Multi-Spectral People Detection from UAVs

- Helen Flynn & Stephen Cameron

- Department of Computer Science,
- University of Oxford, UK

Oxford & UAVs

- OATS: 2005-8
- Aerial Tracking
  - mean-shift filter
- Camera Skills

- Learning target's motion

OATS videos


OATS Flight
Demonstraton
7th April, 2007

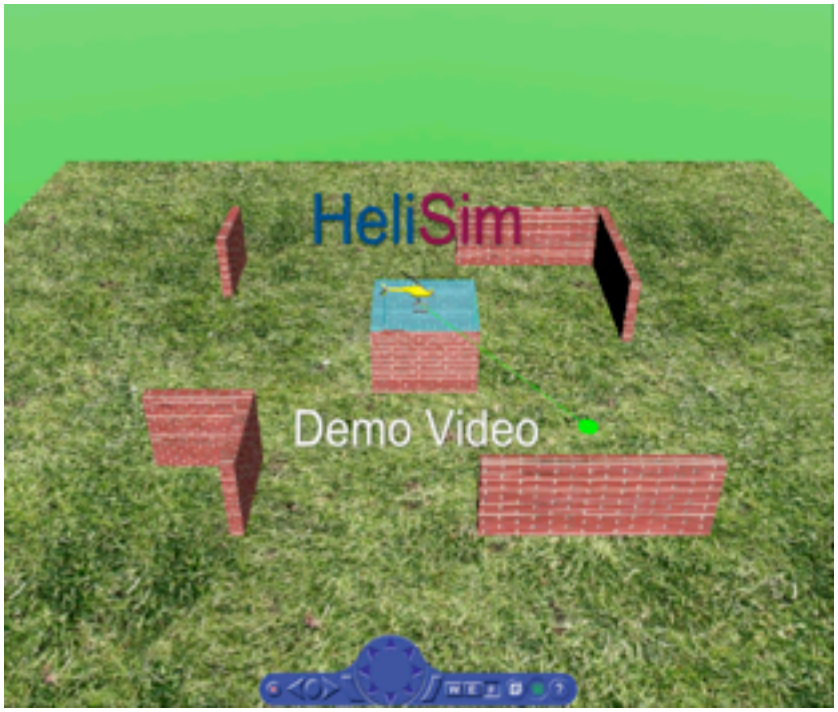
HeliSim
Demo Video

SUAAVE: 2008-2012+

- Joint work with University College London and University of Coleraine (Northern Ireland)
- Focus was on Networking issues with UAVs
  - radio modelling; comms planning; safety
- Example scenario: wilderness search-and-rescue
- My main interest was on planning and vision

Multi-tier Planning

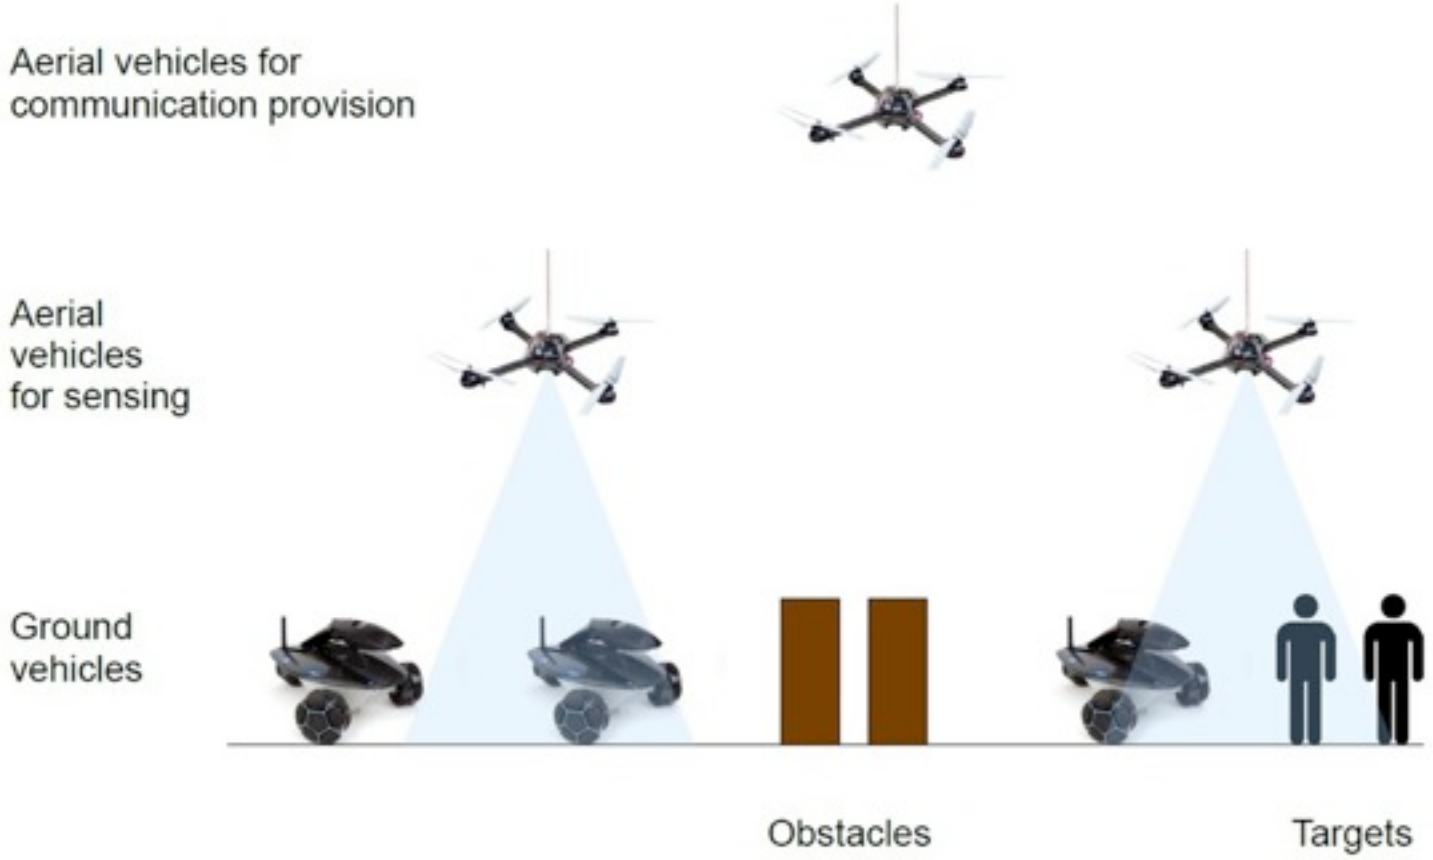

- Task allocation among a mixed fleet of UAVs

AIS: 2012–2015

- **New Foundational Structures for Engineering Verified Multi-UAVs**
- Now the emphasis is on bridging the gap between methods to analyse computer programs, and the real world
- Indoor UAVs used as a challenging test case
- First problem: having UAVs that we can trust indoors

UAV – Localisation

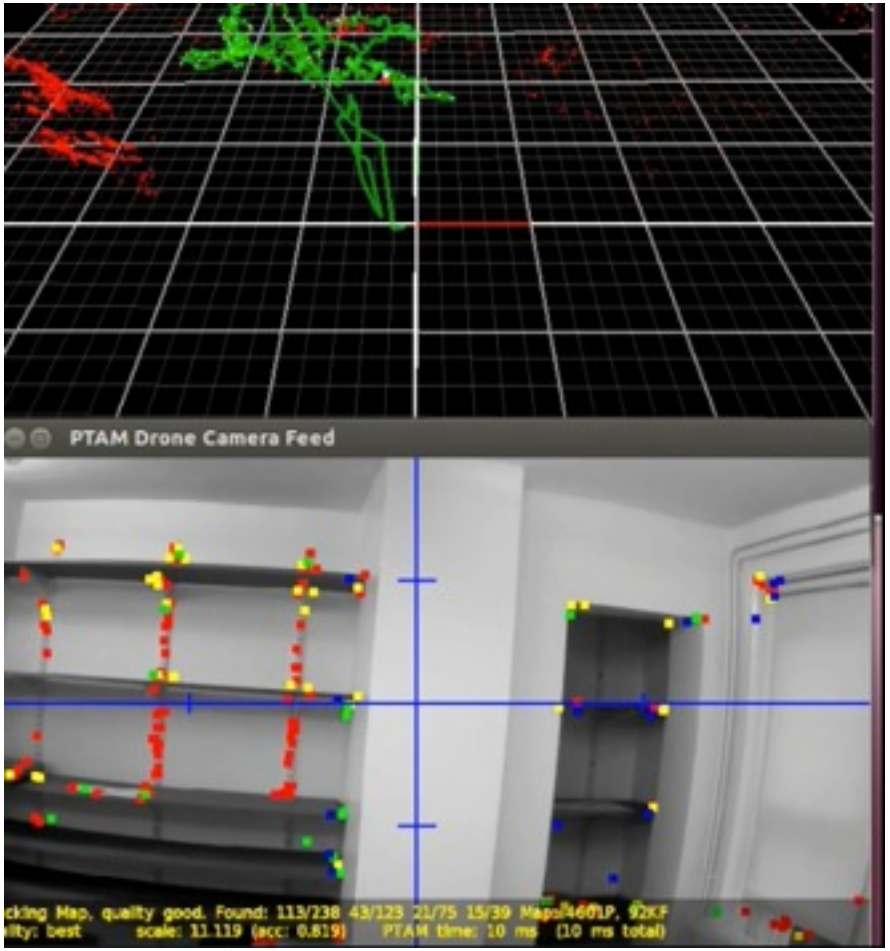

Green trajectory path - Localisation of the UAV

Red feature points – Generation and update of the map

Visual feed from the camera

See Ashu's Talk Tomorrow!

Multi-Spectral Detection

- An adjunct to SUAAVE – Helen's Thesis!
- Scenario: searching for people in a wilderness using a UAV
- People not trying to be seen
- Visually a tough problem
  - occlusions; odd angles; lighting
- Can we use infra-red (IR) to help?

Wilderness

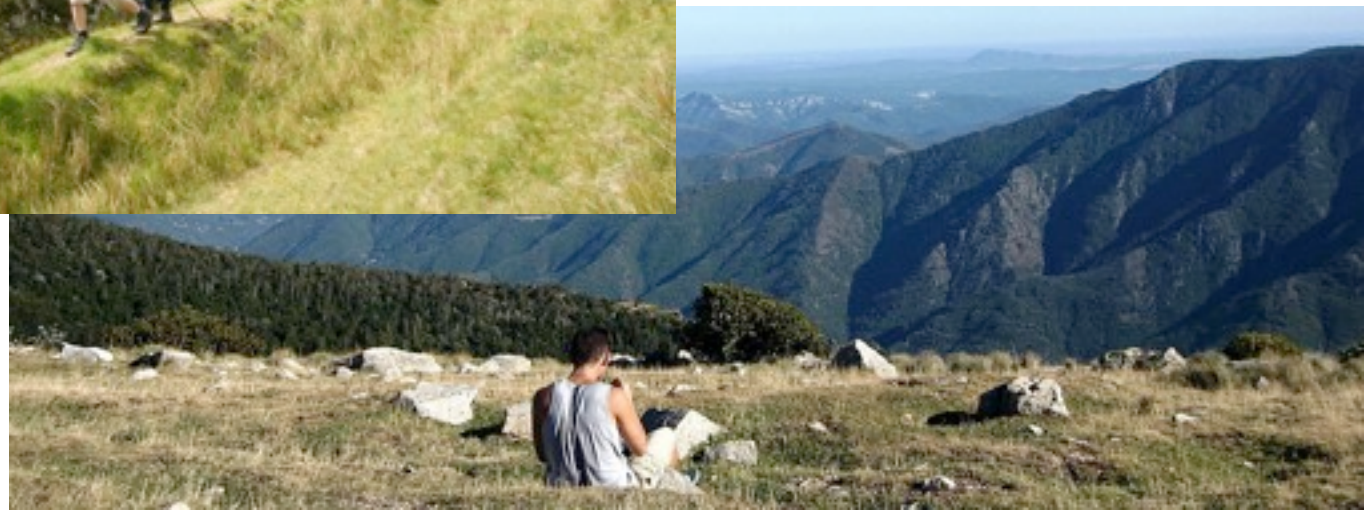

Infra-Red Cameras

- Standard CMOS cameras are IR sensitive
  - but poorly at body temperature – useless
- Cooled cameras very sensitive – and heavy
- New micro-bolometer cameras are ideal
  - lightweight; reasonably robust
  - reasonable (VGA+) resolution
- Expensive – but prices should drop (solid-state, economies of scale)

Thermoteknix Miricle Microcam

- 640x480 IR camera
- c. 150g with 18.8mm (germanium) lens
- Subject to export restrictions!
  - lower resolution units are not

Camera Set-up

- IR camera married with visual light 1280x960 camera with smaller field-of-view
- Concept is that 'hot-spots' in the IR camera are used as targets for close visual inspection

IR Images

Analysis of image histogram used to choose threshold point

People Detection

- Much work in Computer Vision on this
- Largely aimed at
  - CCTV images – people upright
  - Faces
- Our challenging images require a sophisticated detector; we are currently using the Felzenszwalb detector

Felzenszwalb Detector

- It's a part-based detector – looking for arms, head, etc, in proper proximity to each other
- Having detected parts, it combines their estimated positions in a Bayesian framework to give an overall score

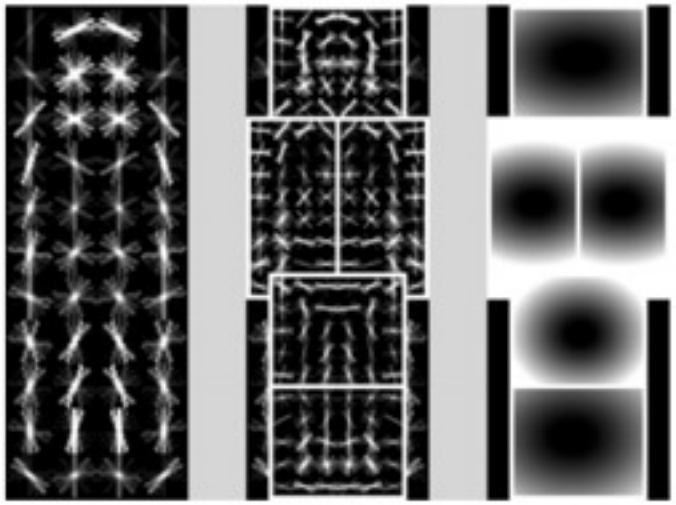

model

Complexity

- Most detectors – like Felzenszwalb – require the scanning of a low-level feature detector...
  - time proportional to area, or maybe using GPU
- ... plus the feature aggregation step, of time at least proportional to area
- We also require (8) different orientations, and multiple scales
- Very slow if applied to large images!

IR + Visual Pipeline

- We use a simple-minded processing of the IR image to generate regions of interest
  - intelligent thresholding + morphological analysis
- Generates regions that are passed to the Felzenszwalb detector
- Current experiments taken from vantage points - equipment expensive!

Examples

Examples

Issues

- So far, not tested in the air
- Have borrow an Ascending Technology Falcon for this purpose – next week?

Results

- Taking images at 15fps from approximately 20m
- Felzenszwalb on entire visual image takes 15-20s
- With IR pre-processing, this drops to 2s
  - with reasonable code optimisation and hardware improvements, should be 'real-time' within a few years

Precision-Recall Curve

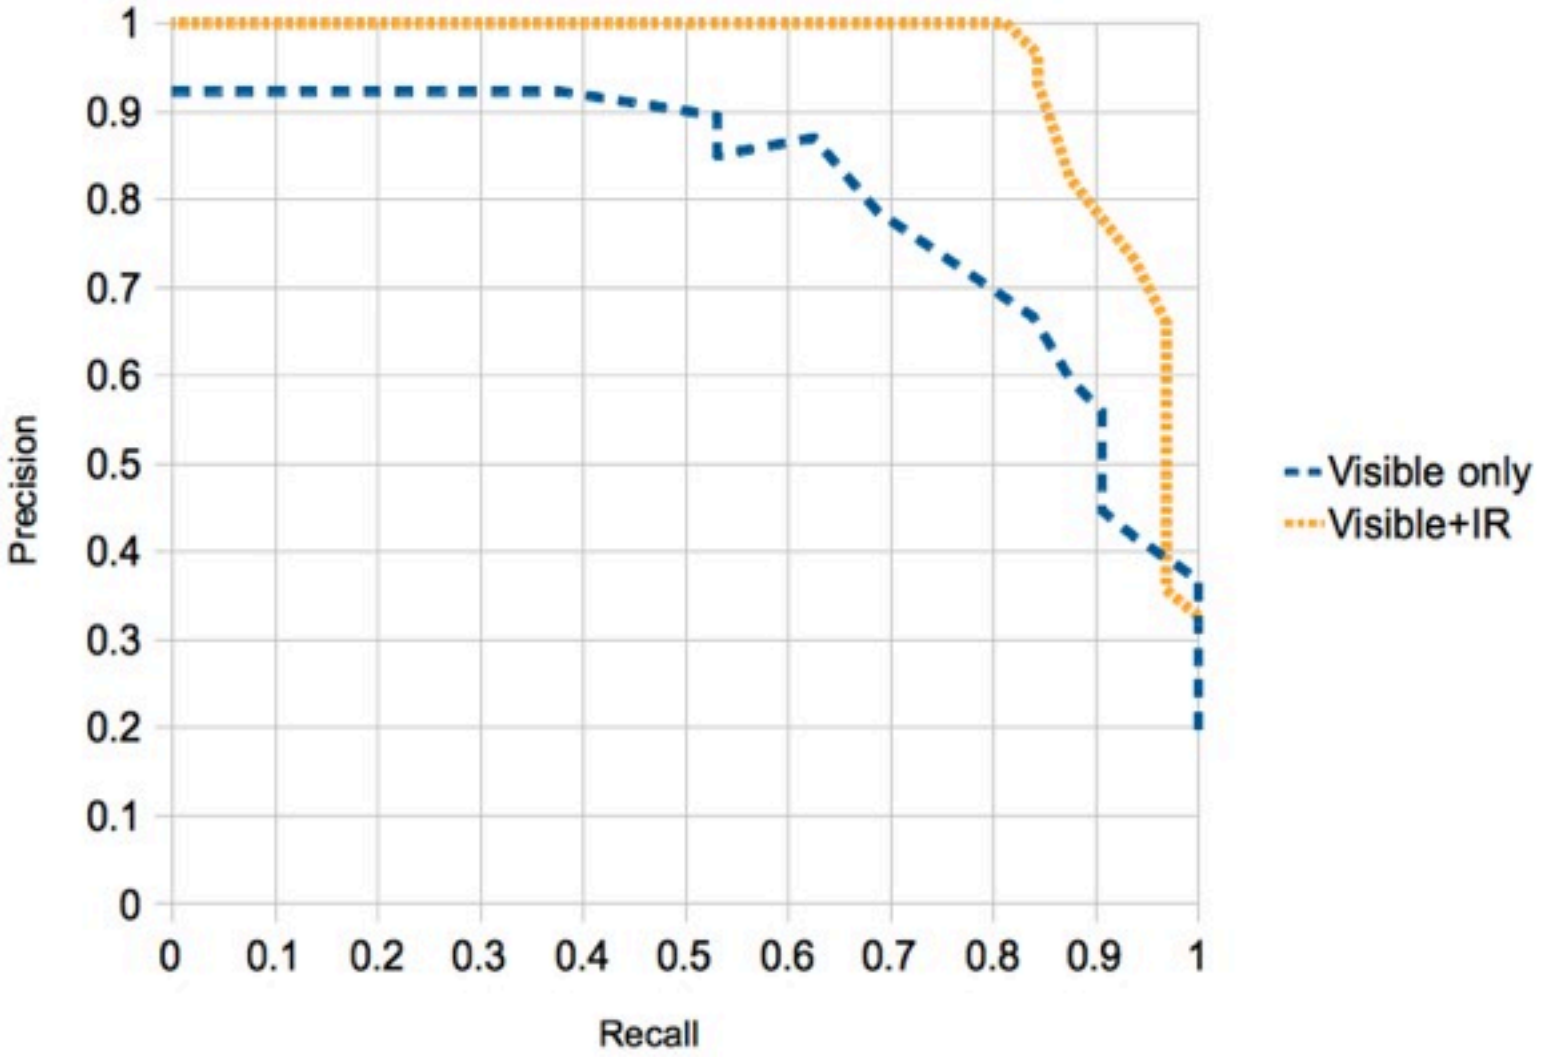

Expected Limitation

- IR analysis is easier with a high temperature contrast between the target and its surrounding
- Under sunny conditions against a hot background, contrast should be lower
- However then the visual contrast should be higher!

Extensions

- We don't yet use estimation of image size
  - should have this information from the height and bearing
- It is possible to run people detectors on the IR image
  - plan to try closing the loop:
  - IR blobs -> Visual detector -> IR detector
- We could do more for tracking

Conclusions

- IR cameras are expensive, and require care
  - but provide useful information, and the price should drop
- Already our results suggest that the approach is useful
- Some extensions in mind to ‘close the loop’
- Should be possible to work in ‘real-time’ soon