Marked Point Process Models in Image Processing: Application to Remote Sensing

J. Zerubia

Joint work with P. Craciun, S. Descamps, X. Descombes, C. Lacoste, F. Lafarge, M. Ortner, G. Perrin, R. Stoica

Ariana research group, http://www-sop.inria.fr/ariana
Ayin research group, https://team.inria.fr/ayin

Bayesian Approach

$$\boxed{P(X \mid Y) = \frac{P(X)P(Y \mid X)}{P(Y)} \propto P(X)P(Y \mid X)}$$

Y : observed data
X : unknown variable (objects, features, …)

$P(Y \mid X)$: likelihood

$P(X)$: prior

$P(X \mid Y)$: posterior

$$\boxed{\text{Estimated } X : X^* = \arg\max_X P(X \mid Y)}$$

Medium Resolution Data: Markov Random Fields

Prior: Markov Random Field on pixel values

Likelihood: conditional independence assumption

$$P(Y \mid X) = \prod_{s \in S} P(y_s \mid x_s)$$

No contextual information in the likelihood:

1 - uncorrelated noise

2 - no texture

Markov Random Fields

$$P(x_s \mid x_t, t \neq s) = P(x_s \mid x_t, t \in v_s)$$

v_s being the neighborhood of s

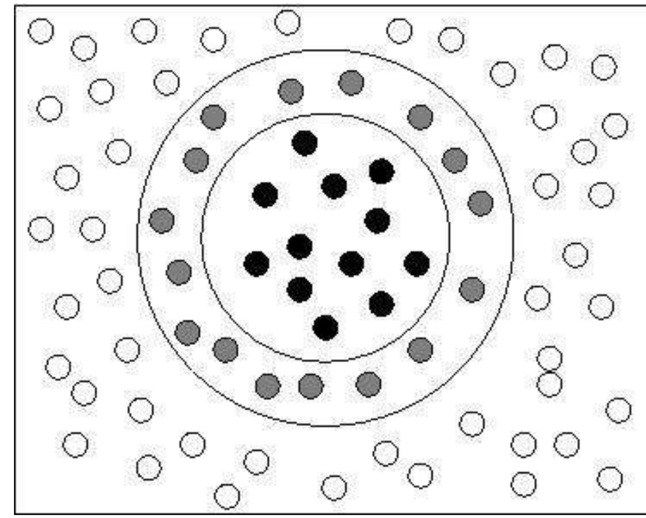

- Contextual Information Modeling
- Link with Statistical Physics: Gibbs Fields

From Context to Geometry

SPOT image © CNES

IKONOS image © Satellite imaging Corporation

IKONOS image © Satellite image Corporation

From Context to Geometry

How to extract structural information from HR images?

SPOT image © CNES

aerial image © IGN

High Resolution Data: From Pixels to Objects

- **Goal:** To model the observed scene as a configuration of objects (roads, rivers, buildings, trees):
  - To take into account data at a macroscopic scale.
  - To take into account the geometry of objects.
  - To take into account relations between objects (macro-texture).
  - Unknown number of objects (MRF on graph impossible).

Solution: Marked point processes

- **Stochastic modeling**: Set of objects in the scene $\equiv$ realization of a **marked point process**, **X**. [Lieshout-00], [Descombes and Zerubia-02]
- **Optimization algorithm:** Monte Carlo sampler (e.g. RJMCMC) + simulated annealing.

Marked Point Processes

- A marked point process X on $\chi = \mathcal{P} \times \mathcal{M}$ is a point process on χ for which the point location is in $\mathcal{P}$ and the marks in $\mathcal{M}$.

- We define X by its probability density f w.r.t. the law $\pi_\nu(.)$ of a Poisson process known as the reference process ($\nu(.)$ is the intensity measure)

Sampling: Birth and Death Algorithm (Geyer/Moller-94)

- **Birth**: with probability ½, randomly propose a new point u in χ to be added to the current configuration x. Let y = x U {u}. Compute the acceptance ratio:

$$R_1(x,y) = \frac{f(y)}{f(x)} \frac{\nu(\chi)}{n(y)}$$

- **Death**: with probability ½, randomly propose a point v to be removed from x. Let y = x /{v}. Compute the acceptance ratio:

$$R_2(x,y) = \frac{f(y)}{f(x)} \frac{n(x)}{\nu(\chi)}$$

- With probability α_i = min{1, R_i}, accept the proposition x_{t+1} = y, otherwise accept the proposition x_{t+1}=x.

Sampling: RJMCMC (Green-95)

- Generalization of Geyer/Moller-94
- Mixture of several proposition kernels:

$$Q(x,.) = \sum_{m} p_m(x) q_m(x,.)$$

- Convergence condition exists.

Sampling: RJMCMC

- **Algorithm**:

*At time **t**:*

1) Select randomly a kernel $\mathbf{q_m}$ *using the discrete law* $\mathbf{(p_m(x))}$

2) *Generate a new configuration **y** with respect to the selected kernel:*
$\mathbf{y} \sim \mathbf{q_m}(\mathbf{x},.)$

3) Compute the acceptance ratio: $\mathbf{R_m}(\mathbf{x},\mathbf{y})$

4) Compute the acceptance rate $\boldsymbol{\alpha}$ *:* $\alpha = \min(\mathbf{1},\mathbf{R_m}(\mathbf{x},\mathbf{y}))$

5) With probability

- $\boldsymbol{\alpha}$ *set:* $\mathbf{X_{t+1}} = \mathbf{y}$
- ***(1-α)*** *set:* $\mathbf{X_{t+1}} = \mathbf{x}$

Optimization Algorithm

- **Goal:** To estimate a configuration maximizing **f(.)**

- **Simulated annealing:**

> *Successive simulations of* $\mathbf{f_t(x)\ n(dx)}$ *using a RJMCMC algorithm with:*
>
> $$f_t(x) = f(x)^{\frac{1}{T_t}}$$
>
> *where* $\mathbf{(T_t)}$ *(≡ temperature) decreases towards zero.*

- Logarithmic decrease ⇒ global maximum.

- **In practice:** geometric decrease.

  At each step, $\mathbf{T_{t+1} = T_t \times c}$, where **c** is a constant close to 1.
  (c=0.99999 or c=0.999999 depending on the difficulty of the detection)

First example: Quality Candy Model for road network extraction

- Objects: segments [Stoica et al. - 04, Lacoste et al. - 05]
- Prior: models the connectivity and the curvature
- Data term

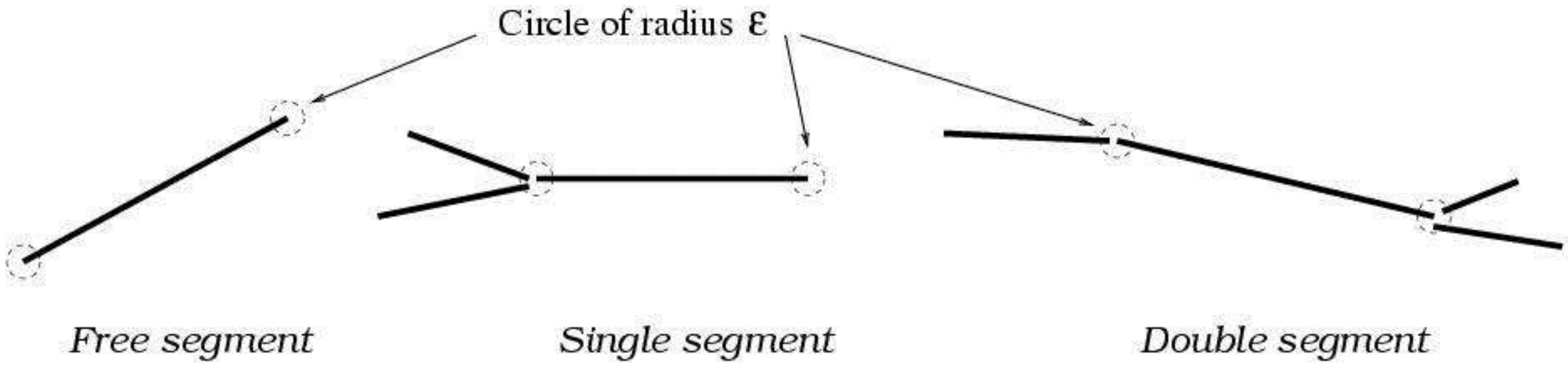

PhDs: R. Stoïca, C. Lacoste in collaboration with IGN and BRGM

First example: Quality Candy Model for road network extraction

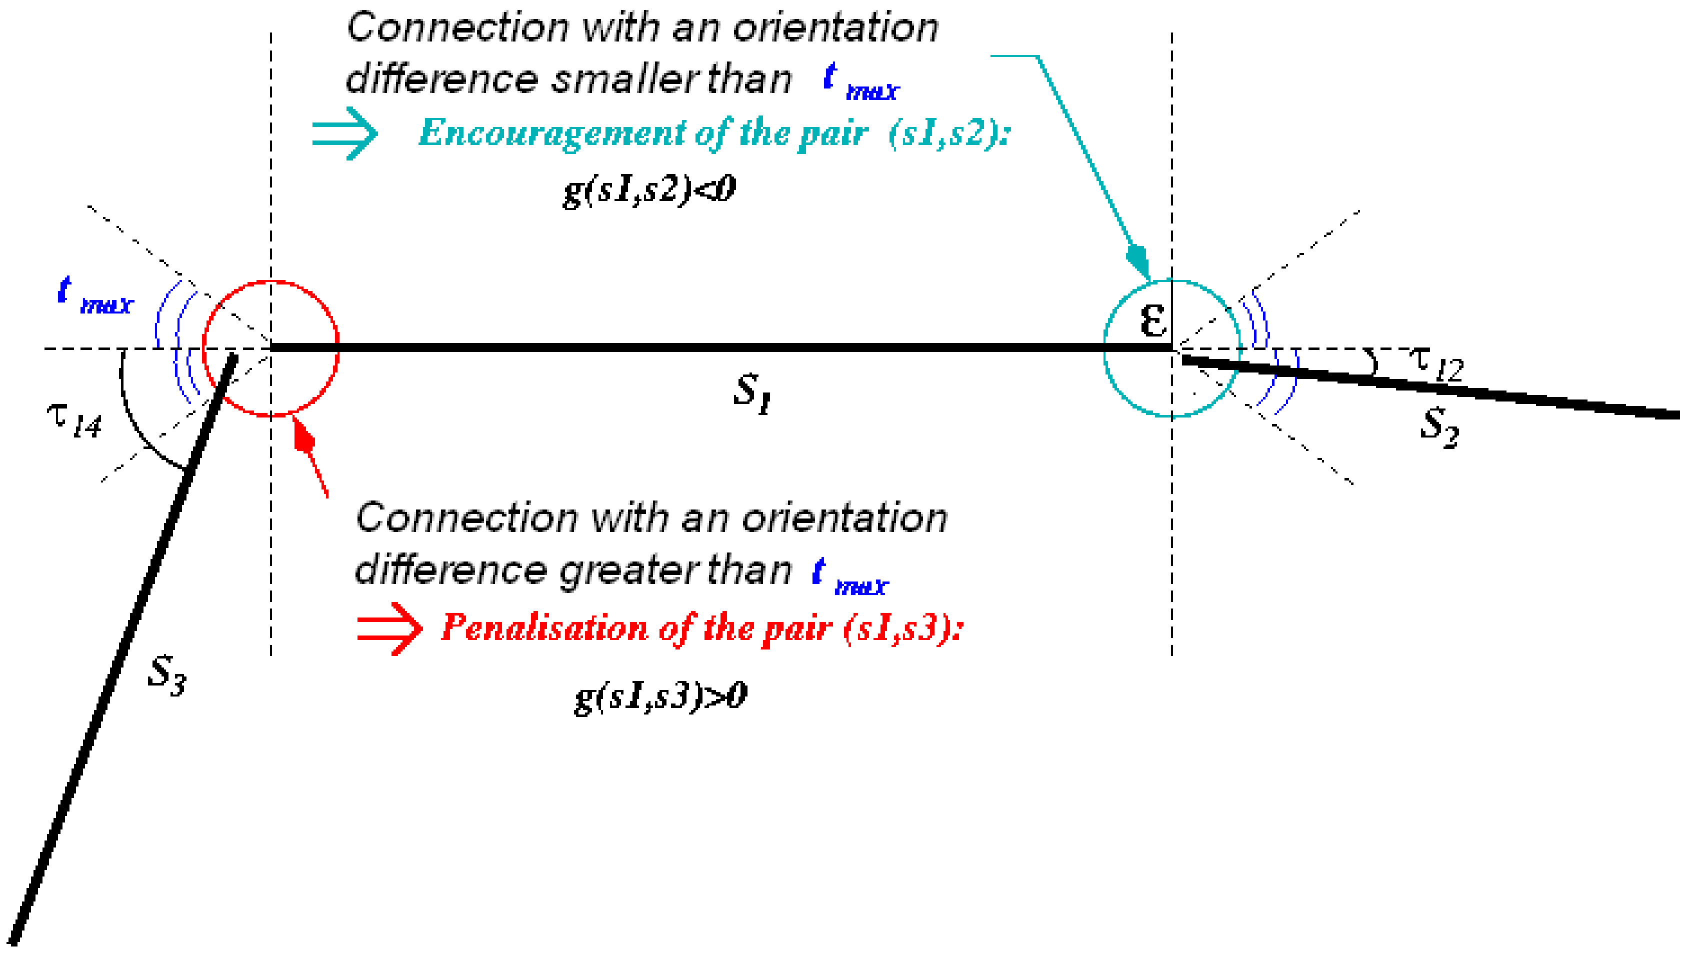

PhDs: R. Stoïca, C. Lacoste in collaboration with IGN and BRGM

First example: Quality Candy Model for road network extraction

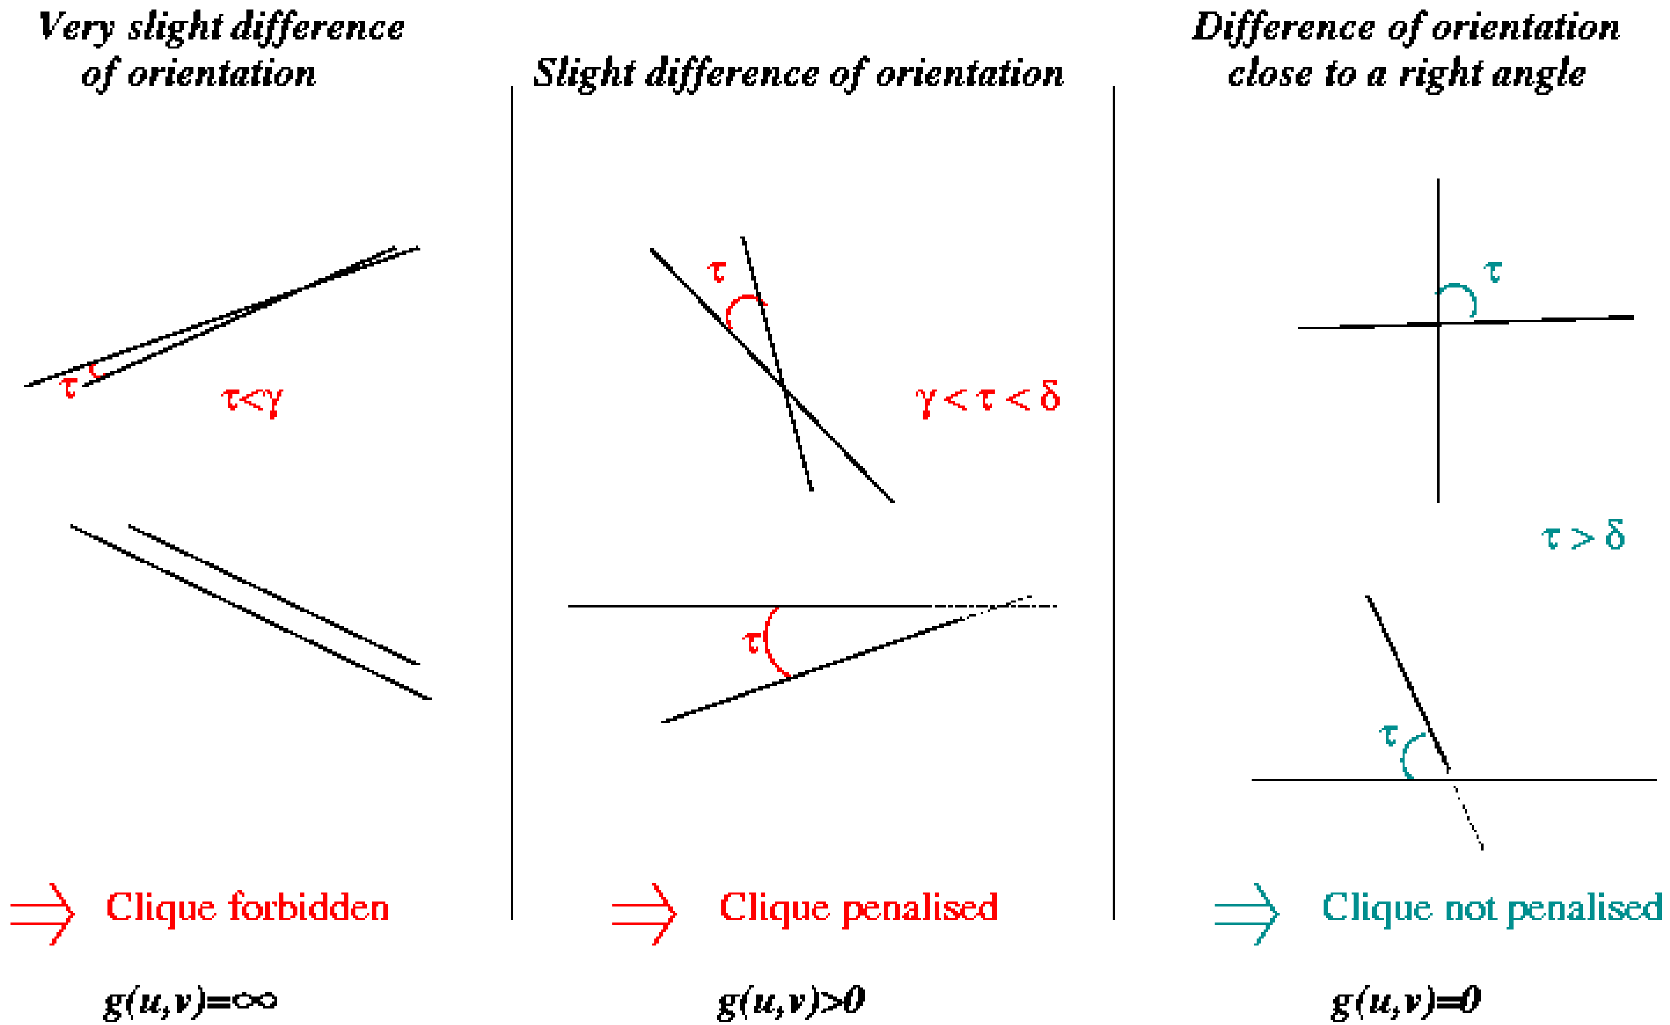

PhDs: R. Stoïca, C. Lacoste in collaboration with IGN and BRGM

First example: Quality Candy Model for road network extraction

- Objects: Segments
- Prior: models the connectivity and the curvature
- First data term: t-test

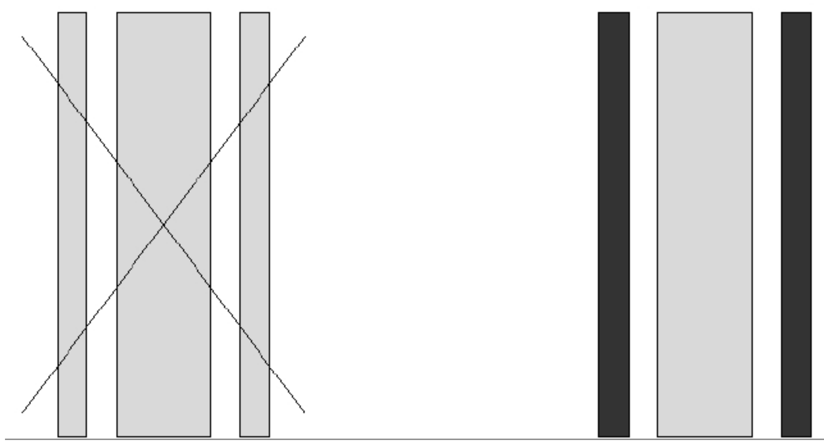

PhDs: R. Stoïca, C. Lacoste in collaboration with IGN and BRGM

First example: Quality Candy Model for road network extraction

- Objects: Segments
- Prior: models the connectivity and the curvature
- Second data term: t-test

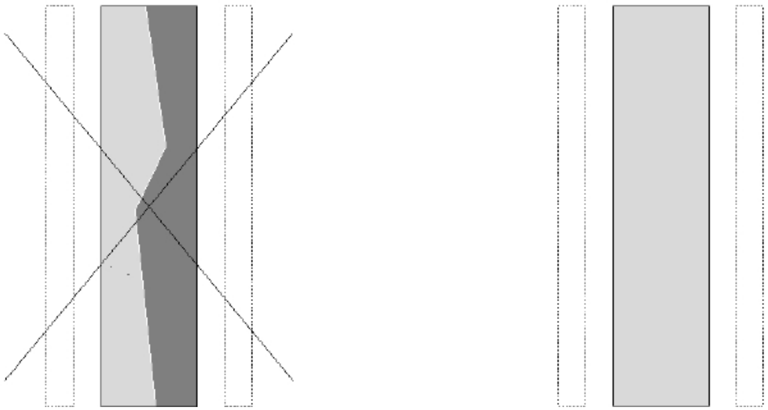

PhDs: R. Stoïca, C. Lacoste in collaboration with IGN and BRGM

Kernels of the RJMCMC algorithm

- Uniform birth and death
- Birth and death in a neighborhood
- Extension/contraction of a segment
- Translation of a segment
- Rotation of a segment

PhDs: R. Stoica, C. Lacoste

Results

PhDs: R. Stoica, C. Lacoste

© IGN

Results

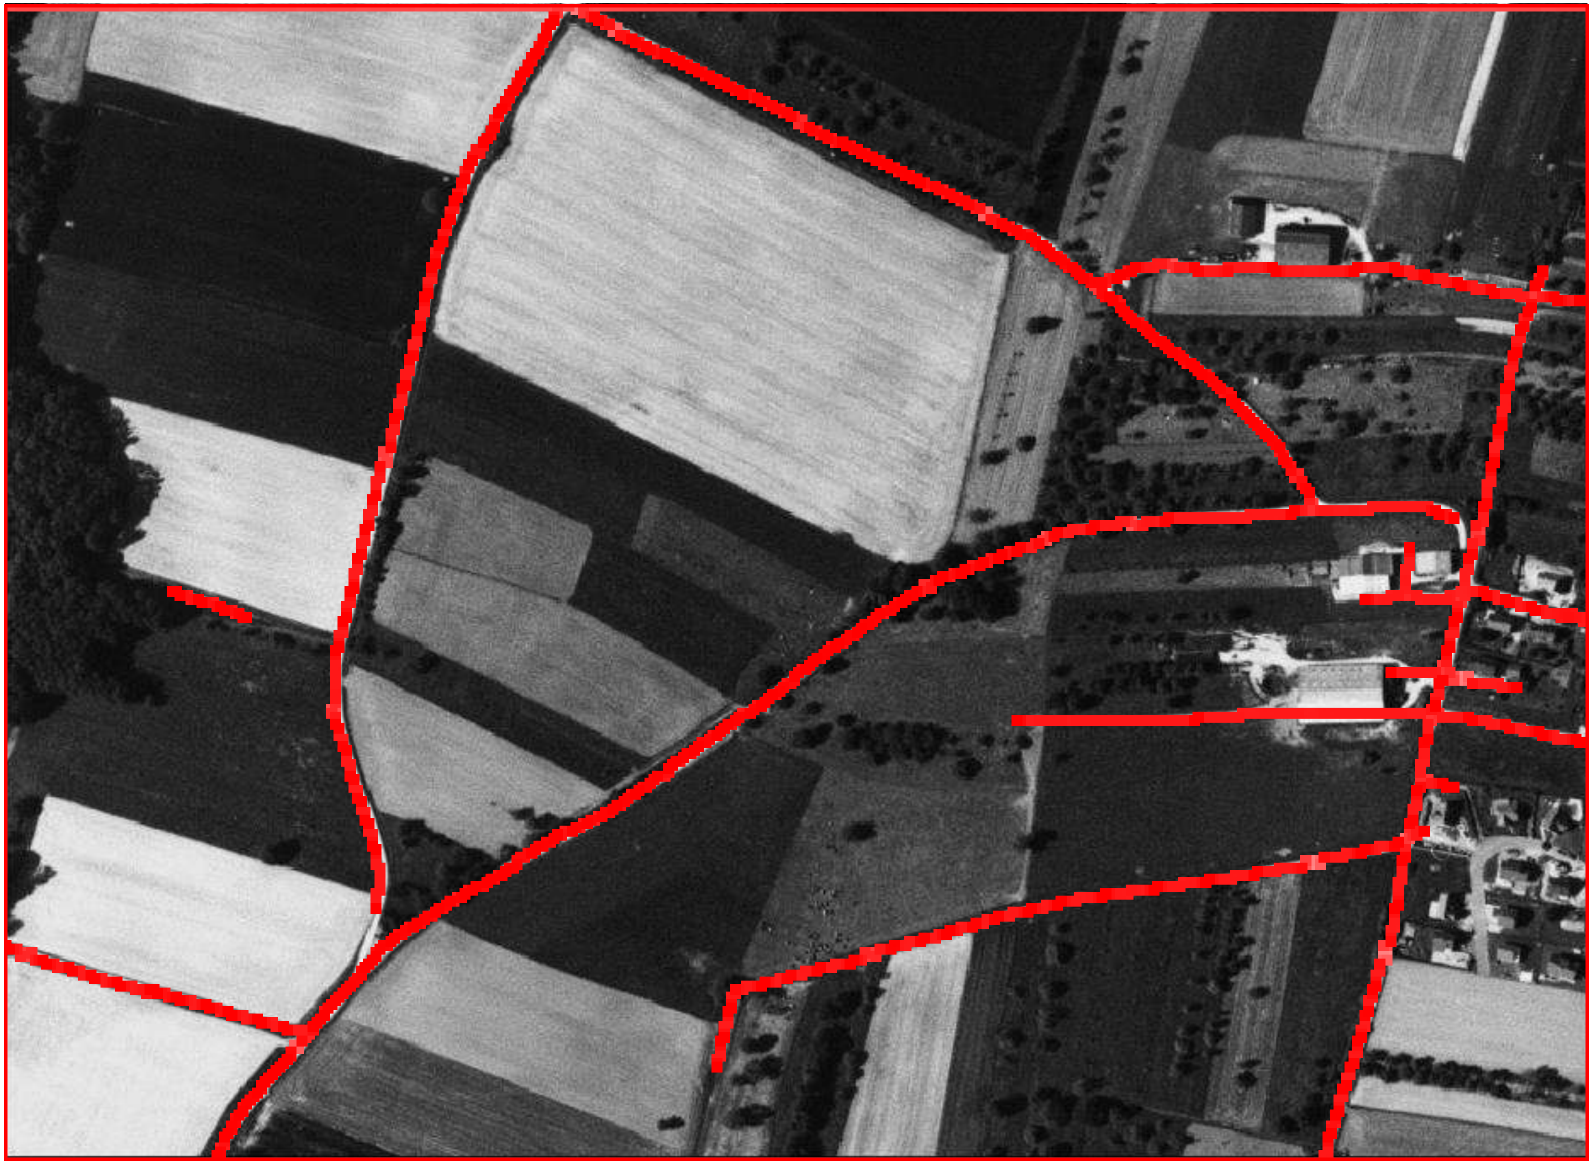

PhDs: R. Stoica, C. Lacoste

© Ariana / INRIA

Results

© ESA

PhDs: R. Stoica, C. Lacoste

Results

PhDs: R. Stoica, C. Lacoste

© Ariana / INRIA

Second example: tree crown extraction
First method

- Object: disc
- Prior: non-overlapping trees
- Data: Gaussian likelihood

$$A_y(S(x)) = \prod_{p \in S(x)} p_{tree}(y_p) \prod_{p \notin S(x)} p_{notree}(y_p)$$

PhD: G. Perrin in collaboration with ECP

Result

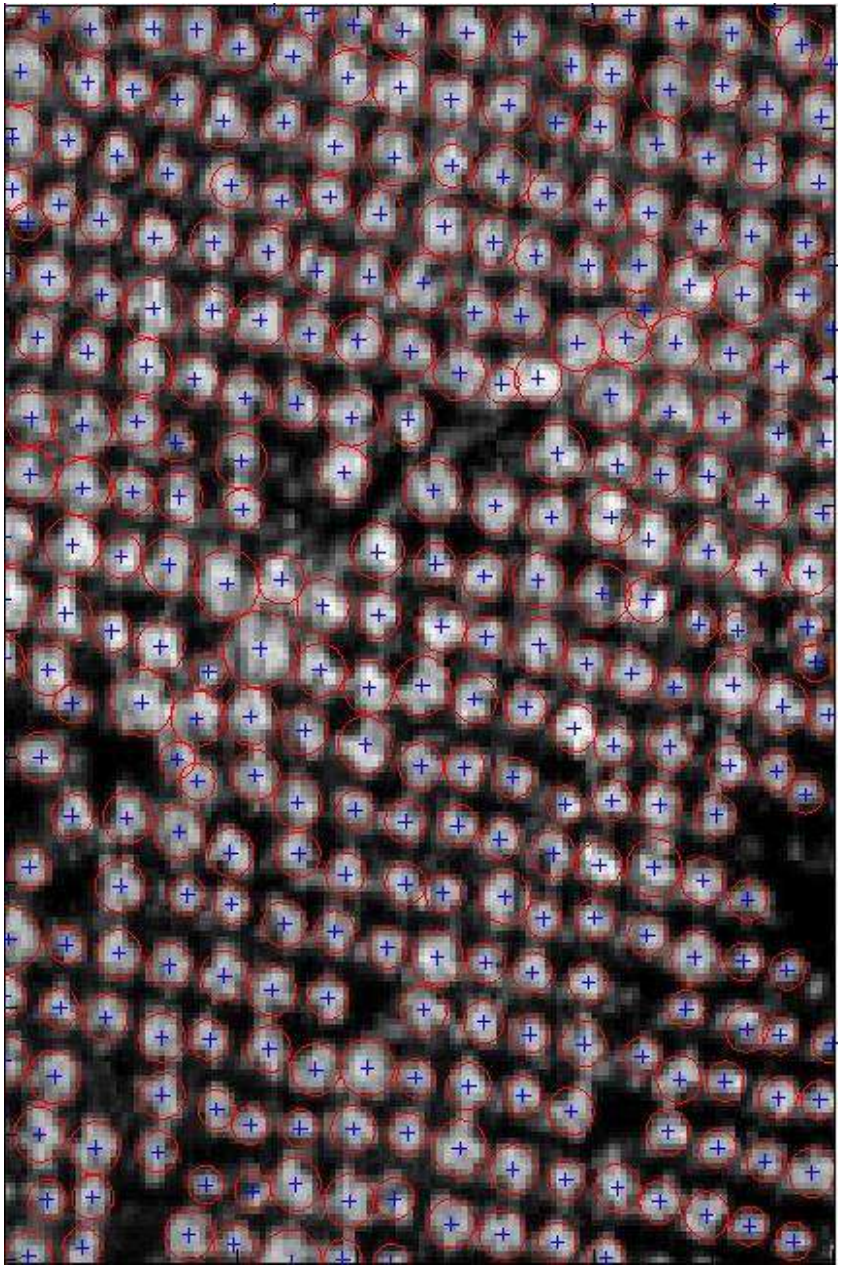

PhD: G. Perrin in collaboration with ECP

Second example: tree crown extraction
Second method

- Marks: ellipses or ellipsoids. [Perrin et al. - 05, Eriksson et al. - 06]

$X = \{X_1, X_2, \ldots, X_n\}$

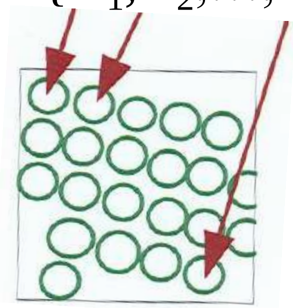

ELLIPSES: 2D model

Dense area: plantation (merged shadows)

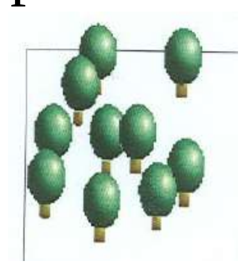

ELLIPSOIDS: 3D model

Sparse vegetation (drop shadows)

PhD: G. Perrin in collaboration with ECP

Density of the process

- Goal: design the density of the MPP in order to make tree configurations be the most likely configurations.
- Minimize the energy: $U(x): \mathrm{f}(x) = \frac{1}{Z}\exp\left(-U(x)\right)$
- Mathematical tools: RJMCMC algorithms + simulated annealing.

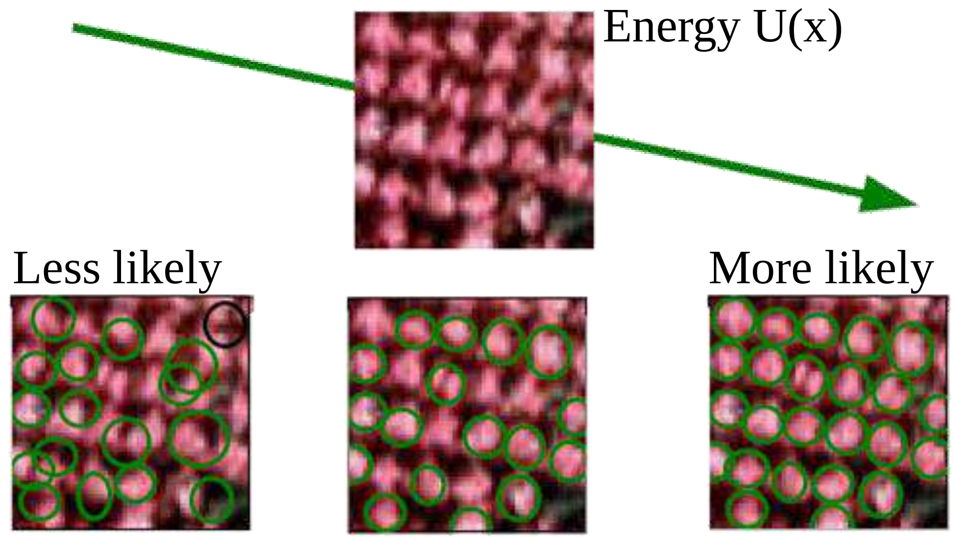

Poplars to be extracted with ellipses

PhD: G. Perrin in collaboration with ECP

Energy of the model

- Regularizing term + data term:

$$U(x) = U_r(x) + U_d(x)$$

- $U_r(x)$: prior term = interactions between objects.

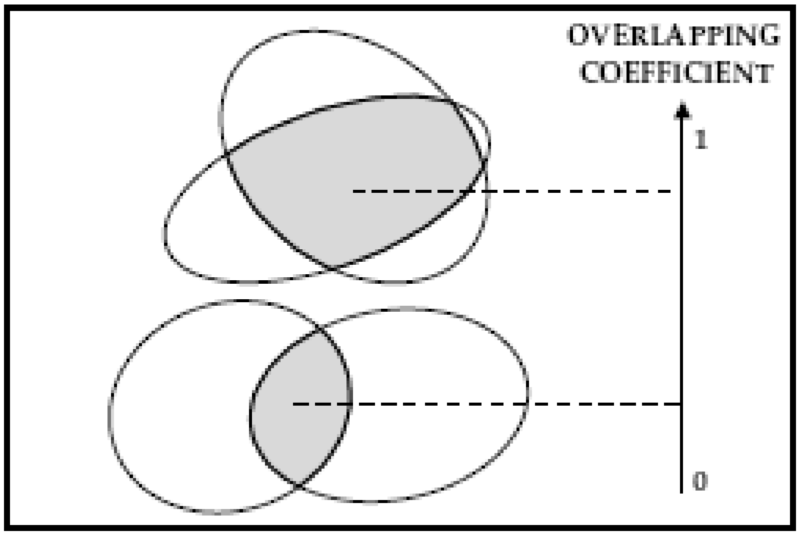

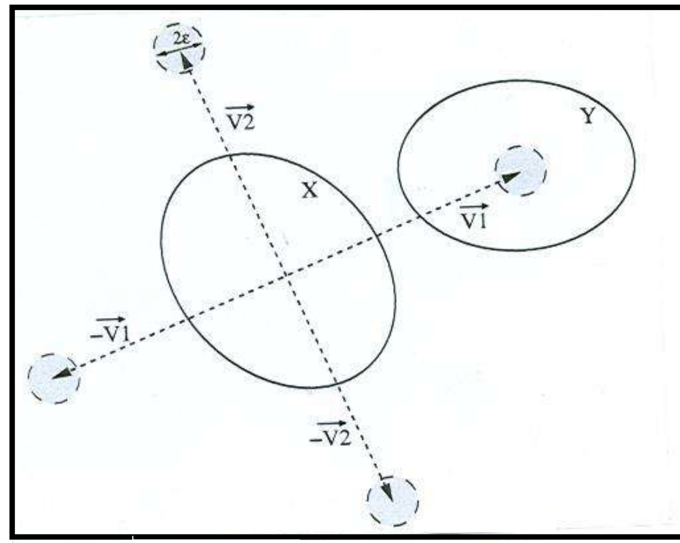

- $U_d(x)$: data term = fitting the object into the image.

$$U_d(x) = \gamma_d \sum_{x_i \in x} U_d(x_i)$$

Data energy term $U_d(x)$

- What is typical of the presence of a tree ?
  - high reflectance in the near infrared.
  - shadow.
  - neighborhood.
- In dense vegetation: merged shadows, shadow area = all around the tree.
- In sparse vegetation: drop shadows, shadow area = in the direction of the sunlight.

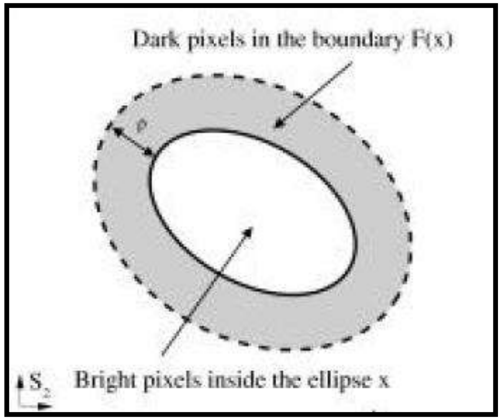

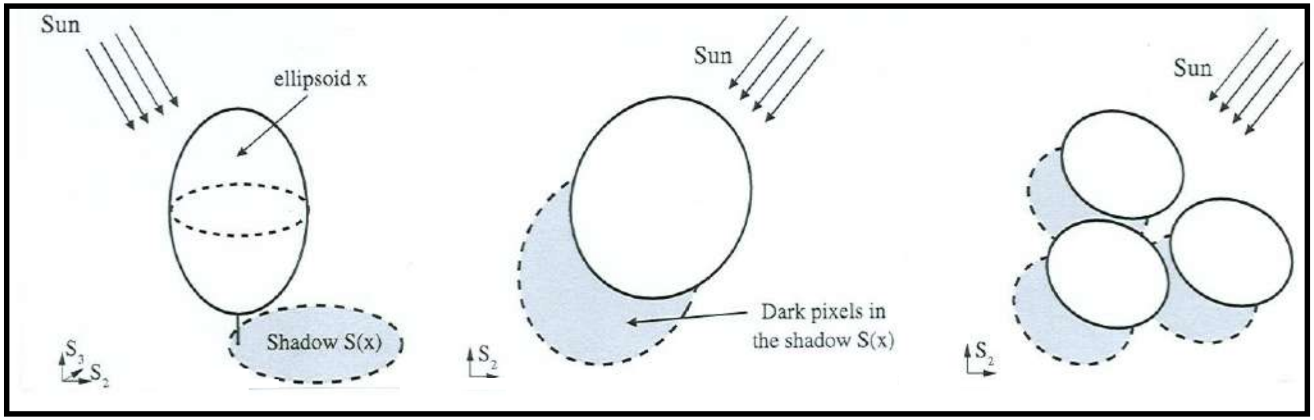

PhD: G. Perrin in collaboration with ECP

Results with the 2D model (1)

Poplar plantation. 1 ha ©IFN.

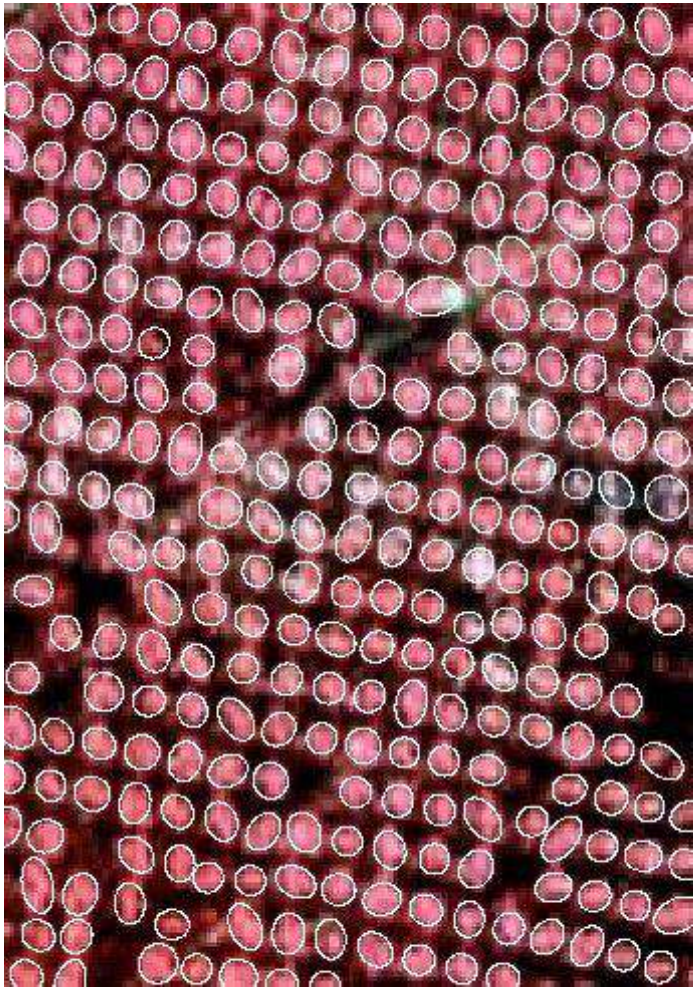

2D model extraction. © Ariana / INRIA

PhD: G. Perrin in collaboration with ECP

Results with the 2D model (2)

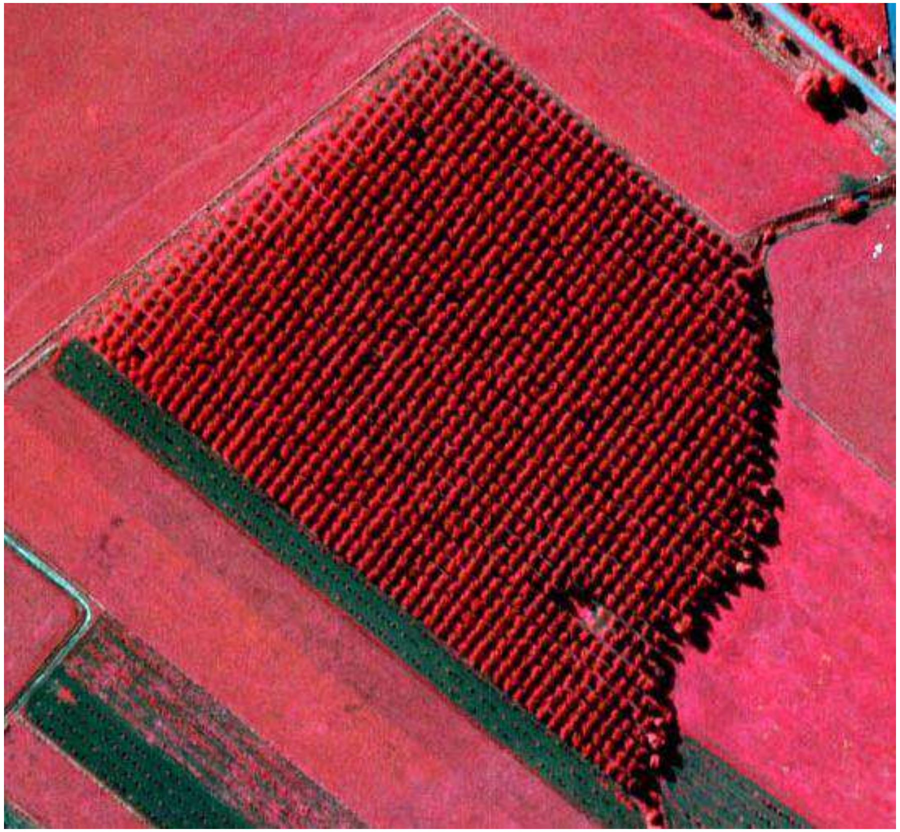

Poplar plantation. 7 ha ©IFN

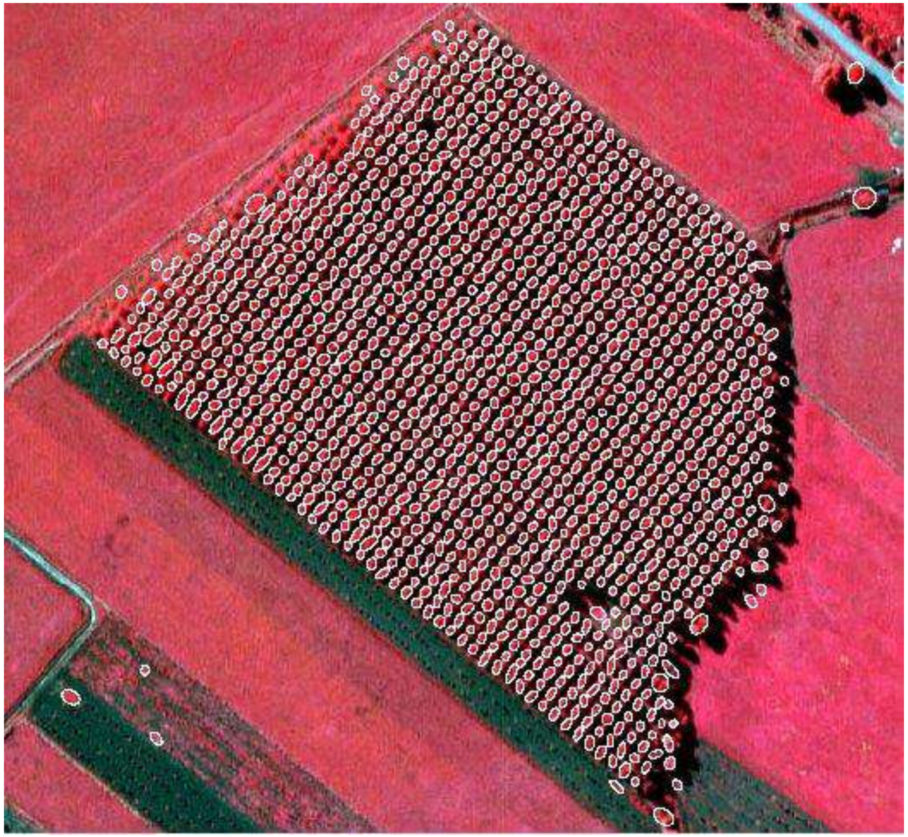

2D model: more than 1300 objects. © Ariana / INRIA

PhD: G. Perrin in collaboration with ECP

Results with the 3D model (1)

• Application: sparse vegetation, trees on the borders of plantations, mixed height stands.

• Hypotheses: the position of the Sun is given, trees close to the nadir and at ground level (no deformation).

• Results: position, crown diameter, approximate height of the tree.

© IFN

© IFN

© IFN

PhD: G. Perrin in collaboration with ECP

Results with the 3D model (2)

- 3D model extraction in sparse vegetation.

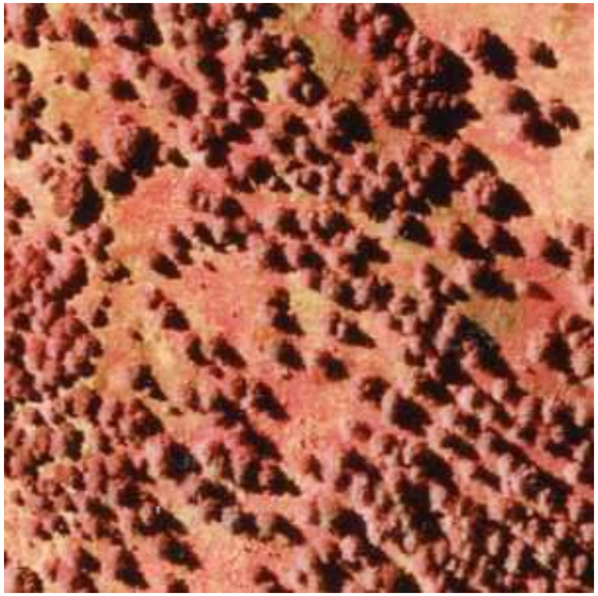

2.5 ha (Alpes Maritimes) © IFN.

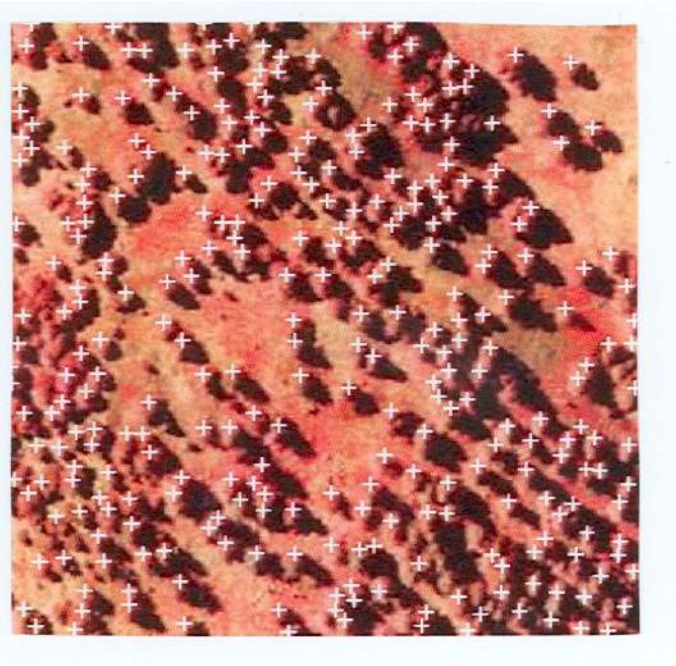

3D model extraction © Ariana / INRIA

PhD: G. Perrin in collaboration with ECP

Results with the 3D model (3)

- Application: density of the sparse vegetation $\approx$ 19%.

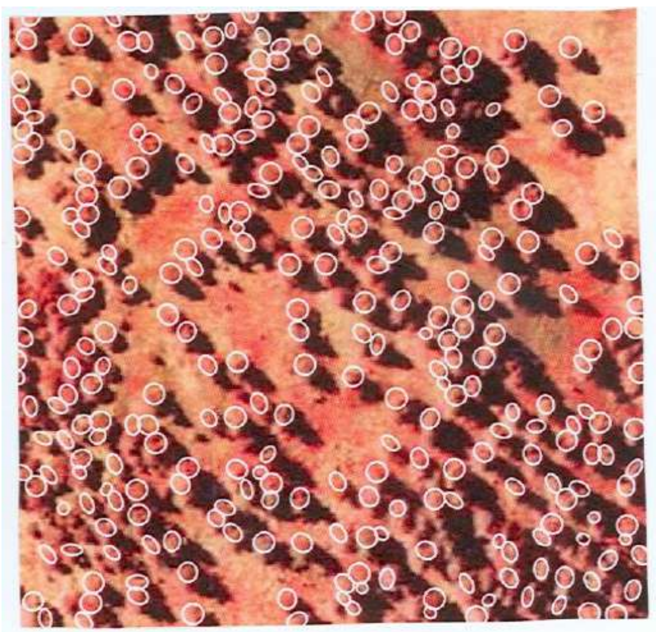

3D model extraction. © Ariana / INRIA

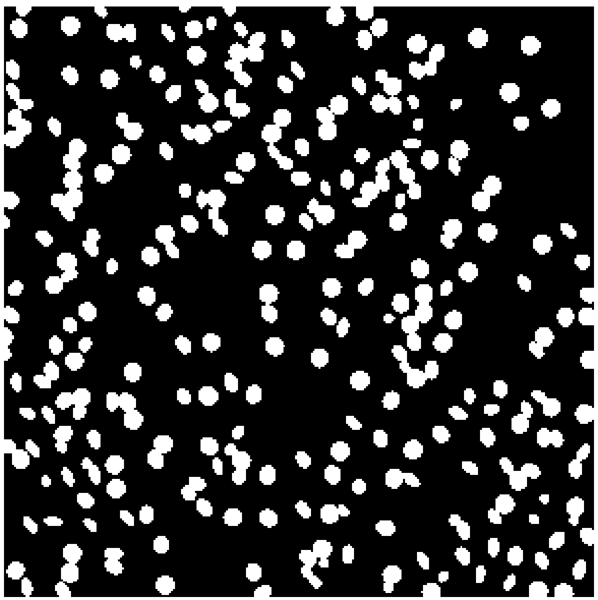

Binary image of the vegetation.

PhD: G. Perrin in collaboration with ECP

Results with the 3D model (4)

- Many objects.
- Information on the timber forest density $\approx$ 15%.

Mixed height stand (3 ha) © IFN.

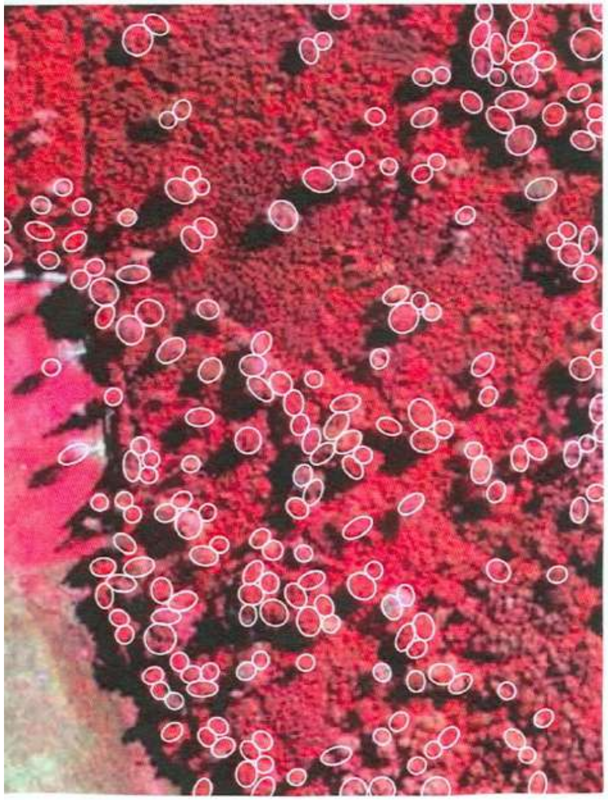

3D model extraction © Ariana / INRIA

PhD: G. Perrin in collaboration with ECP

Third example: Flamingo counting

Previous counting techniques of Flamingos :

✓ Manually, by sampling, counting then extrapolating

✓ Tricky because of the low quality of the aerial images

✓ Time consuming and finally not precise

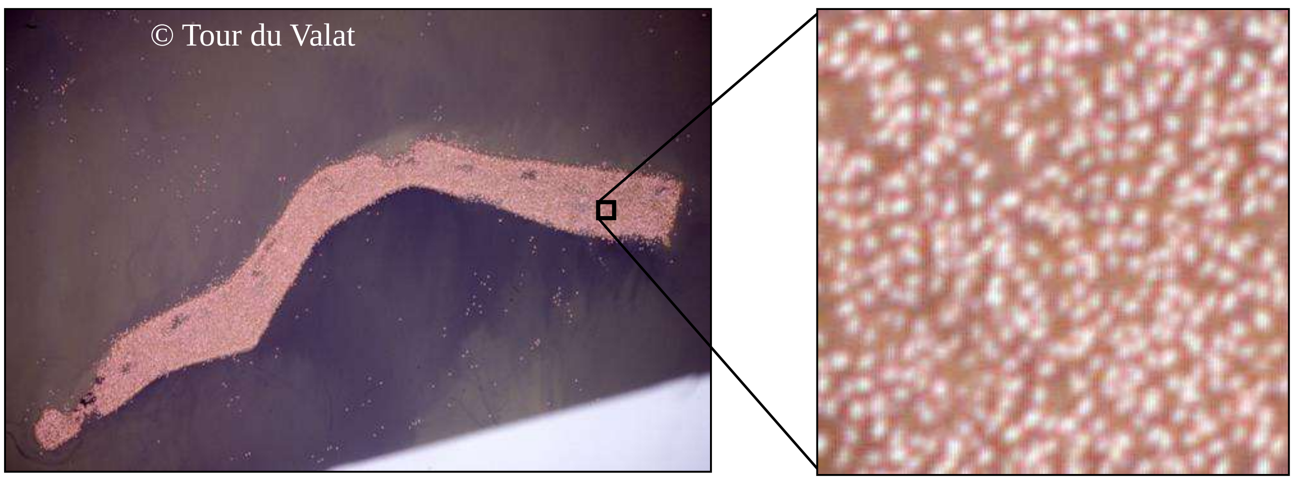

Need to develop a method to count flamingos automatically

MSc: S. Descamps in collaboration with Tour du Valat

Model for the extraction of Flamingos

***A priori* model** : **Interaction between objects [G.Perrin et al., 06]**

→ Penalization of overlapped objects

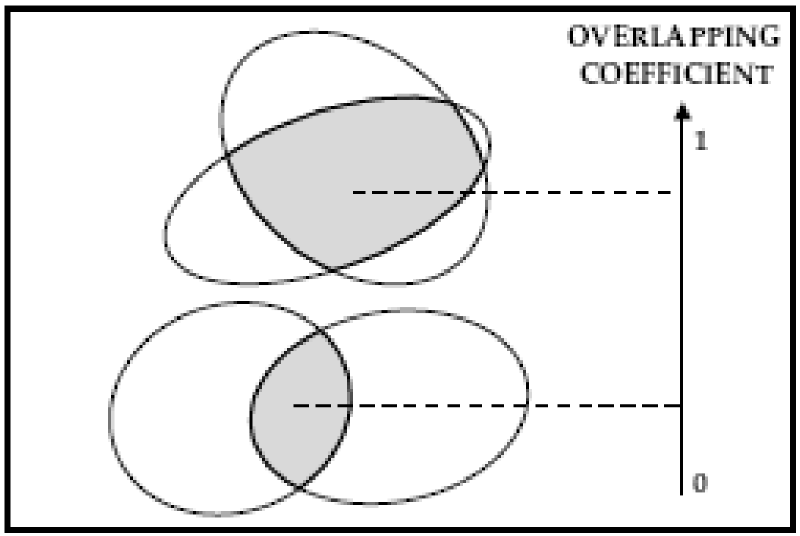

Top: high energy, Bottom: low energy

Model for the extraction of Flamingos

***Data* model : to adapt objects to flamingos**

Flamingos considered as bright ellipses making a contrast with their crowns [Descamps et al., 08]

✓ Modified Bhattacharya distance computation from the pixel distributions in the ellipse and in its crown [G.Perrin et al., 06].

✓ Comparison of the center of the ellipse with the mean value of a flamingo in the image.

We favor good objects, we penalize bad ones:

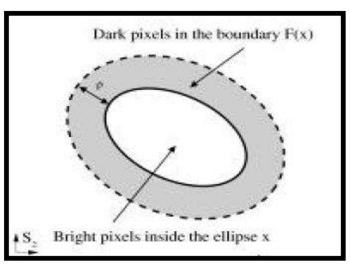

Ellipse and its crown

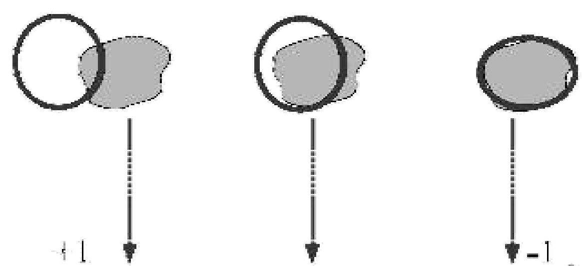

Different levels of energy

MSc: S. Descamps in collaboration with Tour du Valat

Simulation and optimization

Aim: Simulate configurations to reach the minimum of the energy $U(\mathbf{x})$ with the most likely configuration

Method: Multiple birth and death dynamics [X. Descombes et al. 09]

Results

<u>Size estimation of a colony in Camargue, France</u>:

✓ 557 flamingos detected

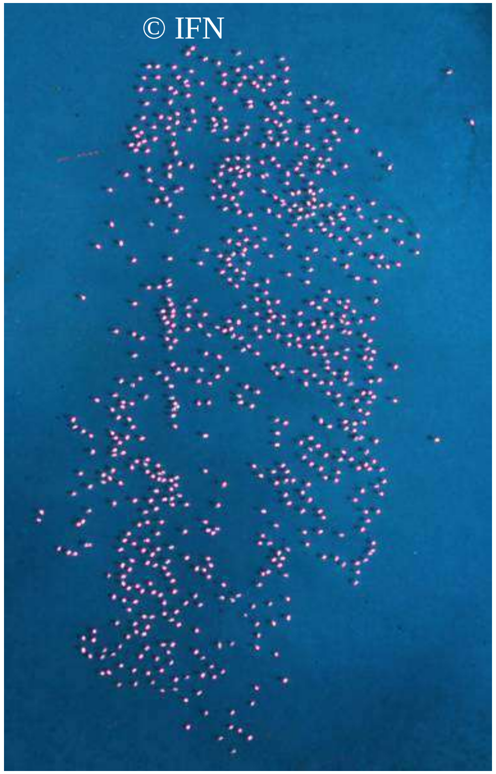

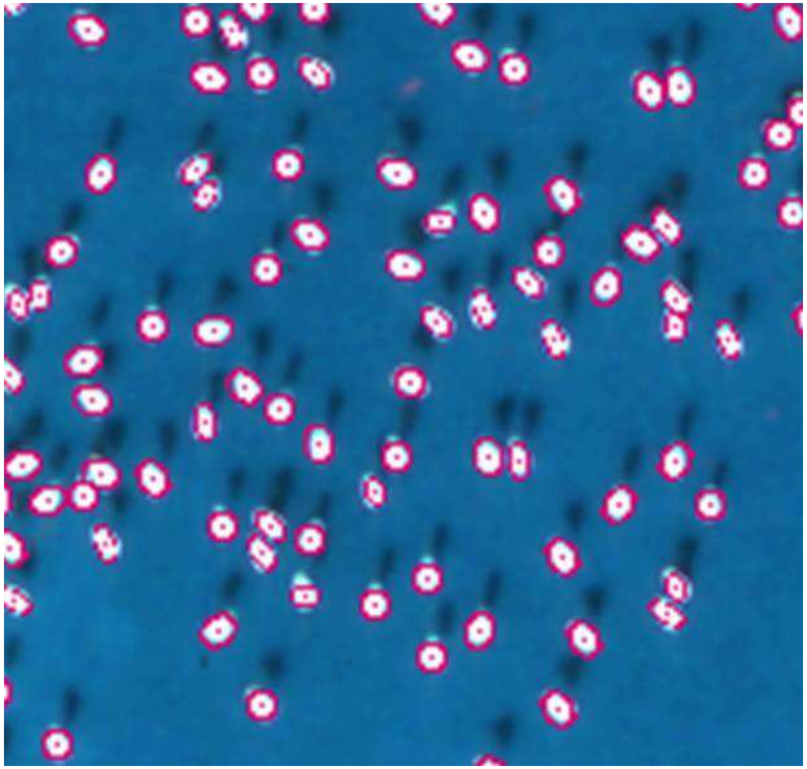

MSc: S. Descamps in collaboration with Tour du Valat

Results

Size estimation of a colony in Turkey (2005):

✓ Low density

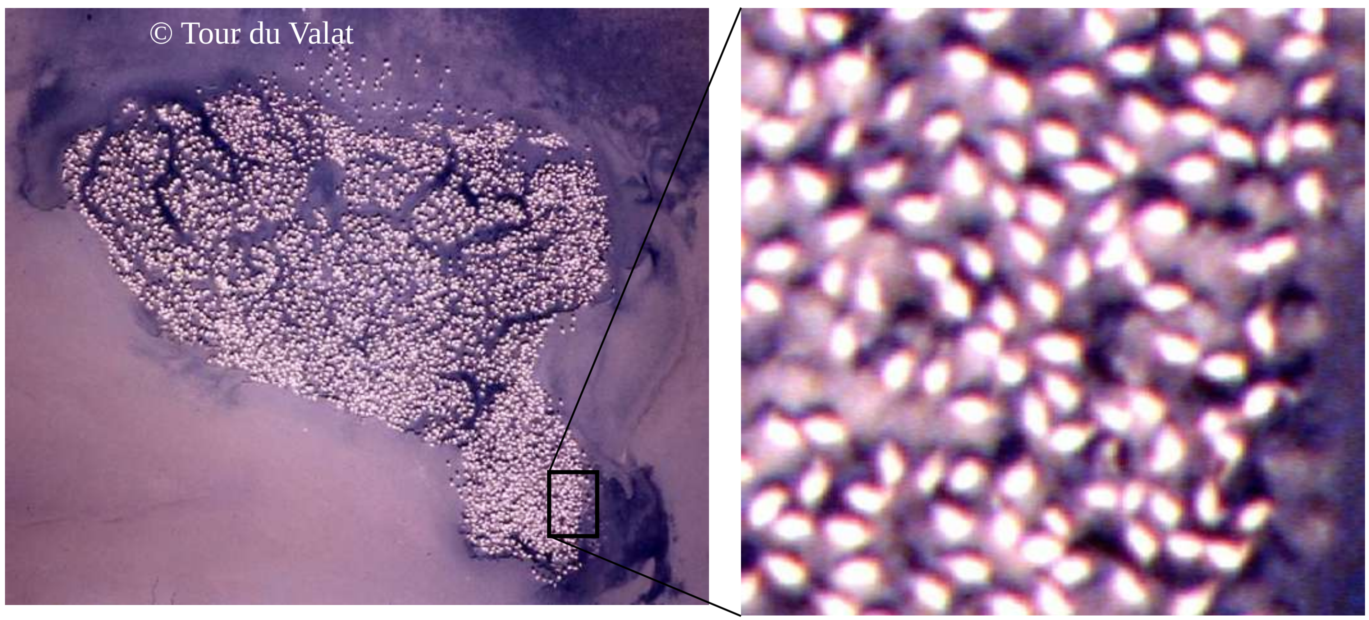

MSc: S. Descamps in collaboration with Tour du Valat

Results

<u>Size estimation of a colony in Turkey (2005):</u>

✓ 3682 flamingos detected for the whole colony (*Tour du Valat = 3684 flamingos*)

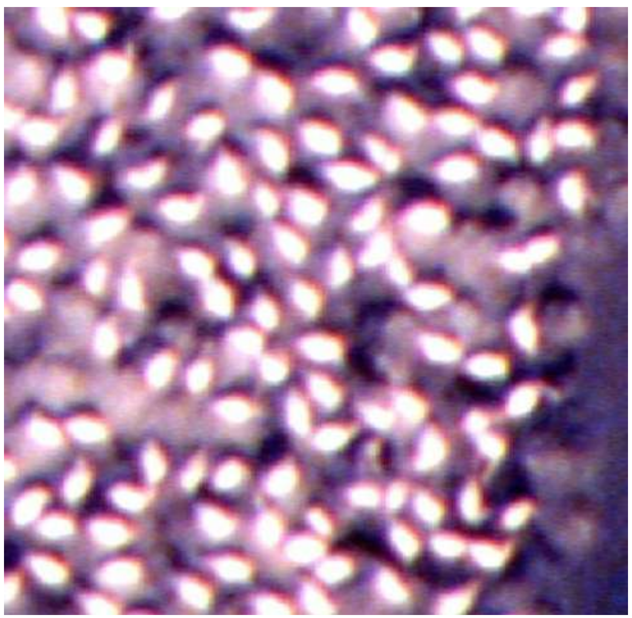

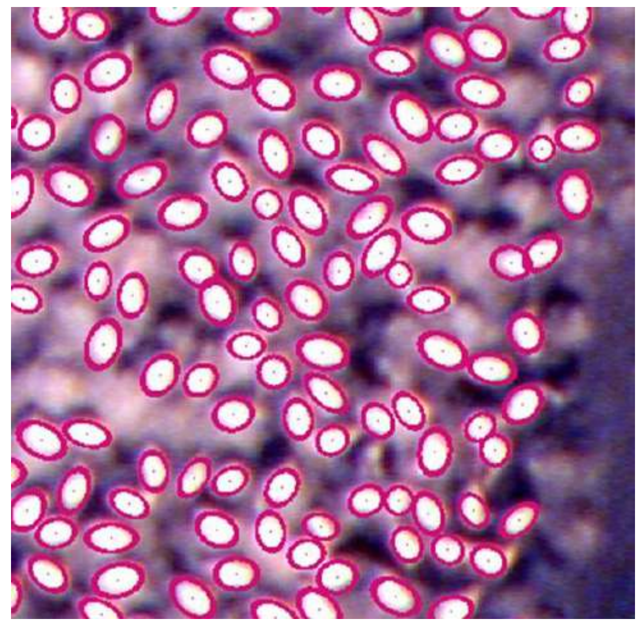

MSc: S. Descamps in collaboration with Tour du Valat

Results

<u>Size estimation of a colony in Mauritania (2004):</u>

✓ High density

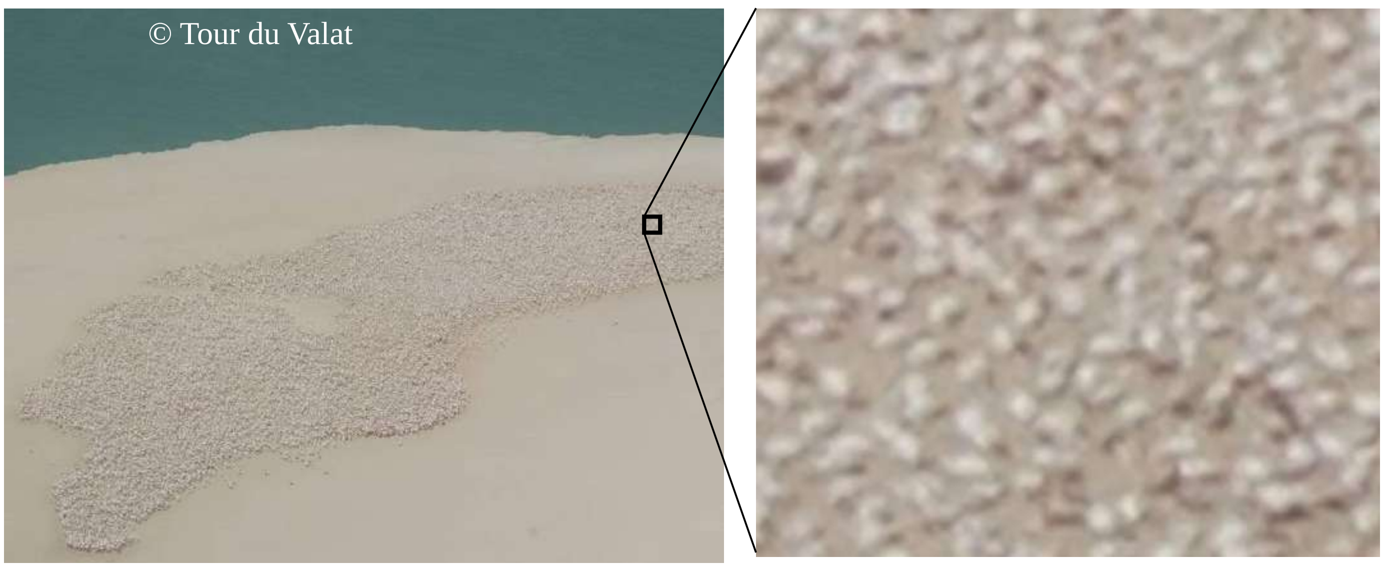

MSc: S. Descamps in collaboration with Tour du Valat

Results

Size estimation of a colony in Mauritania (2004):

✓ 14595 flamingos detected for the whole colony (*Tour du Valat: 13650 flamingos*)

MSc: S. Descamps in collaboration with Tour du Valat

Fourth example: building extraction

Goal: Creation of 3D urban databases

- public (urban planning, disaster recovery …)
- private (wireless telephony, movies …)
- military (operation training, missile guidance …)

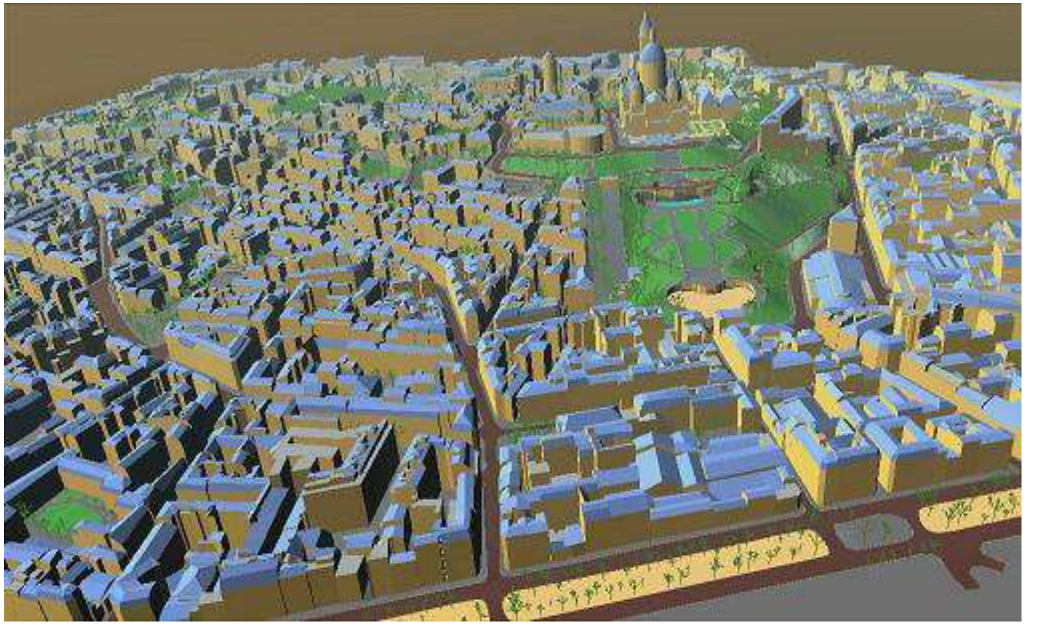

Fourth example: building extraction

Context

- spatial data (PLEIADES simulations, then real data)
- single type of data: a DEM
- automatic (without cadastral maps, without focalisation process)
- dense urban areas

Towards structural modeling

- adapted to data (object approach)
- good compromise generality / robustness
- modular

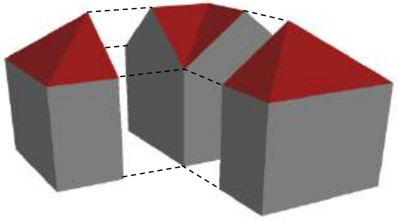

A building = an assembly of simple urban structures

2 stages: 2D extraction, then 3D reconstruction

- computation is greatly reduced

PhDs: M. Ortner, F. Lafarge in collaboration with IGN/CNES/DGA

Stereoscopy

Pair of stereoscopic images

©IGN

3D Information
example: Digital Elevation Model (DEM)
by [Pierrot-Deseilligny et al.,06]

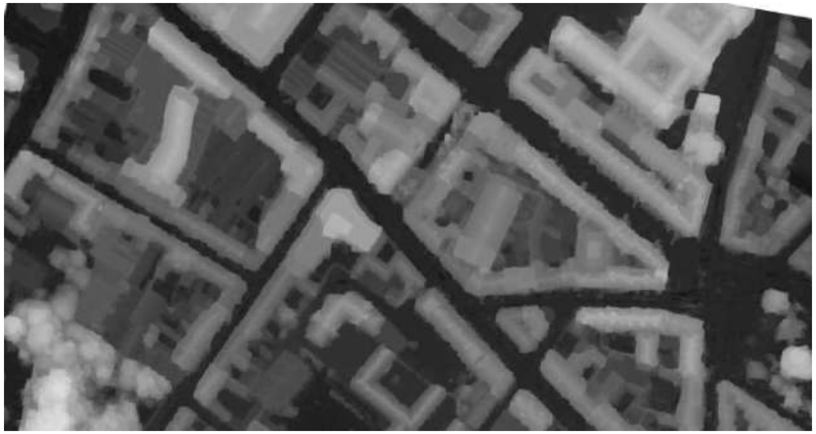

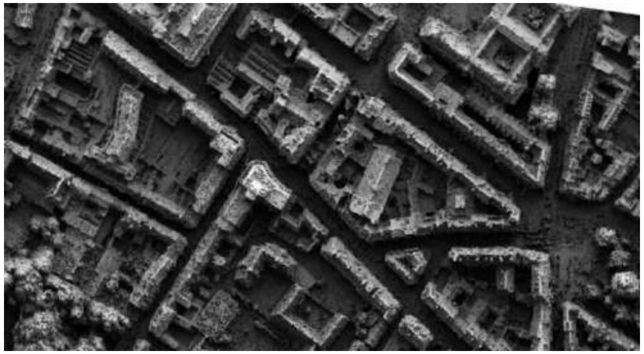

©IGN

Stage 1:
2D extraction of buildings

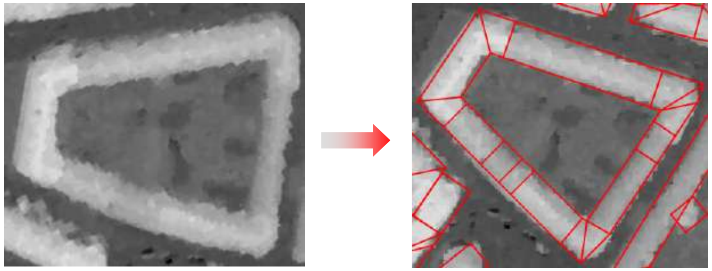

PhDs: M. Ortner, F. Lafarge in collaboration with IGN/CNES/DGA

2D extraction of buildings

Outlines of buildings by marked point processes [Ortner et al. - 04]

- Energy minimization: $U = \rho U_{ext} + U_{int}$

  - U_{ext} : data term
    - coherence between the location of a rectangle and discontinuities in the DEM

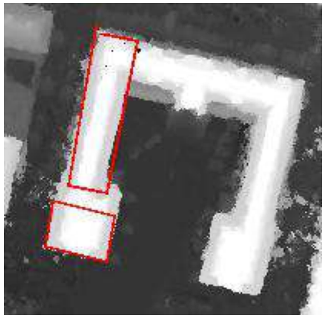

  - U_{int} : regularizing term
    - introduction of prior knowledge about the object layout (alignment, paving, completion)

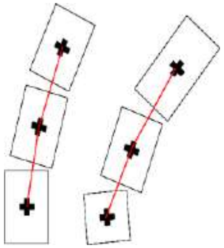

PhDs: M. Ortner, F. Lafarge in collaboration with IGN/CNES/DGA

2D extraction of buildings

Transformation of rectangles into structural supports [Lafarge et al. - 07]

- transformation of rectangles into unspecified quadrilaterals which are ideally connected (without overlapping, with a common edge)

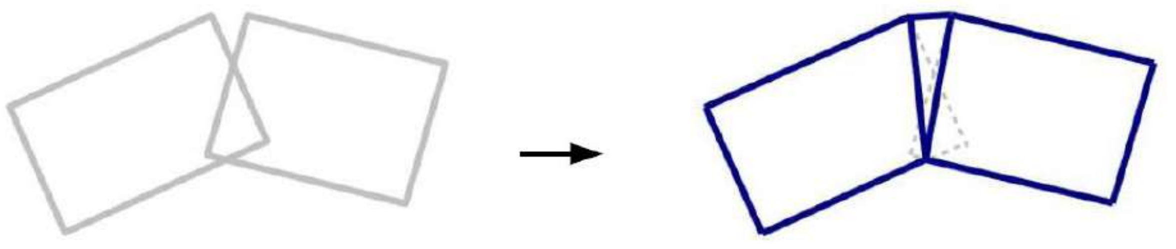

- partitioning of rectangles which represent different urban structures

PhDs: M. Ortner, F. Lafarge in collaboration with IGN/CNES/DGA

2D extraction of buildings

Examples © Ariana / INRIA

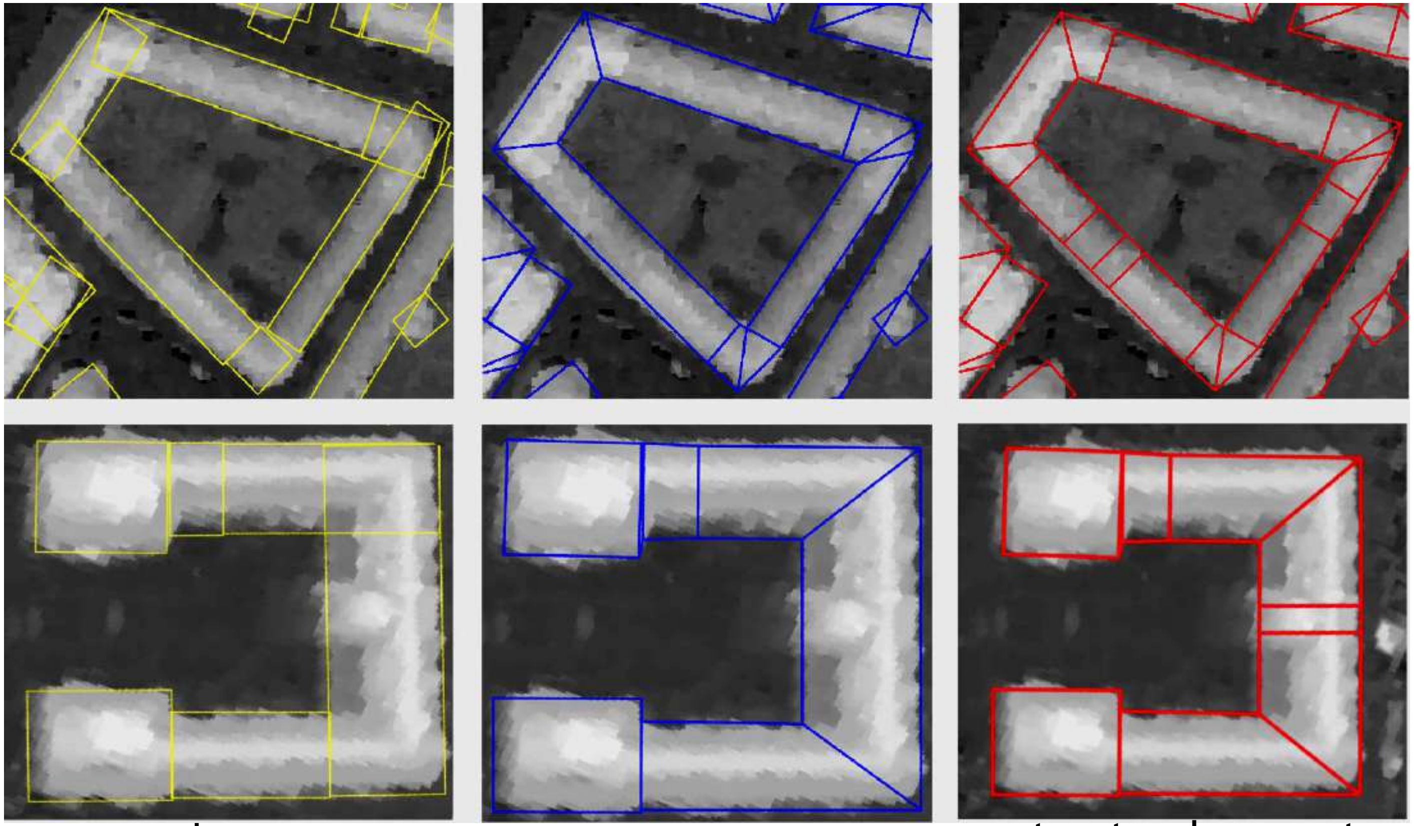

rectangular supports by [Ortner04]

"connected" supports by [Lafarge07]

structural supports by [Lafarge07]

PhDs: M. Ortner, F. Lafarge in collaboration with IGN/CNES/DGA

Stage 2:
3D reconstruction of buildings

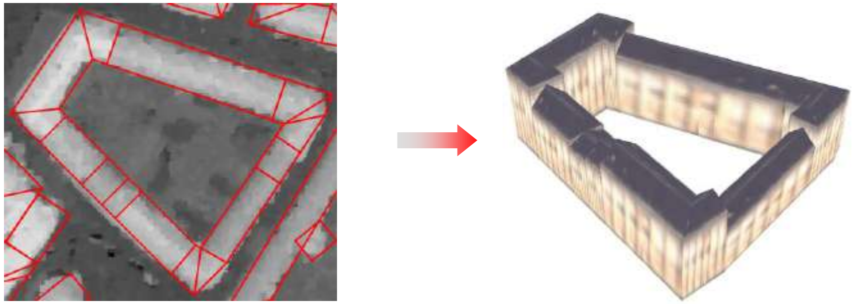

PhDs: M. Ortner, F. Lafarge in collaboration with IGN/CNES/DGA

3D reconstruction of buildings

Library of 3D models
[Lafarge et al. -10]

The roof shapes:

- 9 forms
- 1 to 6 parameters
- includes curved roofs

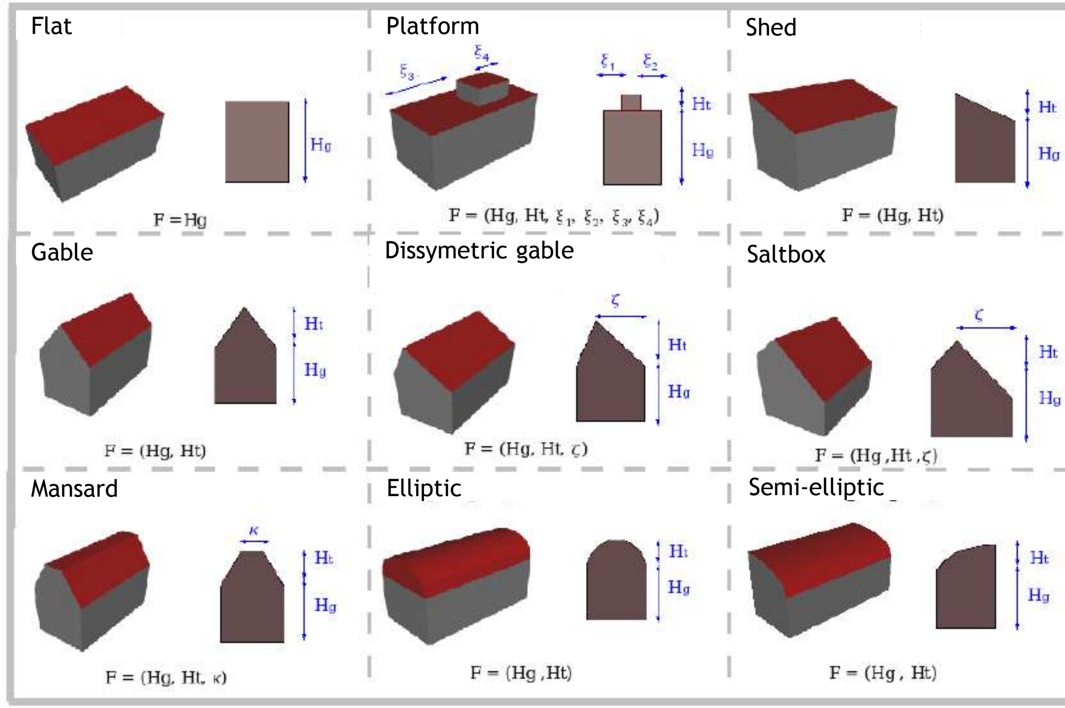

The variants:

- ends and junctions
- orientation of the object

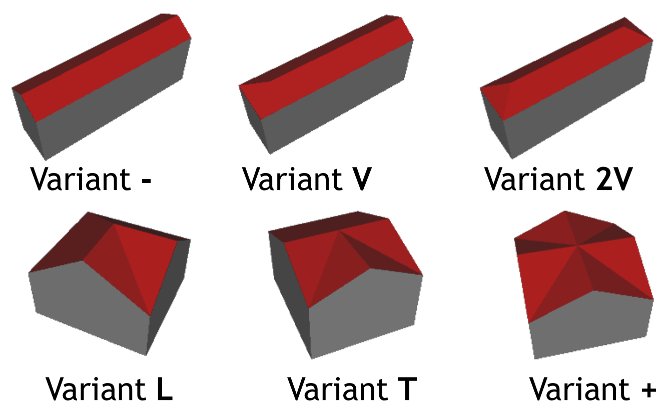

PhDs: M. Ortner, F. Lafarge in collaboration with IGN/CNES/DGA

Inverse problem

Notations

- $\mathcal{Q}$, a configuration of structural supports associated with the DEM Λ
- $\mathcal{Y}$, the data such that $\mathcal{Y} = (\mathcal{Y}_i)_{i \in \mathcal{Q}}$ with $\mathcal{Y}_i = \{\Lambda(s) \in I / s \in S_i\}$
- x, a configuration of 3D objects $x = (x_i)_{i \in \mathcal{Q}}$ where $x_i = (m_i, \theta_i)$ is an objet specified by a model m_i of the library and a parameter set θ_i
- $\mathcal{C}$, the set of 3D object configurations

Inverse problem

- to find the optimal configuration x from the observations $\mathcal{Y}$
- a posteriori density: $h(x) = h(x/\mathcal{Y}) \propto h_p(x)\mathcal{L}(\mathcal{Y}/x)$

Likelihood

Likelihood $\mathcal{L}(\mathcal{Y}/x)$

- to measure the coherence of the observations $\mathcal{Y}$ with an object configuration x

- hypothesis of conditional independence of data:

$$\mathcal{L}(\mathcal{Y}/x) = \prod_{i \in Q} \mathcal{L}(\mathcal{Y}_i/x_i)$$

- use of an altimetric distance between object and DEM:

$$\mathcal{L}(\mathcal{Y}_i/x_i) \propto \exp -\Gamma(\mathcal{S}_{x_i}, \mathcal{Y}_i)$$

where $\mathcal{S}_{x_i}$ corresponds to the roof altitude of object x_i
Γ is the distance (Lp norm)

A priori

A priori $h_p(x)$

- to introduce knowledge w.r.t. the assembling of the objects
  - to compensate for the lack of information contained in the DEM
  - to have realistic buildings
- must be simple (avoid too many tuning parameters)

➡ Solution: a unique type of binary interactions

- Neighboring relationship $\bowtie$ between 2 supports (common edge)
- assembling relation $\sim_a$ between 2 objects
- use of a Gibbs energy: $h_p(x) = \exp - U_p(x)$

A priori

- the assembling relation $\sim_a$ between 2 objects is true if:
  - two objects have the same roof form
  - rooftop orientations are compatible
  - the common edge is not a roof height discontinuity

- A priori expression: $U_p(x) = \beta \sum_{i \bowtie j} \mathbb{1}_{\{x_i \sim_a x_j\}} g(x_i, x_j)$

where $\beta \in \mathbb{R}^+$ is a tuning parameter

g measures the distance between parameters of the objects

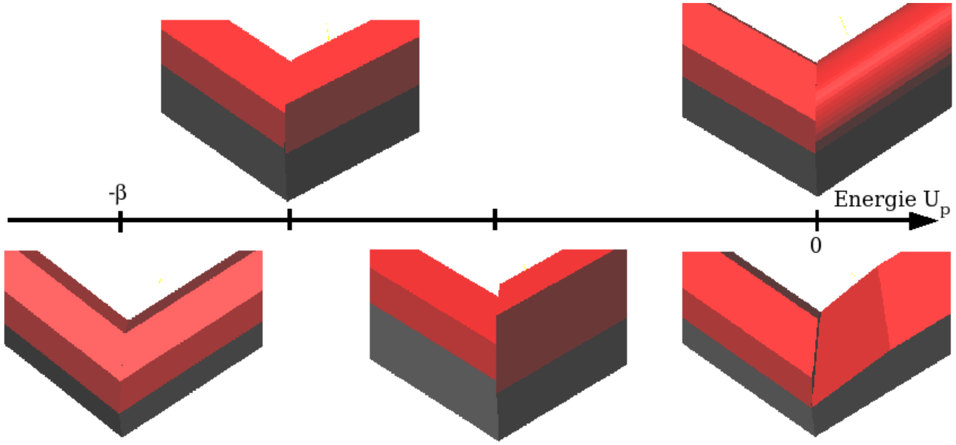

Optimization

MAP estimator: $x_{MAP} = \arg\max_{x \in \mathcal{C}} h(x)$

- non convex optimization problem in a large state space
- $\mathcal{C}$ is a union of spaces of different dimensions

RJMCMC sampler[Green-95]

- consists in simulating a Markov chain $(X_t)_{t \in \mathbb{N}}$ on $\mathcal{C}$ which converges toward a target measure π specified by h

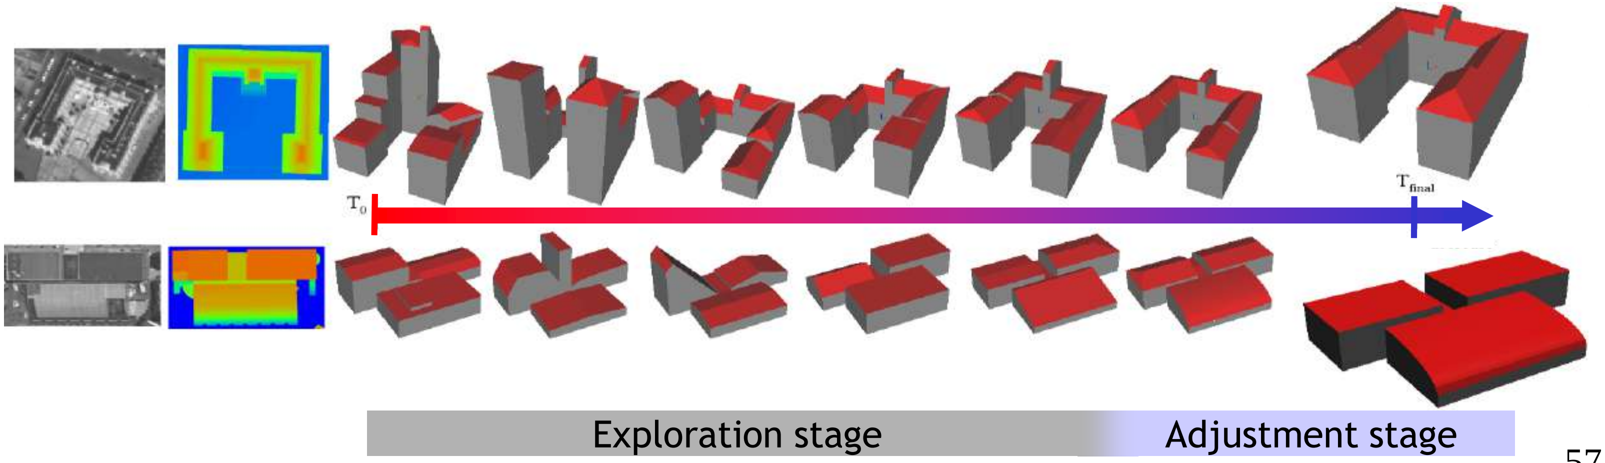

Results

Reconstruction with automatic support extraction

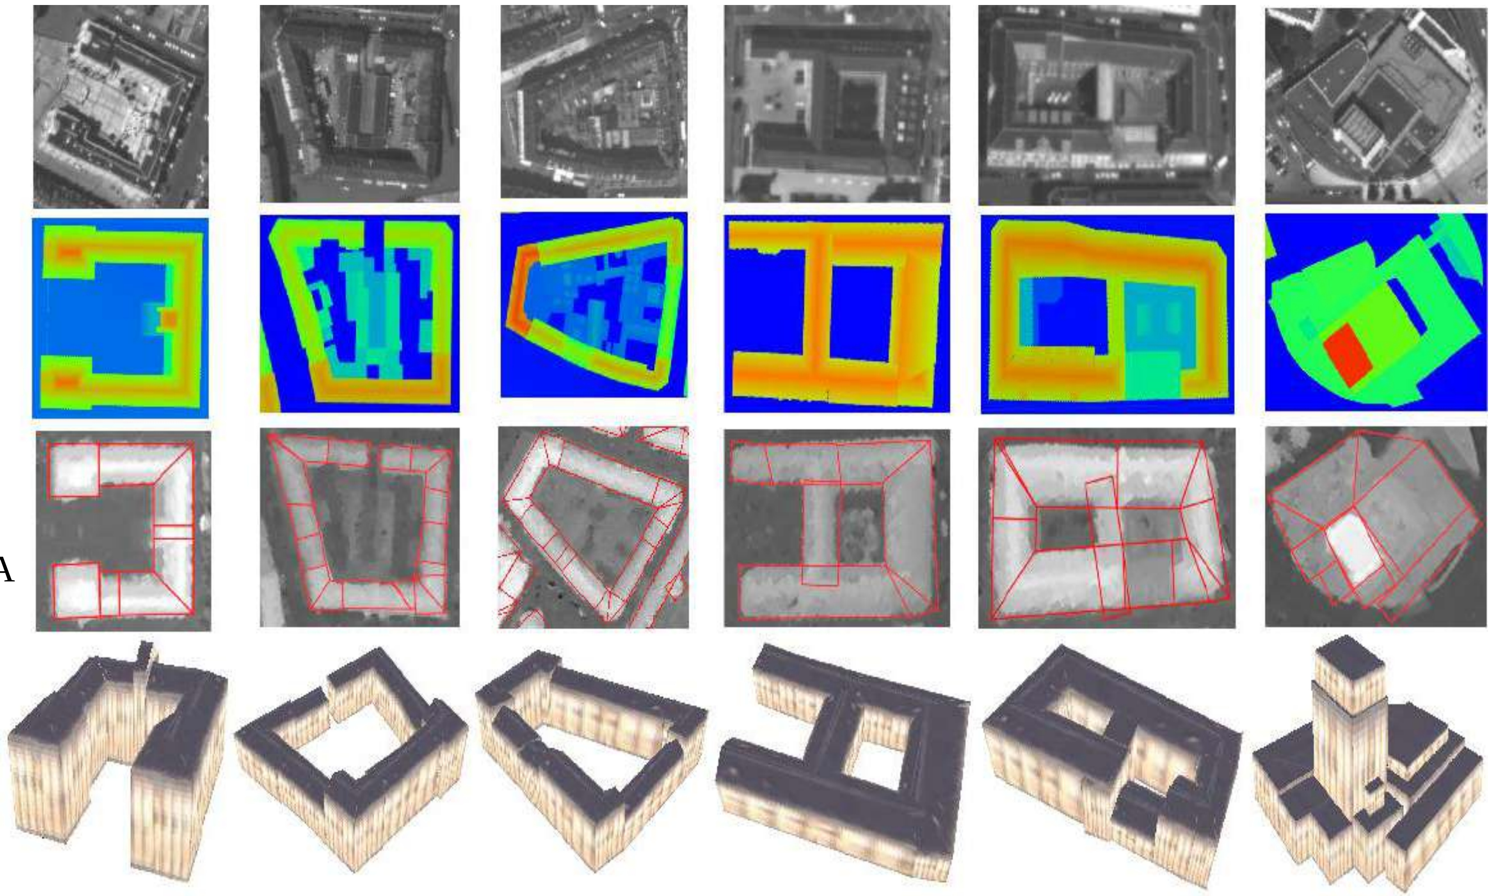

PLEIADES simulations ©CNES

Ground truth ©IGN

Buiding Extraction ©Ariana/INRIA

3D Reconstruction ©IGN/CNES

Results

PLEIADES simulations
© CNES

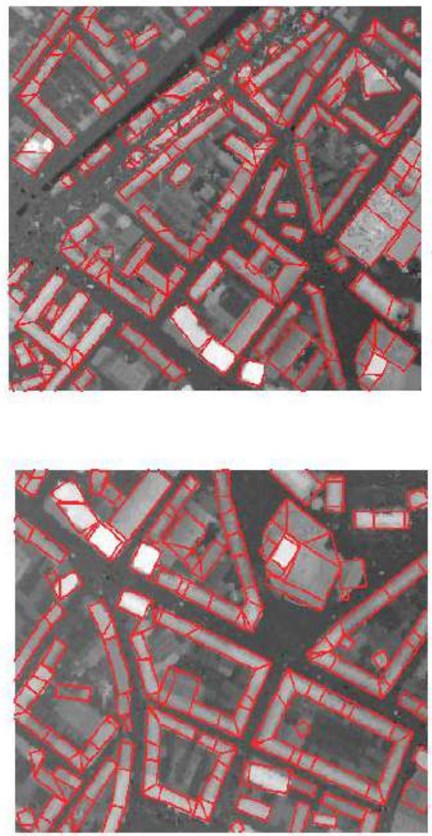

Building Extraction
© Ariana / INRIA

3D Reconstruction
© IGN / CNES

PhDs: M. Ortner, F. Lafarge in collaboration with IGN/CNES/DGA

Results

Reconstruction with interactive support extraction

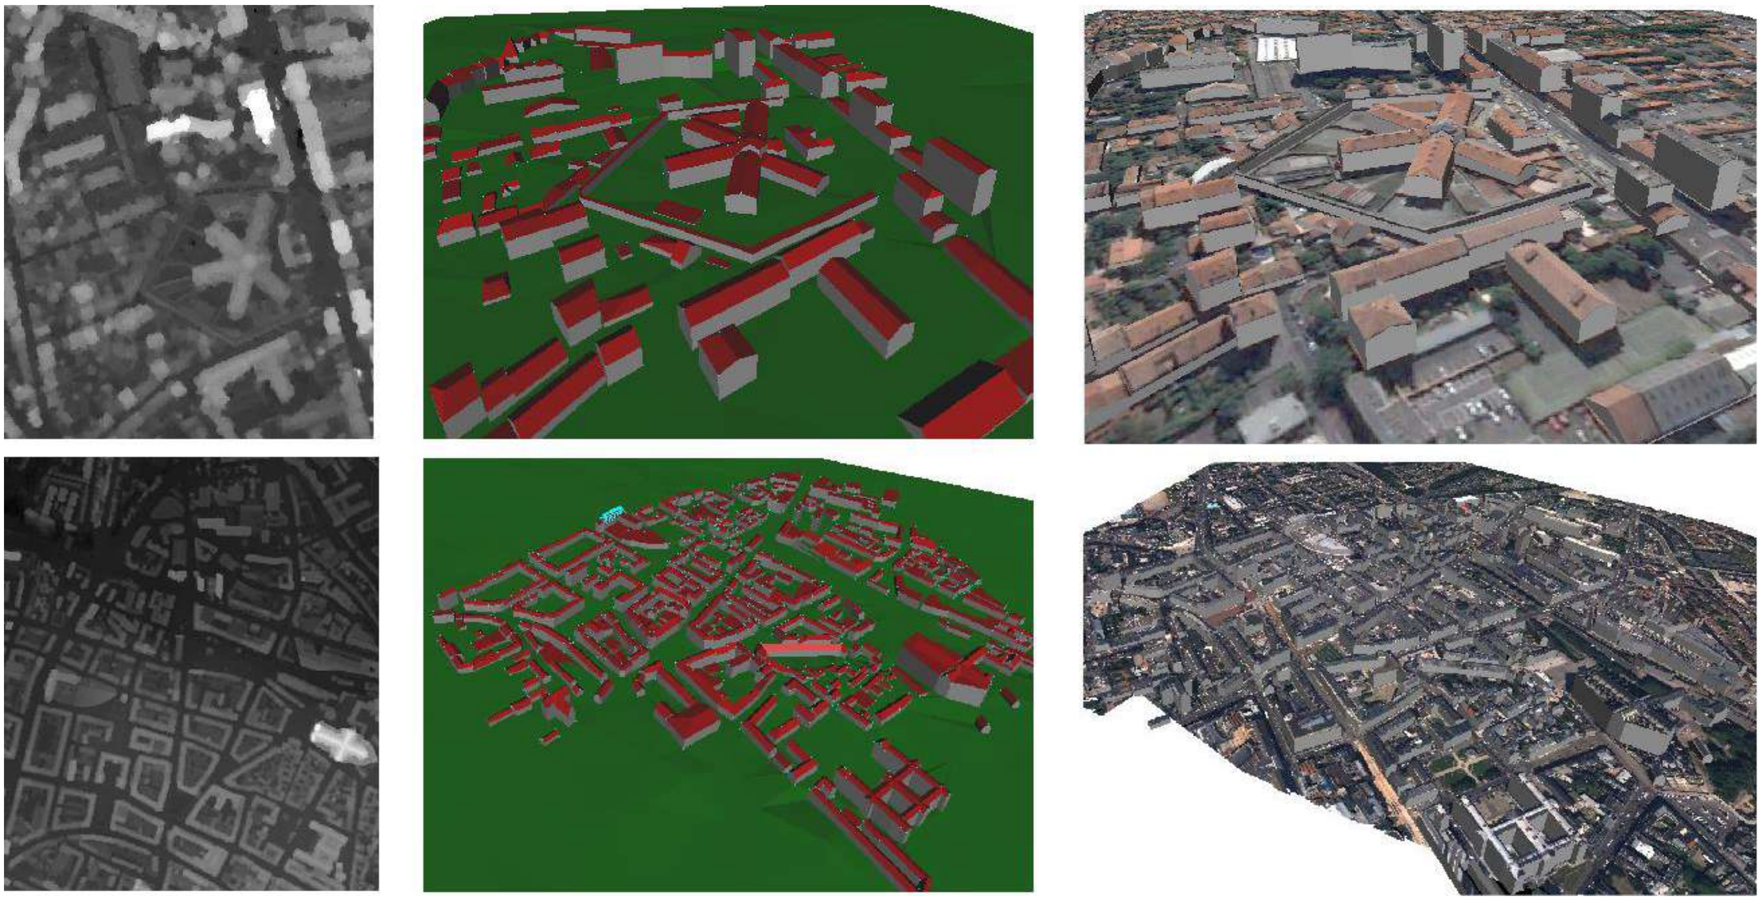

Reconstruction of urban areas (Amiens downtown and St Michel prison in Toulouse) © IGN/CNES

Results

Reconstruction with interactive support extraction

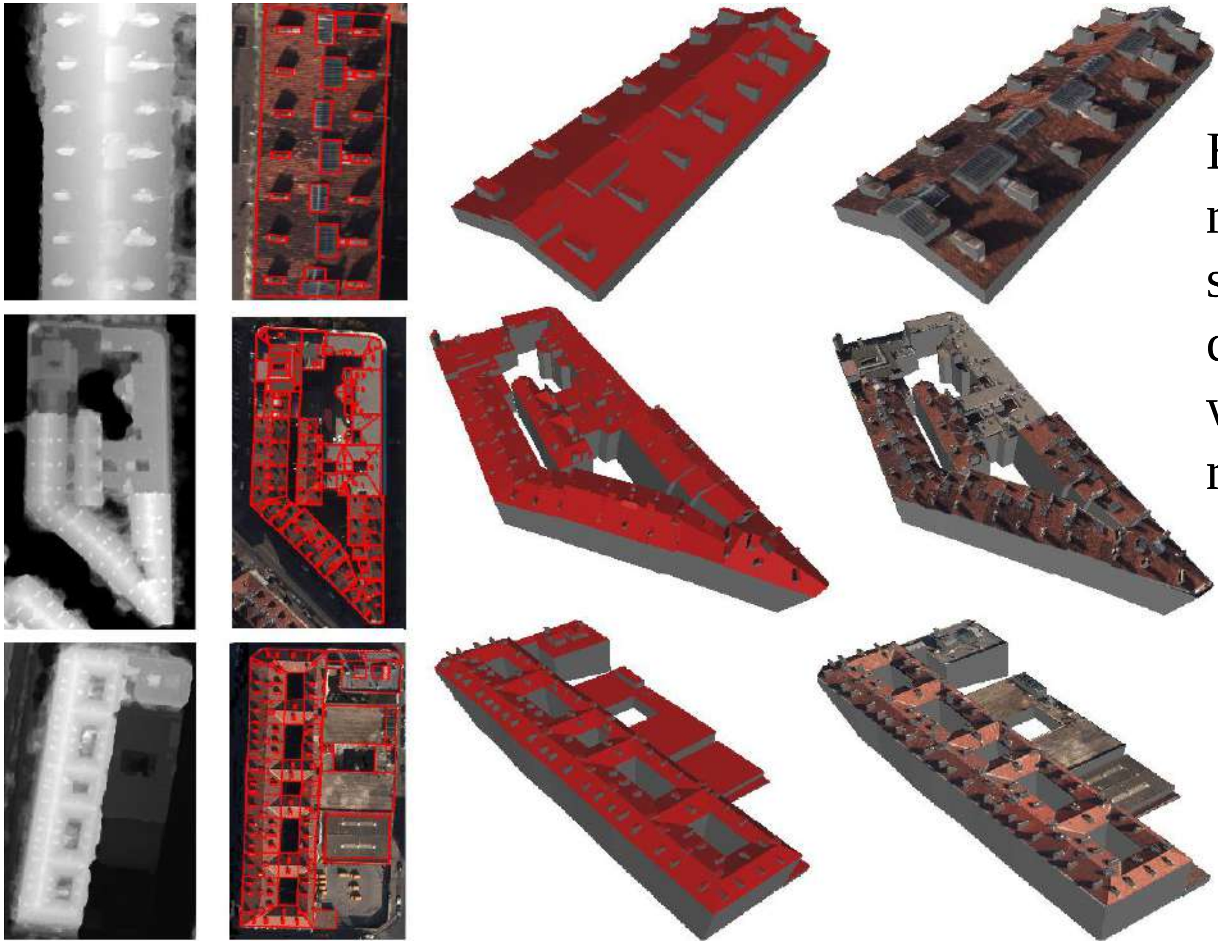

Better reconstruction of superstructures: chimneys, dormer windows, glass roofs...

© IGN/CNES

Reconstruction of buildings with superstructures (Marseille) from 0.1 meter resolution aerial DEM

Fifth example: automatic object detection and tracking

- Context
  - Small object size
  - Large number of objects
  - Shadows
  - Independent camera / object motion
  - Time requirements

PhD: P. Craciun in collaboration with Airbus DS

Optical airborne and space-borne systems

- UAVs (unmanned aerial vehicles)
  - Sub-meter ground sampling resolution imagery
  - Unstable platform

- Low-orbit satellites
  - Sub-meter ground sampling resolution imagery
  - Stable platform
  - High-definition video of up to 90 seconds at 30 frames / second

- Geostationary satellites
  - 1km ground sampling resolution imagery
  - Low temporal frequency …

Multiple Object Tracking (MOT)

- **Goal**: Extract object trajectories throughout a video
- Two sub-problems
  - **Where are the possible targets?** - Detection of targets
  - **Which detection corresponds to each target?** - Solve the data association problem
- Two data-handling approaches
  - **Sequential** – iteratively analyze frames in temporal order
  - **Batch processing** – analyze the entire video at once
- Two main problem solving approaches
  - **Tracking by detection**
  - **Track before detect**

Patterns and stochastic geometry

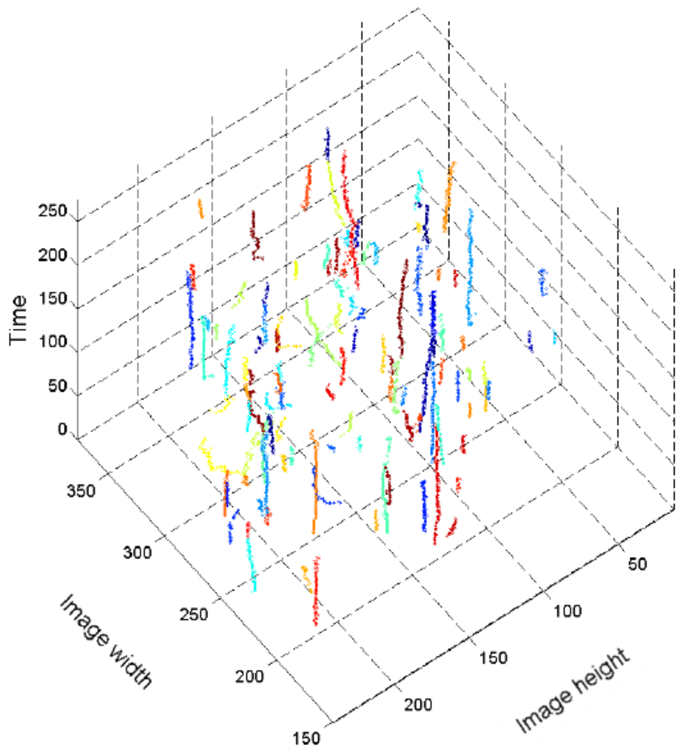

- Object tracking as a spatio-temporal marked point process
- How to model and simulate such a spatio-temporal point process?

PhD: P. Craciun in collaboration with Airbus DS

Overview

- Models
  - Model formulation
  - Quality model vs. Statistical model
- Parameter learning
  - Linear programming
  - Parameter learning as a linear program
- Simulation
  - Standard RJMCMC
  - Parallel implementation of RJMCMC
- Results
- Conclusions and perspectives

Marked point process of ellipses

- Center of the ellipse is a point in the point process
- Marks:
  - Geometric marks: semi-major axis, semi-minor axis, orientation
  - Additional mark: label

$$W = K \times M$$

$$K = [0, I_{h_{max}}] \times [0, I_{w_{max}}] \times \{1, \cdots, T\}$$

$$M = [a_m, a_M] \times [b_m, b_M] \times (-\tfrac{\pi}{2}, \tfrac{\pi}{2}] \times [0, L]$$

$$u = (x_u, y_u, t, a, b, \omega, l)$$

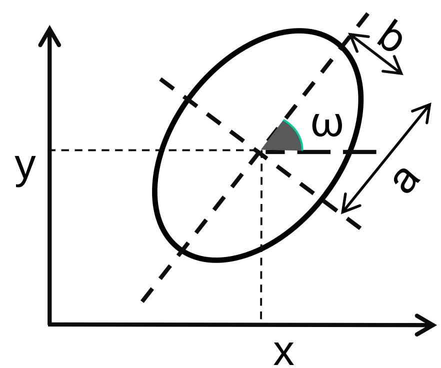

Marked Point Process for Multiple Object Tracking

- Multiple object tracking problem [Craciun et al. - 15]
  - Searching for the most likely configuration $\mathbf{X}$ that fits the given image sequence $\mathbf{Y}$
- Solution
  - $\mathbf{X}$ is a realization of the Gibbs process given by:

$$f_\theta(X = \mathbf{X}|\mathbf{Y}) = \frac{1}{c(\theta|\mathbf{Y})} \exp^{-U_\theta(\mathbf{X},\mathbf{Y})}$$

  - The most likely configuration is given by:

$$X \in arg \max_{\mathbf{X}\in\Omega} f_\theta(X = \mathbf{X}|\mathbf{Y}) = arg \min_{\mathbf{X}\in\Omega} [U_\theta(\mathbf{X},\mathbf{Y})].$$

  - The process energy is composed of two energy terms:

$$U_\theta(\mathbf{X},\mathbf{Y}) = \underbrace{U^{ext}_{\theta_{ext}}}_{\text{External energy}}(\mathbf{X},\mathbf{Y}) + \underbrace{U^{int}_{\theta_{int}}}_{\text{Internal energy}}(\mathbf{X}).$$

Internal energy

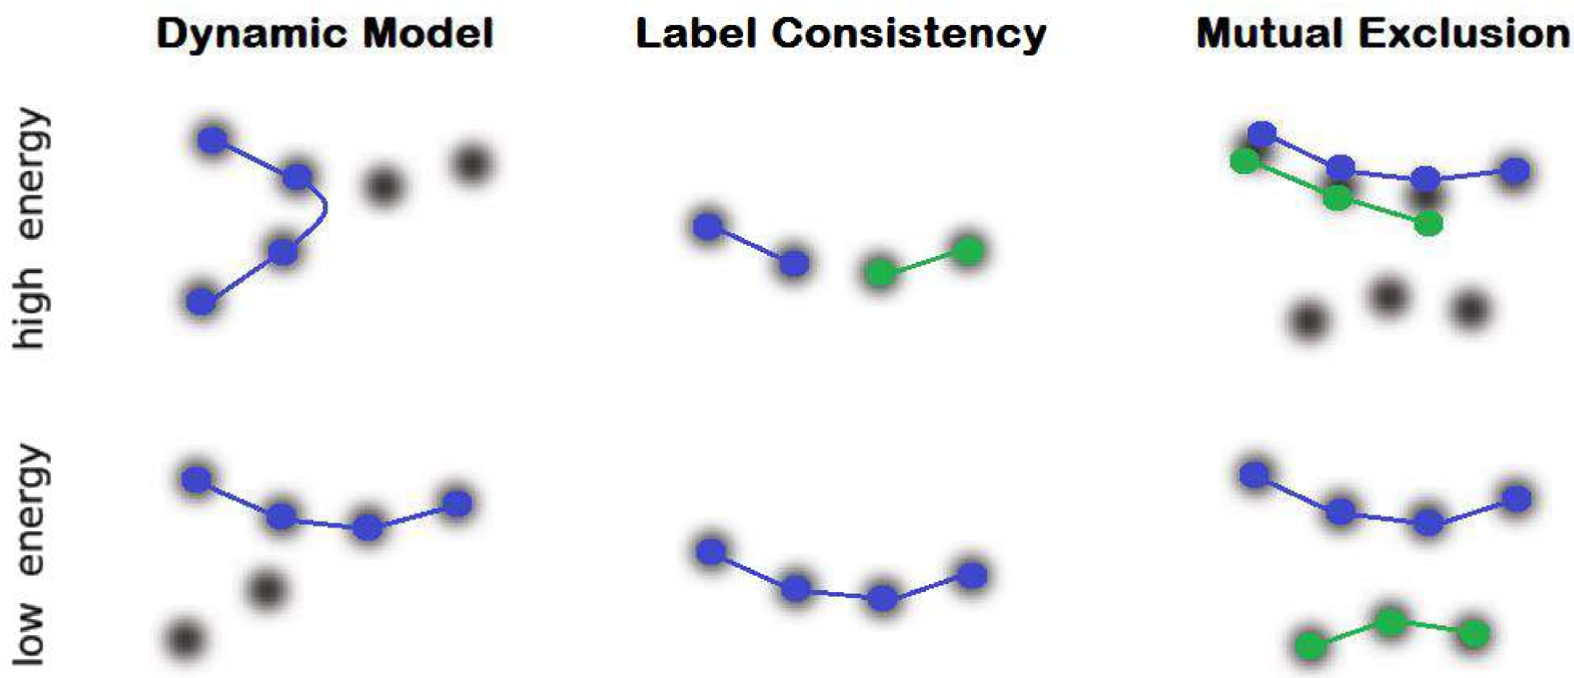

Constant velocity model

Long smooth trajectories

No overlapping objects

PhD: P. Craciun in collaboration with Airbus DS

External energy

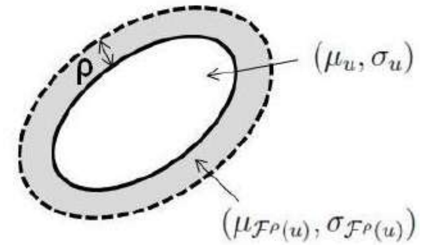

Quality model

- Object evidence through frame differencing

- Contrast distance measure between interior and exterior of ellipse

$$U^{ext}_{\theta_{ext}}(\mathbf{X}|\mathbf{Y}) = \gamma_{ev}\mathcal{E}(u|\mathbf{Y}) + \gamma_{cnt}\sum_{u\in\mathbf{X}}\left(\mathcal{Q}\left(\frac{d_B(u,\mathcal{F}^\rho(u))}{d_0(\mathbf{Y})}\right)\right)$$

Statistical model

- Sliding window
- Two hypotheses:
  - $\boldsymbol{H_0}$: The window covers only the background without any target being present
  - $\boldsymbol{H_1}$: The window is placed in the center of a target
- Neyman-Pearson decision rule

$$U^{ext}_{\theta_{ext}}(\mathbf{X}|\mathbf{Y}) = \gamma_{stat}U^{ext}_{stat}(\mathbf{X}|\mathbf{Y})$$

PhD: P. Craciun in collaboration with Airbus DS

Total energy

Quality model

External energy

$$U_\theta(\mathbf{X},\mathbf{Y}) = \boxed{\gamma_{ev}\mathcal{E}(u|\mathbf{Y}) + \gamma_{cnt}\sum_{u\in\mathbf{X}}\left(\mathcal{Q}\left(\frac{d_B(u,\mathcal{F}^\rho(u))}{d_0(\mathbf{Y})}\right)\right)} + \boxed{\gamma_{dyn}U_{dyn}^{int}(\mathbf{X}) + \gamma_{label}U_{label}^{int}(\mathbf{X}) + \gamma_o U_{overlap}^{int}(\mathbf{X})}$$

Internal energy

Statistical model

External energy

$$U_\theta(\mathbf{X},\mathbf{Y}) = \boxed{\gamma_{stat}U_{stat}^{ext}(\mathbf{X}|\mathbf{Y})} + \boxed{\gamma_{dyn}U_{dyn}^{int}(\mathbf{X}) + \gamma_{label}U_{label}^{int}(\mathbf{X}) + \gamma_o U_{overlap}^{int}(\mathbf{X})}$$

Internal energy

PhD: P. Craciun in collaboration with Airbus DS

Overview

- Models
  - Model formulation
  - Quality model vs. Statistical model
- Parameter learning
  - Linear programming
  - Parameter learning as a linear program
- Simulation
  - Standard RJMCMC
  - Parallel implementation of RJMCMC
- Results
- Conclusions and perspectives

Linear programming

- A linear program has the following form

(1) $$\text{Maximize: } \mathbf{a}^T \mathbf{C}$$

(2) $$\text{Subject to: } A^T \mathbf{C} \leq \mathbf{b}, \quad \mathbf{C} \geq 0$$

Where:

- $\mathbf{a}^T$ – vector of coefficients
- $\mathbf{C}$ – parameter vector
- $A^T \mathbf{C} \leq \mathbf{b}$ – constraints

Objective function

- Quality model energy formulation

$$U_\theta(\mathbf{X},\mathbf{Y}) = \gamma_{ev}\ \mathcal{E}(u|\mathbf{Y}) + \gamma_{cnt}\sum_{u\in\mathbf{X}}\left(\mathcal{Q}\left(\frac{d_B(u,\mathcal{F}^\rho(u))}{d_0(\mathbf{Y})}\right)\right) +$$

$$\gamma_{dyn}\, U_{dyn}^{int}(\mathbf{X}) + \gamma_{label}\, U_{label}^{int}(\mathbf{X}) + \gamma_o\, U_{overlap}^{int}(\mathbf{X})$$

- Objective function

$$\mathbf{a} = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} \qquad \mathbf{C} = \begin{bmatrix} \gamma_{ev} \\ \gamma_{cnt} \\ \gamma_{dyn} \\ \gamma_{label} \\ \gamma_o \end{bmatrix}$$

Gathering constraints

- Only the ratio $\pi(\mathbf{X}')/\pi(\mathbf{X})$ is needs to be computed
- We can create inequalities of the form [Yu 2009]

$$\pi(\mathbf{X}')/\pi(\mathbf{X}) \geq 1$$

- If we have ground truth information

$$\frac{\pi(\mathbf{X}^*)}{\pi(\mathbf{X}_i)} \geq 1$$

- Or more specifically the constraints can be written as

$$f(\mathbf{C}|\mathbf{X}^*) - f(\mathbf{C}|\mathbf{X}_i) \geq 0$$

How many constraints?

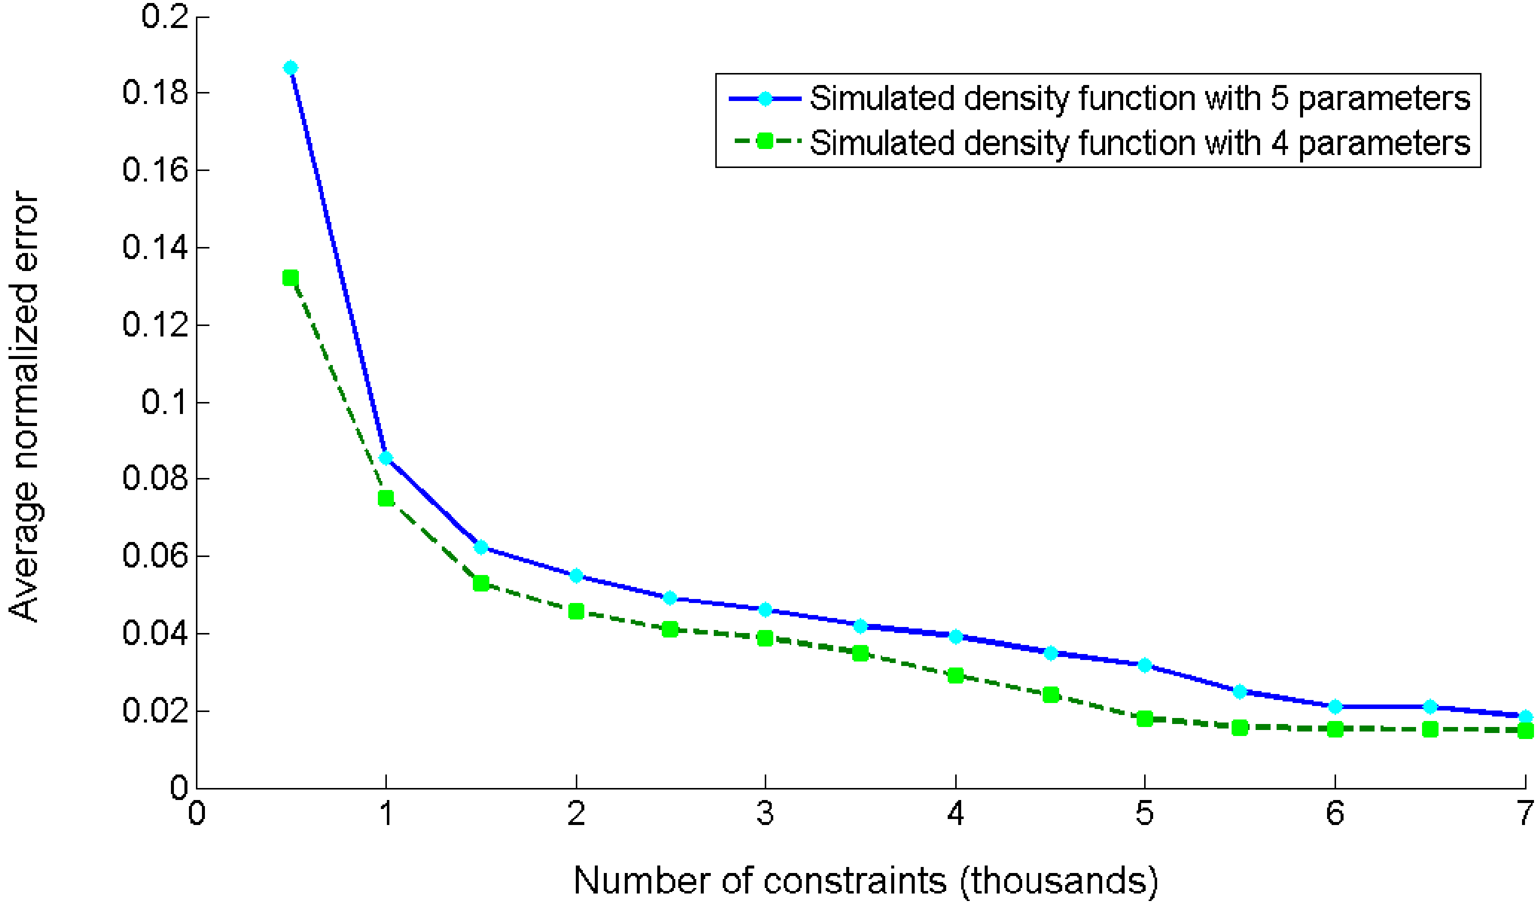

PhD: P. Craciun in collaboration with Airbus DS

Overview

- Models
  - Model formulation
  - Quality model vs. Statistical model
- Parameter learning
  - Linear programming
  - Parameter learning as a linear program
- Simulation
  - Standard RJMCMC
  - Parallel implementation of RJMCMC
- Results
- Conclusions and perspectives

Related samplers

RJMCMC

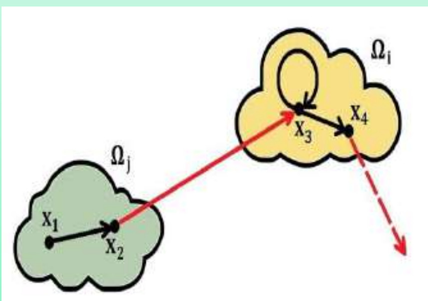

[Green1995]

MBD and MBC

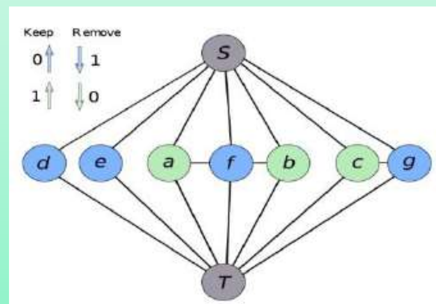

[Descombes2009]
[Gamal2011]

P-RJMCMC

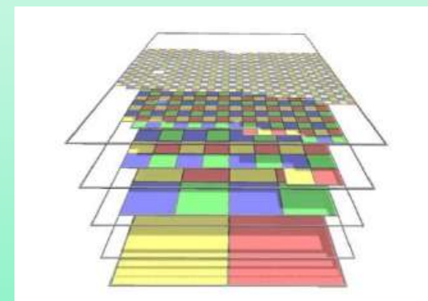

[Verdie2013]

Classic RJMCMC

Standard perturbation kernels

- Birth and Death
  - Birth:
    - Add a new object to the configuration
  - Death:
    - Remove one object from the configuration
- Local transformations

Rotation Translation Scale

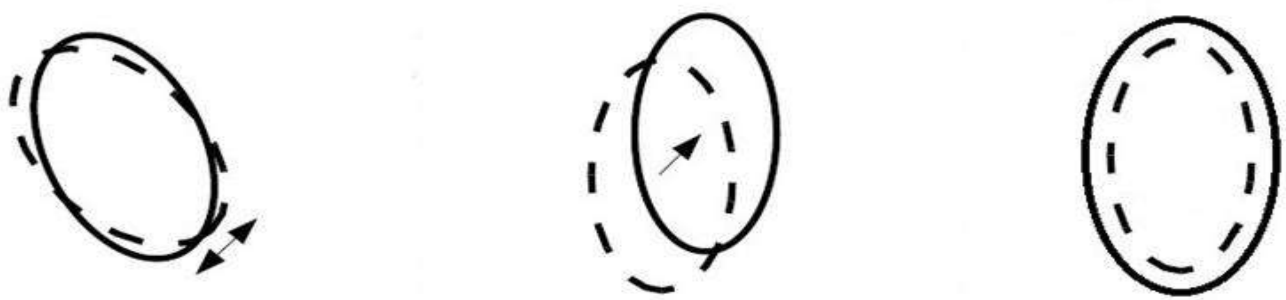

Adding Kalman-inspired births

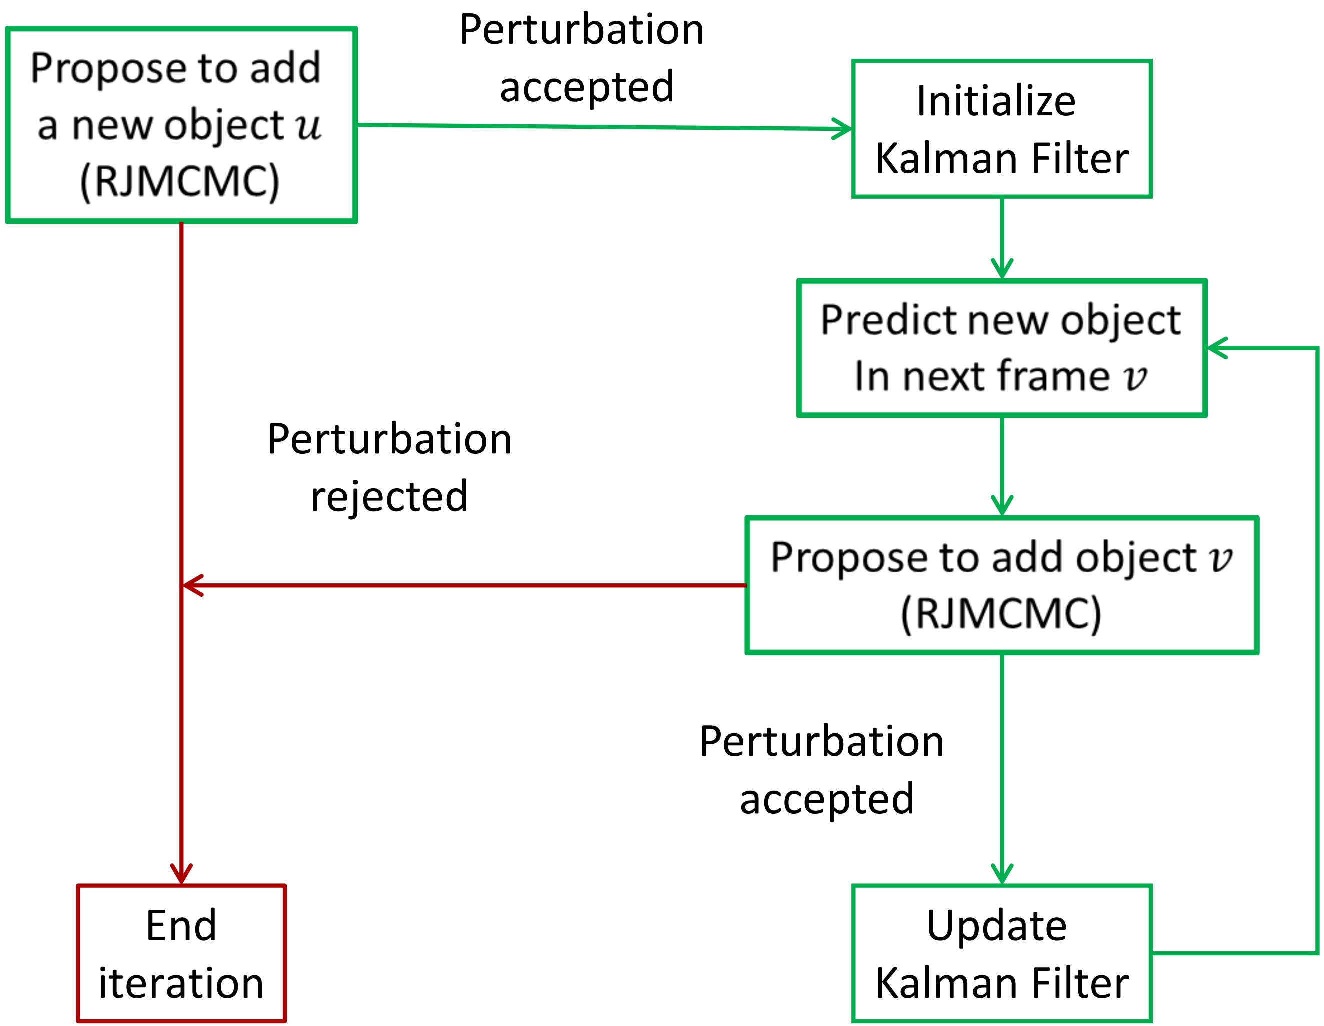

PhD: P. Craciun in collaboration with Airbus DS

Did time efficiency increase?

RJMCMC with Kalman like moves [Craciun et al. - 15]
converges much faster compared to the
standard RJMCMC

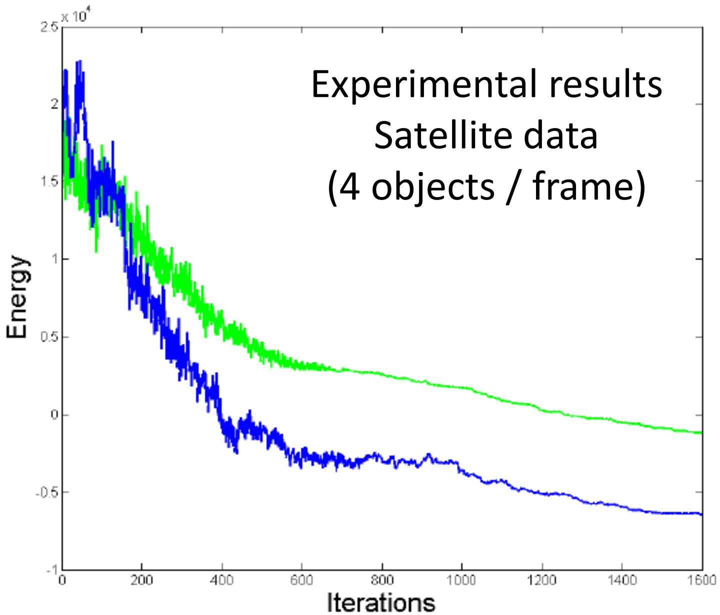

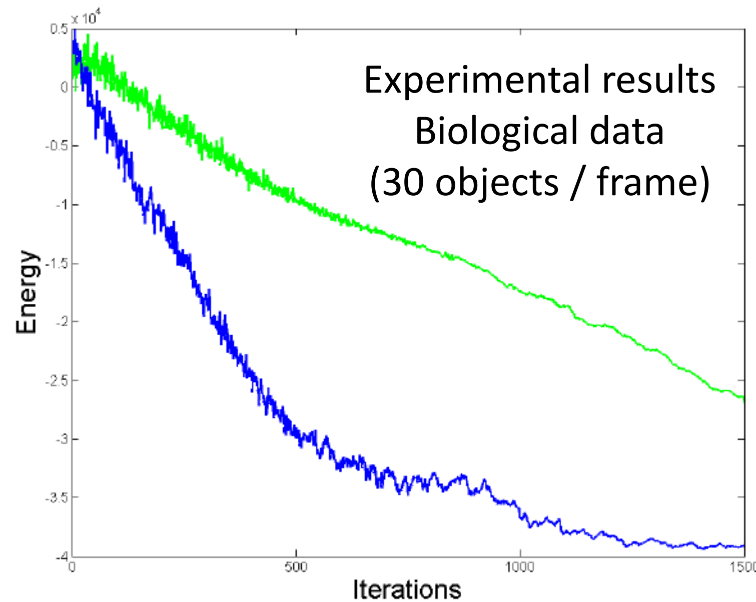

Kalman-inspired births reduce computation times!

Parallel implementation of RJMCMC [Verdie2013]

- Data-driven space partitioning
- Locally conditional independent perturbations

Image with boats © Airbus D&S

Parallel implementation of RJMCMC

- Data-driven space partitioning
- Locally conditional independent perturbations

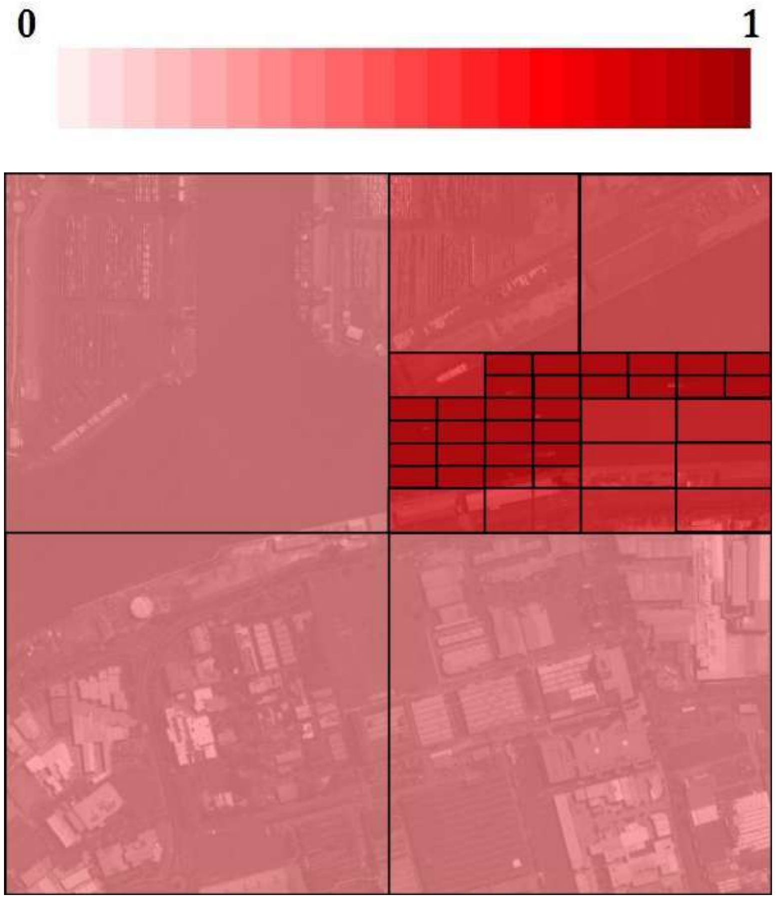

Probability that objects exist in each part of the image

Parallel implementation of RJMCMC

- Data-driven space partitioning
- Locally conditional independent perturbations

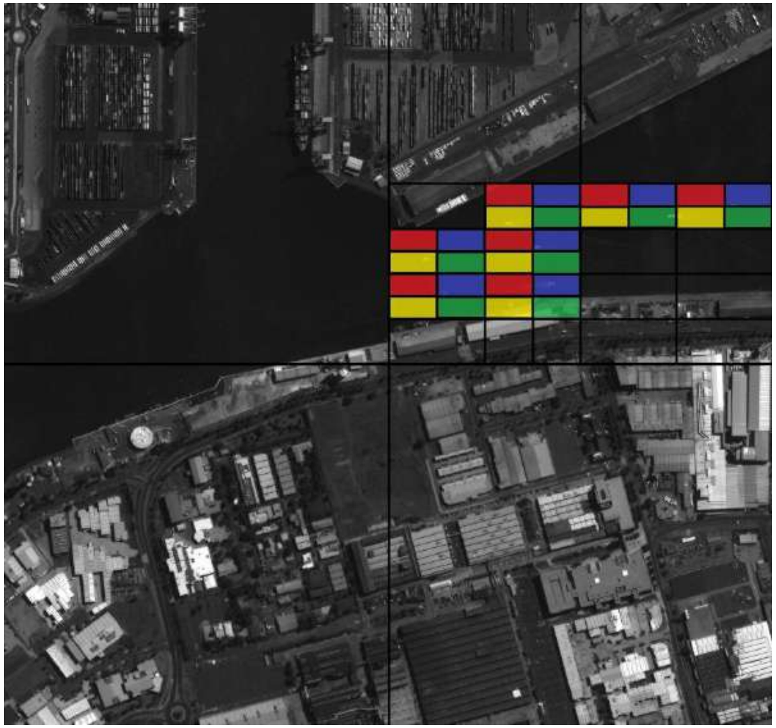

Color coding of quad-tree leafs

Parallel perturbations [Verdie2013]

- A color is randomly chosen
- Perturbations are performed in all cells of the chosen color in parallel

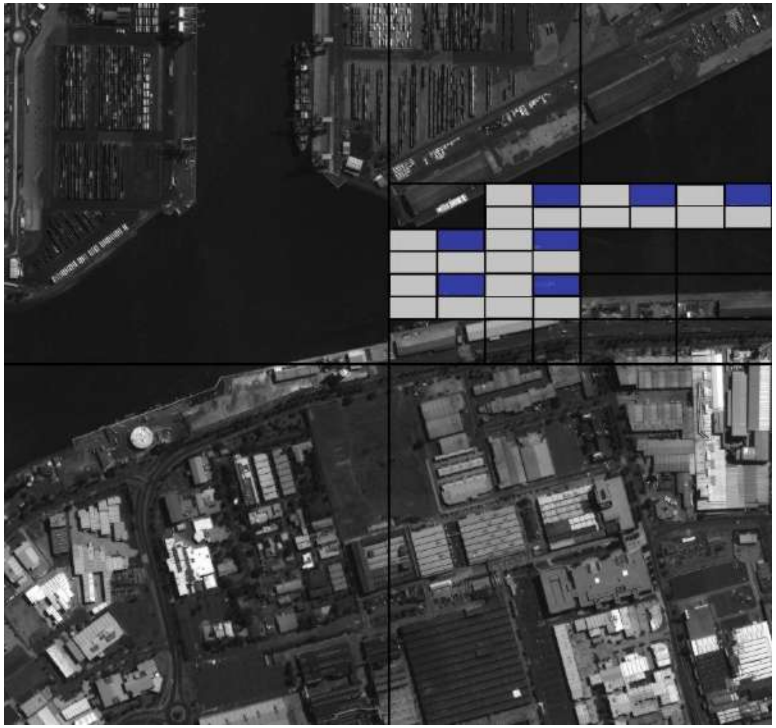

Color blue is randomly chosen

Our improvement to the parallel sampler

Problem

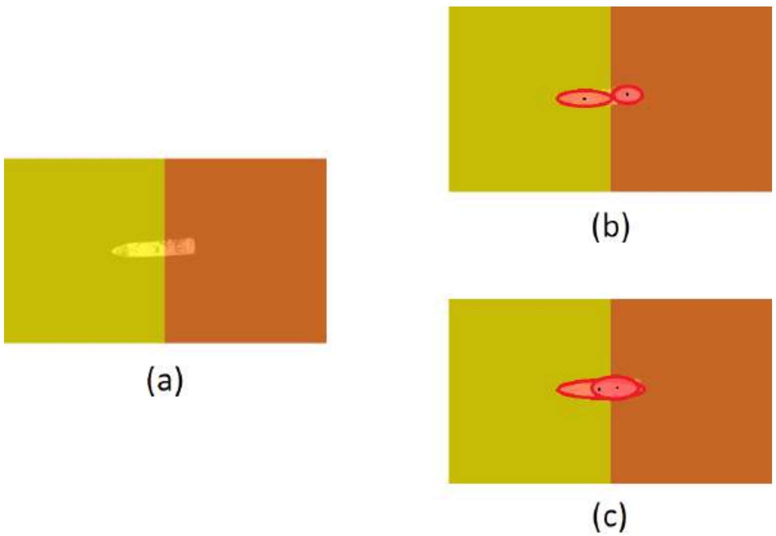

Large boat is split between two neighboring cells

Solution

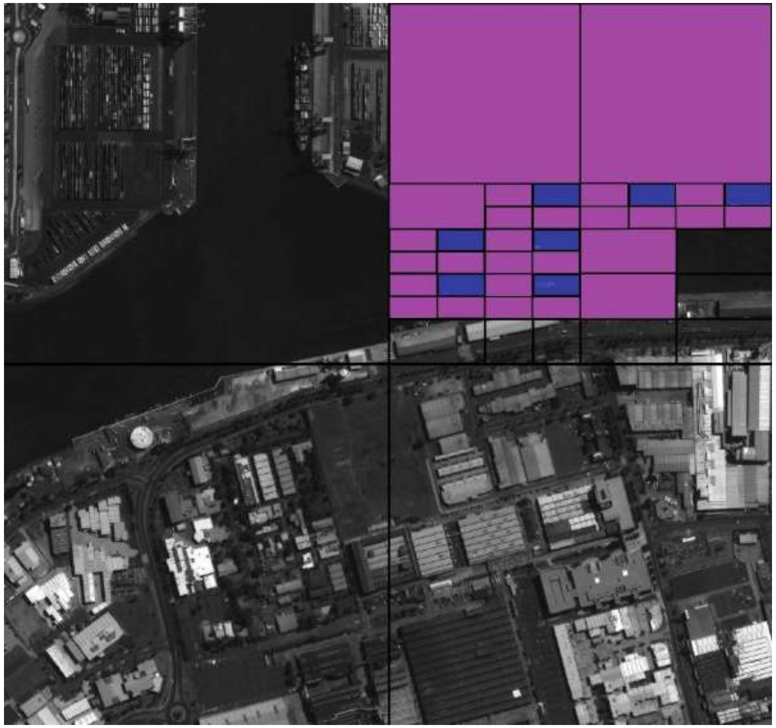

Take the configurations in the neighboring cells into consideration

Did time efficiency increase?

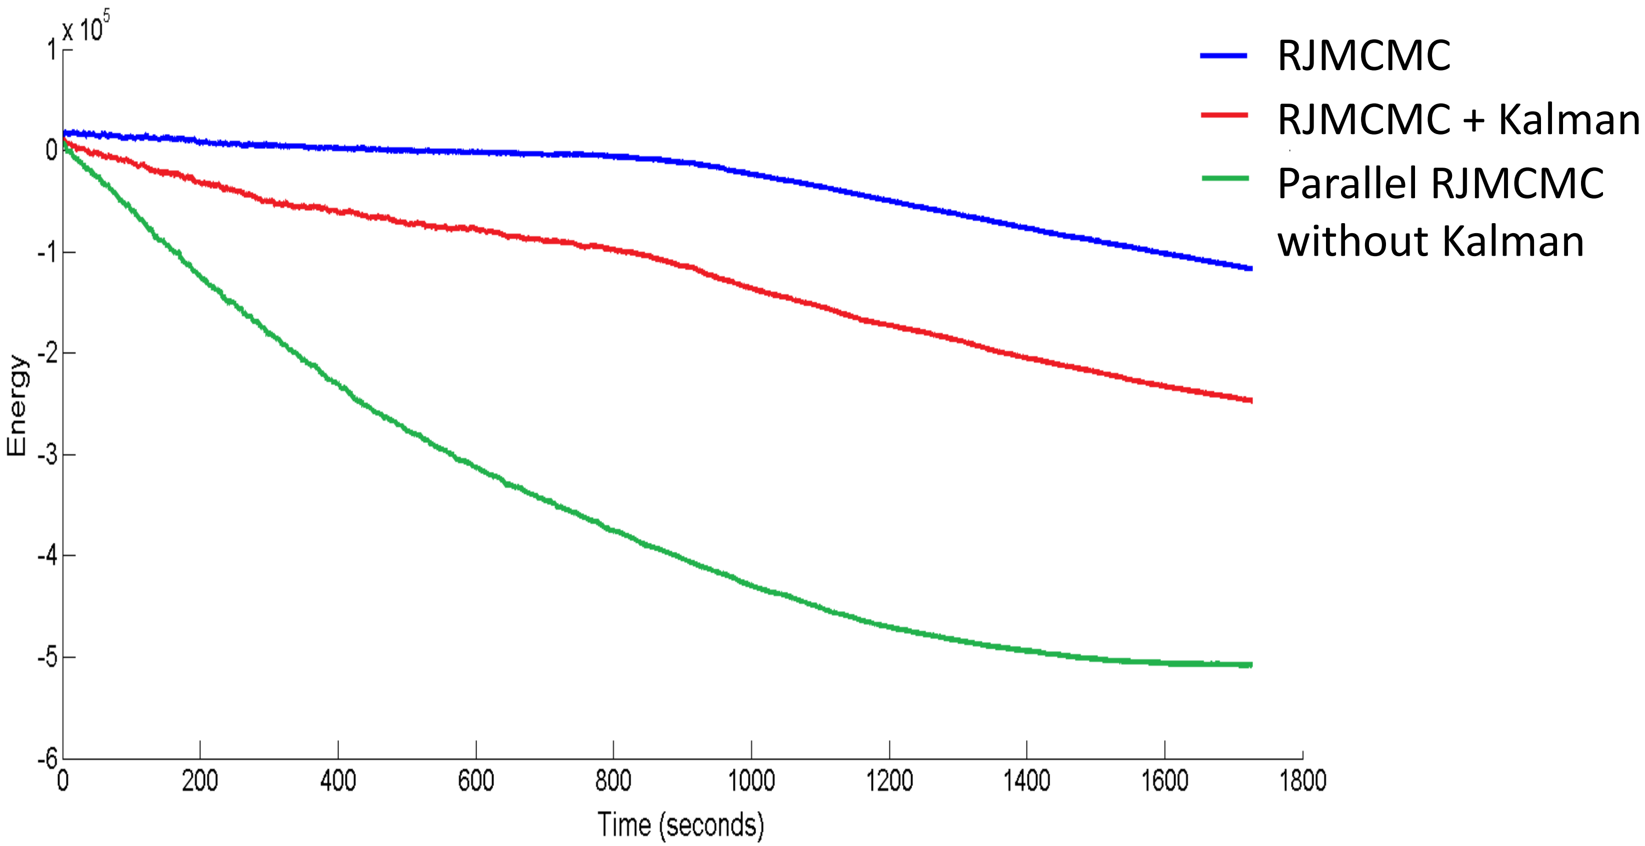

Parallel implementation significantly reduces computation times!

PhD: P. Craciun in collaboration with Airbus DS

Overview

- Models
  - Model formulation
  - Quality model vs. Statistical model
- Parameter learning
  - Linear programming
  - Parameter learning as a linear program
- Simulation
  - Standard RJMCMC
  - Parallel implementation of RJMCMC
- Results
- Conclusions and perspectives

Data sets

- 4 different data sets [Craciun – 15]:
  - Synthetic biological benchmarks (Pasteur Institute)
    - 100 frames / sequence
    - Different levels of noise
  - UAV (unmanned aerial vehicle) data (Public available data set)
  - Satellite data (Airbus Defense and Space)
    - Low temporal frequency (~1-2Hz)
    - High temporal frequency (30Hz)

Synthetic biological benchmarks

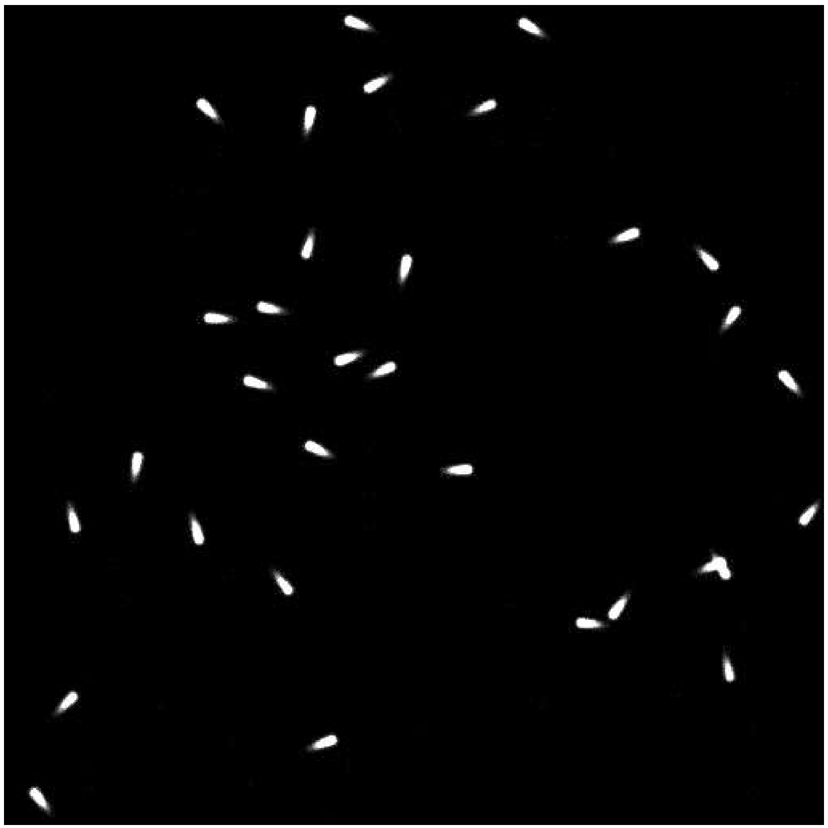

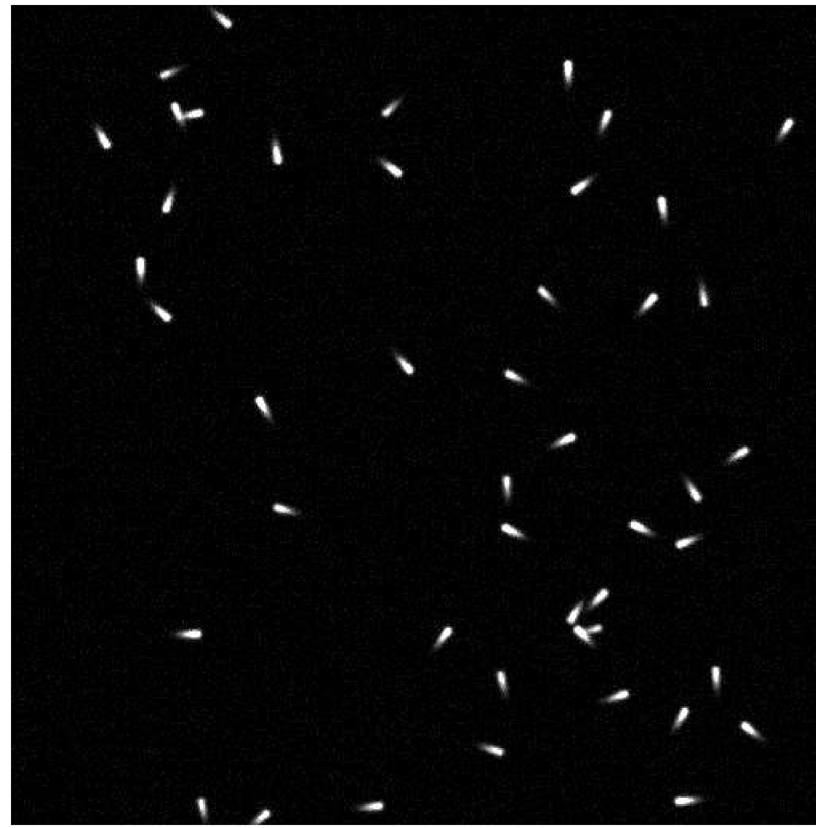

*Generated using the publicly available software ICY courtesy of the Quantitative Analysis Unit at the Pasteur Institute and J.-C. Olivo-Marin, see [Craciun 2016]

Synthetic biological benchmarks

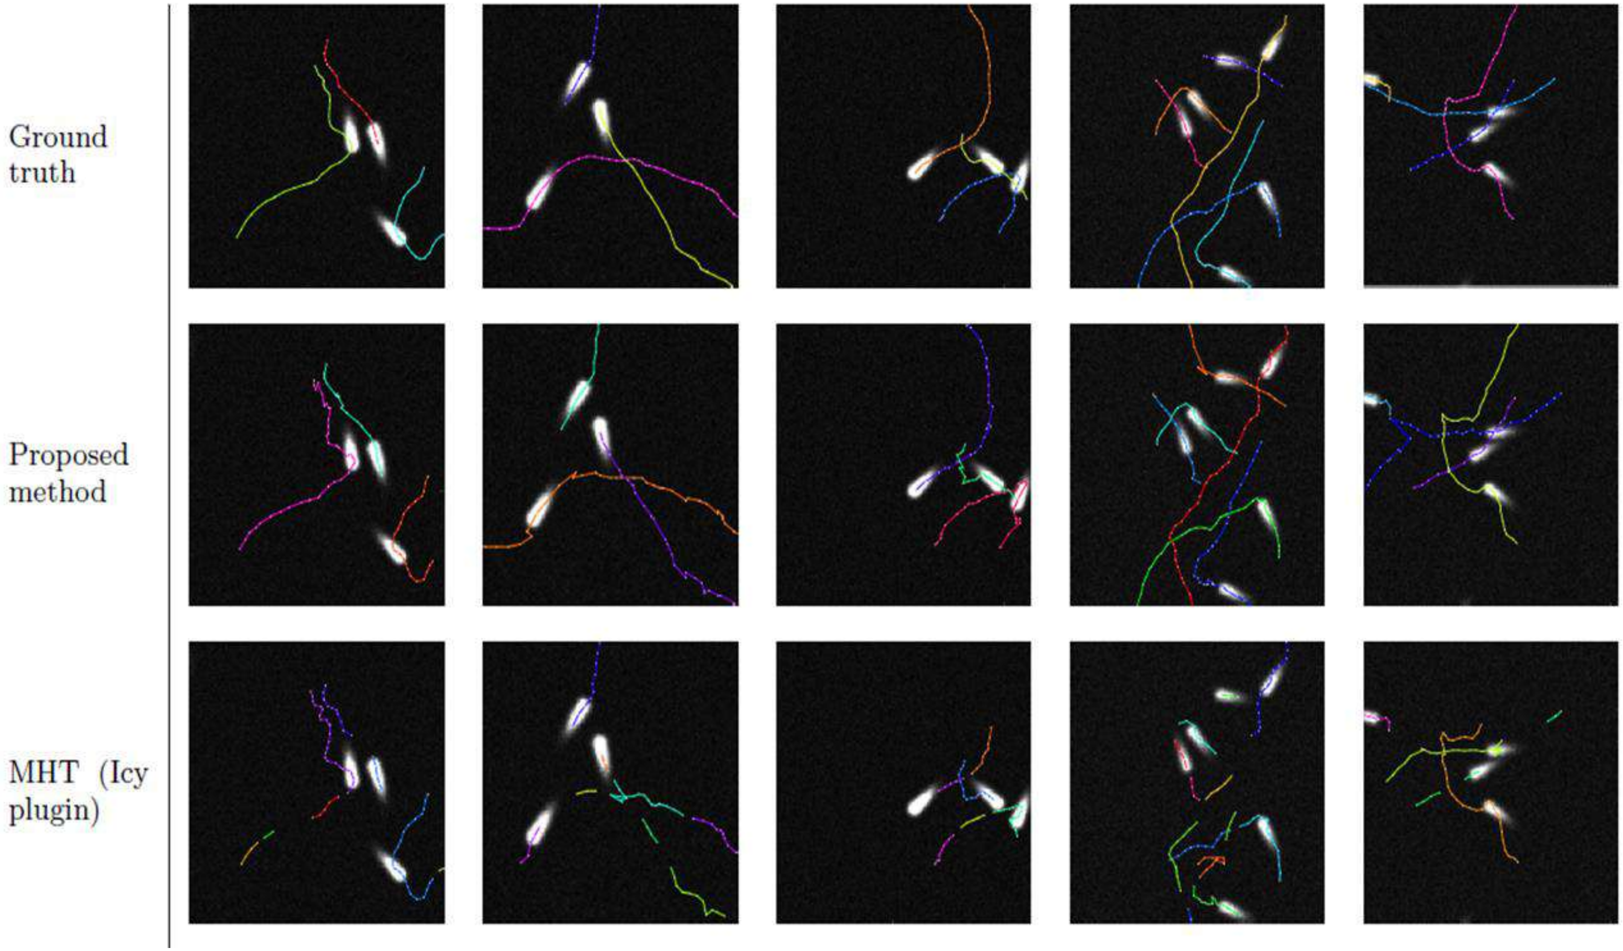

*Generated using the publicly available software ICY courtesy of the Quantitative Analysis Unit at the Pasteur Institute and J.-C. Olivo-Marin [Craciun 2016]

UAV data – low temporal frequency

COLUMBUS LARGE IMAGE FORMAT (CLIF) 2006 data set

Provided by:
The Sensor Data Management System, U.S. AirForce
https://www.sdms.afrl.af.mil

Satellite data – low temporal frequency

Tracking results © INRIA / AYIN

Average computation time: 12 sec / frame on a cluster with 512 cores
Image size: 1600 x 900 pixels

Satellite data – high temporal frequency

Tracking results

© INRIA / AYIN

Average computation time: 8 sec / frame on a cluster with 512 cores
Image size: 1600 x 900 pixels

Overview

- Models
  - Model formulation
  - Quality model vs. Statistical model
- Parameter learning
  - Linear programming
  - Parameter learning as a linear program
- Simulation
  - RJMCMC with Kalman inspired moves
  - Parallel implementation of RJMCMC
- Results
- Conclusions

Critical analysis

Advantages

- Detection of weakly contrasted objects
- Consistent trajectories
- Object interactions modeling
- Robustness to noise and data quality
- Good results on different data sets

Drawbacks

- Real-time processing only in exceptional cases
- Simple shape modeling

Conclusions

- Novel spatio-temporal marked point process model for the detection and tracking of moving objects
- Automatic or semi-automatic parameter estimation using linear programming
- Efficient parallel implementation of the RJMCMC sampler
- Good results on different types of data

Future work

- Detection and tracking of objects with various shapes
- Non-constant velocity model
- Comparison to/combination with Deep Learning techniques (CNN) [Li et al. - 19]
- Large scale sparse optimization for object detection [Boisbunon et al. - 14] and tracking
- On board implementation (mini satellite)

General conclusion

- The marked point process framework extends the application domain of Markov Random Field approaches:
  - Data taken into account at the object level
  - Geometrical information taken into account
- Markov random fields are still an efficient tool (depending on the image resolution)

References

- [Boisbunon2014] A. Boisbunon, R. Flamary, A. Rakotomamonjy, A. Giros and J. Zerubia. Large Scale Sparse Optimization for Object Detection in High Resolution Images. Proc. of MLSP, 2014.
- [Craciun2015] P. Craciun, M. Ortner and J. Zerubia. Joint detection and tracking of moving objects using spatio-temporal marked point processes. *IEEE Winter Conference on Applications of Computer Vision*, Hawaii, USA, Jan 2015.
- [Craciun2015] P. Craciun, M. Ortner and J. Zerubia. Integrating RJMCMC and Kalman filters for multiple object tracking. *GRETSI – Traitement du Signal et des Images*, Lyon, France, Sep 2015.
- [Craciun2015] P. Craciun. Stochastic geometry for automatic multiple object detection and tracking in remotely sensed high resolution image sequences., PhD thesis, Université de Nice - Sophia Antipolis, France, Nov. 2015.
- [Craciun2016] P. Craciun and J. Zerubia. Stochastic Geometry for Multiple Object Tracking in Fluorescence Microscopy. *IEEE International Conference on Image Processing (ICIP)*, Phoenix, USA, Sep 2016.

References

- [Descamps2008] S. Descamps, X. Descombes, A. Bechet and J. Zerubia. Automatic Flamingo Detection Using a Multiple Birth and Death Process. In Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Las Vegas, USA, March 2008.
- [Descombes2002] X. Descombes and J. Zerubia. Marked Point Process in Image Analysis.. IEEE Signal Processing Magazine, 19(5): pages 77-84, 2002.
- [Descombes2009] X. Descombes, R. Minlos, and E. Zhizhina. Object extraction using a stochastic birth-and-death dynamics in continuum. JMIV, 33:347359, 2009.
- [Eriksson2006] M. Eriksson, G. Perrin, X. Descombes and J. Zerubia. A comparative study of three methods for identifying individual tree crowns in aerial images covering different types of forests. In Proc. International Society for Photogrammetry and Remote Sensing (ISPRS), Marne La Vallee, France, July 2006.
- [Gamal2011] A. Gamal Eldin, X. Descombes, G. Charpiat, and J. Zerubia. A fast multiple birthand cut algorithm using belief propagation. Proc. of ICIP, pages 2813-2816, 2011.
- [Green1995] P. Green. Reversible jump Markov Chain Monte Carlo computation and Bayesian model determination. Biometrika, 82(4):711-732, 1995.

References

- [Lacoste2005] C. Lacoste, X. Descombes and J. Zerubia. Point Processes for Unsupervised Line Network Extraction in Remote Sensing. IEEE Trans. Pattern Analysis and Machine Intelligence, 27(10): pages 1568-1579, 2005.
- [Lafarge2008] F. Lafarge, X. Descombes, J. Zerubia and M. Pierrot-Deseilligny. Automatic Building Extraction from DEMs using an Object Approach and Application to the 3D-city Modeling. Journal of Photogrammetry and Remote Sensing, 63(3): pages 365-381, 2008.
- [Lafarge2010] F. Lafarge, X. Descombes, J. Zerubia and M. Pierrot-Deseilligny. Structural approach for building reconstruction from a single DSM. IEEE Trans. Pattern Analysis and Machine Intelligence, 32(1): pages 135-147, 2010.
- [Li2018] T. Li, M. Comer and J. Zerubia. A Connected-Tube MPP Model for Object Detection with Application to Materials and Remotely-Sensed Images. *ICIP 2018 – IEEE International Conference on Image Processing*, Athens, Greece, Oct 2018.
- [Li2019] T. Li, M. Comer and J. Zerubia. Feature extraction and tracking of CNN segmentations for improved road detection from satellite imagery. *ICIP 2019 – IEEE International Conference on Image Processing*, Taipei, Taiwan, Sep 2019.

References

- [Lieshout2000] M.N.M. van Lieshout. Markov Point Processes and Their Applications. Imperial College Press, 2000.
- [Mahler2003] R. Mahler. Multitarget Bayes ltering via rst-order multi target moments. IEEE Trans. on Aerospace and Electronic Systems, 39(4):1152-1178, 2003.
- [Ortner2007] M. Ortner, X. Descombes and J. Zerubia. Building Outline Extraction from Digital Elevation Models using Marked Point Processes. International Journal of Computer Vision, 72(2): pages 107-132, 2007.
- [Pace2011] M. Pace. Stochastic models and methods for multi-object tracking. PhD thesis, Université Sciences et Technologies, Bordeaux, France, 2011.
- [Papi2015] F. Papi, B.T. Vo, M. Bocquel, and B.N. Vo. Generalized labeled multi-Bernoulli approximation of multi-object densities. IEEE Trans. Signal Processing, 63(20): 5487-5497, 2015.
- [Perera2006] A.G.A. Perera, C. Srinivas, A. Hoogs, G. Brooksby, and W. Hu. Multi-object tracking through simultaneous long occlusions and split-merge conditions. Proc. of CVPR, 2006.
- [Perrin2005] G. Perrin, X. Descombes and J. Zerubia. Adaptive Simulated Annealing for Energy Minimization Problem in a Marked Point Process Application. In Proc. Energy Minimization Methods in Computer Vision and Pattern Recognition (EMMCVPR), St Augustine, Florida, USA, November 2005.

References

- [Saleemi2013] I. Saleemi and M. Shah. Multiframe many-many point correspondence for vehicle tracking in high density wide area aerial videos. IJCV, 104(2):198-219, 2013.
- [Stoica2004] R. Stoica, X Descombes and J. Zerubia. Stoica A Gibbs point process for road extraction in remotely sensed images. International Journal of Computer Vision, 57(2): pages 121-136, 2004.
- [Stoyan1987] D. Stoyan,W.S. Kendall, and J. Mecke. Stochastic Geometry and its Applications. John Wiley and Sons, 1987.
- [Stoyan1994] D. Stoyan, and H. Stoyan. Fractals, Random Shapes and Point Fields: Methods of Geometrical Statistics. John Wiley and Sons, 1994.
- [Verdie2013] Y. Verdie. Urban scene modeling from airborne data, PhD thesis, Université de Nice - Sophia Antipolis, France, 2013.
- [Vo2013] B.-T. Vo and B.-N. Vo. Labeled random nite sets and multi-object conjugate priors. IEEE Trans. Signal Processing, 61(13):3460-3475, 2013.
- [Vo2014] B.-N. Vo, B.-T. Vo, and D. Phung. Labeled random finite sets and the Bayes multitarget, tracking filter. IEEE Trans. Signal Processing, 62(24):6554-6567, 2014.
- [Yu2009] Q. Yu and G. Medioni. Multiple-target tracking by spatio-temporal Monte Carlo Markov chain data association. IEEE Trans. PAMI, 31(12):2196-2210, 2009.

Thanks to:

- CNES, IFN, IGN, Tour du Valat, Airbus DS, Satellite Image Corporation, for providing the data.

- DGA, MAE, CNES, IGN, BRGM, Airbus DS, ECP and INRIA, for partial financial support.

Thank you!

Josiane.Zerubia@inria.fr

Hammersley-Clifford Theorem

A MRF verifying a positivity constraint can be written as a Gibbs field:

$$P(X) = \frac{1}{Z} \exp - \left[\sum_{c \in C} V_c (x_s, s \in S) \right]$$

$S = \text{all the pixels}$

$C = \text{all the cliques associated to the neighborhood } v$

Markov process

- A point process density $\boldsymbol{f} : \boldsymbol{N}^f \rightarrow [0,\infty[$

  is Markovian under the neighborhood relation ~ if and only if there exists a measurable function $\phi : \boldsymbol{N}^f \rightarrow [0,\infty[$ such that:

$$\boldsymbol{f}(\boldsymbol{x}) = \alpha \prod_{\text{cliques } \boldsymbol{y} \subseteq \boldsymbol{x}} \phi(\boldsymbol{y})$$

for all $\boldsymbol{x} \in \boldsymbol{N}^f$

Stability

• Condition required for proving the convergence of Markov Chain Monte Carlo sampling methods.

• A point process defined by its f(.) w.r.t. a reference measure π_ν(.) is locally stable if there exists a real number M such that:

$$\boldsymbol{f}\left(\boldsymbol{x} \cup \{\boldsymbol{u}\}\right) \leq \boldsymbol{M}\boldsymbol{f}\left(\boldsymbol{x}\right), \forall \boldsymbol{x} \in \boldsymbol{N}^{f}, \forall \boldsymbol{u} \in \chi$$

Data-association based methods

NN-app

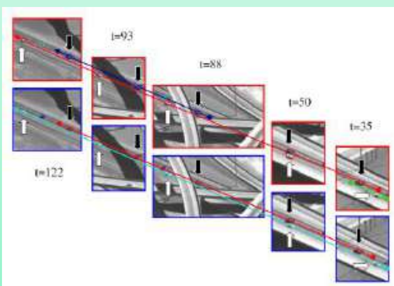

[Perera2006]

DD-MCMCDA

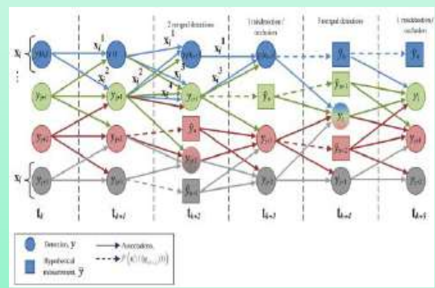

[Yu2008]

MHT

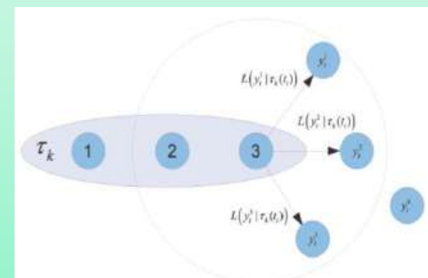

[Saleemi2013]

RFS-based methods

RFS

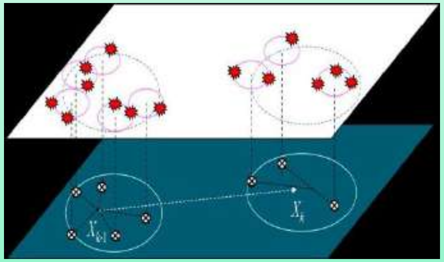

[Mahler2003]

PHD

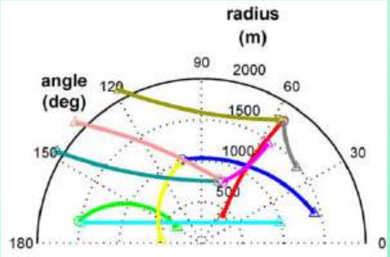

[Vo2005]
[Vo2006]
[Pace2011]

L-RFS / GLMB

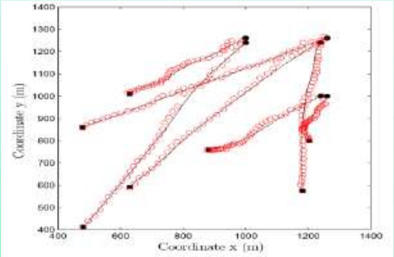

[Vo2013]
[Vo2014]
[Papi2015]