

# Deep Learning: and Deep Data-Science



ROYAL INSTITUTE  
OF TECHNOLOGY

12 May 2015

[roelof@kth.se](mailto:roelof@kth.se)

[www.csc.kth.se/~roelof/](http://www.csc.kth.se/~roelof/)

**Graph Technologies R&D**

[roelof@graph-technologies.com](mailto:roelof@graph-technologies.com)

slides online at:

<https://www.slideshare.net/roelofp/deep-learning-as-a-catdog-detector>

# BUT FIRST...

are you a...

**CAT PERSON?**

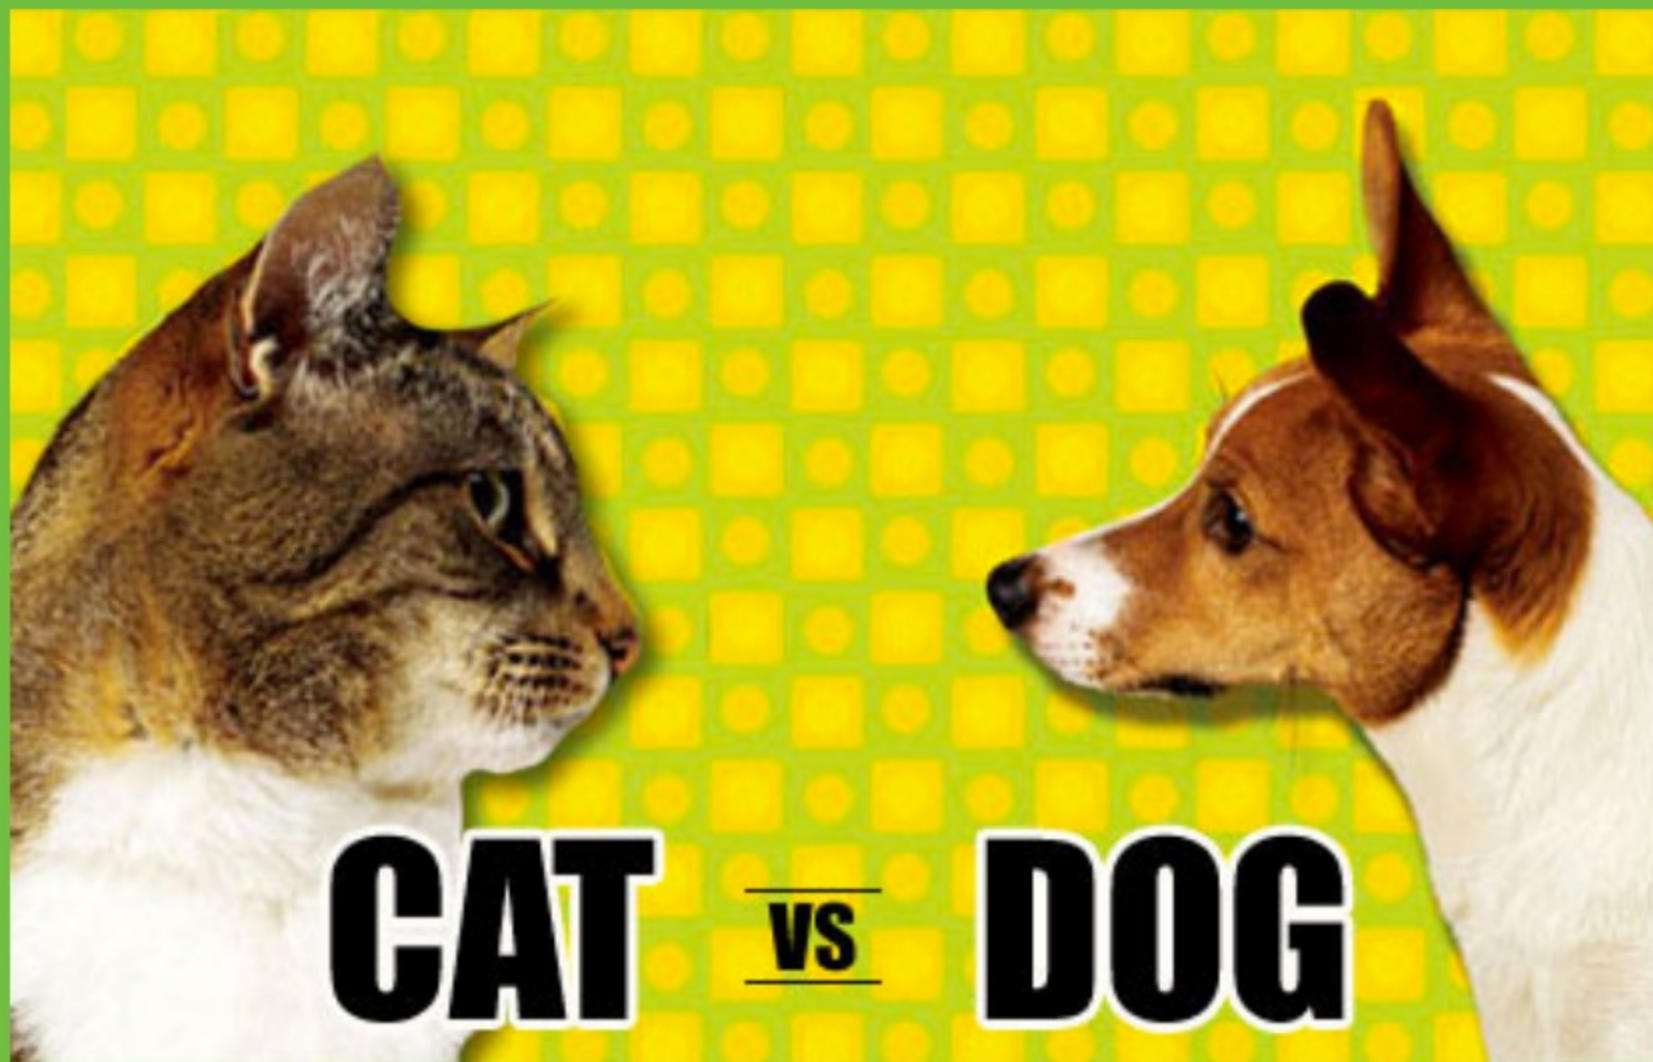


**DOG PERSON?**





in the next few minutes  
we'll be making a



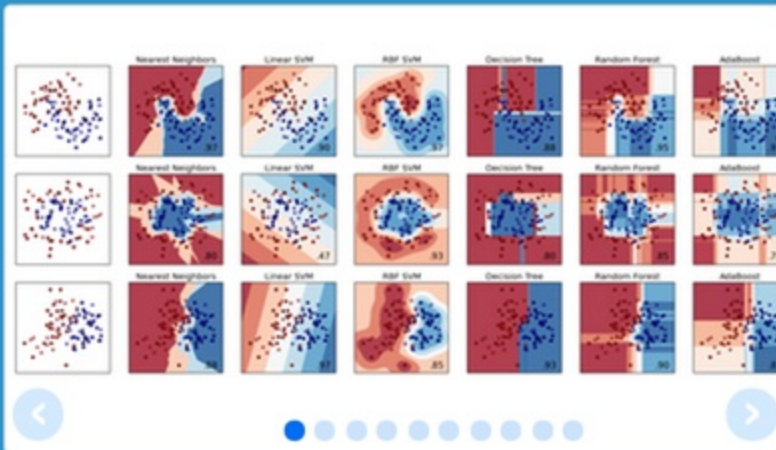
DETECTOR

# main Libraries



- scikit-learn (machine learning)  
<http://scikit-learn.org>
- caffe (deep learning) – for training deep neural nets  
(for today: loading a pre-trained one)  
<http://caffe.berkeleyvision.org>
- theano (efficient gpu-powered math)  
<http://www.deeplearning.net/software/theano/>
- ipython notebook  
<http://ipython.org/notebook.html>





# scikit-learn

Machine Learning in Python

- Simple and efficient tools for data mining and data analysis
- Accessible to everybody, and reusable in various contexts
- Built on NumPy, SciPy, and matplotlib
- Open source, commercially usable - BSD license

## Classification

Identifying to which category an object belongs to.

**Applications:** Spam detection, Image recognition.

**Algorithms:** *SVM, nearest neighbors, random forest, ...* — *Examples*

## Regression

Predicting a continuous-valued attribute associated with an object.

**Applications:** Drug response, Stock prices.

**Algorithms:** *SVR, ridge regression, Lasso, ...* — *Examples*

## Clustering

Automatic grouping of similar objects into sets.

**Applications:** Customer segmentation, Grouping experiment outcomes

**Algorithms:** *k-Means, spectral clustering, mean-shift, ...* — *Examples*

## Dimensionality reduction

Reducing the number of random variables to consider.

**Applications:** Visualization, Increased efficiency

**Algorithms:** *PCA, feature selection, non-negative matrix factorization.* — *Examples*

## Model selection

Comparing, validating and choosing parameters and models.

**Goal:** Improved accuracy via parameter tuning

**Modules:** *grid search, cross validation, metrics.* — *Examples*

## Preprocessing

Feature extraction and normalization.

**Application:** Transforming input data such as text for use with machine learning algorithms.

**Modules:** *preprocessing, feature extraction.* — *Examples*

# Caffe

---


Deep learning framework  
by the [BVLC](#)

Created by

[Yangqing Jia](#)

Lead Developer

[Evan Shelhamer](#)

 [View On GitHub](#)

# Caffe

---

Caffe is a deep learning framework made with expression, speed, and modularity in mind. It is developed by the Berkeley Vision and Learning Center ([BVLC](#)) and by community contributors. [Yangqing Jia](#) created the project during his PhD at UC Berkeley. Caffe is released under the [BSD 2-Clause license](#).

Check out our web image classification [demo](#)!

## Why Caffe?

---

**Expressive architecture** encourages application and innovation. Models and optimization are defined by configuration without hard-coding. Switch between CPU and GPU by setting a single flag to train on a GPU machine then deploy to commodity clusters or mobile devices.

**Extensible code** fosters active development. In Caffe's first year, it has been forked by over 1,000 developers and had many significant changes contributed back. Thanks to these contributors the framework tracks the state-of-the-art in both code and models.

**Speed** makes Caffe perfect for research experiments and industry deployment. Caffe can process **over 60M images per day** with a single NVIDIA K40 GPU\*. That's 1 ms/image for inference and 4 ms/image for learning. We believe that Caffe is the fastest convnet implementation available.

**Community:** Caffe already powers academic research projects, startup prototypes, and even large-scale industrial applications in vision, speech, and multimedia. Join our community of brewers on the [caffe-users group](#) and [Github](#).

\* With the ILSVRC2012-winning [SuperVision](#) model and caching IO. Consult performance [details](#).

## Welcome

Theano is a Python library that allows you to define, optimize, and evaluate mathematical expressions involving multi-dimensional arrays efficiently. Theano features:

- **tight integration with NumPy** – Use `numpy.ndarray` in Theano-compiled functions.
- **transparent use of a GPU** – Perform data-intensive calculations up to 140x faster than with CPU.(float32 only)
- **efficient symbolic differentiation** – Theano does your derivatives for function with one or many inputs.
- **speed and stability optimizations** – Get the right answer for `log(1+x)` even when `x` is really tiny.
- **dynamic C code generation** – Evaluate expressions faster.
- **extensive unit-testing and self-verification** – Detect and diagnose many types of mistake.

Theano has been powering large-scale computationally intensive scientific investigations since 2007. But it is also approachable enough to be used in the classroom (IFT6266 at the University of Montreal).

## News

- We support [cuDNN](#) if it is installed by the user.
- Open Machine Learning Workshop 2014 [presentation](#).
- Colin Raffel [tutorial on Theano](#).
- Ian Goodfellow did a [12h class with exercises on Theano](#).

# theano

## Table Of Contents

[Welcome](#)  
[News](#)  
[Download](#)  
[Status](#)  
[Citing Theano](#)  
[Documentation](#)  
[Community](#)

## Next topic

## Release Notes

## This Page

## Show Source

## Quick search

Go

Enter search terms or a module, class or function name.

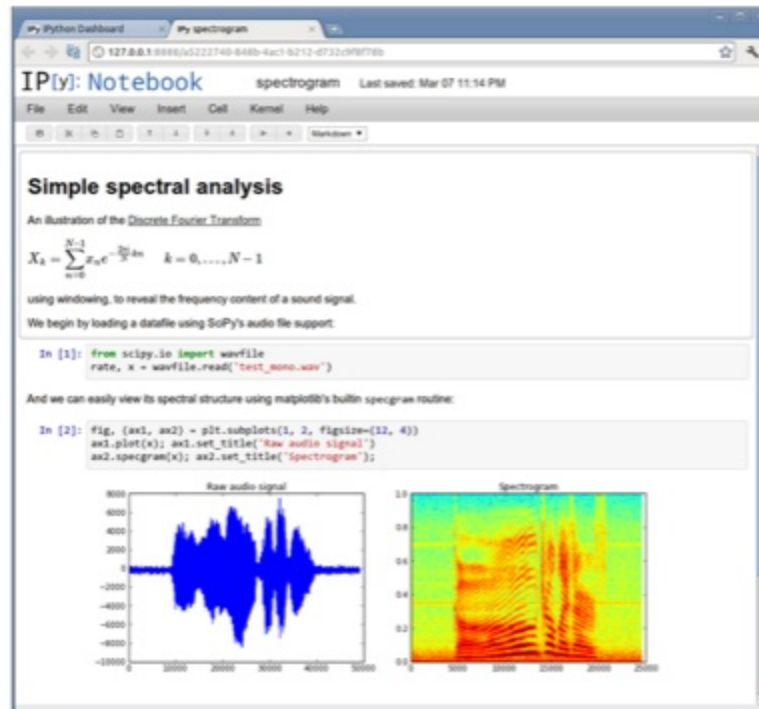


# IP[y]: IPython Interactive Computing

[Install](#) · [Docs](#) · [Videos](#) · [News](#) · [Cite](#) · [Sponsors](#) · [Donate](#)

## The IPython Notebook

The IPython Notebook is an interactive computational environment, in which you can combine code execution, rich text, mathematics, plots and rich media, as shown in this example session:



It aims to be an agile tool for both exploratory computation and data analysis, and provides a platform to support **reproducible research**, since all inputs and outputs may be stored in a one-to-one way in notebook documents.

Google™ Custom Search

## VERSIONS

**Stable**

3.1 – April 2015

[Install](#)

**Development**

4.0.dev

[GitHub](#)

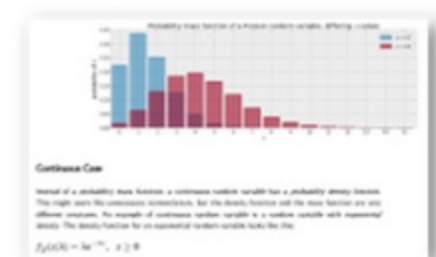
**Offline Docs**

All Versions

[GitHub](#)

## NOTEBOOK VIEWER

Share your notebooks





**Code is ahead, soon...  
I promise :)**

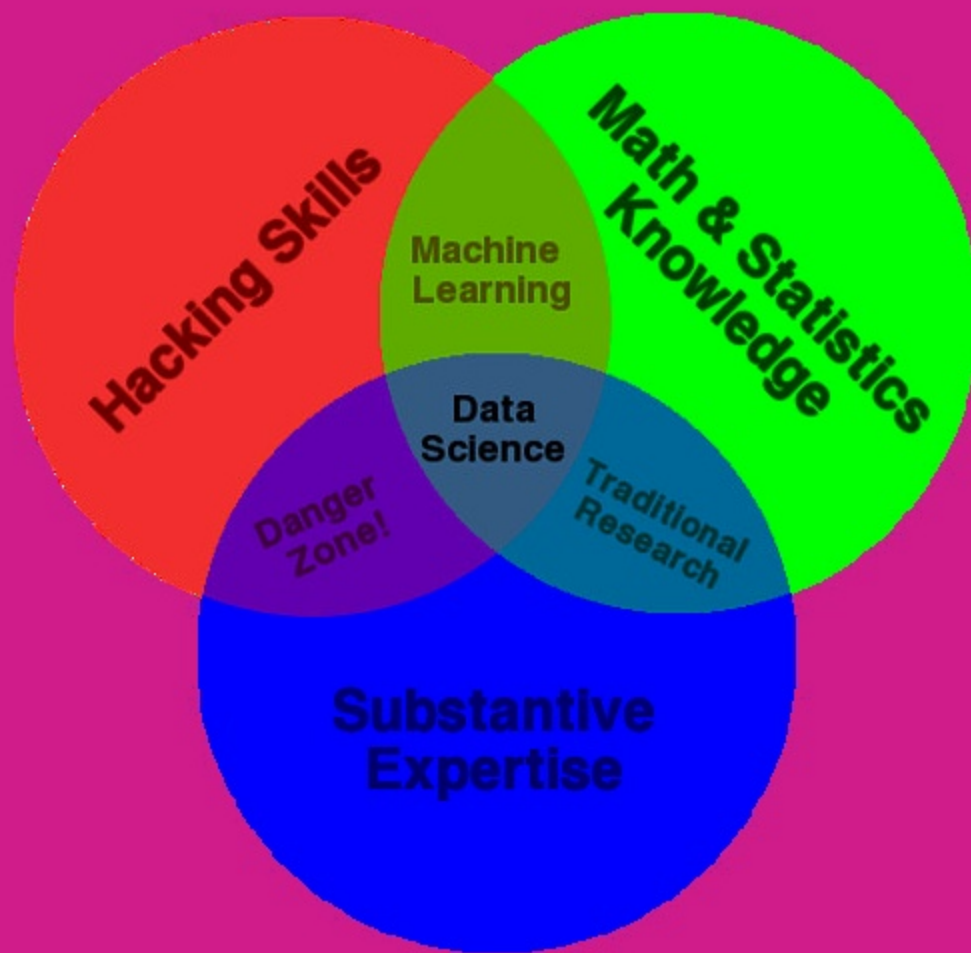
**BEAR  
WITH  
ME**



# Data Science ?

“Data science is clearly a blend of the **hackers' art, statistics** and **machine learning...**”

—Hilary Mason & Chris Wiggins, 2010



(Drew Connoway 2010)



> Features = Awesomeness

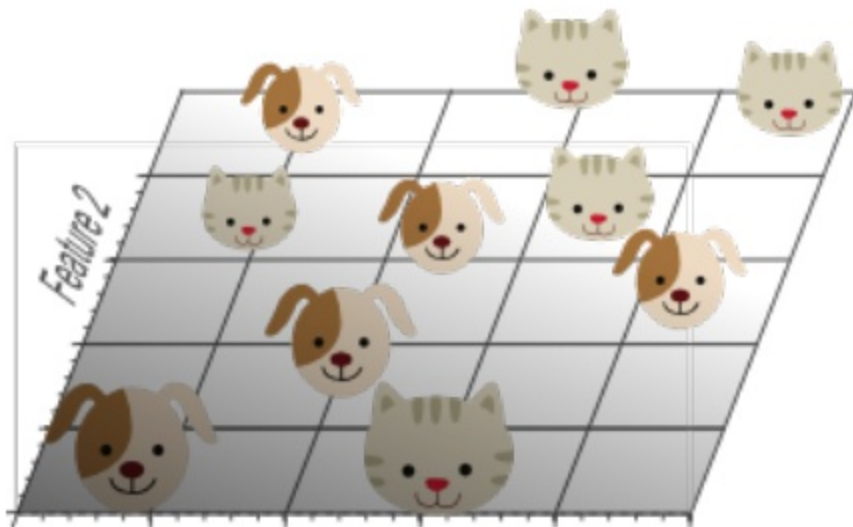


1 feature

# > Features = Awesomeness

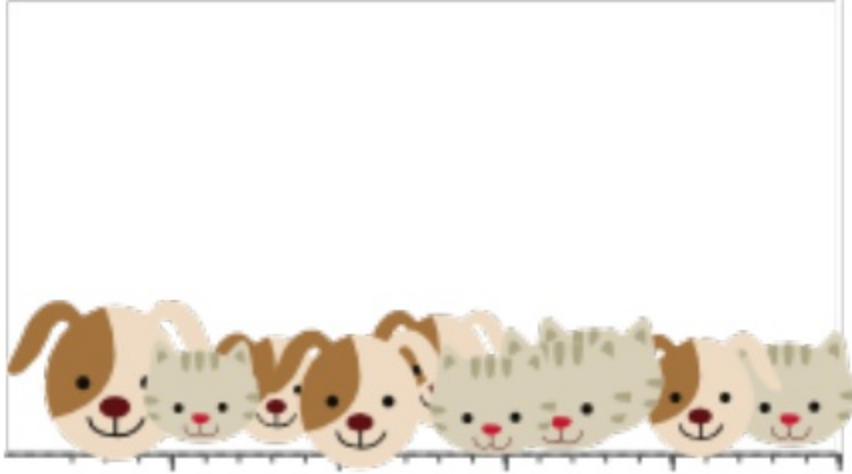


1 feature

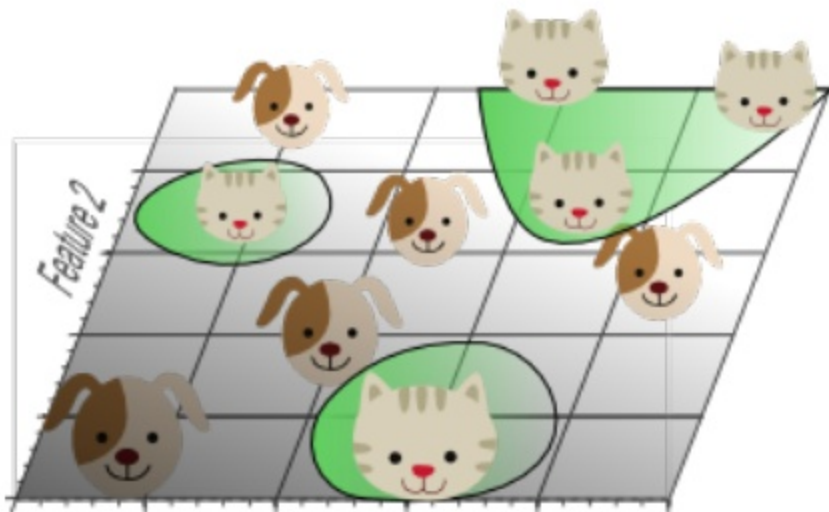


Feature 1

2 features



1 feature



2 features

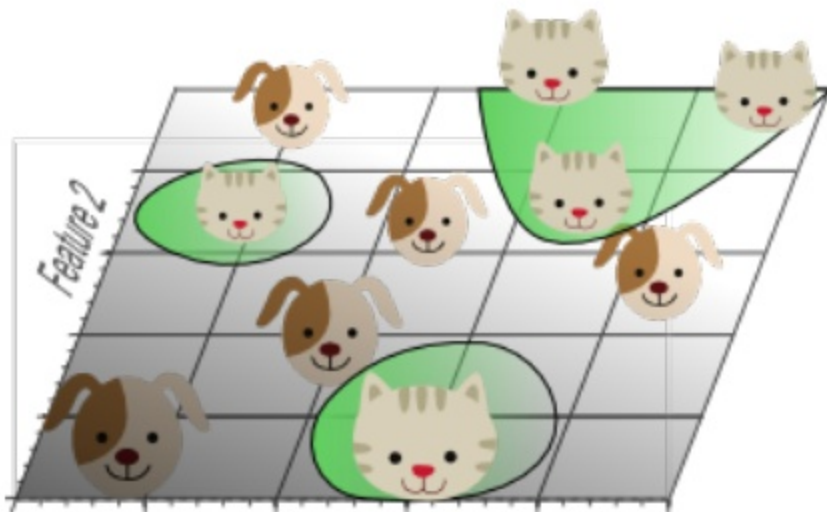
too few features/dimensions = overfitting



# > Features = Awesomeness

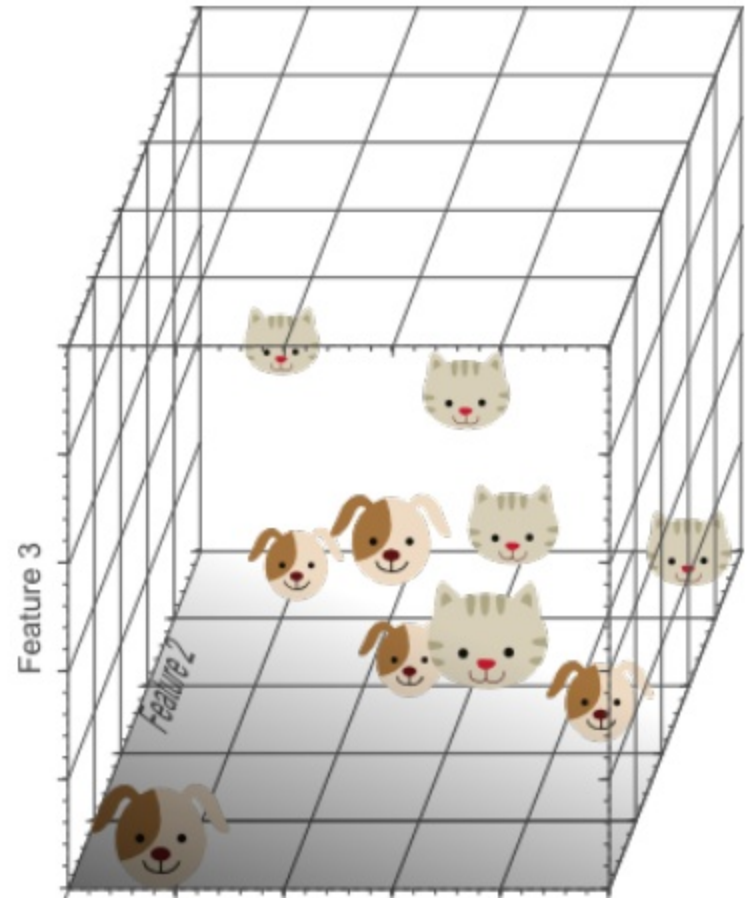


1 feature



Feature 1

2 features



Feature 1

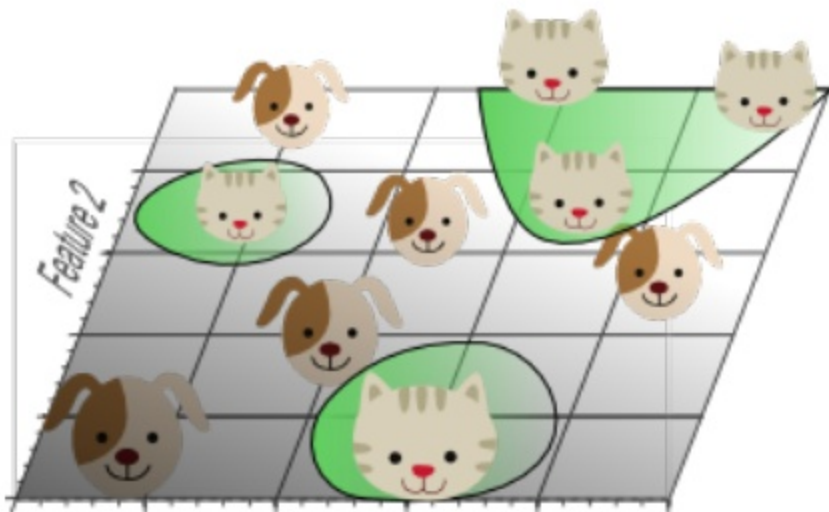
3 features

too few features/dimensions = overfitting

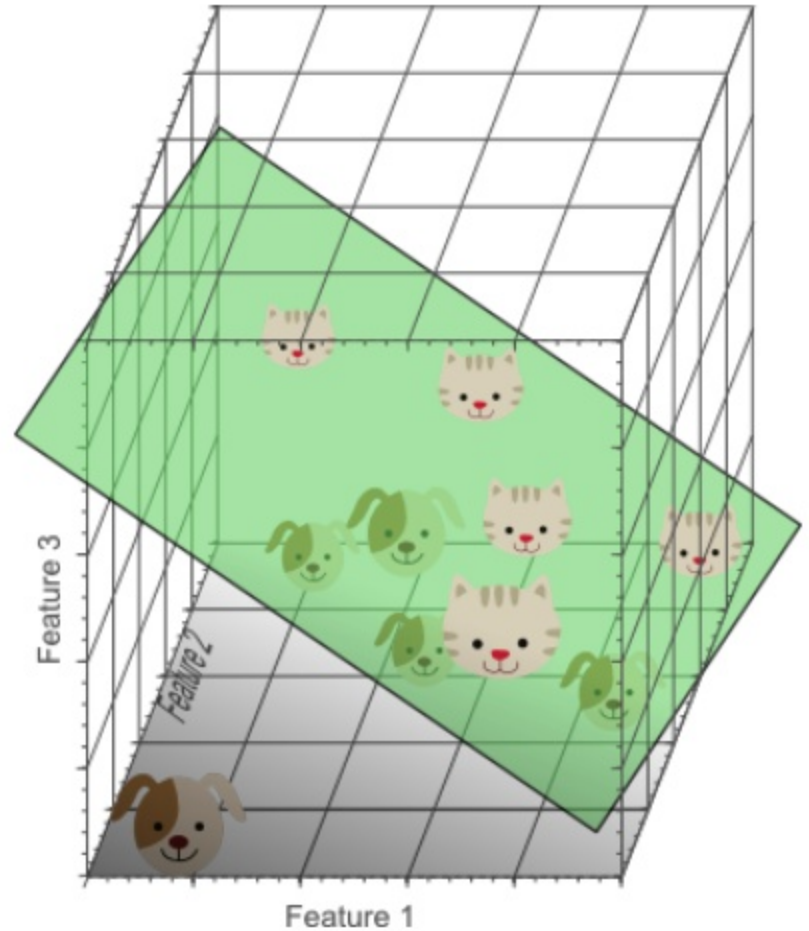
More Features = Awesomeness!



1 feature

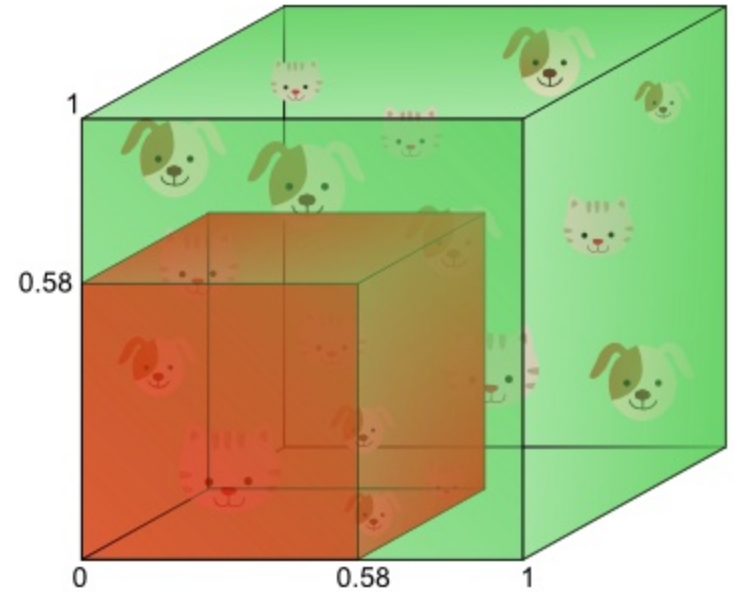
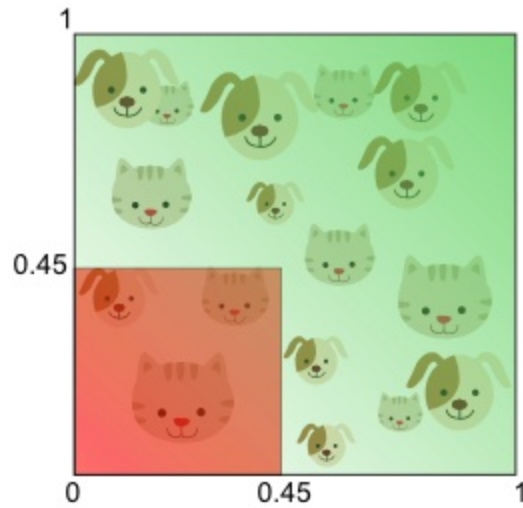


2 features

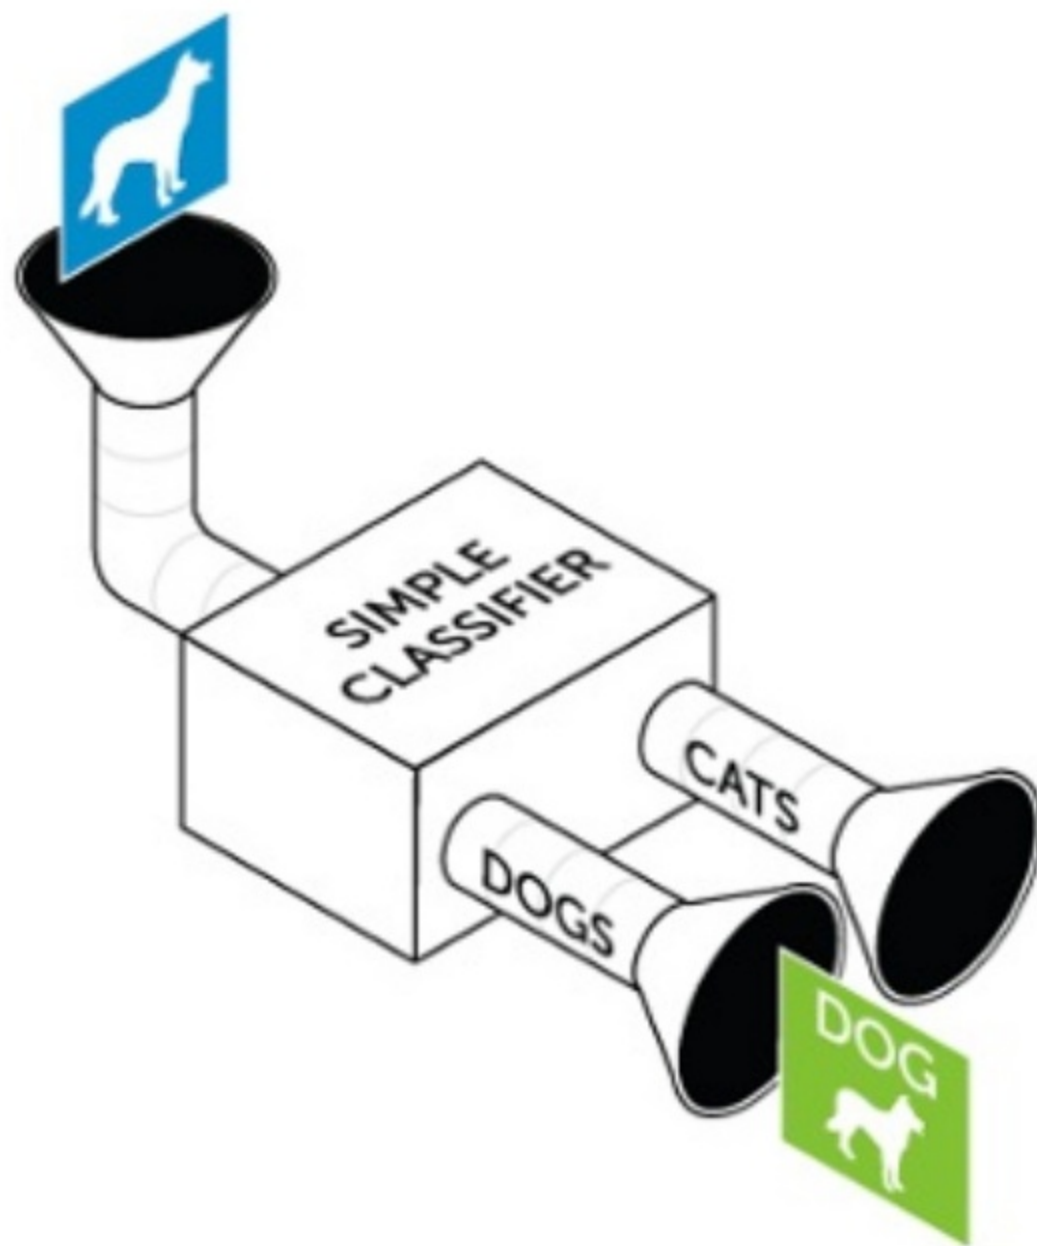


3 features

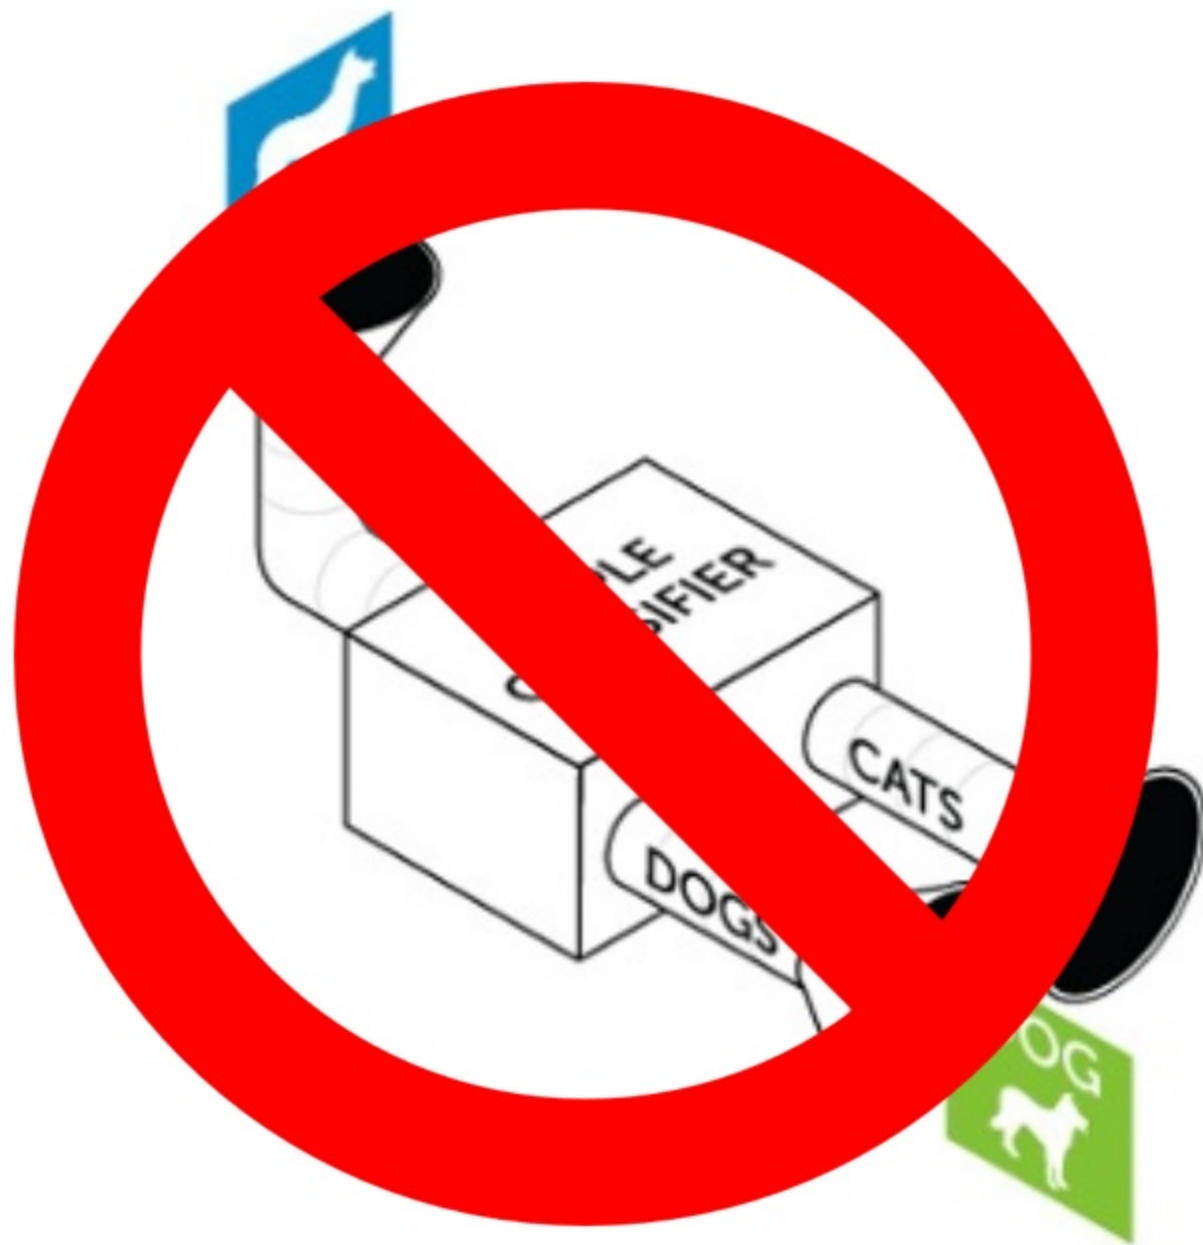
# ++ Data Needs also grow!







(picture by Dato)



(picture by Dato)

# Deep Learning?

- A host of statistical machine learning techniques
- Enables the automatic learning of feature hierarchies
- Generally based on artificial neural networks