

Classifiers!!!

BCH339N Systems Biology / Bioinformatics
Edward Marcotte, Univ of Texas at Austin

Clustering = task of grouping a set of objects in such a way that objects in the same group (a **cluster**) are more similar (in some sense) to each other than to those in other groups (clusters).


VS.

Classification = task of categorizing a new observation, on the basis of a training set of data with observations (or instances) whose categories are known


Adapted from Wikipedia

This is useful because biological systems tend to be modular and often inherited intact across evolution.

(e.g. you tend to have a flagellum or not)



Many such features are possible...



For yeast, $N \sim 6,000$
For human, $N \sim 22,000$

*i.e., a matrix of N
 $\times M$ numbers*

We also needed a measure of the similarity between feature vectors. Here are a few (of many) common distance measures used in clustering.

Names	Formula
Euclidean distance	$\ a - b\ _2 = \sqrt{\sum_i (a_i - b_i)^2}$
Manhattan distance	$\ a - b\ _1 = \sum_i a_i - b_i $
cosine similarity	$\frac{a \cdot b}{\ a\ \ b\ }$

Wikipedia

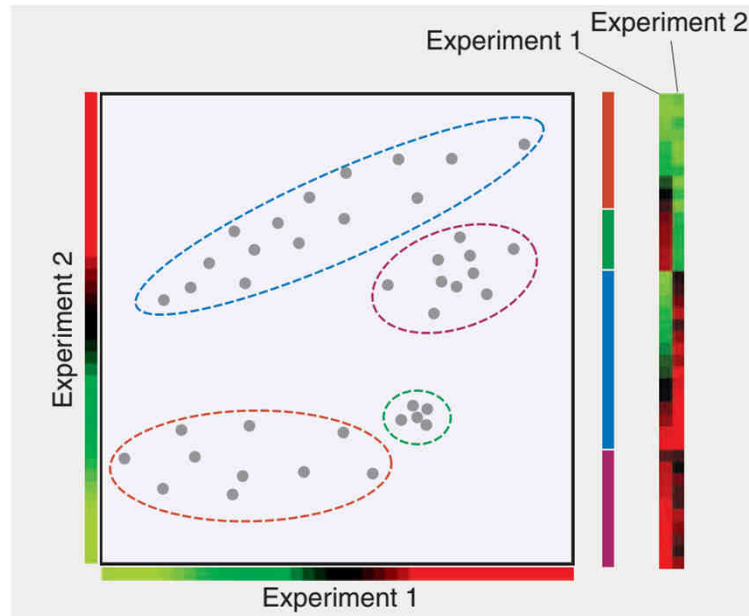
We also needed a measure of the similarity between feature vectors. Here are a few (of many) common distance measures used in clustering.

~~clustering~~
classifying

Names	Formula
Euclidean distance	$\ a - b\ _2 = \sqrt{\sum_i (a_i - b_i)^2}$
Manhattan distance	$\ a - b\ _1 = \sum_i a_i - b_i $
cosine similarity	$\frac{a \cdot b}{\ a\ \ b\ }$

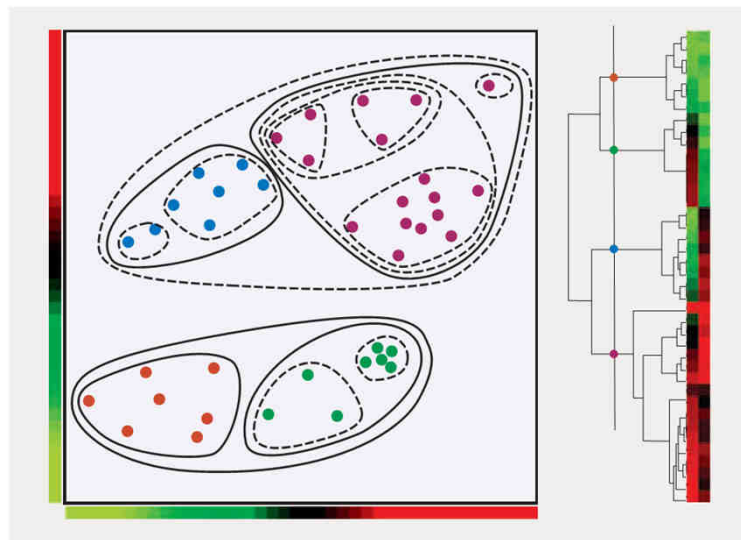
Wikipedia

Clustering refresher: 2-D example




Nature Biotech 23(12):1499-1501 (2005)

Clustering refresher: hierarchical




Nature Biotech 23(12):1499-1501 (2005)

Clustering refresher: SOM



Nature Biotech 23(12):1499-1501 (2005)


Clustering refresher: *k*-means



Nature Biotech 23(12):1499-1501 (2005)

Clustering refresher: k -means

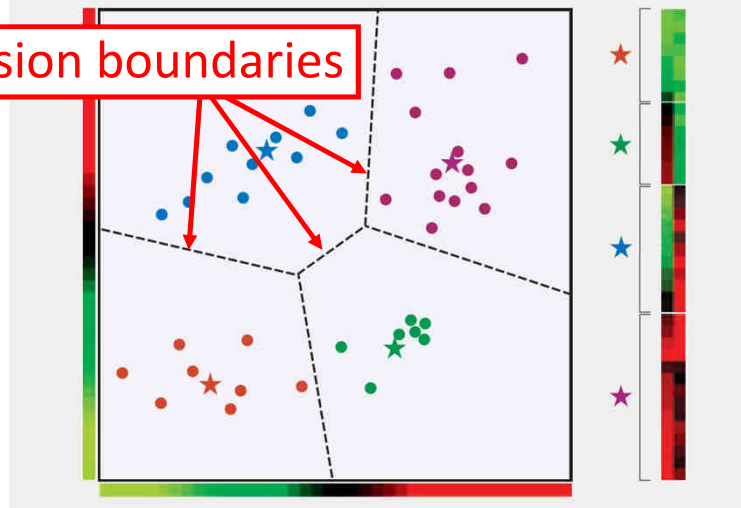
Decision boundaries



Nature Biotech 23(12):1499-1501 (2005)


One of the simplest classifiers uses the same notion of decision boundaries.

Decision boundaries



Nature Biotech 23(12):1499-1501 (2005)

One of the simplest classifiers uses this notion of decision boundaries.




Rather than first clustering, calculate the centroid (mean) of objects with each label.

New observations are classified as belonging to the group whose mean is nearest.

=“minimum distance classifier”

Nature Biotech 23(12):1499-1501 (2005)

One of the simplest classifiers uses this notion of decision boundaries.



For example....

Nature Biotech 23(12):1499-1501 (2005)

**Molecular Classification of
Cancer: Class Discovery and
Class Prediction by Gene
Expression Monitoring**

T. R. Golub,^{1,2*} D. K. Slonim,^{1†} P. Tamayo,¹ C. Huard,¹
M. Gaasenbeek,¹ J. P. Mesirov,¹ H. Coller,¹ M. L. Loh,²
J. R. Downing,³ M. A. Caligiuri,⁴ C. D. Bloomfield,⁴
E. S. Lander^{1,5*}

Let's look at a specific
example:

“Enzyme-based histochemical analyses were introduced in the 1960s to demonstrate that **some leukemias were periodic acid-Schiff positive, whereas others were myeloperoxidase positive...**

This provided the first basis for classification of acute leukemias into those arising from lymphoid precursors (acute lymphoblastic leukemia, ALL), or from myeloid precursors (acute myeloid leukemia, AML).”

15 OCTOBER 1999 VOL 286 SCIENCE

**Molecular Classification of
Cancer: Class Discovery and
Class Prediction by Gene
Expression Monitoring**

T. R. Golub,^{1,2*} D. K. Slonim,^{1†} P. Tamayo,¹ C. Huard,¹
M. Gaasenbeek,¹ J. P. Mesirov,¹ H. Coller,¹ M. L. Loh,²
J. R. Downing,³ M. A. Caligiuri,⁴ C. D. Bloomfield,⁴
E. S. Lander^{1,5*}

Let's look at a specific
example:

“**Distinguishing ALL from AML is critical for successful treatment...**

chemotherapy regimens for ALL generally contain corticosteroids, vincristine, methotrexate, and L-asparaginase, whereas

most AML regimens rely on a backbone of daunorubicin and cytarabine (8).


Although remissions can be achieved using ALL therapy for AML (and vice versa), cure rates are markedly diminished, and unwarranted toxicities are encountered.”

15 OCTOBER 1999 VOL 286 SCIENCE

Molecular Classification of Cancer: Class Discovery and Class Prediction by Gene Expression Monitoring

T. R. Golub,^{1,2*} D. K. Slonim,^{1†} P. Tamayo,¹ C. Huard,¹
 M. Gaasenbeek,¹ J. P. Mesirov,¹ H. Coller,¹ M. L. Loh,²
 J. R. Downing,³ M. A. Caligiuri,⁴ C. D. Bloomfield,⁴
 E. S. Lander^{1,5*}

Let's look at a specific example:




Take labeled samples, find genes whose abundances separate the samples...

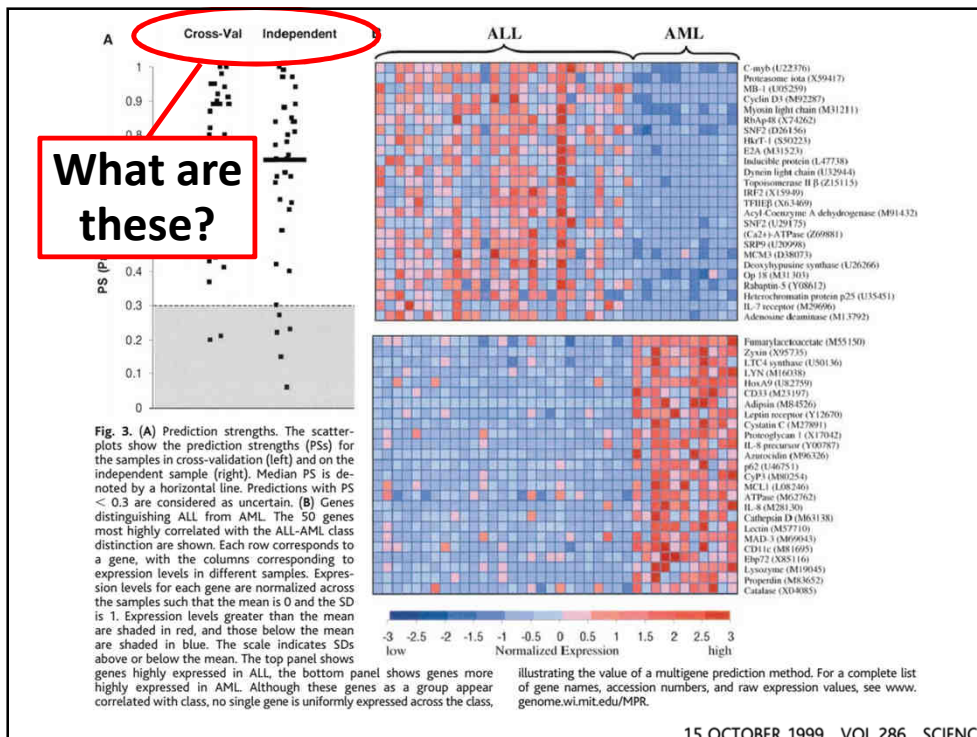
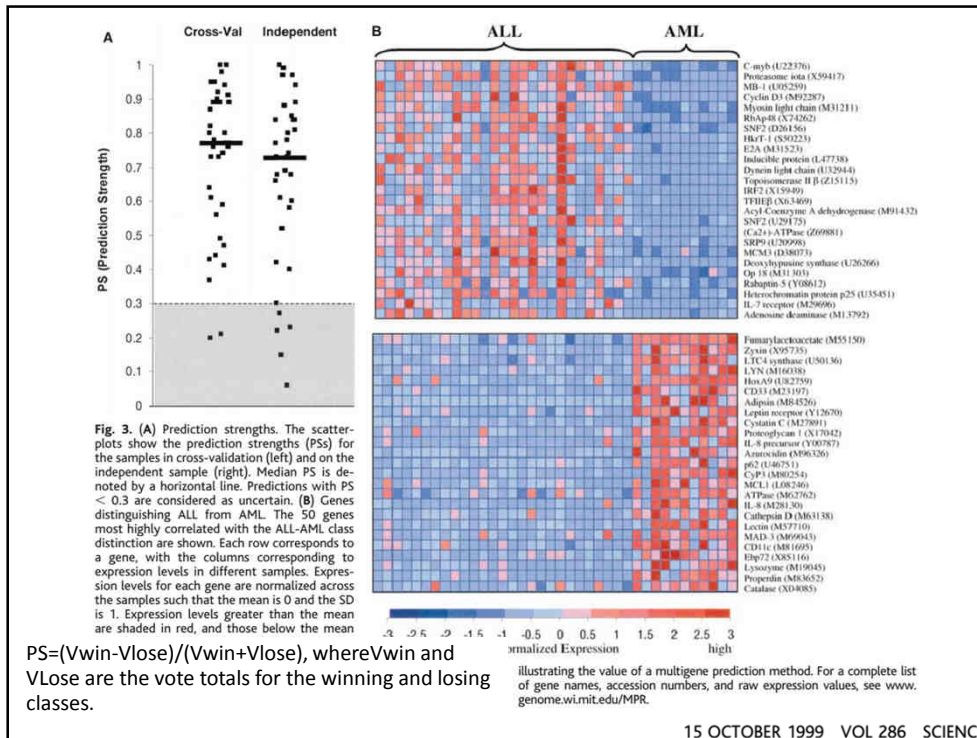
Molecular Classification of Cancer: Class Discovery and Class Prediction by Gene Expression Monitoring

T. R. Golub,^{1,2*} D. K. Slonim,^{1†} P. Tamayo,¹ C. Huard,¹
 M. Gaasenbeek,¹ J. P. Mesirov,¹ H. Coller,¹ M. L. Loh,²
 J. R. Downing,³ M. A. Caligiuri,⁴ C. D. Bloomfield,⁴
 E. S. Lander^{1,5*}

Let's look at a specific example:



Calculate weighted average of indicator genes to assign class of an unknown



Cross-validation

Withhold a sample, build a predictor based only on the remaining samples, and predict the class of the withheld sample.

Repeat this process for each sample, then calculate the cumulative or average error rate.

15 OCTOBER 1999 VOL 286 SCIENCE

X-fold cross-validation **e.g. 3-fold or 10-fold**

Can also withhold $1/X$ (e.g. $1/3$ or $1/10$) of sample, build a predictor based only on the remaining samples, and predict the class of the withheld samples.

Repeat this process X times for each withheld fraction of the sample, then calculate the cumulative or average error rate.

15 OCTOBER 1999 VOL 286 SCIENCE

Independent data

Withhold an entire dataset, build a predictor based only on the remaining samples (**the training data**).

Test the trained classifier on the independent **test data** to give a fully independent measure of performance.

15 OCTOBER 1999 VOL 286 SCIENCE

You already know how to measure how well these algorithms work (way back in our discussion of gene finding!)

True answer:

		Positive	Negative
Algorithm predicts:	Positive	True positive	False positive
	Negative	False negative	True negative

$$\text{Specificity} = \text{TN} / (\text{TN} + \text{FP})$$


$$\text{Sensitivity} = \text{TP} / (\text{TP} + \text{FN})$$

You already know how to measure how well these algorithms work (way back in our discussion of gene finding!)

Sort the data by their classifier score, then step from best to worst and plot the performance:

$$\text{Sensitivity} = \frac{TP}{TP + FN}$$

also called
True Positive
Rate (TPR)



First used in WWII to analyze radar signals (e.g., after attack on Pearl Harbor)

ROC curve
(receiver operator characteristic)


1 - Specificity =
FP / (FP + TN)
also called False Positive Rate (FPR)

Another good option:

Sort the data by their classifier score, then step from best to worst and plot the performance:

$$\text{Precision} = \frac{TP}{TP + FP}$$

also called
positive
predictive value
(PPV)




Precision-recall curve

Recall =
TP / (TP + FN)
(= sensitivity)


Back to our minimum distance classifier...

Would it work well for this data?



Back to our minimum distance classifier...

How about this data? What might?



Back to our minimum distance classifier...

How about this data? What might?

```
XXXXO O O O XXXXO O O O
XXXXO O O O XXXXO O O O
XXXXO O O O XXXXO O O O
XXXXO O O O XXXXO O O O
O O O XXXX O O O XXXX
O O O XXXX O O O XXXX
O O O XXXX O O O XXXX
O O O XXXX O O O XXXX
XXXXO O O O XXXXO O O O
XXXXO O O O XXXXO O O O
XXXXO O O O XXXXO O O O
XXXXO O O O XXXXO O O O
O O O XXXX O O O XXXX
O O O XXXX O O O XXXX
O O O XXXX O O O XXXX
O O O XXXX O O O XXXX
```

This is a great case for something called a **k-nearest neighbors classifier**:

For each new object, calculate the *k* closest data points. Let them vote on the label of the new object.

```
XXXXO O O O XXXXO O O O
XXXXO O O O XXXXO O O O
XXXXO O O O XXXXO O O O
XXXXO O O O XXXXO O O O
O O O XXXX O O O XXXX
O O O XXXX O O O XXXX
O O O XXXX O O O XXXX
O O O XXXX O O O XXXX
XXXXO O O O XXXXO O O O
XXXXO O O O XXXXO O O O
XXXXO O O O XXXXO O O O
XXXXO O O O XXXXO O O O
O O O XXXX O O O XXXX
O O O X X X O O O XXXX
O O O X X X O O O XXXX
O O O XXXX O O O XXXX
O O O XXXX O O O XXXX
```

This is surrounded by O's and will probably be voted to be an O.

This one is surrounded by X's and will probably be voted to be an X.

Back to leukemias.
There was a follow-up study in 2010:

Clinical Utility of Microarray-Based Gene Expression Profiling in the Diagnosis and Subclassification of Leukemia: Report From the International Microarray Innovations in Leukemia Study Group

Torsten Haferlach, Alexander Kohlmann, Lothar Wieczorek, Giuseppe Basso, Geertruy Te Kronnie, Marie-Christine Béné, John De Vos, Jesus M. Hernández, Wolf-Karsten Hofmann, Ken I. Mills, Amanda Gilkes, Sabina Chiaretti, Sheila A. Shurtleff, Thomas J. Kipps, Laura Z. Rassenti, Allen E. Yeoh, Peter R. Papenhausen, Wei-min Liu, P. Mickey Williams, and Robin Foà

- **Tested clinical use of mRNA expression profiling to subtype leukemias into myeloid/lymphoid**
- **Meta-analysis of 11 labs, 3 continents, 3,334 patients**
- **Stage 1 (2,096 patients):**
92.2% classification accuracy for 18 leukemia classes (99.7% median specificity)
- **Stage 2 (1,152 patients):**
95.6% median sensitivity and 99.8% median specificity for 14 subtypes of acute leukemia
- **Microarrays outperformed routine diagnostics in 29 (57%) of 51 discrepant cases**

Conclusion: “Gene expression profiling is a robust technology for the diagnosis of hematologic malignancies with high accuracy”

J Clin Oncol 28:2529-2537. © 2010

In practice, if you want to explore classifiers, I strongly recommend the Weka package:

<http://www.cs.waikato.ac.nz/ml/weka/>




It's free, and easy to install, use, & troubleshoot. It lets you quickly test many alternative (well-vetted) classifiers, all in a proper cross-validated/precision-recall framework.

Here's a nice step-by-step intro for biologists :

Introducing Machine Learning Concepts with WEKA, in *Statistical Genomics, Methods in Molecular Biology*, v. 1418, p. 353-378, 24 March 2016

http://link.springer.com/content/pdf/10.1007%2F978-1-4939-3578-9_17.pdf

There's also a great book to walk you through the entire process.
Highly recommended!!!



In practice, if you want to explore classifiers, I also strongly recommend always testing these classifiers:

Random forests
Support vector machines (SVM)

These two are surprisingly often the best for many biological classification problems. Weka can do both of them.