




THE UNIVERSITY OF BRITISH COLUMBIA

CPSC 425: Computer Vision



Lecture 14: Planar Geometry and RANSAC

Menu for Today

Topics:

- **Planar** Geometry
- **Image Alignment**, Object Recognition
- **RANSAC**

Readings:


- **Today's** Lecture: Szeliski 2.1, 8.1, Forsyth & Ponce 10.4.2

Reminders:

- **Assignment 3**: due tonight
- **Assignment 4**: RANSAC and Panorama Stitching — **available tonight**
- **Photos** at the **end of the class** today

Today's “fun” Example: Im2Calories

ICCV 2015 paper by **Kevin Murphy**
(UBC’s former faculty)



Coincidentally Kevin is also author of one of the most prominent ML books

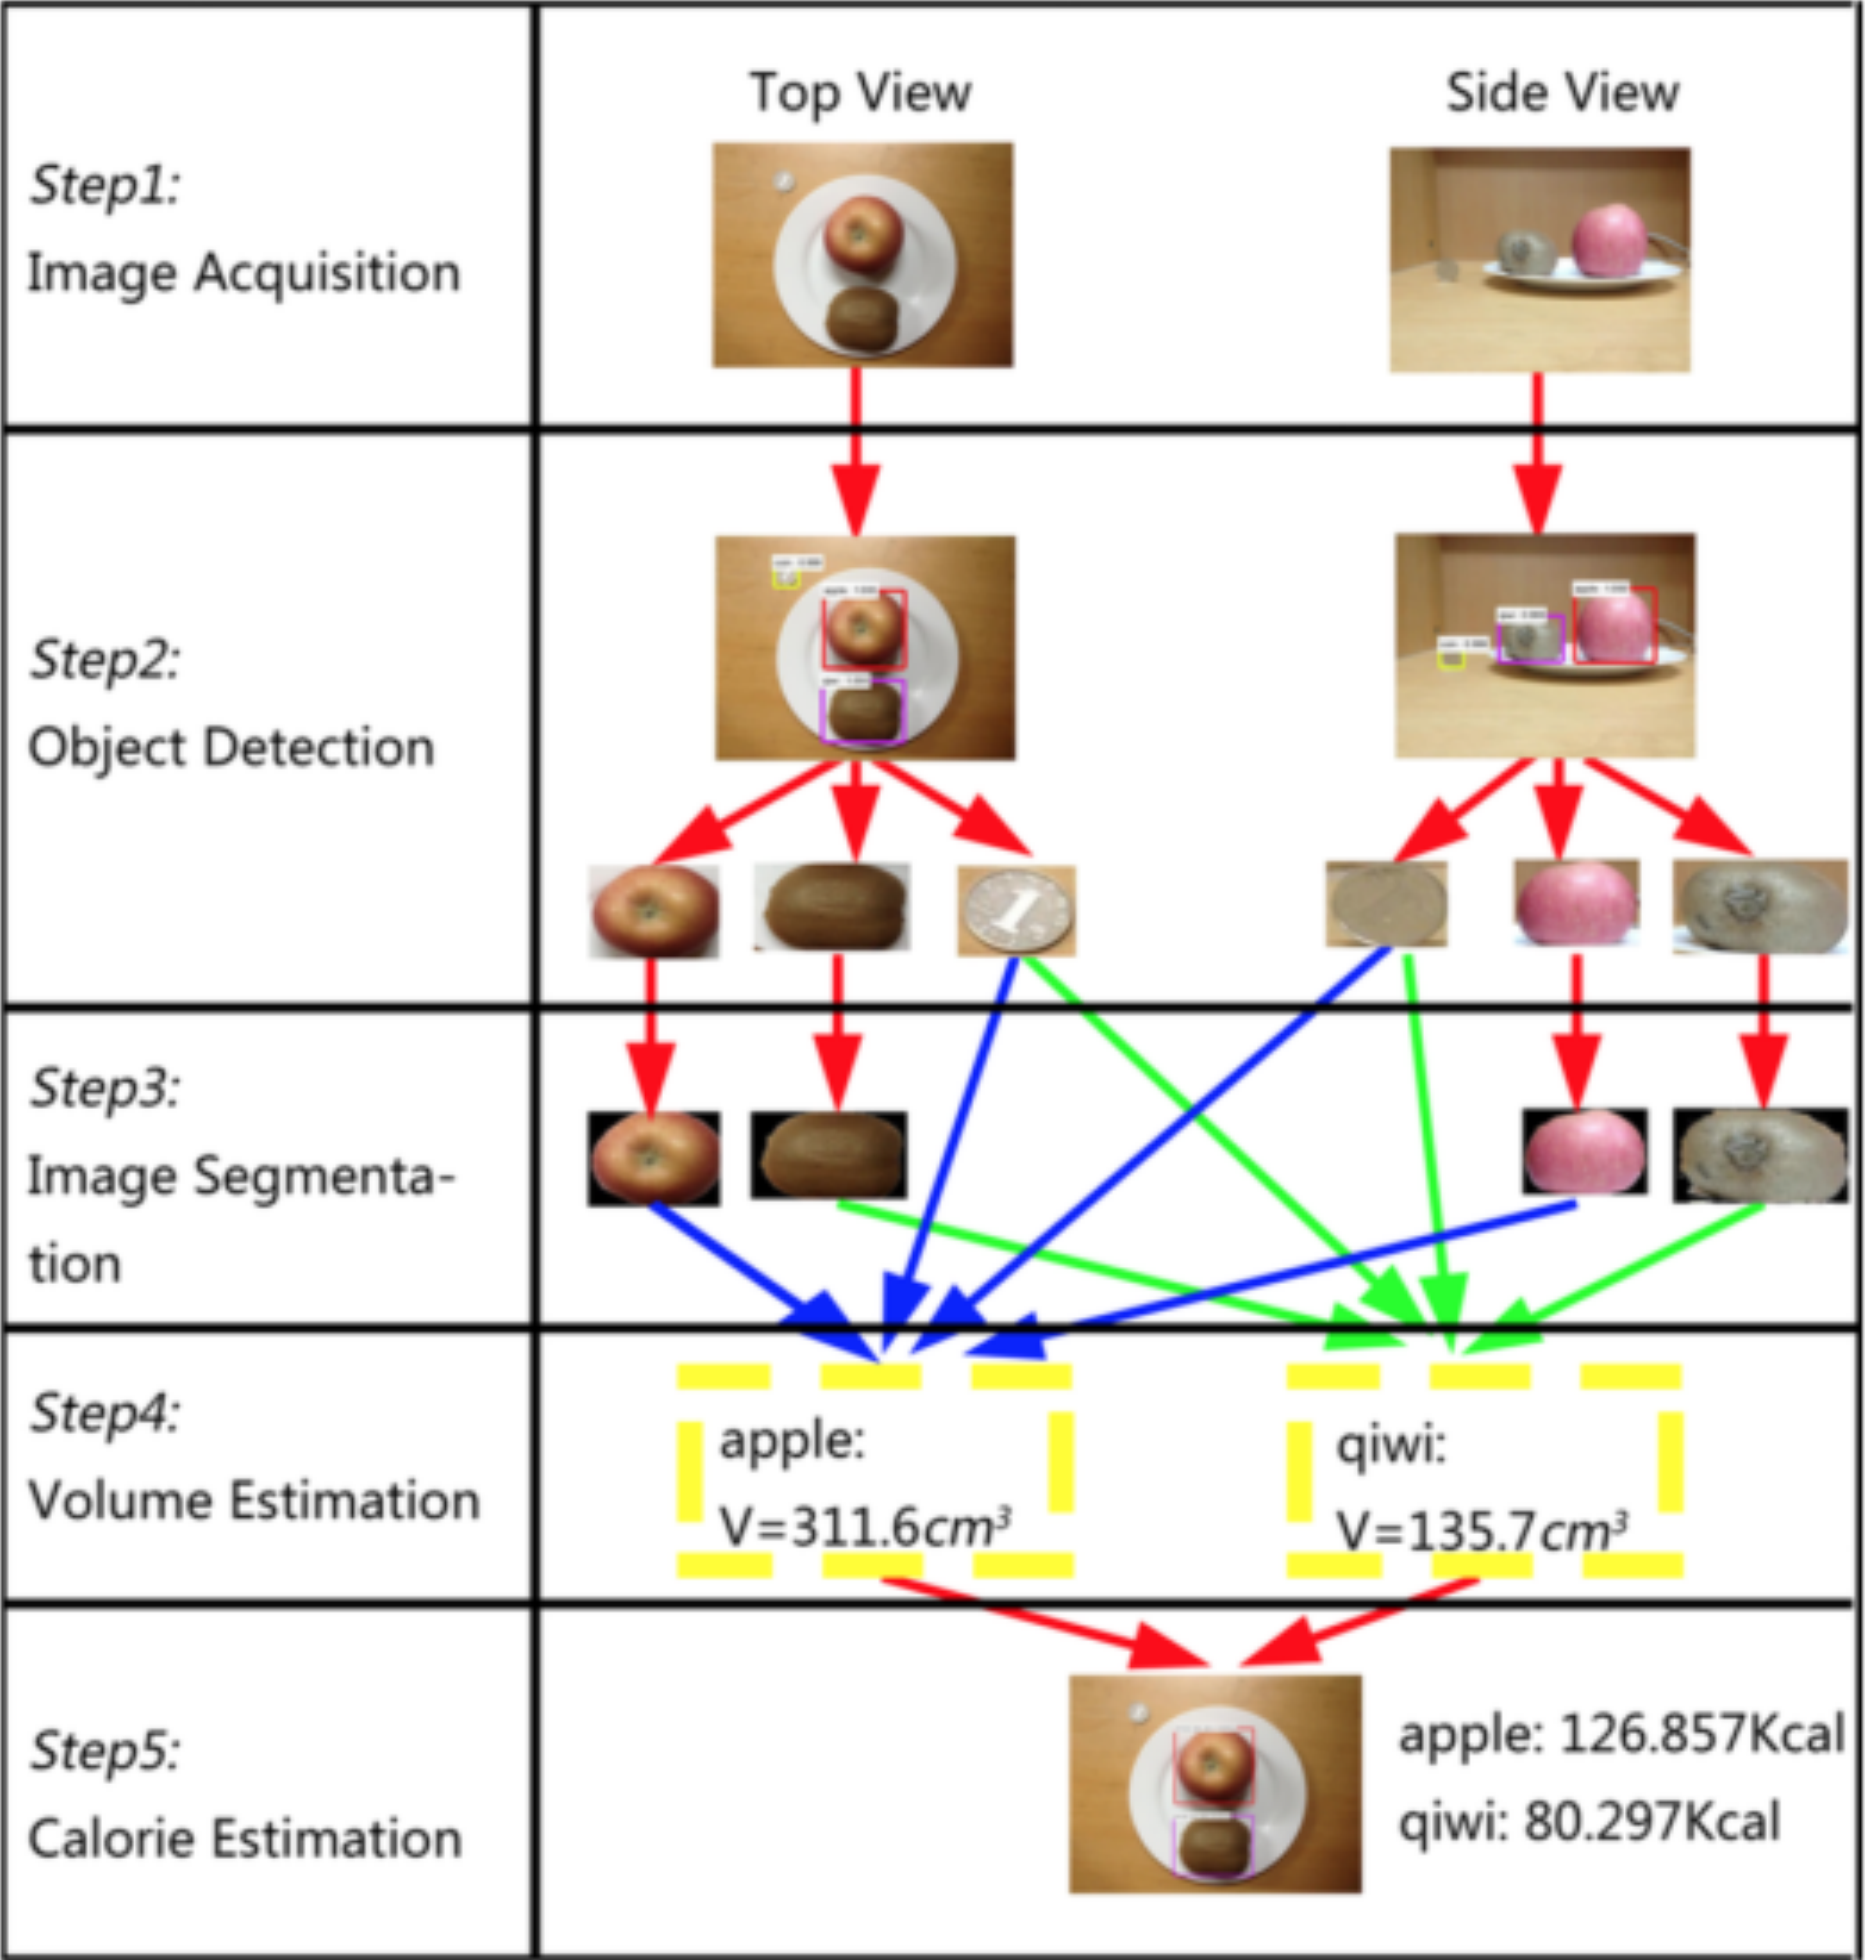


Figure 1: Calorie Estimation Flowchart

Today's “**fun**” Example: Im2Calories

Im2Calories: towards an automated mobile vision food diary

Austin Myers, Nick Johnston, Vivek Rathod, Anoop Korattikara, Alex Gorban Nathan Silberman, Sergio Guadarrama, George Papandreou, Jonathan Huang, Kevin Murphy amyers@umd.edu, (nickj, rathodv, kbanoop, gorban)@google.com (nsilberman, sguada, gpapan, jonathanhuang, kpmurphy)@google.com

Today's “**fun**” Example: Im2Calories

Im2Calories: towards an automated mobile vision food diary

Austin Myers, Nick Johnston, Vivek Rathod, Anoop Korattikara, Alex Gorban Nathan Silberman, Sergio Guadarrama, George Papandreou, Jonathan Huang, Kevin Murphy amyers@umd.edu, (nickj, rathodv, kbanoop, gorban)@google.com (nsilberman, sguada, gpapan, jonathanhuang, kpmurphy)@google.com

Today's “**fun**” Example: Im2Calories

Fun **on-line demo**: <http://www.caloriemama.ai/api>

Lecture 13: Re-Cap

Keypoint is an image location at which a descriptor is computed

- Locally distinct points
- Easily localizable and identifiable



Lecture 13: Re-Cap

Keypoint is an image location at which a descriptor is computed

- Locally distinct points
- Easily localizable and identifiable

The feature **descriptor** summarizes the local structure around the key point

- Allows us to (hopefully) unique matching of keypoints in presence of object pose variations, image and photometric deformations



Lecture 13: Re-Cap

Keypoint is an image location at which a descriptor is computed

- Locally distinct points
- Easily localizable and identifiable

The feature **descriptor** summarizes the local structure around the key point

- Allows us to (hopefully) unique matching of keypoints in presence of object pose variations, image and photometric deformations

Note, for repetitive structure this would still not give us unique matches.



Lecture 13: Re-Cap

- We motivated SIFT for identifying locally distinct keypoints in an image (**detection**)
- SIFT features (**description**) are invariant to translation, rotation, and scale; robust to 3D pose and illumination

1. Multi-scale extrema detection

2. Keypoint localization

3. Orientation assignment

4. Keypoint descriptor

Lecture 13: Re-Cap

Four steps to SIFT feature generation:

1. Scale-space representation and local extrema detection

- use DoG pyramid **Output:** (x, y, s) for each keypoint
- 3 scales/octave, down-sample by factor of 2 each octave

Lecture 13: Re-Cap

Four steps to SIFT feature generation:

1. Scale-space representation and local extrema detection

- use DoG pyramid **Output:** (x, y, s) for each keypoint
- 3 scales/octave, down-sample by factor of 2 each octave

2. Keypoint localization

- select stable keypoints (threshold on magnitude of extremum, ratio of principal curvatures) **Output:** Remove some (weak) keypoints

Lecture 13: Re-Cap

Four steps to SIFT feature generation:

1. Scale-space representation and local extrema detection

- use DoG pyramid **Output:** (x, y, s) for each keypoint
- 3 scales/octave, down-sample by factor of 2 each octave

2. Keypoint localization

- select stable keypoints (threshold on magnitude of extremum, ratio of principal curvatures) **Output:** Remove some (weak) keypoints

3. Keypoint orientation assignment

- based on histogram of local image gradient directions

Output: Orientation for each keypoint

Lecture 13: Re-Cap

Four steps to SIFT feature generation:

1. Scale-space representation and local extrema detection

- use DoG pyramid **Output:** (x, y, s) for each keypoint
- 3 scales/octave, down-sample by factor of 2 each octave

2. Keypoint localization

- select stable keypoints (threshold on magnitude of extremum, ratio of principal curvatures) **Output:** Remove some (weak) keypoints

3. Keypoint orientation assignment

- based on histogram of local image gradient directions

Output: Orientation for each keypoint

4. Keypoint descriptor

- histogram of local gradient directions — vector with $8 \times (4 \times 4) = 128$ dim
- vector normalized (to unit length)

Output: 128D normalized vector characterizing the keypoint region

Lecture 13: Histogram of Oriented Gradients (**HOG**)

Pedestrian detection


128 pixels
16 cells
15 blocks

1 cell step size



$$15 \times 7 \times 4 \times 9 = 3780$$

visualization




64 pixels
8 cells
7 blocks

Redundant representation due to overlapping blocks




Lecture 13: 'Speeded' Up Robust Features



Each cell is represented
by 4 values:

$$\left[\sum d_x, \sum d_y, \sum |d_x|, \sum |d_y| \right]$$

Haar wavelets filters
(Gaussian weighted from center)



How big is the SURF descriptor?
64 dimensions

Lecture 13: Summary

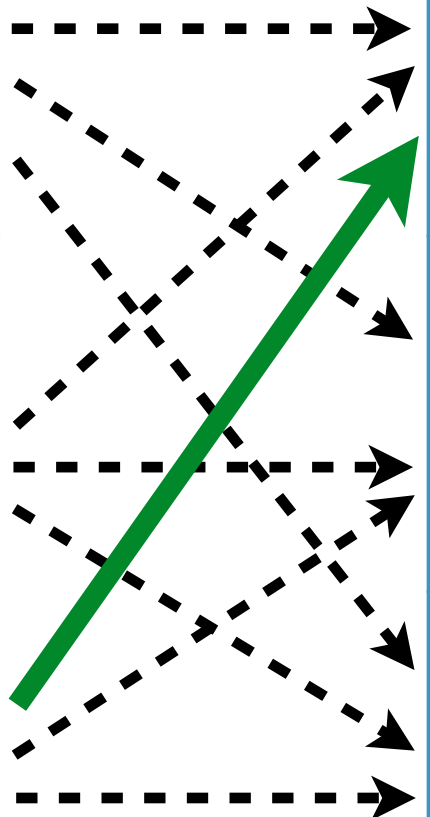
Keypoint Detection Algorithms	Representation		Keypoint Description Algorithms	Representation
Harris Corners	(x,y,s)		SIFT	128D
LoG / Blobs	(x,y,s)		Histogram of Oriented Gradients	3780D
SIFT	(x,y,s,theta)		SURF	64D

Image **Alignment**

Aim: Warp one image to align with another



Image **Alignment**

Aim: Warp one image to align with another using a 2D transformation



Image **Alignment**

Step 1: Find correspondences (matching points) across two images




Image **Alignment**

Step 2: Compute the transformation to align the two images



Image **Alignment**

Not all points will match across two images, we can also reject outliers




Image **Alignment**

Not all points will match across two images, we can also reject outliers



Planar Geometry

- 2D Linear + **Projective** transformations


Euclidean, Similarity, Affine, Homography

- Robust Estimation and **RANSAC**

Estimating 2D transforms with noisy correspondences

2D Transformations

- We will look at a family that can be represented by 3x3 matrices



- This group represents perspective projections of **planar surfaces**

Affine Transformation

- Transformed points are a **linear function** of the input points

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} + \begin{bmatrix} a_{13} \\ a_{23} \end{bmatrix}$$

Affine Transformation

- Transformed points are a **linear function** of the input points

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} + \begin{bmatrix} a_{13} \\ a_{23} \end{bmatrix}$$


- This can be written as a **single matrix multiplication** using **homogeneous** coordinates

$$\begin{bmatrix} x'_1 \\ y'_1 \\ 1 \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ y_1 \\ 1 \end{bmatrix}$$

Linear Transformation


— Consider the action of the unit square under, sample transform

$$\begin{bmatrix} 3 & 1 & 0 \\ 1 & 2 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$



Linear Transformation


— Consider the action of the unit square under, sample transform $\begin{bmatrix} 3 & 1 & 0 \\ 1 & 2 & 0 \\ 0 & 0 & 1 \end{bmatrix}$



$$\begin{bmatrix} 3 & 1 & 0 \\ 1 & 2 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} \begin{bmatrix} 3 \\ 1 \\ 1 \end{bmatrix} \begin{bmatrix} 4 \\ 3 \\ 1 \end{bmatrix}$$


Transformation Points Transformed Points

Linear (or Affine) Transformations




Translation, rotation, scale, shear (parallel lines preserved)

Linear (or Affine) Transformations



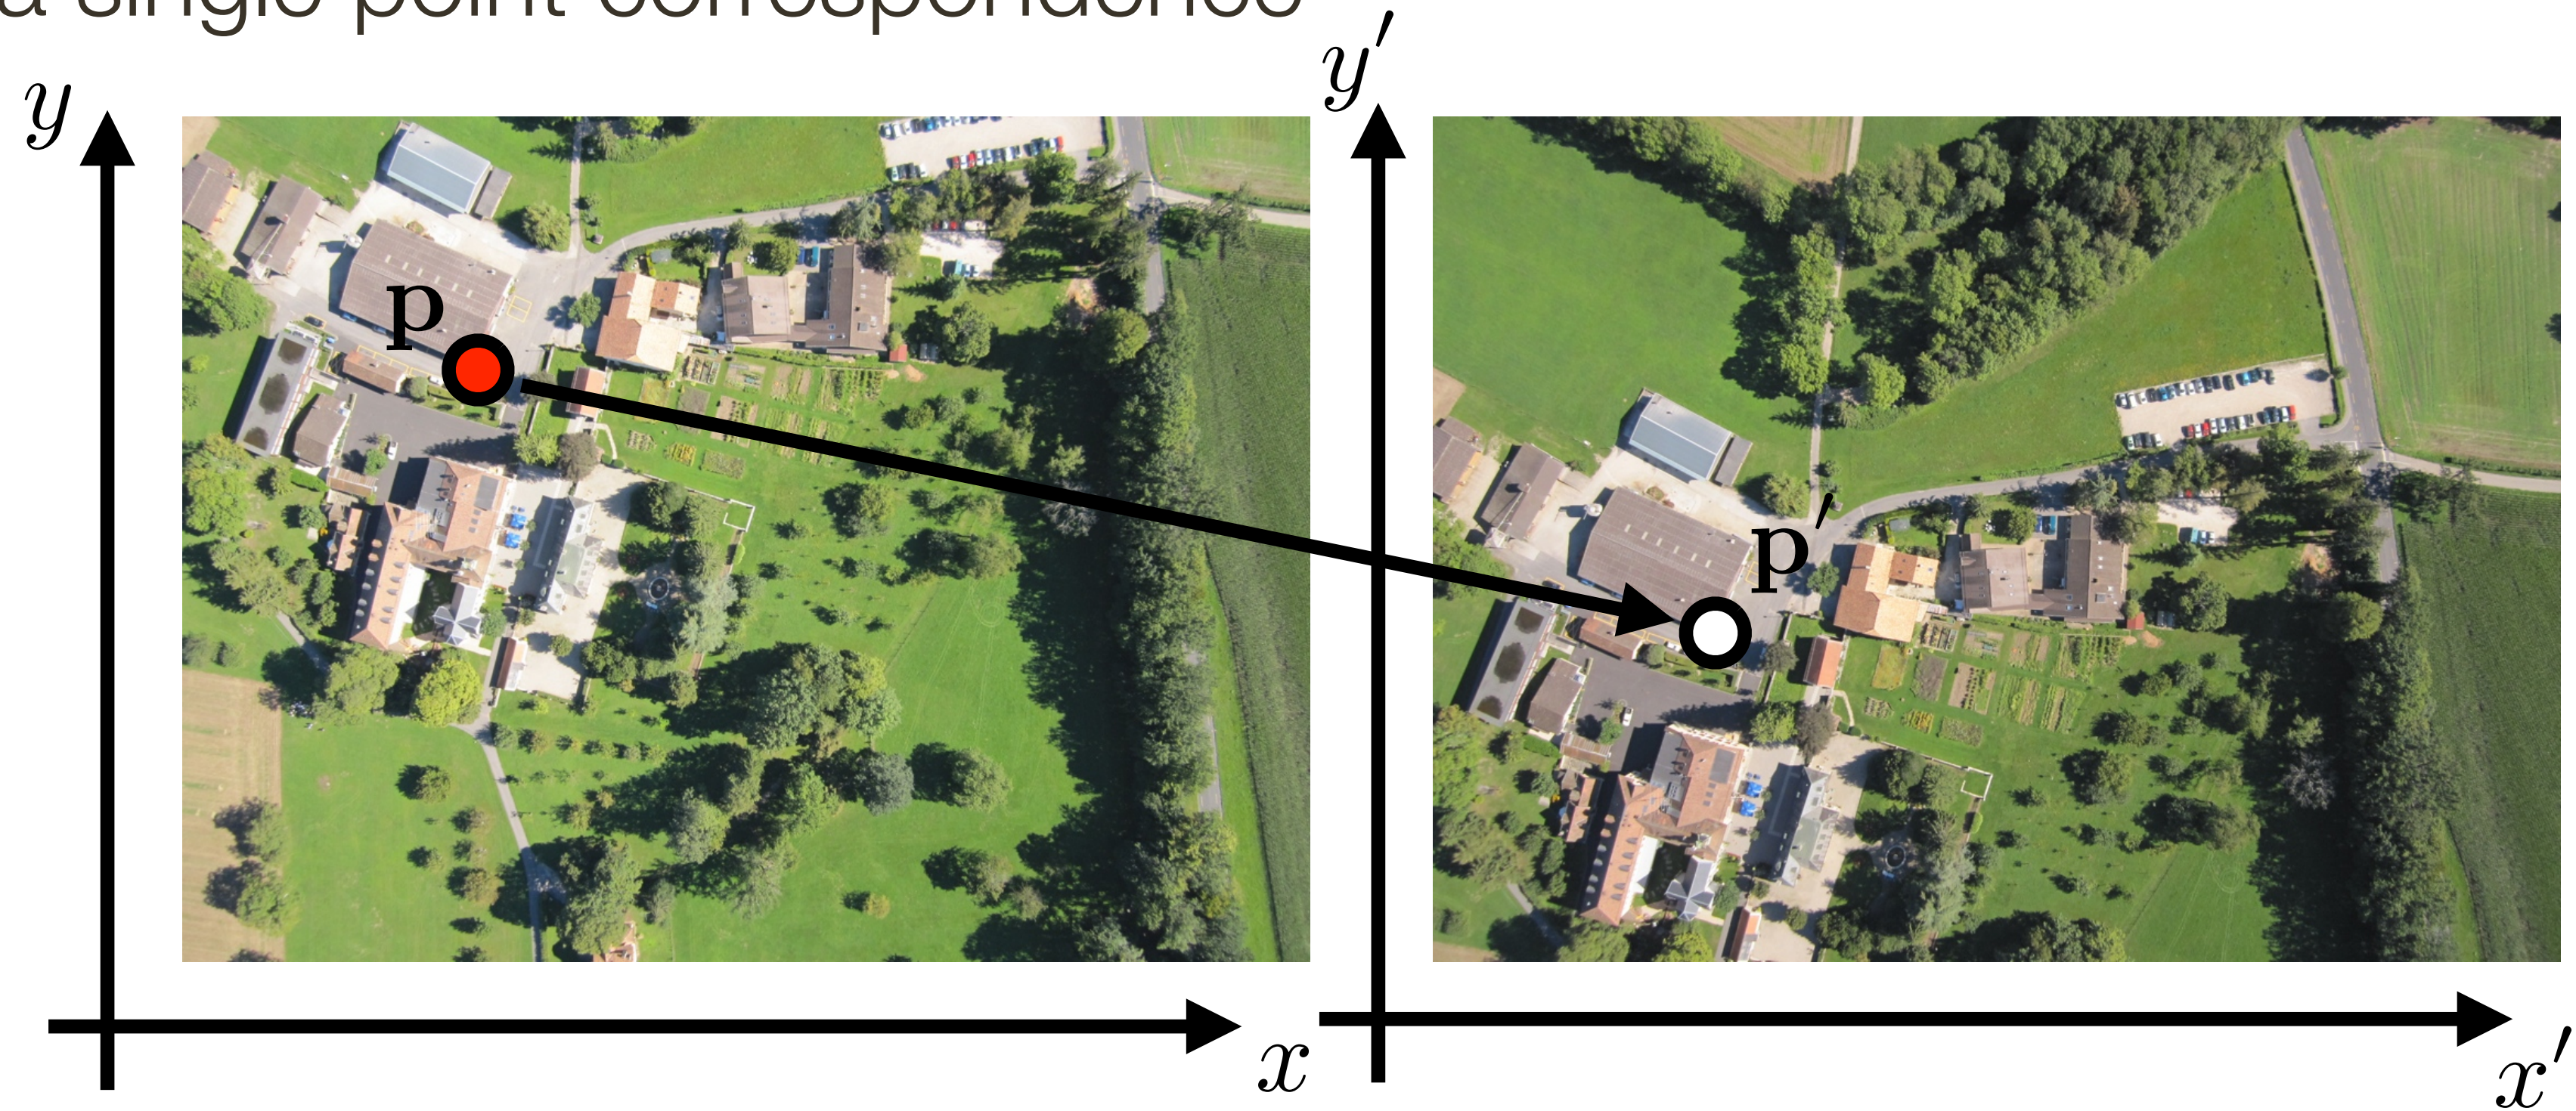
Translation, rotation, scale, shear (parallel lines preserved)



These transforms are not affine (parallel lines not preserved)

Linear (or Affine) Transformations

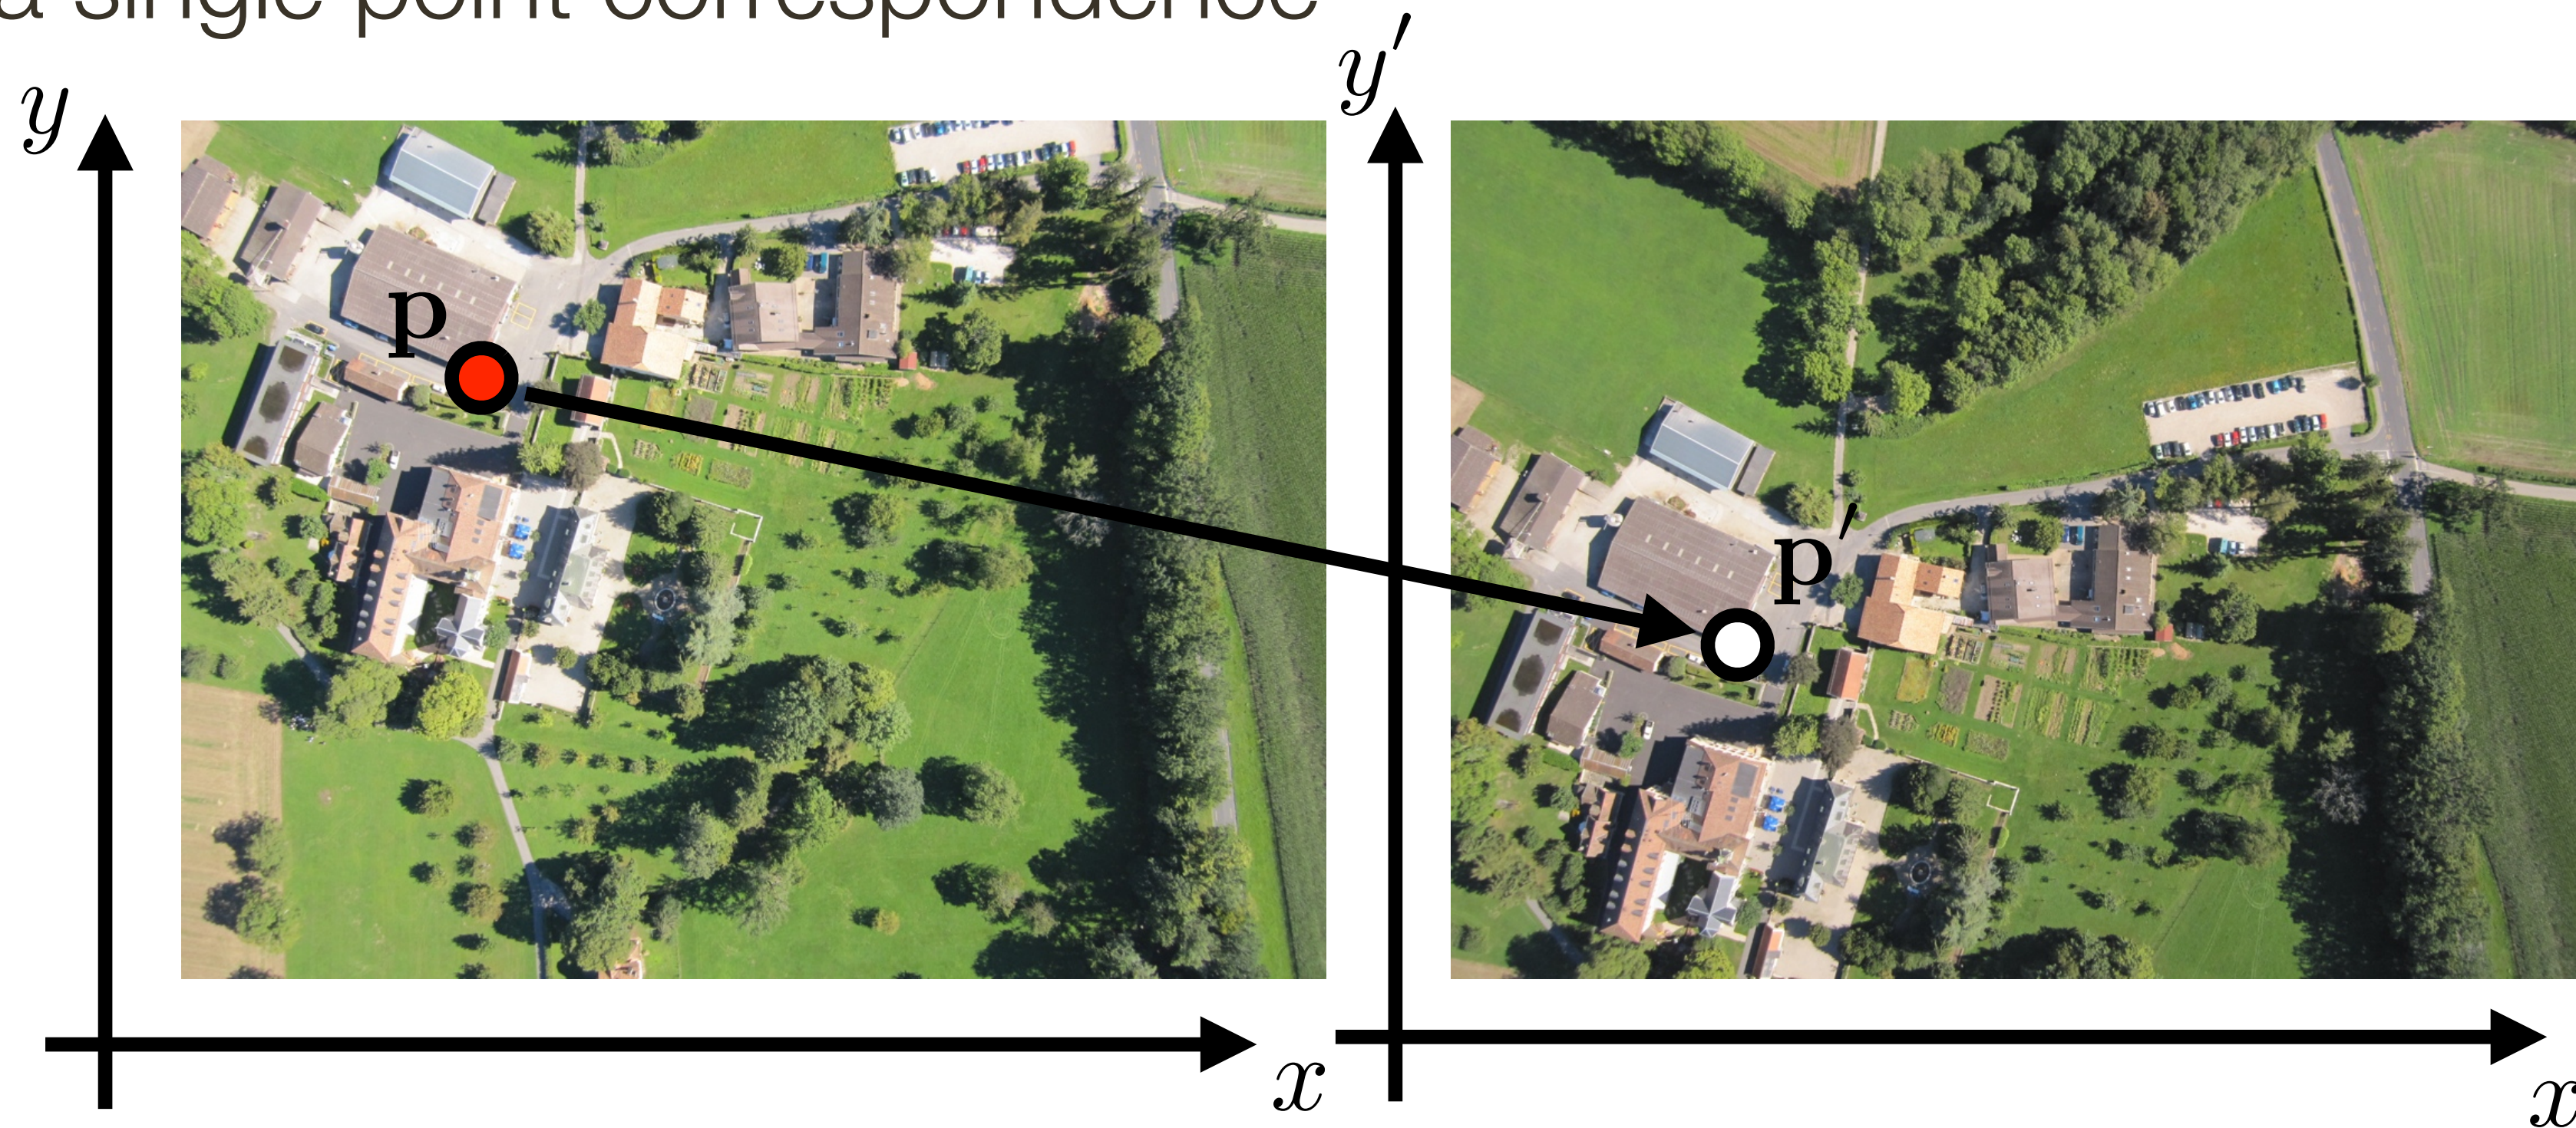
Consider a single point correspondence



$$\begin{bmatrix} x'_1 \\ y'_1 \\ 1 \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ y_1 \\ 1 \end{bmatrix}$$

Linear (or Affine) Transformations

Consider a single point correspondence



$$\begin{bmatrix} x'_1 \\ y'_1 \\ 1 \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ y_1 \\ 1 \end{bmatrix}$$

How many points are needed to solve for \mathbf{a} ?

Computing Affine Transform

Lets compute an affine transform from correspondences:

$$\begin{bmatrix} x'_1 \\ y'_1 \\ 1 \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ y_1 \\ 1 \end{bmatrix}$$

Re-arrange unknowns into a vector

$$\begin{bmatrix} x'_1 \\ y'_1 \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} & a_{21} & a_{22} & a_{23} \end{bmatrix} \begin{bmatrix} 0 & x_1 \\ 0 & y_1 \\ 0 & 1 \\ x_2 & 0 \\ y_2 & 0 \\ 1 & 0 \end{bmatrix}$$

Computing Affine Transform

Linear system in the unknown parameters **a**

$$\begin{bmatrix} x_1 & y_1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x_1 & y_1 & 1 \\ x_2 & y_2 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x_2 & y_2 & 1 \\ x_3 & y_3 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x_3 & y_3 & 1 \end{bmatrix} \begin{bmatrix} a_{11} \\ a_{12} \\ a_{13} \\ a_{21} \\ a_{22} \\ a_{23} \end{bmatrix} = \begin{bmatrix} x'_1 \\ y'_1 \\ x'_2 \\ y'_2 \\ x'_3 \\ y'_3 \end{bmatrix}$$

Of the form

$$\mathbf{M}\mathbf{a} = \mathbf{y}$$

Computing Affine Transform

Linear system in the unknown parameters **a**

$$\begin{bmatrix} x_1 & y_1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x_1 & y_1 & 1 \\ x_2 & y_2 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x_2 & y_2 & 1 \\ x_3 & y_3 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x_3 & y_3 & 1 \end{bmatrix} \begin{bmatrix} a_{11} \\ a_{12} \\ a_{13} \\ a_{21} \\ a_{22} \\ a_{23} \end{bmatrix} = \begin{bmatrix} x'_1 \\ y'_1 \\ x'_2 \\ y'_2 \\ x'_3 \\ y'_3 \end{bmatrix}$$

Of the form

$$\mathbf{M}\mathbf{a} = \mathbf{y}$$

Solve for **a** using Gaussian Elimination

Computing Affine Transform

Once we solve for a transform, we can now map any other points between the two images ... or resample one image in the coordinate system of the other



$$\begin{bmatrix} x'_1 \\ y'_1 \\ 1 \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ y_1 \\ 1 \end{bmatrix}$$

Computing Affine Transform

Once we solve for a transform, we can now map any other points between the two images ... or resample one image in the coordinate system of the other

This allows us to “stitch” the two images



Linear Transformations

Other linear transforms are special cases of **affine** transform:

$$\begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ 0 & 0 & 1 \end{bmatrix}$$

Linear Transformations

Other linear transforms are special cases of **affine** transform:

$$\begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ 0 & 0 & 1 \end{bmatrix}$$

e.g., $\begin{bmatrix} s \cos \theta & s \sin \theta & a_{13} \\ -s \sin \theta & s \cos \theta & a_{23} \\ 0 & 0 & 1 \end{bmatrix}$ **similarity** transform

Linear Transformations

Other linear transforms are special cases of **affine** transform:

$$\begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ 0 & 0 & 1 \end{bmatrix}$$

e.g., $\begin{bmatrix} \cos \theta & \sin \theta & a_{13} \\ -\sin \theta & \cos \theta & a_{23} \\ 0 & 0 & 1 \end{bmatrix}$ **euclidian** transform

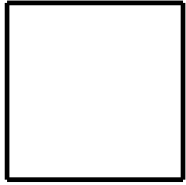
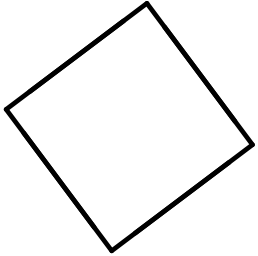
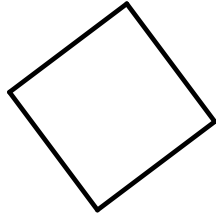
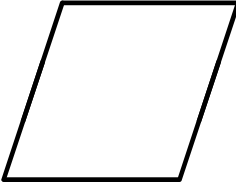
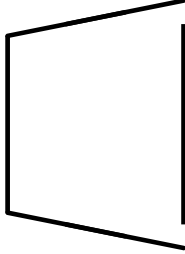
Linear Transformations

Other linear transforms are special cases of **affine** transform:

$$\begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ 0 & 0 & 1 \end{bmatrix}$$


e.g., $\begin{bmatrix} 1 & 0 & a_{13} \\ 0 & 1 & a_{23} \\ 0 & 0 & 1 \end{bmatrix}$ **translation** transform

2D Transformations


Transformation	Matrix	# DoF	Preserves	Icon
translation	$\left[\begin{array}{c c} \mathbf{I} & \mathbf{t} \end{array} \right]_{2 \times 3}$	2	orientation	
rigid (Euclidean)	$\left[\begin{array}{c c} \mathbf{R} & \mathbf{t} \end{array} \right]_{2 \times 3}$	3	lengths	
similarity	$\left[\begin{array}{c c} s\mathbf{R} & \mathbf{t} \end{array} \right]_{2 \times 3}$	4	angles	
affine	$\left[\begin{array}{c} \mathbf{A} \end{array} \right]_{2 \times 3}$	6	parallelism	
projective	$\left[\begin{array}{c} \tilde{\mathbf{H}} \end{array} \right]_{3 \times 3}$	8	straight lines	

Example: Warping with Different Transformations

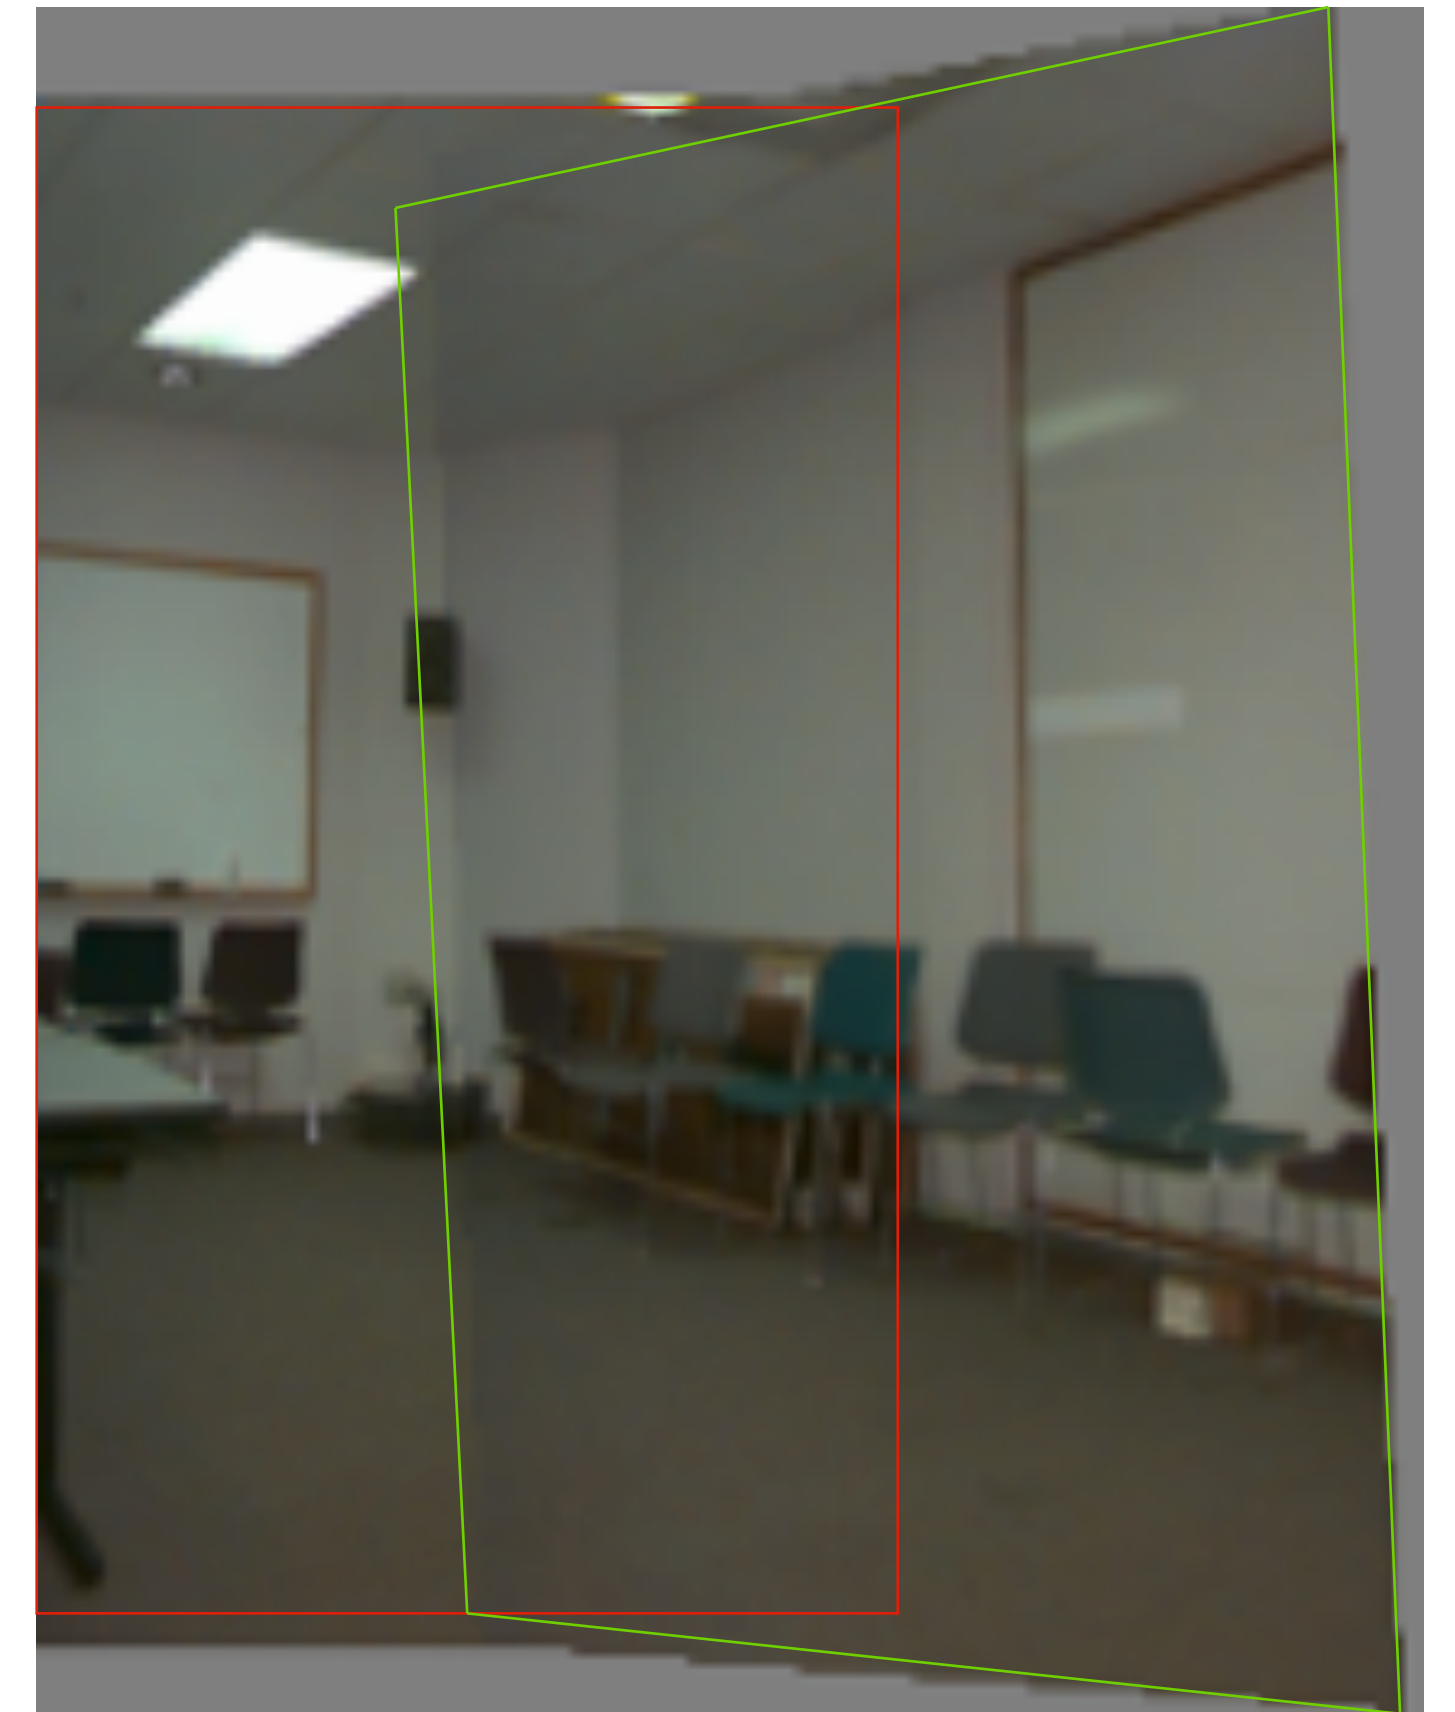
Translation



Affine



Projective
(homography)



Aside: We can use homographies when ...

1.... the scene is planar; or



2.... the scene is very far
or has small (relative)
depth variation → scene
is approximately planar



Aside: We can use homographies when ...

3.... the scene is captured under camera rotation only (no translation or pose change)



Projective Transformation

General 3x3 matrix transformation

$$\begin{bmatrix} x'_1 \\ y'_1 \\ 1 \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix} \begin{bmatrix} x_1 \\ y_1 \\ 1 \end{bmatrix}$$

Projective Transformation

General 3x3 matrix transformation

$$\begin{bmatrix} x'_1 \\ y'_1 \\ 1 \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix} \begin{bmatrix} x_1 \\ y_1 \\ 1 \end{bmatrix}$$

Lets try an example:

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \mathbf{H} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 1 \end{bmatrix} \begin{bmatrix} 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix} = \begin{bmatrix} 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 \\ 1 & 2 & 1 & 2 \end{bmatrix}$$

Transformation Points Transformed Points

Projective Transformation

General 3x3 matrix transformation

$$\begin{bmatrix} x'_1 \\ y'_1 \\ 1 \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix} \begin{bmatrix} x_1 \\ y_1 \\ 1 \end{bmatrix}$$

Lets try an example:

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \mathbf{H} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 1 \end{bmatrix} \begin{bmatrix} 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix} = \begin{bmatrix} 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 \\ 1 & 2 & 1 & 2 \end{bmatrix}$$

Transformation

Points

Transformed Points


Divide by the last row:

$$\begin{bmatrix} 0 & 0 & 1 & 0.5 \\ 0 & 0.5 & 0 & 0.5 \\ 1 & 1 & 1 & 1 \end{bmatrix}$$

Compute **H** from Correspondences

Each match gives 2 equations to solve for **8** parameters

$$\begin{bmatrix} x'_1 \\ y'_1 \\ 1 \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix} \begin{bmatrix} x_1 \\ y_1 \\ 1 \end{bmatrix}$$




→ 4 correspondences to solve for **H** matrix

Solution uses **Singular Value Decomposition** (SVD)

In **Assignment 4** you can compute this using `cv2.findHomography`

Image **Alignment**

Find **corresponding** (matching) points between the image



$$\mathbf{u} = \mathbf{H}\mathbf{x}$$

2 points for Similarity

3 for Affine

4 for Homography

Image **Alignment**

In practice we have many noisy correspondences + **outliers**




Image **Alignment**

In practice we have many noisy correspondences + **outliers**

e.g., for an affine transform we have a linear system in the parameters **a**:

$$\begin{bmatrix} x_1 & y_1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x_1 & y_1 & 1 \\ x_2 & y_2 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x_2 & y_2 & 1 \\ x_3 & y_3 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x_3 & y_3 & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \end{bmatrix} \begin{bmatrix} a_{11} \\ a_{12} \\ a_{13} \\ a_{21} \\ a_{22} \\ a_{23} \\ \vdots \end{bmatrix} = \begin{bmatrix} x'_1 \\ y'_1 \\ x'_2 \\ y'_2 \\ x'_3 \\ y'_3 \\ \vdots \end{bmatrix}$$

It is **overconstrained** (more equations than unknowns) and subject to **outliers** (some rows are completely wrong)

Image **Alignment**

In practice we have many noisy correspondences + **outliers**

e.g., for an affine transform we have a linear system in the parameters **a**:

$$\begin{bmatrix} x_1 & y_1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x_1 & y_1 & 1 \\ x_2 & y_2 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x_2 & y_2 & 1 \\ x_3 & y_3 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x_3 & y_3 & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \end{bmatrix} \begin{bmatrix} a_{11} \\ a_{12} \\ a_{13} \\ a_{21} \\ a_{22} \\ a_{23} \\ \vdots \end{bmatrix} = \begin{bmatrix} x'_1 \\ y'_1 \\ x'_2 \\ y'_2 \\ x'_3 \\ y'_3 \\ \vdots \end{bmatrix}$$

It is **overconstrained** (more equations than unknowns) and subject to **outliers** (some rows are completely wrong)

Let's deal with these problems in a simpler context ...

Fitting a Model to Noisy Data

Suppose we are **fitting a line** to a dataset that consists of 50% outliers

We can fit a line using two points

If we draw pairs of points uniformly at random, what fraction of pairs will consist entirely of 'good' data points (inliers)?

Fitting a Model to Noisy Data

Suppose we are **fitting a line** to a dataset that consists of 50% outliers

We can fit a line using two points

- If we draw pairs of points uniformly at random, then about $1/4$ of these pairs will consist entirely of ‘good’ data points (inliers)
- We can identify these good pairs by noticing that a large collection of other points lie close to the line fitted to the pair
- A better estimate of the line can be obtained by refitting the line to the points that lie close to the line

RANSAC (RANdom **S**Amples **C**onsensus)

1. Randomly choose minimal subset of data points necessary to fit model (a **sample**)
2. Points within some distance threshold, t , of model are a **consensus set**.
Size of consensus set is model's **support**
3. Repeat for N samples; model with biggest support is most robust fit
 - Points within distance t of best model are inliers
 - Fit final model to all inliers

RANSAC (RANdom **S**Amples **C**onsensus)

1. Randomly choose minimal subset of data points necessary to fit model (a **sample**)
2. Points within some distance threshold, t , of model are a **consensus set**.
Size of consensus set is model's **support**
3. Repeat for N samples; model with biggest support is most robust fit
 - Points within distance t of best model are inliers
 - Fit final model to all inliers

RANSAC is very useful for variety of applications

RANSAC (RANdom SAmples Consensus)

1. Randomly choose minimal subset of data points necessary to fit model (a **sample**)
Fitting a Line: 2 points
2. Points within some distance threshold, t , of model are a **consensus set**.
Size of consensus set is model's **support**
3. Repeat for N samples; model with biggest support is most robust fit
 - Points within distance t of best model are inliers
 - Fit final model to all inliers

Example 1: Fitting a Line




Figure Credit: Hartley & Zisserman

Example 1: Fitting a Line




Figure Credit: Hartley & Zisserman

Example 1: Fitting a Line




Figure Credit: Hartley & Zisserman

Algorithm 10.4

This was Algorithm 15.4 in Forsyth & Ponce (1st ed.)

Algorithm 15.4: RANSAC: fitting lines using random sample consensus

Determine:

n — the smallest number of points required

k — the number of iterations required

t — the threshold used to identify a point that fits well

d — the number of nearby points required
to assert a model fits well

Until k iterations have occurred

Draw a sample of n points from the data
uniformly and at random

Fit to that set of n points

For each data point outside the sample

Test the distance from the point to the line
against t ; if the distance from the point to the line
is less than t , the point is close

end

If there are d or more points close to the line
then there is a good fit. Refit the line using all
these points.

end

Use the best fit from this collection, using the
fitting error as a criterion

RANSAC: Fitting Lines Using Random Sample Consensus

RANSAC: How many samples?

Let ω be the fraction of inliers (i.e., points on line)

Let n be the number of points needed to define hypothesis
($n = 2$ for a line in the plane)

Suppose k samples are chosen

The probability that a single sample of n points is correct (all inliers) is

RANSAC: How many samples?

Let ω be the fraction of inliers (i.e., points on line)

Let n be the number of points needed to define hypothesis
($n = 2$ for a line in the plane)

Suppose k samples are chosen

The probability that a single sample of n points is correct (all inliers) is

$$\omega^n$$

The probability that all k samples fail is

RANSAC: How many samples?

Let ω be the fraction of inliers (i.e., points on line)

Let n be the number of points needed to define hypothesis
($n = 2$ for a line in the plane)

Suppose k samples are chosen

The probability that a single sample of n points is correct (all inliers) is

$$\omega^n$$

The probability that all k samples fail is

$$(1 - \omega^n)^k$$

Choose k large enough (to keep this below a target failure rate)

RANSAC: k Samples Chosen ($p = 0.99$)

Sample size	Proportion of outliers						
n	5%	10%	20%	25%	30%	40%	50%
2	2	3	5	6	7	11	17
3	3	4	7	9	11	19	35
4	3	5	9	13	17	34	72
5	4	6	12	17	26	57	146
6	4	7	16	24	37	97	293
7	4	8	20	33	54	163	588
8	5	9	26	44	78	272	1177

Figure Credit: Hartley & Zisserman


After RANSAC

RANSAC divides data into inliers and outliers and yields estimate computed from minimal set of inliers

Improve this initial estimate with estimation over all inliers (e.g., with standard least-squares minimization)

But this may change inliers, so alternate fitting with re-classification as inlier/outlier

Example 2: Fitting a Line



Example 2: Fitting a Line




Figure Credit: Hartley & Zisserman

Image **Alignment** + **RANSAC**

In practice we have many noisy correspondences + **outliers**




Image **Alignment** + RANSAC

RANSAC solution for Similarity Transform (2 points)



Image **Alignment** + RANSAC

RANSAC solution for Similarity Transform (2 points)



Image **Alignment** + **RANSAC**

RANSAC solution for Similarity Transform (2 points)



4 inliers (**red**, **yellow**, **orange**, **brown**),

Image **Alignment** + RANSAC

RANSAC solution for Similarity Transform (2 points)



4 outliers (**blue**, **light blue**, **purple**, **pink**)

Image **Alignment** + **RANSAC**

RANSAC solution for Similarity Transform (2 points)



4 inliers (**red**, **yellow**, **orange**, **brown**),
4 outliers (**blue**, **light blue**, **purple**, **pink**)

Image **Alignment** + RANSAC



RANSAC solution for Similarity Transform (2 points)



choose light blue, purple

Image **Alignment** + RANSAC


RANSAC solution for Similarity Transform (2 points)



warp image

Image **Alignment** + RANSAC


RANSAC solution for Similarity Transform (2 points)



check match distances

Image **Alignment** + RANSAC

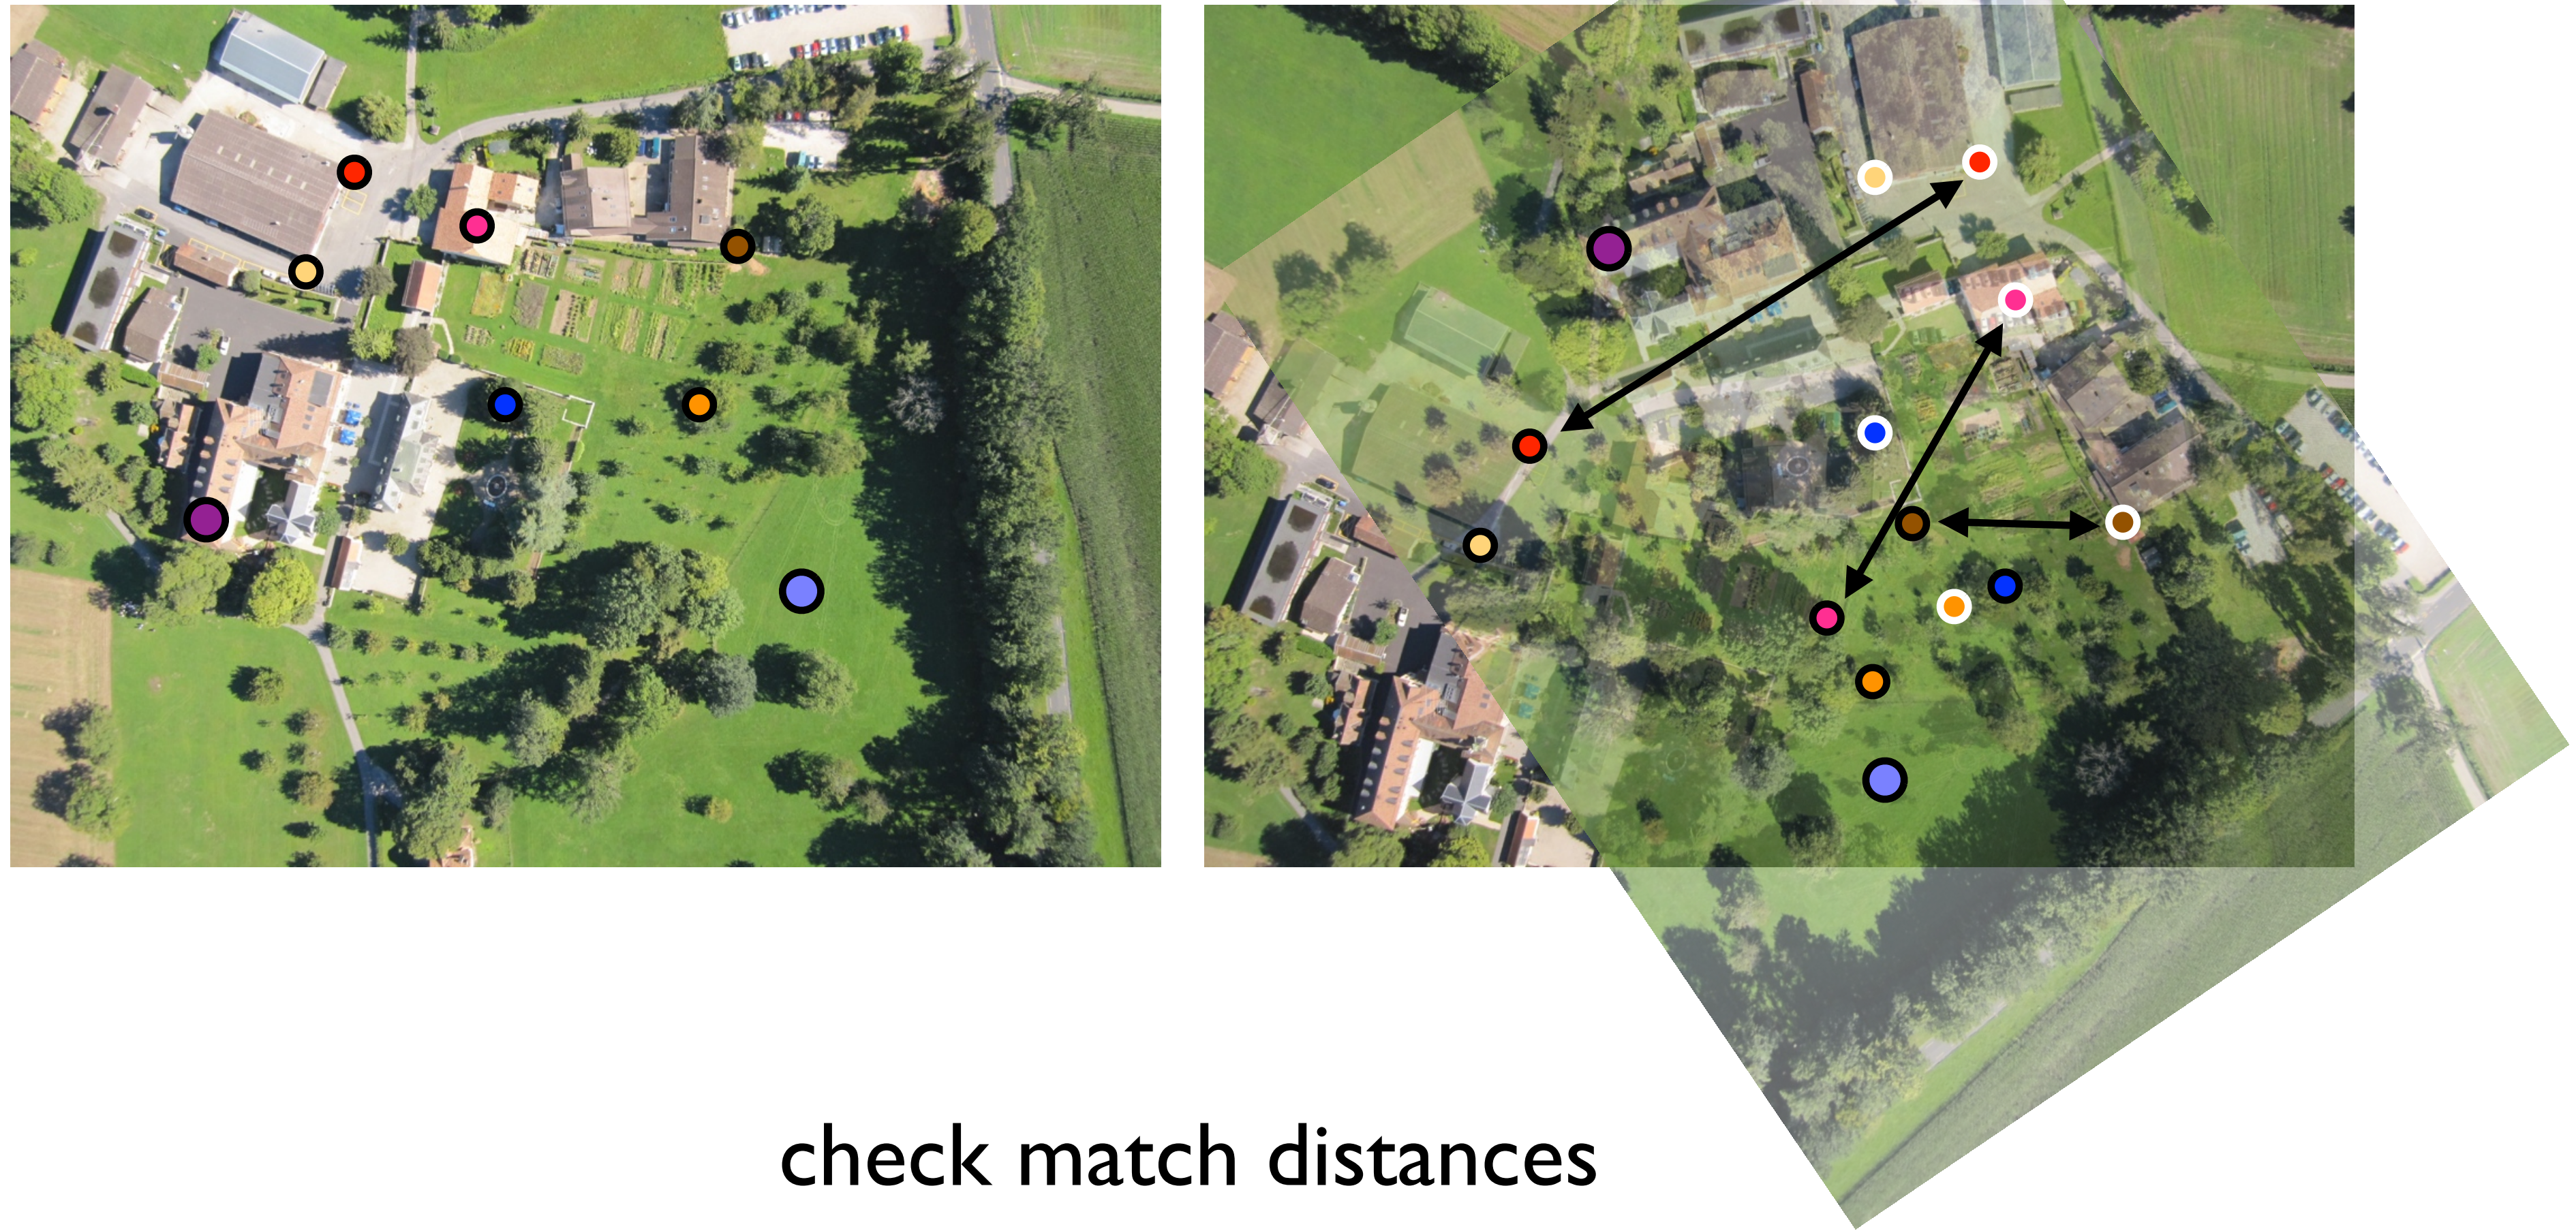
RANSAC solution for Similarity Transform (2 points)



check match distances

Image **Alignment** + RANSAC

RANSAC solution for Similarity Transform (2 points)



check match distances

#inliers = 2

Image **Alignment** + RANSAC

RANSAC solution for Similarity Transform (2 points)



Image **Alignment** + RANSAC

RANSAC solution for Similarity Transform (2 points)



choose **pink**, **blue**

Image **Alignment** + RANSAC

RANSAC solution for Similarity Transform (2 points)



warp image

Image **Alignment** + RANSAC

RANSAC solution for Similarity Transform (2 points)



check match distances

Image **Alignment** + RANSAC

RANSAC solution for Similarity Transform (2 points)



check match distances

Image **Alignment** + RANSAC

RANSAC solution for Similarity Transform (2 points)



check match distances

#inliers = 2

Image **Alignment** + RANSAC

RANSAC solution for Similarity Transform (2 points)



Image **Alignment** + RANSAC

RANSAC solution for Similarity Transform (2 points)



choose **red**, **orange**

Image **Alignment** + RANSAC

RANSAC solution for Similarity Transform (2 points)



warp image

Image **Alignment** + RANSAC

RANSAC solution for Similarity Transform (2 points)



check match distances

Image **Alignment** + RANSAC

RANSAC solution for Similarity Transform (2 points)



Image **Alignment** + RANSAC

RANSAC solution for Similarity Transform (2 points)



check match distances

#inliers = 4

Image **Alignment** + RANSAC

RANSAC solution for Similarity Transform (2 points)



Image **Alignment + RANSAC**

Assignment 4

- 1.** Match feature points between 2 views
- 2.** Select minimal subset of matches*
- 3.** Compute transformation T using minimal subset
- 4.** Check consistency of all points with T — compute projected position and count #inliers with distance $<$ threshold
- 5.** Repeat steps 2-4 to maximize #inliers

* Similarity transform = 2 points, Affine = 3, Homography = 4

RANSAC: k Samples Chosen ($p = 0.99$)

Sample size	Proportion of outliers						
n	5%	10%	20%	25%	30%	40%	50%
2	2	3	5	6	7	11	17
3	3	4	7	9	11	19	35
4	3	5	9	13	17	34	72
5	4	6	12	17	26	57	146
6	4	7	16	24	37	97	293
7	4	8	20	33	54	163	588
8	5	9	26	44	78	272	1177

Figure Credit: Hartley & Zisserman


RANSAC: k Samples Chosen ($p = 0.99$)

Sample size	Proportion of outliers						
n	5%	10%	20%	25%	30%	40%	50%
2	2	3	5	6	7	11	17
3	3	4	7	9	11	19	35
4	3	5	9	13	17	34	72
5	4	6	12	17	26	57	146
6	4	7	16	24	37	97	293
7	4	8	20	33	54	163	588
8	5	9	26	44	78	272	1177

Figure Credit: Hartley & Zisserman


2-view **Rotation** Estimation

Find features + raw matches, use RANSAC to find Similarity




2-view **Rotation** Estimation

Remove outliers, can now solve for R using least squares



2-view **Rotation** Estimation

Final rotation estimation



Object Instance Recognition

Database of planar objects



Instance recognition



Object **Instance Recognition** with SIFT

Match SIFT descriptors between **query image** and a database of known keypoints extracted from **training examples**

- use fast (approximate) nearest neighbour matching
- threshold based on ratio of distances between 1NN and 2NN

Use **RANSAC** to find a **subset of matches** that all agree on an object and geometric transform (e.g., **affine transform**)

Optionally **refine pose estimate** by recomputing the transformation using all the RANSAC inliers

Re-cap RANSAC

RANSAC is a technique to fit data to a model

- divide data into inliers and outliers
- estimate model from minimal set of inliers
- improve model estimate using all inliers
- alternate fitting with re-classification as inlier/outlier

RANSAC is a general method suited for a wide range of model fitting problems

- easy to implement
- easy to estimate/control failure rate

RANSAC only handles a moderate percentage of outliers without cost blowing up