# Language Modeling (and the Noisy Channel) Introduction to Natural Language Processing 

Dr. Jan Hajič

CS Dept., Johns Hopkins Univ.
hajic@cs.jhu.edu
www.cs.jhu.edu/ hajic
June 26, 2014

## The Noisy Channel

- Prototypical case

- Model: probability of error (noise):
- Example: $p(0 \mid 1)=.3 p(1 \mid 1)=.7 p(1 \mid 0)=.4 p(0 \mid 0)=.6$
- The task:
known: the noisy output; want to know; the input (decoding)


## Noisy Channel Applications

- OCR
- straightforward: text $\rightarrow$ print (adds noise), scan $\rightarrow$ image
- Handwriting recognition
- text $\rightarrow$ neurons, muscles ("noise"), scan/digitize $\rightarrow$ image
- Speech recognition (dictation, commands, etc.)
- text $\rightarrow$ conversion to acoustic signal ("noise") $\rightarrow$ acoustic waves
- Machine Translation
- text in target language $\rightarrow$ translation ("noise") $\rightarrow$ source language
- Also: Part of Speech Tagging
- sequence of tags $\rightarrow$ selection of word forms $\rightarrow$ text


## The Golden Rule of OCR, ASR, HR, MT,...

- Recall:

$$
\begin{aligned}
& \mathrm{p}(\mathrm{~A} \mid \mathrm{B})=\mathrm{p}(\mathrm{~B} \mid \mathrm{A}) \quad \mathrm{p}(\mathrm{~A}) / \mathrm{p}(\mathrm{~B}) \quad \text { (Bayes formula) } \\
& A_{\text {best }}=\operatorname{argmax}_{\mathrm{A}} \mathrm{p}(\mathrm{~B} \mid \mathrm{A}) \mathrm{p}(\mathrm{~A}) \quad(\text { The Golden Rule })
\end{aligned}
$$

- $p(B \mid A)$ : the acoustic/image/translation/lexical model
- application-specific name
- will explore later
- $\mathrm{p}(\mathrm{A})$ : language model


## The Perfect Language Model

- Sequence of word forms (forget about tagging for the moment)
- Notation: $\mathrm{A} \sim \mathrm{W}=\left(w_{1}, w_{2}, w_{3}, \ldots, w_{d}\right)$
- The big (modeling) question:

$$
\mathrm{p}(\mathrm{~W})=?
$$

- Well, we know (Bayes/chain rule) $\rightarrow$ ):

$$
\begin{gathered}
\mathrm{p}(\mathrm{~W})=\mathrm{p}\left(w_{1}, w_{2}, w_{3}, \ldots, w_{d}\right)= \\
p\left(w_{1}\right) \times p\left(w_{2} \mid w_{1}\right) \times p\left(w_{3} \mid w_{1}, w_{2}\right) \times \ldots \times p\left(w_{d} \mid w_{1}, w_{2}, \ldots w_{d-1}\right)
\end{gathered}
$$

- Not practical (even short W $\rightarrow$ too many parameters)


## Markov Chain

- Unlimited memory (cf. previous foil):
- for $w_{i}$ we know all its predecessors $w_{1}, w_{2}, w_{3}, \ldots, w_{i-1}$
- Limited memory:
- we disregard "too old" predecessors
- remember only $k$ previous words: $w_{i-k}, w_{i-k+1}, \ldots, w_{i-1}$
- called " $k^{\text {th }}$ order Markov approximation"
-     + stationary character (no change over time):

$$
\mathrm{p}(\mathrm{~W}) \cong \prod_{i=1 . . d} p\left(w_{i} \mid w_{i-k}, w_{i-k+1}, \ldots, w_{i-1}\right), d=|W|
$$

## n-gram Language Models

- $(n-1)^{\text {th }}$ order Markov approximation $\rightarrow \mathrm{n}$-gram LM:

$$
\begin{gathered}
\text { prediction history } \\
\mathrm{p}(\mathrm{~W})={ }_{d f} \prod_{i=1 . . d} p\left(w_{i} \mid w_{i-n+1}, w_{i-n+2}, \ldots, w_{i-1}\right)
\end{gathered}
$$

- In particular (assume vocabulary $|\mathrm{V}|=60 \mathrm{k}$ ):

0-gram LM: uniform model, $\quad p(w)=1 /|\mathrm{V}|, \quad 1$ parameter 1-gram LM: unigram model, $p(w)$,
2-gram LM: bigram model,
$\mathrm{p}\left(w_{i} \mid w_{i-1}\right)$,
$6 \times 10^{4}$ parameters
3-gram LM: trigram model, $\quad \mathrm{p}\left(w_{i} \mid w_{i-2}, w_{i-1}\right)$,
$3.6 \times 10^{9}$ parameters
$2.16 \times 10^{14}$ parameters

## LM: Observations

- How large $n$ ?
- nothing in enough (theoretically)
- but anyway: as much as possible ( $\rightarrow$ close to "perfect" model)
- empirically: 3
- parameter estimation? (reliability, data availability, storage space, ...)
- 4 is too much: $|\mathrm{V}|=60 \mathrm{k} \rightarrow 1.296 \times 10^{19}$ parameters
- but: 6-7 would be (almost) ideal (having enough data): in fact, one can recover original from 7-grams!
- Reliability ~(1/Detail) ( $\rightarrow$ need compromise)
- For now, keep word forms (no "linguistic" processing)


## The Length Issue

- $\forall n ; \Sigma_{w \in \Omega^{n}} p(w)=1 \Rightarrow \Sigma_{n=1 . . \infty} \Sigma_{w \in \Omega^{n}} p(w) \gg 1(\rightarrow \infty)$
- We want to model all sequences of words
- for "fixed" length tasks: no problem - n fixed, sum is 1
- tagging, OCR/handwriting (if words identified ahead of time)
- for "variable" length tasks: have to account for
- discount shorter sentences
- General model: for each sequence of words of length n , define $\mathrm{p}^{\prime}(\mathrm{w})=\lambda_{n} p(w)$ such that $\Sigma_{n-1 . . \infty} \lambda_{n}=1 \Rightarrow$

$$
\Sigma_{n=1 . . \infty} \Sigma_{w \in \Omega^{n}} p^{\prime}(w)=1
$$

e.g. estimate $\lambda_{n}$ from data; or use normal or other distribution

## Parameter Estimation

- Parameter: numerical value needed to compute $p(w \mid h)$
- From data (how else?)
- Data preparation:
- get rid of formating etc. ("text cleaning")
- define words (separate but include punctuation, call it "word")
- define sentence boundaries (insert "words" <s> and </s>)
- letter case: keep, discard, or be smart:
- name recognition
- number type identification (these are huge problems per se!)
- numbers: keep, replace by <num>, or be smart (form ~ punctuation)


## Maximum Likelihood Estimate

- MLE: Relative Frequency...
- ...best predicts the data at hand (the "training data")
- Trigrams from training Data T:
- count sequences of three words in T: $c_{3}\left(w_{i-2}, w i-1, w_{i}\right)$
- (NB: notation: just saying that three words follow each other)
- count sequences of two words in T: $c_{2}\left(w i-1, w_{i}\right)$
- either use $c_{2}(y, z)=\Sigma_{w} c_{3}(y, z, w)$
- or count differently at the beginning (\& end) of the data!

$$
p\left(w_{i} \mid w_{i-2}, w_{i-1}\right)=\text { est. } c_{3}\left(w_{i-2}, w_{i-1}, w_{i}\right) / c_{2}\left(w_{i-2}, w_{i-1}\right)!
$$

## Character Language Model

- Use individual characters instead of words:

$$
p(W)={ }_{d f} \Pi_{i=1 . . d} p\left(c_{i} \mid c_{i-n+1}, c_{i-n+2}, \ldots, c_{i}\right)
$$

- Same formulas etc.
- Might consider 4-grams, 5-grams or even more
- Good only for language comparison)
- Transform cross-entropy between letter- and word-based models:

$$
H_{S}\left(p_{c}\right)=H_{S}\left(p_{w}\right) / \text { avg. \# of characters/word in S }
$$

## LM: an Example

- Training data:
<s> <s> He can buy the can of soda.
- Unigram:

$$
\begin{aligned}
& p_{1}(\mathrm{He})=p_{1}(\text { buy })=p_{1}(\text { the })=p_{1}(\text { of }) p_{1}(\text { soda })=p_{1}(.)=.125 \\
& p_{1}(\text { buy })=.25
\end{aligned}
$$

- Bigram:
$p_{2}(\mathrm{He} \mid<\mathrm{s}>)=1, p_{2}($ can $\mid \mathrm{He})=1, p_{2}($ buy $\mid$ can $)=.5, p_{2}($ of $\mid$ can $)=.5$,
$p_{2}$ (the|buy) $=1, \ldots$
- Trigram:
$p_{3}(\mathrm{He} \mid<s>,<s>)=1, p_{3}($ can $\mid<s>, \mathrm{He})=1, p_{3}($ buy $\mid \mathrm{He}$, can $)=1$,
$p_{3}\left(\right.$ of $\mid$ the, can $=1, \ldots, p_{3}($. |of,soda $)=1$.
- Entropy:

$$
\mathrm{H}\left(p_{1}\right)=2.75, \mathrm{H}\left(p_{2}\right)=.25, \mathrm{H}\left(p_{3}\right)=0 \leftarrow \text { Great?! }
$$

## LM: an Example (The Problem)

- Cross-entropy:
- $S=<s><s>$ It was the greatest buy of all.
- Even $H_{S}\left(p_{1}\right)$ fails $\left(=H_{S}\left(p_{2}\right)=H_{S}\left(p_{3}\right)=\infty\right)$, because:
- all unigrams but $p_{1}$ (the), $p_{1}$ (buy), $p_{1}$ (of) and $p_{1}($.$) are 0$.
- all bigram probabilities are 0 .
- all trigram probabilities are 0.
- We want: to make all (theoretically possible*) probabilities non-zero.
*in fact, all: remeber our graph from day1?

