

# CSC321 Lecture 1: Introduction

Roger Grosse and Nitish Srivastava

January 6, 2015

# Course information

- “Inverted classroom” format
  - Watch Geoff Hinton’s lecture videos for homework
  - During class:
    - More in-depth discussion of the material
    - Question/answer session
    - Working through problems
- Two sections
  - Equivalent content, same assignments and exams
  - Both sections are full, so please attend your own.

# Course information

- Required background
  - Programming experience (recommended)
  - Linear algebra: vector and matrix manipulations
  - Calculus: partial derivatives
  - Probability: Bayes' Rule, Gaussian distribution, mean, variance
- Formal prerequisites:
  - (MAT135H1, MAT136H1)/MAT135Y1/MAT137Y1/MAT157Y1
  - MAT223H1/MAT240H1
  - STA247H1/STA255H1/STA257H1
  - CGPA 3.0, or enrollment in a CSC subject POST
- However, we will be generous about approving waivers
  - We give background pointers each week in case you want to review

# Course information

- Expectations and marking
  - Weekly lecture videos (about 1hr)
  - Quizzes on videos
    - due Monday night **before** classes which cover the content
    - collectively 10% of mark
    - open-book
    - best of 2 attempts
    - aim isn't to evaluate you, but to give you quick feedback and encourage you to think carefully about the material
  - 4 programming assignments (10% each)
    - Python
    - 10-15 lines of code
    - may also involve some mathematical derivations
    - give you a chance to experiment with the algorithms
  - Exams
    - midterm (15%)
    - final (35%)

# Course information

- Textbooks
  - None, but we link to lots of free online resources.
- Tutorials
  - Roughly every other week
  - Programming background; introducing and discussing assignments

# Course information

Course web page: <http://www.cs.toronto.edu/~rgrosse/csc321/>

## CSC321 Winter 2015: Introduction to Neural Networks

---

[Info](#)   [Announcements](#)   [Homework](#)   [Calendar](#)   [Lecture notes](#)

---

### Course information

**Instructors:** Roger Grosse and Nitish Srivastava

**Section 1:** TR 1-2, tutorial R12-1, in BA1200

**Section 2:** T 6-8, tutorial T8-9, in BA1220

Machine learning is a powerful set of techniques that allow computers to learn from data rather than having a human expert program a behavior by hand. Neural networks are a class of machine learning algorithm originally inspired by the brain, but which have recently have seen a lot of success at practical applications. They're at the heart of production systems at companies like Google and Facebook for face recognition, speech-to-text, and language understanding.

[Here](#) are some neat examples of neural net systems developed here at U of T.

Includes detailed course information handout

## Course information

Coursera page (for videos, quizzes, forum):

<https://utoronto.coursera.org/CSC321-003>

- Sign in with your UTorID
- **Not** the same as the version accessible through the main Coursera site
- **Not** the 2013 or 2014 versions — make sure it's 2015

# What is machine learning?

- For many problems, it's difficult to program the correct behavior by hand
  - recommendation systems
  - face recognition
  - spam classification
- We want to program a system which will learn from data, or from experience



# What is machine learning?

- Types of machine learning
  - **Supervised learning:** have labeled examples of the correct behavior
  - **Unsupervised learning:** no labeled examples – instead, looking for interesting patterns in the data
  - **Reinforcement learning:** learning system receives a reward signal, tries to learn to maximize the reward signal

# Supervised learning examples


**Supervised learning:** have labeled examples of the correct behavior

e.g. Handwritten digit classification with the MNIST dataset

- task: given an image of a digit, predict the digit class
- 70,000 images of handwritten digits labeled by humans
- 60,000 used to train the classifier, 10,000 to test its performance
- This dataset is the “fruit fly” of neural net research
- Current best algorithm has only 0.23% error rate!

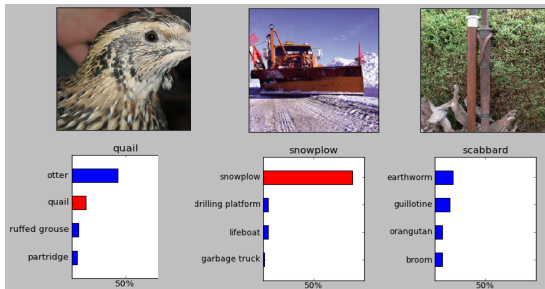
# Supervised learning examples

What makes a "2"?



# Supervised learning examples

## Object recognition



(Krizhevsky and Hinton, 2012)

ImageNet dataset: thousands of categories, millions of labeled images

Lots of variability in viewpoint, lighting, etc.

# Supervised learning examples

## Caption generation



### TAGS:

frisbees frisbee pushups golfers kickball

### Nearest Neighbor Sentence:

- several people that are playing in a frisbee game .

### Top-5 Generated:

- a group of girls are playing a game of frisbee .
- a group of girls are playing a soccer game .
- a group of girls playing on a soccer game .
- a group of people playing a game of frisbee .
- the young people are playing a game of frisbee .

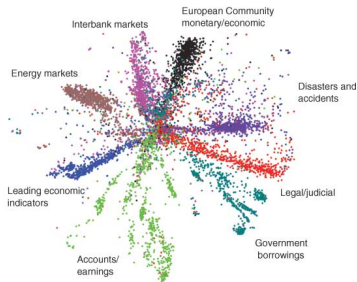
Given: dataset of Flickr images with captions

More examples at <http://deeplearning.cs.toronto.edu/i2t>

# Unsupervised learning examples

**Unsupervised learning:** no labeled examples – instead, looking for interesting patterns in the data

E.g. visualization of documents; algorithm was given 800,000 newswire stories, and learned to represent these documents as points in two-dimensional space



Colors are based on human labels, but these weren't given to the algorithm

# Unsupervised learning examples

## Automatic mouse tracking

- When biologists do behavioral genetics research on mice, it's very time consuming for a person to sit and label everything a mouse does
- The Datta lab at Harvard is building a system for automatically tracking mouse behaviors
- Goal: show the researchers a summary of how much time different mice spend on various behaviors, so they can determine the effects of the genetic manipulations
- One of the major challenges is that we don't know the right "vocabulary" for describing the behaviors — clustering the observations into meaningful groups is an unsupervised learning task
- **(video)**

# Reinforcement learning examples

**Reinforcement learning:** learning system receives a reward signal, tries to learn to maximize the reward signal


e.g. DeepMind's system which learned to play Atari games

- given raw screen as input, plus the score as the reward signal
- <https://www.youtube.com/watch?v=EfGD2qveGdQ>



# What are neural networks?

- Most of the biological details aren't essential, so we use vastly simplified models of neurons.
- While neural nets originally drew inspiration from the brain, nowadays we mostly think about math, statistics, etc.



$$y = g \left( b + \sum_i x_i w_i \right)$$

Diagram illustrating the mathematical representation of a neuron's output. The equation is  $y = g \left( b + \sum_i x_i w_i \right)$ . Labels with arrows point to components: "output" (pink arrow to  $y$ ), "nonlinearity" (red arrow to  $g$ ), "bias" (blue arrow to  $b$ ), "i'th weight" (blue arrow to  $w_i$ ), and "i'th input" (green arrow to  $x_i$ ).

- Neural networks are collections of thousands (or millions) of these simple processing units that together perform useful computations.


# What are neural networks?

- Some near-synonyms for neural networks
  - “Deep learning”
    - Emphasizes that the algorithms often involve hierarchies with many stages of processing

# “Deep learning”

Deep learning: many layers (stages) of processing

E.g. this network which recognizes objects in images:




(Krizhevsky et al., 2012)


Each of the boxes consists of many neurons similar to the one on the previous slide!

# “Deep learning”

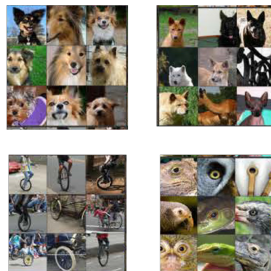
Here are the image regions that most strongly activate various neurons at different layers of the network. (Zeiler and Fergus, 2014)



Layer 1



Layer 2



Layer 5

Higher layers capture more abstract semantic information.

# What are neural networks?

- Some near-synonyms for neural networks
  - “Deep learning”
    - Emphasizes that the algorithms often involve hierarchies with many stages of processing
  - “Representation learning”
    - The algorithms typically map the raw data into some other space which makes the relationships between different things more explicit

# What is a representation?

- In your past computer science courses, you may have learned about various data structures for representing words, documents, etc.
  - arrays of characters
  - dictionaries of word counts
  - tries (i.e. trees of prefixes)

# What is a representation?

- In your past computer science courses, you may have learned about various data structures for representing words, documents, etc.
  - arrays of characters
  - dictionaries of word counts
  - tries (i.e. trees of prefixes)
- How you represent your data determines what questions are easy to answer.
  - E.g. a dict of word counts is good for questions like “What is the most common word in *Hamlet*?”


# What is a representation?

- In your past computer science courses, you may have learned about various data structures for representing words, documents, etc.
  - arrays of characters
  - dictionaries of word counts
  - tries (i.e. trees of prefixes)
- How you represent your data determines what questions are easy to answer.
  - E.g. a dict of word counts is good for questions like “What is the most common word in *Hamlet*?”
- Simple data structures aren't enough to do higher-level semantic reasoning, e.g.
  - Did this reviewer like the book?
  - Alice liked *Harry Potter*. Will she like *The Hunger Games*?
  - Translate this book into French.



# What is a representation?

We will discuss some algorithms for representing words as vectors.



# What is a representation?

- In a good representation, mathematical relationships between the vectors should encode semantic relationships between the things we care about. For instance,
  - Measure similarity between words using the dot product of their vectors (or dissimilarity using Euclidean distance)
  - Represent a web page with the average of its word vectors
  - Complete analogies like “Paris is to France as London is to \_\_\_\_\_” by doing arithmetic on word vectors
- It’s very hard to construct representations like these by hand, so we need to learn them from data
  - This is a big part of what neural nets do, whether it’s supervised, unsupervised, or reinforcement learning!

# Outline of course

- Weeks 1-8: supervised learning
- Weeks 9-11: unsupervised learning
- Week 12: reinforcement learning (maybe)
- Throughout the semester: learning representations of images and text

# Reminders

- No tutorial this week
- Next lecture: linear regression, a simple learning algorithm
- Check out the web site for a detailed schedule
- Reminder: first two quizzes due 11:59pm next Monday