

Solving World Problems

Learning One Example at a Time

Nate Derbinsky

Associate Teaching Professor Director of Teaching Faculty Khoury College of Computer Sciences



My Path to Northeastern

| bit X solutions | 1998-2009 | BitX Solutions, Inc. Founder & President {.gov .edu .org .com} x {desktop web mobile} |
|---------------------------|-----------|--|
| NC STATE UNIVERSITY | 2002-2006 | NC State. BS Computer ScienceDatabases |
| UNIVERSITY OF MICHIGAN | 2006-2012 | U of Michigan. MS/PhD Comp Sci and EngAl+Databases |
| Disnep Research | 2012-2014 | Disney Research. Postdoctoral Associate Machine Learning, Optimization, Robotics |
| | 2014-2017 | Wentworth. Assistant ProfessorCS Education, Service Learning |



Research Interests

Cognitive Systems

Scalable Optimization







AI Applications/Education

Online ML

+Knowledge: 1 Core, World Record Density

ber of Circles

□ TWA: 1 Core △ TWA: 2 Cores × TWA: 4 Cores ∓ TWA: 8 Cores

+Knowledge: 1 Core







Solving World Problems: Learning One Example at a Time

March 18, 2019

Teaching

K-12/ICT-D











UG/Grad

- CS1/2, *DS2000* – OOP, Foundations
- Databases, Web, SE
- AI, Machine Learning
- HTMAA
 - RPi, Arduino



Problem: Signaling Networks



To <u>treat/cure disease</u>, we need to understand the <u>chains of protein</u> <u>interactions</u> that determine <u>how cells</u> <u>process signals</u> from their environment





Problem: Robotic Interaction



To <u>safely operate</u> in distant/dangerous terrain, we need <u>robotic</u> <u>teammates</u> that can <u>autonomously</u> perceive, understand, interact with, and <u>manipulate</u> <u>their environment</u>





Problem: Winning on the Merits



To enact the <u>best</u> <u>policies</u>, we need to understand what makes for a <u>strong argument</u>





Solving World Problems

The common link in solving these and many other complex problems is the potential for using **Machine Learning**, which exploits...

- Big Data: <u>examples</u> from the world
- Cloud Computation: cheap, fast processing
- Algorithms: automatically <u>uncovering value</u>



Machine Learning

- 1. What is it?
 - Why you should care!
- 2. How does it work?
 - Learn your first ML technique!



http://ai.berkeley.edu



What is Machine Learning?

Computer programs that can improve performance with experience





But Wait...

Why Learn?

Many complex tasks are <u>hard to describe</u>, but <u>easy</u> <u>to learn</u> from experience

Why Now?

Data sources and powerful computing are increasingly cheap and plentiful







Natural Language Processing (NLP)



Modern NLP algorithms are typically based on statistical ML







Applications

- Summarization
- Machine Translation
- Speech Processing
- Sentiment Analysis



Computer Vision

Methods for acquiring, processing, analyzing, and understanding images

Applications

- Image search
- Facial recognition
- Object tracking
- Image restoration











Games, Robotics, Medicine, Ads, ...













Solving World Problems: Learning One Example at a Time

March 18, 2019

Jobs!

| Position | Salary |
|-------------------|-----------|
| Data Scientist | \$117,345 |
| Software Engineer | \$103,035 |

| E C C Che New york Eimes | Tech Giants Are Paying Huge Salaries for Scarce A.I. Talent |
|--|--|
| A.I. Researchers Are Making More Than \$1 Million, Even at a Nonprofit By CADE METZ APRIL 19, 2018 | Nearly all big tech compa intelligence project, and t experts millions of dollarTypical A.I. specialists, including both Ph.D.s fresh out of school and people with less education and just a few years of experience, can be paid from \$300,000 to \$500,000 a year or more in salary and company stock, according to nine people who work for major tech companies or have entertained job offers from them. All of them requested anonymity because they did not want to damage their professional prospects. |

"Software Is Eating the World, but AI Is Going to Eat Software" - Jensen Huang (CEO, NVIDIA)

*glassdoor.com, USA National Avg as of March 12, 2019



(Supervised) Machine Learners



α





β



Training Set



γ

Testing Set



?















Machine Learning in Action



https://teachablemachine.withgoogle.com



ML Terminology



example, instance Unit of input

Composed of *features* (or *attributes*)

- In this case, we could represent each digit via raw pixels: 28x28=784-pixel *vector* of greyscale values [0-255]
 - Dimensionality: number of features per instance (|vector|)
- But other *data representations* are possible, and might be advantageous





 In general, the problem of *feature* selection is challenging



Your First ML Technique: kNN

Training

•

Store all examples

Testing

- Find the <u>k nearest</u> <u>neighbors</u> to input
- Vote on output





2D Multiclass Classification





Northeastern University

Honors Welcome Day · March 2019 · Nate Derbinsky

Many Approaches









$SSE = Y^{\intercal}Y - 2Y^{\intercal}XB + B^{\intercal}X^{\intercal}XB$

$\frac{\partial \text{SSE}}{\partial B} = -2X^{\intercal}Y + 2X^{\intercal}XB$

$$0 = -2X^{\mathsf{T}}Y + 2X^{\mathsf{T}}XB$$
$$-2X^{\mathsf{T}}XB = -2X^{\mathsf{T}}Y$$
$$X^{\mathsf{T}}XB = X^{\mathsf{T}}Y$$
$$B = (X^{\mathsf{T}}X)^{-1}X^{\mathsf{T}}Y$$



Solving World Problems: Learning One Example at a Time

4

3

One (of many) Challenges



https://xkcd.com/1838/

What's a "Good" (Unsupervised) Answer?





Northeastern University

Did I Learn Correctly?



Northeastern University

Did I Learn Ethically and Safely?





Pop Quiz!

- Given a dataset of past credit-card transactions (known to be fraudulent or not), build a system to identify future fraud
- 2. If we assume incoming CS1 students are bi-modal, but normally distributed, find the average grades of the two groups
- 3. Build an Atari system that learns game-winning techniques via actually playing and adjusting actions based upon score changes

Reinforcement



March 18, 2019

Supervised

Unsupervised

Learning Signaling Networks



- Using NLP to aggregate and analyze scientific findings
- Given data, learn
 network structure









Learning Robotic Interaction



- Large-scale robotic grasp training
- Inter-planetary!





Learning to Win on the Merits



- Understanding content/style vs strength
- Making new arguments from past experience and data







Opportunities for (Machine) Learning

- 1. Courses
- 2. Research
- 3. Applications



http://ai.berkeley.edu



Northeastern University

Honors Welcome Day · March 2019 · Nate Derbinsky

Fusing Disciplines





Computing @ Northeastern





Some Offerings...

- Artificial Intelligence
- Supervised Machine Learning and Learning Theory
- Unsupervised Machine Learning and Data Mining
- Reinforcement Learning
- Natural Language Processing
- Advanced Machine Learning
- Information Presentation and Visualization
- Robotic Science and Systems
- Pattern Recognition and Computer Vision



Honors Welcome Day · March 2019 · Nate Derbinsky

What Problem Will YOU Learn to Solve?



Thank You :) Questions?



What Problem Will YOU Learn to Solve?



Nate Derbinsky Associate Teaching Professor Director of Teaching Faculty Khoury College of Computer Sciences

> n.derbinsky@northeastern.edu https://derbinsky.info

