

Single-camera SLAM using the SceneLib library

Paul Smith

Robotics Research Group Talk, Nov 2005



Outline

- 1 Introduction
 - The Camera as a Position Sensor
 - Visual SLAM
- 2 Davison's MonoSLAM
 - Overview and Nomenclature
 - Extended Kalman Filter
 - Automatic Map Management
 - Performance
- 3 The SceneLib Libraries
 - Introduction
 - The Scene Library
 - The MonoSLAM Library
 - Applications using SceneLib
- 4 Final Thoughts
 - Final Thoughts



The Camera as a Position Sensor



Aim

- Use a camera as a position sensor

Challenges

- Monocular (no depth)
- Unconstrained
- High acceleration & large rotations
- Usually want **real-time** localisation



Off-line Processing

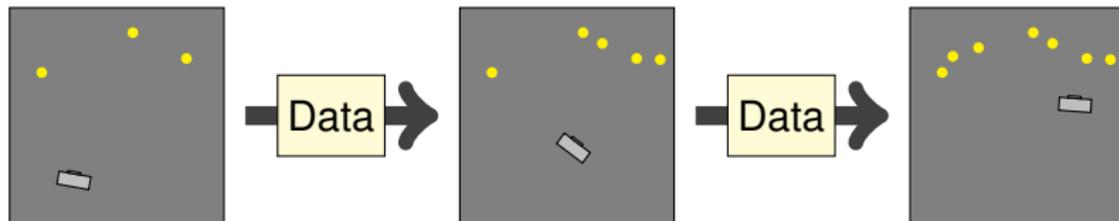


Structure from motion

- Typical **Computer vision** approach
- Bundle adjustment over a long sequence
- Applied to post-production, 3D model reconstruction.



Sequential Real-time Processing



Simultaneous Localisation and Mapping (SLAM)

- Typical **robotics** approach.
- Building a long-term map by propagating and correcting uncertainty
- Mostly used in simplified 2D environments with specialised sensors such as laser range-finders.



Classic Approaches to Visual SLAM

Davison, ICCV 2003

- Traditional SLAM approach (Extended Kalman Filter)
- Maintains full camera and feature covariance
- Limited to Gaussian uncertainty only

Nistér, ICCV 2003

- Structure-from-motion approach (Preemptive RANSAC)
- Frame-to-frame motion only
- Drift: No repeatable localisation



Classic Approaches to Visual SLAM

Davison, ICCV 2003

- Traditional SLAM approach (Extended Kalman Filter)
- Maintains full camera and feature covariance
- Limited to Gaussian uncertainty only

Nistér, ICCV 2003

- Structure-from-motion approach (Preemptive RANSAC)
- Frame-to-frame motion only
- Drift: No repeatable localisation



Recent Approaches to Visual SLAM

Pupilli & Calway, BMVC 2005

- Traditional SLAM approach (Particle Filter)
- Greater robustness: handles multi-modal cases
- New features not rigorously initialised

Eade & Drummond, 2006

- FastSLAM approach (Particle Filter/Kalman Filter)
- Particle per camera hypothesis, Kalman filter for features
- Allows larger maps: update $O(M \log N)$ instead of $O(N^2)$



Recent Approaches to Visual SLAM

Pupilli & Calway, BMVC 2005

- Traditional SLAM approach (Particle Filter)
- Greater robustness: handles multi-modal cases
- New features not rigorously initialised

Eade & Drummond, 2006

- FastSLAM approach (Particle Filter/Kalman Filter)
- Particle per camera hypothesis, Kalman filter for features
- Allows larger maps: update $O(M \log N)$ instead of $O(N^2)$

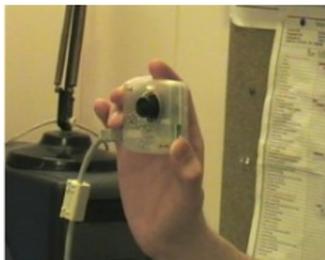


Outline

- 1 Introduction
 - The Camera as a Position Sensor
 - Visual SLAM
- 2 Davison's MonoSLAM
 - Overview and Nomenclature
 - Extended Kalman Filter
 - Automatic Map Management
 - Performance
- 3 The SceneLib Libraries
 - Introduction
 - The Scene Library
 - The MonoSLAM Library
 - Applications using SceneLib
- 4 Final Thoughts
 - Final Thoughts



Davison's MonoSLAM: Overview

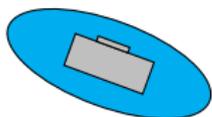


Main features

- Initialisation with known target
- Extended Kalman Filter
 - 'Constant velocity' motion model
 - Image patch features with Active Search
- Automatic Map Measurement
- Particle filter for initialisation of new features



Nomenclature



The camera

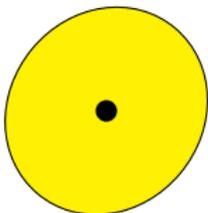
- The **camera position state** \mathbf{x}_p is its 3D position and orientation

$$\mathbf{x}_p = \begin{pmatrix} \mathbf{r}^W \\ \mathbf{q}^{WR} \end{pmatrix} = \begin{pmatrix} x \\ y \\ z \\ q_0 \\ q_x \\ q_y \\ q_z \end{pmatrix}$$

- The **camera state** \mathbf{x}_v contains \mathbf{x}_p plus optional additional state information (e.g. velocity and angular velocity)



Nomenclature



Features

- Each **feature** \mathbf{y}_i is a 3D position vector

$$\mathbf{y}_i = \begin{pmatrix} x_i \\ y_i \\ z_i \end{pmatrix}$$



Nomenclature

System State

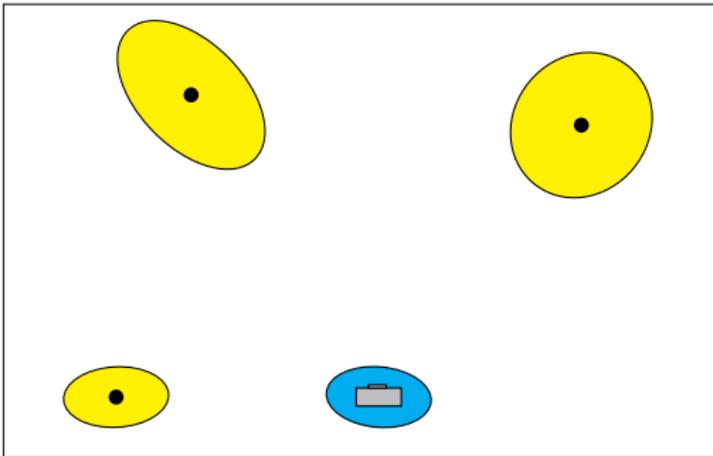
$$\mathbf{x} = \begin{pmatrix} \mathbf{x}_v \\ \mathbf{y}_1 \\ \mathbf{y}_2 \\ \vdots \end{pmatrix} \quad \mathbf{P} = \begin{bmatrix} \mathbf{P}_{xx} & \mathbf{P}_{xy_1} & \mathbf{P}_{xy_2} & \dots \\ \mathbf{P}_{y_1x} & \mathbf{P}_{y_1y_1} & \mathbf{P}_{y_1y_2} & \dots \\ \mathbf{P}_{y_2x} & \mathbf{P}_{y_2y_1} & \mathbf{P}_{y_2y_2} & \dots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

- PDF over camera and feature state is modelled as a single **multi-variate Gaussian** and we can use the Extended Kalman Filter.



Extended Kalman Filter: Prediction Step

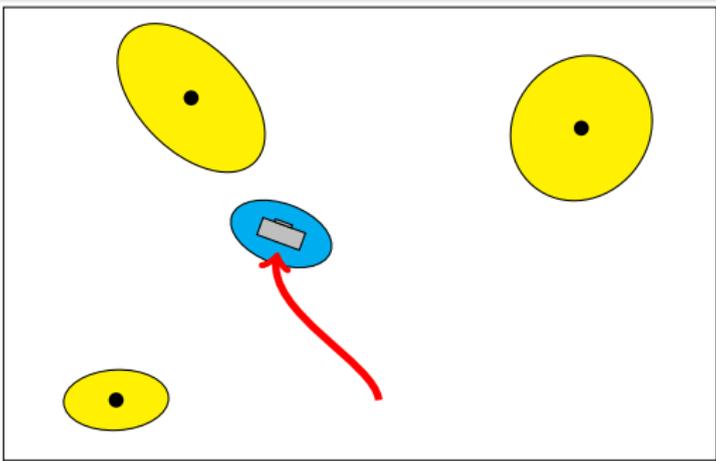
- ## Time Update
- 1 Estimate new location
 - 2 Add process noise



Extended Kalman Filter: Prediction Step

Time Update

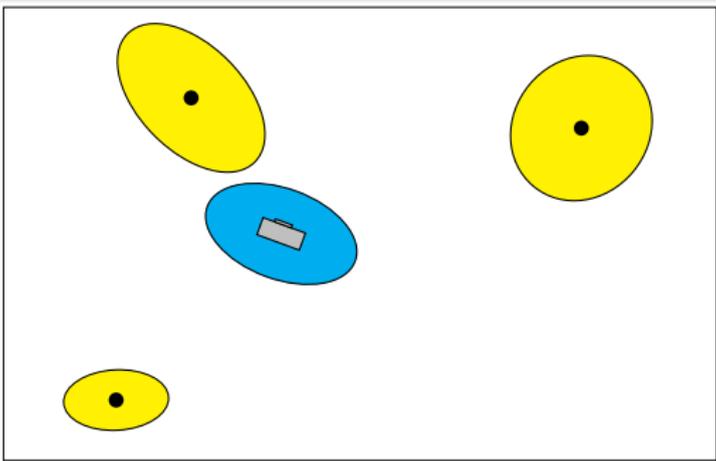
- 1 Estimate new location
- 2 Add process noise



Extended Kalman Filter: Prediction Step

Time Update

- 1 Estimate new location
- 2 Add process noise



Extended Kalman Filter: Prediction Step

Time Update

- 1 Project the state ahead

$$\hat{\mathbf{x}}_{\text{new}} = \mathbf{f}(\hat{\mathbf{x}}, \mathbf{u})$$

- 2 Project the error covariance ahead

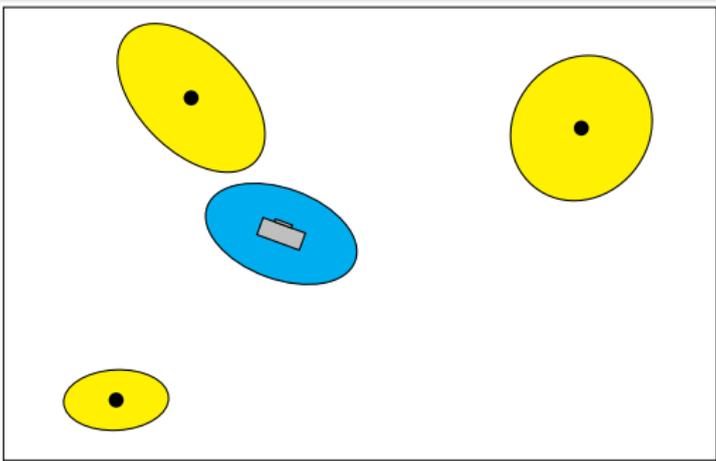
$$\mathbf{P}_{\text{new}} = \frac{\partial \mathbf{f}}{\partial \mathbf{x}} \mathbf{P} \frac{\partial \mathbf{f}^T}{\partial \mathbf{x}} + \mathbf{Q}$$



Extended Kalman Filter: Update Step

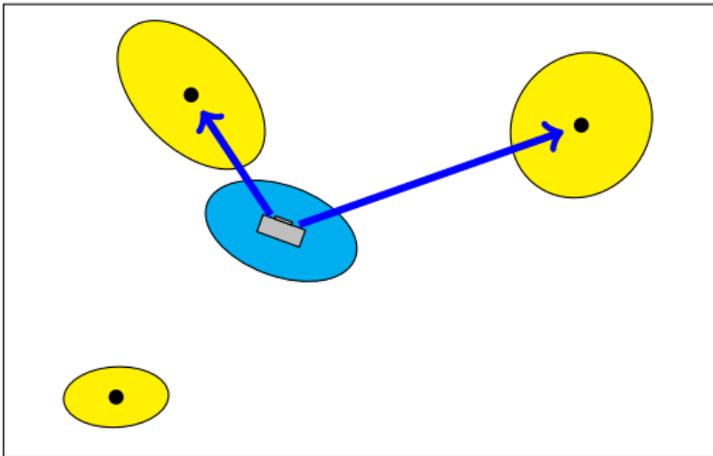
Measurement Update

- 1 Measure feature(s)
- 2 Update positions and uncertainties



Extended Kalman Filter: Update Step

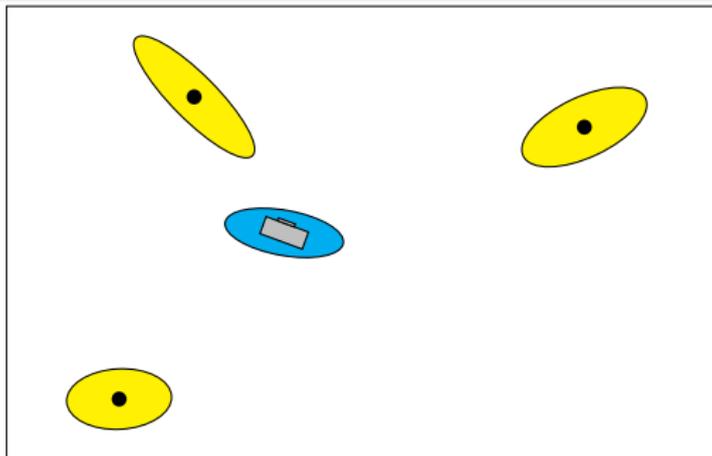
- ## Measurement Update
- 1 Measure feature(s)
 - 2 Update positions and uncertainties



Extended Kalman Filter: Update Step

Measurement Update

- 1 Measure feature(s)
- 2 Update positions and uncertainties



Extended Kalman Filter: Update Step

Measurement Update

- 1 Make measurement \mathbf{z} to give the **innovation** ν

$$\nu = \mathbf{z} - \mathbf{h}(\hat{\mathbf{x}})$$

- 2 Calculate **innovation covariance** \mathbf{S} and **Kalman gain** \mathbf{W}

$$\mathbf{S} = \frac{\partial \mathbf{h}}{\partial \mathbf{x}} \mathbf{P} \frac{\partial \mathbf{h}^T}{\partial \mathbf{x}} + \mathbf{R}$$

$$\mathbf{W} = \mathbf{P} \frac{\partial \mathbf{h}^T}{\partial \mathbf{x}} \mathbf{S}^{-1}$$

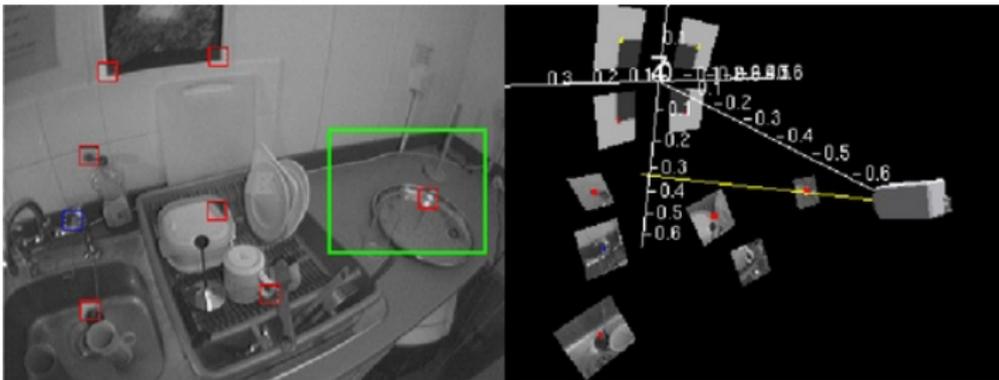
- 3 Update estimate and error covariance

$$\hat{\mathbf{x}}_{\text{new}} = \hat{\mathbf{x}} + \mathbf{W}\nu$$

$$\mathbf{P}_{\text{new}} = \mathbf{P} - \mathbf{W}\mathbf{S}\mathbf{W}^T$$



Measurement Step: Image Features and the Map

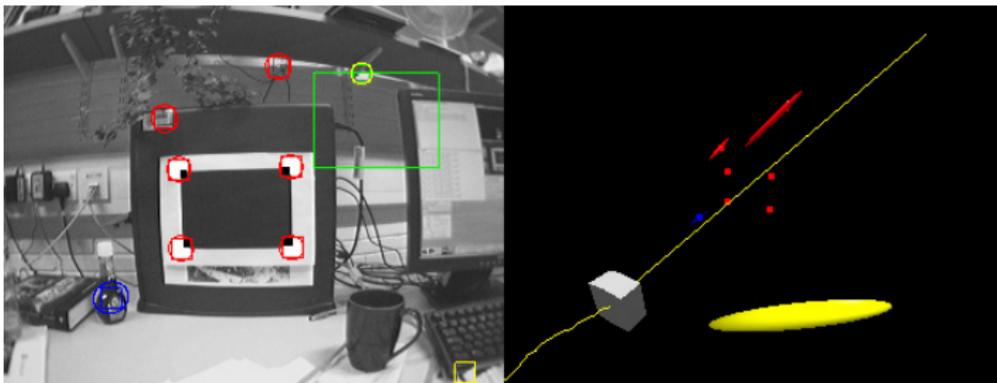


- Feature measurements are the locations of **salient image patches**.
- Patches are detected **once** to serve as **long-term visual landmarks**.
- **Sparse** set of landmarks gradually accumulated and **stored indefinitely**.

Monocular Feature Initialisation with Depth Particles

A new feature has unknown depth

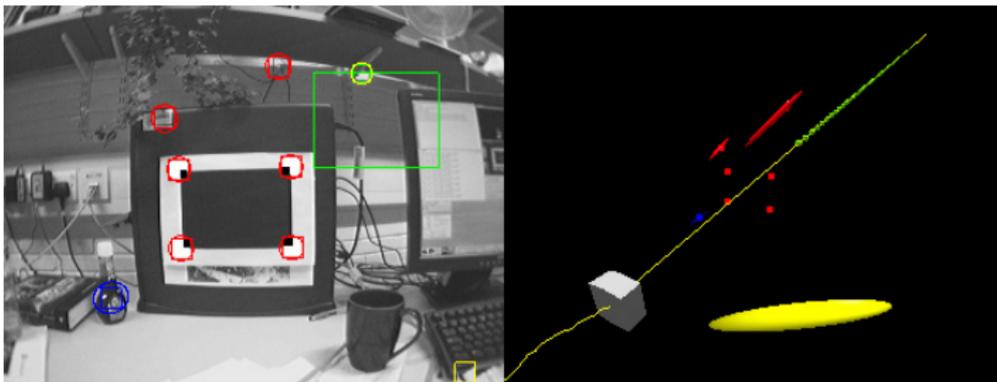
- 1 Populate the line with 100 particles, spaced uniformly between 0.5m and 5m from the camera.
- 2 Match each particle in successive frames to find probability of that depth.
- 3 When depth covariance is small,



Monocular Feature Initialisation with Depth Particles

A new feature has unknown depth

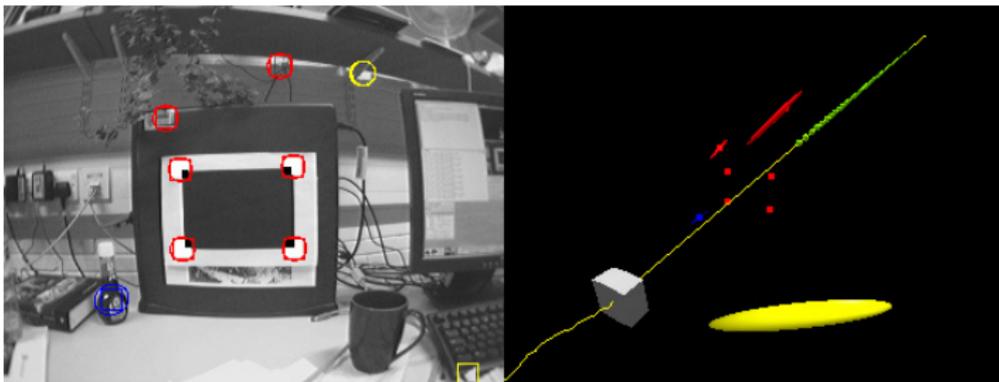
- 1 Populate the line with 100 particles, spaced uniformly between 0.5m and 5m from the camera.
- 2 Match each particle in successive frames to find probability of that depth.
- 3 When depth covariance is small,



Monocular Feature Initialisation with Depth Particles

A new feature has unknown depth

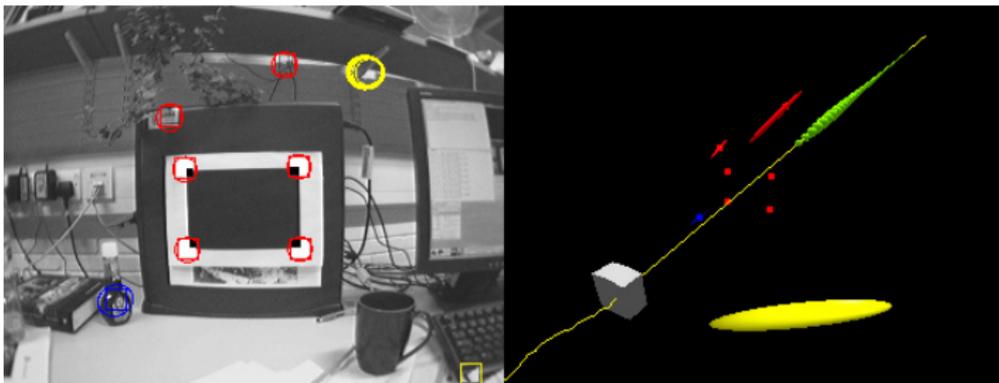
- 1 Populate the line with 100 particles, spaced uniformly between 0.5m and 5m from the camera.
- 2 Match each particle in successive frames to find probability of that depth.
- 3 When depth covariance is small,



Monocular Feature Initialisation with Depth Particles

A new feature has unknown depth

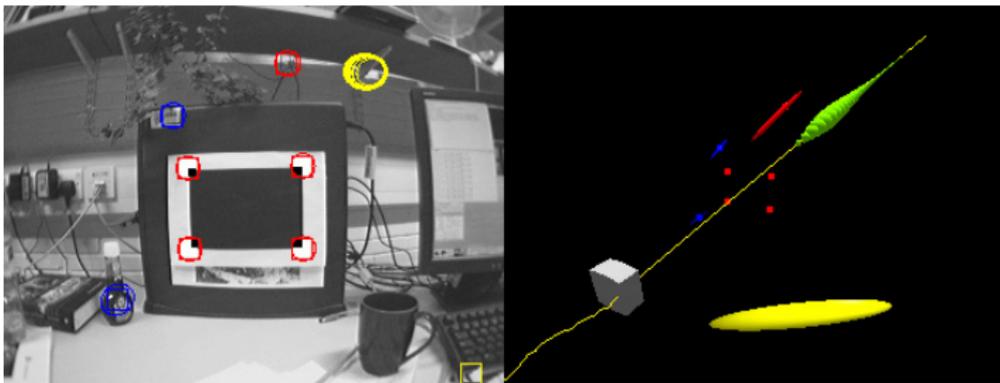
- 1 Populate the line with 100 particles, spaced uniformly between 0.5m and 5m from the camera.
- 2 Match each particle in successive frames to find probability of that depth.
- 3 When depth covariance is small,



Monocular Feature Initialisation with Depth Particles

A new feature has unknown depth

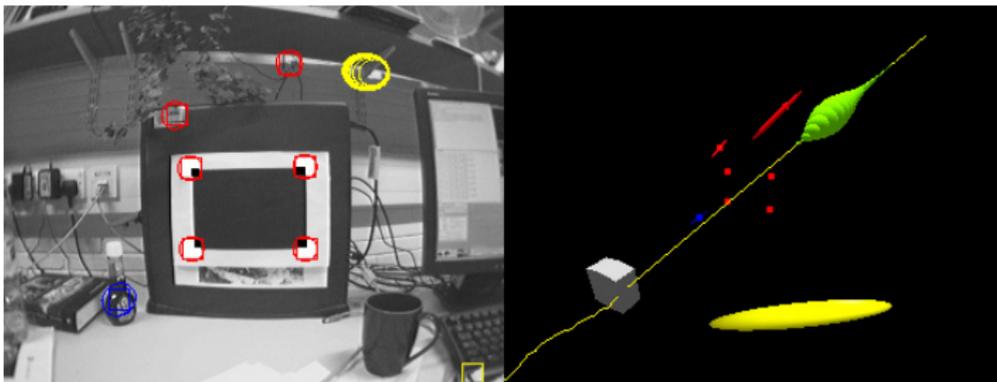
- 1 Populate the line with 100 particles, spaced uniformly between 0.5m and 5m from the camera.
- 2 Match each particle in successive frames to find probability of that depth.
- 3 When depth covariance is small,



Monocular Feature Initialisation with Depth Particles

A new feature has unknown depth

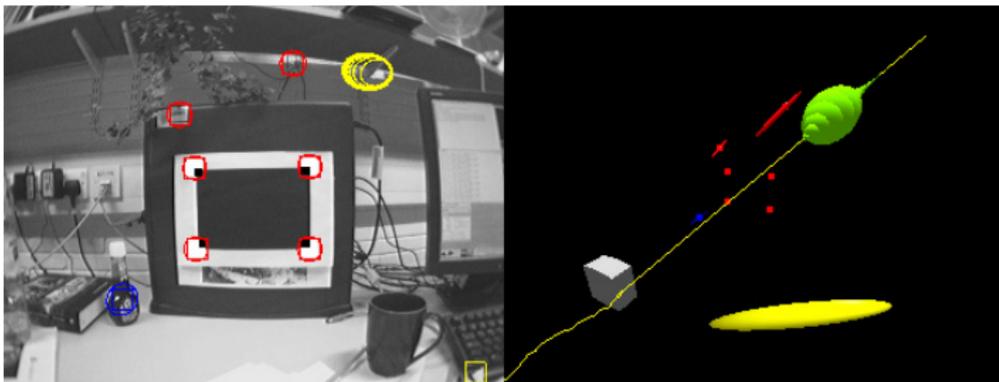
- 1 Populate the line with 100 particles, spaced uniformly between 0.5m and 5m from the camera.
- 2 Match each particle in successive frames to find probability of that depth.
- 3 When depth covariance is small,



Monocular Feature Initialisation with Depth Particles

A new feature has unknown depth

- 1 Populate the line with 100 particles, spaced uniformly between 0.5m and 5m from the camera.
- 2 Match each particle in successive frames to find probability of that depth.
- 3 When depth covariance is small,



Feature initialisation

- New features need adding to the state vector and covariance matrix.

Increasing the state size dynamically

$$\mathbf{x}_{\text{new}} = \begin{pmatrix} \mathbf{x}_v \\ \mathbf{y}_1 \\ \mathbf{y}_2 \\ \mathbf{y}_i \end{pmatrix}$$

$$\mathbf{P} = \begin{bmatrix} \mathbf{P}_{xx} & \mathbf{P}_{xy_1} & \mathbf{P}_{xy_2} & \mathbf{P}_{xx} \frac{\partial \mathbf{y}_i}{\partial \mathbf{x}_v}^T \\ \mathbf{P}_{y_1x} & \mathbf{P}_{y_1y_1} & \mathbf{P}_{y_1y_2} & \mathbf{P}_{y_1x} \frac{\partial \mathbf{y}_i}{\partial \mathbf{x}_v}^T \\ \mathbf{P}_{y_2x} & \mathbf{P}_{y_2y_1} & \mathbf{P}_{y_2y_2} & \mathbf{P}_{y_2x} \frac{\partial \mathbf{y}_i}{\partial \mathbf{x}_v}^T \\ \frac{\partial \mathbf{y}_i}{\partial \mathbf{x}_v} \mathbf{P}_{xx} & \frac{\partial \mathbf{y}_i}{\partial \mathbf{x}_v} \mathbf{P}_{xy_1} & \frac{\partial \mathbf{y}_i}{\partial \mathbf{x}_v} \mathbf{P}_{xy_2} & \frac{\partial \mathbf{y}_i}{\partial \mathbf{x}_v} \mathbf{P}_{xx} \frac{\partial \mathbf{y}_i}{\partial \mathbf{x}_v}^T + \frac{\partial \mathbf{y}_i}{\partial \mathbf{h}_G} \mathbf{R}_L \frac{\partial \mathbf{y}_i}{\partial \mathbf{h}_G}^T \end{bmatrix}$$



Feature deletion

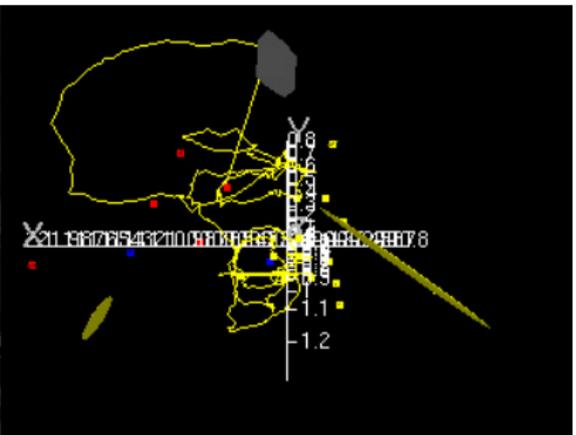
- Delete a feature if more than half of attempted measurements fail.

Reducing the state size dynamically

$$\begin{array}{c}
 \begin{pmatrix} \mathbf{x}_v \\ \mathbf{y}_1 \\ \mathbf{y}_2 \\ \mathbf{y}_3 \end{pmatrix} \rightarrow \begin{pmatrix} \mathbf{x}_v \\ \mathbf{y}_1 \\ \mathbf{y}_3 \end{pmatrix} \\
 \begin{bmatrix} \mathbf{P}_{xx} & \mathbf{P}_{xy_1} & \mathbf{P}_{xy_2} & \mathbf{P}_{xy_3} \\ \mathbf{P}_{y_1x} & \mathbf{P}_{y_1y_1} & \mathbf{P}_{y_1y_2} & \mathbf{P}_{y_1y_3} \\ \mathbf{P}_{y_2x} & \mathbf{P}_{y_2y_1} & \mathbf{P}_{y_2y_2} & \mathbf{P}_{y_2y_3} \\ \mathbf{P}_{y_3x} & \mathbf{P}_{y_3y_1} & \mathbf{P}_{y_3y_2} & \mathbf{P}_{y_3y_3} \end{bmatrix} \rightarrow \begin{bmatrix} \mathbf{P}_{xx} & \mathbf{P}_{xy_1} & \mathbf{P}_{xy_3} \\ \mathbf{P}_{y_1x} & \mathbf{P}_{y_1y_1} & \mathbf{P}_{y_1y_3} \\ \mathbf{P}_{y_3x} & \mathbf{P}_{y_3y_1} & \mathbf{P}_{y_3y_3} \end{bmatrix}
 \end{array}$$



Example Sequence



What Enables It To Run In Real-time?

Timings

Image loading and administration	1ms
Image correlation searches	2ms
Kalman Filter update	1ms
Feature initialisation search	3ms
Graphical rendering	5ms
Total	12ms

Easily manages 30Hz processing on a 3.4GHz desktop PC using C++, Linux, OpenGL

Main time-saving features

- Automatic **map management** criteria to maintain a sufficient but sparse map
- **Active search** guided by uncertainty



What Enables It To Run In Real-time?

Timings

Image loading and administration	1ms
Image correlation searches	2ms
Kalman Filter update	1ms
Feature initialisation search	3ms
Graphical rendering	5ms
Total	12ms

Easily manages 30Hz processing on a 3.4GHz desktop PC using C++, Linux, OpenGL

Main time-saving features

- Automatic **map management** criteria to maintain a sufficient but sparse map
- **Active search** guided by uncertainty



Outline

- 1 Introduction
 - The Camera as a Position Sensor
 - Visual SLAM
- 2 Davison's MonoSLAM
 - Overview and Nomenclature
 - Extended Kalman Filter
 - Automatic Map Management
 - Performance
- 3 The SceneLib Libraries
 - Introduction
 - The Scene Library
 - The MonoSLAM Library
 - Applications using SceneLib
- 4 Final Thoughts
 - Final Thoughts



The SceneLib Libraries

- A complete basic Davison MonoSLAM system
- Written in standard C++
- Three libraries and an application

SceneLib A generic SLAM library. Base classes for motion models, features, measurements and the Kalman Filter.

SceneImproc Image processing for MonoSLAM (i.e. feature detection and correlation)

MonoSLAM Specific motion and feature-measurement models for single-camera SLAM, and a control class.

MonoSLAMGlow Application based on GLOW/GLUT library which uses the libraries.



Obtaining the SceneLib libraries

- All files available from the Active Vision CVS repository
- Will also need VW34 library (soon to be VW35)

Getting and building files from CVS

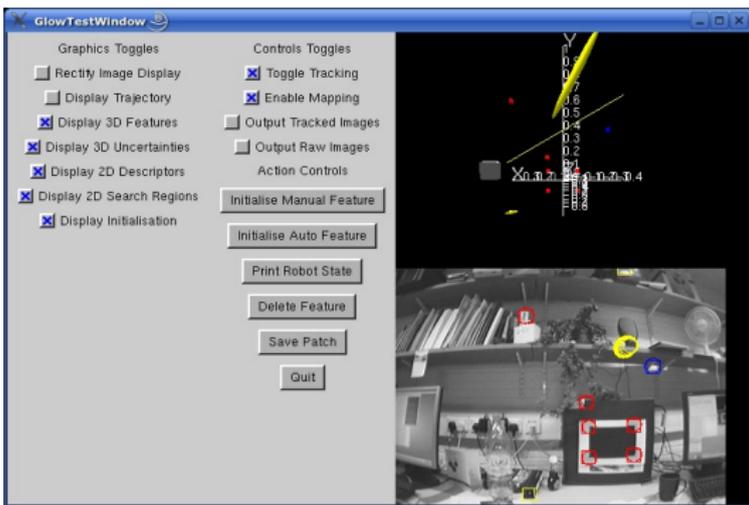
```
cvs -d/data/lav-local/common/cvsroot co VW34
cvs -d/data/lav-local/common/cvsroot co SceneLib
cvs -d/data/lav-local/common/cvsroot co MonoSLAMGlow
cd VW34
./bootstrap
./configure
make
cd ../SceneLib
./configure
make
cd ../MonoSLAMGlow
make
```



Running the MonoSLAMGlow application

Running MonoSLAMGlow

```
./scenerob
```

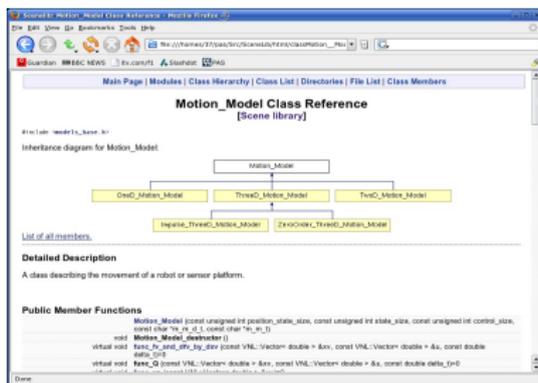


- Comment out `-D_REALTIME_` in Makefile to use previously saved image sequence

SceneLib Documentation

Making and viewing documentation

```
cd SceneLib
make docs
firefox html/index.html
```



● See also SceneLib/Docs/models.tex.



The Scene library

- **Scene** is a generic SLAM library.

Main Scene classes

Scene_Single Stores the full system state and manages features.

Feature Stores and manages a feature's state vector and covariances

Motion_Model Base class for all motion models.

Feature_Measurement_Model Base class for all feature and measurement models.

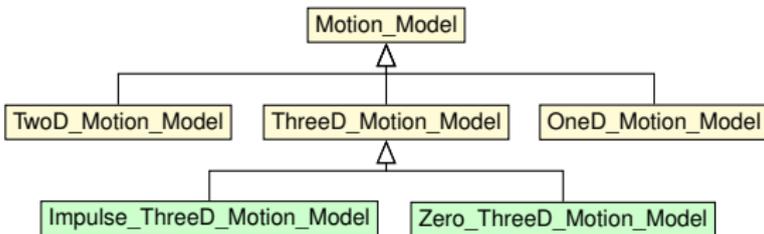
Sim_Or_Rob Base class for something that can make measurements.

Kalman Friend class of *Scene_Single* that implements a Kalman filter.



Motion model classes

- A *Motion_Model* knows how to perform the state update
 $\hat{\mathbf{x}}_{v\text{new}} = \mathbf{f}_v(\hat{\mathbf{x}}_v, \mathbf{u})$
- It stores no state itself



Main functions

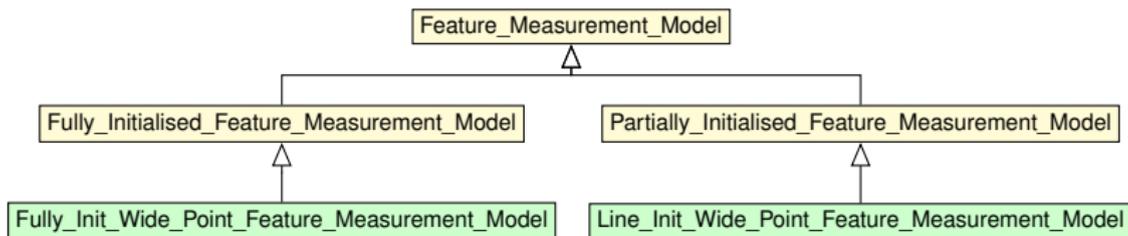
func_fv_and_dfv_by_dxv(xv,u,delta_t) Calculates the new camera state $\hat{\mathbf{x}}_{v\text{new}}$ and Jacobian $\frac{\partial \mathbf{f}_v}{\partial \mathbf{x}_v}$

func_Qi(xv, u, delta_t) Calculates the covariance \mathbf{Q}_i of the process noise.



Feature measurement classes

- A *Feature_Measurement_Model* knows how to manage a feature's state and predict a measurement $\mathbf{h}_i = \mathbf{h}(\mathbf{y}_i, \mathbf{x}_p)$
- It stores no state itself
- Subclassed into two special types:
 - Fully_Initialised_Feature_Measurement_Model* A complete feature in the SLAM map.
 - Partially_Initialised_Feature_Measurement_Model* A feature currently being initialised.



Fully-initialised Features

- A *Fully Initialised Feature Measurement Model* understands the state of a full feature and determines how it is measured.
- It stores no state itself.

Main functions

`func_hi_and_dhi_by_dxp_and_dhi_by_dyi(yi, xp)` Calculates the expected measurement vector \mathbf{h}_i and its Jacobians $\frac{\partial \mathbf{h}_i}{\partial \mathbf{x}_p}$ and $\frac{\partial \mathbf{h}_i}{\partial \mathbf{y}_i}$.

`func_Ri(hi)` Calculates the measurement noise \mathbf{R}_i .

`visibility_test(xp, yi, xp_orig, hi)` Decides whether a feature should be measured.

`func_Si(Pxx, Pxyi, Pyiyi, dhi_by_dxv, dhi_by_dyi, Ri)`
Calculates the innovation covariance \mathbf{S}_i .



Partially-initialised Features

- A *Partially_Initialised_Feature_Measurement_Model* handles a feature which still has some free parameters λ .
- Each partially-initialised feature references a *FeatureInitInfo*
- *FeatureInitInfo* stores a vector of *Particles* representing possible λ s and their probabilities.
- A *Partially_Initialised_Feature_Measurement_Model* can convert its feature into a fully-initialised feature.



The MonoSLAM library

The MonoSLAM library provides

- Specialisations of Scene base classes:
 - *Impulse_ThreeD_Motion_Model* and *ZeroOrder_ThreeD_Motion_Model* motion models.
 - *Fully_Init_Wide_Point_Feature_Measurement_Model* and *Line_Init_Wide_Point_Feature_Measurement_Model*.
 - *Robot* (derived from *Sim_Or_Rob*) to handle image feature measurement.
- A *MonoSLAM* class to provide the main interface.
- Functions to draw the two graphical displays.



The MonoSLAM Class

- The *MonoSLAM* class provides the basic SLAM functionality

Main functions

GoOneStep(image, delta_t, currently_mapping_flag) Step the system onto the next frame. (*image* is ignored, and instead it must be set using *Scene_Single::load_new_image()*).

- The *MonoSLAMInterface* class provides full control and feedback functions.

Main functions

GetScene() Get the *Scene_Single* class.

GetRobot() Get the *Robot* class.

plus >40 other *Get()* and *Set()* functions.



The MonoSLAMGlow Application

The screenshot shows the MonoSLAMGlow application window. The interface is divided into several sections:

- Graphics Toggles:**
 - Rectify Image Display
 - Display Trajectory
 - Display 3D Features
 - Display 3D Uncertainties
 - Display 2D Descriptors
 - Display 2D Search Regions
 - Display Initialisation
- Controls Toggles:**
 - Toggle Tracking
 - Enable Mapping
 - Output Tracked Images
 - Output Raw Images
- Action Controls:**
 - Initialise Manual Feature
 - Initialise Auto Feature
 - Print Robot State
 - Delete Feature
 - Save Patch
 - Quit

The main view is split into two panels:

- Top Panel:** A 2D plot showing feature positions and search regions. A yellow line indicates a feature that is not used. A red line indicates a successful match. A blue line indicates a failed match. The plot includes a coordinate system with axes labeled 'x' and 'y'.
- Bottom Panel:** A grayscale camera view of a room. Red squares highlight feature patches in the scene. A yellow square highlights an unused feature. A blue circle highlights a failed match. A yellow circle highlights a successful match.

Legend:

- Failed match (blue line)
- Successful match (red line)
- Unused feature (yellow line)



Live Demonstration

GlowTestWindow

Graphics Toggles

- Rectify Image Display
- Display Trajectory
- Display 3D Features
- Display 3D Uncertainties
- Display 2D Descriptors
- Display 2D Search Regions
- Display Initialisation

Controls Toggles

- Toggle Tracking
- Enable Mapping
- Output Tracked Images
- Output Raw Images

Action Controls

Initialise Manual Feature

Initialise Auto Feature

Print Robot State

Delete Feature

Save Patch

Quit

File Sequencer

Main Controls

Continuous Next Stop

Outline

- 1 Introduction
 - The Camera as a Position Sensor
 - Visual SLAM
- 2 Davison's MonoSLAM
 - Overview and Nomenclature
 - Extended Kalman Filter
 - Automatic Map Management
 - Performance
- 3 The SceneLib Libraries
 - Introduction
 - The Scene Library
 - The MonoSLAM Library
 - Applications using SceneLib
- 4 Final Thoughts
 - Final Thoughts



Final Thoughts: EKF-based Visual SLAM

Observations

- The Davison visual SLAM system works!
- It works reliably enough for a live demo.
- It needs no real hidden tricks needed to make it work.

Discussion

- Need **more, better** features to track
 - Faster initialisation (fewer particles?)
 - Use full-frame fast feature detection
- **Better initialisation**: how do we deal with points at infinity?
- **Motion model**: How do we get smoother, better tracks?
- **Loop closing**



Final Thoughts: EKF-based Visual SLAM

Observations

- The Davison visual SLAM system works!
- It works reliably enough for a live demo.
- It needs no real hidden tricks needed to make it work.

Discussion

- Need **more, better** features to track
 - Faster initialisation (fewer particles?)
 - Use full-frame fast feature detection
- **Better initialisation**: how do we deal with points at infinity?
- **Motion model**: How do we get smoother, better tracks?
- **Loop closing**



Final Thoughts: The SceneLib library

Observations

- The SceneLib libraries are (reasonably) well-designed and (reasonably) well-documented
- They make it easy to write a Davison-style visual SLAM application

Discussion

- Are they useful to the Active Vision group?
- Can we all use them and get the benefits in code sharing that that will bring?
- Can we at least all use the same nomenclature and colours?



Final Thoughts: The SceneLib library

Observations

- The SceneLib libraries are (reasonably) well-designed and (reasonably) well-documented
- They make it easy to write a Davison-style visual SLAM application

Discussion

- Are they useful to the Active Vision group?
- Can we all use them and get the benefits in code sharing that that will bring?
- Can we at least all use the same nomenclature and colours?

