

DSC190: Machine Learning with Few Labels

Overview

Zhiting Hu

Lecture 1, September 23, 2021

UC San Diego

HALICIOĞLU DATA SCIENCE INSTITUTE

Logistics

- Class webpage: <http://zhiting.ucsd.edu/teaching/dsc190fall2021>

DSC190-Fall2021

Logistics Lectures Homework Project



Machine Learning with Few Labels

DSC 190 • Fall 2021 • UC San Diego

Machine learning is about computational methods that enable machines to learn concepts from experience. Many of the successful results of machine learning rely on supervised learning with massive amount of data labels. However, in many real problems we do not have enough labeled data, but instead have access to other forms of experience, such as structured knowledge, constraints, feedback signals from environment, auxiliary models from related tasks, etc. This course focuses on those learning settings with few labels, where one has to go beyond supervised learning and use other learning methods. This course is designed to give students a holistic understanding of related problems and methodologies (such as zero/few-shot learning, self/weakly-supervised learning, transfer learning, meta-learning, reinforcement learning, adversarial learning, knowledge constrained learning, panoramic learning), different possible perspectives of formulating the same problems, the underlying connections between the diversity of algorithms, and open questions in the field. Students will read, present, and discuss papers, and complete course projects.

Logistics



Instructor: **Zhiting Hu**

Email: zhh019@ucsd.edu

Office hours: Tuesday 3:00-4:00pm

Location: SDSC E247



TA: **Meng Song**

Email: mes050@eng.ucsd.edu

Office hours: Wednesday 2:30-3:30pm

Location: CSE 4109

- Canvas
- Discussion forum: Piazza
- Homework & writeup submission: Gradescope

Logistics: grading

- 2 Homework assignments (30% of grade)
- Paper presentation (20%)
- Course project (46%)
- Participation (4%)

Logistics: grading

- 2 Homework assignments (30% of grade)
 - Theory exercises, implementation exercises
 - 3 total late days without penalty
- Paper presentation (20%)
- Course project (46%)
- Participation (4%)

Logistics: grading

- 2 Homework assignments (30% of grade)
- Paper presentation (20%)
 - Each student will give an oral presentation on a research paper
 - 10 mins = 8 mins presentation + 2 mins QA
 - Discuss both strengths and limitations of the paper
 - Sign up in a google sheet (TBA)
 - Starting October 26
- Course project (46%)
- Participation (4%)

Logistics: grading

- 2 Homework assignments (30% of grade)
- Paper presentation (20%)
- Course project (46%)
 - 3 or 4-member team to be formed and sign up in a google sheet (TBA)
 - Designed to be as similar as possible to researching and writing a conference-style paper:
 - Due to tight timeline, fine to use synthetic/toy data for proof-of-concept experiments + explanation of theory/intuition of why your approach is likely to work
 - **Proposal** : 2 pages excluding references (10%) -- **Due Oct 12**
 - Overview of project idea, literature review, potential datasets and evaluation, milestones
 - **Midway Report** : 4-5 pages excluding references (20%)
 - **Presentation** : oral presentation, 15-20mins (20%)
 - **Final Report** : 6-8 pages excluding references (50%)


Logistics: grading

- 2 Homework assignments (30% of grade)
- Paper presentation (20%)
- Course project (46%)
- Participation (4%)
 - Contribution to discussion on Piazza
 - Complete mid-quarter evaluation
 - Any constructive suggestions



Machine Learning

- Computational methods that enable machines to learn concepts and improve performance from **experiences**.

Experiences of all kinds



Type-2 diabetes is 90% more common than type-1



Data examples

Rules/Constraints


Knowledge graphs

Rewards

Auxiliary agents



Adversaries




Master classes

...


And all combinations thereof

Experiences of all kinds



Data examples


Type-2



Auxiliary agents




Adversaries



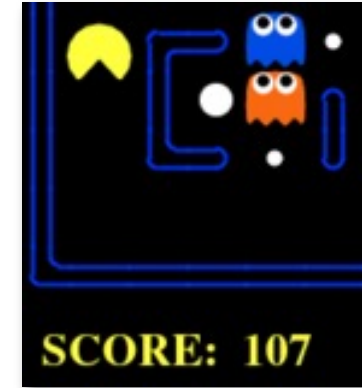

Master classes

ations thereof

Experiences of all kinds



Type-2 diabetes is 90% more common than type-1



Data examples

Rules/Constraints


Knowledge graphs

Rewards

Auxiliary agents



Adversaries



Master classes

...

And all combinations thereof

Experiences: (massive) data examples



Image classification



Machine translation



Language modeling
(BERT, GPT-2, **GPT-3**, ...)

45TB of text data: CommonCrawl, WebText, Wikipedia, corpus of books, ...

Experiences: (massive) data examples

TECH ARTIFICIAL INTELLIGENCE

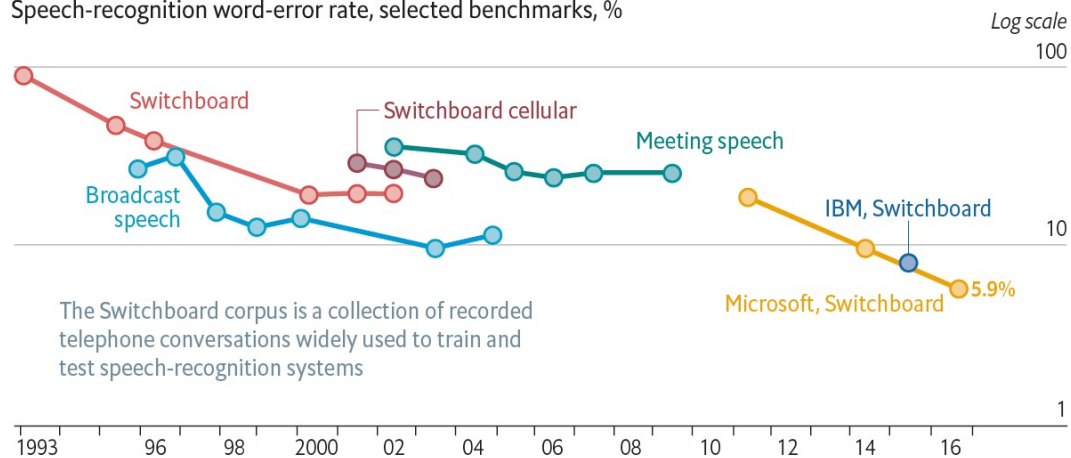
OpenAI's text-generating system GPT-3 is now spewing out 4.5 billion words a day

Robot-generated writing looks set to be the next big thing

By James Vincent | Mar 29, 2021, 8:24am EDT

Loud and clear

Speech-recognition word-error rate, selected benchmarks, %



The Switchboard corpus is a collection of recorded telephone conversations widely used to train and test speech-recognition systems

Sources: Microsoft; research papers

Speak easy

Human scorers' rating* of Google Translate and human translation

Translation method | Phrase-based† | Neural-network† | Human



Input sentence Pour l'ancienne secrétaire d'Etat, il s'agit de faire oublier un mois de cafouillages et de convaincre l'auditoire que M. Trump n'a pas l'étoffe d'un président

Phrase-based†

For the former secretary of state, this is to forget a month of bungling and convince the audience that Mr Trump has not the makings of a president

Neural-network†

For the former secretary of state, it is a question of forgetting a month of muddles and convincing the audience that Mr Trump does not have the stuff of a president

Human

The former secretary of state has to put behind her a month of setbacks and convince the audience that Mr Trump does not have what it takes to be a president


Source: Google

*0=completely nonsense translation, 6=perfect translation †Machine translation

Problems with few data (labels)

- Privacy, security issues

Assistive diagnosis



“The heart size and mediastinal contours appear within normal limits. There is blunting of the right lateral costophrenic sulcus which could be secondary to a small effusion versus scarring ...”

Normal findings

Abnormal findings

Problems with few data (labels)

- Expensive to collect/annotate
- Controllable content generation

Controlling sentiment

Pos The film is *full of imagination!*



Neg The film is *strictly routine!*

Controlling writing style

Plain

LeBron James *contributed* 26 points, 8 rebounds, 7 assists.



Elaborate

LeBron James *rounded out the box score with an all around impressive performance, scoring* 26 points, *grabbing* 8 rebounds and *dishing out* 7 assists.

Problems with few data (labels)

- Expensive to collect/annotate
- Controllable content generation



Source image

Generated images under different poses

Applications: virtual clothing try-on system

Problems with few data (labels)

- Expensive to collect/annotate

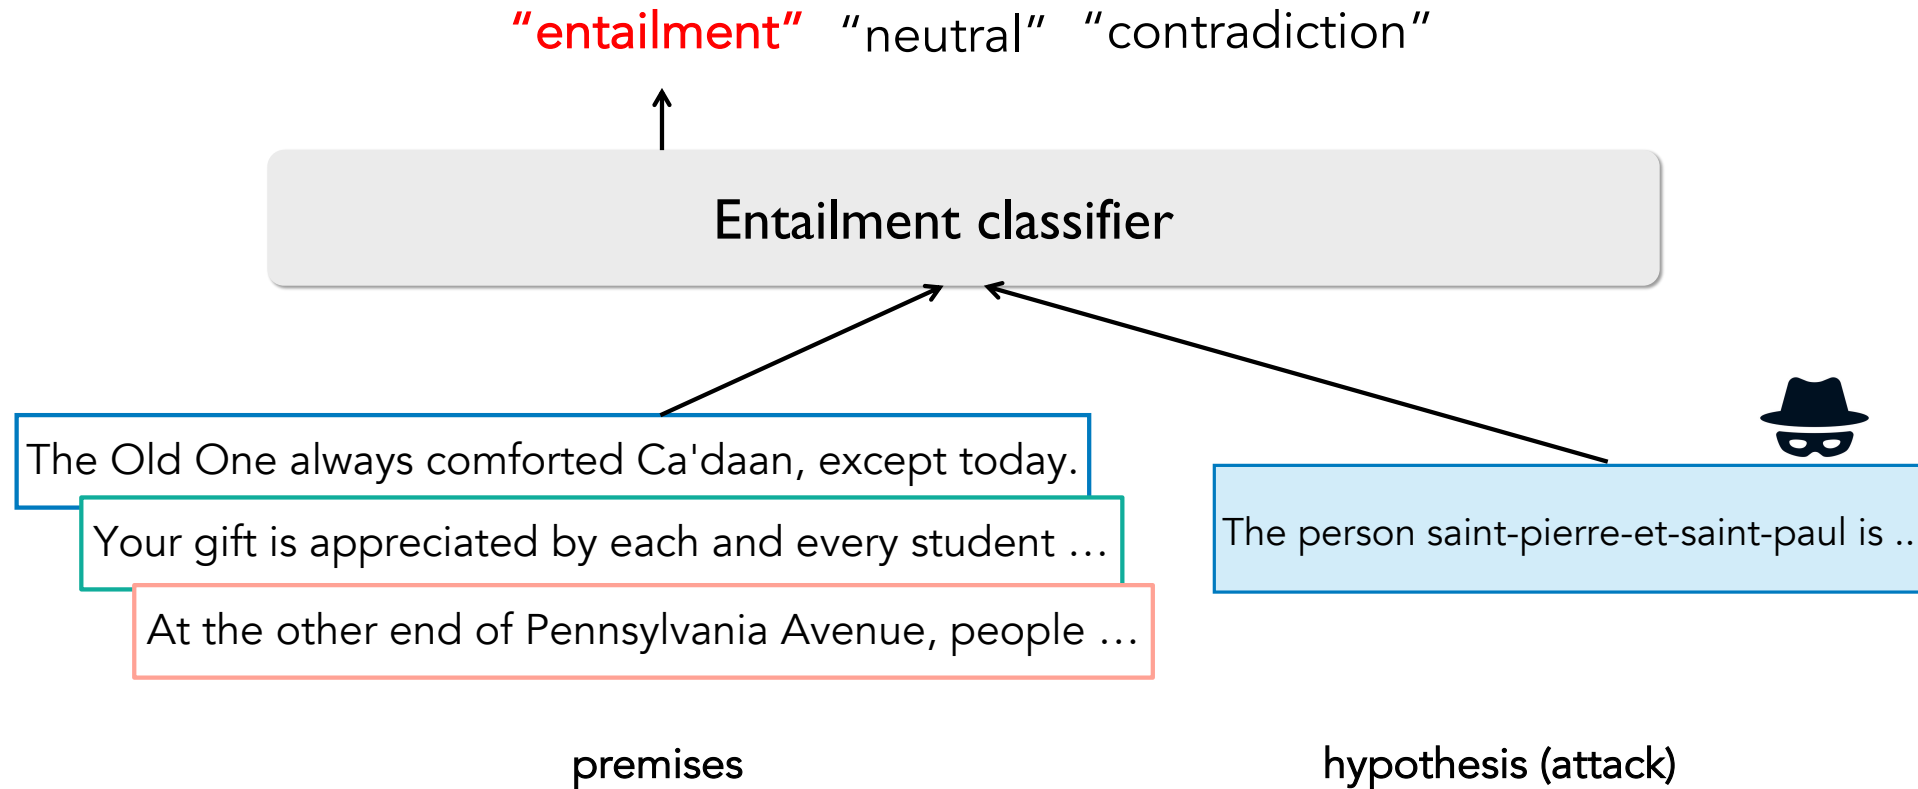
Robotic control



Problems with few data (labels)

- Difficult / expertise-demanding to annotate

Adversarial attack

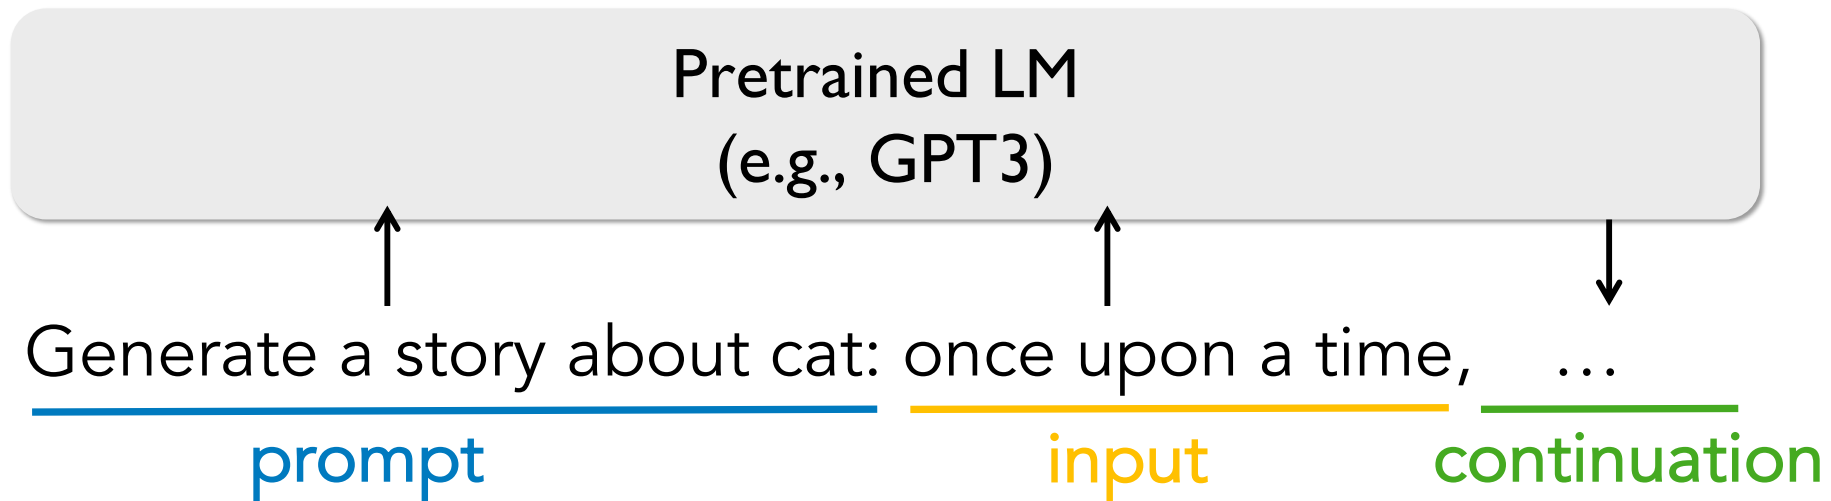


Applications: test model robustness

Problems with few data (labels)

- Difficult / expertise-demanding to annotate

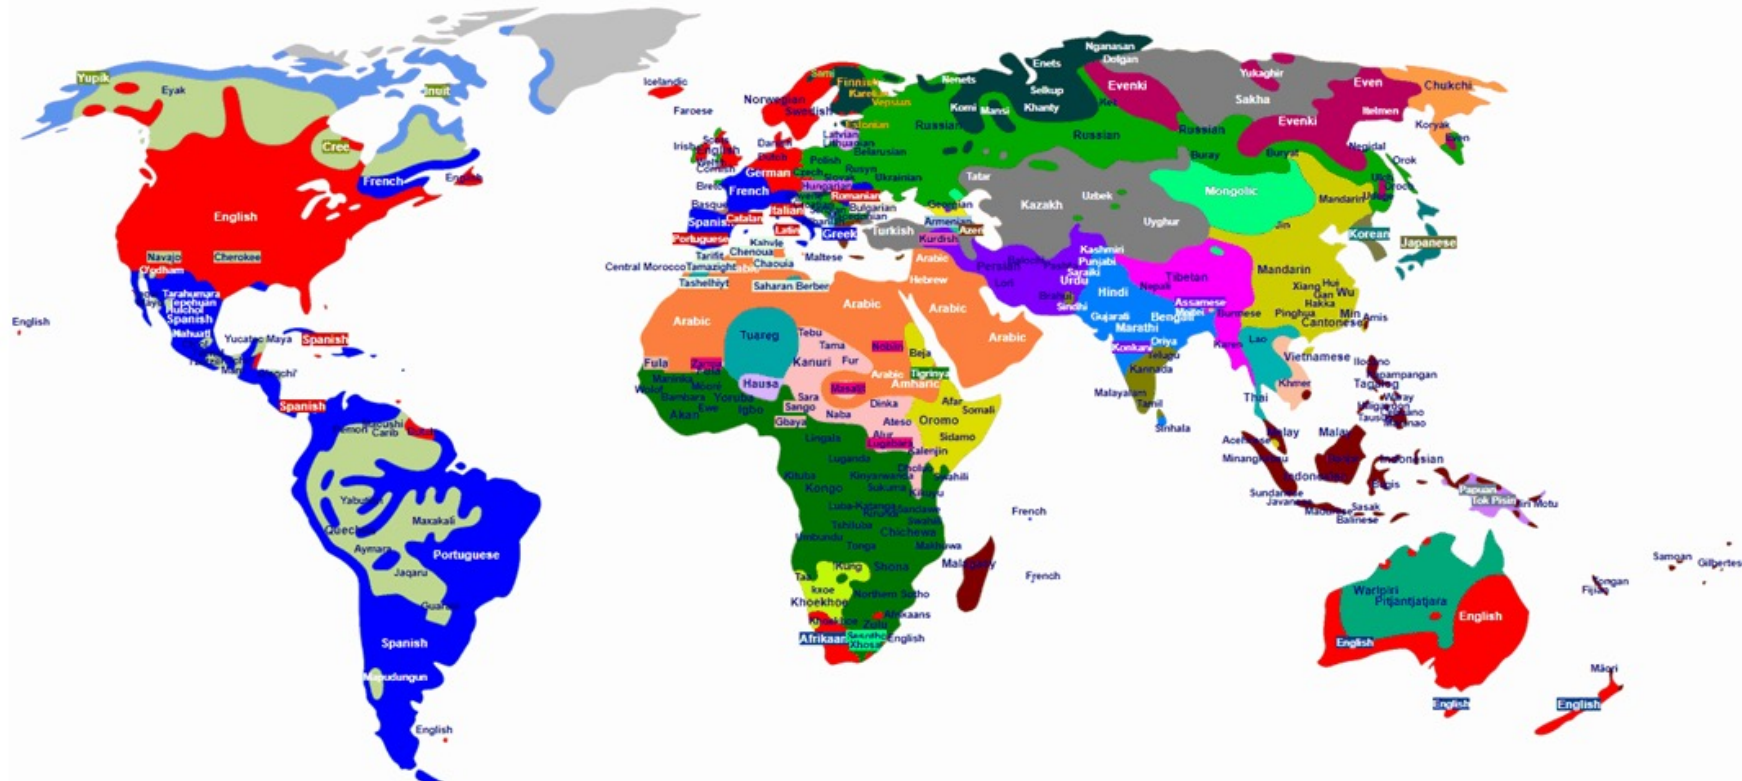
Prompt generation: automatically generating prompts to steer pretrained LMs



Problems with few data (labels)


- Specific domain Low-resource languages

~7K languages in the world




Problems with few data (labels)

- Specific domain Low-resource languages



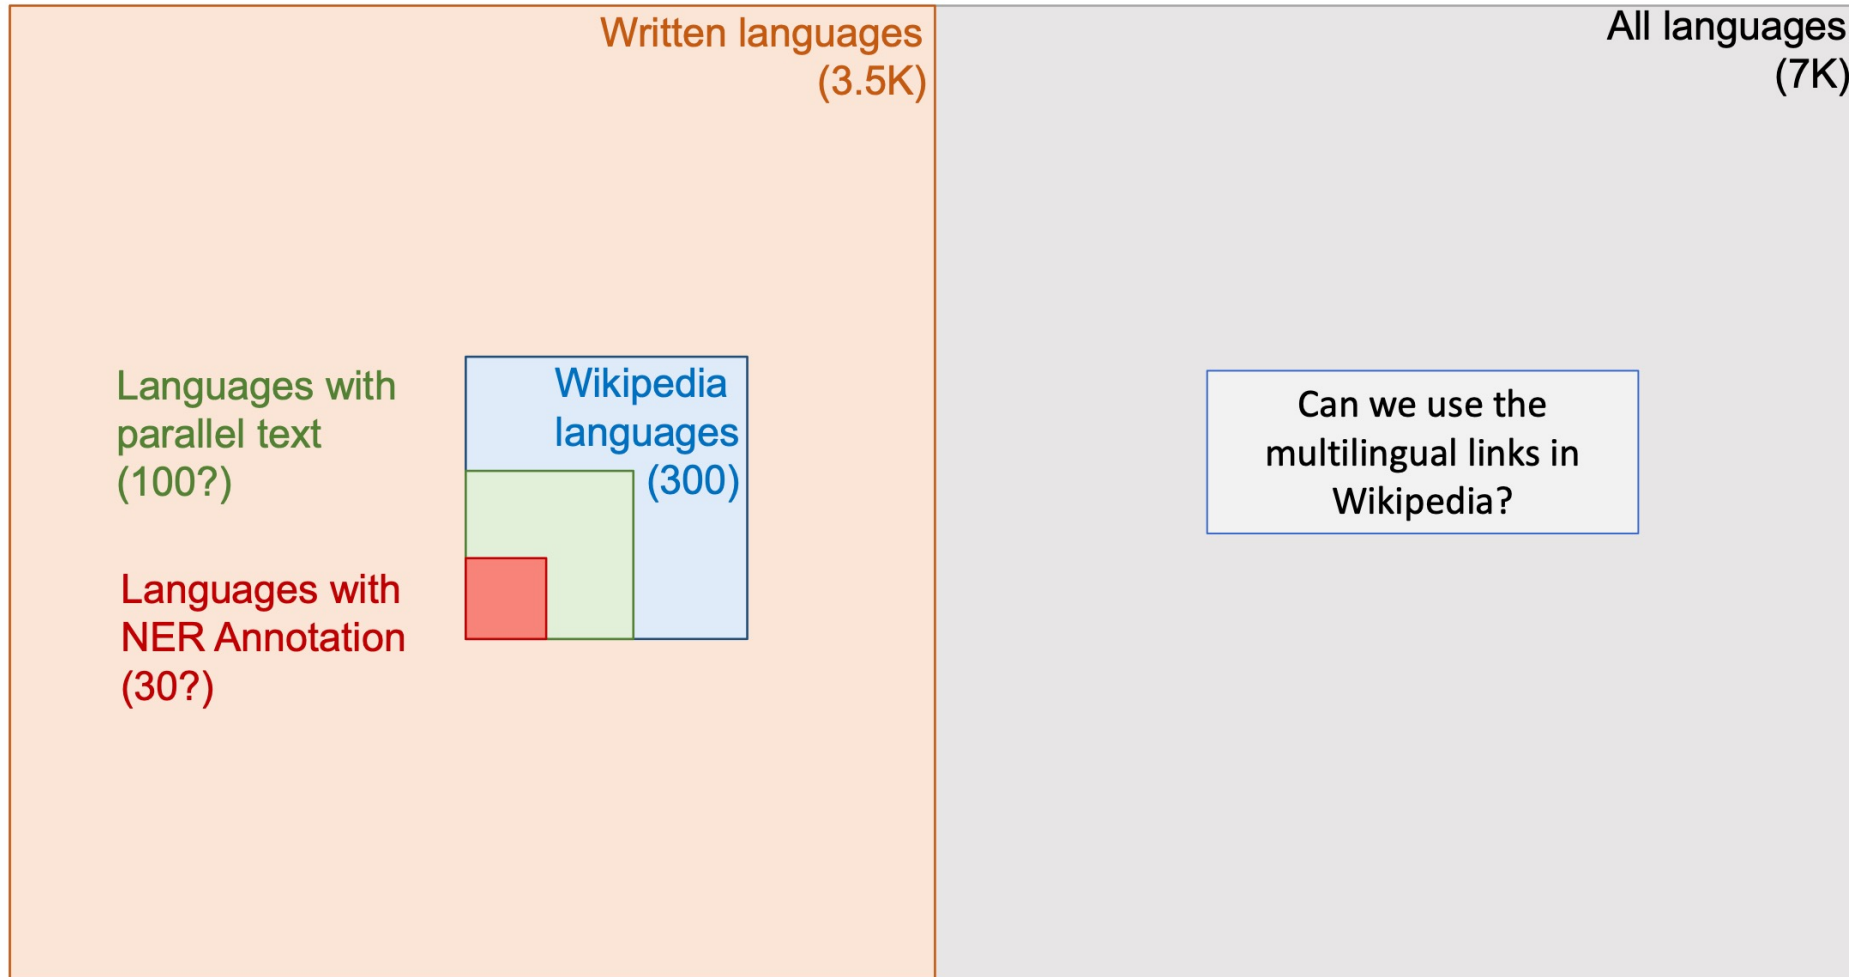
Problems with few data (labels)

- Specific domain Low-resource languages



Problems with few data (labels)

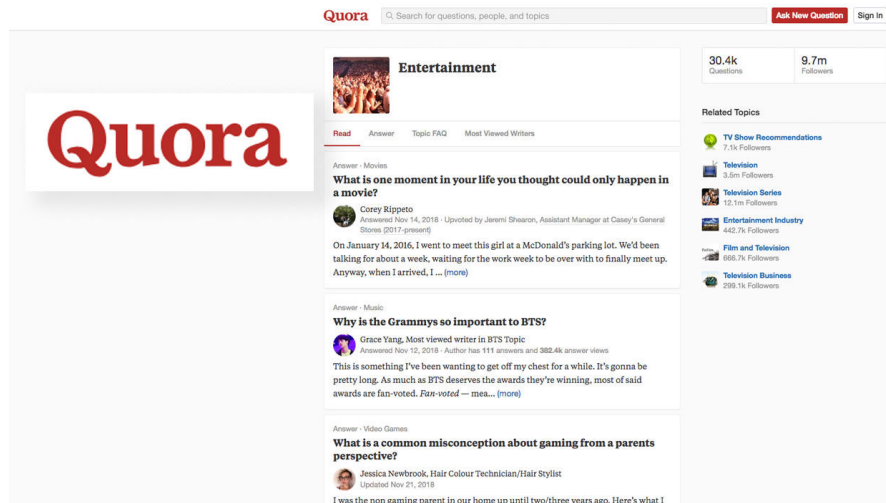
- Specific domain Low-resource languages




Problems with few data (labels)

- Specific domain

Question answering



QA based on car manual?



Problems with few data (labels)

- Privacy, security issues
- Expensive to collect/annotate
- Difficult / expertise-demanding to annotate
- Specific domain

Machine learning solutions given few data (labels)


- How can we make more efficient use of the data?
 - Clean but small-size
 - Noisy
 - Out-of-domain
- Can we incorporate other types of experiences in learning?



Data examples

Type-2 diabetes is 90% more common than type-1

Rules/Constraints



Knowledge graphs




Rewards



Auxiliary agents



Adversaries




Master classes

... And all combinations thereof

Components of a ML solution (roughly)


- Loss
- Experience
- Optimization solver
- Model architecture



Components of a ML solution (roughly)

- Loss
- Experience
- Optimization solver
- Model architecture

This course does *not* discuss model architecture



Components of a ML solution (roughly)

- Loss
- Experience
- Optimization solver
- Model architecture

This course does *not* discuss model architecture

Model of certain architecture whose parameters are the subject to be learned, $p_{\theta}(\mathbf{x}, \mathbf{y})$ or $p_{\theta}(\mathbf{y}|\mathbf{x})$

- Neural networks
- Graphical models
- Compositional architectures


Components of a ML solution (roughly)

- Loss
- Experience
- Optimization solver
- Model architecture

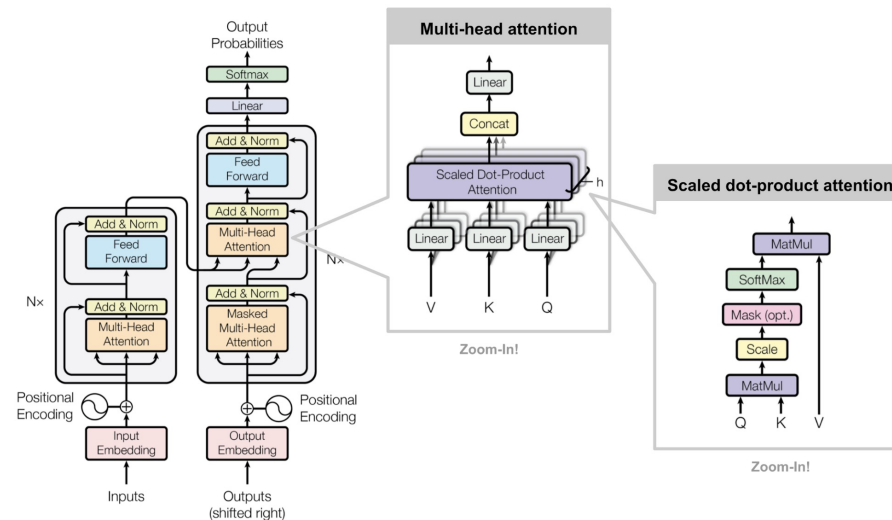
This course does *not* discuss model architecture

Model of certain architecture whose parameters are the subject to be learned, $p_{\theta}(\mathbf{x}, \mathbf{y})$ or $p_{\theta}(\mathbf{y}|\mathbf{x})$

- Neural networks
- Graphical models
- Compositional architectures



Convolutional networks



Transformers

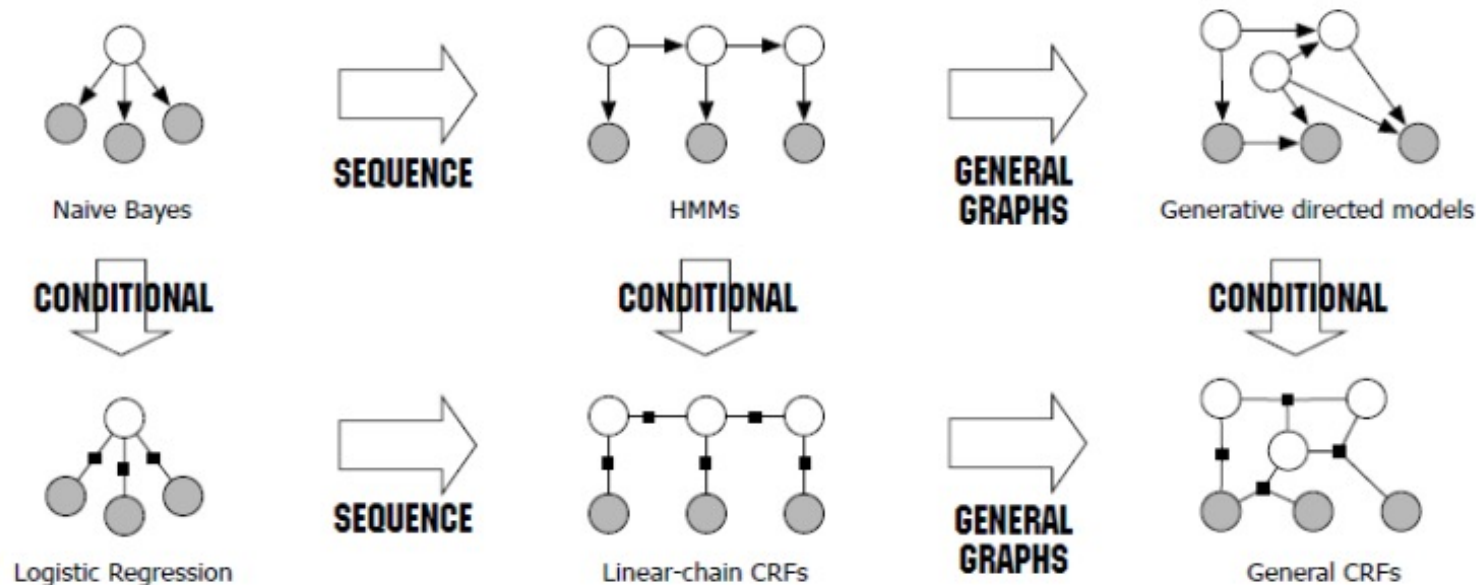
Components of a ML solution (roughly)

- Loss
- Experience
- Optimization solver
- **Model architecture**

This course does *not* discuss model architecture

Model of certain architecture whose parameters are the subject to be learned, $p_{\theta}(\mathbf{x}, \mathbf{y})$ or $p_{\theta}(\mathbf{y}|\mathbf{x})$

- Neural networks
- Graphical models
- Compositional architectures




Components of a ML solution (roughly)

- Loss
- Experience
- Optimization solver
- Model architecture

This course discusses *a little* about optimization


Assuming you know basic procedures:

- (Stochastic) gradient descent
- Backpropagation
- Lagrange multiplier
- ...




Components of a ML solution (roughly)

- Loss This course discusses *a lot* of loss & experience
- Experience
- Optimization solver Core of most learning algorithms
- Model architecture



Machine learning solutions given few data (labels)


- (1) How can we make more efficient use of the data?
 - Clean but small-size, Noisy, Out-of-domain
- (2) Can we incorporate other types of experiences in learning?



Data examples

Type-2 diabetes is 90% more common than type-1

Rules/Constraints



Knowledge graphs




Rewards



Auxiliary agents



Adversaries



Master classes

... And all combinations thereof

Machine learning solutions given few data (labels)


- (1) How can we make more efficient use of the data?
 - Clean but small-size, Noisy, Out-of-domain, ...
- Algorithms
 - **Supervised learning:** MLE, maximum entropy principle
 - **Unsupervised learning:** EM, variational inference, VAEs
 - **Self-supervised learning:** successful instances, e.g., BERT, GPT-3, contrastive learning, applications to downstream tasks
 - **Distant/weakly supervised learning:** successful instances
 - **Data manipulation:** augmentation, re-weighting, curriculum learning, ...
 - **Meta-learning**

Mostly first half of the course

Machine learning solutions given few data (labels)

- (2) Can we incorporate other types of experiences in learning?

- Learning from auxiliary models, e.g., adversarial models:
 - Generative adversarial learning (GANs and variants), co-training, ...
- Learning from structured knowledge
 - Posterior regularization, constraint-driven learning, ...
- Learning from rewards
 - Reinforcement learning: model-free vs model-based, policy-based vs value-based, on-policy vs off-policy, extrinsic reward vs intrinsic reward, ...
- Learning in dynamic environment *(not covered)*
 - Online learning, lifelong/continual learning, ...



Algorithm marketplace

Designs Driven by: experience, task, loss function, training procedure ...



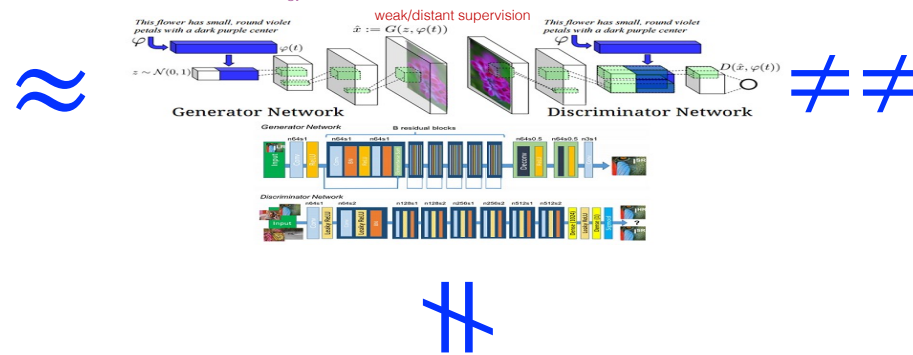


maximum likelihood estimation reinforcement learning as inference
data re-weighting inverse RL active learning
policy optimization
data augmentation reward-augmented maximum likelihood
label smoothing imitation learning softmax policy gradient
actor-critic adversarial domain adaptation
GANs posterior regularization
knowledge distillation intrinsic reward constraint-driven learning
prediction minimization generalized expectation
regularized Bayes learning from measurements
energy-based GANs
weak/distant supervision

Where we are now? Where we want to be?

- Alchemy vs chemistry

maximum likelihood estimation reinforcement learning as inference
 data re-weighting policy optimization inverse RL active learning
 data augmentation reward-augmented maximum likelihood
 label smoothing imitation learning softmax policy gradient
 actor-critic GANs adversarial domain adaptation
 knowledge distillation posterior regularization
 intrinsic reward constraint-driven learning
 prediction minimization generalized expectation
 regularized Bayes learning from measurements
 energy-based GANs



© Mind Juice Media Inc. All rights reserved


Quest for more standardized, unified ML principles

Machine Learning 3: 253–259, 1989
© 1989 Kluwer Academic Publishers – Manufactured in The Netherlands

EDITORIAL

Toward a Unified Science of Machine Learning

[P. Langley,
1989]



REVIEW ————— Communicated by Steven Nowlan

A Unifying Review of Linear Gaussian Models

Sam Roweis*


Computation and Neural Systems, California Institute of Technology, Pasadena, CA 91125, U.S.A.

Zoubin Ghahramani*

Department of Computer Science, University of Toronto, Toronto, Canada

Physics in the 1800's

- Electricity & magnetism:
 - Coulomb's law, Ampère, Faraday, ...
- Theory of light beams:
 - Particle theory: Isaac Newton, Laplace, Plank
 - Wave theory: Grimaldi, Chris Huygens, Thomas Young, Maxwell
- Law of gravity
 - Aristotle, Galileo, Newton, ...



"Standard equations" in Physics

Maxwell's Eqns:
original form

$e + \frac{df}{dx} + \frac{dg}{dy} + \frac{dh}{dz} = 0$	(1) Gauss' Law
$\mu\alpha = \frac{dH}{dy} - \frac{dG}{dz}$ $\mu\beta = \frac{dF}{dz} - \frac{dH}{dx}$ $\mu\gamma = \frac{dG}{dx} - \frac{dF}{dy}$	(2) Equivalent to Gauss' Law for magnetism
$P = \mu \left(\gamma \frac{dy}{dt} - \beta \frac{dz}{dt} \right) - \frac{dF}{dt} - \frac{d\Psi}{dz}$ $Q = \mu \left(\alpha \frac{dz}{dt} - \gamma \frac{dx}{dt} \right) - \frac{dG}{dt} - \frac{d\Psi}{dy}$ $R = \mu \left(\beta \frac{dx}{dt} - \alpha \frac{dy}{dt} \right) - \frac{dH}{dt} - \frac{d\Psi}{dx}$	(3) Faraday's Law (with the Lorentz Force and Poisson's Law)
$\frac{dy}{dy} - \frac{d\beta}{dz} = 4\pi p'$ $\frac{d\alpha}{dz} - \frac{d\gamma}{dx} = 4\pi q'$ $\frac{d\beta}{dx} - \frac{d\alpha}{dy} = 4\pi r'$	(4) Ampère-Maxwell Law
$P = -\xi p \quad Q = -\xi q \quad R = -\xi r$	Ohm's Law
$P = kf \quad Q = kg \quad R = kh$	The electric elasticity equation ($\mathbf{E} = \mathbf{D}/\epsilon$)
$\frac{de}{dt} + \frac{dp}{dx} + \frac{dq}{dy} + \frac{dr}{dz} = 0$	Continuity of charge


Maxwell's Eqns simplified w/ rotational symmetry

$$\nabla \cdot \mathbf{D} = \rho_V$$

$$\nabla \cdot \mathbf{B} = 0$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}$$


$$\nabla \times \mathbf{H} = \frac{\partial \mathbf{D}}{\partial t} + \mathbf{J}$$



Maxwell's Eqns further simplified w/ symmetry of special relativity

$$\epsilon^{uvk\lambda} \partial_v F_{k\lambda} = 0$$

$$\partial_v F^{uv} = \frac{4\pi}{c} j^u$$




Standard Model w/ Yang-Mills theory and US(3) symmetry

$$\mathcal{L}_{gf} = -\frac{1}{2} \text{Tr}(F^2)$$

$$= -\frac{1}{4} F^{\alpha\mu\nu} F_{\mu\nu}^{\alpha}$$



Unification of fundamental forces?



Diverse electro-magnetic theories



1861

1910s

1970s



A "standardized formalism" of ML



Data examples

Type-2 diabetes
is 90% more
common than
type-1

Constraints




Rewards



Auxiliary agents



Adversaries



Imitation

$$\min_{q, \theta} - \mathbb{H} + \mathbb{D} - \mathbb{E}$$

Uncertainty Divergence Experience

- Panoramically learn from all types of experiences
- Subsumes many existing algorithms as special cases

Questions?