

System_calibration

The calibration process involves collecting camera data of the moving wand, ensuring that the Calibrate parameters have desired values, and running the Calibrate program from the Function group in the AMASS shell. The calibration does not require any parameters of a technical nature, and the ones required are mostly self explanatory.

During calibration the program determines distortion corrections for the image of each camera, but it can only do so for parts of the image containing wand data that is also recorded by at least one other camera. Hence parts of the image where no wand markers are recorded, or where wand markers are recorded only by that particular camera, cannot be calibrated. The internal parameters for each camera include a mask that specifies the calibrated parts of each camera's image. Any marker image falling outside of the area defined by the mask will be discarded when the calibration parameters are applied to tracking. Extrapolation of the calibrated region is not supported because the errors become very large very quickly, and defeat one of the primary goals of the software. The fact of the matter is that *we cannot calibrate parts of the image for which we have no data*.

During the calibration the wand is identified by the software, provided the calibration is sufficiently accurate. If a wand cannot be identified it means that the input camera data were not sufficiently complete to enable an accurate calibration, or that the markers on the wand have become twisted and the wand is not recognized by the software (See [Marker Alignment](#)). In that case the calibration is aborted, ensuring that a successful calibration is always a good calibration.

Note. If you have trouble with the calibration process failing, please contact support@c-motion.com and send the .cen files that you collected for the calibration. Without these cen files it is very difficult to determine the cause of the calibration failure.

Parameters

The *Calibrate* page displays all user modifiable parameters for the calibration. On inspection it will be seen that not one of the parameters refers to the calibration computation itself, relieving the user from making difficult and confusing decisions with respect to the complex calculations that are involved in the calibration process.



mm /unit

In AMASS all distance measures are expressed in terms of a single unit of length whose size is specified through this parameter. The user should enter their size of the desired length units in millimeters. For example, if it is desired to have all output in centimeters then the required value is 10.0. The default unit length is millimeters, with a value of 1.0.

Units name

This entry allows the user to enter the name of the length units being used. It accepts a name length of up to 20 characters, and the default name is millimeters (mm).

Frame Increment

In situations where a large number of frames are collected for a calibration, or in captures involving high frame rates, using every frame in the calibration calculation may be inappropriate. In that case the user can choose to use fewer frames from the *.cen* file, e.g. every second or every third frame. For example, entering a value of 5 will result in every 5th frame being used for the computation.

Maximum frames

This entry specifies the maximum frame number to be used for the calibration. An entry of zero implies all frames recorded in the *.cen* file. Note that the number of frames actually used will be influenced by the *Frame Increment* entry.

Cameras to use

Occasionally it may be useful to exclude one or more cameras from the calibration because a camera is malfunctioning, a poorly calibrating camera is influencing the calibration of the other cameras, or you may wish to investigate the contribution of a particular camera on the calibration. The check boxes in this group can be used to ignore data from any camera during the calibration.

ViewLin

Check this box if you want ViewLin to run automatically after every calibration.

Reference marker file

This entry holds the full name of the file that contains the x, y, and z coordinates of the static Reference markers that will be used to translate the calibration data to the desired Global Coordinate System. You may browse for an existing *.ref* file or create a new one.

To create a new Reference marker file select *Browse...*, and in the browse window type in a file name, select a destination folder, and then click *Open*. A window will pop up asking if you wish to create a new file. If you answer Yes a NotePad window will open allowing you to enter the coordinates of your reference markers. After the comment line type in the x, y, z coordinates of each marker, one marker per line, and then save the file.

The location of each marker is specified in the “desired” laboratory coordinate system.

Each row of the ref file contains the x y x components of the marker.

There is no specified order to the rows, which means that you can list the markers in any order that you want. AMASS creates a laboratory coordinate system, such that the markers are located at the specified x y z values.

If, for example, the centroids of all markers are at a height of 20 mm, the z-component of each row should have the value 20.

No Reference Marker File

The system will calibrate without any reference marker file, i.e. .ref file and provide a default global system defined by the first three camera locations. Any time there is a ref file with recognized marker distribution, the calibration and subsequent data are output in the ref file coordinate system.

To clear the file name select *Browse...* followed by *Cancel* in the browse window.

Example .ref file

An example reference file for an L-Frame that has four markers.

Comment line 30.0 20.0 49.5 130.0 20.0 49.5 260.0 20.0 49.5 20.0 260.0 49.5

The wand

An AMASS calibration wand must be used to perform the calibration. It consists of a tensioned cable on which is mounted six markers spaced more or less evenly. A seventh marker is attached to the frame tensioning the cable. The function of extra marker is to better define the orientation of the wand about the axis of the cable, but it is not used in the final calibration computation and has no effect on the final system accuracy. The tensioned cable scheme generally provides the maximum unobstructed views of the markers, and the tension in the cable is sufficient to make it behave as a rigid body. The resulting calibration can thus be made free of wand distortion and reduced marker obscuring artifacts, providing the highest degree of confidence in the resulting camera calibration parameters.



Even though the cable markers are quite evenly spaced, there are subtle differences in their locations in relation to each other. After an initial crude calibration the relative marker position measurements are refined sufficiently to enable the subtle differences to be measured, and the program attempts to identify the marker distribution as belonging to one of the wands distributed by ADTech. If the wand cannot be identified then it is assumed that there were some deficiencies in the wand motions, camera setup, or data capture, and the calibration is aborted. This technique ensures that a successful calibration guarantees high accuracy.

Marker Alignment



The black dot on each marker must point towards the wand structure.

The marker wand is identified by the distribution of markers on the wand. If the cable becomes twisted, the distribution of markers becomes different than the expected distribution, and the wand is unrecognized by AMASS.

Dots on the wand should all be facing the structure. If the markers are not aligned correctly, release the tension in the wand. With the pin disengaged, twist the wand arm by 360 degrees either left or right to get the correct dot alignment, then engage the pin into the slot.

To properly align the dots, the wand must be twisted in the correct direction. While facing the expanded wand, with the fixed arm to the left and the rotating free arm to the right, note the direction the dots are twisting. If the dots are rotating away (to the front up back down) rotate the free end towards you. If the dots are twisting away (to the back up front down), rotate the free end away from you. Depending on how twisted the wand has become, it may be necessary to rotate the wand multiple times until all dots are facing the wand structure.

Wand size

The size of the volume that you wish to calibrate will determine the wand size that is needed. The largest wand currently available is 1000mm between outside markers, and it typically uses 25.4mm markers. Wands scaled to $\frac{1}{2}$, $\frac{1}{4}$, and $\frac{1}{8}$ th sizes are also available, as are custom sized wands by special order.

Ideally wand markers should span $\frac{1}{3}$ to $\frac{1}{2}$ of the each camera's horizontal field-of-view. If a camera does not record sufficient number of frames of the entire wand then it will typically not calibrate. On the other hand if the wand is very small in a camera's image then that camera's calibration accuracy will be compromised. Think of it in terms of using a 1.0 meter ruler to measure a kilometer vs. a 100 meter ruler to measure the same distance. The larger ruler will provide a more accurate measure because the smaller number of operations will reduce the accumulation of errors that influence the result.

It should be noted that the calibration does not assume any specific lens model for optical distortions and is capable of computing accurate parameters even for lenses that show extreme distortions (of at least up to 10%).

Data capture

Preparation

Wand motions

The motions of the wand are critical in providing a successful (and good) calibration because the wand image data constitute the only input to the calibration. The motions should be smooth and continuous, and at a comfortable speed. It is also important to rotate the wand about the handle as it is moved, but not continuously in one direction. Movement speed will be somewhat dependent on the frame rate of the camera system. For systems with frame rates 100 fps or above, the wand can be moved fairly quickly, but for 30 fps or less the motions should be slower.

In the calibration parameters the user may also specify that every 2nd, 3rd, etc. frame may be processed. This is useful for systems with very high framing rates where using every frame essentially provides redundant data for determining the calibration parameters. Fast, sudden changes in direction should be avoided because that may interfere with the wand tracking, and cause flexion and vibrations of the wand. Before performing a calibration the user should mark out the limits of the desired measurement volume so that wand may be moved to cover all of it.

The goal of the movement is to have the wand in all orientations at every location in the 3D volume. This is obviously not possible, but the user should endeavor to rotate the wand in all 3D orientations throughout the motion capture volume.

Number of frames

The number of frames required for a good calibration will depend on a several factors including the camera system's frame rate and the size of the volume being calibrated. For standard video systems collecting data at rates of 50 - 100 fps, between 1000 and 2000 frames are recommended. This number may need to be increased if the cameras observe many "junk" markers, there are many marker merges in the camera images, or more time is required to move the wand through all parts of the measurement volume. For very high frame rates you may wish to specify the use of every other, or every third frame for the calibration computation. Do not use a calibration performed at one frame rate for tracking data collected at a different frame rate because quite often camera characteristics depend upon the frame rate settings.

Reference markers

From the above discussion it should be obvious that complete coverage of the measurement volume with the wand is essential for reliable measurements. Of special concern should be static reference markers locations, because if they do not fall inside of the calibration volume they will not be reconstructed and cannot be used to define the Global Coordinate System. If they are at floor level, care should be taken to ensure that the wand is moved close to the floor and in the vicinity of every static marker.

See [Reference Marker File](#)

Camera interdependence

This section points out an important aspect of the calibration process. The only input to the calibration is the image marker locations digitized from all cameras. Initially the system has no knowledge of camera locations or wand locations. Each camera's location can be estimated from the image coordinates of the markers it sees and the assumed location of those markers in the 3D space.

However, we need to use all of the other cameras to compute the markers' locations in 3D space, meaning that errors in another camera can affect the calibration of our camera. Through an iterative process and use of data from many frames the computation deletes outliers and converges on a set of calibration parameters (camera positions and orientations, wand locations, camera geometry, etc.) that most accurately fits the input data. The process of excluding bad input data is not clear cut and is influenced by a several factors, but chiefly the number of cameras being calibrated. With more cameras it is easier to spot data that is not correct, and a calibration with just two cameras is typically the least accurate.

The recommendation is that problems with individual cameras that demonstrate large residuals, or low ray or frame counts, should be rectified, and the data capture for the calibration should be repeated. Alternately, the camera can be disconnected from the system, or excluded from the calibration by use of the camera settings in the Calibration section of the Project file.

Cameras seeing other cameras

AMASS allows cameras to see other cameras without masking the 2D imageae.

This situation is, however, not ideal because extra markers will be reconstructed at the camera strobe locations.

AMASS attempts to exclude "strobe markers".

If you track with a minimum of 2 cameras, a lot of these camera strobes appear as multiple short segments.

If you track with a minimum of 3 cameras, all of the short stuff goes away, and we recommend using that when necessary.

ViewLin

A successful calibration's results may be viewed by the ViewLin program. The user is presented with the screen shown below. The main graphics view shows the area of a camera's image that was calibrated, with the grid indicating the lens distortions present for the particular camera.



Items in group box

The first line indicates the name of the file the calibration created. The first item in the second line displays the time of capture of the .cen file used for the calibration, or the time of creation of the data file from which the .cen file was derived by use of the AMASS Convert facility. The other items displayed in the header group box are defined as follows:

Frames used: The ***number of frames actually used / the total number of frames available*** for the calibration (also expressed as a percentage). Frames may be excluded for various reasons including the wand being out of view to most cameras, too many "junk" markers, and inability of the

program to track the marker images in the camera views.

Rays used: This is a similar measure but pertains to the total number of rays (a ray is defined as the line from a marker to the center of a camera's lens) that were measured. Again a ray can be excluded for a number of reasons such as; a marker is partially obscured, it is only observed by one camera, it cannot be identified as being part of the wand, it could not be tracked in a camera's image, it was too close to a "junk" marker to be identified with certainty, or its average residual value was too large.

Point residuals: This is an average of all point residuals (in user distance units) for all wand markers used for the calibration. A ray residual is defined as the shortest distance between a ray and the calculated location of the point belonging to that ray. The location of a point (marker center) is computed by finding the point that has the smallest RMS distance to all rays assigned to that marker. So this value is an indicator of how closely the measured ray directions for all accepted markers intersect, and provides a good indication of the accuracy of the calibration.

Wand residuals: The wand may be considered to be a rigid body having a much greater extent than a single marker. In certain circumstances a calibration could produce very low residuals in every location where individual wand markers are found, yet the connection between distant markers may be incorrectly characterized resulting effectively in a distorted 3D space. The Wand residual is the RMS value of all residuals (in user distance units) obtained when all markers on the wand (which is assumed to be rigid) are simultaneously fitted to their rays in each frame. Its value will always be larger than the Point residual but it should not be much greater. A much larger wand residual would indicate that the distance measurements between markers further apart than the length of the wand would probably have significant errors.

Wand length SD: This value presents the standard deviation of the distance between the two outermost markers on the wand over all frames in which the markers could be identified. It is perhaps a better indicator of the "linearity" of the calibrated volume. Again, the length units are user units.

It should be noted that the residual values are a function of the accuracy of the calibration and also the noise inherent in the camera image data. The noise comes from the discrete nature of camera sensors as well as the processing electronics and software. Using cameras with low resolution sensors, especially when imaging small markers, can be a principal source of noise and the major component of the presented residual values.

Here it should also be pointed out that a calibration with a smaller number of frames may have smaller residual values than one with a large number of frames, yet the calibration will usually be less accurate. The reason is that it is easier to fit the many camera system parameters to a smaller number of input data point because the calibration is minimizing the overall residuals of all rays. It biases the resulting calibration to produce very low residuals at the participating points, at the expense of locations that are elsewhere in the volume where the residuals may be quite high.

Camera list box

This area provides some results for individual cameras, as well as letting the user select the camera whose linearity grid is to be displayed in the main graphics area.

Rays: This is the number of rays from a camera that were used for the system calibration. Low values in relation to other cameras indicate that a camera did not successfully image many wand markers and that its calibration might be suspect.

Residual: is the average residual (in user units) for all rays that could be used from the camera. Again, it should not be much larger than for the other cameras, and it is somewhat dependent upon the resolution of the camera's sensor.

% Image: The percentage of the camera's image area that was successfully calibrated. Any marker image falling outside the shown grid area will be ignored during tracking when this calibration is used.

Reference points

This list box provides a report on the reference markers that were used to define the Global Coordinate System in which all marker trajectories will be expressed.

Point #: The reference marker number as defined by the position of its coordinates in the supplied .ref file.

Point residual: The RMS residual of all rays (one from each camera) that could make a measurement on that marker. Units are user units.

Body residual: The RMS residual of all rays to all reference points when the reference points are treated as a rigid body whose dimensions are specified in the .ref file. It is an indication of the accuracy of the marker coordinates provided in the .ref file. If the coordinates of one or more points have large errors then they may not be identified, and cannot be used in computing the transformation to the Global Coordinate System. A minimum of three non-symmetric points must be correctly identified for the transformation to be calculated. Large residuals here do not affect the calibration or tracking accuracy, but could cause the resulting marker coordinates to be rotated or translated from the user's desired coordinate system. One useful technique is to track the static markers using the calibration and use the data to calculate their separations so that their positions in the desired Global Coordinate System can be specified more correctly.

If reference marker coordinates are provided but less than three reference points are found, the calibration program provides the option of saving the calibration. In that case the resulting GCS is defined as follows; camera 1 is at the origin, the line from camera 1 to camera 2 is the x-axis, and camera 3 lies in the positive y of the x-y plane. If there are only two cameras then orientation of the x-y plane is arbitrary.

Button Controls

Close: Closes the ViewLin program.

Next: Closes the ViewLin program and starts calibration of the next .cen file if multiple .cen files were selected for calibration in the AMASS shell, or display the next calibration file if multiple .cal files were selected in the AMASS shell.

Exit: Terminates the ViewLin program.

Apply: Writes the name of the currently displayed calibration file to the current project file's tracking parameters, to be used for the next tracking computation.

Delete: Deletes the currently displayed calibration file.

From:

<https://has-motion.com/wiki/> - Wiki Documentation Site

Permanent link:

https://has-motion.com/wiki/doku.php?id=other:amass:documentation:system_calibration&rev=1721157628

Last update: **2024/07/16 19:20**

