What Can We Learn from Discrete Images about the Continuous World?

Ullrich Köthe

Multidimensional Image Processing
Interdisciplinary Center for Scientific Computing (IWR)
University of Heidelberg
April 2008

Multidimensional Image Processing IWR, Univ. of Heidelberg

The Continuous and the Discrete

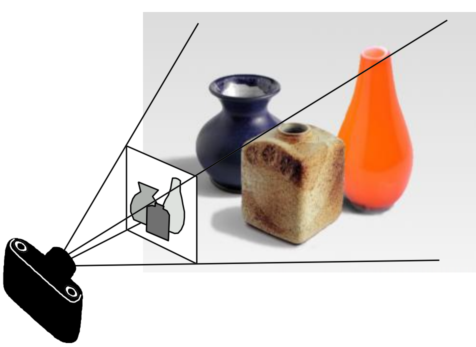

The world is continuous
(as far as image analysis is concerned)
infinite amount of data

but the computer is discrete
finite amount of data

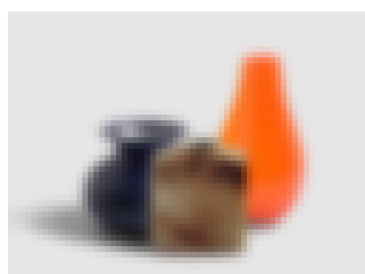

How can we ever be sure that we didn't lose the information of interest?

Take-Home Messages

- It is a good idea to analyze both domains conjointly
  - formally establish correspondences
  - build models in whatever domain is more convenient, knowing that results remain valid in the other domain
- Some useful results already exist
  - signal-theoretic and geometric sampling theorems
  - joint error analysis

Joint work with Peer Stelldinger and Hans Meine

Many thanks to H.S. Stiehl, B. Neumann, V. Kaynig, N. Boëtius, G. Kedenburg, F. Hamprecht, DGCI

Traditional Approaches To Spatial Discretization

- Heuristic Approach
  - develop an algorithm
  - show experimentally that it succeeds sufficiently often
- Physics Approach:
  - theory is derived in the continuous domain
  - discretization is left to another discipline (numerical analysis) – not part of the core theory
  - correctness proofs in form of asymptotic convergence theorems – not applicable to fixed size images
- Approach of Digital Geometry
  - model and prove everything in the discrete domain
  - but: the relation of the original image to the real world is not considered – we may have already lost the relevant information

Example: Discretization of a Differential Equation

Anisotropic Diffusion (structure-enhancing smoothing)

$$\frac{\partial u}{\partial t} = \nabla \cdot (D \nabla u) \quad \text{where } D \text{ is a diffusion tensor}$$

original (no noise)

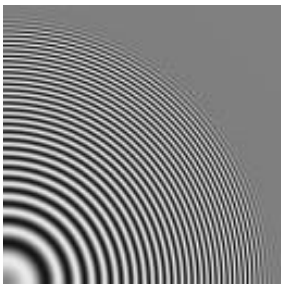

standard discretization

based on standard finite differences

$$[1 \quad -1]$$

Weickert & Scharr: "A scheme for coherence-enhancing diffusion filtering with optimized rotation invariance", J. Visual Communication and Image Representation, 13(1/2):103-118, 2002

Example: Discretization of a Differential Equation

Anisotropic Diffusion (structure-enhancing smoothing)

$$\frac{\partial u}{\partial t} = \nabla \cdot (D \nabla u) \quad \text{where } D \text{ is a diffusion tensor}$$

original (no noise)

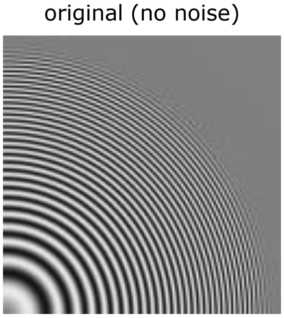

standard discretization

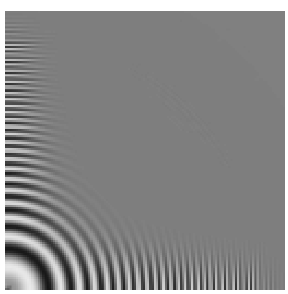

improved discretization

based on finite differences optimized for rotational invariance

$$\frac{1}{32}\begin{bmatrix} 3 & 0 & -3 \\ 10 & 0 & -10 \\ 3 & 0 & -3 \end{bmatrix}$$

Weickert & Scharr: "A scheme for coherence-enhancing diffusion filtering with optimized rotation invariance", J. Visual Communication and Image Representation, 13(1/2):103-118, 2002

Example: Discretization of a Differential Equation

Anisotropic Diffusion (structure-enhancing smoothing)

$$\frac{\partial u}{\partial t} = \nabla \cdot (D \nabla u) \quad \text{where } D \text{ is a diffusion tensor}$$

original (no noise)

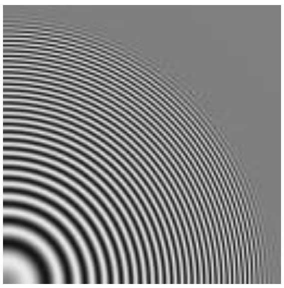

standard discretization

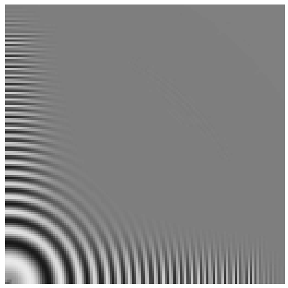

improved discretization

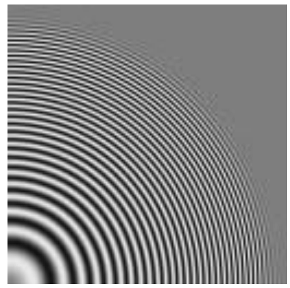

Weickert & Scharr: "A scheme for coherence-enhancing diffusion filtering with optimized rotation invariance", J. Visual Communication and Image Representation, 13(1/2):103-118, 2002

Example: Discretization of a Differential Equation

Anisotropic Diffusion (structure-enhancing smoothing)

$$\frac{\partial u}{\partial t} = \nabla \cdot (D \nabla u) \quad \text{where } D \text{ is a diffusion tensor}$$

original (noisy)

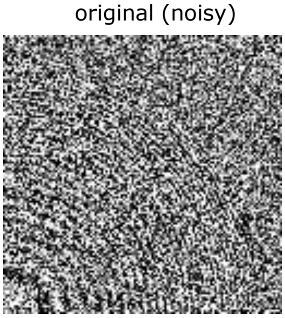

standard discretization

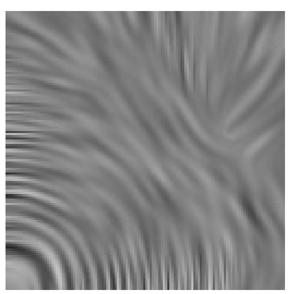

improved discretization

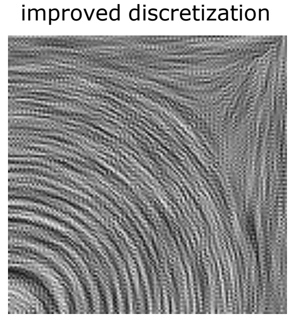

Weickert & Scharr: "A scheme for coherence-enhancing diffusion filtering with optimized rotation invariance", J. Visual Communication and Image Representation, 13(1/2):103-118, 2002

Example: Sampling Artifacts in the Original Image

- Moiré effect

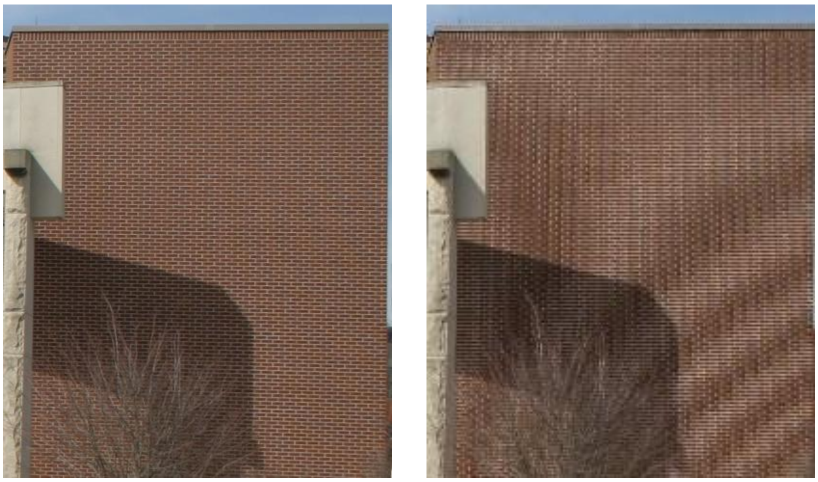

9

Multidimensional Image Processing IWR, Univ. of Heidelberg

Example: Sampling Artifacts in the Original Image

- Moiré effect

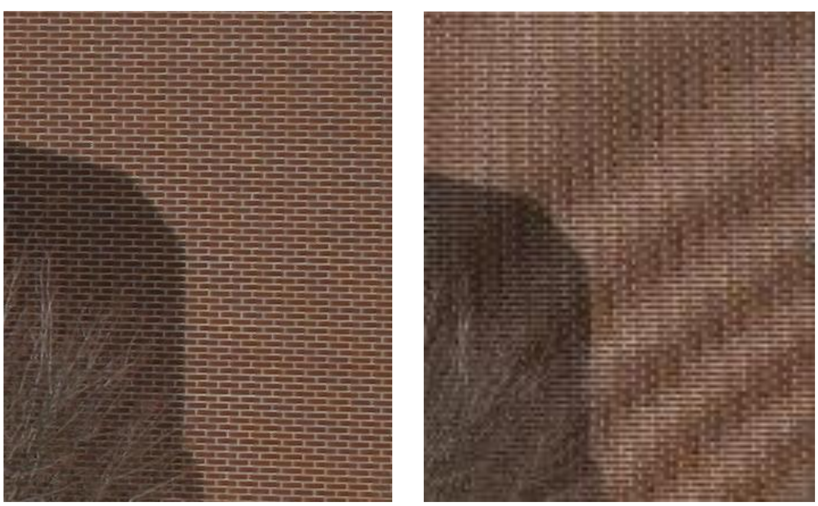

10

Example: Sampling Artifacts in the Original Image

- color Moiré effect

Multidimensional Image Processing IWR, Univ. of Heidelberg

Example: Sampling Artifacts in the Original Image

- color Moiré effect

Example: Sampling Artifacts in the Original Image

- color Moiré effect

staircasing artifacts

13

Example: Sampling Artifacts in the Original Image

- color Moiré effect

staircasing artifacts

14

Example: Wrong Character

- Non-trivial shape changes when sampling is too coarse

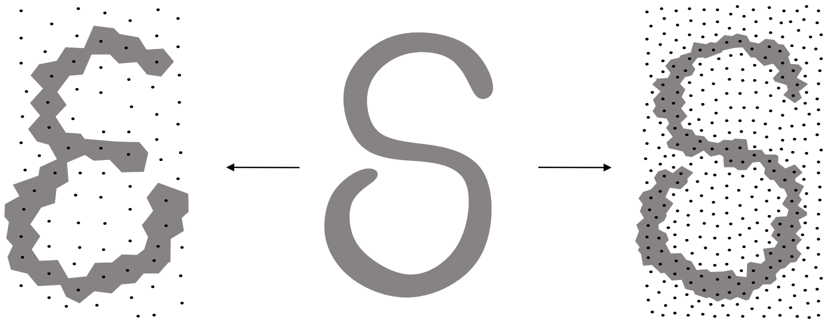

Stelldinger & Köthe: "Towards a general sampling theory for shape preservation", Image and Vision Computing, 23(2):237-248, 2005

15

Multidimensional Image Processing IWR, Univ. of Heidelberg

Cameras as Linear Systems

Distinguish ideal geometric image (infinite resolution) and real digital image (finite)

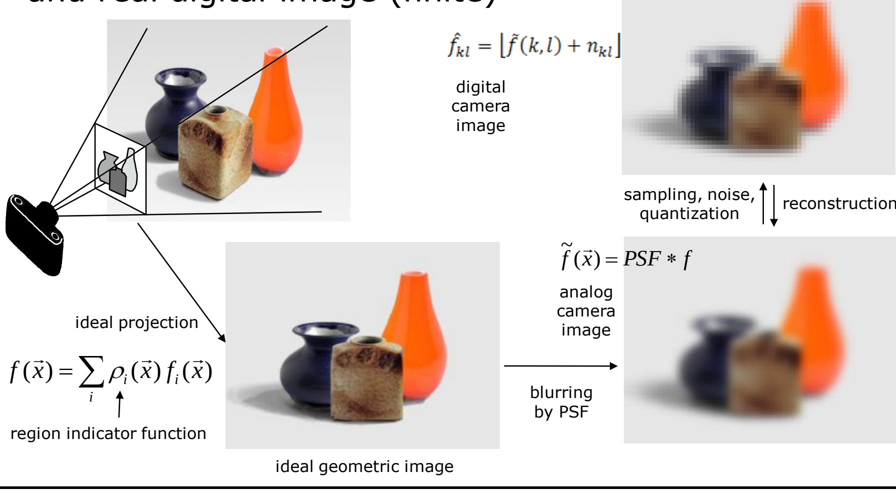

16

Cameras as Linear Systems

Distinguish ideal geometric image (infinite resolution) and real digital image (finite)

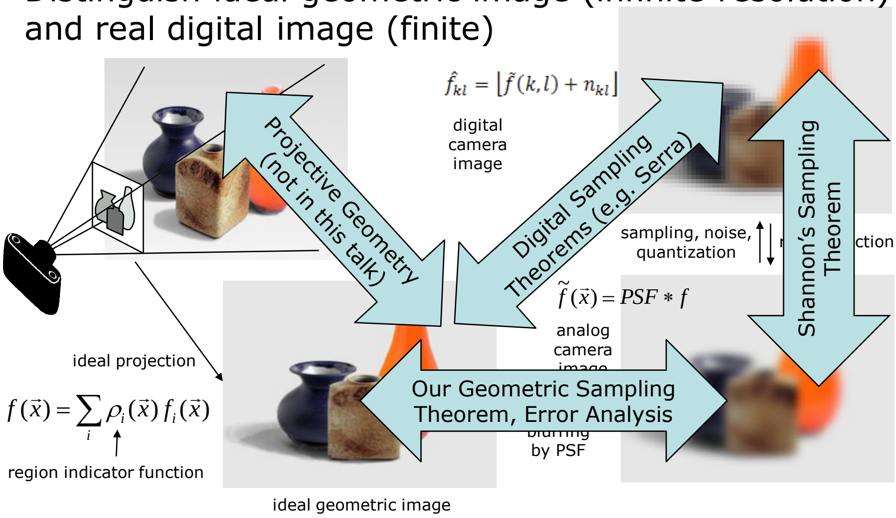

Shannon's Sampling Theorem (1)

- Analog signal can be exactly reconstructed from samples when spectrum is band limited

$$F(u_1,u_2)=0 \text{ for } |u_1|,|u_2| \geq u_{\text{Nyquist}}$$

and sampling is sufficiently dense $u_{\text{sampling}} \geq 2u_{\text{Nyquist}}$

- Holds approximately for real cameras
  - effective Nyquist frequency of MTF: $\approx$100 lp/mm
  - sampling frequency: $\approx$136 lp/mm
  - residual aliasing can be modeled as noise (Moiré effect is rare)

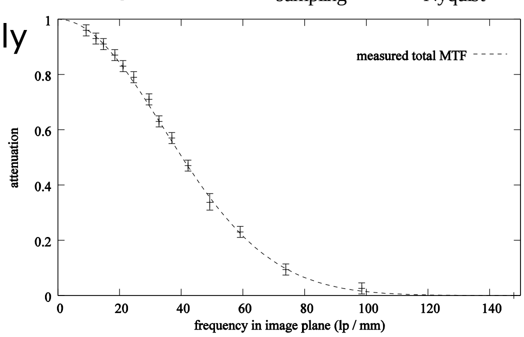

Shannon's Sampling Theorem (2)

- Reconstruction of the analog camera image possible
  - point spread function (PSF) of real cameras enforces band limit (typical: Gaussian PSFs with σ = 0.45 ... 0.8)
  - small errors due to noise, quantization, finite size, and slight under-sampling
  - reconstruction by sinc interpolation or spline interpolation
    - high reconstruction quality for orders above 3 or 5

  ⇨ can work in either continuous or discrete domain

36 rotations by 10 degrees:

high errors for nearest neighbor and linear interpolation

very low error for 5th order spline

Unser et al.: "B-Spline Signal Processing", Trans. Signal Processing, 41(2), pp. 821-848, 1993

Multidimensional Image Processing IWR, Univ. of Heidelberg

Shannon's Sampling Theorem (3)

- Surprising fact: many imaging textbooks introduce sampling theorem, but few draw consequences
- E.g: non-linear operators *increase* bandwidth
  - image gradient (linear filter) doesn't change bandwidth:

$$\nabla f = \begin{pmatrix} f_x \\ f_y \end{pmatrix} = \begin{pmatrix} g_x * f \\ g_y * f \end{pmatrix}$$

  - but: image gradient magnitude

$$\|\nabla f\|^2 = f_x^2 + f_y^2 \;\circ\!\!-\!\!\bullet\; F_x * F_x + F_y * F_y$$

  doubles the bandwidth, because multiplication corresponds to convolution in Fourier domain

  ⇨ gradient sampling rate must double $u_{\text{sampling, grad}} \geq 2u_{\text{sampling, original}}$

Köthe: "Edge and Junction Detection with an Improved Structure Tensor", DAGM '03, pp. 25-32, 2003

Shannon's Sampling Theorem (4)

- Visible aliasing in gradient image

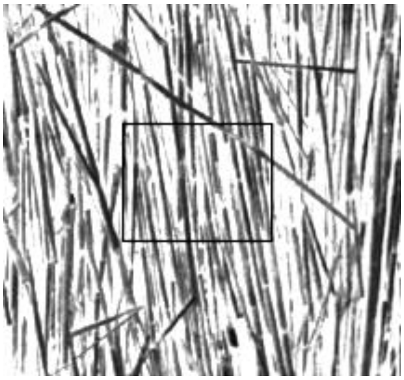

gradient at original sampling rate

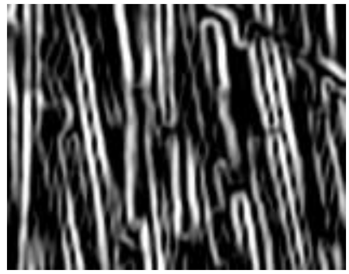

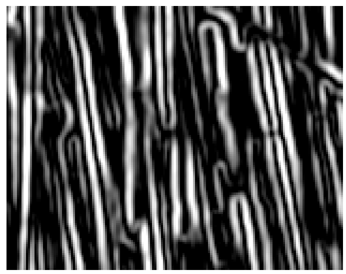

gradient at doubled sampling rate

Shannon's Sampling Theorem (5)

- Gradient based edge detector:

  Errors due to insufficient resolution clearly visible

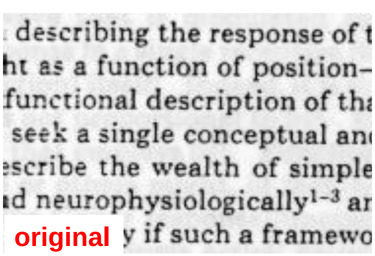
describing the response of t
ht as a function of position–
functional description of tha
seek a single conceptual an
escribe the wealth of simple
d neurophysiologically[1-3] ar
y if such a framewo

original

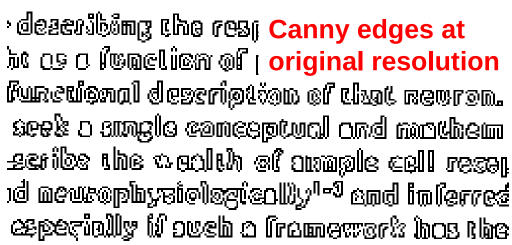
describing the res
ht as a function of
functional description of that neuron.
seek a single conceptual and mathem
scribe the wealth of simple cell recep
d neurophysiologically[1-3] and inferred
especially if such a framework has the

Canny edges at original resolution

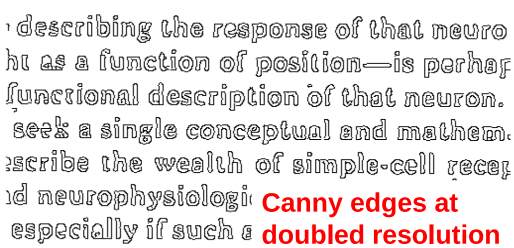
describing the response of that neuro
ht as a function of position—is perhap
functional description of that neuron.
seek a single conceptual and mathem
scribe the wealth of simple-cell recep
d neurophysiologi
especially if such a

Canny edges at doubled resolution

Limitations of Shannon's Sampling Theorem

- Sampling theorem gives no guarantees about image geometry, e.g.
  - number of edges (zero-crossings) per unit area
  - number and arrangement of critical points
  - only limited on average
  - example: construct band-limited function (minimum wavelength = 2) with 9 critical points (4 minima, 1 maximum, 4 saddles) in single facet [0,1]x[0,1]

$$s_1(x,k) = \mathrm{si}(x+k) + \frac{4}{5}\mathrm{si}(x-k-1)$$

$$s_2(x,k) = s_1(x,k) / s_1(0.5,k)$$

$$s_3(x) = s_2(x,1) - s_2(x,2) - 0.05$$

$$f(\vec{x}) = s_3(x)s_3(y)$$

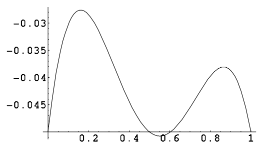

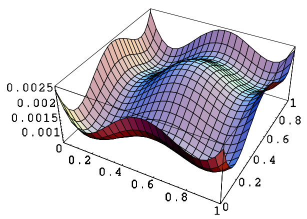

Stelldinger & Köthe: unpublished manuscript

Digital Sampling Theorems

Short cut between geometric image and digital image

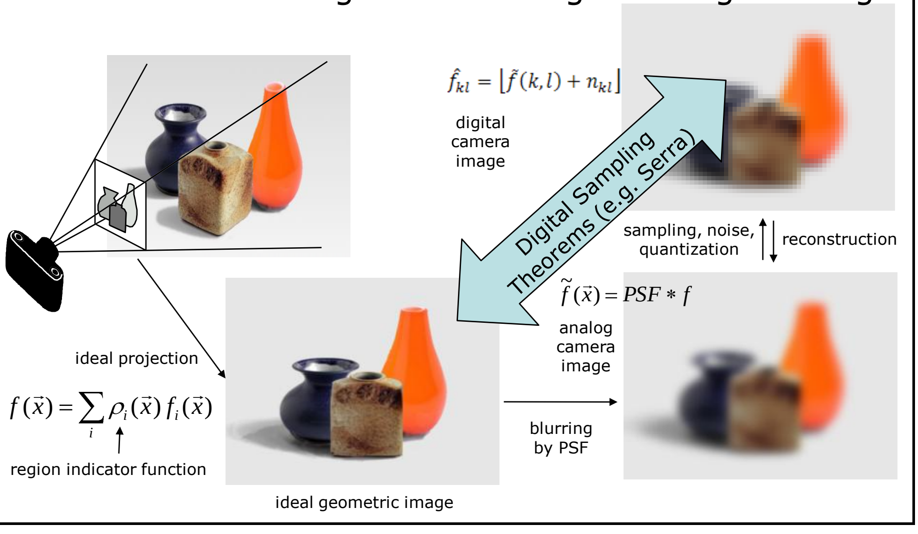

Digital Sampling Theorems

- Consider ideal geometric image as given:
  - assumptions about region shapes (e.g. *r*-regular)
  - model of digitization process (e.g. Gauss digitization)
  - image reconstruction method from digitization (e.g. nearest-neighbor reconstruction)
  - criteria for successful reconstruction (e.g. homeomorphism between original and reconstruction)
- Proof that criteria are fulfilled under these conditions

Digital Sampling Theorems

- Serra-Pavlidis Theorem (1982)
  - *r*-regular shapes: morphologically open and closed w.r.t. ball of radius *r* (equiv.: osculating *r*-balls)

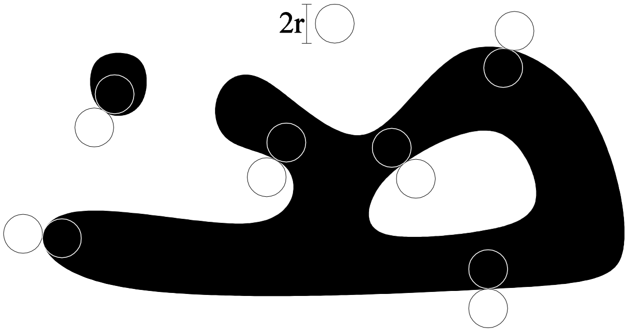

Digital Sampling Theorems

- Serra-Pavlidis Theorem (1982)
  - *r*-regular shapes: morphologically open and closed w.r.t. ball of radius *r* (equiv.: osculating *r*-balls)
  - subset digitization with *s*-grid

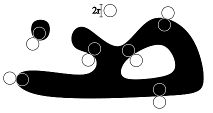

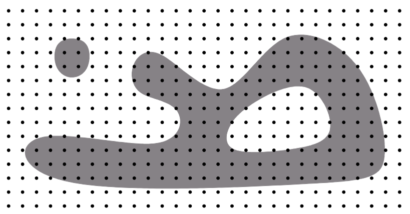

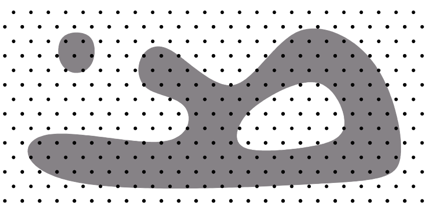

Multidimensional Image Processing IWR, Univ. of Heidelberg

Digital Sampling Theorems

- Serra-Pavlidis Theorem (1982)
  - *r*-regular shapes: morphologically open and closed w.r.t. ball of radius *r* (equiv.: osculating *r*-balls)
  - subset digitization with *s*-grid
  - nearest-neighbor reconstruction

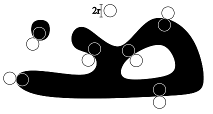

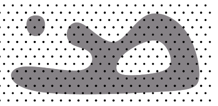

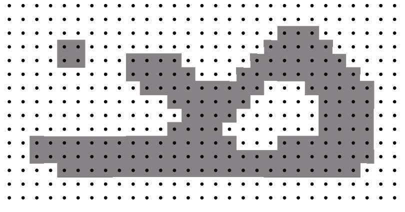

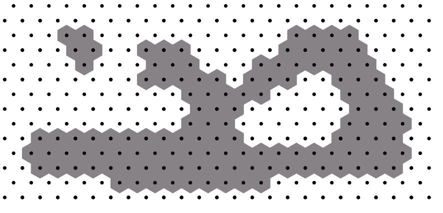

Digital Sampling Theorems

- Serra-Pavlidis Theorem (1982)
  - *r*-regular shapes: morphologically open and closed w.r.t. ball of radius *r* (equiv.: osculating *r*-balls)
  - subset digitization with *s*-grid
  - nearest-neighbor reconstruction
  - criterion: homeomorphism between original and reconstruction
- Proof: shape preserved if $s<r$

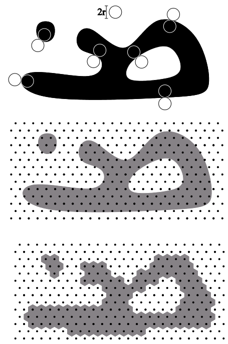

Serra: "Image Analysis and Mathematical Morphology", Academic Press, 1982
Pavlidis: "Algorithms for Graphics and Image Processing", Computer Science Press, 1982

Digital Sampling Theorems

- Various extensions, e.g.
  - Latecki et al.:
    - μ-digitization (arbitrary threshold μ)
  - Stelldinger, Köthe:
    - extension to irregular grids
    - stronger guarantee (*r*-homeomorphism)
    - shapes may be blurred with disc-shaped PSF
- But too restrictive
  - *r*-regularity forbids corners and junctions, implies binary shapes
  - homeomorphism cannot be guaranteed in non-binary images

Latecki, Conrad, Gross: "Preserving Topology by a Digitization Process", JMIV 8:131-159, 1998
Stelldinger & Köthe: "Towards a general sampling theory for shape preservation", IMAVIS, 2005

Geometric Sampling Theorem

Describe difference between geometric image and *reconstructed analog image*

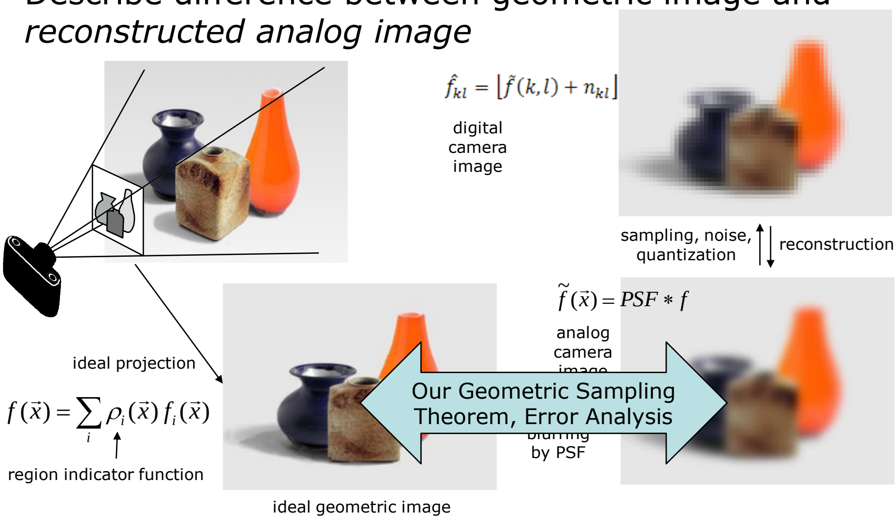

Multidimensional Image Processing IWR, Univ. of Heidelberg

Homeomorphy vs. Homotopy

- Requirement of homeomorphism between original shape and reconstruction is too strong

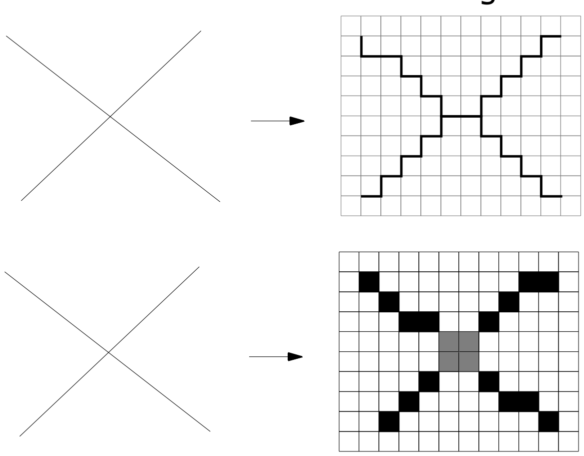

Homeomorphy vs. Homotopy

- Requirement of homeomorphism between original shape and reconstruction is too strong

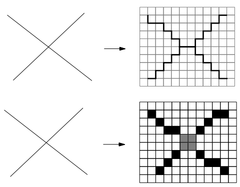

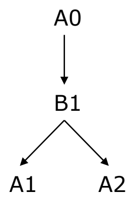

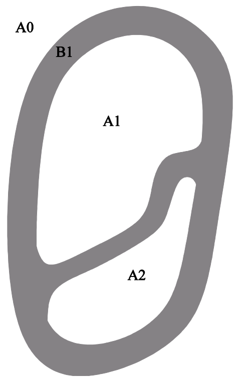

- Only require homotopy equivalence
  – isomorphic homotopy trees
  – alternating region and boundary levels

33

Meine, Köthe, Stelldinger: "A Topological Sampling Theorem for Robust Boundary Reconstruction and Image Segmentation", Discrete Applied Mathematics, to appear, 2008

r-Stable Shapes

- Most scenes are not binary
- Most shapes have corners, occlusion causes junctions
  – new condition: partition must be *r*-stable
  = homotopy equivalence after *r*-dilation of boundary (no waists!)
  – regions must contain 2*r* circle

shape is r_1-stable, but not r_2-stable

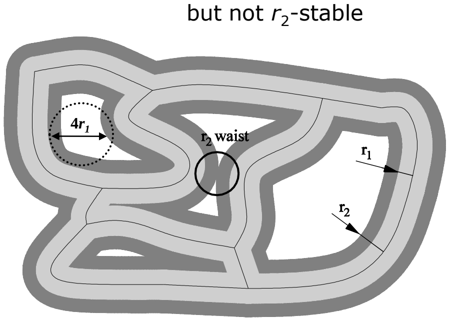

34

Meine, Köthe, Stelldinger: "A Topological Sampling Theorem for Robust Boundary Reconstruction and Image Segmentation", Discrete Applied Mathematics, to appear, 2008

Non-Perfect Boundary Sampling

- Real samplings and edge detectors are not perfect
  - systematic distortions: smoothing by PSF and filters
  - stochastic distortions: noise, round-off error

  - do not sample regions, but boundaries („edge points")
  - detected edge points characterized by two kinds of errors:
    - *p* (maximum distance from true contour to edge point)
    - *q* (maximum distance from edge point to true contour)

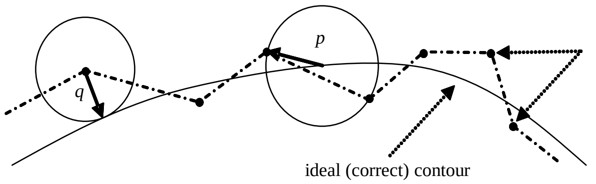

35

Meine, Köthe, Stelldinger: "A Topological Sampling Theorem for Robust Boundary Reconstruction and Image Segmentation", Discrete Applied Mathematics, to appear, 2008

Edge Point Linking

(α,β)-reconstruction: region growing on Delaunay triangulation of edge points

(generalization of α-shapes from laser range scanning)

36

Edge Point Linking

(α,β)-reconstruction: region growing on Delaunay triangulation of edge points

(generalization of α-shapes from laser range scanning)

1. detect edge points

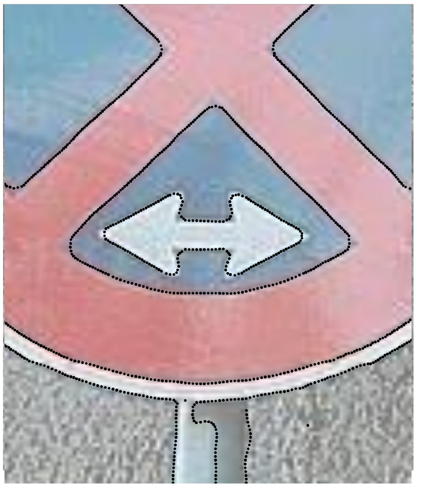

37

Edge Point Linking

(α,β)-reconstruction: region growing on Delaunay triangulation of edge points

(generalization of α-shapes from laser range scanning)

1. detect edge points
2. create Delaunay triangulation of edge points

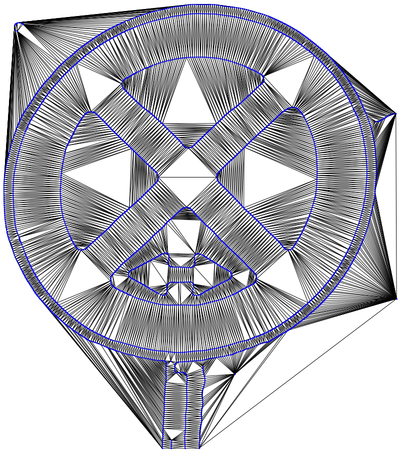

38

Edge Point Linking

(α,β)-reconstruction: region growing on Delaunay triangulation of edge points

(generalization of α-shapes from laser range scanning)

1. detect edge points
2. create Delaunay triangulation of edge points
3. seeds: connected sets of large triangles (radius $\beta > 2p+q$)
4. region growing: remove triangles with radius $\alpha > \max(p,q)$

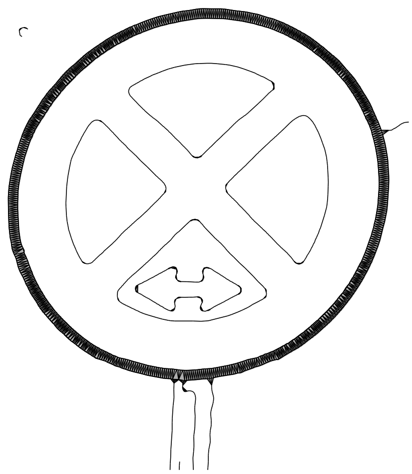

Multidimensional Image Processing IWR, Univ. of Heidelberg

Geometric Sampling Theorem

- Digitization model
  - shapes are r-stable
  - boundary sampled with errors p and q
  - (α,β)-reconstruction with $\alpha > \max(p,q)$ and $\beta > 2p+q$
  - requirement: homotopy equivalence, similar geometry
- Proved guarantee: Shapes are preserved

Meine, Köthe, Stelldinger: "A Topological Sampling Theorem for Robust Boundary Reconstruction and Image Segmentation", Discrete Applied Mathematics, to appear, 2008

Geometric Sampling Theorem

- Self diagnosis: larger errors result in thick boundaries (= uncertain edge positions)

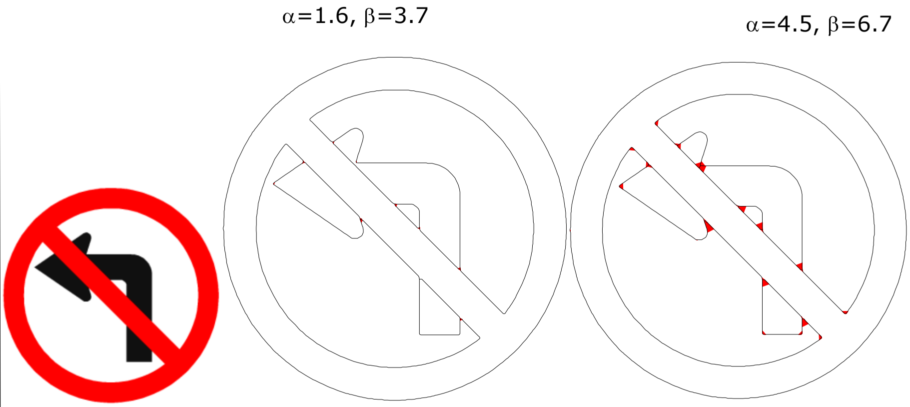

Multidimensional Image Processing IWR, Univ. of Heidelberg

Geometric Sampling Theorem

- Self diagnosis: larger errors result in thick boundaries (= uncertain edge positions)

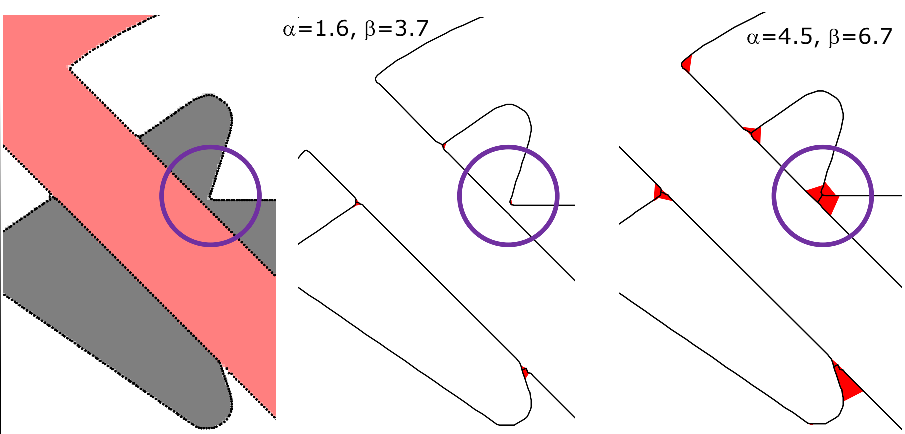

Geometric Sampling Theorem

- Higher errors = segmentation result is less certain

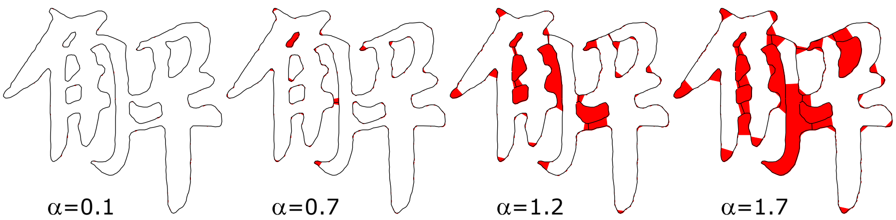

Geometric Sampling Theorem

- thick boundary can be thinned
  - remove long edges as long as topology is preserved

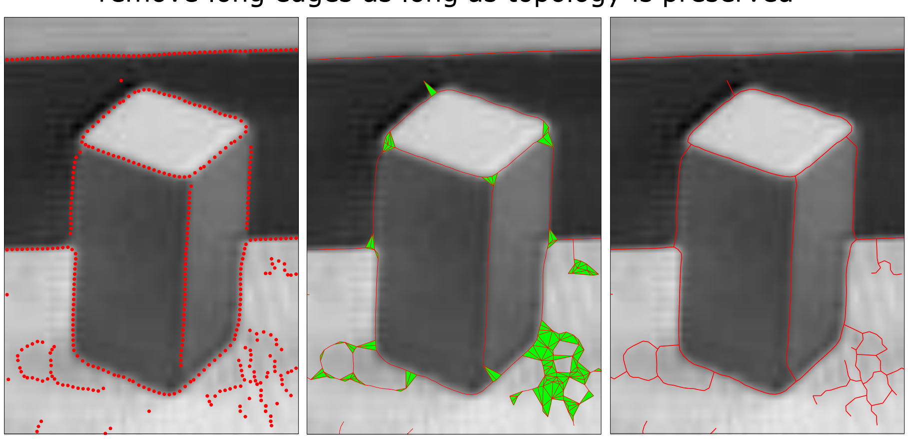

From Sampling Theorems to Error Analysis

- Geometric sampling theorem establishes connection between geometric projection and detected edge points
- Next task: determine errors to show which shapes are actually preserved in practice
  - systematic errors due to blurring (PSF, filters) and approximations in edge modeling
  - statistic errors due to noise and sampling round-off
- Basic distinction:
  pixel-accurate and subpixel-accurate edge detectors
  - pixel-accurate can be interpreted as subpixel-accurate plus rounding to nearest grid coordinate

Multidimensional Image Processing IWR, Univ. of Heidelberg

From Sampling Theorems to Error Analysis

Error sources in different analysis stages

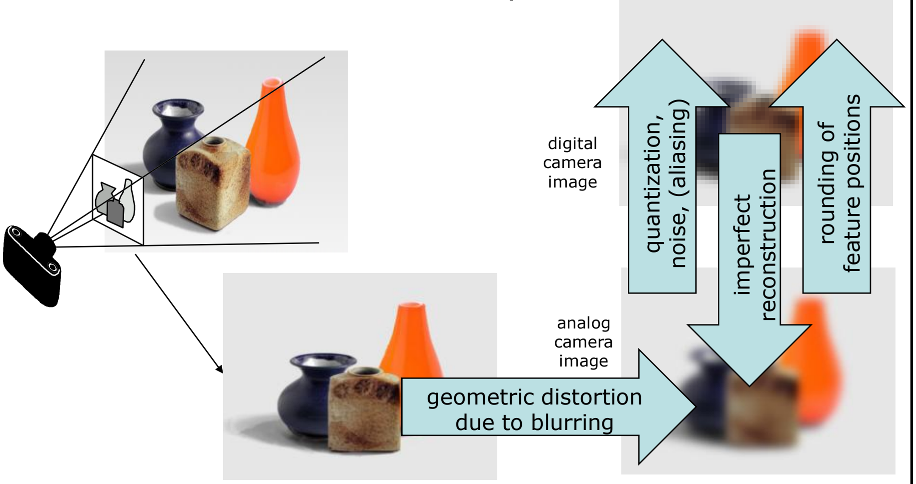

ideal geometric image

Pixel-Accurate vs. Subpixel Operators

Difference is very noticeable – rules of thumb:

- localization error of pixel-accurate operators 5 times as big
- much bigger differences for derivatives (tangents, curvature)

pixel-accurate thresholding

subpixel-accurate thresholding

details

47

Herzog: „Analyse historischer chinesischer Manuskripte", Diploma thesis, University of Hamburg, 2007

Rounding Errors of Pixel-Accurate Edge Detectors

- Example: 8-connected digital straight segments (grid-intersection digitization)
  - round grid-intersection points to nearest pixel center

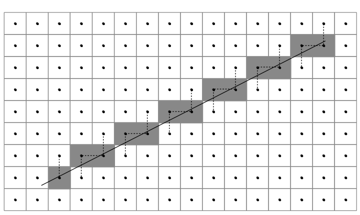

  - maximum rounding errors: $p = \sqrt{2}, \quad q = \frac{1}{2}$

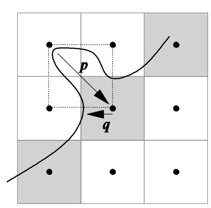

- similar errors for other DSS types and inter-pixel edges
- are added to other errors

48

From Sampling Theorems to Error Analysis

Statistical errors of subpixel edge detectors

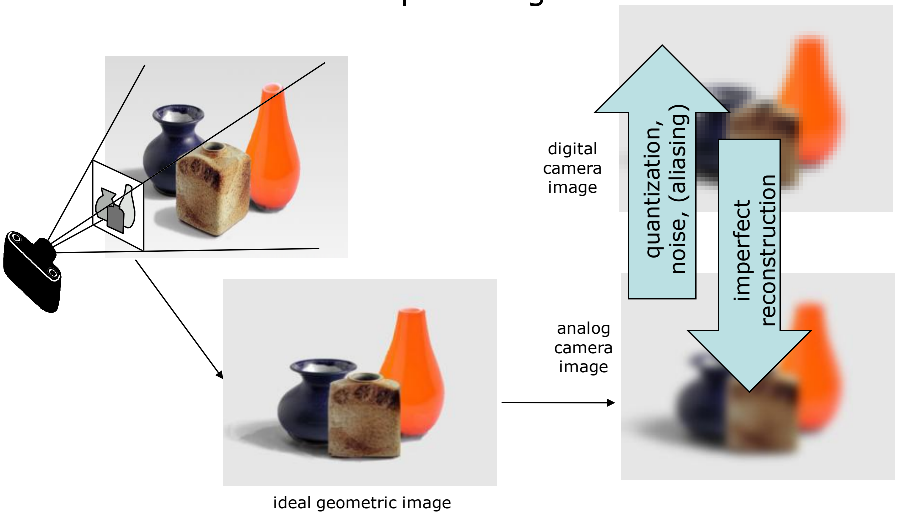

Subpixel-Accurate Image Analysis

- gray values also encode geometry! (due to PSF)
- many operators can take advantage of this additional geometric information
- some tasks cannot be reliably performed with pixel accurate methods
  - example: detection of critical points (local minima, maxima, saddles) on a Gaussian blob

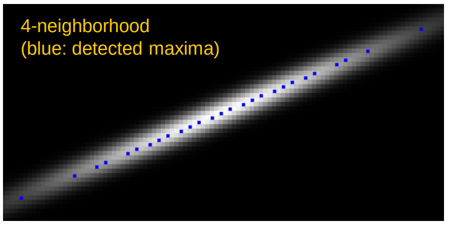

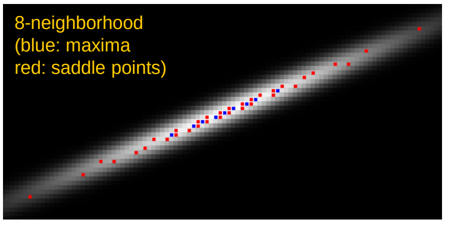

Stelldinger & Köthe: unpublished manuscript

Subpixel-Accurate Image Analysis

- gray values also encode geometry! (due to PSF)
- many operators can take advantage of this additional geometric information
- some tasks cannot be reliably performed with pixel accurate methods
  - example: detection of critical points (local minima, maxima, saddles) on a Gaussian blob

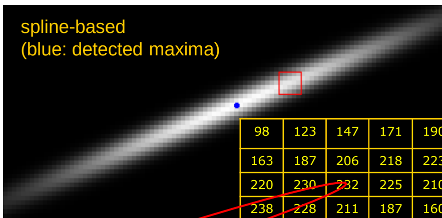

98	123	147	171	190
163	187	206	218	223
220	230	232	225	210
238	228	211	187	160
208	183	155	126	99

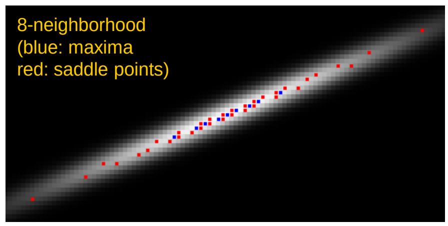

Stelldinger & Köthe: unpublished manuscript

Subpixel Edge Detectors (1)

- Subpixel watersheds
  - interpolate to double resolution
  - compute gradient magnitude
  - determine saddle points of spline

$$\delta\vec{x}^{(n)} = -\mathbf{H}(x^{(n)})^{-1}\nabla f(x^{(n)}) \text{ and } \det(\mathbf{H}(x^{(\infty)})) < 0$$

  - follow flow line to maximum of spline

$$\frac{\partial\vec{x}(t)}{\partial t} = \nabla f(\vec{x}(t)) \quad \text{(Runge-Kutta)}$$

  - suppress unwanted watersheds (e.g. threshold on *f*)

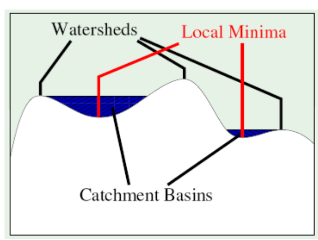

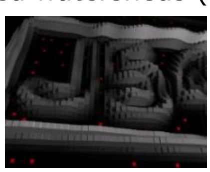

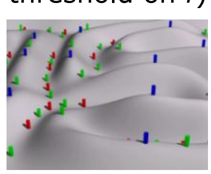

animations courtesy of Hans Meine

Steger: "Subpixel-Precise Extraction of Watersheds", ICCV '99, vol. II, pp. 884-890, 1999
Meine, Köthe: "Image Segmentation with the Exact Watershed Transform", VIIP '05, pp. 400-405, 2005

Subpixel Edge Detectors (2)

- Subpixel Haralick edge:
  - determine oriented 2nd derivative on oversampled image
  - find crossings between zero-contour of spline and grid lines (polynomial root finder)
  - complete by contour following along zero-crossings of spline (predictor-corrector method)
    - find next point along gradient
    - correct perpendicular onto zero contour (Newton)
- Likewise other subpixel zero-crossings (e.g. thresholding)

$$\frac{\partial \vec{x}(t)}{\partial t} = \pm \nabla f(\vec{x}(t))^{\perp}$$

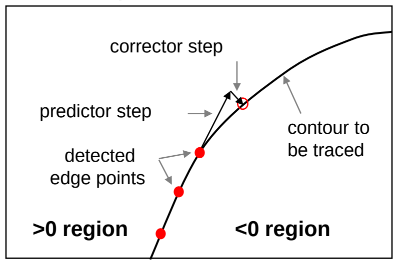

Allgower, Georg: "Numerical path following", In: Handbook of Numerical Analysis, vol. 5, pp. 3-207, 1997
Meine, Köthe: "Image Segmentation with the Exact Watershed Transform", VIIP '05, pp. 400-405, 2005

Statistical Errors of Subpixel Edge Detectors

very good agreement between theory and experiments

- reconstruction error can be neglected
- localization error due to noise (mainly sensor noise):

$$\mathrm{StdDev}[x] = SNR^{-1} \frac{\sqrt{6}}{4} \frac{\sigma^3}{\sigma_{\mathrm{filter}}^3} \qquad \text{with } \sigma = \sqrt{\sigma_{\mathrm{PSF}}^2 + \sigma_{\mathrm{filter}}^2} \approx 1\,\mathrm{pixel}$$

e.g. SNR=10 ⇨ maximum errors $p \approx q \approx 3\,\mathrm{StdDev}[x] \approx 0.35\,\mathrm{pixel}$

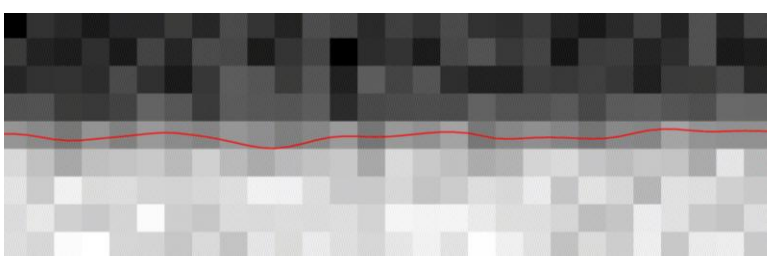

From Sampling Theorems to Error Analysis

Systematic errors of subpixel edge detectors

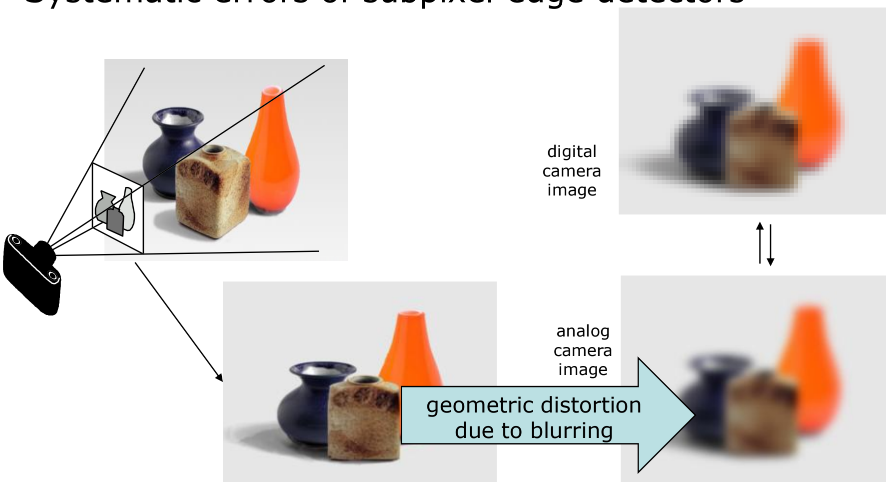

Systematic Errors of Subpixel Edge Detectors

very good agreement between theory and experiments

– curved step edges: curvature bias

$$\hat{R} - R \approx -\frac{\sigma^2}{2R} \leq 0.3\,\text{pixel}$$

– corner bias: $\approx 0.7\sigma$ for 90°,
$\approx 2.2\sigma$ for 15°,
but diverges as angle tends to zero

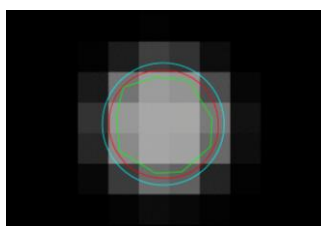

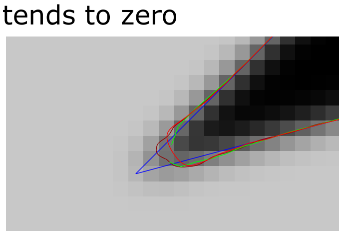

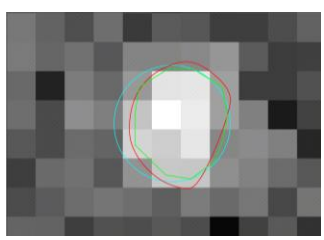

blue: true contour, red: subpixel watersheds, green: subpixel Canny edges

Systematic Errors of Subpixel Edge Detectors

- Unfortunately, errors near junctions are much higher
  - subpixel WS:
    - split-up into junctions of degree 3
    - distortions (increasing with low contrast)
  - subpixel Canny and Haralick
    - gaps at junctions
  - max. errors (experimental): $p \approx 3,\ q \approx 2$

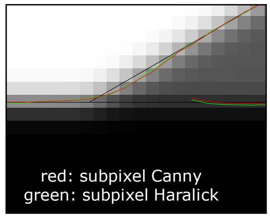

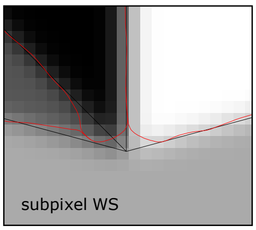

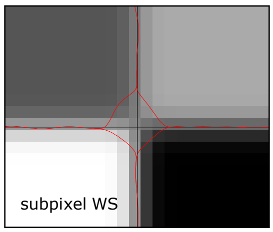

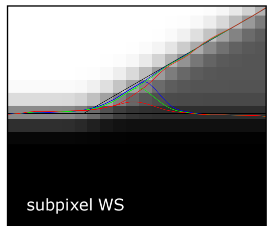

Multidimensional Image Processing IWR, Univ. of Heidelberg

Conclusions

- Conjoint analysis of continuous and discrete domains leads to much more reliable algorithms
  - Sampling theorems and error propagation determine what properties are preserved during digitization and reconstruction
  - Can determine necessary accuracy and optimal settings from image and scene properties
  - Sub-pixel accurate algorithms have 1/5 the error of pixel accurate ones
  - Algorithms work as predicted without additional tuning
  - More fun than trial-and-error

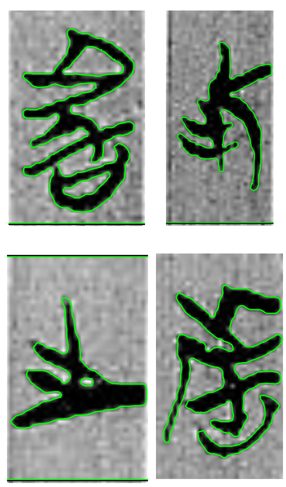

Open Problems

- Segmentation not always accurate enough
  - large errors near junctions
  - false positives or negatives due to noise and shading
- Generalizations of geometric sampling theorems to
  - 3D and higher dimensions
  - interest points and matching
  - object recognition

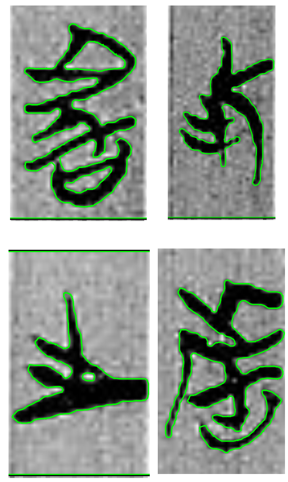

Thank you very much!