4.3 Line Detection

Outline

- Line detection & Hough maps
- Line segment detection
- Vanishing points and Manhattan worlds

Outline

- **Line detection & Hough maps**
- Line segment detection
- Vanishing points and Manhattan worlds

The Built Environment

- The built environment is largely piecewise planar.
- Boundaries, creases and surface markings thus often project as straight lines.
- Line detection is thus a core computer vision problem.

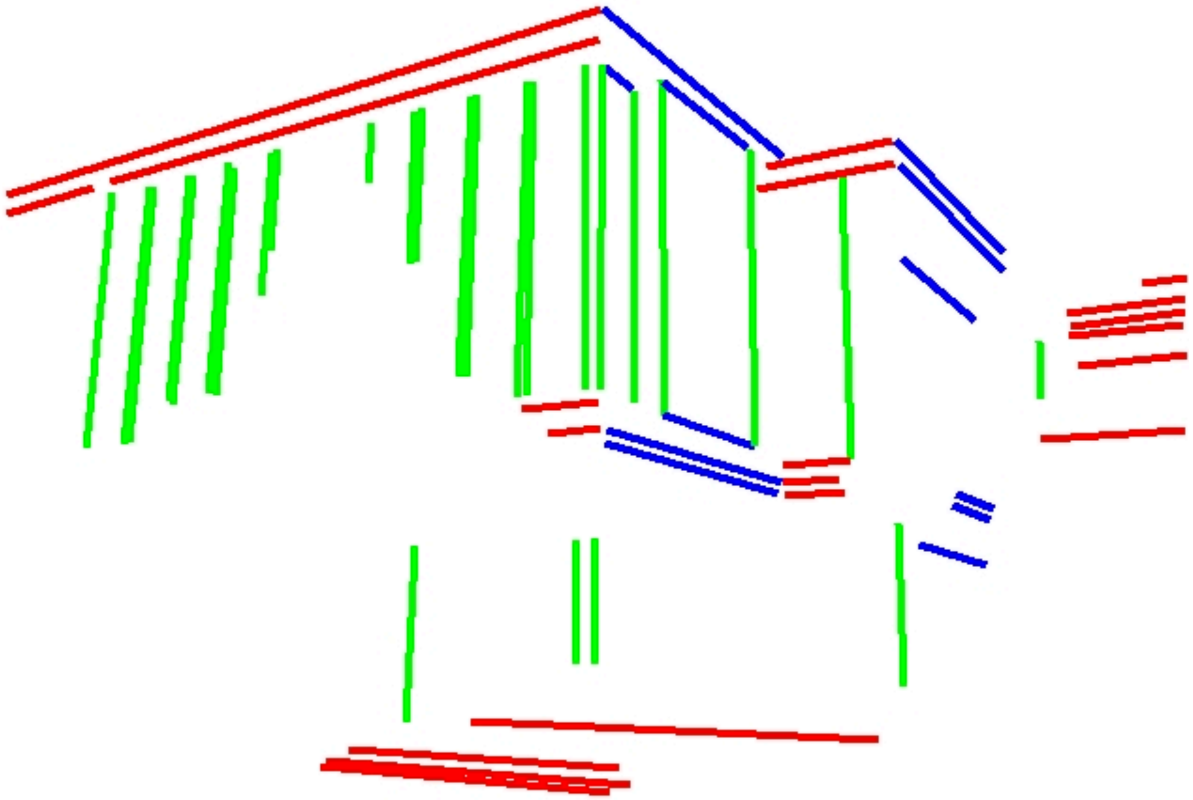

Line Detection

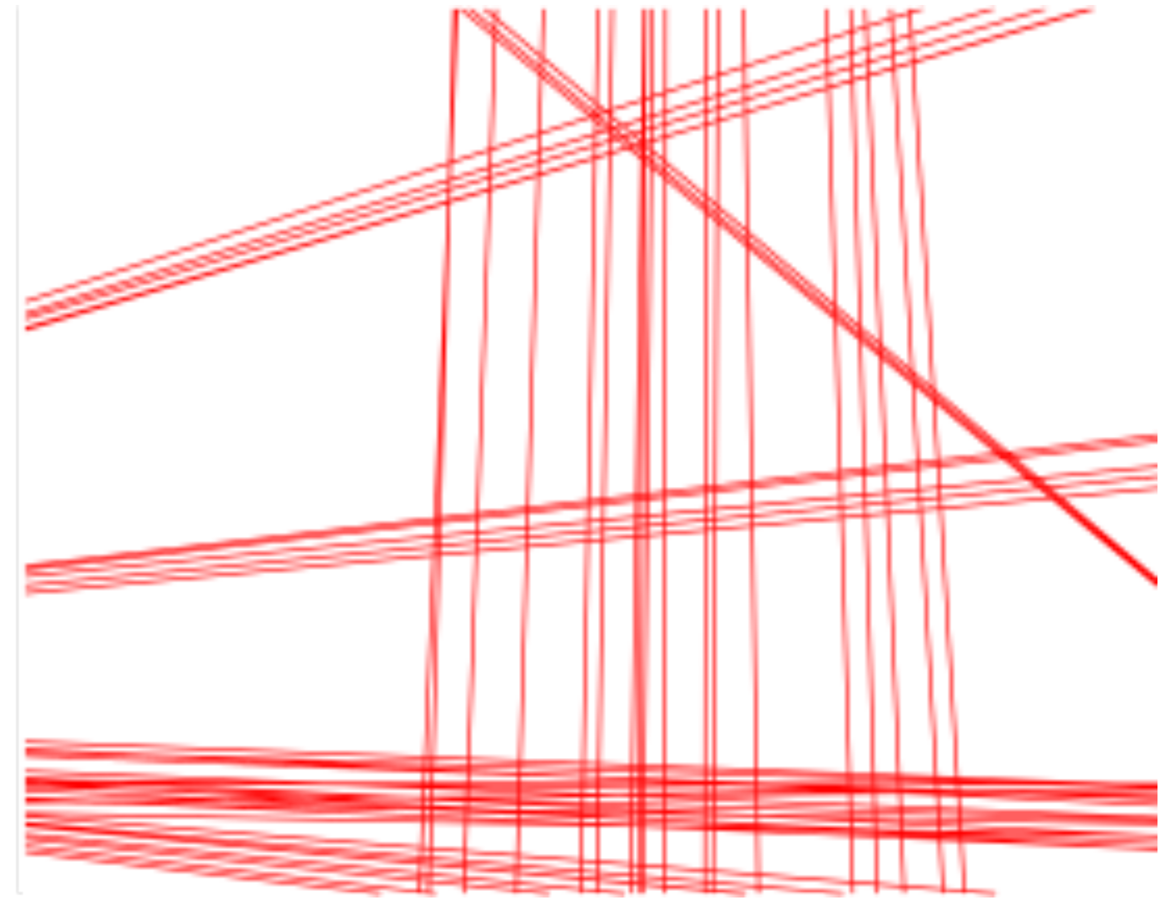

The Hough Transform

- We represent a line by the angle θ of its normal vector from the positive x axis and its *signed* distance ρ from the origin.
- We are free to choose the origin and the direction of the y axis.
  - In MATLAB, the origin is at the top left of the image and the y axis points down

$$-\pi/2 < \theta \leq \pi/2$$

$$\rho \in \mathbb{R}$$

$$\hat{n}_x x + \hat{n}_y y - d = 0$$

$$\cos\theta x + \sin\theta y - \rho = 0$$

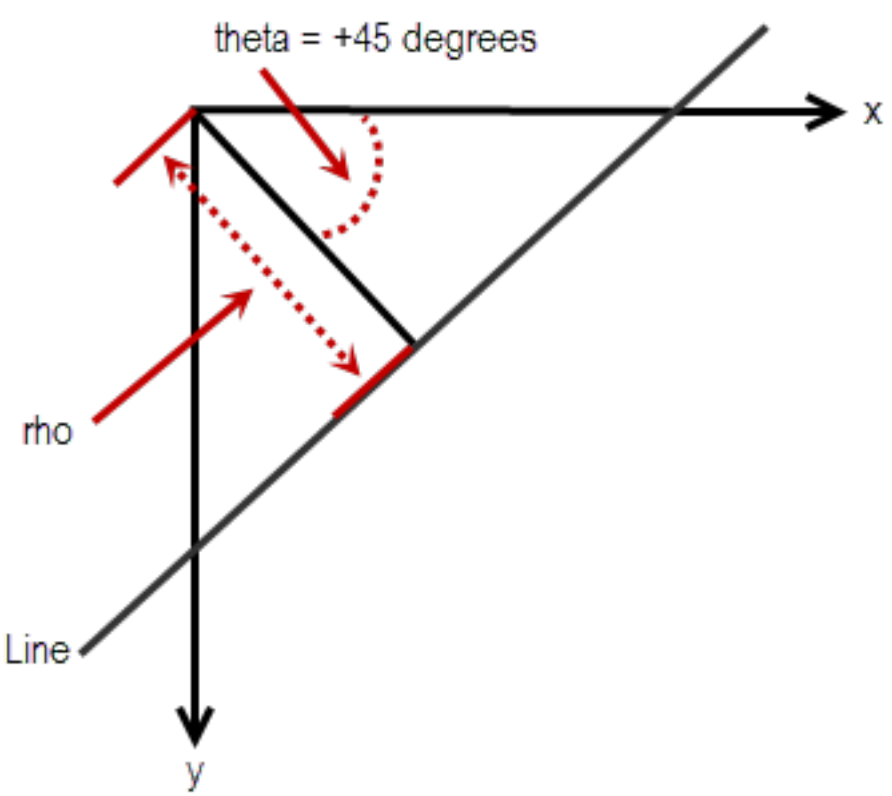

The Hough Transform

- In the textbook, the origin is at the centre of the image and the y axis points up.
- Also, Szeliski defines the normal vector to point in the direction of the luminance gradient, which means that $-\pi < \theta \leq \pi$.
- This is unconventional.

$$\hat{n}_x x + \hat{n}_y y - d = 0$$

$$\cos\theta x + \sin\theta y - \rho = 0$$

$$\rho \in \mathbb{R}$$

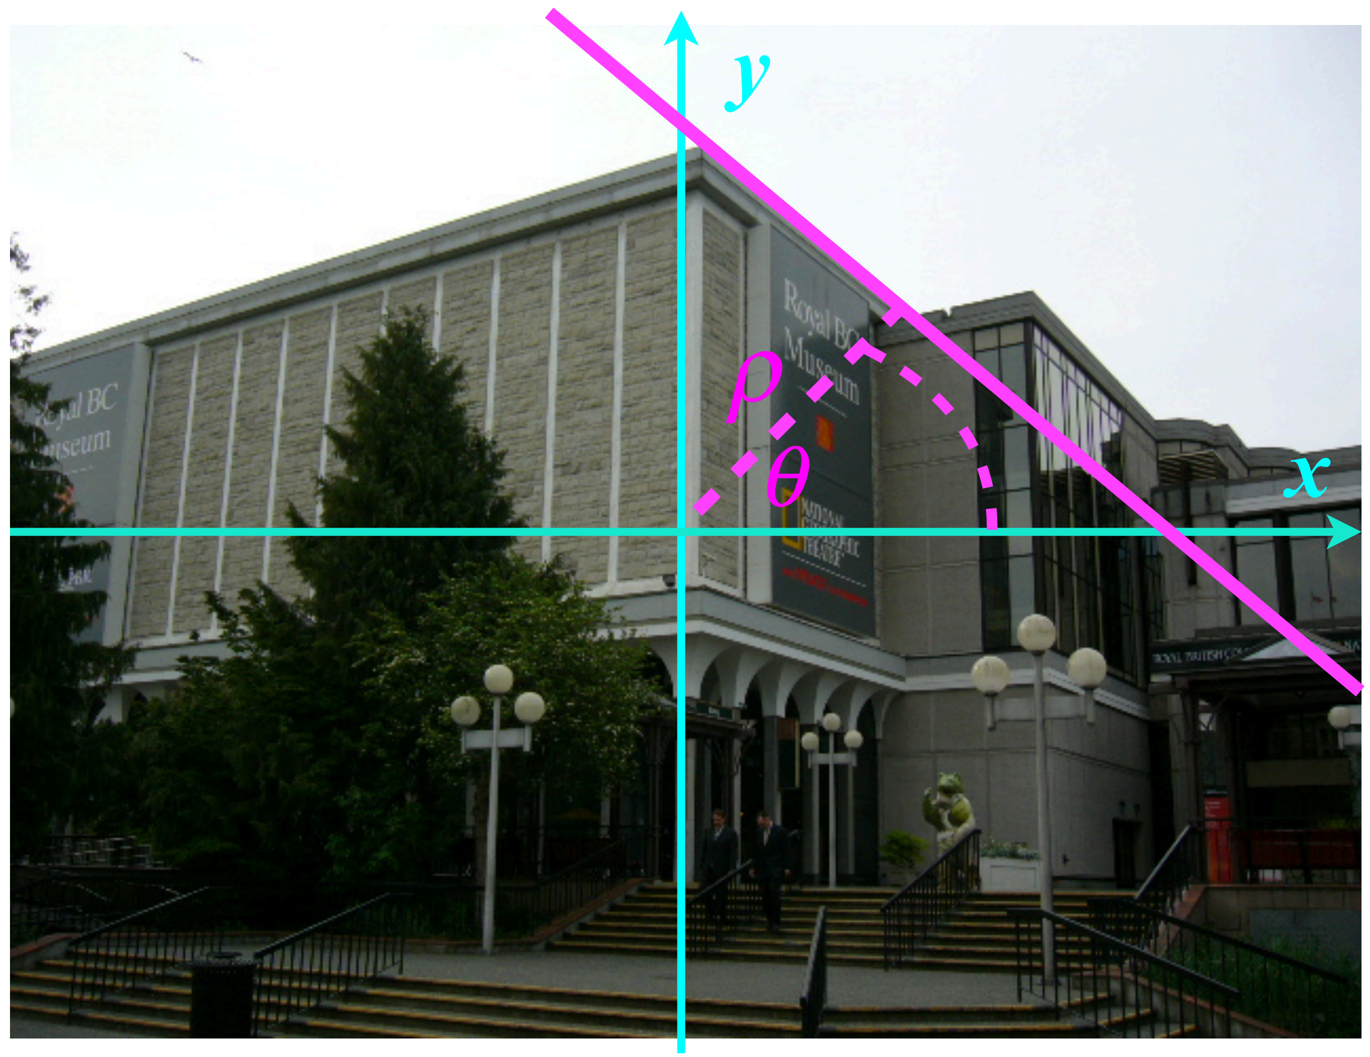

Houghing Points

- Note that a point (x, y) in the image maps to a sinusoid in Hough space.

$$\cos\theta x + \sin\theta y - \rho = 0$$

$$\text{Let } x = r\cos\theta_0, y = r\sin\theta_0$$

$$\text{Then } \rho = r\cos\theta\cos\theta_0 + r\sin\theta\sin\theta_0 = r\cos(\theta - \theta_0)$$

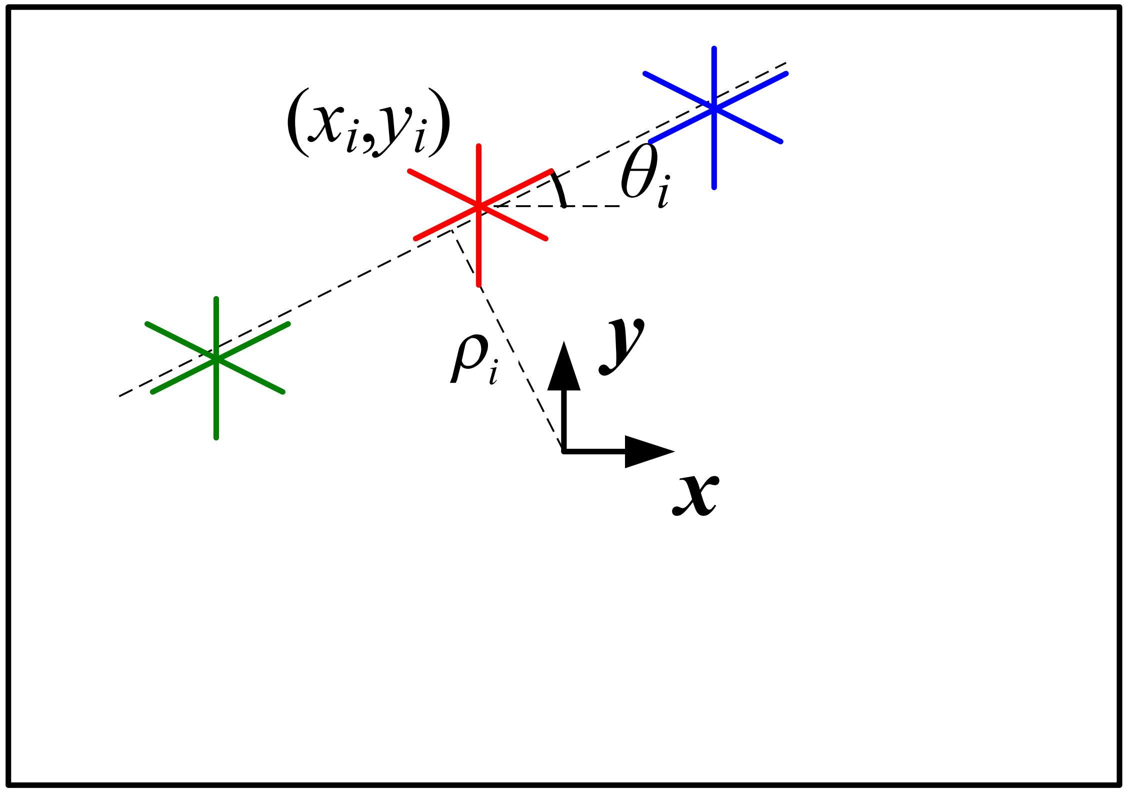

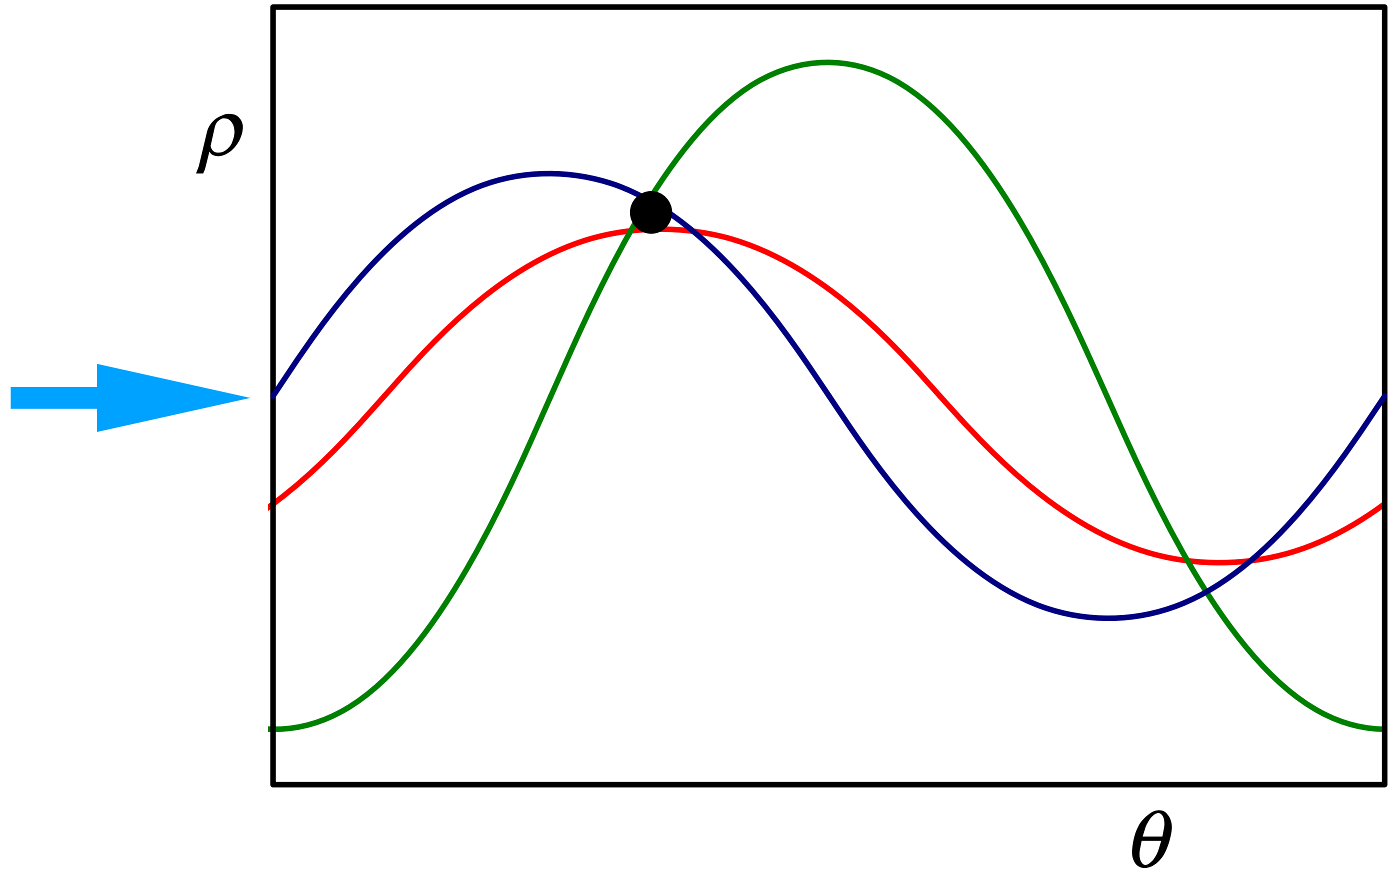

Houghing Points

- Each point on a line in the image generates a different sinusoid in the Hough map.
- The intersection of these sinusoids in the Hough map identifies the line.

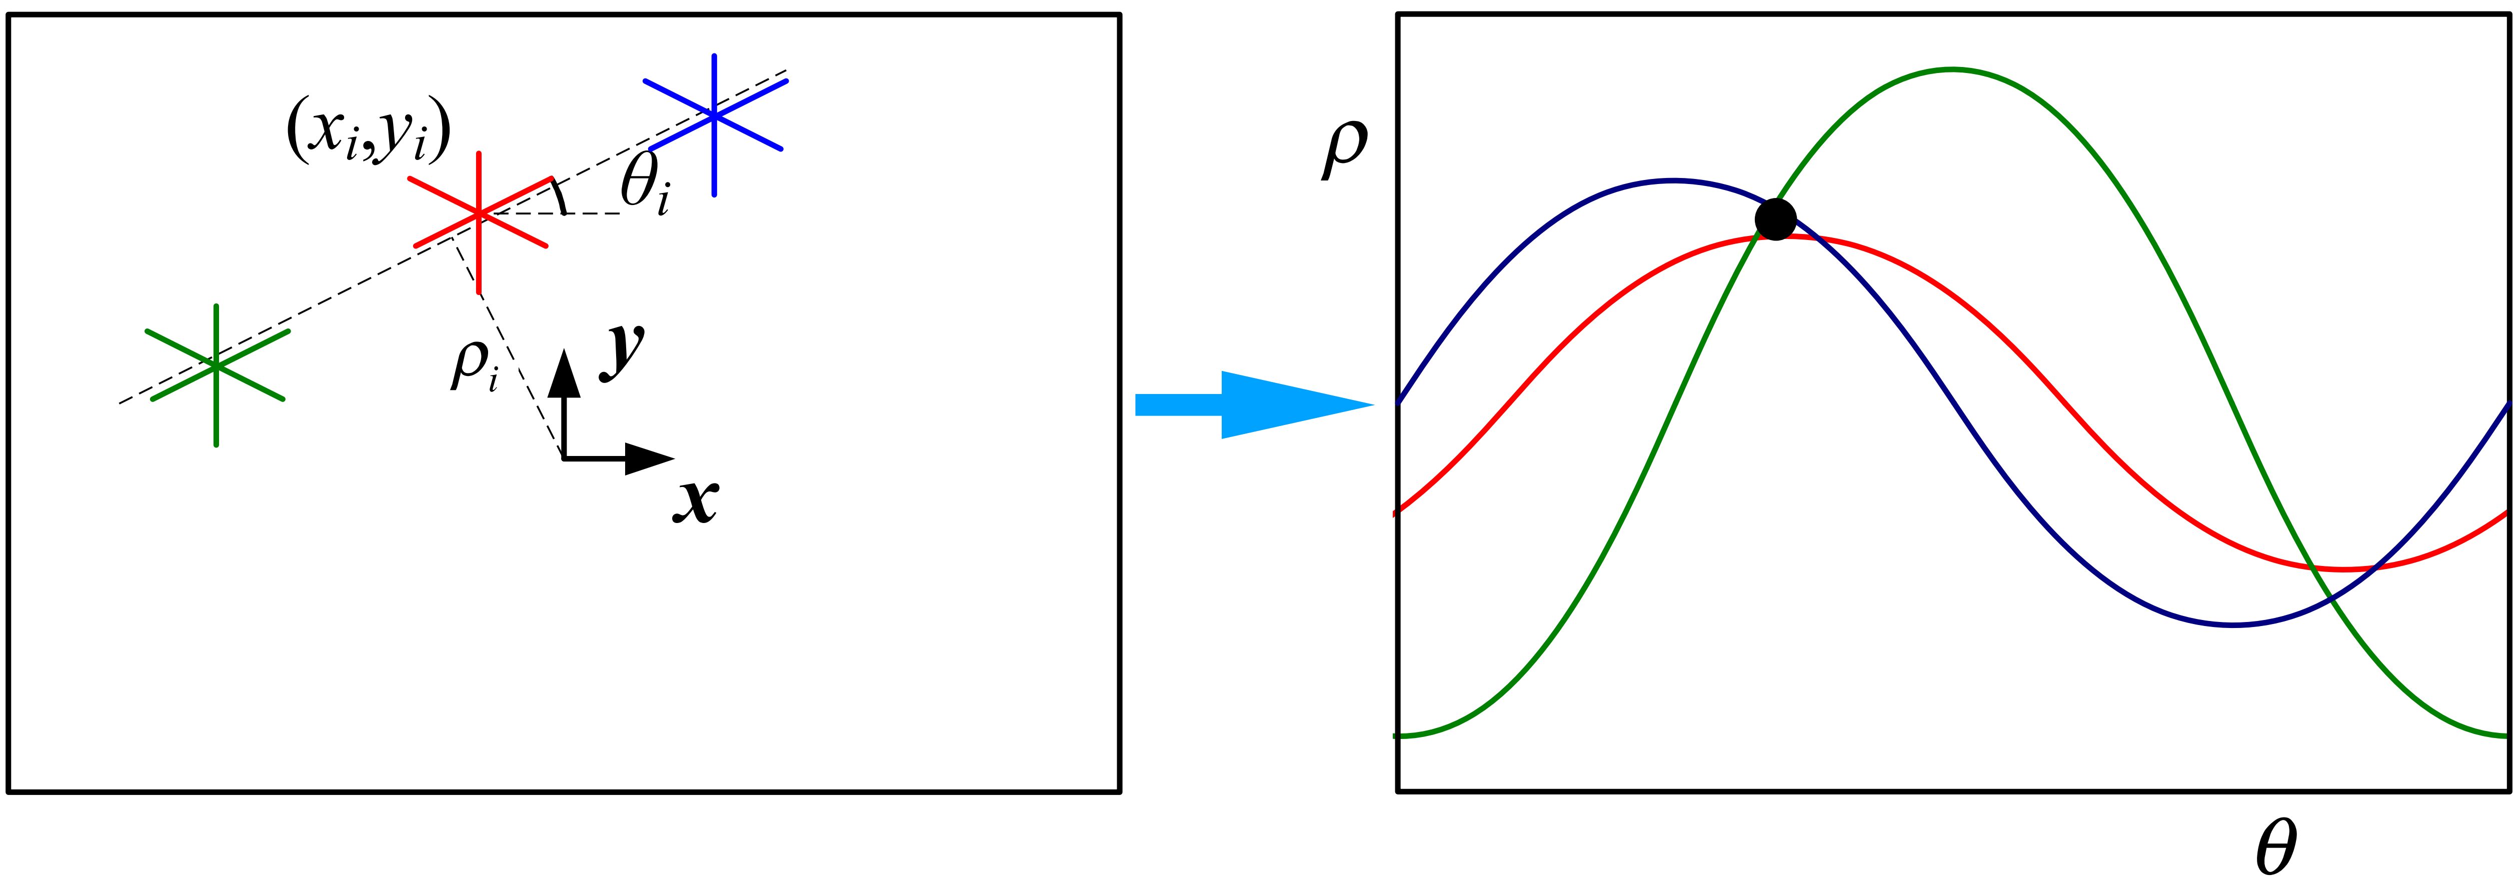

Houghing Edges

- Edge detectors estimate the location and also the orientation of the edge.
- If the orientation estimate is perfect, the edge exactly identifies the line as a point in the Hough map.
- In practice, noise in the location and orientation of the edges on a line results in a distribution of 'votes' in the Hough map.

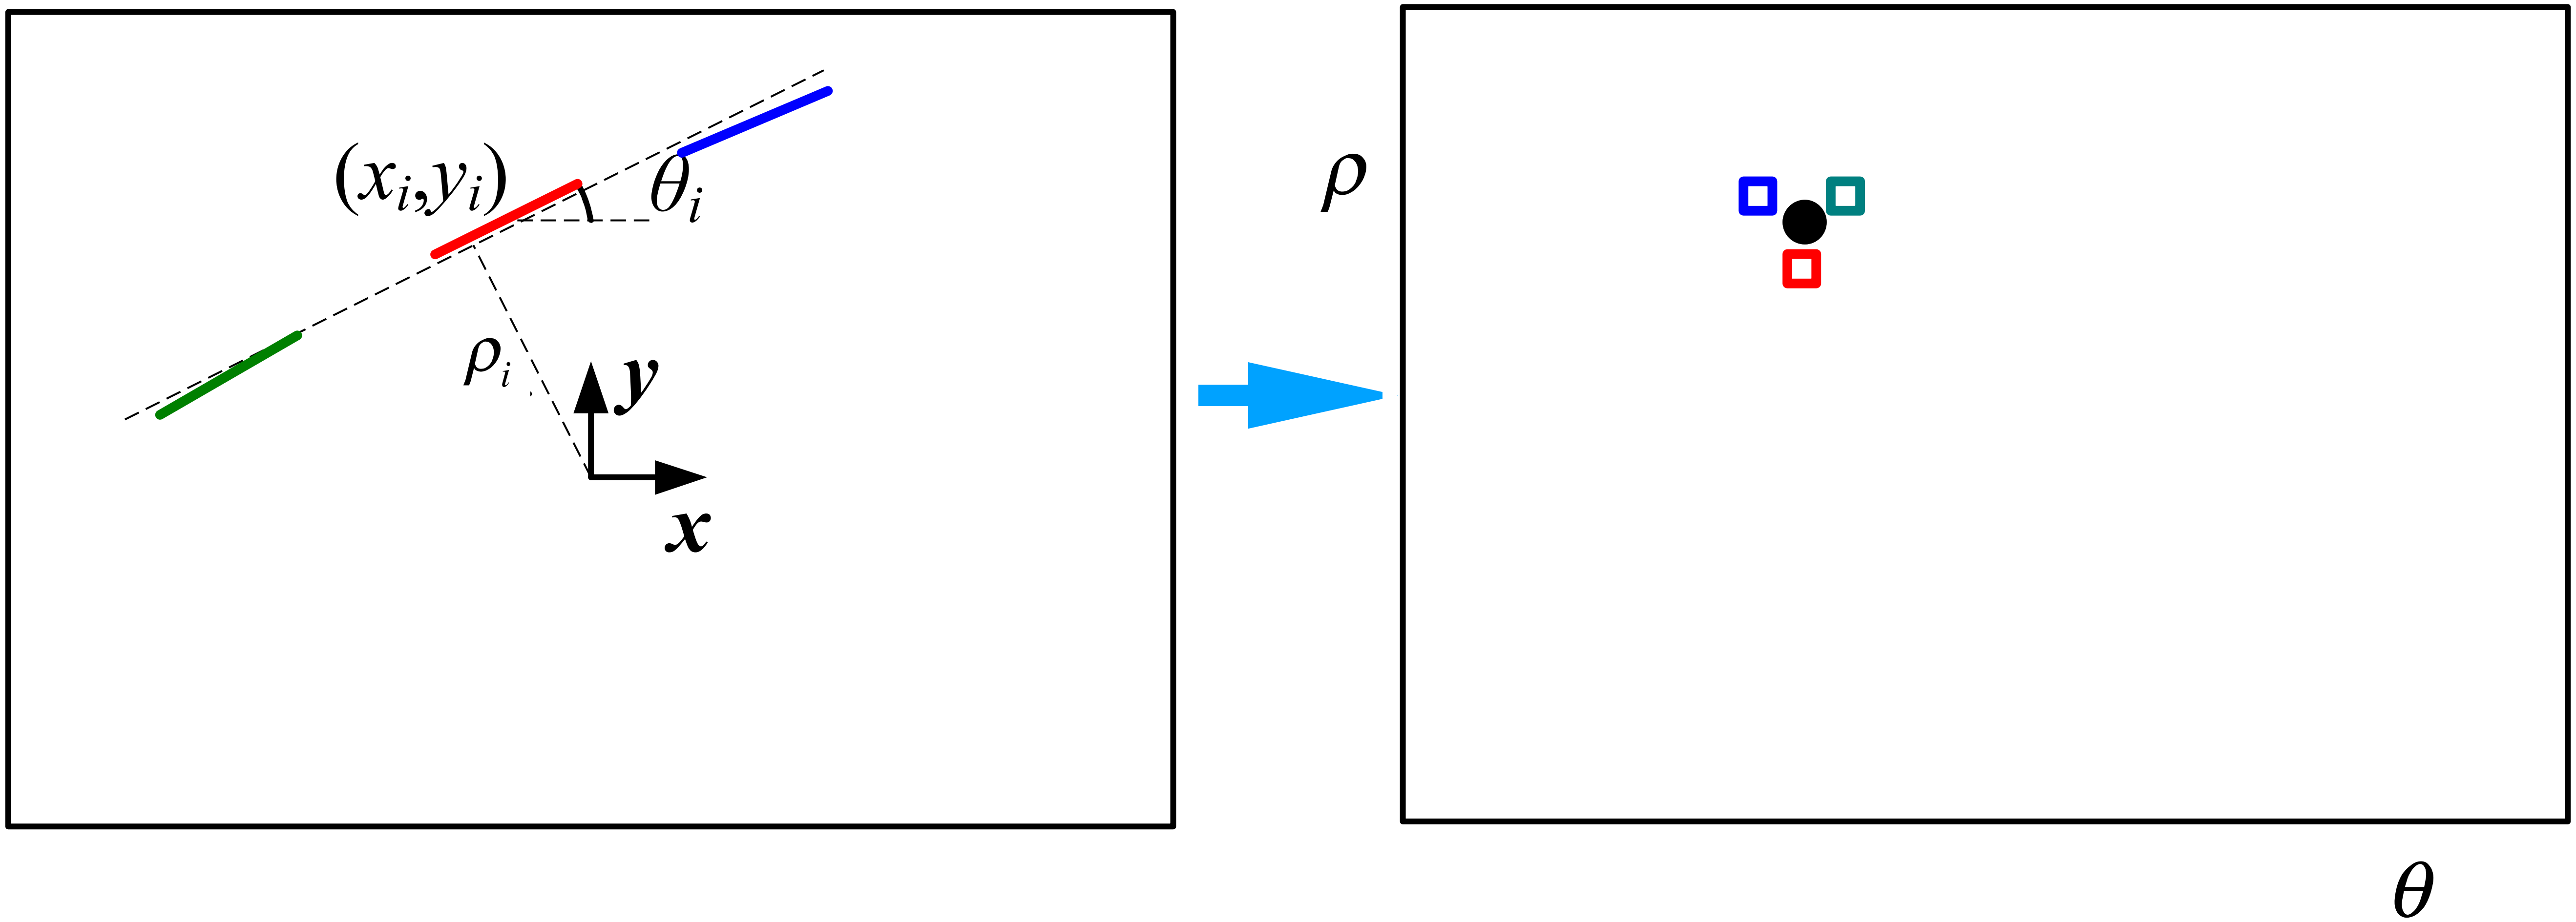

Probabilistic Houghing

- The main challenge with Hough methods is that this distribution may generate multiple peaks in the Hough map, leading to the detection of spurious lines.
- Probabilistic Houghing (Tal & Elder 2012) solves this problem by:
  - Propagating uncertainty in location and orientation of edges to the Hough map.
  - Subtracting edges contributing to a peak from the Hough map when a line is detected.

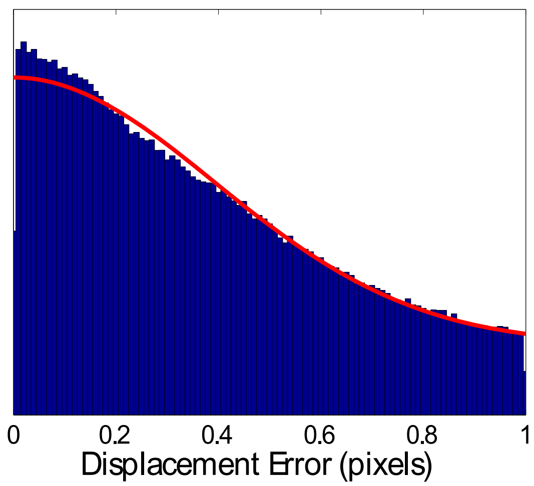

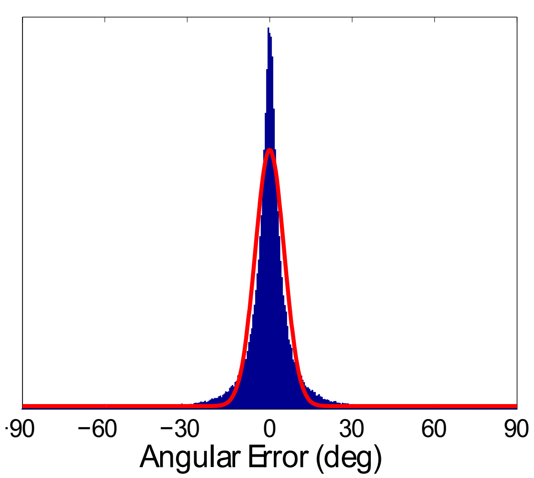

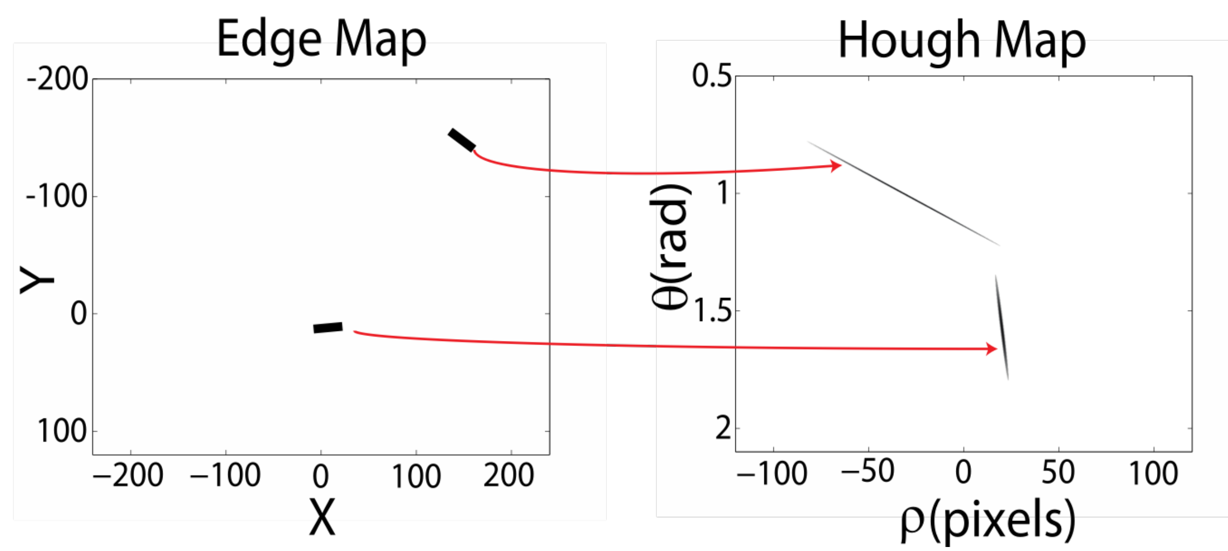

Tal & Elder, ACCV 2012

Ron Tal J. Elder

Edge Subtraction

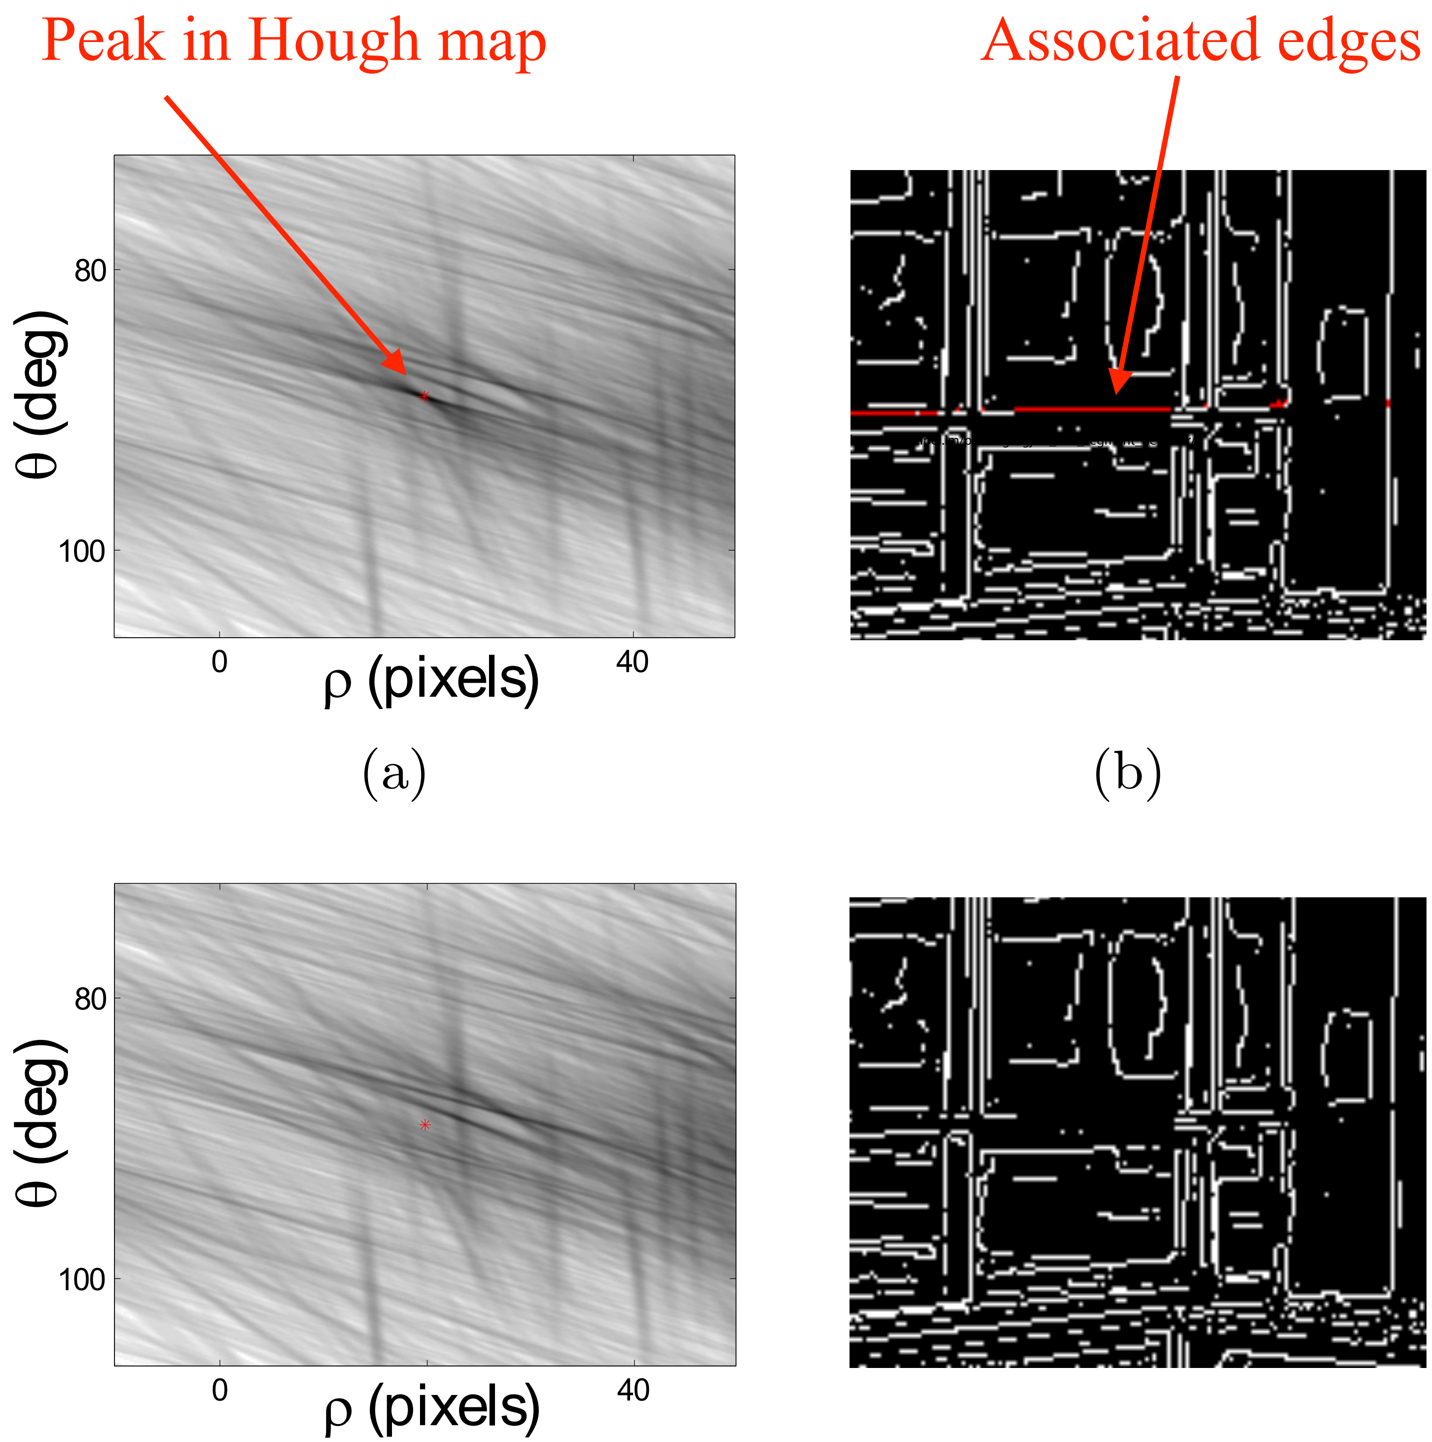

Hough and edge maps after subtraction

Quantitative Results

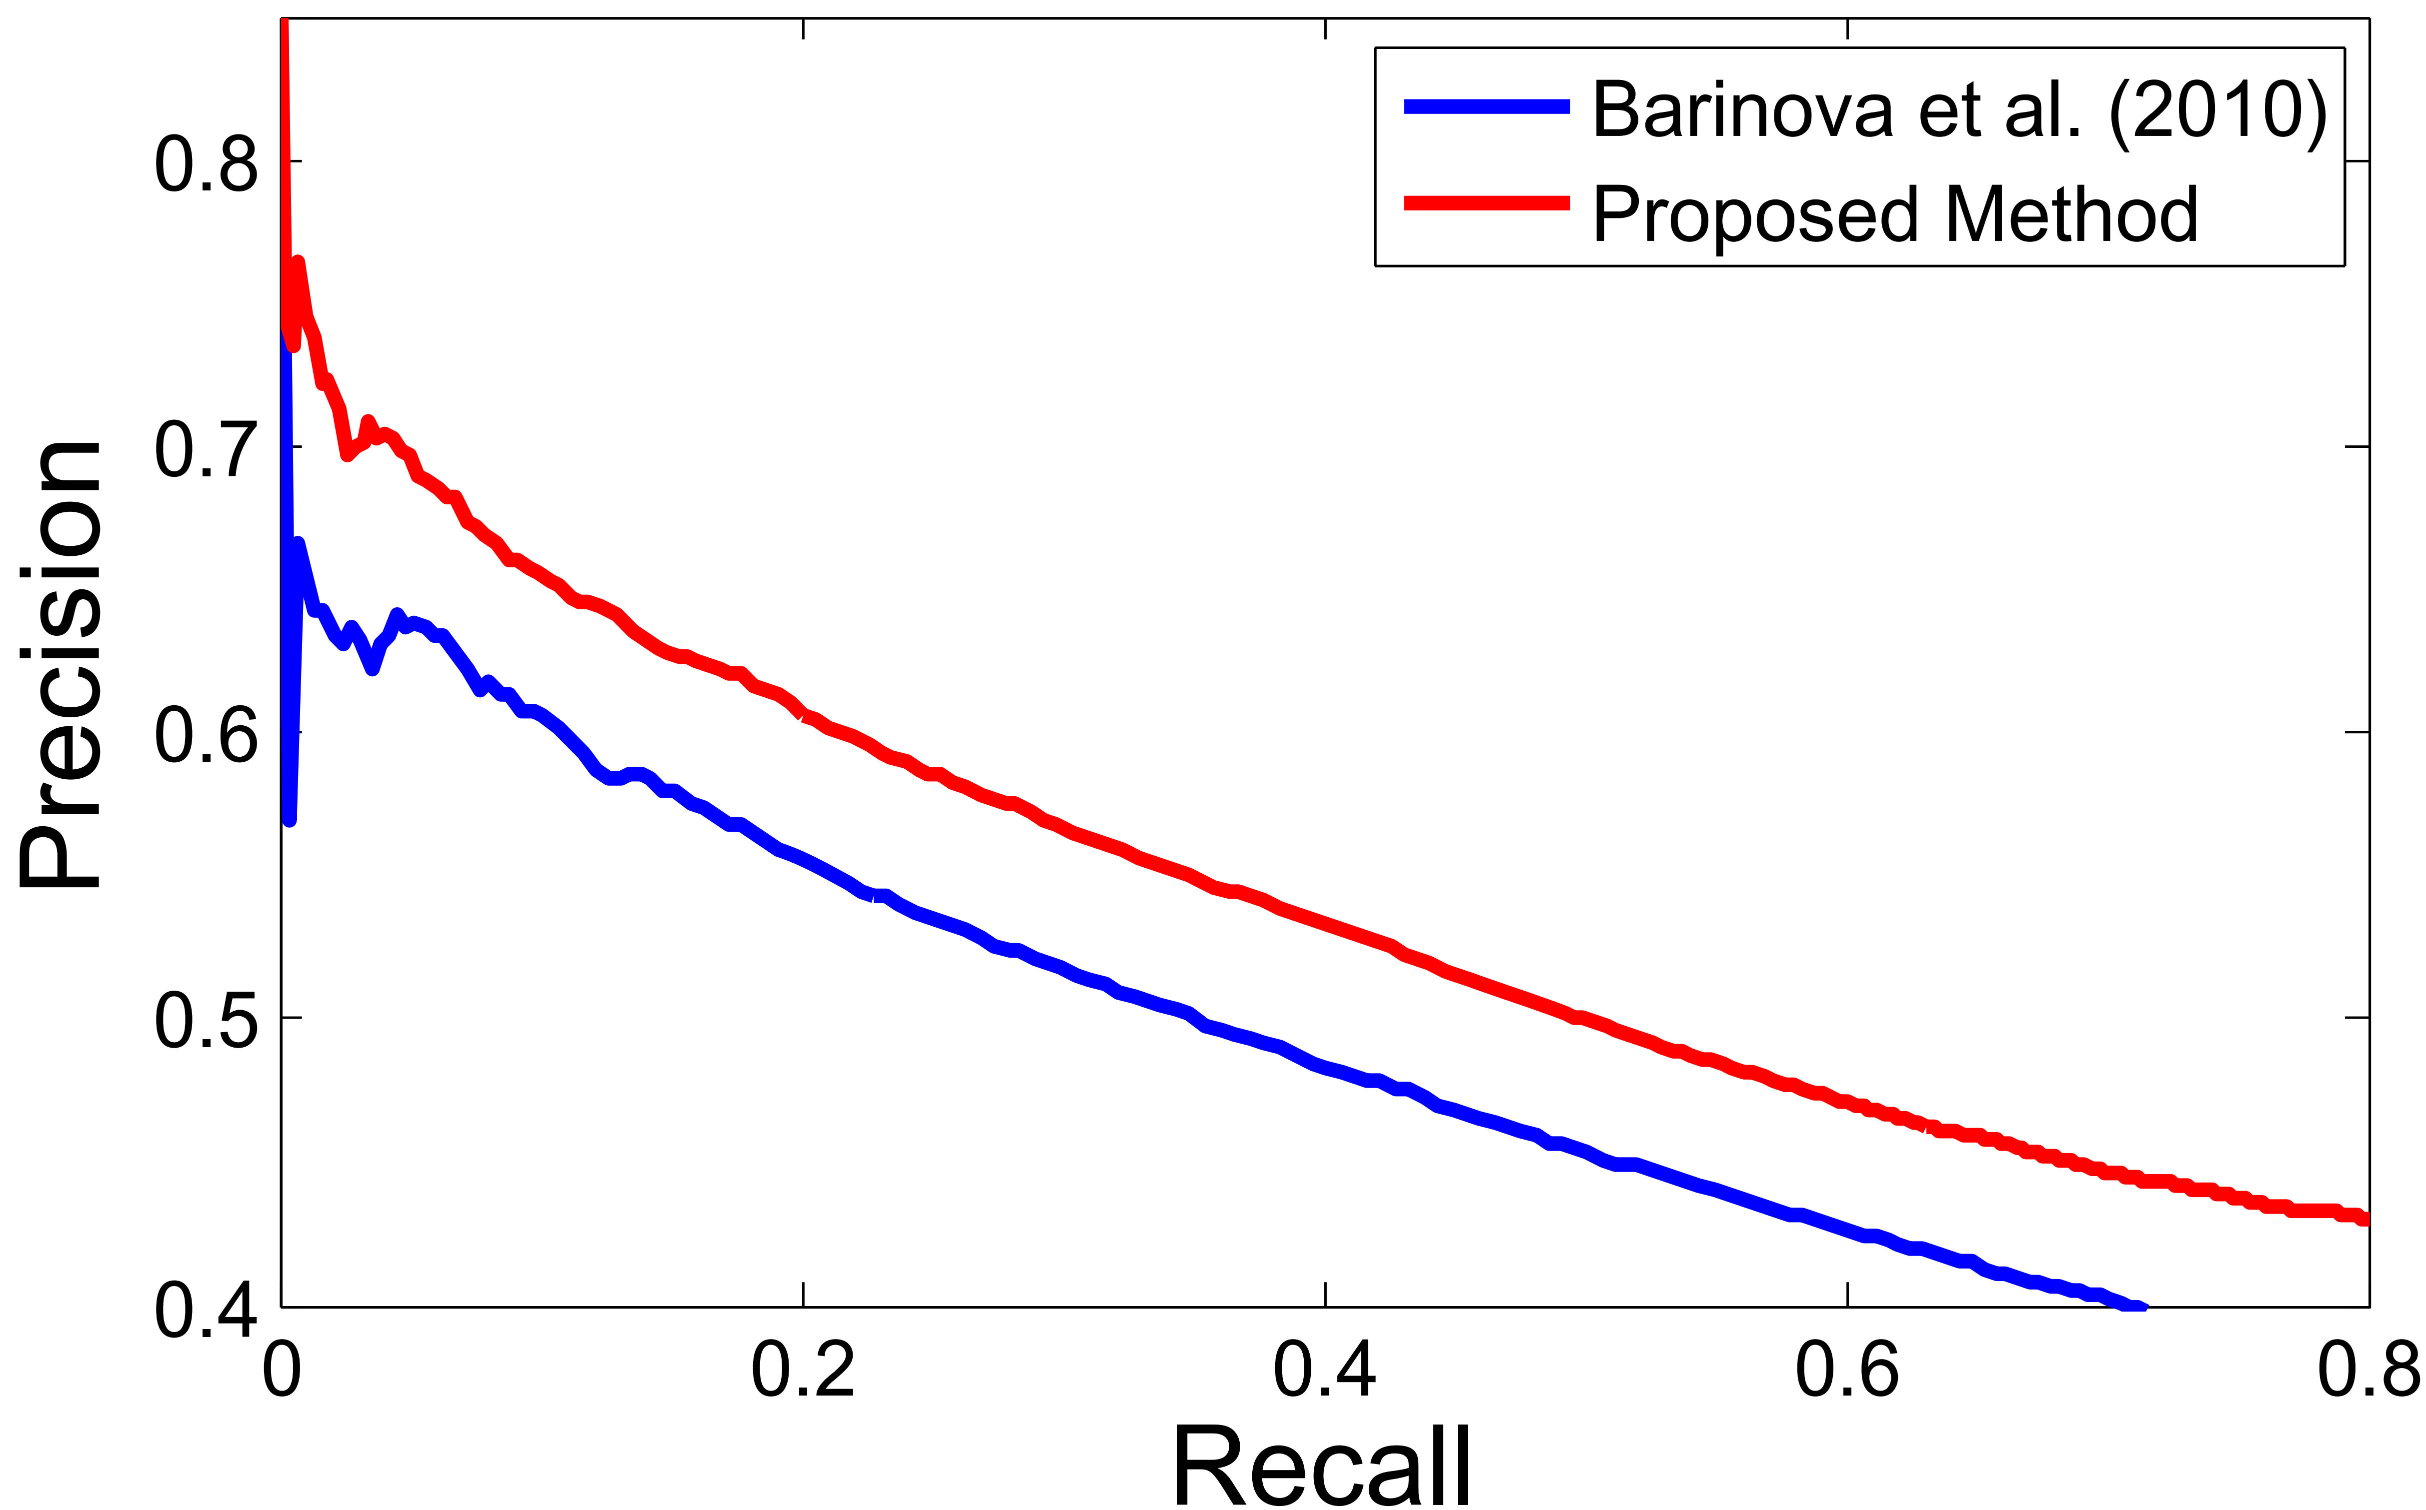

Qualitative Results

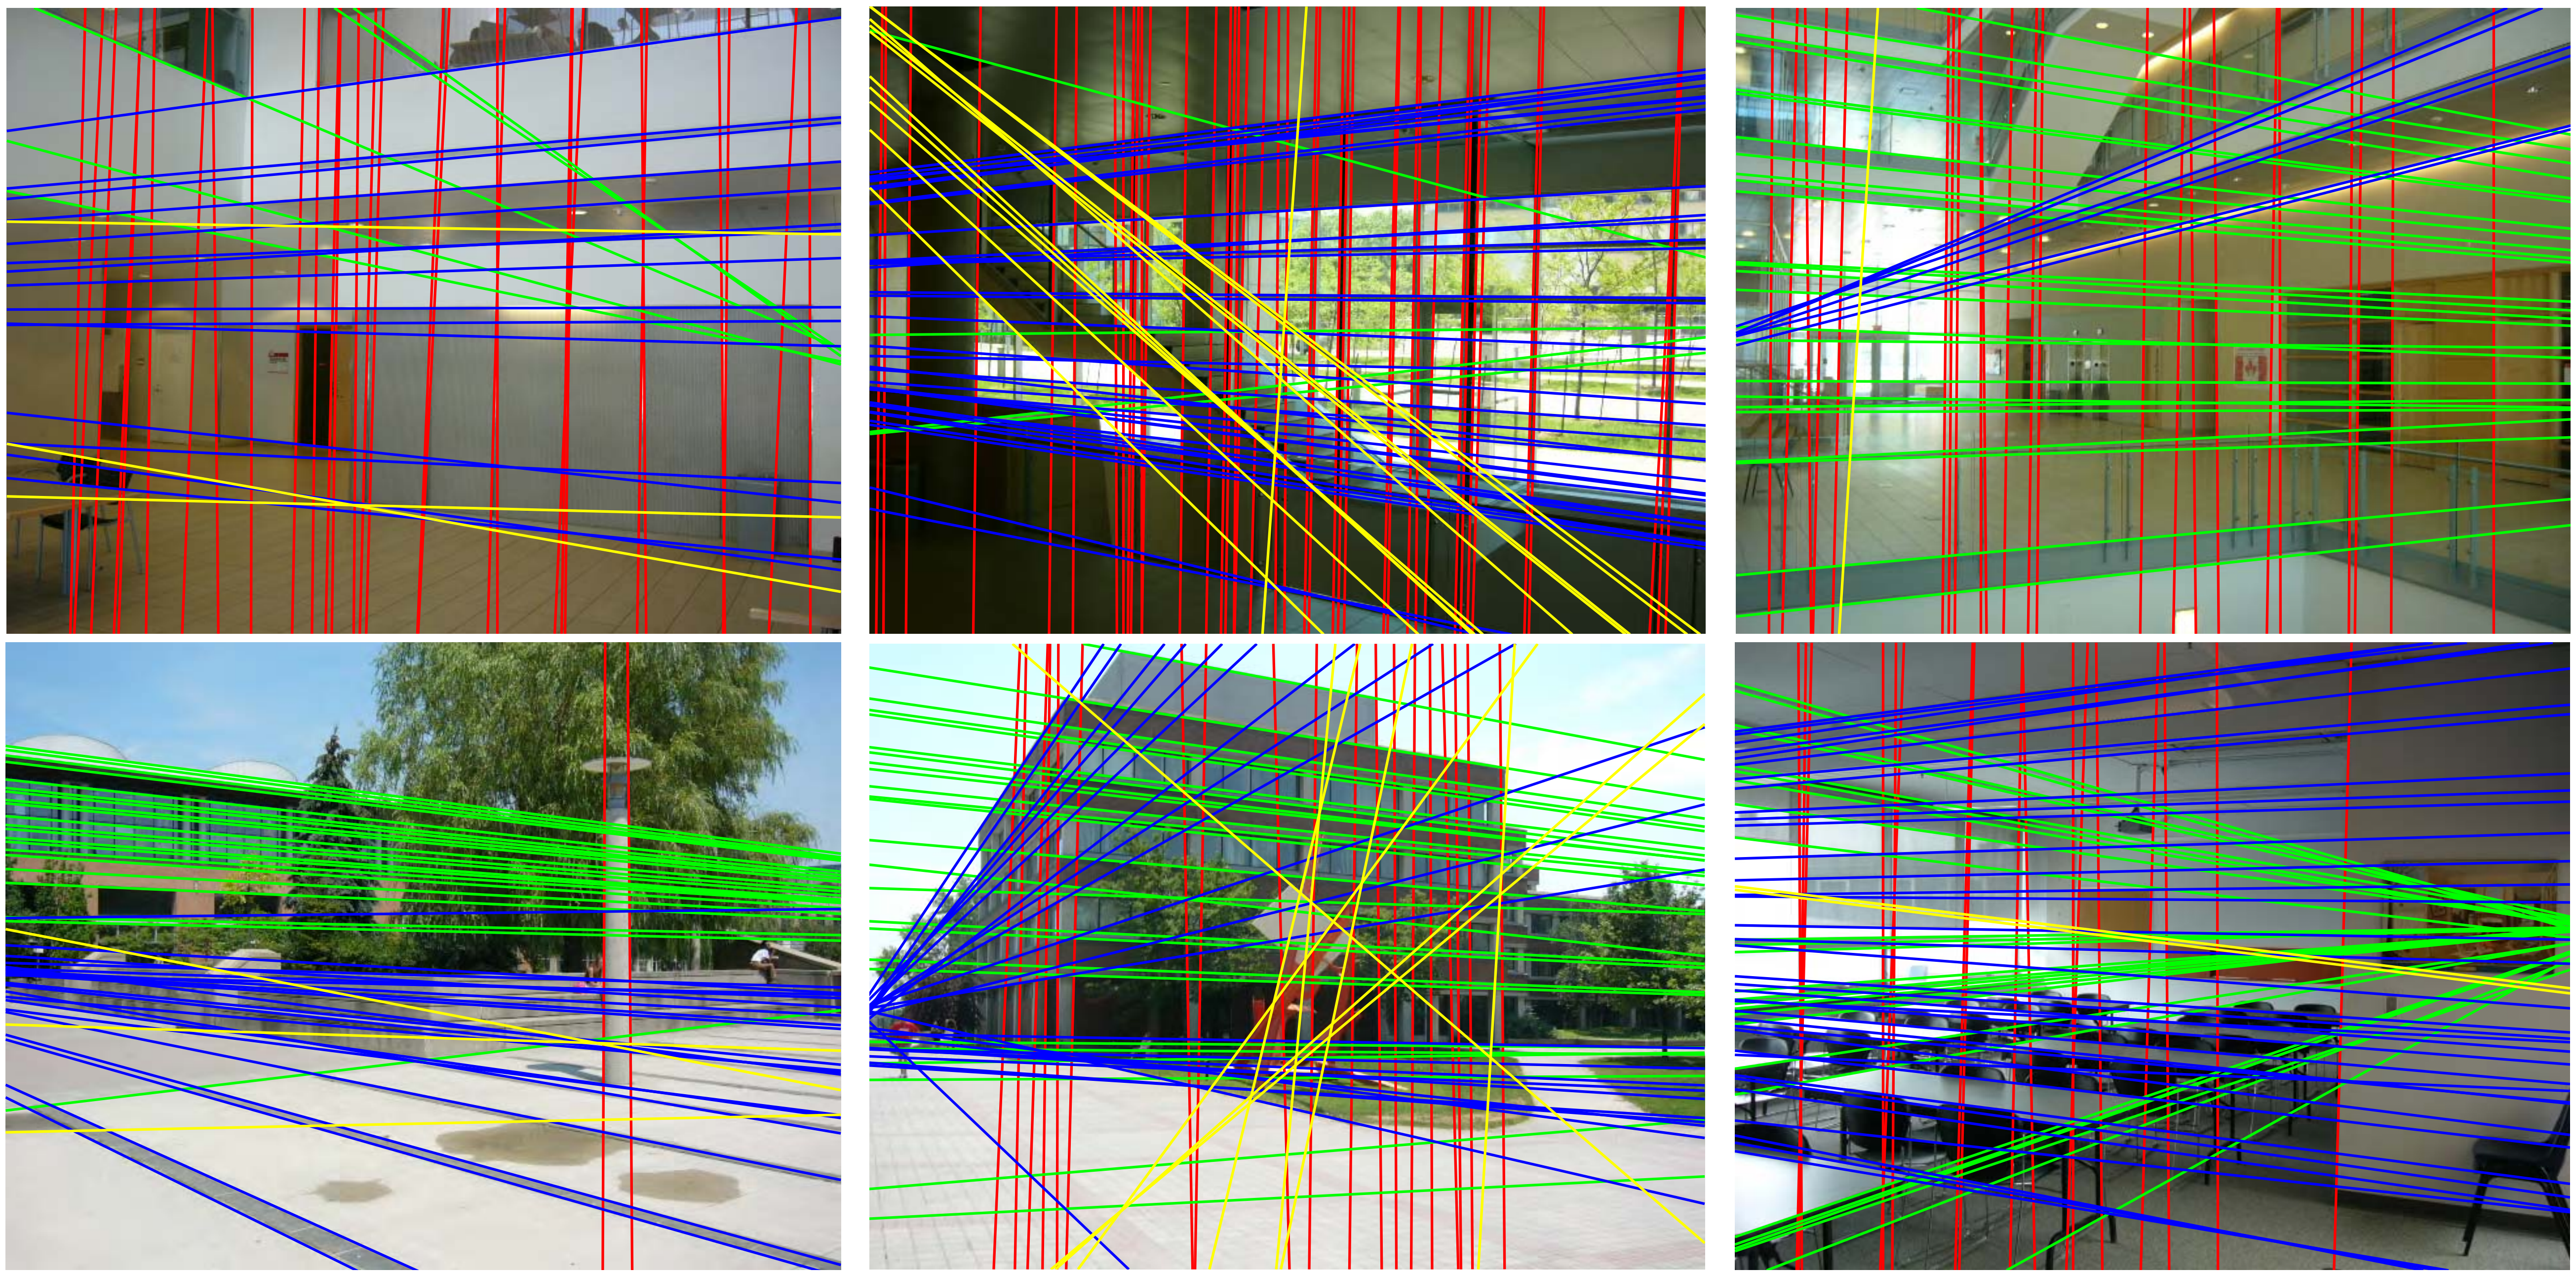

Outline

- Line detection & Hough maps
- **Line segment detection**
- Vanishing points and Manhattan worlds

Step 2: Line Segment Detection

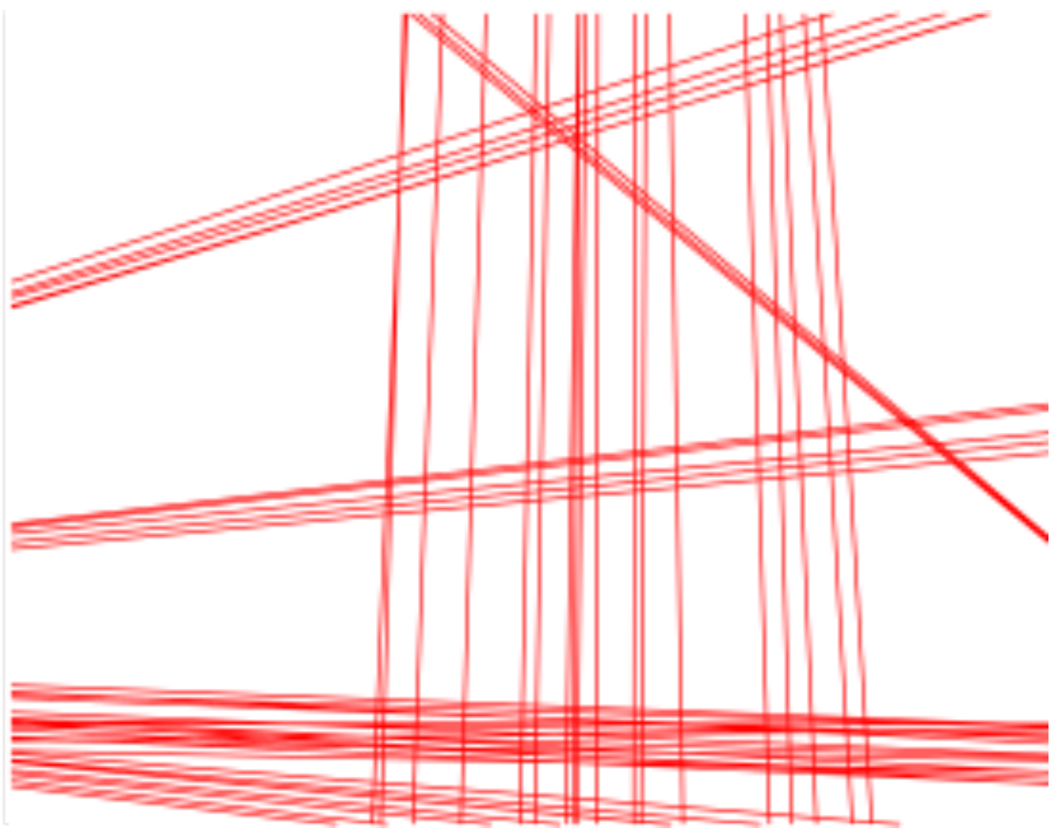

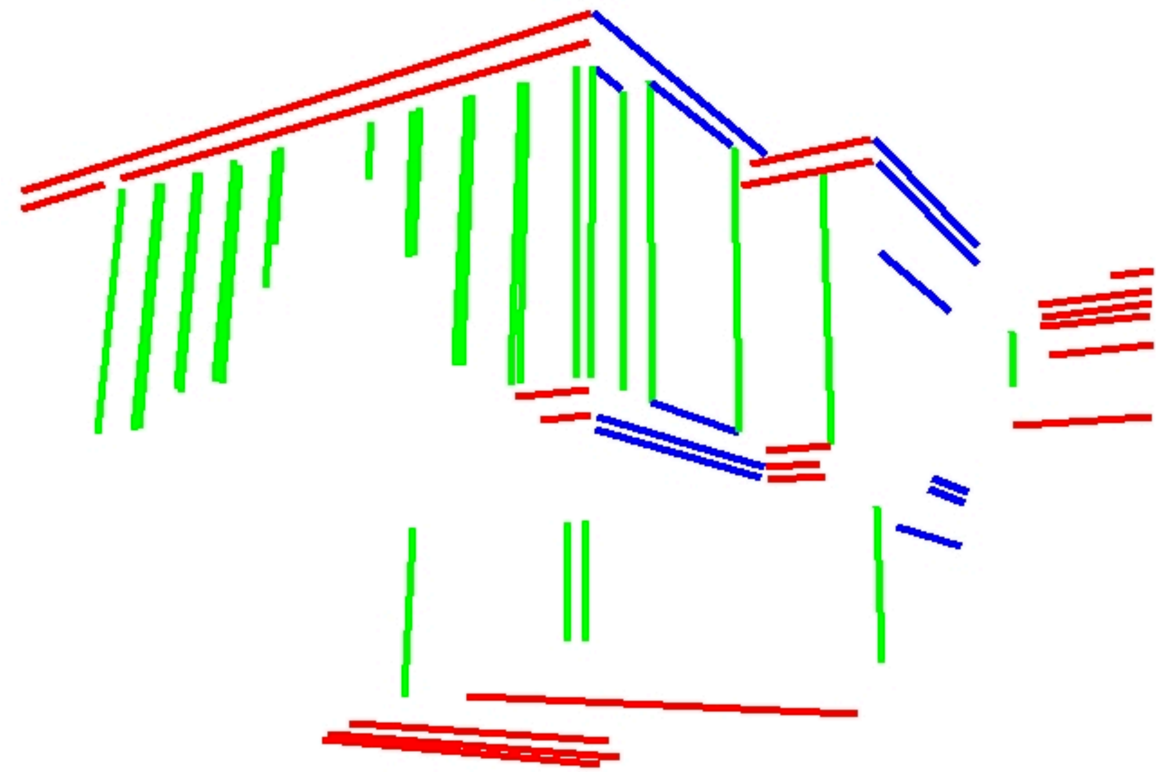

Almazen et al, CVPR 2017

Emilio Almazen Nielsen

Line Segmentation Detection - MCMLSD

- Prior approaches tend to be either global (Hough methods) or local (perceptual grouping).
- Global methods can yield more accurate line extraction by accumulating more complete evidence.
- Local methods tend to be better at localizing segment endpoints.
- MCMLSD† combines the virtues of the two:
  - **Step 1.** Use probabilistic Hough method to identify global lines.
  - **Step 2.** Partition each line into maximum probability segments, using dynamic programming, in linear time.

†Markov Chain Marginal Line Segment Detector

Dynamic Programming Solution

- Factoring of posterior → optimal substructure property → dynamic programming solution

- Max probability configuration ←→ sequence minimizing cost

- Can be computed sequentially in $O(N)$ time using dynamic programming.

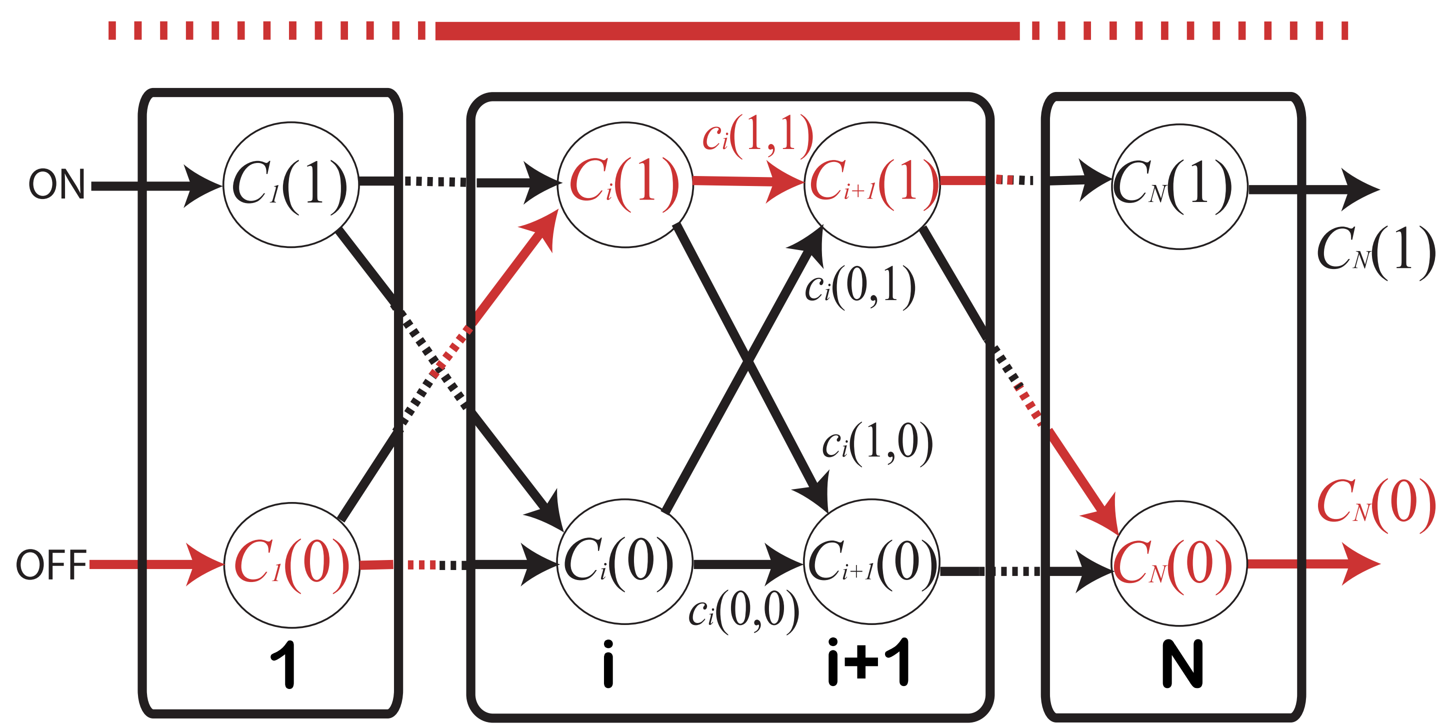

Qualitative Results

LSD (Gioi 2008)

MCMLSD

PPHT	SSWMS	LSD	MCMLSD

Quantitative Evaluation

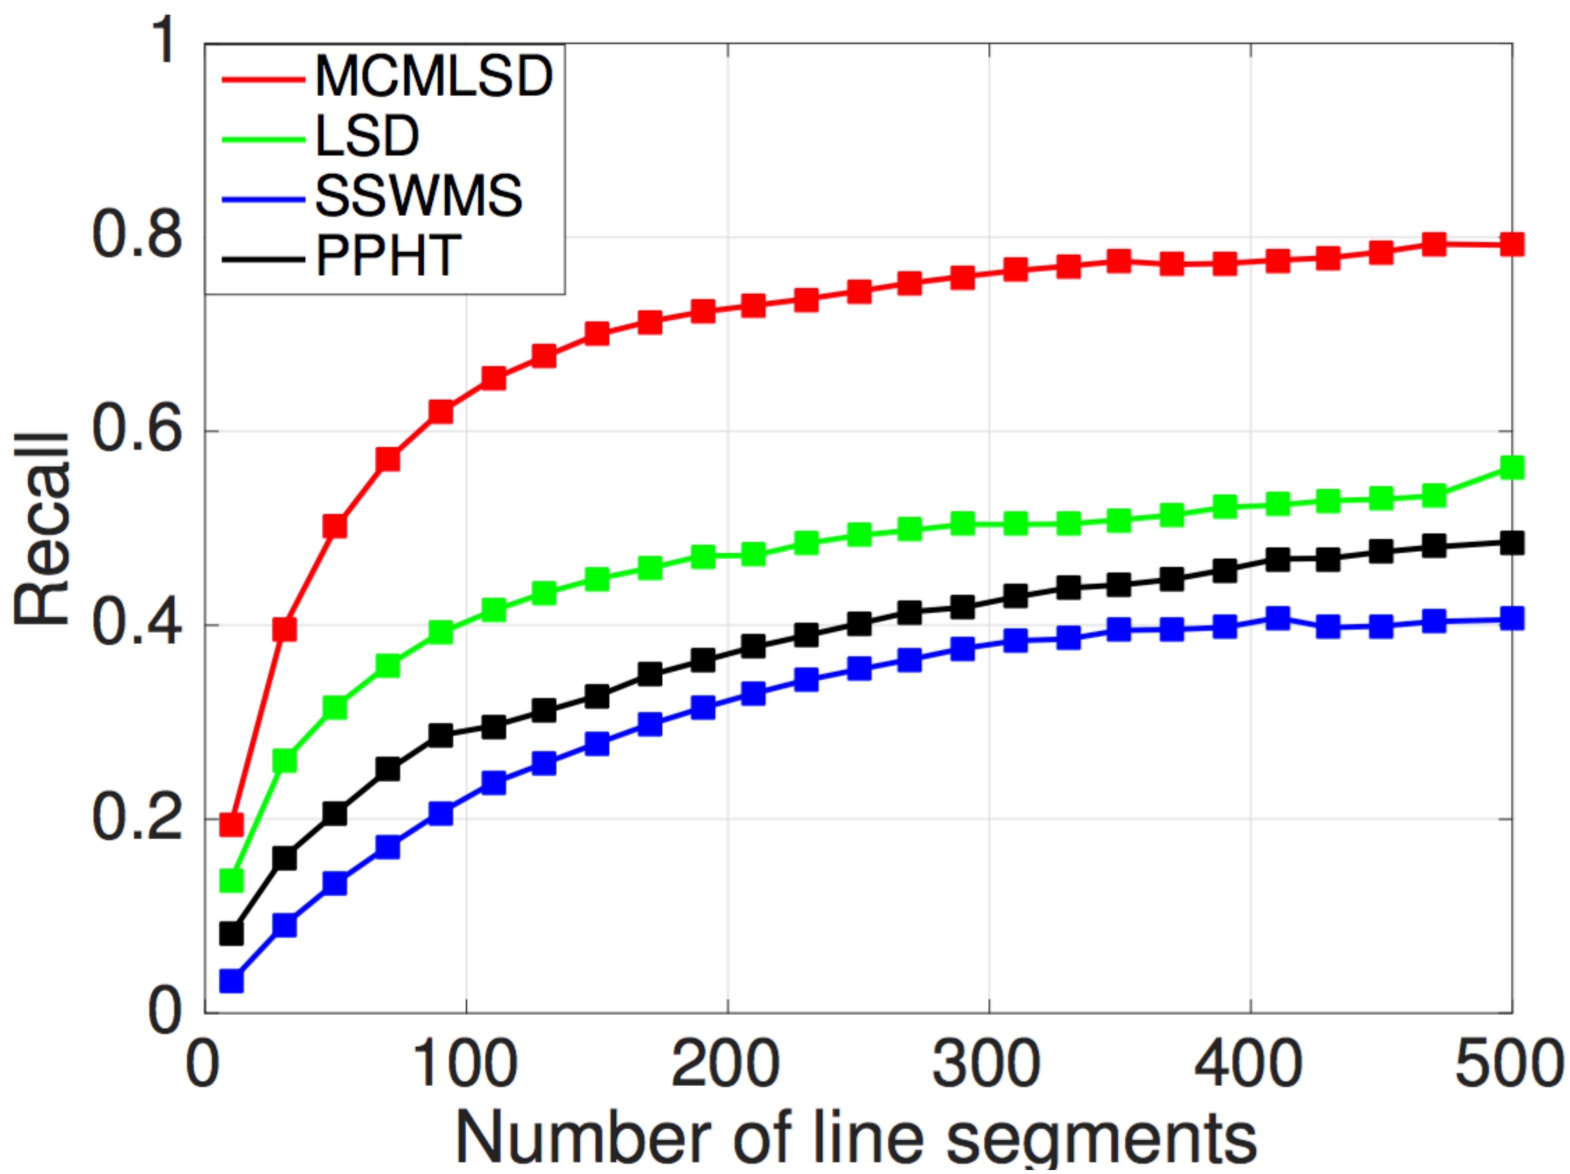

End of Lecture
Nov 5, 2018

Outline

- ❖ Line detection & Hough maps
- ❖ Line segment detection
- ❖ **Linear perspective, Vanishing points and Manhattan worlds**

Linear Perspective

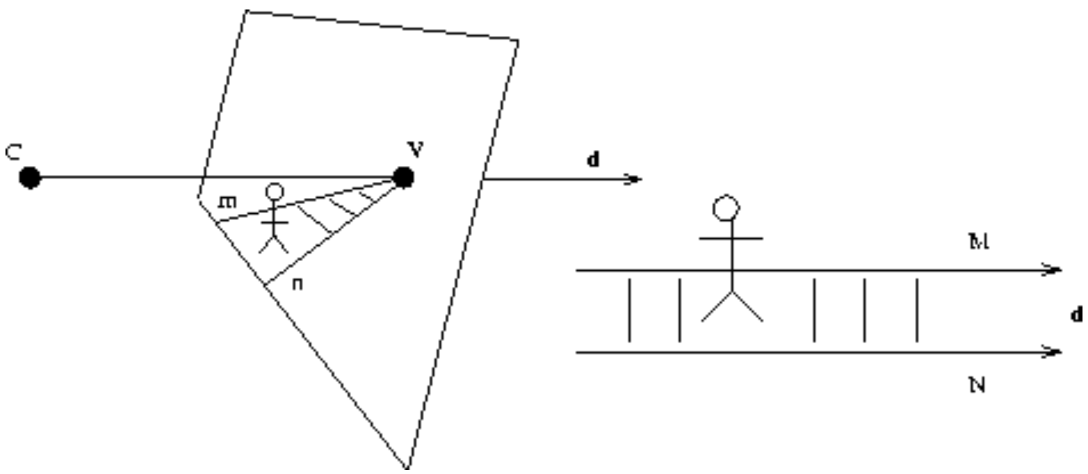

Linear Perspective in Art

- The discovery of the laws of linear perspective is a defining achievement of the early Renaissance.

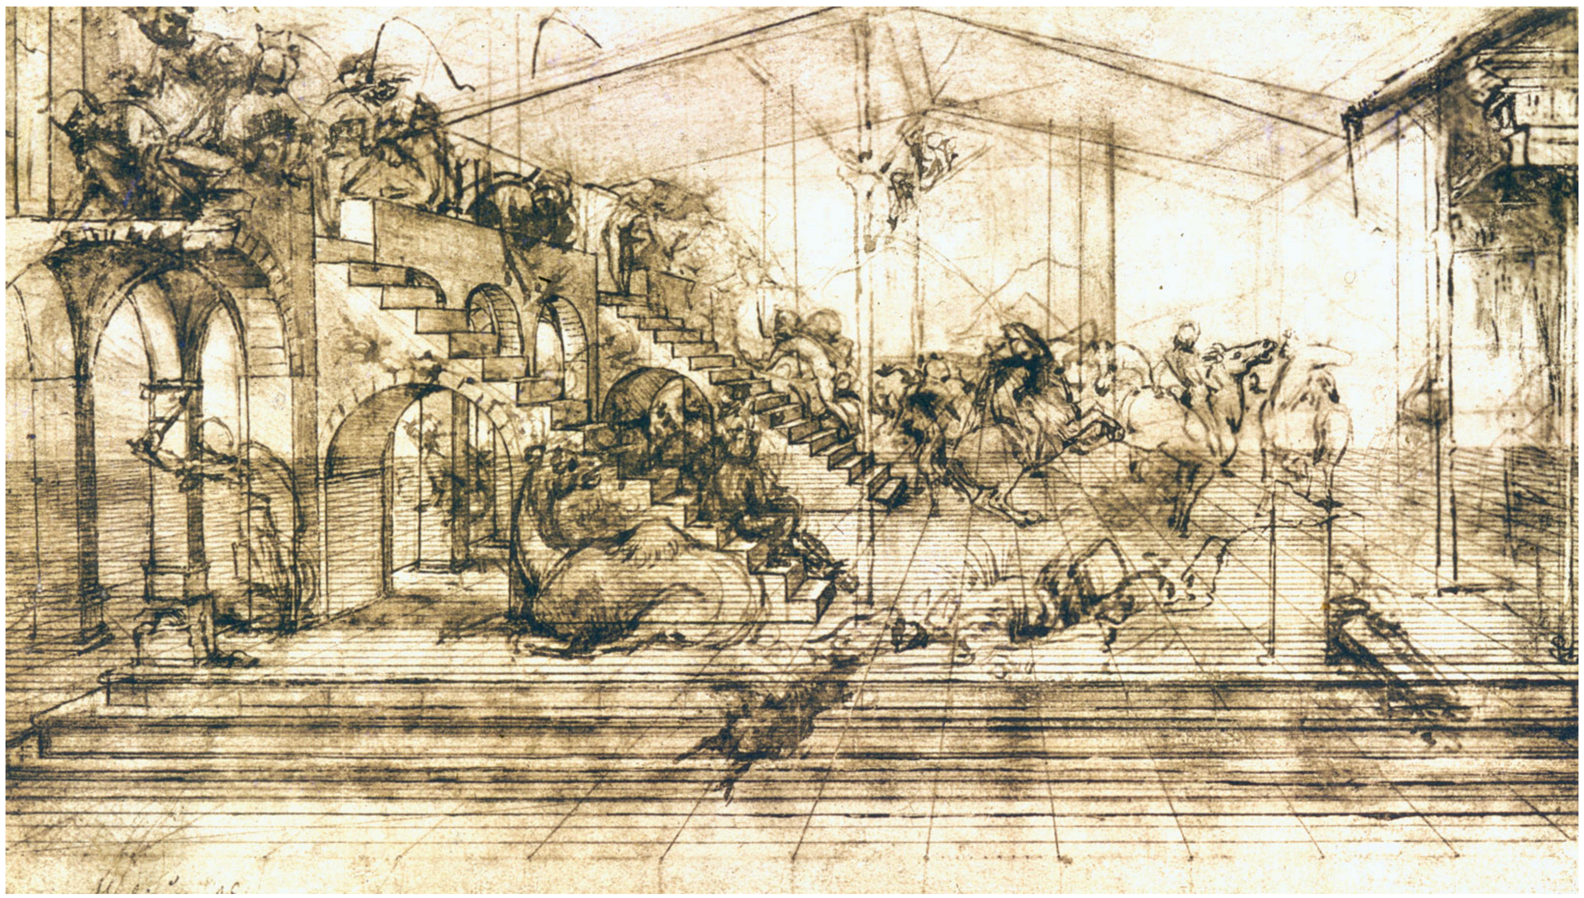

Adoration of the Magi, Leonardo da Vinci
c. 1481

Filippo Brunelleschi
1377 - 1446

The Forward Problem: Early Attempts

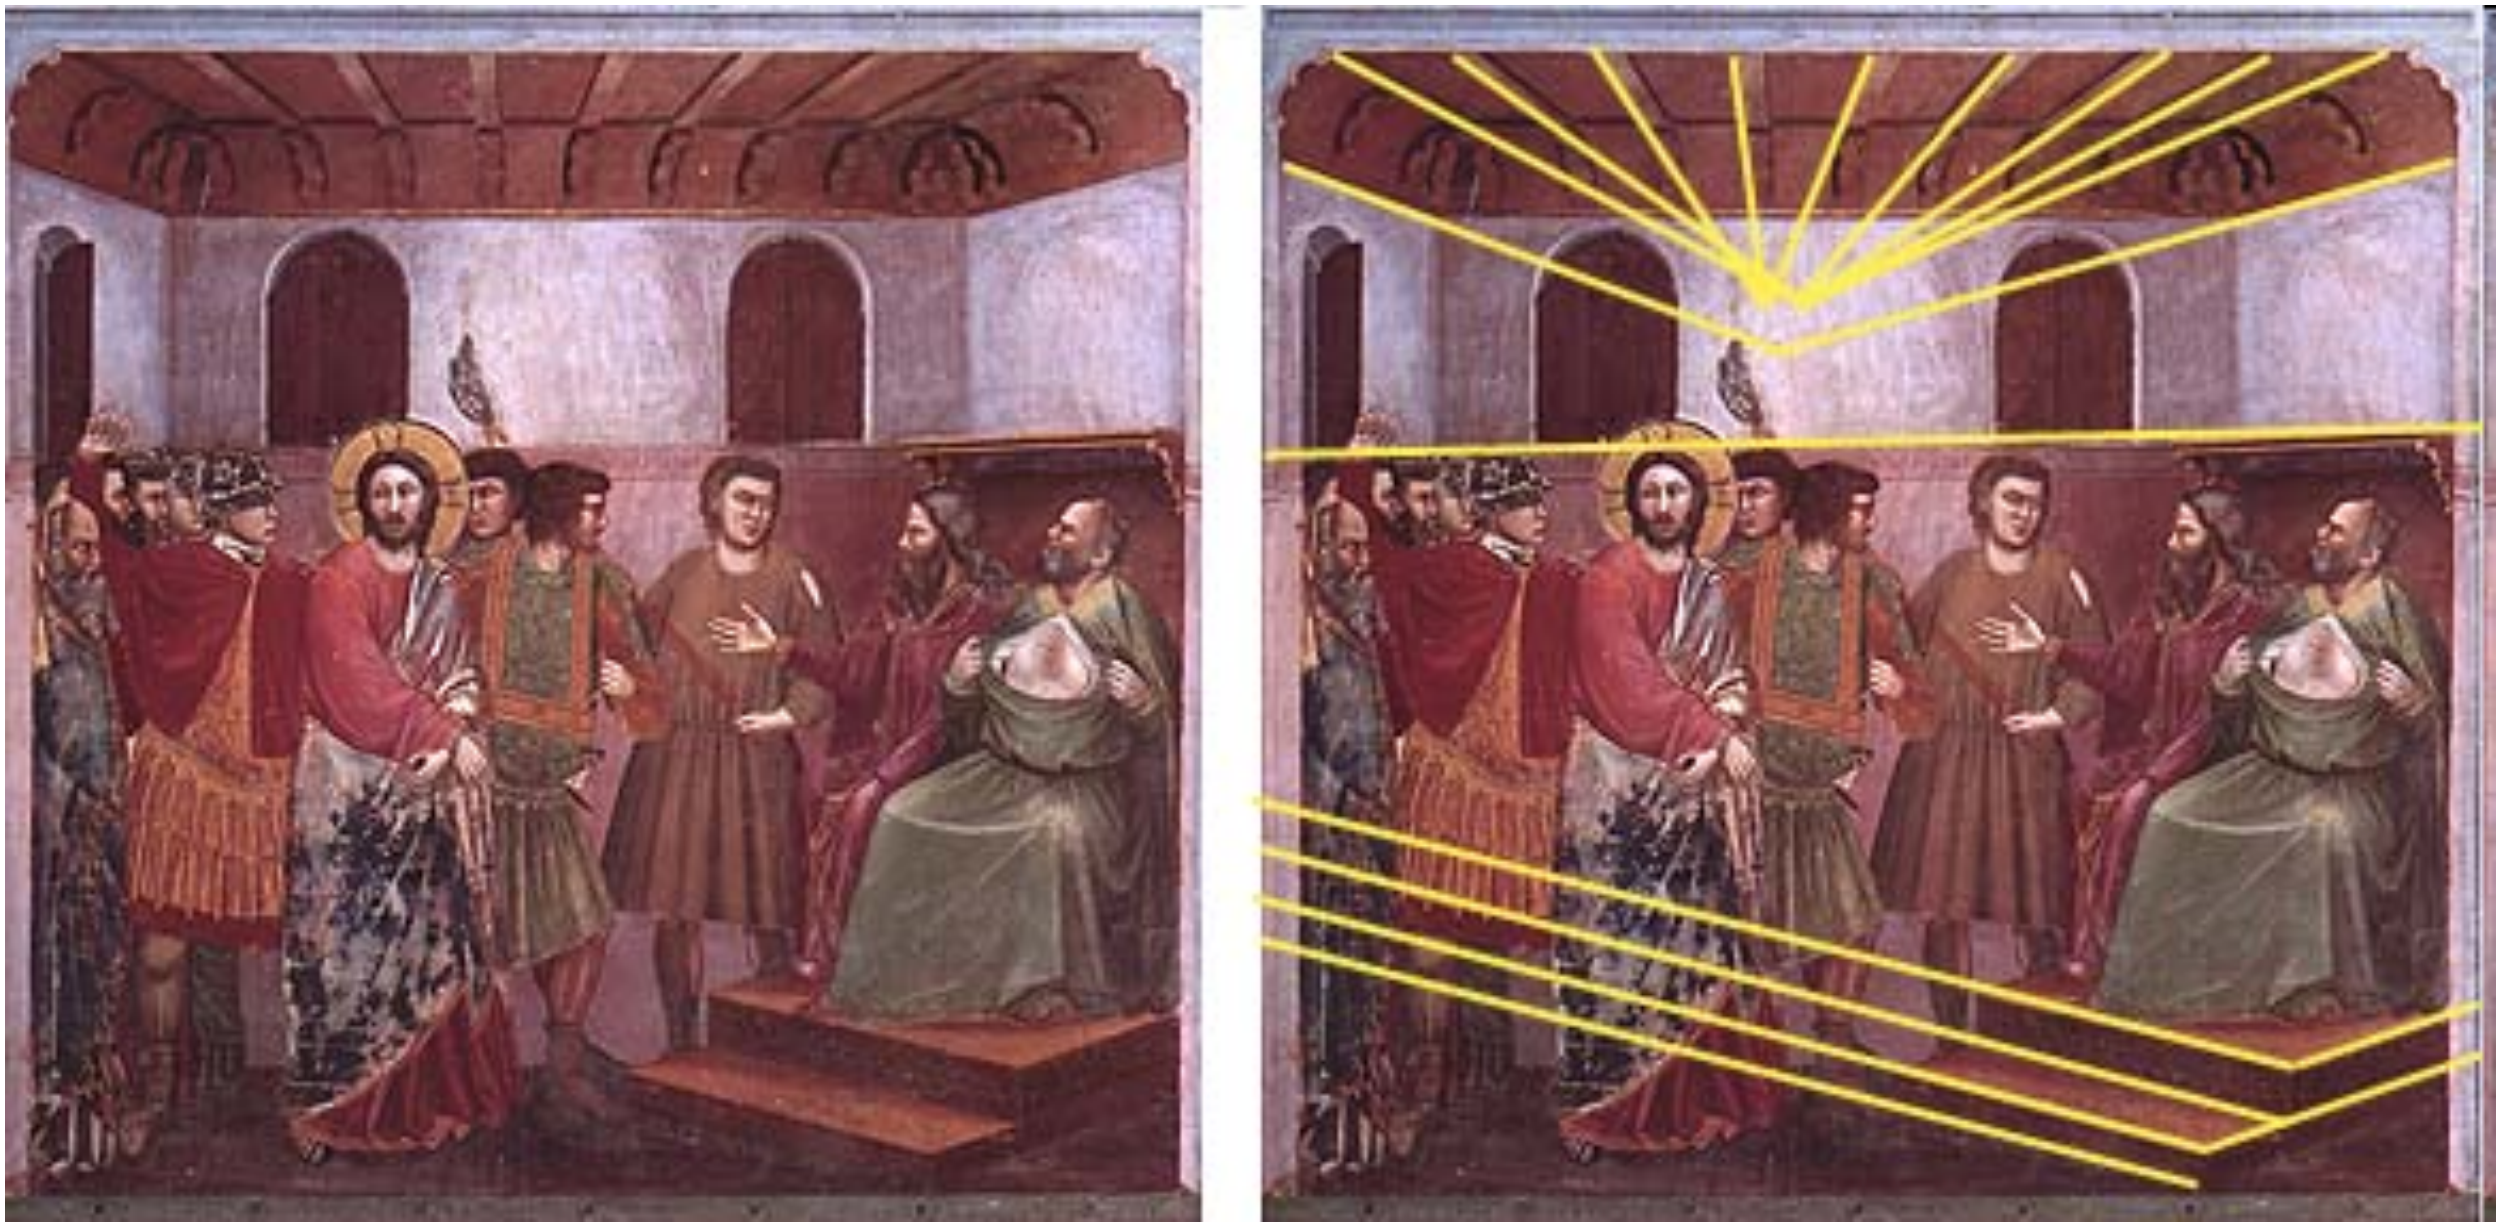

‘Jesus Before the Caïf’, by Giotto (1305).

From CW Tyler, Perspective as a Geometric Tool that Launched the Renaissance

Progress

Follower of Pietro Lorenzetti, *Madonna and Child Enthroned with Angels* 1360/70

Success

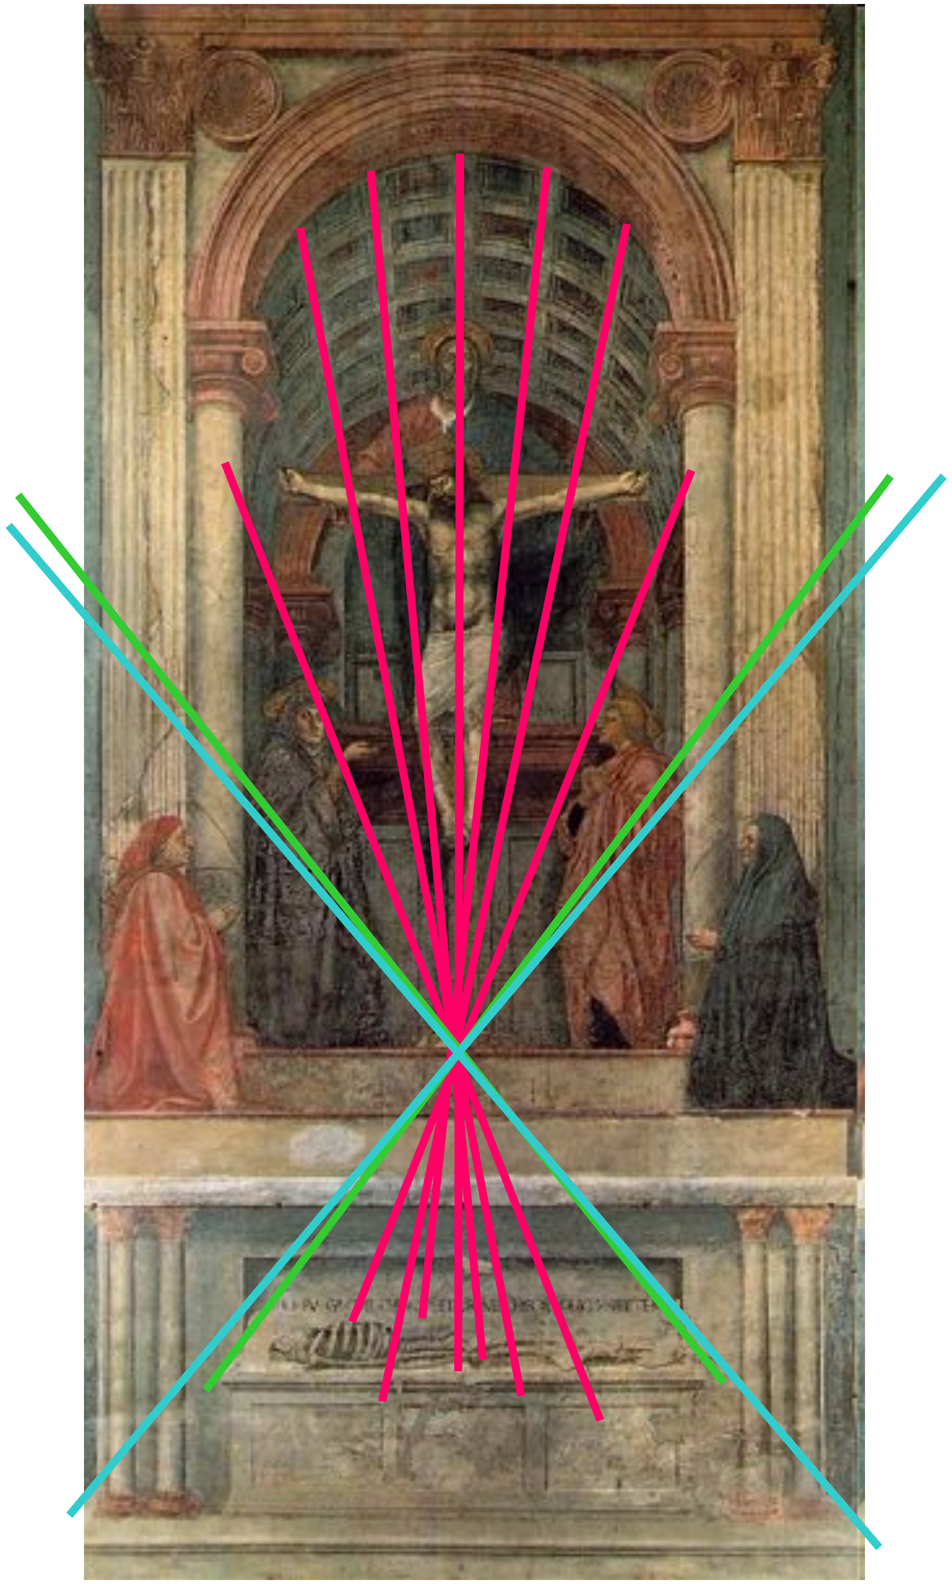

Masaccio's *Trinity* (1427–28)
(Church of Santa Maria Novella, Florence)

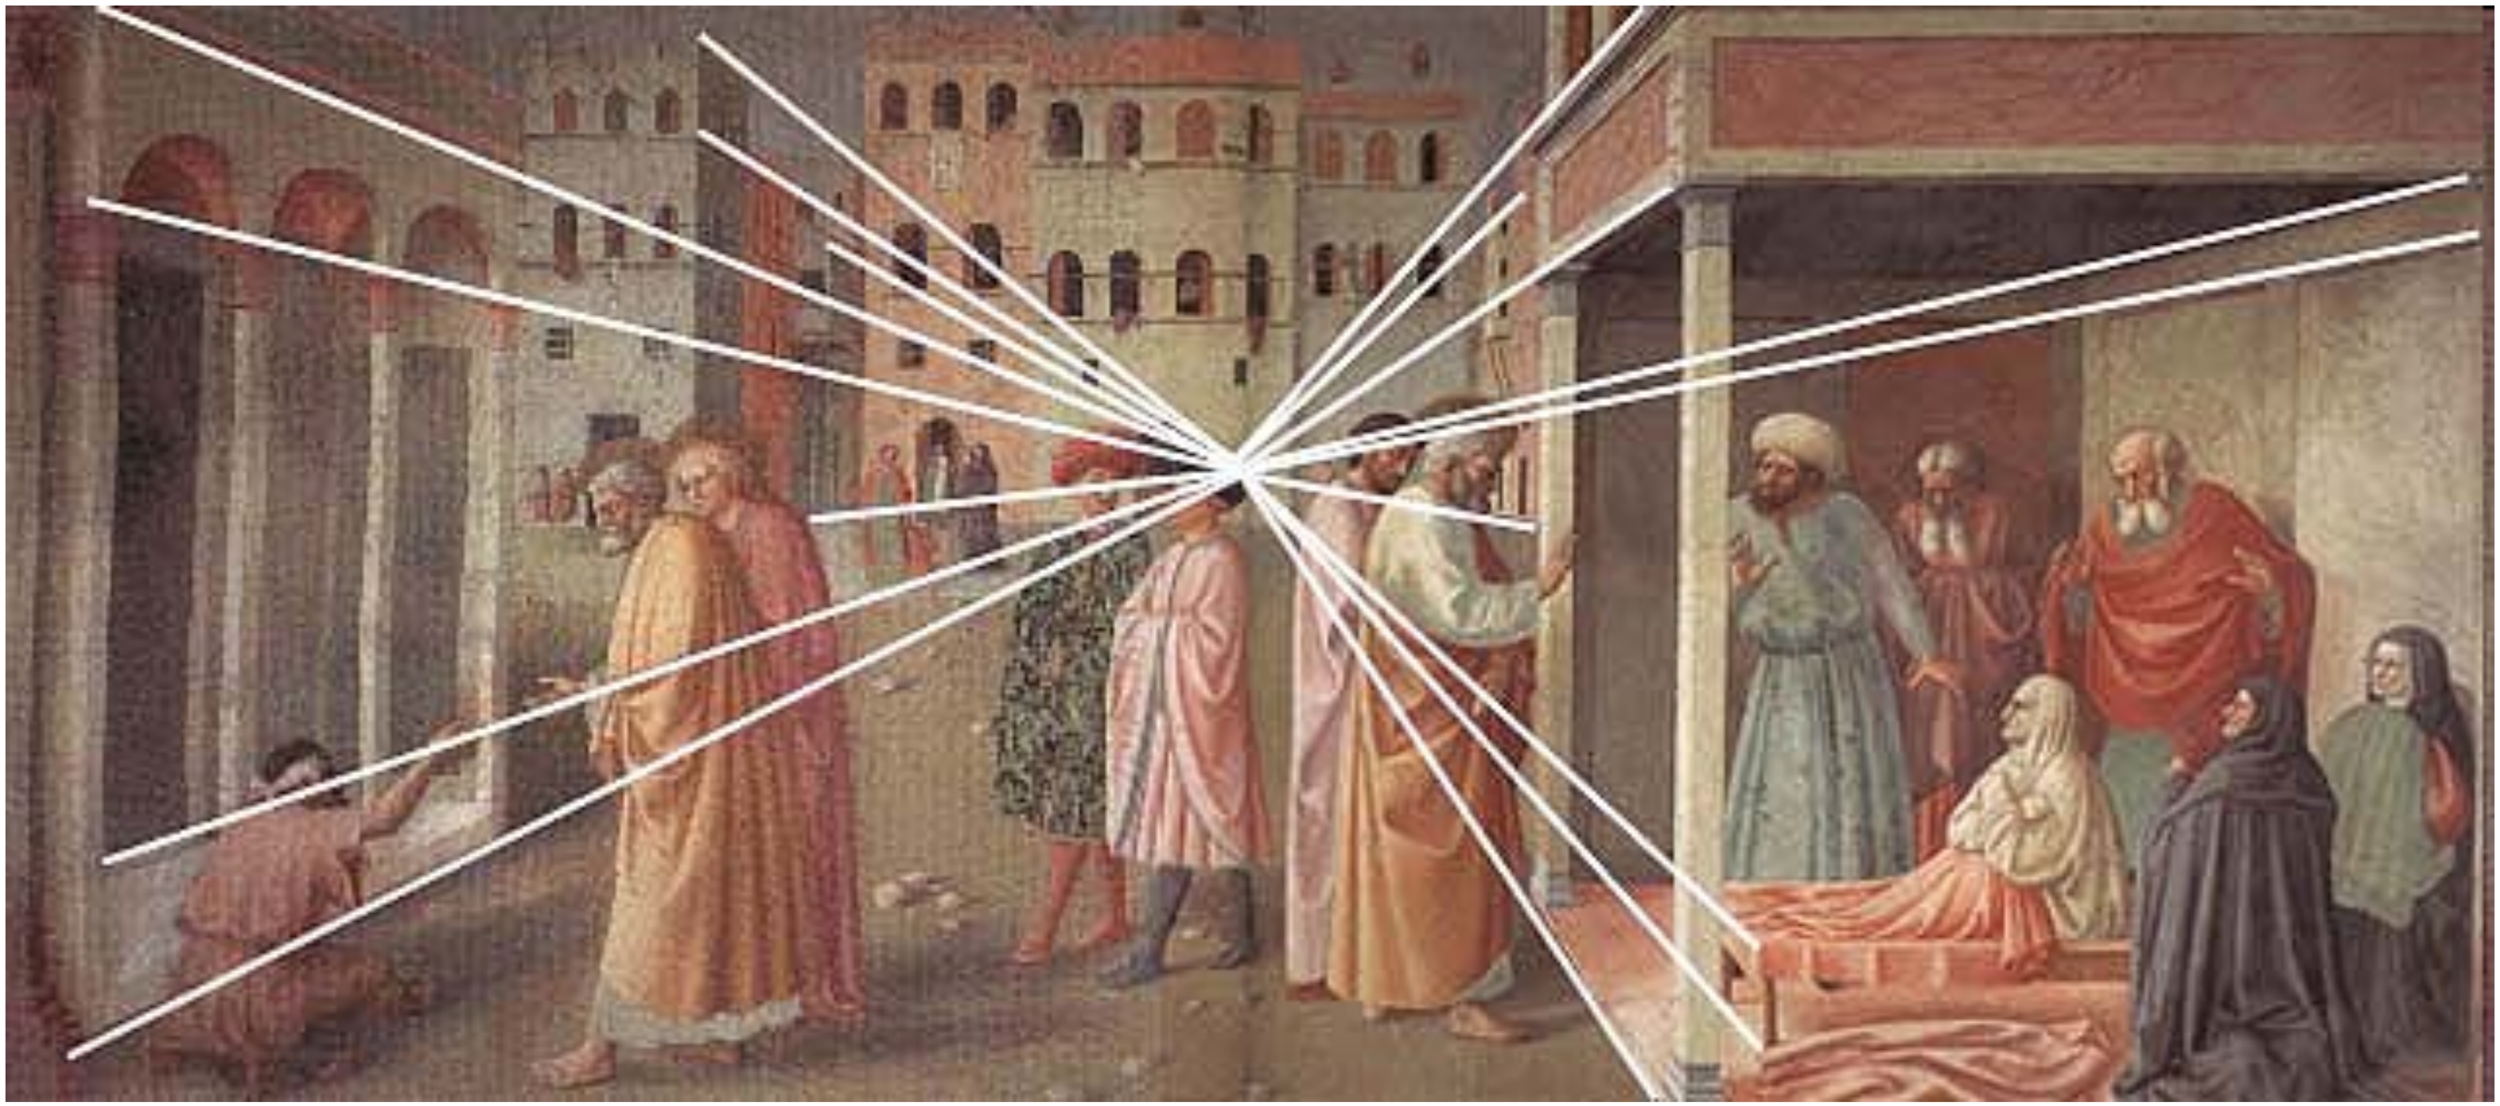

The Healing of the Cripple and the Raising of Tabitha', by Masolino (1425).
Brancacci Chapel, Santa Maria del Carmine, Florence

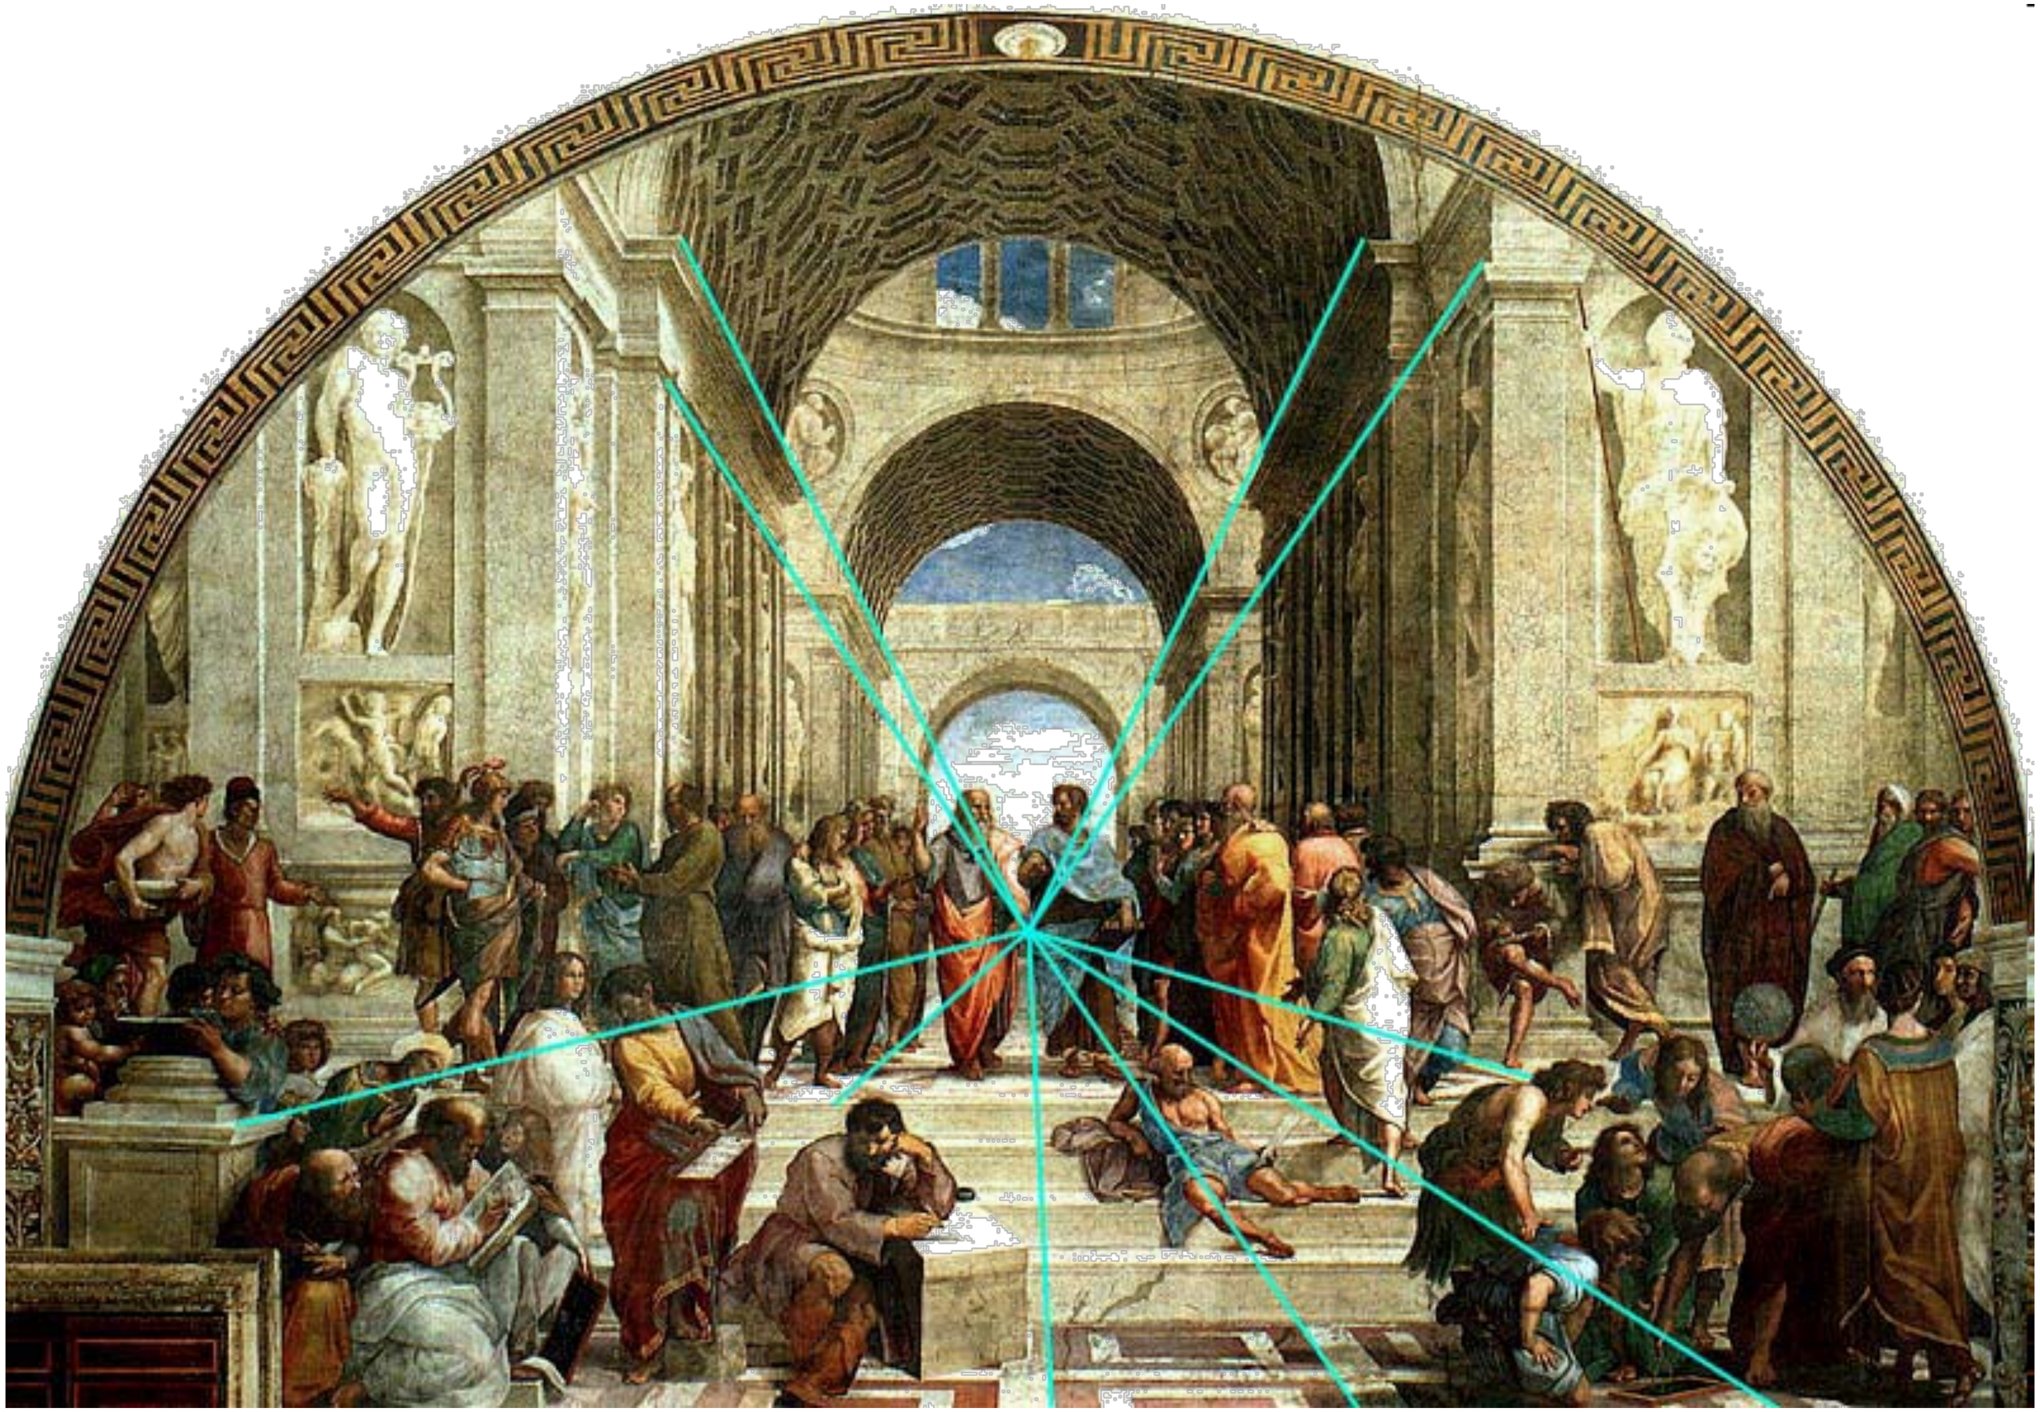

‘The School of Athens’ by Raphael (1518), Stanze di Raffaello, in the Apostolic Palace in the Vatican.

Linear Perspective in Computer Vision

- How can we use linear perspective in computer vision?

The Manhattan World

- The Manhattan world is a model of the 3D environment that assumes that structure in the scene aligned with a 3D orthogonal Cartesian coordinate frame.

Manhattan Frame Estimation

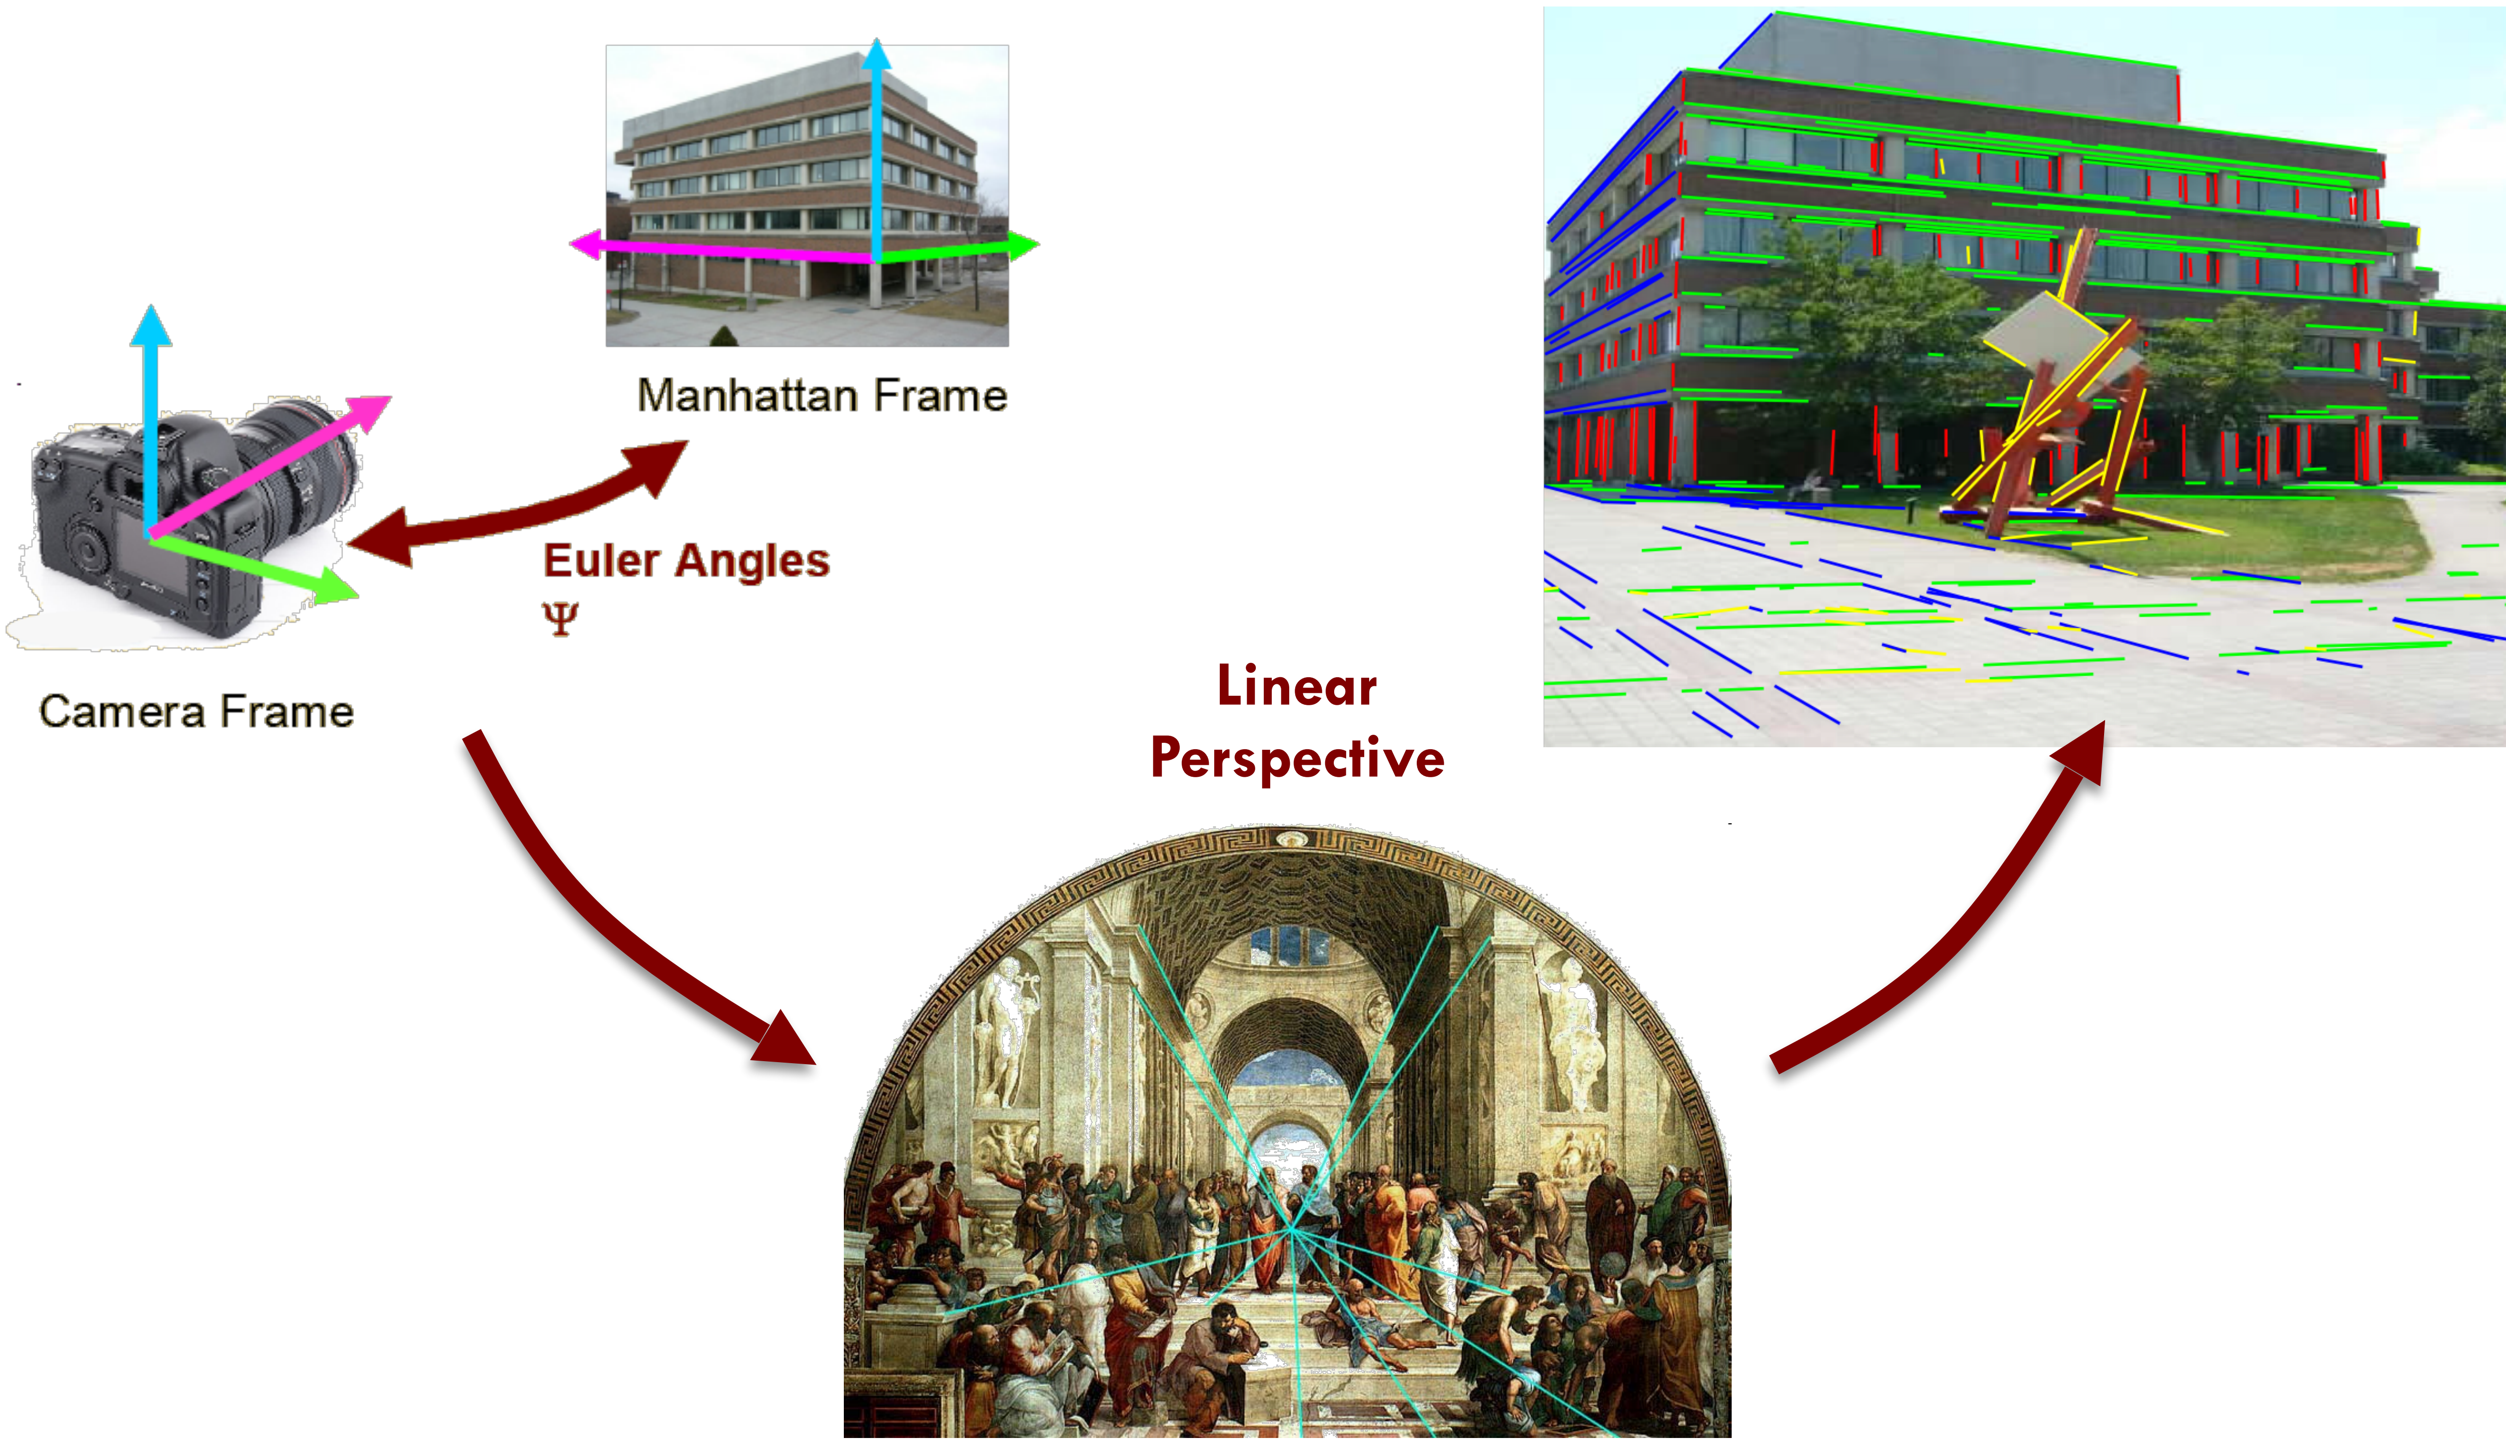

End of Lecture
Nov 7, 2018

3-Point Perspective

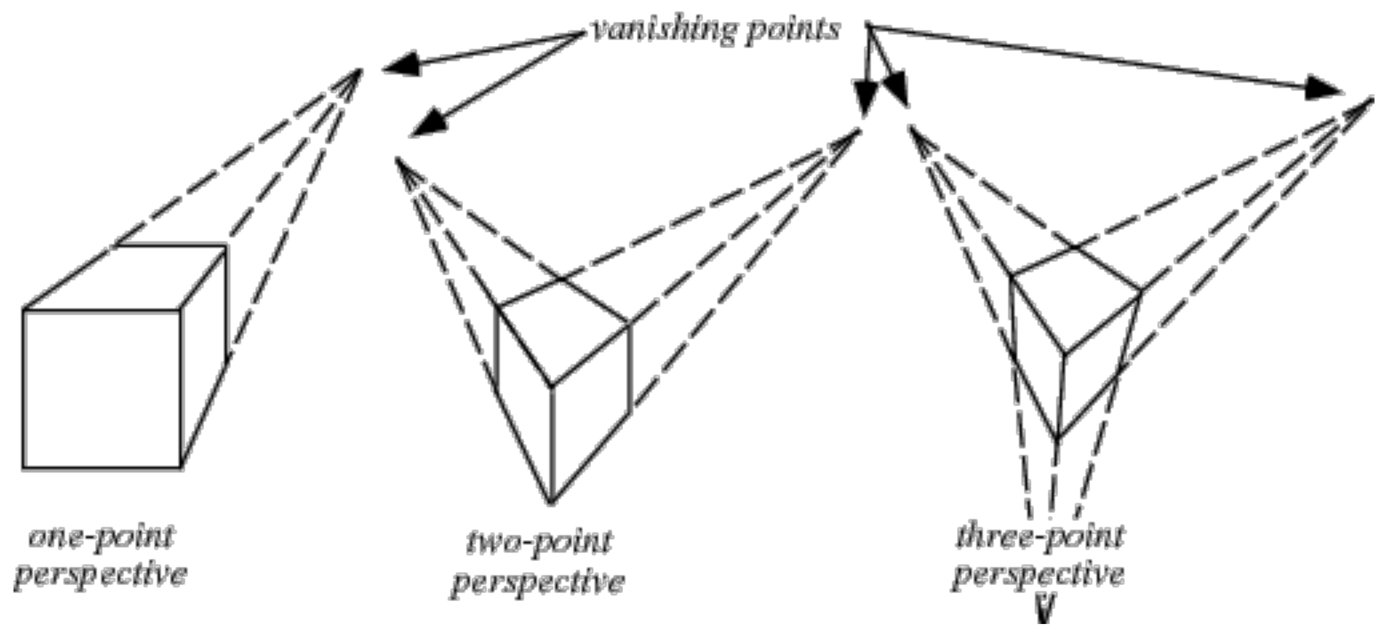

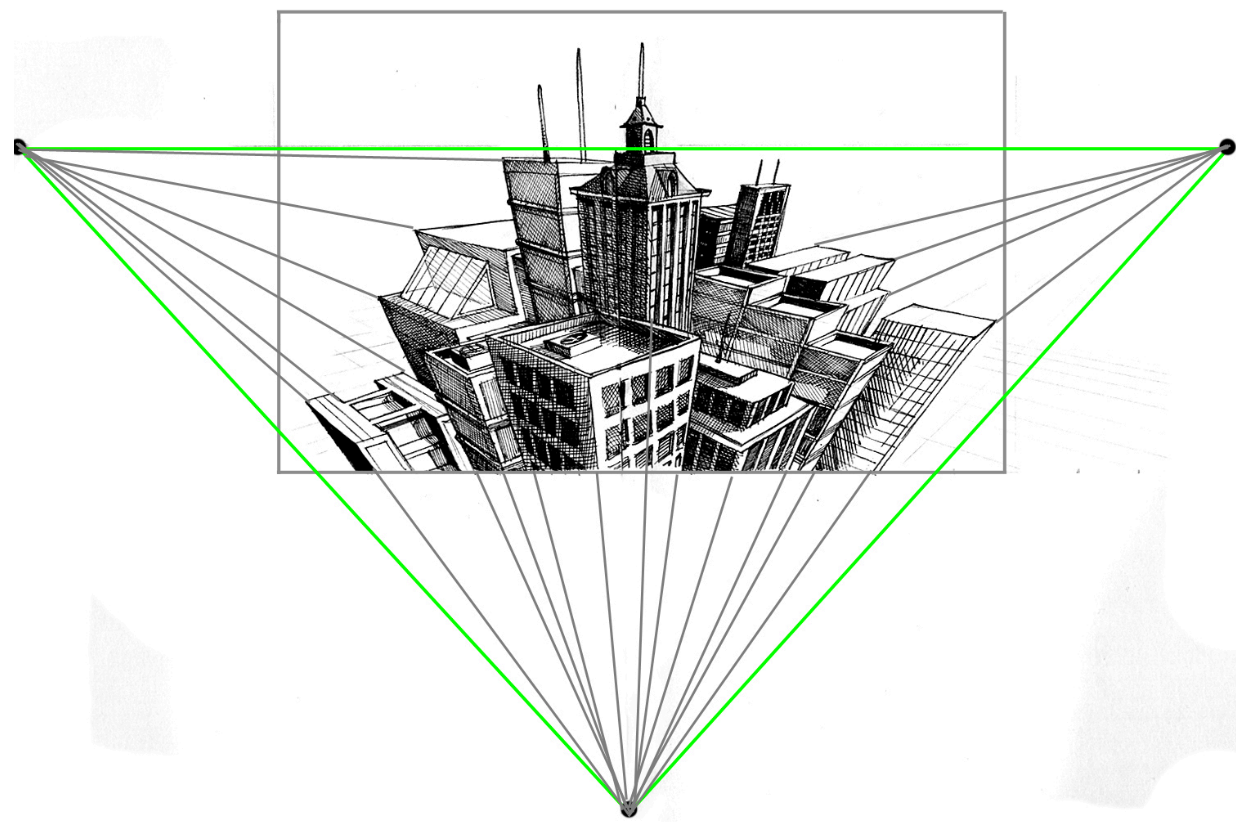

Vanishing Points and the Manhattan Frame

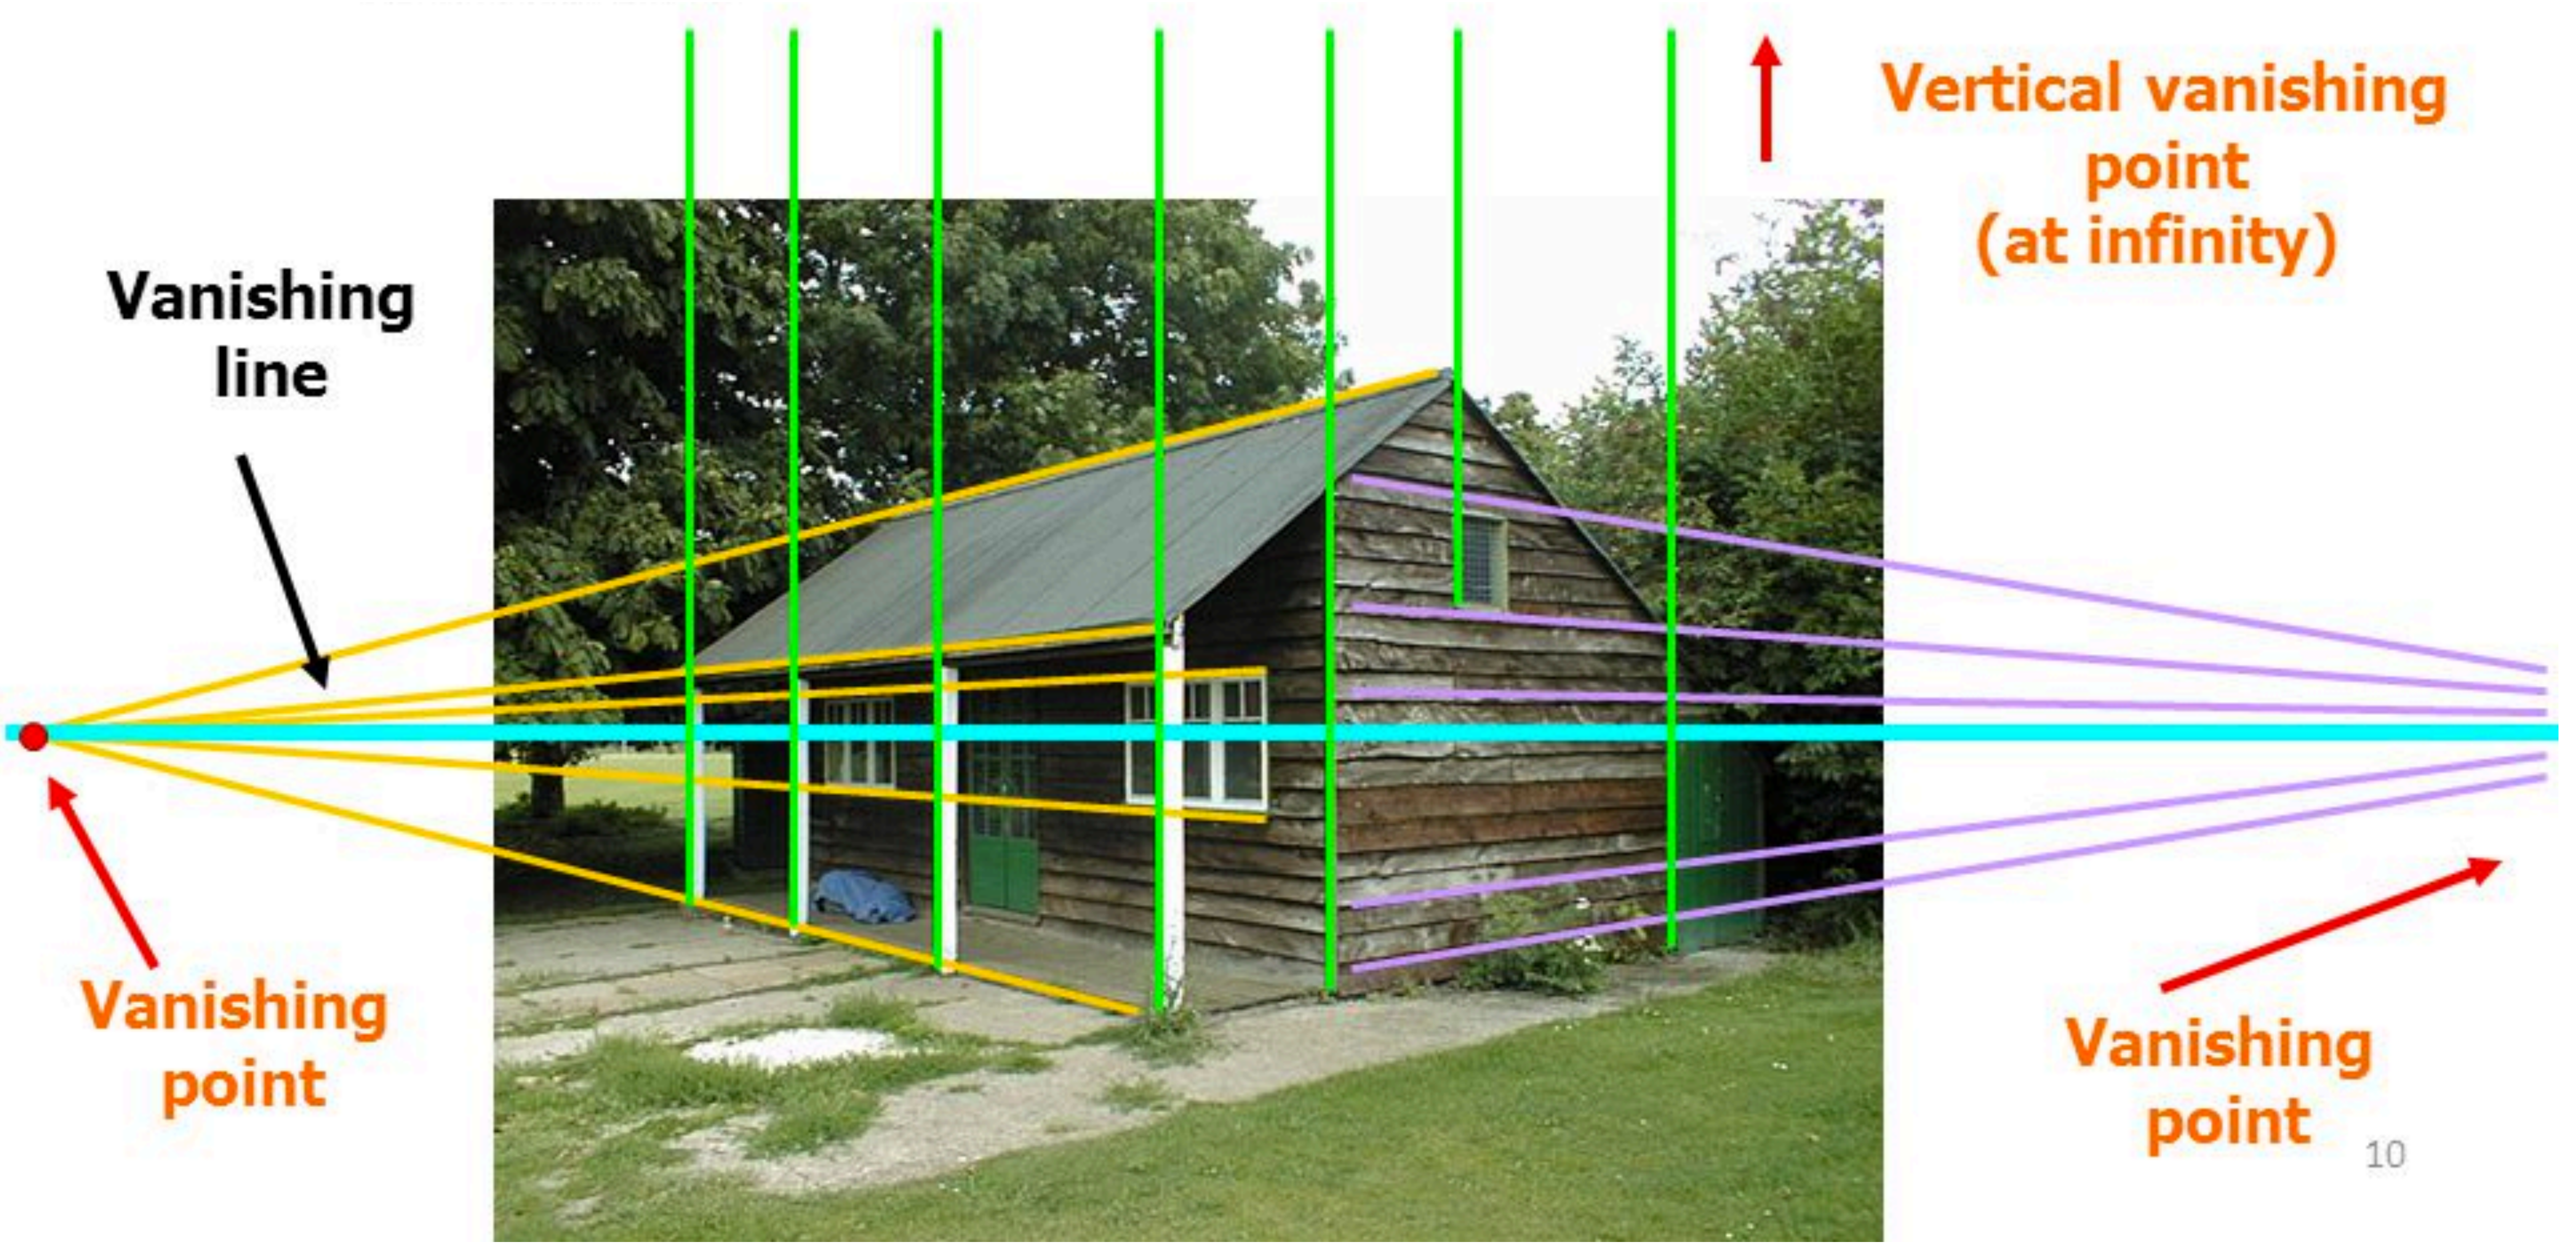

Estimating the Manhattan Frame

- To estimate the 3D rotation of the Manhattan world frame relative to the camera we need to find the vanishing points.
- This rotation has 3 degrees of freedom.
  - e.g., the axis of rotation (2 dof) and the angle of rotation (1 dof)
- The locations of Manhattan vanishing points in the images are determined by:
  - The camera rotation (3 dof)
  - The focal length (1 dof)
  - The principal point (2 dof)
- How many Manhattan vanishing points are needed to estimate the camera rotation
  - If focal length and principal point are known?
  - If focal length and principal point are unknown?

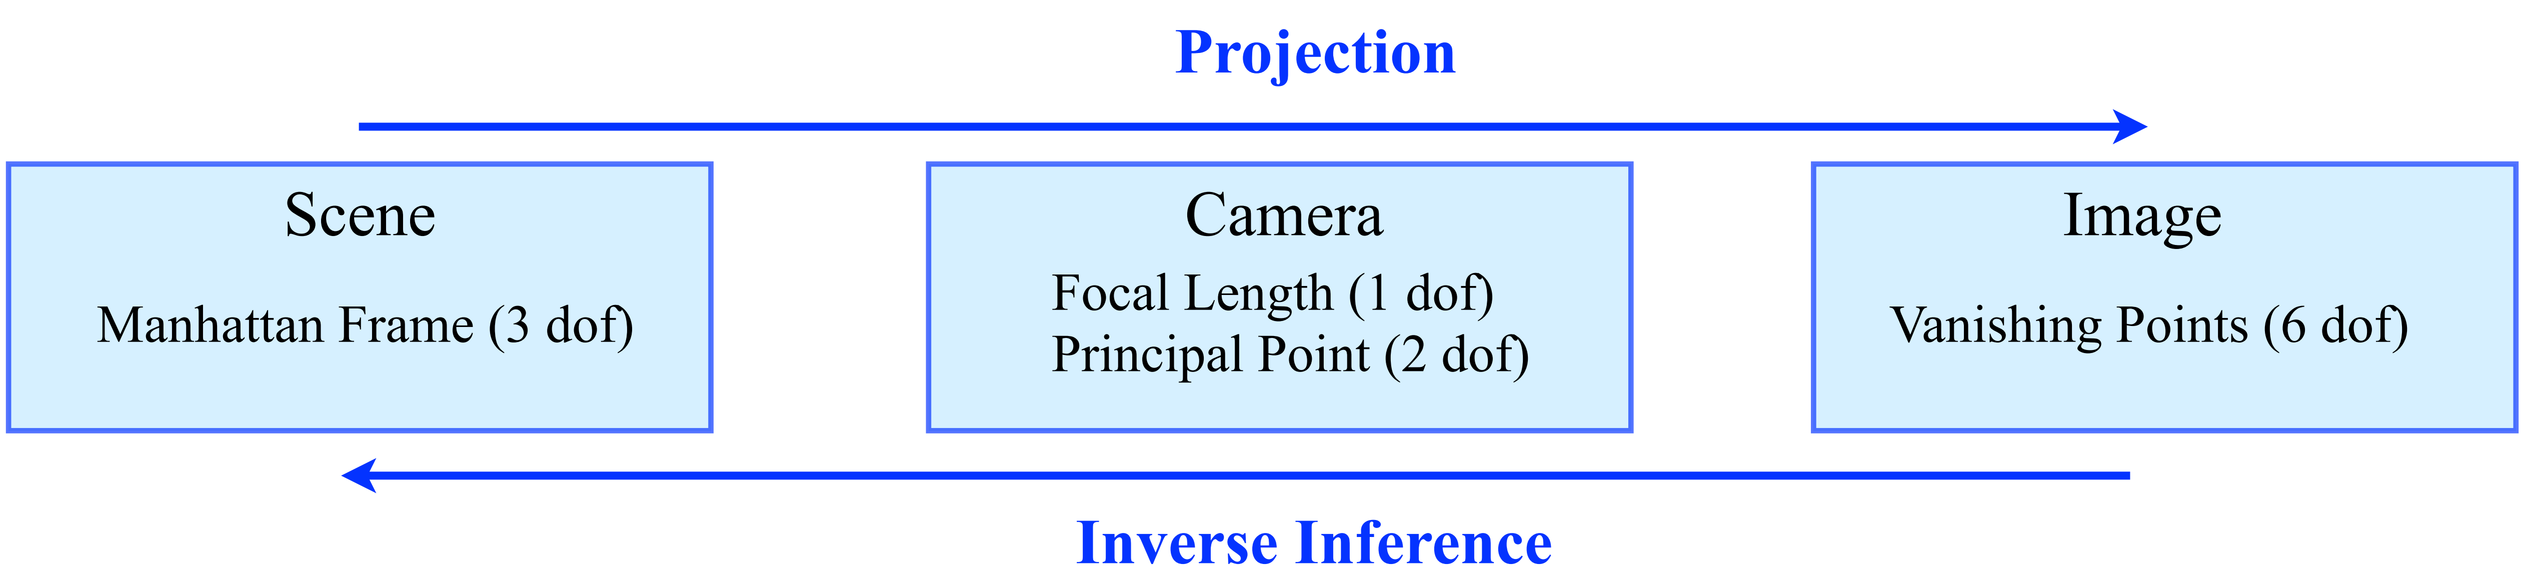

Estimating the Manhattan Frame

- Prior Work
  - Couglan & Yuille 1999, 2003
  - Deutscher et al 2002
  - Schindler & Dellaert 2004
  - Kosecka & Zhang 2002
- Limitations
  - Methods can be slow
  - No standardized database
  - No systematic evaluation and comparison

York Urban Database (2008)

- www.elderlab.yorku.ca/YorkUrbanDB
  - 102 images of urban Toronto scenes
  - 12,122 labelled Manhattan line segments
  - Estimates of ground truth Manhattan frame for each image (estimated accuracy ~1.5 deg)

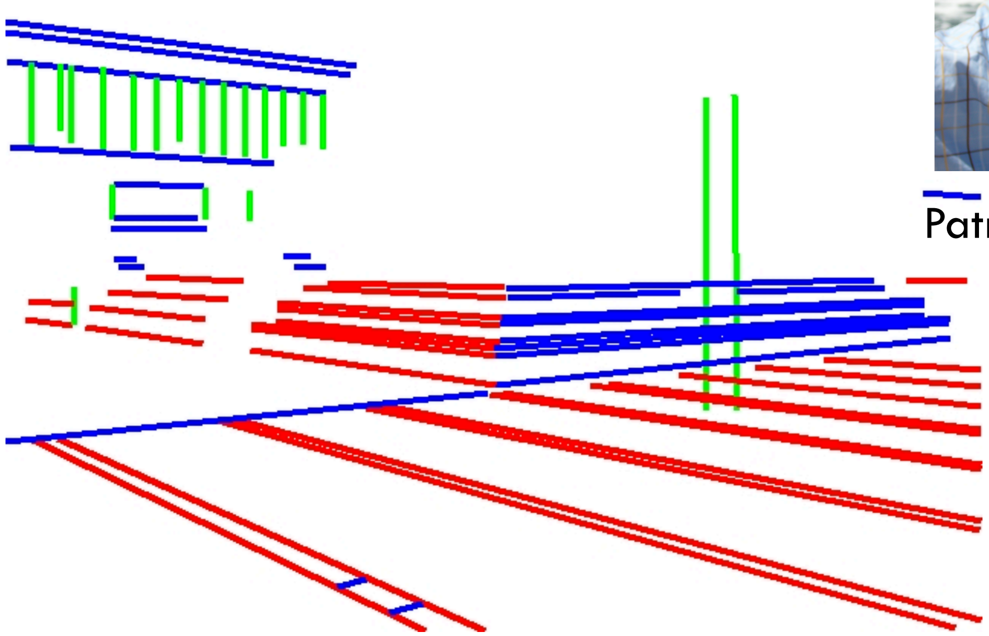

Patrick Denis

Denis, Elder & Estrada, ECCV 2008

Estimating Vanishing Points in the Image

- Evidence for a hypothesized vanishing point can be obtained by measuring angular deviations $\Delta\phi$ of local oriented observations from the predicted direction.

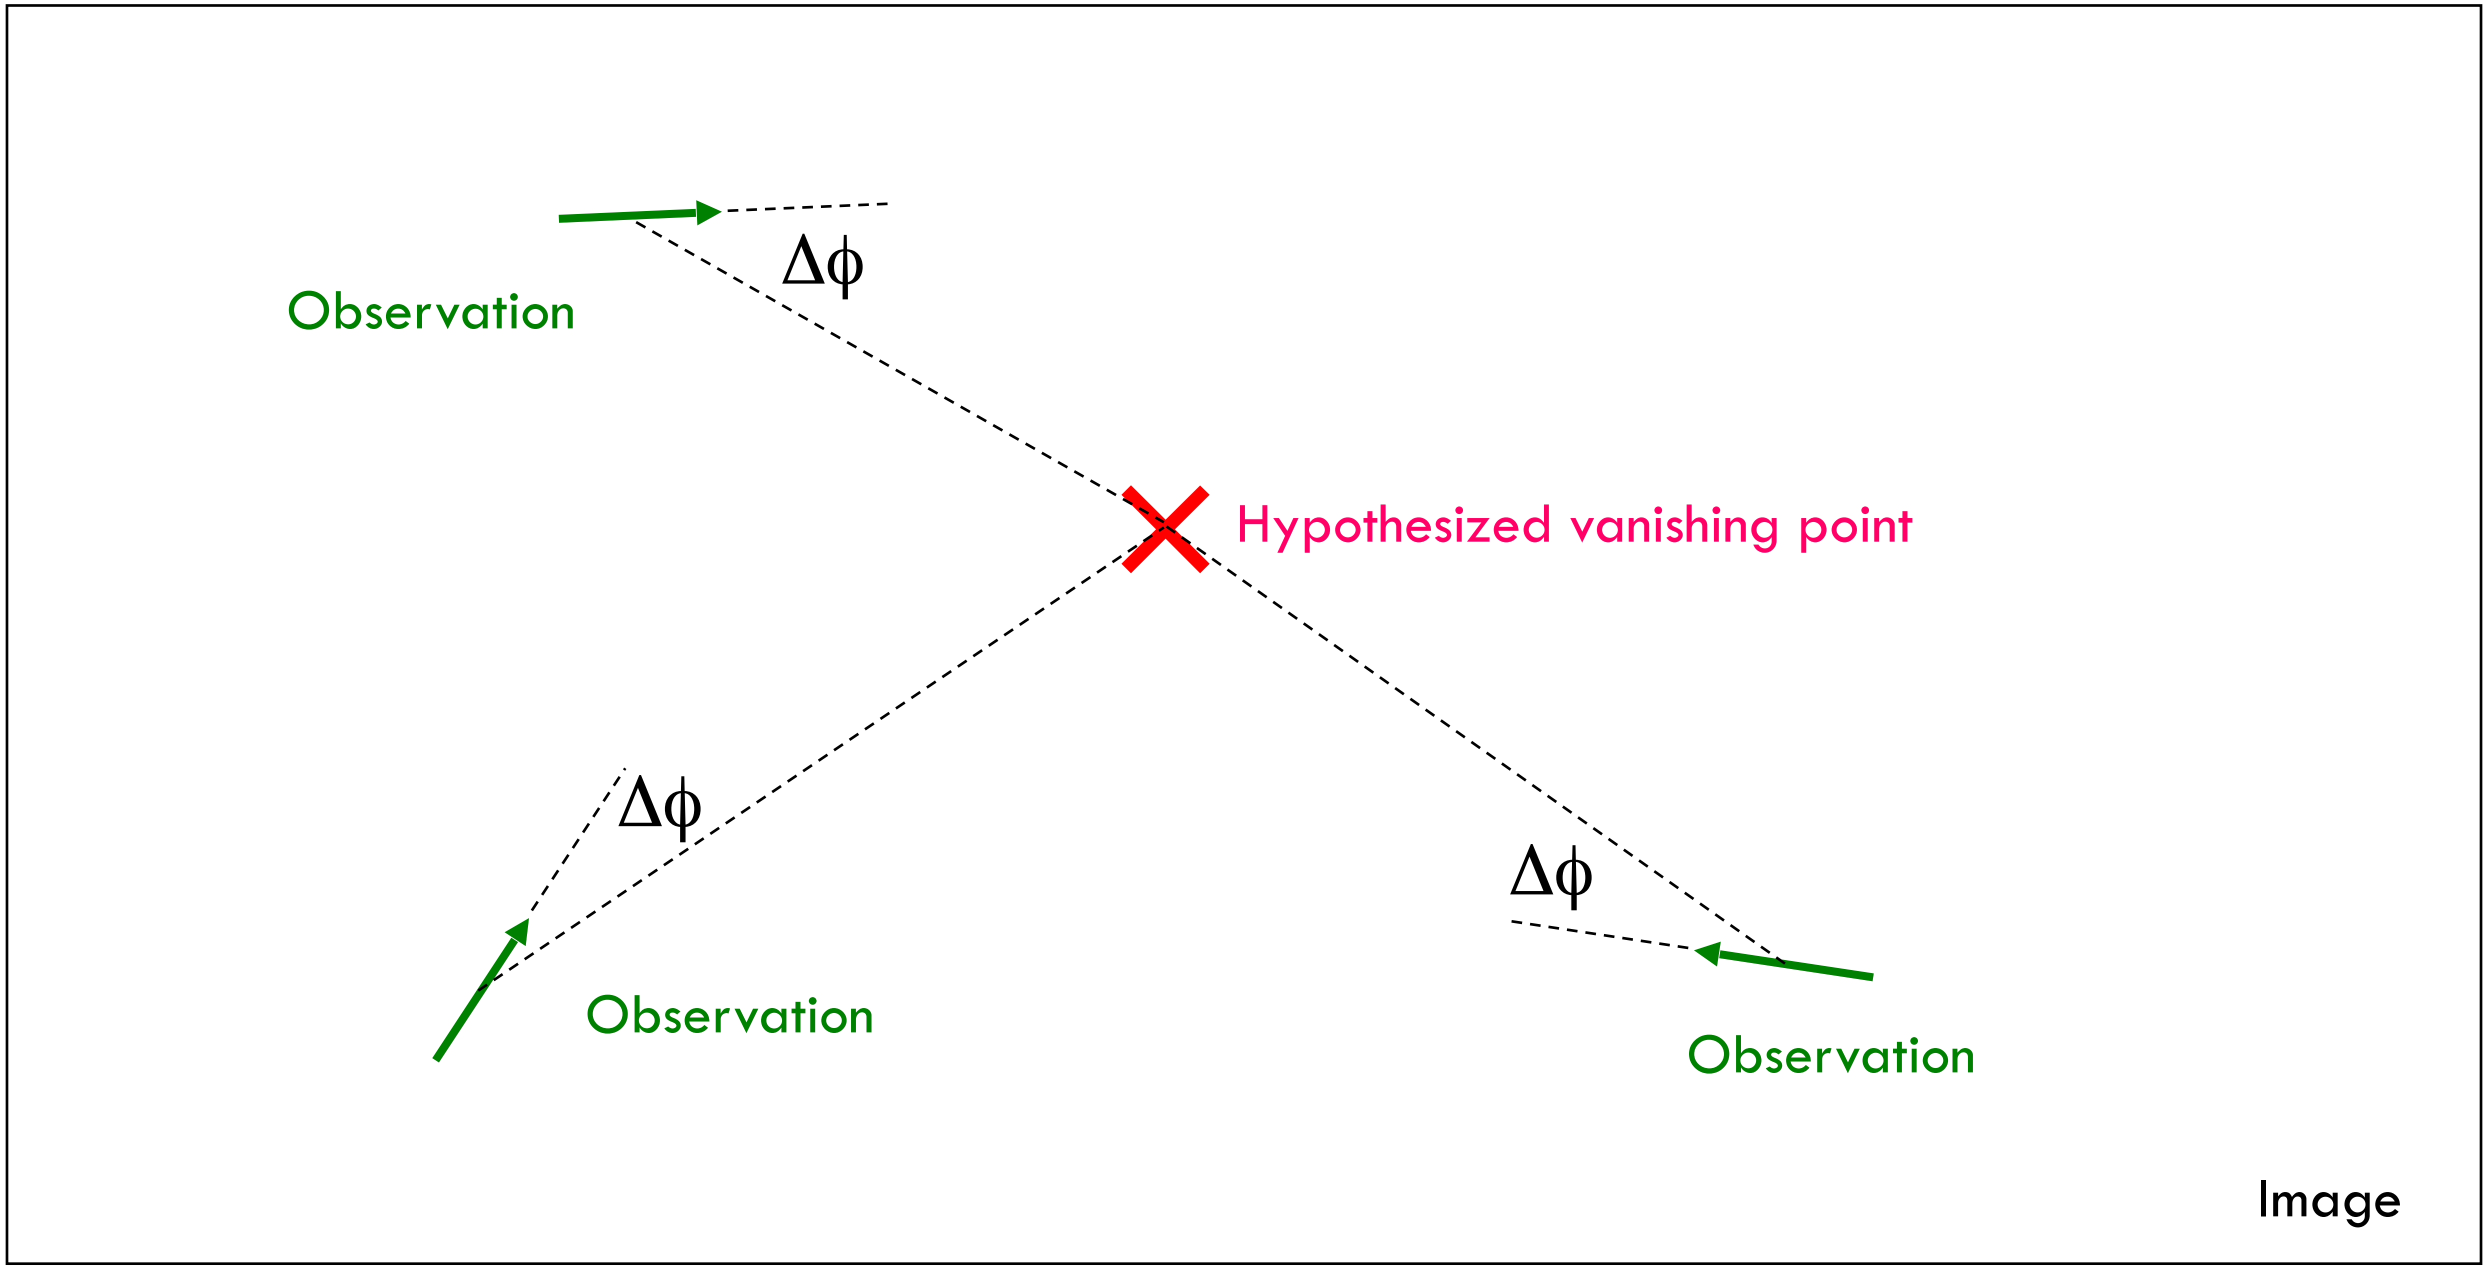

Estimating Vanishing Points on the Gauss Sphere

- Definitions:
  - The **Gauss sphere** is the collection of viewing directions centred on the optical centre.
  - The **interpretation plane** is the 3D plane on which an oriented element must lie

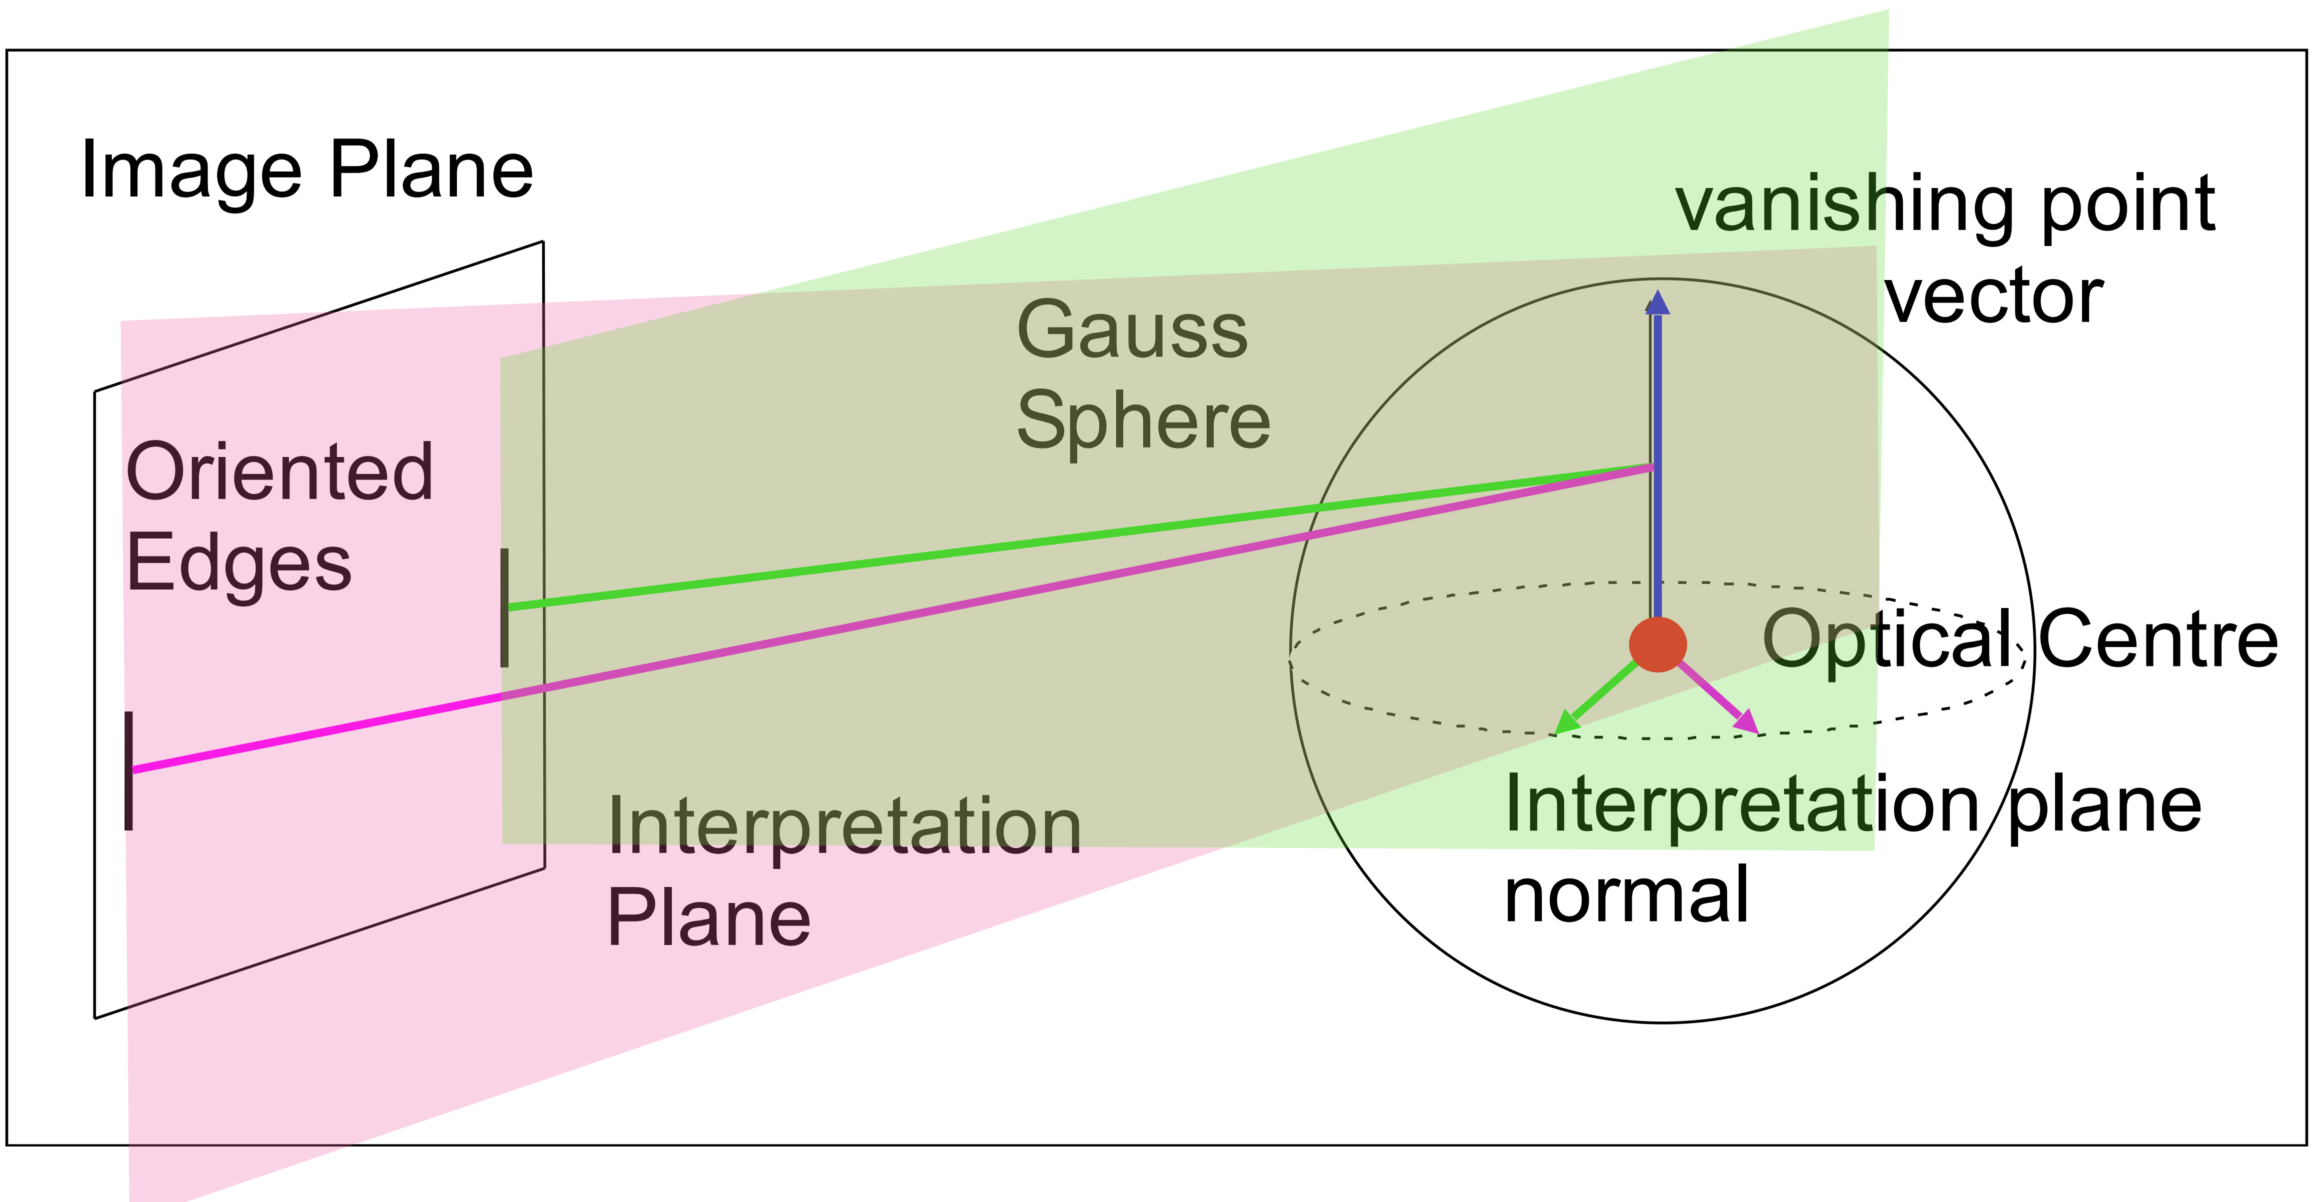

Finding Vanishing Points on the Gauss Sphere

- A family of parallel 3D lines should generate interpretation plane normals distributed over a circle in the Gauss sphere.
- The normal to this circle is the 3D vanishing point direction.
- In practice, due to noise these interpretation plane normals are distributed over a circular band

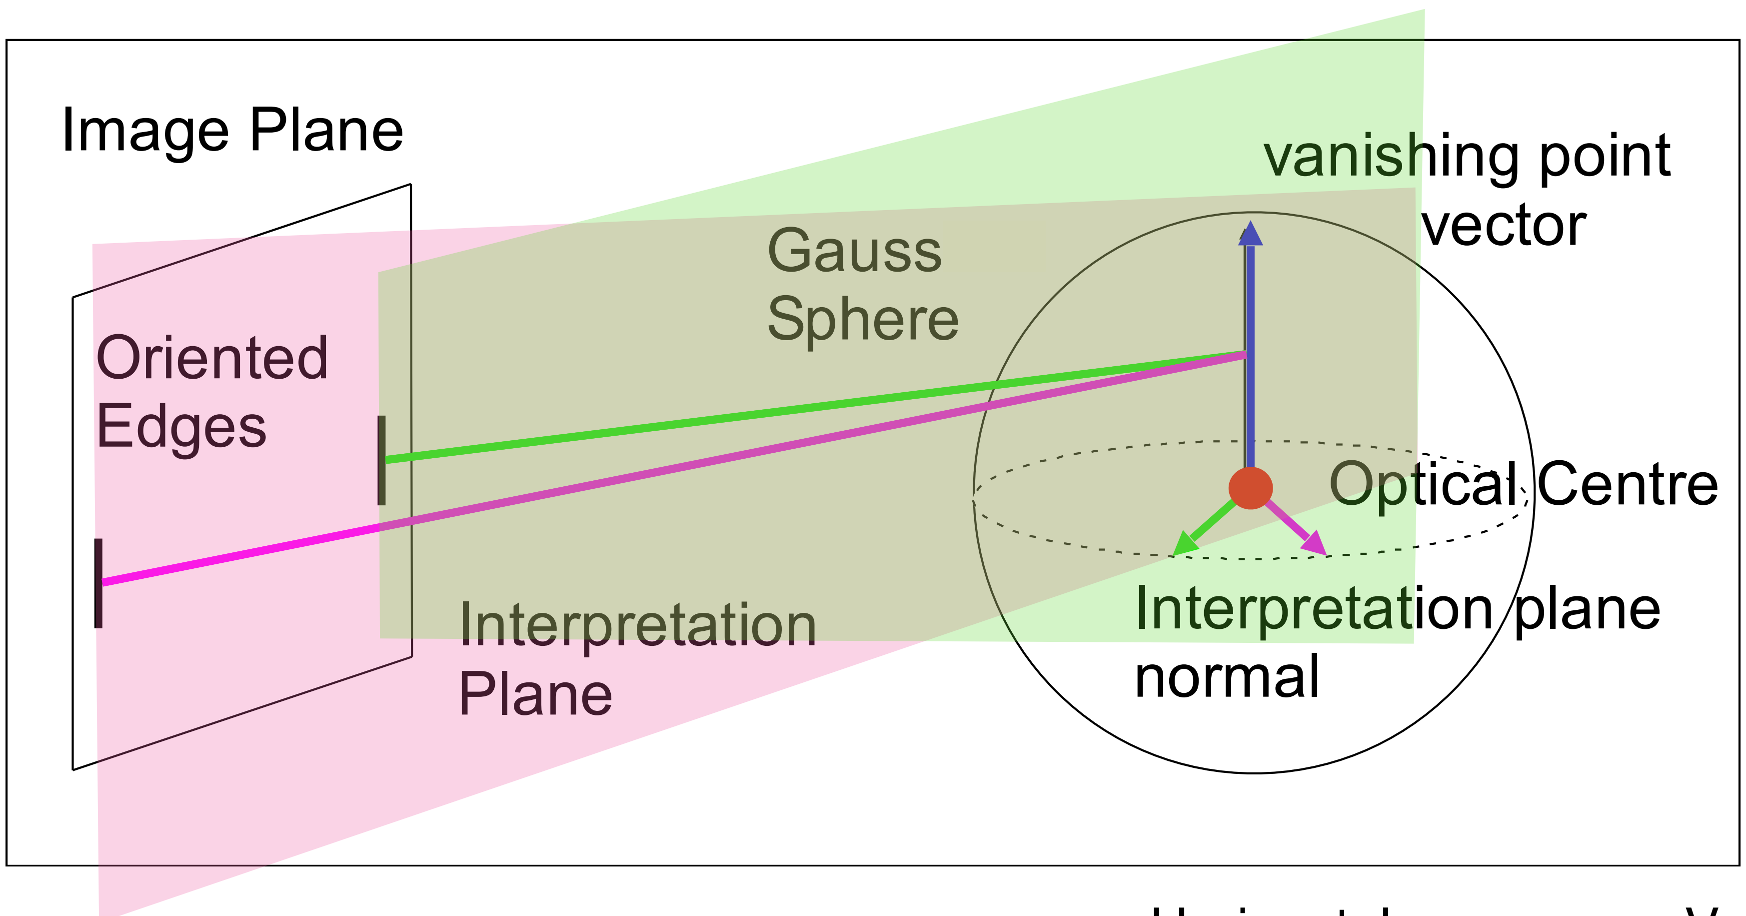

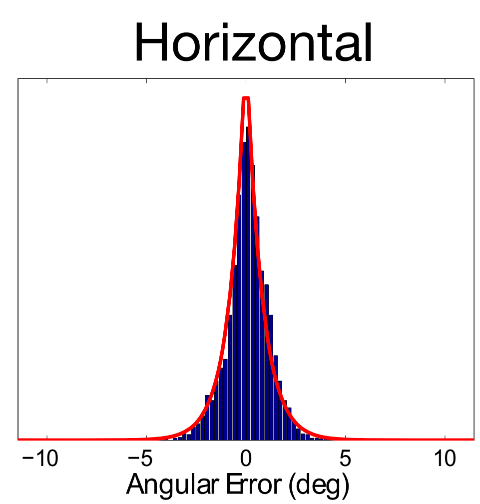

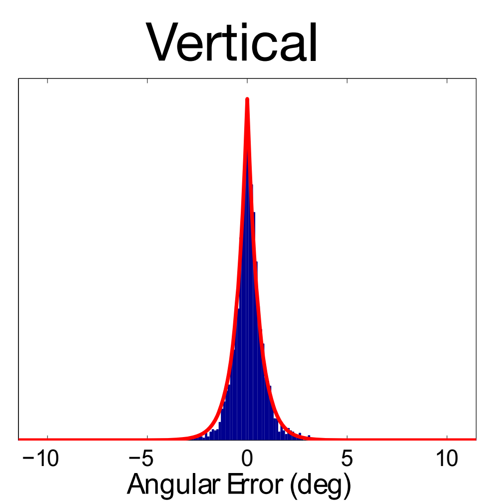

Mixture Model

- Each oriented element E_u in the image is generated by one of four possible kinds of scene structure:
  - $m_{1\text{-}3}$: a line in one of the three Manhattan directions
  - m_4: non-Manhattan structure
- The likelihoods of these elements are co-determined by:
  - The causal process ($m_{1\text{-}4}$)
  - The rotation Ψ of the Manhattan frame relative to the camera

$$\Psi^* = \arg\max_{\Psi} \sum_{\vec{u}} \log P\left(E_{\vec{u}} \mid \Psi\right)$$

where

$$P\left(E_{\vec{u}} \mid \Psi\right) = \sum_{m_{\vec{u}}} \underbrace{P\left(E_{\vec{u}} \mid m_{\vec{u}}, \Psi\right)}_{\text{Likelihood}} \underbrace{P\left(m_{\vec{u}}\right)}_{\text{Prior}}$$

Learn from labelled training data

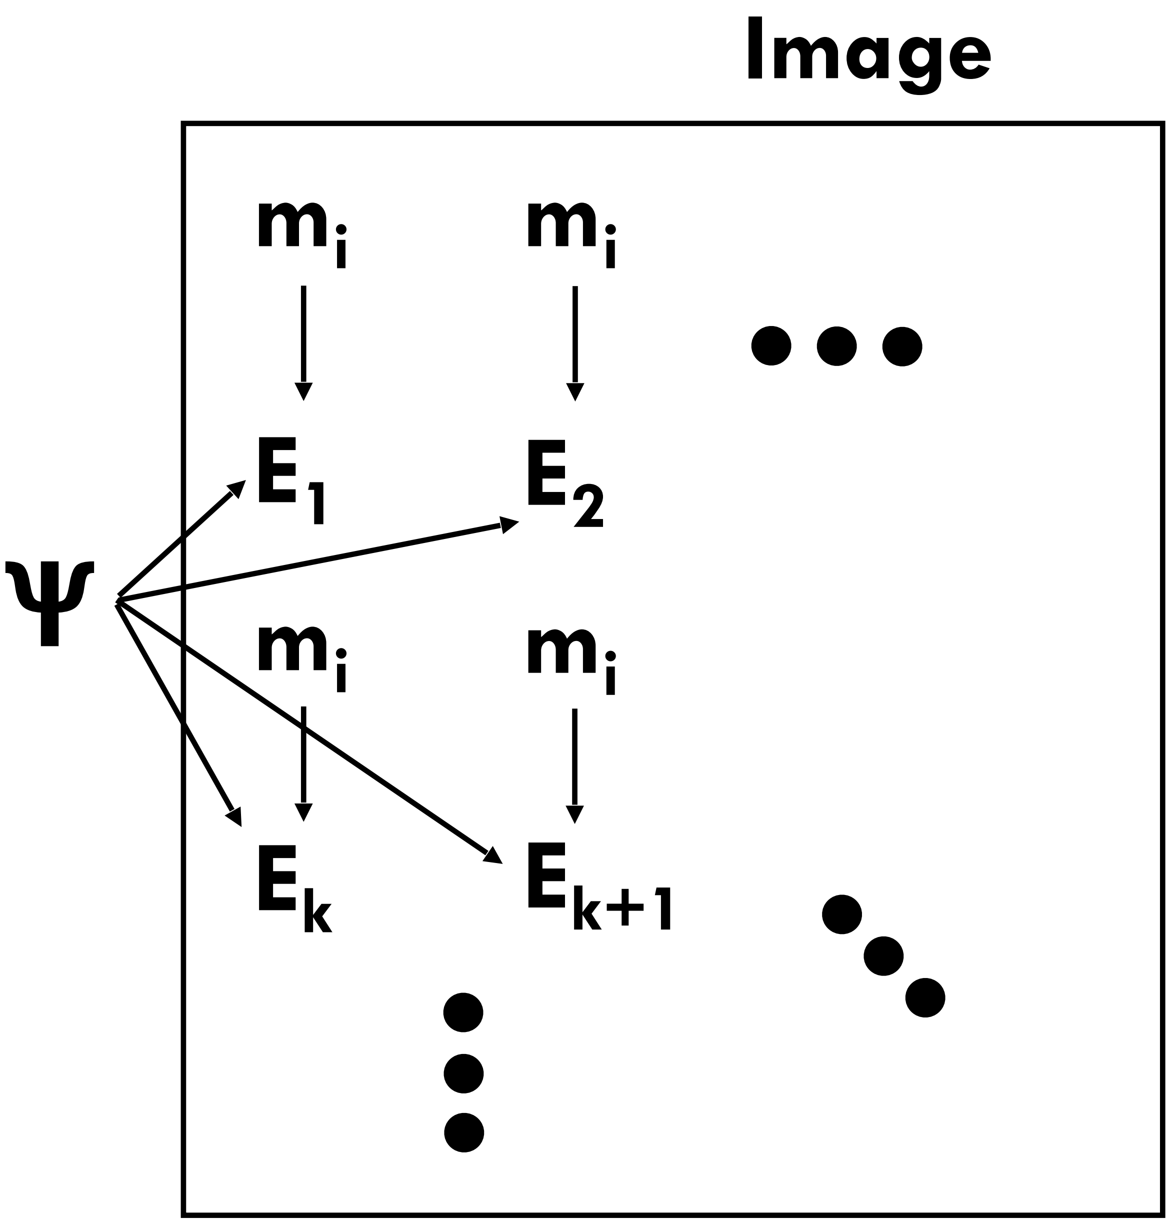

Searching for Ψ*

- Search Methods
  - Coarse-to-Fine (Coughlan & Yuille 2001)
  - Quasi-Newton
  - EM (e.g., Schindler & Dellaert, 2004)
  - Quasi-EM

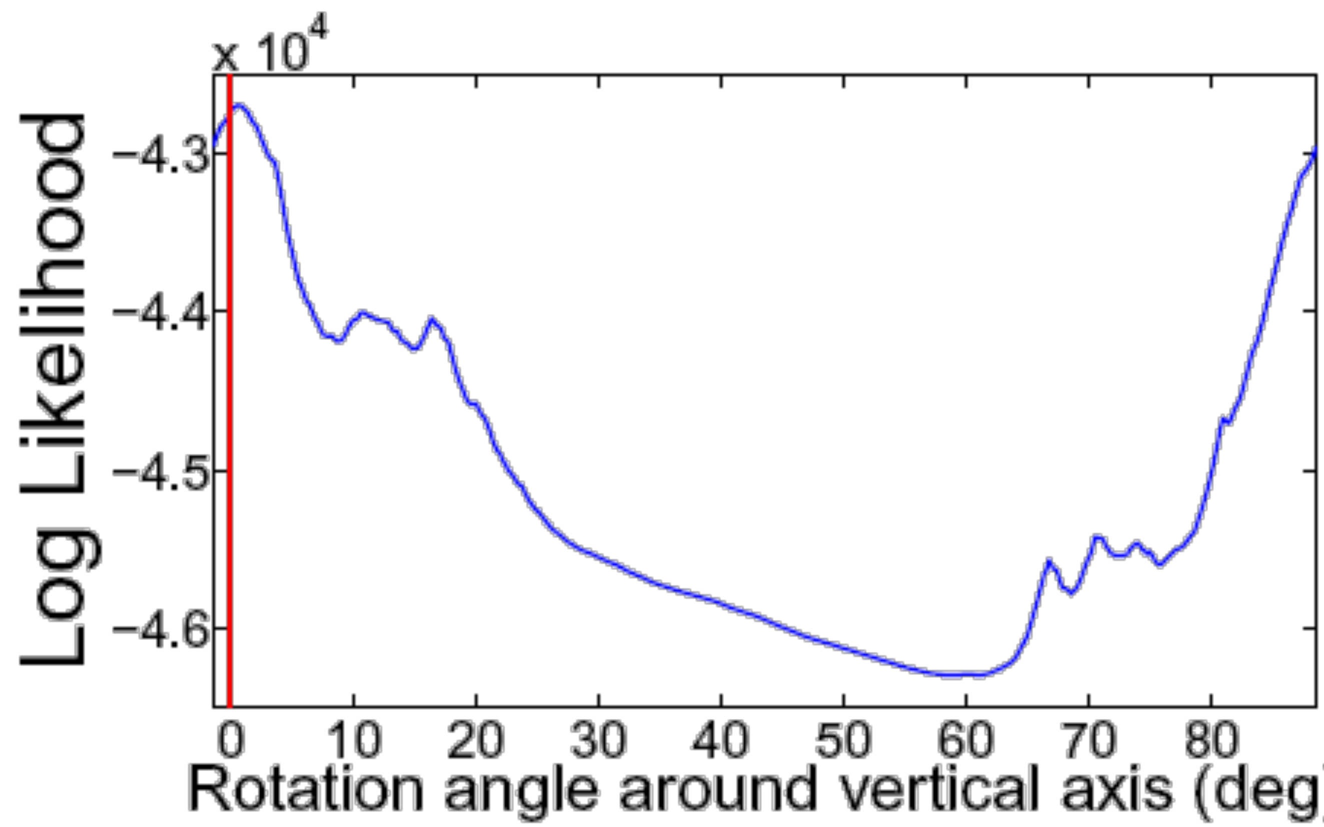

Oriented Elements

- What oriented elements in the image should be used to estimate vanishing points?

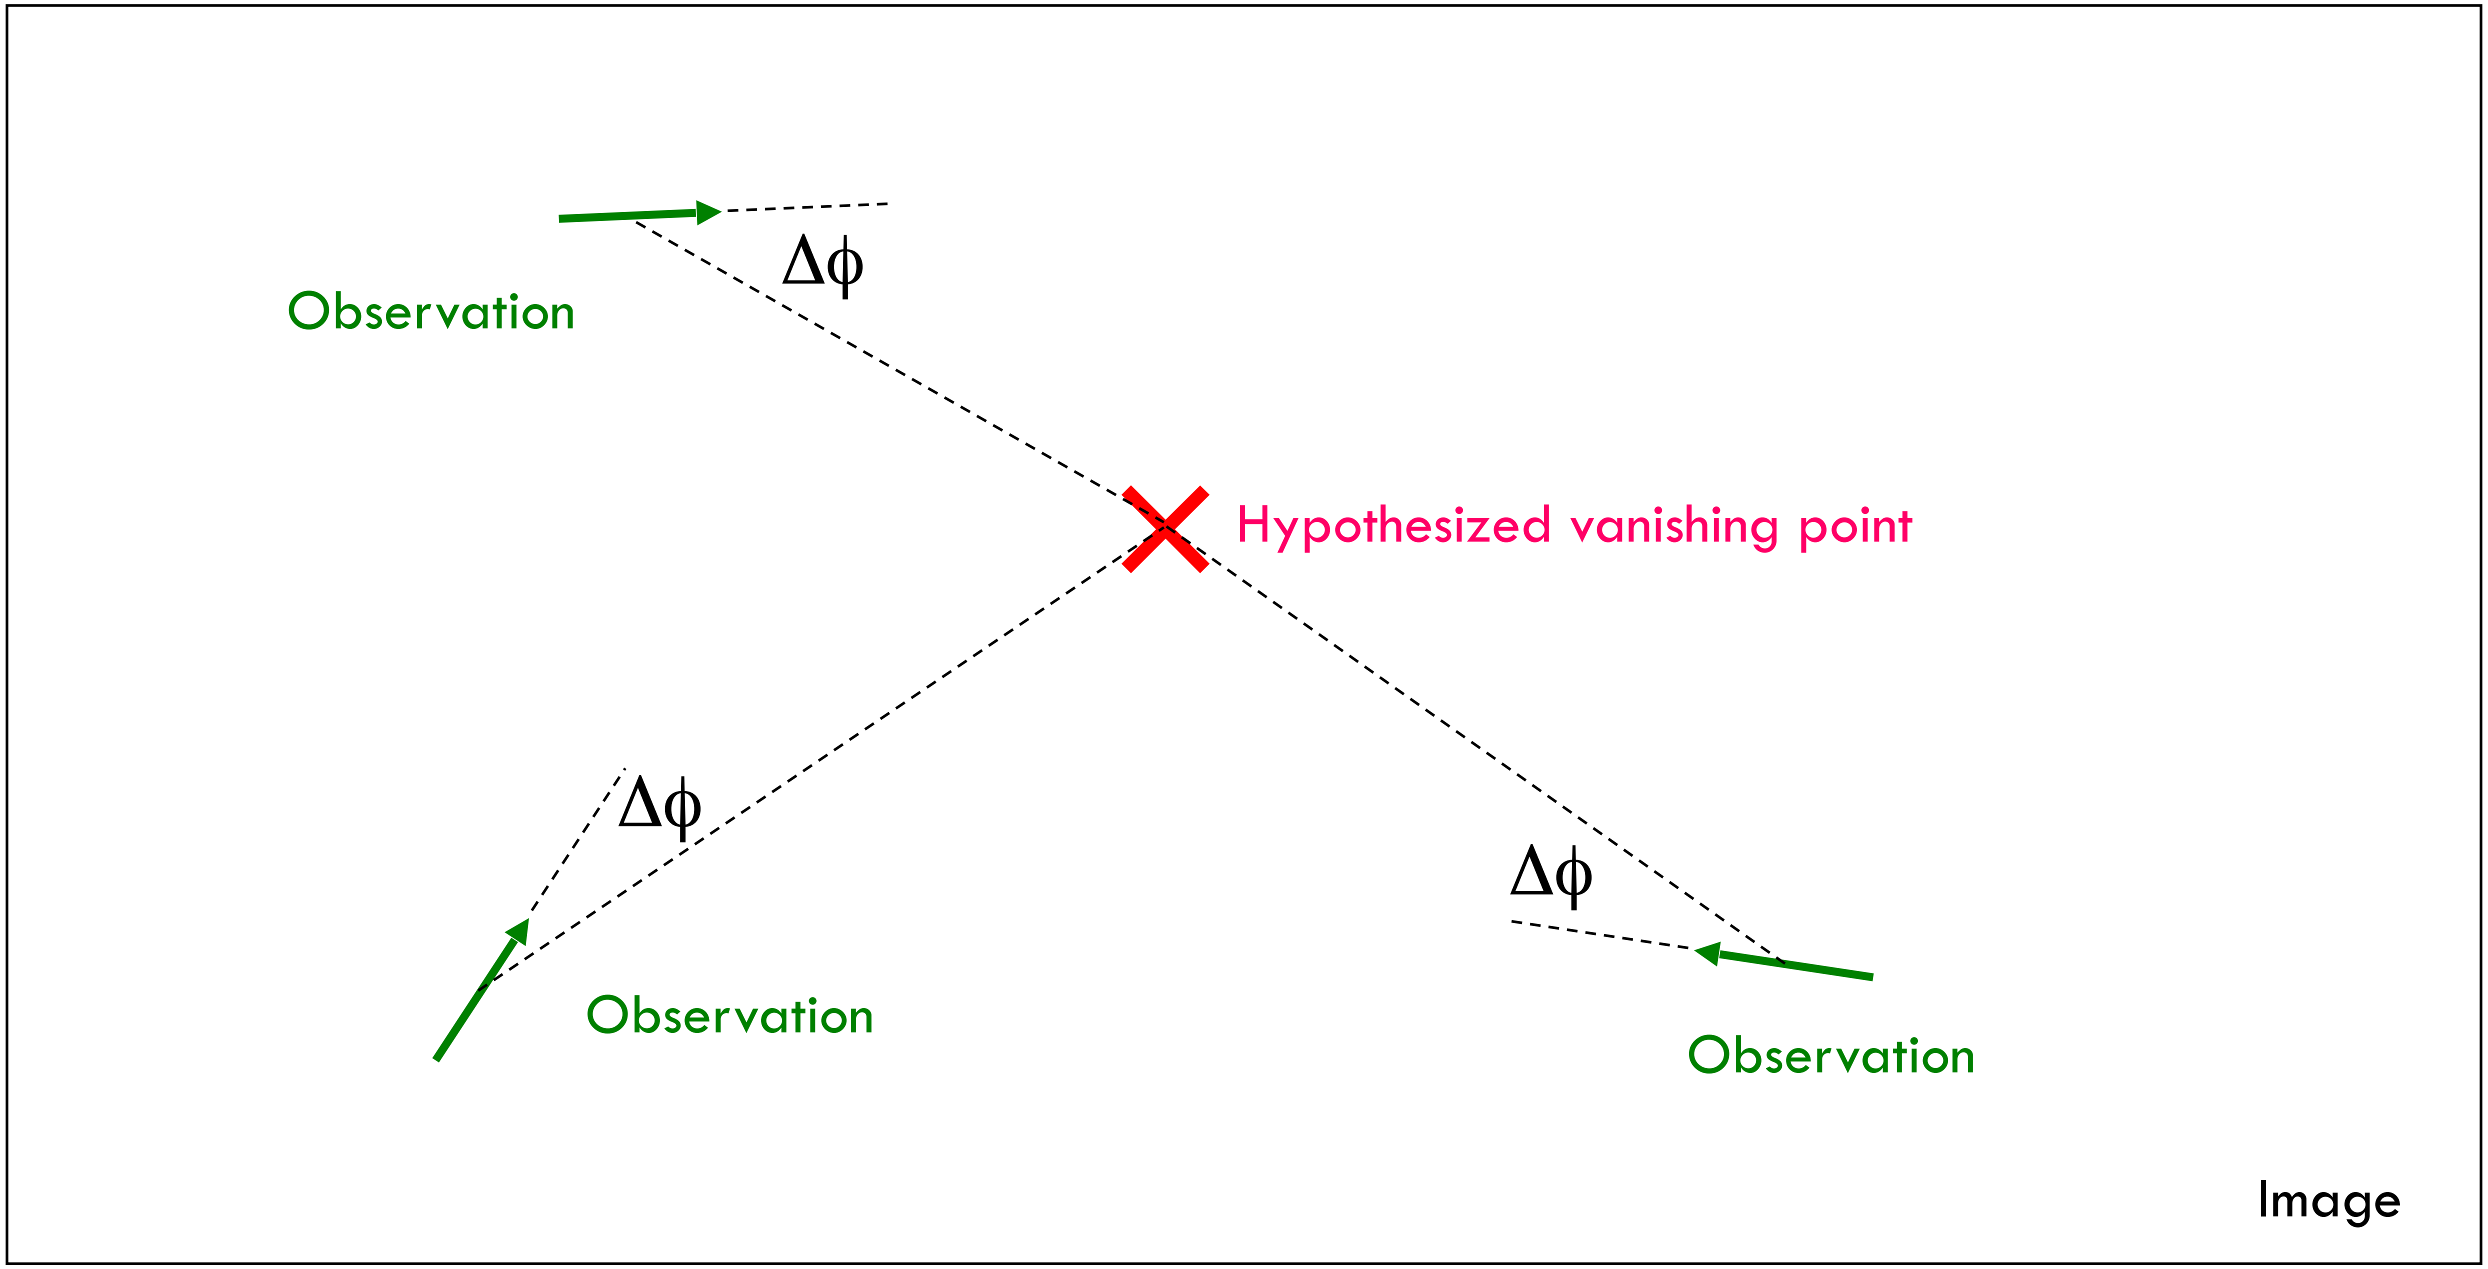

Dense Gradient Map

- In the 1990s/2000s it was thought best to use dense image features to maximize the statistical power of inference.

- Coughlan & Yuille thus proposed to use the dense gradient map, defined at every pixel.

Coughlan & Yuille, 1999

But These are Highly Redundant

Elder 1999
Elder & Goldberg 2001

The Input Space: Squeezing out Redundancy

Image

Gradient Map
(Coughlan & Yuille, 1999)

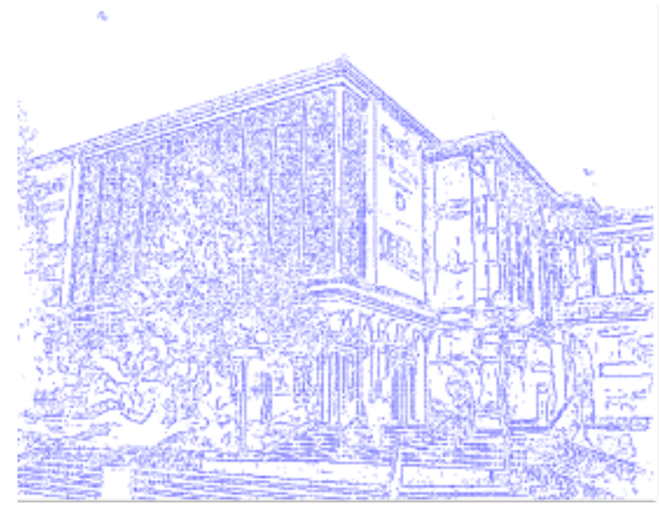

Edge Map
(Denis et al, 2008)

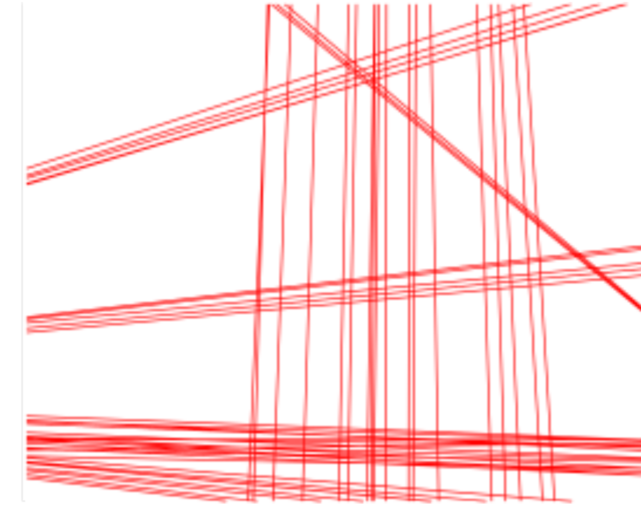

Line Map
(Tal et al, 2012)

Decreasing Redundancy

Sparse Intermediate Representations

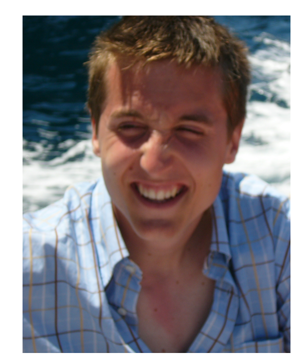

Patrick Denis
Tamgam Systems

Ron Tal
Coinbase

Denis, Elder & Estrada, ECCV 2008
Tal & Elder, ACCV 2012

Manhattan Algorithm Results

Denis, Elder & Estrada, ECCV 2008
Tal & Elder, ACCV 2012

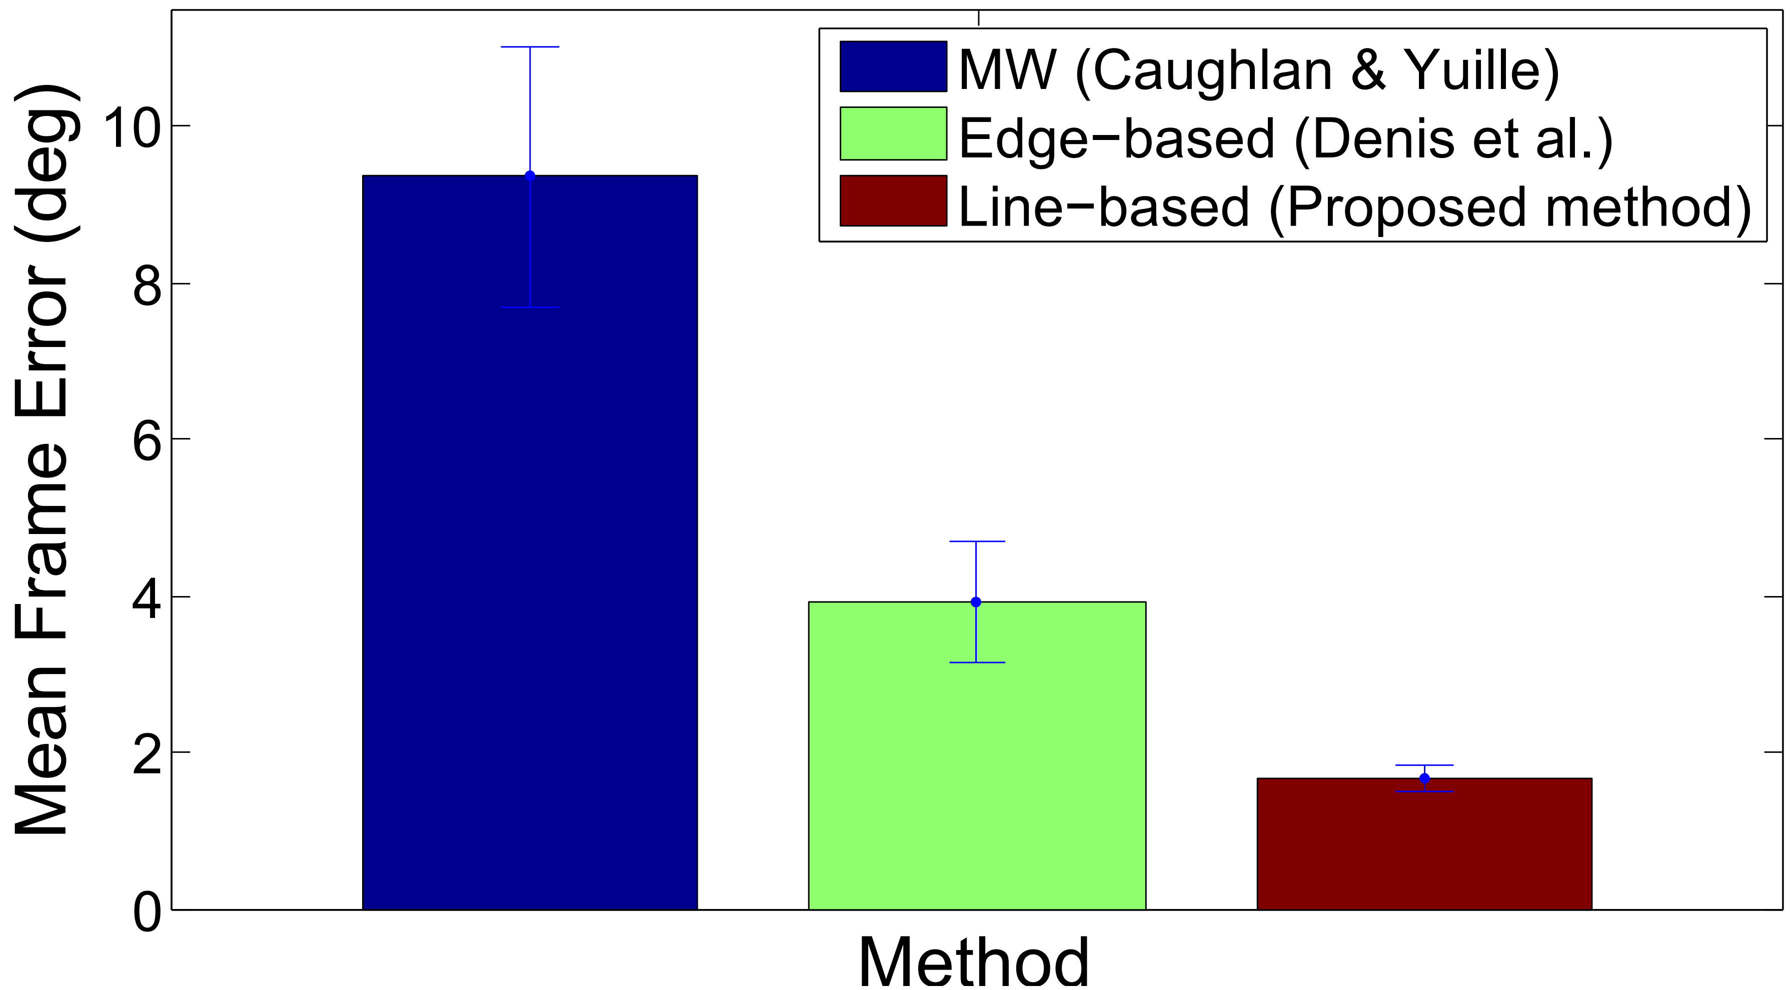

Qualitative Results

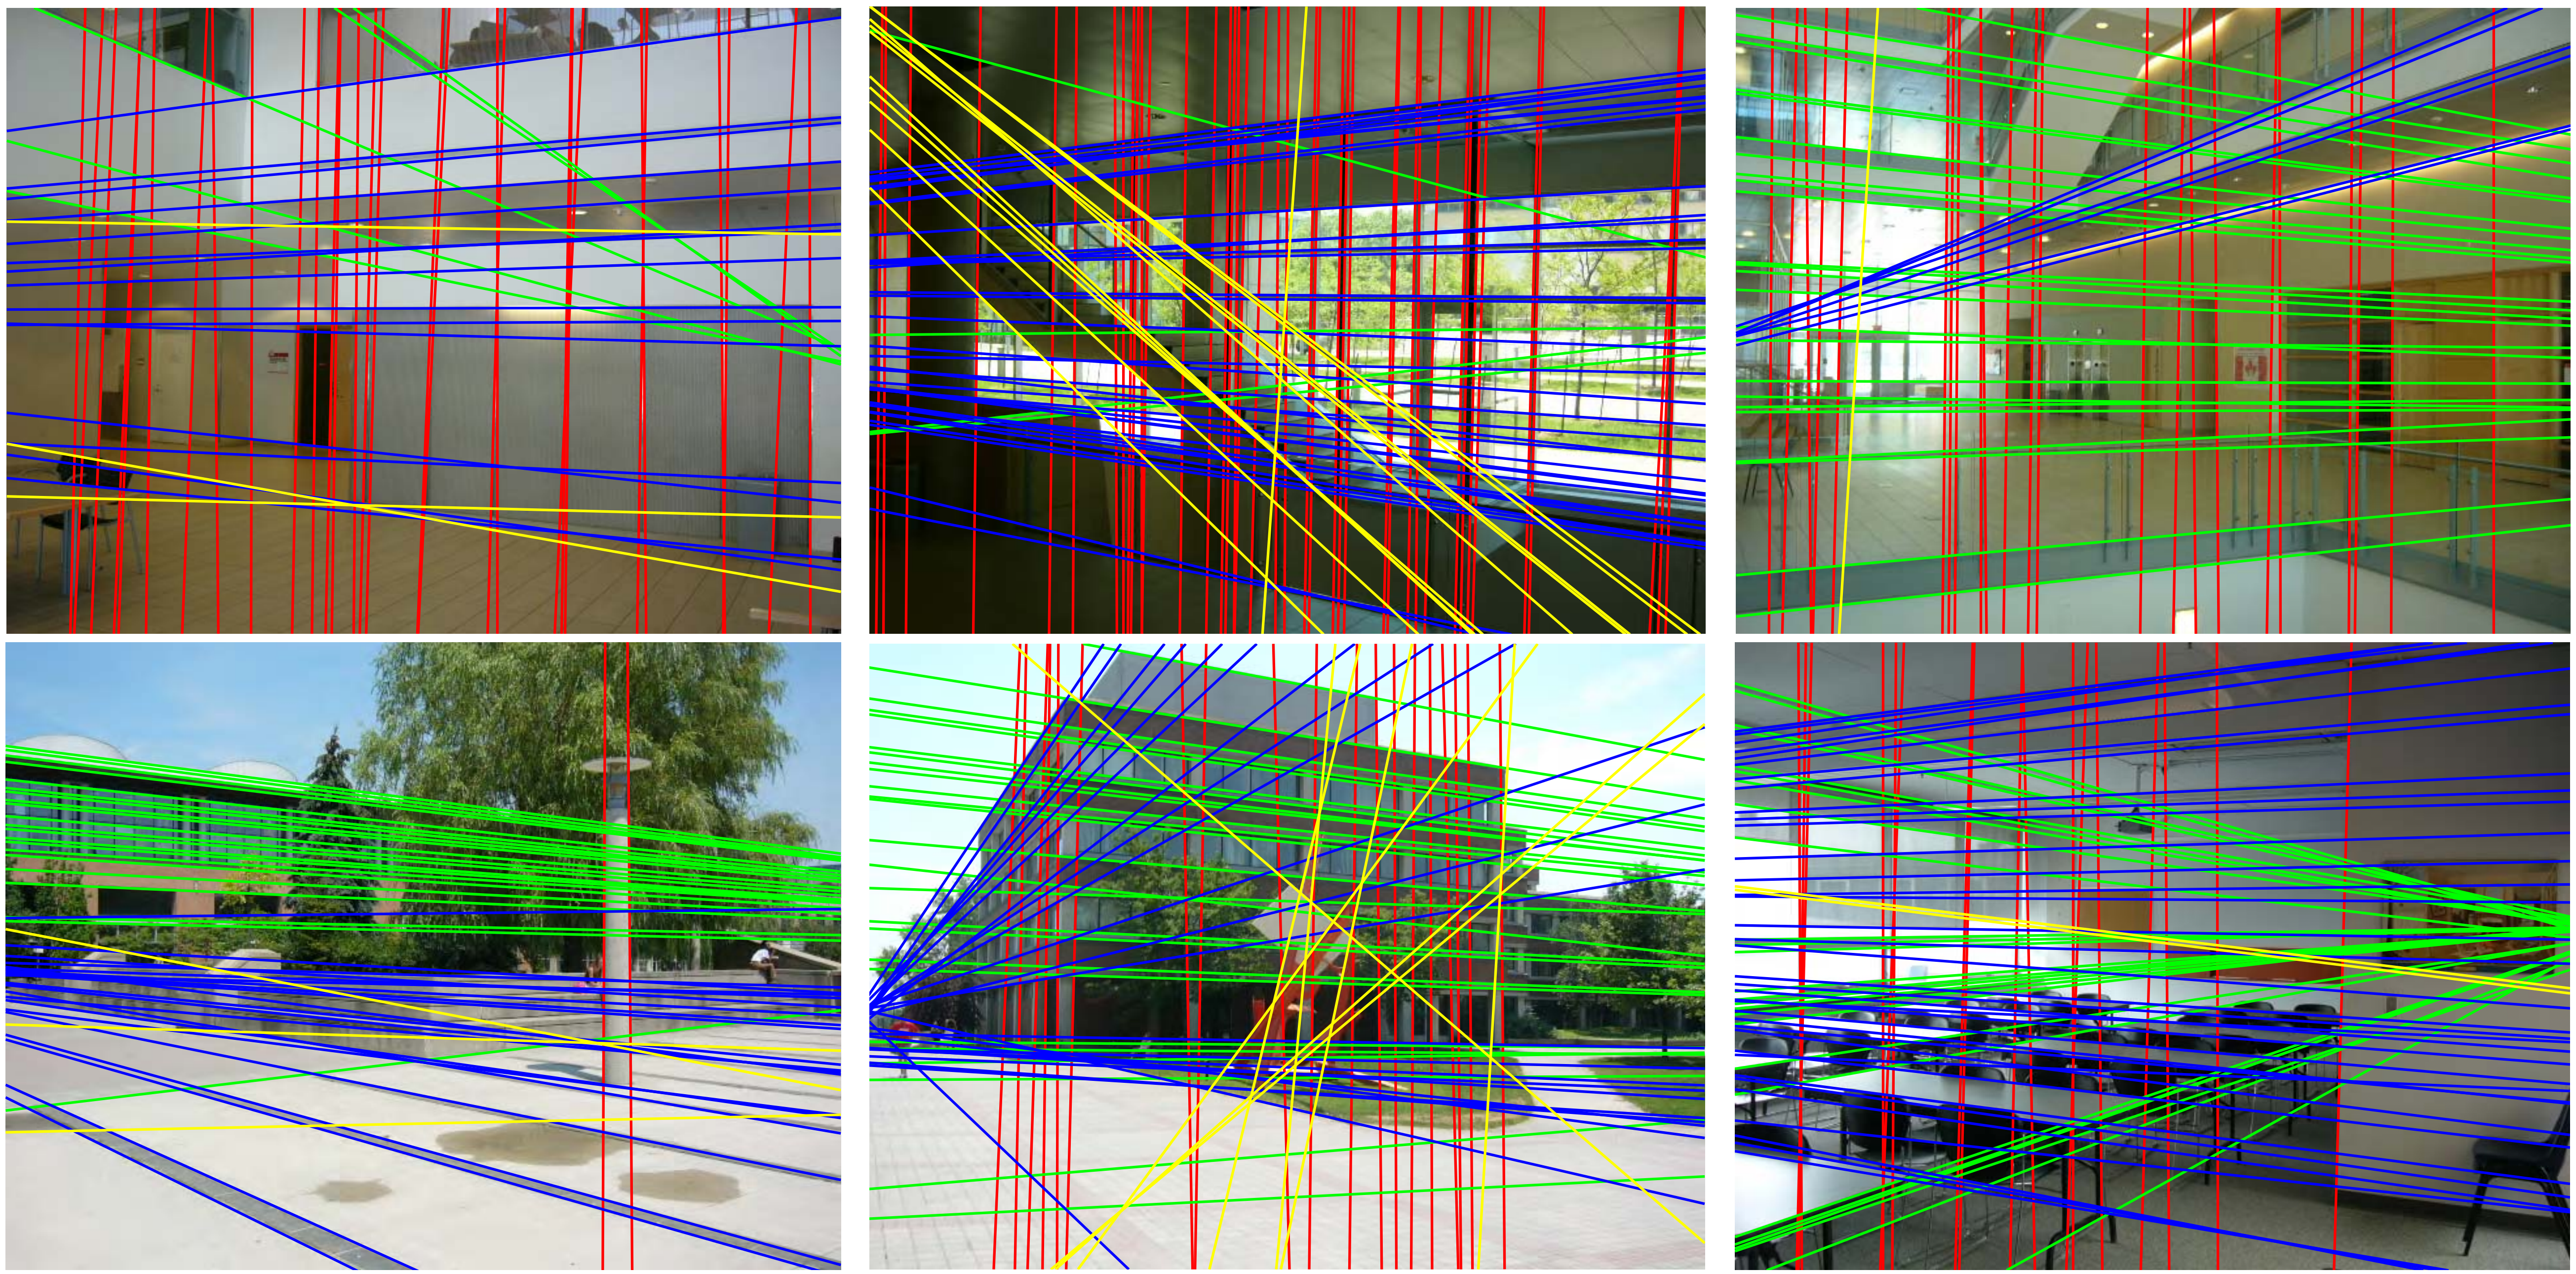

Curvilinear Perspective?

Outline

- Line detection & Hough maps
- Line segment detection
- Vanishing points and Manhattan worlds