

Image Warping



<http://www.jeffrey-martin.com>

Image Warping (a.k.a. Domain Transforms)

- Parametric transformations
 - Linear transformations of images via 2×2 matrices
(a crash course on basic linear algebra)
 - Affine transformations
 - Homographies (3×3 transformation matrices)
- Estimation of parametric transformations (from corresponding points)
- Forward and inverse warps
 - bilinear interpolation

Image Warping

point processing: change **range** of image $g(x) = T(f(x))$




image warping: change **domain** of image $g(x) = f(T(x))$




Image Warping

point processing: change **range** of image $g(x) = T(f(x))$





image warping: change **domain** of image $g(x) = f(T(x))$



iClicker moment: Can image warping change intensity histogram?

A: Yes

B: No

Parametric (global) warping

Examples of parametric warps:



translation



rotation



aspect



affine



perspective




cylindrical

Parametric (global) warping



$$\mathbf{p} = (x, y)$$



$$\mathbf{p}' = (x', y')$$

Transformation T is a coordinate-changing machine:

What does it mean that T is global?

- the same transform for any point p
- described by just a few numbers (parameters)

Example: linear transforms T (represented by matrix \mathbf{M}): $\mathbf{p}' = \mathbf{M}\mathbf{p}$

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \mathbf{M} \begin{bmatrix} x \\ y \end{bmatrix}$$


$$g(p') = \boxed{g(\mathbf{M} \cdot \mathbf{p}) = f(\mathbf{p})} = f(\mathbf{M}^{-1} \cdot \mathbf{p}')$$

inverse warp (slide 56) forward warp (slide 54)

Scaling


Scaling a coordinate means multiplying each of its components by a scalar

Uniform scaling means this scalar is the same for all components:



Scaling

Non-uniform scaling: different scalars per component:



Scaling

Scaling operation:

$$x' = ax$$


$$y' = by$$

Or, in matrix form:

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \underbrace{\begin{bmatrix} a & 0 \\ 0 & b \end{bmatrix}}_{\text{scaling matrix } S} \begin{bmatrix} x \\ y \end{bmatrix}$$


What's inverse of S?

2-D Rotation (around coordinate center)



$$\begin{aligned}x' &= x \cos(\theta) - y \sin(\theta) \\y' &= x \sin(\theta) + y \cos(\theta)\end{aligned}$$

2-D Rotation (if you do not remember derivation)



$$x = r \cos(\phi)$$

$$y = r \sin(\phi)$$

$$x' = r \cos(\phi + \theta)$$

$$y' = r \sin(\phi + \theta)$$

Trig Identity...

$$x' = r \cos(\phi) \cos(\theta) - r \sin(\phi) \sin(\theta)$$

$$y' = r \cos(\phi) \sin(\theta) + r \sin(\phi) \cos(\theta)$$

Substitute...

$$x' = x \cos(\theta) - y \sin(\theta)$$

$$y' = x \sin(\theta) + y \cos(\theta)$$

2-D Rotation

This is easy to capture in matrix form:

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \underbrace{\begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix}}_{\mathbf{R}} \begin{bmatrix} x \\ y \end{bmatrix}$$

Even though $\sin(\theta)$ and $\cos(\theta)$ are nonlinear functions of θ ,

- **x' is a linear combination of x and y**
- **y' is a linear combination of x and y**

What is the inverse transformation?

- Rotation by $-\theta$
- For rotation matrices $\mathbf{R}^{-1} = \mathbf{R}^T$

2x2 Matrices

What types of transformations can be represented with a 2x2 matrix?

2D Identity?

$$\begin{aligned} x' &= x \\ y' &= y \end{aligned} \quad \begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$

2D Scale around (0,0)?

$$\begin{aligned} x' &= s_x * x \\ y' &= s_y * y \end{aligned} \quad \begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} s_x & 0 \\ 0 & s_y \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$

2x2 Matrices

What types of transformations can be represented with a 2x2 matrix?

2D Rotate around (0,0)?

$$\begin{aligned}x' &= \cos \Theta * x - \sin \Theta * y \\y' &= \sin \Theta * x + \cos \Theta * y\end{aligned}$$

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} \cos \Theta & -\sin \Theta \\ \sin \Theta & \cos \Theta \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$

2D Shear?

$$\begin{aligned}x' &= x + sh_x * y \\y' &= y\end{aligned}$$

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} 1 & sh_x \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$

2x2 Matrices

What types of transformations can be represented with a 2x2 matrix?

2D Mirror about Y axis?

$$\begin{aligned}x' &= -x \\ y' &= y\end{aligned}$$

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$

2D Mirror over (0,0)?

$$\begin{aligned}x' &= -x \\ y' &= -y\end{aligned}$$

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} -1 & 0 \\ 0 & -1 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$

2x2 Matrices

What types of transformations can be represented with a 2x2 matrix?

2D Translation?

$$\begin{aligned} \mathbf{x}' &= \mathbf{x} + \mathbf{t}_x \\ \mathbf{y}' &= \mathbf{y} + \mathbf{t}_y \end{aligned} \quad \text{NO!}$$

origin maps to origin

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

(for any M)

All 2D Linear Transformations

Linear transformations are combinations of ...

- Scale,
- Rotation,
- Shear, and
- Mirror

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$

iClicker moment:

Can a (non-degenerate) linear warp change image intensity histogram?

A: Yes

B: No

Properties of linear transformations:

- Origin maps to origin
- Lines map to lines
- Parallel lines remain parallel
- Distance or length ratios are preserved **on parallel lines**
 - scaling of length/distances depends on (line) orientation only (see next slide)
- Ratios of areas are preserved
- Closed under composition

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} i & j \\ k & l \end{bmatrix} \begin{bmatrix} s & q \\ r & t \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$

see pp. 40-41 of Hartley and Zisserman “Multiple View Geometry” (2nd edition)



All 2D Linear Transformations

Decomposition into basic geometric transformations (follows from SVD for 2x2 matrices)

$$A = \begin{bmatrix} a & b \\ c & d \end{bmatrix} = R_{\theta} \cdot \left(R_{-\Phi} \cdot \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} \cdot R_{\Phi} \right)$$

rotation
by angle θ


deformation
(non-isotropic scaling & reflection)



Scaling directions are always orthogonal.

Areas are scaled (homogeneously over a plane)
by a factor of $\det A = |\lambda_1 \lambda_2|$

Linear Transformation as **Space Deformation**




$$\begin{array}{ccc} \mathbf{q} & T & \mathbf{p} \\ \left[\begin{array}{c} q_x \\ q_y \end{array} \right] = \left[\begin{array}{cc} u_x & v_x \\ u_y & v_y \end{array} \right] \left[\begin{array}{c} 4 \\ 3 \end{array} \right] & & \end{array}$$

coordinates of \mathbf{q} in basis \mathbf{i}, \mathbf{j} coordinates of \mathbf{p} in basis \mathbf{i}, \mathbf{j}

$\mathbf{q} = q_x \mathbf{i} + q_y \mathbf{j}$ $\mathbf{p} = 4\mathbf{i} + 3\mathbf{j}$

point \mathbf{p} is transformed into new point \mathbf{q}
fixed basis, point transformed (coordinates change)

Linear Transformation as Change of Basis



now interpret the columns of matrix T
as some vectors \mathbf{u} and \mathbf{v} (their coordinates in basis \mathbf{i}, \mathbf{j})


coordinates of \mathbf{q} in basis \mathbf{i}, \mathbf{j} $\mathbf{q} = q_x \mathbf{i} + q_y \mathbf{j}$

$$\begin{bmatrix} q_x \\ q_y \end{bmatrix} = \begin{bmatrix} u_x & v_x \\ u_y & v_y \end{bmatrix} \begin{bmatrix} 4 \\ 3 \end{bmatrix}$$

coordinates of \mathbf{q} in basis \mathbf{u}, \mathbf{v}
 $\mathbf{q} = 4\mathbf{u} + 3\mathbf{v}$

Indeed, $\mathbf{q} = 4 \cdot (\overbrace{u_x \cdot \mathbf{i} + u_y \cdot \mathbf{j}}^{\mathbf{u}}) + 3 \cdot (\overbrace{v_x \cdot \mathbf{i} + v_y \cdot \mathbf{j}}^{\mathbf{v}}) = (\overbrace{4 \cdot u_x + 3 \cdot v_x}^{q_x}) \cdot \mathbf{i} + (\overbrace{4 \cdot u_y + 3 \cdot v_y}^{q_y}) \cdot \mathbf{j}$

Linear Transformation as Change of Basis



now interpret the columns of matrix T
as some vectors \mathbf{u} and \mathbf{v} (their coordinates in basis \mathbf{i}, \mathbf{j})


coordinates of \mathbf{q} in basis \mathbf{i}, \mathbf{j} $\mathbf{q} = q_x \mathbf{i} + q_y \mathbf{j}$

$$\begin{bmatrix} q_x \\ q_y \end{bmatrix} = \begin{bmatrix} u_x & v_x \\ u_y & v_y \end{bmatrix} \begin{bmatrix} 4 \\ 3 \end{bmatrix}$$

coordinates of \mathbf{q} in basis \mathbf{u}, \mathbf{v}
 $\mathbf{q} = 4\mathbf{u} + 3\mathbf{v}$

point \mathbf{q} represented in different coordinate systems
basis transformed, fixed points (coordinates still change)

Linear Transformation as Change of Basis



now interpret the columns of matrix T
as some vectors \mathbf{u} and \mathbf{v} (their coordinates in basis \mathbf{i}, \mathbf{j})


$$\begin{bmatrix} q_x \\ q_y \end{bmatrix} = \begin{bmatrix} u_x & v_x \\ u_y & v_y \end{bmatrix} \begin{bmatrix} 4 \\ 3 \end{bmatrix}$$

Any matrix can be seen as a (linear) coordinate system basis!!!

Question: What's the inverse matrix T^{-1} ?

$$\begin{bmatrix} 4 \\ 3 \end{bmatrix} = \begin{bmatrix} ? & ? \\ ? & ? \end{bmatrix} \begin{bmatrix} q_x \\ q_y \end{bmatrix}$$

Linear Transformation as Change of Basis



now interpret the columns of matrix T
as some vectors \mathbf{u} and \mathbf{v} (their coordinates in basis \mathbf{i}, \mathbf{j})

$$\begin{bmatrix} q_x \\ q_y \end{bmatrix} = \begin{bmatrix} u_x & v_x \\ u_y & v_y \end{bmatrix} \begin{bmatrix} 4 \\ 3 \end{bmatrix}$$

Any matrix can be seen as a (linear) coordinate system basis!!!

Question: What's the inverse matrix T^{-1} ?


$$\begin{bmatrix} 4 \\ 3 \end{bmatrix} = \begin{bmatrix} i_x & j_x \\ i_y & j_y \end{bmatrix} \begin{bmatrix} q_x \\ q_y \end{bmatrix}$$

coordinates of \mathbf{i} and \mathbf{j}
in basis \mathbf{u}, \mathbf{v}

$$\mathbf{i} = i_x \cdot \mathbf{u} + i_y \cdot \mathbf{v}$$

$$\mathbf{j} = j_x \cdot \mathbf{u} + j_y \cdot \mathbf{v}$$


Linear Transformation as **Change of Basis**



Any matrix can be seen as a (linear) coordinate system basis!!!

Question: What is T if both coordinate systems have **ortho-normal basis**?

Linear Transformation as Change of Basis



Any matrix can be seen as a (linear) coordinate system basis!!!

Question: What is T if both coordinate systems have **ortho-normal basis**?

Then matrix T represents rotation, reflection, or their combination (rotoreflexion) of the coordinate basis

Towards Homogeneous Coordinates

Q: Can we represent translation by matrix multiplication?

$$\mathbf{x}' = \mathbf{x} + \mathbf{t}_x$$

$$\mathbf{y}' = \mathbf{y} + \mathbf{t}_y$$

very simple, but
not a *linear* transformation in 2D

$$T(p+q) \neq T(p) + T(q)$$

$$T(\lambda p) \neq \lambda T(p)$$

Answer: Yes, using **homogeneous coordinates** and **3x3 matrices**

Homogeneous coordinates

- represent coordinates in 2 dimensions with a 3-vector

$$\begin{bmatrix} x \\ y \end{bmatrix} \xrightarrow{\text{homogeneous coordinates}} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} x + t_x \\ y + t_y \\ 1 \end{bmatrix}$$

Translation
matrix (3x3)


Translation

Example of translation


Homogeneous Coordinates



$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} x + t_x \\ y + t_y \\ 1 \end{bmatrix}$$



$$\begin{aligned} t_x &= 2 \\ t_y &= 1 \end{aligned}$$




Homogeneous Coordinates (in general)

Add a 3rd coordinate to every 2D point

- (x, y, w) represents a point at location $(x/w, y/w)$
- $(0, 0, 0)$ is not allowed

Advantages of
homogeneous
coordinate system:

- simple matrix representation of
many useful transformations



$(w \cdot 2, w \cdot 1, w)$
represent the same 2D point
for any value of w


Homogeneous Coordinates (in general)

Add a 3rd coordinate to every 2D point

- (x, y, w) represents a point at location $(x/w, y/w)$
- $(0, 0, 0)$ is not allowed
- $(x, y, 0)$ represents a *point at infinity*

Advantages of
homogeneous
coordinate system:

- simple matrix representation of many useful transformations
- allows to expand R^2 with “*points at infinity*” using finite numerical representation (like $\pm\infty$ for R^1)



Basic 2D Transformations via 3x3 matrices

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

Translate

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} s_x & 0 & 0 \\ 0 & s_y & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

Scale

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} \cos \Theta & -\sin \Theta & 0 \\ \sin \Theta & \cos \Theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

Rotate

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & sh_x & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

Shear

all of the above are special cases of
a general **Affine Transformation**:

$$\begin{bmatrix} x' \\ y' \\ w \end{bmatrix} = \begin{bmatrix} a & b & c \\ d & e & f \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w \end{bmatrix}$$

Composing Affine Transformations

Example:

$$\begin{bmatrix} a & b & c \\ d & e & f \\ 0 & 0 & 1 \end{bmatrix}$$

composition of any affine
transforms is still affine

(as easy to check)

$$\begin{bmatrix} x' \\ y' \\ w' \end{bmatrix} = \left(\begin{bmatrix} 1 & 0 & tx \\ 0 & 1 & ty \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \cos\Theta & -\sin\Theta & 0 \\ \sin\Theta & \cos\Theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} sx & 0 & 0 \\ 0 & sy & 0 \\ 0 & 0 & 1 \end{bmatrix} \right) \begin{bmatrix} x \\ y \\ w \end{bmatrix}$$

$\mathbf{p}' = T(t_x, t_y)$

$R(\Theta)$

$S(s_x, s_y)$

\mathbf{p}

In general: any affine transformation is a combination of translation, rotation/reflection, and anisotropic scaling

Affine Transformations

Affine transformations are combinations of ...

- Linear 2D transformations, and
- Translations

$$\begin{bmatrix} x' \\ y' \\ w \end{bmatrix} = \begin{bmatrix} a & b & c \\ d & e & f \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w \end{bmatrix}$$

Properties of affine transformations:

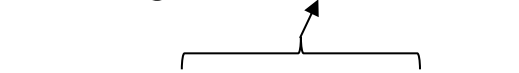
- Origin does not necessarily map to origin (**new** compared to 2x2 matrices)
- Lines map to lines
- Parallel lines remain parallel
- Length/distance ratios are preserved on parallel lines
- Ratios of areas are preserved
- Closed under composition

Projective Transformations (a.k.a. *homographies*)

transformations in homogeneous coordinate space via general 3x3 matrices

Projective transformations ...

- Affine transformations, and
- Projective warps

$$\begin{bmatrix} x' \\ y' \\ w' \end{bmatrix} = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$



Properties of projective transformations:

- Origin does not necessarily map to origin
- Lines map to lines (indeed, line of hom. points \mathbf{p} means $\mathbf{a} \cdot \mathbf{p} = 0$ for some \mathbf{a} . Then, $\mathbf{b} \cdot \mathbf{H}\mathbf{p} = 0$ for $\mathbf{b} = \mathbf{a}\mathbf{H}^{-1}$)
- Parallel lines do not necessarily remain parallel
- Non-parallel lines may become parallel
- Distance/length or area ratios are not preserved
- Closed under composition

Projective Transformations (a.k.a. *homographies*)

- Parallel lines do not necessarily remain parallel
- Non-parallel lines may become parallel

$$\begin{bmatrix} 3 \\ 2 \\ 0 \end{bmatrix} = \overbrace{\begin{bmatrix} a & b & c \\ d & e & f \\ -1 & 1 & 1 \end{bmatrix}}^H \begin{bmatrix} 6 \\ 5 \\ 1 \end{bmatrix}$$




NOTE: “finite” point may transform to “point at infinity”

Projective Transformations (a.k.a. *homographies*)

- Parallel lines do not necessarily remain parallel
- Non-parallel lines may become parallel

$$\begin{bmatrix} 12 \\ 10 \\ 2 \end{bmatrix} = \overbrace{\begin{bmatrix} a' & b' & c' \\ d' & e' & f' \\ g' & h' & i' \end{bmatrix}}^{H^{-1}} \begin{bmatrix} 3 \\ 2 \\ 0 \end{bmatrix}$$




NOTE: or “point at infinity” may transform to “finite” point

Projective Transformations (a.k.a. *homographies*)

- Distance/length or area ratios are not preserved

$$\begin{bmatrix} 3 \\ 2 \\ 0 \end{bmatrix} = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix} \begin{bmatrix} 6 \\ 5 \\ 1 \end{bmatrix}$$



Example: distance $|B'C'|$ remains finite, while $|A'B'|$ is infinite

Projective Transformations (a.k.a. *homographies*)


General property to keep in mind (Theorem 2.10 from Hartley&Zisserman)

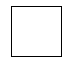
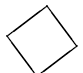


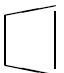
An invertible mapping h from a (homogeneous) plane P^2 onto P^2 preserves straight lines **iff** there exists a non-singular 3×3 matrix H s.t.

$$h(x) = H \cdot x \quad \text{for any } x \in P^2$$

That is, any transformation of a plane onto a plane that preserves straight lines must be a *homography*.

2D image transformations



Name	Matrix	# D.O.F.	Preserves:	Icon
translation	$\begin{bmatrix} \mathbf{I} & \mathbf{t} \end{bmatrix}_{2 \times 3}$	2	orientation + ...	
rigid (Euclidean)	$\begin{bmatrix} \mathbf{R} & \mathbf{t} \end{bmatrix}_{2 \times 3}$	3	lengths + ...	
similarity	$\begin{bmatrix} s\mathbf{R} & \mathbf{t} \end{bmatrix}_{2 \times 3}$	4	angles + ...	
affine	$\begin{bmatrix} \mathbf{A} \end{bmatrix}_{2 \times 3}$	6	parallelism + ...	
projective	$\begin{bmatrix} \tilde{\mathbf{H}} \end{bmatrix}_{3 \times 3}$	8	straight lines	


See Hartley and
Zisserman,
p. 44

These transformations are a nested set of groups

- Closed under composition and inverse is a member

iClicker moment

Q: What best describes the **transformation between two monsters** in this image?



A: translation


B: translation + scale

C: projective

Remaining parts of this lecture

- Estimation of parametric transformations (from corresponding points)
- Forward and inverse warps

Recovering Parametric Transformations




What if we know f and g and want to recover transform T ?

- e.g. to better align images (**image registration**)
- willing to let user provide correspondences

Q: How many pairs of corresponding points do we need?

Translation: # correspondences?



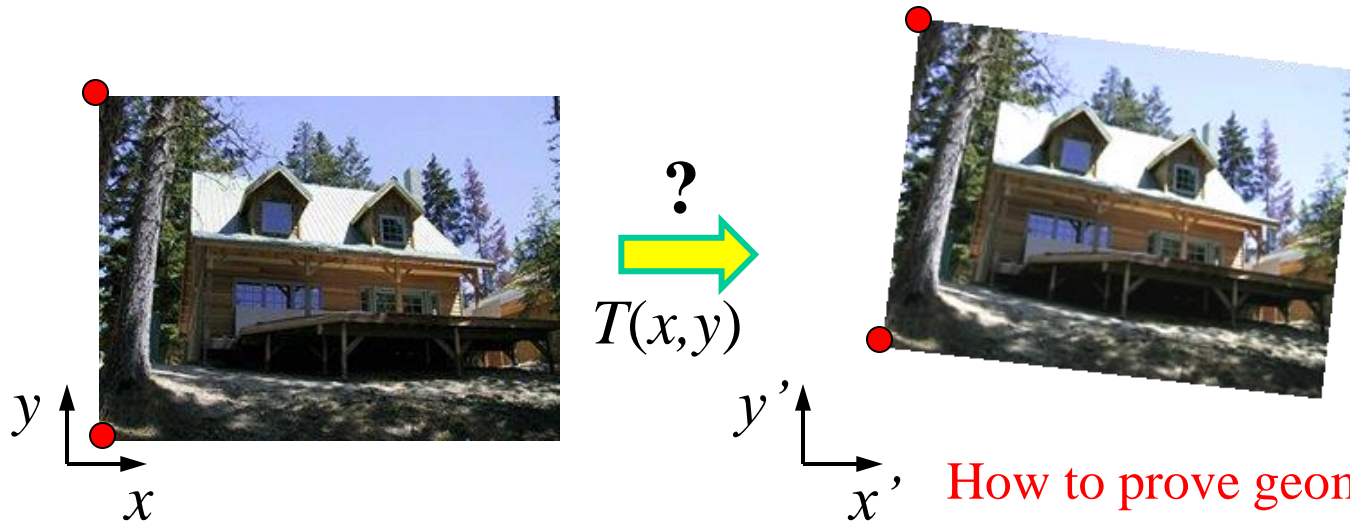
How many correspondences needed for translation?

How many Degrees of Freedom (DOF)?

What is the transformation matrix?

$$\mathbf{M} = \begin{bmatrix} 1 & 0 & c_x \\ 0 & 1 & c_y \\ 0 & 0 & 1 \end{bmatrix}$$

Euclidian: # correspondences?



How to prove geometrically
that 2 pairs is enough?

(use rigid transformation invariants
to map an arbitrary point)


How many correspondences needed for translation+rotation?

How many DOF?

Transformation matrix?

$$\mathbf{M} = \begin{bmatrix} \cos \theta & -\sin \theta & c_x \\ \sin \theta & \cos \theta & c_y \\ 0 & 0 & 1 \end{bmatrix}$$

Affine: # correspondences?



How many correspondences needed for affine?

How many DOF?

Transformation matrix?

$$\mathbf{M} = \begin{bmatrix} a & b & c \\ c & d & f \\ 0 & 0 & 1 \end{bmatrix}$$

Algebraic point of view

$$\mathbf{p}'_i = M \mathbf{p}_i \quad \begin{matrix} \mathbf{p}'_i & M & \mathbf{p}_i \\ \left[\begin{array}{c} x'_i \\ y'_i \\ 1 \end{array} \right] & = & \left[\begin{array}{ccc} a & b & c \\ d & e & f \\ 0 & 0 & 1 \end{array} \right] \left[\begin{array}{c} x_i \\ y_i \\ 1 \end{array} \right] \end{matrix}$$

for any given pair of
corresponding points

$(\mathbf{p}_i, \mathbf{p}'_i)$

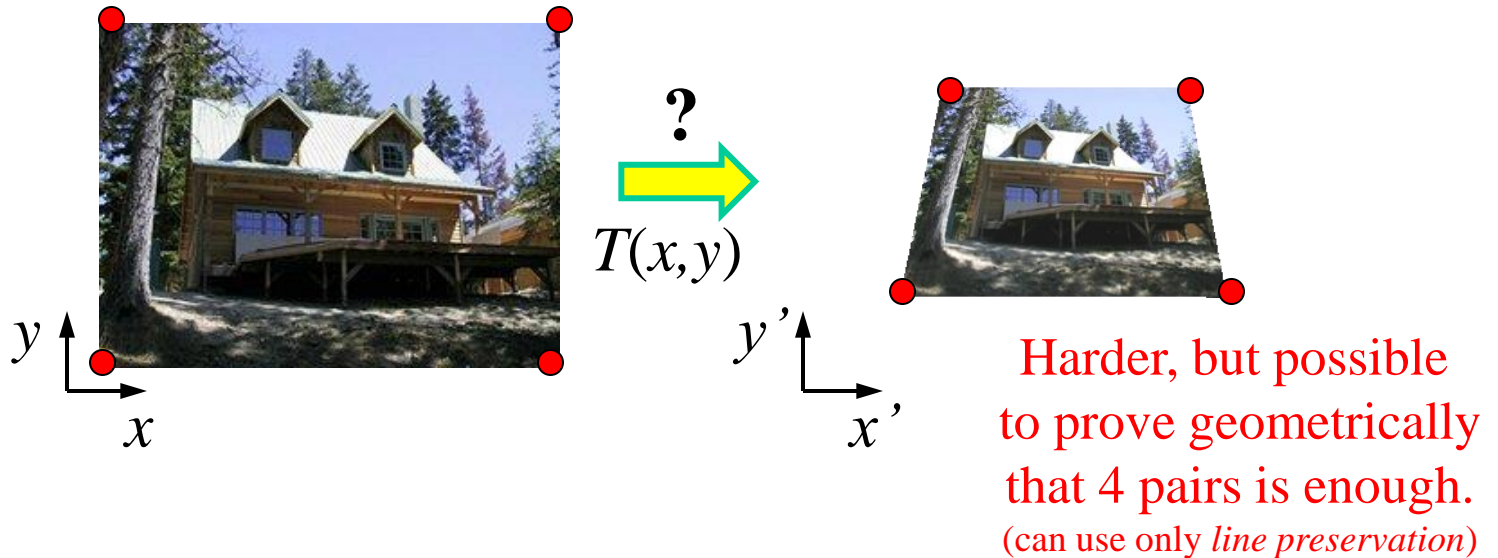
$$\Rightarrow \begin{cases} x'_i = ax_i + by_i + c \\ y'_i = dx_i + ey_i + f \end{cases}$$

**6 unknown
parameters
(variables)**

Each pair of corresponding points $(\mathbf{p}_i, \mathbf{p}'_i)$ gives
two linear equations w.r.t 6 unknown coefficients of matrix M
with known point coordinates for \mathbf{p}_i and \mathbf{p}'_i

3 pairs of corresponding points give $3 \times 2 (=6)$ linear equations
allowing to **resolve 6 unknown parameters**

Projective: # correspondences?



How many correspondences needed for projective?


How many DOF?

Transformation matrix?

Projective: # correspondences?

4 matches is enough to map all **other points**
(*informal* geometric proof based on line preservation)


Optional
slides



Projective: # correspondences?

4 matches is enough to map all other points
(*informal* geometric proof based on line preservation)

Optional
slides




Projective: # correspondences?


4 matches is enough to map all **other points**
(*informal* geometric proof based on line preservation)

Optional
slides

Keep **recursively** subdividing quadrilaterals A, B, C, D
into smaller quadrilaterals while computing more matching
pairs of points and gradually increasing their density



Projective: # correspondences?



How many correspondences needed for projective? **=4**

How many DOF?

Easy to check that 4 pairs give only $4 \times 2 (=8)$ equations!
What about 9 unknowns?

Transformation matrix?





Homographies have only 8 DOF since scale is irrelevant
(multiplying M by any factor does not change the actual transformation).

More on estimating homographies
from 4 matching pairs of point - later in Topic 5.

$$\mathbf{M} = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix}$$

Example: warping triangles

(e.g. “*texture mapping*” for deformable mesh models)



Given two triangles: ABC and A'B'C' in 2D (3 corresponding pairs)

Need to find a **simple parametric transform T** to transfer all pixels from one to the other ?

Common answer: **affine**


(solve 6 linear equations with 6 unknowns)

$$\mathbf{M} = \begin{bmatrix} a & b & c \\ c & d & f \\ 0 & 0 & 1 \end{bmatrix}$$

NOTE: Mesh has many triangles, and each may get a different affine transformation. **The overall (continuous) transformation of the mesh is piecewise-affine with six parameters per triangle.**

Image warping

assume a given transform T , e.g. rotation or projection

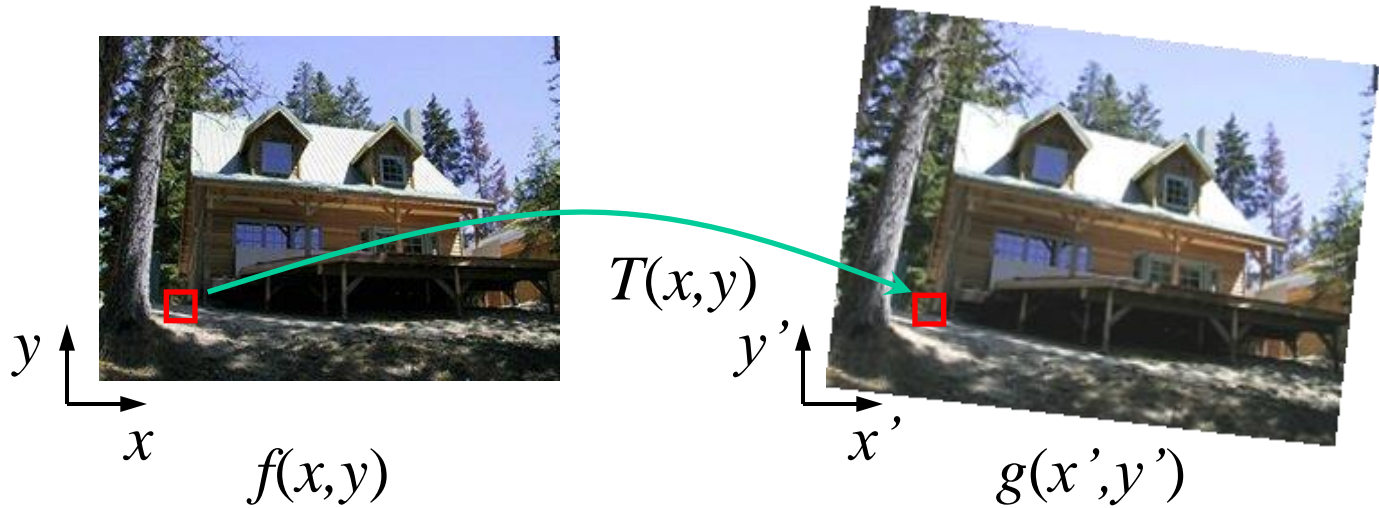


NOTE: in practice,
one should consider
the **canvas bounds**
of the new image

How to generate the transformed image g ?

- e.g.
- **panorama stitching** (next topic)
 - **texture mapping** (3D reconstruction)
 - **novel view generation** (special effects, virtual/augmented reality)
 - **data augmentation** (network training)

Image warping




COMMENT: for simplicity, the slides ignore the bounds of the new image's canvas, but you cannot do so in the homework assignments.

Given a coordinate transform $(x',y') = T(x,y)$ and a source image $f(x,y)$, how do we compute a transformed image $g(x',y') = f(x,y)$?

I. *forward warping*

II. *inverse warping*

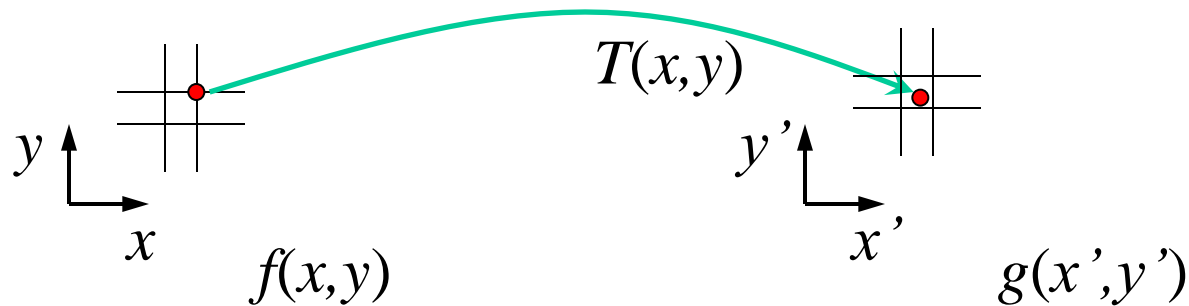
Forward warping



Traverse the pixels (x,y) in the first image.

Send each pixel (x,y) in the first image to its corresponding location $(x',y') = T(x,y)$ in the second image

Forward warping




Traverse the pixels (x, y) in the first image.

Send each pixel (x, y) in the first image to its corresponding location $(x', y') = T(x, y)$ in the second image

Q: what if pixel lands “between” two pixels?

A: distribute color among neighboring pixels (x', y')
 – Known as “splatting”

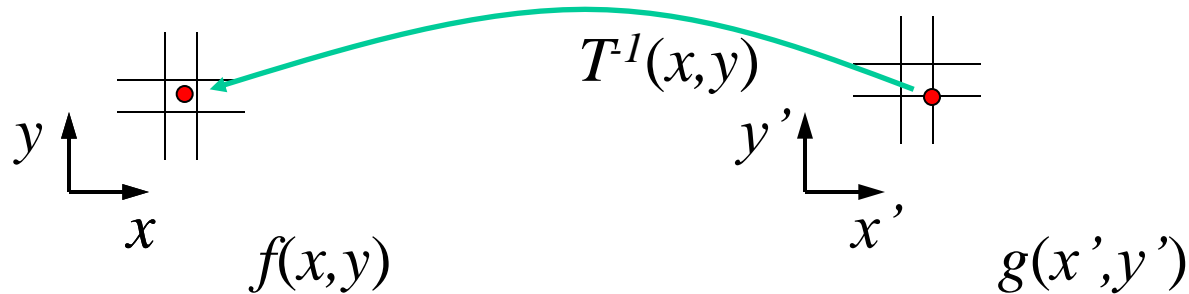
Inverse warping



Traverse the pixels (x', y') in the second image.

Get each pixel (x', y') in the second image from its corresponding location $(x, y) = T^{-1}(x', y')$ in the first image

Inverse warping



Traverse the pixels (x', y') in the second image.

Get each pixel (x', y') in the second image from its corresponding location $(x, y) = T^{-1}(x', y')$ in the first image

Q: what if pixel comes from “between” two pixels?

A: *Interpolate* color value from neighbors

- nearest neighbor, bilinear, Gaussian, bicubic

Linear interpolation in vector spaces


Any point between P and Q can
be obtained as
a linear combination

$$\lambda P + (1-\lambda) Q$$

NOTE: linear combination

$$\sum_i \lambda_i V_i \quad \text{for } V_i \in \mathcal{R}^N$$

is called **convex combination**
if $\sum_i \lambda_i = 1, \lambda_i \geq 0$.




e.g. $P + 0.5v = P + 0.5(Q - P)$
 $= 0.5P + 0.5Q$

Linear interpolation for functions

Assume 1D image (scan line)
with intensity $f(P)$ and $f(Q)$
for 2 pixels P and Q

Linear interpolation of function
 f between P and Q :


$$f(\lambda P + (1-\lambda) Q) = \lambda f(P) + (1-\lambda) f(Q)$$




In fact, any linear function on $[P, Q]$
must satisfy the equation above
(by definition of *linear functions*)

Bilinear interpolation (2 variate image intensity function)

Sampling of f at (x,y) :



pixels viewed as points in 2D



pixels viewed as square regions in 2D

$$\begin{aligned}
 f(x, y) = & (1-a)(1-b) f[i, j] \\
 & + a(1-b) f[i+1, j] \\
 & + ab f[i+1, j+1] \\
 & + (1-a)b f[i, j+1]
 \end{aligned}$$

Interpolated intensity at (x,y) can be seen as a weighted average of 4 near-by pixels intensities where weights based on overlap area


Forward vs. inverse warping

Q: which is better?

A: if area distortion is large, both have problems (blur, holes, aliasing)

when image areas (locally) **stretch** or **shrink** a lot, which happens when the determinant of the Jacobian for transformation $T: R^2 \rightarrow R^2$ significantly differs from 1. For linear warps $x'=Mx$ this corresponds to the determinant of M .

Example for 2D transformation $T: R^2 \rightarrow R^2$



example: blur due to significant area “stretch”


Forward vs. inverse warping

Q: which is better?


A: if area distortion is large, both have problems (blur, holes, aliasing)

when image areas (locally) **stretch** or **shrink** a lot, which happens when the determinant of the Jacobian for transformation $T: R^2 \rightarrow R^2$ significantly differs from 1. For linear warps $x'=Mx$ this corresponds to the determinant of M .

Example for 2D transformation $T: R^2 \rightarrow R^2$



Zwicker & Pfister, Trans. on Visualization and Comp. Graphics, 2002



example: aliasing due to significant area “shrink”

Forward vs. inverse warping



optional
slide

Q: which is better?

A: if area distortion is large, both have problems (blur, holes, aliasing)

when image areas (locally) **stretch** or **shrink** a lot, which happens when the determinant of the Jacobian for transformation $T: R^2 \rightarrow R^2$ significantly differs from 1. For linear warps $x'=Mx$ this corresponds to the determinant of M .

Example for 1D transformation $T: R^1 \rightarrow R^1$



inverse warping in python

Bug Warning: students often specify the transform from the input image to the output image instead of its inverse

`skimage.transform.warp (input_image, inverse_map,...)`



Second argument must be a function transforming coordinates in the output image into their corresponding coordinates in the input image.