Institute of Biomedical Engineering
Department of Engineering

Non-rigid image registration through efficient discrete optimization

Mattias P. Heinrich, Mark Jenkinson, J. Michael Brady and Julia A. Schnabel
University of Oxford

Medical Image Understanding and Analysis 2011

2011, July 14th

Outline

- Introduction and Motivation
- Discrete optimization on Markov Random Field (MRF)
- Efficient single-resolution implementation of belief propagation
- Comparison to continuous and multi-resolution discrete optimization
- Results for 2D multimodal data and clinical 3D CT images
- Conclusion

Image Registration Overview

- Transformation Model (Regularization)
  - Rigid, affine, …
  - B-spline (bending energy), radial basis functions, …
  - non-parametric (elastic, diffusive, Demon's)
- Similarity metric
  - sum of absolute differences (SAD), sum of squared differences (SSD) (single-modal)
  - normalized cross-correlation (NCC), mutual information (MI) (multi-modal)
- Optimization methods
  - Continuous, gradient-based: Steepest descent, Gauss-Newton
  - Discrete optimization

Motivation for discrete optimization

- Requirements for a good optimization method?
  - Independent from choice of transformation model / regularization
  - Independent from choice of similarity metric
  - Global optimum of cost function can be found
  - Low computational complexity

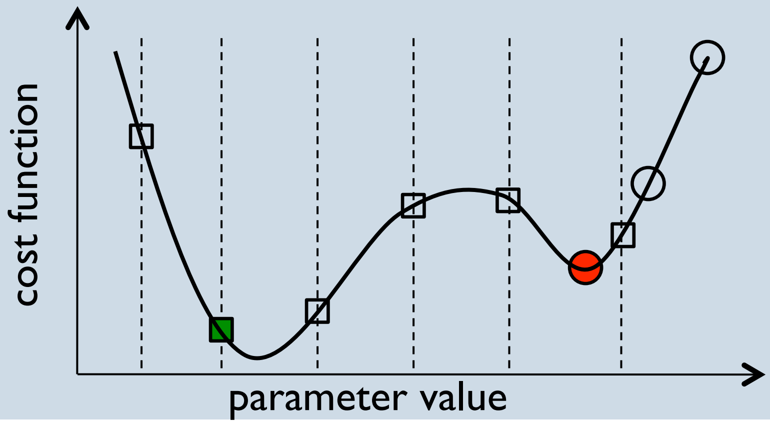

- Limitations of gradient-based continuous optimization ○
  - Regularization term and similarity metric must be differentiable functions
  - usually fast and with subpixel accuracy
  - only local optimum can be found (multi-resolution helps)
- Properties of discrete optimization □
  - Any function can be used: non-differentiable, modulus, truncated, …
  - no need for multi-resolution, warping, interpolation, …
  - Approximate <u>global</u> optimum can be reached (no subpixel accuracy)

Discrete optimisation

- Formulate problem on graph (MRF)
  - nodes correspond to pixels p, edges connect nodes
  - each node has a set of labels
  - labels represent displacements
    - which have associated probabilities
- Labeling using message passing
  - we want to select labeling f_p with highest probability
  - cost function $S_p(f_p) + \alpha R(f_p, f_q)$
    - similarity metric $S_p(f_p)$ for each node independently
    - regularization term $R(f_p, f_q)$ depending on connected nodes
  - several possible schemes to pass information between nodes
    - belief propagation, tree-reweighted message passing, iterative graph-cut

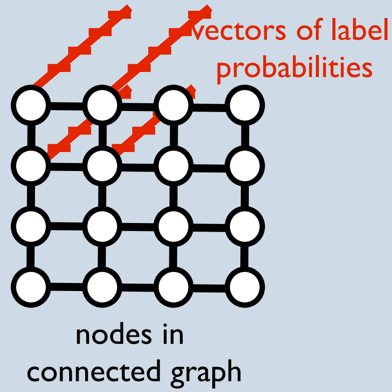

- <u>Problem</u>: Complexity is proportional to number of labels

  (large number of labels needed for 3D medical registration) we will come to this later …

Similarity cost: discrete sampling

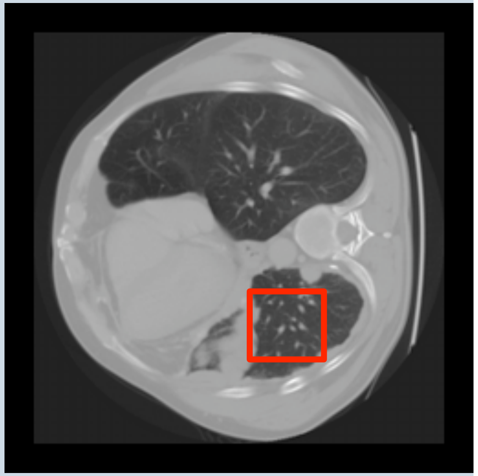

- Similarity metric: absolute differences
  - given a pixel p and the label space f_p for images I and J we get:
    - $S_p(f_p)=|I(p)-J(p+f_p)|$
  - label space is f_p ={$\triangle$x=±7px,$\triangle$y±7px} (2D displacements)

  - repeat sampling for all nodes and store label costs (negative log probabilities) in a 1D vector for each node

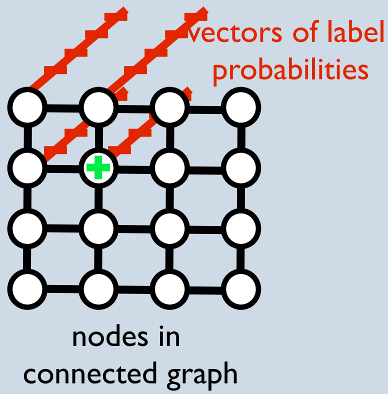

  - <u>next</u>: probability vectors will be updated during message passing

Regularization: inference on graph

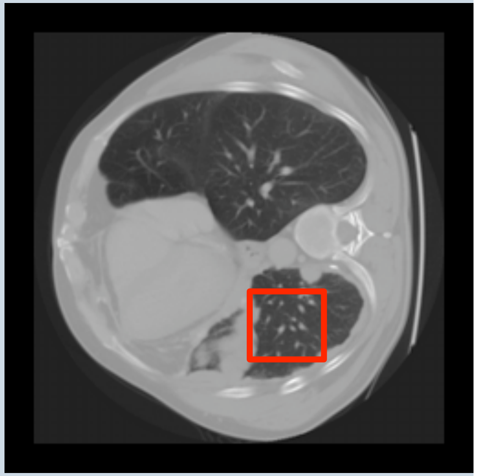

- discrete sampling of similarity/data term $S_p(f_p)$
  - local minimum may not be correct displacement
- Inference on graph using message passing
  - calculate messages for nodes and pass them to all neighbouring nodes (belief propagation)
  - message updates include regularization term $R(f_p,f_q)$

$$m^t_{p\to q}(f_q) = min_{f_p}\left(S_p(f_p)+R(f_p,f_q)+\sum_{s\in N(p)\setminus q} m^{t-1}_{s\to p}(f_p)\right)$$

message updates

incoming and outgoing messages

one message per label

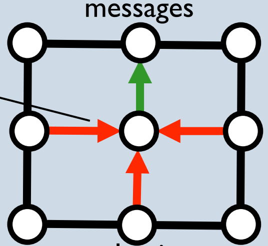

nodes in connected graph (hidden probability vectors)

Deformation field (label with highest probability)

without regularization

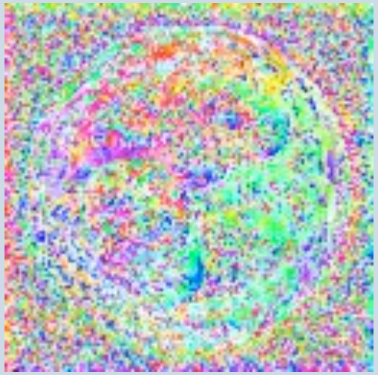

with regularization

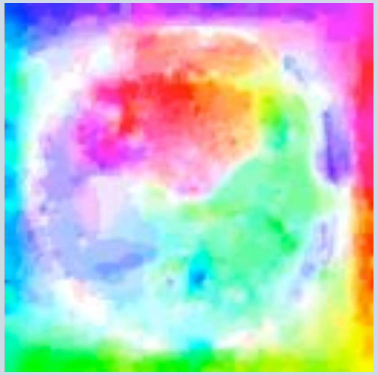

angle representation

Efficient belief propagation I

- Hierarchical BP: efficient propagation using levels
  - group blocks of nodes together (coarse-to-fine strategy)
    - – similarity costs are calculated pixel-wise and aggregated
  - pass messages only between groups
  - faster convergence (less iterations)
  - *Felzenszwalb, "Efficient belief propagation" IJCV2006*

only one label vector per group

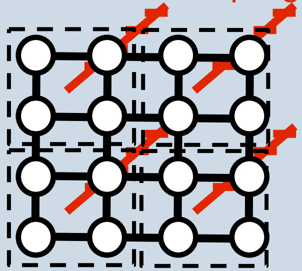

groups of nodes in connected graph

Convergence of standard and hierarchical BP

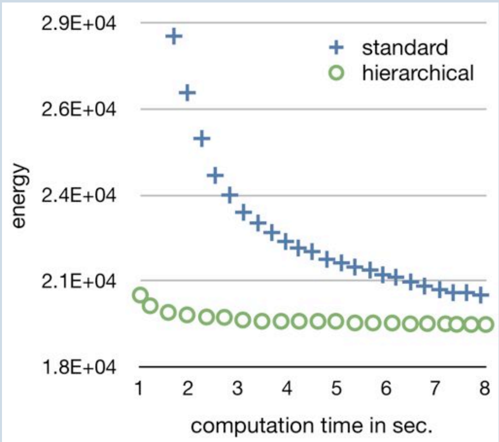

(plot begins after first iteration)

Problem: label space stays the same size for finest level

typical label space for 3D medical registration is 40^3=64'000

memory requirement for messages would be 10 TBytes

Efficient belief propagation II

- Constant space BP (csbp): hierarchical label reduction
  - reduce number of labels to *k* after initial similarity term computation
  - divide number of possible labels by 2 after all iterations for one level
  - pick labels with smallest cost (of message vector)
  - allows to keep multiple hypotheses in highest resolution
    - better accuracy than multi-resolution methods
  - *Yang, "A constant space belief propagation …" CVPR2010*

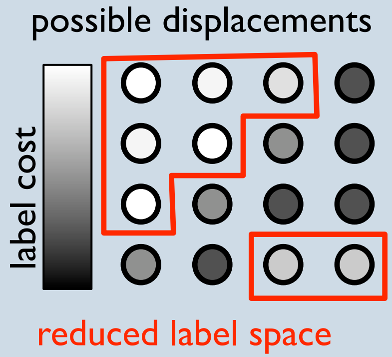

- Complexity comparison with hierarchical BP

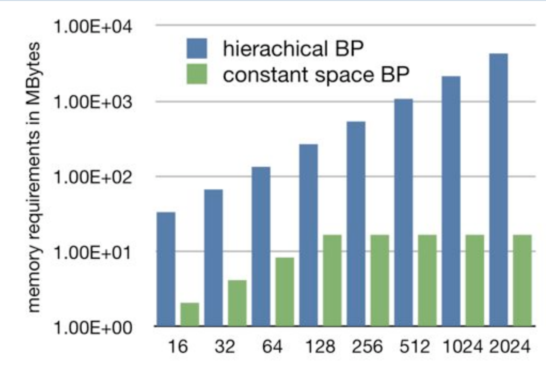

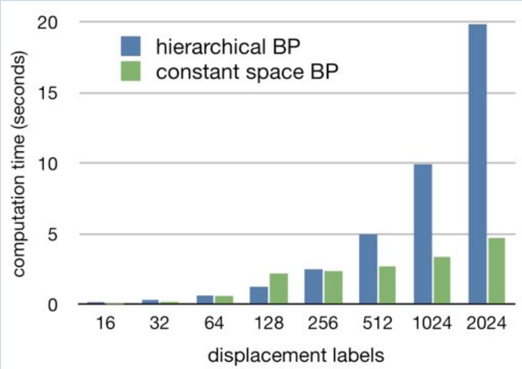

Example for 2D image with 256x256 pixels and increasing number of labels (k=128)

- Constant memory requirements and complexity for all levels

Comparison methods

- Gauss-Newton continuous optimization with diffusion regularization

  [4] Heinrich et al. in MICCAI 2011

  - multi-resolution, multiple warps
  - finite differences to approximate non-trivial cost functions
- drop: discrete optimization with B-splines

  [3] Glocker et al. in Med. Image Anal. 2008

  - multi-resolution, iterative refinements
  - limited deformation space (control points)
  - graph-cuts and linear programming
- csbp: proposed method
  - single-resolution, hierarchical coarse-to-fine grouping for faster message passing
  - hierarchical reduction of label space, constant complexity
  - no resampling or interpolation used

Non-local shape descriptor: A new similarity metric for deformable multi-modal registration

Mattias P. Heinrich[1,2,*], Mark Jenkinson[2], Manav Bhushan[1,2], Tahreema Matin[3], Fergus V. Gleeson[3], Sir J. Michael Brady[4], and Julia A. Schnabel[1]

[1] Institute of Biomedical Engineering, University of Oxford, UK
[2] Oxford University Centre for Functional MRI of the Brain, UK
[3] Department of Radiology Churchill Hospital, Oxford, UK
[4] Department of Radiation Oncology and Biology, University of Oxford, UK
mattias.heinrich@eng.ox.ac.uk, julia.schnabel@eng.ox.ac.uk
http://users.ox.ac.uk/~shil3388

ELSEVIER

Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

Dense image registration through MRFs and efficient linear programming ☆

Ben Glocker [a,b,*], Nikos Komodakis [a,c], Georgios Tziritas [c], Nassir Navab [b], Nikos Paragios [a,d]

HEINRICH et al.: NON-RIGID REGISTRATION THROUGH DISCRETE OPTIMIZATION 1

Non-rigid image registration through efficient discrete optimization

Mattias P. Heinrich[1]
mattias.heinrich@eng.ox.ac.uk
Mark Jenkinson[2]
Sir J. Michael Brady[3]
Julia A. Schnabel[1]

[1] Institute of Biomedical Engineering
Department of Engineering Science
University of Oxford, UK
[2] Oxford Centre of FMRI of the Brain
University of Oxford, UK
[3] Department of Radiation Oncology
and Biology, University of Oxford, UK

2D multimodal experiments

- 2 different colour channels of cryosection of *Visible Human* (intrinsically aligned)
  - synthetic deformation on source image (using B-spline grid) TRE: A 3.1mm B 6.4mm

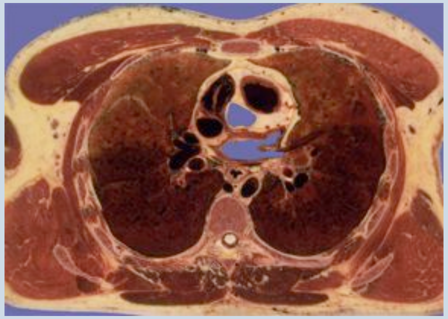

original colour scan

target image (green)

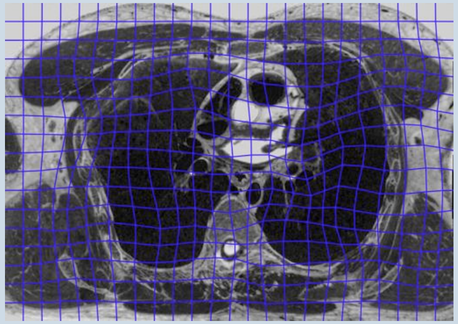

deformed source image (blue)

- Subpixel accuracy is obtained by applying a smoothing spline
- Quantitative results measured as target registration error (TRE) in mm

		before reg.	continuous		drop		csbp+bspline	
		TRE	TRE	time	TRE	time	TRE	time
singlemodal	SSD A	3.14 (2.13)	0.18 (0.22)	11	0.14 (0.21)	8	0.15 (0.17)	22.4
	SSD B	6.44 (4.51)	1.08 (2.05)	11	0.15 (0.25)	8.3	0.17 (0.18)	22.4
multimodal	NCC A	3.18 (2.37)	0.62 (0.79)	25.4	0.45 (0.72)	9	0.22 (0.20)	47.5
	NCC B	6.44 (4.51)	0.92 (1.19)	25.4	1.42 (2.28)	9.9	0.52 (1.14)	47.5

Real clinical example, 4D-CT registration

- Breathing cycle CT
  - 5 scans[1]
  - 5 phases per scan
- Challenges in lung images
  - varying contrast (air density)
  - large deformation of small features
  - discontinuous motion
- Registrations are performed between inhale and exhale

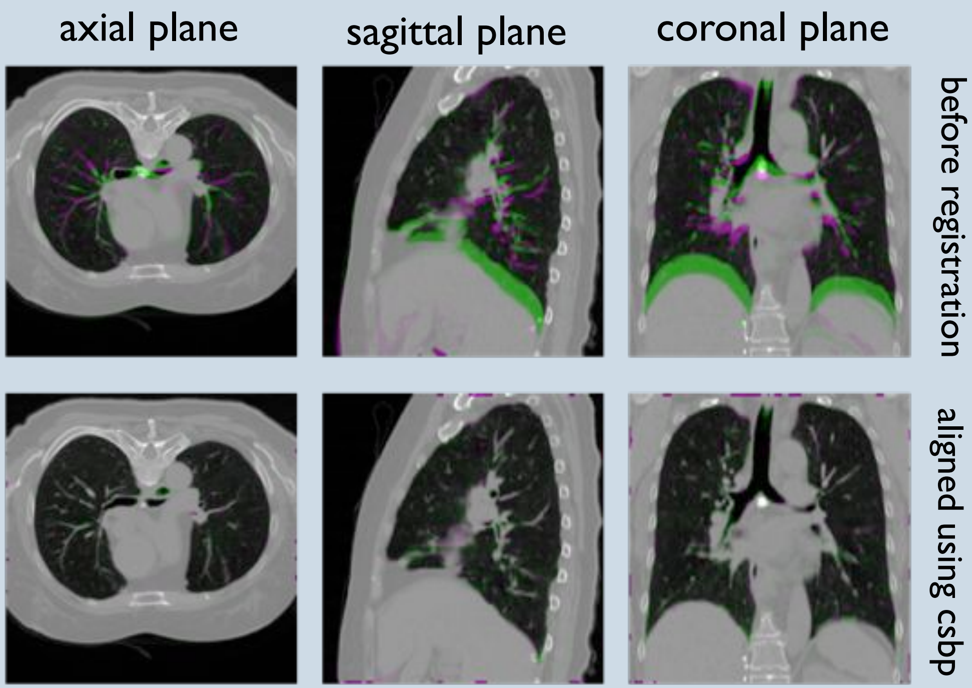

target image (magenta), moving image (green)

[1] Dataset is publicly available at dir-lab.com

Real clinical example, 4D-CT registration

- 300 manually selected anatomical landmarks
  - densely distributed over lung
- Target registration error (TRE) is lowest for proposed method (csbp)
- most improvement for large deformations (4,5)

TRE (std) in mm	before reg.	continuous	drop	4DLTMd	csbp
Case 1	4.01 (2.91)	1.020 (0.499)	1.000 (0.519)	0.97 (1.02)	0.829 (0.944)
Case 2	4.65 (4.09)	1.177 (0.828)	1.048 (0.603)	0.86 (1.08)	0.842 (0.953)
Case 3	6.73 (4.21)	1.950 (1.794)	1.378 (1.018)	1.01 (1.17)	0.995 (1.058)
Case 4	9.42 (4.81)	1.895 (1.884)	1.578 (1.389)	1.40 (1.57)	1.271 (1.241)
Case 5	7.10 (5.15)	2.504 (2.491)	1.858 (2.180)	1.67 (1.79)	1.256 (1.520)

- 4DLTMd has been so far the best performing algorithm for this dataset

  Castillo, „Four dimensional deformable image registration using trajectory modeling“ Phys Med Biol, 2010.

Error plot of TRE (size of circles)

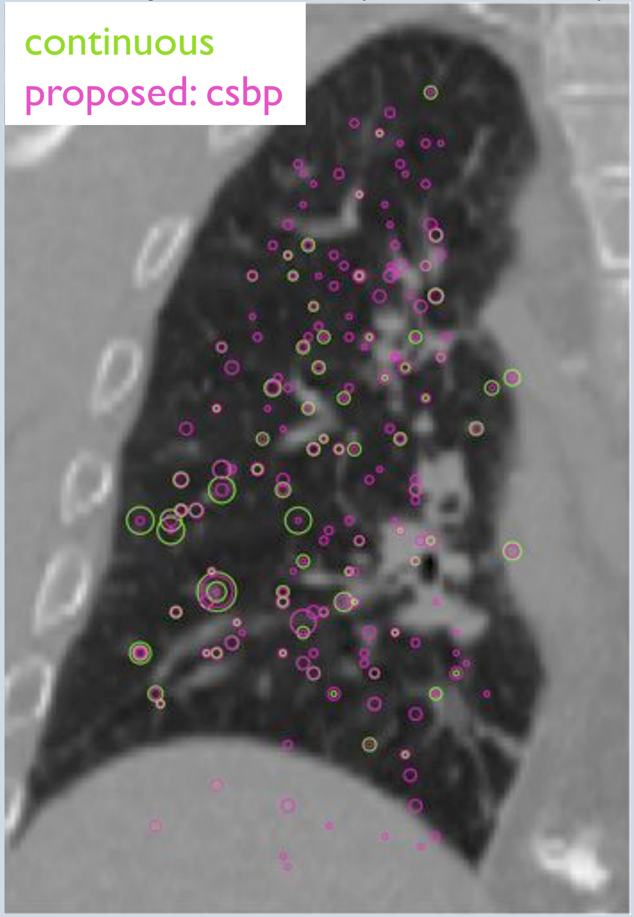

Conclusion

- Advantages over continuous optimization
  - independent of similarity metric (preferably point-wise metrics)
  - independent of regularization term (we only covered 1^{st} order fields)
  - robustly finds approximately global minimum
- Advantages over multi-resolution discrete optimization *(drop)*
  - uses only one image resolution (small features are preserved)
  - removes need of warping schemes and interpolation (partial volume effects)
  - label space reduction keeps complexity constant over levels
- Experimental results
  - increased accuracy for 2D multimodal data and clinical 3D CT images

- Future work
  - symmetric, inverse-consistent formulation
  - more validation on clinical multimodal scans

Conclusion

- Acknowledgements
  - We would like to thank EPSRC and Cancer Research UK for funding this work within the Oxford Cancer Imaging Centre.
  - J.A.S. acknowledges funding from EPSRC EP/H050892/1.

- Questions?