

# **Image Filtering, Warping, and Morphing**

Michael Kazhdan

(601.457/657)



# Outline

- Image Filtering
- Image Warping
- Image Morphing



# Image Filtering

- What about the case when the modification that we would like to make to a pixel depends on the pixels around it?
  - Blurring
  - Edge Detection
  - Etc.



# Multi-Pixel Operations

## Stationary/Local Filtering

A general approach is to:


1. Define a *mask* of weights telling us how values at adjacent pixels should be combined to generate the new value.
2. Apply the (same) mask at every pixel.\*

\*Care is needed around at boundary pixels.



# Blurring

- To blur across pixels, define a mask:
  - Whose entries sum to one
  - Whose value is larger near the center of the mask
  - Whose values are non-negative



Original




Blur

$$\text{Filter} = \begin{bmatrix} 1/16 & 2/16 & 1/16 \\ 2/16 & 4/16 & 2/16 \\ 1/16 & 2/16 & 1/16 \end{bmatrix}$$

# Blurring



Pixel(x,y): red = 36  
green = 36  
blue = 0




Original

$$\text{Filter} = \begin{bmatrix} 1/16 & 2/16 & 1/16 \\ 2/16 & 4/16 & 2/16 \\ 1/16 & 2/16 & 1/16 \end{bmatrix}$$

# Blurring



**Pixel(x,y): red = 36**  
green = 36  
blue = 0



Original

	x - 1	x	x + 1
y - 1	36	109	146
y	32	36	109
y + 1	32	36	73

**Pixel(x,y).red and its  
red neighbors**


$$\text{Filter} = \begin{bmatrix} \frac{1}{16} & \frac{2}{16} & \frac{1}{16} \\ \frac{2}{16} & \frac{4}{16} & \frac{2}{16} \\ \frac{1}{16} & \frac{2}{16} & \frac{1}{16} \end{bmatrix}$$



# Blurring

**New value for Pixel(x,y).red =**

$$\begin{aligned} & (36 * 1/16) + (109 * 2/16) + (146 * 1/16) \\ & (32 * 2/16) + (36 * 4/16) + (109 * 2/16) \\ & (32 * 1/16) + (36 * 2/16) + (73 * 1/16) \end{aligned}$$



Original

	x - 1	x	x + 1
y - 1	36	109	146
y	32	36	109
y + 1	32	36	73

Filter = 
$$\begin{bmatrix} 1/16 & 2/16 & 1/16 \\ 2/16 & 4/16 & 2/16 \\ 1/16 & 2/16 & 1/16 \end{bmatrix}$$


**Pixel(x,y).red and its  
red neighbors**



# Blurring



New value for Pixel(x,y).red = 62.69



Original

	x - 1	x	x + 1
y - 1	36	109	146
y	32	36	109
y + 1	32	36	73


Pixel(x,y).red and its  
red neighbors

$$\text{Filter} = \begin{bmatrix} \frac{1}{16} & \frac{2}{16} & \frac{1}{16} \\ \frac{2}{16} & \frac{4}{16} & \frac{2}{16} \\ \frac{1}{16} & \frac{2}{16} & \frac{1}{16} \end{bmatrix}$$

# Blurring



New value for Pixel(x,y).red = 63



Original



Blur

$$\begin{bmatrix} 1/16 & 2/16 & 1/16 \\ 2/16 & 4/16 & 2/16 \\ 1/16 & 2/16 & 1/16 \end{bmatrix}$$

Note: Though we quantized, we could have used another techniques (e.g. dithering)



# Blurring


- Repeat for each color channel of each pixel.
- Keep source and destination separate to avoid drift.
- For boundary pixels, not all neighbors are used:
  - » Normalize the mask so the values sum to one, or
  - » Assume that the exterior values are black, or
  - » Assume the exterior values can be obtained by reflecting the image across the boundary, or
  - » Assume...



# Blurring

- Masks can have arbitrary size:
  - Can expand smaller masks by zero-padding

$$\begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 2 & 1 & 0 \\ 0 & 2 & 4 & 2 & 0 \\ 0 & 1 & 2 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} / 16 \iff \begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix} / 16$$



Original



Narrow Blur




# Blurring

- Masks can have arbitrary size:
  - Can expand smaller masks by zero-padding
  - Can use more non-zero entries to get a wider blur

$$\begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 2 & 1 & 0 \\ 0 & 2 & 4 & 2 & 0 \\ 0 & 1 & 2 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} / 16$$

$$\begin{bmatrix} 0 & 1 & 2 & 1 & 0 \\ 1 & 2 & 4 & 2 & 1 \\ 2 & 4 & 8 & 4 & 2 \\ 1 & 2 & 4 & 2 & 1 \\ 0 & 1 & 2 & 1 & 0 \end{bmatrix} / 48$$



Original



Narrow Blur



Wide Blur



# Blurring

A general way for defining the entries of a mask of size  $(2r + 1) \times (2r + 1)$  is to evaluate a Gaussian:

$$\text{GaussianMask}[i][j] \sim e^{-\frac{(i-r)^2 + (j-r)^2}{4r^2}}$$
$$i, j \in [0, 2r]$$

- $r$  is the (integer) mask radius
- $0 \leq i \leq 2r$  is the horizontal position in the mask
- $0 \leq j \leq 2r$  is the vertical position in the mask
- **Don't forget to normalize!**

Note:

The center of the mask is at index  $(r, r)$ .




# Edge Detection

An edge is where the image is far from constant:

⇒ The difference between the value at the pixel and the average value of neighboring pixels is large (in absolute value)

Define a mask whose:

- Entries sum to zero
- Value is one at the center pixel



$$\text{Filter} = \frac{1}{8} \begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}$$





# Edge Detection


This is sometimes referred to as a **Laplacian** filter.

Pixels with large absolute values correspond to edges:

- Positive values: “upper” edges
- Negative values: “lower” edges

Define a mask whose:

- Entries sum to zero
- Value is one at the center pixel




$$\text{Filter} = \frac{1}{8} \begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}$$



# Edge Detection



Pixel(x,y): red = 36  
green = 36  
blue = 0



Original

$$\text{Filter} = \frac{1}{8} \begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}$$

# Edge Detection



**Pixel(x,y): red = 36**  
green = 36  
blue = 0



Original

	x - 1	x	x + 1
y - 1	36	109	146
y	32	36	109
y + 1	32	36	73

**Pixel(x,y).red and its  
red neighbors**

$$\text{Filter} = \frac{1}{8} \begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}$$




# Edge Detection

**New value for Pixel(x,y).red =**

$$(36 * -1/8) + (109 * -1/8) + (146 * -1/8)$$

$$(32 * -1/8) + (36 * 1) + (109 * -1/8)$$

$$(32 * -1/8) + (36 * -1/8) + (73 * -1/8)$$



Original

	x - 1	x	x + 1
y - 1	36	109	146
y	32	36	109
y + 1	32	36	73


Pixel(x,y).red and its  
red neighbors

$$\text{Filter} = \frac{1}{8} \begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}$$



# Edge Detection

New value for Pixel(x,y).red = -285/8



Original

	x - 1	x	x + 1
y - 1	36	109	146
y	32	36	109
y + 1	32	36	73


Pixel(x,y).red and its  
red neighbors

$$\text{Filter} = \frac{1}{8} \begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}$$



# Edge Detection

New value for Pixel(x,y).red = -35.625



Original



Detected Edges

$$\text{Filter} = \frac{1}{8} \begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}$$

Note:

Output values **are not colors**.

⇒ Need to find a way to remap for visualization.




# Outline

- Image Filtering
- Image Warping
- Image Morphing



# Image Warping

- Move pixels of image
  - Mapping
  - Resampling



Source image

→  
Warp



Destination image




# Overview

- Mapping
  - Forward
  - Inverse
- Resampling
  - Point sampling
  - Triangle filter
  - Gaussian filter



# Mapping


- Define transformation
  - Describe the destination  $(x, y) = \Phi(u, v)$  for every location  $(u, v)$  in the source





# Example Mappings

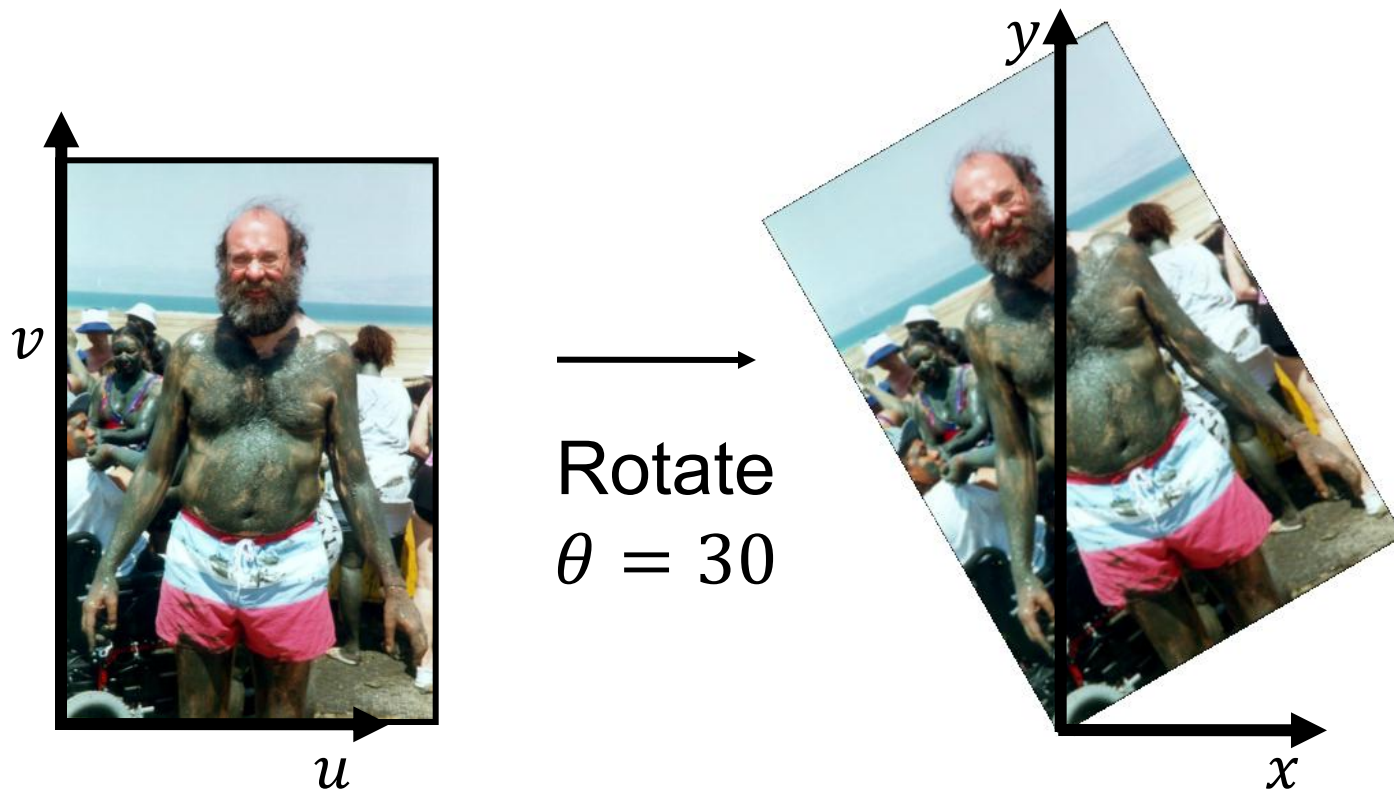
- Scale by  $\sigma$ :
  - $\Phi(u, v) = (\sigma u, \sigma v)$





# Example Mappings


- Rotate by  $\theta$  degrees:
  - $\Phi(u, v) = (u \cos \theta - v \sin \theta, u \sin \theta + v \cos \theta)$






# Example Mappings

- Shear in  $x$  by  $\sigma_x$ :
  - $\Phi(u, v) = (u + \sigma_x \cdot v, v)$




- Shear in  $y$  by  $\sigma_y$ :
  - $\Phi(u, v) = (u, v + \sigma_y \cdot u)$





# Other Mappings


- Any function of  $u$  and  $v$ :
  - $\Phi(u, v) = \dots$



Fish-eye



“Swirl”




“Rain”



# Image Warping Implementation I

- Forward mapping:


```
for( j=0 ; j<srcHeight ; j++ )  
  for( i=0 ; i<srcWidth ; i++ )  
    (x,y) =  $\Phi$ (i,j);  
    dst(x,y) = src(i,j);
```





# Forward Mapping

- Iterate over source image








# Forward Mapping – BAD!

- Iterate over source image

Multiple source pixels  
can map to the same  
destination pixel








# Forward Mapping – BAD!

- Iterate over source image






# Image Warping Implementation II

- Inverse mapping:


```
for( j=0 ; j<dstHeight ; j++ )  
  for( i=0 ; i<dstWidth ; i++ )  
    (u,v) =  $\Phi^{-1}(i,j)$  ;  
    dst(i,j) = src(u,v) ;
```





# Image Warping Implementation II

- Inverse mapping:
  - ✓ A single value assigned to each destination pixel
  - ✗ Must resample source






# Resampling

- Evaluate source image at  $(u, v) = \Phi^{-1}(i, j)$

$(u, v)$  does not usually have integer coordinates





# Overview


- Mapping
  - Forward
  - Inverse
- Resampling
  - Nearest Point Sampling
  - Bilinear Sampling
  - Gaussian Sampling



# Nearest Point Sampling

- Take value at closest pixel (indexed by center):

```
int intU = floor(u+0.5);  
int intV = floor(v+0.5);  
dst(i,j) = src(intU,intV);
```








# Bilinear Sampling

- Bilinearly interpolate four closest source pixels

$\text{dst}(i, j) = \text{Weighted average of source at}$   
 $(u_1, v_1), (u_2, v_1), (u_1, v_2), \text{ and } (u_2, v_2)$



Rotate  
+  
Translate






# Linear Sampling

- Linearly interpolate two closest source pixels

$\text{dst}(i)$  = linear interpolation of  $u_1$  and  $u_2$



Given  $i$ , the pixel location in the target:

$$u = \Phi^{-1}(i)$$

$$u_1 = \text{floor}(u);$$

$$u_2 = u_1 + 1;$$

$$du = u - u_1;$$

$$\text{dst}(i) = \text{src}(u_1) * (1-du) + \text{src}(u_2) * du;$$





# Bilinear Sampling

- Bilinearly interpolate four closest source pixels

$a$  = linear interpolation of  $\text{src}(u_1, v_1)$  and  $\text{src}(u_2, v_1)$

$b$  = linear interpolation of  $\text{src}(u_1, v_2)$  and  $\text{src}(u_2, v_2)$

$\text{dst}(i, j)$  = linear interpolation of  $a$  and  $b$

$$(u, v) = \Phi^{-1}(i, j)$$

$u_1 = \text{floor}(u)$  ,  $u_2 = u_1 + 1$ ;

$v_1 = \text{floor}(v)$  ,  $v_2 = v_1 + 1$ ;


$du = u - u_1$ ;

$a = \text{src}(u_1, v_1) * (1 - du)$   
 $+ \text{src}(u_2, v_1) * du$ ;

$b = \text{src}(u_1, v_2) * (1 - du)$   
 $+ \text{src}(u_2, v_2) * du$ ;

$dv = v - v_1$ ;

$\text{dst}(i, j) = a * (1 - dv) + b * dv$ ;





# Bilinear Sampling

- Bilinearly interpolate four closest source pixels

$a$  = linear interpolation of  $\text{src}(u_1, v_1)$  and  $\text{src}(u_2, v_1)$

$b$  = linear interpolation of  $\text{src}(u_1, v_2)$  and  $\text{src}(u_2, v_2)$

$\text{dst}(i, j)$  = linear interpolation of  $a$  and  $b$

$$(u, v) = \Phi^{-1}(i, j)$$

$$(u_1, v_2) \quad (u, v_2) \quad (u_2, v_2)$$

$$u_1 =$$

$$v_1 =$$

$$du =$$

$$a =$$

$$+ \text{src}(u_2, v_1) * (du);$$

$$b = \text{src}(u_1, v_2) * (1 - du)$$

$$+ \text{src}(u_2, v_2) * du;$$

$$dv = v - v_1;$$

$$\text{dst}(i, j) = a * (1 - dv) + b * dv;$$

Make sure to test that the pixels  
 $(u_1, v_1)$ ,  $(u_2, v_2)$ ,  $(u_1, v_2)$ , and  $(u_2, v_1)$   
are within the image.




# Gaussian Sampling

- Compute weighted sum of pixel neighborhood:
  - The blending weights are the normalized values of a Gaussian function.

## Note:

In contrast to the blurring filter, this doesn't assume that the center is at an integer location.

⇒ Can't precompute a stencil in advance







# Filtering Methods Comparison

- Trade-offs
  - Jagged edges versus blurring
  - Computation speed

We'll talk more about trade-offs next time.



Nearest



Bilinear




Gaussian



# Image Warping Implementation

- Inverse mapping:

```
for( j=0 ; j<dstHeight ; j++ )  
  for( i=0 ; i<dstWidth ; i++ )  
    (u,v) =  $\Phi^{-1}(i,j)$  ;  
    dst(i,j) = resample_src(u,v,r) ;
```






# Image Warping Implementation

- Inverse mapping:

```
for( j=0 ; j<dstHeight ; j++ )  
  for( i=0 ; i<dstWidth ; i++ )  
    (u,v) =  $\Phi^{-1}(i,j)$  ;  
    dst(i,j) = resample_src(u,v,r) ;
```





# Example: Scale

Scale( src , dst ,  $\sigma$  ):

$r \cong ?$ ;


```
for( j=0 ; j<dstHeight ; j++ )
```

```
  for( i=0 ; i<dstWidth ; i++ )
```

```
    (u,v) = (i,j) /  $\sigma$ ;
```

```
    dst(i,j) = resample_src(u,v,r) ;
```

$$r = \frac{1}{\sigma}$$





# Example: Rotate

Rotate( src , dst ,  $\theta$  ):

$r \cong ?$ ;

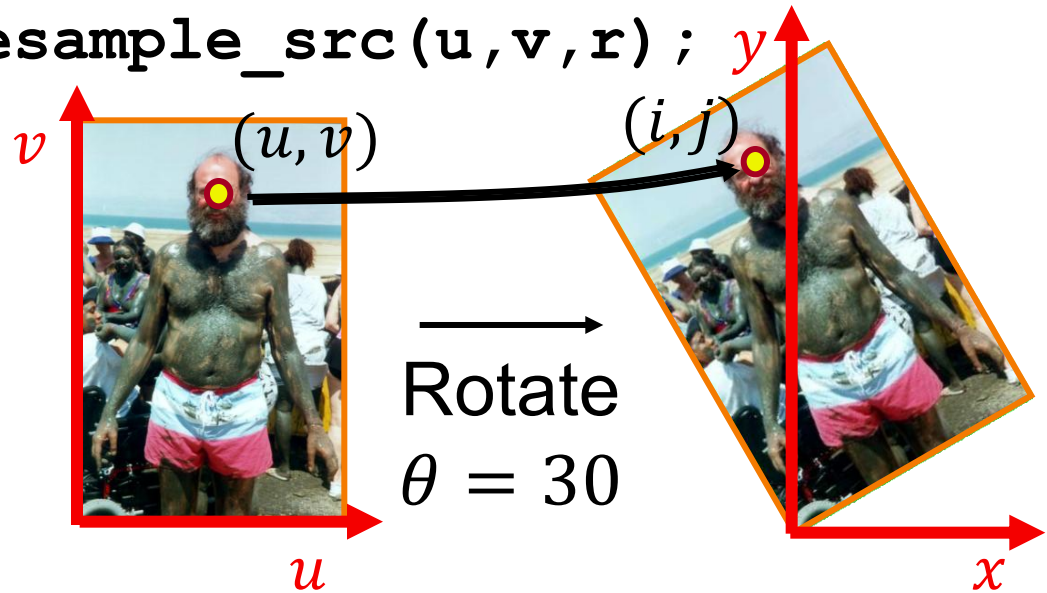
for(  $j=0$  ;  $j<\text{dstHeight}$  ;  $j++$  )

for(  $i=0$  ;  $i<\text{dstWidth}$  ;  $i++$  )

$(u,v) = ( i*\cos(-\theta) - j*\sin(-\theta) ,$   
 $i*\sin(-\theta) + j*\cos(-\theta) )$  ;

$\text{dst}(x,y) = \text{resample\_src}(u,v,r)$  ;

$$r = 1$$








# Example:

General( src , dst ,  $\Phi$  ):

```
r  $\cong$  ?;  
for( j=0 ; j<dstHeight ; j++ )  
  for( i=0 ; i<dstWidth ; i++ )  
    (u,v) =  $\Phi^{-1}(i,j)$  ;  
    dst(i,j) = resample_src(u,v,r) ;
```

$r = ?$





# Example:

General( src , dst ,  $\Phi$  ):

```
r  $\cong$  ? ;  
for( j=0 ; j<dstHeight ; j++ )  
    for( i=0 ; i<dstWidth ; i++ )  
        (u,v) =  $\Phi^{-1}(i,j)$  ;  
        dst(i,j) = resample_src(u,v,r) ;
```

Instead of using a fixed radius circle to sample the source, we can:

1. Have the radius change,
2. Use an ellipse.

For example, the parameters can be determined by looking at the derivative/Jacobian of  $\Phi$ .

# Outline

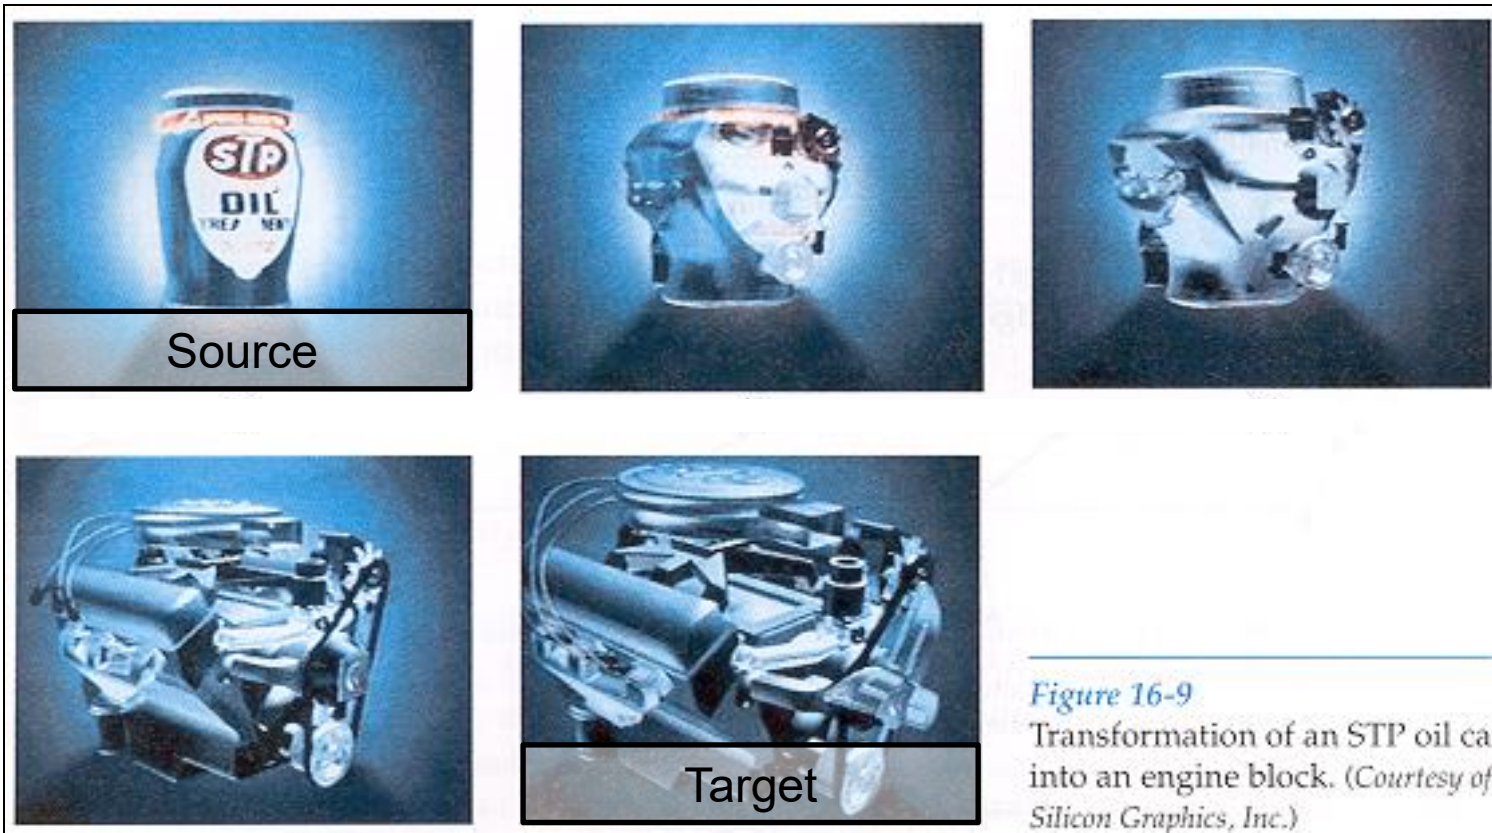
- Image Filtering
- Image Warping
- Image Morphing





# Image Morphing

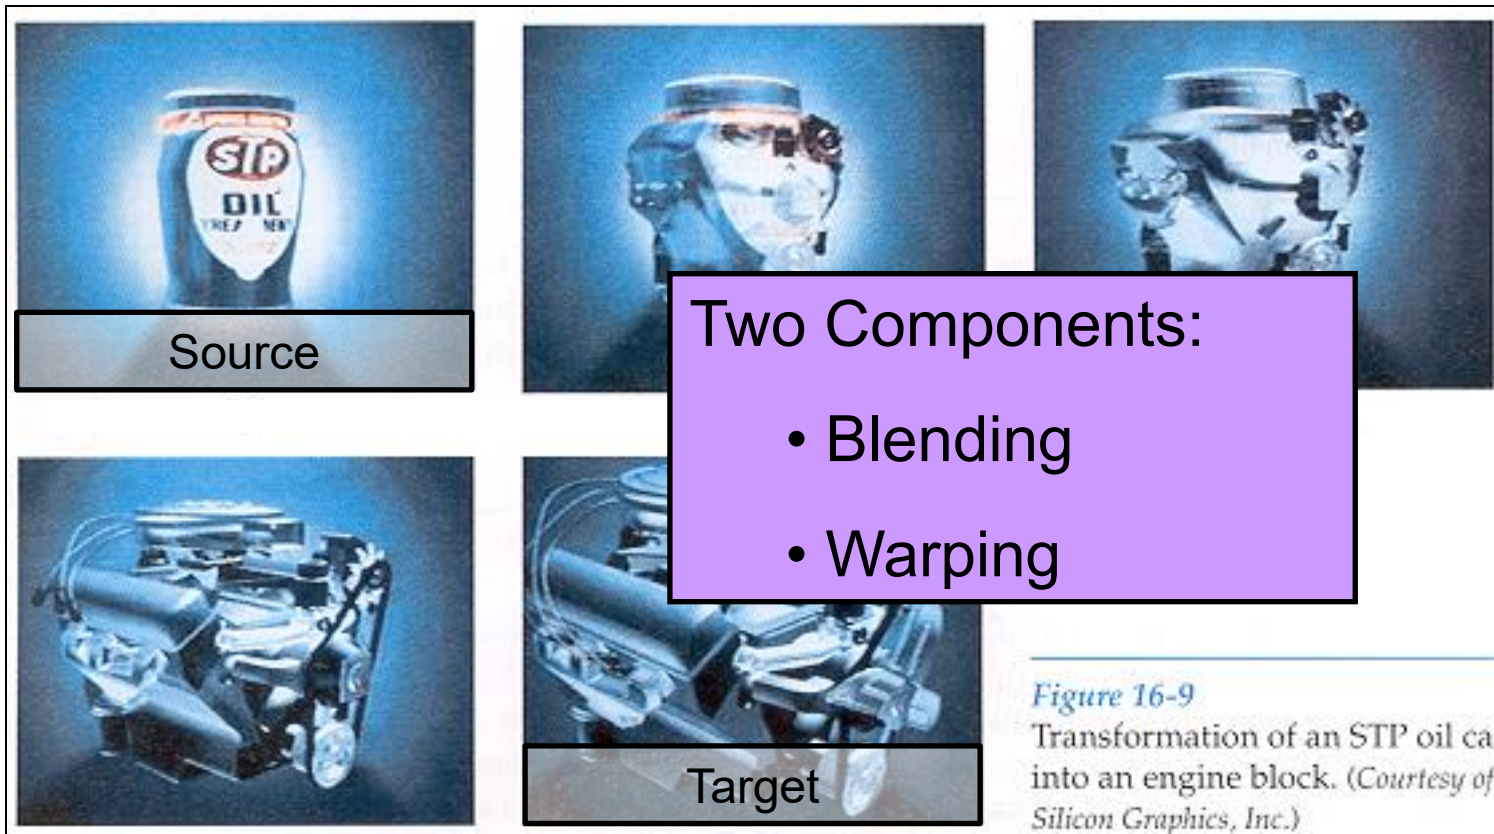
- Animate transition between two images





# Image Morphing

- Animate transition between two images





# Image Morphing

Recall (Inner Product):


1. For two vectors  $\vec{v}_1 = (x_1, y_1)$  and  $\vec{v}_2 = (x_2, y_2)$ , the inner product between  $\vec{v}_1$  and  $\vec{v}_2$  is:

$$\langle \vec{v}_1, \vec{v}_2 \rangle \equiv x_1 \cdot x_2 + y_1 \cdot y_2 = \|\vec{v}_1\| \cdot \|\vec{v}_2\| \cdot \cos \theta$$

⇒ Gives the cosine of the angle between the two, scaled by the product of their lengths.

⇒ Dividing by the length of  $v_2$  gives the signed length of the projection of  $v_1$  onto the line through  $v_2$ :

$$\|\vec{v}_1\| \cdot \cos \theta = \frac{\langle \vec{v}_1, \vec{v}_2 \rangle}{\|\vec{v}_2\|}$$







# Image Morphing

Recall (Inner Product):

1. For two vectors  $\vec{v}_1 = (x_1, y_1)$  and  $\vec{v}_2 = (x_2, y_2)$ , the inner product between  $\vec{v}_1$  and  $\vec{v}_2$  is:

$$\langle \vec{v}_1, \vec{v}_2 \rangle \equiv x_1 \cdot x_2 + y_1 \cdot y_2 = \|\vec{v}_1\| \cdot \|\vec{v}_2\| \cdot \cos \theta$$

2. For a vector  $\vec{v} = (x, y)$ , its  $90^\circ$  (CCW) rotation is:  
$$\vec{v}^\perp = (-y, x)$$

This is the vector:

- With the same length as  $\vec{v}$ , and
- Which is perpendicular to  $\vec{v}$ :

$$\langle \vec{v}, \vec{v}^\perp \rangle = 0$$

In 2D, the perpendicular is unique up to sign (disambiguated with CCW vs CW).

In higher dimensions, there is a continuum of perpendicular vectors.



# Image Morphing

Recall (Inner Product):

1. For two vectors  $\vec{v}_1 = (x_1, y_1)$  and  $\vec{v}_2 = (x_2, y_2)$ , the inner product between  $\vec{v}_1$  and  $\vec{v}_2$  is:

$$\langle \vec{v}_1, \vec{v}_2 \rangle \equiv x_1 \cdot x_2 + y_1 \cdot y_2 = \|\vec{v}_1\| \cdot \|\vec{v}_2\| \cdot \cos \theta$$

2. For a vector  $\vec{v} = (x, y)$ , its  $90^\circ$  (CCW) rotation is:  
$$\vec{v}^\perp = (-y, x)$$

3. The inner-product of two vectors is *isometry-invariant*:  
If  $\mathbf{R} \in \mathbb{R}^{2 \times 2}$  is a rotation/reflection, then:

$$\langle \mathbf{R}(\vec{v}_1), \mathbf{R}(\vec{v}_2) \rangle = \langle \vec{v}_1, \vec{v}_2 \rangle$$

$\Rightarrow$  In particular:

$$\langle \vec{v}_1^\perp, \vec{v}_2^\perp \rangle = \langle \vec{v}_1, \vec{v}_2 \rangle$$






# Image Morphing: Blending

Blend **colors** using an  $\alpha$ -blend ( $\alpha \in [0,1]$ ):




$$\text{blend}(i, j, \alpha) = (1 - \alpha) \cdot \text{Img}_0(i, j) + \alpha \cdot \text{Img}_1(i, j)$$





# Image Morphing: Warping

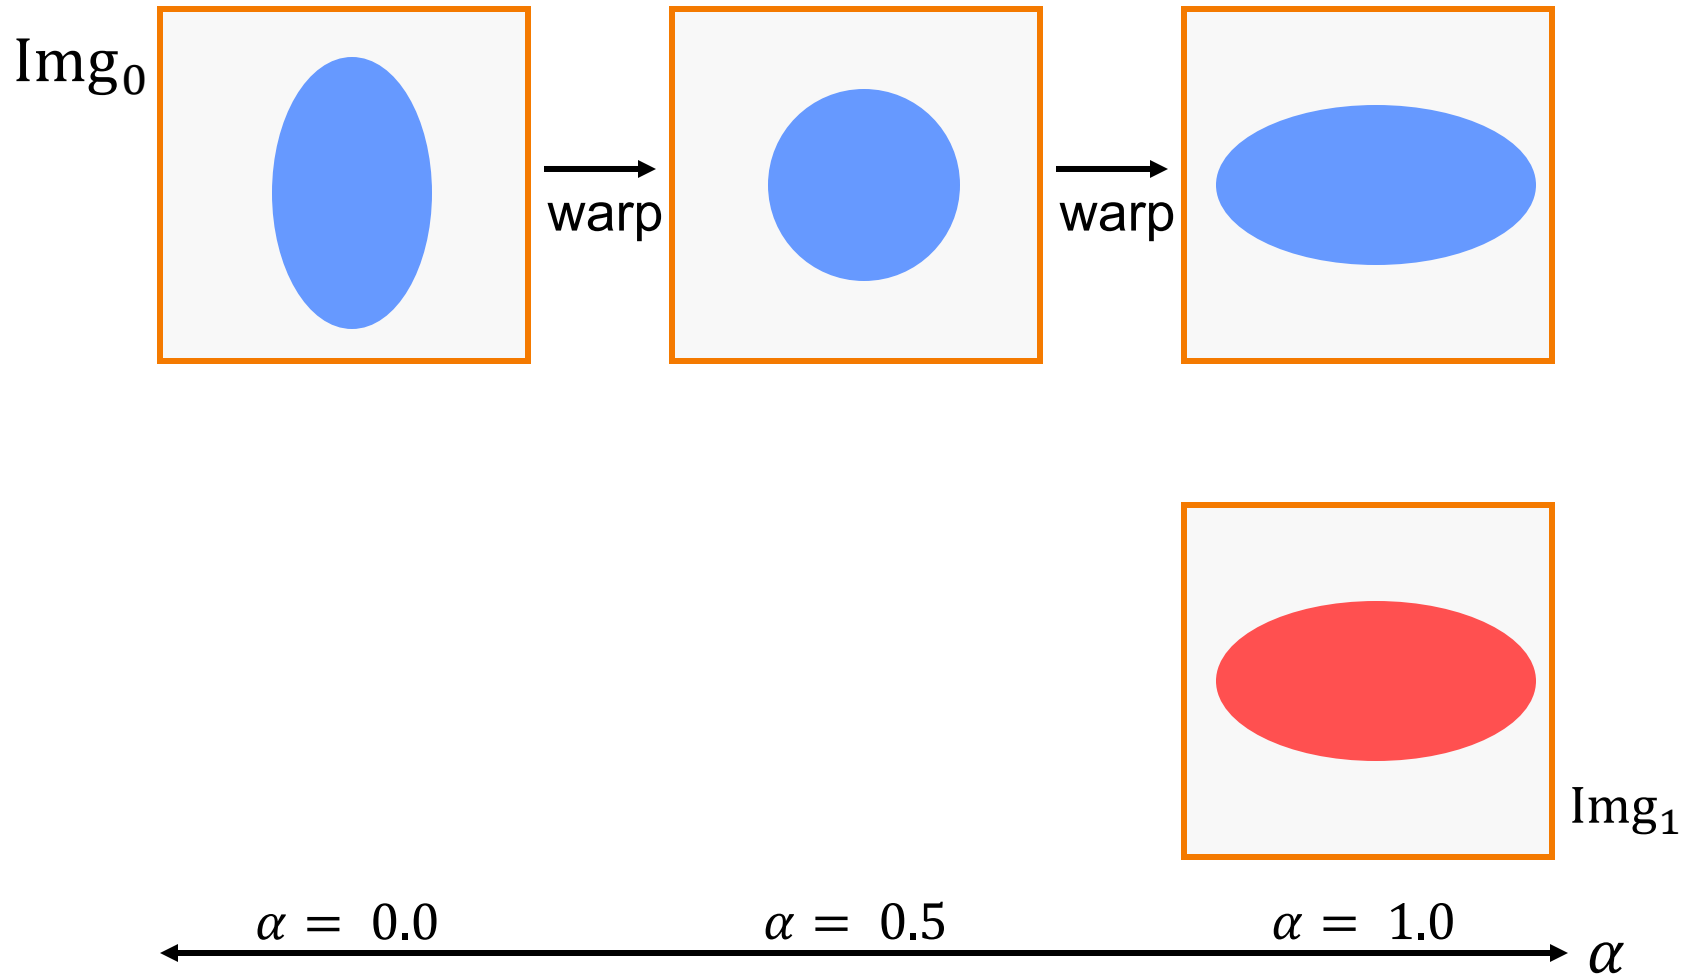
Deform  $\text{Img}_0$  so its **shape** matches that of  $\text{Img}_1$ ...





# Image Morphing: Warping

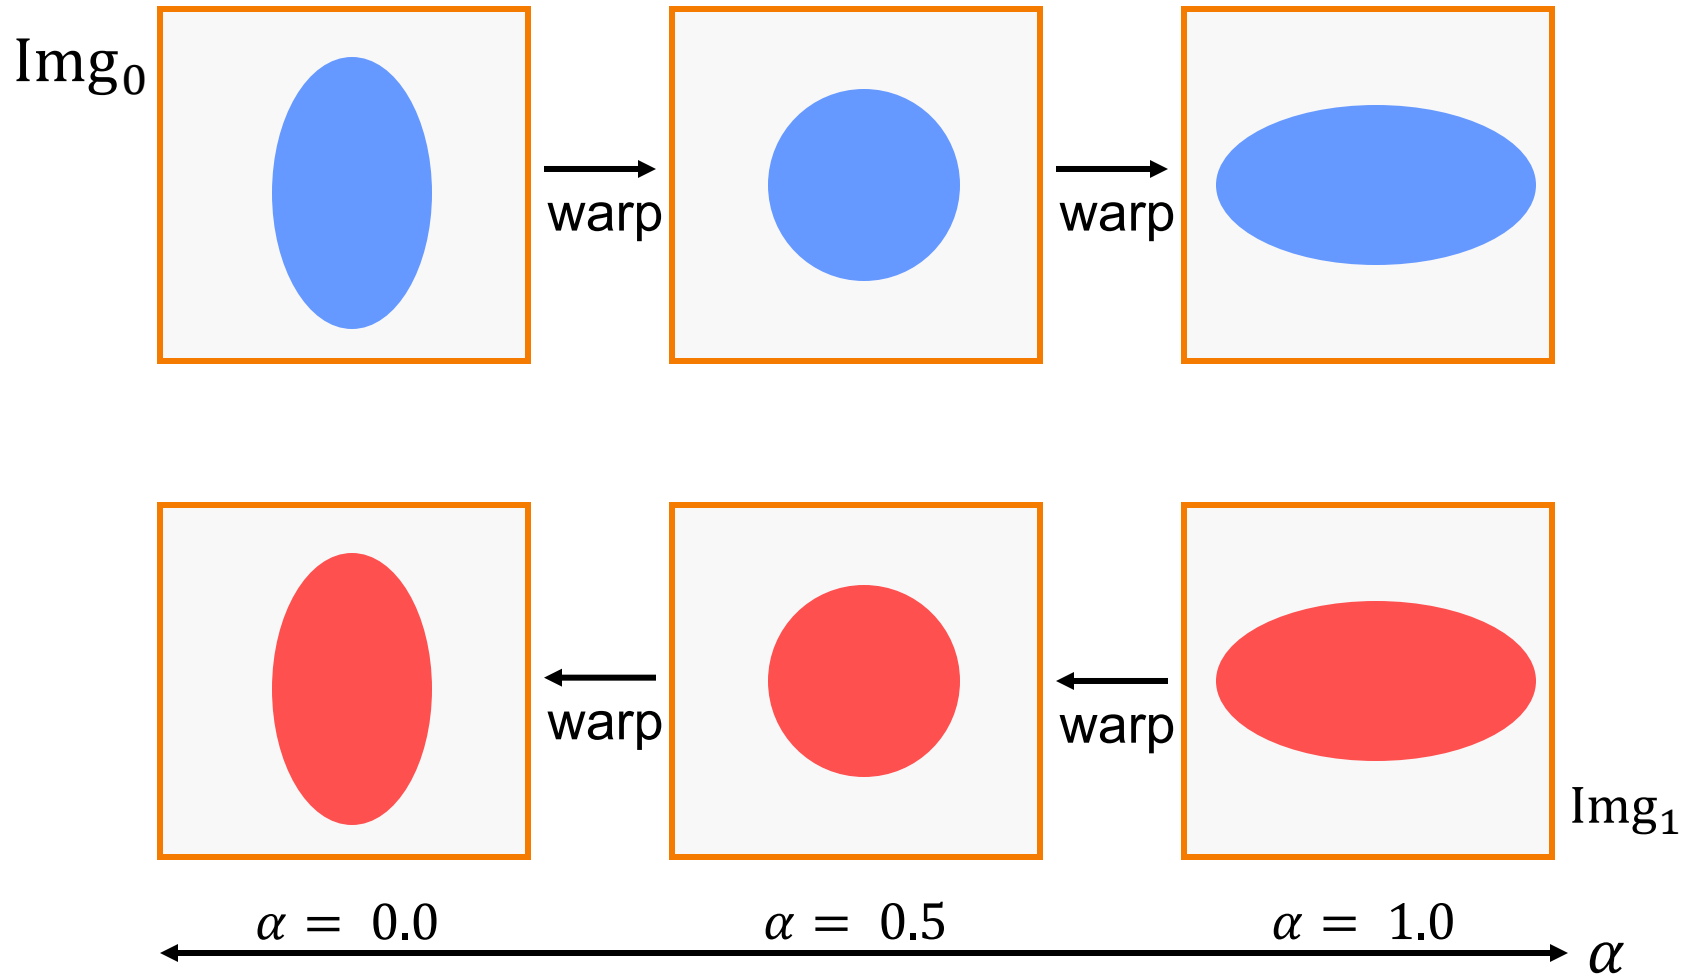
Deform  $\text{Img}_0$  so its **shape** matches that of  $\text{Img}_1$ ...





# Image Morphing: Warping


Deform  $\text{Img}_1$  so its **shape** matches that of  $\text{Img}_0$ ...



# Image Morphing: Warping + Blending



... then blend **colors**





# Image Morphing: Warping

- The warping step is the hard one
  - Aim to align features in images



How do we specify the mapping for the warp?

*Silicon Graphics, Inc.)*

# Feature-Based Warping [Beier & Neeley, 1992]




## Challenge:


- Given  $p$  in the destination, what is the corresponding source location?

## Approach:

- Define a pair of corresponding lines in the source and target
- Describe  $p$  relative to the destination line
- Map the description to the source



Source image



Destination image

$u$  is a signed fraction  
 $v$  is a signed length (in pixels)

# Feature-Based Warping [Beier & Neeley, 1992]



How do we calculate  $v$  (perp. pixel distance)?

Recall:


The signed length of  $\vec{v}_1$  projected onto  $\vec{v}_2$  is:

$$\frac{\langle \vec{v}_1, \vec{v}_2 \rangle}{\|\vec{v}_2\|}$$

In our case we want:

- The signed length of the direction from the start of the segment to  $p$ :

$$\vec{v}_1 = p - s$$





# Feature-Based Warping [Beier & Neeley, 1992]



How do we calculate  $v$  (perp. pixel distance)?

Recall:

The signed length of  $\vec{v}_1$  projected onto  $\vec{v}_2$  is:

$$\frac{\langle \vec{v}_1, \vec{v}_2 \rangle}{\|\vec{v}_2\|}$$


In our case we want:

- The signed length of the direction from the start of the segment to  $p$ :

$$\vec{v}_1 = p - s$$

- Projected onto the perpendicular of the direction from the start of the segment to the end:

$$\vec{v}_2 = (t - s)^\perp$$



# Feature-Based Warping [Beier & Neeley, 1992]



How do we calculate  $v$  (perp. pixel distance)?

Recall:

The signed length of  $\vec{v}_1$  projected onto  $\vec{v}_2$  is:

$$\frac{\langle \vec{v}_1, \vec{v}_2 \rangle}{\|\vec{v}_2\|}$$


In our case we want:

- The signed length of the direction from the start of the segment to  $p$ :

$$\vec{v}_1 = p - s$$

- Projected onto the perpendicular of the direction from the start of the segment to the end:

$$\vec{v}_2 = (t - s)^\perp$$




# Feature-Based Warping [Beier & Neeley, 1992]



How do we calculate  $v$  (perp. pixel distance)?

This gives:

$$v = \frac{\langle \overbrace{p - s}^{\vec{v}_1}, \overbrace{(t - s)^\perp}^{\vec{v}_2} \rangle}{\|(t - s)^\perp\|}$$



# Feature-Based Warping [Beier & Neeley, 1992]




How do we calculate  $u$  (par. **fractional** distance)?

Similarly:

- The signed projected length is:

$$\tilde{u} = \frac{\overbrace{\langle p - s, t - s \rangle}^{\vec{v}_1} \overbrace{\langle t - s, t - s \rangle}^{\vec{v}_2}}{\|t - s\|}$$



# Feature-Based Warping [Beier & Neeley, 1992]




How do we calculate  $u$  (par. **fractional** distance)?

Similarly :

- The **fractional** signed projected length is obtained by dividing by the length of the segment:


$$u = \frac{\overbrace{\langle p - s, t - s \rangle}^{\vec{v}_1}}{\underbrace{\|t - s\|}_{\vec{v}_2}} \cdot \frac{1}{\|t - s\|}$$





# Warping with One Line Pair

What happens to the “F”?




Translation!



# Warping with One Line Pair

What happens to the “F”?




Non-uniform scale!



# Warping with One Line Pair

What happens to the “F”?




Rotation!





# Warping with One Line Pair

What happens to the “F”?




*What types of affine transformations can't be specified?*



# Warping with One Line Pair


Can't specify arbitrary scales, skews, mirrors, angular changes...





# Warping with Multiple Line Pairs

Use weighted combination of points defined by each pair of corresponding lines

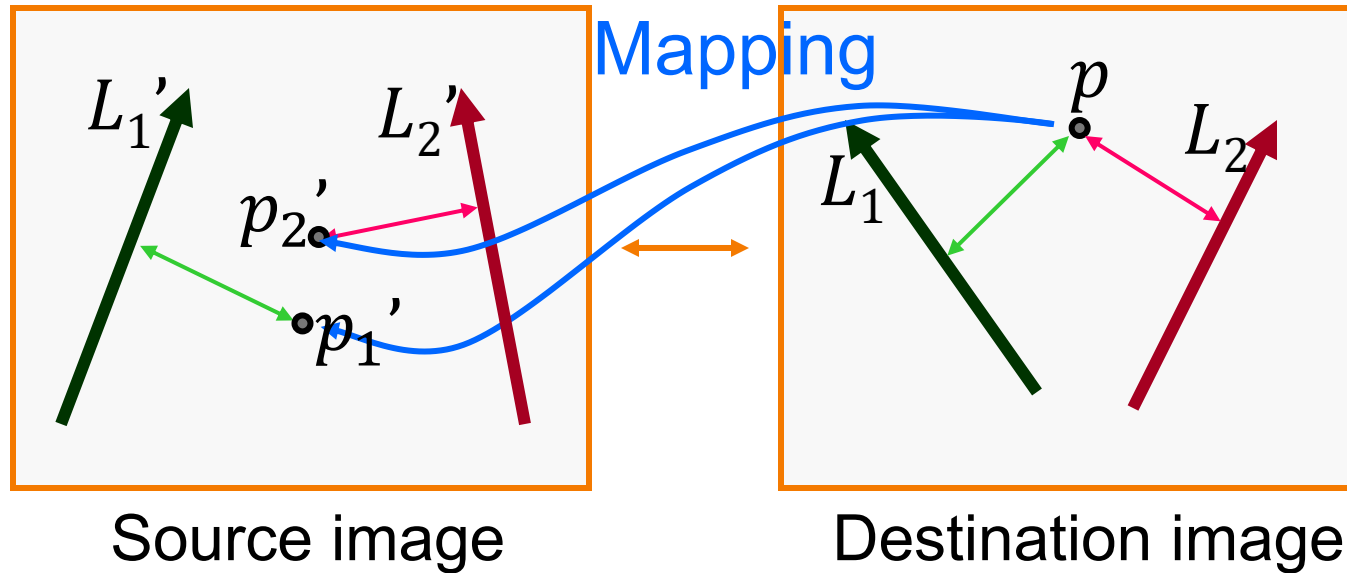


Beier & Neeley, Figure 4



# Warping with Multiple Line Pairs


Use weighted combination of points defined by each pair of corresponding lines





# Warping with Multiple Line Pairs

Use weighted combination of points defined by each pair of corresponding lines



$p'$  is a weighted average



# Weighting Effect of Each Line Pair

Given a set of line pairs  $\{L_{in}[0], \dots, L_{in}[N]\}$  and  $\{L_{out}[0], \dots, L_{out}[N]\}$ , to weight the contribution of each line pair, [Beier & Neely, 1992] use:

$$\text{weight}[i](p) \sim \left( \frac{\text{length}[i]^c}{a + \text{dist}[i](p)} \right)^b$$

where:

- $\text{length}[i]$  is the length of line  $L_{out}[i]$
- $\text{dist}[i](p)$  is the distance from  $p$  to  $L_{out}[i]$
- $a$  (small),  $b \in [0.5, 2.0]$ ,  $c \in [0.0, 1.0]$  are constants that control the warp


Note: The “~” indicates “up to a constant”. Need to normalize so weights sum to 1.



# Feature-Based Warping

How do we calculate the unsigned distance from a point  $p$  to the line segment from  $s$  to  $t$ ?

$$\text{dist}(p) = \begin{cases} |v| & \text{if } u \in [0,1] \\ \|p - s\| & \text{if } u < 0 \\ \|p - t\| & \text{if } u > 1 \end{cases}$$





# Warping

## Input:

Source image

Set of corresponding line segment pairs in source and target

## Goal:

Warp the source image to a target image

## Solution:

Iterating over each pixel in the target

- For each pair of line segments
  - » Compute the corresponding position in the source
  - » Compute the weights
- Average to get the final source position
- Sample the source (at the source position) to get the target color





# Warping: Pseudocode

```
Warp(  $Img_{src}$  ,  $L_{src}[N]$  ,  $L_{tgt}[N]$  )  
{  
    foreach target pixel  $p_{tgt}$ :  
         $p_{src} = (0,0)$   
         $sum = 0$   
        for  $i = 0$  to  $N$ :  
             $q_{src} = p_{tgt}$  transformed by (  $L_{src}[i]$  ,  $L_{tgt}[i]$  )  
             $p_{src} += q_{src} * weight[i]( p_{tgt} )$   
             $sum += weight[i]( p_{tgt} )$   
         $p_{src} /= sum$   
         $Img_{tgt}(p_{tgt}) = Img_{src}( p_{src} )$   
    return  $Img_{tgt}$   
}
```



# Morphing

## Input:

Source and target images

Set of corresponding line segment pairs in source and target

Interpolation time  $\alpha \in [0,1]$

## Goal:

The morph of the source to the target at time  $\alpha$

## Solution:

- Compute the  $\alpha$ -weighted average of the of line segments (by averaging the end-points)
- **Warp** the source using the source and averaged line segments
- **Warp** the target using the target and averaged line segments
- Compute the  $\alpha$ -**blend** of the warped images



# Morphing: Pseudocode

```
Morph(  $Img_0$  ,  $L_0[N]$  ,  $Img_1$  ,  $L_1[N]$  ,  $\alpha$  )  
{  
    foreach  $i \in \{1, \dots, N\}$ :  
         $L_\alpha[i]$  = line  $\alpha$ -th of the way from  $L_0[i]$  to  $L_1[i]$   
  
     $Warp_0$  = Warp(  $Img_0$  ,  $L_0[]$  ,  $L_\alpha[]$  )  
     $Warp_1$  = Warp(  $Img_1$  ,  $L_1[]$  ,  $L_\alpha[]$  )  
    }  
  
    return  $(1-\alpha) * Warp_0 + \alpha * Warp_1$   
}
```

# [Beier & Neely, 1992] Example ( $\alpha = 0.5$ )



Img<sub>0</sub>





Figure 7

Warp<sub>0</sub>

Img<sub>1</sub>




Warp<sub>1</sub>

# [Beier & Neely, 1992] Example ( $\alpha = 0.5$ )




Img<sub>0</sub>



Warp<sub>0</sub>

Img<sub>1</sub>



Warp<sub>1</sub>

# [Beier & Neely, 1992] Example ( $\alpha = 0.5$ )



Img<sub>0</sub>




Figure 7




Figure 10

Warp<sub>0</sub>

Result






Figure 8

Img<sub>1</sub>





Warp<sub>1</sub>

# [Beier & Neely, 1992] Example ( $\alpha = 0.5$ )





Img<sub>0</sub>



Warp<sub>0</sub>

Result



Warp<sub>1</sub>

Img<sub>1</sub>



# Full Morph Animation: Pseudocode



```
Animate(  $Img_0$  ,  $L_0[N]$  ,  $Img_1$  ,  $L_1[N]$  ,  $Imgs_{out}[T+1]$  )  
{  
    foreach  $t \in \{0, \dots, T\}$ :  
         $Imgs_{out}[t] = \text{Morph}(Img_0 , L_0[N] , Img_1 , L_1[N] , t/T )$   
}
```





# Morphing

Check out Michael Jackson's "Black or White" video at:

<https://www.youtube.com/watch?v=pTFE8cirkdQ>



Or the earlier Plymouth Voyager commercial at:

<https://www.youtube.com/watch?v=0b939O7dGqQ>