

# Introduction

He He

New York University

September 5, 2020

# Logistics



He He



Gaomin Wu



Shivesh Ganju



Gauri Dhawan

- ▶ Best way to communicate: Piazza (**remember to sign up**).
- ▶ Lectures and office hours will be online (sadly).
- ▶ Let us know if you have accessibility needs.

# What you'll be able to do by the end of this course

- ▶ Understand the core problems and challenges in NLP
- ▶ Formalize tasks as statistical learning problems
- ▶ Have a toolbox for solving different families of NLP problems
- ▶ Gain hands-on experience in building NLP systems

# What we expect you to know

- ▶ **Linear algebra:** vector, dot product, gradient computation etc.
- ▶ **Probability and statistics:** conditional probability, expectation, Bayes rule etc.
- ▶ **Basic machine learning:** loss function, gradient descent etc.
- ▶ **Programming:** read and write Python code, use Numpy (and MXNet)

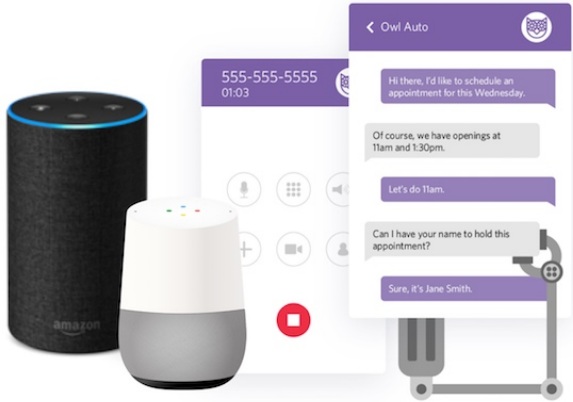
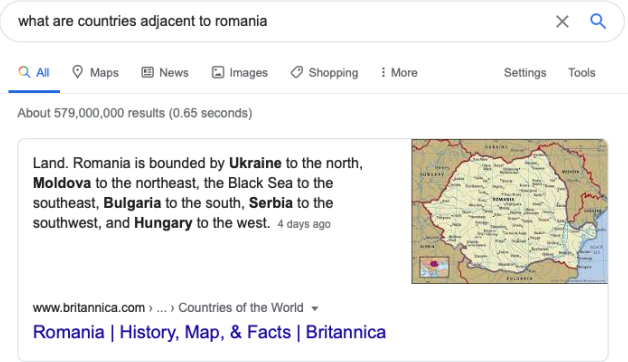
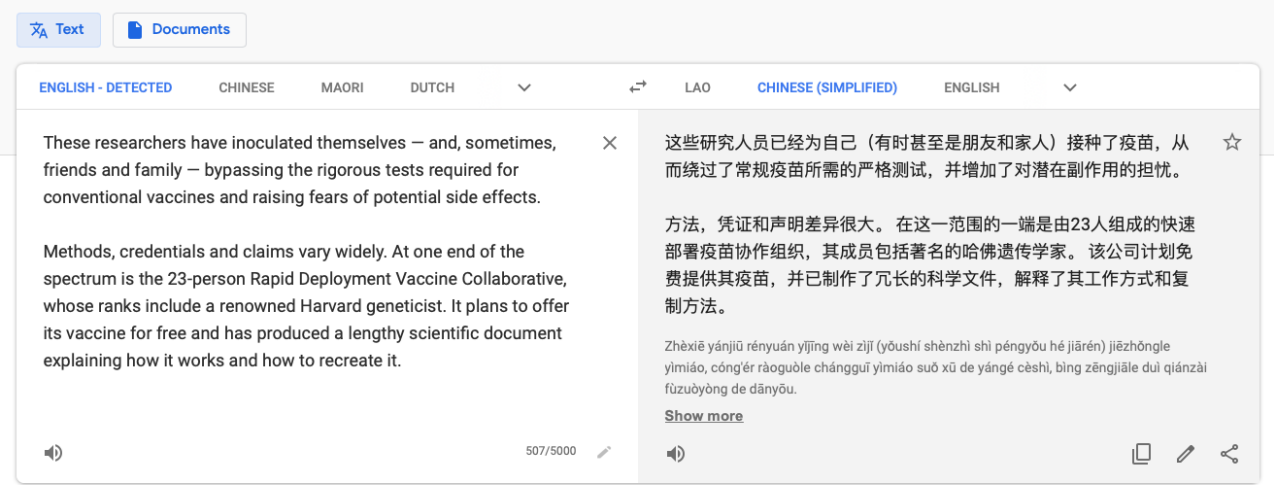
# Table of Contents

1. A brief history

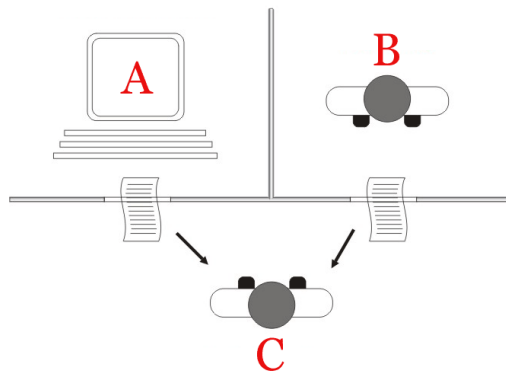
2. Challenges in NLP

3. Course overview

# Products powered by NLP technologies



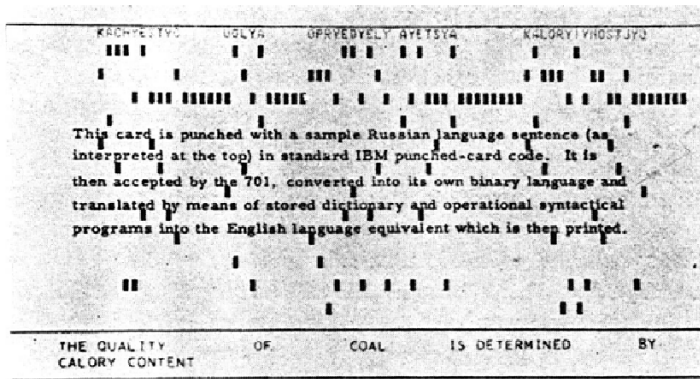
# The imitation game



*“I believe that in about **fifty years**’ time it will be possible to programme computers, with a **storage capacity of about  $10^9$** , to make them play the imitation game so well that an average interrogator will not have more than 70 percent chance of making the right identification after five minutes of questioning.” Turing (1950)*

# The Georgetown-IBM experiment

- ▶ The program:



- ▶ 250 words

- ▶ 6 grammar rules, e.g.

“If first code is 110, is third code associated with preceding complete word equal to 21? If so, reverse order of appearance of words in output (i.e., word carrying 21 should follow that carrying 110)—otherwise, retain order.”



# ELIZA

- ▶ Built by Joseph Weizenbaum at MIT in 1964 to demonstrate the **superficiality** of human-machine communication.
- ▶ Surprisingly, people were convinced that ELIZA had human intelligence.

Human: Well, my boyfriend made me come here.

ELIZA: Your boyfriend made you come here?

Human: He says I'm depressed much of the time.

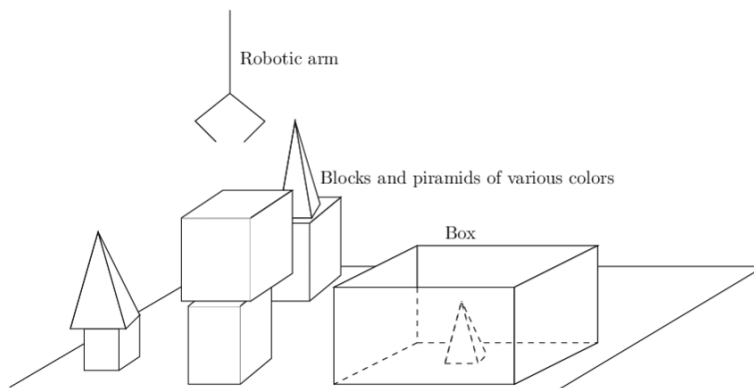
ELIZA: I am sorry to hear you are depressed.

Human: It's true. I'm unhappy.

ELIZA: Do you think coming here will help you not to be unhappy?

# SHRDLU

- ▶ Built by Terry Winograd at MIT in 1968.



Person: Pick up a big red block.

Computer: OK.

Person: Grasp the pyramid.

Computer: I DON'T UNDERSTAND WHICH PYRAMID YOU MEAN.

Person (changing their mind): Find a block which is taller than the one you are holding and put it into the box.

Computer: BY "IT", I ASSUME YOU MEAN THE BLOCK WHICH IS TALLER THAN THE ONE I AM HOLDING.

# Limitations of early systems

- ▶ Optimism in the 50's and 60's
  - “Within the very near future—much less than twenty-five years—we shall have the technical capability of substituting machines for any and all human functions in organizations.”*
- ▶ Disappointing results due to
  - ▶ **Limited computation:** hardware limitation
  - ▶ **Combinatorial explosion:** algorithms are intractable in realistic settings
  - ▶ **Underestimated complexity:** ambiguity, commonsense knowledge etc.

# The rise of statistical learning in the 80's

- ▶ Notable progress in MT from IBM (neglected knowlege of linguistics).
- ▶ HMMs widely used for speech recognition.  
*“Every time I fire a linguist, the performance of the speech recognizer goes up.”*—Frederick Jelinek.
- ▶ Machine learning is the main driving force of NLP today.

# Table of Contents

1. A brief history

2. Challenges in NLP

3. Course overview

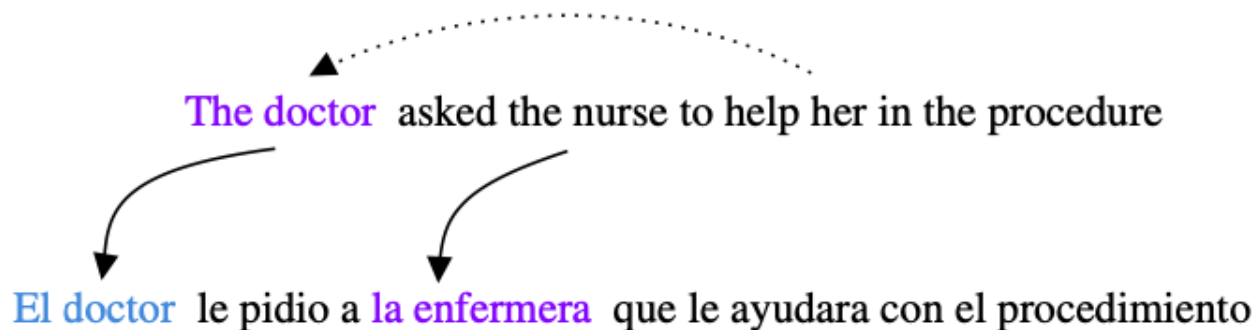
# Are we there yet?

Predictions are not robust to benign perturbations [Ribeiro+ 2020].

Test case	Expected	Predicted	Pass?
<b>A</b> Testing <b>Negation</b> with <i>MFT</i>	Labels: negative, positive, neutral		
<b>Template:</b> I {NEGATION} {POS_VERB} the {THING}.			
I can't say I recommend the food.	neg	pos	X
I didn't love the flight.	neg	neutral	X
...			
Failure rate = 76.4%			

# Are we there yet?

MT systems are prone to gender-biased translation errors [Stanovsky+2019].



# Are we there yet?

QA models are easily distracted by irrelevant sentences [Jia+ 2017].

**Article:** Super Bowl 50

**Paragraph:** *“Peyton Manning became the first quarterback ever to lead two different teams to multiple Super Bowls. He is also the oldest quarterback ever to play in a Super Bowl at age 39. The past record was held by John Elway, who led the Broncos to victory in Super Bowl XXXIII at age 38 and is currently Denver’s Executive Vice President of Football Operations and General Manager. Quarterback Jeff Dean had jersey number 37 in Champ Bowl XXXIV.”*

**Question:** *“What is the name of the quarterback who was 38 in Super Bowl XXXIII?”*

**Original Prediction:** John Elway

**Prediction under adversary:** Jeff Dean



# Why is language hard?

sparsity

ambiguity

context-dependent

# Why is language hard?

## ▶ Discrete

### ▶ How to define metrics?

I work **at** NYU. vs I work **for** NYU.

This is good. vs This is **actually** good.

### ▶ How to define transformations?

The food is okay. → The food is awesome!

They made a brief return to → They returned.

Cambridge to drop the book.

# Why is language hard?

## ▶ **Compositional**

- ▶ How to generalize when we don't see all possible combinations?

Vocabulary (from [Lake+ 2018]):

{jump, walk, turn, once, twice, left, right, before, after,  
and}

Sentences:

jump

jump left

jump left and walk right

jump left after walk right once before turn left twice

...

# Why is language hard?

## ▶ Sparse

- ▶ How to handle the long tail?

BoA's financial assistant Erica:

The bank “learned [that] there are over 2,000 different ways to ask us to move money.”

Zipf's law:

$$\text{word frequency} \propto \frac{1}{\text{rank}}$$

# Why is language hard?

## ▶ **Ambiguous**

- ▶ How to interpret meaning in context?

Bass: fish? guitar? frequency?

I shot an elephant in my pajamas: who is in the pajamas?

The spirit is willing but the flesh is weak.

→ The vodka is strong but the meat is rotten.

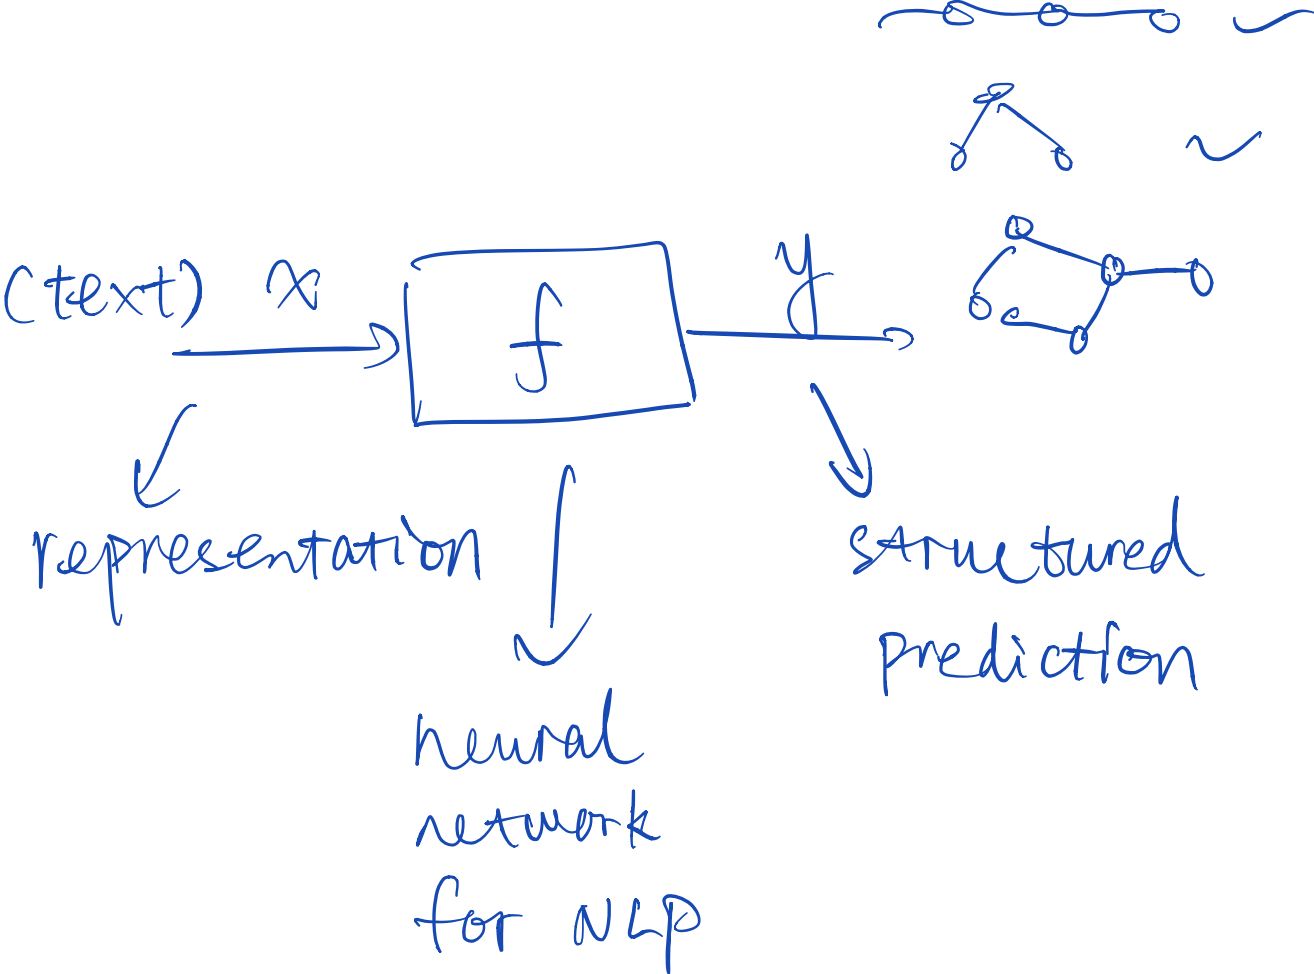
# Table of Contents

1. A brief history

2. Challenges in NLP

3. Course overview

# Overview



# Representation of text

**Symbolic:** a set of objects (concepts)

a white kitten = {white, a, kitten, noun phrase}

**Distributed:** a list of components (properties)

a white kitten = [COLOR=white, SIZE=small  
OBJECT=kitten, NUMBER=1]

Pros and cons?



# Structured prediction: sequences

- ▶ Named entity recognition

New York University is a private research university based in New  
ORG

York City. It is founded in 1831 by Albert Gallatin.  
LOC YEAR PEOPLE

CT of the maxillofacial area showed no facial bone fracture. CT of  
test symptom

the brain showed no acute changes.  
test symptom


# Structured prediction: sequences

## ► Anaphora resolution

John had a great evening meeting with his high school friends.



The city councilmen refused the demonstrators a permit because they feared violence.



# Structured prediction: sequences

- ▶ Semantic role labeling (slot filling)

I would like to book a ticket from New York to San Francisco on Christmas eve.

action=book\_ticket

departure city=

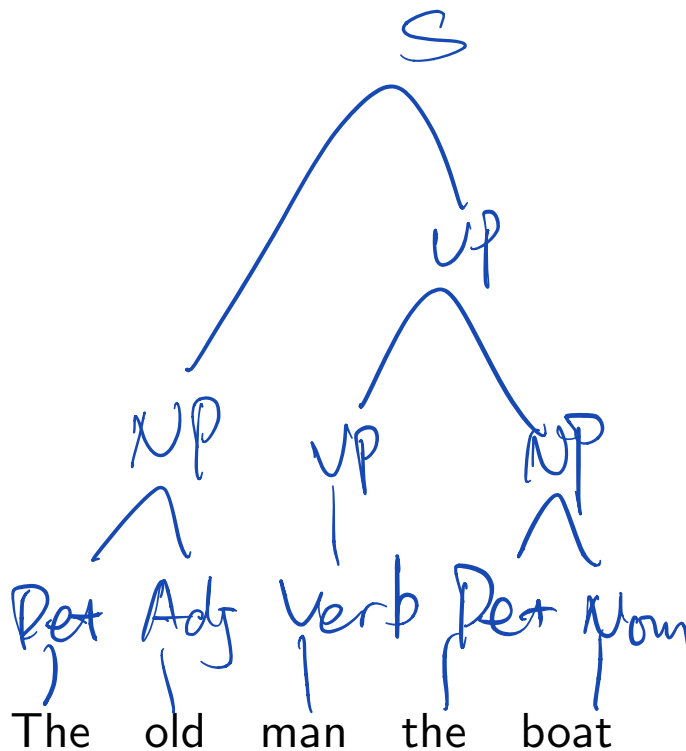
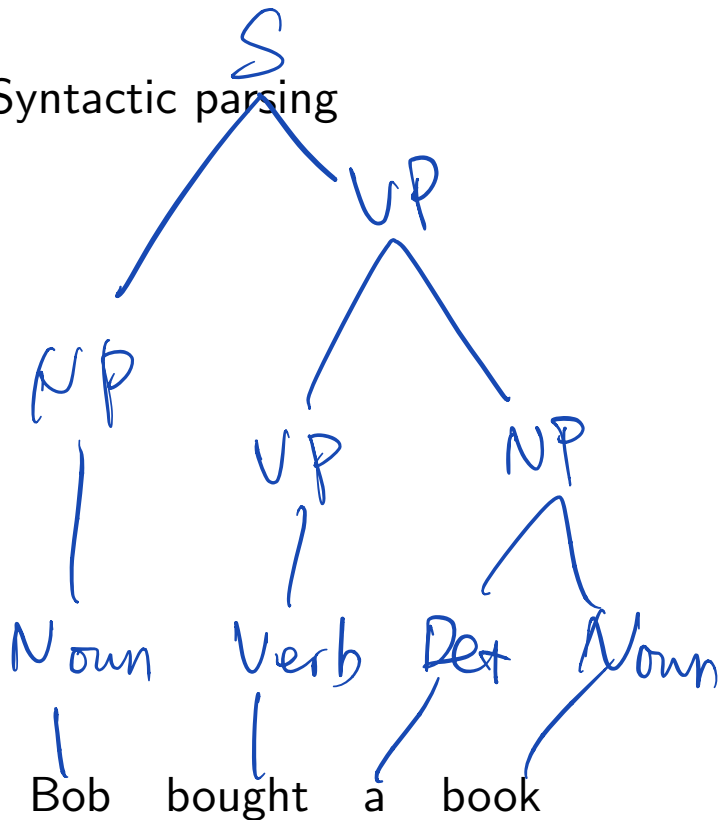
destination city=

date=

time=

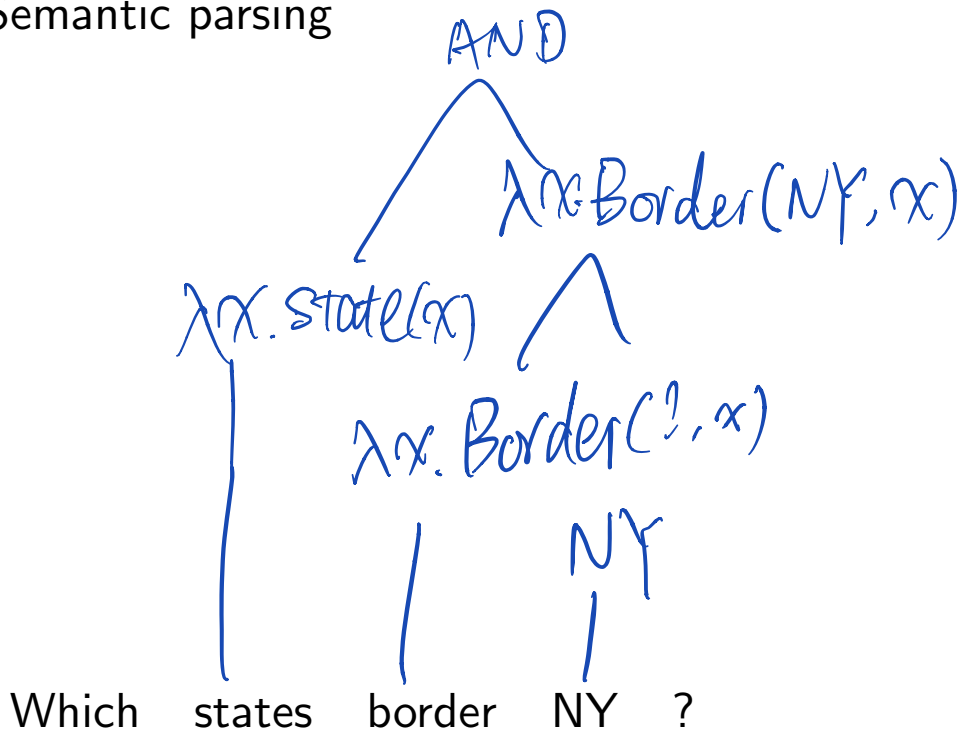
# Structured prediction: trees

► Syntactic parsing



# Structured prediction: trees

► Semantic parsing



# Text generation

- ▶ Machine translation

爱屋及乌 → Love me, love my dog

- ▶ Data-to-text

Date	min	max	
tomorrow	21°C	29°C	→ Tomorrow's temperature will be between 21 and 29 degrees.

- ▶ Summarization

The Justice Department plans to bring an antitrust case against Google as soon as this month, after Attorney General William P. Barr overruled career lawyers who said they needed more time to build a strong case against one of the world's wealthiest, most formidable technology companies, according to five people briefed on internal department conversations.

→ Justice Dept. plans to file antitrust charges against Google in coming weeks.

# Predict structures

- ▶ Modeling

How to model interactions among substructures?

- ▶ Learning

How to efficiently learn the model parameters?

$$\text{score}(X, y)$$

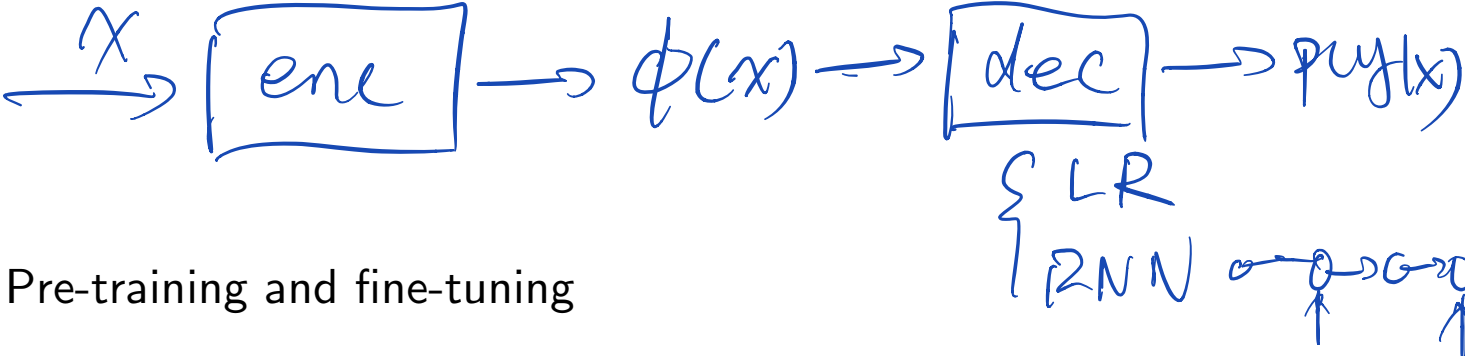
- ▶ Inference

How to efficiently find the best structure given a learned model?

$$\arg \max_{y \in Y} \text{score}(X, y)$$

# Neural networks for NLP

- ▶ Encoder-decoder models



- ▶ Pre-training and fine-tuning



# Beyond individual sentences: discourse

What makes a collection of sentences **coherent**?

John took a train from Paris to Istanbul. He likes spinach.

Jane took a train from Paris to Istanbul. She had to attend a conference.

Transcript of Mrs. Nalini Chidambaram 1  
(Conversation between Mrs. Nalini Chidambaram and Dr. K. Kathirvel)

Mrs. Nalini Chidambaram: Bank only had symbolic possession and they took physical possession with the help of gondas because they knew abt Honorable Supreme Court order. Don't tell anybody that I am involved.

Dr. K. Kathirvel: Okay madam.

Mrs. Nalini Chidambaram: You should not scandalise it.

Dr. K. Kathirvel: Okay madam.

Mrs. Nalini Chidambaram: If it had been some other politician then they would have put you behind bars' usins colof.

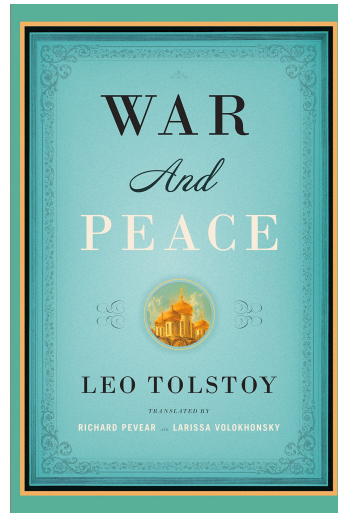
Mrs. Nalini Chidambaram: Sir, I called you because if you get a favourable court order you also take possession in the same way as they did.

Dr. K. Kathirvel: No Madam, its wrong.

Mrs. Nalini Chidambaram: I am angrm so an telling this to you.

Dr. K. Kathirvel: What ever happens I will not do that, and if I do that there will be no difference between them and me.

Mrs. Nalini Chidambaram: Its somebody else's property and they don't want to give up.



OPEN ACCESS Freely available online

PLoS ONE

## A Role for Parasites in Stabilising the Fig-Pollinator Mutualism

Derek W. Dunn<sup>1,2,3</sup>, Simon T. Saper<sup>1,2</sup>, Jo Ribby<sup>1</sup>, Ruth Chan<sup>1</sup>, Ross H. Crozier<sup>1</sup>, Douglas W. Yu<sup>1</sup>, James M. Cook<sup>1,2,4,5,6\*</sup>

**1** Division of Biology, University of California, Davis, California, United States of America, **2** School of Biological Sciences, University of Reading, Reading, United Kingdom, **3** School of Biological Sciences, University of East Anglia, Norwich, United Kingdom, **4** School of Marine and Coastal Biology, Swinburne University of Technology, Hawthorn, Victoria, Australia, **5** School of Biological Sciences, University of Exeter, Exeter, United Kingdom, **6** School of Biological Sciences, University of East Anglia, Norwich, United Kingdom

**Mutations are interspecific interactions in which both players benefit. Explaining their maintenance is problematic, because cheaters should outcompete cooperative conspecifics, leading to mutualism instability. Monoecious figs (*Ficus*) are pollinated by host-specific wasps (*Sphecidae*), whose larvae eat ovules in their "ovules". Typically, female pollinating wasps oviposit directly into *Ficus* ovules from inside the receptive syconium. Across *Ficus* species, there is a widely documented segregation of pollinator guilds to inner ovules and seeds in outer ovules. This pattern suggests that wasps avoid, or are prevented from ovipositing into, inner ovules, and this results in mutualism stability. However, the mechanism preventing wasps from ovipositing into inner ovules remains unknown. We report that in *Ficus* figs, pollinators attempting to enter ovules are vulnerable to attack by parasitic wasps that oviposit from outside the syconium. Parasitism risk decreases towards the centre of the syconium, where inner ovules provide empty space for pollinator oviposition. We suggest that this resulting gradient in oviposition stability is likely to contribute to directing the pollinator to avoid inner ovules, and by forcing wasps to focus on a subset of ovules, reduces their galling rates. This previously unacknowledged mechanism may therefore contribute to mutualism persistence independent of additional factors that involve plant defences against pollinator oviposition, or physiological constraints on pollinators that prevent oviposition in all available ovules.**

**Citation:** Dunn DW, Saper ST, Ribby J, Chan R, Crozier RH, et al. (2008) A Role for Parasites in Stabilising the Fig-Pollinator Mutualism. *PLoS ONE* 3(3): e2811. doi:10.1371/journal.pone.0028111

### Introduction

In a long-term debate in which at the heart of the individual gene [1], explaining the existence of cooperation, such as mutualism, is a major scientific challenge. Mutualisms are interspecific ecological interactions characterized by net positive benefits to both partners [2]. This usually involves costly investments for each. What factors then prevent one partner from exploiting mutualistic resources from the other or another mutualist ally [3–7]? In some mutualisms, the larger, more socially partner, manipulates the other by diverting benefits to cooperative individuals and costs to cheaters [8–7]. However, a general consensus on mutualism persistence has only recently been formulated, and this clearly does not have a high trade-off cost of cooperating is one important factor [8].

The wasp (*Ficus*) and their host-specific, specialist pollinating wasps are a classic example of an obligate mutualism [9,11]. The wasp pollinate the tree, and the tree provides resources for wasp offspring to mature in *Ficus*. Female wasps push their way through a specialised entrance into receptive syconia (colloquially, "figs"), which are enclosed inflorescences. The wasps then pollinate the tree while depositing their eggs individually into ovules. Thus, each egg laid costs the tree one seed, but spurs emergence of the female wasp offspring, dispersing that tree's pollen. Trees need to produce both wasps and seeds for the mutualism to persist, but natural selection should favour wasps that exploit the maximum number of fig ovules from the chest nuts, resulting in a conflict of interest between wasp and tree. However, the mutualism has persisted for at least 60 million years and has radiated into more than 750 species pairs [12]. The mechanisms preventing wasps

from overexploiting fig resources unknown, despite intensive study over 60 decades.

Within receptive syconia, the lengths of floral tubes are highly variable [13,14], and occupying pollinators (most commonly female *Ficus*) with shorter tubes than ovules [15–18]. Style and pedicel lengths of flowers are negatively correlated. Shorter-styled ovules develop into ovules or galls (others a wasp is present near the syconium inner cavity, while most long-styled ovules develop into seeds) the more so [19,20] (Figure 1). These patterns have been shown to reflect the competitive preferences of females, and are unlikely to be the result of greater elongation of pedicels containing eggs during ovulation, maintenance, because no receptive syconia pollinator eggs are usually present in short-styled inner ovules [19]. These ecological observations have been tied to four, not necessarily mutually exclusive, mechanisms that have been proposed to stabilise this fig-pollinator mutualism: (1) Unstable seeds—inner ovules may be detested biochemically or physically against oviposition and larval development [21]. However, no such mechanisms have yet been identified. (2) Short oviposition—pollinator oviposition may be too short to fully penetrate the long styles of

**Authors' Summary:** A *Ficus* tree and its pollinator, a wasp, have a mutualistic relationship. The wasp pollinates the tree, and the tree provides resources for wasp offspring to mature in *Ficus*. Female wasps push their way through a specialised entrance into receptive syconia (colloquially, "figs"), which are enclosed inflorescences. The wasps then pollinate the tree while depositing their eggs individually into ovules. Thus, each egg laid costs the tree one seed, but spurs emergence of the female wasp offspring, dispersing that tree's pollen. Trees need to produce both wasps and seeds for the mutualism to persist, but natural selection should favour wasps that exploit the maximum number of fig ovules from the chest nuts, resulting in a conflict of interest between wasp and tree. However, the mutualism has persisted for at least 60 million years and has radiated into more than 750 species pairs [12]. The mechanisms preventing wasps

from overexploiting fig resources unknown, despite intensive study over 60 decades.

**Abstract:** In a long-term debate in which at the heart of the individual gene [1], explaining the existence of cooperation, such as mutualism, is a major scientific challenge. Mutualisms are interspecific ecological interactions characterized by net positive benefits to both partners [2]. This usually involves costly investments for each. What factors then prevent one partner from exploiting mutualistic resources from the other or another mutualist ally [3–7]? In some mutualisms, the larger, more socially partner, manipulates the other by diverting benefits to cooperative individuals and costs to cheaters [8–7]. However, a general consensus on mutualism persistence has only recently been formulated, and this clearly does not have a high trade-off cost of cooperating is one important factor [8].

PLoS Biology | www.plosbiology.org

0600

March 2008 | Volume 6 | Issue 3 | e2811

# Beyond individual sentences: grounding

Connect language to the world

*“Can you bring me an apple?”*



# Project

- ▶ Related to NLP (doesn't have to be in the scope of this course)
- ▶ New algorithms or models for existing problems
- ▶ Novel applications of NLP techniques
- ▶ Analysis of well-known approaches that leads to new insight
- ▶ Overall rule: should increase our knowledge in some way