appropriate layer(s) to encode.


4. Encode the source image. Encoding a layer with just the area in the flexible mask is then pretty straightforward.
We see that the circles are correctly segmented in the black layer with the [grayscale single threshold segmenter](#):



5. Select all objects of the coded image.
6. Select objects of interest by filtering out objects that are too small.
7. Display the blob feature by iterating over the selected objects to display the chosen feature.

2. EasyGauge - Measuring down to Sub-Pixel

EasyGauge library controls dimensions. It accurately determines position, orientation, curvature and size of parts. It can interact graphically to place and size gauges, combine them in grouped hierarchies, and store and retrieve them with all their parameters.



Workflow

The gauge model can be built programmatically or in a graphical editor, then "played" in the final application.

Choose the workflow that matches the complexity of your model and the accuracy required: uncalibrated, calibrated or grouped.

Uncalibrated Gauging: for a simple model

EasyGauge basic use is straightforward.

1. Create a gauge object that corresponds to the required measurement.
2. Change the parameters whose default values are not appropriate.
3. Invoke the desired measurement function.
4. Read the resulting position parameters.

Uncalibrated gauging is easy to implement but has several drawbacks:

- measurements are performed in pixels, not millimeters.
- measurement models are not portable: gauge positions and sizes must be reworked if viewing conditions change.
- optical distortion or perspective causes inaccurate measurements.

Calibrated gauging: for one or two simple measurement sites

Calibrated gauging is more accurate, and measures the inspected parts independently of the viewing conditions.

All measurements are taken in the calibrated units, with any distortion implicitly compensated.

Refer to Calibration to learn how to master field-of-view calibration.

1. Create a calibrator object.
2. Place it on the inspected scene.
3. Adjust calibration parameters.
4. Attach a gauge.

Complex Gauging

Gauges can be grouped (see Gauge Manipulation Processes) and attached to another item:


- **Attaching gauges to an `EFrameShape`** object moves the gauges with the frame (translation and/or rotation), the application program must adjust the frame position to track the inspected part.
- **Attaching gauges to another gauge** moves them according to the measured position of the supporting gauge. For example, if gauges are attached to a common rectangle gauge that is detecting the outline of a part, all gauges automatically track the part when the rectangle outline is fitted.

If using several measurement sites, you can save the complete model, with calibration modes, coefficients, and attached gauges, in a single file.

Gauge definitions

Point gauge

You can select the most relevant transition points along a line segment probe that crosses one or several objects edges. Crosswise and lengthwise filtering can be activated for noise reduction.



Point location. Contrast-based selection


Rectangle Gauge

The placement of a **rectangle gauge** is defined by its nominal position (given by the coordinates of its **center**), its **nominal size** and its **rotation angle**.

Each side of a rectangle can have its own transition detection parameters, and can be set to active or inactive with the `ActiveEdges` property. When a side is active:

- setting the value of a parameter only applies to the currently active sides1.
- getting the value of a parameter yields a result only when the value of this property is the same for all active sides.
- only active sides are used for measurement and model fitting.

These rules allow to set different parameters for different sides, and measure parallel sides or a corner point instead of the whole rectangle. The four sides are denoted by letters "x", "y", "XX" and "YY" respectively.



Naming conventions for the sides of a rectangle gauge


Usage

Define and position the gauge, then use [Measure](#) to fit the lines.

To obtain the [rectangle properties](#), set [ActualShape](#) to **TRUE** to return the fitted line (**TRUE** value) (default is **FALSE**).

Alternatively, [MeasuredRectangle](#) provides the results as an [ERectangle](#) object.

For instance, you can accurately locate the four corners (landmarks) of a rectangle using a rectangle fitting gauge.



Locating a rectangle's corners

Wedge gauge


The placement of a [wedge gauge](#) is defined by its nominal position (given by the coordinates of its [center](#)), its nominal [inner](#) and [outer radius](#) ([inner](#) and [outer diameter](#)), its [breadth](#) (difference between radii), the [angular position](#) from where it extends and its [angular amplitude](#).

The `Set` member can distinguish between a full ring, a sector of a ring and a disk.

Each side of a wedge can have its own transition detection parameters and can be set to active or inactive with the `ActiveEdges` property. When a side is active, this means that:

- setting the value of a parameter only applies to the currently active sides;
- getting the value of a parameter yields a result only when the value of this parameter is the same for all active sides;
- only active sides are used for measurement and model fitting.

So different sides can have different parameters, and you can measure parallel arcs or oblique sides, or a corner point, instead of the whole wedge. The four sides are denoted by letters "a", "r", "AA" and "RR" respectively.



Naming conventions for the sides of a wedge gauge

Usage


Define and position the gauge, then use `Measure` to fit the lines.

To obtain the `wedge properties`, set the `ActualShape` property to **TRUE** to return the fitted line (instead of the nominal line position **FALSE**, default).

Alternatively, `MeasuredWedge` provides the results as an `EWedge` object.

Line gauge

The placement of a `line gauge` is defined by its `center coordinates`, its `length` and its `angle` with respect to the X-axis. To constrain the line slope value, set `Angle` and `KnownAngle`.



Line fitting

Usage


Define and position the gauge, then use [Measure](#) to fit the lines. To obtain the [line properties](#), set the [ActualShape](#) property to **TRUE** to return the fitted line (**TRUE** value) (instead of the nominal line position **FALSE** value, default).

Alternatively, [MeasuredLine](#) provides the results as an [ELine](#) object.

Circle gauge

The placement of a [circle gauge](#) is defined by its nominal position (given by the coordinates of its [center](#)), its nominal [diameter](#) (or [radius](#)), the angular position from where it extends and its angular [amplitude](#).

The Set member can distinguish between a full circle and an arc (the arc amplitude must be specified).



Circle fitting


Usage

Once the gauge has been defined and positioned, use [Measure](#) to trigger the circle fitting operation. To obtain the measurement results, set the [ActualShape](#) mode to **TRUE**. The [ActualShape](#) mode determines whether an inquiry returns the fitted circle (**TRUE** value) or the nominal circle position (**FALSE** value, default). The requested information is then retrieved by means of the [circle properties](#).

Alternatively, [MeasuredCircle](#) provides the results as an [ECircle](#) object.


Find transition points using peak analysis

Finds the position of all transition points along a line segment probe that crosses one or several objects edges, and allows selecting the most relevant ones. Crosswise and lengthwise filtering can be activated for noise reduction.



Point location. Contrast-based selection

Point Location principle



Point location principle (left) and S-shaped curve and its derivative (right)

On a linear profile extracted from an image, an edge appears as a transition from dark to light (or vice versa). When plotting pixel values along the gauge, this transition appears as an S-shaped curve. The first derivative of this curve exhibits a peak around the transition point. The better the contrast, the sharper the transition and the higher the peak.

EasyGauge extracts the pixel values along a profile (red curve) then uses peak analysis to determine the transition location. All the pixel values in the peak [area](#)¹ are used to compute the transition location.

- Sub-pixel accuracy is only possible if the transition is surrounded by almost uniform regions of at least 2 pixels wide.
- [BWB](#)² transitions have an increasing profile curve and the peak takes positive values. Otherwise, the curve decreases and the peak extends negatively.
- You cannot normally detect peaks using the default threshold value (20) as BWB or WBW transitions base the peak analysis on the gray level profile along the EPointGauge (or sample path) and not its first derivative.

[EPointGauge](#) contains all point measurement parameters, with default values that detect reasonably contrasted edges.

EPointGauge parameters

[Center](#): Nominal point position (will normally be different before and after measurement).

[Tolerance](#): Tolerance value and gauge orientations.

[TransitionType](#), [TransitionChoice](#), [TransitionIndex](#): Peak selection strategies.

¹Area between the derivative curve and a horizontal user-defined threshold level

²Black / White / Black

Threshold: Noise immunity.

MinAmplitude, MinArea: Peak strength.

Thickness, Smoothing: Local filter widths.

RectangularSamplingArea Sets sampling area (rectangular by default) to transverse filtering mode.

Measure: Measures the object.

- In single transition mode, **Valid** returns True when an appropriate point was found. To obtain measurement results, set **ActualShape** to True so that **Center** returns the located point. (False default value returns nominal point position).

- In multiple transition mode, **NumMeasuredPoints** returns the number of points found, **GetMeasuredPoint** returns an **EPoint** object which contains located point information. An integer index between 0 and **GetNumMeasuredPoints**-1 must be passed.


GetMeasuredPeak: Returns **EPeak** containing the peak's **Area** and **Amplitude**, and the delimiting coordinates along the probe segment (**Start**, **Length** and **Center** values).

Select Peaks to improve edge precision

The threshold level is very important:

- Too high can cause significant peaks to be missed, and insufficient pixel values to achieve good precision.
- Too low can cause false peaks because of noise.


To resolve this dilemma, the EasyGauge peak selection mechanism can reject low contrast or false edges: transition strength is measured by peak *amplitude* and *area*. Every edge measurement determines peak amplitude and area. If either value falls below the **minimum amplitude** or **minimum area**, the peak is disregarded and no point is assumed at that location.



Threshold level selection (left) and Peak amplitude and area (right)


Multiple versus single transition

EasyGauge can measure several edge points in a single go and retrieve all results afterwards while in *multiple transition mode*.



Multiple transition (left) versus single transition (right)


You can select the single most relevant transition based on 4 criteria: the highest peak, the peak with the largest area, the peak closest to the gauge center, or the N-th peak encountered starting from one tip of the gauge.



Best area (first image) and best amplitude choices (2nd image), closest (3rd image) and 3rd from the start (4th image)

Positive or negative peak selection


Peak selection can also be refined by choosing the transition polarity: White to Black or Black to White (i.e. positive or negative peak), or indifferent.



Black to white, white to black or indifferent polarities

Pre-filtering

Pre-filtering the image locally can reduce noise effects. Transverse (lengthwise) filtering averages several parallel lines when sampling the image. Longitudinal (crosswise) uniform filtering can also be applied to the resulting profile curve.




Thick point gauge for filtering

Transverse Filtering

Transverse filtering places parallel line segments in either a parallelogram or a rectangle (default). This behavior can be toggled. Parallelogram mode is faster than rectangular if the angle is close to 0° or 90°, or thickness is less than 5. If thickness=1, no difference exists between the two modes. **thickness** determines the number of parallel lines.

sampling area is the smallest region containing all the parallel line segments.



Rectangular sampling area (left) and Parallelogram sampling area (right)

Point Probe Position

The expected **nominal** position of a point gauge is specified by its **center**, orientation **angle** with respect to the X-axis, and length **tolerance** that the point position can vary.

The results are the coordinates of the located points (the **actual** location) and the strength of the transition (amplitude and area).

Low values indicate a weak edge, possibly corresponding to an unreliable or inaccurate measurement.

Tuning Point Measurement Parameters for unclear edges

The EasyGauge default parameters and working modes are good for clear edges. More complex situations may need parameter tuning.

1. Set the gauge point location and tolerance.

The center position and orientation are easy to decide based on a sample image or on coordinate considerations. The tolerance depends on the edge position variations. A larger tolerance increases the likelihood of hitting an edge, but it may be a false edge or extraneous feature.






2. Decide whether noise reduction is required. Lay the gauge over the desired location and observe the profile curve and its derivative (play with the filtering parameters while looking at the plotted curve). The curve regularity gives an indication of the spread of the gray-level values.

When these coefficients are set, the gray-level profile will not change anymore.

3. Set the threshold value to be low enough for useful parts of the peaks to cover enough pixels (to achieve better sub-pixel accuracy), but not lower than the ambient image noise.

4. Remove weak or false edges using the list of peak amplitudes and areas. Plotting these values along with good and extraneous peaks can help find appropriate peak rejection limits.

5. Choose whether all transition points are needed or just the most relevant. If all are required, they can be queried one after another. Otherwise, a point selection strategy should be chosen based on strength, order or transition polarity (black to white and/or conversely).



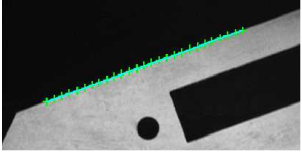
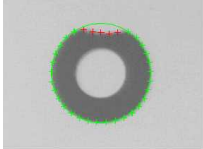
Find shapes using geometric models

[ELineGauge](#), [ECircleGauge](#), [ERectangleGauge](#), or [EWedgeGauge](#) predefined geometric models can be fit over the edges of an object. The targeted edge must be defined, and points sampled along it at regularly spaced point measurement gauges. Model fitting in the least square sense can be applied.

<p>Line: Measures position and orientation of straight edges.</p>	
<p>Circle: Measures position and curvature of a circle or arc.</p>	
<p>Rectangle: Measures position, orientation and size of a rectangle.</p>	
<p>Wedge: Measures position, orientation and size of a ring/ disk sector / curvilinear rectangle.</p>	

All gauge types share these common features:

<p>Point sampling</p> <p>Point gauges are placed along the edges and point measurement carried out at regularly spaced spots, which can be adjusted differently per side in rectangle and wedge gauges. All point measurement parameters and operating modes are available.</p> <p>SamplingStep sets the spacing of point location gauges along the model.</p>	<p>Sampling paths and sampled points</p>
---	--

<p><code>NumSamples</code> returns the number of points sampled during the model fitting operation.</p>	
<p>Model fitting</p> <p>The model is adjusted to minimize error residue and provide the best edge parameter estimates. Rectangles and wedges have parallelism and concentricity constraints. Image shows sampled points and fitted line.</p>	
<p>Outlier rejection</p> <p>After model fitting, some points will be too far away from the fitted model and may harm location accuracy. EasyGauge can tag them as outliers to be ignored using the <code>FilteringThreshold</code> property.</p> <p>The outlier elimination process can be repeated several times using <code>NumFilteringPasses</code>.</p> <p>The number of valid sample points remaining after a model fitting operation is kept in <code>NumValidSamples</code>.</p> <p>The average distance of these points to the fitted model is returned by <code>AverageDistance</code>.</p>	

Gauge Manipulation: Draw, Drag, Plot, Group

EasyGauge provides means to graphically interact with gauges to place and size them, combine them as a hierarchy of grouped items, and store/retrieve them and all working parameters to/from model files.

Draw

`Draw` gives a graphical representation of a gauge. Drawing is done with the current pen in the device context associated to the desired window. Depending on the operation, handles may be displayed.

Drag

An operator can drag a gauge interactively over an image. Several dragging handles are available.

- `HitTest` determines when the mouse cursor is over a handle. When it is, the cursor shape should be changed for feedback, and a drag can take place.
- `Drag` moves the handle and the corresponding gauge accordingly.

Plot

EasyGauge can `Plot` gray-level values along the sampled paths and/or its derivative - useful for parameter tuning.

Point measurement gauges can plot after calling `Measure`.

Model fitting gauges can plot after calling `MeasureSample` with an index argument that lies between **0** and **GetNumSamples-1** (included).

To view the corresponding sampling path, use method `Draw` with mode **EDrawingMode_SampledPath**.


Group

Measurement gauges can be grouped (their relative placement remains fixed) to form a dedicated tool that can be moved (translated and rotated) to follow the movement of inspected items / probes before computing measurements.

`Attach` associates a gauge to a mother gauge or `EFrameShape` object.

`NumDaughters`, `GetDaughter`, or `Mother` retrieves information relative to attached daughters or mother.


`Detach`, `DetachDaughters` dissociates the gauge or daughters from the mother.



Calibration and Transformation


Field-of-view calibration

Calibration establishes the relationship between real-world point coordinates and image pixels. A simple calibration model computes faster, a repeatable part position is easier to locate.




The Raw sensor coordinate system starts from upper left and extends rightwards and downwards.

The range of abscissae is **0** to **width-1** and the range of ordinates is **0** to **height-1** where integer coordinate values correspond to pixel centers.



The Centered sensor coordinate system starts at the center ($[(width-1)/2, (height-1)/2]$ in the Raw system) and extends rightwards and upwards.



The real world 3D coordinates are defined in a 2D reference frame tied to a reference plane. The origin and direction of the axis are normally aligned with major features of the inspected parts.



Before World-to-Sensor Transform










Before converting from world to sensor coordinates, sources of distortion should be eliminated:

- adjust sweep frequency or scanning speed to avoid non-square pixels.
- adjust optical alignment to minimize perspective effect. The field of view should be parallel to the sensor plane.

- use long focal distances and good quality lenses to minimize Optical distortion.
- use appropriate scale factor based on lens magnification, observation distance and focusing.
- minimize skew and translation effects by secure fixtures, and part-movement / acquisition-triggering synchronization.

Effects of World-to-Sensor Transform

- **No calibration.** World and sensor coordinates are identical.
- **Translated calibration:** The coordinate origin can be moved. World coordinates correspond to pixel units.
- **Isotropic scaling** (square pixels). A scale factor converts pixel values to physical measurements.
- **Anisotropic scaling** (non-square pixels). Uses two scale factors with pixel aspect ratio (X / Y) in the range $[-4/3, -3/4]$ (or $[3/4, 4/3]$). Pixels are always displayed as square, so the image appears stretched.
- **Scaled and skewed** (square pixels). Real-world axis aligns with rotated inspected part using translation, rotation and scaling.
- **Scaled and skewed** (non-square pixels). Distortion is apparent. Occurs when camera scan speed does not match pixel spacing.
- **Perspective distortion** causes further away objects to look smaller; lines remain straight but angles are not preserved.
- **Optical distortion** causes cushion or barrel appearance of rectangles.
- **Combined distortions** result in a complex, non linear, transform from real-world to sensor spaces.



Calibration using EWorldShape

The EWorldShape object can calibrate the whole field of view (in given imaging conditions with fixed camera placement and lens magnification), if the optical setup is modified.

EWorldShape computes appropriate calibration coefficients and transforms measurement gauges that are tied to it.

It can set world-to-sensor transform parameters, perform conversions from and to either coordinate system, determine unknown calibration parameters, and save the parameters of a given transform for later reuse.

After calibration EWorldShape can perform coordinate transform for arbitrary points using SensorToWorld and WorldToSensor to:

- measure non-square pixels and rotated coordinate axis.
- correct perspective and optical distortion, with no performance loss.

There are several ways to obtain the calibration coefficients:


Estimate (feasible if no distortion correction is required and accuracy requirements are low)

To estimate the calibration coefficients either locate the limits of the field of view and divide the image resolution by the field of view size, or use the following procedure:

1. Take a picture of the part to be inspected or a calibration target (e.g. rectangle).
2. Locate feature points such as corners in the image (by the eye) and determine their coordinates in pixel units —let (i, j) .
3. Use the Euclidean distance formula to derive the calibration coefficient: $c = \frac{\sqrt{(i_1 - i_0)^2 + (j_1 - j_0)^2}}{D}$ where c is a calibration coefficient, in pixels per unit, and D is the world distance between the corresponding points, in units.
4. For non-square pixels repeat this operation for pairs of horizontal and vertical points.

To estimate a skew angle, apply this formula to two points on the X-axis in the world system:

$$\theta = \arctan \frac{j_1 - j_0}{i_1 - i_0}$$



Estimating scale factors and skew angle

When the calibration coefficients are available, use SetSensor to adjust them and set the calibration mode, or set them individually using: SetSensorSize, SetFieldSize, SetResolution, SetCenter, SetAngle.

Pass a set of reference points (landmarks) to a calibration function

Locate at least 4 landmarks and obtain their coordinates in sensor (using image processing) and world coordinate systems (actual measurements). More landmarks give more accurate calibration.

The resulting pixels aspect ratio (X resolution/Y resolution) must be in the range $[-4/3, -3/4]$ (or $[3/4, 4/3]$).


Analyze a Calibration target

A calibration target can be automatically analyzed to get an appropriate set of landmarks. It is an easy way to achieve automatic calibration, provided an appropriate procedure is available to extract the desired landmark point coordinates.

Open eVision relies on the use of a specific target holding a rectangular grid of symmetrical dots (of any shape) with no other object on the grid.

Dot Grid based calibration example


1. Grab an image of the calibration target in such a way that it covers the whole field of view (or restrict the image of view to an ROI where only dots are visible).
2. Apply blob analysis to extract the coordinates of the centers of the dots, as can be done by [EasyObject](#).
3. Pass all points detected to [AddPoint](#) (sensor coordinates only).
4. Call [RebuildGrid](#) to reconstruct a grid to calibrate a field of view using an iterative algorithm which computes the world coordinates of each dot.
 - a. The grid points nearest to the gravity center (g) of grid points are selected (g_1 and g_2) to form the first reference oriented segment, of length A .



- b. Starting from the extremity of the reference segment (g_2), the algorithm determines 3 tolerance areas (white squares in the figure), in perpendicular directions. The tolerance areas are centered at a distance A (length of the reference segment) from (g_2). They are square, with a side-length of A .

The algorithm searches for 1 neighboring point, in each of the 3 tolerance areas.


The grid will be correctly calibrated if each tolerance area contains a neighboring point.



- c. The 3 perpendicular segments are the references of the next iterative searches. The algorithm goes back to step 2.
5. Call `Calibrate` the landmark approach.

If the grid exhibits too much distortion, grid reconstruction does not work as expected. The following errors could happen:

1. A tolerance area does not contain a neighboring point (red square in the figure).
2. A tolerance area contains more than one neighboring point.
3. The point in the tolerance area is not the correct one. For instance, the point might be diagonally connected (red point in the figure).



Advanced Features

The field-of-view calibration model can be tuned using these parameters:

Sensor Width and Height

The `sensor width` and `sensor height` give the logical image size, in pixels (always integers).

Field-of-View Width and Height

The `field-of-view (f-o-v) width` and `height` give the actual image size, in length units, i.e. the size of the rectangle corresponding to the image edges in the world space. These values are related to the pixel resolution by the following equations:

$$\begin{aligned} \text{f-o-v width} &= \text{pixel width} * \text{sensor width} \\ \text{f-o-v height} &= \text{pixel height} * \text{sensor height} \end{aligned}$$


or

$$\begin{aligned} \text{sensor width} &= \text{f-o-v width} * \text{horizontal resolution} \\ \text{sensor height} &= \text{f-o-v height} * \text{vertical resolution} \end{aligned}$$

By default pixel height is not specified, the pixels are assumed to be square (pixel width = pixel height).

Scale

Ratio




Anisotropic aspect ratio

Center Abscissa and Ordinate

The **center abscissa** (x) and **ordinate** (y) indicate the image origin point (world coordinates $(0,0)$). Default is the image center.

Skew Angle

The **skew angle** is the angle formed by the real-world reference frame (X-axis) and the image edge (horizontal). The default is no skew.




Skew angle

Note: When the pixels are not square, the *EWorldShape* object can convert the angle between the world and sensor spaces.

X and Y Tilt Angles


The **X** and **Y** tilt angles describe the viewing plane direction. They correspond to the required rotations around X and Y axis that bring the Z axis parallel to the optical axis.



Tilt X and tilt Y angles

Perspective Strength


The [perspective strength](#) gives a relative measure of the perspective effect. The shorter the focal length, the larger the value.



Weak and strong perspective

Distortion Strength


[Distortion strength](#) and [GetDistortionStrength2](#) give a relative measure of radial distortion in the image corners, i.e. the ratio of image diagonal length with and without distortion.



Positive and negative distortion

[Calibration mode](#), expressed as a combination of options, can be accessed via [CalibrationModes](#).

Effect of the Calibration Coefficients



No calibration coefficient: All coefficients combined.

Unwarp an image


An **EWorldShape** object manages a field-of-view calibration context. Such an object is able to represent the relationship between world coordinates (physical units) and sensor coordinates (pixels), and account for the distortions inherent in the image formation process.

Image calibration is an important process in quantitative measurement applications. It establishes the relation between the location of points in an image (pixel indices) and the actual positions of those points in the real world, on the inspected item.

Calibration can be setup by providing explicit calibration parameters of the calibration model, or a set of known points (landmarks), or a calibration target.


The goal of calibration is twofold:

- To gain independence with respect to the viewing conditions (part placement in the field of view, lens magnification, sensor resolution, ...), letting you describe the inspected item once for all using absolute measurements.




Single model versus multiple viewing conditions

- To correct some distortion related to the imaging process (perspective effect, optical aberrations, ...).



Removal of image distortion

The pixel indices in an image are usually integer numbers, but fractional values can occur when using sub-pixel methods. They are normally obtained by processing an image and locating known feature points. These values are called sensor coordinates.




Feature point in sensor space

The world coordinates describe the location of points on the inspected item are expressed in an appropriate length measurement unit.

The world coordinates are actual dimensions, usually gathered from design drawings or by mechanical measurements.

They require a reference frame to be defined.




Reference frame in world space

Unwarp

Unwarp an image using `Unwarp`, `SetupUnwarp` and `UnwarpAfterSetup`.

Using a lookup table before unwarping may speed up the process.




Distorted vs. Unwarped image

3. EasyMatch - Matching Area Patterns

EasyMatch learns a pattern and finds exact matches:


1. The pattern is learned by defining an ROI that contains the object to be matched. This ROI is created after iteratively learning from several images which contain the object.
2. The parameters are tuned to ensure the pattern is found reliably.
3. Images can now be searched for one or more occurrences of the pattern, which may be translated, rotated or scaled.




Learning and Matching a pattern

Workflow

Learning workflow



Matching workflow



Learning Process

Select an image containing the pattern/ROI to be searched for and call `LearnPattern`.

The resulting pattern can be saved as a model for later use. You can repeat this process to search for and save multiple patterns.

Best pattern characteristics

- **repeatable**, you need to know if it can translate or rotate or scale.
- **represent the object to be located.**

It should:


- Keep the same appearance whatever the lighting conditions.

- Remain at a fixed location with respect to the part.
- Be rigid and not change shape.
- **exhibit good contrast in small and large scale.** It should be distinctly visible from a distance, and on a reduced image.
- **not be invariant under the degrees of freedom to be measured.** For instance, a pattern of black and white horizontal stripes cannot detect horizontal translation; a cog wheel cannot help measure large rotations.
- **have a neutral background.** If objects around the pattern in the ROI may change, this area should be neutralized by means of "don't care" pixels or a mask.
- **have contrasted margin around the objects** so that foreground and background intensities are seen.

Customize Parameters

Parameters can be tuned to minimize processing time, but it still takes longer than EasyFind as the entire selected area is matched.

- **DontCareThreshold:** If don't care areas are required, the corresponding pixels must hold a value below the `DontCareThreshold`.
If all the background can be ignored, merely adjusting the `DontCareThreshold` to the right thresholding value can do.
Otherwise, when the don't care area is unrelated to the threshold pattern image, the `DontCareThreshold` should be set to 1 and all pixels belonging to the don't care area should be set to black (value 0).
- **MinReducedArea:** To improve time performance, EasyMatch sub-samples the pattern. This parameter stipulates the minimum size of the pattern (as its area in pixels) during sub-sampling. The smaller the value, the faster the matching process, but, set too low, it decreases the matching process reliability.
The value of `MinReducedArea` is computed automatically if `AdvancedLearning` is enabled (default behavior). Setting explicitly `MinReducedArea` will disable `AdvancedLearning`. A value of 64 is usually a good compromise.
- **AdvancedLearning:** If the pattern is defined as a ROI of an image, `AdvancedLearning` optimizes learning parameters, such as `MinReducedArea`, by using the whole image context. `AdvancedLearning` is enabled by default, as it leads to better results in case of tiled or periodic images. If `MinReducedArea` is set explicitly, `AdvancedLearning` is disabled. Please note that as `AdvancedLearning` changes the number of pixels in the pattern, it can have a significant impact on the matching process duration.
- **FilteringMode:** If the image has sharp gray-level transitions, it is better to choose a low-pass kernel instead of the usual uniform kernel.



Learning a pattern

Matching Process

For each new image, one or more occurrences of the pattern is searched for, allowing it to translate, rotate or scale, using a single function call:

- `Match`: receives the target image/ROI as its argument and locates the desired occurrences of the pattern.


You can set these parameters:

- Rotation range: `MinAngle`, `MaxAngle`.
- Scaling range: `MinScale`, `MaxScale`.
- Anisotropic scaling range: `MinScaleX`, `MaxScaleX`, `MinScaleY`, `MaxScaleY`.

The following functions return the result of the matching:

- `NumPositions` returns the number of good matches found. A good match is defined as having a score higher than prescribed value (the `MinScore` threshold value).
- `GetPosition` returns the coordinates of the N-th good match. The positions are sorted by decreasing score.

If you want to match several patterns against the same image, create an `EMatcher` object for each pattern.



Matching a pattern

Advanced Features

The best way to speed up this process is to minimize rotation and scaling, and limit the number of occurrences searched for.


- Learning time:
 - Optimize number of searches: Searching all positions takes too long, so a sequence of searches is performed at various scales (reductions). The coarsest reduction is quick and approximate. Subsequent reductions work in a close neighborhood to improve location, drastically reducing the number of positions to be tried. The location accuracy is given by 2^K , where K is the reduction number.
 - `MinReducedArea`. Indicates how small the pattern can be made for rough location.
- Matching time:
 - Correlation mode (way to compare the pattern and the image): `CorrelationMode`. **Can be standard, offset-normalized, gain-normalized and fully normalized:** the correlation is computed on continuous tone values. Normalization copes with variable light conditions, automatically adjusting the contrast and/or intensity of the pattern before comparison.
 - Contrast mode (way to deal with contrast inversions): `ContrastMode`. Lighting effects can cause an object to appear with inverted contrast, you can choose whether to keep inverted instances or not, and whether to match positive occurrences only, negative occurrences only or both.
 - Maximum positions (expected number of matches): `MaxPositions`, `MaxInitialPositions`. You can compel EasyMatch to consider more instances than needed at the coarse stage using the `MaxInitialPositions` parameter (this number is progressively reduced to reach `MaxPositions` in the final stage).
 - Minimum score (under which match is considered as false and is discarded): `MinScore`, `InitialMinScore`.
 - Sub-pixel accuracy: `Interpolate`. The accuracy with which the pattern is measured can be chosen (the less accurate, the faster). By default, the position parameters for each degree of freedom are computed with a precision of a pixel. Lower precision can be enforced. One tenth-of-a-pixel accuracy can be achieved.
 - Number of reduction steps: `FinalReduction`. Can speed up matching when coarse location is sufficient, range `[0...NumReductions-1]`.
 - Non-square pixels: `GetPixelDimensions`, `SetPixelDimensions`. When images are acquired with non-square pixels, rotated objects appear skewed. Taking the pixel aspect ratio into account can compensate for this effect.
 - "Don't care" pixels (ignored for correlation score) below the `DontCareThreshold` value. When the pattern is inscribed in a rectangular ROI, some parts of the ROI can be ignored by setting the pixels values below a threshold level. The same feature can be used if parts of the template change from sample to sample.

4. EasyFind - Matching Geometric Patterns


4.1. Workflow


EasyFind learns a reference model from a pattern, which is used to find similar patterns in other images and retrieve information about these instances.

It is quick and robust, and very tolerant of noise, blur, occlusion, missing parts and changes in illumination.



Workflow






Feature points definition

A feature point is a pair of coordinates (X, Y) and a type (Edge, Transition or Region).

EasyFind uses feature points to find instances in a search field.




- **Edge feature points:** an abrupt change of gray level between two regions indicates an edge at this location in the search fields.
- **Transition feature points:** a smooth change of gray level between two regions indicates a transition area in their neighborhood (represented by dots in the blue area of the example above). The size of the neighborhood can be modified.
- **Region feature points:** identify 2 regions of roughly uniform gray levels:
 - dark region (represented by a family of dots in the red area of the example above),
 - bright region (represented by a family of dots in the green area).

4.2. Learning Process

EasyFind supports various pattern types (Consistent edges, Thin structures, or Contrasting regions).

During the learning process, EasyFind computes for itself a feature model which is a set of all extracted feature points from a bitmap representation of the pattern.

EasyFind only needs this feature model to start the finding function, but you can create your own optimal model.



The optimal model depends on the type of pattern being searched for.


Consistent Edges

Models must be well contrasted with sharp edges. They should be substantially different from the rest of the expected search fields. Can be scaled or rotated, very robust to: blurring, noise, occlusion, illumination variation (point-by-point scores improves robustness and computation time of the finding phase).

Good for models with consistent edges, sharp contrast transitions, regions delineated by well defined edges that are in approximately the same place for each instance in all search fields.


Choose the similar points with `EPatternFinder::ContrastMode` property:

- **PointByPointNormal**: if points share the same contrast polarity.
- **PointbyPointInverse**: if points exhibit opposite contrast polarity.
- **PointByPointAny**: regardless of their respective contrast polarity.




Contrasting Regions (defined by region feature points and transition feature points)

Ideally, models should be 50% bright area and 50% dark area. There can be more than one dark and/or bright region.



Cannot be scaled or rotated, robust to: blurring, noise, illumination variation (but not occlusion).


Good for models with inconsistent transitions or edges, or with several regions delimited by transitions or edges.



Thin Structures (defined by edge feature points)

Can be scaled or rotated, robust to: blurring, noise, occlusion, illumination variation. Edges must be consistent between thin elements and regions, and the contrast should be the same for each thin element.


Good for models containing thin elements.



Check the Learned model is correct

EasyFind can draw the extracted feature points on the model using the `DrawModel` method of the `PatternFinder` object.

On the examples, edge feature points appear as green points for Consistent Edges and Thin Structures, and crosses for Contrasting Regions.




4.3. Finding Process

You can optimize the finding process by setting and saving some parameters in a configuration file. You simply load this file and run to see what EasyFind has found in the reported information.


Maximum number of expected instances

Set the maximum number of instances that EasyFind should return. In this example the number was three.



Angles and scales (thin structure and consistent edges pattern types)


Ranges have a bias and a tolerance. For instance, for an angle bias of 20° and an angle tolerance of 5° , EasyFind returns instances with an angle between 15° and 25° with respect to the learned model ($20^\circ \pm 5^\circ$).



4.4. Advanced Features

Find partial patterns


EasyFind can locate instances of thin structures and consistent edges that are partially out of the search field, if the extension of search field is set to > 0 pixels.



Tune Parameters

These parameters can be tuned for all models:

- Adjust the **Light Balance** using model drawing to preview the model so that it fits the useful parts of the pattern, then learn the model again.




Bad extract (left) - Adjust light balance (center) - Good extract (right)

- Set the gray-level threshold (this overrides light balancing) and learn the model again.
- Move the **pivot** to a specific place in the model like a corner or a hole. The pivot is the location returned by EasyFind when it finds an instance (the center by default).




- **Thin Structures** may benefit from tuning these parameters:
 - Automatic (thin elements with the same contrast between them and their neighboring regions)
 - Thin elements darker than the neighborhood
 - Thin elements brighter than the neighborhood
- **Contrasting regions** may benefit from tuning these parameters:

- **Transition Thickness:** Transition feature points lie in the *transition band* (blue area in the examples below) which is where the transition occurs. The transition thickness parameter should allow the borders of different instances to stay within the transition band, so it is recommended to make it equal to the biggest variation among instances. Two examples of thickness, and the recommended thickness are below:



- **Ignored areas.** Zero values indicate ignored areas. 255 values indicate areas taken into account. For example: If the text in the center of the model differs from the instance, you can indicate that EasyFind must not extract feature points from this part of the model.



5. Golden Template Validation (EChecker)

EChecker requires two processing steps, Training and Inspection. These two operations are totally independent and can be programmed in separate applications.

Training involves pre-processing the set of reference images to compute the acceptance ranges at each pixel and to store them in two threshold images for use with EasyObject. You can test that the Golden Template is good and tune the learning process parameters if necessary. The training can be done once for all and the results archived in a Golden template file for later use.

Inspection involves processing an image, realigning and normalizing, detecting pixels that fall out of the acceptance intervals, checking the quality, and then tuning the inspection parameters if necessary.

The following sections present the relevant API functions for use in the training and inspection steps.

Training

Reference images are preprocessed to compute the pixel acceptance ranges, and store them in two threshold images for use with EasyObject (legacy). The training results are archived in a model file for later use.

Two modes of operation are provided, depending on whether reference images are stored on disk or are acquired and processed on the fly.

These operations cannot be used during on-line operation.

To define a model, several operations are performed, in this order:

1. On the first reference image, one or two ROIs are placed to define the location patterns (fiducial marks or landmarks).

These ROIs are surrounded by two others to define the possible movement freedom, and are used as search areas for pattern matching. ROIs can be placed interactively using dragging handles.


EChecker manages the dragging operations.

- **Draw** graphically renders the ROIs and handles.
- **HitTest** detects the presence of the cursor over one of the handles.
- **Drag** moves a handle.

These functions operate on all three ROIs (pattern, search and inspected). You will quickly notice that:


- Dragging a pattern ROI causes the corresponding search ROI to adjust automatically so tolerances remain constant.

- Dragging a search ROI causes its size to adjust symmetrically with respect to the pattern.
- Adjusting a search area also sets the inspected ROI to the largest available space in the image.




Pattern ROIs and search ROIs

- If the training images are on disk, the list of file names can be registered and used later to execute learning in a single command (*batch learning*, as opposed to *on-the-fly learning*).
- Another ROI is defined to delimit the area to be inspected. This area must only include pixels of the rigid part (that moves with the fiducial marks), and not the background.



Inspected ROI

- All images are processed and are averaged using Statistical training. It uses **realignment** to deal with displacement of the inspected part in the field of view, and gray-level normalization to deal with global illumination changes.
- Ideally at least 16 images should be used in **learning passes** to create the low and high threshold images. The user can strengthen or loosen the acceptance intervals using a global tolerance parameter. Call **Learn** (ELearningMode_Ready) . Property **RelativeTolerance** adjusts the acceptance ranges.
- The model can be saved to a single file including all relevant information, i.e. placement of the various ROIs, fiducial pattern images, gray-level normalization parameters and the threshold images.



Components of an EChecker model

Check the reference images for reliability.

After the alignment ROIs have been set, they should be checked for reliability of fiducial location. The best way is to load the training images, display them and locate the patterns in them, by means of member function [Register](#). If location fails, different corrective actions can be taken, depending on the problem:

- The choice of pattern is poor. Define other ROIs with more stable contents.
- The search areas are too tight, so that occurrences of the pattern are found along edges. Enlarge the search area.
- The image is insufficiently representative (it has defects). It is better to withdraw it from the learning set.

Note: *The first call to member [Register](#) invokes the [Learn](#) members of the alignment [EasyMatch](#) contexts, i.e. training on the patterns is achieved. Unless member [Learn](#) ([ELearningMode_Reset](#)) is called, these patterns will be used for all subsequent alignment operations. The first image to be used serves both as a reference for defining the alignment pattern and contrast measurement. It is called the mother image.*

If the learning images are saved on disk: Every time a file is successfully loaded and is accepted as a reference image, the [EChecker](#) object can add the file path name to a list. Later, all files can be processed in a single command.

More learning

After the ROI placement and pattern learning have been performed (`Register` operations), training still requires two passes:

- The **"average" pass** is needed to compute an ideal, noise-free, image that reveals the central tendency of the part image.
For each image, realign and normalize (`Register`). If the operation is successful (good pattern location), call `Learn` (`ELearningMode_Average`) for immediate processing (on-the-fly learning), or `AddPathName` for deferred processing (batch learning).
- The **"deviation" pass** measures variations around the average image.
For each image, realign and normalize (`Register`). If the operation was successful, call `Learn` (`ELearningMode_RmsDeviation`) (enhances the large deviations), or `Learn` (`ELearningMode_AbsDeviation`) for immediate processing (more robust and is recommended in most circumstances).

In principle, the images shown during those two passes should be the same.

If you do not want to archive them, two distinct sets of images can be used (on-the-fly learning). These sets need not even be of the same size.

A learning set size of at least 16 images is recommended.

Alternatively, calling `BatchLearn` will perform the two required passes for all images added to the file list.

Adjust thresholding


`Learn` (`ELearningMode_Ready`) . Property `RelativeTolerance` can be adjusted to adjust the acceptance ranges.

Inspection

Inspection is straightforward: the sample image is realigned with the model file and gray-level normalized.

Note: *The inspected ROI must be positioned on the mother image. It can be interactively positioned in the same way the pattern and search areas are. This ROI can be set at the same time as the others. Changing the search tolerance will reset the inspected ROI to the largest available area.*

`EChecker` can tune the global `RelativeTolerance` property at inspection time, thus changing lower and upper threshold images used by `EasyObject` (legacy). This is the only `EChecker` parameter that can be changed after learning.




The resulting image is passed to an `ECodedImage` object for blob analysis; which groups neighboring pixels to form blobs, discards the smaller blobs (usually noise), measures geometric characteristics (location, size, orientation, ...) and others, see `EasyObject`.

This image is then compared to the lower and upper reference images (a comparison is made pixel by pixel between the sample and the template to detect pixels that fall out of acceptance intervals). Defective pixels are handled by standard `EasyObject` functions.

Compare using `EChecker`

`EChecker` finds visible differences between templates and samples, and reports them in a defect map which highlights significant differences. `EasyObject` blob analysis tools can then locate the defects and qualify them in terms of extent, orientation, lightness and so on. It is ideally suited when inspected items are rigid (shape does not vary) and illumination is uniform, ensuring repeatable visual appearance of the image, so that point to point image comparison makes sense.





Reference image vs. sample image with defects

Statistical training

Grabbing several reference images optimizes assessment of normal gray-level variations and acceptance intervals:

- consecutive images of the same part without any change (static test) gives a gray-level distribution that corresponds to noise distribution.
- consecutive images of different defect-free parts reveals variations due to the parts.



Accepted gray-level ranges

5.1. Image Comparison

The best way to compare images is using [EChecker](#) to combine a set of images and determine the range of acceptable values at each pixel.

Images can also be compared by subtracting 2 images (using Arithmetic and Logic functions) to get an absolute value, then thresholding this difference to highlight the non-zero pixels where the images differ.

EChecker performs:

alignment

Pattern matching measures movement of the inspected part in terms of translation, rotation and/or scaling (as described in the [EasyMatch](#) chapter).


When the patterns have been located, the image is realigned so that the inspected part is brought to the same position as in the reference images.

A single matching pattern easily handles translation.

Two matching patterns provide better accuracy for rotation measurement.

When fiducial marks are available, they can be used as landmarks for accurate and repeatable location.

[EChecker](#) hosts one or two [matching contexts](#) internally.



Realignment using fiducial marks

The alignment method has three shortcomings:

1) Unavoidable variations

Tolerance must be introduced as noise can make identical images have slight variations in color or shape.



Comparing two noisy images for strict equality.
Black pixels of the result image are non equal pixels.

2) Placement of inspected parts is rarely repeatable

Slight misplacement means the compared pixels no longer correspond, the resulting effect is especially noticeable around edges.



Comparing misaligned images

3) Lighting cannot be separated from parts



Comparing images in different lighting conditions

PART V
TEXT IDENTIFICATION TOOLS

1. EasyOCR - Reading Texts

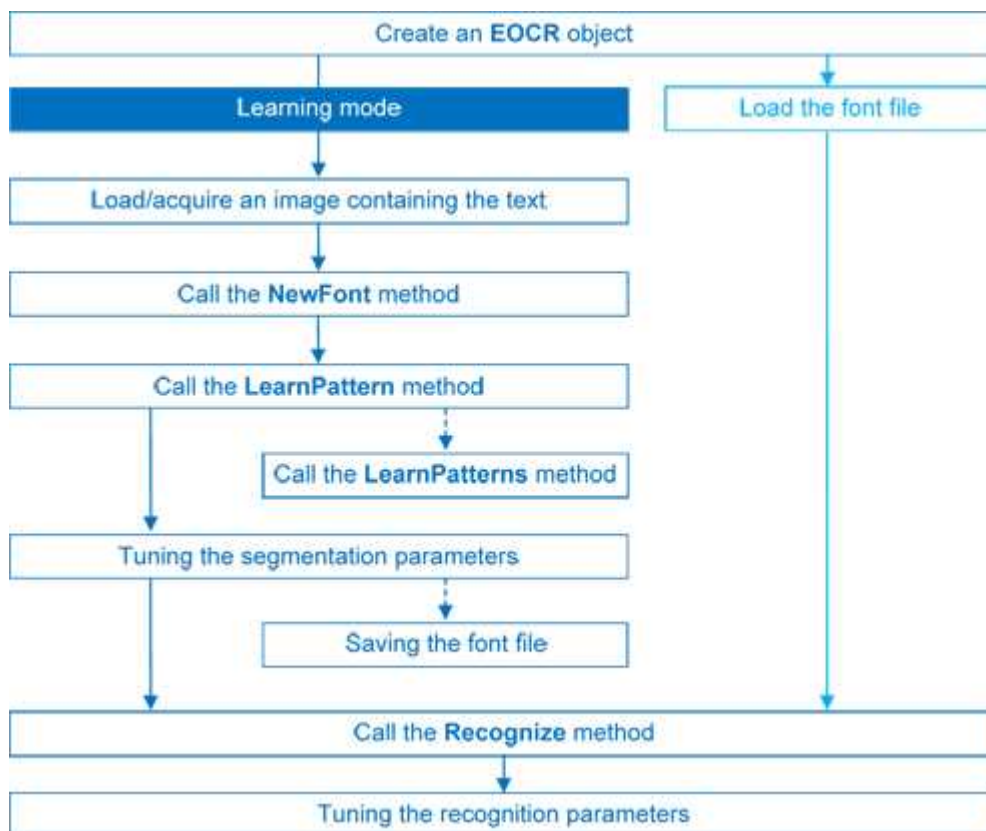
EasyOCR optical character recognition library reads short texts (such as serial numbers, part numbers and dates).

It uses font files (pre-defined OCR-A, OCR-B and Semi standard fonts, or other learned fonts) with a template matching algorithm that can recognize even badly printed, broken or connected characters of any size.

There are 4 steps to recognizing characters:



Workflow



Learning Process

You can learn characters to create font file if required.

Characters are presented one by one to EasyOCR which analyzes them and builds a database of characters called a font. Each character has a numeric code (usually its ASCII code) and belongs to a [character class](#) (which may be used in the recognition process).


Font files are created as follows:

1. [NewFont](#) clears the current font.
2. [LearnPattern](#) or [LearnPatterns](#) adds the patterns from the source image to the font. Patterns are ordered by their index value, as assigned by the [FindAllChars](#) process. The patterns in a font are stored as a small array of pixels, by default 5 pixels wide and 9 pixels high. This size can be changed before learning, using parameters [PatternWidth](#) and [PatternHeight](#).
3. [RemovePattern](#) removes unwanted patterns (optional).
4. [Save](#) writes the contents of the font to a disk file with parameter values: [NoiseArea](#), [MaxCharWidth](#), [MaxCharHeight](#), [MinCharWidth](#), [MinCharHeight](#), [CharSpacing](#), [TextColor](#).

Segmenting

For learning as well as recognition, EasyOCR segments the characters, i.e. locates the characters and determines their bounding box. This is done by means of blob analysis (thresholding followed by a grouping of pixels of the same color, as is done by EasyObject). After blobs have been found, they can be filtered to remove unwanted features (small blobs of noise, large extraneous objects, ...).

1. EasyOCR analyses the blobs to locate the characters and their bounding box, using one of two [segmentation modes](#):
 - **keep objects** mode: one blob corresponds to one character.
 - **repaste objects** mode: the blobs are grouped into characters of a nominal size. This is useful when characters are broken or made up of several parts. When a blob is too large to be considered a single character, it can be split automatically using [CutLargeChars](#).



Character segmentation by blob grouping


2. Filters remove very large and very small unwanted features.

3. EasyOCR processes the character image to normalize the size into a bounding box, extracts relevant features, and stores them in the font file. The patterns in a font are stored as arrays of pixels defined by `PatternWidth` and `PatternHeight` (by default 5 pixels wide and 9 pixels high).

Segmentation parameters

Segmentation parameters must be the same during learning and recognition. Good segmentation improves recognition.

- The `Threshold` parameter helps separate the text from the background. A too high value thickens black characters on white background and may cause merging, a too small value makes parts disappear. If the lighting conditions are very variable, automatic thresholding is a good choice.



Too high threshold value (left), Threshold adjustment (middle), Too low threshold value (right)

- `NoiseArea`: Blob areas smaller than this value are discarded. Make sure small character features are preserved (i.e., the dot over an "i" letter).
- `MaxCharWidth`, `MaxCharHeight`: Maximum character size. If a blob does not fit in a rectangle with these dimensions, it is discarded or split into several parts using vertical cutting lines. If several blobs fit in a rectangle with these dimensions, they are grouped together.
- `MinCharWidth`, `MinCharHeight`: Minimum character size. If a blob or a group of blobs fits in a rectangle with these dimensions, it is discarded.
- `CharSpacing`: The width of the smallest gap between adjacent letters. If it is larger than `MaxCharWidth` it has no effect. If the gap between two characters is wider than this, they are treated as different characters. This stops thin characters being incorrectly grouped together.
- `RemoveBorder`: Blobs near image/ROI edges cannot normally be exploited for character recognition. By default, they are discarded.

Recognition

The characters are compared to a set of patterns, called a **font**. A character is recognized by finding the best match between a character and a pattern in the font. After the character has been located, it is normalized in size (stretched to fit in a predefined rectangle) for matching. The normalized character is compared to each normalized template in the font database and the best matches are returned.

1. **Load**: reads a pre-recorded font from a disk file.
2. **BuildObjects**: The image is segmented into **objects** or blobs (connected components) which help find the **characters**. This step can be bypassed if the exact position of the characters is known. If the character isolation process is bypassed, you must specify the known locations of the characters: [AddChar](#) and [EmptyChars](#).
3. **FindAllChars**: selects the objects considered as characters and sorts them from top to bottom then left to right.
4. **ReadText**: performs the matching and filters characters if the marking structure is fixed or a character set filter was provided.

Character recognition: The characters are compared to a set of patterns, called a **font**. The best match is stretched to fit in a predefined rectangle and compared to each normalized template in the font database.

A **Character set filter** can improve recognition reliability and run time by restricting the range of characters to be compared. For instance, if a marking always consists of two uppercase letters followed by five digits, the last of which is always even, it is possible to assign each character a class (maximum 32 classes) then set the character filter to allow the following classes at recognition time: two uppercase, four even or odd digits, one even digit.

Steps 2 to 4 can be repeated at will to process other images or ROIs. The [Recognize](#) method can be used as well.


Additional information, such as geometric position of the detected characters, can be obtained using: [CharGetOrgX](#), [CharGetOrgY](#), [CharGetWidth](#), [CharGetHeight](#), ...

[CompareAspectRatio](#) makes character and font comparison sensitive to the difference between narrow and wide characters. It improves recognition when characters look like each other after size normalization.

Recognition parameters

- [MaxCharWidth](#), [MaxCharHeight](#): if a blob does not fit within a rectangle with these dimensions, it is not considered as a possible character (too large) and is discarded. Furthermore, if several blobs fit in a rectangle with these dimensions, they are grouped together, forming a single character. The outer rectangle size should be chosen such that it can contain the largest character from the font, enlarged by a small safety margin.
- [MinCharWidth](#), [MinCharHeight](#): if a blob or a group of blobs does fit in a rectangle with these dimensions, it is not considered as a possible character (too small) and is discarded. The inner rectangle size should be chosen such that it is contained in the smallest character from the font, shrunk by a small safety margin.
- [RemoveNarrowOrFlat](#): Small characters are discarded if they are narrow **or** flat. By default they are discarded when they are both narrow **and** flat.
- [CharSpacing](#): if two blobs are separated by a vertical gap wider than this value, they are considered to belong to different characters. This feature is useful to avoid the grouping of thin characters that would fit in the outer rectangle. Its value should be set to the width of the smallest gap between adjacent letters. If it is set to a large value (larger than [MaxCharWidth](#)), it has no effect.

- **CutLargeChars**: when a blob or grouping of blobs is larger than `MaxCharWidth`, it is discarded. When enabled, the blob is split into as many parts as necessary to fit and the amount of white space to be inserted between the split blobs is set by `RelativeSpacing`. This is an attempt to separate touching characters.
- **RelativeSpacing**: when the `CutLargeChars` mode is enabled, setting this value allows specifying the amount of white space that should be inserted between the split parts of the blobs.



Invalid recognition settings

Advanced tuning

These recognition parameters can be tuned to optimize recognition:

CompareAspectRatio: when this setting is on, EasyOCR is less tolerant of size and takes into account the measured aspect ratio. Using this mode improves the recognition when characters look similar after size normalization as it enforces the difference between narrow and wide characters.

Filtering the characters (in the `ReadText` method), can be used if the marking structure is fixed. When objects are larger than the `MaxCharWidth` property, they can be split into as many parts as needed, using vertical cutting lines.

ESegmentationMode, character isolation mode defines how characters are isolated:

- **Keep objects** mode: a character is a blob; no attempt is made to group blobs, thus damaged characters cannot be handled and small features such as accents and dots may be discarded by the minimum character size criterion.
- **Repaste objects** mode: blobs are grouped to form distinct **characters** if they fit in the maximum character size and are not separated by a vertical gap, thus preserving accents and dots.

2. EasyOCR2 - Reading Texts (Improved)

EasyOCR2 is an optical recognition library designed to read short texts such as serial numbers, expiry dates or lot codes printed on labels or on parts.

It uses an innovative segmentation method to detect blobs in the image, and then places textboxes over the detected blobs following a user-defined topology (number of lines, words and characters in the text). These methods support text rotation up to 360 degrees, can handle non-uniform illumination, textured backgrounds, as well as dot-printed or fragmented characters.

A character type (letter / digit / symbol) can be specified for each character in the text, improving recognition rate and speed. The character database that is used for recognition can be learned from sample images or read from a TrueType font (.ttf) file.

Text recognition with EasyOCR2 follows four phases:



Input image (left) and fitting textboxes (right)



Image segmentation (left) and recognition (right)

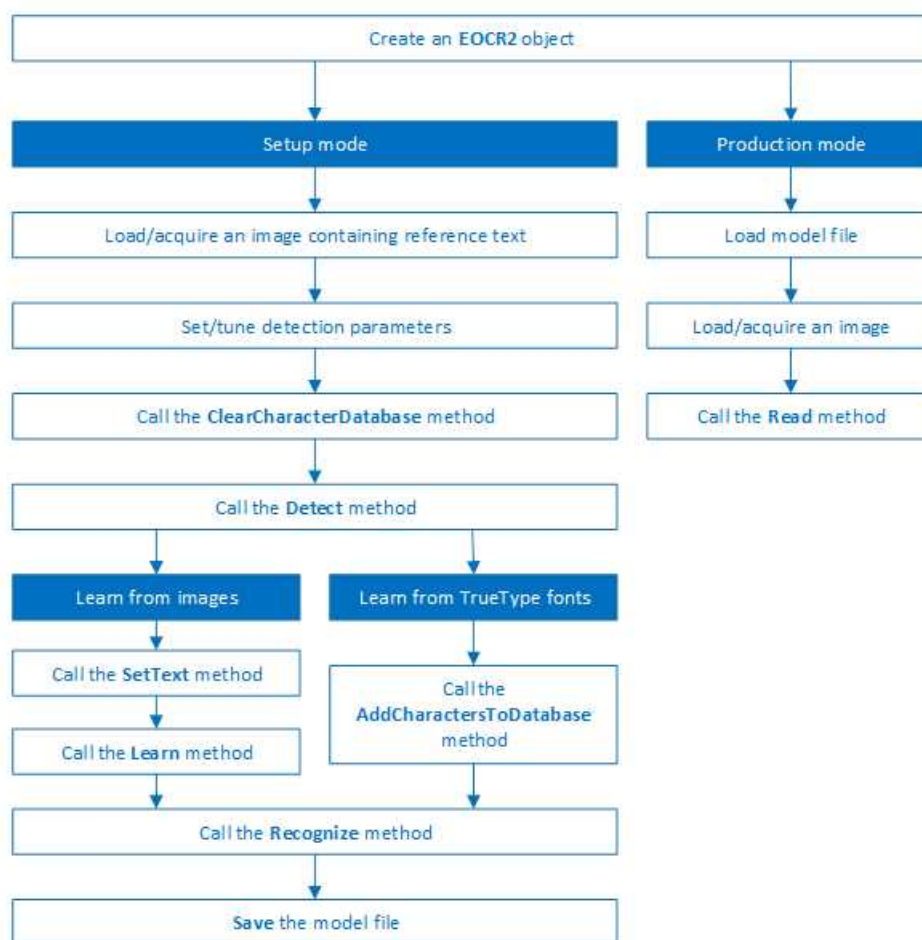
EasyOCR2 vs EasyOCR

EasyOCR2 will give better results than EasyOCR when dealing with:

- Unknown text rotation

- Dotted or fragmented characters
- Non-uniform illumination or textured backgrounds
- When TrueType font files are available that match the text to be read, EasyOCR2 allows the user to use those font files directly for recognition, while EasyOCR does not.
- When none of the above are relevant to the application, the user may prefer to use EasyOCR to EasyOCR2 due to its superior computational speed.

Workflow



Detection

EasyOCR2 finds characters in an image as follows:

1. EasyOCR2 segments the image, finding blobs that represent (parts of) the characters.
2. Blobs that are too large or too small to be considered part of a character are filtered out.
3. EasyOCR2 fits character boxes to the detected blobs according to a given `topology` and `detectionMethod`.

The topology describes the structure of the text in the image, defining the number of lines, the number of words per line and the number of characters per word.

4. EasyOCR2 extracts the pixels inside each character box from the image.

The resulting character-images can be used to learn or recognize the characters.

A workflow detecting text in an image could be as follows:

- a. Set the required detection parameters.
- b. Alternatively, call `Load` to read a pre-made model (.o2m) file containing detection parameters from disk.
- c. Call `Detect` to extract the text from the image.

The method `Detect` will return an `EOCR2Text` structure that contains a textbox and a bitmap image for each character, hierarchically stored in `EOCR2Line` -> `EOCR2Word` -> `EOCR2Char` structures.


See example in code snippet: [Detecting Characters](#)

FIXED WIDTH


An example of a fixed-width font, processed with the `detectionMethod` `'EOCR2DetectionMethod_FixedWidth'`

PROPORTIONAL


An example of a proportional font, processed with the `detectionMethod` `'EOCR2DetectionMethod_Proportional'`




The `textAngle` estimate for this image is slightly off when `NumDetectionPasses=1`



The `textAngle` estimate is better when `NumDetectionPasses=2`



For this dotted text, setting 'CharsMaxFragmentation' to 0.1 leads to incomplete segmentation results



Setting 'CharsMaxFragmentation' to 0.01 gives better segmentation results

Detection parameters

Required parameters

- The parameter `Topology` tells the box-fitting method how to structure the textboxes it fits to the detected blobs. Using a modified version of Regex expressions, the topology determines the number of lines in the text, the number of words per line and the number of characters per word. The section **Recognition Parameters** contains an extensive explanation of the syntax for the Topology.
- The parameter `CharsWidthRange` tells the segmentation and detection methods how wide the characters in the image can be.
- The parameter `CharsHeight` tells the segmentation and detection methods how high the characters in the image can be.
- The parameter `TextPolarity` tells the segmentation method whether it should look for light characters on a dark background or vice versa.

Advanced parameters for segmentation (optional):

- The `CharsMaxFragmentation` parameter tells the segmentation algorithm how small blobs can be to be considered (part of) a character. The minimum allowed area of a blob is given by:

$$\text{minArea} = \text{CharsMaxFragmentation} * \text{CharsHeight} * \text{min}(\text{CharsWidthRange})$$

This parameter should be set between 0 and 1, the default setting is 0.1.

- The `MaxVariation` parameter determines how stable a blob in the image should be in order to be considered a potential character. A region with clearly defined edges is generally considered stable while a blurry region is not. A high setting allows detection of blobs that are more unstable, a low setting allows only very stable blobs. This parameter should be set between 0 and 1, the default setting is 0.25.
- The `DetectionDelta` parameter determines the range of grayscale values used to determine the stability of a blob. A low setting will make the algorithm more sensitive to noise; a high setting will make the algorithm insensitive to blobs with low contrast to the background. This parameter should be set between 1 and 127, the default setting is 12.

Advanced parameters for detection (optional)

- The parameter `DetectionMethod` selects the algorithm used for fitting. The setting `EOCR2DetectionMethod_FixedWidth` (default) is optimized for texts with fixed width fonts (including dotted text), the setting `EOCR2DetectionMethod_Proportional` is optimized for texts with proportional fonts.
- The `TextAngleRange` parameter tells the box-fitting method how the text in the image is oriented. It will test the following range of rotation angles:

$$\text{min}(\text{TextAngleRange}) \leq \text{angle} \leq \text{max}(\text{TextAngleRange})$$

where angles are defined with respect to the horizontal. The unit for the angles (degrees/radians/revolutions/grades) can be set using `easy::SetAngleUnit()`.

The default setting for this parameter is [-20, 20] degrees.

- The parameter `NumDetectionPasses` determines how many passes are made to fit textboxes to the detected blobs. The initial pass will fit textboxes to all detected blobs. Subsequent passes will select only those blobs that are covered by the textboxes from the previous pass and fit textboxes to that subset of blobs, potentially resulting in a more optimal fit. This parameter should be set to either 1 or 2, the default setting is 1.

Advanced parameters, specific for the setting `EOCR2DetectionMethod_FixedWidth`

- The `RelativeSpacesWidthRange` parameter tells the box-fitting method how wide the spaces between words may be. It will test the following range of spaces:

$$\text{min}(\text{SpacesWidthRange}) * \text{charWidth} \leq \text{space} \leq \text{max}(\text{SpacesWidthRange}) * \text{charWidth}$$

- The parameter `CharsWidthBias` biases the optimization toward wider or narrower character boxes.

- The parameter `CharsSpacingBias` biases the optimization toward smaller or larger spacing between character boxes.

Additional remarks

- When the setting `EOCR2DetectionMethod_FixedWidth` is selected, all character boxes will have the same width and they do not necessarily have to fit tightly around the characters.
- When the setting `EOCR2DetectionMethod_Proportional` is selected, the character boxes will fit tightly around the characters, if any character falls outside the range of allowed character widths, the detection will fail.

Learning

In order to recognize characters, EasyOCR2 requires a database of known reference characters. We may generate this character database from images and/or from TrueType system fonts.

A workflow to build a character database could be as follows:

- Set the required detection parameters or call `Load` to read the model (.o2m) file from disk.
- Optionally, call `ClearCharacterDatabase` to clear the current character database.
- Call `Detect` to extract the text from the image.
- Call `SetText` in the extracted text structure to set the correct value for each character.
- Call `Learn` to add the detected characters and their correct value to the current character database.
- Call `SaveCharacterDatabase` to save the current character database to disk.
- Alternatively, call `Save` to save the model file to disk, including the detection parameters and the created character database.

See example in code snippet: [Learning Characters](#)

Recognition

EasyOCR2 recognizes characters using a classifier that is trained on the character database. For each input character, the classifier will calculate a score for all candidate outputs, the candidate with the highest score will be returned as the recognition result. Through the `Topology` parameter, prior information about each character can be passed to the classifier, reducing the number of candidates and improving the recognition rate.

The production workflow for recognizing text from images could be as follows:

- Call `Load` to read the model (.o2m) file from disk. The model file contains all detection parameters, as well as the topology and the reference character database.
- Load or acquire the image.
- Call `Read` to detect and recognize the characters.

- Alternatively, call `Detect` to extract the text from the image, followed by `Recognize` to recognize the extracted text. This allows the user to modify elements of the detected text before recognition if so desired.

The methods `Read` and `Recognize` will return a string with the recognition results. To access more in-depth information about the results, one may call `ReadText`. This returns an `EOCR2Text` structure that contains the coordinates and sizes of each textbox as well as a bitmap image and a list of recognition scores for each character.

See example in code snippet: [Reading Characters](#)

Recognition parameters

The `Topology` parameter specifies the structure of the text (number of lines/words/characters) as well as the type of characters in the text. The recognition method will limit the number of candidates for each character based on the given topology.

It uses modified regular expression wildcards:

- “.” (dot) represents any character (not including a space).
- “L” represents an alphabetic character.
 - “Lu” represents an uppercase alphabetic character.
 - “Ll” represents a lowercase alphabetic character.
- “N” represents a digit.
- “P” represents the punctuation characters: ! “ # % & ‘ () * , - . / : ; < > ? @ [\] _ { | } ~
- “S” represents the symbols: \$ + - < = > | ~
- “\n” represents a line break.
- “ ” (space) represents a space between two words.

Combinations can be made, for example: `[LN]` represents an alpha-numeric character. To specify multiple characters, simply add `{n}` at the end for `n` characters. If the amount of characters is uncertain, specify `{n,m}` for a minimum of `n` characters and a maximum of `m` characters.

The topology `“[LuN]{3,5}PN{4} \n .{5} LL”` represents a text comprised of 2 lines:

- The first line has 1 word composed of 3 to 5 uppercase alpha-numeric characters, followed by a punctuation character and 4 digits.
- The second line has 2 words. The first word comprises of 5 wildcard characters, the second word has 2 letters (upper- or lowercase).

The topology `“L{3}P N{6} \n L{3}P NPN{4}”` represents a text with 2 lines:

- The first line has 2 words. The first word has 3 uppercase letters followed by a punctuation mark, the second word has 6 digits.
- The second line also has two words. The first word has 3 uppercase letters followed by a punctuation mark. The second word has 2 digits, followed by a punctuation mark and 4 additional digits.

The topology `“.{10} \n .{7} \n .{5} .{5} \n .{5} .{7}”` represents a text with 4 lines:

- The first line contains a single word of 10 (ASCII) characters
- The second line contains a single word of 7 characters



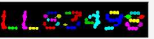



- The third line contains two words, each of 5 characters.
- The fourth line contains two words of 5 and 7 characters respectively.

3. EasyOCV - Validating Texts

Optical Character Verification compares a geometric pattern, a *sample*, with a predefined model, a *template*, while taking into account relative displacement of the constituent parts. For example, a printed part number may be checked for: correct placement with respect to the component body, sufficient contrast, good character shapes, or absence of inking defects.

Workflow

From the raw image to the final model, the model definition follows a logical sequence of steps.


					
Raw image	Thresholded image	Free blobs	Free characters	Single text	Final model

1. Threshold the image to separate the foreground (marking) from the background.
2. Free blobs: Blob analysis is performed and the detected objects may be selected depending on their size to generate a set of candidates. (Blobs can be selected or unselected manually, to add features of unusual size or delete spurious objects.)
3. Free characters: The blobs are aggregated into boxes corresponding to the characters. (Manual correction is possible.)
4. Create single text: The characters are aggregated into boxes corresponding to the texts. (Manual correction is possible.)
5. Create final model: The system will analyze the shape of the template and generate the required data structures that represent it and allow fast inspection. It will also perform some quality measurement on the template for later comparison with the sample. At this stage, only texts and their constituent characters are stored. The free characters and free objects are discarded.
6. Before the template can be saved, the inspection parameters must be defined . These include:
 - the allowed ranges for the location parameters of the texts with respect to their nominal position in the inspected ROI. These parameters correspond to translation and rotation (scaling in both horizontal and vertical directions, and shearing).
 - the allowed ranges of character location with respect to their nominal position in the texts (during translation).
 - the allowed ranges of the quality ratings with respect to their nominal values. These parameters can be chosen among the area of the character background and foreground, the accumulated gray level of these areas, and a similarity coefficient.

Learning Process

Model structure

Inspection takes place in a rectangular ROI. The marking is a set of texts, made of characters, which are made of blobs. Typically there is only one text, and each character is a blob.



In a simple application, a ready-made template (Open eVision Studio provides a comprehensive template editor) is used in the inspection phase.

EasyOCV requires training on correctly printed marks to create a good quality template. Using a single image to create a template has drawbacks:


- the chosen image may itself have small, unnoticed defects and not be fully representative of the whole population.
- a single image gives no insight on the random variations between acceptable samples, and gives no way to adjust quality indicator tolerances.
- the default quality tolerances may not give the sharpness you would like to detect.

Statistics will help you define acceptance criteria. Quality indicator reference values are computed from the template image.

template design considerations

- **Should marking be one piece of text or several?** This depends on the possible movements; marks that move together should be considered a single text. During inspection, texts can move independently of each other.
- **How much can the text position vary in terms of translation, rotation** [and possibly scaling and/or shearing]?
- **Can the text be decomposed into characters ?** A character is the smallest part of a marking that can be inspected in isolation. A small displacement from its nominal position can be allowed and measured.
- **Can individual characters move with respect to their containing text ?**


- **What movement may occur ?** Placement repeatability must be evaluated. If the part travels along a guide, sometimes only translation occurs which can be handled by a single alignment pattern (fiducial). If rotation or scaling can occur, two alignment patterns are preferable.



Bad: rotation not handled


- **Are the alignment patterns fixed, well-contrasted features, not subject to degradation, that move rigidly with the inspected part ?**

When two are used, they should be located as far apart as possible for optimal accuracy. The pattern ROIs should not contain extraneous features likely to change from sample to sample. The patterns should be small so that rotation and scaling has little impact, but large enough to contain information at different scales.



Bad: the location pattern is not repeatable

- Is the search area as small as possible to reduce search time and avoid false matches, but big enough to contain the match?
Check on a representative set of images that location by pattern matching never fails by touching the search area edges.



Bad: too tight search areas

- Does the inspected ROI on the mother image surround all areas where defects may be detected, but not raise false alarms?




Bad: undue inspection of the background

Access the template components


The **OCV context** contains a list of **texts**. Each text contains a list of **characters**.

To access a specific text, you provide its **index**, traversing the hierarchy from the OCV context. To address a particular character, you provide the indexes of both the text and the character.



Access by index (second character of first text)

You can access and modify several components collectively. Every text and every character has a boolean **selected** property, so that in a single operation, you can read/write the value of a property for all currently selected elements. When reading a property, a value is only returned if it is identical for all selected elements.



Access a group of selected components

Scattering operations over selected components is very handy for interactive editors that list and modify parameters for single items or groups of items.

3.1. Learning Passes

After ROI placement and pattern learning (`Register` operations), training still requires two passes:

- To compute an **average** ideal, noise-free, image that reveals the central tendency of the part image.
For each image, realign and normalize (`Register`).
If the operation is successful (good pattern location), call `Learn` (`ELearningMode_Average`) for immediate processing (on-the-fly learning), or `AddPathName` for deferred processing (batch learning).
- To measure **deviations** around the average image.
For each image, realign and normalize (`Register`).
If the operation is successful, call `Learn` (`ELearningMode_AbsDeviation`) for immediate processing and recommended method, or `Learn` (`ELearningMode_RmsDeviation`) for enhancing large deviations.

The images used can be the same for both passes, or two distinct sets of images of different sizes can be used (on-the-fly learning). `BatchLearn` performs both passes for all images in the file list.

A learning set size of at least 16 images is recommended.

Inspect and compare image with model

Normally when you inspect the sample image and compare with the model, text-level inspection is sufficient.

Character-level inspection is more detailed and complex.

The inspection process involves two operations.


- **Locate:** A region of interest is scanned and the best match is found between it and the template, using all desired "[Degrees of Freedom](#)" [below](#).
- **Score:** Every sample character is compared to the corresponding template character. Character quality indicators are computed, collated into "[Quality Indicators](#)" [on page 176 quality indicators](#) of the text, and compared to acceptance intervals. Unacceptable values have their corresponding characters flagged, a diagnostic code is generated, and global diagnostics summarizing all text and character defects are issued.

3.2. Degrees of Freedom

Degrees of freedom can compensate for misalignment and distortion. Each degree of freedom increases the running time, so use them sparingly. In many cases, text and character translation are sufficient. When large amplitude skewing is possible, text translation + skewing can be used. Care must be exercised when combining the other degrees of freedom.

Text translation


All texts can be moved horizontally ($ShiftX$) and vertically ($ShiftY$) in a specified range.



Text translation

Text skewing


All texts can be rotated about the center of their bounding box using the angle defined by $Skew$.



Text skewing

Character translation


All characters can be moved individually horizontally ($ShiftX$) and vertically ($ShiftY$) with respect to their nominal position.



Character translation

Text X/Y-scaling (Advanced - only use if necessary)

All texts can be re-scaled horizontally $ScaleX$ and vertically $ScaleY$, while the center of their bounding box remains fixed. Re-scaling can be isotropic (both scale factors are identical) or anisotropic.




Anisotropic text scaling

Note: By default, text scaling is supposed to be isotropic. If you think your application might generate anisotropic scaling for a text, you must set the *IsotropicScaling* parameters of the *EOCVText* corresponding object to **FALSE** before inspecting. Working with isotropic scaling results in an effective time saving during inspection.

Text shearing (Advanced - only use if necessary)

All texts can be sheared using the angle parameter *Shear*, i.e. become italic, while the center of their bounding box remains fixed.



Text shearing

Degrees of Freedom parameters

The degrees of freedom for location are specified with respect to the position at learning time *nominal position*. The characters are remembered relative to the corresponding text center. The text centers are remembered relative to the template image/ROI center. The default skew and shear angles are **0** and scale factors are **1**.

The parameters must vary within a range of $\text{Bias} \pm \text{Tolerance}$. If bias is **0**, the range is centered around the nominal value.

The number of positions tried for each degree of freedom is specified as follows:

- Translation:** Every integer value in the ranges $\text{ShiftXBias} \pm \text{ShiftXTolerance}$ and $\text{ShiftYBias} \pm \text{ShiftYTolerance}$ is tried. However, for efficiency reasons, the *ShiftXStride* and *ShiftYStride* parameters can be set to a value larger than **1**, so that a gross location pass with the specified stride is followed by a finer one with unit stride. (The expected speed-up is on the order of the square of the *Stride* parameter.) Anyway, choosing too large a value may cause mismatches when local maxima are present. (A value on the order of a fraction of the character size is recommended.)

- **Skewing, scaling and shearing:** the `SkewCount`, `ScaleXCount`, `ScaleYCount` and `ShearCount` parameters indicate the number of values tried for each degree of freedom. The execution time increases as the product of these counts. When a degree of freedom is not used, its count must be left as **1**. When `SkewCount` is set to **0**, EasyOCV automatically chooses an appropriate count value.

3.3. Quality Indicators

After locating the model, the inspection process compares the sample with the template, and rates the resemblance at a character level.

The parameters computed for the **template** serve as a reference and are compared to those computed on the **sample**.

When the template image is binarized, the marking appears as white foreground on black background in the character's bounding box. The bounding box is the tightest rectangle that wholly contains an item, with a safety margin.




Template image, foreground and background (white pixels) Sample image, foreground and background (white pixels)

Area-Based Quality Indicators

When the sample image is rated, thresholding also separates white and black pixels. The foreground (or background) sample areas are defined as the count of the white(or black) sample pixels in the foreground (or background) region of a character. This is not the same as the total count of white and black pixels in the sample.

The difference between the template and sample areas is the area of defects in the character foreground [or background].




Foreground and background areas (white pixels)

Note: The area-based indicators rely on thresholding of the sample image. If necessary, the threshold level must be compensated for a change in intensity (automatic thresholding).

Gray Sum-Based Quality Indicators

A different measure of the amount of light reflected by the marking is given by sums rather than

counts: the foreground [background] sample sum is defined as the sum of the gray-level values of all pixels of the foreground [background] region of the sample image. The foreground [background] template sum is the same feature computed on the template image to provide a reference value. Optionally, the sums are normalized with respect to the reference foreground and background average gray-levels to compensate for possible changes in gain (contrast) and offset (intensity).



Foreground and background sums

Note: *The sum-based indicators do not rely on thresholding of the sample image, but the reference foreground and background gray-levels may take into account changes in gain and offset. Characters are accepted or rejected by comparing the indicator values of the sample and template: if the difference is larger than the specified tolerance, a defect is reported. A smaller foreground value indicates under-printing or missing character parts. A larger background value indicates over-printing or spurious character parts. Character mismatches provoke both kinds of anomalies.*

Correlation-Based Quality Indicators

Normalized correlation rates mismatches between two images. The correlation parameter is a global score in range **0** to **1**, which is implicitly corrected for a change in gain and offset.

Note: *The correlation-based quality indicator should be as close as possible to **1**. It is not sensitive to changes in gain or offset.*

Reporting

Defects are reported in three ways:

- explanatory diagnostics are given for each inspected character and text;
- the items for which diagnostics are reported are highlighted on the display;
- the relevant items are drawn as a box crossed by its main diagonal.

3.4. Advanced Features

EasyOCV can accumulate the results of a series of consecutive inspections: quality indicators average values and standard deviations. EasyOCV functions can use these statistics to define acceptance criteria and automatically adjust position and quality tolerance parameters, so you can better control the manufacturing process.

- Average values show the long-term trend behavior of the system and detect marking drifts.

- Standard deviations show process repeatability and detect appearance of slack. It is customary to select a tolerance value that is a small multiple of the observed standard deviation (± 2 sigma or ± 3 sigma criterion).

3.5. Programming with EasyOCV

Introduction

Writing an inspection application for the production line can be simple if no operator intervention is required to adjust parameters, and if no in-situ learning phase is necessary:

- A model can be loaded once for all.
- Every acquired image is inspected.
- The inspection results are graphically displayed on top of the image.
- A diagnostics report of quality indicators and statistics can be generated.

A more advanced inspection application may allow the operator to modify parameters:

- Dialog boxes must be provided to edit parameters.
- Parameters may be changed globally (same value everywhere), or the operator may select the texts and characters on which to work.

In a complex application, model edition before learning must be made possible which leads to more complex programming:

- select and unselect items.
- display and modify parameters.
- load and save to a model file.

The required steps to create an inspection application, from simplest to advanced are:

1. Set ROI for inspection

The ROI should be placed center on the marking, in a position that is repeatable with respect to the marking substrate. This is achieved either when the position of the inspected object is known and stable, or when the object has been located by pattern matching or edge measurement. .

The ROI :

- Defines the effective search area for text using their centers and location parameters (ShiftX, ShiftY, ...). This way, inspection is not confined to the inspected ROI.
- Locates the marking, then evaluates global contrast in this ROI (centered on the marking with the learning time dimensions). The ROI is not used to evaluate global contrast of the marking.

2. Inspect

Uses a threshold level to determine the global contrast of the marking. Automatic thresholding can be used.

3. Draw inspected items

The inspected texts and characters can be represented by their bounding box, at the position determined by the location process.

Selected and unselected items can be drawn in different colors.

Items with detected defects appear crossed by their main diagonal.

4. Retrieve Diagnostics and Quality Indicators

Global inspection diagnostics summarize all defects found on the marking and raise alarms.

Detailed diagnostic reports are available for each text and character, and all measured quality indicators can be retrieved.

- Retrieving text diagnostics and parameters involves a loop where all texts are visited.
- Retrieving character diagnostics and parameters involves a double loop where all texts and all characters of all texts are visited.

5. Set Inspection Parameters

During operation working parameters can be adjusted in various ways:

- **Global change:** a parameter value may be set for all texts and/or all characters. This is straightforward and requires a single call to `ScatterTextsParameters`, `ScatterTextsCharsParameters`, but before calling, make sure that the parameters you don't want to change are set to an undefined value.
- **Custom change:** the values can be adjusted individually using a user-defined rule. This approach is similar to the retrieval of parameters using indexed access.
- **Selective change:** parameter values can be set for texts or characters in a selected state by interactively selecting the text or characters.

Selecting Items Interactively

To retrieve or modify parameters, individually or grouped, the operator must have the ability to select them using a mouse. EasyOCV provides a general selection/de-selection mechanism: several functions can toggle the state of all/selected/unselected items in a given rectangle.

Note: *The rectangle is usually obtained by a dragging operation. A degenerate rectangle (reduced to a single point) can be used to handle point clicking.*

Since the toggling mechanism combined with the possible rectangle extent and current selection mode is tricky, let us give a few examples.

Assume a model of three texts in the following states: Selected, Selected, Unselected.

- Using a rectangle that contains all three of them will set them to states Unselected, Unselected, Selected (SSU -> UUS).
- Using a rectangle that touches the first of them will set the states to Unselected, Selected, Unselected (SSU -> USU).

Now consider the same operations applied to the selected texts only.

- Using a rectangle that contains all three texts will set them to states `Unselected`, `Unselected`, `Unselected` (`SSU` -> `UUU`).
- Using a rectangle that touches the first of them will set the states to `Unselected`, `Selected`, `Unselected` (`SSU` -> `USU`).

Now consider the same operations applied to the unselected texts only.

- Using a rectangle that contains all three texts will set them to states `Selected`, `Selected`, `Selected` (`SSU` -> `SSS`).
- Using a rectangle touches the first of them will leave their states unchanged (`SSU` -> `SSU`).

Note: *This selection mechanism applies to texts and characters at inspection time (`SelectSample...`). It also applies to free objects, free characters and texts during the model edition phases (`SelectTemplate...`).*

Compute Inspection Statistics

Using EasyOCV, gathering statistical information on the process is possible. For each measured parameter (location parameters and quality indicators), the average and standard deviation can be estimated from a number of samples.

The procedure is straightforward: after an image has been inspected, one can request that the measured parameters be taken into account as valid samples by calling `UpdateStatistics`. (If, for any reason, the sample is to be rejected, just do not call `UpdateStatistics`.) After at least two sample images have been processed, the average and standard deviations can be obtained. The standard mechanisms for text and character parameters retrieval can be used.

To compute the statistics afresh on new samples, start by calling `ClearStatistics`. The number of samples accumulated so far is given by `StatisticsCount`.

Adjust Inspection Parameters from Statistics

Statistics may be used to adjust location and quality indicators. When making adjustments to individual texts or characters, the selection mechanism described above is applicable.

To adjust quality ranges indicators, call `AdjustTextsQualityRanges`, `AdjustCharsQualityRanges`. If you do not want to adjust the quality range of a particular indicator, you should de-activate it by setting `UsedQualityIndicators`.

- the bias value of each indicator is assigned the average value of the inspected samples (provided they had been added to statistics).
- the indicator tolerance is assigned s times the standard deviation, where s is a security factor to provide.

To adjust location parameters, call `AdjustTextsLocationRanges`, `AdjustCharsLocationRanges`. When adjusting location parameters, you must specify minimum and maximum values and a security factor may also be specified.

Interactively Edit a Model

Writing an OCV model editor requires a good understanding of windowed applications design. In particular, it is important to know how to manage the mouse cursor movements, when and

how to refresh the display, handle dragging of selection rectangles and the like. It is out of the scope of this documentation to explain these features which are deeply related to Windows programming. Also note that the level of functionality, from blind -no operator intervention- to full fledged editing is a matter of taste and of programming skill.

Recall the steps in defining the model structure:

1. An `ECodedImage` object is used to segment the image into blobs (`BuildObjects` method).
2. Possibly, blob selection by all means provided in `EasyObject` (legacy), is performed (`SelectObjectsUsingFeature` or `SelectObjectsUsingPosition`). In particular, small blobs generated by noise should be unselected.
3. The selected blobs are passed to the OCV object and enter the free objects list (`EasyObject` (legacy)'s blobs become OCV's free objects, or `TemplateObjects`).
4. At this point, the objects in the free list can be selected/unselected interactively.
5. The (selected) free objects are then used to generate free characters, using one of the available grouping policy (free objects become free characters, or `TemplateChars`).
6. At this point, the free characters can be selected/unselected interactively.
7. The (selected) free characters are then used to generate texts. The default policy is to group all free characters in a single piece of text (free characters become texts, or `TemplateTexts`; these texts now contain embedded characters, or `TemplateTextChars`).

In the simplest form of a model editor, steps 4 and 6 can be skipped, meaning that all free objects and all free characters will enter the model. A better editor will allow withdrawal of unwanted items and explicit grouping (steps 4 and 6). An even more powerful editor should allow grouping as well as ungrouping (backwards from 3 to 2, from 5 to 4, from 7 to 6).

At any time, the following operations can be handled:

- The model components can be selected/unselected interactively (using `SelectTemplateObjects`, `SelectTemplateChars`, `SelectTemplateTexts`).
- New items can be grouped to form new higher level items (using `CreateTemplateObjects`, `CreateTemplateChars`, `CreateTemplateTexts`).
- Items can be ungrouped by destroying the higher level item (using `DeleteTemplateTexts`, `DeleteTemplateChars`, `DeleteTemplateObjects`).

A clean way to organize the editor is to define a sequence of separate phases dealing with objects, free objects, free characters and texts.

Advanced Features: Change contrast, location mode, Quality indicators and resample characters

Contrast Parameters

Image contrast is an important factor during both learning and inspection.

The background and foreground information must be separated using an appropriate threshold, that may be determined automatically. After a threshold is given, the average gray level of the background and foreground are computed separately, and become the reference gray levels. These are used to measure the image contrast and normalize the gray level quality indicators if needed.

The **background** and **foreground** reference gray levels are computed for both the **template** and **sample** images. See the [EOCVChar](#) properties. The sample contrast may then be compared to the templates contrast reference value, to diagnose an over-contrasted or under-contrasted image before further analysis.

Note: *The template threshold directly influences the thickness of the blobs in the model.*

Note: *The sample threshold influences the reference gray levels of the sample image.*

Note: *When gray level normalization is used, it influences location score, gray-level sums, and blob thickness in the sample image, which has immediate consequences on the sample areas.*

Location Modes

For location of the model components, a search process is used during which EasyOCV tries to find the character edges in the sample image, possibly transformed. Four location modes are provided: **raw**, **binarized**, **gradient** and **Laplacian**. See enumeration constants [ELocationMode](#).

Experience reveals that binarized and gradient modes are the most reliable at locating components.. Additionally, the gradient mode is not sensitive to the threshold level. Use of the Laplacian mode is not recommended.

In case you experience location problems, you should try another location mode.

Location Score

It is highly recommended to keep the default [reduction of location scores](#) option turned on. Reduction consists of dividing a raw location score by the number of points its computation required. Thus, the calculated scores do not depend on the number of used points any longer; the number of used points may be decreased without degrading localization to save potential time, if necessary.

Location scores may also be [normalized](#). This option is useful if the lighting conditions of sample images are not the same, or when template and sample images have obviously different reference gray levels. The action of this option is equivalent to performing a global contrast correction on the sample image (or ROI). Location scores do not depend on reference gray levels, so are more reliable.

Finally, [AccurateTextsLocationScores](#) provides an alternative way to compute text location score. During the location process, EasyOCV tries first to locate texts using their contours, which are sets of fixed points one from each other. This rigid definition of text contour has the drawback of making the library return poor location score values if the sample characters have moved from their nominal positions, and may result in a false alarm.

If the [AccurateTextLocationScores](#) property is turned on, a text location score will be computed as sum or average of the characters location scores that form the text (depending on the state of the [ReduceLocationScore](#) property). This way, texts location scores become independent from the position of the characters they contain.

Used Quality Indicators

For a given inspection case, not all quality indicators are relevant. For instance, it sometimes suffices to use the location scores alone to detect absence of a given marking.

To avoid false alarms raised by unused quality indicators for which the tolerances have not been adjusted, and to avoid unnecessary processing, it is important to activate only the quality indicators in use through the `UsedQualityIndicators` property.

Character Resampling

Normally, when text is inspected with rotation, scaling and/or shearing, some resampling must be performed to compute the quality indicators on the separate characters. If the angles remain small and the scale factors remain very close to unity, this resampling can be avoided by setting the `ResampleChars` parameter to **FALSE**.

PART VI
CODE IDENTIFICATION TOOLS

1. EasyBarCode - Reading Bar Codes


1.1. Reading Bar Codes



Bar code (EAN 13 symbology)

[EasyBarCode](#) can locate and read bar codes automatically. Location can be performed manually for prototyping or when automatic mode results are unsatisfactory.

Workflow



Bar code definition

A bar code is a 2D pattern of parallel bars and spaces of varying thickness that represents a character string. It is arranged according to an encoding convention (**symbology**) that specifies the character set and encoding rules.

- The bar code may be black ink on white background or inverted (white ink on black background).
- The bar code should be preceded and followed by a quiet zone of at least ten times the module width (smallest bar or space thickness).
- Bars should be surrounded below and above by a quiet zone of a few pixels.
- Bar and space widths must be greater than or equal to 2 pixels.

symbolologies

A symbology defines the way a bar code is encoded.

Symbologies can be enabled in [StandardSymbologies](#) or [AdditionalSymbologies](#) parameters.

The standard symbologies are enabled by default:

- Code 39
- Code 128
- Code 2/5 5 Interleaved
- Codabar
- EAN 13*
- EAN 128
- MSI
- UPC A*
- UPC E

Note: * EAN 13 and UPC A only differ by the layout of surrounding digits.

Additional symbologies that are supported:

- ADS Anker
- Binary code
- Code 11
- Code 13
- Code 32
- Code 39 Extended (a super-set of Code 39)
- Code 39 Reduced (a subset of Code 39)
- Code 93
- Code 93 Extended
- Code 412 SEMI
- Code 2/5 3 Bars Datalogic
- Code 2/5 3 Bars Matrix
- Code 2/5 5 Bars IATA
- Code 2/5 5 Bars Industry
- Code 2/5 5 Compressed
- Code 2/5 5 Inverted
- Code BCD Matrix
- Code C.I.P
- Code STK
- EAN 8
- IBM Delta Distance A


- Plessey
- Telepen

Checksum

A checksum character enables the reader to check the barcode validity depending on the symbology:

- The checksum may be mandatory and must be checked by the reader.
- The checksum may be mandatory but may not need to be checked.
- The checksum and its verification may both be optional.

`VerifyChecksum` enables or disables (default) checksum verification.



Bar code structure (Code 39)

Read a bar code

The **Automatic** mode reading algorithm locates a bar code in the field of view and [Reads](#) it. If several bar codes are present, only one is located, like a straightforward hand-held bar code reader.

Before reading, the decoding symbologies must be specified in the `StandardSymbologies`, or `AdditionalSymbologies` properties.

Mono-symbology mode reads the bar code using the expected symbology type(s) and reports the encoded information (if readable) or the reason for failure (if not readable). There is only one interpretation for the character string.



Decoded bar code


Note: When the bar code contains `\0x00` characters, the `std::string::c_str` method should not be used (since C-strings are terminated by the `\0x00` character). An iterator over the characters should be used instead of a C-string.

Advanced features

Locate and Read bar code manually

If automatic localization fails or for prototyping purposes, the user can provide the **bar code position** and **reading area** to manually locate the code.

- **Bar code position** can be provided graphically by a bounding box around the bar code or by its parameters. If several symbols appear in the image, they can be processed one after the other.
- The **reading area** of the bar code is the area that is read. It should be wider than the bar code bounding box width, and less high than the bar code bounding box height. It may also be rotated relatively to the bar code bounding box, to take into account slanting bars (Advanced mode!).



Bounding box — graphical appearance (manual location)

Reading area — graphical appearance (manual location)

Read all interpretations (multi-symbology mode)

Use `Detect` to report the number of possible symbologies in the `NumEnabledSymbologies` property, and list the data contents by decreasing likeliness.

Then call the `Decode` method in a loop, using `GetDecodedSymbology` to walk through the list of successful symbologies in decreasing order of likelihood.

1.2. Reading Mail Bar Codes



Mail bar code example


Specifications

The Mail Bar code Reader:

- Detects and decodes postal 4-state bar codes.

- Supports multiple mail bar codes in an image.
- Supports various symbologies.
- Supports the 4 main bar code orientations, with a tolerance of 3°.
- Detects bars that are at least 3 pixels wide.

Workflow



4-state bar codes

A 4-state bar code is a special kind of bar code where data is encoded on the height and position of the bars rather than their width.

Each bar can have one of 4 possible states:

- Short and centered
- Medium and elevated
- Medium and lowered
- Full height



Mail bar code symbologies

The symbology of a mail bar code specifies how to decode the bar code and how to interpret its contents.

Every country uses its own flavor of mail bar code, or symbology. Some countries, like the US, even use multiple symbologies.

As of now, the Open eVision Mail Bar code Reader supports the following symbologies:

- US: PLANET, POSTNET and Intelligent Mail

- Japan: Japan Post

Mail bar code orientation

The Open eVision Mail Bar code Reader is designed to be used in mail-handling machines. As such it is optimized to handle the 4 main orientations you encounter in such machines:

- No Rotation: The mail barcode is horizontal and read from left to right
- Rotated 90° to the right: The mail barcode is vertical and read from top to bottom
- Rotated 90° to the left: The mail barcode is vertical and read from bottom to top
- Rotated 180°: The mail barcode is upside down, horizontal, and read from right to left.

For each of these orientations, an additional rotation of -3 to 3 degrees is allowed.

Checksum

Some symbologies specify the presence of a checksum in the bar code data.

This checksum is an additional character computed from all others encoded characters. It enables the reader to check the decoded character string coherence.

- The Mail Bar code Reader allows the user to verify or not the checksum for all enabled symbologies.
- By default, checksum is not controlled.
- To enable or disable checksum verification for all enabled symbologies, set the `ValidateChecksum` property.

Reading the mail bar codes in an image

To read all the mail barcodes in a given image:

1. Create an `EMailBarcodeReader` object.
2. Optionally, select the relevant symbologies using the `ExpectedSymbologies` property.
By default, Mail Bar code Reader will consider all supported symbologies.
3. Optionally, select the relevant orientations using the `ExpectedOrientations` property.
By default, Mail Bar code Reader will test all supported orientations.
4. Call `Read` on the source image or ROI.

Each mail bar code detected is returned as an `EMailBarcode` object.

5. Each `EMailBarcode` objects contains the following information:
 - The decoded string, using the `Text` property.
 - The decoded string, split up in semantic parts, using the `ComponentStrings` property.
 - The bar code orientation, using the `Orientation` property.
 - The bar code position, using the `Position` property.




US Intelligent Mail bar code with highlighted position and decoded information

Advanced parameters

The advanced parameters of the `EMailBarcodeReader` object are:

- `EnableDottedBarcodes` activates the support for dotted barcodes (barcodes whose bars are printed with dots).

By default, this property is set to false.



Dotted Mail Barcode

- `EnableClutteredBarcodes` activates the support for cluttered barcodes (barcodes in which some bars are connected).

By default, this property is set to true.



Cluttered Mail Barcode

- `ValidateChecksum` activates the validation of the bar codes checksums, if present.

By default, this property is set to false.

2. EasyMatrixCode - Reading Matrix Codes

2.1. EasyMatrixCode vs EasyMatrixCode2

Introducing EasyMatrixCode2

Important: *EasyMatrixCode2 is currently in beta release, so performance and features are not a complete image of the finished product.*

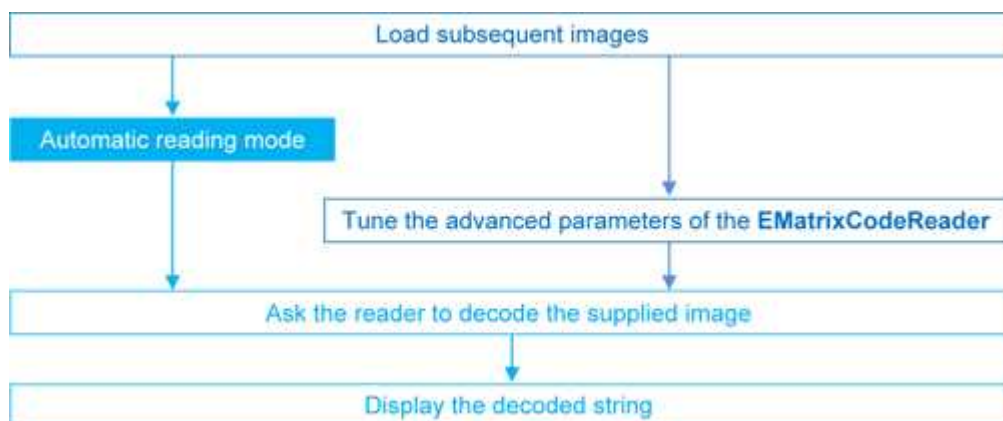
Starting with release 2.5, Open eVision introduces a new data matrix code reader functionality, named "EasyMatrixCode2". In the future, EasyMatrixCode2 should replace the current EasyMatrixCode (called "EasyMatrixCode1" from now on) with the following benefits:

- Ability to read multiple data matrix codes in an image.
- Improved consistency of reading and grading results.
- Improved consistency of processing time.
- Improved handling of deformed data matrix codes (this feature is not available yet and will be in a later release).

Note: *The relevant classes are stored in the EasyMatrixCode2 name space.*

Workflow

EasyMatrixCode1




EasyMatrixCode2

1. Load the image.
2. Read the data matrix codes in the image using `EMatrixCodeReader.Read()`.
3. Loop on found data matrix codes.
4. Display the decoded text.

2.2. Specifications

Data matrix code definition


- A data matrix code is a two-dimensional rectangular array of black and white cells which conveys a string of characters (digits, letters and special characters).
 - It is encoded to achieve maximum packing.
 - Each cell corresponds to a bit of information.
 - Additional redundant bits allow error correction for robust reading of degraded symbols.
- A data matrix code is located using the **Finder pattern**:
 - The bottom and left edges of a Data Matrix code contain only black cells.
 - The top and right edges have alternating cells.



- A data matrix code is characterized by:
 - Its **logical size** (number of cells).
 - Its **encoding type**: ECC 000 (odd symbol sizes, deprecated) or ECC 200 (even symbol sizes)..

Note: The data matrix code definition is provided by ISO/IEC and approved as standard ISO/IEC 16022.

Specifications



ECC 200, 26x26 cells data matrix code (left) and finder pattern (right)

In a single read operation, [EasyMatrixCode](#) locates, unscrambles, decodes, reads and grades the quality of grayscale 2D data matrix codes of any size, contrast, location and orientation (even viewed from the back on a transparent medium), providing they meet the following specifications:

- Minimum cell (= module) size: 3x3 pixels
- Maximum stretching (ratio between cell width and height): 2
- Minimum quiet zone (blank zone around the matrix code) width:
 - EasyMatrixCode1: 3 pixels
 - EasyMatrixCode2: 1 module

2.3. Reading a Matrix Code

Reading a matrix code

You can read matrix codes in an image automatically, using the [EMatrixCodeReader::Read\(\)](#) method. This method returns [EMatrixCode](#) instances, one per read data matrix code.

These instances include, among other things, for the found data matrix code:

- Its content (decoded string) ,
- Its position,
- Its logical size,
- Its encoding type,
- Methods to draw the data matrix code on the source image.

Note:

- *EasyMatrixCode1* reads a single data matrix code in the image.
- *EasyMatrixCode2* reads all the data matrix codes in the image at once.

Computing the print quality

To compute print quality indicators as defined by BC11, ISO 15415, ISO/IEC TR 29158 (formerly known as AIM DPM-1-2006) and SEMI T10-0701 standards, retrieve the grades with the `GetIso15415GradingParameters`, `GetIso29158GradingParameters` and `GetSemiT10GradingParameters` accessors of the `EMatrixCode` class.

Note:

- With `EasyMatrixCode1`, enable the grading before using the `Read` method.
- With `EasyMatrixCode2`, the grading is always enabled.

Reading GS1 data matrix codes

`EasyMatrixCode` is able to find and decode GS1-compliant data matrix codes.

The GS1 standard adds semantic identifiers to the contents of a data matrix code, so that they are interpreted in an easy and consistent way.

GS1-compliant content is structured as follows:

```
]d2 [GS1] {Id1} {Value1} [GS1] {Id2} {Value2}...
```

where:

- “[d2]” is the string identifying a GS1-compliant stream,
- [GS1] is the GS1 escape character (0x1d),
- {Id} is an application identifier,
- {Value} is the value associated to that identifier.

Example

The string:

```
]d2 [GS1] 11180112 [GS1] 15190101
```

can be interpreted as follows:

- It contains two GS1 parts, 11180112 and 15190101.
- The first (11180112) is composed of the identifier 11 and the value 180112, meaning that the product has a production date (the meaning of identifier 11) of January 12th, 2018.
- The second (15190101) is composed of the identifier 15 and the value 190101, meaning that the product has a best before date (the meaning of identifier 15) of January 1st, 2019.

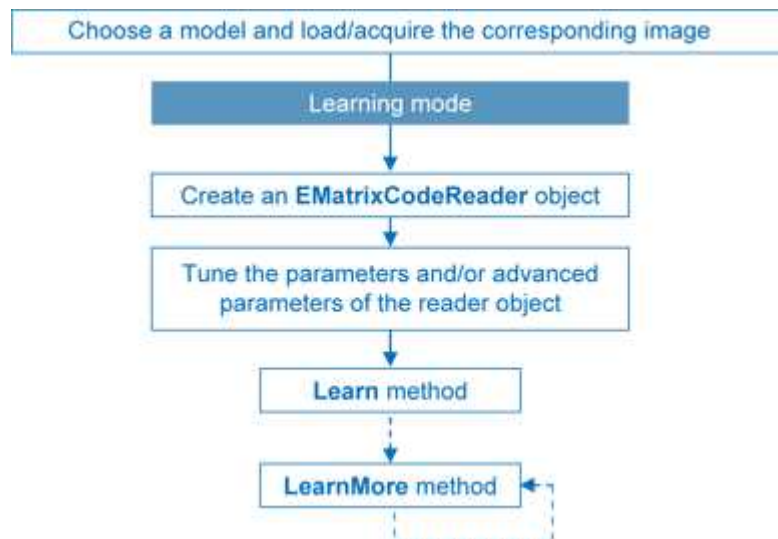
For more information, see <https://www.gs1.org/>

2.4. Learning a Matrix Code

Important: *This topic applies to EasyMatrixCode1 only.*

To search for specific features and speed up your processing, learn a Matrix code model.

Workflow



1. Load the image of the matrix code you want to learn.
2. Learn the model:
 - Use the `Learn` method with `Contrast`, `Family`, `Flipping`, `Logical Size` parameters.
 - If you need to learn several matrix codes, use `LearnMore` and pass additional sample images.
 - Call `Learn` to replace `EMatrixCodeReader` parameters (calling `Learn` several times does not accumulate results, while `LearnMore` does).
3. Tune `search parameters` to be efficient and either:
 - Read only matrix codes that match a sample matrix code,
 - Or read only matrix codes that have the same properties (`Contrast`, `Family`, `Flipping`, `Logical Size`) as the learned one,
 - Or disregard a search parameter of the learned matrix code `SetLearnMaskElement`, for example to read only unflipped matrixcodes. Just remove the default parameters, then add new ones.
4. Ask `EMatrixCodeReader` to decode the supplied image.
5. Display the `decoded string`.
6. Save the state of the reader object using `EMatrixCodeReader::Save`.

Restoring the state of an EMatrixCodeReader

To restore the state of an EMatrixCodeReader and use it to read a matrix code:

1. Load an image.
2. Restore the reader state from the given file using `EMatrixCodeReader::Load`.
3. Read the image.
4. Display the decoded string.

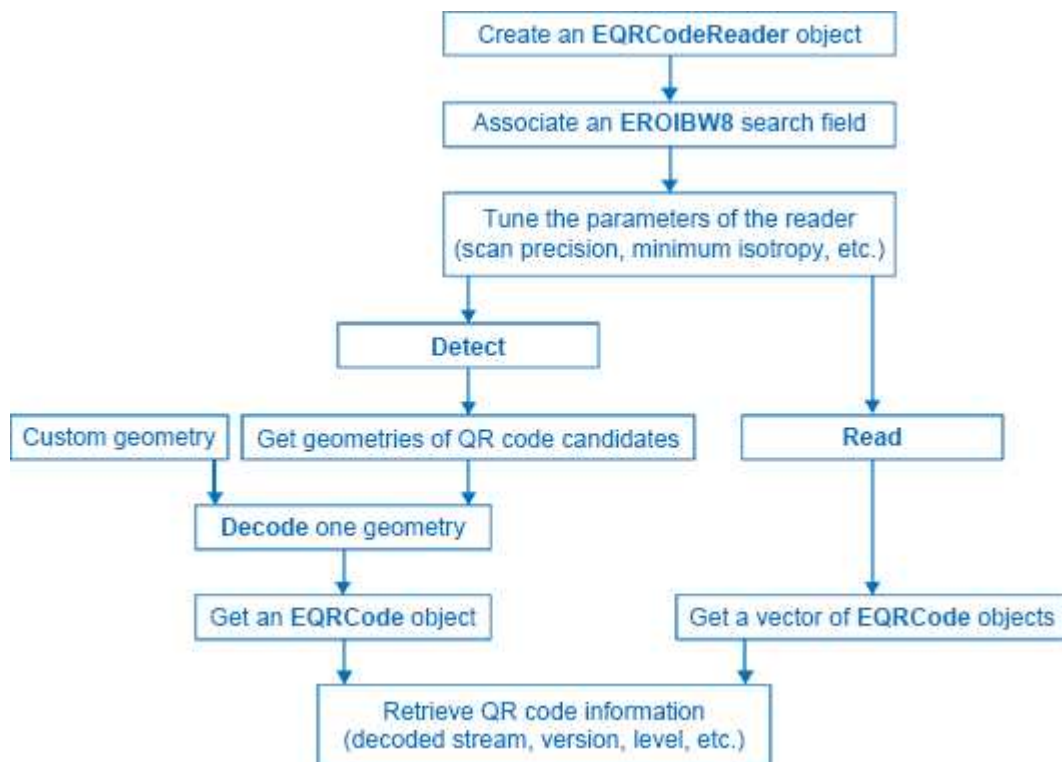
3. EasyQRCode - Reading QR Codes



EasyQRCode detects QR (Quick Response) codes in an image, decodes them, and returns their data.

Error detection and correction algorithms ensure that poorly-printed or distorted QR codes can still be read correctly.

Workflow



QR code definition

A QR code is a square array of dark and light dots. One dot (or "**module**") represents one bit of information.

QR codes contain various types of data and can be different models, versions, and levels. They always contain a message, metadata about alignment, size, format, and error correction bits. They comply with the international standard ISO/IEC 18004 (1, 2 and 2005).


QR code structure

The QR code symbol consists of an **encoding region**, containing data and error correction codewords, and of **function patterns**, containing symbol metadata and position data.


A QR code must be structured with the following elements:

- **Quiet zone:** blank margin around the QR code
- **Finder patterns:** recognizable zones identifying a QR code
- **Extension patterns:** markers for the alignment of the QR code (model 1)
- **Alignment patterns:** markers for the alignment of the QR code (models 2 and 2005)
- **Timing Patterns:** data giving the module size (in pixels)
- **Format information:** zones providing the QR code level
- **Version information:** data giving the QR code size, for instance 25 x 25 modules (models 2 and 2005)
- **Data contents and error correction codewords:** the primary information carried by the symbol, with additional information for error correction


Variants of this structure exist, according to the model, format, or version of the QR code. For instance, model 1 QR codes do not feature alignment patterns but extension patterns. Micro QR codes include only one finder pattern, and no alignment pattern. EasyQRCode can read all of them.



Structure of a model 1 QR code symbol



Structure of a QR code 2005 symbol



Structure of a Micro QR code symbol

Data types

The QR code data can be any mix of these types:

- Numeric data (0-9)
- Alphanumeric data (0-9, A-Z, /, \$, %, etc.)
- Byte data
- Kanji characters

Models (Standards)

- **Model 1:** original QR code international standard, with versions ranging from 1 to 14. Note that the "version" of a QR code is the symbol size (in number of modules). It does not relate to the version of the standard, which is called the "model".
- **Model 2:** improvement of model 1. It provides versions from 1 to 40. It defines alignment patterns to improve reading of distorted QR codes, or QR codes printed on curved surfaces.
- **Model 2005:** improvement of model 2, including white-on-black QR codes, and mirror symbol orientation.
- **Micro QR codes:** (not yet supported) smaller QR codes, from version M1 to version M4. They have been introduced to save printing space.

Versions (Symbol Size)

- **QR codes:** from version **1** (21 x 21 modules) to version **40** (177 x 177 modules), with an increment of +4 x +4 modules (version 2: 25 x 25 modules, version 3: 29 x 29 modules, ..., version 39: 173 x 173 modules).

- **Micro QR codes:** (not yet supported) version **M1** (11 x 11 modules), version **M2** (13 x 13 modules), version **M3** (15 x 15 modules), version **M4** (17 x 17 modules).



Examples of QR codes

From left to right:

Micro QR code, version M3, 15 x 15 modules,
 Model 2 QR code, version 4, 33 x 33 modules, 67-114 characters,
 Model 2 QR code, version 40, 177 x 177 modules, 1852-4296 characters

Levels (Error Correction)

QR codes contain error correction data. The standard offers the following levels of error correction:

- **L:** (low) about 7% of codewords can be restored
- **M:** (medium) 15%
- **Q:** (quality) 25%
- **H:** (high) 30% (not available for Micro QR codes)

QR code geometry

When the QR code reader finds an array of dots that could match a QR code, it returns the "geometry" of this QR code candidate.

A **QR code geometry** is a set of points. It contains the coordinates of the **corners** of the QR code **quadrilateral** (bottom left, top left, top right, bottom right), and the coordinates of the **finder pattern centers** (bottom left, top left, top right).



QR code geometry

Read a QR code

Reading a QR code returns information about [QR codes](#) for which [detection](#) and [decoding](#) were successful.

This is equivalent to detecting and decoding all QR codes in the given search field (see advanced features).

Advanced features

Detect a QR code

1. Set a [search field](#) on an [EROIBW8](#) image or tune the parameters to restrict the numbers of operations to process.
2. The QR code reader scans the image and searches for 3 finder patterns that could match a QR code, with the following requirements:
 - Minimum quiet zone (blank zone around the QR code) width: 3 pixels.
 - Minimum module size: 3 x 3 pixels.
 - Minimum isotropy: 0.5.
 - Maximum corner deformation: 15° (corner angles can range from 75° to 105°).
3. The reader returns QR code candidates, or the result of a [detection](#), as a vector of geometries.

decode a QR code

1. The QR code reader decodes a QR candidate and returns the [QR code: model](#), [version](#), [level](#), [geometry](#) and [decoded stream](#) of data.

The [decoded stream](#) class consists of a [coding mode](#) (basic, FNC1/GS1, or FNC1/AIM), and an [application indicator](#) (if the coding mode is FNC1/AIM, otherwise 0). The [decoded data](#) can be accessed from [each part of the decoded stream](#), according to its [encoding](#) (numeric, alphanumeric, byte, or Kanji). You can also get the [raw bit stream](#) (the bit data after unmasking and error correction, but before decoding as a vector of bytes).

2. The reader can report the amount of [unused error correction](#).
 - Close to 1, very few errors were corrected when decoding the data. The decoding is highly reliable, and the QR code is of good quality.
 - Close to 0, many errors were corrected when decoding the data. The decoding is reliable, but the QR code quality is poor.
 - -1, error correction failed. Decoding was not performed.

Tuning parameters

Scan precision: You can change the scan precision to scan the search field with a fine (recommended for small QR codes), or coarse (recommended on medium to large QR codes) precision.

Minimum score: The QR code reader searches for this QR code finder pattern:



A perfect match returns a pattern finder score of 1.

Less accurate matches return lower scores.

The minimum score allowed by default is 0.65 - you can tune this.

Minimum isotropy: The isotropy of a QR code represents its rectangular deformation.

Perfectly square QR codes have an isotropy of 1 (short side divided by long side, whether the rectangle is vertical or horizontal).

EasyQRCode can detect rectangle QR codes with an isotropy down to 0.5. The default **minimum isotropy** is 0.8, it can be tuned from 0 to 1.



Square and rectangular QR codes (isotropy = 1, 0.5, and 0.5 from left to right)

Model and version: The QR code reader searches for QR codes of all models, and all versions.

You can shorten the process by specifying the QR code **model(s)** and a range of versions (from 1 to 40) to be searched for.

PART VII
3D TOOLS

1. Understanding 3D Concepts

1.1. Basic Concepts

Easy3D

Easy3D is a set of tools for solving computer vision problem using 3D acquisition and processing. Easy3D supports laser line triangulation for fast and precise acquisition of depth maps.

Depth maps are gray scale images where each pixel represents a displacement in the third dimension. Because of the acquisition procedure, they are usually not dimensionally correct. So, while Open eVision 2D image operators are compatible with depth maps, you should not use them for processes requiring precise measurements.


Easy3D provides a calibration tool to generate corrected, metric point clouds and meshes from depth maps. Most 3D operators work on point clouds or meshes. The included export functions to the standard PCD file format allows integration with other 3D tools.

Easy3D also allows the computation of ZMaps. A ZMap is the projection of a point cloud on a given reference plane. Like depth maps, ZMaps are gray scale images, but are also dimensionally correct. As such, they can be used with all Open eVision 2D functions.

All the Easy3D tools are placed in the Easy3D namespace.

3D representation

Open eVision uses a right-handed cartesian 3D coordinate system. In this system, each 3D point is represented by its 3 coordinates X, Y and Z.



Open eVision provides different containers to store 3D objects :

- Depth maps
- Point clouds
- Meshes
- ZMaps

Depth map

A depth map is a way to represent a 3D object using a 2D grayscale image where each pixel (u, v) in the image contains a third coordinate as its gray value.




The grayscale values of a depth map do not necessarily represent a Z metric coordinate. In the context of a laser triangulation setup, these values represent the displacement of the laser line profile, which is not the physical height of the 3D surface.

A depth map contains a gray scale image coded on 8 or 16 bits. One specific gray value, called the undefined value, is reserved for the representation of invalid pixels. By default, this value is 0 for all depth map types.

The calibration process aims to convert the depth map representation to real, metric 3D representations such as point clouds or meshes.

Point cloud

A point cloud is a set of 3D points (x, y and z coordinates) representing the scanned object in the world metric space.



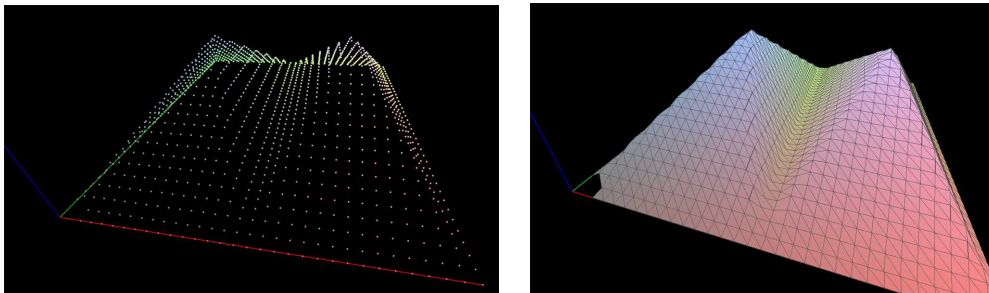
In addition to the calibration process included in Easy3D, point clouds can be produced using various 3D acquisition techniques, like stereo reconstruction or time of flight cameras.

Mesh

A Mesh is a geometric representation of a 3D surface, a set of connected 3D points.

In an `EMesh` object, 3 points are connected to define a triangle.

This kind of 3D representation is also called a "triangle mesh".



A point cloud and the corresponding mesh (displayed with Open eVision `E3DViewer`)

An `EMesh` object contains a point cloud and the indexes of the vertices of all mesh triangles.


`EMesh` uses a metric space representation that can be generated from a depth map and that can be used to produce a `ZMap`.

ZMap

ZMaps are another representation for 3D data.

- They are grayscale images like depth maps but represent metric and corrected 3D points.
- They are convenient representations for measurement and matching.
- They are compatible with most of the 2D processing functions.

ZMaps are generated by the projection of a point cloud or a mesh onto an arbitrary 3D plane.



A depth map and the corresponding ZMap

A ZMap contains an image in which each pixel value represents a positive distance from the reference plane.

Use the method `AsEImage()` to obtain a reference to the contained image.


A ZMap also contains the following information:

- The transformation from the World coordinates to the ZMap coordinates.
- The size of a pixel, called the "resolution".

Like in a depth map, a specific pixel value is reserved to represent invalid pixels. To get this pixel value, use the method `GetUndefinedValue()`.

1.2. Laser Triangulation

In a laser-line triangulation system, a laser line is projected on the object to measure. A camera is looking at the laser line from a different point of view. The line deformation observed by the camera contains the shape information of the measured object.



The scanning of the object consists in moving it under the laser line and recording multiple images.


From the scanning you can reconstruct its 3D shape.



Occlusions

Using the laser triangulation method, the laser may be unable to reach some parts of the object or the camera may be unable to view them. This is called occlusion.

- On the left illustration, the camera does not see the bottom of the hole, inducing camera occlusion.
- On the right illustration, the laser does not reach the bottom of the hole, inducing laser occlusion.




You can limit or avoid occlusions by using advanced scanning methods, for example by using two cameras or two lasers.

1.3. The Laser Line 3D Acquisition Pipeline


The 3D acquisition pipeline starts with the acquisition of a laser line profile and ends up with the point cloud, mesh or ZMap.

The source material for 3D processing is the depth map, coming from a Coaxlink Quad 3D-LLE or generated from a list of images. Two types of depth map are available, one for each different pixel coding scheme (8- or 16-bit).



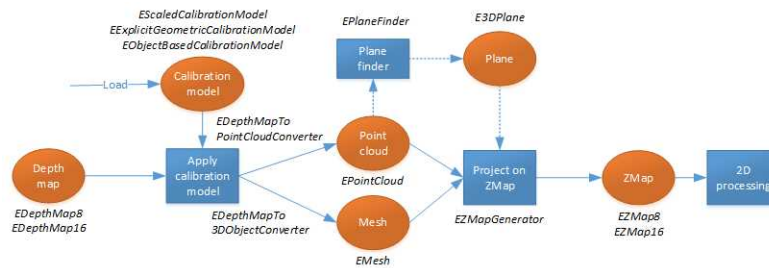
The generation of a depth map, from a hardware or a software source

Some processing methods can use the depth map directly, but most measurement and matching processes need metric, distortion-free representations. Calibration of the laser triangulation setup is therefore required. Calibration is used to turn the depth map into a point cloud or mesh expressed in a metric space that we call “world space”.



The generation of an object based calibration model, from a scan of the reference object

A point cloud is a list of 3D points, expressed in a world space coordinate system. The point cloud can be projected on a plane, producing a ZMap, which is a convenient and effective representation for 2D processing with a metric scale.



The workflow from the depth map to the ZMap

The following sections describe the classes and methods useful for a 3D workflow. The [Use Case - Measuring a Remote Controller](#) goes through this processing pipeline.

2. Easy3D - Using 3D Toolset

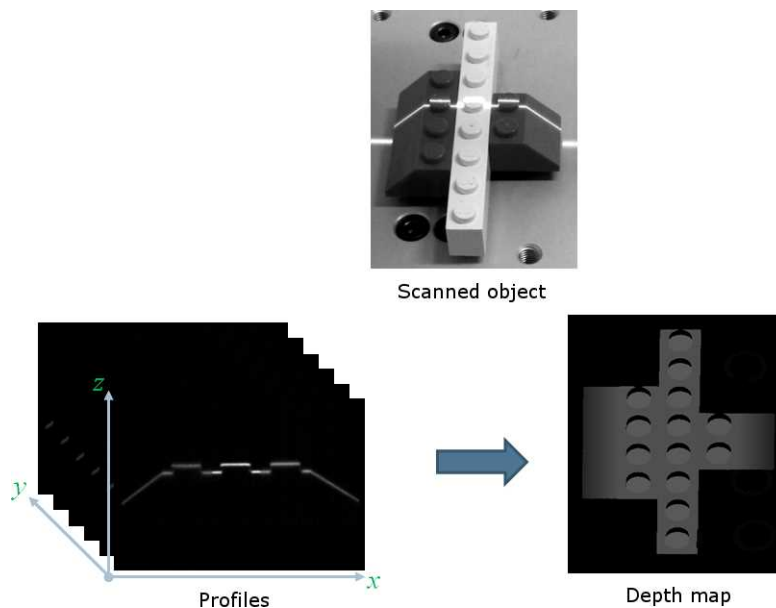
2.1. Laser Line Extraction

A Laser Line Extraction (LLE) algorithm is required to create a depth map from a sequence of profiles of the object captured by the camera sensor.

The objective of an LLE algorithm is to measure the line position along a vertical profile in every column of a sensor frame, within a user-defined region of interest (ROI).

For every step of the object position, the detection analyzes each column of a frame individually and produces a row of output positions, stored as gray values.

The figure below illustrates a depth map generation.




The `ELaserLineExtractor` class provides the laser line extraction functionality in Open eVision. It implements several algorithms to extract the laser line:

- **Maximum detection** returns the position of the pixel of maximum intensity. It's the fastest method but it doesn't support sub-pixel precision.
- **Peak detection** approach detects local maxima. If several maxima are detected, the one with the highest intensity is returned. The position is returned with sub-pixel precision.
- **Center of gravity** algorithm is suitable when the laser line is spread over several pixels. The position is returned with sub-pixel precision.

You can also set a threshold to exclude pixels with low intensity.


The line position returned by the laser line extraction algorithms is relative to the bottom of the region of interest. So, values in the depth map range from 0 (bottom of the ROI) to the height of the ROI.



Low-pass linear filter

Optionally, you can apply a low-pass linear filter in front of the line extraction in order to reduce noise and high frequencies in the image.

The low-pass filter applies a convolution operator on a 1 x 3 sliding window. The 3 elements of the convolution kernel (A, B and C) are configurable, accepting any positive integer. The figure below illustrates the positioning of the convolution kernel elements within a given ROI.



You can activate the low-pass filter for any of the laser line extraction methods with the method `ELaserLineExtractor::SetEnableSmoothing(true/false)`. Parameters A, B and C are set with `ELaserLineExtractor::SetSmoothingParameters(A, B, C)`.

2.2. Calibration

The calibration is used to apply the transformation between a depth map and a point cloud or a mesh.

There are 3 ways to setup this conversion:

- Apply a simple scale on the pixel coordinates of the depth map (`EScaleCalibrationModel` class)
- Use the explicit geometric model (`EExplicitGeometricCalibrationModel` class)
- Use the object-based calibration approach (`EObjectBasedCalibrationModel` class)

These models share the same base class `ECalibrationModel` and exposes the method `Apply()`, which is used to apply the conversion between a depth map pixel and a 3D point. It takes as input the coordinates of one point in a depth map and it returns the coordinates of the corresponding point in the 3D space.

The method `Apply` is not aware of the possible mirroring of the corresponding depth map and cannot make use of `EDepthMap::EAxisSystemType` (see below). If necessary (when the corresponding depth map is vertically mirrored) the y coordinates should be flipped before calling the `Apply` method.

- The class `EDepthMapToPointCloudConverter` generates a point cloud from a depth map, using one of the calibration models.
- The class `EDepthMapToMeshConverter` generates a mesh from a depth map, using one of the calibration models.

By convention:

- The origin of the referential is the lower-left corner of the depth map.
- The center of the first pixel at the lower-left corner is at $x = 0.5$ and $y = 0.5$.
- The center of the pixel at the upper-right corner is at $x = \text{width} - 0.5$ and $y = \text{height} - 0.5$ where width is the width of the depth map and height is its height.

Mirrored depth maps

By default, Easy3D considers that the origin of the 3D axis of the depth map is the bottom left of the internal image buffer, and the Y axis is pointing up. This means that the depth map image is not seen as vertically mirrored compared to the real world image of the scanned object.

Nevertheless, depending on your acquisition setup this mirroring can happen (for example if the direction of the scan is inverted).

If this is your case, you can set the `EDepthMap::EAxisSystemType` to `EAxisSystem_UpperLeftCorned`, meaning that the origin of the 3D axis is on the upper left corner and the Y axis is pointing down.

This value changes the behavior of the methods :

- `EObjectBasedCalibrationGenerator.Compute`

- `EDepthMapToPointCloudConverter.Convert`
- `EDepthMapToMeshConverter.Convert`

Scale calibration

The scale model (`EScaleCalibrationModel`) only applies a simple factor on the X, Y and Z axis. These factors are the only parameters of `EScaleCalibrationModel`.


For depth maps coming from laser triangulation setup, this transformation does not produce corrected, metric points. It's main use is to display depth maps as 3D data with the `E3DViewer` class.

Explicit geometric calibration

The explicit geometric model (`EExplicitGeometricCalibrationModel`) defines a simple and ideal laser triangulation setup. The explicit calibration makes some strong assumptions on the setup geometry and can only be used when a minimum set of parameters are known:


- The angles of the camera and the laser plane, in the counter clockwise direction. The camera angle must be positive.
- The height of the camera above the scanned object.
- The field of view of the camera defined by the sensor size (mm) and the optical focal length (mm).
- The physical distance between two line scans of the depth map (depends on acquisition rate and motion speed).
- The size of the image and the ROI origin used in laser line extraction (between the top (0) and the bottom (height) of the image).

Use the "[Easy3D_Setup_Configuration.xlsx](#)" spreadsheet to compute and check your setup configuration and parameters.



Explicit calibration setup with camera angle, laser angle and camera height

The setup of an explicit geometric calibration uses the constructor of the `EExplicitGeometricCalibration` class.




Object-based calibration

Object-based calibration gives real world, metric, coordinates from an arbitrary laser triangulation setup. From the scan of a reference object, the calibration process tries to calculate all the parameters required for the transformation to the world space (position and attributes of the camera, position of the laser plane, relative motion of the object, optical distortion...).


In the current version of Open eVision, 2 types of reference objects are supported, the double pyramid and the truncated pyramids. The calibration process operates only from acquisitions of such objects only.

Depending on the target object size, a calibration object must be manufactured as a scaled version of one of the reference objects described below.

CAD files in various standard formats are available on request.




First reference object: the double pyramid




Second reference object: the truncated pyramids

- The class `EObjectBasedCalibrationModel` is the container for the object based calibration model.
- The class `EObjectBasedCalibrationGenerator` performs the computation of such a model.
 - Its `Compute()` method takes a depth map as input and returns a calibration model.
 - The computation of the calibration model may take some time.
 - Use the `Save()` method to store the calibration model for later use.
- The real size of the calibration object is set relatively to the dimensions of the figure above, by the `SetCalibrationObjectScale()` method.
 - The dimensions of the figures above are expressed in millimeters.
 - For example, if the double pyramid calibration object has a physical width of 100 mm, the method should be called with a factor of 10 ($10 \times 10 = 100$ mm).
- The calibration object model must be set with the method `SetCalibrationObjectType()`. The possible values are `DoublePyramid` and `TruncatedDoublePyramid`.

If possible, prefer the truncated pyramids model as it does not show symmetry and thus prevents possible orientation confusion.



A manufactured double pyramid under the laser triangulation device and the corresponding captured depth map



Truncated pyramids of different sizes

Recommendations for the acquisition of the calibration object

To ensure a reliable and precise calibration, the following conditions must be satisfied during the acquisition of the calibration object:

- All the calibration object faces must be visible on the depth map (that is a constraint on the camera and laser orientation)
- The calibration object must be horizontal (relative to the motion axis)
- There must be no other object higher than the calibration object in the depth map
- The depth map must have at least 200 x 200 pixels resolution
- The calibration object must cover at least 50% of the defined pixel of the depth map

2.3. Point Cloud

Coordinates Transformations

Affine Transforms


Affine transforms allow you to reposition the point cloud inside the 3D space.

Open eVision provides you with the following basic transformations:

- Rotation around the X, Y or Z axis
- Translation along the X, Y and/or Z axis
- Scaling, around the origin, and either isotropic (the same in all directions) or anisotropic (different along the different axis)

It also provides you with projection transformations, both orthographic and perspective:

- An orthographic projection transforms a volume of space in the shape of a rectangular parallelepiped (and the points it contains) into the canonical view (a cubic space of size 2 and centered on the origin).



- A perspective projection transforms a volume of space in the shape of a frustum (basically a truncated pyramid) into the canonical view. This projection allow you to simulate the perspective effect given by an eye or a camera.

Reducing a Point Cloud

Cropping

Cropping allows you to exclude points from the point cloud based on geometrical considerations.

Open eVision provides the following cropping functions:

- `ESimpleCropper`: simple cropping on the X, Y and/or Z coordinates (aligned rectangle 3D region)
- `ERectangularCropper`: cropping the points outside (or inside) an oriented rectangular parallelepiped
- `ESphericalCropper`: cropping the points outside (or inside) a sphere.
- `EPlaneCropper`: cropping the points depending on their position with respect to a plane

These classes produces a new point cloud with the selected points.

Decimation

The random decimator, `ERandomDecimator`, decimates a point cloud by copying a specified number of points, randomly selected, to a new point cloud.

Specify the number of points to keep as parameter of the constructor.

```
EPointCloud pc;
pc.LoadPCD("c:\\images\\data.pcd");
// Explicitely decimate the point cloud
ERandomDecimator decimator(5000);
```

```
EPointCloud pcDecimated;
decimator.Decimate(pc, pcDecimated);
pcDecimated.SavePCD("c:\\images\\decimatedData.pcd");
```


Managing Planes

E3DPlane

A plane can be represented as an `E3DPlane` object.


This plane is characterized by:

- Its normal which is a vector of norm 1, perpendicular to the plane.
- Its signed distance from the origin, which is the smallest distance from the origin to the plane. The signed distance is positive when the vector binding the origin to the closest point on the plane has the same direction as the normal and is negative when it has the opposite direction.



Once a plane is defined, you can measure the signed distance between this plane and any point in the space (using the method `DistanceTo()`):

- A positive distance means that the vector connecting the plane to the point has the same direction as the normal.
- A negative distance means that the vector has the opposite direction.




EPlaneFinder

You can search for a plane in a point cloud using the object `EPlaneFinder` object.

The main parameters of this object are:

- The maximum distance between the searched plane and a point that belongs to this plane.
- The expected ratio between the numbers of inliers and the total number of points in the point cloud.
 - An **inlier** is a point that belongs to a plane (closer than this maximum distance).
 - An **outlier** is a point that is not an inlier.

The picture below illustrates how points of the space are classified as inliers (in green) and outliers (in red) according to their distance to the searched plane.

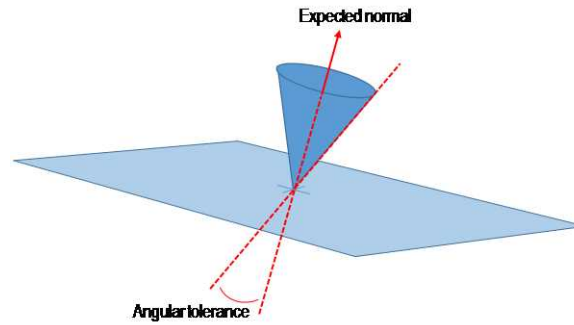


A `EPlaneFinder` object produces a `E3DPlane` object. The algorithm searches for a plane containing as many inliers as possible. This plane is the biggest plane if the samples are evenly distributed.

The maximum distance between the plane and the inliers is a mandatory parameter: it should include the deviation due to the noise but also take warpage into account.

The parameter that specifies the ratio of inliers with respect to the total number of points has a default value of 0.3, meaning that we estimate that about 30% of the points belong to the plane. This parameter is not as critical as the maximum distance but it affects the maximum time the algorithm will spend in searching a plane as well as its robustness.

Optionally, you can specify the expected normal vector to the plane to search. In that case, you should also specify an angular tolerance with respect to this expected direction.



When an expected normal is specified, the algorithm only searches for a plane that satisfies the condition. Setting this condition might speed up the plane search.

Finally, it is important to note that, by default, the `EPlaneFinder` decimates the input point cloud to accelerate the search. The default decimator reduces the input point cloud to 10000 points. Alternatively, you can disable this decimation, or you can decimate a point cloud explicitly, by using an `ERandomDecimator` object and use the decimated point cloud as input for the `EPlaneFinder`. In this case you should disable the default decimator.

Once the main plane is found, a fit is done on all the inliers points and the result is returned (see `EPlaneFitter` below).

EPlaneFitter

The `EPlaneFitter` operator computes a fit on all the points of a point cloud and returns a `E3DPlane` object.

Aligning


EPrincipalAxisExtractor

The `EPrincipalAxisExtractor` computes the “principal axis” of an object from a point cloud (`EPointCloud`) and returns a `E3DTransformMatrix` containing a solid transformation that defines a new orthogonal basis.

This new orthogonal basis has the following characteristics:

- The center is the center of gravity of the point cloud.
- The axis are oriented along the “principal axis” of the object. This is the result of the “PCA” calculation (principal axis analysis).
- The directions of the axis are selected so that the new basis is as close as possible of the basis defined by the reference transformation.

The next figure illustrates the orientation of the principal axis of an object.




The principal axis extraction is done using the `Extract()` method that takes a `EPointCloud` as input and returns an `E3DTransformMatrix`. Optionally, you can pass 3 other output parameters by reference to retrieve the value of the standard deviation along the 3 principal axis.

You can use the returned `E3DTransformMatrix` object to transform the 3D coordinates of a point. For example, apply the transformation matrix to the origin (0, 0, 0) to return the center of gravity of the object.

Specification of a reference transformation

The reference transformation is an optional parameter of the `EPrincipalAxisExtractor` object. It defines a reference basis used to select an orthogonal basis out of the principal axis. The selected basis will be the closest to the reference basis.

*If no reference transformation was supplied, the default reference basis is $((0, 0, 1), (0, 1, 0), (0, 0, 1))$, that corresponds to the identity transformation. On the figure below, the default reference basis determines the direction of the axis **ex**, **ey** and **ez**.*



EFeaturesAligner


A `EFeaturesAligner` object finds the best transformation that maps a list of points to another list of points.

- The first list of points is called the "model". It is stored in the `EFeaturesAligner` object.
- The second list of points is called "measured points". It is passed as a parameter to the `Compute()` method. If successful, the result of this method is a `E3DTransformMatrix` object.
- The 2 lists should form matching pairs. In other words, the first point of the first list matches the first point of the second list, the second point of the first list matches the second point of the second list, and so on...

With the Polarity parameter, you can define which transformation is returned. It can be either:

- The one that moves one point from the first list (the model) to the second list of points (the measured points) if the polarity parameter is set to `EAlignmentPolarity_ModelToMeasured` (default).
- The one that moves a point from the second list (the measured points) to the first list (the model) if the polarity parameter is set to `EAlignmentPolarity_MeasuredToModel`.

The figure below illustrates the computation of the alignment transformation. In this example a model is aligned to an object using the coordinates of their corners.



Once the transformation is computed, use the method `GetOrthoBasis` of the `E3DTransformMatrix` object to get the basis (**ex**, **ey**, **ez**) and the center point **t** that defines the new basis.

You can also apply the computed transformation on any 3D point as illustrated in the code below.

```
EFeaturesAligner alignTool;
E3DTransformMatrix alignBase;
E3DPoint ex, ey, ez, t;
std::vector<E3DPoint> model3d;
std::vector<E3DPoint> points3d;
// add points to model3d and points3d
// ...
alignTool.SetModelPoints(model3d);
alignBase = alignTool.Compute(points3d);
// Get the orthogonal basis and store it in ex, ey, ez and t
alignBase.GetOrthoBasis(ex, ey, ez, t);
// Applying the transformation on point P1, results in point P1b
E3DPoint P1 = E3DPoint(...);
E3DPoint P1b = alignBase*P1;
```

As you can see, the application of the transformation on a point is simply done by multiplying the transformation matrix by the point (as done in the example above).

On the other hand, if you need to transform a point cloud or a list of points, it is more efficient to use the `ApplyTransform()` method of an `EAffineTransformer` object.

2.4. Mesh

A mesh is a geometric representation of a 3D surface. The surface is defined by a triangle mesh connecting the 3D points. Like a point cloud, a mesh is expressed in the metric space.

Like a point cloud, you can generate a mesh from a depth map and use it to produce a ZMap.

Generation

An `EMesh` object is generated from a depth map using the `EDepthMapToMeshConverter` class.

Like `EDepthMapToPointCloudConverter`, this class uses a calibration model to transform the depth map pixels to 3D world positions. In addition, the depth map pixel connectivity is used to build the triangle mesh. Adjacent pixels produce surface triangles.

Use `SetCalibrationModel()` to select a calibration model and the method `Convert()` to generate an `EMesh` from an 8 bits or 16 bits depth map.

Access and usage

In an `EMesh` object the 3D world positions are stored as an `EPointCloud` (accessible through the method `GetPointCloud()`). The triangle mesh is stored as an array of point indexes, where 3 consecutive indexes define a triangle. The method `GetTriangleIndexes()` provides a read-only access to the triangle mesh.

You can use either the Open eVision proprietary format to save and load `EMesh` objects using the `Save()` and `Load()` methods, or use the STL standard file format ([https://en.wikipedia.org/wiki/STL_\(file_format\)](https://en.wikipedia.org/wiki/STL_(file_format))) using the `SaveSTL()` and `LoadSTL()` methods which respectively write to and read from ASCII STL files.


You can use an `EMesh` to produce a ZMap (see "[Generating a ZMap](#)" on the next page). Because an `EMesh` represents a surface, the so generated ZMap can show better continuity and less undefined pixels.

2.5. ZMap

Generating a ZMap

A ZMap is the projection of a point cloud or a mesh on a reference plane, with the distance coded as gray scale values:

- They are grayscale images, compatible with all Open eVision 2D libraries.
- They are distortion free, with affine transformation from/to metric coordinate system.



A depth map (left) and the corresponding ZMap (right), with default generation parameters and undefined pixel filling enabled

All Open eVision 2D processing are available on ZMaps: filtering, thresholding, blob extraction, measuring with EasyGauge, model matching with EasyFind or EasyMatch...

The `EZMapGenerator` class implements the conversion from a point cloud or a mesh to a ZMap. With all parameters at default value, the `Convert()` method automatically chooses the projection plane, the orientation, the map size and the resolution.


Several methods are available to further control the conversion:

- `SetReferencePlane()` defines a world space projection plane. The values of the ZMap pixels are the distance of the point cloud to that reference plane.
By default, the reference plane crosses the origin and is perpendicular to the world Z axis. The plane is defined as a `E3DPlane` object.
- `SetOrientationVector()` sets a world space vector representing the expected direction of the X (width) axis of the ZMap.
The orientation vector allows to “rotate” the object around the normal of the reference plane.
- `SetOrigin()` specifies the world position that is on the ZMap lower left pixel (0, 0).
- `SetMapSize()` defines the resolution (number of pixels in X and Y axis) of the generated ZMap.

- `SetMapXYResolution()` adjusts the X and Y resolution of the ZMap pixels, in world space unit per pixel (for example mm/pixel). This value is used to compute the ZMap size (width and height), depending on the projected size of the point cloud on the reference plane.
- `SetMapZResolution()` sets the Z resolution, in world space unit per pixel unit (gray value). The Z resolution is used to compute the transformation of the distance to the reference plane to the integer 8, 16 or 32 bits pixel value.
- `EnableFillMode()` and `SetFillMode()` control the options used to fill the "hole" in the ZMap. A hole exists when no 3D point is projected in the ZMap at a pixel position.

The methods `SetReferencePlane()`, `SetOrientationVector()` and `SetOrigin()` are used to setup the transformation between the world space and the ZMap space. This transformation is rigid (distances are kept).

Alternatively, it is possible to directly set that transformation with the method `SetWorldToZMapTransform()` using a rigid matrix as parameter. In that case, the reference plane, the orientation vector and the origin parameters are ignored.




The projection of a point cloud on a ZMap, showing 3 coordinate systems: the world space, the ZMap space and the pixel space.

The method `Convert()` performs the effective projection of a point cloud (`EPointCloud`) or a 3D object (`EMesh`) to the 8 bits and 16 bits ZMap.

When generating a ZMap from a point cloud, only individual points are projected on the ZMap. Depending on the point cloud density and the ZMap resolution, some regions of the ZMap may remain “undefined”. To get around this problem, adjust the resolution of the ZMap (`SetMapXYResolution` method) to remove “holes” on the ZMap.

By default, the point cloud to ZMap converter performs a filling algorithm. This process tries to replace undefined pixels with locally interpolated values.



Left: high resolution ZMap, the pixel scale exceeds the point cloud density
Center: the same generator parameters with the filling enabled
Right: a reduced ZMap scale/resolution, without filling

As a mesh defines a surface, its triangles are projected onto the ZMap plane. Thus, the generated image shows better continuity and less undefined pixels. However, the generation of a ZMap from an `EMesh` is slower than from an `EPointCloud`.

Managing the Coordinates


Coordinate systems on a ZMap

A ZMap has multiple coordinate systems:

- The **world space** system is the original, metric space from which the ZMap has been generated. Point clouds and meshes are expressed in the world coordinate system.
- The **ZMap space** is defined by a rigid transformation of the **world space**. The basis linked to this transformation is attached to the lower left corner of the ZMap.
- The **image space** is the system attached to the image representation of the ZMap. Its origin is the upper left corner of the ZMap and its unit length is one pixel along the X and Y axis.

The transformations between:

- The **image space** and the **ZMap space** include a scale factor.
- The **ZMap space** and the **world space** are solid transformations.



EZMap

The `EZMap` object exposes a set of methods to convert coordinates between world, ZMap and image spaces:

- `ImageToZMap` converts a 2D position in the image to ZMap coordinates.
- `ZMapToImage` is the reciprocal operation and converts a ZMap position to an image position.
- `ZMapToWorld` is a method to transform positions from the 3D ZMap space to the 3D world space. The world space is the original point cloud or mesh space.
- `WorldToZMap` is the reciprocal operation, converting from world space to ZMap.
- `ImageToWorld` and `WorldToImage` combine the functions above to transform directly from image space to world space (or the other way).

These methods only perform geometric transformations between the various coordinate systems and do not access the actual ZMap gray scale values.

The functions that access the pixel values are:

- `GetWorldPositionFromPixelPosition()` is a method transforming the actual pixel value at integer position (u, v) to the original world space. This method queries the ZMap internal representation to get the pixel value w and transform the pixel space (u, v, w) coordinates to a world space position.
- `GetPixelPositionFromWorldPosition()` is a method to get a pixel value from a world position. The world position is projected on the ZMap and the pixel value is returned. If the world position is outside the ZMap domain, the method returns `FALSE`.

Static Methods

EFilters class


The `EFilters` class contains static methods used to apply filters to ZMaps or depth maps.

RemoveNoise

The `RemoveNoise()` method removes outliers from a depth map or a ZMap.

- It takes a depth map or a ZMap as input and generates a depth map or a ZMap respectively. The undefined points are not taken into account.
- It is based on a square moving kernel. The size of the kernel is $(2 \times \text{halfKernelSize} + 1)$ where `halfKernelSize` is a parameter of the method.
- The `threshold` parameter is scaled with regard to the Z resolution of the filtered depth map or ZMap.
- There are 3 variations of this filter, depending on the `method` parameter:
 - `ENoiseRemovalMethod_AbsoluteDifferenceFromMean` removes a point when it deviates from the average in the neighborhood, including itself. The threshold is an absolute difference.
 - `ENoiseRemovalMethod_RelativeDifferenceFromMean` removes a point when it deviates from the average in the neighborhood, including itself. The threshold is a multiple of the standard deviation.
 - `ENoiseRemovalMethod_HighStandardDeviation` removes a point when the standard deviation in the neighborhood, including itself, is higher than a defined threshold.

Example: Removing points showing a high standard deviation



The code below removes pixels with a standard deviation higher than a defined threshold.

```
// Load the ZMap data
EZMap16 zmap;
zmap.Load(...);
// Compute the filtered ZMap. The new ZMap is called filteredZmap
// The size of the kernel is 7x7, the threshold is 30.0
EZMap16 filteredZmap;
filteredZmap.SetSize(zmap);
EFilters::RemoveNoise(zmap, filteredZmap, ENoiseRemovalMethod_HighStandardDeviation, 3, 30.0, 0.0);
```

EStatistics class

The `EStatistics` class contains static methods used to compute statistics on ZMaps or depth maps.

ComputeAverageMap

The `ComputeAverageMap()` method computes the local average map.

- You can use this method as a low-pass filter.
- Undefined points are not taken into account.
- This method is based on a square moving kernel. The size of the kernel is $(2 \times \text{halfKernelSize} + 1)$ where `halfKernelSize` is a parameter of the method.

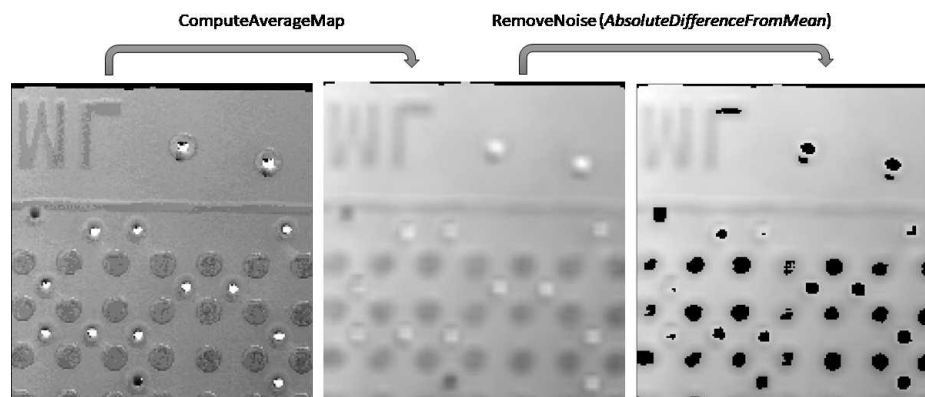
ComputeStandardDeviationMap

The `ComputeStandardDeviationMap()` method computes a map of the local standard deviation.

- You can use this method to determine visually the threshold value to use with the `RemoveNoise()` method when using the `ENoiseRemovalMethod_HighStandardDeviation` setting.

Note: Be aware, however, that in the generated map, a pixel with the value 0 can either be undefined or have a standard deviation equal to zero.

Example: Using a low pass filter on a ZMap, then removing points showing a deviation larger than a defined threshold



The code below first applies a low pass filter, then removes from the result the pixels showing a deviation from the neighborhood larger than the defined threshold.

```
// Load the ZMap data
EZMap16 zmap;
zmap.Load(...);
// Compute the filtered ZMap. The new ZMap is called averagedZMap
// The size of the kernel is 7x7, the threshold is 30.0
EZMap16 averagedZMap;
averagedZMap.SetSize(zmap);
EStatistics::ComputeAverageMap(zmap, averagedZMap, 3, 0.2);
// Compute the filtered ZMap. From averagedZMap, compute filteredZMap
// The size of the kernel is 31x31, the threshold is 20.0
EZMap16 filteredZMap;
filteredZMap.SetSize(zmap);
EFilters::RemoveNoise(averagedZMap, filteredZMap, ENoiseRemovalMethod_AbsoluteDifferenceFromMean,
15, 20.0, 0.2);
```

ComputePixelStatistics

The `ComputePixelStatistics()` method returns basic statistical information about pixel values:

- Minimum
- Maximum
- Average
- Standard deviation
- Number of valid (not undefined) pixels).

Use an `ERegion` object to specify the region of the ZMap or depth map used to compute the statistics.

ComputeStatistics

The `ComputeStatistics()` method returns the same information as the `ComputePixelStatistics()` method, but scaled with respect of the Z resolution.

Use an `ERegion` object to specify the region of the ZMap or depth map used to compute the statistics.

2.6. 3D Viewer


Use the `E3DViewer` class to easily create an interactive 3D display. The viewer displays point clouds, meshes and ZMaps.

You can create `E3DViewer` as a child of an existing window or without a parent. In that last case, a new window is created.

Note: As `E3DViewer` uses `OpenGL` interface, it requires a compatible display device.

Call the `ConfigureRenderSource()` method with a valid 3D geometry to display it. At each call, `ConfigureRenderSource()` replaces the current displayed object. The supported classes are `EPointCloud`, `EMesh` and `EZMap8/EZMap16`.

When you configure a new render source with `ConfigureRenderSource()`, the view point is automatically adapted to display the whole object.




E3DViewer in action: point cloud display (left) and 3D object display (right)

To display the geometry in false colors:

- Use the `GenerateColors()` method that computes RGB colors from the position of the vertices.
- It supports various predefined color ramps.
- Use the `SetColors()` method to use custom colors (one `EC24` entry is requested for each render source vertex).

Use the methods `SetPointSize()`, `SetWireframeMode()` and `SetRenderDecimationLevel()` to adjust the rendering attributes.



E3DViewer in action: wire frame enable (left) and `HueFromZ` color ramp (right)

In the 3D navigation window, use the mouse as follows:

- Press the left button to rotate the image horizontally and vertically.
- Press the right button to translate the image horizontally and vertically.
- Use the wheel to zoom in and out.

In addition, use the following keys:

- Press **R** to reset the viewer.
- Press **W** to show or hide the triangle edges (in wire frame mode).
- Press **+** and **-** to increase or decrease the point size.

Use methods `SetViewTarget()`, `SetViewingAngle()` and `SetViewingDistance()` to change the view point programmatically.