


Machine Learning

ML

Artificial Intelligence → ML

Machine Learning



How does a human learn?

What is Machine Learning?

- A system that can improve on task T , with respect to performance measure P , after observing experience E .
- Task: Distinguish  
- Experience: Labeled instances of Deer & Horses
- Performance: Accuracy

Machine learning applications



Task? Performance? Experience?

Central challenge in machine learning

How can we build computer systems that automatically improve with experience, and what laws govern learning in general?

- Statistics
 - What can be inferred from a set of data, with what reliability?
- Computer science
 - How can we build computers to solve problems, and which problems are tractable/intractable?
- Human learning
 - What mechanisms explain learning in humans, and what teaching strategies are most effective?

Eric Xing 2006

Why use machine learning?

- Data comes in too fast for humans to process
 - Every credit card transaction
 - Every e-mail message
- Data set is just too large for humans to process
 - Protein folding
 - Sloan Digital Sky Survey
- Machines can make decisions faster
 - Once trained, many models predict almost instantly
- Personalization / adaptation
 - Speech recognition
 - ...



sds.org

What's going on in the wizard's head?



- What are the **concepts** or **models** being learned?
- We'll talk about three kinds
 - Rules
 - Linear models
 - Memory-based

Rule-based models

- A sequence of if-then
 - Just like Matlab's **if-else**!
- Need to express rules explicitly
- Example: grammar/spelling checker

Linear models

- Represent a problem as a set of **features**; each feature gets a number of points
- Example: is this document about soccer?
 - Contains "soccer": 50 points
 - Contains "basketball": -50 points
 - Contains "Beckham": 100 points
 - Contains "Posh Spice": -100 points
 - ...
 - If total number of points > 0, say "yes"

Memory-based models

- The model is the training data!



Which is closer?

Very, very hard to write an algorithm but data is easy to collect

How do we get these models?

- Labeled training data
 - Humans have to
 - Collect data (emails, pictures of animals, ...)
 - Label the data (spam vs. not-spam; deer vs. Horse, ...)



- There are many algorithms
 - We'll discuss one: Naïve Bayes

Experimentation

- Need to train and evaluate
- Split data into a **training** set and a **test** set
 - Train the wizard on the training data
 - Evaluate the wizard on the test data
- Lots of data is needed to get a wise wizard, so why not use the whole data set for training?
 - If you evaluate the wizard on (any part of the) training data, it's like letting the wizard "cheat" on a test


Measuring the wizard's skill

- Simplest measure is **accuracy** – on what fraction of the test cases does the wizard predict correctly?
- Accuracy is not a good measure in some cases
 - E.g., credit card fraud
 - Very rare event → always negative = 99.9% accurate
 - Better measure: **false positive rate**, **false negative rate**
- **Precision** and **recall** (remember them?)

Other machine learning topics

- Active learning
 - Labels are expensive!
 - Remember CAPTCHA?
 - Image labeling game

reCAPTCHA




Is OCR "learning" when you "teach" the system how to read/recognize the characters in the word?

Image labeling game



Other machine learning topics

- Active learning
 - Labels are expensive!
 - Remember CAPTCHA?
 - Image labeling game
- Unsupervised learning
 - Learning without the correct answers
- Theory
 - How much data do you need to learn something?
 - What kinds of concepts can you learn?



- Also known as “junk mail” or “unsolicited bulk” e-mail”
 - Unsolicited – you didn't ask for it
 - Bulk – sent to lots and lots of people, not just you
- Typical legal definition: **unsolicited commercial** email from someone without a **pre-existing** business relationship
- Huge problem
 - 50% of all e-mail sent is spam

Much of this material is from Joshua Goodman's SPAM tutorial

What are the recent spam topics

Why is there spam?

- **Money!**
- Almost free advertisement
- Cost of sending spam ~0.01 cent per message.
- If 1 in 100,000 people buy, and I earn \$11 for that purchase, then I make a profit!

Spamming techniques to defeat filters

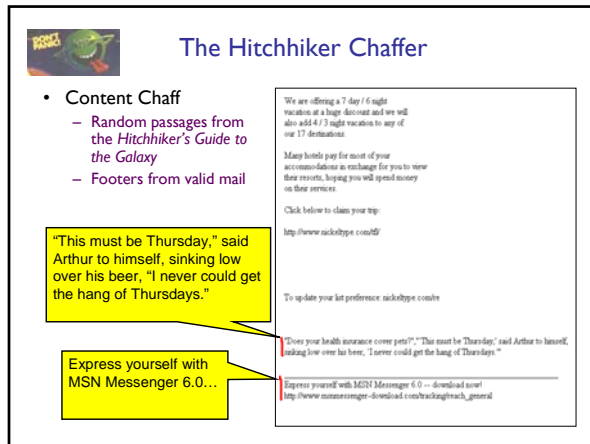
- Content in image
- Chaff
 - Text chaff
 - Content chaff
- Obscuring words



Weather Report Guy

- Content in Image
- Good Word Chaff

Weather, Sunny, High 82, Low 81, Favorite...



The Hitchhiker Chaffer

- Content Chaff
 - Random passages from the *Hitchhiker's Guide to the Galaxy*
 - Footers from valid mail

"This must be Thursday," said Arthur to himself, sinking low over his beer, "I never could get the hang of Thursdays."

Express yourself with MSN Messenger 6.0...

We are offering a 7 day / 6 night vacation at a huge discount and we will also add 4 / 3 night vacations to any of our 17 destinations.

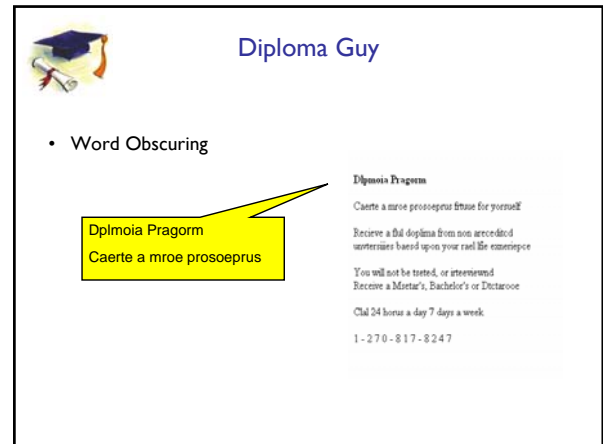
Many hotels pay for most of your accommodations in exchange for you to view their resorts, hoping you will spend money on their services.

Click below to claim your trip:
<http://www.sacktype.com/ff/>

To update your list preference: [sacktype.com/ff/](http://www.sacktype.com/ff/)

"Does your health insurance cover pets?"; "This must be Thursday," said Arthur to himself, sinking low over his beer, "I never could get the hang of Thursdays."

Express yourself with MSN Messenger 6.0 -- download now!
http://www.msnmessenger-download.com/tracking/track_general



Diploma Guy

- Word Obscuring

Diplmoia Pragorm
Caerte a mroe prosoeprus

Diplmoia Pragorm

Caerte a mroe prosoeprus fime for yornelf

Recieve a fdl diploma from non accredited universities based upon your real life experiece

You will not be tested, or interarround
Recieve a Master's, Bachelor's or Dictatoroo

Chld 24 horsu a day 7 days a week

1 - 270 - 817 - 8247

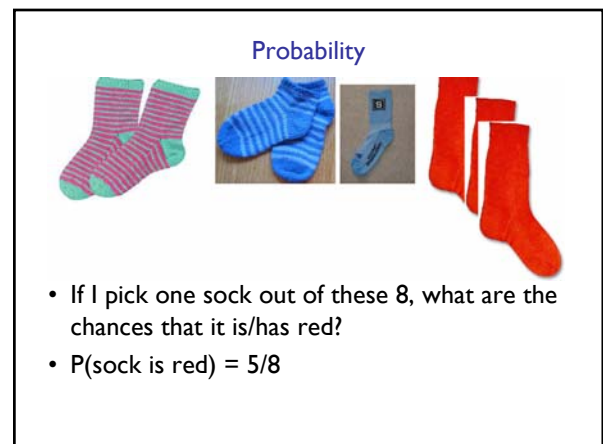
Besides email, where else do you get spam?

Besides email, where else do you get spam?

- "old media"
 - Physical junk mail
 - Phone calls
- Instant messenger
- Chat rooms
- Popups
- Link spam

Naïve Bayes spam filtering

- Most common kind of spam filter
 - Although many filters include rules and other features
- Who is Bayes?
 - 18th century mathematician
 - Let's learn some probability



Probability

If I pick one sock out of these 8, what are the chances that it is/has red?

$P(\text{sock is red}) = 5/8$

Conditional Probability



- $P(\text{sock is red}) = 5/8$
- What if I decide I'm going to pick a solid sock?
- $P(\text{sock is red} \mid \text{solid pattern}) = 3/4$

Bayes' Theorem



- Thomas Bayes

- $P(\text{red} \mid \text{solid}) = P(\text{solid} \mid \text{red}) * P(\text{red}) / P(\text{solid})$
- $3/4 = 3/5 * 5/8 / 4/8$



Naïve Bayes machine learning for spam classification

- Use vector space model of messages
- Decide which is bigger:
 - $P(\text{message is spam} \mid \text{words in message})$
 - or
 - $P(\text{message is not spam} \mid \text{words in message})$
- Use Bayes' rule

Computation of the "probabilities"

- Simplify the computation by
 - computing a score instead of the probability
 - computing $\log(P)$ instead of P
- Prediction of **spam** if

$$\text{score}(\text{spam}) > \text{score}(\text{not-spam})$$
 Otherwise predict **not-spam**
- Why is it "naïve"?
 - Assume the probability of the words in the message to be independent

What are some solutions for ending spam?

- Filtering
 - Machine learning
 - Blackhole lists (IP filtering)
 - Whitelisting
- Postage
 - Money
 - Turing tests
 - Other computation