

CSC 411: Lecture 01: Introduction

Class based on Raquel Urtasun & Rich Zemel's lectures

Sanja Fidler

University of Toronto

Jan 11, 2016

Today

- Administration details
- Why is machine learning so cool?

The Team

- **Instructor:**



Sanja Fidler (fidler@cs.toronto.edu)

- **Office:** 283B in Pratt

- **Office hours:** Mon 1.15-2.30pm, or by appointment

- **TAs:**



Shenlong Wang (slwang@cs.toronto.edu)



Ladislav Rampasek (rampasek@cs.toronto.edu)



Boris Ivanovic (boris.ivanovic@mail.utoronto.ca)

- Liberal wrt waiving pre-requisites
 - ▶ But it is up to you to determine if you have the appropriate background
- Do I have the appropriate background?
 - ▶ **Linear algebra:** vector/matrix manipulations, properties
 - ▶ **Calculus:** partial derivatives
 - ▶ **Probability:** common distributions; Bayes Rule
 - ▶ **Statistics:** mean/median/mode; maximum likelihood
 - ▶ Sheldon Ross: A First Course in Probability

Course Information

- **Class:** Mondays and Wednesday at noon-1pm in LM158
- **Tutorials:** Fridays, same hour as lecture, same classroom
- **Class Website:**
<http://www.cs.toronto.edu/~fidler/teaching/2015/CSC411.html>
- The class will use Piazza for **announcements** and **discussions**:
<https://piazza.com/utoronto.ca/winter2016/csc411/home>
- First time, sign up here:
<https://piazza.com/utoronto.ca/winter2016/csc411>
- Your grade will **not depend on your participation on Piazza**. It's just a good way for asking questions, discussing with your instructor, TAs and your peers

Textbook(s)

- Christopher Bishop: "*Pattern Recognition and Machine Learning*", 2006

Textbook(s)

- Christopher Bishop: "*Pattern Recognition and Machine Learning*", 2006
- Other Textbooks:
 - ▶ Kevin Murphy: "*Machine Learning: a Probabilistic Perspective*"
 - ▶ David Mackay: "*Information Theory, Inference, and Learning Algorithms*"
 - ▶ Ethem Alpaydin: "*Introduction to Machine Learning*", 2nd edition, 2010.

Requirements

- Do the [readings!](#)

Requirements

- Do the [readings!](#)
- [Assignments:](#)
 - ▶ Three assignments, first two worth 12.5% each, last one worth 15%, for a total of 40%
 - ▶ Programming: take Matlab/Python code and extend it
 - ▶ Derivations: pen(cil)-and-paper

Requirements

- Do the [readings!](#)
- [Assignments:](#)
 - ▶ Three assignments, first two worth 12.5% each, last one worth 15%, for a total of 40%
 - ▶ Programming: take Matlab/Python code and extend it
 - ▶ Derivations: pen(cil)-and-paper
- [Mid-term:](#)
 - ▶ One hour exam on Feb 29th
 - ▶ Worth 25% of course mark

Requirements

- Do the [readings!](#)
- [Assignments:](#)
 - ▶ Three assignments, first two worth 12.5% each, last one worth 15%, for a total of 40%
 - ▶ Programming: take Matlab/Python code and extend it
 - ▶ Derivations: pen(cil)-and-paper
- [Mid-term:](#)
 - ▶ One hour exam on Feb 29th
 - ▶ Worth 25% of course mark
- [Final:](#)
 - ▶ Focused on second half of course
 - ▶ Worth 35% of course mark

More on Assignments

- **Collaboration** on the assignments is not allowed. Each student is responsible for his/her own work. Discussion of assignments should be limited to clarification of the handout itself, and should not involve any sharing of pseudocode or code or simulation results. Violation of this policy is grounds for a semester grade of F, in accordance with university regulations.

More on Assignments

- **Collaboration** on the assignments is not allowed. Each student is responsible for his/her own work. Discussion of assignments should be limited to clarification of the handout itself, and should not involve any sharing of pseudocode or code or simulation results. Violation of this policy is grounds for a semester grade of F, in accordance with university regulations.
- The **schedule** of assignments is included in the syllabus. Assignments are due at the beginning of class/tutorial on the due date.

More on Assignments

- **Collaboration** on the assignments is not allowed. Each student is responsible for his/her own work. Discussion of assignments should be limited to clarification of the handout itself, and should not involve any sharing of pseudocode or code or simulation results. Violation of this policy is grounds for a semester grade of F, in accordance with university regulations.
- The **schedule** of assignments is included in the syllabus. Assignments are due at the beginning of class/tutorial on the due date.
- Assignments handed in **late** but before 5 pm of that day will be penalized by 5% (i.e., total points multiplied by 0.95); a late penalty of 10% per day will be assessed thereafter.

More on Assignments

- **Collaboration** on the assignments is not allowed. Each student is responsible for his/her own work. Discussion of assignments should be limited to clarification of the handout itself, and should not involve any sharing of pseudocode or code or simulation results. Violation of this policy is grounds for a semester grade of F, in accordance with university regulations.
- The **schedule** of assignments is included in the syllabus. Assignments are due at the beginning of class/tutorial on the due date.
- Assignments handed in **late** but before 5 pm of that day will be penalized by 5% (i.e., total points multiplied by 0.95); a late penalty of 10% per day will be assessed thereafter.
- **Extensions** will be granted only in special situations, and you will need a Student Medical Certificate or a written request approved by the instructor at least one week before the due date.

More on Assignments

- **Collaboration** on the assignments is not allowed. Each student is responsible for his/her own work. Discussion of assignments should be limited to clarification of the handout itself, and should not involve any sharing of pseudocode or code or simulation results. Violation of this policy is grounds for a semester grade of F, in accordance with university regulations.
- The **schedule** of assignments is included in the syllabus. Assignments are due at the beginning of class/tutorial on the due date.
- Assignments handed in **late** but before 5 pm of that day will be penalized by 5% (i.e., total points multiplied by 0.95); a late penalty of 10% per day will be assessed thereafter.
- **Extensions** will be granted only in special situations, and you will need a Student Medical Certificate or a written request approved by the instructor at least one week before the due date.
- Final assignment is a **bake-off**: competition between ML algorithms. We will give you some data for training a ML system, and you will try to develop the best method. We will then determine which system performs best on unseen test data.

Calendar

Date	Topic	Assignments
Jan 11	Introduction	
Jan 13	Linear Regression	
Jan 15	<i>Probability for ML & Linear regression</i>	
Jan 18	Linear Classification	
Jan 20	Logistic Regression	
Jan 22	<i>Optimization for ML</i>	
Jan 25	Nonparametric Methods	
Jan 27	Decision Trees	
Jan 29	<i>kNN & Decision Trees</i>	Asst 1 Out
Feb 1	Multi-class Classification	
Feb 3	Probabilistic Classifiers	
Feb 5	Probabilistic Classifiers II	
Feb 8	Neural Networks I	
Feb 10	Neural Networks II	Asst 1 In
Feb 22	<i>Naive Bayes and Gaussian Bayes Classifier</i>	
Feb 24	<i>Neural Networks Tutorial</i>	
Feb 26	<i>Mid-term review</i>	
Feb 29	MIDTERM	

Date	Topic	Assignments
Mar 2	Clustering	Assit 2 Out
Mar 4	<i>Clustering</i>	
Mar 7	Mixture of Gaussians & EM	
Mar 9	PCA & Autoencoders	
Mar 11	<i>PCA Tutorial</i>	
Mar 14	Kernels and Margins	Asst 2 In
Mar 16	Support Vector Machines	
Mar 18	<i>SVM Tutorial</i>	Asst3 Out
Mar 21	Ensemble Methods I	
Mar 23	Ensemble Methods II	
Mar 28	Bayesian Methods	
Mar 30	Reinforcement Learning I	
Apr 1	<i>Bagging & Boosting</i>	
Apr 4	Reinforcement Learning II	
Apr 6	Final & Wrap-up	Ass 3 In

What is Machine Learning?

- How can we solve a specific problem?

What is Machine Learning?




- How can we solve a specific problem?
 - ▶ As computer scientists we write a program that encodes a set of rules that are useful to solve the problem



Figure: How can we make a robot cook?

What is Machine Learning?

- How can we solve a specific problem?
 - ▶ As computer scientists we write a program that encodes a set of rules that are useful to solve the problem
 - ▶ In many cases is very difficult to specify those rules, e.g., given a picture determine whether there is a cat in the image



What is Machine Learning?

- How can we solve a specific problem?
 - ▶ As computer scientists we write a program that encodes a set of rules that are useful to solve the problem
 - ▶ In many cases is very difficult to specify those rules, e.g., given a picture determine whether there is a cat in the image
- Learning systems are not directly programmed to solve a problem, instead develop own program based on:

What is Machine Learning?

- How can we solve a specific problem?
 - ▶ As computer scientists we write a program that encodes a set of rules that are useful to solve the problem
 - ▶ In many cases is very difficult to specify those rules, e.g., given a picture determine whether there is a cat in the image
- Learning systems are not directly programmed to solve a problem, instead develop own program based on:
 - ▶ Examples of how they should behave

What is Machine Learning?

- How can we solve a specific problem?
 - ▶ As computer scientists we write a program that encodes a set of rules that are useful to solve the problem
 - ▶ In many cases is very difficult to specify those rules, e.g., given a picture determine whether there is a cat in the image
- Learning systems are not directly programmed to solve a problem, instead develop own program based on:
 - ▶ Examples of how they should behave
 - ▶ From trial-and-error experience trying to solve the problem


What is Machine Learning?

- How can we solve a specific problem?
 - ▶ As computer scientists we write a program that encodes a set of rules that are useful to solve the problem
 - ▶ In many cases is very difficult to specify those rules, e.g., given a picture determine whether there is a cat in the image
- Learning systems are not directly programmed to solve a problem, instead develop own program based on:
 - ▶ Examples of how they should behave
 - ▶ From trial-and-error experience trying to solve the problem
- Different than standard CS:
 - ▶ Want to implement unknown function, only have access to sample input-output pairs (training examples)

What is Machine Learning?

- How can we solve a specific problem?
 - ▶ As computer scientists we write a program that encodes a set of rules that are useful to solve the problem
 - ▶ In many cases is very difficult to specify those rules, e.g., given a picture determine whether there is a cat in the image
- Learning systems are not directly programmed to solve a problem, instead develop own program based on:
 - ▶ Examples of how they should behave
 - ▶ From trial-and-error experience trying to solve the problem
- Different than standard CS:
 - ▶ Want to implement unknown function, only have access to sample input-output pairs (training examples)
- Learning simply means incorporating information from the training examples into the system

Tasks that requires machine learning: What makes a 2?



Robots learn to cook by watching YouTube

When it comes to learning how to cook, it turns out that robots may not be so different from humans after all... or are they?

Sci-Tech

January 20, 2015

4:26 PM PST



by **Michelle Starr**

[@riding_red](#)



When it comes to teaching robots how to do things, there are some very key differences. A human knows what you mean when you say "I need a cup". A robot needs to be taught that that means it has to turn around, go to the cupboard, open it, take out the cup, close the cupboard, turn back around, return to you, manoeuvre the cup over the bench, and release the cup.



John T. Conzoli, UMD

This is one of the key parts of figuring out machine learning: How can you program a robot so that it can intuit that a plastic cup, a glass and a mug may all be classified under the general term "cup"? How can you design a robot that is able to teach itself?

One way, as researchers at the University of Maryland Institute for Advanced Computer Studies are finding out, is YouTube. More specifically, cooking tutorials on YouTube. By watching these videos, robots can learn the complicated series of grasping and manipulation motions required for

Why use learning?

- It is very hard to write programs that solve problems like recognizing a handwritten digit

Why use learning?

- It is very hard to write programs that solve problems like recognizing a handwritten digit
 - ▶ What distinguishes a 2 from a 7?

Why use learning?

- It is very hard to write programs that solve problems like recognizing a handwritten digit
 - ▶ What distinguishes a 2 from a 7?
 - ▶ How does our brain do it?

Why use learning?

- It is very hard to write programs that solve problems like recognizing a handwritten digit
 - ▶ What distinguishes a 2 from a 7?
 - ▶ How does our brain do it?
- Instead of writing a program by hand, we collect examples that specify the correct output for a given input

Why use learning?

- It is very hard to write programs that solve problems like recognizing a handwritten digit
 - ▶ What distinguishes a 2 from a 7?
 - ▶ How does our brain do it?
- Instead of writing a program by hand, we collect examples that specify the correct output for a given input
- A machine learning algorithm then takes these examples and produces a program that does the job

Why use learning?

- It is very hard to write programs that solve problems like recognizing a handwritten digit
 - ▶ What distinguishes a 2 from a 7?
 - ▶ How does our brain do it?
- Instead of writing a program by hand, we collect examples that specify the correct output for a given input
- A machine learning algorithm then takes these examples and produces a program that does the job
 - ▶ The program produced by the learning algorithm may look very different from a typical hand-written program. It may contain millions of numbers.


Why use learning?

- It is very hard to write programs that solve problems like recognizing a handwritten digit
 - ▶ What distinguishes a 2 from a 7?
 - ▶ How does our brain do it?
- Instead of writing a program by hand, we collect examples that specify the correct output for a given input
- A machine learning algorithm then takes these examples and produces a program that does the job
 - ▶ The program produced by the learning algorithm may look very different from a typical hand-written program. It may contain millions of numbers.
 - ▶ If we do it right, the program works for new cases as well as the ones we trained it on.

Learning algorithms are useful in many tasks

1. **Classification:** Determine which discrete category the example is

Examples of Classification



What digit is this?

Examples of Classification



Is this a dog?

Examples of Classification



what about this one?

Examples of Classification



Am I going to pass the exam?

Examples of Classification



Do I have diabetes?

Learning algorithms are useful in many tasks

1. **Classification:** Determine which discrete category the example is
2. **Recognizing patterns:** Speech Recognition, facial identity, etc

Examples of Recognizing patterns



Figure: Siri: <https://www.youtube.com/watch?v=8ciagGASro0>

Examples of Recognizing patterns




Figure: Photomath: <https://photomath.net/>

Learning algorithms are useful in other tasks

1. **Classification:** Determine which discrete category the example is
2. **Recognizing patterns:** Speech Recognition, facial identity, etc
3. **Recommender Systems:** Noisy data, commercial pay-off (e.g., Amazon, Netflix).

Examples of Recommendation systems

The screenshot shows the Netflix interface for the movie 'Despicable Me'. At the top, the Netflix logo is on the left, followed by a 'Browse' dropdown menu and 'DVD' text. A search bar contains 'despi' with a close button. On the right, there are notification and profile icons for 'Raquel'. Below the navigation bar is a movie poster for 'Despicable Me' featuring Gru and the Minions. The main content area displays the movie title 'Despicable Me' in large white text, followed by a 4-star rating, the year '2010', a 'G' rating icon, and a runtime of '1h 34m'. A synopsis follows: 'Villainous Gru hatches a plan to steal the moon from the sky. But he has a tough time staying on task after three orphans land in his care.' Below the synopsis, it lists the cast: 'Starring: Steve Carell, Jason Segel, Russell Brand' and genres: 'Genres: Children & Family Movies, Movies for ages 5 to 7, Movies for ages 8 to 10'. A quote from a user says 'This movie is: Goofy'. A 'MY LIST' button with a plus icon is visible. The background of the main area is a large image of two Minions with a play button overlay. At the bottom, there are three tabs: 'OVERVIEW' (which is underlined in red), 'MORE LIKE THIS', and 'DETAILS'.

NETFLIX Browse ▾ DVD

despi X

Raquel ▾

Despicable Me

★★★★☆ 2010 [G] 1h 34m

Villainous Gru hatches a plan to steal the moon from the sky. But he has a tough time staying on task after three orphans land in his care.

Starring: Steve Carell, Jason Segel, Russell Brand
Genres: Children & Family Movies, Movies for ages 5 to 7, Movies for ages 8 to 10
This movie is: Goofy

Steve Carell, Jason Segel, Russell Brand and Kristen Wiig lend their voices to this animated box office hit.

+ MY LIST

OVERVIEW MORE LIKE THIS DETAILS

Examples of Recommendation systems

NETFLIX Browse ▾ DVD

SEARCH Search 🔔 Raquel ▾

Homeland

★★★★★ 2013 TV-MA 3 Seasons

In this riveting Emmy-winning drama, a CIA agent suspects that a Marine who just returned after years in captivity has been turned into a terrorist.

Starring: Claire Danes, Mandy Patinkin, Damian Lewis
Genres: TV Shows, TV Action & Adventure, TV Dramas
This show is: Suspenseful

🏆 Claire Danes and Damian Lewis both won Emmys and Golden Globes for their performances in this intense series.

⊕ MY LIST

OVERVIEW EPISODES MORE LIKE THIS DETAILS

Examples of Recommendation systems

NETFLIX Browse ▾ DVD

🔍 I-robot X 🔔 🌟 Raquel ▾

Explore titles related to: I, Robot | Robot Chicken Star Wars: Episode II

Titles related to I, Robot

The image displays a grid of movie posters recommended by Netflix. The posters are arranged in three rows of five. The first row includes 'I Am Legend', 'Minority Report', 'Hellboy', 'Star Trek Into Darkness', and 'Deep Impact'. The second row includes 'Galaxy Quest', 'Total Recall', 'Terminator 2: Judgment Day', 'Jack Reacher', and 'Shooter'. The third row includes 'Men in Black 3', 'GI Joe: Retaliation', 'Enemy of the State', 'Parallels', and 'Paycheck'.

Learning algorithms are useful in other tasks

1. **Classification:** Determine which discrete category the example is
2. **Recognizing patterns:** Speech Recognition, facial identity, etc
3. **Recommender Systems:** Noisy data, commercial pay-off (e.g., Amazon, Netflix).
4. **Information retrieval:** Find documents or images with similar content

Examples of Information Retrieval



About 24,500 results (0.52 seconds)

[\[PDF\] CSC 411 MACHINE LEARNING and DATA MINING ...](#)

www.cs.toronto.edu/~zemel/documents/411/syl.pdf

CSC 411. MACHINE LEARNING and DATA MINING. Lectures: Monday, Wednesday 12-1 (section 1), 3-4 (section 2). Lecture Room: MP 134 (section 1); Bahen ...

[Professor Richard Zemel - Department of Computer Science](#)

www.cs.toronto.edu/~zemel/

Image Question Answering: A Visual Semantic Embedding Model and a New Dataset . Mengye Ren, Ryan Kiros, Richard Zemel. ICML 2015 Deep Learning ...

[Course Offerings](#) - [Research Interests](#) - [Students & Post Docs](#) - [Contact Info](#)

[UofT Machine Learning | Course](#)

learning.cs.toronto.edu/courses

CSC 411, Machine Learning and Data Mining (Raquel Urtasun and Richard Zemel); STA 4513, Statistical models of networks, graphs, and other relational ...

[CSC 411: Machine Learning and Data Mining](#)

www.cs.utoronto.ca/~radford/csc411.F06/

CSC 411: Machine Learning and Data Mining (Sept-Dec 2006). Note: The test on December 8 at 3pm will be held in BA B024, not the usual lecture/tutorial room.

[Worth taking CSC321 before CSC411? : UofT - Reddit](#)

https://www.reddit.com/r/UofT/.../worth_taking_csc321_before_csc411/

Jul 11, 2014 - However, CSC411 doesn't have CSC321 as a prerequisite, and it is not even ... Also, if I were to go straight for CSC411/412 without completing



Examples of Information Retrieval

Google artificial intelligence

Web News **Images** Videos Books More Search tools

SafeSearch Settings

Robot Movie Computer Computer Fifth Ge



Examples of Information Retrieval



artificial intelligence



Web

News

Images

Videos

Books

More ▾

Search tools

About 2,830,000 results (0.29 seconds)

'Artificial Intelligence is as dangerous as NUCLEAR ...




www.dailymail.co.uk/.../Artificial-Intelligence-dangero...

Jul 17, 2015

Artificial intelligence has the potential to be as dangerous to mankind as nuclear weapons, a leading pioneer ...

Rise of Future Technology | Artificial Intelligence - New ...



www.youtube.com/watch?v=YUvDBGYk17Y ▾

Dec 6, 2014 - Uploaded by Incredible Documentaries

Rise of Future Technology | **Artificial Intelligence** - New Documentary(2015)

Why You Shouldn't Fear Artificial Intelligence - YouTube



www.youtube.com/watch?v=uEWGjQ0nTm4 ▾

Jan 19, 2015 - Uploaded by DNews

Stephen Hawking and Elon Musk have warned us of the dangers of **Artificial Intelligence**, but is AI really ...

Artificial Intelligence - YouTube



www.youtube.com/watch?v=9TRv0cXUVQw

Aug 17, 2015 - Uploaded by The School of Life

Should we be scared of **artificial intelligence** and all it will bring us? Not so long as we remember to make sure ...

Examples of Information Retrieval



artificial intelligence



Web

News

Images

Videos

Books

More ▾

Search tools

About 32,400 results (0.42 seconds)

Artificial Intelligence: A Modern Approach



<https://books.google.ca/books?isbn=0136042597>

Stuart Jonathan Russell, Peter Norvig - 2010 - Snippet view - [More editions](#)

The revision of this best-selling text offers the most comprehensive, up-to-date introduction to the theory and practice of artificial intelligence.

Artificial Intelligence: A Modern Approach



<https://books.google.ca/books?isbn=1292024208>

Stuart Jonathan Russell, Peter Norvig - 2013 - No preview - [More editions](#)

In this third edition, the authors have updated the treatment of all major areas.

Artificial Intelligence: A Modern Approach



<https://books.google.ca/books?isbn=1405824824>

Stuart J. Russell, Peter Norvig, John Canny - 2005 - No preview - [More editions](#)

Artificial Intelligence for Games



<https://books.google.ca/books?isbn=0123747317>

Ian Millington, John Funge - 2009 - Preview - [More editions](#)

Creating robust artificial intelligence is one of the greatest challenges for game developers, yet the commercial success of a game is often dependent

Learning algorithms are useful in other tasks

1. **Classification:** Determine which discrete category the example is
2. **Recognizing patterns:** Speech Recognition, facial identity, etc
3. **Recommender Systems:** Noisy data, commercial pay-off (e.g., Amazon, Netflix).
4. **Information retrieval:** Find documents or images with similar content
5. **Computer vision:** detection, segmentation, depth estimation, optical flow, etc

Computer Vision






Figure: Kinect: <https://www.youtube.com/watch?v=op82fDRRqSY>

Computer Vision



[Gatys, Ecker, Bethge. A Neural Algorithm of Artistic Style. Arxiv'15.]

Learning algorithms are useful in other tasks

1. **Classification:** Determine which discrete category the example is
2. **Recognizing patterns:** Speech Recognition, facial identity, etc
3. **Recommender Systems:** Noisy data, commercial pay-off (e.g., Amazon, Netflix).
4. **Information retrieval:** Find documents or images with similar content
5. **Computer vision:** detection, segmentation, depth estimation, optical flow, etc
6. **Robotics:** perception, planning, etc

Autonomous Driving



Flying Robots



Figure: Video: <https://www.youtube.com/watch?v=YQIMGV5vtd4>

Learning algorithms are useful in other tasks

1. **Classification:** Determine which discrete category the example is
2. **Recognizing patterns:** Speech Recognition, facial identity, etc
3. **Recommender Systems:** Noisy data, commercial pay-off (e.g., Amazon, Netflix).
4. **Information retrieval:** Find documents or images with similar content
5. **Computer vision:** detection, segmentation, depth estimation, optical flow, etc
6. **Robotics:** perception, planning, etc
7. **Learning to play games**

Playing Games: Atari



Figure: Video: <https://www.youtube.com/watch?v=V1eYniJ0Rnk>

Playing Games: Super Mario



Figure: Video: https://www.youtube.com/watch?v=wfL4L_14U9A

Learning algorithms are useful in other tasks

1. **Classification:** Determine which discrete category the example is
2. **Recognizing patterns:** Speech Recognition, facial identity, etc
3. **Recommender Systems:** Noisy data, commercial pay-off (e.g., Amazon, Netflix).
4. **Information retrieval:** Find documents or images with similar content
5. **Computer vision:** detection, segmentation, depth estimation, optical flow, etc
6. **Robotics:** perception, planning, etc
7. **Learning to play games**
8. **Recognizing anomalies:** Unusual sequences of credit card transactions, panic situation at an airport

Learning algorithms are useful in other tasks

1. **Classification:** Determine which discrete category the example is
2. **Recognizing patterns:** Speech Recognition, facial identity, etc
3. **Recommender Systems:** Noisy data, commercial pay-off (e.g., Amazon, Netflix).
4. **Information retrieval:** Find documents or images with similar content
5. **Computer vision:** detection, segmentation, depth estimation, optical flow, etc
6. **Robotics:** perception, planning, etc
7. **Learning to play games**
8. **Recognizing anomalies:** Unusual sequences of credit card transactions, panic situation at an airport
9. **Spam filtering, fraud detection:** The enemy adapts so we must adapt too

Learning algorithms are useful in other tasks

1. **Classification:** Determine which discrete category the example is
2. **Recognizing patterns:** Speech Recognition, facial identity, etc
3. **Recommender Systems:** Noisy data, commercial pay-off (e.g., Amazon, Netflix).
4. **Information retrieval:** Find documents or images with similar content
5. **Computer vision:** detection, segmentation, depth estimation, optical flow, etc
6. **Robotics:** perception, planning, etc
7. **Learning to play games**
8. **Recognizing anomalies:** Unusual sequences of credit card transactions, panic situation at an airport
9. **Spam filtering, fraud detection:** The enemy adapts so we must adapt too
10. **Many more!**

Types of learning tasks

- **Supervised**: correct output known for each training example

Types of learning tasks

- **Supervised**: correct output known for each training example
 - ▶ Learn to predict output when given an input vector

Types of learning tasks

- **Supervised**: correct output known for each training example
 - ▶ Learn to predict output when given an input vector
 - ▶ **Classification**: 1-of-N output (speech recognition, object recognition, medical diagnosis)

Types of learning tasks

- **Supervised**: correct output known for each training example
 - ▶ Learn to predict output when given an input vector
 - ▶ **Classification**: 1-of-N output (speech recognition, object recognition, medical diagnosis)
 - ▶ **Regression**: real-valued output (predicting market prices, customer rating)

Types of learning tasks

- **Supervised**: correct output known for each training example
 - ▶ Learn to predict output when given an input vector
 - ▶ **Classification**: 1-of-N output (speech recognition, object recognition, medical diagnosis)
 - ▶ **Regression**: real-valued output (predicting market prices, customer rating)
- **Unsupervised learning**
 - ▶ Create an internal representation of the input, capturing regularities/structure in data

Types of learning tasks

- **Supervised**: correct output known for each training example
 - ▶ Learn to predict output when given an input vector
 - ▶ **Classification**: 1-of-N output (speech recognition, object recognition, medical diagnosis)
 - ▶ **Regression**: real-valued output (predicting market prices, customer rating)
- **Unsupervised learning**
 - ▶ Create an internal representation of the input, capturing regularities/structure in data
 - ▶ Examples: form clusters; extract features

Types of learning tasks

- **Supervised**: correct output known for each training example
 - ▶ Learn to predict output when given an input vector
 - ▶ **Classification**: 1-of-N output (speech recognition, object recognition, medical diagnosis)
 - ▶ **Regression**: real-valued output (predicting market prices, customer rating)
- **Unsupervised learning**
 - ▶ Create an internal representation of the input, capturing regularities/structure in data
 - ▶ Examples: form clusters; extract features
 - ▶ How do we know if a representation is good?

Types of learning tasks

- **Supervised**: correct output known for each training example
 - ▶ Learn to predict output when given an input vector
 - ▶ **Classification**: 1-of-N output (speech recognition, object recognition, medical diagnosis)
 - ▶ **Regression**: real-valued output (predicting market prices, customer rating)
- **Unsupervised learning**
 - ▶ Create an internal representation of the input, capturing regularities/structure in data
 - ▶ Examples: form clusters; extract features
 - ▶ How do we know if a representation is good?
- **Reinforcement learning**
 - ▶ Learn action to maximize payoff
 - ▶ Not much information in a payoff signal
 - ▶ Payoff is often delayed

Machine Learning vs Data Mining

- **Data-mining**: Typically using very simple machine learning techniques on very large databases because computers are too slow to do anything more interesting with ten billion examples

Machine Learning vs Data Mining

- **Data-mining**: Typically using very simple machine learning techniques on very large databases because computers are too slow to do anything more interesting with ten billion examples
- Previously used in a negative sense – misguided statistical procedure of looking for all kinds of relationships in the data until finally find one

Machine Learning vs Data Mining

- **Data-mining**: Typically using very simple machine learning techniques on very large databases because computers are too slow to do anything more interesting with ten billion examples
- Previously used in a negative sense – misguided statistical procedure of looking for all kinds of relationships in the data until finally find one
- Now lines are blurred: many ML problems involve tons of data

Machine Learning vs Data Mining

- **Data-mining**: Typically using very simple machine learning techniques on very large databases because computers are too slow to do anything more interesting with ten billion examples
- Previously used in a negative sense – misguided statistical procedure of looking for all kinds of relationships in the data until finally find one
- Now lines are blurred: many ML problems involve tons of data
- But problems with AI flavor (e.g., recognition, robot navigation) still domain of ML

Machine Learning vs Statistics

- ML uses statistical theory to build models

Machine Learning vs Statistics

- ML uses statistical theory to build models
- A lot of ML is rediscovery of things statisticians already knew; often disguised by differences in terminology

Machine Learning vs Statistics

- ML uses statistical theory to build models
- A lot of ML is rediscovery of things statisticians already knew; often disguised by differences in terminology
- But the emphasis is very different:
 - ▶ **Good piece of statistics:** Clever proof that relatively simple estimation procedure is asymptotically unbiased.

Machine Learning vs Statistics

- ML uses statistical theory to build models
- A lot of ML is rediscovery of things statisticians already knew; often disguised by differences in terminology
- But the emphasis is very different:
 - ▶ **Good piece of statistics:** Clever proof that relatively simple estimation procedure is asymptotically unbiased.
 - ▶ **Good piece of ML:** Demo that a complicated algorithm produces impressive results on a specific task.

Machine Learning vs Statistics

- ML uses statistical theory to build models
- A lot of ML is rediscovery of things statisticians already knew; often disguised by differences in terminology
- But the emphasis is very different:
 - ▶ **Good piece of statistics:** Clever proof that relatively simple estimation procedure is asymptotically unbiased.
 - ▶ **Good piece of ML:** Demo that a complicated algorithm produces impressive results on a specific task.
- Can view ML as applying computational techniques to statistical problems. But go beyond typical statistics problems, with different aims (speed vs. accuracy).

Cultural gap (Tibshirani)

MACHINE LEARNING

- weights
- learning
- generalization
- supervised learning
- unsupervised learning
- large grant: \$1,000,000
- conference location: Snowbird, French Alps

STATISTICS

- parameters
- fitting
- test set performance
- regression/classification
- density estimation, clustering
- large grant: \$50,000
- conference location: Las Vegas in August

Course Survey

Please complete the following survey this week:

https://docs.google.com/forms/d/106xRNnKp87GrDM74tkvOMhMIJmwz271TgWdYb6ZitK0/viewform?usp=send_form

Initial Case Study

- What grade will I get in this course?

Initial Case Study

- What grade will I get in this course?
- Data: entry survey and marks from previous years

Initial Case Study

- What grade will I get in this course?
- Data: entry survey and marks from previous years
- Process the data
 - ▶ Split into **training set**; **test set**
 - ▶ Determine representation of input features; output

Initial Case Study

- What grade will I get in this course?
- Data: entry survey and marks from previous years
- Process the data
 - ▶ Split into **training set**; **test set**
 - ▶ Determine representation of input features; output
- Choose form of model: **linear regression**

Initial Case Study

- What grade will I get in this course?
- Data: entry survey and marks from previous years
- Process the data
 - ▶ Split into **training set**; **test set**
 - ▶ Determine representation of input features; output
- Choose form of model: **linear regression**
- Decide how to evaluate the system's performance: **objective function**

Initial Case Study

- What grade will I get in this course?
- Data: entry survey and marks from previous years
- Process the data
 - ▶ Split into **training set**; **test set**
 - ▶ Determine representation of input features; output
- Choose form of model: **linear regression**
- Decide how to evaluate the system's performance: **objective function**
- Set model parameters to optimize performance

Initial Case Study

- What grade will I get in this course?
- Data: entry survey and marks from previous years
- Process the data
 - ▶ Split into **training set**; **test set**
 - ▶ Determine representation of input features; output
- Choose form of model: **linear regression**
- Decide how to evaluate the system's performance: **objective function**
- Set model parameters to optimize performance
- Evaluate on test set: **generalization**