

Computer Vision 4. Interest Points

Lars Schmidt-Thieme

Information Systems and Machine Learning Lab (ISMLL) Institute for Computer Science University of Hildesheim, Germany

・ロト・日本・王王・王王 シック

Outline



- 1. Smoothing, Image Derivatives, Convolutions
- 2. Edges, Corners, and Interest Points
- 3. Image Patch Descriptors
- 4. Interest Point Matching
- 5. A Simple Application: Image Stitching

Lars Schmidt-Thieme, Information Systems and Machine Learning Lab (ISMLL), University of Hildesheim, Germany

◆□▶ ◆□▶ ★∃▶ ★∃▶ ★目★ 少へつ

Outline



1. Smoothing, Image Derivatives, Convolutions

- 2. Edges, Corners, and Interest Points
- 3. Image Patch Descriptors
- 4. Interest Point Matching
- 5. A Simple Application: Image Stitching



Smoothing / Blurring / Averaging

Smoothing: Replace each pixel by the weighted average of its surrounding patch:

$$egin{aligned} & I_{ ext{smooth}}(x,y;w) := \sum_{\Delta x,\Delta y} w(-\Delta x,-\Delta y) I(x+\Delta x,y+\Delta y) \ & = \sum_{x',y'} w(x-x',y-y') I(x',y') \end{aligned}$$

- **padding** with 0 at the image boundaries.
- example: box kernel

Gaussian Kernels

► Precomputed weights: (clipped) Gaussian density values

$$egin{aligned} & ilde{w}(\Delta x,\Delta y) := egin{cases} \mathcal{N}(\sqrt{\Delta x^2 + \Delta y^2}; 0, \sigma^2), & ext{if } |\Delta x| \leq \mathcal{K}, |\Delta y| \leq \mathcal{K} \ 0, & ext{else} \end{aligned}$$
 $&w(\Delta x,\Delta y) := rac{ ilde{w}(\Delta x,\Delta y)}{\sum_{\Delta x',\Delta y'} ilde{w}(\Delta x,\Delta y)} \end{aligned}$

- ► clipped: small support, window size K.
- example ($K = 2, \sigma^2 = 1$):

$$w_{-2:2,-2:2} := \begin{pmatrix} 0.003 & 0.013 & 0.022 & 0.013 & 0.003 \\ 0.013 & 0.060 & 0.098 & 0.060 & 0.013 \\ 0.022 & 0.098 & 0.162 & 0.098 & 0.022 \\ 0.013 & 0.060 & 0.098 & 0.060 & 0.013 \\ 0.003 & 0.013 & 0.022 & 0.013 & 0.003 \end{pmatrix}$$

Note: $\mathcal{N}(x; \mu, \sigma^2) := \frac{1}{2\sigma^2} e^{-\frac{1}{2\sigma^2}}$. Lars Schmidt-Thieme, Information and Machine Learning Lab (ISMLL), University of Hildesheim, Germany



E OQA

Image Derivatives



Image Derivative: How does the intensity values change in x or y direction?

$$I_X(x,y) := I(x,y) - I(x-1,y)$$

$$I_Y(x,y) := I(x,y) - I(x,y-1)$$

or symmetric

$$I_X(x,y) := 2I(x,y) - I(x-1,y) - I(x+1,y)$$

$$I_Y(x,y) := 2I(x,y) - I(x,y-1) - (x,y-2)$$

Lars Schmidt-Thieme, Information Systems and Machine Learning Lab (ISMLL), University of Hildesheim, Germany

◆□▶ ◆□▶ ◆□▶ ◆□▶ ◆□▶ ◆□

Convolutions



convolution: an image where each pixel (x, y) represents the weighted sum around (x, y) in image I weighted with w:

$$(w * I)(x, y) := \sum_{x', y'} w(x - x', y - y')I(x', y')$$

► Examples:





Convolutions / Associativity

Convolutions are associative:

$$I * (J * K) = (I * J) * K$$

Example:

First smooth an image with Gaussian w from slide 2, then compute its x-derivative with $\begin{pmatrix} -1 & 2 & -1 \end{pmatrix}$: \rightarrow just convolve with $\begin{pmatrix} -1 & 2 & -1 \end{pmatrix} * w$

$$\left(\begin{array}{cccc} -1 & 2 & -1 \end{array}\right) * w = \left(\begin{array}{cccc} -0.007 & 0.002 & 0.017 & 0.002 & -0.007 \\ -0.033 & 0.008 & 0.077 & 0.008 & -0.033 \\ -0.054 & 0.077 & 0.128 & 0.077 & -0.054 \\ -0.033 & 0.008 & 0.077 & 0.008 & -0.033 \\ -0.007 & 0.002 & 0.017 & 0.002 & -0.007 \end{array}\right)$$

- (日) - (四) - (三) - (三) - (四) - (2



Outline



1. Smoothing, Image Derivatives, Convolutions

2. Edges, Corners, and Interest Points

- 3. Image Patch Descriptors
- 4. Interest Point Matching
- 5. A Simple Application: Image Stitching

シック 正明 《西王《西王《臣》《日》

Edges, Corners, and Interest Points



- good candidates for points that are easy to recognize and match in two images are
 - points on edges
 - corners
 - i.e., points with sudden intensity changes.
- ▶ two stage approach: given an image $I \in \mathbb{R}^{N \times M}$,
 - 1. compute an interestingness measure $i \in \mathbb{R}^{N \times M}$ for points,
 - 2. select a useful set of points $p_1, \ldots, p_K \in [N] \times [M]$
 - with high interestingness measure
 - not too close to each other.
- ▶ many names: corners, interest points, keypoints, salient points, ...

Note: $[N] := \{1, ..., N\}.$

ふつん 비居 《王》《王》《四》《日》



Gradient Magnitude (Canny Edge Detector)

Simply use the magnitude of the gradient as interestingness measure:

$$i(x,y) = \sqrt{(D_X * I)(x,y)^2 + (D_Y * I)(x,y)^2}$$

• D_X, D_Y : differentiation kernels, e.g.,

$$D_X := \begin{pmatrix} -1 & 2 & -1 \end{pmatrix}, \quad D_Y := \begin{pmatrix} -1 \\ 2 \\ -1 \end{pmatrix}$$

Lars Schmidt-Thieme, Information Systems and Machine Learning Lab (ISMLL), University of Hildesheim, Germany

◆□▶ ◆□▶ ★ヨ▶ ★ヨ▶ ★□▶ ◆○



Laplacian of Gaussian and Difference of Gaussian

Further simple interestingness measures:

• Laplacian of Gaussian (LoG):

$$i(x, y) = (((D_X * D_X + D_Y * D_Y) * G) * I)(x, y)$$

- uses second order information
- Difference of two Gaussians (DoG):

$$i(x,y) = ((G_{\sigma_1} - G_{\sigma_2}) * I)(x,y), \quad \sigma_1 \neq \sigma_2$$

- uses variations at different scales
- often interpreted as limit of Laplacian of Gaussians

$$((D_X * D_X + D_Y * D_Y) * G_{\sigma}) * I \approx \frac{\sigma}{\Delta\sigma} ((G_{\sigma + \Delta\sigma} - G_{\sigma - \Delta\sigma}) * I)$$

Harris Corner Detector

Represent a corner by its patch surrounding it, represent such a patch by a weight function

$$w:[N] imes[M] o\mathbb{R},$$

i.e.,

$$w(x,y) := \begin{cases} 1, & \text{if } |x - x_0| < 3 \text{ and } |y - y_0| < 3 \\ 0, & \text{else} \end{cases}$$

for a rectangular patch of size 5 centered around (x_0, y_0) .

A point is easy to identify, if its minimum in the autocorrelation surface is pronounced:

$$E(\Delta x, \Delta y; w) := \sum_{x,y} w(x,y) (I(x + \Delta x, y + \Delta y) - I(x,y))^2$$

Lars Schmidt-Thieme, Information Systems and Machine Learning Lab (ISMLL), University of Hildesheim, Germany



◆□▶ ◆□▶ ★ヨ▶ ★ヨ▶ ★□▶ ◆○

Shiversiter Fildesheiff









Note: left to right: flower bed, roof edge, cloud.

[Sze11, p. 187] ロト イクト イミト イミト ミニ クへで

<u>-</u>.

Harris Corner Detector

$$E(\Delta x, \Delta y; w) := \sum_{x,y} w(x,y) (I(x + \Delta x, y + \Delta y) - I(x,y))^2$$

with Hessian at minimum:

$$H(0,0;w) \approx 2 \sum_{x,y} w(x,y) \nabla I|_{(x,y)} \nabla I|_{(x,y)}^{T}, \quad \text{for } \frac{\partial^2 I}{\partial^2(x,y)} := 0$$
$$= 2w * \begin{pmatrix} (I_X)^2 & I_X I_Y \\ I_X I_Y & (I_Y)^2 \end{pmatrix},$$
$$I_X(x,y) := I(x+1,y) - I(x,y) \approx \frac{\partial I}{\partial x}(x,y)$$
$$I_Y(x,y) := I(x,y+1) - I(x,y) \approx \frac{\partial I}{\partial y}(x,y)$$

Note: $I * J(x, y) := \sum_{x', y'} I(x - x', y - y') J(x', y')$ convolution of two images.

Harris Corner Detector

use SVD to assess steepness

$$H = U \begin{pmatrix} \sigma_1 & 0 \\ 0 & \sigma_2 \end{pmatrix} U^T, \quad \sigma_1 \ge \sigma_2 \ge 0, UU^T = I$$

and define interestingness measure:

$$\begin{split} i_{\mathsf{Shi-Tomasi}}(x,y) &:= \sigma_2 \\ i_{\mathsf{Harris}}(x,y) &:= \sigma_1 \sigma_2 - \alpha (\sigma_1 + \sigma_2)^2 = \det H - \alpha \operatorname{trace}(H)^2, \qquad \alpha := 0.06 \\ i_{\mathsf{Triggs}}(x,y) &:= \sigma_2 - \alpha \sigma_1, \qquad \alpha := 0.05 \\ i_{\mathsf{Brown}}(x,y) &:= \sigma_1 \sigma_2 / (\sigma_1 + \sigma_2) = \det H / \operatorname{trace}(H) \end{split}$$

- the larger $\sigma_{1:2}$, the steeper the autocorrelation surface *E*.
- ► Harris and Brown avoid computing σ₁, σ₂ explicitly (which requires computing a square root).



Harris Corner Detector / Algorithm



1: procedure INTERESTPOINTS-HARRIS($I \in \mathbb{R}^{N \times M}$; $w \in \mathbb{R}^{-K:K \times -L:L}$, $\alpha \in \mathbb{R}$) 2: $I_X := D_X * I$ $I_{\mathbf{V}} := D_{\mathbf{V}} * I$ 3: $I_X^2 := I_X \cdot I_X$ 4: $I_Y^2 := I_Y \cdot I_Y$ 5: 6: $I_X I_Y := I_X \cdot I_Y$ $\triangleright \text{ compute } H(x, y) = \begin{pmatrix} A(x, y) & C(x, y) \\ C(x, y) & B(x, y) \end{pmatrix}$ $A := w * I_{x}^{2}$ 7: $B := w * I_{v}^{2}$ 8: $C := w * I_X I_Y$ 9: $i := A \cdot B - C \cdot C - \alpha (A + B) \cdot (A + B)$ 10: 11: return i • D_X, D_Y : differentiation kernels, e.g., $D_X := \begin{pmatrix} -1 & 2 & -1 \end{pmatrix}, D_Y := \begin{pmatrix} -1 \\ 2 \\ 1 \end{pmatrix}.$ Note: \cdot denotes the element/pixelwise product. <=> <=> <=> = = <0 < ○



Interest Points at Different Scales (SIFT Detector)

Interest points also can be identified at different scales in parallel:

$$i(p,s) := (G_{\sigma_{s+1}} * I - G_{\sigma_s} * I), \quad s \in [S]$$

where

$$\sigma_1 > \sigma_2 > \cdots > \sigma_S$$

where $S \in \mathbb{N}$ is the **number of scale levels**

- ► Often scale levels are grouped by octaves:
 - \blacktriangleright each octave is represented by a downsampling by a factor 2
 - scales within an octave are σ_s := 2^{s/S₀}σ
 (with S₀ the number of scale levels within an ocatve)

《日》《聞》《言》《言》《言》 (四)



Interest Points at Different Scales (SIFT Detector)



Lars Schmidt-Thieme, Information Systems and Machine Learning Lab (ISMLL), University of Hildesheim, Germany

[Sze11, p. 216]

Shiversiter

Non-Maximum Suppression

- Often neighbors of interest points have similar high interestingness, yielding redundant close-by interest points.
- Keep only interest points that are local maxima in their neighborhood:

$$i'(p) := egin{cases} i(p), & ext{if } i(p) > i(p') \ orall p' \in N(p) \ 0, & ext{else} \end{cases}, \quad p \in [N] imes [M]$$

with neighborhood

$$\begin{split} &N_{K}(p) := \{ p' \in [N] \times [M] \mid |p_{x} - p'_{x}| \leq K, |p_{y} - p'_{y}| \leq K, p' \neq p \} & \text{recta} \\ &N_{K}(p) := \{ p' \in [N] \times [M] \mid ||p - p'|| \leq K, p' \neq p \} & \text{circu} \end{split}$$



Non-Maximum Suppression / At Different Scale

Non-Maximum Suppression also can be extended to work on interest points at different scale:

$$egin{aligned} &\mathcal{N}_{\mathcal{K}}(p,s) := \{(p',s') \in [\mathcal{N}] imes [\mathcal{M}] imes [\mathcal{S}] \mid |p_x - p'_x| \leq \mathcal{K}, |p_y - p'_y| \leq \mathcal{K}, \ &|s - s'| \leq 1, p'
eq p \} \end{aligned}$$

◇▷▷ 비로 《표》《표》 《팀》 《□》

Computer Vision 2. Edges, Corners, and Interest Points

Universiter Hildeshein

Non-Maximum Suppression / At Different Scale

- ► SIFT refines interest points by further steps:
 - localization at sub-pixel granularity

Lars Schmidt-Thieme, Information Systems and Machine Learning Lab (ISMLL), University of Hildesheim, Germany

◆□▶ ◆□▶ ◆□▶ ◆□▶ ◆□▶ ◆□

Outline



- 1. Smoothing, Image Derivatives, Convolutions
- 2. Edges, Corners, and Interest Points
- 3. Image Patch Descriptors
- 4. Interest Point Matching
- 5. A Simple Application: Image Stitching

《日》《聞》《王》《王》《王》《曰》

Histograms



SIFT descriptors



<ロ> < 団> < 団> < 三> < 三> 三三 のへで

Computer Vision 3. Image Patch Descriptors

Histograms of Oriented Gradients (HoG)



▲□▶▲□▶▲三▶▲三▶ 三三 釣べ⊙

Lars Schmidt-Thieme, Information Systems and Machine Learning Lab (ISMLL), University of Hildesheim, Germany

Computer Vision 3. Image Patch Descriptors

Bag of Words Descriptors



Lars Schmidt-Thieme, Information Systems and Machine Learning Lab (ISMLL), University of Hildesheim, Germany

Computer Vision 3. Image Patch Descriptors

?? Dimensionality Reduction



◇▷▷ 비로 세로 / 네트 / 세 / 10 / 10

Lars Schmidt-Thieme, Information Systems and Machine Learning Lab (ISMLL), University of Hildesheim, Germany

Outline



- 1. Smoothing, Image Derivatives, Convolutions
- 2. Edges, Corners, and Interest Points
- 3. Image Patch Descriptors
- 4. Interest Point Matching
- 5. A Simple Application: Image Stitching

《日》《聞》《臣》《臣》 필급 '오오오

Computer Vision 4. Interest Point Matching

. . .



Outline



- 1. Smoothing, Image Derivatives, Convolutions
- 2. Edges, Corners, and Interest Points
- 3. Image Patch Descriptors
- 4. Interest Point Matching
- 5. A Simple Application: Image Stitching

もしゃ 本部 ネル キャット キャット

Computer Vision 5. A Simple Application: Image Stitching

. . .



シック・目目 (日本)(日本)(日本)(日本)

Lars Schmidt-Thieme, Information Systems and Machine Learning Lab (ISMLL), University of Hildesheim, Germany

Summary

. . .



シック 単則 エヨト エヨト エロト

Further Readings



- ▶ Interest points and patch descriptors: [Pri12, ch. 13], [Sze11, ch. 4].
- ▶ Image stitching: [Sze11, ch. 9].

- (日) - (四) - (三) - (三) - (四) - (0) - (0)

References



Simon JD Prince.

Computer vision: models, learning, and inference. Cambridge University Press, 2012.



Richard Szeliski.

Computer vision: algorithms and applications. Springer Science & Business Media, 2011.



うどう 正則 スポッスポッス モッ