Camera Calibration

Readings in Vision, Graphics and Interactive Systems

Christian Lindequist Larsen

October 16, 2009

A Picture Says More Than a Thousand Words

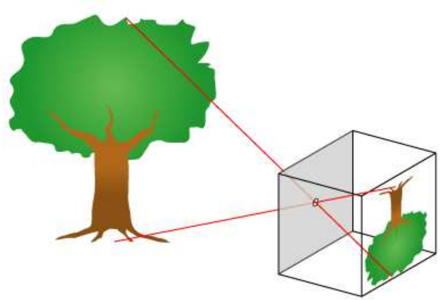

What is Camera Calibration?

- Camera calibration "is the process of finding the true parameters of the camera that produced a given photograph or video," *Wikipedia*

What is Camera Calibration?

- Camera calibration "is the process of finding the true parameters of the camera that produced a given photograph or video," *Wikipedia*
- Two sets of parameters are determined:
  - *Intrinsic* parameters
  - *Extrinsic* parameters

Camera Calibration Parameters

- **Intrinsic parameters**
  - how does the camera depict the scene?

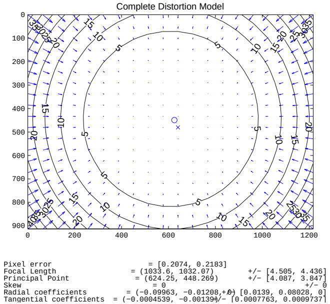

Camera Calibration Parameters

- **Intrinsic parameters**
  - how does the camera depict the scene?
- **Extrinsic parameters**
  - where is the camera located?

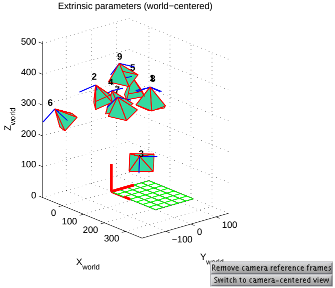

Motivation and Applications

- Calibration is necessary whenever we want to measure anything accurately from images.
- Accurate measurements are vital for many computer vision systems.

Motivation and Applications

- Calibration is necessary whenever we want to measure anything accurately from images.
- Accurate measurements are vital for many computer vision systems.
- Applications include:
  - Augmented reality
  - Robots picking up objects
  - Aerial photography

Augmented Reality

ARToolKit

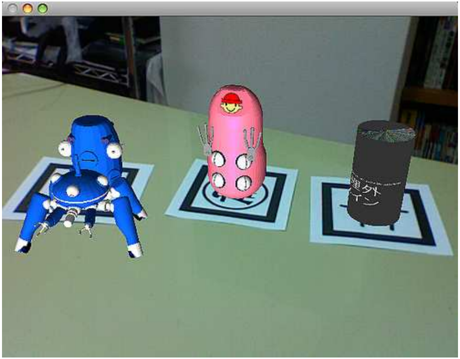

Augmented Reality

Interactive Assembly Guide using Augmented Reality

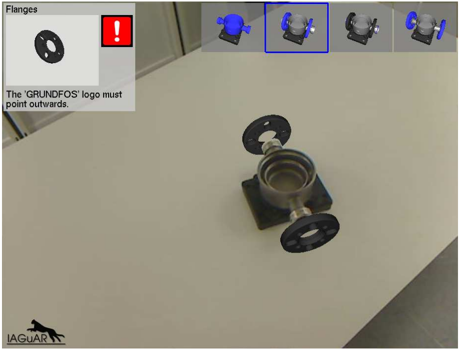

Robots Picking Up Objects

Bin Picking

Camera Models

Basic Pinhole Model

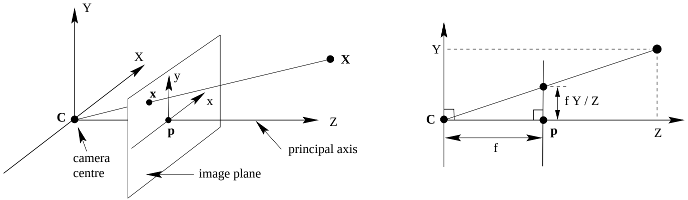

- Central projection of points in space onto the image plane.
- Centre of projection is $\mathbf{C} = \mathbf{O}$, and image plane is $Z = f$.

Basic Pinhole Model

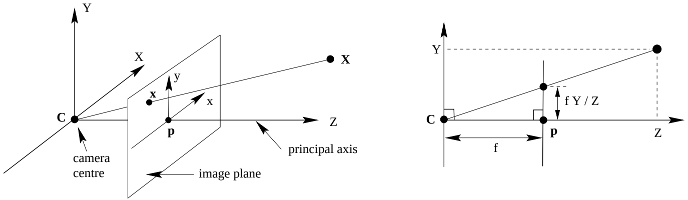

- Central projection of points in space onto the image plane.
- Centre of projection is $\mathbf{C} = \mathbf{O}$, and image plane is $Z = f$.
- A point in space $\mathbf{X} = [X \ \ Y \ \ Z]^\mathsf{T}$ is mapped to the image plane at the intersection with the line from $\mathbf{C}$ to $\mathbf{X}$. By similar triangles

$$\mathbf{X} = [X \ \ Y \ \ Z]^\mathsf{T} \mapsto [fX/Z \ \ \ fY/Z]^\mathsf{T} = \mathbf{x}.$$

Central Projection using Homogeneous Coordinates

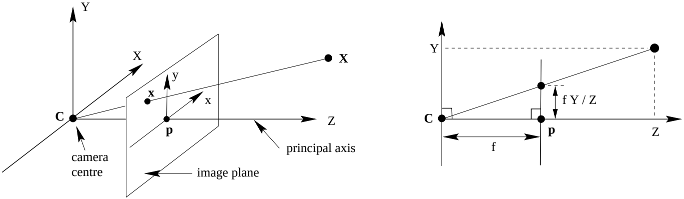

- Representing world and image coordinates by homogeneous vectors, this becomes a linear mapping:

$$\begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} \mapsto \begin{bmatrix} fX \\ fY \\ Z \end{bmatrix} = \begin{bmatrix} f & & & 0 \\ & f & & 0 \\ & & 1 & 0 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}$$

Central Projection using Homogeneous Coordinates

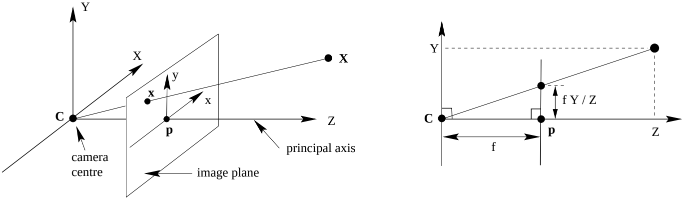

- Representing world and image coordinates by homogeneous vectors, this becomes a linear mapping:

$$\begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} \mapsto \begin{bmatrix} fX \\ fY \\ Z \end{bmatrix} = \begin{bmatrix} f & & & 0 \\ & f & & 0 \\ & & 1 & 0 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}$$

- In homogeneous coordinates $[fX \; fY \; Z]^\mathsf{T} \sim [fX/Z \;\; fY/Z \;\; 1]^\mathsf{T}$.

Central Projection using Homogeneous Coordinates

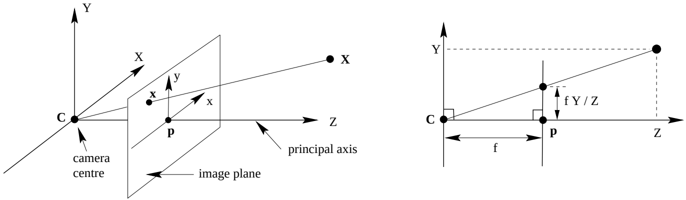

- With $\mathbf{X} = [X \ \ Y \ \ Z \ \ 1]^\mathsf{T}$, and $\mathbf{x} = [fX \ \ fY \ \ Z]^\mathsf{T}$, we have the *camera projection matrix*

$$\mathsf{P} = \begin{bmatrix} f & & & 0 \\ & f & & 0 \\ & & 1 & 0 \end{bmatrix} = \mathrm{diag}(f, f, 1)[\,\mathsf{I} \,|\, \mathbf{0}\,].$$

- The mapping from world to image coordinates is now simply

$$\mathbf{x} = \mathsf{P}\mathbf{X}.$$

Principal Point Offset

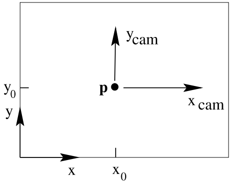

- In general the origin of coordinates in the image plane is not the principal point, $\mathbf{p} = [p_x \;\; p_y]^\mathsf{T}$. Thus there is a mapping

$$[X \;\; Y \;\; Z]^\mathsf{T} \mapsto [fX/Z + p_x \quad fY/Z + p_y]^\mathsf{T}.$$

- Expressed in homogeneous coordinates, this becomes:

$$\begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} \mapsto \begin{bmatrix} fX + Zp_x \\ fY + Zp_y \\ Z \end{bmatrix} = \begin{bmatrix} f & & p_x & 0 \\ & f & p_y & 0 \\ & & 1 & 0 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}$$

Principal Point Offset

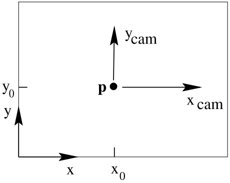

- Letting $P = K[\,\mathbf{I}\,|\,\mathbf{0}\,]$, where K is the *camera calibration matrix*

$$\mathrm{K} = \begin{bmatrix} f & & p_x \\ & f & p_y \\ & & 1 \end{bmatrix},$$

- the mapping from camera space $\mathbf{X}_{\mathrm{cam}}$ to image coordinates $\mathbf{x}$ becomes:

$$\mathbf{x} = \mathrm{P}\mathbf{X}_{\mathrm{cam}} = \mathrm{K}[\,\mathbf{I}\,|\,\mathbf{0}\,]\mathbf{X}_{\mathrm{cam}}$$

Camera Rotation and Translation

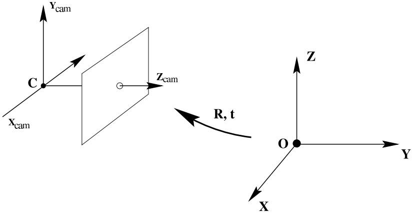

- Points are expressed in the *world coordinate frame*, which is related to the *camera coordinate frame* via a rotation R (3×3 matrix) and a translation $\mathbf{t}$.
- In homogeneous coordinates, this becomes:

$$\mathbf{X}_{\text{cam}} = \begin{bmatrix} \mathrm{R} & \mathbf{t} \\ \mathbf{0} & 1 \end{bmatrix} \mathbf{X}$$

Camera Rotation and Translation

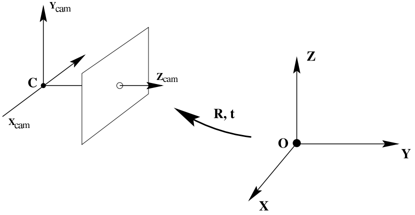

- This leads to a camera projection matrix of the form

$$P = K[R \,|\, \mathbf{t}].$$

- A direct mapping from world to image coordinates is now

$$\mathbf{x} = P\mathbf{X} = K[R \,|\, \mathbf{t}]\mathbf{X}.$$

CCD Cameras

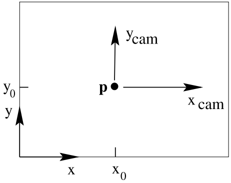

- In the previous slides, image coordinates were assumed to have equal scale on both axes.
- CCD cameras may have non-square pixels, and if image coordinates are measured in pixels, this changes the calibration matrix.

CCD Cameras

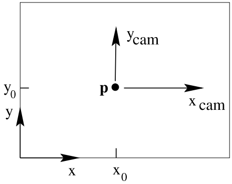

- Let m_x and m_y be the number of pixels per unit distance in the x and y directions, then

$$K = \begin{bmatrix} \alpha_x & & x_0 \\ & \alpha_y & y_0 \\ & & 1 \end{bmatrix},$$

where $\alpha_x = f m_x$, and $\alpha_y = f m_y$ represent the focal length in pixel dimensions. Similarly $x_0 = m_x p_x$, and $y_0 = m_y p_y$ represent the principal point in pixel dimensions.

Finite Projective Camera

- For added generality, a skew parameter s is included, and finally

$$K = \begin{bmatrix} \alpha_x & s & x_0 \\ & \alpha_y & y_0 \\ & & 1 \end{bmatrix}.$$

- If pixels in the CCD array are skewed, $s \neq 0$ meaning the x- and y-axes are not perpendicular.
- In general $s = 0$ for CCD cameras.

Lens Distortion

- A real camera may also have lens distortion, which can not be described by the camera projection matrix.
- This distortion is described by 5 distortion coefficients represented as a column vector $\mathbf{d}$.

Lens Distortion

- A real camera may also have lens distortion, which can not be described by the camera projection matrix.
- This distortion is described by 5 distortion coefficients represented as a column vector **d**.
- Letting $\mathbf{X}_{\text{cam}} = [X_c \ Y_c \ Z_c]^{\mathsf{T}}$, we have the normalized (pinhole) image projection:

$$\mathbf{x}_n = \begin{bmatrix} X_c/Z_c \\ Y_c/Z_c \end{bmatrix} = \begin{bmatrix} x_n \\ y_n \end{bmatrix}$$

Lens Distortion

- With $r^2 = {x_n}^2 + {y_n}^2$, the distorted normalized image coordinate becomes:

$$\mathbf{x}_d = \begin{bmatrix} x_d \\ y_d \end{bmatrix} = (1 + d_1 r^2 + d_2 r^4 + d_5 r^6)\mathbf{x}_n + \mathbf{dx},$$

- where **dx** is the tangential distortion vector:

$$\mathbf{dx} = \begin{bmatrix} 2d_3 x_n y_n + d_4(r^2 + 2{x_n}^2) \\ d_3(r^2 + 2{y_n}^2) + 2d_4 x_n y_n \end{bmatrix}$$

Lens Distortion

- Finally, the pixel coordinate $\mathbf{x}_p = [x_p \ y_p]^\mathsf{T}$ of the projection of $\mathbf{X}_{\text{cam}}$ when distortion is applied is:

$$\begin{bmatrix} x_p \\ y_p \\ 1 \end{bmatrix} = \mathsf{K} \begin{bmatrix} x_d \\ y_d \\ 1 \end{bmatrix}$$

- In other words,

$$\mathbf{x}_p = \begin{bmatrix} \alpha_x x_d + s\, y_d + x_0 \\ \alpha_y y_d + y_0 \end{bmatrix}.$$

Summary of Camera Models

- For *finite projective cameras* points in space $\mathbf{X}$ are mapped onto image coordinates $\mathbf{x}$ through the linear transformation

$$\mathbf{x} = \mathrm{P}\mathbf{X},$$

where P is the 3×4 *camera projection matrix*.

- Defining the *camera calibration matrix*

$$\mathrm{K} = \begin{bmatrix} \alpha_x & s & x_0 \\ & \alpha_y & y_0 \\ & & 1 \end{bmatrix},$$

we can express P as:

$$\mathrm{P} = \mathrm{K}[\,\mathrm{R} \,|\, \mathbf{t}\,],$$

where R and $\mathbf{t}$ are the rotation and translation transforming points from world to camera space.

Summary of Camera Models

- The *intrinsic* parameters of the camera are:
  - the calibration matrix K, and
  - the lens distortion parameters $d_i, i = 1, 2, \ldots, 5$.
- The *extrinsic* parameters are:
  - the rotation matrix R, and
  - the translation vector $\mathbf{t}$.

Camera Calibration

Camera Calibration Toolbox for Matlab

- Recall that we want to find the true parameters of the camera, i.e. the intrinsic and extrinsic parameters.
- This process is greatly simplified using the Camera Calibration Toolbox for Matlab.
- The following is a minimal introduction on how to use the toolbox. Really good tutorials are available on its homepage.

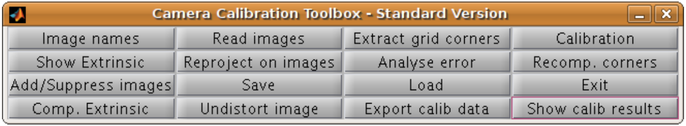

Intrinsic Calibration

Steps involved:

1. Load input images with calibration pattern.
2. Extract grid corners on all images.
3. Compute calibration parameters.
4. Inspect the result.

Loading Calibration Images

- Best results are achieved with more images, e.g. 20, of the calibration pattern from different angles.
- Images are loaded using the *Image names* button in the GUI of the calibration toolbox.
- In this example 9 images are loaded.

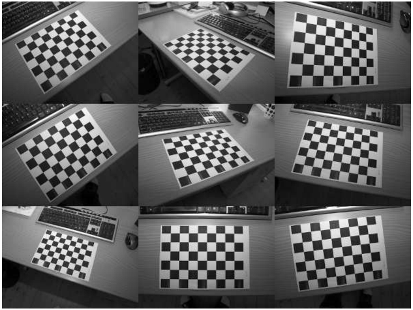

Extracting Grid Corners

- Press *Extract grid corners* in the GUI.
- Click the four extreme grid corners (first click defines the origin).

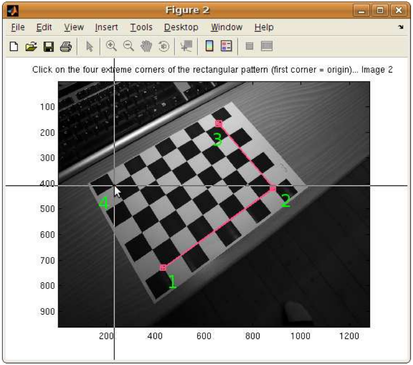

Extracting Grid Corners

- Press *Extract grid corners* in the GUI.
- Click the four extreme grid corners (first click defines the origin).
- Guessed corner locations are marked.

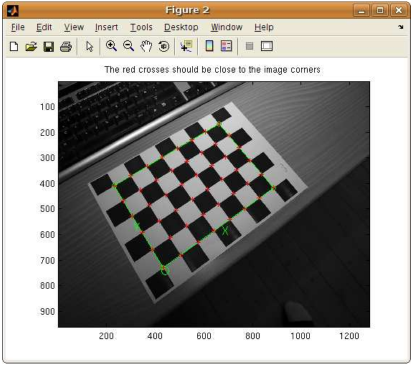

Extracting Grid Corners

- Press *Extract grid corners* in the GUI.
- Click the four extreme grid corners (first click defines the origin).
- Guessed corner locations are marked.
- Optimized corner locations are found.

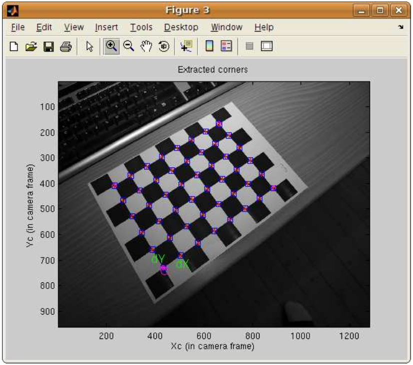

Extracting Grid Corners

- Press *Extract grid corners* in the GUI.
- Click the four extreme grid corners (first click defines the origin).
- Guessed corner locations are marked.
- Optimized corner locations are found.
- Repeat for all images.

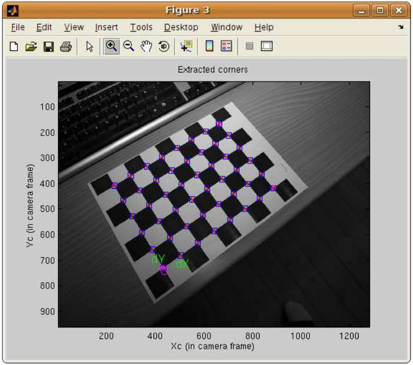

Computing Calibration Parameters

- Press the *Calibration* button in the GUI to start calibration.

```
Aspect ratio optimized (est_aspect_ratio = 1) -> both components of fc are estimated (
Principal point optimized (center_optim=1) - (DEFAULT). To reject principal point, set
Skew not optimized (est_alpha=0) - (DEFAULT)
Distortion not fully estimated (defined by the variable est_dist):
     Sixth order distortion not estimated (est_dist(5)=0) - (DEFAULT) .
Initialization of the principal point at the center of the image.
Initialization of the intrinsic parameters using the vanishing points of planar patter

Initialization of the intrinsic parameters - Number of images: 9


Calibration parameters after initialization:

Focal Length:          fc = [ 1007.79386   1007.79386 ]
Principal point:       cc = [ 639.50000   479.50000 ]
Skew:             alpha_c = [ 0.00000 ]   => angle of pixel = 90.00000 degrees
Distortion:            kc = [ 0.00000   0.00000   0.00000   0.00000   0.00000 ]

Main calibration optimization procedure - Number of images: 9
Gradient descent iterations: 1...2...3...4...5...6...7...8...9...10...11...12...13...1
Estimation of uncertainties...done


Calibration results after optimization (with uncertainties):

Focal Length:          fc = [ 1033.59588   1032.06788 ] □ [ 4.50528   4.43636 ]
Principal point:       cc = [ 624.24952   448.26889 ] □ [ 4.08749   3.84683 ]
Skew:             alpha_c = [ 0.00000 ] □ [ 0.00000  ]   => angle of pixel axes = 90.0
Distortion:            kc = [ -0.09963   -0.01208   -0.00045   -0.00139   0.00000 ] □ [
Pixel error:          err = [ 0.20743   0.21826 ]

Note: The numerical errors are approximately three times the standard deviations (for
```

Calibration Results

- After calibration, the intrinsic parameters of the camera are available.
- The variable KK contains the camera calibration matrix:

$$\text{KK} = \begin{bmatrix} 1033.6 & 0 & 624.25 \\ 0 & 1032.1 & 448.27 \\ 0 & 0 & 1 \end{bmatrix}$$

- The distortion coefficients are available in **kc**:

$$\mathbf{kc} = [-0.0996\ -0.0121\ -0.0005\ -0.0014\ 0]^T$$

Calibration Results

- Visualizing the effect of distortions on the pixel image (run the script *visualize_distortions* from the Matlab prompt):

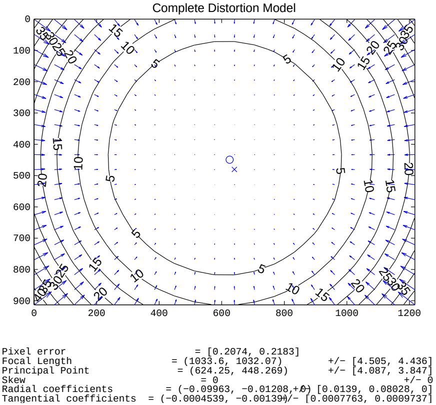

Calibration Results

- Visualization of extrinsic parameters for calibration images (press *Show extrinsic* in the GUI):

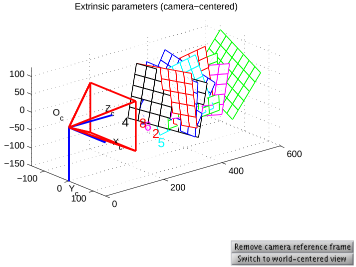

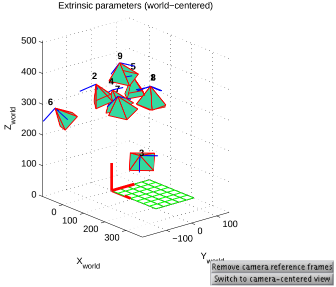

Extrinsic Calibration

- Once the intrinsic parameters of the camera have been found, extrinsic parameters can be computed for other images from the same camera.
- Press the *Comp. extrinsic* button in the GUI.

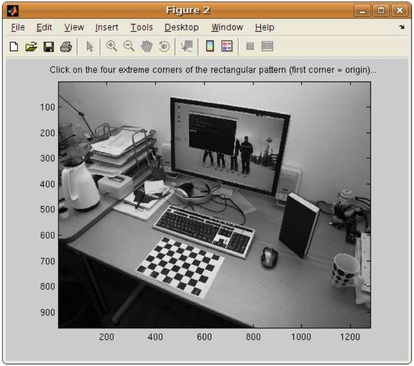

Extrinsic Calibration

- Once the intrinsic parameters of the camera have been found, extrinsic parameters can be computed for other images from the same camera.
- Press the *Comp. extrinsic* button in the GUI.
- As before, click the four extreme corners of the grid.

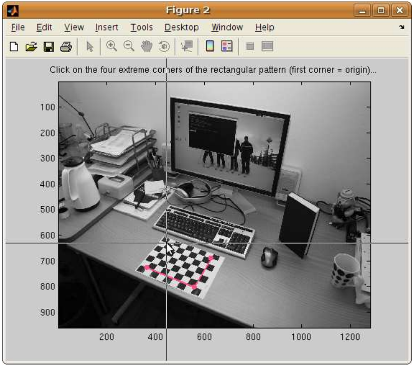

Extrinsic Calibration

- Once the intrinsic parameters of the camera have been found, extrinsic parameters can be computed for other images from the same camera.
- Press the *Comp. extrinsic* button in the GUI.
- As before, click the four extreme corners of the grid.
- The result is shown as an overlay on the image.

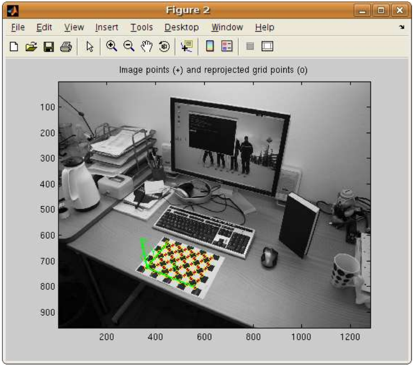

Extrinsic Calibration

- Recall that the extrinsic parameters are a rotation matrix R, and a translation vector **t**.
- The variable Rc_ext now contains the rotation matrix:

$$\text{Rc\_ext} = \begin{bmatrix} 0.9145 & 0.4046 & 0.0093 \\ 0.2263 & -0.4921 & -0.8406 \\ -0.3355 & 0.7708 & -0.5415 \end{bmatrix}$$

- The translation vector is available in **Tc_ext**:

$$\mathbf{Tc\_ext} = [-250.2 \; 267.5 \; 980.4]^T$$

Complete Calibration of a Single Camera

- Having computed both the intrinsic and extrinsic parameters of the camera, the mapping between world and image coordinates is known.
- Using the finite projective camera model, the camera projection matrix of the calibrated camera becomes:

$$\mathrm{P} = \mathrm{K}[\,\mathrm{R}\,|\,\mathbf{t}\,]$$

- Disregarding lens distortion, the world coordinate $\mathbf{X}$ is mapped to the image coordinate $\mathbf{x}$ by the linear transformation

$$\mathbf{x} = \mathrm{P}\mathbf{X}.$$

Image Rectification

- Using the toolbox, images can be rectified (undistorted). Here is an example:

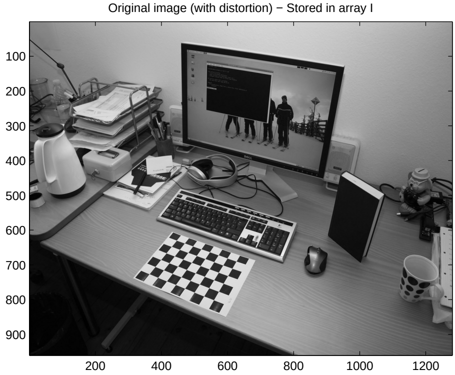

Image Rectification

- Using the toolbox, images can be rectified (undistorted). Here is an example:

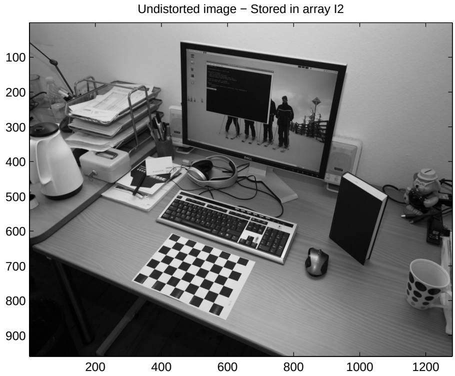

Summary of Camera Calibration

- Estimation of the true parameters of a camera.
- Intrinsic parameters are computed based on a set of calibration images.
  - Camera calibration matrix is stored in KK.
  - Distortion coefficients are stored in **kc**.
- Afterwards extrinsic parameters can be calculated for other images from the same camera.
  - Rotation matrix is stored in Rc_ext.
  - Translation is stored in **Tc_ext**.

Any questions?

References

- Multiple View Geometry in Computer Vision (2nd Edition), *Richard Hartley and Andrew Zisserman*, chapters 6 and 9.
- Camera Calibration Toolbox for Matlab, `http://www.vision.caltech.edu/bouguetj/calib_doc/`
- Camera Resectioning (Calibration), *Wikipedia*, `http://en.wikipedia.org/wiki/Camera_calibration`