Computer vision: models, learning and inference

Chapter 14

The pinhole camera

Structure

- Pinhole camera model
- Three geometric problems
- Homogeneous coordinates
- Solving the problems
  - Exterior orientation problem
  - Camera calibration
  - 3D reconstruction
- Applications

Motivation

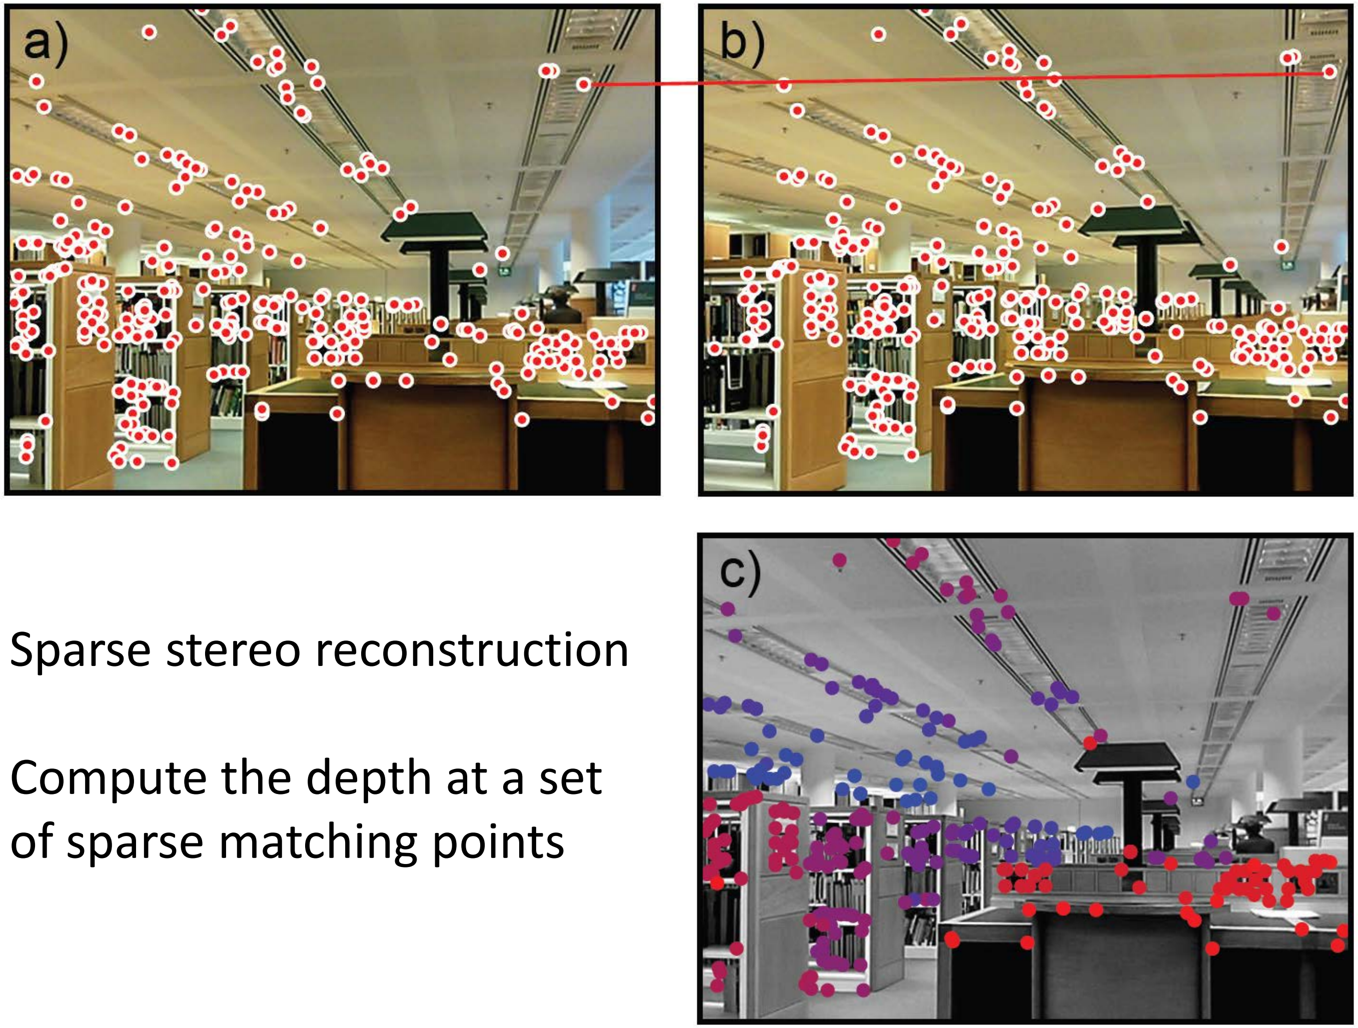

Sparse stereo reconstruction

Compute the depth at a set of sparse matching points

Pinhole camera

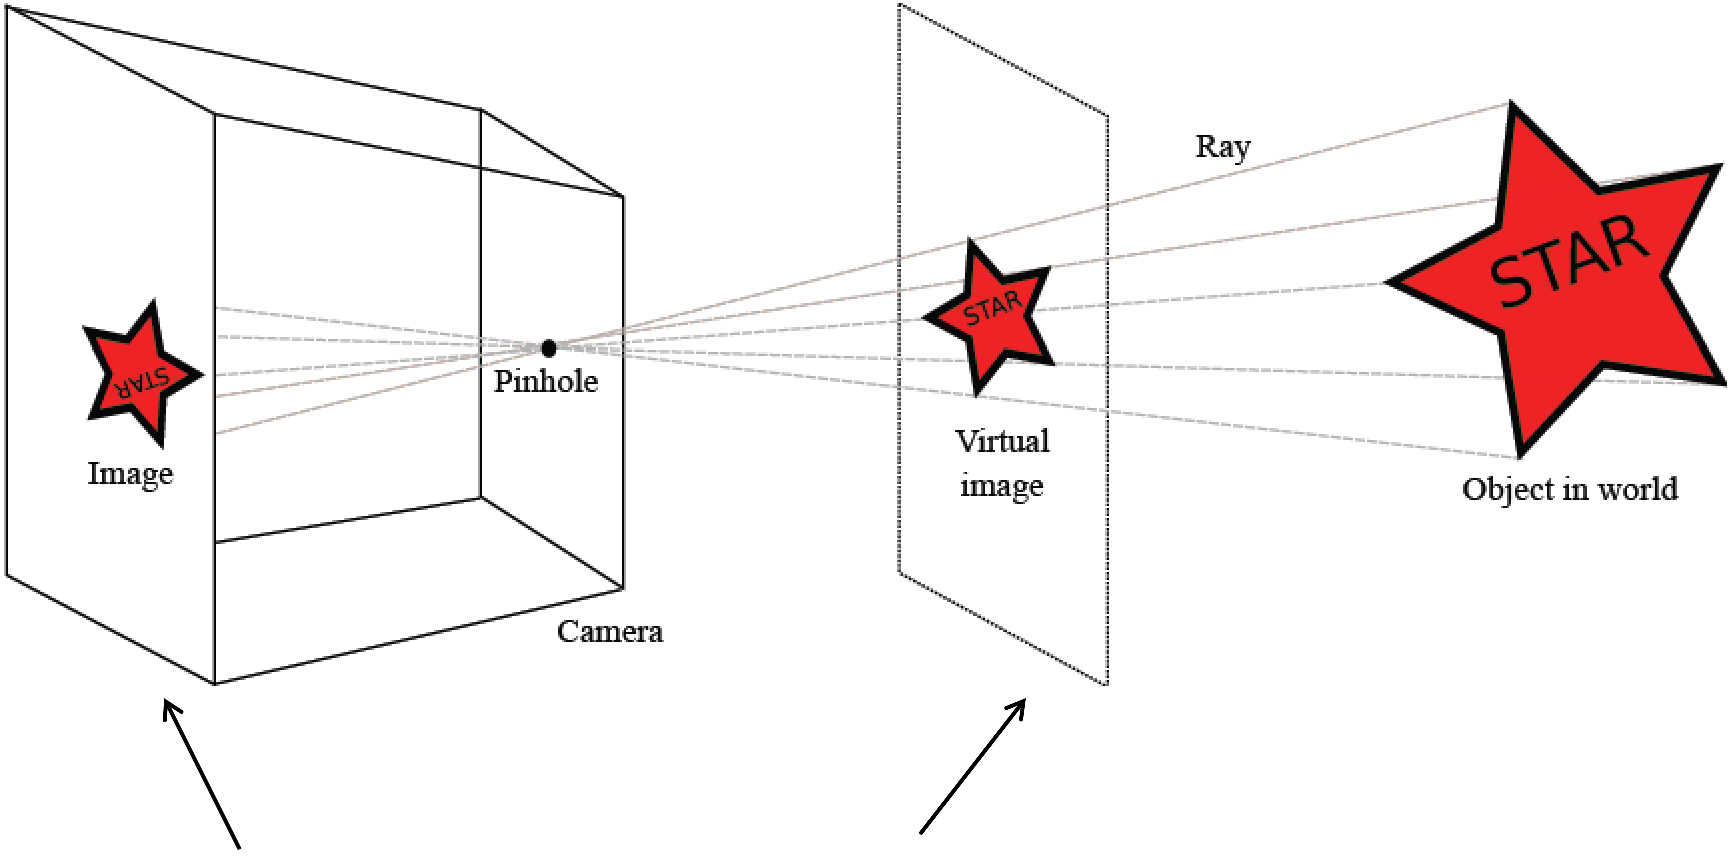

Real camera image is inverted

Instead model impossible but more convenient virtual image

Making Your Own Room With a View | National Geographic
https://www.youtube.com/watch?v=gvzpu0Q9RTU

Pinhole camera terminology

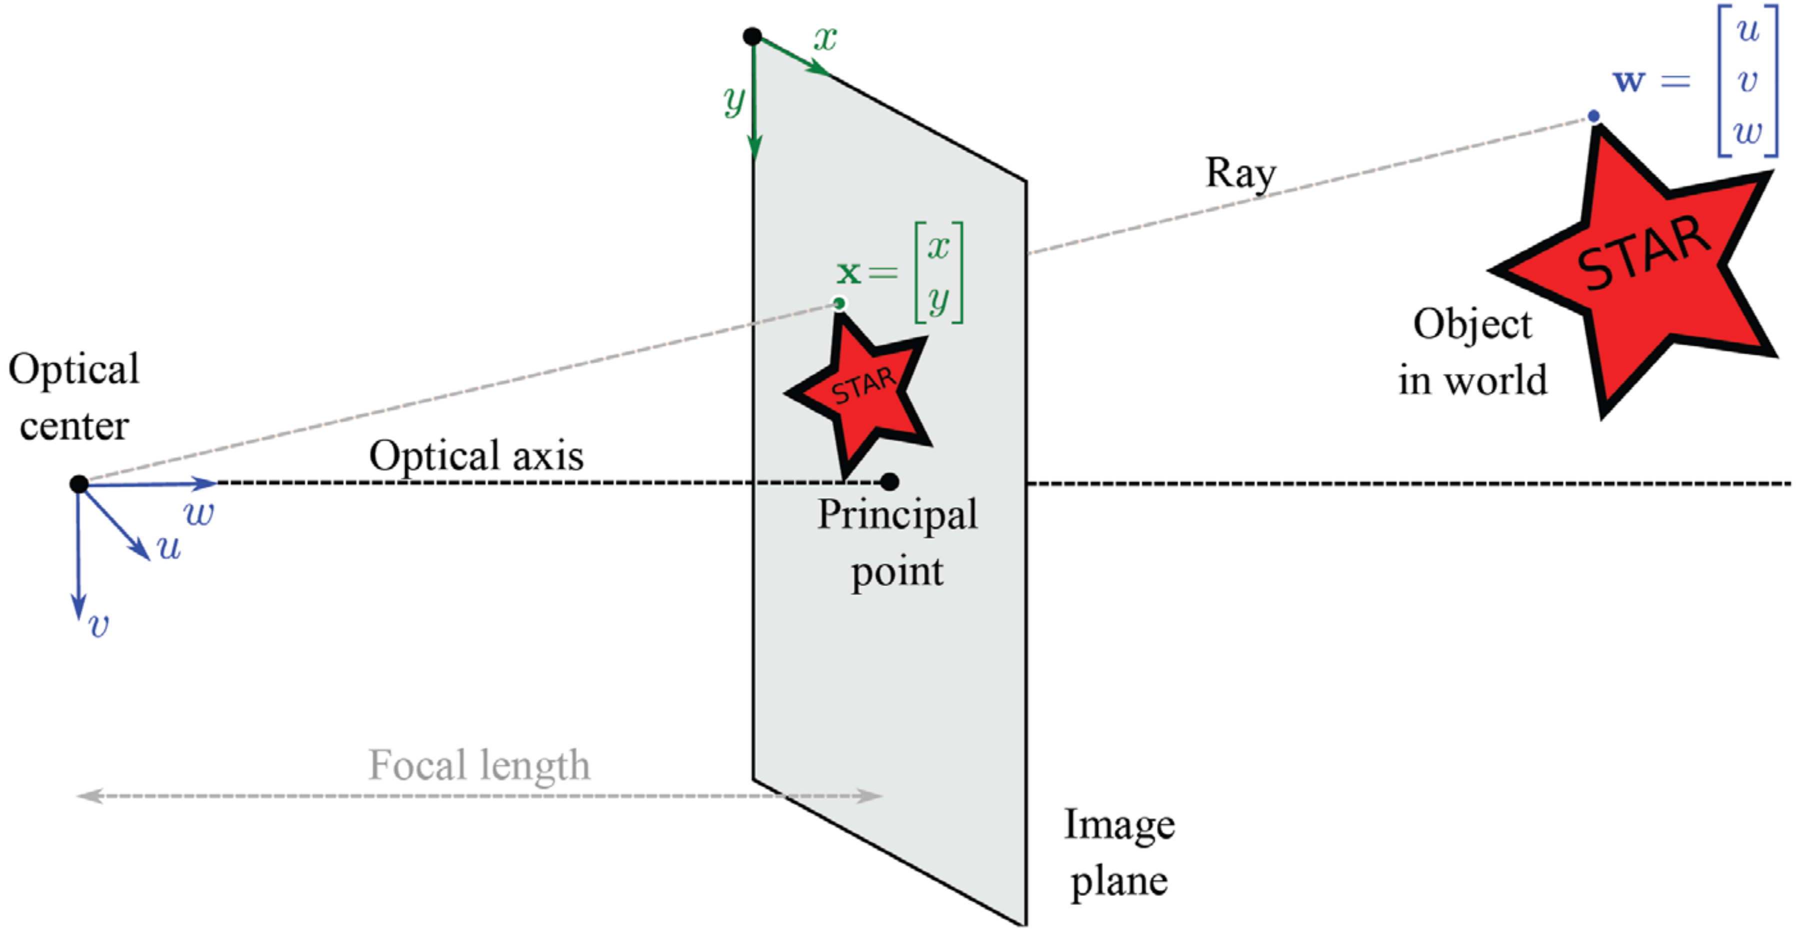
x
y
$\mathbf{w} = \begin{bmatrix} u \\ v \\ w \end{bmatrix}$
Ray
$\mathbf{x} = \begin{bmatrix} x \\ y \end{bmatrix}$
STAR
STAR
Object in world
Optical center
Optical axis
w
u
v
Principal point
Focal length
Image plane

Normalized Camera

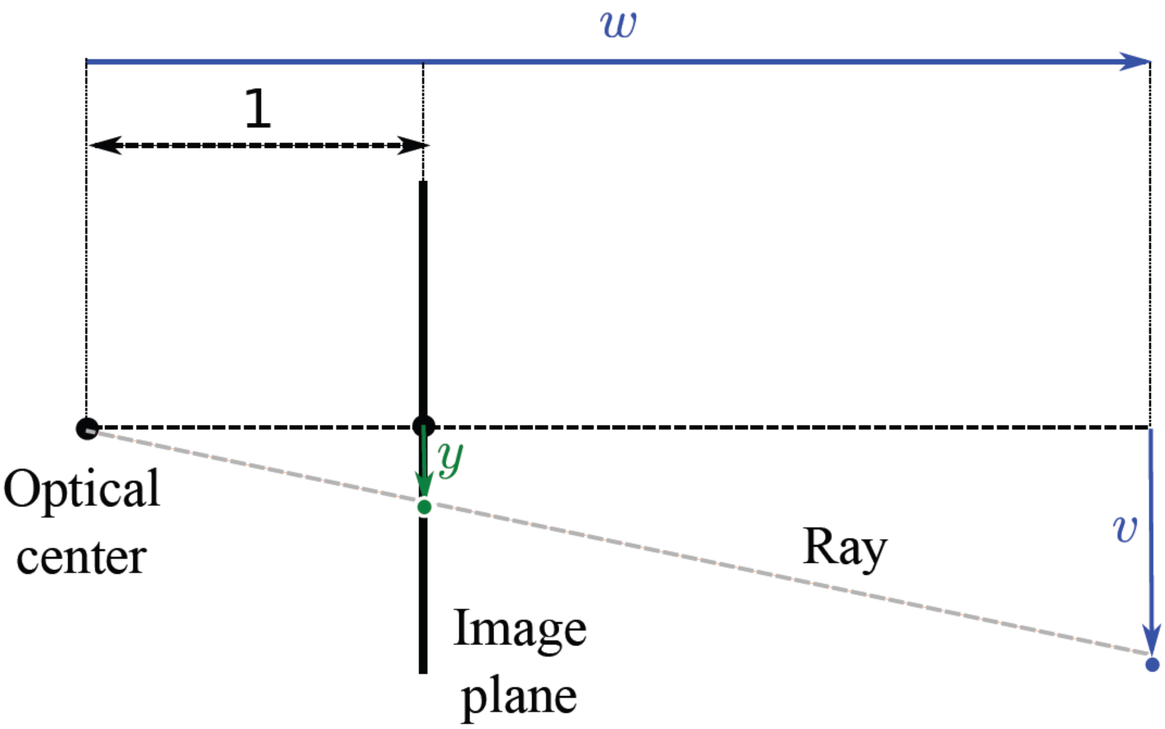

By similar triangles:

$$x = \frac{u}{w} \qquad y = \frac{v}{w}$$

Focal length parameters

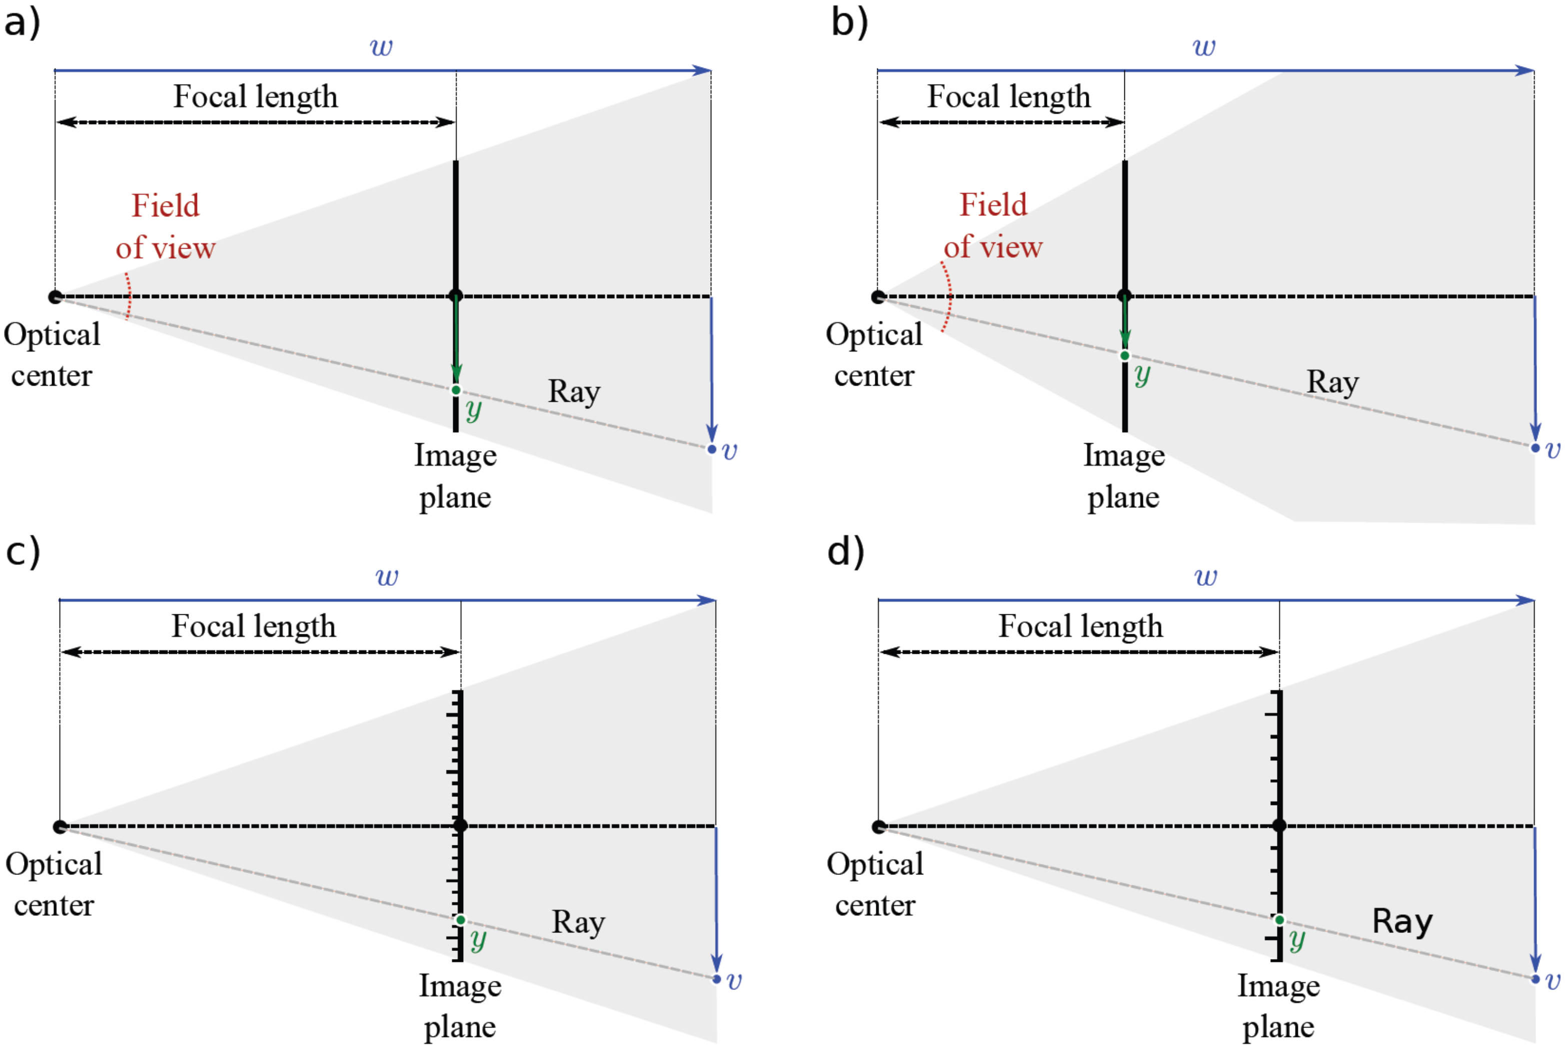

Focal length parameters

Can model both

- the effect of the distance to the focal plane
- the density of the receptors

with a single focal length parameter ϕ

$$x = \frac{\phi u}{w} \qquad y = \frac{\phi v}{w}$$

In practice, the receptors may not be square:

$$x = \frac{\phi_x u}{w} \qquad y = \frac{\phi_y v}{w}$$

So use different focal length parameter for x and y dims

Offset parameters

- Current model assumes that pixel (0,0) is where the principal ray strikes the image plane (i.e. the center)
- Model offset to center

$$
\begin{aligned}
x &= \frac{\phi_x u}{w} + \delta_x \\
y &= \frac{\phi_y v}{w} + \delta_y
\end{aligned}
$$

Skew parameter

- Finally, add skew parameter
- Accounts for image plane being not exactly perpendicular to the principal ray

$$
\begin{aligned}
x &= \frac{\phi_x u + \gamma v}{w} + \delta_x \\
y &= \frac{\phi_y v}{w} + \delta_y
\end{aligned}
$$

Position and orientation of camera

- Position $\mathbf{w}=(u,v,w)^T$ of point in the world is generally not expressed in the frame of reference of the camera.
- Transform using 3D transformation

$$\begin{bmatrix} u' \\ v' \\ w' \end{bmatrix} = \begin{bmatrix} \omega_{11} & \omega_{12} & \omega_{13} \\ \omega_{21} & \omega_{22} & \omega_{23} \\ \omega_{31} & \omega_{32} & \omega_{33} \end{bmatrix} \begin{bmatrix} u \\ v \\ w \end{bmatrix} + \begin{bmatrix} \tau_x \\ \tau_y \\ \tau_z \end{bmatrix}$$

or

$$\mathbf{w}' = \boldsymbol{\Omega}\mathbf{w} + \boldsymbol{\tau}$$

Point in frame of reference of camera

Point in frame of reference of world

Complete pinhole camera model

$$
\begin{aligned}
x &= \frac{\phi_x(\omega_{11}u+\omega_{12}v+\omega_{13}w+\tau_x)+\gamma(\omega_{21}u+\omega_{22}v+\omega_{23}w+\tau_y)}{\omega_{31}u+\omega_{32}v+\omega_{33}w+\tau_z}+\delta_x\\
y &= \frac{\phi_y(\omega_{21}u+\omega_{22}v+\omega_{23}w+\tau_y)}{\omega_{31}u+\omega_{32}v+\omega_{33}w+\tau_z}+\delta_y.
\end{aligned}
$$

- Intrinsic parameters

$$\{\phi_x, \phi_y, \gamma, \delta_x, \delta_y\}$$

(stored as intrinsic matrix)

$$\boldsymbol{\Lambda} = \begin{bmatrix}\phi_x & \gamma & \delta_x\\ 0 & \phi_y & \delta_y\\ 0 & 0 & 1\end{bmatrix}$$

- Extrinsic parameters

$$\{\boldsymbol{\Omega}, \boldsymbol{\tau}\}$$

Complete pinhole camera model

$$\begin{aligned} x &= \frac{\phi_x(\omega_{11}u+\omega_{12}v+\omega_{13}w+\tau_x)+\gamma(\omega_{21}u+\omega_{22}v+\omega_{23}w+\tau_y)}{\omega_{31}u+\omega_{32}v+\omega_{33}w+\tau_z}+\delta_x \\ y &= \frac{\phi_y(\omega_{21}u+\omega_{22}v+\omega_{23}w+\tau_y)}{\omega_{31}u+\omega_{32}v+\omega_{33}w+\tau_z}+\delta_y. \end{aligned}$$

For short:

$$\mathbf{x} = \mathbf{pinhole}[\mathbf{w}, \boldsymbol{\Lambda}, \boldsymbol{\Omega}, \boldsymbol{\tau}]$$

Add noise – uncertainty in localizing feature in image

$$Pr(\mathbf{x}|\mathbf{w}, \boldsymbol{\Lambda}, \boldsymbol{\Omega}, \boldsymbol{\tau}) = \text{Norm}_{\mathbf{x}}\left[\mathbf{pinhole}[\mathbf{w}, \boldsymbol{\Lambda}, \boldsymbol{\Omega}, \boldsymbol{\tau}], \sigma^2\mathbf{I}\right]$$

Radial distortion

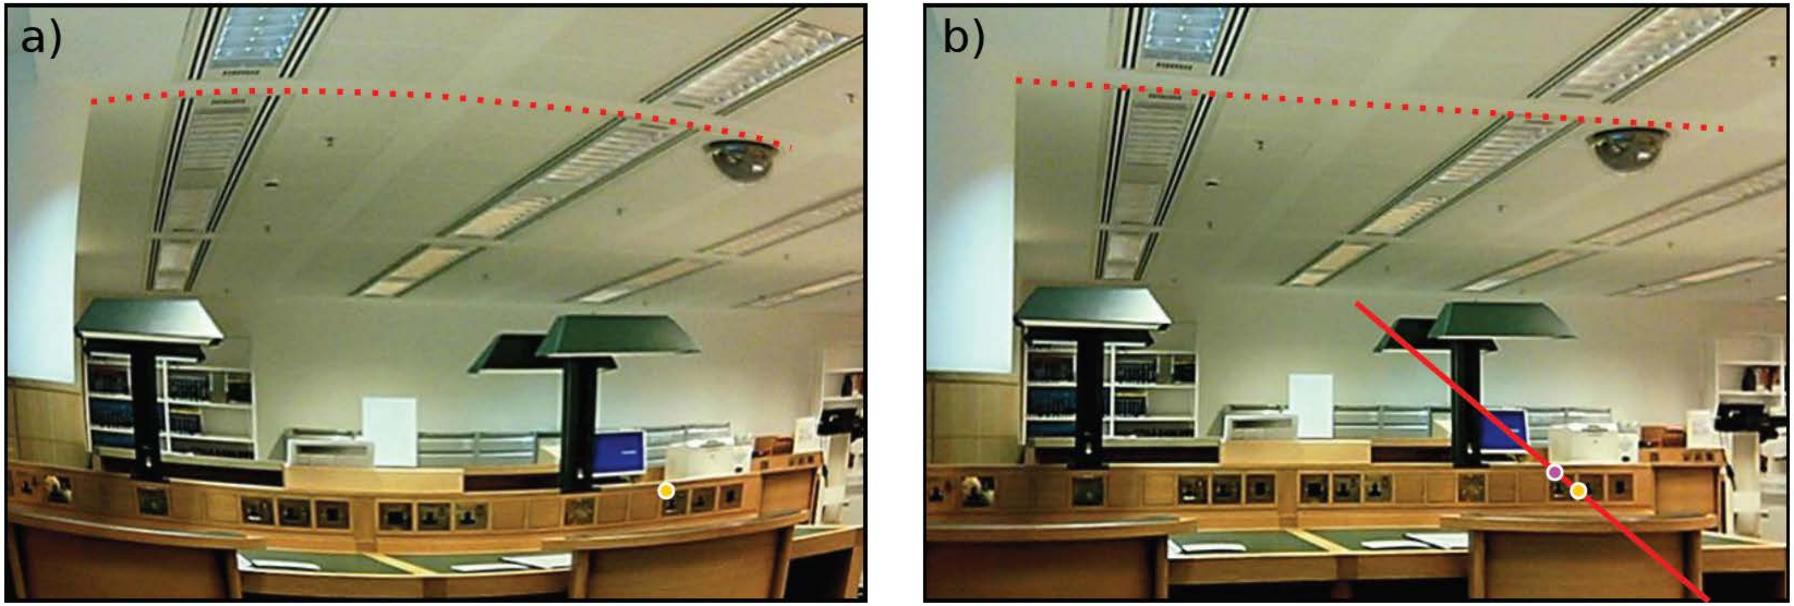

$$
\begin{aligned}
x' &= x(1+\beta_1 r^2+\beta_2 r^4) \\
y' &= y(1+\beta_1 r^2+\beta_2 r^4)
\end{aligned}
$$

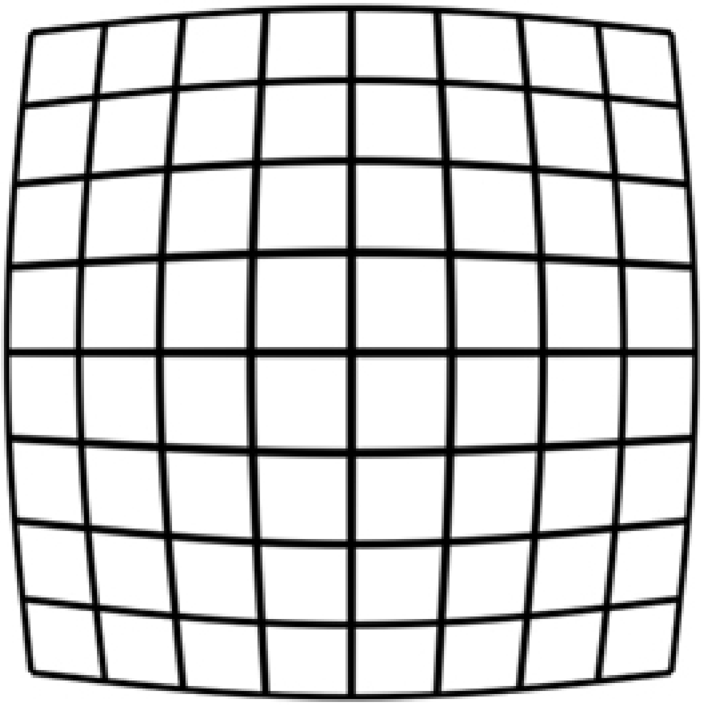

Barrel Distortion

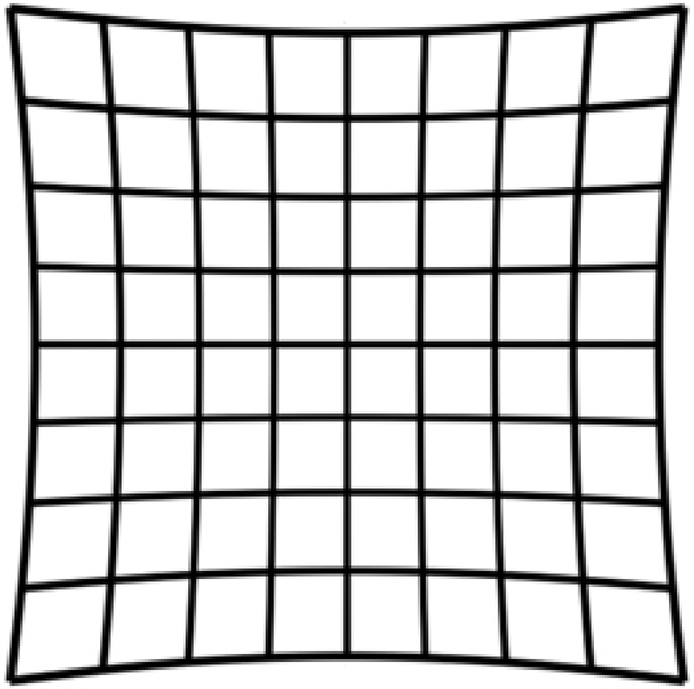

Pincushion Distortion

Structure

- Pinhole camera model
- Three geometric problems
- Homogeneous coordinates
- Solving the problems
  - Exterior orientation problem
  - Camera calibration
  - 3D reconstruction
- Applications

Problem 1: Learning extrinsic parameters (exterior orientation)

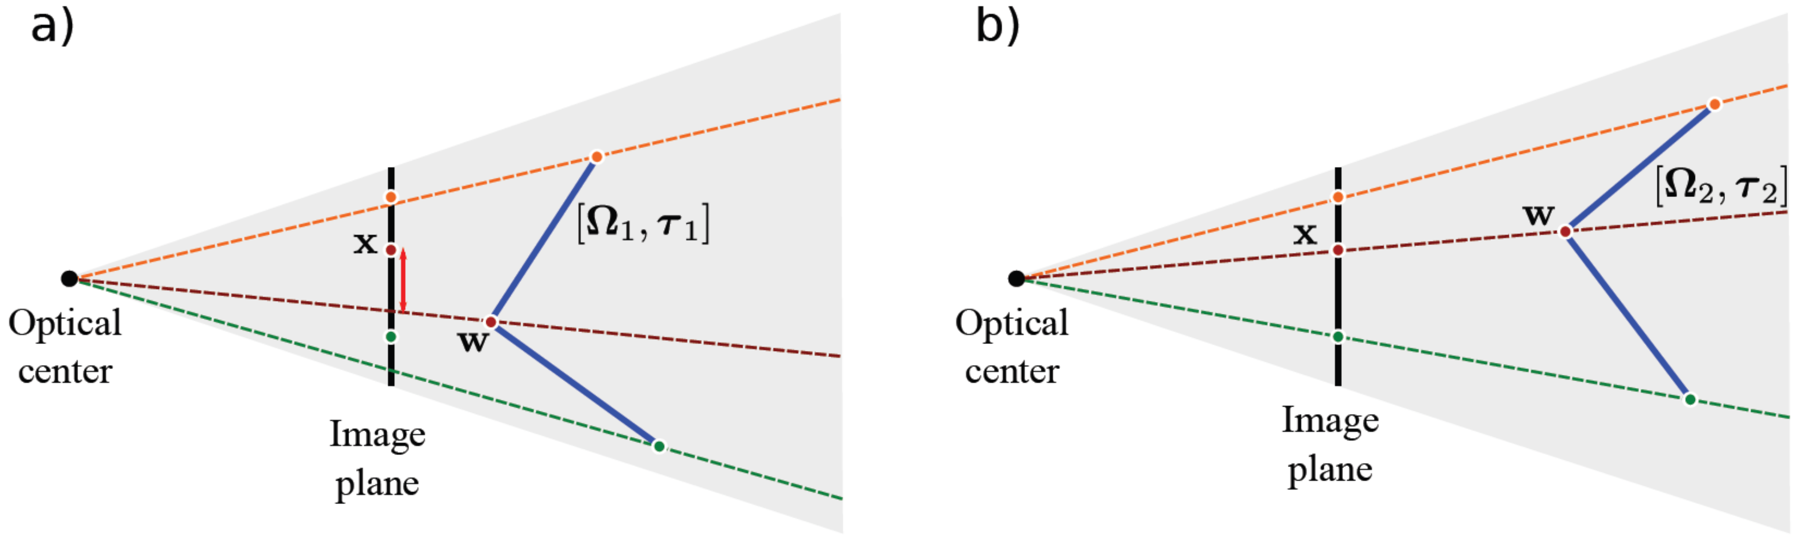

Use maximum likelihood:

$$\hat{\boldsymbol{\Omega}}, \hat{\boldsymbol{\tau}} = \underset{\boldsymbol{\Omega}, \boldsymbol{\tau}}{\operatorname{argmax}} \sum_{i=1}^{I} \log \left[Pr(\mathbf{x}_i | \mathbf{w}_i, \boldsymbol{\Lambda}, \boldsymbol{\Omega}, \boldsymbol{\tau}) \right]$$

Problem 2 – Learning intrinsic parameters (calibration)

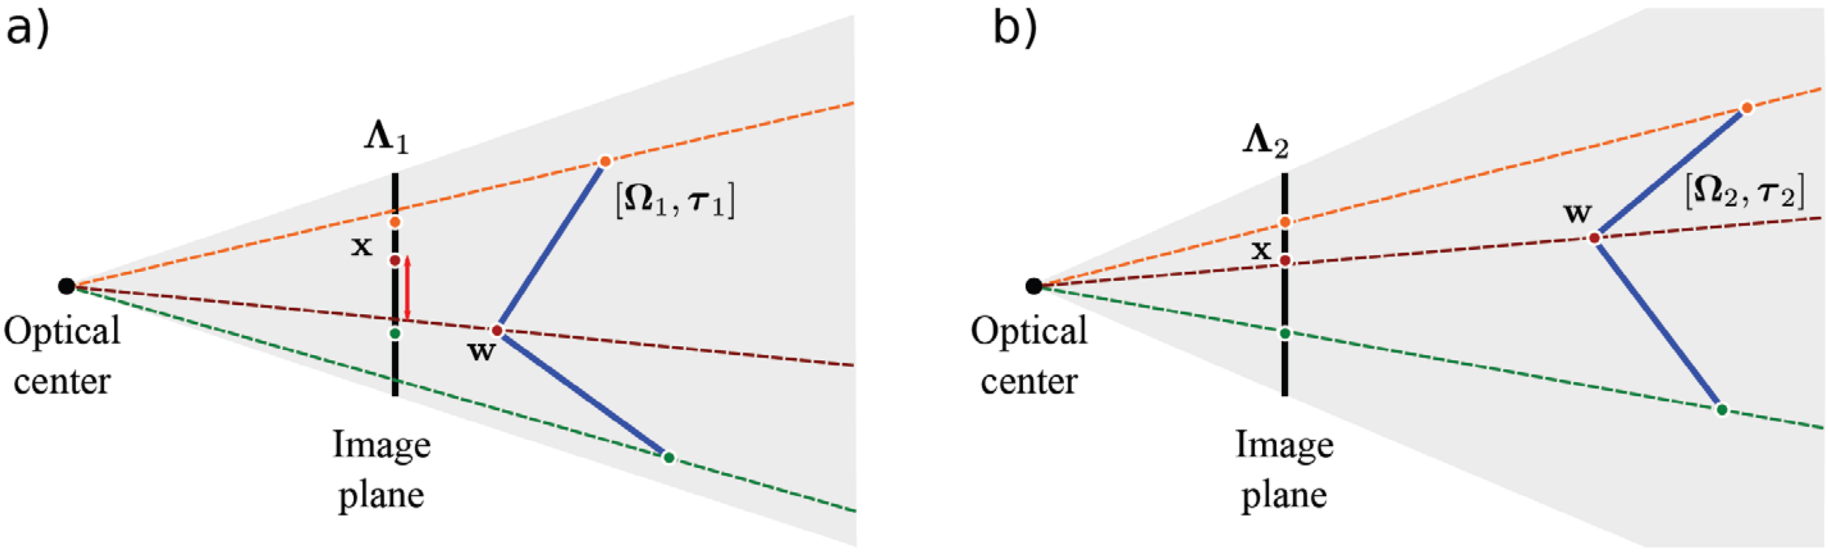

Use maximum likelihood:

$$\hat{\Lambda} = \underset{\Lambda}{\operatorname{argmax}} \left[\max_{\Omega, \tau} \sum_{i=1}^{I} \log \left[Pr(\mathbf{x}_i | \mathbf{w}_i, \Lambda, \Omega, \tau) \right] \right]$$

Calibration

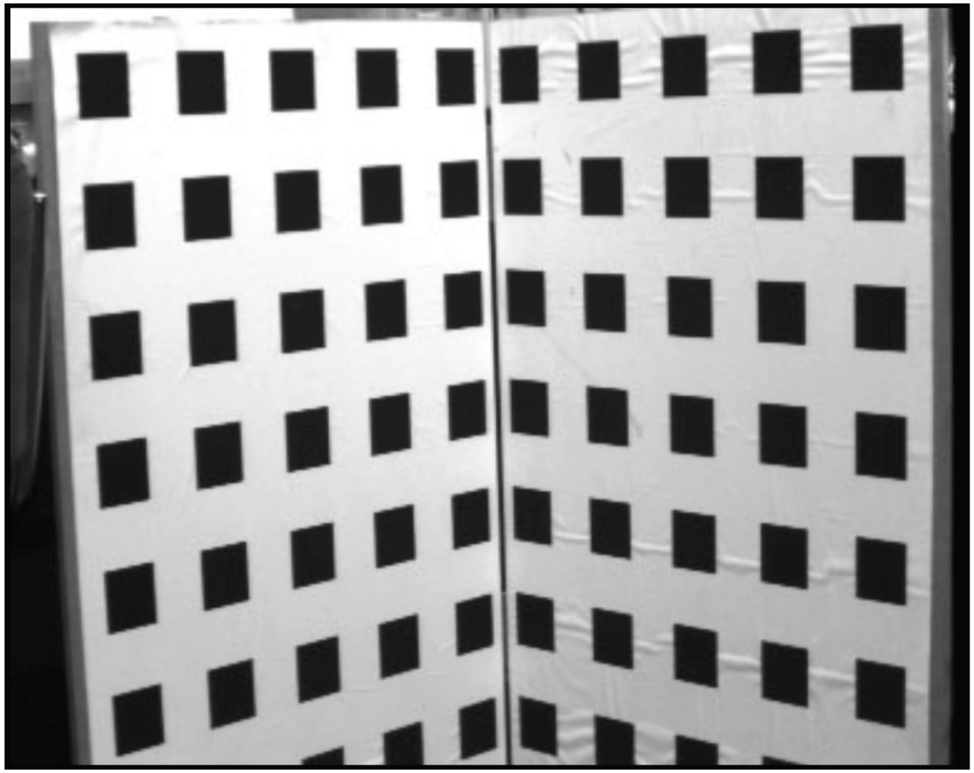

- Use 3D target with known 3D points

Problem 3 – Inferring 3D points (triangulation / reconstruction)

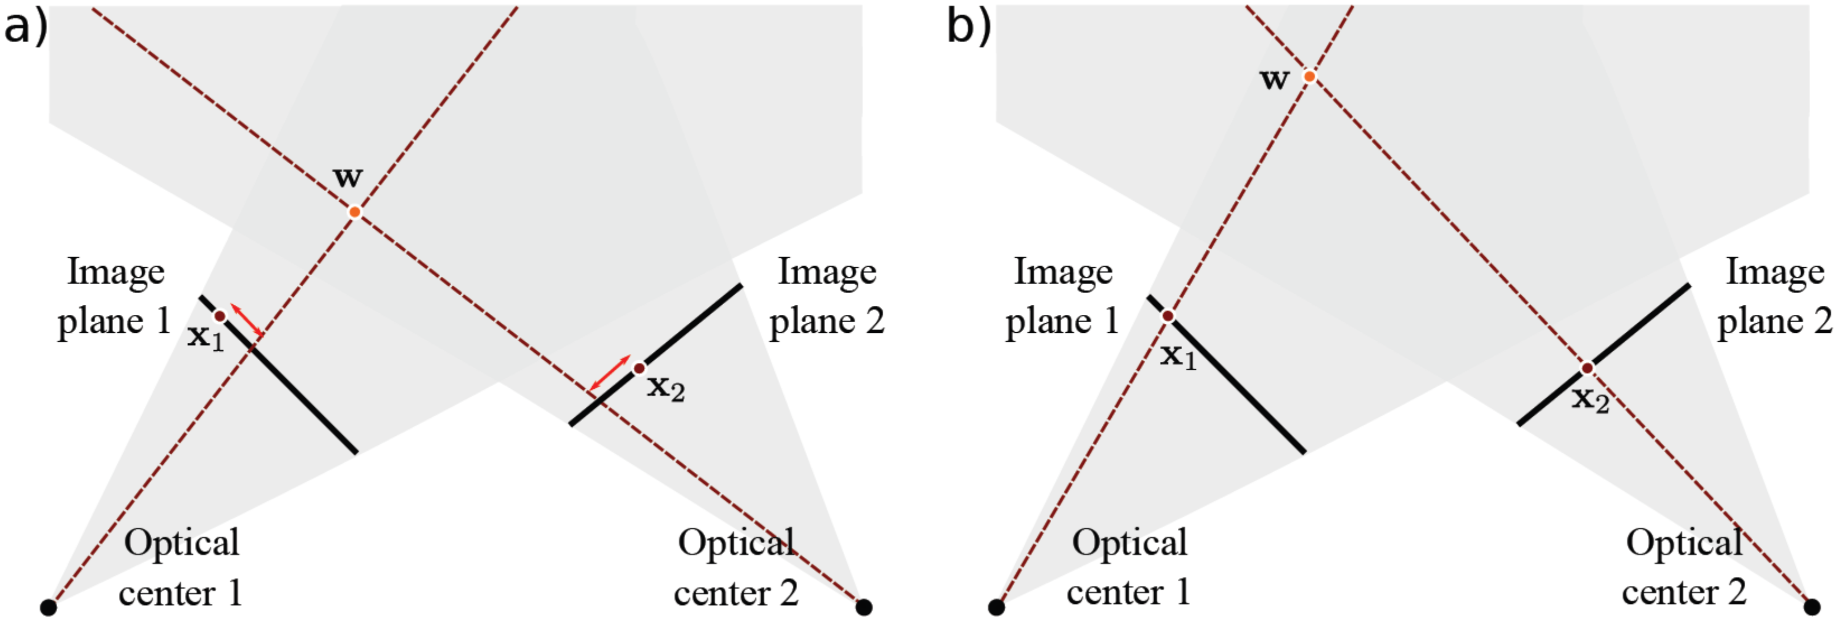

Use maximum likelihood:

$$\hat{\mathbf{w}} = \underset{\mathbf{w}}{\operatorname{argmax}} \left(\sum_{j=1}^{J} \log[Pr(\mathbf{x}_j|\mathbf{w}, \boldsymbol{\Lambda}_j, \boldsymbol{\Omega}_j, \boldsymbol{\tau}_j)] \right)$$

Solving the problems

- None of these problems can be solved in closed form
- Can apply non-linear optimization to find best solution but slow and prone to local minima
- Solution – convert to a new representation (homogeneous coordinates) where we can solve in closed form.
- Caution! We are not solving the true problem – finding global minimum of wrong* problem. But can use as starting point for non-linear optimization of true problem

*= We'll first minimize algebraic error, instead of geometric error (see Minimizing Algebraic Error in Geometric Estimation Problems, Hartley 1998)

Structure

- Pinhole camera model
- Three geometric problems
- Homogeneous coordinates
- Solving the problems
  - Exterior orientation problem
  - Camera calibration
  - 3D reconstruction
- Applications

Homogeneous coordinates

Convert 2D coordinate to 3D

$$\tilde{\mathbf{x}} = \begin{bmatrix} \lambda x \\ \lambda y \\ \lambda \end{bmatrix}$$

To convert back

$$x = \frac{\tilde{x}}{\tilde{z}} \qquad y = \frac{\tilde{y}}{\tilde{z}}$$

Geometric interpretation of homogeneous coordinates

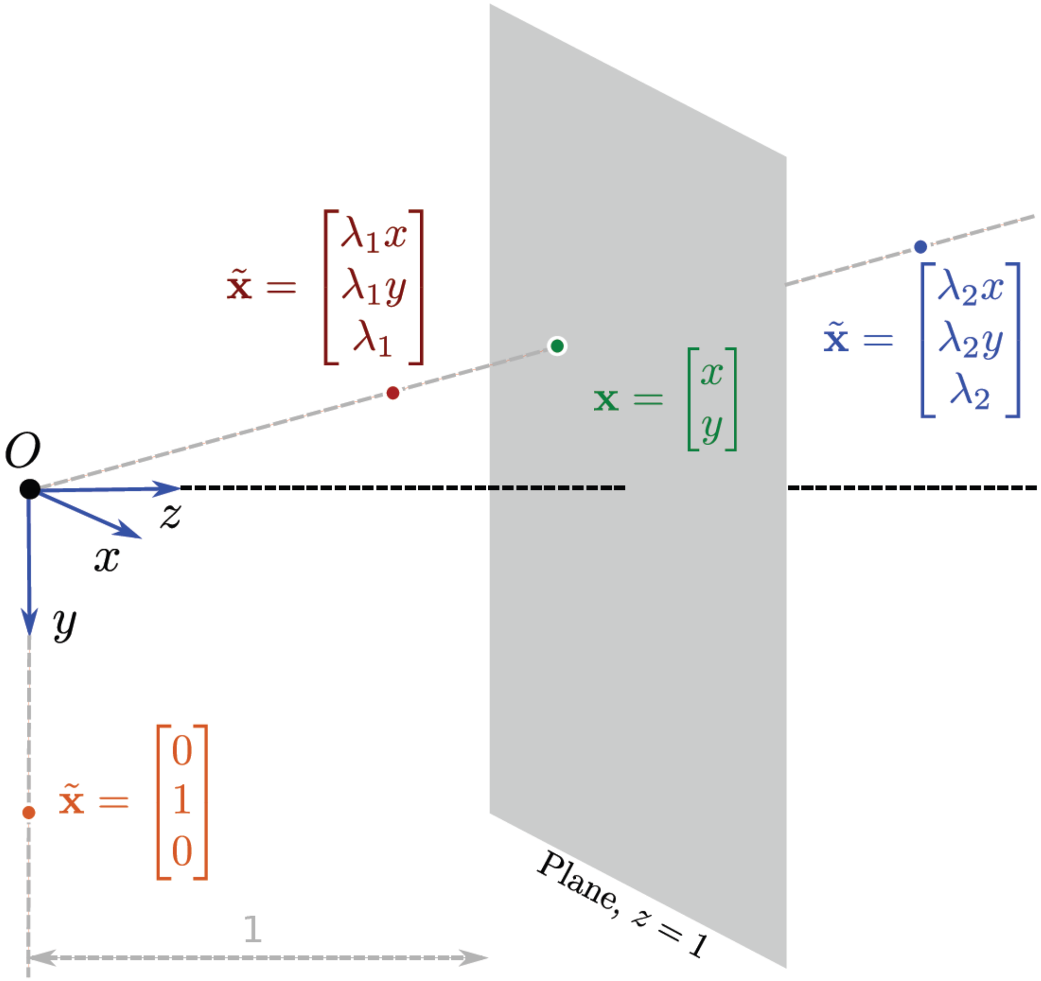
$\tilde{\mathbf{x}} = \begin{bmatrix} \lambda_1 x \\ \lambda_1 y \\ \lambda_1 \end{bmatrix}$
$\mathbf{x} = \begin{bmatrix} x \\ y \end{bmatrix}$
$\tilde{\mathbf{x}} = \begin{bmatrix} \lambda_2 x \\ \lambda_2 y \\ \lambda_2 \end{bmatrix}$
O
z
x
y
$\tilde{\mathbf{x}} = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}$
1
Plane, $z = 1$

Pinhole camera in homogeneous coordinates

Camera model:

$$\begin{aligned} x &= \frac{\phi_x u + \gamma v}{w} + \delta_x \\ y &= \frac{\phi_y v}{w} + \delta_y, \end{aligned}$$

In homogeneous coordinates:

$$\lambda \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} \phi_x & \gamma & \delta_x & 0 \\ 0 & \phi_y & \delta_y & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} u \\ v \\ w \\ 1 \end{bmatrix}$$

(linear!)

Pinhole camera in homogeneous coordinates

$$\lambda \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} \phi_x & \gamma & \delta_x & 0 \\ 0 & \phi_y & \delta_y & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} u \\ v \\ w \\ 1 \end{bmatrix}$$

Writing out these three equations

$$\begin{aligned} \lambda x &= \phi_x u + \gamma v + \delta_x w \\ \lambda y &= \phi_y v + \delta_y w \\ \lambda &= w. \end{aligned}$$

Eliminate λ to retrieve original equations

Adding in extrinsic parameters

$$\lambda \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} \phi_x & \gamma & \delta_x & 0 \\ 0 & \phi_y & \delta_y & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} \omega_{11} & \omega_{12} & \omega_{13} & \tau_x \\ \omega_{21} & \omega_{22} & \omega_{23} & \tau_y \\ \omega_{31} & \omega_{32} & \omega_{33} & \tau_z \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} u \\ v \\ w \\ 1 \end{bmatrix}$$

Or for short:

$$\lambda\tilde{\mathbf{x}} = \begin{bmatrix} \boldsymbol{\Lambda} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \boldsymbol{\Omega} & \boldsymbol{\tau} \\ \mathbf{0}^T & 1 \end{bmatrix} \tilde{\mathbf{w}}$$

Or even shorter:

$$\lambda\tilde{\mathbf{x}} = \boldsymbol{\Lambda} \begin{bmatrix} \boldsymbol{\Omega} & \boldsymbol{\tau} \end{bmatrix} \tilde{\mathbf{w}}$$

Structure

- Pinhole camera model
- Three geometric problems
- Homogeneous coordinates
- Solving the problems
  - Exterior orientation problem
  - Camera calibration
  - 3D reconstruction
- Applications

Problem 1: Learning extrinsic parameters (exterior orientation)

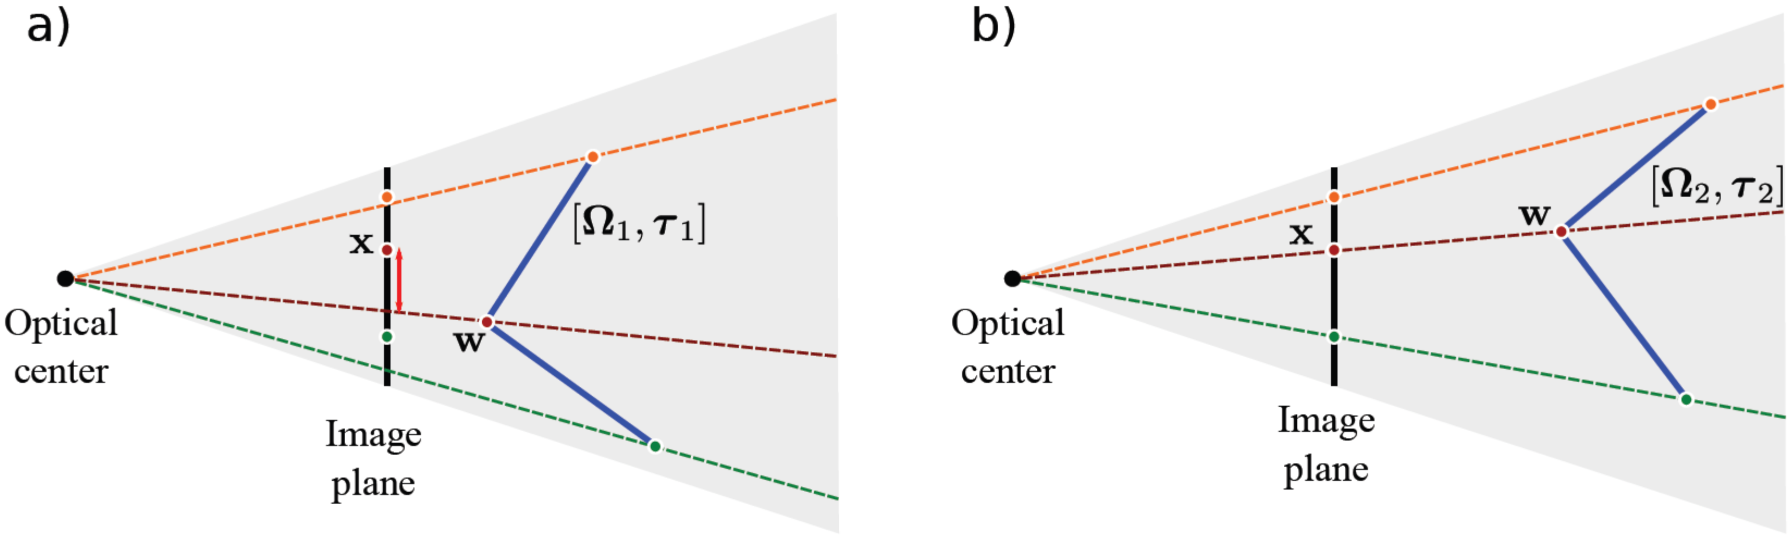

Use maximum likelihood:

$$\hat{\boldsymbol{\Omega}}, \hat{\boldsymbol{\tau}} = \underset{\boldsymbol{\Omega}, \boldsymbol{\tau}}{\operatorname{argmax}} \sum_{i=1}^{I} \log \left[Pr(\mathbf{x}_i | \mathbf{w}_i, \boldsymbol{\Lambda}, \boldsymbol{\Omega}, \boldsymbol{\tau}) \right]$$

Exterior orientation

Start with camera equation in homogeneous coordinates

$$\lambda_i \begin{bmatrix} x_i \\ y_i \\ 1 \end{bmatrix} = \begin{bmatrix} \phi_x & \gamma & \delta_x \\ 0 & \phi_y & \delta_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \omega_{11} & \omega_{12} & \omega_{13} & \tau_x \\ \omega_{21} & \omega_{22} & \omega_{23} & \tau_y \\ \omega_{31} & \omega_{32} & \omega_{33} & \tau_z \end{bmatrix} \begin{bmatrix} u_i \\ v_i \\ w_i \\ 1 \end{bmatrix}$$

Pre-multiply both sides by inverse of camera calibration matrix

$$\lambda_i \begin{bmatrix} x'_i \\ y'_i \\ 1 \end{bmatrix} = \begin{bmatrix} \omega_{11} & \omega_{12} & \omega_{13} & \tau_x \\ \omega_{21} & \omega_{22} & \omega_{23} & \tau_y \\ \omega_{31} & \omega_{32} & \omega_{33} & \tau_z \end{bmatrix} \begin{bmatrix} u_i \\ v_i \\ w_i \\ 1 \end{bmatrix}$$

Exterior orientation

$$\lambda_i \begin{bmatrix} x'_i \\ y'_i \\ 1 \end{bmatrix} = \begin{bmatrix} \omega_{11} & \omega_{12} & \omega_{13} & \tau_x \\ \omega_{21} & \omega_{22} & \omega_{23} & \tau_y \\ \omega_{31} & \omega_{32} & \omega_{33} & \tau_z \end{bmatrix} \begin{bmatrix} u_i \\ v_i \\ w_i \\ 1 \end{bmatrix}$$

The third equation gives us an expression for λ

$$\lambda_i = \omega_{31} u_i + \omega_{32} v_i + \omega_{33} w_i + \tau_z$$

Substitute back into first two lines

$$\begin{bmatrix} (\omega_{31} u_i + \omega_{32} v_i + \omega_{33} w_i + \tau_z) x'_i \\ (\omega_{31} u_i + \omega_{32} v_i + \omega_{33} w_i + \tau_z) y'_i \end{bmatrix} = \begin{bmatrix} \omega_{11} & \omega_{12} & \omega_{13} & \tau_x \\ \omega_{21} & \omega_{22} & \omega_{23} & \tau_y \end{bmatrix} \begin{bmatrix} u_i \\ v_i \\ w_i \\ 1 \end{bmatrix}$$

Exterior orientation

$$\begin{bmatrix}(\omega_{31}u_i+\omega_{32}v_i+\omega_{33}w_i+\tau_z)x'_i\\(\omega_{31}u_i+\omega_{32}v_i+\omega_{33}w_i+\tau_z)y'_i\end{bmatrix}=\begin{bmatrix}\omega_{11}&\omega_{12}&\omega_{13}&\tau_x\\\omega_{21}&\omega_{22}&\omega_{23}&\tau_y\end{bmatrix}\begin{bmatrix}u_i\\v_i\\w_i\\1\end{bmatrix}$$

Linear equation – two equations per point – form system of equations

$$\begin{bmatrix}u_1&v_1&w_1&1&0&0&0&0&-u_1x'_1&-v_1x'_1&-w_1x'_1&-x'_1\\0&0&0&0&u_1&v_1&w_1&1&-u_1y'_1&-v_1y'_1&-w_1y'_1&-y'_1\\u_2&v_2&w_2&1&0&0&0&0&-u_2x'_2&-v_2x'_2&-w_2x'_2&-x'_2\\0&0&0&0&u_2&v_2&w_2&1&-u_2y'_2&-v_2y'_2&-w_2y'_2&-y'_2\\\vdots&\vdots&\vdots&\vdots&\vdots&\vdots&\vdots&\vdots&\vdots&\vdots&\vdots&\vdots\\u_I&v_I&w_I&1&0&0&0&0&-u_Ix'_I&-v_Ix'_I&-w_Ix'_I&-x'_I\\0&0&0&0&u_I&v_I&w_I&1&-u_Iy'_I&-v_Iy'_I&-w_Iy'_I&-y'_I\end{bmatrix}\begin{bmatrix}\omega_{11}\\\omega_{12}\\\omega_{13}\\\tau_x\\\omega_{21}\\\omega_{22}\\\omega_{23}\\\tau_y\\\omega_{31}\\\omega_{32}\\\omega_{33}\\\tau_z\end{bmatrix}=\mathbf{0}$$

Exterior orientation

$$\begin{bmatrix} u_1 & v_1 & w_1 & 1 & 0 & 0 & 0 & 0 & -u_1x'_1 & -v_1x'_1 & -w_1x'_1 & -x'_1 \\ 0 & 0 & 0 & 0 & u_1 & v_1 & w_1 & 1 & -u_1y'_1 & -v_1y'_1 & -w_1y'_1 & -y'_1 \\ u_2 & v_2 & w_2 & 1 & 0 & 0 & 0 & 0 & -u_2x'_2 & -v_2x'_2 & -w_2x'_2 & -x'_2 \\ 0 & 0 & 0 & 0 & u_2 & v_2 & w_2 & 1 & -u_2y'_2 & -v_2y'_2 & -w_2y'_2 & -y'_2 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ u_I & v_I & w_I & 1 & 0 & 0 & 0 & 0 & -u_Ix'_I & -v_Ix'_I & -w_Ix'_I & -x'_I \\ 0 & 0 & 0 & 0 & u_I & v_I & w_I & 1 & -u_Iy'_I & -v_Iy'_I & -w_Iy'_I & -y'_I \end{bmatrix} \begin{bmatrix} \omega_{11} \\ \omega_{12} \\ \omega_{13} \\ \tau_x \\ \omega_{21} \\ \omega_{22} \\ \omega_{23} \\ \tau_y \\ \omega_{31} \\ \omega_{32} \\ \omega_{33} \\ \tau_z \end{bmatrix} = \mathbf{0}$$

Minimum direction problem of the form $\mathbf{Ab} = \mathbf{0}$,
Find minimum of $|\mathbf{Ab}|^2$ subject to $|\mathbf{b}| = 1$.

To solve, compute the SVD $\mathbf{A} = \mathbf{ULV}^T$ and then set $\hat{\mathbf{b}}$ to the last column of $\mathbf{V}$.

Exterior orientation

Now we extract the values of $\boldsymbol{\Omega}$ and $\boldsymbol{\tau}$ from $\mathbf{V}$.

Problem: the scale is arbitrary and the rows and columns of the rotation matrix may not be orthogonal.

Solution: compute SVD $\boldsymbol{\Omega} = \mathbf{ULV}^T$ and then choose $\hat{\boldsymbol{\Omega}} = \mathbf{UV}^T$.

Use the ratio between the rotation matrix before and after to rescale

$$\hat{\boldsymbol{\tau}} = \sum_{m=1}^{3}\sum_{n=1}^{3}\frac{\hat{\Omega}_{mn}}{\Omega_{mn}}\boldsymbol{\tau}$$

Use these estimates for start of non-linear optimisation.

Structure

- Pinhole camera model
- Three geometric problems
- Homogeneous coordinates
- Solving the problems
  - Exterior orientation problem
  - Camera calibration
  - 3D reconstruction
- Applications

Problem 2 – Learning intrinsic parameters (calibration)

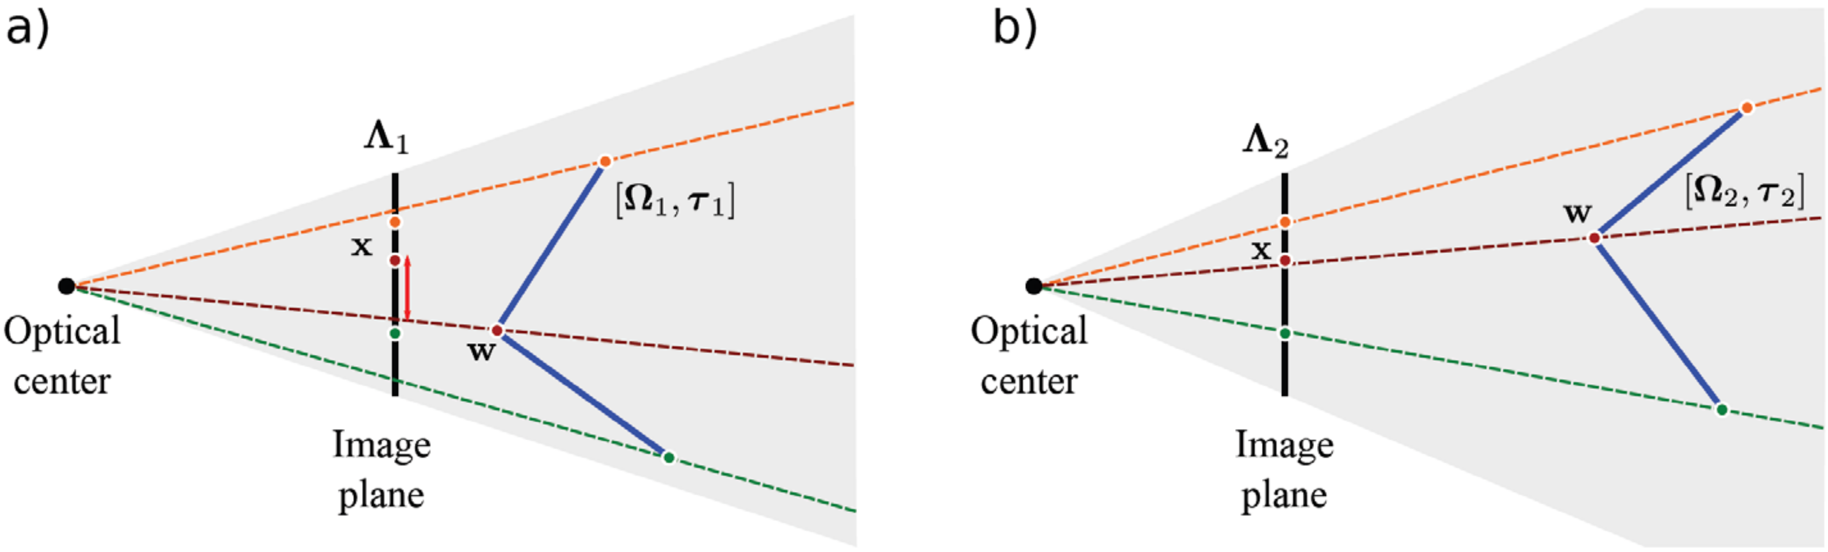

Use maximum likelihood:

$$\hat{\mathbf{\Lambda}} = \underset{\mathbf{\Lambda}}{\operatorname{argmax}} \left[\max_{\mathbf{\Omega}, \boldsymbol{\tau}} \sum_{i=1}^{I} \log\left[Pr(\mathbf{x}_i | \mathbf{w}_i, \mathbf{\Lambda}, \mathbf{\Omega}, \boldsymbol{\tau}) \right] \right]$$

Calibration

$$\hat{\boldsymbol\Lambda} = \underset{\boldsymbol\Lambda}{\operatorname{argmax}}\left[\max_{\boldsymbol\Omega,\boldsymbol\tau}\sum_{i=1}^{I}\log\left[Pr(\mathbf{x}_i|\mathbf{w}_i,\boldsymbol\Lambda,\boldsymbol\Omega,\boldsymbol\tau)\right]\right]$$

One approach (not very efficient) is to alternately

- Optimize extrinsic parameters for fixed intrinsic

$$\hat{\boldsymbol\Omega},\hat{\boldsymbol\tau} = \underset{\boldsymbol\Omega,\boldsymbol\tau}{\operatorname{argmax}}\left[\sum_{i=1}^{I}\log\left[Pr(\mathbf{x}_i|\mathbf{w}_i,\boldsymbol\Lambda,\boldsymbol\Omega,\boldsymbol\tau)\right]\right]$$

- Optimize intrinsic parameters for fixed extrinsic

$$\hat{\boldsymbol\Lambda} = \underset{\boldsymbol\Lambda}{\operatorname{argmax}}\left[\sum_{i=1}^{I}\log\left[Pr(\mathbf{x}_i|\mathbf{w}_i,\boldsymbol\Lambda,\boldsymbol\Omega,\boldsymbol\tau)\right]\right]$$

(Then use non-linear optimization)

Intrinsic parameters

$$\hat{\boldsymbol{\Lambda}} = \underset{\boldsymbol{\Lambda}}{\mathrm{argmax}}\left[\sum_{i=1}^{I}\log\left[Pr(\mathbf{x}_i|\mathbf{w}_i,\boldsymbol{\Lambda},\boldsymbol{\Omega},\boldsymbol{\tau})\right]\right]$$

Maximum likelihood approach

$$\begin{aligned}\hat{\boldsymbol{\Lambda}} &= \underset{\boldsymbol{\Lambda}}{\mathrm{argmax}}\left[\sum_{i=1}^{I}\log\left[\mathrm{Norm}_{\mathbf{x}_i}\left[\mathbf{pinhole}[\mathbf{w}_i,\boldsymbol{\Lambda},\boldsymbol{\Omega},\boldsymbol{\tau}],\sigma^2\mathbf{I}\right]\right]\right]\\ &= \underset{\boldsymbol{\Lambda}}{\mathrm{argmin}}\left[\sum_{i=1}^{I}(\mathbf{x}_i-\mathbf{pinhole}[\mathbf{w}_i,\boldsymbol{\Lambda},\boldsymbol{\Omega},\boldsymbol{\tau}])^T(\mathbf{x}_i-\mathbf{pinhole}[\mathbf{w}_i,\boldsymbol{\Lambda},\boldsymbol{\Omega},\boldsymbol{\tau}])\right]\end{aligned}$$

This is a least squares problem.

Intrinsic parameters

$$
\begin{aligned}
\hat{\boldsymbol{\Lambda}} &= \underset{\boldsymbol{\Lambda}}{\operatorname{argmax}} \left[\sum_{i=1}^{I} \log \left[\operatorname{Norm}_{\mathbf{x}_i} \left[\mathbf{pinhole}[\mathbf{w}_i, \boldsymbol{\Lambda}, \boldsymbol{\Omega}, \boldsymbol{\tau}], \sigma^2 \mathbf{I} \right] \right] \right] \\
&= \underset{\boldsymbol{\Lambda}}{\operatorname{argmin}} \left[\sum_{i=1}^{I} (\mathbf{x}_i - \mathbf{pinhole}[\mathbf{w}_i, \boldsymbol{\Lambda}, \boldsymbol{\Omega}, \boldsymbol{\tau}])^T (\mathbf{x}_i - \mathbf{pinhole}[\mathbf{w}_i, \boldsymbol{\Lambda}, \boldsymbol{\Omega}, \boldsymbol{\tau}]) \right]
\end{aligned}
$$

The function $\mathbf{pinhole}[\bullet, \bullet, \bullet, \bullet]$ is linear w.r.t. intrinsic parameters. Can be written in form $\mathbf{A}_i \mathbf{h}$

$$
\mathrm{A}_i = \begin{bmatrix} \frac{\omega_{11} u_i + \omega_{12} v_i + \omega_{13} w_i + \tau_x}{\omega_{31} u_i + \omega_{32} v_i + \omega_{33} w_i + \tau_z} & \frac{\omega_{21} u_i + \omega_{22} v_i + \omega_{23} w_i + \tau_x}{\omega_{31} u_i + \omega_{32} v_i + \omega_{33} w_i + \tau_z} & 1 & 0 & 0 \\ 0 & 0 & 0 & \frac{\omega_{21} u_i + \omega_{22} v_i + \omega_{23} w_i + \tau_y}{\omega_{31} u_i + \omega_{32} v_i + \omega_{33} w_i + \tau_z} & 1 \end{bmatrix}
$$

$$
\mathbf{h} = [\phi_x, \gamma, \delta_x, \phi_y, \delta_y]^T
$$

Now solve least squares problem

$$
\hat{\mathbf{h}} = \underset{\mathbf{h}}{\operatorname{argmin}} \left[\sum_{i=1}^{I} (\mathbf{A}_i \mathbf{h} - \mathbf{x}_i)^T (\mathbf{A}_i \mathbf{h} - \mathbf{x}_i) \right]
$$

Structure

- Pinhole camera model
- Three geometric problems
- Homogeneous coordinates
- Solving the problems
  - Exterior orientation problem
  - Camera calibration
  - 3D reconstruction
- Applications

Problem 3 – Inferring 3D points (triangulation / reconstruction)

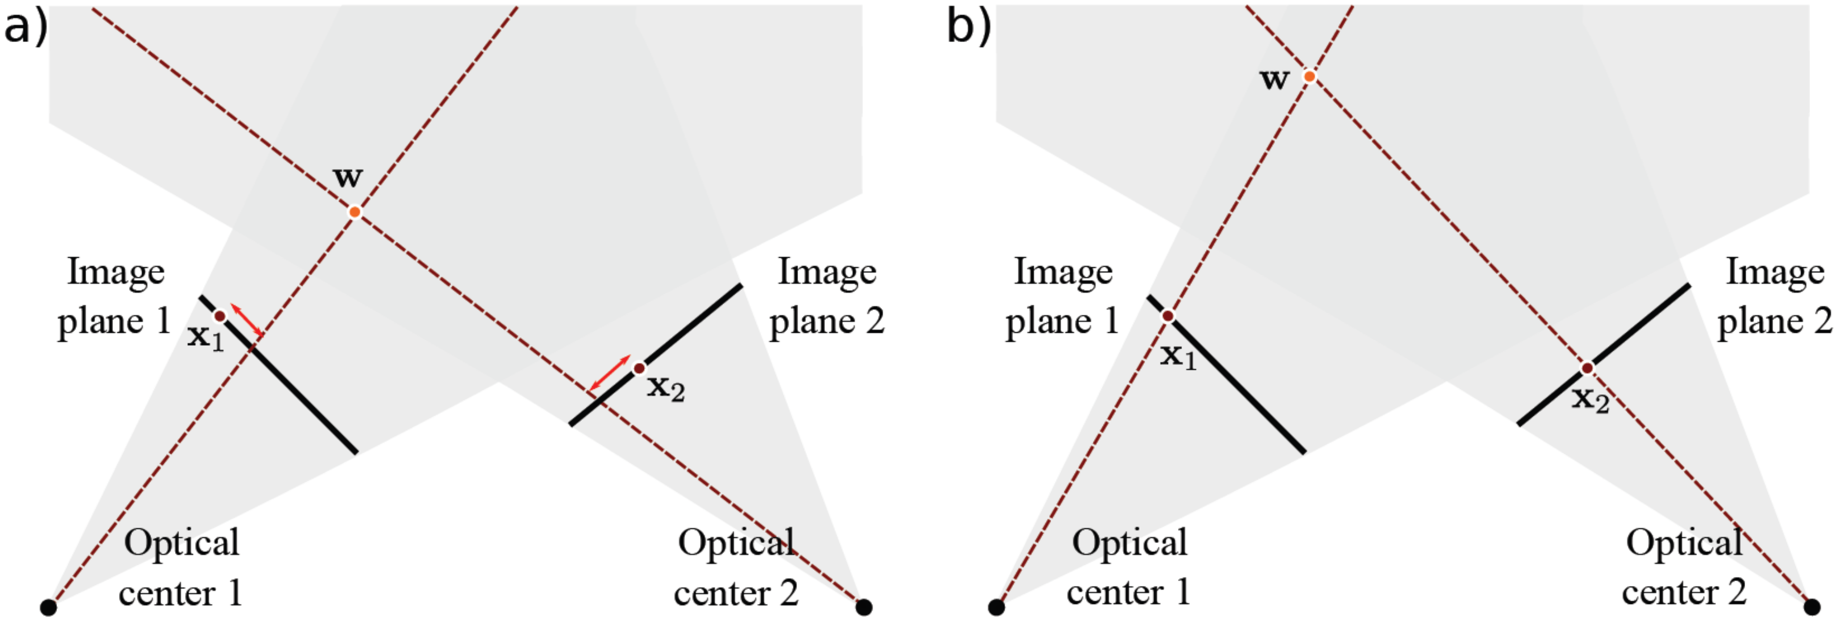

Use maximum likelihood:

$$\hat{\mathbf{w}} = \underset{\mathbf{w}}{\operatorname{argmax}} \left[\sum_{j=1}^{J} \log[Pr(\mathbf{x}_j|\mathbf{w}, \boldsymbol{\Lambda}_j, \boldsymbol{\Omega}_j, \boldsymbol{\tau}_j)] \right]$$

Reconstruction

$$\hat{\mathbf{w}} = \underset{\mathbf{w}}{\operatorname{argmax}}\left[\sum_{j=1}^{J}\log[Pr(\mathbf{x}_j|\mathbf{w},\boldsymbol{\Lambda}_j,\boldsymbol{\Omega}_j,\boldsymbol{\tau}_j)]\right]$$

Write j^th pinhole camera in homogeneous coordinates:

$$\lambda_j\begin{bmatrix}x_j\\y_j\\1\end{bmatrix} = \begin{bmatrix}\phi_{xj} & \gamma_j & \delta_{xj}\\0 & \phi_{yj} & \delta_{yj}\\0 & 0 & 1\end{bmatrix}\begin{bmatrix}\omega_{11j} & \omega_{12j} & \omega_{13j} & \tau_{xj}\\\omega_{21j} & \omega_{22j} & \omega_{23j} & \tau_{yj}\\\omega_{31j} & \omega_{32j} & \omega_{33j} & \tau_{zj}\end{bmatrix}\begin{bmatrix}u\\v\\w\\1\end{bmatrix}$$

Pre-multiply with inverse of intrinsic matrix

$$\lambda_j\begin{bmatrix}x'_j\\y'_j\\1\end{bmatrix} = \begin{bmatrix}\omega_{11j} & \omega_{12j} & \omega_{13j} & \tau_{xj}\\\omega_{21j} & \omega_{22j} & \omega_{23j} & \tau_{yj}\\\omega_{31j} & \omega_{32j} & \omega_{33j} & \tau_{zj}\end{bmatrix}\begin{bmatrix}u\\v\\w\\1\end{bmatrix}$$

Reconstruction

$$\lambda_j \begin{bmatrix} x'_j \\ y'_j \\ 1 \end{bmatrix} = \begin{bmatrix} \omega_{11j} & \omega_{12j} & \omega_{13j} & \tau_{xj} \\ \omega_{21j} & \omega_{22j} & \omega_{23j} & \tau_{yj} \\ \omega_{31j} & \omega_{32j} & \omega_{33j} & \tau_{zj} \end{bmatrix} \begin{bmatrix} u \\ v \\ w \\ 1 \end{bmatrix}$$

Last equations gives $\lambda_j = \omega_{31j}u + \omega_{32j}v + \omega_{33j}w + \tau_{zj}$

Substitute back into first two equations

$$\begin{bmatrix} (\omega_{31j}u + \omega_{32j}v + \omega_{33j}w + \tau_{zj})x'_j \\ (\omega_{31j}u + \omega_{32j}v + \omega_{33j}w + \tau_{zj})y'_j \end{bmatrix} = \begin{bmatrix} \omega_{11j} & \omega_{12j} & \omega_{13j} & \tau_{xj} \\ \omega_{21j} & \omega_{22j} & \omega_{23j} & \tau_{yj} \end{bmatrix} \begin{bmatrix} u \\ v \\ w \\ 1 \end{bmatrix}$$

Re-arranging get two linear equations for [u,v,w]

$$\begin{bmatrix} \omega_{31j}x'_j - \omega_{11j} & \omega_{32j}x'_j - \omega_{12j} & \omega_{33j}x'_j - \omega_{13j} \\ \omega_{31j}y'_j - \omega_{21j} & \omega_{32j}y'_j - \omega_{22j} & \omega_{33j}y'_j - \omega_{23j} \end{bmatrix} \begin{bmatrix} u \\ v \\ w \end{bmatrix} = \begin{bmatrix} \tau_{xj} - \tau_{zj}x'_j \\ \tau_{yj} - \tau_{zj}y'_j \end{bmatrix}$$

Solve using >1 cameras and then use non-linear optimization

Structure

- Pinhole camera model
- Three geometric problems
- Homogeneous coordinates
- Solving the problems
  - Exterior orientation problem
  - Camera calibration
  - 3D reconstruction
- Applications

Depth from structured light

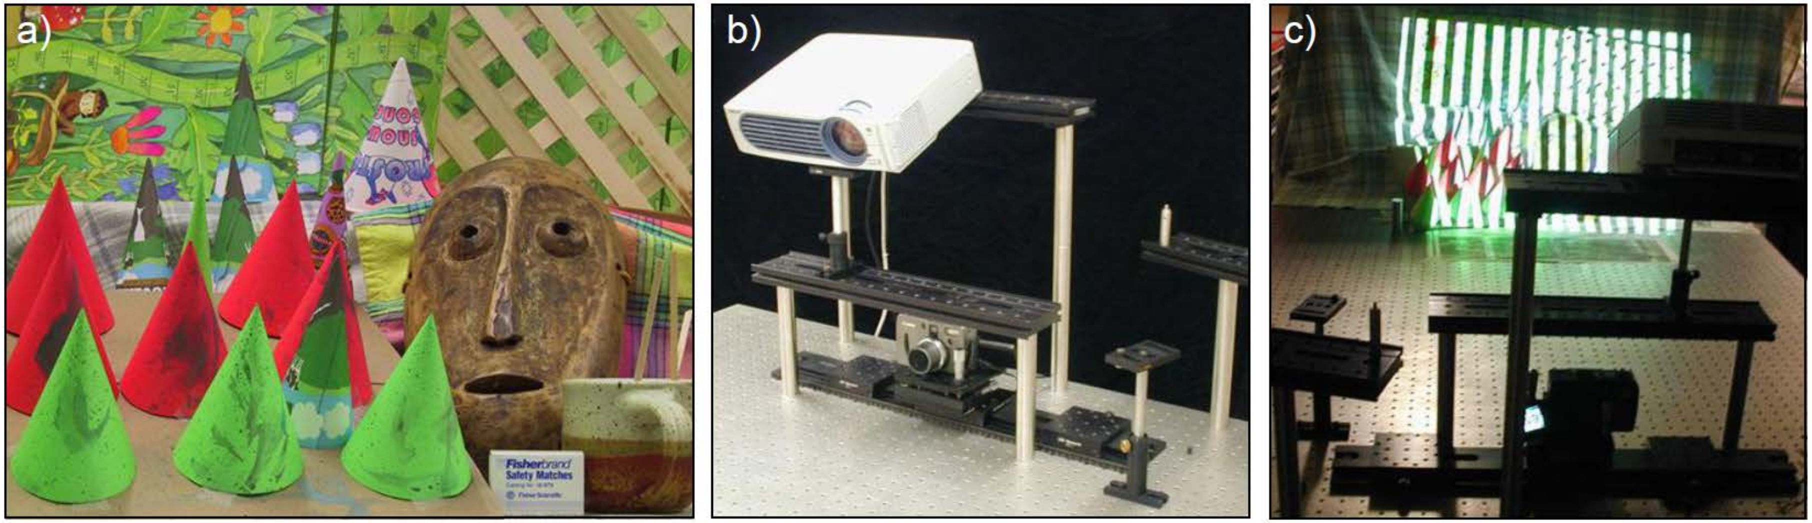

Figure 13.12 Depth maps from structured light. a) A three-dimensional scene that we wish to capture. b) The capture hardware consists of a projector and a camera, which both view the scene from different positions. c) The projector is used to illuminate the scene and the camera records the pattern of illumination from its viewpoint. The resulting images contain information that can be used to compute a 3D reconstruction. Adapted from Scharstein & Szeliski (2003). ©2003 IEEE.

Depth from structured light

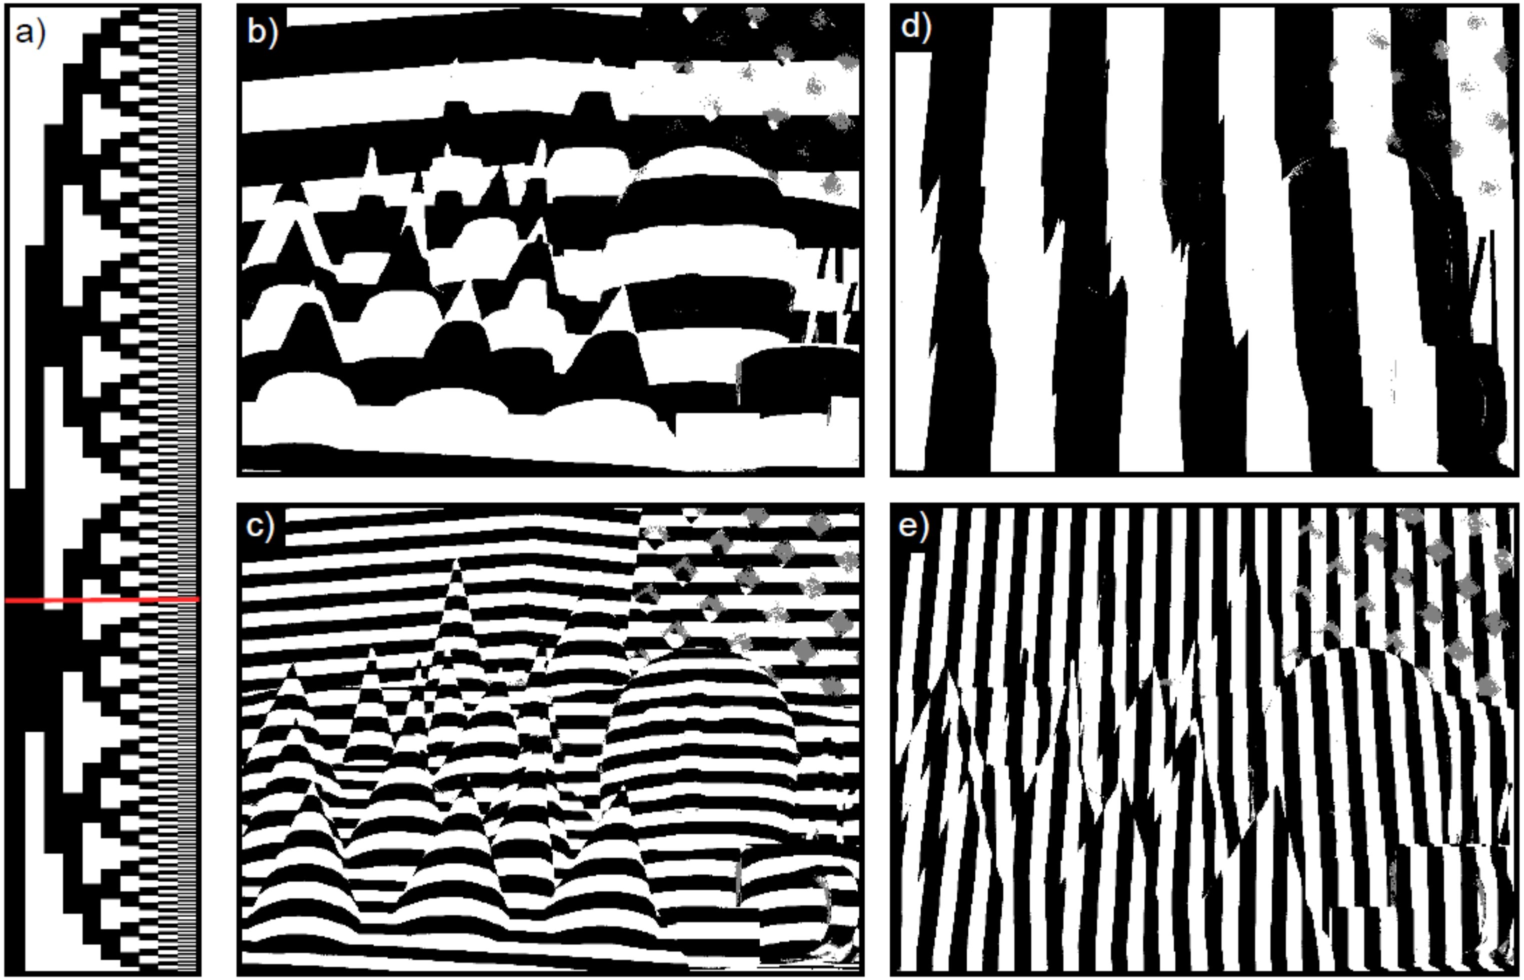

Depth from structured light

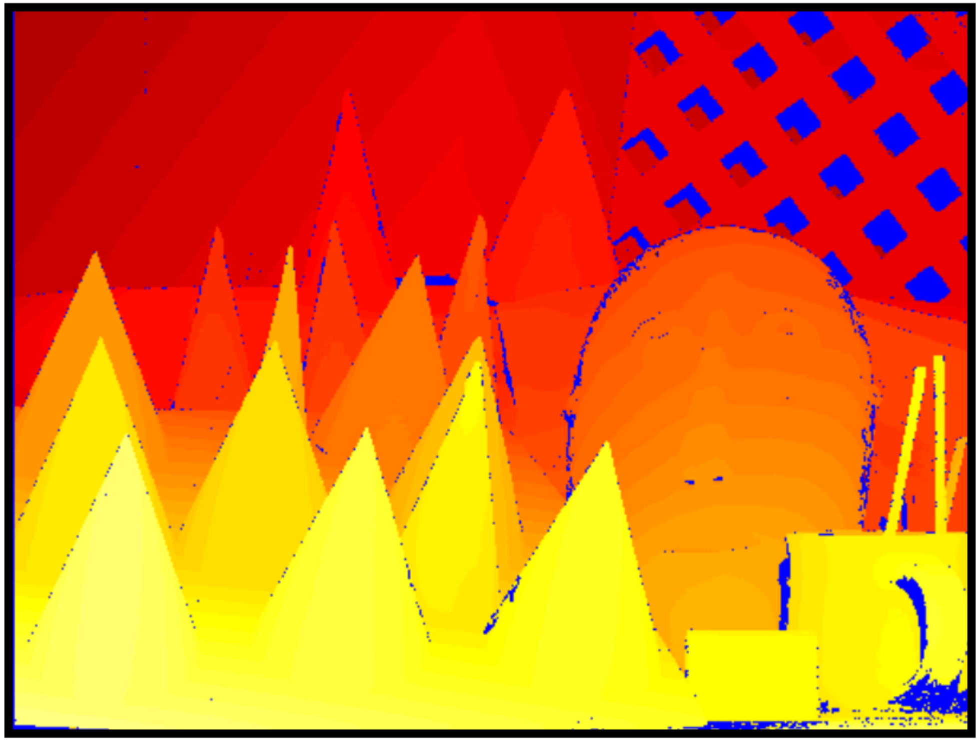

Figure 13.14 Recovered depth map for scene in figure 13.12 using the structured light method. Pixels marked as blue are places where the depth is uncertain: these include positions in the image that were occluded with respect to the projector and so no light was cast onto them. Scharstein & Szeliski (2003) also captured the scene with two cameras under normal illumination; they subsequently used the depth map from the structured light as ground truth data for assessing stereo vision algorithms. Adapted from Scharstein & Szeliski (2003).

Shape from silhouette

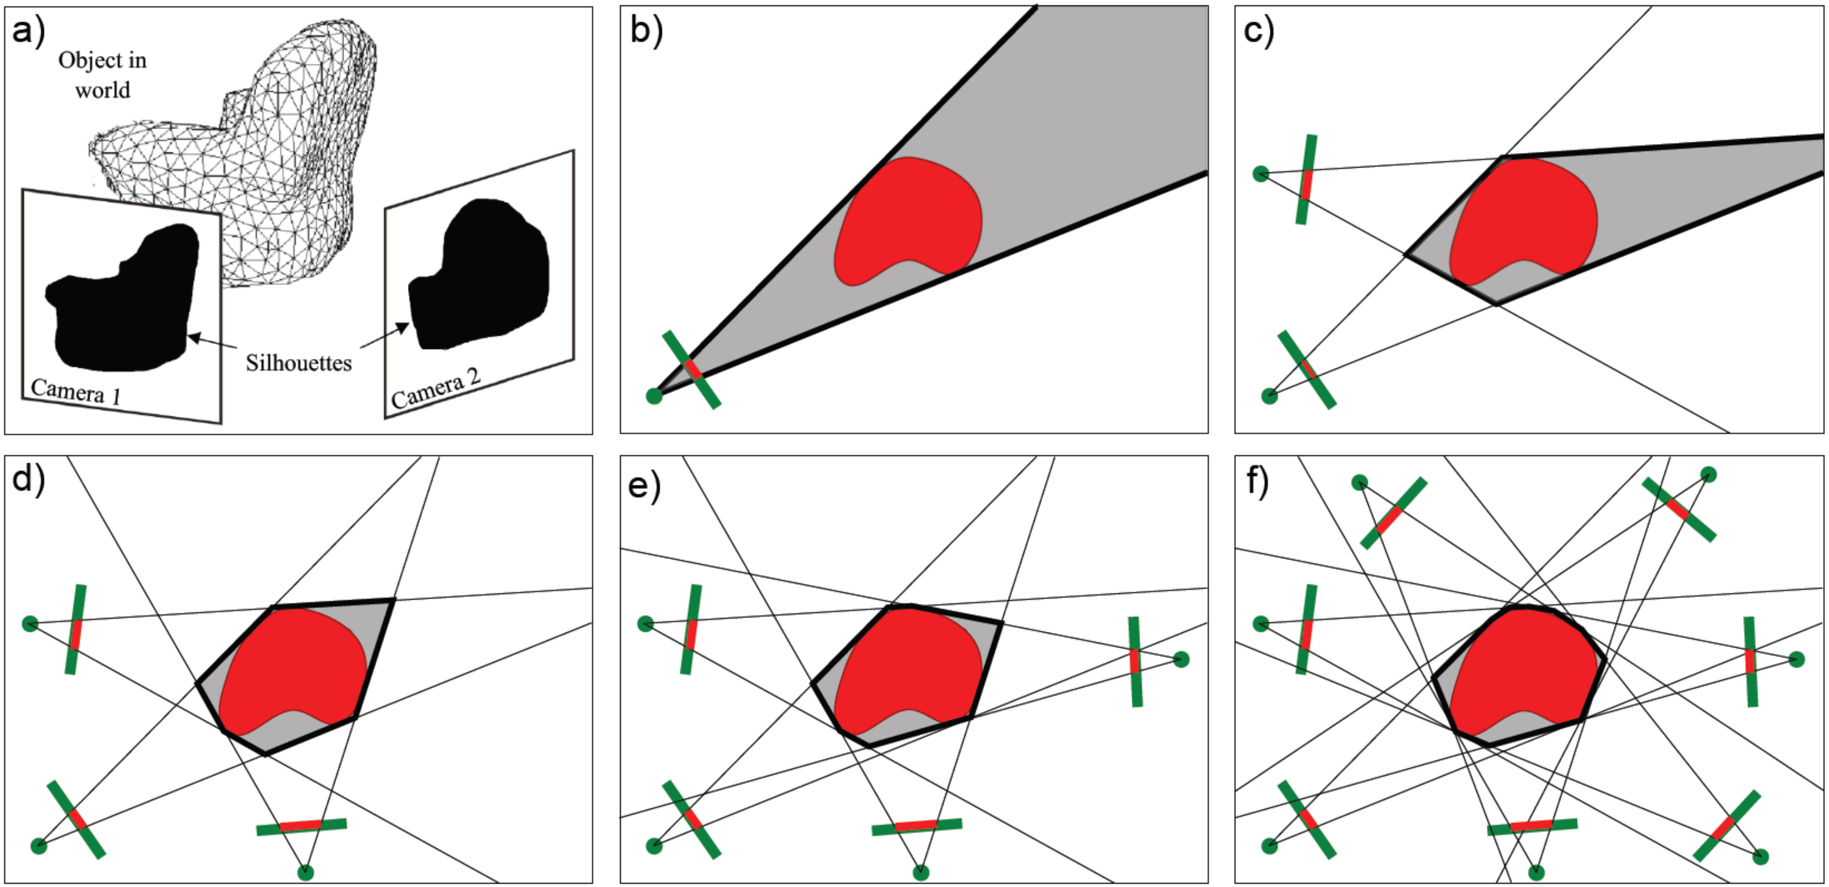

Shape from silhouette

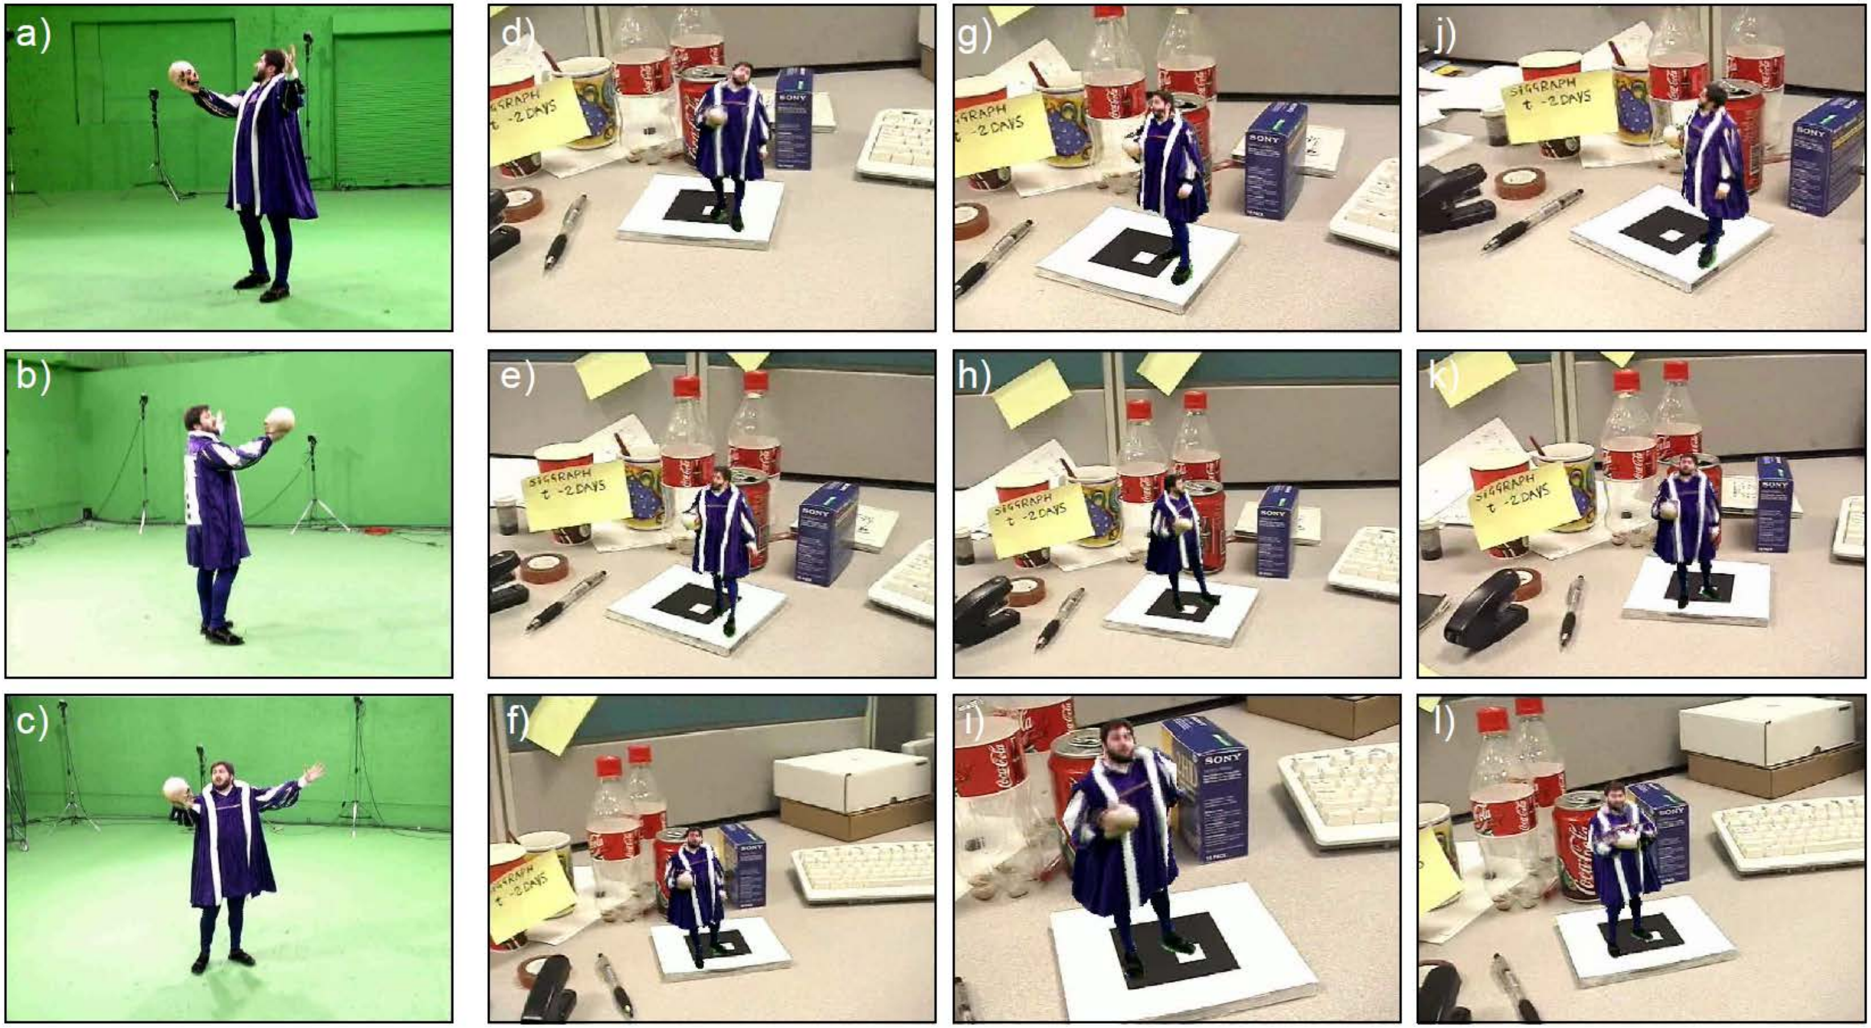

Shape from silhouette

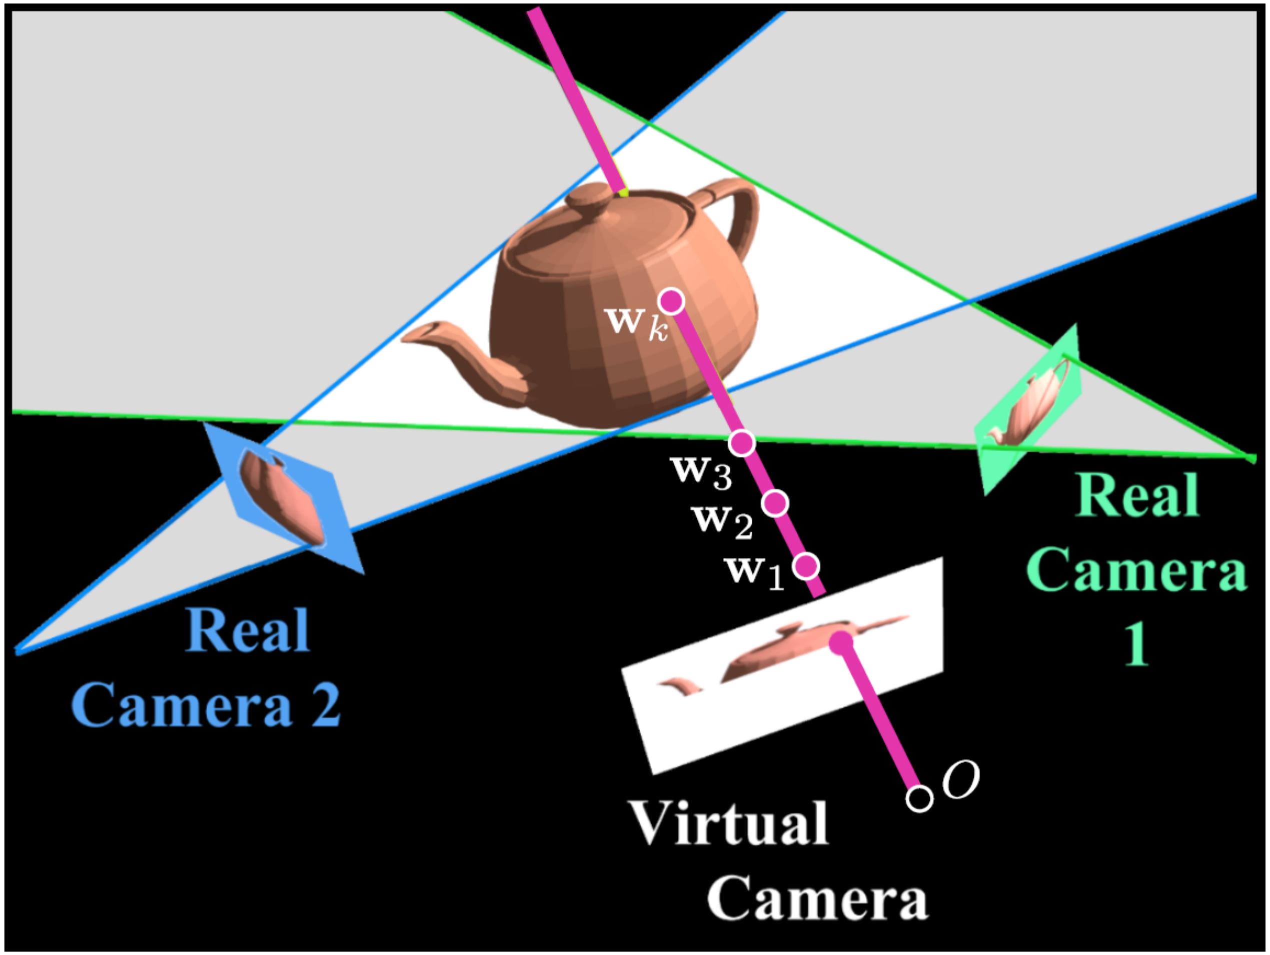

Conclusion

- Pinhole camera model is a non-linear function that takes points in 3D world and finds where they map to in image
- Parameterized by intrinsic and extrinsic matrices
- Difficult to estimate intrinsic/extrinsic/depth because non-linear
- Use homogeneous coordinates where we can get closed form solutions (initial sol’ns only)