Efficient Space-Variant Blind Deconvolution

Stefan Harmeling

joint work with

Michael Hirsch, Suvrit Sra, Christian Schuler,
and Bernhard Schölkopf

Max Planck Institute for Biological Cybernetics, Tübingen, Germany

MPI FOR BIOLOGICAL CYBERNETICS

NIPS Workshop on
Numerical Mathematics Challenges in Machine Learning (NUMML)
Whistler, 11. December 2010

What is Space-Variant Convolution?

Space-Invariant Convolution:

$$y_i = \sum_{j=0}^{k-1} a_j x_{i-j}$$

- a is called space-invariant *Point Spread Function (PSF)*.
- a does not depend on the location in y.

Space-Invariant Convolution:

$$y_i = \sum_{j=0}^{k-1} a_j x_{i-j}$$

- a is called space-invariant *Point Spread Function (PSF)*.
- a does not depend on the location in y.

Space-Variant Convolution:

$$y_i = \sum_{j=0}^{k-1} b_{ij} x_{i-j}$$

- b is called space-variant *Point Spread Function (PSF)*.
- b does depend on the location in image y.

Space-Invariant Convolution:

$$y_i = \sum_{j=0}^{k-1} a_j x_{i-j}$$

- a is called space-invariant *Point Spread Function (PSF)*.
- a does not depend on the location in y.

Space-Variant Convolution:

$$y_i = \sum_{j=0}^{k-1} b_{ij} x_{i-j}$$

- b is called space-variant *Point Spread Function (PSF)*.
- b does depend on the location in image y.

For an observed image y:

- x is the true underlying image to be recovered.
- The PSF fully describes how x is transformed.

Examples of Space-Variant Convolutions

Example: Camera Shake

Example: Camera Shake

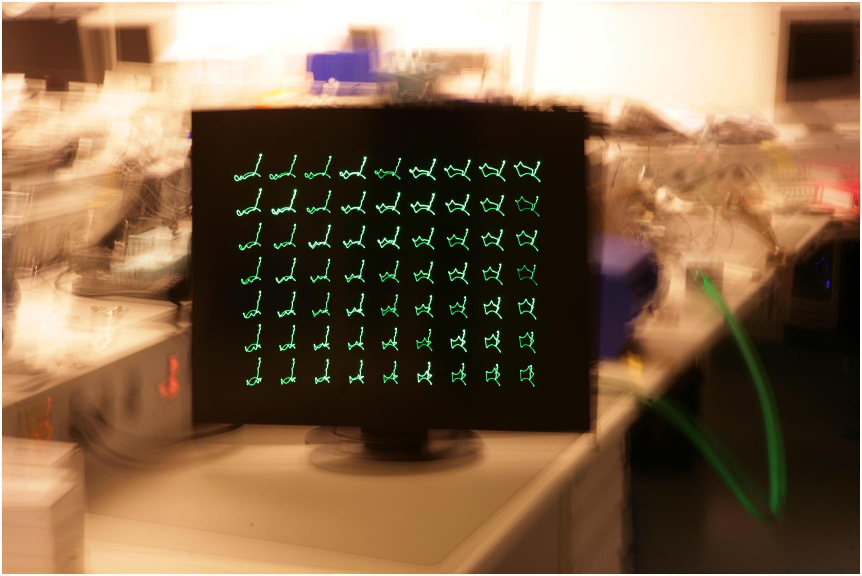

Example: Camera Shake

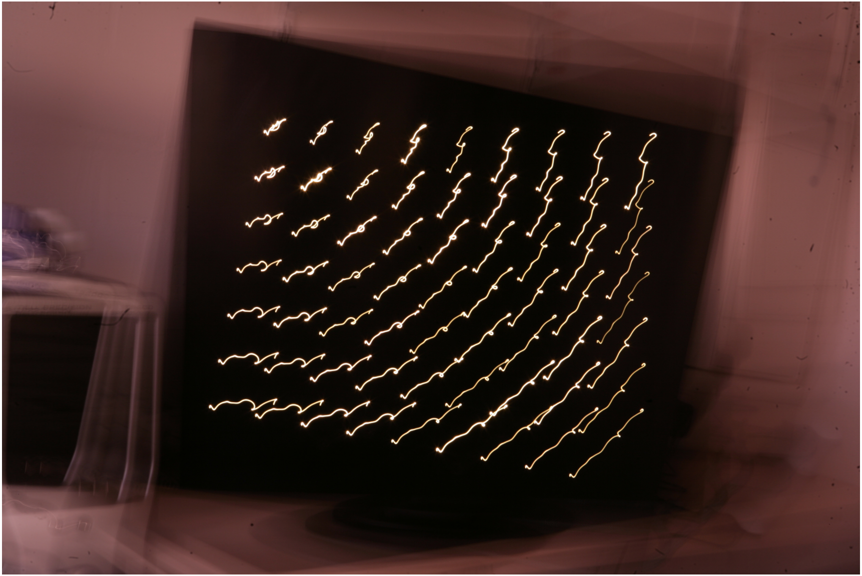

Example: Lens Abberations

Example: Lens Abberations

Example: Lens Abberations (true PSF)

Example: Lens Abberations (cartoon PSF)

Example: Air Turbulence (Exhaust Vent)

Assuming smoothness, how can we implement

Space-Variant Convolution efficiently?

Time-Invariant Convolution revisited

Time-Invariant Convolution revisited

Fast convolution by Hadamard product in Fourier space:

$$y = a * x = Ax = \underbrace{Z_y^{\mathsf{T}} F^{\mathsf{H}} \, \mathrm{Diag}(F Z_a a) F x}_{\text{our matrix notation for 1D}}$$

Time-Invariant Convolution revisited

Fast convolution by Hadamard product in Fourier space:

$$y = a * x = Ax = \underbrace{Z_y^{\mathsf{T}} F^{\mathsf{H}}\, \mathrm{Diag}(F Z_a a) F x}_{\text{our matrix notation for 1D}}$$

- However: this is inefficient if x is long and a is short.

Time-Invariant Convolution revisited

Fast convolution by Hadamard product in Fourier space:

$$y = a * x = Ax = \underbrace{Z_y^{\mathsf{T}} F^{\mathsf{H}}\,\mathrm{Diag}(FZ_a a)Fx}_{\text{our matrix notation for 1D}}$$

- However: this is inefficient if x is long and a is short.

HIGH-SPEED CONVOLUTION AND CORRELATION*

Thomas G. Stockham, Jr.
Massachusetts Institute of Technology, Project MAC
Cambridge, Massachusetts

In Proceedings of the April 26-28, 1966, Spring joint computer conference, pages 229-233. ACM, 1966.

Time-Invariant Convolution revisited

Fast convolution by Hadamard product in Fourier space:

$$y = a * x = Ax = \underbrace{Z_y^{\mathsf{T}} F^{\mathsf{H}} \operatorname{Diag}(F Z_a a) F x}_{\text{our matrix notation for 1D}}$$

- However: this is inefficient if x is long and a is short.

Stockham, 1966: Faster convolution by overlap-add:

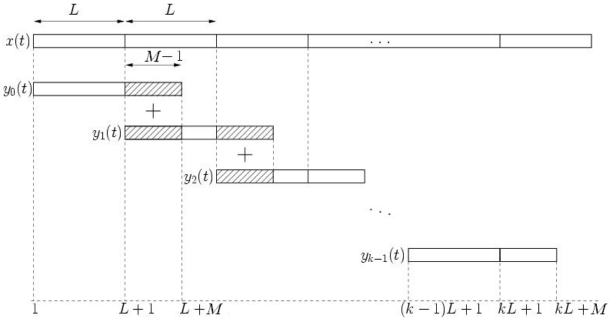

http://en.wikipedia.org/wiki/Overlap-add_method

Time-Invariant Convolution revisited

Fast convolution by Hadamard product in Fourier space:

$$y = a * x = Ax = \underbrace{Z_y^{\mathsf{T}} F^{\mathsf{H}} \,\mathrm{Diag}(FZ_a a) Fx}_{\text{our matrix notation for 1D}}$$

- However: this is inefficient if x is long and a is short.

Stockham, 1966: Faster convolution by overlap-add:

- In our matrix notations:

$$y = Z_y^{\mathsf{T}} \sum_{r=0}^{p-1} D_r^{\mathsf{T}} F^{\mathsf{H}} \,\mathrm{Diag}(FZ_a a) FZ_x C_r x$$

From Time-Invariant to Time-Variant Convolution

From Time-Invariant to Time-Variant Convolution

IEEE TRANSACTIONS ON ACOUSTICS, SPEECH, AND SIGNAL PROCESSING, VOL. ASSP-25, NO. 3, JUNE 1977 235

Short Term Spectral Analysis, Synthesis, and Modification by Discrete Fourier Transform

JONT B. ALLEN

From Time-Invariant to Time-Variant Convolution

IEEE TRANSACTIONS ON ACOUSTICS, SPEECH, AND SIGNAL PROCESSING, VOL. ASSP-25, NO. 3, JUNE 1977 235

Short Term Spectral Analysis, Synthesis, and Modification by Discrete Fourier Transform

JONT B. ALLEN

In our notation:

$$y = \underbrace{Z_y^{\mathrm{T}} \sum_{r=0}^{p-1} D_r^{\mathrm{T}} F^{\mathrm{H}}}_{\text{Synthesis}} \underbrace{\mathrm{Diag}(F Z_a a^{(r)})}_{\text{Modification}} \underbrace{F Z_x \,\mathrm{Diag}(w^{(r)}) C_r x}_{\text{Analysis}}$$

- Stockham's overlap-add is short-time Fourier transform analysis and synthesis.

From Time-Invariant to Time-Variant Convolution

IEEE TRANSACTIONS ON ACOUSTICS, SPEECH, AND SIGNAL PROCESSING, VOL. ASSP-25, NO. 3, JUNE 1977 235

Short Term Spectral Analysis, Synthesis, and Modification by Discrete Fourier Transform

JONT B. ALLEN

In our notation:

$$y = \underbrace{Z_y^{\mathrm{T}} \sum_{r=0}^{p-1} D_r^{\mathrm{T}} F^{\mathrm{H}}}_{\text{Synthesis}} \underbrace{\mathrm{Diag}(F Z_a a^{(r)})}_{\text{Modification}} \underbrace{F Z_x \, \mathrm{Diag}(w^{(r)}) C_r x}_{\text{Analysis}}$$

- Stockham's overlap-add is short-time Fourier transform analysis and synthesis.

Differences:

- Filters $a^{(0)}, \ldots, a^{(p-1)}$ lead to time-variant filtering.

From Time-Invariant to Time-Variant Convolution

IEEE TRANSACTIONS ON ACOUSTICS, SPEECH, AND SIGNAL PROCESSING, VOL. ASSP-25, NO. 3, JUNE 1977 235

Short Term Spectral Analysis, Synthesis, and Modification by Discrete Fourier Transform

JONT B. ALLEN

In our notation:

$$y = \underbrace{Z_y^{\mathrm{T}} \sum_{r=0}^{p-1} D_r^{\mathrm{T}} F^{\mathrm{H}}}_{\text{Synthesis}} \underbrace{\mathrm{Diag}(F Z_a a^{(r)})}_{\text{Modification}} \underbrace{F Z_x \, \mathrm{Diag}(w^{(r)}) C_r x}_{\text{Analysis}}$$

- Stockham's overlap-add is short-time Fourier transform analysis and synthesis.

Differences:

- Filters $a^{(0)}, \ldots, a^{(p-1)}$ lead to time-variant filtering.
- Weights $w^{(0)}, \ldots, w^{(p-1)}$ smoothly fade segments in/out.

Towards Space-Variant Convolution

time-invariant $\xrightarrow{\text{Allen, 1977}}$ time-variant

Towards Space-Variant Convolution

$$\begin{tikzcd}[column sep=huge] \text{time-invariant} \arrow[r, "\text{Allen, 1977}"] \arrow[d, "\text{Hinman et al, 1984}"'] & \text{time-variant} \\ \text{space-invariant} & \end{tikzcd}$$

SHORT-SPACE FOURIER TRANSFORM IMAGE PROCESSING

Brian L. Hinman, Jeffrey G. Bernstein, and David H. Staelin

Massachusetts Institute of Technology
Research Laboratory of Electronics
Cambridge, Massachusetts 02139

In Proceddings of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), 1984.

Towards Space-Variant Convolution

$$\begin{array}{ccc} \text{time-invariant} & \xrightarrow{\text{Allen, 1977}} & \text{time-variant} \\ \Big\downarrow{\scriptstyle \text{Hinman et al, 1984}} & & \\ \text{space-invariant} & & \end{array}$$

Hinman, 1984, in our notation (no change to Stockham):

$$y = Z_y^{\mathsf{T}} \sum_{r=0}^{p-1} D_r^{\mathsf{T}} F^{\mathsf{H}} \,\mathrm{Diag}(F Z_a a) F Z_x C_r x$$

Towards Space-Variant Convolution

$$\begin{tikzcd} \text{time-invariant} \arrow[r, "\text{Allen, 1977}"] \arrow[d, "\text{Hinman et al, 1984}"'] & \text{time-variant} \\ \text{space-invariant} & \end{tikzcd}$$

Hinman, 1984, in our notation (no change to Stockham):

$$y = Z_y^{\mathrm{T}} \sum_{r=0}^{p-1} D_r^{\mathrm{T}} F^{\mathrm{H}} \, \mathrm{Diag}(F Z_a a) F Z_x C_r x$$

- However: implement from right to left using 2D operations.

Towards Space-Variant Convolution

$$\begin{array}{ccc} \text{time-invariant} & \xrightarrow{\text{Allen, 1977}} & \text{time-variant} \\ {\scriptstyle \text{Hinman et al, 1984}}\Big\downarrow & & \Big\downarrow \\ \text{space-invariant} & \longrightarrow & \text{space-variant} \end{array}$$

Space-variant convolution in our notation:

$$y = Z_y^{\mathrm{T}} \sum_{r=0}^{p-1} D_r^{\mathrm{T}} F^{\mathrm{H}} \,\mathrm{Diag}(F Z_a a^{(r)}) F Z_x \,\mathrm{Diag}(w^{(r)}) C_r x$$

- However: implement from right to left using 2D operations.
- Appropriately chosen filters and weightings lead to space-variant convolution.

Cartoon of our Space-Variant Convolution

$$y = \underbrace{Z_y^{\mathsf{T}} \sum_{r=0}^{p-1} D_r^{\mathsf{T}} F^{\mathsf{H}} \,\mathrm{Diag}(F Z_a a^{(r)}) F \, Z_x \,\mathrm{Diag}(w^{(r)}) C_r}_{A} \; x$$

[Illustration from Michael Hirsch.]

Cartoon of our Space-Variant Convolution

$$y = \underbrace{Z_y^{\mathsf{T}} \sum_{r=0}^{p-1} D_r^{\mathsf{T}} F^{\mathsf{H}} \,\mathrm{Diag}(F Z_a a^{(r)}) F \, Z_x \,\mathrm{Diag}(w^{(r)}) C_r}_{A} \; x$$

$$= \underbrace{Z_y^{\mathsf{T}} \sum_{r=0}^{p-1} D_r^{\mathsf{T}} F^{\mathsf{H}} \,\mathrm{Diag}(F Z_x \,\mathrm{Diag}(w^{(r)}) C_r x) F Z_a B_r}_{X} \; a$$

[Illustration from Michael Hirsch.]

Cartoon of our Space-Variant Convolution

$$y = \underbrace{Z_y^{\mathsf{T}} \sum_{r=0}^{p-1} D_r^{\mathsf{T}} F^{\mathsf{H}} \,\mathrm{Diag}(F Z_a a^{(r)}) F \, Z_x \,\mathrm{Diag}(w^{(r)}) C_r}_{A} \; x$$

$$= \underbrace{Z_y^{\mathsf{T}} \sum_{r=0}^{p-1} D_r^{\mathsf{T}} F^{\mathsf{H}} \,\mathrm{Diag}(F Z_x \,\mathrm{Diag}(w^{(r)}) C_r x) F Z_a B_r}_{X} \; a$$

Fast Matrix-Vector-Multiplications (MVMs):

- For A and X: implement from right to left.

[Illustration from Michael Hirsch.]

Cartoon of our Space-Variant Convolution

$$y = \underbrace{Z_y^{\mathrm{T}} \sum_{r=0}^{p-1} D_r^{\mathrm{T}} F^{\mathrm{H}} \operatorname{Diag}(F Z_a a^{(r)}) F \, Z_x \operatorname{Diag}(w^{(r)}) C_r}_{A} \; x$$

$$= \underbrace{Z_y^{\mathrm{T}} \sum_{r=0}^{p-1} D_r^{\mathrm{T}} F^{\mathrm{H}} \operatorname{Diag}(F Z_x \operatorname{Diag}(w^{(r)}) C_r x) F Z_a B_r}_{X} \; a$$

Fast Matrix-Vector-Multiplications (MVMs):

- For A and X: implement from right to left.
- For A^{T} and X^{T}: implement from left to right.

[Illustration from Michael Hirsch.]

Space-Variant Deconvolution (single image)

Space-Variant Deconvolution (single image)

Non-Blind Deconvolution:

- Task: given observed image y and filters $a^{(0)}, \ldots, a^{(p-1)}$, estimate true image x.

$$\{y, a^{(0)}, \ldots, a^{(p-1)}\} \xrightarrow{\text{estimate}} x$$

- Gradients require MVMs with A and A^{T}.

Space-Variant Deconvolution (single image)

Non-Blind Deconvolution:

- Task: given observed image y and filters $a^{(0)}, \ldots, a^{(p-1)}$, estimate true image x.

$$\{y, a^{(0)}, \ldots, a^{(p-1)}\} \xrightarrow{\text{estimate}} x$$

- Gradients require MVMs with A and A^T.

Blind Deconvolution:

- Task: given only the observed image y, estimate true image x and possibly filters $a^{(0)}, \ldots, a^{(p-1)}$.

$$y \xrightarrow{\text{estimate}} x$$

- Gradients require MVMs with A, A^T, X, and X^T.

Properties of our Space-Variant Convolution

Properties of our Space-Variant Convolution

Runtime Complexity for MVMs with A, A^{T}, X, and X^{T}:

- For image size n, overlap o, and patch-size p, the run-time is $O(no \log p)$.
- Not much slower than usual convolution, $O(n \log p)$.

Properties of our Space-Variant Convolution

Runtime Complexity for MVMs with A, A^T, X, and X^T:

- For image size n, overlap o, and patch-size p, the run-time is $O(no \log p)$.
- Not much slower than usual convolution, $O(n \log p)$.

Expressivity:

- Is the PSF piecewise constant?

Properties of our Space-Variant Convolution

Runtime Complexity for MVMs with A, A^T, X, and X^T:

- For image size n, overlap o, and patch-size p, the run-time is $O(no \log p)$.
- Not much slower than usual convolution, $O(n \log p)$.

Expressivity:

- Is the PSF piecewise constant?

Not if the patches overlap.

Properties of our Space-Variant Convolution

Runtime Complexity for MVMs with A, A^{T}, X, and X^{T}:

- For image size n, overlap o, and patch-size p, the run-time is $O(no \log p)$.
- Not much slower than usual convolution, $O(n \log p)$.

Expressivity:

- Is the PSF piecewise constant?

  Not if the patches overlap.
- Is the PSF piecewise linear?

Properties of our Space-Variant Convolution

Runtime Complexity for MVMs with A, A^T, X, and X^T:

- For image size n, overlap o, and patch-size p, the run-time is $O(no \log p)$.
- Not much slower than usual convolution, $O(n \log p)$.

Expressivity:

- Is the PSF piecewise constant?

  Not if the patches overlap.
- Is the PSF piecewise linear?

  Not if the window functions are nonlinear.

Properties of our Space-Variant Convolution

Runtime Complexity for MVMs with A, A^T, X, and X^T:

- For image size n, overlap o, and patch-size p, the run-time is $O(no \log p)$.
- Not much slower than usual convolution, $O(n \log p)$.

Expressivity:

- Is the PSF piecewise constant?

  Not if the patches overlap.
- Is the PSF piecewise linear?

  Not if the window functions are nonlinear.
- Is the PSF differentiable along the image?

Properties of our Space-Variant Convolution

Runtime Complexity for MVMs with A, A^{T}, X, and X^{T}:

- For image size n, overlap o, and patch-size p, the run-time is $O(no \log p)$.
- Not much slower than usual convolution, $O(n \log p)$.

Expressivity:

- Is the PSF piecewise constant?

  Not if the patches overlap.
- Is the PSF piecewise linear?

  Not if the window functions are nonlinear.
- Is the PSF differentiable along the image?

  Yes, if the window functions are chosen carefully.

Properties of our Space-Variant Convolution

Runtime Complexity for MVMs with A, A^T, X, and X^T:

- For image size n, overlap o, and patch-size p, the run-time is $O(no \log p)$.
- Not much slower than usual convolution, $O(n \log p)$.

Expressivity:

- Is the PSF piecewise constant?

  Not if the patches overlap.

- Is the PSF piecewise linear?

  Not if the window functions are nonlinear.

- Is the PSF differentiable along the image?

  Yes, if the window functions are chosen carefully.

Limitation:

- True PSF must vary smoothly, otherwise we might get artefacts.

Approximation of Camera Shake

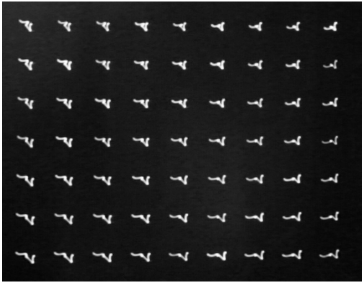

camera shake

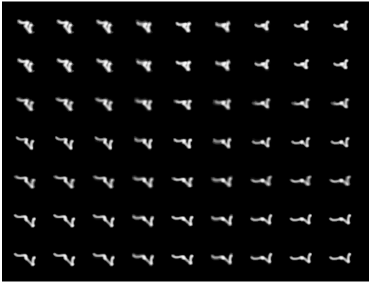

approximation $y = Ax$

with filters $a =$ 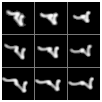

Removing Air Turbulence (multiple frames)

Hirsch, Sra, Schölkopf, Harmeling, CVPR 2010

Removing Camera Shake (single frame)

shaken image

Removing Camera Shake (single frame)

Harmeling, Hirsch, Schölkopf, NIPS, 2010

Removing Camera Shake (single frame)

Cho and Lee, Siggraph Asia, 2009

Removing Lens Abberations (non-blind)

distorted image taken with single lens element

Removing Lens Abberations (non-blind)

our result (submitted)

Related Work to Space-Variant Convolution

Nagy et O'Leary, 1998:

- Space-variant deconv. with rect. and triang. windows.
- No X and X^{T}, so only non-blind deconvolution.

Tai et al, 2009:

- Projective motion path, only for camera shake.
- No X and X^{T}, so only non-blind deconvolution.

Whyte et al, 2010, and Gupta et al, 2010:

- Blind deconvolution for camera shake using homographies.
- Only for camera shake.

Joshi et al, 2010

- Blind deconvolution for camera shake using sensor data.
- Only for camera shake.

Summary

Summary

Efficient space-variant convolution:

- Based on overlap-add and short-space Fourier analysis and synthesis.
- Can be written as MVM: $y = Ax = Xa$.
- MVMs with A, A^T, X, X^T in $O(on \log p)$ allow blind and non-blind deconvolution.

Summary

Efficient space-variant convolution:

- Based on overlap-add and short-space Fourier analysis and synthesis.
- Can be written as MVM: $y = Ax = Xa$.
- MVMs with A, A^{T}, X, X^{T} in $O(on \log p)$ allow blind and non-blind deconvolution.

Results on removing space-variant blurs:

- Camera Shake
- Lens Abberations
- Air Turbulence

Summary

Efficient space-variant convolution:

- Based on overlap-add and short-space Fourier analysis and synthesis.
- Can be written as MVM: $y = Ax = Xa$.
- MVMs with A, A^{T}, X, X^{T} in $O(on \log p)$ allow blind and non-blind deconvolution.

Results on removing space-variant blurs:

- Camera Shake
- Lens Abberations
- Air Turbulence

Many more applications:

- Astronomy: atmospheric turbulence
- Microscopy: depth-varying blurs
- Biomedical imaging: e.g. non-rigid registration in MRI

Summary

Efficient space-variant convolution:

- Based on overlap-add and short-space Fourier analysis and synthesis.
- Can be written as MVM: $y = Ax = Xa$.
- MVMs with A, A^{T}, X, X^{T} in $O(on \log p)$ allow blind and non-blind deconvolution.

Results on removing space-variant blurs:

- Camera Shake
- Lens Abberations
- Air Turbulence

Many more applications:

- Astronomy: atmospheric turbulence
- Microscopy: depth-varying blurs
- Biomedical imaging: e.g. non-rigid registration in MRI

Thanks for your attention!