Further applications of voting

D.A. Forsyth

University of Illinois at Urbana Champaign

Big point about hough transform

- Hough transform:
  - points try to agree on what the line is
  - Mostly, not best way to find lines
    - but can be useful to get line segments to join up

- Forcing hypotheses to agree can be helpful

Instance level classification
(and detection)

Instance level classification is the problem of determining whether a particular object is present in an image. If it is there (wherever it appears) the image is labelled with that object. Instance classification is rather different than *category level categorization*, where one must determine whether any instance of a particular category is present. So, for example, if you have to tell whether your two-year old tabby cat is in a picture, you are doing instance level classification. If you have to tell whether there is a cat in the image, you are doing category level categorization.

Example: Find the book

- Get pictures of book covers, and attach book name

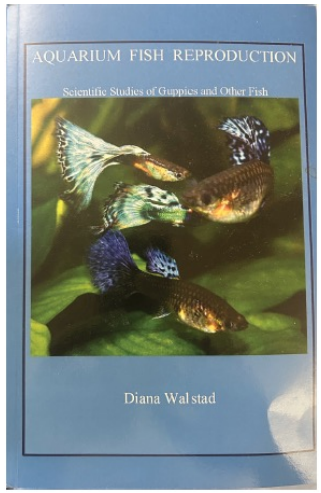

- Now find the book cover in some other image

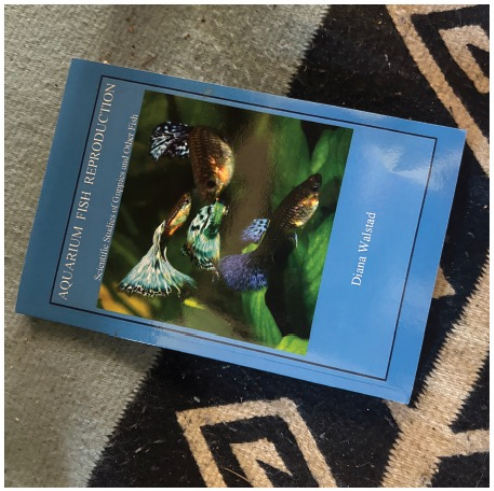

Can’t do this with SSD

- Changes in color caused by changes in lighting

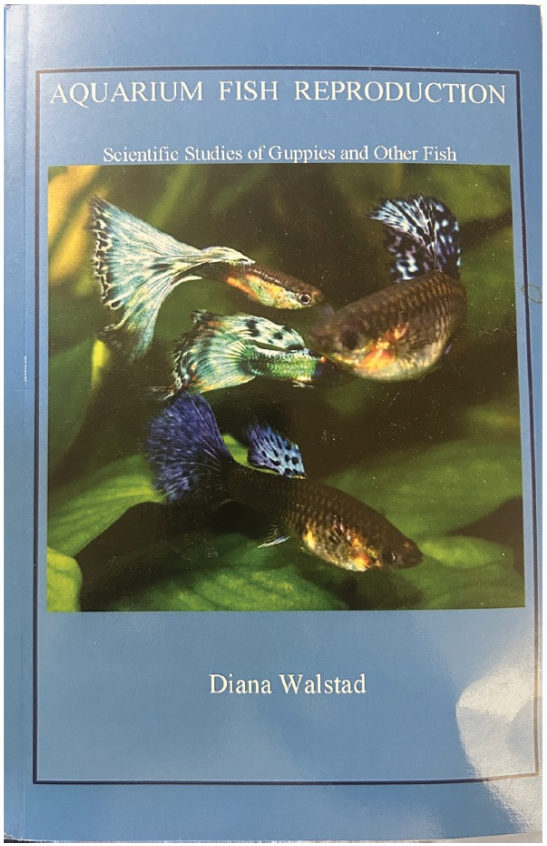

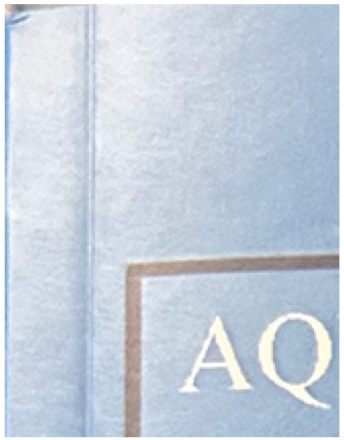

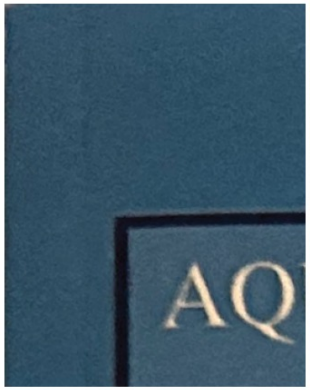

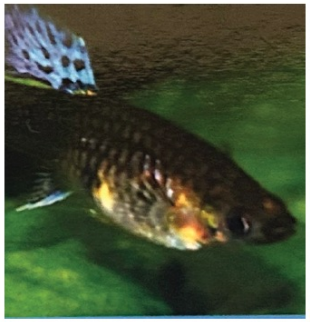

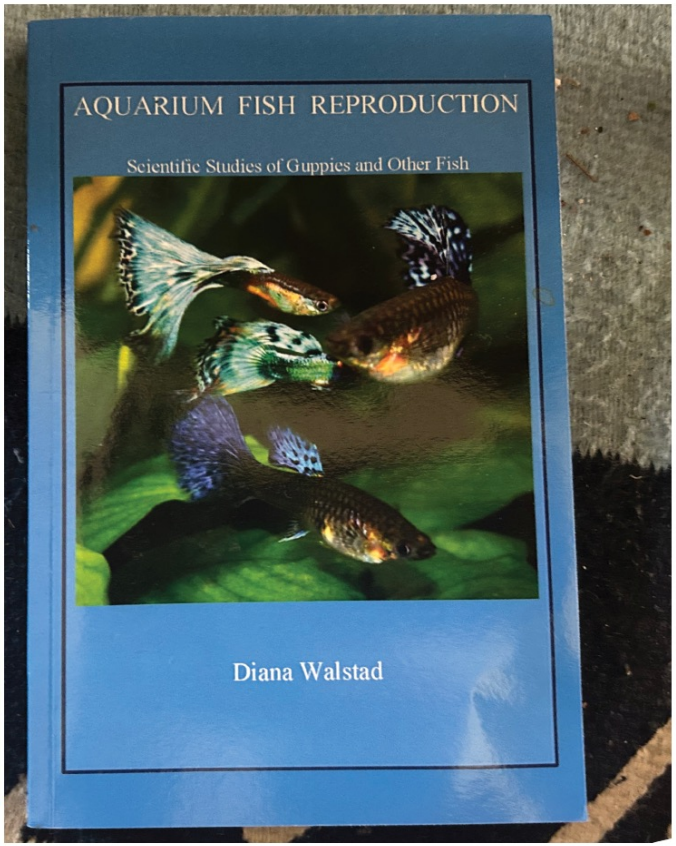

Can’t do this with SSD

- Changes in appearance caused by rotation

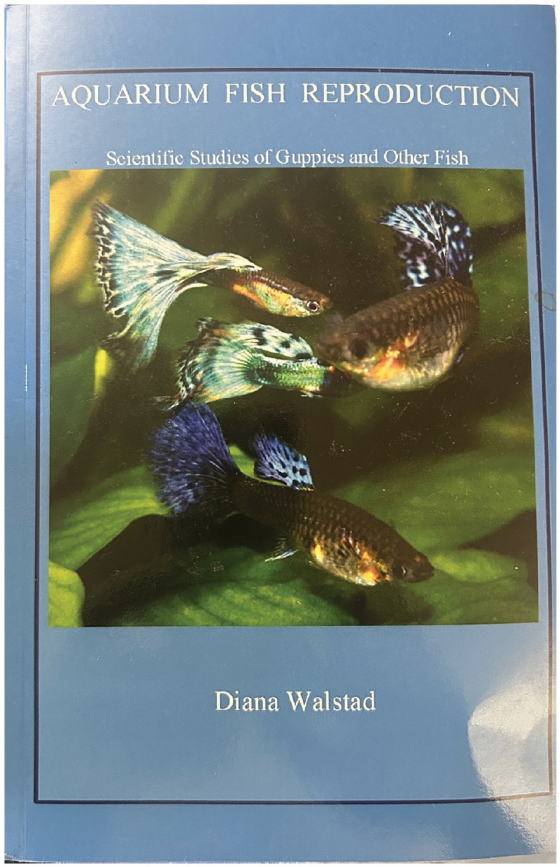

Using interest points

- Make a tree of cover image interest points (with labels)

You can do this cause interest point descriptors are vectors of fixed length. They're nearby when interest points are similar.

- For test image, pass interest points down tree, vote

- Vote how?:
- most common book in leaf
- one vote for each book in leaf
- vote for ANN interest point label
- Issues:
- every patch votes, so there must be junk votes (easily fixed)
- tree doesn't know you're using it for labelling, may not be efficient
- ~~rotation (perhaps insert patches at many rotations?)~~

Using interest points

- Make a tree of cover image interest points (with labels)

  You can do this cause interest point descriptors are vectors of fixed length. They're nearby when interest points are similar.

- For test image, pass interest points down tree, vote
  - Vote how?:
    - most common book in leaf
    - one vote for each book in leaf
    - vote for ANN interest point label
- Issues:
  - every patch votes, so there must be junk votes (easily fixed)
  - tree doesn't know you're using it for labelling, may not be efficient

Generalized Hough transform

- Each interest point has a coordinate system,

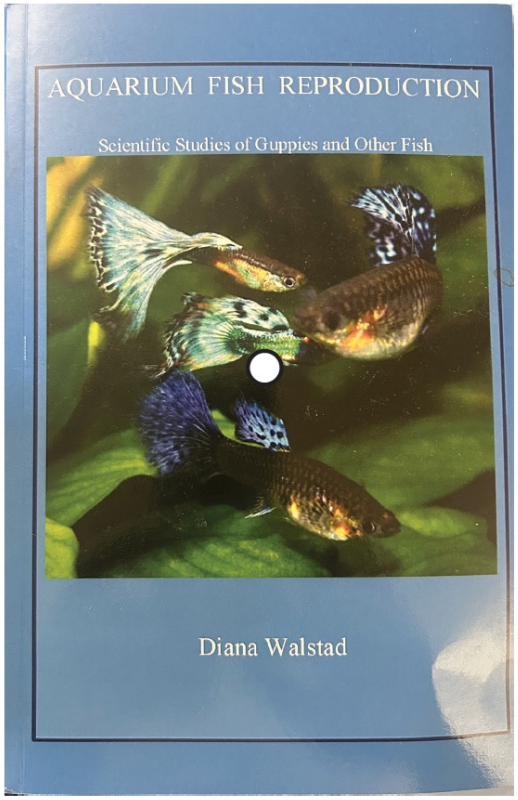

Interest points

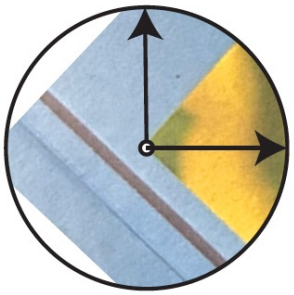

Center in interest point coordinate system

Generalized Hough transform

- Allow an interest point to vote for a book only if another one agrees with center

Generalized Hough transform

- Allow an interest point to vote for a book only if another one agrees with center

- How?
  - build an accumulator array for book centers,
  - accumulate votes
  - scan it to check for “big” votes – do they agree on book label?

Generalized Hough transform

- Interest points know other stuff – the outline of the book

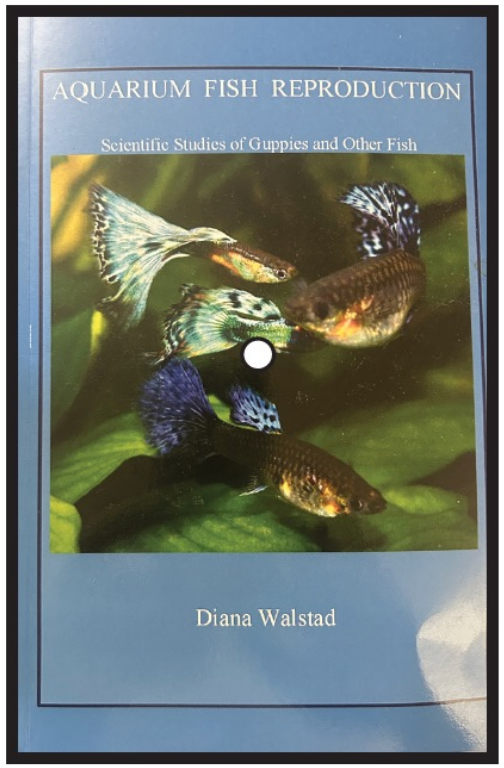

Interest points

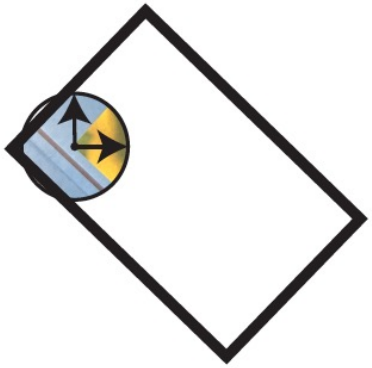

Bounding box in interest point coordinate system

We now have…

- A primitive classifier
- What book is in this image?

- A primitive detector
- What books are in this image, and where are they?

- There are more accurate technologies,
- tend to be slower, require more compute

- This one is extremely fast and very efficient, particularly with an improved tree

The tree

- Current construction:
- Hierarchical k-means
- This breaks up each node into leaves whose elements are close to one another.
- But this isn't what we want
- Decision tree:
- Break each node into leaves where leaves are "informative" about label
- eg in binary tree, labels on the left are all different from labels on the right

Decision trees

- For us:
- tree is now binary
- th

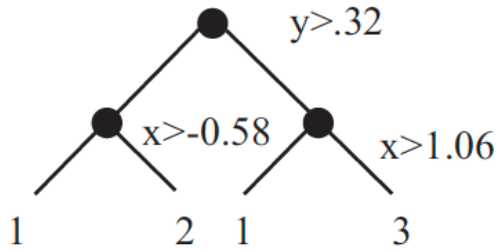

As tree

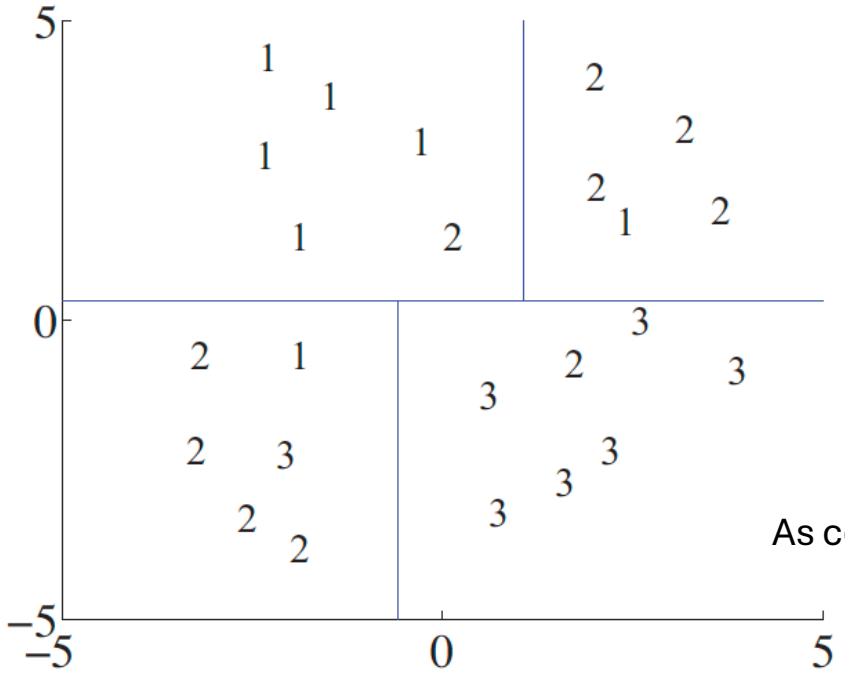

As cells in 2D