

Matching in 2D

- Find an instance of an object class or a specific object in an image



- Find correspondences between points or other features in two (or more) images





Symbolic vs Geometric Matching

- **Symbolic Matching:** the model is a symbolic structure, usually a graph. The image is also represented as a graph, and we use different types of graph matching.
- **Geometric Matching:** the model has specific geometric features and measurements. We try to find a transformation (ie. translation, rotation, scale, skew) from the model to a piece of the image.

The Consistent Labeling Formalism for Symbolic Matching

- A **part** (unit) is a structure in the scene, such as a region or segment or corner.
- A **label** is a symbol assigned to identify the part.
- An **N-ary relation** is a set of N-tuples defined over a set of parts or a set of labels.
- An **assignment** is a mapping from parts to labels.

Example



What are the relationships?

What is the best assignment of model labels to image features?

Consistent Labeling Definition

Given:

1. a set of units P
2. a set of labels for those units L
3. a relation RP over set P
4. a relation RL over set L


A consistent labeling f is a mapping $f: P \rightarrow L$ satisfying

if $(p_i, p_j) \in RP$, then $(f(p_i), f(p_j)) \in RL$

which means that a consistent labeling preserves relationships.

Abstract Example

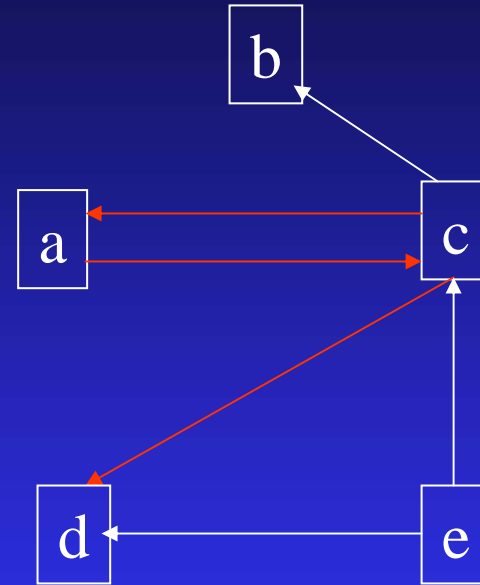
binary relation RP



$$P = \{1, 2, 3\}$$

$$RP = \{(1, 2), (2, 1), (2, 3)\}$$

binary relation RL



$$L = \{a, b, c, d, e\}$$

$$RL = \{(a, c), (c, a), (c, b), (c, d), (e, c), (e, d)\}$$

One consistent labeling is $\{(1, a), (2, c), (3, d)\}$

House Example




Image 1 **P**




Image 2 **L**

$$P = \{S1, S2, S3, S4, S5, S6, S7, S8, S9, S10, S11\}.$$

$$L = \{Sa, Sb, Sc, Sd, Se, Sf, Sg, Sh, Si, Sj, Sk, Sl, Sm\}.$$

$$R_P = \{ (S1, S2), (S1, S5), (S1, S6), (S2, S3), (S2, S4), (S3, S4), (S3, S9), (S4, S5), (S4, S7), (S4, S11), (S5, S6), (S5, S7), (S5, S11), (S6, S8), (S6, S11), (S7, S9), (S7, S10), (S7, S11), (S8, S10), (S8, S11), (S9, S10) \}.$$

$$R_L = \{ (Sa, Sb), (Sa, Sj), (Sa, Sn), (Sb, Sc), (Sb, Sd), (Sb, Sn), (Sc, Sd), (Sd, Se), (Sd, Sf), (Sd, Sg), (Se, Sf), (Se, Sg), (Sf, Sg), (Sf, Sl), (Sf, Sm), (Sg, Sh), (Sg, Si), (Sg, Sn), (Sh, Si), (Sh, Sk), (Sh, Sl), (Sh, Sn), (Si, Sj), (Si, Sk), (Si, Sn), (Sj, Sk), (Sk, Sl), (Sl, Sm) \}.$$

RP and RL are connection relations.

$$f(S1) = S_j$$

$$f(S4) = S_n$$

$$f(S7) = S_g$$

$$f(S10) = S_f$$

$$f(S2) = S_a$$

$$f(S5) = S_i$$

$$f(S8) = S_l$$

$$f(S11) = S_h$$

$$f(S3) = S_b$$

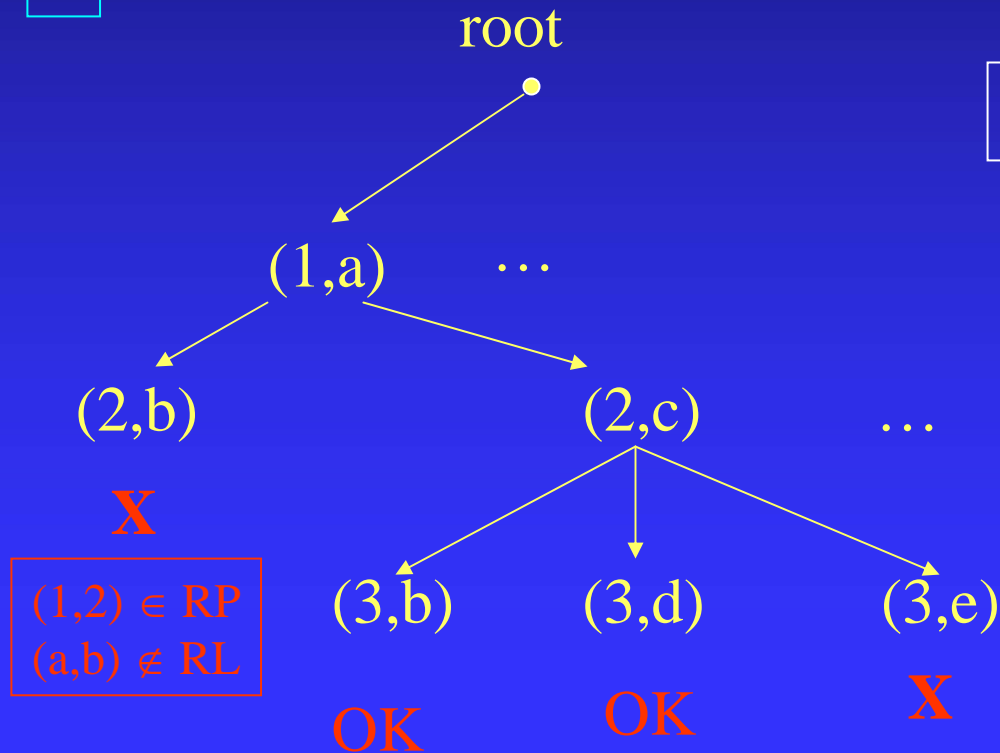
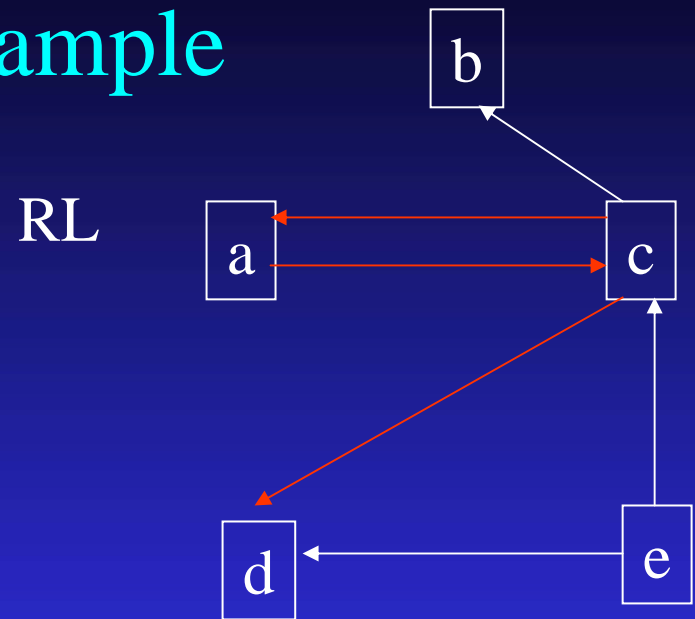

$$f(S6) = S_k$$

$$f(S9) = S_d$$

1. Interpretation Tree

- An **interpretation tree** is a tree that represents all assignments of labels to parts.
- Each path from the root node to a leaf represents a (partial) assignment of labels to parts.
- Every path terminates as either
 1. a complete consistent labeling
 2. a failed partial assignment

Interpretation Tree Example



Tree Search for Binary Relations

Procedure Treesearch(P,L,RP,RL,f)

/* parts P, labels L, part graph RP, label graph RL,
and mapping f */

take the first part p in P

for each label l in L

{

create mapping f' with everything in f plus (p,l)

check the consistency of the new pair (p,l)

with all pairs already in f

if there are no inconsistencies

if we have used all the parts just output f'

else Treesearch(P-{p},L-{l}.RP,RL,f')


}

Checking Consistency

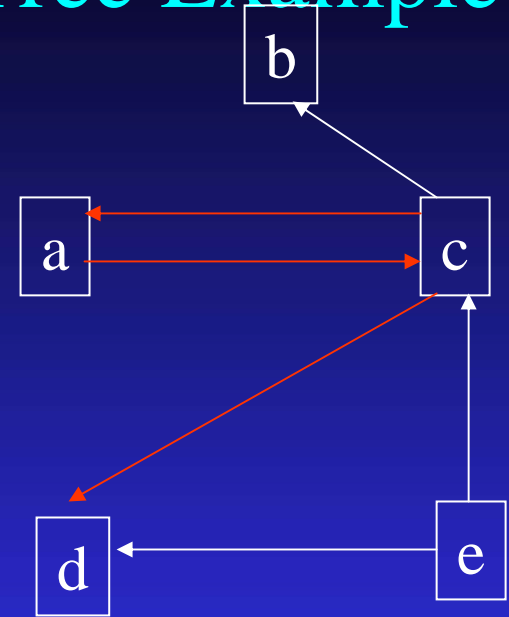
- Suppose $f = \{(p_1, l_1), (p_2, l_2), \dots, (p_k, l_k)\}$.
- Suppose the new pair is (p, l) .
- For each of the “old parts” p_i in $\{p_1, \dots, p_k\}$
 - if (p, p_i) is in RP then (l, l_i) must be in RL
 - if (p_i, p) is in RP then (l_i, l) must be in RL
- If even one test fails, the pair (p, l) is not consistent with the mapping f .

Back to the Interpretation Tree Example

RP




RL



$P = \{1, 2, 3\}$
 $L = \{a, b, c, d, e\}$
 $f = \{\}$

root



Tree Search Algorithm (Very General for N-tuple Relations)

```
procedure Interpretation_Tree_Search( $P, L, R_P, R_L, f$ );
{
 $p := \text{first}(P)$ ;
for each  $l$  in  $L$ 
{
 $f' = f \cup \{(p, l)\}$ ; /* add part-label to interpretation */
OK = true;
for each N-tuple  $(p_1, \dots, p_N)$  in  $R_P$  containing component  $p$ 
and whose other components are all in domain( $f$ )
/* check on relations */
if  $(f(p_1), \dots, f(p_N))$  is not in  $R_L$  then
{
OK = false;
break;
}
if OK then
{
 $P' = \text{rest}(P)$ ;
if isempty( $P'$ ) then output( $f'$ );
else Interpretation_Tree_Search( $P', L, R_P, R_L, f'$ );
}
}
}
```

2. Discrete Relaxation

- Discrete relaxation is an **alternative to** (or addition to) the interpretation tree search.
- Relaxation is an **iterative** technique with polynomial time complexity.
- Relaxation uses **local constraints** at each iteration.
- It can be implemented on parallel machines.

How Discrete Relaxation Works

1. Each unit is assigned a set of **initial possible labels**.
2. All **relations are checked** to see if some pairs of labels are impossible for certain pairs of units.
3. **Inconsistent labels are removed** from the label sets.
4. If any labels have been filtered out then another pass is executed else the relaxation part is done.
5. If there is more than one labeling left, a tree search can be used to find each of them.

Example of Discrete Relaxation

RP

P_i

L1 L2 L3

X

P_j

L6 L8

RL

L1

L2

L3



L8

L6

There is no label in P_j's label set that is connected to L2 in P_i's label set. L2 is inconsistent and filtered out.

3. Relational Distance Matching

- A fully consistent labeling is unrealistic.
- An image may have missing and extra features; required relationships may not always hold.
- Instead of looking for a consistent labeling, we can look for the **best mapping from P to L**, the one that preserves the most relationships.



Preliminary Definitions


Def: A **relational description** DP is a sequence of relations over a set of primitives P.

- Let $DA = \{R_1, \dots, R_I\}$ be a relational description over A.
- Let $DB = \{S_1, \dots, S_I\}$ be a relational description over B.
- Let f be a 1-1, onto mapping from A to B.
- For any relation R, the composition $R \circ f$ is given by

$$R \circ f = \{(b_1, \dots, b_n) \mid (a_1, \dots, a_n) \text{ is in } R \text{ and } f(a_i) = (b_i), i=1, n\}$$

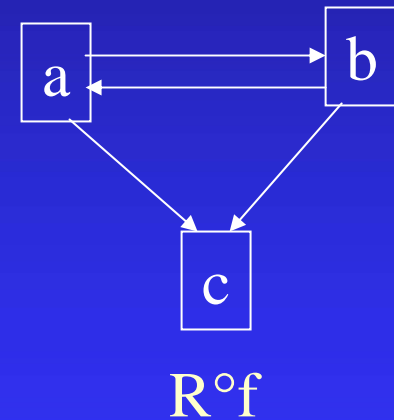
Example of Composition

$$R \circ f = \{(b_1, \dots, b_n) \mid (a_1, \dots, a_n) \text{ is in } R \text{ and } f(a_i) = (b_i), i=1, n\}$$



| | |
|---|---|
| 1 | a |
| 2 | b |
| 3 | c |

f



$R \circ f$ is an isomorphic copy of R with nodes renamed by f.

Relational Distance Definition

Let DA be a relational description over set A ,
 DB be a relational description over set B ,
and $f : A \rightarrow B$.

- The **structural error of f** for R_i in DA and S_i in DB is

$$E_S^i(f) = |R_i \circ f - S_i| + |S_i \circ f^{-1} - R_i|$$

- The **total error of f** with respect to DA and DB is



$$E(f) = \sum_i E_S^i(f)$$

- The **relational distance $GD(DA, DB)$** is given by

$$GD(DA, DB) = \min E(f)$$

$f : A \rightarrow B, f$ 1-1 and onto

Example



What is the best mapping?

What is the error of the best mapping?

Example

Let $f = \{(1,a),(2,b),(3,c),(4,d)\}$



$$\begin{aligned} |R \circ f - S| &= |\{(a,b)(b,c)(c,d)(d,b)\} - \{(a,b)(b,c)(c,b)(d,b)\}| \\ &= |\{(c,d)\}| = 1 \end{aligned}$$

$$\begin{aligned} |S \circ f^{-1} - R| &= |\{(1,2)(2,3)(3,2)(4,2)\} - \{(1,2)(2,3)(3,4)(4,2)\}| \\ &= |\{(3,2)\}| = 1 \end{aligned}$$

$$E(f) = 1+1 = 2$$

Is there a better mapping?

How to find Relational Distance

Branch and Bound Tree Search

```
Procedure BBTreeSearch(P,L,RP,RL,f,ferr,bestmap,besterr)
/* partial mapping f has error ferr (so far); bestmap is the best
   full mapping so far and has error besterr */
take the first part p in P
  for each label l in L {
    create mapping f' with everything in f plus (p,l)
    compute newerr = ferr + the error from (p,l)
    if newerr < besterr
      if we've used all the parts {bestmap=f; besterr=newerr}
    else BBTreeSearch(P-{p},L-{l}.RP,RL,f',newerr,
                     bestmap,besterr)
  }
```

Variations

- Different weights on different relations
- Normalize error by dividing by total possible
- Attributed relational distance for attributed relations
- Penalizing for NIL mappings

4. Continuous Relaxation

- In discrete relaxation, a label for a unit is either possible or not.
- In continuous relaxation, each (unit, label) pair has a probability.
- Every label for unit i has a prior probability.
- A set of **compatibility coefficients** $C = \{c_{ij}\}$ gives the influence that the label of unit i has on the label of unit j .
- The relationship R is replaced by a set of **unit/label compatibilities** where $r_{ij}(l, l')$ is the compatibility of label l for part i with label l' for part j .
- An **iterative process updates the probability** of each label for each unit in terms of its previous probability and the compatibilities of its current labels and those of other units that influence it.

Geometric Matching



engine model

Is there an engine in the image?
If so, where is it located?




image containing an
instance of the model

How can the engine in the image differ from that in the model?

2D Affine Transformations


1. translation

2. rotation

3. scale

4. skew

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



Point Representation and Transformations

Normal Coordinates for a 2D Point


$$P = [x, y]^t = \begin{bmatrix} x \\ y \end{bmatrix}$$

Homogeneous Coordinates

$$P = [sx, sy, s]^t \quad \text{where } s \text{ is a scale factor}$$

Scaling


$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} c_x & 0 \\ 0 & c_y \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} c_x * x \\ c_y * y \end{bmatrix}$$



scaling by
a factor of 2
about (0,0)

Rotation

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





rotate point

rotate axes

Translation

2 X 2 matrix doesn't work for translation!
Here's where we need homogeneous coordinates.

$$\begin{pmatrix} x' \\ y' \\ 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & x_0 \\ 0 & 1 & y_0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ 1 \end{pmatrix} = \begin{pmatrix} x + x_0 \\ y + y_0 \\ 1 \end{pmatrix}$$



Rotation, Scaling and Translation


$$\begin{pmatrix} x_w \\ y_w \\ 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & x_0 \\ 0 & 1 & y_0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} s & 0 & 0 \\ 0 & s & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x_i \\ y_i \\ 1 \end{pmatrix}$$

T S R




T_R

2D Model and 3 Matching Images of a Boeing Airplane Part



a) Part Model



b) Horizontal Image




c) Rotated Image



d) Rotated and Skewed Image

Computing Affine Transformations between Sets of Matching Points



Given 3 matching pairs of points, the affine transformation can be computed through solving a simple matrix equation.

$$\begin{bmatrix} u1 & u2 & u3 \\ v1 & v2 & v3 \\ 1 & 1 & 1 \end{bmatrix} = \begin{bmatrix} a11 & a12 & a13 \\ a21 & a22 & a23 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x1 & x2 & x3 \\ y1 & y2 & y3 \\ 1 & 1 & 1 \end{bmatrix}$$

A More Robust Approach

Using only 3 points is dangerous, because if even one is off, the transformation can be far from correct.

Instead, use many ($n = 10$ or more) pairs of matching control points to determine a **least squares estimate** of the six parameters of the affine transformation.

Error(a_{11} , a_{12} , a_{13} , a_{21} , a_{22} , a_{23}) =

$$\sum_{j=1,n} \left((a_{11} * x_j + a_{12} * y_j + a_{13} - u_j)^2 + (a_{21} * x_j + a_{22} * y_j + a_{23} - v_j)^2 \right)$$

The Equations to Solve

$$\varepsilon(a_{11}, a_{12}, a_{13}, a_{21}, a_{22}, a_{23}) = \sum_{j=1}^n ((a_{11}x_j + a_{12}y_j + a_{13} - u_j)^2 + (a_{21}x_j + a_{22}y_j + a_{23} - v_j)^2) \quad (11.16)$$

Taking the six partial derivatives of the error function with respect to each of the six variables and setting this expression to zero gives us the six equations represented in matrix form in Equation 11.17 .

$$\begin{bmatrix} \Sigma x_j^2 & \Sigma x_j y_j & \Sigma x_j & 0 & 0 & 0 \\ \Sigma x_j y_j & \Sigma y_j^2 & \Sigma y_j & 0 & 0 & 0 \\ \Sigma x_j & \Sigma y_j & \Sigma 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & \Sigma x_j^2 & \Sigma x_j y_j & \Sigma x_j \\ 0 & 0 & 0 & \Sigma x_j y_j & \Sigma y_j^2 & \Sigma y_j \\ 0 & 0 & 0 & \Sigma x_j & \Sigma y_j & \Sigma 1 \end{bmatrix} \begin{bmatrix} a_{11} \\ a_{12} \\ a_{13} \\ a_{21} \\ a_{22} \\ a_{23} \end{bmatrix} = \begin{bmatrix} \Sigma u_j x_j \\ \Sigma u_j y_j \\ \Sigma u_j \\ \Sigma v_j x_j \\ \Sigma v_j y_j \\ \Sigma v_j \end{bmatrix} \quad (11.17)$$


What is this for?

Many 2D matching techniques use it.

1. Local-Feature Focus Method
2. Pose Clustering
3. Geometric Hashing

Local-Feature-Focus Method

- Each model has a set of features (interesting points).
 - The focus features are the particularly detectable features, usually representing several different areas of the model.
 - Each focus feature has a set of nearby features that can be used, along with the focus feature, to compute the transformation.



LFF Algorithm

Let G be the set of detected image features.

Let F_m be focus features of the model.




Let $S(f)$ be the nearby features for feature f .

for each focus feature F_m



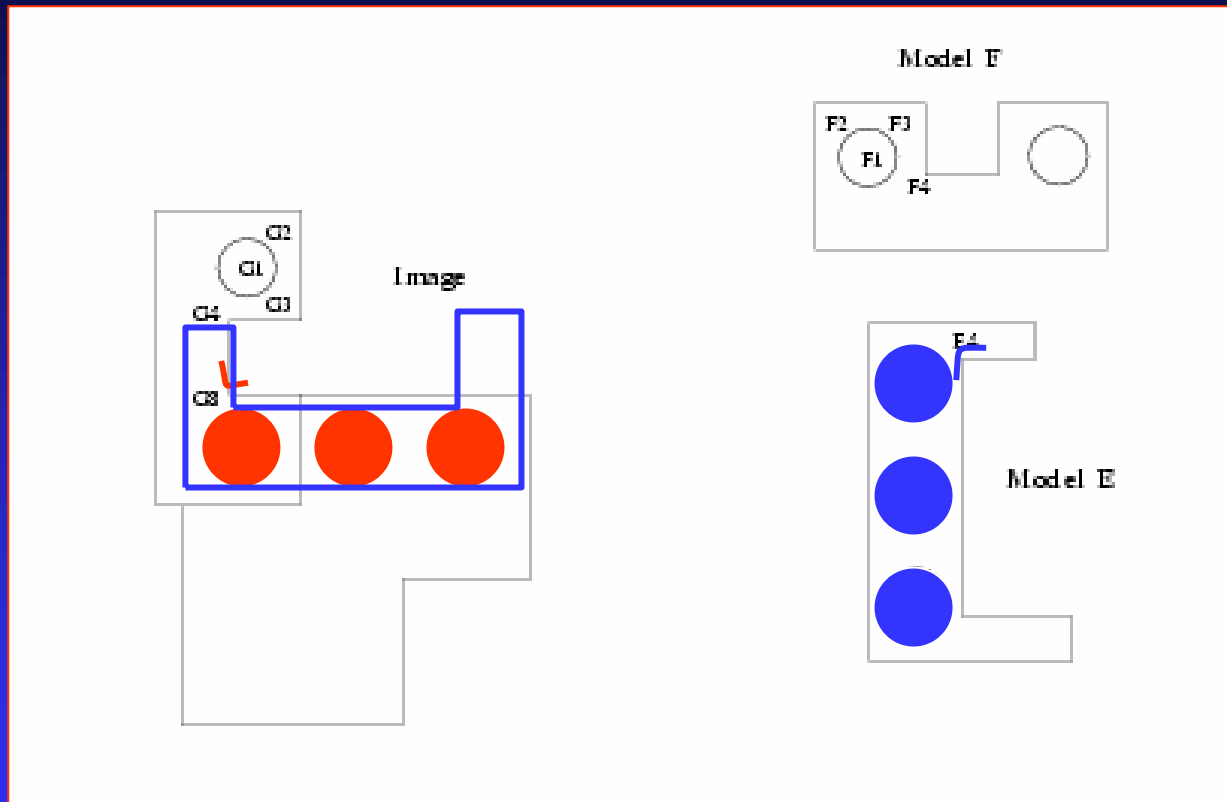
for each image feature G_i of the same type as F_m

1. find the maximal subgraph S_m of $S(F_m)$ that matches a subgraph S_i of $S(G_i)$.
2. Compute transformation T that maps the points of each feature of S_m to the corresponding one of S_i .
3. Apply T to the line segments of the model.
4. If enough transformed segments find evidence in the image, return(T)

Example Match 1: Good Match



Example Match 2: Poor Match



Pose Clustering

Let T be a transformation aligning model M with image object O .

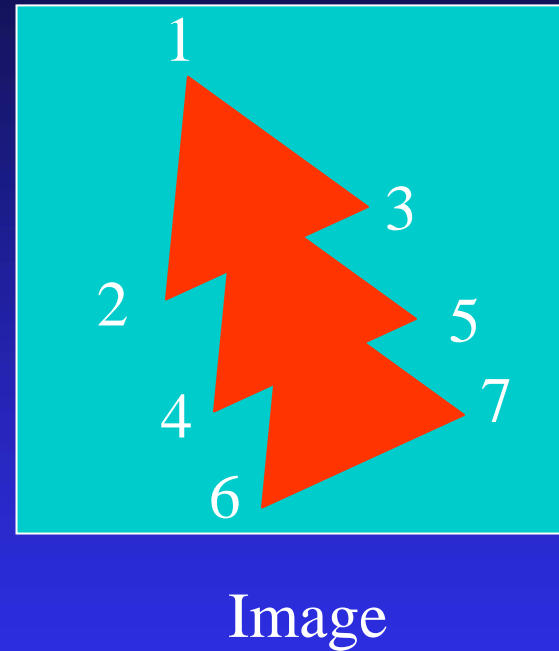
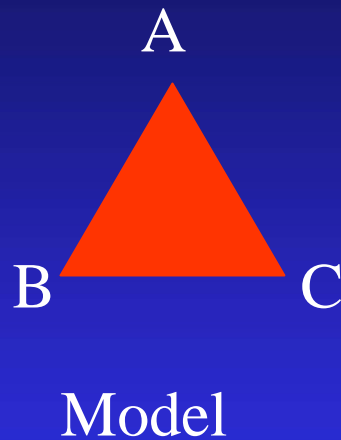
The **pose** of object O is its location and orientation, **defined by T** .

The idea of pose clustering is to **compute lots of possible pose transformations**, each based on 2 points from the model and 2 hypothesized corresponding points from the image.*

Then **cluster all the transformations** in pose space and try to **verify** the large clusters.

* This is not a full affine transformation, just RST.


Pose Clustering




Correct Match: $\text{mapping} = \{ (1,A), (2,B), (3,C) \}$

There will be some votes for $(B,C) \rightarrow (4,5)$, $(B,C) \rightarrow (6,7)$ etc.


Pose Clustering Applied to Detecting a Particular Airplane



(a) original airfield image



(b) model of object




(c) detections matching model

Geometric Hashing

- This method was developed for the case where there is a **whole database of models** to try to find in an image.
- It trades:
 - a large amount of offline preprocessing and a large amount of space
- for potentially fast online
 - object recognition**
 - pose detection**

Theory Behind Geometric Hashing


- A **model M** is a an ordered set of feature points.



$$M = \langle P_1, P_2, P_3, P_4, P_5, P_6, P_7, P_8 \rangle$$

- An **affine basis** is any subset $E = \{e_{00}, e_{01}, e_{10}\}$ of noncollinear points of **M**.
- For basis E , any point $x \in M$ can be represented in **affine coordinates** (ξ, η) .

$$x = \xi(e_{10} - e_{00}) + \eta(e_{01} - e_{00}) + e_{00}$$



Affine Transform

If x is represented in affine coordinates (ξ, η) .

$$\mathbf{x} = \xi(\mathbf{e}_{10} - \mathbf{e}_{00}) + \eta(\mathbf{e}_{01} - \mathbf{e}_{00}) + \mathbf{e}_{00}$$

and we apply affine transform T to point x , we get



$$T\mathbf{x} = \xi(T\mathbf{e}_{10} - T\mathbf{e}_{00}) + \eta(T\mathbf{e}_{01} - T\mathbf{e}_{00}) + T\mathbf{e}_{00}$$

In both cases, x has the same coordinates (ξ, η) .

Example

original object

transformed object



Offline Preprocessing

For each model M

{

 Extract feature point set FM

 for each noncollinear triple E of FM (**basis**)

 for each **other point** x of FM

 {

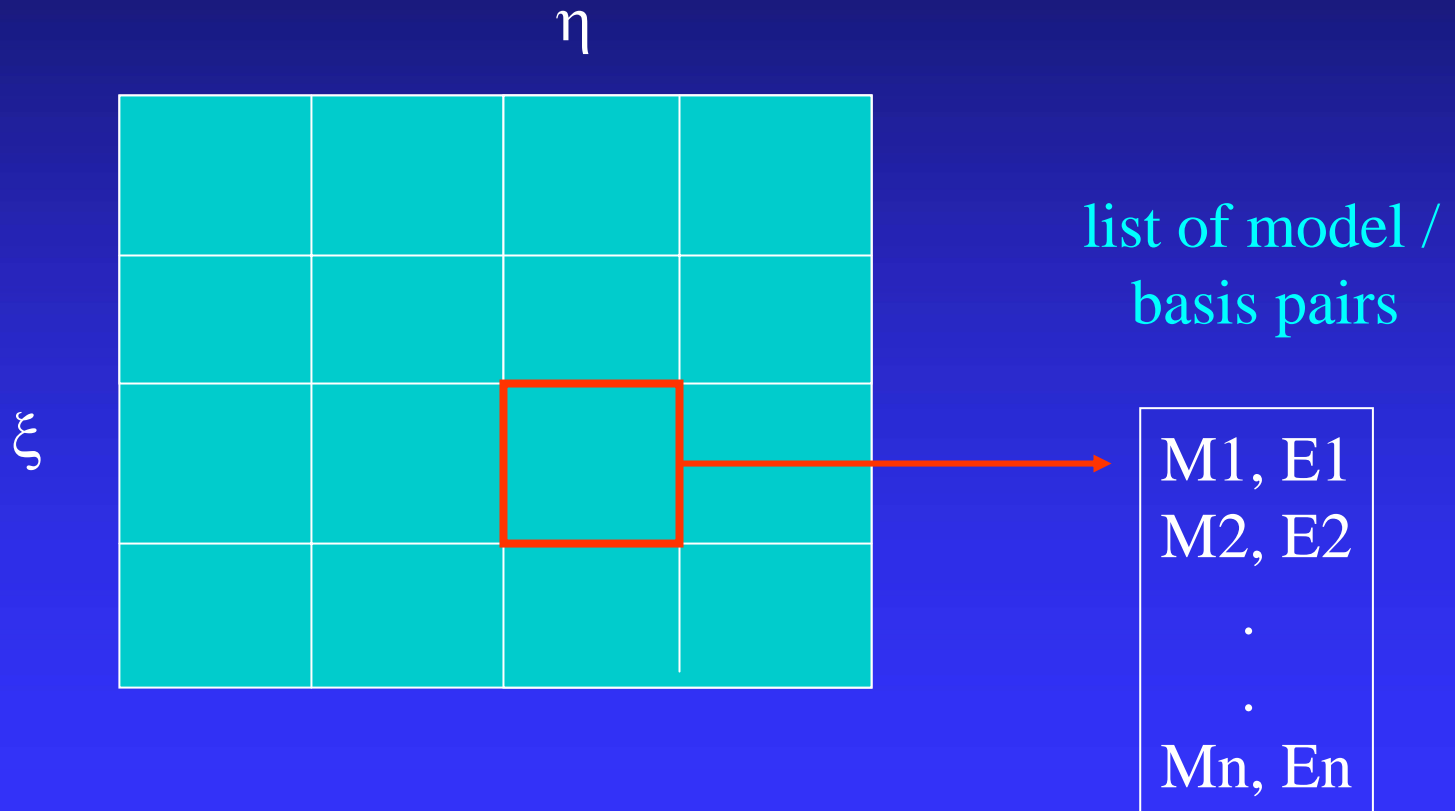
calculate (ξ, η) for x with respect to E

store (M,E) in hash table H at index (ξ, η)

 }

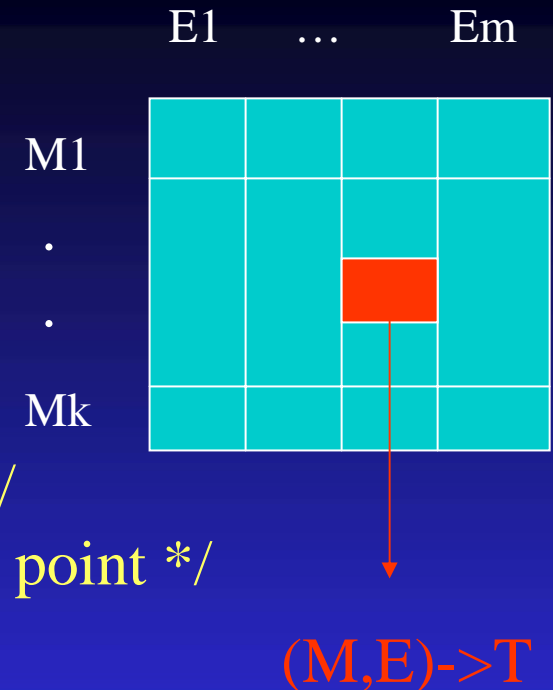
 }

Hash Table



Online Recognition

```
initialize accumulator A to all zero
extract feature points from image
for each basis triple F /* one basis */
  for each other point v /* each image point */
    {
      calculate  $(\xi, \eta)$  for v with respect to F
      retrieve list L from hash table at index  $(\xi, \eta)$ 
      for each pair (M,E) of L
         $A[M,E] = A[M,E] + 1$ 
      }
  find peaks in accumulator array A
  for each peak (M,E) in A
    calculate and try to verify  $T \ni: F = TE$ 
```



Verification

How well does the transformed model line up with the image.

- compare positions of feature points
- compare full line or curve segments

Whole segments work better, allow less **halucination**, but there's a higher cost in execution time.