E1 216 COMPUTER VISION

LECTURE 06: IMAGE FEATURES

Venu Madhav Govindu
Department of Electrical Engineering
Indian Institute of Science, Bengaluru

2026

Image Features

- Earlier we looked at the **geometry** and **radiometry** of image formation
- In this lecture we shall look at **image features**
- Image features are low-level processing (early vision)

Image Features

What's a feature?

- Encodes *some* meaningful information from an image
- Features can be *local* or *global* properties
- Features need to be *detectable*
- Notion of *meaningful* is context and application dependent
- Too vague and general?

Image Features

What's a feature then?

- **Points**
- Straight Lines
- Parametric Curves
- Edges/Contours
- Texture
- Regions (superpixels)
- Global features

Above list is hierarchical but not exhaustive

Edge Detection

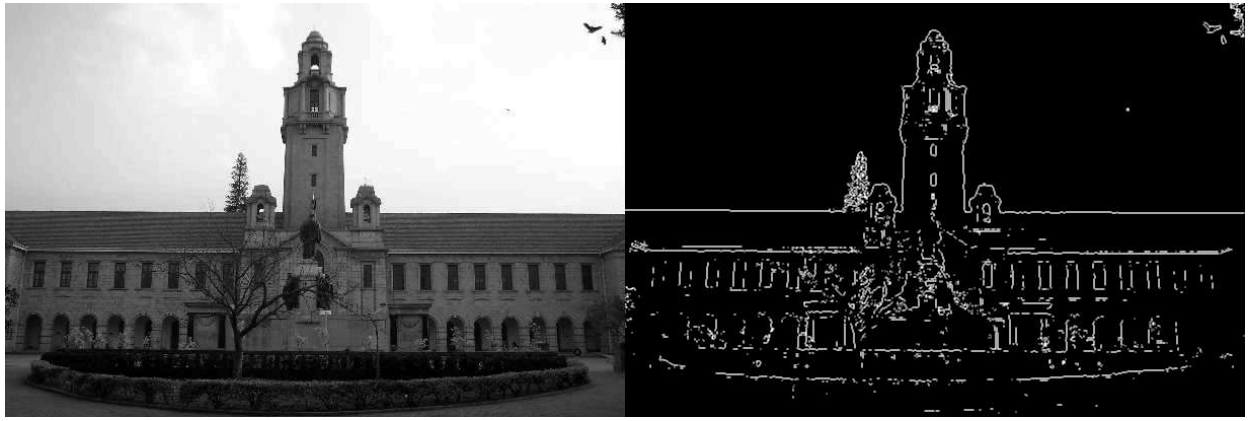

What's an edge ?

Edges are pixels at or around which the image values undergo a sharp variation

Trucco & Verri, p. 69

Edge Detection

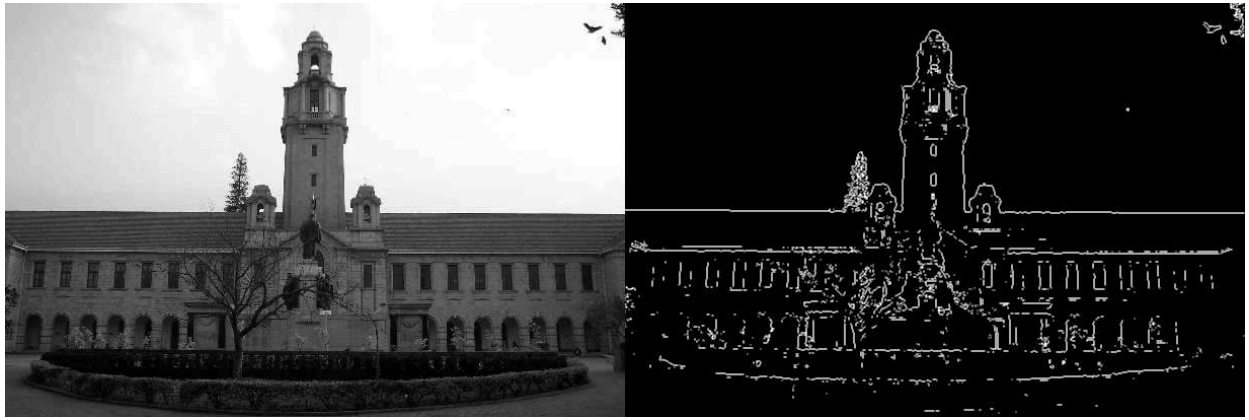

What's an edge?

- A rapid change in brightness level
- Perceptual boundary between regions
- Edge strength and shape can vary
- Detection and localisation are key issues
- Need to counter effects of noise
- Lots of issues, well understood

Edge Detection

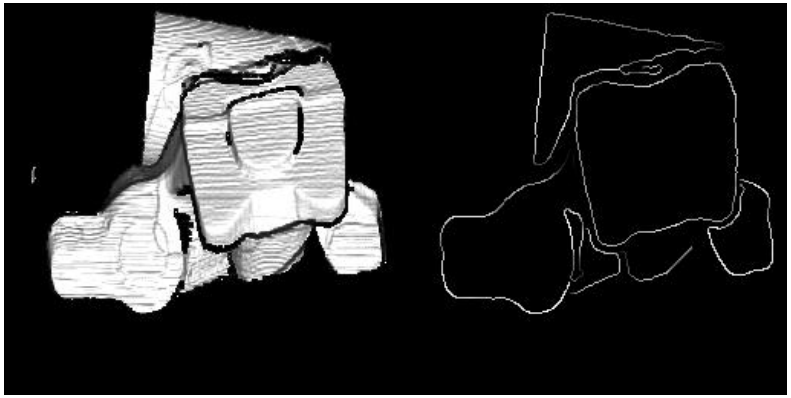

Why are we interested in edges?

- Often correspond to boundaries between regions
- Basic elements for further processing - stereopsis, calibration, motion
- Grouping of edges results in meaningful percepts
- Line drawings are succinct summaries

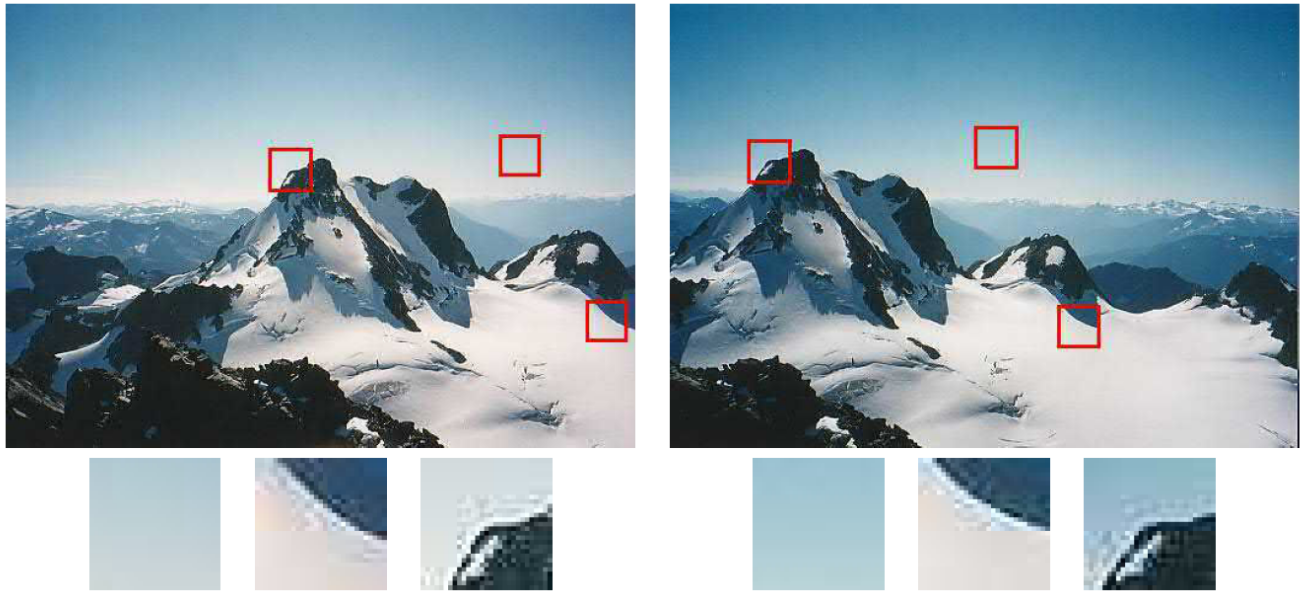

What are corners?

- Corners are point features
- Simple geometry
- Matchable across images
- How?
-

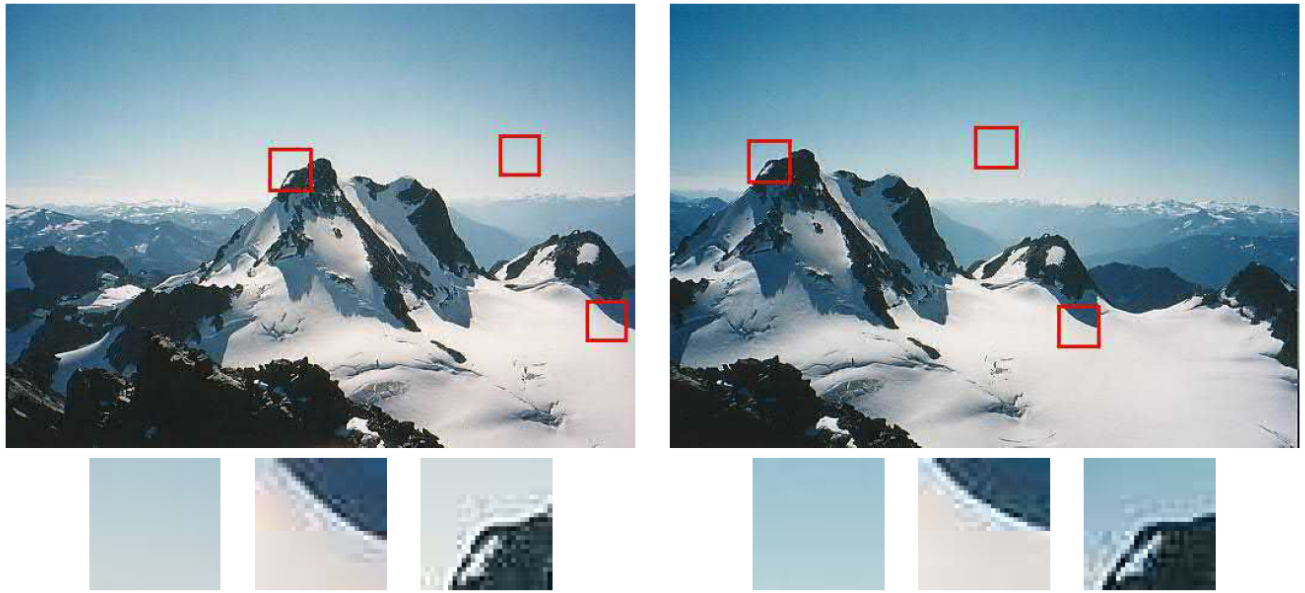

What are corners?

- Corners are point features
- Simple geometry
- Matchable across images
- How? Patches
- All patches equally good?

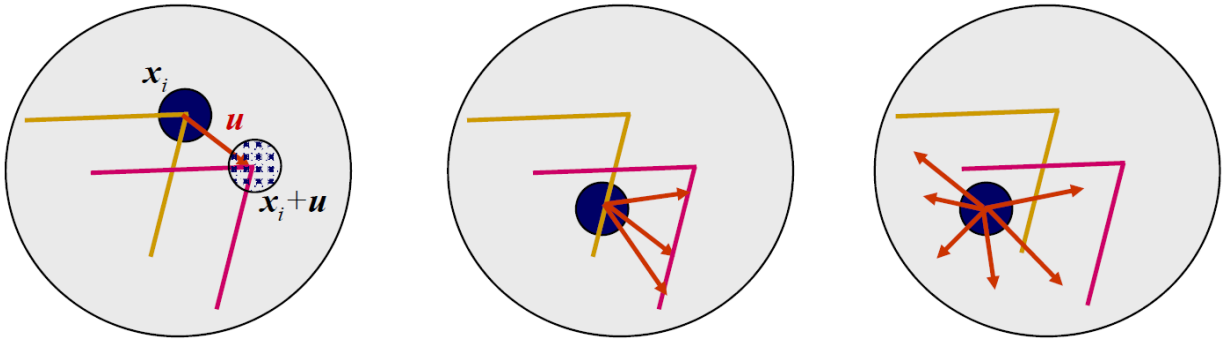

What are corners?

- Corners are point features
- Simple geometry
- Matchable across images
- All patches equally good?
- Discuss feature description later

Szeliski 2nd Edition

Corner Detection

What are corners?

- Edges are easy to define, hard to compute
- Corners are easy to compute
- An *image* corner need not have a physical interpretation
- Corners occur where patterns of intensities represent a perceptual corner

Corner Detection

Corner Detection

Why corners ?

- Corner points represent sharp changes/discontinuties
- Corners are ubiquitous
- Corners are "relatively" stable to viewpoint variation
- Corners can be matched across images
- Corners can be easily tracked across a video sequence
- Vision problems can be easily defined and modelled using point features
- Multiview geometry of point features is well developed
- Computations using point features is efficient

Corner Detection

Uses of corners ?

- Geometrically simple and well-defined
- Succinct representation of information
- Efficient computations
- Attach features to corners (matching, discrimination)
- Point matches for geometry
- Image retrieval (similar ones)
- Comparing image similarity
- Fundamental in computer vision systems

Corner Detection

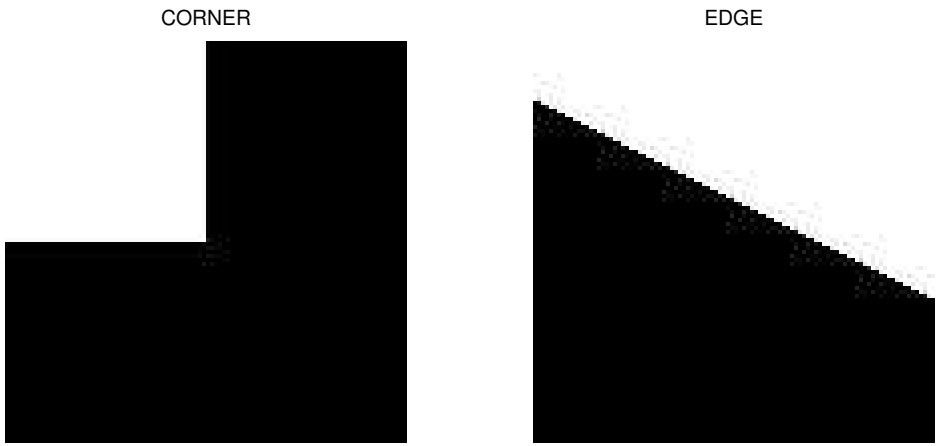

So what are corners ?

- An edge represents intensity variation in one direction
- Intensity variation in orthogonal directions is much lower
- A corner represents intensity variation in both X and Y directions
- For image $f(x, y)$, $\frac{\partial f}{\partial x}$ and $\frac{\partial f}{\partial y}$ for corners are large

Corner Detection

Notion of large $\frac{\partial f}{\partial x}$ and $\frac{\partial f}{\partial y}$ suggests an approach

Structure Tensor

Denote $f_x = \frac{\partial f}{\partial x}$ and $f_y = \frac{\partial f}{\partial y}$

$$C = w_G(\sigma) * \begin{bmatrix} f_x^2 & f_x f_y \\ f_x f_y & f_y^2 \end{bmatrix} = \begin{bmatrix} \widehat{f_x^2} & \widehat{f_x f_y} \\ \widehat{f_x f_y} & \widehat{f_y^2} \end{bmatrix}$$

- Matrix is known as *local structure matrix* or *tensor*
- Encodes the structural variation of intensities around a point
- Corner-ness can be easily defined using matrix C
- The local structured is spatially smoothed using a Gaussian mask (denoted by hat)

Image Corners

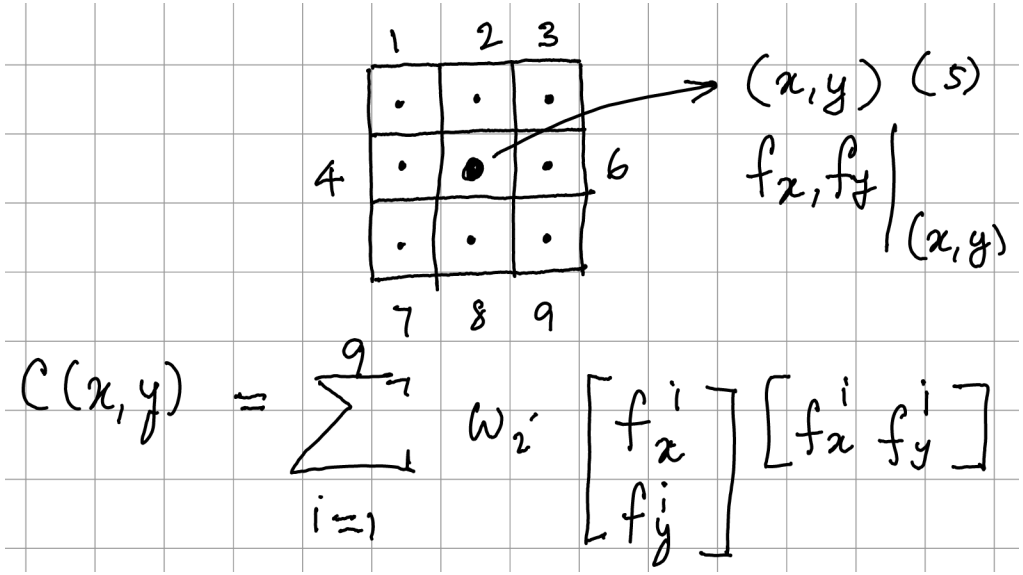

Structure Tensor

- Aggregate information over $N \times N$ neighbourhood of (x, y)
- Size of N ?
- Weighting of individual contributions
- $w_i = \frac{1}{N^2}$ (Uniform)
- $w_i \propto \exp\{-\frac{1}{\sigma^2}((x - x_o)^2 + (y - y_o)^2)\}$ (Gaussian)
- Form of w_i not very important
- Careful computation of (f_x, f_y) is essential. Why ?

Image Corners

$$C = \begin{bmatrix} \widehat{f_x^2} & \widehat{f_x f_y} \\ \widehat{f_x f_y} & \widehat{f_y^2} \end{bmatrix}$$

Properties

- For a single point, C is rank-1
- The structure emerges out of local averaging
- Matrix C is *symmetric*
- Matrix C is *positive-definite*
- Use these properties to detect corners

Image Corners

- Symmetry of $C \Rightarrow C$ is diagonalisable via a rotation U
- $C = UDU^T$ (**Proof?**)
- **Exercise**: What happens to f_x and f_y under rotation ? C ?
- Should we care about the rotation ? Why ?
- Diagonalised matrix will be of the form

$$D = C' = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix}$$

- Positive-definite means the eigen-values are non-negative
  Wlog assume $\lambda_1 \geq \lambda_2 \geq 0$

Corner Detection

Harris Corner Detector

- Proposed in 1988
- Quite popular till recent improvements
- Defines a measure of *corner strength*
- Detection using a threshold
- Uses neighbours to discard weak corners

Local Structure Matrix : Interpretation

- You may wonder why the given structure matrix C
- Consider the local auto-correlation function of the image
- For a shift of $(\Delta x, \Delta y)$ we have

$$R(\Delta x, \Delta y) = \sum_{W} \left[I(x, y) - I(x + \Delta x, y + \Delta y) \right]^2$$

- Taking Taylor series for a small shift $(\Delta x, \Delta y)$ we have

$$I(x + \Delta x, y + \Delta y) \approx I(x, y) + \begin{bmatrix} \frac{\partial I}{\partial x} & \frac{\partial I}{\partial y} \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix}$$

Harris Corner Detector

$$
\begin{aligned}
I(x+\Delta x, y+\Delta y) &\approx I(x,y) + \begin{bmatrix} \frac{\partial I}{\partial x} & \frac{\partial I}{\partial y} \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} \\
\Rightarrow R(x,y) &= \sum_W \left(\begin{bmatrix} \frac{\partial I}{\partial x} & \frac{\partial I}{\partial y} \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} \right)^2
\end{aligned}
$$

Rearranging, we have

$$
\begin{aligned}
R(x,y) &= \sum_W \begin{bmatrix} \Delta x & \Delta y \end{bmatrix} \begin{bmatrix} \sum_W {I_x}^2 & \sum_W I_x I_y \\ \sum_W I_x I_y & \sum_W {I_y}^2 \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} \\
\Rightarrow C &= \begin{bmatrix} \sum_W {I_x}^2 & \sum_W I_x I_y \\ \sum_W I_x I_y & \sum_W {I_y}^2 \end{bmatrix}
\end{aligned}
$$

Interpretation of C matrix

- Both λ_1 and λ_2 small
  'flat' image function
- One λ large and the other small
  Auto-correlation function is ridge shaped
  Image intensity changes in one direction only
  This is the behaviour of an edge
- Both λ are large
  sharp fall in auto-correlation in any direction
  characteristic of a **corner**

Harris Corner Detector

Corner Strength

- Defined as

$$H(x,y) = |C| - \alpha Tr(C)^2$$

where $Tr()$ is the trace of a matrix.

This form of H means we do not need to explicitly compute λ_1 and λ_2

- For diagonalised matrix C we have

$$\begin{aligned} H(x,y) &= \lambda_1\lambda_2 - \alpha(\lambda_1 + \lambda_2)^2 \\ H(x,y) &= {\lambda_1}^2(\kappa - \alpha(1+\kappa)^2) \end{aligned}$$

where $\kappa = \frac{\lambda_2}{\lambda_1}$

- Corner detected when $H(x,y) \geq H_0$ (threshold)

Analysis

Assuming $H > 0$ we have

$$0 \leq \alpha \leq \frac{\kappa}{(1+\kappa)^2} \leq 0.25$$

For small κ we have

$$H \approx {\lambda_1}^2(\kappa - \alpha), \alpha \lesssim \kappa$$

Harris Corner Detector

$$H(x, y) = \lambda_1\lambda_2 - \alpha(\lambda_1 + \lambda_2)^2$$

- Larger $\alpha \Rightarrow$ smaller $H \Rightarrow$ less sensitive detector: less corners detected.
- Smaller $\alpha \Rightarrow$ larger $H \Rightarrow$ more sensitive detector: more corners detected.

Usually, H_{thr} is set close to zero and fixed, while α is a variable parameter.

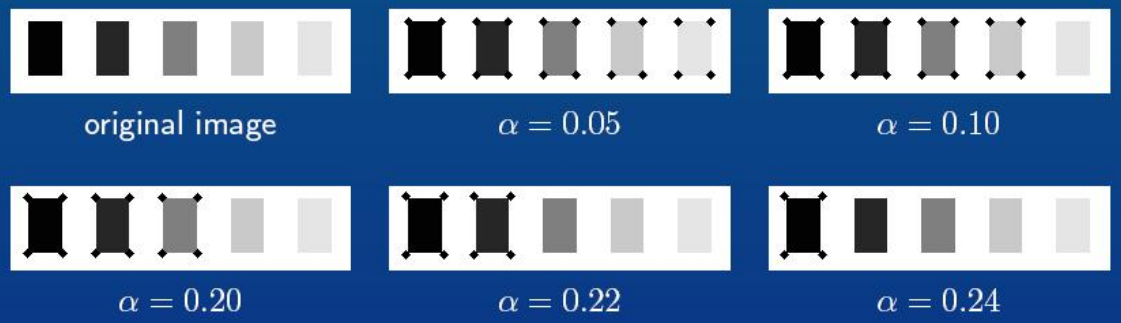

Corner detection by Harris operator: influence of α. ($H_{thr} = 0$.)

Taken from slides of Dmitrij Csetverikov

Refer to slides on **Harris Corner Detector** linked on page for this lecture

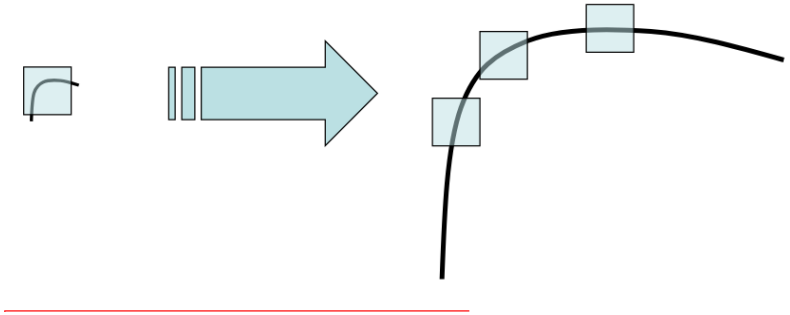

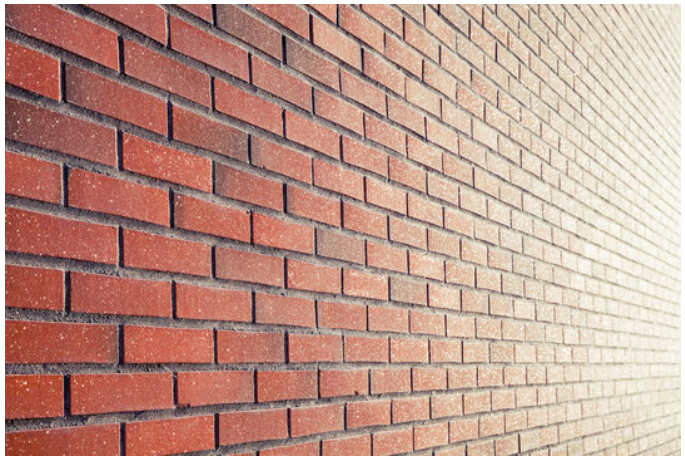

Role of Perspective Projection

- Scale of Harris corner detection?
- Multiple scales in same image?
- Only scale?
- Perspective plays a role
- What do we want?
-

adapted from K. Graumann; stock.adobe.com

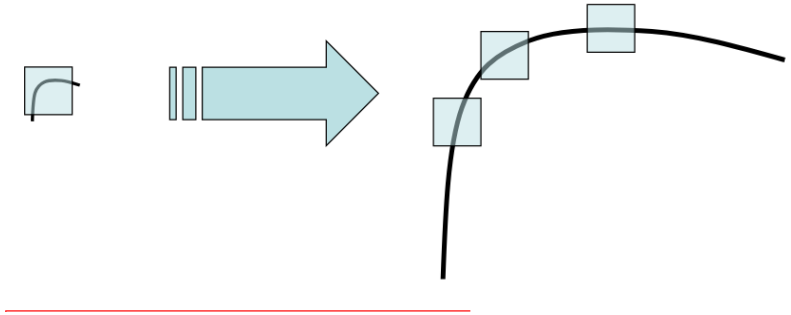

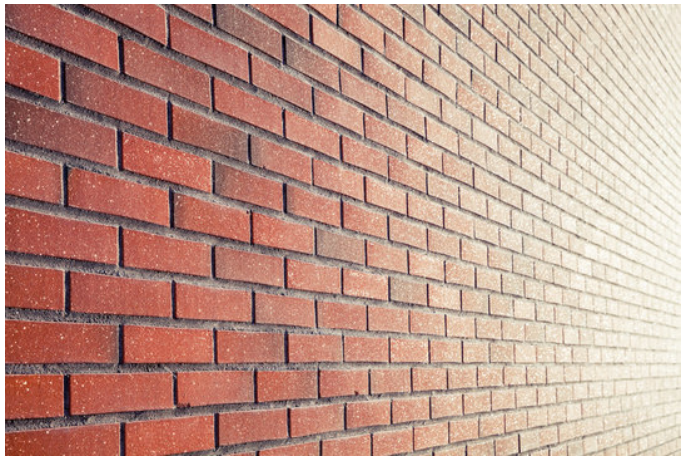

Role of Perspective Projection

- Scale of Harris corner detection?
- Multiple scales in same image?
- Only scale?
- Perspective plays a role
- What do we want?
- Geometric Invariance

adapted from K. Graumann; stock.adobe.com

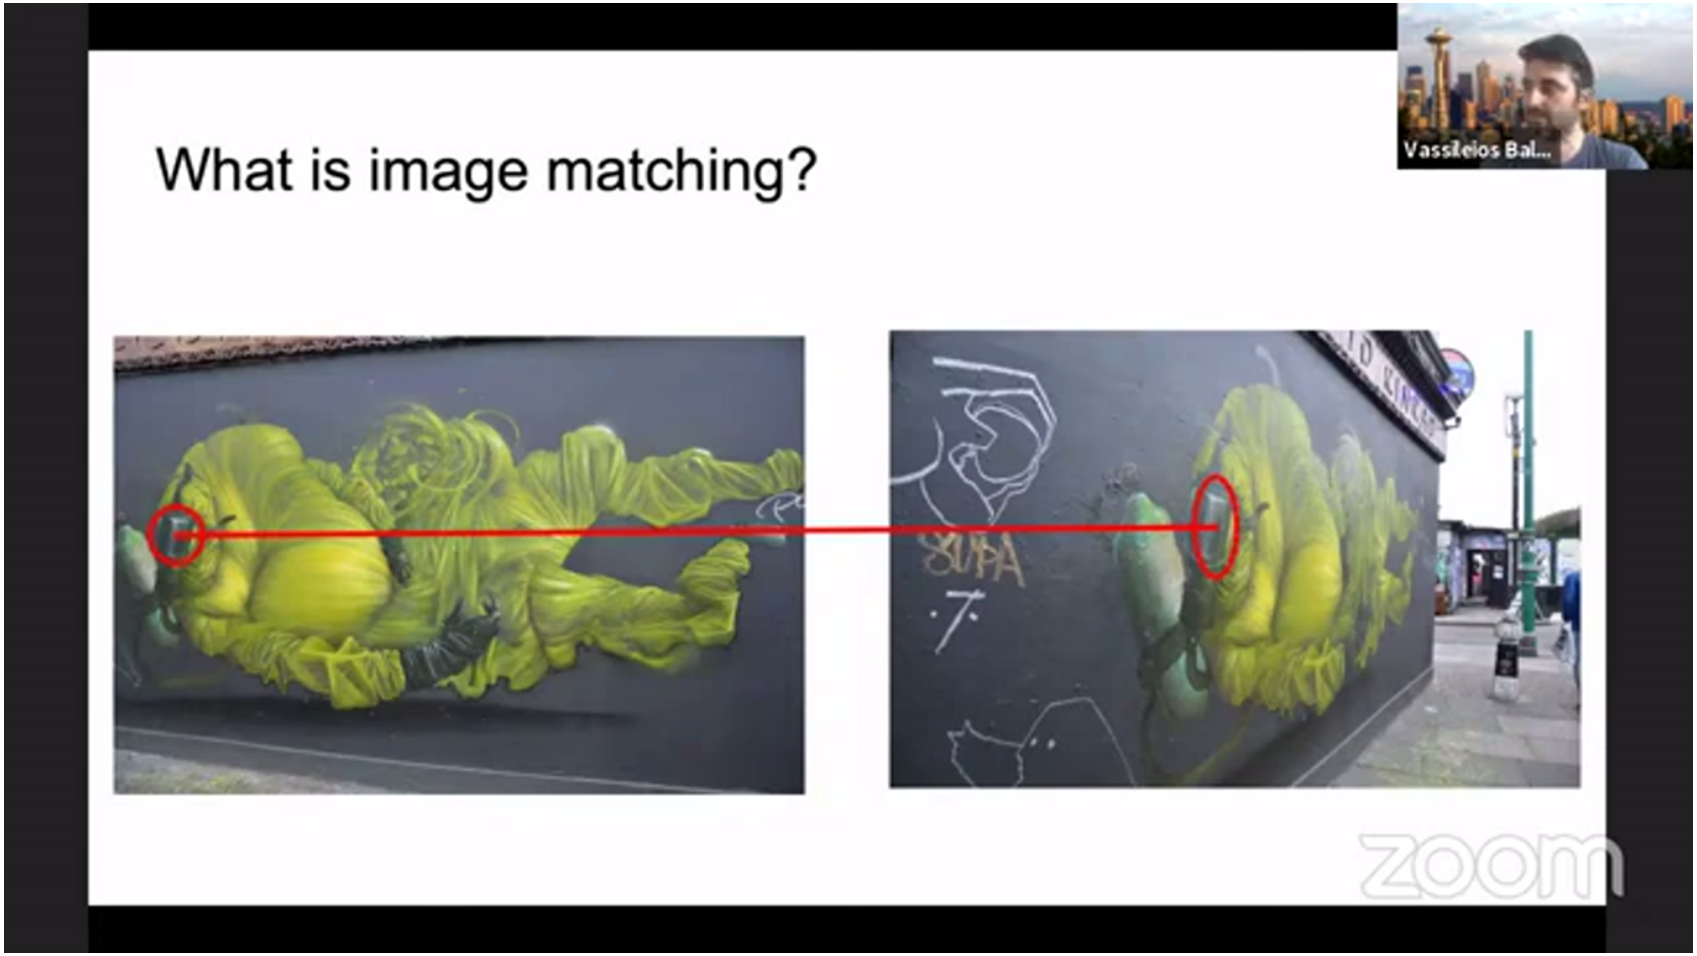
What is image matching?
Vassileios Bal...
zoom

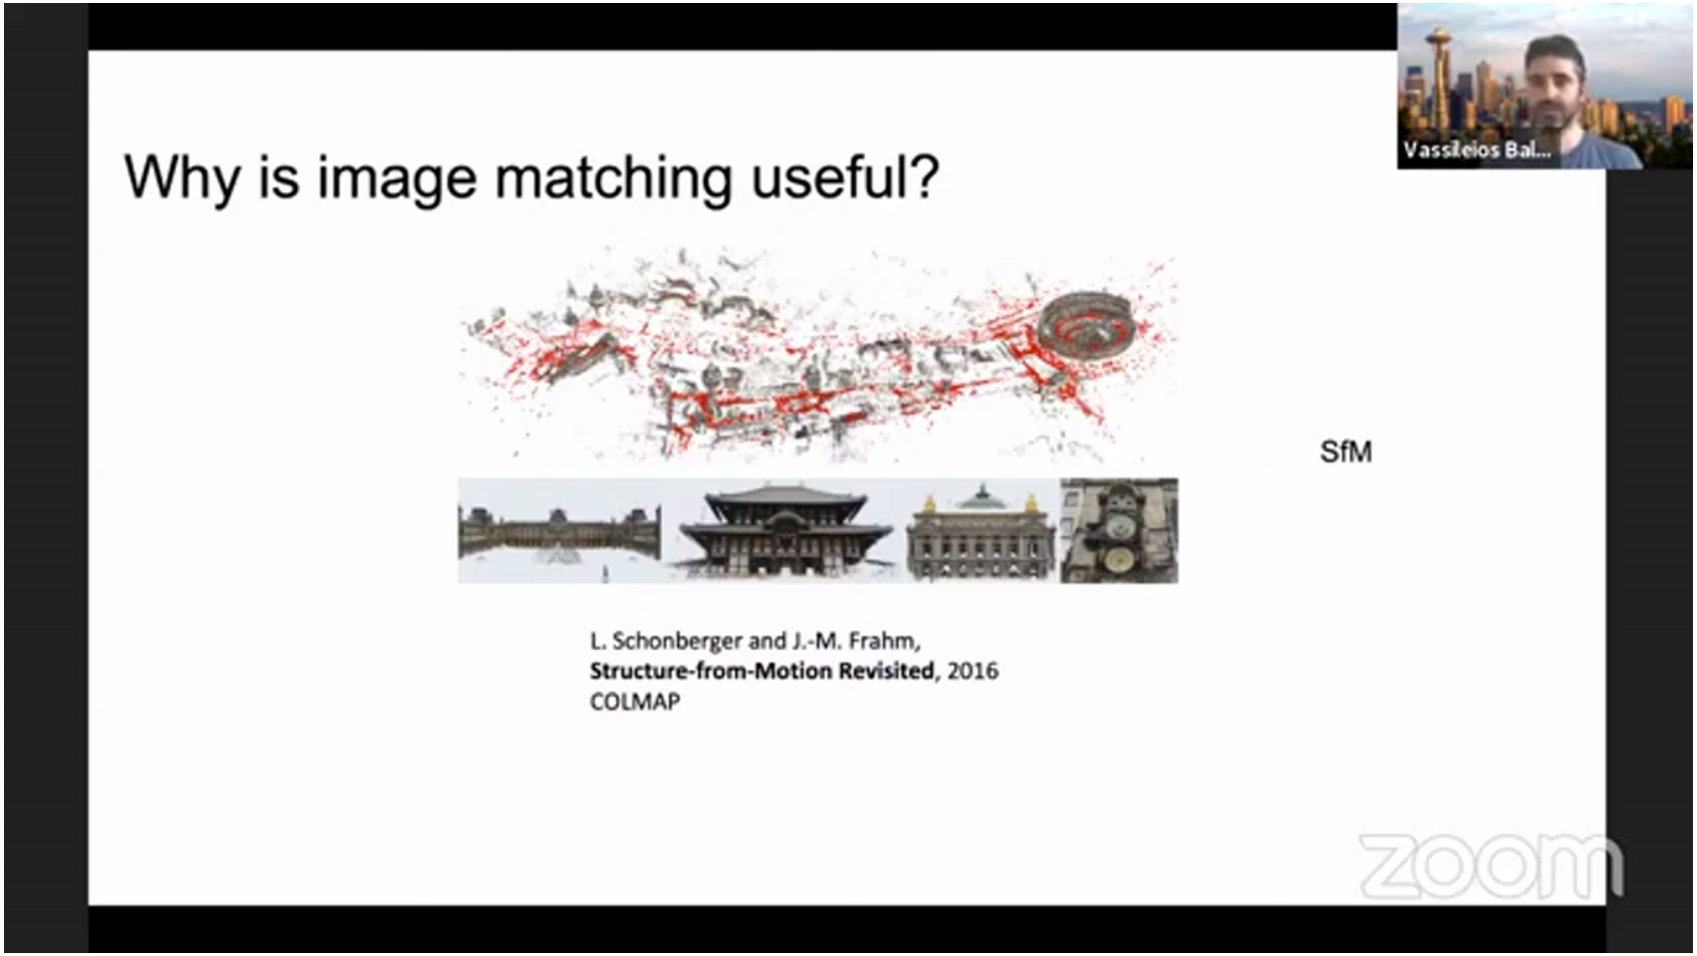
Vassileios Bal...
Why is image matching useful?
SfM
L. Schonberger and J.-M. Frahm,
Structure-from-Motion Revisited, 2016
COLMAP
zoom

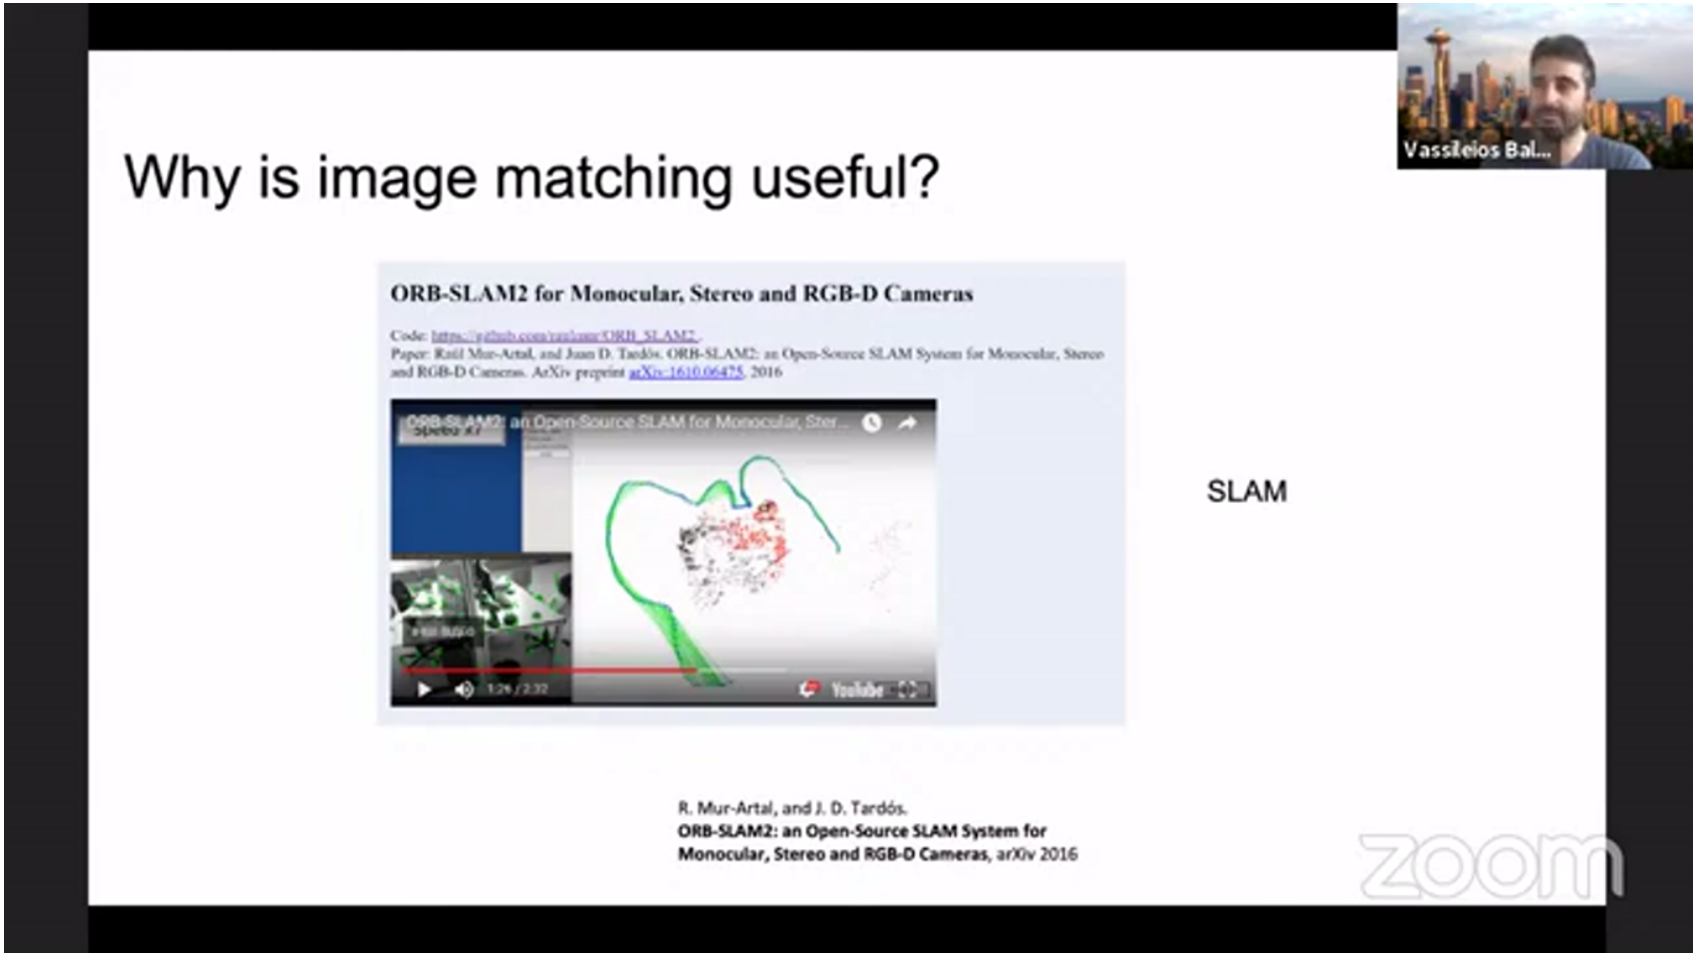
Vassileios Bal...
Why is image matching useful?
ORB-SLAM2 for Monocular, Stereo and RGB-D Cameras
Code: https://github.com/raulmur/ORB_SLAM2.
Paper: Raúl Mur-Artal, and Juan D. Tardós. ORB-SLAM2: an Open-Source SLAM System for Monocular, Stereo and RGB-D Cameras. ArXiv preprint arXiv:1610.06475. 2016
ORB-SLAM2: an Open-Source SLAM for Monocular, Ster...
SLAM
R. Mur-Artal, and J. D. Tardós.
ORB-SLAM2: an Open-Source SLAM System for Monocular, Stereo and RGB-D Cameras, arXiv 2016
zoom

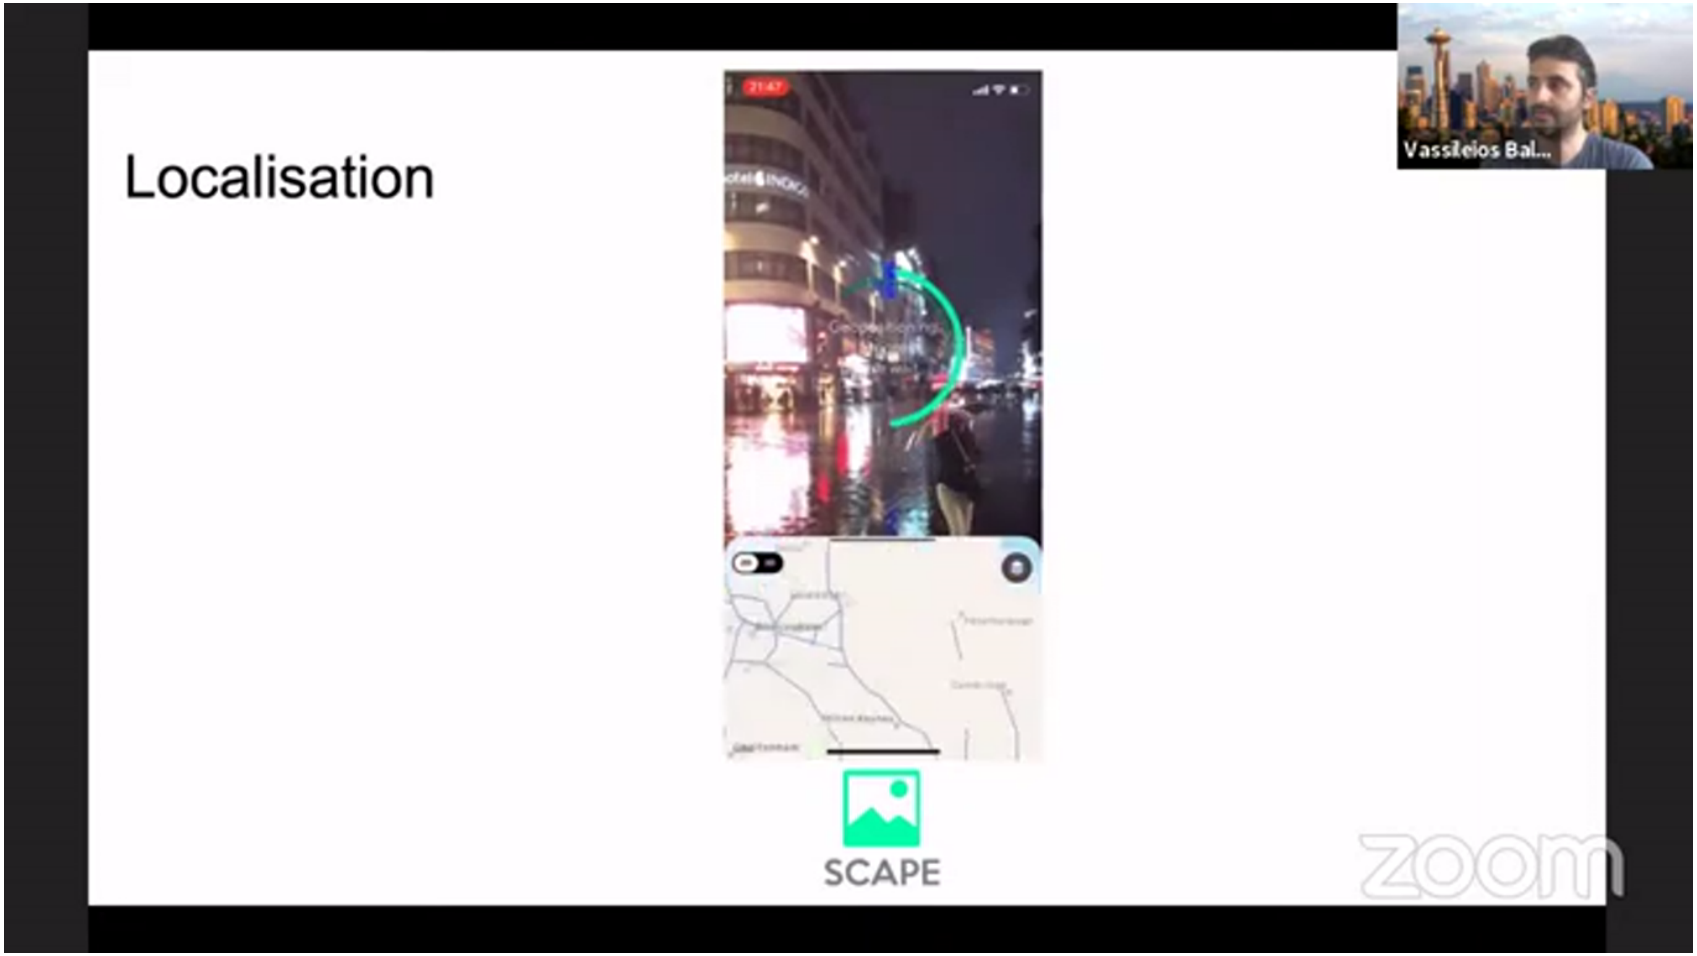
Localisation
SCAPE
Vassileios Bal...
zoom

CVPR 2020 Tutorial on 'Local Features'

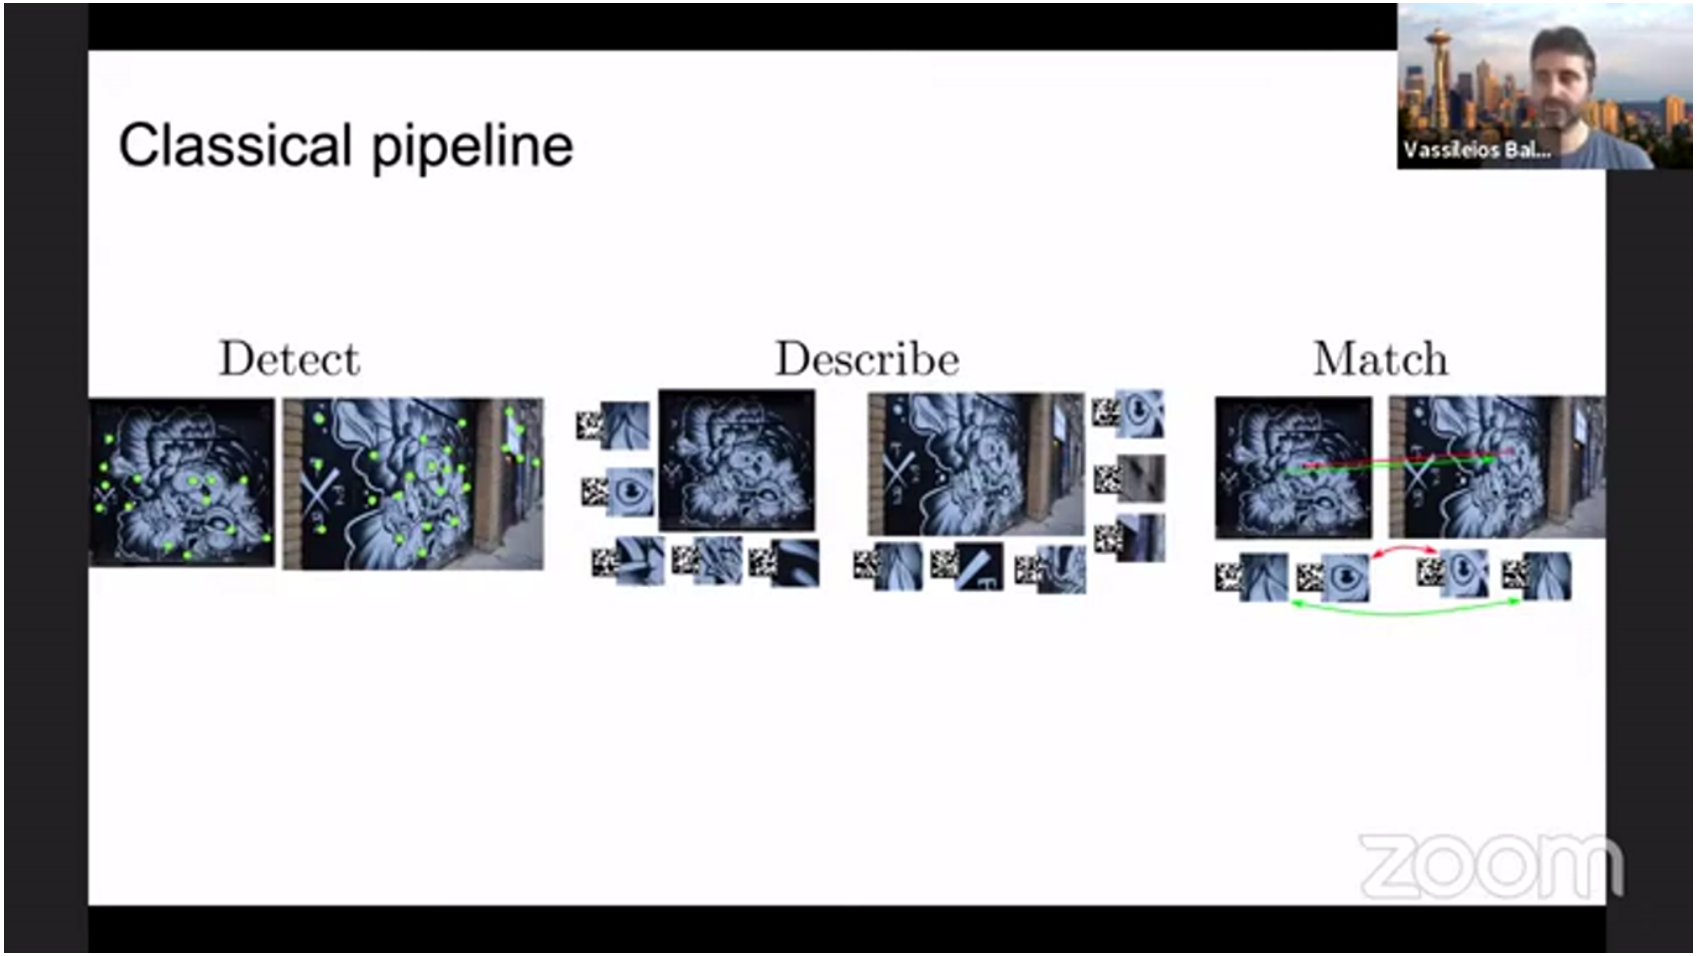

CVPR 2020 Tutorial on ‘Local Features’

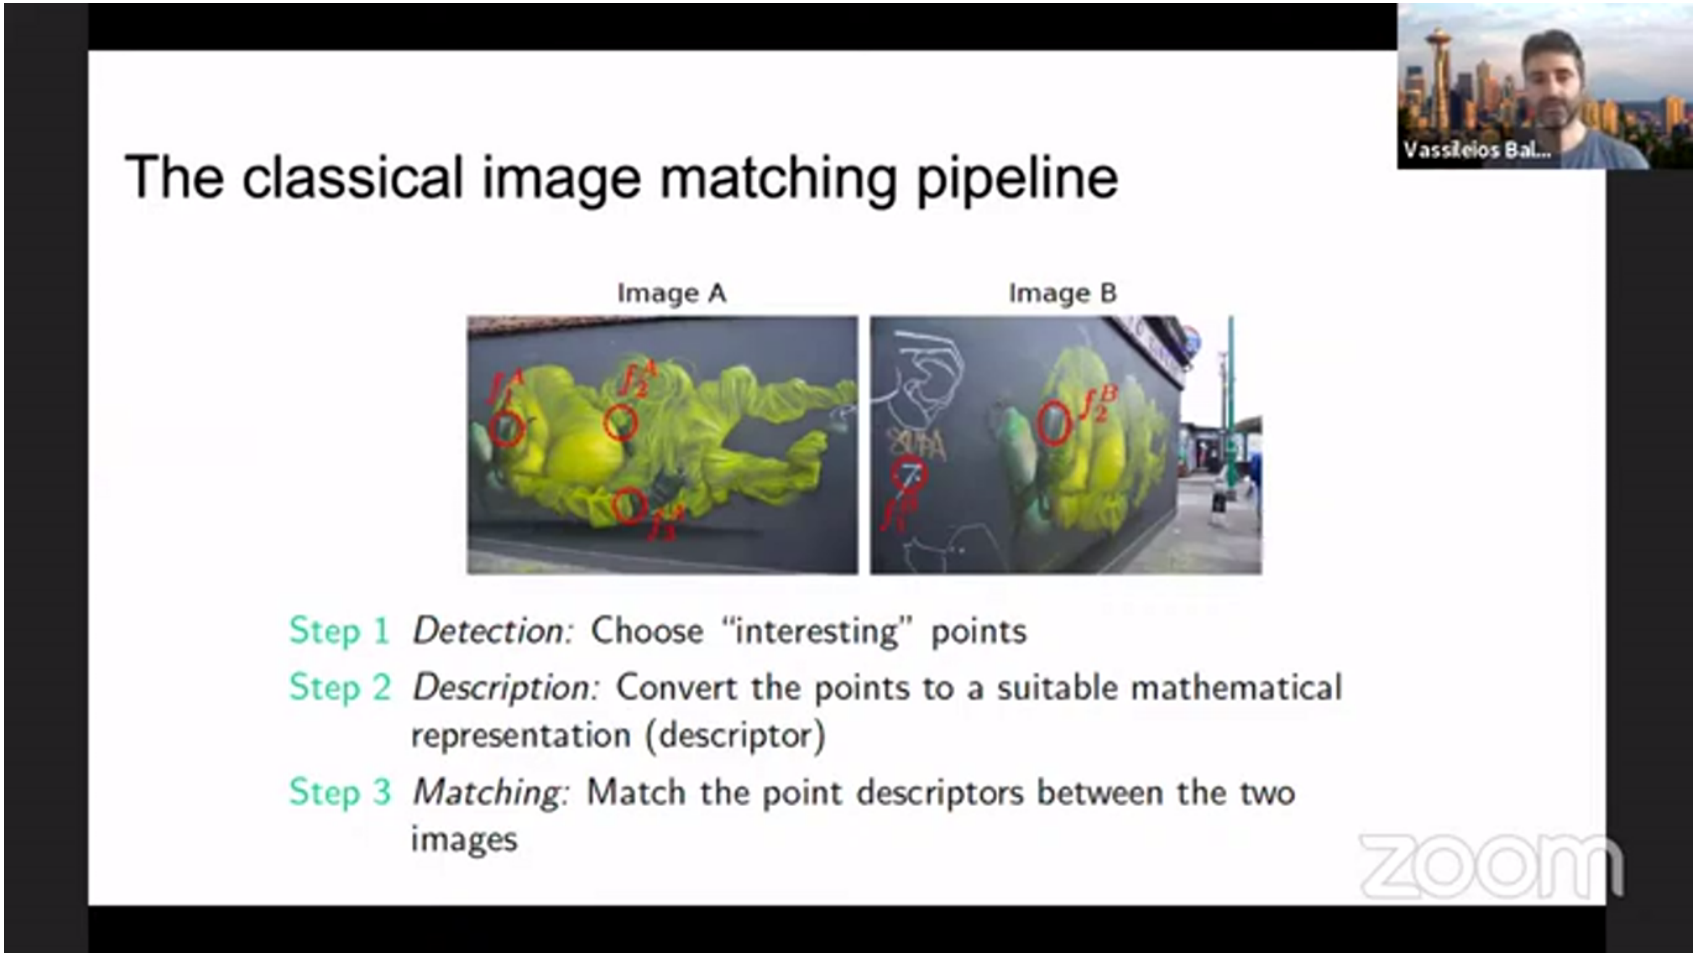
Vassileios Bal...
The classical image matching pipeline
Image A
Image B
Step 1 *Detection:* Choose "interesting" points
Step 2 *Description:* Convert the points to a suitable mathematical representation (descriptor)
Step 3 *Matching:* Match the point descriptors between the two images
zoom

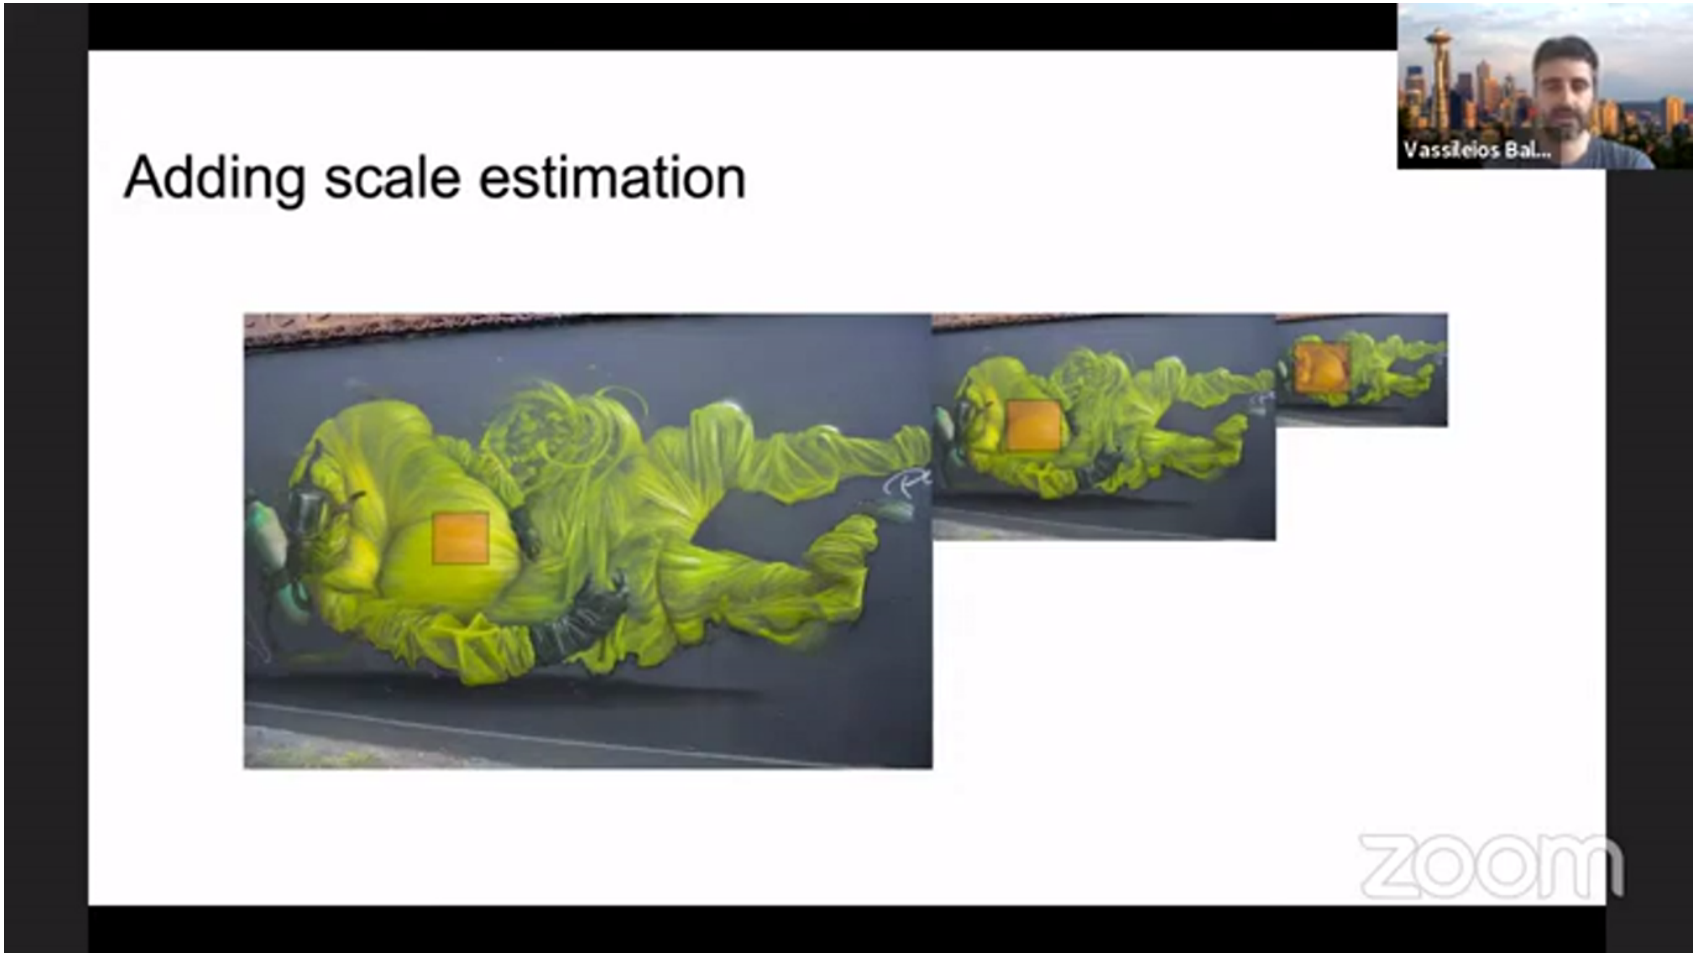
Vassileios Bal...
Adding scale estimation
zoom

Refer to slides on **SIFT** linked on page for this lecture

Image Features

Issues not addressed

- SIFT is *scale* invariant
- SIFT used as generic feature representation
- Affine invariance achieved by other features
- Faster versions of SIFT using integral images etc. (SURF)
- Histogram of Oriented Gradients (HOG) feature used for detection
- Key bottleneck : Feature matching in high-dim (128 dim)
- k-d tree representation
- Approximate nearest neighbour search in high-dim
- Binary features
- Feature descriptor is *local* $\Rightarrow$ false matches
- Need outlier removal for geometry estimation
- **Recent progress**: End-to-end learning
- CVPR 2020 Tutorial: 'Local Features: From SIFT to Differentiable Methods'