Neural network language models

Steve Renals

Automatic Speech Recognition – ASR Lecture 15
7,10 March 2016

n-gram language modelling

- The problem: estimate the probability of a sequence of T words, $P(w_1, w_2, \ldots, w_T) = P(w_1^T)$
- Decompose as conditional probabilities, with n-gram approximation:

$$P(w_1^T) = \prod_{t=1}^{T} P(w_t \mid w_1^{t-1}) \approx P(w_t \mid w_{t-(n-1)}^{t-1})$$

- Many possible word sequences – consider vocab size $|V| = 100\,000$: 10^{20} possible 4-grams. Smooth n-gram estimates using smaller context size
- Problems with n-gram langfuage models
  1. Curse of dimensionality – model size increases exponentially with context size
  2. Probability estimation in a high-dimensional discrete space – not smooth, small changes in discrete context may result in large changes in probability estimate
  3. Do not take word similarity into account

Distributed representation for language modelling

- Each word is associated with a learned *distributed representation* (feature vector)
- Use a neural network to estimate the conditional probability of the next word given the the distributed representations of the context words
- Learn the distributed representations and the weights of the conditional probability estimate jointly by maximising the log likelihood of the training data
- Similar words (distributionally) will have similar feature vectors — small change in feature vector will result in small change in probability estimate (since the NN is a smooth function)

Neural Probabilistic Language Model

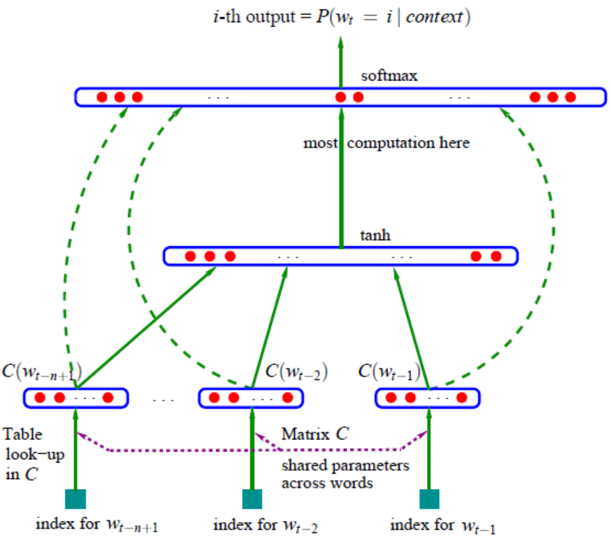

Bengio et al (2006)

Neural Probabilistic Language Model

- Train using stochastic gradient ascent to maximise log likelihood
- Number of free parameters (weights) scales
  - Linearly with vocabulary size
  - Linearly with context size
- Can be (linearly) interpolated with n-gram model
- Perplexity results on AP News (14M words training). $|V| = 18k$

model	**n**	**perplexity**
NPLM(100,60)	6	109
n-gram (KN)	3	127
n-gram (KN)	4	119
n-gram (KN)	5	117

Recurrent Neural Network (RNN) LM

- Rather than fixed input context, *recurrently connected* hidden units provide memory
- Model learns "how to remember" from the data
- Recurrent hidden layer allows clustering of variable length histories

RNN LM

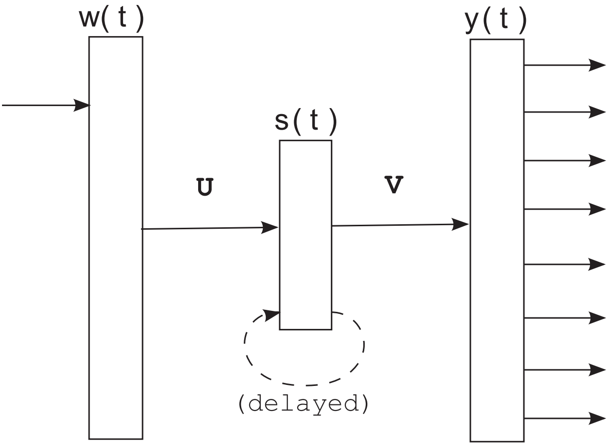

Mikolov (2011)

RNN training: back-propagation through time

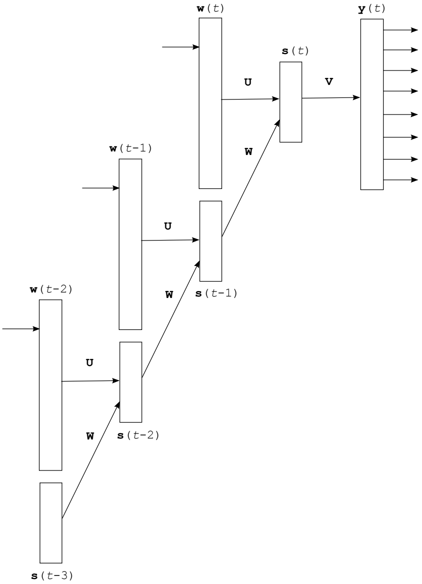

Reducing computation at the output layer

Majority of the weights (hence majority of the computation) is in the output layer – potentially V units wide, where V is vocabulary size

1. Model fewer words
   - **Shortlist**: use the NN to model only the most frequent words
2. Structure the output layer
   - **Factorization of the output layer**: first estimate the probability over word classes then over words within the selected class
   - **Hiearchical softmax**: structure the output layer as a binary tree
3. Efficiently estimate the normalised outputs
   - **Noise contrastive estimation**: train each output unit as an independent binary classifier

Shortlists

- Reduce computation by only including the s most frequent words at the output — the *shortlist* (S) (full vocabulary still used for context)
- Use an n-gram model to estimate probabilities of words not in the shortlist
- Neural network thus redistributes probability for the words in the shortlist

$$P_S(h_t) = \sum_{w \in S} P(w|h_t)$$

$$P(w_t|h_t) = \begin{cases} P_{NN}(w_t|h_t)P_S(h_t) & \text{if} w_t \in S \\ P_{KN}(w_t|h_t) & \text{else} \end{cases}$$

- In a $|V| = 50k$ task a 1024 word shortlist covers 89% of 4-grams, 4096 words covers 97%

NPLM — ASR results

Speech recognition results on Switchboard

7M / 12M / 27M words in domain data.

500M words background data (broadcast news)

Vocab size $|V| = 51k$, Shortlist size $|S| = 12k$

WER/%

in-domain words	7M	12M	27M
KN (in-domain)	25.3	23.0	20.0
NN (in-domain)	24.5	22.2	19.1
KN (+b/g)	24.1	22.3	19.3
NN (+b/g)	23.7	21.8	18.9

Factorised RNN LM

Mikolov 2011

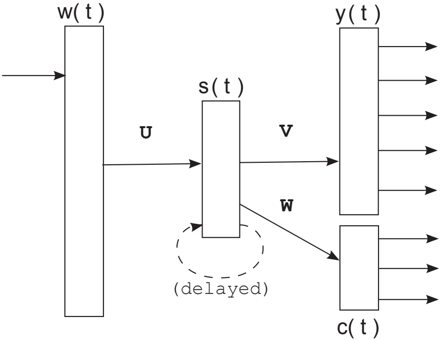

$$P(w_i|hist) = P(c_i|s(t))P(w_i|c_i, s(t))$$

1. Compute a probability distribution over C classes
2. Compute a probability distribution over $V' \leq V$ words in the class

Factorised RNN LM

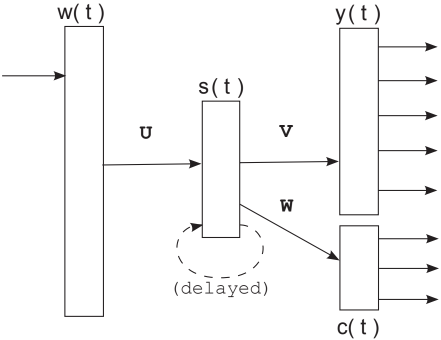

$$P(w_i|hist) = P(c_i|s(t))P(w_i|c_i, s(t))$$

- Instead of doing softmax over V elements, only $C + V'$ outputs have to be computed, and the softmax function is applied separately to the classes and the words in that class.
- C is constant; V can be variable.

Classes for a factorised RNN LM

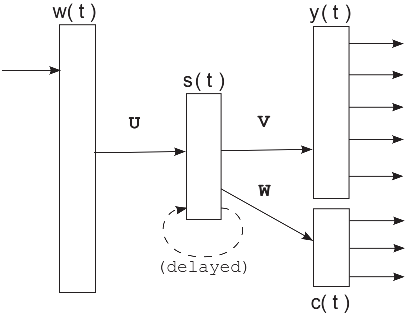

$$P(w_i|hist) = P(c_i|s(t))P(w_i|c_i, s(t))$$

- Each word is assigned to a single class based on unigram probabilities – “frequency binning”
- Most frequent words in class 1: V' is small for that class
- Rarest words in class C: V' is large for that class but the words are infrequent

Perplexity Results

Table 2. *Comparison of different neural network architectures on Penn Corpus (1M words) and Switchboard (4M words).*

	Penn Corpus		Switchboard	
Model	NN	NN+KN	NN	NN+KN
KN5 (baseline)	-	141	-	92.9
feedforward NN	141	118	85.1	77.5
RNN trained by BP	137	113	81.3	75.4
RNN trained by BPTT	123	106	77.5	72.5

Factorised output layer an speedup training by 25x for 100k vocabulary, with minimal effect on perplexity

Hierarchical Softmax

Bengio 2006

- Push the class-based factorization idea to the limit
- Class-based factorization is 1-level structuring
- "Classes of classes" – 2 level structuring
- Balanced binary tree – n level structuring ($n \sim \log_2 V$), each leaf is a word
- Each node of the tree is a 2-class classifier – make probabilistic binary decisions
- Need only consider the $\log_2 V$ nodes on the path from the root to the leaf for each word

$$P(w|hist) = \prod_{j=1}^{n} P(b_j(v)|b_1(v), \ldots, b_{j-1}(v), hist)$$

Noise contrastive estimation (NCE)

Chen et al (2015) (not the original source, but a clear application to ASR)

- Aim: avoid directly computing the normalisation term (denominator) in softmax (involves summing over V units)
- Method: treat each output unit separately, as sigmoid classifier between the observed data (for that word) and a "noise" distribution
- Assume that data for a history h generated by a mixture of an RNNLM distribution $P_{RNN}(\cdot|h)$ and a noise distribution $P_n(\cdot|h)$ – typically, P_n is a unigram
- Each node computes the posterior probability of whether a word sample w comes from the RNNLM or the noise:

$$P(C_w^{RNN} = 1|w, h) = \frac{P_{RNN}(w|h)}{P_{RNN}(w|h) + kP_n(w|h)}$$

$$P(C_w^n = 1|w, h) = 1 - P(C_w^{RNN} = 1|w, h)$$

Noise contrastive estimation (NCE)

- NCE training minimises this cost function:

$$E = -\frac{1}{N_w} \sum_{i=1}^{N_w} \left(\ln P(C_w^{RNN} = 1 | w, h) + \sum_{j=1}^{k} \ln P(C_w^n = 1 | w, h) \right)$$

- k samples drawn from the unigram noise distribution for the current word; typically $k \sim 10$
- The RNNLM distribution is given by

$$P_{RNN}(w_i | h) = \frac{\exp(\mathbf{v}\,\mathbf{s}(t))}{Z}$$

The normalisation term Z is learned; in practice it may be set to a constant for all contexts

State-of-the-art (2016)

Jozefowicz et al (2016), "Exploring the Limits of Language Modeling", http://arxiv.org/abs/1602.02410. (Google)

- Experiments on One Billion Word Benchmark data set, with 800k vocabulary
- Large-scale language modeling experiments, comparing
  - 5-gram model (interpolated Kneser-Ney smoothing)
  - RNN with sigmoid transfer functions
  - Various LSTM recurrent network models
    - "Size matters... The best models are the largest we were able to fit into a GPU memory."
    - Best performing model had two LSTM recurrent layers with 8192 and 1024 units ($\sim$ 1.8 billion parameters)
  - RNN models used a variant of NCE at the output layer
  - Also obtained more compact and slightly better performing models using convolutional layers over characters at input and output

Results: Single models

Model	Test Perplexity
Sigmoid-RNN-2048 (Ji et al., 2015a)	68.3
Interpolated KN 5-gram, 1.1B n-grams (Chelba et al., 2013)	67.6
Sparse Non-Negative Matrix LM (Shazeer et al., 2015)	52.9
RNN-1024 + MaxEnt 9-gram features (Chelba et al., 2013)	51.3
LSTM-512-512	54.1
LSTM-1024-512	48.2
LSTM-2048-512	43.7
LSTM-8192-2048 (No Dropout)	37.9
LSTM-8192-2048 (50% Dropout)	32.2
2-Layer LSTM-8192-1024 (BIG LSTM)	30.6
BIG LSTM+CNN Inputs	**30.0**
BIG LSTM+CNN Inputs + CNN Softmax	39.8
BIG LSTM+CNN Inputs + CNN Softmax + 128-dim correction	35.8
BIG LSTM+CNN Inputs + Char LSTM predictions	47.9

Results: Ensembles of models

Model	Test Perplexity
Large Ensemble (Chelba et al., 2013)	43.8
RNN+KN-5 (Williams et al., 2015)	42.4
RNN+KN-5 (Ji et al., 2015a)	42.0
RNN+SNM10-skip (Shazeer et al., 2015)	41.3
Large Ensemble (Shazeer et al., 2015)	41.0
Our 10 best LSTM models (equal weights)	26.3
Our 10 best LSTM models (optimal weights)	26.1
10 LSTMs + KN-5 (equal weights)	25.3
10 LSTMs + KN-5 (optimal weights)	25.1
10 LSTMs + SNM10-skip (Shazeer et al., 2015)	**23.7**

Reading

- Y Bengio et al (2006), "Neural probabilistic language models" (sections 6.1, 6.2, 6.3, 6.6, 6.7, 6.8), Studies in Fuzziness and Soft Computing Volume 194, Springer, chapter 6. http://link.springer.com/chapter/10.1007/3-540-33486-6_6
- T Mikolov et al (2011), "Extensions of recurrent neural network language model", ICASSP–2011. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5947611
- X Chen et al (2015), "Recurrent neural network language model training with noise contrastive estimation for speech recognition", ICASSP-2015. http://mi.eng.cam.ac.uk/~xc257/papers/ICASSP2015-rnnlm-nce.pdf
- R Jozefowicz et al (2016), "Exploring the Limits of Language Modeling", http://arxiv.org/abs/1602.02410.