Visual mapping using learned structural priors

Osian Haines – José Martínez-Carranza – Andrew Calway
Department of Computer Science, University of Bristol

University of BRISTOL

Introduction

We investigate a new approach to **vision based mapping**, using **single image plane detection** to derive strong priors for initialisation of map features. Our detection algorithm **learns from a large training set**, to relate appearance to structure, and we use its detection in a real-time visual odometry system to **quickly add large structures** to the map as it is built.

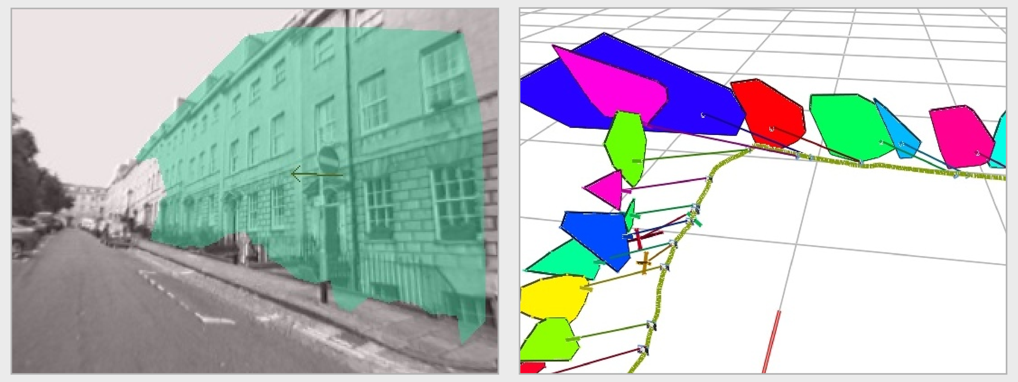

Using plane detection for rapid initialisation of 3D structure

By detecting planes from a single image, using our detecion algorithm [3], we can save time compared to conventional methods which need to infer geometry from parallax. This initialises features in a plane-based visual odometry system [4], showing that having a prior on plane location allows faster convergence, leading to more structured and accurate maps.

Planar visual odometry

As a baseline visual odometry (VO) system, we use the inverse depth plane parameterisation (IDPP) VO of Martínez-Carranza and Calway [4].

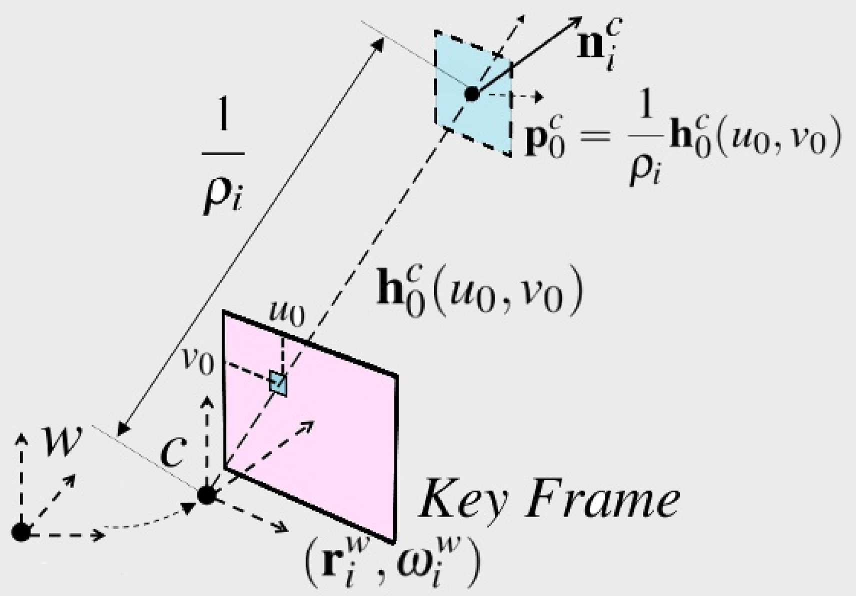

Inverse Depth Plane Parameterisation

- Undelayed initialisation of planar features
- Planes are grown according to geometric information
- Unified plane and point representation – planes collapse into points when appropriate
- Features are represented $\mathbf{m}_i = [\mathbf{r}_i, \boldsymbol{\omega}_i, \mathbf{n}_i, \rho_i]$, being position, orientation, plane normal and **inverse depth**
- Robust estimation using 1-point RANSAC [2]

email: haines@cs.bris.ac.uk
web: www.cs.bris.ac.uk/~haines/

Single image plane detection

The core idea of our VO system is that having a means of finding structure from a single image would be very useful. As such, we use our previously developed single-image plane detection algorithm (SIPD) [3], which can find planes and estimate their orientation using only a single image.

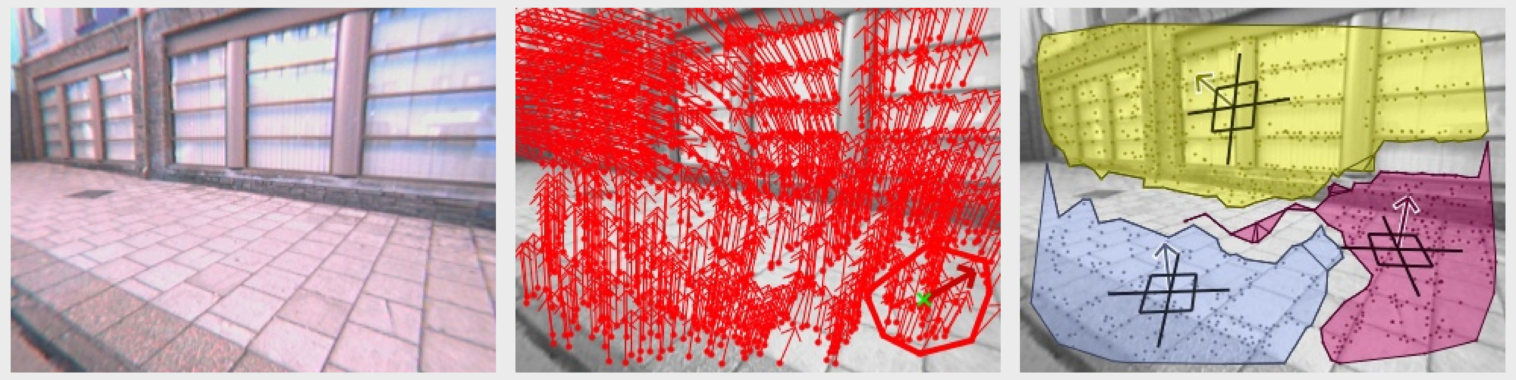

The process of single-image plane detection

- Based on classifying local regions as planar/non-planar
- Can estimate 3D orientation of planes
- Trained on a large amount of labelled training data
- Uses grandient and colour features in spatiograms [1]
- Relevance vector machine for classification and regression
- Segmentation of salient points with an MRF

Examples of plane detection

VO with plane priors

The key part of this work is the combination of IDPP with SIPD, to allow instant initialisation of planar structures, without having to grow from all possible seed points.

- The plane detector is run repeatedly in a separate thread
- It runs on the keyframes stored by IDPP, which have an estimated camera pose
- Planes are only initialised on detected plane centroids
- A mask created from the detected planes guides plane growth
- Planes are good enough to leave in the map after the camera has moved on, to create a rough 3D model

To ICRA paper...

Results

- Our preliminary experiments show this to be able to initialise good planes much more quickly than VO
- Increases localisation accuracy over the baseline system
- Creates fewer planes, which are larger and in more appropriate locations (avoids growing in non-planar regions)
- We observed an increased mapping speed, up to 60 fps

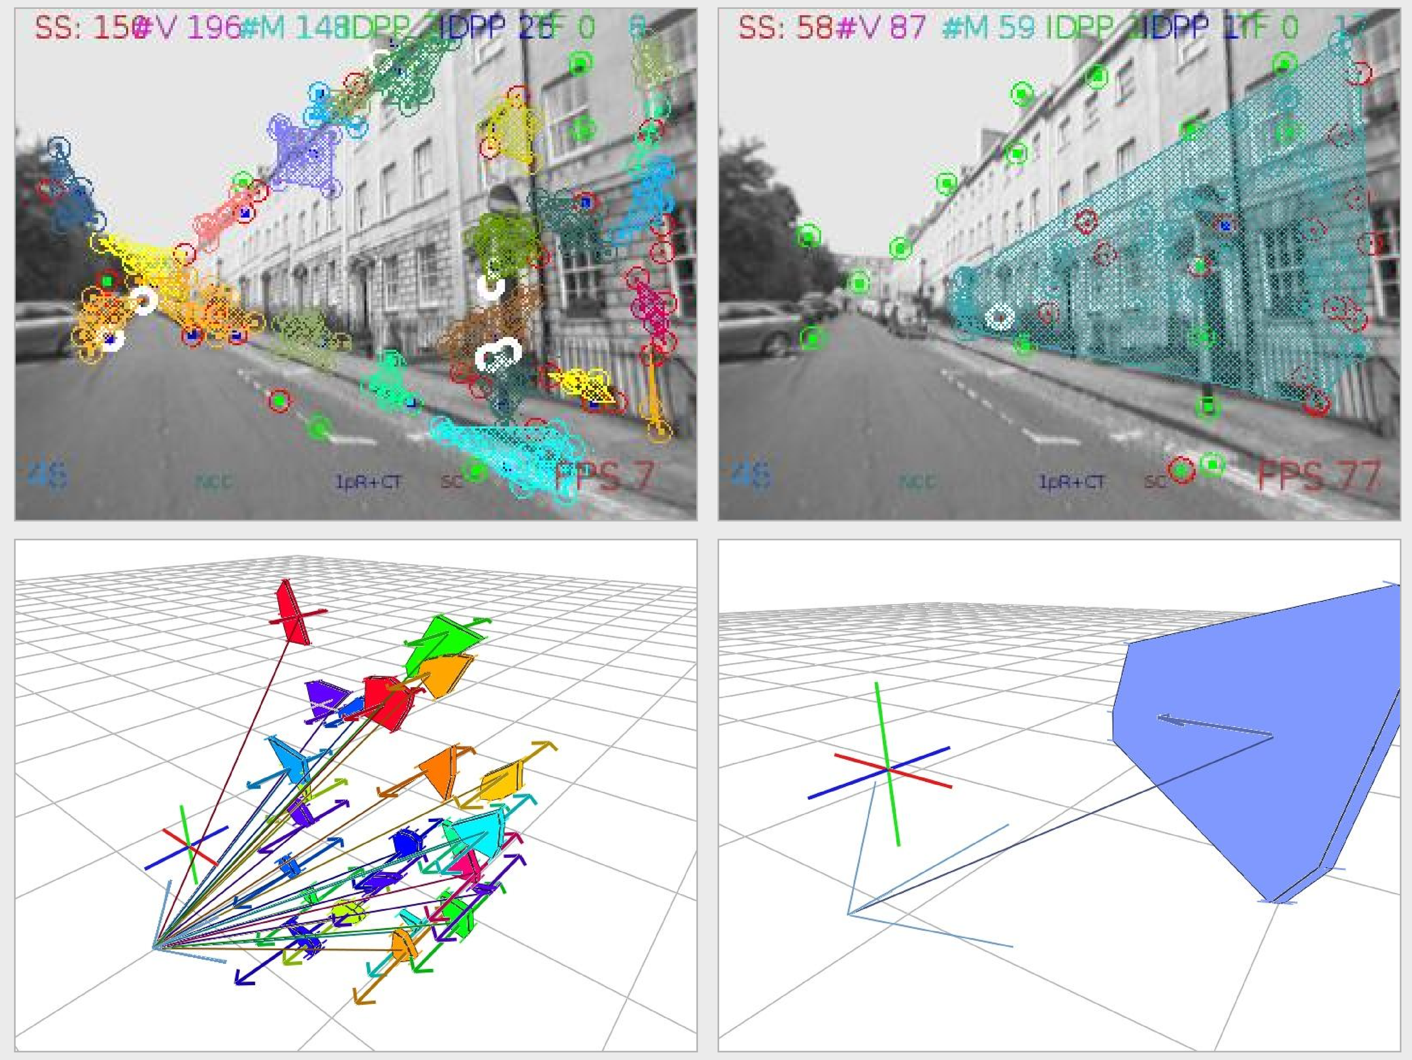

PDVO (right) allows large planes to be initialised quickly – after only 40 frames, it has a much better estimate of the structure

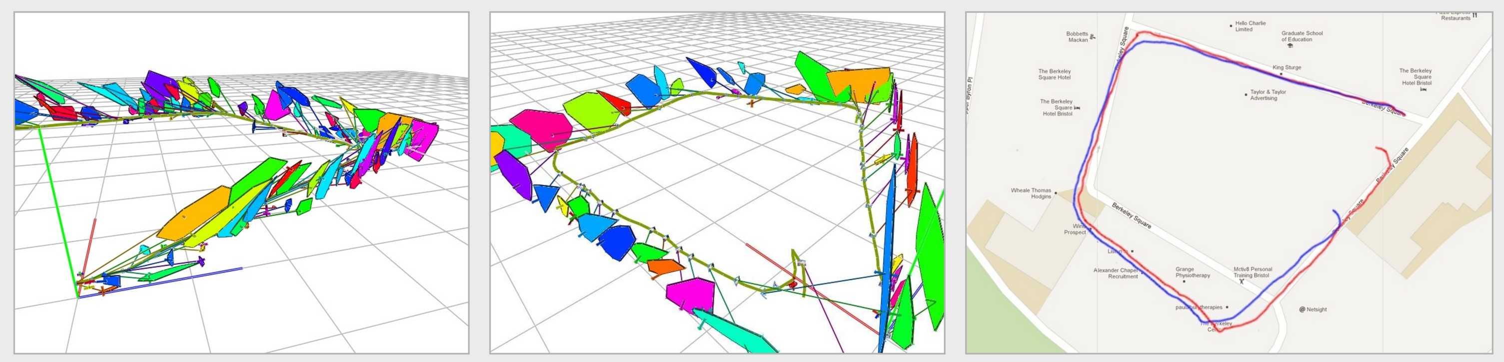

Visualisation, from original IDPP (left) and PDVO (centre); and a comparison of trajectories

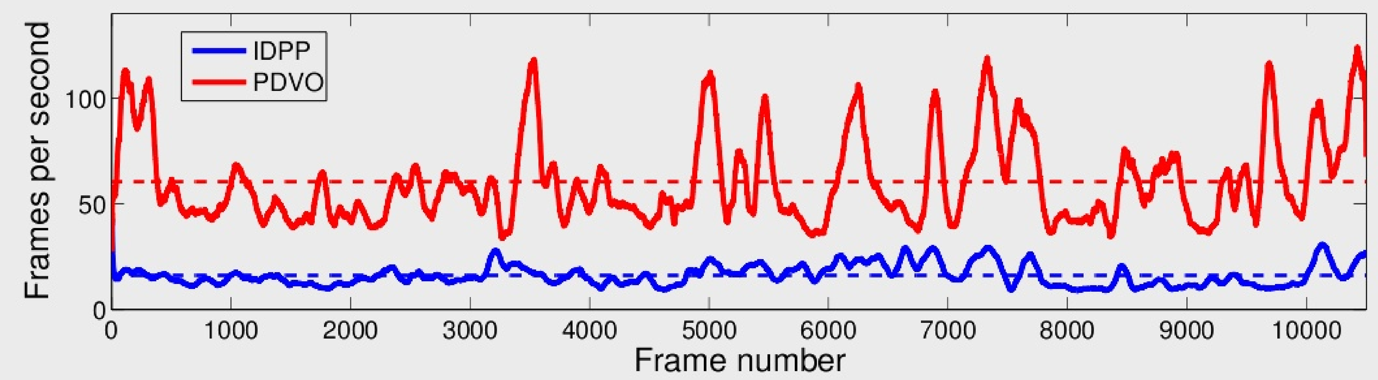

Time (frames per second) for each of the methods (with mean)

References

[1] S.T. Birchfield and S. Rangarajan. Spatiograms versus histograms for region -based tracking. In CVPR, 2005.

[2] J. Civera, O.G. Grasa, A.J. Davison, and JMM Montiel. 1-Point RANSAC for EKF-based structure from motion. In ICIRS, 2009.

[3] O. Haines and A. Calway. Detecting planes and estimating their orientation from a single image. In BMVC, 2012.

[4] J. Martınez-Carranza and A. Calway. Efficient visual odometry using a structure-driven temporal map. In ICRA, 2012.

Detecting planes and estimating their orientation from a single image (BMVC 2012)

Osian Haines and Andrew Calway – Department of Computer Science, University of Bristol

University of BRISTOL

Introduction

We propose an algorithm to detect planes in a single image, in outdoor urban evnrionments. We use machine learning methods, to learn from a labelled training set the relationship between appearance and structure, and show that this can work in a variety of scenes.

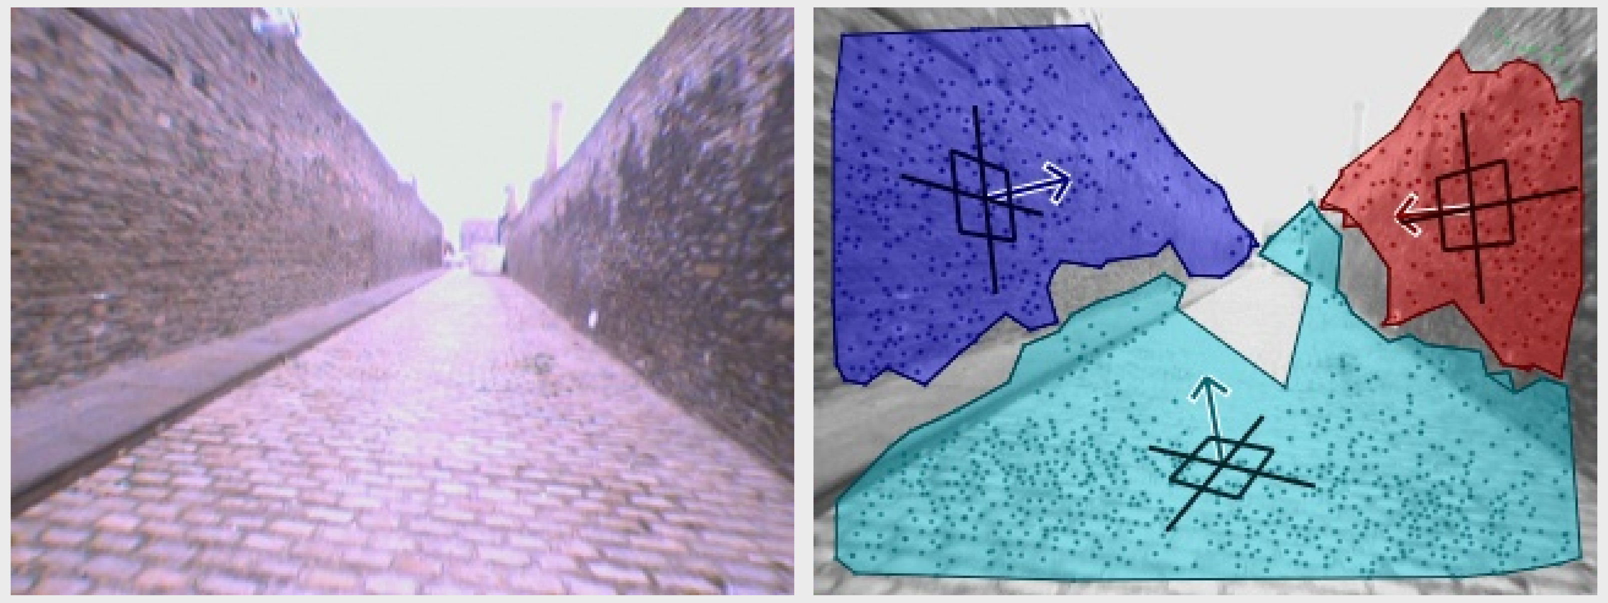

- Aim to detect planes in a single image
- Does not rely on geometric information
- Inspired by human perception – based on prior experience
- Achieves this by learning from examples in a training set

Overview

- Represent salient points using gradient and colour features
- Use plane estimation [2] on multiple overlapping windows
- Robust estimate of planarity and orientation at each point
- Segment into distinct planar regions (MRF)

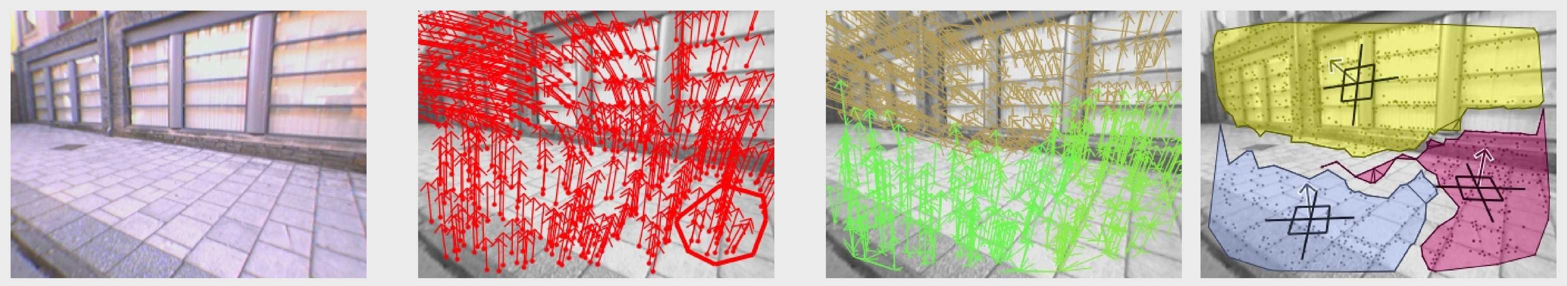

Input | Local plane estimate | Segmentation | Detected planes

References

- S.T. Birchfield and S. Rangarajan. Spatiograms versus histograms for region -based tracking. In CVPR, 2005.
- O. Haines and A. Calway. Estimating planar structure in single images by learning from examples. In ICPRAM, 2012.
- M.E. Tipping. Sparse bayesian learning and the relevance vector machine. Journal of Machine Learning Research, 1, 2001.

email: haines@cs.bris.ac.uk

Plane Estimation

- For one image region: classify (plane or non-plane), and estimate 3D orientation
- Represent regions using bag of words
- Dimensionality reduction using latent semantic analysis
- Encode spatial distribution information with spatiograms [1]
- Relevance Vector Machine [3] for classification and regression

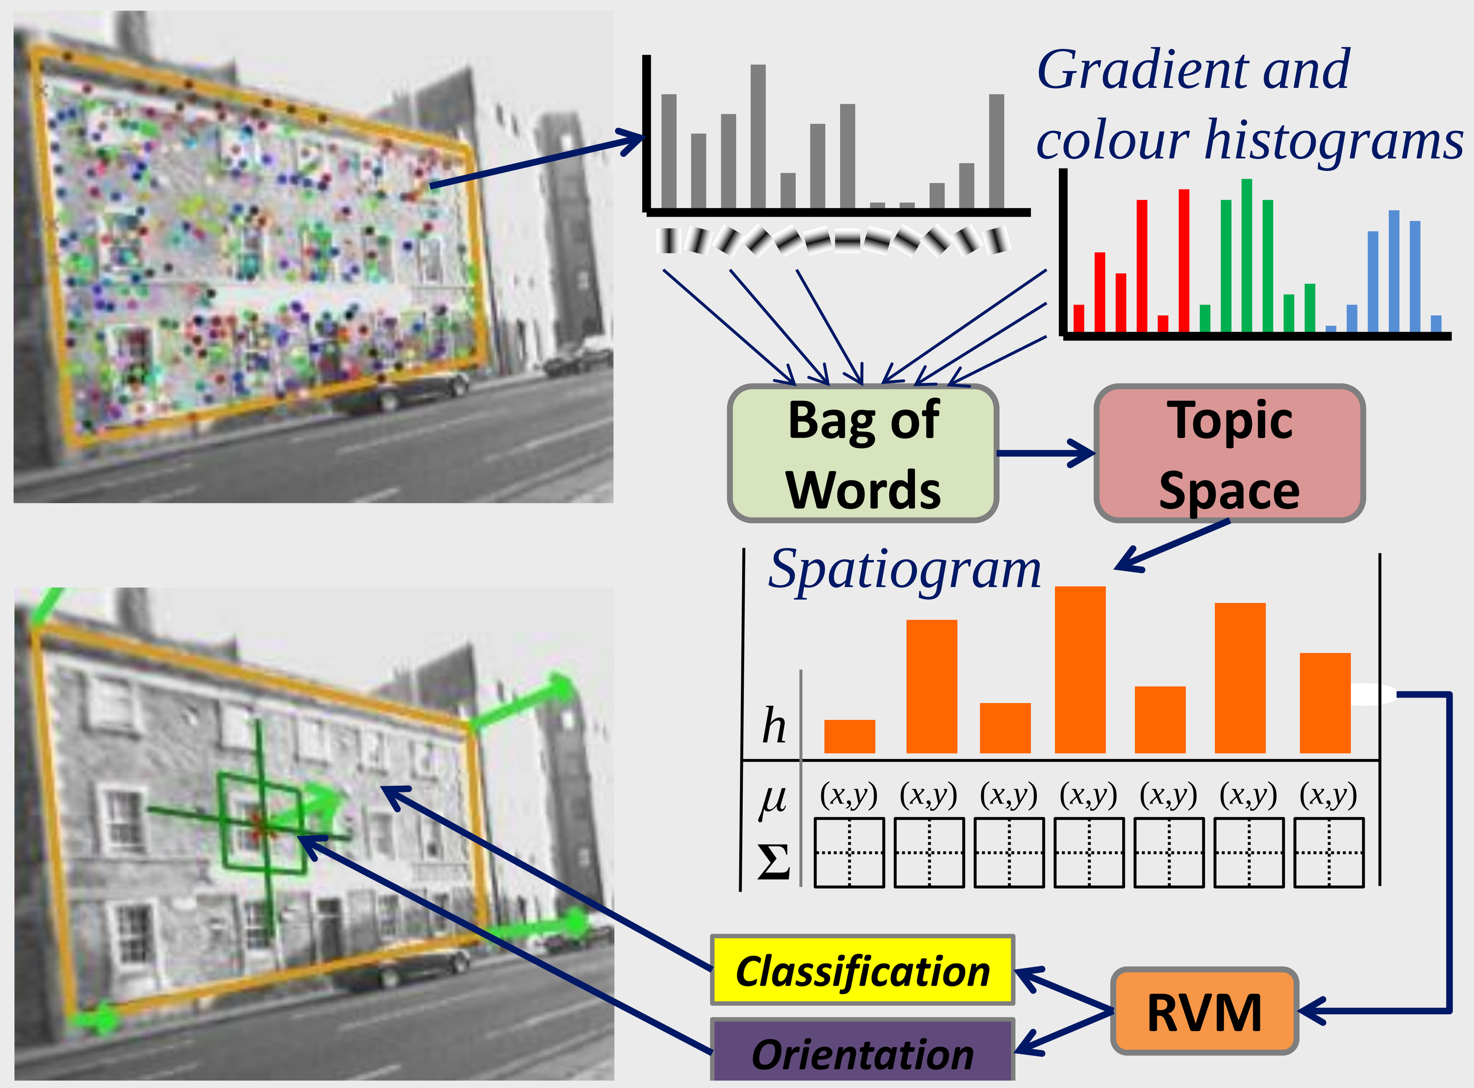

Local Plane Estimate

- Plane estimation (see bottom left) works for one region
- But cannot detect planes
- Apply plane estimation to overlapping windows

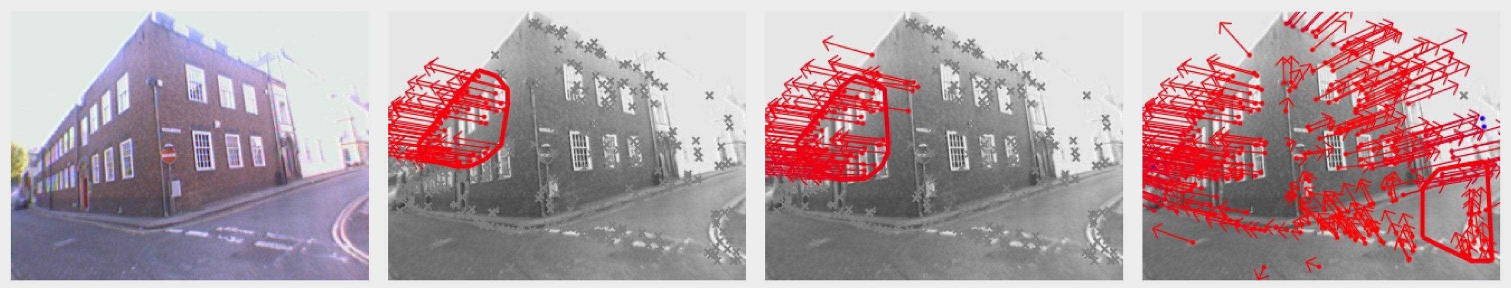

- Obtain local plane estimate for each point
- Not a plane detection – no boundaries
- Used for segmentation

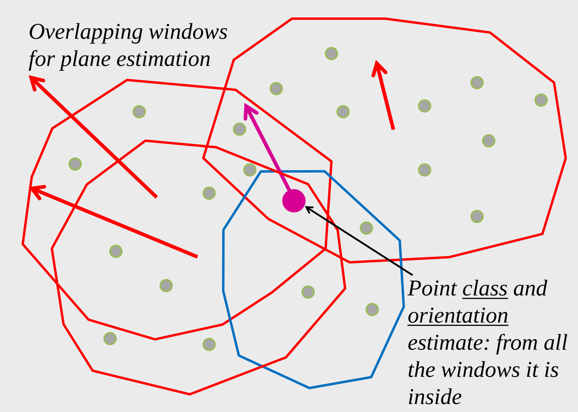

web: www.cs.bris.ac.uk/~haines/

Segmentation

1. Segment planes from non-planes

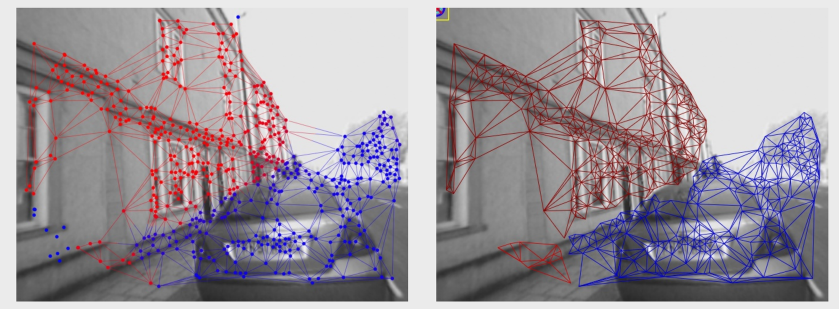

2. Find orientation of dominant planes (mean shift)
3. Segment planes by orientation

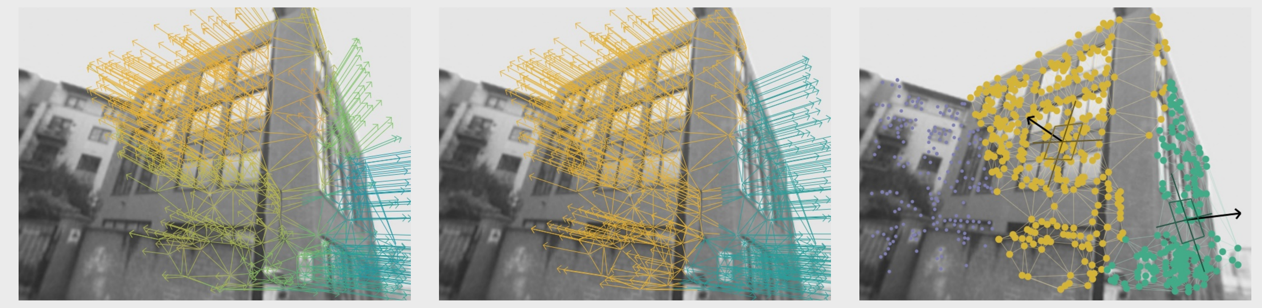

4. Apply plane estimation to plane segments: final plane detection

Results

- Tested on independent data set
- Point classification accuracy of **88%**, orientation error of **18**°

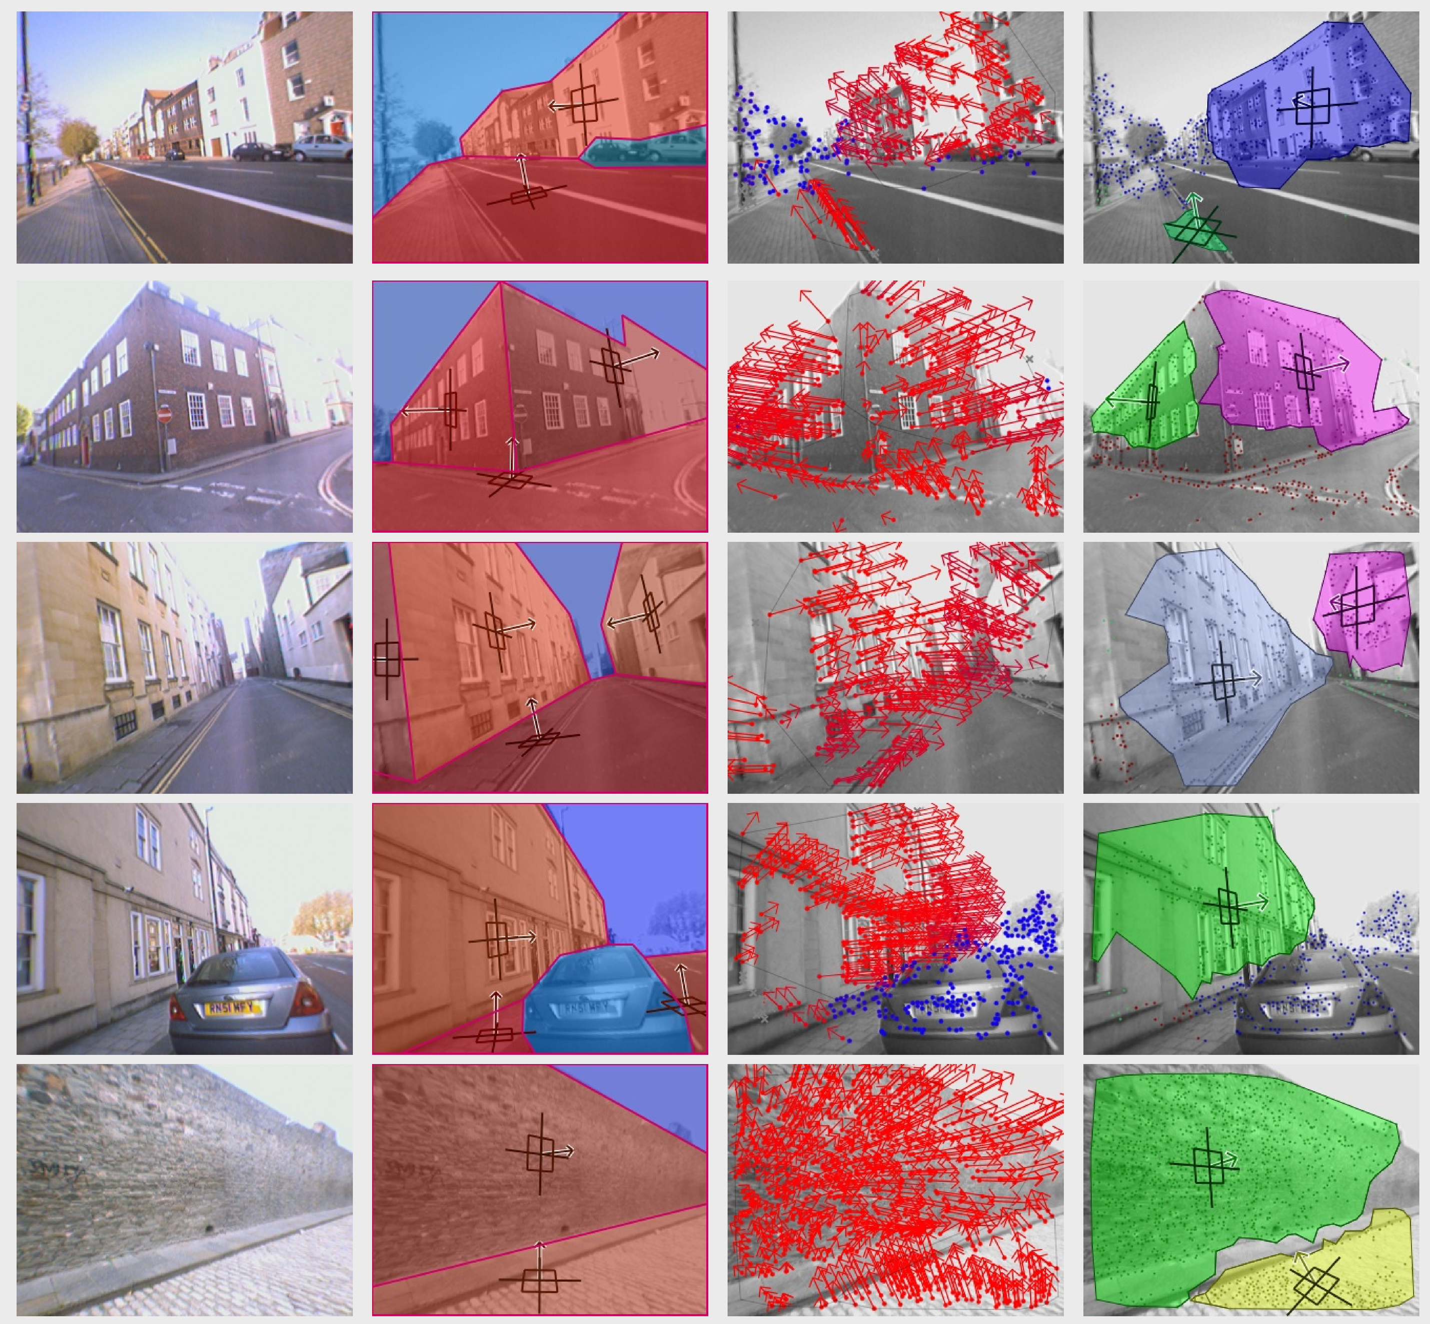

Example results (input, ground truth, local plane estimate, plane detection)

More on Plane Detection...