

CIS 519/419

Applied Machine Learning

www.seas.upenn.edu/~cis519

Dan Roth

danroth@seas.upenn.edu

<http://www.cis.upenn.edu/~danroth/>

461C, 3401 Walnut

Slides were created by Dan Roth (for CIS519/419 at Penn or CS446 at UIUC), Eric Eaton for CIS519/419 at Penn, or from other authors who have made their ML slides available.

CIS(4,5)19: Applied Machine Learning

- **Tuesday, Thursday:** 1:30pm-3:00pm 101 Levine

Registration to Class

- **Office hours:** Tue/Thur 4:30-5:30 pm [my office]
- 9 TAs
- **Assignments:** 5 Problems set (Python Programming)
 - Weekly (light) on-line quizzes
- **Weekly Discussion Sessions**
- **Mid Term Exam**
- **[Project]**
- **Final**
- **No real textbook:**
 - Mitchell/Flach/Other Books/ **Lecture notes** /Literature

Go to the web site

Be on Piazza

CIS519: Today

- What is Learning?
- Who are you?
- What is CIS519 about?
- The Badges Game [...](#)

An Owed to the Spelling Checker

- I have a spelling checker, it came with my PC
- It plane lee marks four my revue
- Miss steaks aye can knot sea.
- Eye ran this poem threw it, your sure reel glad two no.
- Its vary polished in it's weigh
- My checker tolled me sew.
- A checker is a bless sing, it freeze yew lodes of thyme.
- It helps me right awl stiles two reed
- And aides me when aye rime.
- Each frays come posed up on my screen
- Eye trussed to bee a joule...

Machine learning is everywhere

Mail thinks this message is Junk Mail. Load Images Not Junk

From: IEEE World Congress on Multimedia <iccsa2013@yahoo.com> [Hide](#)
Subject: First CFP Submission :15 August, 2013 World Congress on Multimedia & Computer science: October 04-06, 2013, Hammamet, Tunisia
Date: July 17, 2013 9:01:20 AM CDT
To: Julia Hockenmaier
Reply-To: iccsa2013@yahoo.com

7 Attachments, 374 KB Save Quick Look

First- Call For Papers Submission : 15th of



World Congress on Multimedia and Computer science (WCMCS' 2013)
October 04-06, 2013, Hammamet, Tunisia

Iberostar Sa



Azure Machine Learning - Microsoft.com
Powerful Machine Learning Service Will All the Benefits of the Cloud.

Predictive Sales - 6sense.com
6Sense Tells You Who Will Buy, When They'll Buy, and How Much.

Machine Learning Mitchell - Amazon.com
Books to Satisfy Anyone from New Learners to Computer Geeks.

Amazon.com: Machine Learning Books
www.amazon.com/Machine-Learning_Books/?pf_rd_p=... - Amazon.com -
Results 1 - 10 - A catalogue record for Machine Learning from a great selection
at Books Store.

INTRODUCTION TO MACHINE LEARNING - Alex Smola
alex-smola.org/downloads/thisbook.pdf -
This page has pointers to my draft book on Machine Learning and to its...

Machine Learning Books - MachineLearning - Reddit
www.reddit.com/r/MachineLearning/wiki/MachineLearning_Books -
Results 1 - 10 - I have been collecting machine learning books over the past couple
months. It seems that machine learning professors are good about posting...

FAQ: What machine learning book should I start with? - Ge...
benmabey.com/wiki/faq-what-machine-learning-book-should-i-start-with-h...
Oct 7, 2011 - What makes a good starting book is depending on your background
(specifically) ... It gives the reader a resursion on several machine learning

What are some good books on machine learning? - Quora
www.quora.com/What-are-some-good-books-on-machine-learning? -
Answer 1 of 15: This answer attempts the very ambitious problem of producing an
approximately complete list. Please leave comments and tell me what's wrong?

Google's Machine Learning
cloud.google.com/compute/docs/leverage-gcp-machine-learning/
Leverage Google's Machine Learning
Algorithm Using Prediction API

Excel Data Mining Course
www.celestec.com/learn-how-to-get-predictive-insight/
Learn how to get predictive insight
using Excel's data mining features.

Revolutionizing Industry
www.ge.com/learn-how-the-industrial-internet-&-
ge-are-changing-the-way-we-work/

Recommended for You
LOOK INSIDE! **Semantics, Set Linguistics**
by Kate Kearns (May 15)
In Stock
List Price: \$40.00
Price: \$37.31
68 used & new from \$14.00
Add to Cart Add to Wish List

Because you purchased...
Meaning: A Slim Guide to Semantics (Oxford Linguistics) (Paperback)
by Paul Elbourne (Author)

Shop for machine learning books on Google

Introduction to Machine Learning
\$46.12
Barnes & No...

Pattern Recognition
\$55.62
Textbooks.com

Machine Learning A
\$81.40
Barnes & No...

Foundations of Machine Learning
\$62.23
Barnes & No...

Machine Learning for
\$18.05
Bookbyte

Machine Learning
\$128.70
Nebo.com

Machine Learning
\$72.45
Barnes & No...

Machine Learning An
\$35.07
ValueBooks...

Rate this item
 ☆☆☆☆
 I own it
 Not interested

This was a gift
 Don't use for recommendations

Help | C

From: English - detected To: Chinese (Simplified)

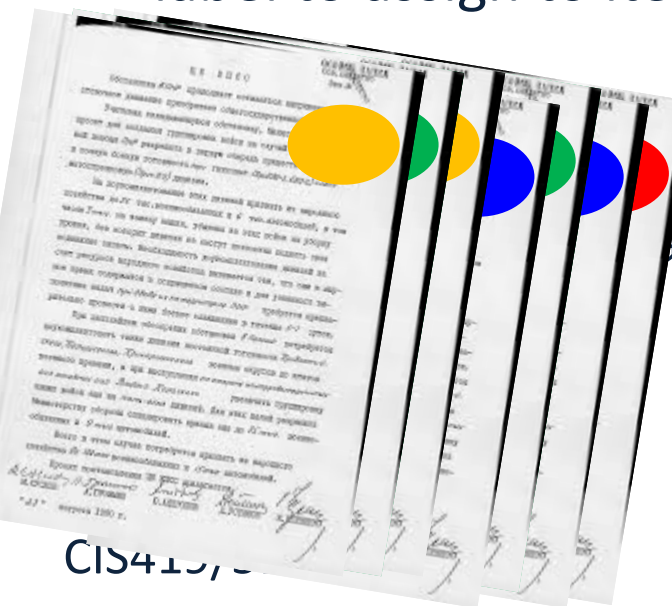
The blue fox jumps over the hedge

蓝狐跨越对冲

Translate

Applications: Spam Detection

- This is a **binary classification task**: Assign **one of two labels** (i.e. yes/no) to the input (here, an email message)
- Classification requires **a model** (a classifier) to determine which label to assign to items.




Documents

- Documents
- Sentences
- Phrases
- Images
- Medical records
-

Labels

- Politics, Sports, Finance
- Positive, Negative
- Person, Location
- cats, dogs, snakes
- Admit again soon/Not



Comprehension

(ENGLAND, June, 1989) - Christopher Robin is alive and well. He lives in England. He is the same person that you read about in the book, Winnie the Pooh. As a boy, Chris lived in a pretty home called Cotchfield Farm. When Chris was three years old, his father wrote a poem about him. The poem was printed in a magazine for others to read. Mr. Robin then wrote a book. He made up a fairy tale land where Chris lived. His friends were animals. There was a bear called Winnie the Pooh. There was also an owl and a young pig, called a piglet. All the animals were stuffed toys that Chris owned. Mr. Robin made them come to life with his words. The places in the story were all near Cotchfield Farm. Winnie the Pooh was written in 1925. Children still love to read about Christopher Robin and his animal friends. Most people don't know he is a real person who is grown now. He has written two books of his own. They tell what it is like to be famous.

1. Christopher Robin was born in England.
2. Winnie the Pooh is a title of a book.
3. Christopher Robin's dad was a magician.
4. Christopher Robin must be at least 65 now.

This is an Inference Problem; where is the learning?



Learning

- **Learning is at the core of**
 - Understanding High Level Cognition
 - Performing knowledge intensive inferences
 - Building adaptive, intelligent systems
 - Dealing with messy, real world data
 - Analytics
- **Learning has multiple purposes**
 - Knowledge Acquisition
 - Integration of various knowledge sources to ensure robust behavior
 - Adaptation (human, systems)
 - Decision Making (Predictions)


Learning = Generalization

H. Simon -

“Learning denotes changes in the system that are adaptive in the sense that they enable the system to do the task or tasks drawn from the same population more efficiently and more effectively the next time.”

The ability to perform a task in a situation which has never been encountered before

Learning = Generalization



Mail thinks this message is junk mail.

Not junk

- The learner has to be able to classify items it has never seen before.

Learning = Generalization

The ability to perform a task in a situation
which has never been encountered before

- Classification
 - Medical diagnosis; credit card applications; hand-written letters; ad selection; sentiment assignment,...
- Planning and acting
 - Game playing (chess, backgammon, go); driving a car
- Skills
 - (A robot) balancing a pole; playing tennis
- Common sense reasoning
 - Natural language interactions

What does the algorithm
get as input?
(features)

Generalization depends on the **Representation** as much as
it depends on the **Algorithm** used.

Same Population?



New Zealand

In ~~New York State~~, the longest period of daylight occurs during the month of _____.

Why Study Machine Learning?

- “A breakthrough in machine learning would be worth ten Microsofts”
-Bill Gates, Chairman, Microsoft
- “Machine learning is the next Internet”
-Tony Tether, Former Director, DARPA
- Machine learning is the hot new thing”
-John Hennessy, President, Stanford
- “Web rankings today are mostly a matter of machine learning”
-Prabhakar Raghavan, Dir. Research, Yahoo
- “Machine learning is going to result in a real revolution”
-Greg Papadopoulos, CTO, Sun
- “Machine learning is today’s discontinuity”
-Jerry Yang, CEO, Yahoo

Why Study Learning?

- Computer systems with new capabilities.
- Understand human and biological learning
- Understanding teaching better.
- Time is right.
 - Initial algorithms and theory in place.
 - Growing amounts of on-line data
 - Computational power available.

 - Necessity: many things we want to do cannot be done by “programming”.
 - (Think about all the examples given earlier)

Learning is the future

- Learning techniques will be a basis for every application that involves a connection to the messy real world
- Basic learning algorithms are ready for use in applications today
- Prospects for broader future applications make for exciting fundamental research and development opportunities
- Many unresolved issues – Theory and Systems
 - While it's hot, there are many things we don't know how to do

Work in Machine Learning

- Artificial Intelligence; Theory; Experimental CS
- Makes Use of:
 - Probability and Statistics; Linear Algebra; Theory of Computation;
- Related to:
 - Philosophy, Psychology (cognitive, developmental), Neurobiology, Linguistics, Vision, Robotics,....
- Has applications in:
 - AI (Natural Language; Vision; Planning; Robotics; HCI)

Very active field

What to teach?

The fundamental paradigms

Some of the most important algorithmic ideas

Modeling

And: what we
don't know

Course Overview

- Introduction: Basic problems and questions
- A detailed example: Linear classifiers; key algorithmic idea
- Two Basic Paradigms:
 - Discriminative Learning & Generative/Probabilistic Learning
- Learning Protocols:
 - Supervised; Unsupervised; Semi-supervised
- Algorithms
 - Gradient Descent
 - Decision Trees
 - Linear Representations: (Perceptron; SVMs; Kernels)
 - Neural Networks/Deep Learning
 - Probabilistic Representations (naïve Bayes)
 - Unsupervised /Semi supervised: EM
 - Clustering; Dimensionality Reduction
- Modeling; Evaluation; Real world challenges
- Ethics

CIS(4,5)19: Applied Machine Learning

- Tuesday, Thursday: 1:30pm-3:00pm 101 Levine

Registration to Class

- Office hours: Tue/Thur 4:30-5:30 pm [my office]
- 9 TAs
- Assignments: 5 Problems set (Python Programming)
 - Weekly (light) on-line quizzes
- Weekly Discussion Sessions
- Mid Term Exam
- [Project]
- Final
- No real textbook:
 - Mitchell/Flach/Other Books/ **Lecture notes** /Literature

Go to the web site

Be on Piazza

CIS519: Machine Learning

- What do you need to know:
 - Some exposure to:
 - Theory of Computation
 - Probability Theory
 - Linear Algebra
 - Programming (Python)

- Homework 0

Participate, Ask Questions

CIS 519: Policies

- Cheating
 - **No.**
 - We take it very seriously.
- Homework:
 - Collaboration is encouraged
 - But, you have to write your own solution/code.
- Late Policy:
 - You have a credit of 4 days (4*24hours); That's it.
- Grading:
 - Possible separate for 419/519.
 - 40% - homework; ; 20%-final; 15%-midterm; 5% Quizzes
 - [Projects: 20%]
- Questions?

[Class' Web Page](#)

**Note also the Schedule Page
and our Notes**

CIS519 on the web

- Check our class website:
 - Schedule, slides, videos, policies
 - <http://www.seas.upenn.edu/~cis519/spring2018/>
 - Sign up, participate in our Piazza forum:
 - Announcements and discussions
 - <http://piazza.com/upenn/spring2018/cis419519>
 - Check out our team
 - Office hours
 - [Optional] Discussion Sessions
- Scribing the Class [Good writers; Latex]?

What is Learning

- The Badges Game.....
 - This is an example of the key learning protocol: supervised learning
- First question: Are you sure you got it?
 - Why?
- Issues:
 - Prediction or Modeling?
 - Representation
 - Problem setting
 - Background Knowledge
 - When did learning take place?
 - Algorithm