UNIVERSITY OF CAMBRIDGE

Computer vision for human-machine interaction: False starts and new horizons

Roberto Cipolla

Department of Engineering

Research team

http://www.eng.cam.ac.uk/~cipolla/people.html

UNIVERSITY OF CAMBRIDGE

False starts? (1992)

If the only tool you have is a hammer you tend to see every problem as a nail

UNIVERSITY OF CAMBRIDGE

Model-based tracking

3D hand model

UNIVERSITY OF CAMBRIDGE

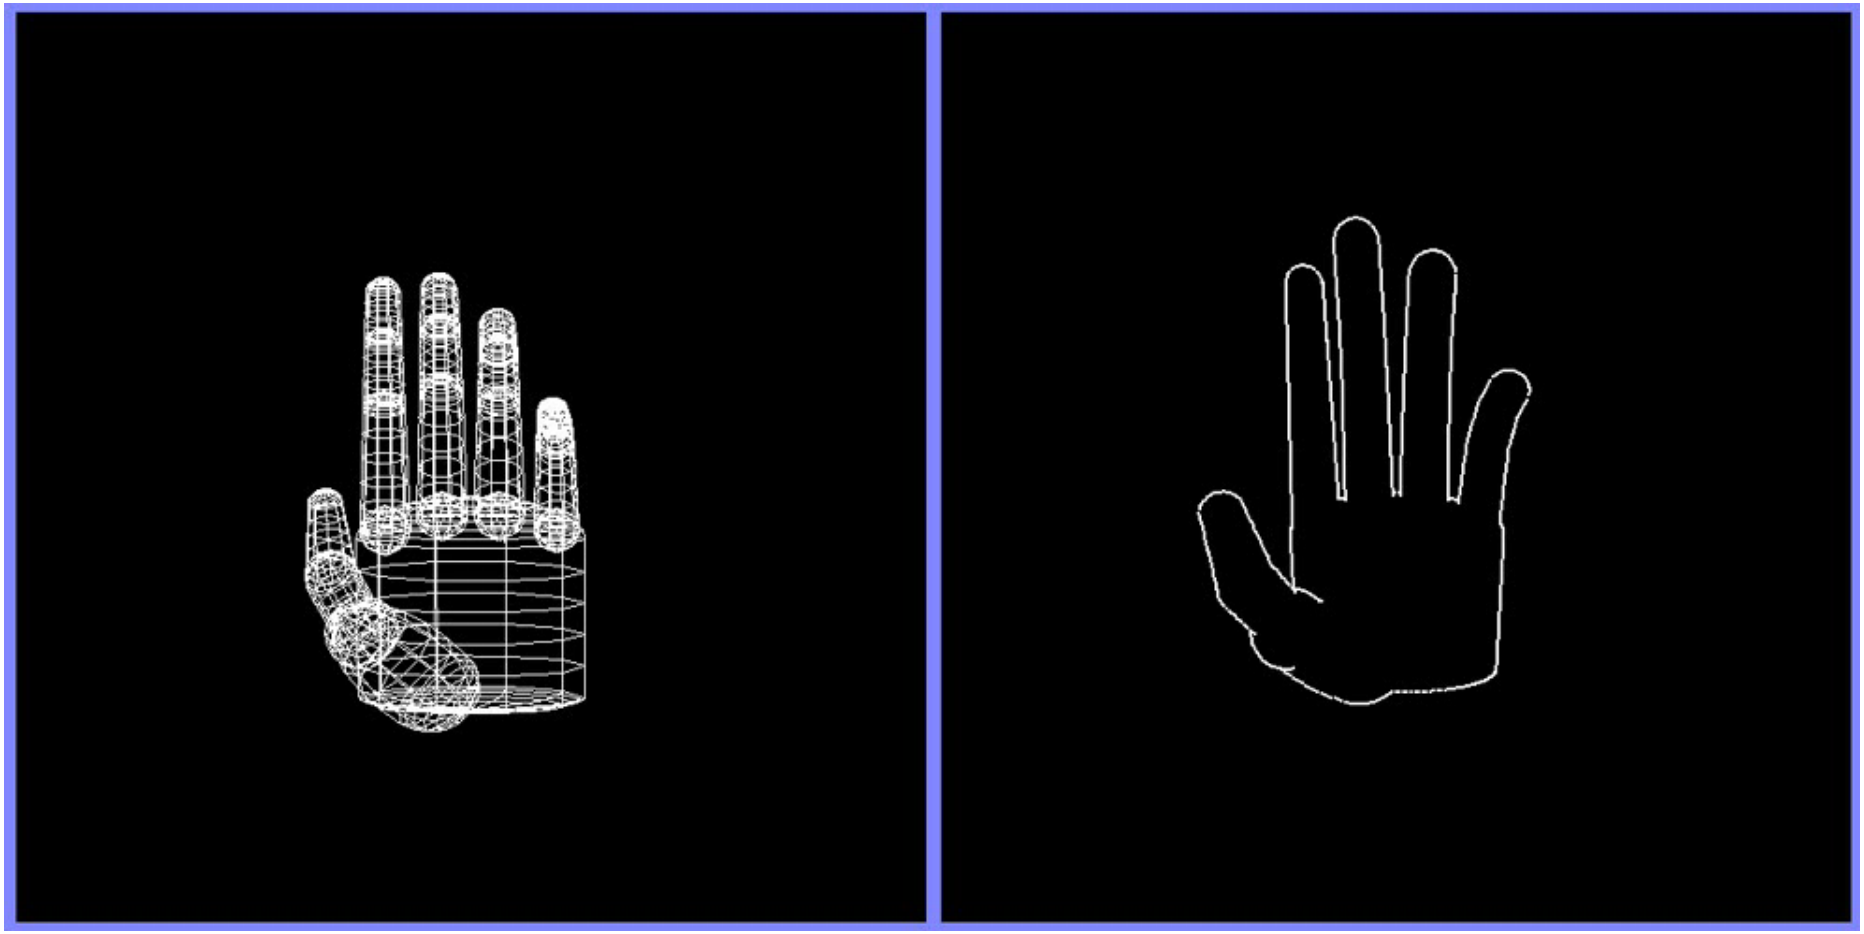

3D hand model

UNIVERSITY OF CAMBRIDGE

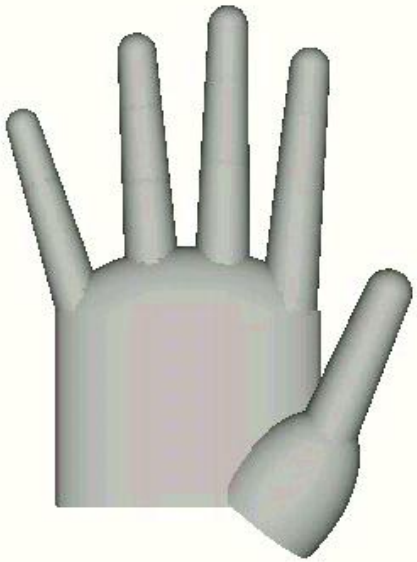

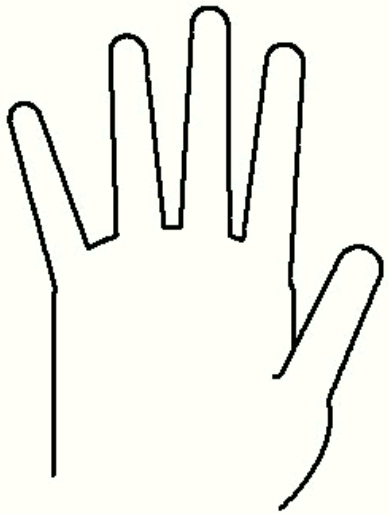

Model-based tracking

UNIVERSITY OF CAMBRIDGE

3D tracking, 6/7 DOF

- **Model:** 3D quadrics
- **Cost Function:** Edges or colour-edges
- **Tracking:** Unscented Kalman filtering
- Single or dual view
- Single hypothesis filter, no recovery strategy

UNIVERSITY OF CAMBRIDGE

Tracking as detection

Pedestrian detection

UNIVERSITY OF CAMBRIDGE

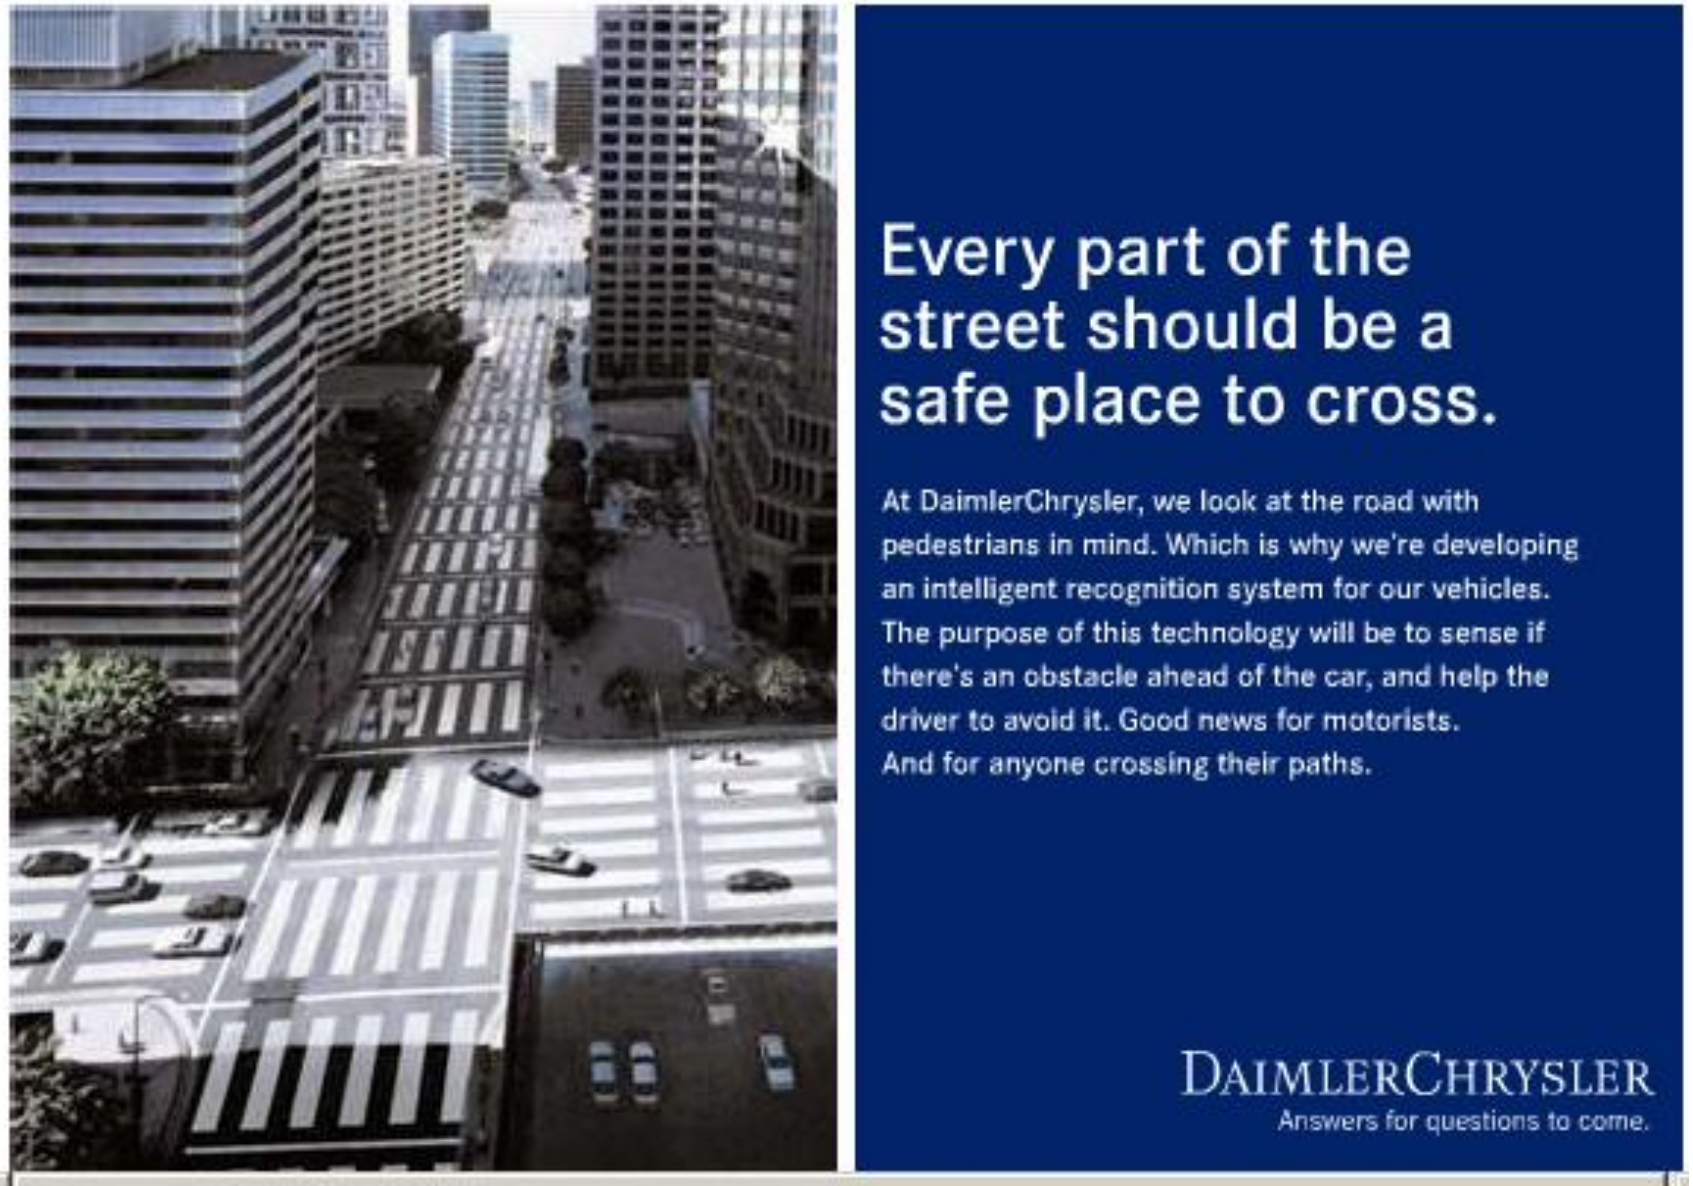

Hierarchical matching with trees

UNIVERSITY OF CAMBRIDGE

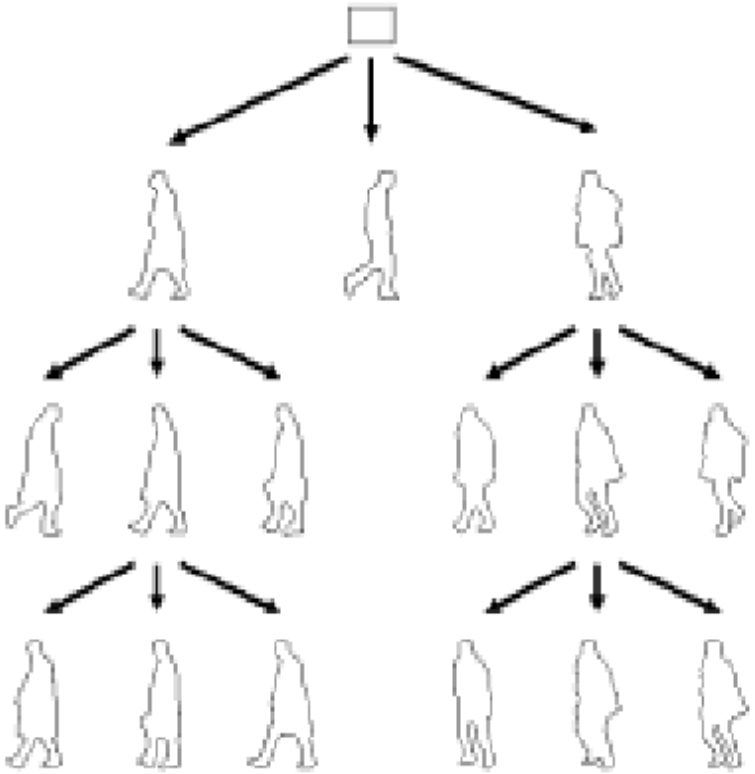

Pedestrian detection

UNIVERSITY OF CAMBRIDGE

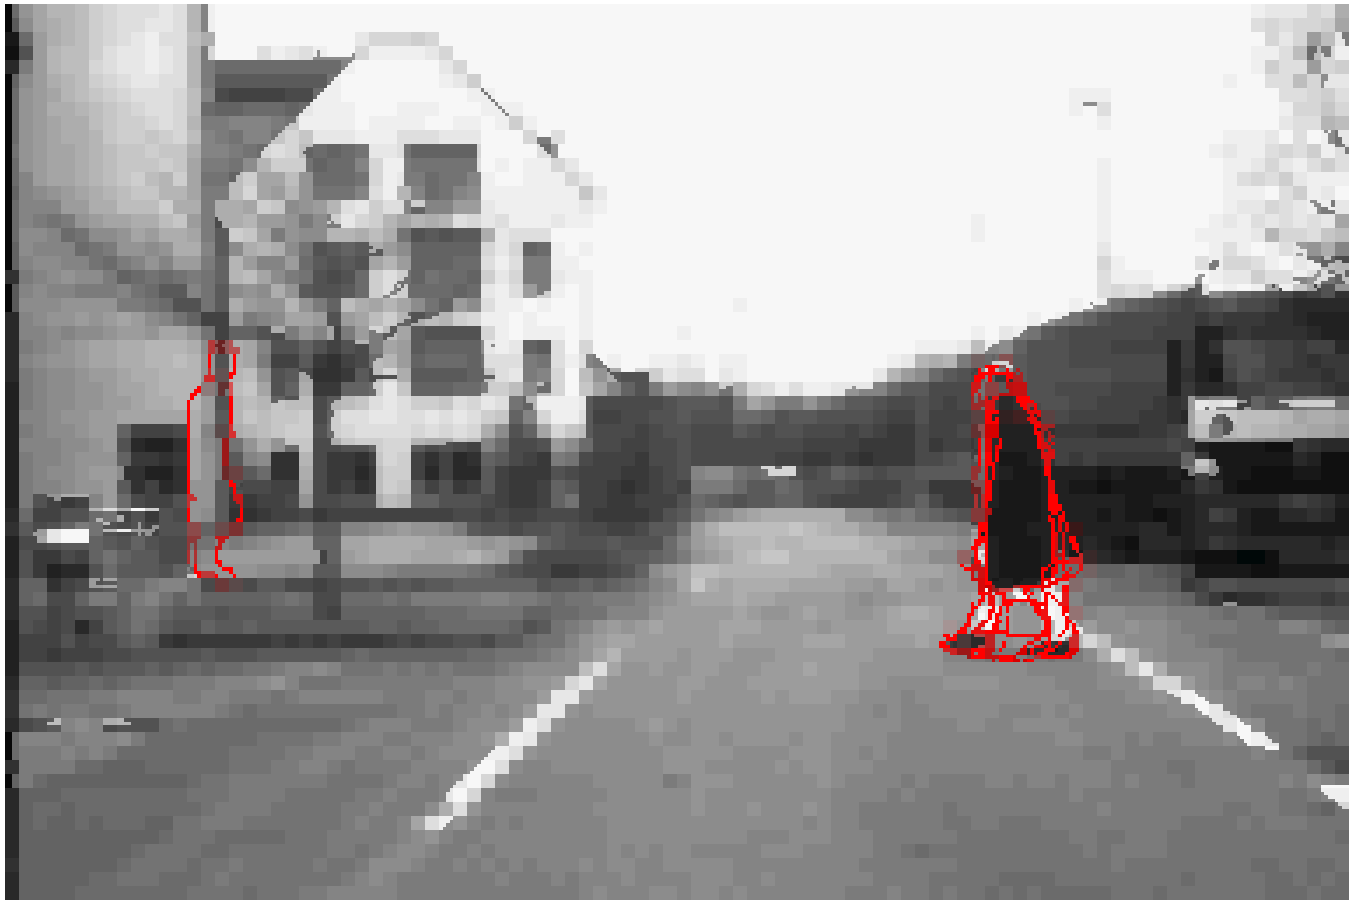

Hand detection system

UNIVERSITY OF CAMBRIDGE

UNIVERSITY OF CAMBRIDGE

Tree-based bayesian filtering

Stenger et al 2003

3D hand model

UNIVERSITY OF CAMBRIDGE

- Used as generative model
- Constructed from 35 truncated quadrics (ellipsoids, cones)
- Efficient contour projection
- 27 degrees of freedom

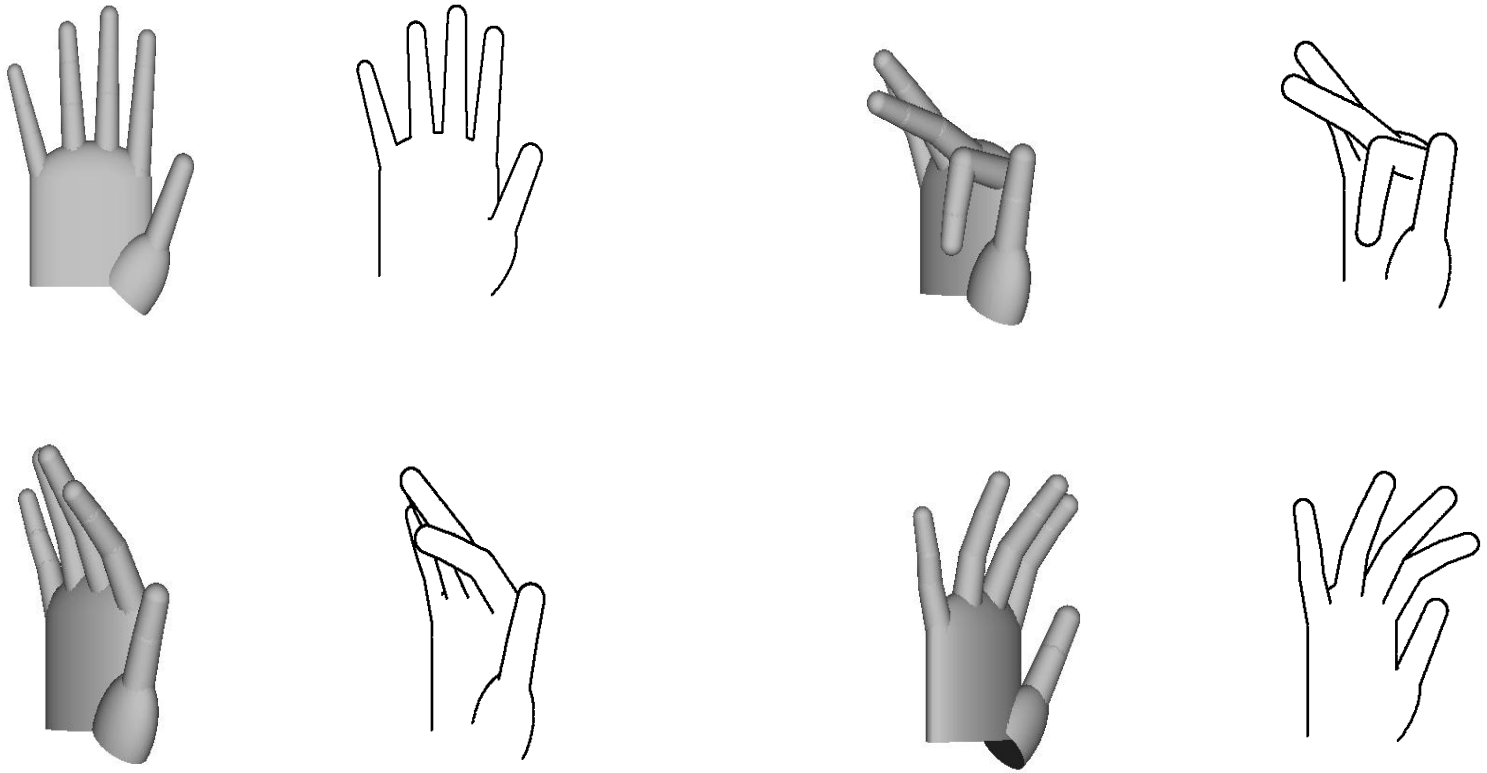

Template-based Detection

UNIVERSITY OF CAMBRIDGE

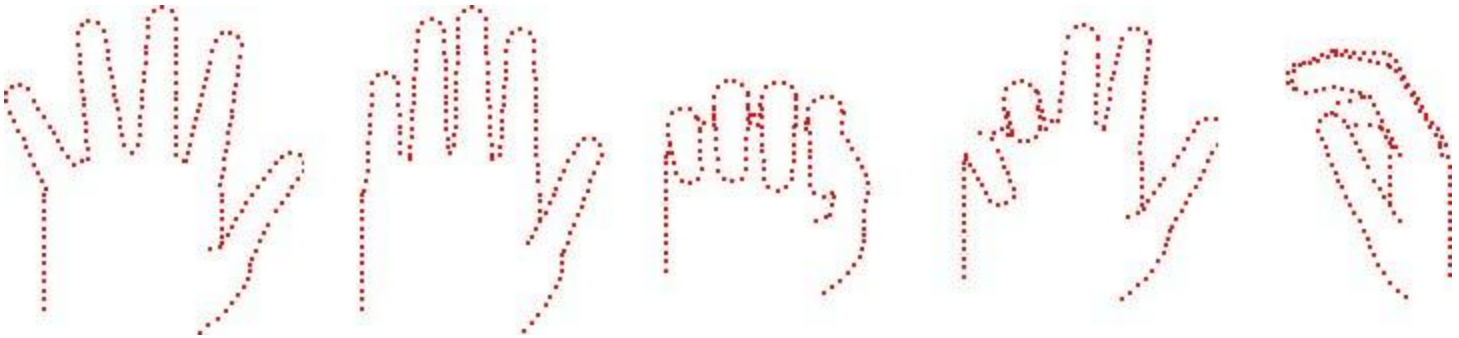

- Large number of templates are generated off-line to handle global motion and finger articulation.

- Need for
  - Inexpensive template-matching function
    - Distance Transform and Chamfer Matching
  - Efficient search structure
    - Bayesian Tree structure

Matching oriented edges

UNIVERSITY OF CAMBRIDGE

Input Image

Edge Detection

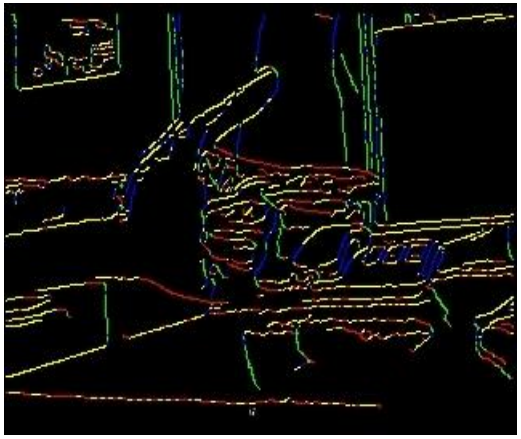

Robust Edge Matching

Using Chamfer Distance

3D Model

Projected Contours

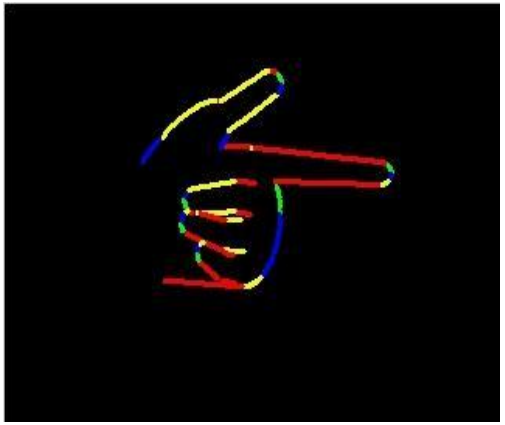

Skin colour features

UNIVERSITY OF CAMBRIDGE

Input Image

3D Model

Skin Colour Model

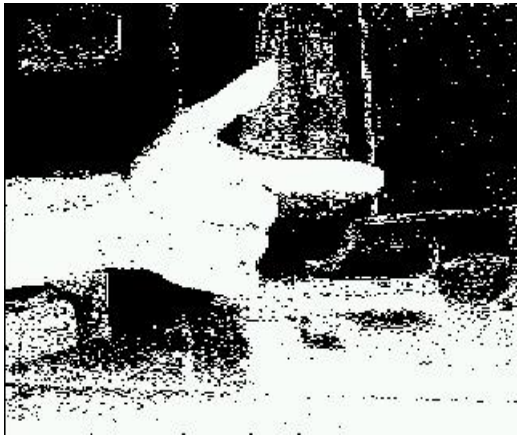

Efficient Template Matching

Projected Silhouette

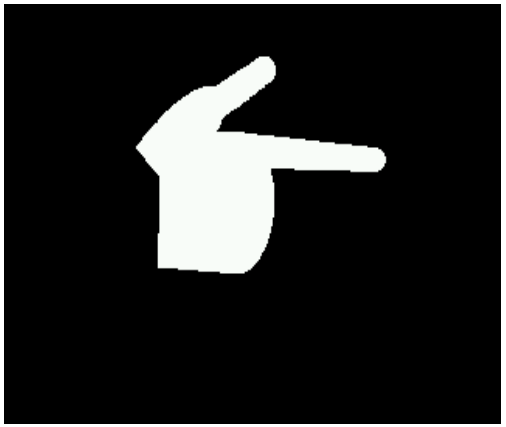

Matching Multiple Templates

UNIVERSITY OF CAMBRIDGE

- Use tree structure to efficiently match many templates (>10,000)
- Arrange templates in tree based on their similarity
- Traverse tree using breadth-first search, several ‘active’ leaves possible

Grid-based partitioning of parameter space

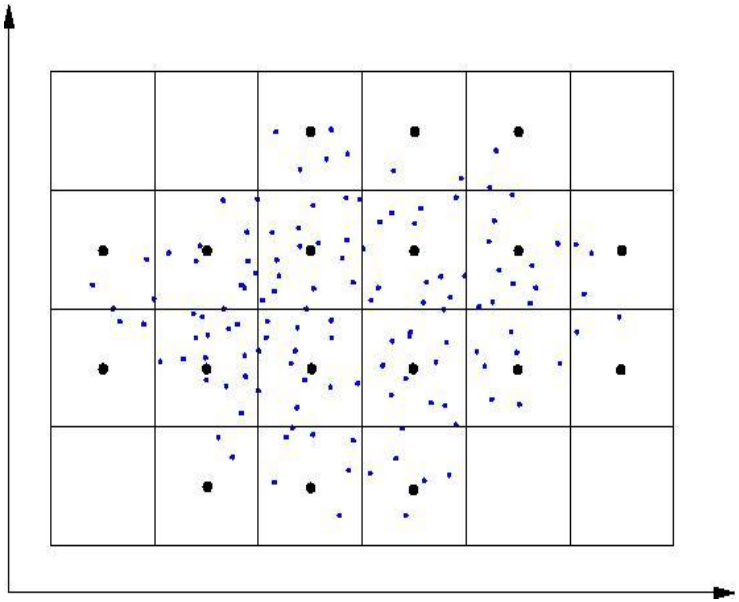

Search Tree

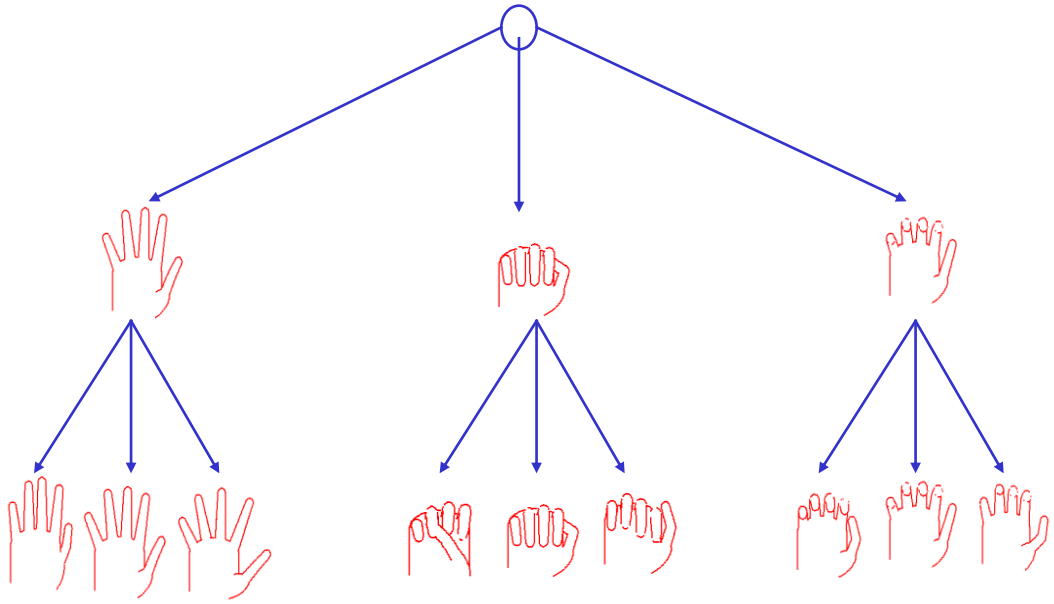

Bayesian-Tree

UNIVERSITY OF CAMBRIDGE

State space partitioning

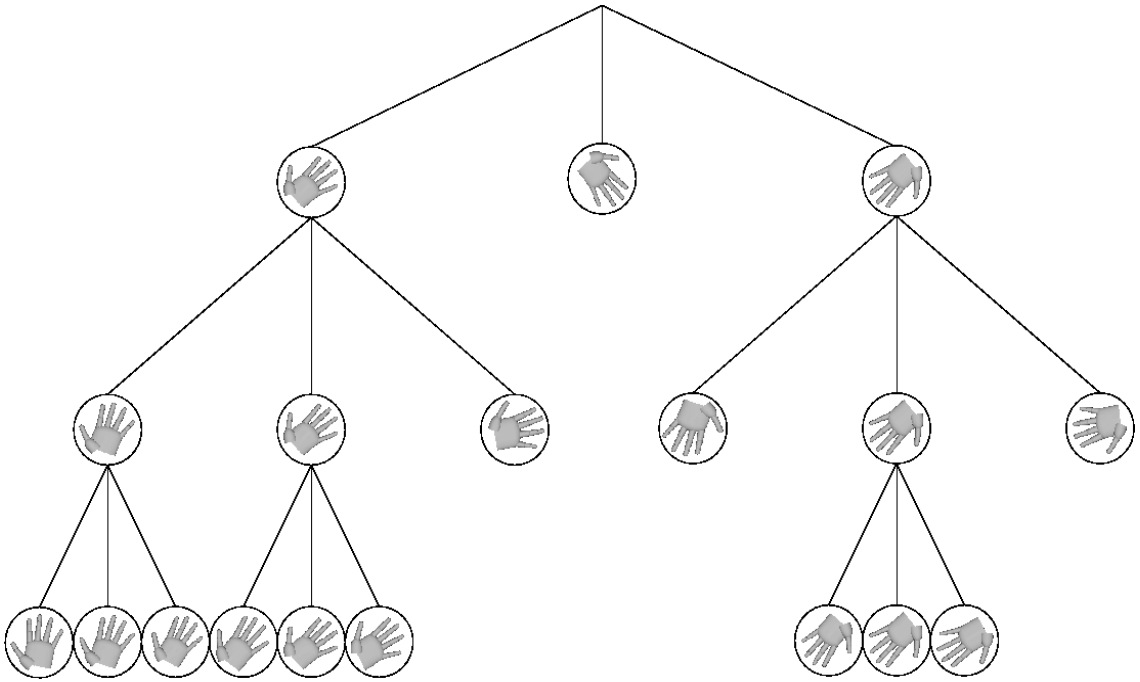

Estimation of posterior pdf

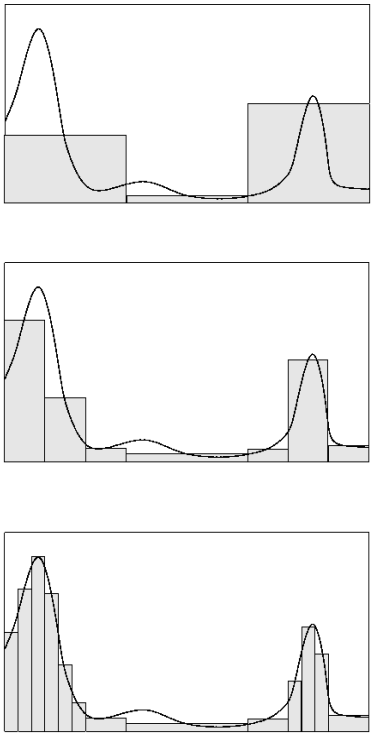

- The search-tree is brought into a Bayesian framework by adding the prior knowledge from previous frame.
- The Bayesian-Tree can be thought as approximating the posterior probability at different resolutions.

Tracking - 3D mouse

UNIVERSITY OF CAMBRIDGE

Rotating in clutter

UNIVERSITY OF CAMBRIDGE

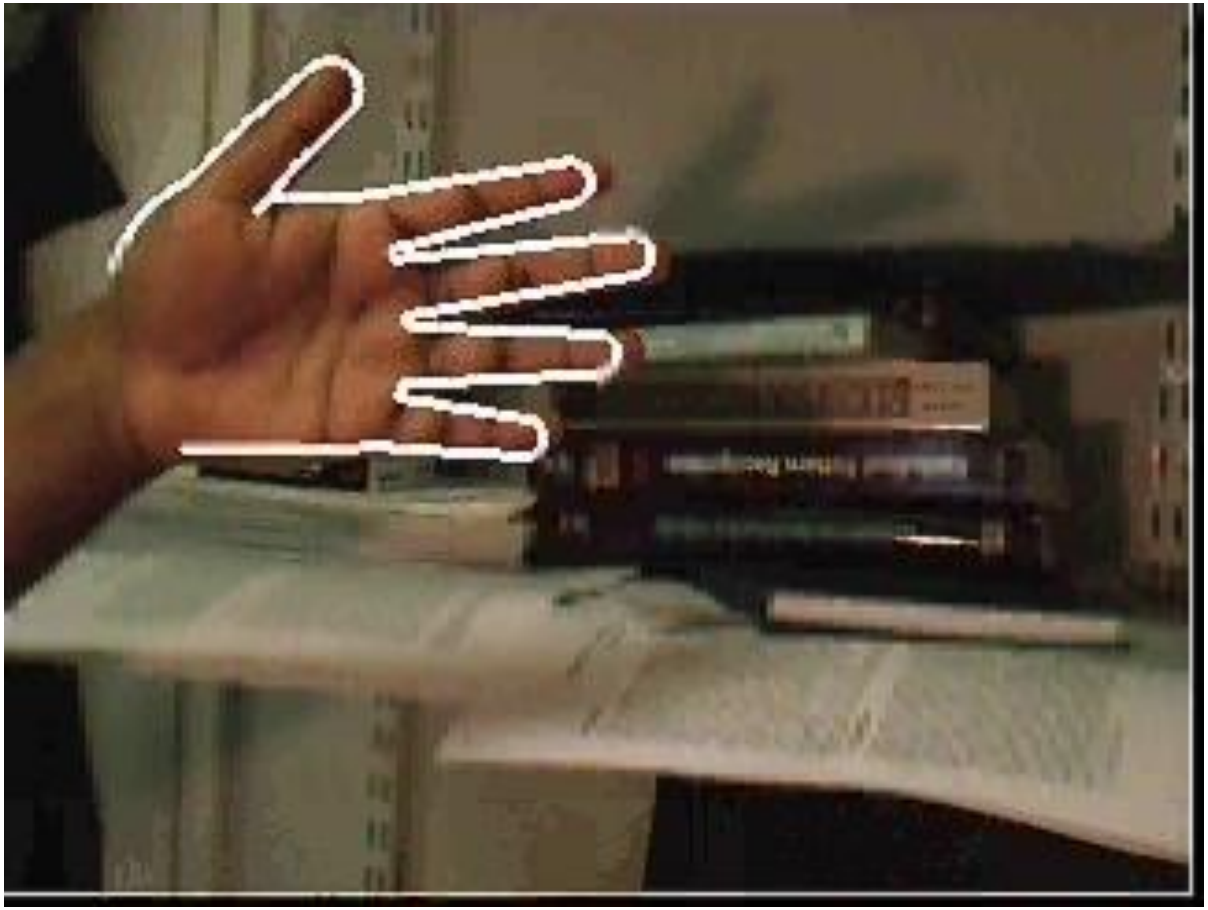

Opening and closing

UNIVERSITY OF CAMBRIDGE

UNIVERSITY OF CAMBRIDGE

Detecting people

Ramanan Navaratnam et al.

Pose Detection

UNIVERSITY OF CAMBRIDGE

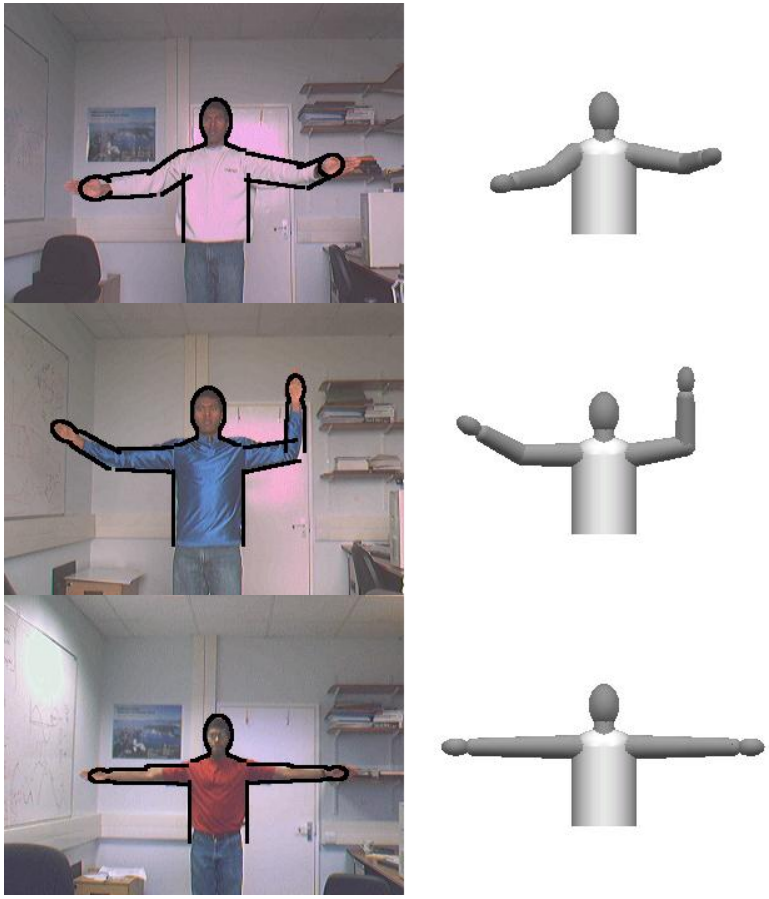

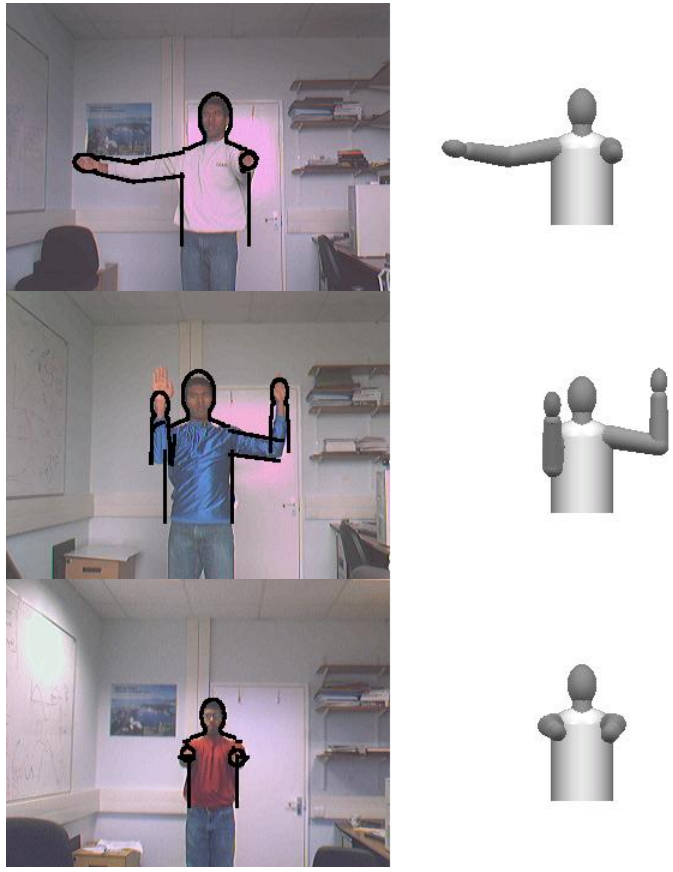

Pose Detection

UNIVERSITY OF CAMBRIDGE

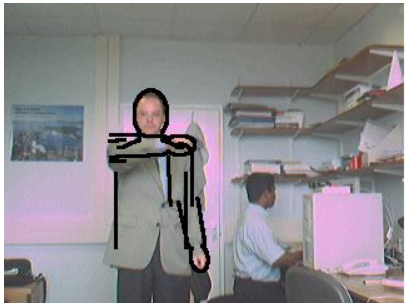

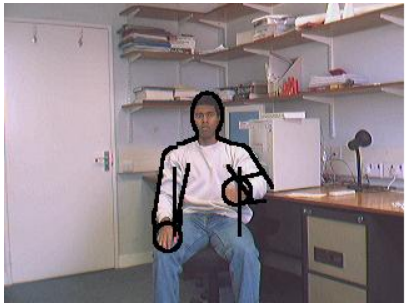

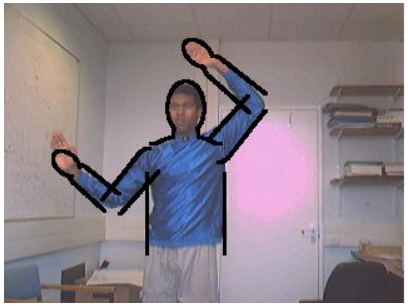


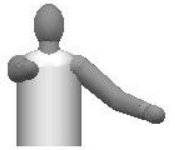
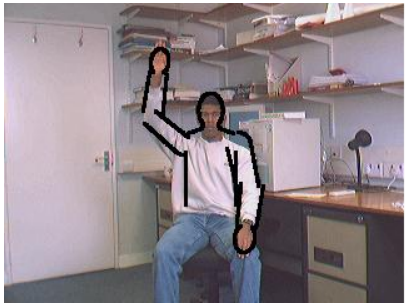

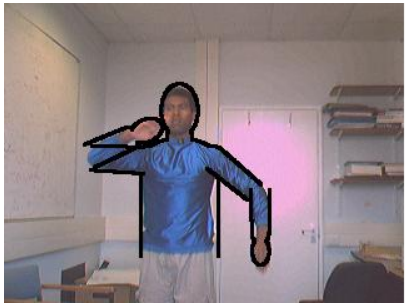


A Tracked Sequence

UNIVERSITY OF CAMBRIDGE

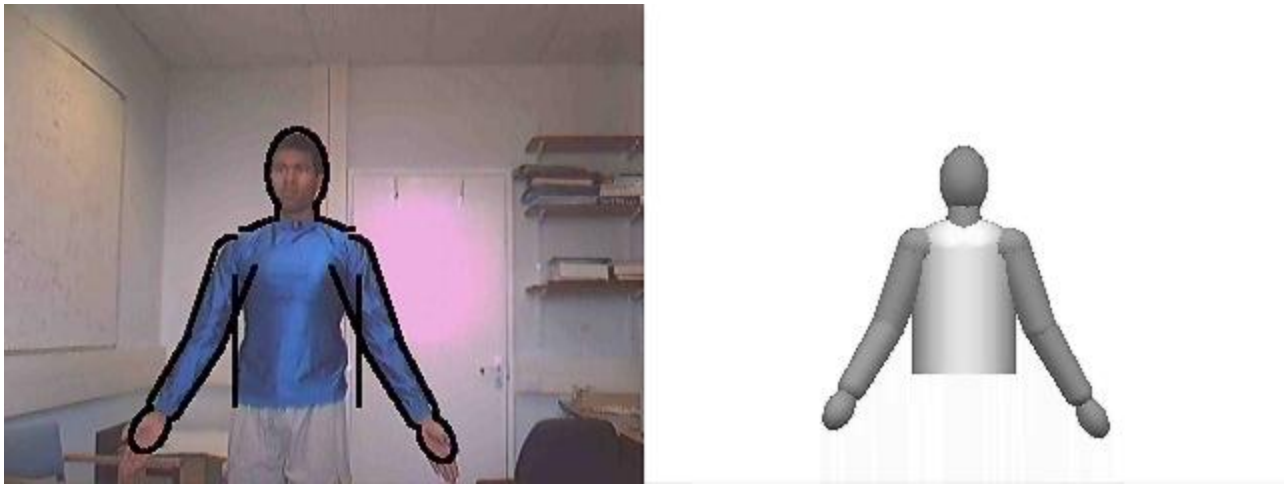

UNIVERSITY OF CAMBRIDGE

Boosted Chamfer Features for Object Detection

Jamie Shotton et al.

Supervised learning

UNIVERSITY OF CAMBRIDGE

- Learn to recognise images of a particular class, localised in space and scale
- i.e. find the horse/cow/car etc!

Desired Results

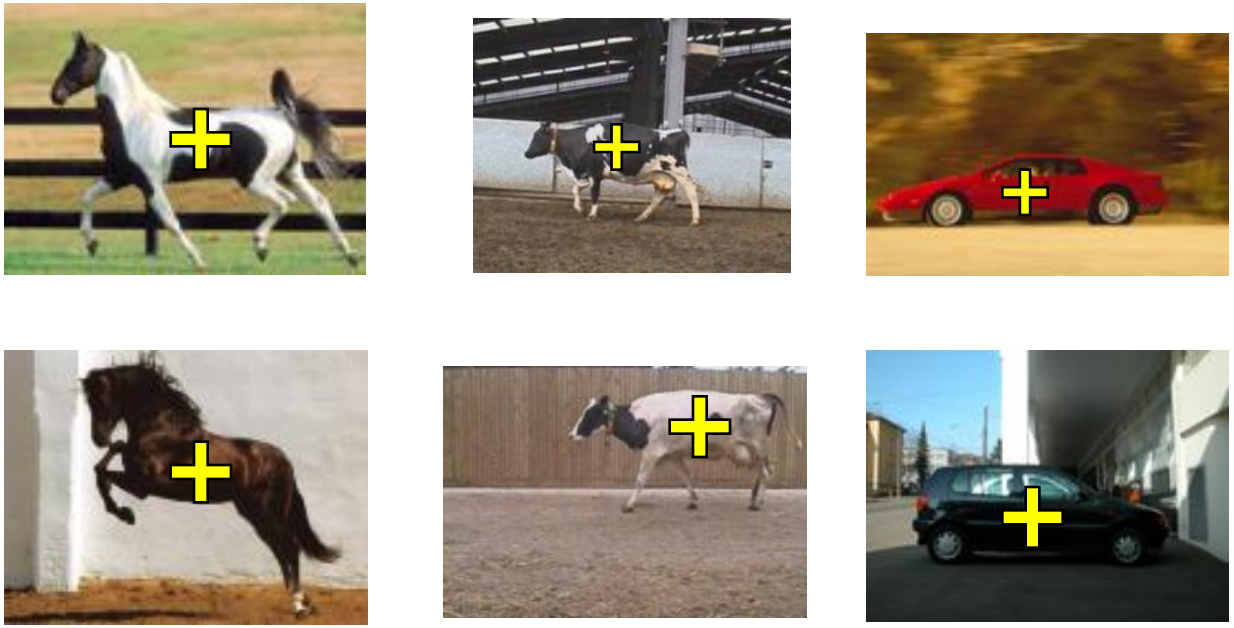

Object Model

UNIVERSITY OF CAMBRIDGE

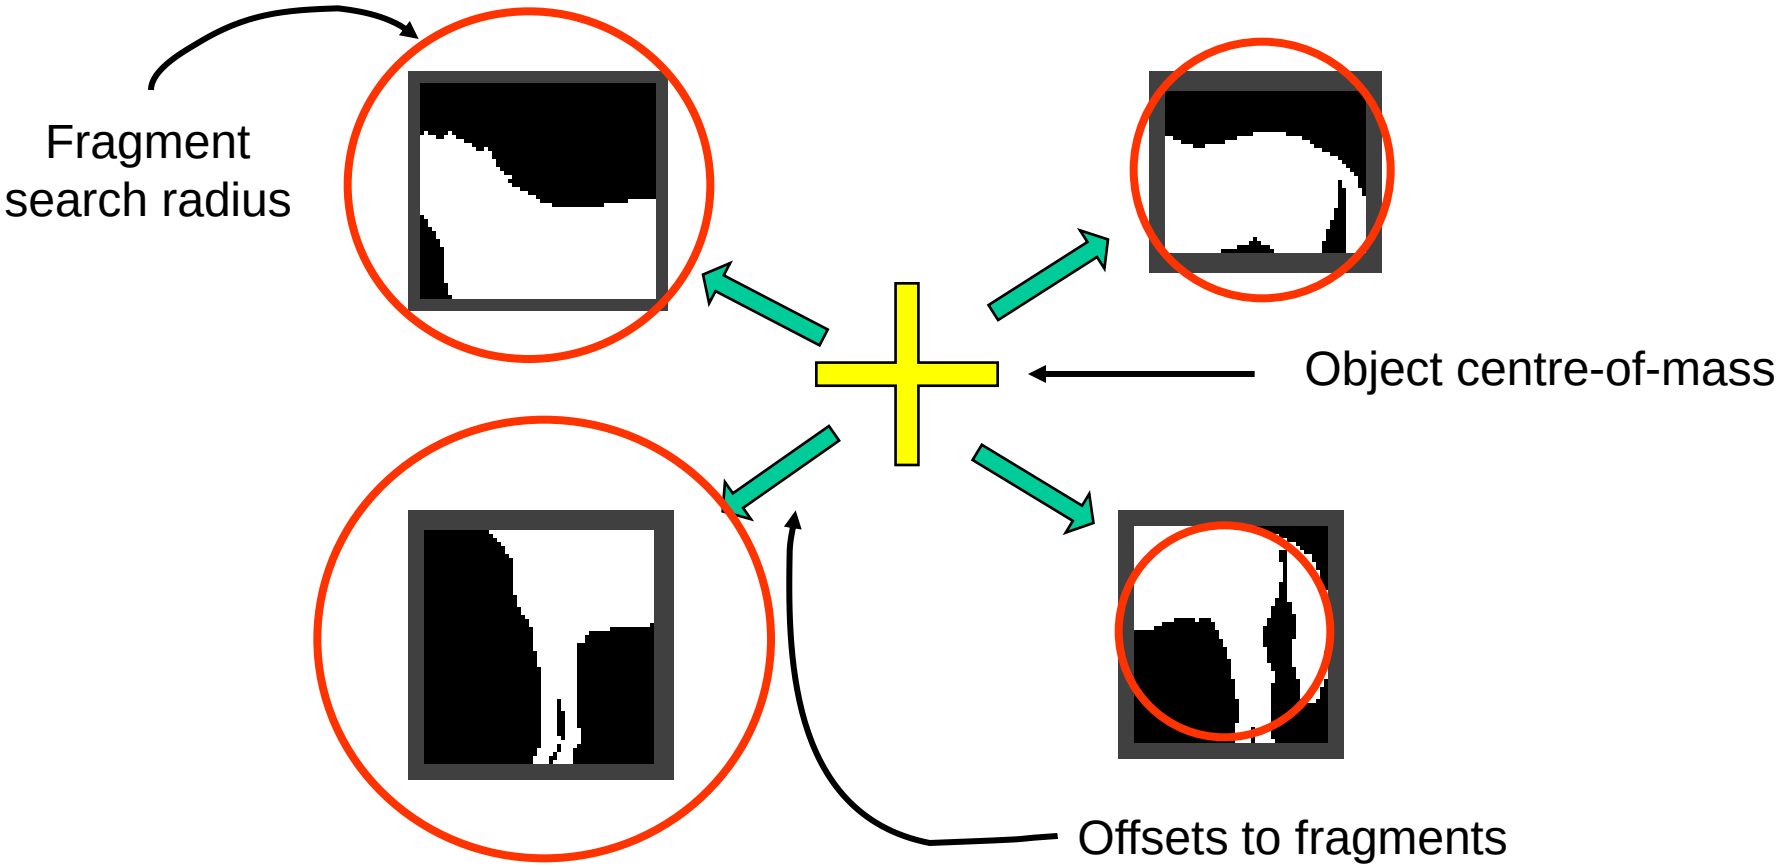

Results - classification maps

UNIVERSITY OF CAMBRIDGE

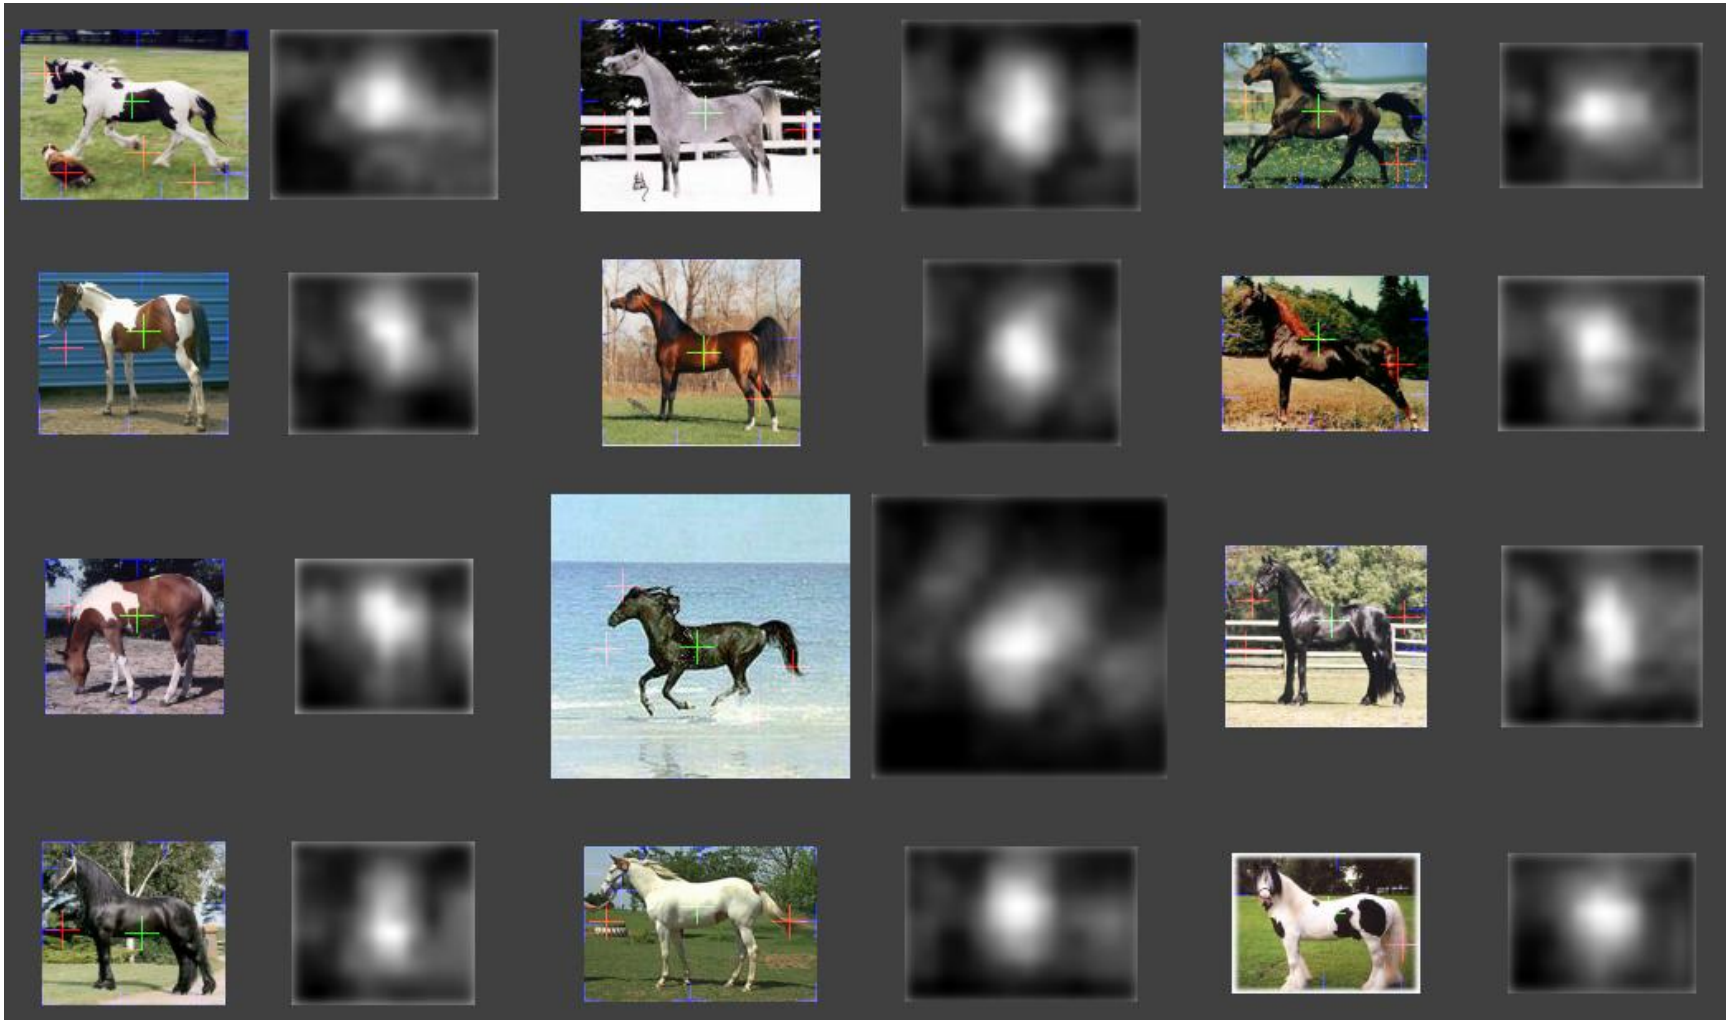

ROC curve & fragment visualisation

UNIVERSITY OF CAMBRIDGE

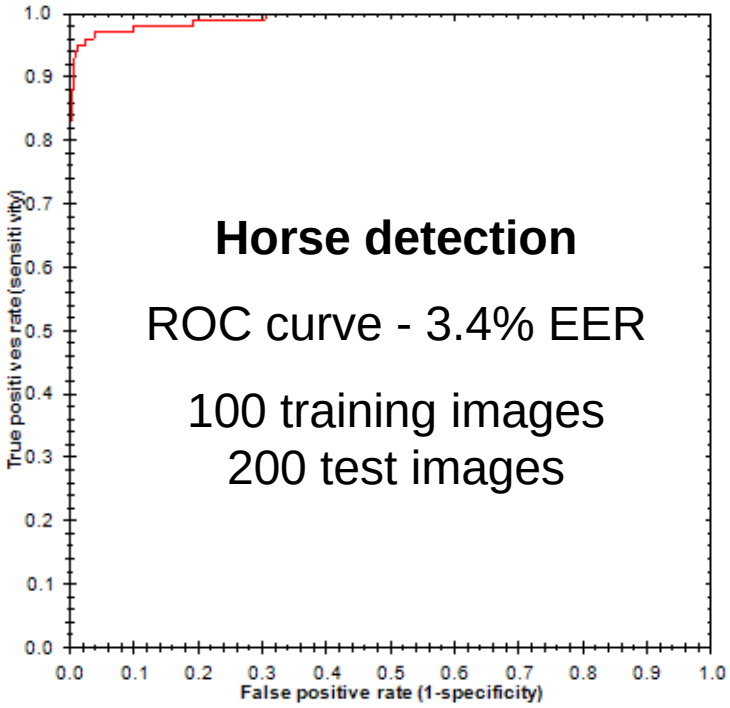

Horse detection

ROC curve - 3.4% EER

100 training images
200 test images

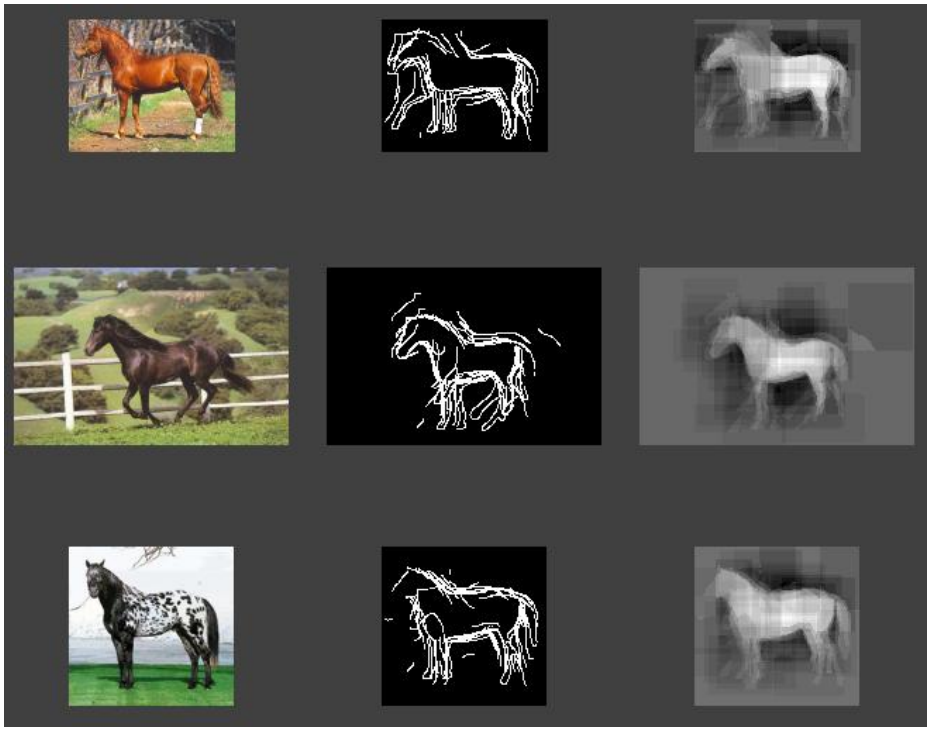

UNIVERSITY OF CAMBRIDGE

New horizons

UNIVERSITY OF CAMBRIDGE

Statistical learning and inference

Ollie Williams et al 2003

Robust Face Tracking

UNIVERSITY OF CAMBRIDGE

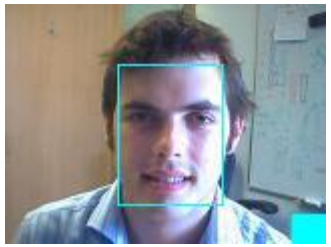

- Self starting
- Self recovering
- Efficient

object detector

exploit temporal coherence

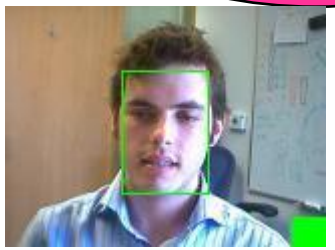

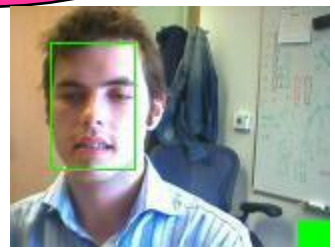

Creating a Training Set

UNIVERSITY OF CAMBRIDGE

- Select a few “seed” stills
- Simulate translation, scaling and rotation
  - → labelled training set

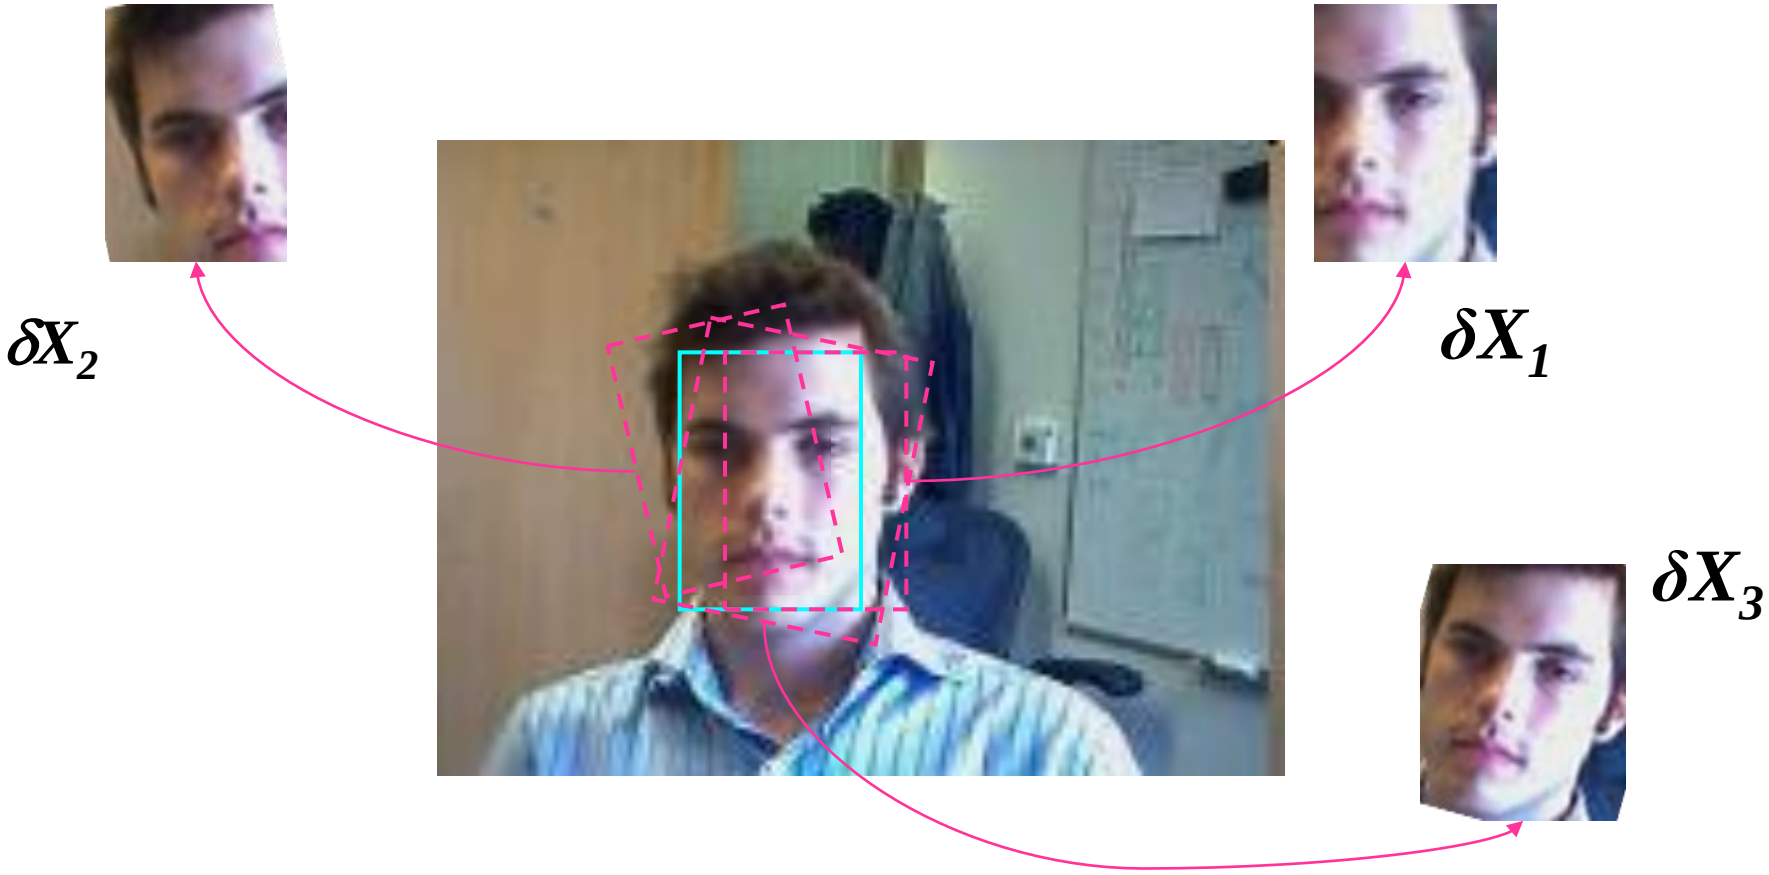

RVM Tracking

UNIVERSITY OF CAMBRIDGE

$\{z_i\}$ 

$\{t_i\}$ δX_1 δX_2 δX_3 δX_N

4D RVM learns mapping

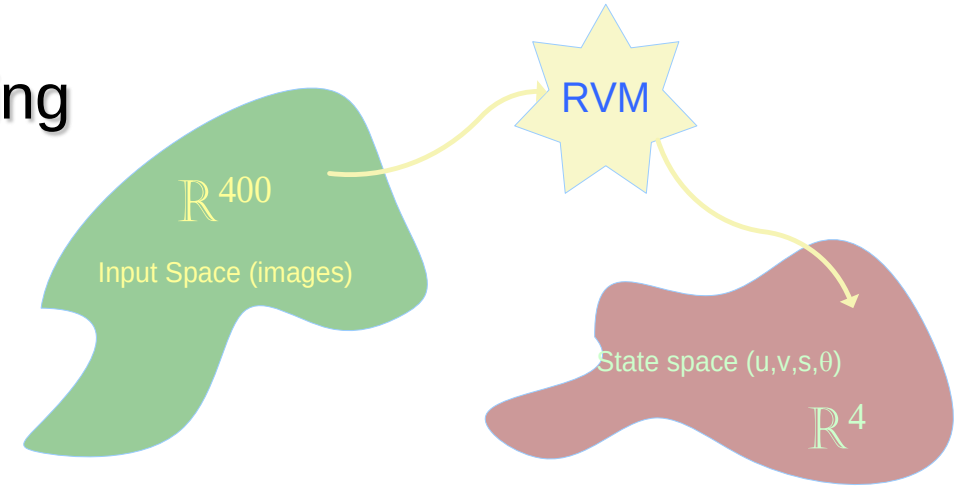

Detecting frontal faces

UNIVERSITY OF CAMBRIDGE

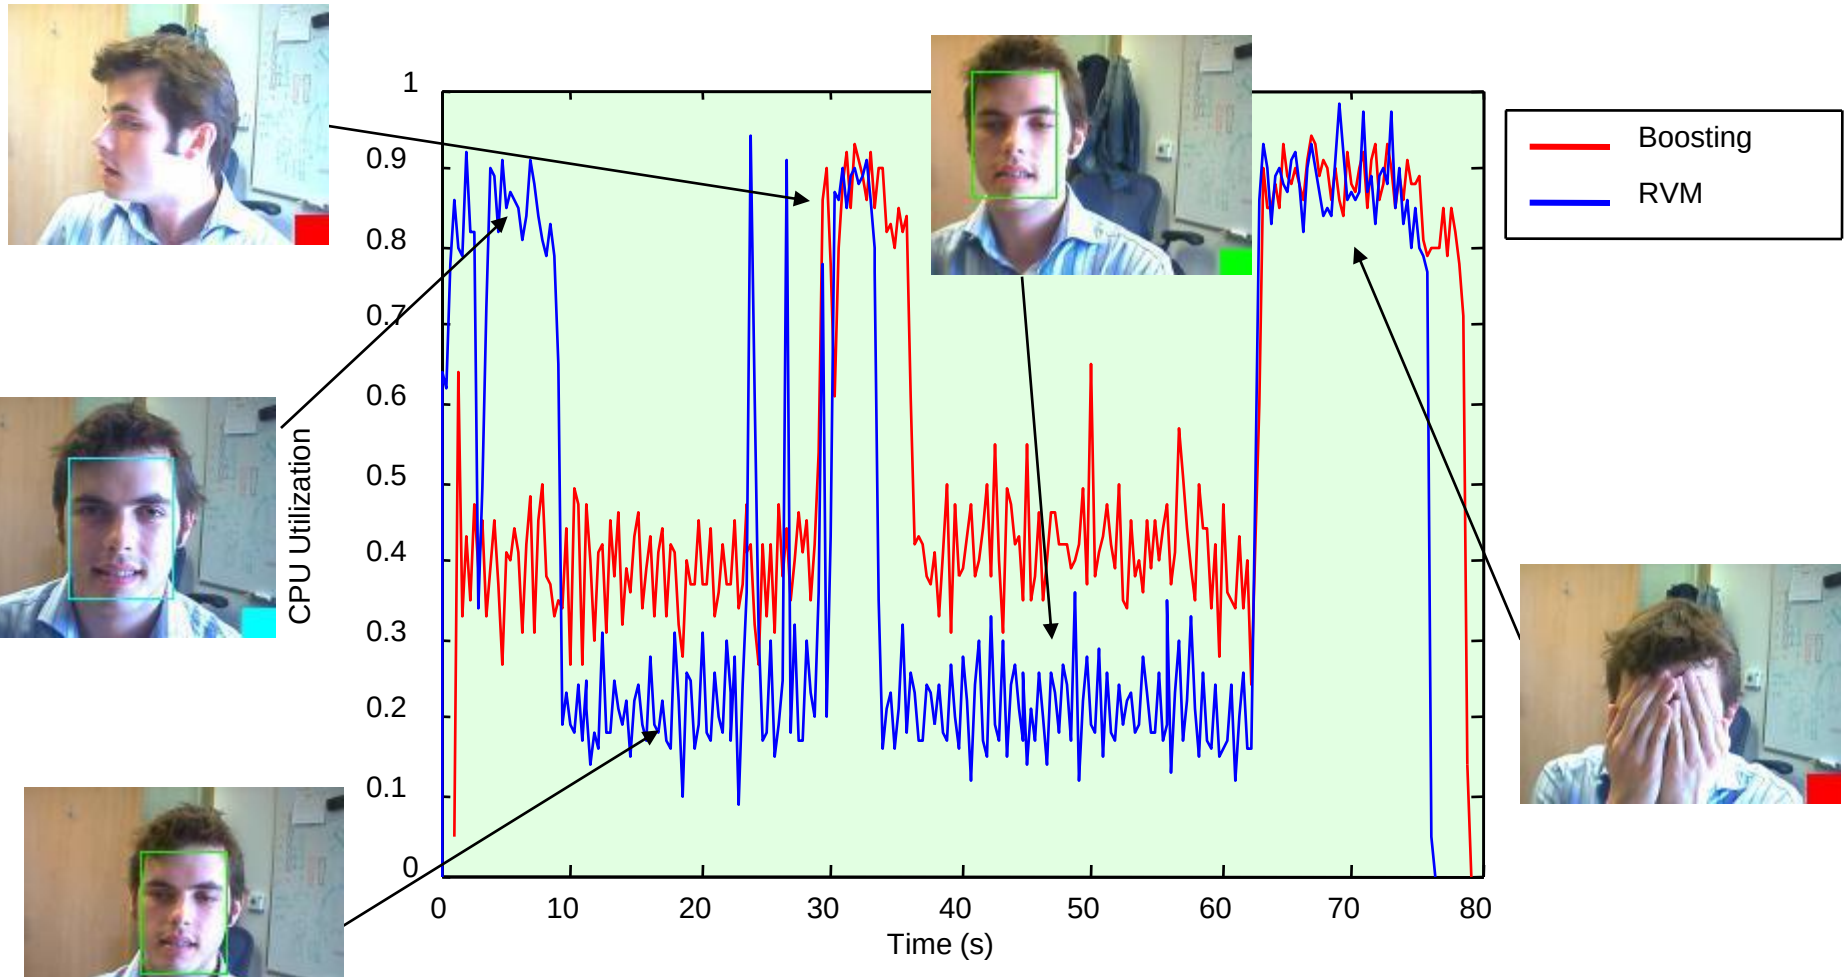

Automatic Camera Management

UNIVERSITY OF CAMBRIDGE

- Use position/scale information to control digital **pan** and **zoom**

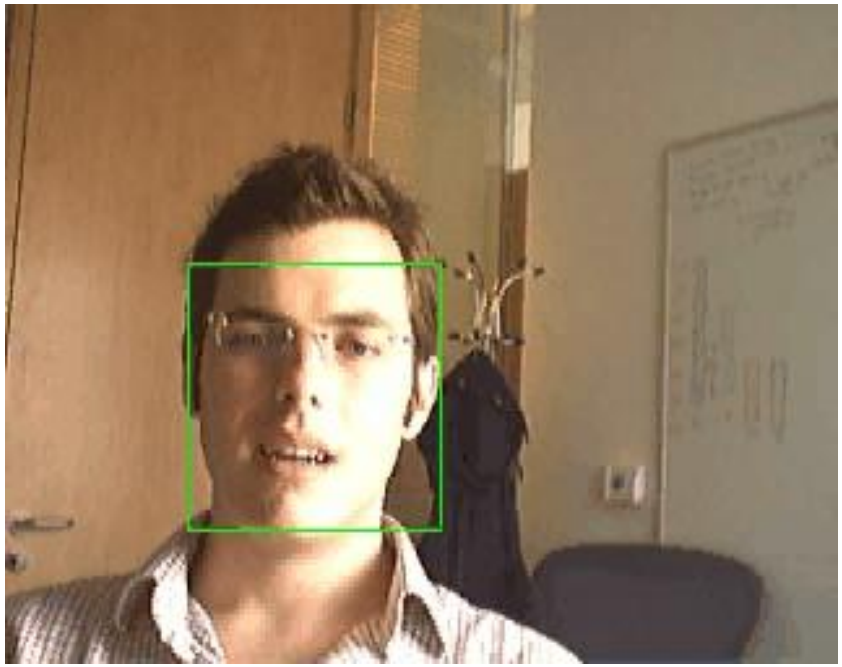

Severe Illumination Change

UNIVERSITY OF CAMBRIDGE

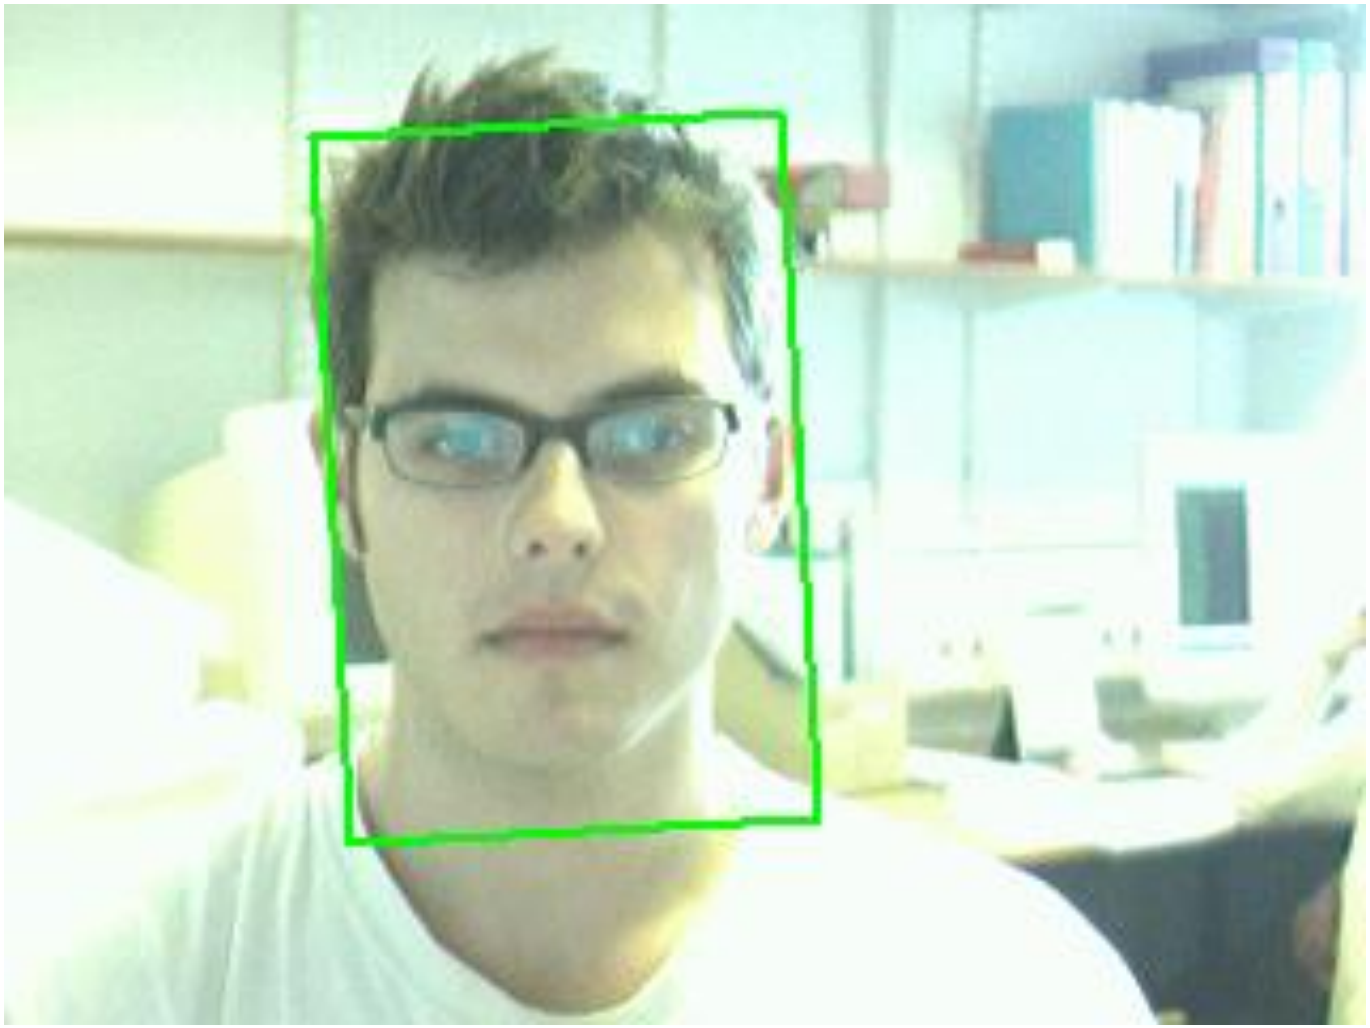

UNIVERSITY OF CAMBRIDGE

Where am I?
Image-based localisation

Johansson and Cipolla 2002

Robertson and Cipolla 2004

The goal – where am I?

UNIVERSITY OF CAMBRIDGE

User takes a picture of a nearby building. System tells you what you are looking at and exactly where you are on a map.

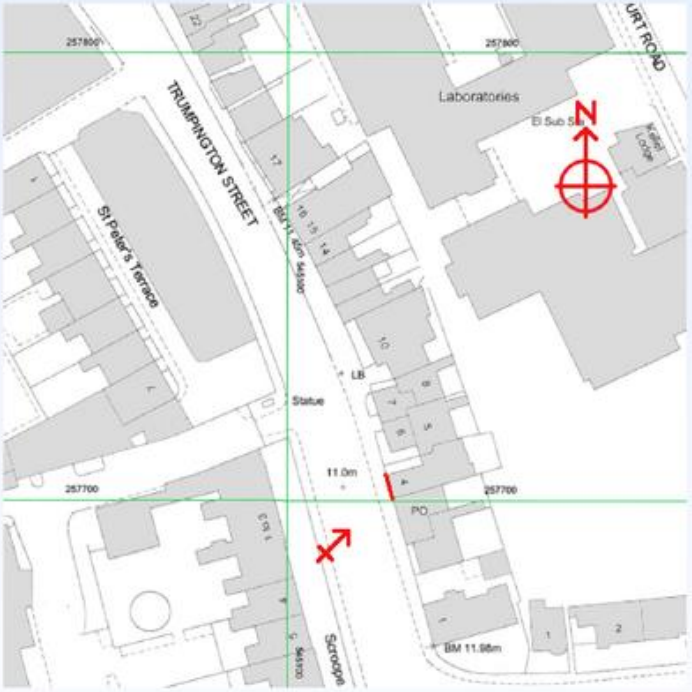

The problem

UNIVERSITY OF CAMBRIDGE

?

Why difficult?

UNIVERSITY OF CAMBRIDGE

Extreme perspective distortion

Differences in colour / lighting conditions

Occlusion

Unconstrained matching

UNIVERSITY OF CAMBRIDGE

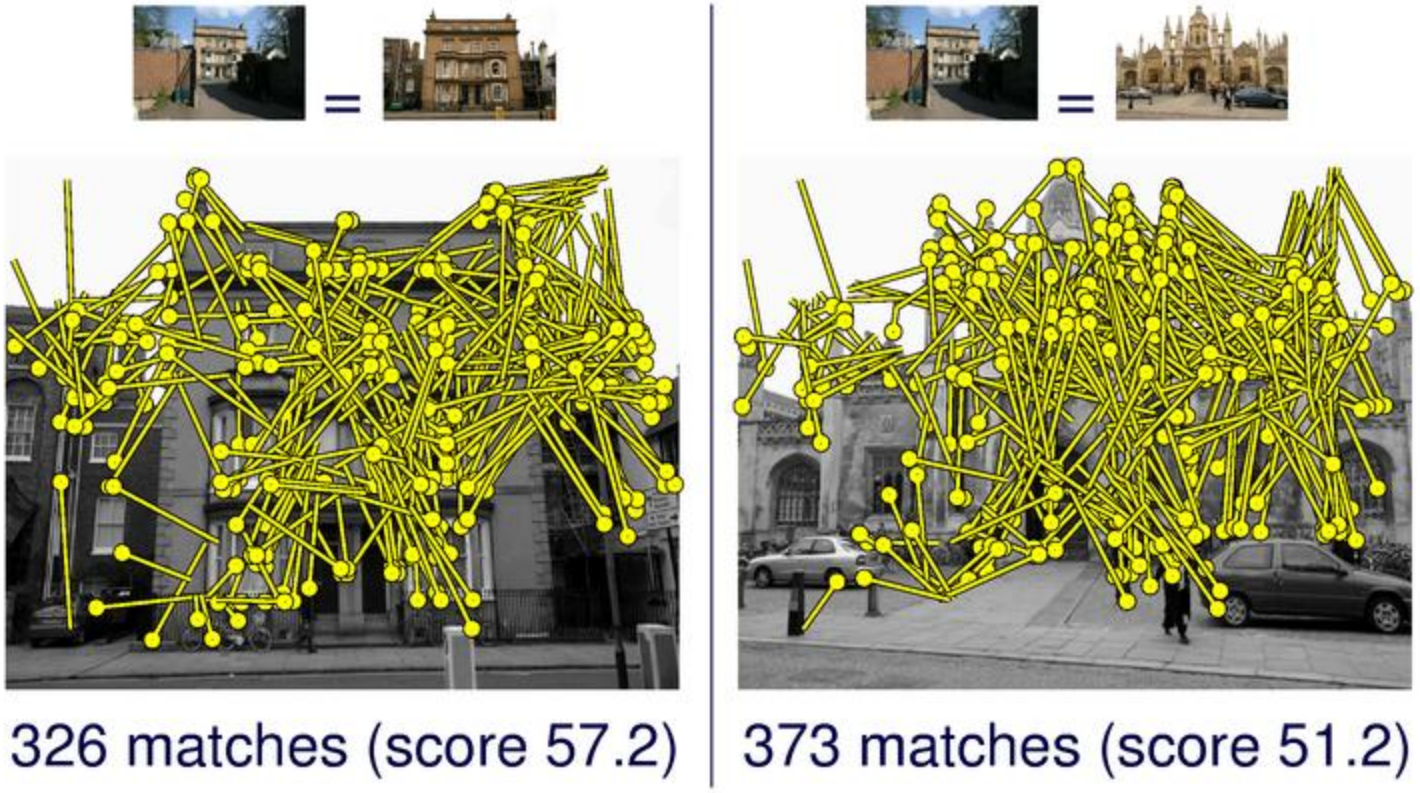

326 matches (score 57.2)

373 matches (score 51.2)

Constrained matching

UNIVERSITY OF CAMBRIDGE

- Building façades are roughly planar
- They contain many horizontal and vertical features
- We can use this to get a “front view” (rectified image)
- Front-views are related by translation and scale only

Constrained matching

UNIVERSITY OF CAMBRIDGE

Constrained matching

UNIVERSITY OF CAMBRIDGE

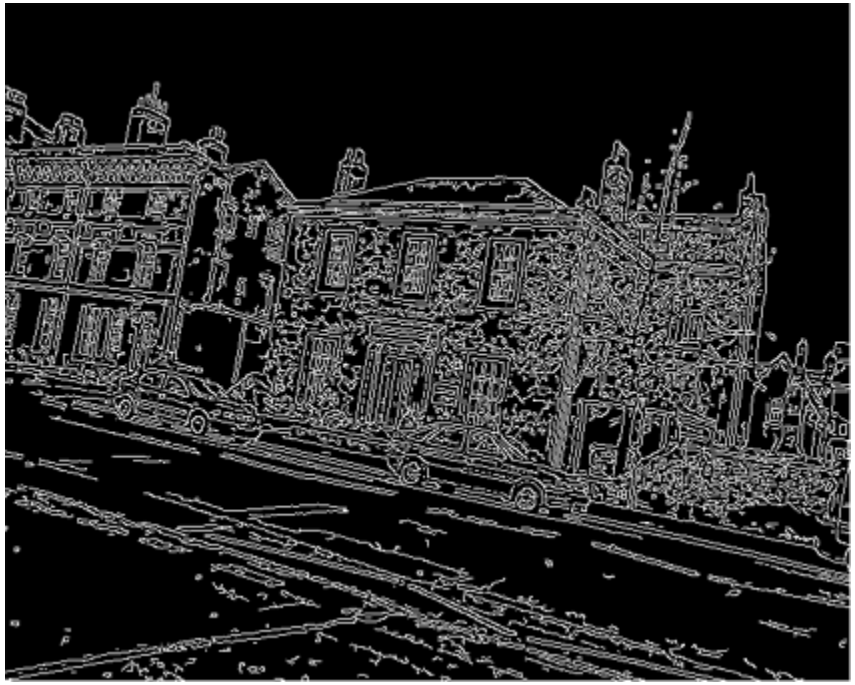

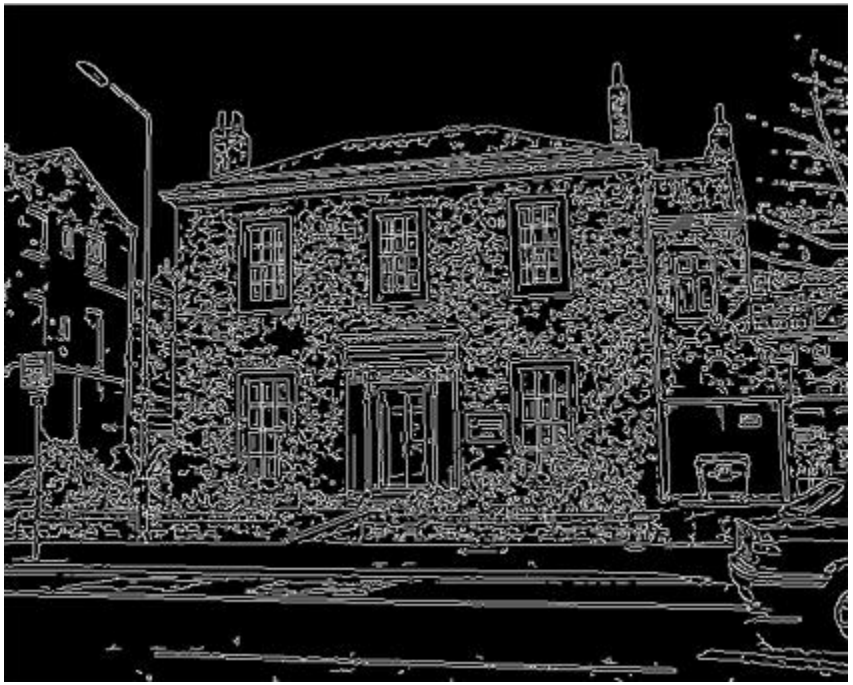

Constrained matching

UNIVERSITY OF CAMBRIDGE

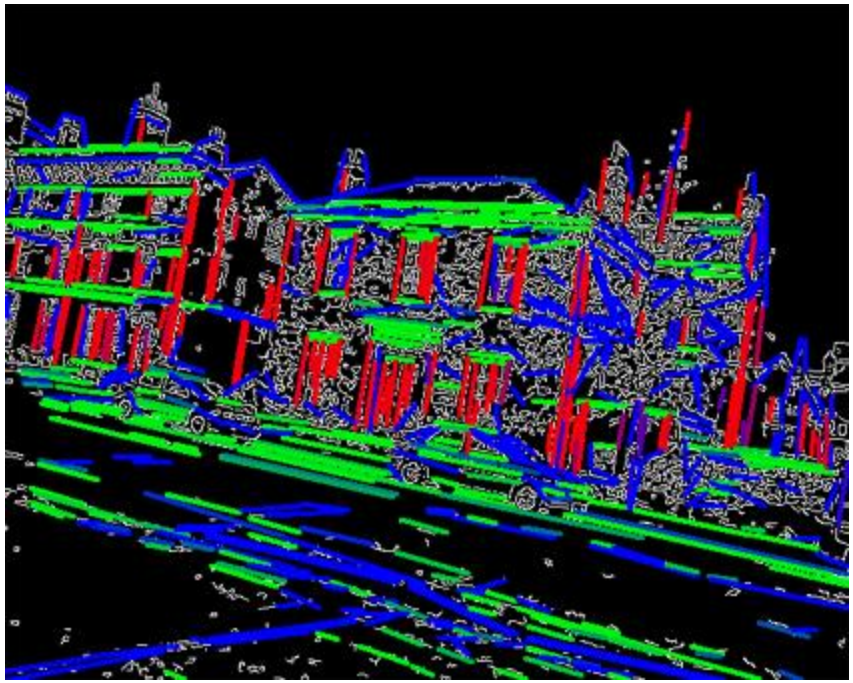

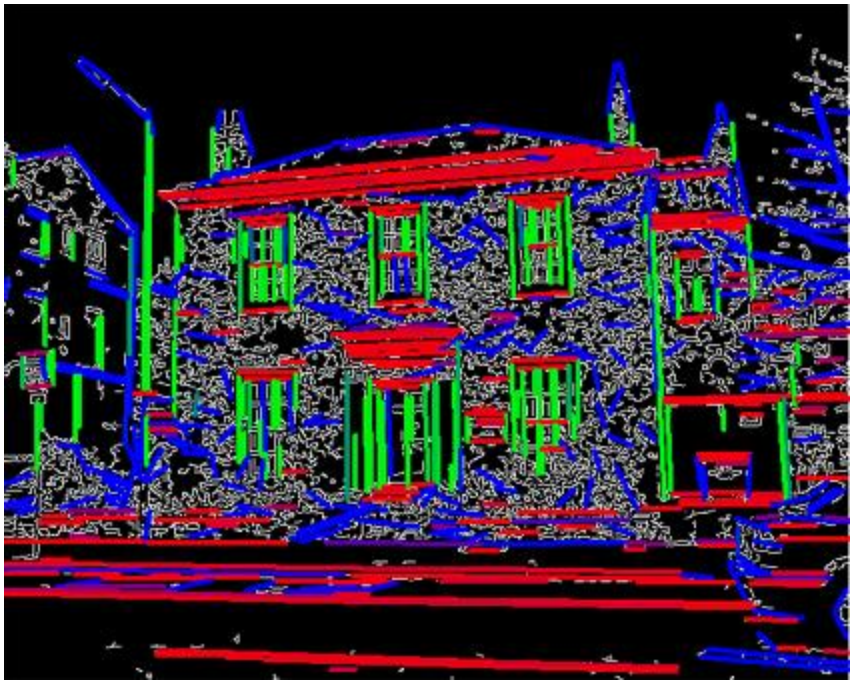

Constrained matching

UNIVERSITY OF CAMBRIDGE

Constrained matching

UNIVERSITY OF CAMBRIDGE

Overview of solution

UNIVERSITY OF CAMBRIDGE

1 vanishing point detection
2 image rectification
3 database search
4 viewpoint determination

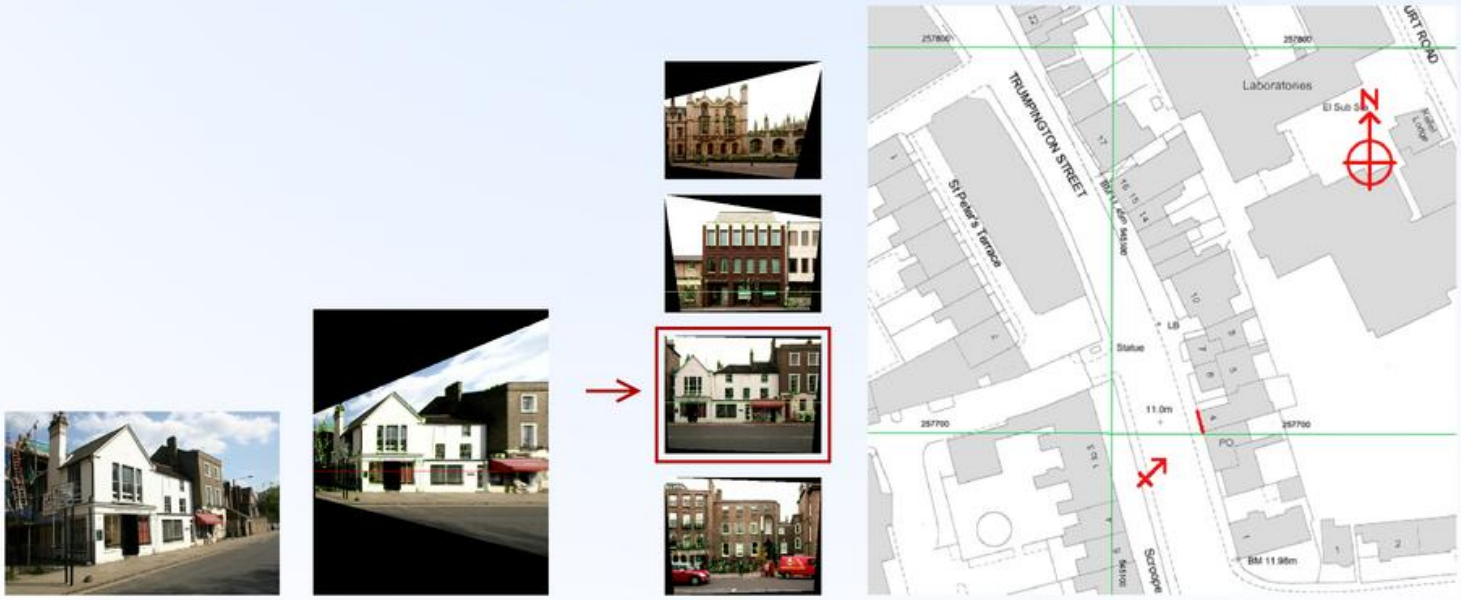

UNIVERSITY OF CAMBRIDGE

Rectification

Rectification

UNIVERSITY OF CAMBRIDGE

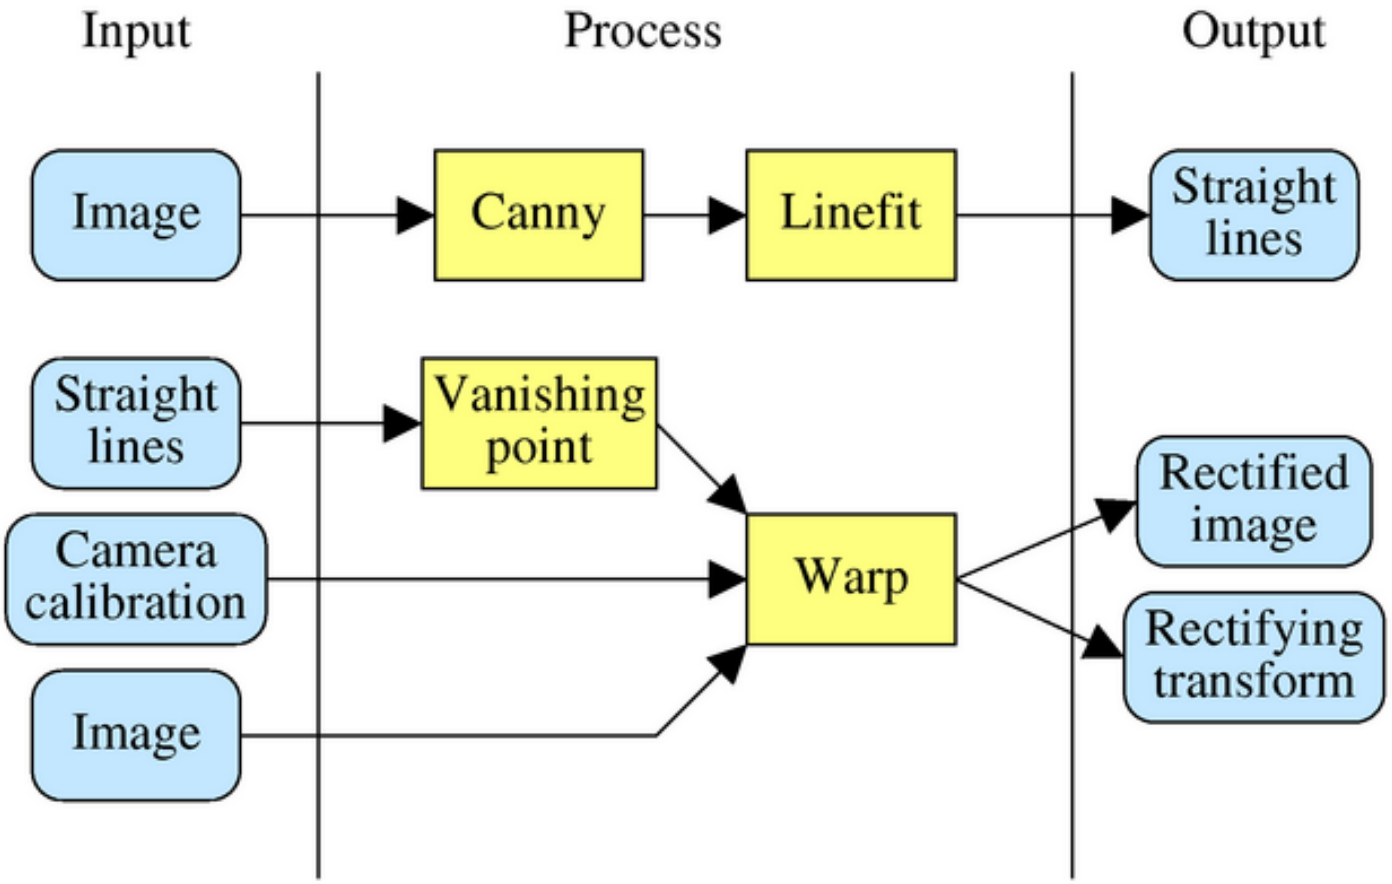

Detection of straight lines

UNIVERSITY OF CAMBRIDGE

Detect straight lines:

Finding vanishing points

UNIVERSITY OF CAMBRIDGE

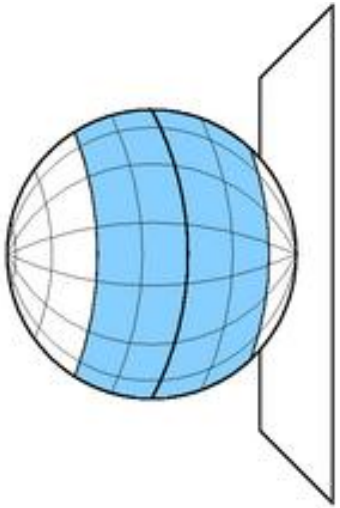

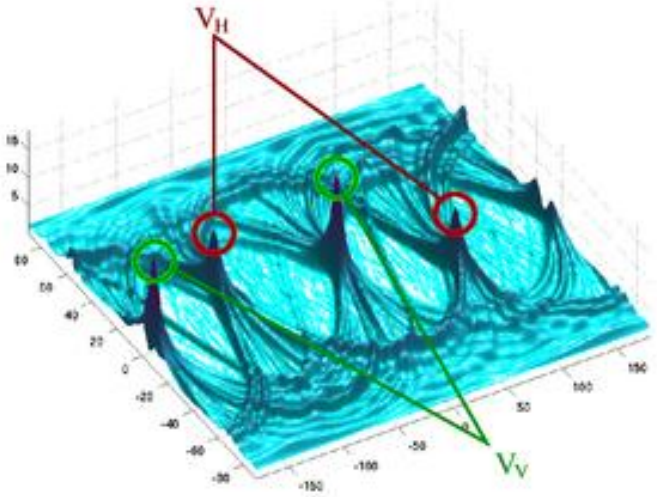

Find vertical and horizontal lines

UNIVERSITY OF CAMBRIDGE

Allocate all lines as vertical, horizontal or “clutter”

Rectification by homography

UNIVERSITY OF CAMBRIDGE

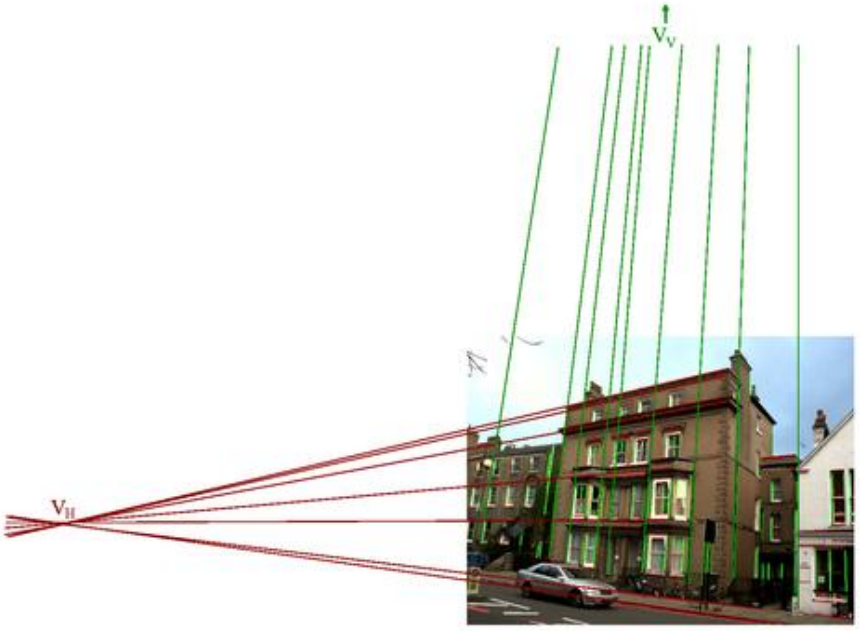

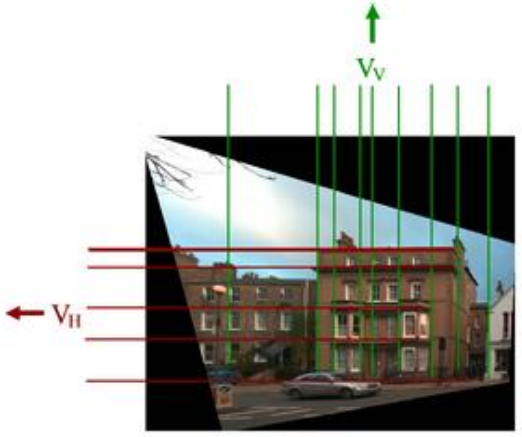

Align horizon

UNIVERSITY OF CAMBRIDGE

Only difference is now scale + x translation

UNIVERSITY OF CAMBRIDGE

Matching

Matching

UNIVERSITY OF CAMBRIDGE

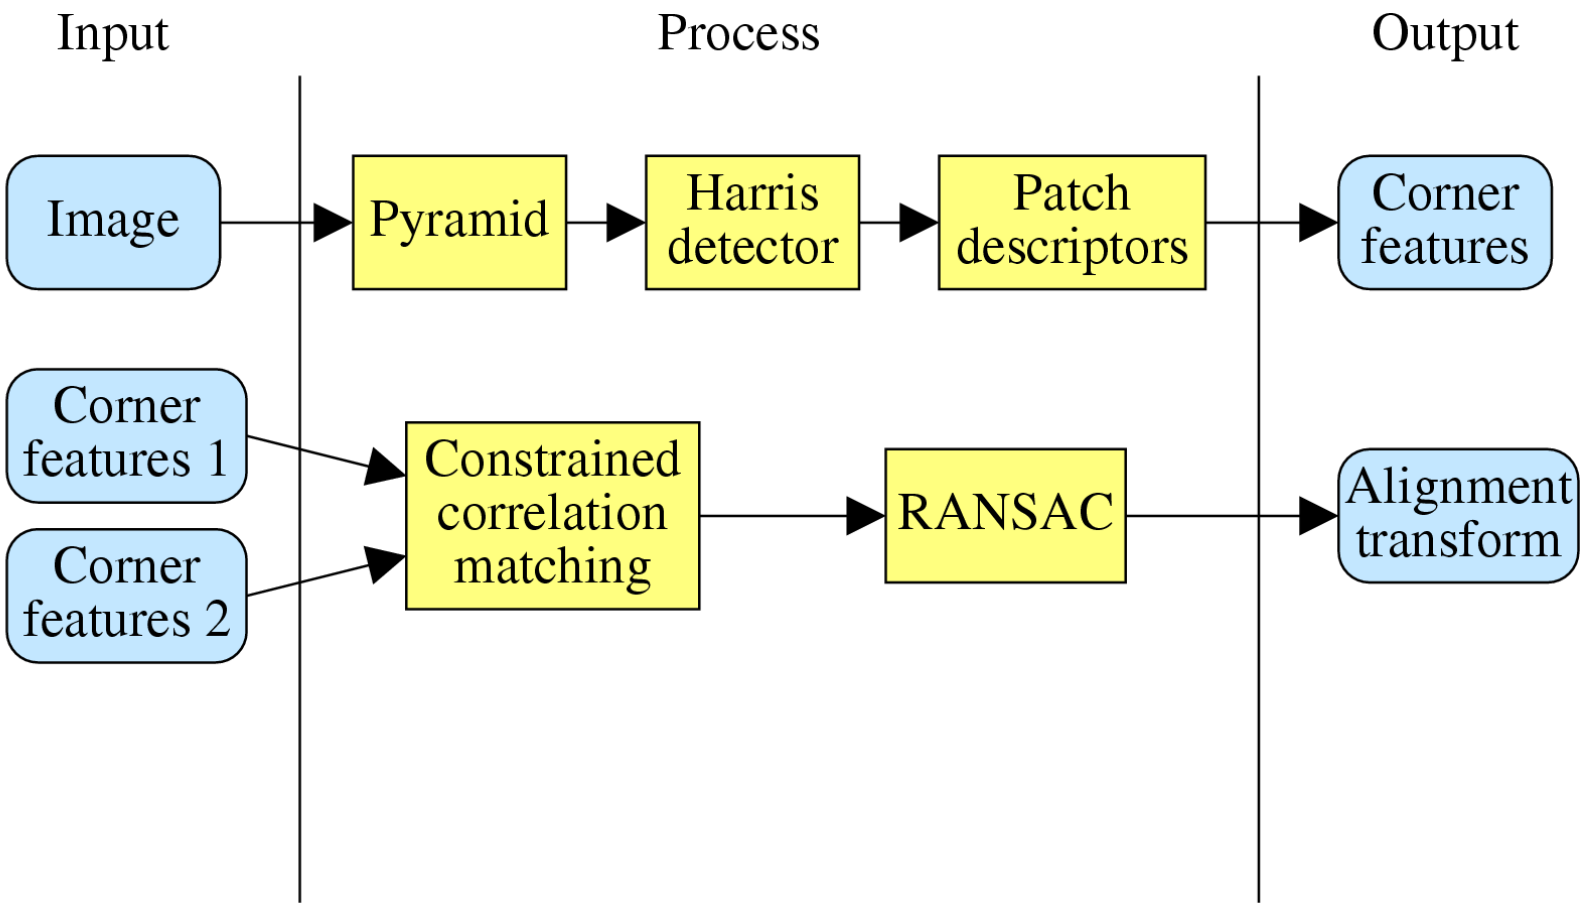

Matching

UNIVERSITY OF CAMBRIDGE

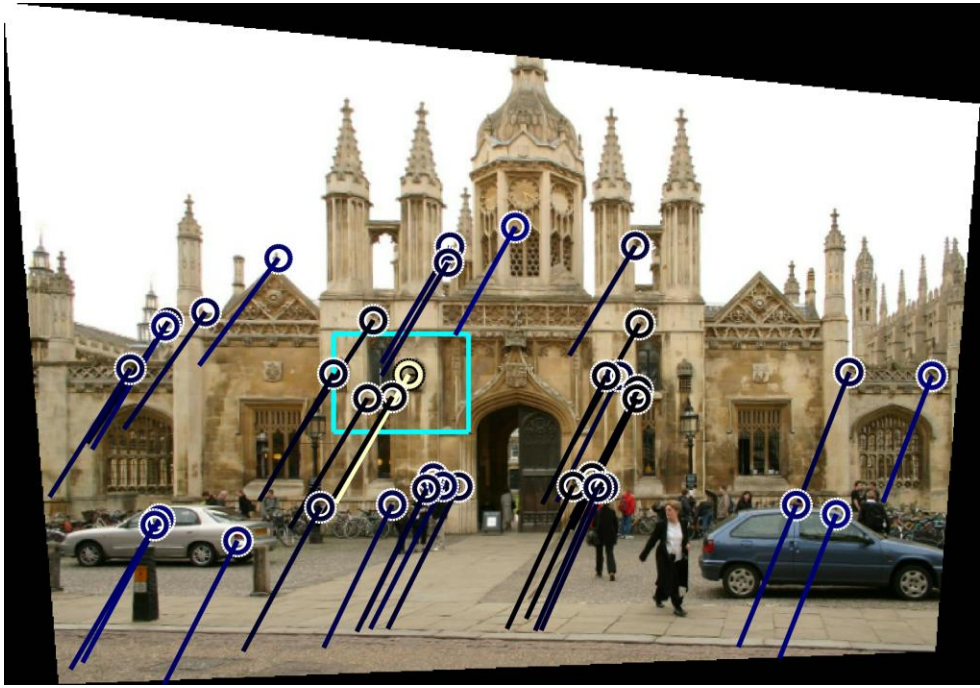

Matching

UNIVERSITY OF CAMBRIDGE

With only 2 params (s,t_x), can search rather than RANSAC.

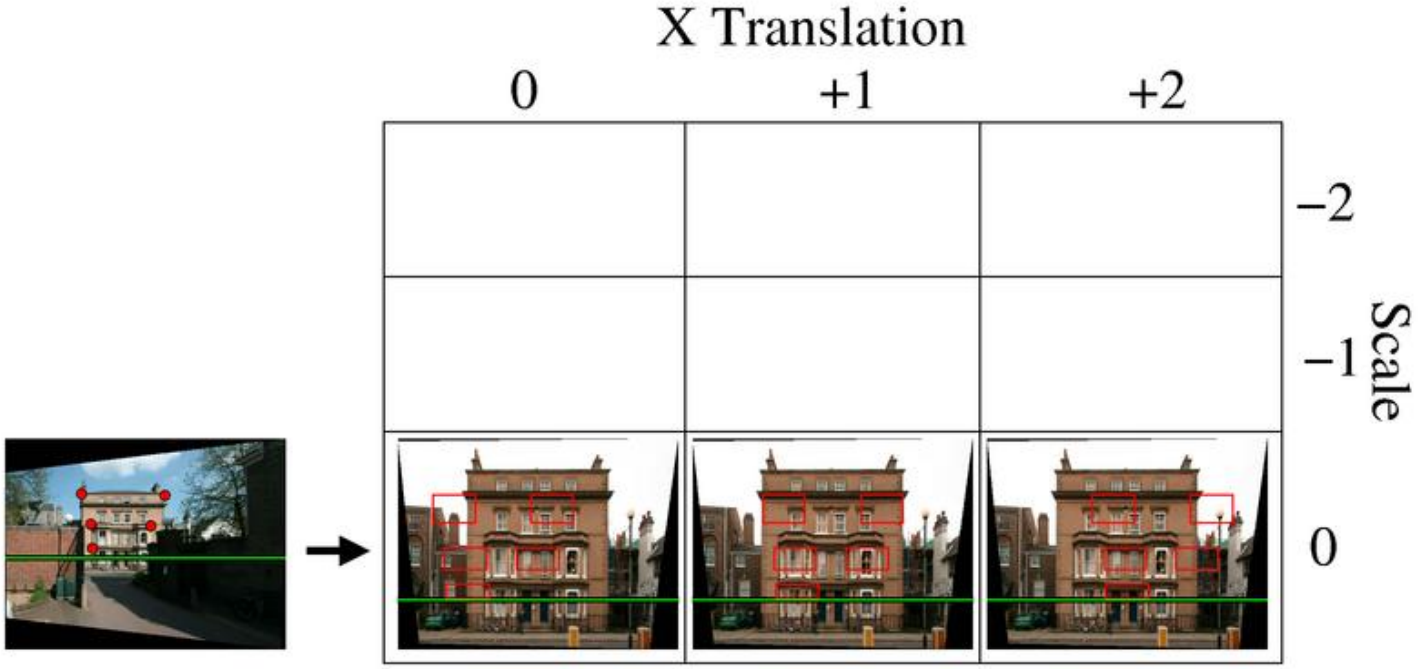

Matching

UNIVERSITY OF CAMBRIDGE

With only 2 params (s,t_x), can search rather than RANSAC.

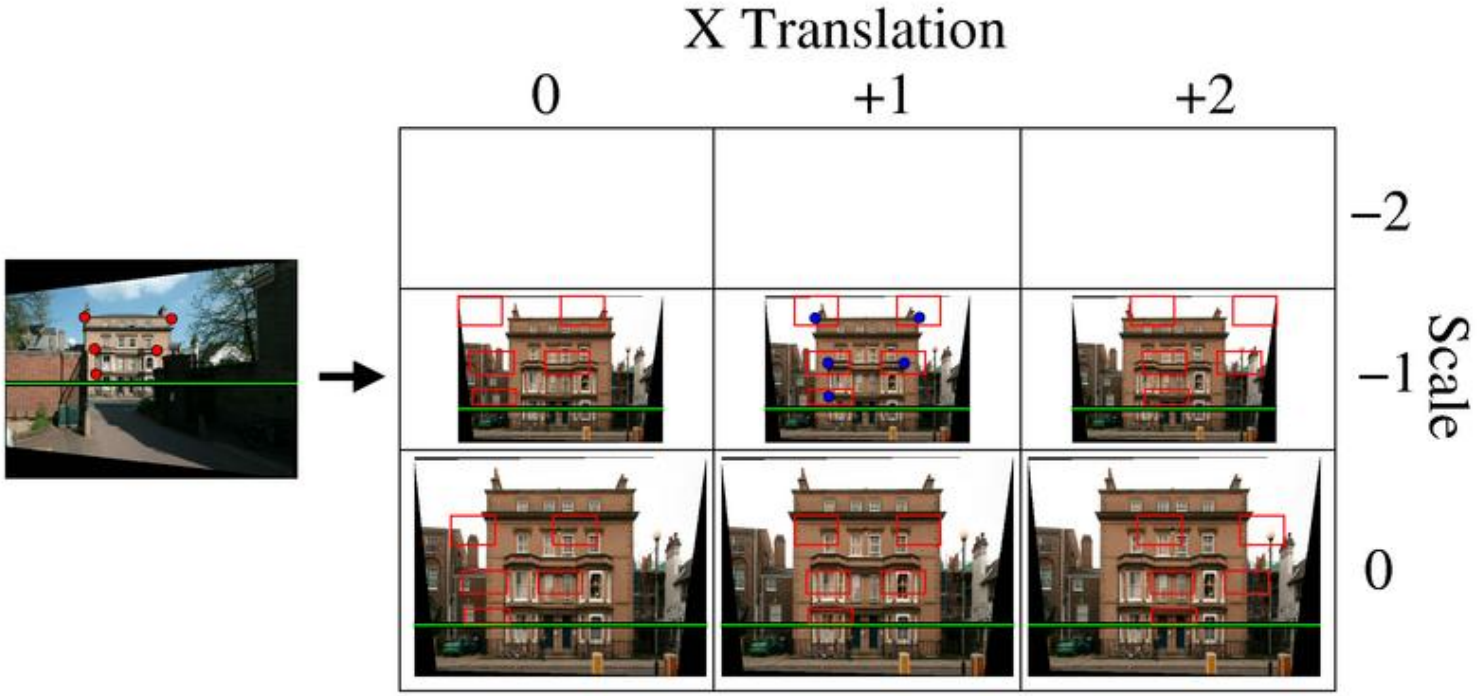

Matching

UNIVERSITY OF CAMBRIDGE

With only 2 params (s,t_x), can search rather than RANSAC.

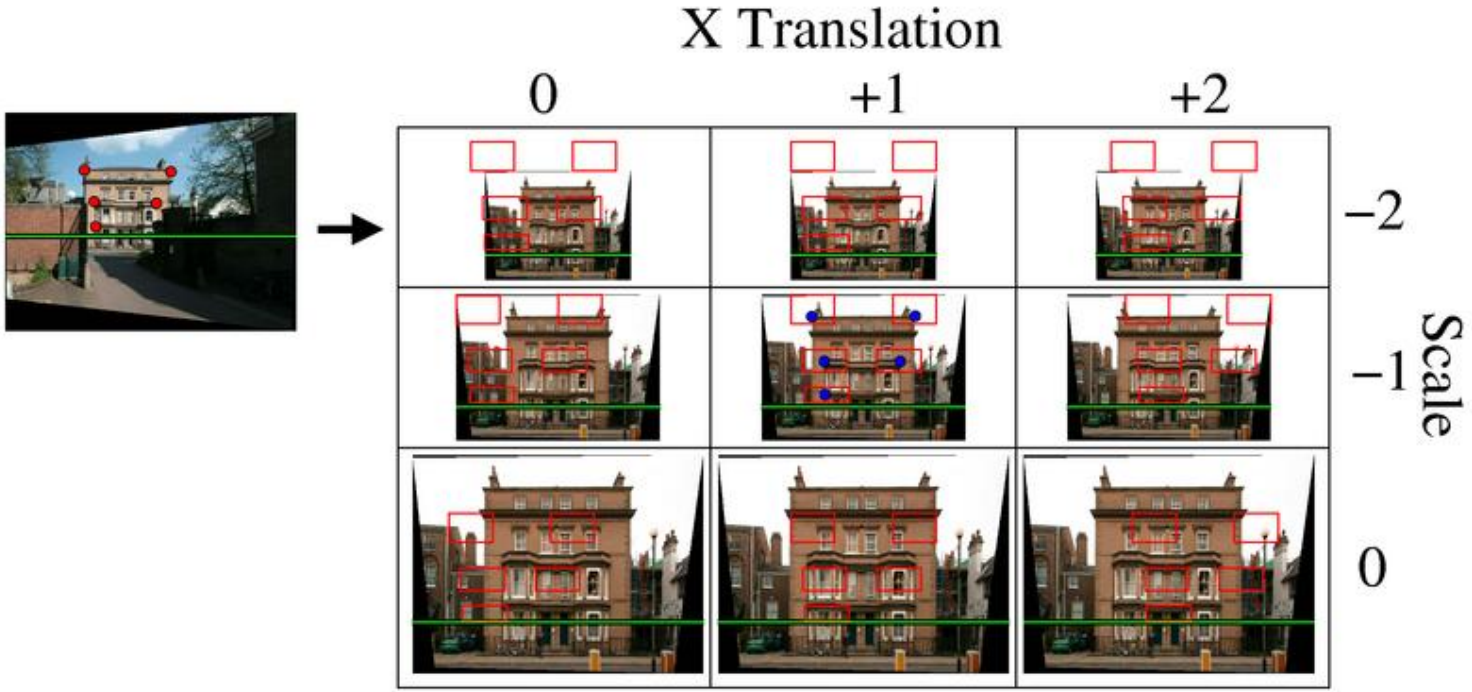

Examples over wide baselines

UNIVERSITY OF CAMBRIDGE

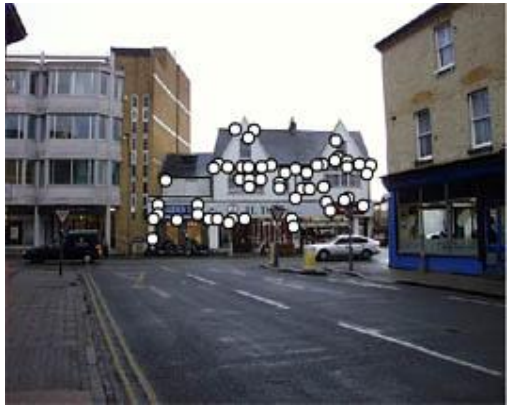

Summary of matching

UNIVERSITY OF CAMBRIDGE

Camera pose estimation - localisation

Localisation

UNIVERSITY OF CAMBRIDGE

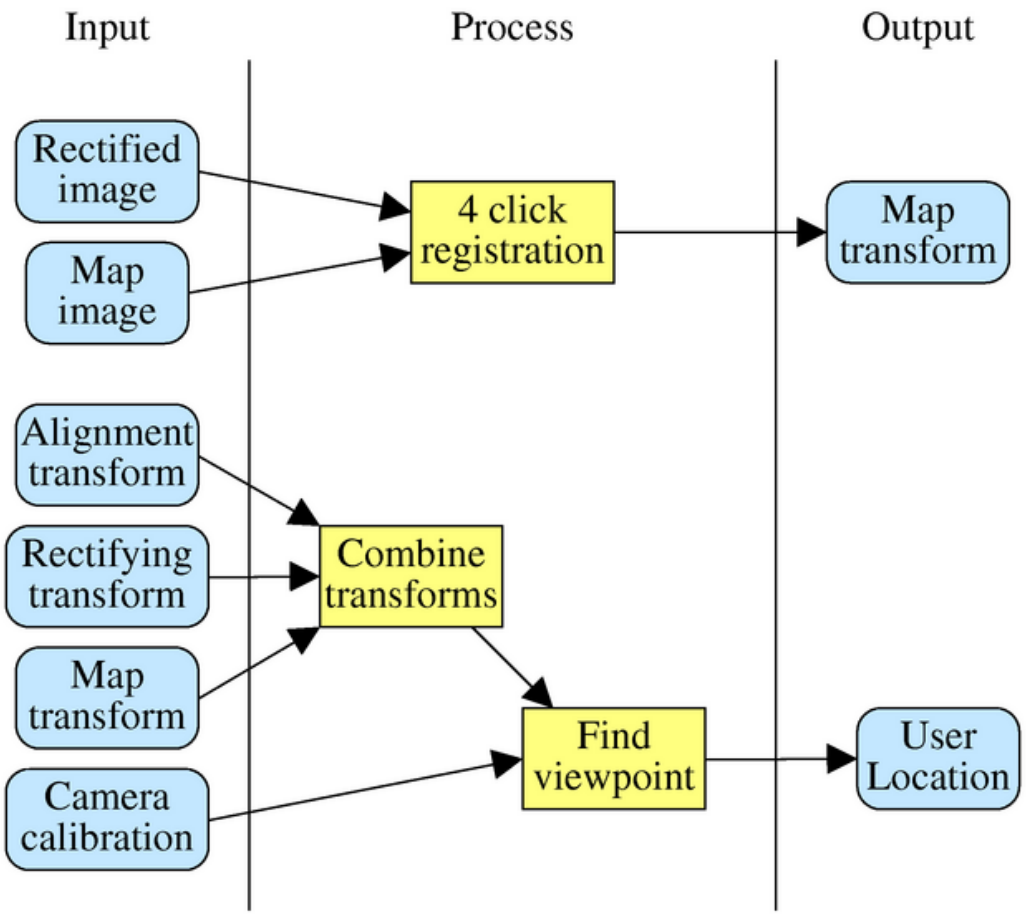

Register database view

UNIVERSITY OF CAMBRIDGE

First align database view to map

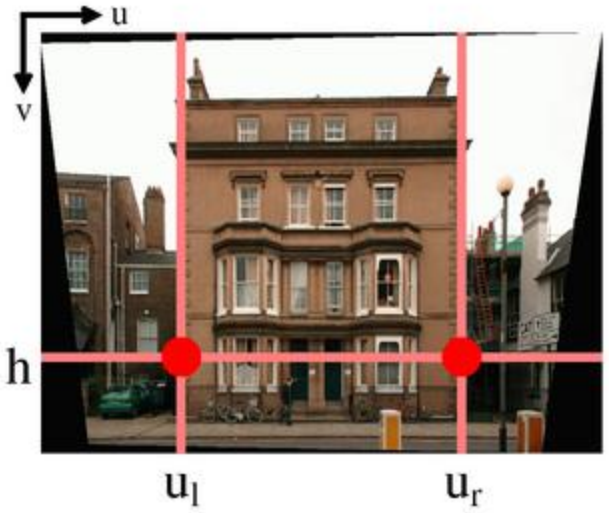

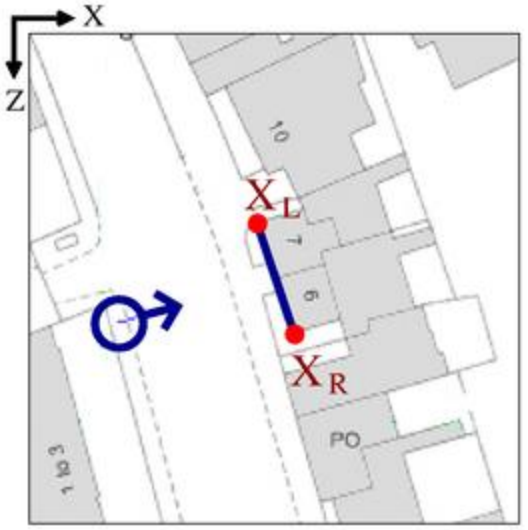

Localisation

UNIVERSITY OF CAMBRIDGE

Knowing the rectifying homography ($H_\perp$), the alignment (H_A), and the database view registration, can work backwards to find user:

Rectifying rotation $R_\perp$ gives the angle from perpendicular and focal length the distance to camera.

Localisation of query view

UNIVERSITY OF CAMBRIDGE

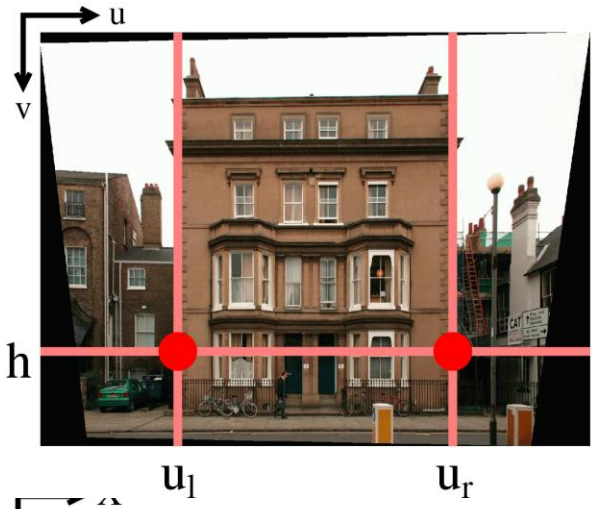

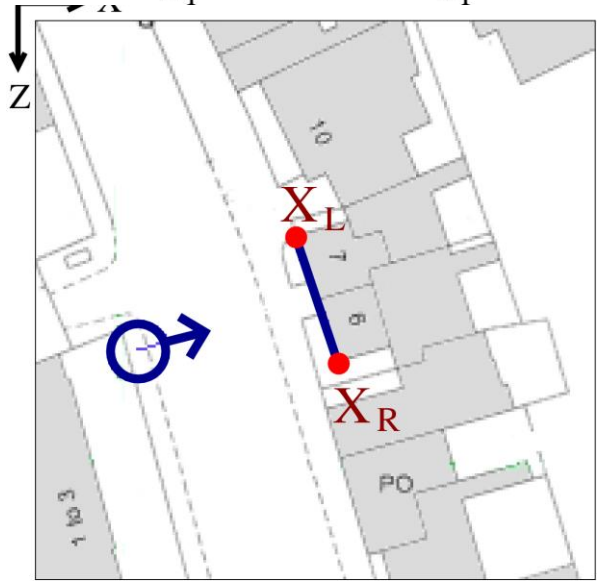

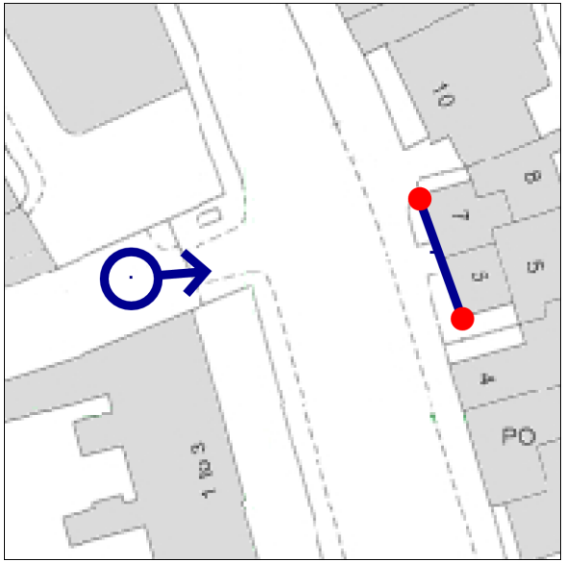

Localisation

UNIVERSITY OF CAMBRIDGE

Summary:

- Using geometric information generic matching is reduced to a 2 DOF search problem
- We are also able to find the camera (ie user) position and orientation

Evaluation

Evaluation

UNIVERSITY OF CAMBRIDGE

Evaluation

UNIVERSITY OF CAMBRIDGE

Image-based localisation

UNIVERSITY OF CAMBRIDGE

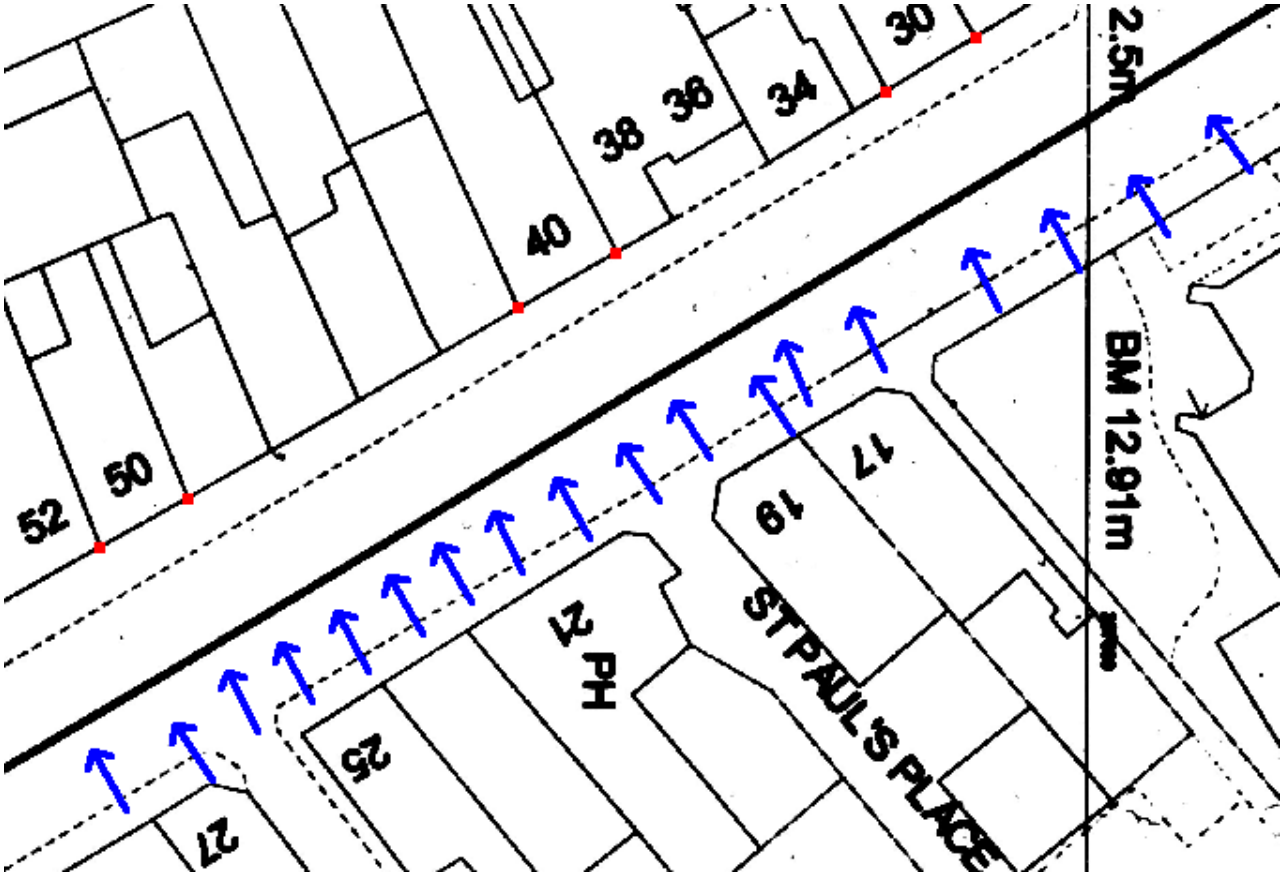

Image-based localisation

UNIVERSITY OF CAMBRIDGE

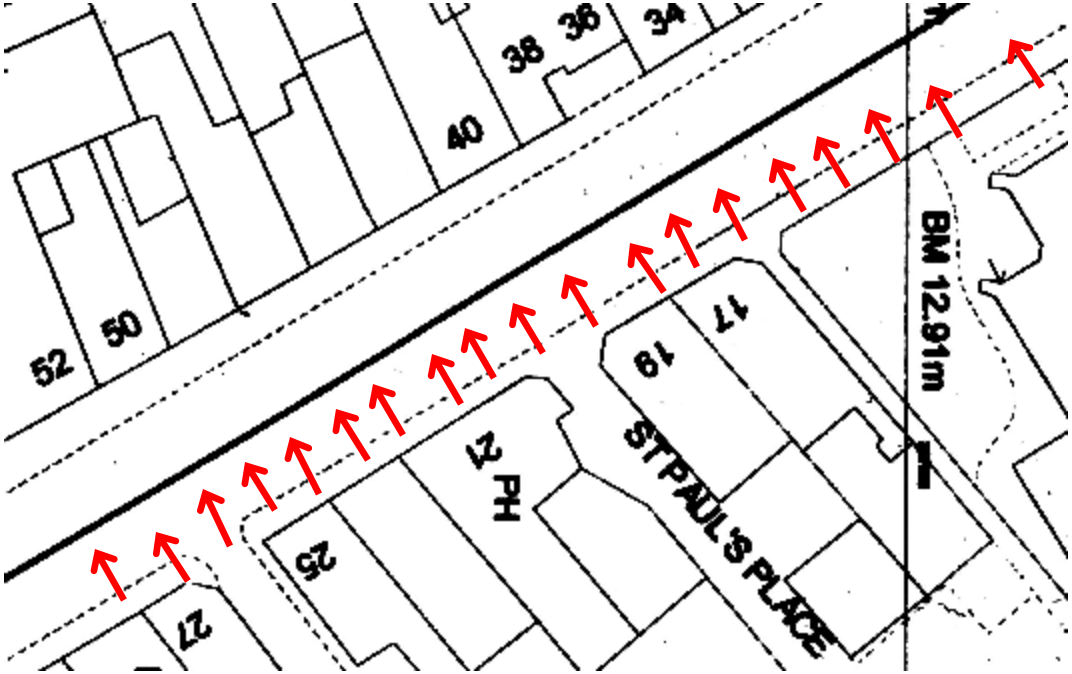

Image-based localisation

UNIVERSITY OF CAMBRIDGE

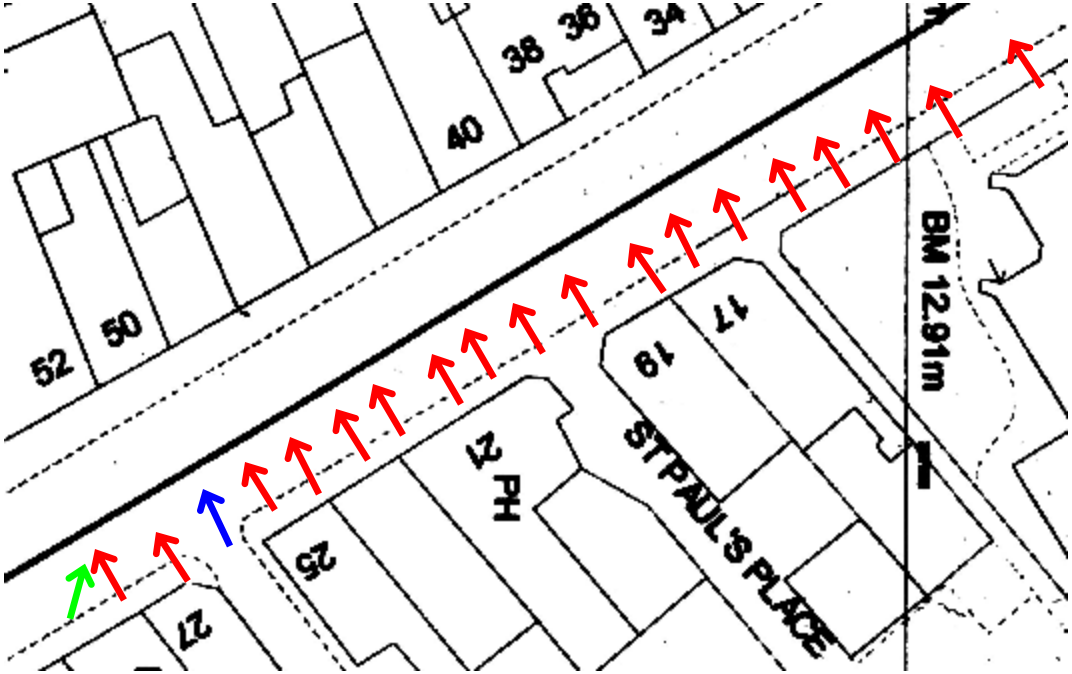

Image-based localisation

UNIVERSITY OF CAMBRIDGE

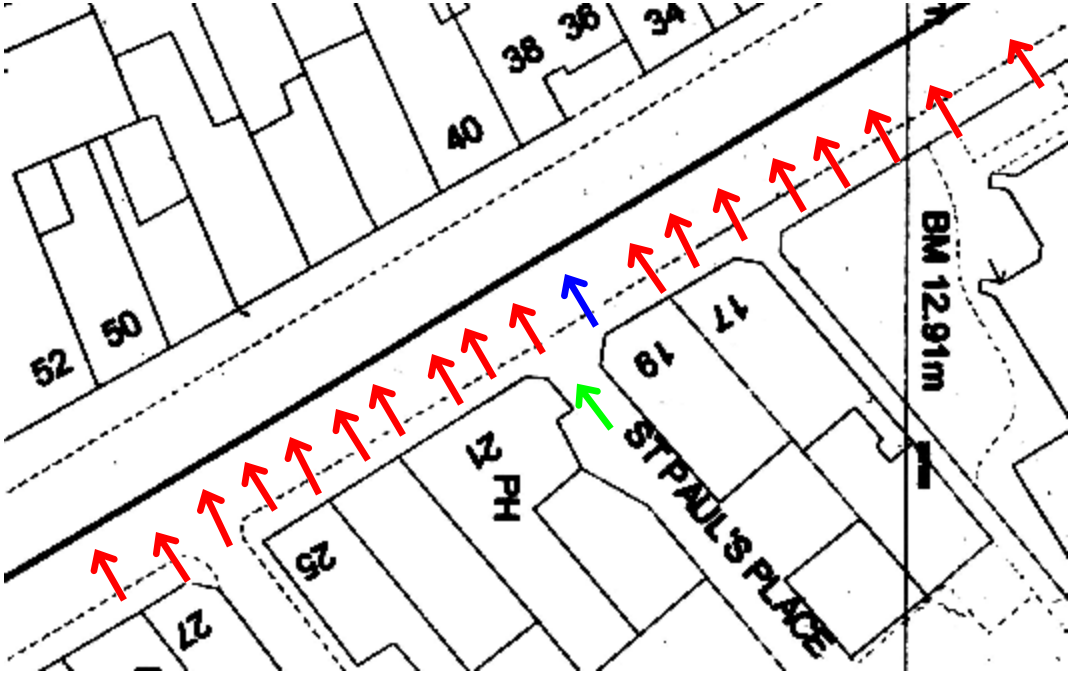

Image-based localisation

UNIVERSITY OF CAMBRIDGE

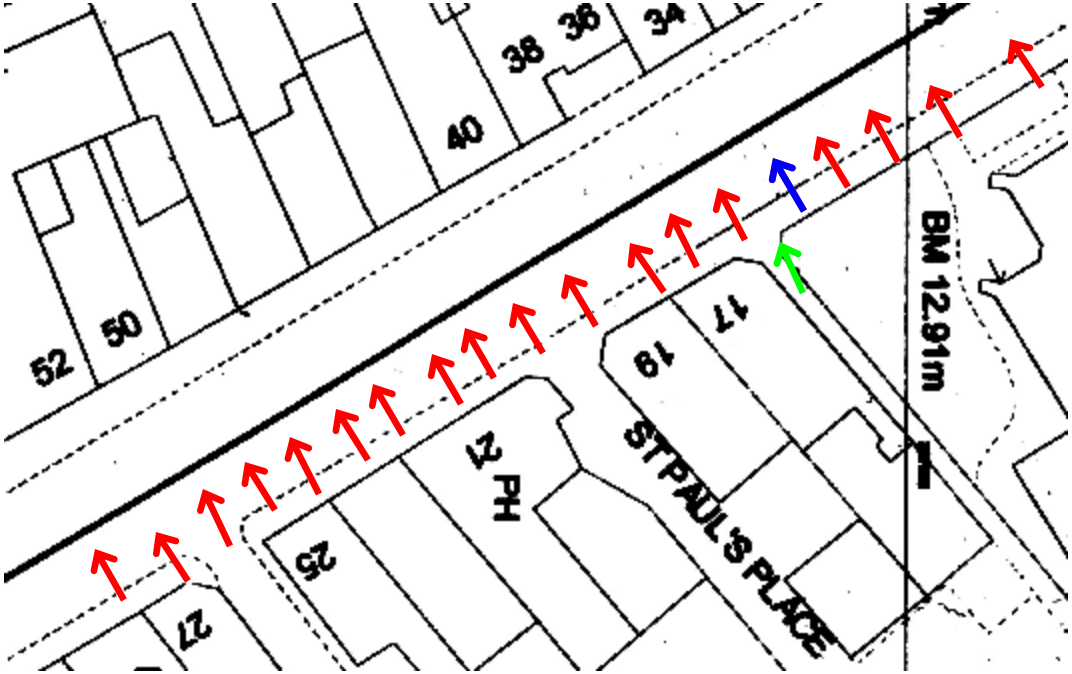

UNIVERSITY OF CAMBRIDGE

Recognizing pictures

Recognition of pictures

UNIVERSITY OF CAMBRIDGE

Conclusions

UNIVERSITY OF CAMBRIDGE

- New tools and a vibrant research community
- New application areas with mobile phones
  - Where am I?
  - What am I looking at?
- Technology is ripe for adaptation and exploitation