




the Power of Machine Vision

Solution Guide I

Basics



How to use HALCON's machine vision methods, Version 10.0.4

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without prior written permission of the publisher.

Edition 1	June 2007	(HALCON 8.0)
Edition 1a	October 2007	(HALCON 8.0.1)
Edition 1b	April 2008	(HALCON 8.0.2)
Edition 2	December 2008	(HALCON 9.0)
Edition 2a	June 2009	(HALCON 9.0.1)
Edition 2b	March 2010	(HALCON 9.0.2)
Edition 3	October 2010	(HALCON 10.0)
Edition 3a	April 2011	(HALCON 10.0.1)
Edition 3b	November 2011	(HALCON 10.0.2)

Copyright © 2007-2013 by MVTec Software GmbH, München, Germany



Protected by the following patents: US 7,062,093, US 7,239,929, US 7,751,625, US 7,953,290, US 7,953,291, US 8,260,059, US 8,379,014. Further patents pending.

Microsoft, Windows, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Microsoft .NET, Visual C++, Visual Basic, and ActiveX are either trademarks or registered trademarks of Microsoft Corporation.

All other nationally and internationally recognized trademarks and tradenames are hereby recognized.

Contents

1	Guide to HALCON Methods	5
2	Image Acquisition	13
3	Region Of Interest	19
4	Blob Analysis	31
5	1D Measuring	49
6	Edge Extraction (Pixel-Precise)	63
7	Edge Extraction (Subpixel-Precise)	73
8	Contour Processing	83
9	Matching	97
10	3D Matching	127
11	Variation Model	143
12	Classification	155
13	Color Processing	173
14	Texture Analysis	189
15	Bar Code	211
16	Data Code	233
17	OCR	245
18	Stereo Vision	277
19	Visualization	289

Chapter 1

Guide to HALCON Methods

This manual introduces you to important machine vision methods. To guide you from your specific application to the sections of the documentation to read, this section lists common application areas and the methods used for them.

Generally, a lot of applications use the following methods:

- [Image Acquisition](#) on page 13 for accessing images via an image acquisition device or via file.
- [Visualization](#) on page 289 for the visualization of, e.g., artificially created images or results of an image processing task.
- [Region of interest](#) on page 19 for reducing the search space for a following image processing task.
- Morphology (Reference Manual, chapter “[Morphology](#)”), e.g., for the elimination of small gaps or protrusions from regions or from structures in gray value images.

Other methods are more specific and thus are suited for specific application areas. Additionally, some application areas are part of another application area. To make the relations more obvious, for the following application areas the corresponding methods and related application areas are listed:

- Color Inspection ([page 6](#))
- Completeness Check ([page 6](#))
- Identification ([page 6](#))
- Measuring and Comparison 2D ([page 7](#))
- Measuring and Comparison 3D ([page 7](#))
- Object Recognition 2D ([page 8](#))
- Object Recognition 3D ([page 8](#))

- Position Recognition 2D ([page 8](#))
- Position Recognition 3D ([page 9](#))
- Print Inspection ([page 9](#))
- Quality Inspection ([page 10](#))
- Robot Vision ([page 10](#))
- Security System ([page 10](#))
- Surface Inspection ([page 11](#))
- Texture Inspection ([page 11](#))

1.1 Color Inspection

For color inspection, see the descriptions for [Color Processing](#) on [page 173](#).

1.2 Completeness Check

Completeness checks can be realized by different means. Common approaches are:

- Object and position recognition 2D/3D (see [page 8](#) ff), which is suitable, e.g., when inspecting objects on an assembly line.
- [Variation Model](#) on [page 143](#), which compares images containing similar objects and returns the difference between them considering a certain tolerance at the object's border.

1.3 Identification

Dependent on the symbols or objects you have to identify, the following methods are suitable:

- Identify symbols or characters
 - [Bar Code](#) on [page 211](#)
 - [Data Code](#) on [page 233](#)
 - [OCR](#) on [page 245](#)
- Identify general objects
 - Object and position recognition 2D/3D (see [page 8](#) ff)

1.4 Measuring and Comparison 2D

For measuring 2D features in images, several approaches are available. In the Solution Guide III-B, [section 4.1](#) on page 29, a graph leads you from the specific features you want to measure and the appearance of the objects in the image to the suitable measuring approach. Generally, the following approaches are common:

- [Blob Analysis](#) on page 31 for objects that consist of regions of similar gray value, color, or texture.
- [Contour Processing](#) on page 83 for objects that are represented by clear-cut edges. The contours can be obtained by different means:
 - [Edge Filtering](#) on page 63 if pixel precision is sufficient.
 - [Edge and Line Extraction](#) on page 73 if subpixel precision is needed.
- [Matching](#) on page 97 for objects that can be represented by a template. Matching comprises different approaches. For detailed information about matching see the [Solution Guide II-B](#).
- [1D Measuring](#) on page 49 if you want to obtain the positions, distances, or angles of edges that are measured along a line or an arc. More detailed information can be found in the [Solution Guide III-A](#).

1.5 Measuring and Comparison 3D

For measuring in 3D, the following approaches are available:

- The approaches used for measuring and comparison 2D (see [page 7](#)) in combination with a camera calibration (see Solution Guide III-C, [section 3.2](#) on page 40) for measuring objects that are viewed by a single camera and that lie in a single plane.
- Pose estimation (Solution Guide III-C, [chapter 4](#) on page 75) for the estimation of the poses of 3D objects that are viewed by a single camera and for which knowledge about their 3D model (e.g., known points, known circular or rectangular shape) is available.

3D reconstruction is an important subcategory of 3D measuring and comprises the following methods:

- Stereo for measuring in images obtained by a binocular or multi-view [stereo system](#) on page 277. Further information can be found in the Solution Guide III-C, [chapter 5](#) on page 105.
- Laser triangulation using the sheet-of-light technique (Solution Guide III-C, [chapter 6](#) on page 139) for measuring height profiles of an object by triangulating the camera view with a projected light line.
- Depth from focus (Solution Guide III-C, [chapter 7](#) on page 157) for getting depth information from a sequence of images of the same object but with different focus positions.
- Photometric Stereo (Reference Manual, chapter “[3D Reconstruction ▸ Photometric Stereo](#)”) for getting information about an object’s shape because of its shading behavior (e.g., by the operator [phot_stereo](#)).

1.6 Object Recognition 2D

For finding specific objects in images, various methods are available. Common approaches comprise:

- [Blob Analysis](#) on page 31 for objects that are represented by regions of similar gray value, color, or texture.
- [Contour Processing](#) on page 83 for objects that are represented by clear-cut edges. The contours can be obtained by different means:
 - [Edge Filtering](#) on page 63 if pixel precision is sufficient.
 - [Edge and Line Extraction](#) on page 73 if subpixel precision is needed.
- [Matching](#) on page 97 for objects that can be represented by a template. Matching comprises different approaches. For detailed information about matching see the [Solution Guide II-B](#).
- [Classification](#) on page 155 for the recognition of objects by a classification using, e.g., Gaussian mixture models, neural nets, or support vector machines. For more detailed information about classification see the [Solution Guide II-D](#).
- [Color Processing](#) on page 173 for the recognition of objects that can be separated from the background by their color.
- [Texture Analysis](#) on page 189 for the recognition of objects that can be separated from the background by their specific texture.
- Movement detection (see [section 1.13](#) on page 10) for the recognition of moving objects.

1.7 Object Recognition 3D

For the recognition of 3D objects that are described by a 3D Computer Aided Design (CAD) model, see the descriptions for [3D Matching](#) on page 127. For the recognition of planar objects that can be oriented arbitrarily in the 3D space, see the descriptions for perspective, deformable matching and descriptor-based matching in the chapter about [Matching](#) on page 97 for the uncalibrated case and in the [Solution Guide III-C, chapter 4](#) on page 75 for the calibrated case.

1.8 Position Recognition 2D

In parts, the approaches for measuring 2D features are suitable to get the position of objects. In the [Solution Guide III-B, section 4.1](#) on page 29, a graph leads you from several specific features, amongst others the object position, and the appearance of the objects in the image to the suitable approach. For position recognition, in particular the following approaches are common:

- [Blob Analysis](#) on page 31 for objects that consist of regions of similar gray value, color, or texture.

- [Contour Processing](#) on page 83 for objects that are represented by clear-cut edges. The contours can be obtained by different means:
 - [Edge Filtering](#) on page 63 if pixel precision is sufficient.
 - [Edge and Line Extraction](#) on page 73 if subpixel precision is needed.
- [Matching](#) on page 97 for objects that can be represented by a template. Matching comprises different approaches. For detailed information about matching see the [Solution Guide II-B](#).
- [1D Measuring](#) on page 49 if you want to obtain the positions of edges that are measured along a line or an arc. More detailed information can be found in the [Solution Guide III-A](#).

1.9 Position Recognition 3D

For the position recognition of 3D objects, the following approaches are available:

- The approaches used for measuring and comparison 2D (see [page 7](#)) in combination with a camera calibration (see [Solution Guide III-C, section 3.2](#) on page 40) for measuring objects that are viewed by a single camera, and which lie in a single plane.
- Pose estimation ([Solution Guide III-C, chapter 4](#) on page 75) for the estimation of the poses of 3D objects that are viewed by a single camera and for which knowledge about their 3D model (e.g., known points, known circular or rectangular shape) is available.
- Stereo for measuring positions in images obtained by a [binocular stereo system](#) on page 277 or a multi-view stereo system. Further information can be found in the [Solution Guide III-C, chapter 5](#) on page 105.
- [3D matching](#) on page 127 for objects that are searched for based on a 3D Computer Aided Design (CAD) model.

1.10 Print Inspection

For print inspection, the suitable method depends on the type of print you need to inspect:

- Optical character verification (Reference Manual, chapter [“Inspection > OCV”](#)) for the verification of characters.
- [Data Code](#) on page 233 for inspecting the quality of a printed 2D data code symbol. Further information can be found in the [Solution Guide II-C, section 6](#) on page 51.
- [Variation Model](#) on page 143, which compares images containing similar objects and returns the difference between them considering a certain tolerance at the object’s border.
- Component-based matching (part of [Matching](#) on page 97), if the print is built by several components which may vary in their relation to each other (orientation and distance) within a specified tolerance.

1.11 Quality Inspection

How to inspect the quality of an object depends on the features describing the quality. The following application areas and the methods used by them can be applied for quality inspection:

- Surface inspection (see [page 11](#)) if, e.g., scratches in an object's surface have to be detected.
- Completeness check (see [page 6](#)) if you need to check an object for missing parts.
- Measuring and comparison 2D/3D (see [page 7](#) ff) if an object has to fulfill certain requirements related to its area, position, orientation, dimension, or number of parts.

Additionally, [Classification](#) on [page 155](#) can be used to check the color or texture of objects. For more detailed information about classification see the [Solution Guide II-D](#).

1.12 Robot Vision

For robot vision, you can combine the approaches for object and position recognition 2D/3D (see [page 8](#) ff) with the hand-eye calibration described in the [Solution Guide III-C](#), [chapter 8](#) on [page 173](#).

1.13 Security System

For a sequence analysis or movement detection you can use, e.g., one of the following approaches:

- [Blob Analysis](#) on [page 31](#), e.g., by using the operator `dyn_threshold` to obtain the difference between two images and thus detect moving objects. The approach is fast, but detects objects only as long as they are moving.
- Background estimator (Reference Manual, chapter "[Tools > Background Estimator](#)") for the recognition of moving objects even if they stop temporarily. It adapts to global changes concerning, e.g., illumination.
- Optical flow (Reference Manual, chapter "[Filters > Match](#)") for the recognition of moving objects even if they stop temporarily. Because of complex calculations, it is slower than the other approaches, but additionally returns the velocity for each object.
- Kalman filter (Reference Manual, chapter "[Tools > Kalman Filter](#)") can be applied after the recognition of moving objects to predict the future positions of objects.

For the recognition of, e.g., irises or faces, [Classification](#) on [page 155](#) may be suitable. For more detailed information about classification see the [Solution Guide II-D](#).

Examples solving different tasks relevant for security systems can be found via the Open Example dialog inside HDevelop for the category `Industry/Surveillance` and `Security`.

1.14 Surface Inspection

For surface inspection, several approaches are available:

- Via comparison with a reference image
 - [Variation Model](#) on page 143, which compares images containing similar objects and returns the difference between them. This is suited especially if you need to detect irregularities that are placed inside the object area, whereas small irregularities at the object's border can be tolerated.
- Via comparison with a reference pattern or color
 - [Texture Analysis](#) on page 189
 - [Color Processing](#) on page 173
 - [Classification](#) on page 155 (for more detailed information about classification see the [Solution Guide II-D](#))
- Via defect description for uniformly structured surfaces
 - [Blob Analysis](#) on page 31 for objects that consist of regions of similar gray value, color, or texture.
 - [Contour Processing](#) on page 83 for objects that are represented by clear-cut edges. The contours can be obtained by an [Edge Filtering](#) on page 63 if pixel precision is sufficient, or by an [Edge and Line Extraction](#) on page 73 if subpixel precision is needed.

If a single image is not suited to cover the object to inspect, several approaches exist to combine images after their acquisition:

- Calibrated mosaicking (Solution Guide III-C, [chapter 9](#) on page 187) for a high-precision mosaicking of a discrete number of overlapping images obtained by two or more cameras.
- Uncalibrated mosaicking (Solution Guide III-C, [chapter 10](#) on page 199) for a less precise mosaicking of a discrete number of overlapping images obtained as an image sequence.
- Combination of lines obtained by a line scan camera to so-called pages (Solution Guide II-A, [section 6.6](#) on page 46) for inspecting continuous material on an assembly line.

1.15 Texture Inspection

Common approaches for texture inspection are:

- Fast Fourier Transformation (see [section 14.6](#) on page 202)
- [Classification](#) on page 155 (for more detailed information about classification see the [Solution Guide II-D](#))
- [Texture Analysis](#) on page 189

Chapter 2

Image Acquisition

Obviously, the acquisition of images is a task that must be solved in all machine vision applications. Unfortunately, this task mainly consists of interacting with special, non-standardized hardware in the form of the image acquisition device, e.g., a frame grabber board or an IEEE 1394 camera. To let you concentrate on the actual machine vision problem, HALCON provides you with interfaces performing this interaction for a large number of image acquisition devices (see <http://www.mvtec.com/halcon/framegrabber> for the latest information).

Within your HALCON application, the task of image acquisition is thus reduced to a few lines of code, i.e., a few operator calls. What's more, this simplicity is not achieved at the cost of limiting the available functionality: Using HALCON, you can acquire images from various configurations of acquisition devices and cameras in different timing modes.

Besides acquiring images from cameras, HALCON also allows you to input images that were stored in files (supported formats: BMP, TIFF, GIF, JPEG, PNG, PNM, PCX, XWD). Of course, you can also store acquired images in files.


2.1 Basic Concept

Acquiring images with HALCON basically consists of three steps. Reading images from files is even simpler: It consists of a single call to the operator `read_image`.

2.1.1 Open Image Acquisition Device

If you want to acquire images from a frame grabber board or an image acquisition device like an IEEE 1394 camera, the first step is to connect to this device. HALCON relieves you of all device-specific details; all you need to do is to call the operator `open_framegrabber`, specifying the name of the corresponding image acquisition interface.

There is also a "virtual" image acquisition interface called `File`. As its name suggests, this "frame grabber" reads images from files, and also from so-called image sequence files. The latter are HALCON-specific files, typically with the extension `.seq`; they contain a list of image file names, separated by new



lines (you can create it easily using a text editor). If you connect to such a sequence, subsequent calls to `grab_image` return the images in the sequence specified in the file. Alternatively, you can also read all images from a specific directory. Then, you do not have to create a sequence file, but simply specify the directory name instead of the sequence file as value for the parameter 'CameraType'. Now, subsequent calls to `grab_image` return the images found in the specified image directory. Both approaches are useful if you want to test your application with a sequence of image files and later switch to a real image acquisition device.

2.1.2 Acquire Image(s)

Having connected to the device, you acquire images by simply calling `grab_image`.

To load an image from disk, you use `read_image`. Images are searched for in the current directory and in the directories specified in the environment variable `HALCONIMAGES`.

2.1.3 Close Image Acquisition Device

At the end of the application, you close the connection to the image acquisition device to free its resources with the operator `close_framegrabber`.

2.1.4 A First Example

As already remarked, acquiring images from file corresponds to a single operator call:

```
read_image (Image, 'particle')
```

The following code processes images read from an image sequence file:


```

SequenceName := 'datacode/ecc200/ecc200_cpu_light.seq'
open_framegrabber ('File', 1, 1, 0, 0, 0, 0, 'default', -1, 'default', -1, \
                  'default', SequenceName, 'default', -1, -1, FGHandle)
while (1)
    grab_image (Image, FGHandle)
    ... process image ...
endwhile

```

2.2 Extended Concept

In real applications, it is typically not enough to tell the camera to acquire an image; instead, it may be important that images are acquired at the correct moment or rate, and that the camera and the image acquisition interface are configured suitably. Therefore, HALCON allows to further parameterize the acquisition process. In HDevelop, an assistant is provided via the menu item **Assistants > Image Acquisition** that assists you when selecting your image source, adjusting the parameters, and generating suitable code.



2.2.1 Open Image Acquisition Device

When connecting to your image acquisition device with `open_framegrabber`, the main parameter is the name of the corresponding HALCON image acquisition interface. As a result, you obtain a so-called handle, with which you can access the device later, e.g., to acquire images with `grab_image` or `grab_image_async`.

With other parameters of `open_framegrabber` you can describe the configuration of image acquisition device(s) and camera(s), which is necessary when using more complex configurations, e.g., multiple cameras connected to different ports on different frame grabber boards. Further parameters allow you to specify the desired image format (size, resolution, pixel type, color space). For most of these parameters

there are default values that are used if you specify the values 'default' (string parameters) or -1 (numeric parameters).

With the operator `info_framegrabber` you can query information like the version number of the interface or the available boards, port numbers, and camera types.

Detailed information about the parameters of `open_framegrabber` can be found in the Solution Guide II-A (configuring the connection: [section 3](#) on page 11; configuring the acquired image: [section 4](#) on page 19).

2.2.2 Set Parameters

As described above, you already set parameters when connecting to the image acquisition device with `open_framegrabber`. These parameters (configuration of image_acquisition device(s) / camera(s) and image size etc.) are the so-called *general parameters*, because they are common to almost all image acquisition interfaces. However, image acquisition devices differ widely regarding the provided functionality, leading to many more *special parameters*. These parameters can be customized with the operator `set_framegrabber_param`.

With the operator `get_framegrabber_param` you can query the current values of the common and special parameters.

Detailed information about setting parameters can be found in the Solution Guide II-A in [section 4](#) on page 19.

2.2.3 Acquire Image(s)

Actually, in a typical machine vision application you will not use the operator `grab_image` to acquire images, but `grab_image_async`. The difference between these two operators is the following: If you acquire and process images in a loop, `grab_image` always requests the acquisition of a new image and then blocks the program until the acquisition has finished. Then, the image is processed, and afterwards, the program waits for the next image. When using `grab_image_async`, in contrast, images are acquired and processed in parallel: While an image is processed, the next image is already being acquired. This, of course, leads to a significant speedup of the applications.

HALCON offers many more modes of acquiring images, e.g., triggering the acquisition by external signals or acquiring images simultaneously from multiple cameras. Detailed information about the various modes of acquiring images can be found in the Solution Guide II-A in [section 5](#) on page 23.

2.3 Programming Examples

Example programs for all provided image acquisition interfaces can be found in the directory %HALCONEXAMPLES%\hdevelop\Image\Acquisition. Further examples are described in the [Solution Guide II-A](#).

2.4 Selecting Operators

2.4.1 Open Image Acquisition Device

Standard:

`open_framegrabber`

Advanced:

`info_framegrabber`

2.4.2 Set Parameters

Standard:

`set_framegrabber_param, get_framegrabber_param`

2.4.3 Acquire Image(s)

Standard:

`read_image, grab_image, grab_image_async`

2.4.4 Close Image Acquisition Device

Standard:

`close_framegrabber`

2.5 Tips & Tricks

2.5.1 Direct Access to External Images in Memory

You can also pass externally created images, i.e., the raw image matrix in the computer's memory, to HALCON using the operators `gen_image1`, `gen_image3`, `gen_image1_extern` or `gen_image3_extern`. For an example see the Solution Guide II-A, [section 6.2](#) on page 38.

2.5.2 Unsupported Image Acquisition Devices

If you want to use an image acquisition device that is currently not supported by HALCON, i.e., for which no HALCON image acquisition interface exists, you can create your own interface; how to do this is described in detail in the [Image Acquisition Interface Programmer's Manual](#).

Chapter 3


Region Of Interest

The concept of regions of interest (ROIs) is essential for machine vision in general and for HALCON in particular. The aim is to focus the processing on a specific part of the image. This approach combines region information with the image matrix: Only the image part corresponding to the region remains relevant, which reduces the number of pixels to be processed.

The advantages of using ROIs are manifold. First of all, it is a very good method to speed up a process because fewer pixels need to be processed. Furthermore, it focuses processing, e.g., a gray value feature is usually calculated only for a part of the image. Finally, ROIs are used to define templates, e.g., for matching. HALCON allows to make full use of the concept of ROIs because it enables using arbitrary shapes for the regions. This means that you are not limited to standard shapes like rectangles or polygons, but can really use *any* form - the best one to solve a given problem.

3.1 Basic Concept

Making use of ROIs is split into two simple parts: creating regions and combining them with the image.



3.1.1 Create Region

HALCON provides many ways to create regions, which can then be used as ROIs. The traditional way is to generate standard shapes like circles, ellipses, rectangles, or polygons. In addition, regions can be derived by converting them from other data types like XLD, by segmenting an image, or by user interaction.

3.1.2 Create ROI

By combining a region with an image, the region assumes the role of an ROI, i.e., it defines which part of the image must be processed. In HALCON, the ROI is also called the domain of the image. This term comes from mathematics where an image can be treated as a function that maps coordinates to gray values. An ROI reduces the domain of this function from the complete image to the relevant part. Therefore, the operator to combine regions and images is called `reduce_domain`. This simple operator fulfills the desired task in almost all applications.

3.1.3 A First Example

As an example for the basic concept, the following program shows all important steps to make use of an ROI. The image is acquired from file. Inside the image, only a circular part around the center should be processed. To achieve this, a circular region is generated with `gen_circle`. This region is combined with the image using `reduce_domain`. This has the effect that only the pixels of the ROI are processed when calling an operator. If, e.g., the operator `edges_sub_pix` is applied to this image, the subpixel accurate contours are extracted only inside the circle. To make this visible, some visualization operators are added to the end of the example program.

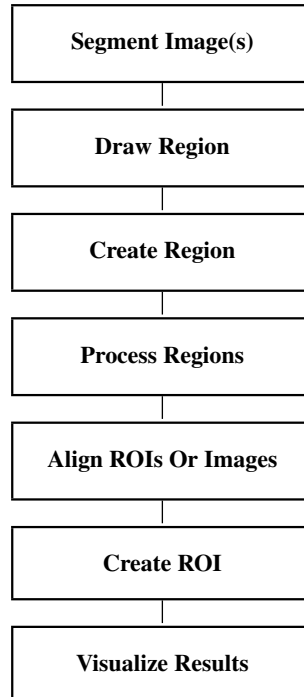
```
read_image (Image, 'mreut')
gen_circle (ROI, 256, 256, 200)
reduce_domain (Image, ROI, ImageReduced)
edges_sub_pix (ImageReduced, Edges, 'lanser2', 0.5, 20, 40)
dev_display (Image)
dev_display (ROI)
dev_display (Edges)
```



Figure 3.1: Processing the image only within the circular ROI.

3.2 Extended Concept

When we take a closer look at ROIs, extra steps become important if an application needs to be more flexible.



3.2.1 Segment Image(s)

Very typical for HALCON is the creation of ROIs by a segmentation step. Instead of having a predefined ROI, the parts of the image that are relevant for further processing are extracted from the image using image processing methods. This approach is possible because ROIs are nothing else but normal HALCON regions, and therefore share all their advantages like efficient processing and arbitrary shapes (see Quick Guide in [section 2.1.2.2](#) on page 17 for more information on HALCON regions). The segmentation of regions used for ROIs follows the same approach as standard blob analysis. For more details, please refer to the [description of this step](#) on page 32.

3.2.2 Draw Region

The standard way to specify ROIs is to draw the shape interactively using the mouse. To make this easy, HALCON provides special operators for standard shapes and free-form shapes. All operators for this kind of action start with the prefix `draw_`. The drawing is performed by making use of the left mouse button (drawing, picking, and dragging) and finished by clicking the right mouse button. For each such

draw-operator HALCON provides operators to generate regions by using the returned parameters (see the description of the step [Create Region](#) on page 19). Operators for mouse interaction can be found in the reference manual in the chapter “[Graphics ▸ Drawing](#)”. More information on user interaction can be also found in the chapter [Visualization](#) on page 289.

3.2.3 Create Region

The standard way is to generate regions based on the coordinates and dimensions returned by a user interaction or by coordinate values stored in a file. In this case, operators like [gen_circle](#), [gen_rectangle2](#), or [gen_region_polygon_filled](#) are used. More advanced are special shapes used to guide a preprocessing step to save execution time. Typical examples for this are grids of lines or dots or checker boards. With these shapes, the images can be covered in a systematic way and checked for specific object occurrences. If you want to segment, e.g., blobs of a given minimum size it is sufficient to use in a first step a search grid that is finer than the minimum object size to locate fragments. In a second step these fragments are dilated ([dilation_rectangle1](#)) and the segmentation method is called once again, now within this enlarged area. If the objects cover only a relatively small area of the image this approach can speed up the process significantly.

3.2.4 Process Regions

Sometimes the shape of a given ROI, either generated from the program or defined by the user, does not fulfill the requirements. Here, HALCON provides many operators to modify the shape to adapt it accordingly. Often used operators are, e.g., [fill_up](#) to fill holes inside the region, [shape_trans](#) to apply a general transformation like the convex hull or the smallest rectangle, or morphological operators like [erosion_circle](#) to make the region smaller or [closing_circle](#) to fill gaps. For more details, please refer to the [description of this step](#) on page 35.

3.2.5 Align ROIs Or Images

Sometimes the coordinates of an ROI depend on the position of another object in the image. If the object moves, the ROI must be moved (aligned) accordingly. This is achieved by first locating the object using template matching. Based on the determined position and the orientation, the coordinates of the ROIs are then transformed.

How to perform alignment using shape-based matching is described in the Solution Guide II-B in [section 2.4.3.2](#) on page 42.

3.2.6 Create ROI

This step combines the region and the image to make use of the region as the domain of the image. The standard method that is recommended to be used is [reduce_domain](#). It has the advantage of being safe and having a simple semantics. [rectangle1_domain](#) is a shortcut for generating rectangular ROIs (instead of calling [gen_rectangle1](#) and [reduce_domain](#) in sequence). For advanced applications [change_domain](#) can be used as a slightly faster version than [reduce_domain](#). This operator does not

perform an intersection with the existing domain and does not check if the region is outside the image - which will cause a system crash when applying an operator to the data afterwards if the region lies partly outside the image. If the programmer ensures that the input region is well defined, this is a way to save (a little) execution time.

3.2.7 Visualize Results

Finally, you might want to display the ROIs or the reduced images. With the operator `get_domain`, the region currently used by the image can be accessed and displayed (and processed) like any other region. When displaying an image, e.g., with `disp_image`, only the defined pixels are displayed. Pixels in the graphics window outside the domain of the image will not be modified.

For detailed information see the [description of this method](#) on page 289.

3.3 Programming Examples

This section gives a brief introduction to programming ROIs in HALCON. Two examples show the principles of region generation, combining these with images, and then processing the data.

3.3.1 Processing inside a User Defined Region

Example: `solution_guide/basics/critical_points.hdev`

[Figure 3.2](#) shows an image with marks that are used for a camera calibration in a 3D application. Here, we assume that the marks must be extracted in a given part of the image only.




Figure 3.2: (a) Original image with drawn ROI; (b) reduced image with extracted points.

To achieve this, the user draws a region of interest with the mouse. The corresponding operator is `draw_region`. It has the window handle returned by `dev_open_window` as input and returns a region

when the right mouse button is pressed. The operator `reduce_domain` combines this region with the image.

```
draw_region (Region, WindowHandle)
reduce_domain (Image, Region, ImageReduced)
```

When calling the point extraction operator `critical_points_sub_pix` on this reduced image, only points inside the ROI are found. The final part of the program shows how to display these points overlaid on the image.

```
critical_points_sub_pix (ImageReduced, 'facet', 1.5, 8, RowMin, ColumnMin, \
                        RowMax, ColumnMax, RowSaddle, ColSaddle)
dev_clear_window ()
dev_display (ImageReduced)
dev_set_color ('yellow')
for i := 0 to |RowSaddle|-1 by 1
    gen_cross_contour_xld (Cross, RowSaddle[i], ColSaddle[i], 25, 0.785398)
    dev_display (Cross)
endfor
```

3.3.2 Interactive Partial Filtering of an Image

Example: `solution_guide/basics/median_interactive.hdev`

The task is to filter an image with a median filter only at the points where the user clicks with the mouse into the image, i.e., in the graphics window displaying the image.



Figure 3.3: Partially filtered image.

To do this, a loop is used inside which the mouse position is continuously requested with `get_mposition`. Because this operator throws an exception if the mouse is outside the graphics window the call is protected with `dev_set_check`.

```
Button := 0
while (Button # 4)
  Row := -1
  Column := -1
  dev_set_check ('~give_error')
  get_mposition (WindowHandle, Row, Column, Button)
  dev_set_check ('give_error')
```

If the mouse is over the window, a circular region is displayed, which shows where the filter would be applied.

```
if (Row >= 0 and Column >= 0)
  gen_circle (Circle, Row, Column, 20)
  boundary (Circle, RegionBorder, 'inner')
  dev_display (RegionBorder)
```

If the left mouse button is pressed, `median_image` must be applied in the local neighborhood of the current mouse position. This is done by generating a circle with `gen_circle` and then calling `reduce_domain`.

```
if (Button = 1)
  reduce_domain (Image, Circle, ImageReduced)
```

Now, the filter is called with this reduced image and the result is painted back into the input image for possible repetitive filtering. The loop will be terminated when the right mouse button is clicked.

```
median_image (ImageReduced, ImageMedian, 'circle', 5, \
              'mirrored')
overpaint_gray (Image, ImageMedian)
endif
```

3.3.3 Inspecting the Contours of a Tool

Example: `hdevelop/Applications/Measuring-2D/circles.hdev`

The task of this example is to inspect the contours of the tool depicted in [figure 3.4](#).

Because the subpixel-precise contour extraction is time-consuming, in a first step an ROI is created via a standard blob analysis (see [Blob Analysis](#) on page 31): With a threshold operator the object to be measured is extracted. This region is converted to its boundary, omitting the pixels at the image border.

```
fast_threshold (Image, Region, 0, 120, 7)
boundary (Region, RegionBorder, 'inner')
clip_region_rel (RegionBorder, RegionClipped, 5, 5, 5, 5)
```




Figure 3.4: Fitting circles to the contours of the tool.

The result is a small region close to the edge of the object. The boundary of the region, i.e., the edge, is dilated to serve as the ROI for the edge extraction. Now, the subpixel-precise edge extractor is called and the contour is segmented into straight lines and circular arcs.

```
dilation_circle (RegionClipped, RegionDilation, 2.5)
reduce_domain (Image, RegionDilation, ImageReduced)
edges_sub_pix (ImageReduced, Edges, 'lanser2', 0.5, 40, 60)
segment_contours_xld (Edges, ContoursSplit, 'lines_circles', 5, 4, 3)
```

For the segments that represent arcs, the corresponding circle parameters are determined. For inspection purposes, circles with the same parameters are generated and overlaid on the image (see also [Contour Processing](#) on page 83).

```
get_contour_global_attrib_xld (ObjectSelected, 'cont_approx', Attrib)
if (Attrib > 0)
    fit_circle_contour_xld (ObjectSelected, 'ahuber', -1, 2, 0, 3, 2, \
                          Row, Column, Radius, StartPhi, EndPhi, \
                          PointOrder)
    gen_circle_contour_xld (ContCircle, Row, Column, Radius, 0, \
                          rad(360), 'positive', 1.0)
    dev_display (ContCircle)
endif
```

3.4 Selecting Operators

3.4.1 Segment Image(s)

Please refer to the [detailed operator list for the step Segment Image\(s\)](#) on page 44.

3.4.2 Draw Region

Standard:

`draw_circle`, `draw_rectangle1`, `draw_rectangle2`, `draw_region`

Advanced:

`draw_circle_mod`, `draw_rectangle1_mod`, `draw_rectangle2_mod`, `draw_xld`, `draw_xld_mod`

3.4.3 Create Region

Standard:

`gen_circle`, `gen_ellipse`, `gen_rectangle2`, `gen_region_polygon_filled`

Advanced:

`gen_checker_region`, `gen_grid_region`

More operators to generate regions can be found in the reference manual in chapter “Regions ▸ Creation”.

3.4.4 Process Regions

Please refer to the [detailed operator list for the step Process Regions](#) on page 44.

3.4.5 Align ROIs Or Images

Operators for aligning ROIs or images are described in the [Solution Guide II-B](#).

3.4.6 Create ROI

Standard:

`reduce_domain`, `rectangle1_domain`

Advanced:

`change_domain`, `full_domain`, `add_channels`

3.4.7 Visualize Results

Please refer to the [operator list for the method Visualization](#) (see section 19.4 on page 297).

3.5 Relation to Other Methods

One class of operators does not follow the standard rules for ROI handling: the operators of the measure tool (see the [description of this method](#) on page 49). Here, the ROI is defined during the creation of a tool ([gen_measure_arc](#) and [gen_measure_rectangle2](#)) by specifying the coordinates as numeric values. The domain defined for the image will be ignored in this case.

3.6 Tips & Tricks

3.6.1 Reuse ROI

If an ROI is used multiple times it is useful to save the region to file and load it at the beginning of the application. Storing to file is done using [write_region](#), loading with [read_region](#).

3.6.2 Effect of ROI Shape on Speed Up

ROIs are a perfect way to save execution time: The smaller the ROI, the faster the application. This can be used as a general rule. If we consider this in more detail, we also need to think about the shape of ROIs. Because ROIs are based on the HALCON regions they use runlength encoding (see Quick Guide in [section 2.1.2.2](#) on page 17. for more information on HALCON regions). This type of encoding is perfect if the runs are long. Therefore, a horizontal line can be both stored and processed more efficiently than a vertical line. This holds as well for the processing time of ROIs. Obviously this type of overhead is very small and can only be of importance with very fast operators like [threshold](#).

3.7 Advanced Topics

3.7.1 Filter masks and ROIs

If a filter is applied with a reduced domain, the result along the ROI boundary might be surprising because gray values lying outside the boundary are used as input for the filter process. To understand this, you must consider the definition of domains in this context: A domain defines for a filter for which input pixels output pixels must be calculated. But pixels outside the domain (which lie within the image matrix) can be used for processing. If this behavior is not desired, the operator [expand_domain_gray](#) can be used to propagate the border gray values outward. The result looks as if the boundary is copied multiple times with larger and larger distances.

3.7.2 Binary Images

In some applications, it might be necessary to use ROIs that are available as binary images. To convert these to HALCON regions, you must use [gen_image1](#) to convert them into a HALCON image, followed by [threshold](#) to generate the region. The conversion back can easily be achieved using [region_to_bin](#) followed by [get_image_pointer1](#). It is also possible to import binary image files using [read_region](#).

3.7.3 Problems caused by Gray Values outside of the Image Domain

Operators that have an image as output (e.g., most filter operators) only return pixels that correspond to the pixels of the input domain.

For performance reasons, pixels that lie outside of the image domain become 'undefined'. This means that they are either assigned a random gray value or a gray value of 0 if the system parameter `init_new_image` is set to `true` (= default). In certain cases these 'undefined' pixels might lead to problems which can, however, easily be handled.

Expanding the resulting image to the full domain with `full_domain`, will lead to artifacts appearing outside of the former image domain. Another cause of problems are undefined values outside of the domain if more than one filter is used in sequence as the filters consider the undefined values close to the domain border as well. This means that with every following filter the error increases, starting from the border to the middle. The following paragraphs present four strategies for solving those problems.

Errors caused by undefined pixels can easily be prevented by, e.g., choosing a domain that is considerably larger than the image part that is actually needed and reducing the image domain (e.g., with the operators `erosion_rectangle1` and `reduce_domain`) by half of the filter width before applying the next filter. In doing so, those parts of the image containing incorrect values are cut off and therefore will not increase the error for the next filter operation.

Another option is to choose the domain exactly the size of the interesting part within the image and then calling the operator `expand_domain_gray` before applying a filter. This operator copies the pixels inside of the border to the outside of the border and therefore avoids errors caused by pixels that are undefined outside of the domain. Subsequently, the domain can again be reduced to its original size. This process should be repeated for every following filter operation. Note, however, that this option increases the runtime significantly.

If runtime is not an issue, the operator `full_domain` can be called before applying the first filter to the image. That way the whole image is defined as domain and undefined pixels can be avoided completely.

Another possibility of achieving an image without undefined pixels is by calling the operator `crop_domain` before applying a filter. The operator `crop_domain` crops the image to the size of the domain which means that the domain then covers a complete smaller image. Note, however, that for the cropped image the coordinate system has changed in relation to the original image which will influence all following applications depending on the image coordinate system (e.g., calculating the center of gravity).

Chapter 4

Blob Analysis

The idea of blob analysis is quite easy: In an image the pixels of the relevant objects (also called foreground) can be identified by their gray value. For example, [figure 4.1](#) shows tissue particles in a liquid. These particles are bright and the liquid (background) is dark. By selecting bright pixels (thresholding) the particles can be detected easily. In many applications this simple condition of dark and bright pixels no longer holds, but the same results can be achieved with extra pre-processing or alternative methods for pixel selection / grouping.





Figure 4.1: Basic idea of blob analysis: (a) original image, (b) extracted blobs with calculated center points.

The advantage of blob analysis is the extreme flexibility that comes from the huge number of operators that HALCON offers in this context. Furthermore, these methods typically have a very high performance.

Methods known from blob analysis can also be combined with many other vision tasks, e.g., as a pre-processing step for a flexible generation of regions of interest.

4.1 Basic Concept

Blob analysis mainly consists of three parts:



4.1.1 Acquire Image(s)

First, an image is acquired.

For detailed information see the [description of this method](#) on page 13.

4.1.2 Segment Image(s)

Having acquired the image, the task is to select the foreground pixels. This is also called segmentation. The result of this process typically is referred to as blobs (binary large objects). In HALCON, the data type is called a region.

4.1.3 Extract Features

In the final step, features like the area (i.e., the number of pixels), the center of gravity, or the orientation are calculated.

4.1.4 A First Example

An example for this basic concept is the following program, which belongs to the example explained above. Here, the image is acquired from file. All pixels that are brighter than 120 are selected using [threshold](#). Then, an extra step is introduced which is not so obvious: The operator [connection](#) separates the set of all bright pixels into so called connected components. The effect of this step is that we now have multiple regions instead of the single region that is returned by [threshold](#). The last step

of this program is the calculation of some features. Here, the operator `area_center` determines the size (number of pixels) and the center of gravity. Please note that `area_center` returns multiple values for all three feature parameters (one value for each connected component).

```
read_image (Image, 'particle')
threshold (Image, BrightPixels, 120, 255)
connection (BrightPixels, Particles)
area_center (Particles, Area, Row, Column)
```

4.2 Extended Concept

In many cases the segmentation of blobs will be more advanced than in the above example. Reasons for this are, e.g., clutter or inhomogeneous illumination. Furthermore, postprocessing like transforming the features to real world units or visualization of results are often required.

4.2.1 Use Region Of Interest

Blob analysis can be sped up by using a region of interest. The more the region in which the blobs are searched can be restricted, the faster and more robust the search will be.

For detailed information see the [description of this method](#) on page 19.

4.2.2 Align ROIs Or Images

In some applications, the regions of interest must be aligned relative to another object. Alternatively, the image itself can be aligned, e.g., by rotating or cropping it.

How to perform alignment using shape-based matching is described in the Solution Guide II-B in [section 2.4.3.2](#) on page 42.

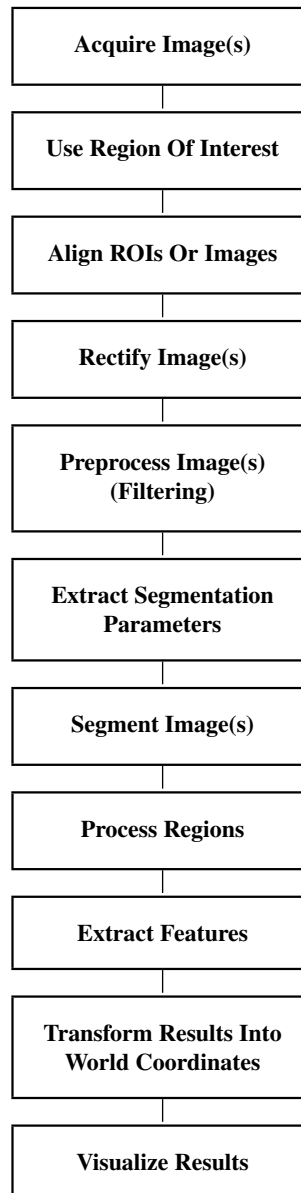
4.2.3 Rectify Image(s)

Similarly to alignment, it may be necessary to rectify the image, e.g., to remove lens distortions or to transform the image into a reference point of view.

Detailed information about rectifying images can be found in the Solution Guide III-C in [section 3.4](#) on page 62.

4.2.4 Preprocess Image(s) (Filtering)

The next important part is the pre-processing of the image. Here, operators like `mean_image` or `gauss_image` can be used to eliminate noise. A fast but less perfect alternative to `gauss_image` is `binomial_filter`. The operator `median_image` is useful for suppressing small spots or thin lines. The operator `anisotrope_diff` is useful for edge-preserving smoothing, and finally `fill_interlace` is used to eliminate defects caused by interlaced cameras.



4.2.5 Extract Segmentation Parameters

Instead of using fixed threshold values, they can be extracted dynamically for each image. One example for this is a gray value histogram that has multiple peaks, one for each object class. Here, you can use the operators `gray_histo_abs` and `histo_to_thresh`.

As an advanced alternative, you can use the operator `intensity` in combination with a reference image that contains only background: During setup, you determine the mean gray value of a background region.

During the inspection, you again determine this mean gray value. If it has changed, you adapt the threshold accordingly.

4.2.6 Segment Image(s)

For the segmentation various methods can be used. The most simple method is [threshold](#), where one or more gray value ranges that belong to the foreground objects are specified. Another very common method is [dyn_threshold](#). Here, a second image is passed as a reference. With this approach, a local threshold instead of a global threshold is used. These local threshold values are stored in the reference image. The reference image can either be static by taking a picture of the empty background or can be determined dynamically with smoothing filters like [mean_image](#).

4.2.7 Process Regions

Once blob regions are segmented, it is often necessary to modify them, e.g., by suppressing small areas, regions of a given orientation, or regions that are close to other regions. In this context, the morphological operators [opening_circle](#) and [opening_rectangle1](#) are often used to suppress noise and [closing_circle](#) and [closing_rectangle1](#) to fill gaps.

Blobs with a specific feature can be selected with [select_shape](#), [select_shape_std](#), and [select_shape_proto](#).

4.2.8 Extract Features

To finalize the image processing, features of the blobs are extracted. The type of features needed depends on the application. A full list can be found in the Reference Manual in the chapters “[Regions > Features](#)” and “[Image > Features](#)”.

4.2.9 Transform Results Into World Coordinates

Features like the area or the center of gravity often must be converted to world coordinates. This can be achieved with the HALCON camera calibration.

How to transform results into world coordinates is described in detail in the Solution Guide III-C in [section 3.3](#) on page 57.

4.2.10 Visualize Results

Finally, you might want to display the images, the blob (regions), and the features.

For detailed information see the [description of this method](#) on page 289.

4.3 Programming Examples

This section gives a brief introduction to using HALCON for blob analysis.

4.3.1 Crystals

Example: `solution_guide/basics/crystal.hdev`

Figure 4.2a shows an image taken in the upper atmosphere with collected samples of crystals. The task is to analyze the objects to determine the frequency of specific shapes. One of the important objects are the hexagonally shaped crystals.




Figure 4.2: Extracting hexagonal crystals: (a) original image with arrows marking the crystals to be extracted, (b) result of the initial segmentation, (c) finally selected blobs.

First, the image is read from file with `read_image`. The segmentation of objects is performed with a local threshold because of the relatively low contrast of the crystals combined with a non-homogeneous background. The background is determined with the average filter `mean_image`. The filter mask size is selected such that it has about three times the width of the dark areas. `dyn_threshold` now compares the smoothed with the original gray values, selecting those pixels that are darker by a contrast of 8 gray values. `connection` separates the objects into connected components. Figure 4.2b shows the result of this initial segmentation.

```
read_image (Image, 'crystal')
mean_image (Image, ImageMean, 21, 21)
dyn_threshold (Image, ImageMean, RegionDynThresh, 8, 'dark')
connection (RegionDynThresh, ConnectedRegions)
```

In the following processing step, the task now is to select only the hexagonally shaped crystals. For this, they are first transformed into their convex hull. This is like putting a rubber band around each region. From these regions, those that are big ([select_shape](#)) and have a given gray value distribution ([select_gray](#)) are selected. The parameters for the selection are determined so that only the relevant crystals remain (see [figure 4.2c](#)).

```
shape_trans (ConnectedRegions, ConvexRegions, 'convex')
select_shape (ConvexRegions, LargeRegions, 'area', 'and', 600, 2000)
select_gray (LargeRegions, Image, Crystals, 'entropy', 'and', 1, 5.6)
```

4.3.2 Atoms

Example: `solution_guide/basics/atoms.hdev`

Specialized microscopes are able to determine the rough location of single atoms. This is useful, e.g., to analyse the grid change of crystals at a p-n-junction. A segmentation that works perfectly well on images like these is the watershed method. Here, each dark basin is returned as a single region. Because at the outer part of the image atoms are only partially visible, the first task is to extract only those that are not close to the image border. Finally, the irregularity is extracted. This is done by looking for those atoms that have an abnormal (squeezed) shape (see [figure 4.3](#)).

```
gauss_image (Image, ImageGauss, 5)
watersheds (ImageGauss, Basins, Watersheds)
select_shape (Basins, SelectedRegions1, 'column1', 'and', 2, Width-1)
select_shape (SelectedRegions1, SelectedRegions2, 'row1', 'and', 2, \
    Height-1)
select_shape (SelectedRegions2, SelectedRegions3, 'column2', 'and', 1, \
    Width-3)
select_shape (SelectedRegions3, Inner, 'row2', 'and', 1, Height-3)
select_shape (Inner, Irregular, 'compactness', 'and', 1.45, 3)
```




Figure 4.3: Inspecting atom structure.

4.3.3 Analyzing Particles

Example: `hdevelop/Applications/Measuring-2D/particle.hdev`

The task of this example is to analyze particles in a liquid. The main difficulty in this application is the presence of two types of objects: big bright objects and small objects with low contrast. In addition, the presence of noise complicates the segmentation.




Figure 4.4: Extracting the small particles: (a) original image, (b) result.

The program segments the two classes of objects separately using two different methods: global and local thresholding. With additional post-processing, the small particles can be extracted in a robust manner.

```
threshold (Image, Large, 110, 255)
dilation_circle (Large, LargeDilation, 7.5)
complement (LargeDilation, NotLarge)
reduce_domain (Image, NotLarge, ParticlesRed)

mean_image (ParticlesRed, Mean, 31, 31)
dyn_threshold (ParticlesRed, Mean, SmallRaw, 3, 'light')
opening_circle (SmallRaw, Small, 2.5)
connection (Small, SmallConnection)
```

4.3.4 Extracting Forest Features from Color Infrared Image

Example: `hdevelop/Applications/Object-Recognition-2D/forest.hdev`




Figure 4.5: (a) Original image; (b) extracted trees and meadows; (c) extracted roads.

The task of this example is to detect different object classes in the color infrared image depicted in [figure 4.5](#): trees (coniferous and deciduous), meadows, and roads.

The image data is a color infrared image, which allows to extract roads very easily because of their specific color. For that, the multi-channel image is split into its three channels and each channel is investigated individually. Further information about decomposing multi-channel images can be found in [Color Processing](#) on page 173.

```
read_image (Forest, 'forest_air1')
decompose3 (Forest, Red, Green, Blue)
threshold (Blue, BlueBright, 80, 255)
connection (BlueBright, BlueBrightConnection)
select_shape (BlueBrightConnection, Path, 'area', 'and', 100, 100000000)
```

Beech trees are segmented in the red channel based on their intensity and minimum size.

```
threshold (Red, RedBright, 120, 255)
connection (RedBright, RedBrightConnection)
select_shape (RedBrightConnection, RedBrightBig, 'area', 'and', 1500, \
    10000000)
closing_circle (RedBrightBig, RedBrightClosing, 7.5)
opening_circle (RedBrightClosing, RedBrightOpening, 9.5)
connection (RedBrightOpening, RedBrightOpeningConnection)
select_shape (RedBrightOpeningConnection, BeechBig, 'area', 'and', 1000, \
    100000000)
select_gray (BeechBig, Blue, Beech, 'mean', 'and', 0, 59)
```

Meadows have similar spectral properties, but are slightly brighter.

```

union1 (Beech, BeechUnion)
complement (BeechUnion, NotBeech)
difference (NotBeech, Path, NotBeechNotPath)
reduce_domain (Red, NotBeechNotPath, NotBeechNotPathRed)
threshold (NotBeechNotPathRed, BrightRest, 150, 255)
connection (BrightRest, BrightRestConnection)
select_shape (BrightRestConnection, Meadow, 'area', 'and', 500, 1000000)

```

The coniferous trees are extracted using the watershed approach with an additional thresholding inside the basins to get rid of the shadow areas.

```

union2 (Path, RedBrightClosing, BeechPath)
smooth_image (Red, RedGauss, 'gauss', 4.0)
invert_image (RedGauss, Invert)
watersheds (Invert, SpruceRed, Watersheds)
select_shape (SpruceRed, SpruceRedLarge, 'area', 'and', 100, 5000)
select_gray (SpruceRedLarge, Red, SpruceRedInitial, 'max', 'and', 100, 200)
gen_empty_obj (LocalThresh)
count_obj (SpruceRedInitial, NumSpruce)
for i := 1 to NumSpruce by 1
    select_obj (SpruceRedInitial, SingleSpruce, i)
    min_max_gray (SingleSpruce, Red, 50, Min, Max, Range)
    reduce_domain (Red, SingleSpruce, SingleSpruceRed)
    threshold (SingleSpruceRed, SingleSpruceBright, Min, 255)
    connection (SingleSpruceBright, SingleSpruceBrightCon)
    select_shape_std (SingleSpruceBrightCon, MaxAreaSpruce, 'max_area', 70)
    concat_obj (MaxAreaSpruce, LocalThresh, LocalThresh)
endfor
opening_circle (LocalThresh, FinalSpruce, 1.5)

```

4.3.5 Checking a Boundary for Fins

Example: `hdevelop/Applications/Measuring-2D/fin.hdev`

The task of this example is to check the outer boundary of a plastic part. In this case, some objects show fins that are not allowed for faultless pieces (see [figure 4.6](#)).

The program first extracts the plastic part and then forms the complement to extract the background region (in which the fin appears as an indentation).

```

bin_threshold (Fin, Dark)
difference (Fin, Dark, Background)

```

This indentation in the background region is then closed using a morphological operator.

```

closing_circle (Background, ClosedBackground, 250)

```

Significant differences between the closed region and the original region are production errors.




Figure 4.6: Boundary with extracted fin.

```
difference (ClosedBackground, Background, RegionDifference)
opening_rectangle1 (RegionDifference, FinRegion, 5, 5)
```

4.3.6 Bonding Balls

Example: `hdevelop/Applications/Completeness-Check/ball.hdev`

The task of this example is to inspect the diameter of the ball bonds depicted in [figure 4.7](#).




Figure 4.7: Measuring the diameter of ball bonds.

The extraction of the ball bonds is a two step approach: First, the die is located by segmenting bright areas and transforming them into their smallest surrounding rectangle.


```
threshold (Bond, Bright, 100, 255)
shape_trans (Bright, Die, 'rectangle2')
```

Now, the processing is focused to the region inside the die using `reduce_domain`. In this ROI, the program checks for dark areas that correspond to wire material.

```
reduce_domain (Bond, Die, DieGrey)
threshold (DieGrey, Wires, 0, 50)
fill_up_shape (Wires, WiresFilled, 'area', 1, 100)
```

After removing irrelevant structures and arranging the bonds in a predefined order, the desired features are extracted.

```
opening_circle (WiresFilled, Balls, 15.5)
connection (Balls, SingleBalls)
select_shape (SingleBalls, IntermediateBalls, 'circularity', 'and', 0.85, \
1.0)
sort_region (IntermediateBalls, FinalBalls, 'first_point', 'true', 'column')
smallest_circle (FinalBalls, Row, Column, Radius)
```

4.3.7 Surface Scratches

Example: `solution_guide/basics/surface_scratch.hdev`

This example detects scratches on a metal surface (see [figure 4.8](#)).




Figure 4.8: Detecting scratches on a metal surface: (a) original image, (b) extracted scratches still partly fractionated, (c) final result with merged scratches.

The main difficulties for the segmentation are the inhomogenous background and the fact that the scratches are thin structures. Both problems can be solved using a local threshold, i.e., the operators `mean_image` and `dyn_threshold`. After `connection`, the small objects that are mainly noise are removed (see [figure 4.8b](#)).

```
mean_image (Image, ImageMean, 7, 7)
dyn_threshold (Image, ImageMean, DarkPixels, 5, 'dark')
connection (DarkPixels, ConnectedRegions)
select_shape (ConnectedRegions, SelectedRegions, 'area', 'and', 10, 1000)
```

The scratches are part of the selection, but if we look closely we see that they are partially fractioned. To solve this problem we combine all fractions again into one big region. By applying `dilation_circle`, neighboring parts with a given maximum distance are now combined. To finally get the correct shape of the scratches - which are now too wide because of the dilation - `skeleton` is used to thin the shape to a width of one pixel.

```
union1 (SelectedRegions, RegionUnion)
dilation_circle (RegionUnion, RegionDilation, 3.5)
skeleton (RegionDilation, Skeleton)
connection (Skeleton, Errors)
```

The last step is to distinguish between small dots and scratches on the surface. This is achieved with `select_shape`, using the size as feature. [Figure 4.8c](#) depicts the result.

```
select_shape (Errors, Scratches, 'area', 'and', 50, 10000)
select_shape (Errors, Dots, 'area', 'and', 1, 50)
```

4.4 Selecting Operators

4.4.1 Acquire Image(s)

Please refer to the [operator list for the method Image Acquisition](#) (see section 2.4 on page 17).

4.4.2 Use Region Of Interest

Please refer to the [operator list for the method Region Of Interest](#) (see section 3.4 on page 26).

4.4.3 Align ROIs Or Images

Operators for rectifying images are described in the [Solution Guide II-B](#).

4.4.4 Rectify Image(s)

Operators for rectifying images are described in the [Solution Guide III-C](#).

4.4.5 Preprocess Image(s) (Filtering)

Standard:

`mean_image`, `gauss_image`, `binomial_filter`, `median_image`

Advanced:

`smooth_image`, `anisotrope_diff`, `fill_interlace`, `rank_image`

4.4.6 Extract Segmentation Parameters

Standard:

`gray_histo_abs`, `histo_to_thresh`

Advanced:

`intensity`

4.4.7 Segment Image(s)

Standard:

`threshold`, `fast_threshold`, `bin_threshold`, `dyn_threshold`, `histo_to_thresh`,
`gray_histo`

Advanced:

`watersheds`, `watersheds_threshold`, `regiongrowing`, `regiongrowing_mean`, `var_threshold`

4.4.8 Process Regions

Standard:

`connection`, `select_shape`, `opening_circle`, `closing_circle`, `opening_rectangle1`,
`closing_rectangle1`, `difference`, `intersection`, `union1`, `shape_trans`, `fill_up`

Advanced:

`select_shape_proto`, `select_gray`, `clip_region`, `sort_region`, `skeleton`,
`partition_dynamic`, `rank_region`

Morphological operators can be found in the Reference Manual in the chapter “[Morphology](#)”.

4.4.9 Extract Features

Standard:

`area_center`, `smallest_rectangle1`, `smallest_rectangle2`, `compactness`, `eccentricity`,
`elliptic_axis`, `area_center_gray`, `intensity`, `min_max_gray`

Advanced:

`diameter_region`, `inner_rectangle1`, `inner_circle`, `gray_histo_abs`, `entropy_gray`

4.4.10 Transform Results Into World Coordinates

Standard:

`image_points_to_world_plane`

Advanced:

`gen_contour_region_xld`, `contour_to_world_plane_xld`

More operators for transforming results into world coordinates are described in the [Solution Guide III-C](#).

4.4.11 Visualize Results

Please refer to the [operator list for the method Visualization](#) (see section 19.4 on page 297).

4.5 Relation to Other Methods

4.5.1 Methods that are Useful for Blob Analysis

Color Processing (see [description](#) on page 173)

Color processing can be considered as an advanced way of blob analysis that uses three color channels instead of one gray value channel. HALCON provides operators for color space transformation, feature extraction, and pixel classification, which can be used in combination with blob analysis.

Texture Analysis (see [description](#) on page 189)

Texture analysis is a method for finding regular or irregular structures, e.g., on the surface of an object and it is therefore useful if texture is a feature of the object to be inspected and simple blob analysis is not sufficient. For texture analysis not only the single gray values but also a larger pixel neighborhood is used. HALCON provides filters that emphasize or suppress specific textures. The result of these filters can then be segmented.

4.5.2 Methods that are Using Blob Analysis

OCR (see [description](#) on page 245)

Blob analysis is typically used as a preprocessing step for OCR to segment the characters.

4.5.3 Alternatives to Blob Analysis

Edge Extraction (Subpixel-Precise) (see [description](#) on page 73)

In blob analysis a region is described by the gray values of its pixels. As an alternative, a region could be described by the change of the gray values at their borders. This approach is called edge detection.

Classification (see [description](#) on page 155)

To select specific gray values, thresholds must be determined. In most cases, fixed values are used or the current value is determined by a feature operator. In some cases it is useful if the system determines the ranges automatically. This can be achieved by using a classifier. In addition, a classifier can also be used to automatically distinguish between good and bad objects based on the extracted features and samples for both classes.

4.6 Tips & Tricks

4.6.1 Connected Components

By default, most HALCON segmentation operators like [threshold](#) return one region even if you see multiple not connected areas on the screen. To transform this region into separated objects (i.e., connected components in the HALCON nomenclature) one has to call [connection](#).

4.6.2 Speed Up

Many online applications require maximum speed. Because of its flexibility, HALCON offers many ways to achieve this goal. Here the most common methods are listed.

- Regions of interest are the standard method to increase the speed by processing only those areas where objects need to be inspected. This can be done using pre-defined regions but also by an online generation of the regions of interest that depend on other objects found in the image.
- If an object has a specific minimum size, the operator [fast_threshold](#) is a fast alternative to [threshold](#). This kind of fast operator can also directly be generated by using operators like [gen_grid_region](#) and [reduce_domain](#) before calling the thresholding operator.
- By default, HALCON performs some data consistency checks. These can be switched off using [set_check](#).
- By default, HALCON initializes new images. Using [set_system](#) with the parameter "init_new_image", this behavior can be changed.

4.7 Advanced Topics

4.7.1 Line Scan Cameras

In general, line scan cameras are treated like normal area sensors. In some cases, however, not single images but an “infinite” sequence of images showing objects, e.g., on a conveyor belt, must be processed. In this case the end of one image is the beginning of the next one. This means that objects that partially lie in both images must be combined into one object. For this purpose HALCON provides the operator [merge_regions_line_scan](#). This operator is called after the segmentation of one image, and combines the current objects with those of previous images. For more information see the [Solution Guide II-A](#).

4.7.2 High Accuracy

Sometimes high accuracy is required. This is difficult with blob analysis because objects are only extracted with integer pixel coordinates. Note, however, that many features that can be calculated for regions, e.g., the center of gravity, will be subpixel-precise. One way to get higher accuracy is to use a higher resolution. This has the effect that the higher number of pixels for each region results in better statistics to estimate features like the center of gravity ([area_center](#)). As an alternative, gray value features (like [area_center_gray](#)) can be used if the object fulfills specific gray value requirements. Here, the higher accuracy comes from the fact that for each pixel 255 values instead of one value (foreground or background) is used. If a very high accuracy is required, you should use the [subpixel-precise edge and line extraction](#) on page 73.

Chapter 5

1D Measuring

The idea of 1D measuring (also called 1D metrology or caliper) is very intuitive: Along a predefined region of interest, edges are located that are mainly perpendicular to the orientation of the region of interest. Here, edges are defined as transitions from dark to bright or from bright to dark.

Based on the extracted edges, you can measure the dimensions of parts. For example, you can measure the width of a part by placing a region of interest over it and locating the edges on its left and the right side. The effect of this can be seen in [figure 5.1a](#), whereas [figure 5.1b](#) shows the corresponding gray value profile.




Figure 5.1: (a) Measuring a fuse wire; (b) gray value profile along the region of measurement with extracted edges.


In addition to these simple rectangular regions of interest, circular arcs can be used to measure, e.g., the widths of the cogs on a cog wheel.

The advantage of the measure approach is its ease of use combined with a short execution time and a very high accuracy. With only a few operators, high-performing applications can be realized.

Alternatively, you can use HDevelop's Measure Assistant, which allows you to perform measurements with just a few mouse clicks. How to measure with this assistant is described in detail in the HDevelop User's Guide, [section 6.4](#) on page 248.

5.1 Basic Concept

Measuring consists of four main steps:



5.1.1 Acquire Image(s)

First, an image is acquired.

For detailed information see the [description of this method](#) on page 13.

5.1.2 Create Measure Object

Having acquired the image, you specify where to measure, i.e., you describe the position, orientation, etc. of the line or arc along which you want to measure. Together with some other parameters, this information is stored in the so-called measure object.

You access the measure object by using a so-called handle. Similarly to a file handle, this handle is needed when working with the tool. Each time the measure tool is executed, this handle is passed as a parameter.

In object-oriented languages like C++ it is possible to use the measure class instead of the low-level approach with handles. Here, creation and destruction are realized with the standard object-oriented methods.

5.1.3 Measure

Then, you can apply the measuring by specifying the measure object and some other vision parameters like, e.g., the minimum contrast. You can find detailed information about this step in the Solution Guide III-A in [chapter 3](#) on page 17.

5.1.4 Destroy Measure Object

When you no longer need the measure object, you destroy it by passing the handle to `close_measure`.

5.2 Extended Concept

In many cases, a measuring application will be more complex than described above. Reasons for this are, e.g., clutter or inhomogeneous illumination. Furthermore, post-processing like transforming the features to real-world units, or visualization of results may be required.

5.2.1 Radiometrically Calibrate Image(s)

To allow high-accuracy measurements, the camera should have a linear response function, i.e., the gray values in the images should depend linearly on the incoming energy. Since some cameras do not have a linear response function, HALCON provides the so-called radiometric calibration (gray value calibration): With the operator `radiometric_self_calibration` you can determine the inverse response function of the camera (offline) and then apply this function to the images using `lut_trans` before performing the measuring.

5.2.2 Align ROIs Or Images

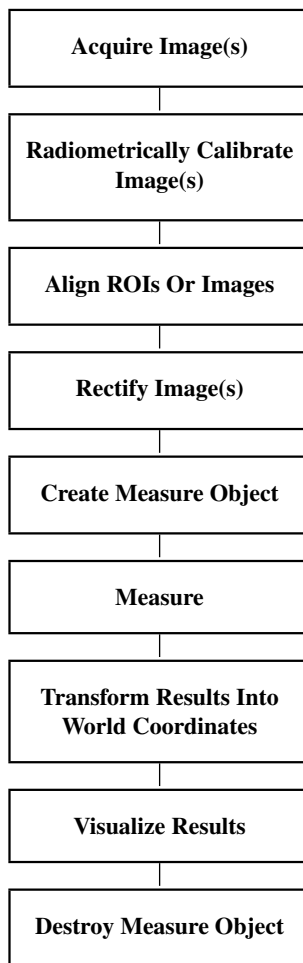
In some applications, the line or arc along which you want to measure, must be aligned relative to another object.

How to perform alignment using shape-based matching is described in the Solution Guide II-B in [section 2.4.3.2](#) on page 42.

5.2.3 Rectify Image(s)

Similarly to alignment, it may be necessary to rectify the image, e.g., to remove lens distortion.

Detailed information about rectifying images can be found in the Solution Guide III-C in [section 3.4](#) on page 62.



5.2.4 Create Measure Object

You can teach the measurement line or arc interactively with operators like [draw_rectangle2](#) or read its parameters from file ([read_string](#)). As an alternative, its coordinates can be generated based on the results of other vision tools like Blob Analysis (see the [description of this method](#) on page 31). In particular, the measurement line or arc may need to be aligned to a certain object as described above.

If the measurement is always performed along the same line or arc, you can create the measure object offline and then use it multiple times before destroying it. However, if you want to align the measurement, the position and orientation of the line or arc will differ for each image. In this case, you must create a new measure object for each image. An exception to this rule is if only the position changes but not the orientation. Then, you can keep the measure object and adapt its position via [translate_measure](#).

Please refer to the Solution Guide III-A, [chapter 2](#) on page 11, for more information.

5.2.5 Transform Results Into World Coordinates

If you have calibrated your vision system, you can easily transform the results of measuring into world coordinates with `image_points_to_world_plane`. How to do this is described in the Solution Guide III-A in [section 3.5](#) on page 26.

This is described in detail in the Solution Guide III-C in [section 3.3](#) on page 57.

5.2.6 Visualize Results

The best way to visualize edge positions is to create (short) XLD line segments with operators like `gen_contour_polygon_xld`.

For detailed information see the [description of this method](#) on page 289.

5.3 Programming Examples

The following examples gives a brief introduction to using the 1D measuring tool of HALCON. The longest parts are pre- and postprocessing; the measurement itself consists only of two operator calls. Further examples are described in the [Solution Guide III-A](#).

5.3.1 Inspecting a Fuse

Example: `solution_guide/basics/fuse.hdev`

Preprocessing consists of the generation of the measurement line. In the example program, this step is accomplished by assigning the measure object's parameters to variables.

```
read_image (Fuse, 'fuse')
Row := 297
Column := 545
Length1 := 80
Length2 := 10
Angle := rad(90)
gen_measure_rectangle2 (Row, Column, Angle, Length1, Length2, Width, Height, \
                        'bilinear', MeasureHandle)
```

Now the actual measurement is performed by applying the measure object to the image. The parameters are chosen such that edges around dark areas are grouped to so called pairs, returning the position of the edges together with the width and the distance of the pairs.

```
measure_pairs (Fuse, MeasureHandle, 1, 1, 'negative', 'all', RowEdgeFirst, \
              ColumnEdgeFirst, AmplitudeFirst, RowEdgeSecond, \
              ColumnEdgeSecond, AmplitudeSecond, IntraDistance, \
              InterDistance)
```




Figure 5.2: Measuring the width of the fuse wire.

The last part of the program displays the results by generating a region with the parameters of the measurement line and converting the edge positions to short XLD contours (see [figure 5.2](#)).

```
for i := 0 to |RowEdgeFirst|-1 by 1
    gen_contour_polygon_xld (EdgeFirst, \
        [-sin(Angle+rad(90))*Length2+RowEdgeFirst[i], \
        -sin(Angle-rad(90))*Length2+RowEdgeFirst[i], \
        [cos(Angle+rad(90))*Length2+ColumnEdgeFirst[i], \
        cos(Angle-rad(90))*Length2+ColumnEdgeFirst[i]])
    gen_contour_polygon_xld (EdgeSecond, \
        [-sin(Angle+rad(90))*Length2+RowEdgeSecond[i], \
        -sin(Angle-rad(90))*Length2+RowEdgeSecond[i], \
        [cos(Angle+rad(90))*Length2+ColumnEdgeSecond[i], \
        cos(Angle-rad(90))*Length2+ColumnEdgeSecond[i]])
    write_string (WindowID, 'width: '+IntraDistance[i]+' pix')
endfor
```

5.3.2 Inspect Cast Part

Example: `hdevelop/Applications/Measuring-2D/measure_arc.hdev`

The task of this example is to inspect the distance between elongated holes of a cast part after chamfering (see [figure 5.3](#)). Note that to achieve best accuracy it would be recommended to use backlight combined with a telecentric lens instead of the depicted setup.

This task can be solved easily by using the measure tool with a circular measurement ROI. The center of the ROI is placed into the center of the cast part; its radius is set to the distance of the elongated holes from the center.




Figure 5.3: Measuring the distance between the holes.

```
Row := 275
Column := 335
Radius := 107
AngleStart := -rad(55)
AngleExtent := rad(170)
gen_measure_arc (Row, Column, Radius, AngleStart, AngleExtent, 10, Width, \
                Height, 'nearest_neighbor', MeasureHandle)
```

Now, the distance between the holes can be measured with a single operator call:

```
measure_pos (Zeiss1, MeasureHandle, 1, 10, 'all', 'all', RowEdge, \
            ColumnEdge, Amplitude, Distance)
```

5.3.3 Inspecting an IC Using Fuzzy Measuring

Example: `hdevelop/Applications/Measuring-2D/fuzzy_measure_pin.hdev`

The task of this example is to inspect the lead width and the lead distance of the IC depicted in [figure 5.4](#).

The illumination conditions in this example are quite difficult. This has the effect that four edges are visible for each lead. Fuzzy rules are used to restrict the measurement to the correct (outer) leads.

```
gen_measure_rectangle2 (Row1, Col1, Phi1, Length1, Length2, Width, Height, \
                        'nearest_neighbor', MeasureHandle1)
create_funct_1d_pairs ([0.0, 0.3], [1.0, 0.0], FuzzyAbsSizeDiffFunction)
set_fuzzy_measure_norm_pair (MeasureHandle1, 11.0, 'size_abs_diff', \
                             FuzzyAbsSizeDiffFunction)
fuzzy_measure_pairs (Image, MeasureHandle1, 1, 30, 0.5, 'positive', \
                    RowEdgeFirst1, ColumnEdgeFirst1, AmplitudeFirst1, \
                    RowEdgeSecond1, ColumnEdgeSecond1, AmplitudeSecond1, \
                    RowEdgeMiddle1, ColumnEdgeMiddle1, FuzzyScore1, \
                    IntraDistance1, InterDistance1)
```




Figure 5.4: Measuring the width and distance of the leads.

5.3.4 Measuring Leads of a Moving IC

Example: `hdevelop/Applications/Measuring-2D/pm_measure_board.hdev`

The task of this example is to measure the positions of the leads of a chip (see [figure 5.5](#)). Because the chip can appear at varying positions and angles, the regions of interest used for the measurement must be aligned.




Figure 5.5: (a) Model image with measurement ROIs; (b) measuring the leads in the aligned ROIs.

In this case, the alignment is achieved by searching for the print on the chip using shape-based matching (see [Matching](#) on page 97).

```

gen_rectangle1 (Rectangle, Row1, Column1, Row2, Column2)
reduce_domain (Image, Rectangle, ImageReduced)
create_shape_model (ImageReduced, 4, 0, rad(360), rad(1), 'none', \
                    'use_polarity', 30, 10, ModelID)

```

After the print has been found, the positions of the measurement ROIs are transformed relative to the position of the print.

```

find_shape_model (ImageCheck, ModelID, 0, rad(360), 0.7, 1, 0.5, \
                  'least_squares', 4, 0.7, RowCheck, ColumnCheck, \
                  AngleCheck, Score)

hom_mat2d_identity (HomMat2DIdentity)
hom_mat2d_translate (HomMat2DIdentity, RowCheck, ColumnCheck, \
                    HomMat2DTranslate)
hom_mat2d_rotate (HomMat2DTranslate, AngleCheck, RowCheck, ColumnCheck, \
                  HomMat2DRotate)
affine_trans_pixel (HomMat2DRotate, Rect1Row, Rect1Col, Rect1RowCheck, \
                    Rect1ColCheck)

```

Then, the measure tools are created and the measurement is applied.

```

gen_measure_rectangle2 (Rect1RowCheck, Rect1ColCheck, AngleCheck, \
                       RectLength1, RectLength2, Width, Height, \
                       'bilinear', MeasureHandle1)
measure_pairs (ImageCheck, MeasureHandle1, 2, 90, 'positive', 'all', \
              RowEdgeFirst1, ColumnEdgeFirst1, AmplitudeFirst1, \
              RowEdgeSecond1, ColumnEdgeSecond1, AmplitudeSecond1, \
              IntraDistance1, InterDistance1)

```

5.3.5 Inspect IC

Example: `hdevelop/Applications/Measuring-2D/measure_pin.hdev`

The task of this example is to inspect major dimensions of an IC (see [figure 5.6](#)).

In the first step the extent of each lead and the distance between the leads is measured. For this, a rectangle that contains the leads is defined (see [figure 5.6a](#)), which is used to generate the measure object. This is used to extract pairs of straight edges that lie perpendicular to the major axis of the rectangle.

```

gen_measure_rectangle2 (Row, Column, Phi, Length1, Length2, Width, Height, \
                       'nearest_neighbor', MeasureHandle)
measure_pairs (Image, MeasureHandle, 1.5, 30, 'negative', 'all', \
              RowEdgeFirst, ColumnEdgeFirst, AmplitudeFirst, \
              RowEdgeSecond, ColumnEdgeSecond, AmplitudeSecond, PinWidth, \
              PinDistance)

```

From the extracted pairs of straight edges, the number of leads, their average width, and the average distance between them is derived.




Figure 5.6: Measuring the dimensions of leads: (a) width of the leads and distance between them; (b) length of the leads.

```
numPins := |PinWidth|
avgPinWidth := sum(PinWidth)/|PinWidth|
avgPinDistance := sum(PinDistance)/|PinDistance|
```

The second part shows the power of the measure tool: The length of the leads is determined. This is possible although each lead has a width of only a few pixels. For this, a new measure object is generated based on a rectangle that contains the leads on two opposite sides of the IC (see [figure 5.6b](#)). The distance between the first and the second found edge is the length of the upper leads, and the distance between the third and the fourth edge is the length of the lower leads.

```
gen_measure_rectangle2 (Row, Column, Phi, Length1, Length2, Width, Height, \
                        'nearest_neighbor', MeasureHandle)
measure_pos (Image, MeasureHandle, 1.5, 30, 'all', 'all', RowEdge, \
            ColumnEdge, Amplitude, Distance)
```

5.4 Selecting Operators

5.4.1 Acquire Image(s)

Please refer to the [operator list for the method Image Acquisition](#) (see section 2.4 on page 17).

5.4.2 Radiometrically Calibrate Image(s)

Standard:

`radiometric_self_calibration, lut_trans`

5.4.3 Align ROIs Or Images

Operators for aligning ROIs or images are described in the [Solution Guide II-B](#).

5.4.4 Rectify Image(s)

Operators for rectifying images are described in the [Solution Guide III-C](#).

5.4.5 Create Measure Object

Standard:

`gen_measure_rectangle2, gen_measure_arc, translate_measure`

5.4.6 Measure

Standard:

`measure_pos, measure_pairs`

Advanced:

`set_fuzzy_measure, fuzzy_measure_pos, fuzzy_measure_pairs, fuzzy_measure_pairing`

5.4.7 Transform Results Into World Coordinates

Standard:

`image_points_to_world_plane`

Advanced:

`gen_contour_region_xld, contour_to_world_plane_xld`

More operators for transforming results into world coordinates are described in the [Solution Guide III-C](#).

5.4.8 Visualize Results

Advanced:

[gen_contour_polygon_xld](#)

Please refer to the [operator list for the method Visualization](#) (see section 19.4 on page 297).

5.4.9 Destroy Measure Object

Standard:

[close_measure](#)

5.5 Relation to Other Methods

5.5.1 Alternatives to 1D Measuring

Edge Extraction (Subpixel-Precise) (see [description](#) on page 73)

A very flexible way to measure parameters of edges is to extract the edge contour with [edges_sub_pix](#). The advantage of this approach is that it can handle free-form shapes. Furthermore, it allows to determine attributes like the edge direction for each edge point.

5.6 Tips & Tricks

5.6.1 Suppress Clutter or Noise

In many applications there is clutter or noise that must be suppressed. The measure operators offer multiple approaches to achieve this. The best one is to increase the threshold for the edge extraction to eliminate faint edges. In addition, the value for the smoothing parameter can be increased to smooth irrelevant edges away.

When grouping edges to pairs, noise edges can lead to an incorrect grouping if they are in the vicinity of the “real” edge and have the same polarity. In such a case you can suppress the noise edges by selecting only the strongest edges of a sequence of consecutive rising and falling edges.

5.6.2 Reuse Measure Object

Because the creation of a measure object needs some time, we recommend to reuse them if possible. If no alignment is needed, the measure object can, for example, be created offline and reused for each image. If the alignment involves only a translation, [translate_measure](#) can be used to correct the position.

5.6.3 Use an Absolute Gray Value Threshold

As an alternative to edge extraction, the measurements can be performed based on an absolute gray value threshold by using the operator [measure_thresh](#). Here, all positions where the gray value crosses the given threshold are selected.

5.7 Advanced Topics

5.7.1 Fuzzy Measuring

In case there are extra edges that do not belong to the measurement, HALCON offers an extended version of measuring: fuzzy measuring. This tool allows to define so-called fuzzy rules, which describe the features of good edges. Possible features are, e.g., the position, the distance, the gray values, or the amplitude of edges. These functions are created with [create_funct_1d_pairs](#) and passed to the tool with [set_fuzzy_measure](#). Based on these rules, the tool will select the most appropriate edges.

The advantage of this approach is the flexibility to deal with extra edges even if a very low minimum threshold or smoothing is used. An example for this approach is the example program [fuzzy_measure_pin.hdev](#) on page 55.

Please refer to the Solution Guide III-A, [chapter 4](#) on page 33, for more information.

5.7.2 Evaluation of Gray Values

To have full control over the evaluation of the gray values along the measurement line or arc, you can use [measure_projection](#). The operator returns the projected gray values as an array of numbers, which can then be further processed with HALCON operators for tuple or function processing (see the chapters “[Tuple](#)” and “[Tools > Function](#)” in the Reference Manual). Please refer to the Solution Guide III-A, [section 3.4](#) on page 22, for more information.

Chapter 6

Edge Extraction (Pixel-Precise)

The traditional way of finding edges, i.e., dark / light transitions in an image, is to apply an edge filter. These filters have the effect to find pixels at the border between light and dark areas. In mathematical terms this means that these filters determine the image gradient. This image gradient is typically returned as the edge amplitude and/or the edge direction. By selecting all pixels with a high edge amplitude, contours between areas can be extracted.

HALCON offers all standard edge filters like the Sobel, Roberts, Robinson, or Frei filters. Besides these, post-processing operators like hysteresis thresholding or non-maximum suppression are provided. In addition, state-of-the-art filters that determine the edge amplitude and edge direction accurately are provided. This enables you to apply the filters in a flexible manner.





Figure 6.1: Result of applying an edge filter: (a) amplitude, (b) direction, (c) extracted edges.

Please note that in addition to this classical approach, HALCON provides advanced operators for

subpixel-precise edge and line extraction (see the [description of this method](#) on page 73) and for successive post-processing and feature extraction.

6.1 Basic Concept

Using edge filters typically consists of three basic steps:



6.1.1 Acquire Image(s)

First, an image is acquired.

For detailed information see the [description of this method](#) on page 13.

6.1.2 Filter Image

On the input image, an edge filter is applied. This operation results in one or two images. The basic result is the edge amplitude, which is typically stored as a byte image, with the gray value of each pixel representing the local edge amplitude. Optionally, the direction of the edges is returned. These values are stored in a so-called direction image, with the values 0...179 representing the angle in degrees divided by two.

6.1.3 Extract Edges

The result of applying the edge filter is an image containing the edge amplitudes. From this image, the edges are extracted by selecting the pixels with a given minimum edge amplitude using a threshold operator. The resulting edges are typically broader than one pixel and therefore have to be thinned. For this step, various methods are available.

6.1.4 Process Edges


Having extracted the edges, they can be further processed: You can convert the edge regions into another data structure for a potential further processing and for extracting features. Besides, you can extract the regions enclosed by the edges.

6.1.5 A First Example

The following program shows an example for the basic concept of edge filters. As an edge filter, `sobel_amp` is applied with the mode `'thin_sum_abs'` to get thin edges together with a 3x3 filter mask. Then, the operator `threshold` is used to extract all pixels with an edge amplitude higher than 20. The resulting region contains some areas where the edge is wider than one pixel. Therefore, the operator `skeleton` is applied to thin all edges completely. The result is depicted in [figure 6.1c](#) on page 63.

```
read_image (Image, 'fuse')
sobel_amp (Image, EdgeAmplitude, 'thin_sum_abs', 3)
threshold (EdgeAmplitude, Region, 20, 255)
skeleton (Region, Skeleton)
```

6.2 Extended Concept



6.2.1 Use Region Of Interest

Edge extraction can be sped up by using a region of interest. The more the region in which edge filtering is performed can be restricted, the faster and more robust the extraction will be.

For detailed information see the [description of this method](#) on page 19.

6.2.2 Filter Image

HALCON offers a wide range of edge filters. One of the most popular filters is the Sobel filter. This is the best of the old-fashioned filters. It combines speed with a reasonable quality. The corresponding operators are called [sobel_amp](#) and [sobel_dir](#).

In contrast, [edges_image](#) provides the state of the art of edge filters. This operator is actually more than just a filter. It includes a thinning of the edges using a non-maximum suppression and a hysteresis threshold for the selection of significant edge points. It also returns the edge direction and the edge amplitude very accurately, which is not the case with the Sobel filter. This operator is recommended if higher quality is more important than a longer execution time. If the images are not noisy or blurred, you can even combine accuracy and speed by using the mode 'sobel_fast' inside [edges_image](#). The corresponding operator to find edges in multi-channel images, e.g., a color image, is [edges_color](#).

6.2.3 Extract Edges

The easiest way to extract the edges from the edge amplitude image is to apply [threshold](#) to select pixels with a high edge amplitude. The result of this step is a region that contains all edge points. With [skeleton](#), these edges can be thinned to a width of one pixel. As an advanced version for [threshold](#), [hysteresis_threshold](#) can be used to eliminate insignificant edges. A further advanced option is to call the operator [nonmax_suppression_dir](#) before [skeleton](#), which in difficult cases may result in more accurate edges. Note that in order to use this operator you must have computed the edge direction image.

In contrast, the advanced filter [edges_image](#) already includes the non-maximum suppression and the hysteresis threshold. Therefore, in this case a simple [threshold](#) suffices to extract edges that are one pixel wide.

If only the edge points as a region are needed, the operator [inspect_shape_model](#) can be used. Here, all steps including edge filtering, non-maximum suppression, and hysteresis thresholding are performed in one step with high efficiency.

6.2.4 Process Edges

If you want to extract the coordinates of edge segments, [split_skeleton_lines](#) is the right choice. This operator must be called for each connected component (result of [connection](#)) and returns all the control points of the line segments. As an alternative, a Hough transform can be used to obtain the line segments. Here, the operators [hough_lines_dir](#) and [hough_lines](#) are available. You can also convert the edge region into XLD contours by using, e.g., the operator [gen_contours_skeleton_xld](#).

The advantage of this approach is the extended set of operators offered for [XLD contour processing](#) on page 83, e.g., for contour segmentation, feature extraction, or approximation.

You can extract the regions enclosed by the edges easily using [background_seg](#). If regions merge because of gaps in the edges, the operators [close_edges](#) or [close_edges_length](#) can be used in advance to close the gaps before regions are extracted. As an alternative, morphological operators like [opening_circle](#) can be applied to the output regions of [background_seg](#). In general, all operators described for the method [Process Regions](#) on page 35 can be applied here as well.

6.2.5 Visualize Results

Finally, you might want to display the images, the edges (regions), and the line segments.

For detailed information see the [description of this method](#) on page 289.

6.3 Programming Examples

This section gives a brief introduction to using HALCON for edge filtering and edge extraction.

6.3.1 Aerial Image Interpretation

Example: `solution_guide/basics/edge_segments.hdev`

[Figure 6.2](#) shows an image taken from an aeroplane. The task is to extract the edges of roads and buildings as a basis for the image interpretation.




Figure 6.2: (a) Extracting edges and (b) approximating them by segments.

The extraction of edges is very simple and reliable when using the operator [edges_image](#). This operator returns both the edge amplitude and the edge direction. Here, the parameters are selected such that a non-maximum suppression (parameter value 'nms') and a hysteresis threshold (threshold values 20 and 40) are performed. The non-maximum suppression has the effect that only pixels in the center of the edge are

returned, together with the corresponding values for the amplitude and the direction. All other pixels are set to zero. Therefore, a threshold with the minimum amplitude of 1 is sufficient here. As a preparation for the next step, the edge contour regions are split up into their connected components.

```
read_image (Image, 'mreut')
edges_image (Image, ImaAmp, ImaDir, 'lanser2', 0.5, 'nms', 20, 40)
threshold (ImaAmp, Region, 1, 255)
connection (Region, ConnectedRegions)
```

The rest of the example program converts the region data into numeric values. To be more precise: the edges are approximated by individual line segments. This is performed by calling [split_skeleton_lines](#) for each connected component. The result of this call are four tuples that contain the start and the end coordinates of the line segments. For display purposes, each of these line segments is converted into an XLD contour.

```
count_obj (ConnectedRegions, Number)
gen_empty_obj (XLDContours)
for i := 1 to Number by 1
    select_obj (ConnectedRegions, SingleEdgeObject, i)
    split_skeleton_lines (SingleEdgeObject, 2, BeginRow, BeginCol, EndRow, \
                        EndCol)
    for k := 0 to |BeginRow|-1 by 1
        gen_contour_polygon_xld (Contour, [BeginRow[k], EndRow[k]], \
                                    [BeginCol[k], EndCol[k]])
        concat_obj (XLDContours, Contour, XLDContours)
    endfor
endfor
dev_display (XLDContours)
```

6.3.2 Segmenting a Color Image

Example: `hdevelop/Filters/Edges/edges_color.hdev`

The task of this example is to segment the color image depicted in [figure 6.3](#).

The example demonstrates the possibilities of a multi-channel edge filter. First, the gray value image is derived from the color information to show that some object borders can no longer be seen. For example, the (green) soccer field cannot be distinguished from the surrounding (red) track.

```
read_image (Image, 'olympic_stadium')
rgb1_to_gray (Image, GrayImage)
```

The color edge filter is applied and the edge amplitude is displayed. If you compare this to the filter result of the gray image the difference can be easily seen.

```
edges_color (Image, ImaAmp, ImaDir, 'canny', 1, 'none', -1, -1)
edges_image (GrayImage, ImaAmpGray, ImaDirGray, 'canny', 1, 'none', -1, -1)
```

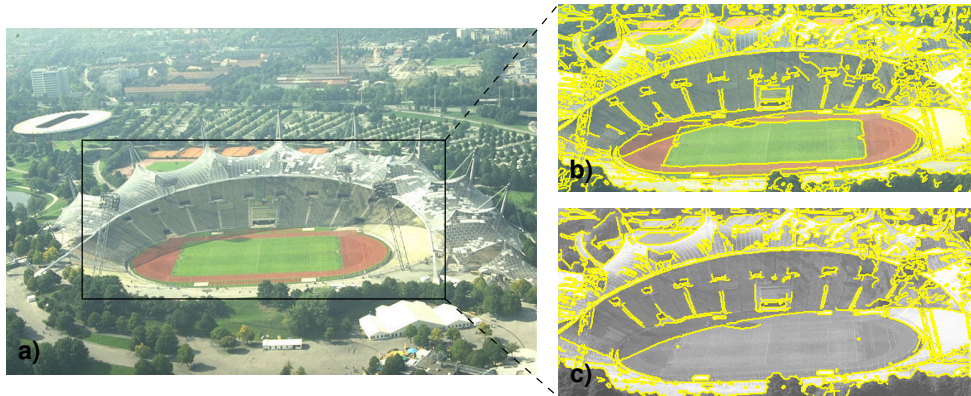


Figure 6.3: (a) Original image; (b) extracted color edges, overlaid on the color image; (c) extracted gray value edges, overlaid on the gray value image.

Finally, the edge segments are extracted for both the color and the gray image and overlaid on the original image.

```
edges_color (Image, ImaAmpHyst, ImaDirHyst, 'canny', 1, 'nms', 20, 40)
threshold (ImaAmpHyst, RegionColor, 1, 255)
skeleton (RegionColor, EdgesColor)
dev_display (Image)
dev_display (EdgesColor)
stop ()
edges_image (GrayImage, ImaAmpGrayHyst, ImaDirGrayHyst, 'canny', 1, 'nms', \
20, 40)
threshold (ImaAmpGrayHyst, RegionGray, 1, 255)
skeleton (RegionGray, EdgesGray)
dev_display (GrayImage)
dev_display (EdgesGray)
```

6.4 Selecting Operators

6.4.1 Acquire Image(s)

Please refer to the [operator list for the method Image Acquisition](#) (see section 2.4 on page 17).

6.4.2 Use Region Of Interest

Please refer to the [operator list for the method Region Of Interest](#) (see section 3.4 on page 26).

6.4.3 Filter Image

Standard:

`sobel_amp, sobel_dir, edges_image`

Advanced:

`derivate_gauss, edges_color`

6.4.4 Extract Edges

Standard:

`threshold, skeleton, inspect_shape_model`

Advanced:

`hysteresis_threshold, nonmax_suppression_dir`

6.4.5 Process Edges

Standard:

`background_seg, close_edges, close_edges_length, opening_circle,
split_skeleton_lines, hough_lines_dir, hough_lines, gen_contours_skeleton_xld`

6.4.6 Visualize Results

Please refer to the [operator list for the method Visualization](#) (see section 19.4 on page 297).

6.5 Relation to Other Methods

6.5.1 Alternatives to Edge Extraction (Pixel-Precise)

Blob Analysis (see [description](#) on page 31)

As an alternative to edge extraction, blob analysis can be used. This approach provides many methods from simple thresholding to region growing and watershed methods.

6.6 Tips & Tricks

6.6.1 Speed Up

Many online applications require maximum speed. Because of its flexibility, HALCON offers many ways to achieve this goal. Here the most common ones are listed.

- Regions of interest are the standard way to reduce the processing to only those areas where objects must be inspected. This can be achieved using pre-defined regions but also by an online generation of the regions of interest that depends on other objects found in the image.
- If high speed is important, the operators `sobel_amp` and `inspect_shape_model` are the preferred choice.
- By default, HALCON initializes new images. Using `set_system('init_new_image', 'false')`, this behavior can be changed to save execution time.

Chapter 7

Edge Extraction (Subpixel-Precise)

In addition to the traditional way of applying an edge filter to get the edge amplitude and thus the edges (see the method [Edge Extraction \(Pixel-Precise\)](#) on page 63), HALCON provides one-step operators that return subpixel-precise XLD contours. Besides this, not only edges but also lines can be extracted. This approach can also be applied to color images.

The advantage of this approach is its ease of use, because only a single operator call is needed. Furthermore, the accuracy and stability of the found contours is extremely high. Finally, HALCON offers a wide set of operators for the post-processing of the extracted contours, which includes, e.g., contour segmentation and fitting of circles, ellipses, and lines.





Figure 7.1: Result of contour extraction: (a) edge contours, (b) line contours.

This chapter covers only the extraction of contours. For information about processing them see the method [Contour Processing](#) on page 83.

7.1 Basic Concept

Extracting contours can easily be performed in a single step. Normally, no other operation is required.



7.1.1 Acquire Image(s)

First, an image is acquired as input for the process.

For detailed information see the [description of this method](#) on page 13.

7.1.2 Extract Edges Or Lines

HALCON offers various operators for the subpixel-accurate extraction of contours. The standard operator is based on the first derivative. It takes the image as input and returns the XLD contours. When using the second derivatives, first a Laplace operator must be executed before the contours along the zero crossings can be extracted. Besides the gray-value-based methods, HALCON provides the latest technology for the extraction of color edges.

Besides the extraction of edges, HALCON provides operators for the extraction of lines. In other systems lines are also called ridges. In contrast to edges, a line consists of two gray value transitions. Thus, a line can be considered as two parallel edges.


7.1.3 A First Example

The following program explains the basic concept of edge extraction. The only operator needed to extract edge contours is `edges_sub_pix`. It has the image as input and returns the XLD contours. Here, the filter 'lanser2' is selected with a medium-sized smoothing mask. The low value for the parameter Low ensures that contours are tracked even along low-contrast parts. To show that the result consists of multiple contours, the 12-color mode for visualization is selected. The result is depicted in [figure 7.1b](#) on page 73.

```
read_image (Image, 'mreut4_3')
edges_sub_pix (Image, Edges, 'lanser2', 0.5, 8, 50)
dev_set_colored (12)
dev_clear_window ()
dev_display (Edges)
```

7.2 Extended Concept

In addition to the extraction, optional steps can be performed.



7.2.1 Radiometrically Calibrate Image(s)

To extract edges or lines with high accuracy, the camera should have a linear response function, i.e., the gray values in the images should depend linearly on the incoming energy. Since some cameras do not have a linear response function, HALCON provides the so-called radiometric calibration (gray value calibration): With the operator `radiometric_self_calibration` you can determine the inverse response function of the camera (offline) and then apply this function to the images using `lut_trans` before performing the edge and line extraction.

7.2.2 Use Region Of Interest

Edge extraction can be sped up by using a region of interest. The more the region in which edges or lines are extracted can be restricted, the faster and more robust the extraction will be.

For detailed information see the [description of this method](#) on page 19.

7.2.3 Extract Edges Or Lines

The most often used operator for edge contour extraction is [edges_sub_pix](#). You can select various filter methods by specifying the corresponding name with the parameter `Filter`. For standard applications, common values are, e.g., 'canny' (based on a Gaussian convolution) or 'lanser2'. The advantage of 'lanser2' is the recursive implementation which has no increase in execution time when using a large smoothing. As a fast version the parameter value 'sobel_fast' can be used, which is recommended as long as the image is not noisy or blurred.

The operator [zero_crossing_sub_pix](#) can be used in combination with a filter like [derivate_gauss](#) with parameter value 'laplace'. The Laplace operator is mainly applied in the medical area.

To extract edges in multi-channel images, e.g., in a color image, HALCON provides the operator [edges_color_sub_pix](#). Similar to [edges_sub_pix](#), the parameter value 'sobel_fast' is recommended for a fast edge extraction as long as the image is not noisy or blurred.

The most commonly used operator for line extraction is [lines_gauss](#). Compared to [lines_facet](#) it is more robust and provides more flexibility. The width of lines that should be extracted is specified by the parameter `Sigma`: The wider the line, the larger the value must be chosen. For very wide lines we recommend to zoom down the image ([zoom_image_factor](#)) in order to reduce the overall execution time.

Like for edges, HALCON provides line extraction also for multi-channel images. The corresponding operator is [lines_color](#).

7.2.4 Determine Contour Attributes

The edge and line extraction operators not only provide the XLD contours but also so-called attributes. Attributes are numerical values; they are associated either with each control point of the contour (called contour attribute) or with each contour as a whole (global contour attribute). The operators [get_contour_attrib_xld](#) and [get_contour_global_attrib_xld](#) enable you to access these values by specifying the attribute name.

The attribute values are returned as tuples of numbers. Typical attributes for edges are, e.g., the edge amplitude and direction. For lines a typical attribute is the line width. The available attributes can be queried for a given contour with [query_contour_attribs_xld](#) and [query_contour_global_attribs_xld](#).

7.2.5 Process XLD Contours

Typically, the task is not finished by just extracting the contours and accessing the attributes. HALCON provides further processing like contour segmentation, feature extraction, or approximation.

For detailed information see the [description of this method](#) on page 83.

7.2.6 Transform Results Into World Coordinates

In many applications the coordinates of contours should be transformed into another coordinate system, e.g., into 3D world coordinates. After you have calibrated your vision system, you can easily perform the transformation with the operator `contour_to_world_plane_xld`. With this approach you can also eliminate lens distortions and perspective distortions.

This is described in detail in the Solution Guide III-C in [section 3.3](#) on page 57.

7.2.7 Visualize Results

Finally, you might want to display the images and the contours.

For detailed information see the [description of this method](#) on page 289.

7.3 Programming Examples

This section gives a brief introduction to using HALCON for edge extraction.

7.3.1 Measuring the Diameter of Drilled Holes

Example: `solution_guide/basics/rim_simple.hdev`

[Figure 7.2](#) shows an image of a car rim. The task is to measure the diameters of the drilled holes.




Figure 7.2: (a) automatically determined ROIs; (b) extracted edges; (c) computed ellipses and diameters.

First, a segmentation step is performed to roughly find the borders of the holes. The actual edge extraction is then performed only in these regions of interest (ROIs). This has two advantages: First, there are many edges in the image that are of no interest for the measurement. By restricting the processing to ROIs you can easily select the relevant objects. Secondly, the contour extraction is time-consuming. Thus, a reduced domain is an efficient way to speed-up the process.

Locating the holes is quite easy: First, all dark pixels are selected. After selecting all those connected components that are circular and have a certain size, only the holes remain. Finally, the regions of interest

are obtained by accessing the borders of the holes and dilating them. The resulting ROIs are depicted in [figure 7.2a](#).

```
threshold (Image, Dark, 0, 128)
connection (Dark, DarkRegions)
select_shape (DarkRegions, Circles, ['circularity','area'], 'and', [0.85, \
    50], [1.0,99999])
boundary (Circles, RegionBorder, 'inner')
dilation_circle (RegionBorder, RegionDilation, 6.5)
union1 (RegionDilation, ROIEdges)
```

Calling `reduce_domain` changes the domain of the image to the prepared region of interest. Now, the edge extractor can be applied (see [figure 7.2b](#)).

```
reduce_domain (Image, ROIEdges, ImageROI)
edges_sub_pix (ImageROI, Edges, 'lanser2', 0.3, 10, 30)
```

The extracted contours are further processed to determine their diameter: With `fit_ellipse_contour_xld`, ellipses are fitted to the contours. In other words, those ellipses are determined that fit the extracted contours as closely as possible. The operator returns the parameters of the ellipses. With the operator `gen_ellipse_contour_xld`, the corresponding ellipses are created and displayed (compare [figure 7.2b](#) and [figure 7.2c](#)). Another option is to use the operator `gen_circle_contour_xld`.

```
fit_ellipse_contour_xld (Edges, 'ftukey', -1, 2, 0, 200, 3, 2, Row, Column, \
    Phi, Ra, Rb, StartPhi, EndPhi, PointOrder)
NumHoles := |Ra|
gen_ellipse_contour_xld (ContEllipse, Row, Column, Phi, Ra, Rb, \
    gen_tuple_const(NumHoles,0), \
    gen_tuple_const(NumHoles,rad(360)), \
    gen_tuple_const(NumHoles,'positive'), 1)
```

The diameters can easily be computed from the ellipse parameters and then be displayed in the image using `write_string` (see [figure 7.2c](#)).

```
for i := 0 to NumHoles-1 by 1
    write_string (WindowID, 'D1=' + 2*Ra[i])
    write_string (WindowID, 'D2=' + 2*Rb[i])
endfor
```

7.3.2 Angiography

Example: `hdevelop/Filters/Lines/lines_gauss.hdev`

The task of this example is to extract the blood vessels in the X-ray image of the heart depicted in [figure 7.3](#). The vessels are emphasized by using a contrast medium. For the diagnosis it is important to extract the width of the vessels to determine locally narrowed parts (stenoses).




Figure 7.3: (a) X-ray image of the heart; (b) extracted blood vessels.

The vessels are extracted using `lines_gauss`. The result of this operator are the centers of the vessels in the form of XLD contours. Besides this, attributes are associated with the contour points, one of which is the local line width. This width is requested and displayed as contours.

```
lines_gauss (Angio, Lines, 2.3, 0.0, 0.7, 'dark', 'true', 'parabolic', \
            'true')
count_obj (Lines, Number)
for I := 1 to Number by 1
    select_obj (Lines, Line, I)
    get_contour_xld (Line, Row, Col)
    get_contour_attrib_xld (Line, 'angle', Angle)
    get_contour_attrib_xld (Line, 'width_left', WidthL)
    get_contour_attrib_xld (Line, 'width_right', WidthR)
    RowR := Row+cos(Angle)*WidthR*sqrt(0.75)
    ColR := Col+sin(Angle)*WidthR*sqrt(0.75)
    RowL := Row-cos(Angle)*WidthL*sqrt(0.75)
    ColL := Col-sin(Angle)*WidthL*sqrt(0.75)
    disp_polygon (WindowID, RowL, ColL)
    disp_polygon (WindowID, RowR, ColR)
endfor
```

7.4 Selecting Operators

7.4.1 Acquire Image(s)

Please refer to the [operator list for the method Image Acquisition](#) (see section 2.4 on page 17).

7.4.2 Radiometrically Calibrate Image(s)

Standard:

`radiometric_self_calibration, lut_trans`

7.4.3 Use Region Of Interest

Please refer to the [operator list for the method Region Of Interest](#) (see section 3.4 on page 26).

7.4.4 Extract Edges Or Lines

Standard:

`edges_sub_pix, derivate_gauss, lines_gauss, lines_facet`

Advanced:

`zero_crossing_sub_pix, edges_color_sub_pix, lines_color`

7.4.5 Determine Contour Attributes

Standard:

`get_contour_attrib_xld, get_contour_global_attrib_xld,
query_contour_attribs_xld, query_contour_global_attribs_xld`

7.4.6 Process XLD Contours

Please refer to the [operator list for the method Contour Processing](#) (see section 8.4 on page 93).

7.4.7 Transform Results Into World Coordinates

Standard:

```
contour_to_world_plane_xld
```

More operators for transforming results into world coordinates are described in the [Solution Guide III-C](#).

7.4.8 Visualize Results

Please refer to the [operator list for the method Visualization](#) (see section 19.4 on page 297).

7.5 Relation to Other Methods

7.5.1 Alternatives to Edge Extraction (Subpixel-Precise)

Subpixel Thresholding

Besides the subpixel-accurate edge and line extractors, HALCON provides a subpixel-accurate threshold operator called `threshold_sub_pix`. If the illumination conditions are stable, this can be a fast alternative.

Subpixel Point Extraction

In addition to the contour-based subpixel-accurate data, HALCON offers subpixel-accurate point operators for various applications. In the reference manual, these operators can be found in the chapter “[Filters > Points](#)”.

Chapter 8

Contour Processing

One of HALCON's powerful tool sets are the subpixel-accurate contours. Contours belong to the data type XLD (see Quick Guide in [section 2.1.2.3](#) on page 18 for more information). These contours are typically the result of some kind of image processing and represent, e.g., the borders of objects. [Figure 8.1a](#) shows such edges overlaid on the original image; [Figure 8.1b](#) zooms into the rectangular area marked in [Figure 8.1a](#) and furthermore highlights the so-called control points of the contours with crosses. Here, you can clearly see the highly accurate positioning of the control points.

HALCON provides operators to perform advanced types of measurements with these contours. For example, the contours can be segmented into lines and circular or elliptic arcs (see [Figure 8.1c](#)). The parameters of these segments, e.g., their angle, center, or radius, can then be determined and used, e.g., in the context of a measuring task.


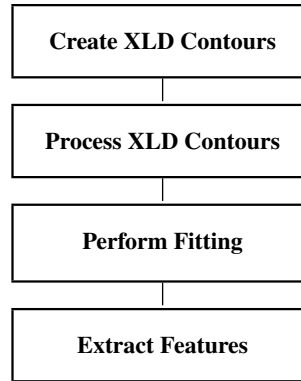


Figure 8.1: XLD contours: (a) edge contours, (b) zoom into rectangular area, (c) segmented lines and elliptic arcs.

The advantage of contour processing is twofold: First, its high accuracy enables reliable measurements. Secondly, the extensive and flexible set of operators provided for this data type enables you to solve problems that cannot be solved with classical methods like 1D measuring. More detailed information about contour processing can be found in the [Solution Guide III-B](#).

8.1 Basic Concept

The processing of contours consists of multiple steps that can be combined in a flexible way.



8.1.1 Create XLD Contours

The most common way to create XLD contours is to apply one of the subpixel-accurate extraction operators described for the method [Extract Edges Or Lines](#) on page 76. As an alternative, an edge filter with some post-processing can be used. The resulting regions are then converted to XLD contours. Please note that this approach is only pixel-accurate. For more information about this approach see the method [Edge Extraction \(Pixel-Precise\)](#) on page 63.

8.1.2 Process XLD Contours

Typically, only certain contours of an object are used for an inspection task. One possibility to restrict the extraction of contours to the desired ones is to use a well-fitting region of interest as, e.g., depicted in [figure 8.2a](#): The rectangular ROI just covers the upper part of the blades. When applying an edge extractor, exactly one contour on each side of the objects is found.

In many cases, however, not only the desired contours are extracted. An example is depicted in [figure 8.2b](#), where the ROI was chosen too large. Thus, the contours must be processed to obtain the desired parts of the contours. In the example, the contours are segmented into parts and only parallel segments with a certain length are selected (see the result in [figure 8.2c](#)).

Another reason for processing contours occurs if the extraction returns unwanted contours caused by noise or texture or if there are gaps between contours because of a low contrast or contour intersections.

8.1.3 Perform Fitting

Having obtained contour segments that represent a line, a rectangle, or a circular or elliptic arc, you can determine the corresponding parameters, e.g., the coordinates of the end points of a line or the center and




Figure 8.2: Selecting the desired contours: (a) exactly fitting ROI, (b) too many contours because of too large ROI, (c) result of post-processing the contours from (b).

radius of a circle, by calling one of the fitting operators. Their goal is to approximate the input contour as closely as possible to a line, rectangle, or a circular or elliptic arc. Because the used minimization algorithms are very advanced and all contour points are used for the process, the parameters can be calculated very reliably.

8.1.4 Extract Features

From both raw contours and processed contour parts features can be determined. Some of these consider the contour as a linear object. Others treat a contour as the outer boundary of an object. Obviously, the center of gravity makes sense only for a closed object, whereas the curvature is a feature of a linear object.

8.1.5 A First Example

The following program is an example for the basic concept of contour processing. It shows how short segments returned by the line extractor can be grouped to longer ones.

First, an image is acquired from file using `read_image`. The task is to extract the roads, which show up as thin bright lines in the image. For this the operator `lines_gauss` is used. When we look at the result of the contour extraction in [figure 8.3a](#), we see that a lot of unwanted small segments are extracted. They can be suppressed easily by calling `select_contours_xld` with a minimum contour length. A further problem is that some roads are split into more than one segment. They can be combined with the operator `union_collinear_contours_xld`. Looking at the result in [figure 8.3b](#), we see that many fragments

have been combined along straight road parts. In curves this method fails because the orientation of the segments differs too much.

```
read_image (Image, 'mreut4_3')
lines_gauss (Image, Lines, 1.5, 2, 8, 'light', 'true', 'bar-shaped', 'true')
select_contours_xld (Lines, LongContours, 'contour_length', 15, 1000, 0, 0)
union_collinear_contours_xld (LongContours, UnionContours, 30, 2, 9, 0.7, \
                             'attr_keep')
```




Figure 8.3: Processing XLD contours, (a) extracted contours, (b) processed contours.

8.2 Extended Concept

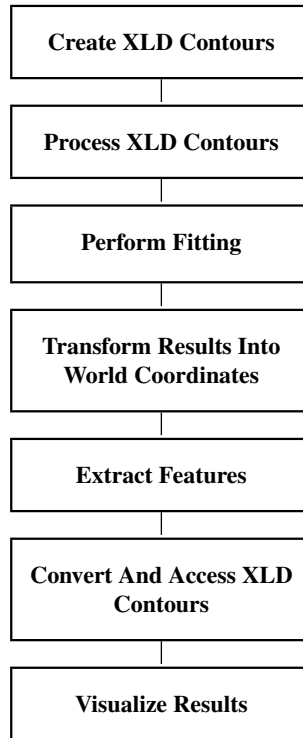
In addition to the standard contour processing, HALCON provides other tools that can be added. Typical examples for these are camera calibration, geometric transformations, or type conversions. With these, the contour methods can be integrated into the overall vision task.

8.2.1 Create XLD Contours

The standard method to create contours is to call a contour extraction operator. Contour extraction for edges is performed with `edges_sub_pix`, `edges_color_sub_pix`, or `zero_crossing_sub_pix`. Lines are extracted using `lines_gauss`, `lines_facet`, or `lines_color`. For subpixel blob analysis the operator `threshold_sub_pix` can be used. These operators are described in more detail with the method [Edge Extraction \(Subpixel-Precise\)](#) on page 73.

If pixel-accuracy is sufficient, you can use an edge filter (like `sobel_amp` or `edges_image`) or a line filter (like `bandpass_image`) followed by thresholding and thinning. The resulting elongated regions are then converted into XLD contours with the operator `gen_contours_skeleton_xld`. For more information on this approach see the method [Edge Extraction \(Pixel-Precise\)](#) on page 63.

Contours can also be synthesized from different sources, e.g., CAD data, user interaction, or measuring. Having obtained the coordinates of the control points from such a source, the operators



`gen_contour_polygon_xld` and `gen_contour_polygon_rounded_xld` convert them to XLD contours. You can also draw XLD contours interactively with the operators `draw_xld` and `draw_xld_mod`.

Finally, the border of regions can be converted into XLD contours. The corresponding operator is called `gen_contour_region_xld`.

8.2.2 Process XLD Contours

The first method to segment contours is to call `segment_contours_xld`. This operator offers various modes: Splitting into line segments, linear and circular segments, or linear and elliptic segments. The individual contour segments can then be selected with `select_obj` and passed to one of the fitting operators described with the step [Perform Fitting](#) on page 84. Whether a contour segment represents a line, a circular, or an elliptic arc can be queried via the global contour attribute 'cont_approx' using the operator `get_contour_global_attrib_xld`.

If only line segments are needed, you can use the combination of `gen_polygons_xld` followed by `split_contours_xld`. The behavior is similar to using `segment_contours_xld`. The main difference is the possible postprocessing: When applying `gen_polygons_xld`, a so-called XLD polygon is generated. This is a different data type, which represents the initial step for grouping of segments to parallel lines.

An important step during contour processing is the suppression of irrelevant contours. This can be accomplished with the operator `select_shape_xld`, which provides almost 30 different shape features.

By specifying the desired minimum and maximum value and possibly combining multiple features, contours can be selected very flexibly. As an alternative, you can use the operator [select_contours_xld](#), which offers typical features of linear structures. Finally, the operator [select_xld_point](#) can be used in combination with mouse functions to interactively select contours.

If there are gaps within a contour, the pieces are treated as separate objects, which makes further processing and feature extraction difficult. You can merge linear segments with the operators [union_collinear_contours_xld](#) or [union_straight_contours_xld](#). Additionally, you can also merge adjacent contours ([union_adjacent_contours_xld](#)), contours that lie on the same circle ([union_cocircular_contours_xld](#)), or contours that are cotangential ([union_cotangential_contours_xld](#)). To handle contours with a complex shape you can first segment them into linear, circular, or elliptic segments (see above).

HALCON also provides an operator for general shape modifications: [shape_trans_xld](#). With this operator you can, e.g., transform the contour into its surrounding circle, convex hull, or surrounding rectangle. Further, for closed contours or polygons, set theoretical operations can be applied to combine contours. For example, with [intersection_closed_contours_xld](#) you can intersect the regions that are enclosed by the closed contours, with [difference_closed_contours_xld](#) you can calculate the difference between the enclosed regions, or with [union2_closed_contours_xld](#) you can merge the enclosed regions.

8.2.3 Perform Fitting

With the operator [fit_line_contour_xld](#) you can determine the parameters of a line segment. The operator provides different optimization methods, most of which are suppressing outliers. It returns the coordinates of the start and the end point of the fitted line segment and the normal form of the line. To visualize the results, you can use the operator [gen_contour_polygon_xld](#).

To fit a rectangle into a contour, the operator [fit_rectangle2_contour_xld](#) can be used. It provides various optimization methods as well. The returned parameters comprise mainly the center position, the extent, and the orientation of the rectangle. To generate the obtained rectangle for a visualization, you can use the operator [gen_rectangle2_contour_xld](#).

For the fitting of circular and elliptic segments the operators [fit_circle_contour_xld](#) and [fit_ellipse_contour_xld](#) are available. They also provide various optimization methods. For a circular segment the center and the radius are returned together with the angle range of the visible part. In addition, a second radius and the orientation of the main axis are returned for elliptic segments. To visualize the results of both operators, you can use either the operator [gen_ellipse_contour_xld](#) or the operator [gen_circle_contour_xld](#).

8.2.4 Transform Results Into World Coordinates

As a post-processing step, it may be necessary to correct the contours, e.g., to remove lens distortions, or to transform the contours into a 3D world coordinate system in order to extract dimensional features in world units. Such a transformation is based on calibrating the camera. After the calibration, you simply call the operator [contour_to_world_plane_xld](#) to transform the contours.

How to transform contours into world coordinates is described in detail in the Solution Guide III-C in [section 3.3](#) on page 57.

8.2.5 Extract Features

HALCON offers various operators to access the feature values. Commonly used shape features are calculated by `area_center_xld`, `compactness_xld`, `convexity_xld`, `eccentricity_xld`, `diameter_xld`, and `orientation_xld`. The hulls of the contours can be determined with `smallest_circle_xld` or `smallest_rectangle2_xld`. Features based on geometric moments are calculated, e.g., by `moments_xld`.

8.2.6 Convert And Access XLD Contours

Finally, it might be necessary to access the raw data of the contours or to convert contours into another data type, e.g., into a region.

You can access the coordinates of the control points with the operator `get_contour_xld`. It returns the row and column coordinates of all control points of a contour in two tuples of floating-point values. In case of a contour array (tuple), you must loop over all the contours and select each one using `select_obj`.

To convert contours to regions, simply call the operator `gen_region_contour_xld`. The operator `paint_xld` paints the contour with anti-aliasing into an image.

The operators for edge and line extraction not only return the XLD contours but also so-called attributes. Attributes are numerical values; they are associated either with each control point (called contour attribute) or with each contour as a whole (global contour attribute). The operators `get_contour_attrib_xld` and `get_contour_global_attrib_xld` enable you to access these values by specifying the attribute name. More information on this topic can be found in the description of the step [Determine Contour Attributes](#) on page 76.

8.2.7 Visualize Results

Finally, you might want to display the images and the contours.

For detailed information see the [description of this method](#) on page 289.

8.3 Programming Examples

This section gives a brief introduction to using HALCON for contour processing.

8.3.1 Measuring Lines and Arcs

Example: `solution_guide/basics/measure_metal_part.hdev`

The first example shows how to segment a contour into lines and (circular) arcs and how to determine the corresponding parameters. [Figure 8.4](#) shows the final result of the fitted primitives overlaid on the input image.




Figure 8.4: Fitted lines and circles.

As the initial step, the contours of the metal part are extracted using the operator `edges_sub_pix`. The resulting contours are segmented into lines and circular arcs and sorted according to the position of their upper left corner.

```

edges_sub_pix (Image, Edges, 'lanser2', 0.5, 40, 90)
segment_contours_xld (Edges, ContoursSplit, 'lines_circles', 6, 4, 4)
sort_contours_xld (ContoursSplit, SortedContours, 'upper_left', 'true', \
                  'column')

```

Then, lines and circles are fitted to the extracted segments. As already noted, the individual segments must be accessed inside a loop. For this, first their total number is determined with `count_obj`. Inside the loop, the individual segments are selected with the operator `select_obj`. Then, their type (line or circular arc) is determined by accessing a global attribute with `get_contour_global_attrib_xld`. Depending on the result, either a circle or a line is fitted. For display purposes, circles and lines are created using the determined parameters. Furthermore, the length of the lines is computed with the operator `distance_pp`.

```

count_obj (SortedContours, NumSegments)
for i := 1 to NumSegments by 1
  select_obj (SortedContours, SingleSegment, i)
  get_contour_global_attrib_xld (SingleSegment, 'cont_approx', Attrib)
  if (Attrib = 1)
    fit_circle_contour_xld (SingleSegment, 'atukey', -1, 2, 0, 5, 2, \
                          Row, Column, Radius, StartPhi, EndPhi, \
                          PointOrder)
    gen_ellipse_contour_xld (ContEllipse, Row, Column, 0, Radius, \
                          Radius, 0, rad(360), 'positive', 1.0)
  else
    fit_line_contour_xld (SingleSegment, 'tukey', -1, 0, 5, 2, RowBegin, \
                      ColBegin, RowEnd, ColEnd, Nr, Nc, Dist)
    gen_contour_polygon_xld (Line, [RowBegin, RowEnd], [ColBegin, ColEnd])
    distance_pp (RowBegin, ColBegin, RowEnd, ColEnd, Length)
  endif
endfor

```

8.3.2 Close gaps in a contour

Example: solution_guide/basics/close_contour_gaps.hdev

The second example demonstrates how to close gaps in an object contour (see [figure 8.5](#)). The example is based on synthetic data. Instead of using a real image, a light gray square on a dark gray background is generated and a part of its boundary is blurred.

```

gen_rectangle1 (Rectangle, 30, 20, 100, 100)
region_to_bin (Rectangle, BinImage, 130, 100, 120, 130)
rectangle1_domain (BinImage, ImageReduced, 20, 48, 40, 52)
mean_image (ImageReduced, SmoothedImage, 15, 15)
paint_gray (SmoothedImage, BinImage, Image)

```




Figure 8.5: Original edges and result of grouping process.

The extraction of contours with `edges_sub_pix` thus results in an interrupted boundary (see [figure 8.5a](#)). Note that the edge extraction is restricted to the inner part of the image, otherwise edges would be extracted at the boundary of the image.

```
rectangle1_domain (BinImage, ImageReduced, 20, 48, 40, 52)
edges_sub_pix (ImageReduced, Edges, 'lanser2', 1.1, 22, 30)
```

A suitable operator for closing gaps in linear segments is `union_collinear_contours_xld`. Before we can apply this operator, some pre-processing is necessary: First, the contours are split into linear segments using `segment_contours_xld`. Then, `regress_contours_xld` is called to determine the regression parameters for each segment. These parameters are stored with each contour and could be accessed with `get_regress_params_xld`. Finally, `union_collinear_contours_xld` is called. Its result is depicted in [figure 8.5b](#) on page 91.

```
segment_contours_xld (Edges, LineSegments, 'lines', 5, 4, 2)
regress_contours_xld (LineSegments, RegressContours, 'no', 1)
union_collinear_contours_xld (RegressContours, UnionContours, 10, 1, 2, 0.1, \
                             'attr_keep')
```

8.3.3 Extract Roads

Example: `hdevelop/Applications/Object-Recognition-2D/roads.hdev`

The task of this example is to extract the roads in the aerial image depicted in [figure 8.6](#)




Figure 8.6: (a) Original image; (b) zoomed image part; (c) extracted roads.

The programs starts by first extracting lines on a reduced scale. These lines correspond very well to roads.

```
threshold (Mreut43, Bright, 160, 255)
reduce_domain (Mreut43, Bright, Mreut43Bright)
lines_gauss (Mreut43Bright, RoadCenters, 1.2, 5, 14, 'light', 'true', \
            'bar-shaped', 'true')
```

To eliminate wrong candidates, edges are extracted on a higher scale. For the road extraction it is assumed that a road consists of two parallel edges with homogeneous gray values and a line segment in between. Using some contour processing operators, this model is refined step by step.

```
edges_image (Part, PartAmp, PartDir, 'mderiche2', 0.3, 'nms', 20, 40)
threshold (PartAmp, EdgeRegion, 1, 255)
clip_region (EdgeRegion, ClippedEdges, 2, 2, PartWidth - 3, PartHeight - 3)
skeleton (ClippedEdges, EdgeSkeleton)
gen_contours_skeleton_xld (EdgeSkeleton, RoadEdges, 1, 'filter')
gen_polygons_xld (RoadEdges, RoadEdgePolygons, 'ramer', 2)
gen_parallel_xld (RoadEdgePolygons, ParallelRoadEdges, 10, 30, 0.15, \
                'true')
mod_parallel_xld (ParallelRoadEdges, Part, ModParallelRoadEdges, \
                ExtParallelRoadEdges, 0.3, 160, 220, 10)
combine_roads_xld (RoadEdgePolygons, ModParallelRoadEdges, \
                ExtParallelRoadEdges, RoadCenterPolygons, RoadSides, \
                rad(40), rad(20), 40, 40)
```

8.4 Selecting Operators

8.4.1 Create XLD Contours

Standard:

```
gen_contour_polygon_xld, gen_contour_region_xld, gen_ellipse_contour_xld,
gen_circle_contour_xld, gen_rectangle2_contour_xld, draw_xld, draw_xld_mod
```

Advanced:

```
gen_contour_polygon_rounded_xld
```

Further operators can be found in the following places: [detailed operator list for the step Extract Edges Or Lines](#) on page 80, [operator list for the method Edge Extraction \(Pixel-Precise\)](#) (see section 6.4 on page 69).

8.4.2 Process XLD Contours

Standard:

```
segment_contours_xld, gen_polygons_xld, split_contours_xld, select_shape_xld,  
select_contours_xld, select_xld_point, union_collinear_contours_xld,  
union_straight_contours_xld, union_adjacent_contours_xld,  
union_cocircular_contours_xld, union_cotangential_contours_xld, shape_trans_xld,  
intersection_closed_contours_xld, intersection_closed_polygons_xld,  
difference_closed_contours_xld, difference_closed_polygons_xld,  
union2_closed_contours_xld, union2_closed_polygons_xld
```

Advanced:

```
union_collinear_contours_ext_xld, symm_difference_closed_contours_xld,  
symm_difference_closed_polygons_xld
```

8.4.3 Perform Fitting

Standard:

```
fit_line_contour_xld, fit_circle_contour_xld, fit_ellipse_contour_xld,  
fit_rectangle2_contour_xld
```

8.4.4 Transform Results Into World Coordinates

Standard:

```
contour_to_world_plane_xld, change_radial_distortion_contours_xld
```

More operators for transforming results into world coordinates are described in the [Solution Guide III-C](#).

8.4.5 Extract Features

Standard:

```
area_center_xld, orientation_xld, smallest_circle_xld, smallest_rectangle1_xld,  
smallest_rectangle2_xld, compactness_xld, convexity_xld, diameter_xld,  
eccentricity_xld
```

A full list of operators can be found in the Reference Manual in the chapter “[XLD ▸ Features](#)”.

8.4.6 Convert And Access XLD Contours

Standard:

```
get_contour_xld, gen_region_contour_xld
```

Advanced:

`paint_xld`

More information on operators for accessing XLD data can be found in the Reference Manual in the chapter “XLD ▷ Access”. Operators for determining contour attributes can be found in the [detailed operator list for the step Determine Contour Attributes](#) on page 80.

8.4.7 Visualize Results

Please refer to the [operator list for the method Visualization](#) (see section 19.4 on page 297).

8.5 Relation to Other Methods

8.5.1 Alternatives to Contour Processing

Line Processing

A very basic alternative to contour processing are the operators for line processing. In this context, lines are treated as tuples of start and end points. The extraction can, e.g., be performed with [detect_edge_segments](#). Of course, XLD polygons can also be converted into this type of lines. Operators for processing this type of lines can be found in the Reference Manual in the chapter “Tools ▷ Lines”.

8.6 Advanced Topics

8.6.1 Line Scan Cameras

In general, line scan cameras are treated like normal area sensors. But in some cases, not single images but an infinite sequence of images showing objects, e.g., on a conveyor belt, have to be processed. In this case the end of one image is the beginning of the next one. This means that contours that partially lie in both images must be combined into one contour. For this purpose HALCON provides the operator [merge_cont_line_scan_xld](#). This operator is called after the processing of one image and combines the current contours with those of previous images. For more information see [Solution Guide II-A](#).

Chapter 9

Matching

The idea of matching is quite simple: In a training image a so-called template is presented. The system derives a model from this template. This model is then used to locate objects that look “similar” to the template in search images. Depending on the selected method, this approach is able to handle changes in illumination, clutter, varying size, position, and rotation, or even relative movement of parts of the template.

The advantage of matching is its ease of use combined with great robustness and flexibility. Matching does not require any kind of segmentation of the desired objects. By some of the matching methods, objects can be located even if they are overlapped by other objects. Furthermore, matching has only a few parameters. Even these can be determined automatically in most cases. This makes this method especially attractive for applications where the end user only has little skills in machine vision.

HALCON offers different methods for matching. The selection depends on the image data and the task to be solved. How to use the individual matching approaches is described in detail in the [Solution Guide II-B](#).

[Figure 9.1](#) summarizes the characteristics of the different matching approaches and helps you to select the appropriate approach for your task.

- The *gray-value-based matching* is the classical method. It can be used if the gray values inside the object do not vary and if there are no missing parts and no clutter. The method can handle single instances of objects, which can appear rotated in the search image. Note that this matching approach should be applied only in the very rare case that an illumination-variant matching is needed.
- The *correlation-based matching* is based on gray values and a normalized cross correlation. In contrast to the classical gray-value-based matching, linear illumination changes can be coped with. The operators are used similar to the operators of shape-based matching. The advantage over shape-based matching is that also texture can be handled. On the other hand, the object may only be rotated but not scaled, and the approach is limited to gray-value images. Additionally, the approach is sensitive to occlusion, clutter, and non-linear illumination changes.
- The *shape-based matching* represents the state of the art in machine vision. Instead of using the gray values, features along contours are extracted and used both for the model generation and the

Search for 2D models (orthogonal view)	Gray value-based Matching	Does not work with clutter, occlusion, illumination changes, scale, or multi-channel images.
	Correlation-based Matching	Invariant to defocus, slight shape deformations, and linear illumination changes. Works good for textured objects. BUT does not work with clutter, occlusion, non-linear illumination changes, scale, or multi-channel images.
	Shape-based Matching	Invariant to clutter, occlusion, non-linear illumination changes, scale, defocus, and slight shape deformations. Works with multi-channel images. Can be applied for multiple models simultaneously. BUT is difficult with some textures.
	Component-based Matching	Parts of objects move relative to each other. Invariant to clutter, occlusion, non-linear illumination changes, and scale. Works with multi-channel images. Can be applied for multiple models simultaneously. BUT is difficult with some textures and does not work with defocus and shape deformations.
	Local deformable Matching	Returns also the deformations of the model instance. Invariant to clutter, occlusion, non-linear illumination changes, scale, and local deformations. Works with multi-channel images.
Search for 2D models (orthogonal or perspective view)	Perspective deformable Matching	Invariant to clutter, occlusion, non-linear illumination changes, scale, defocus, and perspective shape deformations. Works with multi-channel images. BUT is difficult with some textures.
	Descriptor-based Matching	Invariant to clutter, occlusion, non-linear illumination changes, scale, and perspective shape deformations. BUT does not work without texture, in particular distinctive points, and does not work with defocus or multi-channel images.
Search for 3D models	3D Matching	see Solution Guide I, chapter 10 on page 127
Search for corresponding points to combine overlapping images (uncalibrated mosaicking)	Point-based Matching	see Solution Guide III-C, chapter 10 on page 199

Figure 9.1: Guide to the different matching approaches.


matching. This has the effect that this method is invariant to changes in illumination and variations of the object's gray values. It can handle missing object parts, clutter, noise, defocused, and slightly

deformed models. Furthermore, multiple instances can be found and multiple models can be used at the same time. The method allows the objects to be rotated and scaled and can be applied also for multi-channel images.

- The *component-based matching* can be considered as a high-level shape-based matching: The enhancement is that an object can consist of multiple parts that can move (rotate and translate) relative to each other. A simple example of this is a pair of pliers. Logically this is considered as one object, but physically it consists of two parts. The component-based matching allows handling such a compound object in one search step. The advantage is an improved execution time and increased robustness compared to handling the parts as distinct models. But in contrast to the shape-based matching, it can not handle defocused images or slightly deformed models.
- The *local deformable matching* is similar to the shape-based matching, but here significant deformations can be handled and returned. In particular, besides the position and score, the matching can return a rectified version of the significant part of the search image, a vector field that describes the deformations, and the deformed contours of the found model instance.
- The *perspective deformable matching* is also similar to the shape-based matching, but here also strong perspective deformations can be handled and instead of a 2D pose a 2D projective transformation matrix (homography) is returned. In addition, a calibrated version of the perspective deformable matching is available. There, instead of a 2D projective transformation matrix (homography) the 3D pose of the object is returned. Here, the focus is on the uncalibrated case. The calibrated case is described in more detail in the Solution Guide II-B in [section 3.6](#) on page 124.
- The *descriptor-based matching* has the same intention as the perspective deformable matching, i.e., the 2D projective transformation matrix (homography) can be obtained for the uncalibrated case and the 3D pose can be obtained for the calibrated case. The main difference is that points instead of contours are used to create and find the model. Thus, it is especially suitable for highly textured objects but is not suitable for low textured objects with rounded edges. Compared to the perspective deformable matching, it is significantly faster for large search spaces but less accurate. Here, the focus is on the uncalibrated case. The calibrated case is described in more detail in the Solution Guide II-B in [section 3.7](#) on page 136.
- The *3D matching* consists of different methods that use 3D data as template for the matching. They are introduced in more detail in [3D Matching](#) on page 127. Note that if you search for planar objects or object parts in the 3D space, the perspective deformable matching or the descriptor-based matching are alternatives that are faster and more convenient to use.
- The *point-based matching* has the intention to combine two overlapping images. This is done by first extracting significant points in both images. These points are the input for the actual matching process. The result of the matching is a mapping from one image to the other, allowing translation, rotation, scaling, and perspective distortions. This mapping is typically used to combine the two images into a single, larger one. Of course, one image can also be treated as a template and the other image as showing an instance of the object that should be found. The point-based matching is able to handle perspective distortions without calibration, but needs an increased execution time, which comes mainly from the extraction of the significant points.

9.1 Basic Concept

Matching is divided into the following parts:



9.1.1 Acquire Image(s)

Both for training and matching, first an image is acquired.

For detailed information see the [description of this method](#) on page 13.

9.1.2 Create (Train) Model

To create a matching model, first a region of interest that covers the template in the training image must be specified. Only those parts of the image that are really significant and stable should be used for training. The input for the training operator is the reduced image together with control parameters. The handle of the model is the output of the training. The model will then be used for immediate search or stored to file.

9.1.3 Find Model

Having created (or loaded) a model, it can now be used for locating objects in the image. Each method offers specific methods to perform this task. If one or multiple objects are found, their poses (position, rotation, and scaling) or 2D projective transformation matrices (homographies) together with a score are returned. These values can already be the desired result or serve as input for the next step of the vision process, e.g., for aligning regions of interest.

9.1.4 Destroy Model

When you no longer need the matching model you should destroy it. For example, for shape-based matching you destroy it using [clear_shape_model](#).

9.1.5 A First Example

An example for this basic concept is the following program, which shows all necessary steps from model generation to object finding using shape-based matching.




Figure 9.2: Finding all clips in the image.

A training image is acquired from file. A region is generated as region of interest, covering one of the clips in the image. After combining the region with the image, it is used as input for the training operator `create_shape_model`. To keep the example simple, the same image is used to test the matching by searching for all clips with `find_shape_model`. For visualization purposes the model contours are accessed, moved to the corresponding positions, and overlaid on the image. Finally, the model is cleared to release the memory.

```
read_image (Image, 'clip')
gen_rectangle2 (ROI, 124, 181, 0.653, 129, 47)
reduce_domain (Image, ROI, ImageReduced)
create_shape_model (ImageReduced, 0, 0, rad(360), 0, 'no_pregeneration', \
    'use_polarity', 40, 10, ModelID)
find_shape_model (Image, ModelID, 0, rad(360), 0.7, 13, 0.5, \
    'interpolation', 0, 0.9, Row, Column, Angle, Score)
get_shape_model_contours (ModelContours, ModelID, 1)
for i := 0 to |Row|-1 by 1
    vector_angle_to_rigid (0, 0, 0, Row[i], Column[i], Angle[i], HomMat2D)
    affine_trans_contour_xld (ModelContours, ContoursAffinTrans, HomMat2D)
endfor
clear_shape_model (ModelID)
```

9.2 Extended Concept

In many cases, matching will be more complex than in the example above. Reasons for this are, e.g., advanced training using a synthetic template, searching with multiple models at the same time, or using the matching results as input for further processing like alignment.

Figure 9.3 shows the major steps. The first part is offline and consists of the training of the model, using different input sources that depend on the application. The model can then be stored to be loaded again for later use. For the matching itself, one or more models are used to search for objects in images. The results are the poses, i.e., the position, rotation, and scale of the found objects, or the 2D projective transformation matrices (homographies), which describe the relation between the 2D world coordinates and the 2D image coordinates.

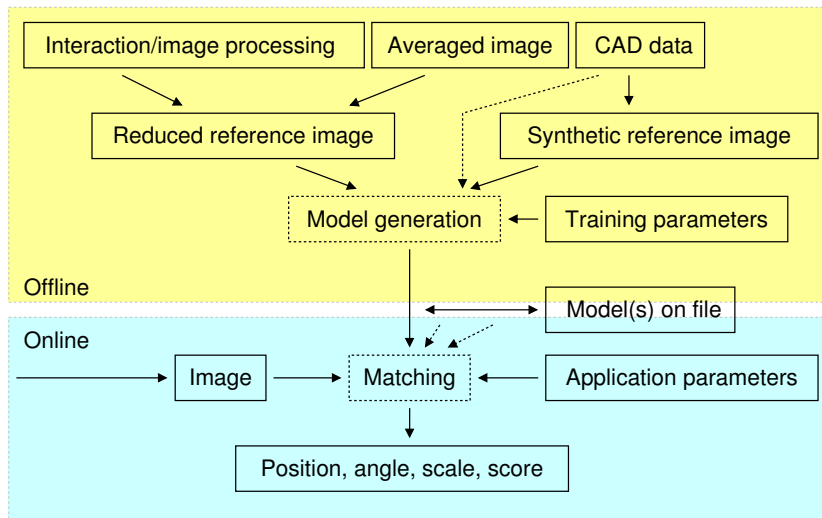


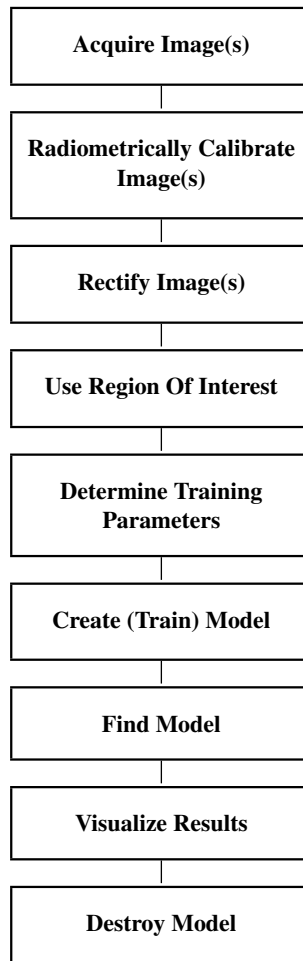
Figure 9.3: Overview of the matching process.

9.2.1 Radiometrically Calibrate Image(s)

To allow high-accuracy matching, the camera should have a linear response function, i.e., the gray values in the images should depend linearly on the incoming energy. Since some cameras do not have a linear response function, HALCON provides the so-called radiometric calibration (gray value calibration): With the operator `radiometric_self_calibration` you can determine the inverse response function of the camera (offline) and then apply this function to the images using `lut_trans` before performing the matching.

9.2.2 Rectify Image(s)

As a preprocessing step, it may be necessary to rectify the image, e.g., to remove lens distortions or to transform the image into a reference point of view. This will allow using matching even if the camera is



looking from the side onto the object plane or the surface is a cylinder.

Detailed information about rectifying images can be found in the Solution Guide III-C in [section 3.4](#) on page 62.

9.2.3 Use Region Of Interest

Regions of interest affect both the model creation and the search. Especially during the creation it is very effective to make use of arbitrarily shaped ROIs and thus mask parts that do not belong to the template. During the search, ROIs can be used to focus the process only on relevant parts, which reduces the execution time. Two important notes on ROIs, which often cause confusion:

- The reference point of the template is defined by the center of gravity of the ROI used *during the creation*, not by the center of gravity of the contours returned for the contour-based matching approaches by [inspect_shape_model](#).

- During the search process, only the reference point of the model must fit into the search ROI, not the complete shape. Therefore, sometimes very small ROIs consisting only of a few pixels can be used.

For an overview on how to construct regions of interest and how to combine them with the image see [Region Of Interest](#) on page 19.

For detailed information see the [description of this method](#) on page 19.

9.2.4 Determine Training Parameters

To make the use of the matching as easy as possible it is necessary to determine the training parameters automatically. Because the shape-based matching is used most frequently, this section focuses on this method. In addition, the component-based matching and the local and perspective deformable matching use very similar parameters.

The operators `create_scaled_shape_model`, `create_shape_model`, and `create_aniso_shape_model` allow to use the value 'auto' for many parameters. This has the effect that the values are determined automatically. The allowed ranges of orientations and scales are typically known from the setup of the system. They can also be determined easily in HDevelop using the assistant provided via the menu item Assistants > Matching. If 'no_pregeneration' is used for the parameter Optimization wide ranges can be used and restricted during the finding process. The other training parameters (Contrast, MinContrast) can be determined also by using the assistant. If an automatic method is needed, please contact your local distributor. Their support can provide application-specific procedures to determine the parameters.

9.2.5 Create (Train) Model

Creating a model differs from method to method, but several rules hold for most of the methods.

- You can use both real and synthetic images for the training. Using real images is easier. However, if no such image is available or if the real images are too noisy you can create a synthetic image, e.g., with the operators `gen_contour_polygon_rounded_xld`, `gen_image_const`, and `paint_xld`. If the training image is too noisy, you can enhance it, e.g., by smoothing noisy areas or by averaging multiple templates into a single one.
- To speed up the matching process, a so-called image pyramid is created, consisting of the original, full-sized image and a set of downsampled images. The model is then created and searched on the different pyramid levels. You can view the image pyramid and the corresponding models for shape-based matching with the matching assistant in HDevelop.
- Each model (except for the point-based matching) is manipulated via a handle. This handle points to the real data, which can consume a large amount of memory. Therefore, it is important to clear the model when it is no longer needed.

Having the general rules in mind, we now show how the creation is performed for the different methods.

- To create a model for *gray-value-based matching*, typically the operator `create_template_rot` is used. The model can then be stored to file with `write_template`.

- To create a model for the *correlation-based matching*, the operator `create_ncc_model` is used. The model can then be stored to file with `write_ncc_model`.
- To create a model for the *shape-based matching*, typically the operator `create_scaled_shape_model` (or `create_aniso_shape_model` for anisotropic scaling) is used. If no scaling is needed, you can use `create_shape_model`. The model can then be stored to file with `write_shape_model`.

In contrast to some other matching methods, you can also directly use XLD contours as shape models. Then, you do not have to generate a synthetical image, but simply create the model with `create_shape_model_xld`, `create_scaled_shape_model_xld`, or `create_aniso_shape_model_xld` instead of `create_shape_model`, `create_scaled_shape_model`, or `create_aniso_shape_model`. If such a contour is stored as a DXF file, you first have to derive the corresponding HALCON XLD contour from the DXF file with `read_contour_xld_dxf`. Note that when using XLD Contours as model, you should specify the polarity for the search images after a first successful match using `set_shape_model_metric`.

- For the *component-based matching*, the model can be created in different ways: If the components and their relations are known, you can use the operator `create_component_model`. If the components are known, but not their relations, you can train the relations using the operator `train_model_components` with a set of training images, and then create the model using `create_trained_component_model`. Finally, if the components themselves are not known, you can determine them with the operator `gen_initial_components`, which then serve as the input for `train_model_components`.

The model can then be stored to file with `write_component_model`. The training components can be saved with `write_training_components`.

- To create a model for the *local deformable matching*, the operator `create_local_deformable_model` is used. The model can then be stored to file with `write_deformable_model`.

Similar to shape-based matching, you can also directly use XLD contours as local deformable models. Then, you do not have to generate a synthetical image, but simply create the model with `create_local_deformable_model_xld` instead of `create_local_deformable_model`. If such a contour is stored as a DXF file, you first have to derive the corresponding HALCON XLD contour from the DXF file with `read_contour_xld_dxf`. Note that when using XLD Contours as model, you should specify the polarity for the search images after a first successful match using `set_local_deformable_model_metric`.

- To create a model for the uncalibrated *perspective deformable matching*, the operator `create_planar_uncalib_deformable_model` is used. To create a model for the calibrated perspective deformable matching, the operator `create_planar_calib_deformable_model` is used. For both, the model can be stored to file with `write_deformable_model`.

Similar to shape-based matching, you can also directly use XLD contours as perspective deformable models. Then, you do not have to generate a synthetical image, but simply create the model with `create_planar_uncalib_deformable_model_xld` or `create_planar_calib_deformable_model_xld` instead of `create_planar_uncalib_deformable_model` or `create_planar_calib_deformable_model`. If such a contour is stored as a DXF file, you first have to derive the corresponding HAL-

CON XLD contour from the DXF file with `read_contour_xld_dxf`. Note that when using XLD Contours as model, you should specify the polarity for the search images after a first successful match using `set_planar_uncalib_deformable_model_metric` or `set_planar_calib_deformable_model_metric`.

- To create a model for the uncalibrated *descriptor-based matching*, the operator `create_uncalib_descriptor_model` is used. To create a model for the calibrated descriptor-based matching, the operator `create_calib_descriptor_model` is used. For both, the model can be stored to file with `write_descriptor_model`.
- For the training of the *point-based matching*, no real model is generated. Here, only coordinates must be determined, which will later be matched: For both the template and the search image, subpixel locations of significant points are extracted. This is done using an operator like `points_foerstner`. The returned positions, together with the image data, will then be input for the matching.

9.2.6 Find Model

Finding an object with matching is an online process that has the image and one or multiple models as input. Depending on the selected method, the search operators offer different ways of controlling the search.

- You can search objects with *gray-value-based matching* using the operator `best_match_rot`. This operator finds single instances of rotated objects. If the objects are not rotated, the operator `fast_match_mg` allows to find all points with a high score. These locations are returned as a region and can be further processed, e.g., with blob analysis (see the [description of this method](#) on page 31).
- To find objects using *correlation-based matching*, the operator `find_ncc_model` is applied. It returns the position, rotation and score of the found instances of a model.
- The standard operator for *shape-based matching* is `find_scaled_shape_model`. It allows locating multiple instances of the same template with rotations and scaling. If multiple models are needed, the operator `find_scaled_shape_models` is used. For both operators, there are variants for anisotropic scaling and without scaling.
- For the *component-based matching*, `find_component_model` is used to locate multiple instances of objects, which may be rotated and whose components may move with respect to each other.
- For the *local deformable matching*, instances of a model are found with `find_local_deformable_model`. It returns the position and the score of the found instances. Optionally, also a rectified image, a vector field that describes the deformations of the found model instance, and the deformed contours of the found model instance can be returned.
- For the uncalibrated *perspective deformable matching*, instances of a model are found with `find_planar_uncalib_deformable_model`. It returns the 2D projective transformation matrix (homography) and the score of the found instances. For the calibrated perspective deformable matching, instances of a model are found with `find_planar_calib_deformable_model`. It returns the 3D pose and the score of the found instances.

- For the uncalibrated *descriptor-based matching*, instances of a model are found with `find_uncalib_descriptor_model`. It returns the 2D projective transformation matrix (homography) and the score of the found instances. For the calibrated descriptor-based matching, instances of a model are found with `find_calib_descriptor_model`. It returns the 3D pose and the score of the found instances.
- The matching operator for the *point-based matching* is `proj_match_points_ransac`. The input are the two images containing the template *and* the object to be located, together with the significant points. The result of the matching process is a 2D mapping from the template image to the object image. This mapping is the input for operators like `gen_projective_mosaic`, which can be used to combine the two images.

9.2.7 Visualize Results

A typical visualization task is to display the model contours overlaid on the image at the found position. An easy way to do this is to access the model contours by calling `get_shape_model_contours` (for the shape-based matching). With `vector_angle_to_rigid` and `disp_xld`, the contours can be displayed. Displaying the results of the component-based matching is more complex. The easiest way is to call `get_found_component_model`, which returns regions, one for each component. These can directly be displayed with `disp_region`. For the local or perspective deformable matching the model contours can be accessed using `get_deformable_model_contours`. For local deformable matching additionally the deformed contours of the found model instance can be displayed. They can be obtained automatically during the matching with `find_local_deformable_model`. The points that define the model for a descriptor-based matching can be queried with `get_descriptor_model_points`.

For detailed information see the [description of this method](#) on page 289.

9.3 Programming Examples

This section gives a brief introduction to using HALCON for template matching.

9.3.1 Creating a Model for the “Green Dot”

Example: `solution_guide/basics/create_model_green_dot.hdev`

This example shows how to use the shape-based matching with objects of varying size. [Figure 9.4](#) depicts the training image, which contains the so-called “green dot”, a symbol used in Germany for recycling packages. The template is not defined by a user interaction but by a segmentation step: Using `threshold` all dark pixels are selected and the connected component with the appropriate size is chosen (`select_shape`). This region is then filled and slightly dilated. [Figure 9.4a](#) depicts the result of the segmentation.

```

threshold (Image, Region, 0, 128)
connection (Region, ConnectedRegions)
select_shape (ConnectedRegions, SelectedRegions, 'area', 'and', 10000, \
                20000)
fill_up (SelectedRegions, RegionFillUp)
dilation_circle (RegionFillUp, RegionDilation, 5.5)

```



Figure 9.4: Creating a model for matching: (a) ROI for the template region; (b) model contours.

The extracted region is then combined with the image ([reduce_domain](#)) to be used as the domain, i.e., as the region of interest. To check whether the value for the parameter `Contrast` has been chosen correctly, [inspect_shape_model](#) is called. The training is finally applied using [create_scaled_shape_model](#). The resulting model contours are depicted in [figure 9.4b](#).

```

reduce_domain (Image, RegionDilation, ImageReduced)
inspect_shape_model (ImageReduced, ModelImages, ModelRegions, 1, 40)
create_scaled_shape_model (ImageReduced, 5, rad(-45), rad(90), 0, 0.8, 1.0, \
                            0, ['none','no_pregeneration'], \
                            'ignore_global_polarity', 40, 10, ModelID)

```

To complete the program, the model is written to file ([write_shape_model](#)); then, the memory of the model is released with [clear_shape_model](#).

```

write_shape_model (ModelID, 'green-dot.shm')
clear_shape_model (ModelID)

```

9.3.2 Locating “Green Dots”

Example: `solution_guide/basics/matching_green_dot.hdev`

In this example, we use the model created in the previous example to locate the so-called “green dots” in a search image. As you can see in [figure 9.5a](#), the search image contains three “green dots” in different orientations and scales. Furthermore, some of them are partly occluded.



Figure 9.5: Using matching to locate rotated and scaled objects: (a) search image; (b) matches.

First, the shape model is read from file.

```
read_shape_model ('green-dot.shm', ModelID)
```

Then, the search operator `find_scaled_shape_model` is executed. The result of the operator are the positions, the orientations, and the scales of the “green dots”. To display the model contours overlaid on the image, the contours are accessed with `get_shape_model_contours`. The for-loop is used to handle the contours for each found location. This is done by generating an appropriate transformation and then moving the contours to the correct position. [Figure 9.5b](#) depicts the result.

```
find_scaled_shape_model (ImageSearch, ModelID, rad(-45), rad(90), 0.8, 1.0, \
                        0.5, 0, 0.5, 'least_squares', 5, 0.8, Row, Column, \
                        Angle, Scale, Score)
get_shape_model_contours (ModelContours, ModelID, 1)
for I := 0 to |Score|-1 by 1
    vector_angle_to_rigid (0, 0, 0, Row[I], Column[I], Angle[I], \
                          HomMat2DRotate)
    hom_mat2d_scale (HomMat2DRotate, Scale[I], Scale[I], Row[I], Column[I], \
                    HomMat2DScale)
    affine_trans_contour_xld (ModelContours, ModelTrans, HomMat2DScale)
    dev_display (ModelTrans)
endfor
```

At the end of the program, the memory of the model is released with `clear_shape_model`.

```
clear_shape_model (ModelID)
```

9.3.3 Distinguishing coins

Example: `solution_guide/basics/matching_coins.hdev`

The task of this example is to distinguish different types of Euro coins, depending on the country of origin. The coins differ on one side, having a specific symbol for each country (see [figure 9.6](#)). Like the examples described before, the task is solved by shape-based matching. For each country symbol, one model is used.




Figure 9.6: Multiple coins that are distinguished.

The program consists of multiple procedures for the different tasks. The main program simply defines the working environment and calls these procedures. The first part is the creation of four models for the four different types of 20 cent coins. Inside a `for`-loop, the training procedure `train_model` is called with the corresponding training images, and the model IDs are collected in a tuple.

```
Names := ['german','italian','greek','spanish']
Models := []
for i := 0 to 3 by 1
  read_image (Image, 'coins/20cent_'+Names[i])
  dev_display (Image)
  train_model (Image, ModelID)
  Models := [Models,ModelID]
endfor
```

In the second part of the main program, the created models are used to recognize the origin of a coin. After applying the matching with the procedure `find_coin`, the result is visualized with the procedure `display_model`.

```
for i := 1 to 13 by 1
  read_image (Image, 'coins/20cent_'+i$.2+'.png')
  find_coin (Image, Models, Row, Column, Angle, Score, Model)
  display_model (Image, Model, Row, Column, Angle, Names, WindowHandle)
endfor
```

For the training, ROIs are generated automatically using the result of the procedure `locate_coin`. This procedure applies `threshold` to extract all bright pixels in the image. From the connected components, the largest one is selected with `select_shape_std`. The selected region contains several holes, which are filled with `shape_trans` using the parameter value 'convex', thereby transforming the region to its convex hull. Because only the inner part of the coins will be used as a template, a circle located in the center of gravity with a fixed radius is generated.

```
threshold (Image, Region, 70, 255)
connection (Region, ConnectedRegions)
select_shape_std (ConnectedRegions, SelectedRegions, 'max_area', 0)
shape_trans (SelectedRegions, RegionTrans, 'convex')
area_center (RegionTrans, Area, Row, Column)
gen_circle (Coin, Row, Column, 120)
```

What remains to do in `train_model` is to determine the contrast parameter and then to create the model using `create_shape_model`. Note that the value 'ignore_local_polarity' is used for the parameter `Metric` because the dark/light transitions can change locally due to the varying illumination.

```
Contrast := 20
HysteresisContrast := [Contrast/2, Contrast+6, 10]
reduce_domain (Image, Coin, ImageReduced)
create_shape_model (ImageReduced, 'auto', 0, rad(360), 'auto', 'none', \
  'ignore_local_polarity', HysteresisContrast, 5, \
  ModelID)
```

The procedure `find_coin` also first calls `locate_coin` to derive a region of interest that is as small as possible. This will speedup the matching process significantly. A circle of radius 30 in the center of the coin is used as the search ROI. Inside this region, `find_shape_models` is called to determine which coin can be seen and to precisely determine its position and orientation.

```
locate_coin (Image, Coin)
area_center (Coin, Area, Row, Column)
gen_circle (Circle, Row, Column, 35)
reduce_domain (Image, Circle, ImageReduced)
find_shape_models (ImageReduced, Models, 0, rad(360), 0.6, 1, 0, \
  'least_squares', 0, 0.9, Row, Column, Angle, Score, \
  Model)
```

The procedure `display_model` accesses the model edges and transforms them according to the found position and orientation and displays them overlaid on the image. As additional feedback, the type of the coin is displayed in the graphics window.

```
get_shape_model_contours (ModelContours, Model, 1)
vector_angle_to_rigid (0, 0, 0, Row, Column, Angle, HomMat2D)
affine_trans_contour_xld (ModelContours, ContoursAffinTrans, HomMat2D)
dev_display (Image)
dev_set_color ('green')
dev_set_line_width (2)
dev_display (ContoursAffinTrans)
set_tposition (WindowHandle, 24, 12)
write_string (WindowHandle, Names[Model])
```

Finally, the model is destroyed.

```
for i := 0 to 3 by 1
    clear_shape_model (Models[i])
endfor
```

9.3.4 Locate Components on a PCB

Example: `hdevelop/Matching/Component-Based/cbm_modules_simple.hdev`

The task of this example is to locate multiple components on a printed circuit board in one step (see [figure 9.7](#)). On a printed circuit board, typically multiple different objects are mounted, whose positions can vary because of the tolerances in the mounting process. To locate all objects quickly and in a robust manner, the component-based matching is used.

In the training step, each object is marked with a region of interest ([figure 9.7a](#)). The possible movement and rotation of the objects is known from the manufacturing process and is passed as a parameter to the training.

```
gen_rectangle2 (ComponentRegions, 318, 109, -1.62, 34, 19)
gen_rectangle2 (Rectangle2, 342, 238, -1.63, 32, 17)
gen_rectangle2 (Rectangle3, 355, 505, 1.41, 25, 17)
gen_rectangle2 (Rectangle4, 247, 448, 0, 14, 8)
gen_rectangle2 (Rectangle5, 237, 537, -1.57, 13, 10)
concat_obj (ComponentRegions, Rectangle2, ComponentRegions)
concat_obj (ComponentRegions, Rectangle3, ComponentRegions)
concat_obj (ComponentRegions, Rectangle4, ComponentRegions)
concat_obj (ComponentRegions, Rectangle5, ComponentRegions)
create_component_model (ModelImage, ComponentRegions, 20, 20, rad(25), 0, \
    rad(360), 15, 40, 15, 10, 0.8, [4,3,3,3,3], 0, \
    'none', 'use_polarity', 'true', ComponentModelID, \
    RootRanking)
```

Now, the component-based matching is able to find all objects in one step ([figure 9.7a](#)), returning the relative positions of each object.

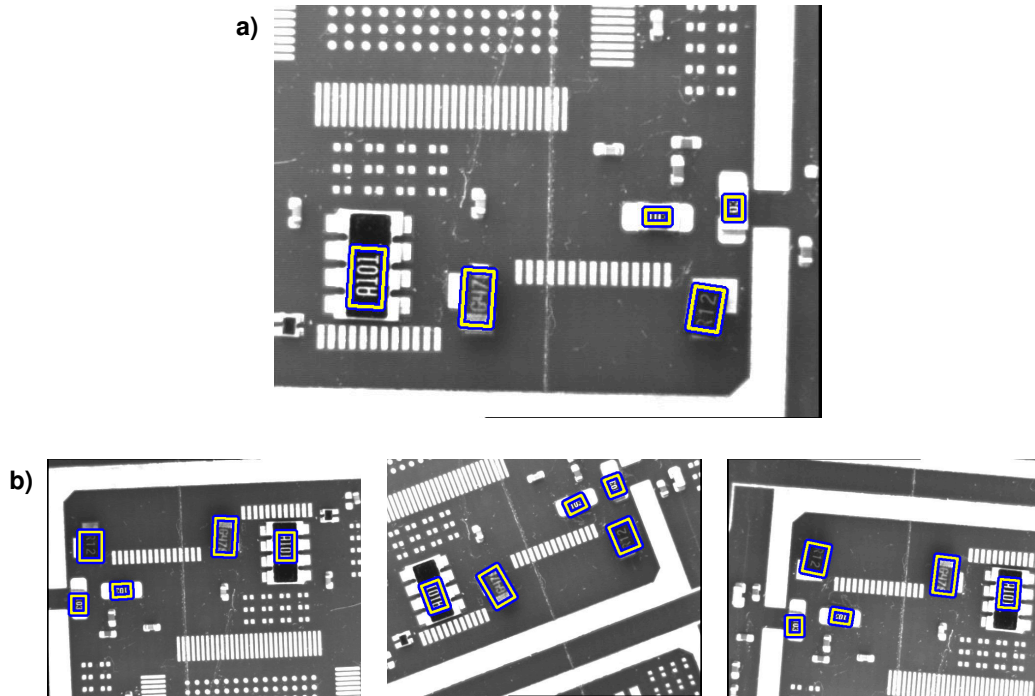


Figure 9.7: Component-based matching: (a) training objects; (b) objects, located in different images where the relation of the objects with respect to each other varies.

```
find_component_model (SearchImage, ComponentModelID, RootRanking, 0, \
    rad(360), 0.5, 0, 0.5, 'stop_search', \
    'search_from_best', 'none', 0.8, 'interpolation', 0, \
    0.8, ModelStart, ModelEnd, Score, RowComp, \
    ColumnComp, AngleComp, ScoreComp, ModelComp)
```

9.3.5 Check the State of a Dip Switch

Example: `hdevelop/Applications/Position-Recognition-2D/cbm_dip_switch.hdev`

The task of this example is to check a dip switch, i.e., to determine the positions of the single switches relative to the casing (see [figure 9.8](#)).

The task is solved using component-based matching. The dip switch consists of 14 components: 12 for the switches and two for the printed text on the casing (see [figure 9.8a](#)). The training is performed using training images that show all possible positions of the switches. From this, the system learns the possible movements of the switches relative to the casing.




Figure 9.8: (a) Model image of the dip switch with ROIs for the components; (b) located dip switch with state of the switches.

```
train_model_components (ModelImage, InitialComponents, TrainingImages, \
                        ModelComponents, 45, 45, 30, 0.95, -1, -1, rad(20), \
                        'speed', 'rigidity', 0.2, 0.5, ComponentTrainingID)
```

To make the recognition more robust, small tolerances are added to the trained movements.

```
modify_component_relations (ComponentTrainingID, 'all', 'all', 0, rad(4))
```

Now, the model can be created.

```
create_trained_component_model (ComponentTrainingID, 0, rad(360), 10, \
                                MinScoreComp, NumLevelsComp, 'auto', \
                                'none', 'use_polarity', 'false', \
                                ComponentModelID, RootRanking)
```

When searching for the dip switch, not only the casing but each single dip together with its relative position is returned. Based on this information, the global status of the switch can easily be derived.

```
find_component_model (SearchImage, ComponentModelID, RootRanking, 0, \
                      rad(360), 0, 0, 0.5, 'stop_search', 'prune_branch', \
                      'none', MinScoreComp, 'least_squares', 0, 0.9, \
                      ModelStart, ModelEnd, Score, RowComp, ColumnComp, \
                      AngleComp, ScoreComp, ModelComp)
```

9.3.6 Locating a Pipe Wrench in Different States

Example: `hdevelop/Applications/Position-Recognition-2D/cbm_pipe_wrench.hdev`




Figure 9.9: (a) Model image with specified ROIs for the components; (b) located pipe wrench in another state.

The task of this example is to locate a pipe wrench based on four predefined ROIs, two for each rigid part (see figure 9.9). Again, the task is solved using component-based matching.

Showing the system multiple example images trains the relative movements of the parts.

```
read_image (ModelImage, 'pipe_wrench/pipe_wrench_model')
gen_rectangle2 (InitialComponentRegions, 298, 363, 1.17, 162, 34)
gen_rectangle2 (Rectangle2, 212, 233, 0.62, 167, 29)
gen_rectangle2 (Rectangle3, 63, 444, -0.26, 50, 27)
gen_rectangle2 (Rectangle4, 120, 473, 0, 33, 20)
concat_obj (InitialComponentRegions, Rectangle2, InitialComponentRegions)
concat_obj (InitialComponentRegions, Rectangle3, InitialComponentRegions)
concat_obj (InitialComponentRegions, Rectangle4, InitialComponentRegions)
gen_empty_obj (TrainingImages)
for i := 1 to 4 by 1
    read_image (TrainingImage, 'pipe_wrench/pipe_wrench_training_'+i)
    concat_obj (TrainingImages, TrainingImage, TrainingImages)
endfor
train_model_components (ModelImage, InitialComponentRegions, TrainingImages, \
    ModelComponents, 22, 60, 30, 0.65, -1, -1, rad(60), \
    'speed', 'rigidity', 0.2, 0.3, ComponentTrainingID)
```

The components and their relations are displayed as follows:

```
get_training_components (ModelComponents, ComponentTrainingID, \
    'model_components', 'model_image', 'false', \
    RowRef, ColumnRef, AngleRef, ScoreRef)
dev_display (ModelComponents)

get_component_relations (Relations, ComponentTrainingID, i, 'model_image', \
    Row, Column, Phi, Length1, Length2, AngleStart, \
    AngleExtent)
dev_display (Relations)
```

Based on the training, the actual component model is created. Component relations are represented in a tree structure.

```
create_trained_component_model (ComponentTrainingID, rad(-90), rad(180), 10, \
                                0.6, 4, 'auto', 'none', 'use_polarity', \
                                'false', ComponentModelID, RootRanking)
get_component_model_tree (Tree, Relations, ComponentModelID, RootRanking, \
                          'model_image', StartNode, EndNode, Row, Column, \
                          Phi, Length1, Length2, AngleStart, AngleExtent)
dev_display (ModelImage)
dev_display (Tree)
dev_display (Relations)
```

Finally, test images are used to locate the pipe wrench. For each image, the model contours are overlaid and the shift and opening angle of the pipe wrench is visualized.

```
find_component_model (SearchImage, ComponentModelID, RootRanking, \
                      rad(-90), rad(180), 0, 0, 1, 'stop_search', \
                      'prune_branch', 'none', 0.6, 'least_squares', 0, \
                      0.7, ModelStart, ModelEnd, Score, RowComp, \
                      ColumnComp, AngleComp, ScoreComp, ModelComp)
dev_display (SearchImage)
NumFound := |ModelStart|
if (NumFound)
    get_found_component_model (FoundComponents, ComponentModelID, \
                              ModelStart, ModelEnd, RowComp, \
                              ColumnComp, AngleComp, ScoreComp, \
                              ModelComp, 0, 'false', RowCompInst, \
                              ColumnCompInst, AngleCompInst, \
                              ScoreCompInst)
    dev_display (FoundComponents)
    visualize_pipe_wrench_match (AngleCompInst, WindowHandle, \
                                RowCompInst, ColumnCompInst, RowRef, \
                                ColumnRef)
endif
```

9.3.7 Creating a Mosaic Image

Example: `hdevelop/Tools/Mosaicking/gen_projective_mosaic.hdev`

The task of this example is to create an image of the elongated printed circuit board depicted in [figure 9.10](#). Using standard image sizes, most of the image would be empty. The solution is to acquire images from multiple viewpoints and then to create a mosaic image using a point-based matching.

Multiple overlapping images of a printed circuit board are the input for the program. In a first step, significant points are extracted in each of these images. These points are the input for the point-based matching. In each step, two successive images are matched. The result of this process is a mapping from one image to the next.




Figure 9.10: Creating a mosaic image from multiple overlapping images.

```
points_foerstner (ImageF, 1, 2, 3, 200, 0.3, 'gauss', 'false', \
    RowJunctionsF, ColJunctionsF, CoRRJunctionsF, \
    CoRCJunctionsF, CoCCJunctionsF, RowAreaF, ColAreaF, \
    CoRRAreaF, CoRCAreaF, CoCCAreaF)
points_foerstner (ImageT, 1, 2, 3, 200, 0.3, 'gauss', 'false', \
    RowJunctionsT, ColJunctionsT, CoRRJunctionsT, \
    CoRCJunctionsT, CoCCJunctionsT, RowAreaT, ColAreaT, \
    CoRRAreaT, CoRCAreaT, CoCCAreaT)
proj_match_points_ransac (ImageF, ImageT, RowJunctionsF, ColJunctionsF, \
    RowJunctionsT, ColJunctionsT, 'ncc', 21, 0, \
    0, 480, 640, 0, 0.5, 'gold_standard', 1, \
    4364537, ProjMatrix, Points1, Points2)
ProjMatrices := [ProjMatrices, ProjMatrix]
```

These mappings are collected and finally used to construct a single high-resolution image of the complete PCB.

```
gen_projective_mosaic (Images, MosaicImage, 2, From, To, ProjMatrices, \
    'default', 'false', MosaicMatrices2D)
```

9.3.8 Locate Brochure Pages

Example: `hdevelop/Applications/Object-Recognition-2D/detect_brochure_pages.hdev`

The task of this example is to locate different pages of a HALCON brochure as depicted in [figure 9.11](#) using descriptor-based matching.

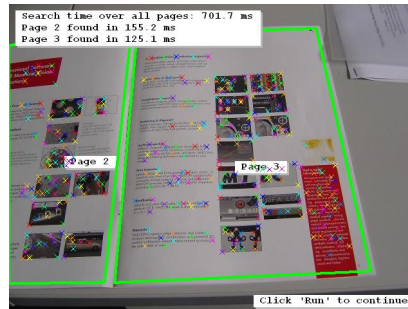


Figure 9.11: Different pages of a HALCON brochure are located via descriptor-based matching.

First, images that contain models of the different pages of the brochure are read, converted into gray value images, and their domains are reduced to a rectangular ROI. From this ROI, the operator `create_uncalib_descriptor_model` creates a descriptor-based model for each model image. That is, the detector, which extracts the interest points of the model, is selected (`harris_binomial`) and parameters for the descriptor, which builds characteristic descriptions of the gray-value neighborhood of the interest points, are set.

```
RowRoi := [10, 10, Height-10, Height-10]
ColRoi := [10, Width-10, Width-10, 10]
gen_rectangle1 (Rectangle, 10, 10, Height-10, Width-10)
for Index := 1 to NumModels by 1
    read_image (Image, 'brochure/brochure_page_' + Index$.2')
    rgb1_to_gray (Image, ImageGray)
    get_image_size (ImageGray, Width, Height)
    reduce_domain (ImageGray, Rectangle, ImageReduced)
    create_uncalib_descriptor_model (ImageReduced, 'harris_binomial', [], \
                                    [], ['min_rot','max_rot','min_scale', \
                                          'max_scale'], [-90,90,0.2,1.1], 42, \
                                    ModelID)
```

The origin of each model is moved from the center of gravity of the model to the center of the rectangle used to create the ROI. This is done with `set_descriptor_model_origin` so that the rectangle can be easily projected correctly onto the model when visualizing the results of the search in a later step. The adapted model is then stored in the tuple `ModelIDs`. The interest points extracted for each model are queried with `get_descriptor_model_points` and the number of extracted points is stored in the tuple `NumPoints`.

```
set_descriptor_model_origin (ModelID, -Height/2, -Width/2)
ModelIDs := [ModelIDs, ModelID]
get_descriptor_model_points (ModelID, 'model', 'all', Row_D, Col_D)
NumPoints := [NumPoints, |Row_D|]
endfor
```

Now, the unknown images are read and transformed into gray value images. For each model, the operator `find_uncalib_descriptor_model` searches for instances of the model in the image.

```
for Index1 := 1 to 12 by 1
  OutputString := []
  read_image (Image, 'brochure/brochure_' + Index1$.2')
  rgb1_to_gray (Image, ImageGray)
  for Index2 := 0 to |ModelIDs|-1 by 1
    find_uncalib_descriptor_model (ImageGray, ModelIDs[Index2], \
                                  'threshold', 800, \
                                  ['min_score_descr', \
                                   'guided_matching'], [0.003,'on'], \
                                  0.25, 1, 'num_points', HomMat2D, \
                                  Score)
```

If a valid 2D projective transformation matrix (homography) and a specific minimum score (that depends on the number of the extracted interest points) was obtained by the search, the points of the model are queried with `get_descriptor_model_points` and displayed by cross contours. Then, the rectangle specified for the creation of the model and its corner points are transformed with the 2D projective transformation matrix (`projective_trans_region` and `projective_trans_pixel`). The transformed corner points are used to calculate the angle between two neighboring edges of the perspectively projected rectangle (`angle_ll`). In the model, the edges are right angled. In the perspective deformed projection they should not deviate more than 20 degrees from the right angle. Thus, only if an angle between 70 and 110 degrees is obtained, the found instance of the model is accepted and the result is displayed.

```
if ((|HomMat2D|>0) and (Score > NumPoints[Index2]/4))
  get_descriptor_model_points (ModelIDs[Index2], 'search', 0, Row, \
                              Col)
  gen_cross_contour_xld (Cross, Row, Col, 6, 0.785398)
  projective_trans_region (Rectangle, TransRegion, HomMat2D, \
                          'bilinear')
  projective_trans_pixel (HomMat2D, RowRoi, ColRoi, RowTrans, \
                          ColTrans)
  angle_ll (RowTrans[2], ColTrans[2], RowTrans[1], ColTrans[1], \
            RowTrans[1], ColTrans[1], RowTrans[0], ColTrans[0], \
            Angle)
  if (Angle > 70 and Angle < 110)
    area_center (TransRegion, Area, Row, Column)
    disp_message (WindowHandle, 'Page '+(Index2+1), 'window', \
                  Row, Column, 'black', 'true')
  endif
endif
endfor
endfor
```

At the end of the program the model is cleared from memory.

```

for Index := 0 to |ModelIDs|-1 by 1
  clear_descriptor_model (ModelIDs[Index])
endfor

```

9.3.9 Locate Road Signs

Example: `hdevelop/Applications/Traffic-Monitoring/detect_road_signs.hdev`

The task of this example is to locate road signs in images as depicted in [figure 9.12](#). Available road signs comprise the attention sign and the dead end road sign.



Figure 9.12: A dead end road sign is located via uncalibrated perspective deformable matching.

The example first reads a synthetical image that shows the model of an attention sign. The first channel of the model is accessed ([access_channel](#)) and zoomed ([zoom_image_factor](#)). Before calling [create_planar_uncalib_deformable_model](#) to create the perspective deformable model, it is checked if the value '3' for the number of levels is suitable. For this, the operator [inspect_shape_model](#) is applied and the result is checked visually. After the creation of the model, the model is stored in the tuple `Models`.

```

read_image (ImageAttentionSign, 'road_signs/attention_road_sign')
access_channel (ImageAttentionSign, Image, Channel[0])
zoom_image_factor (Image, ImageZoomed, 0.1, 0.1, 'weighted')
inspect_shape_model (ImageZoomed, ModelImages, ModelRegions, 3, 20)
create_planar_uncalib_deformable_model (ImageZoomed, 3, 0.0, 0.0, 0.1, \
                                       ScaleRMin[0], ScaleRMax[0], 0.05, \
                                       1.0, 1.0, 0.5, 'none', \
                                       'use_polarity', 'auto', 'auto', [], \
                                       [], ModelID)

Models := ModelID

```

Then, the image containing a dead end road sign is read and the corresponding model is created. Here, the model is not available as artificial model but is derived from an image showing a typical street scenario. Thus, the processing of the image differs in parts from the proceeding used for the synthetical model. Amongst others, the domain of the image is reduced to a rectangle containing the road sign. The model obtained from the call to [create_planar_uncalib_deformable_model](#) is added to the tuple `Models`.

```

read_image (ImageDeadEnd, 'road_signs/dead_end_road_sign')
access_channel (ImageDeadEnd, Image, Channel[1])
gray_closing_shape (Image, ImageClosing, 5, 5, 'octagon')
zoom_image_factor (ImageClosing, ImageZoomed, 0.4, 0.4, 'weighted')
gen_rectangle1 (Rectangle1, 28, 71, 67, 95)
reduce_domain (ImageZoomed, Rectangle1, ImageReduced)
create_planar_uncalib_deformable_model (ImageReduced, 3, 0.0, 0.0, 0.1, \
                                         ScaleRMin[1], ScaleRMax[1], 0.05, \
                                         ScaleRMin[1], ScaleRMax[1], 0.1, \
                                         'none', 'use_polarity', 'auto', \
                                         'auto', [], [], ModelID)

Models := [Models, ModelID]

```

To search for both models in unknown images, a rectangular ROI is created as search space. Within this ROI, the procedure `determine_area_of_interest` automatically determines more specific regions of interest (ROI) for the road signs using a blob analysis. Each model is searched in a specific channel of the image as the attention sign can better be extracted in the blue channel and the dead end road sign is better extracted in the red channel. The channels are specified at the beginning of the program inside the tuple `Channel`. The individual model is then searched in the reduced image with `find_planar_uncalib_deformable_model`.

```

gen_rectangle1 (Rectangle, 115, 0, 360, 640)
for Index := 1 to 16 by 1
    read_image (Image, 'road_signs/street-'+Index$'.02')
    determine_area_of_interest (Image, Rectangle, AreaOfInterest)
    reduce_domain (Image, AreaOfInterest, ImageReduced)
    for Index2 := 0 to |Models|-1 by 1
        access_channel (ImageReduced, ImageChannel, Channel[Index2])
        find_planar_uncalib_deformable_model (ImageChannel, Models[Index2], \
                                              0, 0, ScaleRMin[Index2], \
                                              ScaleRMax[Index2], \
                                              ScaleCMin[Index2], \
                                              ScaleCMax[Index2], 0.8, 1, 0, \
                                              2, 0.4, [], [], HomMat2D, \
                                              Score)
    endfor
endfor

```

If a model could be found, i.e., if a 2D projective transformation matrix was returned, the 2D projective transformation matrix is used to project the contour of the specific model onto the found instance of the model inside the inspected image.

```

    if (|HomMat2D|)
        get_deformable_model_contours (ModelContours, Models[Index2], 1)
        projective_trans_contour_xld (ModelContours, ContoursProjTrans, \
                                       HomMat2D)
        dev_display (ContoursProjTrans)
    endif
endfor
endfor

```

At the end of the program the model is cleared from memory.


```
for Index1 := 0 to 1 by 1
    clear_deformable_model (Models[Index1])
endfor
```

9.4 Selecting Operators

9.4.1 Acquire Image(s)

Please refer to the [operator list for the method Image Acquisition](#) (see section 2.4 on page 17).

9.4.2 Radiometrically Calibrate Image(s)

Standard:

[radiometric_self_calibration](#), [lut_trans](#)

9.4.3 Rectify Image(s)

Operators for rectifying images are described in the [Solution Guide III-C](#).

9.4.4 Use Region Of Interest

Please refer to the [operator list for the method Region Of Interest](#) (see section 3.4 on page 26).

9.4.5 Determine Training Parameters

Standard:

Matching (Correlation-Based) [determine_ncc_model_params](#), [get_ncc_model_params](#)

Matching (Shape-Based) [determine_shape_model_params](#), [get_shape_model_params](#),
[inspect_shape_model](#)

Matching (Local Deformable) [determine_deformable_model_params](#),
[get_deformable_model_params](#)

Matching (Perspective Deformable) [determine_deformable_model_params](#),
[get_deformable_model_params](#)

Matching (Descriptor-Based) [get_descriptor_model_params](#)

9.4.6 Create (Train) Model

Standard:

Matching (Component-Based) `create_component_model`, `gen_initial_components`,
`train_model_components`, `create_trained_component_model`

Matching (Correlation-Based) `create_ncc_model`

Matching (Shape-Based) `create_shape_model`, `create_scaled_shape_model`,
`create_aniso_shape_model`, `create_shape_model_xld`,
`create_scaled_shape_model_xld`, `create_aniso_shape_model_xld`

Matching (Gray-Value-Based) `create_template_rot`

Matching (Point-Based) `points_foerstner`, `points_harris`

Matching (Local Deformable) `create_local_deformable_model`,
`create_local_deformable_model_xld`

Matching (Perspective Deformable) `create_planar_uncalib_deformable_model`,
`create_planar_calib_deformable_model`,
`create_planar_uncalib_deformable_model_xld`,
`create_planar_calib_deformable_model_xld`

Matching (Descriptor-Based) `create_uncalib_descriptor_model`,
`create_calib_descriptor_model`

Advanced:

Matching (Correlation-Based) `set_ncc_model_param`

Matching (Shape-Based) `set_shape_model_param`

Matching (Perspective Deformable) `set_deformable_model_param`

9.4.7 Find Model

Standard:

Matching (Component-Based) `find_component_model`

Matching (Correlation-Based) `find_ncc_model`

Matching (Shape-Based) `find_shape_model`, `find_scaled_shape_model`,
`find_aniso_shape_model`

Matching (Gray-Value-Based) `best_match_mg`

Matching (Point-Based) `proj_match_points_ransac`, `gen_projective_mosaic`

Matching (Local Deformable) `find_local_deformable_model`

Matching (Perspective Deformable) `find_planar_uncalib_deformable_model`,
`find_planar_calib_deformable_model`

Matching (Descriptor-Based) `find_uncalib_descriptor_model`,
`find_calib_descriptor_model`

Advanced:

Matching (Shape-Based) `find_shape_models`, `find_scaled_shape_models`,
`find_aniso_shape_models`

Matching (Gray-Value-Based) `best_match_pre_mg`, `fast_match_mg`

Matching (Point-Based) `proj_match_points_ransac_guided`,
`projective_trans_image_size`

9.4.8 Visualize Results

Please refer to the [operator list for the method Visualization](#) (see section 19.4 on page 297).

9.4.9 Destroy Model

Standard:

Matching (Component-Based) `clear_component_model`, `clear_training_components`

Matching (Correlation-Based) `clear_ncc_model`

Matching (Shape-Based) `clear_shape_model`

Matching (Gray-Value-Based) `clear_template`

Matching (Local Deformable) `clear_deformable_model`

Matching (Perspective Deformable) `clear_deformable_model`

Matching (Descriptor-Based) `clear_descriptor_model`

9.5 Relation to Other Methods

9.5.1 Methods that are Using Matching

1D Measuring (see [description](#) on page 49)

OCR (see [description](#) on page 245)

Variation Model (see [description](#) on page 143)

Bar Code (see [description](#) on page 211)

The pose or 2D projective transformation matrix (homography) returned by the matching operators can be used as the input for a so-called alignment. This means that either the position of an ROI is transformed relative to the movement of a specified object in the image, or the image itself is transformed so that the pixels are moved to the desired position.

A detailed description of alignment can be found in the Solution Guide II-B in [section 2.4.3.2](#) on page 42. For an example of using alignment as a preprocessing for [1D Measuring](#) on page 49 see the description of [pm_measure_board.hdev](#) on page 56. Further examples can be found in the description of the method [Variation Model](#) on page 143.

9.5.2 Alternatives to Matching

Blob Analysis (see [description](#) on page 31)

In some applications, the object to find can be extracted with classical segmentation methods. With the operators [area_center](#) and [orientation_region](#) you can then determine its pose and use this information, e.g., to align an ROI. With this approach, the execution time can be reduced significantly.

If the objects to be found appear only translated but not rotated or scaled, the morphological operators [erosion1](#) and [opening](#) can be used as binary matching methods. Unlike in other vision systems, in HALCON these operators are extremely fast.

9.6 Tips & Tricks

9.6.1 Speed Up

Many online applications require maximum speed. Because of its flexibility, HALCON offers many ways to achieve this goal. Below, the most common ones are listed (note that for most of them the focus is on the shape-based matching and those matching methods that use similar parameters like component-based matching and local or perspective deformable matching):

- Regions of interest are the standard way to increase the speed by processing only those areas where objects need to be inspected. This can be achieved using pre-defined regions, but also by an online generation of the regions of interest that depends on other objects in the image.
- If multiple objects are searched for, it is more efficient to use [find_scaled_shape_models](#) (or [find_aniso_shape_models](#)) for the shape-based matching instead of using [find_scaled_shape_model](#) (or [find_aniso_shape_model](#)) multiple times.

- Increasing the values for the parameters `MinContrast` and `Greediness` will decrease the execution time. However, you must make sure that the relevant objects will still be found.
- Using a lower value for the parameter `Contrast` during the training typically results in a better performance because more pyramid levels can be used. If the contrast is too low, irrelevant contours will also be included into the model, which typically causes a lower recognition rate and a decreased accuracy.
- Very small templates cannot be found as quickly as larger ones. The reason for this is the reduced number of significant contours within an image, which makes it harder to distinguish the template from other structures. Furthermore, with smaller templates fewer pyramid levels can be used.

9.7 Advanced Topics

9.7.1 High Accuracy

Sometimes very high accuracy is required. Here, some of the matching methods offer various interpolation methods to achieve this goal. However, with a poor image quality even a better interpolation method will not improve the accuracy any more - it will just become slower. Here, it is important to consider both the image quality and a reasonable interpolation method.

9.7.2 Use Timeout

For shape-based matching and correlation-based matching, you can set a timeout with the operators `set_shape_model_param` and `set_ncc_model_param`, respectively. Then, the operators for searching the template, e.g., `find_shape_model` or `find_ncc_model`, will return at the latest after the specified time. This mechanism is demonstrated in the example `hdevelop/Matching/Shape-Based/set_shape_model_timeout.hdev`.

Chapter 10

3D Matching

The aim of 3D matching is to find a specific 3D object in search data and determine its pose, i.e., the position and orientation of the object in space. For that, HALCON provides different approaches:

- For the *shape-based 3D matching*, a 3D shape model is generated from a 3D computer aided design (CAD) model, which must be available in one of the supported formats that are listed in the description of [read_object_model_3d](#) in the Reference Manual (e.g., DXF, PLY, or STL). The 3D shape model consists of 2D projections of the 3D object seen from different views. Analogously to the shape-based matching of 2D structures described in the chapter [Matching](#) on page 97, the 3D shape model is used to recognize instances of the object in an image. But here, instead of a 2D position, orientation, and scaling, the 3D pose of each instance is returned.
- For the *surface-based 3D matching*, a surface model is derived either from a CAD model that is available in one of the supported formats that are listed in the description of [read_object_model_3d](#) in the Reference Manual (e.g., DXF, PLY, or STL), or from a 3D object model that is available in OM3 format. The latter can be obtained from images using a 3D reconstruction approach, e.g., [stereo vision](#) on page 277. In contrast to shape-based 3D matching, the object is not searched within an image, but within a 3D scene that is available as a 3D object model. Like shape-based 3D matching, surface-based 3D matching returns the 3D pose of each instance of the object that can be located.

If you need the 3D pose of a planar object or a planar object part, we recommend to use a calibrated perspective matching approach instead of the 3D matching. Available approaches are the calibrated perspective deformable matching and the calibrated descriptor-based matching. Both are significantly faster and more convenient to use.

Note that for the shape-based 3D matching as well as for the calibrated perspective deformable matching and the calibrated descriptor-based matching a camera calibration is necessary. For the surface-based matching, the camera calibration is needed when acquiring the 3D object model(s) from images using a 3D reconstruction approach.

This section goes deeper into the 3D matching. For the calibrated perspective deformable matching and the calibrated descriptor-based matching we refer to the Solution Guide II-B, [section 3.6](#) on page 124

and [section 3.7](#) on page 136. Note that the description here provides you with an overview on the 3D matching approaches. Further details can be found in the Solution Guide III-C, [section 4.2](#) on page 81 for shape-based 3D matching and [section 4.3](#) on page 93 for surface-based 3D matching, and in the Reference Manual.




Figure 10.1: Shape-based 3D matching: (left) original image containing a clamp, (right) 3D model of the found clamp projected in the image and display of its pose.

10.1 Basic Concept

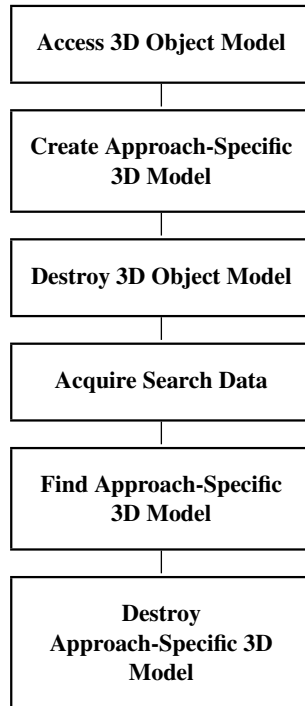
3D matching consists of the following basic steps:

10.1.1 Access 3D Object Model

Before creating the approach-specific 3D model, i.e., a 3D shape model or a surface model, the 3D object model describing the object of interest must be accessed.

- For *shape-based 3D matching* the 3D object model is accessed with the operator [read_object_model_3d](#). Supported CAD formats are, e.g., DXF, PLY, or STL. The complete list can be found in the description of the operator in the Reference Manual. The DXF format (DXF version AC1009, AutoCad release 12) may consist of the DXF entities POLYLINE (Poly-face meshes), 3DFACE, LINE, CIRCLE, ARC, ELLIPSE, SOLID, BLOCK, and INSERT.
- For *surface-based 3D matching* the 3D object model is accessed either from file using the operator [read_object_model_3d](#) or by an online 3D reconstruction. If the 3D object model is accessed from a file, besides the supported CAD formats also the OM3 format can be handled, which is a HALCON-specific format that can be obtained, e.g., by an offline 3D reconstruction.

The description of [read_object_model_3d](#) in the Reference Manual provides additional tips on how to obtain a suitable model.



10.1.2 Create Approach-Specific 3D Model

Having access to the 3D object model, the approach-specific 3D model can be created, which prepares the 3D object model for the specific 3D matching approach.

- For *shape-based 3D matching*, the operator `create_shape_model_3d` creates a 3D shape model. It is generated by computing different views of the 3D object model within a user-specified pose range. The views are obtained by placing virtual cameras around the object model and projecting the 3D object model into the image plane of each camera position. The resulting 2D shape representations of all views are stored in the 3D shape model. To avoid storage and runtime problems, the specified pose range should be restricted as much as possible so that the number of 2D projections that have to be computed and stored in the 3D shape model is minimized. To create the model views correctly, the camera parameters are needed. The camera parameters can be obtained by a camera calibration. How to apply a camera calibration is described in detail in the Solution Guide III-C, [section 3.2](#) on page 40.
- For *surface-based matching*, the operator `create_surface_model` creates a surface model by sampling the 3D object model with a certain distance, so that for an approximate matching, which is the first of three steps the later search consists of, only a reduced set of 3D points must be examined.

10.1.3 Destroy 3D Object Model

After creating the 3D shape model or the surface model, the 3D object model often is not needed anymore and can be destroyed for memory reasons using the operator `clear_object_model_3d`. If the 3D object model is still needed, e.g., for visualization purposes, this step must be moved, e.g., to the end of the application.

10.1.4 Acquire Search Data

If a *shape-based 3D matching* is applied, search images are acquired. For detailed information see the [description of this method](#) on page 13.

If a *surface-based 3D matching* is applied, the search space is not an image but a 3D scene that is available as 3D object model, which can be obtained from images using a 3D reconstruction approach. An overview on the available 3D reconstruction approaches is given in the Solution Guide III-C, [chapter 1](#) on page 9.

10.1.5 Find Approach-Specific 3D Model

With the approach-specific 3D model that was created by `create_shape_model_3d` or `create_surface_model` or that was read from file by `read_shape_model_3d` or `read_surface_model` (see [Re-use Approach-Specific 3D Model](#) on page 133), the object can be searched for in the search data, i.e., in images for *shape-based 3D matching* and in a 3D scene that is represented by a 3D object model for *surface-based matching*.

For the search, the operator `find_shape_model_3d` or `find_surface_model` is applied. Several parameters can be set to control the search process. For detailed information, we recommend to read the descriptions of the operators in the Reference Manual. Both operators return the pose of the matched model instance and a score that describes the quality of the match.

10.1.6 Destroy Approach-Specific 3D Model

When the approach-specific 3D model is not needed anymore, it is destroyed with the operator `clear_shape_model_3d` or `clear_surface_model`, respectively.

10.1.7 A First Example

An example for this basic concept is the following program. It applies a shape-based 3D matching to locate clamps.

The DXF file containing the 3D model of the clamp shown in [Figure 10.1](#) on page 128 is read with `read_object_model_3d`.

```
read_object_model_3d ('clamp_sloped', 'mm', [], [], ObjectModel3DID, \
                    DxfStatus)
```

For the creation of a 3D shape model, the camera parameters are needed. They can be determined by a camera calibration as described in the Solution Guide III-C, [section 3.2](#) on page 40. Here, they are known and just assigned to the variable `CamParam`.

```
CamParam := [0.01221, -2791, 7.3958e-6, 7.4e-6, 308.21, 245.92, 640, 480]
```

The creation of the 3D shape model is applied using `create_shape_model_3d`. There, besides the camera parameters the following further parameters are needed: the reference orientation and the pose range in which the object is expected, i.e., the minimum and maximum longitude and latitude, as well as the minimum and maximum distance between the camera and the center of the object's bounding box (which correspond to the radii of spheres that are built to compute the virtual camera positions). The range for the camera roll angle is set to a full circle. Further information about these parameters can be found in the Reference Manual and in the Solution Guide III-C, [section 4.2](#) on page 81.

```
create_shape_model_3d (ObjectModel3DID, CamParam, RefRotX, RefRotY, \
                      RefRotZ, 'gba', LongitudeMin, LongitudeMax, \
                      LatitudeMin, LatitudeMax, 0, rad(360), DistMin, \
                      DistMax, 10, 'min_face_angle', MinFaceAngle, \
                      ShapeModel3DID)
```

After creating the 3D shape model, the 3D object model is not needed anymore and is destroyed by the operator `clear_object_model_3d`.

```
clear_object_model_3d (ObjectModel3DID)
```

The search images are acquired and the actual matching is applied in each image using the operator `find_shape_model_3d`.

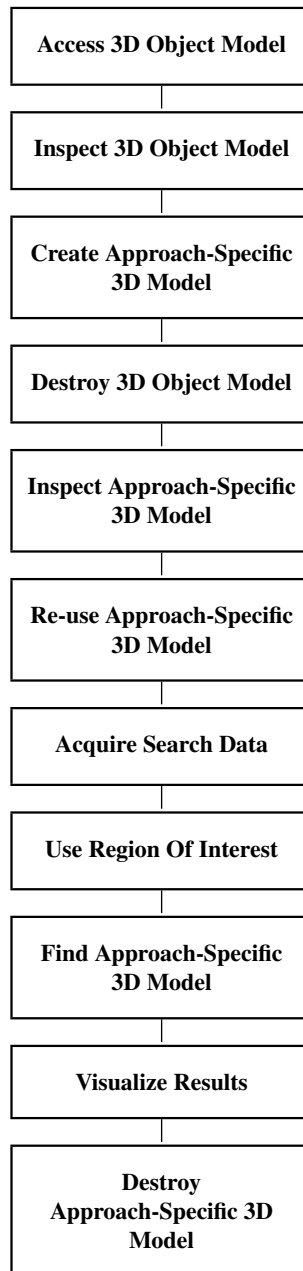
```
read_image (Image, 'clamp_sloped/clamp_sloped_'+ImageNo$'02')
find_shape_model_3d (Image, ShapeModel3DID, 0.7, 0.9, 5, ['num_matches', \
                'pose_refinement'], [2, 'least_squares_very_high'], \
                Pose, CovPose, Score)
```

At the end of the program, the 3D shape model is destroyed by the operator `clear_shape_model_3d`.

```
clear_shape_model_3d (ShapeModel3DID)
```

10.2 Extended Concept

Often, more than the essential steps are necessary. You can, e.g., store the approach-specific model into a file and read it from another application to separate the creation of the model from the actual search process. These and further advanced steps are described in the following sections.



10.2.1 Inspect 3D Object Model

3D object models can be inspected using [get_object_model_3d_params](#). Depending on the parameter specified in `GenParamNames`, you can query different attributes of the 3D object model like the number of points or the parameters of the bounding box. A procedure for the visualization of 3D object models

(disp_object_model_3d) can be found, e.g., with the HDevelop example program `hdevelop/3D-Matching/Surface-Based/find_surface_model.hdev`.

10.2.2 Inspect Approach-Specific 3D Model

Similar to the 3D object model, you can query also the parameters for the approach-specific 3D model:

- The 3D shape model that was created by `create_shape_model_3d` can be inspected using `get_shape_model_3d_params`. Several parameters like the reference orientation of the 3D object model, the camera parameters, or the pose range used for the creation of the 3D shape model can be queried.

The contours of a specific view of the 3D shape model can be accessed with `get_shape_model_3d_contours`. The contours can be used to visualize and rate the 3D shape model and thus decide if certain parameters have to be adjusted within the search process using `find_shape_model_3d` or if the creation of the 3D shape model has to be repeated with adjusted parameters.

- The surface model that was created by `create_surface_model` can be inspected using `get_surface_model_params`. Several parameters like the center of the model or the 3D points of the sampled model versions needed for different steps of the search process can be queried.

10.2.3 Re-use Approach-Specific 3D Model

Because of the complex calculations, the creation of a 3D model can be rather time consuming. Thus, if you need a 3D model more than once, it is recommended to store the approach-specific 3D model into a file rather than to repeat the model creation. To store the model into file you use the operator `write_shape_model_3d` for *shape-based 3D matching* and the operator `write_surface_model` for *surface-based 3D matching*. To read the 3D shape model or 3D surface model from file again to re-use the model in another application, you apply `read_shape_model_3d` or `read_surface_model`, respectively.

10.2.4 Use Region Of Interest

When searching in images, i.e., when using *shape-based 3D matching*, the search can be sped up using a region of interest. The more the region in which the objects are searched can be restricted, the faster and more robust the search will be.

For detailed information see the [description of this method](#) on page 19.

10.2.5 Visualize Results

A typical visualization task for *shape-based 3D matching* is to use the operator `project_object_model_3d` or `project_shape_model_3d` to project the outline of the 3D object model or the 3D shape model into the image. For both operators, an object-camera pose is needed,

which can be determined, e.g., by the operator `create_cam_pose_look_at_point`. Another common visualization task is to display the pose that is obtained by `find_shape_model_3d`.

A typical visualization task for *surface-based 3D matching* is to use the operator `project_object_model_3d` and a known object-camera pose to project the outline of the 3D object model that is used as model into an image. For example, if the 3D object model that is used as search data was reconstructed with a multi-view stereo setup, the model can be projected into one of the used stereo images. This type of visualization is applied, e.g., in the HDevelop example program `hdevelop/Applications/Robot-Vision/locate_pipe_joints_stereo.hdev`. Another common visualization task is to display the 3D object model that is used as model not in an image but within a 3D scene, e.g., within the 3D object model that is used as search data. A procedure that allows to interactively inspect such a 3D scene (`disp_object_model_3d`) is provided, e.g., with the HDevelop example program `hdevelop/3D-Matching/Surface-Based/find_surface_model.hdev`.

For general information about visualizing results see the description of the method [Visualization](#) on page 289.

10.3 Programming Examples

This section shows how to use HALCON for 3D matching.

10.3.1 Recognize 3D Clamps and Their Poses in Images

Example: `hdevelop/Applications/Position-Recognition-3D/3d_matching_clamps.hdev`

This example shows how to apply shape-based 3D matching to find a 3D object, in particular the clamp shown in [Figure 10.1](#) on page 128, in images. Besides the basic steps introduced in the first example, advanced steps are applied and a procedure is provided that helps to interactively determine an appropriate pose range for the creation of the 3D shape model.

At the beginning of the program, you can select by commenting or uncommenting the corresponding code lines whether you want to use an already created 3D shape model or whether you want to create a new one. Further, the camera parameters are assigned to a variable and the 3D object model of the clamps is loaded.

```
ReCreateShapeModel3D := false
* ReCreateShapeModel3D := true
CamParam := [0.01221, -2791, 7.3958e-6, 7.4e-6, 308.21, 245.92, 640, 480]
read_object_model_3d ('clamp_sloped', 'mm', [], [], ObjectModel3DID, \
                    DxfStatus)
```

The procedure `inspect_object_model_3d` can now be used to interactively visualize the model, set the reference pose, and determine further parameters for the creation of the 3D shape model. For example, selected images that were acquired before calling the procedure can be successively added to the visualization by selecting the buttons `Next Image` or `Previous Image`. With the mouse, you can rotate the 3D object model and, after selecting the mouse mode `Move Image`, move the underlying image so that the object model matches the image (see [Figure 10.2](#)). The returned pose for each image can be

added to the pose range by the button Add to Pose Range. When leaving the procedure by selecting Exit, the procedure returns values for a pose range that can be used to create a 3D shape model.

```
gen_empty_obj (Images)
for Index := 1 to 3 by 1
    read_image (Image, 'clamp_sloped/clamp_sloped_'+Index$'02')
    concat_obj (Images, Image, Images)
endfor
inspect_object_model_3d (Images, ObjectModel3DID, CamParam, RefRotX, \
    RefRotY, RefRotZ, LongitudeMin, LongitudeMax, \
    LatitudeMin, LatitudeMax, CamRollMin, CamRollMax, \
    DistMin, DistMax, MinFaceAngle)
```

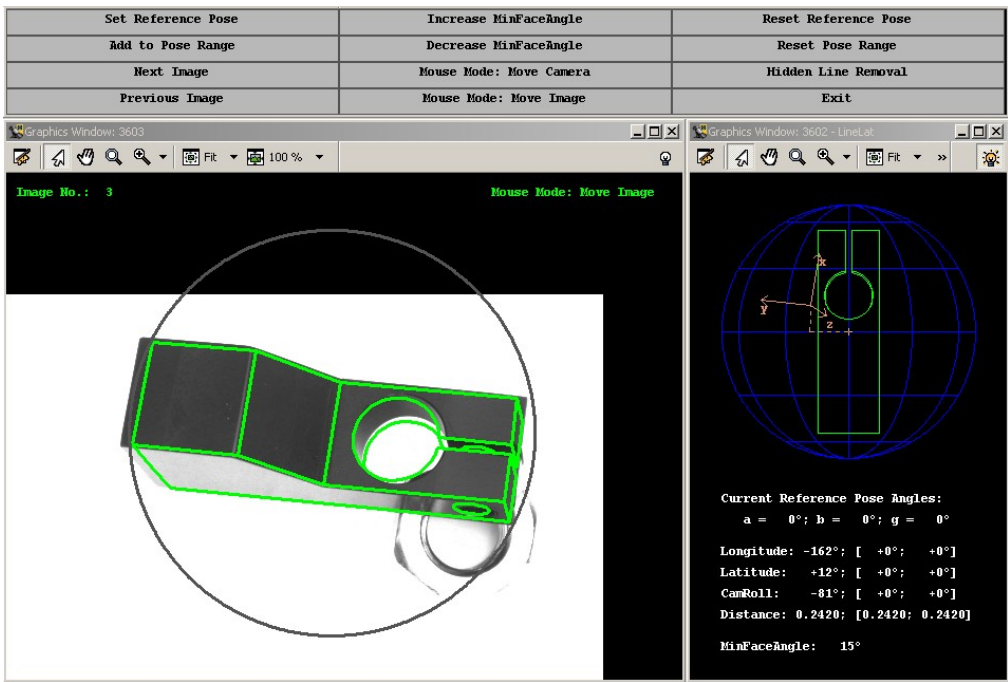


Figure 10.2: Interactive determination of the reference pose and the pose range for the 3D matching of clamps.

Additionally, within the procedure a Hidden Line Removal mode can be selected and the minimum face angle (Increase MinFaceAngle, Decrease MinFaceAngle) can be set. The latter is used, e.g., to suppress edges that are needed to approximate curved structures in the 3D object model but do not appear in the image and thus cannot be matched in the image (see Figure 10.3).

Within the main procedure, the operator `create_shape_model_3d` creates a new 3D shape model if the variable `ReCreateShapeModel3D` was set to true. For this, the parameters determined by the procedure described above are used. The resulting 3D shape model is written into file using `write_shape_model_3d`. If an already existing 3D shape model is used (`ReCreateShapeModel3D` set

to false) and the procedure was used mainly for visualization purposes, the existing 3D shape model is read from file using `read_shape_model_3d`. Since the 3D shape model is available now, we do not need the 3D object model anymore. It is destroyed by `clear_object_model_3d`.

```
if (ReCreateShapeModel3D)
    create_shape_model_3d (ObjectModel3DID, CamParam, RefRotX, RefRotY, \
                          RefRotZ, 'gba', LongitudeMin, LongitudeMax, \
                          LatitudeMin, LatitudeMax, 0, rad(360), DistMin, \
                          DistMax, 10, 'min_face_angle', MinFaceAngle, \
                          ShapeModel3DID)
    write_shape_model_3d (ShapeModel3DID, 'clamp_sloped_user.sm3')
else
    read_shape_model_3d ('clamp_sloped_35.sm3', ShapeModel3DID)
endif
clear_object_model_3d (ObjectModel3DID)
```

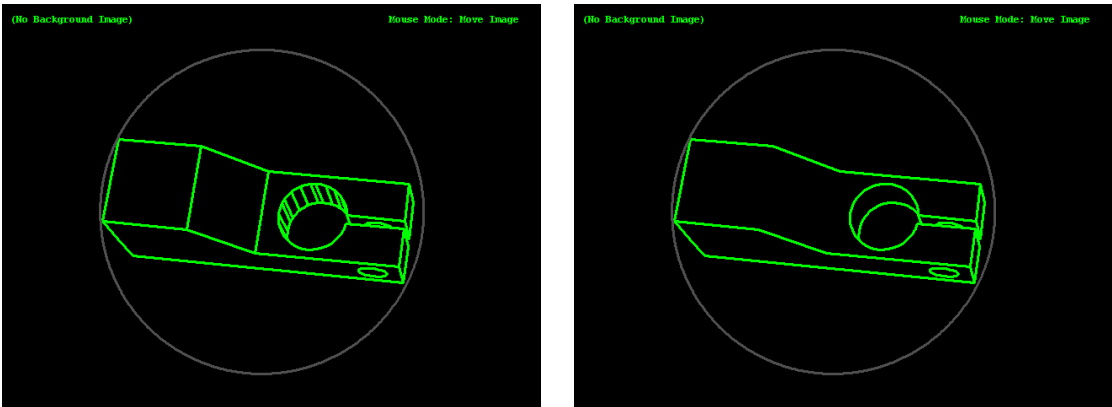


Figure 10.3: Different minimum face angles: (left) 8 degrees, (right) 45 degrees.

For the actual 3D matching, all images that contain the clamps are read and for each image the operator `find_shape_model_3d` searches for instances of the 3D shape model in the image. With the resulting values, the edges of the 3D shape model are projected into the image using `project_shape_model_3d`. The numerical values of the pose are displayed using the procedure `display_match_pose`.

```

for ImageNo := 1 to 35 by 1
  read_image (Image, 'clamp_sloped/clamp_sloped_'+ImageNo$'02')
  find_shape_model_3d (Image, ShapeModel3DID, 0.7, 0.9, 5, ['num_matches', \
    'pose_refinement'], [2,'least_squares_very_high'], \
    Pose, CovPose, Score)
  for I := 0 to |Score|-1 by 1
    PoseI := Pose[I*7:I*7+6]
    CovPoseI := CovPose[I*6:I*6+5]
    ScoreI := Score[I]
    project_shape_model_3d (ModelContours, ShapeModel3DID, CamParam, \
      PoseI, 'true', 0.523599)
    display_match_pose (ShapeModel3DID, PoseI, WindowHandle)
  endfor
endfor

```

Inside the procedure `display_match_pose` the reference point and the camera parameters are queried by `get_shape_model_3d_params` to align the displayed strings in the image.

```

get_shape_model_3d_params (ShapeModel3DID, 'reference_point', \
  ReferencePoint)
get_shape_model_3d_params (ShapeModel3DID, 'cam_param', CamParam)
pose_to_hom_mat3d (Pose, HomMat3D)
affine_trans_point_3d (HomMat3D, ReferencePoint[0], ReferencePoint[1], \
  ReferencePoint[2], X, Y, Z)
project_3d_point (X, Y, Z, CamParam, Row, Column)
set_tposition (WindowHandle, Row, Column - 10)
write_string (WindowHandle, 'Pose:')
set_tposition (WindowHandle, Row+15, Column)
write_string (WindowHandle, 'X: '+(1000*Pose[0])$'4.1f'+ ' mm')
... etc. ...

```

At the end of the program, the shape model is destroyed by the operator `clear_shape_model_3d`.

```

clear_shape_model_3d (ShapeModel3DID)

```

10.3.2 Recognize Pipe Joints and Their Poses in a 3D Scene

Example: `hdevelop/Applications/Robot-Vision/locate_pipe_joints_stereo.hdev`

This example shows how to apply surface-based 3D matching to find a 3D object, in particular a pipe joint, in a 3D scene given as a 3D object model.

First, the model is accessed. It is available as a CAD model in PLY format and is accessed with `read_object_model_3d`. From this 3D object model the corresponding surface model is created using `create_surface_model`.

```

read_object_model_3d ('pipe_joint.ply', 'm', [], [], PipeJointOM3DID, \
  Status)
create_surface_model (PipeJointOM3DID, 0.03, [], [], PipeJointSMID)

```


Note that a visualization with a hidden surface removal is possible only for 3D object models that were prepared for shape-based 3D matching. Thus, although we do not want to apply a shape-based 3D matching we prepare the 3D object model for it to enable a proper visualization.

```
prepare_object_model_3d (PipeJointOM3DID, 'shape_based_matching_3d', 'true', \
                        [], [])
```

In the example, the 3D scene that builds the search data is accessed using multi-view stereo. Note that the following code shows only a part of the stereo reconstruction. The rest is explained in [section 18.3.2](#) on page 284. [Figure 10.4](#) shows the images that are used to reconstruct one of the 3D scenes.

```
for Index := 1 to NumImages by 1
  read_multi_view_stereo_images (Images, ImagePath, ImagePrefix, Index, \
                                NumCameras)
  reconstruct_surface_stereo (Images, StereoModelID, PipeJointPileOM3DID)
```



Figure 10.4: Images used for the multi-view stereo reconstruction of the search data (3D scene).

The matching is applied with `find_surface_model`. It searches the surface model in the 3D scene and returns the poses of the three best matching model instances.

```
NumMatches := 3
Params := ['num_matches', 'pose_ref_scoring_dist_rel']
Values := [NumMatches, 0.02]
find_surface_model (PipeJointSMID, PipeJointPileOM3DID, 0.04, 0.15, \
                  MinScore, 'false', Params, Values, Poses, Scores, \
                  SurfaceMatchingResultID)
```

To visualize the result of each match the 3D object model that was accessed from the CAD model and that was prepared for visualization is transformed with the pose that was returned for the specific match. The transformed 3D object model is then projected into one of the images that were used to reconstruct the 3D scene that was used as search data (see [figure 10.5](#)).




Figure 10.5: Projection of the three best matches into one of the multi-view stereo images.

```
dev_display (Img)
for MatchIndex := 0 to |Scores|-1 by 1
    pose_to_hom_mat3d (Poses[MatchIndex*7:(MatchIndex*7)+6], \
        HomMat3DObjToWorld)
    affine_trans_object_model_3d (PipeJointOM3DID, HomMat3DObjToWorld, \
        ObjectModel3DIDAffineTrans)
    project_object_model_3d (ModelContours, ObjectModel3DIDAffineTrans, \
        CamParam0, WorldPose0, ['data', \
            'hidden_surface_removal'], ['faces', \
            'true'])
    dev_display (ModelContours)
```

When the different 3D object models and the surface model are not needed anymore, they are destroyed by the operators `clear_object_model_3d` or `clear_surface_model`, respectively.

```
clear_object_model_3d (ObjectModel3DIDAffineTrans)
endfor
clear_object_model_3d (PipeJointPileOM3DID)
endfor
clear_surface_model (PipeJointSMID)
clear_object_model_3d (PipeJointOM3DID)
```

10.4 Selecting Operators

10.4.1 Access 3D Object Model

Standard:

`read_object_model_3d`

Advanced:

```
prepare_object_model_3d
```

10.4.2 Inspect 3D Object Model

Standard:

```
get_object_model_3d_params
```

10.4.3 Create Approach-Specific 3D Model

Standard:

```
create_shape_model_3d, create_surface_model
```

Advanced:

```
convert_point_3d_cart_to_spher, convert_point_3d_spher_to_cart
```

10.4.4 Destroy 3D Object Model

Standard:

```
clear_object_model_3d
```

10.4.5 Inspect Approach-Specific 3D Model

Standard:

```
get_shape_model_3d_contours, get_shape_model_3d_params, get_surface_model_params
```

10.4.6 Re-use Approach-Specific 3D Model

Standard:

```
write_shape_model_3d, read_shape_model_3d, write_surface_model,  
read_surface_model
```

10.4.7 Acquire Search Data

For *shape-based 3D matching*, please refer to the [operator list for the method Image Acquisition](#) (see section 2.4 on page 17).

For *surface-based 3D matching*, please refer to the descriptions of the 3D reconstruction approaches in the [Solution Guide III-C](#) and in the Reference Manual. An overview on the available 3D reconstruction approaches is given in the Solution Guide III-C, [chapter 1](#) on page 9.

10.4.8 Use Region Of Interest

Please refer to the [operator list for the method Region Of Interest](#) (see section 3.4 on page 26).

10.4.9 Find Approach-Specific 3D Model

Standard:

[find_shape_model_3d](#), [find_surface_model](#)

10.4.10 Visualize Results

Standard:

[project_shape_model_3d](#), [project_object_model_3d](#)

Further operators can be found in the [operator list for the method Visualization](#) (see section 19.4 on page 297).

10.4.11 Destroy Approach-Specific 3D Model

Standard:

[clear_shape_model_3d](#), [clear_surface_model](#)

10.5 Relation to Other Methods

10.5.1 Alternatives to 3D Matching

Matching (calibrated perspective deformable)

If the 3D pose of planar objects or object parts are searched for, the calibrated perspective, deformable matching is a fast alternative to 3D matching. There, no 3D model has to be provided, as only a 2D model

is needed, which can easily be obtained from training images. As no 2D projections of a 3D model are computed, the calibrated perspective, deformable matching is significantly faster than the shape-based 3D matching. The calibrated perspective, deformable matching is described in the Solution Guide II-B in [section 3.6](#) on page [124](#).

Matching (calibrated descriptor-based)

If the 3D pose of planar objects or object parts are searched for and the objects are textured in a way that distinctive points can be extracted, the calibrated descriptor-based matching is even faster than the calibrated perspective, deformable matching, especially for a large search space. On the other side, it is less accurate. The calibrated descriptor-based matching is described in the Solution Guide II-B in [section 3.7](#) on page [136](#).

Chapter 11

Variation Model

The main principle of a variation model is to compare one or more images to an ideal image to find significant differences. With this, you can, e.g., identify incorrectly manufactured objects by comparing them to correctly manufactured objects. The ideal image is often obtained by a training using several reference images. Besides the ideal image, the training derives information about the allowed gray value variation for each point of the image. This information is stored in a so-called variation image. Both images are used to create a variation model, to which other images can be compared.

The advantage of the variation model is that images can be directly compared by their gray values and the comparison is spatially weighted by the variation image.

11.1 Basic Concept

Variation Model mainly consists of the following parts:


11.1.1 Acquire Image(s)

Both for the training and for the comparison, images are acquired.

For detailed information see the [description of this method](#) on page 13.

11.1.2 Create Variation Model

First, you have to create the variation model used for the image comparison by the operator `create_variation_model`. It stores information that is successively added during the following steps.



11.1.3 Align ROIs Or Images

For training the variation model, all training images must be placed in the same position and with the same orientation. Thus, before training the model the objects must be aligned. Similarly, the images that are to be compared to the variation model must be aligned.

How to perform alignment using shape-based matching is described in the Solution Guide II-B in [section 2.4.3.2](#) on page 42.

11.1.4 Train Variation Model

The variation model can be trained by providing a set of images containing good samples. The operator [train_variation_model](#) uses the training images to calculate an ideal image and a corresponding variation image. The variation image represents the amount of gray value variation, i.e., the tolerance, at every point of the image.

If you do not want to train the model with multiple images, e.g., for memory reasons, you can also use a single image as ideal image, but then you need knowledge about the spatial distribution of the variations. If you know, e.g., that the acceptable tolerances are near the edges of your object, you can create the variation image by applying an edge filter to the ideal image. For this proceeding, the training step is skipped.

11.1.5 Prepare Variation Model

To prepare the variation model for the image comparison, the ideal image and the variation image are converted into two threshold images. If you have trained the model with several training images using `train_variation_model`, you prepare the model with the operator `prepare_variation_model`. If you use a single image as ideal image and create the corresponding variation image manually by filtering the ideal image, e.g., using `sobel_amp`, `edges_image`, or `gray_range_rect`, you have to use `prepare_direct_variation_model` for the preparation because the ideal image and the variation image are not yet connected to the variation model.

The obtained threshold images can be directly read out using the operator `get_thresh_images_variation_model`.

11.1.6 Compare Variation Model

You compare an image to the prepared variation model with the operator `compare_variation_model`. There, the two thresholds obtained during the preparation step (and stored in the variation model) are used to determine a region containing all points of the image that significantly differ from the model. Extended parameter settings for the comparison are available when using `compare_ext_variation_model` instead.

11.1.7 Destroy Variation Model

When you no longer need the variation model, you destroy it with the operator `clear_variation_model`.

11.1.8 A First Example

Example: `hdevelop/Applications/Print-Inspection/print_check.hdev`

An example for this basic concept is the following program. Here, the logos on the pen clips depicted in [figure 11.1](#) are inspected using a variation model. The model is trained successively by a set of images containing correct prints. As the pen clips move from image to image, in each image the pen clip is located and aligned by shape-based matching. After the training, the variation model is prepared for the comparison with other images.

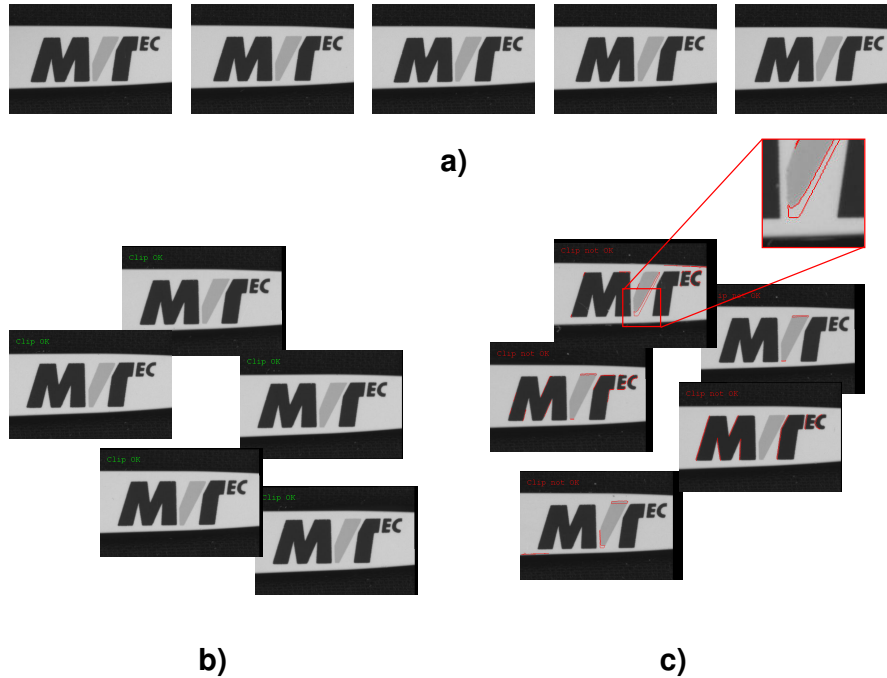


Figure 11.1: Inspection of the logo that is printed on the pen clip: (a) training images; (b) accepted images; (c) rejected images.

```

create_shape_model (ImageReduced, 5, rad(-10), rad(20), 'auto', 'none', \
                    'use_polarity', 20, 10, ShapeModelID)
create_variation_model (Width, Height, 'byte', 'standard', VariationModelID)
for I := 1 to 15 by 1
    read_image (Image, 'pen/pen-'+I$'02d')
    find_shape_model (Image, ShapeModelID, rad(-30), rad(60), 0.5, 1, 0.5, \
                     'least_squares', 0, 0.9, Row, Column, Angle, Score)
    if (|Score| = 1)
        vector_angle_to_rigid (Row, Column, Angle, RowRef, ColumnRef, 0, \
                               HomMat2D)
        affine_trans_image (Image, ImageTrans, HomMat2D, 'constant', \
                             'false')
        train_variation_model (ImageTrans, VariationModelID)
    endif
endfor
get_variation_model (MeanImage, VarImage, VariationModelID)
prepare_variation_model (VariationModelID, 20, 3)

```

During the inspection phase, also the images that are to be inspected are aligned by shape-based matching. Then, the images are compared to the variation model to check the print for errors. The erroneous regions are extracted and displayed. At the end of the program, the handles for the shape model and the variation model are destroyed.

```

find_shape_model (Image, ShapeModelID, rad(-10), rad(20), 0.5, 1, 0.5, \
                  'least_squares', 0, 0.9, Row, Column, Angle, Score)
if (|Score| = 1)
    vector_angle_to_rigid (Row, Column, Angle, RowRef, ColumnRef, 0, \
                          HomMat2D)
    affine_trans_image (Image, ImageTrans, HomMat2D, 'constant', 'false')
    reduce_domain (ImageTrans, RegionROI, ImageReduced)
    compare_variation_model (ImageReduced, RegionDiff, VariationModelID)
    connection (RegionDiff, ConnectedRegions)
    select_shape (ConnectedRegions, RegionsError, 'area', 'and', 20, \
                  1000000)
    dev_display (ImageTrans)
endif
clear_shape_model (ShapeModelID)
clear_variation_model (VariationModelID)

```

11.2 Extended Concept

Often more than the essential steps are necessary. In many cases, after the comparison, a visualization of the result is required. Additionally, you might want to check the quality of the training or you need to save memory. The advanced steps are described in the following sections.

11.2.1 Check Model Quality

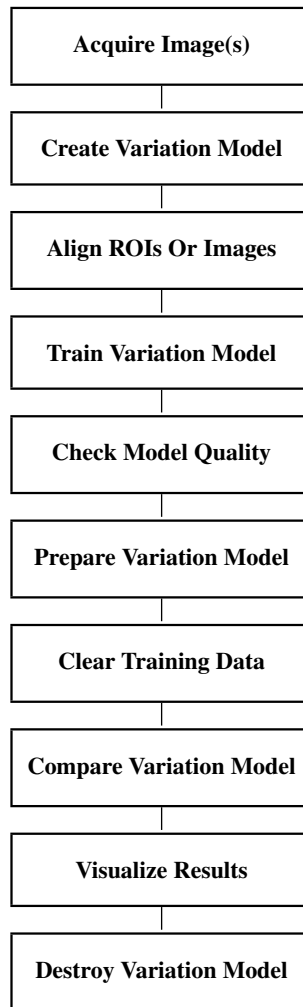
After training a variation model with several images, the image of the ideal object and the corresponding variation image can be queried by the operator `get_variation_model` to check if the images used for the training all contained similar objects. If the variation image contains large variations in areas that should exhibit no variations, this leads to the conclusion that at least one of the training images contained a bad object.

11.2.2 Clear Training Data

After the preparation of the variation model, you can reduce the amount of memory needed for the variation model by applying the operator `clear_train_data_variation_model`. But this is only recommended if you do not need to use the variation model for anything else than for the actual comparison anymore. A further training or the application of `get_variation_model` is not possible after deleting the training data.

11.2.3 Visualize Results

A typical visualization task is to display the image and mark the parts of the image that do not correspond to the model. These parts can be explicitly extracted by applying `connection` to separate the connected



components of the region obtained by the image comparison and afterwards selecting the regions that are within a specific area range by [select_shape](#).

For detailed information see the [description of this method](#) on page 289.

11.3 Programming Examples

This section gives a brief introduction to using HALCON for variation model.

11.3.1 Inspect a Printed Logo Using a Single Reference Image

Example: `solution_guide/basics/variation_model_single.hdev`

This example inspects the same logo as the first example, but instead of training the variation model with multiple images, a single image is used as ideal image.




Figure 11.2: Model for the inspection of the pen clips: a) ideal image; b) variation image; c) inspected image with displayed errors.

From the ideal image a shape model is derived for the later alignment and the variation image is determined manually. To obtain a variation image with tolerances at the edges of the object, the edges in the domain of the corresponding region of interest (ImageReduced) are extracted by the operator `edges_sub_pix`. These edges are converted into regions and the regions are enlarged by a dilation to create a generic image that contains the slightly enlarged boundaries of the object. As the conversion of subpixel-precise edges into pixel-precise regions leads to aliasing, the edges are scaled before applying the conversion. After the region processing, the generic image is zoomed back to the original size using a weighting function, which additionally smoothes the borders of the regions. A further smoothing is realized by the operator `binomial_filter`. The image is used now as variation image (see figure 11.2) and the variation model is prepared for the image comparison.

```

create_shape_model (ImageReduced, 'auto', -rad(10), rad(20), 'auto', 'auto', \
                    'use_polarity', [40,50], 40, ShapeModelID)
edges_sub_pix (ImageReduced, Edges, 'sobel_fast', 0.5, 10, 20)
hom_mat2d_identity (HomMat2DIdentity)
hom_mat2d_scale (HomMat2DIdentity, 4, 4, 0, 0, HomMat2DScale)
affine_trans_contour_xld (Edges, ZoomedEdges, HomMat2DScale)
gen_image_const (VarImageBig, 'byte', 4*Width, 4*Height)
count_obj (ZoomedEdges, NEdges)
for i := 1 to NEdges by 1
    select_obj (ZoomedEdges, ObjectSelected, i)
    get_contour_xld (ObjectSelected, RowEdge, ColEdge)
    gen_region_polygon (Region1, RowEdge, ColEdge)
    dilation_circle (Region1, RegionDilation, 2.5)
    paint_region (RegionDilation, VarImageBig, VarImageBig, 255, 'fill')
endfor
zoom_image_size (VarImageBig, VarImageSmall, Width, Height, 'weighted')
binomial_filter (VarImageSmall, VarImage, 3, 3)
create_variation_model (Width, Height, 'byte', 'direct', VarModelID)
prepare_direct_variation_model (Image, VarImage, VarModelID, 15, 4)

```

During the inspection, in each image to be checked the boundary of the object is searched for to align the image to the ideal image. The aligned image is then compared to the variation model using [compare_ext_variation_model](#). As the mode absolut is set, you can alternatively use the operator [compare_variation_model](#). Differing regions of a certain size are obtained and stored in NDefects.

```

for i := 1 to 30 by 1
    read_image (Image, 'pen/pen-'+i$'02d')
    find_shape_model (Image, ShapeModelID, -rad(10), rad(20), 0.5, 1, 0.5, \
                     'least_squares', 0, 0.9, Row, Column, Angle, Score)
    if (|Score| # 0)
        vector_angle_to_rigid (Row, Column, Angle, ModelRow, ModelColumn, 0, \
                               HomMat2D)
        affine_trans_image (Image, ImageAffinTrans, HomMat2D, 'constant', \
                            'false')
        reduce_domain (ImageAffinTrans, LogoArea, ImageReduced1)
        compare_ext_variation_model (ImageReduced1, RegionDiff, VarModelID, \
                                    'absolute')
        connection (RegionDiff, ConnectedRegions)
        select_shape (ConnectedRegions, SelectedRegions, 'area', 'and', 10, \
                      99999)
        count_obj (SelectedRegions, NDefects)
    endif
endfor

```

Finally, the handle for the variation model is deleted.

```

clear_shape_model (ShapeModelID)
clear_variation_model (VarModelID)

```

11.3.2 Inspect a Printed Logo under Varying Illumination

Example: `solution_guide/basics/variation_model_illumination.hdev`

This example inspects caps of bottles using a variation model (see [figure 11.3](#)). The difficulty here is the changing illumination during the inspection. Like in the example described before, a single image is used as ideal image.




Figure 11.3: Model for the inspection of the caps: a) ideal image; b) variation image; c) results of image comparison for images taken under different illumination conditions.

The variation model here is obtained by filtering the ideal image with the operator `sobel_amp`.

```

create_variation_model (Width, Height, 'byte', 'direct', VariationID)
sobel_amp (ModelImage, VarImage, 'sum_abs', 5)
prepare_direct_variation_model (ModelImage, VarImage, VariationID, [20,25], \
                                [1.6,1.6])

```

To compensate for the changing illumination, inside the procedure `get_grayval_range` the gray value range of the variation model is determined.

```

get_grayval_range (ModelImage, RegionROI, RegionForeground, \
                  RegionBackground, BackgroundGVModel, ForegroundGVModel)

```

Inside the procedure the ideal image is reduced to the domain of the region of interest that encloses the print on the cap. Then, the reduced image is split into foreground and background using `bin_threshold` and `difference`. Finally, for both regions the mean and standard deviation of the gray values is queried by the operator `intensity`.

```

reduce_domain (Image, RegionROI, ImageReduced)
bin_threshold (ImageReduced, RegionBackground)
difference (RegionROI, RegionBackground, RegionForeground)
intensity (RegionForeground, Image, ForegroundGVal, DeviationFG)
intensity (RegionBackground, Image, BackgroundGVal, DeviationBG)

```

The inspection of the caps is realized inside the procedure `inspect_cap`.

```

inspect_cap (rImage, RegionROI, WindowHandle, ModelID, VariationID, \
            RowModel, ColumnModel, BackgroundGVModel, \
            ForegroundGVModel)

```

There, in each image the object is searched for and aligned as described in the previous example. Then, the procedure `get_grayval_range` derives also the gray value range for the ROI of the object to be inspected, so that the image can be adapted to the gray value range of the variation model before comparing it to the variation model. Thus, also images that are taken under varying illumination can be compared. The mode for the image comparison using the operator `compare_ext_variation_model` is set to `light_dark`, so that separate regions for too bright and too dark image points are returned in `RegionDiff`.

```

get_grayval_range (ImageAffinTrans, RegionROI, RegionForegroundImage, \
                  RegionBackgroundImage, BackgroundImage, \
                  ForegroundImage)
Mult := (ForegroundGVModel - BackgroundGVModel) / (ForegroundImage - \
            BackgroundImage)
Add := ForegroundGVModel - Mult*ForegroundImage
scale_image (ImageReduced, ImageScaled, Mult, Add)
compare_ext_variation_model (ImageScaled, RegionDiff, VariationID, \
                            'light_dark')

```

At the end of the program, again the handle for the variation model is deleted.

```
clear_variation_model (VariationID)
```

11.4 Selecting Operators

11.4.1 Acquire Image(s)

Please refer to the [operator list for the method Image Acquisition](#) (see section 2.4 on page 17).

11.4.2 Create Variation Model

Standard:

```
create_variation_model
```

11.4.3 Align ROIs Or Images

Operators for aligning ROIs or images are described in the [Solution Guide II-B](#).

11.4.4 Train Variation Model

Standard:

```
train_variation_model
```

11.4.5 Check Model Quality

Standard:

```
get_variation_model
```

11.4.6 Prepare Variation Model

Standard:

```
prepare_variation_model
```

Advanced:

```
prepare_direct_variation_model, get_thresh_images_variation_model, sobel_amp,  
edges_image, gray_range_rect
```


11.4.7 Clear Training Data

Standard:

```
clear_train_data_variation_model
```

11.4.8 Compare Variation Model

Standard:

```
compare_variation_model
```

Advanced:

```
compare_ext_variation_model
```

11.4.9 Visualize Results

Please refer to the [operator list for the method Visualization](#) (see section [19.4](#) on page [297](#)).

11.4.10 Destroy Variation Model

Standard:

```
clear_variation_model
```

Chapter 12

Classification

Classification is the technical term for the assignment of objects to individual instances of a set of classes. The objects as well as the available classes are described by specific features, e.g., the color of a pixel or the shape of a region. To define the classes, the features have to be specified, e.g., by a training that is based on known objects. After the training, the classifier compares the features of the object with the features associated to the available classes and returns the class with the largest correspondence. Depending on the selected classifier, possibly additional information about the probabilities of the classes or the confidence of the classification is given.

Generally, two approaches for a classification of image data can be distinguished. One approach is based on a pure pixel classification and segments images based on color or texture. The other approach is more general and classifies arbitrary features, i.e., you can additionally classify regions based on region features like, e.g., shape, size, or color. For the programming examples introduced here, the focus is on the first approach. Note that actually the optical character recognition (OCR) provided by HALCON is a further classification approach, for which specific operators are provided. But these are described in more detail in the chapter [OCR](#) on page [245](#).


HALCON provides different classifiers. The most important classifiers are the neural network (multi-layer perceptron or MLP) classifier, a classifier that is based on support vector machines (SVM), and a classifier that is based on Gaussian mixture models (GMM). Additionally, a box classifier is available, but as the GMM classifier leads to comparable results and is more robust, here only the GMM classifier is described. Further, more 'simple' classifiers can be used for image segmentation. These comprise the classifiers for two-dimensional pixel classification with `class_2dim_sup` or `class_2dim_unsup`, and the n-dimensional pixel classification approaches based on hyper-cuboids (`class_ndim_box`), which is related to the box classifier, or hyper-spheres (`class_ndim_norm`), which can be used, e.g., for an euclidean classification. These approaches are less flexible, so the focus of this chapter is on the MLP, SVM, and GMM classifiers. If you want to use one of the 'simple' classifiers, have a look at the corresponding examples.

To decide which classifier to use for a specific task and to choose suitable parameters for the selected classification is rather challenging, as each classification task is different and thus needs different handling. HALCON provides the operators to apply a classification, and this chapter provides some general hints about their usage. Further, the most salient advantages and disadvantages of the different classifiers

are summarized for the step Create Classifier and minor differences between the classifiers, which are related to individual operators, are stated with the corresponding step descriptions or in the corresponding sections of the Reference Manual. Nevertheless, in contrast to other methods described in this manual, for classification no set of fixed rules for the selection of the suitable classifier and the selection of the parameters can be given. In many cases, you have to try out different classifiers for your specific task and with your specific training data. Independent on the selected approach, in almost any case you have to play with the parameters to get a satisfying result. Allow extra time for tests and do not despair if your first tests do not immediately lead to a satisfying result. Classification is complex!

12.1 Basic Concept

Classification consists of the following basic steps:



12.1.1 Acquire Image(s)

When classifying images, both for the generation of the training data and for the classification images must be acquired.

For detailed information see the [description of this method](#) on page 13.

12.1.2 Create Classifier

The first step of a classification is the creation of a new classifier. Here, you have to decide which classifier to apply for your specific classification task. The main advantages and disadvantages of the three classifiers (assuming an optimal tuning of the parameters) are as follows:

- **Gaussian Mixture Models:** The advantage of the GMM classification is that, controlled by the parameter settings, a feature vector that does not match to one of the trained classes can be assigned to a rejection class. Additionally, you can apply a second training that is appended to the first training, e.g., in order to add new training samples. The disadvantage of the GMM classifier is that the recognition rates are not as good as the recognition rates obtained by the MLP or SVM approaches. Further, for the feature vector a length of up to 15 features is recommended, whereas for the MLP and SVM classifiers a feature vector with 500 features is still realistic.
- **Multi-Layer Perceptron:** The MLP classifier leads to good recognition rates and is rather fast at classification. In exchange, the training is not as fast as for the SVM classification, especially for huge training sets. If the classification is time critical but the training can be applied offline, the MLP approach is a good choice. A rejection class is returned, but in comparison to the GMM classifier, it may be influenced by outliers, so that an additional, explicit training of a rejection class is recommended. Additionally, if you want to add additional training samples, you should not append a second training, but repeat the training with both the old and the new training samples.
- **Support Vector Machines:** Compared to the MLP classifier, the SVM classifier leads to slightly better recognition rates and is faster at training, especially for huge training sets. Additionally, a training of new training samples can be simply appended to a previous training. In exchange, the classification is not as fast as for the MLP approach. A rejection class is not obtained automatically.

Dependent on the selected classification approach, you create the classifier using `create_class_gmm`, `create_class_mlp`, or `create_class_svm`.

When creating a classifier, the correct parameter settings are essential for the success of the later classification. These are very specific for the different approaches. To name just a few, the parameter `OutputFunction` should be set to 'softmax' for an MLP classifier, as the classifier can also be used for regression and not only for classification. Additionally, the value for the parameter `NumHidden` has to be adjusted very carefully. In contrast to that, for the SVM approach, the parameter `KernelType` in most cases should be set to 'rbf' and the values of the parameters `Nu` and `KernelParam` strongly influence the classification result and thus should be carefully adjusted. Further, different modes and specific algorithms for the preprocessing of the feature vector can be selected. For details, see the corresponding operator descriptions in the Reference Manual and the [Solution Guide II-D](#).

12.1.3 Train Classifier

The training consists of two important steps: First, representative training samples, i.e., a set of feature vectors for each class, are added to the classifier using the operator `add_sample_class_gmm`, `add_sample_class_mlp`, or `add_sample_class_svm`. If you want to apply a pixel classification, i.e., you want to segment images into regions of different color or texture, you can alternatively add the samples using the operators `add_samples_image_class_gmm`, `add_samples_image_class_mlp`, or `add_samples_image_class_svm`. Then, you can immediately insert a single multi-channel image instead of many feature vectors (one per pixel).

In a second step, the classifier is trained by `train_class_gmm`, `train_class_mlp`, or `train_class_svm`. Note that complex calculations are performed during the training. Thus, depending on the number of training samples and some approach-specific influences, e.g., the size of the MLP used for an MLP classification, the training can take from a few seconds to several hours.

12.1.4 Classify Data

For the general classification you use the operators `classify_class_gmm`, `classify_class_mlp`, or `classify_class_svm`. These use the trained classifier to return the classes that are most probable for the feature vector of the data to classify. In most cases, you need only the best class. Then, you set the parameter `Num` to `'1'`. In a few cases, e.g., when working with overlapping classes, also the second best class may be of interest (`Num='2'`).

Like for the step of adding samples to the training, for a pixel classification, i.e., if an image is to be segmented into regions of different color or texture classes, special operators are provided. For image segmentation, you can use `classify_image_class_gmm`, `classify_image_class_mlp`, or `classify_image_class_svm`. Again, these take images as input, instead of individual feature vectors for each pixel, and return regions as output. Note that for images with a maximum of three channels, you can speed up the image segmentation by using a classification that is based on look-up tables (LUT-accelerated classification). For further information, please refer to the Solution Guide II-D in [section 6.1.6](#) on page 74.

Besides the most probable classes, for the GMM classification the probabilities of the classes, and for the MLP classification the confidence of the classification are returned. Note that the probabilities for the GMM classification are relatively precise, whereas for the confidence of the MLP classification outliers are possible because of the way an MLP is calculated. For example, when classifying positions of objects (rows and columns build the feature vectors), the confidence may be high for positions that are far from the training samples and low in case of overlapping classes.

For the SVM classification, no information about the confidence of the classification or the probabilities of the returned classes are given.

12.1.5 Destroy Classifier

When you no longer need the classifier, you destroy it with the operator `clear_class_gmm`, `clear_class_mlp`, or `clear_class_svm`, dependent on the used classifier. When working with several classifiers of the same type, you can destroy them all in one step using `clear_all_class_gmm`, `clear_all_class_mlp`, or `clear_all_class_svm`.

12.1.6 A First Example

An example for this basic concept is the following program, which segments the image depicted in [figure 12.1](#) into regions of the three classes Sea, Deck, and Walls using a GMM pixel classification. Parts of the image that can not be clearly assigned to one of the specified classes are returned as a rejection class region.

First, the image is read and the sample regions for the classes Sea, Deck, and Walls are defined and listed in the tuple `Classes`.




Figure 12.1: Segmentation of a color image: (a) original image, (b) sample regions for the classes Sea (light gray), Deck (white), and Walls (gray), (c) regions returned by the classification (black marks the rejected regions).

```
read_image (Image, 'patras')
gen_rectangle1 (Sea, 10, 10, 120, 270)
gen_rectangle2 (Deck, [170,400], [350,375], [-0.56192,-0.75139], [64,104], \
[26,11])
union1 (Deck, Deck)
gen_rectangle1 (Walls, 355, 623, 420, 702)
concat_obj (Sea, Deck, Classes)
concat_obj (Classes, Walls, Classes)
```

Then, a new classifier of type Gaussian Mixture Model is created with `create_class_gmm`. The classifier works for a feature vector with three dimensions (the three channels of the color image) and for three classes (Sea, Deck, and Walls). The tuple containing the regions that specify each class to be trained is assigned as a set of samples to the classifier using the operator `add_samples_image_class_gmm`. The samples are used now by the operator `train_class_gmm` to train the classifier.

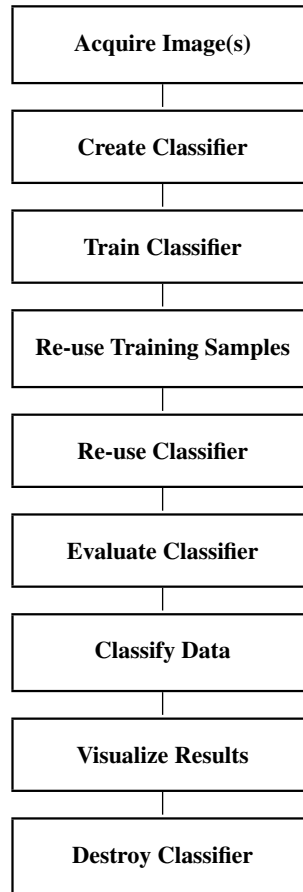
```
create_class_gmm (3, 3, [1,10], 'full', 'none', 2, 42, GMMHandle)
add_samples_image_class_gmm (Image, Classes, GMMHandle, 2.0)
train_class_gmm (GMMHandle, 500, 1e-4, 'uniform', 1e-4, Centers, Iter)
```

Then, the whole image is segmented by `classify_image_class_gmm` into regions of the three trained classes and a rejection class, which consists of the parts that can not be assigned clearly to one of the trained classes. A rejection class is returned automatically for GMM and MLP classification. For the MLP approach the rejection class can be influenced by outliers, so it should be created explicitly like shown in the example `hdevelop/Segmentation/Classification/classify_image_class_mlp.hdev`. At the end of the program, the classifier is destroyed with `clear_class_gmm`.

```
classify_image_class_gmm (Image, ClassRegions, GMMHandle, 0.0001)
clear_class_gmm (GMMHandle)
```

12.2 Extended Concept

When we look more closely at Classification, further operations can be applied. For example, the training samples or the trained classifier can be stored to file to be re-used in a later step, or the trained classifier can be evaluated further.



12.2.1 Train Classifier

If a classification does not lead to a satisfying result, one way to enhance the recognition rate is to add further training samples to the classifier. For the GMM and the SVM approach, you can add new samples to the classifier and train the classifier again, so that only those classes are trained again for which new training samples were added (for details, see the corresponding sections in the Reference Manual). But if the run time of your application is not critical, we recommend to train the classifier anew, using both the old and the new training samples. For the MLP approach a new training should be applied in any case, as the result of appending a second training most likely is not satisfying.

You may want to access the individual training samples if you have the suspicion that one of them is classified incorrectly. To do so, you apply the operators `get_sample_class_gmm`, `get_sample_class_mlp`, or `get_sample_class_svm`. To get the number of available training samples, the operators `get_sample_num_class_gmm`, `get_sample_num_class_mlp`, and `get_sample_num_class_svm` can be applied.

If the training and the classification is applied in the same process, you can save memory by applying the operators `clear_samples_class_gmm`, `clear_samples_class_mlp`, or `clear_samples_class_svm` after the training. These clear the training samples from memory. This step is only recommended if you do not want to access the individual samples in a later step. If the training process and the actual classification is separated (see the step Re-use Classifier), the clearing of the samples is not necessary as the training samples are not stored anyway (unless you store them explicitly, see the step Re-use Training Samples).

For the SVM approach, you can use `reduce_class_svm` to reduce the number of support vectors after the training. The returned classifier is faster than the originally used classifier.

12.2.2 Re-use Training Samples

In many cases it is necessary to access the samples used for the training in a later step, e.g., if you want to repeat the training with some additional samples or if you want to access individual samples to check their correctness. In most applications, you will separate the offline training and the online classification. Thus, in the classification phase the information about the individual samples is not implicitly available anymore. To nevertheless access the information, we recommend to write the used samples to file during the training phase with the operators `write_samples_class_gmm`, `write_samples_class_mlp`, or `write_samples_class_svm`. To read them from file again, you use the operators `read_samples_class_gmm`, `read_samples_class_mlp`, and `read_samples_class_svm`.

12.2.3 Re-use Classifier

In most applications, you will separate the offline training and the online classification. At the end of the training phase, you typically save the classifier to disk by `write_class_gmm`, `write_class_mlp`, or `write_class_svm`. To re-use it for the classification phase, you read the stored classifier by `read_class_gmm`, `read_class_mlp`, or `read_class_svm`, respectively.

If the training and the actual classification are separated, in the classification phase no information about the creation and training of the classifier is present anymore. To nevertheless get information about the parameters used for the creation of the classifier, the operators `get_params_class_gmm`, `get_params_class_mlp`, or `get_params_class_svm` can be applied. If a preprocessing was used for the feature vectors, which was specified with the creation of the classifier, you can also query the information content of the preprocessed feature vectors of the training samples. For this, you use `get_prep_info_class_gmm`, `get_prep_info_class_mlp`, or `get_prep_info_class_svm`.

12.2.4 Evaluate Classifier

After training a classifier, a feature vector can be evaluated. If you need only its most probable class, or perhaps also the second best class, and the related probabilities or confidences, respectively, these

are returned also when applying the actual classification (see Classify Data). But if all probabilities are needed, e.g., to get the distribution of the probabilities for a specific class (see the example description for `hdevelop/Classification/Neural-Nets/class_overlap.hdev` in the Programming Examples section), an evaluation is necessary. For the GMM approach, the evaluation with `evaluate_class_gmm` returns three different probability values: the a-posteriori probability (stored in the variable `ClassProb`), the probability density, and the k-sigma error ellipsoid. For the MLP approach, the result of the evaluation of a feature vector with `evaluate_class_mlp` is stored in the variable `Result`. The values of `ClassProb` and `Result` both can be interpreted as probabilities of the classes. The position of the maximum value corresponds to the class of the feature vector, and the corresponding value is the probability of the class. For the SVM classification, no information about the confidence of the classification or the probabilities of the returned classes are given.

12.2.5 Visualize Results

Finally, you might want to display the result of the classification.

For detailed information see the [description of this method](#) on page 289.

12.3 Programming Examples

This section gives a brief introduction to using HALCON for classification.

12.3.1 Inspection of Plastic Meshes via Texture Classification

Example: `hdevelop/Segmentation/Classification/novelty_detection_svm.hdev`

This example shows how to apply an SVM texture classification in order to detect deviations from the training data (novelties), in this case defects, in a textured material. The image depicted in [figure 12.2](#) shows one of five images used to train the texture of a plastic mesh and the result of the classification for a mesh with defects.

The program reads the first training image, defines a region of interest to avoid problems at the image border that would occur because of the later used texture filters and the specific structure of the mesh, and then creates the SVM classifier `SVMHandle` using `create_class_svm`. There, the mode `'novelty-detection'` is selected. With this mode, feature vectors that are not contained in the training samples are classified as a new class. Thus, novelties in the texture of the plastic mesh can be detected. For novelty detection, always the kernel type `'rbf'` has to be selected, but also for other modes it is recommended in most cases.

```
read_image (Image, 'plastic_mesh/plastic_mesh_01')
gen_rectangle1 (Rectangle, 10, 10, Height/2-11, Width/2-11)
create_class_svm (5, 'rbf', 0.01, 0.0005, 1, 'novelty-detection', \
                  'normalization', 5, SVMHandle)
```

Then, all five training images are read in a loop, zoomed, and transformed into texture images by the procedure `gen_texture_image`.




Figure 12.2: Inspection of a plastic mesh: (a) (zoomed) training image with correct texture (region of interest marked in white), (b) mesh with defects (texture novelties marked in white).

```
for J := 1 to 5 by 1
  read_image (Image, 'plastic_mesh/plastic_mesh_'+J$'02')
  zoom_image_factor (Image, ImageZoomed, 0.5, 0.5, 'constant')
  gen_texture_image (ImageZoomed, ImageTexture)
```

The procedure creates a texture image as follows: Five different texture filters (`texture_laws` with different filter types) are applied to the zoomed image and the resulting five images are used to build a five-channel image (`compose5`). Additionally, the five-channel image is smoothed.

```
texture_laws (Image, ImageEL, 'el', 5, 5)
texture_laws (Image, ImageLE, 'le', 5, 5)
texture_laws (Image, ImageES, 'es', 1, 5)
texture_laws (Image, ImageSE, 'se', 1, 5)
texture_laws (Image, ImageEE, 'ee', 2, 5)
compose5 (ImageEL, ImageLE, ImageES, ImageSE, ImageEE, ImageLaws)
smooth_image (ImageLaws, ImageTexture, 'gauss', 5)
```

After the creation of each texture image, it is added as training sample to the classifier by `add_samples_image_class_svm`.

```
add_samples_image_class_svm (ImageTexture, Rectangle, SVMHandle)
endfor
```

After all training samples were added to the classifier, the training is applied by `train_class_svm`. To enhance the speed of the classification, the number of support vectors is reduced by `reduce_class_svm`, which leads to a new classifier named `SVMHandleReduced`.

```
train_class_svm (SVMHandle, 0.001, 'default')
reduce_class_svm (SVMHandle, 'bottom_up', 2, 0.001, SVMHandleReduced)
```

Now, the images that have to be inspected are read in a loop. Each image is again zoomed and transformed into a texture image. The image is reduced to the part described by the region of interest and the reduced image is classified using the operator `classify_image_class_svm` with the reduced classifier `SVMHandleReduced`.

```
for J := 1 to 14 by 1
    read_image (Image, 'plastic_mesh/plastic_mesh_'+J$'02')
    zoom_image_factor (Image, ImageZoomed, 0.5, 0.5, 'constant')
    gen_texture_image (ImageZoomed, ImageTexture)
    reduce_domain (ImageTexture, Rectangle, ImageTextureReduced)
    classify_image_class_svm (ImageTextureReduced, Errors, SVMHandleReduced)
```

The regions `Errors` that are returned by the classification and which describe the classes with the novelties, can now be visualized in the image. For that, morphological operators are applied to smooth the borders of the error regions, and connected components are extracted. Regions of a certain size are selected and the result of the texture inspection is visualized as text. In particular, the texture of the mesh is classified as being good, if the number of selected regions (`NumErrors`) is 0. Otherwise, the mesh is classified as being erroneous.

```
opening_circle (Errors, ErrorsOpening, 3.5)
closing_circle (ErrorsOpening, ErrorsClosing, 10.5)
connection (ErrorsClosing, ErrorsConnected)
select_shape (ErrorsConnected, FinalErrors, 'area', 'and', 300, 1000000)
count_obj (FinalErrors, NumErrors)
if (NumErrors > 0)
    disp_message (WindowHandle, 'Mesh not OK', 'window', -1, -1, \
        'black', 'true')
else
    disp_message (WindowHandle, 'Mesh OK', 'window', -1, -1, 'black', \
        'true')
endif
if (J < 14)
    disp_continue_message (WindowHandle, 'black', 'true')
endif
endfor
```

At the end of the program, both classifiers are destroyed.

```
clear_class_svm (SVMHandle)
clear_class_svm (SVMHandleReduced)
```

12.3.2 Classification with Overlapping Classes

Example: `hdevelop/Classification/Neural-Nets/class_overlap.hdev`

This artificial example shows how an MLP classification behaves for classes with overlapping feature vectors. Three overlapping classes are defined by 2D feature vectors that consist of the rows and columns of region points. The classes are depicted in [figure 12.3](#).




Figure 12.3: Three overlapping classes (feature vectors consist of rows and columns): (a) class1, (b) class2, (c) class3.

The classes are created as follows: For each class, an elliptic region is generated. Then, noise is added to a generic image, to which a threshold is applied to get three randomly distributed regions consisting of small dots. These regions are intersected with the elliptic regions of the individual classes. The result is an elliptic cluster of small dots for each class.

```
gen_ellipse (RegionClass1, 60, 60, rad(-45), 60, 40)
gen_ellipse (RegionClass2, 70, 130, rad(-145), 70, 30)
gen_ellipse (RegionClass3, 140, 100, rad(100), 55, 40)
gen_image_const (Image, 'byte', 200, 200)
add_noise_white (Image, ImageNoise1, 60)
add_noise_white (Image, ImageNoise2, 60)
add_noise_white (Image, ImageNoise3, 60)
threshold (ImageNoise1, RegionNoise1, 40, 255)
threshold (ImageNoise2, RegionNoise2, 40, 255)
threshold (ImageNoise3, RegionNoise3, 40, 255)
intersection (RegionClass1, RegionNoise1, SamplesClass1)
intersection (RegionClass2, RegionNoise2, SamplesClass2)
intersection (RegionClass3, RegionNoise3, SamplesClass3)
```

For each elliptic cluster, the positions of the region points are obtained by `get_region_points`. Each position builds a feature vector (`Rows[j]`, `Cols[j]`) that is added to the training by `add_sample_class_mlp`. Then, the classifier is trained with `train_class_mlp`.

```

concat_obj (SamplesClass1, SamplesClass2, Samples)
concat_obj (Samples, SamplesClass3, Samples)
create_class_mlp (2, 5, 3, 'softmax', 'normalization', 1, 42, MLPHandle)
for Class := 0 to 2 by 1
    select_obj (Samples, SamplesClass, Class+1)
    get_region_points (SamplesClass, Rows, Cols)
    for J := 0 to |Rows|-1 by 1
        add_sample_class_mlp (MLPHandle, real([Rows[J],Cols[J]]), Class)
    endfor
endfor
train_class_mlp (MLPHandle, 300, 0.01, 0.01, Error, ErrorLog)

```

After the training, the classifier is evaluated. To visualize the probabilities of the classes, we create a probability image for each class. In a loop, the operator `evaluate_class_mlp` returns the probabilities for each pixel position. The first value of the returned tuple shows the probability of the pixel to belong to class 1, the second value shows the probability of the pixel to belong to class 2, etc. The probability value of each class is then used as gray value for the corresponding probability image. Additionally, the result of the classification that is returned by `classify_class_mlp`, i.e., the most probable class for each position, is visualized in a probability image. The probability images showing the probabilities of the individual classes are depicted in [figure 12.4](#). In [figure 12.5a](#) the probabilities of the three classes are united into a three-channel (color) image with `compose3`. [Figure 12.5b](#) shows the result of the classification, i.e., the most probable class for each position in the image.

```

gen_image_const (ProbClass1, 'real', 200, 200)
gen_image_const (ProbClass2, 'real', 200, 200)
gen_image_const (ProbClass3, 'real', 200, 200)
gen_image_const (LabelClass, 'byte', 200, 200)
for R := 0 to 199 by 1
    for C := 0 to 199 by 1
        Features := real([R,C])
        evaluate_class_mlp (MLPHandle, Features, Prob)
        classify_class_mlp (MLPHandle, Features, 1, Class, Confidence)
        set_grayval (ProbClass1, R, C, Prob[0])
        set_grayval (ProbClass2, R, C, Prob[1])
        set_grayval (ProbClass3, R, C, Prob[2])
        set_grayval (LabelClass, R, C, Class)
    endfor
endfor
label_to_region (LabelClass, Classes)
compose3 (ProbClass1, ProbClass2, ProbClass3, Probs)
clear_class_mlp (MLPHandle)

```

Similar examples, but with GMM and SVM classifiers instead of an MLP classifier are `hdevelop/Classification/Gaussian-Mixture-Models/class_overlap_gmm.hdev` and `hdevelop/Classification/Support-Vector-Machines/class_overlap_svm.hdev`. For the GMM example, additionally the k-sigma probability is calculated for each feature vector. This can be used to reject all feature vectors that do not belong to one of the three classes. For the SVM example, the importance of selecting suitable parameter values is shown. For SVM classification, especially the parameter value `Nu`, which specifies the ratio of outliers in the training data set, has to be selected




Figure 12.4: Probabilities for the overlapping classes ("white"=most probable): (a) probability of class 1 (b) probability of class 2, (c) probability of class 3.




Figure 12.5: Probabilities for each pixel: (a) probabilities of the individual classes composed to a three-channel image, (b) confidences returned by the classification.

sufficiently (see [figure 12.6](#)).




Figure 12.6: For SVM classification, the value for the parameter ν has to be selected carefully: (a) $\nu = 0.01$ (too small, wrong classification), (b) $\nu = 0.18$ (still too small, overfitting), (c) $\nu = 0.25$ (good classification).

12.4 Selecting Operators

12.4.1 Acquire Image(s)

Please refer to the [operator list for the method Image Acquisition](#) (see section 2.4 on page 17).

12.4.2 Create Classifier

Standard:

```
create_class_gmm, create_class_mlp, create_class_svm
```

Advanced:

```
create_class_lut_gmm, create_class_lut_mlp, create_class_lut_svm
```

12.4.3 Train Classifier

Standard:

```
add_sample_class_gmm, add_sample_class_mlp, add_sample_class_svm,  
add_samples_image_class_gmm, add_samples_image_class_mlp,  
add_samples_image_class_svm, train_class_gmm, train_class_mlp, train_class_svm
```

Advanced:

```
clear_samples_class_gmm, clear_samples_class_mlp, clear_samples_class_svm,  
get_sample_num_class_gmm, get_sample_num_class_mlp, get_sample_num_class_svm,  
get_sample_class_gmm, get_sample_class_mlp, get_sample_class_svm,  
reduce_class_svm
```

12.4.4 Re-use Training Samples

Standard:

```
write_samples_class_gmm, write_samples_class_mlp, write_samples_class_svm,  
read_samples_class_gmm, read_samples_class_mlp, read_samples_class_svm
```

12.4.5 Re-use Classifier

Standard:

```
write_class_gmm, write_class_mlp, write_class_svm, read_class_gmm, read_class_mlp,  
read_class_svm
```

Advanced:

```
get_params_class_gmm, get_params_class_mlp, get_params_class_svm,  
get_prep_info_class_gmm, get_prep_info_class_mlp, get_prep_info_class_svm
```

12.4.6 Evaluate Classifier

Standard:

```
evaluate_class_gmm, evaluate_class_mlp
```

12.4.7 Classify Data

Standard:

```
classify_class_gmm, classify_class_mlp, classify_class_svm,  
classify_image_class_gmm, classify_image_class_mlp, classify_image_class_svm
```

Advanced:

```
classify_image_class_lut
```

12.4.8 Visualize Results

Please refer to the [operator list for the method Visualization](#) (see section 19.4 on page 297).

12.4.9 Destroy Classifier

Standard:

```
clear_class_gmm, clear_class_mlp, clear_class_svm, clear_all_class_gmm,  
clear_all_class_mlp, clear_all_class_svm
```

Advanced:

```
clear_class_lut
```

12.5 Relation to Other Methods

12.5.1 Methods that are Useful for Classification

Filtering

When applying a pixel classification, a preprocessing, in particular a filtering of the image may be convenient to minimize problems like noise, texture, or overlaid structures. Operators like `mean_image` or `gauss_image` can be used to eliminate noise. A fast but slightly less perfect alternative to `gauss_image` is `binomial_filter`. The operator `median_image` is helpful for suppressing small spots or thin lines and gray value morphology can be used to eliminate noise structures.

Region Of Interest (see [description](#) on page 19)

The concept of domains (the HALCON term for a region of interest) is useful for classification. Domains are used, e.g., to restrict the classification to the image part that has to be classified. For an overview on how to construct regions of interest and how to combine them with the image see the method [Region Of Interest](#) on page 19.

Texture Analysis (see [description](#) on page 189)

Texture analysis is a method for finding regular or irregular structures, e.g., on the surface of an object. For this not only the single gray values but also a larger pixel neighborhood is used. HALCON provides filters that emphasize or suppress specific textures. If your application includes objects with a textured surface it might be helpful to apply texture analysis first before segmenting objects using pixel classifiers.

12.5.2 Methods that are Using Classification

Color Processing (see [description](#) on page 173)

Classification is used for color processing. There, the feature vector for a pixel consists of the gray values connected to each channel of the image. The classification of the pixels due to their color can then be used, e.g., to segment an image into regions of different color classes. With the description of the color processing method you find the example `solution_guide/basics/color_pieces.hdev` which uses an MLP classification to check the completeness of colored game pieces.

OCR (see [description](#) on page 245)

Classification is used for optical character recognition (OCR). HALCON provides OCR-specific operators that classify previously segmented regions of images into character classes. The OCR-specific operators are available for the MLP, SVM, and box classifier. The latter is not recommended anymore as MLP and SVM are more powerful.

12.5.3 Alternatives to Classification

Blob Analysis (see [description](#) on page 31)

Classification is time consuming and rather complex to handle. Thus, it should be applied mainly if other methods fail. One method that may also solve a classification problem is blob analysis. For example, scratches can be separated from a surface by a classification, but in most cases a comparable result can be obtained also by a segmentation of the image via blob analysis, which needs significantly less time and is easy to apply.

Matching (see [description](#) on page 97)

If 2D shapes with fixed outlines are to be classified, in many cases it is more convenient to apply a template matching instead of a classification. A classification may be used if not only the shape but also the color or texture of the object is needed to distinguish the objects of different classes.

Variation Model (see [description](#) on page 143)

If differences between images and a reference image are searched for, the variation model may be used. In contrast to, e.g., template matching, images can be directly compared by their gray values and the comparison is spatially weighted by the variation image.

12.6 Tips & Tricks

12.6.1 OCR for General Classification

The OCR-specific classification operators are developed for optical character recognition but are not limited to characters. Actually, assuming a suitable training, you can classify any shape of a region with OCR. In some cases, it may be more convenient to classify the shapes of regions using the OCR-specific operators instead of the classification operators described here.

12.7 Advanced Topics

12.7.1 Selection of Training Samples

The quality of a classification depends not only on the selected parameters, but also on the set of training samples used to train the classifier. Thus, for a good classification result, you should carefully select your training samples:

- Use as many training samples as possible.
- Use approximately the same number of samples for each class.
- Take care that the samples of a class show some variations. If you do not have enough samples with variations, you can artificially create more of them by slightly changing your available samples.

Chapter 13

Color Processing

The idea of color processing is to take advantage of the additional information encoded in color or multi-spectral images. Processing color images can simplify many machine vision tasks and provide solutions to certain problems that are simply not possible in gray value images. In HALCON the following approaches in color processing can be distinguished: First, the individual channels of a color image can be processed using standard methods like blob analysis. In this approach the channels of the original image have to be decomposed first. An optional color space transformation is often helpful in order to access specific properties of a color image. Secondly, HALCON can process the color image as a whole by calling specialized operators, e.g., for pixel classification. Advanced applications of color processing include lines and edges extraction.




Figure 13.1: Simple color segmentation.

The example illustrated in [figure 13.1](#) shows how to segment blue pieces of plasticine in a color image.

13.1 Basic Concept

Simple color processing, which is using the methods of blob analysis, mainly consists of three parts:



13.1.1 Acquire Image(s)

First, an image is acquired.

For detailed information see the [description of this method](#) on page 13.

13.1.2 Decompose Channels

In order to be able to process the individual channels, RGB color images have to be split up into a red, green, and blue channel by using the operator [decompose3](#).

13.1.3 Process Image (Channels)

Depending on the application, the individual channels can be processed using standard methods described in this chapter. One of the most frequently used methods is blob analysis (see [Blob Analysis](#) on page 31).

13.1.4 A First Example

An example for this basic concept is the following program, which belongs to the example explained above.

Here, an RGB image is acquired from file. The image is split into its channels using [decompose3](#). The red and green channels are subtracted from the blue channel using [sub_image](#). The purpose of this process is to fade out pixels with high values in the other channels, leaving pure blue pixels only. Using [threshold](#), the blue pixels with a certain intensity are selected.




Figure 13.2: Color image decomposed to its red, green, and blue channels.

```
read_image (Image, 'plasticine')
decompose3 (Image, Red, Green, Blue)
sub_image (Blue, Red, RedRemoved, 1, 0)
sub_image (RedRemoved, Green, RedGreenRemoved, 1, 0)
threshold (RedGreenRemoved, BluePixels, 10, 255)
```

13.2 Extended Concept

In many cases the processing of color images will be more advanced than in the above example. Depending on the actual application, the order of the following steps may vary, or some steps may be omitted altogether.

13.2.1 Demosaick Bayer Pattern

If the acquired image is a Bayer image, it can be converted to RGB using the operator `cfa_to_rgb`. The encoding type of the Bayer pattern (the color of the first two pixels of the first row) must be known (see [figure 13.3](#)).

13.2.2 Transform Color Space

The RGB color space is not always the most appropriate starting point to process color images. If this is the case, a transformation to a different color space might be useful. HALCON supports many important color spaces. Namely, the HSV and HSI color spaces are favorable to select distinct colors independent of their intensity. Therefore, color segmentations in these color spaces are very robust under varying illumination. The `il1i2i3` color space qualifies for color classification, whereas the `cielab` color space is a close match to human perception. Typical operators for color transform are `trans_from_rgb` and `trans_to_rgb`. The operators `create_color_trans_lut` and `apply_color_trans_lut` are, however, faster for time-consuming transformations. How time-consuming a transformation is, depends on both the color space and the used hardware. The example `hdevlop/Filters/Color/color_trans_lut.hdev` shows how to use the "faster" operators.

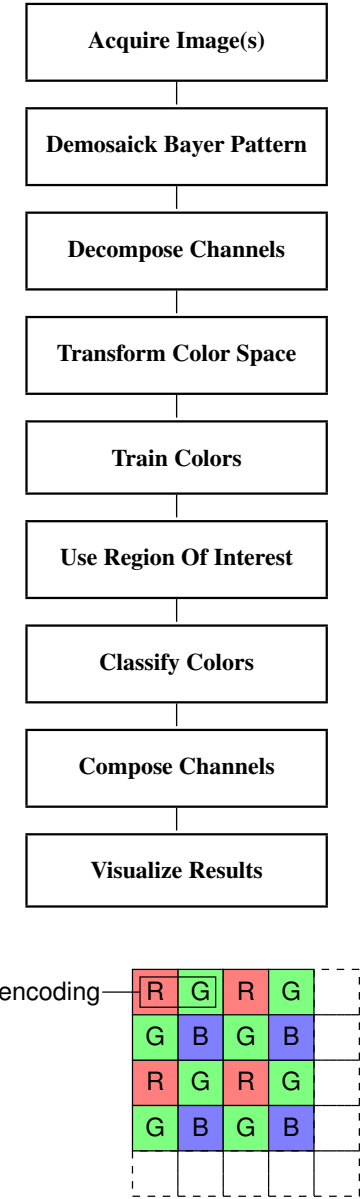


Figure 13.3: Sample Bayer pattern and corresponding encoding.

13.2.3 Train Colors

In order to do color classification the colors that need to be distinguished have to be trained. There are different approaches for full color classification including, e.g., gaussian mixture models (GMM), multilayer perceptron (MLP), and support vector machine (SVM) classification. For further information

about classification, see [Classification](#) on page 155.

13.2.4 Use Region Of Interest

Color processing can be sped up by using a region of interest. The more the region in which the segmentation is performed can be restricted, the faster and more robust it will be.

For detailed information see the [description of this method](#) on page 19.

13.2.5 Classify Colors

The colors trained in the previous step are used in subsequent images to do the actual classification.

13.2.6 Compose Channels

Any number of channels can be joined to a multi-channel image using the operators [compose2](#) through [compose7](#), or [append_channel](#). This way, channels that were processed separately can be composed back to color images for visualization purposes.

13.2.7 Visualize Results

Finally, you might want to display the images, the regions, and the features.

For detailed information see the [description of this method](#) on page 289.

13.3 Programming Examples

This section gives a brief introduction to using HALCON for color processing.

13.3.1 Robust Color Extraction

Example: `solution_guide/basics/color_simple.hdev`

The object of this example is to segment the yellow cable in a color image in a robust manner.

Here, an RGB image is acquired from file. The image is split into its channels using [decompose3](#). Afterwards, a color space transformation from RGB to HSV is performed using [trans_from_rgb](#). This transformation converts the image channels into the separate components hue, saturation and intensity. In the next steps the operator [threshold](#) selects all pixels with a high saturation value, followed by [reduce_domain](#) in the hue channel which effectively filters out pale colors and grays. A histogram of the remaining saturated (vivid) colors is displayed in [figure 13.5](#). Each peak in this histogram corresponds to a distinct color. The corresponding color band is shown below the histogram. Finally, the last [threshold](#) selects the yellowish pixels.




Figure 13.4: Segmentation of a specific color.

```
read_image (Image, 'cable' + i)
decompose3 (Image, Red, Green, Blue)
trans_from_rgb (Red, Green, Blue, Hue, Saturation, Intensity, 'hsv')
threshold (Saturation, HighSaturation, 100, 255)
reduce_domain (Hue, HighSaturation, HueHighSaturation)
threshold (HueHighSaturation, Yellow, 20, 50)
```

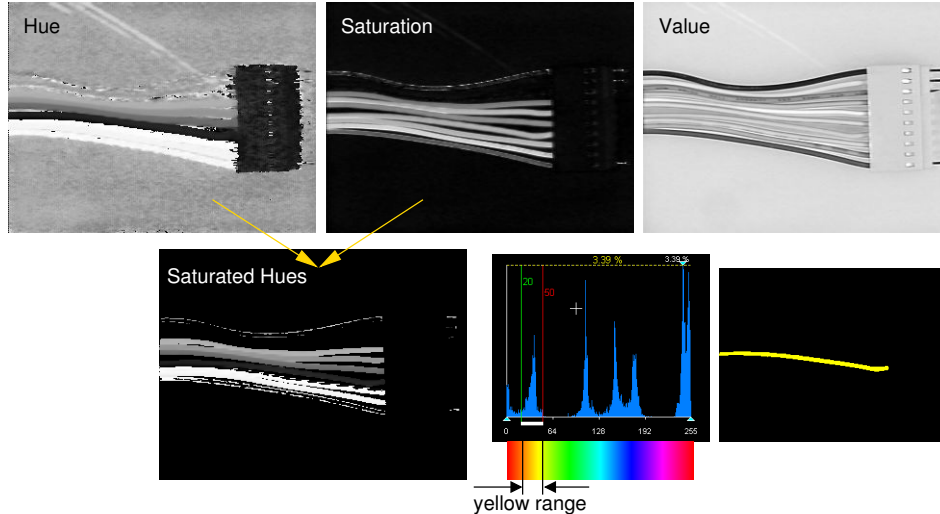


Figure 13.5: Segmentation in HSV color space.

Finding the proper threshold margins is crucial in applications like this. In HDevelop the Gray Histogram tool can be used to determine the values interactively. See the [HDevelop User's Guide](#) for more information. To generate the color band shown in [figure 13.5](#), use the following code snippet:

```

gen_image_gray_ramp (Hue, 0, 1, 128, 32, 128, 256, 64)
gen_image_proto (Hue, White, 255)
trans_to_rgb (Hue, White, White, Red, Green, Blue, 'hsv')
compose3 (Red, Green, Blue, MultiChannelImage)

```

13.3.2 Sorting Fuses

Example: `solution_guide/basics/color_fuses.hdev`



Figure 13.6: Simple classification of fuses by color with varying illumination.

In this example different types of fuses are classified using color images. The applied method is similar to the previous example. A training image has been used to specify ranges of hue for the fuse types that need to be distinguished. The determined ranges are hard-coded in the program.

```

FuseColors := ['Orange','Red','Blue','Yellow','Green']
FuseTypes := [5,10,15,20,30]
* HueRanges: Orange 10-30, Red 0-10...
HueRanges := [10,30,0,10,125,162,30,64,96,128]

```

A sequence of images is acquired from file, converted to the HSV color space, and reduced to contain only saturated colors just like in the previous example. As already mentioned, color selection in this color space is pretty stable under changing illumination. That is why the hard-coded color ranges are sufficient for a reliable classification. However, it has to be kept in mind that a certain degree of color saturation must be guaranteed for the illustrated method to work.

```

decompose3 (Image, Red, Green, Blue)
trans_from_rgb (Red, Green, Blue, Hue, Saturation, Intensity, 'hsv')
threshold (Saturation, Saturated, 60, 255)
reduce_domain (Hue, Saturated, HueSaturated)

```

The classification iterates over the fuse types and checks for sufficiently large areas in the given hue range. This is done using blob analysis. Afterwards, an additional inner loop labels the detected fuses.

```

for Fuse := 0 to |FuseTypes|-1 by 1
  threshold (HueSaturated, CurrentFuse, HueRanges[Fuse*2], \
            HueRanges[Fuse*2+1])
  connection (CurrentFuse, CurrentFuseConn)
  fill_up (CurrentFuseConn, CurrentFuseFill)
  select_shape (CurrentFuseFill, CurrentFuseSel, 'area', 'and', 6000, \
              20000)
  area_center (CurrentFuseSel, FuseArea, Row1, Column1)
  dev_set_color ('magenta')
  for i := 0 to |FuseArea|-1 by 1
    set_tposition (WH, Row1[i], Column1[i])
    write_string (WH, \
                  FuseColors[Fuse] + ' ' + FuseTypes[Fuse] + ' Ampere')
  endfor
  set_tposition (WH, 24*(Fuse+1), 12)
  dev_set_color ('slate blue')
  write_string (WH, FuseColors[Fuse] + ' Fuses: ' + |FuseArea|)
endfor
stop ()

```

13.3.3 Completeness Check of Colored Game Pieces

Example: `solution_guide/basics/color_pieces.hdev`

Completeness checks are very common in machine vision. Usually, packages assembled on a production line have to be inspected for missing items. Before this inspection can be done, the items have to be trained. In the example presented here, a package of game pieces has to be inspected. The game pieces come in three different colors, and the package should contain four of each type. The pieces themselves can be of slightly different shape, so shape-based matching is not an option. The solution to this problem is to classify the game pieces by color. The method applied here is a classification using neural nets (MLP classification).

In the training phase an image is acquired, which contains the different types of game pieces. The task is to specify sample regions for the game pieces and the background using the mouse (see [Figure 13.7a](#)). This is accomplished by looping over the `draw_rectangle1` and `gen_rectangle1` operators to draw and create the corresponding regions. The tuple `Classes`, which will be used for the actual training, is extended each time.

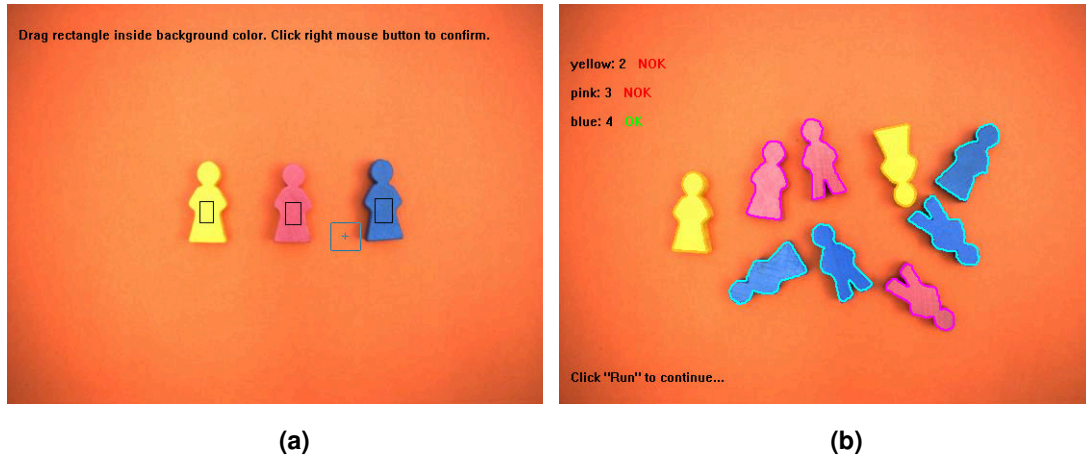


Figure 13.7: Example MLP classification: (a) Training, (b) Result.

```

read_image (Image, ImageRootName+'0')
for i := 1 to 4 by 1
    dev_display (Image)
    dev_display (Classes)
    disp_message (WindowHandle, \
        ['Drag rectangle inside ' + Regions[i-1] + ' color', \
        'Click right mouse button to confirm'], 'window', \
        24, 12, 'black', 'false')
    draw_rectangle1 (WindowHandle, Row1, Column1, Row2, Column2)
    gen_rectangle1 (Rectangle, Row1, Column1, Row2, Column2)
    concat_obj (Classes, Rectangle, Classes)
endfor

```

Once the classes are specified, a multilayer perceptron is created using `create_class_mlp`. With the operator `add_samples_image_class_mlp` the training samples from the image are added to the training data of the multilayer perceptron. The actual training is started with `train_class_mlp`. The duration of the training depends on the complexity and sizes of the training regions.

```

create_class_mlp (3, 7, 4, 'softmax', 'normalization', 3, 42, MLPHandle)
add_samples_image_class_mlp (Image, Classes, MLPHandle)
set_tposition (WindowHandle, 100, 12)
write_string (WindowHandle, 'Training...')
train_class_mlp (MLPHandle, 400, 0.5, 0.01, Error, ErrorLog)

```

After the training has finished, subsequent images are acquired and classified using `classify_image_class_mlp`. The operator returns a classified region.

```

for img := 0 to 3 by 1
    read_image (Image, ImageRootName + img)
    classify_image_class_mlp (Image, ClassRegions, MLPHandle, 0.5)
    stop ()
endfor

```

The returned result is processed further using blob analysis. Each class of the classified region (with the exception of the background class) is accessed using `copy_obj`. The regions of each class are split up using `connection` and reduced to regions of a relevant size (`select_shape`). The remaining few lines of code calculate the number of game pieces found for each class and make a decision whether the result was OK or not.

```

for figure := 1 to 3 by 1
    copy_obj (ClassRegions, ObjectsSelected, figure, 1)
    connection (ObjectsSelected, ConnectedRegions)
    select_shape (ConnectedRegions, SelectedRegions, 'area', 'and', \
        400, 99999)
    count_obj (SelectedRegions, Number)
    dev_set_color (Highlight[figure-1])
    dev_display (SelectedRegions)
    OutString := Regions[figure-1] + ': ' + Number + '    '
    dev_set_color ('green')
    disp_message (WindowHandle, OutString, 'window', 24+30*figure, \
        12, 'black', 'false')
    if (Number#4)
        disp_message (WindowHandle, 'Not OK', 'window', \
            24+30*figure, 120, 'red', 'false')
    else
        disp_message (WindowHandle, 'OK', 'window', 24+30*figure, \
            120, 'green', 'false')
    endif
endfor

```

To illustrate the advantage of using color information, and to compare the classification results, the example program runs an additional training and classification on a converted gray image. As can be seen in [figure 13.8](#) only the yellow region is detected faithfully.

```

rgb1_to_gray (Image, GrayImage)
compose3 (GrayImage, GrayImage, GrayImage, Image)

```

For more complex (and time-consuming) classifications, it is recommended to save the training data to the file system using `write_class_mlp`. Later, the saved data can be re-stored using `read_class_mlp`. As an alternative to the MLP classification you can also apply a classification based on support vector machines (SVM). The corresponding operators are `create_class_svm`, `add_samples_image_class_svm`, `train_class_svm`, `classify_image_class_svm`, `write_class_svm`, and `read_class_svm`. In order to speed up the classification itself, you can apply a segmentation based on look-up tables (for LUT-accelerated classification, please refer to the Solution Guide II-D in [section 6.1.6](#) on page 74) or use a different classification method like Euclidean classification (see `solution_guide/basics/color_pieces_euclid.hdev`).




Figure 13.8: Poor classification result when only using the gray scale image.

13.3.4 Inspect Power Supply Cables

Example: `hdevelop/Filters/Lines/lines_color.hdev`

The task of this example is to locate and inspect the power supply cables depicted in [figure 13.9](#).




Figure 13.9: (a) Original color image with cable centers extracted using the color line extractor; (b) corresponding results when using the gray value image.

The input for the program are sample images of colored power supply cables. The task is to extract the centers of each cable together with the width. This is performed using the subpixel-precise color line extractor. To remove irrelevant structures, contours that are too short are removed.

```
lines_color (Image, Lines, 3.5, 0, 12, 'true', 'false')
select_contours_xld (Lines, LongLines, 'contour_length', 450, 100000, 0, \
0)
```

The cable width is determined by accessing the line width attribute. For display purposes, a contour is generated for each side.

```
count_obj (LongLines, Number)
gen_empty_obj (EdgesL)
gen_empty_obj (EdgesR)
for K := 1 to Number by 1
    select_obj (LongLines, Line, K)
    get_contour_xld (Line, Row, Col)
    get_contour_attrib_xld (Line, 'angle', Angle)
    get_contour_attrib_xld (Line, 'width_right', WidthR)
    get_contour_attrib_xld (Line, 'width_left', WidthL)
    EdgeRR := Row+cos(Angle)*WidthR
    EdgeRC := Col+sin(Angle)*WidthR
    EdgeLR := Row-cos(Angle)*WidthL
    EdgeLC := Col-sin(Angle)*WidthL
    gen_contour_polygon_xld (EdgeR, EdgeRR, EdgeRC)
    gen_contour_polygon_xld (EdgeL, EdgeLR, EdgeLC)
    concat_obj (EdgesL, EdgeL, EdgesL)
    concat_obj (EdgesR, EdgeR, EdgesR)
endfor
```

To compare this result with the classical line extraction approach (see [Edge Extraction \(Subpixel-Precise\)](#) on page 73), a line extractor is also applied to the gray value image. The result is depicted in [figure 13.9b](#). Here, it becomes obvious how hard it is to extract the cable using the luminance only.

```
rgb1_to_gray (Image, GrayImage)
lines_gauss (GrayImage, LinesGray, 3.5, 0.0, 0.7, 'dark', 'true', \
    'bar-shaped', 'false')
```

13.3.5 Locating Board Components by Color

Example: `hdevelop/Applications/Completeness-Check/ic.hdev`

The task of this example is to locate all components on the printed circuit board depicted in [figure 13.10](#).

The input data is a color image, which allows locating components like capacitors and resistors very easily by their significant color: Using a color space transformation, the hue values allow the selection of the corresponding components. The following code extracts the resistors; the extraction of the capacitors is performed along the same lines.

```
decompose3 (Image, Red, Green, Blue)
trans_from_rgb (Red, Green, Blue, Hue, Saturation, Intensity, 'hsv')
threshold (Saturation, Colored, 100, 255)
reduce_domain (Hue, Colored, HueColored)
threshold (HueColored, Red, 10, 19)
connection (Red, RedConnect)
select_shape (RedConnect, RedLarge, 'area', 'and', 150.000000, 99999.000000)
shape_trans (RedLarge, Resistors, 'rectangle2')
```

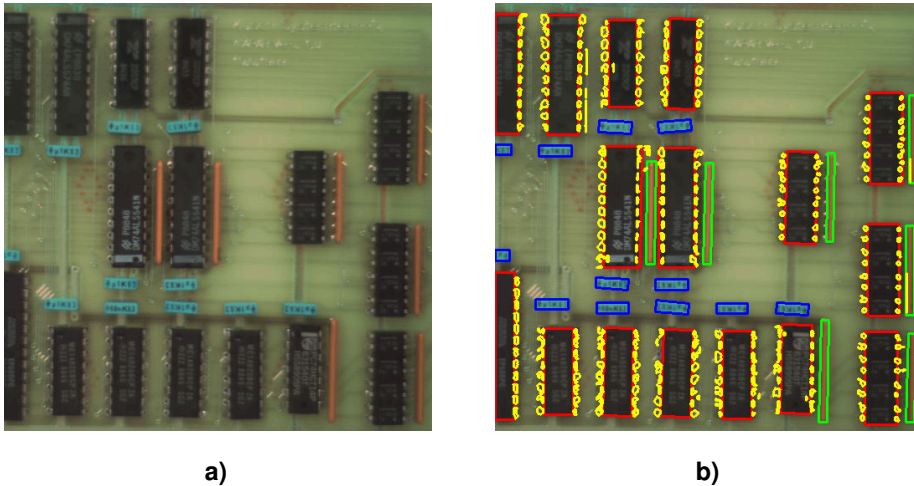



Figure 13.10: (a) Original image; (b) extracted ICs, resistors, and capacitors.

The extraction of the ICs is more difficult because of the bright imprints, which do not allow a simple thresholding in one step. Instead of this, dark areas are selected first, which are then combined using a dilation.

```
threshold (Intensity, Dark, 0, 50)
dilation_rectangle1 (Dark, DarkDilation, 14, 14)
connection (DarkDilation, ICLarge)
```

After this, the segmentation is repeated inside the thus extracted connected components.

```
add_channels (ICLarge, Intensity, ICLargeGray)
threshold (ICLargeGray, ICDark, 0, 50)
shape_trans (ICDark, IC, 'rectangle2')
```

To locate the contact points, small ROIs are generated on the left and right side of each IC.

```
dilation_rectangle1 (IC, ICDilation, 5, 1)
difference (ICDilation, IC, SearchSpace)
dilation_rectangle1 (SearchSpace, SearchSpaceDilation, 14, 1)
union1 (SearchSpaceDilation, SearchSpaceUnion)
```

Inside these areas, locally bright spots are detected.

```
reduce_domain (Intensity, SearchSpaceUnion, SearchGray)
mean_image (SearchGray, SearchMean, 15, 15)
dyn_threshold (SearchGray, SearchMean, PinsRaw, 5.000000, 'light')
connection (PinsRaw, PinsConnect)
fill_up (PinsConnect, PinsFilled)
select_shape (PinsFilled, Pins, 'area', 'and', 10, 100)
```


13.4 Selecting Operators

13.4.1 Acquire Image(s)

Please refer to the [operator list for the method Image Acquisition](#) (see section 2.4 on page 17).

13.4.2 Demosaick Bayer Pattern

Standard:

`cfa_to_rgb`

13.4.3 Decompose Channels

Standard:

`decompose3`, `access_channel`

13.4.4 Transform Color Space

Standard:

`trans_from_rgb`, `trans_to_rgb`, `create_color_trans_lut`, `apply_color_trans_lut`,
`clear_color_trans_lut`, `clear_all_color_trans_luts`

The following color spaces are supported:

- `argyb`
- `cielab`
- `ciexyz`
- `hls`
- `hsi`
- `hsv`
- `ili2i3`
- `ihs`
- `rgb`
- `yiq`
- `yuv`

13.4.5 Train Colors

Standard:

```
train_class_gmm, train_class_mlp, train_class_svm, histo_2dim, learn_ndim_norm,  
learn_ndim_box
```

13.4.6 Use Region Of Interest

Please refer to the [operator list for the method Region Of Interest](#) (see section 3.4 on page 26).

13.4.7 Process Image (Channels)

Standard:

```
smooth_image, mean_image, median_image
```

Advanced:

```
lines_color, edges_color, edges_color_sub_pix
```

Further operators can be found in the [operator list for the method Blob Analysis](#) (see section 4.4 on page 43).

13.4.8 Classify Colors

Standard:

```
classify_image_class_gmm, classify_image_class_mlp, classify_image_class_svm,  
classify_image_class_lut, class_2dim_sup, class_2dim_unsup, class_ndim_norm,  
class_ndim_box
```

13.4.9 Compose Channels

Standard:

```
compose3, append_channel
```

13.4.10 Visualize Results

Standard:

```
disp_color, disp_obj
```

Further operators can be found in the [operator list for the method Visualization](#) (see section 19.4 on page 297).

13.5 Tips & Tricks

13.5.1 Speed Up

Many online applications require maximum speed. Because of its flexibility, HALCON offers many ways to achieve this goal. Here, the most common methods are listed.

- Of the color processing approaches discussed in this section, the simplest is also the fastest (decompose color image and apply blob analysis). If you want to do color classification, you should consider using `class_ndim_norm` or `class_2dim_sup` for maximum performance.
- Regions of interest are the standard method to increase the speed by processing only those areas where objects need to be inspected. This can be done using pre-defined regions but also by an online generation of the regions of interest that depend on other objects found in the image.
- By default, HALCON performs some data consistency checks. These can be switched off using `set_check`.
- By default, HALCON initializes new images. Using `set_system` with the parameter "init_new_image", this behavior can be changed.

13.6 Advanced Topics

13.6.1 Color Edge Extraction

Similar to the operator `edges_image` for gray value images (see [Edge Extraction \(Pixel-Precise\)](#) on page 63), the operator `edges_color` can be applied to find pixel-precise edges in color images. For a subpixel-precise edge extraction in color images, the operator `edges_color_sub_pix` is provided. It corresponds to the gray value based operator `edges_sub_pix` (see [Edge Extraction \(Subpixel-Precise\)](#) on page 73). Processing the detected color edges is the same as for gray value images (see [Contour Processing](#) on page 83).

13.6.2 Color Line Extraction

Similar to color edge extraction, HALCON supports the detection of lines in color images using `lines_color`. For processing the detected lines, please refer to [Edge Extraction \(Subpixel-Precise\)](#) on page 73 and [Contour Processing](#) on page 83.

Chapter 14

Texture Analysis

The concept of texture analysis is based on the use of an object's structural features. In order to thoroughly understand the suitable application of texture analysis, it is important to first comprehend texture itself. Texture designates the structural characteristics of a surface which present themselves as gray value variations in an image. The regularity of a texture varies from absolutely regular to irregular. The basis of a regular texture is an ideal grid structure whereas irregular textures or statistical structures are based upon a random function.




Figure 14.1: a) Regular texture; b) irregular texture.

Textures are composed of so-called texels, a shortform for texture elements, which can easily be recognized in a regular texture. Irregular textures also have recurring elements, which are, however, more difficult to recognize. When working with irregular textures the measures of a texel have to be approximated if necessary. A texel is the smallest recurring structure as depicted in [figure 14.2](#).

Texture analysis is a very useful method for all tasks which cannot simply be solved as their gray value structure is too complex. It can help you to answer two questions:

- Where is a certain texture located?
- What kind of texture is it?





Figure 14.2: Texture elements of extremely regular texture: a) original image; b) texel.

The first task is typically solved by filtering the image with filters that enhance or suppress certain textures. The filtered image is then segmented into regions with similar texture, either with standard blob analysis or by classifying each pixel.

The second task is typically solved by computing texture features for the whole image or for individual regions. Based on these features, the image or region is then classified. As a preprocessing, texture filters can be applied to enhance or suppress certain textures.

14.1 Basic Concept

A simple texture analysis consists of the following steps:



14.1.1 Acquire Image(s)

First, an image is acquired.

For detailed information see the [description of this method](#) on page 13.

14.1.2 Apply Texture Filter

With a texture filter, specific textural structures can be emphasized or suppressed. HALCON's standard texture filter is [texture_laws](#).

14.1.3 Compute Features

The standard HALCON operator for computing texture features is [gen_cooc_matrix](#), which generates the cooccurrence matrix of an image. Other feature related operators are [entropy_gray](#) for calculating the entropy or the anisotropy of an image or region. These calculations are computed on the original image. To determine the amount of a feature, the intensity, with a threshold, the task is performed on the image after filtering.

14.1.4 A First Example

Example: `hdevelop/Applications/Measuring-2D/vessel.hdev`

The example for this basic concept clarifies how texture analysis allows the segmentation of textured images which are impossible to segment using classical segmentation methods like thresholding.




Figure 14.3: a) Original image; b) segmented vessel.

The example illustrated in [figure 14.3](#) shows how to segment a blood vessel in a gray value image using texture analysis.

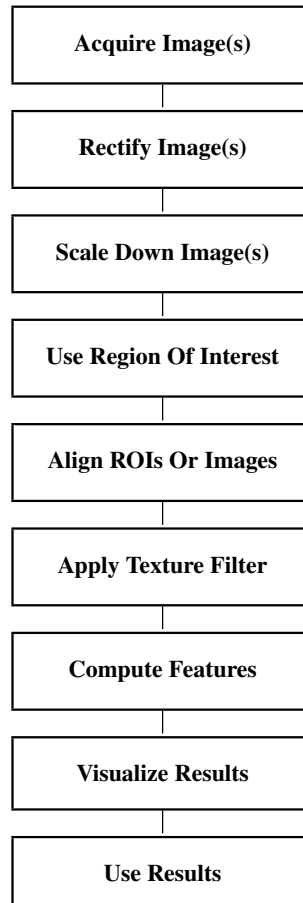
In the gray value image a), a blood vessel has to be segmented and measured. In order to perform this task, [texture_laws](#) enhances the structure of the texture by extracting structural elements of a

certain direction as well as their frequency. Then `mean_image` with the filter size of one texture element, smoothes the image. `bin_threshold` is then able to segment the vessel. Now that the vessel is segmented from its environment in image b), it can easily be measured.

```
read_image (Image, 'vessel')
texture_laws (Image, Texture, 'el', 2, 5)
mean_image (Texture, Energy, 211, 61)
bin_threshold (Energy, Vessel)
```

14.2 Extended Concept

It will, however, sometimes be necessary to perform more advanced steps during texture analysis in order to achieve a state where the image can be, depending on the kind of analysis undertaken, segmented or classified.



14.2.1 Rectify Image(s)

If the camera looks at an angle onto the object to inspect, it may be necessary to rectify the image to remove lens distortion and perspective distortion.

Detailed information about rectifying images can be found in the Solution Guide III-C in [section 3.4](#) on page 62.

14.2.2 Scale Down Image(s)

When both the texel and its smallest substructure are relatively large, you can downscale the image with the operator `zoom_image_factor` to speed up the texture analysis.

You can also create a pyramid of downsampled images by calling `zoom_image_factor` multiple times and then apply the texture filtering or feature computation to the image pyramid. Thus, texture analysis is performed at multiple scales as is shown in example `classify_wood.hdev` on page 197.

14.2.3 Use Region Of Interest

A region of interest can be created to reduce the image domain for which the texture analysis will be performed. This will reduce the processing time. In particular when computing texture features, the region of interest can be the result of a previous segmentation.

For detailed information see the [description of this method](#) on page 19.

14.2.4 Align ROIs Or Images

If you are dealing with a regular texture, you can exploit this fact by rotating the image such that the texture appears aligned to the image axes. Then, both texture filtering and texture features can be applied or computed more specifically.

How to perform alignment using shape-based matching is described in the Solution Guide II-B in [section 2.4.3.2](#) on page 42.

14.2.5 Apply Texture Filter

The operator `texture_laws` actually provides a family of texture filters. More information can be found in [section 14.8](#) on page 209. After calling `texture_laws`, you should smooth the image by using `mean_image` with the mask size set to the size of the texel.

For certain images with simple textures, the operator `dyn_threshold` can be a good alternative to a regular filter as it is faster. Before using this operator, you must call `mean_image` with the mask size set to twice the texel size.

14.2.6 Compute Features

If you want to use the cooccurrence matrix corresponding to different directions, it is more efficient to generate the matrix with `gen_cooc_matrix` and then compute the features with `cooc_feature_matrix`.

Besides the texture features based on the cooccurrence matrix, you can also use the entropy and anisotropy of an image, which are computed by the operator `entropy_gray`, respectively. After the filtering has been performed, the intensity of a feature in the image can be computed using a threshold.

All operators can compute the texture features either for the complete image or for image regions. More information about texture features can be found in [section 14.7](#) on page 203.

As an advanced texture feature, you can also compute the distribution of the image gradient, which shows how the gray values change by first filtering the image with `sobel_amp` and then computing the histogram with `gray_histo_abs`.

14.2.7 Visualize Results

To check the results of texture analysis, you might want to display the images, the regions, and the features. To visualize the cooccurrence matrix, you must compute the matrix and its features in separate steps using `gen_cooc_matrix` and `cooc_feature_matrix`.

Often, details in the texture images and the cooccurrence matrix do not show up clearly in the original gray value image. In such a case, you can try to change the look-up table (LUT).

For detailed information see the [description of this method](#) on page 289.

14.2.8 Use Results

The filtered image is typically used as input for segmentation tasks, either by [Blob Analysis](#) on page 31 or by (pixel) [Classification](#) on page 155. Described examples are `vessel.hdev` on page 191, `novelty_detection_svm.hdev`, and `novelty_detection_dyn_threshold.hdev`. The computed texture features are typically used as input for general [Classification](#) on page 155. A described example is `classify_wood.hdev` on page 197.

14.3 Programming Examples

14.3.1 Detect Defects in a Texture with Novelty Detection

Example: `hdevelop/Segmentation/Classification/novelty_detection_svm.hdev`

A detailed description of this example can be found in the chapter [Classification](#) on page 155.

The task of this example is to detect defects in a texture that do not correspond to the texture of trained good objects. For this, it uses the novelty detection mode of the SVM classifier. In the example

`novelty_detection_dyn_threshold.hdev`, which is described in the next section, the same task is solved with the operator `dyn_threshold`.

The raw errors returned by the SVM are postprocessed to remove insignificant parts of the detected errors. The texture image is generated in a procedure: It is a five-channel image that contains the result of applying five different Laws filters, which basically correspond to first and second derivatives, and smoothing them sufficiently. Laws filters, included in the HALCON operator `texture_laws`, are very useful whenever it is necessary to enhance certain structures in a texture as becomes clear in [figure 14.4](#).

```
texture_laws (Image, ImageEL, 'el', 5, 5)
texture_laws (Image, ImageLE, 'le', 5, 5)
texture_laws (Image, ImageES, 'es', 1, 5)
texture_laws (Image, ImageSE, 'se', 1, 5)
texture_laws (Image, ImageEE, 'ee', 2, 5)
compose5 (ImageEL, ImageLE, ImageES, ImageSE, ImageEE, ImageLaws)
smooth_image (ImageLaws, ImageTexture, 'gauss', 5)
```

The [figure 14.4](#) shows a plastic mesh without defects on the left side and one that is damaged on the right side. Starting with the original image, first the Laws filter `el` is applied which enhances horizontal structures, followed by the Laws filter `le`, stressing vertical structures. The damages in the filtered images d) to f) are quite clearly visible. For more information about the Laws filter read in [section 14.8](#) on page 209 later in this chapter.

14.3.2 Detect Defects in a Web Using Dynamic Thresholding

Example: `solution_guide/basics/novelty_detection_dyn_threshold.hdev`

The task of this example is also to detect defects in a texture that do not correspond to the texture of trained good objects similar to the example `novelty_detection_svm.hdev` on page 194.

The task of this example is to detect defects in a mesh using the operator `dyn_threshold`. In this way, the operator can be used to find textures that differ from the rest of the image. `dyn_threshold` is an operator that is easy to handle. Together with a smoothing filter like `mean_image`, it can extract objects that differ locally from their neighborhood. Note that the operator `dyn_threshold` can only be used successfully for texture analysis if the defects can be detected as darker or brighter than the rest of the texture.




Figure 14.4: Images a) to c) mesh without defects; images d) to f) damaged mesh.

First, the image is smoothed using the HALCON operator `mean_image`. The mesh size determines the size of the extracted objects: The larger the mask size is chosen, the larger the found regions become. As a rule of thumb, the mask size should be about twice the diameter of the objects to be extracted. Subsequently, `dyn_threshold` is performed and connected regions are looked for. The parameter 'area' of the operator `select_shape` makes it possible to find regions that differ in size, e.g., that are too small. Found errors are finally counted and displayed.

```

read_image (Image, 'plastic_mesh/plastic_mesh_'+J$'02')
mean_image (Image, ImageMean, 49, 49)
dyn_threshold (Image, ImageMean, RegionDynThresh, 5, 'dark')
connection (RegionDynThresh, ConnectedRegions)
select_shape (ConnectedRegions, ErrorRegions, 'area', 'and', 500, 99999)
count_obj (ErrorRegions, NumErrors)

```

If the number of the errors exceeds zero, the message 'Mesh not OK' is displayed. Otherwise the web is undamaged and 'Mesh OK' appears as is shown in [figure 14.5](#).

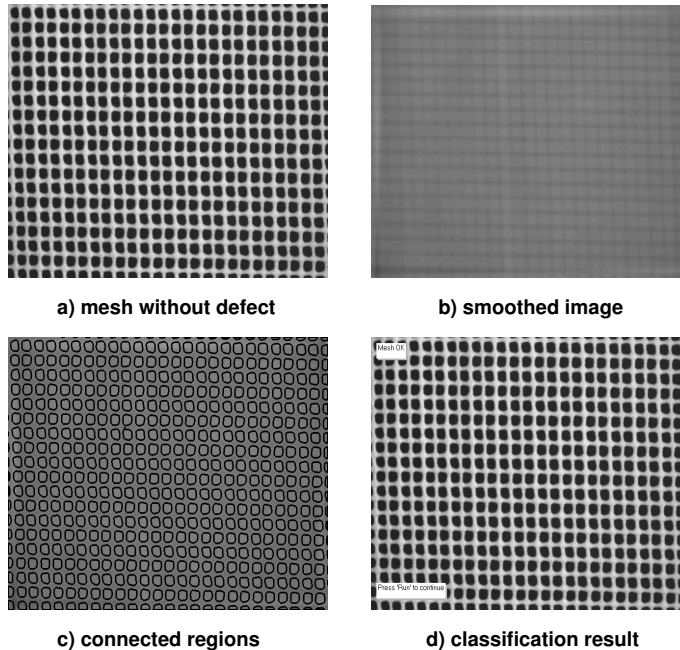


Figure 14.5: a) Faultless mesh b) smoothed image c) connected regions d) classification: Mesh OK. .

14.3.3 Classification of Different Types of Wood

Example: `solution_guide/basics/classify_wood.hdev`

The objective of this example is to classify different types of wood according to their surface texture.

First the different classes of wood, as shown in [figure 14.7](#), are specified.

```
Classes := ['apple', 'beech', 'cherry', 'maple', 'oak', 'walnut']
```

Then, the training is performed, which means that several images are read and trained for each type of wood. Therefore, the SVM learns different wood textures and compares each new wood image to the existing images. It checks which class is most similar and displays the assigned class.

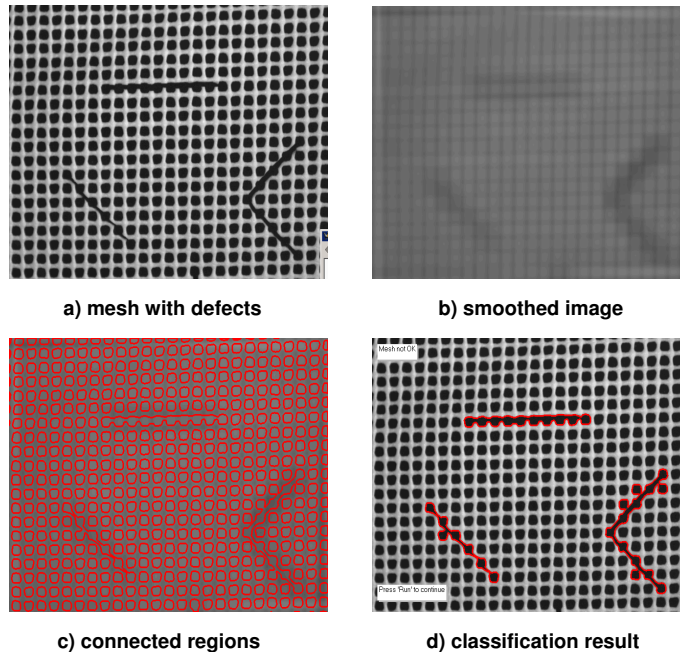


Figure 14.6: a) Defective mesh b) smoothed image c) connected regions d) classification: Mesh not OK. .

For each image, features are computed in a procedure and then passed to the operator `classify_image_class_mlp`.

```
gen_features (Image, FeatureVector)
classify_class_mlp (MLPHandle, FeatureVector, 2, FoundClassIDs, Confidence)
```

The procedure `gen_features` calls the actual feature extraction procedure and then downsamples the image and calls the second procedure with the smaller image again.

```
procedure gen_features (Image, FeatureVector)
  gen_sobel_features (Image, FeatureVector, FeatureVector)
  zoom_image_factor (Image, Zoomed1, 0.5, 0.5, 'constant')
  gen_sobel_features (Zoomed1, FeatureVector, FeatureVector)
```

The procedure `gen_sobel_features` calculates multiple texture features. First the cooccurrence matrix is computed for the directions 0 and 90 degrees.

```
procedure gen_sobel_features (Image, Features, FeaturesExtended)
  * Cooccurrence matrix for 0 deg and 90 deg:
  cooc_feature_image (Image, Image, 6, 0, Energy, Correlation, Homogeneity, \
    Contrast)
  cooc_feature_image (Image, Image, 6, 90, Energy, Correlation, Homogeneity, \
    Contrast)
```

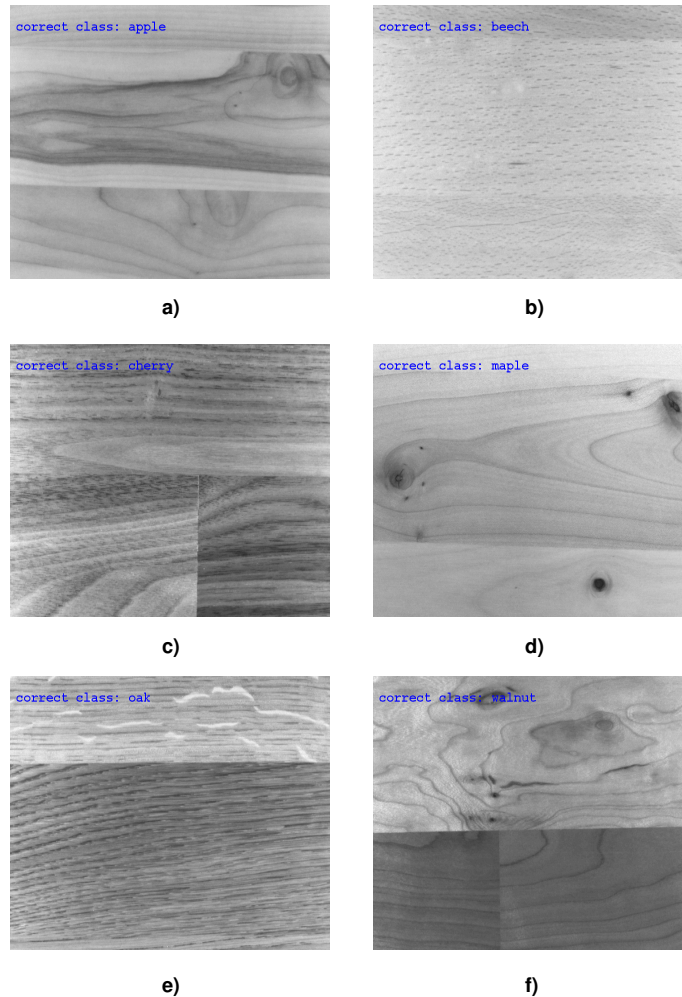


Figure 14.7: Result of wood classification: a) apple, b) beech, c) cherry, d) maple, e) oak, f) walnut .

Furthermore, within the procedure the gray value edges are extracted with the operator `sobel_amp`. The bigger the gray value difference between neighboring pixels is, the brighter these areas are shown in the resulting image. This way, the filtering enables to find and highlight structures in an image. As a feature derived from the sobel image, the absolute gray value histogram of the edge amplitude image is computed.

```
sobel_amp (Image, EdgeAmplitude, 'sum_abs', 3)
gray_histo_abs (EdgeAmplitude, EdgeAmplitude, 8, AbsoluteHistoEdgeAmplitude)
* entropy_gray (Image, Image, Entropy, Anisotropy)
* gray_histo_abs (Image, Image, 8, AbsoluteHistoImage)
```

All calculated features are appended to the input feature vector and returned in the output feature vector.


```
FeaturesExtended := [Features,Energy,Correlation,Homogeneity,Contrast]
FeaturesExtended := [FeaturesExtended,AbsoluteHistoEdgeAmplitude]
```

As already noted, the texture features are computed on multiple pyramid levels by zooming the images with a factor of 0.5 for each level. This makes it possible to analyze a bigger neighborhood. In the example, only 2 pyramid levels are used; to use more for a more complex task, you can activate the corresponding lines. Note, however, that you must then train the classifier anew.

```
zoom_image_factor (Image, Zoomed1, 0.5, 0.5, 'constant')
gen_sobel_features (Zoomed1, FeatureVector, FeatureVector)
zoom_image_factor (Zoomed1, Zoomed2, 0.5, 0.5, 'constant')
* gen_sobel_features (Zoomed2, FeatureVector, FeatureVector)
FeatureVector := real(FeatureVector)
```

Similarly, the procedure `gen_sobel_features` contains deactivated code to compute more features: the entropy, anisotropy, and the absolute gray value histogram of the image.

```
* entropy_gray (Image, Image, Entropy, Anisotropy)
* gray_histo_abs (Image, Image, 8, AbsoluteHistoImage)
* FeaturesExtended := [FeaturesExtended,Entropy,Anisotropy]
* FeaturesExtended := [FeaturesExtended,AbsoluteHistoImage]
```

Note again that if you use more features you must train the classifier anew.

14.4 Selecting Operators

14.4.1 Acquire Image(s)

Please refer to the [operator list for the method Image Acquisition](#) (see section 2.4 on page 17).

14.4.2 Rectify Image(s)

Operators for rectifying images are described in the [Solution Guide III-C](#).

14.4.3 Scale Down Image(s)

Standard:

```
zoom_image_factor
```

14.4.4 Use Region Of Interest

Please refer to the [operator list for the method Region Of Interest](#) (see section 3.4 on page 26).

14.4.5 Align ROIs Or Images

Operators for rectifying images are described in the [Solution Guide II-B](#).

14.4.6 Apply Texture Filter

Standard:

`texture_laws`, `dyn_threshold`, `mean_image`

14.4.7 Compute Features

Standard:

`cooc_feature_image`, `entropy_gray`, `anisotropic_diffusion`

Advanced:

`gen_cooc_matrix`, `cooc_feature_matrix`, `sobel_amp`, `gray_histo_abs`

14.4.8 Visualize Results

Please refer to the [operator list for the method Visualization](#) (see section 19.4 on page 297).

14.4.9 Use Results

Please refer to the following places: [detailed operator list for the step Segment Image\(s\)](#) on page 44, [operator list for the method Classification](#) (see section 12.4 on page 168).

14.5 Relation to Other Methods

14.5.1 Methods that are Using Texture Analysis

Blob Analysis (see [description](#) on page 31)

Blob analysis identifies pixels of relevant objects (also called foreground) by their gray value in an image. Whenever objects cannot simply be segmented by separating dark pixels from bright pixels, those objects can possibly be identified by their texture. Therefore, it is useful to perform a texture analysis first and then use the results to segment the objects.

Classification (see [description](#) on page 155)

Classification is the technical term for the assignment of objects to individual instances of a set of classes and can be a useful method to follow the texture analysis if the object is characterized by a particular texture. To define the classes, this texture has to be specified, e.g., by a training that is based on known objects. After the training, the classifier compares the features of the object with the features associated to the available classes and returns the class with the largest correspondence.

14.6 Advanced Topics

14.6.1 Fast Fourier Transform (FFT)

Fast Fourier transform, short FFT, is another option for texture analysis. This method makes use of the different frequencies in textures. HALCON's FFT exactly determines the amount of each frequency in the image. It enables the construction of any type of linear filters which means that application-oriented filter designs are possible. There is also a huge set of predefined filters as well as advanced Gabor texture filters. With HALCON's FFT it is possible, e.g., to emphasize any structure of a special texture (see [figure 14.8](#)). A convolution can be performed by transforming the image and the filter into the frequency domain, multiplying the two results, and transforming the result back into the spatial domain. Therefore the two reasons for using FFT are:

- For bigger filter sizes it is more efficient to apply the filters in Fourier domain as FFT always takes the same time for any mask size, unlike normal filters.
- It is a great advantage that filters can be customized to remove or emphasize specific frequencies from the image.




Figure 14.8: a) Original image b) with FFT enhanced defect structure.

Standard operators for FFT are:

- `fft_generic`: a generic version of the complex Fourier transform,
- `rft_generic`: real-valued Fourier transform,
- `fft_image`, a shortcut for the forward transformation with `fft_generic`

- and `fft_image_inv`, a shortcut to the backward transformation with `fft_generic`.

14.6.2 Texture Analysis in Color Images

By combining texture analysis with [Color Processing](#) on page 173, even more tasks can be solved efficiently. Typically, the color image is first split into color channels and possibly transferred into another color space. Then, texture filtering is applied and/or texture features are computed.

14.7 More Information About Texture Features

14.7.1 Entropy and Anisotropy (`entropy_gray`)

Entropy and anisotropy describe the gray value distribution of an image independent of the position of the pixels.

The entropy for white noise, as shown in image a) of [figure 14.9](#), is high as white noise comprises of an equal number for all gray values.

Even though binary noise (image b) of [figure 14.9](#) looks similar to white noise at first sight, it is in fact quite different because it only consists of two gray values: black and white. Therefore, the entropy is very low.

Image c) of [figure 14.9](#) shows a gray ramp and image a) of [figure 14.9](#) shows white noise. Both have the same value for entropy even though they look very different. Both images include, however, the same frequency for each gray value. This shows that the position of pixels within an image has no influence on the image's entropy.

Even though real images do not usually show such extreme values, the examples in [figure 14.10](#) still highlight the connection between gray value distribution (as is presented in the histogram) and the entropy value.

Anisotropy determines the symmetry of the gray value distribution. A value of -0,5 means that the histogram curve is symmetric. The value is higher if the image has more bright parts and lower if it has more dark parts as becomes clear from images a) and b) of [figure 14.11](#).




Figure 14.9: a) White noise; b) binary noise; c) gray ramp.

14.7.2 Cooccurrence Matrix ([gen_cooc_matrix](#))

The cooccurrence matrix describes the relationship between the gray value of a pixel and the gray values of its neighbors. The matrix elements contain the probability that two gray values appear next to each other. Slow gray value changes are indicated by high values along the diagonal of the principal axis. Strong contrasts lead to entries far away from the diagonal of the principal axis. The connection between gray values in the image, in the histogram, and finally in the cooccurrence matrix is shown in images a), b), and c) of [figure 14.12](#). Computing the cooccurrence matrix, you can select which neighbors are evaluated. Based on the cooccurrence matrix, four features can be computed:

- energy,
- correlation,
- homogeneity, and
- contrast.

14.7.3 Features of the Cooccurrence Matrix

Energy

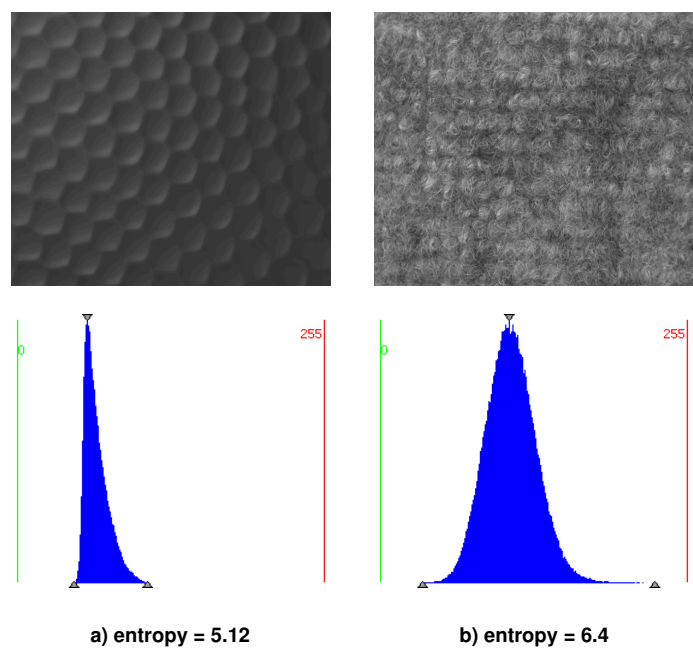


Figure 14.10: Entropy values of a) honeycomb structure and b) dust filter.

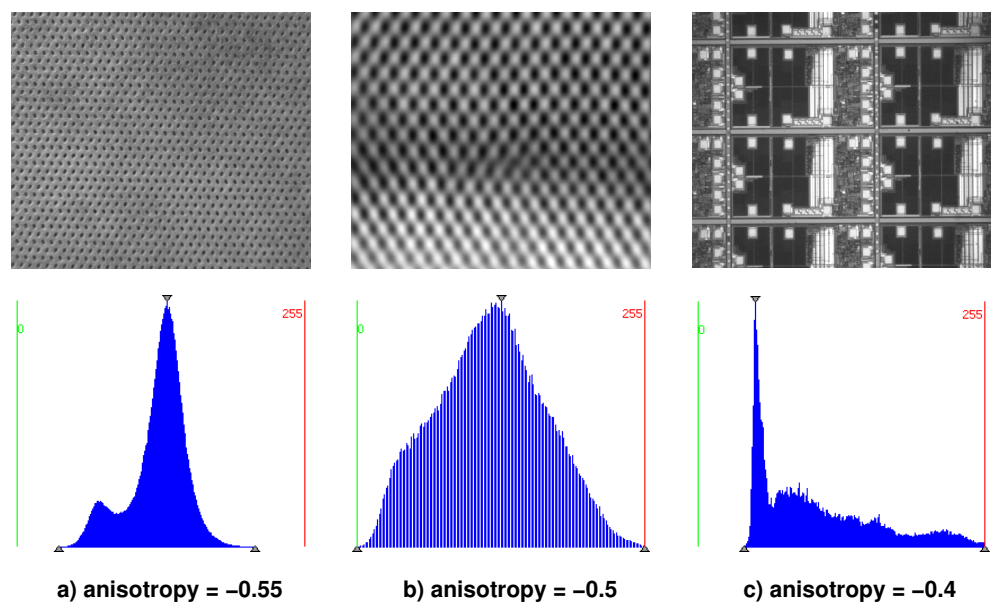


Figure 14.11: Anisotropy in images with different gray value distributions.




Figure 14.12: a) Original image; b) gray value histogram; c) cooccurrence matrix.

Energy is actually the short form for *uniformity of energy* and indicates if all kinds of gray value combinations are present or if certain gray value combinations are dominant. An image with a constant gray value has the highest energy.




Figure 14.13: Energy values.

Image a) in [figure 14.13](#) has a high energy because there are only few gray value combinations, which change always according to the same scheme. Image b) has little energy because the changes between pixels differ, i.e., there is no uniformity.

The textures in [figure 14.14](#) show how parts that contain pixels of the same gray value contribute to the higher energy of image a), whereas gray values in image b) are almost constantly changing and therefore lead to a low energy.

Correlation

Correlation defines the measure of dependency between neighboring values. If the matrix shows a thin line from top left to right bottom, this means that the image is blurry, which is why the values are so close together. A wide, roundish progress indicates a focussed image. Contrasts and little or no correlation are indicated by pixels far away from the main diagonal. The question of correlation between pixels can however sometimes depend on the image's orientation as is depicted in [figure 14.15](#).

The two correlation values of the same image are very different depending on the rotation of the image. 0° causes a low correlation whereas 90° has a high one. In [figure 14.16](#) the correlation values show that




Figure 14.14: Energy values of a) canvas and b) dust filter.




Figure 14.15: Correlation values for different image orientations. Note that images b) and c) are depicted in pseudocolors to enhance the perceptability of the cooccurrence matrix.

image a) with extreme gray value changes has a lower value than image b) with a texture that consists of gray values which are quite similar to each other.




Figure 14.16: Correlation values of a) mesh and b) dust filter.

Homogeneity

Homogeneity is high if large areas in an image have the same gray values.




Figure 14.17: Homogeneity values.

Image a) of [figure 14.17](#) does not include any groups of pixels with the same gray values which makes it very inhomogeneous. Image b) with white noise includes areas, a feature which makes it more homogeneous. A unicolored image has the highest homogeneity value.




Figure 14.18: Homogeneity values of a) mesh and b) grid.

Realistic images, as shown in [figure 14.18](#), clarify these homogeneity rules. Even though both images look partially homogeneous (they both contain bright and dark parts of a certain size but also include some changes from dark to bright) there are little differences. The mesh image is less homogeneous because the structure is three dimensional and therefore little deviations in brightness can cause slightly different gray values on the white mesh. Image b), however, contains bigger parts with a same gray value (the white background) and is also only two dimensional.

Contrast

The contrast value between two pixels is multiplied and squared. High differences lead to steep edges and a high value.

The example images in [figure 14.19](#) depict lines that run in different distances from each other. The




Figure 14.19: Contrast values.

image with the denser lines also has a higher contrast. This shows that the contrast value includes the frequency as well as the intensity.




Figure 14.20: Contrast values of a) grid and b) dust filter.

In [figure 14.20](#), the grid structure in image a) has a higher contrast than the dust filter in image b) because the gray value changes in image a) are more extreme than the gray value changes in image b).

14.8 More Information About Texture Filtering

14.8.1 The Laws Filter ([texture_laws](#))

By applying a Laws filter you can enhance or suppress textural structures in an image based on their spatial frequency, i.e., the size of the texels.

The operator [texture_laws](#) applies texture transformations (according to K. I. Laws) to an image. This is done by convolving the input image with a special filter mask of either 3x3, 5x5, or 7x7 pixels. For most of the filters the resulting gray values must be modified by a shift. This makes the different textures in the output image more comparable to each other, provided suitable filters are used.

The name of the filter is composed of the letters of the two vectors used, where the first letter specifies a convolution in the column direction, while the second letter specifies a convolution in the row direction. The letters themselves specify the convolution. You can select "l", "e", "s", "r", "w", "o" with "l" emphasizing low and "o" emphasizing high frequencies.

Therefore, the second letters detecting low frequencies to high frequencies are "l", "e", "s", "r", "w" and "o". The example figures [figure 14.21](#) and [figure 14.22](#) show how the two common texture filters "le" and "ls" work on two different textures.




Figure 14.21: a) Mesh (original image) b) Lapslacian filter le and c) Lapslacian filter ls.

The Lapslacian filter "le" looks for low frequencies in the image. As can be seen in [figure 14.21](#), there are a few low frequencies in the image. The Lapslacian filter "ls", however, looks for high frequencies, which occur very rarely in the image in places where gray values change from black to white and vice versa. As a result, the filtered image looks quite dark.




Figure 14.22: a) Dust filter (original image) b) Lapslacian filter le and c) Lapslacian filter ls.

The filtered images in [figure 14.22](#) show that the original image contains low frequencies as well as high frequencies as it responds to both the "le" and "ls" filter. The Lapslacian filter is therefore useful to distinguish textures which differ in their frequencies in a certain direction.

Chapter 15


Bar Code

The idea of bar code reading is quite easy. You initialize a bar code model and then execute the operator for reading bar codes. Within this operator, you specify the desired bar code type. That way, you can read different bar code types by successively applying the operator with different parameters for the type, but without the need to create a separate model for each bar code type. The result of the reading is a region that contains the bar code and the decoded string.

The advantage of the HALCON bar code reader is its ease of use. No advanced experience in programming or image processing is required. Only a few operators in a clear and simple order are applied. Furthermore, the bar code reader is very powerful and flexible. An example for this is its ability to read an arbitrary number of bar codes of the same code type in any orientation even if parts are missing.

15.1 Basic Concept

Bar code reading consists mainly of four steps:



15.1.1 Acquire Image(s)

For the online part, i.e., during reading only, images must be acquired.

For detailed information see the [description of this method](#) on page 13.

15.1.2 Create Bar Code Model

You create a bar code model with the operator `create_bar_code_model`. The returned handle provides all necessary information about the structure of the bar code. In most cases, you do not have to adjust any parameters.

15.1.3 Read Bar Code(s)

Bar codes are read using the operator `find_bar_code`. Within this operator, you specify the model and the bar code type that you search for. The operator returns the regions and the decoded strings of all bar codes of the specified type that are found in the image or in the specified region of interest.

15.1.4 Destroy Bar Code Model

When you no longer need the bar code model, you destroy it with the operator `clear_bar_code_model`.

15.1.5 A First Example

As an example for this basic concept, here is a very simple program, that reads the EAN 13 bar code depicted in [figure 15.1](#).

A test image is acquired from file and the bar code model is created with `create_bar_code_model`. Then, the operator `find_bar_code` (with `CodeType` set to 'EAN-13') returns the region and the decoded string for the found bar code. At the end of the program, the bar code model is destroyed using `clear_bar_code_model`.

```
read_image (image, 'barcode/ean13/ean1301')
create_bar_code_model ([], [], BarCodeHandle)
find_bar_code (image, SymbolRegions, BarCodeHandle, 'EAN-13', \
               DecodedDataStrings)
clear_bar_code_model (BarCodeHandle)
```

15.2 Extended Concept

In some cases, bar code reading can be more advanced than in the example above. For example, a preprocessing or rectification of the image as well as the visualization of results can be necessary.



Figure 15.1: Reading a bar code.

15.2.1 Use Region Of Interest

Bar code reading can be sped up by using a region of interest. The more the region in which the code is searched can be restricted, the faster and more robust the search will be.

For detailed information see the [description of this method](#) on page 19.

15.2.2 Preprocess Image(s)

There are various possibilities of preprocessing bar codes. Some of the most common ones are listed below.

If a bar code image is a bit blurry, [emphasize](#) can be used to emphasize high frequency areas of the image and, therefore, let it appear sharper.

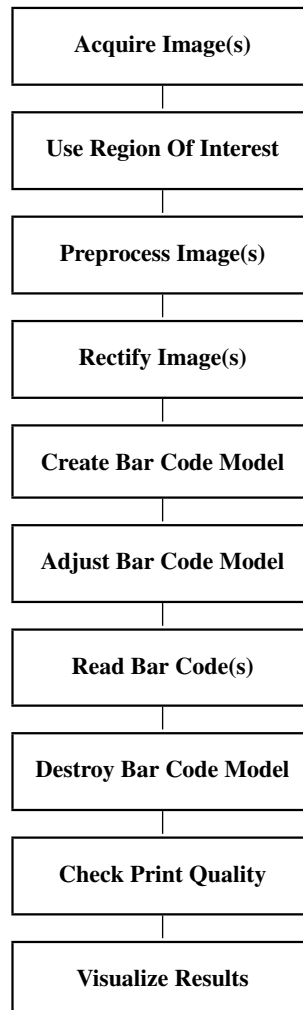
The operator [zoom_image_factor](#) enlarges a bar code that is too small to be read.

Low contrast images can be enhanced by scaling the gray value range with [scale_image](#) (or easier with the external convenience procedure `scale_image_range`).

HALCON expects bar codes to be printed dark on a light background. To read light bar codes on a dark background you must first invert the image using the operator [invert_image](#).

15.2.3 Rectify Image(s)

The bar code reader can handle perspective distortions up to a certain extent. For very strong distortions, it might be necessary to rectify the image before applying the bar code reader.



How to perform the rectification for a bar code that is printed radially on a CD is shown in the description of the example program [circular_barcode.hdev](#) on page 228. Detailed information about compensating distortions caused, e.g., by non-perpendicularly mounted cameras, can be found in the Solution Guide III-C in [section 3.4](#) on page 62.

15.2.4 Create Bar Code Model

Sometimes, a better result can be obtained when adjusting parameters like the range for the element size ('element_size_min', 'element_size_max'). All available parameters can be changed either here, or in a later step using [set_bar_code_param](#).

15.2.5 Adjust Bar Code Model

In most cases no adaptation of parameters is necessary for reading bar codes with HALCON. If the bar code reading fails, however, you can adjust the following parameters with the operator `set_bar_code_param` before applying `find_bar_code`.

Adjust Element Size

In particular, you can adjust the range of the bar code element's width with

- the parameter `'element_size_min'` and
- the parameter `'element_size_max'`.

The description of the example `barcode_typical_cases.hdev` on page 221 shows how an adaption of the element's size can improve the decoding.

Determine Check Character

For code types having a check character, you can specify whether the check character is used to check the result or if it is returned as part of the decoded string (`'check_char'`).

Adapting the Thresholds for the Extraction of Element Edges

Two threshold parameters control the detection of edges within a bounding box. `'meas_thresh'` calculates a relative threshold by detecting the highest and lowest gray values within a rough bar code region and using this dynamic range as a basis for edge detection. `'meas_thresh'` works well unless images contain lots of noise. It also fails if the code is a stacked code, which, due to its structure, contains a white patch within the bounding box. In both cases the calculated gray-value range of the images is small. This results in an unreasonably small value for `'meas_thresh'` and can therefore cause the detection of false edges.

This, however, is prevented by `'meas_thresh_abs'`, which sets an absolute threshold of 5.0 for all images. This means that the difference between the brightest and the darkest value of the element border has to be more than 5.0 to be classified as an edge. Therefore, no false edges are detected in images with a small dynamic range.

Sometimes, a low threshold is needed, e.g., if an image contains parts with a high dynamic range as well as parts with a low one, due to inhomogeneous illumination. Then, `'meas_thresh_abs'` must be turned off, by setting it to 0.0 before decreasing the value of `'meas_thresh'`.

Speedup

There are two cases in which you can improve the performance. The first case concerns images with too many false candidates, which are scanned with the default value of 10 scanlines each, whereas a valid bar code itself is usually decoded after one or two scans. Reducing the number of scanlines therefore leads to a significant speedup. Note that if the bar code cannot be detected after reducing the number of scanlines, the number has to be increased again. The second case concerns images with stacked bar codes (currently RSS-14 Stacked, RSS-14 Stacked Omnidirectional, and RSS Expanded Stacked). The reader evaluates by default the necessary number of scanlines, 20 for RSS-14 Stacked and 55 for RSS Expanded Stacked, which may be higher than necessary. If an RSS Expanded Stacked code does not have the maximum possible number of rows, it is useful to specify the number of scanlines with the parameter `'num_scanlines'` (the reader uses 5 scanlines per row). Consequently, the performance is

increased. Both cases are demonstrated in the example `examples\hdevelop\Identification\Bar-Code\barcode_param_num_scanlines.hdev`. Please refer to the Reference Manual for further information. If the number of bar codes per image is known, decoding time can be decreased by specifying this number with the parameter `'stop_after_result_num'` of the operators [set_bar_code_param](#) or [set_bar_code_param_specific](#). This way, the detection of false codes after all existing bar codes have already been read can be prevented and, as a consequence, processing can be sped up.

Increase Bar Code Detection Robustness

The parameter `'min_identical_scanlines'` specifies how many successfully decoded scanlines with identical result are required to decode a bar code. If this parameter is not set, a bar code is considered as decoded after the first scanline was successfully decoded. It can be useful to increase the value of this parameter when working either with an image containing lots of noise or with various bar codes to prevent the detection of false bar codes. The use of the parameter is demonstrated in the example `hdevelop/Identification/Bar-Code/barcode_param_scanning_control.hdev`.

Adapt Tolerance for Start/Stop Matching Criteria

The `'start_stop_tolerance'` specifies the tolerance for detecting the start- and stop pattern. A less tolerant criterion increases the robustness against false detections, but might reduce the general detection rate. This criterion can be activated by setting the parameter to `'low'`. A `'high'` tolerance is recommended if the boundary of the symbol in the image is clearly visible, the image is not noisy, and the symbol does not contain bar codes of different types. Note that the parameter is currently only implemented for Code 128. How to adapt the parameter `'start_stop_tolerance'` to avoid misdetection is demonstrated in the example `hdevelop/Identification/Bar-Code/barcode_param_scanning_control.hdev`.

To query the values that are currently used for the individual parameters, you can apply the operator [get_bar_code_param](#).

Determine Bar Code Parameters by Training In order to determine the best parameter settings, an automatical training can be performed. The advantage of this training is an increased robustness of the bar code reader. Furthermore, the decoding time can be reduced as a side effect. Parameter settings can be determined for `'element_size_min'`, `'element_size_max'`, `'orientation'`, `'meas_thresh'`, `'meas_thresh_abs'` and `'max_diff_orient'`. To activate the training mode, use the operator [create_bar_code_model](#) to create a model for the training and set the generic parameter `'train'` to a tuple of parameter names that should be trained or `'all'` to train all bar code parameters that are available for the training. To determine all parameters with given properties, e.g., `'trained'`, the operator [query_bar_code_params](#) can be used. How to perform the training is demonstrated in the example `examples\hdevelop\Identification\Bar-Code\barcode_training.hdev`. Please refer to the Reference Manual for further information.

15.2.6 Read Bar Code(s)

The operator [find_bar_code](#) allows you to either find one bar code type or several bar code types in one call. When reading one code type, choose the type - if it is known - as `'CodeType'`.

Autodiscrimination

Autodiscrimination describes the simultaneous decoding of multiple bar code types. It is activated by setting `'CodeType'` to `'auto'` or by choosing the expected bar code types. Note, however, that each allowed bar code type increases the runtime of the operator and too many bar code types may decrease the

reliability of the decoding. You should at least exclude all definitely not occurring bar code types that are scanned before the first expected bar code type occurs or, even better, just scan for the expected bar code types (see [figure 15.2](#)). For more information on autodiscrimination including the decoding order, please refer to the documentation of [find_bar_code](#) in the Reference Manual. Furthermore, the example `examples\hdevelop\Identification\Bar-Code\barcode_autodiscrimination.hdev` demonstrates the practical use of autodiscrimination. The operators [set_bar_code_param_specific](#) and [get_bar_code_param_specific](#) allow to set parameters of the bar code model specifically for certain bar code types, which may especially be interesting for autodiscrimination.




Figure 15.2: Using the 'autodiscrimination' feature, different bar codes within one image - in this case two codes of the type Code 39 and one EAN-13 - can automatically be recognized and read.

Restrict the reading to a part of the image

If the code always appears approximately in the same position within every image, you can speed up your application using the operator [decode_bar_code_rectangle2](#) instead of the operator [find_bar_code](#). Rather than performing a time-consuming search for candidate regions, [decode_bar_code_rectangle2](#) scans the provided region directly for bar codes. The position of the code is given in an arbitrarily oriented rectangle. Note that the direction is important. It needs to be perpendicular to the elements.

Access results

After reading the bar codes with [find_bar_code](#), you can explicitly query the regions of the decoded bar codes or their candidates with the operator [get_bar_code_object](#). The decoded strings or the

underlying reference data (including, e.g., the start/stop characters) can be queried with the operator `get_bar_code_result`.

Diagnostic status messages

In order to control why scanlines are not successfully detected or why the code reading fails completely, diagnostic status messages in human readable format can be returned (see [figure 15.3](#)). With the operator `get_bar_code_result` set to 'status' or 'status_id' and supplying a candidate handle, diagnostic status messages for each possible scanline in the candidate region are returned in a tuple. Use the operator `get_bar_code_object` with the parameter 'scanlines_all' to query the corresponding scanlines. Note that the status of the scanlines of a bar code can only be evaluated if the operator `find_bar_code` or the operator `decode_bar_code_rectangle2` was called before while in 'persistence' mode. Determining the status increases runtime and should therefore only be performed if the debugging information is needed. The example `examples\hdevelop\Identification\Bar-Code\barcode_status.hdev` demonstrates the use of the bar code parameters 'status' and 'status-id'.

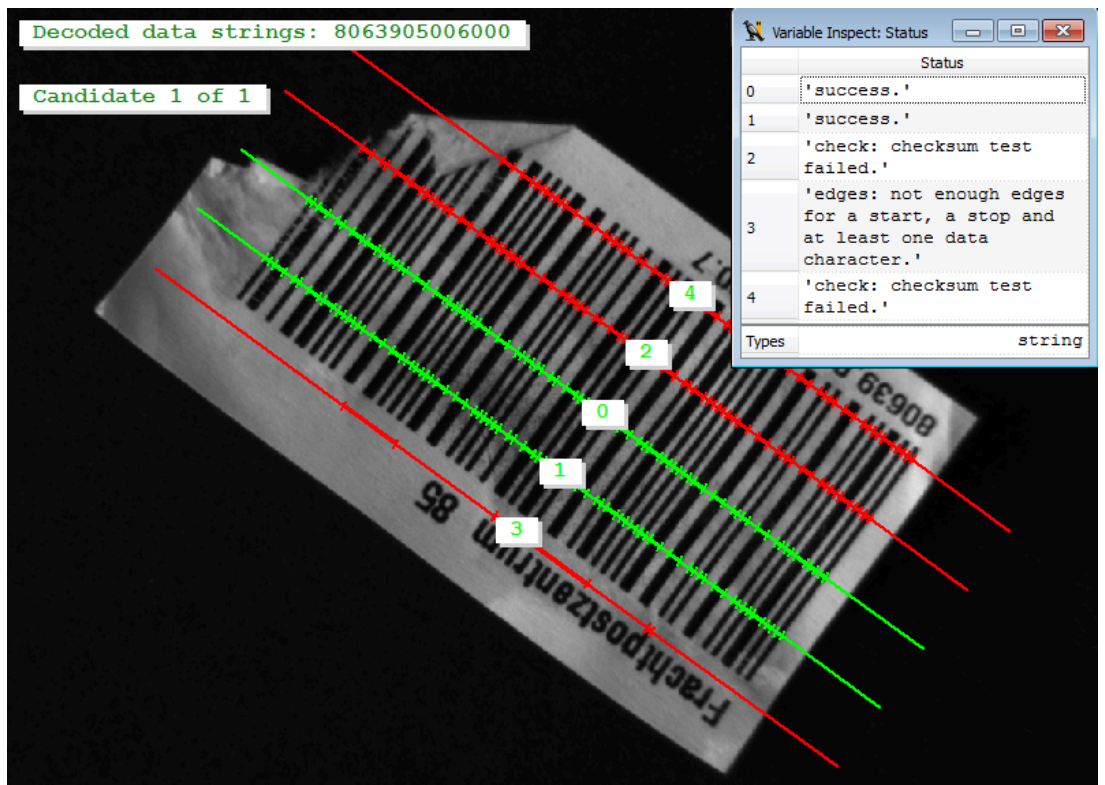


Figure 15.3: Status messages can be used to solve problems that might occur when reading bar codes.

What if the code reading fails?

If the code reading fails, an error message is returned. Depending on the message we have several suggestions on how to solve the problem.

Error Message(s)	Possible Solution
edges: not enough edges detected.	Inspect the region by querying region and scanlines with the operators <code>get_bar_code_object</code> and <code>get_bar_code_result</code> .
edges: not enough edges for a start, a stop and at least one data character.	
edges: too many edges detected.	
edges: center of scanline not within domain.	Change the ROI or use the operator <code>full_domain</code> to search the whole image. If the position of the code is known, the operator <code>decode_barcode_rect2</code> can be used to specify the exact position.
decoding: could not find stop character.	Check the found bar code region or increase the number of scanlines with the parameter <code>num_scanlines</code> of the operator <code>set_bar_code_param(_specific)</code> .
decoding: could not find start and stop characters.	
decoding: number of wide bars of a single character is not equal to 2.	Check if the right code type was specified as 'CodeType' for the operator <code>find_bar_code</code> .
decoding: invalid encoding pattern.	
decoding: invalid mix of character sets.	
decoding: error decoding the reference to a human readable string.	
decoding: could not detect center guard pattern.	
decoding: could not detect left and/or right guard patterns.	
decoding: could not detect add-on guard pattern.	
decoding: could not detect enough finder patterns.	
decoding: no segment found.	
check: detected EAN-13 bar code type instead of specified type.	
check: checksum test failed.	Check if the code even has a checksum. If it does not, set the parameter 'check_char' of the operator <code>set_bar_code_param_specific</code> to 'absent'.
check: check of add-on symbol failed.	Check if the code has an add-on symbol. If it does not, change the 'CodeType' of <code>find_bar_code</code> to the non-add-on variant.
check: symbol region overlaps with another symbol region.	The whole region is skipped because another symbol region that overlaps with this region has already been decoded. This is no error. Set 'stop_after_result_num' of the operator <code>set_bar_code_param</code> or <code>set_bar_code_param_specific</code> to the number of expected codes.

Warning Message(s)	Possible Solution
White spaces too wide.	To improve image quality and prevent errors, please refer to the example program <code>examples\solution_guide\basics\barcode_typical_cases.hdev</code> .
White spaces too narrow.	
Bars too wide.	
Bars too narrow.	
Possible saturation of gray values.	
No composite component found.	Check if the composite code is completely within the image.

If regions are not found properly, you can either improve the image quality itself, which helps to solve many problems, or - if that is not possible - adjust the software to better detect your code. If you know the location of your bar code, you can use the operator `decode_barcode_rect2` to specify the location instead of searching the whole image (see above).

Another reason why regions are not detected can be an element size that is not adjusted to the actual size of the bar code - which is especially a problem, if the code is very big or very small. You can solve the problem by adjusting the parameters `element_size_min` or `element_size_max`. How to adapt these parameters and handle other typical cases is described in the example `examples\solution_guide\basics\barcode_typical_cases.hdev`, which is documented in this chapter (see the example descriptions in [section 15.3](#)). This example presents some typical cases of problems that can occur while reading bar codes and presents some easy solutions.

The error messages (see the table of error messages on [page 219](#)) which lead to a failure of code reading begin either with 'edges', 'decoding' or with 'check'. These categories represent the state within the decoding process in which the code reading failed. First, edges of the bars, i.e. the white-to-black and black-to-white transitions from background to bar and bar to background, are detected. Then the code is decoded according to the specified symbology and finally several tests ('check') are performed to ensure a correct result.

The errors with the codes (0),(8),(9) and (13) ('unknown status', 'decoding: internal error when estimating the maximum string length', 'decoding: internal error when decoding a single character' and 'decoding: error decoding the reference to a human readable string') are internal errors that should not occur. If they do, however, occur, please contact your local HALCON distributor.

Warning messages (see [page 220](#)) occur either with a success message or with an error message. Warnings do not prevent a successful code reading but are merely an indication for possible (future) problems.

15.2.7 Check Print Quality

If you are interested in checking the print quality of your bar code symbols, you can use the HALCON operator `get_bar_code_result`, which grades the symbol according to the standard ISO/IEC 15416. For simple 1D bar codes the print quality is described in a tuple with nine elements:

- Overall Quality,

- Decode,
- Symbol Contrast,
- Minimal Reflectance,
- Minimal Edge Contrast,
- Modulation,
- Defects,
- Decodability,
- and Additional Requirements.

Note that the names and the order of the grades can be queried with the convenience option (`query_isoiec15416_labels`).

For composite bar codes, the print quality is described in a tuple with 24 elements, including components of both bar code print quality grading and data code print quality grading as described for the method [Data Code](#) on page 233 or in more detail in the Solution Guide II-C in [section 6](#) on page 51.

When investigating the reason for quality defects, it can be useful to query the data that was used to calculate the values for the print quality elements. This is possible with the parameter `quality_isoiec15416_values` which returns a tuple with the raw values for the following grades: Symbol Contrast, Minimal Reflectance, Minimal Edge Contrast, Modulation, Defects, and Decodability for simple bar codes, and Symbol Contrast, Minimal Reflectance, Minimal Edge Contrast, Modulation, Defects, Decodability, Rap Contrast, Rap Minimal Reflectance, Rap Minimal Edge Contrast, Rap Modulation, Rap Defects, Rap Decodability, Codeword Yield, and Unused Error Correction for composite bar codes. For grades that are excluded from these lists the operator reports 'N/A'.

For further information on checking print quality of composite bar codes or regular bar codes, please refer to the description of the operator [get_bar_code_result](#) in the reference manual.

15.2.8 Visualize Results

Finally, you might want to display the images, the bar code regions, and the decoded content.

For detailed information see the [description of this method](#) on page 289.

15.3 Programming Examples

This section shows how to use the bar code reader.

15.3.1 How to Read Difficult Barcodes

Example: `solution_guide/basics/barcode_typical_cases.hdev`

When reading bar codes, there are a few conditions you have to consider. These conditions might have to do with the code itself, for which you may, for example, have to use different parameters because it is very big or small, or with the image acquisition environment affecting the image quality. Even though you should always aim for achieving very good images, for technical reasons, it might sometimes not be possible to reach a perfect image quality.

The following example can help you identify simple obstacles, find out at which stage of the bar code reading process they occur, and it subsequently offers suggestions for solving the problem.

The different cases within the example are based on typical defects. There are no problems with reading the bar code for case 0. It just presents the two stages at which intermediate results can be obtained (see [figure 15.4](#)).

You can query candidate regions to find out if the reader detects a bar code candidate at all with the local procedure `disp_bar_code_candidates`. To query scanlines, it is necessary that a candidate region is found. You can observe the scanlines with the local procedure `disp_bar_code_scanlines`. Both procedures use the operator `get_bar_code_object`, setting either the parameter 'candidate_regions', 'scanlines_all' or 'scanlines_valid', respectively.

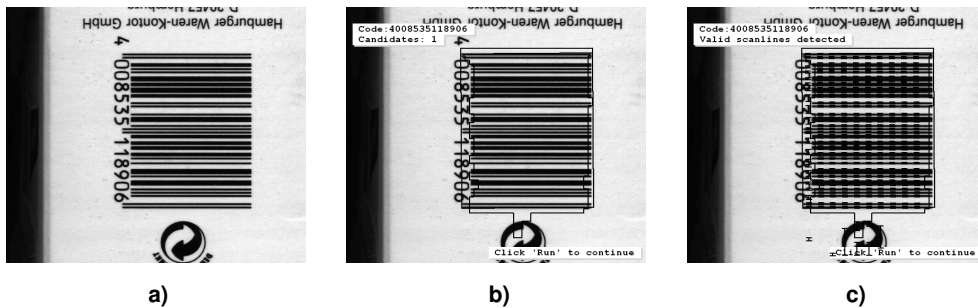


Figure 15.4: a) Original bar code image b) A candidate is detected. c) Valid scanlines are detected.

Case 1 describes a situation where a blurry bar code has to be decoded (see [figure 15.5](#)). A first attempt to read this code fails. Checking at which stage the bar code reading fails, you can see that a candidate is detected but no valid scanlines are found. To make the code readable again, the edges of the bars have to be sharpened. This can be achieved with the operator `emphasize`. By filtering the image, the operator enhances high frequency areas, like edges and corners and therefore makes it look sharper. To calculate the mask width and height, the blur radius, i.e., how many pixels are blurred, has to be measured. The easiest way to measure in an image is to zoom the window by choosing **Zoom Window** under **Visualization** in the window bar. The new window allows you to move the mouse cursor over the image and check row and column in the lower right corner of the **Zoom** window. Depending on your measuring direction, the row or column number, the difference between your start point and your end point equals the size of the blur radius in pixels. As `emphasize` needs uneven numbers as mask sizes, the formula is to multiply this result by two and then add one.

```
BlurRadius := 5
emphasize (Image, ImageEmphasize, BlurRadius*2+1, BlurRadius*2+1, 1)
```




Figure 15.5: a) A blurry bar code cannot be decoded. b) After enhancing the contrast, the code can be read.


Subsequently, it is possible to read the bar code immediately and both candidate and scanlines are found.

Case 2 is very easy to solve as well. The problem here is to keep in mind, that very large bar codes (see [figure 15.6](#)) demand a change of the values for the parameter 'element_size_max'. Otherwise, the default parameter is too small with the result that no candidate is found, as in this case. Adapting the parameter 'element_size_max' with the operator [set_bar_code_param](#) to the size of the elements makes the code readable. How you can measure the width of the elements in HDevelop is described in case 1.

```
set_bar_code_param (BarcodeHandle, 'element_size_max', 32)
```



a)



b)

Figure 15.6: a) The large bar code cannot be read. b) After adapting the corresponding parameter to the element size, the code can be read.

Case 3 describes the opposite problem. A very small bar code has to be decoded (see [figure 15.7](#)). First, it cannot be read because no candidates are found. The solution is very simple: Zooming the image with 'zoom_image_factor' makes it readable again. Generally, zooming is necessary for barcode images with elements that are smaller than two pixels in width.

```
zoom_image_factor (Image, ImageZoomed, 2, 2, 'none')
```

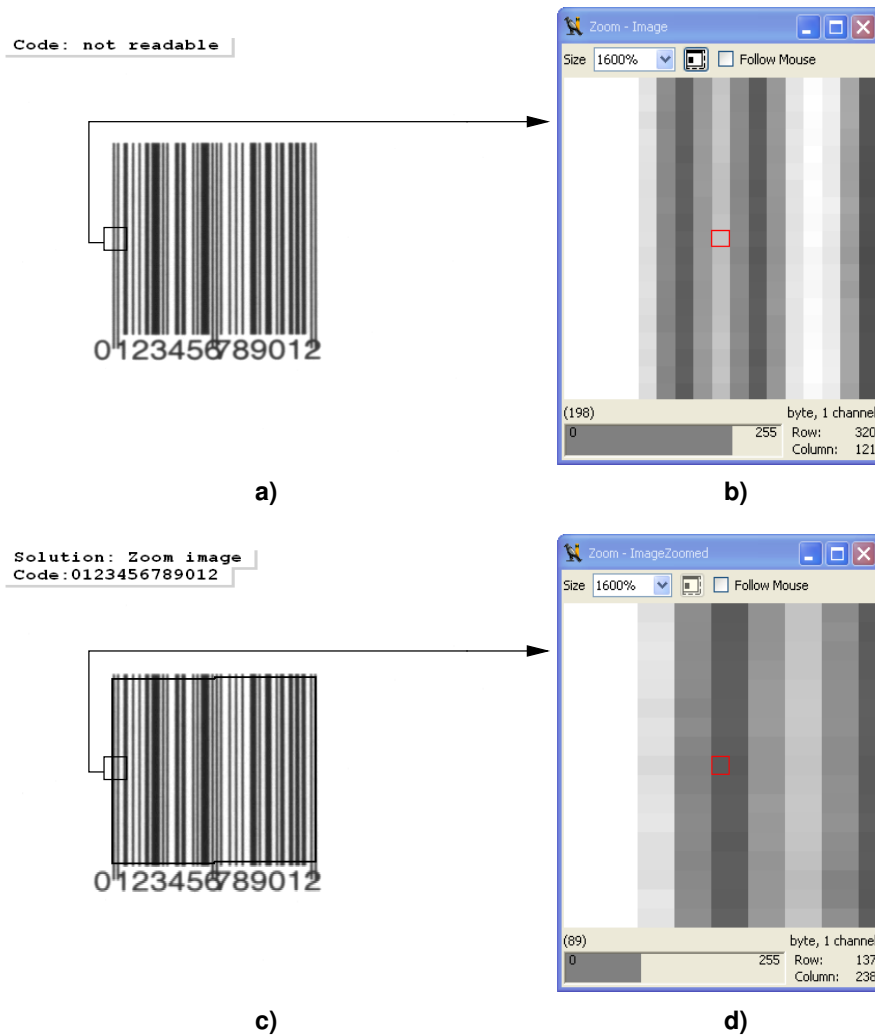


Figure 15.7: a) The small code cannot be read. b) The **Zoom** window shows that the elements are too small. c) Zooming the image makes it possible to read the code. d) Due to the zooming, the element width is larger now.

Case 4 deals with a bar code in a very low contrast image (see [figure 15.8](#)). Subsequently, no candidates are found. However, just one preprocessing step is necessary to overcome this obstacle. The knowledge that there is a defined gray value difference between the code bars and the background leads to the solution. The single gray values lie, however, very close together which makes it impossible for the reader to distinguish the code from the background. The contrast can simply be improved by scaling the gray value range of the image with the convenience procedure `scale_image_range` and therefore achieving a bright background and a black bar code again. To obtain the scaling range, it is useful to check the gray histogram of the image.


```
scale_image_range (Image, ScaledImage, 0, 20)
```

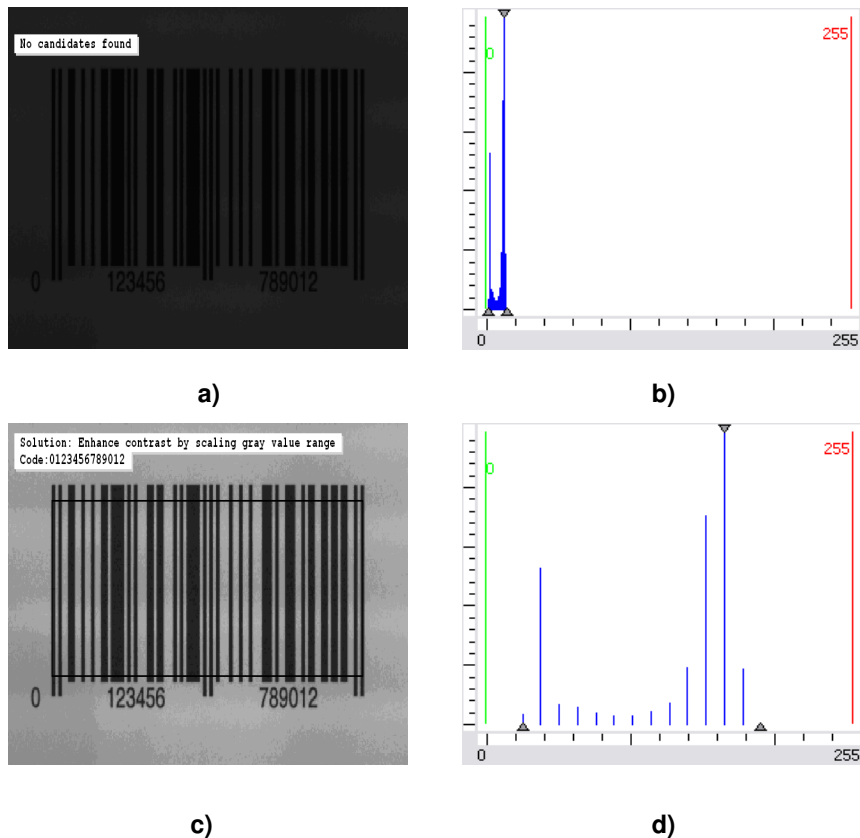


Figure 15.8: a) A bar code with a very low contrast. b) Histogram of the low contrast image. c) Scaling gray values make it possible to decode the bar code. d) The histogram of this image shows how the gray values have been scaled.

After scaling the image, the code can be read and both candidate and scanlines are found.

Case 5 shows a bar code in an overexposed image (see [figure 15.9](#)). Neither candidates nor scanlines of the code can be found. To reverse the effect of an overexposed image, causing a decreased contrast and slimmer code elements, `gray_erosion_shape` can be used. However, the overexposure radius, which is the amount of enlargement of bright areas that is caused by overexposure, is needed first. It can be measured as described for case 1. From this radius the image width and height for `gray_erosion_shape` is calculated by multiplying it with two and adding one.

```
OverexposureRadius := 0.8
gray_erosion_shape (Image, ImageErosion, OverexposureRadius*2+1, \
    OverexposureRadius*2+1, 'octagon')
```



Figure 15.9: a) An overexposed bar code b) `gray_erosion_shape` reverses the effect of the overexposure.

Now that the code is back to its normal shape again, both candidates and scanlines are found and the code can be read.

Case 6 shows another problem that is connected with illumination. The light distribution in the image is inhomogeneous (see figure 15.10). `'disp_bar_code_candidates'` does not detect the "whole" candidate first, because the contrast diminishes in the darker parts of the image. Decreasing the value of `'meas_thresh'` makes it possible to detect and read the whole bar code. It is, however, necessary to disable `'meas_thresh_abs'` first, because `'meas_thresh_abs'` sets a default threshold of 5.0, which is too high in this case.

```
set_bar_code_param (BarcodeHandle, 'meas_thresh_abs', 0.0)
set_bar_code_param (BarcodeHandle, 'meas_thresh', 0.01)
```

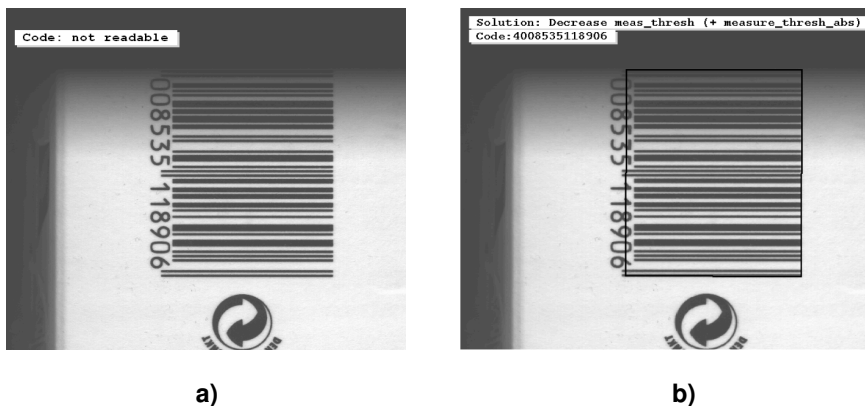


Figure 15.10: a) Bar code with inhomogeneous illumination b) Decreasing `'meas_thresh'` makes the code readable.

15.3.2 Reading a Bar Code on a CD

Example: `hdevelop/Applications/Bar-Codes/circular_barcode.hdev`

Figure 15.11 shows an image of a CD, on which a bar code is printed radially. The task is to read this circular bar code. Because the bar code reader cannot read this kind of print directly, first the image must be transformed such that the elements of the bar code are parallel.




Figure 15.11: (a) Original image with segmented ring; (b) rectified ring with decoded bar code.

The first step is to segment the dark ring on which the bar code is printed. This is performed using `bin_threshold` followed by `connection`. This segmentation returns all dark areas in the image, including the dark ring. To select the ring, `select_shape` is used with corresponding values for the extent.

```
bin_threshold (Image, Region)
connection (Region, ConnectedRegions)
select_shape (ConnectedRegions, Ring, ['width','height'], 'and', [550,550], \
[750,750])
```

After that, the parameters of the outer and the inner circle are determined. The outer circle can be determined directly using `shape_trans`. The inner circle is more complicated to extract. Here, it is calculated by creating the complement region of the ring with appropriate dimensions and then selecting its inner part. With `smallest_circle`, the parameters of the inner and outer circle are determined.

```
shape_trans (Ring, OuterCircle, 'outer_circle')
complement (Ring, RegionComplement)
connection (RegionComplement, ConnectedRegions)
select_shape (ConnectedRegions, InnerCircle, ['width','height'], 'and', \
[450,450], [650,650])
smallest_circle (Ring, Row, Column, OuterRadius)
smallest_circle (InnerCircle, InnerRow, InnerColumn, InnerRadius)
```

The parameters for inner and outer circle are the input for the polar transform (`polar_trans_image_ext`), which transforms the image inside the ring into a rectangular area. Then, the image is inverted to obtain dark bar code elements on a light background (see [figure 15.11b](#)).

```
WidthPolar := 1440
HeightPolar := round(OuterRadius-InnerRadius-10)
polar_trans_image_ext (Image, PolarTransImage, Row, Column, rad(360), 0, \
    OuterRadius-5, InnerRadius+5, WidthPolar, \
    HeightPolar, 'bilinear')
invert_image (PolarTransImage, ImageInvert)
```

To read the bar code, first a model is created with `create_bar_code_model`. As the bar code elements are very thin, parameters are adjusted with the operator `set_bar_code_param`. In particular, the parameter 'element_size_min' is set to 1.5 and the parameter 'meas_thresh' is set to 0.1. Then, `find_bar_code` finds and decodes the bar code, which is of the specified code type 'Code 128'. The bar code model is destroyed with `clear_bar_code_model`.

```
create_bar_code_model ([], [], BarCodeHandle)
set_bar_code_param (BarCodeHandle, 'element_size_min', 1.5)
set_bar_code_param (BarCodeHandle, 'meas_thresh', 0.3)
find_bar_code (ImageZoomed, SymbolRegions, BarCodeHandle, 'Code 128', \
    DecodedDataStrings)
clear_bar_code_model (BarCodeHandle)
```

Finally, the region of the bar code that was returned by `find_bar_code` is transformed back to the original shape of the bar code by `polar_trans_region_inv` and is displayed in the image.

```
polar_trans_region_inv (SymbolRegions, CodeRegionCircular, Row, Column, \
    rad(360), 0, OuterRadius-5, InnerRadius+5, \
    WidthPolar, HeightPolar, Width, Height, \
    'nearest_neighbor')
dev_display (CodeRegionCircular)
```

15.3.3 Checking Bar Code Print Quality

Example: `hdevelop/Identification/Bar-Code/print_quality_isoiec15416.hdev`

This example checks the print quality of bar codes, according to the ISO/IEC 15416 standard, for contrast, minimal reflectance, modulation, minimal edge contrast, defects, decodability and additional requirements.

This program will be explained exemplarily, using the "Contrast" grade, which checks whether the range between the minimal and the maximal value in the reflectance profile is lower or equal to 0.5 of the maximal reflectance value. Otherwise, a value of 0 is assigned as you can see in [figure 15.12](#). Furthermore, the program returns raw values for some quality grades to investigate the reason for quality defects.

First, the data code is detected in the image, then the operator `get_bar_code_result` checks the print quality of the code.

```
find_bar_code (ImageDefect, SymbolRegions, BarCodeHandle, 'EAN-13', \
               DecodedDataStrings)
get_bar_code_result (BarCodeHandle, 0, 'quality_isoiec15416', Quality)
```

Finally, the result is displayed in a message.

```
min_max_gray (SymbolRegions, ImageDefect, 0, Min, Max, Range)
Contrast := Range/255
dev_display (ImageDefect)
dev_set_color ('green')
dev_display (SymbolRegions)
grade_message_text (Labels, Quality, QualityValues, GRADE_CONTRAST, Message)
disp_message (WindowHandle, Message, 'window', 10, -1, 'black', 'true')
```



a)



b)

Figure 15.12: a) The bar code with a high contrast gets the grade 4 and a raw value of 82.45%. b) The defective bar code with a very low contrast gets the grade 0 and a raw value of 15.84%.

The program repeats these checks for the print quality features listed above. For more information concerning composite codes view the example `hdevelop/Identification/Bar-Code/composite_print_quality_isoiec15416.hdev`.

15.4 Selecting Operators

15.4.1 Acquire Image(s)

Please refer to the [operator list for the method Image Acquisition](#) (see section 2.4 on page 17).

15.4.2 Use Region Of Interest

Please refer to the [operator list for the method Region Of Interest](#) (see section 3.4 on page 26).

15.4.3 Preprocess Image(s)

Standard:

`emphasize`, `zoom_image_factor`, `scale_image`, `invert_image`

15.4.4 Rectify Image(s)

Operators for rectifying images are described in the [Solution Guide III-C](#).

15.4.5 Create Bar Code Model

Standard:

`create_bar_code_model`

15.4.6 Adjust Bar Code Model

Standard:

`get_bar_code_param`, `set_bar_code_param`

15.4.7 Read Bar Code(s)

Standard:

`find_bar_code`

Advanced:

`get_bar_code_object`, `get_bar_code_result`

15.4.8 Destroy Bar Code Model

Standard:

`clear_bar_code_model`

15.4.9 Check Print Quality

Standard:

`get_bar_code_result`

15.4.10 Visualize Results

Please refer to the [operator list for the method Visualization](#) (see section 19.4 on page 297).

15.5 Relation to Other Methods

15.5.1 Alternatives to Bar Code

OCR (see [description](#) on page 245)

With some bar codes, e.g., the EAN 13, the content of the bar code is printed in plain text below the elements. Here, OCR can be used, e.g., to check the consistency of the reading process.

15.6 Advanced Topics

15.6.1 Use Timeout

With the operator `set_bar_code_param`, you can set a timeout. Then, the operator `find_bar_code` will return at the latest after the specified time. This mechanism is demonstrated for [Matching \(Shape-Based\)](#) in the example `hdevelop/Matching/Shape-Based/set_shape_model_timeout.hdev`.

Chapter 16

Data Code

Data codes are a special kind of two-dimensional patterns that encode text and numbers. HALCON is able to read the three most popular data codes: Data Matrix ECC 200, QR Code, and PDF417. These codes consist of a so-called finder pattern, which is used to locate the pattern and get basic information about the geometric properties. The code itself contains multiple dots or small squares. Because of the special design of the codes, they can be decoded even if some parts are disturbed.

The advantage of the HALCON data code reader is its ease of use. No advanced experience in programming or image processing is required. Only a few operators in a clear and simple order need to be applied. Furthermore, the data code reader is very powerful and flexible. Examples for this are its ability to read codes in many print styles and the possibility to automatically learn optimal parameters.

16.1 Basic Concept

Data code reading consists mainly of four steps:


16.1.1 Acquire Image(s)

For the online part, i.e., during reading, images are acquired.

For detailed information see the [description of this method](#) on page 13.

16.1.2 Create Data Code Model

First, you create a data code model with the operator `create_data_code_2d_model`. This model provides the reader with all necessary information about the structure of the code. For normal printed codes only the name needs to be provided and HALCON will select suitable default parameters. For special cases, you can modify the model by passing specific parameters.



16.1.3 Read Data Code(s)

To read a data code, just one operator is needed: `find_data_code_2d`. It will locate one or more data codes and decode the content.

16.1.4 Destroy Data Code Model

When you no longer need the data code model, you destroy it with the operator `clear_data_code_2d_model`.

16.1.5 A First Example

As an example for this basic concept, here a very simple program, which reads the data code on the chip depicted in [figure 16.1](#), is discussed.

After reading an image from file, the data code model is generated by calling `create_data_code_2d_model`. As the only required parameter value, the code name 'ECC200' is specified.

```
read_image (Image, 'datacode/ecc200/ecc200_cpu_003')
create_data_code_2d_model ('ECC200', [], [], DataCodeHandle)
```

Then, the data code is read with the operator `find_data_code_2d`.

```
find_data_code_2d (Image, SymbolXLDs, DataCodeHandle, [], [], ResultHandles, \
                  DecodedDataStrings)
```

At the end of the program, the data code model is destroyed using `clear_data_code_2d_model`.



Figure 16.1: Reading a data code.

```
clear_data_code_2d_model (DataCodeHandle)
```

16.2 Extended Concept

In some cases, data code reading can be more advanced than in the example above. Reasons for this are, e.g., parameter optimization for improved execution time. Furthermore, preprocessing like rectification or the visualization of results might be required. The following sections give a brief overview. More detailed information can be found in the [Solution Guide II-C](#).

16.2.1 Acquire Image(s)

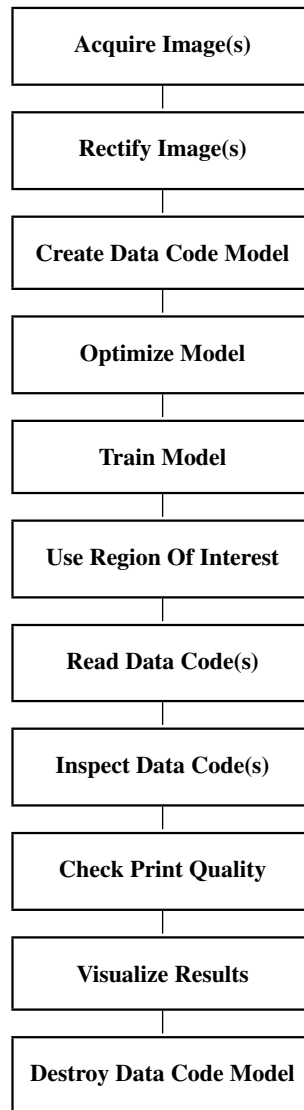
Optionally, additional images can be acquired for parameter optimization (see the description of the step [Optimize Model](#) on page 237).

For detailed information see the [description of this method](#) on page 13.

16.2.2 Rectify Image(s)

HALCON's data code reader is robust against image distortions up to a certain limit. But if the data code is printed on a cylindrical surface or if the camera is tilted relative to the surface, it might be necessary to rectify the image before applying the data code reader.

Detailed information about rectifying images can be found in the Solution Guide III-C in [section 3.4](#) on page 62.



16.2.3 Create Data Code Model

The operator `create_data_code_2d_model` expects the name of the desired code and optionally parameters to specify the geometry and radiometry as input. By default, a parameter set is used that is suitable for data codes that fulfill the following requirements:

- The code must be printed dark on light,
- the contrast must be bigger than 30,
- the size of symbol is smaller or equal to 48x48 modules,

- the width and the height of the modules are between 6 and 20 pixels, and
- there is no or only a small gap in between neighboring modules.

This parameter set is also used if you specify the value `'standard_recognition'` for the parameter `GenParamValues`. In contrast, if you specify the value `'enhanced_recognition'`, a parameter set is used that detects codes that do not follow the rules given above. However, using this parameter set possibly results in a longer processing time.

16.2.4 Optimize Model

Using the default parameters, the data code reader is able to read a wide range of codes. For non-standard codes the parameters can be adapted accordingly. For this, the operator `set_data_code_2d_param` is used.

The easiest way is to use the parameter value `'enhanced_recognition'`, which uses a model that is able to find a very wide range of print styles. An alternative is to specify parameter values separately to adapt the model to the conditions of the used print style.

If a data code symbol is not detected although it is well visible in the image, check whether the symbol's appearance complies with the model. In particular, have a look at the polarity (`'polarity'`: dark-on-light or light-on-dark), the module size (`'module_size'` and `'module_shape'`) and the minimum contrast (`'contrast_min'`). In addition, the parameters `'module_gap'` (allowed gap between modules), `'symbol_size'`, and `'slant_max'` (angle variance of the legs of the finder pattern) should be checked. The current settings of these values can be queried by the operator `get_data_code_2d_param`.

All possible parameter values can be checked in the Reference Manual. Besides this, they can also be queried with the operator `query_data_code_2d_params`.

As an alternative, you can train the model (see below).

16.2.5 Train Model

Instead of modifying the model parameters manually as described above, you can also let HALCON train the model automatically using the operator `find_data_code_2d`. All you need to do is to call this operator with the parameter values `'train'` and `'all'`. Then, HALCON will search for the best parameters needed to extract the given code. It is recommended to apply this to multiple example images to ensure that all variations are covered.

As an alternative, you can execute the finder with normal parameters and request the features of the found symbols with `get_data_code_2d_results`. These values can then be used to change the model with `set_data_code_2d_param`.

16.2.6 Use Region Of Interest

Reading data codes can be sped up by using a region of interest. The more the region in which codes are searched can be restricted, the faster and more robust the search will be.

For detailed information see the [description of this method](#) on page 19.

16.2.7 Read Data Code(s)

The operator `find_data_code_2d` returns for every successfully decoded symbol the surrounding XLD contour in `SymbolXLDs`, a handle to a result structure, which contains additional information about the symbol as well as about the search and decoding process (`ResultHandles`), and the string that is encoded in the symbol (`DecodedDataStrings`). With the result handles and the operators `get_data_code_2d_results` and `get_data_code_2d_objects`, additional data about the extraction process can be accessed.

16.2.8 Inspect Data Code(s)

Using the handles of the successfully decoded symbols returned by `find_data_code_2d`, you can request additional information about the symbol and the finding process using the operators `get_data_code_2d_results` and `get_data_code_2d_objects`. This is useful both for process analysis and for displaying.

In addition, information about rejected candidates can also be queried by requesting the corresponding handles with `get_data_code_2d_results` using, e.g., the parameter values `'all_undecoded'` and `'handle'`.

The operator `get_data_code_2d_results` gives access to several alphanumerical results that were calculated while searching and reading the data code symbols. Besides basic information like the dimensions of the code, its polarity, or the found contrast, also the raw data can be accessed.

The operator `get_data_code_2d_objects` gives access to iconic objects that were created while searching and reading the data code symbols. Possible return values are the surrounding contours or the regions representing the foreground or background modules.

16.2.9 Check Print Quality

If your first aim is not to quickly read the 2D data code symbols but to check how good they were printed, you can query the print quality of a symbol in accordance to the standards ISO/IEC 15415 or AIM DPM-1-2006. For details, see the Solution Guide II-C in [section 6](#) on page 51.

16.2.10 Visualize Results

Finally, you might want to display the images, the data code regions, and the decoded content.

For detailed information see the [description of this method](#) on page 289.

16.3 Programming Examples

This section gives a brief introduction to the programming of the data code reader.

16.3.1 Training a Data Code Model

Example: `solution_guide/basics/ecc200_training_simple.hdev`




Figure 16.2: (a) Dark training image; (b) bright training image; (c) read code with extracted modules.

In this example we show how easy it is to train a data code model, here to allow changes in the illumination of the images. To prepare the reading of the data codes, three major steps are performed: First, the connection to an image sequence is established by calling `open_framegrabber`.

```
SequenceName := 'datacode/ecc200/ecc200_cpu_light.seq'
open_framegrabber ('File', 1, 1, 0, 0, 0, 0, 'default', -1, 'default', -1, \
                  'default', SequenceName, 'default', -1, -1, FGHandle)
```

Then, a model for an ECC 200 is created with `create_data_code_2d_model`.

```
create_data_code_2d_model ('ECC200', [], [], DataCodeHandle)
```

To get the optimal parameters for finding the data code, two sample images are loaded and passed to `find_data_code_2d` with the parameter value 'train'.

```
grab_image (Image, FGHandle)
find_data_code_2d (Image, SymbolXLDs, DataCodeHandle, 'train', 'all', \
                  ResultHandles, DecodedDataStrings)
grab_image (Image, FGHandle)
find_data_code_2d (Image, SymbolXLDs, DataCodeHandle, 'train', 'all', \
                  ResultHandles, DecodedDataStrings)
```

As a preparation for the detection loop, the settings for the find operator are changed with `set_data_code_2d_param` to enable the saving of visual data.

```
set_data_code_2d_param (DataCodeHandle, 'persistence', 1)
```

Inside the while-loop, images are grabbed and `find_data_code_2d` is applied to read the code from the image. After that, iconic results and the read code are visualized. The regions with the foreground and background modules are requested with `get_data_code_2d_objects`.

```
while (1)
    grab_image (Image, FGHandle)
    dev_display (Image)
    find_data_code_2d (Image, SymbolXLDs, DataCodeHandle, [], [], \
        ResultHandles, DecodedDataStrings)
    dev_set_color ('yellow')
    dev_display (SymbolXLDs)
    set_tposition (WindowHandle, 430, 80)
    write_string (WindowHandle, DecodedDataStrings)
    get_data_code_2d_objects (Foreground, DataCodeHandle, ResultHandles[0], \
        'module_1_rois')
    get_data_code_2d_objects (Background, DataCodeHandle, ResultHandles[0], \
        'module_0_rois')

    dev_set_color ('cyan')
    dev_display (Foreground)
    dev_set_color ('magenta')
    dev_display (Background)
endwhile
```

At the end of the program the model for the data code reader is destroyed and the connection to the image sequence is closed.

```
clear_data_code_2d_model (DataCodeHandle)
close_framegrabber (FGHandle)
```

16.3.2 Reading 2D Data Codes on Chips

Example: `hdevelop/Applications/Data-Codes/ecc200_optimized.hdev`

This example program reads 2D data codes of type ECC200, which like the example described before are engraved in chips (see [figure 16.3](#)).

The example shows how to set optimized parameters for efficient data code reading. The code printed on a chip is always light on dark in this application and has a given size and number of modules. Also, the contrast is within a predefined range. By specifying these values for the model, the execution time can be sped up significantly.



Figure 16.3: Decoded data code.

```
create_data_code_2d_model ('Data Matrix ECC 200', [], [], DataCodeHandle)
set_data_code_2d_param (DataCodeHandle, ['module_size_min', \
                                         'module_size_max'], [4,7])
set_data_code_2d_param (DataCodeHandle, 'module_gap', 'no')
set_data_code_2d_param (DataCodeHandle, 'polarity', 'light_on_dark')
set_data_code_2d_param (DataCodeHandle, 'mirrored', 'no')
set_data_code_2d_param (DataCodeHandle, 'contrast_min', 10)
set_data_code_2d_param (DataCodeHandle, 'symbol_size', 18)

find_data_code_2d (Image, SymbolXLDs, DataCodeHandle, [], [], ResultHandles, \
                  DecodedDataStrings)
```

16.4 Selecting Operators

16.4.1 Acquire Image(s)

Please refer to the [operator list for the method Image Acquisition](#) (see section 2.4 on page 17).

16.4.2 Rectify Image(s)

Operators for rectifying images are described in the [Solution Guide III-C](#).

16.4.3 Create Data Code Model

Standard:

[create_data_code_2d_model](#)

16.4.4 Optimize Model

Standard:

`set_data_code_2d_param, get_data_code_2d_param`

Advanced:

`query_data_code_2d_params`

16.4.5 Train Model

Standard:

`find_data_code_2d`

Advanced:

`get_data_code_2d_results, set_data_code_2d_param`

16.4.6 Use Region Of Interest

Please refer to the [operator list for the method Region Of Interest](#) (see section 3.4 on page 26).

16.4.7 Read Data Code(s)

Standard:

`find_data_code_2d`

16.4.8 Inspect Data Code(s)

Standard:

`get_data_code_2d_results, get_data_code_2d_objects`

16.4.9 Check Print Quality

Standard:

`get_data_code_2d_results`

16.4.10 Visualize Results

Please refer to the [operator list for the method Visualization](#) (see section 19.4 on page 297).

16.4.11 Destroy Data Code Model

Standard:

```
clear_data_code_2d_model
```

16.5 Advanced Topics

16.5.1 Use Timeout

With the operator `set_data_code_2d_param`, you can set a timeout. Then, the operator `find_data_code_2d` will return at the latest after the specified time. This mechanism is demonstrated for [Matching \(Shape-Based\)](#) in the example `hdevelop/Matching/Shape-Based/set_shape_model_timeout.hdev`.

Chapter 17

OCR

Optical Character Recognition (OCR) is the technical term for reading, i.e., identifying symbols. In HALCON, OCR is defined as the task to assign an interpretation to regions of an image. These regions typically represent single characters and therefore we consider this as reading single symbols.

In an offline phase, the characters are trained by presenting several samples for each character. In the online phase, the image is segmented to extract the regions representing the characters and then the OCR reader is applied to get the interpretation for each character.

[Figure 17.1](#) shows the principal steps. The first part is offline and consists of collecting training samples and, after that, applying the training. The online part consists of extracting the characters and then reading them.

The advantage of OCR is the flexibility of the training, which allows to select features optimized for an application. Furthermore, the used classifiers are based on the latest neural network technology or on support vector machines, providing both best possible performance.

As a further advantage, HALCON provides you with a set of pretrained fonts, which are based on a large amount of training data from various application areas. These fonts allow you to read text in documents, on pharmaceutical or industrial products, dot prints, and even handwritten numbers. Furthermore, HALCON includes pretrained fonts for OCR-A and OCR-B.

17.1 Basic Concept

The OCR is split into two major parts: training and reading. Each of these major parts requires additional preparation steps:

17.1.1 Acquire Image(s)

Both for the generation of training data and for the OCR itself images must be acquired.

For detailed information see the [description of this method](#) on page 13.

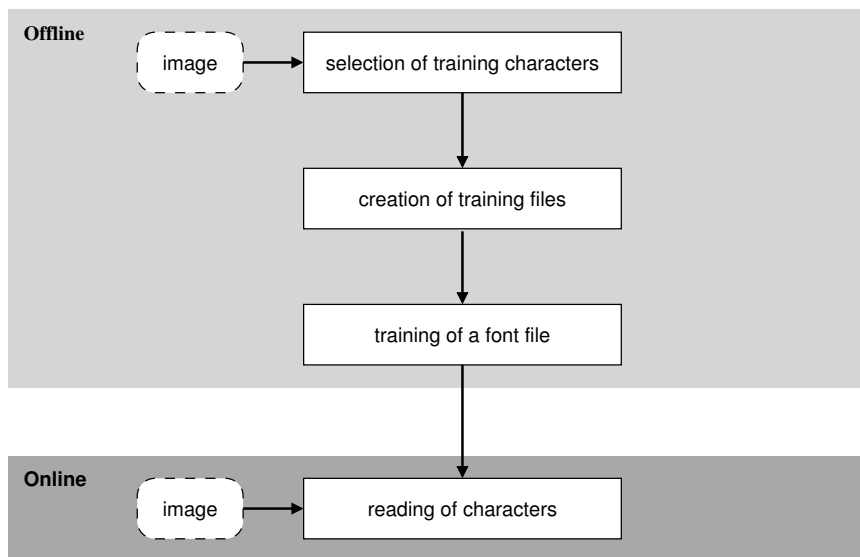



Figure 17.1: Main steps of OCR.



17.1.2 Segment Image(s)

Both for the training samples and for the online reading process, characters must be extracted from the image. This step is called segmentation. This means that the OCR operators like `do_ocr_single_class_svm` do not search for the characters within a given region of interest, but

expect a segmented region, which then will be classified.

If the samples for training are taken from real application images, the same segmentation method will be applied for both training and reading. If the training images are more “artificial”, a simpler method might be used to segment the training images.

17.1.3 Train OCR

The training consists of two important steps: First, a number of samples for each character is selected and stored in so-called training files. In the second step, these files are input for a newly created OCR classifier.

As already noted, HALCON provides pretrained fonts, i.e., ready-to-use classifiers, which already solve many OCR applications. These fonts can be found in the subdirectory `ocr` of the folder where you installed HALCON.

17.1.4 Read Symbol

For reading, you only have to read the trained classifier from disk, segment the characters from the image, and use the segmented characters as input for one of the reading operators that will be discussed later.

17.1.5 Destroy Classifier

When you no longer need the classifier, you destroy it with the operator `clear_ocr_class_mlp` or `clear_ocr_class_svm`, dependent on the used classifier.

17.1.6 A First Example

An example for this basic concept is the following program, which uses one of the pretrained fonts provided by HALCON to read the numbers in the image depicted in [figure 17.2](#).




Figure 17.2: Applying a pretrained classifier.

First, the pretrained font Document_0-9 is read using `read_ocr_class_mlp`. As no file extension is specified, it is searched for a file with the MLP specific extension “.omc” (since HALCON 8.0) or with the extension “.fnt”, which was common for box and MLP classifiers in earlier HALCON versions.

```
FontFile := 'Document_0-9'
read_ocr_class_mlp (FontFile, OCRHandle)
```

Then, the numbers are segmented using `threshold` and `connection`. Because the order is irrelevant here, no further processing is applied.

```
threshold (Image, Region, 0, 125)
connection (Region, Characters)
```

Finally, the numbers are read in a for-loop. The operator `do_ocr_single_class_mlp` takes the single region, the image, and the OCR handle as input. As a result the best and the second best interpretation together with the confidences are returned.

```
count_obj (Characters, Number)
dev_set_color ('white')
for i := 1 to Number by 1
    select_obj (Characters, SingleChar, i)
    do_ocr_single_class_mlp (SingleChar, Image, OCRHandle, 2, Class, \
                           Confidence)
endfor
```

17.2 Extended Concept

When we look more closely at the OCR, many possibilities for adaptation to specific applications become apparent. For example, we have to consider an efficient way of collecting samples as well as correct parameters for the training. In the online phase, an optimal segmentation method is required to extract all characters in a robust manner.

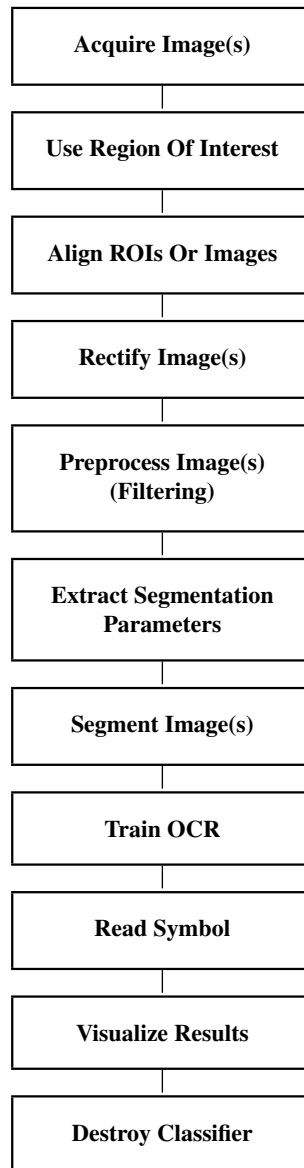
17.2.1 Use Region Of Interest

The OCR can be sped up by using a region of interest. The more the region in which the characters are searched can be restricted, the faster and more robust the search will be.

For detailed information see the [description of this method](#) on page 19.

17.2.2 Align ROIs Or Images

Because the reading of characters is not invariant to rotation, it may be necessary to correct the orientation of the image. This can be achieved either by directly determining the orientation of the text using



the operator `text_line_orientation`, or by locating another object. Then, the part of the image containing the characters is cropped and aligned using the orientation of the text or the found object.

How to perform alignment using shape-based matching is described in the Solution Guide II-B in [section 2.4.3.2](#) on page 42.

17.2.3 Rectify Image(s)

Similarly to alignment, it may be necessary to rectify the image, e.g., to remove perspective distortions. For slanted characters, the slant angle can be determined using the operator [text_line_slant](#).

Detailed information about rectifying images can be found in the Solution Guide III-C in [section 3.4](#) on page 62.

17.2.4 Preprocess Image(s) (Filtering)

Sometimes, the characters may be difficult to extract because of noise, texture, or overlaid structures. Here, operators like [mean_image](#) or [gauss_image](#) can be used to eliminate noise. A fast but slightly less perfect alternative to [gauss_image](#) is [binomial_filter](#). The operator [median_image](#) is helpful for suppressing small spots or thin lines. The operator [dots_image](#) is optimized to emphasize a dot-print while suppressing other structures in the image. Gray value morphology can be used to eliminate noise structures and to adapt the stroke width of characters.

17.2.5 Extract Segmentation Parameters

When segmenting the characters with blob analysis, instead of using fixed threshold values, the values can be extracted dynamically for each image. For more details, please refer to the [description of this step](#) on page 34. Another possibility is the use of the Smart Text Finder, which is described below.

17.2.6 Segment Image(s)

For the segmentation various methods can be used. Using the **Smart Text Finder** enables you to precisely specify a text model and therefore leads to a very robust segmentation. Two other common methods are described under '**General Character Segmentation**'.

Segmentation with the Smart Text Finder

In order to segment images more robustly, the Smart Text Finder allows to specify text features very intuitively. Please view [figure 17.3](#) for overview that also guides you through a segmentation with the Smart Text Finder. For a more practical introduction to the Smart Text Finder, please refer to the HDevelop examples `hdevelop/Applications/OCR/find_text_bottle_label.hdev` or `hdevelop/Applications/OCR/find_text_dongle.hdev` and the corresponding example description [page 261](#).

With the Smart Text Finder, a model, is created first with [create_text_model](#). Segmentation parameters can then be specified with the operator [set_text_model_param](#) and can be queried with [get_text_model_param](#). After this preparation, the text can be segmented with the operator [find_text](#). This operator selects character candidates based on region and gray-value features. These are then further combined to lines which are subsequently tested if they qualify as a text line.

If the text shall be fitted to a line structure, the parameter '[text_line_structure](#)' of the operator [set_text_model_param](#) can be set, which determines the structure, i.e., the number of characters for

each character block of the text that shall be detected. If the text is not horizontally aligned, the operators `text_line_orientation` and `rotate_image` can be used. If the text also includes interpunctuation and separators, set the parameters `'return_interpunctuations'` and `'return_separators'` with the operator `set_text_model_param`.

The result is returned in `TextResultID` which can be queried with `get_text_result` and `get_text_object`, respectively. `get_text_result` queries the `ResultName` of `TextResultID` as control value. `get_text_object` queries the iconic value of text segmentation. To delete the result and the text model, use the operators `clear_text_result` and `clear_text_model`, respectively.

Please refer to the Reference Manual for more information to parameters and values of the operators.

General Character Segmentation

For the common character segmentation you can either use the operator `segment_characters` to get a region containing all character candidates and then apply `select_characters` to select those parts of the region that are candidates for the individual characters, or you use blob analysis. There, the most simple method is the operator `threshold`, with one or more gray value ranges specifying the regions that belong to the foreground objects. Another very common method is `dyn_threshold`. Here, a second image is passed as a reference. With this approach a local instead of a global threshold is used for each position. Further information can be found in the [description of this step for blob analysis](#) on page 32.

17.2.7 Train OCR

[Figure 17.4](#) shows an overview on the generation of the training files: First, the characters from sample images must be extracted using a segmentation method (see above). To each of the single characters a name must be assigned. This can be done either by typing it in, by a programmed input in the case of a well-structured image (having, e.g., multiple samples of each character in different lines), or by reading the character names from file. Then, the regions together with their names are written into training files. The most convenient operator to do this is `append_ocr_trainf`. Before applying the training, we recommend to check the correctness of the training files. This can, e.g., be achieved by using the operator `read_ocr_trainf` combined with visualization operators.

Note that you can also train your own system fonts. By altering and distorting the characters of a font, you can increase the number of different training samples for each class and thus also increase the detection rate. `hdevelop/Applications/OCR/gen_training_file.hdev` is an example program for the training of system fonts.

The actual training is depicted in [figure 17.5](#). First, a new classifier is created. There are three different OCR classifiers available, the box classifier, a neural network (multi-layer perceptron or MLP) classifier, and a classifier that is based on support vector machines (SVM). Note that the box classifier is still available, but not recommended anymore as MLP and SVM are more powerful. Thus, in this manual only MLP and SVM are described further. Dependent on the chosen classifier, you create the classifier using `create_ocr_class_mlp` or `create_ocr_class_svm`. Then, the training is applied using `trainf_ocr_class_mlp` or `trainf_ocr_class_svm`, respectively. After the training, you typically save the classifier to disk for later use by `write_ocr_class_mlp` or `write_ocr_class_svm`.

The difference between the two recommended classifiers is very simple: The MLP classifier is fast at classification, but for a large training set slow at training (compared to the classifier based on SVM). If

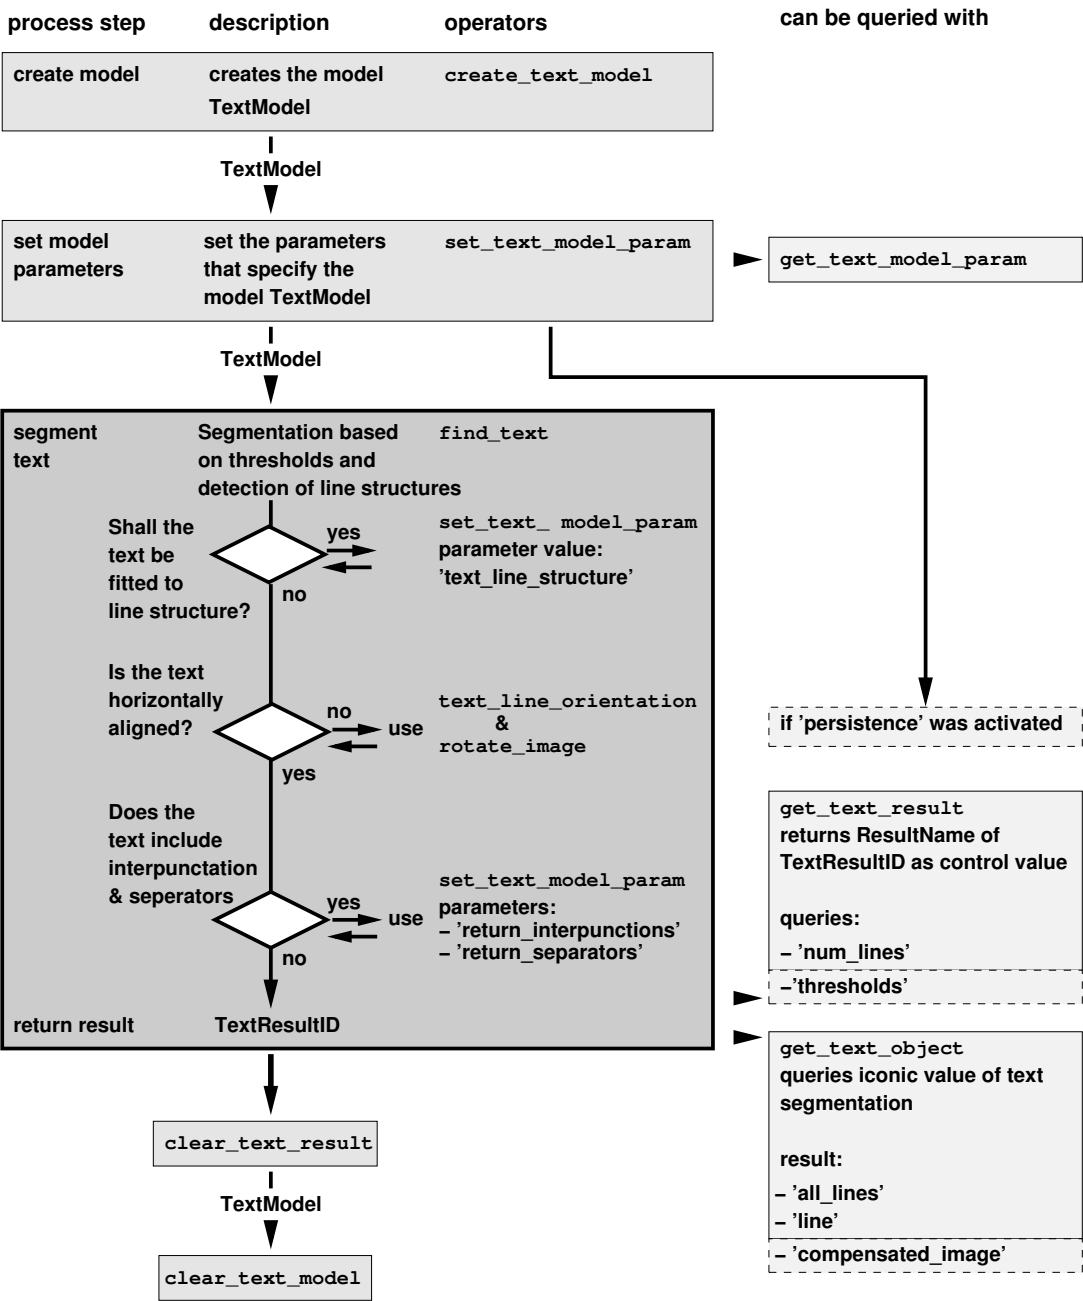


Figure 17.3: Image segmentation with the Smart Text Finder.

the training can be applied offline and thus is not time critical, MLP is a good choice. The classifier based on SVM leads to slightly better recognition rates than the MLP classifier and is faster at training

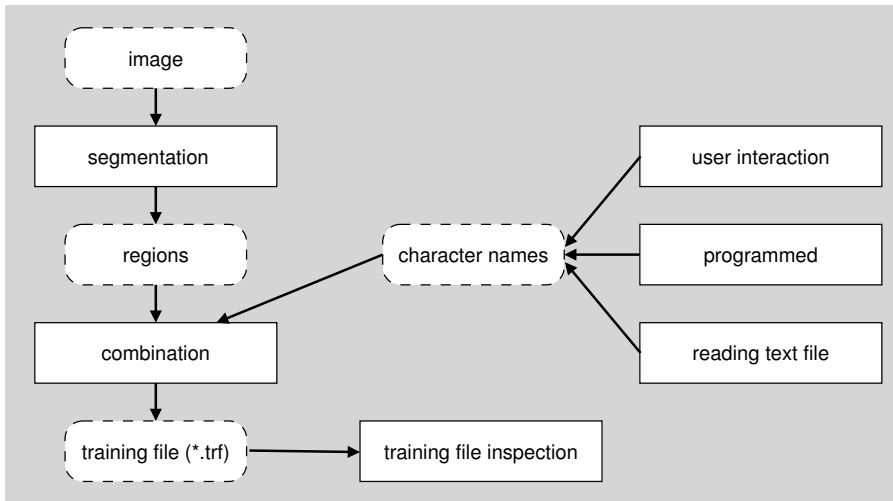


Figure 17.4: Creating training files.

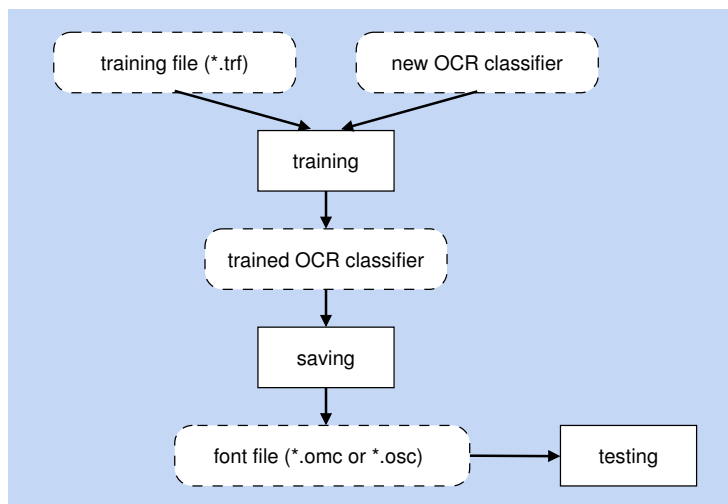


Figure 17.5: Training an OCR classifier.

(especially for large training sets). But compared to the MLP classifier the classification need more time.

17.2.8 Read Symbol

Figure 17.6 shows an overview on the reading process. First, the characters must be extracted using an appropriate segmentation method. Here, you must use a method that returns the characters in a form similar to the ones used for training. After reading the classifier (font file) from file ([read_ocr_class_mlp](#)

or `read_ocr_class_svm`), the classifier can be used for reading.

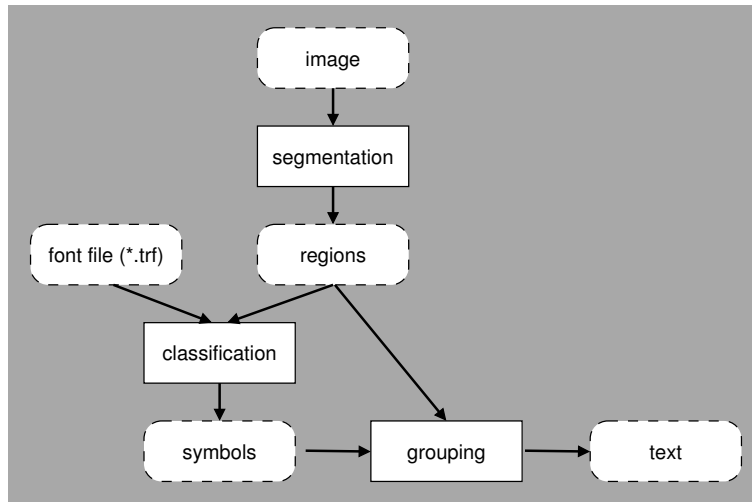


Figure 17.6: Reading characters.

For reading multiple operators are provided: In the easiest case, multiple characters are passed to the reading operators (`do_ocr_multi_class_mlp` or `do_ocr_multi_class_svm`). Here, for each region the corresponding name and the confidence are returned. Sometimes, it can be necessary not only to obtain the characters with the highest confidence but also others with lower confidences. A zero, e.g., might easily be mistaken for the character “O”. This information is returned by the operators `do_ocr_single_class_mlp` and `do_ocr_single_class_svm`.

As a final step it might be necessary to group digits to numbers or characters to words. This can be realized with the region processing operators like those described for the method [blob analysis](#) on page 35.

Additionally, HALCON provides operators for a syntactic and lexicon-based auto-correction. For example, you can use the operator `do_ocr_word_mlp` instead of `do_ocr_multi_class_mlp` to find sets of characters, i.e., words, that match a regular expression or that are stored in a lexicon, which was created by `create_lexicon` or imported by `import_lexicon`.

17.2.9 Visualize Results

Finally, you might want to display the images, the blob (regions), and the result of the reading process.

For detailed information see the [description of this method](#) on page 289.

17.3 Programming Examples

This section gives a brief introduction to using HALCON for OCR. All important steps from training file generation over training to reading are presented.

17.3.1 Generating a Training File

Example: `solution_guide/basics/gen_training_file.hdev`

Figure 17.7 shows a training image from which the characters in the fourth line are used as training samples. For this example image, the segmentation is very simple because the characters are significantly darker than the background. Therefore, `threshold` can be used.




Figure 17.7: Collecting characters for a training file.

The number of the line of characters that is used for training is specified by the variable `TrainingLine`. To select this line, first the operator `closing_rectangle1` is used to combine characters horizontally into lines. These lines are then converted to their connected components with `connection`. Out of all lines the relevant one is selected using `select_obj`. By using `intersection` with the original segmentation and the selected line as input, the characters for training are returned. These are sorted from left to right, using `sort_region`.

```
TrainingLine := 3
threshold (Image, Region, 0, 125)
closing_rectangle1 (Region, RegionClosing, 70, 10)
connection (RegionClosing, Lines)
select_obj (Lines, Training, TrainingLine)
intersection (Training, Region, TrainingChars)
connection (TrainingChars, ConnectedRegions)
sort_region (ConnectedRegions, SortedRegions, 'first_point', 'true', \
            'column')
```

Now, the characters can be stored in the training file. As a preparation step a possibly existing older training file is deleted. Within a loop over all characters the single characters are selected. The variable `Chars` contains the names of the characters as a tuple of strings. With the operator `append_ocr_trainf` the selected regions, together with the gray values and the corresponding name, are added to the training file.

```

Chars := ['0','1','2','3','4','5','6','7','8','9']
TrainFile := 'numbers.trf'
dev_set_check ('~give_error')
delete_file (TrainFile)
dev_set_check ('give_error')
for i := 1 to 10 by 1
    select_obj (SortedRegions, TrainSingle, i)
    append_ocr_trainf (TrainSingle, Image, Chars[i-1], TrainFile)
endfor

```

17.3.2 Creating and Training an OCR Classifier

Example: `solution_guide/basics/simple_training.hdev`

Having prepared the training file, the creation and training of an OCR classifier is very simple. First, the names of the training file and the final font file are determined. Typically, the same base with different extensions is used. We recommend to use “.trf” for training files. For font files i.e., for OCR classifiers, we recommend to use “.obc” for the box classifier (which is not recommended anymore), “.omc” for the neural network classifier, and “.osc” for the classifier based on support vector machines. If no extension is specified during the reading process, for the box or neural network classification it is also searched for files with the extension “.fnt”, which was common for both classifiers in earlier HALCON versions.

To create an OCR classifier, some parameters need to be determined. The most important one is the list of all possible character names. This list can easily be extracted from the training file by using the operator `read_ocr_trainf_names`.

```

TrainFile := 'numbers.trf'
read_ocr_trainf_names (TrainFile, CharacterNames, CharacterCount)

```

Another important parameter is the number of nodes in the hidden layer of the neural network. In this case, it is set to 20. As a rule of thumb, this number should be in the same order as the number of different symbols. Besides these two parameters, here only default values are used for `create_ocr_class_mlp`. The training itself is applied using `trainf_ocr_class_mlp`. We recommend to simply use the default values here as well.

```

NumHidden := 20
create_ocr_class_mlp (8, 10, 'constant', 'default', CharacterNames, \
    NumHidden, 'none', 1, 42, OCRHandle)
trainf_ocr_class_mlp (OCRHandle, TrainFile, 200, 1, 0.01, Error, ErrorLog)

```

Finally, the classifier is stored to disk and the memory is freed.

```

FontFile := 'numbers.omc'
write_ocr_class_mlp (OCRHandle, FontFile)
clear_ocr_class_mlp (OCRHandle)

```

17.3.3 Reading Numbers

Example: `solution_guide/basics/simple_reading.hdev`

This example works similar to the first example described for the basic concept, i.e., it uses a font file to read the numbers in the image depicted in [figure 17.2](#). Instead of a pretrained font that is provided by HALCON, here the font trained in the previous example is used. That means, this time the file `numbers.omc` instead of the file `Document_0-9.omc` is used. The rest of the program corresponds to the first example.

```
FontFile := 'numbers.omc'
read_ocr_class_mlp (FontFile, OCRHandle)
```

17.3.4 "Best Before" Date

Example: `hdevelop/Applications/OCR/bottle.hdev`

The task of this example is to inspect the "best before" date on the bottle depicted in [figure 17.8](#).



Figure 17.8: (a) Original image; (b) read date.

The task is solved in multiple steps. First, dark areas are extracted and post-processed to eliminate structures that are too thin.

```
threshold (Bottle, RawSegmentation, 0, 95)
fill_up_shape (RawSegmentation, RemovedNoise, 'area', 1, 5)
opening_circle (RemovedNoise, ThickStructures, 2.5)
fill_up (ThickStructures, Solid)
```

Next, the region is split into the individual characters; even the characters that are so close that they touch each other can be separated.


```

opening_rectangle1 (Solid, Cut, 1, 7)
connection (Cut, ConnectedPatterns)
intersection (ConnectedPatterns, ThickStructures, NumberCandidates)
select_shape (NumberCandidates, Numbers, 'area', 'and', 300, 9999)
sort_region (Numbers, FinalNumbers, 'first_point', 'true', 'column')

```

Finally, the actual reading is performed.

```

read_ocr_class_mlp (FontName, OCRHandle)
do_ocr_multi_class_mlp (FinalNumbers, Bottle, OCRHandle, RecNum, Confidence)

```

17.3.5 Reading Engraved Text

Example: hdevelop/Applications/OCR/engraved.hdev

The task of this example is to read the engraved text on the metal surface depicted in [figure 17.9](#).




Figure 17.9: (a) Original image; (b) read characters.

The segmentation is solved by using advanced blob analysis: The characters cannot simply be extracted by selecting dark or light pixels. Instead, a simple segmentation would yield only fractions of the characters together with noise objects. Preprocessing the image using gray value morphology allows to segment the real characters.

```

gray_range_rect (Image, ImageResult, 7, 7)
invert_image (ImageResult, ImageInvert)
threshold (ImageResult, Region, 128, 255)
connection (Region, ConnectedRegions)
select_shape (ConnectedRegions, SelectedRegions, 'area', 'and', 1000, 99999)
sort_region (SelectedRegions, SortedRegions, 'first_point', 'true', \
            'column')

```

Finally, the actual reading is performed.

```

read_ocr_class_mlp (FontName, OCRHandle)
for i := 1 to Number by 1
    select_obj (SortedRegions, ObjectSelected, i)
    do_ocr_single_class_mlp (ObjectSelected, ImageInvert, OCRHandle, 1, \
                            Class, Confidence)
endfor
clear_ocr_class_mlp (OCRHandle)

```

17.3.6 Reading Forms

Example: hdevelop/Applications/OCR/ocrcolor.hdev

The task of this example is to extract the symbols in the form. A typical problem is that the symbols are not printed in the correct place, as depicted in [figure 17.10](#).




Figure 17.10: Example images for OCR.

To solve the problem of numbers printed on lines, color is used here: The hue value of the characters differs from the hue of the form. The color classification method is a very simple way to save execution time: In contrast to more difficult color processing problems, here it is sufficient to consider the difference of the red and green channel combined with the intensity.

```

threshold (Green, ForegroundRaw, 0, 220)
sub_image (RedReduced, GreenReduced, ImageSub, 2, 128)
mean_image (ImageSub, ImageMean, 3, 3)
bin_threshold (ImageMean, Cluster1)
difference (Foreground, Cluster1, Cluster2)
concat_obj (Cluster1, Cluster2, Cluster)
opening_circle (Cluster, Opening, 2.5)

```

The selected pixels are grouped and post-processed with morphological operators.

```

closing_rectangle1 (NumberRegion, NumberCand, 1, 20)
difference (Image, NumberCand, NoNumbers)
connection (NumberRegion, NumberParts)
intensity (NumberParts, Green, MeanIntensity, Deviation)
expand_gray_ref (NumberParts, Green, NoNumbers, Numbers, 20, 'image', \
    MeanIntensity, 48)
union1 (Numbers, NumberRegion)
connection (NumberRegion, Numbers)

```

For reading, it is important not to use the gray values of the background because of the changes in color. To solve this, only region features are used for the font, i.e., the regions are used to create an artificial image that is printed dark on light.

```

paint_region (NoNumbers, Green, ImageOCRRaw, 255, 'fill')
paint_region (NumberRegion, ImageOCRRaw, ImageOCR, 0, 'fill')

```

The actual character classification is performed in the artificial image.

```

read_ocr_class_mlp ('Industrial_0-9', OCRHandle)
do_ocr_multi_class_mlp (FinalNumbers, ImageOCR, OCRHandle, RecChar, \
    Confidence)
clear_ocr_class_mlp (OCRHandle)

```

17.3.7 Segment and Select Characters

17.3.7.1 Segment Rotated Characters

Example: `hdevelop/OCR/Segmentation/select_characters.hdev`

This example shows how to easily segment the characters of a rotated dot print using the segmentation operators that are provided especially for OCR (see [figure 17.11](#)).

First, the image is read from file. As the print is rotated, the orientation of the text line is determined via `text_line_orientation`. The obtained angle is used to rotate the image so that the print becomes horizontal.




Figure 17.11: (a) Original image; b) selected characters.

```
read_image (Image, 'dot_print_rotated/dot_print_rotated_'+J$'02d')
text_line_orientation (Image, Image, 50, rad(-30), rad(30), \
    OrientationAngle)
rotate_image (Image, ImageRotate, -OrientationAngle/rad(180)*180, \
    'constant')
```

Then, the operators `segment_characters` and `select_characters` are applied to first segment the region of the complete print and then select those parts of the region that are candidates for individual characters. In contrast to a classical segmentation using blob analysis, here the regions of the individual characters are found although they still consist of components that are not connected.

```
segment_characters (ImageRotate, ImageRotate, ImageForeground, \
    RegionForeground, 'local_auto_shape', 'false', 'true', \
    'medium', 25, 25, 0, 10, UsedThreshold)
select_characters (RegionForeground, RegionCharacters, 'true', \
    'ultra_light', 60, 60, 'false', 'false', 'none', 'true', \
    'wide', 'true', 0, 'completion')
```

The extracted regions can now be used for an OCR application like those described above.

17.3.7.2 Segment Characters with the Smart Text Finder

Example: `hdevelop/Applications/OCR/find_text_dongle.hdev`

This example demonstrates how to segment characters printed in dot print on a dongle with the operator `find_text` before performing OCR. For more background information, please refer to the section about the segmentation with the Smart Text Finder in [section 17.2.6](#) on page 250 and the figure [figure 17.3](#).

First, image and classifier are read.

```
read_image (Image, 'ocr/dongle_01')
read_ocr_class_mlp ('DotPrint', OCRHandle)
```

Then a text model (TextModel) is created with the operator `create_text_model` and the text properties are specified with the operator `set_text_model_param` and TextModel.

```
create_text_model (TextModel)
set_text_model_param (TextModel, 'char_width', 24)
set_text_model_param (TextModel, 'char_height', 32)
set_text_model_param (TextModel, 'is_dotprint', 'true')
set_text_model_param (TextModel, 'max_line_num', 2)
set_text_model_param (TextModel, 'return_interpunctations', 'false')
set_text_model_param (TextModel, 'return_separators', 'false')
set_text_model_param (TextModel, 'stroke_width', 4)
set_text_model_param (TextModel, 'eliminate_horizontal_lines', 'true')
```

The structures of the text lines are defined with the parameter 'text_line_structure'. E.g., 'text_line_structure' '6 1 8' means that the text has three blocks consisting of 6, 1 and 8 character(s). In order to define more than one structure, an index number can be added to the parameter name like, e.g., 'text_line_structure_2'.

Note that for the second line two structures are defined, because sometimes the '/' is classified as separator and sometimes as character. Furthermore, to increase the robustness of the character recognition, a regular expression is defined, which will later be used by `do_ocr_word_mlp`.

```
set_text_model_param (TextModel, 'text_line_structure_0', '6 1 8')
set_text_model_param (TextModel, 'text_line_structure_1', '8 10')
set_text_model_param (TextModel, 'text_line_structure_2', '19')
TextPattern1 := '(FLEXID[0-9][A-Z][0-9]{3}[A-F0-9]{4})'
TextPattern2 := '([A-Z]{3}[0-9]{5}.?[A-Z][0-9]{4}[A-Z][0-9]{4})'
Expression := TextPattern1 + '|' + TextPattern2
```

For preprocessing the domain is reduced to the dark area where the text is assumed to be found.

```
bin_threshold (Image, Region)
opening_rectangle1 (Region, RegionOpening, 400, 50)
connection (RegionOpening, ConnectedRegions)
select_shape_std (ConnectedRegions, SelectedRegion, 'max_area', 70)
reduce_domain (Image, SelectedRegion, ImageReduced)
```

The contrast is then improved with `scale_image_max` and the image is horizontally aligned.

```
scale_image_max (ImageReduced, ImageScaleMax)
text_line_orientation (SelectedRegion, ImageScaleMax, 30, rad(-30), rad(30), \
    OrientationAngle)
rotate_image (ImageScaleMax, ImageRotate, deg(-OrientationAngle), \
    'constant')
```




Figure 17.12: Reading characters on a dongle with the Smart Text Finder.

The text is found and the results are displayed for every segmented region. The OCR uses regular expressions to read the text more robustly.

`get_text_result` returns the number of lines with 'num_lines'. It can also be used to query 'thresholds' if the parameter 'persistence' of the operator `set_text_model_param` was activated. `get_text_object` can be used to return 'all_lines' or in this case with 'line', it queries specific lines. The operator also returns 'compensated_image' if 'persistence' was activated.

```
get_text_result (TextResult, 'num_lines', NumLines)
dev_display (ImageRotate)
for J := 0 to NumLines - 1 by 1
    get_text_object (Line, TextResult, ['line', J])
    do_ocr_word_mlp (Line, ImageRotate, OCRHandle, Expression, 3, 5, Class, \
                    Confidence, Word, Score)
```

The results are then displayed.

```
smallest_rectangle1 (Line, Row1, Column1, Row2, Column2)
count_obj (Line, NumberOfCharacters)
dev_set_colored (6)
dev_display (Line)
dev_set_color ('white')
for K := 1 to NumberOfCharacters by 1
    select_obj (Line, Character, K)
    set_tposition (WindowHandle, Row2[0]+4, Column1[K-1])
endfor
```

Finally, text result, model and handle are deleted.

```
clear_text_result (TextResult)
clear_text_model (TextModel)
clear_ocr_class_mlp (OCRHandle)
```

17.3.8 Syntactic and Lexicon-Based Auto-Correction of OCR Results

Example: `hdevelop/OCR/Neural-Nets/label_word_process_mlp.hdev`

This example reads the "best before" date depicted in [figure 17.13](#). To correct an incorrect OCR result for the upper text line, a lexicon-based auto-correction is used. Errors can occur, e.g., because of the similarity of characters, e.g., between the character O and the number 0. For the second text line, regular expressions are used to ensure that the result has the correct format.

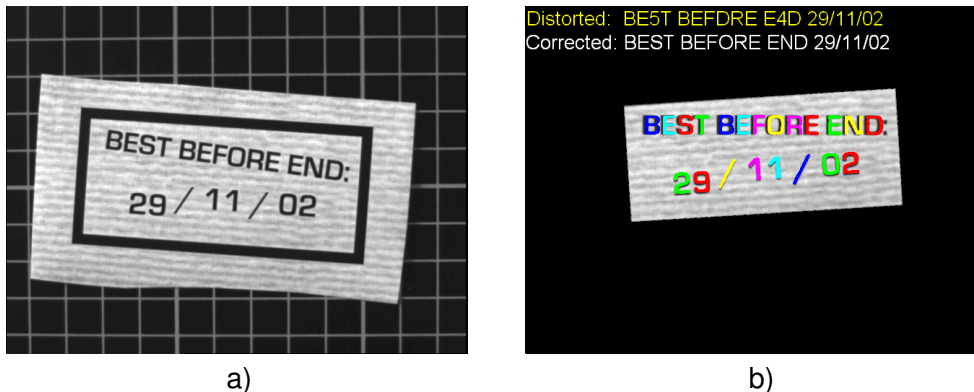


Figure 17.13: (a) Original image; (b) text is corrected using syntactic and lexicon-based auto-correction.

First, the pretrained font `Industrial` is read as preparation for the actual OCR reading. For the upper line of the text, the three expected words are stored in a lexicon that is created with `create_lexicon` and will be used later. Then, the images are read, an ROI for the print is generated and aligned, and the regions for the characters are extracted and stored in the variable `SortedWords` using blob analysis.

```
read_ocr_class_mlp ('Industrial', OCRHandle)
create_lexicon ('label', ['BEST','BEFORE','END'], LexiconHandle)
for i := 1 to 9 by 1
    read_image (Image, 'label/label_0' + i + '.png')
    ... get ROI, align print, and extract regions ...
```

Now, the text line is read twice with the operator `do_ocr_word_mlp`. The first time it is read without syntactic or lexicon-based auto-correction, and the second time the result is corrected by matching it to the lexicon entries.

```

sort_region (CharactersWords, SortedWords, 'character', 'true', 'row')
gen_empty_obj (Word)
Text := ''
OriginalText := ''
for j := 1 to |Column| - 1 by 1
    select_obj (SortedWords, Character, j)
    concat_obj (Word, Character, Word)
    if (j = |Column| or (Column[j] - Column[j-1]) > 30)
        do_ocr_word_mlp (Word, ImageOCR, OCRHandle, '.*', 1, 5, Class, \
            Confidence, WordText, WordScore)
        OriginalText := OriginalText + ' ' + WordText
        do_ocr_word_mlp (Word, ImageOCR, OCRHandle, '<label>', 1, 5, Class, \
            Confidence, WordText, WordScore)
        Text := Text + ' ' + WordText
        gen_empty_obj (Word)
    endif
endfor

```

The second text line, i.e., the actual date, is read using a regular expression that ensures the correct format for the date. This is done to suppress structures that may be extracted as candidates for characters but syntactically do not fit into the searched string.

```

sort_region (CharactersDate, SortedDate, 'character', 'true', 'row')
do_ocr_word_mlp (SortedDate, ImageOCR, OCRHandle, '.*', 5, 5, Class, \
    Confidence, OriginalDateText, DateScore)
do_ocr_word_mlp (SortedDate, ImageOCR, OCRHandle, \
    '^([0-2][0-9]|30|31)/([0-9]|11|12)/0[0-5]$', 5, 5, Class, \
    Confidence, DateText, DateScore)

```


17.4 Selecting Operators

17.4.1 Acquire Image(s)

Please refer to the [operator list for the method Image Acquisition](#) (see section 2.4 on page 17).

17.4.2 Use Region Of Interest

Please refer to the [operator list for the method Region Of Interest](#) (see section 3.4 on page 26).

17.4.3 Align ROIs Or Images

Standard:

[text_line_orientation](#)

Further operators for aligning ROIs or images are described in the [Solution Guide II-B](#).

17.4.4 Rectify Image(s)

Standard:

[text_line_slant](#)

Further operators for rectifying images are described in the [Solution Guide III-C](#).

17.4.5 Preprocess Image(s) (Filtering)

Standard:

[mean_image](#), [gauss_image](#), [binomial_filter](#), [median_image](#), [gray_opening_shape](#),
[gray_closing_shape](#), [dots_image](#), [gray_range_rect](#)

Advanced:

[gray_dilation_shape](#), [gray_erosion_shape](#), [anisotrope_diff](#)

17.4.6 Extract Segmentation Parameters

Please refer to the [detailed operator list for the step Extract Segmentation Parameters](#) on page 44.

17.4.7 Segment Image(s)

Standard:

```
create_text_model, set_text_model_param, get_text_model_param, find_text,  
get_text_result, get_text_object, clear_text_result, clear_text_model,  
clear_all_text_results, clear_all_text_models
```

Advanced:

```
segment_characters, select_characters
```

Further operators can be found in the [detailed operator list for the step Segment Image\(s\)](#) on page 44.

17.4.8 Train OCR

Standard:

```
read_string, append_ocr_trainf, delete_file, read_ocr_trainf,  
create_ocr_class_mlp, trainf_ocr_class_mlp, write_ocr_class_mlp,  
create_ocr_class_svm, trainf_ocr_class_svm, write_ocr_class_svm
```

Advanced:

```
get_prep_info_ocr_class_mlp, get_prep_info_ocr_class_svm
```

17.4.9 Read Symbol

Standard:

```
read_ocr_class_mlp, do_ocr_multi_class_mlp, do_ocr_word_mlp,  
do_ocr_single_class_mlp, read_ocr_class_svm, do_ocr_multi_class_svm,  
do_ocr_word_svm, do_ocr_single_class_svm
```

Advanced:

```
clear_all_lexica, clear_lexicon, create_lexicon, import_lexicon, inspect_lexicon,  
lookup_lexicon, suggest_lexicon, dilation_rectangle1, closing_circle,  
partition_rectangle, partition_dynamic
```

17.4.10 Visualize Results

Please refer to the [operator list for the method Visualization](#) (see section 19.4 on page 297).

17.4.11 Destroy Classifier

Standard:

```
clear_ocr_class_mlp, clear_all_ocr_class_mlp, clear_ocr_class_svm,  
clear_all_ocr_class_svm
```

17.5 Relation to Other Methods

17.5.1 Alternatives to OCR

Matching (see [description](#) on page 97)

As an alternative to classical OCR, matching can be used to read single characters or more complex symbols. In this case, one model for each character must be generated. The advantage of matching is the invariance to rotation. Furthermore, it is not necessary to segment the characters prior to the classification. Therefore, the matching approach should be considered when there is no robust way to separate the characters from the background.

Classification

You can consider the OCR tool as a convenient way of using a classifier. The OCR automatically derives invariant features and passes them to the underlying classifier. If the features offered by the OCR do not fulfill the needs of your application, you can create an “extended” OCR classifier by calculating the features using normal HALCON feature extraction operators and then using them with one of the classifiers that HALCON offers (see the chapters “[Regions > Features](#)” and “[Classification](#)” in the Reference Manual).

17.6 Tips & Tricks

17.6.1 Composed Symbols

Some characters and symbols are composed of multiple sub-symbols, like an “i”, “%”, or “!”. For the OCR, these sub-symbols must be combined into a single region. If you use the operators [segment_characters](#) and [select_characters](#) for the segmentation of the characters, sub-symbols are automatically combined. Otherwise, you can combine them by calling [closing_rectangle1](#) after thresholding, typically using a small width but a larger height. After calling [connection](#) to separate the characters, you use the operator [intersection](#) to get the original segmentation (input parameter 2), while preserving the correct connected components from [connection](#) (input parameter 1).

17.7 Advanced Topics

17.7.1 Line Scan Cameras

In general, line scan cameras are treated like normal area sensors. But in some cases, not single images but an “infinite” sequence of images showing objects, e.g., on a conveyor belt, must be processed. In this case, the end of one image is the beginning of the next one. This means that text or numbers, which partially lie in both images, must be combined into one object. For this purpose, HALCON provides the operator `merge_regions_line_scan`. This operator is called after the segmentation of one image, and combines the current objects with those of previous images. For more information see the [Solution Guide II-A](#).

17.7.2 Circular Prints

In some cases the symbols are not printed as straight lines but along arcs, e.g., on a CD. To read these, the (virtual) center and radius of the corresponding circle are extracted. Using the operator `polar_trans_image_ext`, the image is then unwrapped. To project a region obtained in the unwrapped image back into the original image, you can use the operator `polar_trans_region_inv`.

17.7.3 OCR Features

HALCON offers many different features for the OCR. Most of these are for advanced use only. In most cases it is recommended to use the feature combination 'default'. This combination is based on the gray values within the surrounding rectangle of the character. In case that the background of the characters cannot be used, e.g., if it varies because of texture, the features 'pixel_binary', 'ratio', and 'anisometry' are good combinations. Here, only the region is used, the underlying gray values are ignored.

17.8 Pretrained OCR Fonts

The following sections shortly introduce you to the pretrained OCR fonts provided by HALCON. You can access them in the subdirectory `ocr` of the folder where you installed HALCON. Note that the pretrained fonts were trained with symbols that are printed dark on light. If you want to read light on dark symbols with one of the provided fonts, you can either invert the image with the operator `invert_image`, or, if this does not lead to a satisfying result, preprocess the image by applying first the operator `gen_image_proto` with a light gray value and then `overpaint_region` with the gray value set to 0.

Note that pretrained fonts are not available for the classifier based on SVM.

17.8.1 Nomenclature for the Ready-to-Use OCR Fonts

There are several groups of OCR fonts available. The members of each group differ as they contain different symbol sets. The content of an OCR font is described by its name. For the names of the pretrained OCR fonts the following nomenclature is applied:

The name starts with the group name, e.g., `Document` or `DotPrint`, followed by indicators for the set of symbols contained in the OCR font. The meaning of the indicators is the following:

- 0-9: The OCR font contains the digits 0 to 9.
- A-Z: The OCR font contains the uppercase characters A to Z.
- + : The OCR font contains special characters. The list of special characters varies slightly over the individual OCR fonts. It is given below for each OCR font separately.

If the name of the OCR font does not contain any of the above indicators, typically, the OCR font contains the digits 0 to 9, the uppercase characters A to Z, the lowercase characters a to z, and special characters. Some of the OCR fonts do not contain lowercase characters or the full set of uppercase characters (e.g., MICR). For these OCR fonts, the contained symbols are named explicitly.

17.8.2 Ready-to-Use OCR Font 'Document'

The OCR font `Document` can be used to read characters printed in fonts like Arial, Courier, or Times New Roman. These are typical fonts for printing documents or letters.

Note that the characters I and l of the font Arial cannot be distinguished. That means that an l may be mistaken for an I and vice versa.

Available special characters: - = + < > . # \$ % & () @ *

The following OCR fonts with different symbol sets are available:

- `Document`
- `Document_A-Z+`
- `Document_0-9`
- `Document_0-9A-Z`

17.8.3 Ready-to-Use OCR Font 'DotPrint'

The OCR font `DotPrint` can be used to read characters printed with dot printers (see [figure 17.14](#)).

It contains no lowercase characters.

Available special characters: - / . * :

The following OCR fonts with different symbol sets are available:

- `DotPrint`

- DotPrint_A-Z+
- DotPrint_0-9
- DotPrint_0-9+
- DotPrint_0-9A-Z




Figure 17.14: Examples for dot prints.

17.8.4 Ready-to-Use OCR Font 'HandWritten_0-9'

The OCR font HandWritten_0-9 can be used to read handwritten numbers (see [figure 17.15](#)).

It contains the digits 0-9.

Available special characters: none




Figure 17.15: Examples for handwritten numbers.

17.8.5 Ready-to-Use OCR Font 'Industrial'

The OCR font Industrial can be used to read characters printed in fonts like Arial, OCR-B, or other sans-serif fonts (see [figure 17.16](#)). These fonts are typically used to print, e.g., labels.

Available special characters: - / + . \$ % * €

The following OCR fonts with different symbol sets are available:

- Industrial
- Industrial_A-Z+

- Industrial_0-9
- Industrial_0-9+
- Industrial_0-9A-Z




Figure 17.16: Examples for industrial prints.

17.8.6 Ready-to-Use OCR Font 'MICR'

The OCR font MICR can be used to read characters printed in the font MICR (see [figure 17.17](#)).

It contains the digits 0-9 and the characters A-D.

Available special characters: none



Figure 17.17: The MICR font.

17.8.7 Ready-to-Use OCR Font 'OCR-A'

The OCR font OCR-A can be used to read characters printed in the font OCR-A (see [figure 17.18](#)).

Available special characters: - ? ! / \ { } = + < > . # \$ % & () @ *

The following OCR fonts with different symbol sets are available:

- OCRA
- OCRA_A-Z+
- OCRA_0-9
- OCRA_0-9A-Z

0123456789
 ABCDEFGHIJKLM
 NOPQRSTUVWXYZ
 abcdefghijklm
 nopqrstuvwxy
 z-?!/\+=<>.#\$%&()@*

Figure 17.18: Selected characters of the OCR-A font.

17.8.8 Ready-to-Use OCR Font 'OCR-B'

The OCR font OCR-B can be used to read characters printed in the font OCR-B (see [figure 17.19](#)).

Available special characters: - ? ! / \ { } = + < > . # \$ % & () @ *

The following OCR fonts with different symbol sets are available:

- OCRB
- OCRB_A-Z+
- OCRB_0-9
- OCRB_0-9A-Z
- OCRB_passport (contains digits, uppercase characters, and the special characters + and <)

0123456789
 ABCDEFGHIJKLM
 NOPQRSTUVWXYZ
 abcdefghijklm
 nopqrstuvwxy
 z-?!/\+=<>.#\$%&()@*

Figure 17.19: Selected characters of the OCR-B font.

17.8.9 Ready-to-Use OCR Font 'Pharma'

The OCR font Pharma can be used to read characters printed in fonts like Arial, OCR-B, and other fonts that are typically used in the pharmaceutical industry (see [figure 17.20](#)).

This OCR font contains no lowercase characters.

Available special characters: - / . () :

The following OCR fonts with different symbol sets are available:

- Pharma
- Pharma_0-9
- Pharma_0-9A-Z
- Pharma_0-9+




Figure 17.20: Examples for pharmaceutic labels.

17.8.10 Ready-to-Use OCR Font 'SEMI'

The OCR font 'SEMI'

The OCR font SEMI can be used to read characters printed in the SEMI font which consists of characters which are designed to be easily distinguished from each other. It has a limited set of characters which can be viewed in [figure 17.21](#). This OCR font contains no lowercase characters.

Available special characters: - .

ABCDEFGHIJKLMNO
PQRSTUVWXYZ-
0123456789.

Figure 17.21: Character set of SEMI font.

XB0225066244F5
7ICEN033MMD2
SI165A352110B3

Figure 17.22: Examples for SEMI font. Note that these images were inverted before the training was applied! That is, this font is nevertheless pretrained only for dark on light symbols.

Chapter 18

Stereo Vision

The basic principle of stereo vision is that 3D coordinates of object points are determined from two or more images that are acquired simultaneously from different points of view. HALCON provides two stereo methods: binocular stereo and multi-view stereo.

Binocular stereo uses exactly two cameras and returns disparity images, distance images, or 3D coordinates. [Figure 18.1](#), e.g., shows a stereo image pair of a board and the resulting height map of the board's components.




Figure 18.1: Basic principle of binocular stereo. Top: stereo image pair; Bottom: height map.

Multi-view stereo can also use more than two cameras. It can be used to either reconstruct surfaces that are returned as 3D object models or to reconstruct individual points. [Figure 18.2](#), e.g., shows the images of a multi-view stereo camera system that is used to reconstruct the surface of pipe joints and the resulting 3D object model.




Figure 18.2: Left: images of a multi-view stereo camera system; right: reconstructed surface of pipe joints.

The advantage of stereo vision is that 3D information of the surface of arbitrarily shaped objects can be determined from images. Stereo vision can also be combined with other vision methods, e.g., as a preprocessing step for blob analysis, which can be used to extract specific objects or object parts from the depth image, or surface-based 3D matching, which locates objects that are provided as 3D models in the reconstructed surface.

For detailed information about stereo vision, please refer to the Solution Guide III-C, [chapter 5](#) on page [105](#).

18.1 Basic Concept

The derivation of 3D information with a stereo camera system consists of four main steps:

18.1.1 Acquire Calibration Image(s)

A number of stereo calibration images is acquired. Each image shows the HALCON calibration plate in a different position and orientation.


18.1.2 Calibrate Stereo Camera System

Using the previously acquired calibration images, the stereo camera system is calibrated. For this, the image coordinates of the calibration marks must be extracted from the calibration images. Then, the parameters of the stereo setup are determined.

For **binocular stereo**, the parameters are then used to create the rectification maps for the rectification of the stereo images.

For **multi-view stereo**, a so-called stereo model is created, which contains the camera parameters and further information.

The calibration process is described in detail in the Solution Guide III-C in [section 5.2](#) on page [110](#).



18.1.3 Acquire Stereo Image(s)

Stereo images are simultaneously acquired with the calibrated stereo camera system. They show the object for which the 3D information should be reconstructed.

18.1.4 Rectify Image(s)

For **binocular stereo**, the stereo images must be rectified such that corresponding points (conjugate points) lie on the same row in both rectified images. For this, the rectification map that has been determined above must be used.

18.1.5 Reconstruct 3D Information

Now, 3D information can be reconstructed.


With **binocular stereo**, for each point of the first rectified image the conjugate point in the second rectified image is determined (stereo matching). For these points, either the disparity or the distance to the stereo camera system can be calculated and returned as a gray value image. The reference plane to which these distances are related can be changed (see the Solution Guide III-C, [section 5.3.5.3](#) on page 124). It is also possible to derive the distance or the 3D coordinates for selected points for which the disparity is known. What is more, 3D coordinates can be determined from the image coordinates of each pair of conjugate points directly. See the Solution Guide III-C, [section 5.3.5](#) on page 120, for details.

With **multi-view stereo**, you can reconstruct complete surfaces or 3D coordinates for selected points. When reconstructing surfaces, you first set parameters for the reconstruction with

`set_stereo_model_param`, then you apply `set_stereo_model_image_pairs` to define which cameras build pairs, and finally you reconstruct the surface with `reconstruct_surface_stereo`. When reconstructing selected points of an object, you first extract the corresponding points from the stereo images, then you accumulate the correspondence information for all cameras that image the object, and finally you reconstruct the points with `reconstruct_points_stereo`. See the Solution Guide III-C, section 5.4.2 on page 133, for details.

18.2 Extended Concept

In many cases, the derivation of 3D information with a binocular stereo system involves more steps than described above. Reasons for this are, e.g., the need to restrict the stereo reconstruction to an ROI. Furthermore, postprocessing like the transformation of the 3D coordinates into another coordinate system or the visualization of results is often required.



18.2.1 Use Region Of Interest

A region of interest can be created to reduce the image domain for which the stereo matching will be performed. This will reduce the processing time.

For detailed information see the [description of this method](#) on page 19.

18.2.2 Transform Results Into World Coordinates

In some applications, the 3D coordinates must be transformed into a given world coordinate system.

For **binocular stereo**, this step is performed after the reconstruction: Beforehand, the relation between the given world coordinate system and the stereo camera system must be determined. Then, the 3D coordinates can be transformed as requested. How to transform results into world coordinates is described in detail in the Solution Guide III-C in [section 3.3](#) on page 57.

For **multi-view stereo**, the world coordinate system can be specified already when creating the stereo model (see the Solution Guide III-C, [section 5.4.1.2](#) on page 132 for details).

18.2.3 Visualize Results

Finally, you might want to visualize the disparity or distance images or the reconstructed surface (3D object model). We recommend to have a look at the provided examples.

18.3 Programming Examples

This section gives a brief introduction to using HALCON for stereo vision.

18.3.1 Segment the Components of a Board With Binocular Stereo

Example: `hdevelop/Applications/Object-Recognition-2D/board_components.hdev`

[Figure 18.3](#) shows a stereo image pair of a board together with the result of a segmentation of raised objects. The segmentation has been carried out based on the distance image that was derived with binocular stereo.

First, a calibration data model is created and initialized:

```
create_calib_data ('calibration_object', 2, 1, CalibDataID)
set_calib_data_cam_param (CalibDataID, 'all', 'area_scan_division', \
                          StartCamPar)
set_calib_data_calib_object (CalibDataID, 0, CalDescrFile)
```

Then, a number of calibration images must be acquired. The calibration plate must be completely visible in each calibration image. A subset of these calibration images is shown in [figure 18.4](#).




Figure 18.3: Segment board components based on their height.




Figure 18.4: A subset of the calibration images that are used for the calibration of the stereo camera system.

The calibration plate is located and the image coordinates of the calibration marks are extracted from the individual calibration images and stored in the calibration data model.

```

for imgnr := 1 to Number by 1
    find_caltab (ImageL, CaltabL, CalDescrFile, SizeGauss, MarkThresh, \
                MinDiamMarks)
    find_caltab (ImageR, CaltabR, CalDescrFile, SizeGauss, MarkThresh, \
                MinDiamMarks)
    find_marks_and_pose (ImageL, CaltabL, CalDescrFile, StartCamPar, \
                        StartThresh, DeltaThresh, MinThresh, Alpha, \
                        MinContLength, MaxDiamMarks, RCoordL, CCoordL, \
                        StartPoseL)
    find_marks_and_pose (ImageR, CaltabR, CalDescrFile, StartCamPar, \
                        StartThresh, DeltaThresh, MinThresh, Alpha, \
                        MinContLength, MaxDiamMarks, RCoordR, CCoordR, \
                        StartPoseR)
    set_calib_data_observ_points (CalibDataID, 0, 0, imgnr, RCoordL, \
                                CCoordL, 'all', StartPoseL)
    set_calib_data_observ_points (CalibDataID, 1, 0, imgnr, RCoordR, \
                                CCoordR, 'all', StartPoseL)
endfor

```

With this, the actual calibration of the stereo camera system is performed and the results are accessed.

```

calibrate_cameras (CalibDataID, Error)
get_calib_data (CalibDataID, 'camera', 0, 'params', CamParamL)
get_calib_data (CalibDataID, 'camera', 1, 'params', CamParamR)
get_calib_data (CalibDataID, 'camera', 1, 'pose', cLPcR)

```

Now, the rectification maps for the rectification of the stereo image pair can be generated.

```

gen_binocular_rectification_map (MapL, MapR, CamParamL, CamParamR, cLPcR, 1, \
                                'geometric', 'bilinear', RectCamParL, \
                                RectCamParR, CamPoseRectL, CamPoseRectR, \
                                RectLPosRectR)

```

Then, each stereo image pair acquired with the calibrated stereo camera system can be rectified. This has the effect that conjugate points have the same row coordinate in both images. The rectified images are displayed in [figure 18.5](#)

```

map_image (ImageL, MapL, ImageRectifiedL)
map_image (ImageR, MapR, ImageRectifiedR)

```

From the rectified images, a distance image can be derived in which the gray values represent the distance of the respective object point to the stereo camera system. This step is the core of the stereo approach. Here, the stereo matching, i.e., the determination of the conjugate points takes place.

```

binocular_distance (ImageRectifiedL, ImageRectifiedR, DistanceImage, \
                  ScoreImageDistance, RectCamParL, RectCamParR, \
                  RectLPosRectR, 'ncc', MaskWidth, MaskHeight, \
                  TextureThresh, MinDisparity, MaxDisparity, NumLevels, \
                  ScoreThresh, 'left_right_check', 'interpolation')

```



Figure 18.5: Rectified images.

Finally, the distance image can, e.g., be corrected such that a given object plane receives a specified distance value, e.g., zero. Objects that deviate from the given object plane can thus be segmented very easily with a threshold operation.

```
threshold (HeightAboveReferencePlaneReduced, Range1, Height1_Min, \
          Height1_Max)
threshold (HeightAboveReferencePlaneReduced, Range2, Height2_Min, \
          Height2_Max)
threshold (HeightAboveReferencePlaneReduced, Range3, Height3_Min, \
          Height3_Max)
```

18.3.2 Reconstruct the Surface of Pipe Joints With Multi-View Stereo

Example: `hdevelop/Applications/Robot-Vision/locate_pipe_joints_stereo.hdev`

Figure 18.2 shows images from a 4-camera system. Using multi-view stereo, the surface of the pipe joints is reconstructed.

Here, the calibration is performed in an offline step and the saved camera setup model is just loaded.

```
read_camera_setup_model ('../../Calibration/Multi-View/four_camera_setup_model.csm', \
                        CameraSetupModelID)
```

Using the camera setup model, a stereo model for surface reconstruction is created. Afterwards, the camera setup model is destroyed.

```
create_stereo_model (CameraSetupModelID, 'surface_pairwise', [], [], \
                    StereoModelID)
clear_camera_setup_model (CameraSetupModelID)
```

Now, the stereo model is configured by setting several parameters. Some of them adjust the internal call of [binocular_disparity](#). Others adjust, e.g., the subsampling that is used for the reconstruction or the bounding box that restricts the reconstruction to a specific part of the 3D space.

```

set_stereo_model_param (StereoModelID, 'sub_sampling_step', 3)
set_stereo_model_param (StereoModelID, 'rectif_interpolation', 'bilinear')
set_stereo_model_param (StereoModelID, 'rectif_sub_sampling', 1.2)
set_stereo_model_param (StereoModelID, 'binocular_method', 'ncc')
set_stereo_model_param (StereoModelID, 'binocular_num_levels', 1)
set_stereo_model_param (StereoModelID, 'binocular_mask_width', 19)
set_stereo_model_param (StereoModelID, 'binocular_mask_height', 19)
set_stereo_model_param (StereoModelID, 'binocular_texture_thresh', 0)
set_stereo_model_param (StereoModelID, 'binocular_score_thresh', 0.4)
set_stereo_model_param (StereoModelID, 'binocular_filter', \
    'left_right_check')
set_stereo_model_param (StereoModelID, 'binocular_sub_disparity', \
    'interpolation')
set_stereo_model_param (StereoModelID, 'bounding_box', [-0.2,-0.07,-0.075, \
    0.2,0.07, -0.004])

```

Furthermore, the image pairs are specified. That is, it is specified between which cameras the disparity is computed.

```

set_stereo_model_image_pairs (StereoModelID, [0,2], [1,3])

```

Finally, the reconstruction operator is called. It returns the reconstructed surface in form of a 3D object model.

```

reconstruct_surface_stereo (Images, StereoModelID, PipeJointPile0M3DID)

```

In [section 10.3.2](#) on page 137, the reconstructed surface is then used as input for surface-based 3D matching, which recognizes and locates individual pipe joints.

18.4 Selecting Operators

18.4.1 Calibrate Stereo Camera System

Standard:

```

create_calib_data, set_calib_data_cam_param, set_calib_data_calib_object,
find_caltab, find_marks_and_pose, set_calib_data_observ_points,
calibrate_cameras, get_calib_data, gen_binocular_rectification_map,
create_stereo_model

```

18.4.2 Acquire Image(s)

Please refer to the [operator list for the method Image Acquisition](#) (see [section 2.4](#) on page 17).

18.4.3 Rectify Image(s)

Standard:

[map_image](#)

18.4.4 Use Region Of Interest

Standard:

[reduce_domain](#)

Further operators can be found in the [operator list for the method Region Of Interest](#) (see section 3.4 on page 26).

18.4.5 Reconstruct 3D Information

Standard:

[binocular_disparity](#), [binocular_distance](#), [binocular_disparity_mg](#),
[binocular_distance_mg](#), [disparity_image_to_xyz](#), [disparity_to_distance](#),
[disparity_to_point_3d](#), [distance_to_disparity](#), [intersect_lines_of_sight](#),
[set_stereo_model_param](#), [set_stereo_model_image_pairs](#),
[reconstruct_surface_stereo](#), [reconstruct_points_stereo](#)

18.4.6 Transform Results Into World Coordinates

Operators for transforming results into world coordinates are described in the Solution Guide III-C, [section 5.4.1.2](#) on page 132.

18.4.7 Visualize Results

Please refer to the [operator list for the method Visualization](#) (see section 19.4 on page 297).

18.5 Relation to Other Methods

18.5.1 Methods that are Using Stereo Vision

Blob Analysis (see [description](#) on page 31)

The results of **binocular stereo** can be used as input for blob analysis. This may be useful if locally high structures must be extracted that cannot be reliably detected from the original image. By applying blob analysis to the distance image, such structures may be extracted very easily.

3D Matching (Surface-Based) (see [description](#) on page 127)

The 3D object model that is returned by a surface reconstruction with **multi-view stereo** can be used as input for surface-based 3D matching.

3D Primitives Fitting (see Solution Guide III-C, [section 4.4](#) on page 97)

The 3D object model that is returned by a surface reconstruction with **multi-view stereo** can be used as input for 3D primitives fitting if the contained surface is meshed. To get a meshed surface, the parameter 'point_meshing' must have been set for the stereo model with [set_stereo_model_param](#) before reconstructing the surface.

18.6 Tips & Tricks

18.6.1 Speed Up

Many online applications require maximum speed. Although the stereo matching is a very complex task, you can speed up this process by using regions of interest, which are the standard method to increase the speed by processing only those areas where objects must be inspected. This can be done by using pre-defined regions but also by an online generation of the region of interest that depends on other objects found in the image.

Additionally, for the surface extraction with multi-view stereo, a significant speed up can be obtained when selecting the bounding box that restricts the reconstruction to a part of the 3D space as small as possible. The bounding box must be set with [set_stereo_model_param](#) before applying the reconstruction.

18.7 Advanced Topics

18.7.1 High Accuracy

Sometimes very high accuracy is required. To achieve a high distance resolution, i.e., a high accuracy with which the distance of the object surface from the stereo camera system can be determined, special care should be taken of the configuration of the stereo camera setup. The setup should be chosen such that the distance between the cameras as well as the focal length are large and the stereo camera system is placed as close as possible to the object. For more information, please refer to the Solution Guide III-C, [section 5.1.2](#) on page 110.

Chapter 19


Visualization

Displaying data in HALCON is quite easy: In the graphics windows provided by HALCON, all supported data types can be visualized directly using specific display operators. Both the creation of these windows and the displaying requires only little programming effort because the functionality is optimized for the use in machine vision.

Making use of the HALCON visualization provides several advantages. First of all, it saves a lot of time during the development because all important visualization methods are already predefined. Furthermore, the functionality is independent of the operating system: Writing a program under Windows using the visualization of HALCON can be ported easily to Linux because the visualization operators behave identically and therefore only user code making use of operating system functions needs to be rewritten.

19.1 Basic Concept

For visualization there are two important aspects to consider: The graphics windows and the data that must be visualized.



19.1.1 Handling Graphics Windows

HALCON provides an easy-to-use operator to create a graphics window for visualization: [open_window](#). This operator takes all necessary parameters to specify the dimensions, the mode, and the

relation to a potential parent window. As a result, a `WindowHandle` is returned with which you refer to the window when displaying into it or when visualization parameters are changed. Note that the dimension of the window is not limited by the size of the virtual display. Thus, you can work also on systems with multiple screens. The most important operators for controlling a window are `clear_window` to reset it to its background color, `set_part` to specify the display coordinate system and `close_window` when the window is no longer needed. The operator `set_window_param` allows to set different parameters of an open window.

19.1.2 Displaying

For each HALCON data type, specific operators for displaying are provided. The most convenient operator for iconic data (images, regions, and XLD) is `disp_obj`. This operator automatically handles gray or color images, regions, and XLDs. To control the way how data is presented in the window, operators with the prefix `set_` (or `dev_set_` for the visualization in HDevelop) are used. They allow to control the color, the draw mode, the line width, and many other parameters.

19.1.3 A First Example

An example for this basic concept is the following program, which shows how to visualize an image overlaid with a segmentation result. Here, the visualization operators provided in HDevelop are used.




Figure 19.1: Visualizing the segmented clips.


After reading the image from file the dark clips are selected with `bin_threshold`, which automatically selects the threshold value. After determining the connected components and selecting regions with

the appropriate size, the result is visualized. First, the image is displayed. After this, the parameters for regions are set to multi-colors (12) and margin mode. Finally, the regions are displayed with these settings.

```
read_image (Image, 'clip')
bin_threshold (Image, Dark)
connection (Dark, Single)
select_shape (Single, Selected, 'area', 'and', 5000, 10000)
dev_display (Image)
dev_set_colored (12)
dev_set_draw ('margin')
dev_display (Selected)
```

19.2 Extended Concept

In advanced applications it is required to gain complete control over the visualization process. This can include using graphics windows in programs developed with Microsoft Visual Basic or Microsoft C++, changing the behavior of the graphics windows, making use of buffered output, or even using external programs for the visualization. HALCON provides full control over all these topics to provide advanced flexibility, in addition to the ease of use.



19.2.1 Handling Graphics Windows

In this section we consider the concept of graphics windows in more detail.

- The graphics windows are designed such that each one stores all the corresponding parameters in a kind of graphics context. Whenever an operator like `set_draw` is called, this context is modified to be used for this window until the value is overwritten. All operators that modify the graphics context start with the prefix `set_`. The current value of the context can be requested using the corresponding operator with the prefix `get_`, e.g., `get_draw`.

- Besides display parameters, each graphics window has a coordinate system that can be defined with `set_part`. The upper left corner of an image is (0,0), the lower right corner is (height-1,width-1). The size of an image can be requested using `get_image_size`. Please note that unlike in HDevelop this coordinate system must be specified by calling `set_part`. Otherwise, the part of the image that is visualized is undefined.
- The operator `open_window` has a parameter called `Father`. By default, the value 0 is used, which means that the window has no parent and floats as an independent instance. You can construct hierarchies of windows by passing the handle of one window as parent to the other. Besides this, it is also possible to use the parent mechanism to embed the graphics windows into other forms (see [section 19.6](#) on page 299). To set different parameters of an open window, use `set_window_param`.

19.2.2 Displaying

After having opened a graphics window, the returned window handle is used to communicate with it. Typically, first the visualization parameters are specified before data is displayed. To control images only few operators are needed: `set_paint` (for profiles and 3D plots), `set_lut` (for look-up-tables), and `set_part_style` (for the zooming interpolation). To specify the output for regions, many parameters are available. The most important ones are `set_draw` (for filled or margin mode), `set_color` (for the pen color), `set_line_width` (for the pen width), and `set_colored` (for multi-color modes). To display XLD data, the same parameters (except `set_draw`) are used.

With the display mode '3d_plot' you can create an interactive display of a height field in 3D. The operator `update_window_pose` allows to manipulate the pose of such a 3D plot in an intuitive way and the operator `unproject_coordinates` calculates the image coordinates for a point in a 3D plot window. The example `hdevelop/Graphics/Parameters/set_paint_3d_plot.hdev` shows how to use these operators. Note that when working with HDevelop, you can also switch directly to the interactive display of the height field in the 3d_plot mode using the button that is described in the HDevelop User's Guide (see [section 5.7](#) on page 160).

The visualization itself is performed with operators like `disp_image`, `disp_color`, `disp_region`, or `disp_xld`. The most convenient way is to use `disp_obj`, which automatically uses the correct method.

For text output, you first specify the font with `set_font`. The desired position in the graphics window is determined with `set_tposition`. Writing text into the window is performed with `write_string`

For a flicker-free display see [section 19.6](#) on page 299.

19.2.3 Mouse Interaction

The most important thing to know is that HALCON does *not* use an event-driven approach for mouse handling. Each operator is designed such that the mouse interaction starts when the operator is called and finishes when the operator returns. If an event-driven mode is needed one has to use the standard mechanisms provided by the operating system.

Interacting with the mouse mainly involves two tasks:

- The first task is to request the position of the mouse. This can be achieved using `get_mposition` if a pixel-precise result is sufficient, or using `get_mposition_sub_pix` for a subpixel-precise output. These operators return immediately and have the position and the mouse button as result. An alternative are `get_mbutton` and `get_mbutton_sub_pix` for subpixel-precise positions. These operators only return when a mouse button has been clicked.
- The second important action is drawing shapes with the mouse. This is done with special operators whose names start with `draw_`. Operators for many different shapes (like circles or rectangles) are provided. Furthermore, different data types like regions or XLD contours can be used for the result.

19.2.4 Gnuplot

To visualize numeric values, HALCON provides an interface to the public domain software Gnuplot. This tool is especially useful to plot tuples of values like distributions and also provides extended versions of 3D plots of images.

19.3 Programming Examples

This section gives a brief introduction to using the visualization operators provided by HALCON.

19.3.1 Displaying HALCON data structures

Example: `solution_guide/basics/display_operators.hdev`

The example program is designed to show the major features of displaying images, regions, XLD, and text. It consists of a small main program that calls procedures handling the four different data types. The program is written for HDevelop and uses its specific display operators. This is done in a way that naturally ports (e.g., with the automatic export) to other language interfaces like C++, C#, or Visual Basic.

The main procedure contains five procedures: `open_graphics_window`, `display_image`, `display_regions`, `display_xld`, and `display_text`. To switch between the programs of the individual procedures you can use the combo box Procedures in the program window. Following, selected parts of each procedure are explained.

```
read_image (Image, 'fabrik')
open_graphics_window (Image, WindowHandle)
display_image (Image, WindowHandle)
regiongrowing (Image, Regions, 1, 1, 3, 100)
display_regions (Image, Regions, WindowHandle)
edges_sub_pix (Image, Edges, 'lanser2', 0.5, 10, 30)
display_xld (Image, Edges, WindowHandle)
display_text (Image, Regions, WindowHandle)
```

`open_graphics_window` is a support procedure to open a graphics window that has the same size as the image. This is done by calling `get_image_size` to access the image dimensions. Before opening the new window, the existing window is closed. To adapt the coordinate system accordingly, `dev_set_part` is called. This would be done automatically in HDevelop, but for the other programming environments this step is necessary. Finally, the default behavior of HDevelop of displaying each result automatically is switched off. This has the effect that only programmed output will be visible.

```
get_image_size (Image, Width, Height)
dev_close_window ()
dev_open_window (0, 0, Width, Height, 'white', WindowHandle)
dev_set_part (0, 0, Height-1, Width-1)
dev_update_window ('off')
```

Then, the display procedure for images is called: `display_image`. It has the `Image` and the `WindowHandle` as input parameters. First, the window is activated, which again is not needed for HDevelop, but is important for other programming environments. Now, the image is displayed in the graphics window.

```
dev_set_window (WindowHandle)
dev_display (Image)
```

To change the look-up-table (LUT), `dev_set_lut` is called and the effect will become visible after calling `dev_display` once again.

```
dev_set_lut ('temperature')
dev_display (Image)
```

Next, a part of the image is displayed using a so-called 3D plot (see [figure 19.2](#)). Here, the gray values are treated as height information. For this mode another LUT is used.

```
gen_rectangle1 (Rectangle, 358, 298, 387, 329)
dev_set_draw ('margin')
dev_set_color ('yellow')
dev_display (Rectangle)
dev_set_part (358, 298, 387, 329)
dev_set_lut ('twenty_four')
dev_set_paint (['3D-plot_hidden', 7, 1, 110, 160, 600, 0, 0])
dev_display (Image)
```

The procedure to display regions is called `display_regions`. It first displays the image as background and then sets the display parameters for the regions. `dev_set_draw` specifies that only the region border is visualized. `dev_set_colored` activates the multi-color mode, where each region is displayed with different colors (which change cyclically). As an alternative to simply showing the original shape of the regions, HALCON enables you to modify the shape using `dev_set_shape`. In the given example the equivalent ellipses are chosen. The result is depicted in [figure 19.3](#).




Figure 19.2: (a) Original image with ROI; (b) 3D plot of image within the ROI.

```
dev_display (Image)
dev_set_draw ('margin')
dev_set_colored (6)
dev_display (Regions)
dev_display (Image)
dev_set_shape ('ellipse')
dev_display (Regions)
```




Figure 19.3: (a) Regions and (b) their equivalent ellipses.

The procedure `display_xld` first shows all contours overlaid on the image using the multi-color mode. Then, a zoom is defined using `dev_set_part`. This zoom mode allows to inspect the sub-pixel accurate contours easily. To give additional information, contours with a given size are selected and for each of these the control points are extracted using `get_contour_xld`. The coordinates are here returned as tuples of real values. For each of these control points, a cross is generated using

`gen_cross_contour_xld`, which is then overlaid onto the contour.

```
dev_display (Contours)
gen_rectangle1 (Rectangle, 239, 197, 239+17, 197+17)
dev_set_part (239, 197, 239+17, 197+17)
select_shape_xld (Contours, SelectedXLD, 'area', 'and', 2000, 3000)
count_obj (SelectedXLD, Number)
for k := 1 to Number by 1
    select_obj (SelectedXLD, SingleContour, k)
    get_contour_xld (SingleContour, Row, Col)
    for i := 0 to |Row|-1 by 1
        gen_cross_contour_xld (Cross, Row[i], Col[i], 0.8, rad(45))
        dev_display (Cross)
    endfor
endfor
```




Figure 19.4: Subpixel accurate contour control points: (a) original image; (b) zoomed image part.

The last part of the program is a procedure called `display_text`. It shows how to handle the mouse and text output. The task is to click with the mouse into the graphics window to select a specific region and to calculate features, which are then displayed in the graphics window. First, the font is selected. This is the only part of the visualization where the parameters differ between Linux/UNIX and Windows because of the incompatible font name handling. However, as can be seen, the program can still be written so that it ports without problems.

```
OpSystem := environment('OS')
if (OpSystem='Windows_NT')
    set_font (WindowHandle, '-mono-16-*-*-*-*-*')
else
    set_font (WindowHandle, \
        '-adobe-courier-bold-r-normal--17-*-*-*-*-*')
endif
```

The rest of the program consists of a while-loop, which terminates as soon as the right mouse button is

pressed. The mouse operator `get_mbutton` waits until the user clicks with the mouse into the graphics window and then returns the coordinate and the button value. This coordinate is used to select the region that contains this point using `select_region_point`. For this region, the size and the center of gravity are calculated with `area_center`. First, the text cursor is positioned with `set_tposition` and the values are displayed using `write_string`. Here, it can be seen how conveniently strings can be composed using the "+" operator.

```
Button := 0
while (Button # 4)
  get_mbutton (WindowHandle, Row, Column, Button)
  select_region_point (Regions, DestRegions, Row, Column)
  area_center (DestRegions, Area, RowCenter, ColumnCenter)
  if (|Area| > 0)
    set_tposition (WindowHandle, Row, Column)
    dev_set_color ('yellow')
    write_string (WindowHandle, '('+RowCenter+', \
                  '+ColumnCenter+') = '+Area)
  endif
endwhile
```




Figure 19.5: Center of gravity and area of selected region.

19.4 Selecting Operators

19.4.1 Handling Graphics Windows

Standard:

`open_window`, `clear_window`, `close_window`, `set_window_param`

19.4.2 Displaying

Standard:

```
set_paint, set_lut, set_part_style, set_draw, set_line_width, set_color, set_colored,  
disp_obj, set_font, set_tposition, write_string, unproject_coordinates,  
update_window_pose
```

Advanced:

```
disp_image, disp_color, disp_region, disp_xld
```

19.4.3 Mouse Interaction

Standard:

```
get_mposition, get_mposition_sub_pix, get_mbutton, get_mbutton_sub_pix,  
draw_region, draw_circle
```

19.4.4 Gnuplot

Advanced:

```
gnuplot_open_file, gnuplot_open_pipe, gnuplot_plot_ctrl, gnuplot_plot_image
```

19.5 Tips & Tricks

19.5.1 Saving Window Content

HALCON provides an easy way to save the content of a graphics window to a file. This can, e.g., be useful for documentation purposes. The corresponding operator is called `dump_window`. It takes the window handle and the file name as input. The parameter `Device` allows to select amongst different file formats.

19.5.2 Execution Time

Generally, visualization takes time. To reduce time, it is recommended to visualize only when it is really needed for the specific task. When using `HDevEngine` (see the Programmer's Guide in [part VI](#) on page 171), by default the execution of all `dev_*` operators is suppressed. Thus, you can use as much display operators as you need while developing with `HDevelop` but save time when executing the final application from a programming language.

Further influences on the run time concern the graphics card and the bit depth. Choosing a bit depth of 16 instead of 32 in many cases speeds up the program execution significantly. So, if speed is important for you, we recommend to simply try different bit depths, perform typical display operators like `disp_image`, `disp_color`, `disp_region`, or `disp_xld`, and measure the execution time with `count_seconds`. Additionally, although it has generally only a small influence on the run time, cases exist where you can speed up the visualization also by choosing a lower screen resolution.

19.6 Advanced Topics

19.6.1 Programming Environments

The handling of graphics windows differs in the different programming environments. Especially HDevelop has a specific way of working with the visualization.

- Because HDevelop is an interactive environment, the handling of windows must be as easy as possible. In particular, it is important to display data without any need of programming. Therefore, the concept of window handles is used only if needed. Normally, the window where the output has to go to is not specified explicitly. Instead, HDevelop makes use of the activation status of the graphics windows. As a consequence, the display operators of HDevelop do not have a parameter for a window handle. Visualization operators in HDevelop look identical to their counterparts of HALCON except that their name starts with `dev_` and that the parameter for the window handle is suppressed. When exporting the code, this missing handle will automatically be inserted.

The second difference is that the windows in HDevelop have a history to automatically redisplay the data when the window has been resized. This is not the case with standard HALCON graphics windows.

The last difference is that HDevelop automatically sets the coordinate system according to the current image, whereas you must do this explicitly with programmed code when using HALCON. To make an export to, e.g., C++, C#, or Visual Basic transparent, we recommend to use `dev_set_window` when working with multiple graphics windows and to call `dev_set_part` to specify the coordinate system.

- When using HALCON windows in MFC, the usual way is to use the windows as a subwindow of a parent form. This can easily be achieved by using the window handle of the current form as the father. The handle must be converted to a long value that can then be passed to `open_window`. Note that `set_check` is needed to change to exception handling of HALCON in this context.

```
set_window_attr("border_width",0);
set_check("~father");
long lWWindowID = (long)m_hWnd;
open_window(0,0,640,480,lWWindowID,"visible","",&m_lWindowID);
set_check("father");
```

- Opening a HALCON window in Visual Basic is similar to the approach used for MFC. Here, you use the `memberhWnd` of the form or of another subwindow like a picture box. As an alternative, the `HWindowXCtrl` can be used.

```
Call sys.SetCheck("~father")
Call op.OpenWindow(8, 8, 320, 240, form.hWnd, "", "", WindowHandle)
Call sys.SetCheck("father")
```

19.6.2 Flicker-Free Visualization

For a flicker-free visualization, a sequential call of display operators is not suitable, because after each call the new data will immediately be flushed on the visible screen, which may cause flickering results.

Under Linux/UNIX or if you want to create a platform-independent application, this can be fixed by using a buffer window in addition to the visible window. A buffer window is opened with `open_window`, while setting `Mode` to 'buffer'. Then, all drawing operations are performed into the buffer window. Finally, the buffer window contents are copied into the visible window using `copy_rectangle`. An example using a buffer window is provided under `hdevelop/Transformations/2D-Transformations/projective_trans_pixel.hdev`.

When working exclusively with Windows, a faster approach is available. Here, you can call `set_system` with the parameter value 'flush_graphic' set to 'false', perform all display operators - except the last one - and then call `set_system` again with 'flush_graphic' set to 'true'. When finally applying the last display call of the display sequence, the whole data becomes visible in one step. An example using this approach is provided under `solution_guide/basics/median_interactive.hdev`.

19.6.3 Remote Visualization

In some applications, the computer used for processing differs from the computer used for visualization. Such applications can easily be created with HALCON using the socket communication. Operators like `send_image` or `receive_tuple` allow a transparent transfer of the relevant data to the control computer to apply visualization there.

19.6.4 Programmed Visualization

Sometimes it might be necessary not to apply the standard HALCON visualization operators, but to use a self-programmed version. This can be achieved by using the access functions provided for all data types. Examples for these are `get_image_pointer1`, `get_region_runs`, or `get_contour_xld`. Operators like these allow full access to all internal data types. Furthermore, they provide the data in various forms (e.g., runlength encoding, points, or contours) to make further processing easier. Based on this data, a self-programmed visualization can be developed easily.

As an alternative, with `set_window_type` the window type 'pixmap' can be chosen. In this case, all displayed data is painted into an internal buffer that can be accessed with `get_window_pointer3`. The returned pointers reference the three color channels of the buffer. This buffer can then easily be transferred (e.g., to another system) and/or transformed into the desired format. One example for a conversion is to call `gen_image3` to create a HALCON color image.

Index

- 1D measuring
 - relations to other methods, 60
 - basic concept, 50
 - examples, 53
 - extended concept, 51
 - operators, 58
 - overview, 49
- 3D matching
 - relations to other methods, 141
 - basic concept, 128
 - examples, 134
 - extended concept, 131
 - operators, 139
 - overview, 127
- 3D reconstruction, 7
- access external images, 17
- acquire image(s)
 - basic concept, 14
 - extended concept, 16
 - operators, 17
- acquire search data for 3D matching (surface-based)
 - basic concept, 130
 - operators, 141
- adjust bar code model
 - extended concept, 215
 - operators, 231
- align regions of interest or images for OCR
 - extended concept, 248
 - operators, 266
- alignment, 22
- apply texture filter
 - basic concept, 191
 - extended concept, 193
 - operators, 201
- bar code
 - relations to other methods, 232
 - basic concept, 211
 - examples, 221
 - extended concept, 212
 - operators, 230
 - overview, 211
- blob analysis
 - relations to other methods, 45
 - basic concept, 32
 - examples, 36
 - extended concept, 33
 - operators, 43
 - overview, 31
- calibrate multiple cameras
 - basic concept, 278
 - operators, 285
- check print quality of bar code
 - extended concept, 220
 - operators, 232
- check print quality of data code
 - extended concept, 238
 - operators, 242
- check variation model quality
 - extended concept, 147
 - operators, 153
- classification
 - relations to other methods, 169
 - basic concept, 156
 - examples, 162
 - extended concept, 160
 - operators, 168
 - overview, 155
- classify colors
 - extended concept, 177
 - operators, 187
- classify data
 - basic concept, 158
 - operators, 169

- clear training data of variation model
 - extended concept, 147
 - operators, 154
- close image acquisition device
 - basic concept, 14
 - operators, 17
- color inspection, 6
- color processing
 - basic concept, 174
 - examples, 177
 - extended concept, 175
 - operators, 186
 - overview, 173
- compare image with variation model
 - basic concept, 145
 - operators, 154
- completeness check, 6
- compose channels for color processing
 - extended concept, 177
 - operators, 187
- connected components, 46
- contour processing
 - relations to other methods, 95
 - basic concept, 84
 - examples, 89
 - extended concept, 86
 - operators, 93
 - overview, 83
- convert and access XLD contours
 - extended concept, 89
 - operators, 94
- create 3D matching model
 - basic concept, 129
 - operators, 140
- create bar code model
 - basic concept, 212
 - extended concept, 214
 - operators, 231
- create classifier
 - basic concept, 156
 - operators, 168
- create data code model
 - basic concept, 233
 - extended concept, 236
 - operators, 241
- create measure object for 1D measuring
 - basic concept, 50
 - extended concept, 52
 - operators, 59
- create region
 - basic concept, 19
 - extended concept, 22
 - operators, 27
- create region of interest
 - basic concept, 20
 - extended concept, 22
 - operators, 27
- create training samples from system fonts, 251
- create variation model
 - basic concept, 143
 - operators, 153
- create XLD contours
 - basic concept, 84
 - extended concept, 86
 - operators, 93
- data code
 - basic concept, 233
 - examples, 238
 - extended concept, 235
 - operators, 241
 - overview, 233
- decompose channels for color processing
 - basic concept, 174
 - operators, 186
- demosack Bayer pattern for color processing
 - extended concept, 175
 - operators, 186
- destroy 3D matching model
 - basic concept, 130
 - operators, 141
- destroy 3D object model
 - basic concept, 130
 - operators, 140
- destroy bar code model
 - basic concept, 212
 - operators, 231
- destroy classifier
 - basic concept, 158
 - operators, 169
- destroy data code model
 - basic concept, 234
 - operators, 243
- destroy matching model

- basic concept, 100
 - operators, 124
- destroy measure object
 - basic concept, 51
 - operators, 60
- destroy OCR classifier
 - basic concept, 247
 - operators, 268
- destroy variation model
 - basic concept, 145
 - operators, 154
- determine bar code parameters by training, 216
- determine contour attributes for edge extraction (subpixel-precise)
 - extended concept, 76
 - operators, 80
- determine training parameters for matching
 - extended concept, 104
 - operators, 122
- displaying
 - basic concept, 290
 - extended concept, 292
 - operators, 298
- draw region
 - extended concept, 21
 - operators, 27
- edge extraction (pixel-precise)
 - relations to other methods, 70
 - basic concept, 64
 - examples, 67
 - extended concept, 65
 - operators, 69
 - overview, 63
- edge extraction (subpixel-precise)
 - relations to other methods, 81
 - basic concept, 73
 - examples, 77
 - extended concept, 75
 - operators, 80
 - overview, 73
- effect of region of interest shape on speed up, 28
- effects of programming environment on visualization method, 299
- evaluate classifier
 - extended concept, 161
 - operators, 169
- extract color edges, 188
- extract color lines, 188
- extract edges (pixel-precise)
 - basic concept, 64
 - extended concept, 66
 - operators, 70
- extract edges or lines (subpixel-precise)
 - basic concept, 74
 - extended concept, 76
 - operators, 80
- extract features for blob analysis
 - basic concept, 32
 - extended concept, 35
 - operators, 45
- extract features of XLD contours
 - basic concept, 85
 - extended concept, 89
 - operators, 94
- extract segmentation parameters for blob analysis, 250
 - extended concept, 34
 - operators, 44
- Fast Fourier Transform (FFT), 202
- features for optical character recognition (OCR), 269
- features for texture analysis
 - basic concept, 191
 - extended concept, 194
 - operators, 201
- filter image for edge extraction (pixel-precise)
 - basic concept, 64
 - extended concept, 66
 - operators, 70
- find 3D matching model
 - basic concept, 130
 - operators, 141
- find model for matching
 - basic concept, 100
 - extended concept, 106
 - operators, 123
- flicker-free visualization method, 300
- fuzzy measuring, 61
- gnuplot
 - extended concept, 293
 - operators, 298

- gray-value profile, 61
- handling graphics windows
 - basic concept, 289
 - extended concept, 291
 - operators, 297
- handling gray values outside of the image domain, 29
- high-accuracy blob analysis, 47
- high-accuracy stereo vision, 287
- identification, 6
- image acquisition
 - basic concept, 13
 - examples, 16
 - extended concept, 15
 - operators, 17
 - overview, 13
- image filtering with region of interest, 28
- inspect 3D matching model
 - extended concept, 133
 - operators, 140
- inspect 3D object model
 - extended concept, 132
 - operators, 140
- inspect data code(s)
 - extended concept, 238
 - operators, 242
- matching
 - relations to other methods, 125
 - basic concept, 100
 - examples, 107
 - extended concept, 102
 - operators, 122
 - overview, 97
- matching with high accuracy, 126
- measure (1D measuring)
 - basic concept, 51
 - operators, 59
- measure with gray-value threshold (1D measuring), 61
- measuring and comparison 2D, 7
- measuring and comparison 3D, 7
- model creation (training)
 - basic concept, 100
 - extended concept, 104
 - operators, 123
- mouse interaction
 - extended concept, 292
 - operators, 298
- object recognition 2D, 8
- object recognition 3D, 8
- open image acquisition device
 - basic concept, 13
 - extended concept, 15
 - operators, 17
- optical character recognition (OCR)
 - relations to other methods, 268
 - basic concept, 245
 - examples, 254
 - extended concept, 248
 - operators, 266
 - overview, 245
- optimize model of data code
 - extended concept, 237
 - operators, 242
- perform fitting of XLD contours
 - basic concept, 84
 - extended concept, 88
 - operators, 94
- position recognition 2D, 8
- position recognition 3D, 9
- prepare variation model
 - basic concept, 145
 - operators, 153
- preprocess image(s) (filtering) for blob analysis
 - extended concept, 33
 - operators, 44
- preprocess image(s) (filtering) for OCR
 - extended concept, 250
 - operators, 266
- preprocess image(s) bar code
 - extended concept, 213
 - operators, 231
- print inspection, 9
- process edges (pixel-precise)
 - basic concept, 65
 - extended concept, 66
 - operators, 70
- process image (channels) for color processing
 - basic concept, 174
 - operators, 187
- process regions for blob analysis, 22

- extended concept, 35
 - operators, 44
- process XLD contours
 - basic concept, 84
 - extended concept, 76, 87
 - operators, 80, 93
- programmed visualization method, 300
- radiometrically calibrate image(s) for 1D measuring
 - extended concept, 51
 - operators, 59
- radiometrically calibrate image(s) for edge extraction (subpixel-precise)
 - extended concept, 75
 - operators, 80
- radiometrically calibrate image(s) for matching
 - extended concept, 102
 - operators, 122
- re-use 3D matching model
 - extended concept, 133
 - operators, 140
- re-use classifier
 - extended concept, 161
 - operators, 168
- re-use classifier training samples
 - extended concept, 161
 - operators, 168
- re-use measure object, 60
- re-use region of interest, 28
- read 3D object model
 - basic concept, 128
 - operators, 139
- read bar code
 - basic concept, 212
 - extended concept, 216
 - operators, 231
- read circular print, 269
- read composed symbols, 268
- read data code(s)
 - basic concept, 234
 - extended concept, 238
 - operators, 242
- read symbol
 - basic concept, 247
 - extended concept, 253
 - operators, 267
- reconstruct 3D information with stereo
 - basic concept, 279
 - operators, 286
- rectify image(s) for optical character recognition (OCR)
 - extended concept, 250
 - operators, 266
- rectify image(s) for stereo
 - basic concept, 279
 - operators, 286
- region of interest
 - relations to other methods, 28
 - basic concept, 19
 - examples, 23
 - extended concept, 21
 - operators, 26
 - overview, 19
- remote visualization method, 300
- robot vision, 10
- scale down image(s) for texture analysis
 - extended concept, 193
 - operators, 200
- security system, 10
- segment image(s) for blob analysis, 21
 - basic concept, 32
 - extended concept, 35
 - operators, 44
- segment image(s) for optical character recognition (OCR)
 - basic concept, 246
 - extended concept, 250
 - operators, 267
- select classifier training samples, 171
- set parameters for image acquisition
 - extended concept, 16
 - operators, 17
- set timeout for bar code reader, 232
- set timeout for data code reader, 243
- set timeout for matching, 126
- smart text finder, 250
- speed up bar code reader, 215
- speed up blob analysis, 46
- speed up color processing, 188
- speed up edge extraction (pixel-precise), 71
- speed up matching, 125
- speed up stereo vision, 287

- speed up visualization method, 298
- stereo vision
 - relations to other methods, 286
 - basic concept, 278
 - examples, 281
 - extended concept, 280
 - operators, 285
 - overview, 277
- suppress clutter or noise for 1D measuring, 60
- surface inspection, 11
- texture analysis
 - relations to other methods, 201
 - basic concept, 190
 - examples, 194
 - extended concept, 192
 - operators, 200
 - overview, 189
- texture analysis in color images, 203
- texture inspection, 11
- train classifier
 - basic concept, 157
 - extended concept, 160
 - operators, 168
- train colors
 - extended concept, 176
 - operators, 187
- train model of 2D data code
 - extended concept, 237
 - operators, 242
- train optical character recognition (OCR)
 - basic concept, 247
 - extended concept, 251
 - operators, 267
- train variation model
 - basic concept, 144
 - operators, 153
- transform color space
 - extended concept, 175
 - operators, 186
- transform results of 1D measuring into 3D (world) coordinates
 - extended concept, 53
 - operators, 59
- transform results of blob analysis into 3D (world) coordinates
 - extended concept, 35
 - operators, 45
- transform results of contour processing into 3D (world) coordinates
 - extended concept, 88
 - operators, 94
- transform results of edge extraction (subpixel-precise) into 3D (world) coordinates
 - extended concept, 77
 - operators, 81
- unsupported image acquisition device, 18
- use bar code autodiscrimination, 216
- use binary images as region of interest, 28
- use line scan camera for blob analysis, 47
- use line scan camera for contour processing, 95
- use line scan camera for optical character recognition (OCR), 269
- use optische Schriftzeichenerkennung (OCR) for classification, 171
- use region of interest for stereo
 - extended concept, 281
 - operators, 286
- use results of texture analysis
 - extended concept, 194
 - operators, 201
- variation model (image comparison)
 - basic concept, 143
 - examples, 148
 - extended concept, 147
 - operators, 153
 - overview, 143
- visualization
 - basic concept, 289
 - examples, 293
 - extended concept, 291
 - operators, 297
 - overview, 289
- visualize results of 1D measuring
 - extended concept, 53
 - operators, 60
- visualize results of 3D matching
 - extended concept, 133
 - operators, 141
- visualize results of color processing
 - extended concept, 177
 - operators, 187

write window content, [298](#)