A Discrete Search Method for Multi-modal Non-Rigid Image Registration

Alexander Shekhovtsov **Juan D. García-Arteaga** **Tomáš Werner**

Czech Technical University in Prague
Faculty of Electrical Engineering, Department of Cybernetics
Center for Machine Perception
Czech Republic

NORDIA, 2008

Introduction

- We are interested in applying discrete optimization methods to difficult problems
  - helps avoiding local minima
  - is applicable to wider class of models (non-differentiable, discontinuous)
- Non-Rigid Multi-Modal Image Registration is a well-known carefully studied problem

- Outline
  - Review of Maximum Likelihood vs. Mutual Information criteria
  - Discretization of deformation field and optimization
  - Experiments

Non-Rigid Multi-Modal Image Registration

- Example: MRI with different exposure time (simulated [Cocosco et al.-97])

$I : T_1 \to \mathcal{C}_1$ $\qquad$ $d : T_1 \to T_2$ $\qquad$ $J : T_2 \to \mathcal{C}_2$

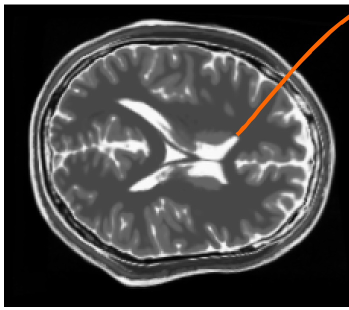

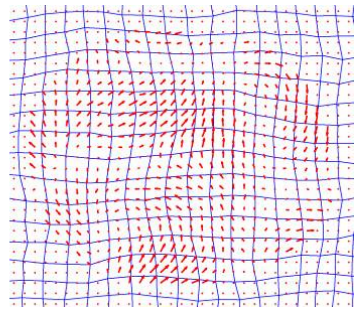

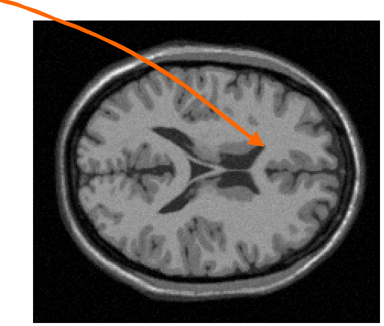

- smooth deformation field d

- non-functional signal dependence between $I(t)$ and $J(d(t))$

$p(c_1|c_2;\theta)$

0.9

0.1

Bayesian Model

- Statistical model [Roche00: Unifying maximum likelihood approaches in medical image registration]

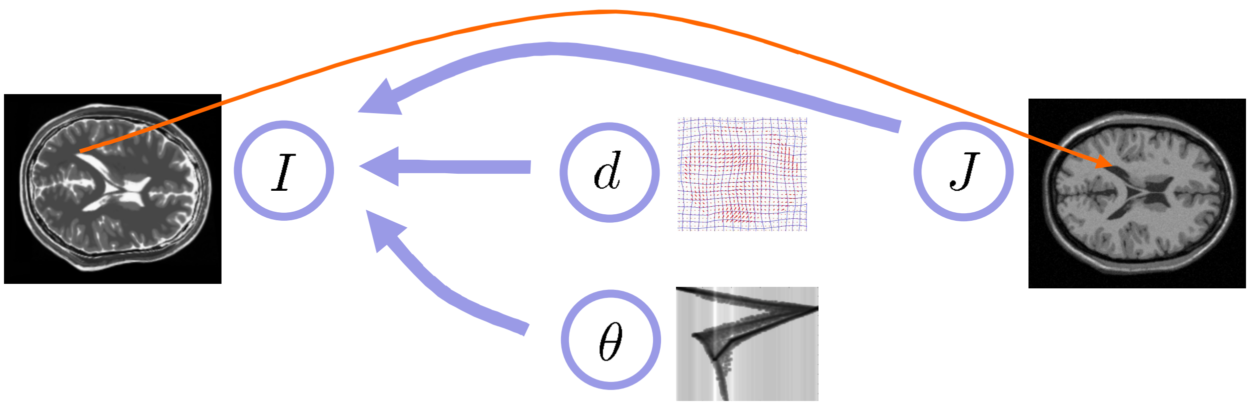

- model: $p(I, d|J; \theta) = p(I|J, d; \theta)p(d)$

- Maximum Likelihood: $$\boxed{(d^*, \theta^*) = \underset{d,\theta}{\operatorname{argmax}}\, p(I|J, d; \theta)p(d)}$$

- conditional independence of pixels:

$$p(I(t)|J, d; \theta) = p(I(t)|J(d(t)); \theta)$$

$$\prod_{t \in T_1} p(I(t)|J(d(t)); \theta)$$

Empirical Mutual Information

- Maximum Likelihood: $\boxed{\max_d \left[\underbrace{\max_\theta \log \prod_{t \in T_1} p(I(t)|J(d(t));\theta)} + \log p(d) \right]}$

- ML formulation is equivalent to maximization of an estimate of mutual information [Roche00, Kim03,...].

$\boxed{\text{const} + |T_1| \hat{I}(n_{\mathcal{C}_1,\mathcal{C}_2})}$

Let d be fixed

$n_{\mathcal{C}_1,\mathcal{C}_2} : \mathcal{C}_1 \times \mathcal{C}_2 \to \mathbb{N}$ – counts of matching colors; $p(c_1, c_2; \theta)$ – unknown;

estimate of mutual information: $\hat{I}(n_{\mathcal{C}_1,\mathcal{C}_2}) = \hat{H}(n_{\mathcal{C}_1}) + \hat{H}(n_{\mathcal{C}_2}) - \hat{H}(n_{\mathcal{C}_1,\mathcal{C}_2})$

estimate of entropy: $\hat{H}(n_{\mathcal{C}_1,\mathcal{C}_2}) = -\sum_{c_1,c_2} \frac{n_{\mathcal{C}_1,\mathcal{C}_2}(c_1,c_2)}{|T_1|} \log p(c_1, c_2; \hat{\theta})$

– uses sample mean and ML estimate $\hat{\theta} = \operatorname*{argmax}_\theta \prod_{c_1,c_2} p(c_1, c_2; \theta)^{n_{c_1,c_2}}$

cf.: $H_{\mathcal{C}_1,\mathcal{C}_2} = -\sum_{c_1,c_2} p(c_1, c_2; \theta) \log p(c_1, c_2; \theta)$

Optimization

- Maximum Likelihood: $$\max_{d,\theta} \left[\log \prod_{t \in T_1} p(I(t)|J(d(t));\theta) + \log p(d) \right]$$

Alternate between d and θ

$p(c_1, c_2; \theta)$ – Gaussian mixture with fixed number of components;
ML estimate of θ is similar to Parzen window estimate.

d – field of discrete displacements of small patches:

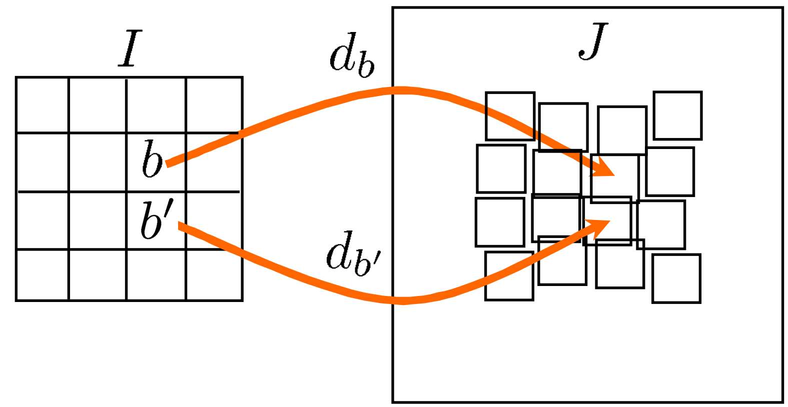

prior in the form:

$$p(d) = \frac{1}{Z} \prod_{bb' \in E} \phi_{bb'}(d_b, d_{b'})$$

ML leads to:

$$d^* = \operatorname*{argmin}_d \left[\sum_{b \in B} q_b(d_b) - \sum_{bb' \in E} \log \phi_{bb'}(d_b, d_{b'}) \right]$$

– MRF energy minimization

Optimization

MRF energy minimization: $$d^* = \operatorname*{argmin}_{d} \left[\sum_{b \in B} q_b(d_b) + \sum_{bb' \in E} q(d_b, d_{b'}) \right]$$

graph $G = (\mathcal{V}, \mathcal{E})$

labeling $d : \mathcal{V} \mapsto \mathcal{L}$

parameters q

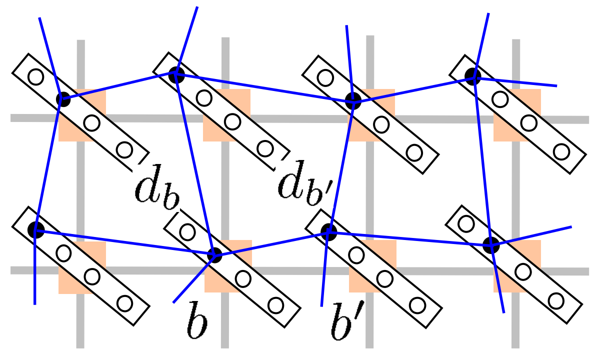

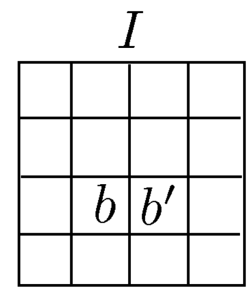

Difficult!

- Apply LP-relaxation technique [Schlesinger76, Koster98, Wainwright03, Chekuri05, Kolmogorov05...]

LP relaxation is difficult!

- Approximate image registration by MRF locally [Glocker et al. 07]
- reformulate in fewer labels [Shekhovtsov et al. CVPR'07]

Reformulate with Fewer Labels

- Standard: for each block a single variable $d_b \in W \times H$

$$E(d|q) = \sum_{b \in B} q_b(d_b) + \sum_{bb' \in E} q(d_b, d_{b'})$$

Cost is encoded on univariate term $q_b(d_b)$

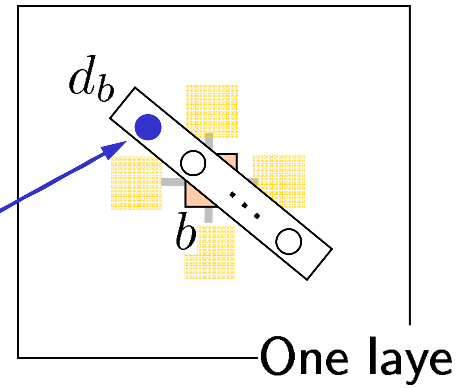

- Decomposed: variables $d_{b1} \in W$ and $d_{b2} \in H$

$$E(d|q) = \sum_{ab \in E^{\mathrm{data}} \cup E^{\mathrm{reg}}} q(d_a, d_b)$$

Cost is encoded on pairwise term $q_{b1,b2}(d_{b1}, d_{b2})$

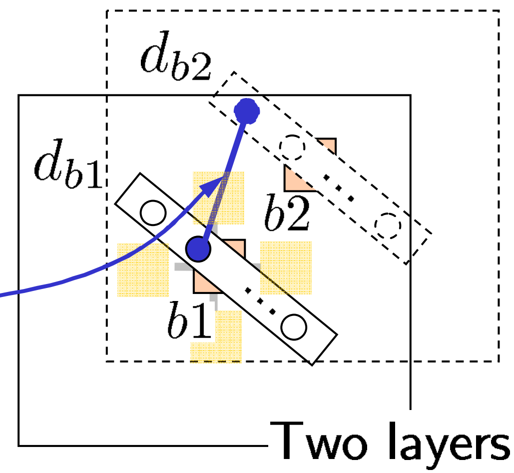

works only for separable regularizers, e.g.: $(\bar{d}_a - \bar{d}_b)^2 = (d_{a1} - d_{b1})^2 + (d_{a2} - d_{b2})^2$

Experiments

- Alternate optimization converges in several iterations

log likelihood

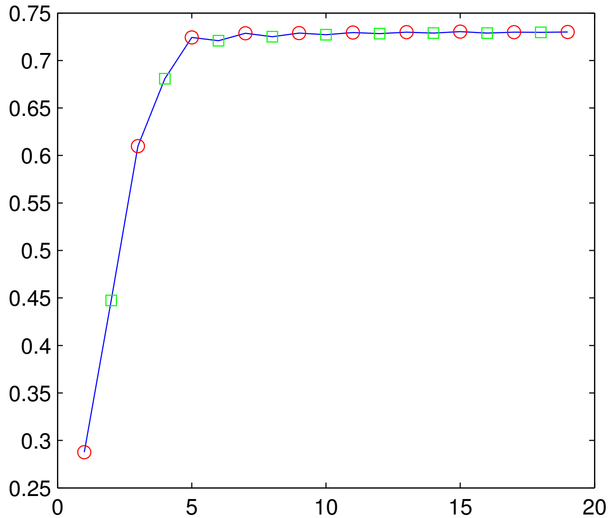

iterations

○ θ - steps

□ d - steps

suboptimal and non-monotonous in d

convergence is slower for larger noise

Concergence of $p(c_1|c_2;\theta)$

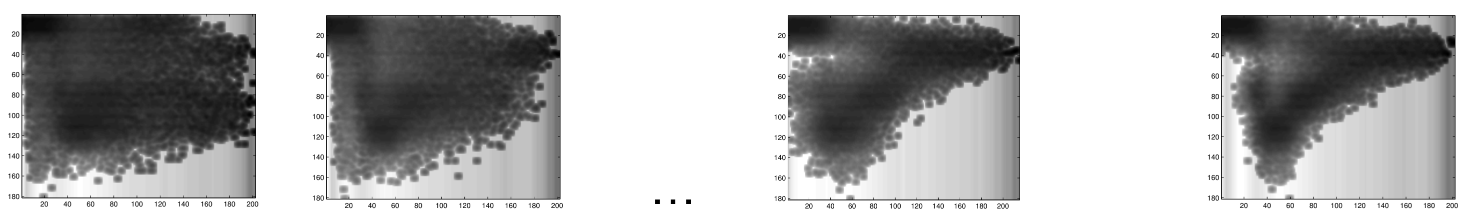

1   2   ...   convergence   from GT d

Experiments

- Generated deformations and noise

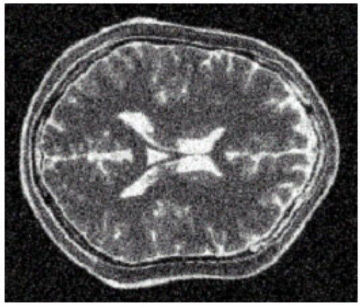

- Compare deformation filed reconstruction accuracy with Image Registration Toolkit (ITK) [Rueckert et al.-99, Schnabel et al-01.]

	$\mathcal{N}(0,\ 0.1^2)$		$\mathcal{B}(2^2) + \mathcal{N}(0,\ 0.1^2)$	
	our	ITK	our	ITK
AE Mean	0.193	0.14	0.201	0.198
AE Median	0.135	0.0908	0.147	0.135
AE Std	0.199	0.172	0.199	0.216
MOD Mean	0.562	0.419	0.591	0.577
MOD Median	0.443	0.307	0.484	0.453
MOD Std	0.441	0.382	0.433	0.458

AE – Angular Error, [deg].

MOD – Magnitude of Difference, [px].

ITK is slightly better and works in about the same speed.

Discussion

- Advantages:
  1) would work for variety of models $p(c_1|c_2;\theta)$
  2) can potentially represent more complex models (e.g. deformations with discontinuities, include segmentation, etc.)
  3) can potentially help avoiding local minima

- Disadvantages:
  1) non-monotonous d-optimization steps
  2) regularization is difficult to model with pairwise MRF

Thank You