 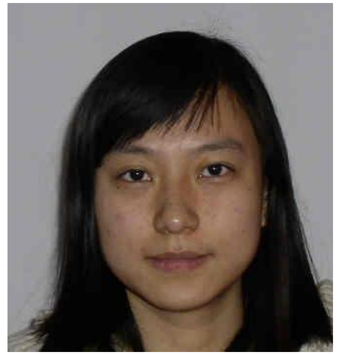  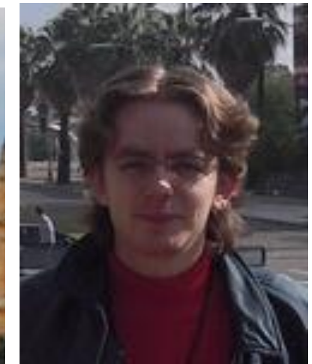

N. Komodakis X. Bo. C. Wang V. Fecamp

Discrete Inference & Learning in Artificial Vision & Medical Imaging

Nikos Paragios

CENTRALE PARIS

inria informatiques mathématiques

Human Vision

- The sensor (iris - diaphragm in a camera, the cornea and the lens are both lens-like objects, the retina is where the image is recorded - CCD sensor)

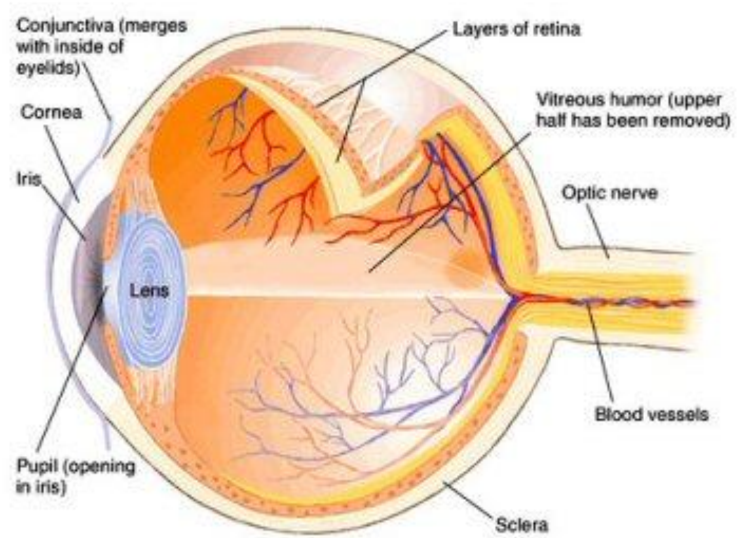

- The processor (information is transferred through the optic nerve to the striate cortex brain part where massive processing is performed towards complete real-time visual scene understanding – almost 50% of the human brain)

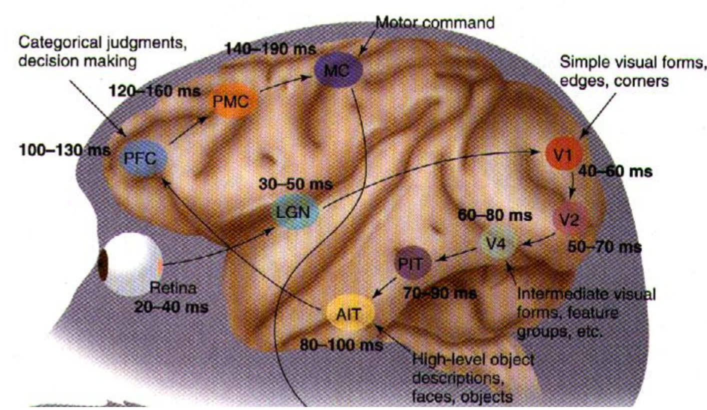

Artificial Vision

- The input (static, video, depth, monochromatic, color high dynamic range, etc sensors)
- The processor (powerful computers exploring input, prior knowledge and models)
- The process (expressing task-specific visual understanding tasks as mathematical inference problems and solve them approximately through computer simulations)

Why artificial vision is so complex?

The input

- Large variety of sensors
- Images/signals of varying quality

The processor

- Even the most powerful individual processor does not match up with a tiny portion of the human brain processing capacities

The mathematical inference

- We are ending up solving problems being ill-defined, ill-posed, non convex, involving non-linear objective functions with numerous local minima

This is what you see

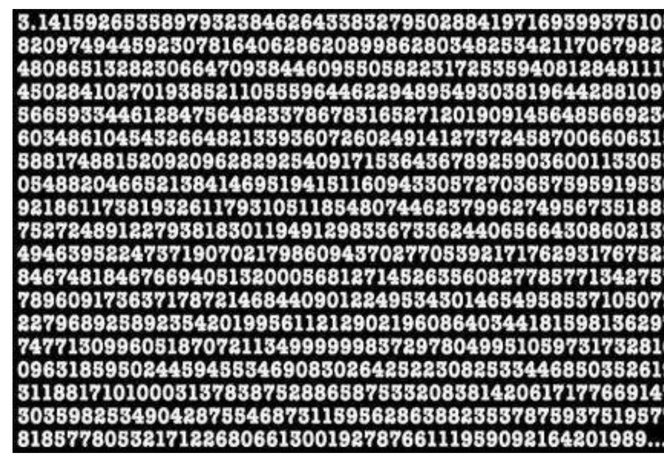

This is what your computer sees

Artificial Vision

Artificial Vision Paradigm

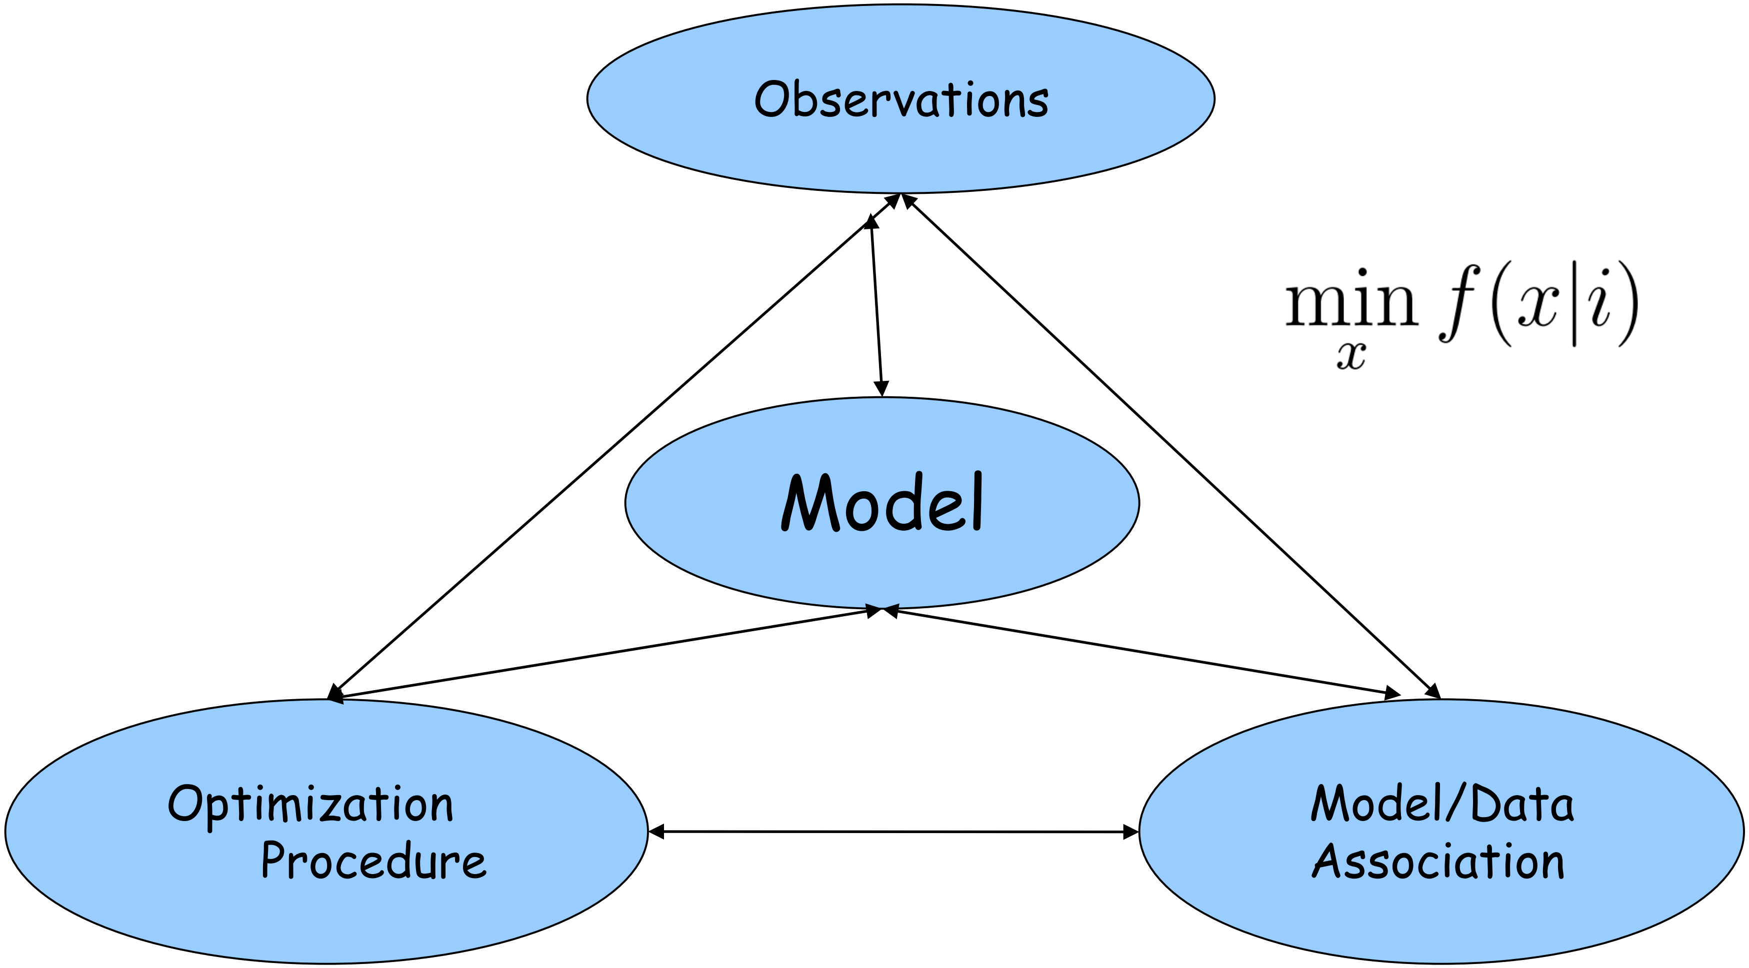

$$\min_x f(x|i)$$

Computer Vision Paradigm

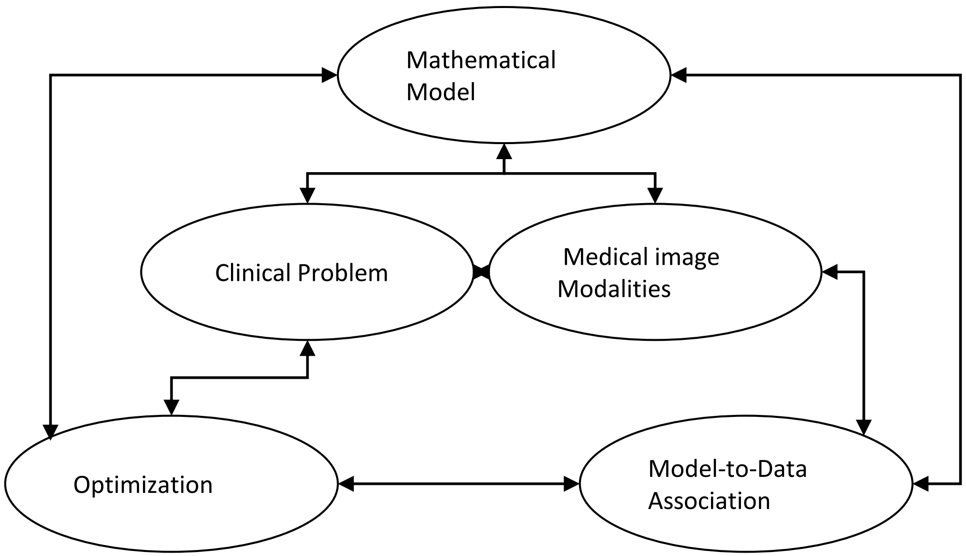

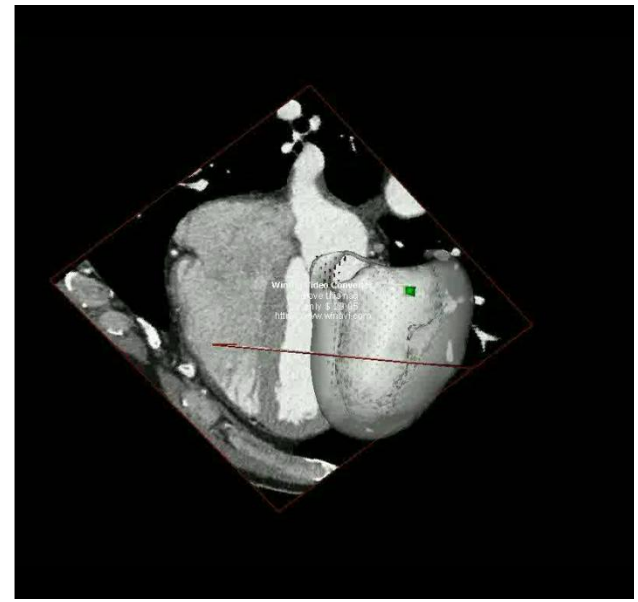

Left Ventricle Segmentation (risk of heart attack)

- Parameters $x = (x_1, \cdots, x_n)$
- Mathematical Model $\partial\mathcal{R} = \pi(x_1, \cdots, x_n)$
- Model-to-data-association $\min_{(x_1,\cdots,x_n)} \int_{\partial\mathcal{R}} f(\partial\mathcal{R}(p)|I)dp$
- Optimization

$$x_i^{\tau+1} = x_i^{\tau+1} + \Delta t \frac{\partial}{\partial x_i} \int_{\partial\mathcal{R}} f(\partial\mathcal{R}(p))I)dp$$

Main Challenges

- Curse of Dimensionality : find a compromise between the expression power of the model and its complexity [finding the right model]
- Curse of Non-linearity: the association of the model parameters and the observations are highly non-linear [finding the right relation between measurements and parameters to be estimated]
- Curse of Non-Convexity: the designed objective function leaves in a high-dimensional non-convex space [finding the right objective function and be able to solve it]
- Curse of Non-Modularity: any solution is hardly portable to another application setting or another problem [do not repeat the process from scratch when moving from one visual task to another]

Discrete Artificial Vision

- Given:
  - Parameters $\mathcal{X}$ from a graph
    $$\mathcal{G} = (\mathcal{X}, \mathcal{E})$$
  - A neighborhood System $\mathcal{E}$
  - Discrete label set $\mathcal{L}$

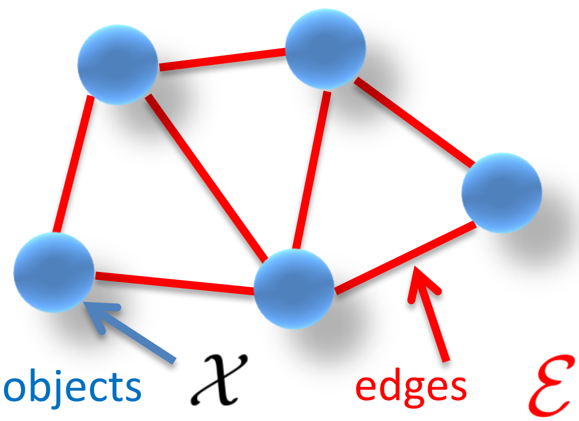

- Assign labels (to objects) that minimize the energy:

$$min_{x_p} \sum_{p \in \mathcal{X}} \underbrace{\Theta_p(x_p)}_{\text{unary potential}} + \underbrace{\Theta_{pq}(x_p, x_q)}_{\text{pairwise potential}}$$

- MRF optimization ubiquitous in vision (and beyond)

MRF hardness

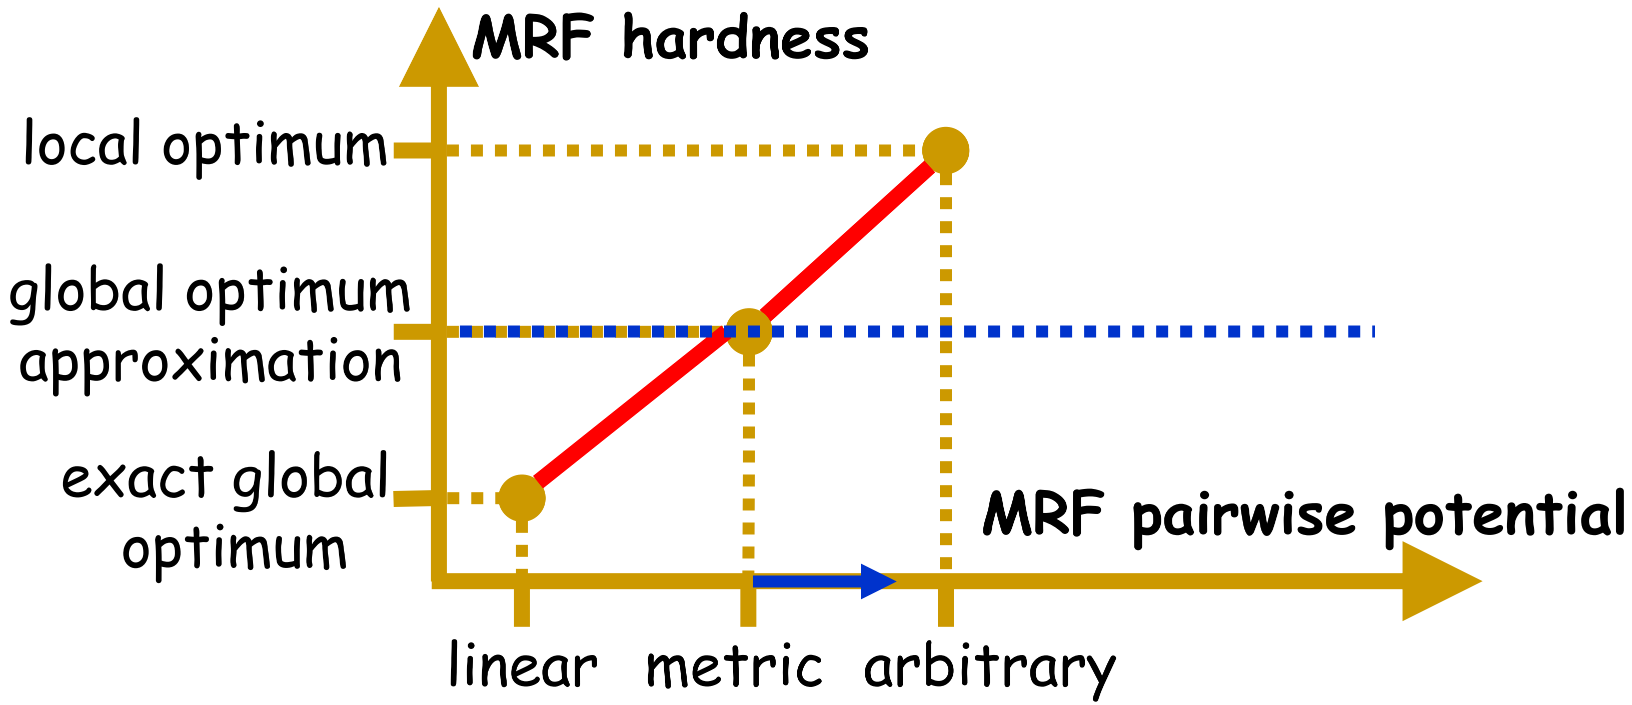

Move left in the horizontal axis,

and remain low in the vertical axis
(i.e., still be able to provide approximately optimal solutions)

But we want to be able to do that efficiently, i.e. Fast – accurately, i.e. Global minimum

MRFs and Optimization

- Deterministic Methods:
  Iterated Conditional Modes/Highest Confidence First
- Non-Deterministic Methods:
  Mean-field and Simulating Annealing, etc
- Graph-cut based techniques such as *a*-expansion:
  Min cut/max flow, etc
- Message-passing techniques:
  Belief Propagation Networks generalized by TRW methods

- The above statement is more or less true for almost all state-of-the-art MRF techniques

Optimization of high-order models

- Hypergraph $\mathcal{G} = (\mathcal{X}, \mathcal{C})$
  - Parameters $\mathcal{X}$
  - Hyperedges/cliques $\mathcal{C}$

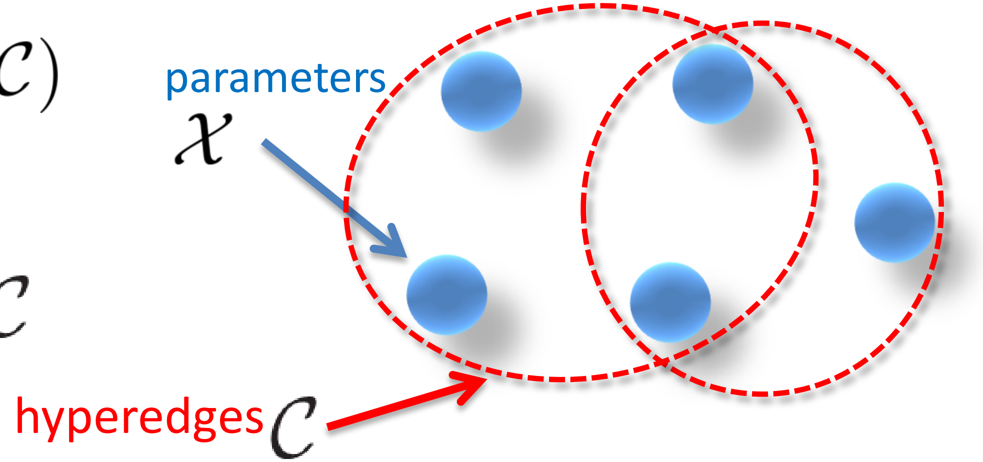

- High-order energy minimization problem

$$min_{x_p} \sum_{p \in \mathcal{X}} \Theta_p(x_p) + \Theta_c(x_p, \cdots, x_q)$$

unary potential (one per node)

high-order potential (one per clique)

MRF optimization via dual-decomposition

Decomposition

- Very successful and widely used technique in optimization.

- The underlying idea behind this technique is surprisingly simple (and yet extremely powerful):

  - decompose your difficult optimization problem into easier subproblems (these are called the **slaves**)

  - extract a solution by cleverly combining the solutions from these subproblems (this is done by a so called **master** program)

Dual decomposition

- The role of the master is simply to coordinate the slaves via messages

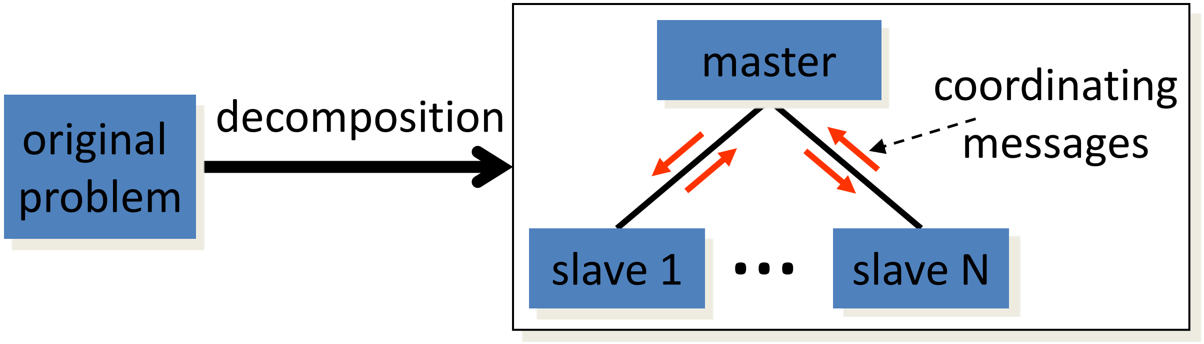

- Depending on whether the primal or a Lagrangian dual problem is decomposed, we talk about **primal** or **dual** decomposition respectively

An illustrating toy example (1/4)

- For instance, consider the following optimization problem (where **x** denotes a vector):

$$\begin{array}{ll} \min_x & \sum_i f^i(x) \\ \text{s.t.} & x \in \mathcal{C} \end{array}$$

- We assume minimizing each $f^i(\cdot)$ separately is easy, but minimizing their sum $\sum_i f^i(\cdot)$ is hard.

- To apply dual decomposition, we will use multiple copies $\mathbf{x}^i$ of the original variables **x**

- Via these auxiliary variables $\{x^i\}$, we will thus transform our problem into:

$$\begin{array}{ll} \min_{\{x^i\},x} & \sum_i f^i(x^i) \\ \text{s.t.} & x^i \in \mathcal{C}, \boxed{x^i = x} \end{array}$$

An illustrating toy example (2/4)

- If coupling constraints $x^i = x$ were absent, problem would decouple. We thus relax them (via Lagrange multipliers $\{\lambda^i\}$) and form the following Lagrangian dual function:

$$g(\{\lambda^i\}) = \min_{\{x^i \in C\},x} \sum_i f^i(x^i) + \sum_i \lambda^i \cdot (x^i - x)$$

Last equality assumes $\{\lambda^i\} \in \Lambda = \{\{\lambda^i\} | \sum_i \lambda^i = 0\}$

because otherwise it holds $g(\{\lambda^i\}) = -\infty$

- The resulting dual problem (i.e., the maximization of the Lagrangian) is now decoupled! Hence, the decomposition principle can be applied to it!

An illustrating toy example (3/4)

- The **i-th slave problem** obviously reduces to:

$$g^i(\lambda^i) = \min_{x^i \in \mathcal{C}} f^i(x^i) + \lambda^i \cdot x^i$$

Easily solved by assumption. Responsible for updating only x^i,
set equal to $\bar{x}^i(\lambda^i) \equiv$ minimizer of i-th slave problem for given λ^i

- The **master problem** thus reduces to:

$$\max_{\{\lambda^i\} \in \Lambda} g(\{\lambda^i\}) = \sum_i g^i(\lambda^i)$$

This is the Lagrangian dual problem, responsible to update $\{\lambda^i\}$
Always convex, hence solvable by projected subgradient method:

$$\lambda^i \leftarrow \left[\lambda^i + \alpha_t \nabla g^i(\lambda^i)\right]_\Lambda$$

$\nabla \equiv$ subgradient w.r.t. λ^i

$[\cdot]_\Lambda \equiv$ projection on feasible set Λ

In this case, it is easy to check that: $\nabla g^i(\lambda^i) = \bar{x}^i(\lambda^i)$

An illustrating toy example (4/4)

- **The master-slaves communication then proceeds as follows:**

1. Master sends current $\{\lambda^i\}$ to the slaves

2. Slaves respond to the master by solving their easy problems and sending back to him the resulting minimizers $\bar{x}^i(\lambda^i)$

3. Master updates each λ^i by setting $\lambda^i \leftarrow \left[\lambda^i + \alpha_t \bar{x}^i(\lambda^i)\right]_\Lambda$

(Steps 1, 2, 3 are repeated until convergence)

Optimizing MRFs via dual decomposition

We can apply a similar idea to the problem of MRF optimization, which can be cast as a linear integer program:

$$\min_{\mathbf{x}} \left[\sum_p \sum_{a \in \mathcal{L}} \theta_p(a) x_p(a) + \sum_{pq} \sum_{a,b \in \mathcal{L}} \theta_{pq}(a,b) x_{pq}(a,b) \right]$$

$$\min_{\mathbf{x}} \; E(\boldsymbol{\theta}; \mathbf{x}) = \boldsymbol{\theta} \cdot \mathbf{x} = \sum_p \boldsymbol{\theta}_p \cdot \mathbf{x}_p + \sum_{pq} \boldsymbol{\theta}_{pq} \cdot \mathbf{x}_{pq}$$

s.t. $\mathbf{x} \in \mathcal{X}$

s.t. $\sum_{a \in \mathcal{L}} x_p(a) = 1,$ ← (only one label assigned per vertex)

$$\left.\begin{array}{l} \sum_{a \in \mathcal{L}} x_{pq}(a,b) = x_q(b), \\ \sum_{b \in \mathcal{L}} x_{pq}(a,b) = x_p(a), \end{array}\right\}$$

← enforce consistency between variables $x_p(a)$, $x_q(b)$ and variable $x_{pq}(a,b)$

$x_p(a) \in \{0,1\}, \quad x_{pq}(a,b) \in \{0,1\}$

$\theta = \{\{\theta_p\}, \{\theta_{pq}\}\}$ is the vector of MRF-parameters consisting of all unary $\theta_p = \{\theta_p(\cdot)\}$ and pairwise $\theta_{pq} = \{\theta_{pq}(\cdot,\cdot)\}$ vectorized potentials.

$\mathbf{x} = \{\{\mathbf{x}_p\}, \{\mathbf{x}_{pq}\}\}$ is the vector of binary MRF-variables consisting of all unary subvectors $\mathbf{x}_p = \{x_p(\cdot)\}$ and pairwise subvectors $\mathbf{x}_{pq} = \{x_{pq}(\cdot,\cdot)\}$

Binary variables:

- $x_p(a) = 1$ ⇔ label a is assigned to node p
- $x_{pq}(a,b) = 1$ ⇔ labels a, b are assigned to nodes p, q

Constraints enforce consistency between variables $\{\mathbf{x}_p\}$ and $\{\mathbf{x}_{pq}\}$

Algorithmic properties

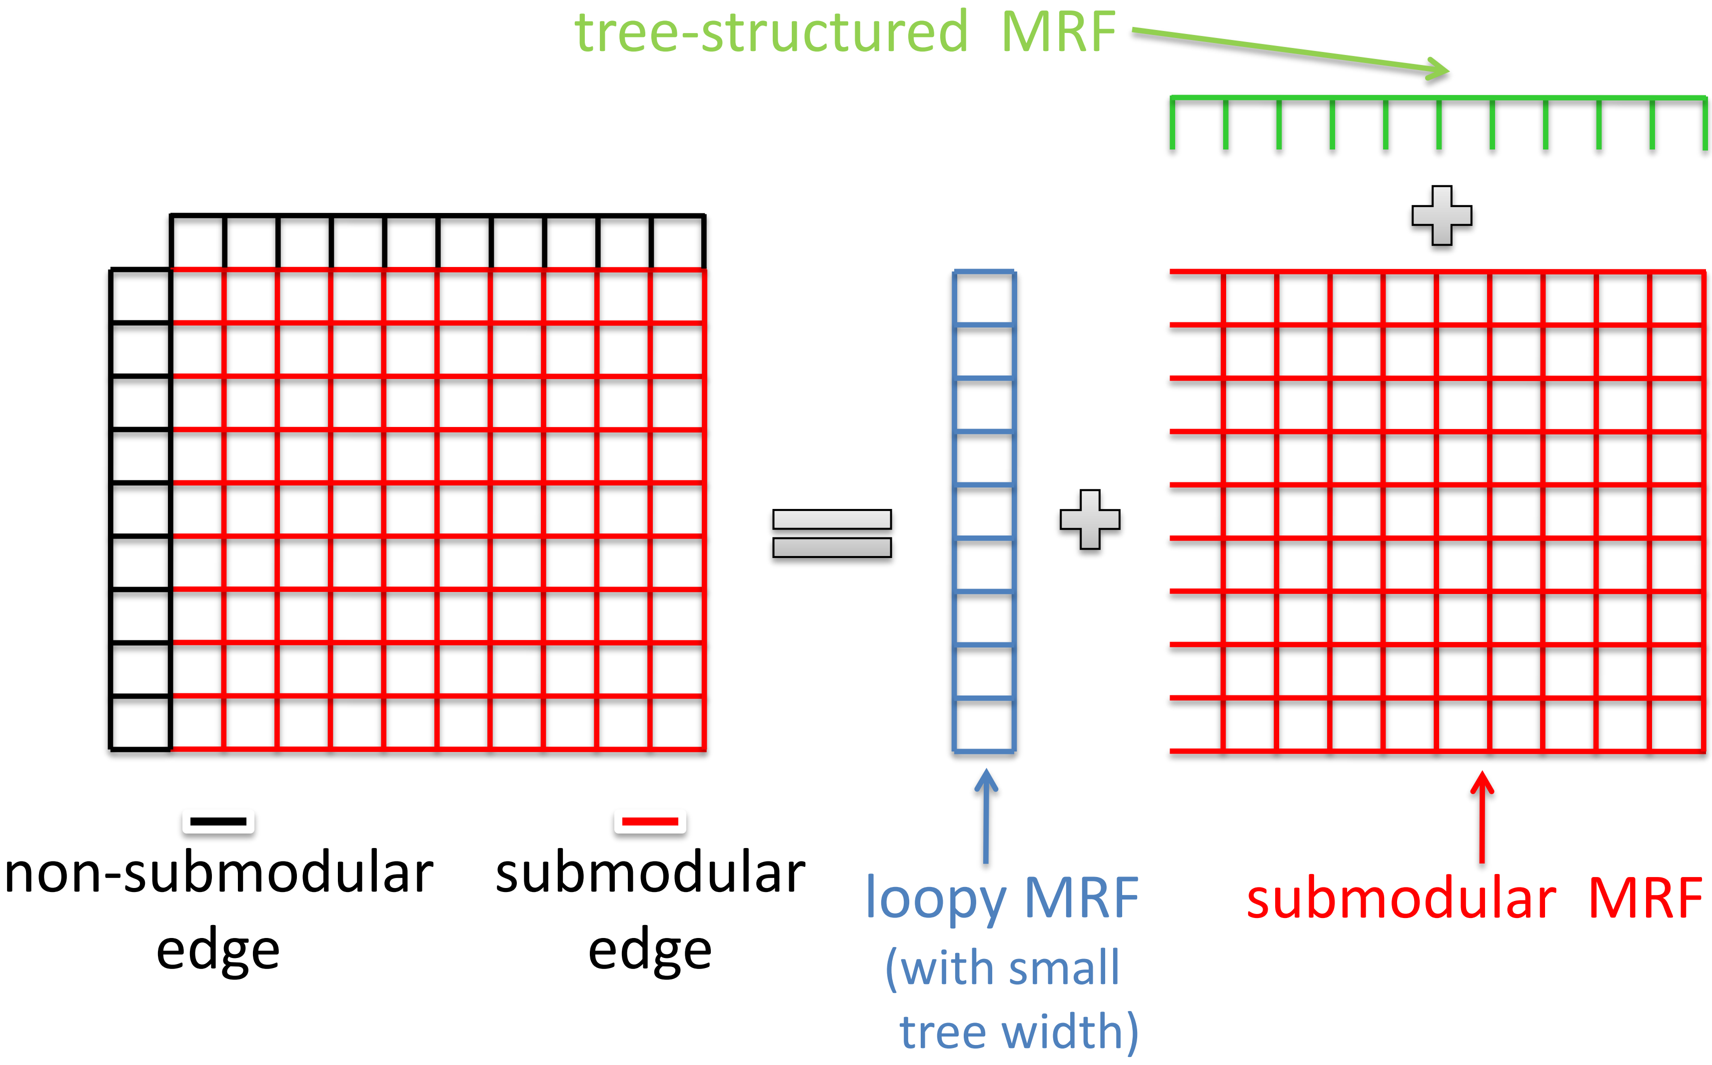

Algorithmic properties

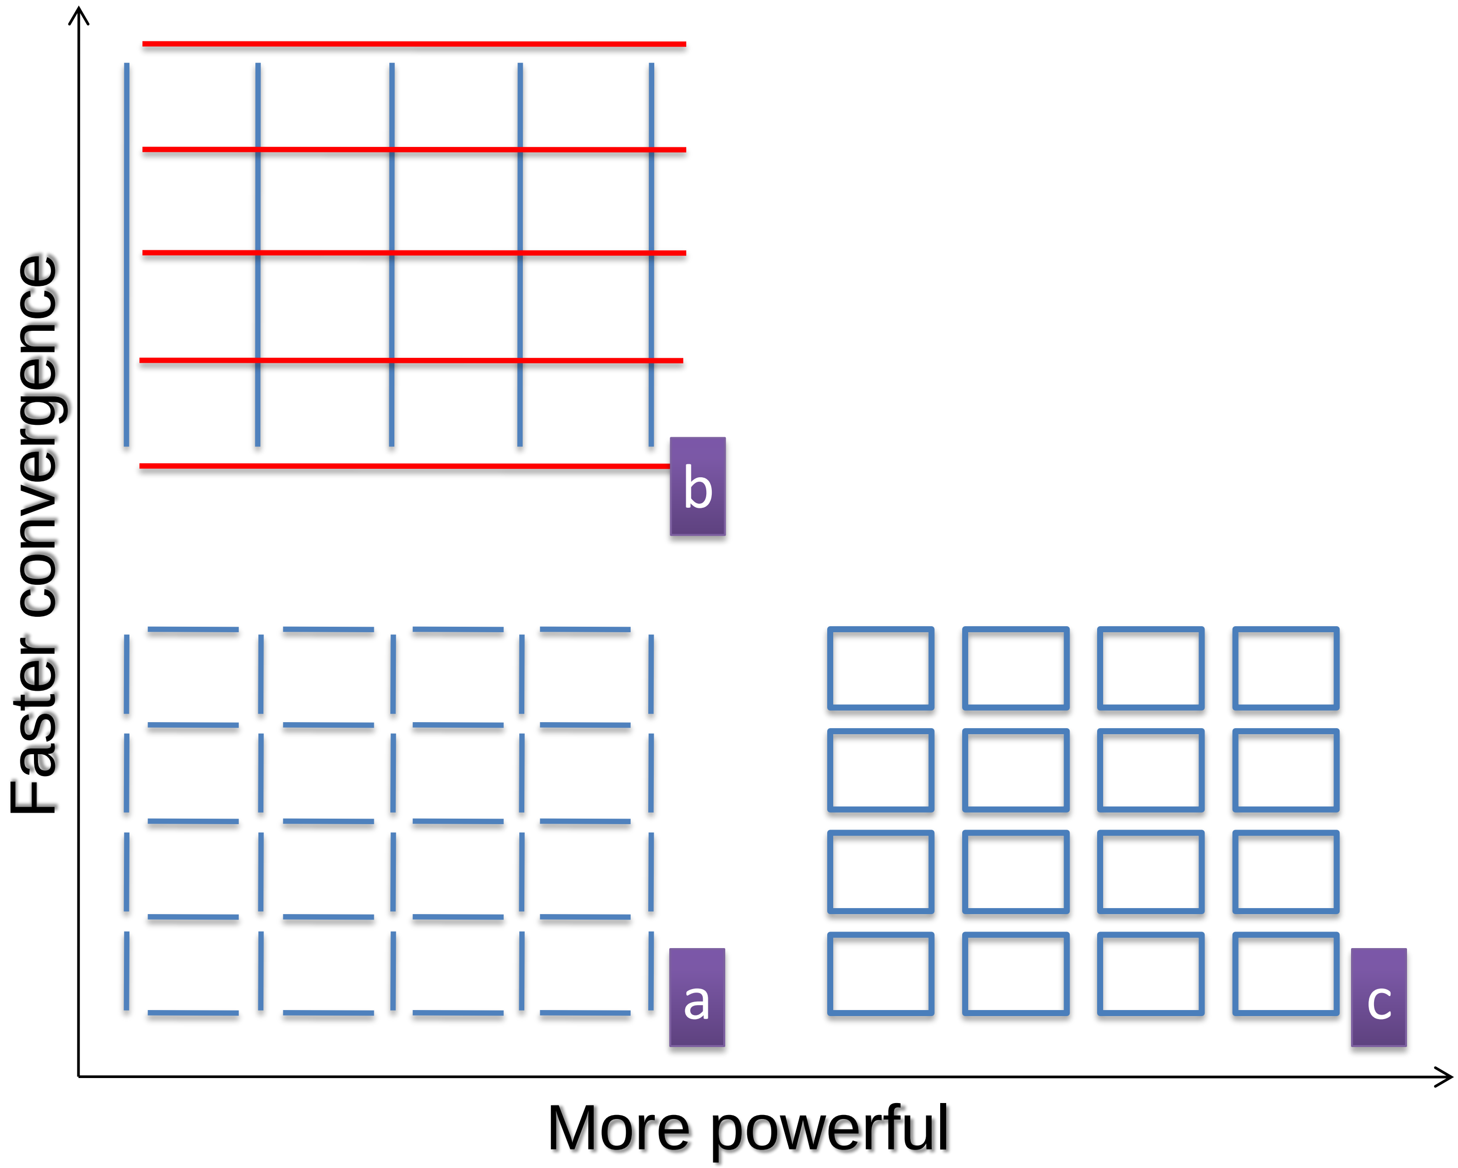

Blind Image Deconvolution

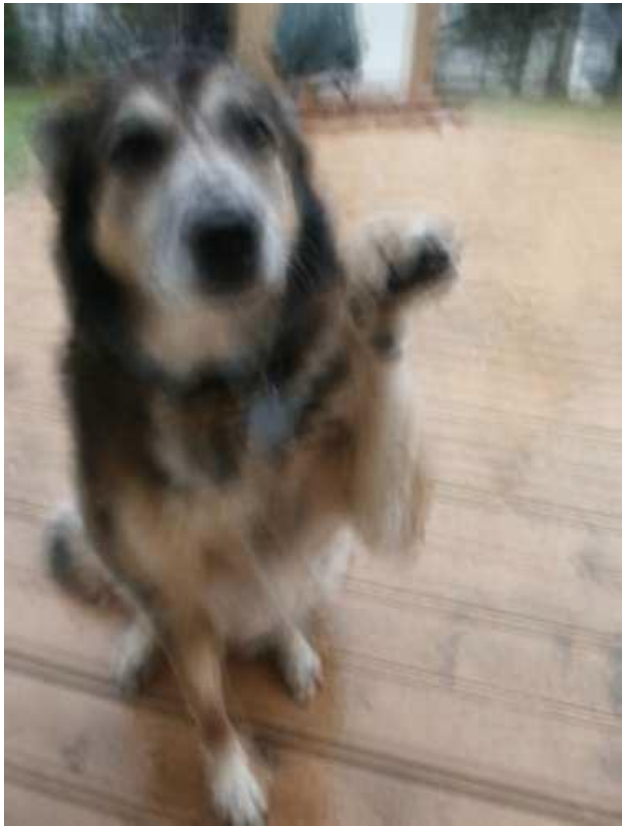

Blurred image generation process

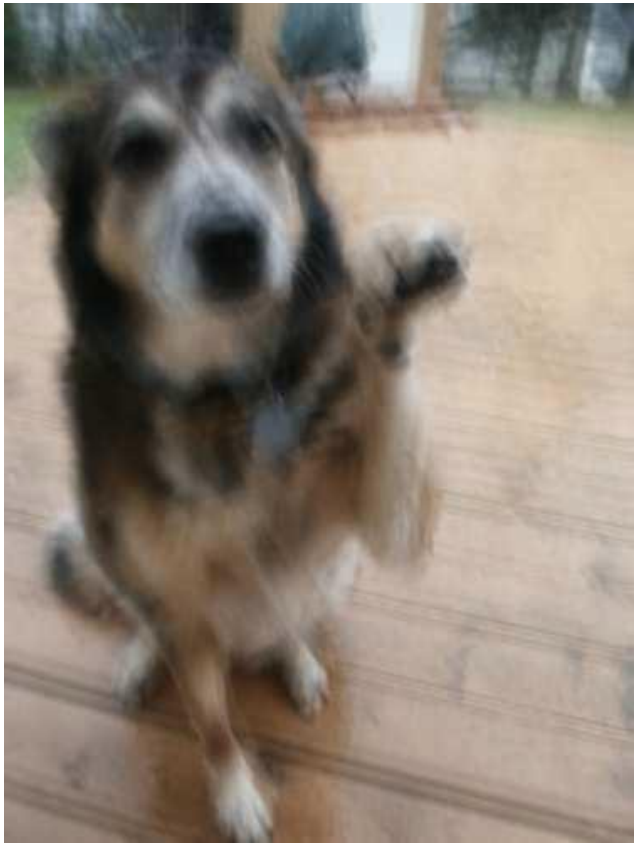

=

Blurred image generation process

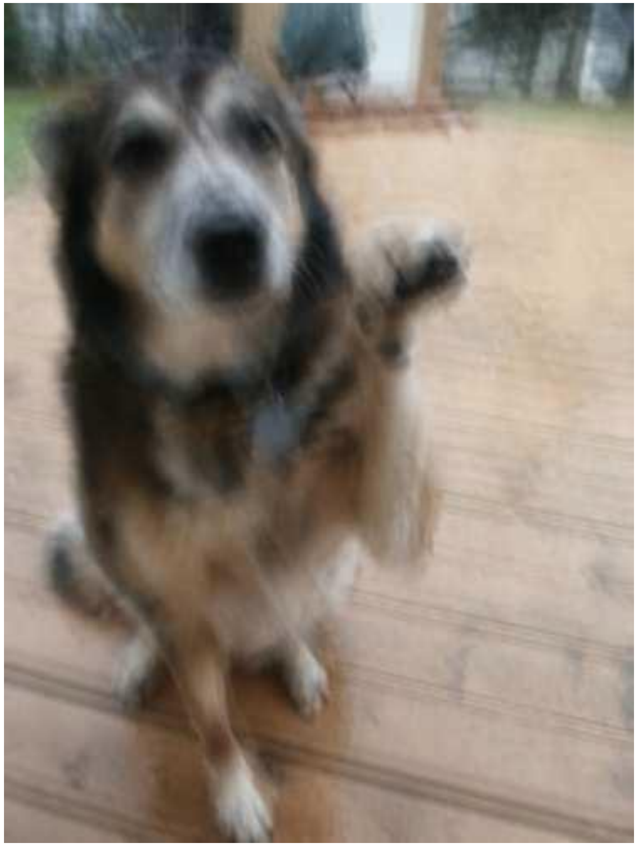

=

⊗

Blurred image generation process

Blurred image generation process

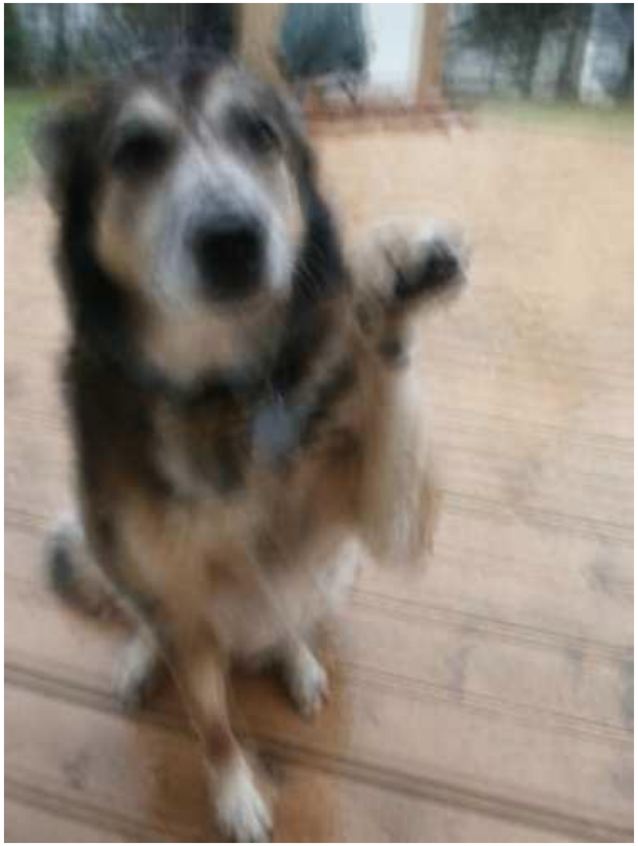

=

⊗

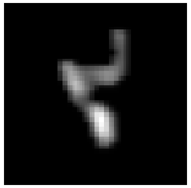

blur kernel = camera motion

Blind Image deconvolution

latent sharp image

blur kernel

$$\mathbf{I} = \mathbf{x} \otimes \mathbf{k} + \mathbf{n}$$

observed blurred image

noise

Goal: given just $\mathbf{I}$ compute both $\mathbf{x}$ and $\mathbf{k}$

High-level idea: how to reduce ill-posedness?

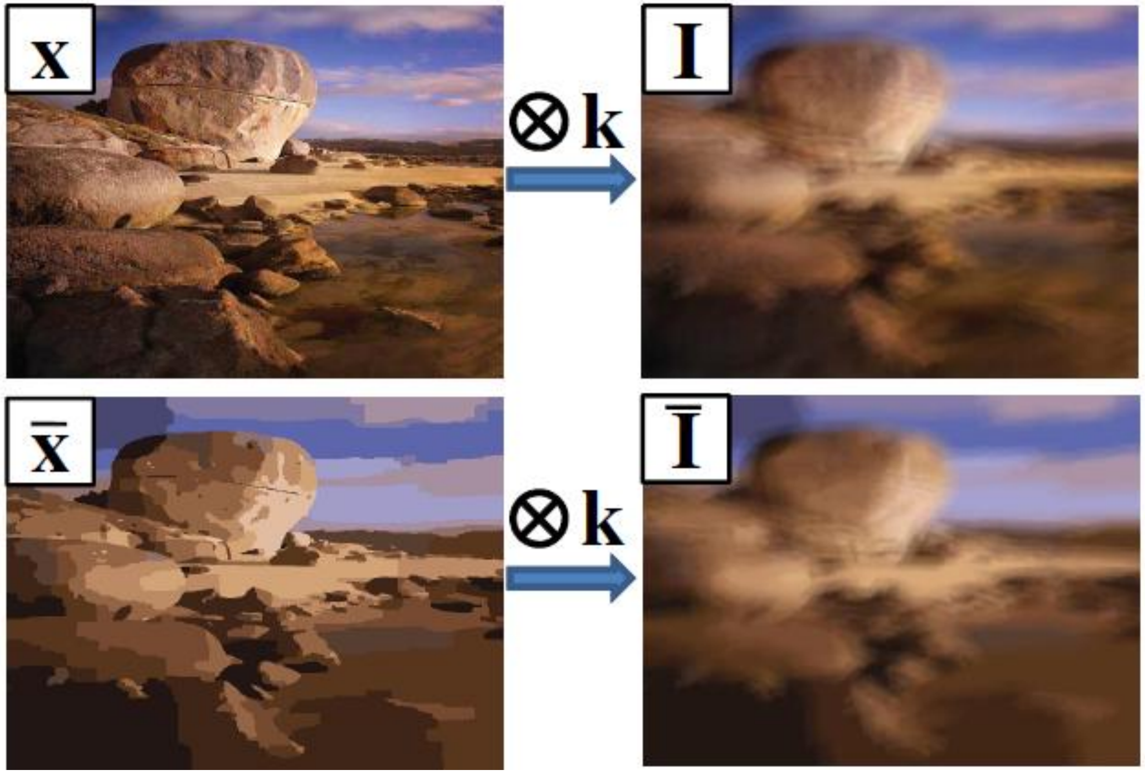

$\bar{\mathbf{x}}$ = quantized version of image $\mathbf{x}$ with just 15 colors (piecewise constant)

Yet both $\mathbf{x}$ and $\bar{\mathbf{x}}$ produce almost same blurry image

IDEA: compute $\bar{\mathbf{x}}$, which has much simpler structure

MRF-based Blind Image deconvolution

MRF-based Blind Image deconvolution

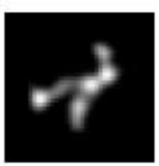

MRF-based Blind Image deconvolution

MRF-based Blind Image deconvolution

MRF-based Blind Image deconvolution

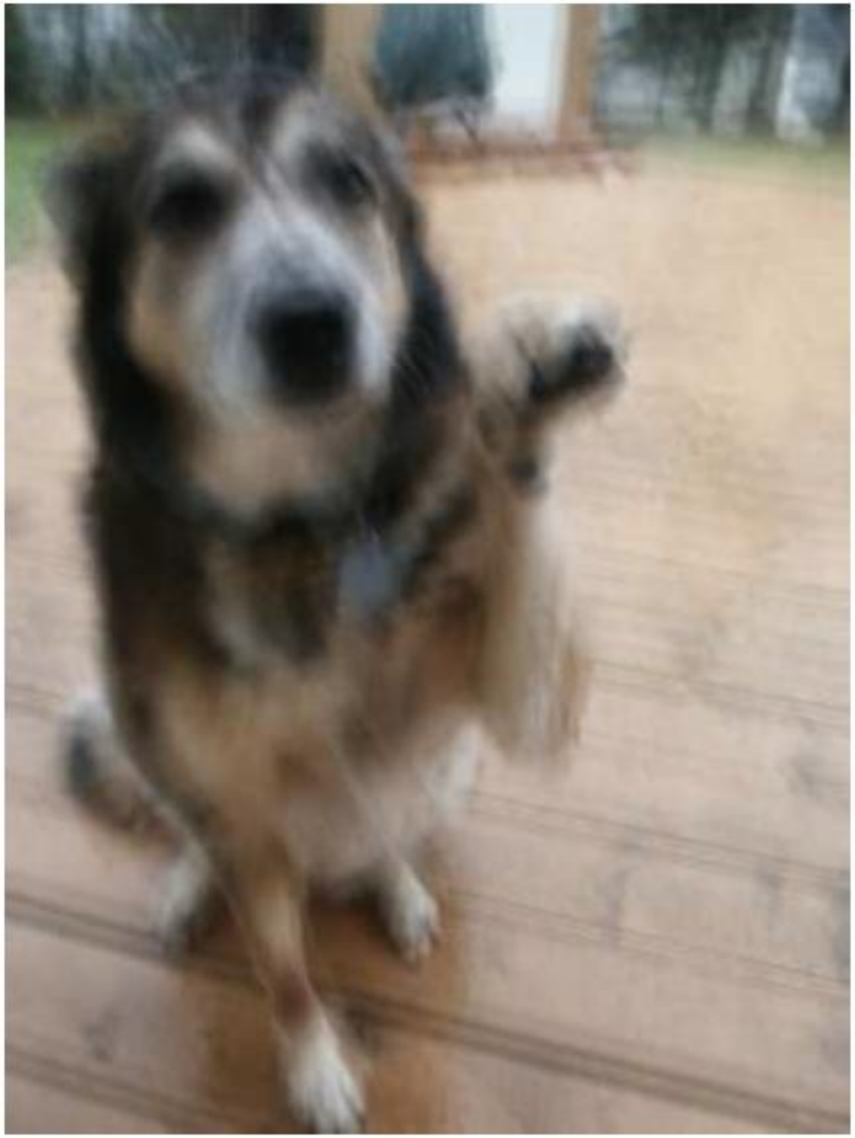

MRF-based Blind Image deconvolution

MRF-based Blind Image deconvolution

MRF-based Blind Image deconvolution

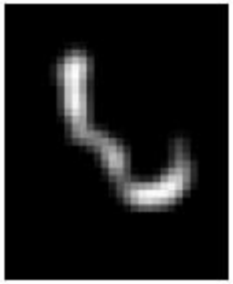

The Image Completion Problem

- Based only on the observed part of an incomplete image, fill its missing part in a visually plausible way

- We want to be able to handle:
  - complex natural images
  - with (possibly) large missing regions
  - in an automatic way (i.e. without user intervention)

- Many applications: photo editing, film post-production, object removal, text removal, image repairing etc.

The Image Completion Problem

- We would also like our method to be able to handle the related problem of texture synthesis
- In texture synthesis, we are given as input a small texture and we want to generate a larger texture of arbitrary size (specified by the user)

Exemplar-based approaches

- **Key idea:** fill missing region by copying exemplars i.e. pixels (or patches) from the observed image part

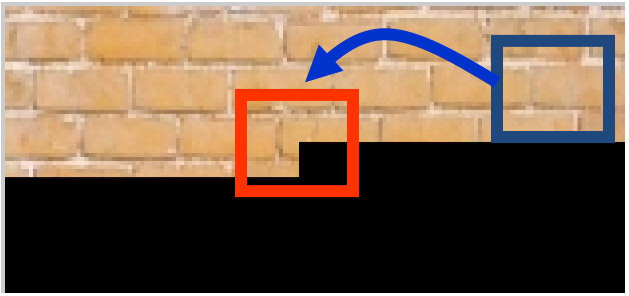

- **Disadvantages:**
  - Successful if missing region consists of only one texture e.g. texture synthesis
  - Greedy approach: image is filled one patch at a time

Image Completion as a Discrete Global Optimization Problem

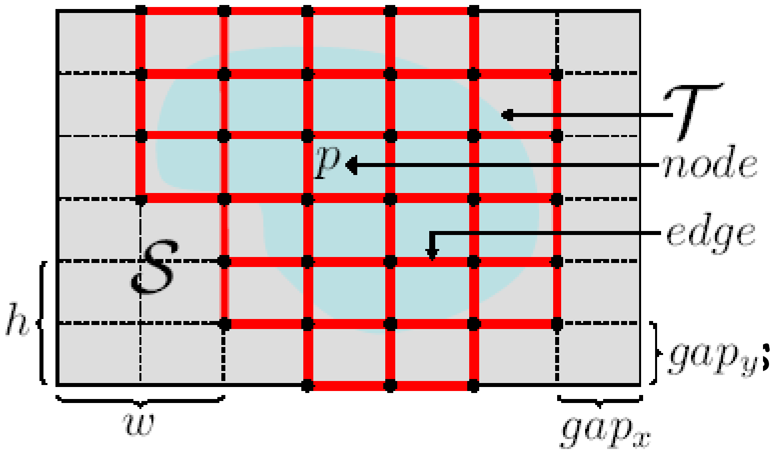

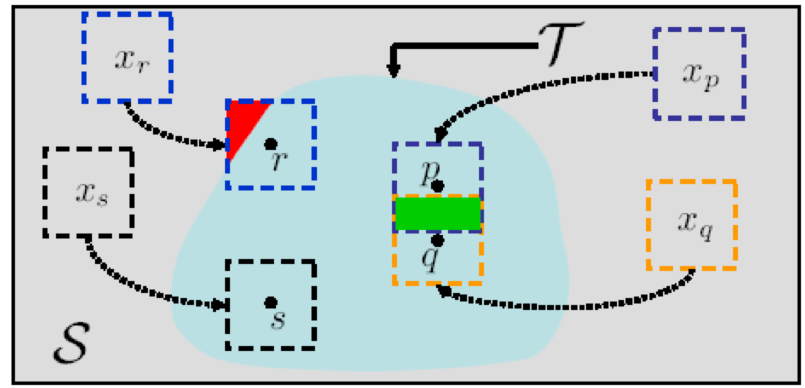

- **Labels L** = all wxh patches from source region S
- **MRF nodes** = all lattice points whose neighborhood intersects target region T
- **potential** $\Theta_p(x_p)$ = how well source patch x_p agrees with source region around p
- **potential** $\Theta_{pq}(x_p, x_q)$ = how well source patches x_p, x_q agree on their overlapping region

 

Pose Invariant Segmentation of the Heart

- Challenges
  - Human variability
  - Complex background
  - Low contrast
  - Noise
- Goal
  - Automatic
  - Robust
  - Pose-invariant !

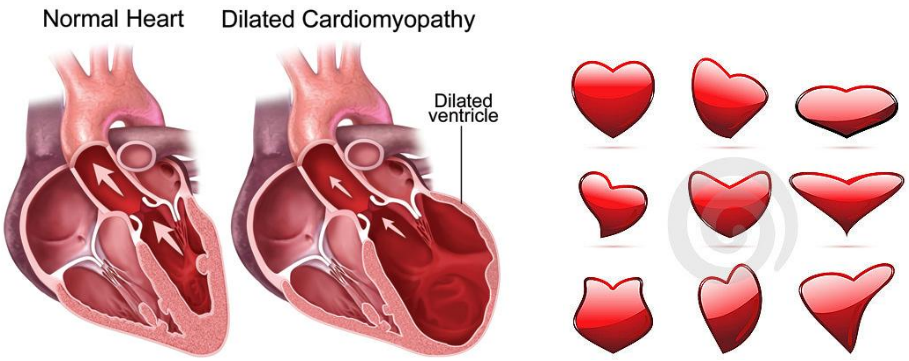

Fig. Human variability

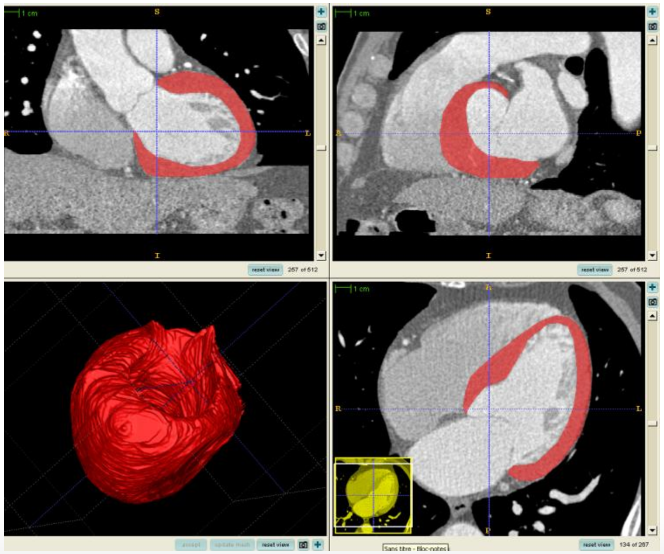

Fig. Manual segmentation on 3D CT images

B. XIANG et al, 3D Cardiac Segmentation with Pose-invariant Higher-order MRFs, ISBI 2012

Shape representation

- Point distribution model

  $X = \{x_1, \cdots, x_n\}$

  $Y \subset X$

- Third-order cliques

  $T = \{(x_i, x_j, x_k)\}$

  $S \subset T$

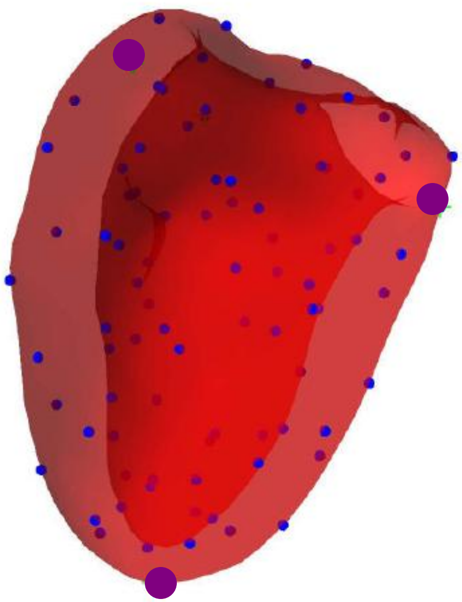

Point distribution model

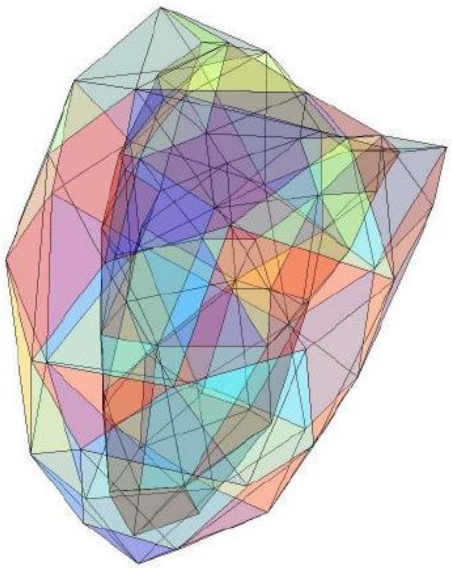

Triangulated mesh

Statistical shape prior

- Local constraints

$$P_{(i,j,k)}(\alpha, \beta)$$

- Global shape

$$P(X) = \frac{1}{Z} \cdot \prod_{c \in C} P_c(\alpha, \beta)$$

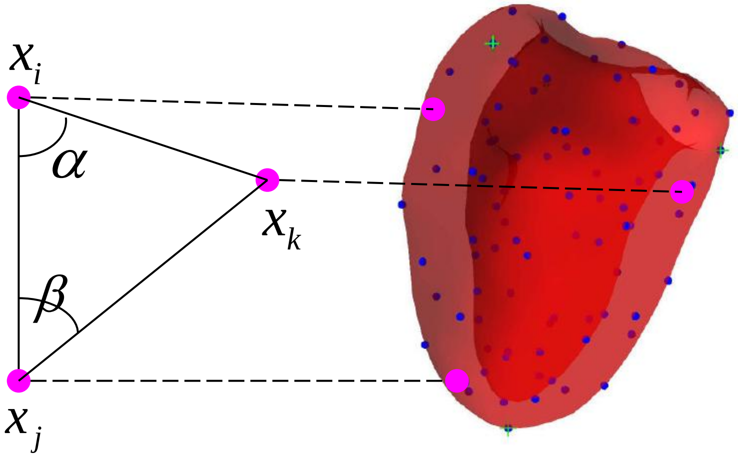

Statistical shape prior

- Local constraints

$$P_{(i,j,k)}(\alpha,\beta)$$

- Global shape

$$P(X)=\frac{1}{Z}\cdot\prod_{c\in C}P_c(\alpha,\beta)$$

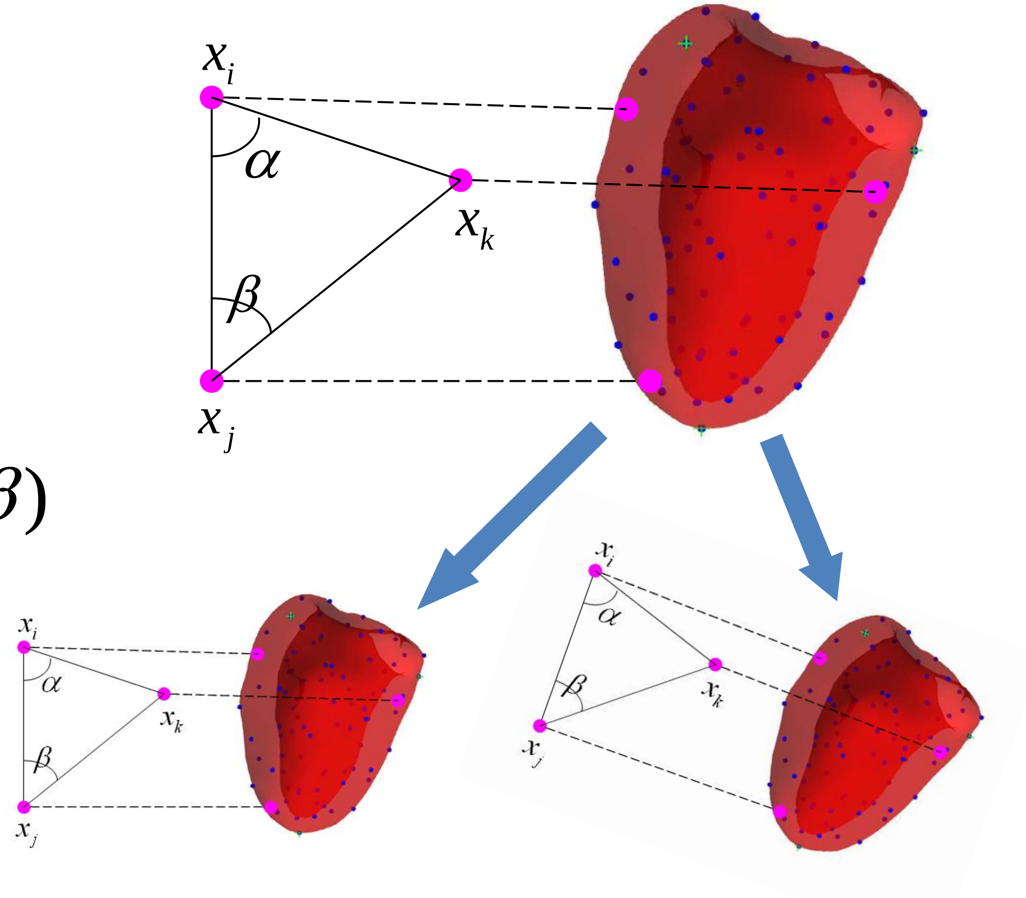

Pose-invariant (i.e. translation, rotation, scale) !

Qualitative Results

Accurate boundaries with low contrast images

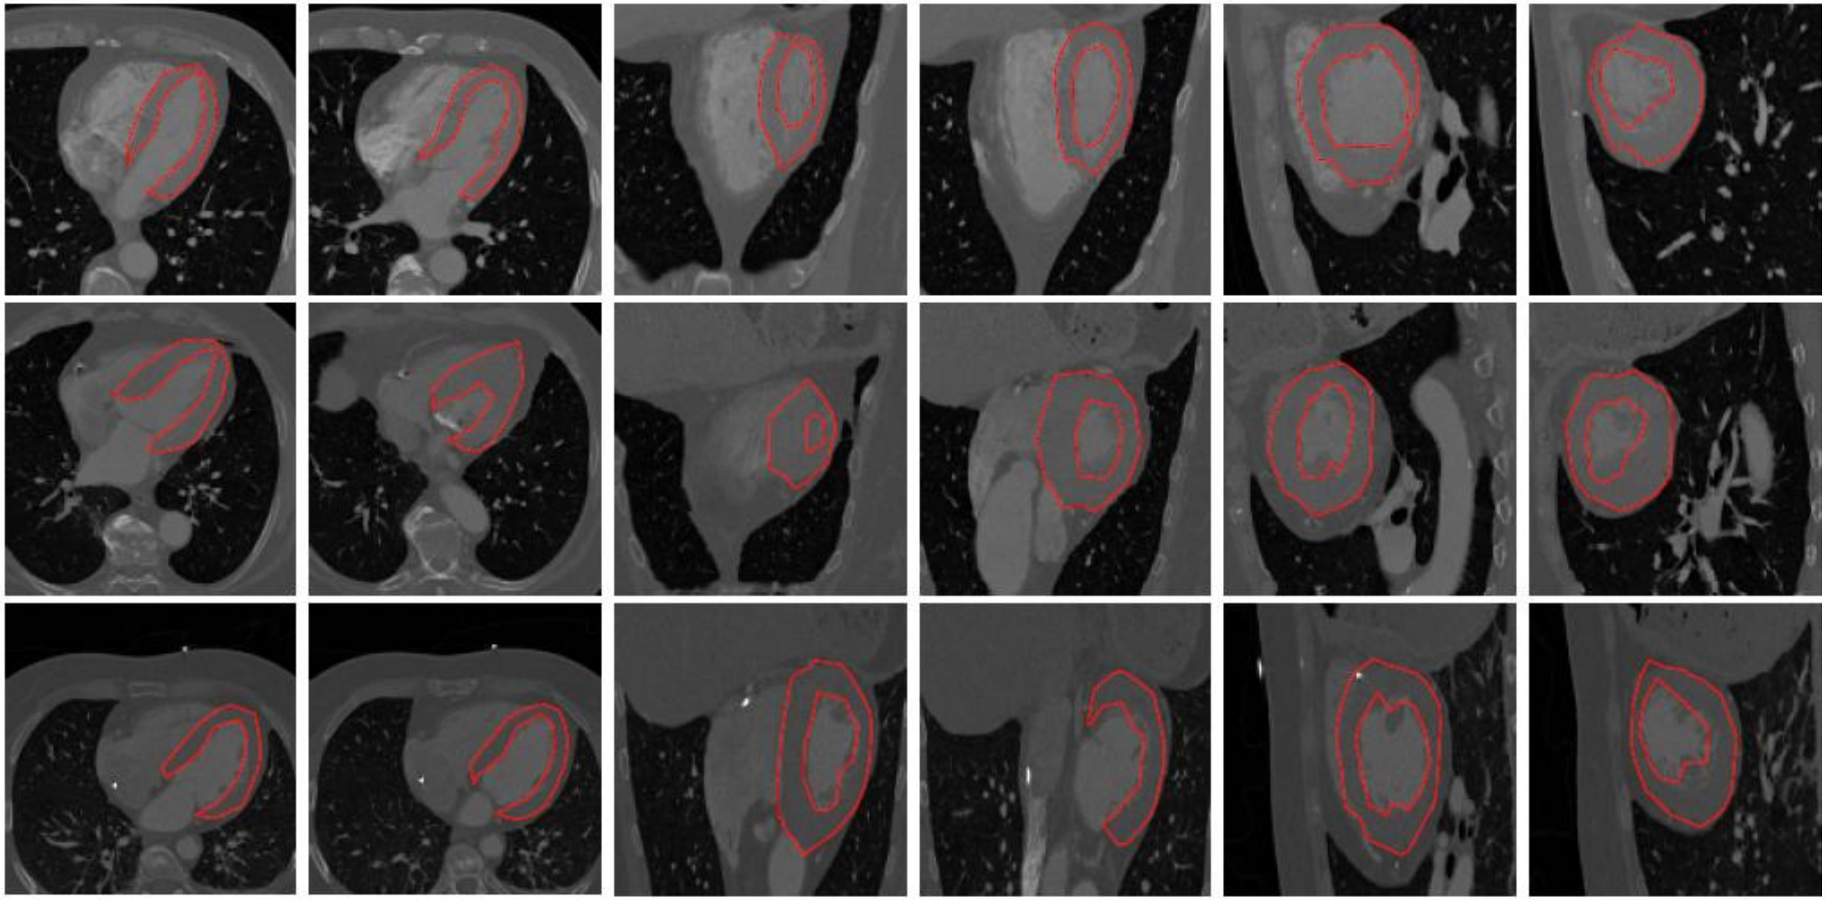

Fig. Segmentation results of 3D CT volumes

Linear Registration Using High Order Graphs

Unary Potentials

- Comparison of a patch from the source to a patch from the target image
- Metric-invariant

$$U_p(l_p) = \rho(B_p, B_{l_p}).$$

High Order Cliques

- 3 aligned points along each axis

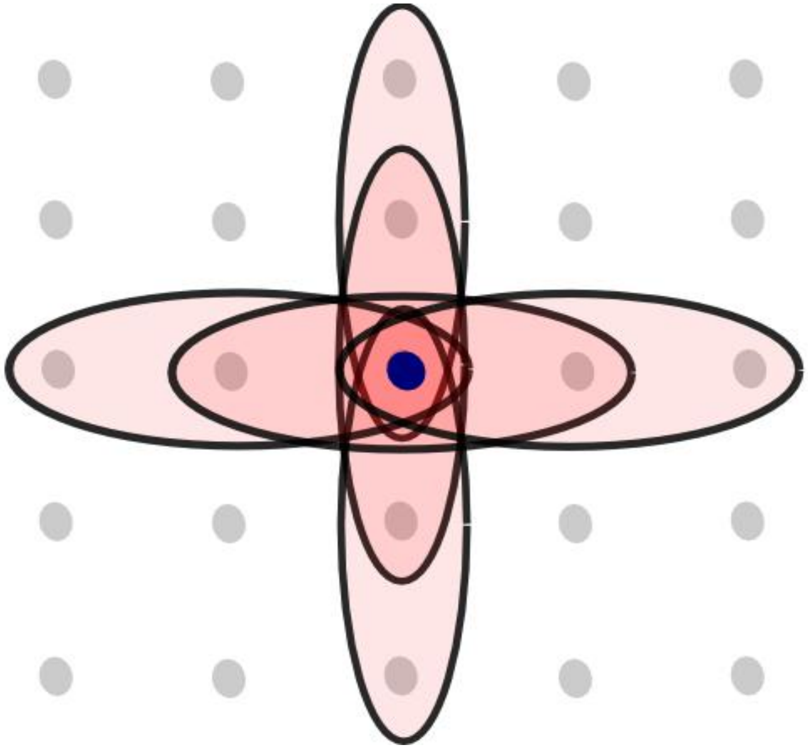

- 1 T-clique

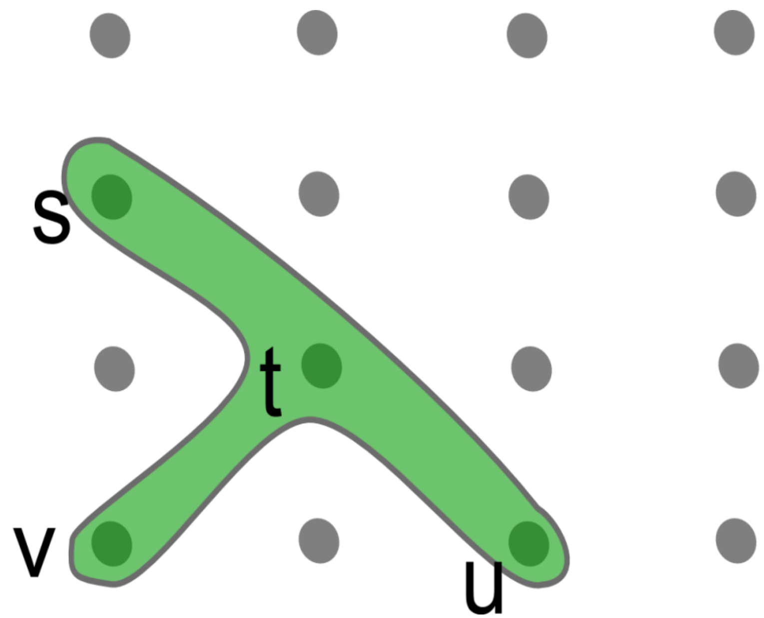

High Order potentials

- Linear transformation :
  - Preserves barycenters
- Condition (P) :

$$\vec{l_p} + \vec{l_r} - 2 * \vec{l_q} = \vec{0}.$$

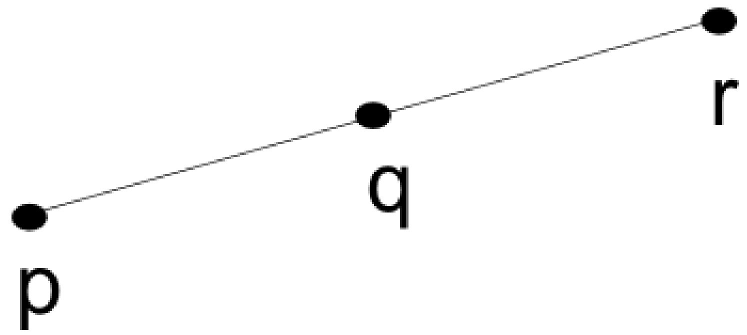

- Easy to compute, only depends on the label

High Order potentials in T-clique

- Check condition (P) for (s,t,u)
- Similarity registration :
  - Images of (s,v,u) form a right isosceles triangle
- Rigid registration :
  - Images of (s,v,u) form a right isosceles triangle, with the same size as (s,v,u) triangle

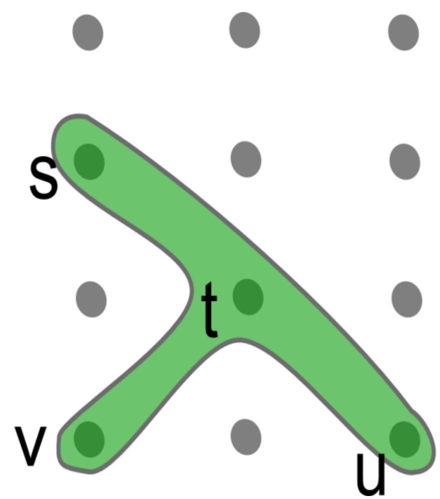

Some results

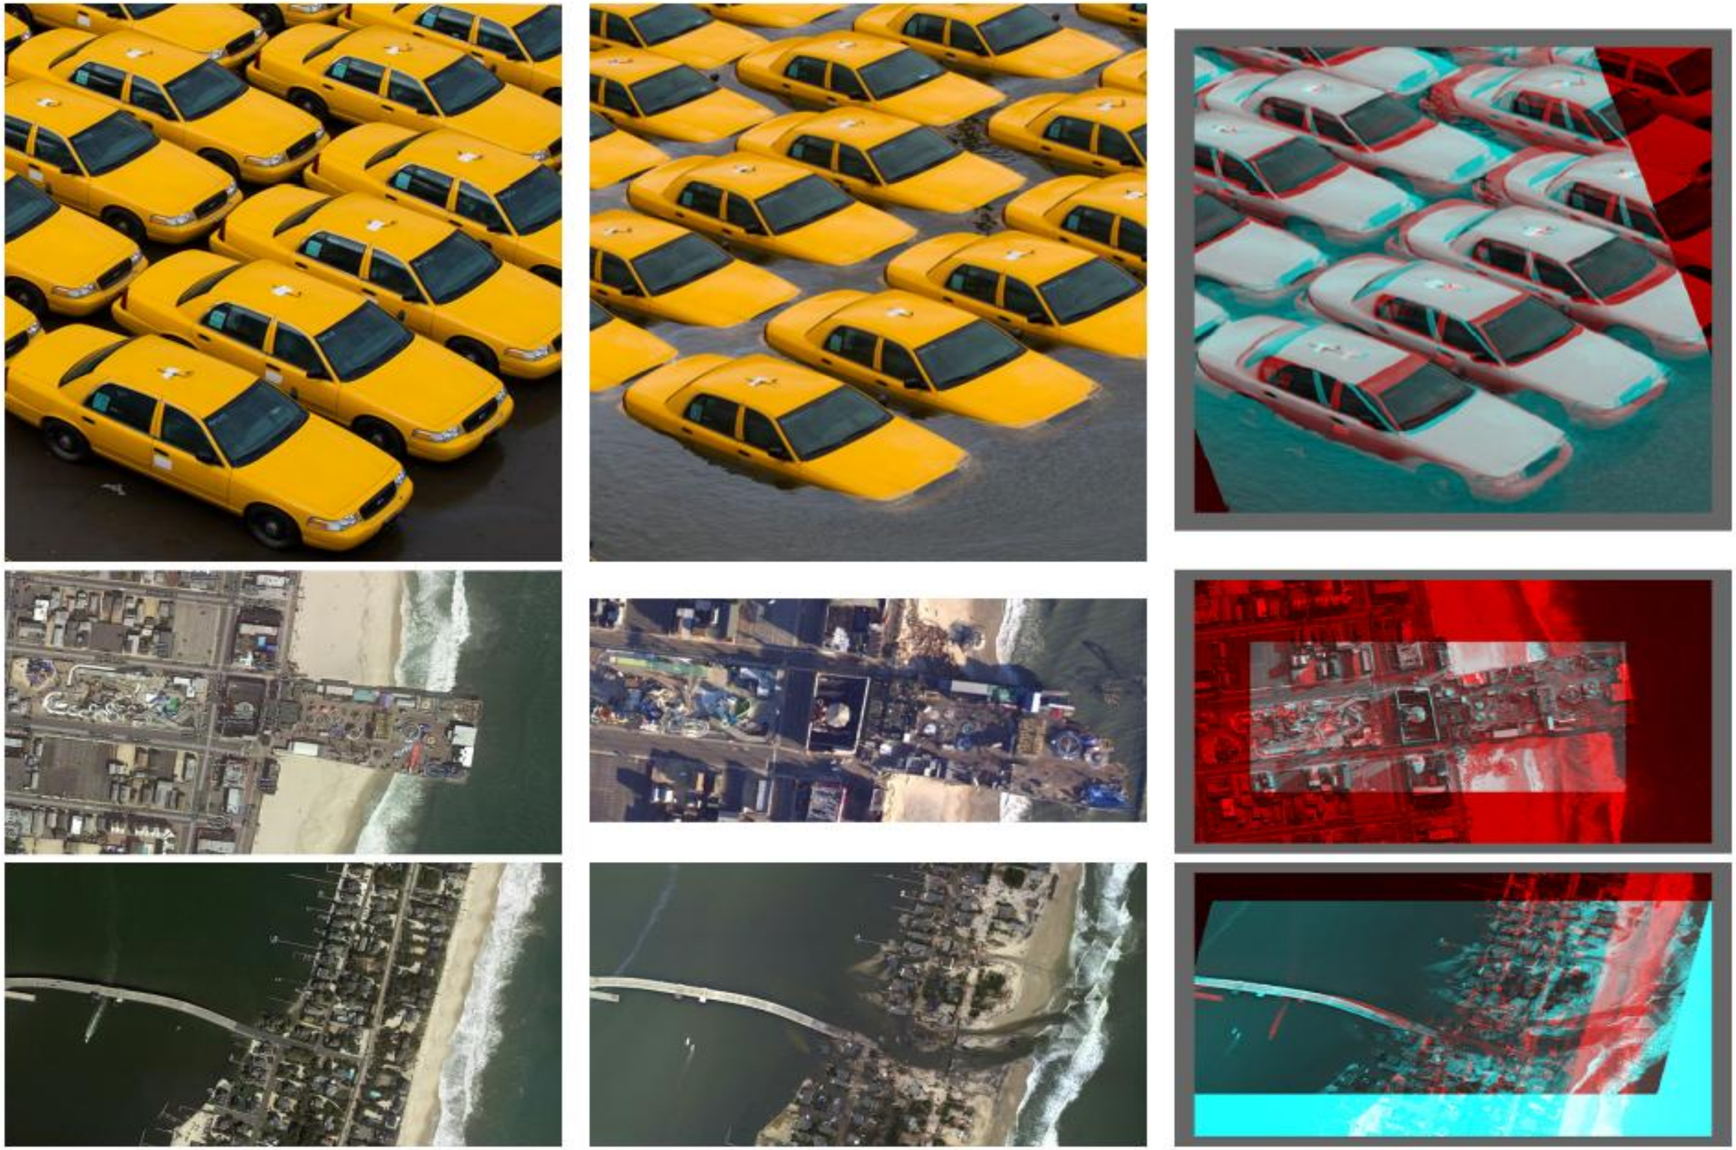

(a) Before Sandy (b) After Sandy (c) After affine registration

HIGHER-ORDER NON-RIGID 3D SURFACE MATCHING

High-order Graph Matching

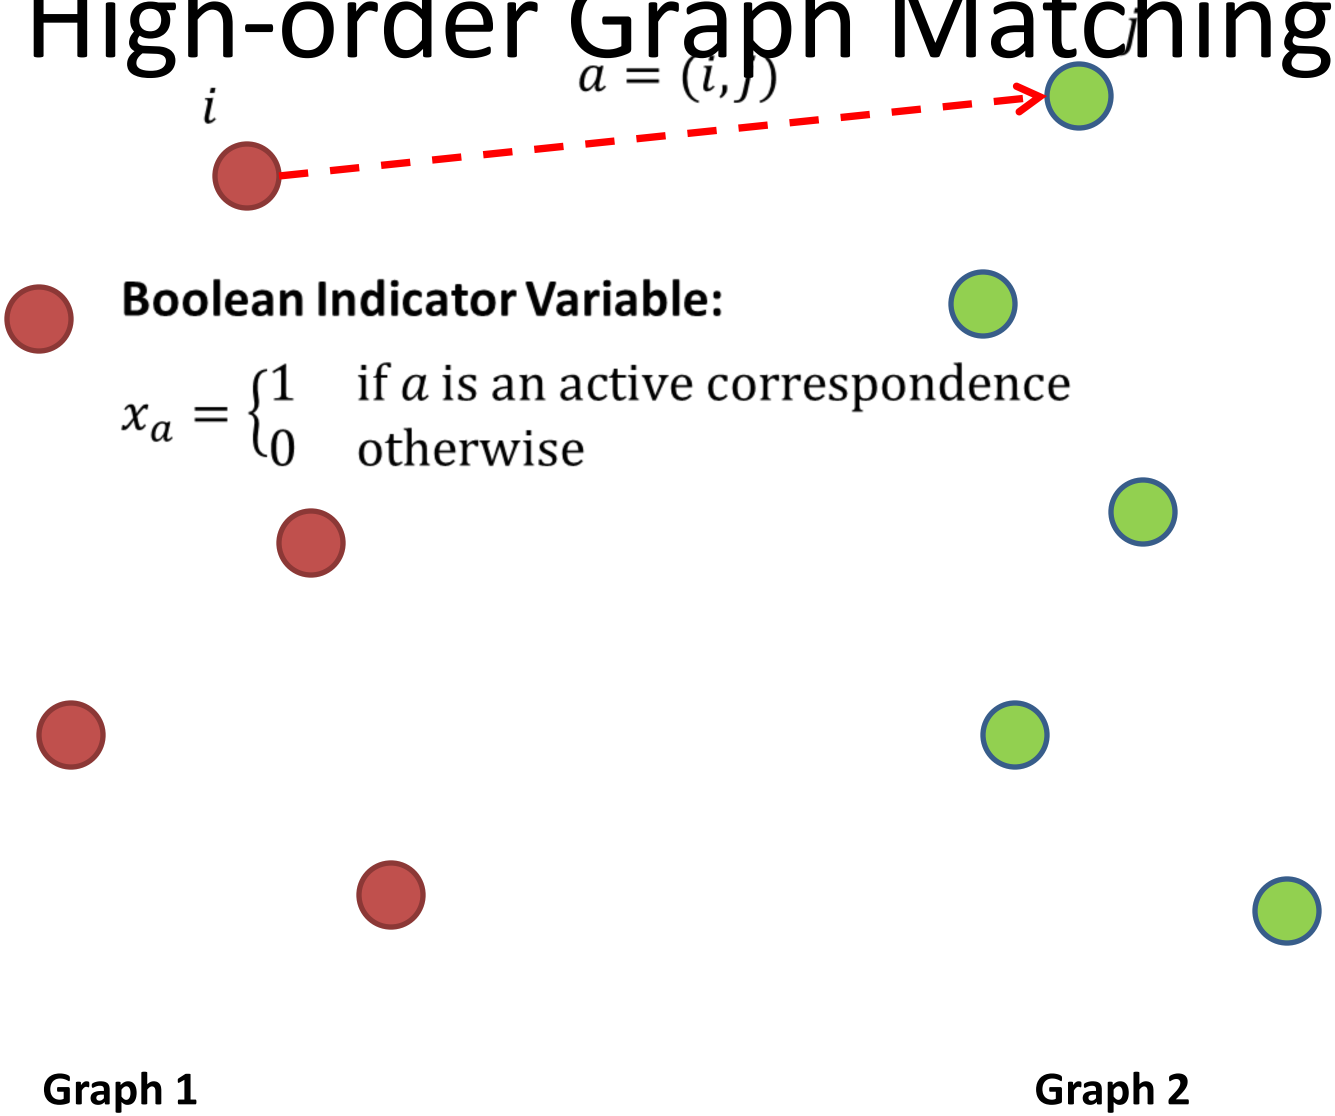

Boolean Indicator Variable:

$$x_a = \begin{cases} 1 & \text{if } a \text{ is an active correspondence} \\ 0 & \text{otherwise} \end{cases}$$

High-order Graph Matching

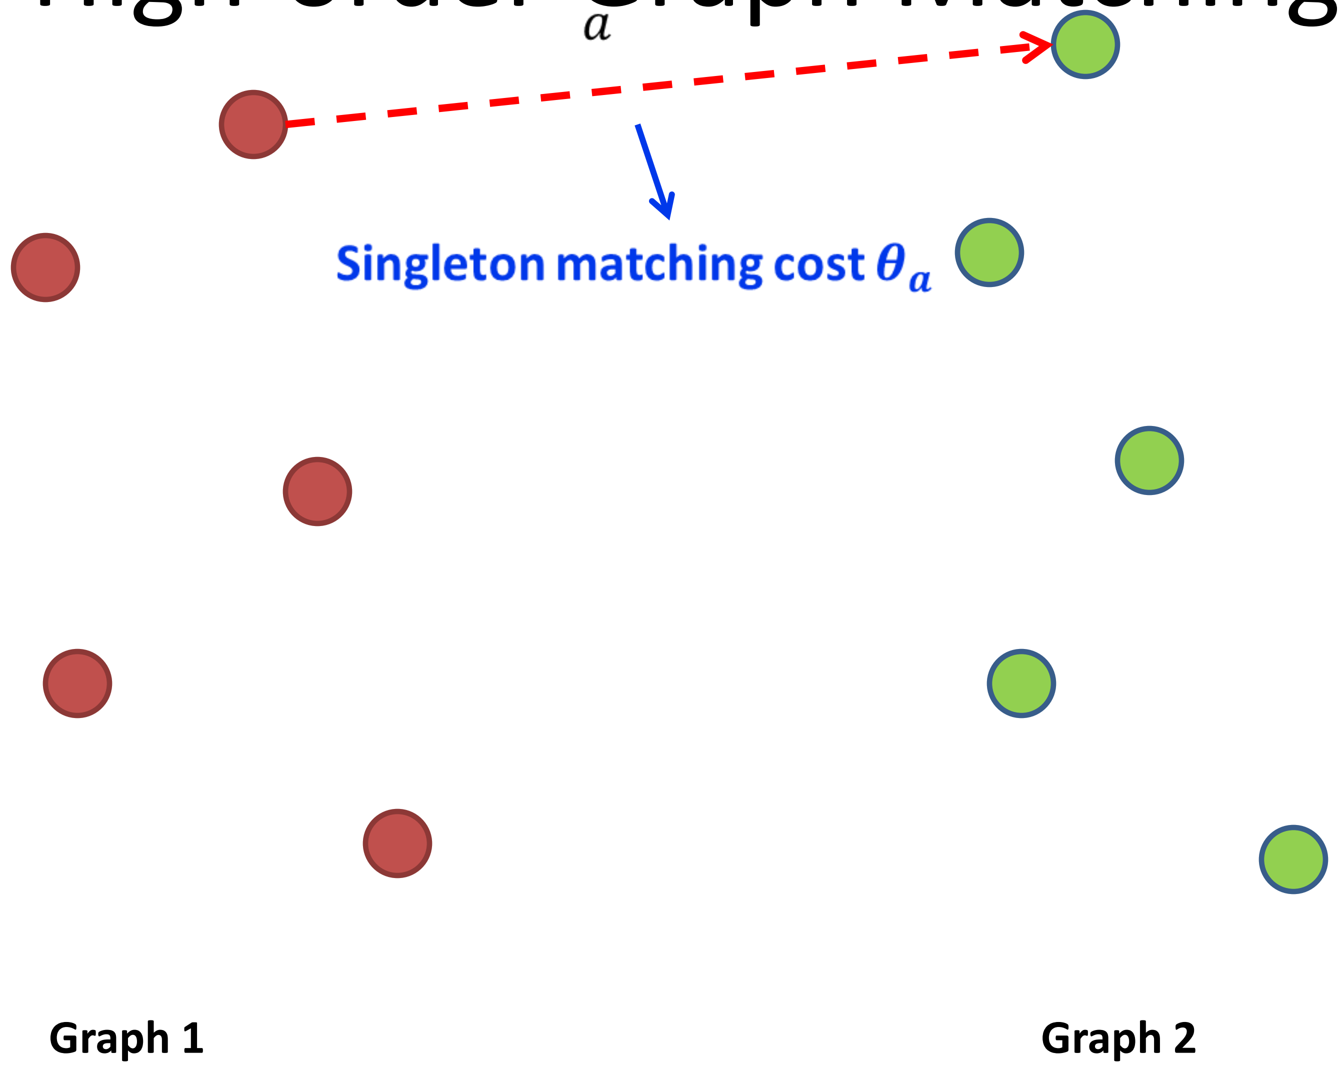

High-order Graph Matching

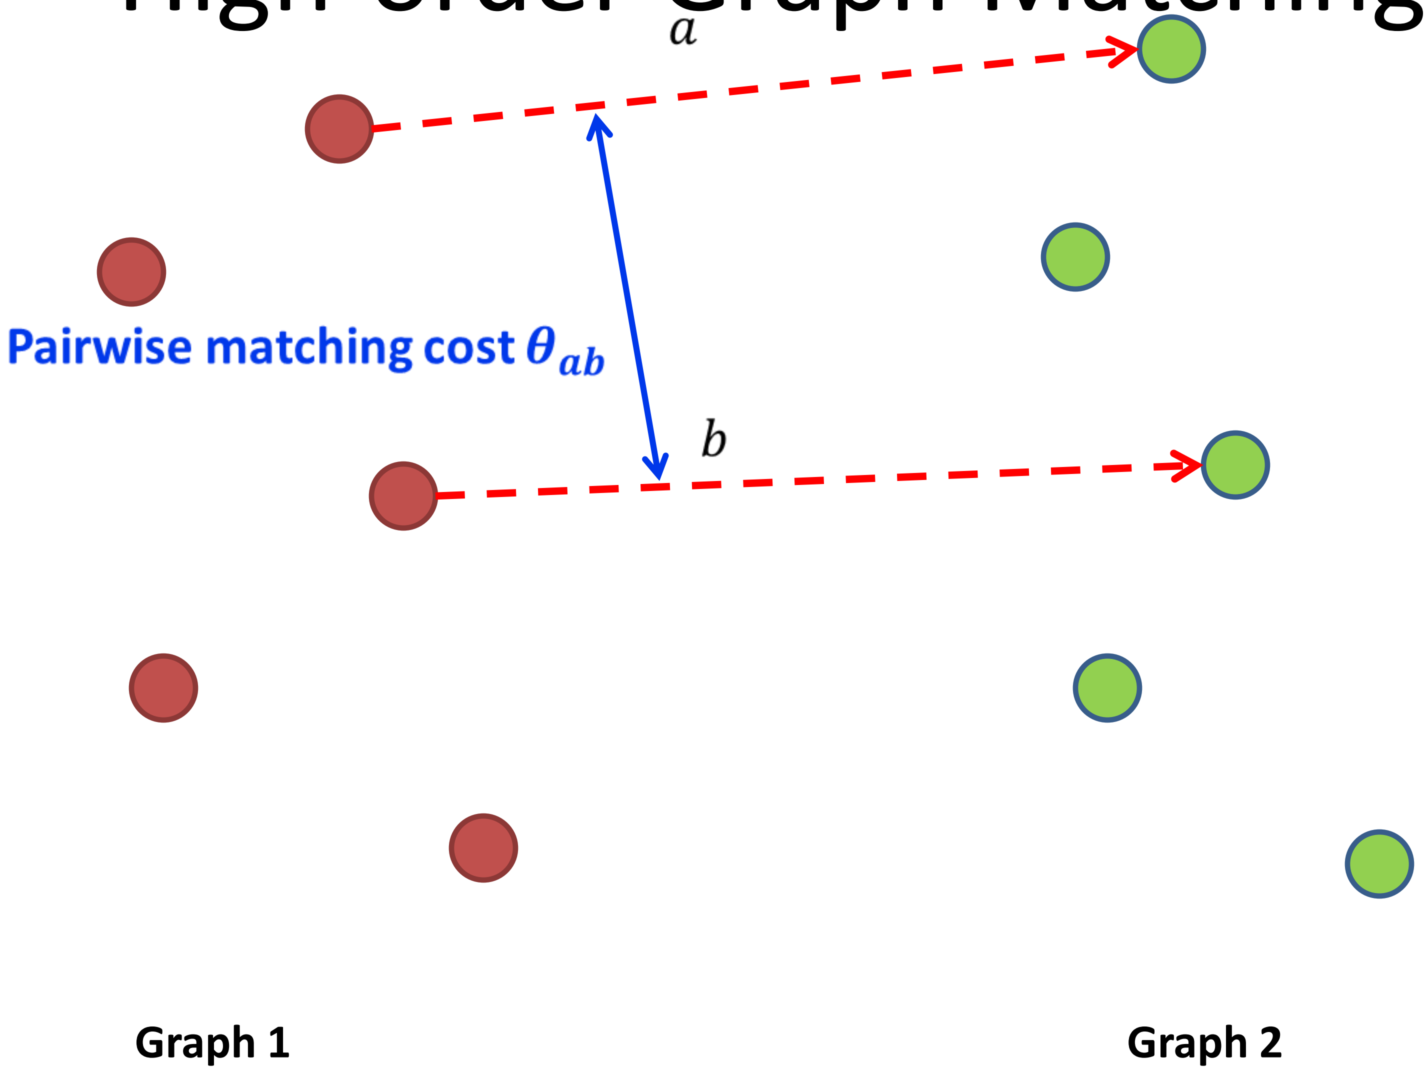

High-order Graph Matching

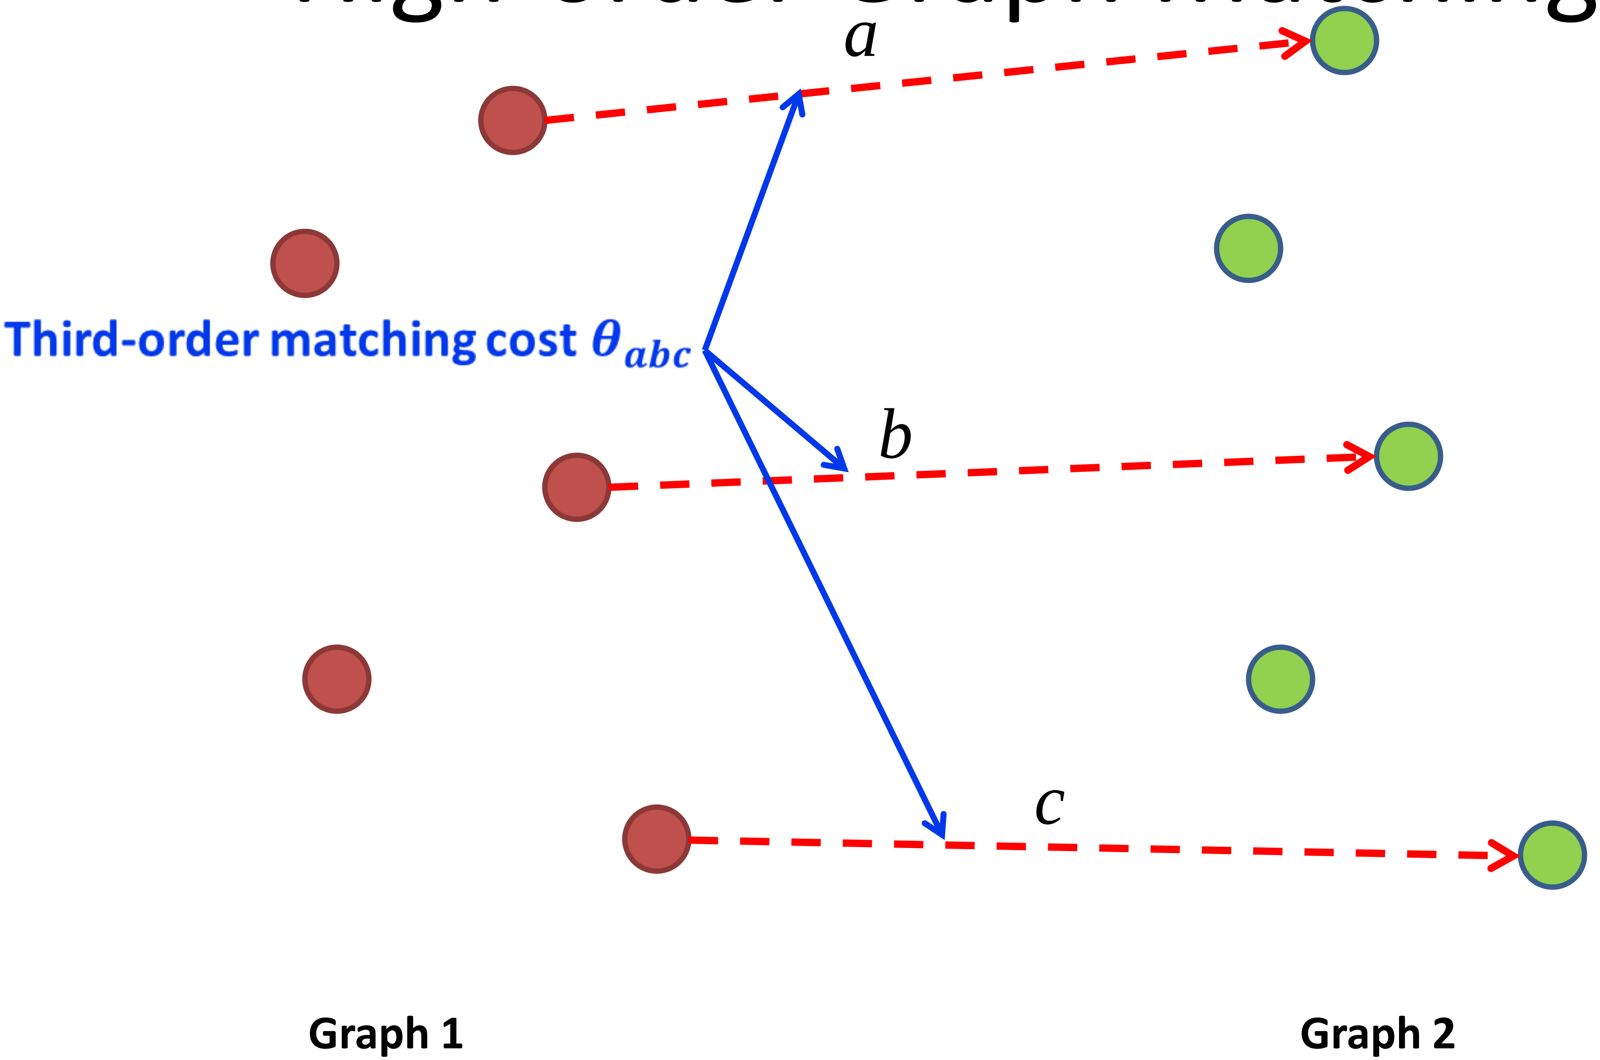

Experimental Results

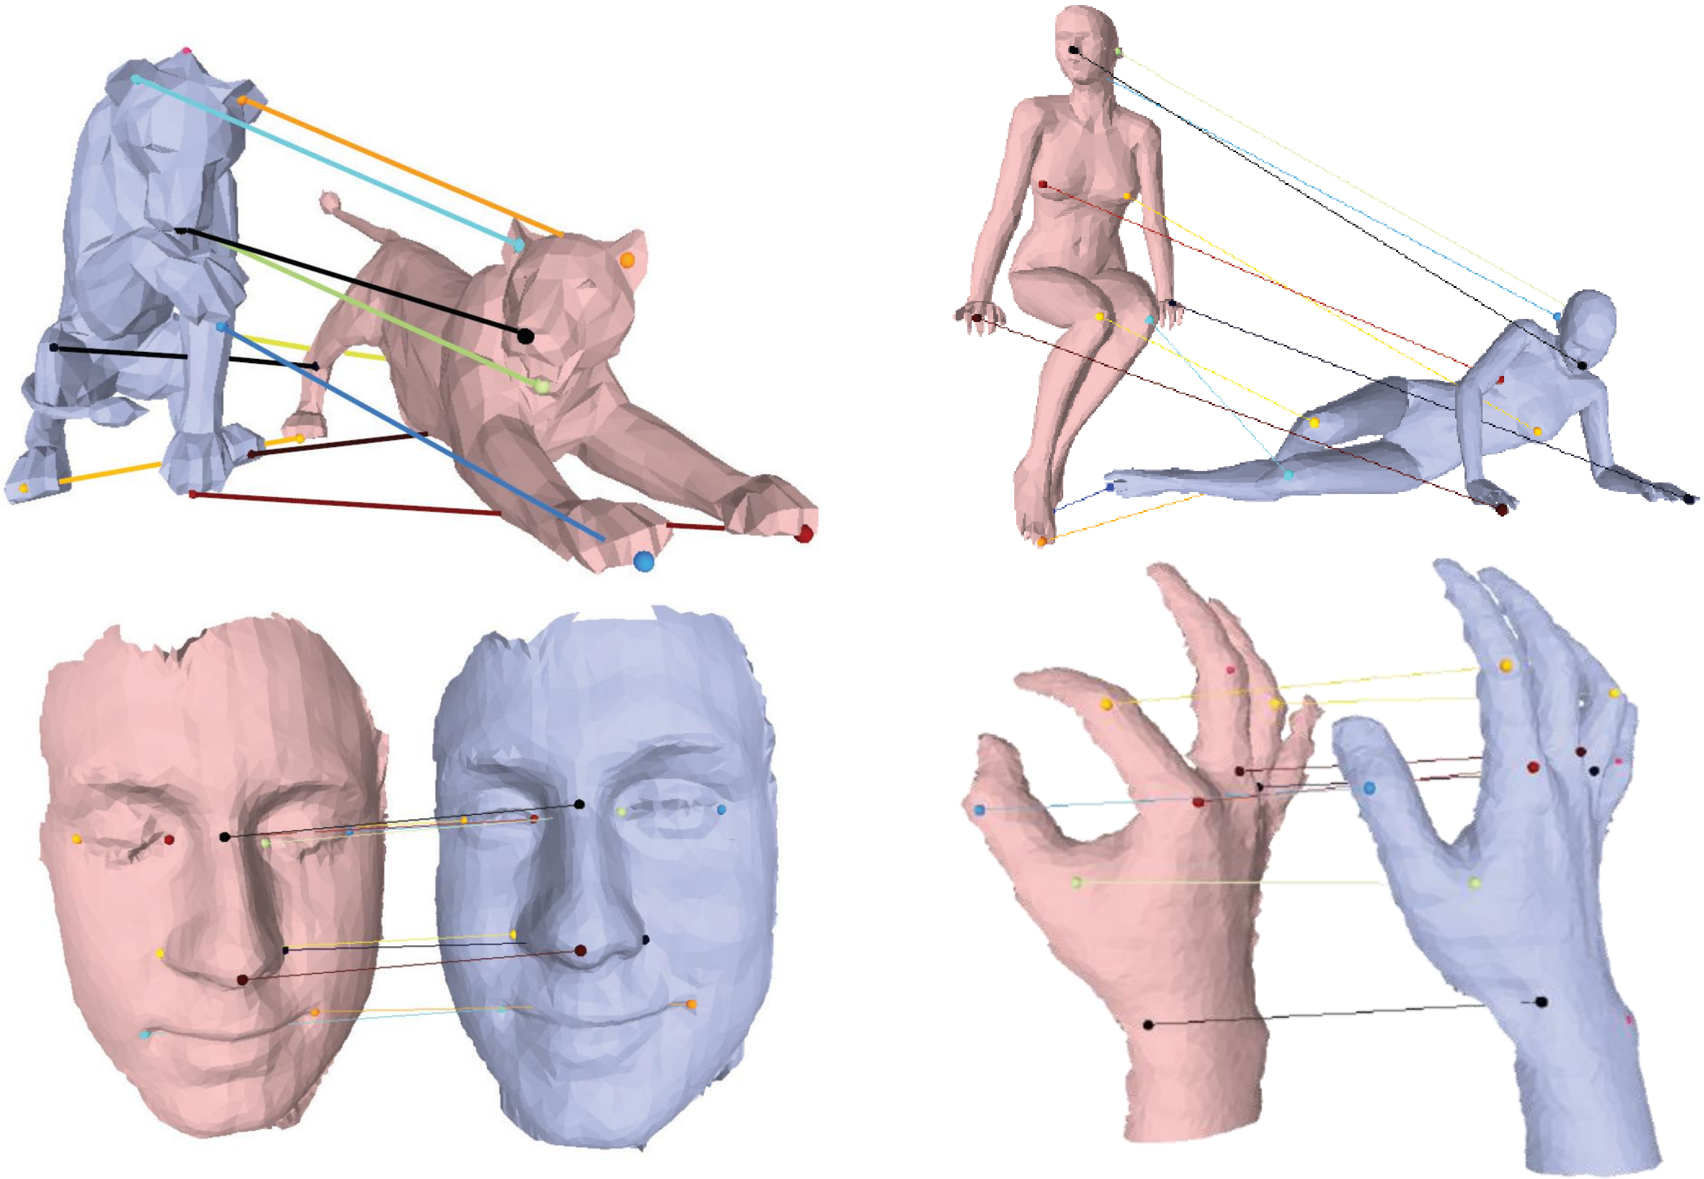

Experimental Results

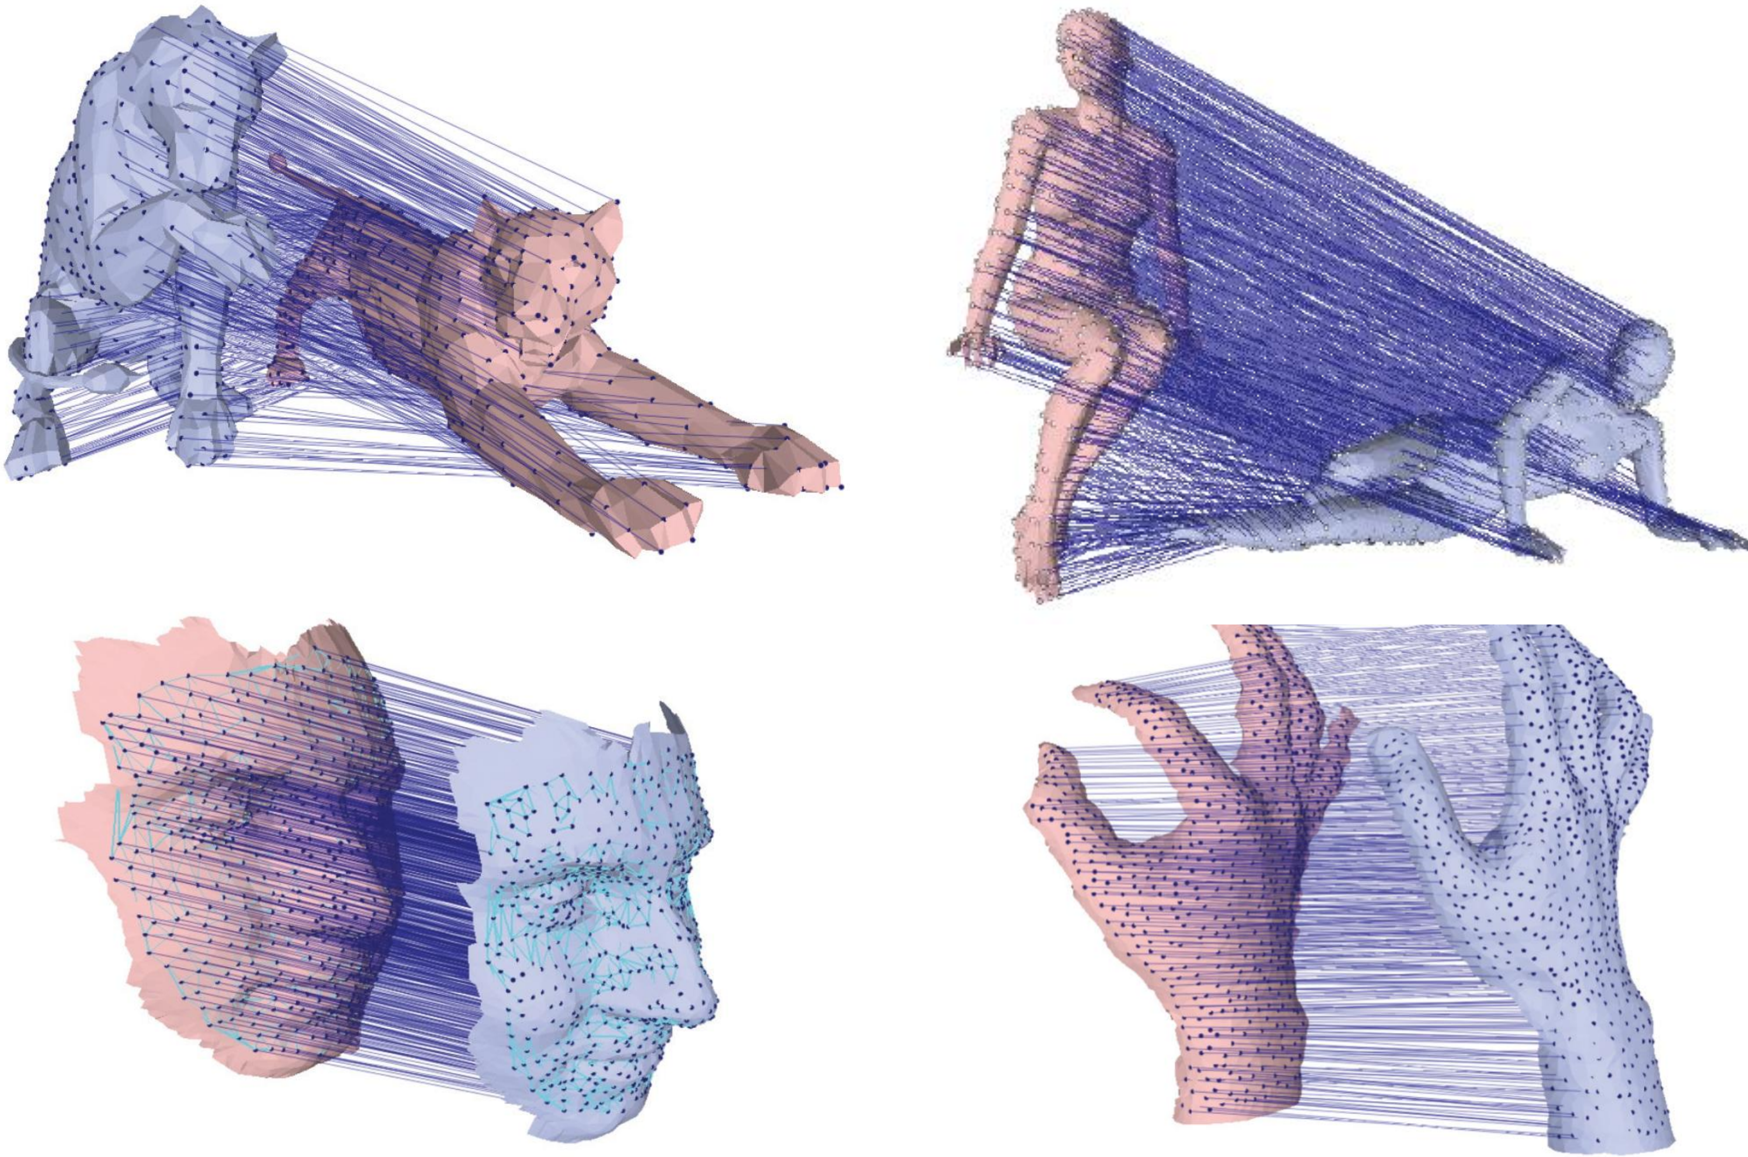

Discrete Artificial Vision

- Given:
  - Parameters $\mathcal{X}$ from a graph $\mathcal{G} = (\mathcal{X}, \mathcal{E})$
  - A neighborhood System $\mathcal{E}$
  - Discrete label set $\mathcal{L}$

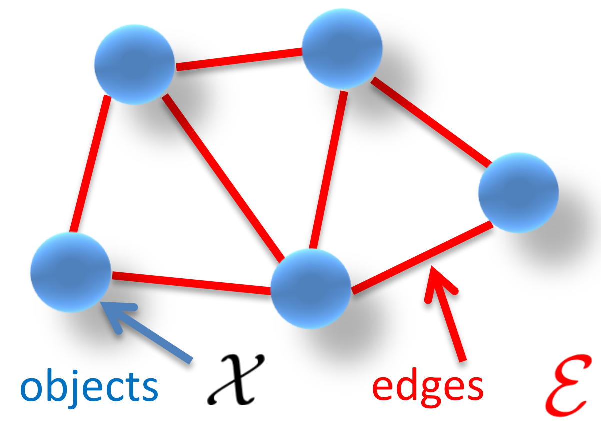

- Assign labels (to objects) that minimize the energy:

$$min_{x_p} \sum_{p \in \mathcal{X}} \underbrace{\Theta_p(x_p)}_{\text{unary potential}} + \underbrace{\Theta_{pq}(x_p, x_q)}_{\text{pairwise potential}}$$

- MRF optimization ubiquitous in vision (and beyond)

Optimization of high-order models

- Hypergraph $\mathcal{G} = (\mathcal{X}, \mathcal{C})$
  - Parameters $\mathcal{X}$
  - Hyperedges/cliques $\mathcal{C}$

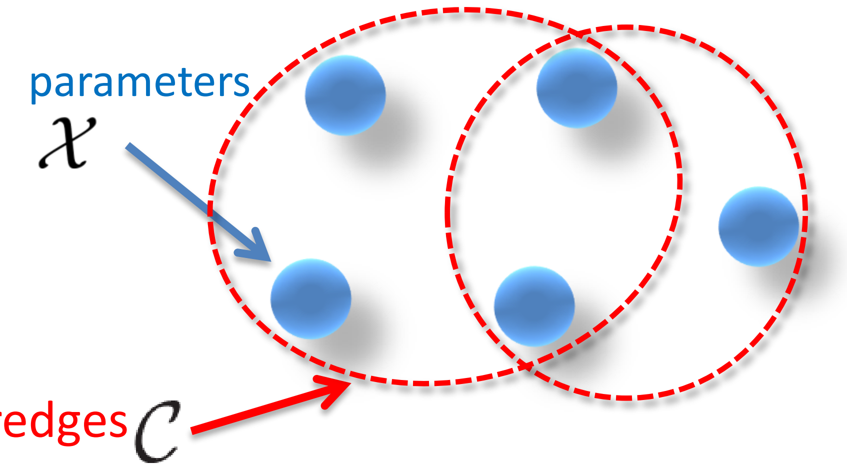

- High-order energy minimization problem

$$min_{x_p} \sum_{p \in \mathcal{X}} \Theta_p(x_p) + \Theta_c(x_p, \cdots, x_q)$$

unary potential
(one per node)

high-order potential
(one per clique)

Conclusions

- Discrete Graphical Models, is a promising answer to artificial vision
  - Curse of Dimensionality : Prior Knowledge either through anatomy of machine learning techniques towards dimensionality reduction
  - Curse of Non-linearity: Model Decomposition / Data association allows direct support estimation of parameter selection from the images
  - Curse of Non-Convexity: Regularization terms / dropping out of constraints can improve the optimality properties of the obtained solution
  - Curse of Non-Modularity: Model/Data Association/Inference Decomposition and use of gradient free methods

Future

- The future belongs to:
  - Higher order **structured** models (**Grammars**)
  - Message Passing Methods running on parallel architectures
  - Grammars [focus was up to now on inference and not on the design of the objective function]
    - Learning their parameters
    - Learning their derivation sequences

Discrete Image Registration

Nikos Paragios

Benjamin Glocker, Aristeidis Sotiras, Nikos Komodakis, Yangming Ou, Christos Davatzikos, Nassir Navab






Outline

1. Context
2. Image Registration
3. Hybrid Registration
4. *Symmetric Hybrid Registration*
5. Group-wise Registration
6. Conclusion

Context

Image Registration: Definition

- Establish correspondences between images.

- Estimate transformation so that the images are aligned.

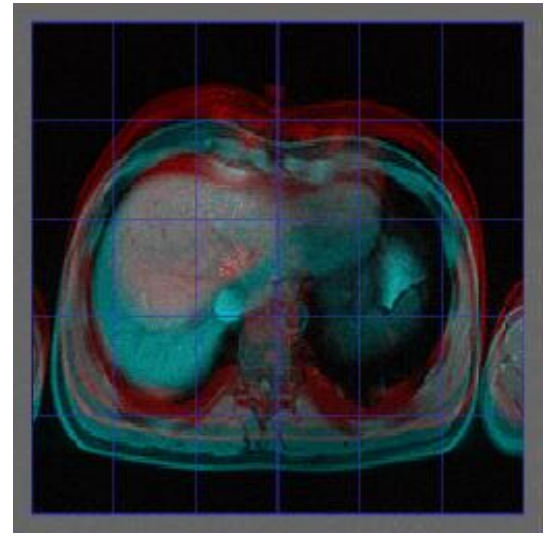

Synonyms:

Image alignment; Image fusion; Image matching; Motion estimation; Optical flow; Image correspondence problem

Challenges

- Important time constraints
- Important volume of data
- Increasing data dimensionality
- Vast range of applications

Efficiency (important time constraints, important volume of data, increasing data dimensionality)

Vast range of applications → Versatility

Image Registration

Image Registration

Energy minimization:

$$\mathbf{\Theta}^* = \arg\min_{\mathbf{\Theta}} M\left(\Phi(\mathrm{T}), \Phi(\mathrm{S})\circ \mathcal{T}(\mathbf{\Theta})\right) + R\left(\mathcal{T}(\mathbf{\Theta})\right)$$

T: fixed image or target

S: moving image or source

$\mathcal{T}$: transformation parameterized by $\mathbf{\Theta}$

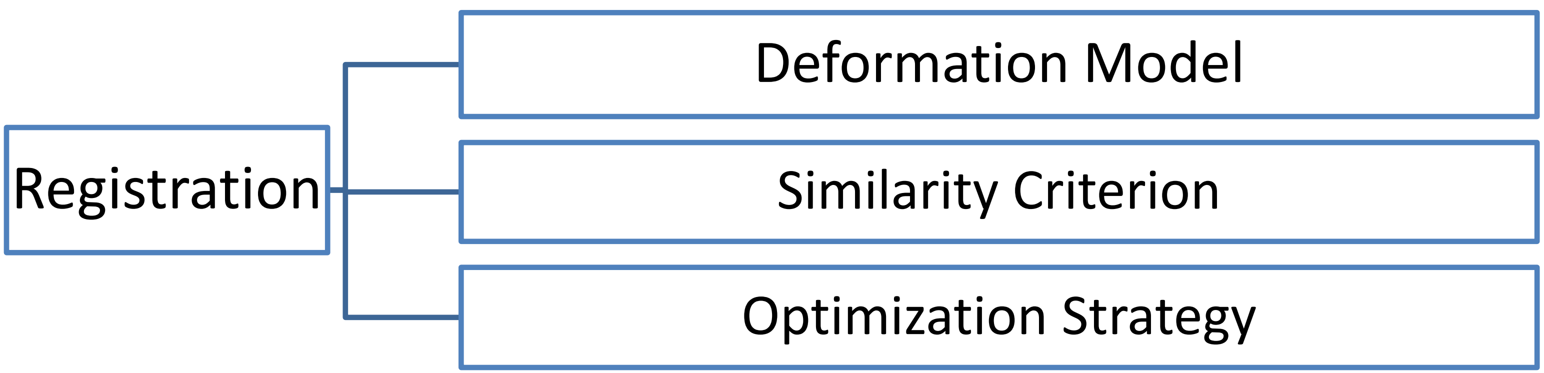

Image Registration: Deformation Model

Physical models

- Elastic **[Bajscy 89, Davatzikos 97]**
- Viscous fluid **[Christensen 96]**
- Diffusion **[Thirion 98, Pennec 99, Vercauteren 07]**
- Diffeomorphisms **[Joshi 00, Beg 05]**

→ **Computational intensive**

Interpolation theory

- Radial Basis Functions **[Bookstein 91, Rohr 01, Rohde 03]**
- Piecewise Affine **[Pitiot 06, Arsigny 05]**
- Free Form Deformations **[Rueckert 99, Rueckert 06]**

→ **Fewer degrees of freedom**

Constraints

- Topology preservation **[Droske 03, Noblet 05, Haber 06, Sdika 08]**
- Volume preservation **[Rohlfing 03, Haber 04, Mansi 11]**
- Rigidity constraints **[Loeckx 04, Staring 07, Modersitzki 07]**

→ **Task specific**

Our strategy

- Cubic *B*-Splines Free Form Deformations
  - Efficient
  - Implicit regularization
  - Topology preservation through hard constraints
  - Coarse-to-fine scheme to capture complex deformations

Image Registration: Similarity Criterion

Iconic

- Mono-m
  - Int
    CC
  - Fea
    [Ou
- Multi-mo
  - Inf
    [W
    [St
    [Pl
    [He
    [Ru
  - Re
    Simulate one modality **[Wein 08, Michel 10]**, Map to a common space **[Maintz 01, Haber 07, Michel 11]**

→ **Dense solution, computational intensive**
→ **Initial conditions, treat all points equally**

Geometric

s: Procrustes
BFs **[Rohr 01]**,
[Guo 06]
nces: Affine
eikkila 04],
nes 04, Tsin
ce functions

ajczyk 05]
eordeanu 05,
Wang 10]

- Spatial transformation + correspondence
  - ICP **[Besl 92]**, EM-ICP **[Granger 02]**, TPS-RPM **[Chui 03]**

→ **Sparse, limited accuracy**
→ **Efficient, robust to large deformations**

Our strategy

- Hybrid

→ **Exploit both intensity and geometric info**
→ **Get the best of both worlds**

Image Registration: Optimization Strategy

Continuous

- GD **[Rueckert 99, Droske 03]**
- Conjugate gradient **[Miller 01, Joshi 07]**
- GN **[Vercauteren 09]**
- LM **[Thevenaz 00, Kybic 03]**
- Stochastic GD **[Klein 07, Balci 07]**

→ Differentiable functions, local search

Discrete

- Graph-based **[Tang 07, Liao 11]**
- Message passing **[Murphy 99, Felzenszwalb 06]**
- Linear-programming approaches **[Glocker 08, Kwon 08, Zikic 10]**

→Non-differentiable functions, global search

Miscellaneous

- Heuristics and Meta-heuristics **[Shen 02, Xue 04, Liu 04]**
- Evolutionary methods **[Klein 07, Santamaria 11]**

→ No optimality guarantees
→ General

Our strategy

- Discrete: MRF formulation

$$E_{MRF} = \sum_{p \in V} U_p(l_p) + \sum_{pq \in E}$$

Image Registration: Optimization Strategy

Continuous

- GD **[Rue**
- Conjugat
- GN **[Ver**
- LM **[The**
- Stochast

→ Differen

Discrete

1]

), Felzenszwalb

hes **[Glocker**

global search

Miscellan

- Heuristics and Meta-heuristics **[Shen 02, Xue 04, Liu 04]**
- Evolutionary methods **[Klein 07, Santamaria 11]**

→ No optimality guarantees
→ General

Our strategy

- Discrete: Belief Propagation **[Alchatzidis 11]**
  - Generality
  - Optimality
  - Per-instance approximation factors
  - Speed

Hybrid registration

Basic Idea of Intensity-based Registration

- Image registration as an optimization problem

$$T^* = \arg\min_T \phi\,(I, J \circ T)$$

- Target and source Image:

$$I, J : \Omega \subset \mathbb{R}^d \mapsto \mathbb{R}$$

- Transformation:

$$T(\mathbf{x}) = \mathbf{x} + D(\mathbf{x})$$

- Image metric:

$$\phi : (I, J) \mapsto \mathbb{R}$$

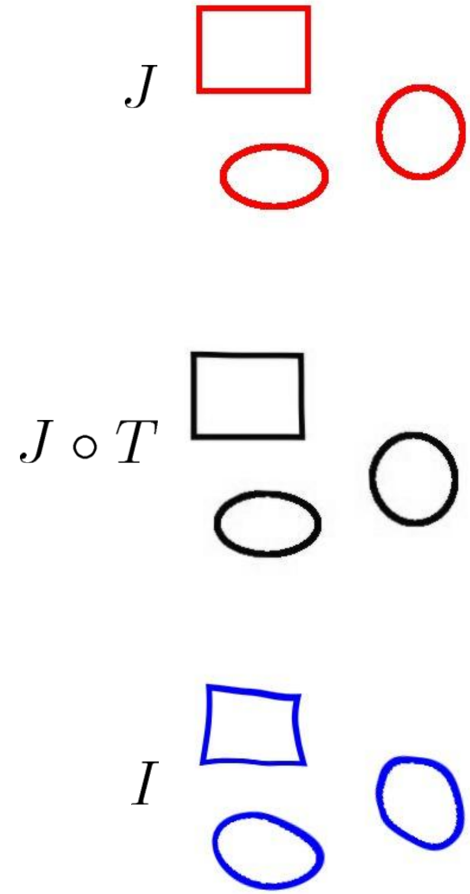

Dimensionality Reduction

- Linear combination of control points

$$D(\mathbf{x}) = \sum_{i}^{M} \eta(\mathbf{x})\, \mathbf{d}_i$$

η : basis functions $\mathbf{d}$: displacements

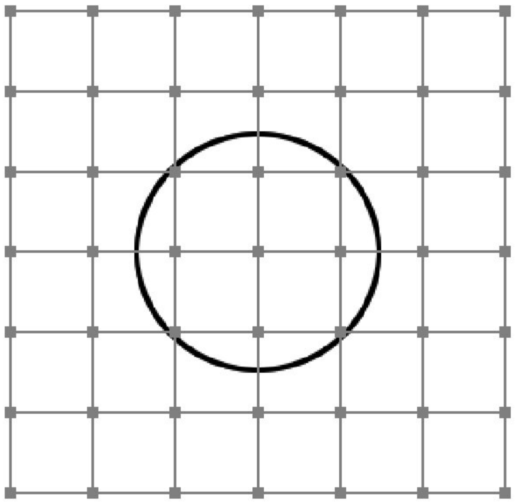

- e.g. Free-Form Deformations (Sederberg et al. 1986; Rueckert et al. 1999)

(Weighted) Block Matching

- Redefinition of data term w.r.t. control lattice

$$E_{\text{data}}(D) = \sum_{i}^{M} \int_{\Omega} \hat{\eta}(\mathbf{x}) \; (I(\mathbf{x}) - J(\mathbf{x} + D(\mathbf{x})))^2 \, d\mathbf{x}$$

$$\text{with} \quad \hat{\eta}(\mathbf{x}) = \frac{\eta(\mathbf{x})}{\int_{\Omega} \eta(\mathbf{y}) \, d\mathbf{y}}$$

- Pixel-wise image metrics weighted by normalized basis functions
  - image points closer to a control point gain more influence on its matching energy
- Statistical image metrics (e.g. mutual information, cross correlation)
  - evaluation of image metric in local patches centered at the control points
  - block size depends on control lattice resolution

Discrete Labeling Problem

- Markov Random Field formulation with pairwise interactions

$$E_{\text{mrf}}(\mathbf{l}) = \sum_{p \in G} V_p(l_p) + \sum_{(p,q) \in N} V_{pq}(l_p, l_q)$$

Edges

Nodes

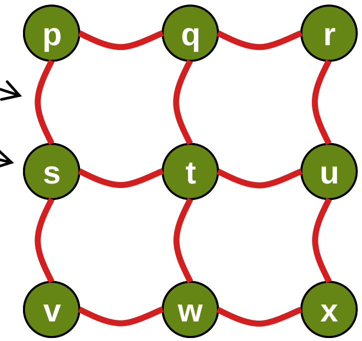

- Unary potentials (matching):

$$V_p(l_p) = \underbrace{\int_{\Omega} \hat{\eta}(\mathbf{x}) \left(I(\mathbf{x}) - J(\mathbf{x} + \mathbf{d}^{l_p}) \right)^2 d\mathbf{x}}_{\text{or any other local image metric}}$$

- Pairwi

$$V_{pq}(l_p, l_q) = \lambda \left\| \mathbf{d}^{l_p} - \mathbf{d}^{l_q} \right\|$$

Deformable Registration by Discrete Optimization

Low-dimensional deformation model (B-Spline FFD)

Update computation:

1. For each control point CP_i

   For a discrete number of displacements d^{l_p}
   evaluate approximative change in similarity measure

$$V_p(l_p) = \underbrace{\int_\Omega \hat{\eta}(\mathbf{x}) \left(I_\mathrm{T}(x) - I_\mathrm{S}(x + d^{l_p})\right)^2 \,\mathrm{d}x}_{\text{or any other local image metric}}$$

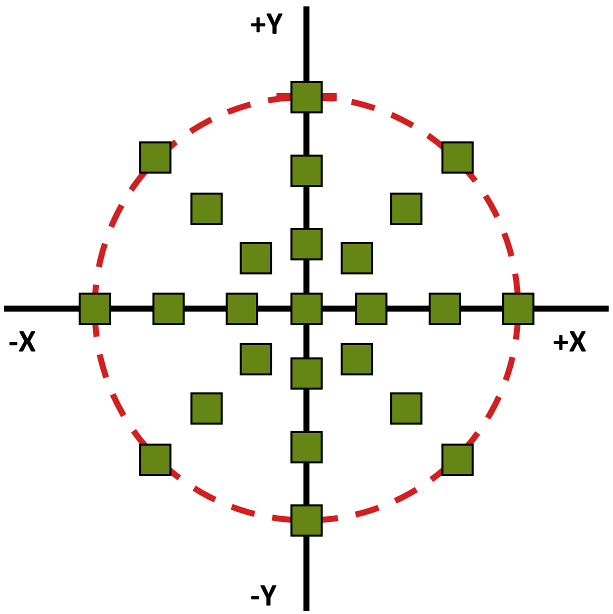

2. Compute approximately optimal combination of the pre-computed displacements w.r.t. chosen regularization with fast and accurate discrete optimization techniques

$$E_\mathrm{mrf}(\mathbf{l}) = \sum_{p \in G} V_p(l_p) + \sum_{(p,q) \in N} V_{pq}(l_p, l_q)$$

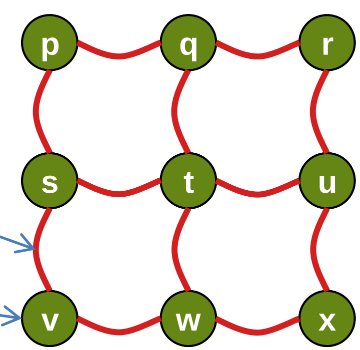

[Glocker 2008]

Deformable Registration by Discrete Optimization

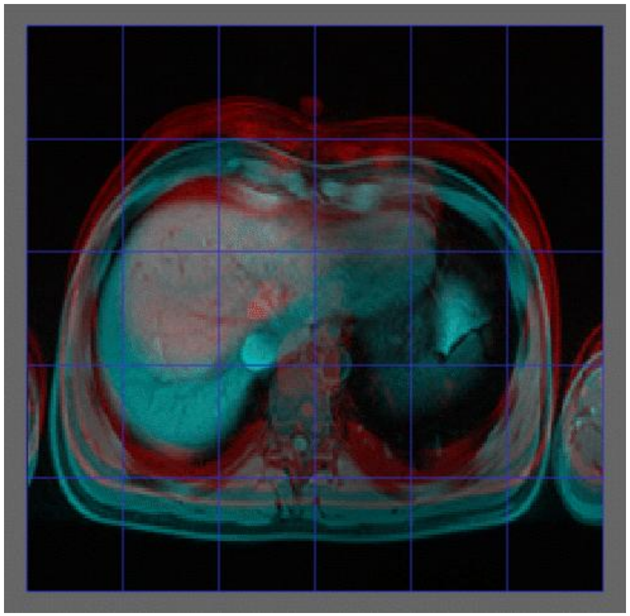

Properties:

– No derivative computation required

– Similar efficiency for any difference measure

– Larger/non-local search range for each CP

→ increased capture range

Related Work

Initialization

- Surface-iconic **[Liu 04, Postelnicu 09, Gibson 09]**
- Landmark-iconic **[Johnson 02, Auzias 11]**
- Segmented structure-iconic **[Camara 07]**

→ Independent solutions
→ robustness, no coupling guarantee

Constraint

- Soft sparse constraints **[Hartkens 02, Hellier 03, Papademetris 04, Rohr 04, Avants 06]**
- Soft dense constraints **[Worz 07, Biesdorf 09, Azar 06]**
- Hard constraint **[Joshi 07]**

→ One-way flow of information

Coupled

- Landmark-iconic **[Cachier 01]**
- Surface-iconic [**Joshi 09]**

→ One objective function
→ Mono-modal
→ Constraints on landmarks
→ Spherical geometries

Our strategy

- Coupled approach
  - Discrete framework
  - One-shot optimization
  - Any similarity criterion
  - Diffeomorphic
  - No constraints on landmark

Hybrid Registration

Input

- Two images
  - Source S
  - Target T
- Two sets of landmarks
  - Landmarks in source $K = \{\kappa_1, \ldots, \kappa_n\}$

Output

- Landmark correspondences T_{geo}
- Dense deformation field T_{ico}
  - Parametrized by Cubic *B*-spline FFDs

➢ Solutions are obtained simultaneously and are consistent

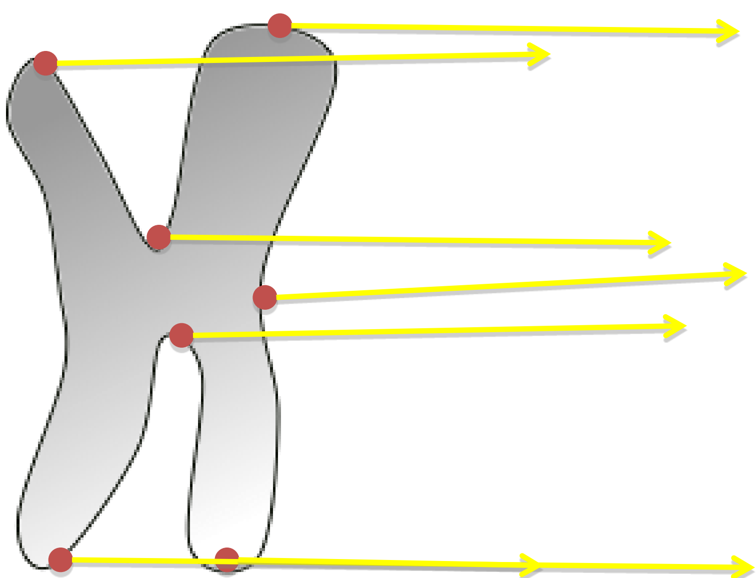

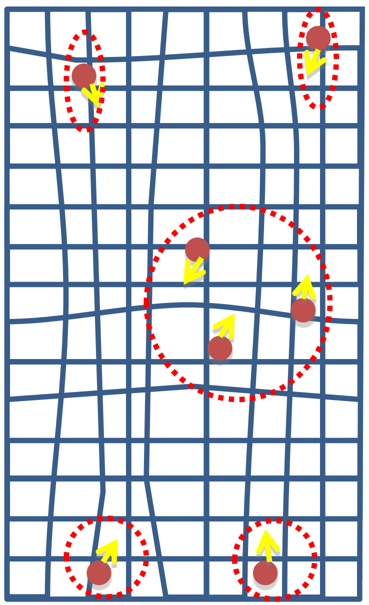

Graph Matching

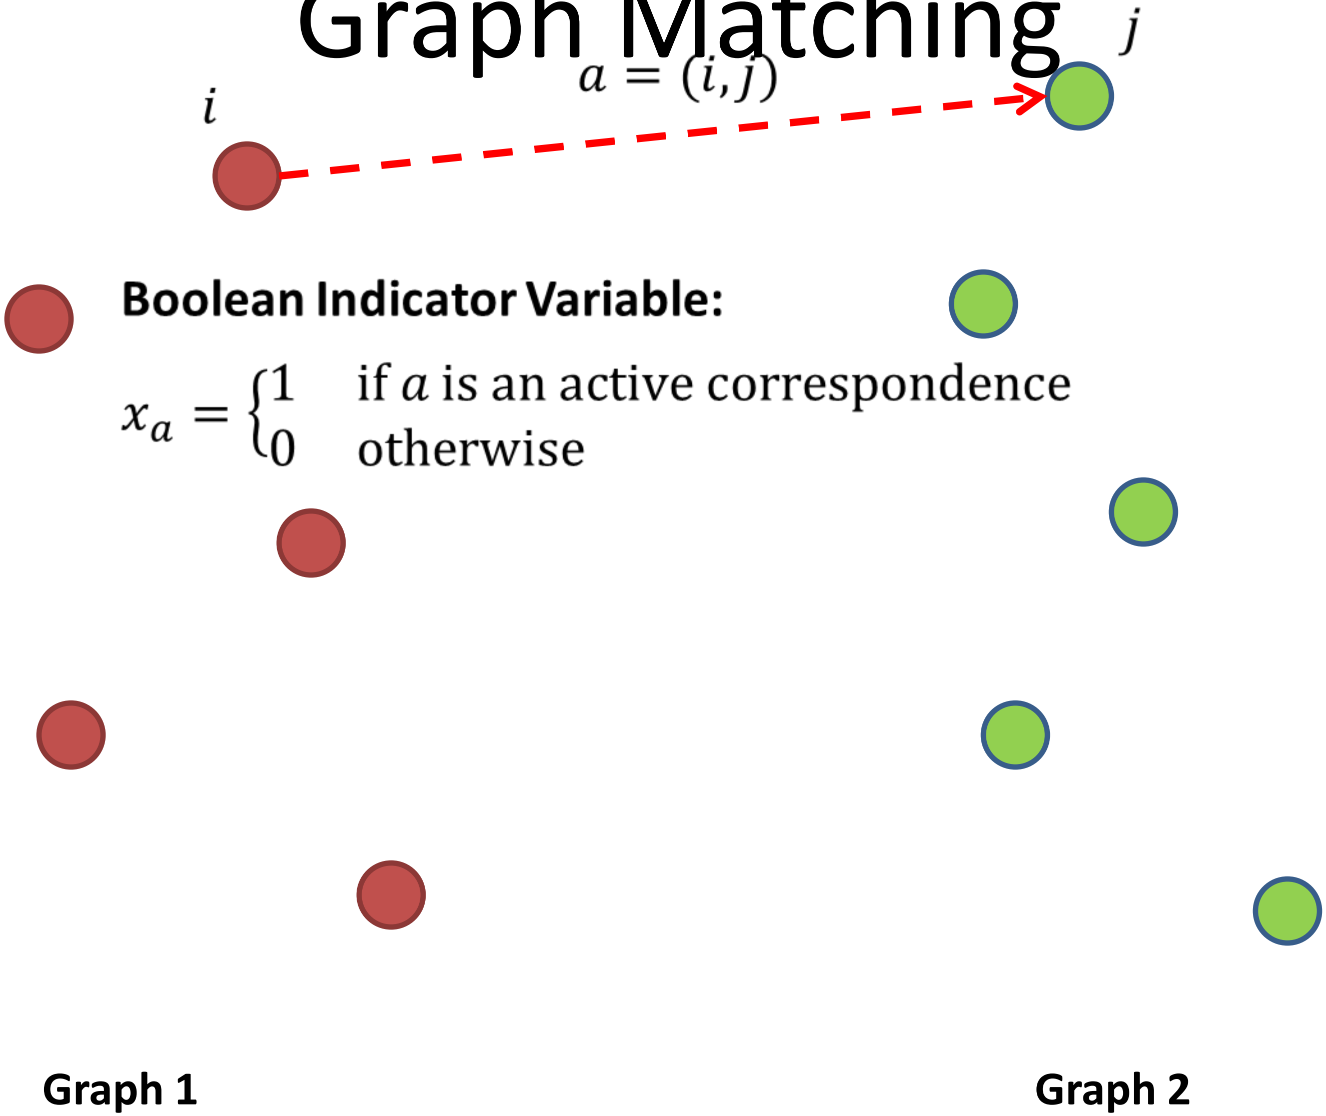

Boolean Indicator Variable:

$$x_a = \begin{cases} 1 & \text{if } a \text{ is an active correspondence} \\ 0 & \text{otherwise} \end{cases}$$

Graph Matching

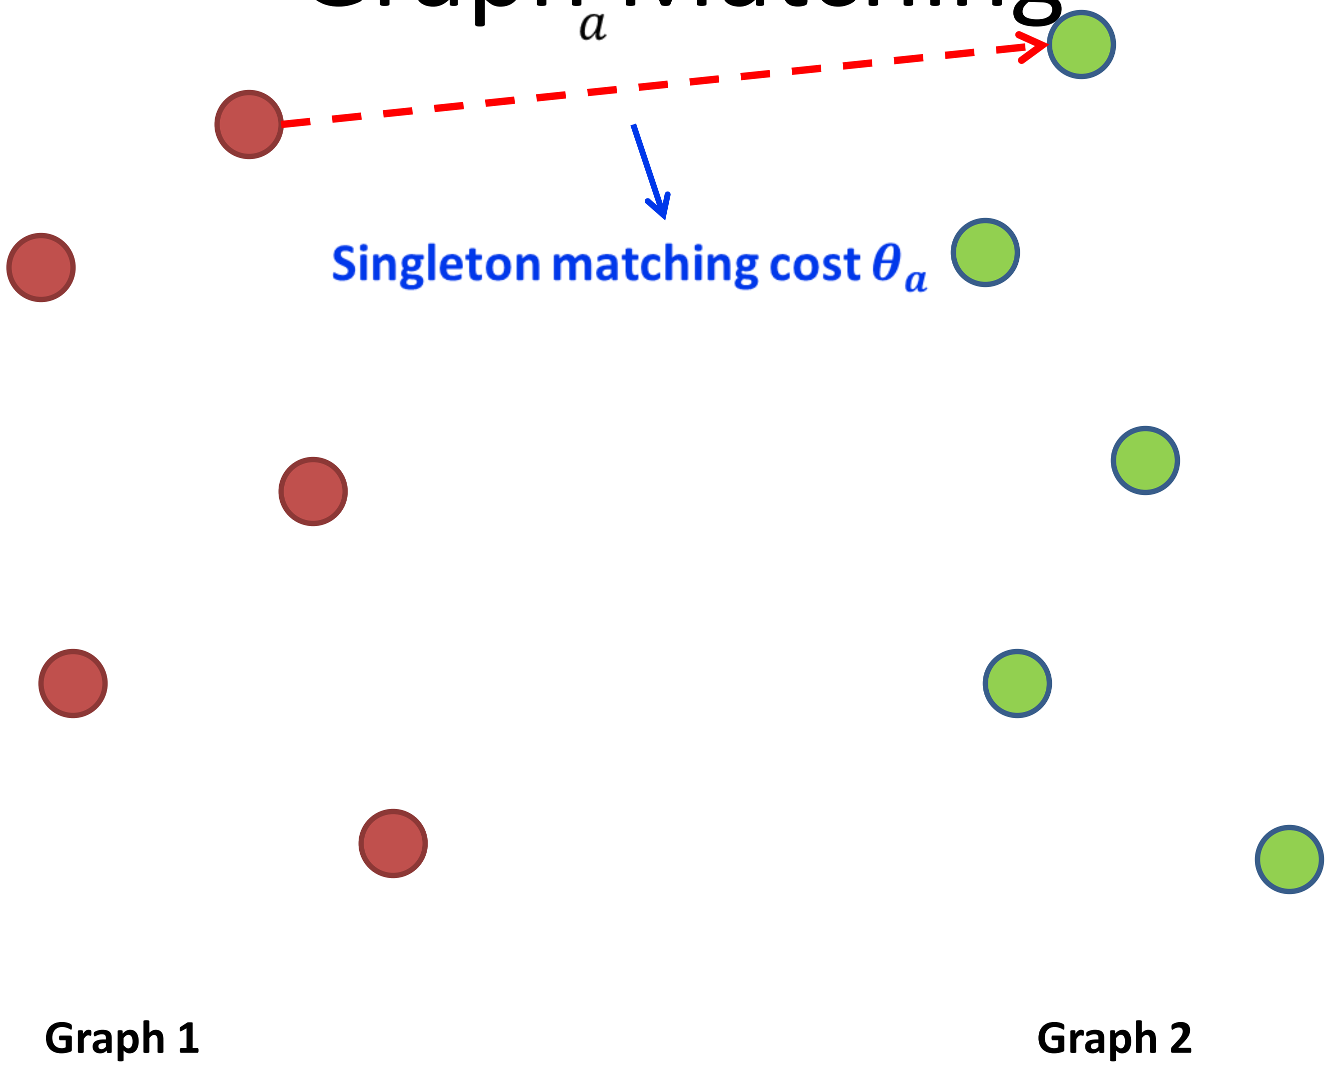

Graph Matching

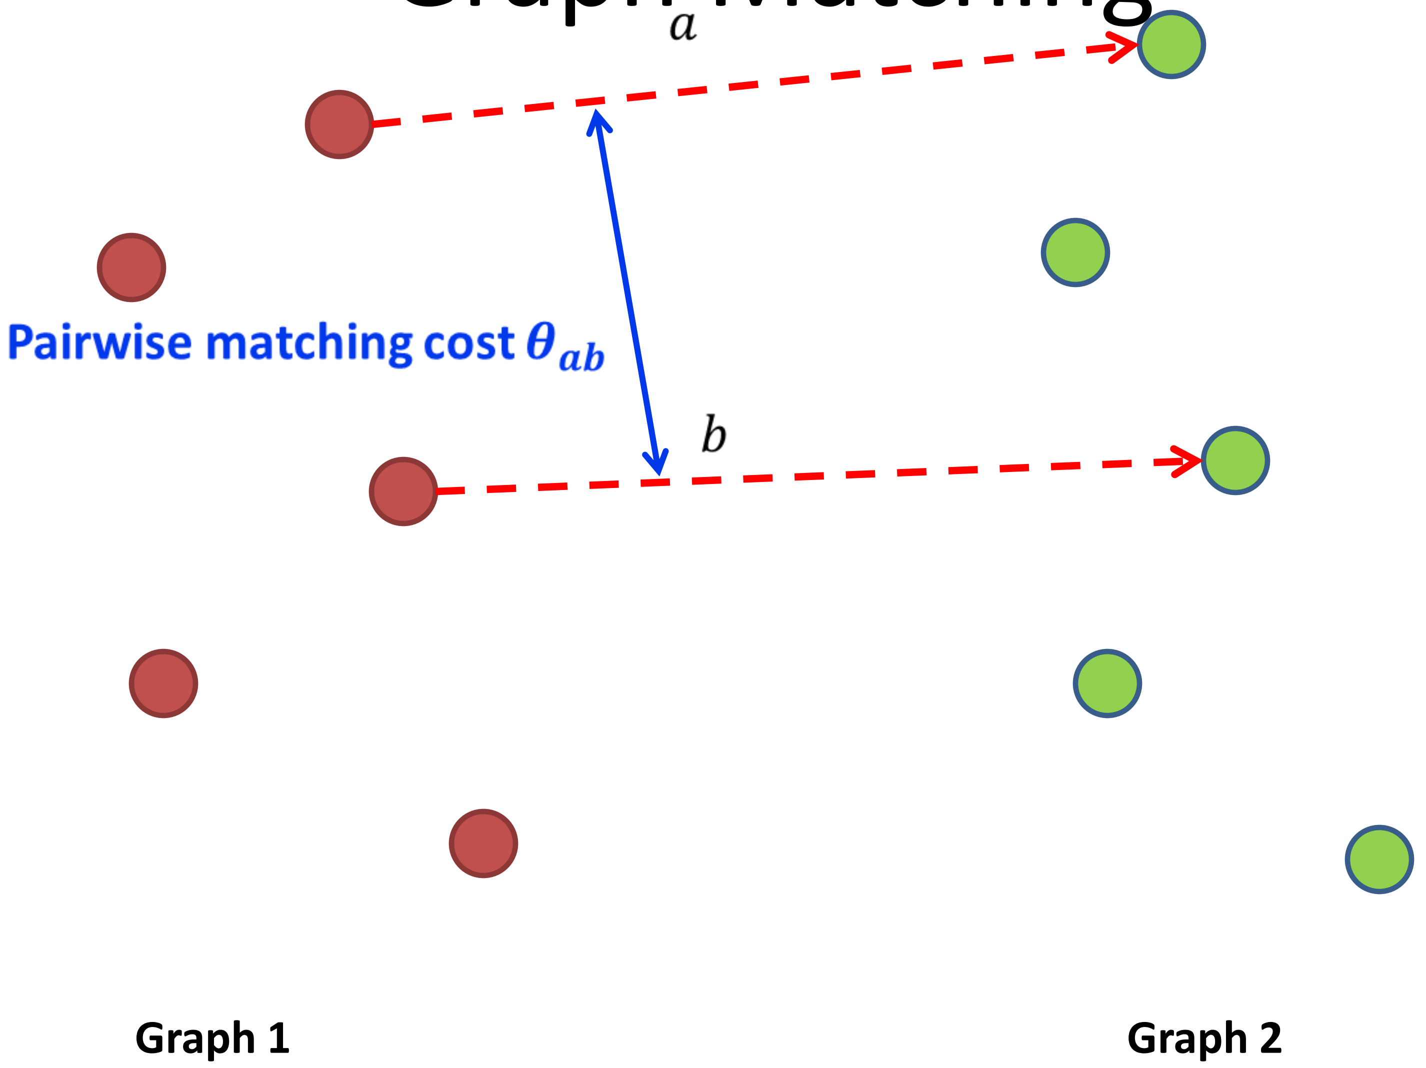

Geometric Part – Graph

$$E_{geo} = \sum_{p \in V_{geo}} U_{geo}(l_p) + \sum_{pq \in E}$$

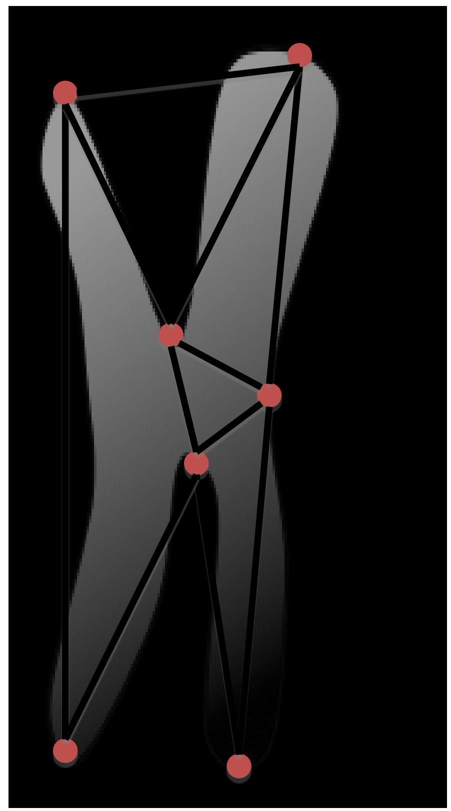

Landmark correspondence

Pair-wise interaction - geometric constraint

Geometric Part – Data Term

$$U\downarrow geo\,(l\downarrow p)=\varrho(\kappa\downarrow p\,,\lambda$$

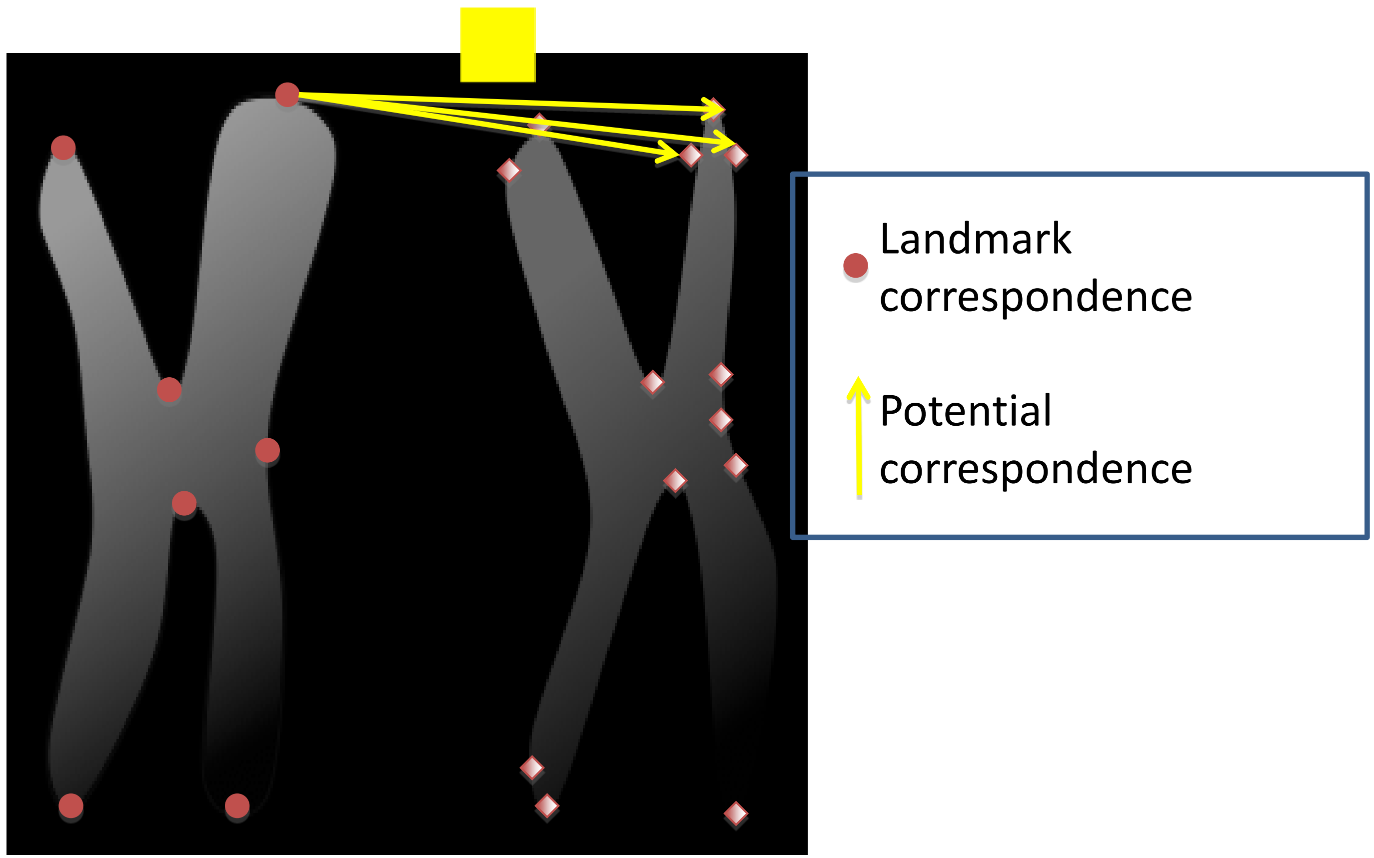

Geometric Part – Regularization Term

$$P_{geo}(l_p, l_q) = \|(\boldsymbol{\lambda}_{l_p} - \boldsymbol{\lambda}_{l_q}) - (\boldsymbol{\kappa}_p$$

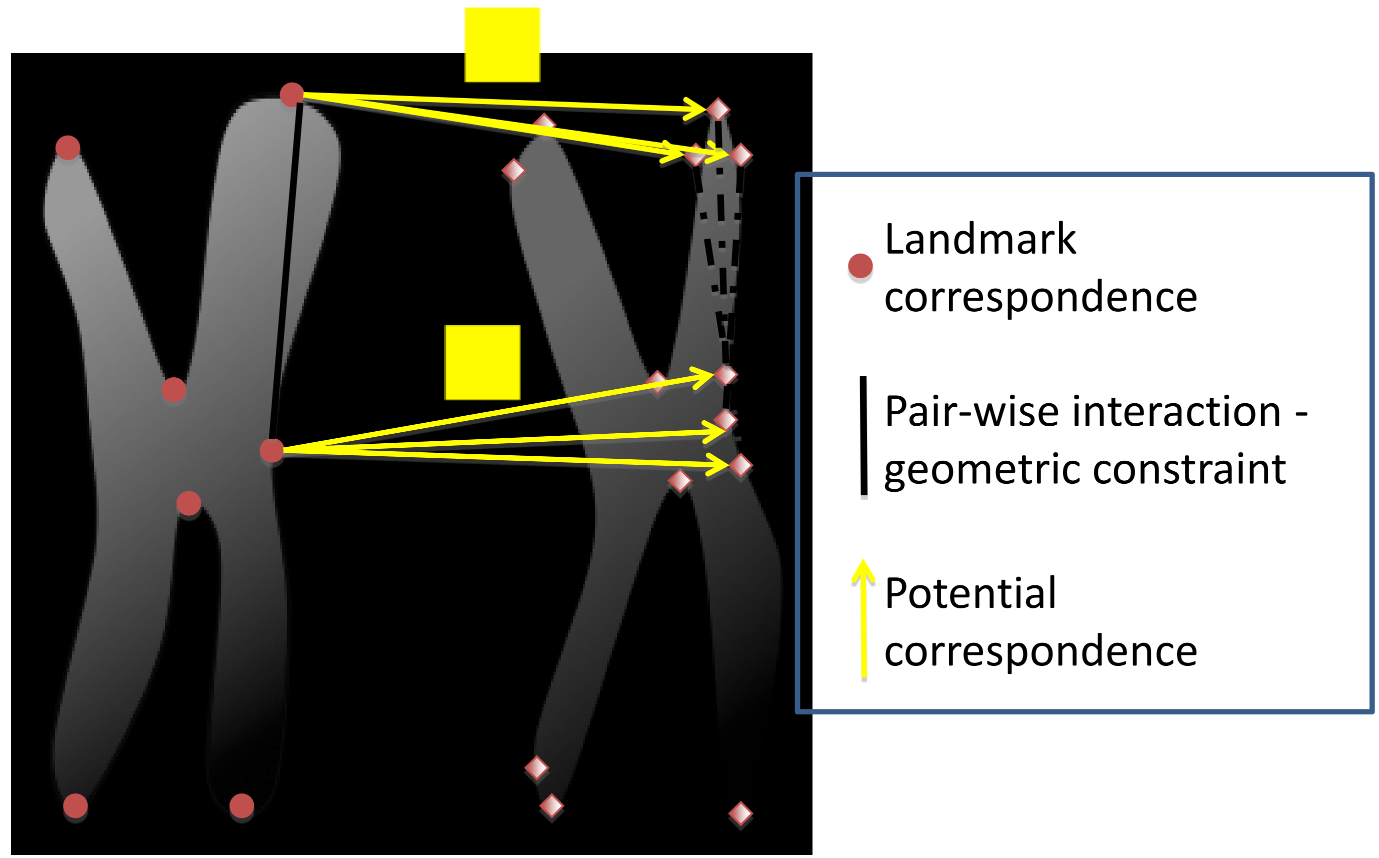

Iconic Part – Graph

$$E_{ico} = \sum_{p \in V_{ico}} U_{ico}(l_p) + \sum_{pq \in E}$$

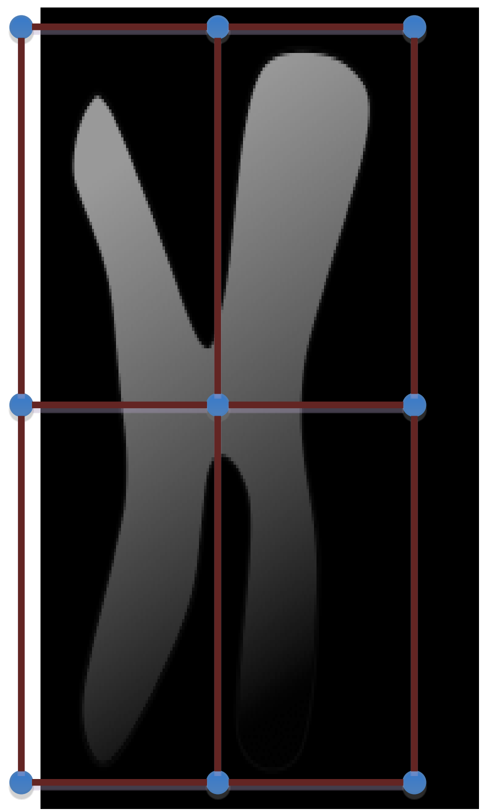

- Deformation grid node displacement
- Pair-wise interaction - smoothness constraint

Iconic Part [Glocker 08]

Deformation grid – Cubic *B*-spline FFD:

$$T(\boldsymbol{x}) = \boldsymbol{x} + \sum_{i=1}^{k} \omega_i \, (x$$

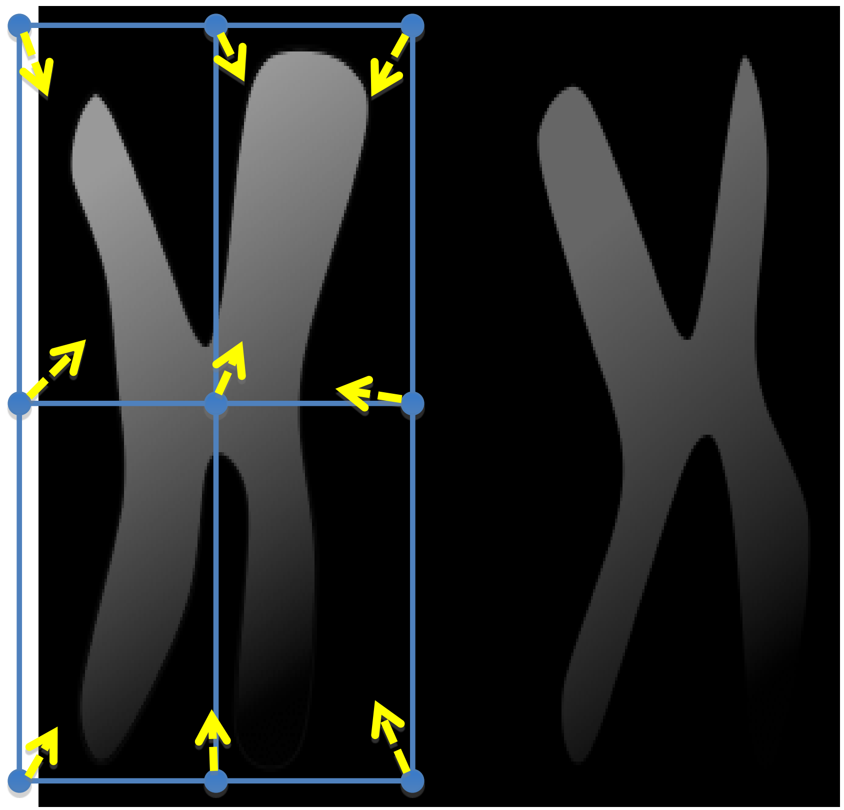

Iconic Part – Data Term

$$U_{ico}(l_p) = \int \omega_p(\boldsymbol{x}) \rho(S \circ T_{ico,l_p}(\boldsymbol{x}), \mathrm{T}($$

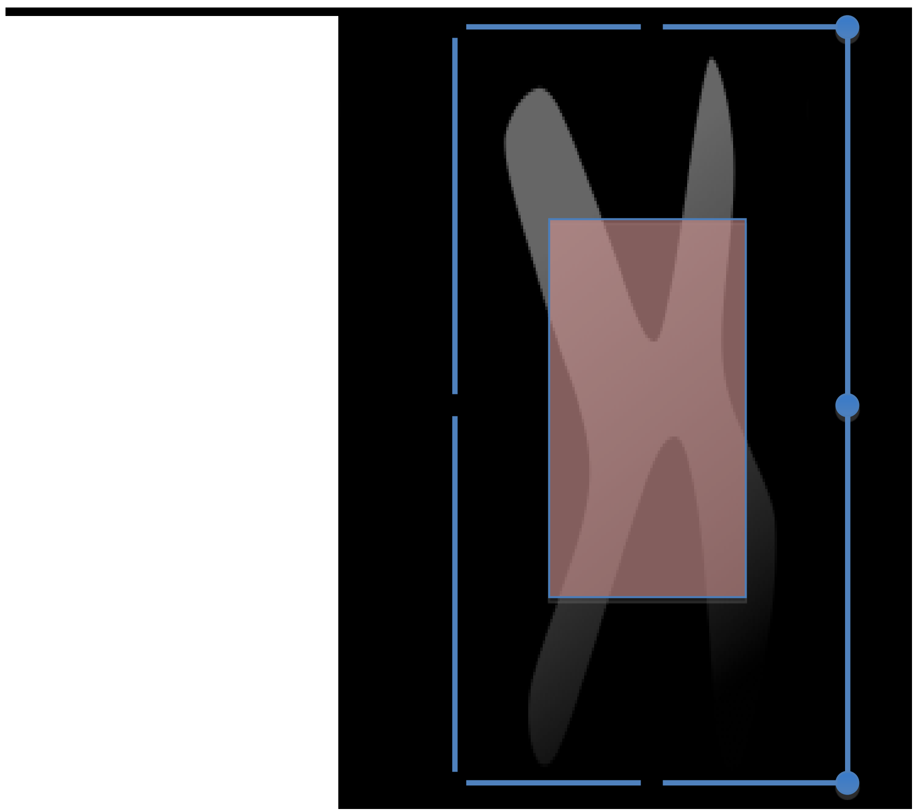

Deformation grid node displacement

Iconic Part – Regularization Term

Elastic regularization: $P_{ico}(l_p, l_q) = \|(\mathbf{d}_p + l_p) - (\mathbf{d}_q +$

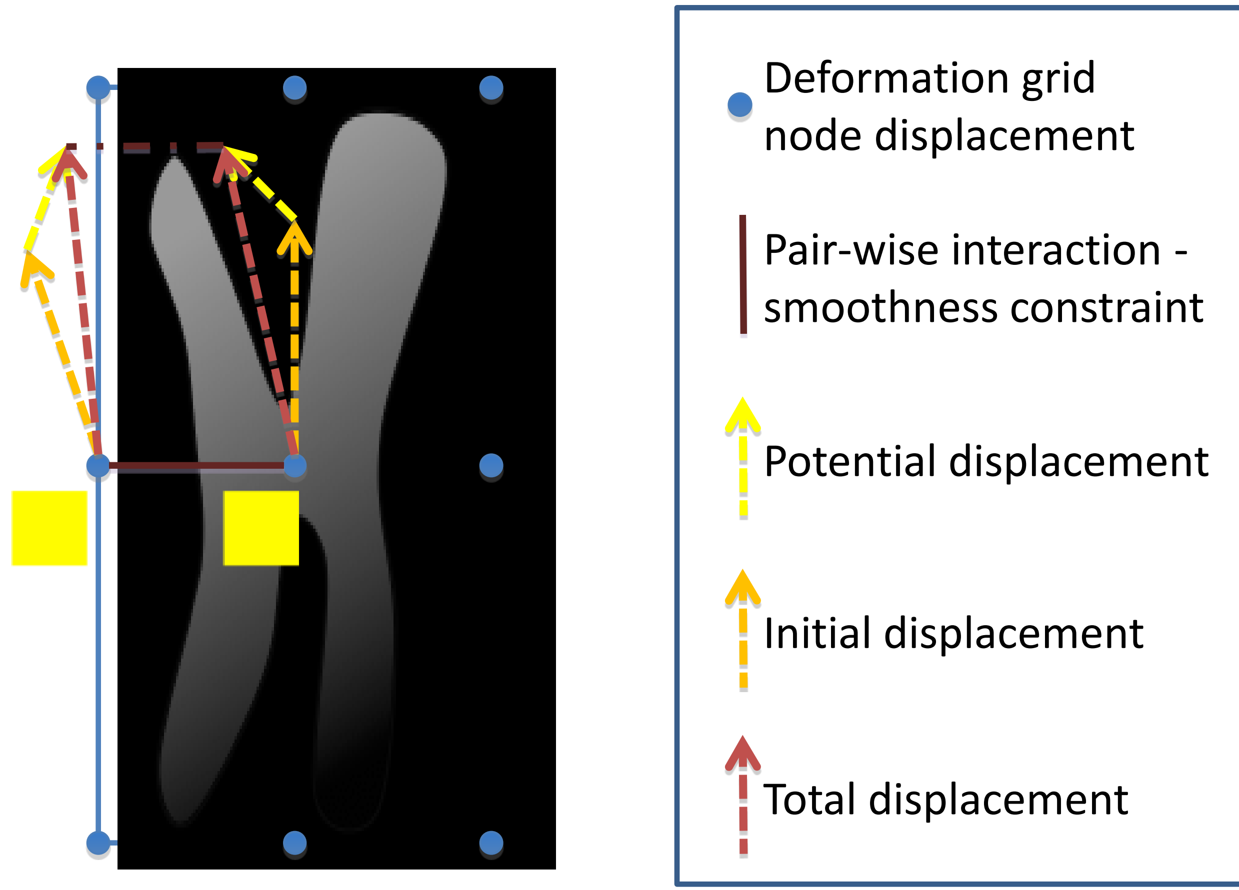

Hybrid Part - Graph

$$E_{coupling}=\sum_{pq\in E_{hyb}} P$$

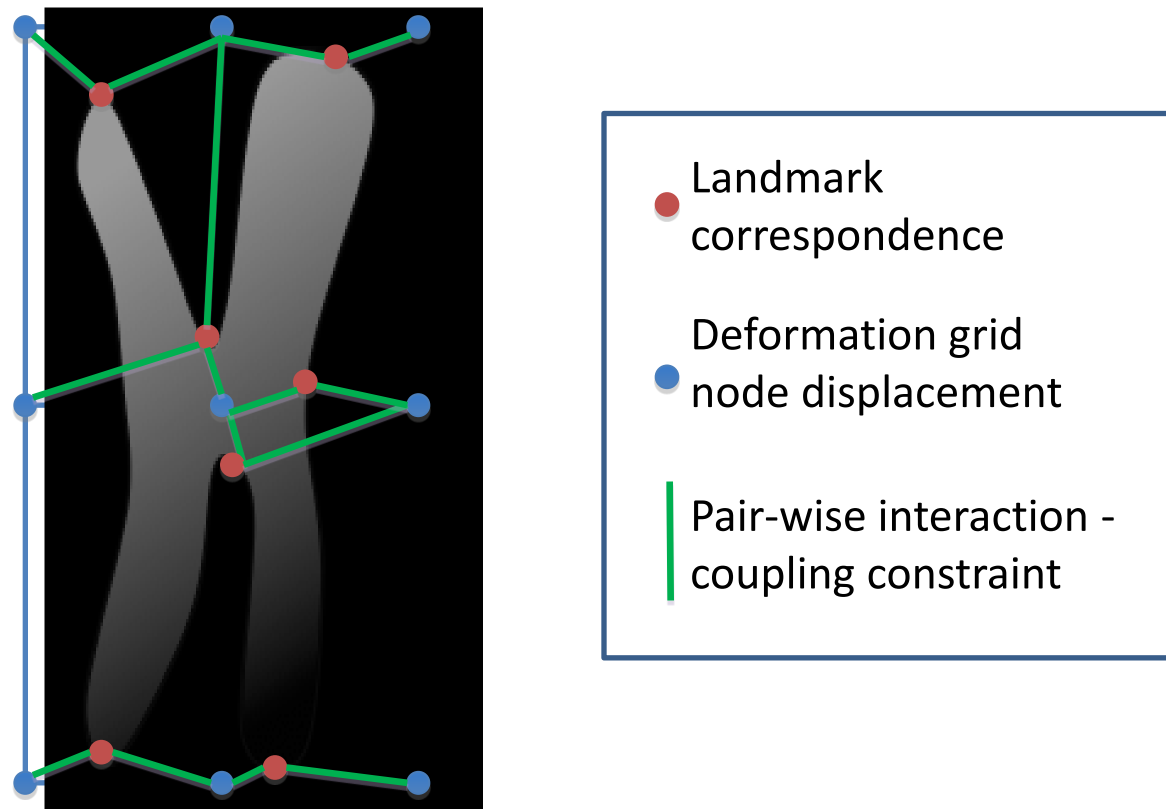

Hybrid Part – Coupling Constraint

$$P_{hyb}(l_p, l_q) = \omega_q(\boldsymbol{\kappa}_p) \|(\boldsymbol{\lambda}_{l_p} - \boldsymbol{\kappa}_p) - (\boldsymbol{a}$$

Landmark correspondence

Deformation grid node displacement

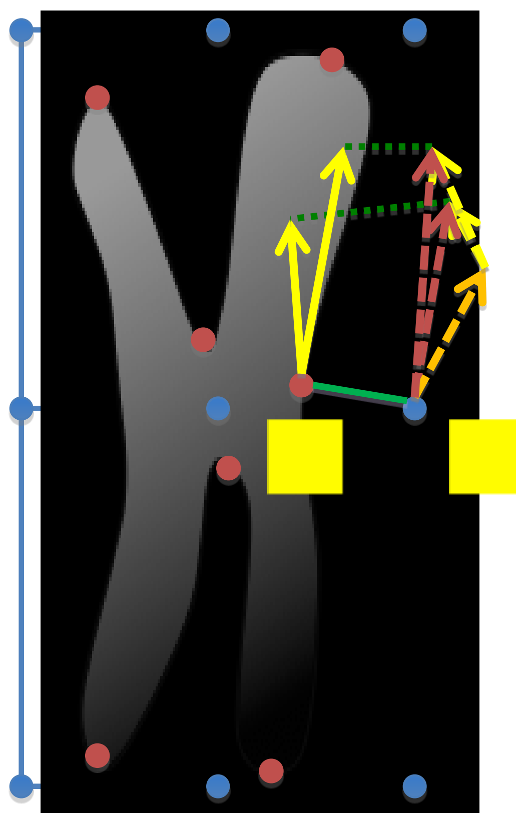

Potential correspondence

Pair-wise interaction - coupling constraint

Potential displacement

Initial displacement

Total displacement

Hybrid Registration - Graph

$$G_{hyb}=(V_{ico}\cup V_{geo}\ ,E_{ico}\cup E_{geo}$$

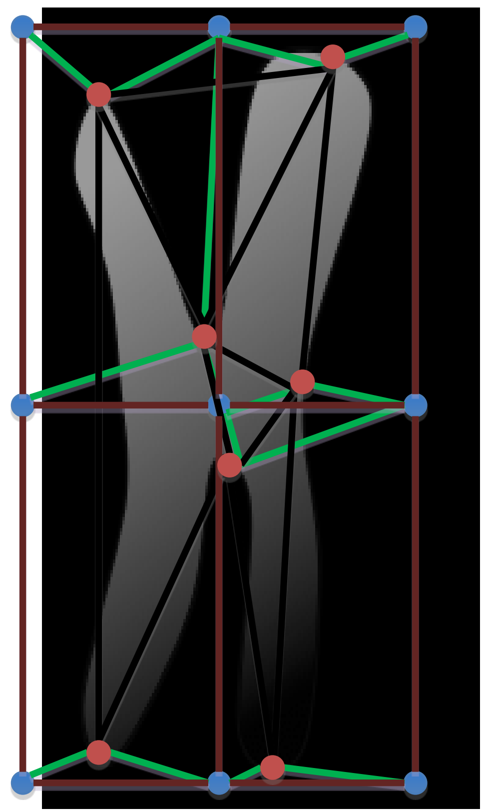

- Landmark correspondence
- Deformation grid node displacement
- Pair-wise interaction - geometric constraint
- Pair-wise interaction - smoothness constraint
- Pair-wise interaction - coupling constraint

Experimental Validation – Qualitative Results

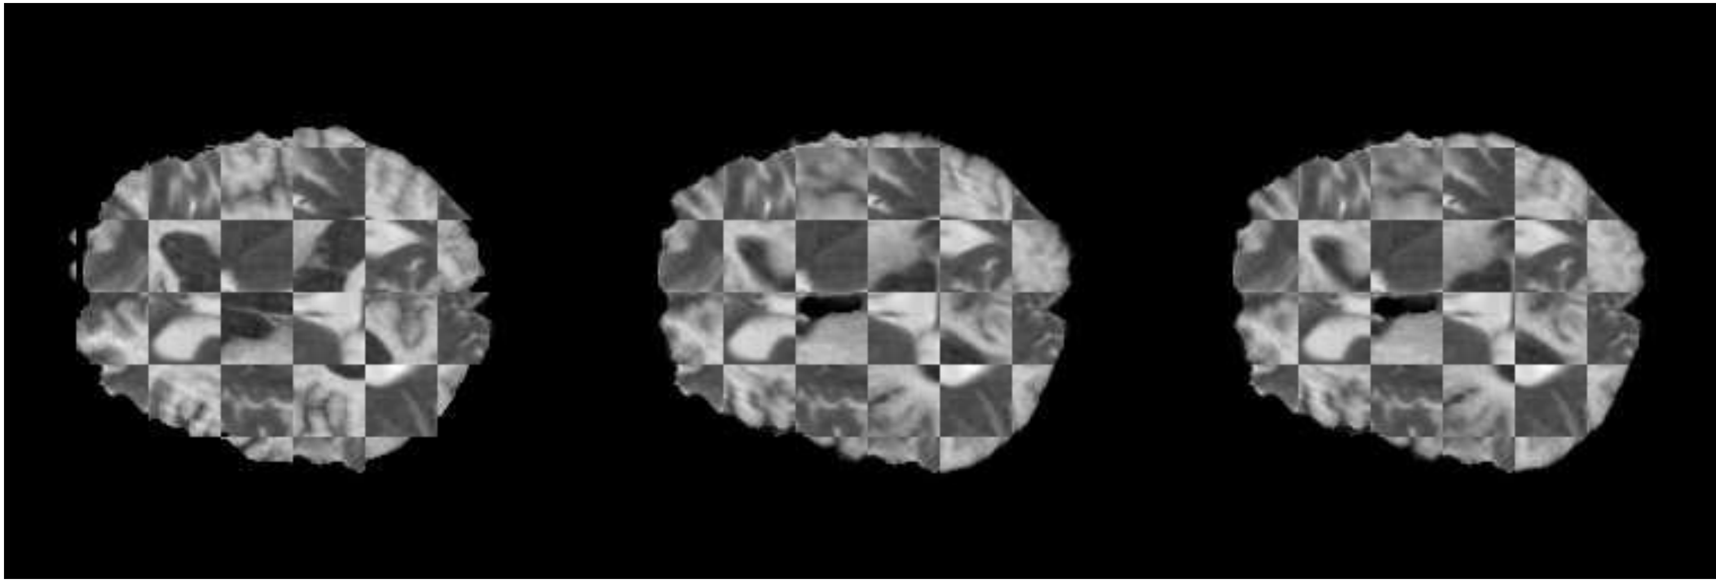

Visual results before (left column) and after registration using the proposed iconic (middle column) and hybrid approach (right column). The results are given in the form of a checkerboard where neighbouring tiles come from different images.

Experimental Validation – Quantitative Results

End point error (in mm)								
#	Iconic (h = 60mm)		Hybrid (h = 60mm)		Iconic (h = 20mm)		Hybrid (h = 20mm)	
	mean	std	mean	std	mean	std	mean	std
1	1,33	0,69	**1,25**	0,59	1,38	1,21	**0,98**	0,61
2	1,32	0,75	**1,18**	0,53	2,46	3,21	**1,06**	0,68
3	1,44	0,97	**1,22**	0,56	2,05	2,40	**1,03**	0,67
4	1,40	0,74	**1,16**	0,50	1,40	1,02	**1,08**	0,69
5	1,23	0,60	**1,15**	0,56	1,38	1,01	**1,03**	0,67
6	1,35	0,74	**1,24**	0,62	1,58	1,39	**1,05**	0,71
7	1,16	0,56	**1,09**	0,50	1,45	1,18	**1,05**	0,67
8	1,29	0,68	**1,23**	0,58	1,93	2,61	**1,11**	0,79
9	1,23	0,62	**1,19**	0,53	1,72	1,89	**1,04**	0,71
10	1,54	1,08	**1,19**	0,58	2,60	3,43	**1,05**	0,73
all	1,33	0,11	**1,19**	0,05	1,79	0,45	**1,05**	0,03

Experimental Validation – Qualitative Results

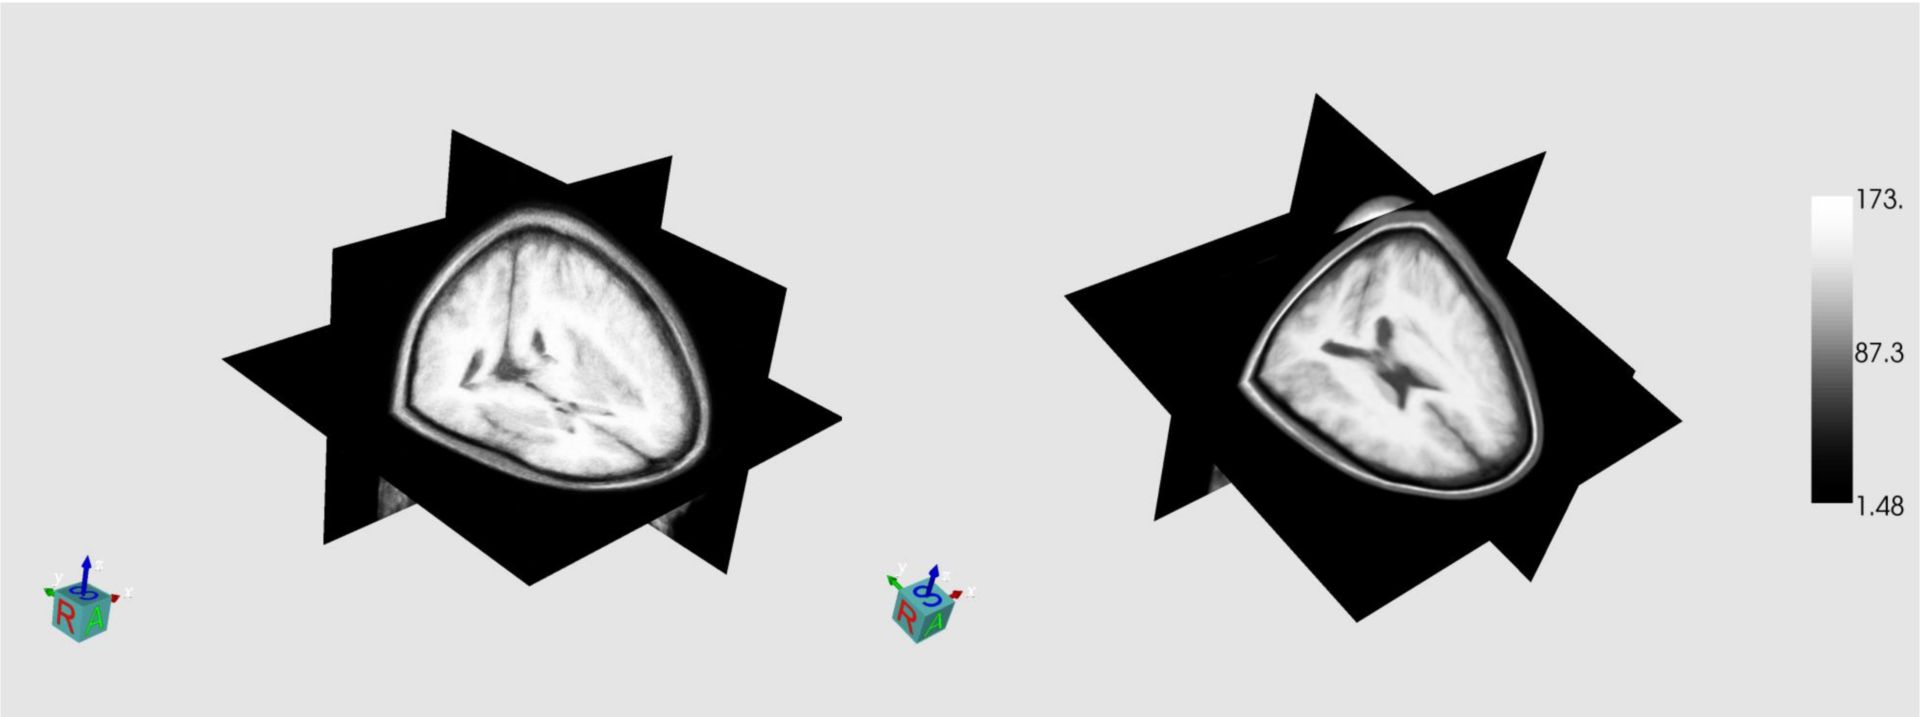

Mean image before and after registration.

Experimental Validation – Data Set [Baker 10]

Dimetrodon

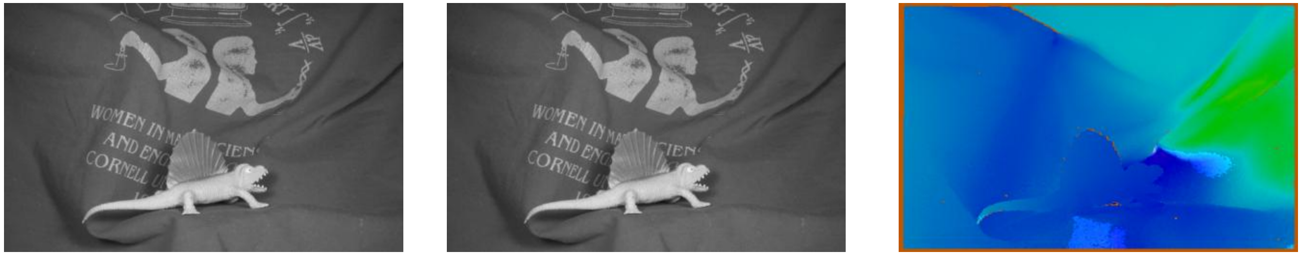

RubberWhale

Urban

Venus

Experimental Validation – Geometric Part

Urban

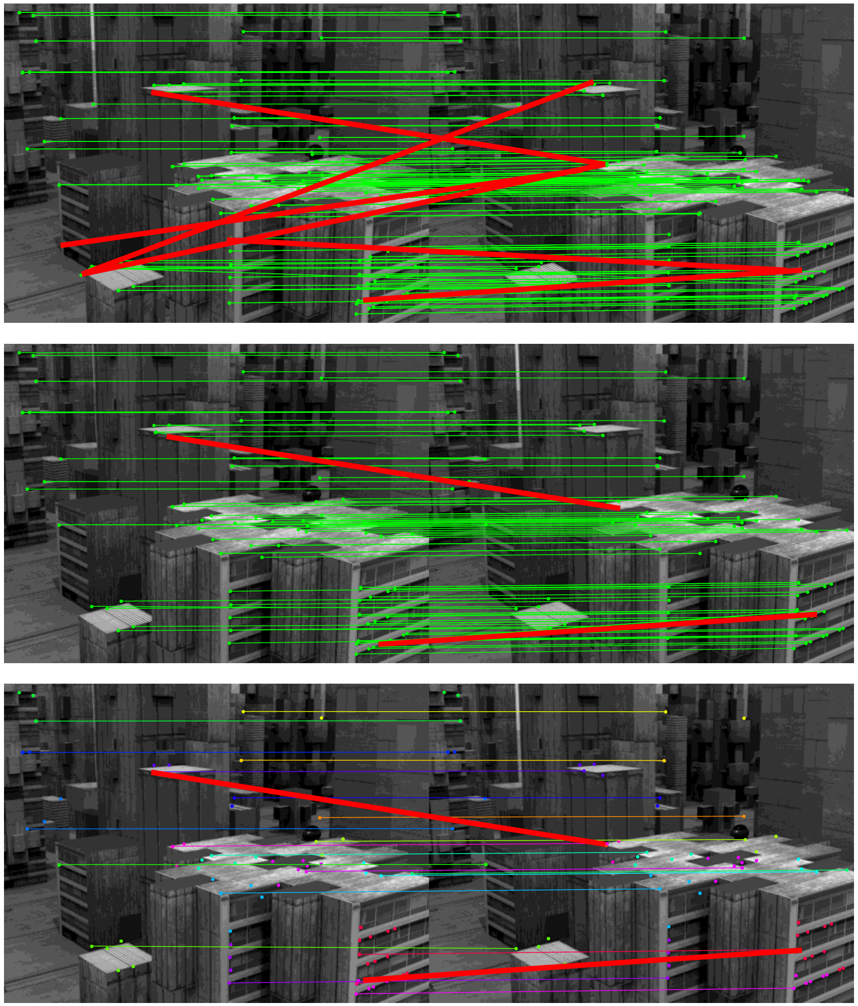

RubberWhale

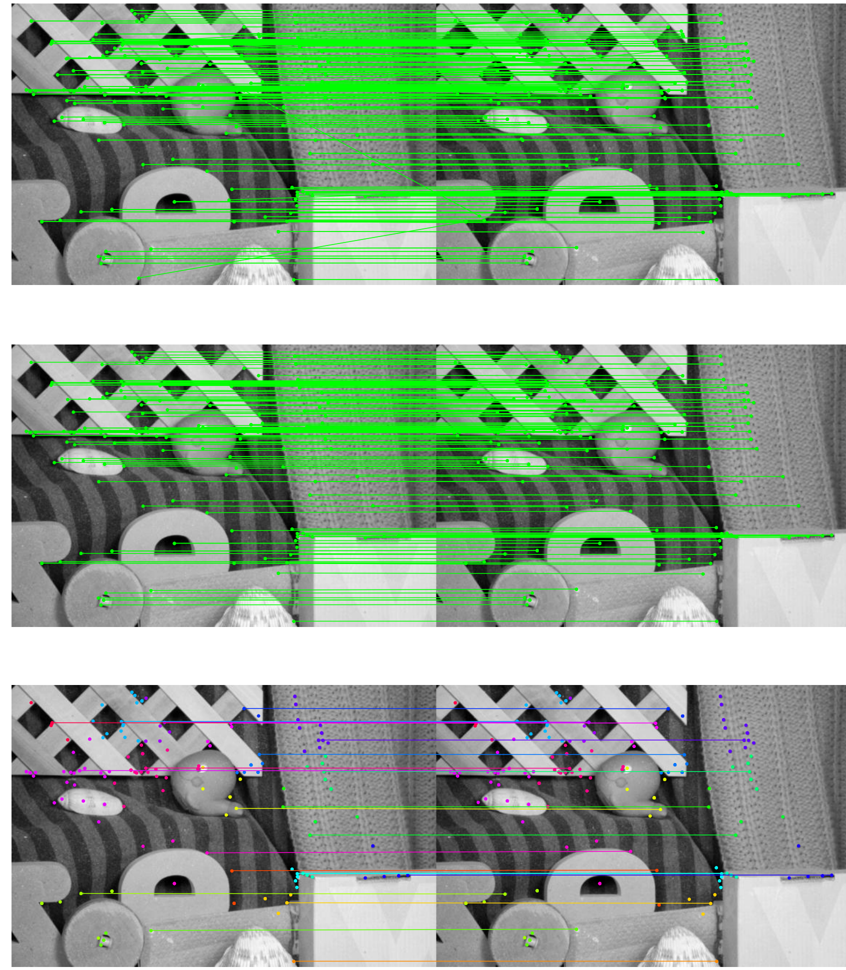

First row: initial correspondences. Second row: correspondences after uniqueness constraints. Third row: correspondences after clustering.

Experimental Validation – Optical Flow

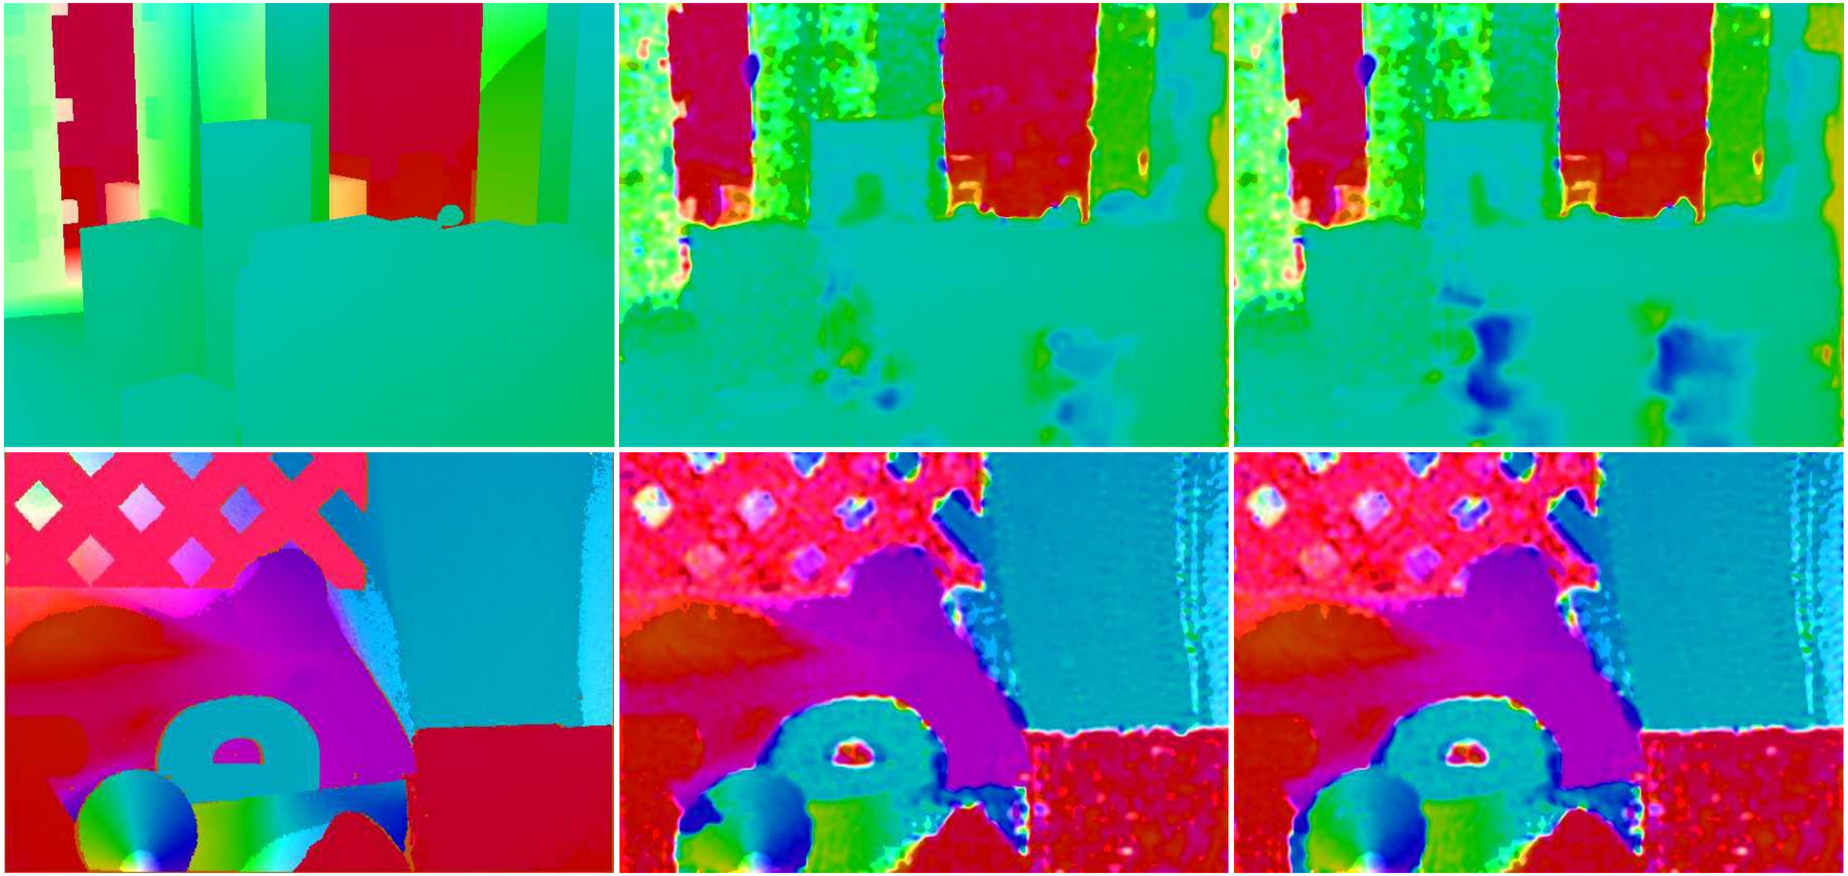

First row: result for Urban sequence. Second row: result for RubberWhale sequence. Left column: ground truth. Middle column: result obtained with the hybrid method. Right column: result obtained with the iconic method.

Experimental Validation – Optical Flow

	Angular error (in degrees)				End point error (in mm)			
Image sequence	Iconic		Coupled		Iconic		Coupled	
	mean	std	mean	std	mean	std	mean	std
Dimetrodon	5,71	4,70	**5,68**	4,71	**0,28**	0,23	0,28	0,24
Grove2	3,92	6,84	**3,90**	6,92	**0,28**	0,44	0,28	0,44
Grove3	**7,88**	15,88	7,97	16,01	**0,82**	1,52	0,83	1,54
Hydrangea	3,73	6,55	**3,63**	6,45	**0,33**	0,49	0,33	0,51
RubberWhale	**6,65**	12,70	7,05	13,91	**0,20**	0,36	0,22	0,45
Urban2	7,95	12,60	**7,46**	12,50	1,51	3,01	**1,27**	2,50
Urban3	9,82	25,89	**8,16**	22,38	1,43	3,11	**1,21**	2,54
Venus	9,00	16,80	**8,97**	9,99	0,58	0,76	**0,56**	0,73

Group-wise registration

Related Work

Template Driven: selection

- Averaging deformations **[Guimond 00]**
- Averaging mean images **[Seghers 04]**
- Intensity reference **[Bhatia 04]**
- Least biased template selection **[Park 05, Hamm 10]**

→ Template introduces bias

Template Driven: construction

- Mean model **[Joshi 04]**
- Geometric median **[Fletcher 09]**
- Geodesic averaging **[Avants 04]**
- Minimum message length criterion **[Cootes 04, Cootes 10]**

→ Template introduces bias

Template-free

- Congealing framework **[Learned-Miller 06, Zollei 05, Balci 07]**
- Summing pairwise differences **[Wachinger 09, Geng 09]**
- Morphological Manifolds **[Baloch 09]**

→ Non-modular w.r.t similarity criterion
→ Not efficient; Only intensity information

Our strategy

- Template-free method
- Graphical model
  - Global statistical similarity criterion
  - Pairwise local comparisons
  - Regularization
- Implicit representation of geometric information

Symmetric Hybrid Registration

Input

- A group of images $\{I_1, \ldots, I_n\}$
- A group of segmentation masks $\{S_1, \ldots, S_n\}$

Output

- A set of transformations $\{T_1, \ldots, T_n\}$, $T_i : \Omega_C \to \Omega_i$
- Unbiased solution $\sum_{i=1}^{n} T_i = 0$

A deformation grid for each image is considered. All grids are isomorphic $(G \sim G_1 \sim \ldots \sim$

Geometric constraint:
$\sum_{i=1}^{n} d_j^i = 0$

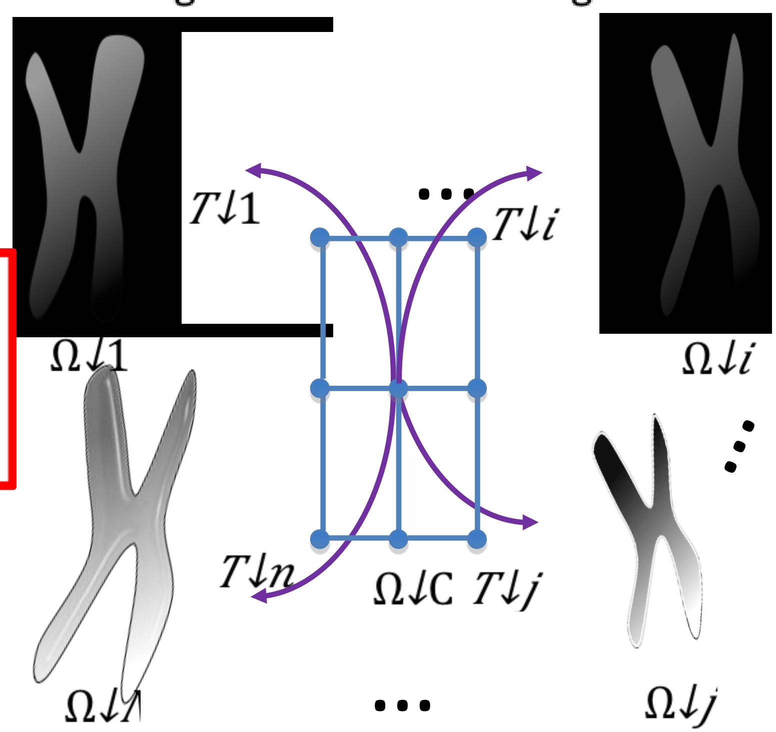

Group-wise Registration - Graph

$$E=E_{global}+E_{local}+E_{smo}$$

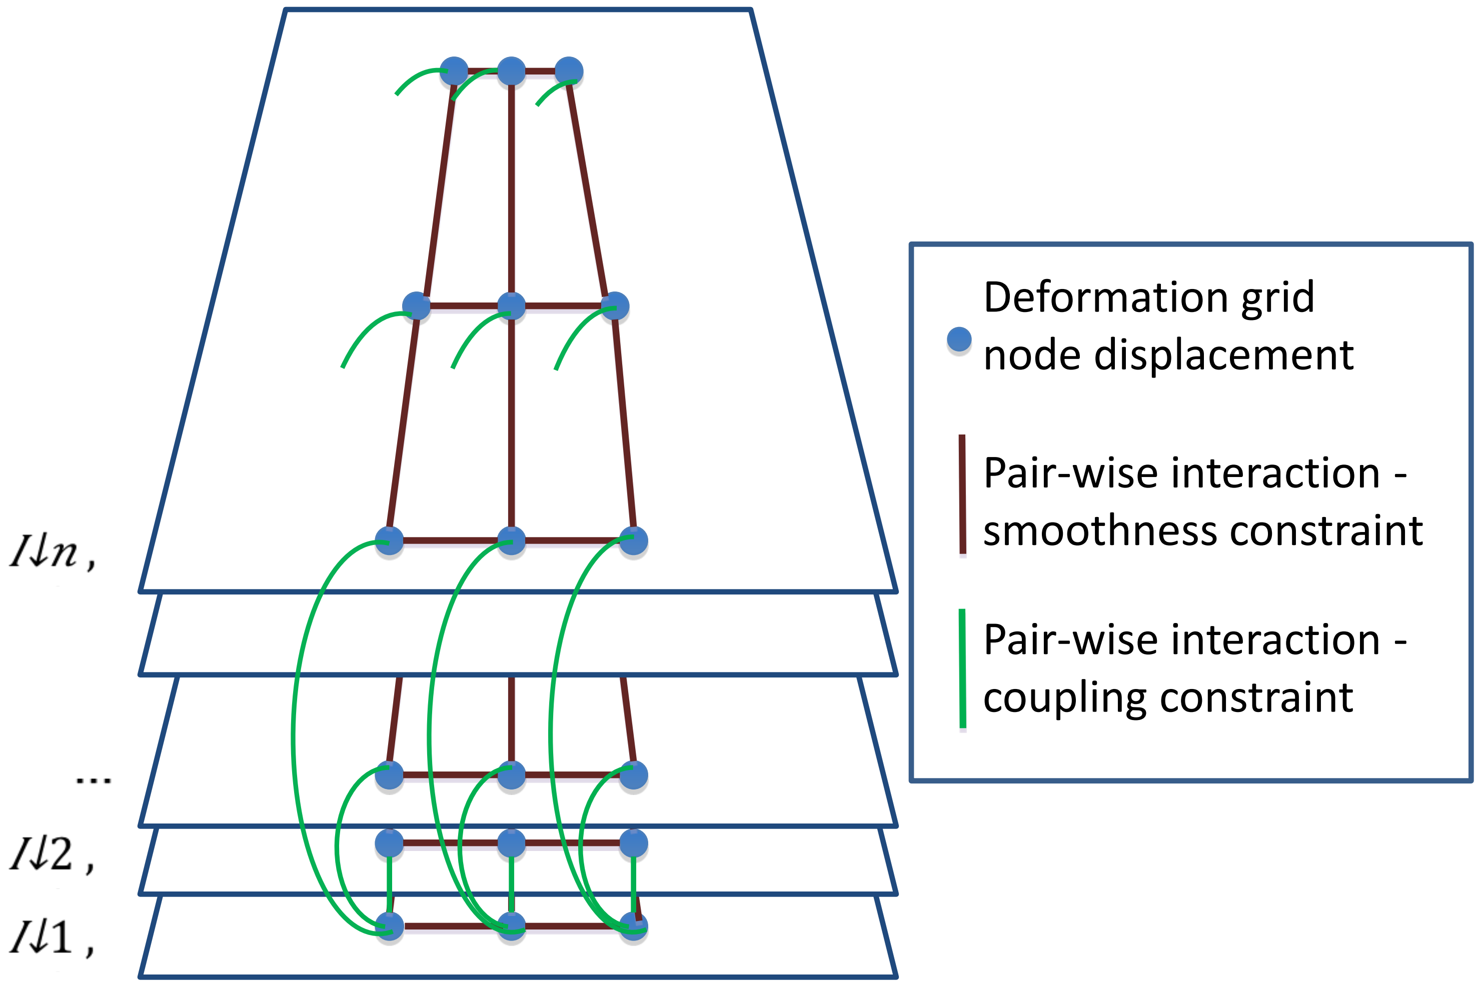

Geometric Information

Geometric information through segmenation mask.

→ **Treat geometric information as iconic one**

→ **No explicit establishment of correspondences**

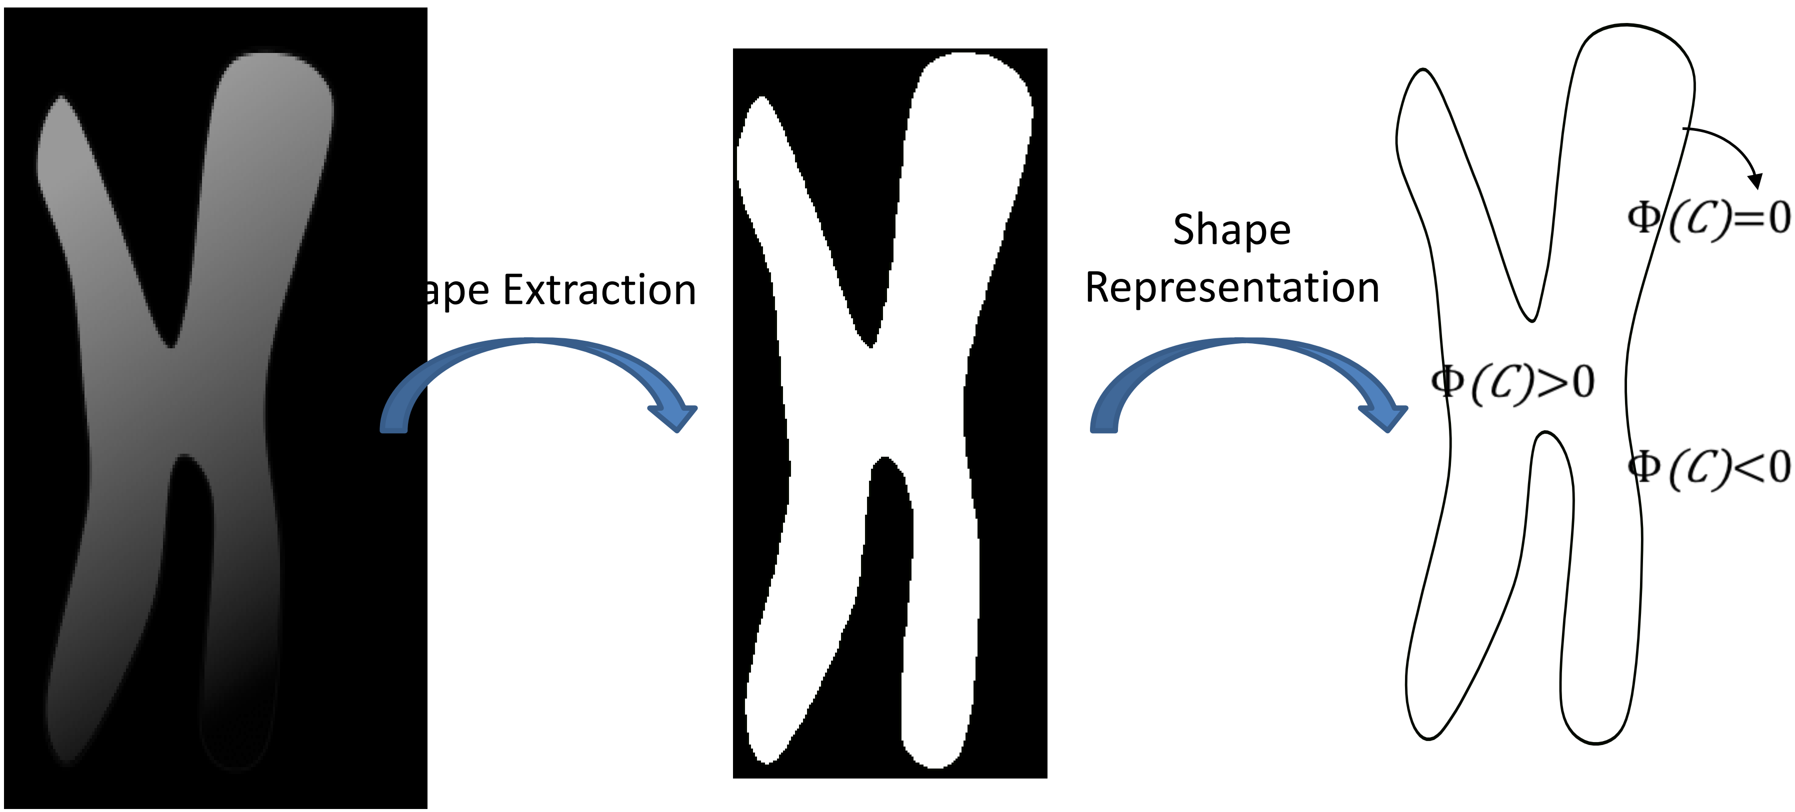

Iconic Part – Local Pairwise Comparison

$$E_{local} = \sum_{i=1}^{n} \sum_{j>i} \rho_{ij} (I$$

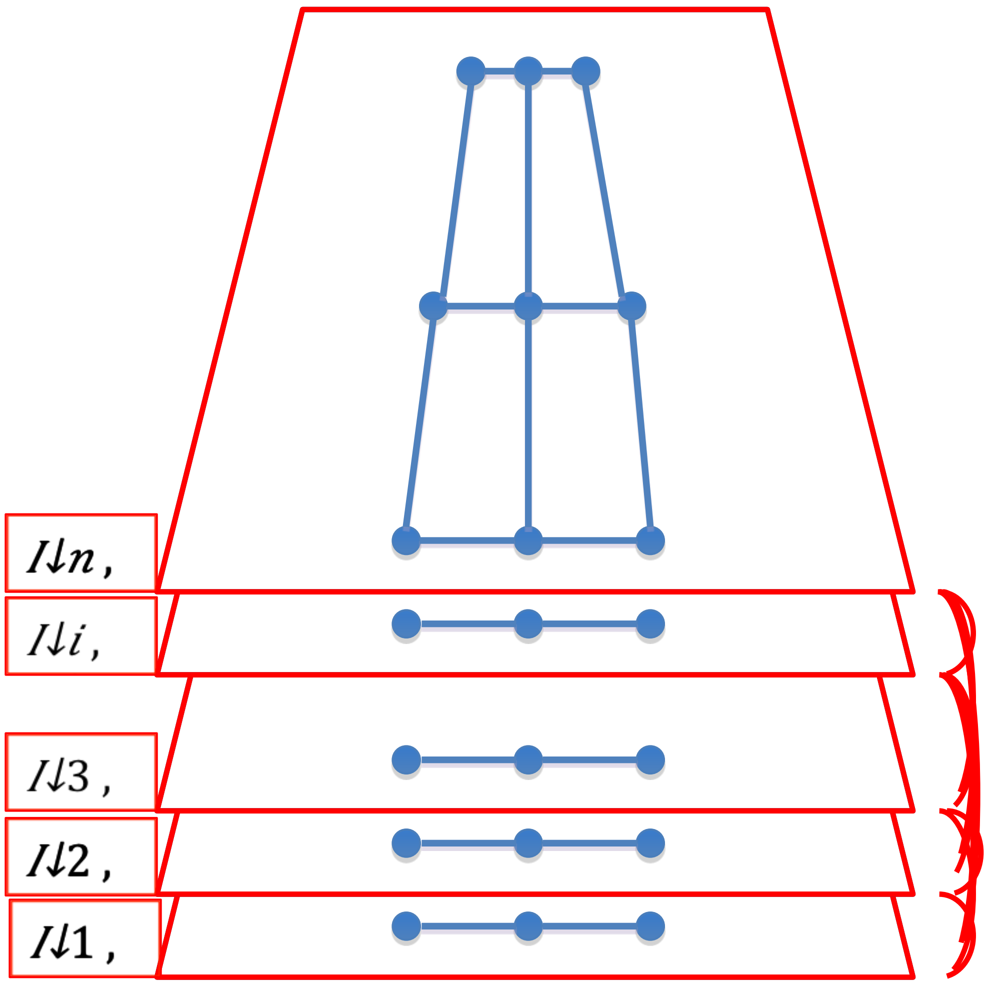

Iconic Part – Local Pairwise Comparison

$$P_{ico,pq}(l_p, l_q) = \int_{\Omega_C} \omega_p(\boldsymbol{x})\, \rho_{ij}(I_i \circ T_{i,l_p}(\boldsymbol{x})),$$

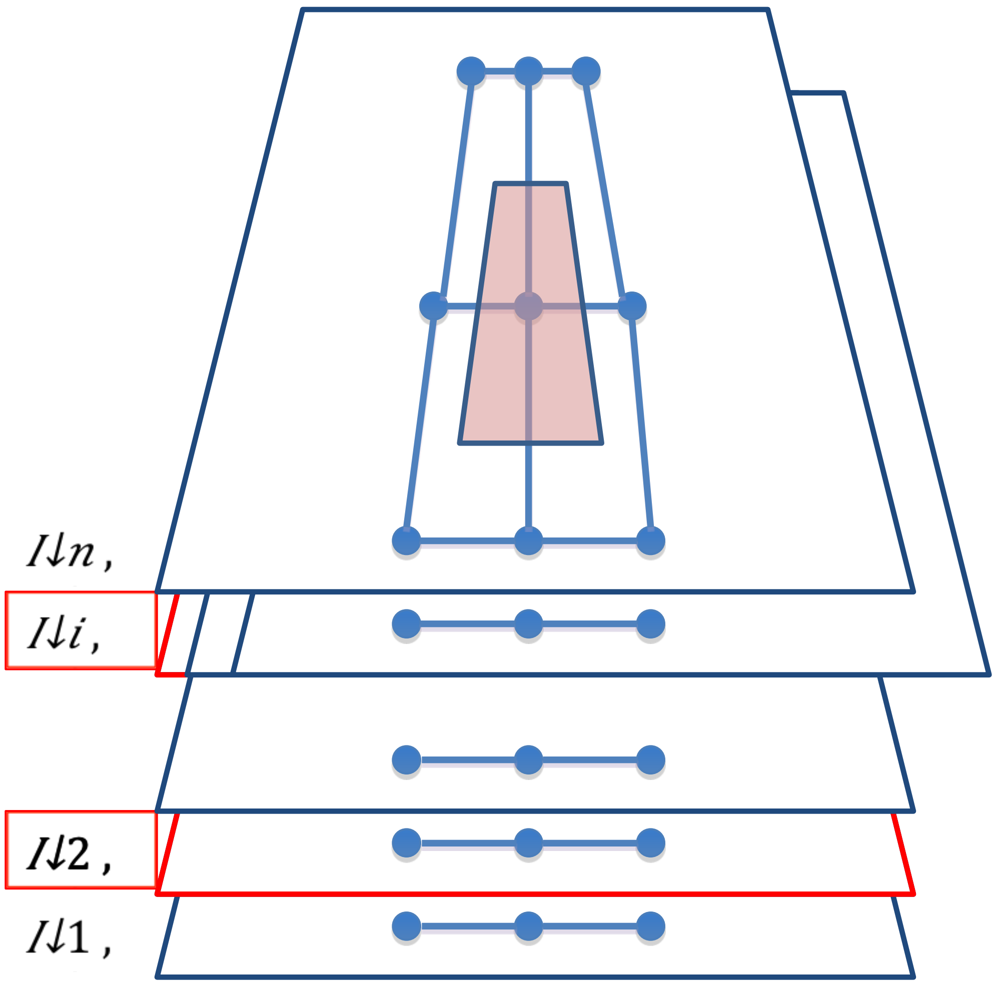

Graph Construction – Inter-layer Edges

$$E_{inter} = \bigcup_{i=1}^{n} ($$

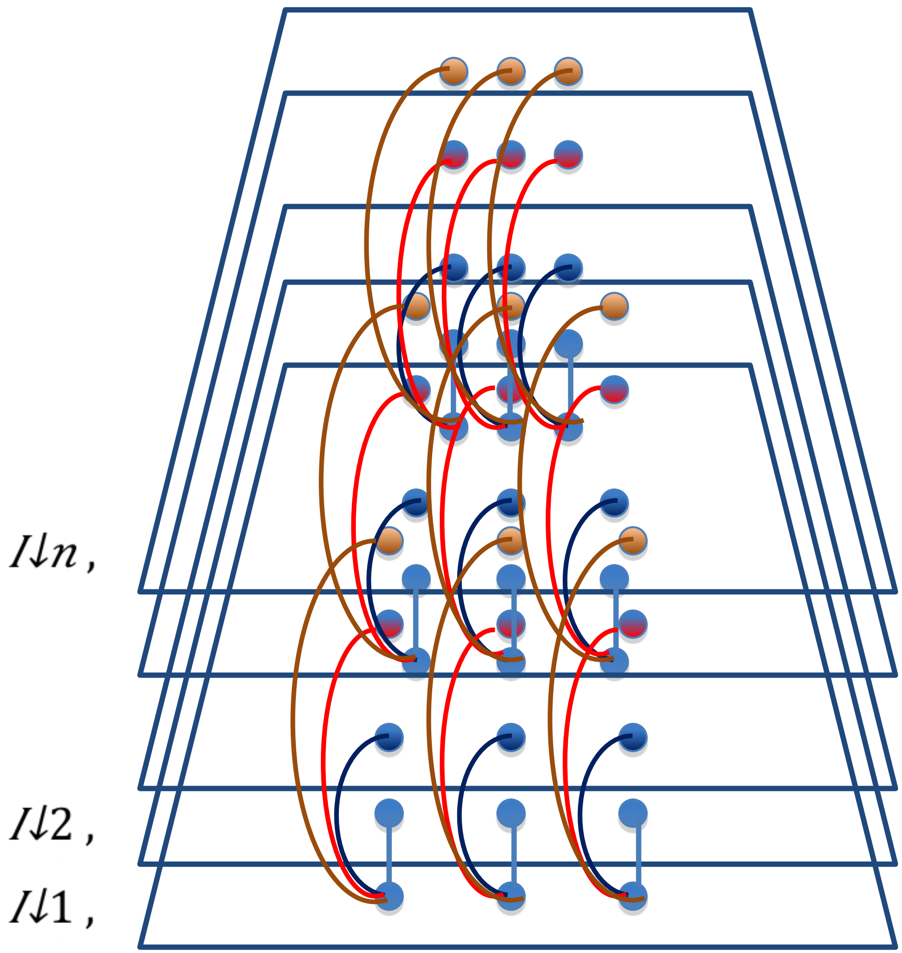

Unified Local Pairwise Comparison

$$P_{hyb,pq}(l_p, l_q) = (1-\alpha)P_{ico,pq}(l_p, l_q) + \alpha P_{geo,pq}(l_p, l_q)$$

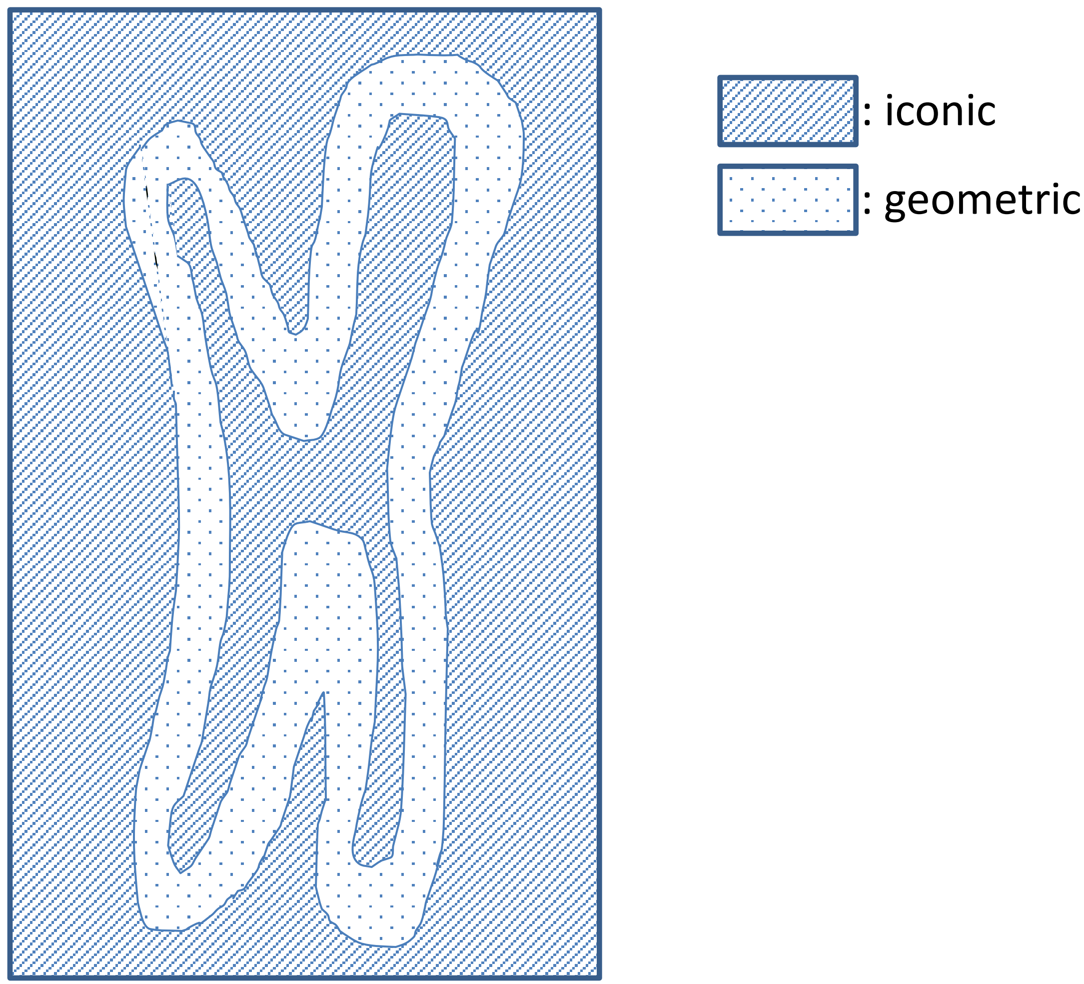

Iconic Part – Global Statistical Criterion

$$M_{global,ico}(I_{p_1},\dots,I_{p_n}) = \int_{\Omega_C} \omega_p(\boldsymbol{x})\,\gamma(\pi(I_1 \circ T_1, I_{p_1}(\boldsymbol{x}),\dots$$

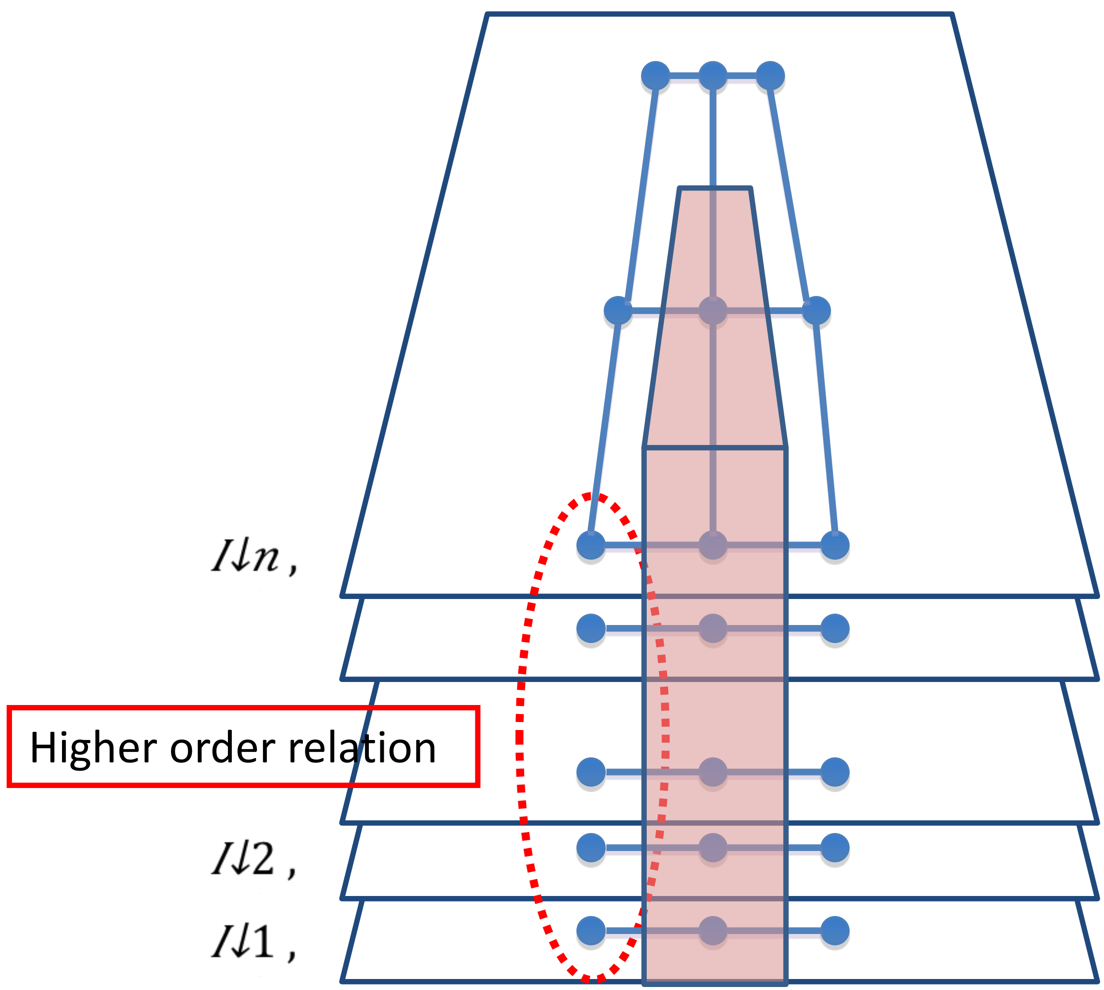

Iconic Part – Global Statistical Criterion

$U\downarrow global,ico,n,p\ (l\downarrow p\ )=\int\Omega\downarrow C\ \uparrow\ \omega\downarrow p\ (\boldsymbol{x})\gamma(\pi(I\downarrow 1\circ T\downarrow 1,l\downarrow p\downarrow 1\ \ \uparrow t-1\ (\boldsymbol{x}),\ldots,I\downarrow n-1\circ T\downarrow n-1,l\downarrow p$

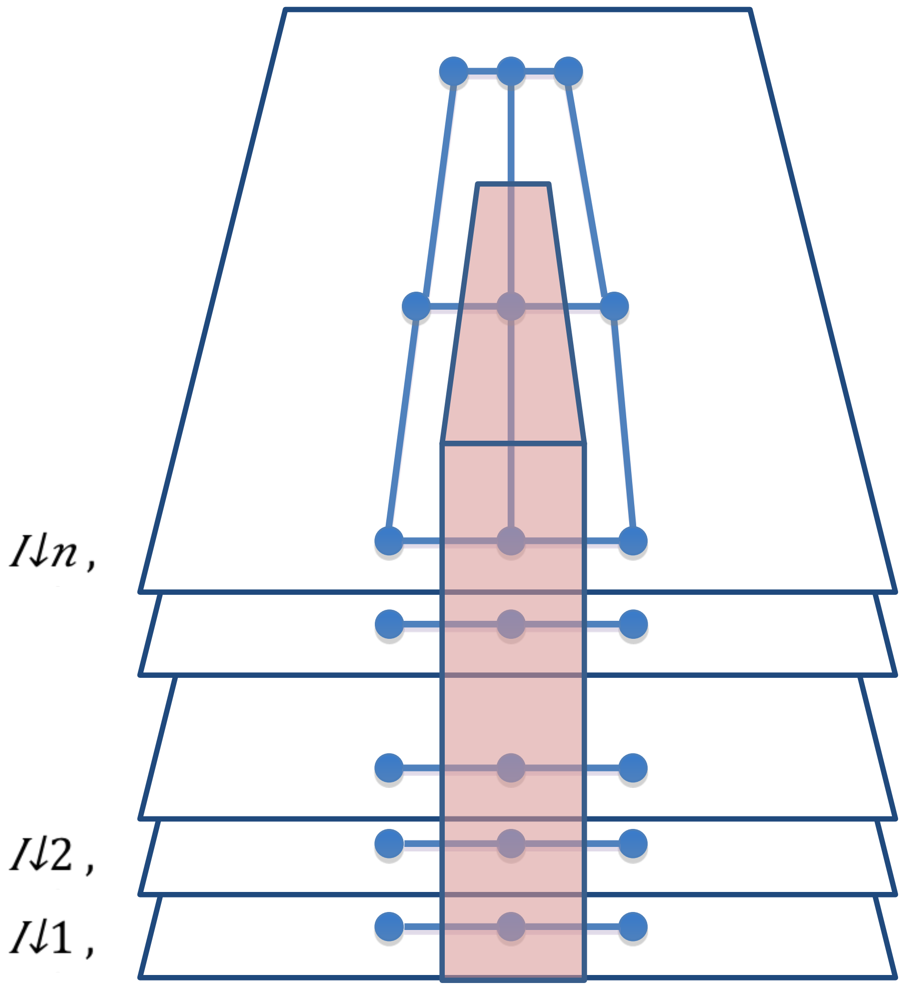

Fluid Regularization – Intra-layer Edges

$E_{intra} = U_{i=1}$ ...

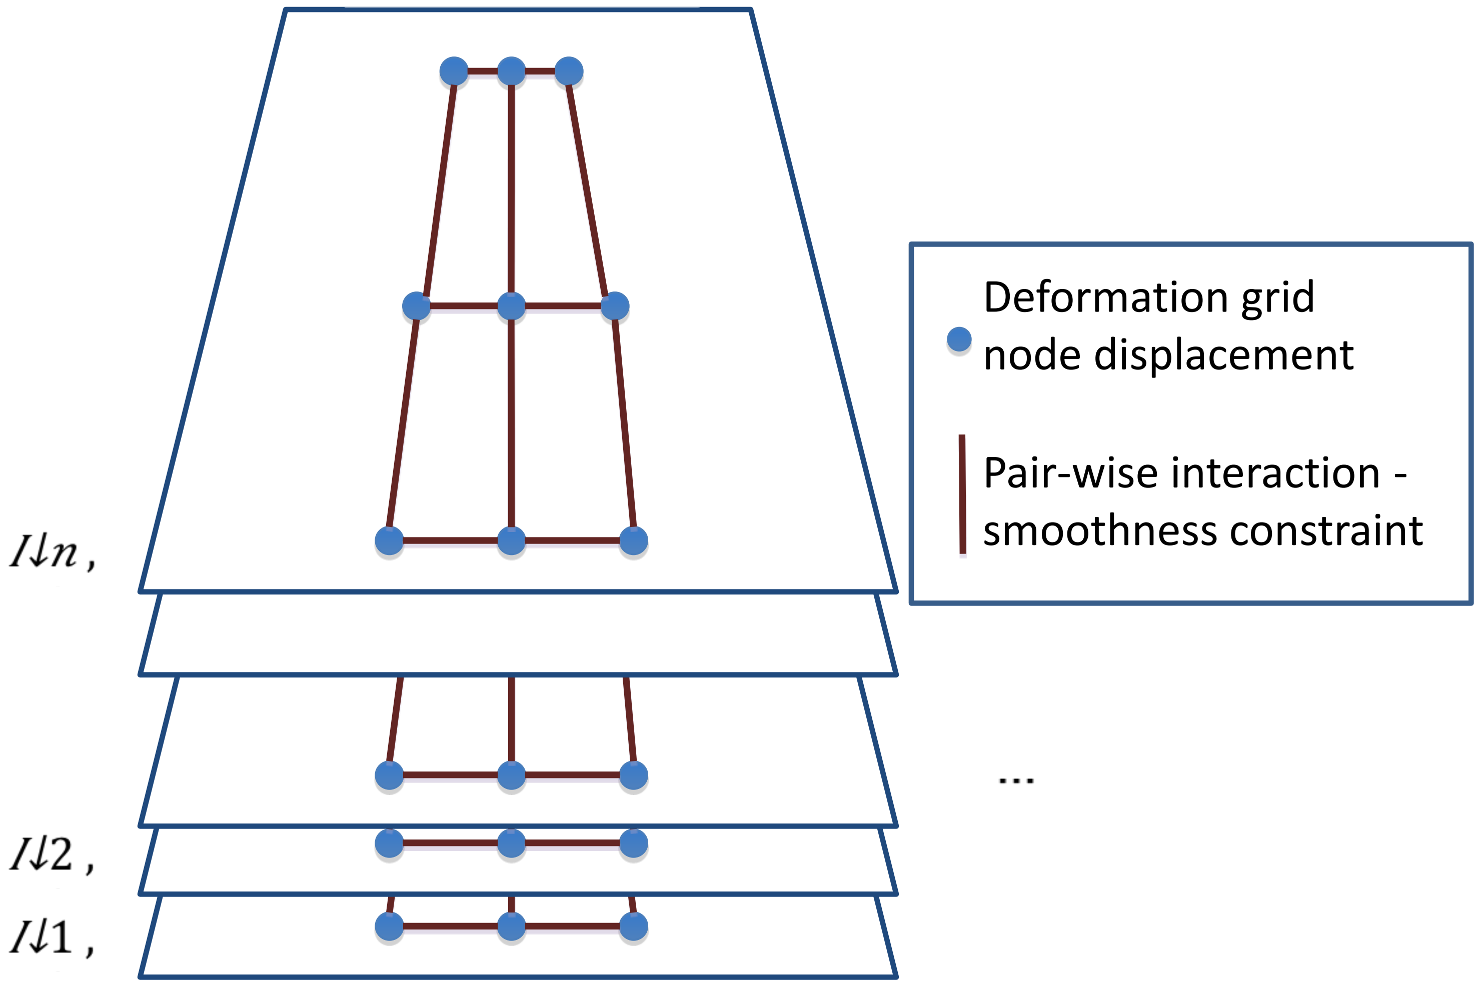

Iconic Part – Regularization Term

Fluid regularization: $P(l_p, l_q) = \|l_p - l_q$

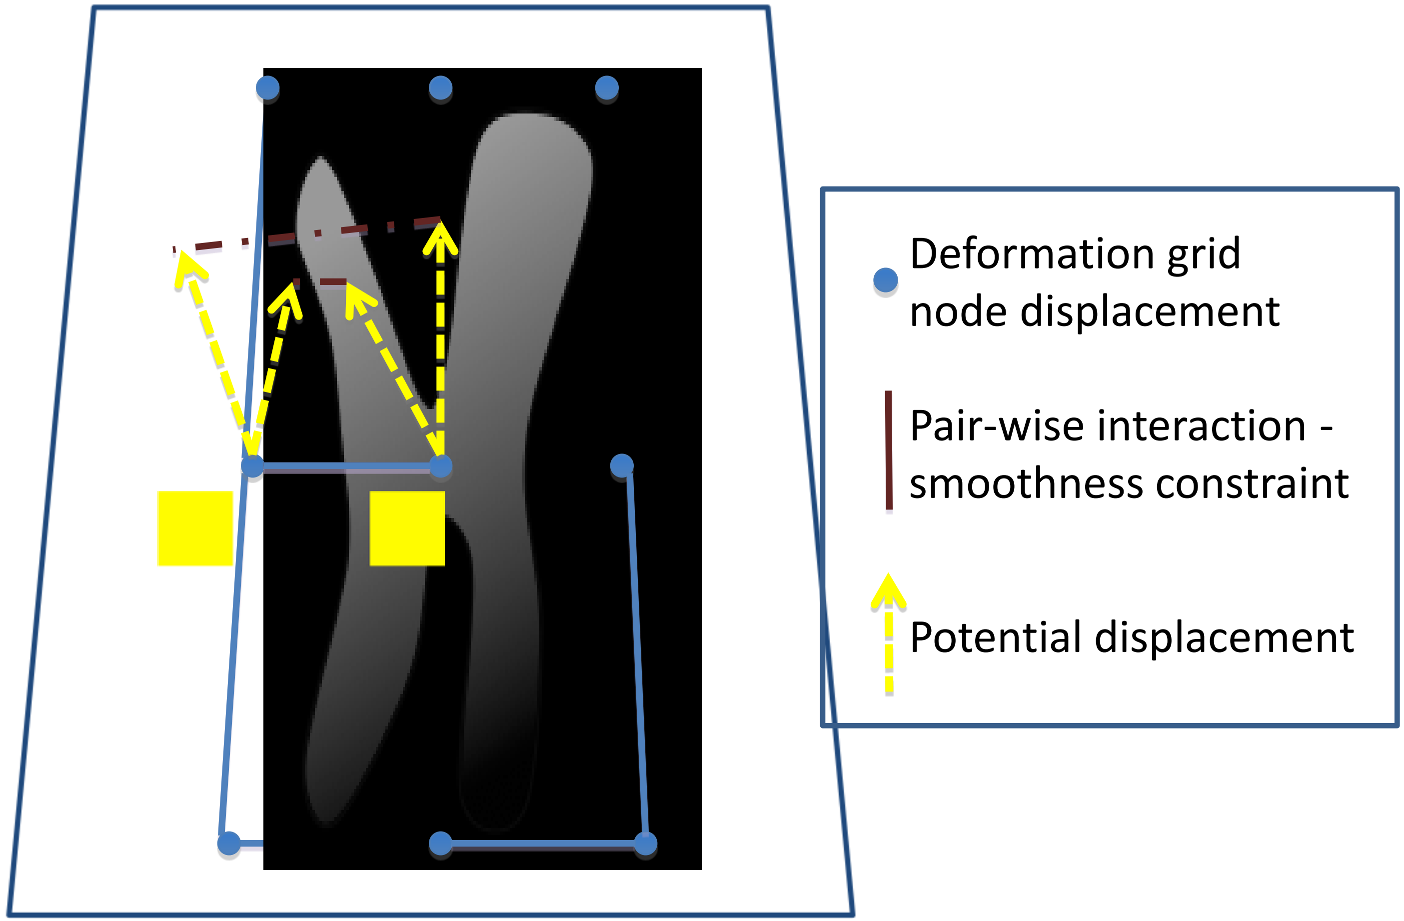

I_i, G_i

Experimental Validation – Data Set [CMA GMH Harvard]

Midway Histogram Equalization **[Delon 04]**

Rerscaled and resampled to equal size and resolution

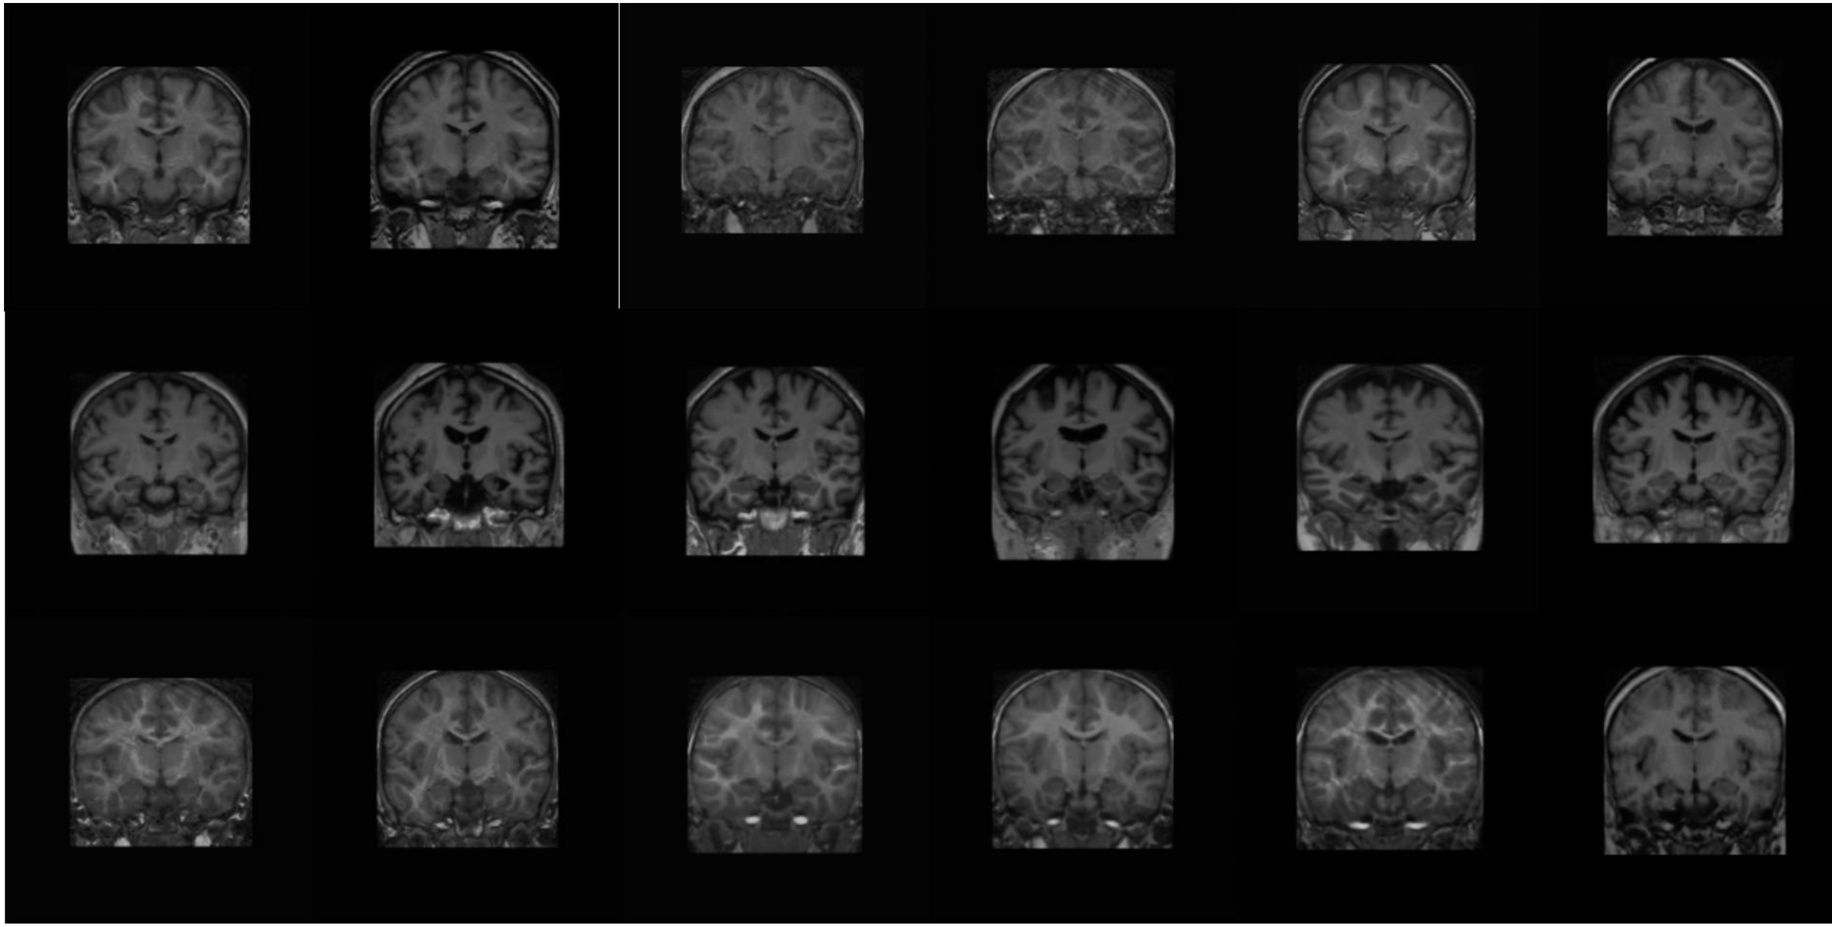

Experimental Validation – Qualitative Results

Initial data

[Balci 07] – low samping rate

[Balci 07] – high samping rate

Proposed - iconic

Proposed - hybrid

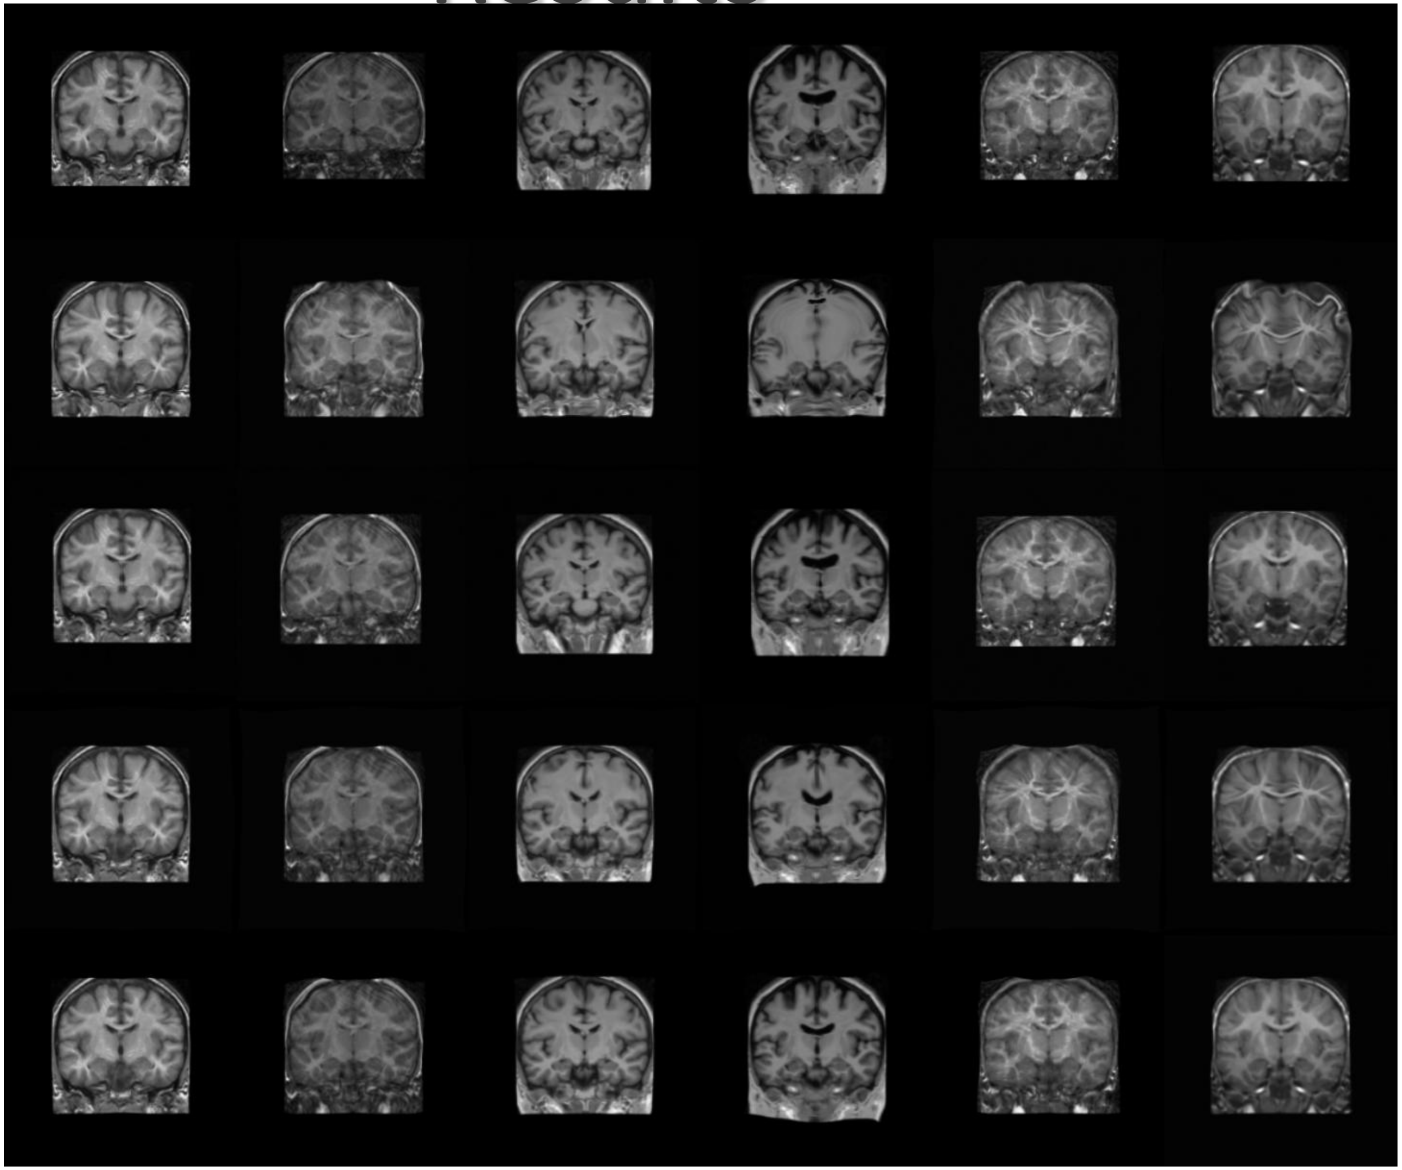

Experimental Validation – Zoom in One Example

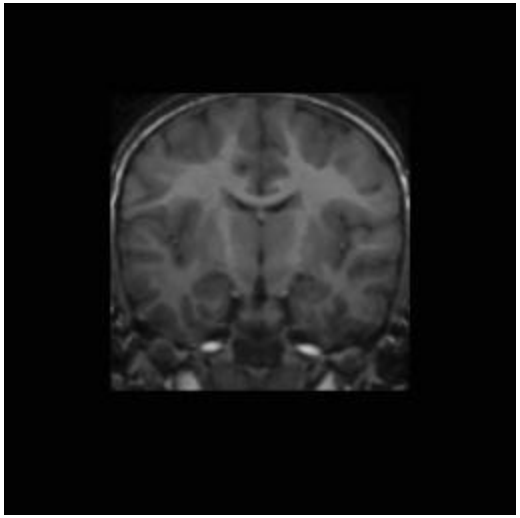

Initial data

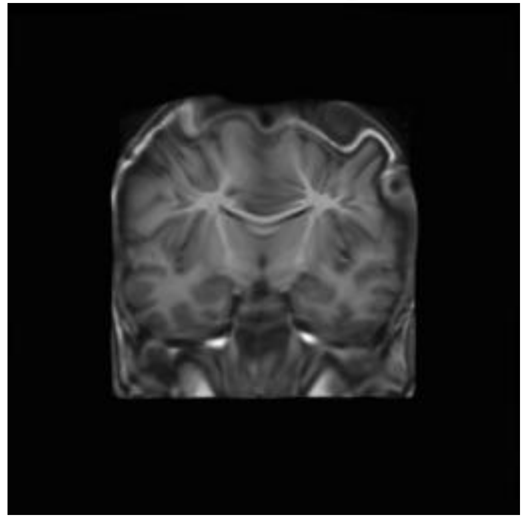

[Balci 07] – LSR

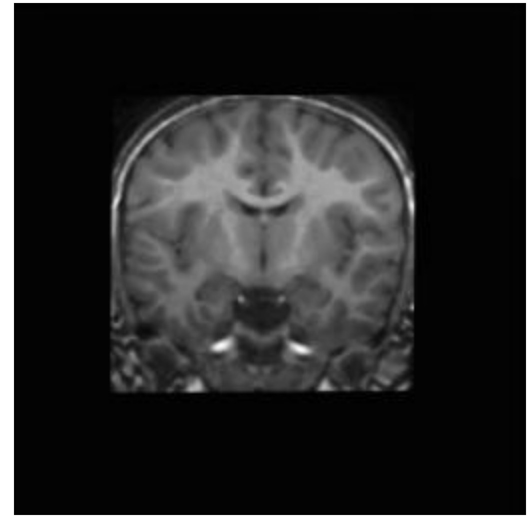

[Balci 07] – HSR

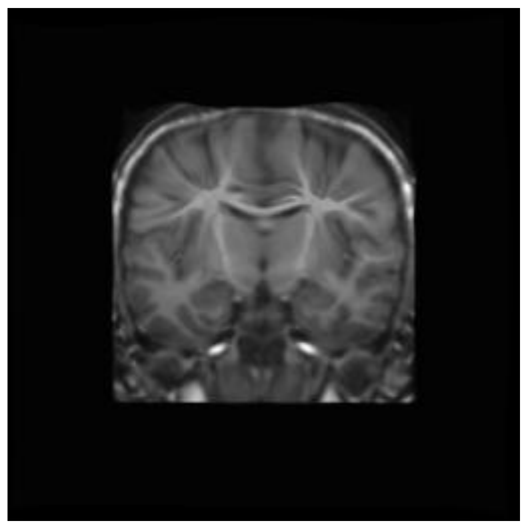

Proposed - iconic

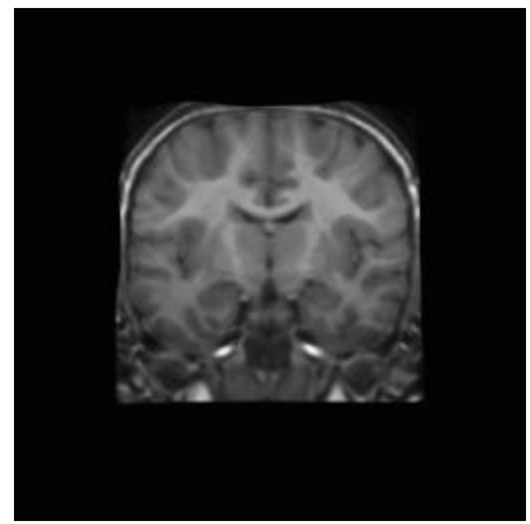

Proposed - hybrid

Experimental Validation – Mean

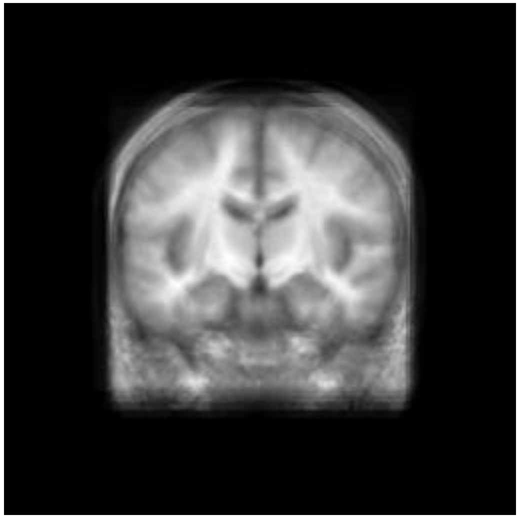

Initial data

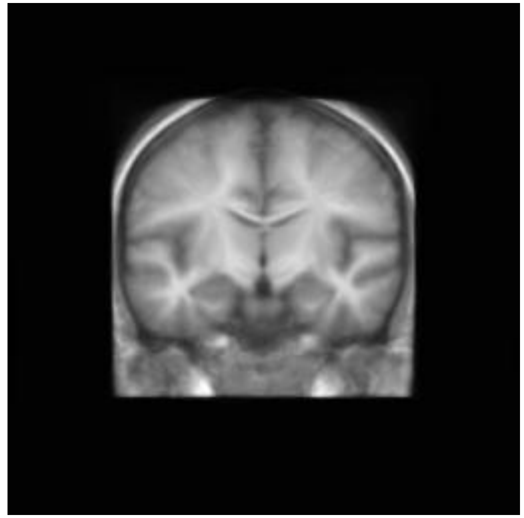

[Balci 07] – LSR

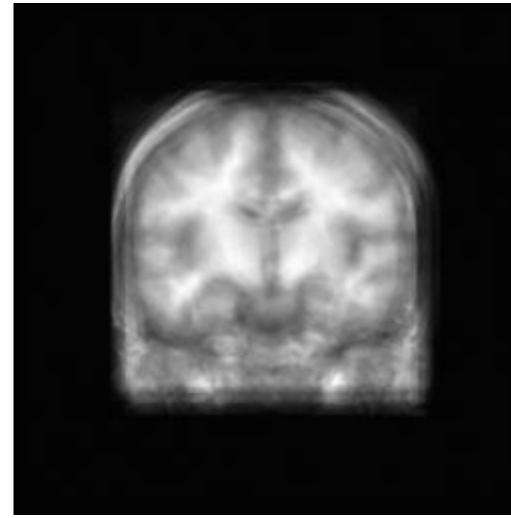

[Balci 07] – HSR

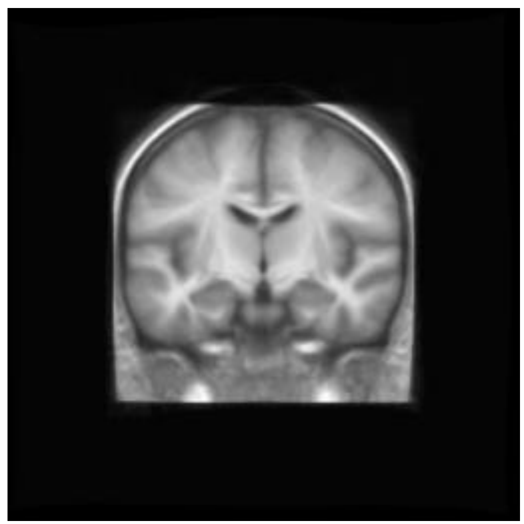

Proposed - iconic

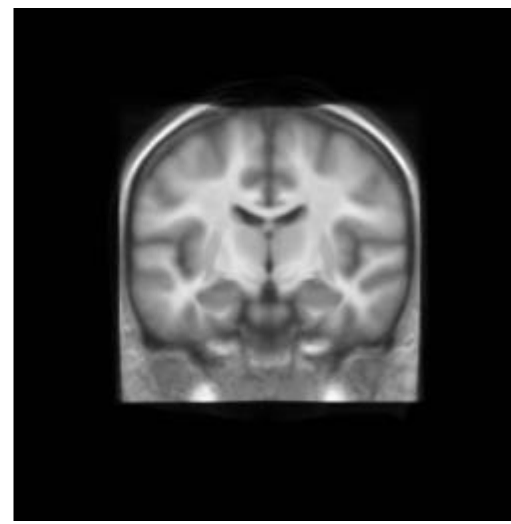

Proposed - hybrid

Experimental Validation – STD

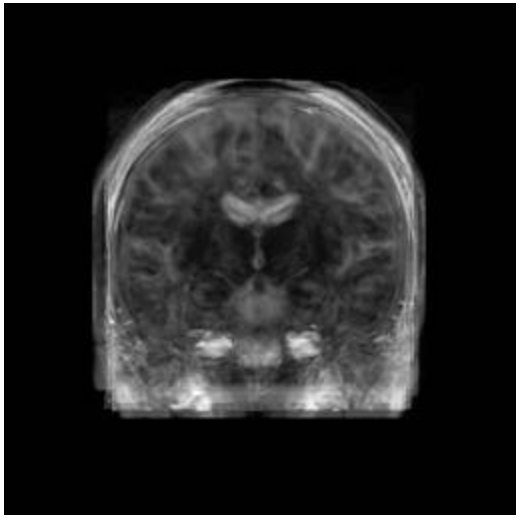

Initial data

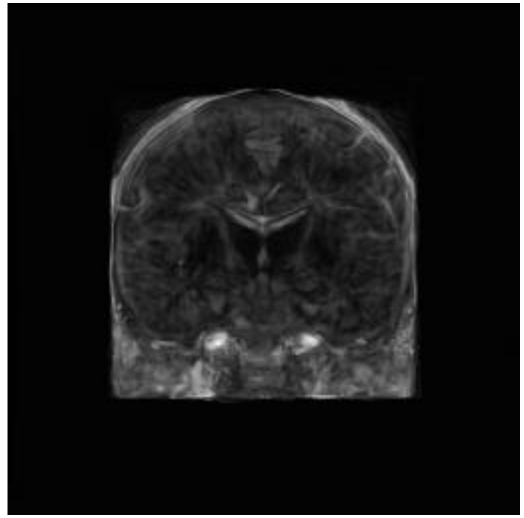

[Balci 07] – LSR

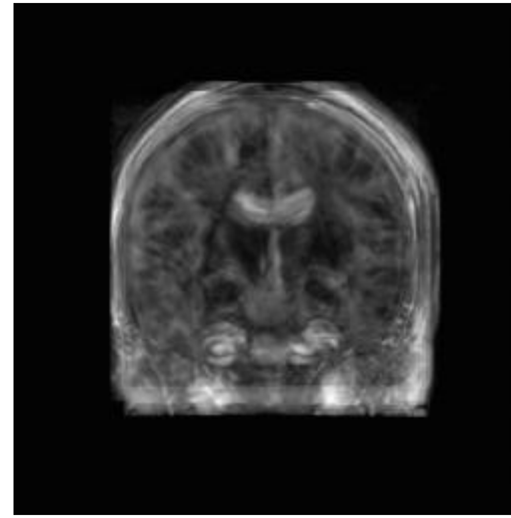

[Balci 07] – HSR

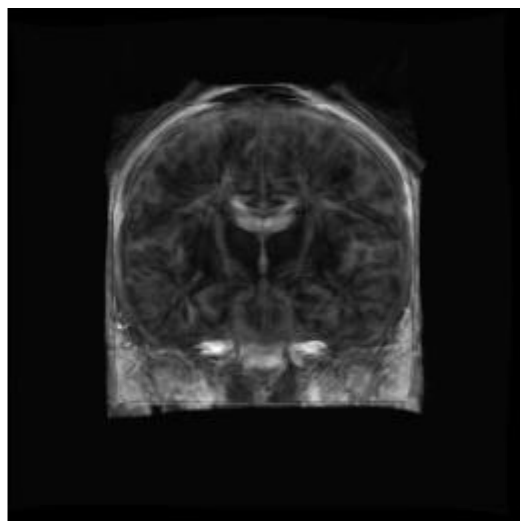

Proposed - iconic

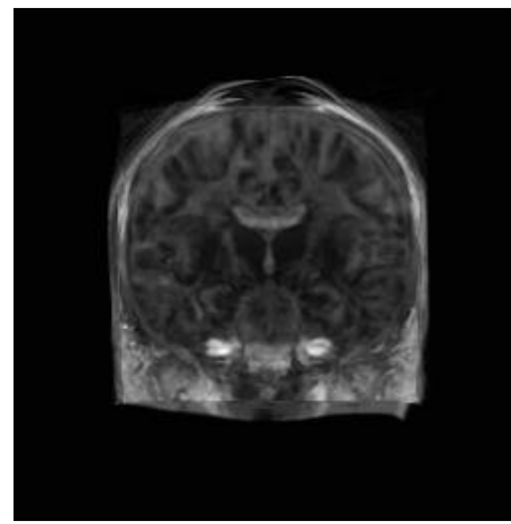

Proposed - hybrid

Conclusion

Contributions

Hybrid Registration

- ✓ Unified objective function
- ✓ One shot optimization
- ✓ Local influence of landmarks
- ✓ No particular assumptions on landmarks or on the nature of the geometric information

Symmetric Coupled Registration

- ✓ Coupled framework
- ✓ Symmetric geometric framework
- ✓ Symmetric iconic registration
- ✓ Robust
- ✓ Interaction between two problems

Group-wise Registration

- ✓ Consider both iconic and geometric information
- ✓ Combine both global and local criteria

Common Properties

- ✓ Diffeomorphic
- ✓ Efficient
- ✓ Versatile
- ✓ Modular w.r.t similarity criterion
- ✓ Modular w.r.t interpolation method

Thank You!

Questions?