

Overview of NLP Tasks and Featurization

Frank Ferraro – ferraro@umbc.edu
CMSC 473/673

Today's Learning Goals

- Define featurization and some examples
- Learn about NLP Tasks at a high-level, e.g.,
 - Document classification
 - Part of speech tagging
 - Syntactic parsing
 - Entity id/coreference

ML Term: “Featurization”

The procedure of extracting **features** for some input

Often viewed as a K-dimensional vector function f of the input language x

$$f(x) = (\underbrace{f_1(x)}_{\text{feature}}, \dots, \underbrace{f_K(x)}_{\text{feature}})$$

Each of these is a feature
(/feature function)

ML Term: “Featurization”

The procedure of extracting **features** for some input

Often viewed as a K-dimensional vector function f of the input language x

$$f(x) = (f_1(x), \dots, f_K(x))$$

In supervised settings, it can equivalently be viewed as a K-dimensional vector function f of the input language x and a potential label y

$$f(x, y) = (f_1(x, y), \dots, f_K(x, y))$$

Features can be thought of as “soft” rules

E.g., POSITIVE sentiments tweets *may* be more likely to have the word “happy”

Defining Appropriate Features

Feature functions help extract useful features (characteristics) of the data

They turn *data* into *numbers*

Features that are not 0 are said to have **fired**

Defining Appropriate Features

Feature functions help extract useful features (characteristics) of the data

They turn *data* into *numbers*

Features that are not 0 are said to have fired

You can define classes of features by *templating* (we'll come back to this!)

Often binary-valued (0 or 1), but can be real-valued

Three Common Types of Featurization in NLP

1. Bag-of-words (or bag-of-characters, bag-of-relations)
2. Linguistically-inspired features
3. Dense features via embeddings

Three Common Types of Featurization in NLP

1. Bag-of-words (or bag-of-characters, bag-of-relations)
 - Identify *unique* sufficient atomic sub-parts (e.g., words in a document)
 - Define simple features over these, e.g.,
 - Binary (0 or 1) → indicating presence
 - Natural numbers → indicating number of times in a context
 - Real-valued → various other score (we'll see examples throughout the semester)
2. Linguistically-inspired features
3. Dense features via embeddings

Three Common Types of Featurization in NLP

1. Bag-of-words (or bag-of-characters, bag-of-relations)
 - Identify **unique** sufficient atomic sub-parts (e.g., words in a document)
 - Define simple features over these, e.g.,
 - Binary (0 or 1) → indicating presence
 - Natural numbers → indicating number of times in a context
 - Real-valued → various other score (we'll see examples throughout the semester)
2. Linguistically-inspired features
 - Define features from words, word spans, or linguistic-based annotations extracted from the document
3. Dense features via embeddings

Three Common Types of Featurization in NLP

1. Bag-of-words (or bag-of-characters, bag-of-relations)
 - Identify **unique** sufficient atomic sub-parts (e.g., words in a document)
 - Define simple features over these, e.g.,
 - Binary (0 or 1) → indicating presence
 - Natural numbers → indicating number of times in a context
 - Real-valued → various other score (we'll see examples throughout the semester)
2. Linguistically-inspired features
 - Define features from words, word spans, or linguistic-based annotations extracted from the document
3. Dense features via embeddings
 - Compute/extract a real-valued vector, e.g., from word2vec, ELMO, BERT, ...

Example: Document Classification via Bag-of-Words Features

Electronic alerts have been used to assist the authorities in moments of chaos and potential danger: after the Boston bombing in 2013, when the Boston suspects were still at large, and last month in Los Angeles, during an active shooter scare at the airport.

TECH

NOT TECH

Let's make a core assumption: the **label** can be predicted from **counts of individual word types**

Example: Document Classification via Bag-of-Words Features

Electronic alerts have been used to assist the authorities in moments of chaos and potential danger: after the Boston bombing in 2013, when the Boston suspects were still at large, and last month in Los Angeles, during an active shooter scare at the airport.

feature extraction

TECH

NOT TECH

With V word types, define V feature functions $f_i(x)$ as

$f_i(x)$ = # of times word type i appears in document x

Core assumption: the label can be predicted from counts of individual word types

Example: Document Classification via Bag-of-Words Features

Electronic alerts have been used to assist the authorities in moments of chaos and potential danger: after the Boston bombing in 2013, when the Boston suspects were still at large, and last month in Los Angeles, during an active shooter scare at the airport.

feature extraction

TECH

NOT TECH

With V word types, define V feature functions $f_i(x)$ as

$f_i(x)$ = # of times word type i appears in document x

$$f(x) = (f_i(x))_i^V$$

Core assumption: the label can be predicted from counts of individual word types

Example: Document Classification via Bag-of-Words Features

Electronic alerts have been used to assist the authorities in moments of chaos and potential danger: after the Boston bombing in 2013, when the Boston suspects were still at large, and last month in Los Angeles, during an active shooter scare at the airport.

TECH

NOT TECH

feature extraction

feature $f_i(x)$	value
alerts	1
assist	1
bombing	1
Boston	2
...	
sniffle	0
...	

Core assumption:
the label can be
predicted from
counts of individual
word types

Example: Document Classification with Bag-of-Words Features

Electronic alerts have been used to assist the authorities in moments of chaos and potential danger: after the Boston bombing in 2013, when the Boston suspects were still at large, and last month in Los Angeles, during an active shooter scare at the airport.

TECH

NOT TECH

feature	weight
alerts	.043
assist	-0.25
bombing	0.8
Boston	-0.00001
...	

w: weights

feature $f_i(x)$	value
alerts	1
assist	1
bombing	1
Boston	2
...	

f(x): “bag of words”

Text Annotation Tasks

1. Classify the entire document
2. Classify word tokens individually
3. Classify word tokens in a sequence
4. Identify phrases (“chunking”)
5. Syntactic annotation (parsing)
6. Semantic annotation

Text Annotation Tasks

1. Classify the entire document (“text categorization”)
2. Classify individual word tokens
3. Classify word tokens in a sequence
4. Identify phrases (“chunking”)
5. Syntactic annotation (parsing)
6. Semantic annotation

Text Classification

Assigning subject
categories, topics, or
genres

Spam detection

Authorship identification

Age/gender identification

Language Identification

Sentiment analysis

...

Text Classification

Assigning subject
categories, topics, or
genres

Spam detection

Authorship identification

Age/gender identification

Language Identification

Sentiment analysis

...

Input:

a document

a fixed set of classes $C = \{c_1, c_2, \dots, c_J\}$

Output: a predicted class c from C

Text Classification

Assigning subject
categories, topics, or
genres

Spam detection

Authorship identification

Age/gender identification

Language Identification

Sentiment analysis

...

Input:

a document linguistic blob

a fixed set of classes $C = \{c_1, c_2, \dots, c_J\}$

Output: a predicted class c from C

Text Classification: Hand-coded Rules?

Assigning subject
categories, topics, or
genres

Spam detection

Authorship identification

Age/gender identification

Language Identification

Sentiment analysis

...

Rules based on combinations of words or other features
spam: black-list-address OR (“dollars” AND “have been selected”)

Accuracy can be high
If rules carefully refined by expert

Building and maintaining these rules is expensive

Can humans faithfully assign uncertainty?

Text Classification: Supervised Machine Learning

Assigning subject
categories, topics, or
genres

Spam detection

Authorship identification

Age/gender identification

Language Identification

Sentiment analysis

...

Input:

a document d

a fixed set of classes $C = \{c_1, c_2, \dots, c_J\}$

A training set of m hand-labeled
documents $(d_1, c_1), \dots, (d_m, c_m)$

Output:

a learned classifier γ that maps
documents to classes

Text Classification: Supervised Machine Learning

Assigning subject
categories, topics, or
genres

Spam detection

Authorship identification

Age/gender identification

Language Identification

Sentiment analysis

...

Input:

a document d

a fixed set of classes $C = \{c_1, c_2, \dots, c_J\}$

A training set of m hand-labeled
documents $(d_1, c_1), \dots, (d_m, c_m)$

Output:

a learned classifier γ that maps
documents to classes

Naïve Bayes

Logistic regression

Support-vector
machines

k-Nearest Neighbors

...

Text Classification: Supervised Machine Learning

Assigning subject
categories, topics, or
genres

Spam detection

Authorship identification

Age/gender identification

Language Identification

Sentiment analysis

...

Input:

a document d

a fixed set of classes $C = \{c_1, c_2, \dots, c_J\}$

A training set of m hand-labeled
documents $(d_1, c_1), \dots, (d_m, c_m)$

Output:

a learned classifier γ that maps
documents to classes

Naïve Bayes

Logistic regression

Support-vector
machines

k-Nearest Neighbors

...

Text Annotation Tasks

1. Classify the entire document (“text categorization”)
2. **Classify individual word tokens**
3. Classify word tokens in a sequence
4. Identify phrases (“chunking”)
5. Syntactic annotation (parsing)
6. Semantic annotation

p(class | token in context)

Word Sense Disambiguation (WSD)

Problem:

The company said the *plant* is still operating ...

⇒ (A) Manufacturing plant or

⇒ (B) Living plant

Training Data: Build a special classifier just for tokens of “plant”

Sense	Context
(1) Manufacturing	... union responses to <i>plant</i> closures
” ”	... computer disk drive <i>plant</i> located in ...
” ”	company manufacturing <i>plant</i> is in Orlando ...
(2) Living	... animal rather than <i>plant</i> tissues can be ...
” ”	... to strain microscopic <i>plant</i> life from the ...
” ”	and Golgi apparatus of <i>plant</i> and animal cells

Test Data:

Sense	Context
???	... vinyl chloride monomer <i>plant</i> , which is ...
???	... molecules found in <i>plant</i> tissue from the ...

p(class | token in context)

WSD for Machine Translation (English → Spanish)

Problem:

... He wrote the last **sentence** two years later ...

⇒ *sentencia* (legal sentence) or

⇒ *frase* (grammatical sentence)

Training Data: Build a special classifier just for tokens of "sentence"

Translation	Context
(1) sentencia	... for a maximum <i>sentence</i> for a young offender ...
” ”	... of the minimum <i>sentence</i> of seven years in jail ...
” ”	... were under the <i>sentence</i> of death at that time ...
(2) frase	... read the second <i>sentence</i> because it is just as ...
” ”	... The next <i>sentence</i> is a very important ...
” ”	... It is the second <i>sentence</i> which I think is at ...

Test Data:

Translation	Context
???	... cannot criticize a <i>sentence</i> handed down by ...
???	... listen to this <i>sentence</i> uttered by a former ...

p(class | token in context)

Accent Restoration in Spanish & French

Problem:

Input: ... déjà travaille cote a cote ...



Output: ... déjà travaillé côte à côte ...

Examples:

... appeler l'autre **cote** de l'atlantique ...

⇒ *côté* (meaning side) or

⇒ *côte* (meaning coast)

... une famille des **pecheurs** ...

⇒ *pêcheurs* (meaning fishermen) or

⇒ *pécheurs* (meaning sinners)

p(class | token in context)

Accent Restoration in Spanish & French

Training Data:

Pattern	Context
(1) côté	... du laisser de <i>cote</i> faute de temps ...
” ”	... appeler l’ autre <i>cote</i> de l’ atlantique ...
” ”	... passe de notre <i>cote</i> de la frontiere ...
(2) côte	... vivre sur notre <i>cote</i> ouest toujours ...
” ”	... creer sur la <i>cote</i> du labrador des ...
” ”	travaillaient cote a <i>cote</i> , ils avaient ...

Test Data:

Pattern	Context
???	... passe de notre <i>cote</i> de la frontiere ...
???	... creer sur la <i>cote</i> du labrador des ...

p(class | token in context)

Capitalization Restoration

Problem:

... FRIED CHICKEN, **TURKEY** SANDWICHES AND FROZEN ...

⇒ *turkey* (the bird) or

⇒ *Turkey* (the country)

Training Data:

Capitalization	Context
(1) turkey	... OF FRIED CHICKEN , TURKEY SANDWICHES AND FROZEN ...
” ”	... NTS A POUND , WHILE TURKEY PRICES ROSE 1.2 CENTS ...
” ”	... PLAY , REAL GRADE-A TURKEY , WHICH ONLY A PRICE ...
(2) Turkey	... INUNDATED EASTERN TURKEY AFTER THE EARLIER ...
” ”	... FEELINGS TOWARD TURKEY SURFACED WHEN GREECE ...
” ”	... THE CONTRACT WITH TURKEY WILL PROVIDE OPPORTU...

Test Data:

Capitalization	Context
???	... NECK LIKE THAT OF A TURKEY ON A CHOPPING BLOCK ...
???	... PROBLEM IS THAT TURKEY IS NOT A EUROPEAN ...

p(class | token in context)

Text-to-Speech Synthesis

Problem:

... slightly elevated *lead* levels ...

⇒ *lɛd* (as in *lead mine*) or

⇒ *li:d* (as in *lead role*)

Training Data:

Pronunciation	Context
(1) <i>lɛd</i>	... it monitors the <i>lead</i> levels in drinking ...
” ”	... conference on <i>lead</i> poisoning in ...
” ”	... strontium and <i>lead</i> isotope zonation ...
(2) <i>li:d</i>	... maintained their <i>lead</i> Thursday over ...
” ”	... to Boston and <i>lead</i> singer for Purple ...
” ”	... Bush a 17-point <i>lead</i> in Texas , only 3 ...

Test Data:

Pronunciation	Context
???	... median blood <i>lead</i> concentration was ..
???	... his double-digit <i>lead</i> nationwide . The ...

p(class | token in context)

Spelling Correction

Problem:

... and he fired presidential **aid/aide** Dick Morris after ...

⇒ *aid* or

⇒ *aide*

Training Data:

Spelling	Context
(1) aid	... and cut the foreign <i>aid/aide</i> budget in fiscal 1996 ...
” ”	... they offered federal <i>aid/aide</i> for flood-ravaged states ...
(2) aide	... fired presidential <i>aid/aide</i> Dick Morris after ...
” ”	... and said the chief <i>aid/aide</i> to Sen. Baker, Mr. John ...

Test Data:

Spelling	Context
???	... said the longtime <i>aid/aide</i> to the Mayor of St. ...
???	... will squander the <i>aid/aide</i> it receives from the ...

What features? Example: “word to [the] left [of correction]”

Word to left	Frequency as Aid	Frequency as Aide
foreign	718	1
federal	297	0
western	146	0
provide	88	0
covert	26	0
oppose	13	0
future	9	0
similar	6	0
presidential	0	63
chief	0	40
longtime	0	26
aids-infected	0	2
sleepy	0	1
disaffected	0	1
indispensable	2	1
practical	2	0
squander	1	0

Spelling correction using an n-gram language model ($n \geq 2$) would use words to left and right to help predict the true word.

Similarly, an HMM would predict a word's class using classes to left and right.

But we'd like to throw in all kinds of other features, too ...

An assortment of possible cues ...

	Position	Collocation	led	li:d
N-grams (word, lemma, part-of-speech)	+1 L	lead <i>level/N</i>	219	0
	-1 W	<i>narrow</i> lead	0	70
	+1 W	lead <i>in</i>	207	898
	-1 W,+1 W	<i>of</i> lead <i>in</i>	162	0
	-1 W,+1 W	<i>the</i> lead <i>in</i>	0	301
	+1 P,+2 P	lead , < <i>NOUN</i> >	234	7
Wide-context collocations	$\pm k$ W	<i>zinc</i> (in $\pm k$ words)	235	0
	$\pm k$ W	<i>copper</i> (in $\pm k$ words)	130	0
Verb-object relationships	-V L	<i>follow/V</i> + lead	0	527
	-V L	<i>take/V</i> + lead	1	665

generates a whole bunch
of potential cues – use
data to find out which
ones work best

Word to left	Frequency as Aid	Frequency as Aide
foreign	718	1
federal	297	0
western	146	0
provide	88	0

An assortment of possible cues ...

	Position	Collocation	l:e:d	l:i:d
N-grams (word, lemma, part-of-speech)	+1 L	lead <i>level/N</i>	219	0
	-1 W	<i>narrow</i> lead	0	70
	+1 W	lead <i>in</i>	207	898
	-1 W,+1 W	<i>of</i> lead <i>in</i>	162	0
	-1 W,+1 W	<i>the</i> lead <i>in</i>	0	301
Wide-context collocations	$\pm k$ W	<i>zinc</i> (in $\pm k$ words)	235	0
	$\pm k$ W	<i>copper</i> (in $\pm k$ words)	130	0
Verb-object relationships	-V L	<i>follow/V</i> + lead	0	527
	-V L	<i>take/V</i> + lead	1	665

This feature is relatively weak, but weak features are still useful, especially since very few features will fire in a given context.

merged ranking
of all cues
of all these types

11.42	<i>follow/V</i> + lead	⇒ l:i:d
11.20	<i>zinc</i> (in $\pm k$ words)	⇒ l:e:d
11.16	lead <i>level/N</i>	⇒ l:e:d
10.66	<i>of</i> lead <i>in</i>	⇒ l:e:d
10.59	<i>the</i> lead <i>in</i>	⇒ l:i:d
10.51	lead <i>role</i>	⇒ l:i:d

Final decision list for *lead* (abbreviated)

List of all features,
ranked by their weight.

(These weights are for a simple
"decision list" model where the
single highest-weighted feature
that fires gets to make the
decision all by itself.

However, a log-linear model,
which adds up the weights of all
features that fire, would be
roughly similar.)

LogL	Evidence	Pronunciation
11.40	<i>follow/V + lead</i>	⇒ li:d
11.20	<i>zinc</i> (in $\pm k$ words)	⇒ lɛd
11.10	<i>lead level/N</i>	⇒ lɛd
10.66	<i>of lead in</i>	⇒ lɛd
10.59	<i>the lead in</i>	⇒ li:d
10.51	<i>lead role</i>	⇒ li:d
10.35	<i>copper</i> (in $\pm k$ words)	⇒ lɛd
10.28	<i>lead time</i>	⇒ li:d
10.24	<i>lead levels</i>	⇒ lɛd
10.16	<i>lead poisoning</i>	⇒ lɛd
8.55	<i>big lead</i>	⇒ li:d
8.49	<i>narrow lead</i>	⇒ li:d
7.76	<i>take/V + lead</i>	⇒ li:d
5.99	<i>lead , NOUN</i>	⇒ lɛd
1.15	<i>lead in</i>	⇒ li:d
	ooo	

Text Annotation Tasks


1. Classify the entire document (“text categorization”)
2. Classify individual word tokens
3. **Classify word tokens in a sequence**
4. Identify phrases (“chunking”)
5. Syntactic annotation (parsing)
6. Semantic annotation

Part of Speech Tagging

- We could treat tagging as a token classification problem
 - Tag each word independently given features of context
 - And features of the word's spelling (suffixes, capitalization)


Sequence Labeling as Classification

Classify each token independently but use as input features, information about the surrounding tokens (sliding window).




Sequence Labeling as Classification

Classify each token independently but use as input features, information about the surrounding tokens (sliding window).



Sequence Labeling as Classification


Classify each token independently but use as input features, information about the surrounding tokens (sliding window).



Sequence Labeling as Classification

Classify each token independently but use as input features, information about the surrounding tokens (sliding window).


John saw the saw and decided to take it to the table.



Sequence Labeling as Classification

Classify each token independently but use as input features, information about the surrounding tokens (sliding window).


John saw the saw and decided to take it to the table.



Sequence Labeling as Classification

Classify each token independently but use as input features, information about the surrounding tokens (sliding window).


John saw the saw and decided to take it to the table.



Sequence Labeling as Classification

Classify each token independently but use as input features, information about the surrounding tokens (sliding window).


John saw the saw and decided to take it to the table.



Sequence Labeling as Classification

Classify each token independently but use as input features, information about the surrounding tokens (sliding window).


John saw the saw and decided to take it to the table.



Sequence Labeling as Classification

Classify each token independently but use as input features, information about the surrounding tokens (sliding window).


John saw the saw and decided to take it to the table.



Sequence Labeling as Classification

Classify each token independently but use as input features, information about the surrounding tokens (sliding window).


John saw the saw and decided to take it to the table.



Sequence Labeling as Classification


Classify each token independently but use as input features, information about the surrounding tokens (sliding window).

John saw the saw and decided to take it to the table.



Sequence Labeling as Classification


Classify each token independently but use as input features, information about the surrounding tokens (sliding window).



Part of Speech Tagging

We'll see
HMMs
later!

Or we could use an HMM:



probs
from tag
bigram
model



Start PN Verb Det Noun Prep Noun Pr

probs from
unigram
replacement



Bill directed a cortege of autos thr


Part of Speech Tagging

- We could treat tagging as a token classification problem
 - Tag each word independently given features of context
 - And features of the word's spelling (suffixes, capitalization)
- Or we could use an HMM:
 - The point of the HMM is basically that the tag of one word might depend on the tags of adjacent words.
- Combine these two ideas??
 - We'd like rich features (e.g., in a **log-linear model**), but we'd also like our feature functions to depend on adjacent tags.
 - So, the problem is to predict **all** tags together.

Supervised Learning Methods


- Easy to build a “yes” or “no” predictor from supervised training data
 - Plenty of software packages to do the learning & prediction
 - Lots of people in NLP never go beyond this 😊
- Similarly, easy to build a system that chooses from a small finite set
 - Basically the same deal
 - But runtime goes up linearly with the size of the set, unless you’re clever (HW3)
- Harder to predict the best string or tree (set is exponentially large or infinite)

Can We {Do Better? {Be More Expressive?



See: CMSC 678 or CMSC 691
(Prob. & Graphical ML)

Can We Use Neural, Recurrent Methods?



Supervised Learning Methods

- Easy to build a “yes” or “no” predictor from supervised training data
 - Plenty of software packages to do the learning & prediction
 - Lots of people in NLP never go beyond this 😊
- Similarly, easy to build a system that chooses from a small finite set
 - Basically the same deal
 - But runtime goes up linearly with the size of the set, unless you’re clever (HW3)
- Harder to predict the best string or tree (set is exponentially large or infinite)
 - In the best case, requires dynamic programming; you might have to write your own code
 - But finite-state or CRF toolkits may find the best string for you
 - And you could modify someone else’s parser to pick the best tree
 - An algorithm for picking the best can usually be turned into a learning algorithm
 - You may need to rely on **approximate** solutions (e.g., via **beam search**)

Text Annotation Tasks

1. Classify the entire document (“text categorization”)
2. Classify individual word tokens
3. Classify word tokens in a sequence
4. Identify phrases (“chunking”)
5. Syntactic annotation (parsing)
6. Semantic annotation

Example: Finding Named Entities

Named entity recognition (NER)

Identify proper names in texts, and classification into a set of predefined categories of interest

- Person names

- Organizations (companies, government organisations, committees, etc)

- Locations (cities, countries, rivers, etc)

- Date and time expressions

- Measures (percent, money, weight etc), email addresses, Web addresses, street addresses, etc.

- Domain-specific: names of drugs, medical conditions, names of ships, bibliographic references etc.

Named Entity Recognition

CHICAGO (AP) — Citing high fuel prices, **United Airlines** said Friday it has increased fares by \$6 per round trip on flights to some cities also served by lower-cost carriers. **American Airlines**, a unit **AMR**, immediately matched the move, spokesman **Tim Wagner** said. **United**, a unit of **UAL**, said the increase took effect Thursday night and applies to most routes where it competes against discount carriers, such as **Chicago** to **Dallas** and **Atlanta** and **Denver** to **San Francisco**, **Los Angeles** and **New York**.

NE Types

Type	Tag	Sample Categories
People	PER	Individuals, fictional characters, small groups
Organization	ORG	Companies, agencies, political parties, religious groups, sports teams
Location	LOC	Physical extents, mountains, lakes, seas
Geo-Political Entity	GPE	Countries, states, provinces, counties
Facility	FAC	Bridges, buildings, airports
Vehicles	VEH	Planes, trains, and automobiles

Type	Example
People	<i>Turing</i> is often considered to be the father of modern computer science.
Organization	The <i>IPCC</i> said it is likely that future tropical cyclones will become more intense.
Location	The <i>Mt. Sanitas</i> loop hike begins at the base of <i>Sunshine Canyon</i> .
Geo-Political Entity	<i>Palo Alto</i> is looking at raising the fees for parking in the University Avenue district.
Facility	Drivers were advised to consider either the <i>Tappan Zee Bridge</i> or the <i>Lincoln Tunnel</i> .
Vehicles	The updated <i>Mini Cooper</i> retains its charm and agility.

Example: Information Extraction

As a task:

Filling slots in a database from sub-segments of text.


October 14, 2002, 4:00 a.m. PT

For years, [Microsoft Corporation](#) [CEO](#) [Bill Gates](#) railed against the economic philosophy of open-source software with Orwellian fervor, denouncing its communal licensing as a "cancer" that stifled technological innovation.

Today, Microsoft claims to "love" the open-source concept, by which software code is made public to encourage improvement and development by outside programmers. Gates himself says Microsoft will gladly disclose its crown jewels--the coveted code behind the Windows operating system--to select customers.

"We can be open source. We love the concept of shared source," said [Bill Veghte](#), a [Microsoft](#) [VP](#). "That's a super-important shift for us in terms of code access."

[Richard Stallman](#), [founder](#) of the [Free Software Foundation](#), countered saying...



NAME	TITLE	ORGANIZATION
Bill Gates	CEO	Microsoft
Bill Veghte	VP	Microsoft
Richard Stallman	founder	Free Soft..

Phrase Types to Identify for IE

Closed set

U.S. states

He was born in Alabama...

The big Wyoming sky...

Regular set

U.S. phone numbers

Phone: (413) 545-1323

The CALD main office can be reached at 412-268-1299

Complex pattern

U.S. postal addresses

University of Arkansas
P.O. Box 140
Hope, AR 71802

Headquarters:
1128 Main Street, 4th Floor
Cincinnati, Ohio 45210

Ambiguous patterns, needing context and many sources of evidence

Person names

...was among the six houses
sold by Hope Feldman that year.

Pawel Opalinski, Software
Engineer at WhizBang Labs.

Identifying phrases

- A key step in IE is to identify relevant phrases
 - Named entities
 - As on previous slides
 - Relationship phrases
 - “said”, “according to”, ...
 - “was born in”, “hails from”, ...
 - “bought”, “hopes to acquire”, “formed a joint agreement with”, ...
 - Simple syntactic chunks (e.g., non-recursive NPs)
 - “Syntactic chunking” sometimes done before (or instead of) parsing
 - Also, “segmentation”: divide Chinese text into words (no spaces)
- So, how do we learn to mark phrases?

Reduce to a tagging problem ...

- The IOB encoding (Ramshaw & Marcus 1995):
 - B_X = “beginning” (first word of an X)
 - I_X = “inside” (non-first word of an X)
 - O = “outside” (not in any phrase)
 - Does not allow overlapping or recursive phrases

...United Airlines said Friday it has increased ...

B_ORG I_ORG O O O O

... the move , spokesman Tim Wagner said ...

O O O O B_PER I_PER O

What if this were tagged as B_ORG instead?

Example applications for IE

- Classified ads
- Restaurant reviews
- Bibliographic citations
- Appointment emails
- Legal opinions
- Papers describing clinical medical studies
- ...

Text Annotation Tasks

1. Classify the entire document (“text categorization”)
2. Classify individual word tokens
3. Classify word tokens in a sequence
4. Identify phrases (“chunking”)
5. Syntactic annotation (parsing)
6. Semantic annotation


Context Free Grammar

$S \rightarrow NP VP$ $PP \rightarrow P NP$
 $NP \rightarrow Det Noun$ $AdjP \rightarrow Adj Noun$
 $NP \rightarrow Noun$ $VP \rightarrow V NP$
 $NP \rightarrow Det AdjP$ $Noun \rightarrow Baltimore$
 $NP \rightarrow NP PP$...

Set of rewrite rules, comprised of terminals and non-terminals

Generate from a Context Free Grammar


$S \rightarrow NP VP$ $PP \rightarrow P NP$
 $NP \rightarrow Det Noun$ $AdjP \rightarrow Adj Noun$
 $NP \rightarrow Noun$ $VP \rightarrow V NP$
 $NP \rightarrow Det AdjP$ $Noun \rightarrow Baltimore$
 $NP \rightarrow NP PP$...



Baltimore is a great city

Assign Structure (**Parse**) with a Context Free Grammar

$S \rightarrow NP VP$ $PP \rightarrow P NP$
 $NP \rightarrow Det Noun$ $AdjP \rightarrow Adj Noun$
 $NP \rightarrow Noun$ $VP \rightarrow V NP$
 $NP \rightarrow Det AdjP$ $Noun \rightarrow Baltimore$
 $NP \rightarrow NP PP$...



Baltimore is a great city

[_S [_{NP} [_{Noun} Baltimore]] [_{VP} [_{Verb} is] [_{NP} a great city]]]

bracket notation

(S (NP (Noun Baltimore))
(VP (V is)
(NP a great city)))

S-expression

```
T = Cell[N][N+1]
```


```
for(j = 1; j ≤ N; ++j) {  
    T[j-1][j].add(X for non-terminal X in G if X → wordj)  
}
```

```
for(width = 2; width ≤ N; ++width) {  
    for(start = 0; start < N - width; ++start) {  
        end = start + width  
        for(mid = start+1; mid < end; ++mid) {  
            for(non-terminal Y : T[start][mid]) {  
                for(non-terminal Z : T[mid][end]) {  
                    T[start][end].add(X for rule X → Y Z : G)  
                }  
            }  
        }  
    }  
}
```



You may get to use a
dynamic program



You can get **dependency-based** syntactic forms



And Go From Syntax to Shallow Semantics



“Open
Information
Extraction”

(input)
↓
she took the midnight train going anywhere
Born in a small town, she took the midnight train
Born in a town, she took the midnight train

she took the midnight train
she took midnight train
...

Angeli et al. (2015)

(she; took; midnight train)

(extracted clause)
↓
she Born in small town
she Born in a town
she Born in town
↓
(she; born in; small town)
(she; born in; town)

<http://corenlp.run/> (constituency & dependency)

<https://github.com/hltcoe/predpatt>

a sampling of
efforts

<http://openie.allenai.org/>

<http://www.cs.rochester.edu/research/knext/browse/> (constituency trees)

<http://rtw.ml.cmu.edu/rtw/>

Supervised Learning Methods

- Easy to build a “yes” or “no” predictor from supervised training data
 - Plenty of software packages to do the learning & prediction
 - Lots of people in NLP never go beyond this 😊
- Similarly, easy to build a system that chooses from a small finite set
 - Basically the same deal
 - But runtime goes up linearly with the size of the set, unless you’re clever (HW3)
- Harder to predict the best string or tree (set is exponentially large or infinite)
 - In the best case, requires dynamic programming; you might have to write your own code
 - But finite-state or CRF toolkits may find the best string for you
 - And you could modify someone else’s parser to pick the best tree
 - An algorithm for picking the best can usually be turned into a learning algorithm
 - You may need to rely on **approximate** solutions (e.g., via **beam search**)
- Hardest if your features look at “non-local” properties of the string or tree
 - Now dynamic programming won’t work (or will be something awful like $O(n^9)$)
 - You need some kind of approximate search
 - Can be harder to turn approximate search into a learning algorithm
 - Still, this is a standard preoccupation of machine learning (“structured prediction,” “graphical models”)

Text Annotation Tasks

1. Classify the entire document (“text categorization”)
2. Classify individual word tokens
3. Classify word tokens in a sequence
4. Identify phrases (“chunking”)
5. Syntactic annotation (parsing)
6. **Semantic annotation**

Semantic Role Labeling (SRL)

- For each predicate (e.g., verb)
 1. find its arguments (e.g., NPs)
 2. determine their **semantic roles**

John drove Mary from Austin to Dallas in his Toyota Prius.

The hammer broke the window.

- **agent**: Actor of an action
- **patient**: Entity affected by the action
- **source**: Origin of the affected entity
- **destination**: Destination of the affected entity
- **instrument**: Tool used in performing action.
- **beneficiary**: Entity for whom action is performed

As usual, can solve as classification ...

- Consider one verb at a time: “bit”
- Classify the role (if any) of each of the 3 NPs

Color Code:

not-a-role

agent


patient

source

destination

instrument

beneficiary




Parse tree paths as classification features

Path feature is

V ↑ **VP** ↑ **S** ↓ **NP**

which tends to
be associated
with **agent** role




Parse tree paths as classification features

Path feature is

V ↑ VP ↑ S ↓ NP ↓ PP ↓ NP

which tends to
be associated
with **no role**



Head words as features

- Some roles prefer to be filled by certain kinds of NPs.
- This can give us useful features for classifying accurately:
 - “**John** ate the spaghetti with chopsticks.” (instrument)
“**John** ate the spaghetti with meatballs.” (patient)
“**John** ate the spaghetti with Mary.”
 - Instruments should be tools
 - Patient of “eat” should be edible
 - “**John** bought the car for \$21K.” (instrument)
“**John** bought the car for Mary.” (beneficiary)
 - Instrument of “buy” should be Money
 - Beneficiaries should be animate (things with desires)
 - “**John** drove Mary to school in the van”
“**John** drove the van to work with Mary.”
 - What do you think?

Uses of Semantic Roles

- Find the answer to a user's question
 - “Who” questions usually want Agents
 - “What” question usually want Patients
 - “How” and “with what” questions usually want Instruments
 - “Where” questions frequently want Sources/Destinations.
 - “For whom” questions usually want Beneficiaries
 - “To whom” questions usually want Destinations
- Generate text
 - Many languages have specific syntactic constructions that must or should be used for specific semantic roles.
- Word sense disambiguation, using selectional restrictions
 - The **bat** ate the **bug**. (what kind of bat? what kind of bug?)
 - Agents (particularly of “eat”) should be animate – animal bat, not baseball bat
 - Patients of “eat” should be edible – animal bug, not software bug
 - John **fired** the secretary.
John **fired** the rifle.
Patients of fire₁ are different than patients of fire₂


Other Current Semantic Annotation Tasks

(similar to SRL)

- PropBank – coarse-grained roles of verbs
- NomBank – similar, but for nouns
- FrameNet – fine-grained roles of any word
- TimeBank – temporal expressions

FrameNet Example

REVENGE FRAME



FrameNet Example

REVENGE FRAME

*triggering words and phrases
(not limited to verbs)*

avenge, revenge, retaliate, get back at, pay back, get even, ...

revenge, vengeance, retaliation, retribution, reprisal, ...

vengeful, retaliatory, retributive; in revenge, in retaliation, ...

take revenge, wreak vengeance, exact retribution, ...

FN2_Desktop

MainActionWindow

Remainder

Removing

Render_nonfunctional

Reparation

Reporting

Request

Reshaping

Residence

Rest

Revenge

- Avenger <F1>
- Injured_Party <>
- Injury <F3>
- Offender <F3>
- Punishment <F12>
- Degree <G>
- Instrument <F3>
- Manner <M>
- Place <F3>
- Time <F2>
- Depictive <D>
- Purpose <F4>
- Result <E>

- avenge.v
- Lemma(V)
- rcoll-brother [1/1]
- rcoll-death [5/12]
 - It will do no good t
 - With this , El Cid a
 - His secret ambition
 - For his distraught f
 - In Article 3 of the
 - The nausea threatene
 - Suddenly he walked b
 - In Scaramouche the m
 - Are you planning t
 - To avenge the death
 - The Trojans wish to
 - Did someone in this
- rcoll-defeat [5/16]
- rcoll-father [0/3]
- rcoll-murder [2/4]
- np-ppagainst [0/1]
- np-ppfor [1/2]
- np-ppon [2/5]

SubCorpus Editor: V-429-s20-rcoll-death (77339)

It will do no good to AVENGE my death by killing him . "

With this , El Cid at once AVENGED the death of his son and once again showed that any attempt to reconquer Valencia was fruitless while he still lived . DNI

His secret ambition was for the Argentine ban to be lifted so he could get to England and AVENGE Pedro 's death by taking out the English and especially one poker-faced Guards Officer . DNI

For his distraught family , only hanging would have AVENGED the death of the father of four .

In Article 3 of the agreement , each had promised to AVENGE the violent death of the other with the blood of the murderer . DNI

Layer	W	i	t	h	,	E	l	C	i	d	a	t	o	n	c	e	a	v	e	n	g	e	t	h	e	d	e	a	t	h	o	f	h	i	s	s	o	n
FE	P	u	n	i	s		A	v	e	n	g	e												I	n	j	u	r	y									
GF	C	o	m	p		E	x	t															O	b	j													
PT	P	P				N	P																N	P														
Other																																						
Verb																																						
Sent																																						

FE

GF

PT

Other

Verb

Sent

Appositive

Comp <F3>

Ext <F1>

Gen <F5>

Head <F4>

Mod <F6>

Obj <F2>

Quant

Slide courtesy Jason Eisner

Generating new text

1. Speech recognition (transcribe as text)
2. Machine translation
3. Text generation from semantics
4. Inflect, analyze, or transliterate words
5. Single- or multi-doc summarization

Deeper Information Extraction

1. Coreference resolution (within a document)
2. Entity linking (across documents)
3. Event extraction and linking
4. Knowledge base population (KBP)
5. Recognizing textual entailment (RTE)

User interfaces

1. Dialogue systems
 - Personal assistance
 - Human-computer collaboration
 - Interactive teaching
2. Language teaching; writing help
3. Question answering
4. Information retrieval

Multimodal interfaces or modeling

1. Sign languages
2. Speech + gestures
3. Images + captions
4. Brain recordings, human reaction times

NLP automates things that humans do well, so that they can be done automatically on more sentences. But this slide is about language analysis that's hard even for humans. Computational linguistics (like comp bio, etc.) can discover underlying patterns in large datasets: things we didn't know!

Discovering Linguistic Structure

1. Decipherment
2. Grammar induction
3. Topic modeling
4. Deep learning of word meanings
5. Language evolution (historical linguistics)
6. Grounded semantics

Today's Learning Goals

- Define featurization and some examples
- Learn about NLP Tasks at a high-level:
 - Document classification
 - Part of speech tagging
 - Syntactic parsing
 - Entity id/coreference