



4.3. Local analysis of 2D image patches (cont)

4.4. Case study: Intelligent Scissors

Today's Topics

4.3. Local analysis of 2D image patches (cont)

4.4. Case study: Intelligent Scissors

Announcements

Marks for A1 are already available on line, through Blackboard

Next week is reading week.

No class, office hours only on Wednesday 11-12.

Topic 4.3:

Local analysis of 2D image patches

- Images as surfaces in 3D
- Directional derivatives
- Image Gradient
 - Painterly rendering
- Edge detection & localization
 - Gradient extrema
 - Laplacian zero-crossings
- Local geometry at image extrema
- The Image Hessian
- Eigenvectors & eigenvalues
- Corner & feature detection
 - Lowe feature detector
 - Harris/Forstner detector

Analysing Special 2D Image Patches

Goal: To mathematically characterize salient image patches



Analysing Special 2D Image Patches

Goal: To mathematically characterize salient image patches

Edges

Gradient
magnitude
and direction.

Zero-crossings
of Laplacian



Analysing Special 2D Image Patches

Goal: To mathematically characterize salient image patches



3. Local analysis of 2D image patches (cont)

4.4. Case study: Intelligent Scissors

Let's start with a demo!




Topic 4.3:

Local analysis of 2D image patches

- Images as surfaces in 3D
- Directional derivatives
- Image Gradient
 - Painterly rendering
- Edge detection & localization
 - Gradient extrema
 - Laplacian zero-crossings
- Local geometry at image extrema
- The Image Hessian
- Eigenvectors & eigenvalues
- Corner & feature detection
 - Lowe feature detector
 - Harris/Forstner detector

Analysing Special 2D Image Patches

Goal: To mathematically characterize salient image patches




Corners




??


Not these corners...



Corners



Corners




What is a corner?

How is this image patch special?

a corner patch is one where two edges intersect

Corners



will the image gradient be useful?

$$\nabla I(x, y) = \left[\frac{\partial I}{\partial x}(x, y) \quad \frac{\partial I}{\partial y}(x, y) \right]$$

(what was the intuition behind the image gradient?)

Reminder: Partial Derivative along x

$$\nabla I(x, y) = \left[\frac{\partial I}{\partial x}(x, y) \quad \frac{\partial I}{\partial y}(x, y) \right]$$

Local metric of image intensity variation in the horizontal direction

$$\left| \frac{\partial I}{\partial x}(x, y) \right|$$



Reminder: Partial Derivative along y

$$\nabla I(x,y) = \left[\frac{\partial I}{\partial x}(x,y) \quad \frac{\partial I}{\partial y}(x,y) \right]$$


Local metric of image intensity variation in the vertical direction

$$\left| \frac{\partial I}{\partial y}(x,y) \right|$$




And the action is at the gradient's...


... extrema (maxima and minima)




$$\nabla I(x, y) = \left[\frac{\partial I}{\partial x}(x, y) \quad \frac{\partial I}{\partial y}(x, y) \right]$$




$$\frac{\partial I}{\partial x}(x_0, y)$$



... extrema (maxima and minima)




$$\nabla I(x, y) = \left[\frac{\partial I}{\partial x}(x, y) \quad \frac{\partial I}{\partial y}(x, y) \right]$$



$$\frac{\partial I}{\partial y}$$

$$\frac{\partial I}{\partial y}(x_0, y)$$



but what happens at a corner?


but what happens at a corner?

and why do we care?

why we care: feature tracking




why we care: motion segmentation



a corner patch is one where two edges intersect


how can they be found using a computer?

Analysis in Neighborhood of Function Extrema




cross-section
along x


local maximum



local minimum




Analysis in Neighborhood




cross-section
along y


local maximum



local minimum





Analysis in Neighborhood of Function Extrema



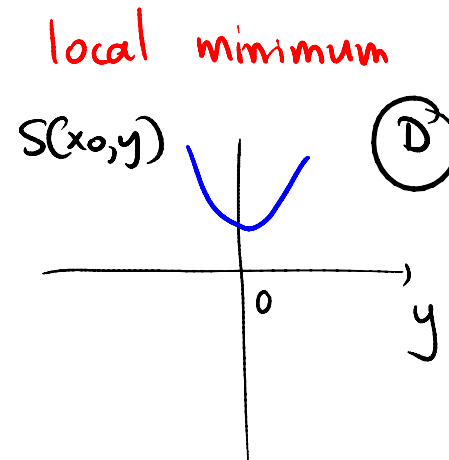

local maximum

local minimum

cross-section
along x



cross-section
along y



4 possible
combinations
of cross
sections along
 x and y
axes

Function extrema

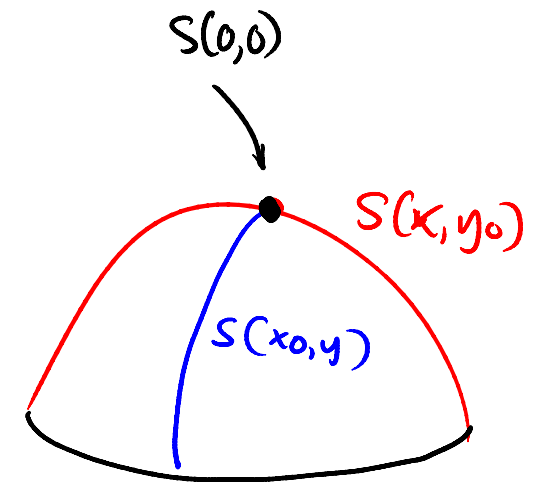

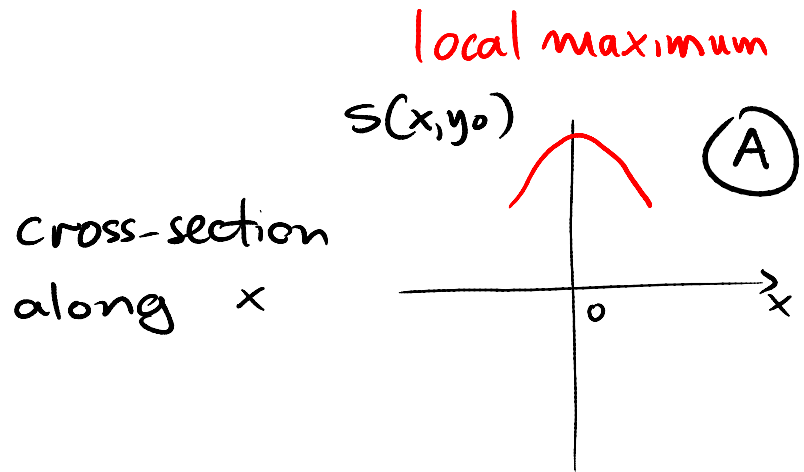
Conditions for gradient extrema:

$$\frac{\partial^2 I}{\partial x^2}(x, y) = 0$$

$$\frac{\partial^2 I}{\partial y^2}(x, y) = 0$$

Case A-C: Elliptical Points

If $\frac{\partial^2 I}{\partial x^2}(x,y) = \frac{\partial^2 I}{\partial y^2}(x,y) = 0$ it could be because:




convex surface point
(a.k.a. elliptic)


Case A-C: Elliptical Points

If $\frac{\partial^2 I}{\partial x^2}(x,y) = \frac{\partial^2 I}{\partial y^2}(x,y) = 0$ it could be because:


local maximum



local maximum




Like in:



Case A-C: Elliptical Points


If $\frac{\partial^2 I}{\partial x^2}(x,y) = \frac{\partial^2 I}{\partial y^2}(x,y) = 0$ it could be because:

local maximum




cross-section
along x

local maximum



cross-section
along y




Case A-D: Hyperbolic Points

If $\frac{\partial^2 I}{\partial x^2}(x,y) = \frac{\partial^2 I}{\partial y^2}(x,y) = 0$ it could be because:



local minimum

cross-section
along x



local maximum

cross-section
along y




saddle point
(a.k.a. hyperbolic)

Case A-D: Hyperbolic Points

If $\frac{\partial^2 I}{\partial x^2}(x,y) = \frac{\partial^2 I}{\partial y^2}(x,y) = 0$ it could be because:


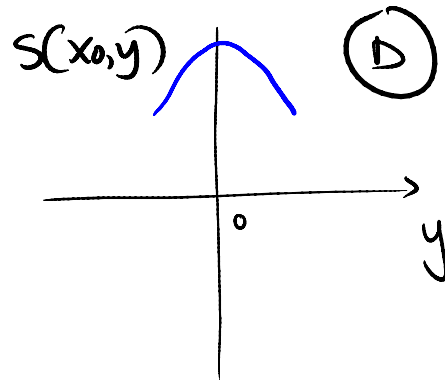
local minimum

cross-section
along x



local maximum

cross-section
along y




saddle point
(a.k.a. hyperbolic)


Cases B-C and B-D

If $\frac{\partial^2 I}{\partial x^2}(x, y) = \frac{\partial^2 I}{\partial y^2}(x, y) = 0$ it could be because:

The other two cases (B-C and B-D) produce "upside-down" versions of these two shapes



convex surface point
(a.k.a. elliptic)




saddle point
(a.k.a. hyperbolic)

Local Geometry Near Surface Extrema

What if we wanted to approximate the image close to these extreme points?

Taylor series expansion of S

$$S(x, y) = S(0, 0) +$$
$$x \frac{\partial S}{\partial x}(0, 0) + y \frac{\partial S}{\partial y}(0, 0)$$
$$+ \frac{1}{2} \left[x^2 \frac{\partial^2 S}{\partial x^2}(0, 0) + 2xy \frac{\partial^2 S}{\partial x \partial y}(0, 0) + y^2 \frac{\partial^2 S}{\partial y^2}(0, 0) \right]$$

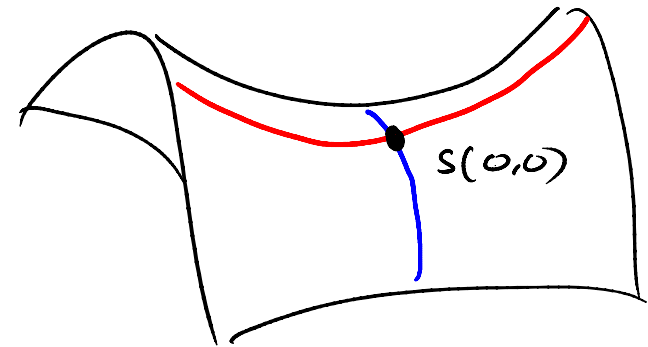


Local Geometry Near Surface Extrema

What if we wanted to approximate the image close to these extreme points?

Taylor series expansion of S

$$S(x,y) = S(0,0) + \frac{1}{2} \left[x^2 \frac{\partial^2 S}{\partial x^2}(0,0) + 2xy \frac{\partial^2 S}{\partial x \partial y}(0,0) + y^2 \frac{\partial^2 S}{\partial y^2}(0,0) \right]$$



Note that local shape is determined by the 2nd derivative only!

Topic 4.2:

Local analysis of 2D image patches

- Images as surfaces in 3D
- Directional derivatives
- Image Gradient
 - Painterly rendering
- Edge detection & localization
 - Gradient extrema
 - Laplacian zero-crossings
- Local geometry at image extrema
- **The Image Hessian**
- Eigenvectors & eigenvalues
- Corner & feature detection
 - Lowe feature detector
 - Harris/Forstner detector

The Hessian Matrix

If $\frac{\partial^2 I}{\partial x^2}(x,y) = \frac{\partial^2 I}{\partial y^2}(x,y) = 0$ at a point (x,y) , then the

Taylor Series expansion of the function at that point is:


$$S(x,y) = S(0,0) + \frac{1}{2} \left[x^2 \frac{\partial^2 S}{\partial x^2}(0,0) + 2xy \frac{\partial^2 S}{\partial x \partial y}(0,0) + y^2 \frac{\partial^2 S}{\partial y^2}(0,0) \right]$$

Hessian
matrix

$$= S(0,0) + \frac{1}{2} \begin{bmatrix} x & y \end{bmatrix} \begin{bmatrix} \frac{\partial^2 S}{\partial x^2} & \frac{\partial^2 S}{\partial x \partial y} \\ \frac{\partial^2 S}{\partial x \partial y} & \frac{\partial^2 S}{\partial y^2} \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$

The Hessian Matrix: Intuition

The Hessian matrix H determines how $S(x,y)$ changes from a unit-length displacement d , in a given direction




$$S(x,y) = S(0,0) + \frac{1}{2} \begin{bmatrix} x & y \end{bmatrix} \begin{bmatrix} \frac{\partial^2 S}{\partial x^2} & \frac{\partial^2 S}{\partial x \partial y} \\ \frac{\partial^2 S}{\partial x \partial y} & \frac{\partial^2 S}{\partial y^2} \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$

$$S(x,y) = S(0,0) + d^T H d \text{ with } d = \begin{bmatrix} x \\ y \end{bmatrix}$$

The Hessian Matrix: Intuition

In other words, the Hessian knows about the local shape of S



$$\begin{bmatrix} \frac{\partial^2 S}{\partial x^2} & \frac{\partial^2 S}{\partial x \partial y} \\ \frac{\partial^2 S}{\partial x \partial y} & \frac{\partial^2 S}{\partial y^2} \end{bmatrix}$$

Local Geometry from the Hessian Matrix

$$S(x, y) = S(0, 0) + \frac{1}{2} \begin{bmatrix} x & y \end{bmatrix} \begin{bmatrix} \frac{\partial^2 S}{\partial x^2} & \frac{\partial^2 S}{\partial x \partial y} \\ \frac{\partial^2 S}{\partial x \partial y} & \frac{\partial^2 S}{\partial y^2} \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$

The **Hessian** (H) defines 2 orthogonal unit vectors v_1, v_2 such that:

$$\cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_1 \iff S(x, y) - S(0, 0) = \lambda_1 = \text{MAXIMAL}$$


$$\cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_2 \iff S(x, y) - S(0, 0) = \lambda_2 = \text{MINIMAL}$$

$$\text{angle} \left(\begin{bmatrix} x \\ y \end{bmatrix}, v_1 \right) = \theta \implies S(x, y) - S(0, 0) = \lambda_1 \cos^2 \theta + \lambda_2 \sin^2 \theta$$

Local Geometry from the Hessian Matrix

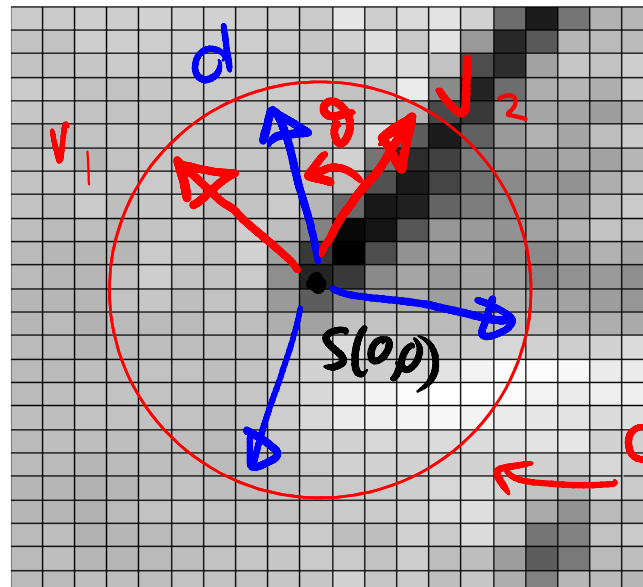
$$\cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_1 \Leftrightarrow S(x,y) - S(0,0) = \lambda_1 = \text{MAXIMAL}$$

$$\cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_2 \Leftrightarrow S(x,y) - S(0,0) = \lambda_2 = \text{MINIMAL}$$



Local Geometry from the Hessian Matrix

$$\text{angle} \left(\begin{bmatrix} x \\ y \end{bmatrix}, v_1 \right) = \theta \Rightarrow S(x, y) - S(0, 0) = \lambda_1 \cos^2 \theta + \lambda_2 \sin^2 \theta$$



circle of
unit radius

Local Geometry from the Hessian Matrix

$$\text{angle} \left(\begin{bmatrix} x \\ y \end{bmatrix}, v_1 \right) = \vartheta \implies S(x, y) - S(0, 0) = \lambda_1 \cos^2 \vartheta + \lambda_2 \sin^2 \vartheta$$

Proof

If $\text{angle}(d, v_1) = \vartheta$, we have $d = \cos \vartheta \cdot v_1 + \sin \vartheta \cdot v_2$

$$\left(\cos \vartheta \cdot v_1^T + \sin \vartheta \cdot v_2^T \right) H \left(\cos \vartheta \cdot v_1 + \sin \vartheta \cdot v_2 \right) =$$

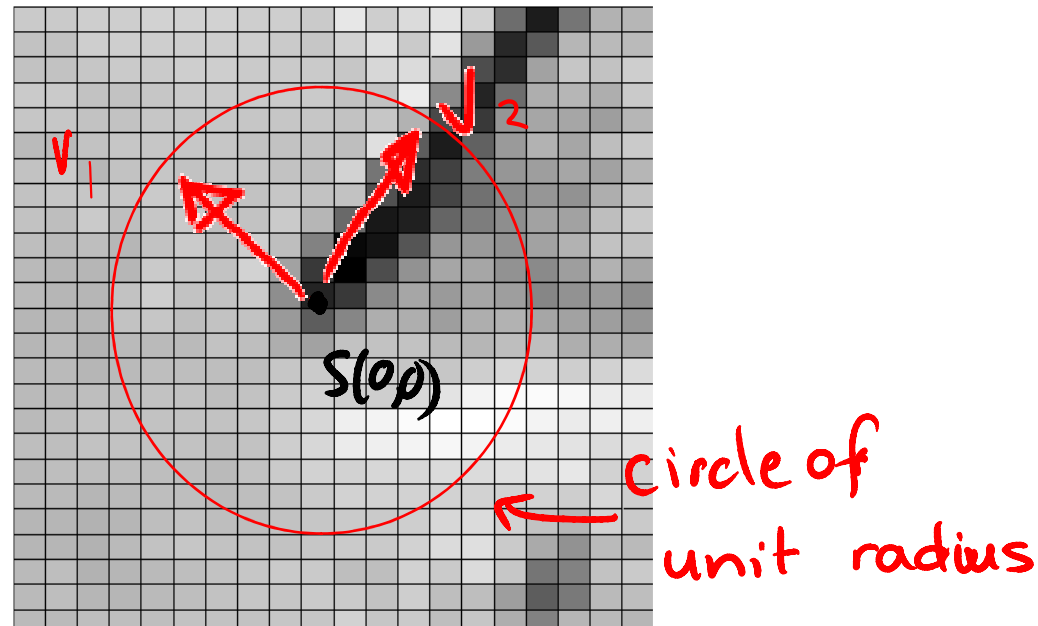
$$\cos^2 \vartheta \cdot v_1^T \overset{\lambda_1 v_1}{(H v_1)} + \cancel{\sin \vartheta \cdot \cos \vartheta \cdot v_2^T \overset{\lambda_1 v_1}{(H v_1)}} + \cancel{\cos \vartheta \cdot \sin \vartheta \cdot v_1^T \overset{\lambda_2 v_2}{(H v_2)}} + \sin^2 \vartheta \cdot v_2^T \overset{\lambda_2 v_2}{(H v_2)} =$$

$$\cos^2 \vartheta \cdot \lambda_1 + \sin^2 \vartheta \cdot \lambda_2$$

Local Geometry from the Hessian Matrix

$$\cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_1 \Leftrightarrow S(x,y) - S(0,0) = \lambda_1 = \text{MAXIMAL}$$

$$\cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_2 \Leftrightarrow S(x,y) - S(0,0) = \lambda_2 = \text{MINIMAL}$$



Local Geometry from the Hessian Matrix

$$\cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_1 \Leftrightarrow S(x,y) - S(0,0) = \lambda_1 = \text{MAXIMAL}$$

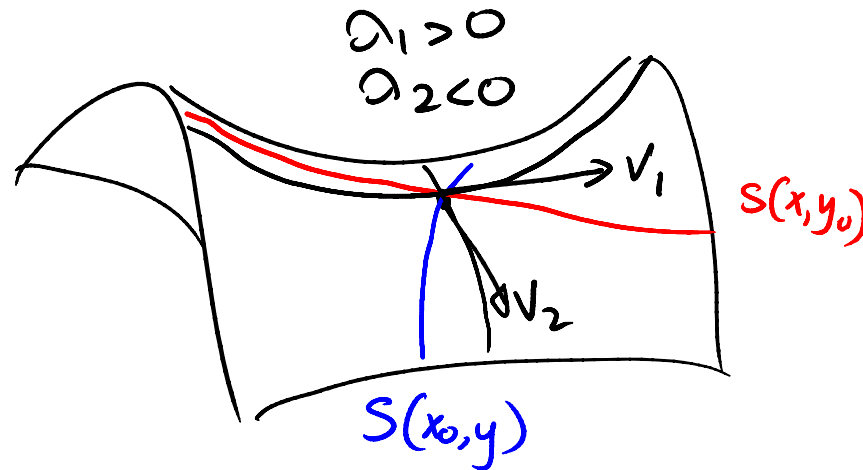
$$\cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_2 \Leftrightarrow S(x,y) - S(0,0) = \lambda_2 = \text{MINIMAL}$$

- v_1, v_2 are unit Eigenvectors of H
- λ_1, λ_2 are Eigenvalues of H

Local Geometry from the Hessian Matrix

$$\cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_1 \Leftrightarrow S(x,y) - S(0,0) = \lambda_1 = \text{MAXIMAL}$$

$$\cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_2 \Leftrightarrow S(x,y) - S(0,0) = \lambda_2 = \text{MINIMAL}$$



v_1 = direction of greatest ascent
 v_2 = direction of smallest ascent
(or greatest descent)

Local Geometry from the Hessian Matrix


$$\cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_1 \Leftrightarrow S(x,y) - S(0,0) = \lambda_1 = \text{MAXIMAL}$$

$$\cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_2 \Leftrightarrow S(x,y) - S(0,0) = \lambda_2 = \text{MINIMAL}$$

$$\lambda_1 < 0$$
$$\lambda_2 < 0$$

v_1 = direction of steepest ascent (or less steep descent)

v_2 = direction of steepest descent (or least steep ascent)



Local Geometry from the Hessian Matrix


$$\begin{bmatrix} x \\ y \end{bmatrix} = v_1 \Leftrightarrow S(x,y) - S(0,0) = \lambda_1 = \text{MAXIMAL}$$

$$\begin{bmatrix} x \\ y \end{bmatrix} = v_2 \Leftrightarrow S(x,y) - S(0,0) = \lambda_2 = \text{MINIMAL}$$

$$\lambda_1 < 0$$
$$\lambda_2 < 0$$


v_1 = direction of steepest ascent (or less steep descent)

v_2 = direction of steepest descent (or least steep ascent)



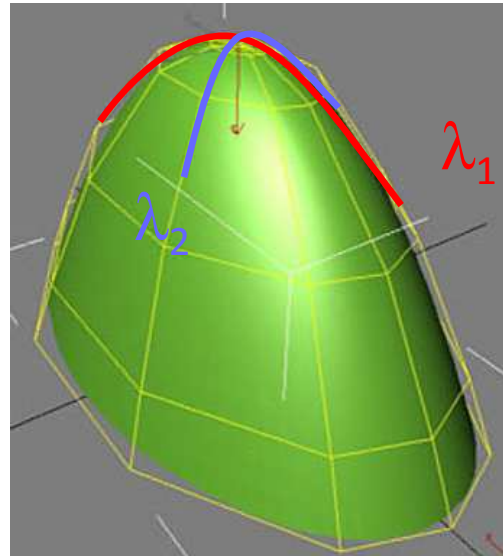
Principal Directions & Curvatures

- v_1 and v_2 are called the principal directions at the surface point $S(0,0)$



Principal Directions & Curvatures

- λ_1 and λ_2 are called the principal curvatures at $S(0,0)$



Topic 4.3:

Local analysis of 2D image patches

- Images as surfaces in 3D
- Directional derivatives
- Image Gradient
 - Painterly rendering
- Edge detection & localization
 - Gradient extrema
 - Laplacian zero-crossings
- Local geometry at image extrema
- The Image Hessian
 - Eigenvectors & eigenvalues
- Corner & feature detection
 - Lowe feature detector
 - Harris/Forstner detector

Background: Eigenvectors & Eigenvalues

Definition:

A non-zero vector \mathbf{v} is an eigenvector of a matrix H if

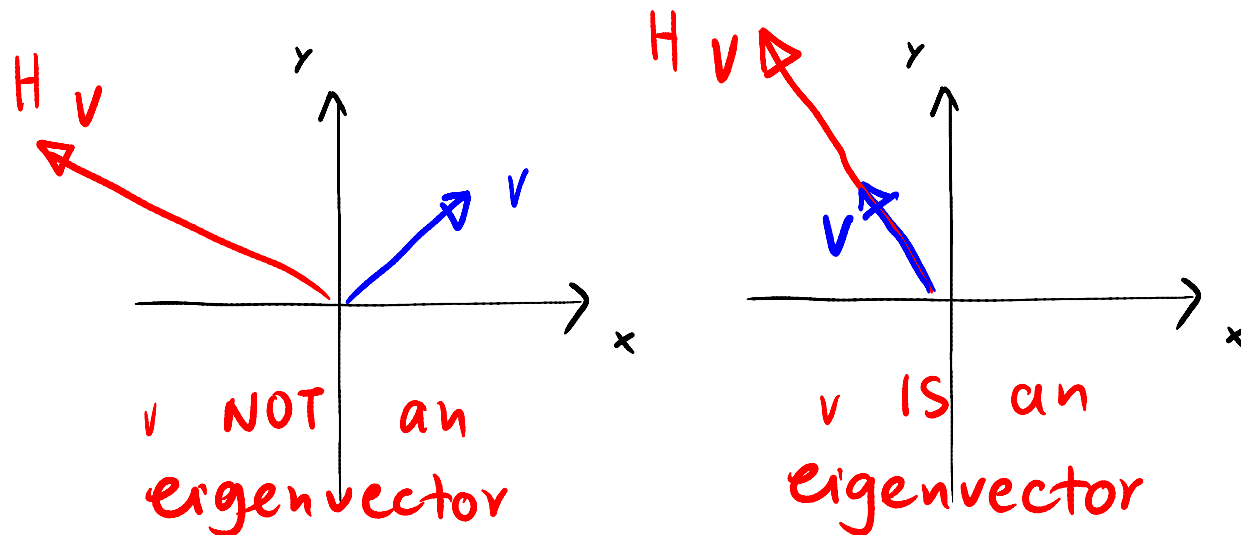
$$H\mathbf{v} = \lambda\mathbf{v}.$$

The scalar λ is the eigenvalue associated to \mathbf{v} .

Background: Eigenvectors & Eigenvalues

A non-zero vector \mathbf{v} is an eigenvector of a matrix H if $H\mathbf{v} = \lambda\mathbf{v}$.
The scalar λ is the eigenvalue associated to \mathbf{v} .


A matrix H transforms a vector \mathbf{v} to the vector $H\mathbf{v}$, so if:



Background: Eigenvectors & Eigenvalues

A non-zero vector \mathbf{v} is an eigenvector of a matrix H if $H\mathbf{v} = \lambda\mathbf{v}$.
The scalar λ is the eigenvalue associated to \mathbf{v} .

A matrix H transforms a vector \mathbf{v} to the vector $H\mathbf{v}$, so if:



This also means that if \mathbf{v} is an eigenvector and $k \neq 0$ is a constant, then $k\mathbf{v}$ is also an eigenvector.

We will think of eigenvectors as unit-eigenvectors $||\mathbf{v}||=1$.

Eigenvectors of Symmetric Matrices

The hessian $H = \begin{bmatrix} \frac{\partial^2 S}{\partial x^2} & \frac{\partial^2 S}{\partial x \partial y} \\ \frac{\partial^2 S}{\partial x \partial y} & \frac{\partial^2 S}{\partial y^2} \end{bmatrix}$ is a **symmetric** matrix, which

means that $H = H^T$ (where H^T denotes the transpose of H).

We said that if v is an eigenvector then $Hv = \lambda v$, but when H is symmetric, then: $v^T H = \lambda v^T$.

Proof: Let $H = \begin{bmatrix} a & c \\ c & b \end{bmatrix}$, $v = \begin{bmatrix} x \\ y \end{bmatrix}$

$$\begin{aligned} H \begin{bmatrix} x \\ y \end{bmatrix} &= \begin{bmatrix} a & c \\ c & b \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} ax+cy \\ cx+by \end{bmatrix} \\ \begin{bmatrix} x & y \end{bmatrix} \begin{bmatrix} a & c \\ c & b \end{bmatrix} &= \begin{bmatrix} ax+cy & cx+by \end{bmatrix} \end{aligned}$$

Eigenvectors of Symmetric Matrices

When the matrix H is symmetric ($H = H^T$ like the Hessian), its associated eigenvectors are orthogonal: $v_1^T v_2 = 0$.

Proof :


$$\lambda_1 v_1^T v_2 = (\lambda_1 v_1^T) v_2 = v_1^T H v_2 = v_1^T v_2 \cdot \lambda_2$$

$$(\lambda_1 - \lambda_2) v_1^T v_2 = 0$$

Since $\lambda_1 \neq \lambda_2$, $v_1^T v_2$ must be 0

BTW, what other matrices are symmetric

The covariance matrix!



$$\Sigma_{ij} = \text{cov}(X_i, X_j) = \text{E}[(X_i - \mu_i)(X_j - \mu_j)]$$

Eigenvalues & the Trace of a Matrix

The trace of a matrix (denoted as $\text{tr}(H)$) is the sum of the diagonal elements.

The sum of the eigenvalues of a matrix H is equal to its trace (regardless of H being symmetric).

For a 2×2 matrix $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$ $a + d = \lambda_1 + \lambda_2$

This is important for computational efficiency

Eigenvalues & the Trace of a Matrix

The product of the eigenvalues of a matrix H is equal to its determinant (also regardless of H being symmetric).

For a 2×2 matrix $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$ $ad - bc = \lambda_1 \cdot \lambda_2$


This is very important for computational efficiency

Topic 4.3:

Local analysis of 2D image patches

- Images as surfaces in 3D
- Directional derivatives
- Image Gradient
 - Painterly rendering
- Edge detection & localization
 - Gradient extrema
 - Laplacian zero-crossings
- Local geometry at image extrema
- The Image Hessian
- Eigenvectors & eigenvalues
- Corner & feature detection
 - Lowe feature detector
 - Harris/Forstner detector

Analysis in Neighborhood of Function Extrema





What type of extrema do we have?



local maximum

local minimum


cross-section
along x



cross-section
along y



Analysis in Neighborhood of Function Extrema





We can determine the type of extremum from $\text{tr}(H)$ and $\det(H)$!!

local maximum

local minimum



cross-section
along x




cross-section
along y

local maximum

local minimum




Analysis in Neighborhood of Function Extrema




We can determine the type of extremum from $\text{tr}(H)$ and $\det(H)$!!

local maximum




cross-section along x


local minimum



locally constant




local maximum




cross-section along y

local minimum



locally constant



The Hessian Matrix: Elliptical Points

H defines 2 orthogonal unit vectors v_1 and v_2 such that:

$$\begin{aligned} \cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_1 &\iff S(x,y) - S(0,0) = \lambda_1 = \text{MAXIMAL} \\ \cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_2 &\iff S(x,y) - S(0,0) = \lambda_2 = \text{MINIMAL} \end{aligned}$$

$$\det(H) = \lambda_1 \cdot \lambda_2 > 0 \quad ?$$

The Hessian Matrix: Elliptical Points

H defines 2 orthogonal unit vectors v_1 and v_2 such that:

$$\begin{aligned} \cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_1 & \iff S(x,y) - S(0,0) = \lambda_1 = \text{MAXIMAL} \\ \cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_2 & \iff S(x,y) - S(0,0) = \lambda_2 = \text{MINIMAL} \end{aligned}$$


$$\det(H) = \lambda_1 \cdot \lambda_2 > 0 \implies \text{sign}(\lambda_1) = \text{sign}(\lambda_2)$$

The Hessian Matrix: Elliptical Points

H defines 2 orthogonal unit vectors v_1 and v_2 such that:


$$\begin{aligned} \cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_1 &\iff S(x,y) - S(0,0) = \lambda_1 = \text{MAXIMAL} \\ \cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_2 &\iff S(x,y) - S(0,0) = \lambda_2 = \text{MINIMAL} \end{aligned}$$

$$\det(H) = \lambda_1 \cdot \lambda_2 > 0 \implies \text{sign}(\lambda_1) = \text{sign}(\lambda_2) \implies \text{Elliptical}$$




The Hessian Matrix: Elliptical Points

$$\det(H) = \lambda_1 \cdot \lambda_2 > 0 \Rightarrow \text{sign}(\lambda_1) = \text{sign}(\lambda_2) \Rightarrow \text{Elliptical}$$




The Hessian Matrix: Elliptical Points

$$\det(H) = \lambda_1 \cdot \lambda_2 > 0 \implies \text{sign}(\lambda_1) = \text{sign}(\lambda_2) \implies \text{Elliptical}$$




$$\text{tr}(H) = \begin{cases} > 0 & \text{case B-D} \\ < 0 & \text{case A-C} \end{cases}$$


local minimum




local maximum



local minimum



local maximum



$$\text{tr}(H) > 0$$

$$\text{tr}(H) < 0$$

The Hessian Matrix: Elliptical Points

H defines 2 orthogonal unit vectors v_1 and v_2 such that:

$$\begin{aligned} \cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_1 &\iff S(x,y) - S(0,0) = \lambda_1 = \text{MAXIMAL} \\ \cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_2 &\iff S(x,y) - S(0,0) = \lambda_2 = \text{MINIMAL} \end{aligned}$$




$$\det(H) = \lambda_1 \cdot \lambda_2 < 0$$

The Hessian Matrix: Elliptical Points

H defines 2 orthogonal unit vectors v_1 and v_2 such that:

$$\begin{aligned} \cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_1 &\iff S(x,y) - S(0,0) = \lambda_1 = \text{MAXIMAL} \\ \cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_2 &\iff S(x,y) - S(0,0) = \lambda_2 = \text{MINIMAL} \end{aligned}$$


$\det(H) = \lambda_1 \cdot \lambda_2 < 0 \implies \text{sign}(\lambda_1) \neq \text{sign}(\lambda_2) \implies \text{saddle points}$




The Hessian Matrix: Elliptical Points

$\det(H) = \lambda_1 \cdot \lambda_2 < 0 \Rightarrow \text{sign}(\lambda_1) \neq \text{sign}(\lambda_2) \Rightarrow$ saddle points


$S(x, y)$




local minimum




local maximum



local maximum



local minimum



The Hessian Matrix: Elliptical Points

H defines 2 orthogonal unit vectors v_1 and v_2 such that:

$$\begin{aligned} \cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_1 &\iff S(x,y) - S(0,0) = \lambda_1 = \text{MAXIMAL} \\ \cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_2 &\iff S(x,y) - S(0,0) = \lambda_2 = \text{MINIMAL} \end{aligned}$$


$$\det(H) = \lambda_1 \cdot \lambda_2 = 0 \quad \text{but} \quad \text{tr}(H) \neq 0$$

The Hessian Matrix: Elliptical Points

H defines 2 orthogonal unit vectors v_1 and v_2 such that:

$$\begin{aligned} \cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_1 &\iff S(x,y) - S(0,0) = \lambda_1 = \text{MAXIMAL} \\ \cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_2 &\iff S(x,y) - S(0,0) = \lambda_2 = \text{MINIMAL} \end{aligned}$$

$$\det(H) = \lambda_1 \cdot \lambda_2 = 0 \quad \text{but} \quad \text{tr}(H) \neq 0$$




E-C, E-D, A-F, B-F

The Hessian Matrix: Elliptical Points

$$\det(H) = \lambda_1 \cdot \lambda_2 = 0 \quad \text{but} \quad \text{tr}(H) \neq 0$$

$S(x, y)$



$$\text{tr}(H) = \begin{cases} > 0 & \text{case E-D, F-B} \\ < 0 & \text{case E-C, F-A} \end{cases}$$

Topic 4.3:

Local analysis of 2D image patches

- Images as surfaces in 3D
- Directional derivatives
- Image Gradient
 - Painterly rendering
- Edge detection & localization
 - Gradient extrema
 - Laplacian zero-crossings
- Local geometry at image extrema
- The Image Hessian
- Corner & feature detection
 - Lowe feature detector
 - Harris/Forstner detector

The (single-scale) Lowe Feature Detector

Goal of Lowe feature detector

Find pixels that are very different from their surroundings.

$I(x,y)$



$$S(x,y) = \nabla^2 I(x,y) = \frac{\partial^2 I}{\partial x^2} + \frac{\partial^2 I}{\partial y^2}$$




The (single-scale) Lowe Feature Detector

Idea: Apply the Hessian eigenvalue analysis to the Laplacian of an image

$I(x,y)$




$$S(x,y) = \nabla^2 I(x,y) = \frac{\partial^2 I}{\partial x^2} + \frac{\partial^2 I}{\partial y^2}$$



The (single-scale) Lowe Feature Detector

Steps:

- ① Identify all pixels that are extrema of $S(x,y)$



$I(x,y)$




$$S(x,y) = \nabla^2 I(x,y) = \frac{\partial^2 I}{\partial x^2} + \frac{\partial^2 I}{\partial y^2}$$



The (single-scale) Lowe Feature Detector

Steps:

- ① Identify all pixels that are extrema of $S(x,y)$
- ② Eliminate all extrema with $|S(x,y)| < \text{threshold}$ ($= 0.03$)



$I(x,y)$




$$S(x,y) = \nabla^2 I(x,y) = \frac{\partial^2 I}{\partial x^2} + \frac{\partial^2 I}{\partial y^2}$$



The (single-scale) Lowe Feature Detector

Steps:

- ① Identify all pixels that are extrema of $S(x,y)$
- ② Eliminate all extrema with $|S(x,y)| < \text{threshold}$ ($= 0.03$)
- ③ Eliminate all extrema whose local shape is near-cylindrical (i.e. "edge-like")



$I(x,y)$



$$S(x,y) = \nabla^2 I(x,y) = \frac{\partial^2 I}{\partial x^2} + \frac{\partial^2 I}{\partial y^2}$$



The (single-scale) Lowe Feature Detector

- ③ Eliminate all extrema whose local shape is near-cylindrical (i.e. "edge-like")


How?

The Hessian Matrix: Elliptical Points

H defines 2 orthogonal unit vectors v_1 and v_2 such that:

$$\begin{aligned} \cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_1 &\iff S(x,y) - S(0,0) = \lambda_1 = \text{MAXIMAL} \\ \cdot \begin{bmatrix} x \\ y \end{bmatrix} = v_2 &\iff S(x,y) - S(0,0) = \lambda_2 = \text{MINIMAL} \end{aligned}$$

$$\det(H) = \lambda_1 \cdot \lambda_2 = 0 \quad \text{but} \quad \text{tr}(H) \neq 0$$



E-C, E-D, A-F, B-F

The (single-scale) Lowe Feature Detector

③ Eliminate all extrema whose local shape is near-cylindrical (i.e. "edge-like")

3a: Compute $\frac{\partial^2 S}{\partial x^2}$, $\frac{\partial^2 S}{\partial y^2}$, $\frac{\partial^2 S}{\partial x \partial y}$

3b: Compute $\text{Det}(H)$ at each pixel

$$\text{det}(H) = \lambda_1 \cdot \lambda_2 = 0$$

in practice:

check $|\lambda_1| \gg |\lambda_2|$ or $|\lambda_2| \gg |\lambda_1|$

The (single-scale) Lowe Feature Detector

③ Eliminate all extrema whose local shape is near-cylindrical (i.e. "edge-like")

3a: Compute $\frac{\partial^2 S}{\partial x^2}$, $\frac{\partial^2 S}{\partial y^2}$, $\frac{\partial^2 S}{\partial x \partial y}$

3b: Compute $\text{Det}(H)$ at each pixel

$$\det(H) = \lambda_1 \cdot \lambda_2 = 0$$

in practice:

check $|\lambda_1| \gg |\lambda_2|$ or $|\lambda_2| \gg |\lambda_1|$

The (single-scale) Lowe Feature Detector

- ③ Eliminate all extrema whose local shape is near-cylindrical (i.e. "edge-like")

But what we really want to know is if λ_1 is much bigger compared terms to λ_2 , as in "r" times bigger, so if we do:

$$\lambda_1 = r \lambda_2$$

We can use the following equation to test for the ratio:

$$\frac{\text{Tr}(H)^2}{\text{Det}(H)} = \frac{(\lambda_2 + r\lambda_2)^2}{r \cdot \lambda_2^2} = \frac{(r+1)^2}{r^2}$$

The (single-scale) Lowe Feature Detector

③ Eliminate all extrema whose local shape is near-cylindrical (i.e. "edge-like")

3a: Compute $\frac{\partial^2 S}{\partial x^2}$, $\frac{\partial^2 S}{\partial y^2}$, $\frac{\partial^2 S}{\partial x \partial y}$

3b: Compute $\text{Det}(H)$ at each pixel

3c: Compute $\text{Tr}(H)$

3d: Keep only pixels with

$$\frac{\text{Tr}(H)^2}{\text{Det}(H)} < \left(\frac{r+1}{r}\right)^2 \quad \text{usually: } r=10$$

The (single-scale) Lowe Feature Detector

Result: Top 500 un-eliminated pixels
(ranked according to $|\nabla^2 I|$)



Topic 4.3:

Local analysis of 2D image patches

- Images as surfaces in 3D
- Directional derivatives
- Image Gradient
 - Painterly rendering
- Edge detection & localization
 - Gradient extrema
 - Laplacian zero-crossings
- Local geometry at image extrema
- The Image Hessian
- Corner & feature detection
 - Lowe feature detector
 - Harris/Forstner detector

The Harris/Forstner Corner Detector

Goal of Harris / Forstner corner detector:

Find pixels that are "very different"
from their neighborhood

Idea:

Equivalent
eigenvalue
analysis on
a slightly
different
 $S(x,y)$
functional



The Harris/Forstner Corner Detector

The Harris/Forstner Corner Detection looks for (dis)-similarity between the center pixel and its neighborhood, literally

$$S(x, y) = \sum_u \sum_v w(u, v) (I(u + x, v + y) - I(u, v))^2$$

The center pixels are more heavily weighted than the rest.

Now, because $I(u + x, v + y)$ can be approximated using a Taylor expansion

$$I(u + x, v + y) \approx I(u, v) + I_x(u, v)x + I_y(u, v)y$$

The Harris/Forstner Corner Detector

The dissimilarity metric

$$S(x, y) = \sum_u \sum_v w(u, v) (I(u + x, v + y) - I(u, v))^2$$

can then be written as the approximation:

$$S(x, y) \approx \sum_u \sum_v w(u, v) (I_x(u, v)x + I_y(u, v)y)^2$$

or equivalently:

$$S(x, y) \approx \begin{pmatrix} x & y \end{pmatrix} A \begin{pmatrix} x \\ y \end{pmatrix}$$

with:
$$A = \sum_u \sum_v w(u, v) \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix}$$

The Harris/Forstner Corner Detector

The Harris/Forstner Corner detector looks for patches that are “very dissimilar” from their neighbors!



The Harris/Forstner Corner Detector

The Harris/Forstner Corner detector looks for patches that are “very dissimilar” from their neighbors!



Topic 4.4:


Application: Intelligent Scissors

- Assigning cost to “links” between pixels
- Contour tracing as a shortest-path problem

Image Scissoring: Motivation

By scissoring portions of one or more images & pasting them together we can create new, composite images

source images



composite image



Image Scissoring: Requirements

Interactive operation

“User is always right”

Scissoring system must allow user to select arbitrary regions

Scissoring operations must be performed efficiently

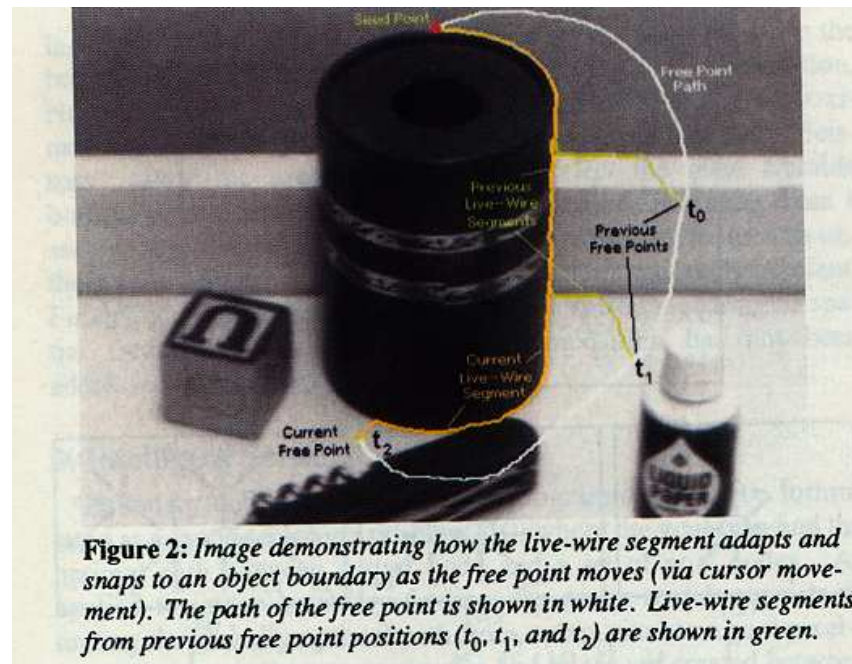
Scissoring interface must be simple & easy to use



Image Scissoring

In contrast the manual approach requires the user to manually delineate every single image pixel that defines the region boundary

What is proposed here is: Intelligent Scissors approach (“livewire”) developed by Eric Mortensen & presented at SIGGRAPH’95

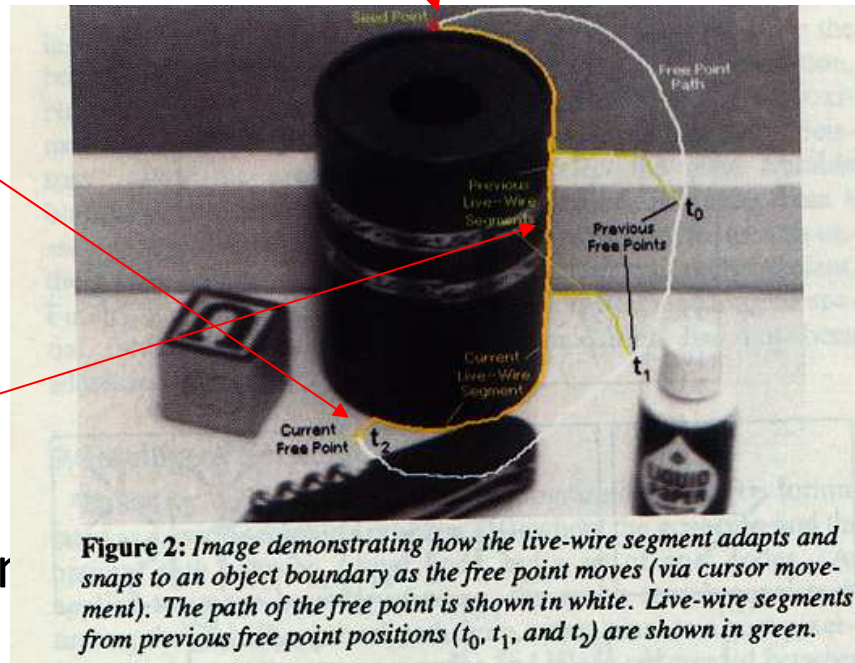


Intelligent Scissors: Operation

User loads image & specifies a “seed” point on boundary that must be outlined

User then positions mouse close to object boundary

System automatically creates a “live-wire” that connects seed & current mouse position & follows boundary as much as possible

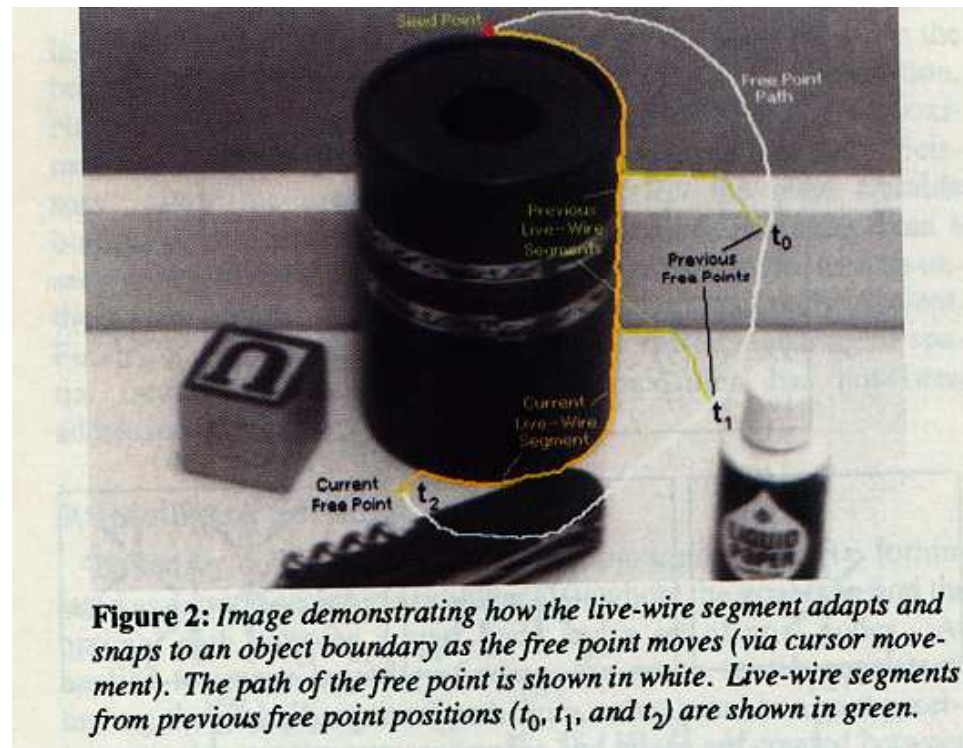


Live wire updated whenever mouse moves

Intelligent Scissors: Basic Idea

Approach answers one basic question:

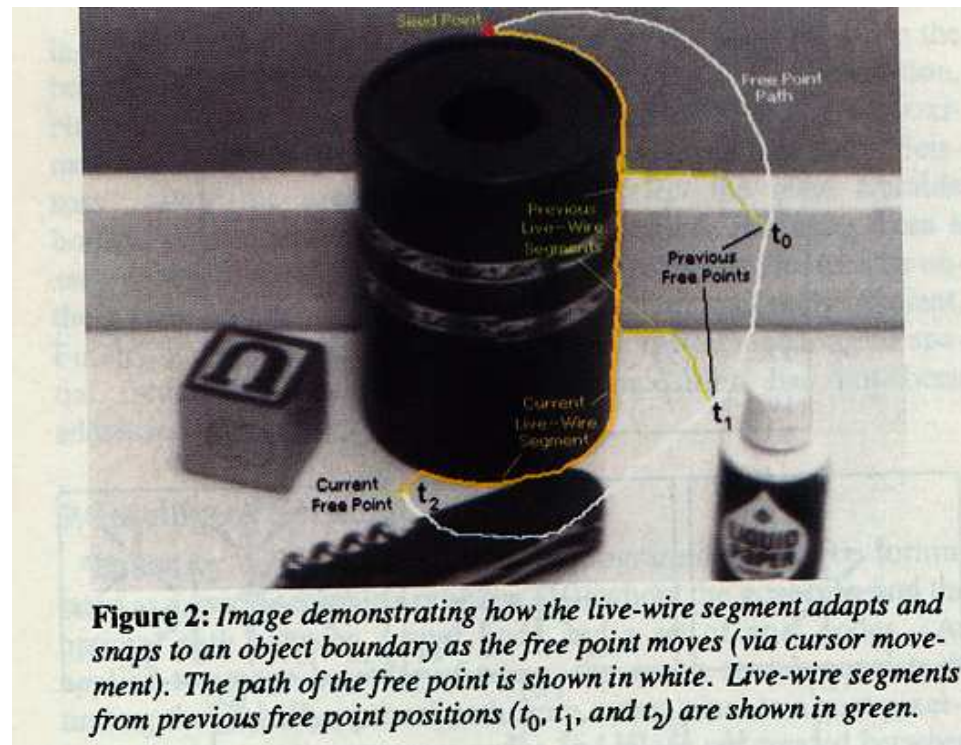
- How should we define a path from seed to mouse that follows an object boundary as closely as possible?




- Answer: Define a path that is as close as possible to image edges

Intelligent Scissors: Basic Idea

Approach taken in intelligent scissors attempts to exploit user interaction while avoiding the need to detect edges corresponding to object boundaries very accurately





Intelligent Scissors: Basic Idea



link



- Every pair of neighboring pixels is called a link and is assigned an “edgeness” weight
- Link weights defined so that pixel links along an edge have very low weights
- To connect seed & mouse positions, choose the **path** that minimizes the total weight of links along the path



Intelligent Scissors: Basic Idea

Two questions must be answered to fully specify the algorithm:


- How do we assign a weight to a link?
- How do we find the lowest-cost path between any two image pixels?



Intelligent Scissors: Weight Assignment

Path: a sequence of adjacent pixels in the image

Link: a pair of adjacent pixels along the path



Given a link defined by pixels p & q , its weight is defined to be

$$l(p, q) = 0.43 f_Z(q) + 0.43 f_D(p, q) + 0.14 f_G(q)$$

$f_Z(q) = 0$ if the Laplacian has a zero-crossing at q

$f_Z(q) = 1$ otherwise

term that penalizes links not consistent with the gradient direction at p and q


$$f_G(q) = 1 - \frac{|\nabla I|}{\max(|\nabla I|)}$$

largest value over the image

Intelligent Scissors: Weight Assignment

Path: a sequence of adjacent pixels in the image

Link: a pair of adjacent pixels along the path



Given a link defined by pixels p & q , its weight is defined to be

$$l(p, q) = 0.43 f_Z(q) + 0.43 f_D(p, q) + 0.14 f_G(q)$$

$$f_G(q) = 1 - \frac{|\nabla I|}{\max(|\nabla I|)}$$

largest value
over the image

Weight Assignment: Gradient Term


Recall:

- can detect edges where 1st derivative is high



$$I_x \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}$$

$$I_y \begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ 1 & 2 & 1 \end{bmatrix}$$



gradient magnitude

$$f_G(q) = 1 - \frac{|\nabla I|}{\max(|\nabla I|)}$$

—————> largest value
over the image


$$|\nabla I| = |(I_x, I_y)| = \sqrt{I_x^2 + I_y^2}$$

- high gradients produce low costs
- scaled by largest gradient in image, to lie in [0,1]

Weight Assignment: Laplacian Term

Path: a sequence of adjacent pixels in the image

Link: a pair of adjacent pixels along the path



Given a link defined by pixels p & q, its weight is defined to be

$$l(p, q) = 0.43 f_Z(q) + 0.43 f_D(p, q) + 0.14 f_G(q)$$


$f_Z(q) = 0$ if the Laplacian has a
zero-crossing at q

$f_Z(q) = 1$ otherwise

Intelligent Scissors: Weight Assignment

Recall:



- can detect edges where 2nd derivative is zero (inflection points)
- find zero-crossings instead (sign change with 8-neighbors)



$f_z(q)$

$$L = I_{xx} + I_{yy}$$

$$\begin{bmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{bmatrix}$$



Laplacian $L = I_{xx} + I_{yy}$

$f_z(q) = 0$ if the Laplacian has a zero-crossing at q


$f_z(q) = 1$ otherwise

- zero-crossings produce low costs
- many zero-crossings are due to noise

Intelligent Scissors: Weight Assignment

Path: a sequence of adjacent pixels in the image

Link: a pair of adjacent pixels along the path



Given a link defined by pixels p & q , its weight is defined to be

$$l(p, q) = 0.43 f_Z(q) + 0.43 f_D(p, q) + 0.14 f_G(q)$$




term that penalizes links not
consistent with the gradient
direction at p and q

Weight Assignment: Direction Term

gradient direction term $f_D(p, q)$

- penalizes paths that do not follow edges in the image
- penalizes sharp changes in path direction (creases)
- normalized to lie in $[0, 1]$



$$f_D(p, q) = \left(\text{angle}(v, \nabla p^\perp) + \text{angle}(v, \nabla q^\perp) \right) / \pi$$

$$\text{angle}(u, w) = \arccos \left(\frac{|u \cdot w|}{|u||w|} \right) \in [0, \pi/2]$$

Intelligent Scissors: Path Optimizer


Path optimization formulated as a graph search algorithm that computes the minimum-cost path from seed to all other image pixels (use Dijkstra's algorithm)

Weights


11	13	12	9	5	8	3	1	2	4	10
14	11	7	4	2	5	8	4	6	3	8
11	6	3	5	7	9	12	11	10	7	4
7	4	6	11	13	18	17	14	8	5	2
6	2	7	10	15	15	21	19	8	3	5
8	3	4	7	9	13	14	15	9	5	6
11	5	2	8	3	4	5	7	2	5	9
12	4	2	1	5	6	3	2	4	8	12
10	9	7	5	9	8	5	3	7	8	15

seed


Step 1




Step 2




Step 3



Step n



Intelligent Scissors: Results



Intelligent Scissors: Results




(a)



Intelligent Scissors: Results

By scissoring portions of one or more images & pasting them together we can create new, composite images

source images



composite image

