COGNEX®

AlignPlus 2D Hand-Eye Calibration Concepts

5/12/2014
Version 1.0

Contents

Introduction

AlignPlus hand-eye calibration allows you to calibrate images obtained with your cameras to the platform on which the object to be inspected is and whose position relative to your cameras moves. The images you obtain using your cameras may exhibit lens distortion and perspective distortion, your motion system input may be offset from the actual position it moves to and it may have systematic errors as well. AlignPlus hand-eye calibration allows you to view and inspect features of objects in an undistorted manner (with physically correct length units) and with placement in the native coordinate system of the motion system. That is, it shows you the object to be inspected in its real physical appearance and it tells you where that object is (in the native coordinate space).

The following are determined regarding your motion system during calibration:

- Origin and axis directions of native coordinate space – the native coordinate space of the motion system is determined as a mapping from the image coordinate space for each camera
- Systematic errors along the axes of movement – systematic motion system errors are corrected

Note: Even a configuration with a perfect motion system (that is, one that performs the movements as required within precision thresholds) and no image distortions requires hand-eye calibration because the mappings between each image coordinate space and the native coordinate space need to be determined. These mappings enable positioning image features in the native coordinate space. In other words, by establishing these mappings, correspondence is established between features found in images taken by the cameras to the physical coordinates of these features in the native coordinate space.

To get basic feedback on your motion stage and identify problems before hand-eye calibration, you perform motion stage validation before hand-eye calibration.

AlignPlus hand-eye calibration is a 2D calibration, it calibrates images in the plane of the platform on which the object to be inspected is. This tool provides no 3D information.

There are two types of scenarios: either the platform moves with the object to be inspected and the cameras are stationary or the cameras move and the platform with the object to be inspected is stationary.

The following figure illustrates this with a single camera.

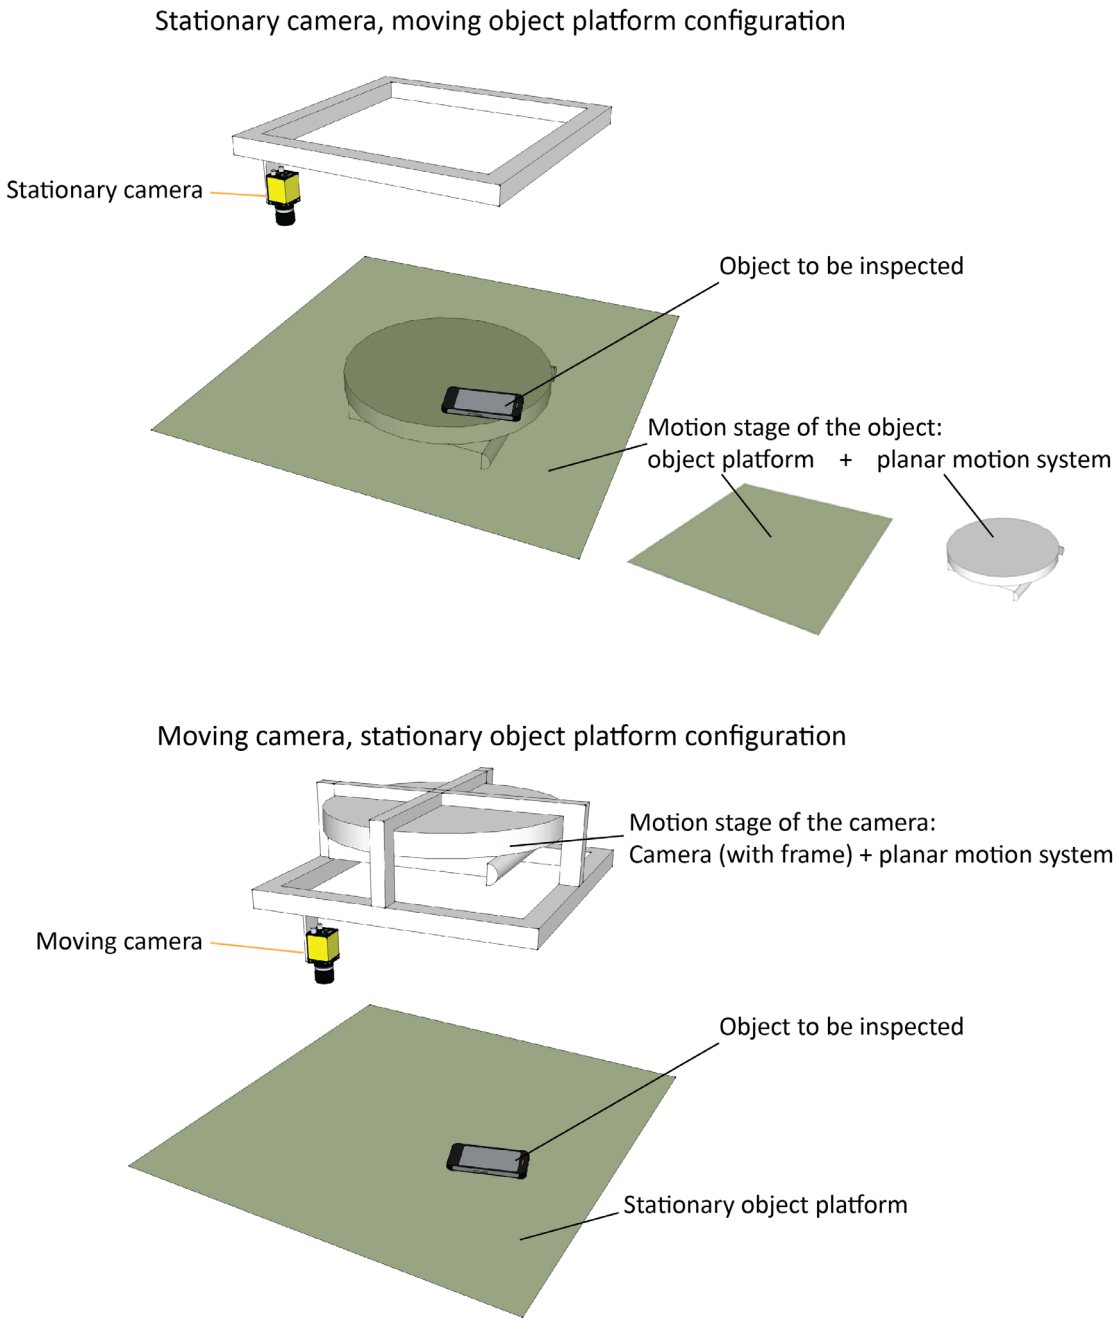

The motion system performs planar movement:

- It can move in the X and Y directions,
- It can perform Theta rotation based on your inputs.

The following figures demonstrate the X and Y movements and the Theta rotation of the motion stage.

Object's motion stage movement – stationary camera configuration

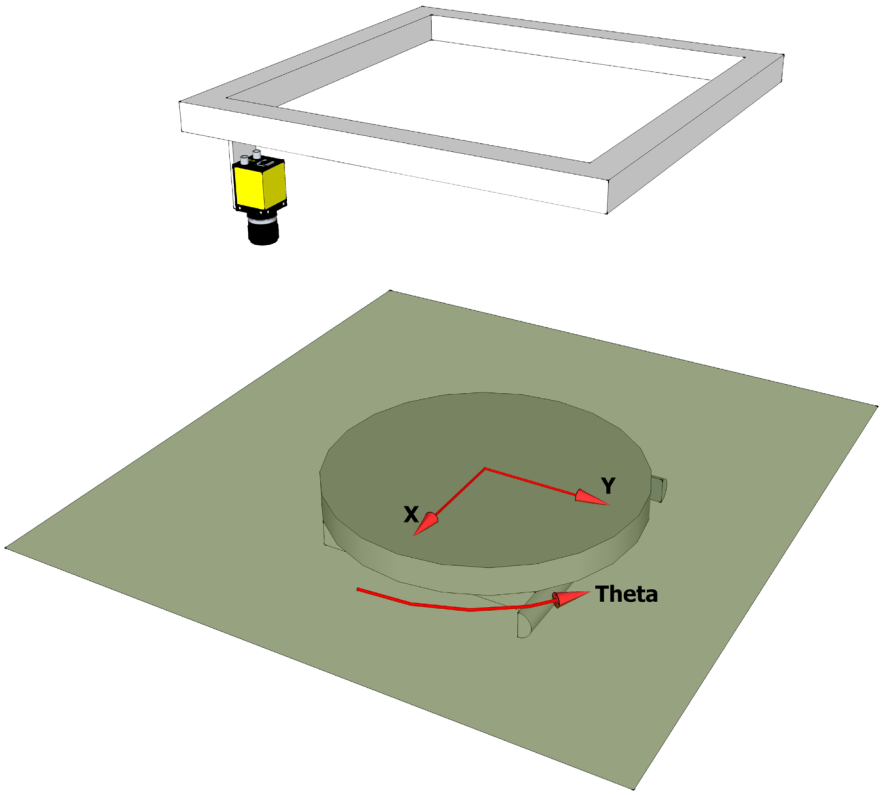

Camera's motion stage movement – moving camera configuration

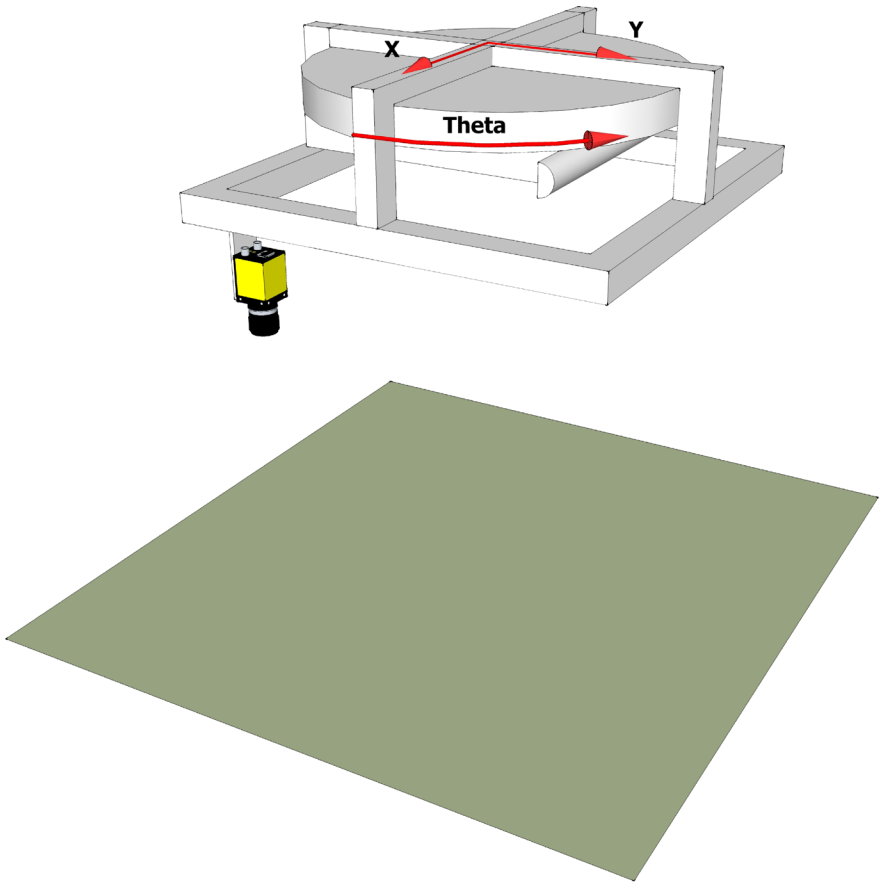

Movement is performed based on your commands and may have errors that you want to track and correct.

Hand-eye calibration performs the following main calibration tasks:

- Removes perspective distortion and lens distortion your images exhibit (image correction). This allows you to perform physically accurate measurements on the features in the images.
- Allows mapping features in your image to the native physical coordinate space of the motion system. This allows you to determine the exact positions of features in the images.
- Allows you to command the stage to the desired position based on calibrated images you obtained with your cameras.

You can use AlignPlus hand-eye calibration, for example, to align a smart phone frame by commanding calibrated position data to the motion system. The following one-camera calibration example shows this use case with commanded stage poses 1, 2, and 3.

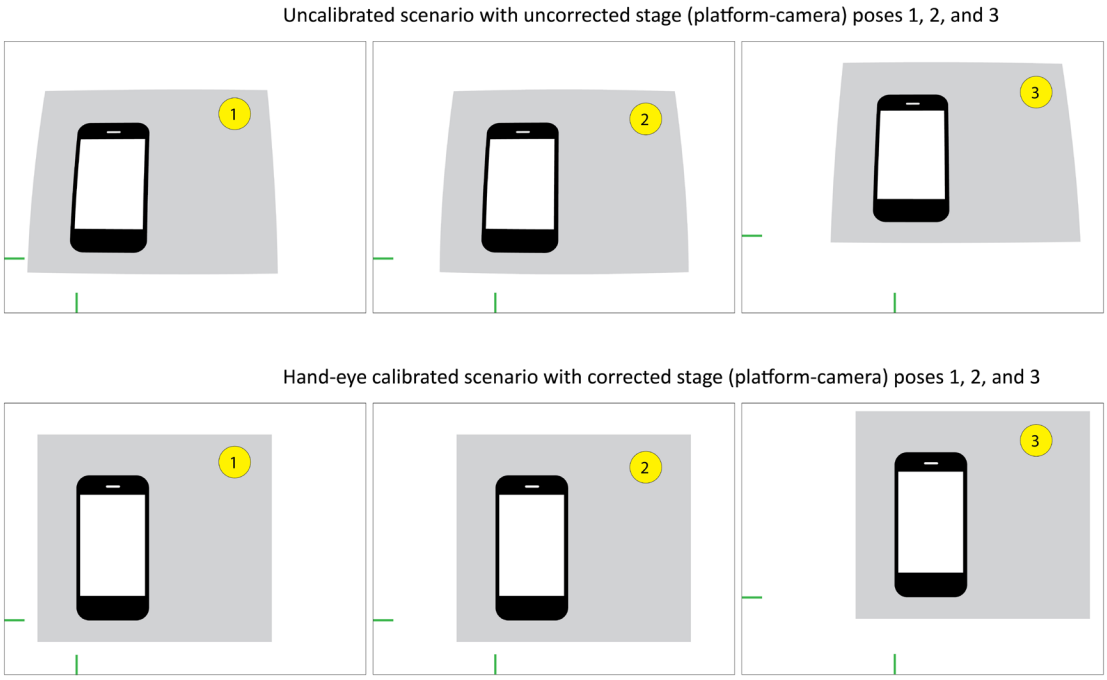

Green lines mark the coordinates of the desired stage (platform-camera) poses in real world coordinates. These are the poses to which you want to place your object to be inspected, that is, align your object. In this example, ideally, these poses mark the lower left corner of the smart phone. The uncalibrated scenario exhibits perspective distortion, camera lens distortion, and systematic stage positioning errors. In the hand-eye calibrated scenario, the poses you command get corrected using the calibration data and as a result your object moves to the right pose in real world coordinates.

How Hand-Eye Calibration is Performed

Train Time

Hand-eye calibration is performed by affixing a calibration plate on the object platform, typically a checkerboard calibration plate with data matrices or a fiducial mark, and moving the stage (either with the object platform or with the camera) while obtaining an image about the plate in each stage pose. Movement includes X and Y translation and Theta rotation.

The calibration plate has features with known dimensions in real world coordinates, these features are matched with those visible in the acquired image and these two types of data (physical location of features and their appearance in the images) are stored as correspondence pairs. The train-time calibration plate image gets corrected (unwarped) based on the relation between the two sets of data in the correspondence pairs, removing perspective distortion and lens distortion from the image (see more in the documentation for checkerboard calibration). Once the images have all been corrected, the commanded stage poses are compared with those visible in the corrected images and a mapping is defined. Calibration results include information to place run-time (or train-time) images in the native coordinate space of the motion system based on this mapping.

AlignPlus hand-eye calibration performs the following calibration tasks in addition to its main tasks:

- Determines the pose of each camera.
- Determines the pose of the calibration target.

Run Time

You can use the transform defined by the correspondence pairs and the pose mapping in the calibration results to calibrate run-time (or train-time) images.

You can use the calibration results to accurately command the motion stage to the desired position.

Calibration Process

Train-Time Calibration Process

The following figure demonstrates the train-time calibration process using a checkerboard calibration plate with a fiducial mark.

Moving stage calibration

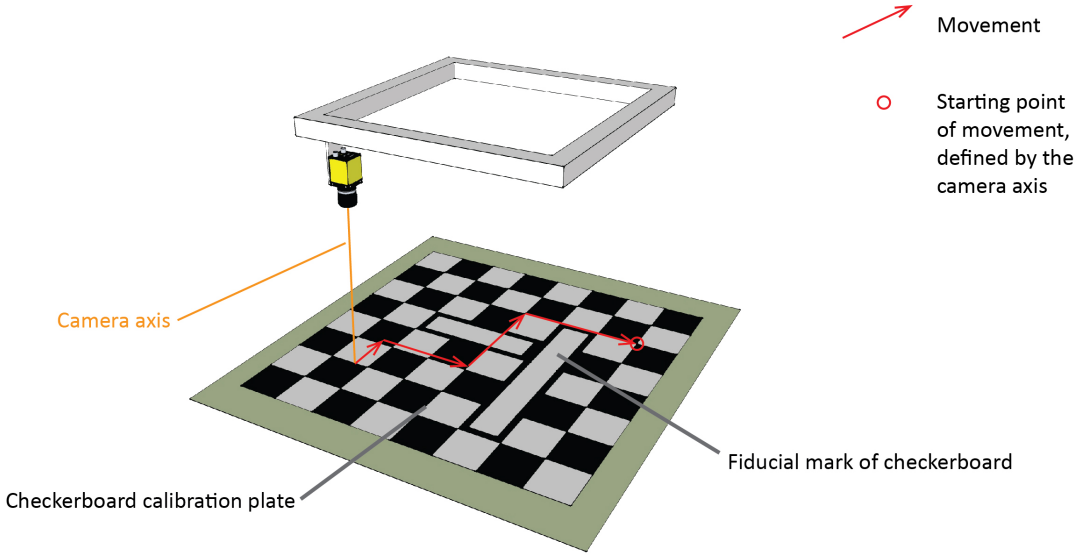

Moving camera calibration

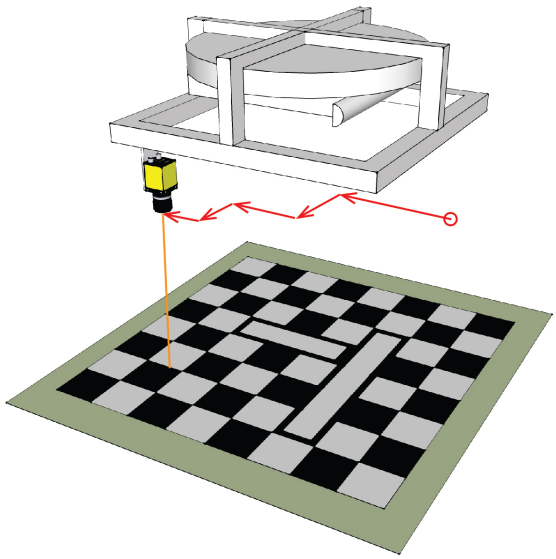

The following figure provides an overview on the train-time process of AlignPlus hand-eye calibration for the typical use case.

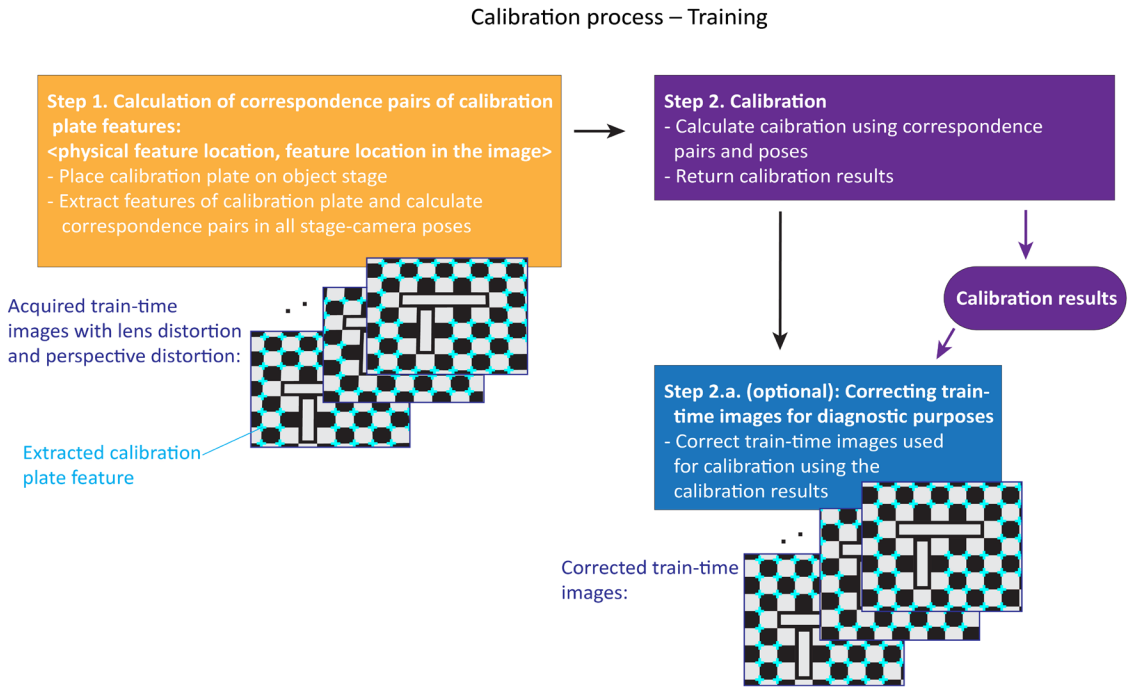

The multiple calibration images are acquired from various commanded movements with one camera as shown above (or with multiple cameras).

Cognex recommends using a checkerboard calibration plate with data matrix codes, especially for multi-camera applications. The following figure shows such a plate with four data matrix codes encoding positions on the plate.

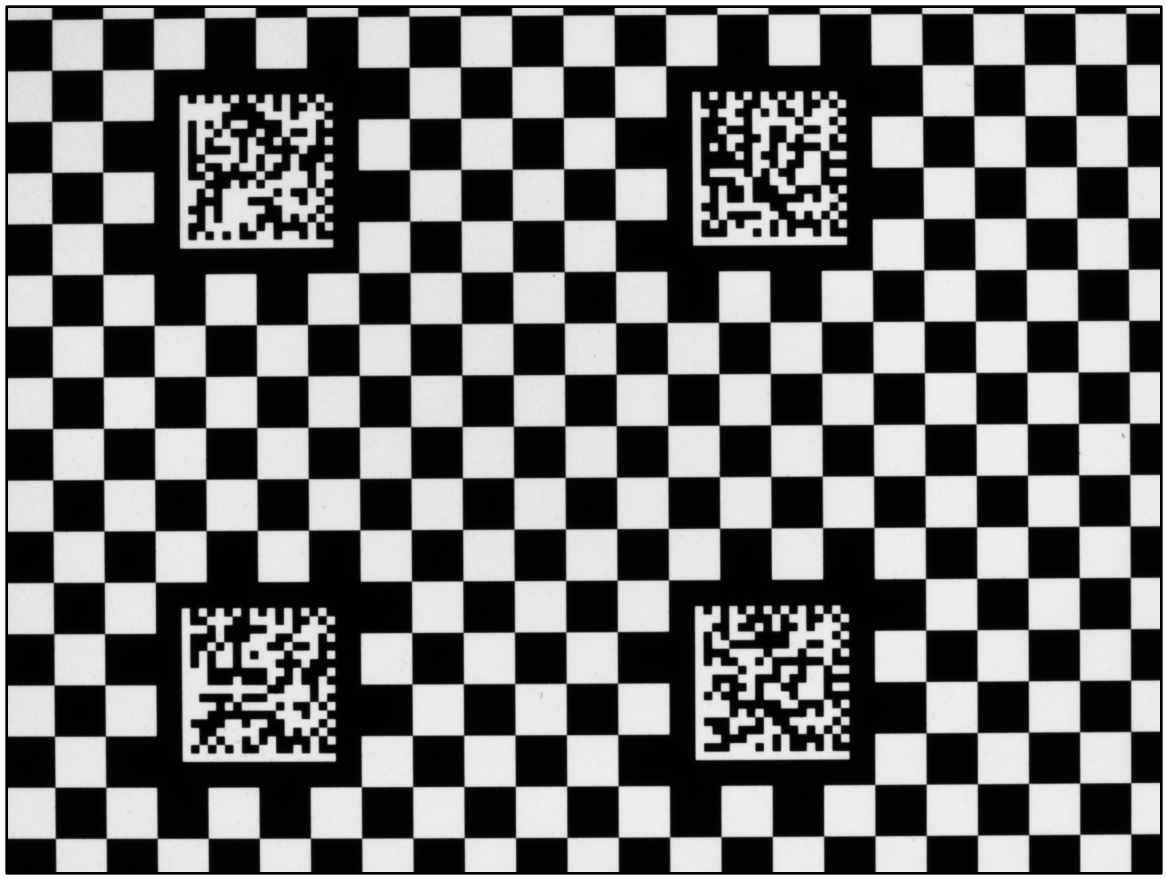

Run-Time Process Using Calibration Results

The following figure provides an overview on the run-time process of AlignPlus hand-eye calibration for the typical use case.

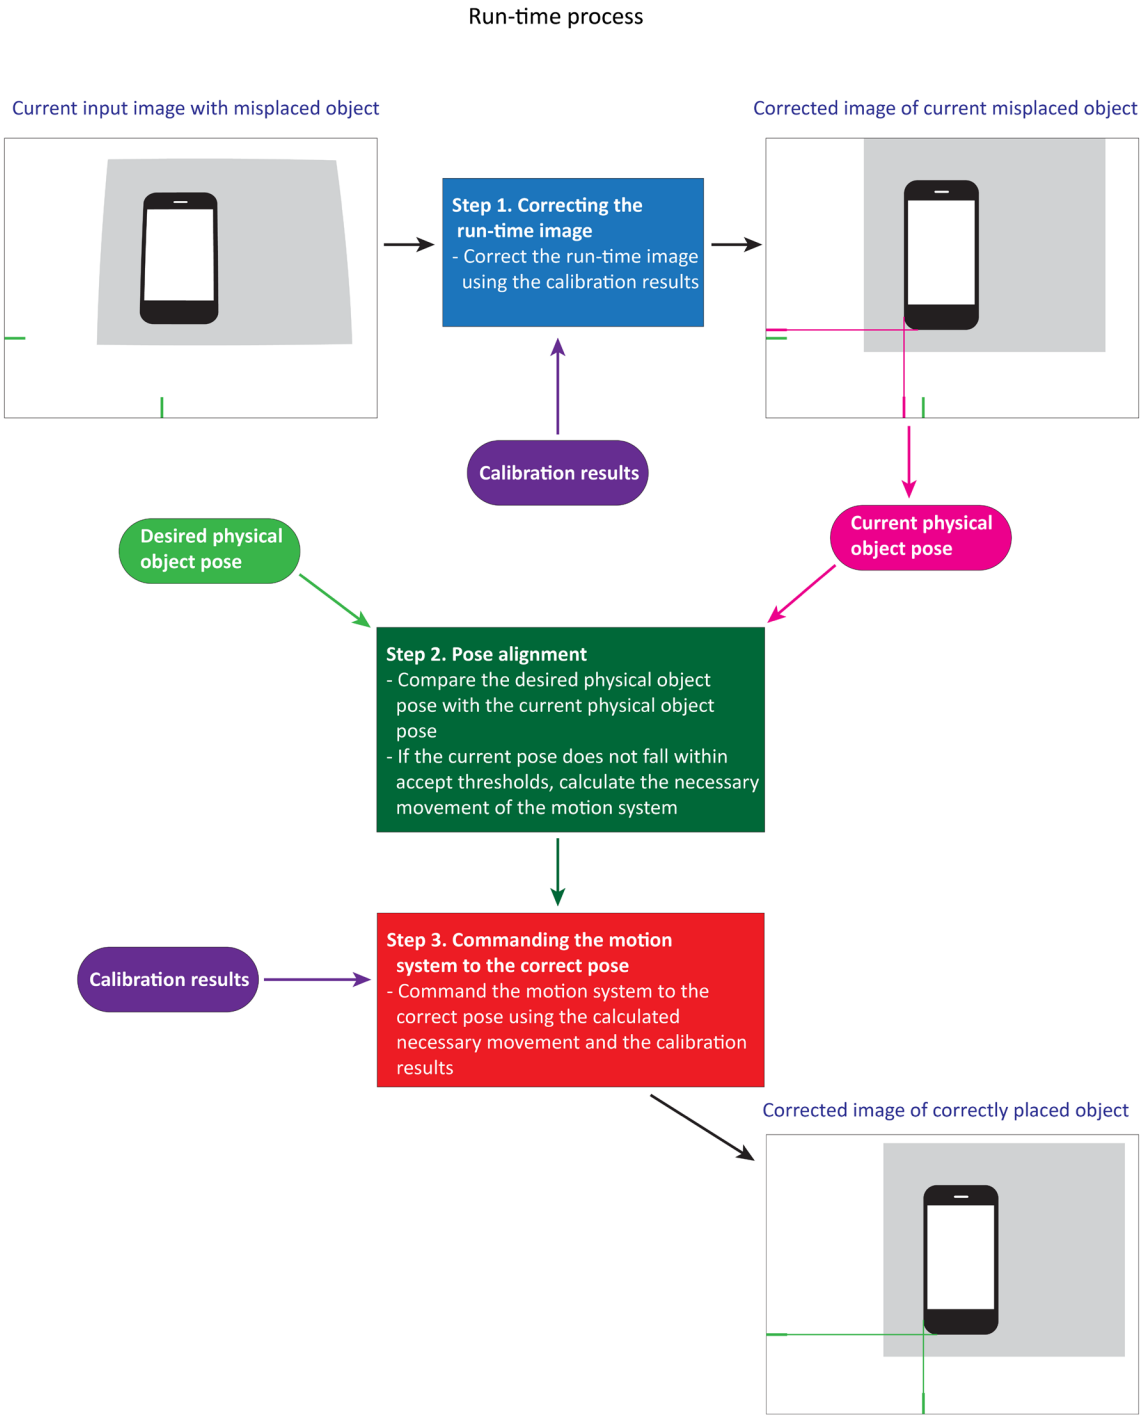

During image correction, the input image's distortions get corrected and the native coordinate space of the motion system is added to the image's coordinate space tree to generate the output image on which you can perform measurements.

Hand-eye Calibration Coordinate Spaces

Home2D is the native coordinate space of the motion system and Stage2D is the coordinate space of the motion stage moving along with it. Home2D and Stage2D are defined in the plane of the stage's motion. Plate2D is the coordinate space of the calibration plate and Camera2D is the coordinate space for a single camera. Plate2D and Camera2D are defined in the plane of the calibration plate.

In an appropriate physical setup, the plane of the object platform, the plane of the calibration plate (which is affixed to the object platform), and the plane of the stage's motion are the same (that is, they are parallel to each other). This plane is the plane of interest. Therefore, coordinate spaces associated with AlignPlus hand-eye calibration described in detail hereinafter are defined in the plane of interest.

Plane of interest marked with green-white stripes

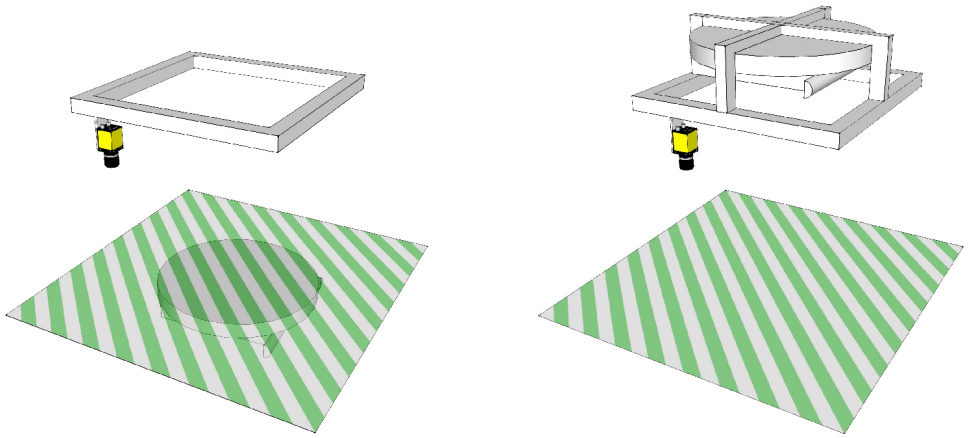

Stationary camera configuration

Moving camera configuration

Raw2D

Raw2D is the pixel space coordinate system of a single camera. There is one instance of such a coordinate system for each camera. Each camera's Raw2D coordinate system is independent of every other camera.

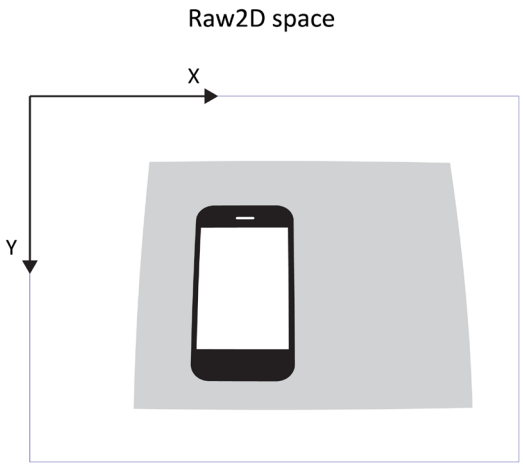

Acquired image before correction

Home2D

Home2D is the base reference space in which all other coordinate spaces and their relationships are described. Home2D is defined by the initial position of the motion stage. Home2D is defined by the X axis of the motion stage and the motion stage's center of rotation. The origin is at the stage's center of rotation when the stage is at the home position. The X axis of Home2D is perfectly aligned with the motion stage's X axis. The Y axis of Home2D is exactly 90 degrees from the X axis and in the general direction of the motion stage's Y axis.

Home2D is the native coordinate space of the motion system.

The following figure shows the Home2D coordinate space in red in the case of stationary camera configuration.

Home2D space – stationary camera configuration

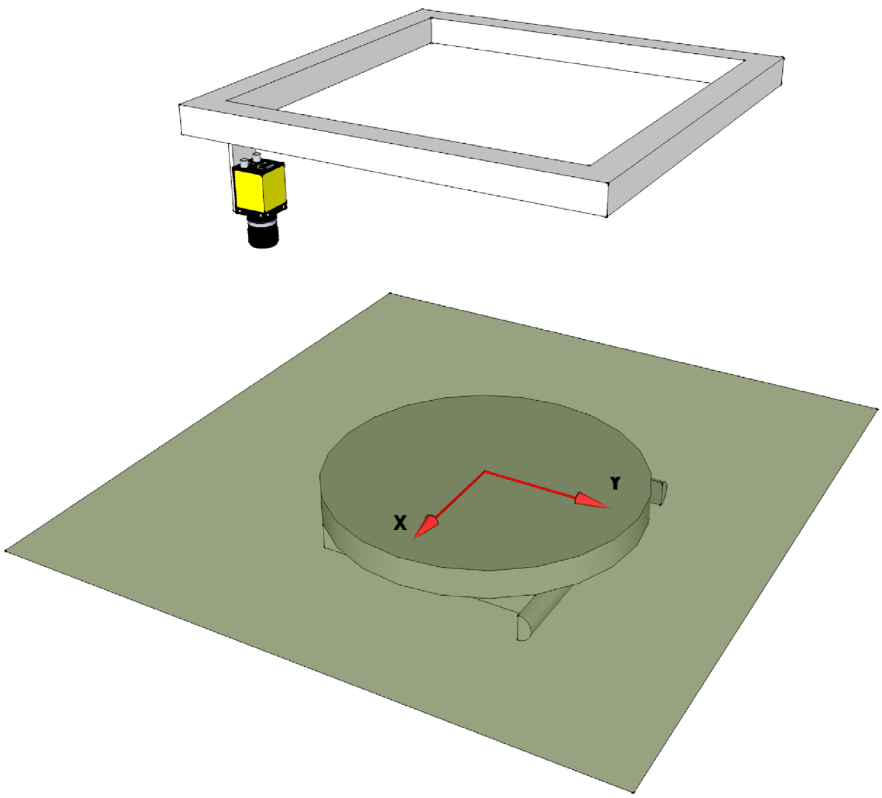

The following figure shows the Home2D coordinate space in red in the case of moving camera configuration.

Home2D space – moving camera configuration

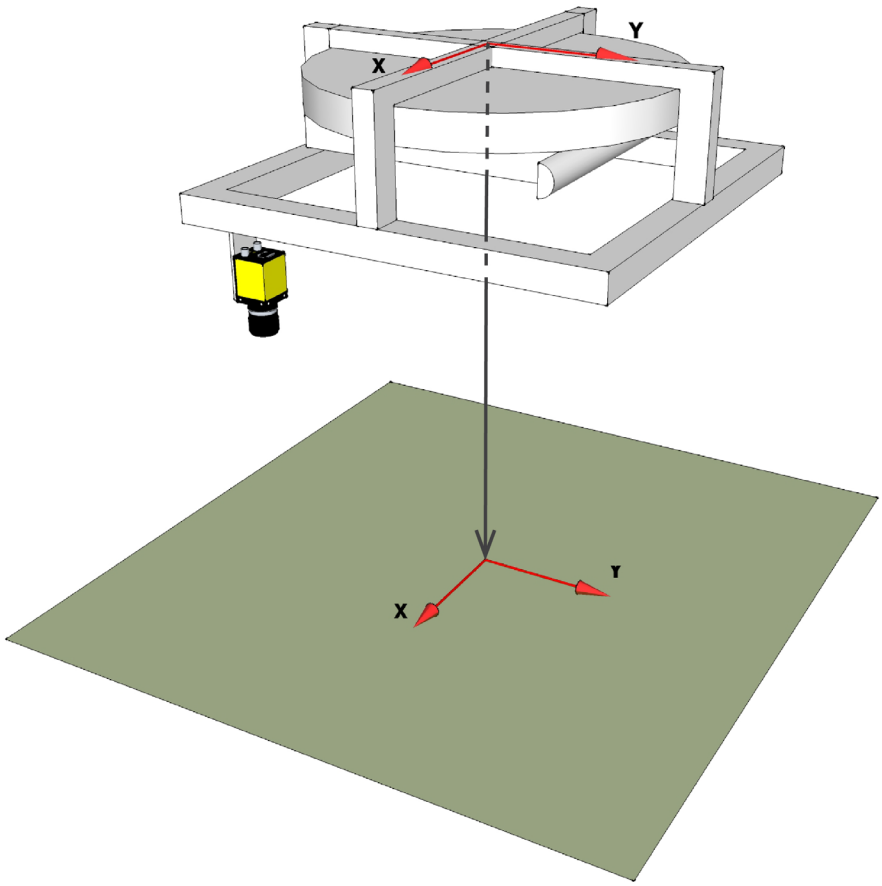

Stage2D

Stage2D is an orthonormal coordinate system that is attached to the motion stage's center of rotation, and moves and rotates along with the motion stage. In the initial position of the stage (that is when X and Y translations, and Theta rotation are zero), Stage2D coincides with Home2D.

The following figure shows the Stage2D coordinate space in blue in the case of stationary camera configuration.

Stage2D space – stationary camera configuration

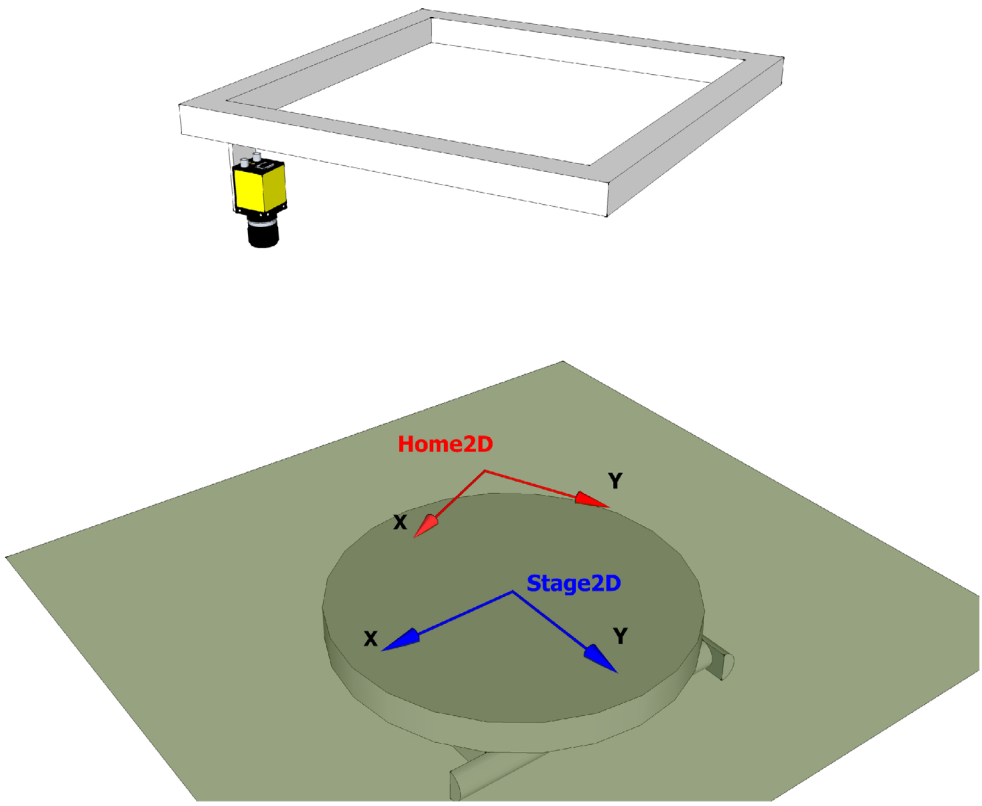

The following figure shows the Stage2D coordinate space in blue in the case of moving camera configuration.

Stage2D space – moving camera configuration

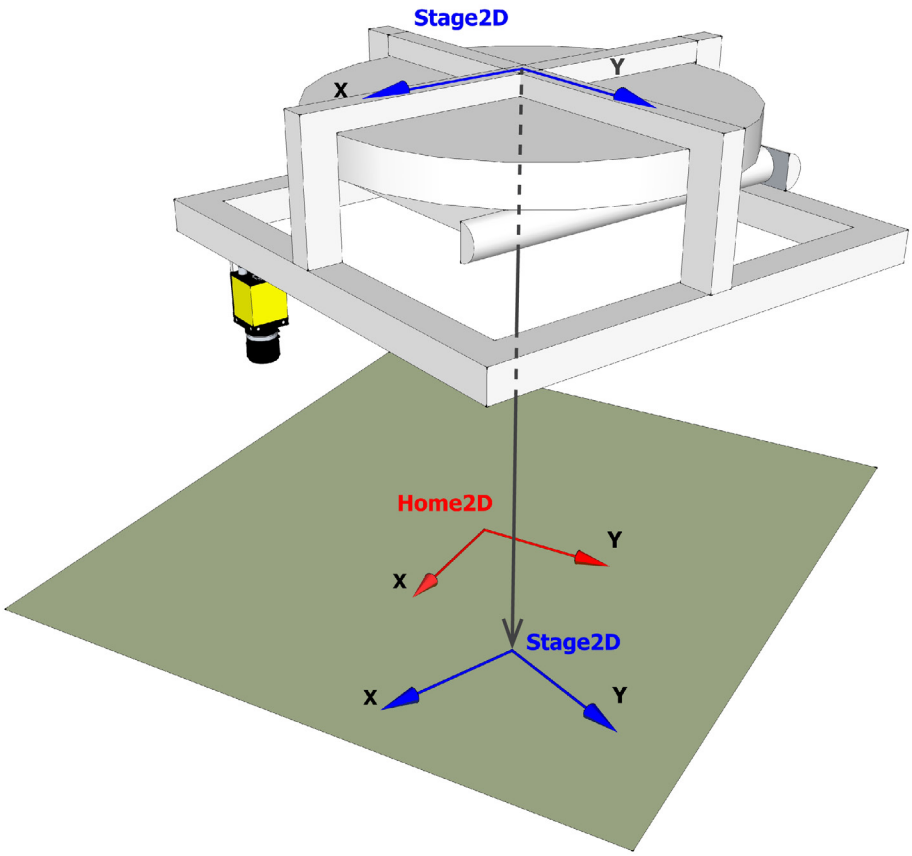

Plate2D

Plate2D is the (global) coordinate system of the calibration plate. At any instance of time, all calibration features viewed by all cameras are described in the same Plate2D coordinate system. However, this space may be moved around in the Home2D coordinates by the motion stage. Plate2D is an orthonormal coordinate system, although its length unit may have a non-identity scale factor from Home2D. Plate2D and Home2D may have different handedness.

The origin of the Plate2D space is defined either by the fiducial mark or the data matrix codes printed on it. See the documentation on checkerboard calibration for details.

The following figure shows the Plate2D coordinate space in dark purple with pink guides to the fiducial mark features the coordinate space is derived from (in the case of stationary camera configuration).

Plate2D space – plate with fiducial mark (stationary camera configuration)

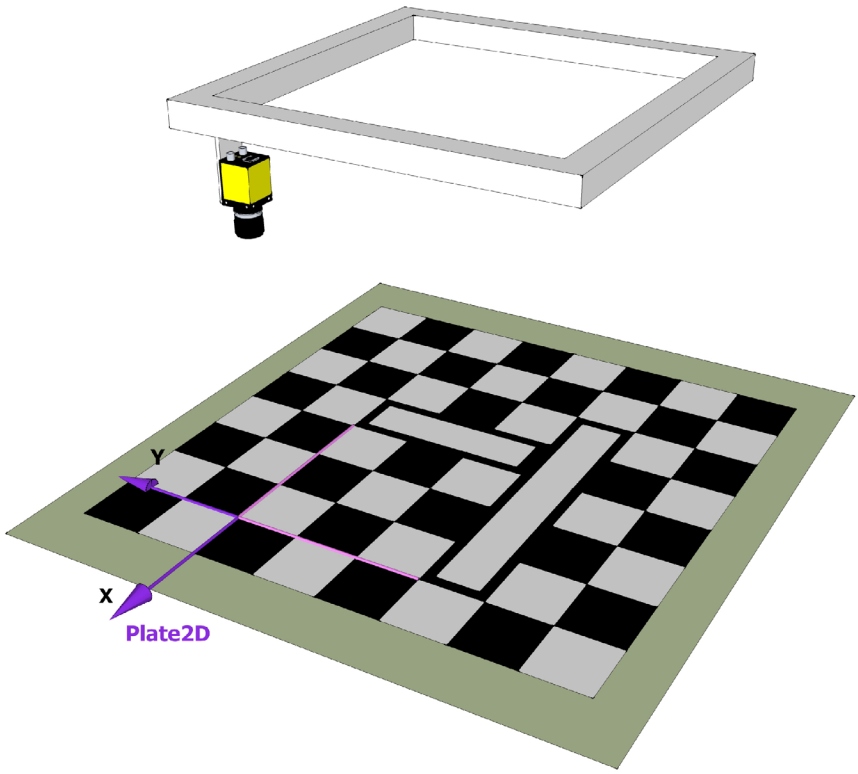

The following figure shows the Plate2D coordinate space for the previously mentioned checkerboard calibration plate with data matrix codes. The positions marked by the data matrix codes are also shown.

Plate2D space – plate with data matrix codes

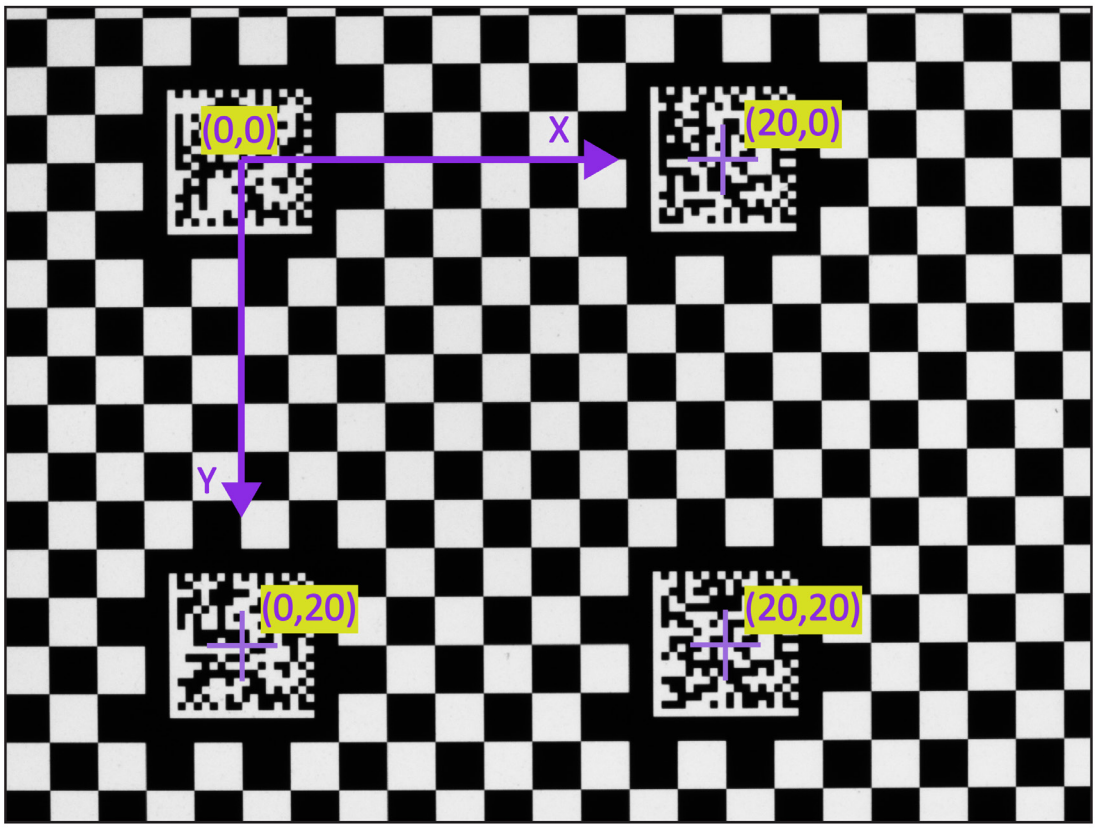

Notes:

- As stated earlier, the calibration plate must be parallel to the stage's motion.
- The calibration plate should be in the plane of the part's features that are used for alignment.

Camera2D

Camera2D is the physical orthonormal coordinate system for a single camera. There is one instance of such a coordinate system for each camera. The origin of Camera2D is at a position in Home2D that corresponds to the center of the camera's image window. Its X axis is parallel to and points in the same direction as the Raw2D X axis. Its Y axis is perpendicular to the X axis and points in the general direction of the Raw2D Y axis.

Camera2D space – moving camera configuration

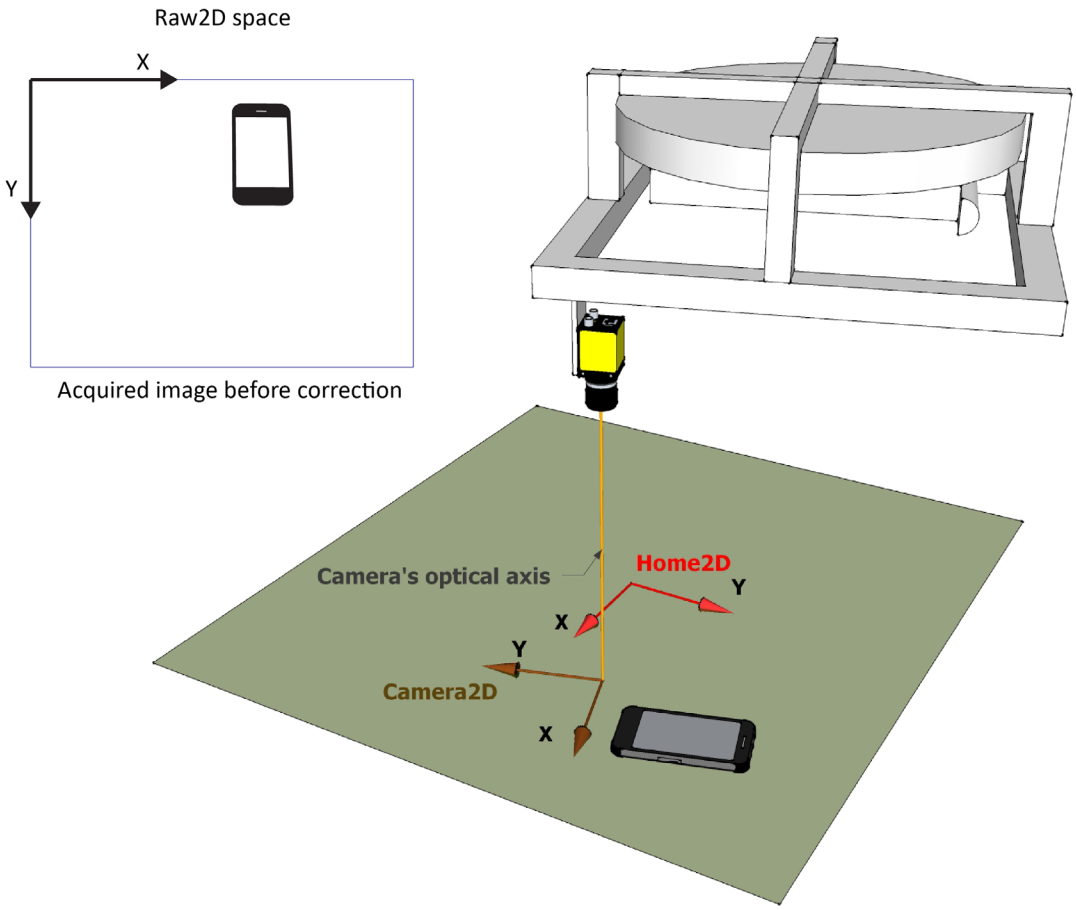

Handedness of Coordinate Spaces

Independent handedness is assumed among Home2D, Camera2D, and Plate2D. The Home2D handedness is determined by the directions of the motion stage's axes. The Plate2D handedness comes with the calibration plate artwork design. The Camera2D handedness is determined by the number of mirrors in the optical path. These are naturally independent of each other, and are treated as such.

Transforms

The following transforms are associated with the previously defined coordinate spaces.

General Transforms

These transforms are valid both in stationary and moving camera configurations.

Home2DFromStage2D

This is the actual physical pose of the motion stage in Home2D. It defines the relationship between Home2D and Stage2D and is represented by a 2D rigid transform.

(In the stationary camera configuration, this is the motion of the motion stage of the object. In the moving camera configuration, this is the motion of the motion stage onto which the cameras are attached – the object platform is stationary in this case.)

UncorrectedHome2DFromStage2D

This is the commanded pose of the motion stage.

Due to systematic errors in the motion stage, the commanded pose of the motion stage may not match the actual physical pose of the motion stage in Home2D (that is, Home2DFromStage2D). UncorrectedHome2DFromStage2D is used to describe the commanded pose of the motion stage because it represents your best guess of the pose of the motion stage prior to calibration.

Raw2DFromHome2D

The mapping between Raw2D and Home2D. In the moving camera configuration, this transform is dependent on the actual UncorrectedHome2DFromStage2D pose.

ConvertUncorrectedHome2DFromStage2DToHome2DFromStage2D() and ConvertHome2DFromStage2DToUncorrectedHome2DFromStage2D() Conversion Functions

These methods are conversion functions: one returns the Home2DFromStage2D transform corresponding to the input UncorrectedHome2DFromStage2D pose and the other one returns the UncorrectedHome2DFromStage2D pose corresponding to the input Home2DFromStage2D transform.

Raw2DFromCamera2D

The transform that maps coordinates in the camera coordinate system (Camera2D) to the image coordinate system (Raw2D).

Stationary Camera Configuration transforms

These transforms are valid in the stationary camera configuration only.

Home2DFromStationaryCamera2D

This is the placement pose of the camera, that is, this specifies how the camera was placed in the system. The placement pose of the cameras is necessary for mapping a physical position in camera coordinates to Home2D. This is a rigid transform that may also flip handedness.

Stage2DFromMovingPlate2D

This is the placement pose of the calibration plate, that is, this specifies how the calibration plate was placed in the system. The placement pose of the calibration target is necessary for mapping a physical position on the calibration target to Home2D. This is a rigid transform that may also flip handedness.

Moving Camera Configuration transforms

These transforms are valid in the moving camera configuration only.

Home2DFromStationaryPlate2D

This is the placement pose of the calibration plate, that is, this specifies how the calibration plate was placed in the system. The placement pose of the calibration target is necessary for mapping a physical position on the calibration target to Home2D. This is a rigid transform that may also flip handedness.

Stage2DFromMovingCamera2D

This is the placement pose of the camera, that is, this specifies how the camera was placed in the system. The placement pose of the cameras is necessary for mapping a physical position in camera coordinates to Home2D. This is a rigid transform that may also flip handedness.

Transforms Overview

The following figures show the relationships between the various coordinate spaces and transforms for multi-camera configurations.

Coordinate spaces and transforms – stationary camera configuration

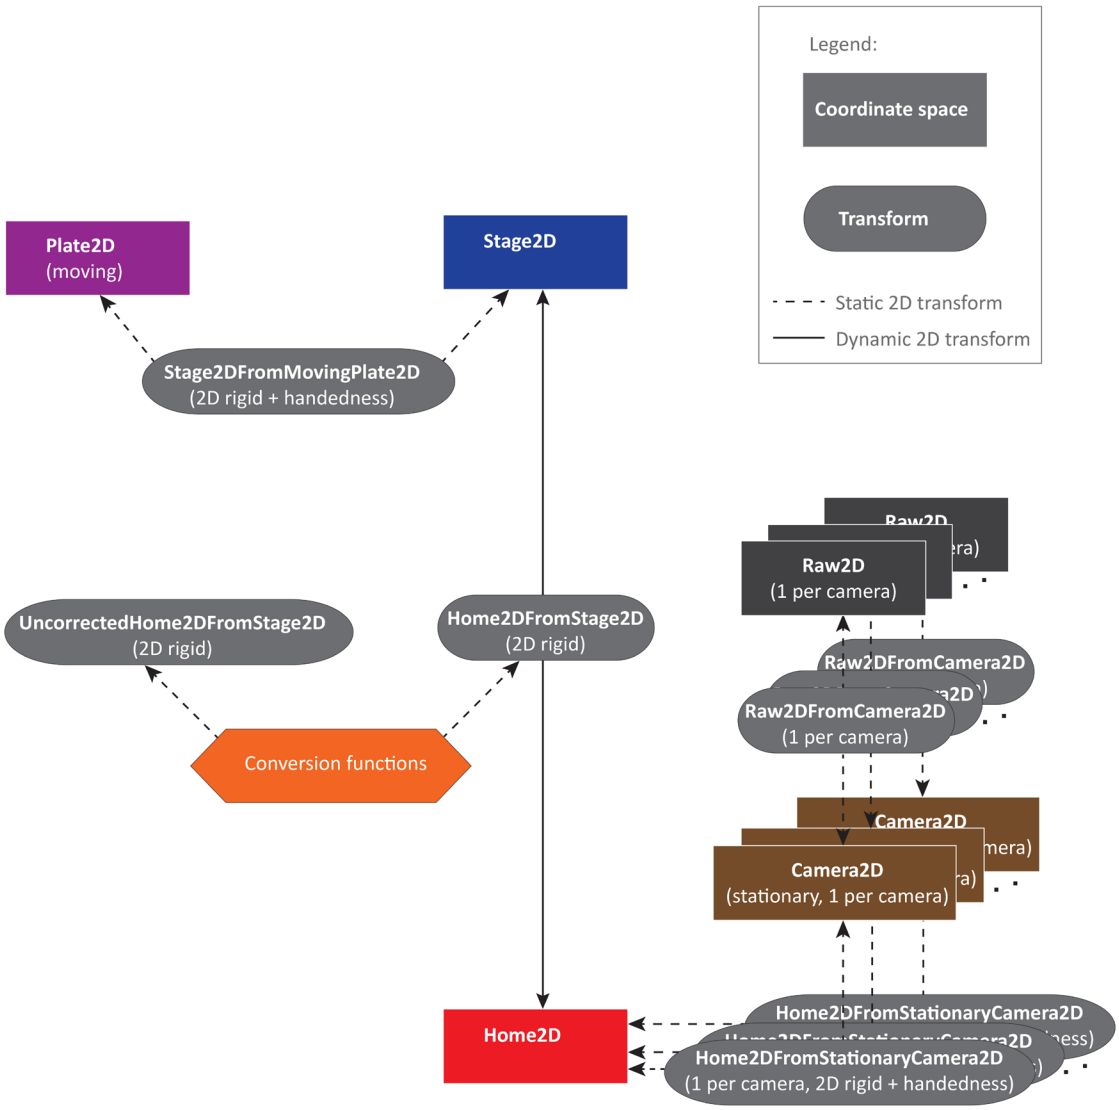

Coordinate spaces and transforms – moving camera configuration

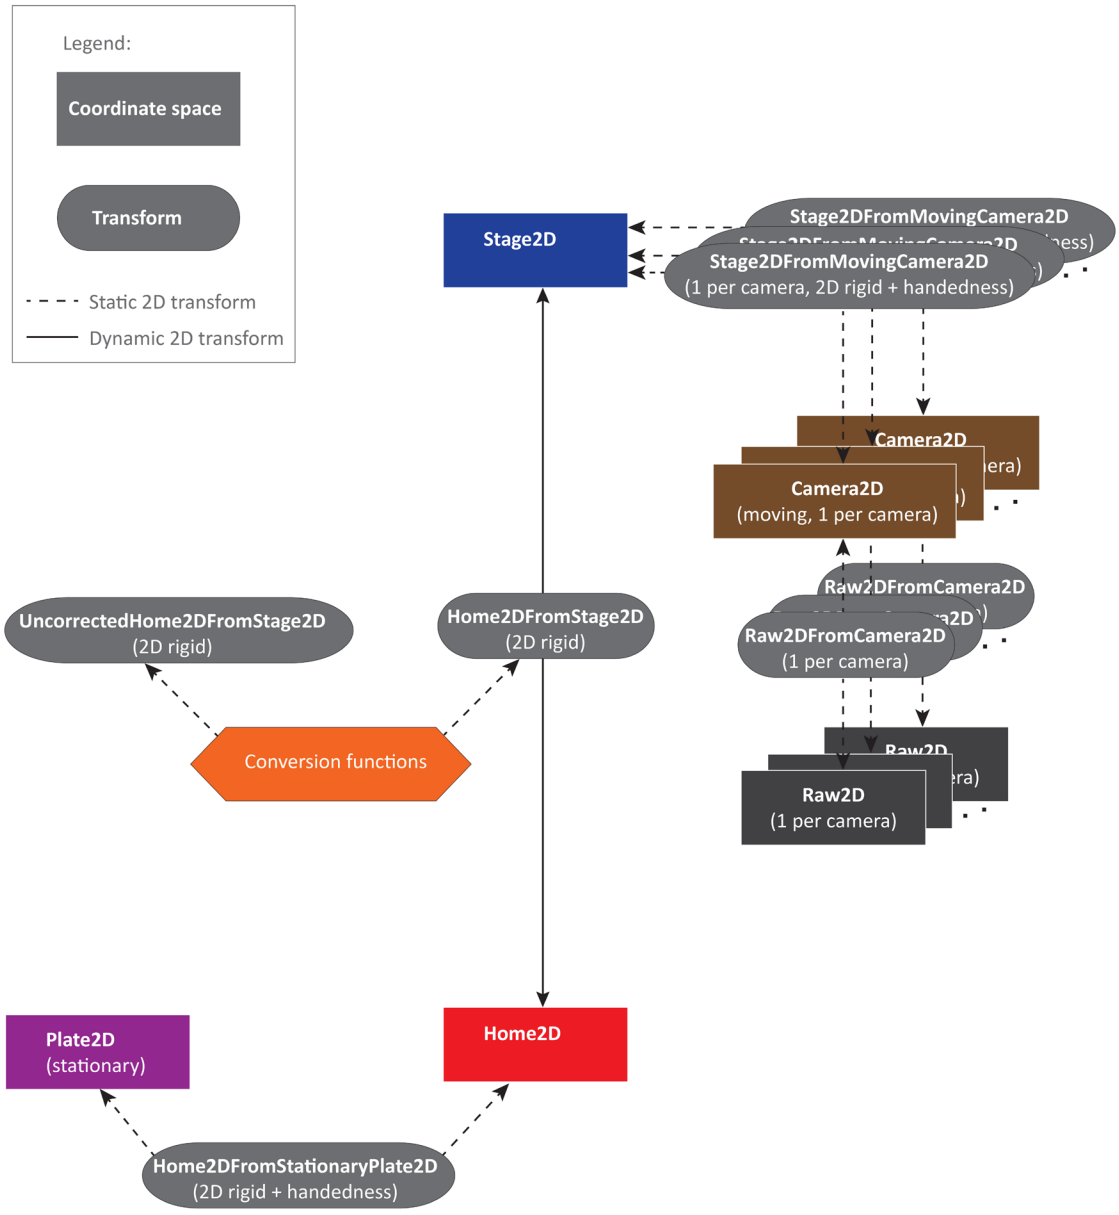

Motion Stage Validation

Motion stage validation gives you basic feedback on your stage and helps identify problems before calibration.

The CogMotionStageValidator class performs motion stage validation prior to performing a hand-eye calibration of the stage in a machine vision system with one or more cameras. The purpose of this class is to verify that the stage moves to its commanded poses (X, Y, Theta), and to characterize certain types of systematic errors in the observed motion.

Cognex strongly recommends performing stage validation before performing hand-eye calibration.

Similar to the hand-eye calibration usage model, this tool uses correspondence data extracted from all views (images) from all cameras, along with the UncorrectedHome2DFromStage2D pose associated with each view. The tool validates the motion stage for all sets of UncorrectedHome2DFromStage2D poses that contain motion corresponding to the metric requested for in the input parameters. The following metrics are available:

- XScale
- YScale
- ThetaScale
- Skew between the X and Y motion axes

You can use the same set of correspondence pairs for both stage validation and hand-eye calibration.

For details on motion stage validation, see the description in the CogMotionStageValidator class in the online documentation.

What AlignPlus Hand-Eye Calibration Does

AlignPlus hand-eye calibration corrects the input image distortions, the correction is performed by CogCalibImageCorrector. Also, it computes the mathematical transform to map image coordinate positions to and from their corresponding coordinates in the coordinate system defined by the motion axes of the stage, that is, the axes of the native coordinate space, which is the Home2D space. This is performed by CogHandEyeCalibrator. This is a 2D transform. No 3D information is provided. After image distortion correction, CogCalibImageCorrector adds the Home2D coordinate space to the corrected image's coordinate space tree based on this transform to generate the output image.

In the case of stationary camera configuration, for you to be able to align your objects visible in your cameras, you need a mapping between the Raw2D pixel space and the Home2D physical space. AlignPlus hand-eye calibration provides this as the explicit Raw2DFromHome2D transform and by adding Home2D to the corrected image's coordinate space tree. This way, the output image is ready for you to consume.

In the case of moving camera configuration, for you to be able to align your objects visible in your cameras, you need a mapping between the Raw2D pixel space and the Home2D physical space. You use the Raw2DFromHome2D transform to provide the needed mapping. This mapping takes into account the current stage position.

Motion Capability Limitations

In order to fully determine the relationship between Raw2D and Home2D, the motion stage or robot arm must be capable of performing both translation and rotation. For motion stages or robots that can perform only translation or only rotation, the tool can recover only part of the relationship between the two coordinate systems. Translational motion provides information to recover the rotation angle between the axes of the two coordinate systems. Rotational motion provides information to recover the translation between the origins of the two coordinate systems.

The CogHandEyeCalibrator.MotionCapability property specifies this information for the tool.

For motion stages or robots capable of only translations, the tool arbitrarily specifies the plate placement pose's translation vector to be (0,0). The consequence is that the origin of Home2D coincides with the origin of Plate2D. This is true for both stationary-camera and moving-camera configurations. To operate the tool in this mode, you need to specify CogHandEyeMotionCapabilityConstants.TranslationOnly1Axis or CogHandEyeMotionCapabilityConstants.TranslationOnly2Axes for the CogHandEyeCalibrator.MotionCapability property.

For motion stages or robots capable of only rotations, the tool arbitrarily specifies the rotation angle of the plate placement pose to be 0. The consequence is that the X-axis of Home2D is

parallel to that of Plate2D when the stage is at home. This is true for both stationary-camera and moving-camera configurations. To operate the tool in this mode, you need to specify CogHandEyeMotionCapabilityConstants.RotationOnly for the CogHandEyeCalibrator.MotionCapability property.

For motion stages or robots capable of both translation and rotation, the tool can automatically recover the full relationship between the coordinate systems. This requires at least three UncorrectedHome2DFromStage2D poses. Among these poses, there must be at least one pair with different translation components. In addition, there must be at least one UncorrectedHome2DFromStage2D pose with rotation components that differ by at least some minimum angle. To operate the tool in this mode, you need to specify CogHandEyeMotionCapabilityConstants.RotationAndTranslation for the CogHandEyeCalibrator.MotionCapability property.

See more in the CogHandEyeMotionCapabilityConstants enumeration description in the online documentation.

Fundamental Requirements and Residuals

In order for the tool to function properly, the following fundamental requirements must be satisfied.

Fixed Placement Pose of the Calibration Target and Cameras

In the stationary camera configuration, the calibration target must be rigidly attached to the motion stage. The attachment point on the target and the orientation of the target are arbitrary. Similarly, in the moving camera configuration the cameras must be rigidly attached to the motion stage, and the attachment point of the cameras and their orientations are arbitrary. In both configurations, once attached, the attachment point and the orientation must remain unchanged throughout the calibration procedure.

As a result of rigidly attaching the calibration target (or cameras), the placement pose remains unchanged as the stage moves the target (or cameras) around. The placement pose is a 2D rigid transform that may also flip the handedness. In most applications, the placement pose cannot be precisely controlled and is unknown prior to running the calibration tool. The only requirement is that the placement pose remains unchanged throughout the calibration procedure.

Precision of Motion Stage

For the tool to function properly, the motion stage must be precise and accurate. Calibration result accuracy may be affected by the precision and accuracy of the motion stage. Calibration results are generally more accurate if the motion stage is more accurate.

The tool estimates the skew and aspect of the motion stage axes. However, in general, the overall calibration is more accurate when the angle between the MotionXAxisHome2D and

MotionYAxisHome2D is close to 90 degrees, and the length of MotionXAxisHome2D is close to the length of MotionYAxisHome2D.

In-Plane Motion

This tool solves a 2D calibration problem. The calibration pattern must be planar, and must remain in the same physical plane while the calibration target is moved to various places in the field-of-view of each camera. The calibration target must be parallel to the stage's motion.

Each camera's optical axis does not need to be perpendicular to this plane. The initial in-plane orientation of the calibration pattern relative to the camera's image coordinate axes, as defined by the placement of the calibration target onto the object platform, can be arbitrary.

Indications of Violations to Requirements

Two things in the tool's result can be used as indicators of whether the above requirements are satisfied, namely residual fitting errors and an internal estimation of the Home2DFromStage2D poses of the calibration target as suggested by image data. Any mismatch between the estimated Home2DFromStage2D poses and the specified input UncorrectedHome2DFromStage2D poses indicates inconsistency in the specified input UncorrectedHome2DFromStage2D poses or a systematic error in the motion stage, and is most likely accompanied by large residual errors. This may be caused by any of the following, which you need to investigate in order to achieve accurate calibration:

- Slippage of the calibration target (violation to requirement about the fixed placement pose of the calibration target and cameras)
- An imprecise motion stage (violation to requirement about the precision of the motion stage)
- Tilt of the calibration target (violation to requirement about in-plane motion)
- An inaccurate calibration pattern on the calibration target

Residual Errors

In both stationary camera and moving camera modes, the tool's results also include measurement of residual errors, which are indicators of overall quality of the calibration. A good calibration should yield small residual errors. Large residual errors usually indicate bad inputs, examples of which include damaged and/or inaccurate calibration targets, inaccurate feature extraction, wrong correspondence, inaccurate motion stage, wrong UncorrectedHome2DFromStage2D pose, inappropriate lens distortion model, and violations to fundamental requirements listed previously.

The tool computes two sets of residuals:

- Single-view residuals, which are computed from the single-view calibrations for each view of each camera. These residuals are computed in Raw2D and Plate2D, by mapping

the correspondence pairs for each view of each camera through the corresponding single-view calibration. Note that the single-view calibrations are not included in the result, only their residuals are. These residuals indicate the quality of the calibration target, feature extractor, correspondence pairs and the selected distortion model. The Raw2D residual errors are in units of pixels, and the Plate2D residual errors are in the physical units of the calibration target. Generally, a residual error of less than 0.1 pixels in Raw2D is considered excellent calibration accuracy.

- Overall residuals, which are computed from the entire hand-eye calibration. They are computed in Raw2D and Home2D by mapping the correspondence pairs through the entire hand-eye calibration (that is, Raw2DFromHome2D). The mapping uses both directions of Raw2DFromHome2D to go to and from Raw2D and Home2D. The unit of the residual errors is pixel in Raw2D and physical units in Home2D. Generally, a residual error of less than 0.25 pixels in Raw2D is considered excellent calibration accuracy.

Note that calibration transforms are known to behave poorly in extrapolation. In other words, for positions within the area spanned by the cloud of input feature points, the accuracy of the calibration transform can be judged by the residual errors. However, for positions away from the point cloud, the accuracy of the calibration transform degrades very quickly. Therefore, Cognex strongly advises against using calibration results on regions beyond what is spanned by the input feature points.

Multi-Camera Hand-Eye Calibration

You can hand-eye calibrate images acquired from multiple cameras. The cameras have to be affixed rigidly to the stationary camera frame in the case of stationary cameras configuration or to the motion stage in the case of the moving cameras configuration. The cameras' relative position to each other must not be altered after calibration training has been performed.

Calibration is performed using one big calibration plate that covers the field of view of all cameras in all calibration positions. Calibration is performed independently for each camera; however, the calibrated images from each camera share the same physical coordinate space. This can be achieved by using a calibration plate that has multiple marks spread throughout its surface, marking physical positions. Such marks can be data matrix codes for example.

Sample application

You can build up your own application based on the AlignPlus hand-eye sample application available here:

http://www.cognex.com/Support/VisionPro/