

SCALABLE LANGUAGE MODEL ADAPTATION FOR SPOKEN DIALOGUE SYSTEMS

Ankur Gandhe, Ariya Rastrow, Bjorn Hoffmeister

Amazon

ABSTRACT

Language models (LM) for interactive speech recognition systems are trained on large amounts of data and the model parameters are optimized on past user data. New application intents and interaction types are released for these systems over time, imposing challenges to adapt the LMs since the existing training data is no longer sufficient to model the future user interactions. It is unclear how to adapt LMs to new application intents without degrading the performance on existing applications. In this paper, we propose a solution to (a) estimate n-gram counts directly from the hand-written grammar for training LMs and (b) use constrained optimization to optimize the system parameters for future use cases, while not degrading the performance on past usage. We evaluated our approach on new applications intents for a personal assistant system and find that the adaptation improves the word error rate by up to 15% on new applications even when there is no adaptation data available for an application.

Index Terms— language modeling, speech recognition, adaptation

1. INTRODUCTION

The quality of the data used for language modeling (LM) directly determines the accuracy of an automatic speech recognition (ASR) system. Most commercial ASR systems use n-gram based LMs trained on large amounts of text data. In spoken dialog systems like *Amazon Alexa* and *Google Home*, transcribed user utterances, collected from live traffic, form a large part of the training data. The transcribed data is also used for optimizing the LM parameters. Due to this, the data for training and tuning the LM parameters is always from the past user usage. When a new application intent is added to the system¹, the past user data is no longer sufficient to model the user interactions expected with the release of the application intent. This mismatch often leads to a reduced recognition quality on recently launched applications.

In the absence of text data for training n-gram models, a common approach is to use a probabilistic context-free grammar (pCFG), similar to one described in [1]. Figure 1 shows an example grammar for a new application intent “cooking

```
i_want_to = (("i" ("want" | "need" |  
                "would" | "like"))) "to");  
action = ("prepare" |  
         "cook" | "bake" );  
food_or_drink = ("food" | DISH_NAME);  
cook_dish = ( i_want_to action food_or_drink)
```

Fig. 1: Example grammar written for an application intent “cooking recipe” in thrax [2] format

recipes”. The grammar is written to encompass the expected user interaction for the new application. The non terminals (*DISH_NAME* in the example) are expanded by specifying a catalog of entities for each non-terminal. Further, weights can be specified for the phrases and entities in the grammar to assign a probability to each sentence generated by the pCFG. In this sense, the pCFG can be regarded as a LM which is specific to a particular application intent.

The LM of an ASR system can be adapted to the new application by combining the pCFG with an existing model. A common approach is to sample data from the grammar [3], [4], or transform out-of-domain sentences into in-domain sentences using a simulator [5] and combine this data with pre-existing training data. However, certain grammars, e.g. those with large non-terminal, can encode exponentially many different paths, and hence, a large amount of data must be generated from these grammars to capture all possible variations. Another approach is to represent the final LM as a union of the grammar and n-gram models, where each of them is represented as an FST [6], [7]. However, rapidly accelerating spoken dialogue systems introduce ten’s of new applications every month, and the final model can become quite large and lead to increased latency. In section 2, we introduce a novel method for extracting n-gram counts directly from the grammar, thereby eliminating the need for generating sampled data. These counts can be then used to train a n-gram based model or as additional data for training neural network based models.

Language model adaptation has also been studied in recent studies. [8] adapted a recurrent neural network LM to a single domain while [9], [10] built a multi-domain neural network model by using domain vectors to modify the model output. However, an important missing aspect of the

¹An example of a new application intent is the ability to ask for cooking recipes

forementioned approaches is the ways of optimizing the LM training for multiple domains while maintaining the performance on past usage. In section 3, we introduce a constrained optimization criteria that minimizes the loss on a new application intents while not degrading the performance on existing applications supported by the spoken dialogue system. Our work can be extended to adapting neural language models by applying the same constrained loss during network training.

2. ESTIMATING N-GRAM MODELS FROM PCFG

A language model estimates the probability of a given sequence of words \vec{w} . N-gram LMs make the assumption that the probability of word w_i depends only on previous $n - 1$ words, so that the probability can be written as:

$$P(\vec{w}) = \prod_{|\vec{w}|} p(w_i | w_{i-1}, w_{i-2}, \dots, w_{i-n+1}) \quad (1)$$

For maximum likelihood estimation, it can be shown that $p(w_i | h) = \frac{C(w_i, h)}{\sum_j C(w_j, h)}$ where $h = w_{i-1}, w_{i-2}, \dots, w_{i-n+1}$ and $C(w_i, h)$ is count of the n-gram w_i, h in the training data. Furthermore, for most commonly used smoothing techniques [11], n-gram counts are sufficient to estimate the n-gram LM. In this paper, we extract fractional n-gram counts directly from a pCFG using the algorithm described in 2.1, and use smoothing techniques that can work with expected counts to estimate the n-gram probabilities.

2.1. Expected n-gram counts

A pCFG can be represented as a weighted finite state transducer (FST) where each arc represents either a word or a non-terminal, and the weights on the arcs are the probabilities associated with them; each non-terminal is also represented as an FST. When a sentence is sampled from the FST, the non-terminal arcs are replaced by the non-terminal FST to generate the final sentence.

Each sequence \vec{w} generated from the FST has an associated probability $p(\vec{w})$. Then, the expected count of an n-gram w_1, w_2, \dots, w_n within this sequence is same as $p(\vec{w})$. The expected count of an n-gram for the entire FST can be calculated by generating all possible sequences and summing their expected counts:

$$\mathbb{E}(C(w_1, \dots, w_n)) = \sum_{\vec{w}_i \in FST} p(\vec{w}_i) (C_i(w_1, \dots, w_n) \in \vec{w}_i) \quad (2)$$

where $C_i(w_1, \dots, w_n)$ is the count of w_1, \dots, w_n in \vec{w}_i . Computing this sum is non-trivial, as it requires enumerating all the sequences \vec{w}_i in FST and calculating the normalized probability $p(\vec{w}_i)$. However, we can use dynamic programming to efficiently compute the n-gram counts as well as the normalized probability for each n-gram.

Each n-gram w_1, \dots, w_n is represented as a sequence of arcs arc_1, \dots, arc_n in an FST and the expected count of the n-gram can be computed by the sum of the probability of all paths that include the corresponding sequence of arcs. If s_1 is the initial state of arc_1 , s_n is the final state of arc_n , then the sum of probabilities for the sequence of arcs is:

$$p(arc_{1 \rightarrow n}) = \alpha_{s_1} * c_{arc_1} * c_{arc_2} \dots * c_{arc_n} * \beta_{s_n} / Z \quad (3)$$

where α_{s_1} and β_{s_n} is the unnormalized sum of the probabilities of the paths ending in state s_1 and starting in s_n respectively, c_{arc_i} is the unnormalized probability of arc arc_i , and Z is the normalization factor (sum of all paths in the FST). The values of α_s, β_s for every state in the FST and the normalization factor Z can be computed efficiently using the forward-backward algorithm [12]. Then, the expected count of the n-gram is same as $p(arc_{1 \rightarrow n})$.

To compute the counts of all n-grams in the FST, we traverse the FST in a forward topological order. For each state, we iterate over all the incoming n-grams w_1, w_2, \dots, w_{n-1} with their corresponding accumulated probability $p(arc_{1 \rightarrow n-1})$. For each out-going arc at the state, with word w_n , we calculate the expected count $\mathbb{E}(w_1, w_2, \dots, w_n)$ using the following iterative equation:

$$p(arc_{1 \rightarrow n}) = p(arc_{1 \rightarrow n-1}) * c_{arc_n} * \beta_{s_n} / \beta_{s_{n-1}} \quad (4)$$

and propagate the n-gram w_1, w_2, \dots, w_n to the next state of the arc² along with its accumulated probability $p(arc_{1 \rightarrow n})$.

2.2. Estimating Language Models

Given the fractional counts extracted from a pCFG, smoothing algorithms that use fractional n-gram counts [13], [14] to build the n-gram model can be used. Using count-of-count methods like Katz smoothing [15] and mod-KN smoothing [11] requires extracting count-of-counts from the pCFG, which can be done in a way similar to extracting fractional n-gram counts. Similarly, the n-gram counts can be used to train a feed-forward neural network or interpolate with an existing neural network model [16]. However, for simplicity, in this paper, we use Katz smoothing on scaled fractional counts to build the final n-gram model.

3. OPTIMIZING INTERPOLATION WEIGHTS

A pre-existing LM can be adapted to the new application intent by linear interpolation with the LM from section 2, where the probability of a word sequence \vec{w} is calculated by using different interpolation weights, λ_i 's, for each LM in the mixture:

$$p_{mix}(w_i | h_i) = \sum_k \lambda_k p_k(w_i | h_i) \quad (5)$$

²Iterating over all n-grams can be expensive for adjacent non-terminals or cyclic grammars; this can be mitigated by applying heuristics such as not propagating forward n-grams with counts lower than a given threshold

such that $\sum_k \lambda_i = 1$. In this approach, the interpolation weights are estimated by minimizing a loss, e.g. perplexity on representative data from the new application intent. When bootstrapping an existing LM with a grammar, however, just optimizing on the application’s data is not sufficient; we need to make sure that the final LM does not degrade on existing applications. Hence, we propose a constrained optimization problem:

$$\begin{aligned} & \underset{\vec{\lambda}}{\text{minimize}} && \text{Loss}(D_{app}|\vec{\lambda}) \\ & \text{subject to} && \text{Loss}(D_{past}) \leq C \end{aligned} \quad (6)$$

where $\vec{\lambda} = \lambda_{app}, 1 - \lambda_{app}$ are interpolation weights for pre-existing LM and application intent LM respectively, $\text{Loss}(D|\vec{\lambda})$ is the loss function being minimized, D_{app} is representative data for the new application, and D_{past} is development set based on past usage.

3.1. Loss functions

The actual loss function to be minimized (or constrained) depends on the adaptation data available for a new application intent. In this paper, we propose to use three loss functions:

3.1.1. Negative squared minimization

When no data is available for a application intent, the only way to optimize the LM for the new intent is to assign maximum possible interpolation weight to the feature LM. Hence, the loss can be expressed as:

$$\text{Loss} = -\lambda_{app}^2 \quad (7)$$

This will maximize the application intent’s interpolation weight until we violate the constraint on past data.

3.1.2. Perplexity minimization

When sample text data, from the new application intent, is available for LM adaptation, one can minimize the perplexity of the interpolated model on the data

$$\text{Loss} = \exp\left(-\frac{\sum_M \log_e p_{mix}(w_i)}{M}\right) \quad (8)$$

where $p_{mix}(w_i)$ is the probability assigned by the interpolated model to the word w_i and M is the total number of words in the data.

3.1.3. Expected WER minimization

When there is transcribed audio available for an application intent, we can directly minimize the expected recognition error rate of the ASR system as explained in [17]. [18] showed that perplexity is often not correlated with WER, and hence it might be better to directly minimize the WER when possible.

Expected WER loss is computed by summing over the recognition error $\mathcal{E}(\vec{w}|\vec{w}_{ref})$ for all hypotheses, \vec{w} s, weighted by their posterior probability $p(\vec{w}|O)$:

$$\text{Loss} = \sum_{\vec{w}} p(\vec{w}|O) \mathcal{E}(\vec{w}|\vec{w}_{ref}) \quad (9)$$

In order to make the sum over all possible hypothesis tractable, we approximate the sum by restricting \vec{w} s to an n-best list of ASR hypothesis.

3.2. Constraint optimization

The constraint on past usage is implemented as a penalty term [19] in the final optimization function:

$$\text{Loss}_{constraint} = \sigma * \max(0, \text{Loss}(D_{past}|\vec{\lambda}) - C)^2 \quad (10)$$

where $C = \text{Loss}(D_{past}|Baseline)$, is the loss of the pre-existing LM, and σ is the penalty coefficient to control the tolerance of the constraint. σ can be set either statically to a high value or changed dynamically with every iteration of the optimization as suggested in [20]. In this paper, we use a static penalty value of 1000. The loss function for the constraint can be either perplexity or expected WER. However, we expect that adding the constraint on WER will lead to better interpolation weights for the new application intent.

3.3. Scaling to multiple applications

The proposed method can easily scale to adding multiple application intent grammars at the same time to an existing n-gram LM. We can extend equation 6 with one loss per application:

$$\text{Loss} = \sum_f^N \text{Loss}_f(D_{app_a}|\vec{\lambda})$$

The choice of loss function and the constraint used is independent of each application intent which enables us to choose the loss function for each intent depending on the availability of data for that application.

4. EXPERIMENTS

The training data for the baseline language model is a combination of transcribed data from users of a personal assistant system as well as other in-house data sources³. We build one LM for each data source and interpolate them into a single LM. The interpolation weights are optimized on a held-out subset of transcribed data from users.

We evaluate our approach on grammars written for new application intents which have little or no coverage in the LM

³Data crawled from reddit.com, amazon forums, and amazon customer reviews



Fig. 2: Figure above shows the perplexity and word error rate results when different amount of data is sampled from the grammars for generating n-gram models. Sampling size=0 shows the baseline perplexity and n-gram count is the method proposed in the paper

Table 1: Table showing improvement in perplexity and word error rate when minimizing different loss functions and different constraints

Constraint	Loss function	Past Data		Stock Price		Flights Info		Recipes	
		PPL	rel. WER	PPL	WER	PPL	WER	PPL	WER
Baseline		32.4	-	41.6	6.80	85.6	6.45	90.9	7.68
PPL	min L2	32.6	0%	41.6	6.82	80.3	6.29	90.9	7.65
	PPL	32.4	0%	37.6	6.67	65.7	5.82	76.3	5.96
	exp-WER	32.7	0%	37.3	6.64	62.4	5.75	73.8	5.78
expWER	min L2	34.4	1%	35.2	6.54	65.3	5.74	83.2	6.49
	PPL	34.5	1%	35.1	6.52	64.3	5.7	81	6.26
	expWER	34.5	1%	35.1	6.54	61.9	5.63	79.2	6.14

training data. We tested with three application intents - *Getting stock prices*, *Getting Flight Information*, and *Asking for recipes*.

For running ASR experiments, we use an in-house ASR system⁴ based on [21, 22]. The acoustic model is a DNN trained on concatenated LFBE and i-vectors, and the baseline LM is a 4-gram models with a vocabulary of 200K trained with modified kneser-ney smoothing [11]. The LMs for the data extracted from application grammars were trained using Katz smoothing, and we limited the number of new vocabulary words (new relative to baseline) to 10K. Each intent was tested for perplexity (PPL) and word error rate (WER) on a held out test set targeted towards the intent containing about 500 to 700 utterances⁵. We also report the PPL and the relative WER degradation on past usage test set, dubbed as *Past Data*.

4.1. n-gram count extraction

In section 2.1, we proposed extracting n-gram counts directly from the pCFG written for a new application intent. In figure 2, we compare the proposed method with a model trained on data sampled from the pCFGs on three application intents. The application LM is interpolated with the baseline model based on weights estimated using constrained optimization with expected word error as the loss and the constraint. We find that the difference in performance of the two techniques varies for different applications - the perplexity difference is small for *Recipes* and *Stock Price* and significantly better for *Flights info*. Similar results are observed on word error rate, although *Stock Price* word error rate degrades with larger number of samples.

We looked at the total number of tri-grams (without replacing non-terminals) in the grammars for the applications intents as well as the total number of non-terminals in the three grammars. *Flight Info* and *Stock Price* application grammars have much larger number of tri-grams in the grammar. However, *Flight Info* has a large number of non-terminals pairs (non-terminals appearing next to each other), which leads to an exponentially large number of n-grams that can be sampled,

⁴Trained on only a subset of data used for our production system

⁵The test set was collected in-house by language experts and is not reflective of actual live usage of the application

	Stock Price	Recipes	Flight Info
# of tri-grams	152k	49k	306k
# of non-terminals	3	12	23
# of non-terminal pair	3	23	203

Table 2: Number of paths and non-terminals in grammar of different intents

and hence we observe a large drop in perplexity and word error rate when using the proposed method.

4.2. Comparing Loss function and constraint

Section 3 describes different methods for estimating interpolation weights, both for the loss to be minimized as well the constraint. Table 1 shows the comparison of different loss functions and constraining on perplexity of baseline LM on the past data. It also compares the same loss functions but with expected WER as the constraint in the optimization. While both constraints ensure that the WER on past data degrades only marginally, constraint of expected WER leads to larger improvements in WER for new features. The improvements vary across the different application intent, and is more than 10% relative for *Flights Info* and *Recipes* applications, which can be attributed to the initial model having a high perplexity on these applications.

We also find that L2 minimization works quite with expected WER constraint and the improvement in WER is comparable to other loss functions. This is quite useful as it shows that we can adapt to new application intents even when there is *no adaptation data* available. The expected WER constraint uses the acoustic scores of each word when calculating the posterior probability, thereby reducing the dependency on the LM scores. This can explain why expected WER is less-strict constraint than perplexity and assigns larger interpolation weights to new application intent LMs.

4.3. Scalability

In the previous experiments, each application intent was optimized separately on its own loss function. However, our proposed method allows us to add as many new application intents into the optimization as we want. In figure 3, we show the word error rate change for the three application intents as more and more application intents are added into the optimization. We find even though the word error rate increases with very new application added, the increase is not significant even with 12 new applications.

5. CONCLUSION

We propose a new method for adapting the language model of a spoken dialogue system, specifically designed to scale

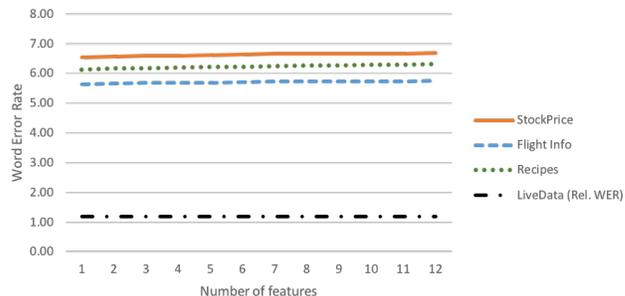


Fig. 3: Word error rate values of three different application intents and past data as more applications are added into the system

to support multiple application intents. The proposed constrained optimization function ensures that the accuracy of the system does not degrade as new applications are added to the system. We also show that using word error rate as a constraint leads to better adaptation for the new applications and removes the need for any adaptation data. In the future, we want to extend the same framework for adaptation of neural network based LMs.

6. REFERENCES

- [1] Daniel Jurafsky, Chuck Wooters, Jonathan Segal, Andreas Stolcke, Eric Fosler, Gary Tajchaman, and Nelson Morgan, “Using a stochastic context-free grammar as a language model for speech recognition,” in *Acoustics, Speech, and Signal Processing, 1995. ICASSP-95., 1995 International Conference on.* IEEE, 1995, vol. 1, pp. 189–192.
- [2] “<http://www.openfst.org/twiki/bin/view/grm/thrax>,” .
- [3] Lucian Galescu, Eric Ringger, and James Allen, “Rapid language model development for new task domains,” 1998.
- [4] Karl Weilhammer, Matthew N Stuttle, and Steve Young, “Bootstrapping language models for dialogue systems,” in *Ninth International Conference on Spoken Language Processing*, 2006.
- [5] Grace Chung, Stephanie Seneff, and Chao Wang, “Automatic induction of language model data for a spoken dialogue system,” in *6th SIGdial Workshop on Discourse and Dialogue*, 2005.
- [6] Ariya Rastrow, Bjorn Hoffmeister, Sri Venkata Surya Siva Rama Krishna, Rohit Krishna Prasad, et al., “Speech recognition with combined grammar and statistical language models,” Sept. 20 2016, US Patent 9,449,598.

- [7] Mehryar Mohri, Fernando Pereira, and Michael Riley, "Weighted finite-state transducers in speech recognition," *Computer Speech & Language*, vol. 16, no. 1, pp. 69–88, 2002.
- [8] Junho Park, Xunying Liu, Mark JF Gales, and Phil C Woodland, "Improved neural network based language modelling and adaptation," in *Eleventh Annual Conference of the International Speech Communication Association*, 2010.
- [9] Tanel Alumäe, "Multi-domain neural network language model.," 2013.
- [10] Ottokar Tilk and Tanel Alumäe, "Multi-domain recurrent neural network language model for medical speech recognition.," in *Baltic HLT*, 2014, pp. 149–152.
- [11] Stanley F Chen and Joshua Goodman, "An empirical study of smoothing techniques for language modeling," *Computer Speech & Language*, vol. 13, no. 4, pp. 359–394, 1999.
- [12] Lawrence R Rabiner, "A tutorial on hidden markov models and selected applications in speech recognition," *Proceedings of the IEEE*, vol. 77, no. 2, pp. 257–286, 1989.
- [13] Hui Zhang and David Chiang, "Kneser-ney smoothing on expected counts," in *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2014, vol. 1, pp. 765–774.
- [14] Vitaly Kuznetsov, Hank Liao, Mehryar Mohri, Michael Riley, and Brian Roark, "Learning n-gram language models from uncertain data.," 2016.
- [15] Slava Katz, "Estimation of probabilities from sparse data for the language model component of a speech recognizer," *IEEE transactions on acoustics, speech, and signal processing*, vol. 35, no. 3, pp. 400–401, 1987.
- [16] Graham Neubig and Chris Dyer, "Generalizing and hybridizing count-based and neural language models," *arXiv preprint arXiv:1606.00499*, 2016.
- [17] Xunying Liu, Mark JF Gales, and Philip C Woodland, "Use of contexts in language model interpolation and adaptation," in *Tenth Annual Conference of the International Speech Communication Association*, 2009.
- [18] Rukmini Iyer, Mari Ostendorf, and Marie Meteer, "Analyzing and predicting language model improvements," in *Automatic Speech Recognition and Understanding, 1997. Proceedings., 1997 IEEE Workshop on*. IEEE, 1997, pp. 254–261.
- [19] Alice E Smith, David W Coit, Thomas Baeck, David Fogel, and Zbigniew Michalewicz, "Penalty functions," .
- [20] Özgür Yeniay, "Penalty function methods for constrained optimization with genetic algorithms," *Mathematical and computational Applications*, vol. 10, no. 1, pp. 45–56, 2005.
- [21] Sree Hari Krishnan Parthasarathi, Bjorn Hoffmeister, Spyros Matsoukas, Arindam Mandal, Nikko Strom, and Sri Garimella, "fMLLR based feature-space speaker adaptation of DNN acoustic models," in *Proceedings Interspeech*, 2015.
- [22] Sri Garimella, Arindam Mandal, Nikko Strom, Bjorn Hoffmeister, Spyros Matsoukas, and Sree Hari Krishnan Parthasarathi, "Robust i-vector based adaptation of DNN acoustic model for speech recognition," in *Proceedings Interspeech*, 2015.