Local All-Pass Image Registration

Thierry Blu

Department of Electronic Engineering
The Chinese University of Hong Kong

Common work with Chris Gilliam (RMIT) and Xinxin Zhang (CUHK)

Geometric deformations

A geometric deformation is a 2D mapping

$$\mathbf{r} \mapsto \mathbf{D}(\mathbf{r}) = \mathbf{r} + \mathbf{u}(\mathbf{r})$$

where $\mathbf{u}(\mathbf{r})$ is the displacement induced by the transformation.

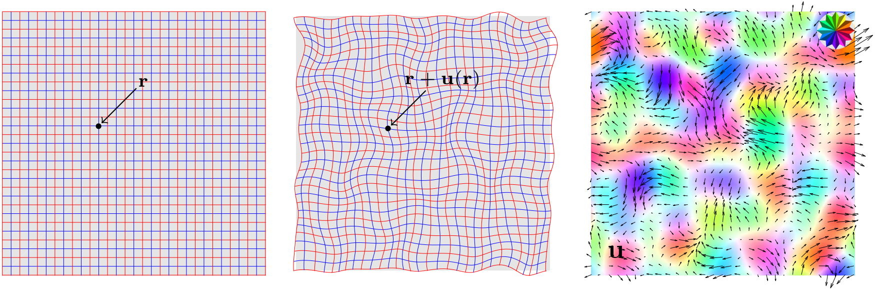

Image deformations

Given a "source" image $I_s(\mathbf{r})$, an image deformation transforms this image into a "target" image $I_t(\mathbf{r})$ according to

- $I_t(\mathbf{r}) = I_s\big(\mathbf{r} + \mathbf{u}(\mathbf{r})\big)$ (brightness consistency, monomodal case)
- $I_t(\mathbf{r}) = \mathscr{F}\Big\{I_s\big(\mathbf{r} + \mathbf{u}(\mathbf{r})\big)\Big\}$ (intensity changes, multimodal case)

Image registration

Given I_s and I_t find $\mathbf{u}(\mathbf{r})$ and, if applicable, find $\mathscr{F}$.

Applications

- Medical applications

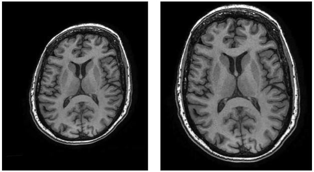

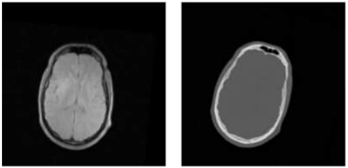

- Remote sensing

- Optical flow estimation (computer vision, tracking etc.)

Registration algorithms

Typical settings

- Global parametric registration (AECC: Evangelidis, PAMI 2008)
- Elastic (local) registration (Demons: Thirion, Medical Image Analysis 1998; Lombaert, Neuroimage 2009)
- Landmark-based registration (Rohr, TMI 2001)

Typical optimization criterion

- Mean-square error (monomodal)
- Mutual information (multimodal)

Idea 1: Shifting is all-pass filtering

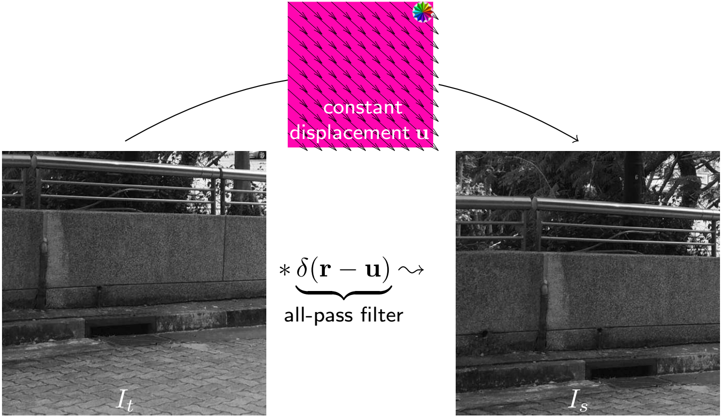

Geometric shifting is equivalent to analytic convolution.

Idea 2: Representation of all-pass filters

All-pass filters $h(\mathbf{r})$ are characterized by $|H(\boldsymbol{\omega})| = 1$. Any real all-pass filter can be expressed as

$$H(\boldsymbol{\omega}) = \frac{P(\boldsymbol{\omega})}{P(-\boldsymbol{\omega})}$$

where $P(\boldsymbol{\omega})$ is the Fourier transform of some real "forward" filter $p(\mathbf{r})$.

The brightness consistency equation $I_t(\mathbf{r}) = I_s(\mathbf{r} + \mathbf{u})$ can then be expressed as

$$\underbrace{p(-\mathbf{r})}_{\text{backward}} * I_t(\mathbf{r}) = \underbrace{p(\mathbf{r})}_{\text{forward}} * I_s(\mathbf{r})$$

Idea 3: Approximation and minimization

The forward filter can be approximated using *three local elementary filters*

$$p(\mathbf{r}) = g(\mathbf{r}) + a\frac{\partial g(\mathbf{r})}{\partial x} + b\frac{\partial g(\mathbf{r})}{\partial y}$$

where $g(\mathbf{r}) = \exp\left(-\frac{\|\mathbf{r}\|^2}{2\sigma^2}\right)$.

Minimizing for the coefficients a and b the MSE

$$\|p(-\mathbf{r}) * I_t(\mathbf{r}) - p(\mathbf{r}) * I_s(\mathbf{r})\|^2$$

gives $\mathbf{u} \approx (2a, 2b)$ with excellent accuracy.

NOTE: solution of a linear system of equations, fast.

Local displacement estimation

Local shift assumption

The elastic displacement can be approximated locally by a shift
$\rightsquigarrow$ "local all-pass" filter.

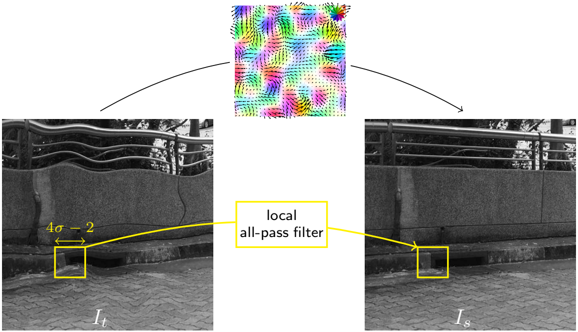

At central pixel: Estimate local all-pass filter, then extract motion from the filter.

Multiscale LAP

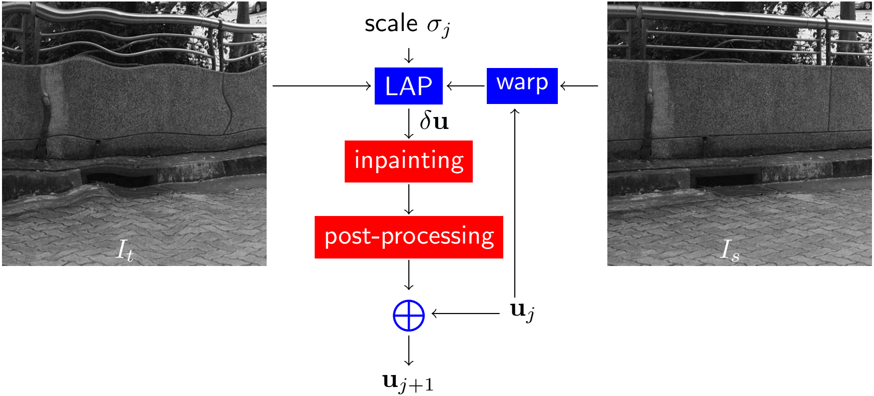

- Estimation of the displacement at different scales $\rightsquigarrow$ change σ_j
- Estimation of erroneous displacements (e.g., too large, boundaries) $\rightsquigarrow$ inpainting
- smoothing

Synthetic result

before alignment

after alignment

Median error $= 0.003$ pixel, mean error $= 0.05$ pixel, computation time 3s (512×512 image).

Real result

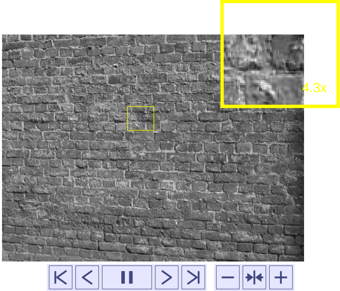

before alignment

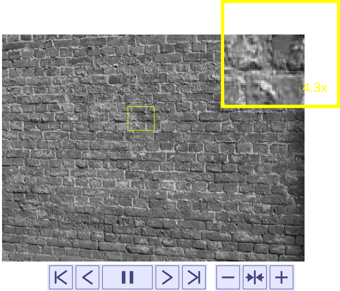

after alignment

ERROR COMPARISON FOR THE PF-LAP AND A SELECTION OF IMAGE REGISTRATION ALGORITHMS ON IMAGES FROM THE OXFORD AFFINE DATASET [61]

	Bikes			Leuven			Wall		
	E_{Med}	E_{Mean}	Time	E_{Med}	E_{Mean}	Time	E_{Med}	E_{Mean}	Time
PF-LAP	0.223	**0.292**	12.10	**0.171**	**0.217**	**10.41**	**0.506**	**0.866**	**11.80**
Demons [21]	0.458	0.702	15.09	0.243	0.395	11.83	1.383	21.66	13.16
bUnwarpJ [24]	**0.220**	0.308	**8.54**	0.189	0.229	11.59	10.57	30.28	30.69
MIRT [26]	0.726	4.228	129.9	0.363	0.813	87.76	0.571	1.874	110.4

* Bold values indicate the best results

Real result

source and target

AECC 12.61s

MIRT 82.32s

Demons 8.75s

LAP 18.82s

LAP with parametric fitting 18.58s

Retinal image registration

source and target

AECC 25.36s

MIRT 219.05s

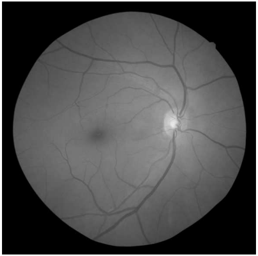

Multispectral registration—green and NIR

source and target

AECC

MIRT

Demons

LAP

LAP with parametric fitting

Extensions: 3D MRI

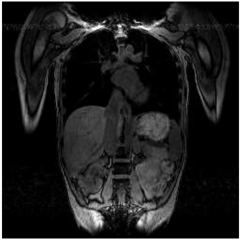

before alignment

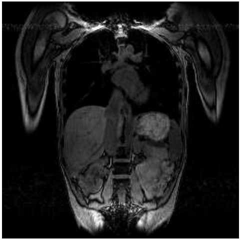

after alignment

	Lung Segmentation[3]		Image Registration	Computation
	Dice-Coefficients[4]	Cross-Correlation	PSNR (dB)	Time
3D LAP	**0.90** (0.01)	**0.97** (0.01)	**39.9**	**36.3**
Elastix[1]	0.87 (0.02)	0.95 (0.02)	37.3	61.6
Demons[2]	0.73 (0.05)	0.92 (0.02)	38.2	434.6

*Image Size = 256 by 256 by 72 voxels

Extensions: video

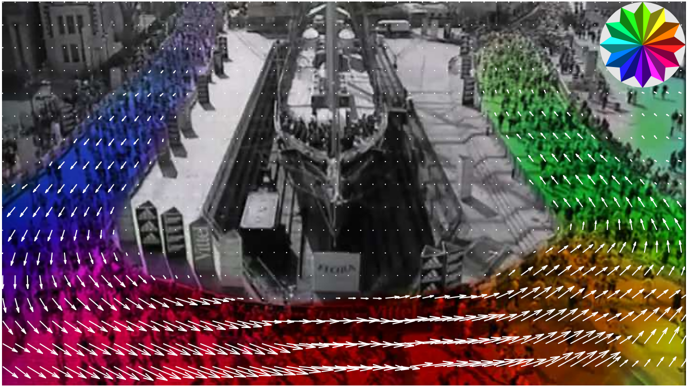

stationary motion overlayed

Conclusion

- Estimated motion using Local All-Pass Filters
  - Shifting by a constant displacement $\Longrightarrow$ All-pass filtering
  - Assume motion is locally constant $\Longrightarrow$ Local All-Pass Filters
  - Fast and efficient implementation
- Applied to Biomedical Images
  - Demonstrated accuracy using synthetic images
  - Accurate removal of respiratory motion from MRI data
  - Motivate the idea of analysing scene dynamics

Main paper: Gilliam, C. & Blu, T.,"*Local All-Pass Geometric Deformations*", IEEE Transactions on Image Processing, Vol. 27 (2), pp. 1010-1025, February 2018.