

Language Modelling

Marcello Federico
FBK-irst Trento, Italy

MT Marathon, Edinburgh, 2012

- Role of LM in ASR and MT
- N-gram Language Models
- Evaluation of Language Models
- Smoothing Schemes
- Discounting Methods
- Class-based LMs
- Maximum-Entropy LMs
- Neural Network LMs
- Toolkits and ARPA file format

Goal: find the words \mathbf{w}^* in a speech signal \mathbf{x} such that:

$$\mathbf{w}^* = \arg \max_{\mathbf{w}} \Pr(\mathbf{x} | \mathbf{w}) \Pr(\mathbf{w}) \quad (1)$$

Problems:

- **language modeling** (LM): estimating $\Pr(\mathbf{w})$
- **acoustic modeling** (AM): estimating $\Pr(\mathbf{x} | \mathbf{w})$
- **search problem:** computing (1)

AM sums over hidden state sequences \mathbf{s} a Markov process of (\mathbf{x}, \mathbf{s}) from \mathbf{w}

$$\Pr(\mathbf{x} | \mathbf{w}) = \sum_{\mathbf{s}} \Pr(\mathbf{x}, \mathbf{s} | \mathbf{w})$$

Hidden Markov Model: hidden states "link" speech frames to words.

Goal: find the English string e translating the foreign text f such that:

$$e^* = \arg \max_e \Pr(f | e) \Pr(e) \quad (2)$$

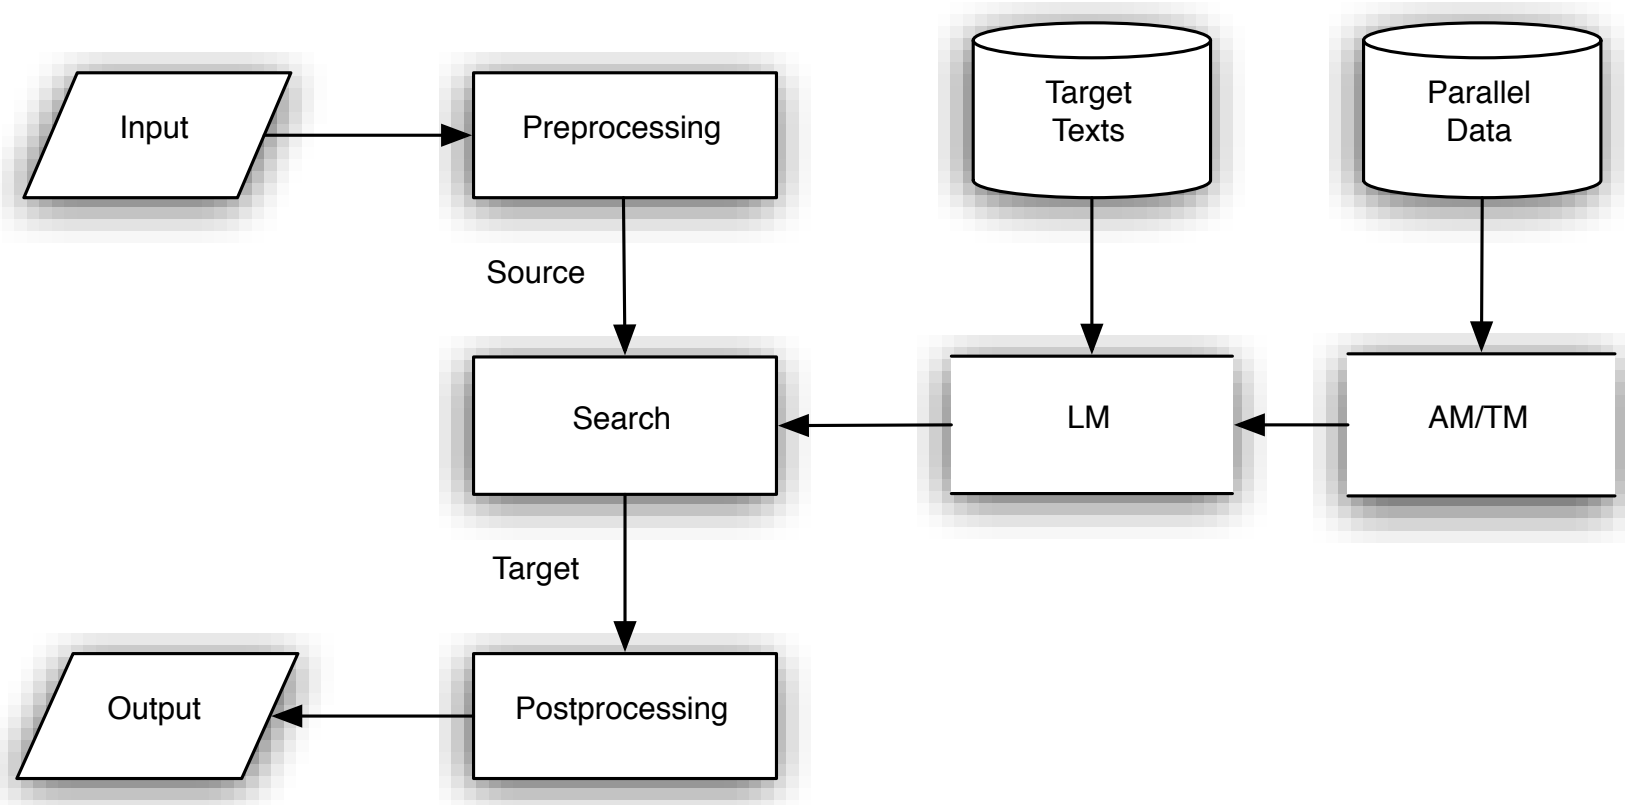
Problems:

- **language modeling** (LM): estimating $\Pr(e)$
- **translation modeling** (TM): estimating $\Pr(f | e)$
- **search problem:** computing (2)

TM sums over hidden alignments a a stochastic process generating (f, a) from e .

$$\Pr(f | e) = \sum_a \Pr(f, a | e)$$

Alignment Models: hidden alignments "link" foreign words with English words.



- **Parallel data** are samples of observations (\mathbf{x}, \mathbf{w}) and (\mathbf{f}, \mathbf{e})
- AM and TM can be **machine-learned** without observing \mathbf{s} and \mathbf{a} .

- Translation hypotheses are ranked by:

$$\mathbf{e}^* = \arg \max_{\mathbf{e}, \mathbf{a}} \sum_i \lambda_i h_i(\mathbf{e}, \mathbf{f}, \mathbf{a})$$

- **Phrases** are finite strings (cf. n-grams)
- Hidden variable \mathbf{a} embeds:
 - **segmentation** of \mathbf{f} and \mathbf{e} into phrases
 - **alignment** of phrases of \mathbf{f} with phrases of \mathbf{e}
- **Feature functions** $h_k()$ include:
 - Translation Model: appropriateness of phrase-pairs
 - Distortion Model: word re-ordering
 - **Language Model**: fluency of target string
 - Length Model: number of target words
- **Role of the LM** is exactly the same as for the noisy channel approach:
 - **to score translations incrementally** generated by the search algorithm!

Given a text $\mathbf{w} = w_1 \dots, w_t, \dots, w_{|\mathbf{w}|}$ we can compute its probability by:

$$\Pr(\mathbf{w}) = \Pr(w_1) \prod_{t=2}^{|\mathbf{w}|} \Pr(w_t | h_t) \quad (3)$$

where $h_t = w_1, \dots, w_{t-1}$ indicates the **history** of word w_t .

- $\Pr(w_t | h_t)$ becomes difficult to estimate as the history h_t grows .
- hence, we take the n -gram **approximation** $h_t \approx w_{t-n+1} \dots w_{t-1}$

e.g. Full history: $\Pr(\text{Parliament} | \text{I declare resumed the session of the European})$

3-gram : $\Pr(\text{Parliament} | \text{the European})$

The choice of n determines the complexity of the LM (# of parameters):

- **bad**: no magic recipe about the optimal order n for a given task
- **good**: language models can be evaluated quite cheaply

- Extrinsic: **impact on task** (e.g. BLEU score for MT)
- Intrinsic: capability of **predicting words**

The **perplexity** (PP) measure is defined as: ¹

$$PP = 2^{LP} \quad \text{where} \quad LP = -\frac{1}{|\mathbf{w}|} \log_2 p(\mathbf{w}) \quad (4)$$

- \mathbf{w} is a **sufficiently long test sample** and $p(\mathbf{w})$ is the LM probability.

Properties:

- $0 \leq PP \leq |V|$ (size of the vocabulary V)
- **predictions** are as good as guessing among PP equally likely options

Good news: there is typical strong correlation between PP and BLEU scores!

¹[Exercise 1. Find PP of 1-gram LM on the sequence T H T H T H T T H T T H for $p(\text{T})=0.3$, $p(\text{H})=0.7$ and $p(\text{H})=0.3$, $p(\text{T})=0.7$. Comment the results.]

Train-set vs. test-set perplexity

For an n -gram LM, the LP quantity can be computed as follows:

$$LP = -\frac{1}{|\mathbf{w}|} \sum_{t=1}^{|\mathbf{w}|} \log_2 p(w_t | h_t).$$

PP is a function of a LM and a text:

- the lower the PP the better the LM
- test-set PP evaluates LM generalization capability
- PP strongly penalizes zero probabilities
- train-set PP measures how good the LM explains training data

Note: train-set PP is strictly related to the train-set **likelihood**.

Estimating n -gram probabilities is not trivial due to:

- **model complexity**: e.g. 10,000 words correspond to 1 trillion 3-grams!
- **data sparseness**: e.g. most 3-grams are rare events even in huge corpora.

Relative frequency estimate: MLE of any discrete conditional distribution is:

$$f(w | x y) = \frac{c(w | x y)}{\sum_w c(w | x y)} = \frac{c(x y w)}{c(x y)}$$

where n -gram counts $c(\cdot)$ are taken over the **training corpus**.

Problem: relative frequencies in general overfit the training data

- if the test sample contains a "new" n -gram **PP** $\rightarrow +\infty$
- this is largely the most frequent case for $n \geq 3$

We need frequency smoothing!

Issue: $f(w | x y) > 0$ only if $c(x y w) > 0$

Idea: for each observed w **take off some fraction of probability** from $f(w | x y)$ and redistribute the total to all words never observed after $x y$.

- the **discounted frequency** $f^*(w | x y)$ satisfies:

$$0 \leq f^*(w | x y) \leq f(w | x y) \quad \forall x, y, w \in V$$

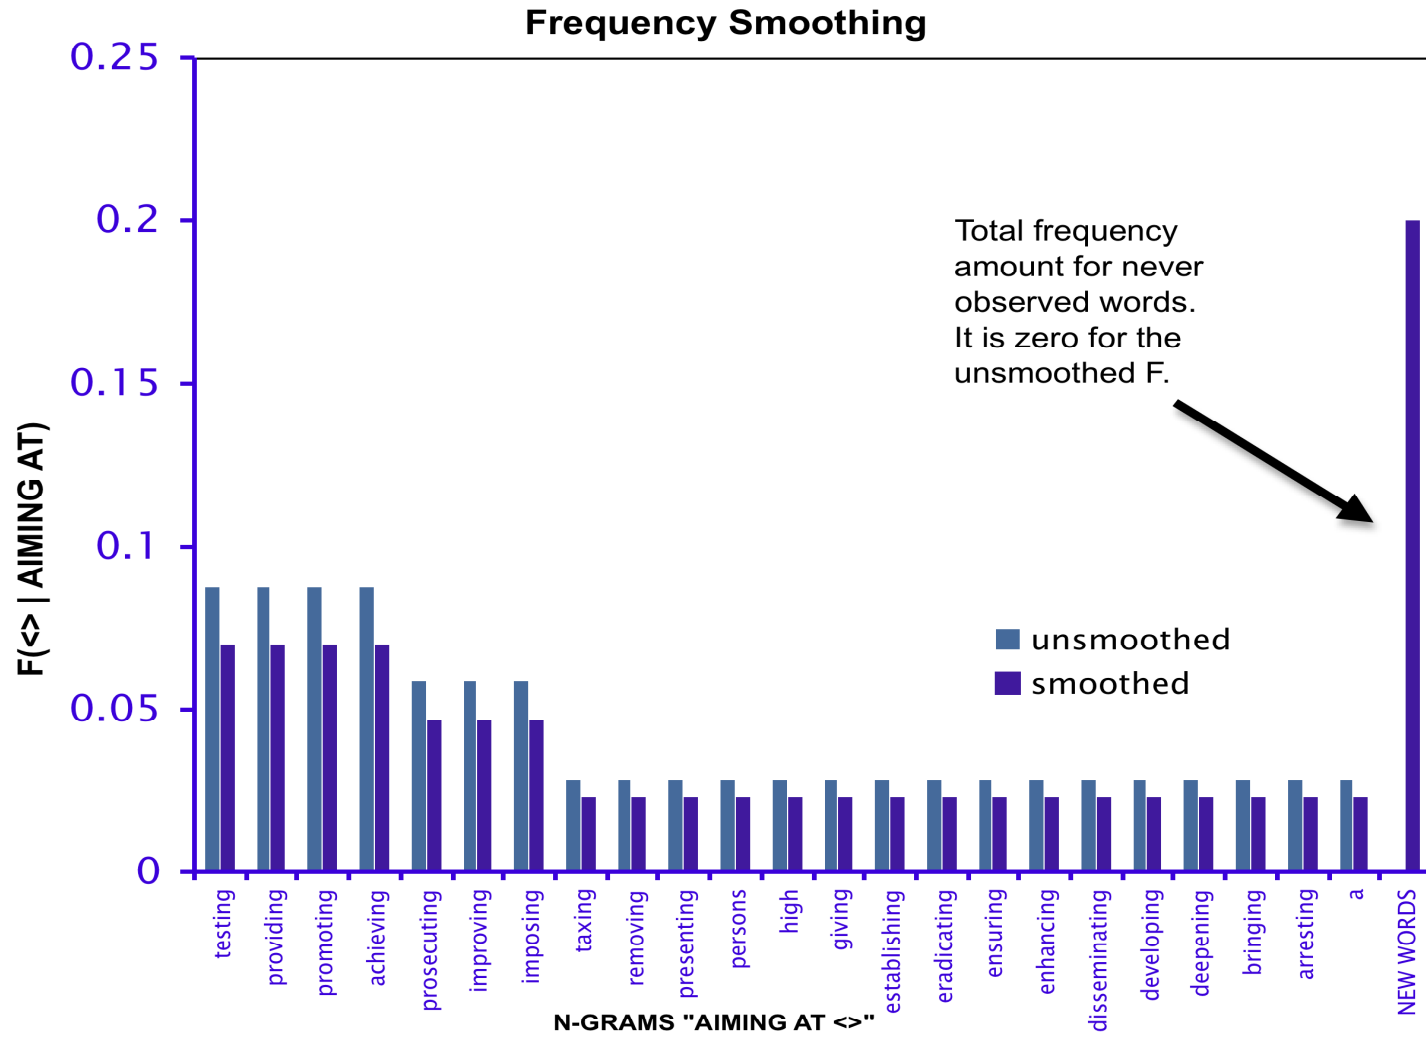
- the total discount is called **zero-frequency probability** $\lambda(x y)$:²

$$\lambda(x y) = 1.0 - \sum_{w \in V} f^*(w | x y)$$

How to redistribute the total discount?

²Notice: by convention $\lambda(x y) = 1$ if $f(w | x y) = 0$ for all w , i.e. $c(x y) = 0$.

Discounting Example



Insight: redistribute $\lambda(x \ y)$ according to the lower-order smoothed frequency.

Two major **hierarchical** schemes to compute the **smoothed frequency** $p(w \mid x \ y)$:

- **Back-off**, i.e. select the best available n -gram approximation:

$$p(w \mid x \ y) = \begin{cases} f^*(w \mid x \ y) & \text{if } f^*(w \mid x \ y) > 0 \\ \alpha_{xy} \times \lambda(x \ y)p(w \mid y) & \text{otherwise} \end{cases} \quad (5)$$

where α_{xy} is an appropriate normalization term.³

- **Interpolation**, i.e. sum up the two approximations:

$$p(w \mid x \ y) = f^*(w \mid x \ y) + \lambda(x \ y)p(w \mid y). \quad (6)$$

Smoothed frequencies are learned bottom-up, starting from 1-grams ...

³[Exercise 2. Find an expression for α_{xy} s.t. $\sum_w p(w \mid x \ y) = 1$.]

Unigram smoothing permits to treat **out-of-vocabulary** (OOV) words in the LM.

Assumptions:

- $|U|$: upper-bound estimate of the size of the **true vocabulary**
- $f^*(w) > 0$ on **observed vocabulary** V , e.g. $f^*(w) = c(w)/(N + |V|)$
- λ : total discount reserved to OOV words, e.g. $\lambda = N/(N + |V|)$

Then: 1-gram back-off/interpolation schemes collapse to:

$$p(w) = \begin{cases} f^*(w) & \text{if } w \in V \\ \lambda \times (|U| - |V|)^{-1} & \text{otherwise} \end{cases} \quad (7)$$

Notice: we introduce approximations when an OOV word o appears:

$$p(w \mid h_1 \ o \ h_2) = p(w \mid h_2) \quad \text{and} \quad p(o \mid h) = p(o)$$

Important: use a common value $|U|$ when comparing/combining different LMs!

Linear interpolation (LI) [Jelinek, 1990]

- **Insight:**
 - learn $\lambda(x y)$ from data with a mixture model
 - MLE on some held-out data (**EM algorithm**)
- **Solution:**

$$f^*(w | xy) = (1 - \lambda([x y]))f(w | xy) \quad \text{and} \quad 0 \leq \lambda([x y]) \leq 1$$

the notation $[x y]$ means that a map is applied to reduce the set of parameters, e.g., according to the frequency of the last word in the history:

$$[x y] = \begin{cases} 0 & \text{if } c(y) \leq k_1 \\ c(y) & \text{if } k_1 < c(y) \leq k_2 \\ y + k_2 & \text{if } k_2 < c(y) \end{cases}$$

- **Pros:** sound and robust
- **Cons:** over-smooths frequent n-grams

Witten-Bell estimate (WB) [Witten and Bell, 1991]

- **Insight:** count how often you would back-off after $x \ y$ in the training data
 - corpus: $x \ y \ u \ x \ x \ y \ t \ t \ x \ y \ u \ w \ x \ y \ w \ x \ y \ t \ u \ x \ y \ u \ x \ y \ t$
 - assume $\lambda(x \ y) \propto$ number of back-offs (i.e. 3)
 - hence $f^*(w \mid x \ y) \propto$ relative frequency (linear discounting)
- **Solution:**


$$\lambda(x \ y) = \frac{n(x \ y \ *)}{c(x \ y) + n(x \ y \ *)} \quad \text{and} \quad f^*(w \mid xy) = \frac{c(x \ y \ w)}{c(x \ y) + n(x \ y \ *)}$$

where $c(x \ y) = \sum_w c(x \ y \ w)$ and $n(x \ y \ *) = |\{w : c(x \ y \ w) > 0\}|$.⁴

- **Pros:** easy to compute, robust for small corpora, works with artificial data.
- **Cons:** underestimates probability of frequent n -grams

⁴[Exercise 3. Compute $f^*(u \mid x \ y)$ with WB on the above artificial text.]

- Interpolation and back-off with WB discounting
- Trigram LMs estimated on the English Europarl corpus
- Logprobs of 3-grams of type aiming at _ observed in training



- Peaks correspond to very probable 2-grams interpolated with f^* respectively: at that, at national, at European
- Back-off performs slightly better than interpolation but costs more

Absolute Discounting (AD) [Ney and Essen, 1991]

- **Insight:**
 - high counts are be more reliable than low counts
 - **subtract a small constant β** ($0 < \beta \leq 1$) from each count
 - estimate β by maximizing the leaving-one-out likelihood of the training data
- **Solution:** (notice: one distinct β for each n-gram order)

$$f^*(w | x y) = \max \left\{ \frac{c(xyw) - \beta}{c(xy)}, 0 \right\} \text{ which gives } \lambda(xy) = \beta \frac{\sum_{w:c(xyw)>1} 1}{c(xy)}$$

where $\beta \approx \frac{n_1}{n_1+2n_2} \leq 1$ and $n_r = |\{x y w : c(x y w) = r\}|$.⁵

- **Pros:** easy to compute, accurate estimate of frequent n -grams.
- **Cons:** problematic with small and artificial samples.

⁵[Exercise 4. Given the text in WB slide find the number of 3-grams, n_1 , n_2 , β , $f^*(w | x y)$ and $\lambda(x y)$]

Kneser-Ney method (KN) [Kneser and Ney, 1995]

- **Insight:** lower order counts are only used in case of back-off
 - estimate frequency of back-offs to $y \ w$ in the training data (cf. WB)
 - corpus: **x y w x t y w t x y w u y w t y w u x y w u u y w**
 - replace $c(x \ y)$ with $n(* \ y \ w) = \#$ of observed back-offs (=3)
- **Solution:** (for 3-gram normal counts)

$$f^*(w | y) = \max \left\{ \frac{n(* \ y \ w) - \beta}{n(* \ y \ *)}, 0 \right\} \text{ which gives } \lambda(y) = \beta \frac{\sum_{w:n(* \ y \ w) > 1} 1}{n(* \ y \ *)}$$

where $n(* \ y \ w) = |\{x : c(x \ y \ w) > 0\}|$ and $n(* \ y \ *) = |\{x \ w : c(x \ y \ w) > 0\}|$

- **Pros:** better back-off probabilities, can be applied to other smoothing methods
- **Cons:** LM cannot be used to compute lower order n -gram probs

Modified Kneser-Ney (MKN) [Chen and Goodman, 1999]

- **Insight:**
 - specific discounting coefficients for infrequent n -grams
 - introduce more parameters and estimate them with leaving-one-out
- **Solution:**


$$f^*(w | x y) = \frac{c(x y w) - \beta(c(x y w))}{c(x y)}$$

where $\beta(0) = 0$, $\beta(1) = D_1$, $\beta(2) = D_2$, $\beta(c) = D_{3+}$ if $c \geq 3$, coefficients are computed from n_r statistics, corrected counts used for lower order n -grams

- **Pros:** see previous + more fine grained smoothing
- **Cons:** see previous + more sensitiveness to noise


Important: LM interpolation with MKN is the **most popular smoothing method**. Under proper training conditions it gives the best PP and BLEU scores!

- Interpolation with WB and MKN discounting (Europarl corpus)
- The plot shows the logprob of observed 3-grams of type `aiming at _`

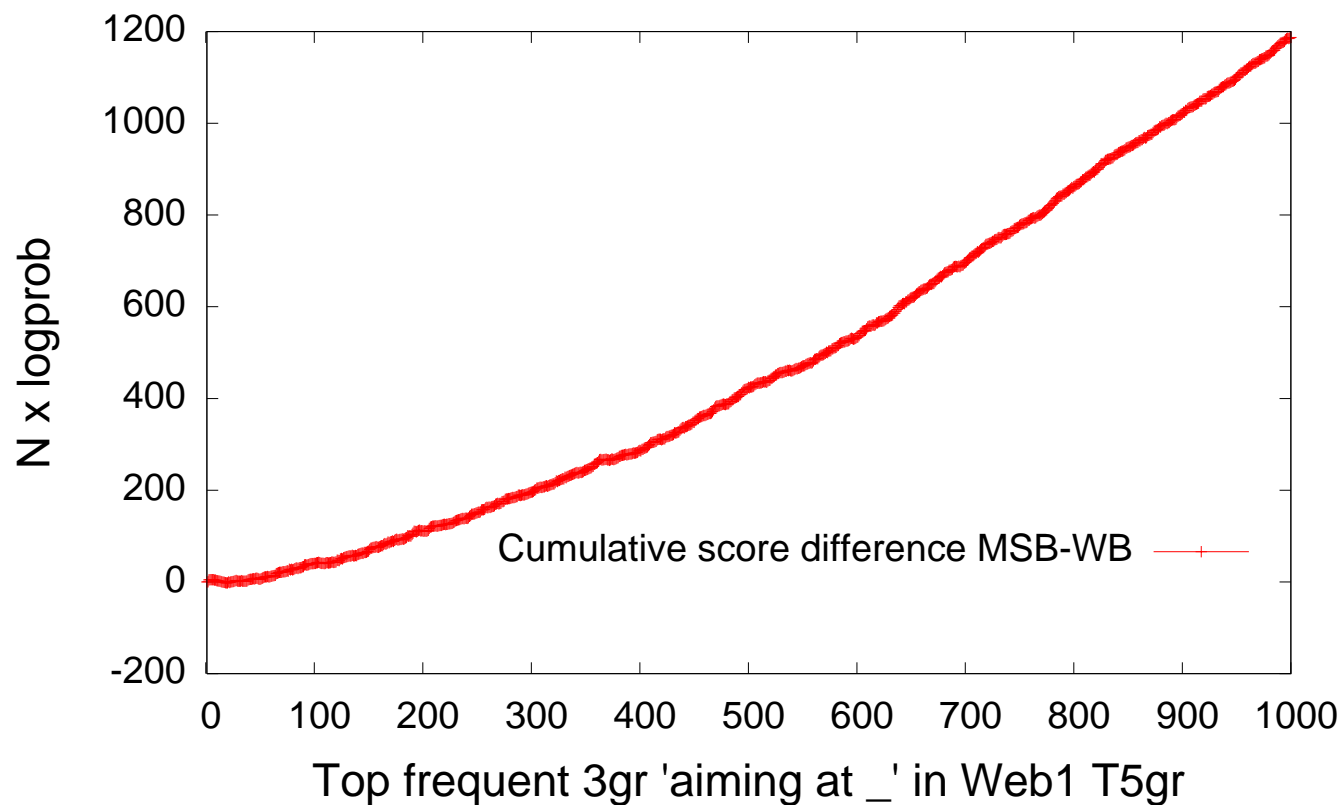


- Notice that for less frequent 3-grams WB assigns higher probability
- We have three very high peaks corresponding to large corrected counts:
 $n(*at\ that)=665$ $n(*\ at\ national)=598$ $n(*\ at\ European)=1118$
- Another interesting peak at rank #26: $n(*\ at\ very)=61$

- Train: interpolation with WB and MKN discounting (Europarl corpus)
- Test: 3-grams of type `aiming at _` (Google 1TWeb corpus)



- Train: interpolation with WB and MKN discounting (Europarl corpus)
- Test: 3-grams of type `aiming at _` (Google 1TWeb corpus)
- Plot: cumulative score differences between MKN and WB on top 1000 3-grams




- **LM Quantization** [Federico and Bertoldi, 2006]
 - **Idea**: one codebook for each n-gram/back-off level
 - **Pros**: improves storage efficiency
 - **Cons**: reduces discriminatory power
 - Experiments with 8bit quantization on ZH-EN NIST task showed:
 - * 2.7% BLEU drop with a 5-gram LM trained on 100M-words
 - * 1.6% BLEU drop with a 5-gram LM trained on 1.7G words.
- **Stupid back-off** [Brants et al., 2007]
 - no discounting, no corrected counts, **no back-off normalization**

$$p(w | x y) = \begin{cases} f(w | x y) & \text{if } f(w | x y) > 0 \\ k \cdot p(w | y) & \text{otherwise} \end{cases} \quad (8)$$

where $k = 0.4$ and $p(w) = c(w)/N$.

Is LM Smoothing Necessary?



From [Brants et al., 2007]. SB=stupid back-off, KN=modified Kneser-Ney

- **Conclusion:** proper smoothing useful up to 1 billion word training data!

- Use **less sparse representation of words** than surface form words
 - e.g. part-of-speech, semantic classes, lemmas, automatic clusters
- Higher chance to match longer n-grams in test sequences
 - allows to model longer dependencies, **to capture more syntax structure**
- For a text w we assume a corresponding class sequence g
 - ambiguous (e.g. POS) or deterministic (word classes)
- Factored LMs can be **integrated into log-linear models** with:
 - a **word-to-class factored model**: $\mathbf{f} \rightarrow \mathbf{e} \rightarrow \mathbf{g}$ with features:

$$h_1(\mathbf{f}, \mathbf{e}), h_2(\mathbf{f}, \mathbf{g}), h_3(\mathbf{f}), h_4(\mathbf{g})$$

- a **word-class joint model** $\mathbf{f} \rightarrow (\mathbf{e}, \mathbf{g})$ with features

$$h_1(\mathbf{f}, \mathbf{e}, \mathbf{g}), h_2(\mathbf{f}), h_3(\mathbf{g})$$

Features of single sequences are log-probs of standard n -gram LMs.

- The n -gram prob is modeled with log-linear model [Rosenfeld, 1996]:

$$p_{\lambda}(w | h) = \frac{\exp(\sum_{r=1}^m \lambda_r h_r(h, w))}{\sum_{w'} \exp(\sum_{r=1}^m \lambda_r h_r(h, w'))} = \frac{1}{Z(h)} \exp\left(\sum_{r=1}^m \lambda_r h_r(h, w)\right)$$

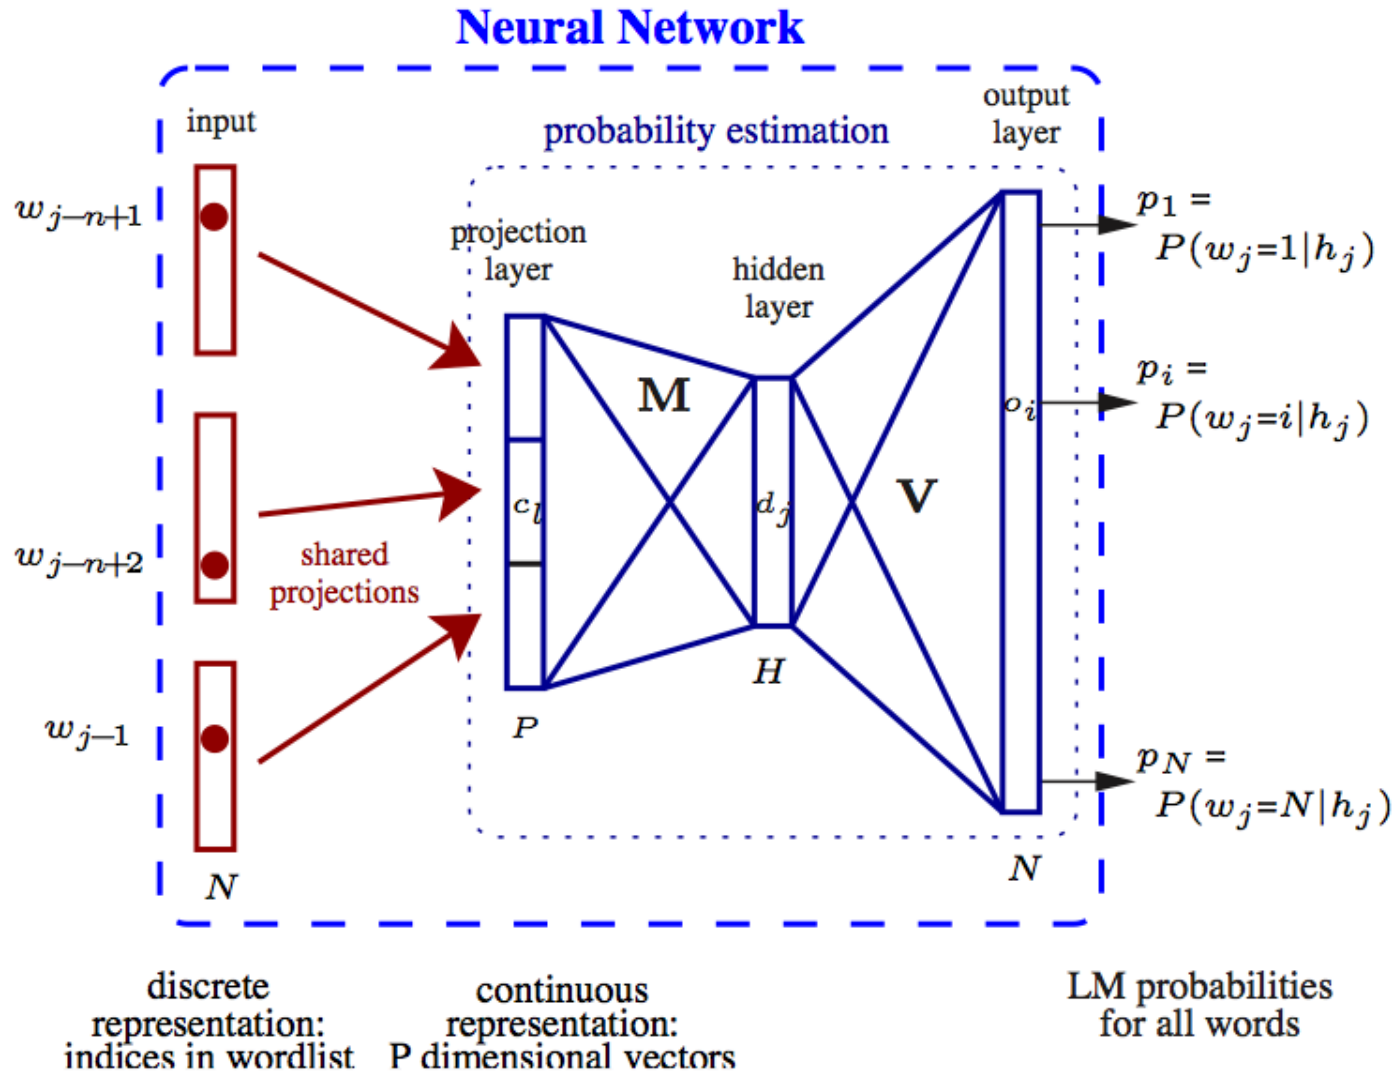
- $h_r(\cdot)$ are **feature functions** (arbitrary statistics), λ_r are **free parameters**
- **Features can model any dependency** between w and h .
- Given feature functions and training sample w , parameters can be estimated [Berger et al., 1996] by maximizing the **posterior log-likelihood**:

$$\hat{\lambda} = \arg \max_{\lambda \in \mathbf{R}^m} \sum_{t=1}^{|\mathbf{w}|} \log p_{\lambda}(w_t | h_t) + \log q(\lambda)$$

- where the second term is a **regularizing Gaussian prior**
- ME n-grams are rarely used: perform comparably but at higher computational costs, because of the partition function $Z(h)$.

- Most promising among recent development on n-gram LMs.
- **Idea:** Map single word into a $|V|$ -dimensional vector space
 - Represent n-gram LM as a **map between vector spaces**
- **Solution:** Learn map with neural network (NN) architecture
 - one hidden layer compress information (projection)
 - second hidden layer performs the n-gram prediction
 - other architectures are possible: e.g. recurrent NN
- **Implementations:**
 - Continuous Space Language Model [Schwenk et al., 2006]
 - Recurrent Neural Network Language Modeling Toolkit ⁶
- **Pros:**
 - Fast run-time, competitive when used jointly with standard model
- **Cons:**
 - Computational cost of training phase
 - Not easy to integrate into search algorithm (used in re-scoring)

⁶<http://rnnlm.sourceforge.net>



(From [Schwenk et al., 2006])

- Availability of large scale corpora has pushed research toward using huge LMs
- MT systems set for evaluations use LMs with over a billion of 5-grams
- Estimating accurate large scale LMs is still computationally costly
- Querying large LMs can be carried out rather efficiently (with adequate RAM)

Available LM toolkits

- SRILM: training and run-time, Moses support, open source (no commercial)
- IRSTLM: training and run-time, Moses support, open source
- KENLM: run-time, Moses support, open source

Interoperability

- The standard for n-gram LM representation is the so-called **ARPA file format**.

Represents both interpolated and back-off n-gram LMs

- format: $\log(\text{smoothed-prob}) :: \text{n-gram} :: \log(\text{back-off weight})$
- computation: look first for smoothed-prob, otherwise back-off

```
ngram 1= 86700
ngram 2= 1948935
ngram 3= 2070512
```

```
\1-grams:
```

```
-2.94351    world    -0.51431
-6.09691    friends  -0.15553
-2.88382    !        -2.38764
```

```
...
```

```
\2-grams:
```

```
-3.91009    world !    -0.3514
-3.91257    hello world -0.2412
-3.87582    hello friends -0.0312
```

```
...
```

```
\3-grams:
```

```
-0.00108    hello world !
-0.00027    hi hello !
```

```
...
```

```
\end\
```

Represents both interpolated and back-off n-gram LMs

- **format**: $\log(\text{smoothed-prob}) :: \text{n-gram} :: \log(\text{back-off weight})$
- **computation**: look first for smoothed-prob, otherwise back-off

```
ngram 1= 86700
ngram 2= 1948935
ngram 3= 2070512
```

Query: $\text{Pr}(! / \text{hello friends })?$

```
\1-grams:
-2.94351    world    -0.51431
-6.09691    friends  -0.15553
-2.88382    !        -2.38764
...

\2-grams:
-3.91009    world !    -0.3514
-3.91257    hello world -0.2412
-3.87582    hello friends -0.0312
...

\3-grams:
-0.00108    hello world !
-0.00027    hi hello !
...

\end\
```

1. look-up $\log\text{Pr}(\text{hello friends } !)$
failed! then back-off
2. look-up $\log\text{Bow}(\text{hello friends })$
res=-0.0312
3. look-up $\log\text{Pr}(\text{friends } !)$
failed! then back-off
4. look-up $\log\text{Bow}(\text{friends })$
res=res-0.15553
5. look-up $\log\text{Pr}(!)$
res=res-2.88382
6. prob= $\exp(\text{res})=0.04640$

References

- [Berger et al., 1996] Berger, A., Pietra, S. A. D., and Pietra, V. J. D. (1996). A maximum entropy approach to natural language processing. Computational Linguistics, 22(1):39–71.
- [Brants et al., 2007] Brants, T., Popat, A. C., Xu, P., Och, F. J., and Dean, J. (2007). Large language models in machine translation. In Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning, pages 858–867, Prague, Czech Republic.
- [Chen and Goodman, 1999] Chen, S. F. and Goodman, J. (1999). An empirical study of smoothing techniques for language modeling. Computer Speech and Language, 4(13):359–393.
- [Federico, 1999] Federico, M. (1999). Usability evaluation of a spoken data-entry interface. In Proceedings of the IEEE International Conference on Multimedia Computing and Systems, volume I, pages 726–731, Florence, Italy.
- [Federico and Bertoldi, 2001] Federico, M. and Bertoldi, N. (2001). Broadcast news LM adaptation using contemporary texts. In Proceedings of the 7th European Conference on Speech Communication and Technology, Aalborg, Denmark.
- [Federico and Bertoldi, 2006] Federico, M. and Bertoldi, N. (2006). How many bits are needed

to store probabilities for phrase-based translation? In Proceedings on the Workshop on Statistical Machine Translation, pages 94–101, New York City. Association for Computational Linguistics.

[Federico et al., 2008] Federico, M., Bertoldi, N., and Cettolo, M. (2008). Iirstlm: an open source toolkit for handling large scale language models. In Proceedings of Interspeech, Brisbane, Australia.

[Franz and McCarley, 2002] Franz, M. and McCarley, J. S. (2002). How many bits are needed to store term frequencies? In Proceedings of ACM SIGIR, pages 377–378, Tampere, Finland.

[Good, 1953] Good, I. J. (1953). The population frequencies of species and the estimation of population parameters. Biometrika, 40:237–264.

[Jelinek, 1990] Jelinek, F. (1990). Self-organized language modeling for speech recognition. In Weibel, A. and Lee, K., editors, Readings in Speech Recognition, pages 450–505. Morgan Kaufmann, Los Altos, CA.

[Katz, 1987] Katz, S. M. (1987). Estimation of probabilities from sparse data for the language model component of a speech recognizer. IEEE Trans. Acoust., Speech and Signal Proc., ASSP-35(3):400–401.

[Kneser and Ney, 1995] Kneser, R. and Ney, H. (1995). Improved backing-off for m-gram language modeling. In Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing, volume 1, pages 181–184, Detroit, MI.

- [Nadas, 1985] Nadas, A. (1985). On Turing's formula formula for word probabilities. IEEE Trans. Acoust., Speech and Signal Proc., ASSP-32:1414–1416.
- [Ney and Essen, 1991] Ney, H. and Essen, U. (1991). On smoothing techniques for bigram-based natural language modelling. In Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing, pages S12.11:825–828, Toronto, Canada.
- [Rosenfeld, 1996] Rosenfeld, R. (1996). A maximum entropy approach to adaptive statistical language modeling. Computer Speech and Language, 10:187–228.
- [Schwenk et al., 2006] Schwenk, H., Dechelotte, D., and Gauvain, J.-L. (2006). Continuous space language models for statistical machine translation. In Proceedings of the COLING/ACL 2006 Main Conference Poster Sessions, pages 723–730, Sydney, Australia. Association for Computational Linguistics.
- [Whittaker and Raj, 2001] Whittaker, E. and Raj, B. (2001). Quantization-based language model compression. In Proceedings of Eurospeech, pages 33–36, Aalborg, Denmark.
- [Witten and Bell, 1991] Witten, I. H. and Bell, T. C. (1991). The zero-frequency problem: Estimating the probabilities of novel events in adaptive text compression. IEEE Trans. Inform. Theory, IT-37(4):1085–1094.