The background is a dark grey-green color with faint, light-colored sketches of various scientific and technical objects. These include a globe, a microscope, a telescope, a satellite, a book, a percentage sign, and various geometric shapes and lines. The sketches are rendered in a style similar to chalk or light pencil on a dark surface.

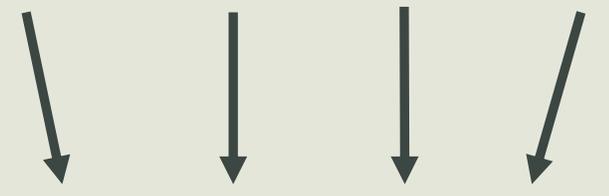
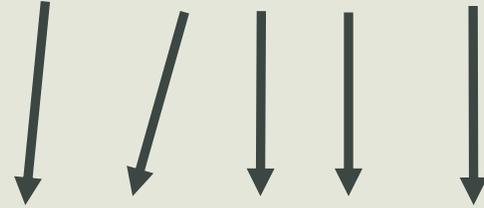
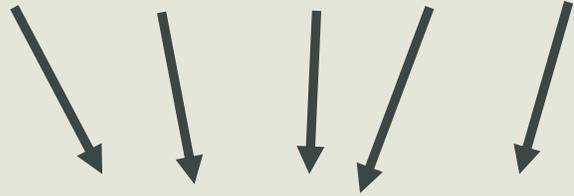
SHARING VISUAL FEATURES FOR MULTICLASS AND MULTIVIEW OBJECT DETECTION

**ANTONIO TORRALBA, KEVIN P. MURPHY,
WILLIAM T. FREEMAN**

THE PROBLEM

- ❖ In the previous papers we learned how to detect and classify images with 1 class.
- ❖ Our Recognition Problem is to recognize many object category with few images per category.
- ❖ How to detect many categories?

FIND EACH CATEGORY



...



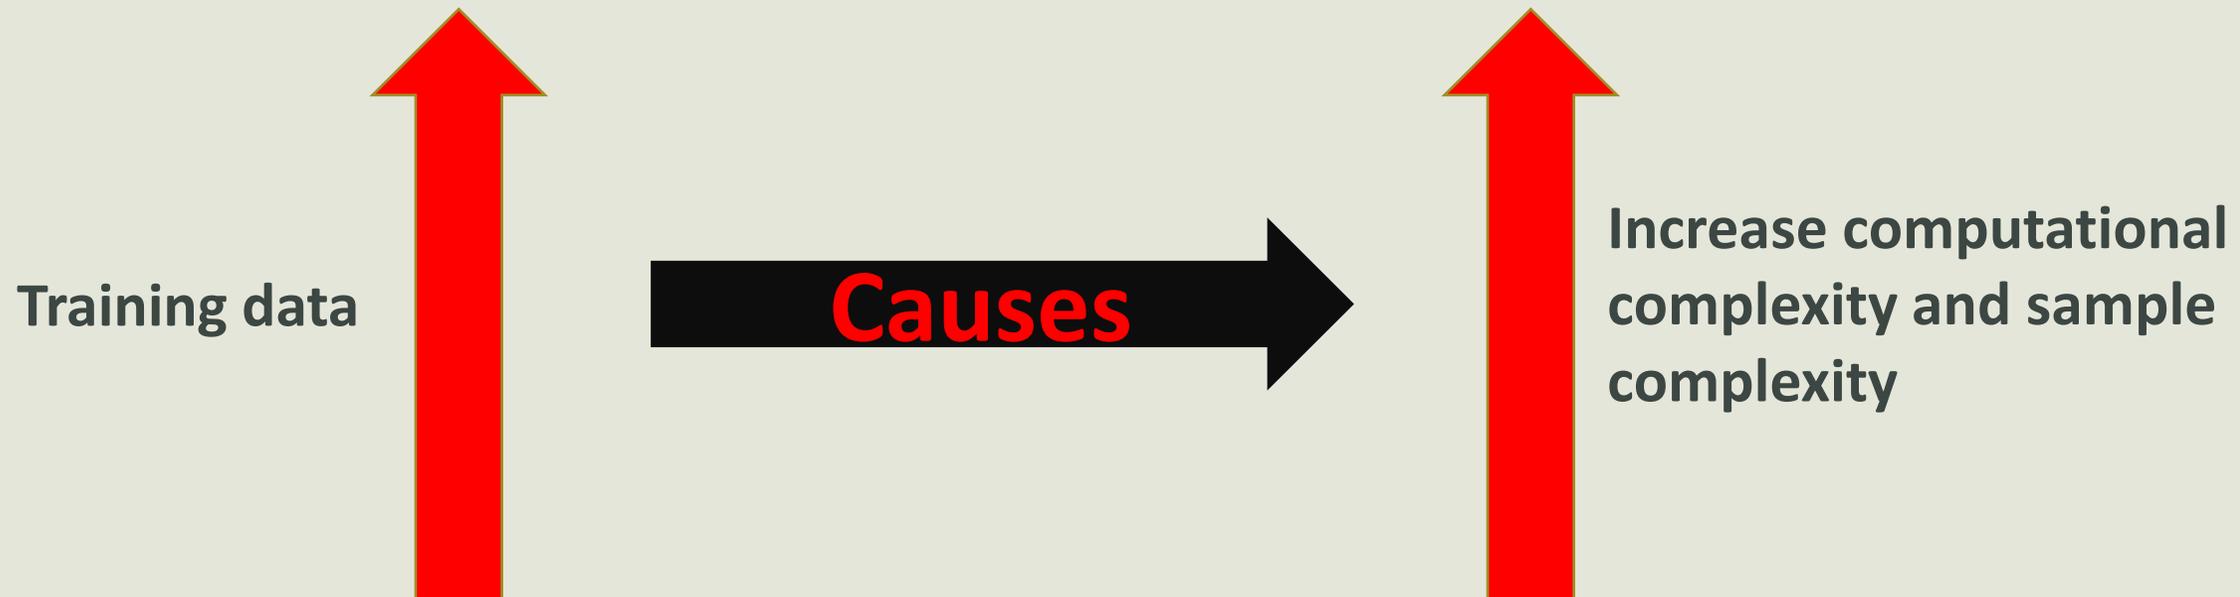
...





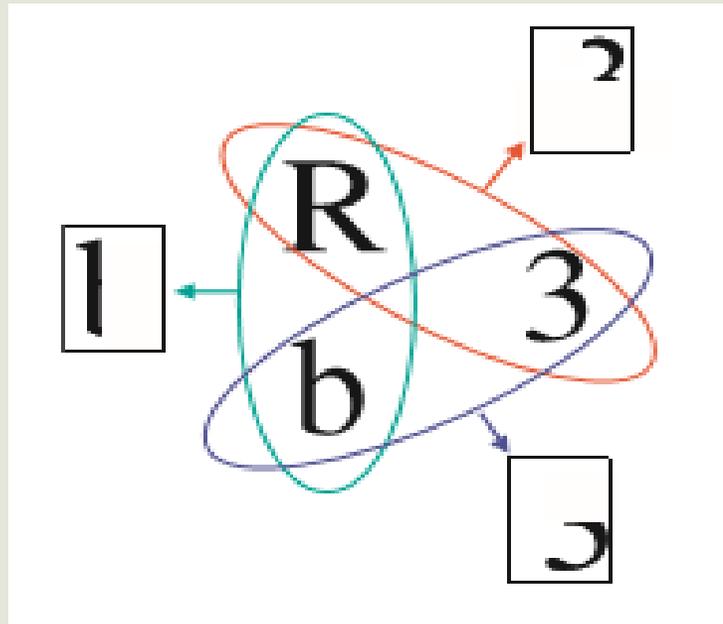
WHY NOT?

- ❖ Require a lot of training data since each classifier require many different images.



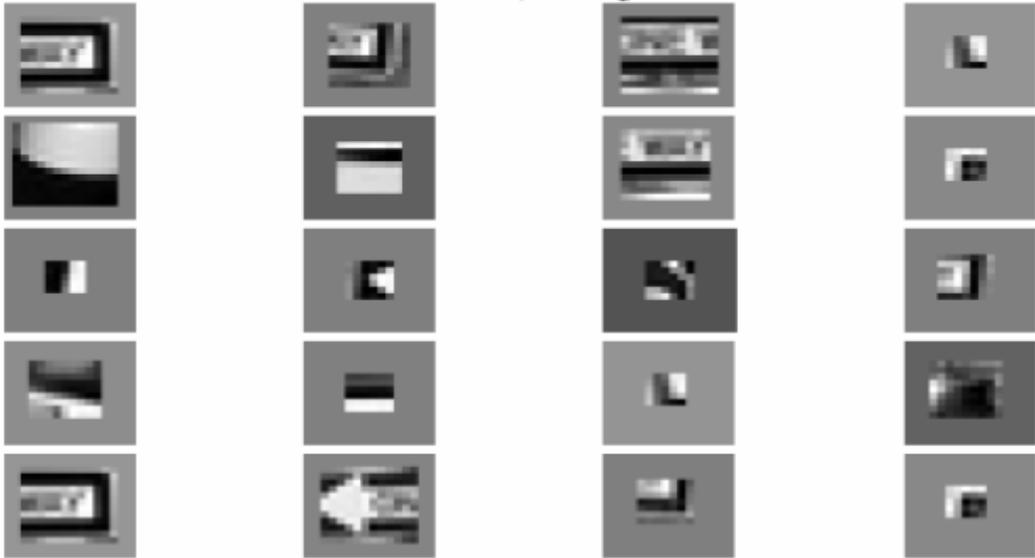
THE SOLUTION - FEATURE SHARING

- Find common features that distinguish a subset of classes against the rest.

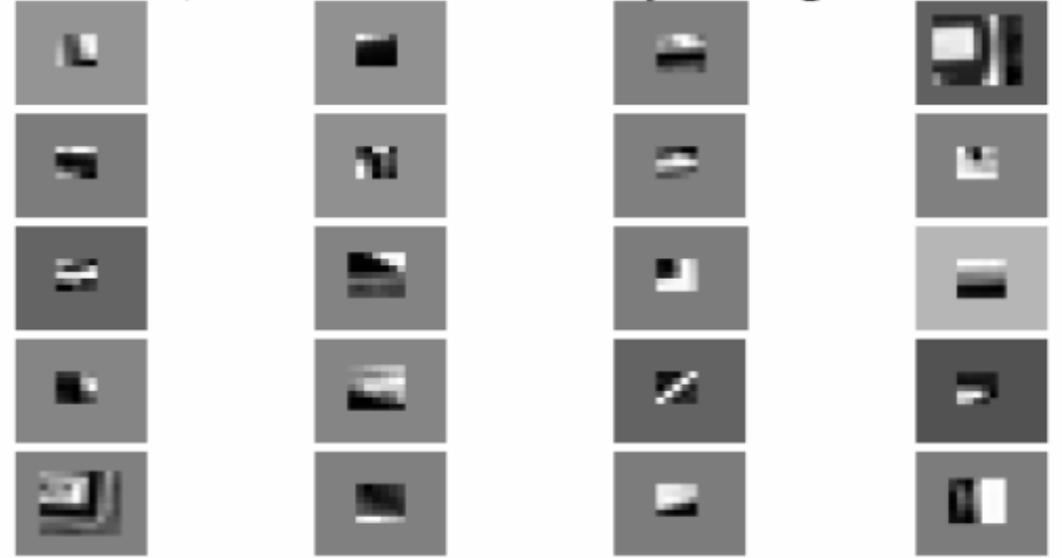




a) Object



b) Selected features by a single detector



c) Selected features when trained jointly

Template-like features

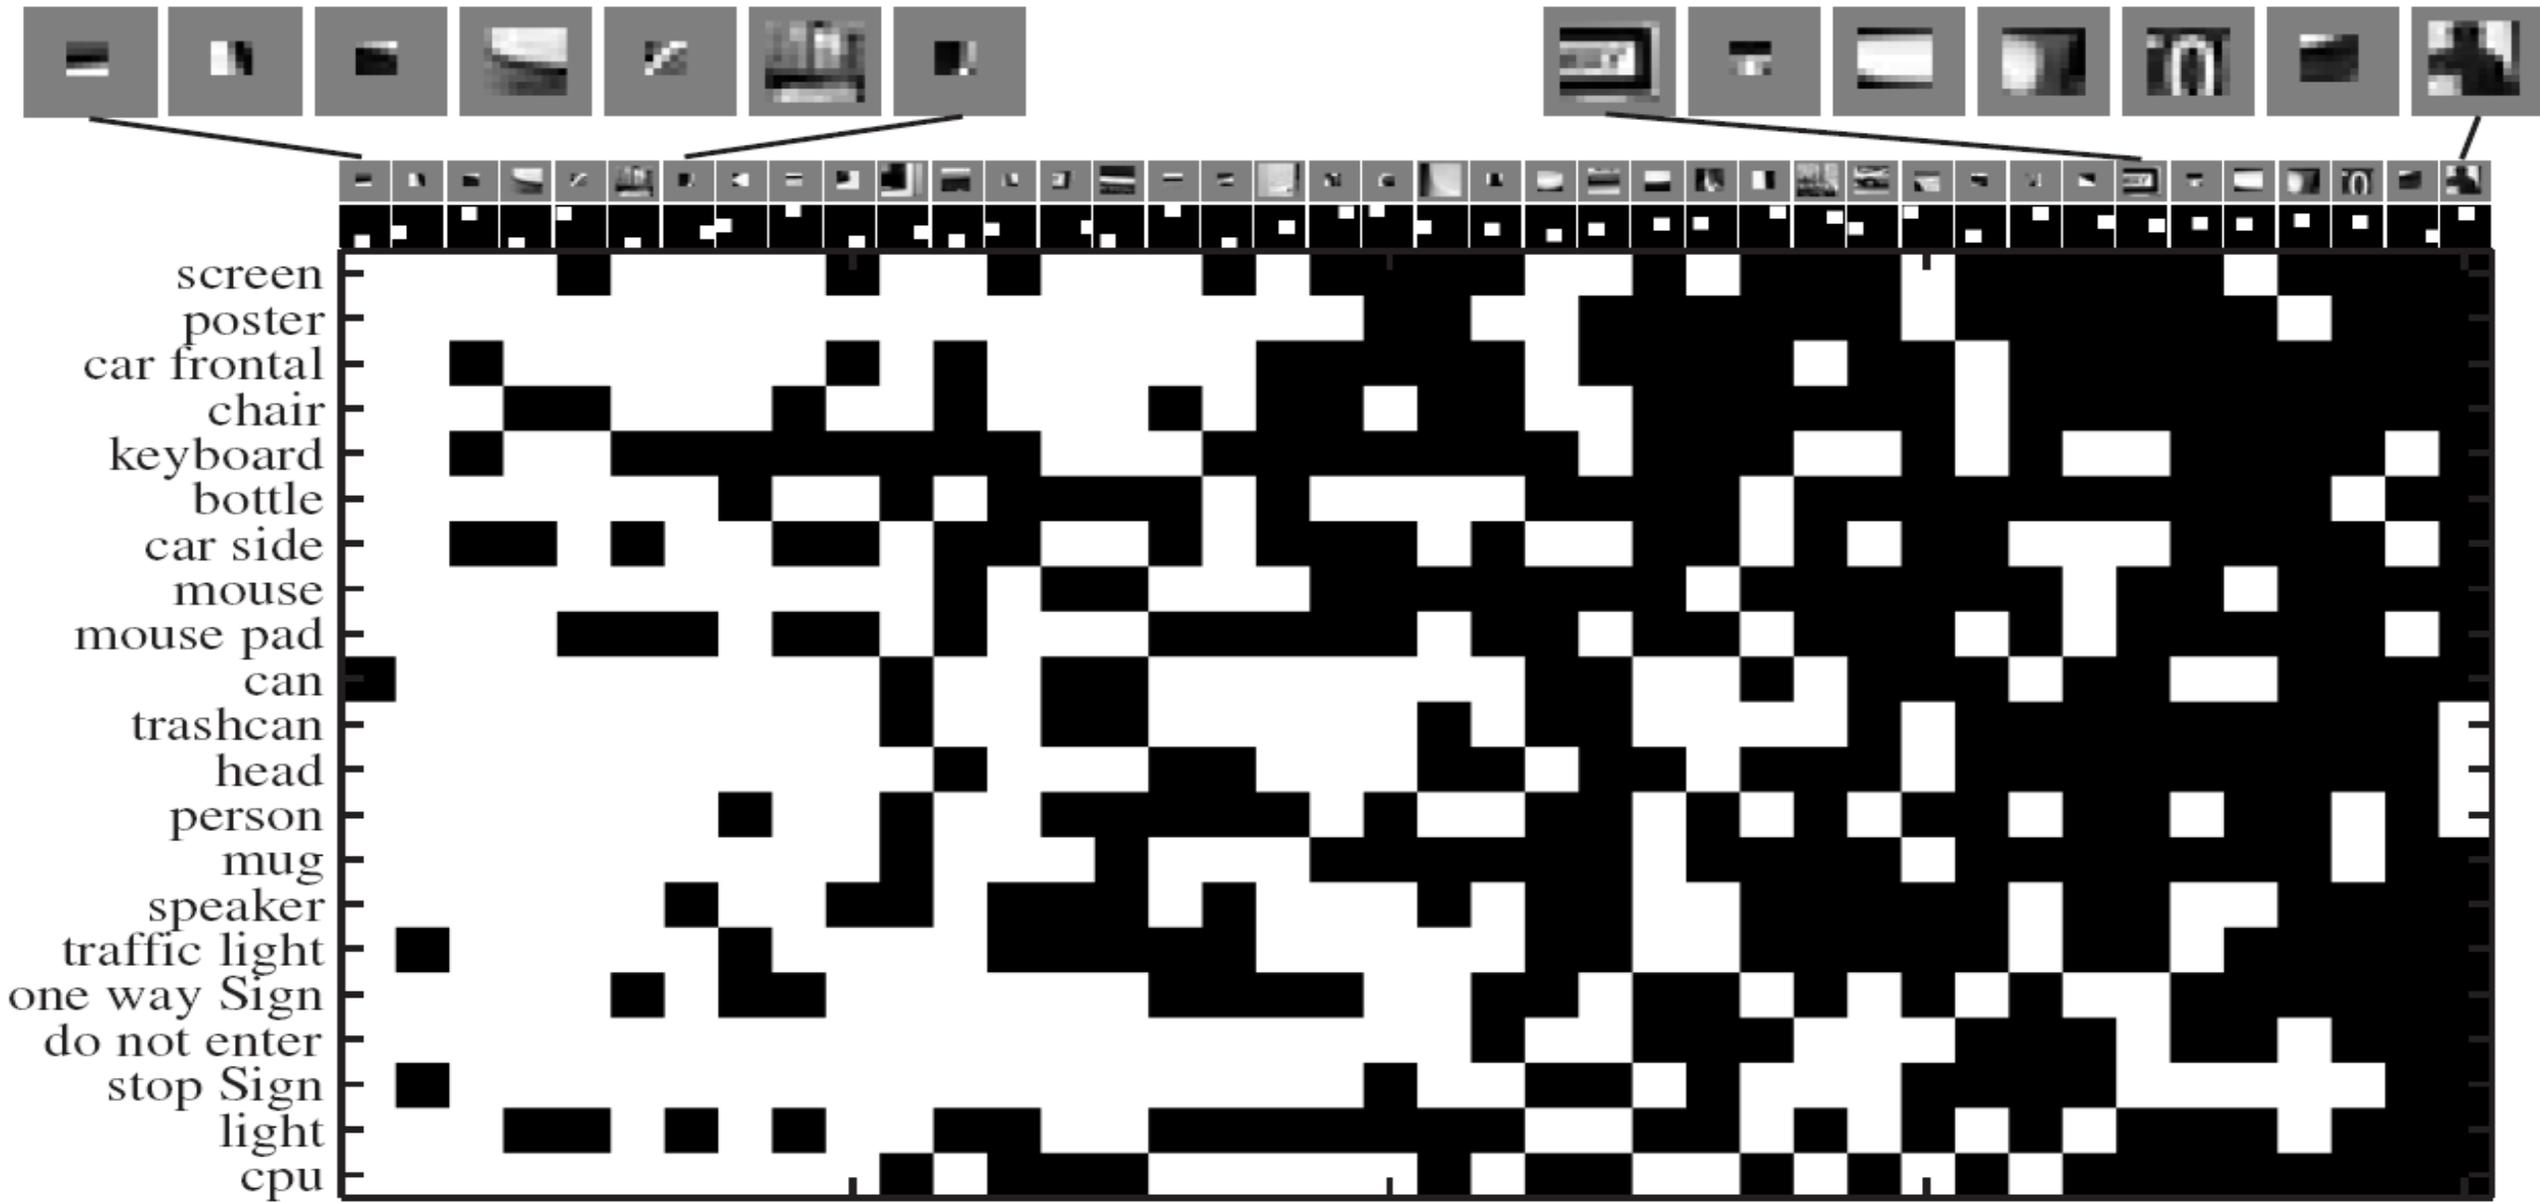
100% accuracy for a **single** object

But too **specific**.

Wavelet-like features,

weaker discriminative power

but shared in **many** classes.



Left part: common features which are being shared in most classes

Right part: specific features which are being shared by only a few classes

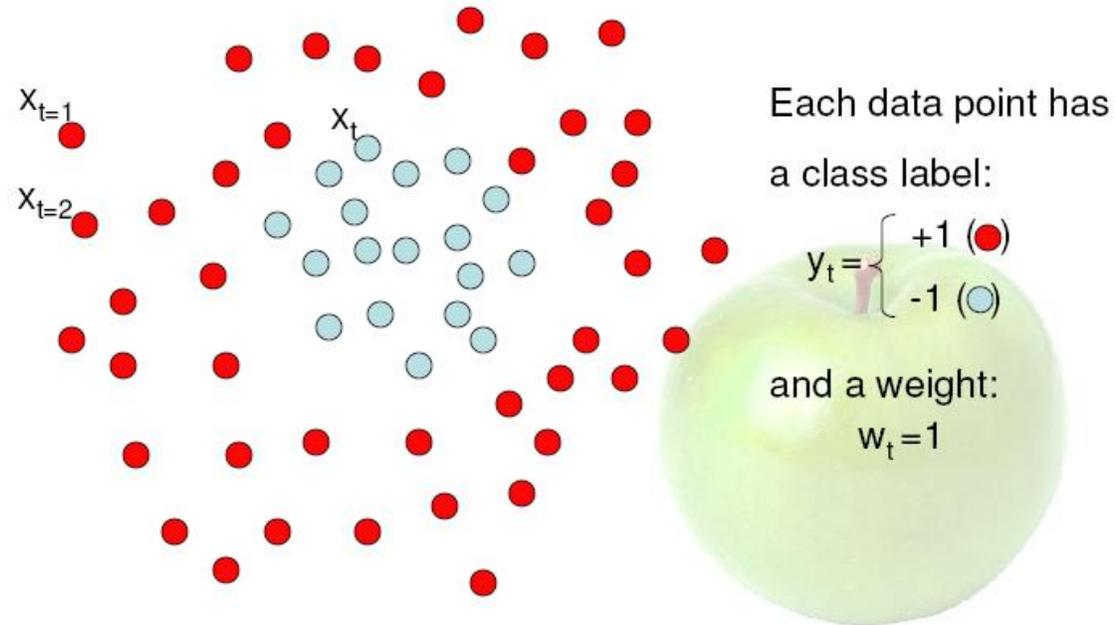
WHAT WE GAIN FROM SHARING FEATURES?

- Essentially more positive samples
- Reuse the data

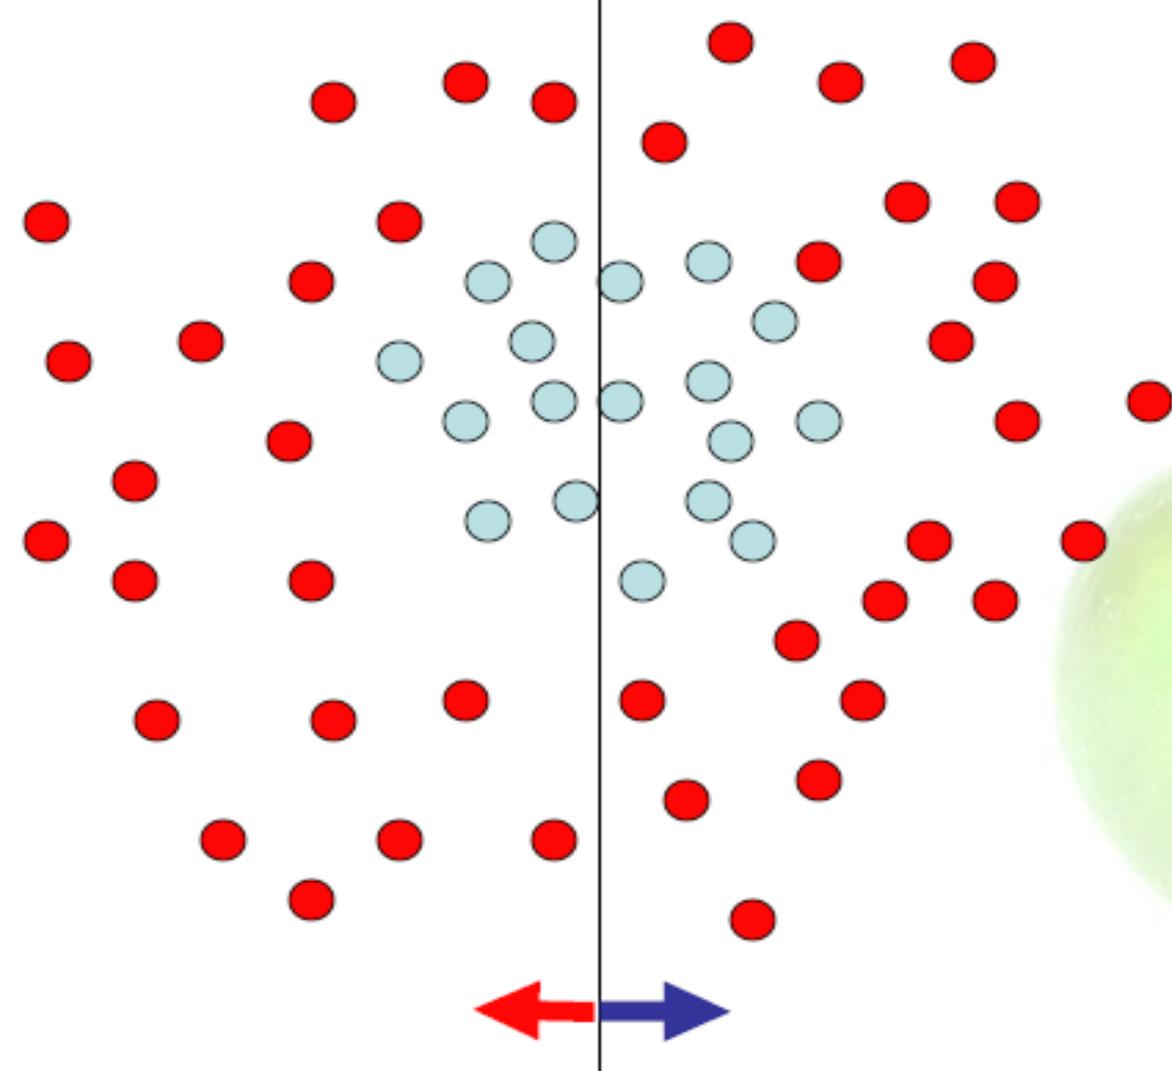
EXAMPLE OF THE STRONG CLASSIFIER

Boosting

- It is a sequential procedure:



Weak learners from the family of lines

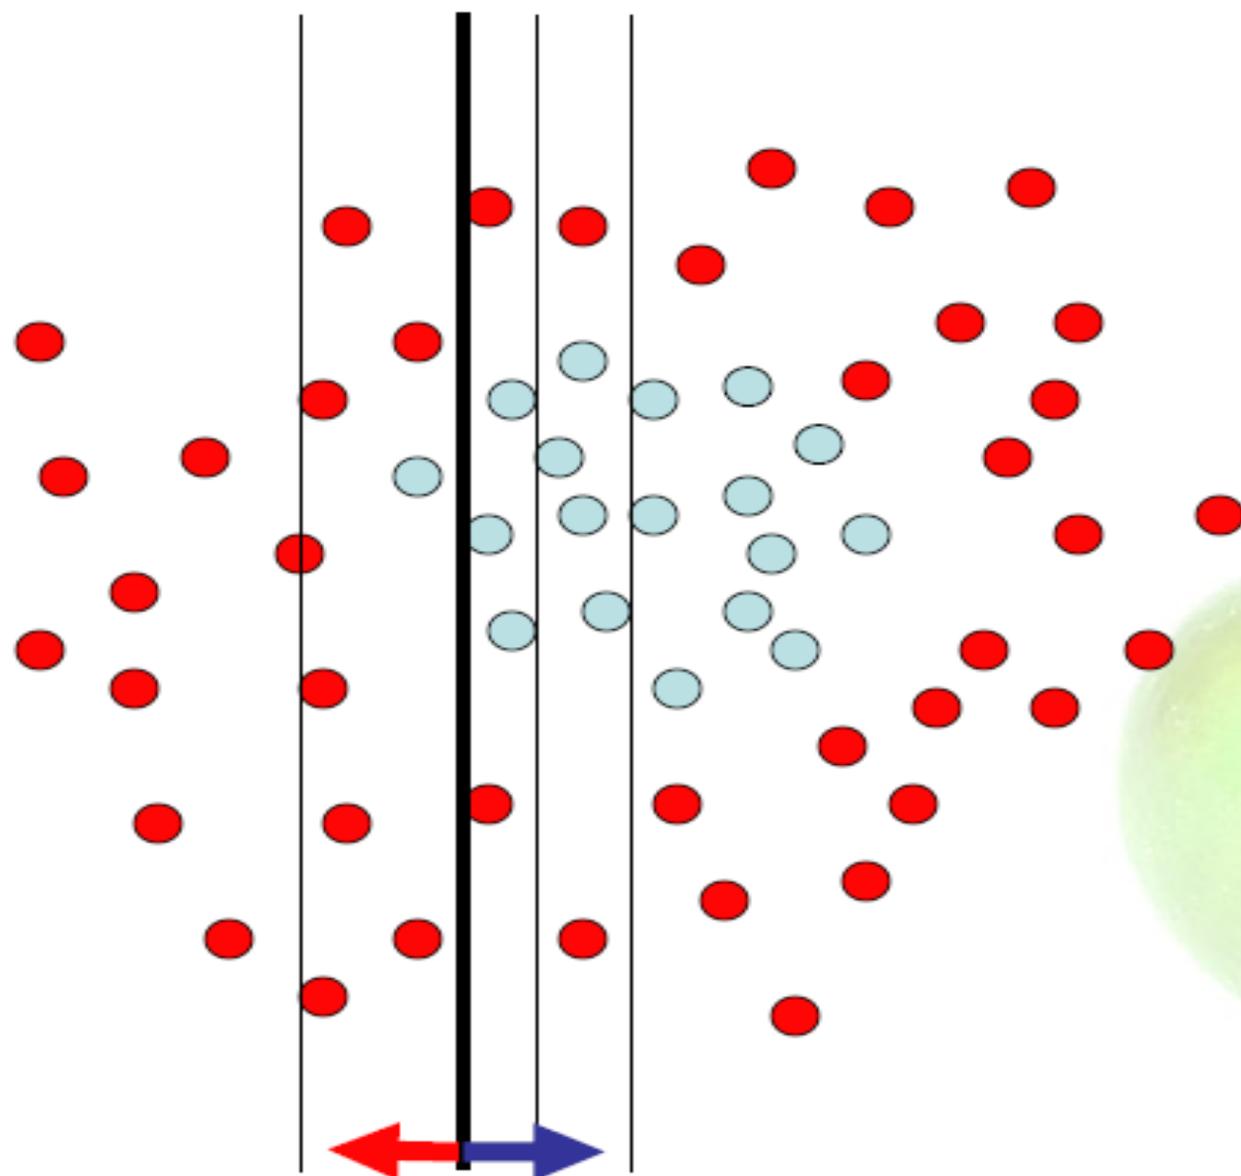


Each data point has a class label:

$$y_t = \begin{cases} +1 (\bullet) \\ -1 (\circ) \end{cases}$$

and a weight:
 $w_t = 1$

$h \Rightarrow p(\text{error}) = 0.5$ it is at chance



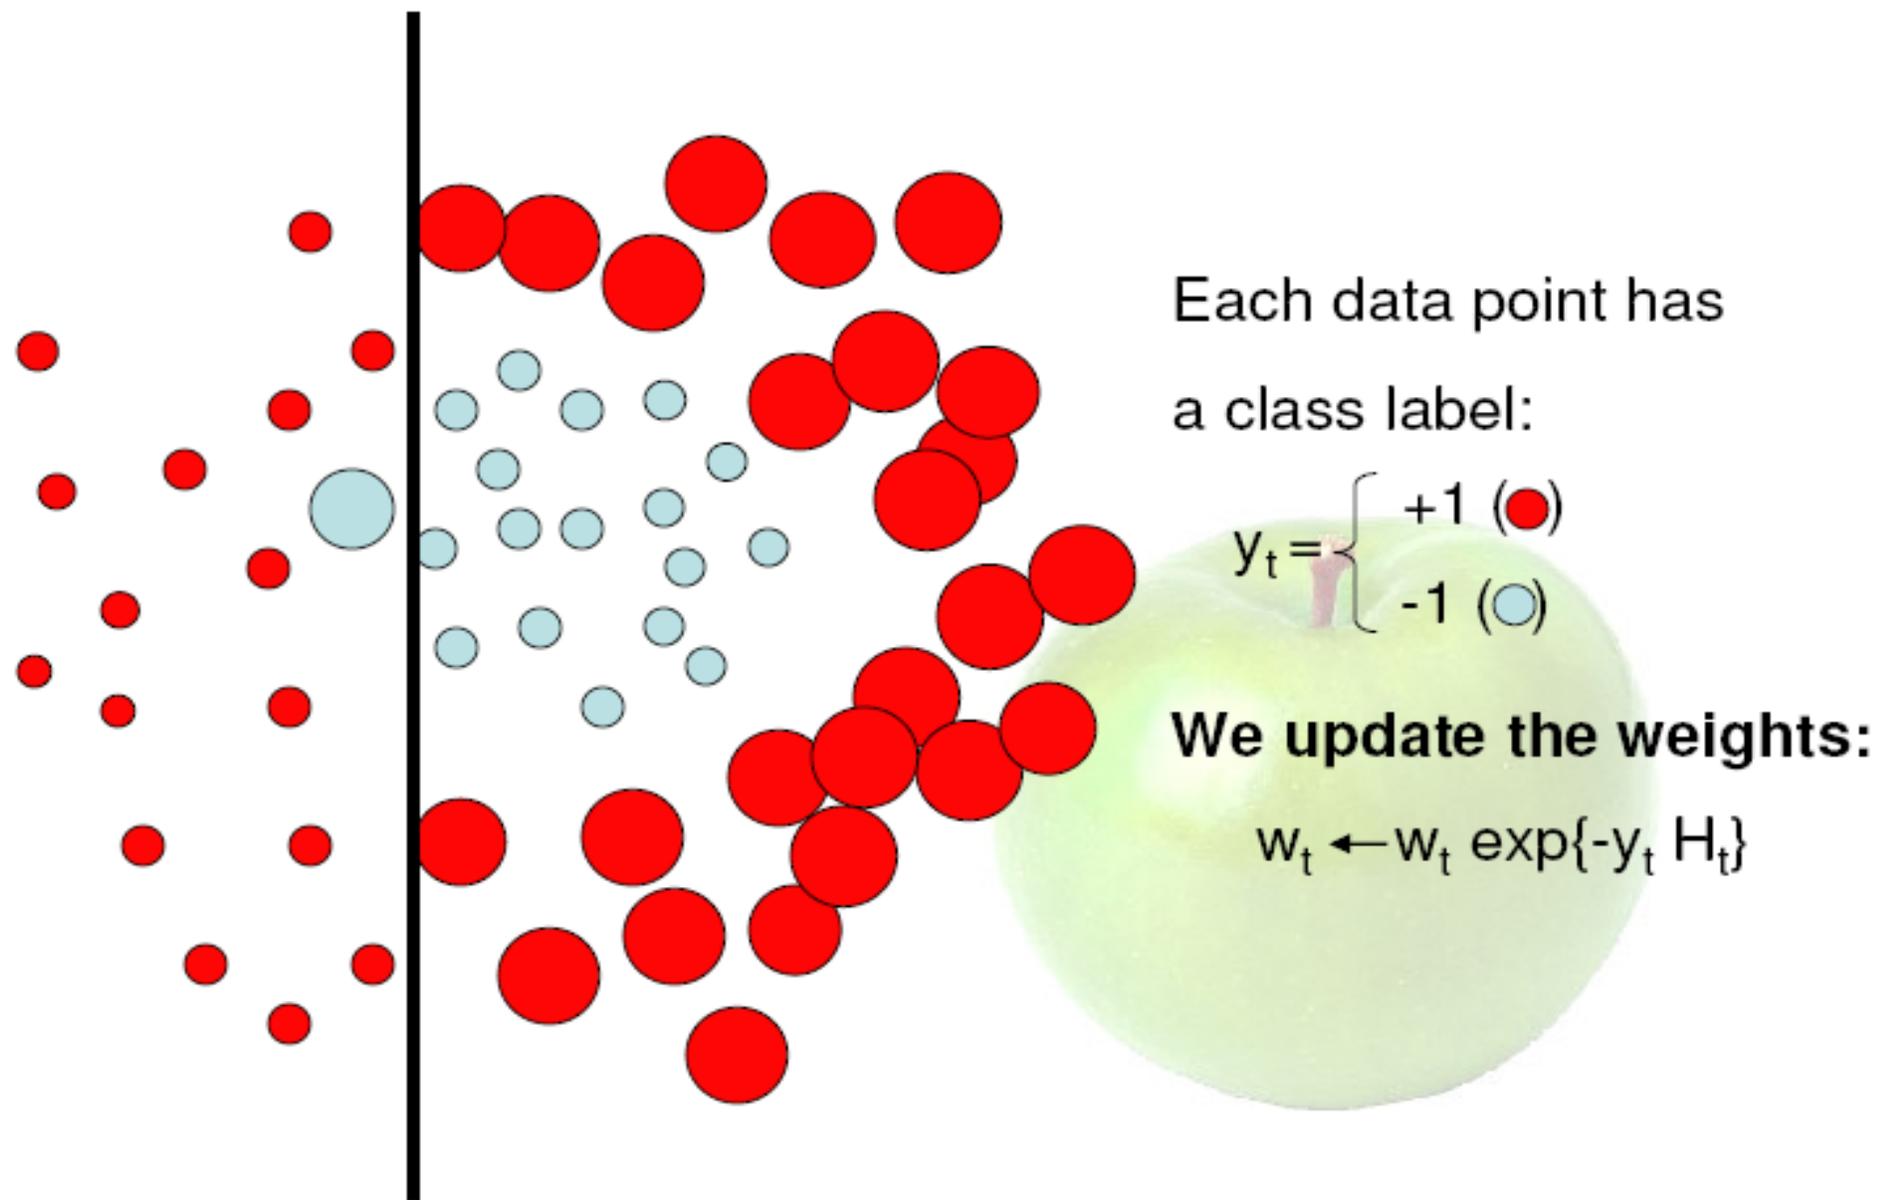
Each data point has a class label:

$$y_t = \begin{cases} +1 & (\bullet) \\ -1 & (\circ) \end{cases}$$

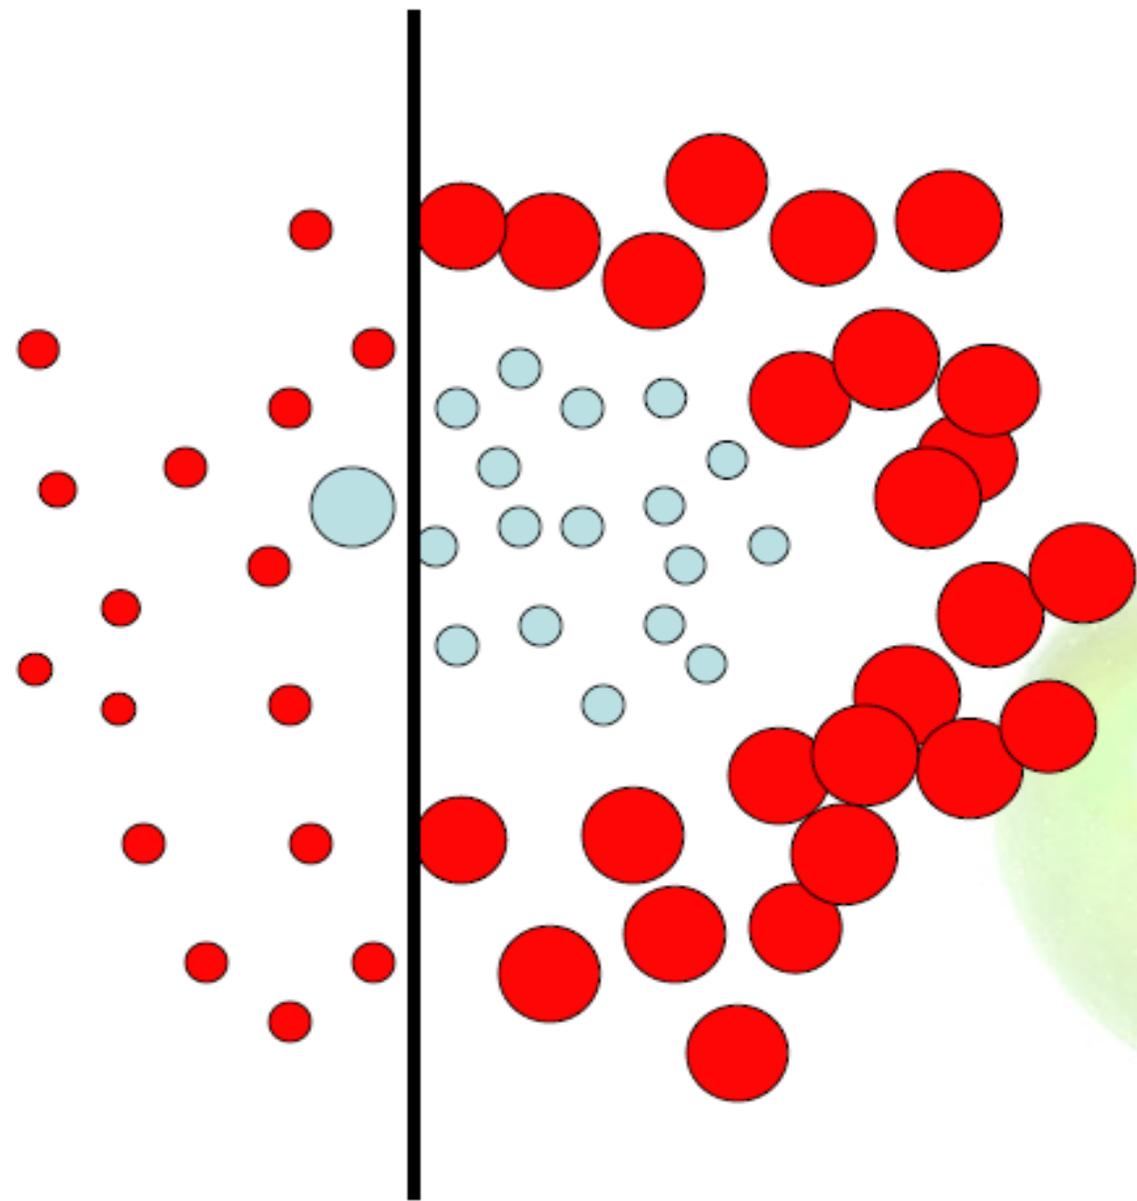
and a weight:
 $w_t = 1$

This one seems to be the best

This is a **'weak classifier'**: It performs slightly better than chance.



We set a new problem for which the previous weak classifier performs at chance again



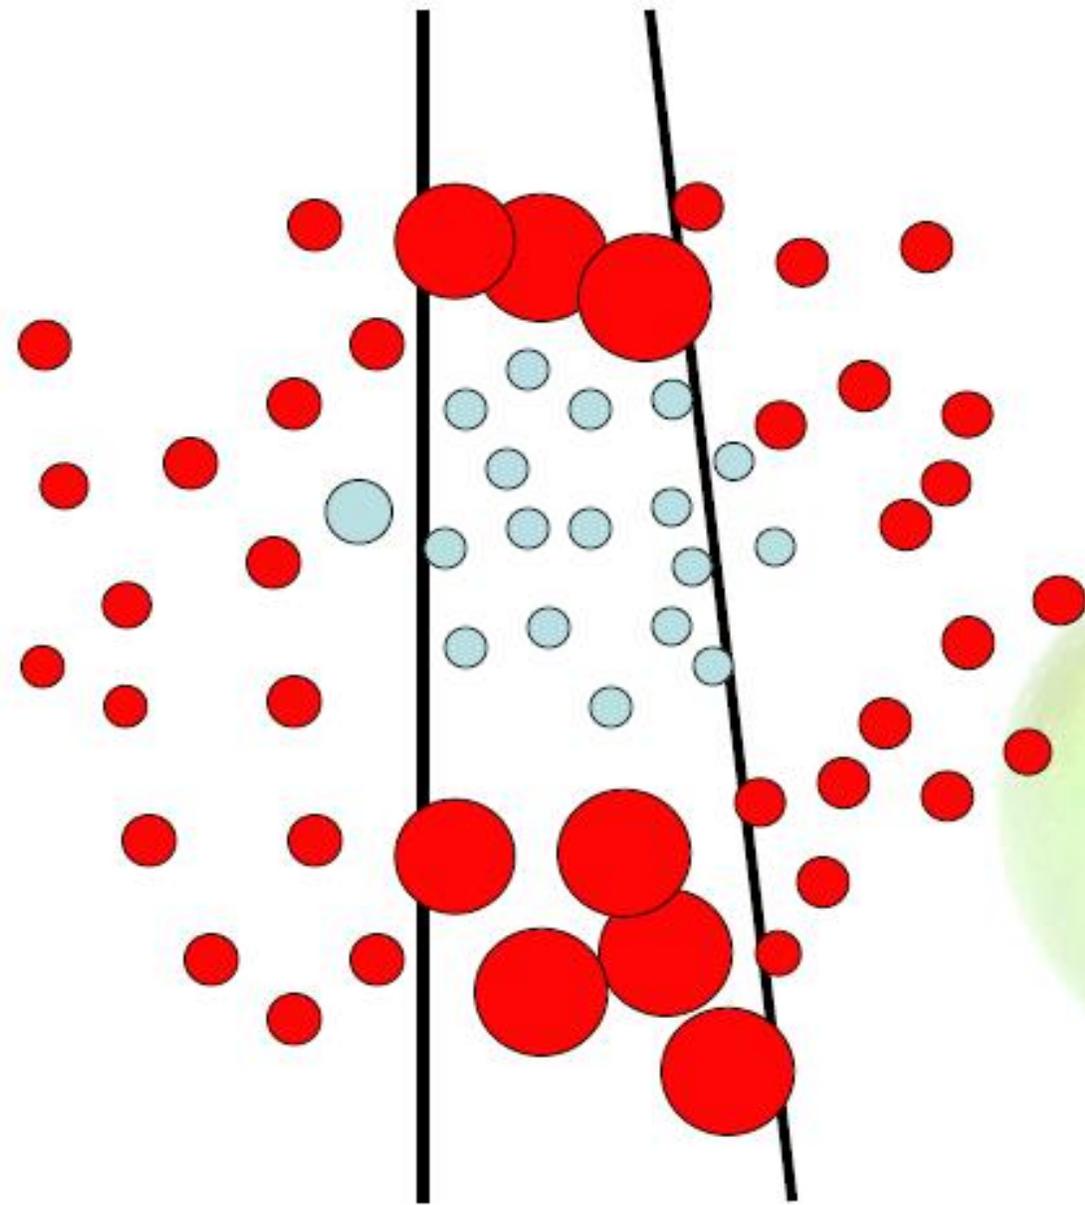
Each data point has
a class label:

$$y_t = \begin{cases} +1 & (\bullet) \\ -1 & (\circ) \end{cases}$$

We update the weights:

$$w_t \leftarrow w_t \exp\{-y_t H_t\}$$

We set a new problem for which the previous weak classifier performs at chance again



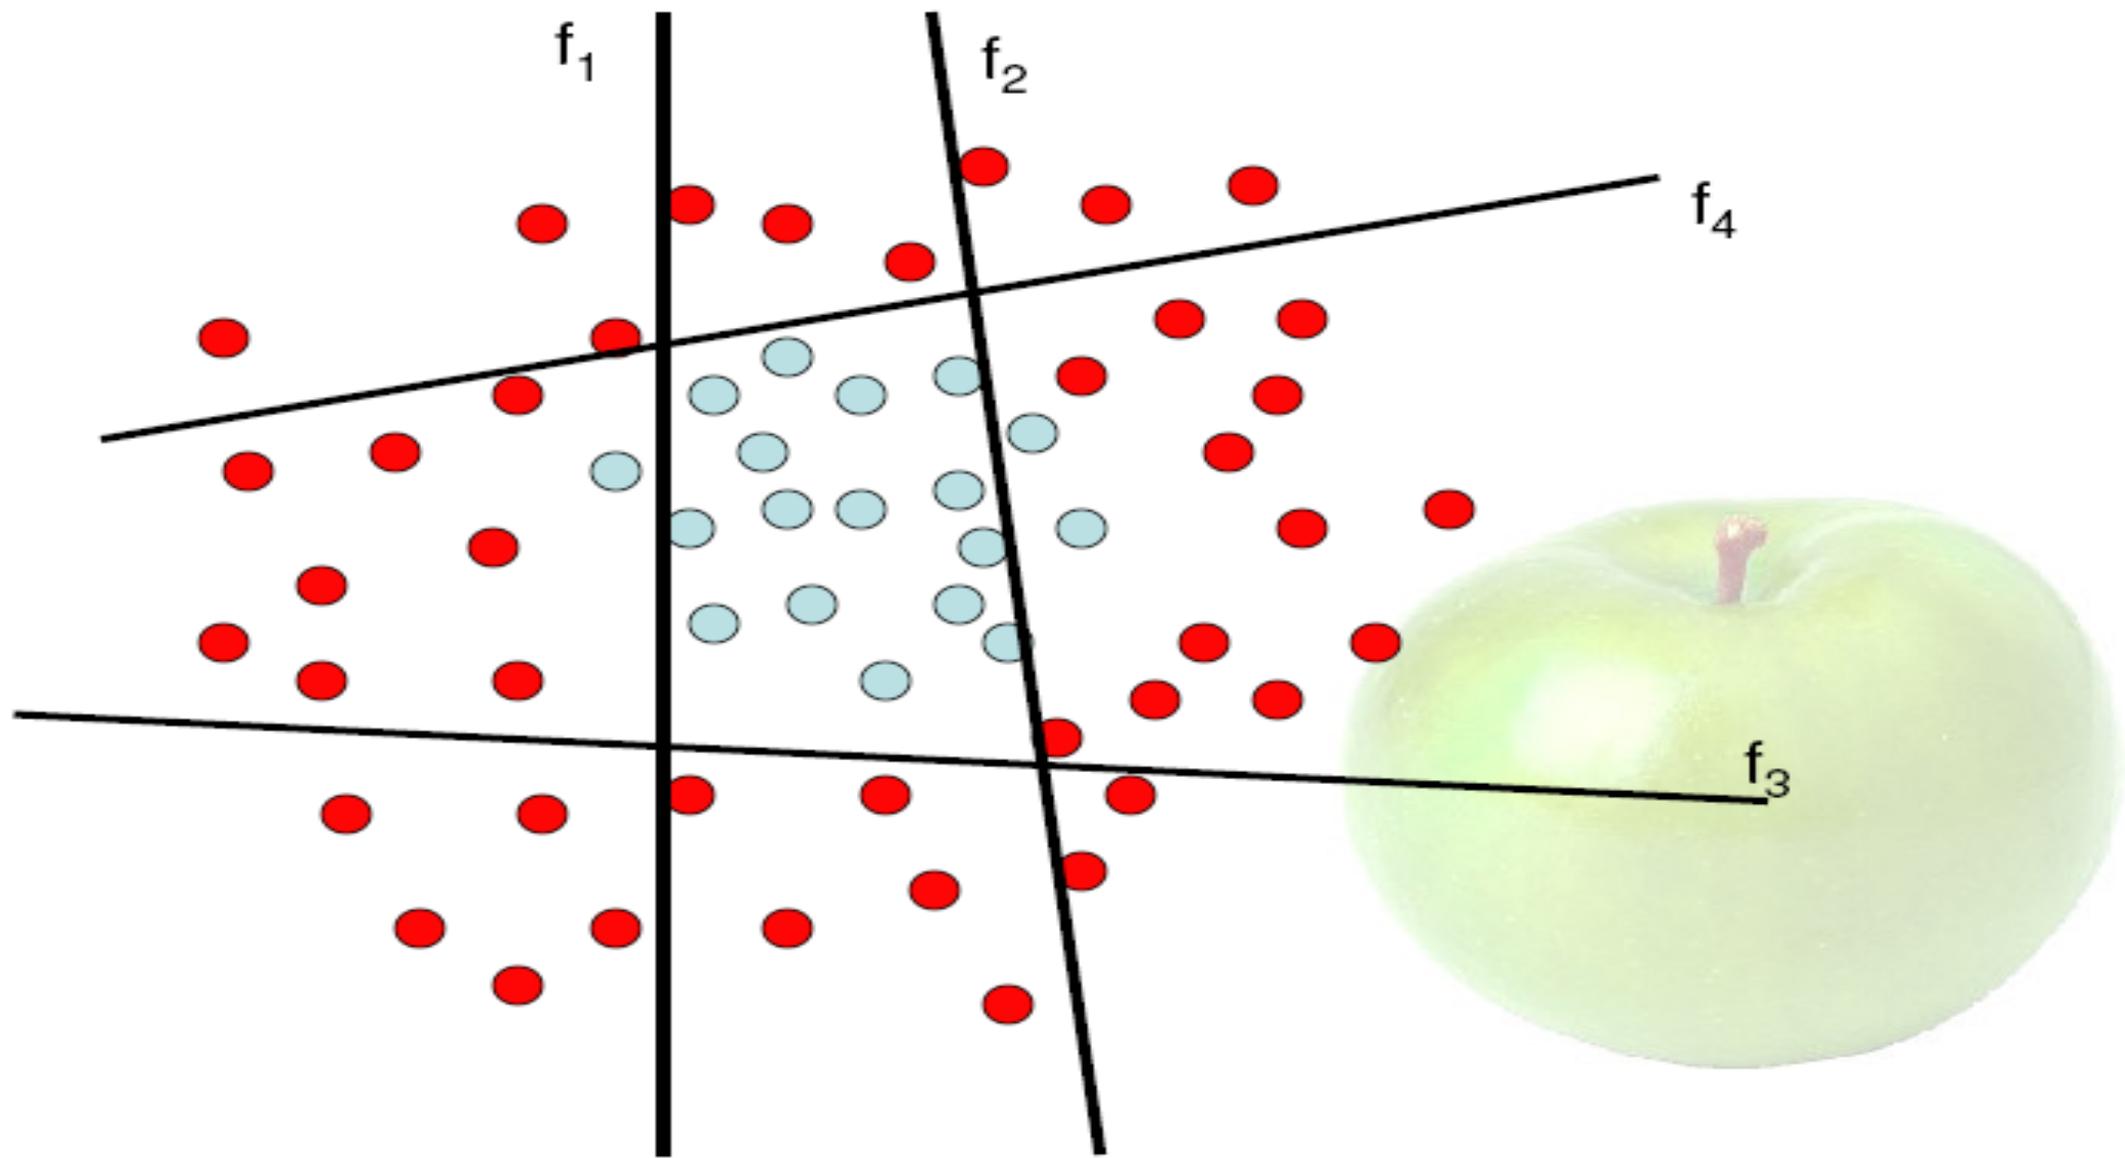
Each data point has
a class label:

$$y_t = \begin{cases} +1 (\bullet) \\ -1 (\circ) \end{cases}$$

We update the weights:

$$w_t \leftarrow w_t \exp\{-y_t H_t\}$$

We set a new problem for which the previous weak classifier performs at chance again



The strong (non-linear) classifier is built as the combination of all the weak (linear) classifiers.

JOINT BOOST - MULTICLASS

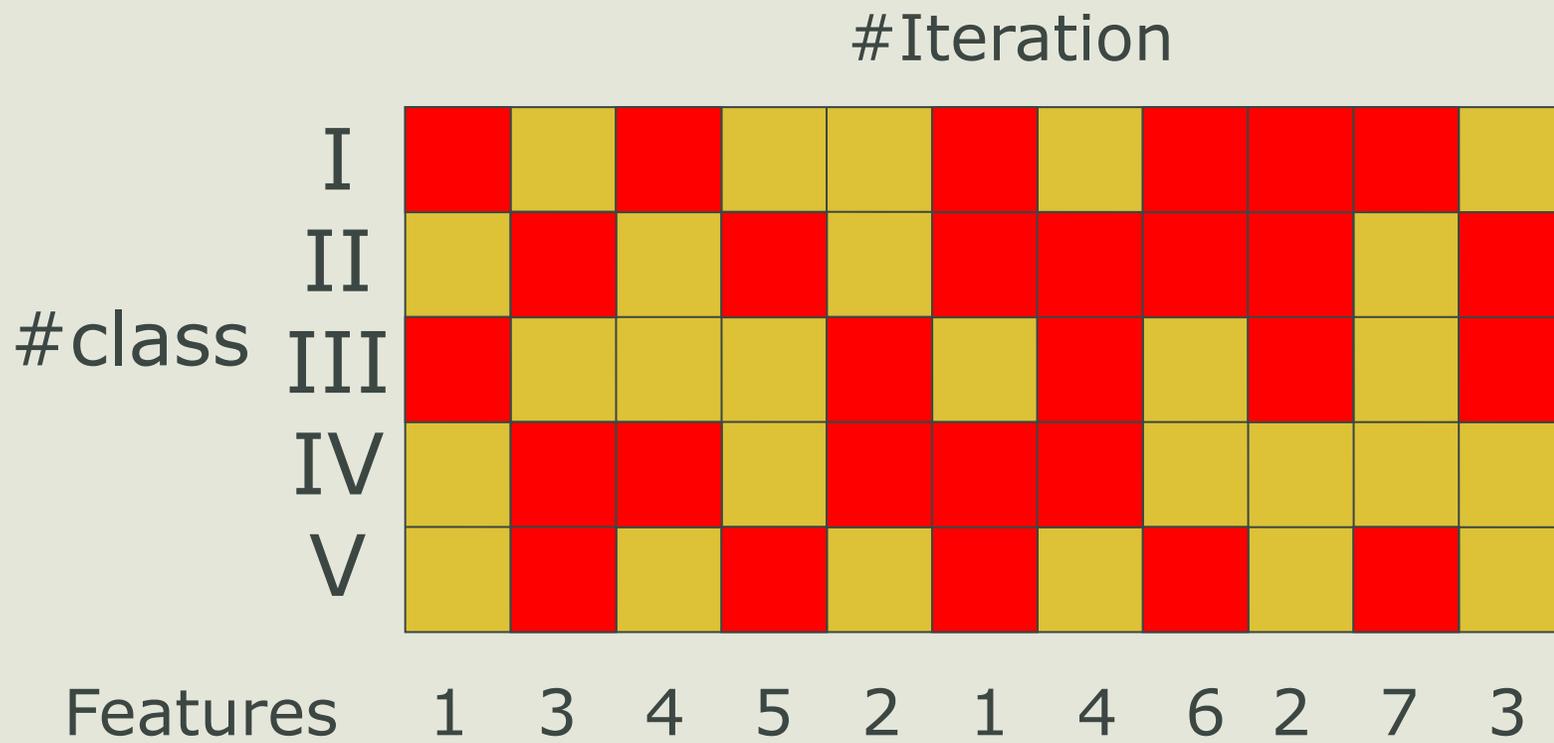
- We can minimize a similar function using one-vs-all strategy

$$H(v, c) = \sum_{m=1}^M h_m(v, c) \quad J = \sum_{c=1}^C E \left[e^{-z^c H(v, c)} \right]$$

- This doesn't work very well, since it is separable in c .
- Put constraints. -> **shared** features!

JOIN BOOST

- ❖ In each iteration, choose
 - ❖ One common feature
 - ❖ A subset of classes that use this feature
- ❖ So that the objective decreases most

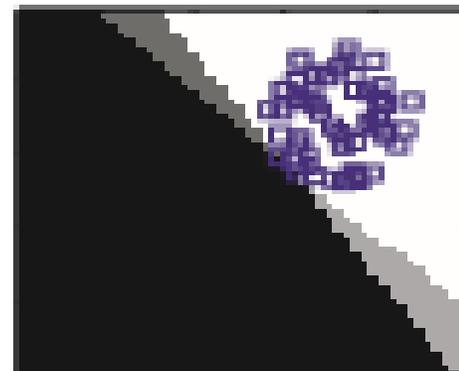
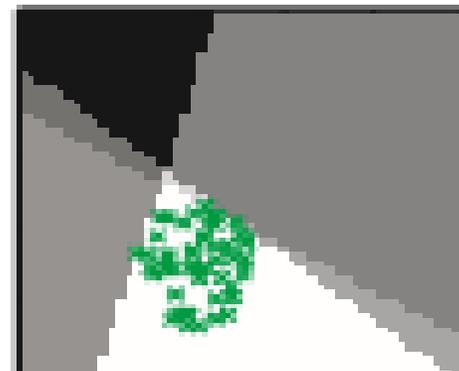
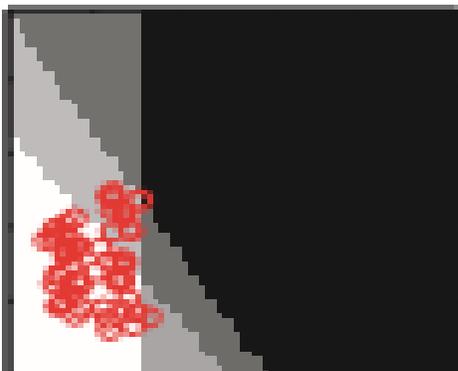
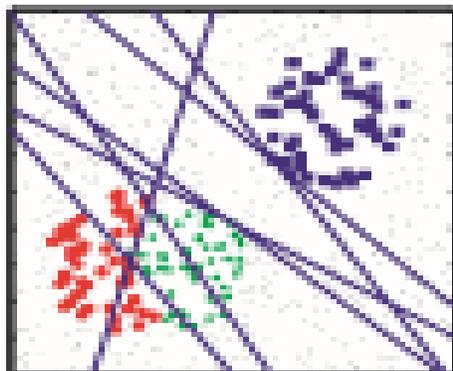
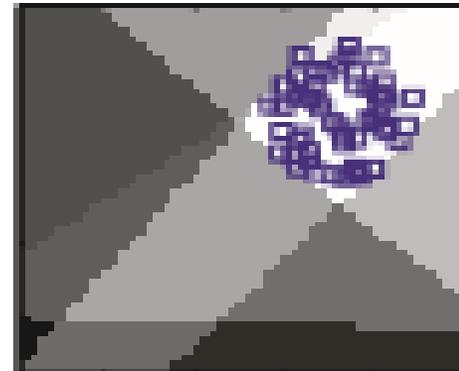
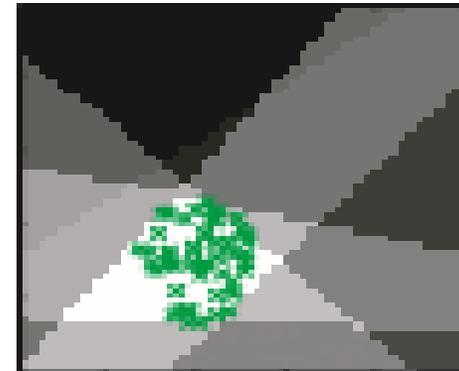
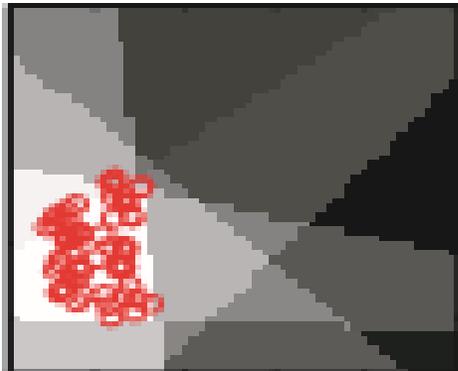
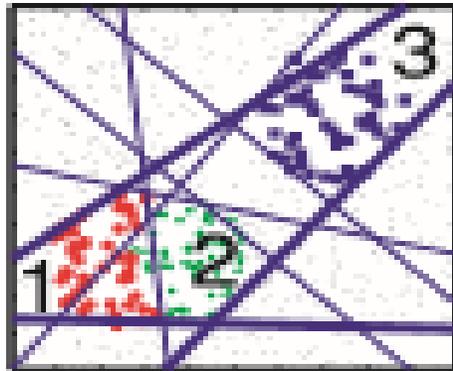


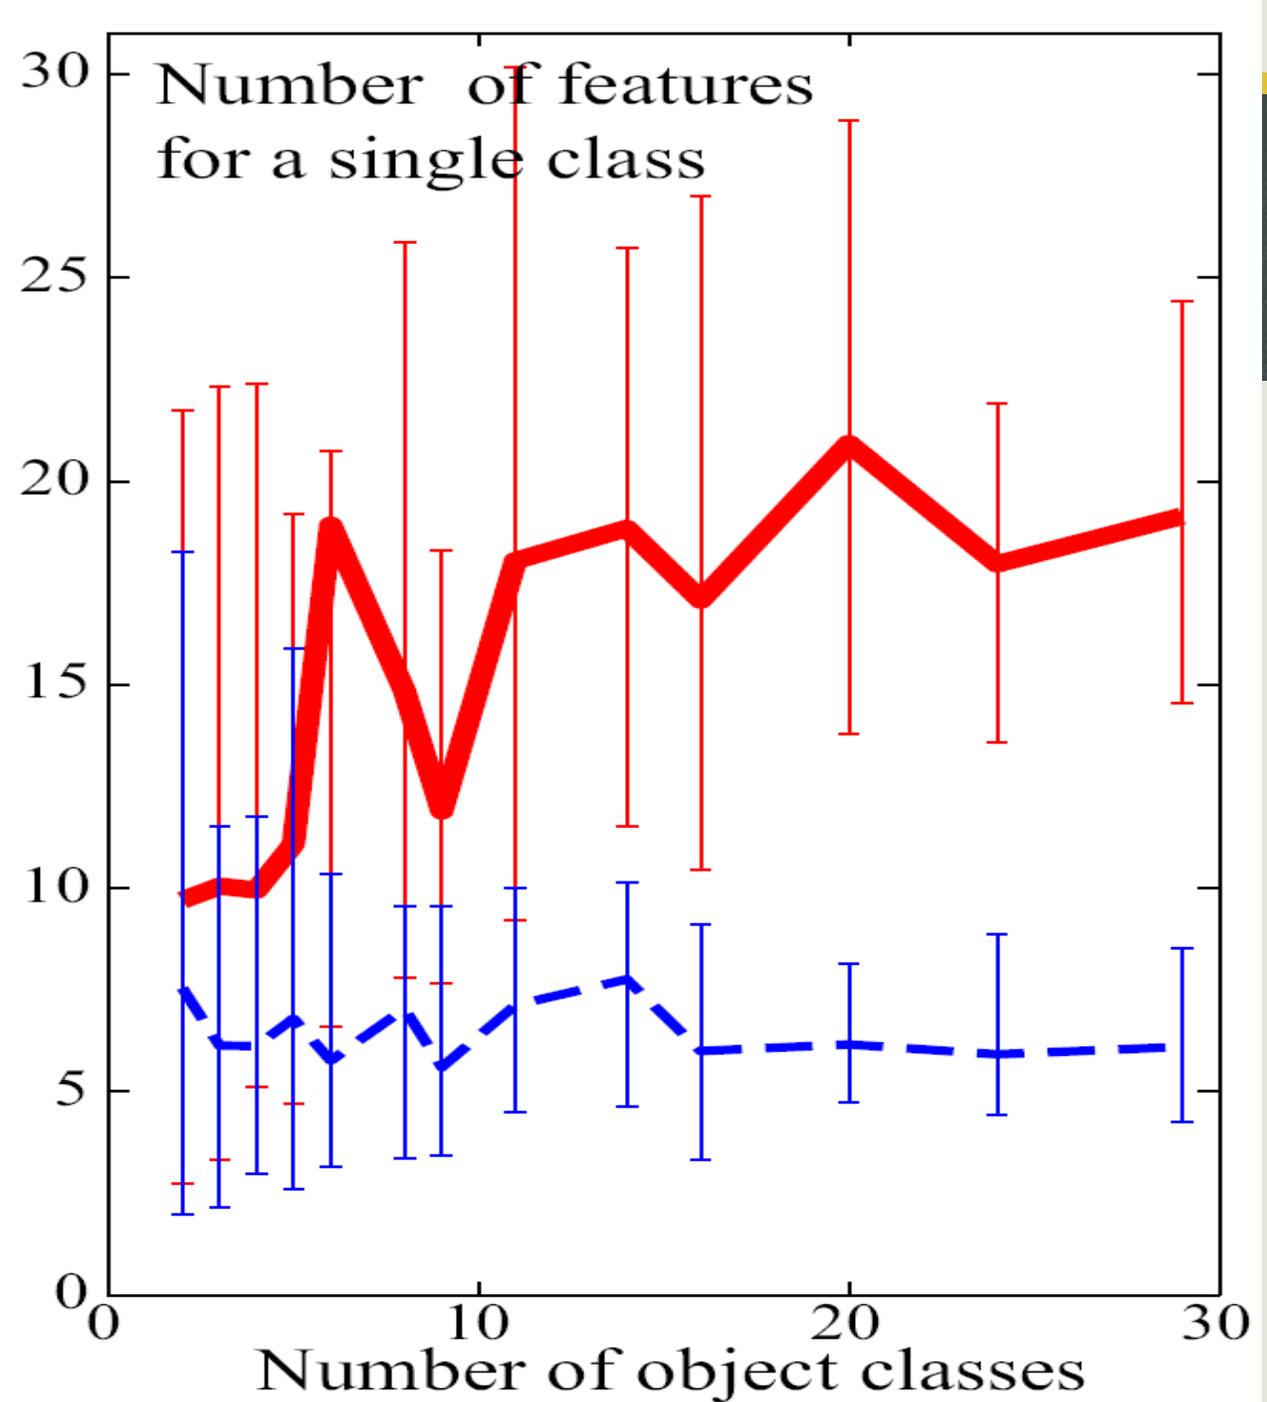
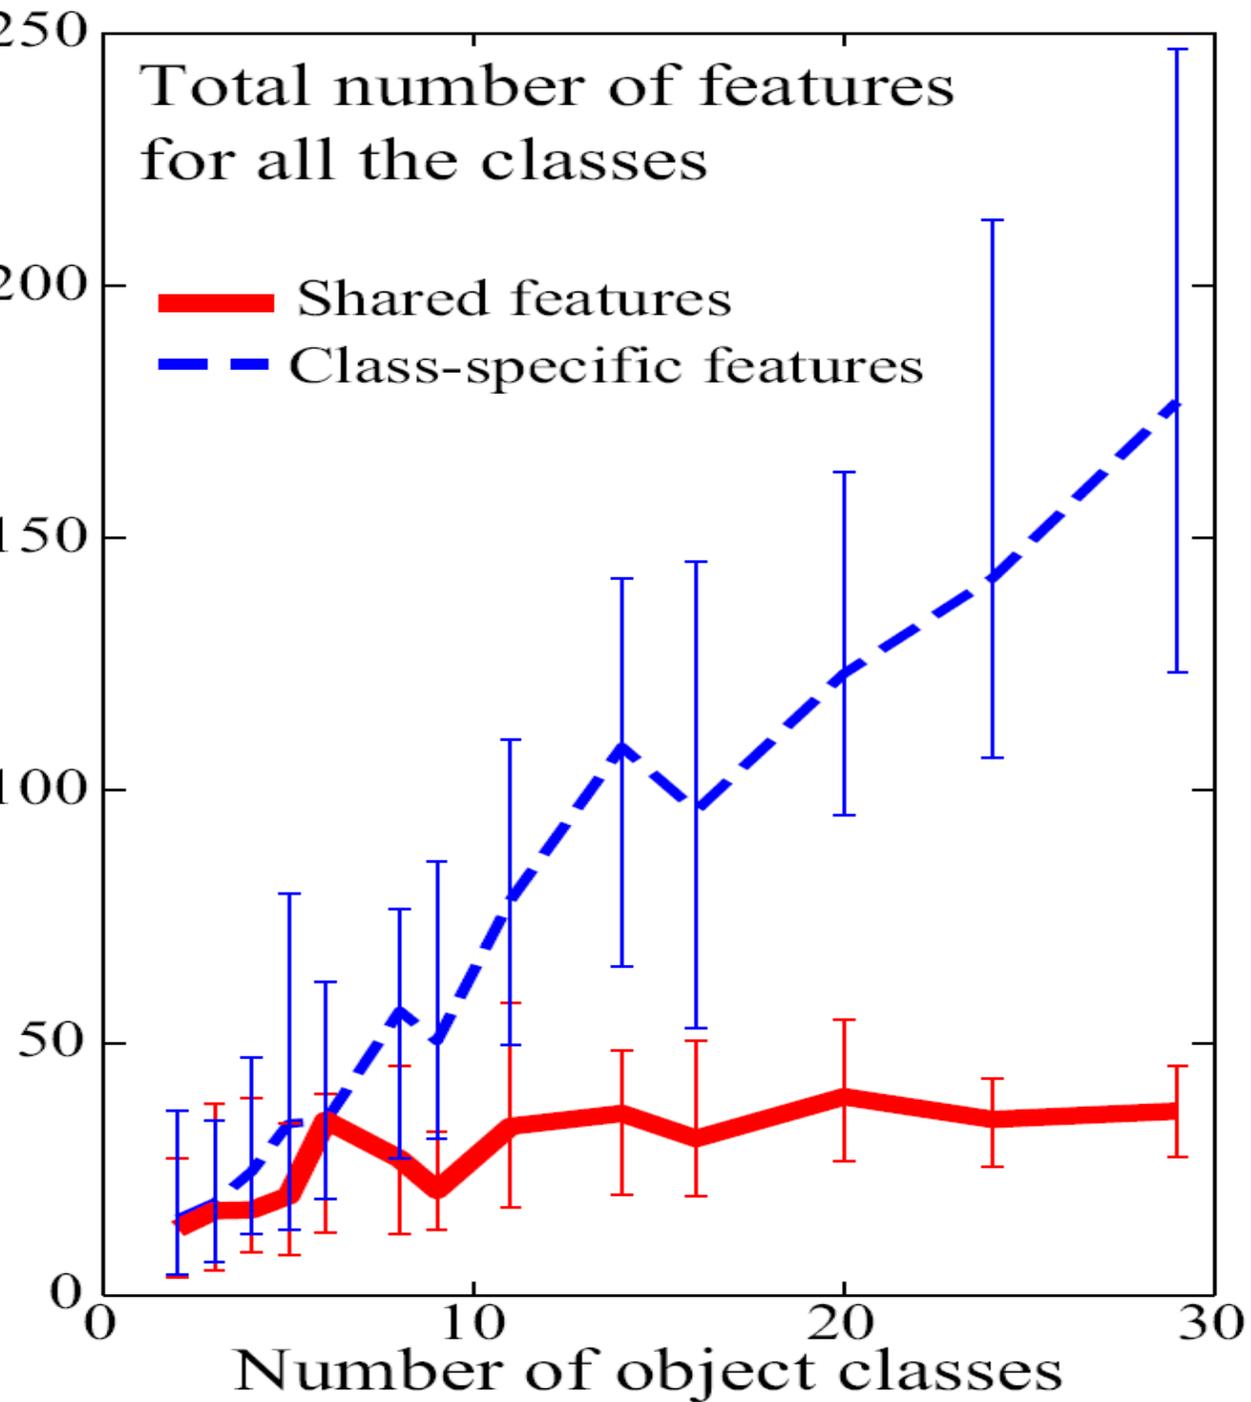
- ❖ Each class may have its own favorite feature
- ❖ a common feature may not be any of them, however it **simultaneously** decreases errors of many classes.

COMPUTATIONAL ISSUE

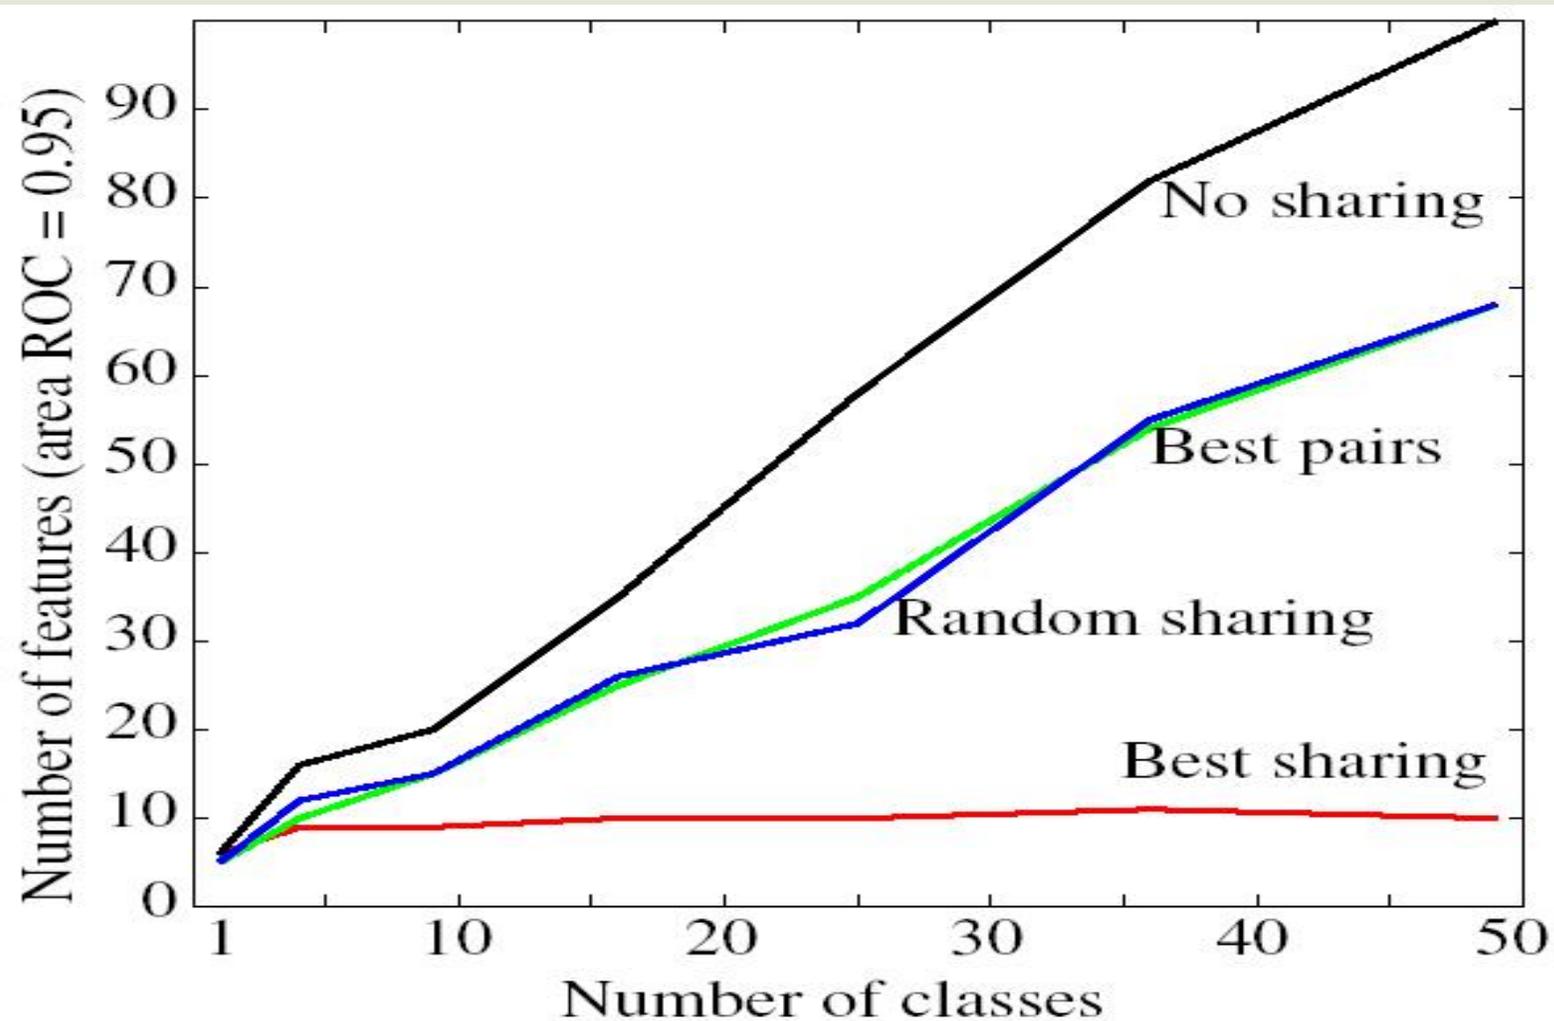
- ❖ Choose the best subset is prohibitive
- ❖ Use greedy approach
 - ❖ Choose one class and one feature so that the objective decreases the most
- ❖ Iteratively add more classes until the objective increases again
 - ❖ Note the common feature may change

JOINT BOOSTING VS. INDEPENDENT





#FEATURES = O(LOG #CLASS)

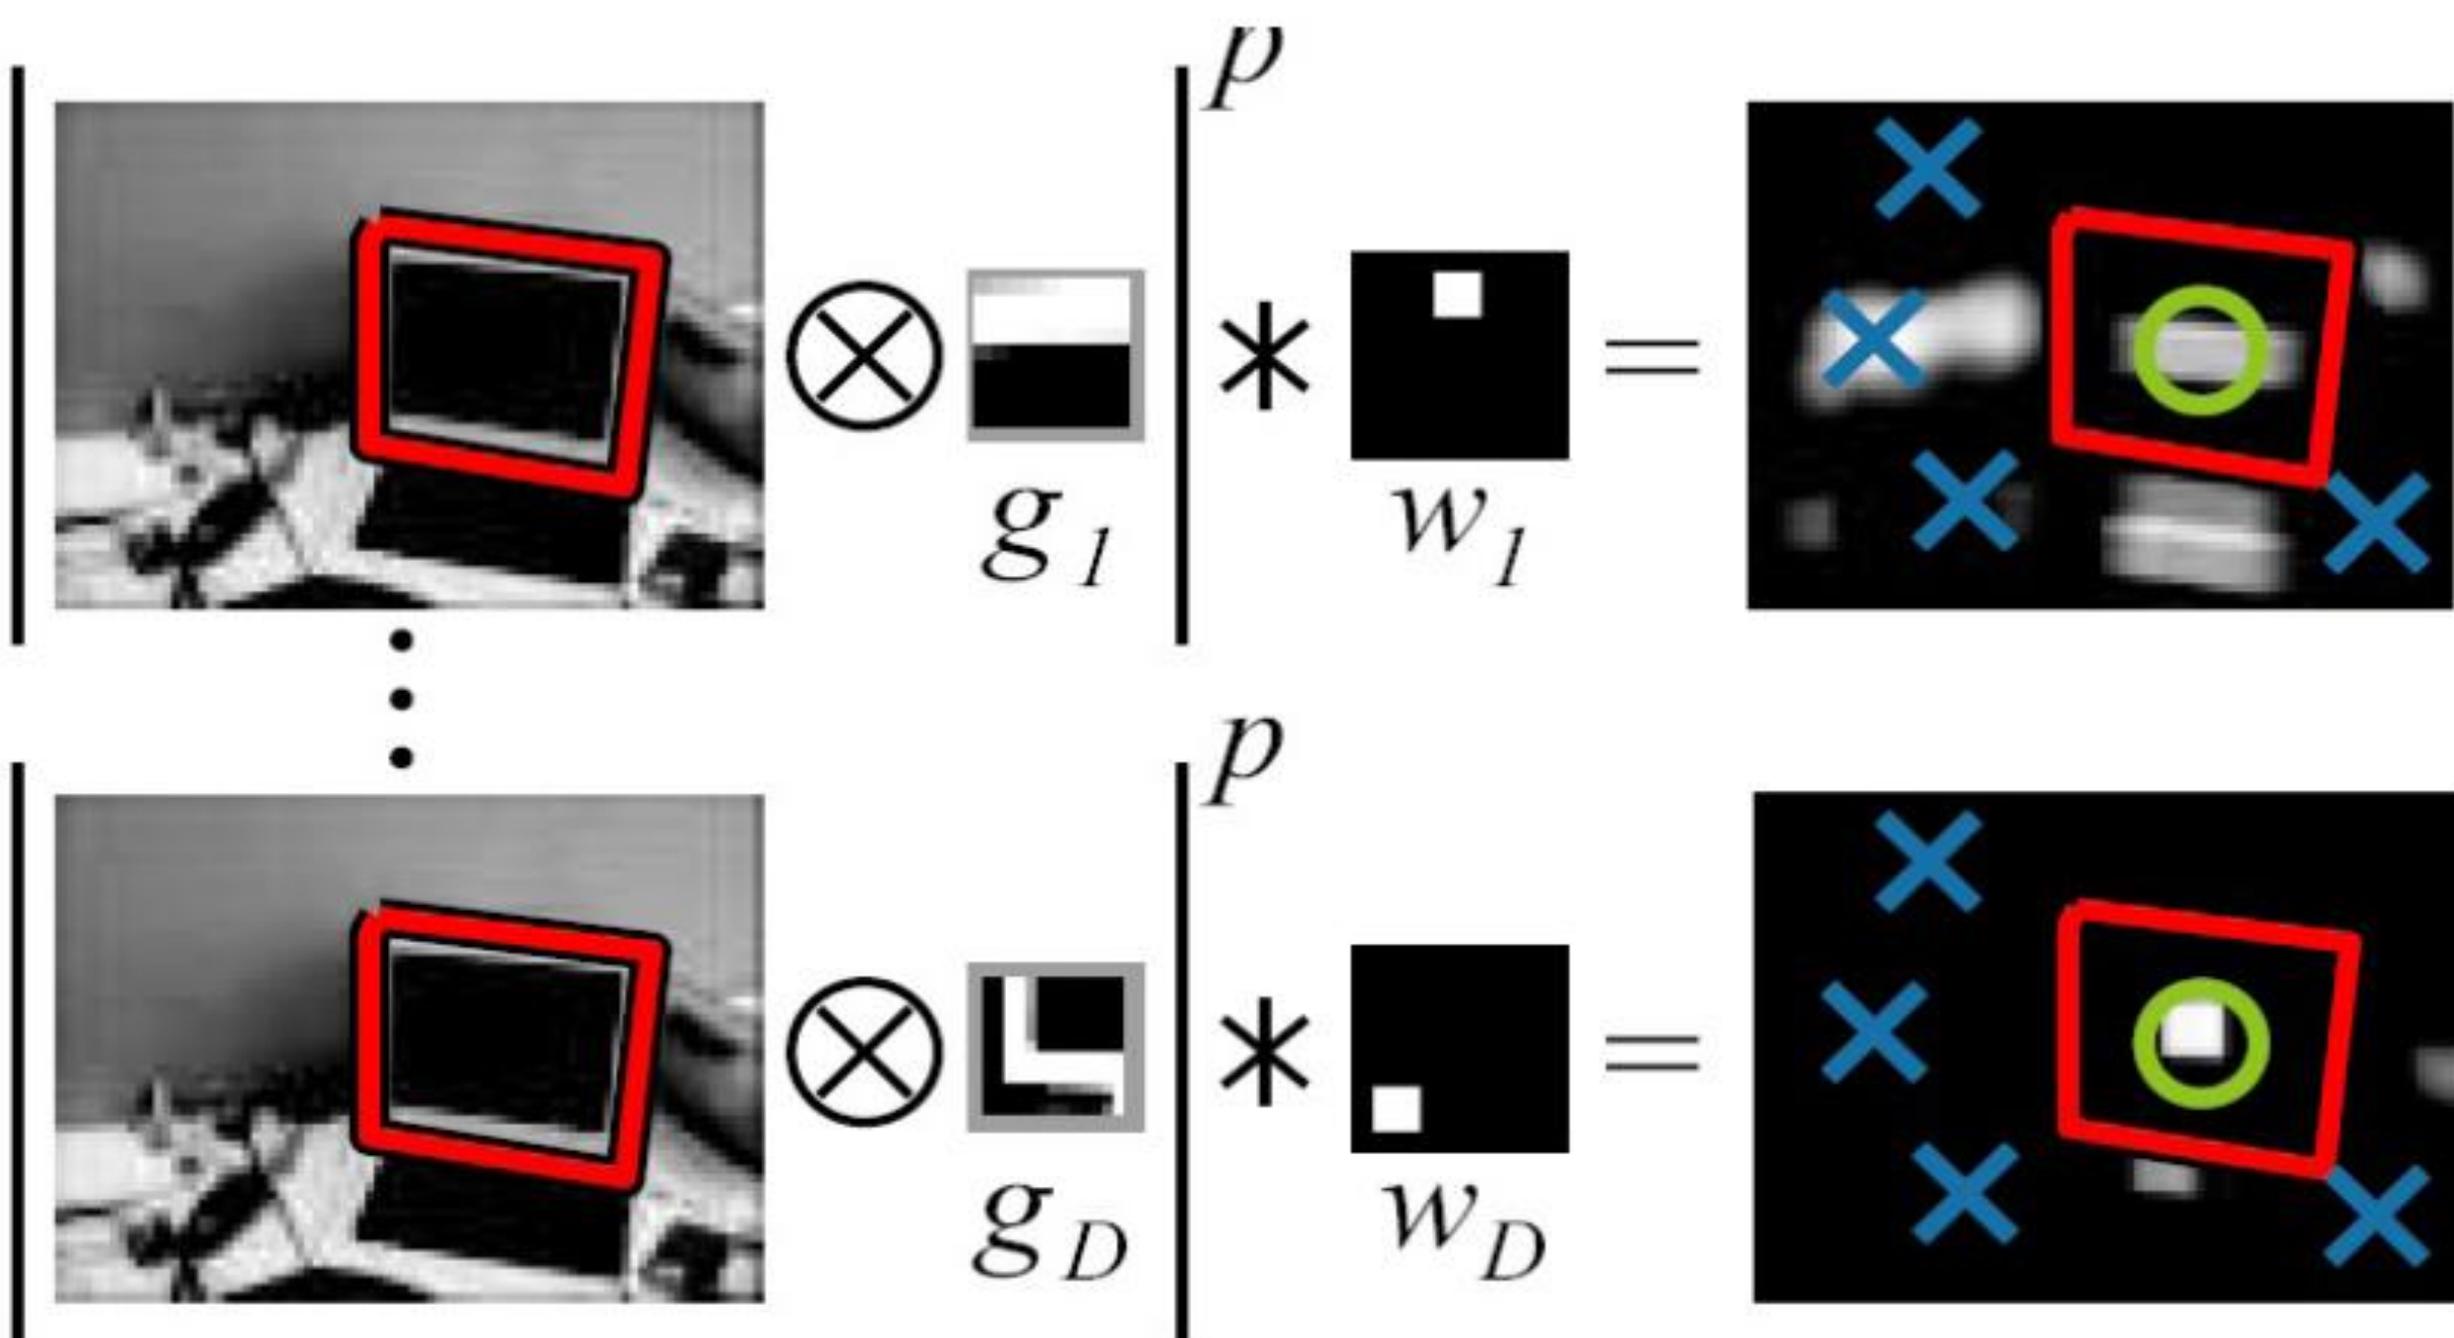


FEATURE THEY USED IN THE PAPER

- ❖ Dictionary
 - ❖ 2000 random sampled patches
 - ❖ Of size from 4x4 to 14x14
 - ❖ no clustering
 - ❖ Each patch is associated with a spatial mask

FEATURES

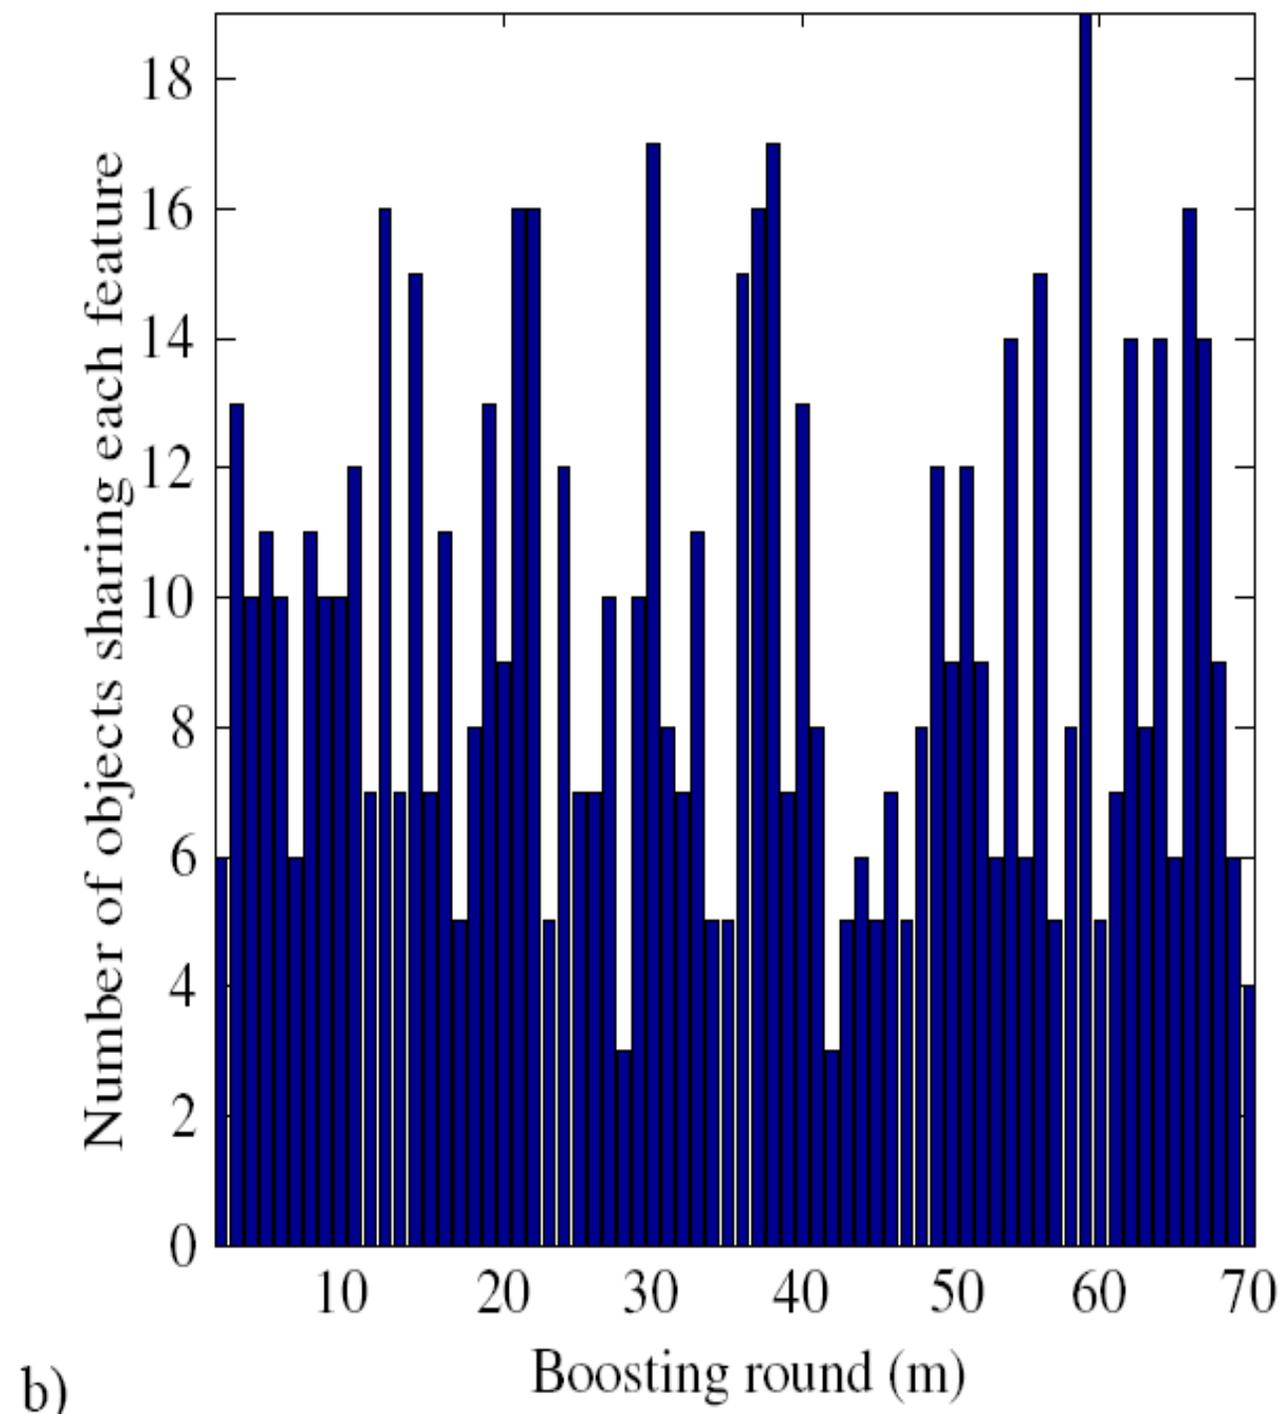
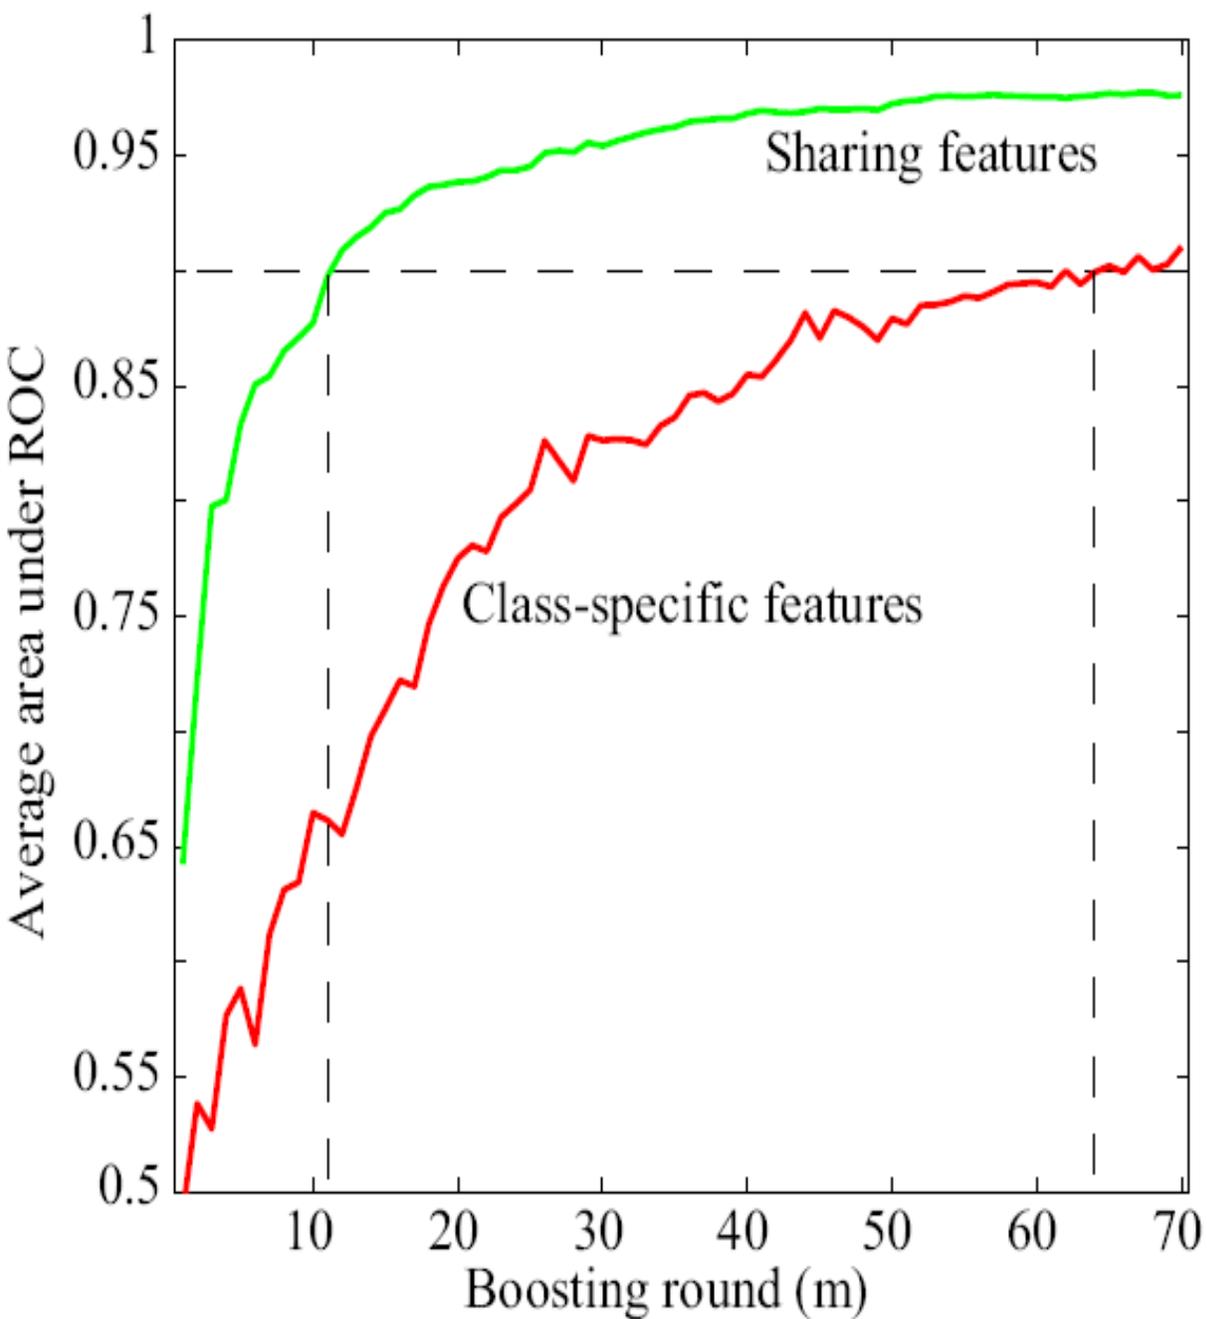
- ❖ Building feature vectors
 - ❖ Normalized correlation with each patch to get response
 - ❖ Raise the response to some power
 - ❖ Large value gets even larger and dominate the response (max operation)
 - ❖ Use spatial mask to align the response to the object center (voting)
 - ❖ Extract response vector at object center



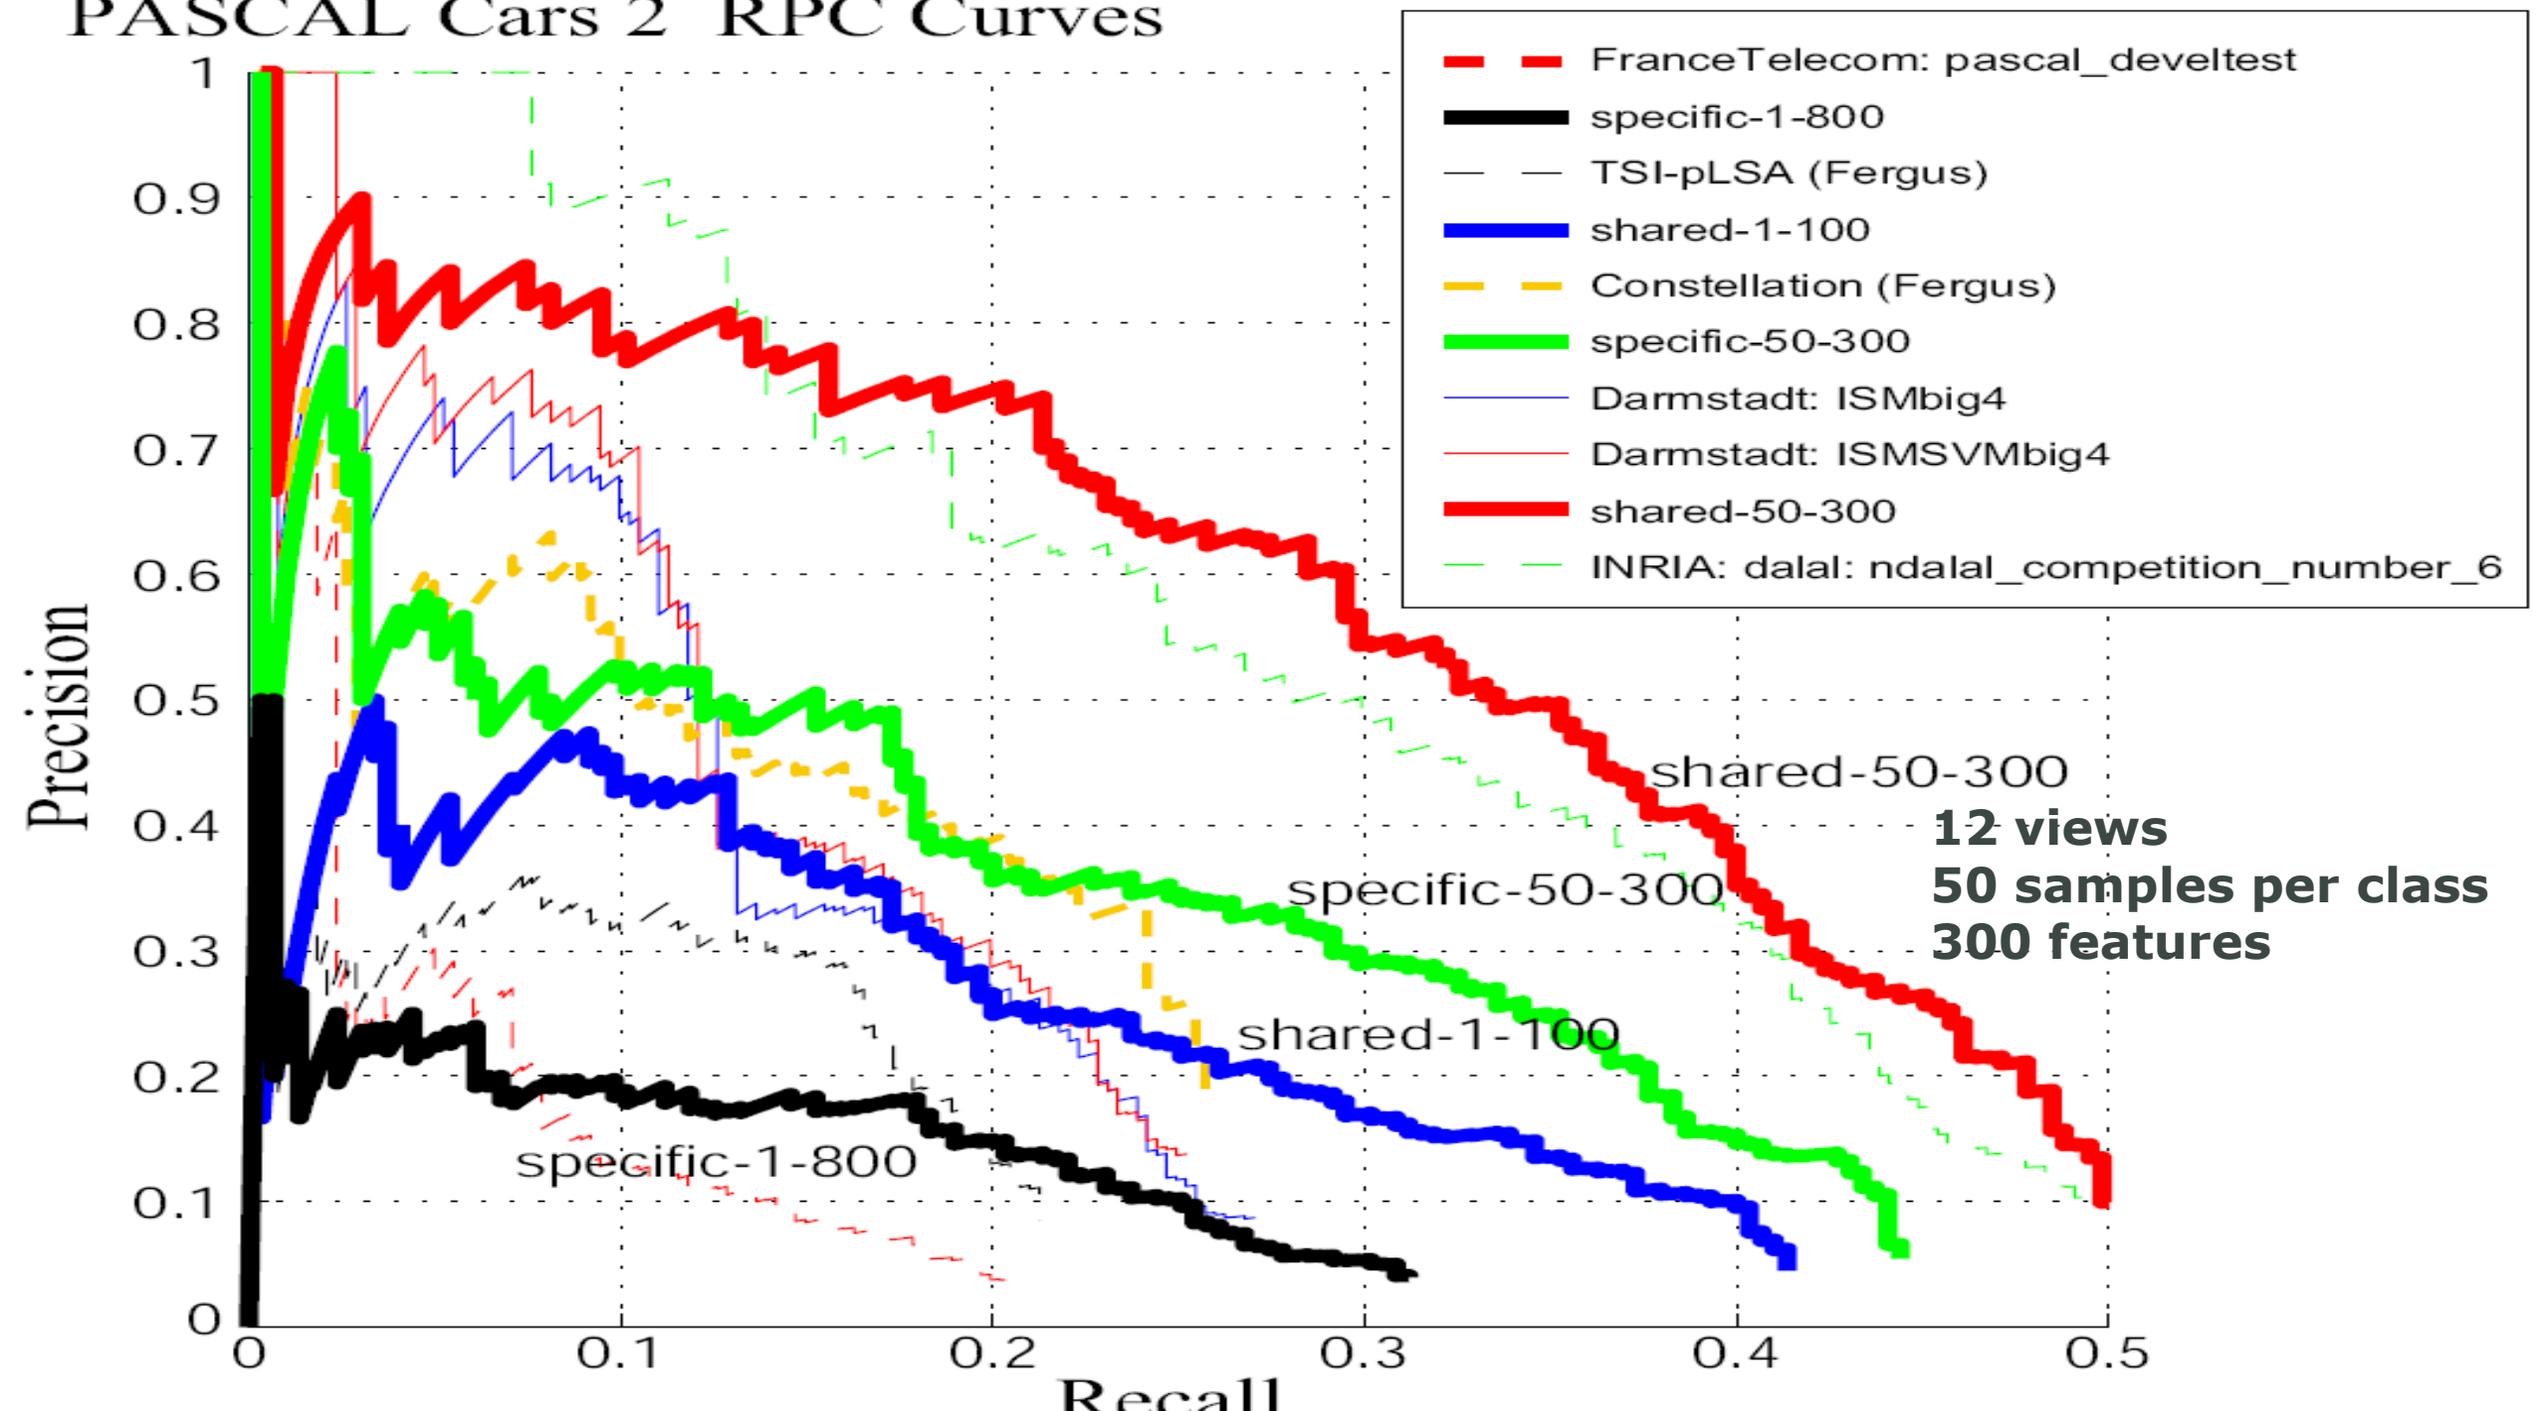
THE DATA

- Multiclass object recognition
 - Dataset: LabelMe
 - 21 objects, 50 samples per object
 - 500 rounds
- Multiview car recognition
 - Train on LabelMe, test on PASCAL
 - 12 views, 50 samples per view
 - 300 rounds





PASCAL Cars 2 RPC Curves



CONCLUSION

- ❖ Joint Boosting indeed works
 - ❖ Especially when the number of images per class is not too small (otherwise NN)
- ❖ Better performance in the presence of
 - ❖ Many classes, each class has only a few samples
 - ❖ Introduce regularization that reduce over fitting
- ❖ Disadvantages
 - ❖ Train slowly, $O(C^2)$.



THE END