The SpaceFusion* project: goals and preliminary results

André Jalobeanu
PASEO Research Group
MIV team @ LSIIT, Illkirch, France

**Projet ANR "Jeunes Chercheurs" 2006-2008*

Outline

- Introduction
- Objectives
  - Astronomy: 2D image (sky map)
  - Remote sensing: 2D reflectance map
  - Small bodies: 3D surface geometry
  - Earth/planetary sciences: reflectance and topography
- Proposed approach
  - Bayesian inference from multiple observations
  - Accurate forward modeling
  - Preliminary results: validation on 1D signal fusion
- Collaborations
  - Deformation fields in Earth Sciences
  - Dempster-Shafer fusion theory

Why multisource data fusion?

Problem: lots of data, same object!

Usually, images are recorded with various:
- **pose parameters (position, orientation)**
- **sensors (resolution, noise, bad pixels)**
- **observing conditions (transparency, seeing)**
- **instruments (PSF, distortions)**

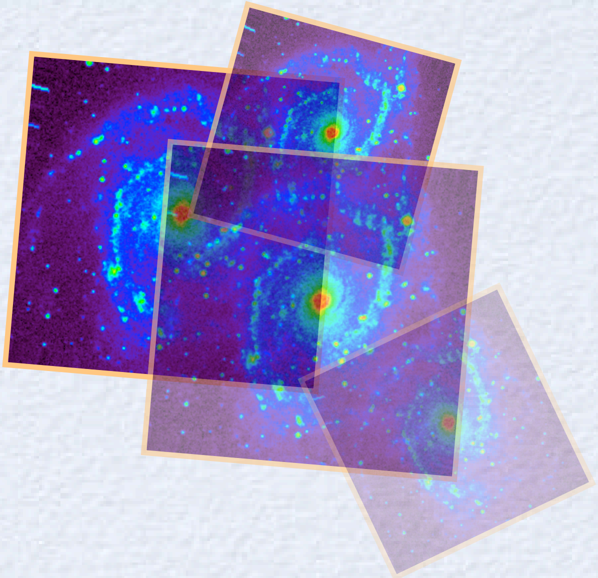

- Multisource data fusion
  - Optimally **combine** all observations into a single model
    Co-add or build a mosaic, depending on the overlap
  - **Preserve** all the information from the original data set
    - *Increase resolution if needed*
    - *Compute the uncertainties*
    - *Reconstruct the 3D geometry if required*
  - **Enhance** the image quality (optional)
    Denoise or deblur depending on the degradation

The SpaceFusion project

Projet ANR “Jeunes Chercheurs”
3-year grant, Jan 2006 - Dec 2008

Name	**Position, lab**	**time**
André Jalobeanu	CR, LSIIT/MIV	90%
Christophe Collet	PU, LSIIT/MIV	40%
Mireille Louys	MCF, LSIIT/MIV	40%
Fabien Salzenstein	MCF, InESS	40%
Françoise Nerry	CR, LSIIT/TRIO	20%
Albert Bijaoui / Eric Slezak	A/AA, OCA	10%
Bernd Vollmer	A, Obs. Strasbourg	10%
~~Mickaël Ferrand~~	~~PhD, LSIIT/MIV~~	~~70%~~
Jorge A. Gutiérrez + ?	PhD+	20 mo total

Objectives

- ☑ Produce a corrected, super-resolved image in astronomy
- ☑ Reconstruct a reflectance function in remote sensing
- ☑ Recover the geometry of small bodies and planetary surfaces
- ☑ Reconstruct both reflectance and topography in Earth/Space Sciences

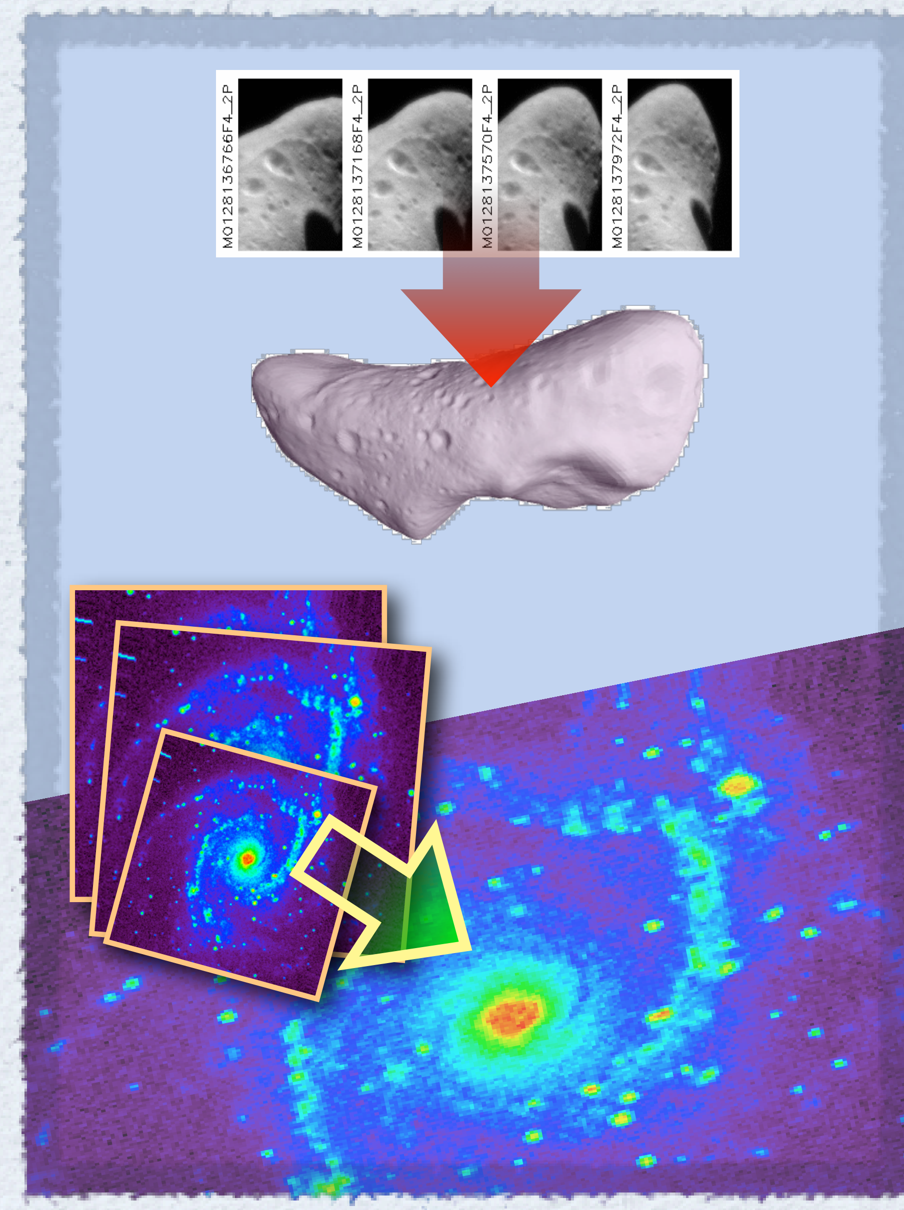

Astronomy: 2D image reconstruction

DeepSkyFusion
Multisource data fusion and 2D super-resolution

Astronomy & Astrophysics

- Input:
  - Multiple images (single band, multispectral or IFS) ***Virtual Observatory***
  - Optical, UV, IR / calibrated or not / missing or corrupted data
- Output:
  - Single model, 2D (image-like), well-sampled
  - Uncertainties (simplified inverse covariance)
  - If applicable, spatial and spectral super-resolution

Remote sensing:
2D reflectance reconstruction (3D space)

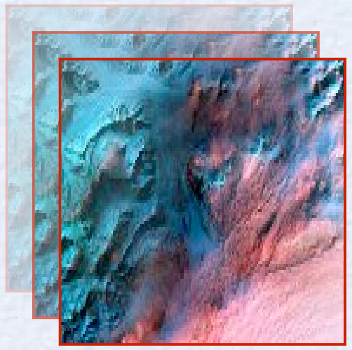

ReflectanceFusion
Multisource data fusion for flat terrain BRDF recovery

Remote Sensing, Planetary Imaging

- Input:
  - Multiple images (single band, multispectral or hyperspectral)
  - Optical, IR / calibrated or not / missing or corrupted data
- Output:
  - Single model, 2D (image-like) reflectance map, well-sampled
  - Uncertainties (simplified inverse covariance)
  - If applicable, spatial and spectral super-resolution

Small bodies:
3D surface recovery (geometry only)

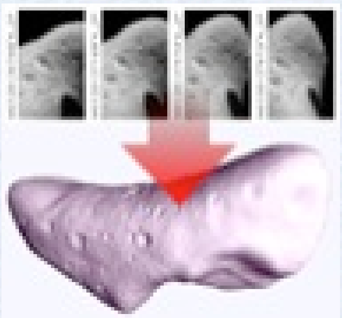

3DShapeInference
3D shape recovery via Bayesian inference

Planetary Imaging (small bodies and planets)

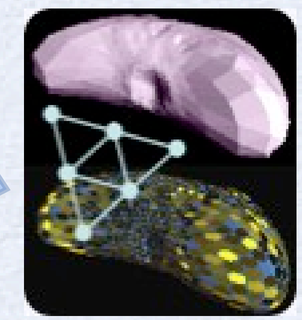

SurfaceModelRender
Accurate rendering and modeling of natural 3D surfaces

- Input:
  - Multiple images (single band)
  - Optical, IR / calibrated or not / missing or corrupted data
- Output:
  - Single model, 3D mesh (planar or spherical topology)
  - Uncertainties (simplified inverse covariance)

Earth & Planetary Sciences: reflectance and topography recovery

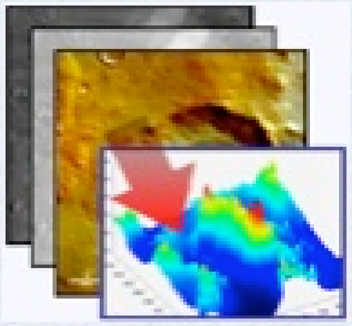

3DSpaceFusion
Multisource data fusion, 3D surface recovery and super-resolution

Planetary Imaging

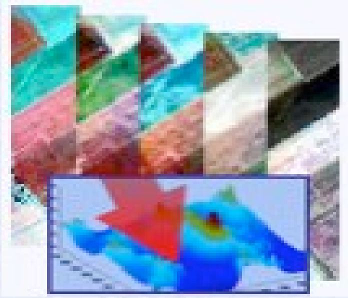

3DEarthFusion
Multisource data fusion, 3D surface recovery, BRDF inference and super-resolution

Remote Sensing

- Input:
  - Multiple images (single band, multispectral or hyperspectral)
  - Optical, IR / calibrated or not / missing or corrupted data
- Output:
  - Single model, 3D mesh + well-sampled reflectance map
  - Uncertainties (simplified inverse covariance)
  - If applicable, spatial and spectral super-resolution (reflectance)

The proposed approach

- Use Bayesian inference to recover a single object from all observations
- Provide uncertainty estimates, allow for recursive data processing
- In 2D: recover a well-sampled image, possibly super-resolved
- Check the validity of this approach in 1D (*first results*)

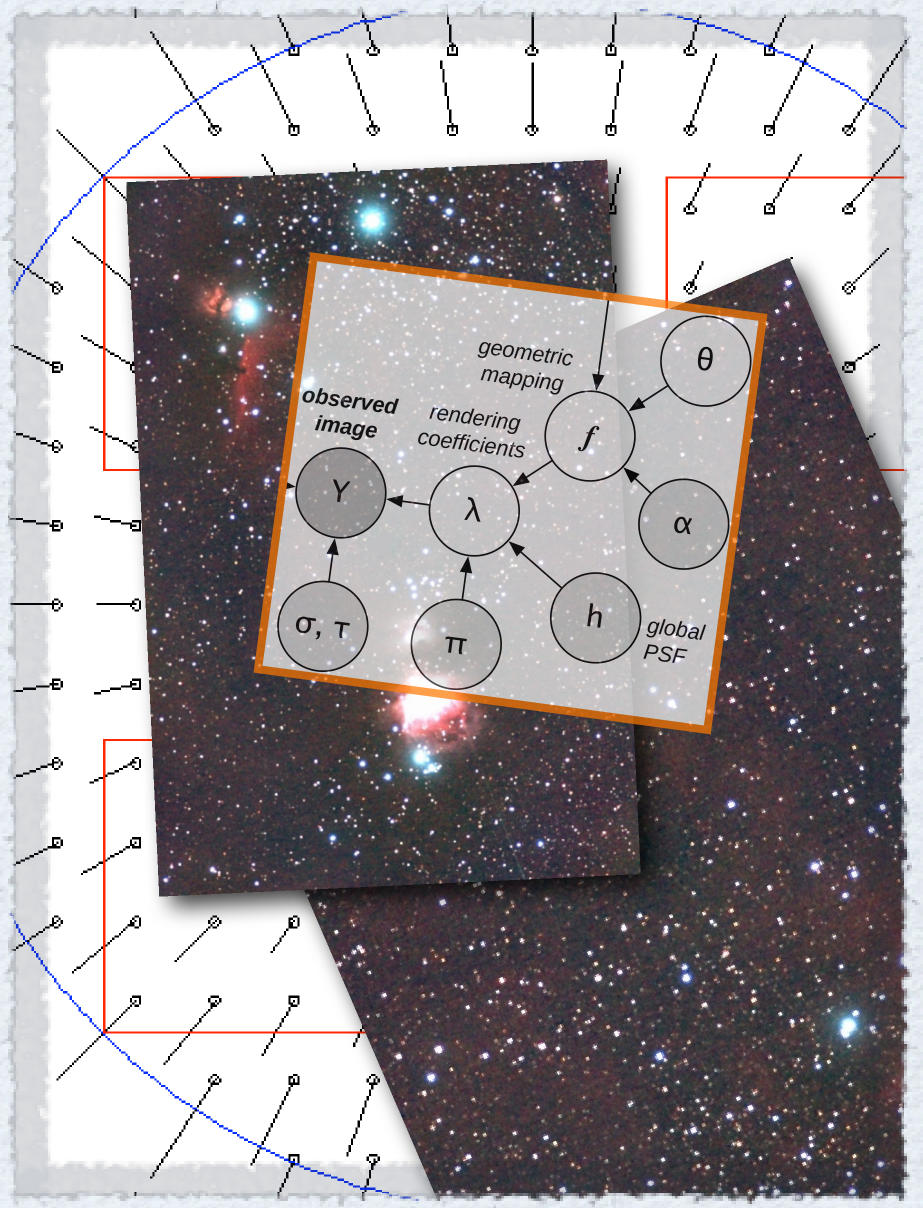

Bayesian Vision

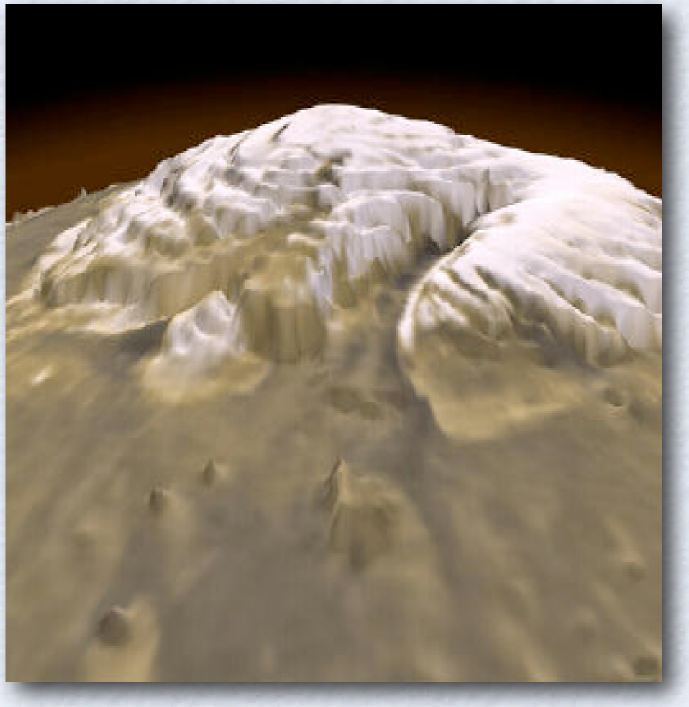

- **Computer vision:**
  model reconstruction from multiple observations,
  inverse problem of rendering

- **Bayesian inference**
  applied to this inverse problem:
  everything is described by random variables

- **Data fusion** into a single model becomes a **parameter estimation** problem

- It can be solved by existing efficient **optimization** techniques

Bayesian inference from multiple observations

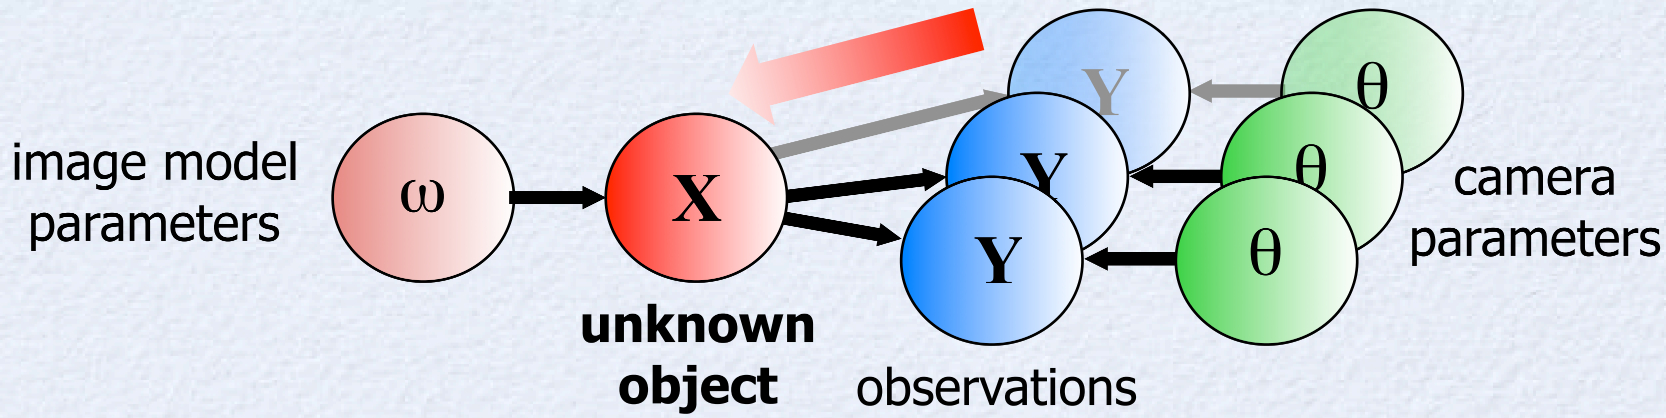

- Probabilistic approach
  - Modeling:
    - **Object modeling** (image, 3D geometry, reflectance map...)
    - **Image formation** = forward model (rendering)
  - Bayesian inference:
    - **Estimate the optimal object given all observations:** mode or mean of the posterior distribution
    - **Integrate w.r.t. all nuisance variables** (marginalization)
    - **Evaluate the uncertainties:** covariance matrix (Gaussian approx. of the posterior distribution)
    - Model **selection** and **assessment**

Accurate forward modeling

- 2D object, 2D space
  - Resampling, account for deformations & PSF (possibly irregular sampling grid: IFS)

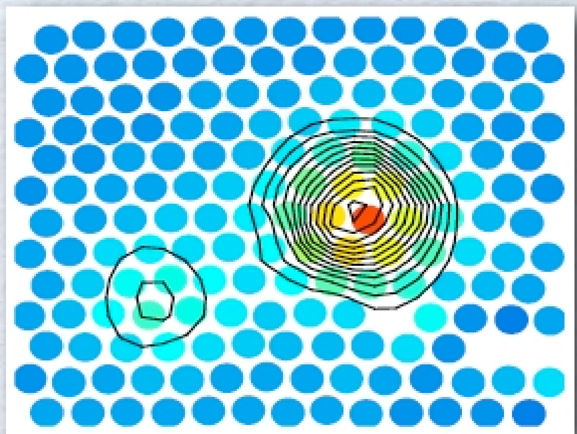

- 2D object, 3D space
  - Resampling, account for perspective transformation, deformations & PSF

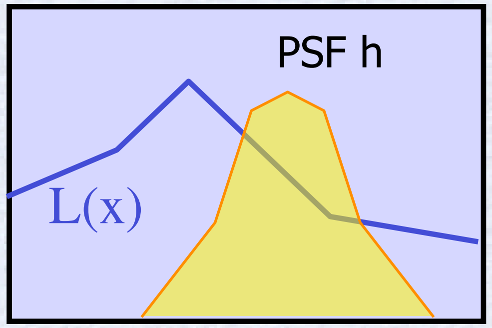

- 3D object, 3D space
  - Rendering in the object space, account for occlusions, shadows, perspective, deformations & PSF

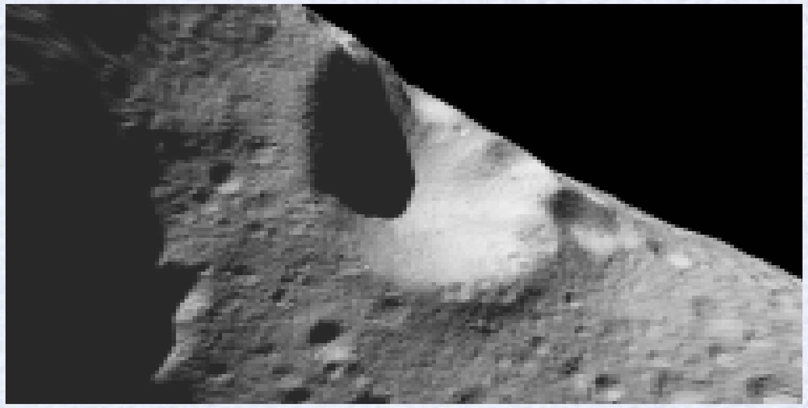

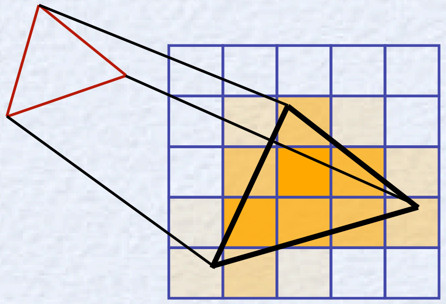

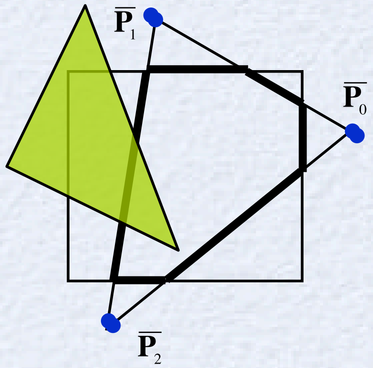

First goal: 2D image reconstruction

Goal: combine N images
(different blur, resolution, FOV, noise...)
into a single object: pixel values + uncertainties

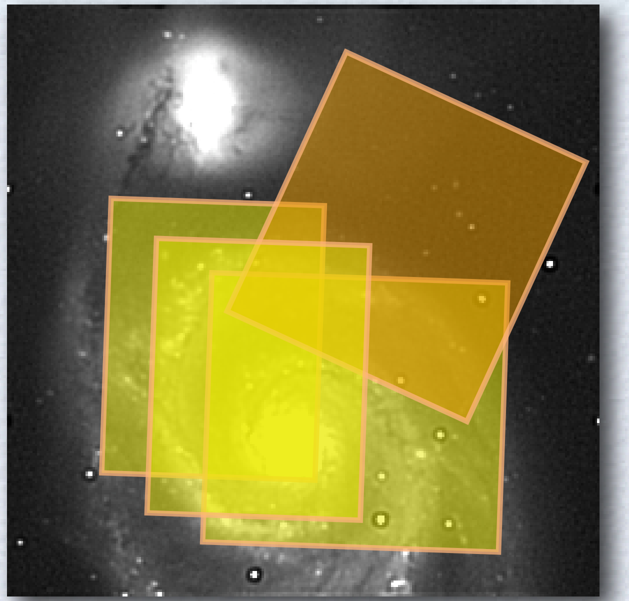

Preserve the information from the original data set: photometry and astrometry

Related problems

- Image registration: external camera parameter estimation
- Image modeling for regularization purposes
- Prior model parameter estimation
- Model selection (*e.g. scene model resolution*)

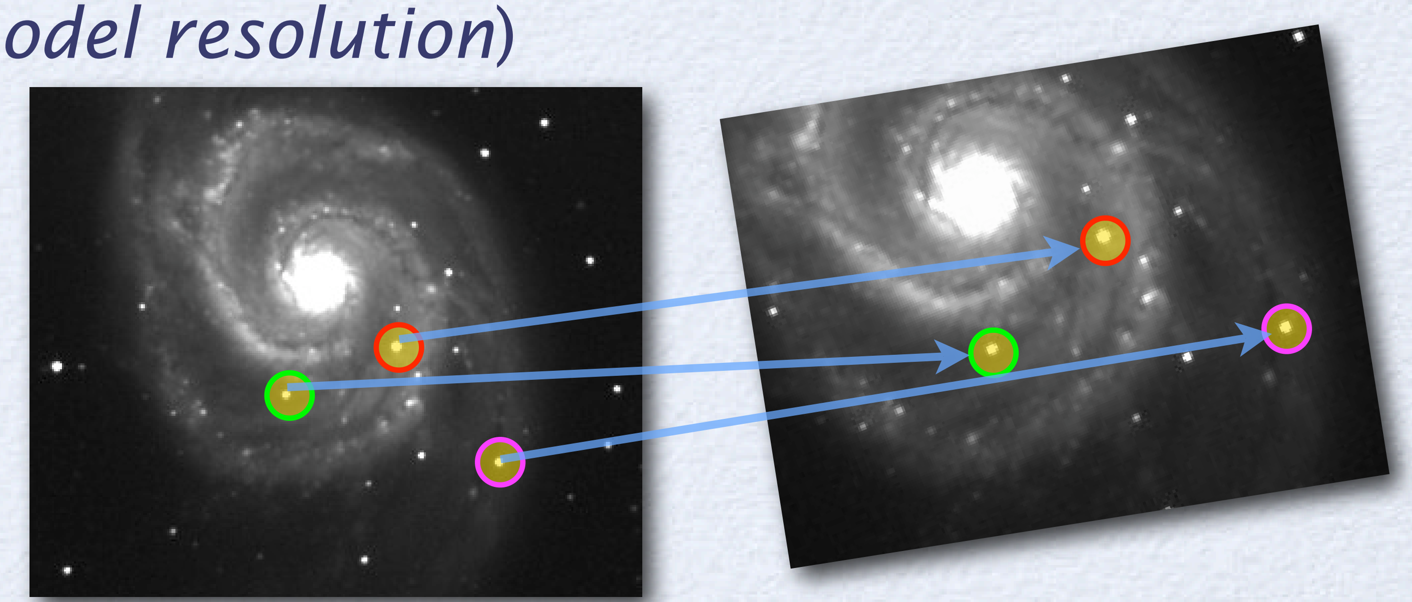

Simplified graphical model

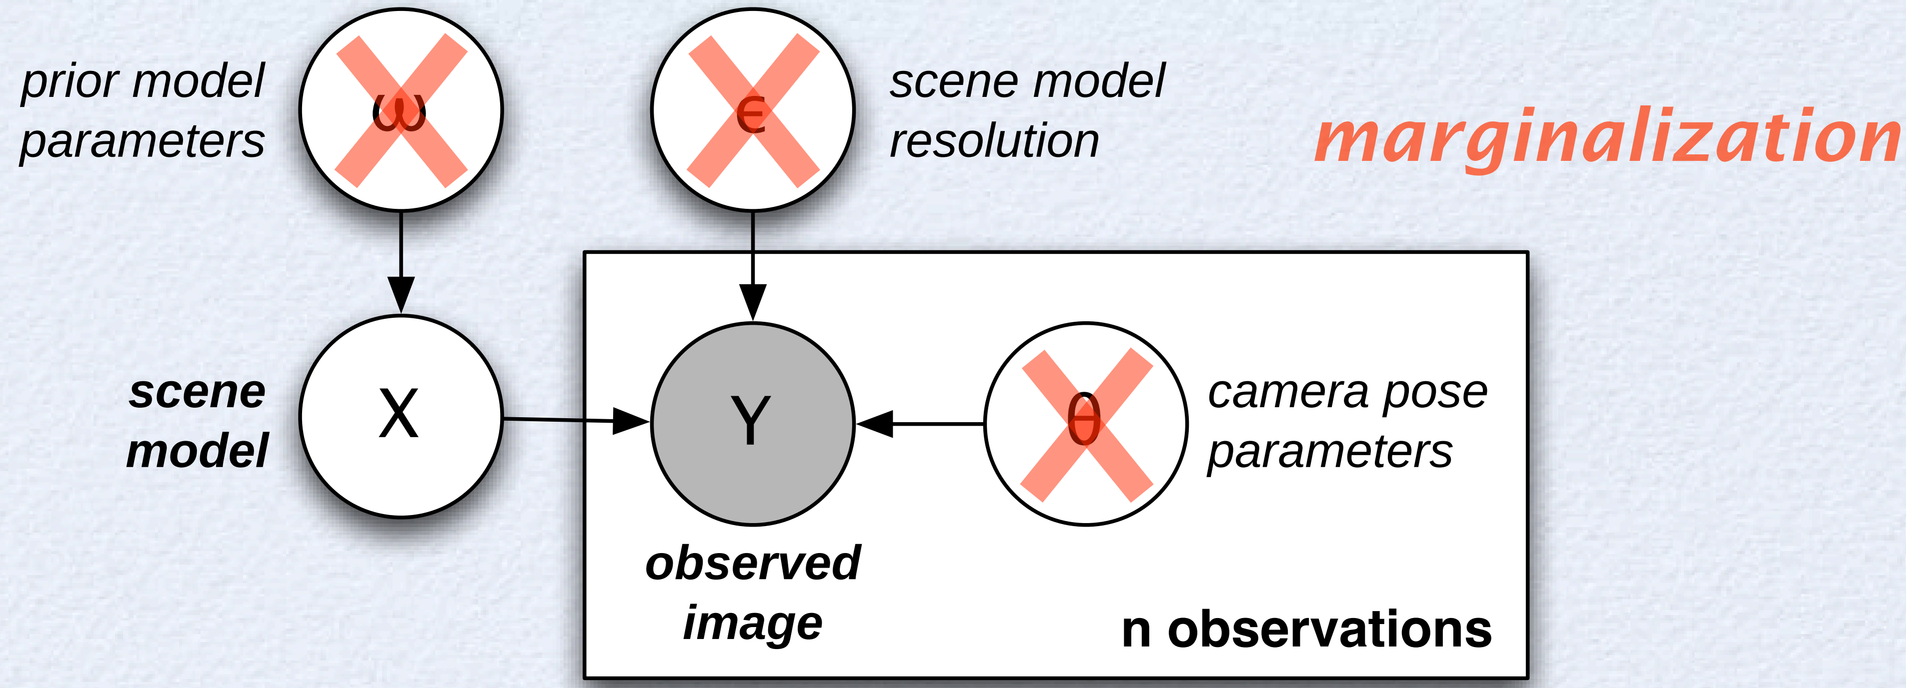

Directed graphical models:

Node = set of random variables
No incoming arrow: prior density

Arrow = dependence
Set of incoming arrows: conditional density

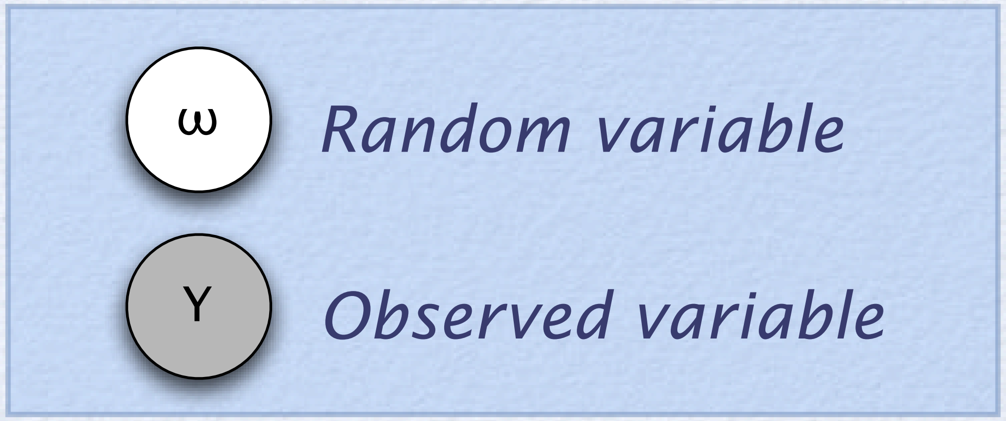

Joint distribution: $P(X,Y,\omega,\theta,\epsilon)=P(\omega)P(\theta)P(\epsilon)P(X|\omega)P(Y|X,\theta,\epsilon)$

Posterior marginal: $P(X|Y)\propto\int P(X,Y,\omega,\theta,\epsilon)\, d\omega d\theta d\epsilon$

Full graphical model

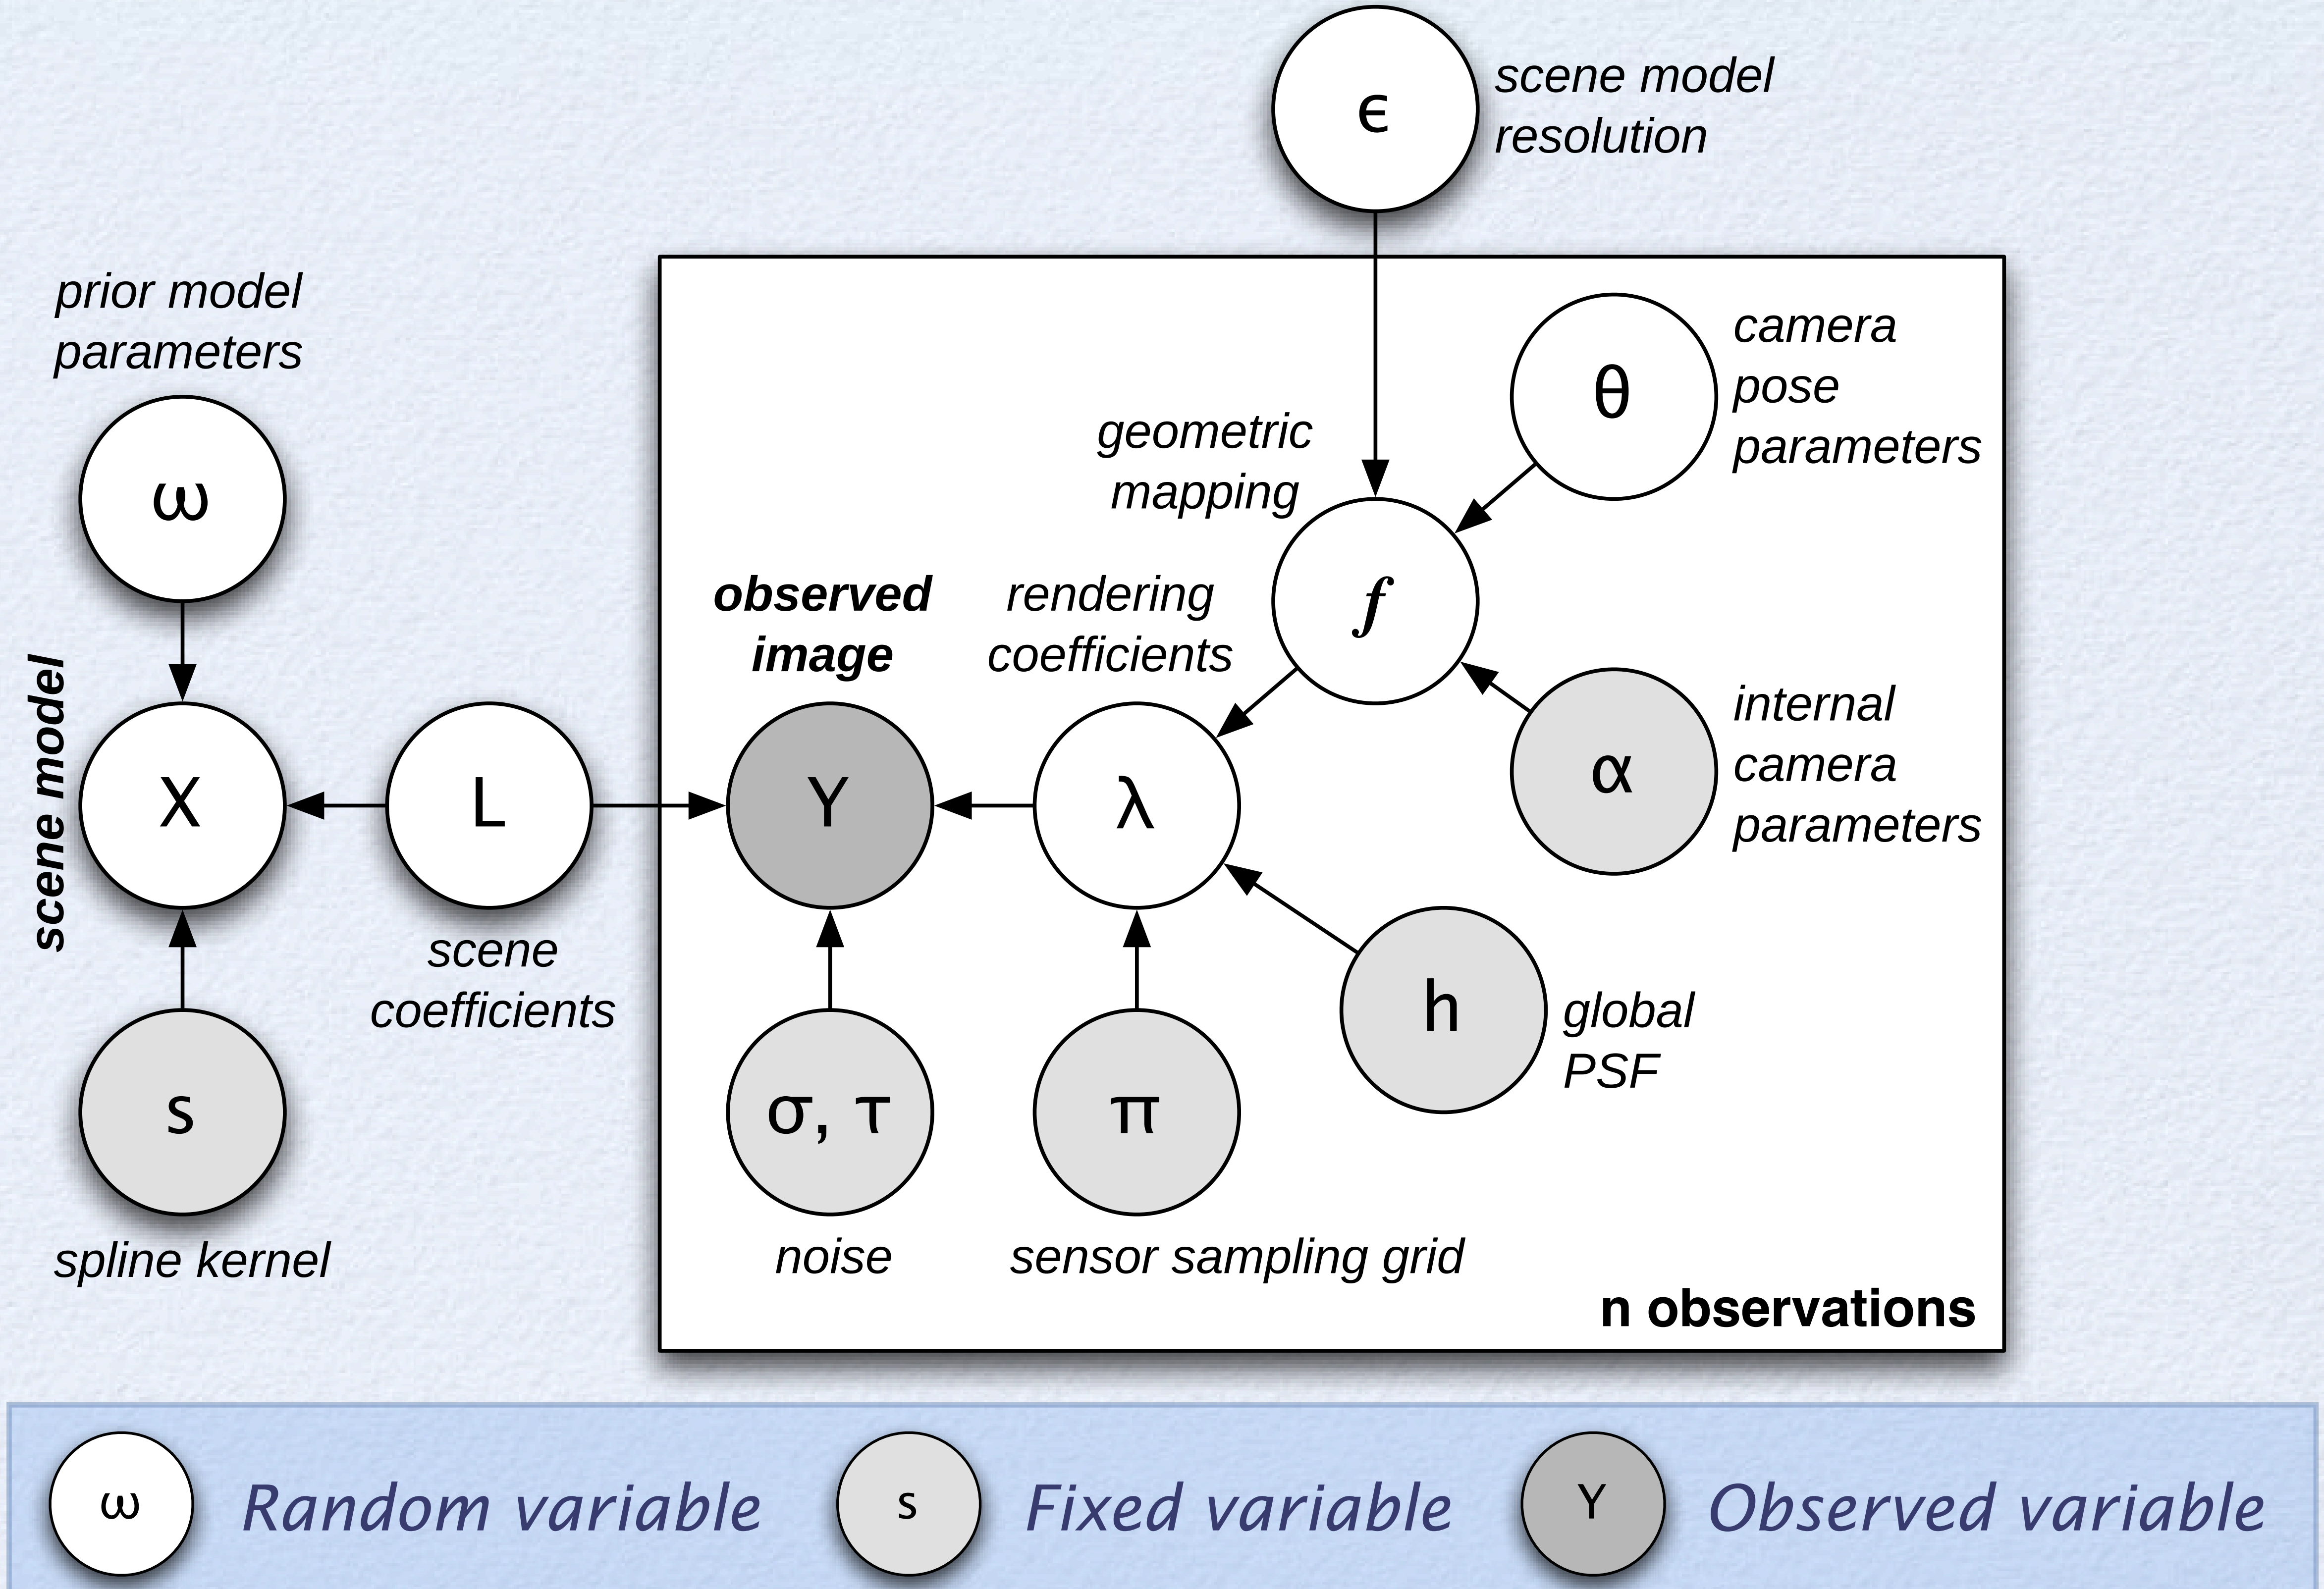

Image model

- Model of the unknown object (*2D image*)
  - Choose an appropriate **parametrization** and topology
    - Rectangular or hexagonal **lattice**
    - Sampling **grid size** ϵ chosen to avoid undersampling

Output pixel size < input PSF FWHM / 2

  - Understand the **sampling** theorem!
    - Don't try to go beyond the Shannon **sampling limit** (frequency cut-off)
    - Choose an correct target: **near-optimal sampling, band-limited**
      *The **BSpline-3 kernel** provides a good approximation*
  - **Constrain** and stabilize this **inverse problem**
    (can be ill-posed in some cases, *e.g. deblurring*)
    - Use **smoothness priors** to avoid noise amplification (oversampled areas will undergo a deconvolution even if we just want data fusion...)
    - Use **efficiently designed prior models** (*e.g. multiscale, wavelets*) to help preserve useful information while filtering the noise, and remain computationally effective

Noise modeling

Probabilistic image formation scheme

- Gauss+Poisson+Quantization noise

  $P(Y_p \mid I_p) = \text{Gauss}(0,a^2) \ast \text{Poisson}(bI_p) \ast U(0,1)$

  - Gaussian noise: thermal
  - Poisson noise: counting process
  - Uniform noise: quantization

  - **Approximation**: Gaussian, spatially variable variance $\sigma I + \tau$ (depends on each sensor, possibly spatially variable σ and τ)

  $P(Y_p \mid I_p) \approx \text{Gauss}(I_p, v_p)$ with $v_p = \tau I_p + \sigma^2$

- Pixel & sensor **indep**. assumption: $P(Y \mid L) = \Pi_{pn} P(Y_p^n \mid L)$

First results: 2X super-resolution (1D signals)

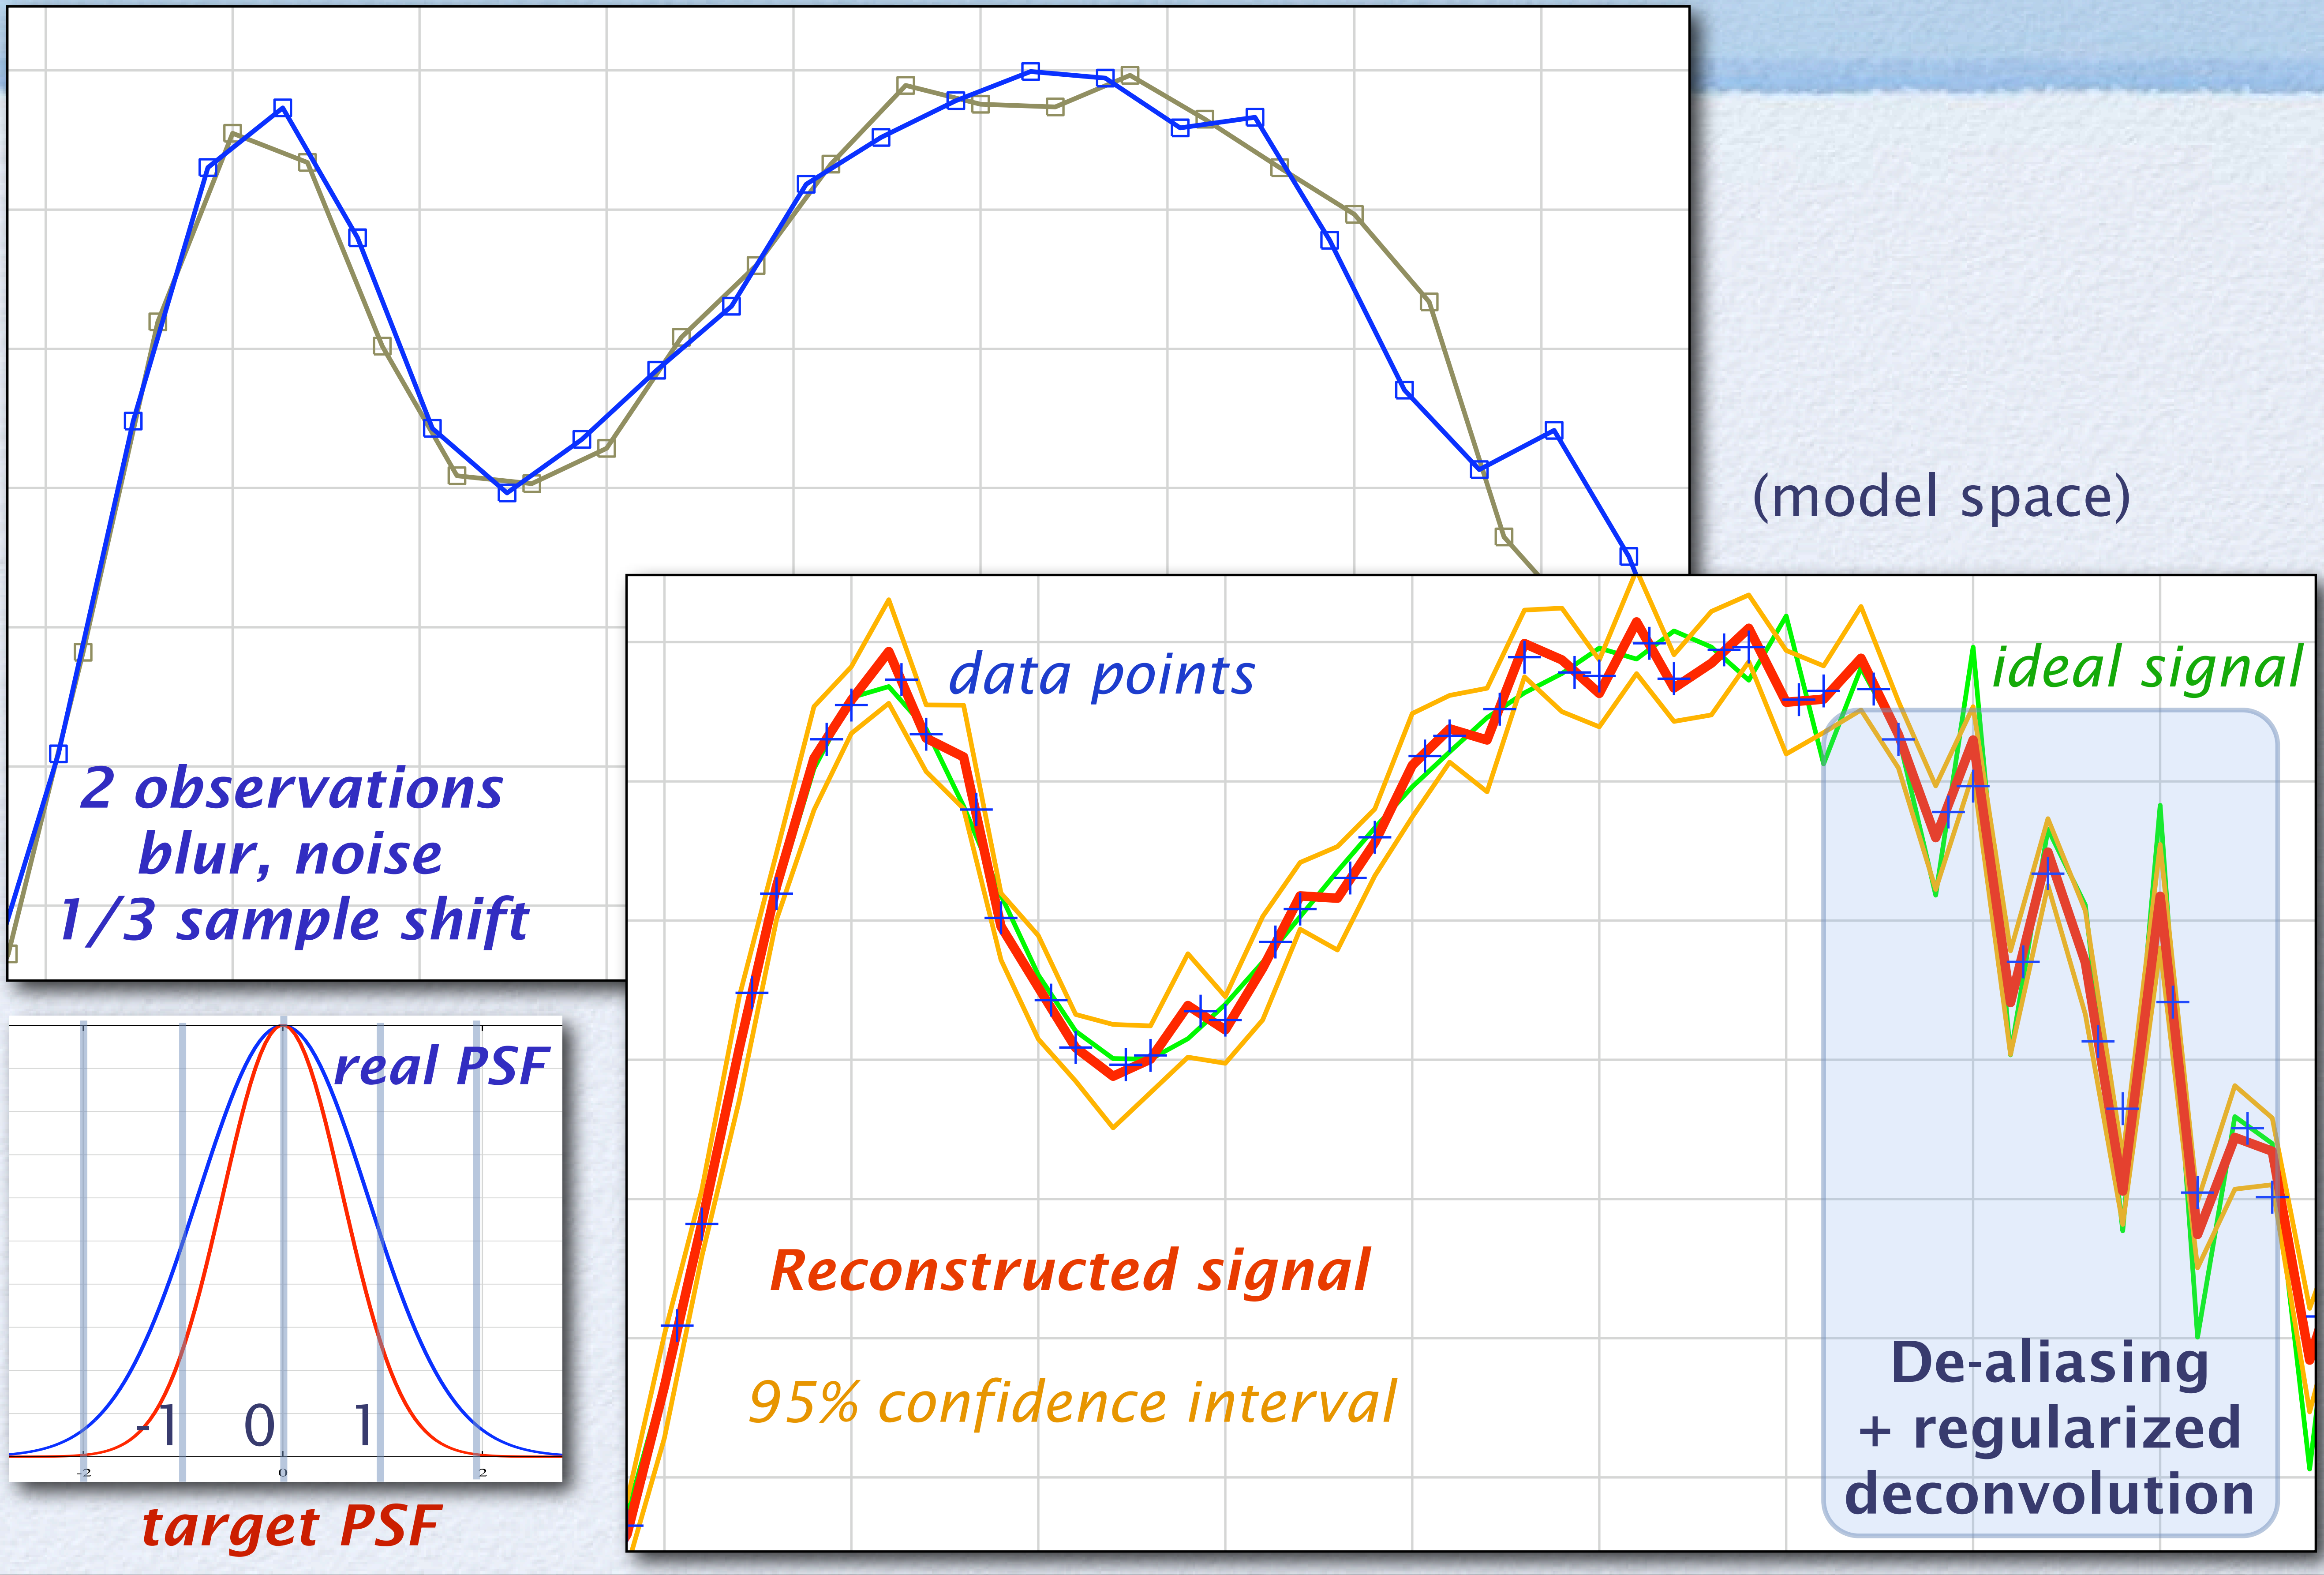

Computing and propagating uncertainties

- Inverse covariance matrix computation
  - **Second derivatives** of the energy U(X) at the optimum

$$[\Sigma_X^{-1}] = S^{-1}[\nabla^2 U]S$$

$$[\Sigma_X^{-1}] = 2\omega D^T D + \sum_{n,p} \frac{1}{v_p^n}(S^{-1}\lambda_p^n)\,(S^{-1}\lambda_p^n)^T$$

BSpline coefficients computation

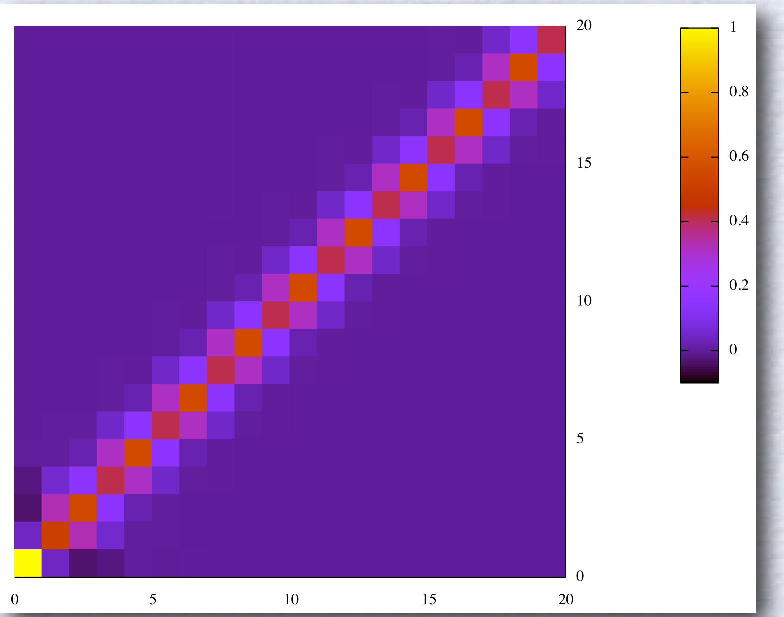

  - **Sparse** matrix

    (interaction range depends on the size of h)

- Recursive processing and uncertainty propagation
  - Use the simplified **posterior** (mean, approx. inv. covariance) as a **prior** density for subsequent data processing
  - Recursive (vs. batch) data fusion: allow for model **updates**

$$\Phi(L)^{(k+1)} = L^T\, S[\tilde{\Sigma}_X^{-1\,(k)}]S\, L$$

Inversion with unknown parameters

- **Full** Bayes P(X | Y) - ***intractable in general***
- **Empirical** Bayes
  - First compute P(ω,θ | Y) - e.g. marginal MAP
  - Plug in the estimate and maximize P(X | Y,ω,θ)
  - **Good approximation** of the full Bayes if **P(ω,θ | Y) is peaked**, otherwise the data is used twice (learning/inference)...
- Parameter inference with **E-M**
  - Goal: P(ω,θ | Y); consider X as the missing data
  - Standard E-M: maximize P(ω,θ | Y), variational E-M: inference
  - **Simpler**, but more sensitive to local optima than exact marginal.
- **Joint** MAP
  - Compute the joint MAP related to P(X,ω,θ | Y)
  - Usually done by alternate optimizations X,ω,θ (sub-optimal)
  - **Simple** but unstable, biased, not recommended

Remarks

Special cases of the proposed framework:

- **Spline** interpolation in the presence of **noise** [Unser & Blu 05]
  - Single observation (no fusion):
    sampling resolution = model resolution
  - Assumed blur kernel = spline kernel
  - Gaussian noise

- **Spline** interpolation and **irregular** sampling? [Arigovindan 05]
  - Similar assumptions (single, spline, Gauss)
  - Irregular sampling in the sensor space

The proposed approach is a generalization to multiple observations, arbitrary noise, arbitrary geometry

Uncertainties are provided, recursive inference is made possible

Collaborations

- Validation: specialists from astronomy and remote sensing
- Inferring deformation fields from satellite images (Earth Sciences)
- Links between Bayesian and Dempster-Shafer theory

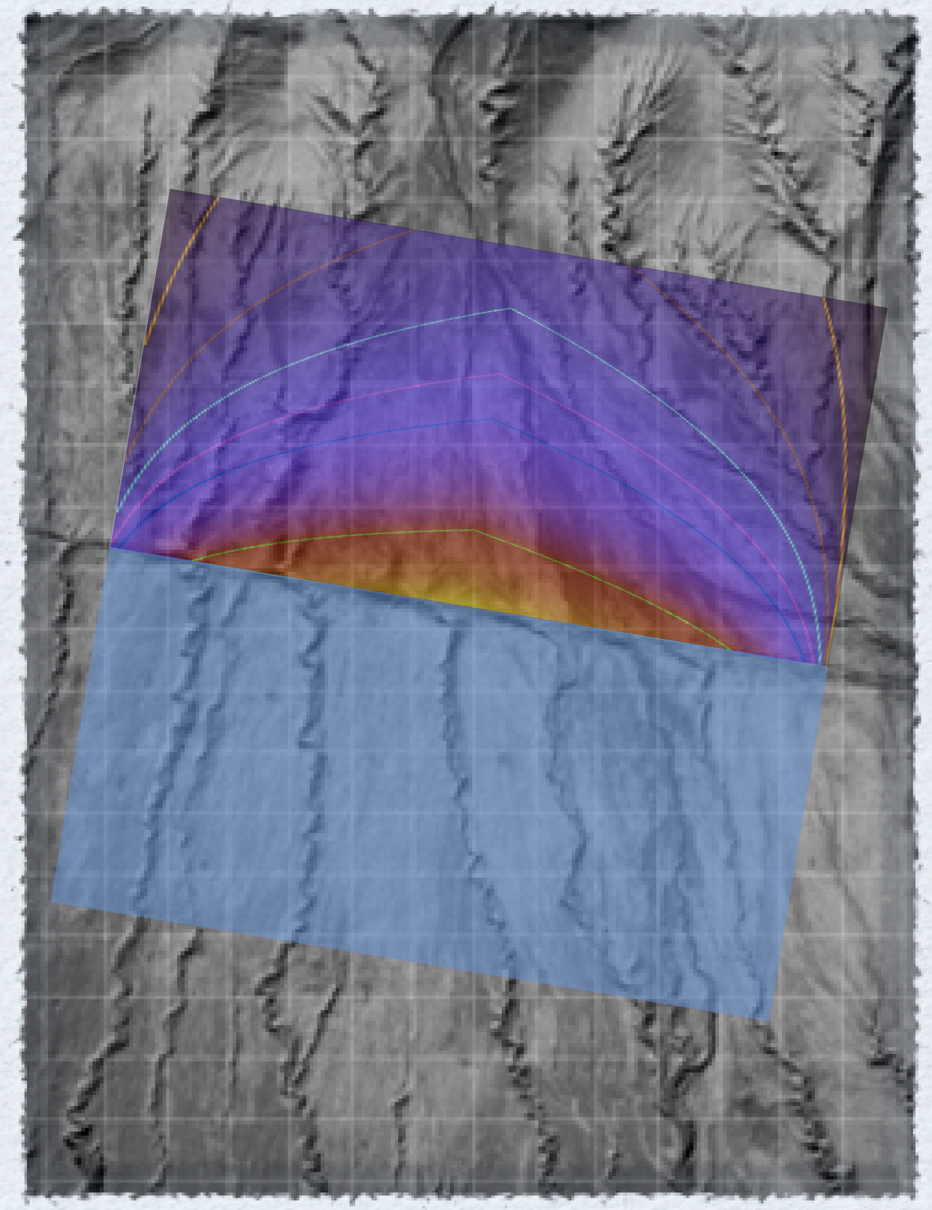

Validation in astronomy & remote sensing

- **Astronomy**
  Check the validity of the models (e.g. priors on images, sensor and instrument physics and geometry), the good match between our goals and the astronomer's needs
  - **Observatoire de Strasbourg**
    B. Vollmer
  - **Observatoire de la Côte d'Azur (OCA)**
    A. Bijaoui, E. Slezak

- **Remote sensing**
  Check the validity of the models (e.g. hyperspectral image and reflectance function models, sensors & PSFs), the good match between our goals and the specialist's needs
  - **LSIIT / TRIO (remote sensing team @ LSIIT)**
    F. Nerry

Deformation fields in Earth Sciences

D. Fitzenz, J. Van der Woerd - **IPG Strasbourg**

- Infer the parameters of the geometric transform

  2 images: one before, one after earthquake or deformation
  - Deformation field = spatially variable translation
  - Challenge: subpixel accuracy (0.1 pixel to detect a 10 cm shift)
  - Use a smoothness prior allowing for discontinuities on segments (faults)

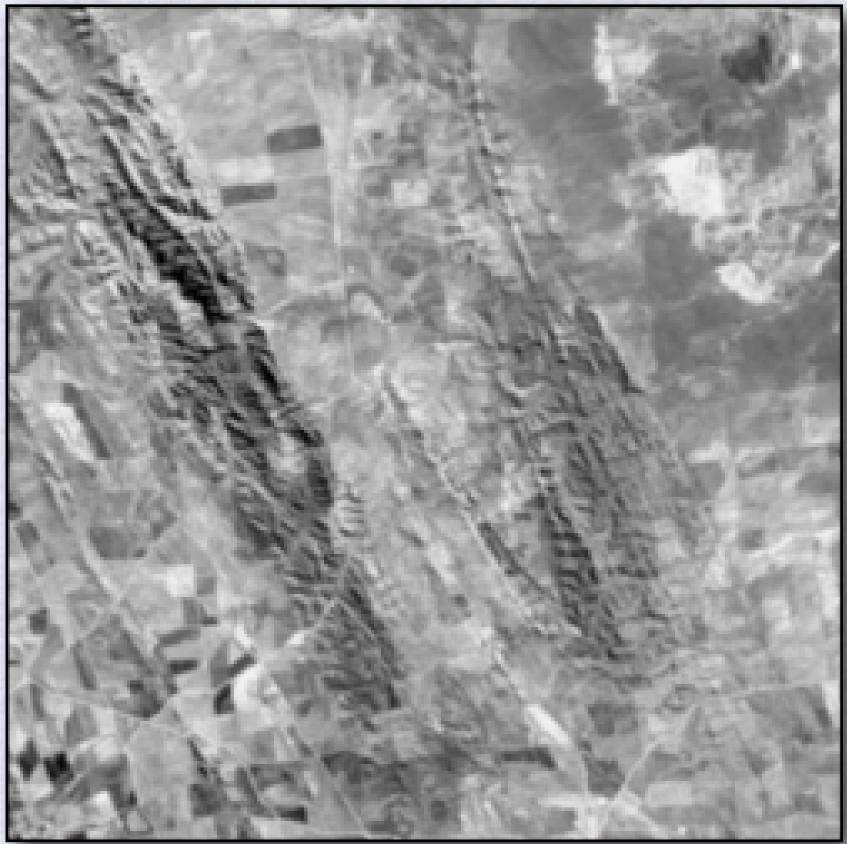

Before EQ (simulation)

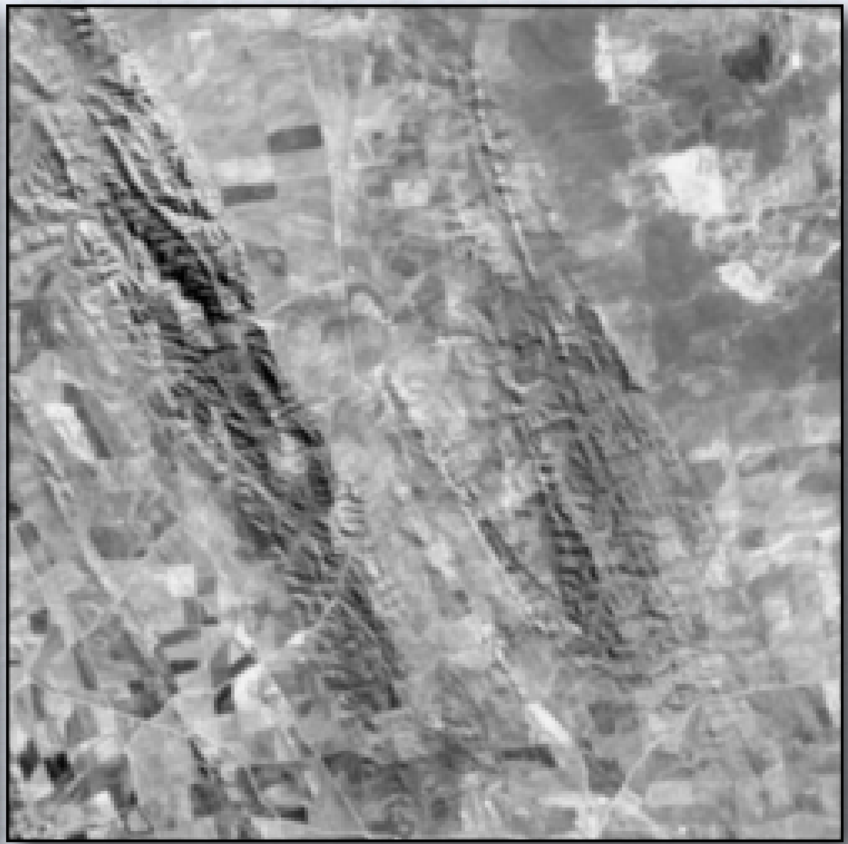

After EQ (simulation)

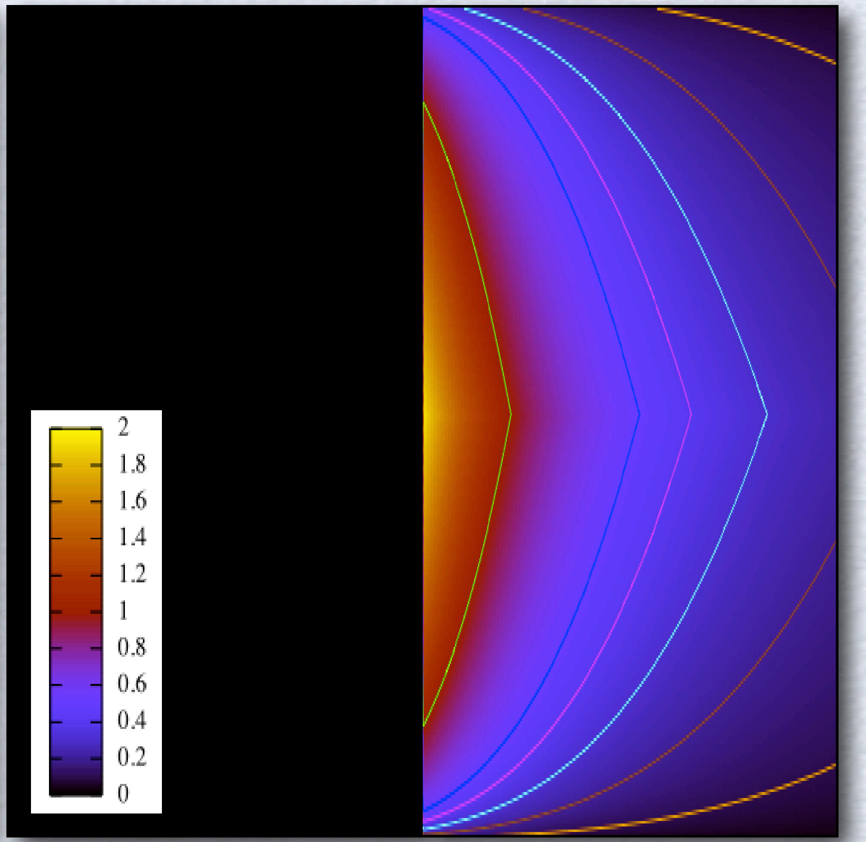

Deformation field (y)

Dempster-Shafer fusion theory: sensor reliability

Pieczinski - **INT Evry**

- Other approaches to data fusion?
  - **Dempster-Shafer** theory of evidence: [Shafer 76]
    defined for **discrete** variables (e.g. hard classification),
    assign/combine degrees of belief (epistemic plausibilities)

    More general than Bayes,
    Better handling of what is "non-informative"

  - How to take into account the global reliability of each sensor?
  - Is the Bayesian approach the best answer to missing data or incomplete model knowledge?
  - Can we switch between the different approaches, and how?
    ...

Conclusions

- **Accomplishments**
  - Bayesian approach to data fusion in 2D (theory)
  - Validation in 1D (bandlimited signal reconstruction)
    - Super-resolution from multiple undersampled observations
    - Uncertainty computation - covariance & inverse covariance matrices
- **To do...**
  - 2D implementation (direct extension of the 1D work)
  - 2D/3D: more complex imaging model, but same approach
  - Full 3D surface recovery:
    - Extension of the 2D curve reconstruction method [MaxEnt04]
    - Forward model (rendering): radial basis functions?
    - Reflectance map inference
  - Validation on real data (Ikonos / VO)