

Nonrigid image registration that preserves topology

Jeffrey A. Fessler

EECS Department
BME Department, Dept. of Radiology
The University of Michigan

KLA-Tencor
2010-03-01

Acknowledgments: Se Young Chun, Matt Jacobson
James Balter, Marc Kessler




Outline

- Introduction
- Image registration
 - Enforcing / encouraging local invertibility (diffeomorphism)
 - IEEE J. Selected Topics in Signal Processing, Feb. 2009
- Motion-compensated image reconstruction
 - Conventional
 - Model based
 - Temporal regularization


Image reconstruction (and registration) toolbox:

<http://www.eecs.umich.edu/~fessler>

Image Reconstruction Overview



X-ray CT Sinogram



X-ray CT Image

Image Reconstruction Formulations

- Static $f(\vec{r})$
 - Dynamic $f(\vec{r}, t)$
 - contrast changes
 - object motion
- (synergy with image registration)

Image Reconstruction Research Topics

- Models for imaging physics
- Models for measurement statistics
- Models for object
- Algorithm formulation
- Algorithm acceleration
- Algorithm analysis / evaluation
- ...

Why



FBP (thin slice)

(120 kVp, 100-150 mA with modulation, 0.4 s rotation, pitch 1.375, slice thickness 0.625 mm)



MBIR

Part 1
Nonrigid image registration
ensuring local invertibility


Image Registration

Many applications, *e.g.*, forensics, remote sensing, medicine ...

- rigid transformations
- nonrigid transformations (warps)

Example: Respiratory motion (application: radiotherapy planning)

Target



Inhale

Source



Exhale

Image registration: Overview

Given two images (or image volumes): $f(\vec{r})$ and $g(\vec{r})$, $\vec{r} = (x, y, z)$, find a spatial transformation $\vec{T}(\vec{r})$, where $\vec{T} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$, such that $f(\vec{r})$ “is similar to” the warped image $g(\vec{T}(\vec{r}))$

Usual steps:

- parameterize by α the spatial transformation: $\vec{T}(\vec{r}; \alpha)$
- choose a similarity measure $\Psi\left(f(\cdot), g(\vec{T}(\cdot))\right)$
- find optimal deformation parameters α numerically:

$$\hat{\alpha} = \arg \max_{\alpha} \Psi\left(f(\cdot), g(\vec{T}(\cdot; \alpha))\right)$$

Challenge: want estimated transformation $\vec{T}(\vec{r}; \hat{\alpha})$ to be plausible. Typically we want it to be *diffeomorphic*, or *topology preserving*, or *invertible*, or at least *locally invertible*.

Image registration: Similarity measures

- sum of squared differences

$$\Psi(f_1, f_2) = - \sum_i |f_1(\mathbf{r}_i) - f_2(\mathbf{r}_i)|^2$$

- correlation

$$\Psi(f_1, f_2) = \sum_i (f_1(\mathbf{r}_i) - \bar{f}_1) (f_2(\mathbf{r}_i) - \bar{f}_2)$$

- mutual information

$$\Psi(f_1, f_2) = \mathbf{MI}(f_1, f_2) = H(f_1) + H(f_2) - H(f_1, f_2)$$

- ...

Image registration: B-spline deformations

Nonrigid spatial transformation:

$$\vec{T}(\vec{r}; \boldsymbol{\alpha}) = \vec{r} + \underbrace{(d^x(\vec{r}; \boldsymbol{\alpha}^x), d^y(\vec{r}; \boldsymbol{\alpha}^y), d^z(\vec{r}; \boldsymbol{\alpha}^z))}_{\text{deformation}},$$

where $\boldsymbol{\alpha} = (\boldsymbol{\alpha}^x, \boldsymbol{\alpha}^y, \boldsymbol{\alpha}^z)$ denotes unknown deformation coefficients.


Tensor-product B-spline deformation model:

$$\begin{aligned}d^x(\vec{r}; \boldsymbol{\alpha}^x) &= \sum_{i,j,k} \alpha_{ijk}^x \beta(x/m_x - i) \beta(y/m_y - j) \beta(z/m_z - k) \\d^y(\vec{r}; \boldsymbol{\alpha}^y) &= \sum_{i,j,k} \alpha_{ijk}^y \beta(x/m_x - i) \beta(y/m_y - j) \beta(z/m_z - k) \\d^z(\vec{r}; \boldsymbol{\alpha}^z) &= \sum_{i,j,k} \alpha_{ijk}^z \beta(x/m_x - i) \beta(y/m_y - j) \beta(z/m_z - k)\end{aligned}$$

m_x, m_y, m_z denote the knot spacing in each dimension.

These spacings determine the spatial scale of the deformation.

Cubic B-spline Kernel



B-spline deformations: Benefits

- differentiable (smooth)
- local support
- recursive filters for computations
- piecewise polynomial
- hierarchical


Nonrigid image registration similarity measures usually have many local maximizers.

To help find a “good” local maximum, one usually uses coarse-to-fine search. This is easy with B-spline deformations.


(Thevenaz & Unser, IEEE T-IP, 2000)

B-spline deformations illustrated

knot locations, $m_x=8$ $m_y=16$




Knot locations




Local support

B-spline deformations illustrated



Invertible



Not invertible

Invertibility

When is $\vec{r} \mapsto \vec{T}(\vec{r}) = \vec{r} + \vec{d}(\vec{r})$ a locally invertible transformation?

By the inverse function theorem, it is sufficient for \vec{T} to

- be continuously differentiable, and
- have positive Jacobian determinant: $\det\left\{\nabla \vec{T}(\vec{r})\right\} > 0$ for all \vec{r} .

Jacobian of transformation/deformation

$$\nabla \vec{T}(\vec{r}) = \nabla \left(\vec{r} + \vec{d}(\vec{r}) \right) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} + \begin{bmatrix} \frac{\partial}{\partial x} d^x & \frac{\partial}{\partial y} d^x & \frac{\partial}{\partial z} d^x \\ \frac{\partial}{\partial x} d^y & \frac{\partial}{\partial y} d^y & \frac{\partial}{\partial z} d^y \\ \frac{\partial}{\partial x} d^z & \frac{\partial}{\partial y} d^z & \frac{\partial}{\partial z} d^z \end{bmatrix}$$

Ensuring local invertibility in image registration

We need to estimate B-spline deformation coefficients α subject to some local invertibility constraint $\alpha \in C$:

$$\hat{\alpha} = \arg \max_{\alpha \in C} \Psi(\alpha).$$

- Ideal sufficient condition ensuring local invertibility for parametric deformation model:

$$\alpha \in C_0 = \left\{ \alpha : \det \left\{ \nabla \vec{T}(\vec{r}; \alpha) \right\} > 0, \forall \vec{r} \in \mathbb{R}^3 \right\}.$$

This condition is very difficult to implement.

- Conventional relaxed local invertibility condition: $C_0 \subset C_1$

$$\alpha \in C_1 = \left\{ \alpha : \det \left\{ \nabla \vec{T}(\vec{r}; \alpha) \right\} > 0, \vec{r} \in \textit{grid points} \right\}.$$

This condition does not ensure local invertibility everywhere. It is also computationally demanding.

We seek simpler *sufficient* conditions for local invertibility: $C \subset C_0$.

Unconstrained vs “constrained” optimization

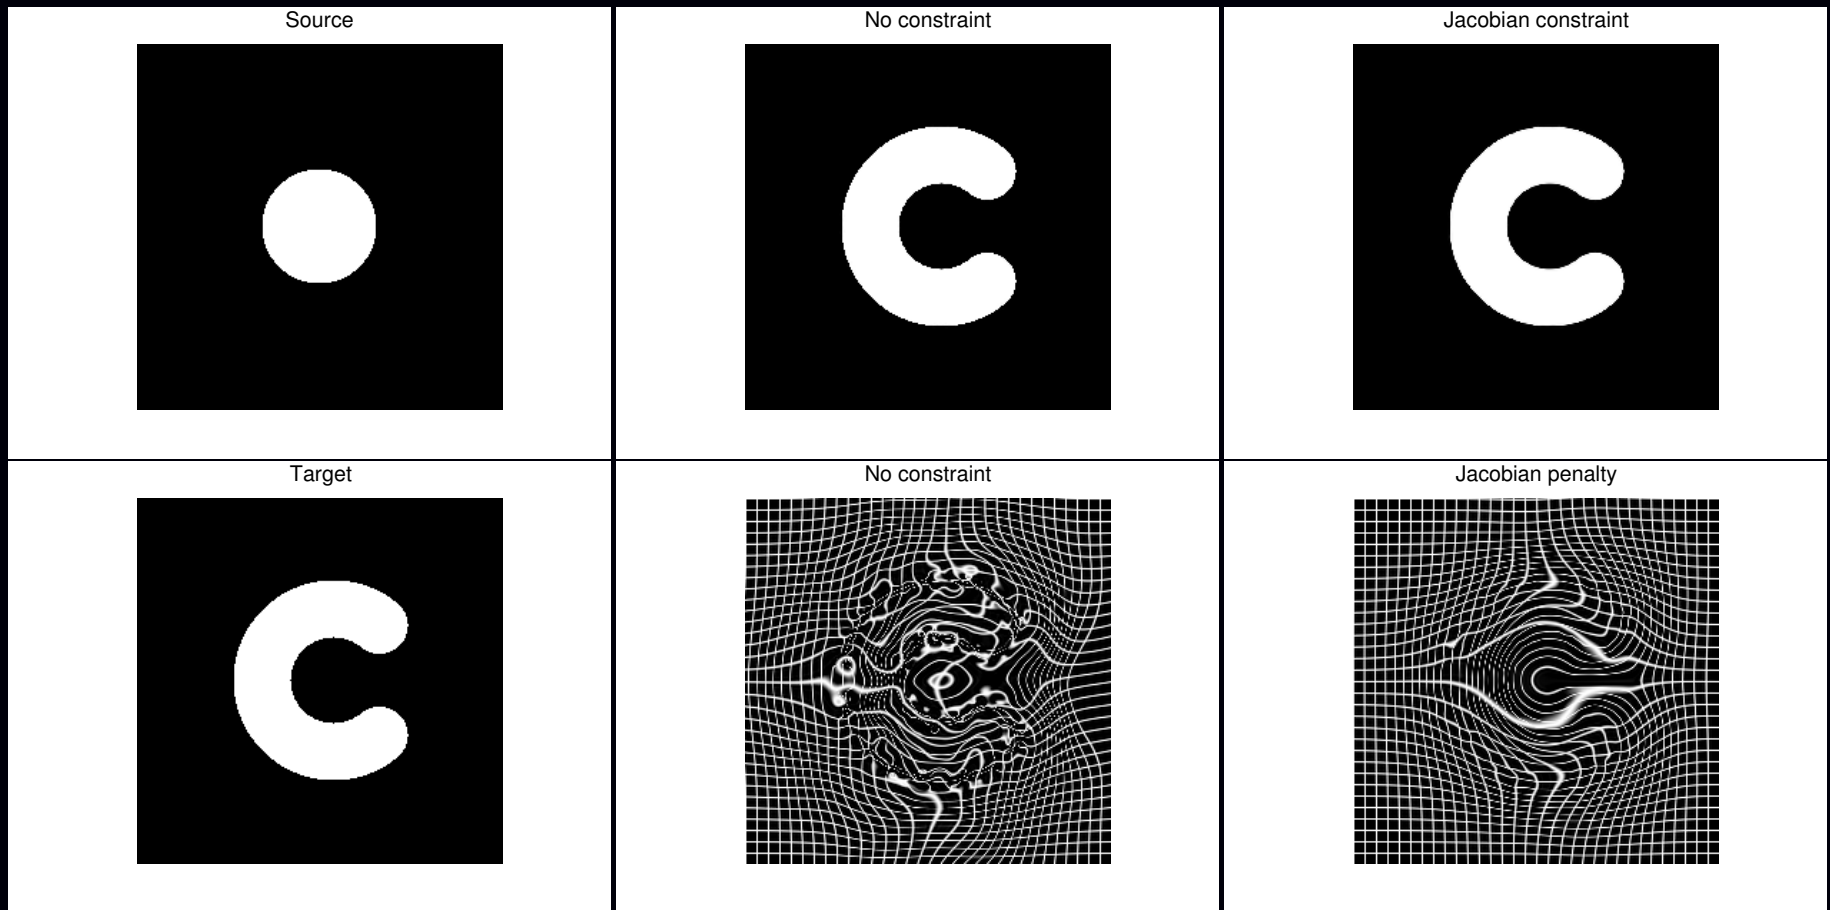


Image registration is an ill-posed problem.

Jacobian constraint on grid required $> 3\times$ computation as unconstrained case.

Nevertheless, some negative Jacobians remain (between grid points) because $C_0 \subset C_1$.

We need a simpler constraint that ensures positive Jacobian determinants everywhere.

A sufficient condition: Box constraints

Simple lower/upper bounds on B-spline coefficients:

$$\boldsymbol{\alpha} \in C_3 = \left\{ \boldsymbol{\alpha} : |\alpha_{ijk}^x| \leq \frac{m_x}{K}, |\alpha_{ijk}^y| \leq \frac{m_y}{K}, |\alpha_{ijk}^z| \leq \frac{m_z}{K}, \forall i, j, k \right\},$$

where $K \approx 2.05$ in 2D and $K \approx 2.48$ in 3D.

Choi *et al.*, 2000; Rueckert *et al.*, MICCAI 2006

Fact: $C_3 \subset C_0$.

So constraining $\boldsymbol{\alpha} \in C_3$ ensures local invertibility everywhere.

Box constraints are particularly simple for optimization.

However, C_3 is a very restrictive set of deformations.

- Maximum displacement is only about half the knot spacing.
- Precludes even simple (large) global translations.

Proposed sufficient condition for invertibility

Theorem (IEEE JSTSP, Feb. 2009):

Suppose $0 \leq k_q < \frac{1}{2}$ for $q \in \{x, y, z\}$. Define the set:

$$\begin{aligned}
 C_4 = \{ \boldsymbol{\alpha} : & -m_x k_x \leq \alpha_{i+1,j,k}^x - \alpha_{i,j,k}^x \leq m_x K_x, \\
 & -m_y k_y \leq \alpha_{i,j+1,k}^y - \alpha_{i,j,k}^y \leq m_y K_y, \\
 & -m_z k_z \leq \alpha_{i,j,k+1}^z - \alpha_{i,j,k}^z \leq m_z K_z, \\
 & |\alpha_{i+1,j,k}^q - \alpha_{i,j,k}^q| \leq m_q k_q \text{ for } q = y, z, \\
 & |\alpha_{i,j+1,k}^q - \alpha_{i,j,k}^q| \leq m_q k_q \text{ for } q = x, z, \\
 & |\alpha_{i,j,k+1}^q - \alpha_{i,j,k}^q| \leq m_q k_q \text{ for } q = x, y, \forall i, j, k \}.
 \end{aligned}$$

If $\boldsymbol{\alpha} \in C_4$, then $\forall \vec{r} \in \mathbb{R}^3$:


$$\begin{aligned}
 1 - (k_x + k_y + k_z) & \leq \det \left\{ \nabla \vec{T}(\vec{r}; \boldsymbol{\alpha}) \right\} \\
 & \leq (1 + K_x)(1 + K_y)(1 + K_z) + (1 + K_x)k_y k_z + k_x(1 + K_y)k_z + k_x k_y(1 + K_z).
 \end{aligned}$$

Corollary:

Choosing $k_x = k_y = k_z = 1/3 - \varepsilon$ ensures that $0 < \det \left\{ \nabla \vec{T}(\vec{r}; \boldsymbol{\alpha}) \right\}, \forall \vec{r}$.

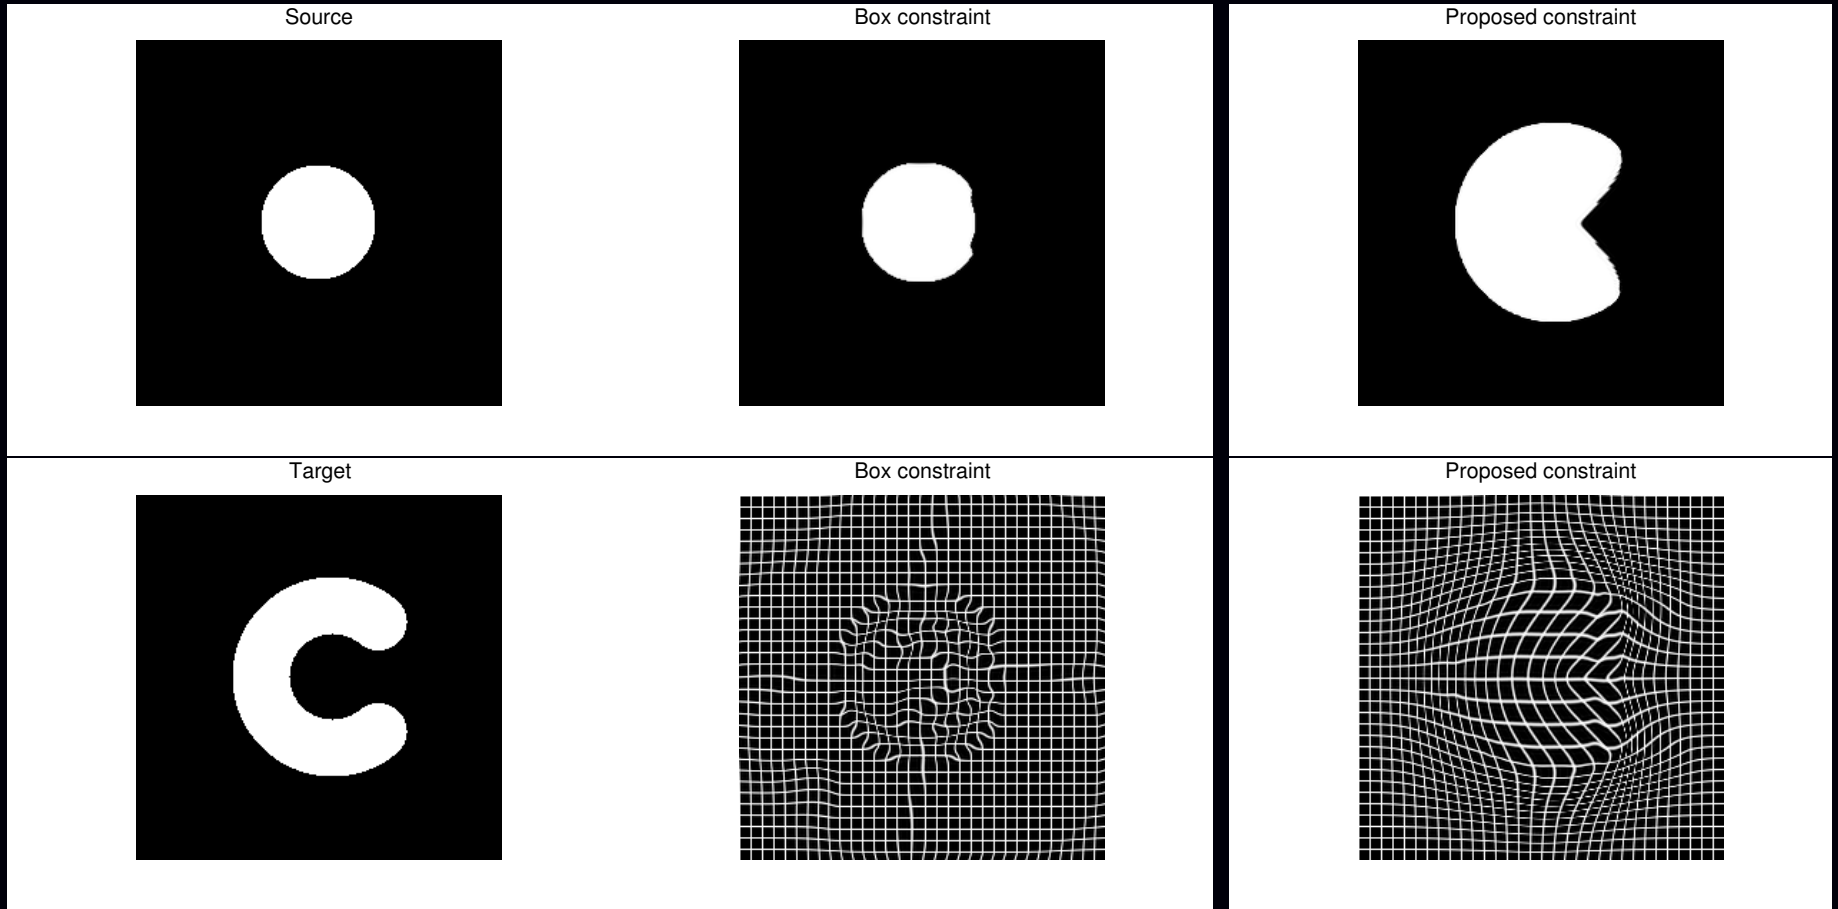
Comparing sufficient conditions

1D example with two coefficients: α_1, α_2 ,
for $n = 2$ (quadratic B-splines)



Limitations of sufficient conditions

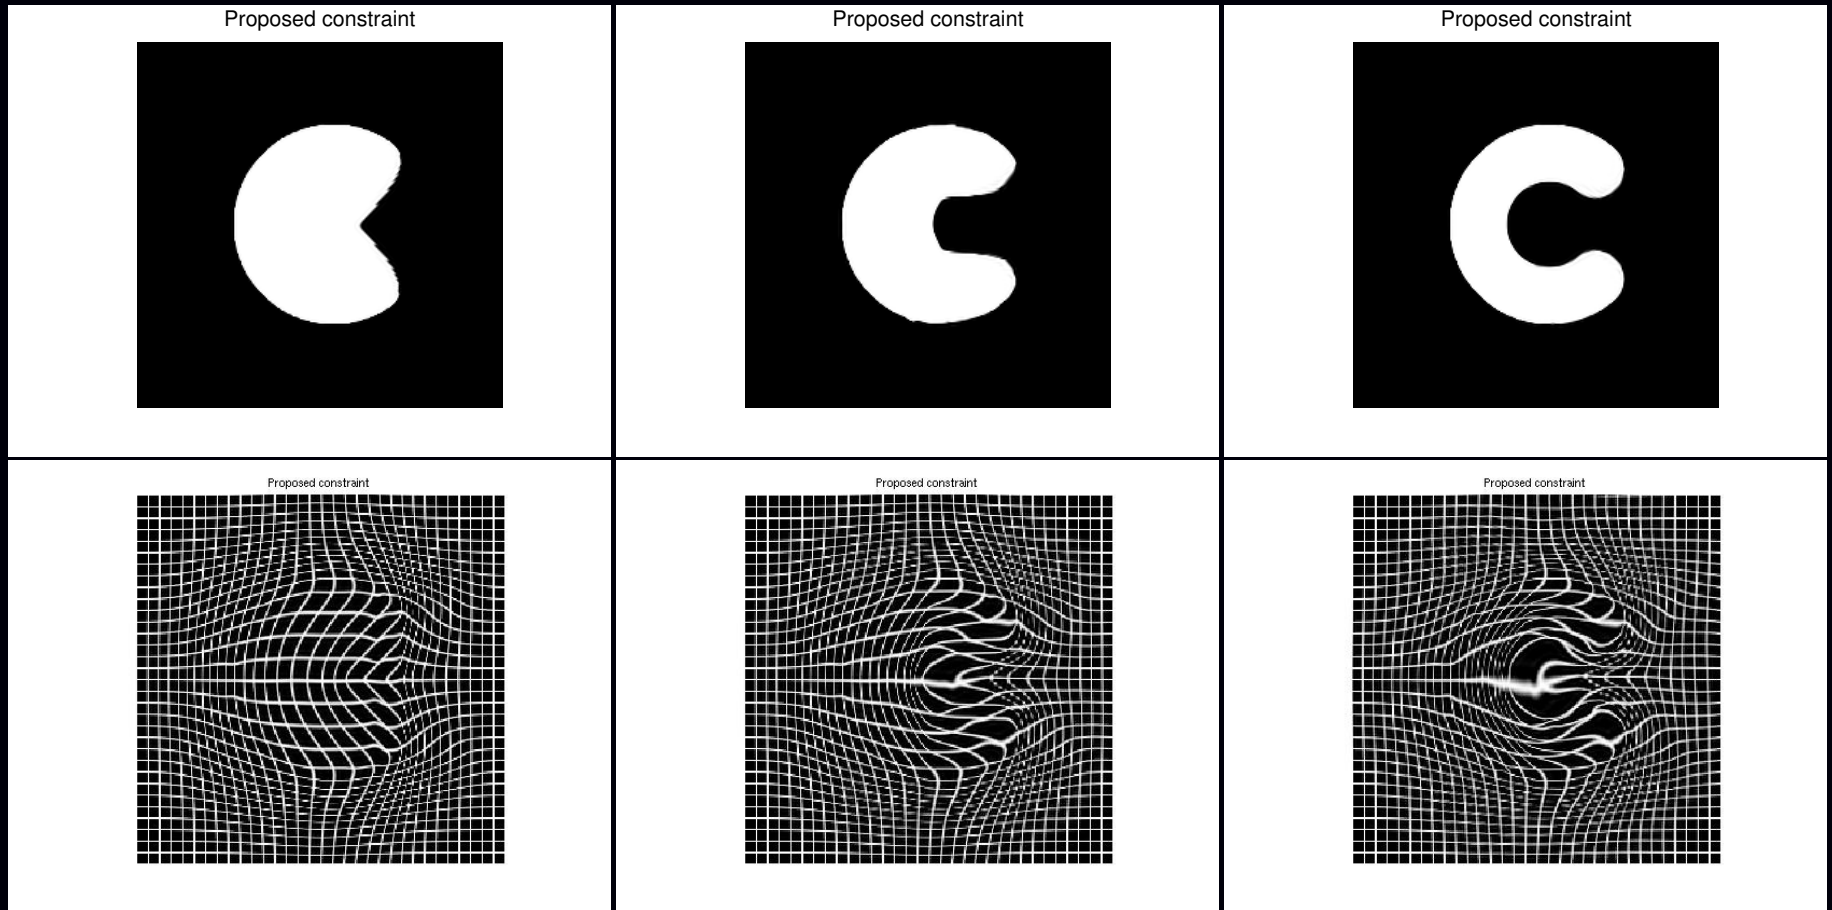
2D simulations using augmented Lagrange multiplier approach to enforce the constraint $\alpha \in C_3$ or $\alpha \in C_4$.



Clearly $C_3 \subset C_0$ and $C_4 \subset C_0$.

Solution: composition of transformations

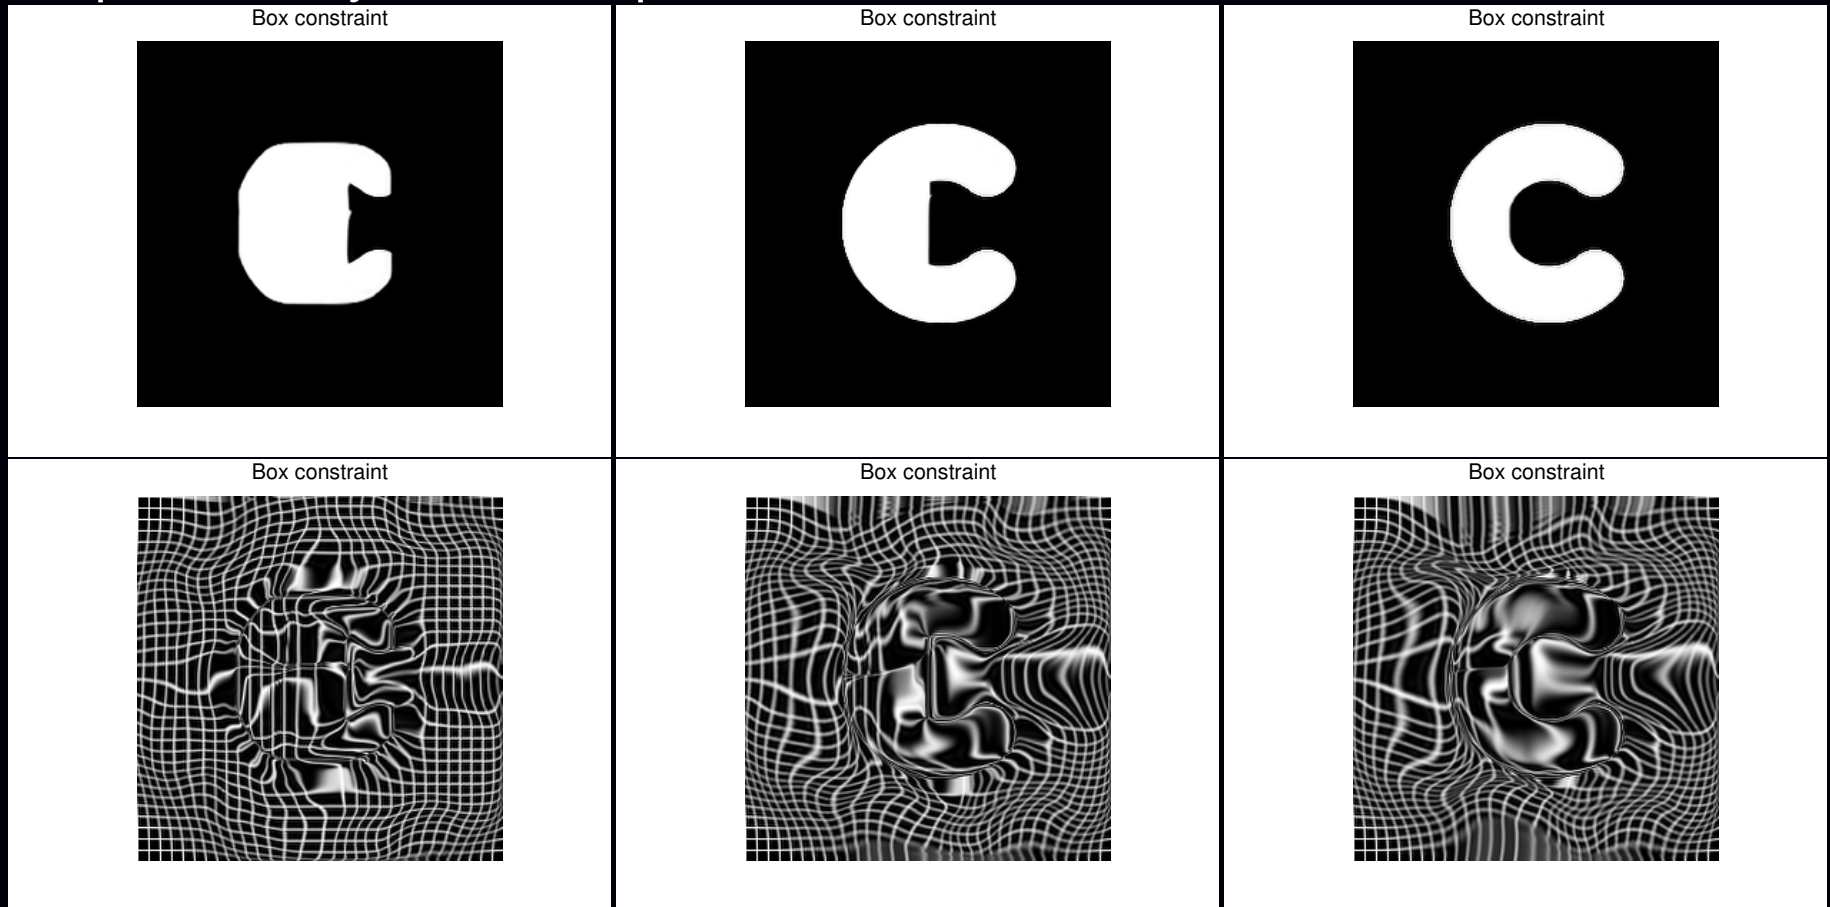
Composing multiple transformations can overcome the limitations of sufficient conditions, e.g., $\vec{T} \triangleq \vec{T}_{\alpha_3} \circ \vec{T}_{\alpha_2} \circ \vec{T}_{\alpha_1}$ where $\alpha_1, \alpha_2, \alpha_3 \in C_4$.



Composition for box constraints

Requires many more compositions:

Rueckert *et al.*, MICCAI 2006



10

20

30

Each of the 30 warps used many augmented Lagrangian iterations.
Tradeoff: simplicity of constraint and its flexibility.

Simplyfing further via regularization

Idea: replace constrained optimization

$$\hat{\alpha} = \arg \max_{\alpha \in C_4} \Psi(\alpha)$$

with simpler unconstrained, but regularized, optimization:

$$\hat{\alpha} = \arg \max_{\alpha} \Psi(\alpha) - \gamma R(\alpha)$$

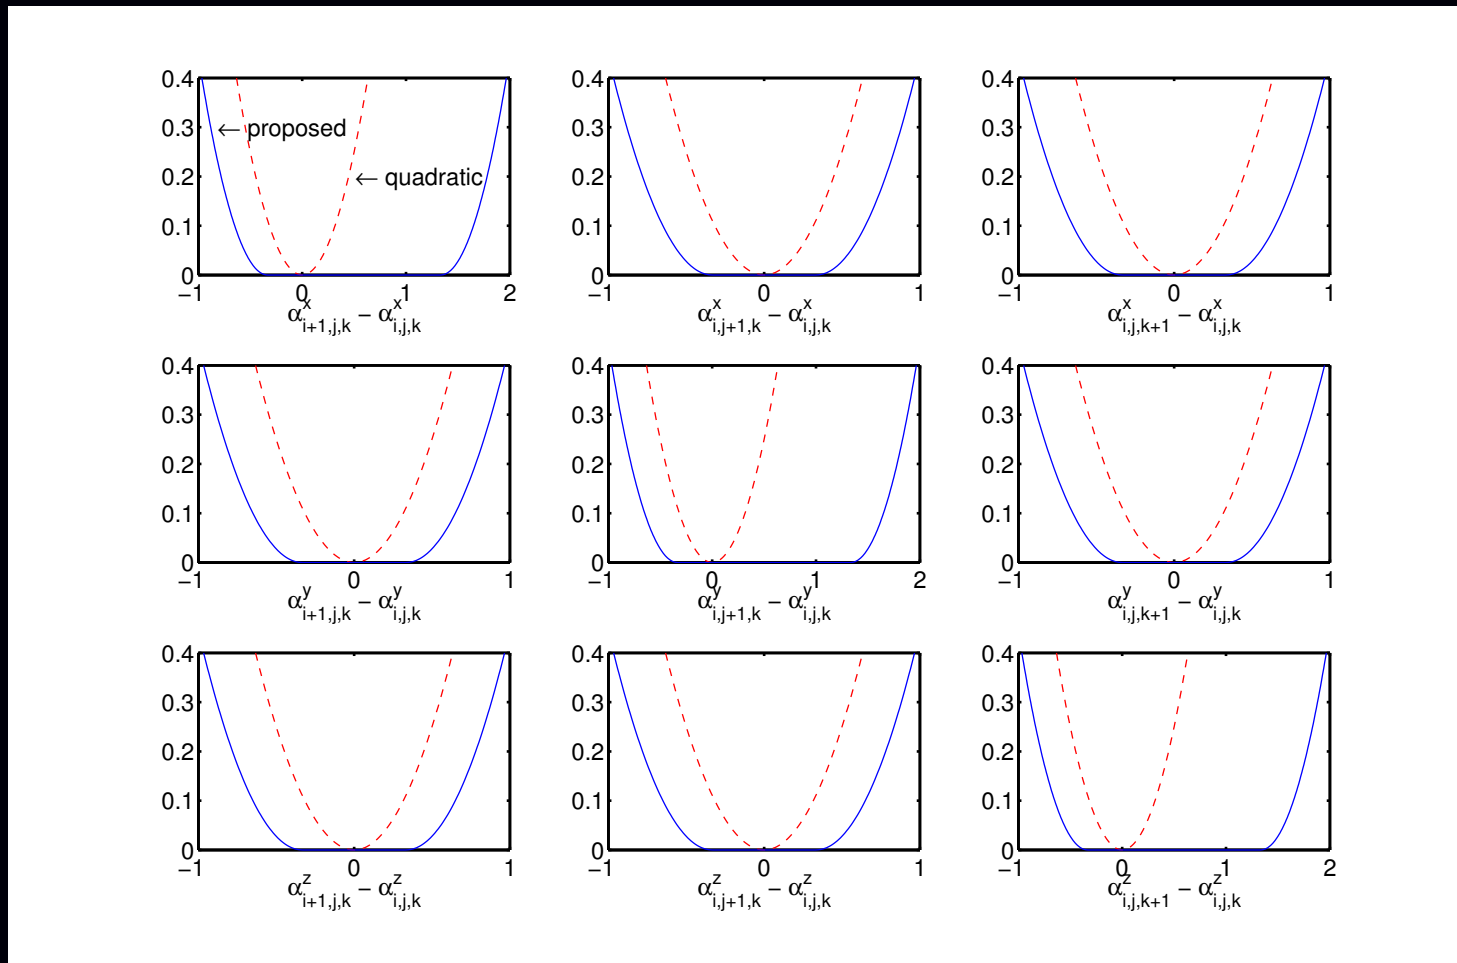
where $R(\alpha)$ is zero if $\alpha \in C_4$ but is “large” otherwise.

This *encourages* local invertibility, but does not enforce it strictly.

The regularization parameter γ controls the tradeoff between

- image similarity
- regularity of the deformation (local invertibility).

Proposed regularizer



Interval constraints in C_4 replaced by piecewise quadratic penalty function of differences of neighboring B-spline coefficients.

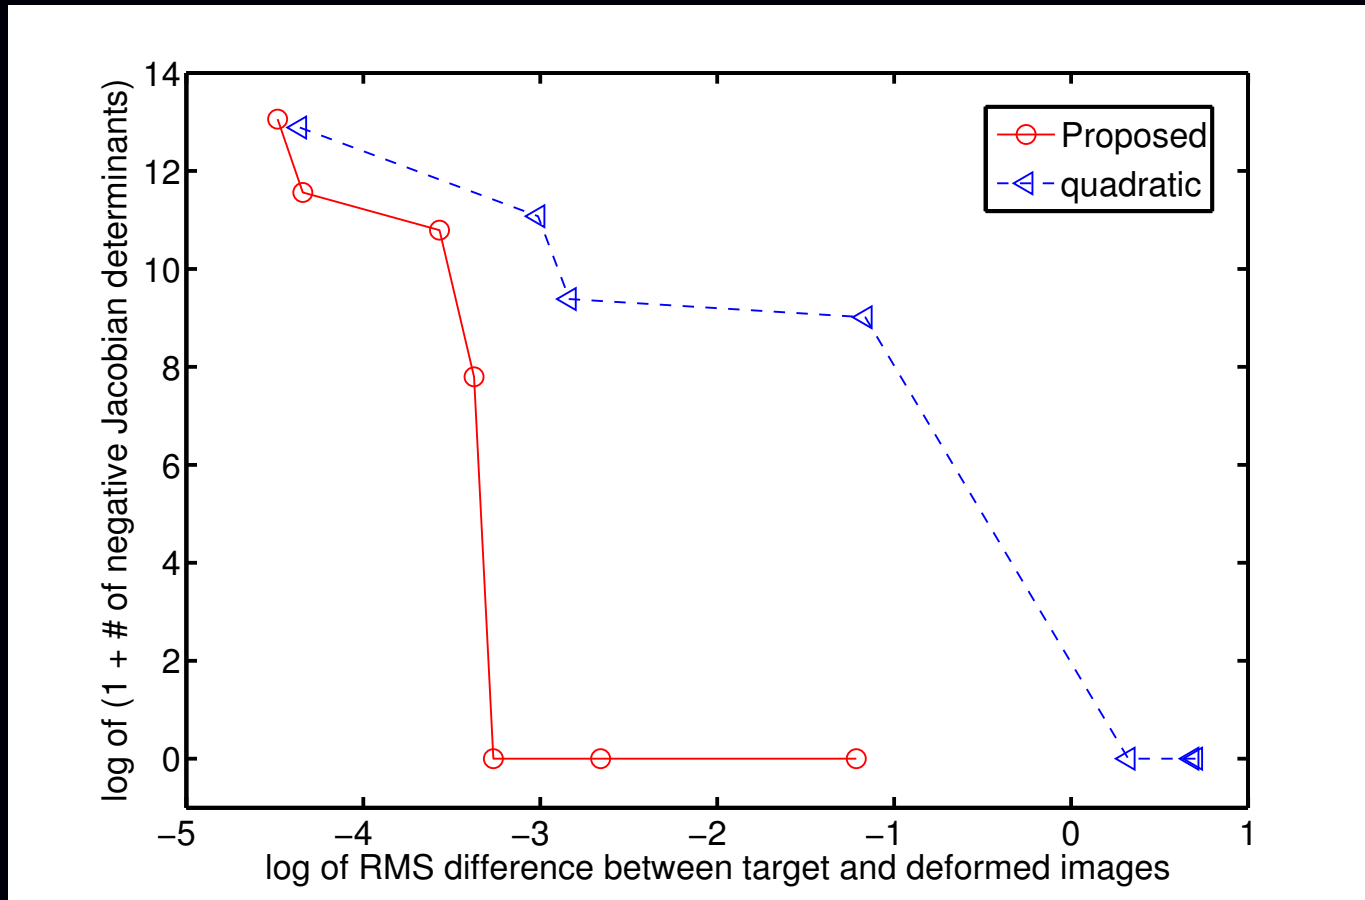
cf. conventional quadratic roughness regularization

$m_x = m_y = m_z = 1, k_x = k_y = k_z = 1/3$, and $K_x = K_y = K_z = 4/3$.

Regularization tradeoffs


As regularization parameter $\gamma \uparrow$

- # of negative Jacobian determinants \downarrow so
- RMS difference between images \uparrow



Proposed regularizer: good image similarity, few negative Jacobian determinants.

Regularization tradeoffs: Details



3D registration of CT inhale/exhale scans

3D CT scans of a cancer patient at exhale and inhale, for radiation treatment planning. $396 \times 256 \times 128$ voxels.

Source: Coronal




Target: Coronal




CT inhale/exhale scans: Sagittal

Source: Sagittal




Target: Sagittal





3D registration results: Unconstrained


No constraint: Sagittal



No constraint: Sagittal




No constraint: Sagittal




Relatively small difference image values,
but many negative Jacobian determinants

3D registration results: Jacobian


Jacobian penalty: Sagittal



Jacobian penalty: Sagittal

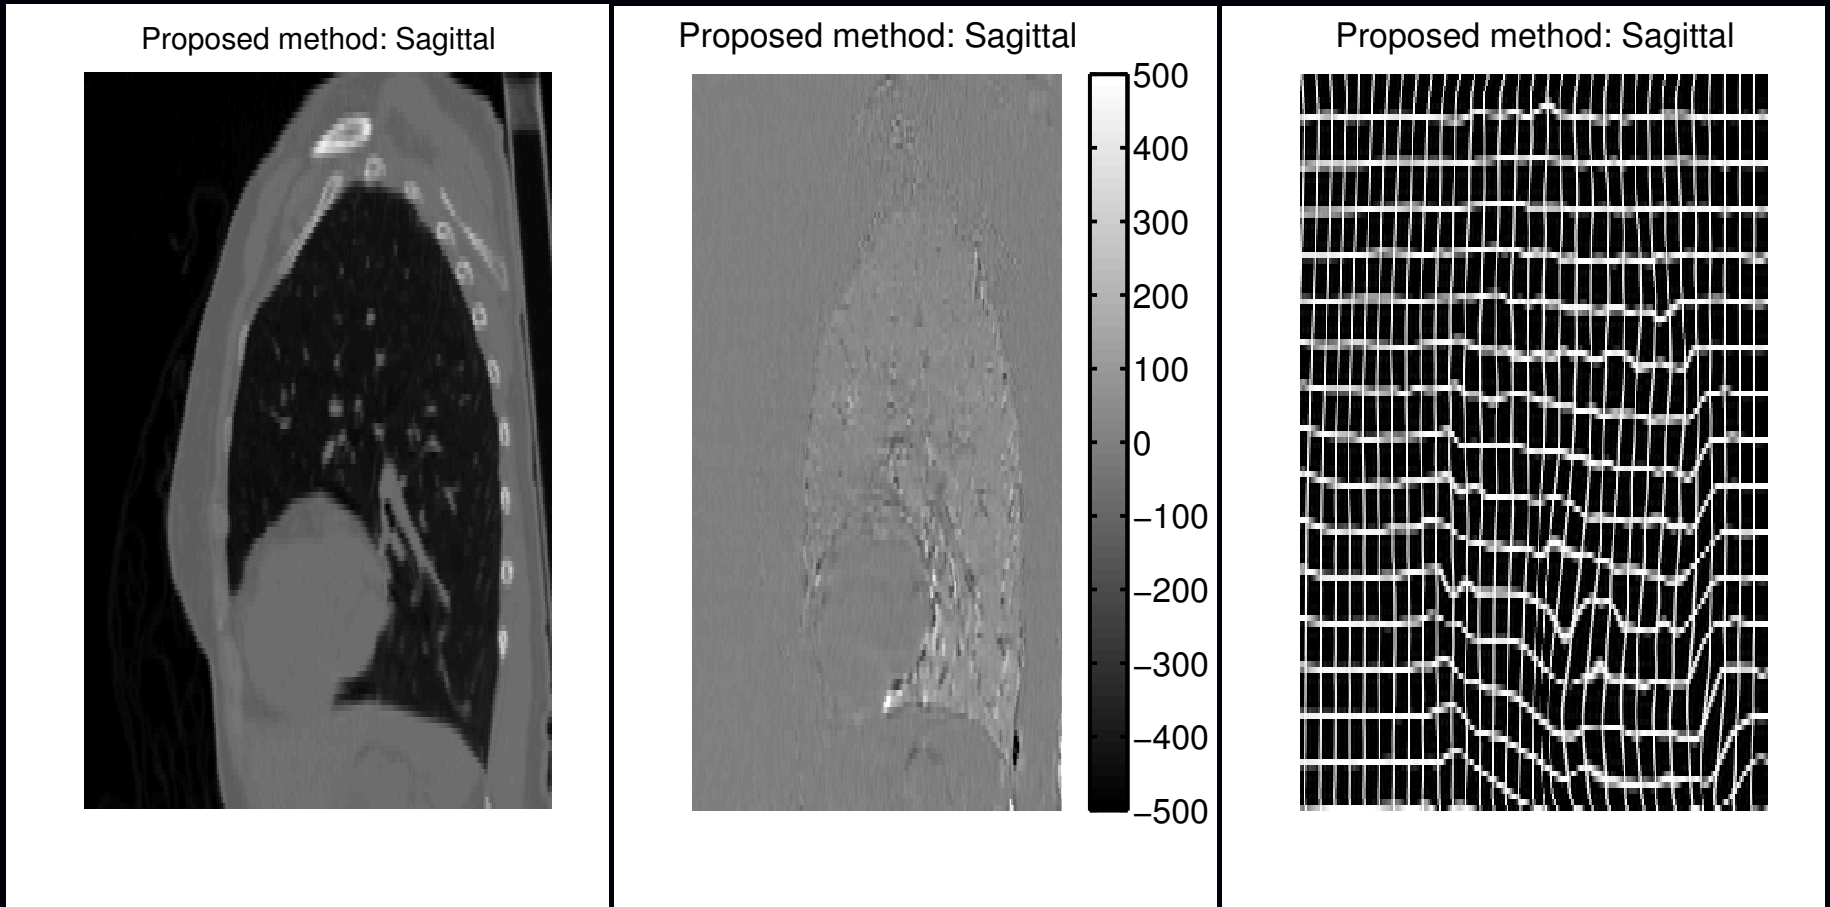


Jacobian penalty: Sagittal



Jacobian penalty based on C_1 (grid)
Better behaved warp and reasonable difference image.
But slow.

3D registration results: Proposed regularizer



Proposed regularizer based on C_4 .

Tailored design of constraint/penalty: $k_x = k_y = 1/4$, $k_z = 1/2$.

Similar warp and difference image, but faster.

Quantification

Method	CPU time (seconds)	RMS difference (HU)	# negative Jacobians
Unconstrained	25.7	19.9	316914
Jacobian penalty	81.1	25.9	0
Proposed penalty	27.4	29.2	0

Computation time per iteration (in seconds) at the finest level


Regularization parameter adjusted empirically in both penalized cases to be the smallest value that yields no negative Jacobian determinants on the voxel grid.

3 multiresolution levels: knot spacings 8 pixels with downsampled images, 8 pixels with original images, 4 pixels with original images. 120 iterations of CG at each level.

Work in progress to compose a couple coarse-scale deformations before refining to fine scale to reduce RMS differences.


Code on web site: <http://www.eecs.umich.edu/~fessler>

Diffeomorphisms: To be or not to be...



Sliding at diaphragm / rib cage interface.
Enforcing smoothness leads to bone warping.

Work in progress...



Requiring the warp to be smooth everywhere seems suboptimal. One possible solution presented at SPIE 2009. Stay tuned...

Summary

- Simple condition for ensuring local invertibility everywhere
- Admits more deformations than conventional box constraints
- Simple regularizer requires comparable computation as unregularized image registration and much less computation than Jacobian determinant constraints / penalties

Open problems

- Rigid structures (bones)
- Sliding tissue interfaces
- Parameter selection
- Computation (GPUs?)
- Performance characterization

Part 2
**Model-based image reconstruction
with motion compensation**

Motion in image reconstruction

Object being scanned: $f(\vec{r}, t)$

Measured data vector: $\mathbf{y} = (y_1, \dots, y_M)$

Static image reconstruction:

- Assume $f(\vec{r}, t) = f(\vec{r}, t_0) = f(\vec{r})$ during scan.
- Estimate $f(\vec{r})$ from measurements \mathbf{y} . (Ill posed.)

Dynamic image reconstruction (“List-mode” data model)

- Assume each data point y_i is acquired instantaneously at a corresponding time instant t_i
- Relate y_i to object at time t_i , *e.g.*,

$$y_i = \underbrace{\int a_i(\vec{r}) f(\vec{r}, t_i) d\vec{r}}_{\text{physics}} + \underbrace{\epsilon_i}_{\text{statistics}}$$

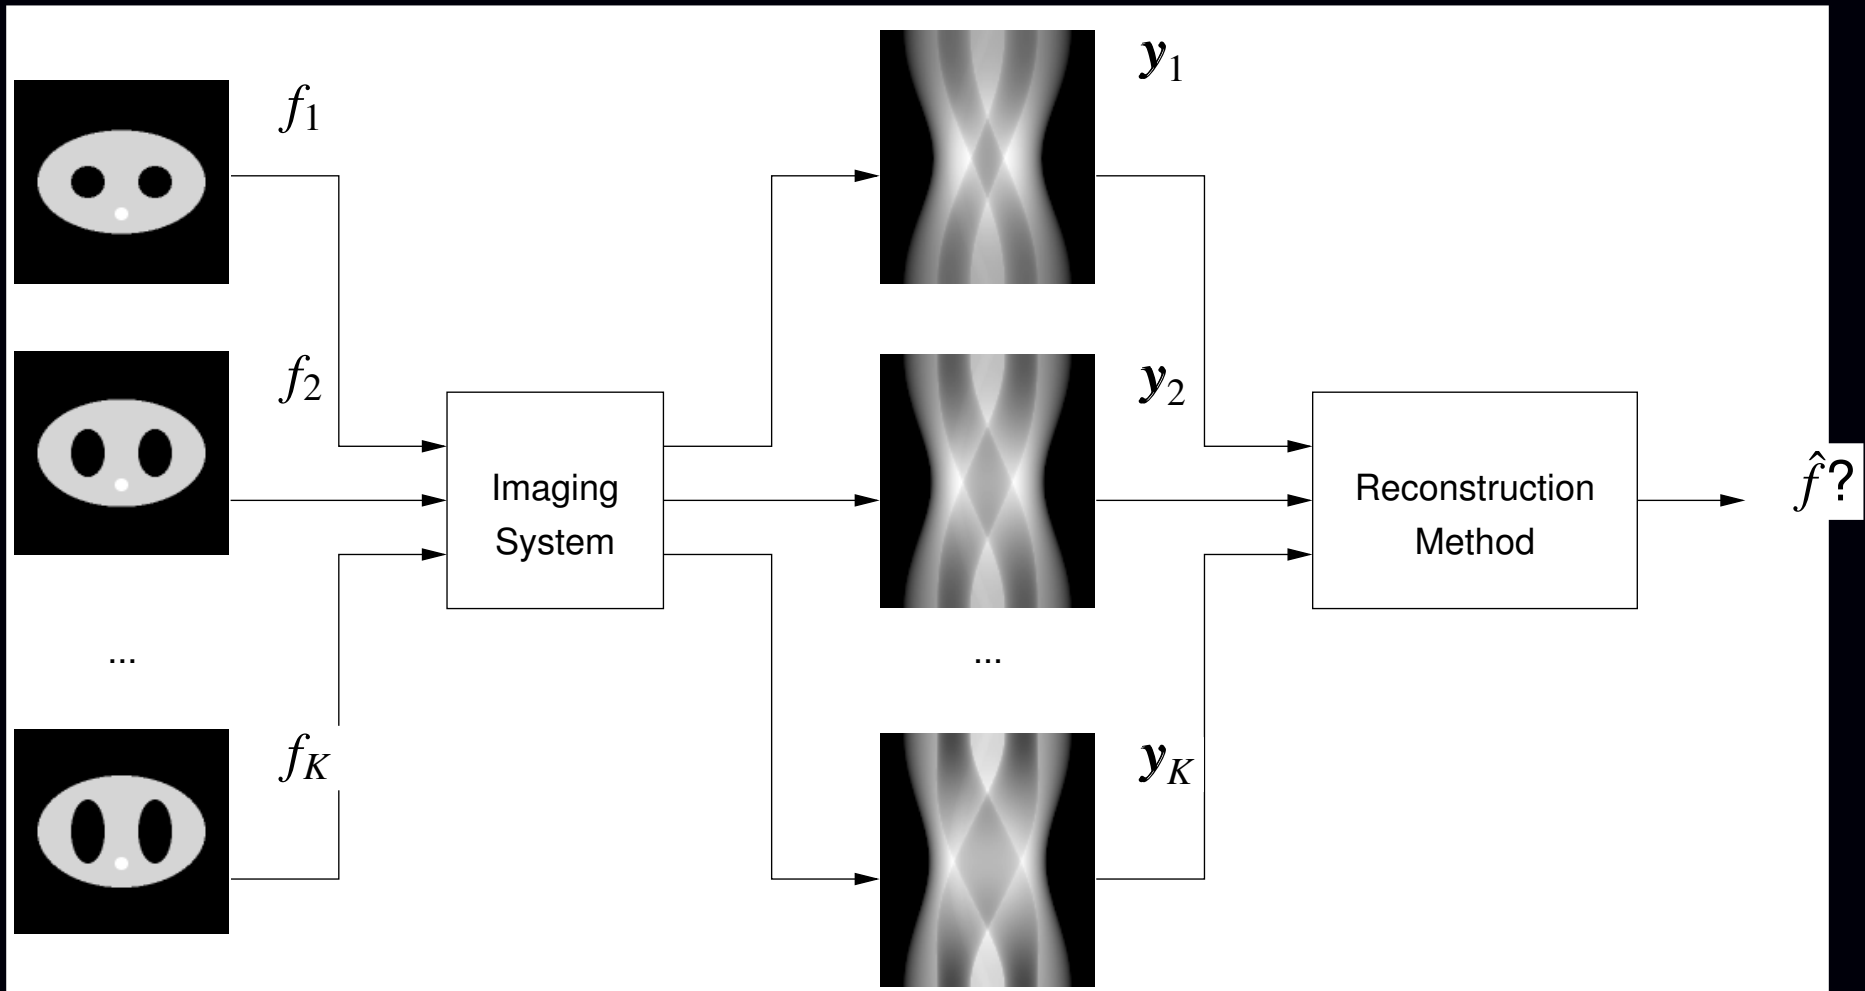
More generally: $p(y_i | f(\cdot, t_i))$.

- Estimate $f(\vec{r}, t)$ from measurements \mathbf{y} . (Even “more” ill posed!)

Gated data model

- Group data into K vectors, e.g., K respiratory phases: $\mathbf{y}_1, \dots, \mathbf{y}_K$
- Assume $f(\vec{r}, t)$ is stationary during k th phase of data acquisition
- Relate \mathbf{y}_k to $f_k(\vec{r}) \triangleq f(\vec{r}, t_k)$ using physics and statistics
- From K data vectors $\mathbf{y}_1, \dots, \mathbf{y}_K$, reconstruct object: ?

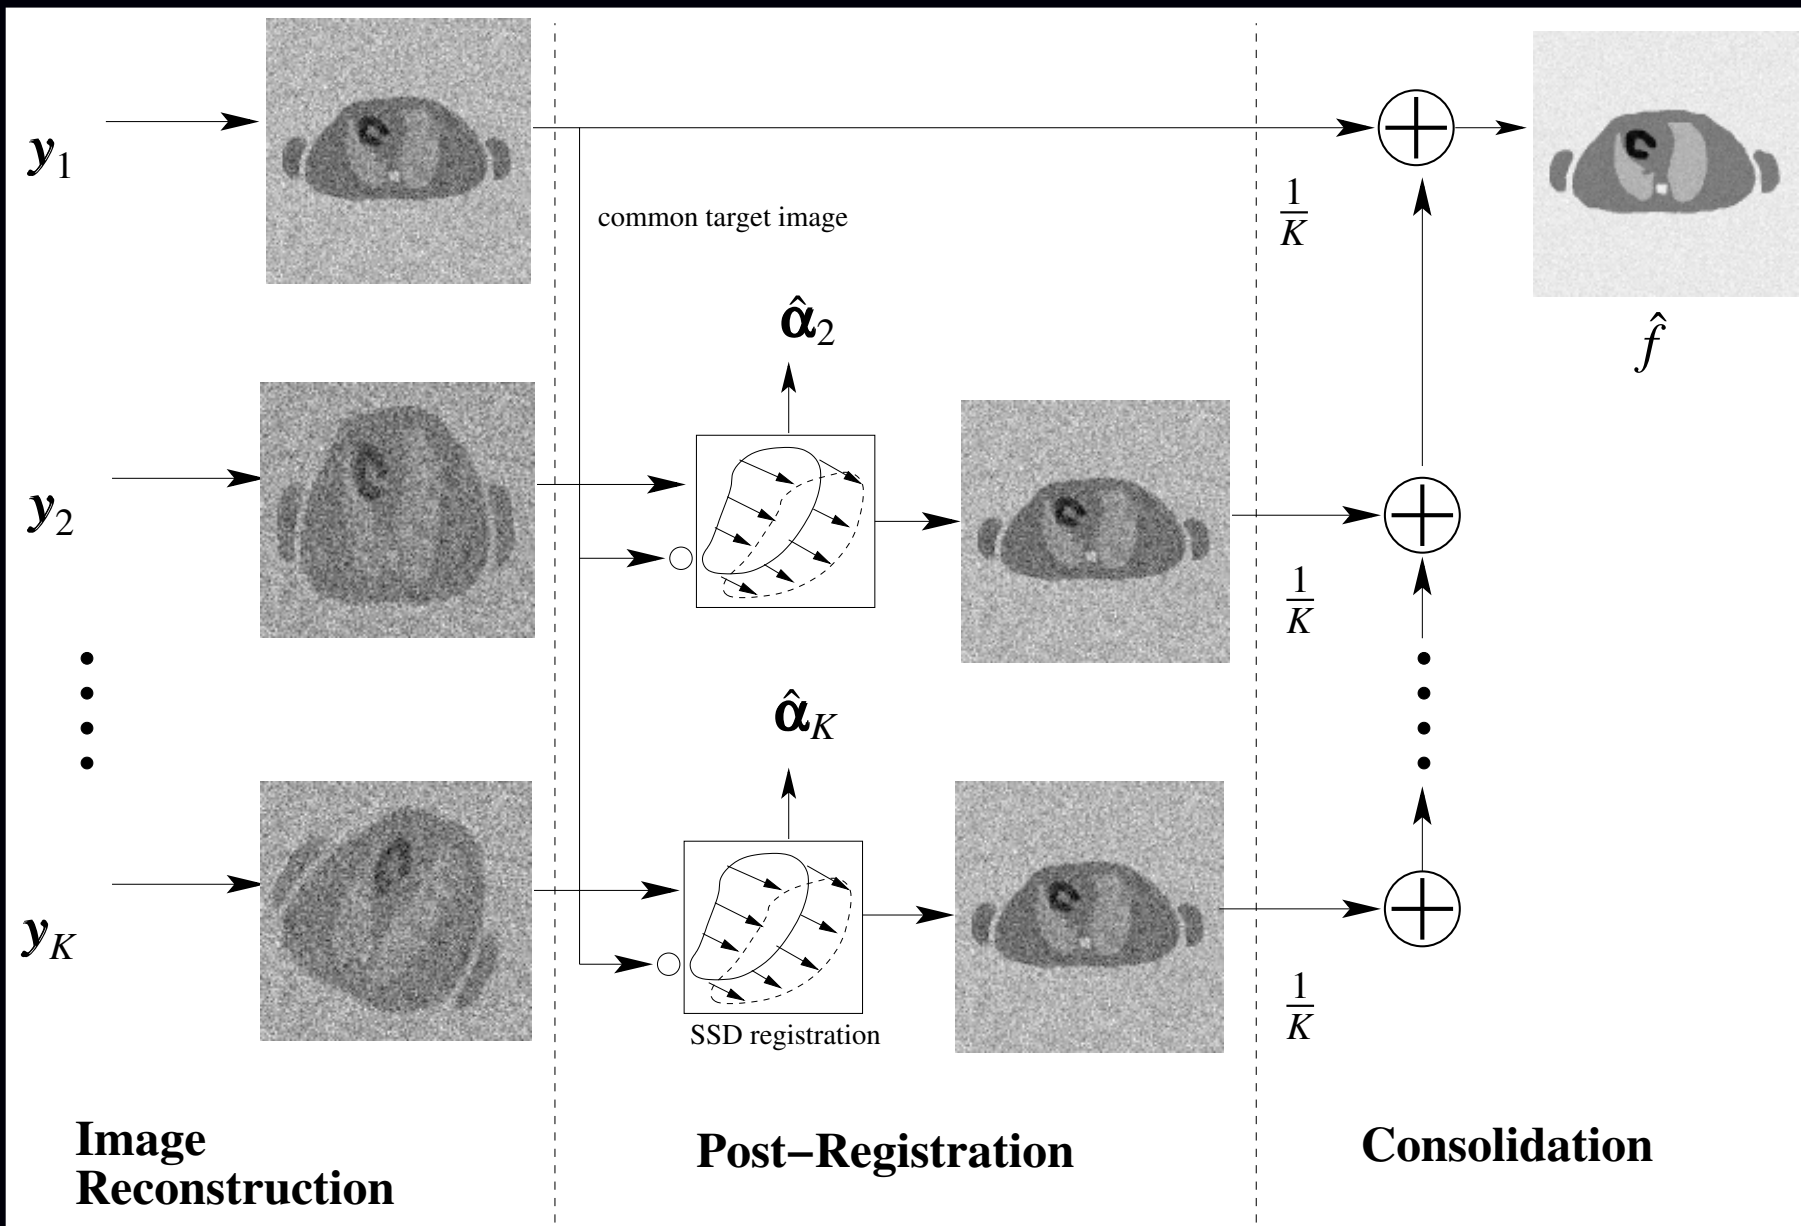
Gated data model: Illustration



Gated data: Image reconstruction options

- Pool all data and ignore motion
 - fast
 - low noise variance
 - motion induced blur
- Frame-wise: reconstruct each gate/frame separately: $\mathbf{y}_k \mapsto \hat{f}_k$
 - simple
 - high noise variance
 - no motion blur (except within-gate motion)
- Frame-wise with post-reconstruction averaging (FW-PRA)
 - map each reconstructed frame onto 1st frame, then average
 - averaging should reduce noise
 - should avoid motion blur if registration is accurate
 - registration accuracy limited by noise in the individual gates
- FW-PRA with motion estimates from a separate modality
 - use another modality (*e.g.*, PET-CT) to estimate motion
 - performance depends on consistency of motion between modalities
 - Thorndyke *et al.*, Med Phys, 2006
- ...

Frame-wise with post-reconstruction averaging



Forward model with motion

Post-reconstruction averaging assumes an implicit model that relates the frames f_2, \dots, f_K to the first frame f_1 .

We now make the (motion) model explicit:

$$f_k = \mathbf{W}(\boldsymbol{\alpha}_k) f_1, \quad k = 2, \dots, K.$$

$\mathbf{W}(\boldsymbol{\alpha}_k)$ is the linear (!) transformation of the image values corresponding to motion $\boldsymbol{\alpha}_k$.

This model suggests additional image reconstruction approaches.

Linear interpolation and Nonrigid deformations

B-spline interpolation model for continuous-space image:

$$f(\vec{r}) = f(x, y, z) = \sum_{n,m,l} c_{nml} \beta(x-n) \beta(y-m) \beta(z-l)$$

Find coefficients $\mathbf{c} = \{c_{nml}\}$ by prefiltering digital image $f[n, m, l]$.

Nonrigid deformation of f :

$$\begin{aligned} g(\vec{r}) &= f\left(\vec{T}(\vec{r}; \boldsymbol{\alpha})\right) \\ &= \sum_{n,m,l} c_{nml} \beta(T^x(\vec{r}; \boldsymbol{\alpha}) - n) \beta(T^y(\vec{r}; \boldsymbol{\alpha}) - m) \beta(T^z(\vec{r}; \boldsymbol{\alpha}) - l). \end{aligned}$$

Resample warped image on grid (of target image):

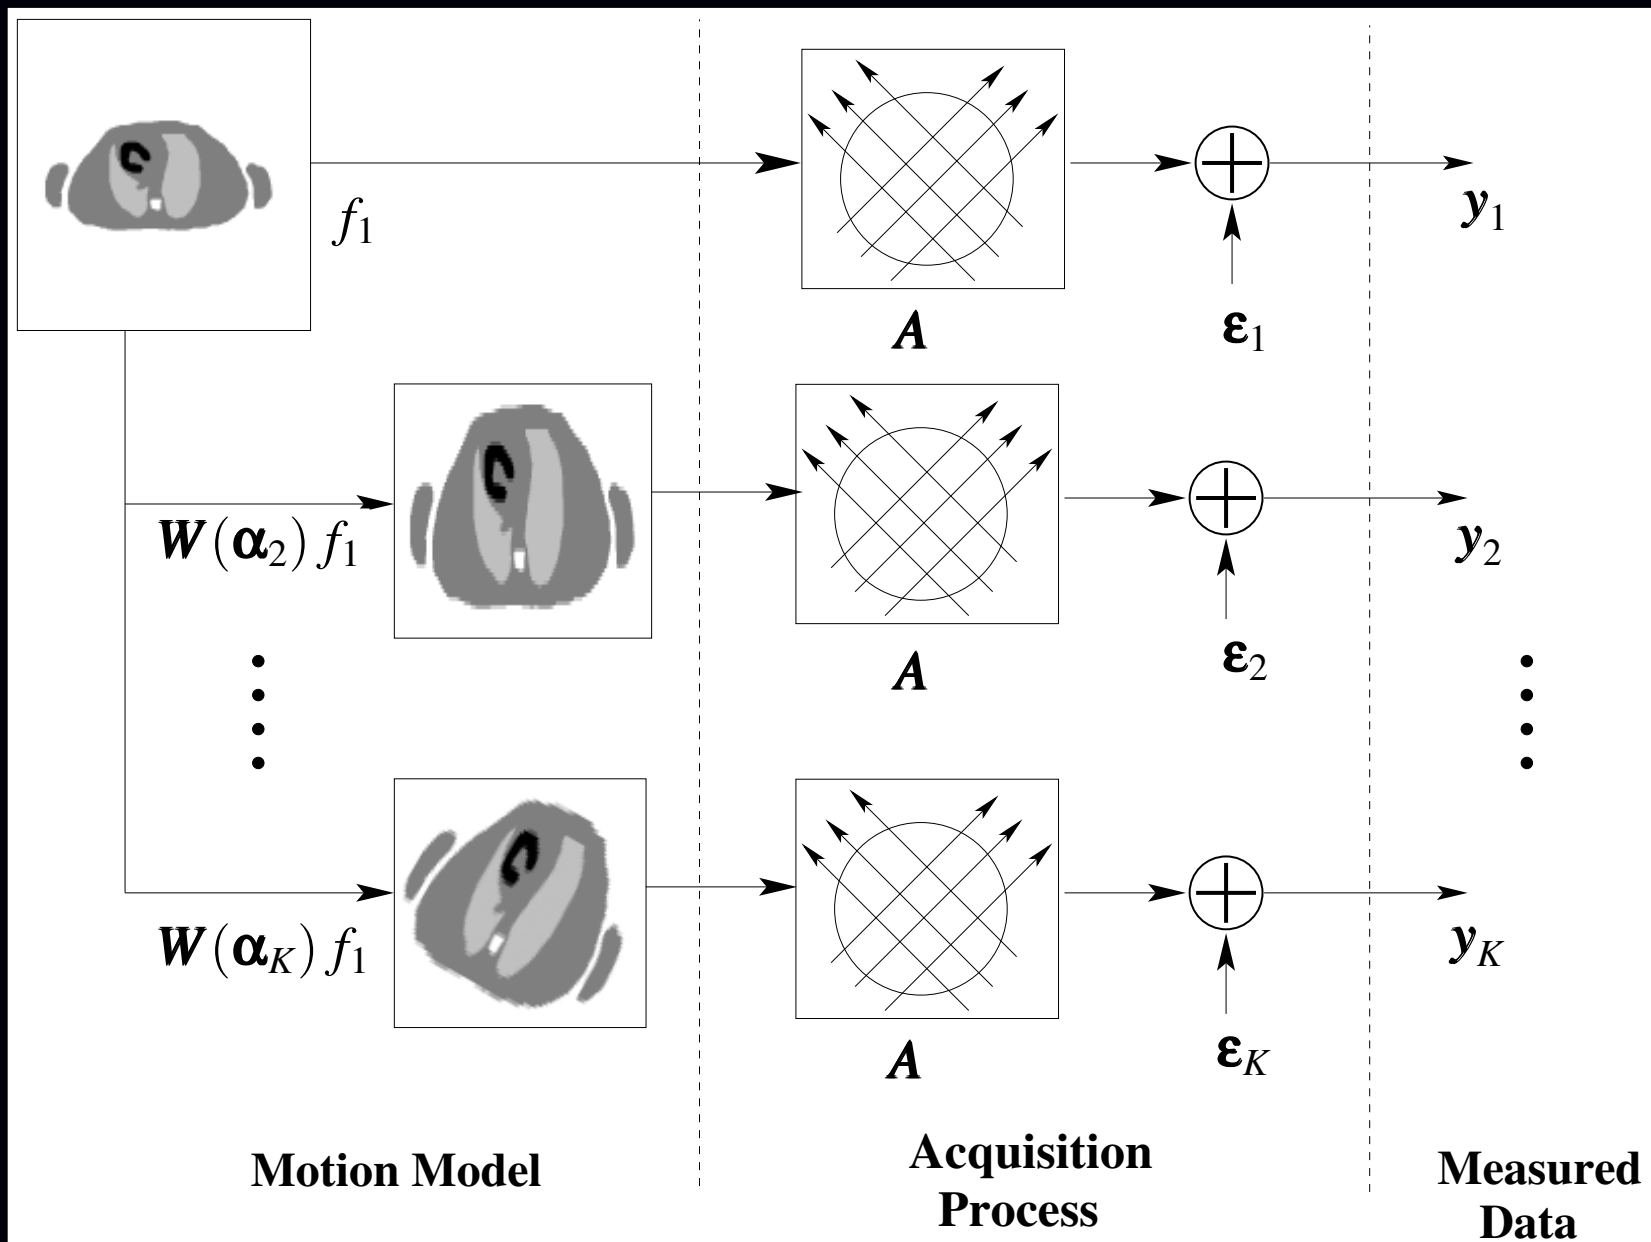
$$g(\vec{r}_j) = \sum_{n,m,l} c_{nml} W_{\vec{r}_j;n,m,l}(\boldsymbol{\alpha}), \quad j = 1, \dots, N$$

$$W_{\vec{r};n,m,l}(\boldsymbol{\alpha}) \triangleq \beta(T^x(\vec{r}; \boldsymbol{\alpha}) - n) \beta(T^y(\vec{r}; \boldsymbol{\alpha}) - m) \beta(T^z(\vec{r}; \boldsymbol{\alpha}) - l)$$

In matrix-vector form, where $\mathbf{g} = \{g(\vec{r}_j)\}$ and $\mathbf{f} = \{f[n, m, l]\}$:

$$\mathbf{g} = \mathbf{W}(\boldsymbol{\alpha})\mathbf{c}, \quad \mathbf{c} = \mathbf{W}^{-1}(\mathbf{0})\mathbf{f} \implies \mathbf{g} = \mathbf{W}(\boldsymbol{\alpha})\mathbf{W}^{-1}(\mathbf{0})\mathbf{f}.$$

Forward model with motion



Gated data: More image reconstruction options

- Model-based image reconstruction with motion compensation
 - given motion estimates, from FW-PRA or from separate modality,
 - compensate for motion in reconstruction process.
 - Qiao *et al.*, PMB 2006; Taguchi *et al.*, SPIE 2007.
- Model-based image reconstruction jointly with registration
 - Jacobson & Fessler, IEEE NSS-MIC 2003, IEEE SSP 2003, ISBI 2006
 - Odille *et al.*, MRM 2008
 - estimate jointly the first frame and the motion from all data
- Model-based image reconstruction with temporal regularization
 - Mair *et al.*, IEEE T-MI, 2006
 - estimate all frames and the motion between frames from all data

Model-based image reconstruction with motion compensation

- Given: motion estimates $\hat{\alpha}_k$ for $k = 2, \dots, K$, from FW approach or from a separate modality,
- model for system physics / statistics: $p(\mathbf{y}_k | f_k) = p(\mathbf{y}_k | \mathbf{W}(\hat{\alpha}_k) f_1)$.

Perform penalized-likelihood (aka MAP) estimation of *one* image:

$$\hat{f}_1 = \arg \max_{f_1} \Psi(f_1; \hat{\alpha}_2, \dots, \hat{\alpha}_K; \mathbf{y}_1, \dots, \mathbf{y}_K)$$

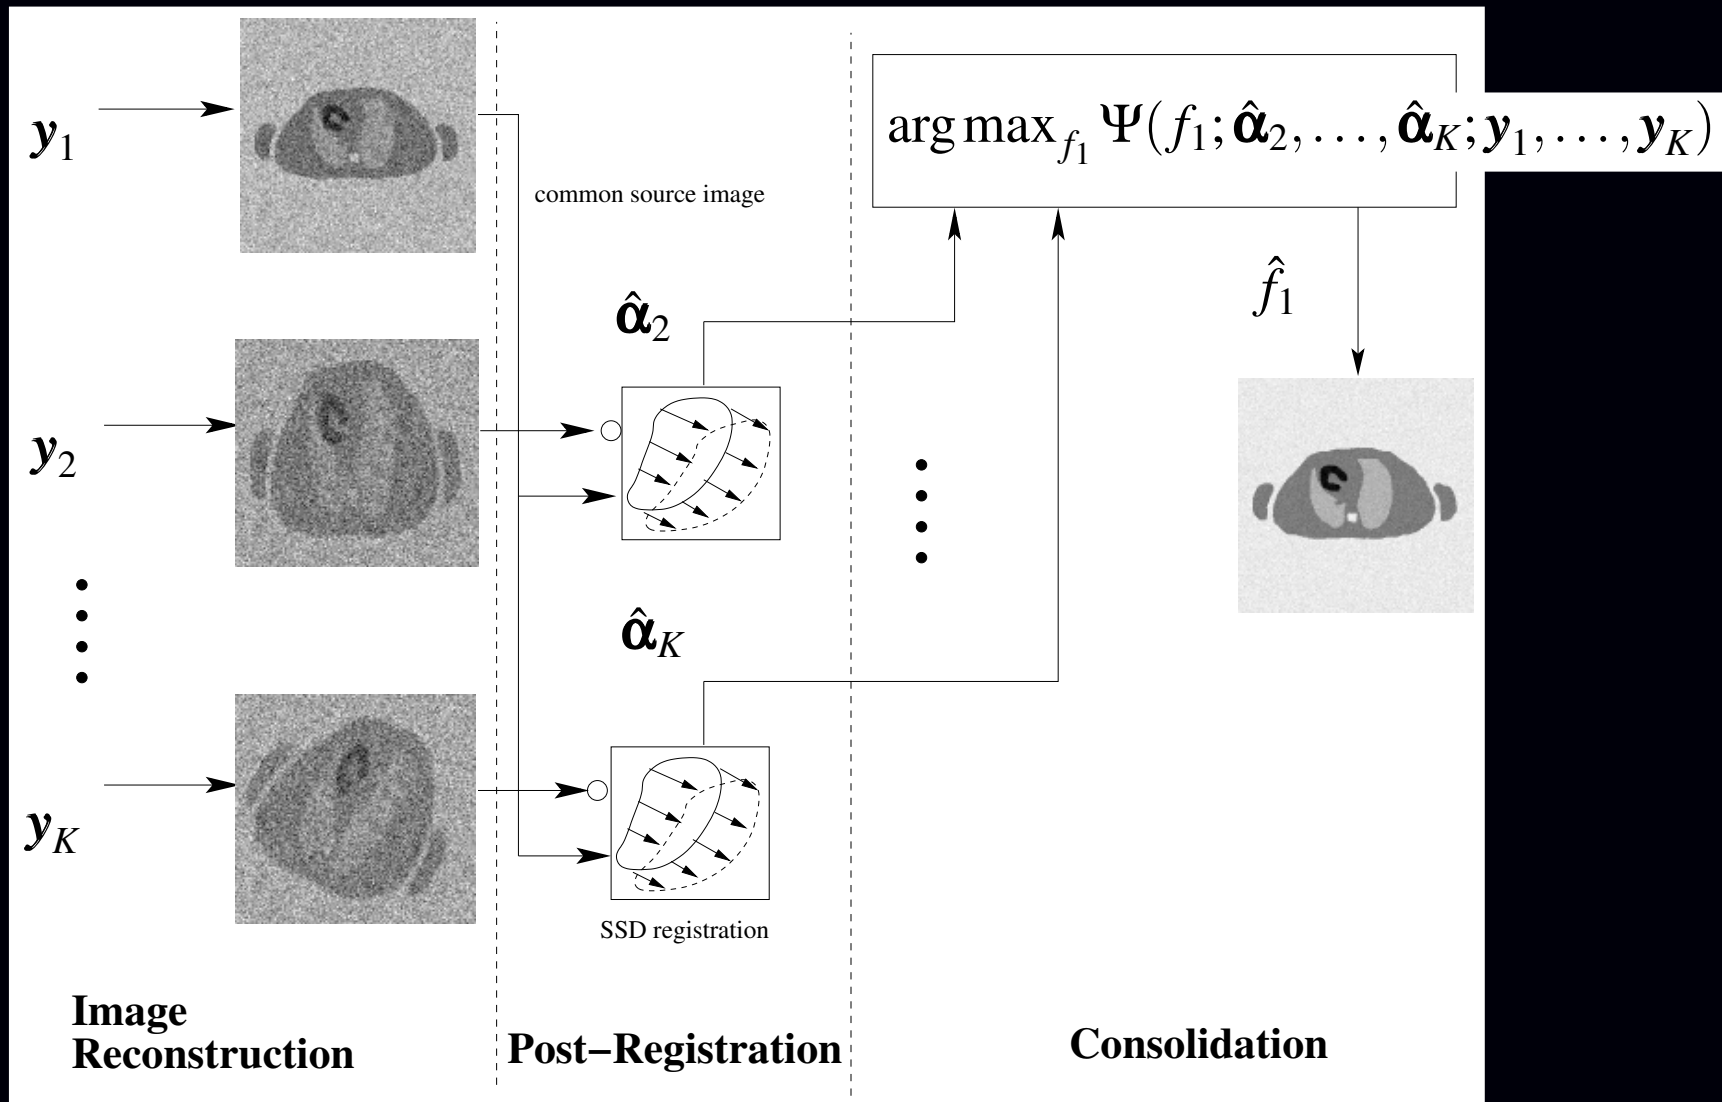
$$\Psi(f_1; \hat{\alpha}_2, \dots, \hat{\alpha}_K; \mathbf{y}_1, \dots, \mathbf{y}_K) \triangleq \sum_{k=1}^K \log p(\mathbf{y}_k | \mathbf{W}(\hat{\alpha}_k) f_1) - \beta R(f_1).$$

$R(f)$ is optional regularization to control noise in ill-posed image reconstruction problems.

For linear model with additive gaussian noise $\mathbf{y}_k = \mathbf{A}_k f_k + \boldsymbol{\varepsilon}_k$:

$$\hat{f}_1 = \arg \min_{f_1} \sum_{k=1}^K \|\mathbf{y}_k - \mathbf{A}_k \mathbf{W}(\hat{\alpha}_k) f_1\|^2 + \beta R(f_1).$$

Motion compensated image reconstruction



Joint image reconstruction / registration

Previous approach used possibly suboptimal motion estimates:

$$\hat{f}_1 = \arg \max_{f_1} \Psi(f_1; \hat{\alpha}_2, \dots, \hat{\alpha}_K; \mathbf{y}_1, \dots, \mathbf{y}_K)$$

Alternative: jointly estimate one image and $K - 1$ deformation parameters:

$$(\hat{f}_1, \hat{\alpha}_2, \dots, \hat{\alpha}_K) = \arg \max_{f_1, \alpha_2, \dots, \alpha_K} \Psi(f_1; \alpha_2, \dots, \alpha_K; \mathbf{y}_1, \dots, \mathbf{y}_K)$$

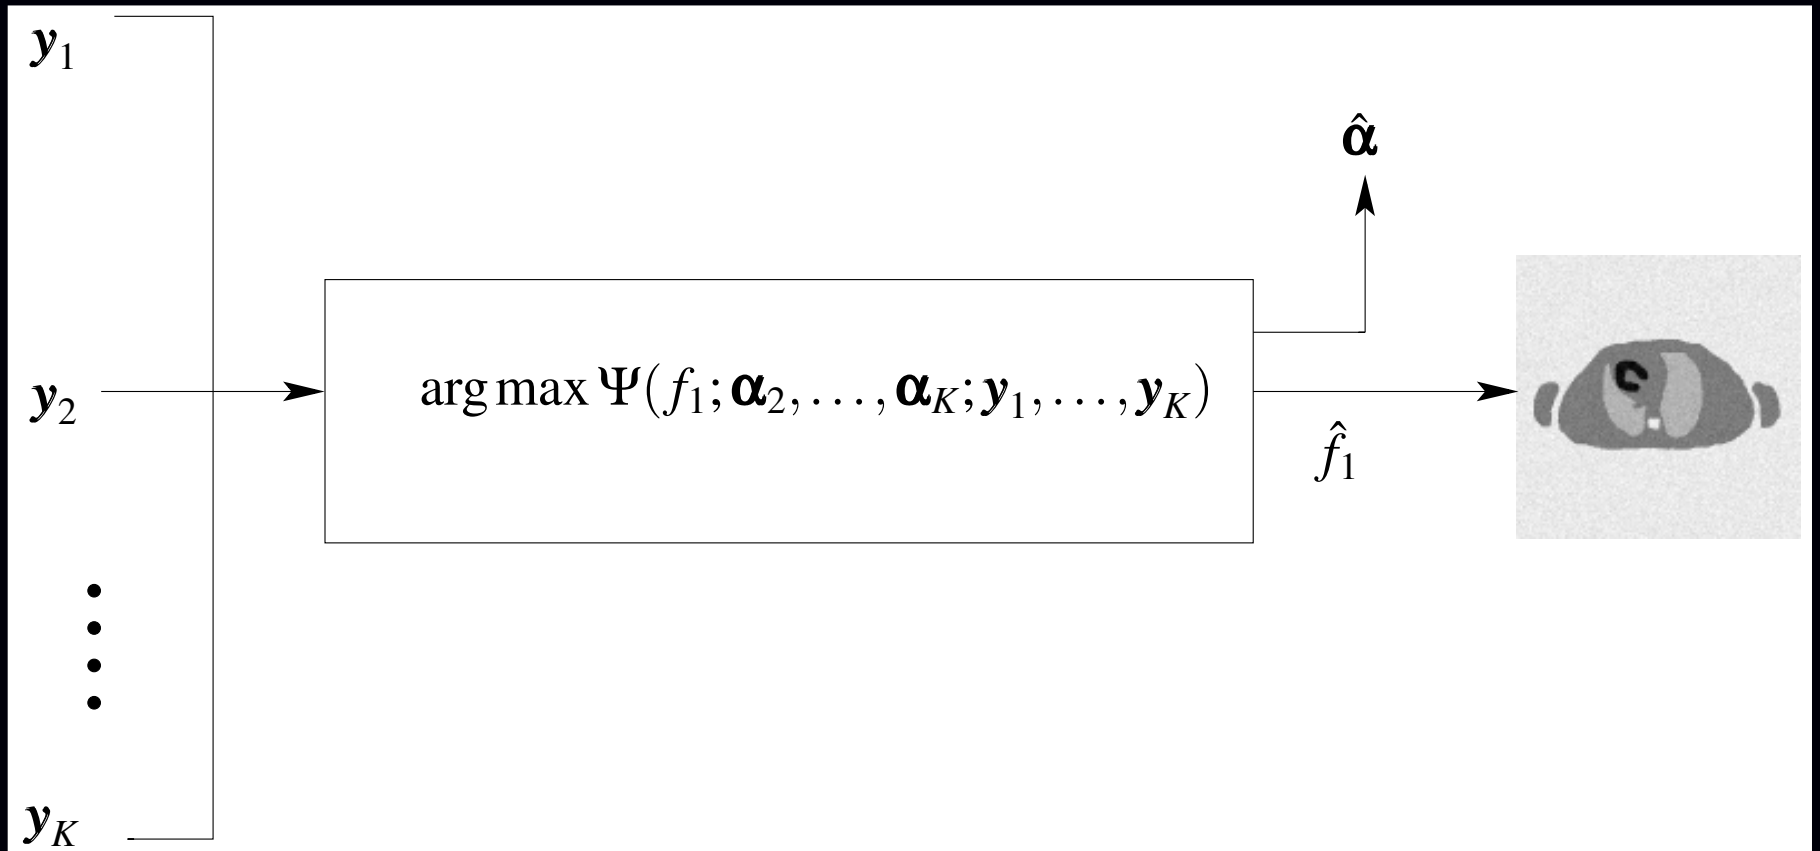
$$\Psi(f_1; \alpha_2, \dots, \alpha_K; \mathbf{y}_1, \dots, \mathbf{y}_K) = \sum_{k=1}^K \log p(\mathbf{y}_k | \mathbf{W}(\alpha_k) f_1) - \beta R(f_1)$$

Natural optimization strategy is to alternate between:

- updating image estimate \hat{f}_1 using current motion parameters,
- updating motion estimates $\{\hat{\alpha}_k\}$ using current image estimate.

Can initialize motion parameters using frame-wise method.

Joint estimation illustrated



Goal: find image estimate and motion parameters that best fit all measured data.

Motion-compensated temporal regularization

Previous joint estimation approach:

$$(\hat{f}_1, \hat{\alpha}_2, \dots, \hat{\alpha}_K) = \arg \max_{f_1, \alpha_2, \dots, \alpha_K} \Psi(f_1; \alpha_2, \dots, \alpha_K; \mathbf{y}_1, \dots, \mathbf{y}_K)$$


Alternative approach based on temporal regularization:

$$(\hat{f}_1, \dots, \hat{f}_K; \hat{\alpha}_2, \dots, \hat{\alpha}_K) = \arg \max_{f_1, \dots, f_K; \alpha_2, \dots, \alpha_K} \Psi(f_1, \dots, f_K; \alpha_2, \dots, \alpha_K; \mathbf{y}_1, \dots, \mathbf{y}_K)$$

$$\Psi(f_1, \dots, f_K; \alpha_2, \dots, \alpha_K; \mathbf{y}_1, \dots, \mathbf{y}_K) = \sum_{k=1}^K \log \mathbf{p}(\mathbf{y}_k | f_k) - \beta R(f_k) - \underbrace{\gamma \sum_{k=2}^K \|f_{k+1} - \mathbf{W}(\alpha_k) f_k\|^2}_{\text{temporal regularization with motion effects}}$$

Pro: no warp in log-likelihood. Con: more unknowns; γ choice?

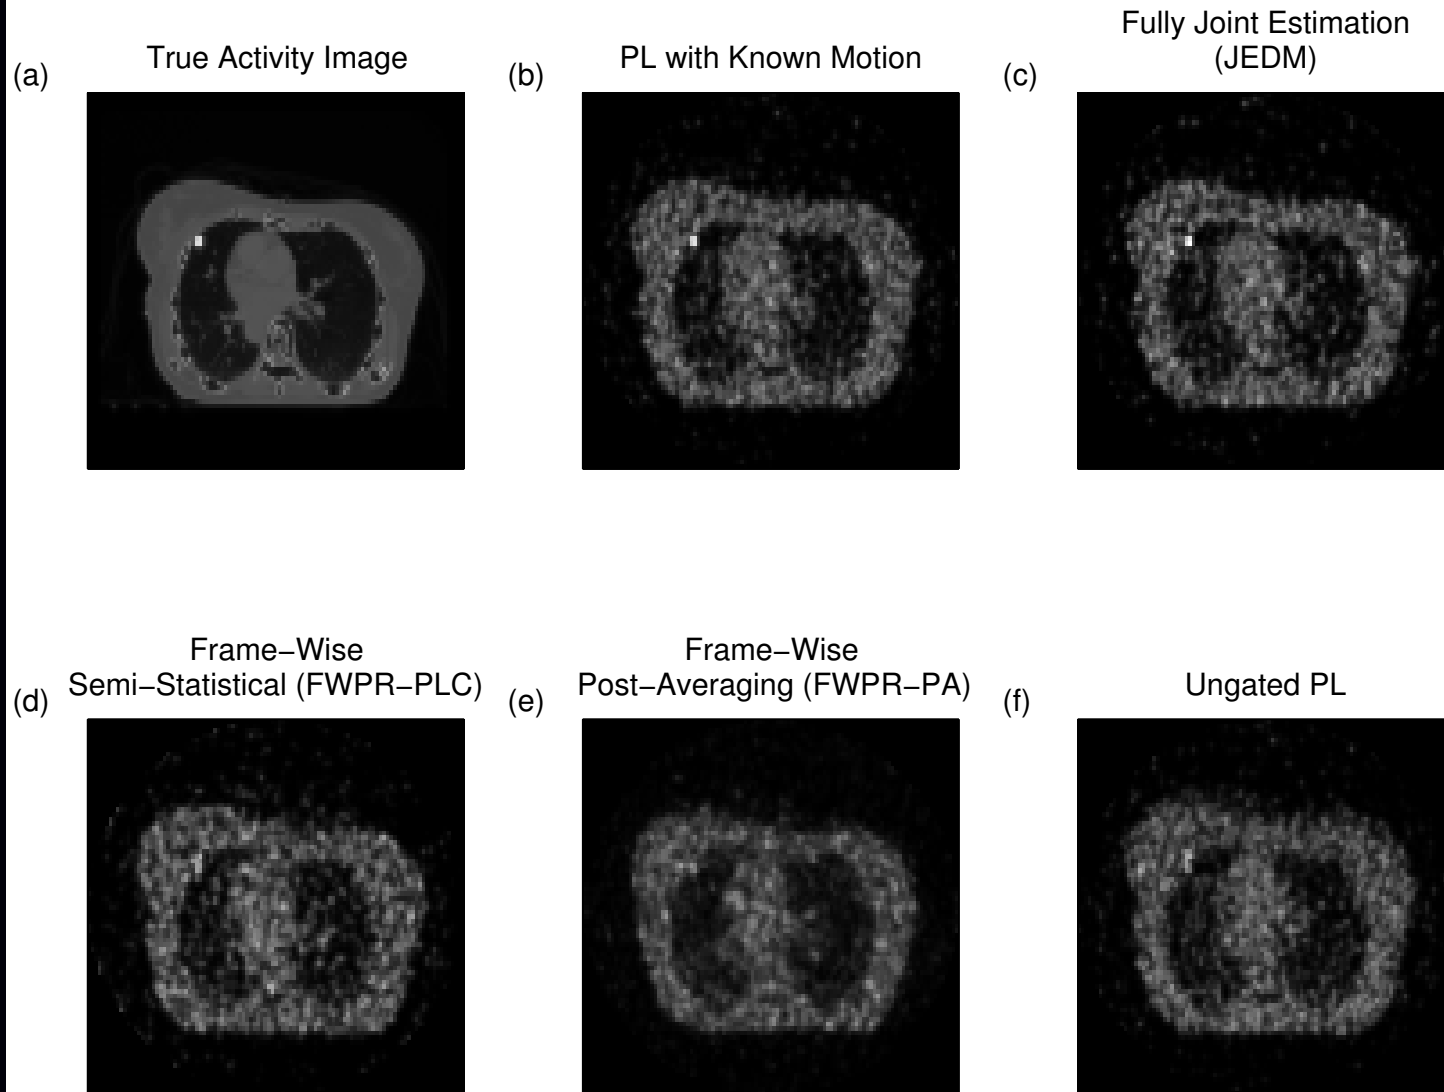
Ten-Gate 3D PET Simulation



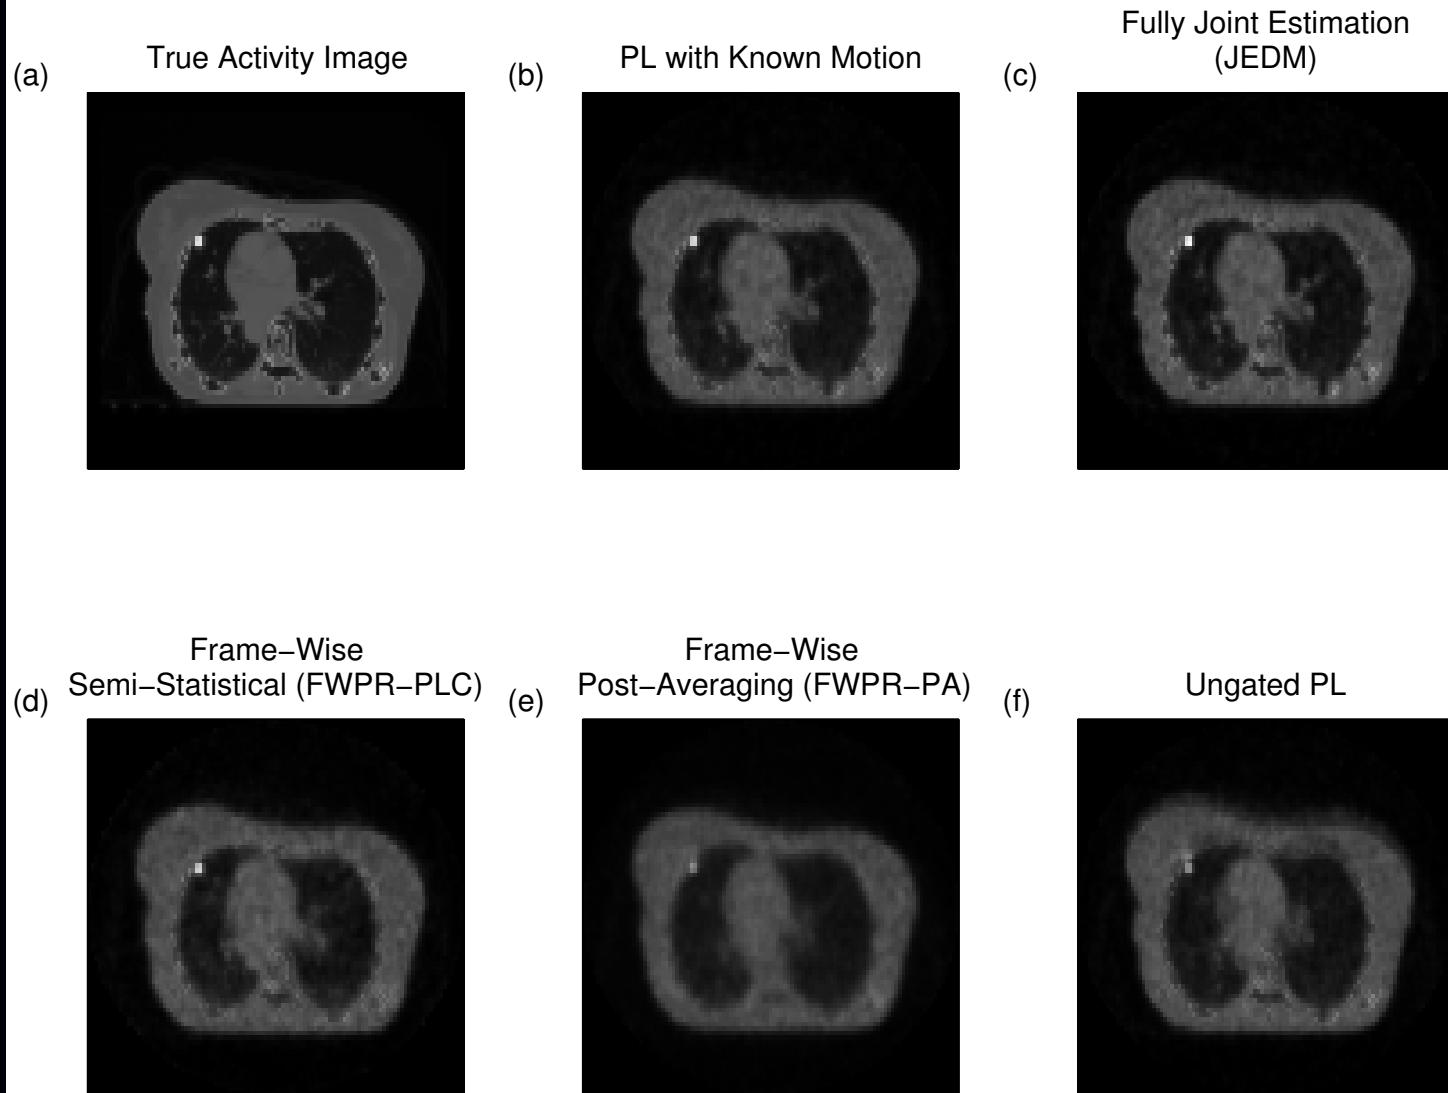
- 80K total counts/axial mm and 30% randoms (ECAT HR+), divided across 10 gates.
- Derived from 17 slices of real thorax anatomy.
- B-spline deformations (11x14x5x3 control grid), derived from helical CT scans at multiple inspirations

(Matt Jacobson, 2006 thesis)


Sample Reconstructed Images



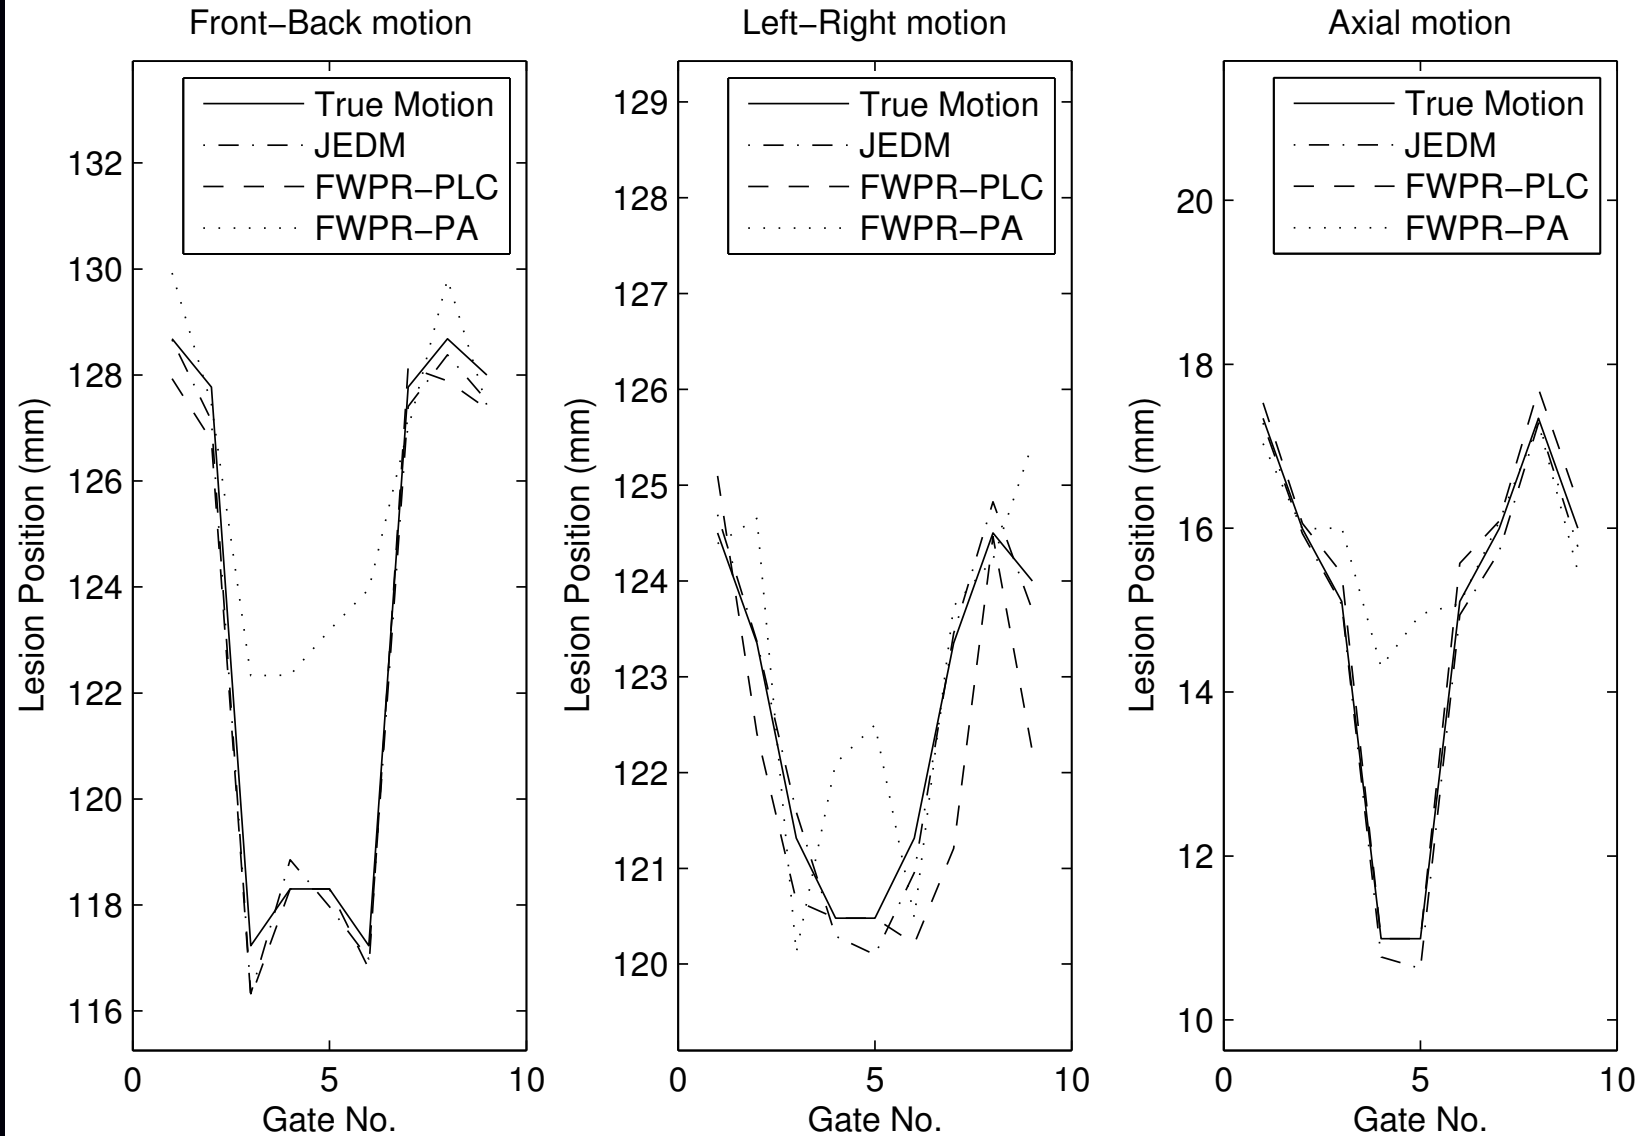
Mean Reconstructed Images



Lesion Recovery Comparison




Motion Tracking Performance



Regularization Using PET-CT Side Info.

Relax regularization strength in neighborhood of lesion.

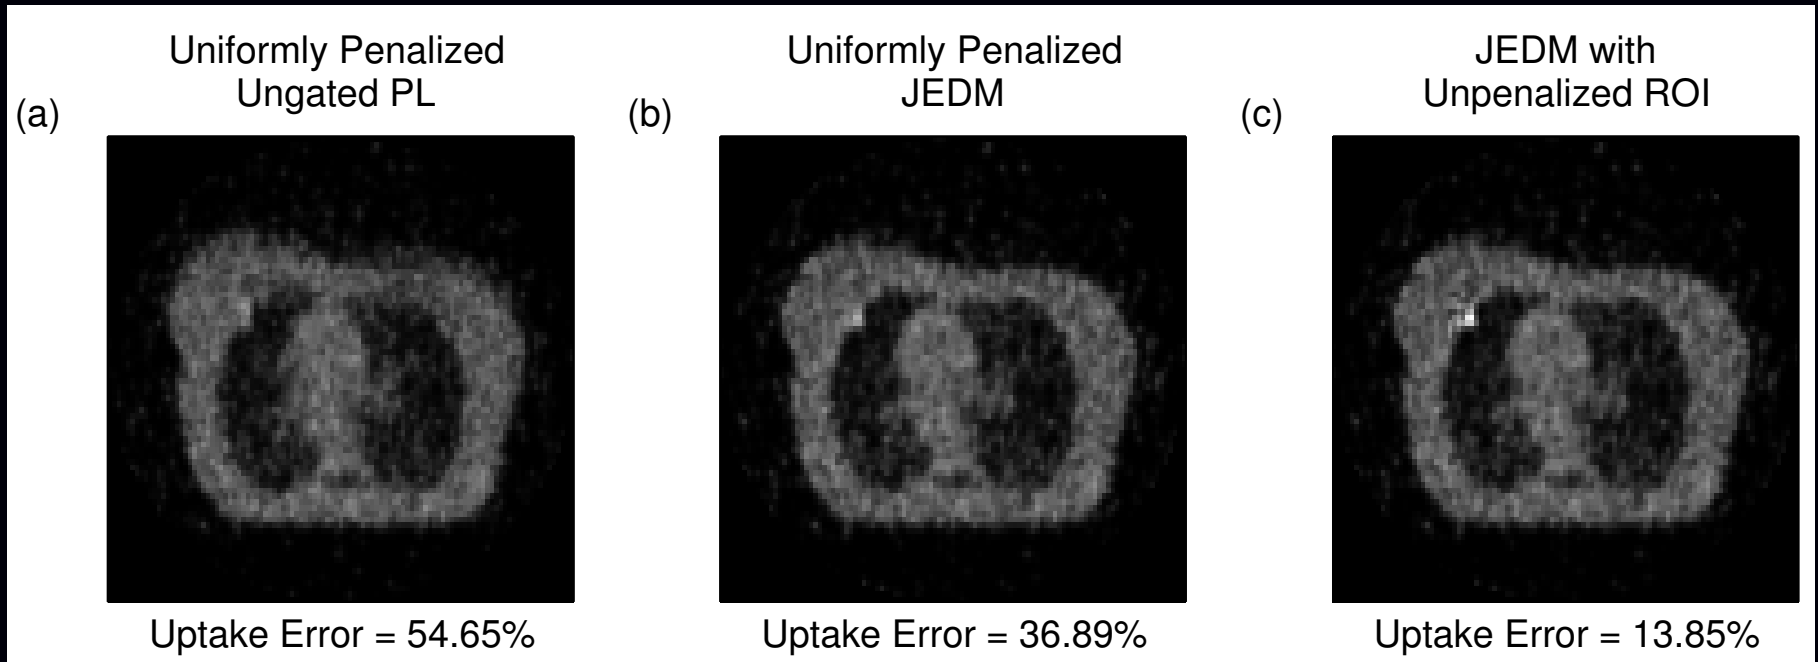
(a)



(b)



Regularization Using PET-CT Side Info. (cont'd)



Very “weak” use of boundary side information
⇒ robust to mis-registration.

Summary

- Several possible methods for motion-compensated image reconstruction
- Model-based approaches such as joint estimation have potential
- Repeated motion estimation steps necessitate simple invertibility regularizers
- More work needed on algorithms, acceleration, evaluation, ...