Alignment and Image Comparison

Erik Learned-Miller
University of Massachusetts, Amherst

Alignment and Image Comparison

Erik Learned-Miller
University of Massachusetts, Amherst

Alignment and Image Comparison

Erik Learned-Miller

University of Massachusetts, Amherst

Lecture I

- Introduction to alignment
- A case study: mutual information alignment

Lecture I

- **Introduction to alignment**
- A case study: mutual information alignment

Examples of Alignment

- Medical image registration
- Face alignment
- Tracking
- Joint alignment (model building)

Medical Image Registration

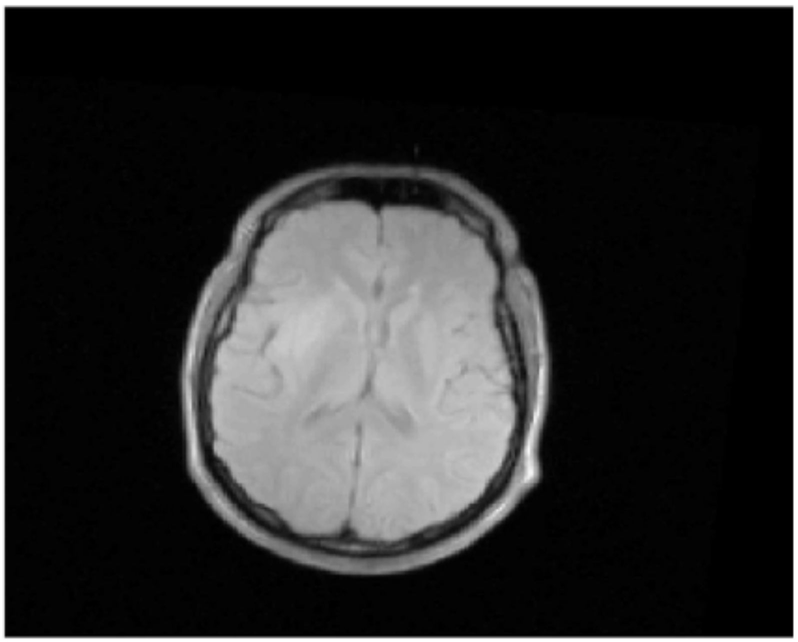

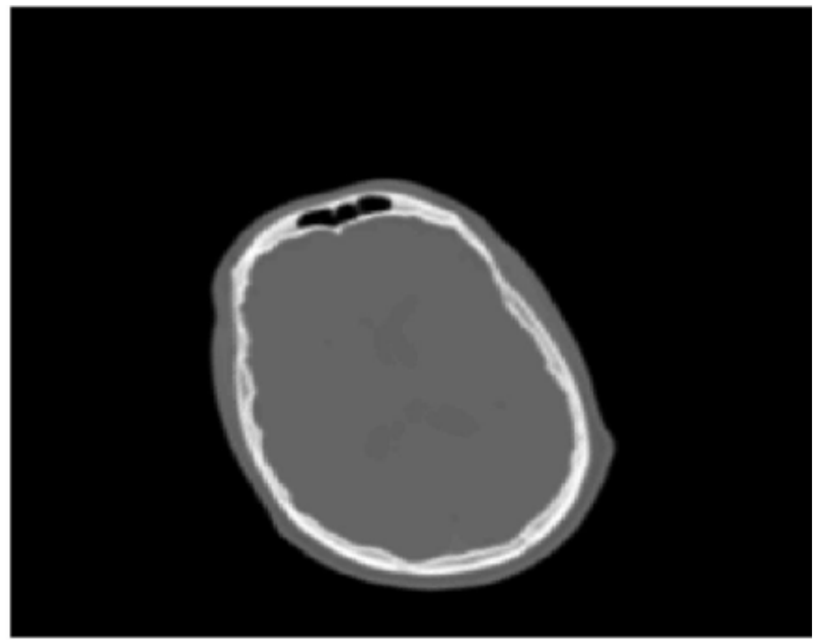

Lilla Zöllei and William Wells. Course materials for HST.582J / 6.555J / 16.456J, Biomedical Signal and Image Processing, Spring 2007.
MIT OpenCourseWare (http://ocw.mit.edu), Massachusetts Institute of Technology. Downloaded on [July 20, 2012].

Medical Image Registration

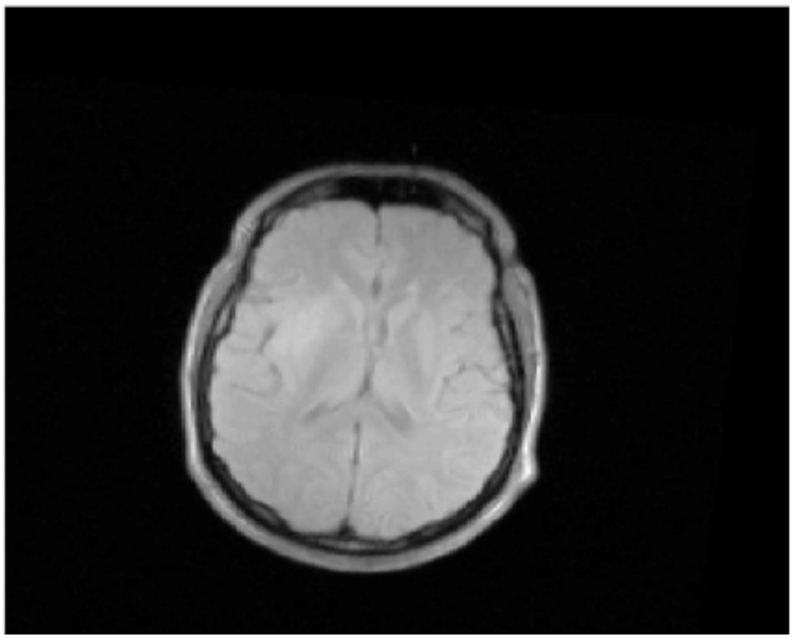

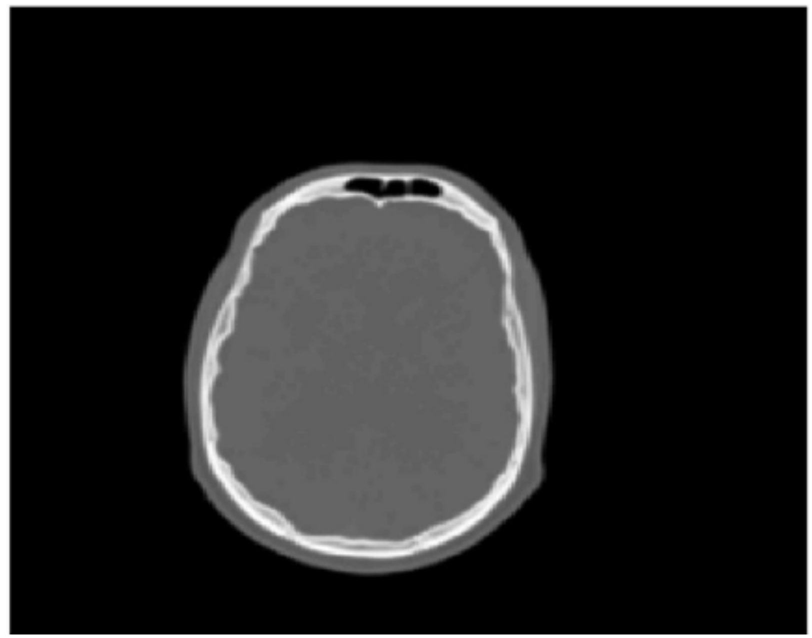

Lilla Zöllei and William Wells. Course materials for HST.582J / 6.555J / 16.456J, Biomedical Signal and Image Processing, Spring 2007.
MIT OpenCourseWare (http://ocw.mit.edu), Massachusetts Institute of Technology. Downloaded on [July 20, 2012].

Medical Image Registration

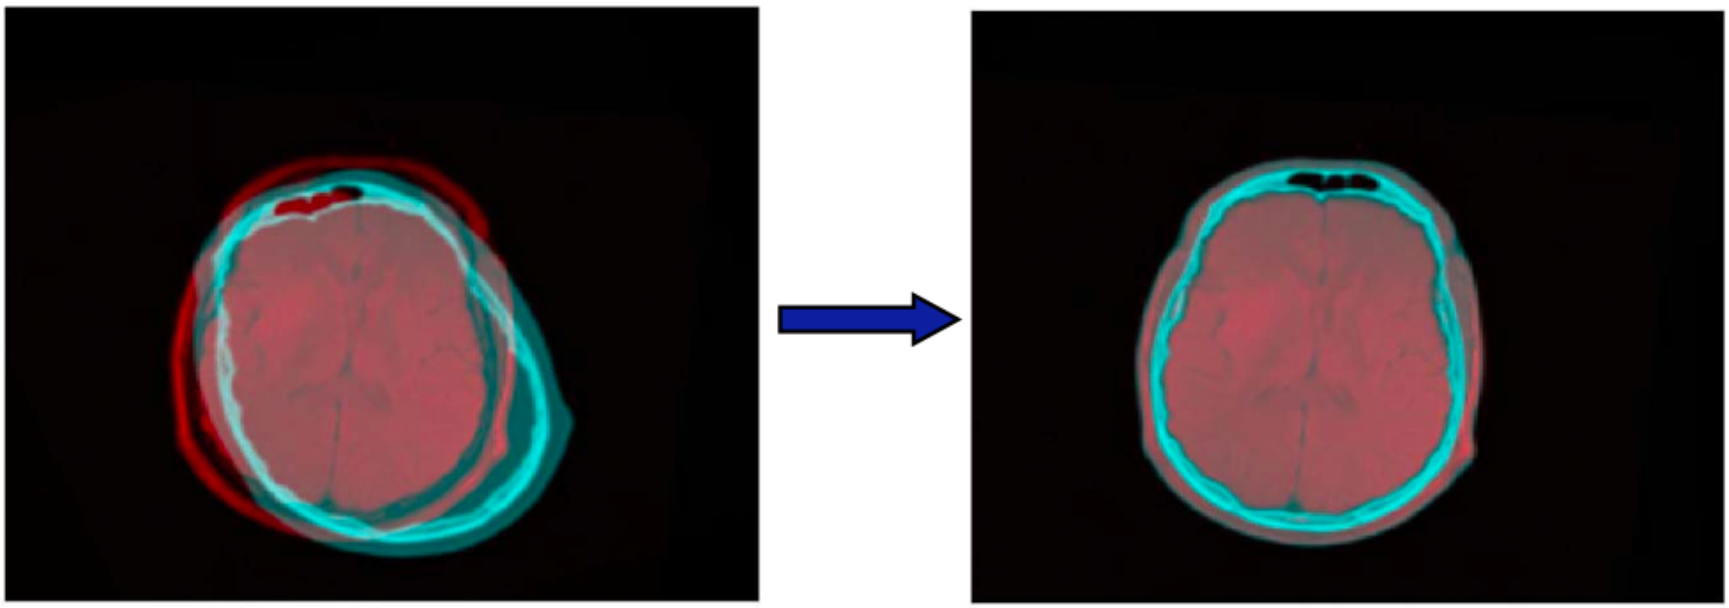

Lilla Zöllei and William Wells. Course materials for HST.582J / 6.555J / 16.456J, Biomedical Signal and Image Processing, Spring 2007.
MIT OpenCourseWare (http://ocw.mit.edu), Massachusetts Institute of Technology. Downloaded on [July 20, 2012].

Face Alignment

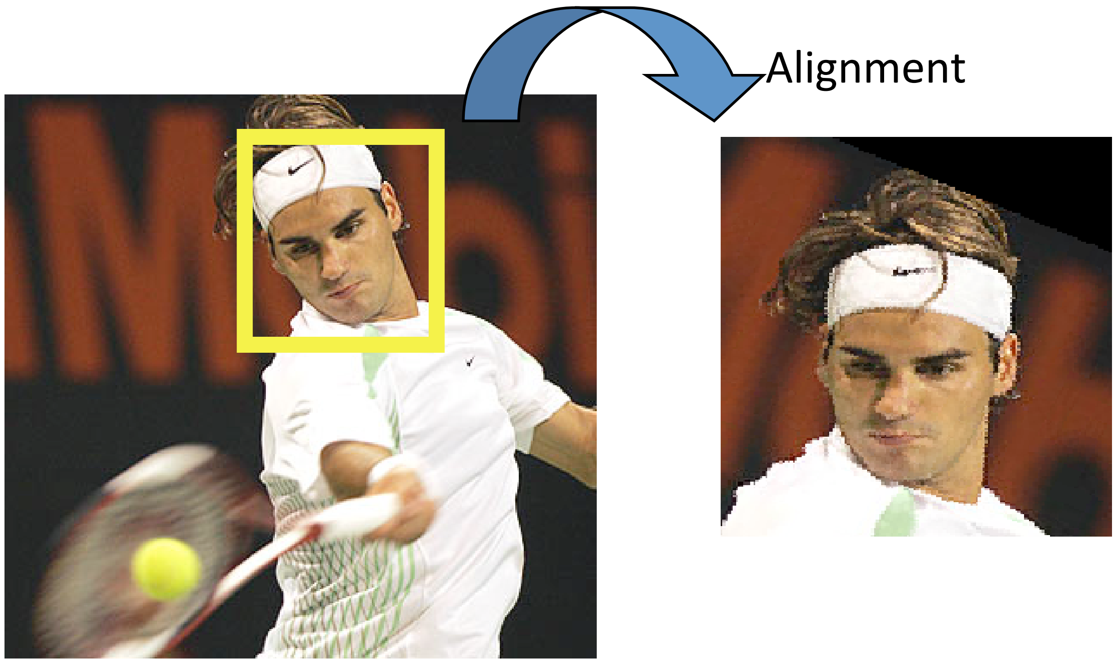

Face Alignment

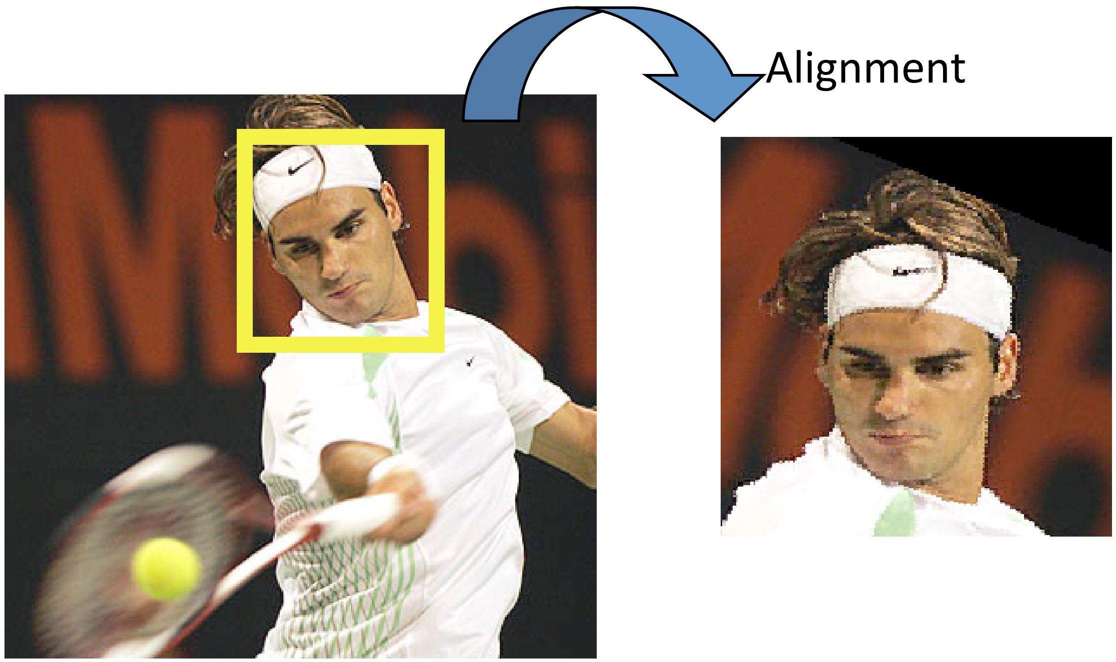

- *Surprisingly important for recognition algorithms...*

Face Alignment

Original pictures...

Face Alignment

After detection...

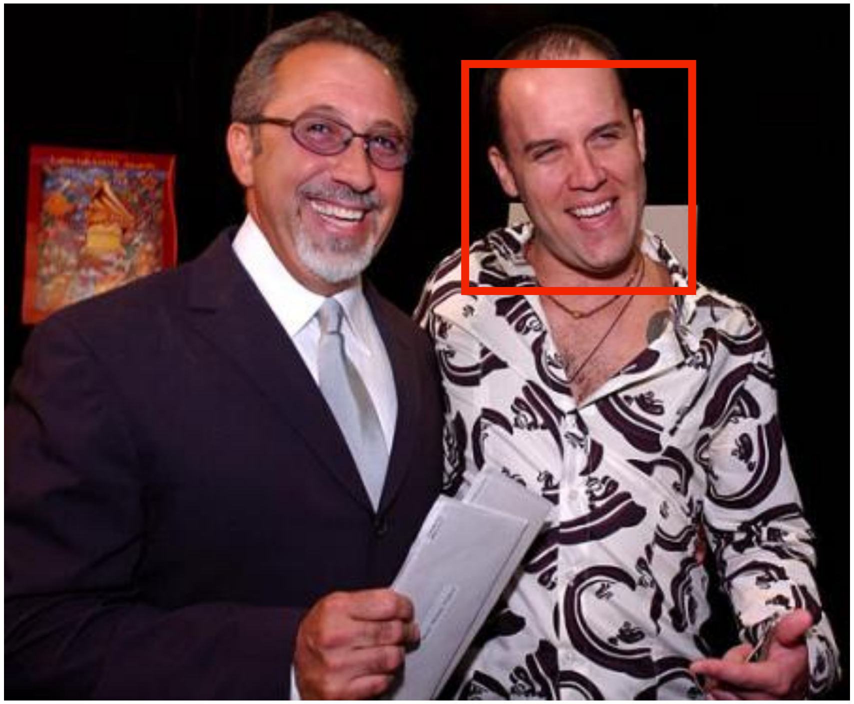

Face Alignment

Cropping...

Face Alignment

Patchwise comparison...

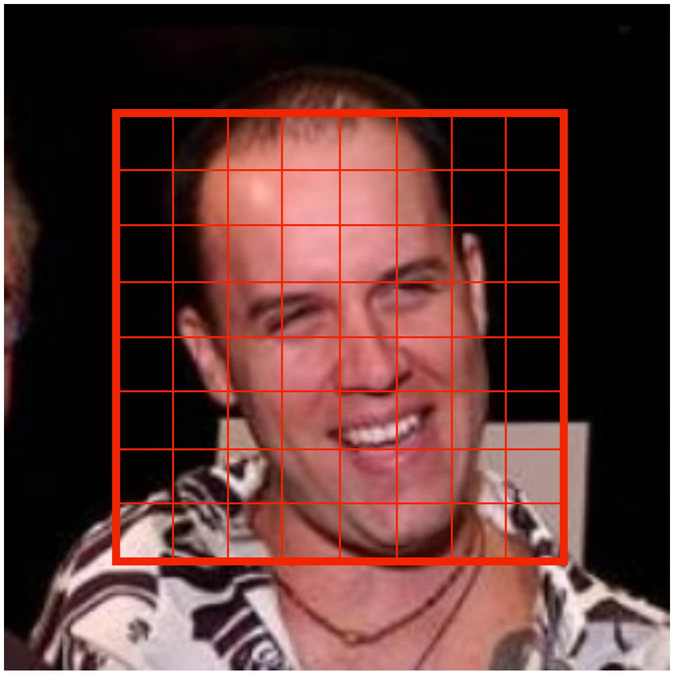

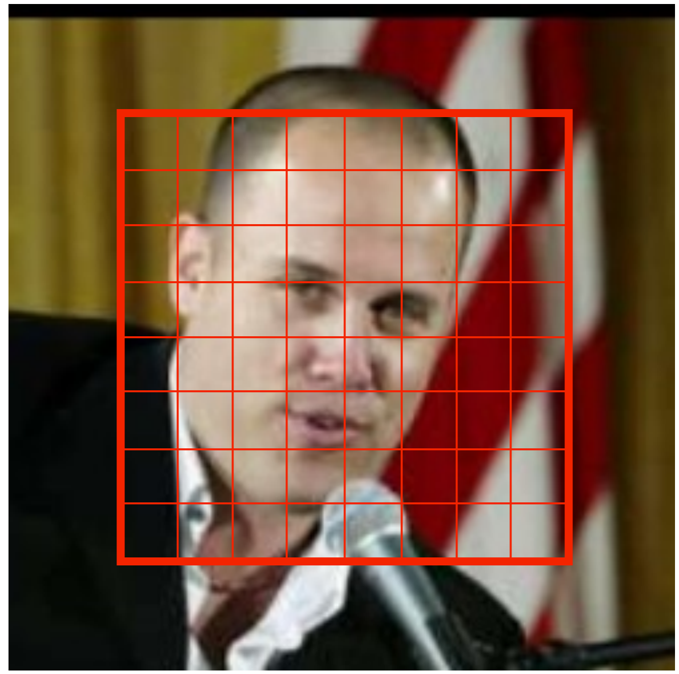

Face Alignment

Differences are too large for successful recognition

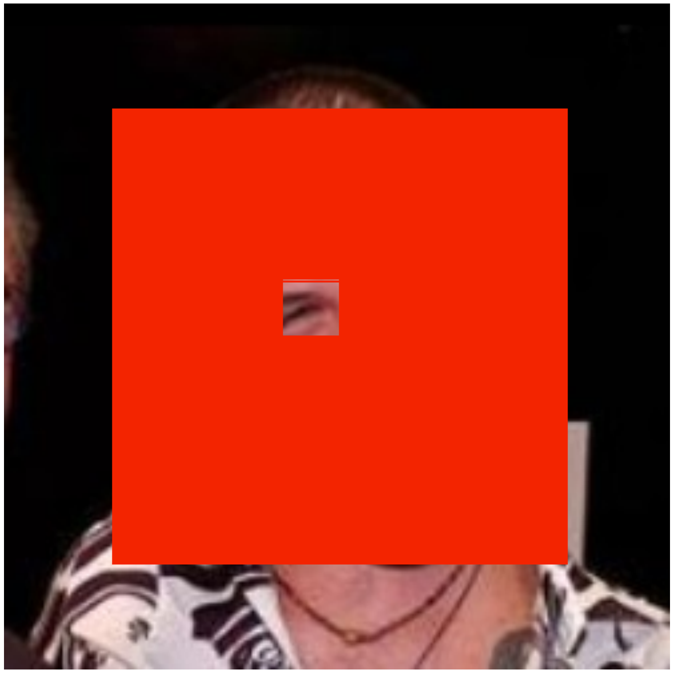

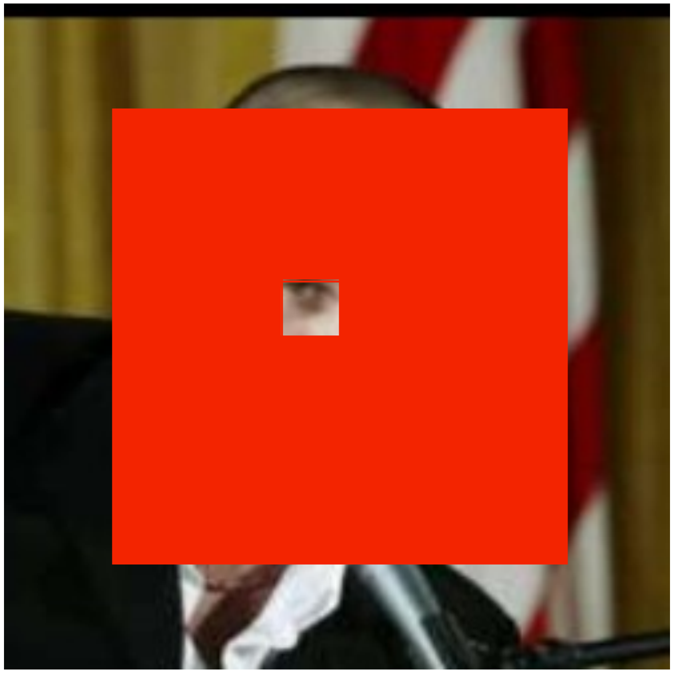

Face Alignment

Cropping...

Face Alignment

Improved alignment

Face Alignment

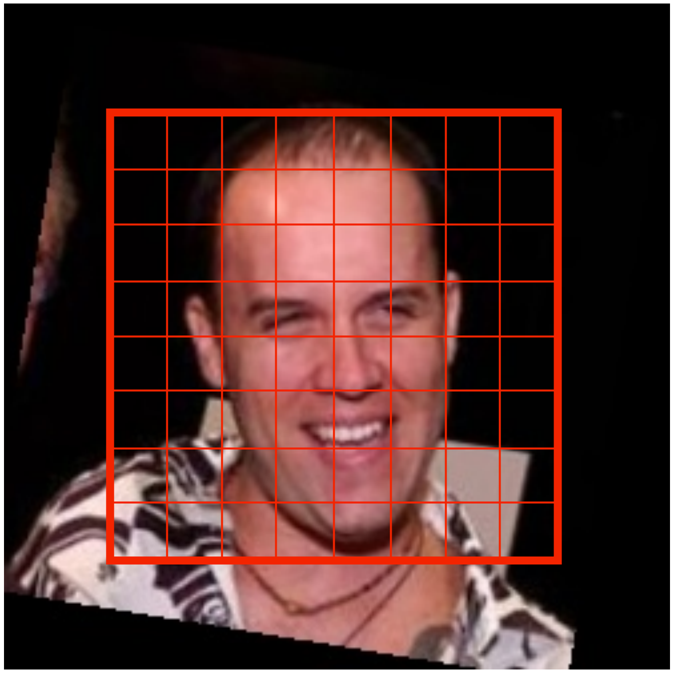

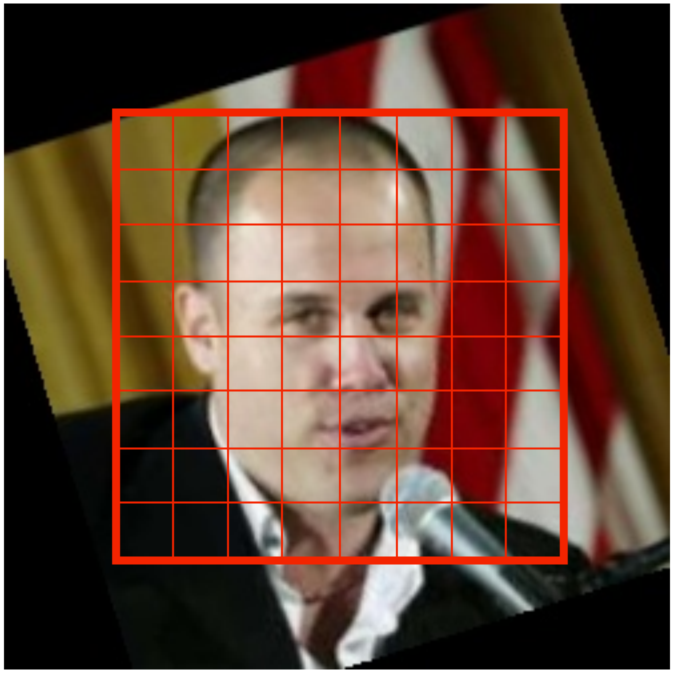

Face Alignment

Recognition greatly improved...

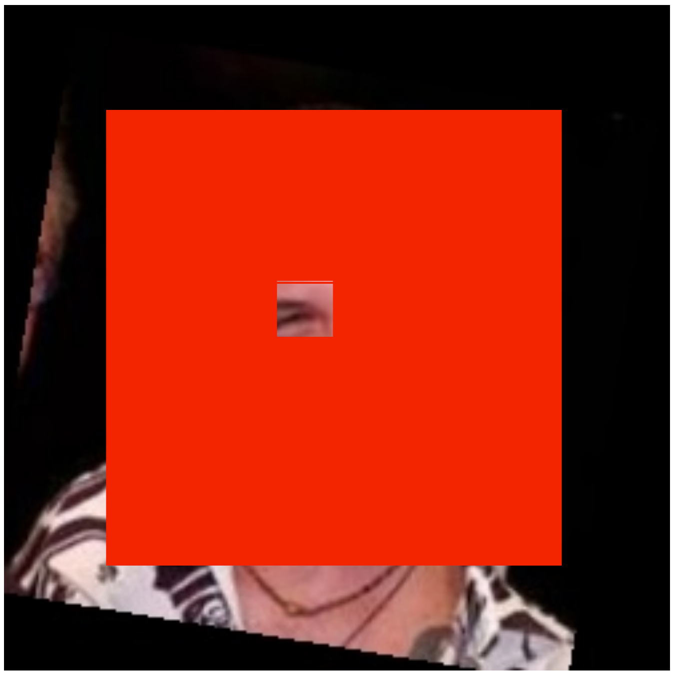

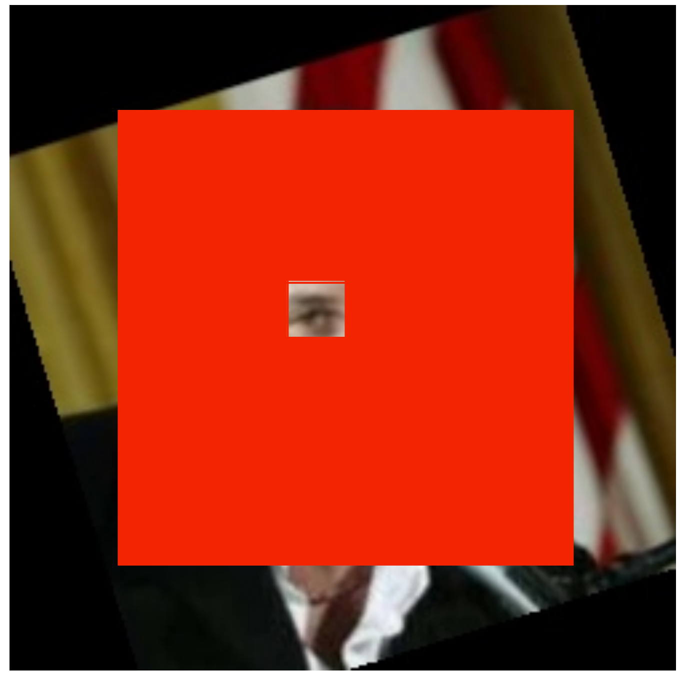

Alignment for Tracking

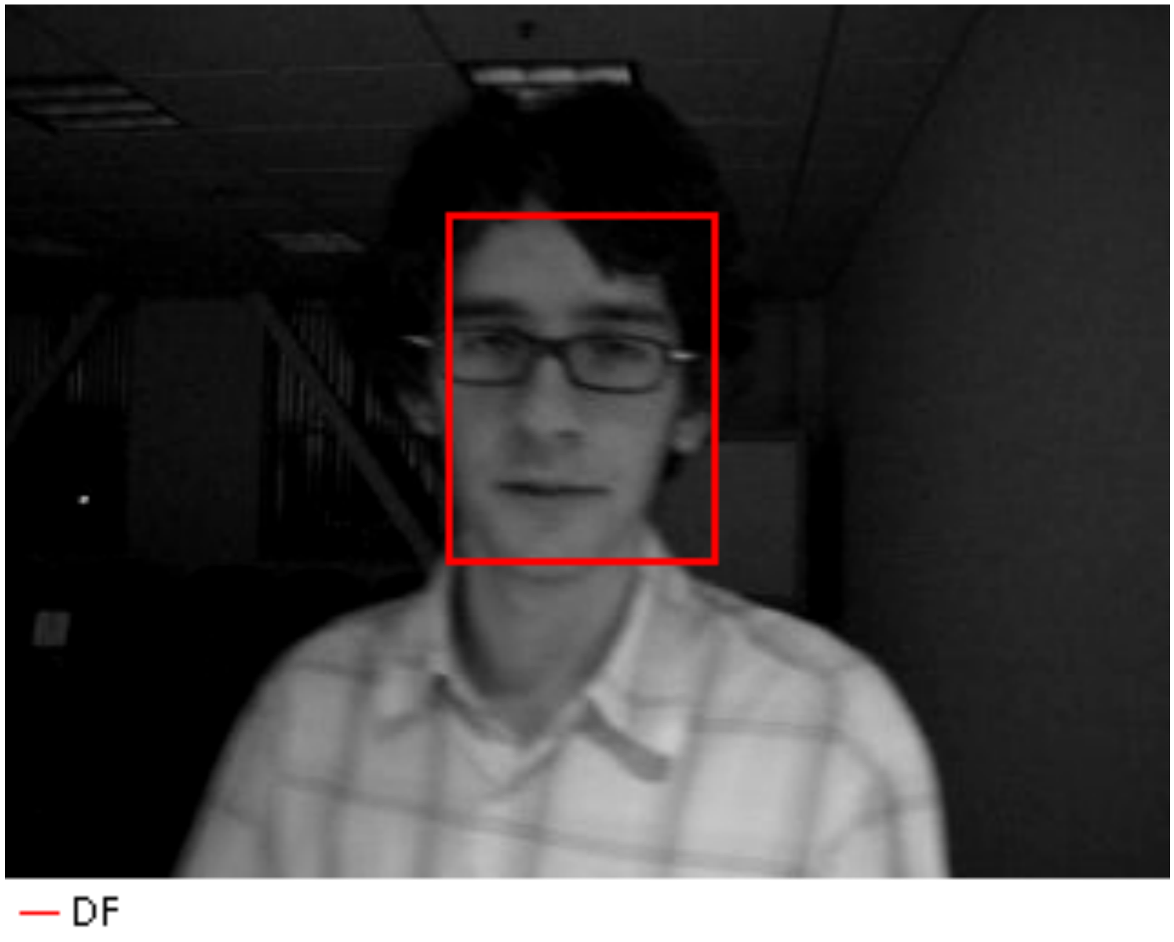

Alignment for Tracking

Frame T

Frame T+d

Alignment for Tracking

Frame T

Frame T+d

Alignment for Tracking

Frame T

Frame T+d

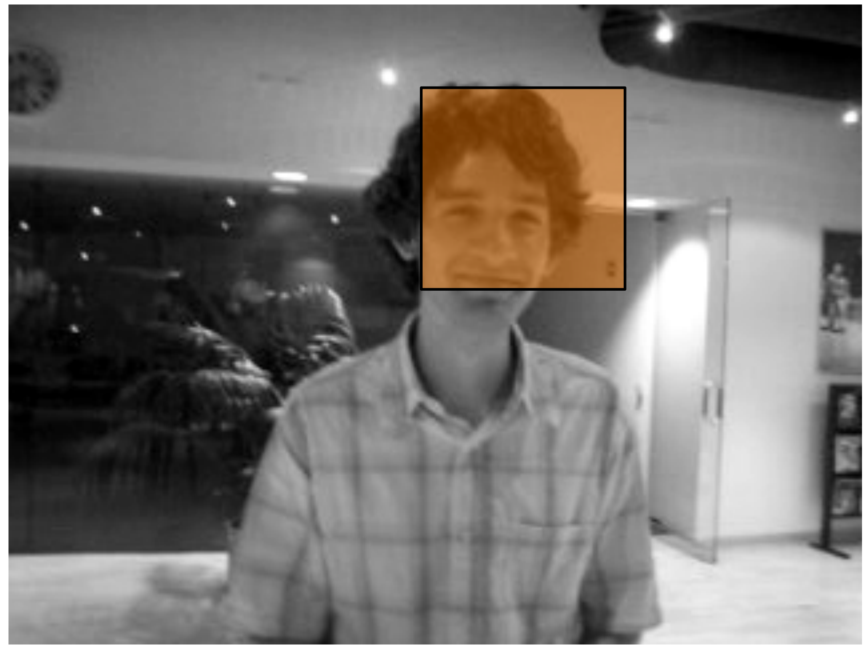

Alignment for Tracking

Frame T

Frame T+d

Alignment for Tracking

Find best match of patch I to image J,
for some set of transformations.

patch I

image J

Joint Alignment

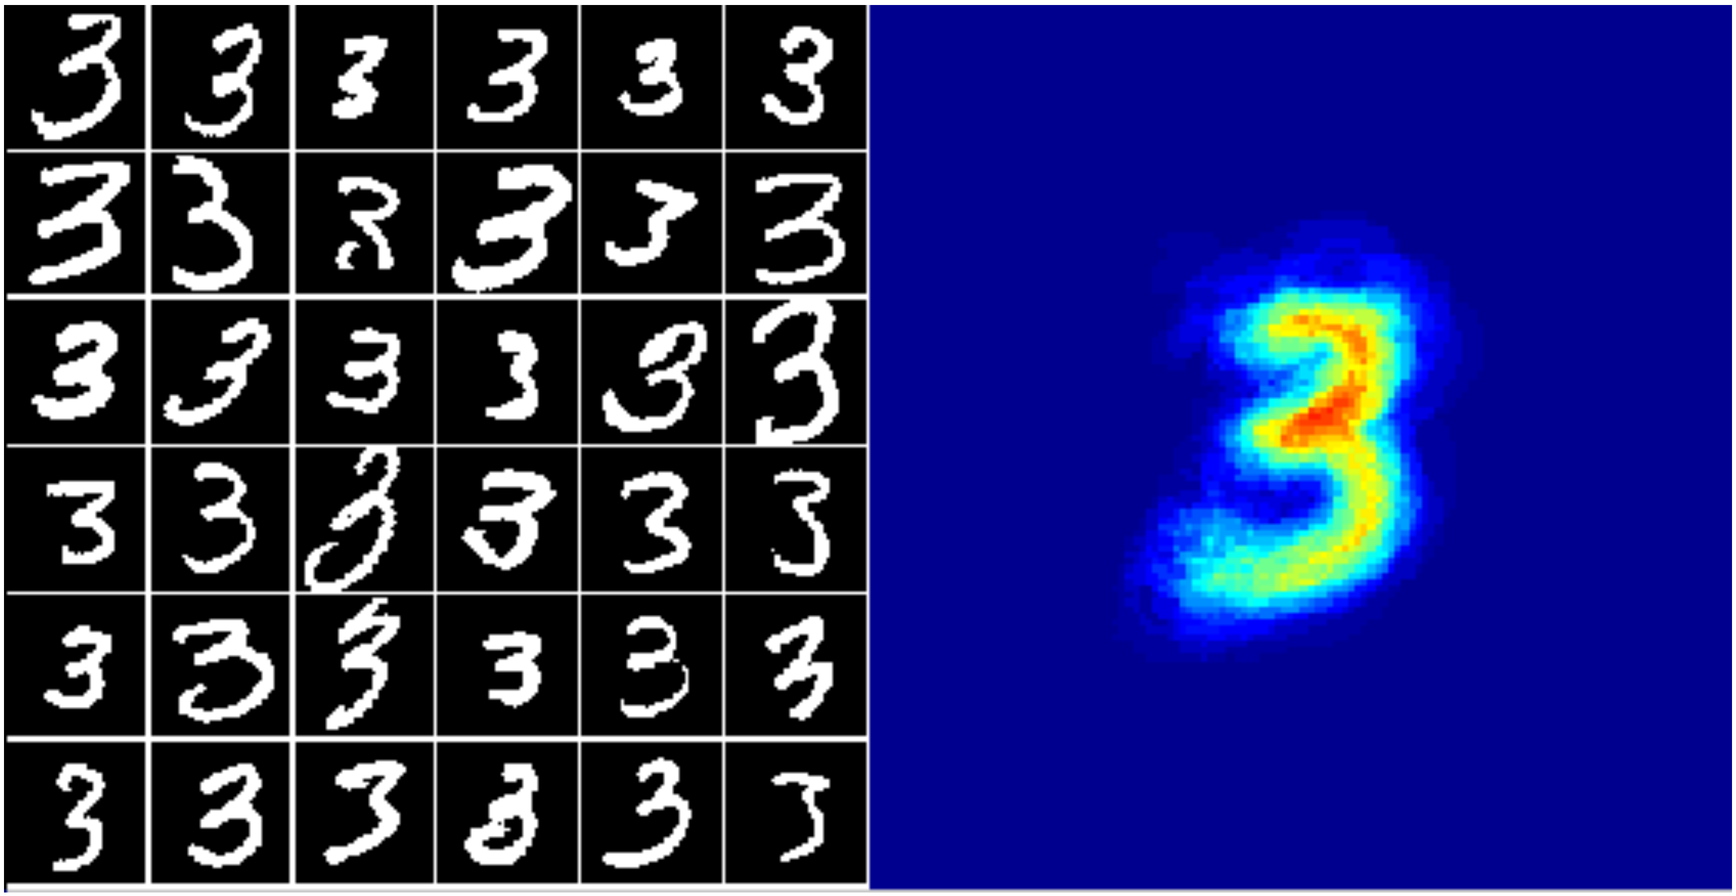

Joint Alignment

Examples of Alignment

- Medical image registration
- Face alignment
- Tracking
- Joint alignment (model building)

Questions for Thought

- How should we define alignment?
- What is the purpose of alignment?
- Is alignment a well-posed problem?
- Does a meaningful alignment always exist?
- How can human recognition be so robust to the alignments of objects?... How does the human visual system solve the alignment problem?

What’s this?

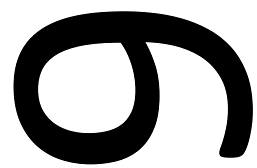

What’s this?

795

What’s this?

765

General Categories of Alignment

- Image to image
  - Align one image to another image as well as possible
    - Example: Medical images within patient MR to CT registration.
- Image to model
  - Align an image to a model for more precise evaluation
    - Example: Character recognition
- Joint alignment (congealing)
  - Align many images to each other simultaneously
    - Example: Build a face model from unaligned images.

General Categories of Alignment

- Image to image
  - Align one image to another image as well as possible
    - Example: Medical images within patient MR to CT registration.
- Image to model
  - Align an image to a model for more precise evaluation
    - Example: Character recognition
- Joint alignment (congealing)
  - Align many images to each other simultaneously
    - Example: Build a face model from unaligned images.

Image to image alignment

- Basic elements:
  - Two images I and J.
  - A family of transformations.
  - An alignment criterion.
- Definition of image to image alignment:
  - Find the transformation of I, T(I), that optimizes the alignment criterion.

Image to image alignment

- Basic elements:
  - Two images I and J.
  - A family of transformations.
  - An alignment criterion.
- Definition of image to image alignment:
  - Find the transformation of I, T(I), that optimizes the alignment criterion.
- Note: there are many other possible definitions
  - Example: transform *both* images.

Families of Transformations

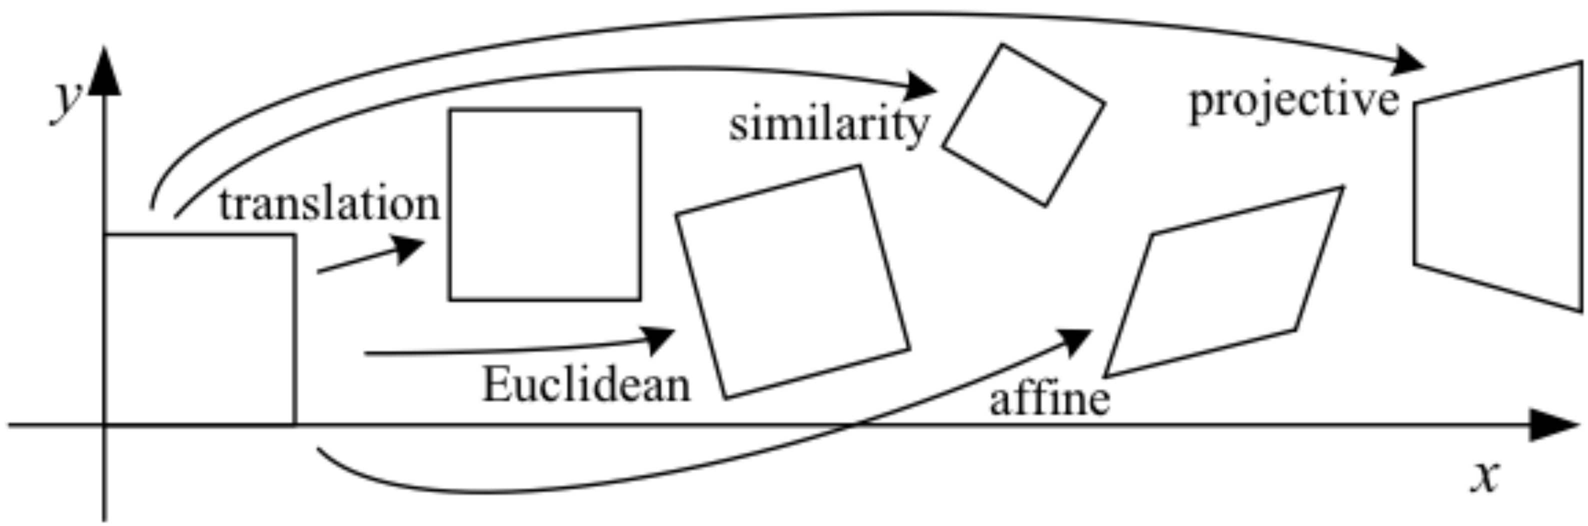

Figure 2.4 Basic set of 2D planar transformations.

From Computer Vision: Algorithms and Applications, by Rick Szeliski

Additional Transformation Families

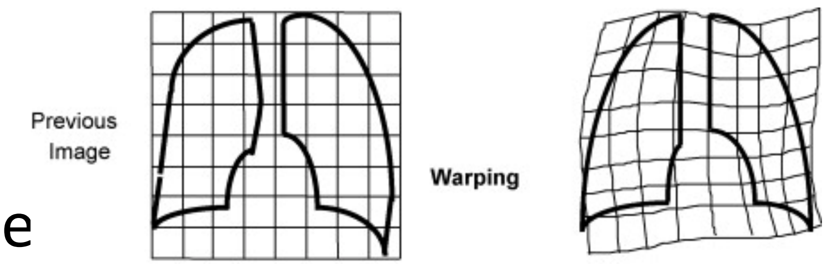

- "Warps"
  - Splines (e.g. cubic spline
    - Polynomials with "control points"
  - Diffeomorphisms:
    - Arbitrary differentiable mappings of coordinate functions
- Discontinuous and non-differentiable mappings
  - Medical images often undergo non-differentiable mappings! Examples:
    - Growth of a brain tumor.
    - Surgical removal of a portion of the brain.

Non-Diffeomorphic Transformations

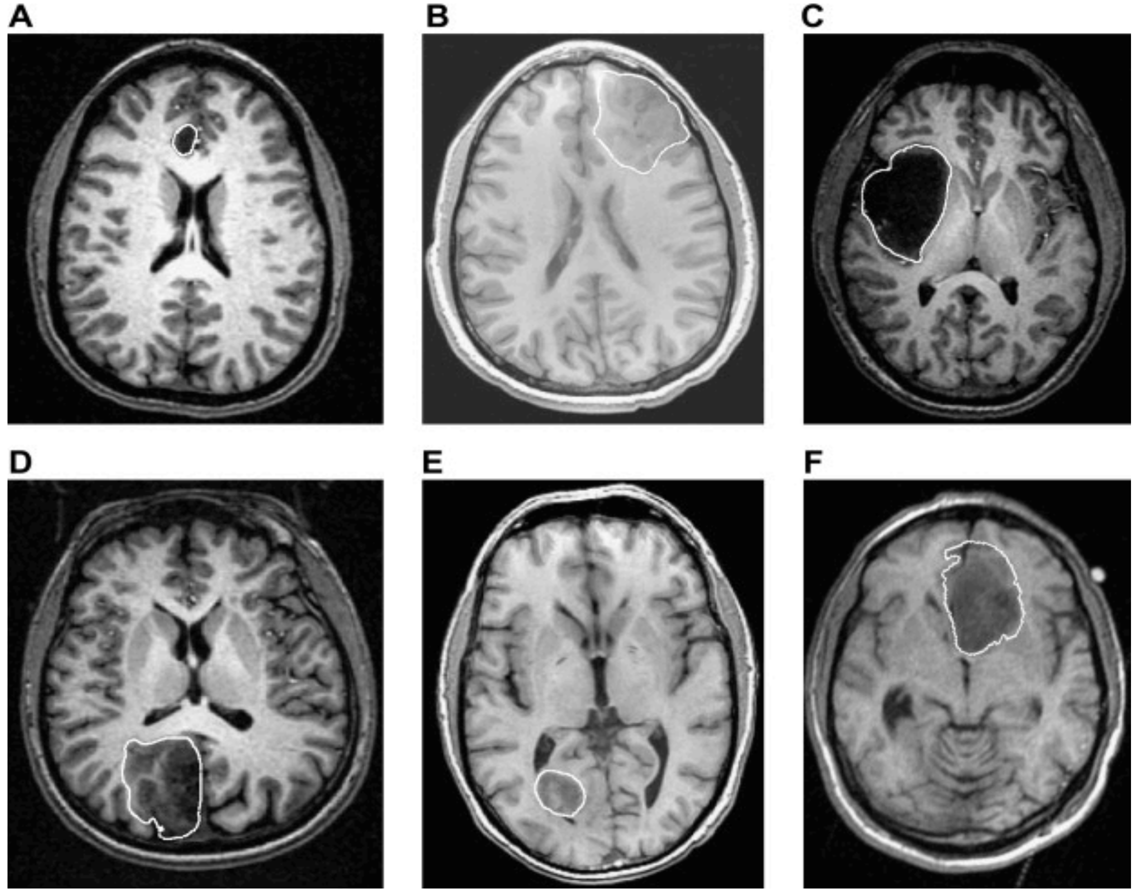

Patient-specific non-linear finite element modelling for predicting soft organ deformation in real-time; Application to non-rigid neuroimage registration

Adam Wittek[a, ,], Grand Joldes[a], Mathieu Couton[a, c, 1], Simon K. Warfield[b],

Additional Transformation Families

- Brightness transformations:
  - Scaling of brightness
  - Brightness offsets
  - Smooth brightness changes
- Important in many applications
  - Example: Correction of MRI inhomogeneity bias

Additional Transformation Families

- Brightness transformations:
  - Scaling of brightness
  - Brightness offsets
  - Smooth brightness changes
- Important in many applications
  - Example: Correction of MRI inhomogeneity bias

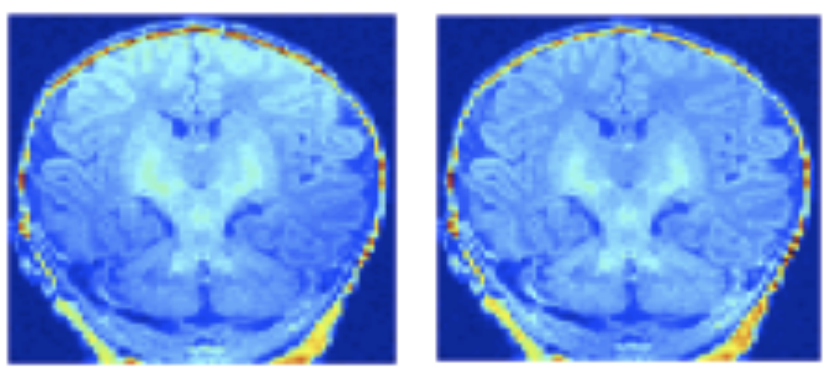

Alignment Criteria

- How to “score” an alignment.
- What should we compare at each location?
  - Pixel colors?
  - Edge features?
  - Complex features?
- Given what we are comparing, what function should we use to compare those things?
- This is an open question!

Alignment Criteria

- Alignment criteria clearly depend upon the ***image representation***:
  - A gray value at each pixel location (grayscale image).
  - A red-blue-green triple at each pixel location (standard color image).
  - An edge strength and orientation at each pixel.
  - Color histograms
  - Histograms of oriented gradients (HOG features).
  - Many other possible representations.

Alignment Criteria

- Some simple criteria:
  - Sum of squared differences of feature values at each pixel (L2 difference):

$$f(I,J) = \sum_{i=1}^{N} (I_i - J_i)^2 \quad or \quad \sqrt{\sum_{i=1}^{N} (I_i - J_i)^2}$$

  - Sum of absolute differences (L1 difference).
  - Normalized correlation
    - Usually used with gray scale representation.

How do we choose an alignment criterion?

- How do we judge whether an alignment criterion is good or bad?
  - Should it match human judgments?
  - Should it have a simple mathematical formulation?
  - What representation is a "good" representation?
    - It may depend upon the task.
- We will address this question in more detail in Feature Unit.

Definition of alignment

- Formal definition of alignment for images I and J:

J_T: transformation of image J by transform T
$\mathcal{T}$: a set of transformations

$$T^* = \underset{T \in \mathcal{T}}{\operatorname{argmin}}\, f(I, J_T)$$

- How should we perform this optimization?

Optimizing the Alignment Criterion

- Exhaustive search
  - Try all possible image transformations!
  - Gets extremely expensive as the family of transformations gets larger.
- Search for "keypoints" and align the keypoints.
  - SIFT based alignment
  - See Szeliski book for excellent treatment.
- Gradient descent ("local search")
  - Slowly change the transformation to improve the alignment score.
    - Depends strongly on "landscape" of alignment function.

Summary of Intro

- Categories of alignment
  - Image to image
  - Image to model
  - Joint image alignment
- Definition of image to image alignment
  - Choose a representation for images
  - Choose a family of transformations
  - Choose a criterion of alignment
  - Optimize alignment over family of transformations

Lecture I

- Introduction to alignment
- **A case study: mutual information alignment**

Alignment by the Maximization of Mutual Information

- Classic example of M.I. alignment:
  - Aligning medical images from different modalities
    - magnetic resonance images
    - computed tomography images
  - Magnetic resonance images
    - Measures proton density (in some cases) of tissue
  - Computed tomography
    - Measures X-ray transparency
- Original work by Viola and Wells, and also by Collignon et al.

CT and MR images

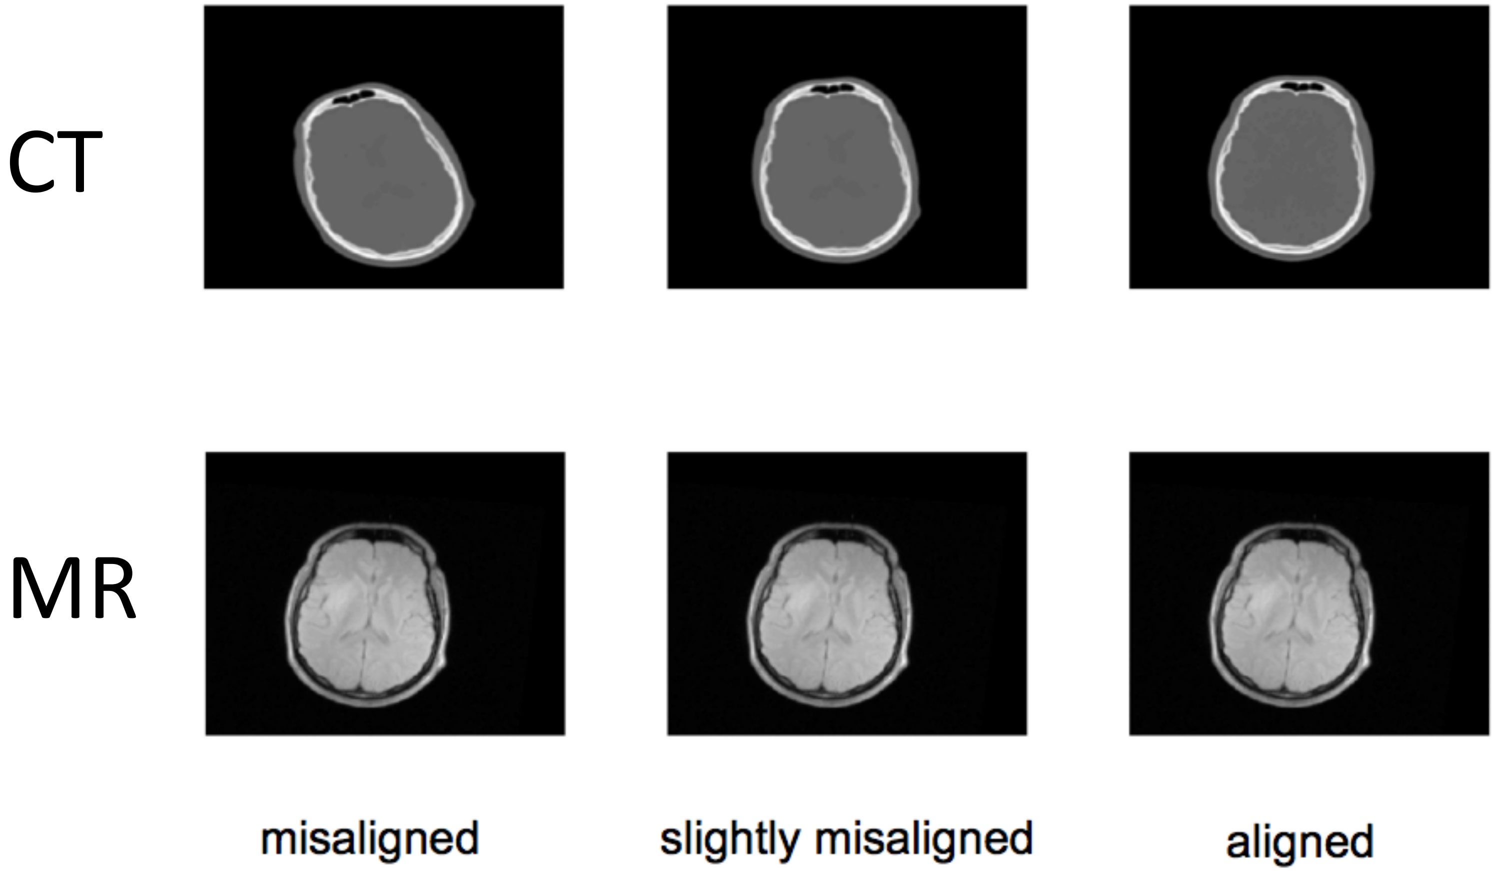

Two Different MR modalities

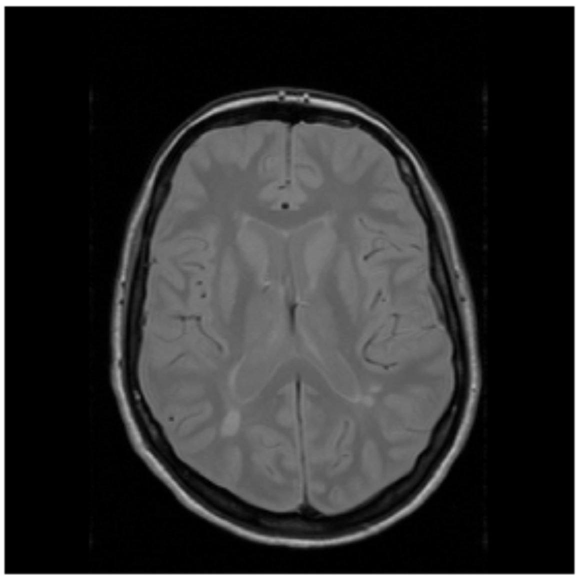

Same anatomy but
left is T_1 weighted,
right is T_2 weighted

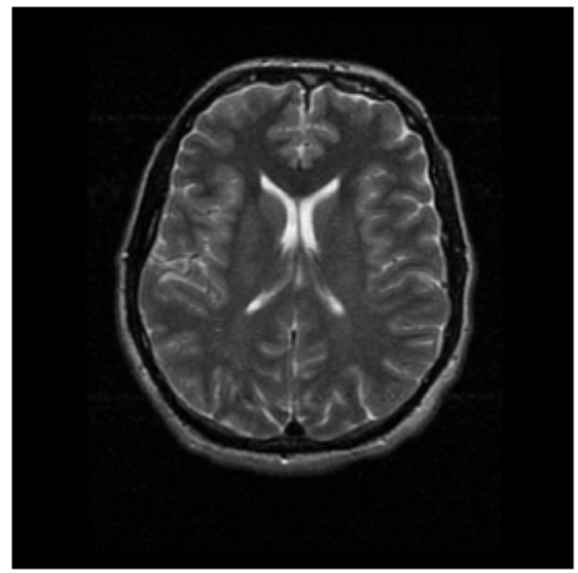

Traditional alignment criteria fail

- When images are registered:
  - L2 error is not low
    - Same tissue has different value in MR and CT
  - Correlation score is not high
    - MR and CT are not linearly related
  - Many other criteria also fail
  - Need a different criterion....

The Random Variable View

- Consider two images I and J of different modalities. Assume for the moment that they are aligned.
- Consider a random pixel location X.
  - Let X_i be the brightness value in image I at X.
  - Let X_j be the brightness value in image J at X.
- X_i and X_j are random variables.
- What can we say about X_i and X_j?

Relationship Between X_i and X_j

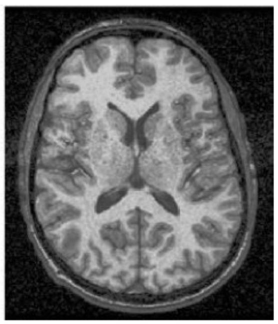

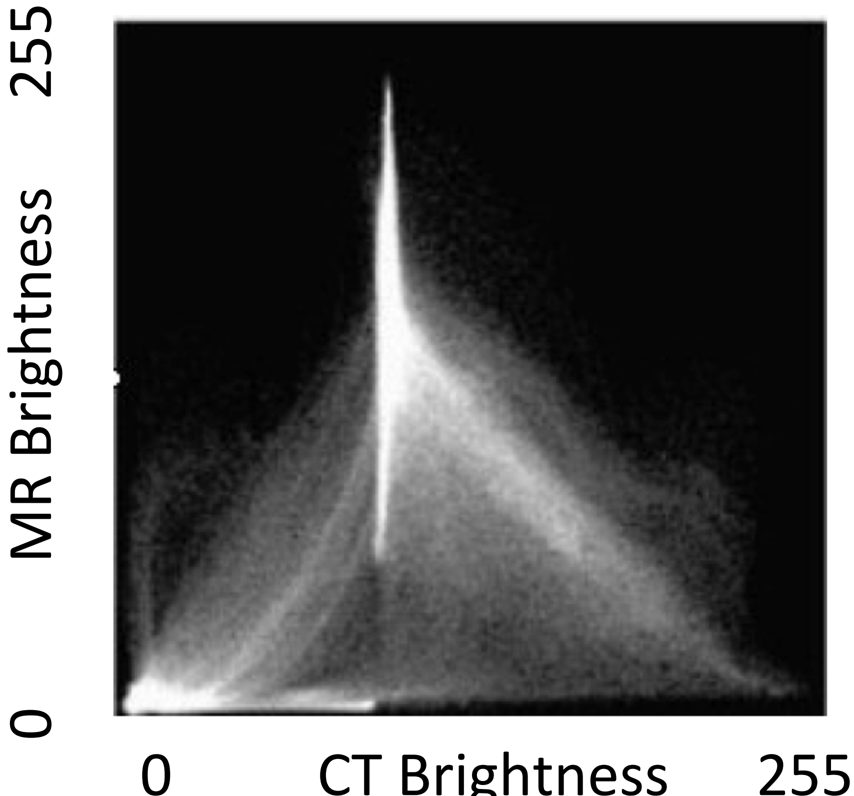

Joint Brightness
Histogram

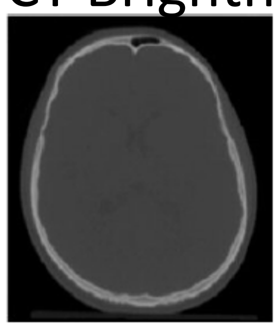

Figures from
Michael Brady
Oxford University

MR and CT images

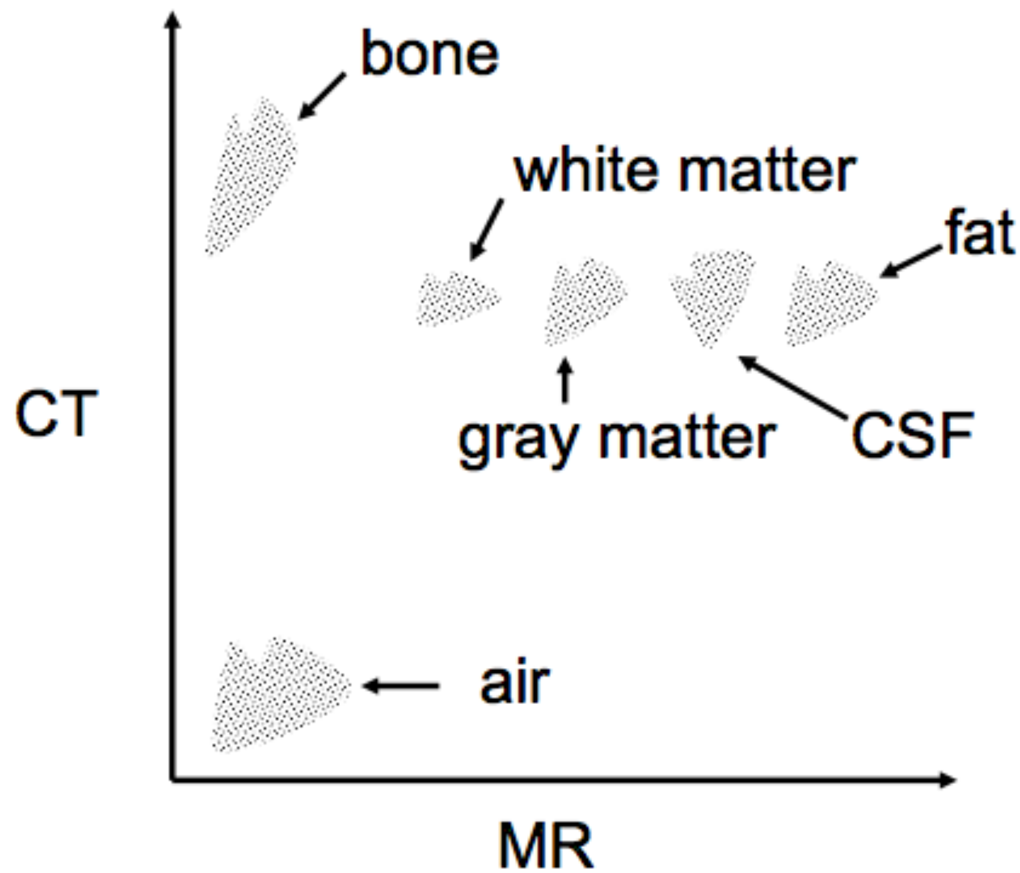

Statistical Dependence

- MR and CT values are NOT linearly related
- MR and CT are not FUNCTIONALLY related
  - There is no function that maps one to another.
- However, they have strong *statistical dependence.*
- When MR and CT pixels are unaligned, the dependence drops.
  - Basic idea: move images around to *maximize statistical dependence*

Joint Distribution as a Function of Displacement

0mm 2mm 5mm

MR-MR

MR-CT

MR-PET

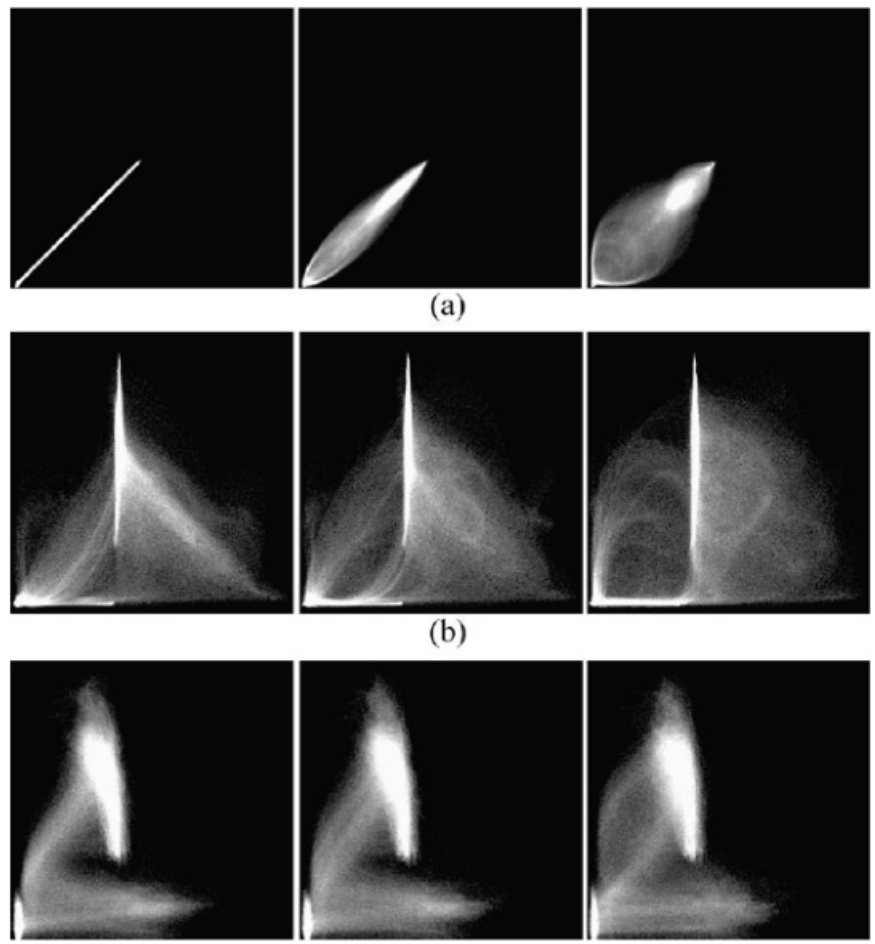

Figures from
Michael Brady
Oxford University

Mutual Information

Two random variables X and Y are *statistically independent* if and only if

$$P(X,Y) = P(X)P(Y)$$

The *mutual information* between two random variables X and Y is

$$I(X,Y) = \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} P(x,y) \log \frac{P(x,y)}{P(x)P(y)}.$$

Mutual Information

- Is 0 only when two variables are independent.
- Goes up as variables become more dependent.
  - Goal: maximize dependence so maximize mutual information.

Example on real data

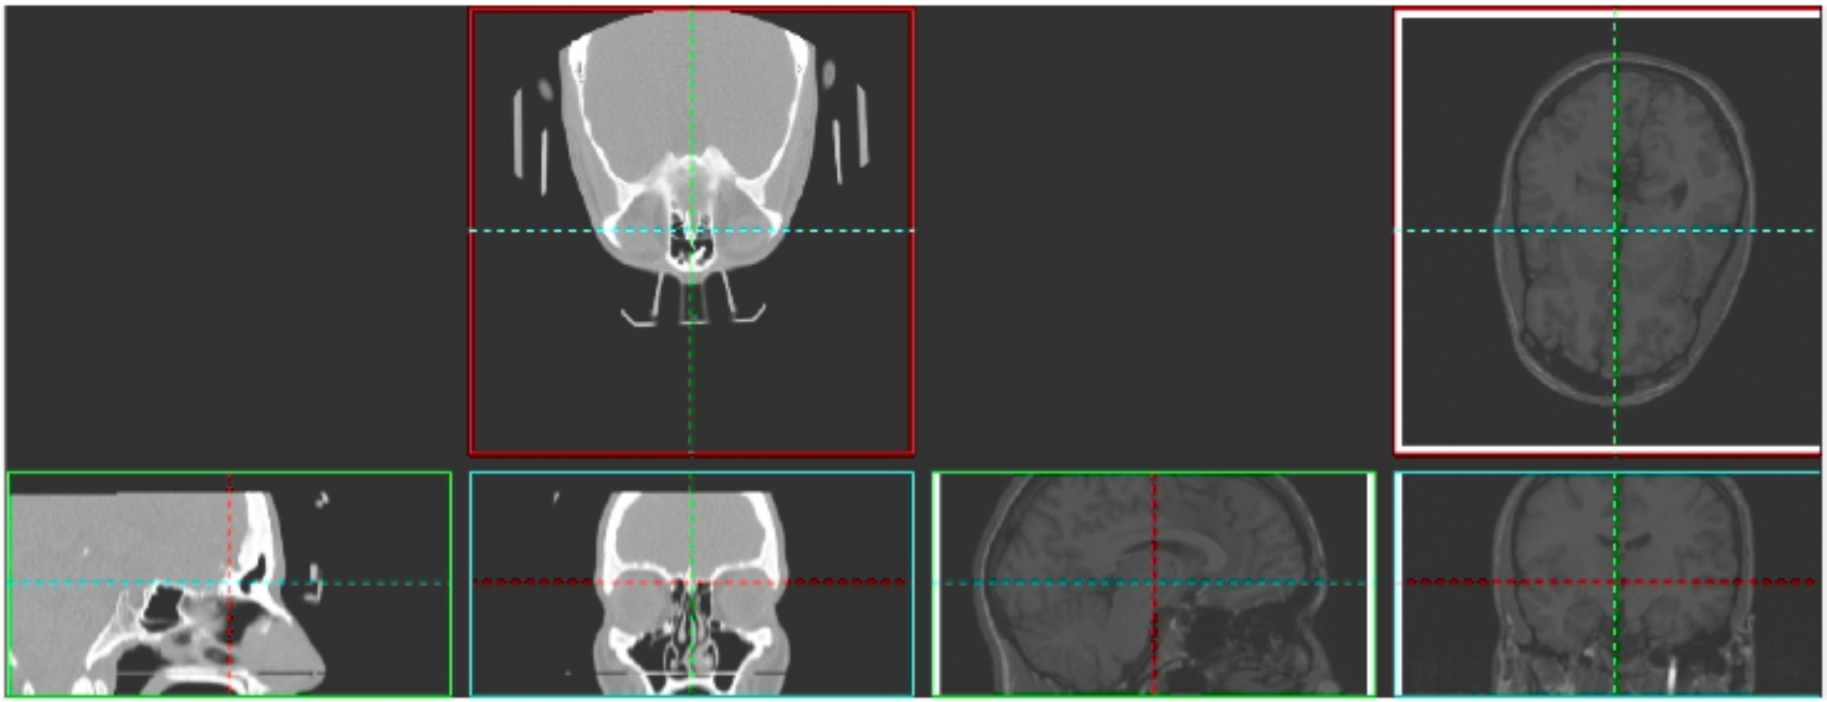

Lilla Zöllei and William Wells. Course materials for HST.582J / 6.555J / 16.456J, Biomedical Signal and Image Processing, Spring 2007.
MIT OpenCourseWare (http://ocw.mit.edu), Massachusetts Institute of Technology. Downloaded on [July 20, 2012].

Example on real data

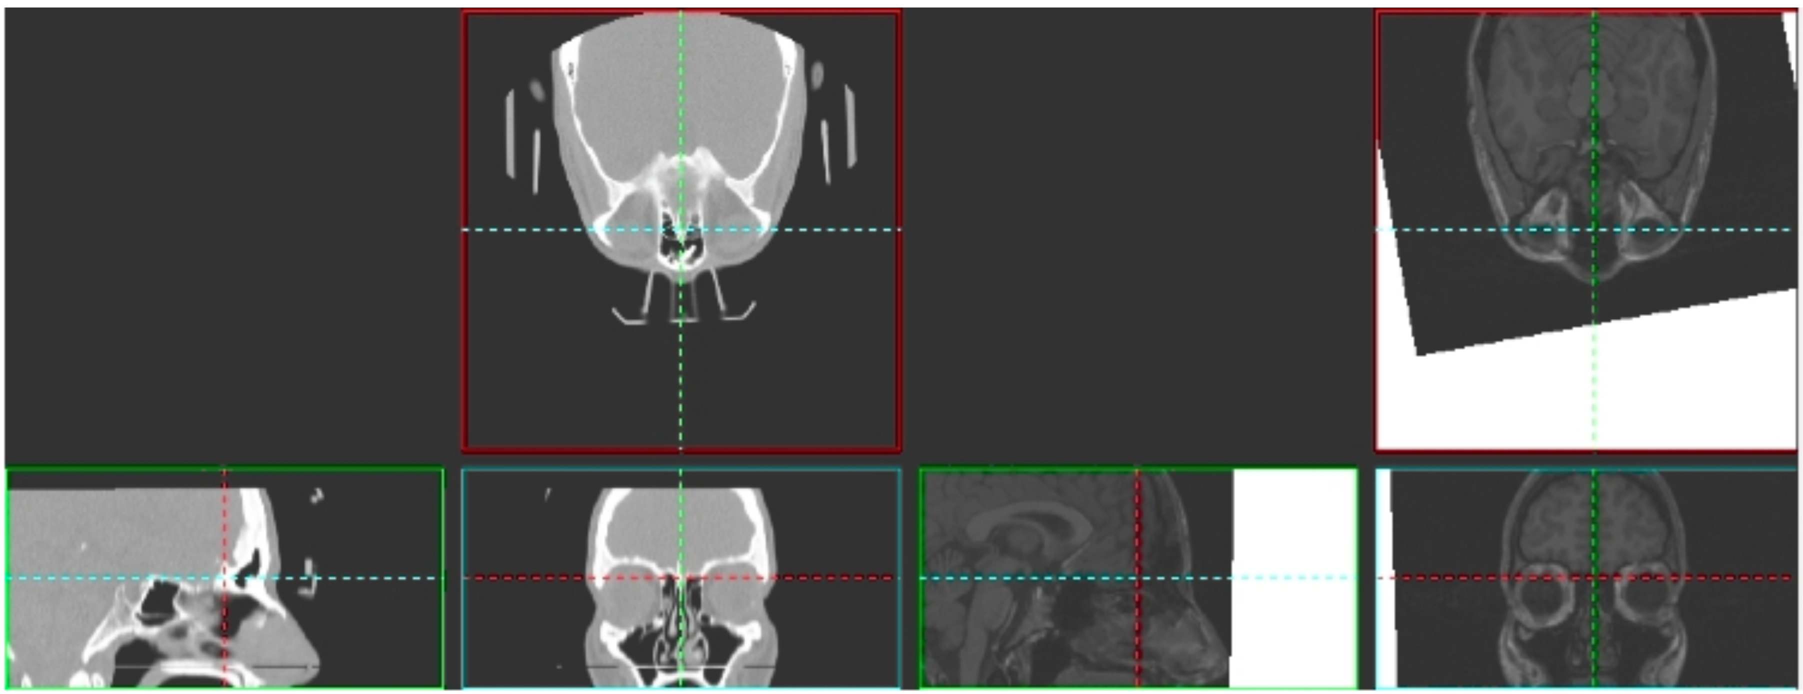

Lilla Zöllei and William Wells. Course materials for HST.582J / 6.555J / 16.456J, Biomedical Signal and Image Processing, Spring 2007.
MIT OpenCourseWare (http://ocw.mit.edu), Massachusetts Institute of Technology. Downloaded on [July 20, 2012].

Families of Transformations

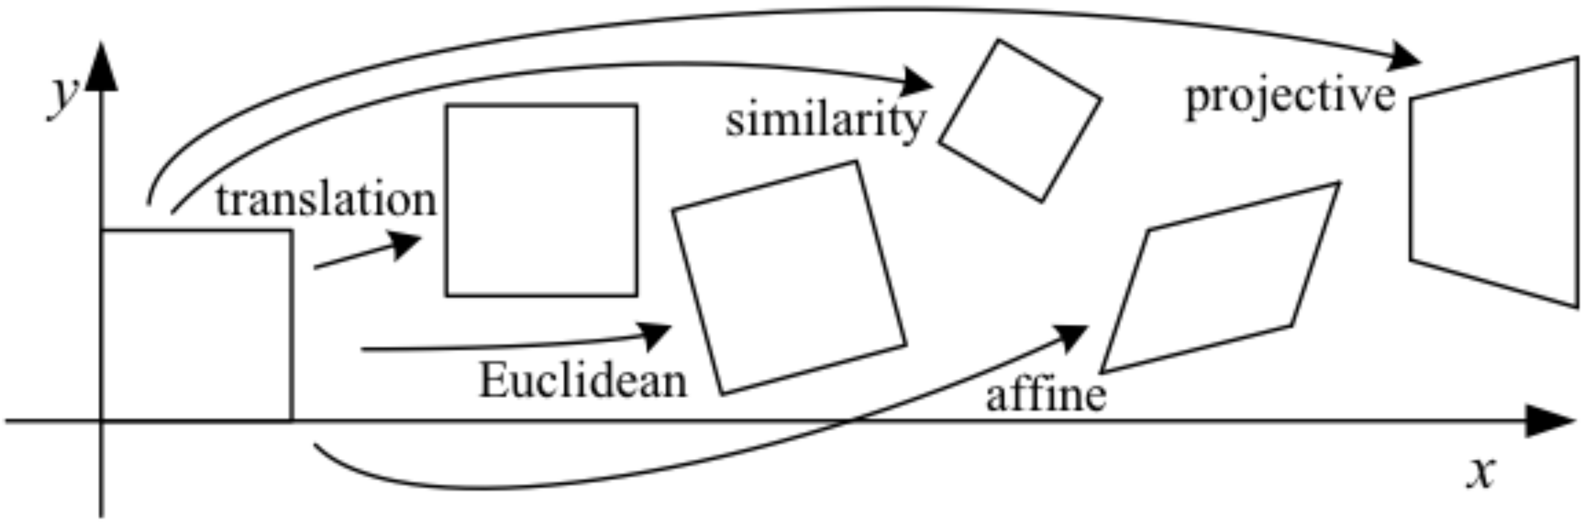

Figure 2.4 Basic set of 2D planar transformations.

From Computer Vision: Algorithms and Applications, by Rick Szeliski

How do we move pixels?

- Think about moving coordinates, not pixels.

Transformations

```
figure(1);
[ox oy]=meshgrid(1:10,1:10);
plot(ox,oy,'.');
line([0 0],[-5 15]);
line([-5 15],[0 0]);
axis([-5,15,-5,15]);
```

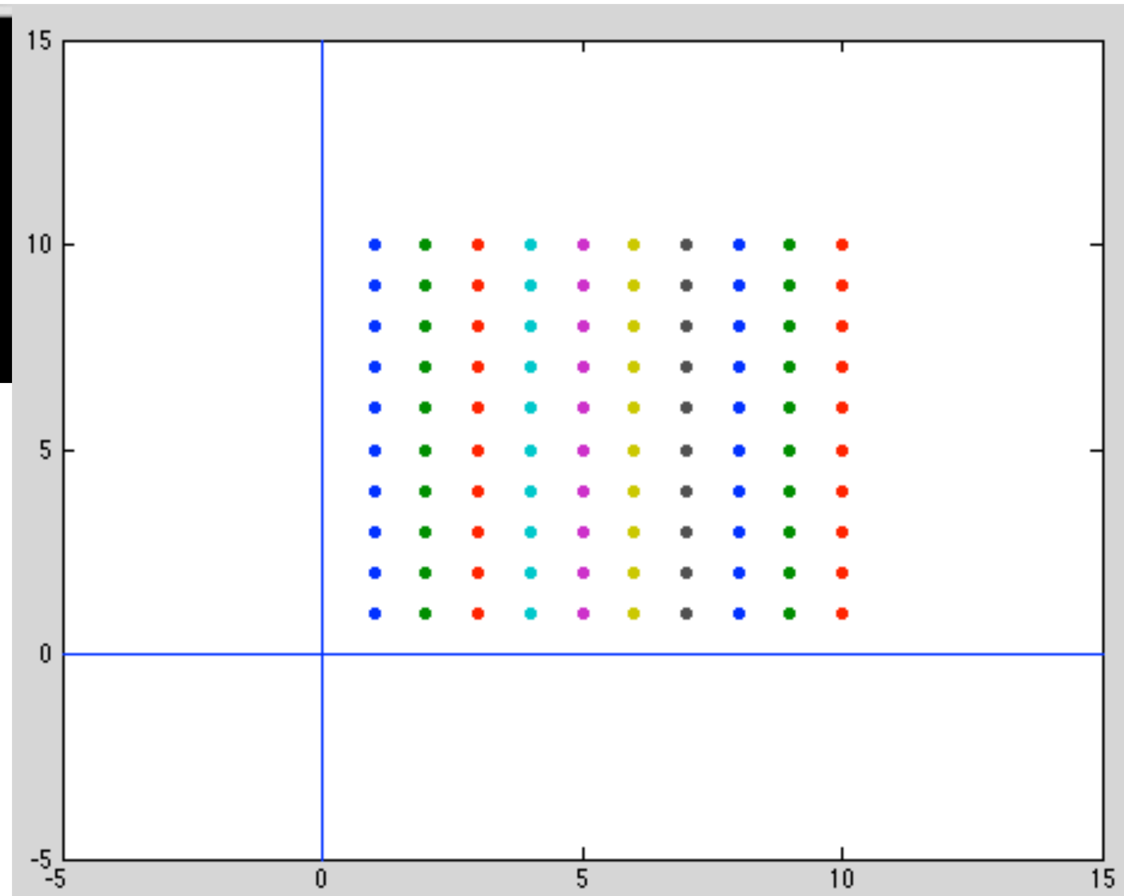

Translation

```
figure(2);
nx=ox+5;
ny=oy;
plot(nx,ny,'.');
axis([-5,15,-5,15]);
line([0 0],[-5 15]);
line([-5 15],[0 0]);
```

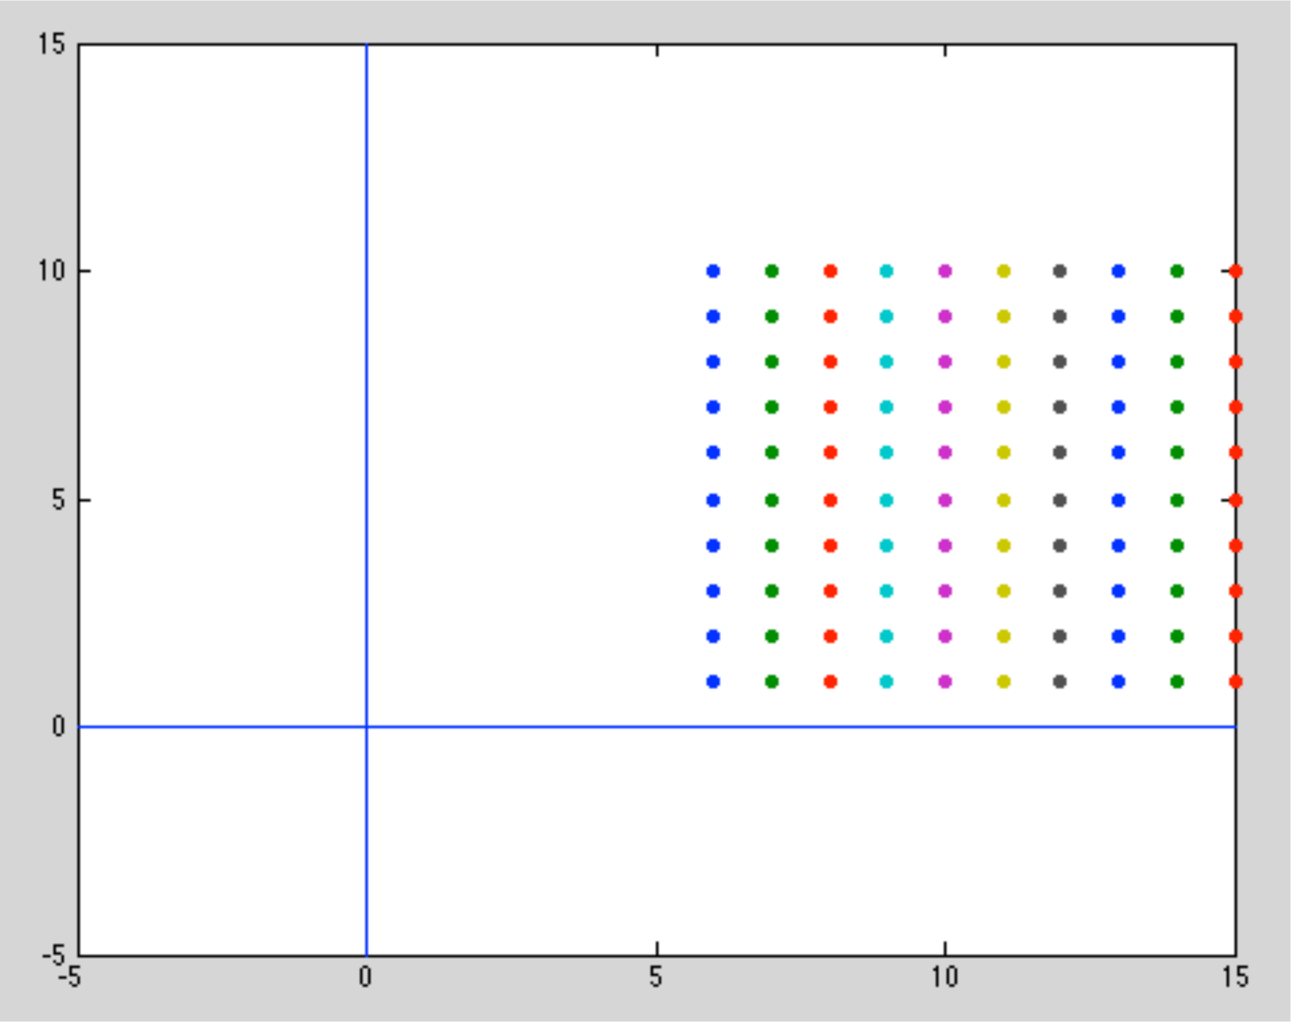

Scaling

```
figure(3);
oxy=[ox(:)'; oy(:)';];
A=[1.6 0; 0 1.6];
nxy=A*oxy;

nx=nxy(1,:);
ny=nxy(2,:);
nx=reshape(nx,[10,10]);
ny=reshape(ny,[10 10]);
plot(nx,ny,'.');
```

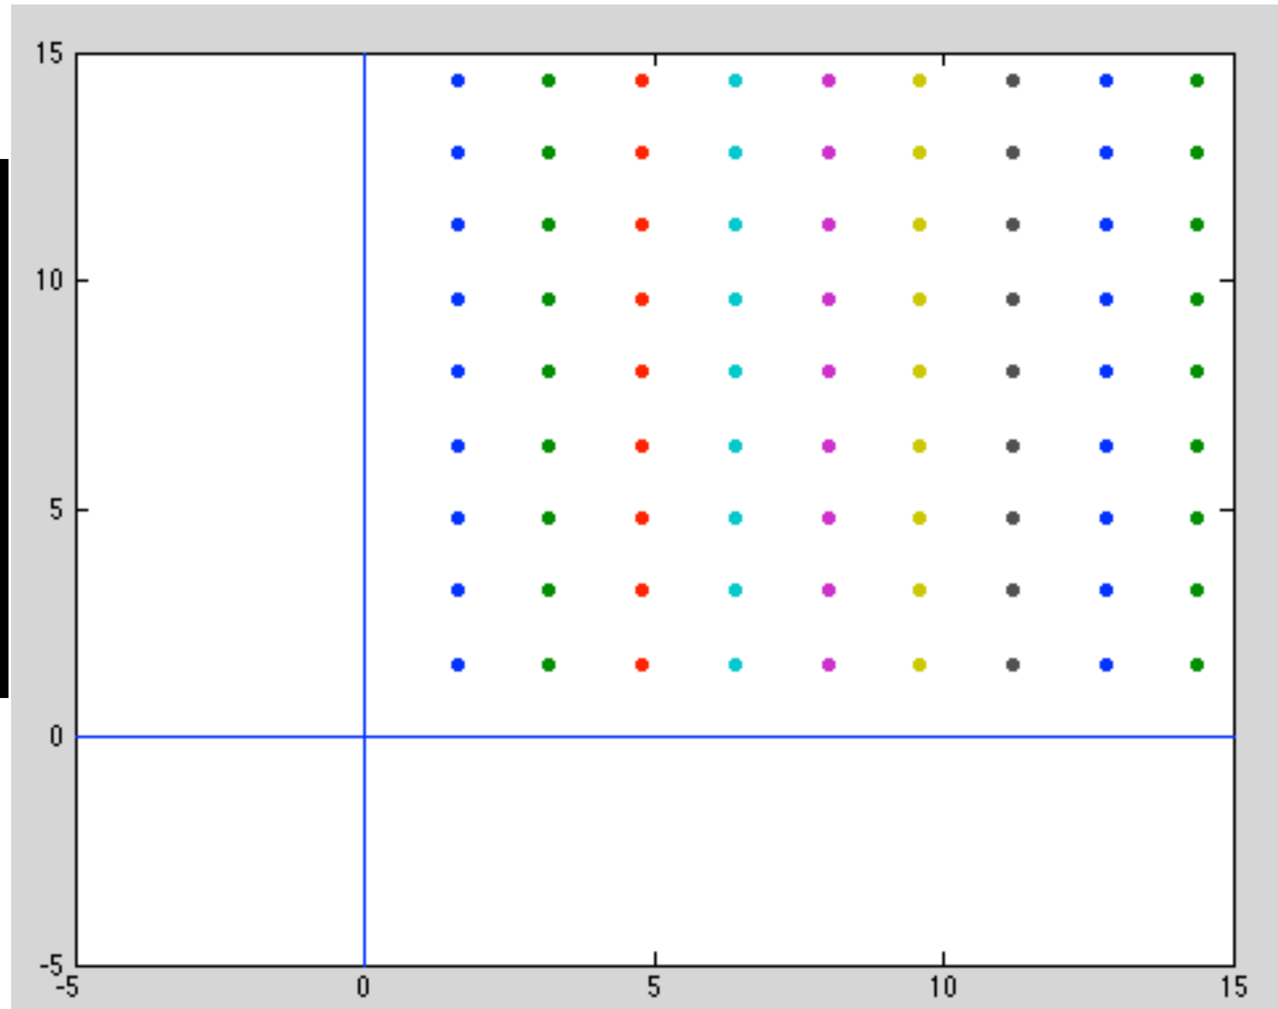

Rotation

```
figure(4);
oxy=[ox(:)'; oy(:)';];
A=[.707 .707; .707 -.707];
nxy=A*oxy;

nx=nxy(1,:);
ny=nxy(2,:);
nx=reshape(nx,[10,10]);
ny=reshape(ny,[10 10]);
plot(nx,ny,'.');
```

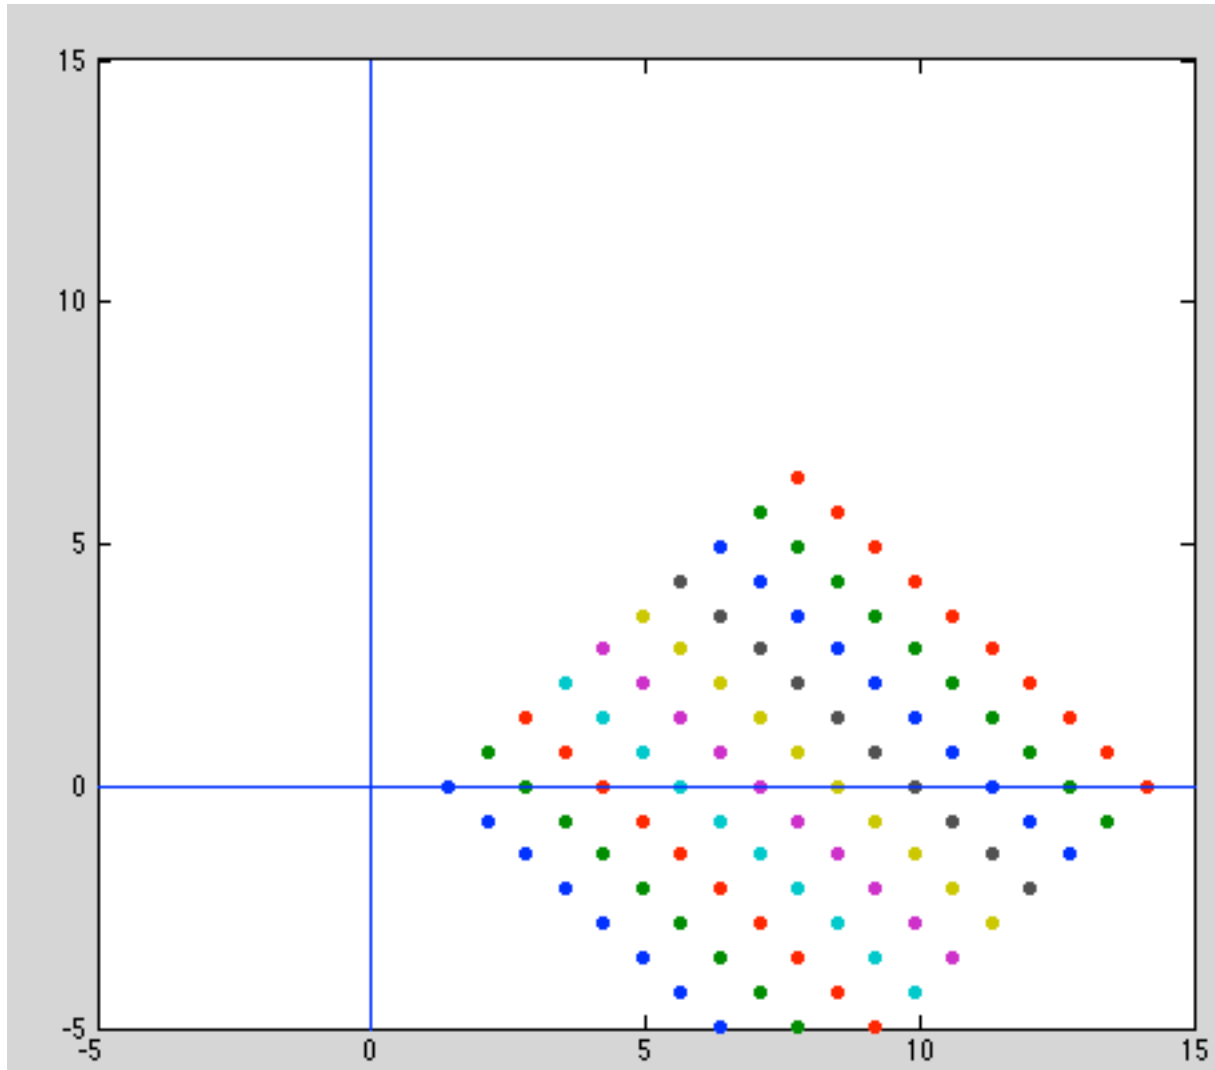

Shear

```
figure(5);
oxy=[ox(:)'; oy(:)';];
A=[1.2 .5; 0 1];
nxy=A*oxy;

nx=nxy(1,:);
ny=nxy(2,:);
nx=reshape(nx,[10,10]);
ny=reshape(ny,[10 10]);
plot(nx,ny,'.');
```

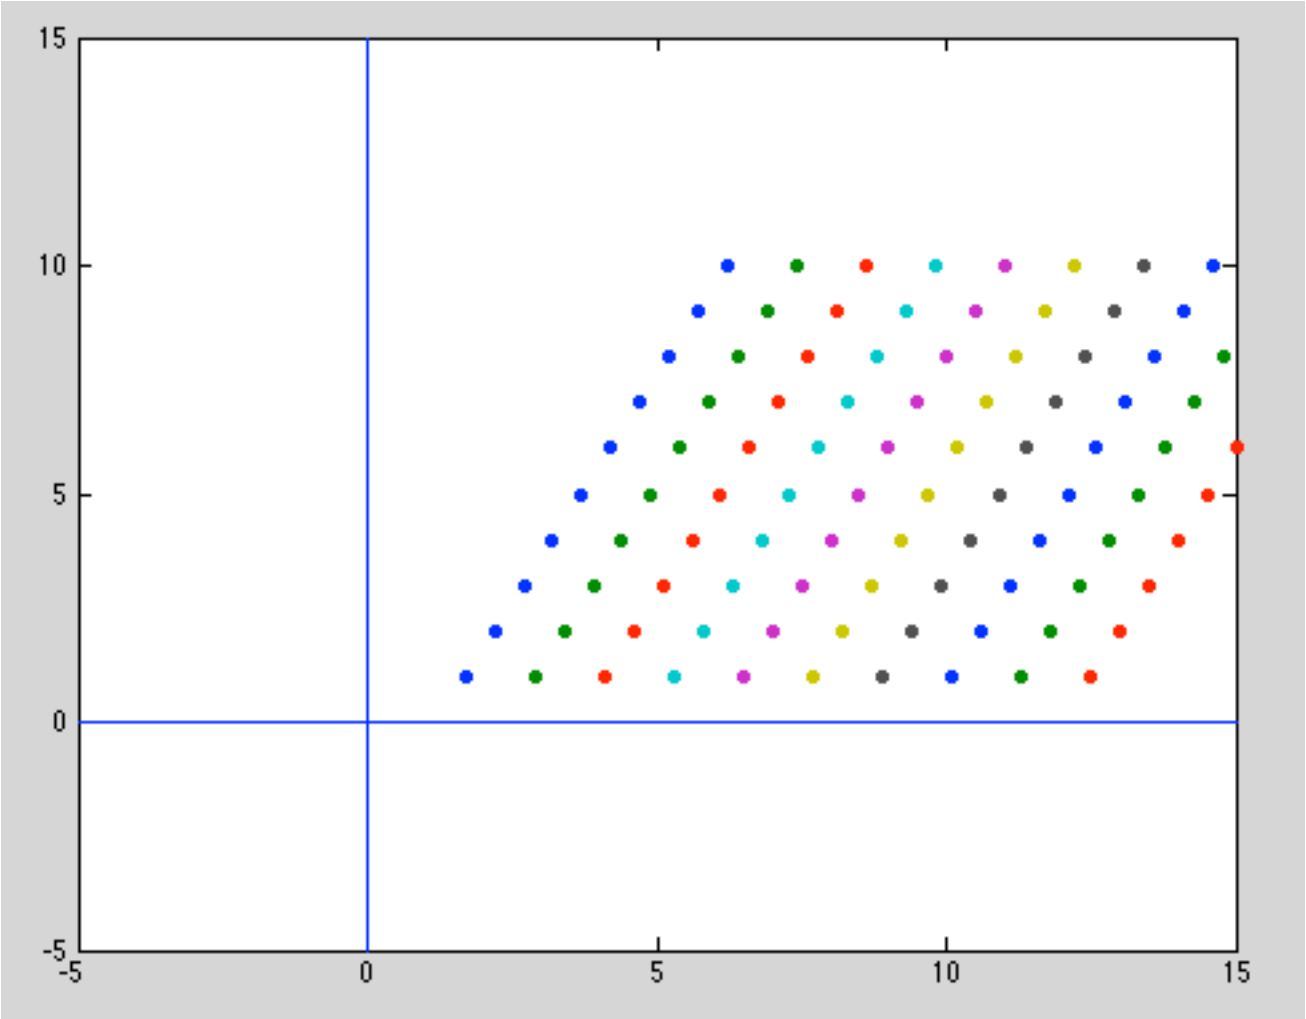

Arbitrary Linear Transformation

```
figure(6);
oxy=[ox(:)'; oy(:)';];
A=[.7 .3; -.5 1.2];
nxy=A*oxy;

nx=nxy(1,:);
ny=nxy(2,:);
nx=reshape(nx,[10,10]);
ny=reshape(ny,[10 10]);
plot(nx,ny,'.');
```

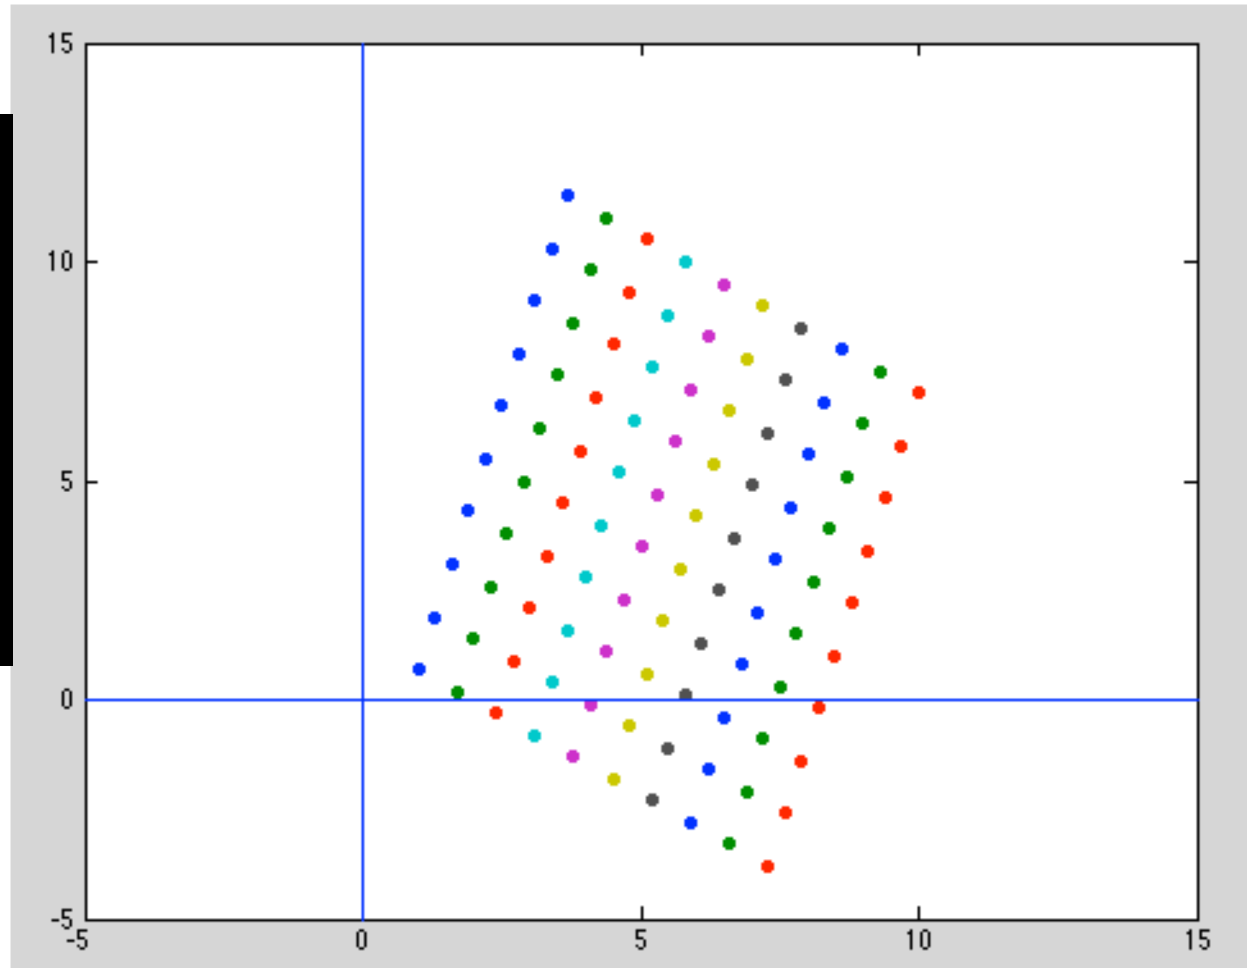

Families of Linear Transformations

Identity:

$$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

Uniform scaling:

$$\begin{bmatrix} s & 0 \\ 0 & s \end{bmatrix}$$

Scaling in x:

$$\begin{bmatrix} s_x & 0 \\ 0 & 1 \end{bmatrix}$$

Rotation by θ radians:

$$\begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix}$$

Shearing in x:

$$\begin{bmatrix} 1 & sh_x \\ 0 & 1 \end{bmatrix}$$

Arbitrary linear transformation:

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix}$$

Families of *Affine* Transformations

Identity:

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

Translation:

$$\begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix}$$

Uniform scaling:

$$\begin{bmatrix} s & 0 & 0 \\ 0 & s & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

Scaling in x:

$$\begin{bmatrix} s_x & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

Rotation by θ radians:

$$\begin{bmatrix} \cos(\theta) & -\sin(\theta) & 0 \\ \sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

Shearing in x:

$$\begin{bmatrix} 1 & sh_x & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

Arbitrary *affine* transformation:

$$\begin{bmatrix} a & b & t_x \\ c & d & t_y \\ 0 & 0 & 1 \end{bmatrix}$$

Mechanics of Transformations: Translations (shifts)

- To shift an image (dx,dy)

  Let (ox,oy) be the coordinates of a pixel in the original image with a particular appearance.

  Let (nx,ny) be the new coordinates, i.e., where we want that pixel.

  For each pixel in old image:

```
nx=ox+dx;        // Compute new pixel pos.
ny=oy+dy;
newIm(ny,nx)=im(oy,ox);
```

Mechanics of Transformations

procedure *forwardWarp*$(f, \boldsymbol{h}, \textbf{out}\ g)$:

For every pixel $\boldsymbol{x}$ in $f(\boldsymbol{x})$

1. Compute the destination location $\boldsymbol{x}' = \boldsymbol{h}(\boldsymbol{x})$.
2. Copy the pixel $f(\boldsymbol{x})$ to $g(\boldsymbol{x}')$.

From Computer Vision: Algorithms and Applications, by Rick Szeliski

Mechanics of Transformations

procedure *forwardWarp*$(f, \boldsymbol{h}, \textbf{out } g)$:

For every pixel $\boldsymbol{x}$ in $f(\boldsymbol{x})$

1. Compute the destination location $\boldsymbol{x}' = \boldsymbol{h}(\boldsymbol{x})$.
2. Copy the pixel $f(\boldsymbol{x})$ to $g(\boldsymbol{x}')$.

Problems:

- Leaves gaps in destination image.
- Interpolation is less intuitive.

Example of Forward Warp (rotation by 45 degrees)

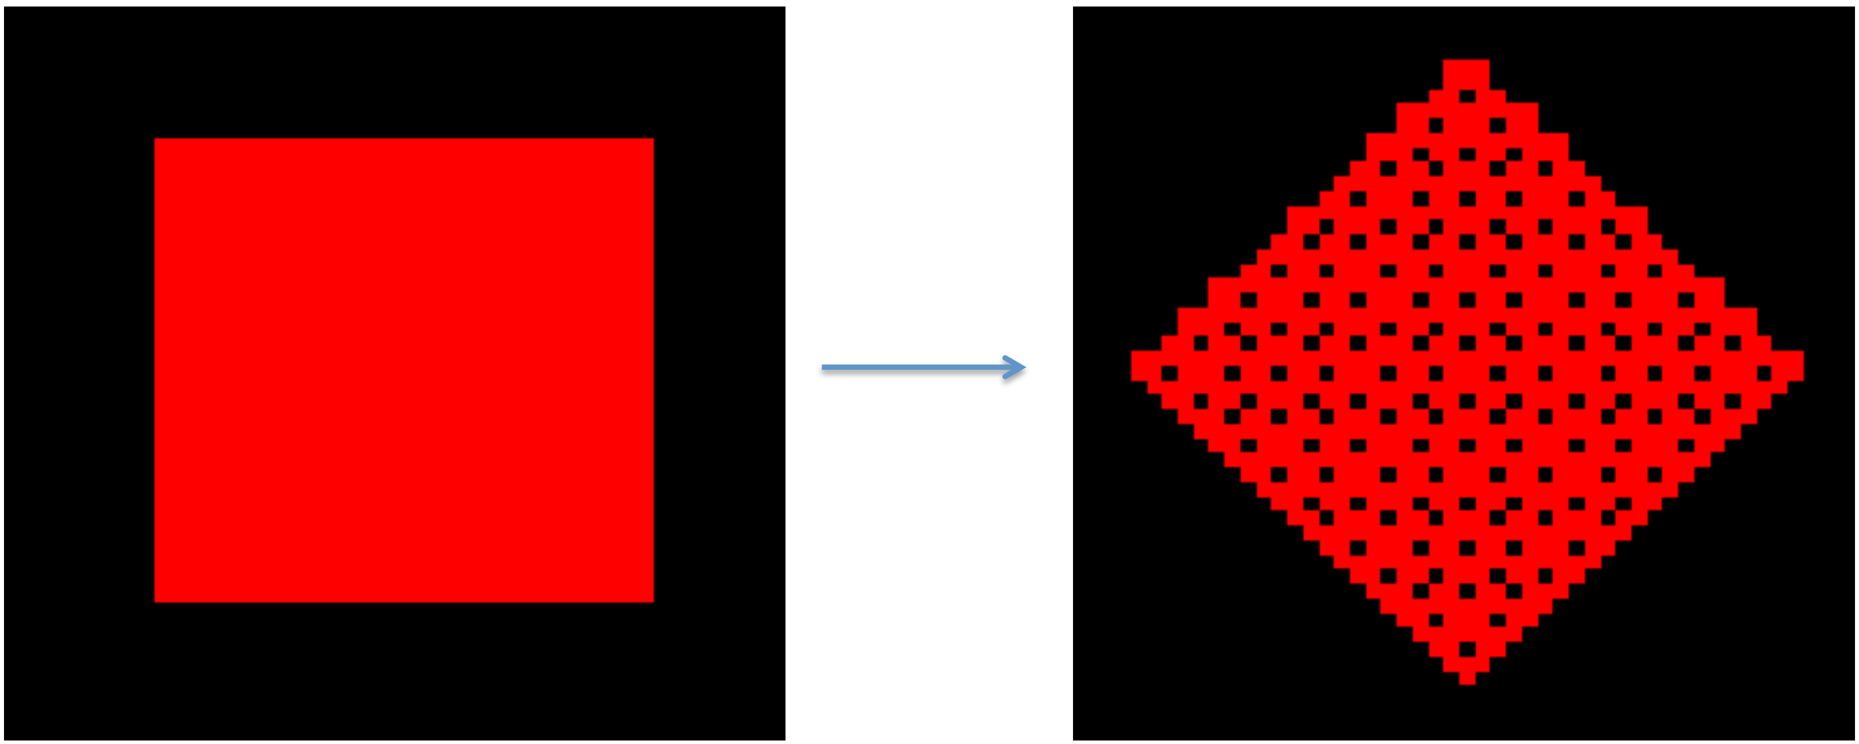

Mechanics of Transformations

procedure *inverseWarp*$(f, \boldsymbol{h}, \textbf{out } g)$:

For every pixel $\boldsymbol{x}'$ in $g(\boldsymbol{x}')$

1. Compute the source location $\boldsymbol{x} = \hat{\boldsymbol{h}}(\boldsymbol{x}')$
2. Resample $f(\boldsymbol{x})$ at location $\boldsymbol{x}$ and copy to $g(\boldsymbol{x}')$

From Computer Vision: Algorithms and Applications, by Rick Szeliski

Example of Reverse Warp

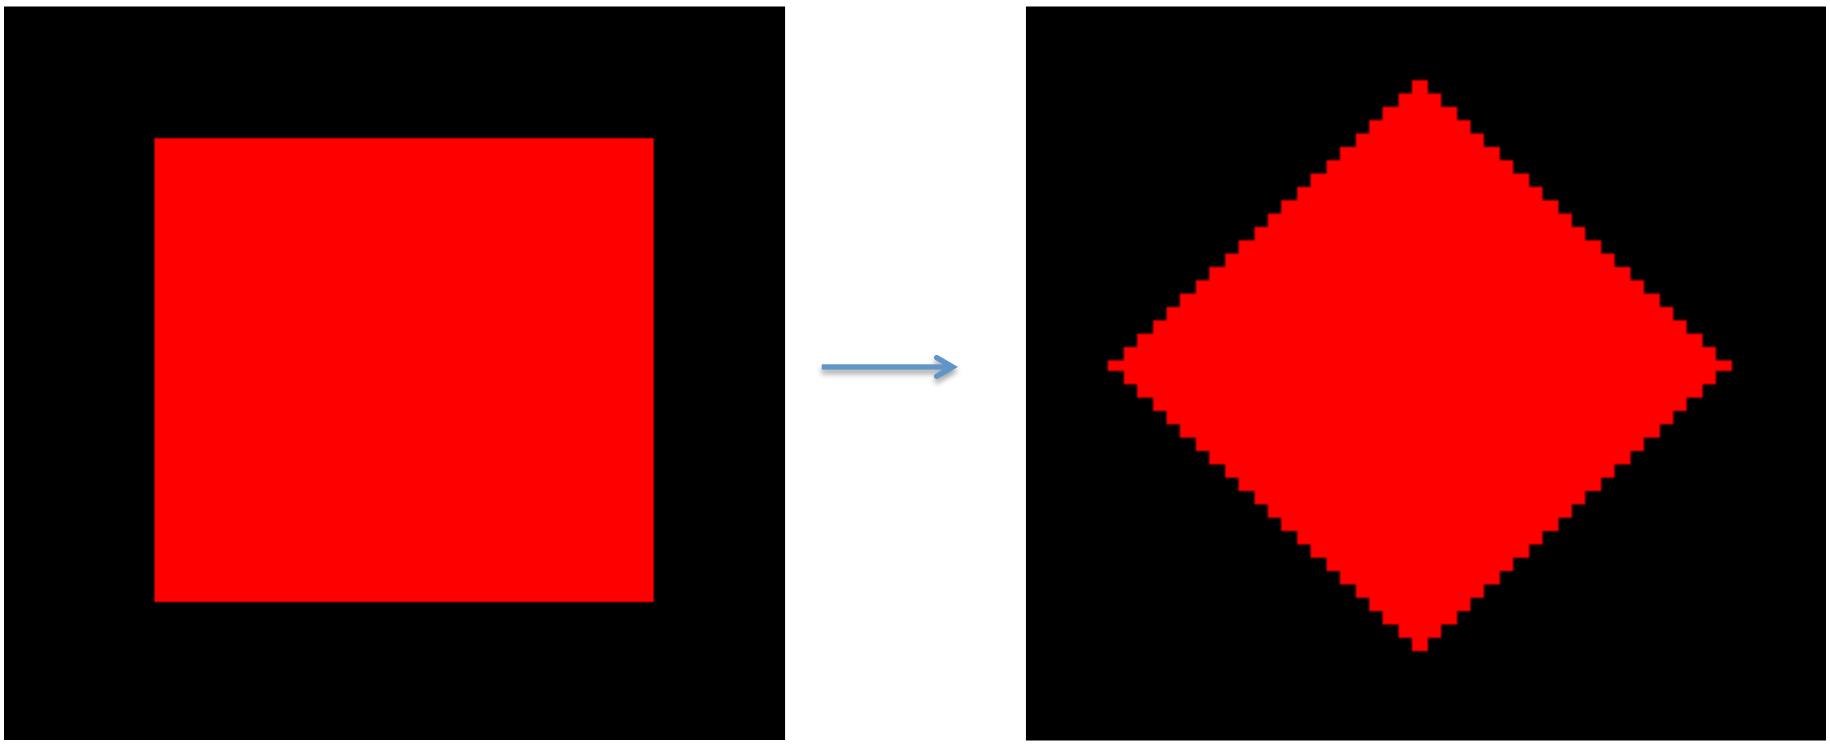

Summary

- Basic elements of alignment
  - Representation
  - Alignment criterion
  - Method of optimization
- Mutual information alignment
  - Criterion address problems of aligning images from different modalities
  - Why not always use mutual information alignment?
    - Chess board example.