Spatio-temporal video segmentation with shape growth or shrinkage constraint

Yuliya Tarabalka[1], Guillaume Charpiat[2],
Ludovic Brucker[3] and Bjoern H. Menze[4]

[1,2]Inria Sophia-Antipolis, AYIN[1] and STARS[2] teams, France
[3]NASA GSFC & Universities Space Research Association, MD, USA
[4]Computer Vision Laboratory, ETH Zurich, Switzerland & Inria Sophia-Antipolis, Asclepios team
e-mail: yuliya.tarabalka@inria.fr

February 5, 2014

Outline

1. Introduction
2. Enforcing shape growth/shrinkage in graph cuts
3. Applications
   - Melting sea ice in satellite images
   - Growing burned areas in satellite data
   - Growing tumor in 3D medical scans
4. Conclusions and perspectives

Motivation: how to segment a melting floe?

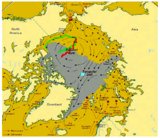

- **Track** multiyear ice floes from **low-resolution** images
  - Advanced Microwave Scanning Radiometer, 6.25 km/pix
- **How to segment moderate-resolution** images?
  - Moderate-Resolution Imaging Spectroradiometer, 250 m/pix

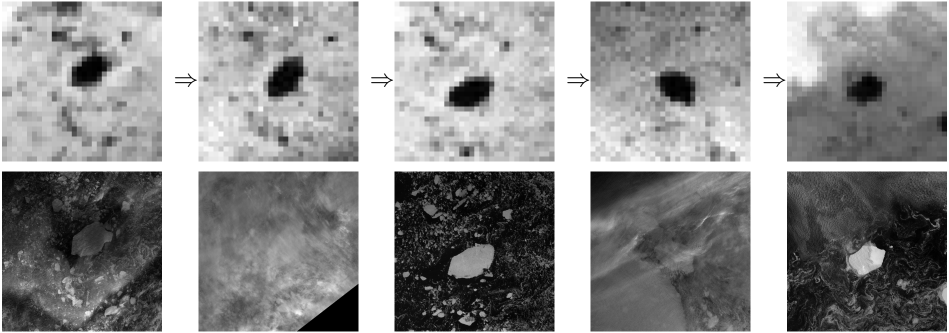

Motivation: how to segment a melting floe?

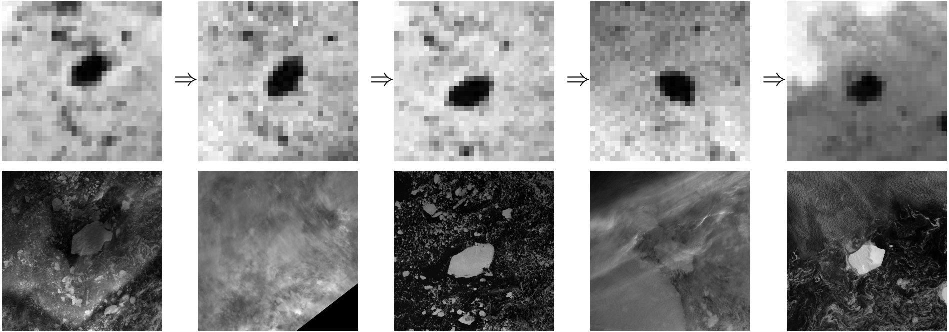

- **Difficulties:**
  - Low signal-to-noise ratio
  - Low contrast between neighboring objects
  - Foreground & background intensity distributions vary significantly over time
  - Foreground can be occluded or undistinguishable from a part of the background
  - Data for some pixels can be missing

How to segment a melting floe?

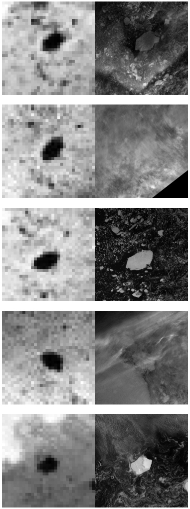

Solution:

- **Temporal coherence** in video sequences = a lot of information, not available for a single image
- Take advantage of both **past and future** data (omniscient approach)

How to exploit temporal coherence?

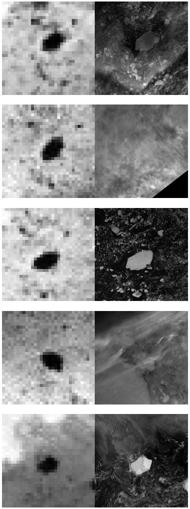

Previous works:	**Our problem:**
• Rely on coherence of foreground/background intensity distributions over time [Shi'98, Grundmann'10]	• Foreground/background intensity distributions vary significantly over time
• Introduce shape priors into image segmentation [Cremers'02, Schoenemann'07]	• Shape prior is unknown • Shape is changing over time
• Smooth 2D+T spatio-temporal volume [Riklin-Raviv'10, Wolz'10]	• Rapid shrinkage events will be underestimated

How to exploit temporal coherence?

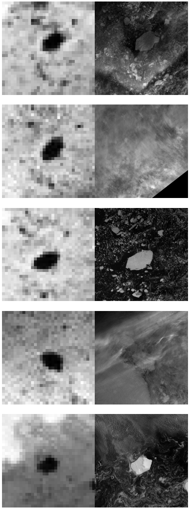

Shape prior information:

- Object **monotonously shrinks** in time (multiyear ice floe can only melt)

$\Downarrow$

Solution:

- Introduce shape **shrinkage constraint** in spatio-temporal video segmentation

Objective

- **Aim:**
  - To segment **monotonously growing or shrinking** shapes,
  - From time sequences of extremely noisy images,
  - In a low computational time
- **Method:**
  - Formulate video segmentation as an optimization problem,
  - Using the spatio-temporal graph of pixels,
  - With shape **growth or shrinkage constraint** expressed with directed infinite links.
  - Globally-optimal solution is computed with a **graph cut**
- **Examples of growing shapes:**

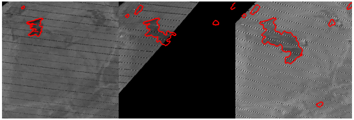

Savanna fires, 2D satellite data

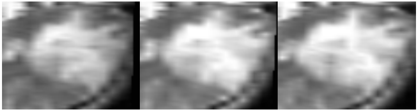

Brain tumor, 3D medical MRI volumes

Outline

Graph cut for image segmentation

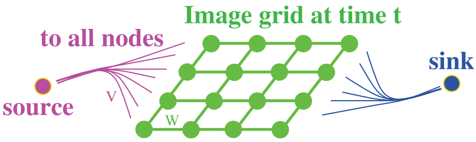

- **Graph-cut** = tool to find **globally-optimal segmentation***:
  1. map an image onto a graph
  2. minimize a criterion of the form:

$$E(L) = \sum_{\text{pixels } i} V_i(L_i) + \sum_{i \sim j} W_{i,j}(L_i, L_j)$$

  - L_i = label of pixel i
  - individual potential $V_i(L_i^t)$ = penalty for a pixel i to have a label L_i
  - $W_{i,j}(L_i, L_j)$ = submodular interaction term between adjacent pixels i and j: $W_{i,j}(0,0) + W_{i,j}(1,1) \leqslant W_{i,j}(0,1) + W_{i,j}(1,0)$

*[Boykov&Kolmogorov'04]

Graph cut for image segmentation

- **Goal:** Compute $T (t \in [1, T])$ segmentation maps $L^t = \{L^t_{(x,y)} \in [0, 1], x = [1..H], y = [1..W]\}$,

$$L^t_{(x,y)} = \begin{cases} 1, & \text{if } (x, y) \in \text{foreground at time t;} \\ 0, & \text{otherwise.} \end{cases}$$

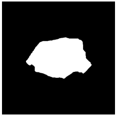

- **Graph-cut segmentation:**
  1. map each image $I(t)$ onto a graph
  2. minimize a submodular energy of the form:

$$E^t(L) = \sum_{\text{pixels } i} V_i^t(L_i^t) + \sum_{i \sim j} W_{i,j}^t(L_i^t, L_j^t)$$

- L_i^t = label of pixel i at time t
- individual potential $V_i^t(L_i^t)$ = penalty for a pixel i to have a label L_i^t
- $W_{i,j}^t(L_i^t, L_j^t)$ = interaction term between adjacent pixels i and j

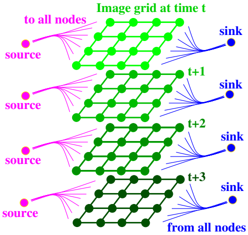

Independent segmentation of T images

Enforcing shape growth

- **Shape growth** = property that the foreground cannot lose any pixel when time advances
- **Enforcing shape growth** (label 1 = foreground, label 0 = background)
  - $\Leftrightarrow$ if $L_i^{t_1} = 1$, then $L_i^{t_2} = 1 \; \forall t_2 > t_1$
  - $\Leftrightarrow$ pair of pixels $((x, y, t), (x, y, t+1))$ cannot have the pair of labels $(1, 0)$
  - $\Leftrightarrow$ **directed infinite link** from each pixel to its predecessor in time

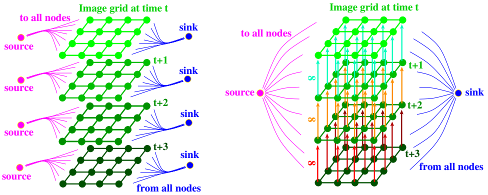

Independent segmentation of T images ➡ Joint segmentation with growth enforcement

t

$t+1$

$t+2$

Graph cut with shape growth constraint

- Segment **jointly** all T images together
  - apply graph cut to the 3D grid $W \times H \times T$
  - with directed infinite links in time
- Criterion minimized: $E = \sum_t E^t$ under the **constraint of shape growth**:

$$E \;=\; \sum_{\text{pixels } i} V_i(L_i) \;+\; \sum_{i \sim j} W_{i,j}(L_i, L_j) \;+\; \infty \sum_t \delta_{L_i^t > L_i^{t+1}}$$

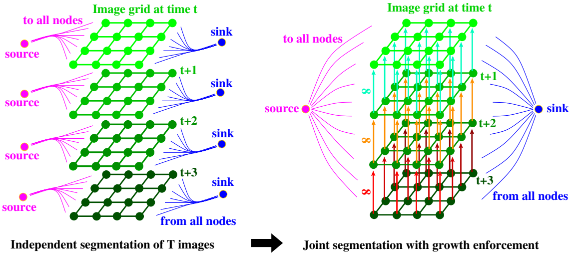

Extensions

- **Shape shrinkage:** reverse the direction of infinite links
  - from each pixel to its successor in time
- **3D shape:** set directed infinite links for all voxel pairs $((x, y, z, t), (x, y, z, t-1))$
- **Encourage, but not impose shape growth:** replace directed infinite links by directed finite links
- **Inter-sequences inclusion constraint:** foreground in one sequence has to be included in foreground of another sequence
  - see figure
- **Weighting frames by reliability**
  - strong level of noise at time $t \rightarrow$ multiply E^t by a small reliability factor < 1

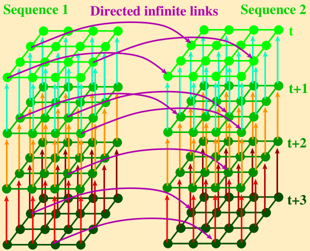

Segmenting jointly two sequences S1 and S2, by enforcing the foreground of S1 to contain the foreground of S2, with directed infinite links between all pixels of coordinates (x, y, t), from S1 towards S2

Complexity

- Precise **theoretical worst case complexity**:
  - Depends on the max-flow algorithm used
  - Ranges from quasi-quadratic to cubic
- **In practice:** computational complexity is typically much faster
  - We used the graph-cut algorithm of Boykov & Kolmogorov
  - Total computational time grows linearly with the number of frames

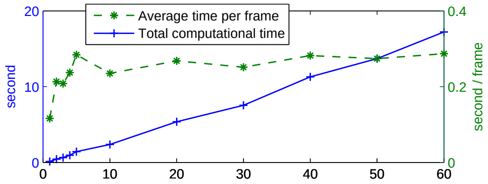

- Memory requirements
  - Long sequences of big images ⇒ graph-cut implementations for massive grids [Delong&Boykov'08]

Rewriting as a multi-label problem

Sequence segmentation with shape growth constraint

- Successive labels $L_i(t)$ of a pixel i over time **might change only once**
  - 0 (background) $\rightarrow$ 1 (foreground)
- This vector of labels $L_i(t)$ is of the form $(0, 0, \ldots, 0, 1, \ldots, 1)$
  - **can be represented by the time index τ_i of the first 1**
  - $\tau_i \in [1, T+1]$, with $T+1$ meaning "never"

$\Updownarrow$

Multi-label problem on a single image

- Can be expressed in the MRF form (slide 10) with:
  - $V_i(\tau_i) \;:=\; \sum_{t<\tau_i} V_i^t(0) + \sum_{t \geqslant \tau_i} V_i^t(1)$
  - $W_{i,j}(\tau_i, \tau_j) \;:=\; \sum_{t<\min(\tau_i,\tau_j)} W_{i,j}^t(0,0) + \sum_{\tau_i \leqslant t < \tau_j} W_{i,j}^t(1,0) + \sum_{\tau_j \leqslant t < \tau_i} W_{i,j}^t(0,1) + \sum_{t \geqslant \max(\tau_i,\tau_j)} W_{i,j}^t(1,1)$

Multi-label problem (τ_i = time index of the first 1)

- Can be expressed in the **MRF form** (slide 10) with:
  - $V_i(\tau_i) := \sum_{t<\tau_i} V_i^t(0) + \sum_{t\geqslant\tau_i} V_i^t(1)$
  - $W_{i,j}(\tau_i,\tau_j) := \sum_{t<\min(\tau_i,\tau_j)} W_{i,j}^t(0,0) + \sum_{\tau_i\leqslant t<\tau_j} W_{i,j}^t(1,0) + \sum_{\tau_j\leqslant t<\tau_i} W_{i,j}^t(0,1) + \sum_{t\geqslant\max(\tau_i,\tau_j)} W_{i,j}^t(1,1)$
- Submodularity of the binary interaction terms W^t in each frame $\Rightarrow$ **submodularity of the multilabel interaction** term W:

$$W_{i,j}(\tau_1,\tau_2) + W_{i,j}(\tau_1',\tau_2') \leqslant W_{i,j}(\tau_1,\tau_2') + W_{i,j}(\tau_1',\tau_2)$$

for all labels satisfying $\tau_1 \leqslant \tau_1'$ and $\tau_2 \leqslant \tau_2'$

$\Updownarrow$

This MRF-based energy can be minimized globally efficiently

Multi-label problem (τ_i = time index of the first 1)

- Can be expressed in the **MRF form** (slide 10) with:
  - $V_i(\tau_i) := \sum_{t<\tau_i} V_i^t(0) + \sum_{t\geqslant\tau_i} V_i^t(1)$
  - $W_{i,j}(\tau_i, \tau_j) := \sum_{t<\min(\tau_i,\tau_j)} W_{i,j}^t(0,0) + \sum_{\tau_i\leqslant t<\tau_j} W_{i,j}^t(1,0) + \sum_{\tau_j\leqslant t<\tau_i} W_{i,j}^t(0,1) + \sum_{t\geqslant\max(\tau_i,\tau_j)} W_{i,j}^t(1,1)$

- **Particular case:** interaction terms W^t do not depend on t
  - Interaction term W can be rewritten as a convex function g of $(\tau_i - \tau_j)$
  - Ishikawa's construction [Ishikawa'03] can be applied

- **Advantages of our formulation:**
  - Interaction terms can depend on t
  - Inclusion constraint can be enforced in spatial or/and time subregions only

Outline

Automated ice floe analysis

- The melting of sea ice is correlated to:
  - increases in sea surface temperature
  - associated climatic changes

$\Downarrow$

- It is important to:
  - monitor sea ice evolution
  - develop methods for automated analysis of satellite measurements

Objective:

- Determine how rapidly a multiyear ice floe can melt
- By analyzing NASA Aqua measurements:
  - Advanced Microwave Scanning Radiometer - Earth Observing System (AMSR-E)
  - Moderate-Resolution Imaging Spectroradiometer (MODIS)

Data set

- **Target:** Arctic multiyear sea ice floe
  - 45-day sequence (mi-August - end of September 2008)
- **AMSR-E data**: 6.25 km/pix, 89 GHz, 32 × 32 pixels
  - multiyear ice has a low microwave emissivity

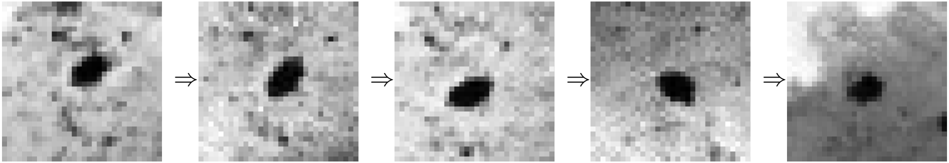

- **MODIS data**: band 1, 250 m/pix, 0.620-0.670 μm, 800 × 800 pixels

Floe detection

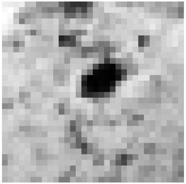

- We denote
  - upscaled AMSR-E images smoothed by Gaussian: $A^t, t \in [1, T]$
  - MODIS images: $I^t, t \in [1, T]$
- On AMSR-E images, multiyear ice is darker than water, young ice and clouds

  $\Rightarrow$ Estimate for each time t:
  - **reliable** region of the **foreground** R_F
  - **reliable** region of the **background** R_B

Floe alignment

- The images **must be aligned**:

$$\forall t, x, y, \qquad L^{t+1}_{(x,y)} = 1 \implies L^{t}_{(x,y)} = 1$$

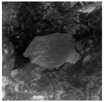

- Compute:
  - histograms of the intensities I^t of the floe, $p^t(I|F)$, and of the background, $p^t(I|B)$
  - map of floe probabilities:

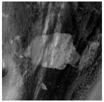

$$p^t(F|I) = \frac{p^t(I|F)P^t(F)}{p^t(I|F)P^t(F) + p^t(I|B)P^t(B)},$$

$$P^t(B) = \frac{A^t - \min_{x,y} A^t_{(x,y)}}{\max_{x,y} A^t_{(x,y)} - \min_{x,y} A^t_{(x,y)}}, \quad P^t(F) = 1 - P^t(B).$$

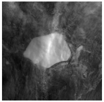

- Align images: exhaustive searching over rigid motions
  - **maximize the correlation between maps of foreground probabilities** at the current and previous moments
- To reach a pixelic alignment precision
  - maximize correlation between $\nabla p^t(F|I)$ and $[\nabla p^{t-1}(F|I) + \nabla p^{t-2}(F|I)]/2$

Graph-cut segmentation with shrinkage constraint

- **Potentials and interaction terms** between neighboring pixels:

$$V_i^t(1) = -ln[p_i^t(F|I)], \quad V_i^t(0) = -ln[p_i^t(B|I)],$$

$$W_{i,j}^t = \delta_{L_i \neq L_j}\, \beta \exp\left[-\frac{(I_i^t - I_j^t)^2}{2\sigma^2}\right],$$

where:

  - $p_i^t(F|I)$ and $p_i^t(B|I)$ are computed from histograms of floe and background reliable regions
  - $\sigma^2 := \mathrm{var}(I^t)$
  - β controls the importance of spatial interaction

- Apply **graph-cut**: minimize

$$E = \sum_{\text{pixels } i} V_i(L_i) + \sum_{i \sim j} W_{i,j}(L_i, L_j) + \infty \sum_t \delta_{L_i^{t+1} > L_i^t}$$

Ice floe segmentation with shrinkage constraint ($\beta = 2$)

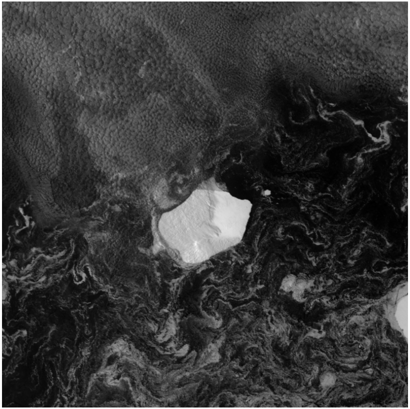

Original MODIS data

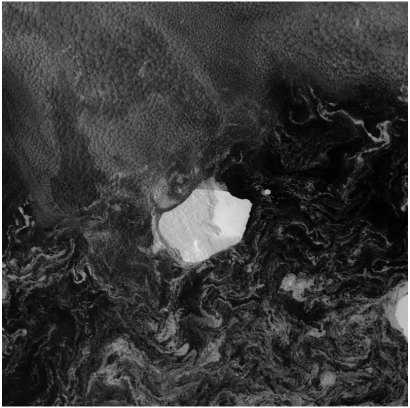

Graph-cut with directed infinite links

Manual segmentation

Dice score (DC) = 0.980 ± 0.007

Comparison with other graph-cut-based methods

- *[w/o]* No temporal links, *i.e.* independent segmentation of each frame
- *[Feedforward]* Foreground pixels of the frame t are marked as seeds with infinite unary costs in the frame $(t+1)$
- *[Bi=const]* Bidirectional temporal links with a constant weight
- *[Bi=variable]* Bidirectional temporal links are computed based on intensity differences between pixels in successive frames [Wolz'10]

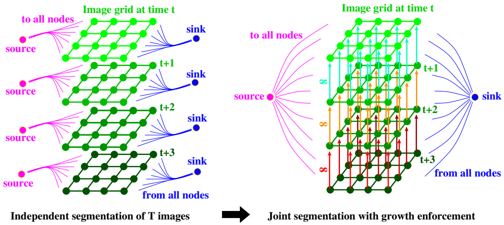

Comparison with other graph-cut-based methods

[w/o]

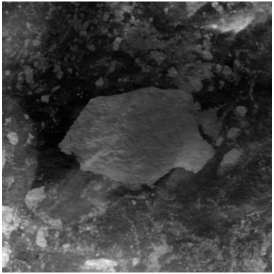

DC = .933 ± .099

[Feedforward]

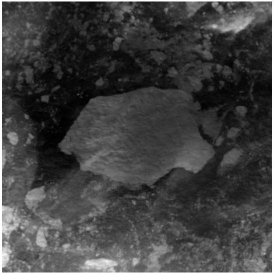

DC = .554 ± .128

[Bi=variable]

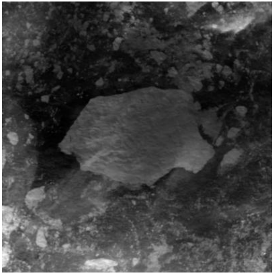

DC = .958 ± .048

- **Conclusion:** These methods are very sensitive to:
  - noise
  - variations of foreground/background intensities

Using temporal links with constant weights

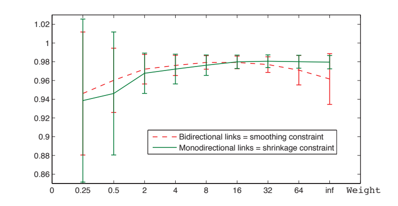

Mean and standard deviation for the dice score as a function of the temporal link's weight, when using mono- and bidirectional temporal links

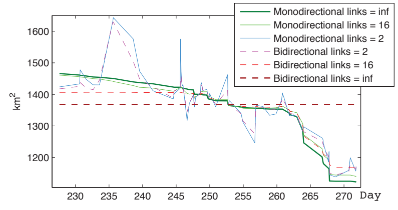

Area of a multiyear ice floe as a function of time, computed by using mono- and bidirectional links with different weights

Conclusions

Advantages of using graph-cut with temporal directed infinite links:

1. Succeeds in segmenting very noisy and low-contrast data
2. Copes well with rapid shrinkage events
3. No parameters needed to quantify temporal coherency

Outline

Automated mapping of burned areas

- **Biomass burning** has a significant impact on a climate system

⇓

- Automated **mapping of burned areas** to:
  - help heal the scars
  - prevent future fires

Objective:

- Segment growing burned areas in time series of images
- By analyzing Terra Moderate Resolution Imaging Spectroradiometer (MODIS) measurements

State of the art

- MODIS Collection 5.1 Direct Broadcast Monthly Burned Area Product (MCD64A1)
  - change detection approach [Giglio 2009]
  - uses MODIS Level 2G (bands 1, 5, 7) and Level 3 daily active fire products
  - spatial filtering within the closest fixed neighborhoods

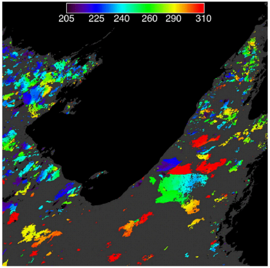

Estimated days of burn, MODIS tile h31v10

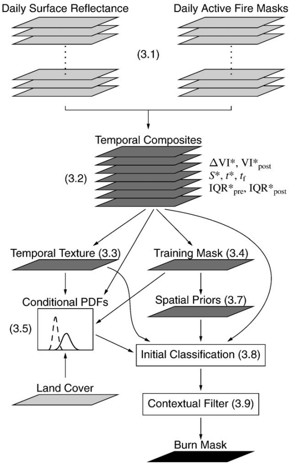

Our objective

- Compute **globally-optimal spatio-temporal** segmentation of **growing burned areas**
- From a time series of very noisy data
  - Cloud contamination, missing data
- Using a new **graph-cut-based** method with **shape growth constraint**

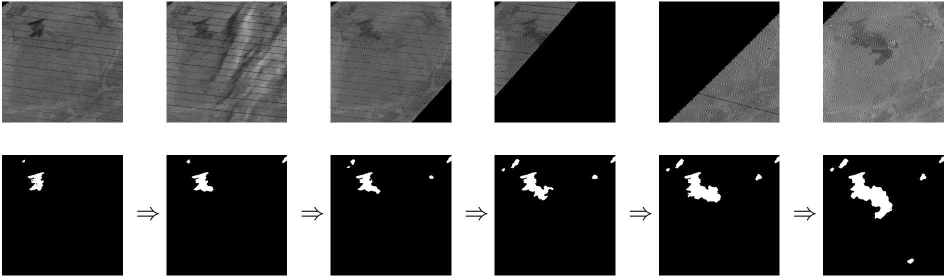

Data set

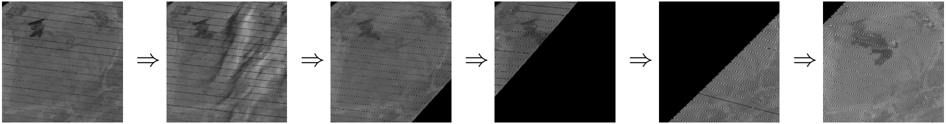

- **Forty days** of **Terra MODIS** Level 2G measurements (MOD09GA)
  - Over tropical savannas in the Northern Australia (tile h31v10)
  - Acquired in September - October 2011 (days 244-283)
  - Band 5 (1.24 μm) 500-m land surface reflectance data
  - $T = 40$ images with spatial dimensions of 400×400 pixels

- **MCD64A1** burned area product
  - Training: computing an initial histogram of burned areas
  - Validation

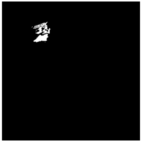

Training mask
(days 213-243)

Proposed spatio-temporal segmentation method

0. **Initialization:**
   - $k := 0$
   - MCD64A1$[t_1 - D, t_1 - 1] \rightarrow$ initial **burned** training mask R_k^B
1. **Unburned** training mask R_k^U = complementary (dilation (R_k^B))

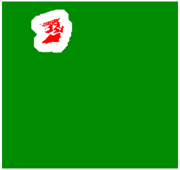

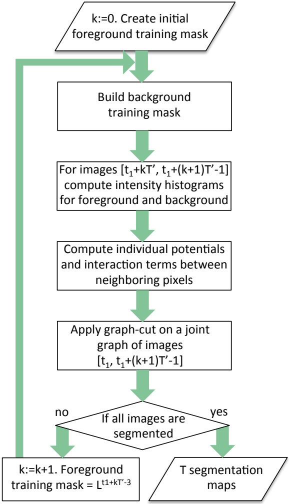

Proposed spatio-temporal segmentation method

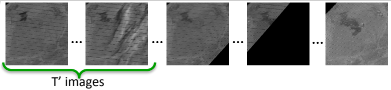

For images $t = [t_1 + kT', t_1 + (k+1)T' - 1]$:

2. Compute intensity **histograms** of MODIS band 5 for burned $p^t(I|B)$ and unburned $p^t(I|U)$ areas
   - using masks R_k^B and R_k^U
3. Compute **individual potentials** and **interaction terms**, assuming $p^t(B) = p^t(U) = 1/2$:

$$V_i^t(1) = -ln[p^t(B|I_i^t)] = -ln\left[\frac{p^t(I_i^t|B)}{p^t(I_i^t|B) + p^t(I_i^t|U)}\right],$$

$$V_i^t(0) = -ln[p^t(U|I_i^t)] = -ln\left[\frac{p^t(I_i^t|U)}{p^t(I_i^t|B) + p^t(I_i^t|U)}\right],$$

$$W_{i,j}^t = \delta_{L_i \neq L_j} \, \beta \, \exp\left[-\frac{(I_i^t - I_j^t)^2}{2\sigma^2}\right],$$

$\sigma^2 := \mathrm{var}(I^t)$.
If I_i^t is missing, $V_i^t(1) = V_i^t(0) = 0$

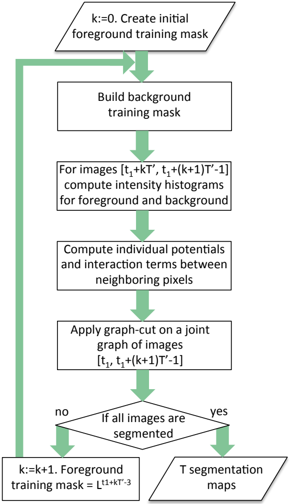

Proposed spatio-temporal segmentation method

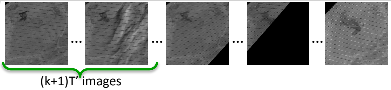

4. Apply graph-cut on a joint graph of images $[t_1, t_1 + (k+1)T' - 1]$

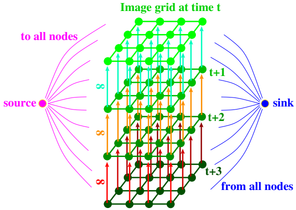

5. If all images are segmented, exit. Otherwise:
   - $k := k + 1$
   - **Burned** training mask $R_k^B = L^{t_1+kT'-3}$
   - Go to step 1 (Consider the next T' images)

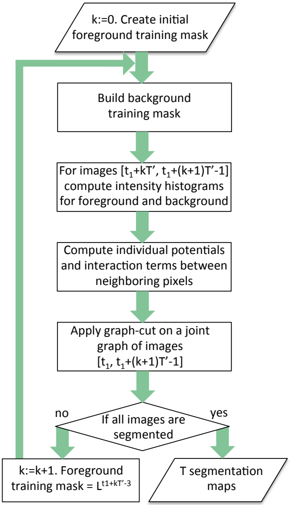

Segmentation results ($\beta = 2$, $T' = 20$)

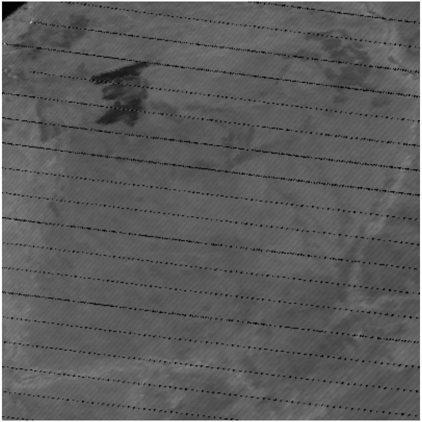

Segmentation results

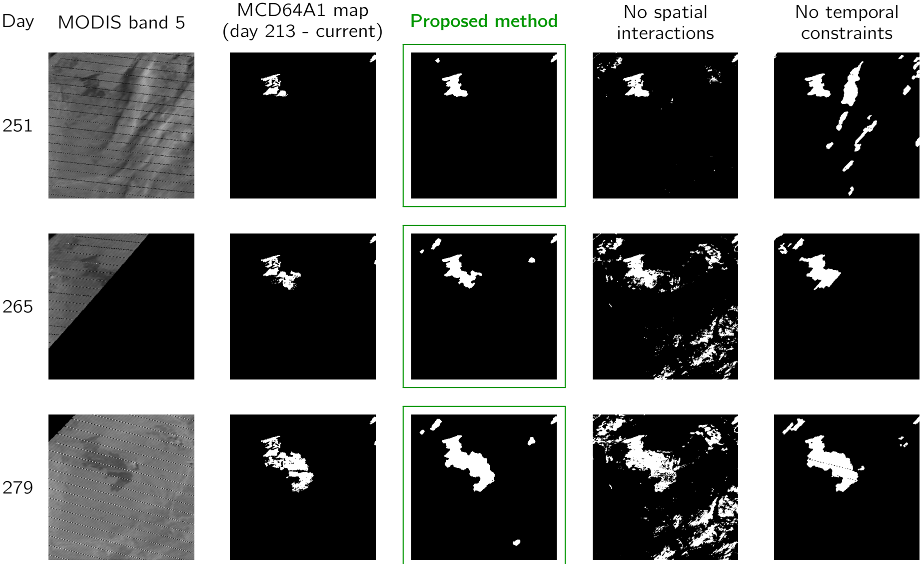

Segmentation results

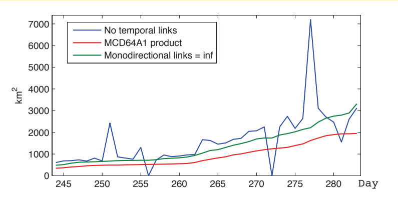

Burned area as a function of time, when using no temporal links, monodirectional infinite links and MCD64A1 product

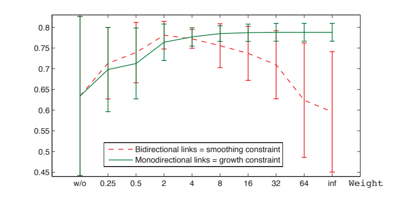

Mean and standard deviation for the dice score (proposed *versus* MCD64A1) as a function of the temporal link's weight, when using mono- and bidirectional temporal links

Segmentation results

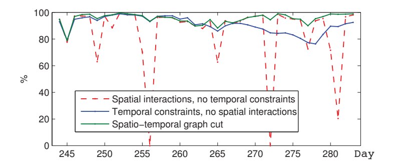

Percentage of pixels identified as burned by the **proposed method** method *AMONG* the pixels identified as burned during [day 244 - current] by MCD64A1

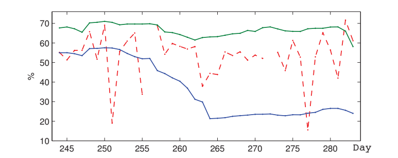

Percentage of pixels identified as burned during [day 213 - current] by MCD64A1 *AMONG* the pixels identified as burned by the **proposed method**

Conclusions and perspectives

Conclusion

- The new method proved to be robust to:
  - noisy and low-contrast images
  - missing data

Perspectives

- Extend the method for segmenting long time series of satellite data

Outline

Task

- **Glioma** is the most frequent primary tumor of the brain
- The tumor is known to grow steadily
- **Objective:** segment lesions from longitudinal sets of multimodal magnetic resonance image (MRI) volumes

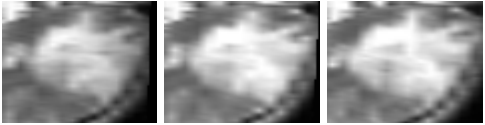

Data description

- **Multimodal image volumes**, each comprising:
  - T1 MRI
  - contrast-enhanced T1 MRI (T1c)
  - T2 MRI
  - T2 FLAIR MRI
- Acquired from **ten patients** initially diagnosed with low grade glioma
- Time series have **3-14 time points**
  - 3-6 months between any two acquisitions
- All image volumes were **rigidly registered**
- **Approximate truth:** three 2D slices intersecting with the tumor center were manually annotated by an expert in every volume

Segmentation with growth constraint

- Segmentations of 3D volumes of each individual data point: **generative model for multimodal brain segmentation** [Menze'10]
  - Models the lesion with a latent atlas class [Riklin-Raviv'10] amending the tissue atlas of the standard EM segmenter [Kapur'96]
- Tumor = Foreground (F), Healthy tissue = Background (B)
  - Changes of the core (visible in T1c) occur within the larger edema regions (visible in T2 or FLAIR)
  - Class transitions: from *healthy* to *edema*, from *edema* to *core*
- **Potential** $V_i^{s,t}(L_i^{s,t})$ of label $L_i^{s,t}$ at voxel i, time point t, and imaging sequence s:

$$V_i^{s,t}(0) = p_{s,t}(F|I^{s,t}),$$

$$V_i^{s,t}(1) = p_{s,t}(B|I^{s,t}) = 1 - p_{s,t}(F|I^{s,t}).$$

We identified tumor subclasses with $p(F|I^{s=T1,t})$ for *core*, and $p(F|I^{s=T2,t})$ with *edema*

Segmentation with growth constraint

- **Potential** $V_i^{s,t}(L_i^{s,t})$ of label $L_i^{s,t}$ at voxel i, time point t, and imaging sequence s:

$$V_i^{s,t}(0) = p_{s,t}(F|I^{s,t}),$$
$$V_i^{s,t}(1) = p_{s,t}(B|I^{s,t}) = 1 - p_{s,t}(F|I^{s,t}).$$

- **3D spatial constraints** in a 26-neighborhood:

$$W_{i,j}^{s,t}(L_i^{s,t}, L_j^{s,t}) = \delta_{L_i \neq L_j}\, \beta\, \frac{\alpha(i,j)}{\alpha_{\text{tot}}}\, \exp\left(-\left(\frac{I^{s,t}(i) - I^{s,t}(j)}{A}\right)^2\right)$$

with $\beta = 0.5$, $\alpha(p,q) = \frac{1}{\text{distance}(p,q)}$, $\alpha_{\text{tot}} = \sum_{q \in \mathcal{N}(\text{pixel } p)} \alpha(p,q)$ and $A = \frac{1}{3}\left(\max I^{s,t} - \min I^{s,t}\right)$

- **Inter-sequence inclusion constraints:**
  - Foregrounds in T1 and T1c modalities are included in the one of T2
  - Foreground in T2 is included in the one of FLAIR

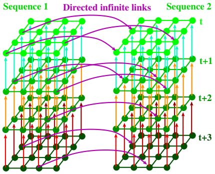

Segmentation results (14 observations)

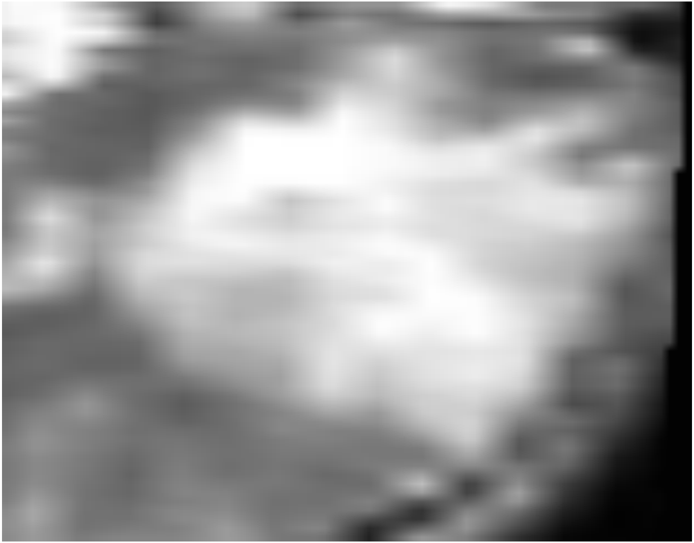

Volume-time plot for a patient with 14 observations

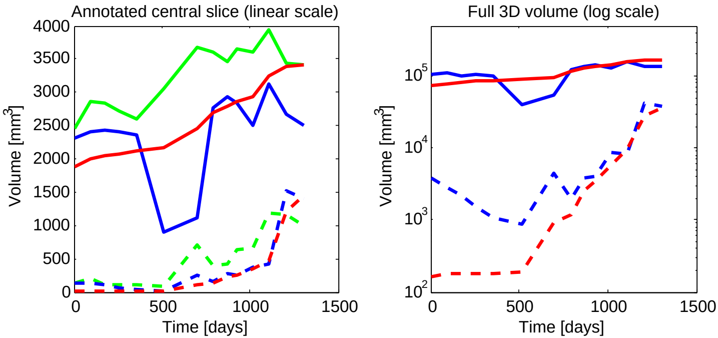

- Solid lines: *edema*
- Dashed lines: *core*

(Red) Proposed segmentation with growth constraint
(Blue) Initial multimodal segmentation [Menze'10]
(Green) Manual segmentation

Two time series of T2 and FLAIR MR image volumes

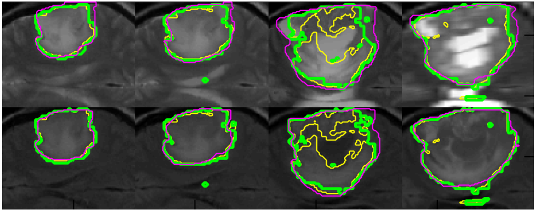

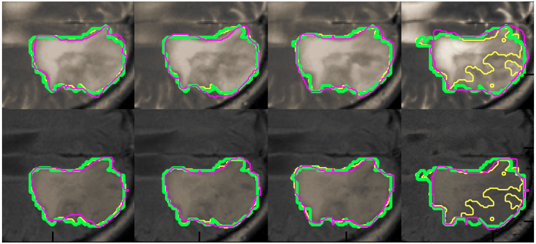

(Green) Proposed segmentation with growth constraint
(Yellow) Initial multimodal segmentation [Menze'10]
(Magenta) Manual evaluation

The tumor grows rapidly between the second and forth scene.

Intensity modifications in the last scene lead to a suboptimal performance of the initial multimodal segmentation

Using temporal links with constant weights

Changes* in the average segmentation performance of the ten image sequences when testing different regularization approaches

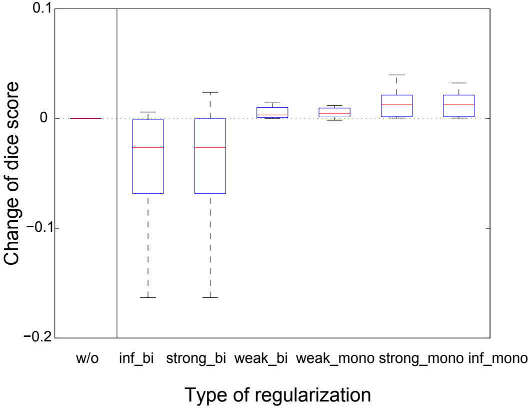

*The box indicates quartiles, the whiskers indicate outliers

Conclusions and perspectives

Conclusions

- The main contribution:
  1. a new framework for segmentation of 2D/3D image time series with the constraint of shape growth/shrinkage,
  2. in order to be able to segment very noisy/low-contrast/incomplete data,
  3. in a very low computational time.
- The new method:
  - proved to be robust to important noise and low-contrast
  - linear complexity in practice

Future works

- Other applications, such as organ development

Thank you for your attention!

Spatio-temporal video segmentation with shape growth or shrinkage constraint

Yuliya Tarabalka[1], Guillaume Charpiat[2],
Ludovic Brucker[3] and Bjoern H. Menze[4]

[1,2]Inria Sophia-Antipolis, AYIN[1] and STARS[2] teams, France
[3]NASA GSFC & Universities Space Research Association, MD, USA
[4]Computer Vision Laboratory, ETH Zurich, Switzerland & Inria Sophia-Antipolis, Asclepios team
e-mail: yuliya.tarabalka@inria.fr

February 5, 2014

INVENTORS FOR THE DIGITAL WORLD

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich