APS360 Fundamentals of AI

Lisa Zhang

Lecture 1; Jan 7, 2019

Agenda

- Introduction
- Motivation
- Logistics

Welcome to APS360!

Introduction

- **Instructor**: Lisa Zhang
- **Email**: lczhang@cs.toronto.edu
  - Please prefix email subject with 'APS360'
- **Office hours**: TBD (Thursday afternoon)

About your instructor

Before I started teaching, I was...

- a masters student doing research in Machine Learning
- a senior data scientist at an advertising technology company
- a startup founder of a data visualization company
- a software developer intern in various Silicon Valley companies, e.g. Facebook, ContextLogic (Wish)

I studied...

- machine learning at UofT (supervised by Prof. Richard Zemel, Prof. Raquel Urtasun)
- pure math at UWaterloo

Ask me about anything outside of class, or empty office hours!

About you

- What is your name?
- What is your area of study?
- Why are you here?

Survey: demographics

- **Year of study**:
  - 75% - 3rd years
  - 20% - 4th years
- **Area of study**:
  - Computuer engineering
  - Industrial engineering
  - Mechanical engineering
  - others

Why did you take this course?

- AI Minor / Certificate Requirement
- For fun, sounds cool, gain knowledge
- Want to work in the field, career prospects, useful
- Interested in AI and machine learning

Previous ML and Neural Networks courses?

- **Most people have not**
- Some took Coursera courses
- Taking ECE421 concurrently

Interest in Machine Learning

- Application to another field: ~80%
- Becoming a data scientist: ~45%
- Machine Learning Research: ~30%

What do you know about AI?

What is the difference between:

- Artificial Intelligence,
- Machine Learning, and
- Deep Learning?

AI vs ML vs DL

Artificial Intelligence: Create intelligent machines that work and act like humans.

Machine Learning: Find an algorithm that automatically learns from example data.

Deep Learning: Using deep neural networks to automatically learn from example data.

Relationship

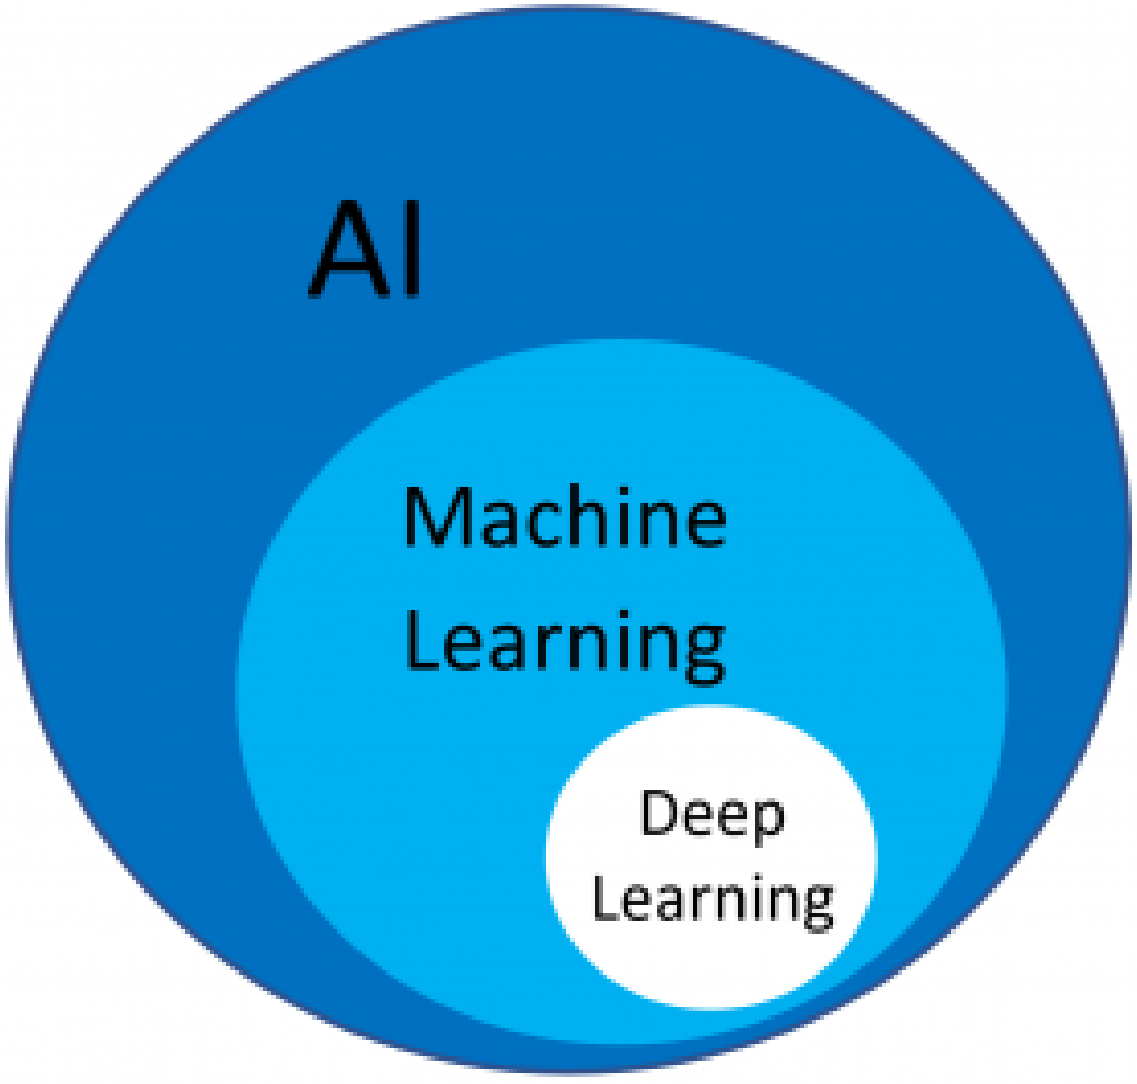

AI Timeline

- 1950's
  - The term "Artificial Intelligence" was first adopted.
  - John McCarthy developed LISP programming language.
  - Alan Turing proposes the Turing Test.
  - Arthur Samuel writes the first game-playing program (checkers)
- 1960's
  - Unimate: first industrial robot on a GM assembly line.
  - Joseph Weizenbaum developes chatbot Eliza
- 1970's - 1980's: "AI Winter"
- 1990's
  - Gerry Tesauro builds a backgammon program using reinforcement learning
  - Deep Blue chess machine defeats chess champion
- 2000's
  - iRobot's Roomba vacuums floors
- 2010's
  - ImageNet (Convolutional Neural Networks)
  - Siri, Watson, Alexa
  - Google Deepmind's AlphaGo defeats Go champion

ARTIFICIAL INTELLIGENCE

Artificial Intelligence captures the imagination of the world.

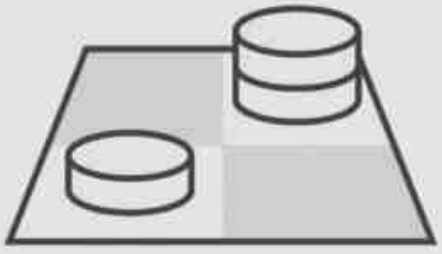

MACHINE LEARNING

Machine learning starts to gain traction.

DEEP LEARNING

Deep learning catapults the industry.

Turing Test Devised
1950

ELIZA
1964 - 1966

Edward Shortliffe writes MYCIN, an Expert or Rule based System, to classify blood disease
1970s

IBM Deep Blue defeats Grand Master Garry Kasparov in chess
1996

ImageNet Feeds Deep Learning
2009

AlphaGo defeats Go champion Lee Sedol
2016

1950s 1960s 1970s 1980s 1990s 2000s 2010s

Why machine learning?

For many problems, it is difficult to program the correct behavior.

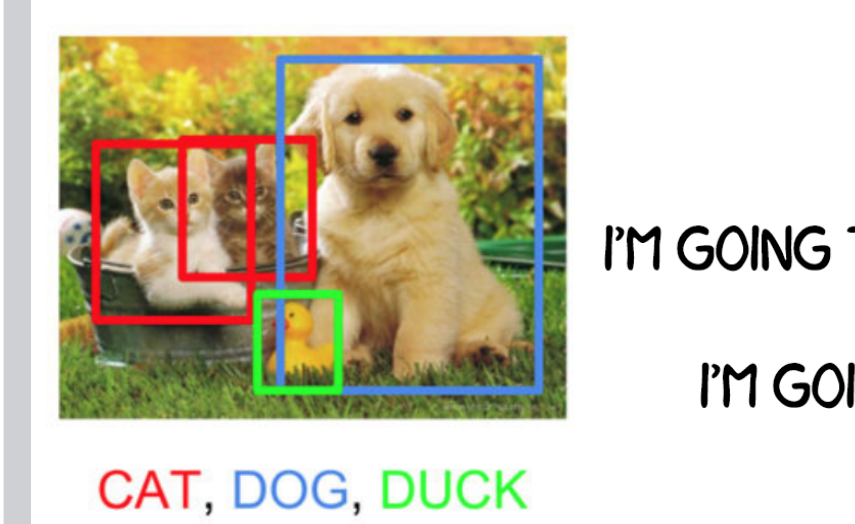

Deep Learning Successes: Object Detection

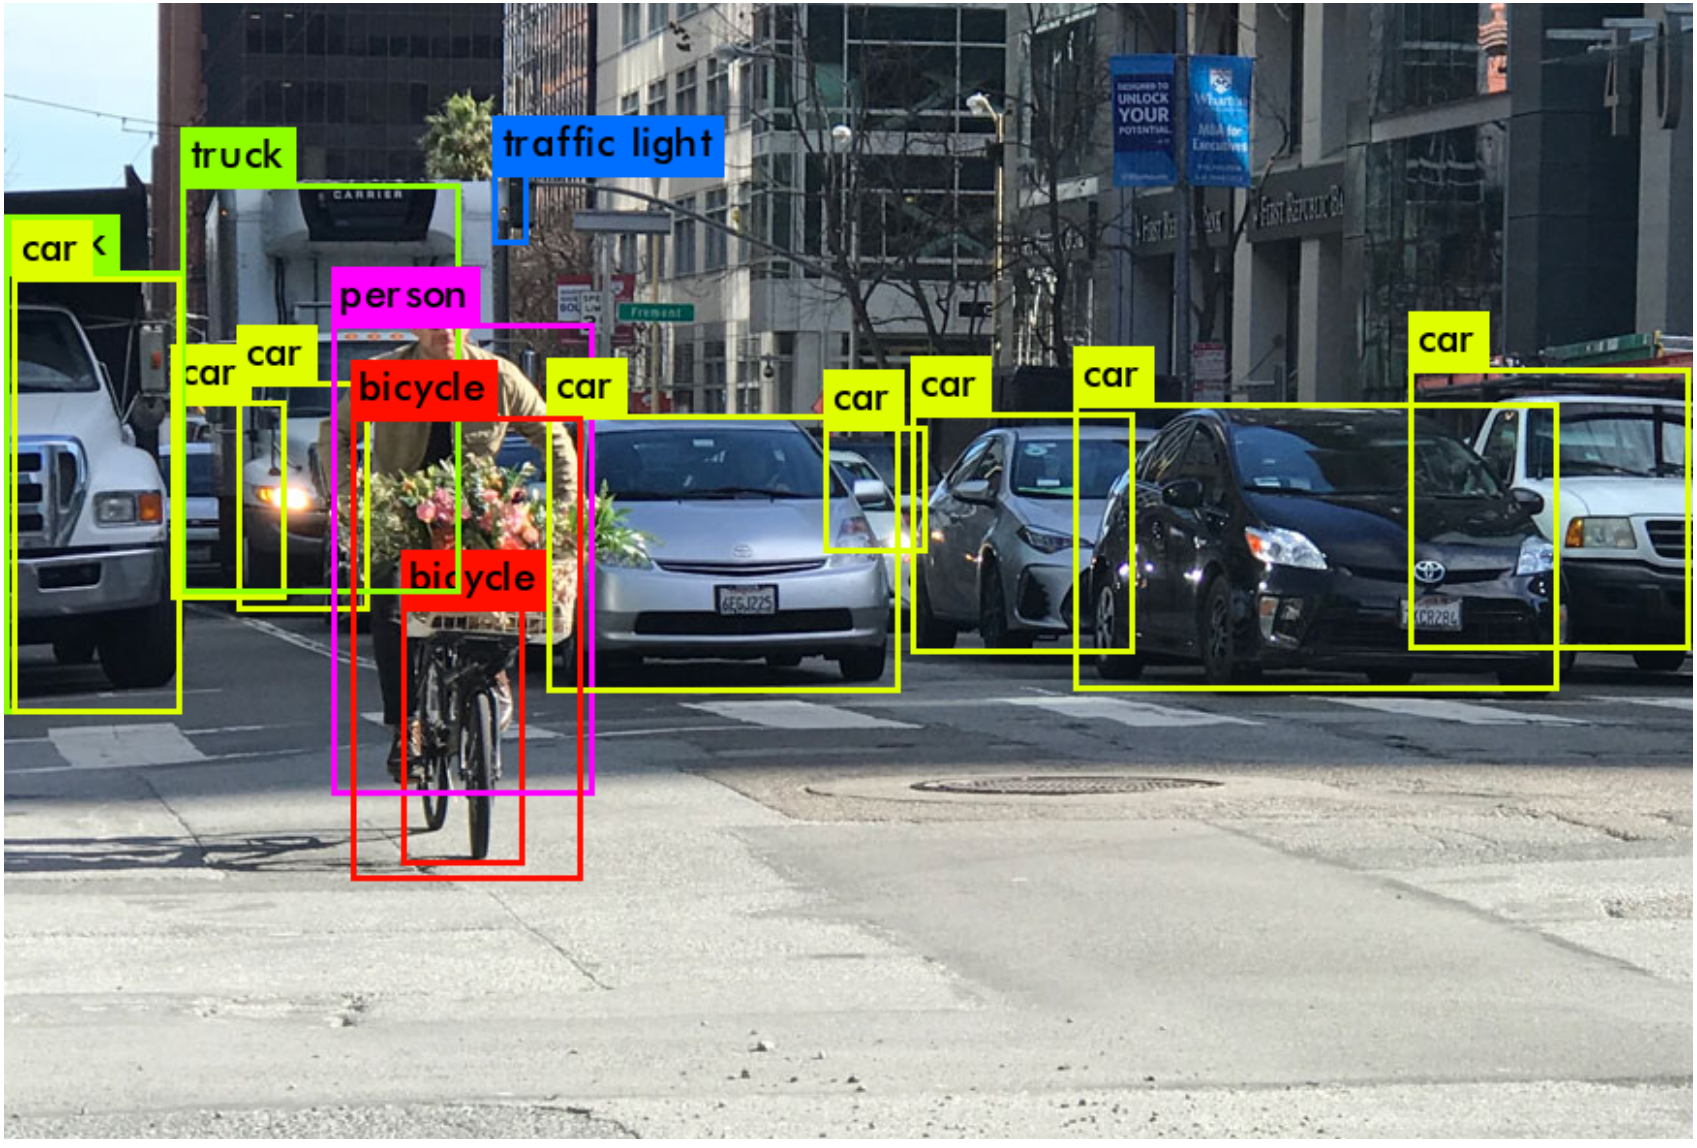

Deep Learning Successes: AlphaGo

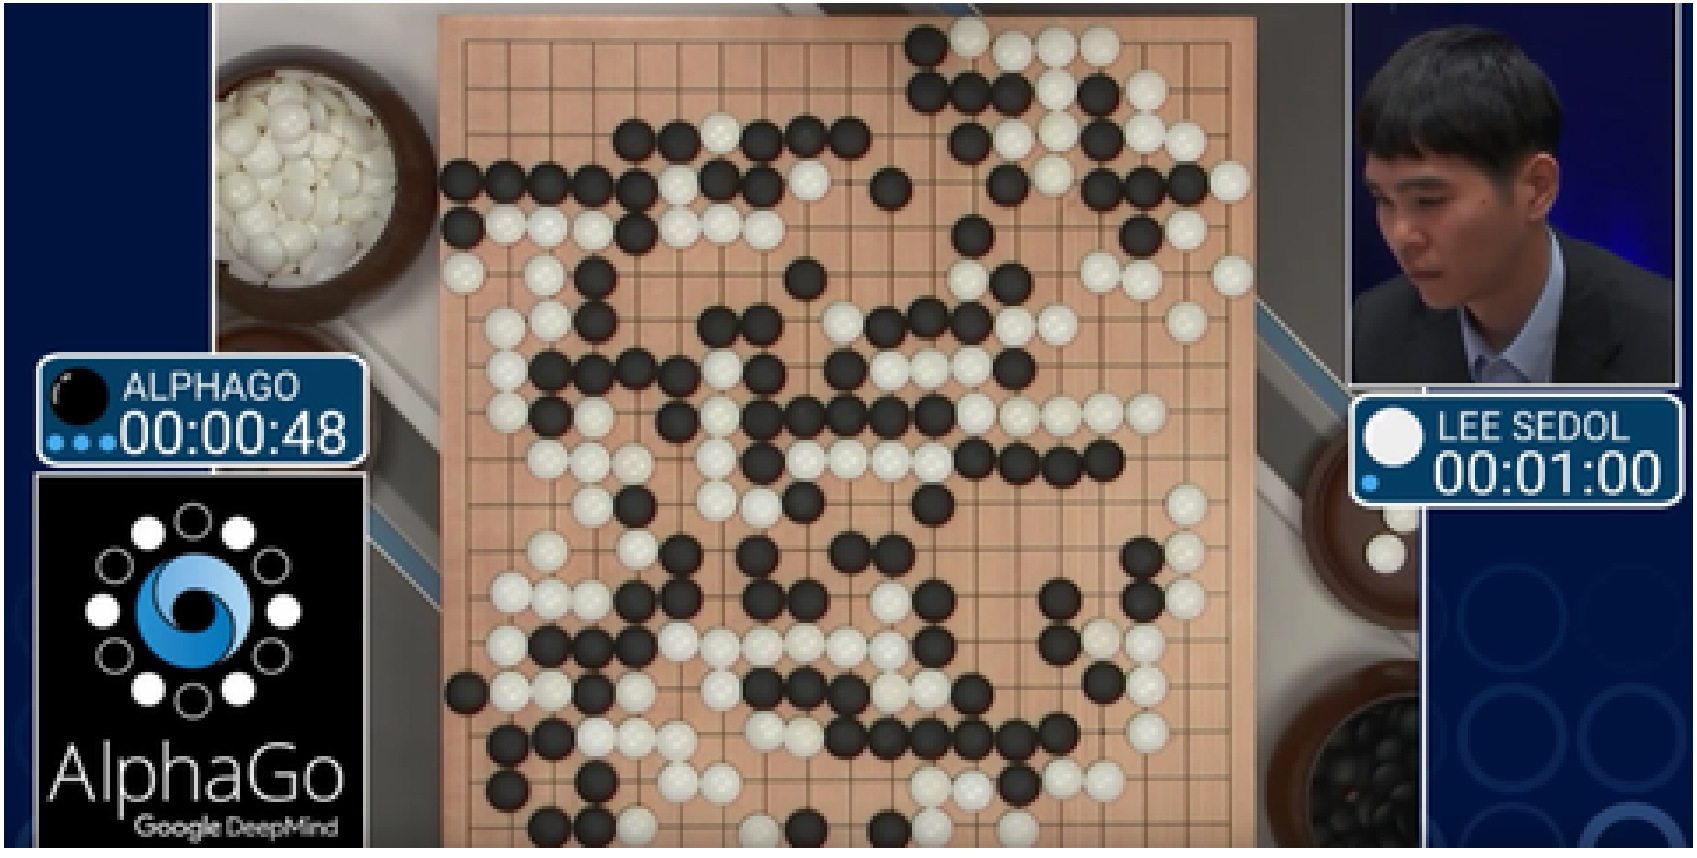

Deep Learning Successes: Style Transfer

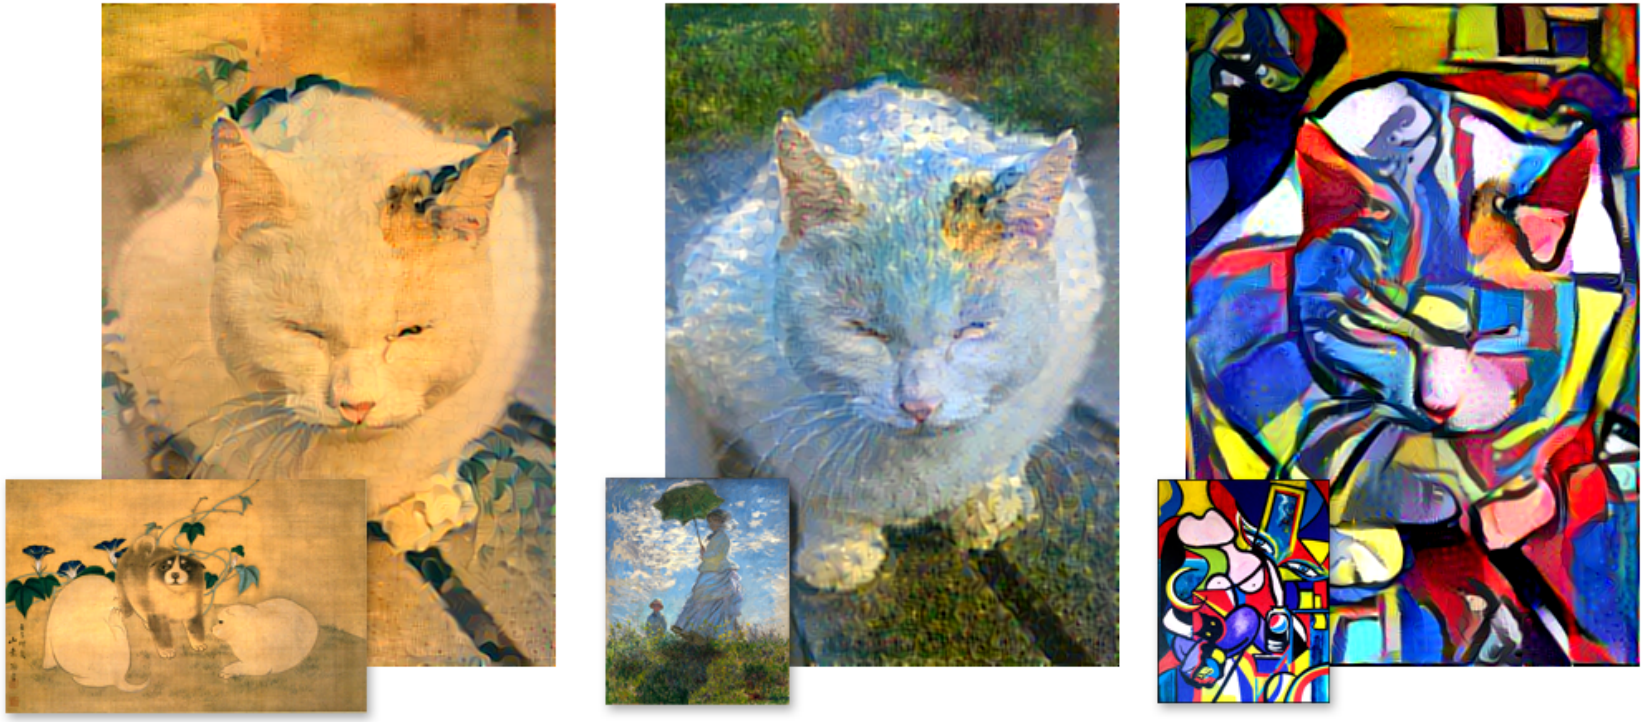

Deep Learning Successes: Style Transfer

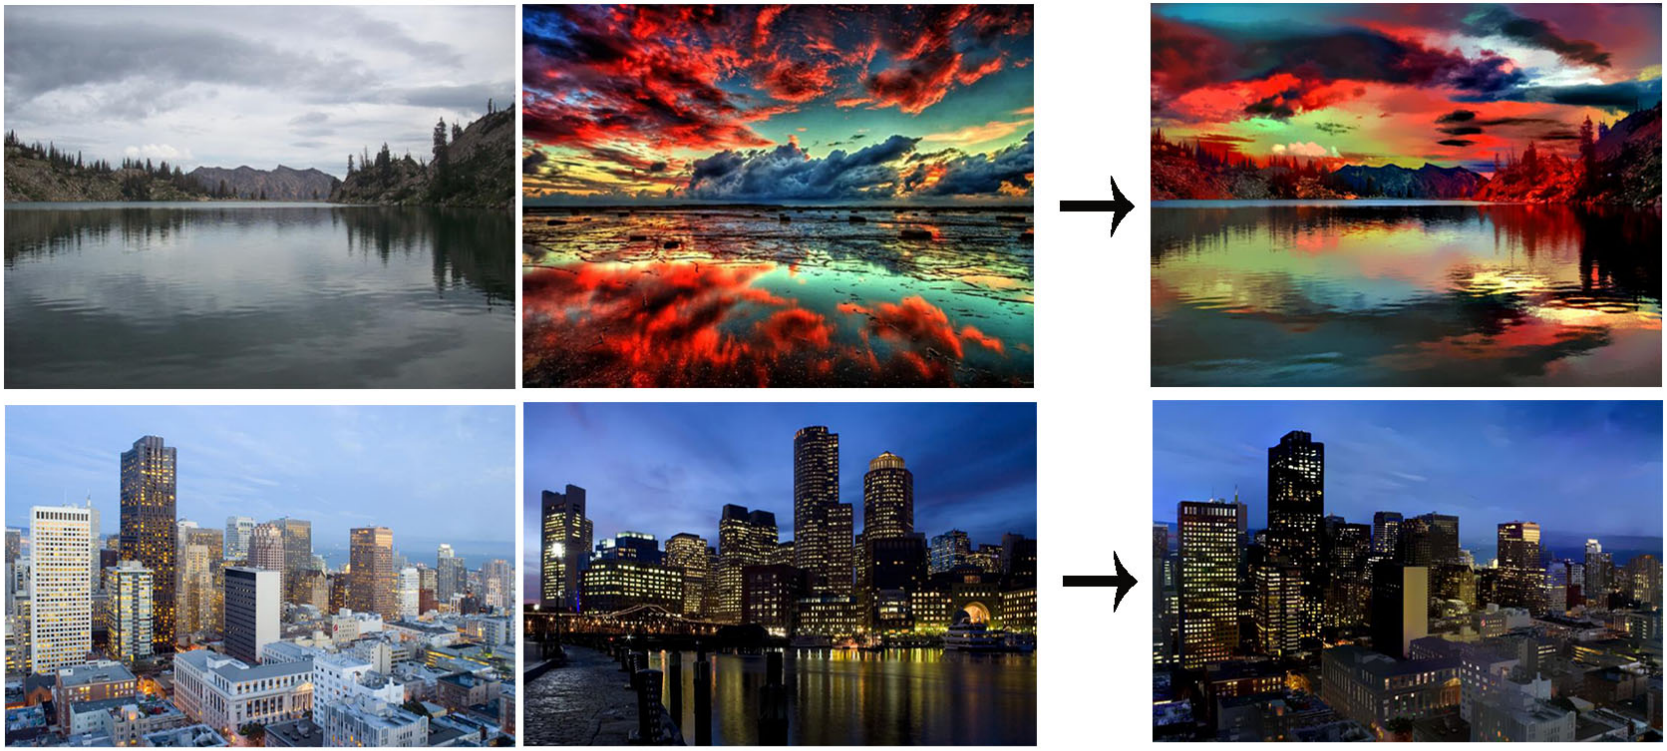

Original Photo

Example Photo

Result

Deep Learning Caveats: Interpretability

Figure 1: from https://xkcd.com/1838/

Deep Learning Caveats: Adversarial Examples

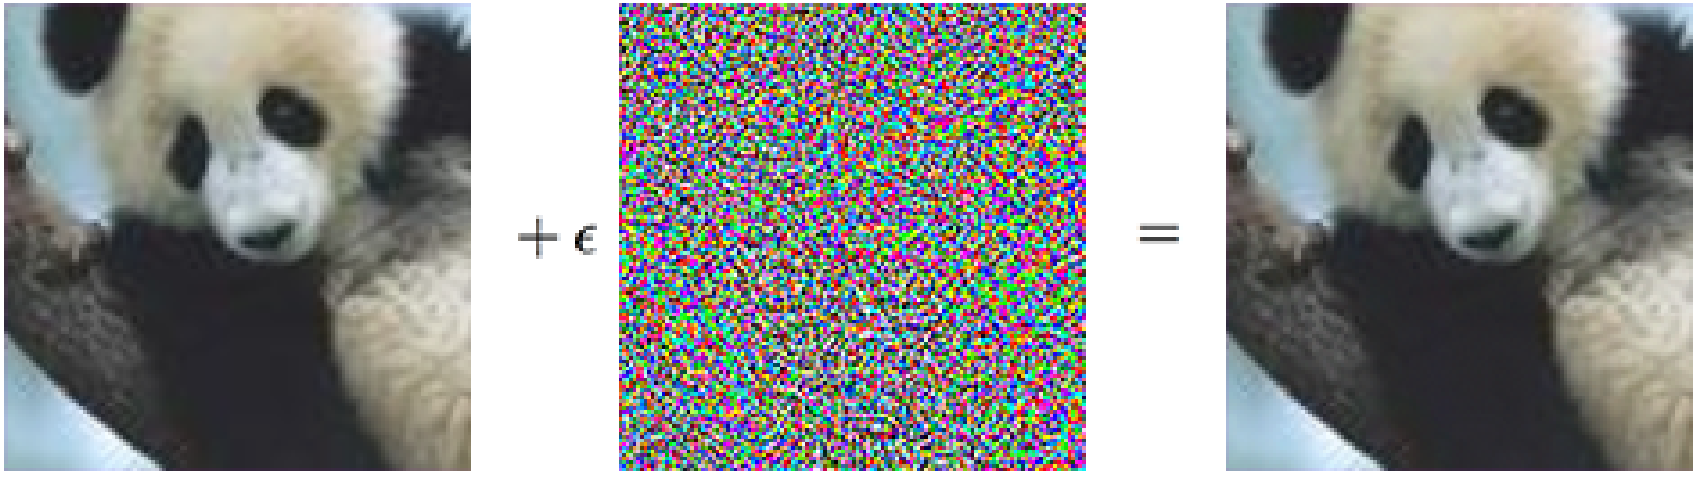

"panda"

57.7% confidence

"gibbon"

99.3% confidence

Deep Learning Caveats: Fairness

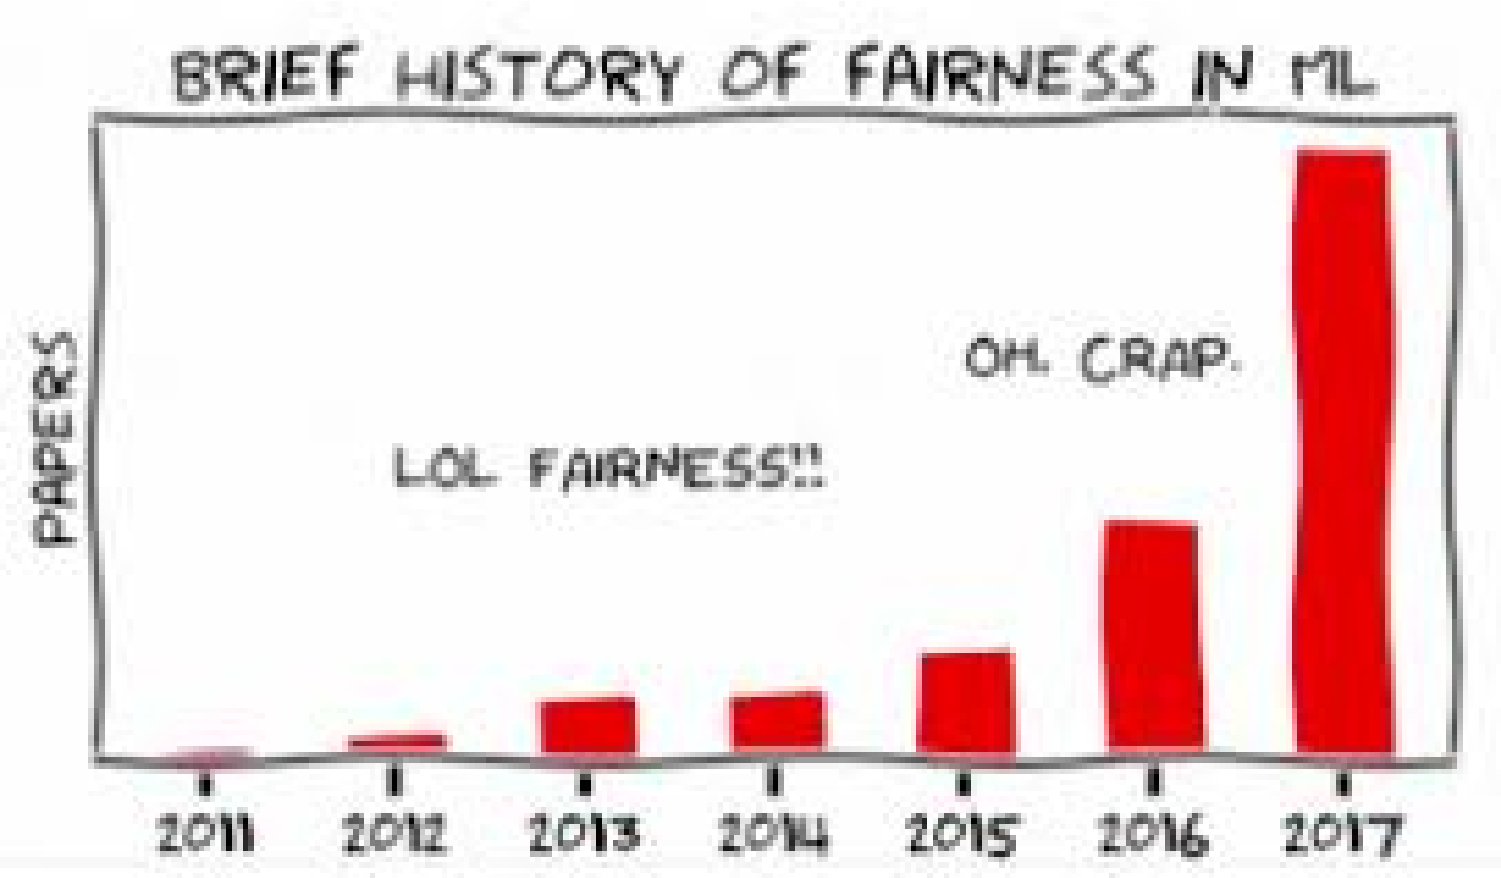

Example

The U.S. military built an AI tool to find suitable combat personnel but had to shut it down because it was discriminating against women

News

Course Coverage

We will focus exclusively on neural networks and deep learning.

Course Philosophy

- Top-down approach
  - Learn by doing
  - Explains the entire system first
  - Details in future courses
- We will introduce very little math
- Focus on implementation and software skills
- Focus on communication skills

Course Website

https://www.cs.toronto.edu/~lczhang/360/

Course Components

- **Lectures**: Monday (1 hr), Thursday (2 hrs)
- **Tutorials**: Monday (1 hr), lead by a TA
- **Assignments**: Weekly assignments in the first half of the course
- **Project**: Implementation project in the second half of the course
- **Readings**:
  - No textbooks
  - Readings are posted weekly.
- Any material covered in lectures / tutorials / readings is fair game for the midterm, and exam.

Teaching Assistants

- **Head Teaching Assistant**:
  - Hojjat Salehinejad
- **Teaching Assistants**: (in alphabetical order)
  - Andrew Jung
  - Bibin Sebastian
  - Kingsley Chang

Grade Breakdown

Tentative grade breakdown:

- **Assignments**: 20% (4% each)
- **Project**: 30%
- **Midterm**: 15%
- **Exam**: 35%

Assignments

- One per week in the first half of the course
- Done **individually** – must be your own work
- Due Sunday nights 9pm
- First one is due **end of week 2** (not week 1)
  - but you should do it sooner, ideally **by Thursday**

Late Policy:

- Penalty of 20% (a late assignment with a grade of 78% will count as 58%)
- Accepted up to 24 hours past the deadline.

Software

- Python 3.6
- NumPy
- PyTorch
- Jupyter Notebooks

All assignment handouts will be Jupyter notebooks

Course Project

Work in a group of 3 to build a useful machine learning system.

Everyone must contribute to all parts of the project to earn a grade.

- **Project Proposal**: 3%
- **Progress Meeting with TA Mentor**: 3%
- **Progress Report**: 4%
- **Presentation**: 10%
- **Project Report**: 10%

Midterm

- Week 7 Thursday Lecture (after reading week)
- Length: 2 hours
- Location: TBD
- No aids permitted

Schedule (Weeks 1-6)

Week	Monday	Thursday
1	Introduction	Pigeons to Neural Networks
2	Terminology	Neural Network Training
3	Define Neural Networks	Convolutional Networks (CNN)
4	Training CNN	CNN Architecture
5	Autoencoders	Unsupervised Learning; word2vec
6	Language Models	Recurrent Neural Networks

Schedule (Weeks 7-12)

Week	Monday	Thursday
7	Distillation	Midterm Test
8	Guest Lecture (TBD)	Generative Adversarial Networks
9	Guest Lecture (TBD)	Reinforcement Learning
10	Guest Lecture (TBD)	Transfer Learning
11	Fairness in AI	Ethics of AI
12	Presentations	Presentations

Course Notes

- The notes that I write will also be Jupyter Notebooks
- Please email me if you have any corrections
- Tell me if you use them, whether it's worth my time to write them

Questions?

What to do this week

- Set up your development environment
- Go to the tutorial (i.e. stay in this room. . .)
- Do assignment 1: even though it's not due for a while, it will help you understand Thursday's lecture