




the Power of Machine Vision

Solution Guide III-C

3D Vision



Machine vision in 3D world coordinates, 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	December 2003	(HALCON 7.0)
Edition 1a	July 2004	(HALCON 7.0.1)
Edition 1b	April 2005	(HALCON 7.0.2)
Edition 2	July 2005	(HALCON 7.1)
Edition 2a	December 2006	(HALCON 7.1.2)
Edition 3	June 2007	(HALCON 8.0)
Edition 3a	October 2007	(HALCON 8.0.1)
Edition 4	December 2008	(HALCON 9.0)
Edition 4a	June 2009	(HALCON 9.0.1)
Edition 4b	March 2010	(HALCON 9.0.2)
Edition 5	October 2010	(HALCON 10.0)
Edition 5a	April 2011	(HALCON 10.0.1)

Copyright © 2003-2013 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.

More information about HALCON can be found at: <http://www.halcon.com/>

About This Manual

Measurements in 3D become more and more important. HALCON provides many methods to perform 3D measurements. This Solution Guide gives you an overview over these methods, and it assists you with the selection and the correct application of the appropriate method.

A short characterization of the various methods is given in [chapter 1](#) on page [9](#). Principles of 3D transformations and poses as well as the description of the camera model can be found in [chapter 2](#) on page [15](#). Afterwards, the methods to perform 3D measurements are described in detail.

The HDevelop example programs that are presented in this Solution Guide can be found in the specified subdirectories of the directory `%HALCONROOT%`.

Contents

1	Introduction	9
2	Basics	15
2.1	3D Transformations and Poses	15
2.1.1	3D Coordinates	15
2.1.2	Translation	16
2.1.3	Rotation	18
2.1.4	Rigid Transformations and Homogeneous Transformation Matrices	22
2.1.5	3D Poses	25
2.2	Camera Model and Parameters	26
2.2.1	Area scan cameras	27
2.2.2	Line scan cameras	33
3	3D Vision in a Specified Plane With a Single Camera	37
3.1	First Example	38
3.2	3D Camera Calibration	40
3.2.1	Creating the Calibration Data Model	41
3.2.2	Specifying Initial Values for the Internal Camera Parameters	41
3.2.3	Describing the Calibration Object	45
3.2.4	Observing the Calibration Object in Multiple Poses (Images)	49
3.2.5	Restricting the Calibration to Specific Parameters	52
3.2.6	Performing the Calibration	52
3.2.7	Accessing the Results of the Calibration	53
3.2.8	Saving the Results and Destroying the Calibration Data Model	56
3.2.9	Troubleshooting	57
3.3	Transforming Image into World Coordinates and Vice Versa	57
3.3.1	The Main Principle	58
3.3.2	World Coordinates for Points	58
3.3.3	World Coordinates for Contours	60
3.3.4	World Coordinates for Regions	60
3.3.5	Transforming World Coordinates into Image Coordinates	60
3.3.6	Compensate for Lens Distortions Only	61
3.4	Rectifying Images	62
3.4.1	Transforming Images into the WCS	62
3.4.2	Compensate for Lens Distortions Only	70

3.5	Inspection of Non-Planar Objects	70
4	3D Position Recognition of Known Objects	75
4.1	Pose Estimation from Points	76
4.2	Pose Estimation Using Shape-Based 3D Matching	81
4.2.1	General Proceeding for Shape-Based 3D Matching	81
4.2.2	Enhance the Shape-Based 3D Matching	86
4.2.3	Tips and Tricks for Problem Handling	89
4.3	Pose Estimation Using Surface-Based 3D Matching	93
4.3.1	General Proceeding for Surface-Based 3D Matching	93
4.4	Pose Estimation Using 3D Primitives Fitting	97
4.5	Pose Estimation Using Calibrated Perspective Deformable Matching	101
4.6	Pose Estimation Using Calibrated Descriptor-Based Matching	101
4.7	Pose Estimation for Circles	102
4.8	Pose Estimation for Rectangles	102
5	3D Vision With a Stereo System	105
5.1	The Principle of Stereo Vision	107
5.1.1	The Setup of a Stereo Camera System	109
5.1.2	Resolution of a Stereo Camera System	110
5.2	Calibrating the Stereo Camera System	110
5.2.1	Creating and Configuring the Calibration Data Model	111
5.2.2	Acquiring Calibration Images	112
5.2.3	Observing the Calibration Object	113
5.2.4	Calibrating the Cameras	113
5.3	Binocular Stereo Vision	113
5.3.1	Correlation-Based Versus Multigrid Stereo	114
5.3.2	Accessing the Calibration Results	115
5.3.3	Acquiring Stereo Images	115
5.3.4	Rectifying the Stereo Images	116
5.3.5	Reconstructing 3D Information	120
5.3.6	Uncalibrated Stereo Vision	130
5.4	Multi-View Stereo Vision	130
5.4.1	Initializing the Stereo Model	131
5.4.2	Reconstructing 3D Information	133
6	Laser Triangulation with Sheet of Light	139
6.1	The Principle of Sheet of Light	139
6.2	The Measurement Setup	141
6.3	Calibrating the Sheet-of-Light Setup	143
6.4	Performing the Measurement	148
6.4.1	Calibrated Sheet-of-Light Measurement	149
6.4.2	Uncalibrated Sheet-of-Light Measurement	151
6.5	Using the Score Image	153
6.6	3D Cameras for Sheet of Light	155
7	Depth from Focus	157
7.1	The Principle of Depth from Focus	158

7.1.1	Speed vs. Accuracy	159
7.2	Setup	160
7.2.1	Camera	160
7.2.2	Illumination	163
7.2.3	Object	164
7.3	Working with Depth from Focus	165
7.3.1	Rules for Taking Images	166
7.3.2	Practical Use of Depth from Focus	167
7.3.3	Volume Measurement with Depth from Focus	168
7.4	Solutions for Typical Problems With DFF	168
7.4.1	Calibrating Aberration	169
7.5	Special Cases	170
7.6	Performing Depth from Focus with a Standard Lens	171
8	Robot Vision	173
8.1	The Principle of Hand-Eye Calibration	174
8.2	Determining Suitable Input Parameters	176
8.2.1	Corresponding World and Image Points	176
8.2.2	Robot Poses	177
8.3	Performing the Calibration	178
8.4	Using the Calibration Data	180
8.4.1	Using the Hand-Eye Calibration for Grasping (3D Alignment)	180
8.4.2	How to Get the 3D Pose of the Object	181
8.4.3	Example Application with a Stationary Camera: Grasping a Nut	182
9	Calibrated Mosaicking	187
9.1	Setup	188
9.2	Calibration	188
9.3	Merging the Individual Images into One Larger Image	190
9.3.1	Definition of the Rectification of the First Image	191
9.3.2	Definition of the Rectification of the Second Image	193
9.3.3	Rectification of the Images	196
10	Uncalibrated Mosaicking	199
10.1	Rules for Taking Images for a Mosaic Image	201
10.2	Definition of Overlapping Image Pairs	202
10.3	Detection of Characteristic Points	207
10.4	Matching of Characteristic Points	208
10.5	Generation of the Mosaic Image	210
10.6	Bundle Adjusted Mosaicking	211
10.7	Spherical Mosaicking	213
11	Rectification of Arbitrary Distortions	215
11.1	Basic Principle	217
11.2	Rules for Taking Images of the Rectification Grid	218
11.3	Machine Vision on Ruled Surfaces	220
11.4	Using Self-Defined Rectification Grids	223

A	The HALCON Calibration Plate	229
B	HDevelop Procedures Used in this Solution Guide	231
B.1	gen_hom_mat3d_from_three_points	231
B.2	parameters_image_to_world_plane_centered	232
B.3	parameters_image_to_world_plane_entire	233
B.4	tilt_correction	234
B.5	visualize_results_of_find_marks_and_pose	234
B.6	display_calplate_coordinate_system	235
B.7	display_3d_coordinate_system	235
B.8	calc_calplate_pose_movingcam	236
B.9	calc_calplate_pose_stationarycam	236
B.10	define_reference_coord_system	237
	Index	239

Chapter 1

Introduction

With HALCON you can perform 3D vision in various ways. The main applications comprise the 3D position recognition and the 3D inspection, which both consist of several different approaches with different characteristics, so that for a wide range of 3D vision tasks a proper solution can be provided. This Solution Guide provides you with detailed information on the available approaches, including also some auxiliary methods that are needed only in specific cases.

What Basic Knowledge Do You Need for 3D Vision?

Typically, you have to calibrate your camera(s) before applying a 3D vision task. Especially, if you want to achieve accurate results, the **camera calibration** is essential, because it is of no use to extract edges with an accuracy of 1/40 pixel if the lens distortion of the uncalibrated camera accounts for a couple of pixels. This also applies if you use cameras with telecentric lenses. But don't be afraid of the calibration process: In HALCON, this can be done with just a few lines of code. To prepare you for the camera calibration, [chapter 2](#) on page 15 introduces you to the details on the camera model and parameters. The actual camera calibration is then described in [chapter 3](#) on page 37.

Using a camera calibration, you can transform image processing results into arbitrary 3D coordinate systems and thus derive metrical information from images, regardless of the position and orientation of the camera with respect to the object. In other words, you can perform inspection tasks in 3D coordinates in specified object planes, which can be oriented arbitrarily with respect to the camera. This is, e.g., useful if the camera cannot be mounted such that it looks perpendicular to the object surface. Thus, besides the pure camera calibration, [chapter 3](#) shows how to apply a general **3D vision task with a single camera in a specified plane**. Additionally, it shows how to **rectify the images** such that they appear as if they were acquired from a camera that has no lens distortions and that looks exactly perpendicular onto the object surface. This is useful for tasks like OCR or the recognition and localization of objects, which rely on images that are not distorted too much with respect to the training images.

Before you develop your application, we recommend to read [chapter 2](#) and [chapter 3](#) and then, depending on the task at hand, to step into the section that describes the 3D vision approach you selected for your specific application.

How Can You Obtain an Object's 3D Position and Orientation?

The position and orientation of 3D objects with respect to a given 3D coordinate system, which is needed, e.g., for pick-and-place applications (3D alignment), can be determined by one of the methods described in [chapter 4](#) on page 75:

- The pose estimation of a known 3D object from **corresponding points** ([section 4.1](#) on page 76) is a rather general approach that includes a camera calibration and the extraction of at least three significant points for which the 3D object coordinates are known. The approach is also known as “mono 3D”.
- HALCON's **3D matching** locates known 3D objects based on a 3D model of the object. In particular, it automatically searches objects that correspond to a 3D model in the search data and determines their 3D poses. The model must be provided, e.g., as a Computer Aided Design (CAD) model. Available approaches are the *shape-based 3D matching* ([section 4.2](#) on page 81) that searches the model in 2D images and the *surface-based 3D matching* ([section 4.3](#) on page 93) that searches the model in a 3D scene, i.e., in a set of 3D points that is available as 3D object model, which can be obtained by a 3D reconstruction approach like stereo or sheet of light. Note that the surface-based matching is also known as “volume matching”, although it only relies on points on the object's surface.
- HALCON's **3D primitives fitting** ([section 4.4](#) on page 97) fits a primitive 3D shape like a cylinder, sphere, or plane into a 3D scene, i.e., into a set of 3D points that is available as a 3D object model, which can be obtained by a 3D reconstruction approach like stereo or sheet of light followed by a 3D segmentation.
- The calibrated **perspective matching** locates perspective distorted planar objects in images based on a 2D model. In particular, it automatically searches objects that correspond to a 2D model in the search images and determines their 3D poses. The model typically is obtained from a representative model image. Available approaches are the *calibrated perspective deformable matching* ([section 4.5](#) on page 101) that describes the model by its contours and the *calibrated descriptor-based matching* ([section 4.6](#) on page 101) that describes the model by a set of distinctive points that are called “interest points”.
- The **circle pose estimation** ([section 4.7](#) on page 102) and **rectangle pose estimation** ([section 4.8](#) on page 102) use the perspective distortions of circles and rectangles to determine the pose of planar objects that contain circles and/or rectangles in a rather convenient way.

How Can You Inspect a 3D Object?

The inspection of 3D objects can be applied by different means. If the inspection in a specified plane is sufficient, you can **use a camera calibration together with a 2D inspection** as is described in [chapter 3](#) on page 37.

If the surface of the 3D object is needed and/or the inspection can not be reduced to a single specified plane, you can **use a 3D reconstruction together with a 3D inspection**. That is, you use the point, surface, or height information returned for a 3D object by a 3D reconstruction and inspect the object, e.g., by comparing it to a reference point, surface, or height.

[Figure 1.1](#) provides you with an overview on the methods that are available for 3D position recognition and 3D inspection.

vary the focus position additional hardware like a translation stage or linear piezo stage is required. Note that depending on the direction in which the focus position is modified, the result corresponds either to a height image or to a distance image. A height image contains the distances between a specific object or measure plane and the object points, whereas the distance image typically contains the distances between the camera and the object points.

- With **photometric stereo** (Reference Manual, chapter “[3D Reconstruction](#) ▸ [Photometric Stereo](#)”) a height image can be obtained using images that are acquired by a single telecentric camera but with at least three different telecentric illumination sources for which the spatial relations to the camera must be known.
- Besides the 3D reconstruction approaches provided by HALCON, you can obtain 3D information also by **specific 3D sensors** like time of flight (TOF) cameras or specific setups that use structured light. These cameras typically are calibrated and return X, Y, and Z images.

Figure 1.2 allows to compare some important features of the different 3D reconstruction approaches like the approach-specific result types.

3D Reconstruction Approach	Hardware Requirements	Object Size	Possible Results
Multi-View Stereo	multiple cameras, calibration object	approx. > 10 cm	3D object model or X, Y, Z coordinates
Binocular Stereo	two cameras, calibration object	approx. > 10 cm	X, Y, Z coordinates, approach-specific disparity image, or Z image
Sheet of Light	camera, laser line projector, unit to move the object, and calibration object	object must fit onto the moving unit	3D object model, X, Y, Z images, or approach-specific disparity image
Depth from Focus	telecentric camera, hardware to variate the focus position	approx. < 2cm	Z image
Photometric Stereo	telecentric camera, at least three telecentric illumination sources	restricted by field of view of telecentric lens	Z image
3D Sensors	special camera like calibrated TOF	approx. 30cm-5m	X, Y, Z images

Figure 1.2: 3D reconstruction: a coarse comparison.

About the 3D Object Model

Note that all (calibrated) 3D reconstruction approaches are suitable to explicitly or implicitly derive a **3D object model**. For example, with multi-view stereo you can explicitly obtain a 3D object model, whereas

with a common 3D sensor X, Y, and Z images and with depth from focus depth images are obtained. X, Y, and Z images implicitly contain the information needed for a 3D object model. Thus, you can derive a 3D object model from X, Y, and Z images using the operator `xyz_to_object_model_3d`. If only a (calibrated) depth image, i.e., a “Z image” is available, you can obtain corresponding X, Y, and Z images by adding artificial X and Y images to the Z image. The artificial images have the same size as the Z image. The X image is created by assigning the column numbers of the Z image to each row of the X image and the Y image is created by assigning the row numbers of the Z image to each column of the Y image. The thus created X, Y, and Z images can then be transformed again into a 3D object model.

Note that depending on the selected reconstruction approach different information can be contained in the 3D object model. Typically, a 3D object model that is obtained from X, Y, and Z images contains the coordinates of the 3D points and the corresponding 2D mapping, whereas a 3D object model that is obtained by multi-view stereo can contain a lot of further information. Generally, you can use the operator `prepare_object_model_3d` to prepare a 3D object model for one of the 3D position recognition approaches that use the 3D object models as input, in particular the surface-based matching and the 3D primitives fitting, as these need specific information from the 3D object model. [Figure 1.3](#) guides you through the different ways how to derive a 3D object model.

How Can You Extend 3D Vision to Robot Vision?

A typical application area for 3D vision is **robot vision**, i.e., using the results of machine vision to command a robot. In such applications you must perform an additional calibration: the so-called hand-eye calibration, which determines the relation between camera and robot coordinates ([chapter 8](#) on page [173](#)). Again, this calibration must be performed only once (offline). Its results allow you to quickly transform machine vision results from camera into robot coordinates.

What Tasks May be Needed Additionally?

If the object that you want to inspect is too large to be covered by one image with the desired resolution, multiple images, each covering only a part of the object, can be combined into one larger **mosaic image**. This can be done either based on a calibrated camera setup with very high precision ([chapter 9](#) on page [187](#)) or highly automated for arbitrary and even varying image configurations ([chapter 10](#) on page [199](#)).

If an image shows distortions that are different to the common perspective distortions or lens distortions, caused, e.g., by a non-flat object surface, the so-called **grid rectification** can be applied to rectify the image ([chapter 11](#) on page [215](#)).

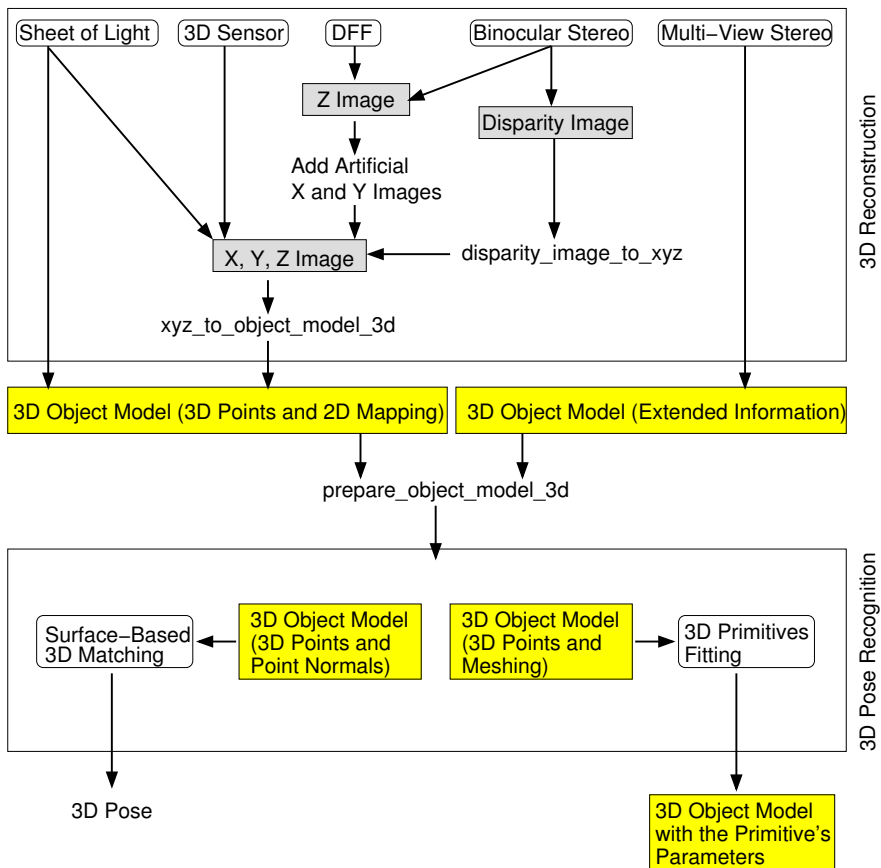


Figure 1.3: Overview on the 3D object model.

Chapter 2

Basics

2.1 3D Transformations and Poses

Before we start explaining how to perform 3D vision with HALCON, we take a closer look at some basic questions regarding the use of 3D coordinates:

- How to describe the transformation (translation and rotation) of points and coordinate systems,
- how to describe the position and orientation of one coordinate system relative to another, and
- how to determine the coordinates of a point in different coordinate systems, i.e., how to transform coordinates between coordinate systems.

In fact, all these tasks can be solved using one and the same means: homogeneous transformation matrices and their more compact equivalent, 3D poses.

2.1.1 3D Coordinates

The position of a 3D point P is described by its three coordinates (x_p, y_p, z_p) . The coordinates can also be interpreted as a 3D vector (indicated by a bold-face lower-case letter). The coordinate system in which the point coordinates are given is indicated to the upper right of a vector or coordinate. For example, the coordinates of the point P in the camera coordinate system (denoted by the letter c) and in the world coordinate system (denoted by the letter w) would be written as:

$$\mathbf{p}^c = \begin{pmatrix} x_p^c \\ y_p^c \\ z_p^c \end{pmatrix} \quad \mathbf{p}^w = \begin{pmatrix} x_p^w \\ y_p^w \\ z_p^w \end{pmatrix}$$

Figure 2.1 depicts an example point lying in a plane where measurements are to be performed and its coordinates in the camera and world coordinate system, respectively.

2.1.2 Translation

2.1.2.1 Translation of Points

In [figure 2.2](#), our example point has been translated along the x -axis of the camera coordinate system.

The coordinates of the resulting point P_2 can be calculated by adding two vectors, the coordinate vector \mathbf{p}_1 of the point and the translation vector \mathbf{t} :

$$\mathbf{p}_2 = \mathbf{p}_1 + \mathbf{t} = \begin{pmatrix} x_{p_1} + x_t \\ y_{p_1} + y_t \\ z_{p_1} + z_t \end{pmatrix} \quad (2.1)$$

Multiple translations are described by adding the translation vectors. This operation is *commutative*, i.e., the sequence of the translations has no influence on the result.

2.1.2.2 Translation of Coordinate Systems

Coordinate systems can be translated just like points. In the example in [figure 2.3](#), the coordinate system c_1 is translated to form a second coordinate system, c_2 . Then, the position of c_2 in c_1 , i.e., the coordinate vector of its origin relative to c_1 ($\mathbf{o}_{c_2}^{c_1}$), is identical to the translation vector:

$$\mathbf{t}^{c_1} = \mathbf{o}_{c_2}^{c_1} \quad (2.2)$$

2.1.2.3 Coordinate Transformations

Let's turn to the question how to transform point coordinates between (translated) coordinate systems. In fact, the translation of a point can also be thought of as translating it together with its local coordinate system. This is depicted in [figure 2.3](#): The coordinate system c_1 , together with the point Q_1 , is translated




Figure 2.1: Coordinates of a point in two different coordinate systems.




Figure 2.2: Translating a point.




Figure 2.3: Translating a coordinate system (and point).

by the vector \mathbf{t} , resulting in the coordinate system c_2 and the point Q_2 . The points Q_1 and Q_2 then have the same coordinates relative to their local coordinate system, i.e., $\mathbf{q}_1^{c1} = \mathbf{q}_2^{c2}$.

If coordinate systems are only translated relative to each other, coordinates can be transformed very easily between them by adding the translation vector:

$$\mathbf{q}_2^{c1} = \mathbf{q}_2^{c2} + \mathbf{t}^{c1} = \mathbf{q}_2^{c2} + \mathbf{o}_{c_2}^{c1} \quad (2.3)$$

In fact, [figure 2.3](#) visualizes this equation: \mathbf{q}_2^{c1} , i.e., the coordinate vector of Q_2 in the coordinate system c_1 , is composed by adding the translation vector \mathbf{t} and the coordinate vector of Q_2 in the coordinate system c_2 (\mathbf{q}_2^{c2}).

The downside of this graphical notation is that, at first glance, the direction of the translation vector appears to be contrary to the direction of the coordinate transformation: The vector points from the coordinate system c_1 to c_2 , but transforms coordinates from the coordinate system c_2 to c_1 . According to this, the coordinates of Q_1 in the coordinate system c_2 , i.e., the inverse transformation, can be obtained by subtracting the translation vector from the coordinates of Q_1 in the coordinate system c_1 :

$$\mathbf{q}_1^{c_2} = \mathbf{q}_1^{c_1} - \mathbf{t}^{c_1} = \mathbf{q}_1^{c_1} - \mathbf{o}_{c_2}^{c_1} \tag{2.4}$$

2.1.2.4 Summary

- Points are translated by adding the translation vector to their coordinate vector. Analogously, coordinate systems are translated by adding the translation vector to the position (coordinate vector) of their origin.
- To transform point coordinates from a translated coordinate system c_2 into the original coordinate system c_1 , you apply the same transformation to the points that was applied to the coordinate system, i.e., you add the translation vector used to translate the coordinate system c_1 into c_2 .
- Multiple translations are described by adding all translation vectors; the sequence of the translations does not affect the result.

2.1.3 Rotation

2.1.3.1 Rotation of Points

In [figure 2.4a](#), the point \mathbf{p}_1 is rotated by -90° around the z^c -axis of the camera coordinate system.




Figure 2.4: Rotate a point: (a) first around the z^c -axis; (b) then around the y^c -axis.

Rotating a point is expressed by multiplying its coordinate vector with a 3×3 rotation matrix \mathbf{R} . A rotation around the z-axis looks as follows:

$$\mathbf{p}_3 = \mathbf{R}_z(\gamma) \cdot \mathbf{p}_1 = \begin{bmatrix} \cos \gamma & -\sin \gamma & 0 \\ \sin \gamma & \cos \gamma & 0 \\ 0 & 0 & 1 \end{bmatrix} \cdot \begin{pmatrix} x_{p_1} \\ y_{p_1} \\ z_{p_1} \end{pmatrix} = \begin{pmatrix} \cos \gamma \cdot x_{p_1} - \sin \gamma \cdot y_{p_1} \\ \sin \gamma \cdot x_{p_1} + \cos \gamma \cdot y_{p_1} \\ z_{p_1} \end{pmatrix} \quad (2.5)$$

Rotations around the x- and y-axis correspond to the following rotation matrices:

$$\mathbf{R}_y(\beta) = \begin{bmatrix} \cos \beta & 0 & \sin \beta \\ 0 & 1 & 0 \\ -\sin \beta & 0 & \cos \beta \end{bmatrix} \quad \mathbf{R}_x(\alpha) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \alpha & -\sin \alpha \\ 0 & \sin \alpha & \cos \alpha \end{bmatrix} \quad (2.6)$$

2.1.3.2 Chain of Rotations

In [figure 2.4b](#), the rotated point is further rotated around the y-axis. Such a chain of rotations can be expressed very elegantly by a chain of rotation matrices:

$$\mathbf{p}_4 = \mathbf{R}_y(\beta) \cdot \mathbf{p}_3 = \mathbf{R}_y(\beta) \cdot \mathbf{R}_z(\gamma) \cdot \mathbf{p}_1 \quad (2.7)$$

Note that in contrast to a multiplication of scalars, the multiplication of matrices is not commutative, i.e., if you change the sequence of the rotation matrices, you get a different result.

2.1.3.3 Rotation of Coordinate Systems

In contrast to points, coordinate systems have an orientation relative to other coordinate systems. This orientation changes when the coordinate system is rotated. For example, in [figure 2.5a](#) the coordinate system c_3 has been rotated around the y-axis of the coordinate system c_1 , resulting in a different orientation of the camera. Note that in order to rotate a coordinate system in your mind's eye, it may help to image the points of the axis vectors being rotated.

Just like the position of a coordinate system can be expressed directly by the translation vector (see [equation 2.2](#) on page 16), the orientation is contained in the rotation matrix: The columns of the rotation matrix correspond to the axis vectors of the rotated coordinate system in coordinates of the original one:

$$\mathbf{R} = \begin{bmatrix} \mathbf{x}_{c_3}^{c_1} & \mathbf{y}_{c_3}^{c_1} & \mathbf{z}_{c_3}^{c_1} \end{bmatrix} \quad (2.8)$$

For example, the axis vectors of the coordinate system c_3 in [figure 2.5a](#) can be determined from the corresponding rotation matrix $\mathbf{R}_y(90^\circ)$ as shown in the following equation; you can easily check the result in the figure.

$$\begin{aligned} \mathbf{R}_y(90^\circ) &= \begin{bmatrix} \cos(90^\circ) & 0 & \sin(90^\circ) \\ 0 & 1 & 0 \\ -\sin(90^\circ) & 0 & \cos(90^\circ) \end{bmatrix} = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ -1 & 0 & 0 \end{bmatrix} \\ \Rightarrow \mathbf{x}_{c_3}^{c_1} &= \begin{pmatrix} 0 \\ 0 \\ -1 \end{pmatrix} \quad \mathbf{y}_{c_3}^{c_1} = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} \quad \mathbf{z}_{c_3}^{c_1} = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \end{aligned}$$




Figure 2.5: Rotate coordinate system: (a) first around the y^{c1} -axis; (b) then around the z^{c3} -axis.

2.1.3.4 Coordinate Transformations

Like in the case of translation, to transform point coordinates from a rotated coordinate system c_3 into the original coordinate system c_1 , you apply the same transformation to the points that was applied to the coordinate system c_3 , i.e., you multiply the point coordinates with the rotation matrix used to rotate the coordinate system c_1 into c_3 :

$$\mathbf{q}_3^{c1} = {}^{c1}\mathbf{R}_{c3} \cdot \mathbf{q}_3^{c3} \tag{2.9}$$

This is depicted in [figure 2.5](#) also for a chain of rotations, which corresponds to the following equation:

$$\mathbf{q}_4^{c1} = {}^{c1}\mathbf{R}_{c3} \cdot {}^{c3}\mathbf{R}_{c4} \cdot \mathbf{q}_4^{c4} = \mathbf{R}_y(\beta) \cdot \mathbf{R}_z(\gamma) \cdot \mathbf{q}_4^{c4} = {}^{c1}\mathbf{R}_{c4} \cdot \mathbf{q}_4^{c4} \tag{2.10}$$

2.1.3.5 In Which Sequence and Around Which Axes are Rotations Performed?

If you compare the chains of rotations in [figure 2.4](#) and [figure 2.5](#) and the corresponding [equations 2.7](#) and [2.10](#), you will note that two different sequences of rotations are described by the same chain of rotation matrices: In [figure 2.4](#), the point was rotated *first* around the z -axis and *then* around the y -axis, whereas in [figure 2.5](#) the coordinate system is rotated *first* around the y -axis and *then* around the z -axis. Yet, both are described by the chain $\mathbf{R}_y(\beta) \cdot \mathbf{R}_z(\gamma)$!

The solution to this seemingly paradox situation is that in the two examples the chain of rotation matrices can be “read” in different directions: In [figure 2.4](#) it is read from the right to left, and in [figure 2.5](#) from left to the right.

However, there still must be a difference between the two sequences because, as we already mentioned, the multiplication of rotation matrices is not commutative. This difference lies in the second question in the title, i.e., around which axes the rotations are performed.

Let's start with the second rotation of the coordinate system in [figure 2.5b](#). Here, there are two possible sets of axes to rotate around: those of the “old” coordinate system c_1 and those of the already rotated, “new” coordinate system c_3 . In the example, the second rotation is performed around the “new” z-axis.

In contrast, when rotating points as in [figure 2.4](#), there is only one set of axes around which to rotate: those of the “old” coordinate system.

From this, we derive the following rules:

- When reading a chain from the **left to right**, rotations are performed around the “**new**” axes.
- When reading a chain from the **right to left**, rotations are performed around the “**old**” axes.

As already remarked, point rotation chains are always read from right to left. In the case of coordinate systems, you have the choice how to read a rotation chain. In most cases, however, it is more intuitive to read them from left to right.

[Figure 2.6](#) shows that the two reading directions really yield the same result.

2.1.3.6 Summary

- Points are rotated by multiplying their coordinate vector with a rotation matrix.
- If you rotate a coordinate system, the rotation matrix describes its resulting orientation: The column vectors of the matrix correspond to the axis vectors of the rotated coordinate system in coordinates of the original one.
- To transform point coordinates from a rotated coordinate system c_3 into the original coordinate system c_1 , you apply the same transformation to the points that was applied to the coordinate system, i.e., you multiply them with the rotation matrix that was used to rotate the coordinate system c_1 into c_3 .

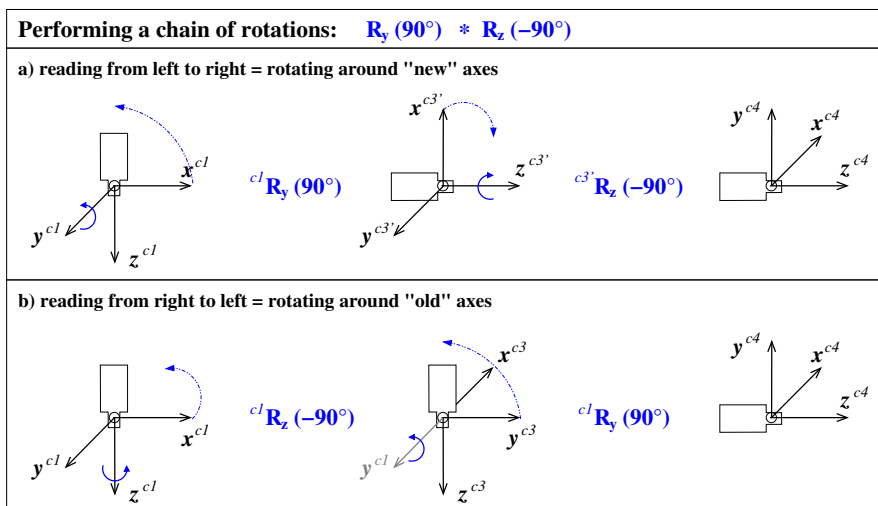


Figure 2.6: Performing a chain of rotations (a) from left to the right, or (b) from right to left.




Figure 2.7: Combining the translation from figure 2.2 on page 17 and the rotation of figure 2.4 on page 18 to form a rigid transformation.

- Multiple rotations are described by a chain of rotation matrices, which can be read in two directions. When read from left to right, rotations are performed around the “new” axes; when read from right to left, the rotations are performed around the “old” axes.

2.1.4 Rigid Transformations and Homogeneous Transformation Matrices

2.1.4.1 Rigid Transformation of Points

If you combine translation and rotation, you get a so-called rigid transformation. For example, in figure 2.7, the translation and rotation of the point from figures 2.2 and 2.4 are combined. Such a transformation is described as follows:

$$p_5 = R \cdot p_1 + t \tag{2.11}$$

For multiple transformations, such equations quickly become confusing, as the following example with two transformations shows:

$$p_6 = R_a \cdot (R_b \cdot p_1 + t_b) + t_a = R_a \cdot R_b \cdot p_1 + R_a \cdot t_b + t_a \tag{2.12}$$

An elegant alternative is to use so-called *homogeneous transformation matrices* and the corresponding homogeneous vectors. A homogeneous transformation matrix H contains both the rotation matrix and the translation vector. For example, the rigid transformation from equation 2.11 can be rewritten as follows:

$$\begin{pmatrix} p_5 \\ 1 \end{pmatrix} = \begin{bmatrix} R & t \\ 0 & 1 \end{bmatrix} \cdot \begin{pmatrix} p_1 \\ 1 \end{pmatrix} = \begin{pmatrix} R \cdot p_1 + t \\ 1 \end{pmatrix} = H \cdot p_1 \tag{2.13}$$

The usefulness of this notation becomes apparent when dealing with sequences of rigid transformations, which can be expressed as chains of homogeneous transformation matrices, similarly to the rotation chains:

$$\mathbf{H}_1 \cdot \mathbf{H}_2 = \begin{bmatrix} \mathbf{R}_a & \mathbf{t}_a \\ 000 & 1 \end{bmatrix} \cdot \begin{bmatrix} \mathbf{R}_b & \mathbf{t}_b \\ 000 & 1 \end{bmatrix} = \begin{bmatrix} \mathbf{R}_a \cdot \mathbf{R}_b & \mathbf{R}_a \cdot \mathbf{t}_b + \mathbf{t}_a \\ 000 & 1 \end{bmatrix} \quad (2.14)$$

As explained for chains of rotations, chains of rigid transformation can be read in two directions. When reading from left to right, the transformations are performed around the “new” axes, when read from right to left around the “old” axes.

In fact, a rigid transformation is already a chain, since it consists of a translation and a rotation:

$$\mathbf{H} = \begin{bmatrix} \mathbf{R} & \mathbf{t} \\ 000 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & \mathbf{t} \\ 0 & 1 & 0 & \\ 0 & 0 & 1 & \\ 0 & 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} \mathbf{R} & 0 \\ 000 & 1 \end{bmatrix} = \mathbf{H}(\mathbf{t}) \cdot \mathbf{H}(\mathbf{R}) \quad (2.15)$$

If the rotation is composed of multiple rotations around axes as in [figure 2.7](#), the individual rotations can also be written as homogeneous transformation matrices:

$$\mathbf{H} = \begin{bmatrix} \mathbf{R}_y(\beta) \cdot \mathbf{R}_z(\gamma) & \mathbf{t} \\ 000 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & \mathbf{t} \\ 0 & 1 & 0 & \\ 0 & 0 & 1 & \\ 0 & 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} \mathbf{R}_y(\beta) & 0 \\ 000 & 1 \end{bmatrix} \cdot \begin{bmatrix} \mathbf{R}_z(\gamma) & 0 \\ 000 & 1 \end{bmatrix}$$

Reading this chain from right to left, you can follow the transformation of the point in [figure 2.7](#): First, it is rotated around the z-axis, then around the (“old”) y-axis, and finally it is translated.

2.1.4.2 Rigid Transformation of Coordinate Systems

Rigid transformations of coordinate systems work along the same lines as described for a separate translation and rotation. This means that the homogeneous transformation matrix ${}^{c_1}\mathbf{H}_{c_5}$ describes the transformation of the coordinate system c_1 into the coordinate system c_5 . At the same time, it describes the position and orientation of coordinate system c_5 relative to coordinate system c_1 : Its column vectors contain the coordinates of the axis vectors and the origin.

$${}^{c_1}\mathbf{H}_{c_5} = \begin{bmatrix} \mathbf{x}_{c_5}^{c_1} & \mathbf{y}_{c_5}^{c_1} & \mathbf{z}_{c_5}^{c_1} & \mathbf{o}_{c_5}^{c_1} \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (2.16)$$

As already noted for rotations, chains of rigid transformations of coordinate systems are typically read from left to right. Thus, the chain above can be read as first translating the coordinate system, then rotating it around its “new” y-axis, and finally rotating it around its “newest” z-axis.

2.1.4.3 Coordinate Transformations

As described for the separate translation and the rotation, to transform point coordinates from a rigidly transformed coordinate system c_5 into the original coordinate system c_1 , you apply the same transformation to the points that was applied to the coordinate system c_5 , i.e., you multiply the point coordinates with the homogeneous transformation matrix:

$$\begin{pmatrix} \mathbf{p}_5^{c_1} \\ 1 \end{pmatrix} = {}^{c_1}\mathbf{H}_{c_5} \cdot \begin{pmatrix} \mathbf{p}_5^{c_5} \\ 1 \end{pmatrix} \quad (2.17)$$

Typically, you leave out the homogeneous vectors if there is no danger of confusion and simply write:

$$\mathbf{p}_5^{c_1} = {}^{c_1}\mathbf{H}_{c_5} \cdot \mathbf{p}_5^{c_5} \quad (2.18)$$

2.1.4.4 Summary

- Rigid transformations consist of a rotation and a translation. They are described very elegantly by homogeneous transformation matrices, which contain both the rotation matrix and the translation vector.
- Points are transformed by multiplying their coordinate vector with the homogeneous transformation matrix.
- If you transform a coordinate system, the homogeneous transformation matrix describes the coordinate system's resulting position and orientation: The column vectors of the matrix correspond to the axis vectors and the origin of the coordinate system in coordinates of the original one. Thus, you could say that a homogeneous transformation matrix “is” the position and orientation of a coordinate system.
- To transform point coordinates from a rigidly transformed coordinate system c_5 into the original coordinate system c_1 , you apply the same transformation to the points that was applied to the coordinate system, i.e., you multiply them with the homogeneous transformation matrix that was used to transform the coordinate system c_1 into c_5 .
- Multiple rigid transformations are described by a chain of transformation matrices, which can be read in two directions. When read from left to the right, rotations are performed around the “new” axes; when read from the right to left, the transformations are performed around the “old” axes.

2.1.4.5 HALCON Operators

As we already anticipated at the beginning of [section 2.1](#) on page 15, homogeneous transformation matrices are the answer to all our questions regarding the use of 3D coordinates. Because of this, they form the basis for HALCON's operators for 3D transformations. Below, you find a brief overview of the relevant operators. For more details follow the links into the Reference Manual.

- [hom_mat3d_identity](#) creates the identity transformation
- [hom_mat3d_translate](#) translates along the “old” axes: $\mathbf{H}_2 = \mathbf{H}(t) \cdot \mathbf{H}_1$

- `hom_mat3d_translate_local` translates along the “new” axes: $\mathbf{H}_2 = \mathbf{H}_1 \cdot \mathbf{H}(\mathbf{t})$
- `hom_mat3d_rotate` rotates around the “old” axes: $\mathbf{H}_2 = \mathbf{H}(\mathbf{R}) \cdot \mathbf{H}_1$
- `hom_mat3d_rotate_local` rotates around the “new” axes: $\mathbf{H}_2 = \mathbf{H}_1 \cdot \mathbf{H}(\mathbf{R})$
- `hom_mat3d_compose` multiplies two transformation matrices: $\mathbf{H}_3 = \mathbf{H}_1 \cdot \mathbf{H}_2$
- `hom_mat3d_invert` inverts a transformation matrix: $\mathbf{H}_2 = \mathbf{H}_1^{-1}$
- `affine_trans_point_3d` transforms a point using a transformation matrix: $\mathbf{p}_2 = \mathbf{H}_0 \cdot \mathbf{p}_1$

2.1.5 3D Poses

Homogeneous transformation matrices are a very elegant means of describing transformations, but their content, i.e., the elements of the matrix, are often difficult to read, especially the rotation part. This problem is alleviated by using so-called *3D poses*.

A 3D pose is nothing more than an easier-to-understand representation of a rigid transformation: Instead of the 12 elements of the homogeneous transformation matrix, a pose describes the rigid transformation with 6 parameters, 3 for the rotation and 3 for the translation: (`TransX`, `TransY`, `TransZ`, `RotX`, `RotY`, `RotZ`). The main principle behind poses is that even a rotation around an arbitrary axis can always be represented by a sequence of three rotations around the axes of a coordinate system.

In HALCON, you create 3D poses with `create_pose`; to transform between poses and homogeneous matrices you can use `hom_mat3d_to_pose` and `pose_to_hom_mat3d`.

2.1.5.1 Sequence of Rotations

However, there is more than one way to represent an arbitrary rotation by three parameters. This is reflected by the HALCON operator `create_pose`, which lets you choose between different pose types with the parameter `OrderOfRotation`. If you pass the value `'gba'`, the rotation is described by the following chain of rotations:

$$\mathbf{R}_{gba} = \mathbf{R}_x(\text{RotX}) \cdot \mathbf{R}_y(\text{RotY}) \cdot \mathbf{R}_z(\text{RotZ}) \quad (2.19)$$

You may also choose the inverse order by passing the value `'abg'`:

$$\mathbf{R}_{abg} = \mathbf{R}_z(\text{RotZ}) \cdot \mathbf{R}_y(\text{RotY}) \cdot \mathbf{R}_x(\text{RotX}) \quad (2.20)$$

For example, the transformation discussed in the previous sections can be represented by the homogeneous transformation matrix

$$\mathbf{H} = \begin{bmatrix} \mathbf{R}_y(\beta) \cdot \mathbf{R}_z(\gamma) & \mathbf{t} \\ 0 & 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} \cos \beta \cdot \cos \gamma & -\cos \beta \cdot \sin \gamma & \sin \beta & x_t \\ \sin \gamma & \cos \gamma & 0 & y_t \\ -\sin \beta \cdot \cos \gamma & \sin \beta \cdot \sin \gamma & \cos \beta & z_t \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

The corresponding pose with the rotation order 'gba' is much easier to read:

$$(\text{TransX} = x_t, \text{TransY} = y_t, \text{TransZ} = z_t, \text{RotX} = 0, \text{RotY} = 90^\circ, \text{RotZ} = -90^\circ)$$

If you look closely at [figure 2.5](#) on page 20, you can see that the rotation can also be described by the sequence $\mathbf{R}_z(-90^\circ) \cdot \mathbf{R}_x(-90^\circ)$. Thus, the transformation can also be described by the following pose with the rotation order 'abg':

$$(\text{TransX} = x_t, \text{TransY} = y_t, \text{TransZ} = z_t, \text{RotX} = -90^\circ, \text{RotY} = 0, \text{RotZ} = -90^\circ)$$

2.1.5.2 HALCON Operators

Below, the relevant HALCON operators for dealing with 3D poses are briefly described. For more details follow the links into the Reference Manual.

- [create_pose](#) creates a pose
- [hom_mat3d_to_pose](#) converts a homogeneous transformation matrix into a pose
- [pose_to_hom_mat3d](#) converts a pose into a homogeneous transformation matrix
- [convert_pose_type](#) changes the pose type
- [write_pose](#) writes a pose into a file
- [read_pose](#) reads a pose from a file
- [set_origin_pose](#) translates a pose along its “new” axes

2.1.5.3 How to Determine the Pose of a Coordinate System

The previous sections showed how to describe known transformations using translation vectors, rotation matrices, homogeneous transformation matrices, or poses. Sometimes, however, there is another task: How to describe the position and orientation of a coordinate system with a pose.

[Figure 2.8](#) shows how to proceed for a rather simple example. The task is to determine the pose of the world coordinate system from [figure 2.1](#) on page 16 relative to the camera coordinate system.

In such a case, we recommend to build up the rigid transformation from individual translations and rotations from left to right. Thus, in [figure 2.8](#) the camera coordinate system is first translated such that its origin coincides with that of the world coordinate system. Now, the y-axes of the two coordinate systems coincide; after rotating the (translated) camera coordinate system around its (new) y-axis by 180° , it has the correct orientation.

2.2 Camera Model and Parameters

If you want to derive *accurate* world coordinates from your imagery, you first have to calibrate your camera. To calibrate a camera, a model for the mapping of the 3D points of the world to the 2D image generated by the camera, lens, and frame grabber is necessary.




Figure 2.8: Determining the pose of the world coordinate system in camera coordinates.

HALCON supports the calibration of two different kinds of cameras: area scan cameras and line scan cameras. While area scan cameras acquire the image in one step, line scan cameras generate the image line by line (see Solution Guide II-A, [section 6.6](#) on page 46). Therefore, the line scan camera must move relative to the object during the acquisition process.

Two different types of lenses are relevant for machine vision tasks. The first type of lens effects a perspective projection of the world coordinates into the image, just like the human eye does. With this type of lens, objects become smaller in the image the farther they are away from the camera. This combination of camera and lens is called a *pinhole camera model* because the perspective projection can also be achieved if a small hole is drilled in a thin planar object and this plane is held parallel in front of another plane (the image plane).

The second type of lens that is relevant for machine vision is called a telecentric lens. Its major difference is that it effects a parallel projection of the world coordinates onto the image plane (for a certain range of distances of the object from the camera). This means that objects have the same size in the image independent of their distance to the camera. This combination of camera and lens is called a *telecentric camera model*.

In the following, first the camera model for area scan cameras is described in detail, then, the camera model for line scan cameras is explained.

2.2.1 Area scan cameras

[Figure 2.9](#) displays the perspective projection effected by a pinhole camera graphically. The world point P is projected through the optical center of the lens to the point P' in the image plane, which is located at a distance of f (the focal length) behind the optical center. Actually, the term “focal length” is not quite correct and would be appropriate only for an infinite object distance. To simplify matters, in the following always the term “focal length” is used even if the “image distance” is meant. Note that the focal length and thus the focus must not be changed after applying the camera calibration.

Although the image plane in reality lies behind the optical center of the lens, it is easier to pretend that it lies at a distance of f in front of the optical center, as shown in [figure 2.10](#). This causes the image




Figure 2.9: Perspective projection by a pinhole camera.

coordinate system to be aligned with the pixel coordinate system (row coordinates increase downward

and column coordinates to the right) and simplifies most calculations.

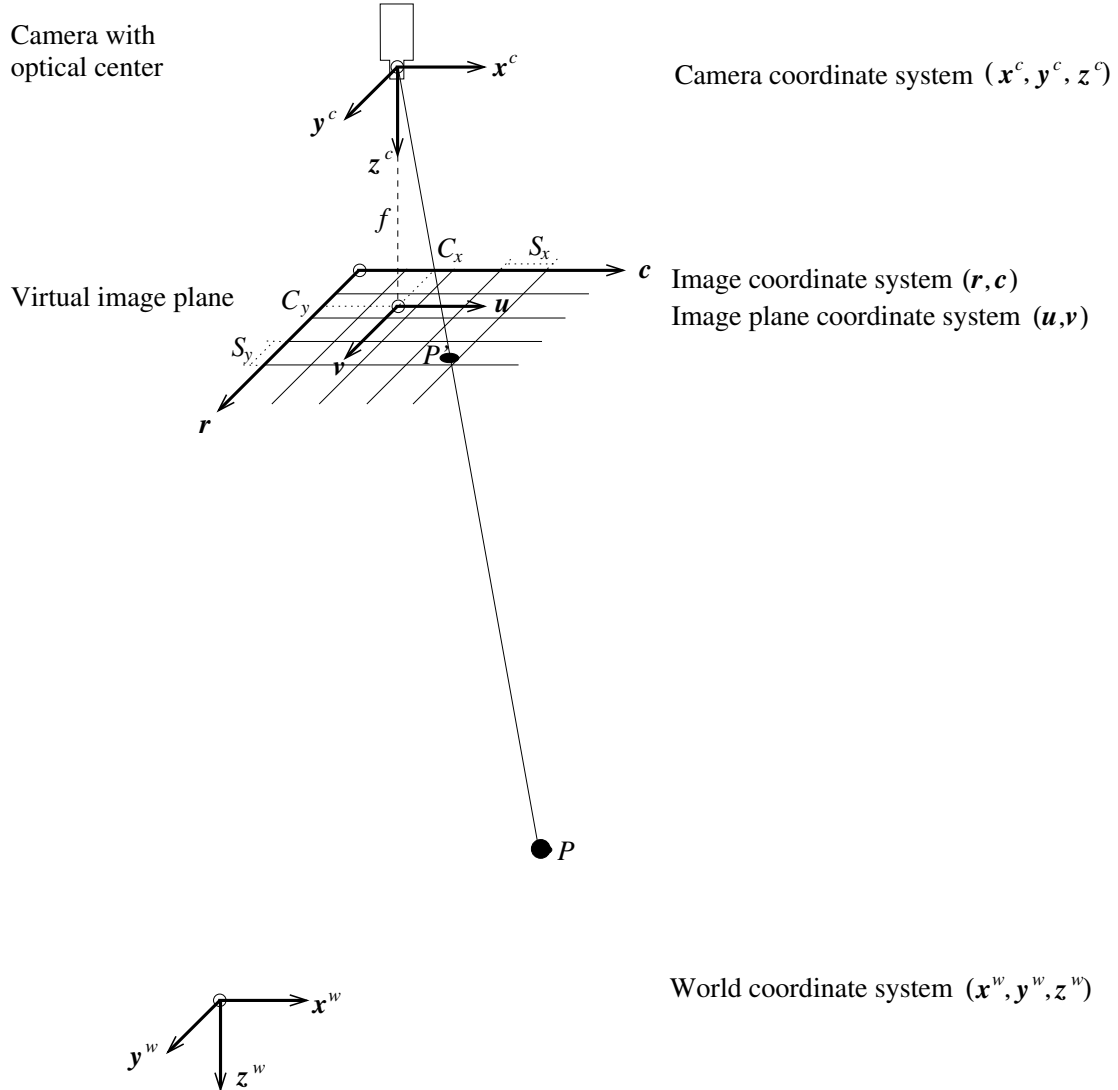


Figure 2.10: Image plane and virtual image plane.

2.2.1.1 Transformation into Camera Coordinates (External Camera Parameters)

With this, we are now ready to describe the projection of objects in 3D world coordinates to the 2D image plane and the corresponding camera parameters. First, we should note that the points P are given in a world coordinate system (WCS). To make the projection into the image plane possible, they need to be transformed into the camera coordinate system (CCS). The CCS is defined so that its x and y axes

are parallel to the column and row axes of the image, respectively, and the z axis is perpendicular to the image plane.

The transformation from the WCS to the CCS is a rigid transformation, which can be expressed by a pose or, equivalently, by the homogeneous transformation matrix ${}^c\mathbf{H}_w$. Therefore, the camera coordinates $\mathbf{p}^c = (\mathbf{x}^c, \mathbf{y}^c, \mathbf{z}^c)^T$ of point P can be calculated from its world coordinates $\mathbf{p}^w = (\mathbf{x}^w, \mathbf{y}^w, \mathbf{z}^w)^T$ simply by

$$\mathbf{p}^c = {}^c\mathbf{H}_w \cdot \mathbf{p}^w \quad (2.21)$$

The six parameters of this transformation (the three translations t_x , t_y , and t_z and the three rotations α , β , and γ) are called the *external camera parameters* because they determine the position of the camera with respect to the world. In HALCON, they are stored as a pose, i.e. together with a code that describes the order of translation and rotations.

2.2.1.2 Projection

The next step is the projection of the 3D point given in the CCS into the image plane coordinate system (IPCS). For the pinhole camera model, the projection is a perspective projection, which is given by

$$\begin{pmatrix} u \\ v \end{pmatrix} = \frac{f}{z^c} \begin{pmatrix} \mathbf{x}^c \\ \mathbf{y}^c \end{pmatrix} \quad (2.22)$$

For the telecentric camera model, the projection is a parallel projection, which is given by

$$\begin{pmatrix} u \\ v \end{pmatrix} = \begin{pmatrix} \mathbf{x}^c \\ \mathbf{y}^c \end{pmatrix} \quad (2.23)$$

As can be seen, there is no focal length f for telecentric cameras. Furthermore, the distance z of the object to the camera has no influence on the image coordinates.

2.2.1.3 Lens Distortion

After the projection into the image plane, the lens distortion causes the coordinates $(u, v)^T$ to be modified. If no lens distortions were present, the projected point P' would lie on a straight line from P through the optical center, indicated by the dotted line in [figure 2.11](#). Lens distortions cause the point P' to lie at a different position.

The lens distortion is a transformation that can be modeled in the image plane alone, i.e., 3D information is unnecessary. In HALCON, the distortions can be modeled either by the division model or by the polynomial model.

The division model uses one parameter (κ) to model the radial distortions. The following equations transform the distorted image plane coordinates into undistorted image plane coordinates if the division model is used:

$$u = \frac{\tilde{u}}{1 + \kappa(\tilde{u}^2 + \tilde{v}^2)} \quad \text{and} \quad v = \frac{\tilde{v}}{1 + \kappa(\tilde{u}^2 + \tilde{v}^2)} \quad (2.24)$$




Figure 2.11: Schematic illustration of the effect of the lens distortion.

These equations can be inverted analytically, which leads to the following equations that transform undistorted coordinates into distorted coordinates if the division model is used:

$$\tilde{u} = \frac{2u}{1 + \sqrt{1 - 4\kappa(u^2 + v^2)}} \quad \text{and} \quad \tilde{v} = \frac{2v}{1 + \sqrt{1 - 4\kappa(u^2 + v^2)}} \quad (2.25)$$

The parameter κ models the magnitude of the radial distortions. If κ is negative, the distortion is barrel-shaped, while for positive κ it is pincushion-shaped (see [figure 2.12](#)).

The polynomial model uses three parameters (K_1, K_2, K_3) to model the radial distortions, and two parameters (P_1, P_2) to model the decentering distortions. The following equations transform the distorted image plane coordinates into undistorted image plane coordinates if the polynomial model is used:

$$u = \tilde{u} + \tilde{u}(K_1 r^2 + K_2 r^4 + K_3 r^6) + 2P_1 \tilde{u} \tilde{v} + P_2 (r^2 + 2\tilde{u}^2) \quad (2.26)$$

$$v = \tilde{v} + \tilde{v}(K_1 r^2 + K_2 r^4 + K_3 r^6) + P_1 (r^2 + 2\tilde{v}^2) + 2P_2 \tilde{u} \tilde{v} \quad (2.27)$$

with $r = \sqrt{\tilde{u}^2 + \tilde{v}^2}$. These equations cannot be inverted analytically. Therefore, distorted image plane coordinates must be calculated from undistorted image plane coordinates numerically.

The parameters $K_1, K_2,$ and K_3 model the magnitude of the radial distortions and the parameters P_1 and P_2 model the magnitude of the decentering distortions. Some examples for the kind of distortions that can be modeled with the polynomial model are shown in [figure 2.13](#).




Figure 2.12: Effect of radial distortions modeled with the division model with $\kappa > 0$ (left), $\kappa = 0$ (middle), and $\kappa < 0$ (right).

2.2.1.4 Transformation into Pixel Coordinates

Finally, the point $(\tilde{u}, \tilde{v})^T$ is transformed from the image plane coordinate system into the image coordinate system (the pixel coordinate system):

$$\begin{pmatrix} r \\ c \end{pmatrix} = \begin{pmatrix} \frac{\tilde{v}}{S_y} + C_y \\ \frac{\tilde{u}}{S_x} + C_x \end{pmatrix} \quad (2.28)$$

Here, S_x and S_y are scaling factors. For pinhole cameras, they represent the horizontal and vertical distance of the sensors elements on the CCD chip of the camera. For cameras with telecentric lenses, they represent the size of a pixel in world coordinates (not taking into account the lens distortions). The point $(C_x, C_y)^T$ is the principal point of the image. For pinhole cameras, this is the perpendicular projection of the optical center onto the image plane, i.e., the point in the image from which a ray through the optical center is perpendicular to the image plane. It also defines the center of the radial distortions. For telecentric cameras, no optical center exists. Therefore, the principal point is solely defined by the radial distortions.

The six parameters $(f, \kappa, S_x, S_y, C_x, C_y)$ or the ten parameters $(f, K_1, K_2, K_3, P_1, P_2, S_x, S_y, C_x, C_y)$ of the pinhole camera and the five parameters $(\kappa, S_x, S_y, C_x, C_y)$ or the nine parameters $(K_1, K_2, K_3, P_1, P_2, S_x, S_y, C_x, C_y)$ of the telecentric camera are called the *internal camera parameters* because they determine the projection from 3D to 2D performed by the camera.

Note that, in HALCON, the internal camera parameters always contain the focal length and, in addition to the above described parameters, the width (*NumColumns*) and the height (*NumRows*) of the image. The differentiation among the two camera models (pinhole and telecentric) is done based on the value of the focal length. If it has a positive value, a pinhole camera with the given focal length is assumed. If the focal length is set to zero, the telecentric camera model is used. The differentiation among the two models for the lens distortion (division model and polynomial model) is done based on the number of values of the internal camera parameters. If eight values are passed, the division model is used. If twelve values are passed, the polynomial model is used.

With this, we can see that camera calibration is the process of determining the internal camera parameters $(f, \kappa, S_x, S_y, C_x, C_y)$ and the external camera parameters $(t_x, t_y, t_z, \alpha, \beta, \gamma)$.

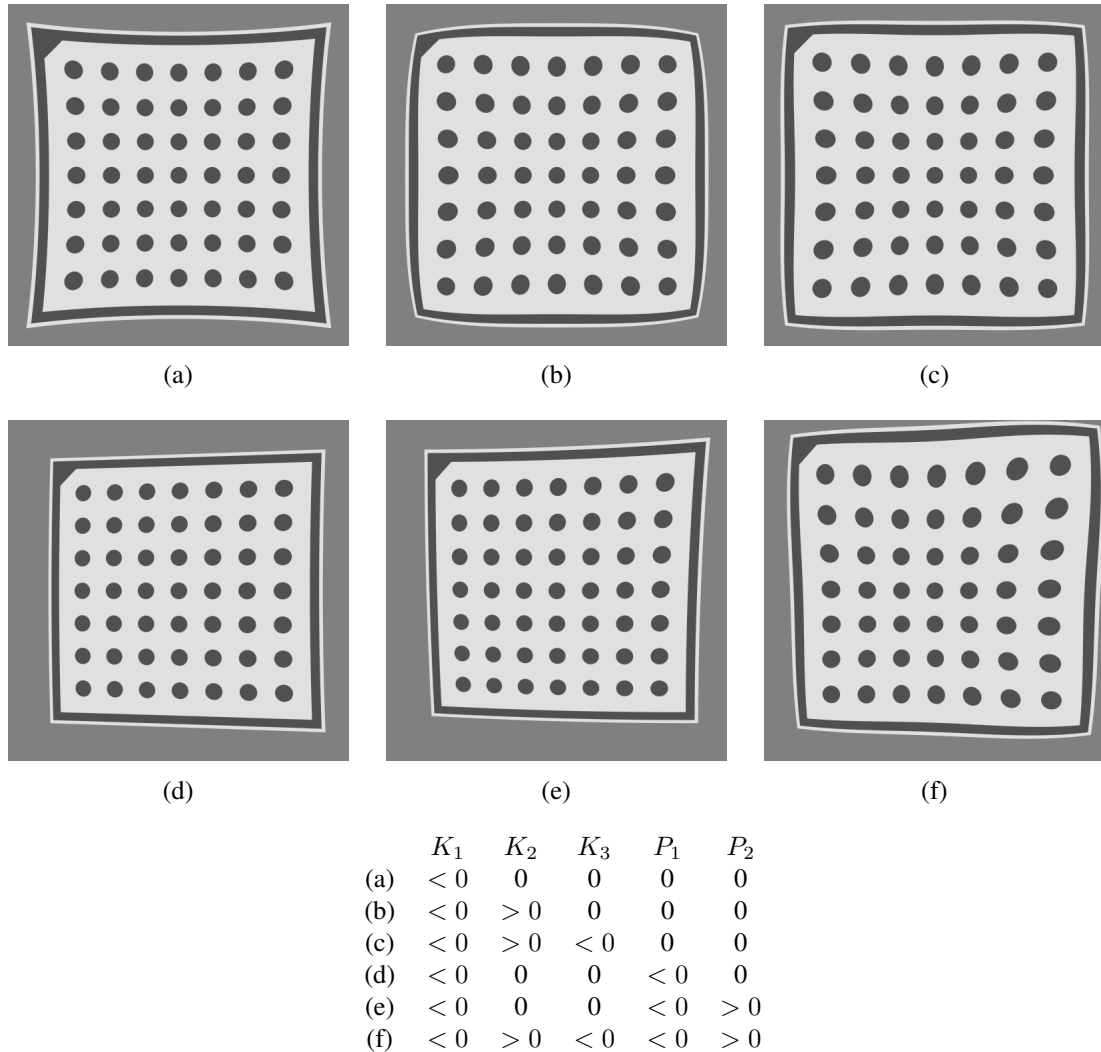


Figure 2.13: Effect of distortions modeled with the polynomial model with different values for the parameters K_1 , K_2 , K_3 , P_1 , and P_2 .

2.2.2 Line scan cameras

A line scan camera has only a one-dimensional line of sensor elements, i.e., to acquire an image, the camera must move relative to the object (see [figure 2.14](#) on page 34). This means that the camera moves over a fixed object, the object travels in front of a fixed camera, or camera and object are both moving.

The relative motion between the camera and the object is modeled in HALCON as part of the internal camera parameters. In HALCON, the following assumptions for this motion are made:

1. the camera moves — relative to the object — with constant velocity along a straight line




Figure 2.14: Principle of line scan image acquisition.

2. the orientation of the camera is constant with respect to the object
3. the motion is equal for all images

The motion is described by the motion vector $V = (V_x, V_y, V_z)^T$, which must be given in [meters/scanline] in the camera coordinate system. The motion vector describes the motion of the camera, i.e., it assumes a fixed object. In fact, this is equivalent to the assumption of a fixed camera with the object traveling along $-V$.

The camera coordinate system of line scan cameras is defined as follows (see [figure 2.15](#)): The origin of the coordinate system is the center of projection. The z-axis is identical to the optical axis and it is directed so that the visible points have positive z coordinates. The y-axis is perpendicular to the sensor line and to the z-axis. It is directed so that the motion vector has a positive y-component, i.e., if a fixed object is assumed, the y-axis points in the direction in which the camera is moving. The x-axis is perpendicular to the y- and z-axis, so that the x-, y-, and z-axis form a right-handed coordinate system.

Similarly to area scan cameras, the projection of a point given in world coordinates into the image is modeled in two steps: First, the point is transformed into the camera coordinate system. Then, it is projected into the image.

As the camera moves over the object during the image acquisition, also the camera coordinate system moves relative to the object, i.e., each image line has been imaged from a different position. This means that there would be an individual pose for each image line. To make things easier, in HALCON all transformations from world coordinates into camera coordinates and vice versa are based on the pose of the first image line only. The motion V is taken into account during the projection of the point p^c into the image.

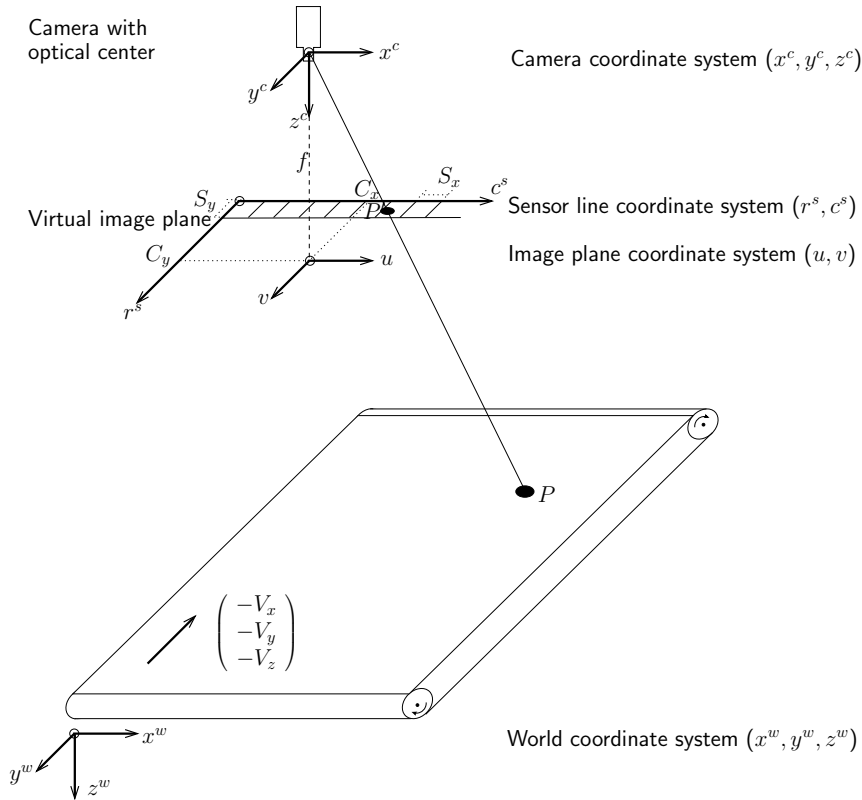


Figure 2.15: Coordinate systems in regard to a line scan camera.

The transformation from the WCS to the CCS of the first image line is a rigid transformation, which can be expressed by a pose or, equivalently, by the homogeneous transformation matrix ${}^c\mathbf{H}_w$. Therefore, the camera coordinates $\mathbf{p}^c = (x^c, y^c, z^c)^T$ of point P can be calculated from its world coordinates $\mathbf{p}^w = (x^w, y^w, z^w)^T$ simply by

$$\mathbf{p}^c = {}^c\mathbf{H}_w \cdot \mathbf{p}^w \quad (2.29)$$

The six parameters of this transformation (the three translations t_x , t_y , and t_z and the three rotations α , β , and γ) are called the *external camera parameters* because they determine the position of the camera with respect to the world. In HALCON, they are stored as a pose, i.e. together with a code that describes the order of translation and rotations.

For line scan cameras, the projection of the point \mathbf{p}^c that is given in the camera coordinate system of the first image line into a (sub-)pixel $[r,c]$ in the image is defined as follows:

Assuming

$$\mathbf{p}^c = \begin{pmatrix} x \\ y \\ z \end{pmatrix},$$

the following set of equations must be solved for m , \tilde{u} , and t :

$$\begin{aligned} m \cdot D \cdot \tilde{u} &= x - t \cdot V_x \\ -m \cdot D \cdot p_v &= y - t \cdot V_y \\ m \cdot f &= z - t \cdot V_z \end{aligned}$$

with

$$\begin{aligned} D &= \frac{1}{1 + \kappa(\tilde{u}^2 + p_v^2)} \\ p_v &= S_y \cdot C_y \end{aligned}$$

This already includes the compensation for radial distortions.

Finally, the point is transformed into the image coordinate system, i.e., the pixel coordinate system:

$$c = \frac{\tilde{u}}{S_x} + C_x \quad \text{and} \quad r = t$$

S_x and S_y are scaling factors. S_x represents the distance of the sensor elements on the CCD line, S_y is the extent of the sensor elements in y-direction. The point $(C_x, C_y)^T$ is the principal point. Note that in contrast to area scan images, $(C_x, C_y)^T$ does not define the position of the principal point in image coordinates. It rather describes the relative position of the principal point with respect to the sensor line.

The nine parameters $(f, \kappa, S_x, S_y, C_x, C_y, V_x, V_y, V_z)$ of the pinhole line scan camera are called the *internal camera parameters* because they determine the projection from 3D to 2D performed by the camera.

As for area scan cameras, the calibration of a line scan camera is the process of determining the internal camera parameters $(f, \kappa, S_x, S_y, C_x, C_y, V_x, V_y, V_z)$ and the external camera parameters $(t_x, t_y, t_z, \alpha, \beta, \gamma)$ of the first image line.

Chapter 3

3D Vision in a Specified Plane With a Single Camera

In HALCON it is easy to obtain undistorted measurements in world coordinates from images. In general, this can only be done if two or more images of the same object are taken at the same time with cameras at different spatial positions. This is the so-called *stereo* approach (see [chapter 5](#) on page 105).

In industrial inspection, we often have only one camera available and time constraints do not allow us to use the expensive process of finding corresponding points in the stereo images (the so-called *stereo matching* process).

Nevertheless, it is possible to obtain measurements in world coordinates for objects acquired through telecentric lenses and objects that lie in a known plane, e.g., on an assembly line, for pinhole cameras. Both of these tasks can be solved by intersecting an optical ray (also called line of sight) with a plane.

With this, it is possible to measure objects that lie in a plane, even when the plane is tilted with respect to the optical axis. The only prerequisite is that the camera has been *calibrated*. In HALCON, the calibration process is very easy as can be seen in the example described in [section 3.1](#), which introduces the operators that are necessary for the calibration process.

The easiest way to perform the calibration is to use the HALCON standard calibration plates. You just need to take a few images of the calibration plate (see [figure 3.1](#) for an example), where in one image the calibration plate has been placed directly on the measurement plane.

The sections that follow the example show how to

- calibrate a camera ([section 3.2](#) on page 40),
- transform image into world coordinates ([section 3.3](#) on page 57), and
- rectify images, i.e., remove perspective and/or lens distortions from an image ([section 3.4](#) on page 62).

Finally, we briefly look at the case of inspecting a non-planar object ([section 3.5](#) on page 70).




Figure 3.1: The HALCON calibration plate.

3.1 First Example

The HDevelop example `solution_guide\3d_vision\camera_calibration_multi_image.hdev` shows how easy it is to calibrate a camera and use the calibration results to transform measurements into 3D world coordinates.

First, we specify general parameters for the calibration.

```
create_calib_data ('calibration_object', 1, 1, CalibDataID)
set_calib_data_cam_param (CalibDataID, 0, 'area_scan_division', StartCamPar)
set_calib_data_calib_object (CalibDataID, 0, CaltabName)
```

Then, images of the calibration plate are read, the marks on the plate are extracted, and their coordinates are added to the calibration data.

```
for I := 1 to NumImages by 1
  read_image (Image, ImgPath+'calib_'+I$'02d')
  find_caltab (Image, Caltab, CaltabName, SizeGauss, MarkThresh, \
              MinDiamMarks)
  find_marks_and_pose (Image, Caltab, CaltabName, StartCamPar, \
                      StartThresh, DeltaThresh, MinThresh, Alpha, \
                      MinContLength, MaxDiamMarks, RCoord, CCoord, \
                      StartPose)
  set_calib_data_observ_points (CalibDataID, 0, 0, I, RCoord, CCoord, \
                               'all', StartPose)
endfor
```

Now, we perform the actual calibration with the operator `calibrate_cameras`.

```
calibrate_cameras (CalibDataID, Errors)
```

Afterwards, we can access the calibration results, i.e., the internal camera parameters and the pose of the calibration plate in a reference image.

```
get_calib_data (CalibDataID, 'camera', 0, 'params', CamParam)
get_calib_data (CalibDataID, 'calib_obj_pose', [0, NumImage], 'pose', Pose)
```

This pose is used as the external camera parameters, i.e., the pose of the 3D world coordinate system in camera coordinates. In the example, the world coordinate system is located on the caliper rule (see [figure 3.2](#)). To compensate for the thickness of the calibration plate, the pose is moved by the corresponding value.

```
set_origin_pose (Pose, 0, 0, 0.00075, Pose)
```

Now, we perform the measurement.

```
gen_measure_rectangle2 (195, 226, -0.526627271434, 174, 3, 652, 494, \
    'bilinear', MeasureHandle)
measure_pairs (Image, MeasureHandle, 0.5, 5, 'all', 'all', RowEdgeFirst, \
    ColumnEdgeFirst, AmplitudeFirst, RowEdgeSecond, \
    ColumnEdgeSecond, AmplitudeSecond, IntraDistance, \
    InterDistance)
Row := (RowEdgeFirst+RowEdgeSecond)/2.0
Col := (ColumnEdgeFirst+ColumnEdgeSecond)/2.0
```

With the internal and external camera parameters, the measurement results are transformed into 3D world coordinates with the operator [image_points_to_world_plane](#).

```
image_points_to_world_plane (CamParam, Pose, Row, Col, 'mm', X1, Y1)
```




Figure 3.2: Marks on the caliper rule are measured, the results are transformed into 3D world coordinates with the calibration results.

3.2 3D Camera Calibration

This section describes the process of 3D camera calibration in detail. Note that you can easily calibrate your camera with the help of HDevelop's Calibration Assistant (see the HDevelop User's Guide, [section 6.2](#) on page 200 for more details).

The calibration process consists of three steps:

- the **preparation**,
- the **calibration**, and
- the **access of the results**.

After accessing the results, you can store them and then release the memory of the calibration data model (see [section 3.2.8](#) on page 56).

Preparations Before Performing the Calibration

All information for the calibration is passed in the so-called *calibration data model*. In particular, you

- create the model and specify basic information like the number of cameras to calibrate ([section 3.2.1](#)),
- specify initial values for the internal camera parameters ([section 3.2.2](#)),
- describe the calibration object ([section 3.2.3](#) on page 45),
- observe the calibration object in multiple poses (images) and store the extracted information ([section 3.2.4](#) on page 49), and
- optionally restrict the calibration to certain parameters, keeping the others fixed ([section 3.2.5](#) on page 52).

Performing the Actual Calibration

In HALCON, you calibrate single or multiple cameras with the operator `calibrate_cameras`, which needs the calibration data model as input (see [section 3.2.6](#) on page 52 for more information).

Accessing the Results of the Calibration

`calibrate_cameras` again stores its results, in particular the calibrated camera parameters and poses of the calibration objects, in the calibration data model. You can access them (and all other calibration parameters) with the operator `get_calib_data` (see [section 3.2.7](#) on page 53).

3.2.1 Creating the Calibration Data Model

You create a calibration data model with the operator `create_calib_data`, specifying the number of cameras and the numbers of calibration objects. When using a single camera, you typically use a single calibration object as well.

```
create_calib_data ('calibration_object', 1, 1, CalibDataID)
```

Then, you proceed to

- specify initial values for the internal camera parameters ([section 3.2.2](#)) and
- describe the calibration object ([section 3.2.3](#) on page 45).

3.2.2 Specifying Initial Values for the Internal Camera Parameters

You set internal camera parameters with the operator `set_calib_data_cam_param`.

```
StartCamPar := [0.016,0,0.0000074,0.0000074,326,247,652,494]  
set_calib_data_cam_param (CalibDataID, 0, 'area_scan_division', StartCamPar)
```

Besides the calibration data model, the operator needs the following parameters as input:

- **CameraIdx**: the index of the camera (0 for a single camera),
- **CameraType**: the camera type, and
- **CameraParam**: a tuple with values for the internal camera parameters.

Below, you find

- a list of the supported camera types and their parameters ([section 3.2.2.1](#)),
- tips how to choose the suitable distortion model ([section 3.2.2.2](#)),
- special tips for area scan cameras ([section 3.2.2.3](#) on page 43), and
- special tips for line scan cameras ([section 3.2.2.4](#) on page 44).

3.2.2.1 Camera Type and Camera Parameters

The values of `CameraType` and the sequence of the values of `CameraParam` depend on the used camera and the selected model for the lens distortion. Please, refer to [section 2.2](#) on page 26 for the description of the underlying camera models and to [section 3.2.2.2](#) for tips that help you to decide which distortion model to use. The values of `CameraType` and `CameraParam` are set as follows:

- **Projective area-scan sensors and a division model for the lens distortion:**

CameraType: *'area_scan_division'*

CameraParam: [Focus, Kappa, Sx, Sy, Cx, Cy, ImageWidth, ImageHeight]

- **Projective area-scan sensors and a polynomial model for the lens distortion:**

CameraType: *'area_scan_polynomial'*

CameraParam: [Focus, K1, K2, K3, P1, P2, Sx, Sy, Cx, Cy, ImageWidth, ImageHeight]

- **Telecentric area-scan sensors and a division model for the lens distortion:**

CameraType: *'area_scan_telecentric_division'*

CameraParam: [0, Kappa, Sx, Sy, Cx, Cy, ImageWidth, ImageHeight]

- **Telecentric area-scan sensors and a polynomial model for the lens distortion:**

CameraType: *'area_scan_telecentric_polynomial'*

CameraParam: [0, K1, K2, K3, P1, P2, Sx, Sy, Cx, Cy, ImageWidth, ImageHeight]

- **Line scan cameras:**

CameraType: *'line_scan'*

CameraParam: [Focus, Kappa, Sx, Sy, Cx, Cy, ImageWidth, ImageHeight, Vx, Vy, Vz]

Note that in addition to the internal camera parameters, the width (ImageWidth) and height (ImageHeight) of the image must be given.

For most of the internal camera parameters, initial values can be determined from the specifications of the CCD sensor and the lens. The following sections contain additional tips on how to find suitable initial values for

- area scan cameras ([section 3.2.2.3](#)) and
- line scan cameras ([section 3.2.2.4](#) on page 44).

3.2.2.2 Which Distortion Model to Use

For area scan cameras, two distortion models can be used: The division model and the polynomial model. The division model uses one parameter to model the radial distortions while the polynomial model uses five parameters to model radial and decentering distortions (see [section 2.2.1](#) on page 27).

The advantages of the division model are that the distortions can be applied faster, especially the inverse distortions, i.e., if world coordinates are projected into the image plane. Furthermore, if only few calibration images are used or if the field of view is not covered sufficiently, the division model typically yields more stable results than the polynomial model. The main advantage of the polynomial model is that it can model the distortions more accurately because it uses higher order terms to model the radial distortions and because it also models the decentering distortions. Note that the polynomial model cannot be inverted analytically. Therefore, the inverse distortions must be calculated iteratively, which is slower than the calculation of the inverse distortions with the (analytically invertible) division model.

Typically, the division model should be used for the calibration. If the accuracy of the calibration is not high enough, the polynomial model can be used. But note that the calibration sequence used for the polynomial model must provide a complete coverage of the area in which measurements will later be performed. The distortions may be modeled inaccurately outside of the area that was covered by the calibration plate. This holds for the image border as well as for areas inside the field of view that were not covered by the calibration plate.

3.2.2.3 Tips for Area Scan Cameras

In the following, some hints for the determination of the initial values for the internal camera parameters of an **area scan camera** are given:

Focus f : The initial value is the nominal focal length of the used lens, e.g., 0.016 m. For telecentric cameras you must pass 0.

κ : Use 0.0 as initial value (omitted if the polynomial model is used).
Or:

K1, K2, K3, P1, P2: Use the initial value 0.0 for each of the five coefficients (omitted if the division model is used).

S_x: For pinhole cameras, the initial value for the horizontal distance between two neighboring CCD cells depends on the dimension of the used CCD chip of the camera (see technical specifications of the camera). Generally, common CCD chips are either 1/3"-Chips (e.g., SONY XC-73, SONY XC-777), 1/2"-Chips (e.g., SONY XC-999, Panasonic WV-CD50), or 2/3"-Chips (e.g., SONY DXC-151, SONY XC-77). Notice: The value of S_x increases if the image is sub-sampled! Appropriate initial values are:

	Full image (640*480)	Subsampling (320*240)
1/3"-Chip	0.0000055 m	0.0000110 m
1/2"-Chip	0.0000086 m	0.0000172 m
2/3"-Chip	0.0000110 m	0.0000220 m

The value for S_x is calibrated, since the video signal of a CCD camera normally is not sampled pixel-synchronously.

S_y: Since most off-the-shelf cameras have square pixels, the same values for S_y are valid as for S_x . In contrast to S_x the value for S_y will *not* be calibrated for pinhole cameras because the video signal of a CCD camera normally is sampled line-synchronously. Thus, the initial value is equal to the final value. Appropriate initial values are:

	Full image (640*480)	Subsampling (320*240)
1/3"-Chip	0.0000055 m	0.0000110 m
1/2"-Chip	0.0000086 m	0.0000172 m
2/3"-Chip	0.0000110 m	0.0000220 m

Cx and Cy: Initial values for the coordinates of the principal point are the coordinates of the image center, i.e., half the image width and half the image height. Notice: The values of Cx and Cy decrease if the image is subsampled! Appropriate initial values are, for example:

	Full image (640*480)	Subsampling (320*240)
Cx	320.0	160.0
Cy	240.0	120.0

ImageWidth and ImageHeight: These two parameters are set by the used frame grabber and therefore are not calibrated. Appropriate initial values are, for example:

	Full image (640*480)	Subsampling (320*240)
ImageWidth	640	320
ImageHeight	480	240

3.2.2.4 Tips for Line Scan Cameras

In the following, some hints for the determination of the initial values for the internal camera parameters of a **line scan camera** are given:

Focus f: The initial value is the nominal focal length of the used lens, e.g., 0.035 m.

κ : Use 0.0 as initial value.

Sx: The initial value for the horizontal distance between two neighboring sensor elements can be taken from the technical specifications of the camera. Typical initial values are $7 \cdot 10^{-6}$ m, $10 \cdot 10^{-6}$ m, and $14 \cdot 10^{-6}$ m. Notice: The value of Sx increases if the image is subsampled!

Sy: The initial value for the size of a cell in the direction perpendicular to the sensor line can also be taken from the technical specifications of the camera. Typical initial values are $7 \cdot 10^{-6}$ m, $10 \cdot 10^{-6}$ m, and $14 \cdot 10^{-6}$ m. Notice: The value of Sy increases if the image is subsampled! In contrast to Sx, the value for Sy will NOT be calibrated for line scan cameras because it cannot be determined separately from the parameter Cy.

Cx: The initial value for the x-coordinate of the principal point is half the image width. Notice: The values of Cx decreases if the image is subsampled! Appropriate initial values are:

Image width:	1024	2048	4096	8192
Cx:	512	1024	2048	4096

Cy: Normally, the initial value for the y-coordinate of the principal point can be set to 0.

ImageWidth and ImageHeight: These two parameters are determined by the used frame grabber and therefore are not calibrated.

V_x, V_y, V_z: The initial values for the x-, y-, and z-component of the motion vector depend on the image acquisition setup. Assuming a fixed camera that looks perpendicularly onto a conveyor belt, such that the y-axis of the camera coordinate system is anti-parallel to the moving direction of the conveyor belt (see [figure 3.3](#) on page 46), the initial values are V_x = V_z = 0. The initial value for V_y can then be determined, e.g., from a line scan image of an object with known size (e.g., calibration plate or ruler):

$$V_y = L[m]/L[row]$$

with:

$$L[m] = \text{Length of the object in object coordinates [meter]}$$

$$L[row] = \text{Length of the object in image coordinates [rows]}$$

If, compared to the above setup, the camera is rotated 30 degrees around its optical axis, i.e., around the z-axis of the camera coordinate system ([figure 3.4](#) on page 47), the above determined initial values must be changed as follows:

$$V_{zx} = \sin(30^\circ) \cdot V_y$$

$$V_{zy} = \cos(30^\circ) \cdot V_y$$

$$V_{zz} = V_z = 0$$

If, compared to the first setup, the camera is rotated -20 degrees around the x-axis of the camera coordinate system ([figure 3.5](#) on page 48), the following initial values result:

$$V_{xx} = V_x = 0$$

$$V_{xy} = \cos(-20^\circ) \cdot V_y$$

$$V_{xz} = \sin(-20^\circ) \cdot V_y$$

The quality of the initial values for V_x, V_y, and V_z are crucial for the success of the whole calibration. If they are not accurate enough, the calibration may fail. [Section 3.2.9.1](#) on page 57 provides you with tips what to do in this case.

3.2.3 Describing the Calibration Object

With the operator `set_calib_data_calib_object` you specify the needed information about the calibration object.

If you are using the HALCON calibration plate, the name of the corresponding description file is sufficient.

```
set_calib_data_calib_object (CalibDataID, 0, CalibName)
```




Figure 3.3: Line scan camera looking perpendicularly onto a conveyor belt.

How to obtain a HALCON calibration plate is explained in [section 3.2.3.1](#).

However, you can also use your own calibration object. Then, you pass the coordinates of the markers instead of the file name to the operator (see [section 3.2.3.2](#) for more information).

3.2.3.1 How to Obtain a Suitable Calibration Plate

The simplest method to determine the camera parameters of a CCD camera is to use the HALCON calibration plate (see [figure 3.6](#) on page 48 for examples). In this case, the whole process of finding the calibration plate, extracting the calibration marks, and determining the correspondences between the extracted calibration marks and the respective 3D world coordinates can be carried out automatically. Even more important, these calibration plates are highly accurate, up to ± 150 nm (nanometers), which is a prerequisite for high accuracy applications. Therefore, we recommend to obtain such a calibration plate from the local distributor from which you purchased HALCON.

The calibration plates are available in different materials (ceramics for front light and glass for back light applications) and sizes (e.g., $0.65 \times 0.65 \text{ mm}^2$, $10 \times 10 \text{ mm}^2$, $200 \times 200 \text{ mm}^2$). Thus, you can choose the one that is optimal for your application. A calibration plate should cover at very least the ninth part of the image area. To improve the calibration quality, we recommend, however, a plate that covers at least a quarter of the image area. Detailed information about the available materials, sizes, and the accuracy can be obtained from your distributor.




Figure 3.4: Line scan camera rotated around the optical axis.

Each calibration plate comes with a description file. Place this file in the subdirectory `calib` of the folder where you installed HALCON. Then, you can use its file name directly in the operator `caltab_points` (see section 3.2.4 on page 49).

For test purposes, you can create a calibration plate yourself with the operator `gen_caltab`. Print the resulting PostScript file and mount it on a planar and rigid surface, e.g., an aluminum plate or a solid cardboard. If you do not mount the printout on a planar and rigid surface, you will not get meaningful results by HALCON's camera calibration as the operator `gen_caltab` assumes that the calibration marks lie within a plane. Such self-made calibration plates should only be used for test purposes as you will not achieve the high accuracy that can be obtained with an original HALCON calibration plate. Note that the printing process is typically not accurate enough to create calibration plates smaller than 3 cm.

3.2.3.2 Using Your Own Calibration Object

With HALCON, you are not restricted to using a planar calibration object like the HALCON calibration plate. You can use a 3D calibration object or even arbitrary characteristic points (natural landmarks). The only requirement is that the 3D world position of the model points is known with high accuracy.

Then, you simply pass the 3D coordinates of all points (markers) of the calibration object as a tuple in the parameter `CalibObjDescr` of `set_calib_data_calib_object`. All x , y , and z coordinates of all points must be packed sequentially in the tuple in the form $[X, Y, Z]$.




Figure 3.5: Line scan camera rotated around the x-axis.





Figure 3.6: Examples of calibration plates used by HALCON.

Note however, that if you use your own calibration object, you cannot use the operators `find_caltab` and `find_marks_and_pose` anymore. Instead, you must determine the 2D locations of the model points and the correspondence to the 3D points by yourself.

3.2.4 Observing the Calibration Object in Multiple Poses (Images)

The main input data for the calibration are the so-called observations. For this, the calibration object is placed in different poses. For each pose, the camera acquires an image. In this image, the markers of the calibration object are extracted, and their (pixel) coordinates are stored in the calibration data model with the operator `set_calib_data_observ_points`, together with the indices of camera, calibration object, and calibration object pose and with a tuple containing the indices of the corresponding markers.

If you are using the HALCON calibration plate, you can easily extract the markers using the operators `find_caltab` and `find_marks_and_pose` and pass 'all' instead of the list of marker correspondences to the parameter `Index`:

```
for I := 1 to NumImages by 1
  ... acquire image ...
  find_caltab (Image, Caltab, CaltabName, SizeGauss, MarkThresh, \
              MinDiamMarks)
  find_marks_and_pose (Image, Caltab, CaltabName, StartCamPar, \
                      StartThresh, DeltaThresh, MinThresh, Alpha, \
                      MinContLength, MaxDiamMarks, RCoord, CCoord, \
                      StartPose)
  set_calib_data_observ_points (CalibDataID, 0, 0, I, RCoord, CCoord, \
                               'all', StartPose)
endfor
```

If you are using your own calibration object, you must extract its markers and determine the correspondences by yourself.

The following sections contain

- rules for acquiring suitable calibration images ([section 3.2.4.1](#)) and
- detailed information about the operators `find_caltab` and `find_marks_and_pose` ([section 3.2.4.2](#)).

3.2.4.1 Rules for Acquiring Calibration Images

If you want to achieve accurate results, please follow the rules given in this section:

- Use a clean calibration plate.
- Cover the whole field of view with multiple images, i.e., place the calibration plate in all areas of the field of view at least once.
- Vary the orientations of the calibration plate. This includes rotations around the x- and y-axis of the calibration plate, such that the perspective distortions of the calibration pattern are clearly visible. Without some tilted calibration plates the focal length can not be calculated properly (a tilt angle of approximately 45 degrees is recommended).
- Use at least 10 – 15 images.

- Use an illumination where the background is darker than the calibration plate.
- The bright parts of the calibration plate should have a gray value of at least 100.
- The contrast between the bright and the dark parts of the calibration plate should be more than 100 gray values.
- Use an illumination where the calibration plate is homogeneous.
- The images should not be overexposed (the gray values of the bright parts of the image should be strictly below 255).
- The diameter of a circle should be at least 10 pixels and the size of the calibration plate should be at least a quarter of the whole image.
- The calibration plate should be completely visible inside the image.
- The images should contain as little noise as possible.
- The images should be sharply focused, i.e., transitions between objects should be clearly delimited.

Note that a good calibration result can be obtained only for a homogeneous distribution of calibration marks within the field of view of the camera. You can imagine the part of the 3D space that corresponds to the field of view as a calibration volume like shown in [figure 3.7](#). There, two poses of calibration plates and the positions of their calibration marks, when seen from different views, are illustrated. You can see, e.g., in the view from side 1, that large parts are not covered by marks. To get a homogeneous distribution of the marks and thus enable a good calibration result, you have to place the calibration plates in your other images so that the empty parts of the calibration volume are minimized for all views. Be aware that when having very small calibration plates (compared to the field of view), this means that it may be necessary to use significantly more than the recommended 10 – 15 calibration images.

If only one image is used for the calibration process or if the orientations of the calibration plate do not vary over the different calibration images it is not possible to determine both the focal length and the pose of the camera correctly; only the ratio between the focal length and the distance between calibration plate and camera can be determined in this case. Nevertheless, it is possible to measure world coordinates in the plane of the calibration plate but it is not possible to adapt the camera parameters in order to measure in another plane, e.g., the plane onto which the calibration plate was placed.

The accuracy of the resulting world coordinates depends — apart from the measurement accuracy in the image — very much on the number of images used for the calibration process. The more images (with significantly different calibration plate poses) are used, the more accurate results will be achieved.

3.2.4.2 Additional Information About `find_caltab` and `find_marks_and_pose`

About `find_caltab`

The operator `find_caltab` searches for the calibration plate based on the knowledge that it appears bright with dark calibration marks on it. The parameter `SizeGauss` determines the size of the Gauss filter that is used to smooth the input image. A larger value leads to a stronger smoothing, which might be necessary if the image is very noisy. After smoothing the image, a thresholding operator with the




Figure 3.7: Investigation of the calibration volume: (left) calibration volume with two calibration plate poses and (right) the corresponding distribution of calibration marks when seen from different views. For a good calibration result, the areas without calibration marks (which are especially large in the view from side 1) have to be minimized by a cautious selection of the further calibration plate poses.

minimum gray value `MarkThresh` and the maximum gray value 255 is applied to find the calibration plate. Therefore, `MarkThresh` should be set to a gray value that is lower than that of the white parts of the calibration plate, but preferably higher than that of any other large bright region in the image. Among the regions resulting from the threshold operation, the most convex region with an almost correct number of holes (corresponding to the dark marks of the calibration plate) is selected. Holes with a diameter smaller than `MinDiamMarks` are eliminated to reduce the impact of noise. The number of marks is read from the calibration plate description file `CalTabDescrFile`.

About `find_marks_and_pose`

The operator `find_marks_and_pose` extracts the calibration marks and precisely determines their image coordinates. Therefore, in the input image (`Image`) an edge detector is applied to the region `CalTabRegion`. This region is typically obtained with the operator `find_caltab`. The edge detector can be controlled via the parameter `Alpha`. Larger values for `Alpha` lead to a higher sensitivity of the edge detector with respect to small details, but also to a decreased robustness against noise.

Within the internally resulting edge image, closed contours are extracted. For the detection of the contours a threshold operator is applied to the amplitude of the edges. All points with a high amplitude (i.e., borders of marks) are selected. First, the threshold value is set to `StartThresh`. If the search for the closed contours or the successive pose estimation fails, this threshold value is successively decreased by `DeltaThresh` down to the minimum value specified in `MinThresh`.

The number of closed contours must correspond to the number of calibration marks as described in

the calibration plate description file `CalTabDescrFile` and the contours must have an elliptical shape. Contours shorter than `MinContLength` are discarded, just as contours enclosing regions with a diameter larger than `MaxDiamMarks` (e.g., the border of the calibration plate).

The image coordinates of the calibration marks are determined by applying `find_marks_and_pose` for each image separately. They must be concatenated such that all row coordinates are together in one tuple and all column coordinates are in a second tuple.

The length of these tuples depends on the number of calibration marks and on the number of calibration images. Assuming a calibration plate with m calibration marks and l calibration images, the tuples containing all the image coordinates of the calibration marks have a length of $m \cdot l$, because they contain the coordinates of the m calibration marks extracted from each of the l images. The order of the values is “image by image”, i.e., the first m values correspond to the coordinates of the m calibration marks extracted from the first image. The order of the coordinates of the calibration marks of a single image corresponds to the order in which they appear in the parameters `X`, `Y`, and `Z` that are returned by the operator `caltab_points`. The next m values correspond to the marks extracted from the second image, etc.

Note that the order of all the parameter values must be followed strictly. Therefore, it is very important that each calibration mark is extracted in each image.

3.2.5 Restricting the Calibration to Specific Parameters

If certain camera parameters are already known, you can exclude them from the calibration with the operator `set_calib_data`. Analogously, you can restrict the calibration to certain parameters.

Please refer to the Reference Manual for more information and a short example.

3.2.6 Performing the Calibration

After preparing the calibration data model as described in the previous sections, you perform the calibration by calling the operator `calibrate_cameras`, using the calibration data model as input.

```
calibrate_cameras (CalibDataID, Errors)
```

As a direct result, only the calibration error is returned. It corresponds to the average distance (in pixels) between the backprojected calibration points and their extracted image coordinates. An error of up to 1.5 pixels indicates that the calibration was successful. You can further analyse the quality of the calibration results with the operator `get_calib_data` (see the Reference Manual for details).

The main results of the calibration, e.g., the internal camera parameters, are stored in the calibration data model. How to access them is described in the following section.

If the calibration fails, please refer to [section 3.2.9](#) on page 57 for additional information.

3.2.7 Accessing the Results of the Calibration

The main results of the operator `calibrate_cameras` comprise the internal camera parameters and the pose of the calibration plate in each of the images from which the corresponding points were determined. The operator stores them in the calibration data model. You can access them with the operator `get_calib_data`.

The example program `solution_guide\3d_vision\camera_calibration_interior.hdev` shows how to access the internal camera parameters and write them into a file.

```
get_calib_data (CalibDataID, 'camera', 0, 'params', CamParam)
write_cam_par (CamParam, 'camera_parameters.dat')
```

The external camera parameters cannot be queried directly, because the needed information about the world coordinate system is not stored in the calibration data model. However, if the calibration plate was placed directly on the measurement plane, its pose can be used to easily derive the external camera parameters, which are the pose of the measurement plane. This is described in the following section.

3.2.7.1 Determining the External Camera Parameters

The external camera parameters describe the relation between the measurement plane and the camera, i.e., only if the external parameters are known it is possible to transform coordinates from the camera coordinate system (CCS) into the coordinate system of the measurement plane and vice versa. In HALCON, the measurement plane is defined as the plane $z = 0$ of the world coordinate system (WCS).

The external camera parameters can be determined in different ways:

1. Use the pose obtained from one of the calibration images in which the calibration plate is placed directly on the measurement plane. In this case, you just need to access this pose with the operator `get_calib_data`.
2. Separate the determination of the internal camera parameters from the determination of the external camera parameters by using an additional image in which the calibration plate is placed directly on the measurement plane. Apply `find_caltab` and `find_marks_and_pose` to extract the calibration marks and the pose.
3. Determine the correspondences between 3D world points and their projections in the image by yourself and then call `vector_to_pose`.

If you only need to accurately measure the dimensions of an object, regardless of the absolute position of the object in a given coordinate system, one of the first two cases can be used.

The latter two cases have the advantage that the external camera parameters can be determined independently from the internal camera parameters. This is more flexible and might be useful if the measurements should be done in several planes from a single camera or if it is not possible to calibrate the camera in situ.

In the following, the different cases are described in more detail.

Placing the Calibration Plate on the Measurement Plane in One of the Calibration Images

The **first** case is the easiest way of determining the external parameters. The calibration plate must be placed directly on the measurement plane, e.g., the assembly line, in one of the (many) images used for the determination of the internal parameters.

Since the pose of the calibration plate is determined by the operator `calibrate_cameras`, you can simply access its pose with the operator `get_calib_data`. This way, internal and external parameters are determined in one single calibration step as is shown in the HDevelop example program `solution_guide\3d_vision\camera_calibration_multi_image.hdev`. Here, the pose of the calibration plate (calibration object index 0) in the eleventh calibration image is determined. Please note that each pose consists of seven values.

```
NumImage := 11
get_calib_data (CalibDataID, 'calib_obj_pose', [0, NumImage], 'pose', Pose)
```

The resulting pose would be the true pose of the measurement plane if the calibration plate were infinitely thin. Because real calibration plates have a thickness $d > 0$, the pose of the calibration plate is shifted by an amount $-d$ perpendicular to the measurement plane, i.e., along the z axis of the WCS. To correct this, we need to shift the pose by d along the z axis of the WCS. To perform this shift, the operator `set_origin_pose` can be used.

```
set_origin_pose (Pose, 0, 0, 0.00075, Pose)
```

In general, the calibration plate can be oriented arbitrarily within the WCS as long as the spatial relation between the calibration plate and the measurement plane is known (see [figure 3.8](#)). Then, to derive the pose of the measurement plane from the pose of the calibration plate, a rigid transformation is necessary. In the following example, the pose of the calibration plate is adapted by a translation along the y axis followed by a rotation around the x axis.

```
pose_to_hom_mat3d (FinalPose, HomMat3D)
hom_mat3d_translate_local (HomMat3D, 0, 3.2, 0, HomMat3DTranslate)
hom_mat3d_rotate_local (HomMat3DTranslate, rad(-14), 'x', HomMat3DAdapted)
hom_mat3d_to_pose (HomMat3DAdapted, PoseAdapted)
```

Placing the Calibration Plate on the Measurement Plane in a Separate Image

If the advantages of using the HALCON calibration plate should be combined with the flexibility given by the separation of the internal and external camera parameters the **second** method for the determination of the external camera parameters can be used.

First, the camera is calibrated as described in the previous sections. This can be done, e.g., prior to the mounting of the camera at its final usage site.

Then, after setting up the camera at its final usage site, the external parameters can be determined. The only thing to be done is to take an additional image in which the calibration plate is placed directly on the measurement plane. From this image the external parameters can be determined as is shown in the HDevelop example program `solution_guide\3d_vision\camera_calibration_exterior.hdev`.




Figure 3.8: Relation between calibration plate and measurement plane.

There, the internal camera parameters, the image in which the calibration plate was placed directly on the measurement plane, and the world coordinates of the calibration marks are read from file.

```
read_cam_par ('camera_parameters.dat', CamParam)
read_image (Image, ImgPath+'calib_11')
```

Then, the calibration marks and the pose of the calibration plate are extracted.

```
find_caltab (Image, Caltab, CaltabName, SizeGauss, MarkThresh, MinDiamMarks)
find_marks_and_pose (Image, Caltab, CaltabName, CamParam, StartThresh, \
    DeltaThresh, MinThresh, Alpha, MinContLength, \
    MaxDiamMarks, RCoord, CCoord, PoseForCalibrationPlate)
```

Finally, to take the thickness of the calibration plate into account, the z value of the origin given by the camera pose is translated by the thickness of the calibration plate.

```
set_origin_pose (PoseForCalibrationPlate, 0, 0, 0.00075, \
                PoseForCalibrationPlate)
```

Note that it is very important to fix the focus of your camera if you want to separate the calibration process into two steps as described in this section, because changing the focus is equivalent to changing the focal length, which is part of the internal parameters.

Using Known 3D Points and Their Corresponding Image Points

If it is necessary to perform the measurements within a given world coordinate system, the **third** case for the determination of the external camera parameters can be used. Here, you need to know the 3D world coordinates of at least three points that do not lie on a straight line. Then, you must determine the corresponding image coordinates of the projections of these points. Now, the operator `vector_to_pose` can be used for the determination of the external camera parameters.

An example for this possibility of determining the external parameters is given in the following program. First, the world coordinates of three points are set.

```
X := [0,50,100,80]
Y := [5,0,5,0]
Z := [0,0,0,0]
```

Then, the image coordinates of the projections of these points in the image are determined. In this example, they are simply set to some approximate values. In reality, they should be determined with subpixel accuracy since they define the external camera parameters.

```
RCoord := [414,227,85,128]
CCoord := [119,318,550,448]
```

Finally, the operator `vector_to_pose` is called with the correspondences and the internal camera parameters.

```
vector_to_pose (X, Y, Z, RCoord, CCoord, CamParam, 'iterative', 'error', \
              FinalPose, Errors)
```

Again, it is very important to fix the focus of your camera because changing the focus is equivalent to changing the focal length, which is part of the internal parameters.

3.2.8 Saving the Results and Destroying the Calibration Data Model

After accessing the results (and perhaps storing them using the operators `write_cam_par` and `write_pose`), you can destroy the calibration data model with the operator `clear_calib_data`.

3.2.9 Troubleshooting

Below, you find information for the case that the calibration of a line scan camera fails.

3.2.9.1 Problems With Calibrating Line Scan Cameras

In general, the procedure for the calibration of line scan cameras is identical to the one for the calibration of area scan cameras.

However, line scan imaging suffers from a high degree of parameter correlation. For example, any small rotation of the linear array around the x-axis of the camera coordinate system can be compensated by changing the y-component of the translation vector of the respective pose. Even the focal length is correlated with the scale factor Sx and with the z-component of the translation vector of the pose, i.e., with the distance of the object from the camera.

The consequences of these correlations for the calibration of line scan cameras are that some parameters cannot be determined with high absolute accuracy. Nevertheless, the set of parameters is determined consistently, what means that the world coordinates can be measured with high accuracy.

Another consequence of the parameter correlations is that the calibration may fail in some cases where the start values for the internal camera parameters are not accurate enough. If this happens, try the following approach: In many cases, the start values for the motion vector are the most difficult to set. To achieve better start values for the parameters Vx , Vy , and Vz , reduce the number of parameters to be estimated such that the camera calibration succeeds. Try first to estimate the parameters Vx , Vy , Vz , α , β , γ , tx , ty , and tz by calling `set_calib_data` with `ItemType = 'camera'`, `ItemIdx = 'general'`, `DataName = 'calib_settings'`, and `DataValue = ['vx', 'vy', 'vz', 'alpha', 'beta', 'gamma', 'transx', 'transy', 'transz']` and if this does not work, try `DataValue = ['vx', 'vy', 'vz', 'transx', 'transy', 'transz']`. Then, determine the whole set of parameters using the above determined values for Vx , Vy , and Vz as start values.

If none of the above proposed tips works, try to determine better start values directly from the camera setup. If possible, change the setup such that it is easier to determine appropriate start values, e.g., mount the camera such that it looks approximately perpendicularly onto the conveyor belt (see [figure 3.3](#) on page 46).

If the calibration plates lie in a plane, you may get the error 8440 (“Camera calibration did not converge”), because in this case the parameters are even more correlated. A possible solution may be to exclude the rotation around the z axis, i.e., ‘alpha’ from the calibration with the operator `set_calib_data`.

3.3 Transforming Image into World Coordinates and Vice Versa

In this section, you learn how to obtain world coordinates from images based on the calibration data. On the one hand, it is possible to process the images as usual and then to transform the extraction results into the world coordinate system. In many cases, this will be the most efficient way of obtaining world coordinates. On the other hand, some applications may require that the segmentation itself must be

carried out in images that are already transformed into the world coordinate system (see [section 3.4](#) on page 62).

In general, the segmentation process reduces the amount of data that needs to be processed. Therefore, rectifying the segmentation results is faster than rectifying the underlying image. What is more, it is often better to perform the segmentation process directly on the original images because smoothing or aliasing effects may occur in the rectified image, which could disturb the segmentation and may lead to inaccurate results. These arguments suggest to rectify the segmentation results instead of the images.

In the following, first some general remarks on the underlying principle of the transformation of image coordinates into world coordinates are given. Then, it is described how to transform points, contours, and regions into the world coordinate system. Finally, we show that it is possible to transform world coordinates into image coordinates as well, e.g., in order to visualize information given in the world coordinate system.

3.3.1 The Main Principle

Given the image coordinates of one point, the goal is to determine the world coordinates of the corresponding point in the measurement plane. For this, the line of sight, i.e., a straight line from the optical center of the camera through the given point in the image plane, must be intersected with the measurement plane (see [figure 3.9](#)).

The calibration data is necessary to transform the image coordinates into camera coordinates and finally into world coordinates.

All these calculations are performed by the operators of the family `..._to_world_plane`.

Again, please remember that in HALCON the measurement plane is defined as the plane $z = 0$ with respect to the world coordinate system. This means that all points returned by the operators of the family `..._to_world_plane` have a z -coordinate equal to zero, i.e., they lie in the plane $z = 0$ of the world coordinate system.

3.3.2 World Coordinates for Points

The world coordinates of an image point (r, c) can be determined using the operator [image_points_to_world_plane](#). In the following code example, the row and column coordinates of pitch lines are transformed into world coordinates.

```
image_points_to_world_plane (CamParam, FinalPose, RowPitchLine, \
                             ColPitchLine, 1, X1, Y1)
```

As input, the operator requires the internal and external camera parameters as well as the row and column coordinates of the point(s) to be transformed.

Additionally, the unit in which the resulting world coordinates are to be given is specified by the parameter [Scale](#) (see also the description of the operator [image_to_world_plane](#) in [section 3.4.1](#) on page

set the unit of the resulting coordinates directly by setting the parameter `Scale` to `'m'` (corresponding to the value `1.0`, which could be set alternatively for the parameter `Scale`), `'cm'` (`0.01`), `'mm'` (`0.001`), `'microns'` (`1e-6`), or `'μm'` (again, `1e-6`). Then, if the parameter `Scale` is set to, e.g., `'m'`, the resulting coordinates are given in meters. If, e.g., the coordinates of the calibration target are given in μm and the resulting coordinates have to be given in millimeters, the parameter `Scale` must be set to:

$$\text{Scale} = \frac{\text{mm}}{\mu\text{m}} = \frac{1 \cdot 10^{-3} \text{ m}}{1 \cdot 10^{-6} \text{ m}} = 1000 \quad (3.2)$$

3.3.3 World Coordinates for Contours

If you want to convert an XLD object containing pixel coordinates into world coordinates, the operator `contour_to_world_plane_xld` can be used. Its parameters are similar to those of the operator `image_points_to_world_plane`, as can be seen from the following example program.

```
lines_gauss (ImageReduced, Lines, 1, 3, 8, 'dark', 'true', 'bar-shaped', \
            'true')
contour_to_world_plane_xld (Lines, ContoursTrans, CamParam, PoseAdapted, 1)
```

3.3.4 World Coordinates for Regions

In HALCON, regions cannot be transformed directly into the world coordinate system. Instead, you must first convert them into XLD contours using the operator `gen_contour_region_xld`, then apply the transformation to these XLD contours as described in the previous section.

If the regions have holes and if these holes would influence your further calculations, set the parameter `Mode` of the operator `gen_contour_region_xld` to `'border_holes'`. Then, in addition to the outer border of the input region the operator `gen_contour_region_xld` returns the contours of all holes.

3.3.5 Transforming World Coordinates into Image Coordinates

In this section, the transformation between image coordinates and world coordinates is performed in the opposite direction, i.e., from world coordinates to image coordinates. This is useful if you want to visualize information given in world coordinates or it may be helpful for the definition of meaningful regions of interest (ROI).

First, the world coordinates must be transformed into the camera coordinate system. For this, the homogeneous transformation matrix ${}^{CCS}\mathbf{H}_{WCS}$ is needed, which can easily be derived from the pose of the measurement plane with respect to the camera by the operator `pose_to_hom_mat3d`. The transformation itself can be carried out using the operator `affine_trans_point_3d`. Then, the 3D coordinates, now given in the camera coordinate system, can be projected into the image plane with the operator `project_3d_point`. An example program is given in the following:

There, the world coordinates of four points defining a rectangle in the WCS are defined.

```
ROI_X_WCS := [-2,-2,112,112]
ROI_Y_WCS := [0,0.5,0.5,0]
ROI_Z_WCS := [0,0,0,0]
```

Then, the transformation matrix ${}^{CCS}\mathbf{H}_{WCS}$ is derived from the respective pose.

```
pose_to_hom_mat3d (FinalPose, CCS_HomMat_WCS)
```

With this transformation matrix, the world points are transformed into the camera coordinate system.

```
affine_trans_point_3d (CCS_HomMat_WCS, ROI_X_WCS, ROI_Y_WCS, ROI_Z_WCS, \
                      CCS_RectangleX, CCS_RectangleY, CCS_RectangleZ)
```

Finally, the points are projected into the image coordinate system.

```
project_3d_point (CCS_RectangleX, CCS_RectangleY, CCS_RectangleZ, CamParam, \
                 RectangleRow, RectangleCol)
```

3.3.6 Compensate for Lens Distortions Only

All operators discussed above automatically compensate for lens distortions. In some cases, you might want to compensate for lens distortions only without transforming results or images into world coordinates.

Note that in the following, only the compensation for radial distortions using the division model is described. The compensation for radial and decentering distortions using the polynomial model is done analogously by replacing $\kappa = 0$ with $K_1 = K_2 = K_3 = P_1 = P_2 = 0$.

The procedure is to specify the original internal camera parameters and those of a virtual camera that does not produce lens distortions, i.e., with $\kappa = 0$.

The easiest way to obtain the internal camera parameters of the virtual camera would be to simply set κ to zero. This can be done directly by changing the respective value of the internal camera parameters.

```
CamParVirtualFixed := CamParOriginal
CamParVirtualFixed[1] := 0
```

Alternatively, the operator `change_radial_distortion_cam_par` can be used with the parameter `Mode` set to `'fixed'` and the parameter `Kappa` set to `0`.

```
change_radial_distortion_cam_par ('fixed', CamParOriginal, 0, \
                                 CamParVirtualFixed)
```

Then, for the rectification of the segmentation results, the HALCON operator `change_radial_distortion_contours_xld` can be used, which requires as input parameters the original and the virtual internal camera parameters. If you want to change the lens distortion of image coordinates (Row, Col), you can alternatively use `change_radial_distortion_points`.

```
change_radial_distortion_contours_xld (Edges, EdgesRectifiedFixed, \
                                       CamParOriginal, CamParVirtualFixed)
```

The rectification of the segmentation results changes the visible part of the scene (see [figure 3.10b](#)). To obtain virtual camera parameters such that the whole image content lies within the visible part of the scene, the parameter `Mode` of the operator `change_radial_distortion_cam_par` must be set to `'fullsize'` (see [figure 3.10c](#)). Again, to eliminate the lens distortions, the parameter `Kappa` must be set to 0, or all coefficients of the polynomial model must be set to zero, respectively.

```
change_radial_distortion_cam_par ('fullsize', CamParOriginal, 0, \
                                  CamParVirtualFullsize)
```

If the lens distortions are eliminated in the image itself using the rectification procedure described in [section 3.4.2](#) on page 70, the mode `'fullsize'` may lead to undefined pixels in the rectified image. The mode `'adaptive'` (see [figure 3.10d](#)) slightly reduces the visible part of the scene to prevent such undefined pixels.

```
change_radial_distortion_cam_par ('adaptive', CamParOriginal, 0, \
                                  CamParVirtualAdaptive)
```

The mode `'preserve_resolution'` (see [figure 3.10e](#)) works similar to the mode `'fullsize'` but prevents undefined pixels by additionally increasing the size of the modified image so that the image resolution does not decrease in any part of the image.

```
change_radial_distortion_cam_par ('preserve_resolution', CamParOriginal, 0, \
                                  CamParVirtualPreservedResolution)
```

Note that this compensation for lens distortions is not possible for line scan images because of the acquisition geometry of line scan cameras. To eliminate radial distortions from segmentation results of line scan images, the segmentation results must be transformed into the WCS (see [section 3.3.2](#) on page 58, [section 3.3.3](#) on page 60, and [section 3.3.4](#) on page 60).

3.4 Rectifying Images

For applications like blob analysis or OCR, it may be necessary to have undistorted images. Imagine that an OCR has been trained based on undistorted image data. Then, it will not be able to recognize characters in heavily distorted images. In such a case, the image data must be rectified, i.e., the lens and perspective distortions must be eliminated before the OCR can be applied.

3.4.1 Transforming Images into the WCS

The operator `image_to_world_plane` rectifies an image by transforming it into the measurement plane, i.e., the plane $z = 0$ of the WCS. The rectified image shows no lens and no perspective distortions. It

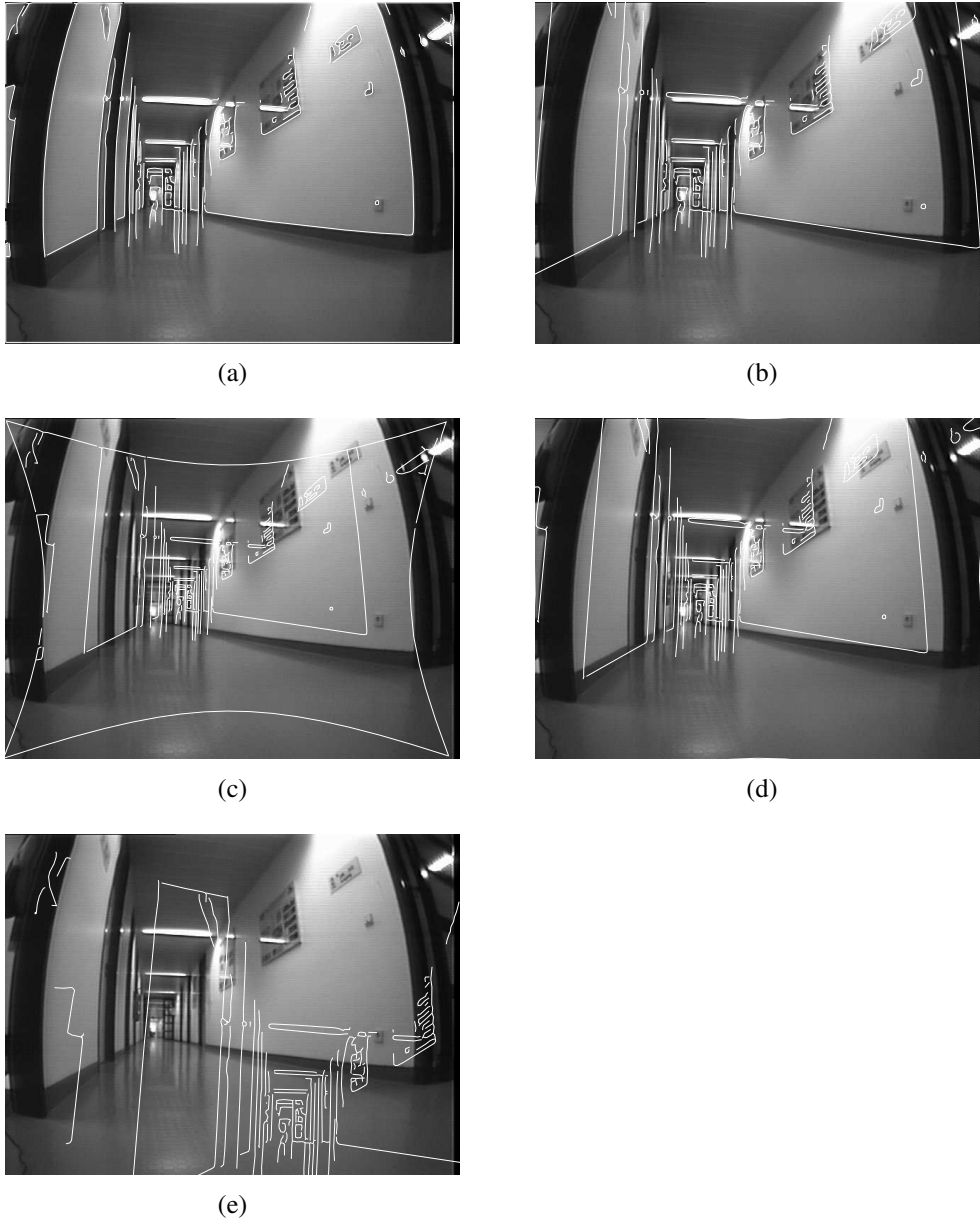


Figure 3.10: Eliminating radial distortions: The original image overlaid with (a) edges extracted from the original image; (b) edges rectified by setting κ to zero; (c) edges rectified with mode 'fullsize'; (d) edges rectified with mode 'adaptive'; (e) edges rectified with mode 'preserved_resolution'.

corresponds to an image captured by a camera that produces no lens distortions and that looks perpendicularly to the measurement plane.

```
image_to_world_plane (Image, ImageMapped, CamParam, PoseForCenteredImage, \
                    WidthMappedImage, HeightMappedImage, \
                    ScaleForCenteredImage, 'bilinear')
```

If more than one image must be rectified, a projection map can be determined with the operator `gen_image_to_world_plane_map`, which is used analogously to the operator `image_to_world_plane`, followed by the actual transformation of the images, which is carried out by the operator `map_image`.

```
gen_image_to_world_plane_map (Map, CamParam, PoseForCenteredImage, \
                             WidthOriginalImage, HeightOriginalImage, \
                             WidthMappedImage, HeightMappedImage, \
                             ScaleForCenteredImage, 'bilinear')
map_image (Image, Map, ImageMapped)
```

The size of the rectified image can be chosen with the parameters `Width` and `Height` for the operator `image_to_world_plane` and with the parameters `WidthMapped` and `HeightMapped` for the operator `gen_image_to_world_plane_map`. The size of the rectified image must be given in pixels.

The pixel size of the rectified image is specified by the parameter `Scale` (see also the description of the operator `image_points_to_world_plane` in section 3.3.2 on page 58). This parameter is the ratio between the pixel size of the rectified image and the unit in which the world coordinates of the calibration target are given (equation 3.3).

$$\text{Scale} = \frac{\text{pixel size of rectified image}}{\text{unit of world coordinates of calibration target}} \quad (3.3)$$

In many cases the coordinates of the calibration targets are given in meters. In this case, it is possible to set the pixel size directly by setting the parameter `Scale` to `'m'` (corresponding to the value 1.0 , which could be set alternatively for the parameter `Scale`), `'cm'` (0.01), `'mm'` (0.001), `'microns'` ($1e-6$), or `'μm'` (again, $1e-6$). Then, if the parameter `Scale` is set to, e.g., `'μm'`, one pixel of the rectified image has a size that corresponds to an area of $1 \mu\text{m} \times 1 \mu\text{m}$ in the world. The parameter `Scale` should be chosen such that in the center of the area of interest the pixel size of the input image and of the rectified image is similar. Large scale differences would lead to aliasing or smoothing effects. See below for examples of how the scale can be determined.

The parameter `Interpolation` specifies whether bilinear interpolation (`'bilinear'`) should be applied between the pixels in the input image or whether the gray value of the nearest neighboring pixel (`'none'`) should be used.

The rectified image `ImageWorld` is positioned such that its upper left corner is located exactly at the origin of the WCS and that its column axis is parallel to the x-axis of the WCS. Since the WCS is defined by the external camera parameters `CamPose` the position of the rectified image `ImageWorld` can be translated by applying the operator `set_origin_pose` to the external camera parameters. Arbitrary transformations can be applied to the external camera parameters based on homogeneous transformation matrices. See below for examples of how the external camera parameters can be set.

In figure 3.11, the WCS has been defined such that the upper left corner of the rectified image corresponds to the upper left corner of the input image. To illustrate this, in figure 3.11, the full domain of the rectified

image, transformed into the virtual image plane of the input image, is displayed. As can be seen, the upper left corner of the input image and of the projection of the rectified image are identical.

Note that it is also possible to define the WCS such that the rectified image does not lie or lies only partly within the imaged area. The domain of the rectified image is set such that it contains only those pixels that lie within the imaged area, i.e., for which gray value information is available. In [figure 3.12](#), the WCS has been defined such that the upper part of the rectified image lies outside the imaged area. To illustrate this, the part of the rectified image for which no gray value information is available is displayed dark gray. Also in [figure 3.12](#), the full domain of the rectified image, transformed into the virtual image plane of the input image, is displayed. It can be seen that for the upper part of the rectified image no image information is available.

If several images must be rectified using the same camera parameters the operator `gen_image_to_world_plane_map` in combination with `map_image` is much more efficient than the operator `image_to_world_plane` because the transformation must be determined only once. In this case, a projection map that describes the transformation between the image plane and the world plane is generated first by the operator `gen_image_to_world_plane_map`. Then, this map is used by the operator `map_image` to transform the image very efficiently.

The following example from `solution_guide\3d_vision\transform_image_into_wcs.hdev` shows how to perform the transformation of images into the world coordinate system using the operators `gen_image_to_world_plane_map` together with `map_image` as well as the operator `image_to_world_plane`.

In the first part of the example program the parameters `Scale` and `CamPose` are set such that a given point appears in the center of the rectified image and that in the surroundings of this point the scale of the rectified image is similar to the scale of the original image.




Figure 3.11: Projection of the image into the measurement plane.




Figure 3.12: Projection of the image into the measurement plane with part of the rectified image lying outside the image area.

First, the size of the rectified image is defined.

```
WidthMappedImage := 652
HeightMappedImage := 494
```

Then, the scale is determined based on the ratio of the distance between points in the WCS and of the respective distance in the ICS.

```
Dist_ICS := 1
image_points_to_world_plane (CamParam, Pose, CenterRow, CenterCol, 1, \
                             CenterX, CenterY)
image_points_to_world_plane (CamParam, Pose, CenterRow+Dist_ICS, CenterCol, \
                             1, BelowCenterX, BelowCenterY)
image_points_to_world_plane (CamParam, Pose, CenterRow, CenterCol+Dist_ICS, \
                             1, RightOfCenterX, RightOfCenterY)
distance_pp (CenterY, CenterX, BelowCenterY, BelowCenterX, \
             Dist_WCS_Vertical)
distance_pp (CenterY, CenterX, RightOfCenterY, RightOfCenterX, \
             Dist_WCS_Horizontal)
ScaleVertical := Dist_WCS_Vertical/Dist_ICS
ScaleHorizontal := Dist_WCS_Horizontal/Dist_ICS
ScaleForCenteredImage := (ScaleVertical+ScaleHorizontal)/2.0
```

Now, the pose of the measurement plane is modified such that a given point will be displayed in the center of the rectified image.

```
DX := CenterX-ScaleForCenteredImage*WidthMappedImage/2.0
DY := CenterY-ScaleForCenteredImage*HeightMappedImage/2.0
DZ := 0
set_origin_pose (Pose, DX, DY, DZ, PoseForCenteredImage)
```

These calculations are implemented in the HDevelop procedure `parameters_image_to_world_plane_centered`.

```
procedure parameters_image_to_world_plane_centered (: : CamParam, Pose,
                                                    CenterRow, CenterCol,
                                                    WidthMappedImage,
                                                    HeightMappedImage:
                                                    ScaleForCenteredImage,
                                                    PoseForCenteredImage)
```

which is part of the HDevelop example program

`solution_guide\3d_vision\transform_image_into_wcs.hdev` (see [appendix B.2](#) on page 232).

Finally, the image can be transformed.

```

gen_image_to_world_plane_map (Map, CamParam, PoseForCenteredImage, \
                             WidthOriginalImage, HeightOriginalImage, \
                             WidthMappedImage, HeightMappedImage, \
                             ScaleForCenteredImage, 'bilinear')
map_image (Image, Map, ImageMapped)

```

The second part of the example program `solution_guide\3d_vision\transform_image_into_wcs.hdev` shows how to set the parameters `Scale` and `CamPose` such that the entire image is visible in the rectified image.

First, the image coordinates of the border of the original image are transformed into world coordinates.

```

full_domain (Image, ImageFull)
get_domain (ImageFull, Domain)
gen_contour_region_xld (Domain, ImageBorder, 'border')
contour_to_world_plane_xld (ImageBorder, ImageBorderWCS, CamParam, Pose, 1)

```

Then, the extent of the image in world coordinates is determined.

```

smallest_rectangle1_xld (ImageBorderWCS, MinY, MinX, MaxY, MaxX)
ExtentX := MaxX-MinX
ExtentY := MaxY-MinY

```

The scale is the ratio of the extent of the image in world coordinates and of the size of the rectified image.

```

ScaleX := ExtentX/WidthMappedImage
ScaleY := ExtentY/HeightMappedImage

```

Now, the maximum value must be selected as the final scale.

```

ScaleForEntireImage := max([ScaleX,ScaleY])

```

Finally, the origin of the pose must be translated appropriately.

```

set_origin_pose (Pose, MinX, MinY, 0, PoseForEntireImage)

```

These calculations are implemented in the HDevelop procedure

```

procedure parameters_image_to_world_plane_entire (Image: : CamParam, Pose,
                                                  WidthMappedImage,
                                                  HeightMappedImage:
                                                  ScaleForEntireImage,
                                                  PoseForEntireImage)

```

which is part of the example program `solution_guide\3d_vision\transform_image_into_wcs.hdev` (see [appendix B.3](#) on page 233).

If the object is not planar the projection map that is needed by the operator `map_image` may be determined by the operator `gen_grid_rectification_map`, which is described in [section 11.3](#) on page 220.

If only the lens distortions should be eliminated the projection map can be determined by the operator `gen_radial_distortion_map`, which is described in the following section.

3.4.2 Compensate for Lens Distortions Only

The principle of the compensation for lens distortions has already be described in [section 3.3.6](#) on page 61.

If only one image must be rectified the operator `change_radial_distortion_image` can be used. It is used analogously to the operator `change_radial_distortion_contours_xld` described in [section 3.3.6](#), with the only exception that a region of interest (ROI) can be defined with the parameter `Region`.

```
change_radial_distortion_image (GrayImage, ROI, ImageRectifiedAdaptive, \
                               CamParOriginal, CamParVirtualAdaptive)
```

Again, the internal parameters of the virtual camera that does not show lens distortions can be determined by setting κ to zero for the division model or K_1 , K_2 , K_3 , P_1 , and P_2 to zero for the polynomial model (see [figure 3.13b](#)). Alternatively, the internal parameters of the virtual camera can be obtained by using the operator `change_radial_distortion_cam_par` with the parameter `Mode` set to `'fixed'` (equivalent to setting κ or the coefficients of the polynomial model to zero; see [figure 3.13b](#)), `'adaptive'` (see [figure 3.13c](#)), `'fullsize'` (see [figure 3.13d](#)), or `'preserve_resolution'` (see [figure 3.13e](#)).

If more than one image must be rectified, a projection map can be determined with the operator `gen_radial_distortion_map`, which is used analogously to the operator `change_radial_distortion_image`, followed by the actual transformation of the images, which is carried out by the operator `map_image`, described in [section 3.4.1](#) on page 62. If a ROI is to be specified, it must be rectified separately (see [section 3.3.4](#) on page 60).

```
gen_radial_distortion_map (MapFixed, CamParOriginal, CamParVirtualFixed, \
                          'bilinear')
map_image (GrayImage, MapFixed, ImageRectifiedFixed)
```

Note that this compensation for lens distortions is not possible for line scan images because of the acquisition geometry of line scan cameras. To eliminate radial distortions from line scan images, the images must be transformed into the WCS (see [section 3.4.1](#) on page 62).

3.5 Inspection of Non-Planar Objects

Note that the measurements described so far will only be accurate if the object to be measured is planar, i.e., if it has a flat surface. If this is not the case the perspective projection of the pinhole camera (see [equation 2.22](#) on page 30) will make the parts of the object that lie closer to the camera appear bigger than

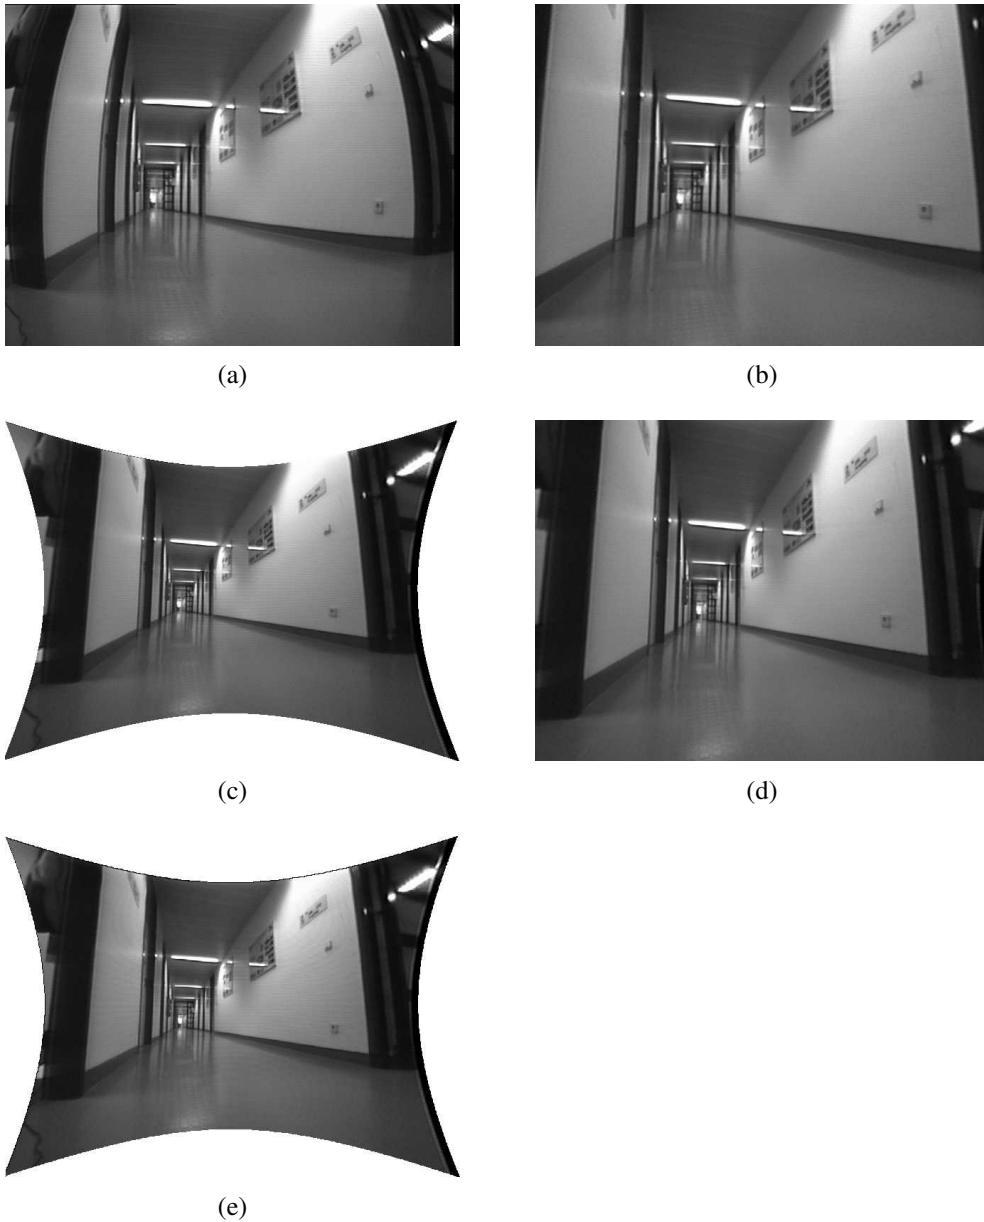


Figure 3.13: Eliminating radial distortions: (a) The original image; (b) the image rectified by setting κ to zero; (c) the image rectified with mode *'fullsize'*; (d) the image rectified with mode *'adaptive'*; (e) the image rectified with mode *'preserved_resolution'*.

the parts that lie farther away. In addition, the respective world coordinates are displaced systematically. If you want to measure the top side of objects with a flat surface that have a significant thickness that is equal for all objects it is best to place the calibration plate onto one of these objects during calibration.

With this, you can make sure that the optical rays are intersected with the correct plane.

The displacement that results from deviations of the object surface from the measurement plane can be estimated very easily. Figure 3.14 shows a vertical section of a typical measurement configuration. The measurement plane is drawn as a thick line, the object surface as a dotted line. Note that the object surface does not correspond to the measurement plane in this case. The deviation of the object surface from the measurement plane is indicated by Δz , the distance of the projection center from the measurement plane by z , and the displacement by Δr . The point N indicates the perpendicular projection of the projection center (PC) onto the measurement plane.




Figure 3.14: Displacement Δr caused by a deviation of the object surface from the measurement plane.

For the determination of the world coordinates of point Q , which lies on the object surface, the optical ray from the projection center of the camera through Q' , which is the projection of Q into the image plane, is intersected with the measurement plane. For this reason, the operators of the family `..._to_world_plane` do not return the world coordinates of Q , but the world coordinates of point P , which is the perspective projection of point Q' onto the measurement plane.

If we know the distance r from P to N , the distance z , which is the shortest distance from the projection center to the measurement plane, and the deviation Δz of the object's surface from the measurement plane, the displacement Δr can be calculated by:

$$\Delta r = \Delta z \cdot \frac{r}{z} \quad (3.4)$$

Often, it will be sufficient to have just a rough estimate for the value of Δr . Then, the values r , z , and Δz can be approximately determined directly from the measurement setup.

If you need to determine Δr more precisely, you first have to calibrate the camera. Then you have to select a point Q' in the image for which you want to know the displacement Δr . The transformation of Q'

into the WCS using the operator `image_points_to_world_plane` yields the world coordinates of point P . Now, you need to derive the world coordinates of the point N . An easy way to do this is to transform the camera coordinates of the projection center PC , which are $(0, 0, 0)^T$, into the world coordinate system, using the operator `affine_trans_point_3d`. To derive the homogeneous transformation matrix ${}^{WCS}H_{CCS}$ needed for the above mentioned transformation, first, generate the homogeneous transformation matrix ${}^{CCS}H_{WCS}$ from the pose of the measurement plane via the operator `pose_to_hom_mat3d` and then, invert the resulting homogeneous transformation matrix (`hom_mat3d_invert`). Because N is the perpendicular projection of PC onto the measurement plane, its x and y world coordinates are equal to the respective world coordinates of PC and its z coordinate is equal to zero. Now, r and z can be derived as follows: r is the distance from P to N , which can be calculated by the operator `distance_pp`; z is simply the z coordinate of PC , given in the WCS.

The following HALCON program (`solution_guide\3d_vision\height_displacement.hdev`) shows how to implement this approach. First, the camera parameters are read from file.

```
read_cam_par ('camera_parameters.dat', CamParam)
read_pose ('pose_from_three_points.dat', Pose)
```

Then, the deviation of the object surface from the measurement plane is set.

```
DeltaZ := 2
```

Finally, the displacement is calculated, according to the method described above.

```
get_mbutton (WindowHandle, RowQ, ColumnQ, Button)
image_points_to_world_plane (CamParam, Pose, RowQ, ColumnQ, 1, WCS_PX, \
                             WCS_PY)
pose_to_hom_mat3d (Pose, CCS_HomMat_WCS)
hom_mat3d_invert (CCS_HomMat_WCS, WCS_HomMat_CCS)
affine_trans_point_3d (WCS_HomMat_CCS, 0, 0, 0, WCS_PCX, WCS_PCY, WCS_PCZ)
distance_pp (WCS_PX, WCS_PY, WCS_PCX, WCS_PCY, r)
z := fabs(WCS_PCZ)
DeltaR := DeltaZ*r/z
```

Assuming a constant Δz , the following conclusions can be drawn for Δr :

- Δr increases with increasing r .
- If the measurement plane is more or less perpendicular to the optical axis, Δr increases towards the image borders.
- At the point N , Δr is always equal to zero.
- Δr increases the more the measurement plane is tilted with respect to the optical axis.

The maximum acceptable deviation of the object's surface from the measurement plane, given a maximum value for the resulting displacement, can be derived by the following formula:

$$\Delta z = \Delta r \cdot \frac{z}{r} \quad (3.5)$$

The values for r and z can be determined as described above.

If you want to inspect an object that has a surface that consists of several parallel planes you can first use [equation 3.5](#) to evaluate if the measurement errors stemming from the displacements are acceptable within your project or not. If the displacements are too large, you can calibrate the camera such that the measurement plane corresponds to, e.g., the uppermost plane of the object. Now, you can derive a pose for each plane, which is parallel to the uppermost plane simply by applying the operator [set_origin_pose](#). This approach is also useful if objects of different thickness may appear on the assembly line. If it is possible to classify these objects into classes corresponding to their thickness, you can select the appropriate pose for each object. Thus, it is possible to derive accurate world coordinates for each object.

Note that if the plane in which the object lies is severely tilted with respect to the optical axis, and if the object has a significant thickness, the camera will likely see some parts of the object that you do not want to measure. For example, if you want to measure the top side of a cube and the plane is tilted, you will see the side walls of the cube as well, and therefore might measure the wrong dimensions. Therefore, it is usually best to align the camera so that its optical axis is perpendicular to the plane in which the objects are measured. If the objects do not have significant thickness, you can measure them accurately even if the plane is tilted.

What is more, it is even possible to derive world coordinates for an object's surface that consists of several non-parallel planes if the relation between the individual planes is known. In this case, you may define the relative pose of the tilted plane with respect to an already known measurement plane.

```
RelPose := [0,3.2,0,-14,0,0,0]
```

Then, you can transform the known pose of the measurement plane into the pose of the tilted plane.

```
pose_to_hom_mat3d (FinalPose, HomMat3D)
pose_to_hom_mat3d (RelPose, HomMat3DRel)
hom_mat3d_compose (HomMat3D, HomMat3DRel, HomMat3DAdapted)
hom_mat3d_to_pose (HomMat3DAdapted, PoseAdapted)
```

Alternatively, you can use the operators of the family `hom_mat3d_..._local` to adapt the pose.

```
hom_mat3d_translate_local (HomMat3D, 0, 3.2, 0, HomMat3DTranslate)
hom_mat3d_rotate_local (HomMat3DTranslate, rad(-14), 'x', HomMat3DAdapted)
hom_mat3d_to_pose (HomMat3DAdapted, PoseAdapted)
```

Now, you can obtain world coordinates for points lying on the tilted plane, as well.

```
contour_to_world_plane_xld (Lines, ContoursTrans, CamParam, PoseAdapted, 1)
```

If the object is too complex to be approximated by planes, or if the relations between the planes are not known, it is not possible to perform precise measurements in world coordinates using the methods described in this section. In this case, it is necessary to use two cameras and to apply the HALCON stereo operators described in [chapter 5](#) on page 105.

Chapter 4

3D Position Recognition of Known Objects

Estimating the 3D pose of an object is an important task in many application areas, e.g., during completeness checks or for 3D alignment in robot vision applications (see [section 8.4.1](#) on page 180). HALCON provides multiple methods to determine the position or pose of known 3D objects.

The most general approach determines the pose of a known 3D object using at least three **corresponding points**, i.e., points with known 3D object coordinates for which the corresponding image coordinates are extracted. The approach is also known as “mono 3D” ([section 4.1](#)).

If a model of a known 3D object is available, **3D matching** can be applied to locate the object. The available 3D matching approaches perform a full 3D object recognition, i.e., they not only estimate a pose but first locate the object in the respective search data. The following approaches are available:

- *Shape-based 3D matching* ([section 4.2](#) on page 81) can be used to locate a complex 3D object in a single 2D image. The model of the 3D object must be available as a Computer Aided Design (CAD) model in, e.g., DXF, OFF, or PLY format (see the Reference Manual entry of [read_object_model_3d](#) for details about the supported formats) and the object needs “hard geometric edges” to be recognized.
- *Surface-based 3D matching* ([section 4.3](#) on page 93) can be used to quickly locate a complex 3D object in a 3D scene, i.e., in a set of 3D points that is available as a so-called 3D object model (see also [section 1](#) on page 12). The model of the 3D object must be available also as a 3D object model and can be obtained either from a CAD model (see the Reference Manual entry of [read_object_model_3d](#) for details about the supported formats) or from a reference 3D scene that is obtained by a 3D reconstruction approach, e.g., stereo or sheet of light. Here, the object may also consist of a “smooth surface”. Note that this approach is also known as “volume matching”.

If the poses of simple 3D shapes like cylinders, spheres, or planes, which are called “3D primitives”, are searched in a 3D scene that is available as a 3D object model, the **3D primitives fitting** ([section 4.4](#) on page 97) can be used. There, the 3D scene is segmented into sub-parts so that into each sub-part a

primitive of a selected type can be fitted. For each sub-part the fitting returns the parameters of the best fitting primitive, e.g., the pose for a fitted plane.

Sometimes, a full 3D object recognition or 3D matching is not necessary because you can estimate the pose of the object with simpler means. For example, if the object contains a characteristic planar part, you can estimate its 3D pose from a single image using **perspective matching**. Similar to the 3D matching approaches, they locate the object before they estimate its pose. Two approaches are available:

- The *calibrated perspective deformable matching* determines the 3D pose of a planar object that was defined by a template object by using automatically derived contours of the object (section 4.5 on page 101).
- The *calibrated descriptor-based matching* determines the 3D pose of a planar object that was defined by a template object by using automatically derived distinctive points of the object, which are called “interest points” (section 4.6 on page 101).

If a circle or rectangle is contained in the plane for which the 3D pose is needed, the pose estimation can be applied also by a simple **circle pose** (section 4.7 on page 102) or **rectangle pose** (section 4.8 on page 102) **estimation**. There, the circle or rectangle must be extracted from the image and the internal camera parameters as well as the dimensions of the circle or rectangle must be known.

An example application for pose estimation in a robot vision system is described in section 8.4.3 on page 182. Note that an introduction to the different 3D matching approaches can be found also in the Solution Guide I, chapter 10 on page 127. The approaches for the perspective matching are described in more detail in the Solution Guide II-B.

4.1 Pose Estimation from Points

If the internal camera parameters are known, the pose of an object can be determined by a call of the operator `vector_to_pose`.

The individual steps are illustrated based on the example program `solution_guide\3d_vision\pose_of_known_3d_object.hdev`, which determines the pose of a metal part with respect to a given world coordinate system.

First, the camera must be calibrated, i.e., the internal camera parameters and, if the pose of the object is to be determined relative to a given world coordinate system, the external camera parameters must be determined. See section 3.2 on page 40 for a detailed description of the calibration process. The world coordinate system can either be identical to the calibration plate coordinate system belonging to the calibration plate from one of the calibration images, or it can be modified such that it fits to some given reference coordinate system (figure 4.1). This can be achieved, e.g., by using the operator `set_origin_pose`

```
set_origin_pose (PoseOfWCS, -0.0568, 0.0372, 0, PoseOfWCS)
```

or if other transformations than translations are necessary, via homogeneous transformation matrices (section 2.1 on page 15).

```
pose_to_hom_mat3d (PoseOfWCS, camHwcs)
hom_mat3d_rotate_local (camHwcs, rad(180), 'x', camHwcs)
hom_mat3d_to_pose (camHwcs, PoseOfWCS)
```

With the homogeneous transformation matrix ${}^c\mathbf{H}_w$, which corresponds to the pose of the world coordinate system, world coordinates can be transformed into camera coordinates.




Figure 4.1: Determination of the pose of the world coordinate system.

Then, the pose of the object can be determined from at least three points (control points) for which both the 3D object coordinates and the 2D image coordinates are known.

The 3D coordinates of the control points need to be determined only once. They must be given in a coordinate system that is attached to the object. You should choose points that can be extracted easily and accurately from the images. The 3D coordinates of the control points are then stored in three tuples, one for the x coordinates, one for the y coordinates, and the last one for the z coordinates.

In each image from which the pose of the object should be determined, the control points must be extracted. This task depends heavily on the object and on the possible poses of the object. If it is known that the object will not be tilted with respect to the camera the detection can, e.g., be carried out by shape-based matching (for a detailed description of shape-based matching, please refer to the Solution Guide II-B, [section 3.3](#) on page 64).

Once the image coordinates of the control points are determined, they must be stored in two tuples that contain the row and the column coordinates, respectively. Note that the 2D image coordinates of the control points must be stored in the same order as the 3D coordinates.

In the example program, the centers of the three holes of the metal part are used as control points. Their image coordinates are determined with the HDevelop procedure [determine_control_points](#),

```
procedure determine_control_points (Image: Intersections: : RowCenter,
                                   ColCenter)
```

which is part of the example program `solution_guide\3d_vision\pose_of_known_3d_object.hdev`.

Now, we simply call the operator [vector_to_pose](#), passing the 3D object coordinates and the 2D image coordinates of the control points together with the internal camera parameters.

```
vector_to_pose (ControlX, ControlY, ControlZ, RowCenter, ColCenter, \
               CamParam, 'iterative', 'error', PoseOfObject, Errors)
```

If both the pose of the world coordinate system and the pose of the object coordinate system are known with respect to the camera coordinate system (see [figure 4.2](#)), it is easy to determine the transformation matrices for the transformation of object coordinates into world coordinates and vice versa:

$${}^w\mathbf{H}_o = {}^w\mathbf{H}_c \cdot {}^c\mathbf{H}_o \quad (4.1)$$

$$= ({}^c\mathbf{H}_w)^{-1} \cdot {}^c\mathbf{H}_o \quad (4.2)$$

where ${}^w\mathbf{H}_o$ is the homogeneous transformation matrix for the transformation of object coordinates into world coordinates and ${}^c\mathbf{H}_w$ and ${}^c\mathbf{H}_o$ are the homogeneous transformation matrices corresponding to the pose of the world coordinate system and the pose of the object coordinate system, respectively, each with respect to the camera coordinate system.

The transformation matrix for the transformation of world coordinates into object coordinates can be derived by:

$${}^o\mathbf{H}_w = ({}^w\mathbf{H}_o)^{-1} \quad (4.3)$$

The calculations described above can be implemented in HDevelop as follows. First, the homogeneous transformation matrices are derived from the respective poses.




Figure 4.2: Pose of the object coordinate system and transformation between object coordinates and world coordinates.

```
pose_to_hom_mat3d (PoseOfWCS, camHwcs)
pose_to_hom_mat3d (PoseOfObject, camHobj)
```

Then, the transformation matrix for the transformation of object coordinates into world coordinates is derived.

```
hom_mat3d_invert (camHwcs, wcsHcam)
hom_mat3d_compose (wcsHcam, camHobj, wcsHobj)
```

Now, known object coordinates can be transformed into world coordinates with `affine_trans_point_3d`.

```
affine_trans_point_3d(wcsHobj, CornersXObj, CornersYObj, CornersZObj, \
                    CornersXWCS, CornersYWCS, CornersZWCS)
```

In the example program `solution_guide\3d_vision\pose_of_known_3d_object.hdev`, the world coordinates of the four corners of the rectangular hole of the metal part are determined from their respective object coordinates. The object coordinate system and the world coordinate system are visualized as well as the respective coordinates for the four points (see [figure 4.3](#)).




Figure 4.3: Object coordinates and world coordinates for the four corners of the rectangular hole of the metal part.

4.2 Pose Estimation Using Shape-Based 3D Matching

For the 3D pose estimation with shape-based 3D matching, a 3D shape model is generated from a 3D computer aided design (CAD) model. The 3D shape model consists of 2D projections of the 3D object seen from different views. To restrain the needed memory and runtime for the shape-based 3D matching, you should restrict the allowed pose range of the shape model and thus minimize the number of 2D projections that have to be computed and stored in the 3D shape model. Analogously to the shape-based matching of 2D structures described in the Solution Guide II-B, [section 3.3](#) on page 64, the 3D shape model is used to recognize instances of the object in the image. But here, instead of a 2D position, orientation, and scaling, the 3D pose of each instance is returned.

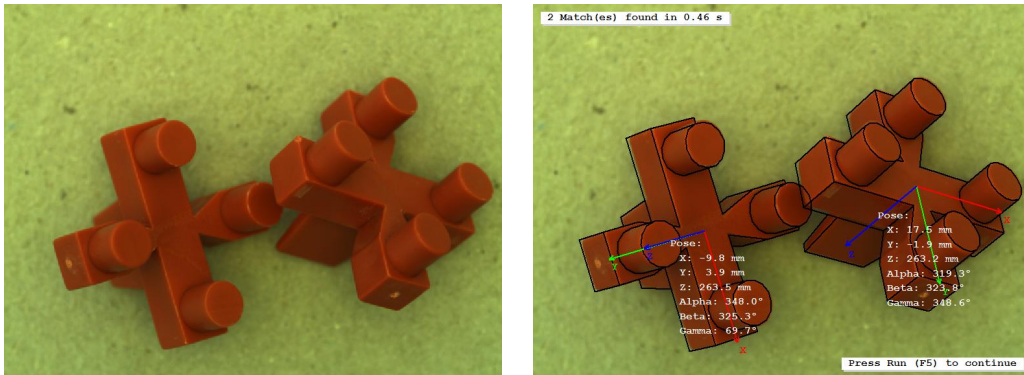


Figure 4.4: Shape-based 3D matching: (left) original image containing two tile spacers, (right) 3D shape model of the tile spacer projected in the image with the poses of the found model instances.

If you need the 3D pose of a planar object or a planar object part, we recommend to use the calibrated perspective deformable matching (see [section 4.5](#) on page 101) or the calibrated descriptor-based matching (see [section 4.6](#) on page 101). Both are significantly faster, because no 2D projections of the model must be computed. Instead, a single 2D model is derived from images.

In the following, the general proceeding for shape-based 3D Matching is introduced (see [section 4.2.1](#)), it is shown how to generally enhance the matching robustness and speed (see [section 4.2.2](#) on page 86), and tips and tricks for the handling of specific problems are provided (see [section 4.2.3](#) on page 89).

4.2.1 General Proceeding for Shape-Based 3D Matching

Shape-based 3D matching consists of the following basic steps:

- the 3D object model is accessed from file,
- the 3D shape model is created from it,
- the 3D object model is destroyed,
- the 3D shape model is used to search the object in search images, and

- the 3D shape model is destroyed.

An example for the shape-based 3D matching of tile spacers is the HDevelop example program `hdevelop\3D-Matching\Shape-Based\create_shape_model_3d_lowest_model_level.hdev` (see [figure 4.4](#)).

Step 1: Read the 3D object model

The 3D object model describing the search object is loaded in HALCON with the operator `read_object_model_3d`. It must be available as a CAD model in one of the supported formats, e.g., DXF, STL, or PLY. The list of supported CAD formats and tips on how to obtain a suitable model, including the specification of the requirements the CAD models must fulfill, are provided with the description of the operator in the Reference Manual. For additional tips on handling selected CAD formats, please contact your distributor.

```
read_object_model_3d ('tile_spacer.dxf', 0.0001, [], [], ObjectModel3DID, \
                    DXFStatus)
```

Step 2: Create the 3D shape model

The 3D shape model is created with the operator `create_shape_model_3d`. It needs the 3D object model and camera parameters as input. Additionally, a set of parameters has to be adjusted. The camera parameters can be obtained by a camera calibration as is described in detail in [section 3.2](#) on page 40. In the example, the camera parameters are known and just assigned to the variable `CamParam`. Before creating the 3D shape model, it is recommended to prepare the 3D object model for the shape-based 3D matching using `prepare_object_model_3d`. Otherwise, the preparation is applied internally within `create_shape_model_3d`, which may slow down the application if the same 3D object model is used several times.

```
CamParam := [0.0269462, -354.842, 1.27964e-005, 1.28e-005, 254.24, 201.977, \
            512, 384]
prepare_object_model_3d (ObjectModel3DID, 'shape_based_matching_3d', 'true', \
                        [], [])
create_shape_model_3d (ObjectModel3DID, CamParam, 0, 0, 0, 'gba', -rad(55), \
                    rad(55), -rad(55), rad(55), 0, rad(360), 0.26, 0.27, \
                    10, 'lowest_model_level', 3, ShapeModel3DID)
```

The 3D shape model 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.

An important task is to specify the pose range. To ease this task, imagine a sphere that surrounds the object. On the surface of the sphere, a camera is placed that looks at the object. Now, the pose range can be defined by restricting the position of the camera to a part of the sphere's surface. Additionally, the minimum and maximum distance of the camera to the object, i.e., the radii of different spheres, must be specified.

In the following, the position of the sphere relative to the object and the definition of the surface part are described. The position of the sphere is defined by placing its center at the center of the object's

bounding box. The radius of the sphere corresponds to the distance of the camera to the center of the object. To define a specific part of the sphere's surface, the geographical coordinates longitude (λ) and latitude (φ) are used (see figure 4.5a). For these, minimum and maximum values are specified so that a quadrilateral on the sphere is obtained (see figure 4.5b).




Figure 4.5: Geographic coordinate system: (a) the geographical coordinates longitude (λ) and latitude (φ) describe positions on the sphere's surface, (b) the minimum and maximum values for λ and φ describe a quadrilateral on the sphere, which defines the pose range.

To describe the orientation of the geographical coordinate system we introduce an object-centered coordinate system. This is obtained by moving the origin of the object coordinate system (CAD model) to the center of the sphere. The xz -plane of the object-centered coordinate system defines the equatorial plane of the geographical coordinate system. The north pole lies on the negative y -axis. The origin of the geographical coordinate system ($\lambda = \varphi = 0^\circ$), i.e., the intersection of the equator with the zero meridian, lies in the negative z -axis.

To illustrate the above description, let us assume that we have the object model shown in figure 4.6a. For illustrative purpose additionally the object coordinate system is visualized. In figure 4.6b, the corresponding geographical coordinate system is shown together with a camera placed at its origin ($\lambda = \varphi = 0^\circ$). Consequently, this camera view corresponds to a bottom view of the object.

Note that the coordinate system introduced here is only used to specify the pose range. The pose resulting from the shape-based 3D matching always refers to the original object coordinate system used in the CAD file and not to the center of the object's bounding box.

In most cases, the specification of the pose range can be simplified by changing the origin of the geographical coordinate system (i.e., the orientation of the sphere) such that it coincides with a mean viewing direction of the real camera to the object. This can be achieved by rotating the object-centered coordinate system, which defines the geographical coordinate system as described above. The rotation can be specified by passing the rotation angles to the parameters `RefRotX`, `RefRotY`, and `RefRotZ` of the operator `create_shape_model_3d`. That is, you can specify the pose range either by adjusting the longitude and latitude, leaving the origin of the sphere at its initial position, or by rotating the object-centered

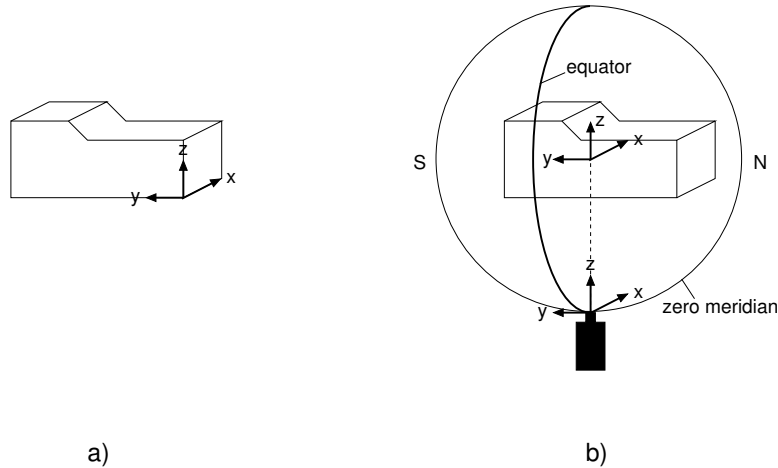


Figure 4.6: Object coordinate systems: (a) original object coordinate system of the CAD model, (b) object-centered geographical coordinate system obtained by moving the object to the center of the sphere: the camera is placed on the sphere's surface at the position $\lambda = \varphi = 0^\circ$ and the object is placed at the center of the sphere. Note that in contrast to the previous image, the sphere is tilted, i.e., it is visualized in a way that the equator is vertical and the zero meridian is completely visible.

coordinate system to get a mean reference view and then specify the pose range in a more intuitive way (see [figure 4.7](#) on page 85), which is recommended in most cases.

[Figure 4.8](#) shows a rotation-symmetric object, for which a reference view is specified. Rotation-symmetric objects have the specific advantage that they have the same appearance from all directions that are perpendicular to their rotation axis. Thus, if the mean reference view is selected such that the camera view is perpendicular to the rotation axis (and the rotation axis crosses the poles of the sphere), the longitude range can collapse to a single value so that the number of calculated 2D projections is significantly reduced. That is, less memory is needed and the matching becomes faster. Here, the initial z-axis of the object-centered coordinate system corresponds to the rotation axis of the object. To obtain the intended reference view, the orientation of the sphere is changed by rotating the object-centered coordinate system by -90° around its x-axis. Now, the pose range for φ is specified like described before and the pose range for λ is restricted to 0° .

Besides the definition of the pose range, also the camera roll angle, i.e., the allowed range for the rotation of the virtual camera around its z-axis, must be set. In most cases, it is recommended to allow a full circle for the camera roll angle. For details, we recommend to read the description of the operator [create_shape_model_3d](#) in the Reference Manual.

Step 3: Destroy the 3D object model

After creating the 3D shape 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.

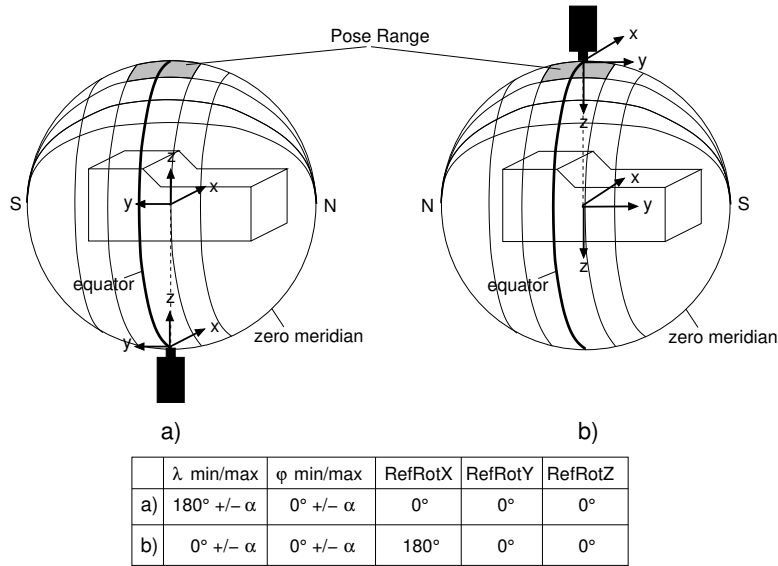


Figure 4.7: Specify pose range: (a) only by longitude and latitude, or (b) by additionally rotating the object-centered coordinate system to a reference view. Note that because of the rotation around the x-axis, the positions of the poles have changed.

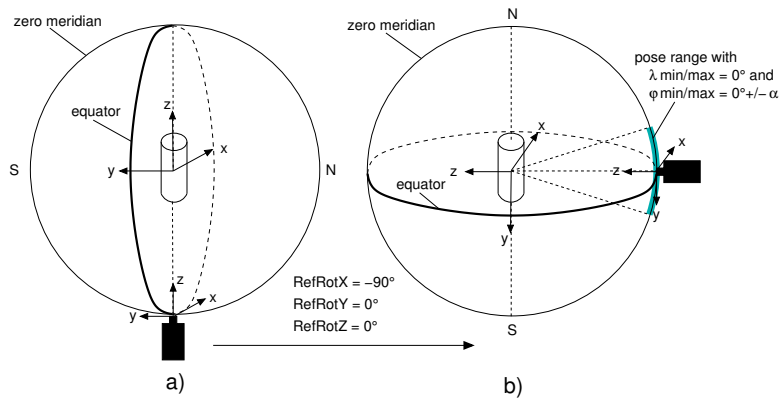


Figure 4.8: Change the origin of the geographical coordinate system for (a) a rotation-symmetric object with the z-axis of the object-centered coordinate system corresponding to the rotation axis such that (b) the z-axis becomes perpendicular to the rotation axis.

```
clear_object_model_3d (ObjectModel3DID)
```

Step 4: Find the 3D shape model in search images

With the 3D shape model that was created by `create_shape_model_3d` or read from file by `read_shape_model_3d`, the object can be searched for in images. For the search, the operator

`find_shape_model_3d` is applied.

```
for I := 1 to NumImages by 1
  read_image (Image, 'tile_spacers/tile_spacers_color_'+I$'02')
  find_shape_model_3d (Image, ShapeModel3DID, 0.7, 0.9, 0, 'num_matches', \
                      3, Pose, CovPose, Score)
endfor
```

Several parameters can be set to control the search process. For detailed information, we recommend to read the description of the operator in the Reference Manual. The operator returns the pose of the matching model, the standard deviation of the pose, and the score of the found instances of the 3D shape model that describes how much of the model is visible in the image.

Step 5: Destroy the 3D shape model

When the 3D shape model is not needed anymore, it is destroyed with the operator `clear_shape_model_3d`.

```
clear_shape_model_3d (ShapeModel3DID)
```

Besides the basic steps, it is often required to inspect the 3D object model or the 3D shape model, to re-use the 3D shape model, or to visualize the result of the matching. These steps are described in the Solution Guide I, [chapter 10](#) on page 127.

4.2.2 Enhance the Shape-Based 3D Matching

The following sections generally show how to enhance the robustness ([section 4.2.2.1](#)) and speed ([section 4.2.2.2](#)) of shape-based 3D matching.

4.2.2.1 Enhance the Robustness

For a robust shape-based 3D matching it is important that the edges of the object are clearly visible in the image. Thus, the following general tips may help you to enhance your application already when acquiring the images of the object:

- If possible, use a background with a good contrast to the object, so that the background can be clearly separated from the object.
- Carefully adjust the lighting for the image acquisition.

To get clearly visible edges of the object in your images, take special care of the lighting conditions. In general, the edges that are included in the 3D shape model should also be visible in the image. The edges that are included in the model can be adjusted with the generic parameter `min_face_angle` of `create_shape_model_3d`. The effect of this parameter can be inspected by visualizing the resulting edges with the procedure `inspect_object_model_3d`, which can be found in the example program `hdevelop\Applications\Position-Recognition-3D\3d_matching_clamps.hdev`.

- If possible, use multi-channel images.

Multi-channel, e.g., color images contain more information and thus typically lead to a more robust edge extraction. An especially robust edge extraction can be obtained if color images are used and the object is illuminated from different directions by three differently colored, typically red, green and blue, light sources (see [figure 4.9](#)).




Figure 4.9: A fuse illuminated from three different directions by three differently colored light sources.

4.2.2.2 Enhance the Speed

There are several means to speed up the shape-based 3D matching:

- Use a homogeneous background during the image acquisition.

Generally, you should try to adjust the lighting for the image acquisition so that the edges of the object but no surface texture are visible in the images.

- Reduce the resolution of the image.

A reduced resolution of the image can speed up the online as well as the offline phase of the shape-based 3D matching, because less 2D projections have to be generated. Note that if you reduce the resolution of the images (e.g., from 1 Megapixel to 640x480 pixels), you must also change the camera parameters used for the creation of the 3D shape model by adapting `Sx`, `Sy`, `Cx`, `Cy`, `ImageWidth`, and `ImageHeight` (see the description of `write_cam_par` in the Reference Manual). For example, if the image is scaled down by a factor of 0.5, `Sx` and `Sy` must be multiplied by 2, whereas `Cx`, `Cy`, `ImageWidth`, and `ImageHeight` must be multiplied by 0.5.

- Use a region of Interest.

Using a region of interest, you can speed up the search. 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 Solution Guide I, [chapter 3](#) on page 19 or the Solution Guide II-B, [section 2.1.2](#) on page 21.

- Restrict the pose range.

The more the pose range is restricted while creating a 3D shape model with `create_shape_model_3d`, the faster is the search process. But note that only those object instances are found that correspond to the selected pose range.

- Eliminate unnecessary edges from the 3D shape model.

If edges are contained in the 3D shape model that are not visible in the search image, the performance, i.e., the robustness and speed, of the shape-based 3D matching decreases. Especially for objects that contain curved surfaces, which are approximated by multiple planar faces in the 3D object model, the generic parameter `min_face_angle` should be adjusted within `create_shape_model_3d` to eliminate unnecessary edges from the 3D shape model. To check which edges are visible with a specific minimum face angle, you can display the corresponding contours with the operator `project_object_model_3d`.

- Select the value for the number of used pyramid levels as large as possible.

In `create_shape_model_3d` as well as in `find_shape_model_3d`, the number of used pyramid levels can be set by the parameter `num_levels`. To speed up the search process, the value should be as large as possible, but the object should be still recognizable in the model. To check a view on the object in a specific pyramid level, you can query the corresponding contours by the operator `get_shape_model_3d_contours`.

- Disable the pregeneration of the model views on lower pyramid levels.

When specifying a large pose range, the number of model views on lower pyramid levels may become very large, which leads to a slow model generation and high memory consumption. To speed up the model generation, you can disable the pregeneration of the model views on lower pyramid levels using the generic parameter `'lowest_model_level'` within `create_shape_model_3d`. Note that you nevertheless obtain a high accuracy during the search with `find_shape_model_3d`, because the pose is still refined (but now on the fly) on the original pyramid level.

- Select the value for `MinScore` as large as possible.

In `find_shape_model_3d` you can adjust the parameter `MinScore`. A large `MinScore` speeds up the search process, but allows less invisible edges for the object, so that possibly some objects are not recognized.

- Select the value for `Greediness` as large as possible.

In `find_shape_model_3d` you can adjust the parameter `Greediness`. A large `Greediness` speeds up the search process. But because the search becomes less robust, possibly some objects are not recognized.

- Adjust the pose refinement.

In `find_shape_model_3d` you can adjust the generic parameter `pose_refinement`. If it is set to `none`, the search is fast but the pose is determined with limited accuracy. A tradeoff between runtime and accuracy is to set the pose refinement to `least_squares_high`. For complex models with a large number of faces it is reasonable to speed up the pose refinement by splitting it such that some of the needed calculations are already performed during the creation of the model. Thus, the generic parameter `'fast_pose_refinement'` of `create_shape_model_3d` is by default set to `'true'`. This leads to a faster matching but also to a higher memory consumption. If the storage is more critical than the speed of the matching, you can set the parameter to `'false'`.

- Check the setting of `border_model`.

In `find_shape_model_3d` you can adjust the generic parameter `border_model`. If it is set to `true`, also objects that extend beyond the image borders can be found. Because this is rather time-consuming, if you only search for objects that are completely contained in the image, we recommend to leave the default value `false` unchanged.

- Speed up the visualization.

When visualizing the 3D object model or the 3D shape model by projecting them into the image using `project_object_model_3d` or `project_shape_model_3d`, respectively, you can speed up the visualization by setting the parameter `HiddenLineRemoval` to `false`. Then, also those edges of the model are visualized that would be hidden by faces.

4.2.3 Tips and Tricks for Problem Handling

In [section 4.2.2](#) on page 86, general tips to enhance the robustness and speed of shape-based 3D matching were listed. In the following, the focus is on the handling of specific problems, i.e., possible reasons for an unsuccessful, erroneous, or very slow recognition or model generation are introduced. In particular, possible reasons for problems and tips to solve them are provided for the cases that

- the model generation is very slow (see [section 4.2.3.1](#)),
- the recognition is not successful, i.e., the object is not found (see [section 4.2.3.2](#)),
- the object is found, but the estimated object pose is wrong (see [section 4.2.3.3](#) on page 91),
- the object is found in the right pose, but the pose is estimated with low accuracy (see [section 4.2.3.4](#) on page 91), or
- the recognition is successful but very slow ([section 4.2.3.5](#) on page 92).

4.2.3.1 The Model Generation is Very Slow

The computation time of the model generation increases quadratically with the number of faces in the CAD model. Thus, if the model generation with `create_shape_model_3d` is very slow, most possibly the CAD model is too complex. Note that in most cases, a very coarse model is sufficient to locate a 3D object with shape-based 3D matching. Thus, in case of a very slow model generation, you should eliminate all unimportant details from your model using suitable CAD software. Alternatively, you can also increase the value of the generic parameter `'lowest_model_level'` to work with a coarser model. Additionally, only the part of the object that is relevant for the search should be contained in the model. Thus, sometimes further modifications of the CAD model using suitable CAD software might be necessary.

4.2.3.2 The Object is not Found

If the object is not found with `find_shape_model_3d`, typically the reason can be found in at least one of the following problems:

- The value of the parameter `MinScore` was chosen too large. Check if the object can be found if matches with a smaller score are accepted. But note that the possibility of finding false matches increases with a decreasing `MinScore` value.
- The value of the parameter `Greediness` was chosen too large. A large value leads to a fast but less robust search. Check if the object can be found with a smaller `Greediness` value.
- The value of the parameter `NumLevels` was estimated or chosen too large. Note that the shape representation of the views on the highest pyramid level must still be recognizable and must contain enough model points. You can visually check the views on the specific pyramid levels using `get_shape_model_3d_contours`.
- The value of the generic parameter `'min_face_angle'` was chosen too small when generating the 3D shape model with `create_shape_model_3d`. Thus, the 3D shape model contains also non-visible object edges. Note that after creating a model, you can visually check the model edges on the specific pyramid levels using `get_shape_model_3d_contours`. Before creating a model, you can display the contours of the underlying 3D object model with a specific minimum face angle using the operator `project_object_model_3d`.
- The chosen reference pose or the pose range are not correct. Check whether the created views cover the desired pose range. Note that you can query the number of created views for each pyramid level using `get_shape_model_3d_params` (setting `GenParamNames` to `'num_views_per_level'`). Then, you can visually check selected views out of this range of views using `get_shape_model_3d_contours`.
- The value of the parameter `MinContrast` was chosen too large when generating the 3D shape model with `create_shape_model_3d`. Thus, edges that belong to the object are not extracted in the search image and the matching score decreases. Check the value of `MinContrast` by extracting edges in the search image using `edges_image` with the parameter `Filter` set to `'sobel_fast'` and the parameters `Low` and `High` set to the value of `MinContrast`.
- The camera parameters are not accurate enough. Thus, the projected model and the imaged object do not accurately fit together. If possible, improve the accuracy of the calibration. Otherwise, stop the search on a higher pyramid level where the differences between the projected model and the imaged object are small enough. To stop the search on a higher pyramid level, the parameter `NumLevels` is set as a tuple consisting of the highest and lowest used pyramid levels.
- The projected model and the imaged object do not accurately fit together because the CAD file is not modeled accurately enough or the objects do not exactly correspond to the model, e.g., because of tolerances at the fabrication. In the first case, if possible, improve the accuracy of the CAD model. Otherwise, stop the search on a higher pyramid level where the differences between the projected model and the imaged object are small enough. To stop the search on a higher pyramid level, the parameter `NumLevels` is set as a tuple consisting of the highest and lowest used pyramid levels.
- Some of the object edges are no “sharp” edges. Model the round edges in the CAD model or stop the search on a higher pyramid level where the differences between the projected model and the imaged object are small enough. To stop the search on a higher pyramid level, the parameter `NumLevels` is set as a tuple consisting of the highest and lowest used pyramid levels.
- The object edges are not visible in the image. To enhance the visibility of the edges during the

image acquisition, follow the advices for a robust shape-based matching that are given in [section 4.2.2.1](#) on page 86.

4.2.3.3 The Object is Found in a Wrong Pose

If the object is found with `find_shape_model_3d`, but the estimated pose is wrong, typically the reason can be found in at least one of the following problems:

- The value of the parameter `MinScore` was chosen too small so that false matches could be found. Check if the object still can be found with a larger `MinScore` value.
- The value of the generic parameter `'min_face_angle'` was chosen too large when generating the 3D shape model with `create_shape_model_3d`. Thus, the 3D shape model does not contain all visible object edges. Note that after creating a model, you can visually check the model edges on the specific pyramid levels using `get_shape_model_3d_contours`. Before creating a model, you can display the contours of the underlying 3D object model with a specific minimum face angle using the operator `project_object_model_3d`.
- The value of the parameter `MinContrast` was chosen too small when generating the 3D shape model with `create_shape_model_3d`. This leads to clutter edges in the search image. Check the value of `MinContrast` by extracting edges in the search image using `edges_image` with the parameter `Filter` set to `'sobel_fast'` and the parameters `Low` and `High` set to `MinContrast`.
- The background contains too much clutter. To solve this problem for one-channel images, you can set `'metric'` to `'ignore_part_polarity'`. But if possible, you should use another background, which is more homogeneous. If the background cannot be changed, you can try to optimize its appearance in the images by using a diffuse lighting source and by adjusting the light's direction. Note that you can check the background by extracting edges in the search image using `edges_image` with the parameter `Filter` set to `'sobel_fast'` and the parameters `Low` and `High` set to the value of `MinContrast`.
- The pose range contains degenerated views, i.e., views that do not significantly represent the object anymore. For example, if a cube is viewed exactly orthogonal to one of its faces, the view collapses from a perspective representation of a 3D cube to a simple square. Such a view may lead to many false matches. An even more extreme example is an exact side-view on a flat object. There, the view may collapse to a straight line. Such a view would lead to many matches in any search image, even if the actual 3D object is not contained. Therefore, you should limit the pose range as much as possible and ensure that no degenerated views are contained in the pose range.
- There are too many clutter edges around the object, which is typical for some bin picking applications where objects touch or overlap each other. Try to separate the objects, e.g., by dumping the objects onto a plane and/or shaking the bin or plane to equally distribute the objects.

4.2.3.4 The Object Pose is Estimated With Low Accuracy

If the object is found with `find_shape_model_3d`, but the accuracy of the estimated object pose is low, typically the reason can be found in at least one of the following problems:

- Some of the object edges are no “sharp” edges. Model the round edges in the CAD model or stop the search on a higher pyramid level where the differences between the projected model and the imaged object are small enough. To stop the search on a higher pyramid level, the parameter `NumLevels` is set as a tuple consisting of the highest and lowest used pyramid levels.
- The camera parameters are not accurate enough. Thus, the projected model and the imaged object do not accurately fit together. If possible, improve the accuracy of the calibration.
- The projected model and the imaged object do not accurately fit together because the CAD file is not modeled accurately enough or the objects do not exactly correspond to the model, e.g., because of tolerances at the fabrication. In the first case, if possible, improve the accuracy of the CAD model.

4.2.3.5 The Object Recognition is Very Slow

If the object is found with `find_shape_model_3d`, but the recognition is very slow, typically the reason can be found in at least one of the following problems:

- The value of the parameter `Greediness` was chosen too small. A larger value speeds up the search, but because the search becomes less robust, possibly some objects are not recognized. Thus, you should check if the object can still be found with a larger `Greediness` value.
- The value of the parameter `MinScore` was chosen too small so that many false matches could be found. Check if the object still can be found with a larger `MinScore` value.
- The value of the parameter `MinContrast` was chosen too small when generating the 3D shape model with `create_shape_model_3d`. This leads to clutter edges in the search image. These clutter edges generate many matching candidates that must be examined. Check the value of `MinContrast` by extracting edges in the search image using `edges_image` with the parameter `Filter` set to `'sobel_fast'` and the parameters `Low` and `High` set to the value of `MinContrast`.
- The CAD model is very complex and `'fast_pose_refinement'` was set to `'false'` when generating the 3D shape model with `create_shape_model_3d`. As the computation time of the least-squares refinement increases quadratically with the number of faces in the CAD model, you can speed up the recognition by eliminating all important details from the model using your CAD software and setting `'fast_pose_refinement'` to `'true'`.
- The image size, especially the extent of the search object in pixels, is very large so that the number of model views on lower pyramid levels becomes very large, too. If you cannot work with images of a reduced resolution as is described in [section 4.2.2.2](#) on page 87, we strongly recommend to restrict the search in the image to an ROI and to increase the value of `'lowest_model_level'` when generating the 3D shape model with `create_shape_model_3d`.
- The pose range is too large. When generating the 3D shape model with `create_shape_model_3d`, restrict the pose range as much as possible, especially take care of `DistMin`. Additionally, for objects with multiple stable poses you should use multiple models. That is, instead of generating one model that covers the complete pose range, multiple models each covering one stable pose should be used. Note that you can also restrict the pose range by exploiting rotational symmetries of the object (see [section 4.2.1](#) on page 84).

- The background contains too much clutter so that the clutter edges in the search image lead to many matching candidates that must be examined. To solve this problem for one-channel images, you can set 'metric' to 'ignore_part_polarity'. But if possible, you should use another background, which is more homogeneous. If the background cannot be changed, you can try to optimize its appearance in the images by using a diffuse light source and by adjusting the light's direction. Note that you can check the background by extracting edges in the search image using `edges_image` with the parameter `Filter` set to 'sobel_fast' and the parameters `Low` and `High` set to the value of `MinContrast`.

4.3 Pose Estimation Using Surface-Based 3D Matching

For the 3D pose estimation with surface-based matching, a surface model is generated from a 3D object model that was obtained either from a 3D computer aided design (CAD) model or from a 3D reconstruction approach, e.g., stereo vision ([chapter 5](#) on page 105) or sheet of light ([chapter 6](#) on page 139). The surface model consists of a set of 3D points and the points' normal vectors. That is, the corresponding information must be provided (at least implicitly) by the 3D object model. In contrast to the shape-based 3D matching, the instances of the object are not located in images but in a 3D scene, i.e., in a set of 3D points that is provided as another 3D object model and which can be obtained by a 3D reconstruction approach, too.

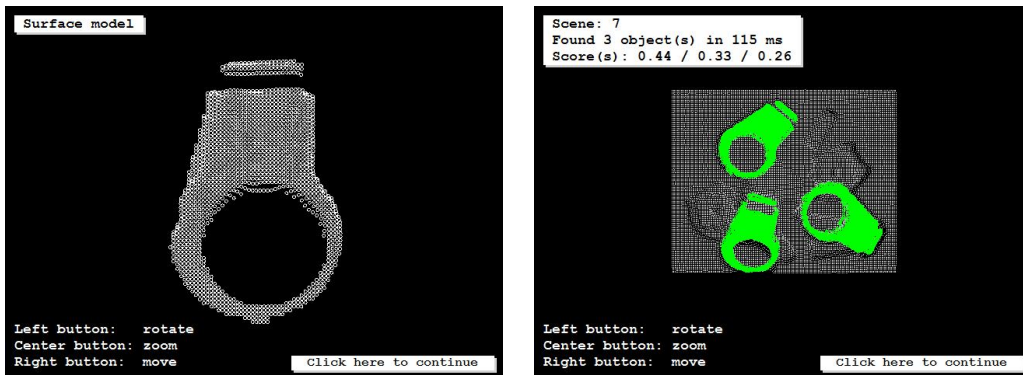


Figure 4.10: Surface-based matching: (left) 3D model of engine cover, (right) model instances found in a 3D scene.

In the following, the general proceeding for surface-based 3D matching is introduced (see [section 4.3.1](#)).

4.3.1 General Proceeding for Surface-Based 3D Matching

Surface-based 3D matching consists of the following basic steps:

- access the 3D object model needed for the creation of the surface model,
- create the surface model from it,

- access the 3D object model that represents the search data,
- use the surface model to search the object in the search data, and
- destroy the handles for the matching result, all 3D object models, and the surface model.

An example for the surface-based 3D matching of engine covers (see [figure 4.10](#) on page 93) is the HDevelop example program `hdevelop\3D-Matching\Surface-Based\find_surface_model.hdev`.

Step 1: Access the 3D object model needed for the creation of the surface model

In contrast to shape-based 3D matching, for surface-based matching the 3D object model does not need to be available as CAD model but can also be derived by a 3D reconstruction. If a 3D object model is available as a CAD model or if the 3D object model was saved to file after an offline 3D reconstruction, it can be accessed using `read_object_model_3d`. Suitable file formats are, e.g., DXF, OFF, PLY, or OM3. The format OM3 is a HALCON-specific format for 3D object models that can be derived from the results of HALCON's 3D reconstruction approaches (see [figure 1.3](#) on page 14 in [chapter 1](#)). Tips on how to obtain a suitable model, including the specification of the requirements the CAD models must fulfill, are provided with the description of the operator `read_object_model_3d` in the Reference Manual.

In the example program, the 3D object model is not read from file but is derived from X, Y, and Z images that were obtained by a specific 3D sensor: a “time-of-flight” (TOF) camera. To get the 3D object model of a single engine cover, a region containing a single engine cover is extracted by a blob analysis from the Z image (see [figure 4.11](#)) and the Z image is reduced to the corresponding ROI.

```
read_image (Image, ImagePath + 'engine_cover_xyz_01')
decompose3 (Image, Xm, Ym, Zm)
threshold (Zm, ModelZ, 0, 650)
connection (ModelZ, ConnectedModel)
select_obj (ConnectedModel, ModelROI, [10,9])
union1 (ModelROI, ModelROI)
reduce_domain (Xm, ModelROI, Xm)
```

The X, Y, and Z images are now transformed into a 3D object model using the operator `xyz_to_object_model_3d`. This 3D object model is needed to create the surface model that will serve as model for the matching.

```
xyz_to_object_model_3d (Xm, Ym, Zm, ObjectModel3DModel)
```

Note that for surface-based matching information about the coordinates of the 3D points and their normals is needed. Thus, if a 3D object model is obtained from a CAD model or from multi-view stereo, the normals or alternatively a triangular or polygon mesh must be contained in the 3D object model. If a 3D object model is obtained from X, Y, and Z images, like in the example program, it contains the coordinates of the 3D points and a 2D mapping. From this 2D mapping, `create_surface_model`, i.e., the operator that is used to create the surface model in the next step, can automatically derive the needed information.

Step 2: Create the surface model

The operator `create_surface_model` creates a surface model by sampling the 3D object model with a certain distance. The sampling distance can be adjusted with the parameter `RelSamplingDistance`.




Figure 4.11: A single engine cover extracted from the Z image.

Note that a smaller value leads to a slower but more robust matching, whereas a larger value speeds up the matching but at the same time decreases the robustness.

```
create_surface_model (ObjectModel3DModel, 0.03, [], [], SFM)
```

Step 3: Access the 3D object model that represents the search data

Similar to the 3D object model that was needed for the creation of the surface model, the 3D object models in which the object of interest are searched must be accessed, i.e., they are read from file using [read_object_model_3d](#) or they are derived online by a 3D reconstruction. In the example program, the 3D object models are derived again from X, Y, and Z images using the operator [xyz_to_object_model_3d](#). Note that a blob analysis is applied to remove the background plane from the search data.

```
NumImages := 10
for Index := 2 to NumImages by 1
  read_image (Image, ImagePath + 'engine_cover_xyz_' + Index$'02')
  decompose3 (Image, X, Y, Z)
  threshold (Z, SceneGood, 0, 666)
  reduce_domain (X, SceneGood, XReduced)
  xyz_to_object_model_3d (XReduced, Y, Z, ObjectModel3DSceneReduced)
```

Step 4: Find the surface model in the search data

With the surface model that was created by [create_surface_model](#) or read from file by [read_surface_model](#), the object can be searched for in the search data using the operator [find_surface_model](#).

```

find_surface_model (SFM, ObjectModel3DSceneReduced, 0.05, 0.2, 0.15, \
                    'true', ['num_matches'], [10], Pose, Score, \
                    SurfaceMatchingResultID)

```

Several parameters can be set to control the search process. For detailed information, we recommend to read the description of the operator in the Reference Manual. For each model instance that could be located by the matching, the operator returns the pose and a score that describes the quality of the match.

After the matching, the result can be visualized. Here, for each match that exceeds a certain score, the 3D object model is transformed with the corresponding pose and the transformed 3D object models of a specific search scene are stored in the tuple `ObjectModel3DResult`.

```

ObjectModel3DResult := []
for Index2 := 0 to |Score|-1 by 1
  if (Score[Index2]<0.11)
    continue
  endif
  CPose := Pose[Index2*7:Index2*7+6]
  pose_to_hom_mat3d (CPose, CHomMat3D)
  affine_trans_object_model_3d (ObjectModel3DModel, CHomMat3D, \
                                ObjectModel3DAffineTrans)
  ObjectModel3DResult := [ObjectModel3DResult, \
                          ObjectModel3DAffineTrans]
endfor

```

The visualization is then applied using the procedure `disp_object_model_3d`. Actually, the transformed models are displayed together with the original 3D scene (see [figure 4.10](#) on page 93), i.e., with the 3D object model that is obtained from the X, Y, and Z images from which also the search data was obtained, but this time no blob analysis is applied to remove the background plane.

```

xyz_to_object_model_3d (X, Y, Z, ObjectModel3DScene)
NumResult := |ObjectModel3DResult|
tuple_gen_const (NumResult, 'green', Colors)
tuple_gen_const (NumResult, 'circle', Shapes)
tuple_gen_const (NumResult, 3, Radii)
disp_object_model_3d (EmptyObject, [ObjectModel3DScene, \
                                   ObjectModel3DResult], CamParam, [], 'false', \
                    ['gray',Colors], [], ['circle',Shapes], [0, \
                    Radii], [], [], Message, [], Instructions, \
                    PoseOut)

```

Step 5: Destroy the handles for the matching result, all 3D object models, and the surface model

The 3D object models that were accessed or obtained during a search or visualization process are destroyed from memory before searching the model in another 3D scene. The same applies for the handle of the matching result that was returned by the matching. The 3D object models are destroyed with `clear_object_model_3d` and the handle of the result is destroyed with `clear_surface_matching_result`.


```

clear_object_model_3d (ObjectModel3DSceneReduced)
clear_object_model_3d (ObjectModel3DScene)
clear_object_model_3d (ObjectModel3DResult)
clear_surface_matching_result (SurfaceMatchingResultID)
endfor

```

When the 3D object model of the model and the surface model that was created from it are not needed anymore, they are destroyed, too. The surface model is destroyed with `clear_surface_model`.

```

clear_object_model_3d (ObjectModel3DModel)
clear_surface_model (SFM)

```

Besides the basic steps, it is often required to inspect the 3D object model, to re-use the surface model, or to visualize the result of the matching. These steps are described in the Solution Guide I, [chapter 10](#) on page [127](#).

4.4 Pose Estimation Using 3D Primitives Fitting

3D primitives fitting determines amongst others the positions or 3D poses of simple 3D shapes, so-called “3D primitives”, that are fitted into segmented parts of a 3D scene. The 3D scene is a set of 3D points that is available as 3D object model. It can be obtained, e.g., by stereo vision ([chapter 5](#) on page [105](#)) or sheet of light ([section 6.1](#) on page [139](#)). The available types of 3D primitives comprise a sphere, a cylinder, and a plane. When fitting 3D primitives into the segmented 3D data, the results comprise a radius and position for a sphere, a radius and 3D pose for a cylinder, and a 3D pose for a plane.

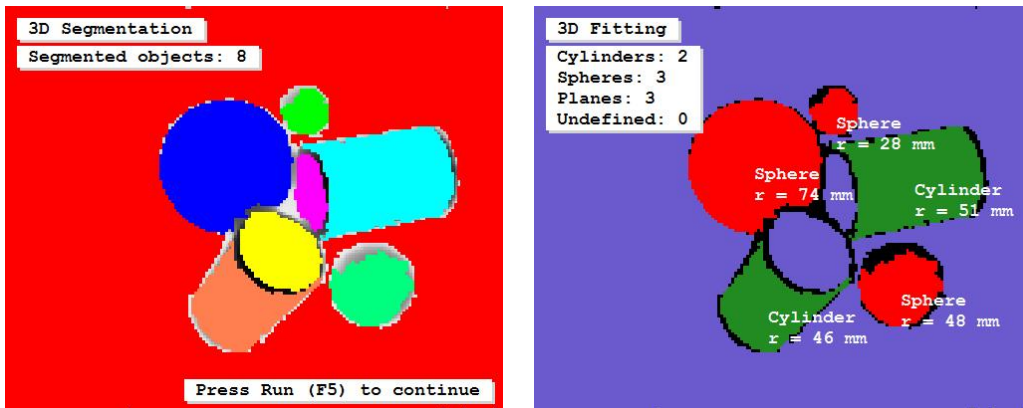


Figure 4.12: 3D primitives fitting: (left) segmentation of 3D object model, (right) result of the fitting.

The fitting typically consists of two steps. The first step is the **segmentation** of the 3D scene into sub-sets of neighbored 3D points that may correspond to selected types of 3D primitives (see [figure 4.12](#), left). The segmentation can be applied by different means, depending on the input data. In any case, the resulting sub-sets of 3D points must be available as 3D object models that contain, at least implicitly,

the coordinates of the 3D points and their meshing. The second step is the actual **fitting** (see figure 4.12, right). There, the operator `fit_primitives_object_model_3d` tries to find the best fitting 3D primitive for an individual 3D object model. The result is another 3D object model from which the parameters of the successfully fitted primitive can be queried.

An example for a 3D primitives fitting that uses a simple 2D segmentation to derive the 3D object model of a single cylinder from a 3D scene that is provided by X, Y, and Z images is the HDevelop example program `hdevelop\3D-Tools\3D-Segmentation\fit_primitives_object_model_3d.hdev`.



Figure 4.13: Height data used for 3D primitives fitting: (from left to right) X, Y, and Z image.

The individual channels of the image are accessed with `access_channel` (see figure 4.13). Then, a threshold is applied to the Z image to separate the cylinder from the background (see figure 4.14). The corresponding ROI is created with `reduce_domain`. From the reduced image channels, a 3D object model is created using `xyz_to_object_model_3d`.

```
read_image (XYZ, '3d_machine_vision/segmentation/3d_primitives_xyz_02.tif')
access_channel (XYZ, X, 1)
access_channel (XYZ, Y, 2)
access_channel (XYZ, Z, 3)
threshold (Z, Region, 0.0, 0.83)
reduce_domain (X, Region, XTmp)
xyz_to_object_model_3d (XTmp, Y, Z, ObjectModel3DID)
```

A 3D primitive of the 'primitive_type' 'cylinder' is fitted into the 3D object model using the operator `fit_primitives_object_model_3d`.

```
ParFitting := ['primitive_type', 'fitting_algorithm', 'min_radius', \
               'max_radius', 'output_xyz_mapping']
ValFitting := [ 'cylinder', 'least_squares_huber', 0.01, \
                0.1, 'true' ]
fit_primitives_object_model_3d (ObjectModel3DID, ParFitting, ValFitting, \
                               ObjectModel3DOutID)
```

The result of the fitting is a handle for a 3D object model from which information like the primitive's parameters can be queried with `get_object_model_3d_params`.

Often, a simple 2D segmentation is not suitable, e.g., because the 3D data is not derived from Y, X, and Z images or because the 3D scene consists of several 3D structures that touch or



Figure 4.14: 3D primitives fitting: (left) simple threshold for 2D segmentation, (right) result returned by the fitting.

overlap. Then, a 3D segmentation using the operator `segment_object_model_3d` has to be applied as is shown in the HDevelop example program `hdevelop\3D-Tools\3D-Segmentation\segment_object_model_3d.hdev`. There, the individual objects can not be separated by a simple 2D segmentation, because they overlap (see [figure 4.15](#)). The 3D data is accessed again from X, Y, and Z images, but this time, the 3D object model is created without a preceding 2D segmentation.

```
read_image (XYZ, '3d_machine_vision/segmentation/3d_primitives_xyz_01.tif')
access_channel (XYZ, X, 1)
access_channel (XYZ, Y, 2)
access_channel (XYZ, Z, 3)
xyz_to_object_model_3d (X, Y, Z, ObjectModel3DID)
```




Figure 4.15: Height data used for 3D primitives fitting: (from left to right) X, Y, and Z image.

Instead, the derived 3D object model is prepared for a 3D segmentation by calling the operator `prepare_object_model_3d`. Note that, if you do not explicitly prepare the 3D object model, the operator is called internally during the segmentation with `segment_object_model_3d`, which may slow down the application if the same 3D object model is used several times.

```
prepare_object_model_3d (ObjectModel3DID, 'segmentation', 'false', \
                        'max_area_holes', 100)
```

The operator `segment_object_model_3d` segments the 3D object model into different sub-sets of 3D points that have similar characteristics like the same orientation of the normals or the same curvature of the underlying surface. By default, the operator not only segments the 3D scene but already tries to fit primitives of the selected types into the segments (see [figure 4.12](#) on page 97). Thus, no separate calls to `fit_primitives_object_model_3d` are needed. Here, 'primitive_type' is set to 'all'. That is, for each sub-set of 3D points the best fitting type of primitive is returned (see [figure 4.12](#) on page 97).

```
ParSegmentation := ['max_orientation_diff', 'max_curvature_diff', \
                   'output_xyz_mapping', 'min_area']
ValSegmentation := [ 0.13, 0.11, \
                    'true', 150 ]
ParFitting := ['primitive_type', 'fitting_algorithm', 'min_radius', \
              'max_radius']
ValFitting := [ 'all', 'least_squares_huber', 0.01, \
               0.1 ]
segment_object_model_3d (ObjectModel3DID, [ParSegmentation, ParFitting], \
                       [ValSegmentation, ValFitting], ObjectModel3DOutID)
```

The result of the combined segmentation and fitting is a tuple of handles for the 3D object models that represent the different sub-sets of 3D points. From each of these 3D object models, information like the success of the fitting, the primitives' types, and the primitives' parameters can be queried with `get_object_model_3d_params`. Querying the primitive's parameters, a tuple is returned. The size and content of this tuple depends on the type of the primitive. In particular, for a cylinder the tuple contains seven values (three for the position, three for the orientation, and one for the radius), for a sphere it contains four values (three for the position and one for the radius), and for a plane it contains again four values (three for the unit normal vector and one for the orthogonal distance of the plane from the origin of the coordinate system). How to access a single parameter from such a tuple is shown exemplarily for the radius of a cylinder.

```
for Index := 0 to |ObjectModel3DOutID|-1 by 1
  get_object_model_3d_params (ObjectModel3DOutID[Index], \
                             'has_primitive_data', ParamValue)
  if (ParamValue='true')
    get_object_model_3d_params (ObjectModel3DOutID[Index], \
                               'primitive_parameter', GenParamValuesP)
    get_object_model_3d_params (ObjectModel3DOutID[Index], \
                               'primitive_type', ParamValue)
    if (ParamValue='cylinder')
      RadiusCylinder := GenParamValuesP[6]
    endif
  endif
endfor
```

When the different 3D object models are not needed anymore, they are destroyed.

```
clear_object_model_3d (ObjectModel3DID)
for Index := 0 to |ObjectModel3DOutID|-1 by 1
    clear_object_model_3d (ObjectModel3DOutID[Index])
endfor
```

4.5 Pose Estimation Using Calibrated Perspective Deformable Matching

The perspective deformable matching finds and locates objects that are similar to a template model in an image. This matching approach uses the contours of the object in the images and is independent of the perspective view on the object, i.e., perspective deformations are considered when searching the model in unknown images. The perspective deformable matching can be applied either for a calibrated camera, then the 3D pose of the object is returned, or for an uncalibrated camera, then only the 2D projective transformation matrix (homography) is returned.

The perspective deformable matching is suitable for all planar objects or planar object parts that are clearly distinguishable by their contours. Compared to the shape-based 3D matching (see [section 4.2](#) on page 81 or the Solution Guide I, [chapter 10](#) on page 127), there is no need to pregenerate different views of an object. Thus, it is significantly faster. Hence, if you search for planar perspective deformed objects, we recommend the perspective deformable matching.

The Solution Guide II-B, [section 3.6](#) on page 124 shows in detail how to apply the approach. [Figure 4.16](#) shows the poses of engine parts obtained by the HDevelop example `hdevelop\Applications\Position-Recognition-3D\locate_engine_parts.hdev`.

4.6 Pose Estimation Using Calibrated Descriptor-Based Matching

Similar to the perspective deformable matching, the descriptor-based matching finds and locates objects that are similar to a template model in an image. Again, the matching can be applied either for a calibrated camera, then the 3D pose of an object is returned, or for an uncalibrated camera, then only the 2D projective transformation matrix (homography) is returned.

The essential difference between the descriptor-based and the perspective deformable matching is that the descriptor-based matching is not based on contours but uses distinctive object points, so-called interest points, to describe the template and to find the model in the image.

Note that the calibrated descriptor-based matching is suitable mainly to determine the 3D pose of planar objects with characteristic texture and distinctive object points. For low-textured objects with rounded edges you should select one of the other pose estimation approaches. Further, the descriptor-based matching is less accurate than the perspective deformable matching. But on the other hand, it is significantly faster if a large search space, e.g., caused by a large scale range, is used.

The Solution Guide II-B, [section 3.7](#) on page 136 shows in detail how to apply the approach. [Figure 4.17](#) on page 103 shows the 3D pose of a cookie box obtained by the HDevelop example `hdevelop\Applications\Object-Recognition-2D\locate_cookie_box.hdev`. A comparison between the




Figure 4.16: 3D poses of engine parts obtained by a calibrated perspective deformable matching.

calibrated descriptor-based matching and pose estimation methods that use manually extracted correspondences (see [section 4.1](#) on page 76) is provided by the HDevelop example `hdevelop\Matching\Descriptor-Based\pose_from_point_correspondences.hdev`.

4.7 Pose Estimation for Circles

HALCON offers an alternative approach to estimate the pose of 3D circles, which can be applied with less effort than the previously described approaches. It is based on the known geometrical behavior of perspectively distorted circles. In particular, 3D circles are represented as ellipses in the image. Using the extracted 2D ellipse of a 3D circle together with the internal camera parameters and the known radius of the circle, the two possible 3D poses of the circle (having the same position but opposite orientations) can be obtained easily using the operator `get_circle_pose`. The HDevelop examples `hdevelop\Transformations\Poses\get_circle_pose.hdev` and `hdevelop\Applications\Position-Recognition-3D\3d_position_of_circles.hdev` show in detail how to apply the approach.

4.8 Pose Estimation for Rectangles

Additionally to the pose estimation for 3D circles, also the poses of 3D rectangles can be estimated with an approach that can be applied with less effort than the general approach. It is based on the



Figure 4.17: 3D pose of a cookie box label obtained by a calibrated descriptor-based matching.

known geometrical behavior of perspective distorted rectangles. In particular, a contour is segmented into four line segments and their intersections are considered as the corners of a quadrangular contour. Using the extracted 2D quadrangle of the 3D rectangle together with the internal camera parameters and the known size of the rectangle, the four (or eight in case of a square) possible 3D poses of the rectangle can be obtained easily using the operator `get_rectangle_pose`. The HDevelop examples `hdevelop\Applications\Position-Recognition-3D\get_rectangle_pose_barcode.hdev` and `hdevelop\Applications\Position-Recognition-3D\3d_position_of_rectangle.hdev` show in detail how to apply the approach.

Chapter 5

3D Vision With a Stereo System

With a stereo system, i.e., with two or more cameras, you can derive 3D information of the surface of arbitrarily shaped objects. Possible results are distance images, 3D coordinates, or 3D surfaces.

Typical applications of stereo vision comprise, but are not limited to, completeness checks, inspection of ball grid arrays, etc. Reconstructing surfaces can also serve as a preprocessing step for surface-based 3D matching (see [section 4.3](#) on page 93) or 3D primitives fitting (see [section 4.4](#) on page 97).

[Figure 5.1](#) shows a surface reconstructed from four stereo images. [Figure 5.2](#) shows an image of a binocular stereo camera system, the resulting stereo image pair, and the height map that has been derived from the reconstructed distance image.




Figure 5.1: Left: images of a 4-camera system; right: reconstructed surface (3D object model).

HALCON actually provides two stereo methods:

- **Binocular stereo** ([section 5.3](#) on page 113) uses exactly two cameras. The result is a disparity image, a distance image, or 3D coordinates. The latter are returned either for selected points or for the complete view.




Figure 5.2: Top: stereo camera system; center: stereo image pair; bottom: height map.

- **Multi-view stereo** (section 5.4 on page 130) can use more than two cameras. Thus, it is able to image the whole 3D object and not just the surface of a specific view. It returns results in form of 3D surfaces (given as 3D object models, see page 12) or 3D coordinates of selected points.

Before describing the two methods, we first take a look at some general topics. In particular, we

- introduce you to the principle of stereo vision (section 5.1) and
- show how to calibrate a stereo system (section 5.2 on page 110).

5.1 The Principle of Stereo Vision

Assume the simplified configuration of two parallel looking 1D cameras with identical internal parameters as shown in figure 5.3. Furthermore, the basis, i.e., the straight line connecting the two optical centers of the two cameras, is assumed to coincide with the x-axis of the first camera.




Figure 5.3: Vertical section of a binocular stereo camera system.

Then, the image plane coordinates of the projections of the point $P(x^c, z^c)$ into the two images can be expressed by

$$u_1 = f \frac{x^c}{z^c} \quad (5.1)$$

$$u_2 = f \frac{x^c - b}{z^c} \quad (5.2)$$

where f is the focal length and b the length of the basis.

The pair of image points that results from the projection of one object point into the two images is often referred to as *conjugate points* or *homologous points*.

The difference between the two image locations of the *conjugate points* is called the *disparity* d :

$$d = (u_2 - u_1) = -\frac{f \cdot b}{z^c} \quad (5.3)$$

Given the camera parameters and the image coordinates of two conjugate points, the z^c coordinate of the corresponding object point P , i.e., its distance from the stereo camera system, can be computed by

$$z^c = -\frac{f \cdot b}{d} \quad (5.4)$$

Note that the internal camera parameters of both cameras and the relative pose of the second camera in relation to the first camera are necessary to determine the distance of P from the stereo camera system.

Thus, the tasks to be solved for stereo vision are

1. to determine the camera parameters and
2. to determine conjugate points.

The first task is achieved by the calibration of the stereo camera system, which is described in [section 5.2](#). This calibration is quite similar to the calibration of a single camera, described in [section 3.2](#) on page 40, in fact, it even uses the same operators.

The second task is the so-called stereo matching process, which in HALCON is just a call of the operator `binocular_disparity` (or `binocular_distance`, respectively) for the correlation-based stereo and a call of the operator `binocular_disparity_mg` (or `binocular_distance_mg`, respectively) for multigrid stereo. These operators are described in [section 5.3.5](#), together with the operators doing all the necessary calculations to obtain world coordinates from the stereo images.

The multi-view surface reconstruction described in [section 5.4.2.1](#) on page 134 extends the basic stereo vision principle to more than one image pair. There, the matching step is “hidden” in the reconstruction operators, which use the operator `binocular_disparity` internally to compute disparity images of the individual stereo image pairs.

5.1.1 The Setup of a Stereo Camera System

The basic stereo camera system consists of two cameras looking at the same object from different positions (see [figure 5.4](#)). Multi-view stereo systems have additional cameras, but the principle stays the same.




Figure 5.4: Stereo camera system (r and y axes point towards the reader).

It is very important to ensure that neither the internal camera parameters (e.g., the focal length) nor the relative pose between the cameras changes during the calibration process or between the calibration process and the ensuing application of the calibrated stereo camera system, because the calibration remains valid only as long as the cameras preserve their relative pose. Therefore, it is advisable to mount the cameras on a *stable* platform.

5.1.2 Resolution of a Stereo Camera System

The manner in which the cameras are placed influences the accuracy of the results that is achievable with the stereo camera system.

The distance resolution Δz , i.e., the accuracy with which the distance z of the object surface from the stereo camera system can be determined, can be expressed by

$$\Delta z = \frac{z^2}{f \cdot b} \cdot \Delta d \quad (5.5)$$

To achieve a high distance resolution, the setup should be chosen such that the length b of the basis as well as the focal length f are large, and that the stereo camera system is placed as close as possible to the object. In addition, the distance resolution depends directly on the accuracy Δd with which the disparities can be determined. Typically, the disparities can be determined with an accuracy of $1/5$ up to $1/10$ pixel, which corresponds to approximately $1 \mu\text{m}$ for a camera with $7.4 \mu\text{m}$ pixel size.

In [figure 5.5](#), the distance resolutions that are achievable in the ideal case are plotted as a function of the distance for four different configurations of focal lengths and base lines, assuming Δd to be $1 \mu\text{m}$.

Note that if the ratio between b and z is very large, problems during the stereo matching process may occur, because the two images of the stereo pair differ too much. The maximum reasonable ratio b/z depends on the surface characteristics of the object. In general, objects with little height differences can be imaged with a higher ratio b/z , whereas objects with larger height differences should be imaged with a smaller ratio b/z .

If this is difficult to ensure in your application when using two cameras, you should consider using a multi-view stereo system, i.e., more than two cameras.

5.2 Calibrating the Stereo Camera System

As mentioned above, the calibration of the stereo camera system is very similar to the calibration of a single camera ([section 3.2](#) on page 40). The major difference is that it is not sufficient to view the calibration plate from a single camera, but it must be completely visible in at least some of the overlapping parts of the images that are taken by each pair of neighbored cameras (see [section 5.2.2](#) on page 112 for details).

In this section only a brief description of the calibration process is given. More details can be found in [section 3.2](#) on page 40. Here, the stereo-specific parts of the calibration process are described in depth. In particular, it is shown

- how to create and configure the calibration data model for multiple cameras ([section 5.2.1](#)),
- how to acquire suitable calibration images ([section 5.2.2](#) on page 112),
- how to add the observation data of multiple cameras to the calibration data model ([section 5.2.3](#) on page 113), and




Figure 5.5: Distance resolution plotted over the distance ($\Delta d = 1 \mu m$).

- how to perform the calibration (section 5.2.4 on page 113).

The code fragments used for the following descriptions belong to the HDevelop example program `hdevelop\Calibration\Multi-View\calibrate_cameras.hdev`.

5.2.1 Creating and Configuring the Calibration Data Model

As described in section 3.2.1 on page 41, the first step of preparing for a calibration is to create the calibration data model with the operator `create_calib_data`. In the example, four cameras are used with one calibration object.

```
NumCameras := 4
NumCalibObjects := 1
create_calib_data ('calibration_object', NumCameras, NumCalibObjects, \
                  CalibDataID)
```

Then, the initial camera parameters are set with the operator `set_calib_data_cam_param` (see also section 3.2.2 on page 41). The parameter `CameraIdx` is set to 'all', so that the values are set for all cameras.

```
StartCamPar := [0.0085, 0.0, 0.0, 0.0, 0.0, 0.0, 6e-6, 6e-6, Width*.5, \
                Height*.5, Width, Height]
set_calib_data_cam_param (CalibDataID, 'all', CameraType, StartCamPar)
```

Finally, the calibration object is described with the operator `set_calib_data_calib_object` (see [section 3.2.3](#) on page 45).

```
CaltabDescr := 'caltab_100mm.descr'
set_calib_data_calib_object (CalibDataID, 0, CaltabDescr)
```

5.2.2 Acquiring Calibration Images

For the calibration of the stereo camera system, each camera acquires multiple images of one or more calibration objects in different poses. Note that it is not necessary that the calibration object is always visible in all poses for each camera. The only requirement is that the cameras can be “connected” in a chain by the calibration object poses, e.g., that camera 0 and 1 observe the pose 0 of the calibration object, camera 1 and 2 observe pose 3, and camera 2 and 3 observe pose 8 (see [figure 5.6](#)).




Figure 5.6: Cameras are connected in a chain.

Note that when taking the images for the calibration, the rules for taking the calibration images for the single camera calibration (see [section 3.2.4.1](#) on page 49) apply accordingly.

Note that you must not change the camera setup between the acquisition of the calibration images and the acquisition of the stereo images of the object to be investigated. How to obtain suitable stereo image pairs for the object to investigate is shown in [section 5.3.3](#) on page 115.

5.2.3 Observing the Calibration Object

As in the case of the single camera calibration, the main input data for the camera calibration are the observed points of the calibration objects in the camera images (see [section 3.2.4](#) on page 49). You add an observation to the calibration data model with the operator `set_calib_data_observ_points`. In the example, this operator is called within a double loop over all poses and all cameras. The code raises a warning if the calibration plate could not be found in an image.

```
for PoseIndex := 0 to NumPoses-1 by 1
  for CameraIndex := 0 to NumCameras-1 by 1
    read_image (Image, FileName)
    IgnoreImg := 0
    try
      find_caltab (Image, Caltab, CaltabDescr, 3, 82, 5)
      find_marks_and_pose (Image, Caltab, CaltabDescr, StartCamPar, \
                          128, 10, 18, 0.9, 15, 100, RCoord, CCoord, \
                          StartPose)
```

5.2.4 Calibrating the Cameras

The actual calibration of the stereo camera system is carried out with the operator `calibrate_cameras` (see [section 3.2.6](#) on page 52).

```
calibrate_cameras (CalibDataID, Error)
```

The calibration stores its results in the calibration data model. How to access them is described in separate sections for binocular stereo ([section 5.3.2](#) on page 115) and multi-view stereo ([section 5.4.1.1](#) on page 131) because the methods use the results in different forms.

5.3 Binocular Stereo Vision

With the operators for binocular stereo vision, you can compute disparity and distance images and 3D coordinates using two cameras. In fact, HALCON provides two binocular stereo methods: correlation-based and multigrid stereo (see [section 5.3.1](#) for the differences).

The following sections show how to

- access the results of the calibration ([section 5.3.2](#) on page 115),
- rectify the stereo images ([section 5.3.4](#) on page 116), and
- reconstruct 3D information ([section 5.3.5](#) on page 120).

If not stated otherwise, the example programs used in this section are

- `stereo_calibration.dev`,
- `height_above_reference_plane_from_stereo.dev`, and
- `3d_information_for_selected_points.dev`.

They can be found in the directory `solution_guide\3d_vision`.

As an alternative to fully calibrated stereo, HALCON offers the so-called uncalibrated stereo vision. Here, the relation between the cameras is determined from the scene itself, i.e., without needing a special calibration object. Please refer to [section 5.3.6](#) on page 130 for more information about this method.

5.3.1 Correlation-Based Versus Multigrid Stereo

Two approaches for stereo matching are available, the traditionally used correlation-based stereo matching and the multigrid stereo matching. The correlation-based stereo matching uses correlation techniques to find corresponding points and thus to determine the disparities or distances for the observed image points. The disparities or distances are calculated with the operators `binocular_disparity` or `binocular_distance`, respectively. The correlation-based stereo is characterized by the following advantages and disadvantages:

Most important advantages of correlation-based stereo:

- fast,
- can be automatically parallelized on multi-core or multi-processor hardware, and
- is invariant against gray-value changes.

Most important disadvantage of correlation-based stereo:

- works good only for significantly textured areas. Areas without enough texture cannot be reconstructed.

The multigrid stereo matching uses a variational approach based on multigrid methods. This approach returns disparity and distance values also for image parts that contain no texture (as long as these parts are surrounded by significant structures between which an interpolation of values is possible). The disparities or distances are calculated with the operators `binocular_disparity_mg` or `binocular_distance_mg`, respectively. The multigrid stereo is characterized by the following advantages and disadvantages:

Most important advantages of multigrid stereo:

- interpolates 3D information for areas without texture based on the surrounding areas,
- in particular for edges, the accuracy is higher as for correlation-based stereo, and
- the resolution is higher than for correlation-based stereo, i.e., smaller objects can be reconstructed.

Most important disadvantages of multigrid stereo:

- only partially invariant against gray value changes and
- cannot be parallelized automatically.

5.3.2 Accessing the Calibration Results

As described in [section 3.2.7](#) on page 53, you access the results of the calibration with the operator `get_calib_data`. In addition to the internal camera parameters, we now use the operator also to get the relative pose between the two cameras.

```
get_calib_data (CalibDataID, 'camera', 0, 'params', CamParamL)
get_calib_data (CalibDataID, 'camera', 1, 'params', CamParamR)
get_calib_data (CalibDataID, 'camera', 1, 'pose', cLPcR)
```

If you want to perform the calibration in an offline step, you can save the camera setup model with `write_camera_setup_model` and then destroy the calibration data model with `clear_calib_data`.

```
write_camera_setup_model (CameraSetupModelID, 'stereo_camera_setup.csm')
clear_calib_data (CalibDataID)
```

5.3.3 Acquiring Stereo Images

The following rules help you to acquire suitable stereo image pairs. [Section 5.3.3.1](#) and [section 5.3.3.2](#) provide you with background information to understand the rules.

- Do not change the camera setup between the acquisition of the calibration images and the acquisition of the stereo images of the object to be investigated. How to obtain suitable images for the calibration of a stereo camera system is shown in [section 5.2.2](#) on page 112.
- Ensure a proper illumination of the object, e.g., avoid reflections.
- If the object shows no texture, consider to project texture onto it or use multigrid stereo (see [section 5.3.1](#) on page 114).
- Place the object such that repetitive patterns are not aligned with the rows of the rectified images.

5.3.3.1 Image Texture

The 3D coordinates of each object point are derived by intersecting the lines of sight of the respective conjugate image points. The conjugate points are determined by an automatic matching process. This matching process has some properties that should be accounted for during the image acquisition.

For each point of the first image, the conjugate point in the second image must be determined. This point matching relies on the availability of texture. The conjugate points cannot be determined correctly in areas without sufficient texture ([figure 5.7](#)).

This applies in particular for the correlation-based (binocular) stereo (which is currently the base for multi-view stereo). The multigrid (binocular) stereo, however, can interpolate values in areas without texture. But note that even then, a certain amount of texture must be available to enable the interpolation. In [section 5.3.1](#) on page 114 the differences between both stereo matching approaches are described in more detail.




Figure 5.7: Rectified stereo images and matching result of the correlation-based stereo in a poorly textured area (regions where the matching process failed are displayed white).

5.3.3.2 Repetitive Patterns

If the images contain repetitive patterns, the matching process may be confused, since in this case many points look alike. In order to make the matching process fast and reliable, the stereo images are rectified such that pairs of conjugate points always have identical row coordinates in the rectified images, i.e., that the search space in the second rectified image is reduced to a line. With this, repetitive patterns can disturb the matching process only if they are parallel to the rows of the rectified images (figure 5.8).

5.3.4 Rectifying the Stereo Images

With the internal camera parameters and the relative pose, the stereo images can be *rectified*, so that conjugate points lie on the same row in both rectified images.

Note that it is assumed that the parameters of the *first* image of a pair stem from the *left* image and the parameters of the *second* image stem from the *right* image, whereas the notations 'left' and 'right' refer




Figure 5.8: Rectified stereo images with repetitive patterns aligned to the image rows and cutout of the matching result (regions where the matching process failed are displayed white).

to the line of sight of the two cameras (see [figure 5.4](#) on page 109). If the images are used in the reverse order, they will appear upside down after the rectification.

In [figure 5.9](#) the original images of a stereo pair are shown, where the two cameras are rotated heavily with respect to each other. The corresponding rectified images are displayed in [figure 5.10](#).




Figure 5.9: Original stereo images.

The rectification itself is carried out using the operators [gen_binocular_rectification_map](#) and [map_image](#). The operator [gen_binocular_rectification_map](#) requires the internal camera parameters of both cameras and the relative pose of the second camera in relation to the first camera.




Figure 5.10: Rectified stereo images.

```
gen_binocular_rectification_map (MapL, MapR, CamParamL, CamParamR, cLPcR, 1, \
    'geometric', 'bilinear', RectCamParL, \
    RectCamParR, CamPoseRectL, CamPoseRectR, \
    RectLPosRectR)
```

The parameter `SubSampling` can be used to change the size and resolution of the rectified images with respect to the original images. A value of `1` indicates that the rectified images will have the same size as the original images. Larger values lead to smaller images with a resolution reduced by the given factor, smaller values lead to larger images.

Reducing the image size has the effect that the following stereo matching process runs faster, but also that less details are visible in the result. In general, it is proposed not to use values below `0.5` or above `2`. Otherwise, smoothing or aliasing effects occur, which may disturb the matching process.

The rectification process can be described as projecting the original images onto a common rectified image plane. The method to define this plane can be selected by the parameter `Method`. So far, only the method `'geometric'` can be selected, in which the orientation of the common rectified image plane is defined by the cross product of the base line and the line of intersection of the two original image planes.

The rectified images can be thought of as being acquired by a virtual stereo camera system, called rectified stereo camera system, as displayed in [figure 5.11](#). The optical centers of the rectified cameras are the same as for the real cameras, but the rectified cameras are rotated such that they are looking parallel and that their `x`-axes are collinear. In addition, both rectified cameras have the same focal length. Therefore, the two image planes coincide. Note that the principal point of the rectified images, which is the origin of the image plane coordinate system, may lie outside the image.

The parameter `Interpolation` specifies whether bilinear interpolation (`'bilinear'`) should be applied between the pixels of the input images or whether the gray value of the nearest pixel (`'none'`) should be used. Bilinear interpolation yields smoother rectified images, whereas the use of the nearest neighbor is faster.

The operator returns the rectification maps and the camera parameters of the virtual, rectified cameras.




Figure 5.11: Rectified stereo camera system.

Finally, the operator `map_image` can be applied to both stereo images using the respective rectification map generated by the operator `gen_binocular_rectification_map`.

```
map_image (ImageL, MapL, ImageRectifiedL)
map_image (ImageR, MapR, ImageRectifiedR)
```

If the calibration was erroneous, the rectification will produce wrong results. This can be checked very easily by comparing the row coordinates of conjugate points selected from the two rectified images. If the row coordinates of conjugate points are different within the two rectified images, they are not correctly rectified. In this case, you should check the calibration process carefully.

An incorrectly rectified image pair may look like the one displayed in [figure 5.12](#).




Figure 5.12: Incorrectly rectified stereo images.

5.3.5 Reconstructing 3D Information

There are many possibilities to derive 3D information from rectified stereo images.

Non-metrical information: If only non-metrical information about the surface of an object is needed, it may be sufficient to determine the disparities within the overlapping area of the stereo image pair by using the operator `binocular_disparity` for correlation-based stereo (section 5.3.5.1) or `binocular_disparity_mg` for multigrad stereo (section 5.3.5.2 on page 123). The differences between both stereo matching approaches are described in more detail in section 5.3.1 on page 114.

Distance of the object surface: If metrical information is required, the operator `binocular_distance` (section 5.3.5.3 on page 124) or `binocular_distance_mg` (section 5.3.5.4 on page 126), respectively, can be used to extract the distance of the object surface from the stereo camera system.

3D coordinate images: Having a disparity image of a rectified binocular stereo system, you can additionally derive the corresponding x , y , and z coordinates using `disparity_image_to_xyz` (see section 5.3.5.6 on page 127).

3D coordinates for selected points: To derive metrical information for selected points only, the operators `disparity_to_distance` or `disparity_to_point_3d` can be used. The first of these two operators calculates the distance z of points from the stereo camera system based on their disparity (see section 5.3.5.5 on page 127). The second operator calculates the x , y , and z coordinates from the row and column position of a point in the first rectified image and from its disparity (see section 5.3.5.6 on page 127).

Alternatively, the operator `intersect_lines_of_sight` can be used to calculate the x , y , and z coordinates of selected points (see section 5.3.5.7 on page 128). Then, there is no need to determine the disparities in advance. Only the image coordinates of the conjugate points and the camera parameters are needed. This operator can also handle image coordinates of the original stereo images. Thus, the rectification can be omitted. In exchange, you must determine the conjugate points by yourself.

Note that all operators that deal with disparities or distances require all input to be based on the rectified images. This applies to the image coordinates as well as to the camera parameters.

5.3.5.1 Determining Disparities Using Correlation-Based Stereo

Disparities are an indicator for the distance z of object points from the stereo camera system, since points with equal disparities also have equal distances z (equation 5.4 on page 108).

Therefore, if it is only necessary to know whether there are locally high objects, it is sufficient to derive the disparities. For correlation-based stereo, this is done by using the operator `binocular_disparity`.

```
binocular_disparity (ImageRectifiedL, ImageRectifiedR, DisparityImage, \
                    ScoreImageDisparity, 'ncc', MaskWidth, MaskHeight, \
                    TextureThresh, MinDisparity, MaxDisparity, NumLevels, \
                    ScoreThresh, 'left_right_check', 'interpolation')
```

The operator requires the two *rectified* images as input. The disparities are derived only for those conjugate points that lie within the respective image domain in both images. With this, it is possible to speed up the calculation of the disparities if the image domain of at least one of the two rectified images is reduced to a region of interest, e.g., by using the operator `reduce_domain`.

Several parameters can be used to control the behavior of the matching process that is performed by the operator `binocular_disparity` to determine the conjugate points:

With the parameter `Method`, the matching function is selected. The methods `'sad'` (summed absolute differences) and `'ssd'` (summed squared differences) compare the gray values of the pixels within a matching window directly, whereas the method `'ncc'` (normalized cross correlation) compensates for the mean gray value and its variance within the matching window. Therefore, if the two images differ in brightness and contrast, the method `'ncc'` should be preferred. However, since the internal computations are less complex for the methods `'sad'` and `'ssd'`, they are faster than the method `'ncc'`.

The width and height of the matching window can be set independently with the parameters `MaskWidth` and `MaskHeight`. The values should be odd numbers. Otherwise they will be increased by one. A larger matching window will lead to a smoother disparity image, but may result in the loss of small details. In contrast, the results of a smaller matching window tend to be noisy but they show more spatial details.

Because the matching process relies on the availability of texture, low-textured areas can be excluded from the matching process. The parameter `TextureThresh` defines the minimum allowed variance within the matching window. For areas where the texture is too low, no disparities will be determined.

The parameters `MinDisparity` and `MaxDisparity` define the minimum and maximum disparity values. They are used to restrict the search space for the matching process. If the specified disparity range does not contain the actual range of the disparities, the conjugate points cannot be found correctly. Therefore, the disparities will be incomplete and erroneous. On the other hand, if the disparity range is specified too large, the matching process will be slower and the probability of mismatches increases.

Therefore, it is important to set the parameters `MinDisparity` and `MaxDisparity` carefully. There are several possibilities to determine the appropriate values:

- If you know the minimum and maximum distance of the object from the stereo camera system (section 5.3.5.3 on page 124), you can use the operator `distance_to_disparity` to determine the respective disparity values.
- You can also determine these values directly from the rectified images. For this, you should display the two rectified images and measure the approximate column coordinates of the point N , which is nearest to the stereo camera system (N_{col}^{image1} and N_{col}^{image2}) and of the point F , which is the farthest away (F_{col}^{image1} and F_{col}^{image2}), each in both rectified images. Now, the values for the definition of the disparity range can be calculated as follows:

$$MinDisparity = N_{col}^{image2} - N_{col}^{image1} \quad (5.6)$$

$$MaxDisparity = F_{col}^{image2} - F_{col}^{image1} \quad (5.7)$$

The operator `binocular_disparity` uses image pyramids to improve the matching speed. The disparity range specified by the parameters `MinDisparity` and `MaxDisparity` is only used on the uppermost pyramid level, indicated by the parameter `NumLevels`. Based on the matching results on that level, the disparity range for the matching on the next lower pyramid levels is adapted automatically.

The benefits with respect to the execution time are greatest if the objects have different regions between which the appropriate disparity range varies strongly. However, take care that the value for `NumLevels` is not set too large, as otherwise the matching process may fail because of lack of texture on the uppermost pyramid level.

The parameter `ScoreThresh` specifies which matching scores are acceptable. Points for which the matching score is not acceptable are excluded from the results, i.e., the resulting disparity image has a reduced domain that comprises only the accepted points.

Note that the value for `ScoreThresh` must be set according to the matching function selected via `Method`. The two methods `'sad'` ($0 \leq \text{score} \leq 255$) and `'ssd'` ($0 \leq \text{score} \leq 65025$) return lower matching scores for better matches. In contrast, the method `'ncc'` ($-1 \leq \text{score} \leq 1$) returns higher values for better matches, where a score of zero indicates that the two matching windows are totally different and a score of minus one denotes that the second matching window is exactly inverse to the first matching window.

The parameter `Filter` can be used to activate a downstream filter by which the reliability of the resulting disparities is increased. Currently, it is possible to select the method `'left_right_check'`, which verifies the matching results based on a second matching in the reverse direction. Only if both matching results correspond to each other, the resulting conjugate points are accepted. In some cases, this may lead to gaps in the disparity image, even in well textured areas, as this verification is very strict. If you do not want to verify the matching results based on the `'left_right_check'`, set the parameter `Filter` to `'none'`.

The subpixel refinement of the disparities is switched on by setting the parameter `SubDisparity` to `'interpolation'`. It is switched off by setting the parameter to `'none'`.

The results of the operator `binocular_disparity` are the two images `Disparity` and `Score`, which contain the disparities and the matching score, respectively. In figure 5.13, a rectified stereo image pair is displayed, from which the disparity and score images that are displayed in figure 5.14 were derived.

Both resulting images refer to the image geometry of the first rectified image, i.e., the disparity for the point (r,c) of the first rectified image is the gray value at the position (r,c) of the disparity image. The disparity image can, e.g., be used to extract the components of the board, which would be more difficult in the original images, i.e., without the use of 3D information.




Figure 5.13: Rectified stereo images.




Figure 5.14: Disparity image (left) and score image (right).

In [figure 5.14](#), areas where the matching did not succeed, i.e., undefined regions of the images, are displayed white in the disparity image and black in the score image.

5.3.5.2 Determining Disparities Using Multigrid Stereo

For multigrid stereo, the disparities can be derived with `binocular_disparity_mg`. Similar to the correlation-based approach, the two rectified images are used as input and the disparity image as well as a score image are returned. But as the disparities are obtained by a different algorithm, the parameters that control the behavior of the multigrid stereo matching process are completely different. In particular, the multigrid-specific control parameters are `GrayConstancy`, `GradientConstancy`, `Smoothness`, `InitialGuess`, `CalculateScore`, `MGParamName`, and `MGParamValue`. They are explained in detail in the Reference Manual entry for `binocular_disparity_mg`. The significant differences between correlation-based and multigrid stereo are listed in [section 5.3.1](#) on page 114.

How to determine the disparities of the components on a PCB using `binocular_disparity_mg`

with different levels of accuracy is shown in the HDevelop example program `hdevelop\3D-Reconstruction\Binocular-Stereo\binocular_disparity_mg.hdev`.

There, first the two stereo images are rectified.

```
read_image (ImageL, 'stereo/board/board_l_01')
read_image (ImageR, 'stereo/board/board_r_01')
CamParamL := [0.0130507774353, -665.817817207, 1.4803417027e-5, 1.48e-5, \
              155.89225769, 126.70664978, 320, 240]
CamParamR := [0.0131776504517, -731.860636733, 1.47997569293e-5, 1.48e-5, \
              162.98210144, 119.301040649, 320, 240]
RelPose := [0.153573364258, -0.00373362231255, 0.0447351264954, \
            0.174289124775, 319.843388114, 359.894955219, 0]
gen_binocular_rectification_map (MapL, MapR, CamParamL, CamParamR, RelPose, \
                                1, 'geometric', 'bilinear', RectCamParL, \
                                RectCamParR, CamPoseRectL, CamPoseRectR, \
                                RectLPosRectR)
map_image (ImageL, MapL, ImageRectifiedL)
map_image (ImageR, MapR, ImageRectifiedR)
```

Then, `binocular_disparity_mg` is called four times, each time with a different parameter for the parameter `MGParamValue`, which sets the accuracy that should be achieved. Note that an increasing accuracy also leads to an increasing runtime.

```
DefaultParameters := ['fast', 'fast_accurate', 'accurate', 'very_accurate']
for I := 0 to |DefaultParameters|-1 by 1
    Parameters := DefaultParameters[I]
    binocular_disparity_mg (ImageRectifiedL, ImageRectifiedR, Disparity, \
                          Score, 1, 10, 5, 0, 'false', \
                          'default_parameters', Parameters)
endfor
```

5.3.5.3 Determining Distances Using Correlation-Based Stereo

The distance of an object point from the stereo camera system is defined as its distance from the x - y -plane of the coordinate system of the first rectified camera. For correlation-based stereo, it can be determined by the operator `binocular_distance`, which is used analogously to the operator `binocular_disparity` described in [section 5.3.5.1](#) on page 121.

```
binocular_distance (ImageRectifiedL, ImageRectifiedR, DistanceImage, \
                  ScoreImageDistance, RectCamParL, RectCamParR, \
                  RectLPosRectR, 'ncc', MaskWidth, MaskHeight, \
                  TextureThresh, MinDisparity, MaxDisparity, NumLevels, \
                  ScoreThresh, 'left_right_check', 'interpolation')
```

The three additional parameters, namely the camera parameters of the rectified cameras as well as the relative pose of the second rectified camera in relation to the first rectified camera can be taken directly from the output of the operator `gen_binocular_rectification_map`.

Figure 5.15 shows the distance image and the respective score image for the rectified stereo pair of figure 5.13 on page 123. Because the distance is calculated directly from the disparities and from the camera parameters, the distance image looks similar to the disparity image (figure 5.14). What is more, the score images are identical, since the underlying matching process is identical.




Figure 5.15: Distance image (left) and score image (right).

It can be seen from figure 5.15 that the distance of the board changes continuously from left to right. The reason is that, in general, the x - y -plane of the coordinate system of the first rectified camera will be tilted with respect to the object surface (see figure 5.16).

If it is necessary that one reference plane of the object surface has a constant distance value of, e.g., zero, the tilt can be compensated easily: First, at least three points that lie on the reference plane must be defined. These points are used to determine the orientation of the (tilted) reference plane in the distance image. Therefore, they should be selected such that they enclose the region of interest in the distance image. Then, a distance image of the (tilted) reference plane can be simulated and subtracted from the distance image of the object. Finally, the distance values themselves must be adapted by scaling them with the cosine of the angle between the tilted and the corrected reference plane.

These calculations are carried out in the procedure `tilt_correction`, which is part of the example program `solution_guide\3d_vision\height_above_reference_plane_from_stereo.hdev` (appendix B.4 on page 234).

```
procedure tilt_correction (DistanceImage, RegionDefiningReferencePlane,
                          DistanceImageCorrected):::
```

In principle, this procedure can also be used to correct the disparity image, but note that you must not use the corrected disparity values as input to any operators that derive metric information.

If the reference plane is the ground plane of the object, an inversion of the distance image generates an image that encodes the heights above the ground plane. Such an image is displayed on the left hand side in figure 5.17.

Objects of different height above or below the ground plane can be segmented easily using a simple `threshold` with the minimal and maximal values given directly in units of the world coordinate system,




Figure 5.16: Distances of the object surface from the x - y -plane of the coordinate system of the first rectified camera.

e.g., meters. The image on the right hand side of [figure 5.17](#) shows the results of such a segmentation, which can be carried out based on the corrected distance image or the image of the heights above the ground plane.

5.3.5.4 Determining Distances Using Multigrid Stereo

For multigrid stereo, the distance of an object point from the stereo camera system can be determined by the operator `binocular_distance_mg`, which is used analogously to the operator `binocular_disparity_mg` described in [section 5.3.5.2](#) on page 123, but with the three additional parameters described also for `binocular_distance`. These are the camera parameters of the rectified cameras as well as the relative pose of the second rectified camera in relation to the first rectified camera, which can be taken directly from the output of the operator `gen_binocular_rectification_map`.

How to determine the depths of the components on a PCB with high accuracy is shown in the HDevelop example program `hdevelop\3D-Reconstruction\Binocular-Stereo\binocular_distance_mg.hdev`.




Figure 5.17: Left: Height above the reference plane; Right: Segmentation of high objects (white: 0-0.4 mm, light gray: 0.4-1.5 mm, dark gray: 1.5-2.5 mm, black: 2.5-5 mm).

```
binocular_distance_mg (ImageRectifiedL, ImageRectifiedR, Distance, Score, \
    RectCamParL, RectCamParR, RectLPosRectR, 1, 10, 5, \
    0, 'false', 'default_parameters', 'accurate')
```

5.3.5.5 Determining Distances for Selected Points from the Disparity Image

If only the distances of selected points should be determined, the operator `disparity_to_distance` can be used. It simply transforms given disparity values into the respective distance values. For example, if you want to know the minimum and maximum distance of the object from the stereo camera system you can determine the minimum and maximum disparity values from the disparity image and transform them into distances.

```
min_max_gray (CaltabL, Disparity, 0, MinDisparity, MaxDisparity, Range)
disparity_to_distance (RectCamParL, RectCamParR, RectLPosRectR, \
    [MinDisparity,MaxDisparity], MinMaxDistance)
```

This transformation is constant for the entire rectified image, i.e., all points having the same disparity have the same distance from the x - y -plane of the coordinate system of the first rectified camera. Therefore, besides the camera parameters of the rectified cameras, only the disparity values need to be given.

5.3.5.6 Determining 3D Coordinates from the Disparity Image

If the x , y , and z coordinates of points have to be calculated, different operators are available that derive the corresponding information from the disparity image:

`disparity_to_point_3d` computes the 3D coordinates for specified image coordinates and returns them in three tuples.

```
disparity_to_point_3d (RectCamParL, RectCamParR, RectLPosRectR, RL, CL, \
                      DisparityOfSelectedPoints, X_CCS_FromDisparity, \
                      Y_CCS_FromDisparity, Z_CCS_FromDisparity)
```

[disparity_image_to_xyz](#) computes 3D coordinates for the complete image and returns them in three images (see the HDevelop example program `hdevelop\3D-Reconstruction\Binocular-Stereo\disparity_image_to_xyz.hdev`).

```
disparity_image_to_xyz (DisparityImage, X, Y, Z, RectCamParL, RectCamParR, \
                      RectLPosRectR)
```

The operators require the camera parameters of the two rectified cameras and the relative pose of the cameras.

The x , y , and z coordinates are returned in the coordinate system of the first rectified camera.

5.3.5.7 Determining 3D Coordinates for Selected Points from Point Correspondences

If only the 3D coordinates of selected points should be determined, you can alternatively use the operator [intersect_lines_of_sight](#) to determine the x , y , and z coordinates of points from the image coordinates of the respective conjugate points. Note that you must determine the image coordinates of the conjugate points yourself.

```
intersect_lines_of_sight (RectCamParL, RectCamParR, RectLPosRectR, RL, CL, \
                         RR, CR, X_CCS_FromIntersect, Y_CCS_FromIntersect, \
                         Z_CCS_FromIntersect, Dist)
```

The x , y , and z coordinates are returned in the coordinate system of the first (rectified) camera.

The operator can also handle image coordinates of the original stereo images. Thus, the rectification can be omitted. In this case, the camera parameters of the original stereo cameras have to be given instead of the parameters of the rectified cameras.

It is possible to transform the x , y , and z coordinates determined by the latter two operators from the coordinate system of the first (rectified) camera into a given world coordinate system (WCS), e.g., a coordinate system with respect to the building plan of a factory building. For this, a homogeneous transformation matrix, which describes the transformation between the two coordinate systems, is needed.

This homogeneous transformation matrix can be determined in various ways. The easiest way is to take an image of a HALCON calibration plate with the first camera only. If the 3D coordinates refer to the *rectified* camera coordinate system, the image must be rectified as well. Then, the pose of the calibration plate in relation to the first (rectified) camera can be determined using the operators [find_caltab](#), [find_marks_and_pose](#), and [vector_to_pose](#).


```

find_caltab (ImageRectifiedL, CaltabL, 'caltab_30mm.descr', SizeGauss, \
            MarkThresh, MinDiamMarks)
find_marks_and_pose (ImageRectifiedL, CaltabL, 'caltab_30mm.descr', \
                    RectCamParL, StartThresh, DeltaThresh, MinThresh, \
                    Alpha, MinContLength, MaxDiamMarks, RCoordL, CCoordL, \
                    StartPoseL)
vector_to_pose (X, Y, Z, RCoordL, CCoordL, RectCamParL, 'iterative', \
               'error', PoseOfCalibrationPlate, Quality)

```

The resulting pose can be converted into a homogeneous transformation matrix.

```
pose_to_hom_mat3d (PoseOfCalibrationPlate, HomMat3d_WCS_to_RectCCS)
```

If necessary, the transformation matrix can be modified with the operators [hom_mat3d_rotate_local](#), [hom_mat3d_translate_local](#), and [hom_mat3d_scale_local](#).

```

hom_mat3d_translate_local (HomMat3d_WCS_to_RectCCS, 0.01125, -0.01125, 0, \
                          HomMat3DTranslate)
hom_mat3d_rotate_local (HomMat3DTranslate, rad(180), 'y', \
                       HomMat3d_WCS_to_RectCCS)

```

The homogeneous transformation matrix must be inverted in order to represent the transformation from the (rectified) camera coordinate system into the WCS.

```
hom_mat3d_invert (HomMat3d_WCS_to_RectCCS, HomMat3d_RectCCS_to_WCS)
```

Finally, the 3D coordinates can be transformed using the operator [affine_trans_point_3d](#).

```

affine_trans_point_3d (HomMat3d_RectCCS_to_WCS, X_CCS_FromIntersect, \
                     Y_CCS_FromIntersect, Z_CCS_FromIntersect, X_WCS, \
                     Y_WCS, Z_WCS)

```

The homogeneous transformation matrix can also be determined from three specific points. If the origin of the WCS, a point on its x -axis, and a third point that lies in the x - y -plane, e.g., directly on the y -axis, are given, the transformation matrix can be determined using the procedure `gen_hom_mat3d_from_three_points`, which is part of the HDevelop example program `solution_guide\3d_vision\3d_information_for_selected_points.hdev`.

```

procedure gen_hom_mat3d_from_three_points (Origin, PointOnXAxis,
                                          PointInXYPlane, HomMat3d):::

```

The resulting homogeneous transformation matrix can be used as input for the operator [affine_trans_point_3d](#), as shown above.

5.3.6 Uncalibrated Stereo Vision

Similar to uncalibrated mosaicking (see [chapter 10](#) on page 199), HALCON also provides an “uncalibrated” version of stereo vision, which derives information about the cameras from the scene itself by matching characteristic points. The main advantage of this method is that you need no special calibration object. The main disadvantage of this method is that, without a precisely known calibration object, the accuracy of the result is highly dependent on the content of the scene, i.e., the accuracy of the result degrades if the scene does not contain enough 3D information or if the extracted characteristic points in the two images do not precisely correspond to the same world points, e.g., due to occlusion.

In fact, HALCON provides two versions of uncalibrated stereo: **Without any knowledge about the cameras and about the scene**, you can rectify the stereo images and perform a segmentation similar to the method described in [section 5.3.5.3](#) on page 124. For this, you first use the operator `match_fundamental_matrix_ransac`, which determines the so-called *fundamental matrix*. This matrix models the cameras and their relation. But in contrast to the model described in [section 5.1](#) on page 107, internal and external parameters are not available separately. Thus, no metric information can be derived.

The fundamental matrix is then passed on to the operator `gen_binocular_proj_rectification`, which is the “uncalibrated” version of the operator `gen_binocular_rectification_map`. With the output of this operator, i.e., the rectification maps, you can then proceed to rectify the images as described in [section 5.3.4](#) on page 116. Because the relative pose of the cameras is not known, you cannot generate the distance image and segment it as described in [section 5.3.5.3](#) on page 124. The HDevelop example program `hdevelop\Applications\Object-Recognition-2D\board_segmentation_uncalib.hdev` shows an alternative approach that can be used if the reference plane appears dominant in the images, i.e., if many correspondences are found on it.

Because no calibration object is needed, this method can be used to perform stereo vision with a single camera. Note, however, that the method assumes that there are no lens distortions in the images. Therefore, the accuracy of the results degrades if the lens has significant distortions.

If the internal parameters of the camera are known, you can determine the relative pose between the cameras using the operator `match_rel_pose_ransac` and then use all the stereo methods described for the fully calibrated case. There is, however, a limitation: The determined pose is relative in a second sense, because it can be determined only up to a scale factor. The reason for this is that without any knowledge about the scene, the algorithm cannot know whether the points in the scene are further away or whether the cameras are further apart because the effect in the image is the same in both cases. If you have additional information about the scene, you can solve this ambiguity and determine the “real” relative pose. This method is shown in the HDevelop example program `hdevelop\Applications\3D-Reconstruction\Binocular-Stereo\uncalib_stereo_boxes.hdev`.

5.4 Multi-View Stereo Vision

In comparison to binocular stereo, multi-view stereo allows

- to use more than two cameras and thus to reconstruct 3D information from all around an object and

- to reconstruct surfaces and 3D coordinates of selected points in form of 3D object models.

Internally, it is based on binocular stereo. However, by default it does not return the disparity image as a result.

The following sections show

- how to initialize the stereo model ([section 5.4.1](#)) and
- how to reconstruct 3D information ([section 5.4.2](#) on page 133).

5.4.1 Initializing the Stereo Model

The operators for multi-view stereo use a so-called *stereo model* to encapsulate all needed data. The following sections show

- how to access the calibration results ([section 5.4.1.1](#)),
- how to specify the world coordinate system ([section 5.4.1.2](#)), and
- how to create the stereo model ([section 5.4.1.3](#)).

5.4.1.1 Accessing the Calibration Results

In contrast to binocular stereo (see [section 3.2.7](#) on page 53), for multi-view stereo you access the results of the calibration not separately but in form of a so-called *camera setup model*, which contains the internal camera parameters as well as the relative poses between the cameras. To derive the camera setup model from the calibration data model, which must have been calibrated before with [calibrate_cameras](#) as is described in [section 5.2](#) on page 110, you call the operator [get_calib_data](#) as follows:

```
get_calib_data (CalibDataID, 'model', 'general', 'camera_setup_model', \
               CameraSetupModelID)
```

If you want to perform the calibration in an offline step, you can save the camera setup model with the operator [write_camera_setup_model](#) and then destroy the calibration data model with [clear_calib_data](#) and the camera setup model with [clear_camera_setup_model](#).

```
write_camera_setup_model (CameraSetupModelID, 'four_camera_setup_model.csm')
clear_calib_data (CalibDataID)
clear_camera_setup_model (CameraSetupModelID)
```

5.4.1.2 Specifying the Coordinate System of the Camera Setup

The coordinate system of the stereo camera setup is identical to the coordinate system of the so-called *reference camera* of the setup, which is typically the camera with the index 0 (see the upper part of [figure 5.18](#) on page 133). The poses of the other cameras and the reconstructed coordinates are computed relative to this camera.

You can change the setup's coordinate system with the operator `set_camera_setup_param` in two ways. In particular, you can

- **select another camera as reference camera** by setting `CameraIdx` to 'general' and `ParamName` to 'reference_camera' and passing the index of the camera in `ParamValue` or
- **specify the pose of the desired setup coordinate system** (relative to the current one) by setting `CameraIdx` to 'general' and `ParamName` to 'coord_transf_pose' and passing the pose in `ParamValue`.

The latter case is shown at the bottom of [figure 5.18](#). There, the desired coordinate system is marked by the calibration plate (typically, you would add a rotation to let the z axis point upwards).

How to change the pose of the setup's coordinate system is shown, e.g., in the HDevelop example program `hdevelop\Calibration\Multi-View\calibrate_cameras.hdev`. There, it is moved from the reference camera to the calibration plate with pose 0. For that, pose 0 of the calibration plate relative to the reference camera is accessed with `get_calib_data` using the parameter 'calib_obj_pose' and, to consider the thickness of the calibration plate, the z coordinate of the pose is modified with `set_origin_pose`. Then, the setup's coordinate system is moved into this pose with `set_camera_setup_param`.

```
RefPoseIndex := 0
get_calib_data (CalibDataID, 'calib_obj_pose', [0, RefPoseIndex], 'pose', \
              PoseCam0Idx0)
set_origin_pose (PoseCam0Idx0, 0, 0, CaltabThickness, ReferencePose)
set_camera_setup_param (CameraSetupModelID, 'general', 'coord_transf_pose', \
                      ReferencePose)
```

5.4.1.3 Creating the Stereo Model

After adapting the camera setup model to your requirements, you pass it to the operator `create_stereo_model`, which creates the stereo model. Note that at this point you must already specify whether you want to reconstruct points or surfaces!

To reconstruct surfaces, use the parameter 'surface_pairwise'.

```
create_stereo_model (CameraSetupModelID, 'surface_pairwise', [], [], \
                  StereoModelID)
```

To reconstruct points, call the operator with the parameter 'points_3d'.

```
create_stereo_model (CameraSetupModelID, 'points_3d', [], [], StereoModelID)
```




Figure 5.18: Coordinate systems of a multi-view camera setup: Top: default setup coordinate system located in the reference camera 0; bottom: setup coordinate system moved to the coordinate system of the calibration plate.

5.4.2 Reconstructing 3D Information

With multi-view stereo, you can reconstruct

- the surface of an object ([section 5.4.2.1](#)) or
- 3D coordinates for selected points ([section 5.4.2.2](#) on page 135).

5.4.2.1 Reconstructing Surfaces

The main functionality of multi-view stereo is the reconstruction of surfaces. The following code fragments belong to the HDevelop example program `hdevelop\Applications\Robot-Vision\locate_pipe_joints_stereo.hdev`, which reconstructs pipe joints using four cameras (and afterwards performs surface-based 3D matching to estimate the pose of the individual pipe joints, see Solution Guide I, [section 10.3.2](#) on page 137). [Figure 5.1](#) on page 105 shows the four camera images and the reconstructed surface.

After creating and initializing the stereo model, the model can be configured for the stereo reconstruction using `set_stereo_model_param`. With this operator, several parameters are adjusted that control, e.g., the subsampling that is used for the reconstruction. Additionally, as the multi-view surface reconstruction is based on computing binocular disparity images (see [section 5.3.5.1](#) on page 121), several parameters are used to configure the binocular image rectification and the internally called binocular stereo operators.

```
* -> Subsampling X, Y, Z
set_stereo_model_param (StereoModelID, 'sub_sampling_step', 3)
* -> Interpolation aliasing by binocular image rectification
set_stereo_model_param (StereoModelID, 'rectif_interpolation', 'bilinear')
set_stereo_model_param (StereoModelID, 'rectif_sub_sampling', 1.2)
* -> Binocular disparity parameters
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')
```

Furthermore, the surface reconstruction must be restricted to a specific part of the 3D space, which is realized by the definition of a bounding box. This box is built by the coordinates of its front lower left corner and its back upper right corner. In the program, a camera setup model is used for which the coordinate system was moved to the object, or more precisely, to a calibration plate that was used for the calibration of the scene in which the object was placed. Relative to this calibration plate, the coordinates are specified in meters.

```
set_stereo_model_param (StereoModelID, 'bounding_box', [-0.2,-0.07,-0.075, \
    0.2,0.07, -0.004])
```

Before calling the reconstruction operator, you must call the operator `set_stereo_model_image_pairs` to specify which cameras form pairs, i.e., between which camera images the disparity images are to be computed. In the example, the cameras 0 and 1 and the cameras 2 and 3, respectively, form pairs.

```
set_stereo_model_image_pairs (StereoModelID, [0,2], [1,3])
```

Then, the actual surface reconstruction is applied with the operator `reconstruct_surface_stereo`.

```
reconstruct_surface_stereo (Images, StereoModelID, PipeJointPileOM3DID)
```

If the reconstruction fails, please refer to the Reference Manual entry of `reconstruct_surface_stereo`, which contains detailed information about troubleshooting.

The reconstructed surface is returned as a 3D object model (see [page 12](#)), which by default consists of points and their normals. If you need a surface that contains meshing information, e.g., because you want to apply a 3D primitives fitting (see [section 4.4](#) on [page 97](#)) to the 3D object model, you have to additionally set the parameter 'point_meshing' within `set_stereo_model_param` before building the camera pairs for the reconstruction.

In the example, the 3D object model is visualized by the procedure `disp_object_model_3d`, which allows to interactively rotate, move, and zoom into the model.

```
PoseIn := [0.0, 0.0, 0.5, -30, 0, 180, 0]
if (Index = 1)
    disp_object_model_3d (EmptyObject, PipeJointPileOM3DID, CamParam0, \
        PoseIn, 'false', 'yellow', 'points', \
        'circle', 0, [], 1, \
        'Reconstructed scene in '+ReconsTime$.3+' s', \
        [], Instructions, PoseOut)
endif
```

5.4.2.2 Reconstructing 3D Points

Multi-view stereo also allows to reconstruct the 3D coordinates of selected points. The main advantages of using the multi-view approach in comparison to the binocular variant (see, [section 5.3.5.7](#) on [page 128](#)) are that the reconstruction is more accurate when more than two lines of sight can be taken into account and that points located on different sides of an object can be reconstructed.

After creating and initializing the stereo model, you can directly use multi-view stereo to reconstruct the 3D coordinates from point correspondences. This is shown in the HDevelop example program `hdevelop\3D-Reconstruction\Multi-View\reconstruct_points_stereo.hdev`, where four cameras are used to reconstruct the coordinates of the calibration marks of a calibration plate in three different poses (see [figure 5.19](#)).

The main input for the reconstruction operator `reconstruct_points_stereo` are the corresponding points from the multi-view images. They must be passed as tuples in the parameters

- `Row` (row coordinate of the point),
- `Column` (column coordinate of the point),
- `CameraIdx` (index of the camera), and
- `PointIdx` (index of the point).

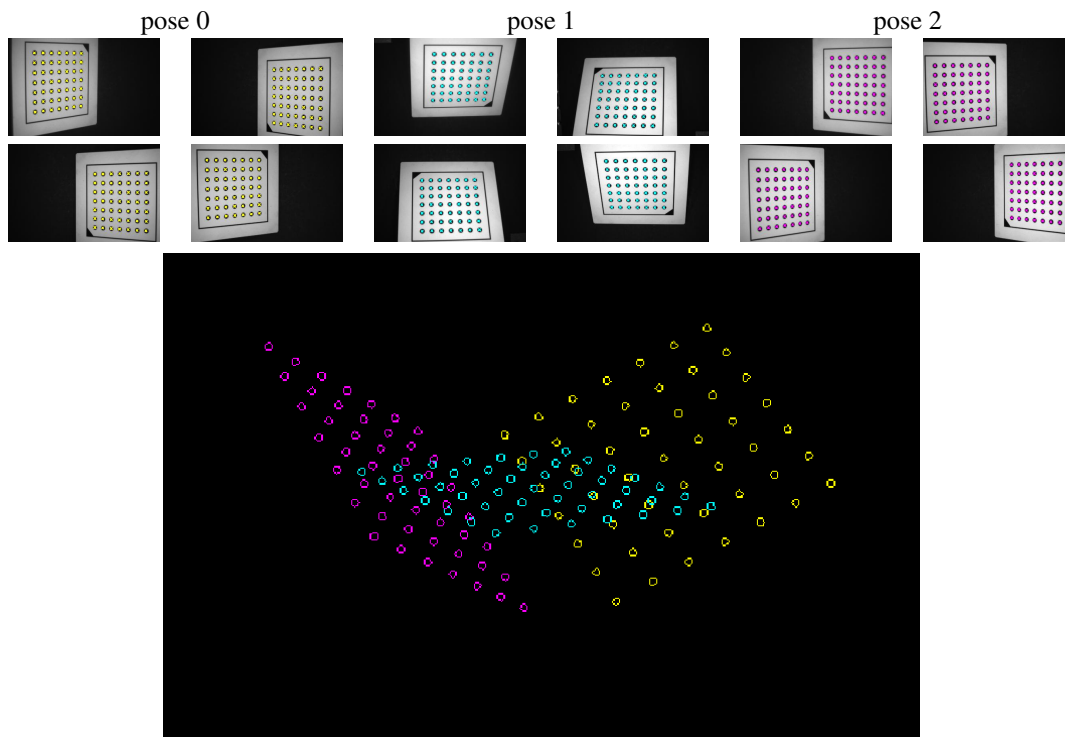


Figure 5.19: Top: images of four cameras of the calibration plate with extracted marks in three different poses; Bottom: reconstructed points.

Generally, you must extract the corresponding points by yourself. In the example, this task is easy, because here the points correspond to the marks of a calibration plate, which can be easily extracted with the operator `find_marks_and_pose`. For other objects than the HALCON calibration plate, the extraction may be a bit more challenging.

In the example, the extraction is realized as follows: First, `caltab_points` is used to derive the number of calibration marks from the description of the used calibration plate. As `find_marks_and_pose` always extracts the calibration marks in the same order, the tuple with the point indices of the correspondence information for a single image can be created using `tuple_gen_sequence`. As the parameter `CameraIdx` must be a tuple of the same length as `Row`, `Column`, and `PointIdx`, a tuple with the same length for which each element is '1' is created, which is used later to assign the correct correspondence values to the camera indices.

```
caltab_points (CaltabDescr, X, Y, Z)
tuple_gen_sequence (0, |X|-1, 1, Indices)
Ones := gen_tuple_const (|X|, 1)
```

Now, for each pose of the calibration plate, empty tuples for the correspondence information are created and filled with the values obtained for each camera that images the calibration plate under this pose. In particular, for each camera, the calibration marks are extracted with `find_caltab` and

`find_marks_and_pose`, and the resulting row and column coordinates are added to the tuples `AllRow` and `AllColumn`. The tuple with the point indices is added to the tuple `AllIndices` and the corresponding elements with the specific camera index are added to the tuple `AllCams`.

```

Objects3D := []
for PoseIndex := 0 to 2 by 1
  AllRow := []
  AllColumn := []
  AllIndices := []
  AllCams := []
  for CameraIndex := 0 to 3 by 1
    get_camera_setup_param (CameraSetupModelID, CameraIndex, 'params', \
                           CameraParam)
    ImageFile := ImgPath + 'multi_view_calib_cam_' + CameraIndex + '_' \
                 + (13+PoseIndex)$'02'
    read_image (Image, ImageFile)
    find_caltab (Image, Caltab, CaltabDescr, 3, 82, 5)
    find_marks_and_pose (Image, Caltab, CaltabDescr, CameraParam, 128, \
                        10, 18, 0.9, 15, 100, Row, Column, Pose)
    AllRow := [AllRow, Row]
    AllColumn := [AllColumn, Column]
    AllIndices := [AllIndices, Indices]
    AllCams := [AllCams, CameraIndex*Ones]
  endfor

```

After accumulating the correspondence information for all cameras that image the calibration plate under the specific pose, the reconstruction is applied with `reconstruct_points_stereo`.

```

reconstruct_points_stereo (StereoModelID, AllRow, AllColumn, [], \
                          AllCams, AllIndices, X, Y, Z, CovWP, \
                          PointIndexOut)

```

It returns tuples with the x , y , and z coordinates and with the index of those points that could be reconstructed, i.e., which were extracted in two or more images. In the example, the coordinates are transformed via x , y , and z images into a 3D object model and the models of all three calibration plate poses are interactively visualized with the procedure `disp_object_model_3d`, which allows to rotate, move, and zoom into the model.

```
gen_image_const (ImageX, 'real', 1, |X|)
gen_image_const (ImageY, 'real', 1, |X|)
gen_image_const (ImageZ, 'real', 1, |X|)
...
    set_grayval (ImageX, Indices, 0*Indices, X)
    set_grayval (ImageY, Indices, 0*Indices, Y)
    set_grayval (ImageZ, Indices, 0*Indices, Z)
    xyz_to_object_model_3d (ImageX, ImageY, ImageZ, ObjectModel3DID)
    Objects3D := [Objects3D, ObjectModel3DID]
    disp_continue_message (WindowHandles[3], 'black', 'true')
    stop ()
endfor
disp_object_model_3d (BackgroundImage, Objects3D, CameraParam, [], 'false', \
                    Colors, 'points', 'circle', 5, [], [], [], [], \
                    Instructions, PoseOut)
```

Please note that the example shows only the main functionality of the stereo point reconstruction. For more information, e.g., about input and output covariances or the influence of the bounding box of the stereo model, please refer to the Reference Manual entry of [reconstruct_points_stereo](#).

Chapter 6

Laser Triangulation with Sheet of Light

Laser triangulation can be used to reconstruct the surface of a 3D object by approximating it via a set of height profiles. HALCON provides operators for a special type of laser triangulation that is called sheet-of-light technique.

6.1 The Principle of Sheet of Light

The basic idea of the sheet-of-light technique is to project a thin luminous straight line, e.g., generated by a laser line projector, onto the surface of the object that is to be reconstructed and then image the projected line with a camera. As shown in [figure 6.1](#) the projection of the laser line builds a plane that is called light plane or sheet of light. The optical axis of the camera and the light plane form an angle α , which is called angle of triangulation. The points of intersection between the laser line and the camera view depend on the height of the object. Thus, if the object onto which the laser line is projected differs in height, the line is not imaged as a straight line but represents a profile of the object. Using this profile, we can obtain the height differences of the object. To reconstruct the whole surface of an object, i.e., to get many height profiles, the object must be moved relative to the measurement system, i.e., the unit built by the laser line projector and the camera.

The sheet-of-light technique can be applied either to a calibrated measurement setup or to the uncalibrated setup. If the setup is calibrated, the measurement returns the disparities, the x , y , and z coordinates of the points that build the profiles in the world coordinate system (WCS, see [figure 6.1](#)), and a 3D object model that is derived from the x , y , and z coordinates. The disparities are returned in form of a disparity image, i.e., the disparities of each measured profile are stored in one row of the disparity image (see [figure 6.2](#) and note that the camera must be oriented such that the profiles are roughly parallel to the rows of the image). The x , y , and z coordinates are also not explicitly returned as values but are expressed as values of pixels within images. That is, three images are returned, one for the x coordinates (X), one for the y coordinates (Y), and one for the z coordinates (Z). The 3D object model contains information




Figure 6.1: Basic principle of sheet of light (light plane marked in gray).

about the 3D coordinates and the corresponding 2D mapping. If the setup is uncalibrated, only the disparity image and a score that describes how reliable the measurement result is can be returned by the measurement.

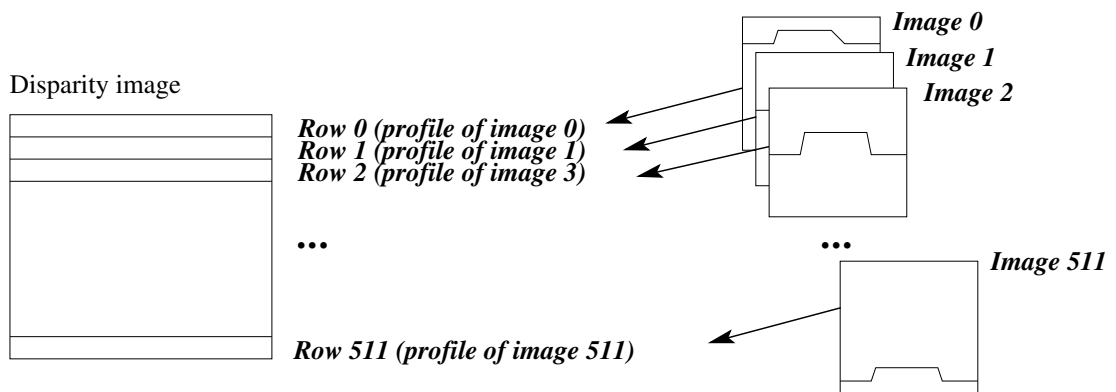


Figure 6.2: Disparity image: the disparity obtained from each profile (or image, respectively) is stored in a row of the disparity image.

Note that the disparity image for sheet of light has not exactly the same meaning as the disparity image described for stereo matching in [section 5.1](#) on page 107 and [section 5.3.5.1](#) on page 121. For stereo,

the disparity describes the difference between the row coordinates of the left and right stereo images, whereas for sheet of light, the disparity is built by the subpixel row values at which the profile was detected.

6.2 The Measurement Setup

The hardware needed for a sheet-of-light measurement consists of a projector that is able to project a thin luminous line, a camera, a positioning system, and the object to measure. In the following, we assume the projector to be a laser line projector and the positioning system to be linear (e.g., a conveyor belt), as these are very common in laser triangulation applications.

The relation between the projector, the camera, and the linear positioning system must not be changed, whereas the position of the object that is transported by the positioning system changes in relation to the projector-camera unit. Since the profile images are processed column by column, the profiles must be oriented roughly horizontal, i.e., roughly parallel to the rows of the image.

The relation between the laser line projector, the camera, and the object to measure can be described by different measurement setups. Figure 6.3 shows the three apparent configurations for the three components. In the first case, the camera view is orthogonal to the object and the light plane is tilted. The second case shows a tilted camera view and an orthogonal light plane. For the third case, both the camera view and the laser line are tilted.

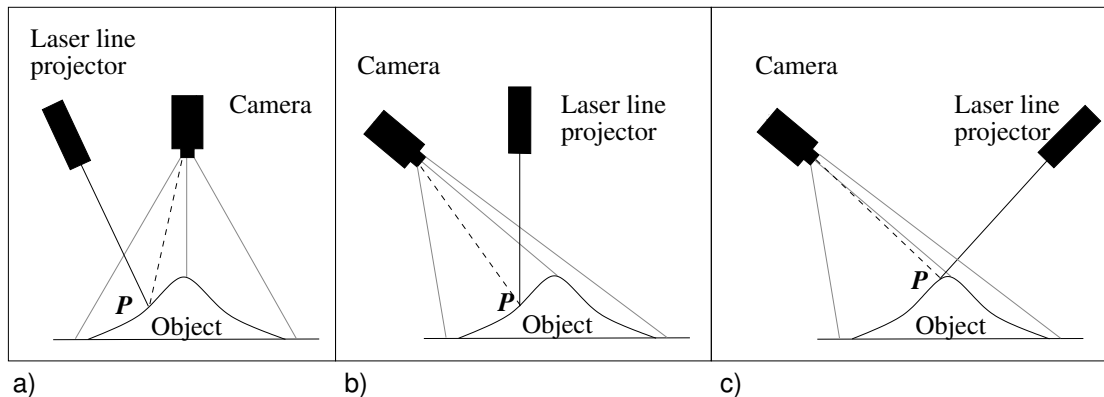


Figure 6.3: Basic configurations possible for a sheet-of-light measurement setup.

Figure 6.4 exemplarily shows a measurement setup as it is used for the examples that will be discussed in the following sections.

Which measurement configuration to use depends on the goal of the measurement and the geometry of the object. The configuration in figure 6.3 (a), e.g., is especially suitable if an orthogonal projection of the object in the image is needed for any reason. Then, a cuboid is imaged as a rectangle. For all configurations in which the camera is not placed orthogonal, it would be imaged as a trapezoid because of the perspective deformations (see figure 6.5).




Figure 6.4: Exemplary setup for a sheet-of-light measurement consisting of a camera, a laser line projector, and a positioning system.




Figure 6.5: With the camera being orthogonal to the object, an orthogonal projection of the object is possible: (left) orthogonal camera view, (right) perspective view.

The most important criterion for the selection of the measurement setup is the geometry of the object. The setup should be selected such that the amount of shadowing effects and occlusions is minimized. Occlusions occur if an object point is illuminated by the laser line but is not visible in the image, because other parts of the object lie between the line of sight of the camera and the object point (see [figure 6.6](#) on page 143, top). Shadowing effects occur if an object point is visible in the image but is not illuminated by the laser line, because other parts of the object lie between the laser projection and the imaged object

point (see [figure 6.6](#) on page 143, bottom).




Figure 6.6: Problems that have to be considered before selecting the measurement setup: (top) occlusions and (bottom) shadowing effects.

For all three setup configurations, the angle of triangulation, i.e., the angle between the light plane and the optical axis of the camera, should be in a range of 30° to 60° to get a good measurement accuracy. If the angle is smaller, the accuracy decreases. If the angle is larger, the accuracy increases, but more problems because of occlusions and shadowing effects are to be expected. Thus, you have to find a trade-off between the accuracy and the completeness of the measurement.

6.3 Calibrating the Sheet-of-Light Setup

This section describes how to calibrate the sheet-of-light measurement setup. If the uncalibrated case is sufficient for your application, you can skip this section and proceed with [section 6.4](#) on page 148.

The calibration of the sheet-of-light setup is applied to get the internal and external camera parameters, the orientation of the light plane in the WCS, and the relative movement of the object between two successive measurements. The calibration consists of the following steps:

1. Calibrate the camera.
2. Determine the orientation of the light plane with respect to the WCS.
3. Calibrate the movement of the object relative to the measurement setup.

The camera is calibrated by a standard camera calibration as described in [section 3.2](#) on page 40. As result, the camera calibration returns the internal camera parameters and the pose of the WCS relative to the camera coordinate system (external camera parameters).

To determine the light plane and its pose, we need at least three corresponding points (see [figure 6.7](#)), in particular two points in the plane of the WCS with ' $z=0$ ' (P1, P2) and one point that differs significantly in z direction (P3). Thus, you place the calibration object, e.g., the standard HALCON calibration plate, once or twice so that it lies in the plane of the WCS with ' $z=0$ ', and once so that a higher position can be viewed, i.e., the plate is either translated in z direction or it is placed in a tilted position. For each position of the calibration plate, you take two images, one showing the calibration plate and one showing the laser line. Note that you have to adapt the lighting in between to get one image with a clearly represented calibration plate and one image that shows a well-defined laser line. The translated or tilted position of the calibration plate should be selected so that the plane that is built by the points P1, P2, and P3 becomes as large as possible. The height difference should be at least as big as the height differences expected for the objects to measure.




Figure 6.7: Position of the third point (P3) obtained by (left) tilted calibration plate or (right) translated calibration plate.

Note that the laser line has to be projected onto the same plane in which the calibration plate is placed. But if possible, it should not be directly projected onto the calibration plate but only near to it. This is because the standard HALCON calibration plate is made of a ceramic that shows a strong volume scattering. This leads to a broadened profile (see [figure 6.8](#)), which results in a poor accuracy. If you use a calibration object that consists of a material with different reflection properties, this might be no problem.

Note further that the three corresponding points described above represent the minimum number of points needed to obtain a plane. To enhance the precision of the calibration, redundancy is needed; thus, we recommended to measure more than three corresponding points. Then, the light plane is approximated by fitting a plane into the obtained point cloud.




Figure 6.8: The white parts of a HALCON calibration plate show a very broad laser line because of volume scattering.

In the final step, the pose describing the movement of the object must be calibrated using two images containing a calibration plate that was moved by the positioning system by a known number of movement steps.

Summarized, we have to acquire:

- a set of images for the camera calibration,
- at least two images that clearly show the laser line in different planes and which correspond to images that were used for the calibration, and
- at least two images that show the calibration plate in the plane of the linear positioning system. Between the acquisition of the first and the second image, the calibration plate must be moved by the positioning system by a known number of movement steps.

The HDevelop example program

`hdevelop\Applications\Measuring-3D\calibrate_sheet_of_light.hdev` shows in detail how to calibrate a sheet-of-light measurement setup.

For the first step, i.e., the camera calibration, initial values for the internal camera parameters and for the thickness of the calibration plate are set.

```
StartParameters := [0.0125, 0.0, 0.0, 0.0, 0.0, 0.0, 0.000006, 0.000006, \
                    376.0, 120.0, 752, 240]
CalTabDescription := 'caltab_30mm.descr'
CalTabThickness := .00063
```

Then, the calibration images are read. These should fulfill the requirements that are described for a camera calibration in [section 3.2.4.1](#) on page 49. Now, for each image, the calibration plates are searched and their marks and poses are derived.

```

NumCalibImages := 20
for Index := 1 to NumCalibImages by 1
  read_image (Image, 'sheet_of_light/connection_rod_calib_'+Index$.2')
  find_caltab (Image, PlateRegion, CalTabDescription, 3, 105, 5)
  find_marks_and_pose (Image, PlateRegion, CalTabDescription, \
                      StartParameters, 128, 10, 18, 0.9, 15, 100, \
                      MarkRows, MarkColumns, EstimatedPose)
  set_calib_data_observ_points (CalibDataID, 0, 0, Index, MarkRows, \
                               MarkColumns, 'all', EstimatedPose)
endfor

```

With the obtained data, the actual camera calibration is performed, so that the internal camera parameters (CameraParameters) and the external camera parameters (camera poses) for all calibration images can be obtained. The internal camera parameters and the camera pose for one of the calibration images are the first two variables that we need for the sheet-of-light measurement that is described in the next section.

```

calibrate_cameras (CalibDataID, Errors)
get_calib_data (CalibDataID, 'camera', 0, 'params', CameraParameters)

```

Note that by selecting the camera pose of one of the calibration images you define the origin of the WCS used for the measurement.

For the second step, i.e., the orientation of the light plane in relation to the WCS, the poses of two of the calibration images are needed. The images show the calibration plates in different heights. The pose of one image is used to define the WCS and the pose of the other image is used to define a temporary coordinate system (TCS). For both images, the origins of the poses are shifted with [set_origin_pose](#) to consider the thickness of the calibration plate.

```

Index := 19
get_calib_data (CalibDataID, 'calib_obj_pose', [0, Index], 'pose', \
               CalTabPose)
set_origin_pose (CalTabPose, 0.0, 0.0, CalTabThickness, CameraPose)
Index := 20
get_calib_data (CalibDataID, 'calib_obj_pose', [0, Index], 'pose', \
               CalTabPose)
set_origin_pose (CalTabPose, 0.0, 0.0, CalTabThickness, TmpCameraPose)

```

For each of the two calibration images a corresponding laser line image was acquired. There, the laser line is clearly projected onto the same plane that contained the calibration plate in the calibration image. With the two laser line images and the poses obtained from the two corresponding calibration images the procedure `compute_3d_coordinates_of_light_line` calculates the 3D coordinates of the points that build the laser lines. The obtained point cloud consists of the points of the light plane in the plane of the WCS with 'z=0' (see P1 and P2 in [figure 6.7](#) on page 144) and the points of the light plane in the plane of the TCS with 'z=0' (see P3 in [figure 6.7](#) on page 144).

```

read_image (ProfileImage1, \
            'sheet_of_light/connection_rod_lightline_019.png')
compute_3d_coordinates_of_light_line (ProfileImage1, MinThreshold, \
                                     CameraParameters, [], CameraPose, \
                                     X19, Y19, Z19)

read_image (ProfileImage2, \
            'sheet_of_light/connection_rod_lightline_020.png')
compute_3d_coordinates_of_light_line (ProfileImage2, MinThreshold, \
                                     CameraParameters, TmpCameraPose, \
                                     CameraPose, X20, Y20, Z20)

```

Now, the procedure `fit_3d_plane_xyz` fits a plane into the point cloud. This plane is the light plane, for which the pose is needed as the third variable for the calibrated sheet-of-light measurement. This pose (`LightPlanePose`) is calculated from the plane using the procedure `get_light_plane_pose`.

```

procedure fit_3d_plane_xyz (X, Y, Z, Ox, Oy, Oz, Nx, Ny, Nz, MeanResidual)
get_light_plane_pose (Ox, Oy, Oz, Nx, Ny, Nz, LightPlanePose)

```

In the third step, i.e., the calibration of the movement of the object in relation to the measurement setup, the calibration plate is moved in discrete steps by the linear positioning system that will be used also for the following measurement. To calibrate the movement of the linear positioning system, two images with different movement states are needed. To enhance the accuracy, we do not use images of two succeeding movement steps but use images with a known number of movement steps between them. Here, the number of movement steps between both images is 19.

```

read_image (CaltabImagePos1, 'sheet_of_light/caltab_at_position_1.png')
read_image (CaltabImagePos20, 'sheet_of_light/caltab_at_position_2.png')
StepNumber := 19

```

Now, for both images the poses of the calibration plates are derived.

```

find_caltab (CaltabImagePos1, CaltabPos1, CalTabDescription, 3, 105, 15)
find_marks_and_pose (CaltabImagePos1, CaltabPos1, CalTabDescription, \
                    CameraParameters, 128, 10, 18, 0.9, 15, 100, \
                    RCoordPos1, CCoordPos1, CameraPosePos1)
find_caltab (CaltabImagePos20, CaltabPos20, CalTabDescription, 3, 105, 15)
find_marks_and_pose (CaltabImagePos20, CaltabPos20, CalTabDescription, \
                    CameraParameters, 128, 10, 18, 0.9, 15, 100, \
                    RCoordPos20, CCoordPos20, CameraPosePos20)

```

Then, the pose that describes the transformation between these two poses, i.e., the transformation needed for 19 movement steps, is calculated (`MovementPoseNSteps`). Note that a rotation is not assumed and therefore all rotational elements are set to 0.

```

pose_to_hom_mat3d (CameraPosePos1, HomMat3DPos1ToCamera)
pose_to_hom_mat3d (CameraPosePos20, HomMat3DPos20ToCamera)
pose_to_hom_mat3d (CameraPose, HomMat3DWorldToCamera)
hom_mat3d_invert (HomMat3DWorldToCamera, HomMat3DCameraToWorld)
hom_mat3d_compose (HomMat3DCameraToWorld, HomMat3DPos1ToCamera, \
                  HomMat3DPos1ToWorld)
hom_mat3d_compose (HomMat3DCameraToWorld, HomMat3DPos20ToCamera, \
                  HomMat3DPos20ToWorld)
affine_trans_point_3d (HomMat3DPos1ToWorld, 0, 0, 0, StartX, StartY, StartZ)
affine_trans_point_3d (HomMat3DPos20ToWorld, 0, 0, 0, EndX, EndY, EndZ)
MovementPoseNSteps := [EndX-StartX, EndY-StartY, EndZ-StartZ, 0, 0, 0, 0]

```

To get the pose for a single movement step (`MovementPose`), the elements of `MovementPoseNSteps` that describe a translation are divided by the number of steps. `MovementPose`, together with the internal and external camera parameters and the pose of the light plane can now be used to apply a calibrated sheet-of-light measurement.

```
MovementPose := MovementPoseNSteps/StepNumber
```

For details about the proceedings inside the stated procedures, we refer to the example.

6.4 Performing the Measurement

A sheet-of-light measurement is applied to get height information for the object to measure. This height information is presented by a disparity image in which each row contains the disparities of one measured profile of the object (see [figure 6.2](#) on page 140), by the images X, Y, and Z that express the x, y, and z coordinates of the measured profiles as values of pixels within images, or by a 3D object model that contains the coordinates of the object's 3D points and the corresponding 2D mapping. The images X, Y, and Z as well as the 3D object model can be obtained only for a calibrated measurement setup, whereas the disparity image can be obtained also for the uncalibrated case. A sheet-of-light measurement consists of the following basic steps:

1. Calibrate the measurement setup (if a calibrated measurement is needed) as described in the previous section.
2. Create a sheet-of-light model with `create_sheet_of_light_model` and set additional parameters with successive calls to `set_sheet_of_light_param`.
3. Acquire images for each profile to measure, e.g., using `grab_image_async`.
4. Measure the profile of each image with `measure_profile_sheet_of_light`.
5. Get the results of the measurement with successive calls to `get_sheet_of_light_result` or, if a 3D object model is required, with a single call to `get_sheet_of_light_result_object_model_3d`.

6. If only the uncalibrated case was applied and a disparity image was obtained, but the x , y , and z coordinates or the 3D object model are still needed, you can subsequently apply a calibration. Then, you have to calibrate the measurement setup like described in the previous section and add the obtained camera parameters to the model with `set_sheet_of_light_param`. Afterwards you call the operator `apply_sheet_of_light_calibration` with the disparity image and the adapted sheet-of-light model. The resulting images that contain the coordinates X , Y , and Z or the 3D object model are queried from the model with `get_sheet_of_light_result` or `get_sheet_of_light_result_object_model_3d`, respectively. How to subsequently apply a sheet of light calibration to a disparity image is shown in the HDevelop example program `hdevelop\Applications\Measuring-3D\calibrate_sheet_of_light.hdev`.
7. Clear the model from memory with `clear_sheet_of_light_model`.

Optionally, you can query all parameters that you have already set for a specific model or that were set by default using `get_sheet_of_light_param`. To query all parameters that can be set for a sheet-of-light model you call `query_sheet_of_light_params`.

6.4.1 Calibrated Sheet-of-Light Measurement

How to apply a calibrated sheet-of-light measurement is shown in the HDevelop example program `hdevelop\Applications\Measuring-3D\reconstruct_connection_rod_calib.hdev`, which measures the object shown in [figure 6.9](#).



Figure 6.9: Object to measure.

The first step is to assign the information obtained by the calibration of the sheet-of-light measurement setup (see previous section) to a set of variables.

```
CamParam := [0.0126389, 658.417, -2.43408e+07, 1.09591e+12, -0.0995053, \
             0.0286728, 6.00046e-06, 6e-06, 388.079, 120.651, 752, 240]
CamPose := [-0.00180691, -6.9929e-005, 0.299885, 0.590058, 0.572089, 180.027, 0]
LightplanePose := [0.00270987, -0.00548833, 0.00843719, 66.9927, 359.72, \
                  0.659447, 0]
MovementPose := [7.78327e-008, 0.000120111, 1.98684e-006, 0, 0, 0, 0]
```

Then, a sheet-of-light model is created for a rectangular region of interest using `create_sheet_of_light_model`. The ROI should be selected as large as necessary but as small as possible. That is, it should approximately be some pixels larger than the width of the object in width and the maximal expected displacement of the laser line caused by the height of the object, i.e., the largest expected disparity, in height.

Now, some parameters are set with `set_sheet_of_light_param`. As a calibrated measurement is applied, the parameter 'calibration' is set to 'xyz'. For an uncalibrated measurement, it would be 'none', which is the default. Further, the variables with the calibration information are passed as values to the corresponding parameters for the internal camera parameters ('camera_parameter'), the external camera parameters ('camera_pose'), the pose of the light plane ('lightplane_pose'), and the movement of the object relative to the measurement setup ('movement_pose').

```
gen_rectangle1 (ProfileRegion, 120, 75, 195, 710)
create_sheet_of_light_model (ProfileRegion, ['min_gray','num_profiles', \
      'ambiguity_solving'], [70,290,'first'], \
      SheetOfLightModelID)
set_sheet_of_light_param (SheetOfLightModelID, 'calibration', 'xyz')
set_sheet_of_light_param (SheetOfLightModelID, 'scale', 'mm')
set_sheet_of_light_param (SheetOfLightModelID, 'camera_parameter', CamParam)
set_sheet_of_light_param (SheetOfLightModelID, 'camera_pose', CamPose)
set_sheet_of_light_param (SheetOfLightModelID, 'lightplane_pose', \
      LightplanePose)
set_sheet_of_light_param (SheetOfLightModelID, 'movement_pose', \
      MovementPose)
```

Then, for each profile to measure an image is acquired (see, e.g., [figure 6.10](#)) to apply the actual measurement. Here, the images for each movement step are read from file with `read_image`. In practice, you will most probably grab the images directly from your image acquisition device using `grab_image_async` (see [Solution Guide II-A](#) for details about image acquisition). For each image, the profile within the rectangular region of interest is measured with `measure_profile_sheet_of_light`, i.e., the disparities for the profile are determined and stored in the sheet-of-light model.

```
for Index := 1 to 290 by 1
  read_image (ProfileImage, 'sheet_of_light/connection_rod_'+Index$.3')
  measure_profile_sheet_of_light (ProfileImage, SheetOfLightModelID, [])
endfor
```




Figure 6.10: Measure profile inside a rectangular ROI.

The default for the number of profiles to measure is 512. You can change it with the parameter 'num_profiles' within `create_sheet_of_light_model` or `set_sheet_of_light_param`. If you

measure more than the specified number of profiles, the value of 'num_profiles' is automatically adapted in the model. Nevertheless, this adaptation requires additional runtime. Thus, we recommend to set 'num_profiles' to a suitable value before starting the measurement. Note that the number of measured profiles defines the number of rows and the width of the ROI used for the measurement defines the number of columns for the result images (i.e., the disparity image, the X, Y, and Z images, and the score image).

After all measurements were performed, the results of the sheet-of-light measurement are queried with calls to `get_sheet_of_light_result` and `get_sheet_of_light_result_object_model_3d`. Here, we query the disparity image (ResultName set to 'disparity'), the images X, Y, and Z (ResultName set to 'x', 'y', and 'z', respectively), and the 3D object model. The images X, Y, and Z are shown in [figure 6.11](#). The 3D object model is interactively displayed using the procedure `disp_object_model_3d` as shown in [figure 6.12](#).

The interpretation of the gray values of the disparity image and the images X, Y, and Z is as follows: black parts are outside of the domain of the resulting image, i.e., they indicate parts for which no 3D information could be reconstructed. For the pixels inside the domain of the image bright parts show low object parts and dark parts show higher object parts. Note that in this example the images are not visualized by their default gray values but are converted additionally by a look-up table so that the images are colored. This is done because the human eye can separate much more colors than gray values. Thus, details can be better distinguished during a visual inspection.

```
get_sheet_of_light_result (Disparity, SheetOfLightModelID, 'disparity')
get_sheet_of_light_result (X, SheetOfLightModelID, 'x')
get_sheet_of_light_result (Y, SheetOfLightModelID, 'y')
get_sheet_of_light_result (Z, SheetOfLightModelID, 'z')
get_sheet_of_light_result_object_model_3d (SheetOfLightModelID, \
                                           ObjectModel3DID)
```

At the end of the measurement, the model is cleared from memory.

```
clear_sheet_of_light_model (SheetOfLightModelID)
```

6.4.2 Uncalibrated Sheet-of-Light Measurement

The uncalibrated sheet-of-light measurement is shown in the HDevelop example program `hdevelop\Applications\Measuring-3D\reconstruct_connection_rod_uncalib.hdev`. Here, no calibration results are needed, so we simply create the model for the specified region of interest and set the few needed parameters directly within `create_sheet_of_light_model`.

```
gen_rectangle1 (ProfileRegion, 120, 75, 195, 710)
create_sheet_of_light_model (ProfileRegion, ['min_gray','num_profiles', \
                                           'ambiguity_solving','score_type'], [70,290, \
                                           'first','width'], SheetOfLightModelID)
```

The actual measurement is applied by the same process used for the calibrated measurement.




Figure 6.11: Result of calibrated sheet-of-light measurement: images representing the (from top to bottom) x, y, and z coordinates of the object.

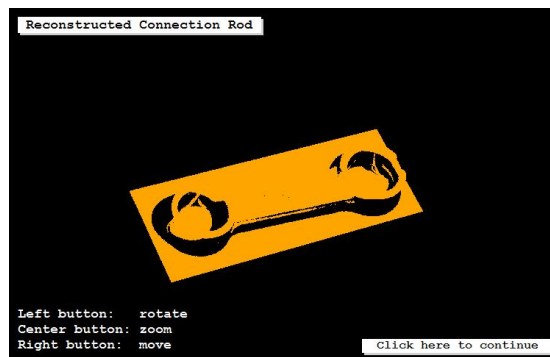


Figure 6.12: Result of calibrated sheet-of-light measurement: 3D object model.

```

for Index := 1 to 290 by 1
  read_image (ProfileImage, 'sheet_of_light/connection_rod_'+Index$.3')
  measure_profile_sheet_of_light (ProfileImage, SheetOfLightModelID, [])
endfor

```

As result, we can only query the disparity image (see [figure 6.13](#)) and the score (see [section 6.5](#)) of the measurement (ResultName set to 'score').


```
get_sheet_of_light_result (Disparity, SheetOfLightModelID, 'disparity')
get_sheet_of_light_result (Score, SheetOfLightModelID, 'score')
```



Figure 6.13: Result of uncalibrated sheet-of-light measurement: disparity image.

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

```
clear_sheet_of_light_model (SheetOfLightModelID)
```

6.5 Using the Score Image

Caused by the specific characteristics of a laser line projector and the general principle of triangulation the results of a sheet-of-light measurement, i.e., the disparities or the calibrated coordinates, sometimes show disturbing artifacts. The score image can be used to detect and partially remove artifacts.

There are two types of artifacts. The first type is caused by the geometry of the surface that is to be reconstructed. As illustrated in [figure 6.14](#), compared to flat and smooth surfaces (e.g., the object in [figure 6.13](#)), curved surfaces with a small radius of curvature and surfaces with a significant slope lead to a broadening of the light line. Furthermore, the light distribution within the profile might be no longer symmetric, which leads to a reduced measurement accuracy.

By using the width of the profile stored in the score image (for each pixel of the disparity, the score value is set to the number of pixels used to determine the disparity value) it is possible to detect artifacts and to reject the corresponding disparities or the corresponding calibrated coordinates. [Figure 6.15](#) shows the score image obtained by the uncalibrated sheet-of-light measurement performed in the HDevelop example program `hdevelop\Applications\Measuring-3D\reconstruct_connection_rod_uncalib.hdev`. The gray values inside the score image indicate the widths of the laser line in each pixel. Thus, artifacts, in this case parts with a significantly broadened laser line, can be recognised easily by their brightness.

In the example, the artifacts are rejected from the disparity image by applying a threshold to the score image (pixels with a value larger than 7.5 are rejected) and reducing the disparity image to the obtained region.

```
threshold (Score, ScoreRegion, 1.5, 7.5)
reduce_domain (Disparity, ScoreRegion, DisparityReduced)
```




Figure 6.14: Curved surfaces with a small radius of curvature and surfaces with a significant slope lead to a broadening of the light line and thus to a low score: (top) small influence of curvature, (bottom) large influence of curvature.



Figure 6.15: Result of uncalibrated sheet-of-light measurement: score image (score_type set to 'width', i.e., bright parts indicate artifacts).

The second type of artifacts is caused by the interaction of the coherent laser light with the surface of the object. Laser light produces disturbing interference patterns when it is projected on a rough textured surface. Those interference patterns are called speckle and can be considered as a non-additive noise, which means that this noise can not be reduced by averaging during the image acquisition. In this case, the only way to increase the measurement accuracy is to use a higher aperture for the image acquisition or a low-speckle line projector. Note that enlarging the aperture for the image acquisition device will also reduce the depth of field which might be an undesired side effect. If your application requires high accuracy, we strongly recommend to use low-speckle projection devices. Note that speckle in most cases

is the limiting factor for the accuracy of laser triangulation systems.

6.6 3D Cameras for Sheet of Light

The proceeding described in the previous sections works for any standard 2D camera that is suitable for machine vision. An alternative is to use specific 3D cameras for which the sheet-of-light measurement is applied inside the camera. These cameras are more expensive than standard 2D cameras, but the sheet-of-light measurement becomes significantly faster because of the reduced CPU load. Using one of these cameras, you simply access the camera with HALCON and basically leave the measurement to the camera.

Generally, we distinguish between cameras with an inbuilt laser, i.e., the camera and the laser are integrated in a single unit, and cameras for which the laser is mounted separately.

If the camera and the laser are integrated in a single unit, the measurement setup is restricted to a fixed angle of triangulation and should be oriented in a defined way. For example, the SICK Ruler camera should be oriented so that the laser is perpendicular to the linear positioning system. Because of the preset measurement setup, the camera and the orientation of the light plane with respect to the world coordinate system are already calibrated. Thus, no further processing with HALCON is needed to obtain calibrated height profiles.

If the camera and the laser are mounted separately, any configuration of the measurement setup is possible (see [section 6.2](#) on page 141), but by default the result of the measurement is uncalibrated. If the result of the measurement is needed in world coordinates, you can either query the uncalibrated data from the camera and subsequently apply a calibration with HALCON as described in [section 6.4](#) on page 148, or, before performing the actual measurement, you apply a calibration that is provided specifically for the selected camera. For the SICK Ranger cameras, e.g., the camera-specific calibration needs the software provided with the camera (the SICK Coordinator tool) and a specific calibration object that has to be purchased separately.

Note that in contrast to the proceeding described in the previous sections, the movement of the linear positioning system is mostly assumed to be known, because the measurement of each profile is triggered by a signal coming from an encoder on the linear positioning system. That is, when working with an encoder and if the thus obtained accuracy is sufficient, it is not necessary to calibrate the distance between two profiles.

Chapter 7

Depth from Focus

Depth from focus (DFF) is a method that enables the reconstruction of 3D surface information from several images taken at different focus distances between camera and object. It allows a highly accurate (up to about 5 micrometers) non-destructive 3D measurement of surfaces. DFF is even more precise than the methods stereo [chapter 5](#) on page 105 and laser triangulation with sheet of light [chapter 6](#) on page 139. Furthermore, the setup requires only a single camera, therefore, it is possibly more compact than, e.g., a stereo setup. DFF requires, however, cameras with telecentric or microscope lenses in order to achieve a (nearly) parallel projection. Therefore, DFF is only suitable for small objects. Examples for suitable objects in semiconductors industry are a ball grid array (BGA) (the result of DFF on a single ball of a BGA can be viewed in [figure 7.1](#)) or solder paste inspection, another application in the engineering sector is the inspection of indexable inserts.




Figure 7.1: This image shows the 3D surface reconstruction of a single ball on a ball grid array.

7.1 The Principle of Depth from Focus

With depth from focus, you can reconstruct the surface of a 3D object based on the knowledge that object points have different distances to the camera and the camera has a limited depth of field. Depending on the distance and the focus, object points are displayed more or less sharply in the image, i.e., only those pixels within the correct distance to the camera are focused. Taking images with various object distances, each object point can be displayed sharply in at least one image. Such a sequence of images is called “focus stack”. By determining in which image an object point is in focus, i.e., sharply imaged, the distance of each object point to the camera can be calculated. This principle is clarified in [figure 7.2](#).




Figure 7.2: This figure shows the focus stack of images on the left side and the corresponding focus value - that is determined for each pixel - on the right side. The best focus point is the image where a pixel has the highest sharpness.

For more information about determining sharpness, please refer to `hdevelop\Applications\FA\determine_sharpness.hdev`. In this example, a flat object is imaged, and the global sharpness of the image is determined. In the case of DFF, the sharpness will be determined for each single pixel. If you need real three-dimensional information, e.g., if you want to further use your DFF results for any 3D measurements or surface-based matching (see [section 4.3](#) on page 93) you need to telecentrically calibrate your system. Information on how to perform a calibration can be found in the chapter about 3D camera calibration ([section 3.2](#) on page 40).

Depth of field (DOF) is a similar term which, however, is not a method but a technical term concerning the camera. The depth of field is the range of distance within which the image is sharp as opposed to the best focus point, which is the point with perfect sharpness in the image. The DOF depends on the pixel size, the aperture (f-number), the focal length in the case of a non-telecentric lens, and the focusing distance. A low depth of field means that only a small slice of the object is sharply imaged (see [figure 7.3](#)), whereas a high depth of field means that a big part or maybe the whole image is sharp. For DFF, a low depth of field leads to a higher precision. In order to obtain a low depth of field, use lenses with a high aperture which is a small f-number on your lens.

A method for 3D surface measurement, similar to DFF, is depth from defocus. This method requires one sharp image of the foreground and one sharp image of the background. The distance of all points that lie between foreground and background is interpolated by their amount of “unsharpness”. As it only




Figure 7.3: a) Only a small slice of the image is sharp. Therefore point P1 is mapped to a single point and is therefore sharp, whereas P2 is mapped to a spot and is consequently blurry. b) If the object is very flat, the whole object can be sharp at the same time, even if the depth of field is not very high.

depends on two images, it may be faster but it is also not as precise as depth from focus. Depth from defocus is not available in HALCON.

7.1.1 Speed vs. Accuracy

As mentioned before, depth from focus allows 3D measurement with a maximum resolution of up to about 5 micrometers. It is therefore perfect for applications that require very high accuracy, however, at the cost of a longer runtime.

7.1.1.1 Depth from Focus

DFF may consume more processing time due to the number of images that may need to be processed. This does, however, depend on the actual number of images that are necessary for the specific task. Generally speaking, the more images need to be processed, the higher the accuracy, the longer the runtime. Therefore, for some applications the runtime will naturally not be very high, because the required accuracy may be lower. Nevertheless, the runtime can also be improved for applications that require a high as well as for those which require a low accuracy. Internally, the speed of HALCON's `depth_from_focus` is sped up using parallelization (only available in the 'local' mode).

7.1.1.2 Depth from Focus vs. Other Methods

Due to the high number of images that have to be processed, the data volume is also high. This leads to a reduced speed compared to other methods. DFF has proved to be very useful in the microscopical

range (small objects that are magnified more than once) and often easier to realize than other methods. In many cases, the setup can be more compact than stereo ([chapter 5](#) on page 105) which needs quite a lot of space due to the two cameras. Nevertheless, the choice of a suitable lens is very important for DFF, and it can counter the advantage of using a single camera. For small objects having dimensions smaller than some ten millimeters, sheet of light ([chapter 6](#) on page 139) might become very expensive because of the thin laser line that is needed and less precise due to speckle. Photometric stereo (see “[3D Reconstruction ▸ Photometric Stereo](#)”) is usually more precise and easier to realize for flat objects without steep geometric edges. For macroscopic measurements (measurement range up to 100 mm), stereo or sheet of light would be quicker as they require less images.

7.2 Setup

Before actually starting your application with depth from focus, it is important to first set up your application environment properly. This section concentrates on the camera and application object setup to prepare the image acquisition. Additional information about equipment and image acquisition can be read in [section 7.3.1](#) on page 166.

7.2.1 Camera

7.2.1.1 Recommended Camera Setup and Adjustments

As performance of depth from focus for your application depends on the used lens, the depth of field and the precision of the movements, it is important to use the right camera with the right adjustments. Please read our recommended camera setup and adjustments for the best results:

1. **If possible, use a camera with a telecentric lens.** In order to perform depth from focus, a camera with a telecentric lens or an almost telecentric lens, e.g., a microscopic imaging system, will produce the best results. Only a telecentric lens enables you to take images with exactly the same field of view at different focus positions. This is important because depth from focus uses image coordinates and for each pixel finds the image in which it is displayed sharply. Only in images taken with a telecentric lens, those pixel are comparable, i.e., in the same position. Therefore, DFF is a good method for measuring small objects, like, e.g., microelectrical workpieces. Good results have been achieved, for example, when performing DFF with the XENOPLAN telecentric lens series by Schneider-Kreuznach. Note however that even when using a telecentric lens, aberration - the effect that not all pixels are in focus on a planar surface - can occur. Aberration influences the accuracy of DFF and should be calibrated. How to calibrate aberration is described in [section 7.4.1](#) on page 169. It is also possible, though not recommended due to the reasons mentioned before, to use DFF with a standard lens. For performing DFF with a standard lens, please read [section 7.6](#) on page 171. A problem that is related to the lens is the correspondence problem. This refers to effect that the position of the pixels outside the optical axis shifts when the distance is changed during image acquisition. Those planes that lie close to the border of the field of view as well as those that lie diagonal to the optical axis are affected most. A small depth of field reduces the effect for planes with other directions. The correspondence problem results in wrong focus distances in the depth image. This effect is minimal when using a telecentric lens. For telecentric lenses, it is

important that the movement causing a change of focus is applied parallel to the optical axis. For more information about this effect see [figure 7.5](#).

2. **Use mirrors to obtain focus images.** For DFF, mirrors are a very good solution because they can be moved very quickly and accurately and - if necessary - can be replaced easily and inexpensively. Moving the camera may be harmful to the camera sensor which suffers under the vibrations. Moving the object might be difficult as the object is probably located on a conveyor belt. All movements for DFF have to be performed very precisely. What such a setup with camera and mirrors can look like is shown in [figure 7.4](#). Note that the focus should only be moved along the optical axis.
3. **Use a low depth of field to achieve a higher accuracy.** This requires, however, more images at different focus positions. In contrast, large distances between the images require a higher depth of field and lead to a less precise height reconstruction. A low depth of field requires - as stated before - images at various focus positions.
4. **Use a high aperture.** The aperture needs to be as open as possible as this reduces the depth of field which is responsible for a higher precision as mentioned before. The highest possible precision is, therefore, limited by the depth of field.




Figure 7.4: Rays from the object first cross a semi-transparent mirror and are then reflected by a mirror mounted on a linear piezo stage. The rays then reach the semi-transparent mirror again and are this time reflected back to the camera. By moving the mirror that is mounted on the linear piezo stage, the distance between the object and the camera can be varied in a controlled way, which makes it possible to acquire a sequence of images with varying focus. .

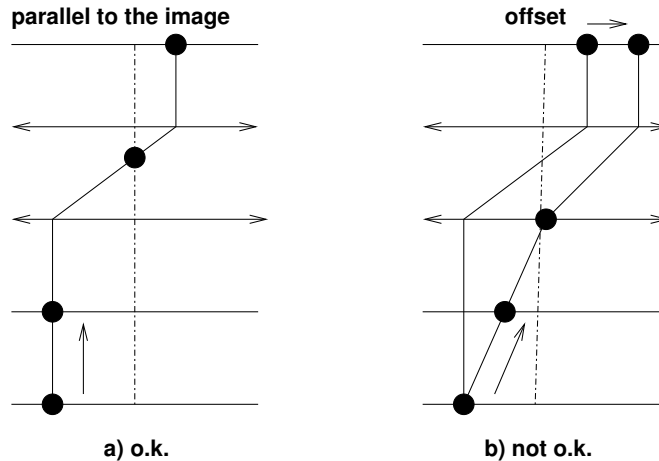


Figure 7.5: This figure illustrates the correspondence problem. In figure a), a change of focus is applied parallel to the optical axis, which is good. If the focus is not changed parallel to the optical axis, a lateral offset occurs, as depicted in figure b).

7.2.1.2 Acquire Measurement Range

There are four rules that define how the measurement range can be acquired.

First rule: The distance range has to exceed the height of the measure object.

Second rule: The range in which the image part changes from sharp over unsharp to sharp should be approximately five images (as depicted in [figure 7.6](#)). Otherwise, areas of the object's surface cannot be determined correctly and, therefore, cannot be measured precisely.

Third rule: As the depth of field of the used lens is fixed, the minimum number of focus positions has to exceed:

$$\frac{\text{focus range in } m}{\text{depth of field of the used lens}}$$

The reason for the shifting of the depth of field area of the lens. This shift needs to be smaller than the depth of field of the lens so that the depth of field areas from two successive images can overlap. The more those image overlap, the higher the achieved precision. The downside is an increased runtime due to the high number of images. You cannot, however, increase precision indefinitely as it is also limited by the camera noise. The smaller the overlap, the smaller the achieved precision, the shorter the runtime.

Fourth rule: If there is a limit to the number of images that can be acquired, for example, because you have limited runtime, the depth of field of the used lens needs to be increased (e.g., by closing the aperture of the lens).

Images from a focus stack can be viewed in [figure 7.7](#).

Depending on the direction of image acquisition, i.e., from the camera or towards the camera, the result of the measurement is either a distance image or a height image.




Figure 7.6: This image depicts an object from which images are taken at different focusing distances (indicated by the dashed lines) for a focus sequence. It also exemplarily shows the depth of field, consisting of five images in which the image with the sharpest pixel is marked by a dot.




Figure 7.7: The images above are part of a focus sequence. They show the same object but are taken at different focus positions. Image a) focuses on the top of the object, image b) is sharp at a medium height and image c) is sharp in the background.

7.2.2 Illumination

7.2.2.1 Choosing Illumination

In order to enhance the surface texture of the object, direct illumination is needed. At the same time, reflections have to be minimized. Therefore, coaxial and lightfield lighting would not work. A suitable lighting could be an illumination from various directions, as depicted in [figure 7.8](#), because lighting that comes from different directions enhances the structure very well and also causes few reflections. As an alternative darkfield lighting is also possible, because the low angle enhances the surface structure - it does, however, result in an image that has quite dark parts in some areas. Other illumination setups can be used as long as the lighting enhances the surface texture and causes as few reflections as possible. A suitable illumination, therefore, highly depends on the application object's features. For more information about lighting, please refer to the Solution Guide II-A, [appendix C.1](#) on page 57.




Figure 7.8: Illumination from various directions leads to good results for depth from focus. It can, for example, be produced with lights that are arranged within a dome.

7.2.2.2 Overexposure

Overexposure is an illumination-connected problem that may occur and reduce the accuracy of the resulting image. It leads to a loss of information in the overexposed region which then reaches saturation (a gray value of 255). Furthermore, overexposure causes the detected false sharp pixels in blurry areas which are the result of high frequencies between saturated image areas and blurry areas. The results of overexposure are visualized in [figure 7.9](#). In order to avoid overexposure, it is recommended to take darker images for your measurements and improve visualization with the operators `scale_image` and `scale_image_max` as presented in the example [section 7.3.2.1](#) on page 167.

7.2.3 Object

In order to successfully perform DFF, it is important to know your application object.

First of all, DFF is a practical method for any small object with textured surface. For objects that are larger than that, there are several other 3D methods available like binocular stereo ([chapter 5](#) on page 105). DFF is can, however, still be used for very high objects that would usually require a very high number of images and thus slow down the measuring process a lot. If the height of such an object is required there are two possibilities of gaining these measuring results.

- One possibility is to take a certain number (maybe 10) images that depict the bottom of the object, then also take a certain number images showing the top of the object. Because you know the distance between the highest image that is taken of the bottom of the object and the first image taken of the top of the object, you can just add this height to the images taken of the top of the object.
- Another possibility is to acquire two sequences: one has to be taken around the bottom of the object and another one at the top of the objects. Within each of the sequences the sharpest level




Figure 7.9: a) The sharp pixel in this figure is overexposed. b) If wider parts of the image are overexposed, instead of detecting the correct location of a sharp pixel, wrong sharp pixels are detected at the border between the saturated image areas and the blurry image areas.

is detected. As the movement of the motor that determines the acquisition of images at different focus levels is known, the distance between those images can be calculated.

d, two image sequences are acquired, one taken around the highest point of measurement and one sequence around the lowest point of measurement. Once, the sharpest pixels at the highest and lowest point are found, the height can be calculated.

Regular (passive) DFF can be used for objects with a textured surface. Fortunately, small objects mostly have a structured surface when observed at a suitable magnification. The exception are perfectly polished, i.e., specular surfaces.

If there is very little texture on the object, it is possible to still perform the so-called active depth from focus. Active DFF compensates for the missing surface structure by projecting a texture on the object. When performing active DFF, it is possible to produce focus images with the projector. It is, however, necessary that - for every image that is taken - the depth of field of the camera is conform with the depth of field of the projector. Active DFF does work for objects with little texture on their surfaces, it does not work for reflecting surfaces, though.

7.3 Working with Depth from Focus

Depth from focus requires a focus stack of images, i.e., images that have each a different distance to the object and in which, therefore, different pixels are sharp. Depth from focus is then applied to this focus stack returning a sharp image, i.e., an image containing only gray values from focused, sharp pixels as these are the relevant pixels for DFF, as well as a depth image showing the three-dimensional shape of the object that is inspected. Those two results can be viewed in [figure 7.10](#).

7.3.1 Rules for Taking Images

There are several rules that should be followed when taking images for measuring with depth from focus. Those rules concern image quality as well as the system setup.

7.3.1.1 Rules for Achieving a High Image Quality

1. **Avoid overexposure.**
2. **Avoid reflections.**
3. **Use a camera with a wide dynamic range.**
4. **Use direct illumination** - preferably from several directions.
5. **Use a camera with low noise.**
6. **The maximum number of images is limited by camera noise.** Taking more images to cover the whole object only improves the precision of the reconstruction to a certain degree which is, amongst others, limited by the camera noise.
7. **As a rule of thumb, the minimum number of images that should be acquired for DFF is 10 and the maximum number is about 150.** While 10 images can be just sufficient for measuring (depending on the required accuracy), more than 150 images will usually not improve accuracy any more. The corresponding formula would be:

$$\frac{\text{object height}}{\text{depth of field}} \times 5$$

7.3.1.2 Rules for the Best Results with your System Setup

1. **Cover the whole distance range.**
2. **Let the depth of field areas overlap** (see [figure 7.6](#)).
3. **The distance between the light source and the object should remain constant.**
4. **The optical system should enable an orthographic projection (parallel projection)**, i.e., use a telecentric lens or a microscope in your system setup. If this is impossible, please read [section 7.6](#) on page [171](#).
5. **The axis of the focusing displacement must be parallel to the optical axis of the lens**, otherwise the object shifts laterally within the focus sequence.
6. **Do not move the camera or object in x or y direction.** DFF does not work on an object that is, e.g., moving on a conveyor belt between images or an object that is moved by any kind of vibration or agitation.

Note that the quality of your measuring results does depend on the quality of the input image and you should therefore aim to achieve the highest possible quality.

7.3.2 Practical Use of Depth from Focus

7.3.2.1 Example Application: Inspecting a PCB with DFF

This section describes an example application where the task is to test if a PCB board is covered by an appropriate amount of soldering paste. This inspection is performed with DFF and can be viewed in the HDevelop example program `hdevelop\Applications\Measuring-3D\measure_solder_paste_dff.hdev` (for more information also refer to the example `hdevelop\Applications\Measuring-3D\measure_bga_dff.hdev` which is a similar program showing how a ball grid array is measured with DFF). An image sequence is acquired and a height map of the single circuits and pads is calculated. This way parts that have no soldering paste can be identified as well as those that are covered sufficiently.

First all necessary images for the focus series have to be acquired and combined to a multi-channel image with the operator `channels_to_image`.

```
read_image (ImageArray, 'dff/focus_pcb_solder_paste_'+Sequence$'02')
channels_to_image (ImageArray, Image)
```

Then, depth from focus is performed, a depth map is calculated and all sharp gray values are selected. Both results can be used for further processing. Depth from focus is performed with the HALCON operator `depth_from_focus`. The channel number is returned for each pixel together with a confidence value, which is an indicator for the quality of the distance value. Pixels with the best focus are chosen. The method can be selected with the parameters `Filter` and `Selection`.

It is striking that the images are very dark, when looking at the images in the example `hdevelop\Applications\Measuring-3D\measure_solder_paste_dff.hdev`. They have been acquired like this on purpose to avoid overexposure. However, to improve the visibility of the object's surface in the image, the operator `scale_image` enhances the sharpness, `scale_image_max` spreads the gray values in the image to improve visibility despite of the darkness. For more information on depth from focus and overexposure please refer to [section 7.2.2.2](#) on page 164. `median_rect` suppresses unwanted outliers.

```
depth_from_focus (Image, Depth, Confidence, 'bandpass', 'next_maximum')
select_grayvalues_from_channels (Image, Depth, SharpenedImage)
scale_image (SharpenedImage, ImageScaled, 4, 0)
scale_image_max (Depth, ImageScaleMax)
median_rect (ImageScaleMax, DepthMean, 25, 25)
```

Finally, the results - the sharp image as well as the 3D plot - are displayed as shown in [figure 7.10](#). A sharp image is reconstructed by selecting the gray value of each pixel that is in focus for each coordinate using the depth image as index table. The focus stack and the depth image are used as input to reconstruct a focused image using the parameters `MultiChannelImage`, `IndexImage` and `Selected`. Sharp gray values can be identified by high frequencies, i.e., high edge amplitudes in the image, i.e., where the gray-value information changes quickly. For each sharp pixel, a confidence score is returned. The amount of sharpness defines the score. Furthermore, a 3D plot of the object in the image is calculated which is helpful for detecting defects.




Figure 7.10: Results of the PCB solder inspection: a) a synthetic sharp image and b) a 3D plot of the object.

```
dev_open_window (0, Width*0.7+5, Width*0.7, Height*0.7, 'black', \
                WindowHandle3D)
dev_set_paint (['3d_plot', 'texture'])
compose2 (DepthMean, ImageScaled, MultiChannelImage)
dev_display (MultiChannelImage)
```

7.3.3 Volume Measurement with Depth from Focus

In contrast to stereo, the height information is not calibrated for depth from focus. The values in the height image are indices of input images. To measure a real world height or volume, the distance in-between these images must be known. The easiest case is that images are taken with the same movement z . If two coordinates differ by an index value of n , the real world distance will be $z \times n$, with the unit of z . The unit in x and y , i.e., the size of a pixel must be known. Make sure that all dimensions are given in the same unit. Volume can be determined with the operator `area_center_gray` by adding up the pixel values which are equal to the height values. The resulting value must be multiplied by x , y and z .

7.4 Solutions for Typical Problems With DFF

There are two main problems that occur with depth from focus: overexposure and reflections.

1. **Overexposure leads to extreme values** (255) and results in peaks within the constant gray-value range of an image. This means that even within parts of the images that are blurry, the difference between neighboring gray values is so dominant that false edges might be detected (this effect is visualized in [figure 7.9](#) on page 165). Adjusting your illumination to minimize reflections can

reduce overexposure For more information on suitable illumination and handling overexposure, please refer to [section 7.2.2](#) on page 163.

2. **Object surfaces may reflect too much** for measuring with depth from focus. One advantage of depth from focus is, however, that you are very close to the object. Therefore, some objects that seem to have a reflecting surfaces might under the microscope show some surface texture after all. Try to minimize reflections with diffuse illumination as described in [section 7.2.2](#) on page 163. Furthermore, aberration occurs for DFF especially when images were taken with a standard lens but it also cannot totally be avoided for telecentric lenses either.

7.4.1 Calibrating Aberration

7.4.1.1 Aberration

Aberration is the effect that, when looking perpendicularly on a planar surface, not all pixels are in focus at the same time. Either the center or the outer part of the image are completely in focus. This effect is illustrated in [figure 7.11](#).




Figure 7.11: Image a) is an example for aberration. It is sharp in the middle and pixels become more blurry the further away they are from the center. Image b) shows a graphic that clarifies what happens when aberration occurs.

Aberration is the curvature of field that effects images taken with cameras using standard lenses but also cannot be completely avoided when using cameras with telecentric lenses. It results in an error in the depth image.

There are two main kinds of optical aberration:

1. **Spherical aberration** occurs when light rays have different focus points depending on their distance to the optical axis, i.e., the center part and the border of the image are not simultaneously sharp, even though a planar object is imaged.

2. **Coma** is the asymmetric accumulation of light intensity for off-axis points. It affects the periphery of the field of view. This kind of aberration is, however, not relevant here and is just mentioned for the sake of completeness.

The following paragraph describes how to calibrate spherical aberration and therefore avoid errors.

7.4.1.2 Setup Aberration Calibration

In order to get an accurate result from your depth from focus application, aberration has to be calibrated. This enables the calculation of correct sharp depth images and also enable further processing of the DFF results.

Rules for setting up your calibration to correct aberration.

1. **A planar surface with reasonable texture is needed** as reference plane.
2. **The camera has to be mounted perpendicular to the surface** (angles can be determined with HALCON's camera calibration if necessary).
3. **It is important that the same distance and camera setting are used** as will be applied during application!
4. **Using depth from focus, the “curvature on the surface”, i.e., the aberration is determined.**
5. **The extracted reference surface is used to correct the later measurements.**
6. **It is recommended to store the reference image to file** for further use.
7. **It is recommended to store the used parameters to file** for further use.
8. **The aberration can be approximated by a paraboloidal function.**
9. **The approximation has the advantage of reducing noise effects** that influence the measurements.
10. **The parameters of the paraboloidal function can be used to generate a reference image.**
11. **By subtracting the reference image from the depth measurement the error caused by the aberration can be corrected.**

7.5 Special Cases

Depth from focus may not always be used to reconstruct the three-dimensional surface of an object in order to subsequently measure it.

It may also be used to obtain a sharp image of the object if this is not otherwise possible due to the setup and continue working with this image.

Another possibility of using the DFF method is simply checking whether an object is present or not. Therefore, the lens is focused on a certain depth. If sharp pixels are found within an image at this certain depth, the object is present, otherwise it is missing. In this case, only one image can be sufficient.

Similarly, it can be used to check whether an object is tilted or not.

7.6 Performing Depth from Focus with a Standard Lens

Even though not recommended, it is possible to perform depth from focus with standard lenses. Using a non-telecentric lens does, however, require some adaptations of the DFF measuring process to ensure the best results.

Note that the accuracy of DFF measurements with a standard lens instead of a telecentric lens is reduced. For DFF with standard lenses, the focal length needs to be as long as possible in order to keep the perspective shift of the points in the field of view small. Furthermore, the depth of field needs to be small.

If the points shift too much within the field of view, measuring with DFF produces an erroneous result.

Calibrate aberration as described in [section 7.4.1](#) on page 169.

Note that the rules that have been described for DFF with a telecentric lens are also valid for DFF with a standard lens.

Chapter 8

Robot Vision

A typical application area for 3D vision is robot vision, i.e., whenever robots are equipped with cameras that supply information about the parts to manufacture. Such systems are also called “hand-eye systems” because the robotic “hand” is guided by mechanical “eyes”.

In order to use the information extracted by the camera, it must be transformed into the coordinate system of the robot. Thus, besides calibrating the camera(s) you must also calibrate the hand-eye system, i.e., determine the transformation between camera and robot coordinates. The following sections explain how to perform this hand-eye calibration with HALCON.




Figure 8.1: Robot vision scenarios: (a) moving camera, (b) stationary camera.

Please note that in order to use HALCON’s hand-eye calibration, **the camera must observe the workspace of the robot.** Figure 8.1 depicts the two possible scenarios: The camera can either be mounted on the robot and is moved to different positions by it, or it can be stationary. If the camera does not observe the workspace of the robot, e.g., if the camera observes parts on a conveyor belt, which are then handled by a robot further down the line, you must determine the relative pose of robot and camera with different means.

The calibration result can be used for different tasks. Typically, the results of machine vision, e.g., the position of a part, are to be transformed from camera into robot coordinates to create the appropriate robot commands, e.g., to grasp the part. Section 8.4 on page 180 describes such an application. Another possible application for hand-eye systems with a moving camera is to transform the information extracted from different camera poses into a common coordinate system.






Figure 8.2: Chain of transformations for a moving camera system.

Note that HALCON’s hand-eye calibration is not restricted to systems with a “hand”, i.e., a manipulator. You can also use it to calibrate cameras mounted on a pan-tilt head or surveillance cameras that rotate to observe a large area. Both systems correspond to a camera mounted on a robot; the calibration then allows you to accumulate visual information from different camera poses.

Further note that, although in the following only systems with a single camera are described, you can of course also use a stereo camera system. In this case, you typically calibrate only the relation of the robot to one of the cameras, because the relation between the cameras is determined by the stereo calibration (see [section 5.2.4](#) on page 113).

8.1 The Principle of Hand-Eye Calibration

Like the camera calibration (see [section 3.2](#) on page 40), the hand-eye calibration is based on providing multiple images of a known calibration object. But in contrast to the camera calibration, here the calibration object is not moved manually but by the robot, which moves either the calibration object in front of a stationary camera or the camera over a stationary calibration object. The pose, i.e., the position and orientation, of the robot in each calibration image **must be known with high accuracy!**

This results in a chain of coordinate transformations (see [figure 8.2](#) and [figure 8.3](#)). In this chain, two transformations (poses) are known: the robot pose ${}^{base}\mathbf{H}_{tool}$ and the pose of the calibration object in camera coordinates ${}^{cam}\mathbf{H}_{cal}$, which is determined from the calibration images before starting the hand-eye calibration. The hand-eye calibration then estimates the other two poses, i.e., the relation between the robot and the camera and between the robot and the calibration object, respectively.

Note that the chain consists of different poses depending on the used scenario. For a *moving camera*, the pose of the robot tool in camera coordinates and the pose of the calibration object in robot base coordinates are determined (see [figure 8.2](#)):

$${}^{cam}\mathbf{H}_{cal} = {}^{cam}\mathbf{H}_{tool} \cdot {}^{tool}\mathbf{H}_{base} \cdot {}^{base}\mathbf{H}_{cal} \quad (8.1)$$

Please beware that in this chain the inverse robot pose, i.e., the pose of the robot base in tool coordinates, is used.




Figure 8.3: Chain of transformations for a stationary camera system.

For a *stationary camera*, the pose of the robot base in camera coordinates and of the calibration object in robot tool coordinates are determined (see [figure 8.3](#)):

$${}^{cam}\mathbf{H}_{cal} = {}^{cam}\mathbf{H}_{base} \cdot {}^{base}\mathbf{H}_{tool} \cdot {}^{tool}\mathbf{H}_{cal} \quad (8.2)$$

The hand-eye calibration is performed with a single call to the operator [hand_eye_calibration](#).

```
hand_eye_calibration (NX, NY, NZ, NRow, NCol, NumPoints, RobotPoses, \
                    CamParam, 'nonlinear', 'error_pose', CameraPose, \
                    CalibrationPose, PoseError)
```

Let's have a brief look at the parameters; the referenced sections contain more detailed information:

- [X, Y, Z, Row, Col, NumPoints](#) (see [section 8.2.1](#))

As for the camera calibration, you must supply 3D model points and their corresponding image points. Note, however, that for the hand-eye calibration the **3D point coordinates must be supplied for each image**, together with the number of 3D points visible in each image. This requirement might seem tedious if you use the standard calibration plate, because then the same points are visible in each image, but it offers more flexibility for users of other calibration objects. !

- [RobotPoses](#) (see [section 8.2.2](#) on page 177)

This parameter contains the poses of the robot in each calibration image.

- [CamParam](#)

In this parameter you pass the internal camera parameters. In the HDevelop example programs `solution_guide\3d_vision\handeye_movingcam_calibration.hdev` and `solution_guide\3d_vision\handeye_stationarycam_calibration.hdev`, the camera is calibrated using the calibration images acquired for the hand-eye calibration. For detailed information about obtaining the internal camera parameters please refer to [section 3.2.7](#) on page 53.

- [Method, QualityType](#)

These parameters allow you to choose between different methods for calibration and result assessment.

- [CameraPose](#), [CalibrationPose](#), [Quality](#)

These are the calibration results and their quality assessment. How to use them in robot vision applications is described in [section 8.4](#) on page 180.

Besides the coordinate systems described above, two others may be of interest in a robot vision application: First, sometimes results must be transformed into a reference (world) coordinate system. You can define such a coordinate system easily based on a calibration image. Secondly, especially if the robot system uses different tools (grippers), it might be useful to place the tool coordinate system used in the calibration at the mounting point of the tools and introduce additional coordinate systems at the gripper (tool center point). The example application in [section 8.4](#) on page 180 shows how to handle both cases.

8.2 Determining Suitable Input Parameters

Below, we show how to determine values for the input parameters of [hand_eye_calibration](#). The code examples stem from the HDevelop programs `solution_guide\3d_vision\handeye_movingcam_calibration.hdev` and `solution_guide\3d_vision\handeye_stationarycam_calibration.hdev`, which perform the calibration of hand-eye systems with a moving and stationary camera, respectively. The programs stop after processing each calibration image; press Run to continue.

8.2.1 Corresponding World and Image Points

As for the camera calibration, you must supply 3D model points and their corresponding image points (compare [section 3.2.4](#) on page 49). First, we create empty tuples to accumulate data from all calibration images.

```
NRow := []
NCol := []
NX := []
NY := []
NZ := []
NumPoints := []
```

When using the standard calibration plate, the 3D model points, i.e., the 3D coordinates of the calibration marks, can be read from the description file.

```
caltab_points (CalTabFile, X, Y, Z)
```



In each calibration image, we then locate the calibration plate and extract the image coordinates of the calibration marks. Please note that for the hand-eye calibration **we strongly recommend to use the asymmetric calibration plate** introduced with HALCON 7.1 (see [section 3.2.3.1](#) on page 46). If even only in a single calibration image the pose of the old, symmetric calibration plate is estimated wrongly because it is rotated by more than 90 degrees, the calibration will fail!


```

for i := 0 to NumImages-1 by 1
  read_image (Image, ImageNameStart+i$'02d')
  find_caltab (Image, Caltab, CalTabFile, SizeGauss, MarkThresh, \
              MinDiamMarks)
  find_marks_and_pose (Image, Caltab, CalTabFile, StartCamParam, \
                      StartThresh, DeltaThresh, MinThresh, Alpha, \
                      MinContLength, MaxDiamMarks, RCoord, CCoord, \
                      StartPose)

```

Finally, the corresponding coordinates are accumulated in the tuples.

```

NRow := [NRow, RCoord]
NCol := [NCol, CCoord]
NX := [NX, X]
NY := [NY, Y]
NZ := [NZ, Z]
NumPoints := [NumPoints, 49]
endfor

```

Note that the 3D coordinates and number of marks per image are accumulated even if they don't change between images. As already explained, the possibility to use different model points in each image is not necessary when using the standard calibration plate, but can be very useful if you use your own calibration object, especially if it is a three-dimensional one.

8.2.2 Robot Poses

For each of the calibration images, you must specify the corresponding pose of the robot. Note that **the accuracy of the poses is critical to obtain an accurate hand-eye calibration**. There are two ways to “feed” the poses into HALCON: In many cases, you will simply read them from the robot control unit and then enter them into your HALCON program manually. For this, you can use the HDevelop example program `solution_guide\3d_vision\handeye_create_robot_poses.hdev`, which lets you input the poses in a text window and writes them into files.

As an alternative, if the robot has a serial or socket interface, you can also send them via this connection to your HALCON program (see the sections “[System > Serial](#)” and “[System > Sockets](#)” in the Reference Manual for more information).

In both cases, you then convert the data into HALCON 3D poses using the operator `create_pose`. As described in [section 2.1.5](#) on page 25 (and in the Reference Manual entry for `create_pose`), you can specify a pose in more than one way, because the orientation can be described by different sequences of rotations. Therefore, you must first check which sequence is used by your robot system. In many cases, it will correspond to

$$\mathbf{R}_{abg} = \mathbf{R}_z(\text{RotZ}) \cdot \mathbf{R}_y(\text{RotY}) \cdot \mathbf{R}_x(\text{RotX}) \quad (8.3)$$

If this is the case, select the value 'abg' for the parameter `OrderOfRotation` of `create_pose`. For the inverse order, select 'gba'.

If your robot system uses yet another sequence, you cannot use `create_pose` but must create a corresponding homogeneous transformation matrix and convert it into a pose using `hom_mat3d_to_pose`.



If, e.g., your robot system uses the following sequence of rotations where the rotations are performed around the z-axis, then around the y-axis, and finally again around the z-axis

$$\mathbf{R}_{zyz} = \mathbf{R}_z(Rl) \cdot \mathbf{R}_y(Rm) \cdot \mathbf{R}_z(Rr) \quad (8.4)$$

the pose can be created with the following code:

```
hom_mat3d_identity (HomMat3DIdentity)
hom_mat3d_translate (HomMat3DIdentity, Tx, Ty, Tz, HomMat3DTranslate)
hom_mat3d_rotate_local (HomMat3DTranslate, rad(Rl), 'z', HomMat3DT_Rl)
hom_mat3d_rotate_local (HomMat3DT_Rl, rad(Rm), 'y', HomMat3DT_Rl_Rm)
hom_mat3d_rotate_local (HomMat3DT_Rl_Rm, rad(Rr), 'z', HomMat3D)
hom_mat3d_to_pose (HomMat3D, Pose)
```

Note that the rotation operators expect angles to be given in radians, whereas `create_pose` expects them in degrees!

The example program `solution_guide\3d_vision\handeye_create_robot_poses.hdev` allows you to enter poses of the three types described above. If your robot system uses yet another sequence of rotations, you can easily extend the program by modifying (or copying and adapting) the code for ZYZ poses.

The HDevelop example programs `solution_guide\3d_vision\handeye_movingcam_calibration.hdev` and `solution_guide\3d_vision\handeye_stationarycam_calibration.hdev` read the robot pose files in the loop of processing the calibration images. Before this, an empty tuple is created to accumulate the poses.

```
RobotPoses := []
```

For each calibration image, the pose of the robot is read from file using `read_pose` and accumulated in the tuple `RobotPoses`.

```
read_pose (DataNameStart+'robot_pose_'+i$'02d'+'.dat', TmpRobotPose)
RobotPoses := [RobotPoses, TmpRobotPose]
```

If you are using a hand-eye system with a moving camera, you must invert the pose (compare the chain of transformations in [figure 8.2](#) on page 174).

```
hom_mat3d_invert (base_H_tool, tool_H_base)
hom_mat3d_to_pose (tool_H_base, RobotPoseInverse)
RobotPoses := [RobotPoses, RobotPoseInverse]
```

8.3 Performing the Calibration

Similar to the camera calibration, the main effort lies in collecting the input data. The calibration itself is performed with a single operator call.

```
hand_eye_calibration (NX, NY, NZ, NRow, NCol, NumPoints, RobotPoses, \
                    CamParam, 'nonlinear', 'error_pose', CameraPose, \
                    CalibrationPose, PoseError)
```

Typically, you then save the calibrated poses in files so that your robot vision application can read them at a later time. The following code does so for a system with a moving camera.

```
write_pose (CameraPose, DataNameStart+'final_pose_cam_tool.dat')
write_pose (CalibrationPose, DataNameStart+'final_pose_base_calplate.dat')
```

Of course, you should check whether the calibration was successful by looking at the output parameter [Quality](#), which is a measure for the accuracy of the pose parameters. With the parameter [QualityType](#), you can specify the type of quality measures (note that not all types are possible for all calibration methods selected with the parameter [Method](#)):

'error_pose' returns the pose error of the complete chain of transformations in form of a tuple with four elements:

- the root-mean-square error of the translational part
- the root-mean-square error of the rotational part
- the maximum translational error
- the maximum rotational error

'standard_deviation' returns the standard deviations of the two poses in a tuple with 12 elements. The first six elements refer to the camera pose and the others to the pose of the calibration points.

'covariance' returns the full 12x12 covariance matrix of both poses.

In the example, the pose error values are returned and displayed.

```
disp_message (WindowHandle, \
             'Quality of the results:  root mean square    maximum', \
             'window', 40, -1, 'black', 'true')
disp_message (WindowHandle, \
             ' translation part:      ' + PoseError[0] + '      ' + PoseError[2], \
             'window', 60, -1, 'black', 'true')
disp_message (WindowHandle, \
             ' rotation part:        ' + PoseError[1] + '      ' + PoseError[3], \
             'window', 80, -1, 'black', 'true')
```

The example programs then visualize the calibrated poses by displaying the coordinate system of the calibration plate in each calibration image. For this, they compute the pose of the calibration plate in camera coordinates based on the calibrated poses. For a moving camera system, this corresponds to the following code (compare [equation 8.1](#) on page 174).

```

* CalplatePose = cam_H_calplate = cam_H_tool * tool_H_base * \
* base_H_calplate
pose_to_hom_mat3d (CalibrationPose, base_H_calplate)
pose_to_hom_mat3d (CameraPose, cam_H_tool)
pose_to_hom_mat3d (RobotPoseInverse, tool_H_base)
hom_mat3d_compose (cam_H_tool, tool_H_base, cam_H_base)
hom_mat3d_compose (cam_H_base, base_H_calplate, cam_H_calplate)
hom_mat3d_to_pose (cam_H_calplate, CalplatePose)

```

This code is encapsulated in a procedure, which is called in a loop over all images.

```

for i := 0 to NumImages-1 by 1
  read_image (Image, ImageNameStart+i*'02d')
  TmpRobotPoseInverse := RobotPoses[i*7:i*7+6]
  calc_calplate_pose_movingcam (CalibrationPose, CameraPose, \
                                TmpRobotPoseInverse, TmpCalplatePose)
  display_calplate_coordinate_system (CalTabFile, TmpCalplatePose, \
                                      CamParam, WindowHandle)
endfor

```

The corresponding procedure for a stationary camera system is listed in [appendix B.9](#) on page 236.

8.4 Using the Calibration Data

Typically, you use the result of the hand-eye calibration to **transform the results of machine vision from camera coordinates into robot coordinates** (${}^{cam}\mathbf{H}_{obj} \rightarrow {}^{base}\mathbf{H}_{obj}$) to generate the appropriate robot commands, e.g., to grasp an object whose position has been determined in an image as in the application described in [section 8.4.3](#) on page 182. For a stationary camera, this transformation corresponds to the following equation (compare [figure 8.3](#) on page 175):

$${}^{base}\mathbf{H}_{obj} = {}^{base}\mathbf{H}_{cam} \cdot {}^{cam}\mathbf{H}_{obj} = (\mathbf{CameraPose})^{-1} \cdot {}^{cam}\mathbf{H}_{obj} \quad (8.5)$$

For a moving camera system, the equation also contains the pose of the robot when acquiring the image of the object ${}^{base}\mathbf{H}_{tool}(\text{acq. pos.})$ (compare [figure 8.2](#) on page 174):

$$\begin{aligned} {}^{base}\mathbf{H}_{obj} &= {}^{base}\mathbf{H}_{tool}(\text{acq. pos.}) \cdot {}^{tool}\mathbf{H}_{cam} \cdot {}^{cam}\mathbf{H}_{obj} \\ &= {}^{base}\mathbf{H}_{tool}(\text{acq. pos.}) \cdot (\mathbf{CameraPose})^{-1} \cdot {}^{cam}\mathbf{H}_{obj} \end{aligned} \quad (8.6)$$

8.4.1 Using the Hand-Eye Calibration for Grasping (3D Alignment)

Grasping an object corresponds to a very simple equation that says “move the robot gripper to the pose of the object” (“grasping pose”). This is also called 3D alignment.

Note that if the tool coordinate system used during hand-eye calibration is not placed at the gripper (tool center point), the equation also contains the transformation between tool and gripper coordinate system.

This transformation cannot be calibrated with the hand-eye calibration but must be measured or taken from the CAD model or technical drawing of the gripper.

$$\begin{aligned} \text{tool} = \text{gripper:} \quad & {}^{base}\mathbf{H}_{tool}(\text{gr. pos.}) = {}^{base}\mathbf{H}_{obj} \\ \text{tool} \neq \text{gripper:} \quad & {}^{base}\mathbf{H}_{tool}(\text{gr. pos.}) \cdot {}^{tool}\mathbf{H}_{gripper} = {}^{base}\mathbf{H}_{obj} \\ & {}^{base}\mathbf{H}_{tool}(\text{gr. pos.}) = {}^{base}\mathbf{H}_{obj} \cdot ({}^{tool}\mathbf{H}_{gripper})^{-1} \end{aligned} \quad (8.7)$$

If we replace ${}^{base}\mathbf{H}_{obj}$ according to [equation 8.5](#) on page 180 and [equation 8.6](#) on page 180, we get the “grasping equation” for a **stationary camera**

$${}^{base}\mathbf{H}_{tool}(\text{gr. pos.}) = (\text{CameraPose})^{-1} \cdot {}^{cam}\mathbf{H}_{obj} \left[({}^{tool}\mathbf{H}_{gripper})^{-1} \right] \quad (8.8)$$

and for a **moving camera**

$${}^{base}\mathbf{H}_{tool}(\text{gr. pos.}) = {}^{base}\mathbf{H}_{tool}(\text{acq. pos.}) \cdot (\text{CameraPose})^{-1} \cdot {}^{cam}\mathbf{H}_{obj} \left[({}^{tool}\mathbf{H}_{gripper})^{-1} \right] \quad (8.9)$$

The notation $\left[({}^{tool}\mathbf{H}_{gripper})^{-1} \right]$ indicates that this part is only necessary if the tool coordinate system is not identical with the gripper coordinate system.

8.4.2 How to Get the 3D Pose of the Object

The 3D pose of the object in camera coordinates (${}^{cam}\mathbf{H}_{obj}$) can stem from different sources:

- With a **binocular stereo system**, you can determine the 3D pose of unknown objects directly (see [chapter 5](#) on page 105).
- For single camera systems, HALCON provides multiple methods. The most powerful one is **shape-based 3D Matching** (see [section 4.2](#) on page 81 or the Solution Guide I, [chapter 10](#) on page 127), which performs a full object recognition, i.e., it not only estimates the pose but first locates the object in the image. If only one, planar side of the object is visible, fast alternatives are the calibrated perspective deformable matching ([section 4.5](#) on page 101) and the calibrated descriptor-based matching ([section 4.6](#) on page 101).
- If a full object recognition is not necessary, you can use **pose estimation** to determine the 3D pose of known objects (see [section 8.4.3](#) for an example application and [chapter 4](#) on page 75 for more details on pose estimation).
- Finally, you can determine the 3D coordinates of unknown objects if **object points lie in a known plane** (see [section 8.4.3](#) for an example application and [section 3.3](#) on page 57 for more details on determining 3D points in a known plane).

Please note that if you want to use the 3D pose for grasping the object, the extracted pose, in particular the orientation, must be identical to the pose of the gripper at the grasping position.

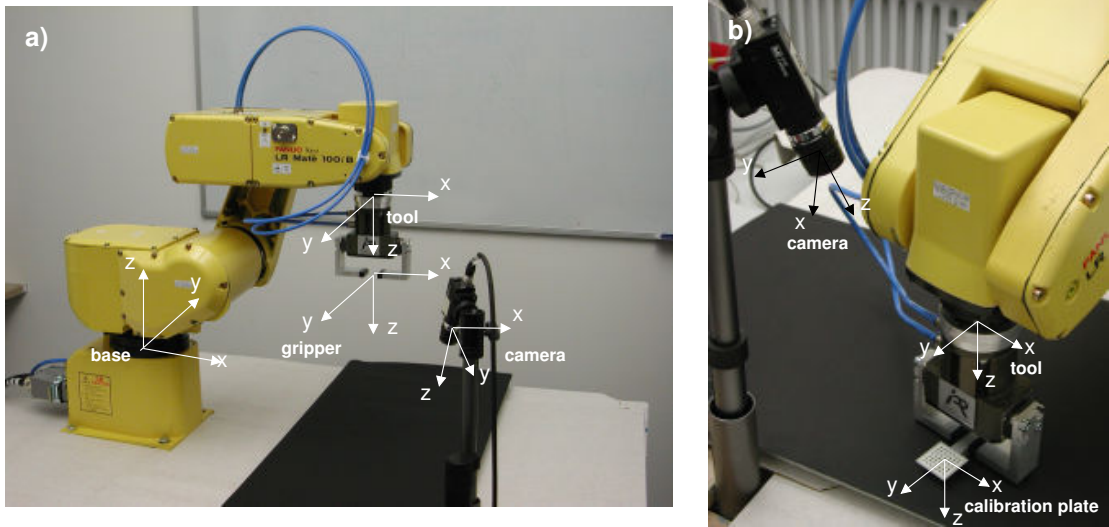


Figure 8.4: Example hand-eye system with a stationary camera: coordinate systems (a) of robot and camera, (b) with calibration plate.

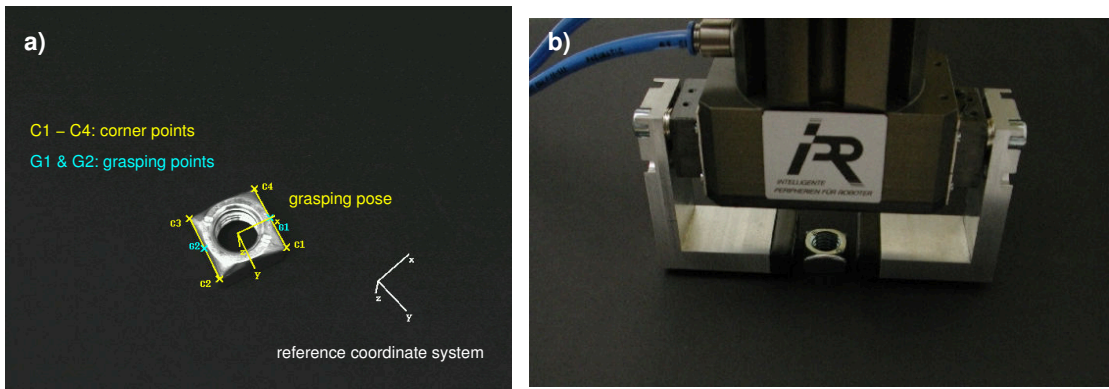


Figure 8.5: (a) Determining the 3D pose for grasping a nut; (b) robot at grasping pose.

8.4.3 Example Application with a Stationary Camera: Grasping a Nut

This section describes an example application realized with the hand-eye system depicted in [figure 8.4](#). The task is to localize a nut and determine a suitable grasping pose for the robot (see [figure 8.5](#)). The HDevelop example program `solution_guide\3d_vision\handeye_stationarycam_grasp_nut.hdev` performs the machine vision part and transforms the resulting pose into robot coordinates using the calibration data determined with `solution_guide\3d_vision\handeye_stationarycam_calibration.hdev` as described in the previous sections. As you will see, using the calibration data is the shortest part of the program, its main part is devoted to machine vision.

Step 1: Read calibration data

First, the calibration data is read from files; for later computations, the poses are converted into homogeneous transformation matrices.

```
read_cam_par (DataNameStart+'final_campar.dat', CamParam)
read_pose (DataNameStart+'final_pose_cam_base.dat', PoseCamBase)
pose_to_hom_mat3d (PoseCamBase, cam_H_base)
read_pose (DataNameStart+'final_pose_tool_calplate.dat', PoseToolCalplate)
pose_to_hom_mat3d (PoseToolCalplate, tool_H_calplate)
```

In the used hand-eye system, the tool coordinate system used in the calibration process is located at the mounting point of the tool; therefore, an additional coordinate system is needed between the fingers of the gripper (see [figure 8.4a](#) on page 182). Its pose in tool coordinates is also read from file.

```
read_pose (DataNameStart+'pose_tool_gripper.dat', PoseToolGripper)
pose_to_hom_mat3d (PoseToolGripper, tool_H_gripper)
```

Step 2: Define reference coordinate system

Now, a reference coordinate system is defined based on one of the calibration images. In this image, the calibration plate has been placed into the plane on top of the nut. This allows to determine the 3D coordinates of extracted image points on the nut with a single camera and without knowing the dimensions of the nut. The code for defining the reference coordinate system is contained in a procedure, which is listed in [appendix B.10](#) on page 237.

```
define_reference_coord_system (ImageNameStart+'calib3cm_00', CamParam, \
                               CalplateFile, WindowHandle, PoseRef)
pose_to_hom_mat3d (PoseRef, cam_H_ref)
```

Step 3: Extract grasping points on the nut

The following code extracts grasping points on two opposite sides of the nut. The nut is found with simple blob analysis; its boundary is converted into XLD contours.

```
threshold (Image, BrightRegion, 60, 255)
connection (BrightRegion, BrightRegions)
select_shape (BrightRegions, Nut, 'area', 'and', 500, 99999)
fill_up (Nut, NutFilled)
gen_contour_region_xld (NutFilled, NutContours, 'border')
```

The contours are then processed to find long, parallel straight line segments; their corners are accumulated in tuples.

```

segment_contours_xld (NutContours, LineSegments, 'lines', 5, 4, 2)
fit_line_contour_xld (LineSegments, 'tukey', -1, 0, 5, 2, RowBegin, \
    ColBegin, RowEnd, ColEnd, Nr, Nc, Dist)
gen_empty_obj (Lines)
for i := 0 to |RowBegin| -1 by 1
    gen_contour_polygon_xld (Contour, [RowBegin[i],RowEnd[i]], [ColBegin[i], \
        ColEnd[i]])
    concat_obj (Lines, Contour, Lines)
endfor
gen_polygons_xld (Lines, Polygon, 'ramer', 2)
gen_parallels_xld (Polygon, ParallelLines, 50, 100, rad(10), 'true')
get_parallels_xld (ParallelLines, Row1, Col1, Length1, Phi1, Row2, Col2, \
    Length2, Phi2)
CornersRow := [Row1[0], Row1[1], Row2[0], Row2[1]]
CornersCol := [Col1[0], Col1[1], Col2[0], Col2[1]]

```

Step 4: Determine the grasping pose in camera coordinates

The grasping pose is calculated in 3D coordinates. For this, the 3D coordinates of the corner points in the reference coordinate system are determined using the operator [image_points_to_world_plane](#). The origin of the grasping pose lies in the middle of the corners.

```

image_points_to_world_plane (CamParam, PoseRef, CornersRow, CornersCol, 'm', \
    CornersX_ref, CornersY_ref)
CenterPointX_ref := sum(CornersX_ref)*0.25
CenterPointY_ref := sum(CornersY_ref)*0.25

```

The grasping pose is oriented almost like the reference coordinate system, only rotated around the z-axis so that it is identical to the gripper coordinate system, i.e., so that the gripper “fingers” are parallel to the sides of the nut. To calculate the rotation angle, first the grasping points in the middle of the sides are determined. Their angle can directly be used as the rotation angle.

```

GraspPointsX_ref := [(CornersX_ref[0]+CornersX_ref[1])*0.5, \
    (CornersX_ref[2]+CornersX_ref[3])*0.5]
GraspPointsY_ref := [(CornersY_ref[0]+CornersY_ref[1])*0.5, \
    (CornersY_ref[2]+CornersY_ref[3])*0.5]
GraspPhiZ_ref := atan((GraspPointsY_ref[1]-GraspPointsY_ref[0])/ \
    (GraspPointsX_ref[1]-GraspPointsX_ref[0]))

```

With the origin and rotation angle, the grasping pose is first determined in the reference coordinate system and then transformed into camera coordinates.

```

hom_mat3d_identity (HomMat3DIdentity)
hom_mat3d_rotate (HomMat3DIdentity, GraspPhiZ_ref, 'z', 0, 0, 0, \
    HomMat3D_RZ_Phi)
hom_mat3d_translate (HomMat3D_RZ_Phi, CenterPointX_ref, CenterPointY_ref, 0, \
    ref_H_grasp)
hom_mat3d_compose (cam_H_ref, ref_H_grasp, cam_H_grasp)

```


Alternatively, the example also shows how to calculate the grasping pose using pose estimation (see [chapter 4](#) on page 75 for a detailed description). This method can be used when points on the object are known. In the example, we specify the 3D coordinates of the corners of the nut.

```
NX := [0.009, -0.009, -0.009, 0.009]
NY := [0.009, 0.009, -0.009, -0.009]
```

The grasping pose is then calculated by simply calling the operator `vector_to_pose`. Before, however, the image coordinates of the corners must be sorted such that the first one lies close to the x-axis of the reference coordinate system. Otherwise, the orientation of the reference coordinate system would differ too much from the grasping pose and the pose estimation would fail.

```
sort_corner_points (CornersRow, CornersCol, WindowHandle, NRow, NCol)
vector_to_pose (NX, NY, NZ, NRow, NCol, CamParam, 'iterative', 'error', \
               PoseCamNut, Quality)
disp_3d_coord_system (WindowHandle, CamParam, PoseCamGripper, 0.01)
```

The result of both methods is displayed in [figure 8.5a](#) on page 182.

Step 5: Transform the grasping pose in robot coordinates

Now comes the moment to use the results of the hand-eye calibration: The grasping pose is transformed into robot coordinates with the formula shown in [equation 8.8](#) on page 181.

```
hom_mat3d_invert (cam_H_base, base_H_cam)
hom_mat3d_compose (base_H_cam, cam_H_grasp, base_H_grasp)
```

As already mentioned, the tool coordinate system used in the calibration process is placed at the mounting point of the tool, not between the fingers of the gripper. Thus, the pose of the tool in gripper coordinates must be added to the chain of transformations to obtain the pose of the tool in base coordinates.

```
hom_mat3d_invert (tool_H_gripper, gripper_H_tool)
hom_mat3d_compose (base_H_grasp, gripper_H_tool, base_H_tool)
```

Step 6: Transform pose type

Finally, the pose is converted into the type used by the robot controller.

```
hom_mat3d_to_pose (base_H_tool, PoseRobotGrasp)
convert_pose_type (PoseRobotGrasp, 'Rp+T', 'abg', 'point', \
                  PoseRobotGrasp_ZYX)
```

[Figure 8.5b](#) on page 182 shows the robot at the grasping pose.

Chapter 9

Calibrated Mosaicking

Some objects are too large to be covered by one single image. Multiple images that cover different parts of the object must be taken in such cases. You can measure precisely across the different images if the cameras are calibrated and their external parameters are known with respect to one common world coordinate system.

It is even possible to merge the individual images into one larger image that covers the whole object. This is done by rectifying the individual images with respect to the same measurement plane (see [section 3.4.1](#) on page 62). In the resulting image, you can measure directly in world coordinates.

Note that the 3D coordinates of objects are derived based on the same principle as described in [chapter 3](#) on page 37, i.e., a measurement plane that coincides with the object surface must be defined. Although two or more cameras are used, this is no stereo approach. For more information on 3D vision with a binocular stereo system, please refer to [chapter 5](#) on page 105.

If the resulting image is not intended to serve for high-precision measurements in world coordinates, you can generate it using the mosaicking approach described in [chapter 10](#) on page 199. With this approach, it is not necessary to calibrate the cameras.

A setup for generating a high-precision mosaic image from two cameras is shown in [figure 9.1](#). The cameras are mounted such that the resulting pair of images has a small overlap. The cameras are first calibrated and then the images are merged together into one larger image. All further explanations within this section refer to such a two-camera setup.

Typically, the following steps must be carried out:

1. Determination of the internal camera parameters for each camera separately.
2. Determination of the external camera parameters, using *one* calibration object, to facilitate that the relation between the cameras can be determined.
3. Merge the images into one larger image that covers the whole object.

The HDevelop example program `solution_guide\3d_vision\two_camera_calibration.hdev` shows how to determine the external camera parameters and how to merge two images into one larger image.

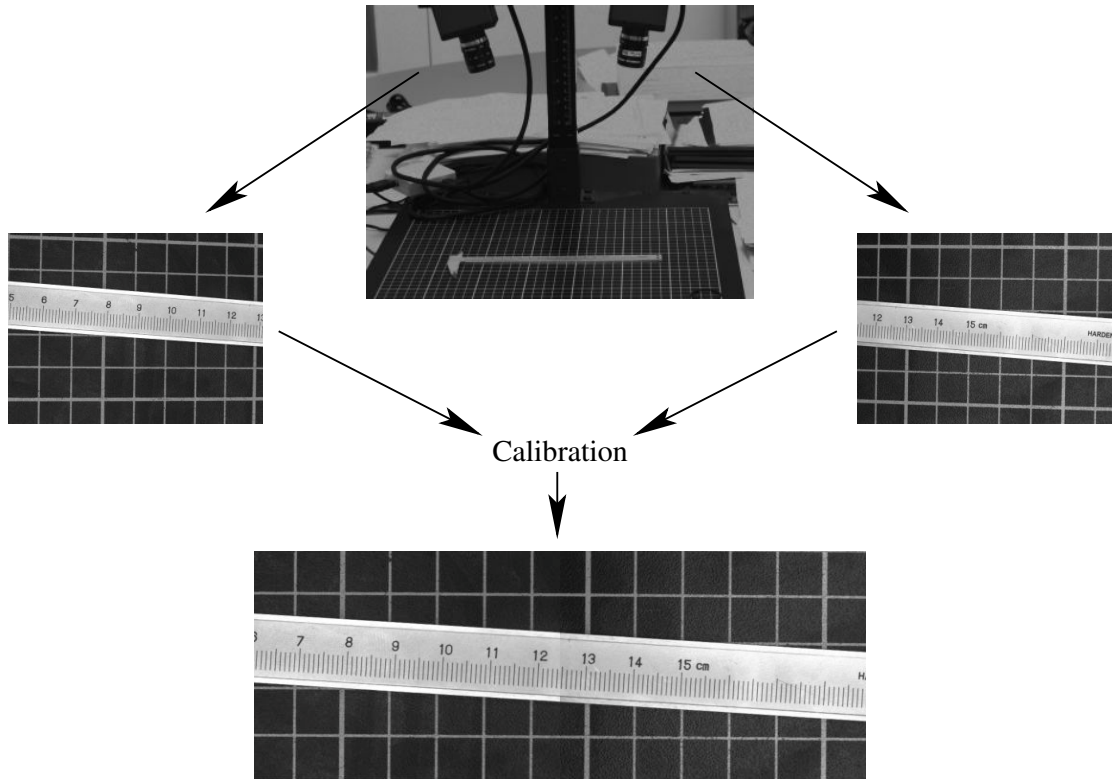


Figure 9.1: Two-camera setup.

9.1 Setup

Two or more cameras must be mounted on a *stable* platform such that each image covers a part of the whole scene. The cameras can have an arbitrary orientation, i.e., it is not necessary that they are looking parallel or perpendicular onto the object surface.

To setup focus, illumination, and overlap appropriately, use a big reference object that covers all fields of view. To permit that the images are merged into one larger image, they must have some overlap (see [figure 9.2](#) for an example). The overlapping area can be even smaller than depicted in [figure 9.2](#), since the overlap is only necessary to ensure that there are no gaps in the resulting combined image.

9.2 Calibration

The calibration of the images can be broken down into two separate steps.

The first step is to determine the internal camera parameters for each of the cameras in use. This can be done for each camera independently, as described in [section 3.2](#) on page 40.




Figure 9.2: Overlapping images.

The second step is to determine the external camera parameters for all cameras. Because the final coordinates should refer to *one* world coordinate system for all images, a big calibration object that appears in all images has to be used. We propose to use a calibration object like the one displayed in [figure 9.3](#), which consists of as many calibration plates as the number of cameras that are used.

For the determination of the external camera parameters, it is sufficient to use one calibration image from each camera only. Note that the calibration object must not be moved in between the acquisition of the individual images. Ideally, the images are acquired simultaneously. The pose of the calibration plates relative to the cameras is then computed with the operator `find_marks_and_pose`.

```
min_max_gray (Image1, Image1, 3, Min, Max, Range)
find_caltab (Image1, Caltab, CaltabName, 3, min([Min+40,200]), 5)
find_marks_and_pose (Image1, Image1, CaltabName, CamParam1, 128, 10, 18, \
                    0.9, 15, 100, RCoord1, CCoord1, Pose1)
min_max_gray (Image2, Image2, 3, Min, Max1, Range1)
find_caltab (Image2, Caltab, CaltabName, 3, min([Min+40,200]), 5)
find_marks_and_pose (Image2, Image2, CaltabName, CamParam2, 128, 10, 18, \
                    0.9, 15, 100, RCoord2, CCoord2, Pose2)
```

The calibration is easy if standard HALCON calibration plates mounted on some kind of carrier plate are used such that in each image one calibration plate is completely visible. An example for such a calibration object for a two-camera setup is given in [figure 9.3](#). The respective calibration images for the determination of the external camera parameters are shown in [figure 9.4](#). Note that the relative position of the calibration plates with respect to each other must be known precisely. This can be done with the pose estimation described in [chapter 4](#) on page 75.

Note also that only the relative position of the calibration marks among each other shows the high accuracy stated in [section 3.2.3.1](#) on page 46 but not the borders of the calibration plate. The rows of calibration marks may be slanted with respect to the border of the calibration plate and even the dis-

tance of the calibration marks from the border of the calibration plate may vary. Therefore, aligning the calibration plates along their boundaries may result in a shift in x - and y -direction with respect to the coordinate system of the calibration plate in its initial position.




Figure 9.3: Calibration object for two-camera setup.




Figure 9.4: Calibration images for two-camera setup.

9.3 Merging the Individual Images into One Larger Image

At first, the individual images must be rectified, i.e., transformed so that they exactly fit together. This can be achieved by using the operators `gen_image_to_world_plane_map` and `map_image`. Then, the mosaic image can be generated by the operator `tile_images`, which tiles multiple images into one larger image. These steps are visualized in [figure 9.5](#).

The operators `gen_image_to_world_plane_map` and `map_image` are described in [section 3.4.1](#) on page 62. In the following, we will only discuss the problem of defining the appropriate image detail, i.e., the position of the upper left corner and the size of the rectified images. Again, the description is based on the two-camera setup.

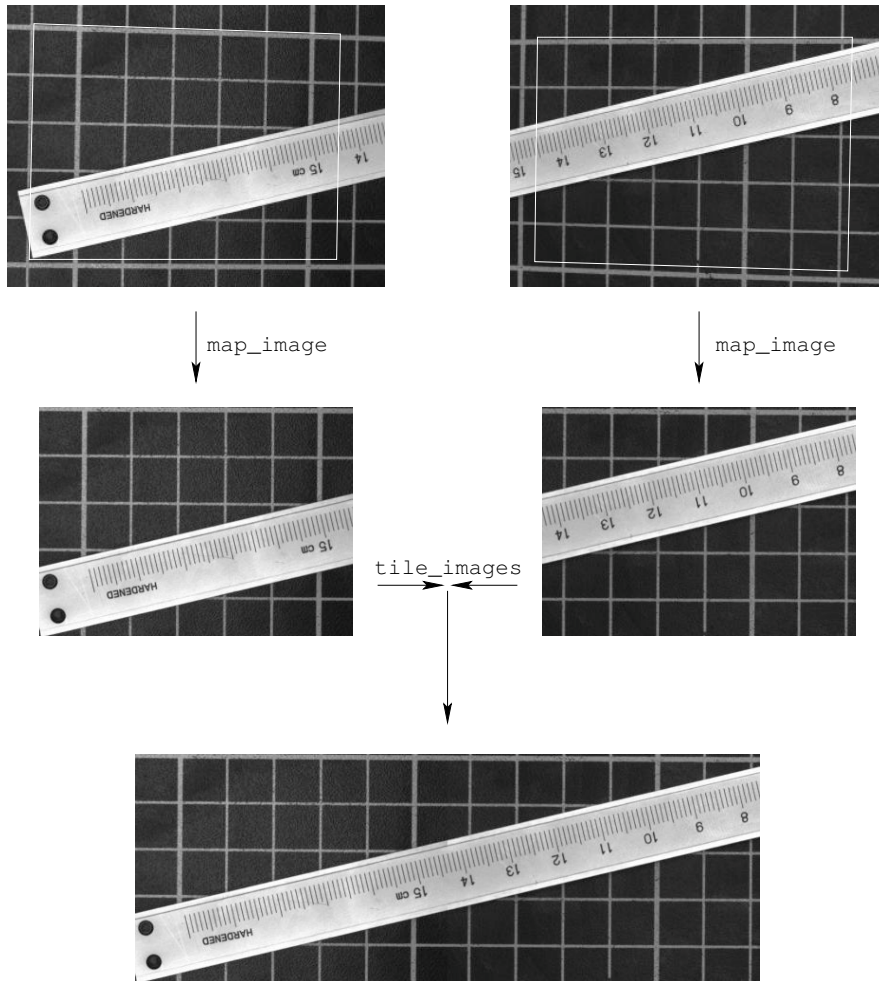


Figure 9.5: Image rectification and tiling.

9.3.1 Definition of the Rectification of the First Image

For the first (here: left) image, the determination of the necessary shift of the pose is straightforward. You can define the upper left corner of the rectified image in image coordinates, e.g., interactively or, as in the example program, based on a preselected border width.

```
ULRow := HeightImage1*BorderInPercent/100.0
ULCol := WidthImage1*BorderInPercent/100.0
```

Then, this point must be transformed into world coordinates.

```
image_points_to_world_plane (CamParam1, Pose1, ULRow, ULCol, 'm', ULX, ULY)
```

The resulting coordinates can be used directly, together with the shift that compensates the thickness of the calibration plate (see [section 3.2.7.1](#) on page 53) to modify the origin of the world coordinate system in the left image.

```
set_origin_pose (Pose1, ULX, ULY, DiffHeight, PoseNewOrigin1)
```

This means that we shift the origin of the world coordinate system from the center of the calibration plate to the position that defines the upper left corner of the rectified image ([figure 9.6](#)).




Figure 9.6: Definition of the upper left corner of the first rectified image.

The size of the rectified image, i.e., its width and height, can be determined from points originally defined in image coordinates, too. In addition, the desired pixel size of the rectified image must be specified.

```
PixelSize := 0.0001
```

For the determination of the height of the rectified image we need to define a point that lies near the lower border of the first image.

```
LowerRow := HeightImage1*(100-BorderInPercent)/100.0
```

Again, this point must be transformed into the world coordinate system.

```
image_points_to_world_plane (CamParam1, Pose1, LowerRow, ULCol, 'm', X1, \
                             LowerY)
```


The height can be determined as the vertical distance between the upper left point and the point near the lower image border, expressed in pixels of the rectified image.

```
HeightRect := int((LowerY-ULY)/PixelSize)
```

Analogously, the width can be determined from a point that lies in the overlapping area of the two images, i.e., near the right border of the first image.

```
RightCol := WidthImage1*(100-OverlapInPercent/2.0)/100.0
image_points_to_world_plane (CamParam1, Pose1, ULRow, RightCol, 'm', RightX, \
                             Y1)
WidthRect := int((RightX-ULX)/PixelSize)
```

Note that the above described definitions of the image points, from which the upper left corner and the size of the rectified image are derived, assume that the x- and y-axes of the world coordinate system are approximately aligned to the column- and row-axes of the first image. This can be achieved by placing the calibration plate in the first image approximately aligned with the image borders. Otherwise, the distances between the above mentioned points make no sense and the upper left corner and the size of the rectified image must be determined in a manner that is adapted for the configuration at hand.

With the shifted pose and the size of the rectified image, the rectification map for the first image can be derived.

```
gen_image_to_world_plane_map (MapSingle1, CamParam1, PoseNewOrigin1, Width, \
                              Height, WidthRect, HeightRect, PixelSize, \
                              'bilinear')
```

9.3.2 Definition of the Rectification of the Second Image

The second image must be rectified such that it fits exactly to the right of the first rectified image. This means that the upper left corner of the second rectified image must be identical with the upper right corner of the first rectified image. Therefore, we need to know the coordinates of the upper right corner of the first rectified image in the coordinate system that is defined by the calibration plate in the second image.

First, we express the upper right corner of the first rectified image in the world coordinate system that is defined by the calibration plate in the first image. It can be determined by a transformation from the origin into the upper left corner of the first rectified image (a translation in the example program) followed by a translation along the upper border of the first rectified image. Together with the shift that compensates the thickness of the calibration plate, this transformation is represented by the homogeneous transformation matrix ${}^{cp1}\mathbf{H}_{ur1}$ (see [figure 9.7](#)), which can be defined in HDevelop by:

```
hom_mat3d_translate_local (HomMat3DIdentity, ULX+PixelSize*WidthRect, ULY, \
                          DiffHeight, cp1Hur1)
```

Then, we need the transformation between the two calibration plates of the calibration object. The homogeneous transformation matrix ${}^{cp1}\mathbf{H}_{cp2}$ describes how the world coordinate system defined by




Figure 9.7: Definition of the upper right corner of the first rectified image.

the calibration plate in the first image is transformed into the world coordinate system defined by the calibration plate in the second image (figure 9.8). This transformation must be known beforehand from a precise measurement of the calibration object.

From these two transformations, it is easy to derive the transformation that transforms the world coordinate system of the second image such that its origin lies in the upper left corner of the second rectified image. For this, the two transformations have to be combined appropriately (see figure 9.9):

$${}^{cp2}\mathbf{H}_{ul2} = {}^{cp2}\mathbf{H}_{cp1} \cdot {}^{cp1}\mathbf{H}_{ur1} \quad (9.1)$$

$$= ({}^{cp1}\mathbf{H}_{cp2})^{-1} \cdot {}^{cp1}\mathbf{H}_{ur1} \quad (9.2)$$

This can be implemented in HDevelop as follows:

```
hom_mat3d_invert (cp1Hcp2, cp2Hcp1)
hom_mat3d_compose (cp2Hcp1, cp1Hur1, cp2Hul2)
```

With this, the pose of the calibration plate in the second image can be modified such that the origin of the world coordinate system lies in the upper left corner of the second rectified image.

```
pose_to_hom_mat3d (Pose2, cam2Hcp2)
hom_mat3d_compose (cam2Hcp2, cp2Hul2, cam2Hul2)
hom_mat3d_to_pose (cam2Hul2, PoseNewOrigin2)
```




Figure 9.8: Transformation between the two world coordinate systems, each defined by the respective calibration plate.




Figure 9.9: Definition of the upper left corner of the second rectified image.

With the resulting new pose and the size of the rectified image, which can be the same as for the first rectified image, the rectification map for the second image can be derived.

```
gen_image_to_world_plane_map (MapSingle2, CamParam2, PoseNewOrigin2, Width, \
                             Height, WidthRect, HeightRect, PixelSize, \
                             'bilinear')
```

9.3.3 Rectification of the Images

Once the rectification maps are created, every image pair from the two-camera setup can be rectified and tiled very efficiently. The resulting mosaic image consists of the two rectified images and covers a part as indicated in [figure 9.10](#).




Figure 9.10: The position of the final mosaic image.

The rectification is carried out by the operator `map_image`.

```
map_image (Image1, MapSingle1, RectifiedImage1)
map_image (Image2, MapSingle2, RectifiedImage2)
```

This transforms the two images displayed in [figure 9.11](#), into the two rectified images that are shown in [figure 9.12](#).

As a preparation for the tiling, the rectified images must be concatenated into one tuple, which then contains both images.

```
concat_obj (RectifiedImage1, RectifiedImage2, Concat)
```

Then the two images can be tiled.




Figure 9.11: Two test images acquired with the two-camera setup.




Figure 9.12: Rectified images.

```
tile_images (Concat, Combined, 2, 'vertical')
```

The resulting mosaic image is displayed in [figure 9.13](#).




Figure 9.13: Mosaic image.

Chapter 10

Uncalibrated Mosaicking

If you need an image of a large object, but the field of view of the camera does not allow to cover the entire object with the desired resolution, you can use image mosaicking to generate a large image of the entire object from a sequence of overlapping images of parts of the object.

An example for such an application is given in [figure 10.1](#). On the left side, six separate images are displayed stacked upon each other. On the right side, the mosaic image generated from the six separate images is shown. Note that the folds visible in the image do not result from the mosaicking. They are due to some degradations on the PCB, which can be seen already in the separate images.

The mosaicking approach described in this section is designed for applications where it is not necessary to achieve the high-precision mosaic images as described in [chapter 9](#) on page 187. The advantages compared to this approach are that no camera calibration is necessary and that the individual images can be arranged automatically.

The example program `solution_guide\3d_vision\mosaicking.hdev` generates the mosaic image displayed in [figure 10.7](#) on page 205. First, the images are read from file and collected in one tuple.

```
gen_empty_obj (Images)
for J := 1 to 10 by 1
    read_image (Image, ImgPath+ImgName+J$'02')
    concat_obj (Images, Image, Images)
endfor
```

Then, the image pairs must be defined, i.e., which image should be mapped to which image.

```
From := [1,2,3,4,6,7,8,9,3]
To := [2,3,4,5,7,8,9,10,8]
```

Now, characteristic points must be extracted from the images, which are then used for the matching between the image pairs. The resulting projective transformation matrices¹ must be accumulated.

¹A projective transformation matrix describes a perspective projection. It consists of 3×3 values. If the last row contains the values $[0,0,1]$, it corresponds to a homogeneous transformation matrix of HALCON and therefore describes an affine transformation.




Figure 10.1: A first example for image mosaicking.


```

Num := |From|
ProjMatrices := []
for J := 0 to Num-1 by 1
  F := From[J]
  T := To[J]
  select_obj (Images, ImageF, F)
  select_obj (Images, ImageT, T)
  points_harris (ImageF, SigmaGrad, SigmaSmooth, Alpha, Threshold, \
    RowFAll, ColFAll)
  points_harris (ImageT, SigmaGrad, SigmaSmooth, Alpha, Threshold, \
    RowTAll, ColTAll)
  proj_match_points_ransac (ImageF, ImageT, RowFAll, ColFAll, RowTAll, \
    ColTAll, 'sad', MaskSize, RowMove, ColMove, \
    RowTolerance, ColTolerance, Rotation, \
    MatchThreshold, 'gold_standard', \
    DistanceThreshold, RandSeed, ProjMatrix, \
    Points1, Points2)
  ProjMatrices := [ProjMatrices, ProjMatrix]
endfor

```

Finally, the image mosaic can be generated.

```

gen_projective_mosaic (Images, MosaicImage, StartImage, From, To, \
  ProjMatrices, StackingOrder, 'false', \
  MosaicMatrices2D)

```

Note that image mosaicking is a tool for a quick and easy generation of large images from several overlapping images. For this task, it is not necessary to calibrate the camera. If you need a high-precision image mosaic, you should use the method described in [chapter 9](#) on page 187.

In the following sections, the individual steps for the generation of a mosaic image are described.

10.1 Rules for Taking Images for a Mosaic Image

The following rules for the acquisition of the separate images should be considered:

- The images must overlap each other.
- The overlapping area of the images must be textured in order to allow the automatic matching process to identify identical points in the images. The lack of texture in some overlapping areas may be overcome by an appropriate definition of the image pairs (see [section 10.2](#)). If the whole object shows little texture, the overlapping areas should be chosen larger.
- Overlapping images must have approximately the same scale. In general, the scale differences should not exceed 5-10 %.
- The images should be radiometrically similar, at least in the overlapping areas, as no radiometric adaption of the images is carried out. Otherwise, i.e., if the brightness differs heavily between neighboring images, the seams between them will be clearly visible as can be seen in [figure 10.2](#).




Figure 10.2: A second example for image mosaicking.

The images are mapped onto a common image plane using a projective transformation. Therefore, to generate a geometrically accurate image mosaic from images of non-flat objects, the separate images must be acquired from approximately the same point of view, i.e., the camera can only be rotated around its optical center (see [figure 10.3](#)).

When dealing with flat objects, it is possible to acquire the images from arbitrary positions and with arbitrary orientations if the scale difference between the overlapping images is not too large ([figure 10.4](#)).

The lens distortions of the images are not compensated by the mosaicking process. Therefore, if lens distortions are present in the images, they cannot be mosaicked with high accuracy, i.e., small distortions at the seams between neighboring images cannot be prevented (see [figure 10.8](#) on page 206). To eliminate this effect, the lens distortions can be compensated before starting the mosaicking process (see [section 3.4.2](#) on page 70).

If processing time is an issue, it is advisable to acquire the images in the same orientation, i.e., neither the camera nor the object should be rotated around the optical axis, too much. Then, the rotation range can be restricted for the matching process (see [section 10.4](#) on page 208).

10.2 Definition of Overlapping Image Pairs

As shown in the introductory example, it is necessary to define the overlapping image pairs between which the transformation is to be determined. The successive matching process will be carried out for these image pairs only.




Figure 10.3: Image acquisition for non-flat objects.




Figure 10.4: Image acquisition for flat objects.

Figure 10.5 shows two configurations of separate images. For configuration (a), the definition of the image pairs is simply (1,2) and (2,3), which can be defined in HDevelop as:

```
From := [1,2]
To := [2,3]
```

In any case, it is important to ensure that each image must be “connected” to all the other images. For example, for configuration (b) of figure 10.5, it is not possible to define the image pairs as (1,2) and (3,4),




Figure 10.5: Two configurations of overlapping images.

only, because images 1 and 2 would not be connected to images 3 and 4. In this case, it would, e.g., be possible to define the three image pairs (1,2), (1,3), and (2,4).

```
From := [1,1,2]
To := [2,3,4]
```

Assuming there is no texture in the overlapping area of image two and four, the matching could be carried out between images three and four instead.

```
From := [1,1,3]
To := [2,3,4]
```

If a larger number of separate images are mosaicked, or, e.g., an image configuration similar to the one displayed in [figure 10.6](#), where there are elongated rows of overlapping images, it is important to thoroughly arrange the image pair configuration. Otherwise it is possible that some images do not fit together precisely. This happens since the transformations between the images cannot be determined with perfect accuracy because of very small errors in the point coordinates due to noise. These errors are propagated from one image to the other.

[Figure 10.7](#) shows such an image sequence of ten images of a BGA and the resulting mosaic image. [Figure 10.8](#) shows a cut-out of that mosaic image. It depicts the seam between image 5 and image 10 for two image pair configurations, using the original images and the images where the lens distortions have been eliminated, respectively. The position of the cut-out is indicated in [figure 10.7](#) by a rectangle.

First, the matching has been carried out in the two image rows separately and the two rows are connected via image pair $1 \rightarrow 6$.

```
From := [1,2,3,4,6,7,8,9,1]
To := [2,3,4,5,7,8,9,10,6]
```

In this configuration the two neighboring images 5 and 10 are connected along a relatively long path ([figure 10.9](#)).




Figure 10.6: A configuration of ten overlapping images.




Figure 10.7: Ten overlapping images and the resulting (rigid) mosaic image.

To improve the geometrical accuracy of the image mosaic, the connections between the two image rows could be established by the image pair (3,8), as visualized in (figure 10.10)).

This can be achieved by defining the image pairs as follows.

```
From := [1,2,3,4,6,7,8,9,3]
To   := [2,3,4,5,7,8,9,10,8]
```

As can be seen in figure 10.8, now the neighboring images fit better.

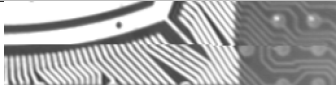
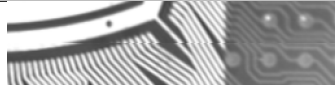
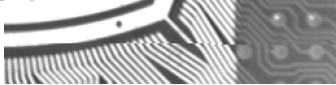
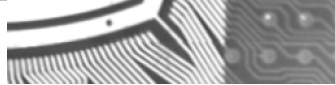
	with lens distortions	lens distortions eliminated
unfavorable configuration		
good configuration		

Figure 10.8: Seam between image 5 and image 10 for various configurations.




Figure 10.9: Unfavorable configuration of image pairs.

Recapitulating, there are three basic rules for the arrangement of the image pairs:

Take care that

1. each image is connected to all the other images.
2. the path along which neighboring images are connected is not too long.
3. the overlapping areas of image pairs are large enough and contain enough texture to ensure a proper matching.

In principle, it is also possible to define more image pairs than required (number of images minus one). However, then it cannot be controlled which pairs are actually used. Therefore, we do not recommend this approach.




Figure 10.10: Good configuration of image pairs.

10.3 Detection of Characteristic Points

HALCON provides you with various operators for the extraction of characteristic points (interest points). The most important of these operators are

- `points_foerstner`
- `points_harris` and `points_harris_binomial`
- `points_lepetit`
- `points_sojka`
- `saddle_points_sub_pix`

All of these operators can determine the coordinates of interest points with subpixel accuracy.

In [figure 10.11](#), a test image together with typical results of these interest operators is displayed.

The operator `points_foerstner` classifies the interest points into two categories: junction-like features and area-like features. The results are very reproducible even in images taken from a different point of view. Therefore, it is very well suited for the extraction of points for the subsequent matching. It is very accurate but computationally the most expensive operator out of the interest operators presented in this section.

The results of the operator `points_harris` are very reproducible, too. Admittedly, the points extracted by the operator `points_harris` are sometimes not meaningful to a human, e.g., they often lie slightly beside a corner or an eye-catching image structure. Nevertheless, it is faster than the operator `points_foerstner`. The operator `points_harris_binomial` detects points of interest using the binomial approximation of the `points_harris` operator. It is therefore faster than `points_harris`.

The operator `points_lepetit` extracts points of interest like corners or blob-like structures from the image. This operator can especially be used for very fast interest point extraction. It is the fastest out of the six operators presented in this section.

The operator `points_sojka` is specialized in the extraction of corner points.

The operator `saddle_points_sub_pix` is designed especially for the extraction of saddle points, i.e., points whose image intensity is minimal along one direction and maximal along a different direction.

The number of interest points influence the execution time and the result of the subsequent matching process. The more interest points are used, the longer the matching takes. If too few points are used the probability of an erroneous result increases.

In most cases, the default parameters of the interest operators need not be changed. Only if too many or too few interest points are found adaptations of the parameters might be necessary. For a description of the parameters, please refer to the respective pages of the Reference Manual (`points_foerstner`, `points_harris`, `points_harris_binomial`, `points_lepetit`, `points_sojka`, `saddle_points_sub_pix`).

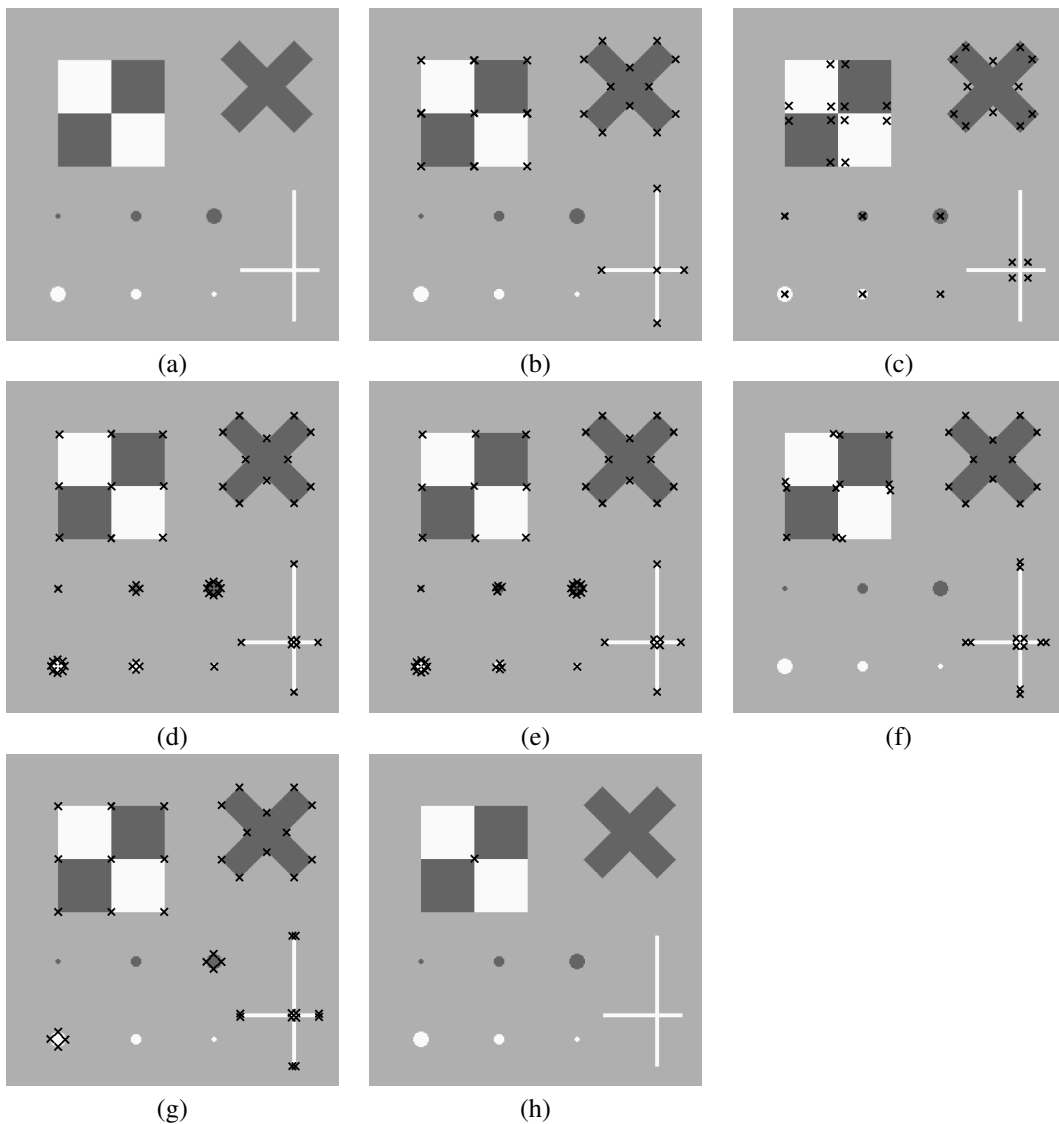


Figure 10.11: Comparison of typical results of interest operators. a) Test image; b) Förstner, junctions; c) Förstner, area; d) Harris; e) Harris, binomial f) Lepetit; g) Sojka; h) Saddle points.

10.4 Matching of Characteristic Points in Overlapping Areas and Determination of the Transformation between the Images

The most demanding task during the generation of an image mosaic is the matching process. The operator `proj_match_points_ransac` is able to perform the matching even if the two images are shifted

and rotated arbitrarily.

```
proj_match_points_ransac (ImageF, ImageT, RowFall, ColFall, RowTall, \
                          ColTall, 'sad', MaskSize, RowMove, ColMove, \
                          RowTolerance, ColTolerance, Rotation, \
                          MatchThreshold, 'gold_standard', \
                          DistanceThreshold, RandSeed, ProjMatrix, Points1, \
                          Points2)
```

The only requirement is that the images should have approximately the same scale. If information about shift and rotation is available it can be used to restrict the search space, which speeds up the matching process and makes it more robust.

In case the matching fails, ensure that there are enough characteristic points and that the search space and the maximum rotation are defined appropriately.

If the images that should be mosaicked contain repetitive patterns, like the two images of a BGA shown in [figure 10.12a](#), it may happen that the matching does not work correctly. In the resulting erroneous mosaic image, the separate images may not fit together or may be heavily distorted. To achieve a correct matching result for such images, it is important to provide initial values for the shift between the images with the parameters [RowMove](#) and [ColMove](#). In addition, the search space should be restricted to an area that contains only one instance of the repetitive pattern, i.e., the values of the parameters [RowTolerance](#) and [ColTolerance](#) should be chosen smaller than the distance between the instances of the repetitive pattern. With this, it is possible to obtain proper mosaic images, even for objects like BGAs (see [figure 10.12b](#)).




Figure 10.12: Separate images (a) and mosaic image (b) of a BGA.

For a detailed description of the other parameters, please refer to the Reference Manual ([proj_match_points_ransac](#)).

The results of the operator `proj_match_points_ransac` are the projective transformation matrix and the two tuples `Points1` and `Points2` that contain the indices of the matched input points from the two images.

The projective transformation matrices resulting from the matching between the image pairs must be accumulated.

```
ProjMatrices := [ProjMatrices, ProjMatrix]
```

Alternatively, if it is known that the mapping between the images is a rigid 2D transformation, the operator `proj_match_points_ransac` can be used to determine the point correspondences only, since it returns the indices of the corresponding points in the tuples `Points1` and `Points2`. With this, the corresponding point coordinates can be selected.

```
RowF := subset(RowFAll, Points1)
ColF := subset(ColFAll, Points1)
RowT := subset(RowTAll, Points2)
ColT := subset(ColTAll, Points2)
```

Then, the rigid transformation between the image pair can be determined with the operator `vector_to_rigid`. Note that we have to add 0.5 to the coordinates to make the extracted pixel positions fit the coordinate system that is used by the operator `gen_projective_mosaic`.

```
vector_to_rigid (RowF+0.5, ColF+0.5, RowT+0.5, ColT+0.5, HomMat2D)
```

Because `gen_projective_mosaic` expects a 3×3 transformation matrix, but `vector_to_rigid` returns a 2×3 matrix, we have to add the last row $[0, 0, 1]$ to the transformation matrix before we can accumulate it.

```
ProjMatrix := [HomMat2D, 0, 0, 1]
ProjMatrices := [ProjMatrices, ProjMatrix]
```

Furthermore, the operator `proj_match_points_ransac_guided` is available. Like `proj_match_points_ransac`, it can be used to calculate the projective transformation matrix between two images by finding correspondences between points. But in contrast to `proj_match_points_ransac`, it is based on a known approximation of the projective transformation matrix. Thus, it can be used, for example, to speed up the matching of very large images by implementing an image-pyramid-based projective matching algorithm. The HDevelop example program `hdevelop\Tools\Mosaicking\mosaicking_pyramid.hdev` shows how to implement the image-pyramid-based approach and compares the runtime for different numbers of pyramid levels.

10.5 Generation of the Mosaic Image

Once the transformations between the image pairs are known the mosaic image can be generated with the operator `gen_projective_mosaic`.

```
gen_projective_mosaic (Images, MosaicImage, StartImage, From, To, \
                      ProjMatrices, StackingOrder, 'false', \
                      MosaicMatrices2D)
```

It requires the images to be given in a tuple. All images are projected into the image plane of a so-called start image. The start image can be defined by its position in the image tuple (starting with I) with the parameter `StartImage`.

Additionally, the image pairs must be specified together with the corresponding transformation matrices.

The order in which the images are added to the mosaic image can be specified with the parameter `StackingOrder`. The first index in this array will end up at the bottom of the image stack while the last one will be on top. If `'default'` is given instead of an array of integers, the canonical order (the order in which the images are given) will be used.

If the domains of the images should be transformed as well, the parameter `TransformRegion` must be set to `'true'`.

The output parameter `MosaicMatrices2D` contains the projective 3×3 transformation matrices for the mapping of the separate images into the mosaic image. These matrices can, e.g., be used to transform features extracted from the separate images into the mosaic image by using the operators `projective_trans_pixel`, `projective_trans_region`, `projective_trans_contour_xld`, or `projective_trans_image`.

10.6 Bundle Adjusted Mosaicking

It is also possible to generate the mosaic based on the matching results of all overlapping image pairs. The transformation matrices between the images are then determined together within one bundle adjustment. For this, the operators `bundle_adjust_mosaic` and `gen_bundle_adjusted_mosaic` are used.

The main advantage of the bundle adjusted mosaicking compared with the mosaicking based on single image pairs is that the bundle adjustment determines the geometry of the mosaic as robustly as possible. Typically, this leads to more accurate results. Another advantage is that there is no need to figure out a good pair configuration, you simply pass the matching results of all overlapping image pairs. What is more, it is possible to define the class of transformations that is used for the transformation between the individual images. A disadvantage of the bundle adjusted mosaicking is that it takes more time to perform the matching between all overlapping image pairs instead of just using a subset. Furthermore, if the matching between two images was erroneous, sometimes the respective image pair is difficult to find in the set of all image pairs.

With this, it is obvious that the bundle adjustment is worthwhile if there are multiple overlaps between the images, i.e., if there are more than $n - 1$ overlapping image pairs, with n being the number of images. Another reason for using the bundle adjusted mosaicking is the possibility to define the class of transformations.

The example program `solution_guide\3d_vision\bundle_adjusted_mosaicking.hdev` shows how to generate the bundle adjusted mosaic from the ten images of the BGA displayed in [figure 10.7](#) on page 205. The design of the program is very similar to that of the example program `solution_guide\3d_vision\mosaicking.hdev`, which is described in the introduction of [chapter 10](#) on page 199. The main differences are that

- the matching is carried out between all overlapping images,
- in addition to the projective transformation matrices also the coordinates of the corresponding points must be accumulated, and
- the operator `gen_projective_mosaic` is replaced with the operators `bundle_adjust_mosaic` and `gen_bundle_adjusted_mosaic`.

First, the matching is carried out between all overlapping image pairs, which can be defined as follows:

```
From := [1,1,1,2,2,2,2,3,3,3,3,4,4,4,4,5,5]
To   := [6,7,2,6,7,8,3,7,8,9,4,8,9,10,5,9,10]
```

In addition to the accumulation of the projective transformation matrices, as described in [section 10.4](#) on page 208, also the coordinates of the corresponding points as well as the number of corresponding points must be accumulated.

```
Rows1 := [Rows1,subset(RowFAll,Points1)]
Cols1 := [Cols1,subset(ColFAll,Points1)]
Rows2 := [Rows2,subset(RowTAll,Points2)]
Cols2 := [Cols2,subset(ColTAll,Points2)]
NumCorrespondences := [NumCorrespondences,|Points1|]
```

This data is needed by the operator `bundle_adjust_mosaic`, which determines the bundle adjusted transformation matrices.

```
bundle_adjust_mosaic (10, StartImage, From, To, ProjMatrices, Rows1, Cols1, \
                    Rows2, Cols2, NumCorrespondences, Transformation, \
                    MosaicMatrices2D, Rows, Cols, Error)
```

The parameter `Transformation` defines the class of transformations that is used for the transformation between the individual images. Possible values are *'projective'*, *'affine'*, *'similarity'*, and *'rigid'*. Thus, if you know, e.g., that the camera looks perpendicular onto a planar object and that the camera movement between the images is restricted to rotations and translations in the object plane, you can choose the transformation class *'rigid'*. If translations may also occur in the direction perpendicular to the object plane, you must use *'similarity'* because this transformation class allows scale differences between the images. If the camera looks tilted onto the object, the transformation class *'projective'* must be used, which can be approximated by the transformation class *'affine'*. [Figure 10.13](#) shows cut-outs of the resulting mosaic images. They depict the seam between image 5 and image 10. The mosaic images have been created using the images where the lens distortions have been eliminated. The position of the cut-out within the whole mosaic image is indicated by the rectangle in [figure 10.7](#) on page 205.

Finally, with the transformation matrices `MosaicMatrices2D`, which are determined by the operator `bundle_adjust_mosaic`, the mosaic can be generated with the operator `gen_bundle_adjusted_mosaic`.

```
gen_bundle_adjusted_mosaic (Images, MosaicImage, MosaicMatrices2D, \
                          StackingOrder, TransformRegion, TransMat2D)
```




Figure 10.13: Seam between image 5 and image 10 for different classes of transformations: (a) projective, (b) affine, (c) similarity, and (d) rigid.

10.7 Spherical Mosaicking

The methods described in the previous sections arranged the images on a plane. As the name suggests, using spherical mosaicking you can arrange them on a sphere instead. Note that this method can only be used if the camera is only rotated around its optical center or zoomed. If the camera movement includes a translation or if the rotation is not performed exactly around the optical center, the resulting mosaic image will not be accurate and can therefore not be used for high-accuracy applications.

To create a spherical mosaic, you first perform the matching as described in the previous sections to determine the projective transformation matrices between the individual images. This information is the input for the operator `stationary_camera_self_calibration`, which determines the internal camera parameters of the camera and the rotation matrices for each image. Based on this information, the operator `gen_spherical_mosaic` then creates the mosaic image. Please have a look at the HDevelop example program `Calibration\Self-Calibration\stationary_camera_self_calibration.hdev` for more information about how to use these operators.

As an alternative, you can map the images on the six sides of a cube using `gen_cube_map_mosaic`. Cube maps are especially useful in computer graphics.

Chapter 11

Rectification of Arbitrary Distortions

For many applications like OCR or bar code reading, distorted images must be rectified prior to the extraction of information. The distortions may be caused by the perspective projection and by the radial lens distortions as well as by the decentering lens distortions, a non-flat object surface, or by any other reason. In the first three cases, i.e., if the object surface is flat and the camera shows only radial or decentering distortions, the rectification can be carried out very precisely as described in [section 3.4.1](#) on page 62. For the remaining cases, a piecewise bilinear rectification can be carried out. In HALCON, this kind of rectification is called grid rectification.

The following example (`solution_guide\3d_vision\grid_rectification_ruled_surface.hdev`) shows how the grid rectification can be used to rectify the image of a cylindrically shaped object ([figure 11.1](#)). In the rectified image ([figure 11.1b](#)), the bar code could be read correctly, which was not possible in the original image ([figure 11.1a](#)). Note that since HALCON 8.0, the bar code reader can manage also the bar code in the original image, but for even more distorted bar codes nevertheless a rectification can be necessary.

The main idea of the grid rectification is that the mapping for the rectification is determined from an image of the object, where the object is covered by a known pattern.

First, this pattern, which is called rectification grid, must be created with the operator `create_rectification_grid`.

```
create_rectification_grid (WidthOfGrid, NumSquares, 'rectification_grid.ps')
```

The resulting PostScript file must be printed. An example for such a rectification grid is shown in [figure 11.2a](#).

Now, the object must be wrapped with the *rectification grid* and an image of the wrapped object must be taken ([figure 11.2b](#)).

From this image, the mapping that describes the transformation from the distorted image into the rectified image can be derived. For this, first, the rectification grid must be extracted. Then, the rectification map is derived from the distorted grid. This can be achieved by the following lines of code:

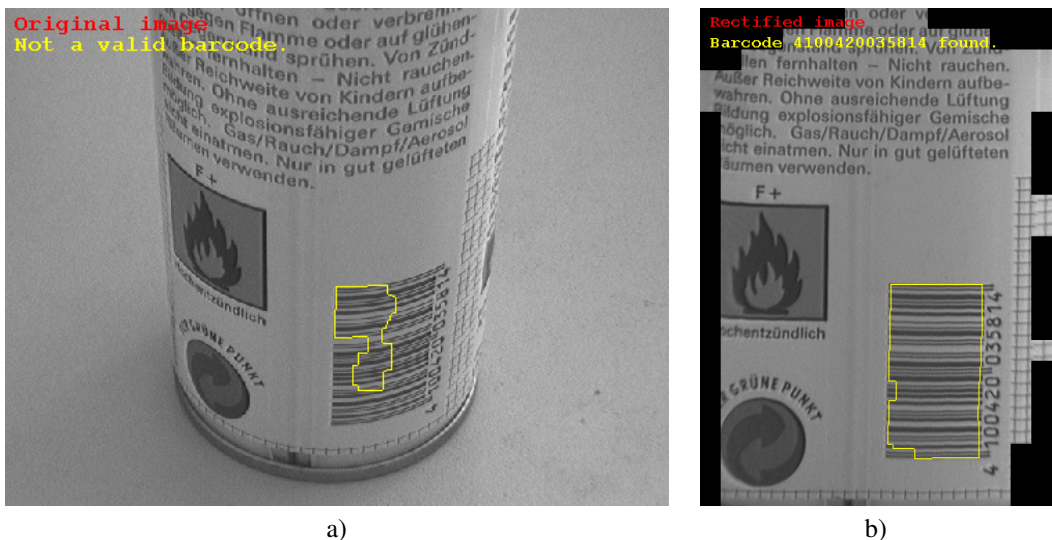


Figure 11.1: Cylindrical object: a) Original image; b) rectified image.




Figure 11.2: a) Example of a rectification grid. b) Cylindrical object wrapped with the rectification grid.

```

find_rectification_grid (Image, GridRegion, MinContrast, Radius)
reduce_domain (Image, GridRegion, ImageReduced)
saddle_points_sub_pix (ImageReduced, 'facet', SigmaSaddlePoints, Threshold, \
    Row, Col)
connect_grid_points (ImageReduced, ConnectingLines, Row, Col, \
    SigmaConnectGridPoints, MaxDist)
gen_grid_rectification_map (ImageReduced, ConnectingLines, Map, Meshes, \
    GridSpacing, 0, Row, Col, 'bilinear')
    
```


Using the derived map, any image that shows the same distortions can be rectified such that the parts that were covered by the rectification grid appear undistorted in the rectified image (figure 11.1b). This mapping is performed by the operator `map_image`.

```
map_image (ImageReduced, Map, ImageMapped)
```

In the following section, the basic principle of the grid rectification is described. Then, some hints for taking images of the rectification grid are given. In section 11.3 on page 220, the use of the involved HALCON operators is described in more detail based on the above example application. Finally, it is described briefly how to use self-defined grids for the generation of rectification maps.

11.1 Basic Principle

The basic principle of the grid rectification is that a mapping from the distorted image into the rectified image is determined from a distorted image of the rectification grid whose undistorted geometry is well known: The black and white fields of the printed rectification grid are squares (figure 11.3).




Figure 11.3: Rectification grid.

In the distorted image, the black and white fields do not appear as squares (figure 11.4a) because of the non-planar object surface, the perspective distortions, and the lens distortions.

To determine the mapping for the rectification of the distorted image, the distorted rectification grid must be extracted. For this, first, the corners of the black and white fields must be extracted with the operator `saddle_points_sub_pix` (figure 11.4b). These corners must be connected along the borders of the black and white fields with the operator `connect_grid_points` (figure 11.4c). Finally, the connecting lines must be combined into meshes (figure 11.4d) with the operator `gen_grid_rectification_map`, which also determines the mapping for the rectification of the distorted image.

If you want to use a self-defined grid, the grid points must be defined by yourself. Then, the operator `gen_arbitrary_distortion_map` can be used to determine the mapping (see section 11.4 on page 223 for an example).

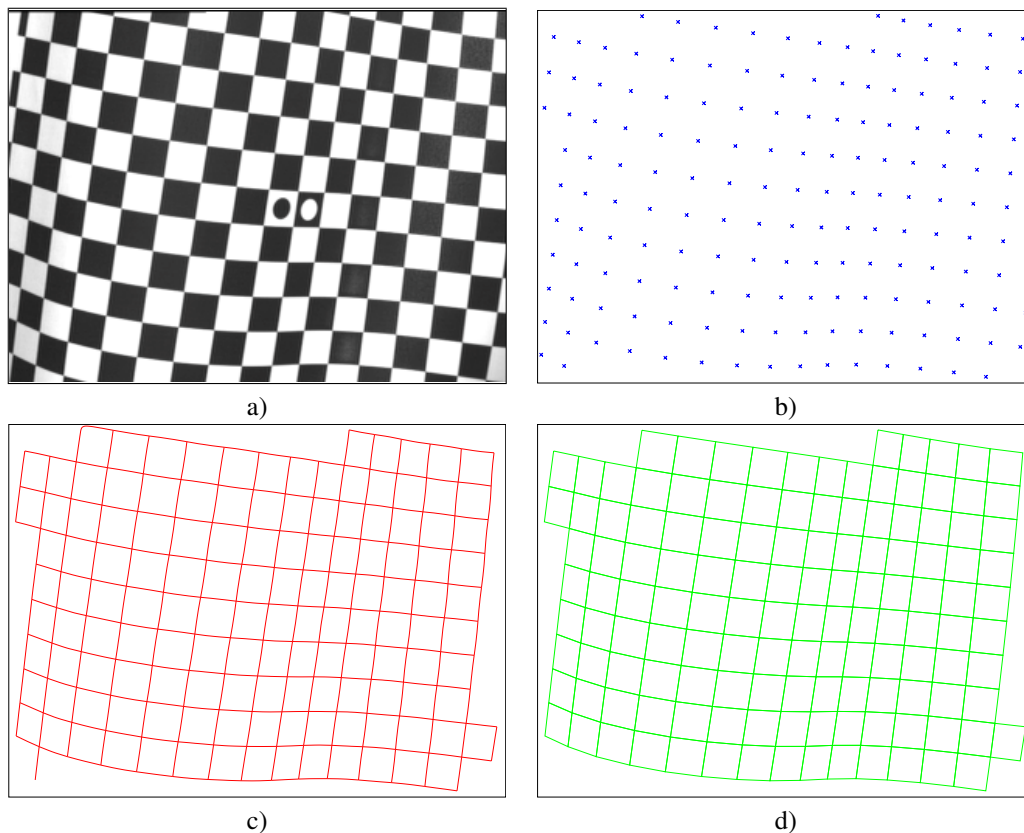


Figure 11.4: Distorted rectification grid: a) Image; b) extracted corners of the black and white fields; c) lines that connect the corners; d) extracted rectification grid.

The mapping is determined such that the distorted rectification grid will be mapped into its original undistorted geometry (figure 11.5). With this mapping, any image that shows the same distortions can be rectified easily with the operator `map_image`. Note that within the meshes a bilinear interpolation is carried out. Therefore, it is important to use a rectification grid with an appropriate grid size (see section 11.2 for details).

11.2 Rules for Taking Images of the Rectification Grid

If you want to achieve accurate results, please follow the rules given in this section:

- The image must not be overexposed or underexposed: otherwise, the extraction of the corners of the black and white fields of the rectification grid may fail.
- The contrast between the bright and the dark fields of the rectification grid should be as high as possible.




Figure 11.5: Mapping of the distorted rectification grid (a) into the undistorted rectification grid (b).

- Ensure that the rectification grid is homogeneously illuminated.
- The images should contain as little noise as possible.
- The border length of the black and white fields should be at least 10 pixels.

In addition to these few rules for the taking of the images of the rectification grid, it is very important to use a rectification grid with an appropriate grid size because the mapping is determined such that within the meshes of the rectification grid a bilinear interpolation is applied. Because of this, non-linear distortions within the meshes cannot be eliminated.

The use of a rectification grid that is too coarse (figure 11.6a), i.e., whose grid size is too large, leads to errors in the rectified image (figure 11.6b).




Figure 11.6: Cylindrical object covered with a very coarse rectification grid: a) Distorted image; b) rectified image.

If it is necessary to fold the rectification grid, it should be folded along the borders of the black and white fields (figure 11.7a), the rectified image (figure 11.7b) will contain distortions because of the bilinear interpolation within the meshes.




Figure 11.7: Rectification grid folded across the borders of the black and white fields: a) Distorted image; b) rectified image.

11.3 Machine Vision on Ruled Surfaces

In this section, the rectification of images of ruled surfaces is described in detail. Again, the example of the cylindrically shaped object (`solution_guide\3d_vision\grid_rectification_ruled_surface.hdev`) is used to explain the involved operators.

First, the operator `create_rectification_grid` is used to create a suitable rectification grid.

```
create_rectification_grid (WidthOfGrid, NumSquares, 'rectification_grid.ps')
```

The parameter `WidthOfGrid` defines the effectively usable size of the rectification grid in meters and the parameter `NumSquares` sets the number of squares (black and white fields) per row. The rectification grid is written to the PostScript file that is specified by the parameter `GridFile`.

To determine the mapping, an image of the rectification grid, wrapped around the object, must be taken as described in [section 11.2](#) on page 218. [Figure 11.8a](#) shows an image of a cylindrical object and [figure 11.8b](#) shows the same object wrapped by the rectification grid.

Then, the rectification grid is searched in this image with the operator `find_rectification_grid`.

```
find_rectification_grid (Image, GridRegion, MinContrast, Radius)
```

The operator `find_rectification_grid` extracts image areas with a contrast of at least `MinContrast` and fills up the holes in these areas. Note that in this case, contrast is defined as the gray value difference of neighboring pixels in a slightly smoothed copy of the image (Gaussian smoothing with $\sigma = 1.0$). Therefore, the value for the parameter `MinContrast` must be set significantly lower than the gray value difference between the black and white fields of the rectification grid. Small areas of high contrast are then eliminated by an opening with the radius `Radius`. The resulting region is used to restrict the search space for the following steps with the operator `reduce_domain` (see [figure 11.9a](#)).

```
reduce_domain (Image, GridRegion, ImageReduced)
```




Figure 11.8: Cylindrical object: a) Without and b) with rectification grid.

The corners of the black and white fields appear as saddle points in the image. They can be extracted with the operator `saddle_points_sub_pix` (see figure 11.9b).

```
saddle_points_sub_pix (ImageReduced, 'facet', SigmaSaddlePoints, Threshold, \
                      Row, Col)
```

The parameter `Sigma` controls the amount of Gaussian smoothing that is carried out before the actual extraction of the saddle points. Which point is accepted as a saddle point is based on the value of the parameter `Threshold`. If `Threshold` is set to higher values, fewer but more distinct saddle points are returned than if `Threshold` is set to lower values. The filter method that is used for the extraction of the saddle points can be selected by the parameter `Filter`. It can be set to `'facet'` or `'gauss'`. The method `'facet'` is slightly faster. The method `'gauss'` is slightly more accurate but tends to be more sensitive to noise.




Figure 11.9: Distorted rectification grid: a) Image reduced to the extracted area of the rectification grid; b) extracted corners of the black and white fields; c) lines that connect the corners.

To generate a representation of the distorted rectification grid, the extracted saddle points must be connected along the borders of the black and white fields ([figure 11.9c](#)). This is done with the operator `connect_grid_points`.

```
connect_grid_points (ImageReduced, ConnectingLines, Row, Col, \
                    SigmaConnectGridPoints, MaxDist)
```

Again, the parameter `Sigma` controls the amount of Gaussian smoothing that is carried out before the extraction of the borders of the black and white fields. When a tuple of three values $[\textit{sigma_min}, \textit{sigma_max}, \textit{sigma_step}]$ is passed instead of only one value, the operator `connect_grid_points` tests every sigma within the given range from $\textit{sigma_min}$ to $\textit{sigma_max}$ with a step size of $\textit{sigma_step}$ and chooses the sigma that causes the largest number of connecting lines. The same happens when a tuple of only two values $\textit{sigma_min}$ and $\textit{sigma_max}$ is passed. However, in this case a fixed step size of 0.05 is used. The parameter `MaxDist` defines the maximum distance with which an edge may be linked to the respectively closest saddle point. This helps to overcome the problem that edge detectors typically return inaccurate results in the proximity of edge junctions. [Figure 11.10](#) shows the connecting lines if the parameter `MaxDist` has been selected inappropriately: In [figure 11.10a](#), `MaxDist` has been selected too small, whereas in [figure 11.10b](#), it has been selected too large.




Figure 11.10: Connecting lines: Parameter `MaxDist` selected a) too small and b) too large.

Then, the rectification map is determined from the distorted grid with the operator `gen_grid_rectification_map`.

```
gen_grid_rectification_map (ImageReduced, ConnectingLines, Map, Meshes, \
                           GridSpacing, 0, Row, Col, 'bilinear')
```

The parameter `GridSpacing` defines the size of the grid meshes in the rectified image. Each of the black and white fields is projected onto a square of `GridSpacing` \times `GridSpacing` pixels. The parameter `Rotation` controls the orientation of the rectified image. The rectified image can be rotated by 0, 90, 180, or 270 degrees, or it is rotated such that the black circular mark is left of the white circular mark if `Rotation` is set to 'auto'.

Using the derived rectification map, any image that shows the same distortions can be rectified very fast with the operator `map_image` (see [figure 11.11](#)). Note that the objects must appear at exactly the same

position in the distorted images.

```
map_image (ImageReduced, Map, ImageMapped)
```

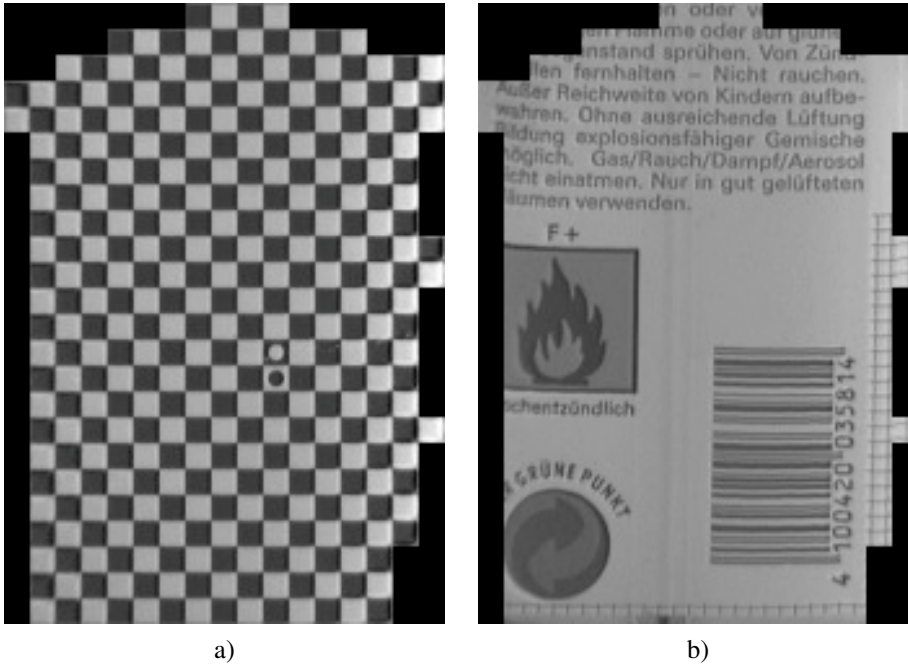


Figure 11.11: Rectified images: a) Rectification grid; b) object.

11.4 Using Self-Defined Rectification Grids

Up to now, we have used the predefined rectification grid together with the appropriate operators for its segmentation. In this section, an alternative to this approach is presented. You can arbitrarily define the rectification grid by yourself, but note that in this case you must also carry out the segmentation by yourself.

This example shows how the grid rectification can be used to generate arbitrary distortion maps based on self-defined grids.

The example application is a print inspection. It is assumed that some parts are missing and that smudges are present. In addition, lines may be vertically shifted, e.g., due to an inaccurate paper transport, i.e., distortions in the vertical direction of the printed document may be present. These distortions should not result in a rejection of the tested document. Therefore, it is not possible to simply compute the difference image between a reference image and the image that must be checked.

Figure 11.12a shows the reference image and figure 11.12b the test image that must be checked.




Figure 11.12: Images of one page of a document: a) Reference image; b) test image that must be checked.

In a first step, the displacements between the lines in the reference document and the test document are determined. With this, the rectification grid is defined. The resulting rectification map is applied to the reference image to transform it into the geometry of the test image. Now, the difference image of the mapped reference image and the test image can be computed.

The HDevelop example program `solution_guide\3d_vision\grid_rectification_arbitrary_distortion.hdev` uses the component-based matching to determine corresponding points in the reference image and the test image. First, the component model is generated with the operator `create_component_model`. Then, the corresponding points are searched in the test image with the operator `find_component_model`.

Based on the corresponding points of the reference and the test image (`RowRef`, `ColRef`, `RowTest`, and `ColTest`), the coordinates of the grid points of the distorted grid are determined. In this example, the row and column coordinates can be determined independently from each other because only the row coordinates are distorted. Note that the upper left grid point of the undistorted grid is assumed to have the coordinates $(-0.5, -0.5)$. This means that the corresponding grid point of the distorted grid will be mapped to the point $(-0.5, -0.5)$. Because there are only vertical distortions in this example, the column coordinates of the distorted grid are equidistant, starting at the value -0.5 .


```

GridSpacing := 10
ColShift := mean(ColTest-ColRef)
RefGridColValues := []
for HelpCol := -0.5 to WidthTest+GridSpacing by GridSpacing
  RefGridColValues := [RefGridColValues, HelpCol+ColShift]
endfor

```

The row coordinates of the distorted grid are determined by a linear interpolation between the above determined pairs of corresponding row coordinates.

```

MinValue := 0
MaxValue := HeightTest+GridSpacing
sample_corresponding_values (RowTest, RowRef-0.5, MinValue, MaxValue, \
  GridSpacing, RefGridRowValues)

```

The interpolation is performed within the procedure which is part of the HDevelop example program `solution_guide\3d_vision\grid_rectification_arbitrary_distortion.hdev`.

```

procedure sample_corresponding_values (Values, CorrespondingValues,
  MinValue, MaxValue,
  InterpolationInterval,
  SampledCorrespondingValues):::

```

Now, the distorted grid is generated row by row.

```

RefGridRow := []
RefGridCol := []
Ones := gen_tuple_const(|RefGridColValues|, 1)
for r := 0 to |RefGridRowValues|-1 by 1
  RefGridRow := [RefGridRow, RefGridRowValues[r]*Ones]
  RefGridCol := [RefGridCol, RefGridColValues]
endfor

```

The operator `gen_arbitrary_distortion_map` uses this distorted grid to derive the rectification map that maps the reference image into the geometry of the test image¹.

```

gen_arbitrary_distortion_map (Map, GridSpacing, RefGridRow, RefGridCol, \
  |RefGridColValues|, WidthRef, HeightRef, \
  'bilinear')

```

With this rectification map, the reference image can be transformed into the geometry of the test image. Note that the size of the mapped image depends on the number of grid cells and on the size of one grid cell, which must be defined by the parameter `GridSpacing`. Possibly, the size of the mapped reference image must be adapted to the size of the test image.

¹In this case, the reference image is mapped into the geometry of the test image to facilitate the marking of the differences in the test image. Obviously, the rectification grid can also be defined such that the test image is mapped into the geometry of the reference image.

```
map_image (ImageRef, Map, ImageMapped)
crop_part (ImageMapped, ImagePart, 0, 0, WidthTest, HeightTest)
```

Finally, the test image can be subtracted from the mapped reference image.

```
sub_image (ImagePart, ImageTest, ImageSub, 1, 128)
```

Figure 11.13 shows the resulting difference image. In this case, missing parts appear dark while the smudges appear bright.

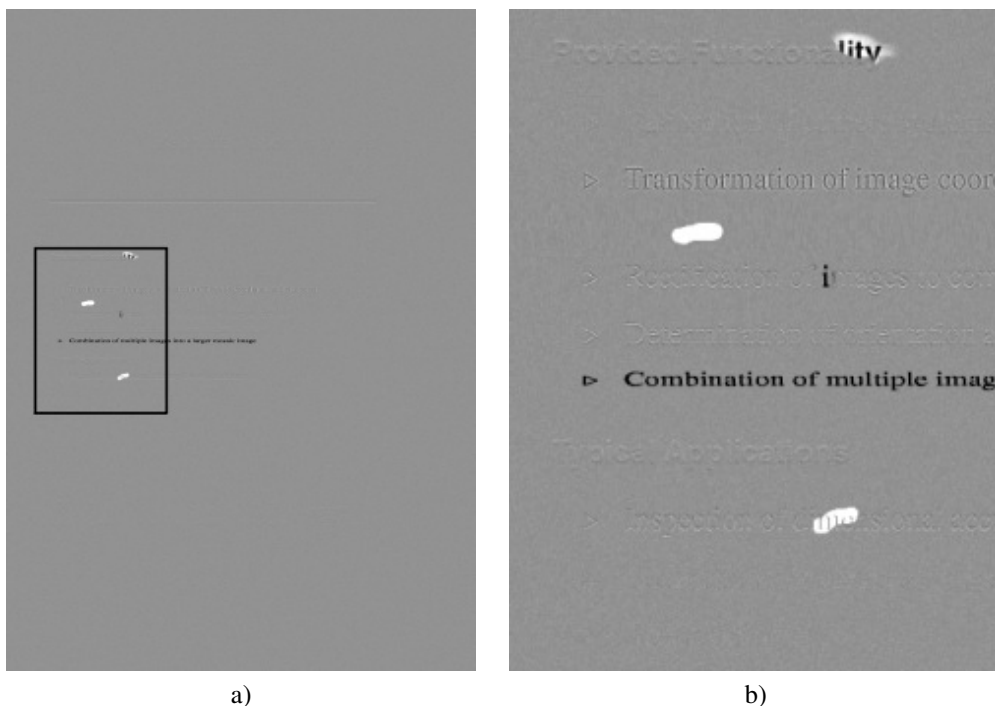


Figure 11.13: Difference image: a) The entire image overlaid with a rectangle that indicates the position of the cut-out. b) A cut-out.

The differences between the test image and the reference image can now be extracted easily from the difference image with the operator `threshold`. If the difference image is not needed, e.g., for visualization purposes, the differences can be derived directly from the test image and the reference image with the operator `dyn_threshold`.

Figure 11.14 shows the differences in a cut-out of the reference image (figure 11.14a) and of the test image (figure 11.14b). The calibration marks near the left border of figure 11.14b indicate the vertical position of the components that were used for the determination of the corresponding points. Vertical shifts of the components with respect to the reference image are indicated by a vertical line of the respective length that is attached to the respective calibration mark. All other differences that could be detected between the test image and the reference image are encircled.

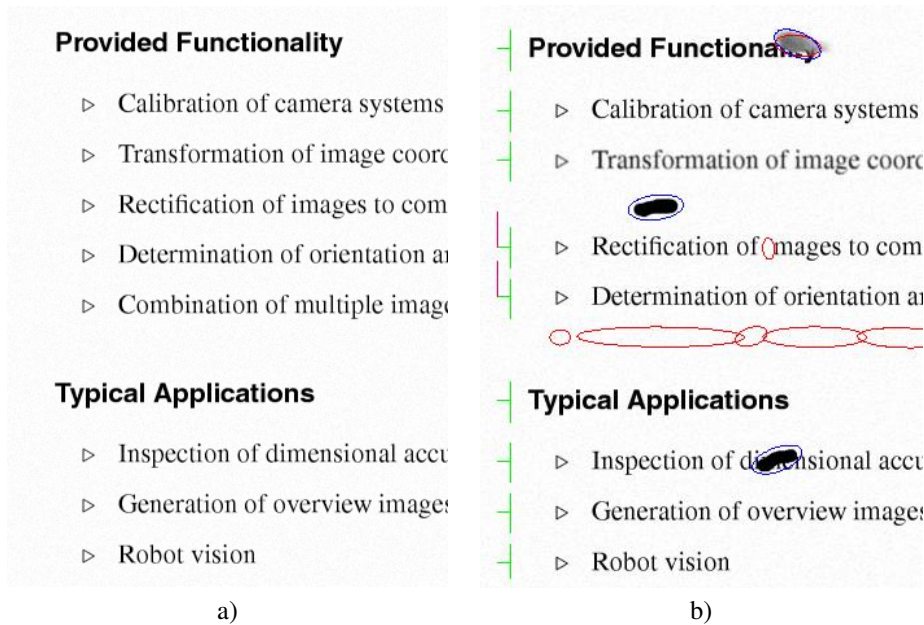


Figure 11.14: Cut-out of the reference and the checked test image with the differences marked in the test image: a) Reference image; b) checked test image.

Appendix A

The HALCON Calibration Plate

Figure A.1a shows a HALCON calibration plate. Note that it has an asymmetric pattern in the upper left corner. This pattern ensures that the pose of the calibration plate can be determined uniquely.

Old calibration plates do not have this pattern (see figure A.1b). This may lead to problems if, e.g., a stereo camera or hand-eye system must be calibrated because the poses must be determined uniquely for this. To overcome this problem, you can make an asymmetric calibration plate out of your old calibration plate by marking one corner. Pay attention that the asymmetric pattern is not too close to the circular calibration marks because otherwise it could have an influence on the geometric accuracy of the calibration result.

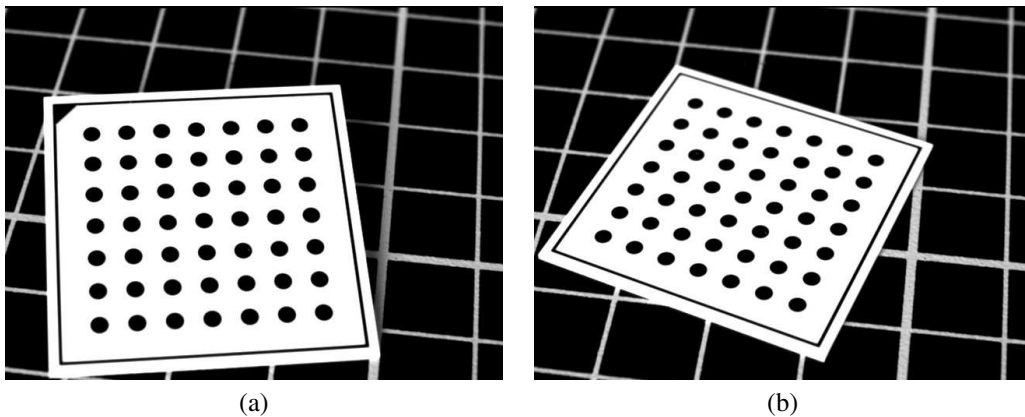


Figure A.1: (a) The HALCON calibration plate with the asymmetric pattern in the upper left corner; (b) an old calibration plate that does not have the asymmetric pattern.

There are two different types of calibration plate description files, which typically lie in the subdirectory `calib` of the folder where you installed HALCON: The standard description files for calibration plates that have the asymmetric pattern and the old description files for calibration plates without the asymmetric pattern. The old description files are indicated with the suffix *old* (orientation-less description).

The behavior of the operator `find_marks_and_pose` depends on the combination of the used calibration plate and the specified calibration plate description file:

Calibration plate	Description file	Behavior of <code>find_marks_and_pose</code>
asymmetric	asymmetric	The pose will be determined uniquely.
asymmetric	old	The pose will be determined such that the x-axis points to the right and the y-axis points downwards.
old	asymmetric	The operator <code>find_marks_and_pose</code> returns an error because it cannot find the asymmetric pattern.
old	old	The pose will be determined such that the x-axis points to the right and the y-axis points downwards.

Appendix B

HDevelop Procedures Used in this Solution Guide

B.1 gen_hom_mat3d_from_three_points

```

procedure gen_hom_mat3d_from_three_points (Origin, PointOnXAxis,
                                          PointInXYPlane, HomMat3d):::
XAxis := [PointOnXAxis[0]-Origin[0],PointOnXAxis[1]-Origin[1], \
          PointOnXAxis[2]-Origin[2]]
XAxisNorm := XAxis/sqrt(sum(XAxis*XAxis))
VectorInXYPlane := [PointInXYPlane[0]-Origin[0], \
                    PointInXYPlane[1]-Origin[1], \
                    PointInXYPlane[2]-Origin[2]]
cross_product (XAxisNorm, VectorInXYPlane, ZAxis)
ZAxisNorm := ZAxis/sqrt(sum(ZAxis*ZAxis))
cross_product (ZAxisNorm, XAxisNorm, YAxisNorm)
HomMat3d_WCS_to_RectCCS := [XAxisNorm[0],YAxisNorm[0],ZAxisNorm[0], \
                           Origin[0],XAxisNorm[1],YAxisNorm[1], \
                           ZAxisNorm[1],Origin[1],XAxisNorm[2], \
                           YAxisNorm[2],ZAxisNorm[2],Origin[2]]
hom_mat3d_invert (HomMat3d_WCS_to_RectCCS, HomMat3d)
return ()

```

This procedure uses the procedure

```

procedure cross_product (V1, V2, CrossProduct)
CrossProduct := [V1[1]*V2[2]-V1[2]*V2[1],V1[2]*V2[0]-V1[0]*V2[2], \
                 V1[0]*V2[1]-V1[1]*V2[0]]
return ()

```

B.2 parameters_image_to_world_plane_centered

```

procedure parameters_image_to_world_plane_centered (CamParam, Pose,
                                                    CenterRow, CenterCol,
                                                    WidthMappedImage,
                                                    HeightMappedImage,
                                                    ScaleForCenteredImage,
                                                    PoseForCenteredImage):::
* Determine the scale for the mapping
* (here, the scale is determined such that in the
* surroundings of the given point the image scale of the
* mapped image is similar to the image scale of the original image)
Dist_ICS := 1
image_points_to_world_plane (CamParam, Pose, CenterRow, CenterCol, 1, \
                             CenterX, CenterY)
image_points_to_world_plane (CamParam, Pose, CenterRow+Dist_ICS, CenterCol, \
                             1, BelowCenterX, BelowCenterY)
image_points_to_world_plane (CamParam, Pose, CenterRow, CenterCol+Dist_ICS, \
                             1, RightOfCenterX, RightOfCenterY)
distance_pp (CenterY, CenterX, BelowCenterY, BelowCenterX, \
            Dist_WCS_Vertical)
distance_pp (CenterY, CenterX, RightOfCenterY, RightOfCenterX, \
            Dist_WCS_Horizontal)
ScaleVertical := Dist_WCS_Vertical/Dist_ICS
ScaleHorizontal := Dist_WCS_Horizontal/Dist_ICS
ScaleForCenteredImage := (ScaleVertical+ScaleHorizontal)/2.0
* Determine the parameters for set_origin_pose such
* that the point given via get_mbutton will be in the center of the
* mapped image
DX := CenterX-ScaleForCenteredImage*WidthMappedImage/2.0
DY := CenterY-ScaleForCenteredImage*HeightMappedImage/2.0
DZ := 0
set_origin_pose (Pose, DX, DY, DZ, PoseForCenteredImage)
return ()

```


B.3 parameters_image_to_world_plane_entire

```
procedure parameters_image_to_world_plane_entire (Image, CamParam, Pose,
                                                WidthMappedImage,
                                                HeightMappedImage,
                                                ScaleForEntireImage,
                                                PoseForEntireImage):::
* Transform the image border into the WCS (scale = 1)
full_domain (Image, ImageFull)
get_domain (ImageFull, Domain)
gen_contour_region_xld (Domain, ImageBorder, 'border')
contour_to_world_plane_xld (ImageBorder, ImageBorderWCS, CamParam, Pose, 1)
smallest_rectangle1_xld (ImageBorderWCS, MinY, MinX, MaxY, MaxX)
* Determine the scale of the mapping
ExtentX := MaxX-MinX
ExtentY := MaxY-MinY
ScaleX := ExtentX/WidthMappedImage
ScaleY := ExtentY/HeightMappedImage
ScaleForEntireImage := max([ScaleX,ScaleY])
* Shift the pose by the minimum X and Y coordinates
set_origin_pose (Pose, MinX, MinY, 0, PoseForEntireImage)
return ()
```

B.4 tilt_correction

```

procedure tilt_correction (DistanceImage, RegionDefiningReferencePlane,
                          DistanceImageCorrected):::
* Reduce the given region, which defines the reference plane
* to the domain of the distance image
get_domain (DistanceImage, Domain)
intersection (RegionDefiningReferencePlane, Domain, \
              RegionDefiningReferencePlane)
* Determine the parameters of the reference plane
moments_gray_plane (RegionDefiningReferencePlane, DistanceImage, MRow, MCol, \
                    Alpha, Beta, Mean)
* Generate a distance image of the reference plane
get_image_pointer1 (DistanceImage, Pointer, Type, Width, Height)
area_center (RegionDefiningReferencePlane, Area, Row, Column)
gen_image_surface_first_order (ReferencePlaneDistance, Type, Alpha, Beta, \
                               Mean, Row, Column, Width, Height)
* Subtract the distance image of the reference plane
* from the distance image of the object
sub_image (DistanceImage, ReferencePlaneDistance, DistanceImageWithoutTilt, \
          1, 0)
* Determine the scale factor for the reduction of the distance values
CosGamma := 1.0/sqrt(Alpha*Alpha+Beta*Beta+1)
* Reduce the distance values
scale_image (DistanceImageWithoutTilt, DistanceImageCorrected, CosGamma, 0)
return ()

```

B.5 visualize_results_of_find_marks_and_pose

```

procedure visualize_results_of_find_marks_and_pose (Image, WindowHandle,
                                                    RCoord, CCoord, Pose,
                                                    CamPar, CalTabFile):::

dev_set_window (WindowHandle)
dev_display (Image)
dev_set_color ('yellow')
gen_cross_contour_xld (Cross, RCoord, CCoord, 6, 0)
dev_display (Cross)
display_calplate_coordinate_system (CalTabFile, Pose, CamPar, WindowHandle)
return ()

```

B.6 display_calplate_coordinate_system

```

procedure display_calplate_coordinate_system (CalTabFile, Pose, CamPar,
                                             WindowHandle):::
  caltab_points (CalTabFile, X, Y, Z)
  * arrow should point to farthest marks
  ArrowLength := abs(X[0])
  display_3d_coordinate_system (Pose, CamPar, ArrowLength, WindowHandle, \
                               'blue')
  return ()

```

B.7 display_3d_coordinate_system

```

procedure display_3d_coordinate_system (Pose, CamPar, ArrowLength,
                                       WindowHandle, Color):::
  pose_to_hom_mat3d (Pose, HomMat3D)
  * store coordinates of the arrows in tuples
  * sequence: origin, x-axis, y-axis, z-axis
  ArrowsXCoords := [0, ArrowLength, 0, 0]
  ArrowsYCoords := [0, 0, ArrowLength, 0]
  ArrowsZCoords := [0, 0, 0, ArrowLength]
  * transform arrow points into camera coordinates
  affine_trans_point_3d (HomMat3D, ArrowsXCoords, ArrowsYCoords, \
                        ArrowsZCoords, ArrowsXCoords_cam, ArrowsYCoords_cam, \
                        ArrowsZCoords_cam)
  * get the image coordinates
  project_3d_point (ArrowsXCoords_cam, ArrowsYCoords_cam, ArrowsZCoords_cam, \
                  CamPar, ArrowsRows, ArrowsCols)
  * display the coordinate system
  dev_set_color (Color)
  gen_contour_polygon_xld (XAxis, [ArrowsRows[0], ArrowsRows[1]], \
                          [ArrowsCols[0], ArrowsCols[1]])
  dev_display (XAxis)
  set_tposition (WindowHandle, ArrowsRows[1], ArrowsCols[1])
  write_string (WindowHandle, 'x')
  gen_contour_polygon_xld (YAxis, [ArrowsRows[0], ArrowsRows[2]], \
                          [ArrowsCols[0], ArrowsCols[2]])
  dev_display (YAxis)
  set_tposition (WindowHandle, ArrowsRows[2], ArrowsCols[2])
  write_string (WindowHandle, 'y')
  gen_contour_polygon_xld (ZAxis, [ArrowsRows[0], ArrowsRows[3]], \
                          [ArrowsCols[0], ArrowsCols[3]])
  dev_display (ZAxis)
  set_tposition (WindowHandle, ArrowsRows[3], ArrowsCols[3])
  write_string (WindowHandle, 'z')
  return ()

```

B.8 calc_calplate_pose_movingcam

```
procedure calc_calplate_pose_movingcam (CalibrationPose, CameraPose,
                                         RobotPoseInverse, CalplatePose)::
* CalplatePose = cam_H_calplate = cam_H_tool * tool_H_base * \
* base_H_calplate
pose_to_hom_mat3d (CalibrationPose, base_H_calplate)
pose_to_hom_mat3d (CameraPose, cam_H_tool)
pose_to_hom_mat3d (RobotPoseInverse, tool_H_base)
hom_mat3d_compose (cam_H_tool, tool_H_base, cam_H_base)
hom_mat3d_compose (cam_H_base, base_H_calplate, cam_H_calplate)
hom_mat3d_to_pose (cam_H_calplate, CalplatePose)
return ()
```

B.9 calc_calplate_pose_stationarycam

```
procedure calc_calplate_pose_stationarycam (CalibrationPose, CameraPose,
                                             RobotPose, CalplatePose)::
* CalplatePose = cam_H_calplate = cam_H_base * base_H_tool * \
* tool_H_calplate
pose_to_hom_mat3d (CalibrationPose, tool_H_calplate)
pose_to_hom_mat3d (CameraPose, cam_H_base)
pose_to_hom_mat3d (RobotPose, base_H_tool)
hom_mat3d_compose (cam_H_base, base_H_tool, cam_H_tool)
hom_mat3d_compose (cam_H_tool, tool_H_calplate, cam_H_calplate)
hom_mat3d_to_pose (cam_H_calplate, CalplatePose)
return ()
```

B.10 define_reference_coord_system

```
procedure define_reference_coord_system (ImageName, CamParam, CalplateFile,
                                       WindowHandle, PoseCamRef)::
  read_image (RefImage, ImageName)
  dev_display (RefImage)
  caltab_points (CalplateFile, X, Y, Z)
  * parameter settings for find_caltab and find_marks_and_pose
  SizeGauss := 3
  MarkThresh := 100
  MinDiamMarks := 5
  StartThresh := 128
  DeltaThresh := 10
  MinThresh := 18
  Alpha := 0.9
  MinContLength := 15
  MaxDiamMarks := 100
  find_caltab (RefImage, Caltab, CalplateFile, SizeGauss, MarkThresh, \
              MinDiamMarks)
  find_marks_and_pose (RefImage, Caltab, CalplateFile, CamParam, StartThresh, \
                      DeltaThresh, MinThresh, Alpha, MinContLength, \
                      MaxDiamMarks, RCoord, CCoord, PoseCamRef)
  disp_3d_coord_system (WindowHandle, CamParam, PoseCamRef, 0.01)
  return ()
```


Index

- 2D projective matrix from RANSAC point matching, 208
- 3D affine transformation of point, 60
- 3D alignment, 180
- 3D coordinates, 15
- 3D coordinates from binocular stereo disparity, 127
- 3D coordinates from multi-view stereo images, 135
- 3D coordinates with sheet of light, 151
- 3D distance from binocular stereo disparity, 127
- 3D homogeneous transformation matrix, 22
- 3D inspection, 10
- 3D measurement plane, 53
- 3D object model, 12
- 3D pose (position and orientation), 25
- 3D pose estimation, 75
- 3D pose of circle, 102
- 3D pose of rectangle, 102
- 3D reconstruction
 - guide, 11
- 3D reconstruction with binocular stereo, 120
- 3D reconstruction with multi-view stereo, 133
- 3D reconstruction with sheet of light (laser triangulation), 148
- 3D rotation, 18
- 3D transformation, 15
- 3D translation, 16
- 3D vision, 9
- 3D vision with single camera, 37
 - first example, 38
- access 3D matching model (surface-based), 95
- accuracy of 3D measuring with single camera, 70
- acquire image for grid rectification, 218
- acquire images for camera calibration, 49
- acquire images for depth from focus, 166
- acquire images for stereo camera calibration, 112
- acquire images for uncalibrated mosaicking, 201
- affine 3D transformation of points, 24
- area scan camera model, 27
- binocular stereo (uncalibrated), 130
- bundle adjust mosaic images, 211
- calibrate aberration for depth from focus, 169
- calibrate camera for mosaicking, 188
- calibrate camera parameters, 52
- calibrate external camera parameters, 53
- calibrate hand-eye system parameters, 178
- calibrate internal camera parameters, 52
- calibrate line scan camera parameters, 57
- calibrate multiple cameras, 110
- calibrate sheet-of-light system parameters, 143
- calibrated camera setup model, 131
- calibrated external camera parameters, 53
- calibrated internal camera parameters, 53
- calibration plate, 229
- camera calibration, 40
- camera calibration results, 53
- camera calibration troubleshooting, 57
- camera coordinate system (3D), 29
- camera model (3D), 26
- camera scale factor, 32
- change lens distortion, 61
- change radial distortion of points, 61
- change radial distortion of XLD contours, 61, 70
- check success of camera calibration, 52
- connect points of rectification grid, 217
- convert 3D pose into 3D homogeneous matrix, 26, 60, 72
- convert 3D pose type, 26
- correlation-based stereo, 114

- corresponding points for hand-eye calibration, 176
- create 3D homogeneous identity matrix, 24
- create 3D matching model (shape-based), 82
- create 3D matching model (surface-based), 94
- create 3D object model from 3D coordinates, 94
- create 3D pose, 25
- create calibration plate, 47
- create camera calibration data model, 41
- create camera setup model, 132
- create mosaic image, 210, 211
- create rectification grid, 220
- create sheet-of-light model, 149
- create spherical mosaic, 213
- create XLD contour of region, 60

- depth from focus, 157, 158
 - example, 167
- destroy 3D matching model (shape-based), 86
- destroy 3D object model, 84
- destroy camera calibration data model, 56
- destroy sheet-of-light model, 151
- disparity image with correlation-based stereo, 121
- disparity image with multigrid stereo, 123
- division model of lens distortion, 30

- external camera parameters, 29
- extract points for uncalibrated mosaicking, 207

- find 3D matching model (shape-based), 86
- find 3D matching model (surface-based), 95
- find calibration plate, 50
- find calibration plate marks and 3D pose, 50
- focal length, 30
- fundamental matrix from RANSAC point matching, 130

- get 3D object pose, 10
- get measurement range for depth from focus, 162

- grid rectification, 215
 - background information, 217

- hand-eye calibration, 174

- image center point, 32
- image plane, 27

- image plane coordinate system, 27
- internal camera parameters, 27

- lens distortion models
 - guide, 42
- lens distortion of camera parameters, 61, 70
- lens distortion of image, 70
- line scan camera model, 33

- mapping for grid rectification, 217
- mapping from image coordinates to 3D coordinates, 62
- mapping to change lens distortion, 70
- mapping to rectify arbitrary distortion, 217
- match points for uncalibrated mosaicking, 208
- measure volume with depth from focus, 168
- mosaicking (image stitching) calibrated, 187
- mosaicking (image stitching) uncalibrated, 199
- multigrid stereo, 114
- multiply 3D homogeneous matrix, 26

- obtain calibration plate, 46

- parallel projection, 30
- perspective projection, 30
- pinhole camera, 26
- polynomial model of lens distortion, 30
- pose estimation for 3D alignment, 181
- pose estimation from 3D matching (shape-based), 81
- pose estimation from 3D matching (surface-based), 93
- pose estimation from matching (descriptor-based), 101
- pose estimation from matching (perspective deformable), 101
- pose estimation from point correspondences, 76
- pose estimation from primitives fitting, 97
- prepare 3D object model, 82
- problem handling for 3D matching (shape-based), 89

- read 3D model (shape-based), 85
- read 3D object model, 82, 94
- read 3D pose, 26
- reconstruct 3D distance image with correlation-based stereo, 124
- reconstruct 3D distance image with multigrid stereo, 126

- reconstruct 3D distance with focus images, 167
- reconstruct 3D information with sheet of light, 149
- reconstruct 3D information with stereo (binocular), 120
- reconstruct 3D information with stereo (multi-view), 133
- reconstruct 3D point from lines of sight, 128
- reconstruct uncalibrated 3D information via sheet of light, 151
- rectification, 61
- rectify image of ruled surface, 220
- rectify image with user-specific rectification grid, 223
- rectify image(s), 62
- rectify image(s) for stereo, 116
- rectify images for mosaicking, 196
- relative camera pose from RANSAC point matching, 130
- remove artifacts of sheet-of-light results, 153
- resolution of stereo vision, 110
- rigid 3D transformation, 22
- robot poses for hand-eye calibration, 177
- robot vision, 13, 173

- select illumination for depth from focus, 163
- self-calibrate projective camera parameters, 213
- set 3D coordinate system of camera setup model, 132
- set calibration object for camera calibration, 45
- set camera calibration parameters, 52
- set image pairs of stereo model, 134
- set initial camera parameters for camera calibration, 41
- set mosaicking image pairs, 202
- set observed points for camera calibration, 49
- set origin of 3D pose, 26
- set sheet-of-light model parameter, 150
- set up calibrated mosaicking, 188
- set up camera for depth from focus, 160
- set up depth from focus application, 160
- set up sheet-of-light system, 141
- set up stereo camera system, 109
- sheet of light (laser triangulation), 139
- sheet-of-light result, 151
- solve depth from focus problems, 168
- special applications for depth from focus, 170

- speed up 3D matching (shape-based), 86
- speed up depth from focus, 159
- standard lens for depth from focus, 171
- stereo (binocular), 113
- stereo (multi-view), 130
- stereo vision
 - background information, 107
 - overview, 105
- suitable objects for depth from focus, 164

- telecentric camera, 26
- tile images, 196
- transform 3D coordinates into pixel coordinates (projection), 60
- transform 3D point into image coordinates, 30
- transform 3D point into pixel coordinates (projection), 60
- transform 3D shape model into pixel coordinates, 86
- transform image coordinates into 3D coordinates, 57
- transform image into 3D coordinates, 62
- transform image plane coordinates into pixel coordinates, 32
- transform pixel coordinates into 3D coordinates, 58, 64, 72
- transform region into 3D coordinates, 60
- transform XLD contour into 3D coordinates, 60
- translate 3D homogeneous matrix, 24
- translate 3D homogeneous matrix around local axes, 24
- translate 3D pose, 64

- use 3D camera for sheet-of-light measuring, 155
- use hand-eye system parameters, 180
- user-specific calibration object, 47

- world coordinate system (3D), 29
- write 3D pose, 26

