Lost at See

Using Random Walks for Scale Spaces in Sea Surface Satellite Image Analysis

Florian Sobieczky

Data Science Association at the University of Denver (CUD)
Wednesday, 21st of January, 2015

This talk is dedicated to the crew of the Niña, lost at sea since June 2014.

Retired University of Colorado professor Evi Nemeth, a computer engineer renowned in tech circles for her expertise in UNIX and Linux systems, is reportedly lost at sea between New Zealand and Australia.

The New Zealand Herald reported Thursday that Nemeth, 73, and six other people aboard a classic American racing yacht were last heard from June 4 and that rescue crews have "grave concerns" for the fate of those onboard the 84-year-old wooden schooner that set sail May 29 from New Zealand's Bay of Islands.

Evi Nemeth

The yacht, named Nina, was bound for Newcastle, Australia, the Herald reported.

Figure : Article from Dayly Camera, 6/27/2013

Satellite Imagery used for Sea Surface observation

Figure : 239 + 162 victims of airplane accident overseas

Motivation and Problem Formulation

Scale Spaces

Scale Spaces

Statistics of Natural Images

Fast detectors of Sea Surface Objects

Satellite Imagery used for Sea Surface observation

Figure : Image from sample gallery of Digital Globe.

What do satellite sea surface images look like?

Satellite Imagery used for Sea Surface observation

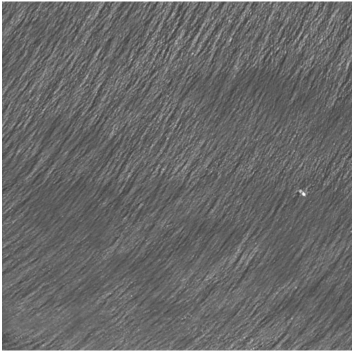

Figure : Image from sample gallery of Digital Globe.

How can boats, debris, or other objects be detected?

Satellite Imagery used for Sea Surface observation

Figure : Image from sample gallery of Skytruth (Copyright Google 2007).

What are the statistical properties of ‘natural ocean pictures’?
How do objects appear ‘untypical’ in these statistics?

Satellite Imagery used for Sea Surface observation

Figure : Image from sample gallery of Skytruth (Copyright Google 2007).

What are the statistical properties of 'natural ocean pictures'?
How do objects appear 'untypical' in these statistics?

Amount of Data

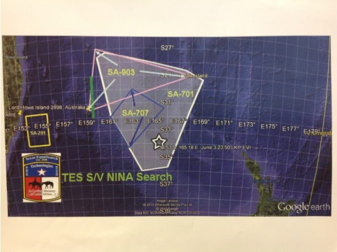

Figure : Image from Texas Equusearch

- Typical GEO-TIFF file size (e.g. samples from DIGITAL Globe): $10000x10000$ pixels, corresponding to $(3km)^2$.

Amount of Data

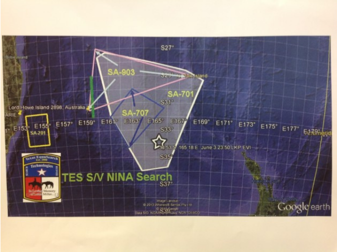

Figure : Image from Texas Equusearch

- Typical GEO-TIFF file size (e.g. samples from DIGITAL Globe): $10000x10000$ pixels, corresponding to $(3km)^2$.
- Search area above $\simeq 555.000(km)^2 \simeq 61.000$ GEOTiffs

Amount of Data

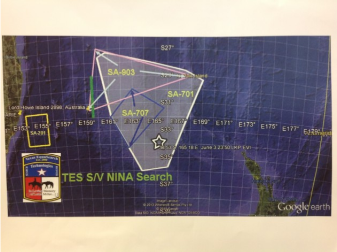

Figure : Image from Texas Equusearch

- Typical GEO-TIFF file size (e.g. samples from DIGITAL Globe): 10000x10000 pixels, corresponding to $(3km)^2$.
- Search area above $\simeq 555.000(km)^2 \simeq 61.000$ GEOTiffs
- about 2.5 Million 500x500 Pixel images $\simeq$ 1 month (1 sec/pic of size 500 x 500 pixels)

Scale Spaces

- Space of images $\mathcal{S} : D \to \mathbb{R}$
  (e.g. $D = (\epsilon\mathbb{Z})^2 \bigcap [0,1]^2$, and $R = \{0, \dots, 255\}$
  or $D = \mathbb{R}^2$, and $R = \mathbb{R}$)

Scale Spaces

- Space of images $\mathcal{S} : D \to \mathbb{R}$
(e.g. $D = (\epsilon\mathbb{Z})^2 \bigcap [0,1]^2$, and $R = \{0, \ldots, 255\}$
or $D = \mathbb{R}^2$, and $R = \mathbb{R}$)
- A scale space is an indexed family of images (Index: Scale-Parameter; e.g. $\epsilon > 0$) relative to some (inital) image
- $\Phi_t : \mathcal{S} \times \mathcal{S} \to \mathcal{S}$, with $t \in (0, \infty) =:$ 'Scale parameter'

Scale Spaces

- Space of images $\mathcal{S} : D \to \mathbb{R}$
(e.g. $D = (\epsilon\mathbb{Z})^2 \bigcap [0,1]^2$, and $R = \{0, \dots, 255\}$
or $D = \mathbb{R}^2$, and $R = \mathbb{R}$)
- A scale space is an indexed family of images (Index: Scale-Parameter; e.g. $\epsilon > 0$) relative to some (inital) image
- $\Phi_t : \mathcal{S} \times \mathcal{S} \to \mathcal{S}$, with $t \in (0, \infty) =:$ 'Scale parameter'
- V discrete (Lindeberg 94) or continuous (Iijima '62, Otsu '81, Witkin 83, Koenderink 84)

Scale Spaces

- Space of images $\mathcal{S} : D \to \mathbb{R}$
  (e.g. $D = (\epsilon\mathbb{Z})^2 \bigcap [0,1]^2$, and $R = \{0, \dots, 255\}$
  or $D = \mathbb{R}^2$, and $R = \mathbb{R}$)
- A scale space is an indexed family of images (Index: Scale-Parameter; e.g. $\epsilon > 0$) relative to some (inital) image
- $\Phi_t : \mathcal{S} \times \mathcal{S} \to \mathcal{S}$, with $t \in (0, \infty) =$: 'Scale parameter'
- V discrete (Lindeberg 94) or continuous (Iijima '62, Otsu '81, Witkin 83, Koenderink 84)
- Idea: Splitting up information of image into different scales which label different 'derived images' according to different degree of detail (Burt 81, Crowley 81, Witkin 83)

Scale Spaces

- Space of images $\mathcal{S} : D \to \mathbb{R}$
  (e.g. $D = (\epsilon\mathbb{Z})^2 \bigcap [0,1]^2$, and $R = \{0, \ldots, 255\}$
  or $D = \mathbb{R}^2$, and $R = \mathbb{R}$)
- A scale space is an indexed family of images (Index: Scale-Parameter; e.g. $\epsilon > 0$) relative to some (inital) image
- $\Phi_t : \mathcal{S} \times \mathcal{S} \to \mathcal{S}$, with $t \in (0, \infty) =:$ 'Scale parameter'
- V discrete (Lindeberg 94) or continuous (Iijima '62, Otsu '81, Witkin 83, Koenderink 84)
- Idea: Splitting up information of image into different scales which label different 'derived images' according to different degree of detail (Burt 81, Crowley 81, Witkin 83)
- Typical Applications: Scale-Detection, Feature-Recognition, Edge-Detection, Image-Registration, Object-Classification

Scale Spaces: Examples

- Gaussian Scale Space: $\Phi_t[g] = \phi_t * g(x)$ with $\phi_t(x) = \exp(-x^2/t)/(2\pi t^2)$.

Scale Spaces: Examples

- Gaussian Scale Space: $\Phi_t[g] = \phi_t * g(x)$ with $\phi_t(x) = \exp(-x^2/t)/(2\pi t^2)$.
- $\Phi_t[g](x) = \mathbb{E}_x[g(B_t)]$ with B_t Brownian Motion

Scale Spaces: Examples

- Gaussian Scale Space: $\Phi_t[g] = \phi_t * g(x)$ with $\phi_t(x) = \exp(-x^2/t)/(2\pi t^2)$.
- $\Phi_t[g](x) = \mathbb{E}_x[g(B_t)]$ with B_t Brownian Motion
- $\Phi_t[g](x) = \mathbb{E}_x[g(X_t)]$ where X_t is a diffusion, so

$$\Phi_t[g](x) = \int_\Omega g(y)K(x, y, t)dy$$

where $K(x, y, t) \neq K(x - y, 0, t)$, so this Scale Space **doesn't** have translational invariance (not a convolution). (see Leo Grady: 'Random Walks for Image Segmentation', IEEE Tr. PAMI, 2006)

Scale Spaces: Examples

- Gaussian Scale Space: $\Phi_t[g] = \phi_t * g(x)$ with $\phi_t(x) = \exp(-x^2/t)/(2\pi t^2)$.
- $\Phi_t[g](x) \ = \ \mathbb{E}_x[g(B_t)]$ with B_t Brownian Motion
- $\Phi_t[g](x) \ = \ \mathbb{E}_x[g(X_t)]$ where X_t is a diffusion, so

$$\Phi_t[g](x) = \int_\Omega g(y)K(x,y,t)dy$$

where $K(x,y,t) \neq K(x-y,0,t)$, so this Scale Space **doesn't** have translational invariance (not a convolution). (see Leo Grady: 'Random Walks for Image Segmentation', IEEE Tr. PAMI, 2006)

Scale Spaces: Examples

- Gaussian Scale Space: $\Phi_t[g] = \phi_t * g(x)$ with $\phi_t(x) = \exp(-x^2/t)/(2\pi t^2)$.
- $\Phi_t[g](x) = \mathbb{E}_x[g(B_t)]$ with B_t Brownian Motion
- $\Phi_t[g](x) = \mathbb{E}_x[g(X_t)]$ where X_t is a diffusion, so

  $$\Phi_t[g](x) = \int_\Omega g(y)K(x,y,t)dy$$

  where $K(x,y,t) \neq K(x-y,0,t)$, so this Scale Space **doesn't** have translational invariance (not a convolution). (see Leo Grady: 'Random Walks for Image Segmentation', IEEE Tr. PAMI, 2006)
- Non-linear: $u(t,x) = \Phi_t[g](x) = \int_\Omega g(y)K(u(t,y);x,y,t)dy$ (see J. Weickert: 'Anisotropic Diffusion', Teubner, 1998)

Properties of Scale Spaces

Figure : San Francisco at Golden Gate Bridge, Sample of Digital Globe

- Usually defined by ‘Scale-Space Axioms’, Causality, Linearity, Scale Invariance, Semi-group property, Isotropy, Homogeneity, ...

Properties of Scale Spaces

Figure : San Francisco at Golden Gate Bridge, Sample of Digital Globe

- Usually defined by 'Scale-Space Axioms', Causality, Linearity, Scale Invariance, Semi-group property, Isotropy, Homogeneity, ...
- Unique solution fulfilling all properties: Gaussian Scale Space

Two scale spaces for edge-preserving smoothing

Figure : GIMP's 'Selective Gaussian Blurr' (top row) and Random Walk Smoothing (bottom row). Original: Left Column. Random Walk is the 'Delayed Random Walk' after 2, 3 and 4 steps, with threshold of 20 greyvalues out of 256. Gaussian blurr with comparable removal of noise sooner deteriorates fine detail.

L. Grady's Model:

L. Grady's Model:

- In Grady's Model, graph Laplacian L is set up for a weighted lattice graph, wheights $\sim \exp(-a|f(x) - f(y)|)$, $x, y \in V$.

L. Grady's Model:

- In Grady's Model, graph Laplacian L is set up for a weighted lattice graph, wheights $\sim \exp(-a|f(x) - f(y)|)$, $x, y \in V$.
- Several exit points are defined (RW is 'killed' there), one for each Segment: 'Boundary of the Graph'.

L. Grady's Model:

- In Grady's Model, graph Laplacian L is set up for a weighted lattice graph, wheights $\sim \exp(-a|f(x) - f(y)|)$, $x, y \in V$.
- Several exit points are defined (RW is 'killed' there), one for each Segment: 'Boundary of the Graph'.
- Each exit carries label.

L. Grady's Model:

- In Grady's Model, graph Laplacian L is set up for a weighted lattice graph, wheights $\sim \exp(-a|f(x)-f(y)|)$, $x, y \in V$.
- Several exit points are defined (RW is 'killed' there), one for each Segment: 'Boundary of the Graph'.
- Each exit carries label.
- Instead of computing the eigenvectors, for each point $x \in V$ and time $t > 0$ the exit measure (harmonic measure) is computed

L. Grady's Model:

- In Grady's Model, graph Laplacian L is set up for a weighted lattice graph, wheights $\sim \exp(-a|f(x) - f(y)|)$, $x, y \in V$.
- Several exit points are defined (RW is 'killed' there), one for each Segment: 'Boundary of the Graph'.
- Each exit carries label.
- Instead of computing the eigenvectors, for each point $x \in V$ and time $t > 0$ the exit measure (harmonic measure) is computed
- Point x obtains label of exit with highest exit measure.

L. Grady's Model:

- In Grady's Model, graph Laplacian L is set up for a weighted lattice graph, wheights $\sim \exp(-a|f(x) - f(y)|)$, $x, y \in V$.
- Several exit points are defined (RW is 'killed' there), one for each Segment: 'Boundary of the Graph'.
- Each exit carries label.
- Instead of computing the eigenvectors, for each point $x \in V$ and time $t > 0$ the exit measure (harmonic measure) is computed
- Point x obtains label of exit with highest exit measure.
- Advantage: L with boundary is invertible (RW properly substochastic):
  Solving a linear system, instead of computing eigenvectors

L. Grady's Model:

- In Grady's Model, graph Laplacian L is set up for a weighted lattice graph, wheights $\sim \exp(-a|f(x) - f(y)|)$, $x, y \in V$.
- Several exit points are defined (RW is 'killed' there), one for each Segment: 'Boundary of the Graph'.
- Each exit carries label.
- Instead of computing the eigenvectors, for each point $x \in V$ and time $t > 0$ the exit measure (harmonic measure) is computed
- Point x obtains label of exit with highest exit measure.
- Advantage: L with boundary is invertible (RW properly substochastic):
  Solving a linear system, instead of computing eigenvectors
- Solves 'Bottleneck' Problem.

L. Grady's Model

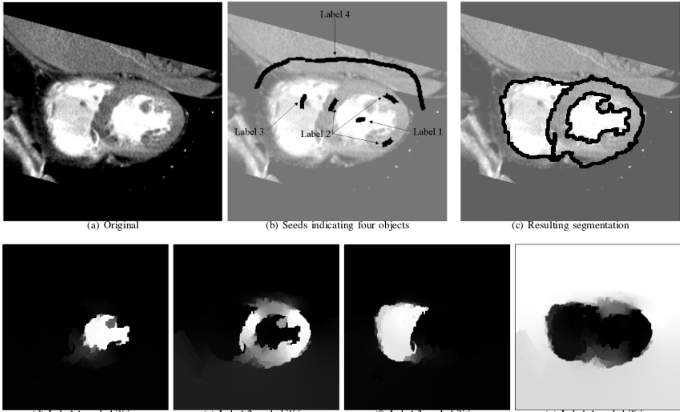

(a) Original (b) Seeds indicating four objects (c) Resulting segmentation
(d) Label 1 probabilities (e) Label 2 probabilities (f) Label 3 probabilities (g) Label 4 probabilities

Figure : Result of a Segmentation using Seeds (=Starting points of Random Walks) and the Harmonic Measure (=Hitting measure)

Wavelet-coefficients

- Wavelets form a complete l^2-basis for S

Wavelet-coefficients

- Wavelets form a complete l^2-basis for S
- Wavelets are localized (their support is a bounded region in image)

Wavelet-coefficients

- Wavelets form a complete l^2-basis for S
- Wavelets are localized (their support is a bounded region in image)
- (Simplest) 'Haar'-wavelet –
$g(x) = \frac{1}{2}(g(x) + g(x+1)) \ + \ g(x) - g(x+1)$

Wavelet-coefficients

- Wavelets form a complete l^2-basis for S
- Wavelets are localized (their support is a bounded region in image)
- (Simplest) 'Haar'-wavelet –
  $g(x) = \frac{1}{2}(g(x) + g(x+1)) \ + \ g(x) - g(x+1)$
- Observation: Coefficients of 1-st subband follow distribution on 'natural images'

Wavelet-coefficients

- Wavelets form a complete l^2-basis for S
- Wavelets are localized (their support is a bounded region in image)
- (Simplest) ‘Haar’-wavelet – $g(x) = \frac{1}{2}(g(x) + g(x+1)) \ + \ g(x) - g(x+1)$
- Observation: Coefficients of 1-st subband follow distribution on ‘natural images’
- Idea: Use to determine ‘untypical’ features in image

Wavelet-coefficients

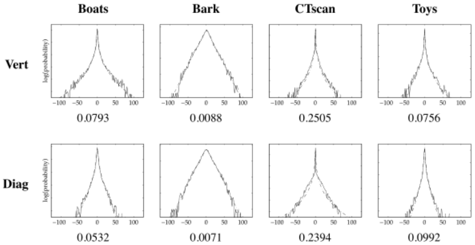

Figure : From Buccigrossi, Simoncelli: 'Image Compression via Joint Statistical Characterization in the Wavelet Domain': Measured Distribution of discrete Gradient (= coefficient of First Sub-band) $g(x+1) - g(x)$: Natural Images have usually a wider Peak...

Wavelet-coefficients

Figure : The Images ‘Bark’, ‘Boats’, ‘CTScan’, and ‘Toys’

What to look for

- The ‘naked eye’ is still one of the best filters.

What to look for

- The ‘naked eye’ is still one of the best filters.
- Task is to filter out ‘any strange object’ to reduce the amount of data that has to be viewed.

What to look for

- The ‘naked eye’ is still one of the best filters.
- Task is to filter out ‘any strange object’ to reduce the amount of data that has to be viewed.
- Use changes of the function across 1-dim. Profiles

What to look for

- The ‘naked eye’ is still one of the best filters.
- Task is to filter out ‘any strange object’ to reduce the amount of data that has to be viewed.
- Use changes of the function across 1-dim. Profiles
- Use untypical statistic of objects

What to look for

- The ‘naked eye’ is still one of the best filters.
- Task is to filter out ‘any strange object’ to reduce the amount of data that has to be viewed.
- Use changes of the function across 1-dim. Profiles
- Use untypical statistic of objects
- Possibly use Scale Space to make features more pronounced

Sea Surface Object detection on Satellite Images

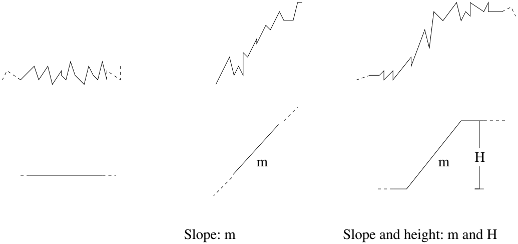

Figure : Models for features to detect, and their smoothed version.

Example for an object to be reviewed by naked eye

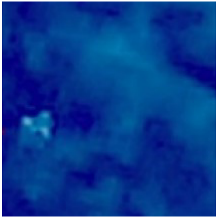

Figure : Example of possible object (life-raft etc.) shown in the media in connection with the search for flight MH370

Example for an object to be reviewed by naked eye

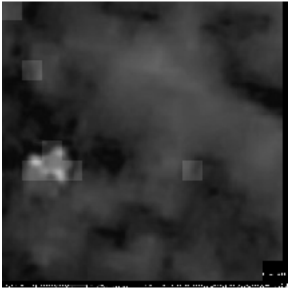

Figure : Object can be easily recognized by strong gradient around object, nowhere else to be found in the image

Wavelet-coefficients

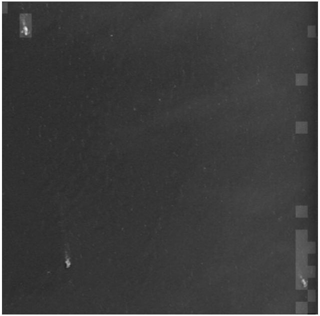

Figure : If wave crests are too pronounced, looking for largest gradients will confuse filter.

Quantization of picture: Work on Bitplanes alone

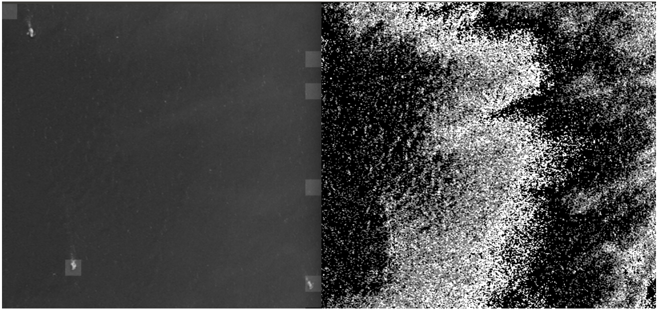

Figure : Solution: Look at quantized picture, and take gradient then. This increases the focus on parts in which the large gradients belong to object0-boundaries (due to repetition)

Example for an object to be reviewed by naked eye

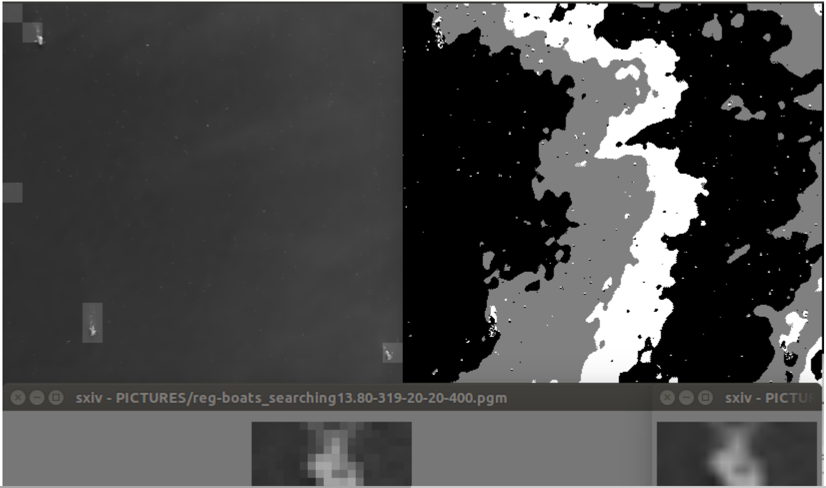

The Nina

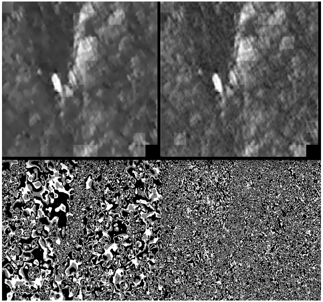

Figure : On Oct 15, 2013, Media (e.g.Dayly Mail.com) reported Texas EquuSearch found satellite image well fitting the 'Nina' at -28.784317, 164.45064.

The Nina

Figure : Thank you for your interest!