Edge Landmarks in Monocular SLAM

Ethan Eade
Dr. Tom Drummond

Machine Intelligence Laboratory
Cambridge University
ee231@eng.cam.ac.uk

Simultaneous Localisation And Mapping

The One-Slide Version

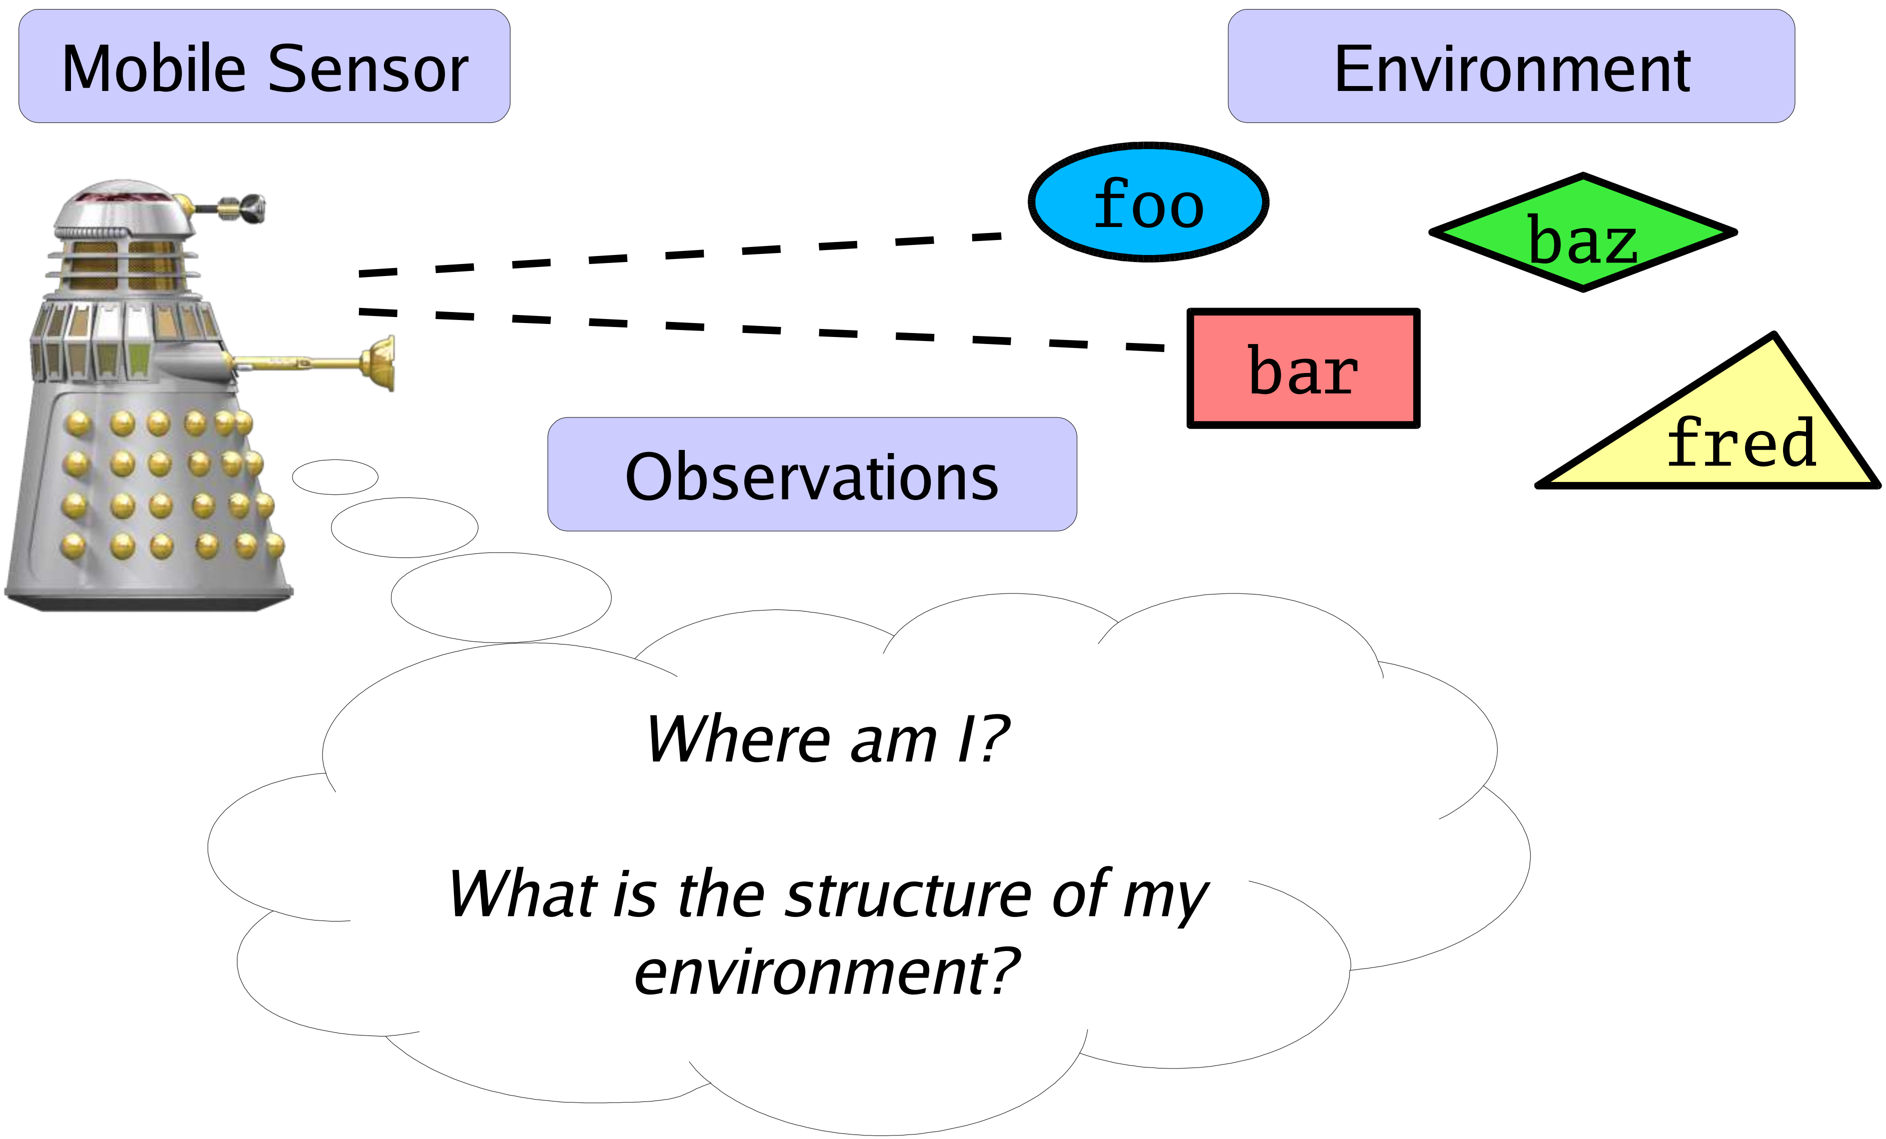

Ethan Eade ee231@cam.ac.uk

Monocular SLAM with Points

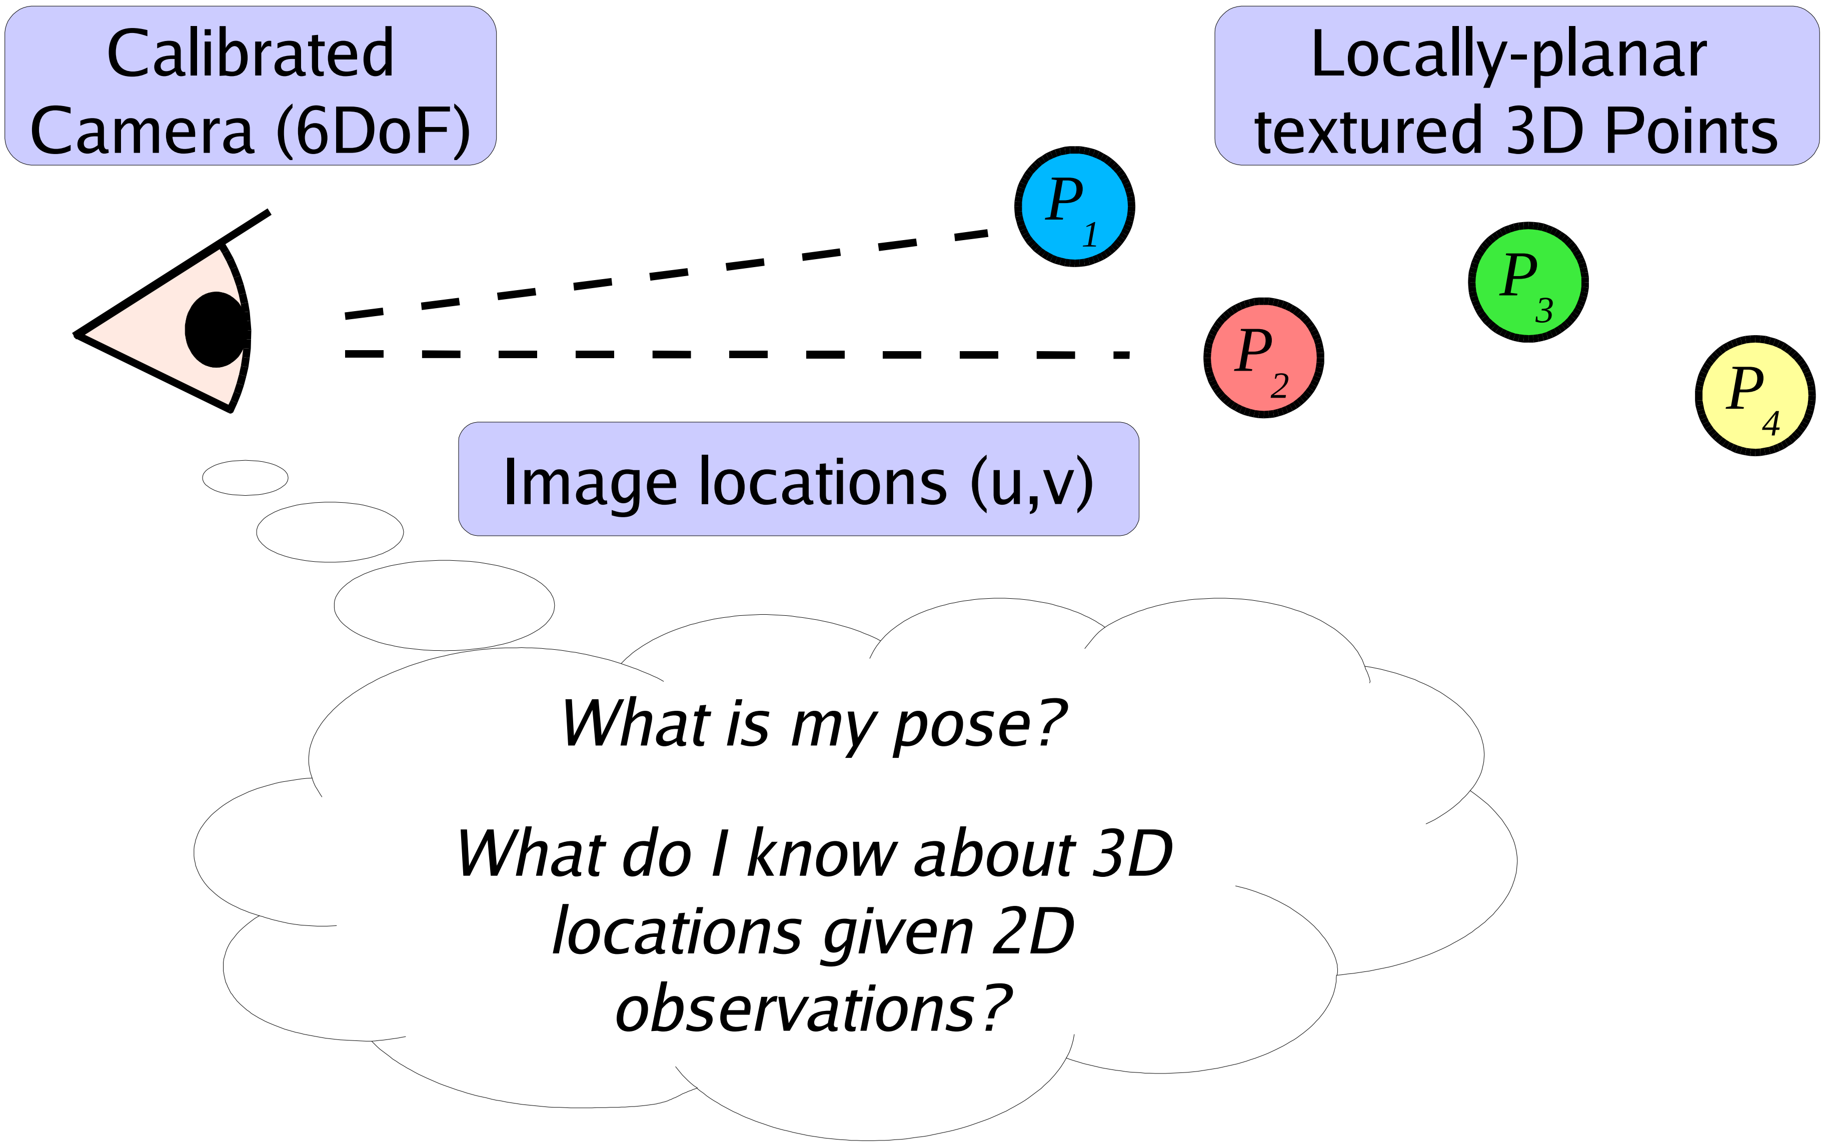

Ethan Eade ee231@cam.ac.uk

Monocular SLAM with Edges

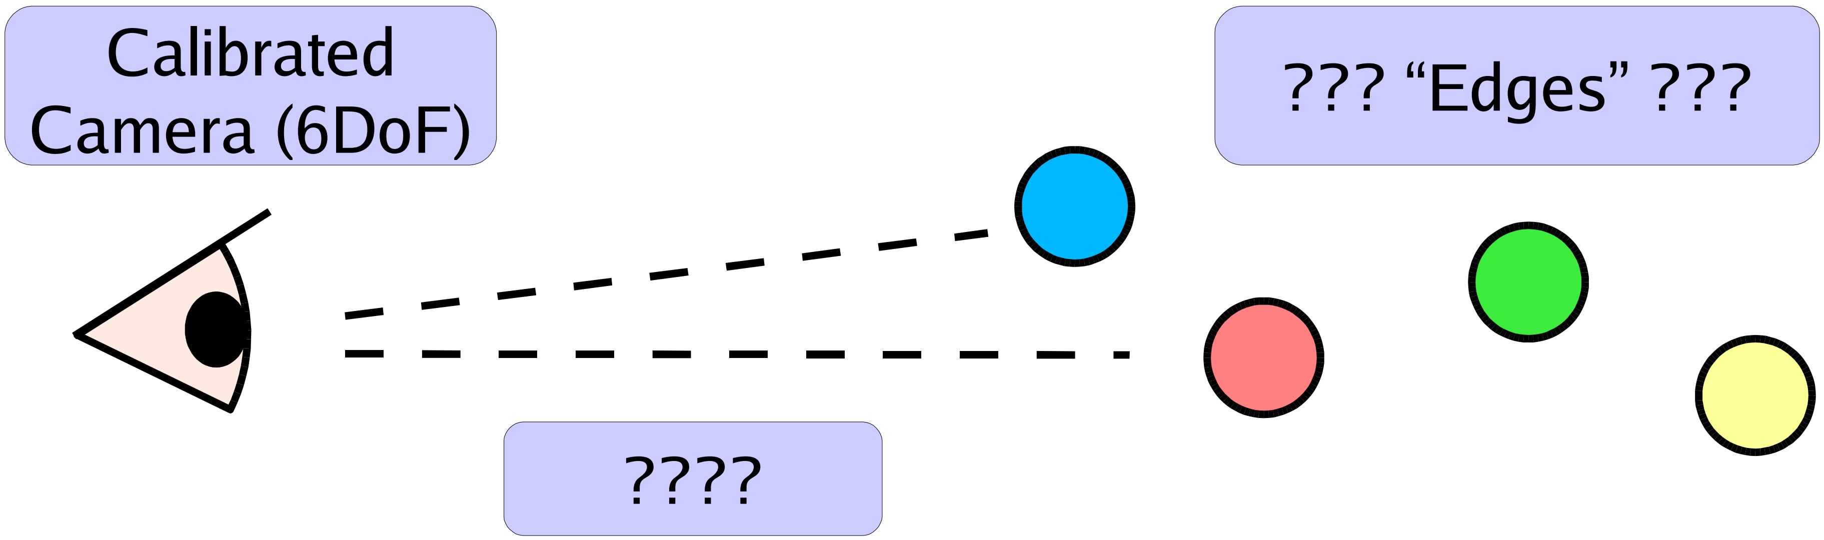

What is an edge in the world?

What is an edge in the image?

How should we {`select|parametrise|observe|estimate`} edges?

Monocular SLAM with Edges

Required

- Definition
- Representation
- Observation
- Selection
- Initialisation

Why Bother?

- Build a richer map
- Depend less on textured environment
- Benefit from more information

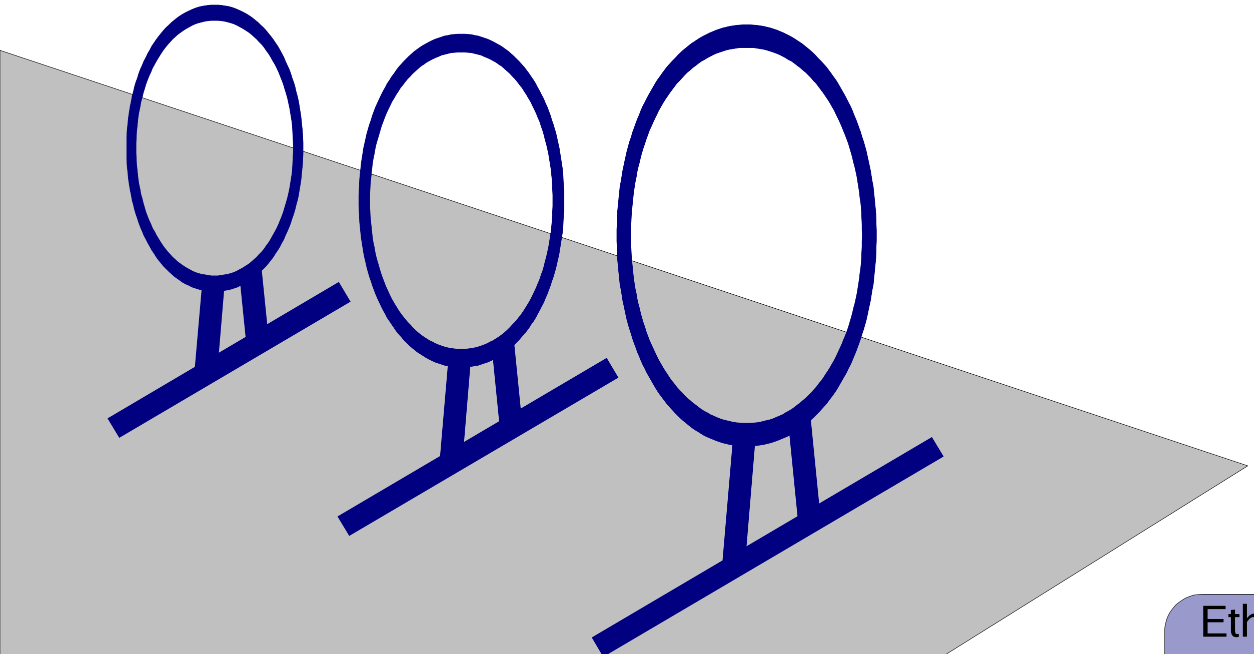

Monocular SLAM with Edges

Required

- Definition
- Representation
- Observation
- Selection
- Initialisation

Why Bother?

- Build a richer map
- Depend less on textured environment
- Benefit from more information

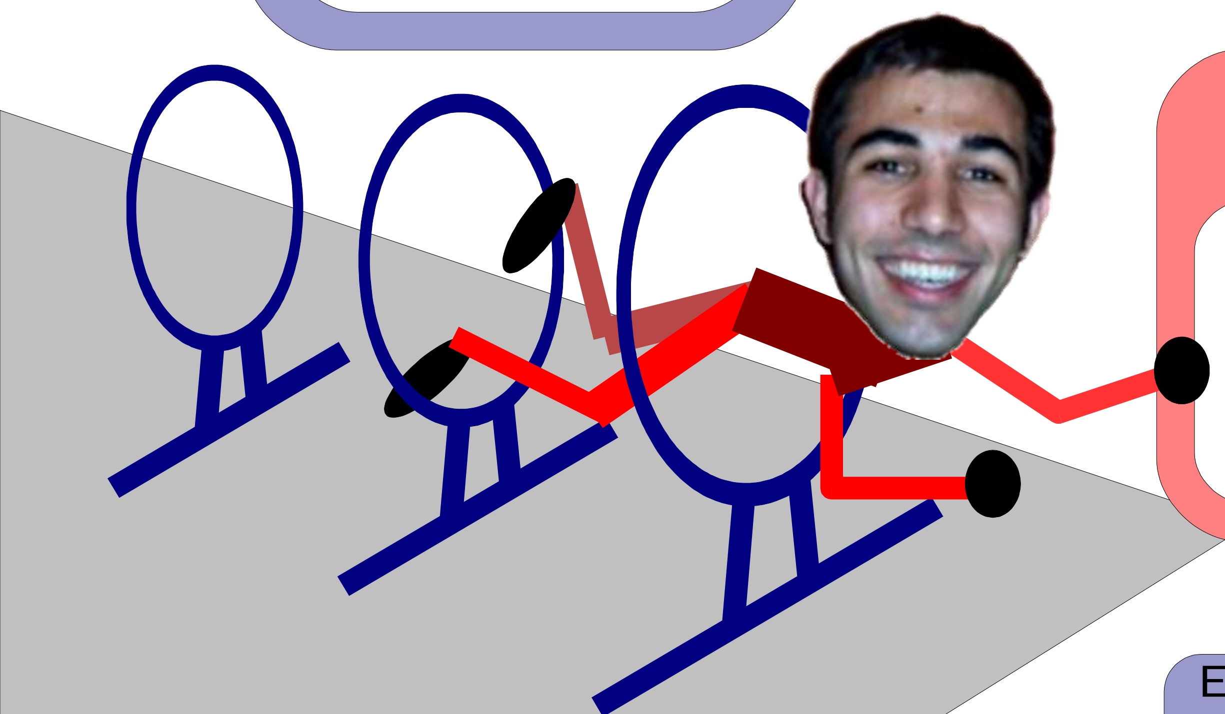

Contributions

Define ***edgelets***; show how to represent, observe, select, and initialise them

Ethan Eade ee231@cam.ac.uk

Monocular SLAM with Edges

Required

- Definition
- Representation
- Observation
- Selection
- Initialisation

Why Bother?

- Build a richer map
- Depend less on textured environment
- Benefit from more information

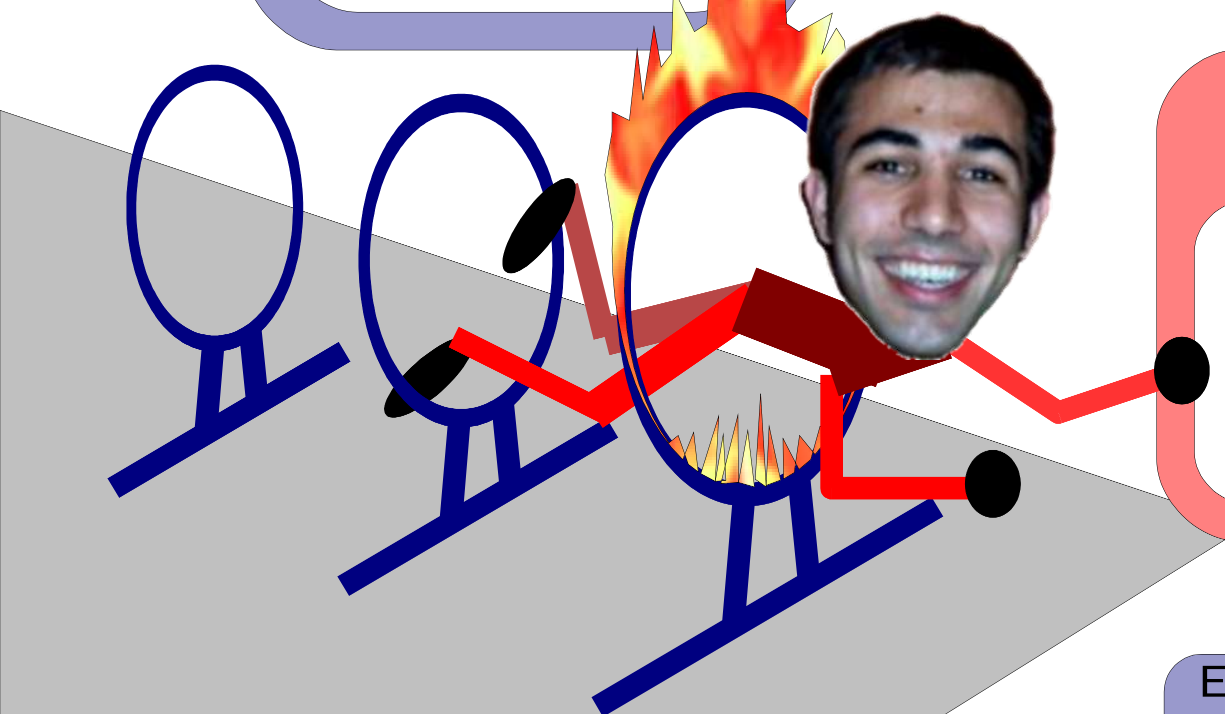

Contributions

Define ***edgelets***; show how to represent, observe, select, and initialise them

Definition

edgelet:

A locally-straight piece of 1D structure that appears as an intensity change.

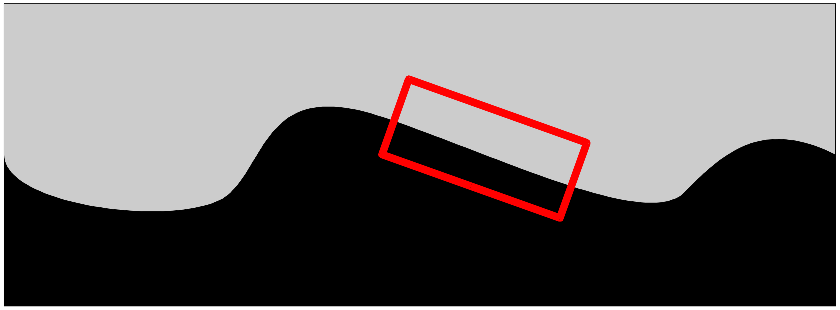

An edgelet...

- might be part of a longer straight segment
- might be part of a longer, slowly curving or piecewise-straight segment
- appears as a short straight edge in images
- is more than an edgel

Representation

An edgelet describes the derivative of the 3D edge $L(t)$ at a chosen point:

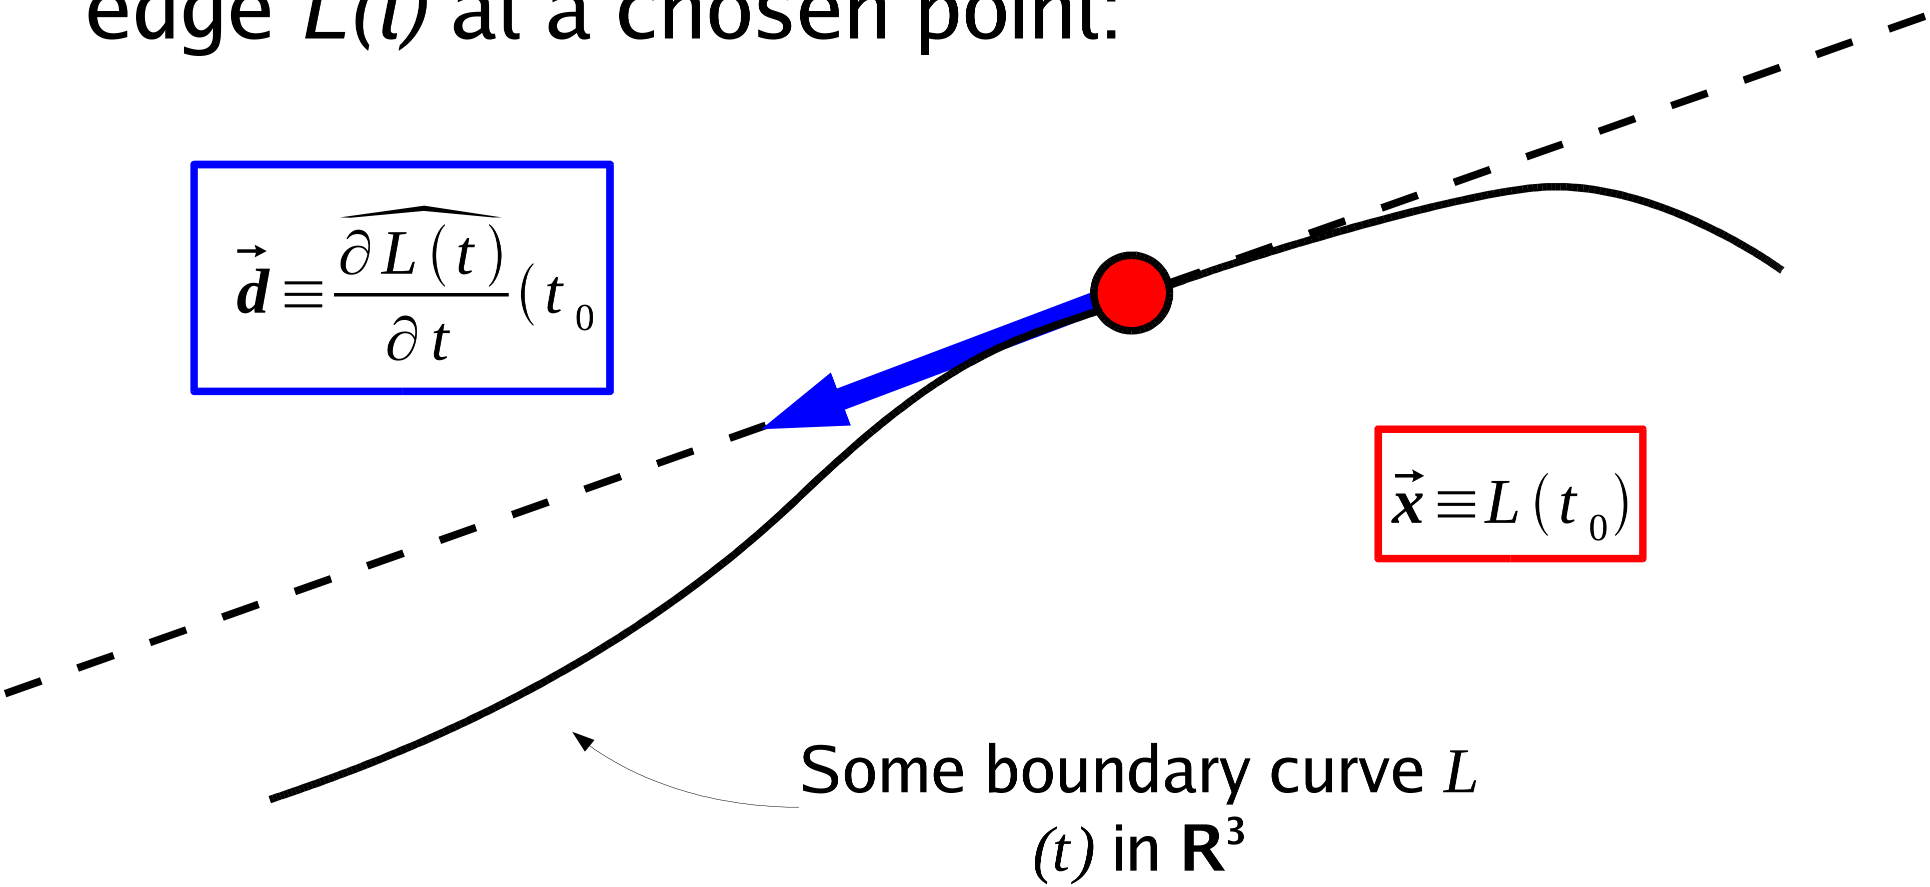

Ethan Eade ee231@cam.ac.uk

This Slide is Boring but Important

Point	Edgelet
Locally planar structure	Locally linear structure
3D state: point	5D state: point, direction
2D texture	1D intensity change
2D observations (image point)	2D observations (image line)
Distinctive, rich descriptor	1-bit descriptor

Observation

Project uncertainty into image, and search for straight segments in a gated region.

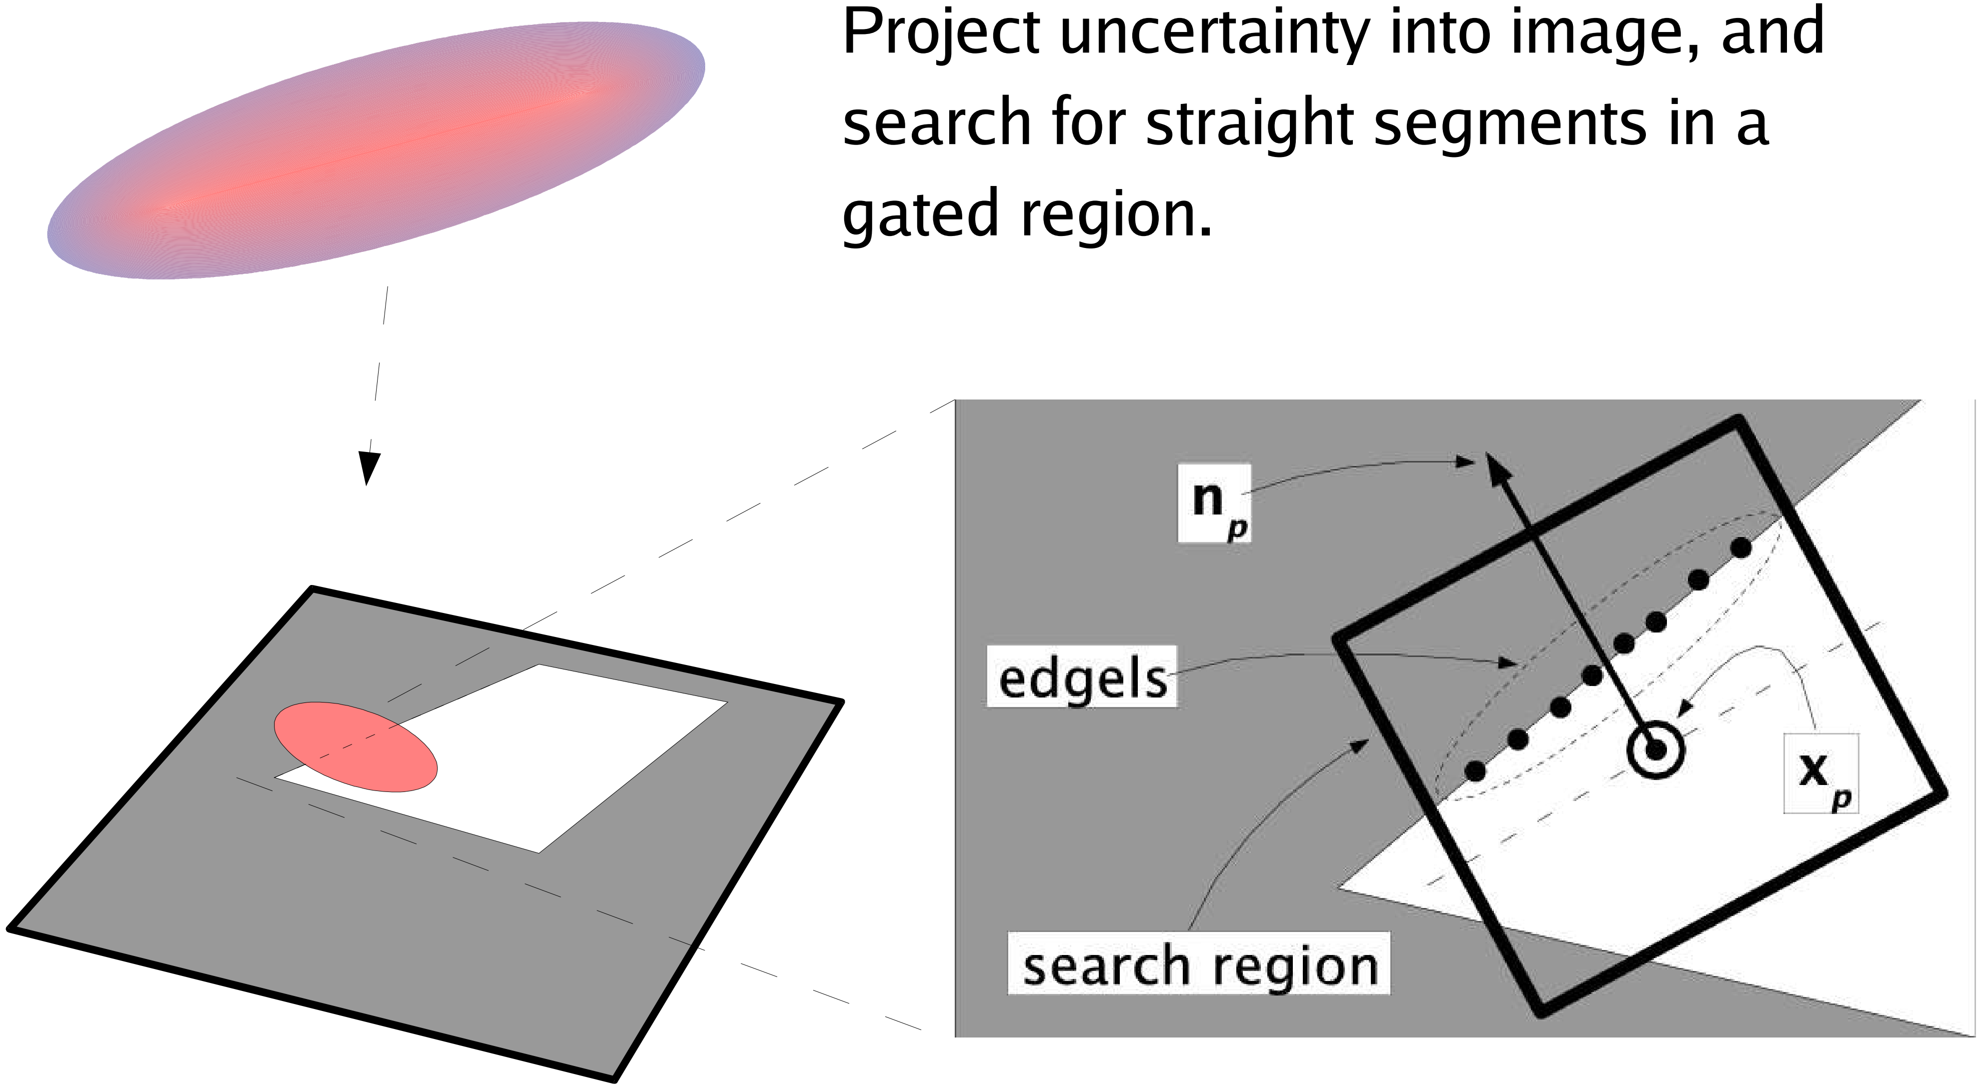

Observation

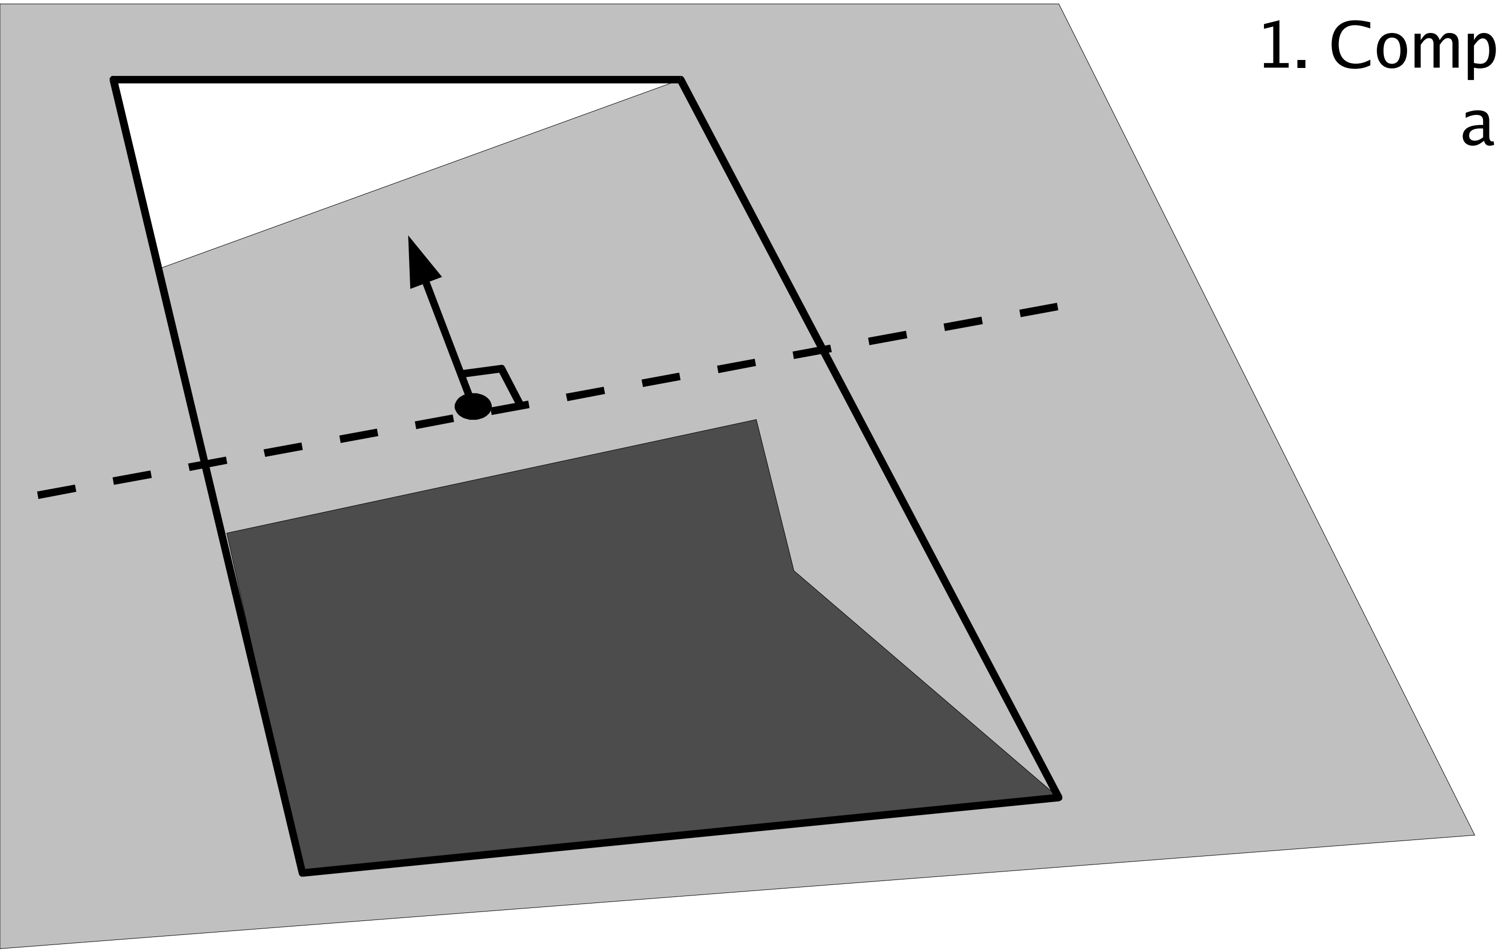

1. Compute search region and direction

Observation

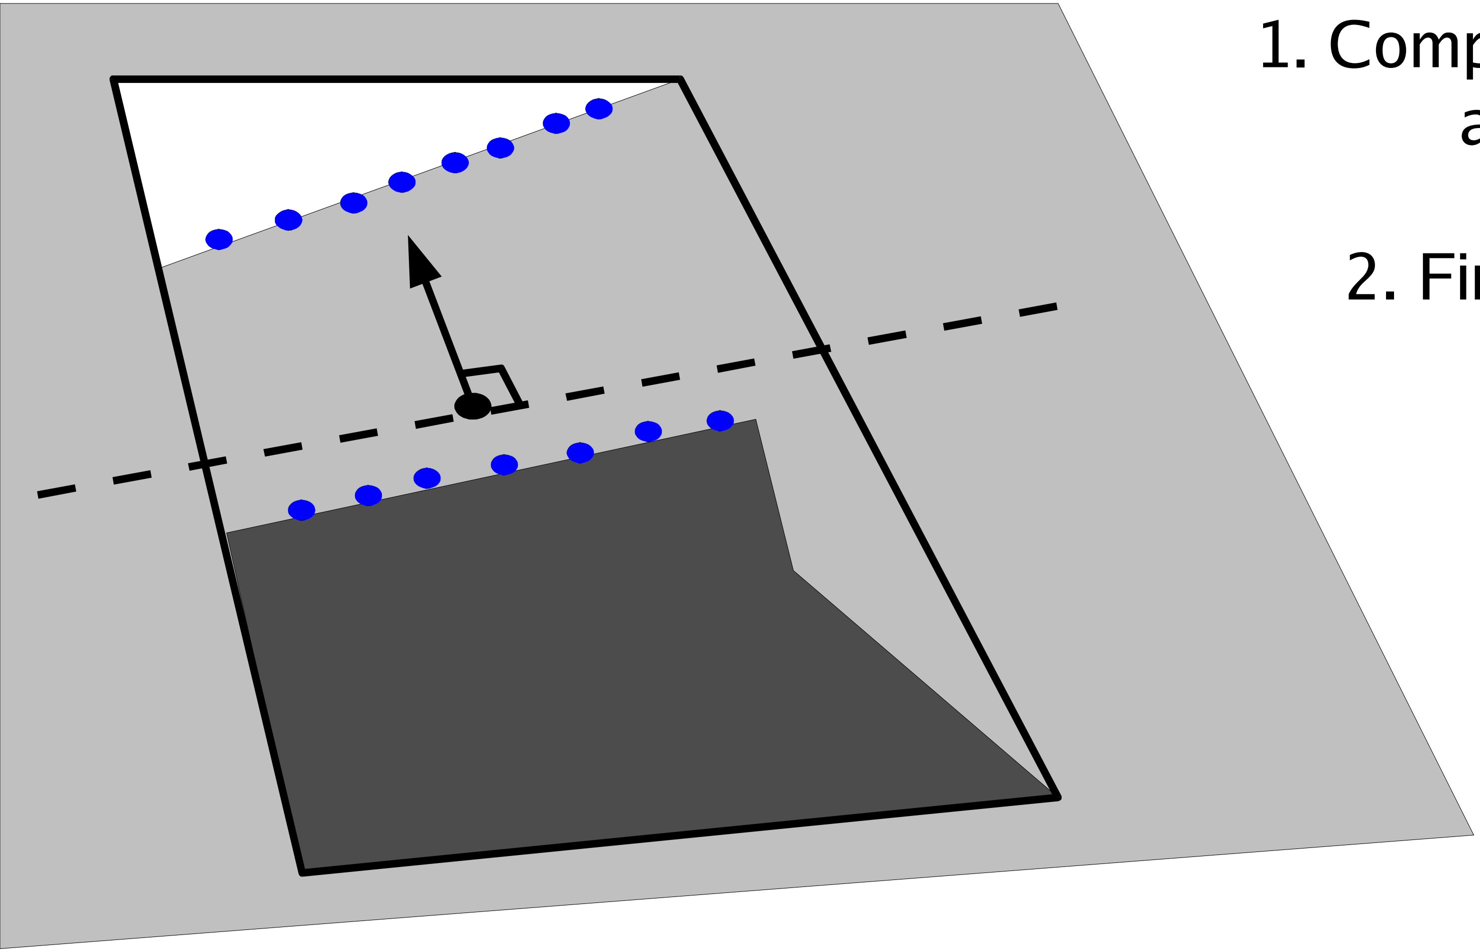

1. Compute search region and direction

2. Find edgels in region

Observation

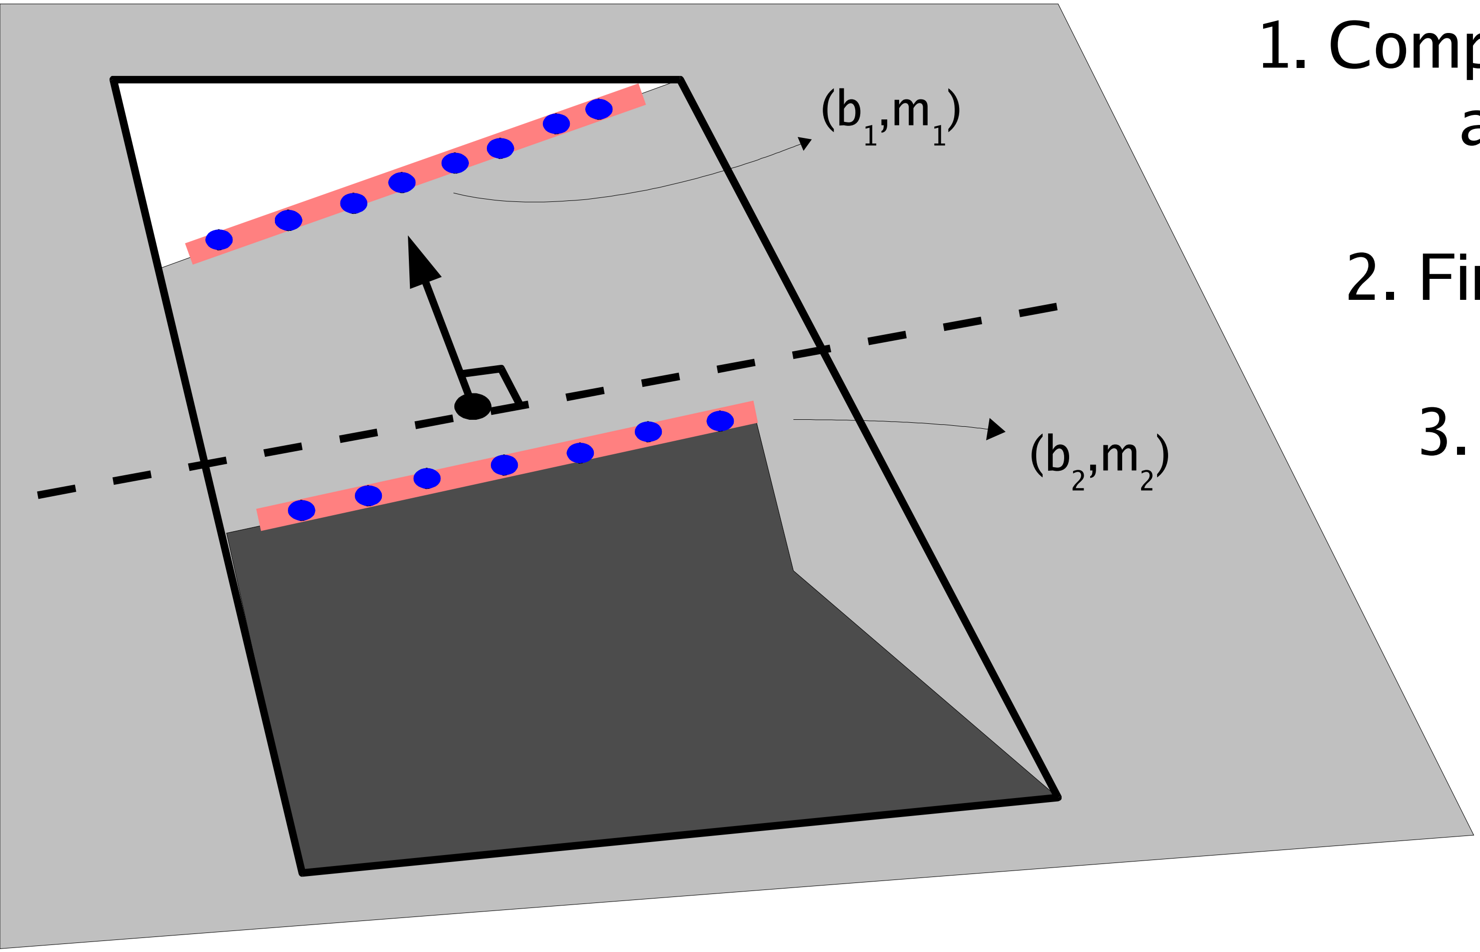

1. Compute search region and direction
2. Find edgels in region
3. Group edgels into straight segments

Observation

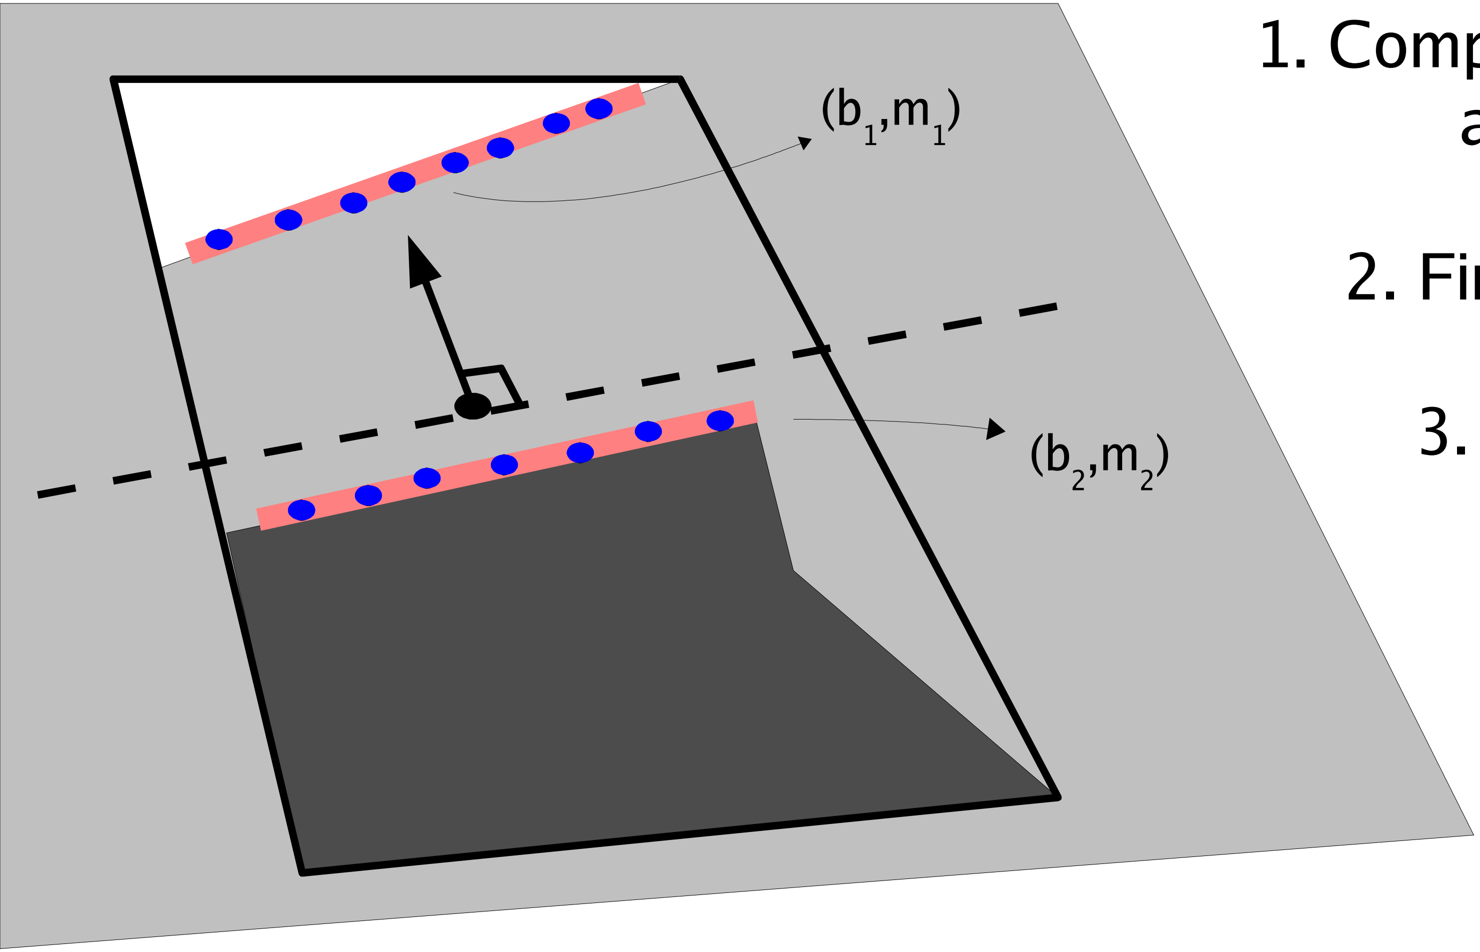

1. Compute search region and direction
2. Find edgels in region
3. Group edgels into straight segments

Multiple possible matches

Observation

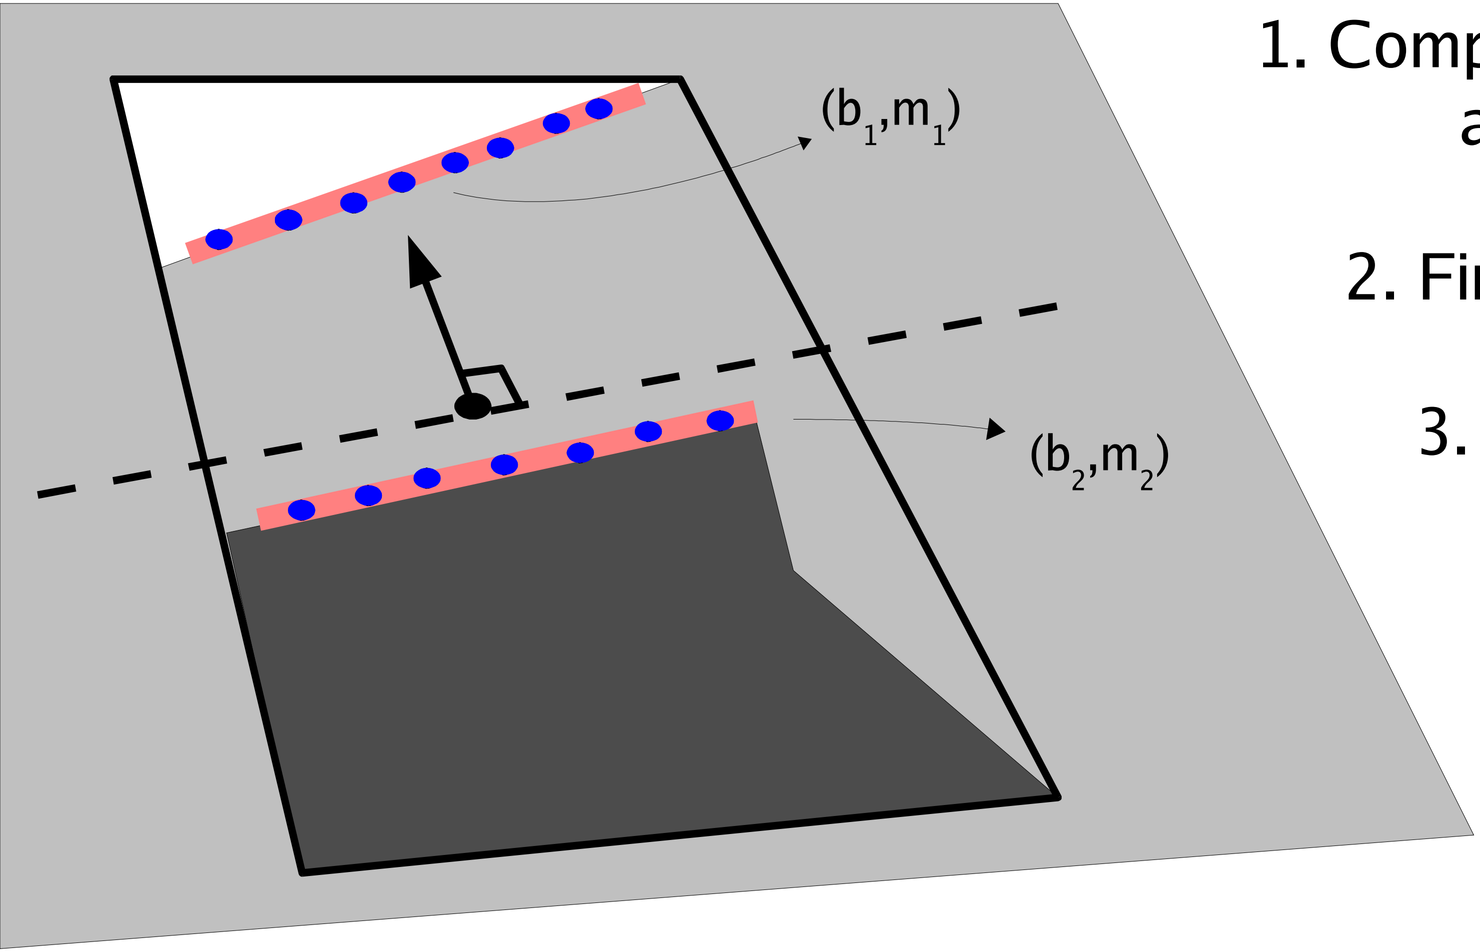

1. Compute search region and direction
2. Find edgels in region
3. Group edgels into straight segments

Multiple possible matches

Collect all observations each frame and find the maximum-likelihood data association using MLESAC

Selection: Finding new Edgelets

Data association is not straightforward for edgelets. How can we make our job easier?

We want short edge segments that are

- straight
- separated from other segments of similar orientation

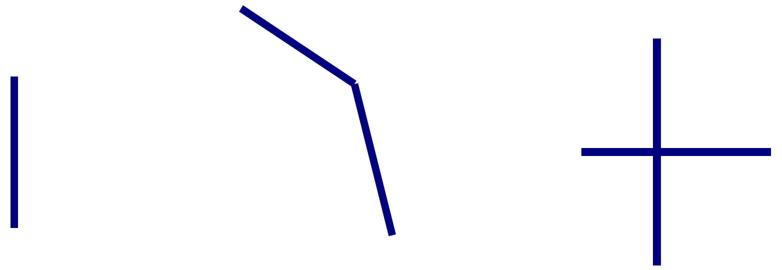

Good

Bad

Ethan Eade ee231@cam.ac.uk

Selection: Finding new Edgelets

1. Find Canny-like edgels (no linking)
2. Divide the edgels with a grid.
   For each cell:
   a) Find the dominant gradient direction *d*
   b) Find edgels in agreement with *d*
   c) Threshold on variance of position in direction *d*

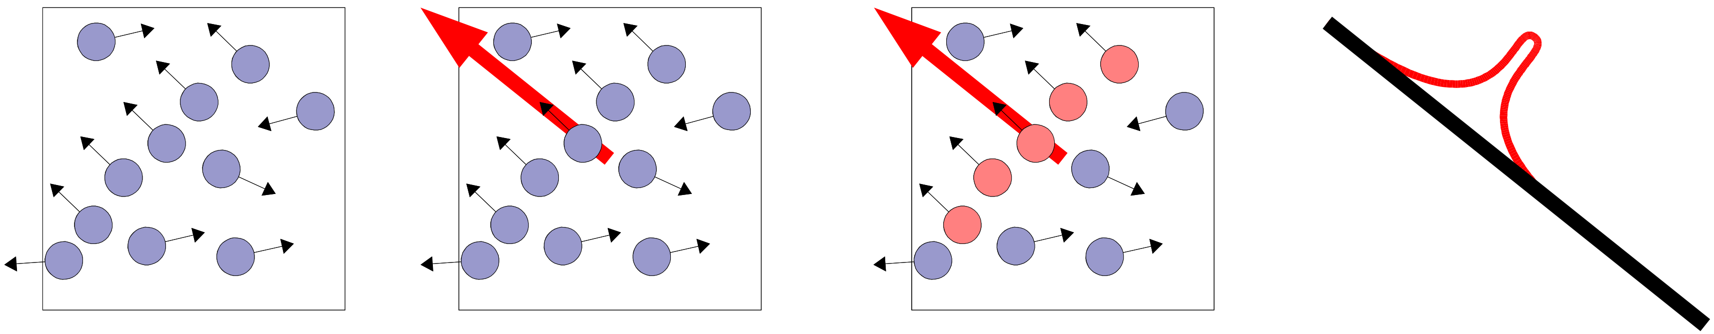

Initialisation of New Edgelets

- Highly non-Gaussian distributions in Cartesian coordinates

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} \qquad \begin{pmatrix} dx \\ dy \\ dz \end{pmatrix}$$

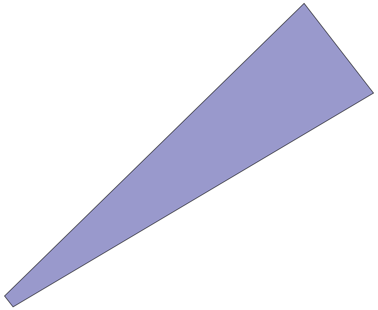

Ethan Eade ee231@cam.ac.uk

Initialisation of New Edgelets

- Highly non-Gaussian distributions in Cartesian coordinates

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} \quad \begin{pmatrix} dx \\ dy \\ dz \end{pmatrix}$$

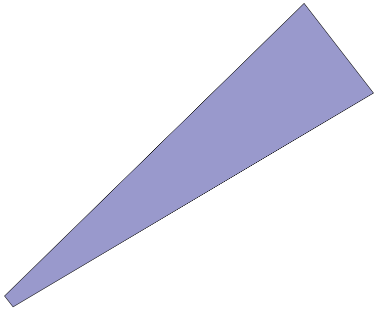

- Inverse-depth coordinates make the observation model nearly linear...

$$\begin{pmatrix} u \\ v \\ q \equiv \dfrac{1}{z} \end{pmatrix} \quad \begin{pmatrix} du \\ dv \\ dq \end{pmatrix} \quad C_0 \equiv (R_{0,} T_0)$$

Initialisation of New Edgelets

- Highly non-Gaussian distributions in Cartesian coordinates

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} \quad \begin{pmatrix} dx \\ dy \\ dz \end{pmatrix}$$

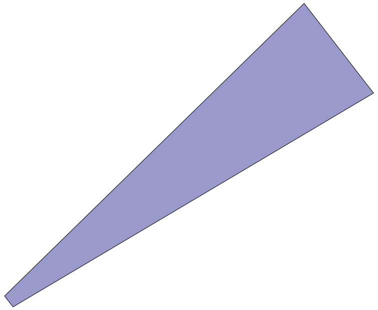

- Inverse-depth coordinates make the observation model nearly linear...

$$\begin{pmatrix} u \\ v \\ q \equiv \frac{1}{z} \end{pmatrix} \quad \begin{pmatrix} du \\ dv \\ dq \end{pmatrix} \quad C_0 \equiv (R_{0,} T_0)$$

- ... so the estimate is better represented by a Gaussian.

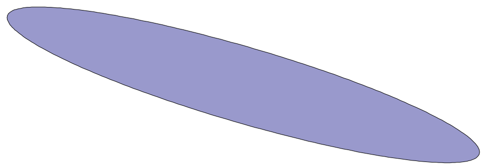

Problem solved.

Ethan Eade ee231@cam.ac.uk

Initialisation of New Edgelets

- Highly non-Gaussian distributions in Cartesian coordinates

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} \quad \begin{pmatrix} dx \\ dy \\ dz \end{pmatrix}$$

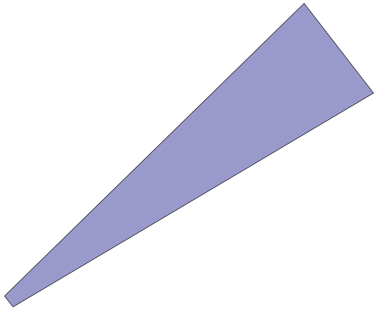

- Inverse-depth coordinates make the observation model nearly linear...

$$\begin{pmatrix} u \\ v \\ q \equiv \frac{1}{z} \end{pmatrix} \quad \begin{pmatrix} du \\ dv \\ dq \end{pmatrix} \quad C_0 \equiv (R_0, T_0)$$

- ... so the estimate is better represented by a Gaussian.

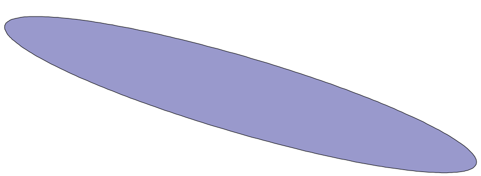

Problem solved. ← **This is a fib**

Ethan Eade ee231@cam.ac.uk

The author demonstrates
~~failure modes of~~
the system

Thank you

Life-cycle of an Edgelet

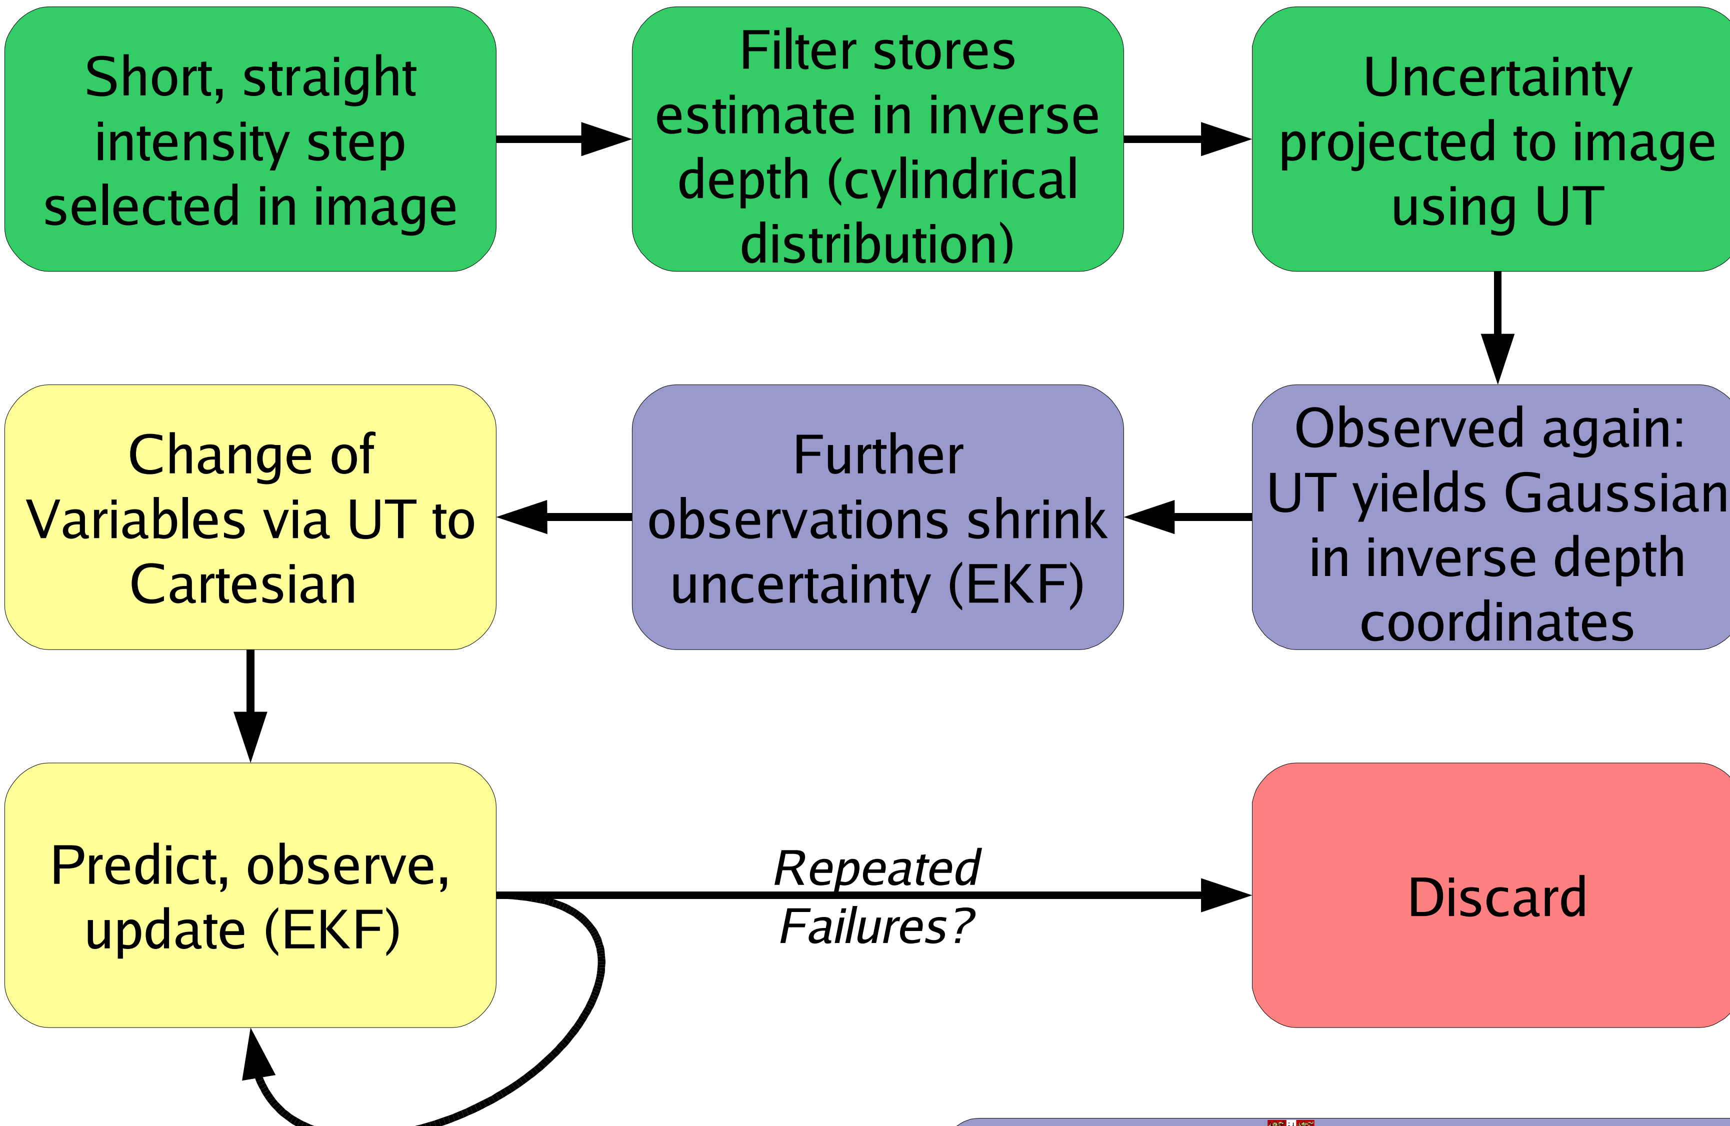

Ethan Eade ee231@cam.ac.uk