

# 15-494/694: Cognitive Robotics

Dave Touretzky

Lecture 15:

Machine learning with  
scikit-learn

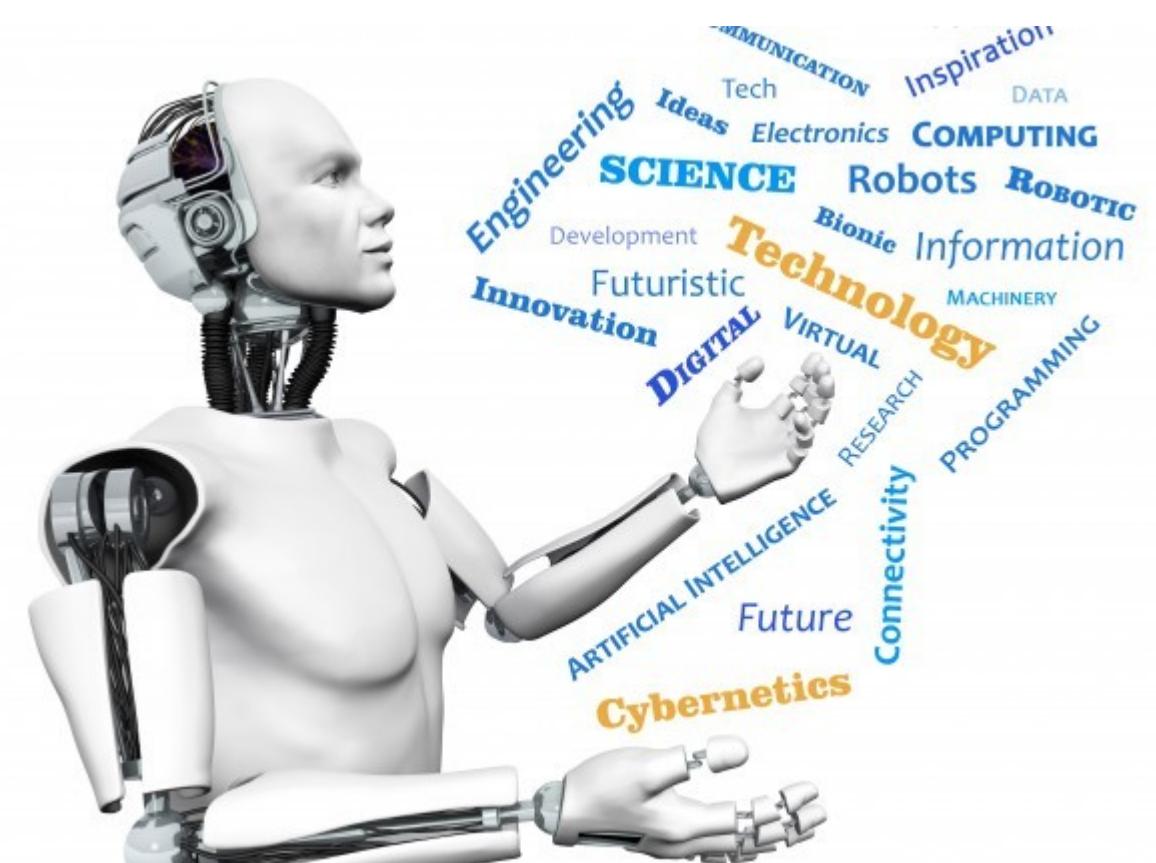


Image from <http://www.futuristgerd.com/2015/09/10>

# Machine Learning

- ML is a branch of Artificial Intelligence.
- “Learning” does not mean human-like learning.
- It means extracting information from data. This has many uses in robotics.
- Types of learning algorithm:
  - Supervised (labeled data)
  - Unsupervised (unlabeled data)
  - Reinforcement

# Supervised Learning

- For each training point, there is a desired output value.
- Error measure: difference between actual output and desired output.
  - Example: sum-squared error

$$E = \sum (d_i - y_i)^2$$

- Learning adjusts the model parameters to reduce the error.

# Unsupervised Learning

- Data points are unlabeled: there is no “correct” answer.
- Learning discovers structure in the data.
- Examples:
  - Clustering: finding categories.
  - Dimensionality reduction: finding key features and relationships between features. Useful for data compression.

# Reinforcement Learning

- Model is trained via a reinforcement signal that tells it how well it is doing.
- We don't tell it the right answer, just reward it when it does well.
- Example:
  - Learning to play a game by reinforcing wins. Program can learn by playing against itself.

# Supervised Learning: Classification

- Desired outputs are binary.
- Examples:
  - Tell “spam” from “not spam”.
  - Distinguish images containing cats from those without cats.
  - Recognize handwritten digits 0-9.

# Supervised Learning: Regression

- Desired outputs are continuous, possibly vectors.
- Examples:
  - Predict stock prices.
  - Calculate inverse kinematics solutions for a non-linear robot.

# Parametric vs. Non-Parametric Models

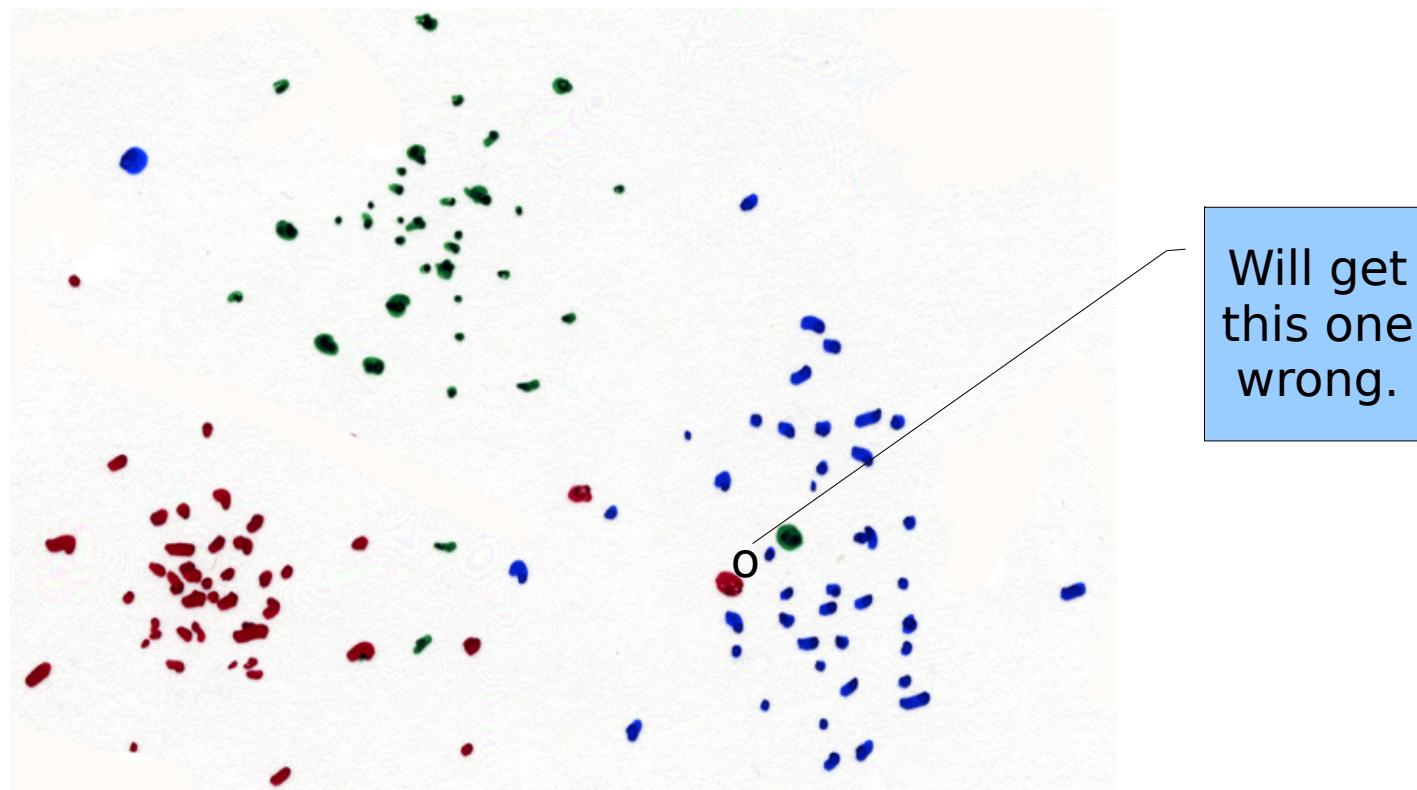
- Parametric models describe data using equations with a small number of parameters.
  - Example: Gaussian distribution.
  - Parameters are mean  $\mu$  and variance  $\sigma$
- Pros: compact representation; easy to test new data points.
- Cons: what if your data doesn't fit the equation?

# Parametric vs. Non-Parametric Models

- Non-parametric models don't make any assumptions about the distribution of the data. The data represents itself.
- Pros: “training” is instantaneous. Can represent arbitrary distributions.
- Cons: can take a lot of memory to store all the data, and classifying new points can be slow.

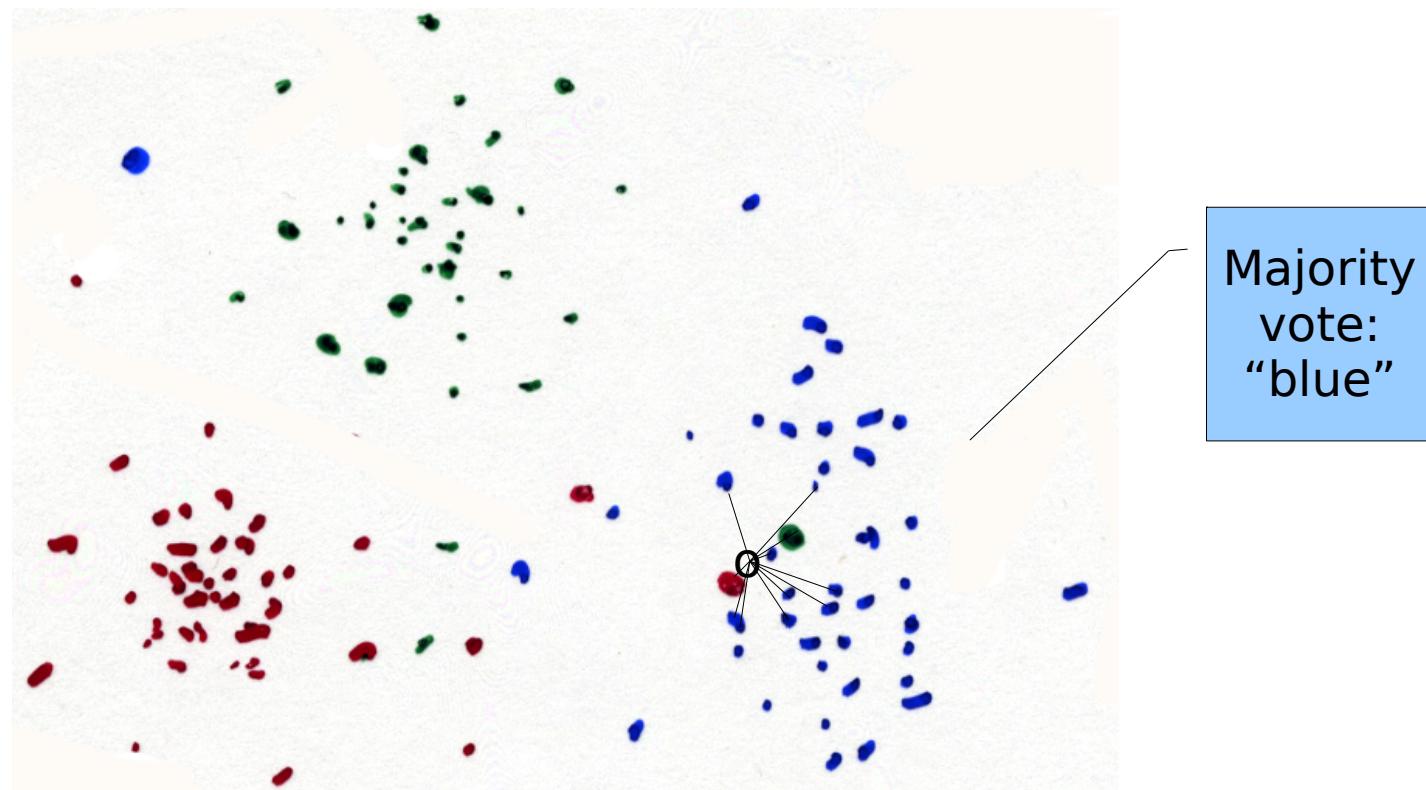
# Nearest-Neighbor Classifier

- Simplest non-parametric classifier.
- Noisy data can be a problem.



# k-Nearest-Neighbor

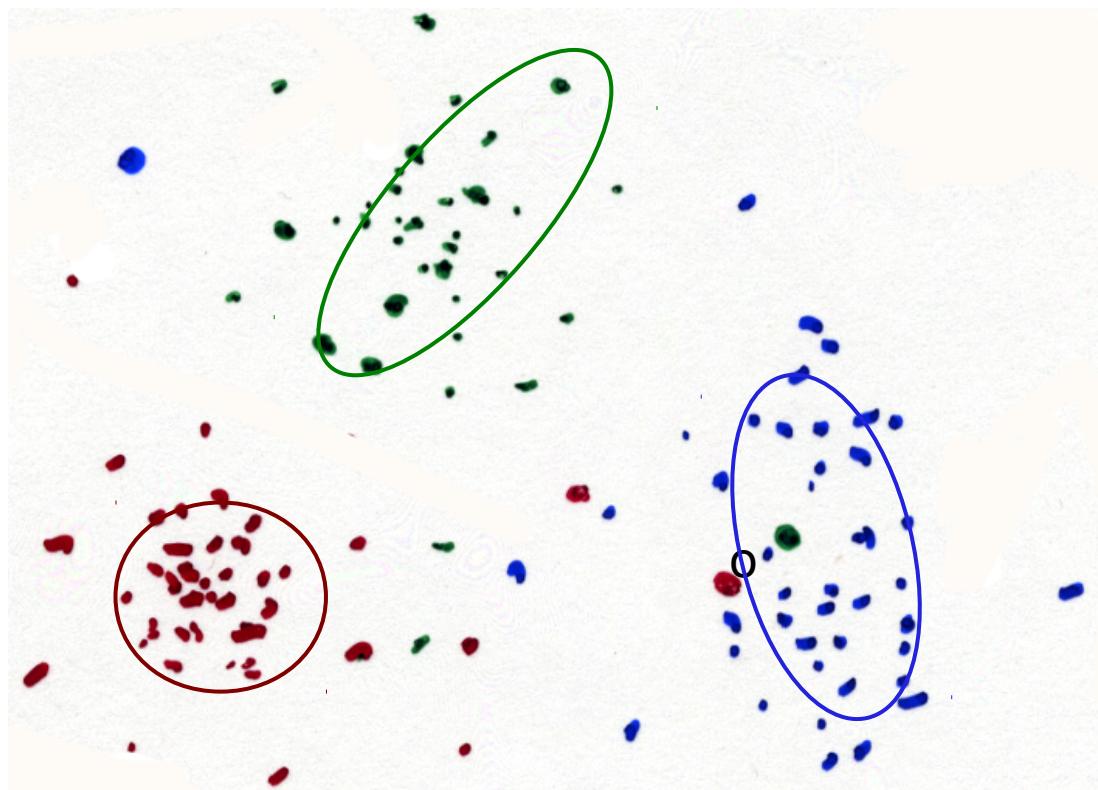
- Still a non-parametric classifier.
- Majority vote to find the correct class.



# Gaussian Mixture Model

- Parametric model: Gaussian distributions.

$$p(x \mid \mu, \sigma) = e^{\frac{-(x-\mu)^2}{\sigma^2}}$$



How do we  
find the  
correct  
values of  
the  
parameters?

Learning  
algorithm!

# scikit-learn

- Open source collection of machine learning algorithms implemented in Python.
- Documentation at [scikit-learn.org](http://scikit-learn.org)
- Installed on the lab machines:  
`import sklearn`

# Sample Learning Problem: Color Classes

- Assume objects come in a small number of colors.
- We want to know what the colors are.
  - This is a clustering problem.
- Given a new object, we want to determine its color class.
  - This is a classification problem.

# Training Data: RGB Values



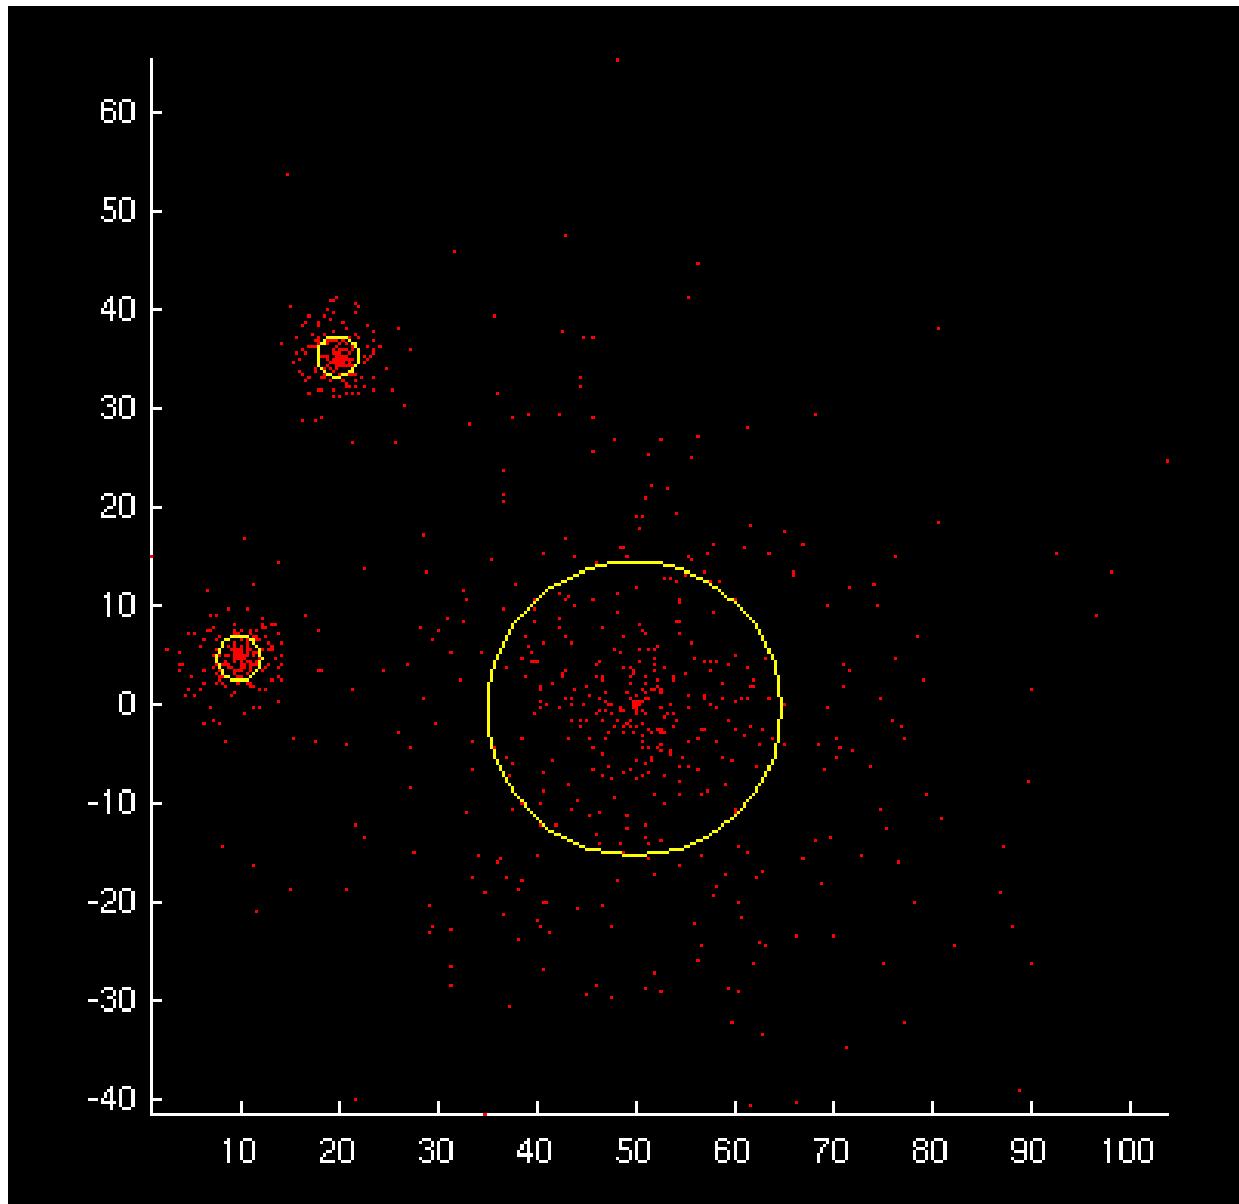
# Feature Space

- Three-dimensional feature space: RGB.
- $320 \times 240 = 76,800$  data points per image.
- Let's assume that each color class can be modeled as a gaussian distribution:
  - Mean color  $\mu$
  - Covariance matrix  $\Sigma$

# Expectation-Maximization Algorithm

- Unsupervised learning algorithm for finding clusters in data.
- Learns the  $\mu$  and  $\Sigma$  parameters for a set of gaussians.
- You must guess the number of classes.
- Runs quickly but can get stuck in local minima.

# E-M Clustering



# E-M Algorithm

- Expectation step:
  - For each point  $\mathbf{x}$ , for each gaussian  $(\mu_i, \Sigma_i)$ , calculate the likelihood of  $\mathbf{x}$  having been generated by the  $i$ -th gaussian:  $P(\mathbf{x}|\mu_i, \Sigma_i)$ .
- Maximization step:
  - For each gaussian, recalculate its mean and covariance  $\mu_i, \Sigma_i$  based on the likelihood-weighted data points.
- Repeat for several iterations.

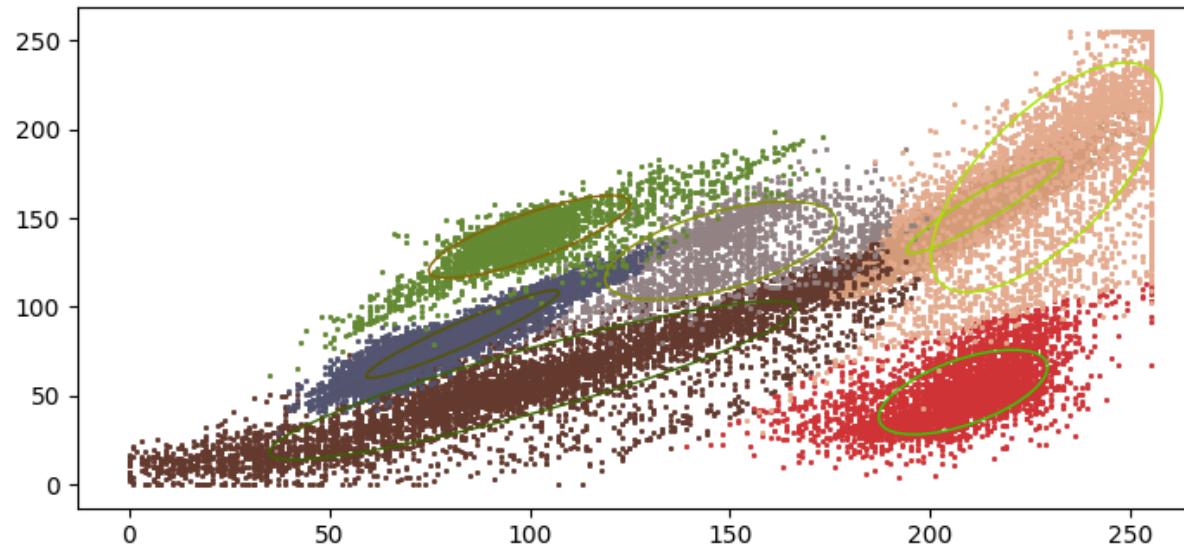
# E-M in scikit-learn

```
gmm = GaussianMixture(n_components=7)  
gmm.fit(data)  
means = gmm.means_  
covariances = gmm.covariances_  
  
classes = gmm.predict(new_data)
```

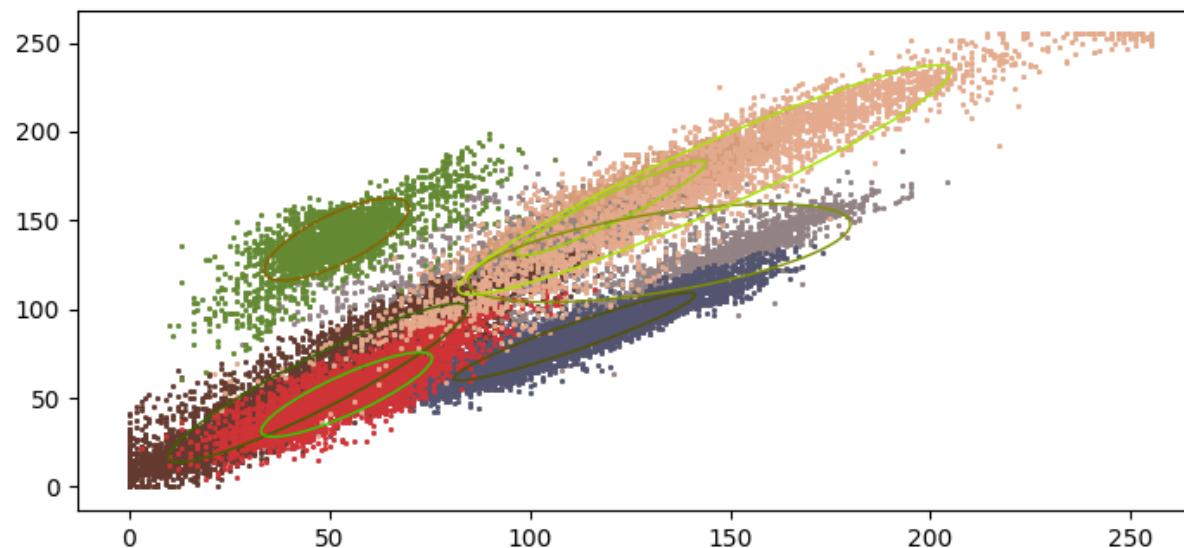
# Color Classification: Visualizing the Result

- Each color class is modeled as an ellipsoid in 3D space (RGB space).
- Too hard to plot. So instead:
  - Generate R-vs-G and B-vs-G plots.
  - Draw the ellipses in feature space determined by the covariance values.

# Scatter Plots With Gaussian Ellipses



Red  
vs.  
Green

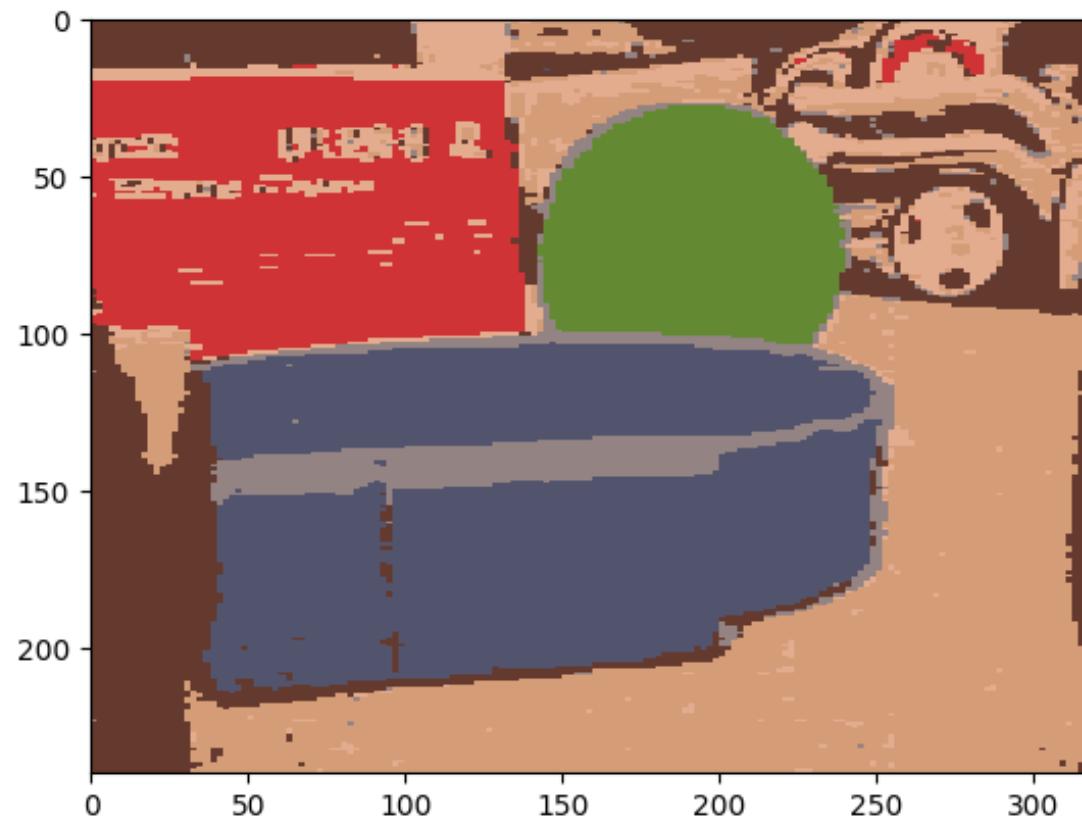


Blue  
vs.  
Green

# Classification

- Once we've learned the color classes, we can assign a class to each pixel.
  - This gives us a color-quantized image.
- Can then use these color classes to classify new images the same way.

# Classified Image



# Refinements

- Can detect local minima by checking the density of points near the mean of the gaussian.
- Split/merge EM can reallocate gaussians if some are being wasted and others are spread between two clusters.
- BayesianGaussianMixture class in scikit-learn can infer the number of effective components.