



Space-Efficient Representation of Entity-centric Query Language Models

Christophe Van Gysel, Mirko Hannemann, Ernest Pusateri, Youssef Oualil, Ilya Oparin

Apple

{cvangysel, mirko_hannemann, epusateri, youalil, ioparin}@apple.com

Abstract

Virtual assistants make use of automatic speech recognition (ASR) to help users answer entity-centric queries. However, spoken entity recognition is a difficult problem, due to the large number of frequently-changing named entities. In addition, resources available for recognition are constrained when ASR is performed on-device. In this work, we investigate the use of probabilistic grammars as language models within the finite-state transducer (FST) framework. We introduce a deterministic approximation to probabilistic grammars that avoids the explicit expansion of non-terminals at model creation time, integrates directly with the FST framework, and is complementary to n-gram models. We obtain a 10% relative word error rate improvement on long tail entity queries compared to when a similarly-sized n-gram model is used without our method.

Index Terms: ASR, language modeling, spoken entities, FSTs

1. Introduction

Virtual assistants (VAs) are becoming increasingly popular [1] to help users accomplish a variety of tasks [2]. VA queries can be informational [3], such as “*how old is Joe Biden*”, or navigational/transaccional (e.g., “*show me Dune*”), and are often centered around a named entity. Spoken entity recognition is a difficult problem [4, 5, 6, 7, 8], since the space of named entities is large, dynamic, and suffers from a long tail. In many use-cases, entity-oriented queries follow templates where the same template is applicable to entities of the same class [9, 10]. Consider the previous example query “*how old is Joe Biden*”, where segment “*how old is*” references a specific entity attribute and “*Joe Biden*” is an entity that belongs to the class of “famous people”. We can now generalize this example into a probabilistic grammar that consists of the sole template “*how old is X*” and X is a non-terminal that expands to a weighted list of entity names. An additional benefit of this approach is that entities can be extracted from external knowledge sources [4] that correlate with user queries and allow the VA to accurately recognize emerging entities that are not yet prevalent in training data.

In this paper we investigate the use of entity-centric probabilistic grammars as language models (LMs) within Automated Speech Recognition (ASR) to improve the recognition of entity-oriented queries that are common in VA usage. We are particularly interested in increasing the ASR system’s coverage of entities, while minimizing model size—i.e., we wish to reduce the word error-rate of all entity-oriented queries represented by the grammar. Hence, methods are trained using the full grammar and evaluated on their ability to cover that same grammar.

Using grammars for ASR is not a new idea. Stolcke et al. [11] presented an algorithm to estimate n-gram probabilities directly from rule-based grammars. In [12], the authors use grammars to encode semantic rules within limited-vocabulary

The template/entity lists used in this paper are available at github.com/apple/ml-interspeech2022-phi_rtn.

#	template	prob.	#	entity	prob.
(1)	play \$entity	0.4	(a)	hip hop rap	$2.7 \cdot 10^{-3}$
(2)	\$entity	0.2	(b)	Adele	$8.0 \cdot 10^{-5}$
(3)	hey VA \$entity	0.1	(c)	Drake	$7.9 \cdot 10^{-5}$
(4)	hey VA play \$entity	0.1	(d)	NBA YoungBoy	$7.4 \cdot 10^{-5}$
(5)	VA play \$entity	0.1	(e)	The Beatles	$6.3 \cdot 10^{-5}$
(6)	show me \$entity	0.1	(f)	play on Canada	$9.6 \cdot 10^{-9}$

(A) Template queries.

(B) Non-terminal \$entity.

Figure 1: Grammar of entity-centric media player queries.

ASR using a dynamic ad-hoc system. More recently, Gandhe et al. [13] estimate n-gram LMs directly from grammars to improve ASR for new application intents in VAs. Our goal is to integrate the probabilistic grammar within the finite-state transducer (FST) framework due to its ability to efficiently integrate acoustic and language knowledge sources [14], where LM and acoustic components are composed dynamically at runtime [15]. The dynamic interpolation of LMs is necessary when different LM components are updated at different cadences. Reducing the size of model updates, by only updating some components of an ASR system, is essential when ASR occurs on-device and bandwidth/disk space is costly. For VAs, the entity query distribution changes faster than the general voice query distribution, and hence, requires more frequent updates. In addition, in order to guarantee efficient decoding, the approaches we consider are limited to those that guarantee deterministic (i.e., unlike a dynamically-generated class-based LM).

Our research questions (RQs) are: (RQ1) Can we improve the long tail entity coverage of LMs while operating under model size constraints? (RQ2) Do size-constrained improvements in entity coverage translate into ASR quality improvements? We contribute: (1) analysis on prevalence of entity-oriented queries within VA usage, (2) a modeling approach— ϕ -RTN (“*phi RTN*”)—that provides a deterministic approximation to a prob. grammar compatible with the FST framework and complementary to n-gram models, and (3) their experimental evaluation.

2. Methodology

We build an entity-centric LM of spoken VA queries that is complementary to a general LM trained on transcribed VA queries. Our goal is to improve the speech recognition of tail entities, while taking into account resource constraints (i.e., model size). ASR LMs assign a non-zero probability $P(w_k | w_1, \dots, w_{k-1})$ to word w_k preceded by left-context w_1, \dots, w_{k-1} , with words w_i members of a fixed-size vocabulary V . During ASR decoding, the LM probability is combined with acoustic information to differentiate between competing hypotheses. We will now proceed as follows. First, we describe how entity-oriented queries can be described as probabilistic grammars (§2.1). Next, we provide background on non-deterministic graph grammar representations and their shortcomings (§2.2). Finally, we describe our approach, ϕ -RTN, that provides an approximation to FST-compatible graph-based grammars (§2.3).

2.1. Entity-centric Query Grammars

An analysis of a month of a random sample anonymized U.S. English query logs from a popular VA shows that over **15%** of media player queries that instructed the VA to play a song, album or artist follow the “*play \$entity*” template. Hence, a significant portion of media player queries can be represented using a probabilistic grammar. Estimating a grammar from entity-centric queries extracted from usage logs—that will subsequently be used for LM estimation—has the advantage that it allows us to keep the query templates static, while updating the weighted entity list using external knowledge sources that correlate with usage. In turn, this allows the VA to recognize emerging entities that are not yet included in transcribed training data. Denote T and E as the sets of query templates and entities, resp., an example grammar G —consisting of templates T with slots for entities E —that we wish to support is depicted in Fig. 1. The templates in Fig. 1A were extracted from usage logs and subsequently inspected by domain experts, whereas the entity feed in Fig. 1B was extracted from an external knowledge source. We are particularly interested in improving LM coverage for tail entity queries. We define queries whose joint probability falls below the median as tail queries.

2.2. Recursive Transitive Networks

Encoding the grammar depicted in Fig. 1 as a static FST by expanding every occurrence of $\$entity$ in the templates (Fig. 1A) leads to excessively large models—proportional to the cross-product, $|T| \cdot |E|$. Recursive Transitive Networks (RTNs) [16] offer an alternative representation to encode grammars—and more generally, class-based LMs [9]—that do not store the cross-product. Instead, an RTN consists of a family of FSTs—each associated with a non-terminal—where non-terminal labels on arcs within the constituent FSTs are recursively replaced by the FST they reference. Non-terminal symbol S indicates the root non-terminal that belongs to the constituent FST that is explored first. Grammar G (Fig. 1) can be represented using a RTN with a single level of recursion. Unfortunately, RTNs, when used as decoding graphs in a FST decoder, are generally non-deterministic [17, §4.5]. An FST is non-deterministic if at a given state there are multiple outgoing arcs that share the same input symbol. Non-determinism occurs within RTNs since there exist multiple paths through the graph that result in the same token sequence. Within the grammar depicted in Fig. 1, this would be the case within the state that corresponds to context “*hey VA*”—i.e., the wake word contributed by templates (3) and (4)—since the next token “*play*” can either match template (4) or entity (f). Non-determinism is detrimental to efficient ASR decoding as multiple paths can lead to the same hypothesis, and therefore should be avoided.

2.3. Deterministic Approximate RTNs

A key characteristic of entity-oriented VA queries is that carrier phrases (i.e., the templates excluding the non-terminal) are often short and use a limited vocabulary of frequent tokens (< 100 unique; see §3.1), whereas the entity names typically consist of ample infrequent tokens ($> 10^4$ unique; see §3.1). We can use this observation to relax the harsh conditions necessary for determinism when using RTNs for modeling entity-oriented queries by imposing a precedence of regular symbols over non-terminal symbols. More specifically, when we are at state s in our RTN after observing i tokens w_1, \dots, w_i and need to transition to the next state s' after observing the next token

w_{i+1} , we will first attempt to match a regular symbol on any arc leaving state s , and only explore entering a non-terminal FST, or exiting the current non-terminal FST, if a match cannot be found. Within the FST framework, this behavior can be implemented using ϕ -transitions [18], i.e., transitions followed when the requested symbol cannot be matched at the current state, that are typically used to implement back-off n-gram models [19] using FSTs.

We will now proceed by describing the topology of our model, named ϕ -RTN, as a two-level RTN, and explain how our model is made deterministic—approximating the actual RTN—by following precedence rules using ϕ -transitions. In addition, we discuss how we ensure that our model is correctly normalized. Later, at the end of §4, we discuss limitations/drawbacks of our approach—and means to alleviate them. **Topology.** We employ a two-level RTN, where root non-terminal S is associated with a rigid grammar FST F_T estimated on templates (e.g., Fig. 1A). The transition probabilities at each state are multiplied with discounting factor $\bar{\alpha} = 1 - \alpha$ ($0 < \alpha < 1$) to allow for out-of-domain utterances. The left-over mass, α , is distributed over all unseen tokens at the state through the use of a back-off arc that leads to a self-loop unigram state $s_{P(w)}$ that encodes a unigram probability distribution $P(w), \forall w \in V$. Template FST F_T references non-terminal $\$entity$, associated with entity FST F_E . FST F_E is modeled as a regular, non-backoff n -gram model over entity names (e.g., Fig. 1B), with probabilities scaled by $\bar{\alpha}$ to allow for out-of-domain entities. The left-over mass at each state in F_E will be used for transitioning out of F_E , back to F_T , and possibly $s_{P(w)}$ in the case of out-of-domain entities (see below).

Determinism. We address the RTN determinization issue discussed in §2.2 by prioritizing regular symbols over non-terminal entry/exit. More specifically, within F_T , when observing non-terminal $\$entity$ after context w_1, \dots, w_i , with corresponding state s , the corresponding sub-FST F_E can be reached by following the ϕ -transition from state s . Note that in the absence of $\$entity$, the ϕ -transition leads to the unigram state $s_{P(w)}$. When entering F_E , we remember at which state we need to return to in F_T . Within F_E , ϕ -transitions are used to exit sub-FST F_E and return to template FST F_T . Final states in F_E are removed, and instead, at each state s in F_E , the FST state’s final probability mass is combined with the leftover mass α and distributed over regular symbols not observed at s .

Normalization. To ensure that the resulting model is properly normalized, i.e. $\sum_{w \in V} P(w|s) = 1$ for all states s , we follow a strategy similar to back-off n-gram models [19]. More precisely, we assign a weight to the ϕ -transition leaving state s such that the probability mass assigned to unobserved events at s , when following ϕ -transitions recursively, are scaled to fit into the leftover mass at state s . For ϕ -transitions leading from F_T to F_E , the weights can be computed statically when the model is constructed—since we know the start state in F_E the ϕ -transition leads to, and the state from F_E that leads back to F_T after following ϕ once more. However, computation of the weights for ϕ -transitions that exit from F_E to F_T (excl. start state in F_E) is more difficult—due to the dependency between the state in F_E from which we are exiting and the state in F_T we are returning to. If we were to compute these ϕ -transition weights statically, we would need to store $\mathcal{O}(\text{NumStates}(F_T) \cdot \text{NumStates}(F_E))$ weights. To avoid the prohibitively expensive storage costs, we compute the weights of ϕ -transitions leading from F_E back to F_T partially at runtime. Partial computation of the weights is achieved by pre-computing the marginal unigram probability,

$\sum_{w \in \text{OutgoingSymbols}(s)} P(w)$, for all explicit symbols defined at each state s in F_E when the model is constructed. At runtime, since we know the leftover mass at state s in F_E equals $\alpha + P(\text{final} | s)$, we thus only need to consider explicit events defined at state s' in F_T that the ϕ -transition leaving F_E leads to, in order to update the partial weight. Since in our setting, F_T is sparse, the computation involved in this operation is negligible. Note that ϕ -RTN does not support templates with two or more consecutive non-terminals since that would lead to significantly more computation.

3. Experimental setup

3.1. Entity-centric Query Grammar

We focus on media player-type queries, where users instruct a VA to interact with audio content (e.g., songs, playlists). Our media player query grammar consists of a weighted list of 293 templates (77 unique tokens) that reference an entity slot and a weighted list of 2 608 460 entities (230 321 unique tokens) that can fill the slot. The templates were derived from high-frequency use-cases in a representative sample of query logs by domain experts with prior probabilities proportional to their presence in user requests. As shown in §2.1, a small number of templates can model a significant proportion of user queries. The list of media entities was extracted from the catalog of a popular media service, with probabilities based on interactions.

3.2. Evaluation Sets

For evaluation, we partition the set of expanded synthetic queries according to the rank percentiles of their joint probabilities (i.e., $P(\text{template}) \cdot P(\text{entity})$): head (top-10%), torso (between top-10% and the median), and tail (bottom-50%). Subsequently we sample 10k queries from each stratum for (1) evaluating the trade-off between perplexity and query LM size across the compared approaches (see below), and (2) word error rate (WER) evaluation where we linearly interpolate the query LMs with our main ASR LM, with weight coefficients 0.05 and 0.95, resp., on audio generated from the randomly sampled queries with a Neural Text-To-Speech system described in [20] to measure the effectiveness of our approach. The 95/5% weight distribution was chosen to allow the query LM to influence recognition while not dominating it. Within experiments not included in this paper, we verified that a 5% weight is sufficient for any model to maximize its impact without negatively impacting out-of-domain utterances. In addition, we also evaluate WER using the same setup on a uniform random sample of VA queries to ensure there is no regression. For (hyper-)parameter optimization, we build a development set by first sampling 10k queries from each stratum—that do not overlap with the test set—and taking the union. When interpolating multiple query LMs with the main ASR LM, the 0.05 weight is divided across the query LMs by maximizing likelihood on the dev. set using L-BFGS-B [21].

3.3. Approaches under comparison

We compare the following approaches that provide deterministic decoding graphs, where we sweep over hyperparameters that influence model size, for the task of recognition of spoken entity-centric VA queries: (1) We construct ϕ -RTN models directly from the grammar of templates and entities ($n = 2, 3, 4$). (2) We train back-off n-gram LMs ($n = 2, 3, 4$) using SRILM with Witten-Bell smoothing, and subsequently perform entropy

pruning [22] (with threshold $\theta \in \{4^{-i} \mid 4 \leq i < 20\} \cup \{0\}$) to reduce the model size, on the full set of expanded synthetic queries (§1, §2.1). Witten-Bell smoothing was chosen since it puts less reliance on the absolute counts than other smoothing methods (e.g., Good-Turing) and is therefore more suitable for synthetic data. The grammar-generated phrases are weighted according to pseudo-counts that were obtained by rescaling their probabilities such that the lowest pseudo-count equals 1. To avoid the filtering of n-grams with low pseudo-counts, we disabled the n-gram discounting cutoffs (i.e., $\text{gt}n\text{min} = 1 \forall n$). After n-gram model generation, we convert the obtained ARPA model to OpenFST [23] models using Kaldi’s `arpa2fst` tool [24]. To evaluate the space requirement of the model, we compare multiple FST formats: (a) free-form formats, such as the mutable `vector` and the immutable space-optimized `compact` acceptor formats, (b) the specialized `ngram` format [25] that uses the LOUDS encoding to efficiently represent n-gram back-off models. For the free-form formats (`vector`, `compact`), we apply FST minimization after `arpa2fst` in order to further reduce FST size. For generation of `ngram` FSTs, we skip the removal of redundant states during `arpa2fst`, such that the intermediate output FST retains the necessary information about back-off, and finally use the `fstconvert` utility to obtain a `ngram` FST. Note that all FSTs generated from the same n-gram model will be qualitatively identical and only differ in disk space requirements. **ASR system.** We used a CNN acoustic model [26], a 4-gram LM with Good-Turing smoothing in the 1st pass, and the same LM—in addition to the query LMs—interpolated with a Neural LM in the 2nd pass using a dynamic weighting scheme. To build the 4-gram LM, component 4-gram models built from data sources (> 10B manually/auto. transcribed anonymized VA requests) were combined [27] on held out manually transcribed VA data [28].

4. Results & Discussion

We answer the research questions asked in §1 as follows.

RQ1: Fig. 2 shows the trade-off between perplexity and model size for the various approaches (§3.3) used to encode back-off LMs (`vector`, `compact`, `ngram`), and our method, ϕ -RTN (`phi-rtn`) on head, torso and tail dev. sets (§3.2). When considering the back-off LMs for each dev. set, regardless of FST format, the trade-off curves all follow similar shapes and only differ by a near-constant offset on the x-axis. For the back-off LMs, the perplexity/size curve oscillates since we sweep over both n-gram order and pruning thresholds—combined within a single curve per approach. On the head of the distribution (Fig. 2a), back-off LMs quickly reach low perplexity at small model sizes ($\sim 2^7$ MB). However, when we consider the torso/tail (Fig. 2b, 2c), the curve becomes less sharp and only converges to the minimum when using large models ($\sim 2^{10}$ MB) with nearly no pruning. Across all back-off LM FST formats, the `ngram` format is most economical, followed by `compact` and then `vector`. Our method, `phi-rtn`, provides approximately the same coverage as `ngram` on the head set (Fig. 2a) at $\sim 2^6$ MB. However, on the tail set (Fig. 2c), `phi-rtn` obtains about an order of magnitude lower perplexity, and consequently, better coverage, at a model size of $\sim 2^6$ MB compared to the back-off LMs. The lower perplexity of back-off LMs on the head set, and their higher perplexity on the tail, is due to the following dual-effect: (a) back-off smoothing methods rely on absolute counts to evaluate the reliability of empirical prob. estimates, and hence, are generally biased towards high-prob. events, and (b) entropy pruning favors the removal of

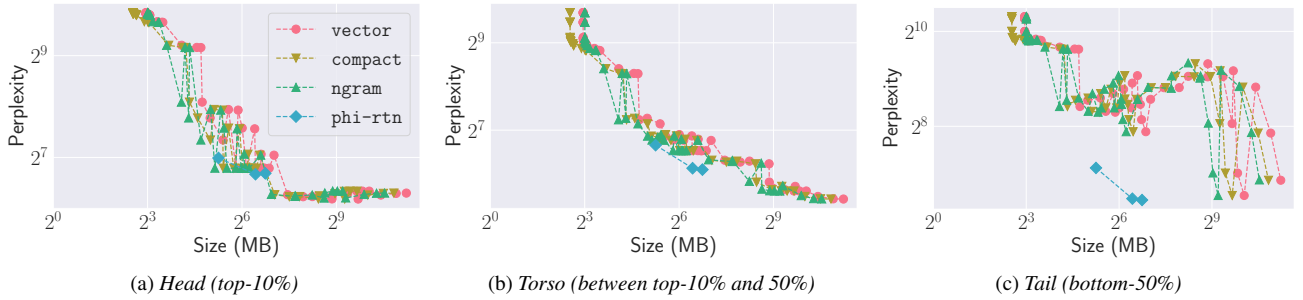


Figure 2: Model size vs. perplexity trade-off on dev. sets (§3.2) (closer to bottom-left is better) as we sweep over hyper-params (§3.3).

low-prob. events. Even in the case that we were to modify the smoothing/pruning methods to be less favorable w.r.t. the head of the distribution, the size of back-off LMs is proportional to the number of explicitly-defined probs., and hence, the curves in Fig. 2c would not change significantly. We answer our first RQ as follows: while back-off LMs are unsuitable to improve long tail entity coverage when operating under model size constraints, ϕ -RTN provides significantly better coverage. Since the worst-case model size of ϕ -RTN is asymptotically bounded by the number of unique entities, we can choose the entity set to be included according to the target model size.

Table 1: Overview of methods under comparison for WER evaluation, how they were selected, and their hyper-parameters.

Name	Selection criterion	Hyper-params
phi-rtm	Based on Fig. 2	$n = 3$
ngram ($<$)	$\sim \frac{1}{2} \times$ size phi-rtm	$n = 3, \theta = 6e-8$
ngram (\sim)	\sim size phi-rtm	$n = 4, \theta = 1e-8$
ngram ($>$)	$\sim 2 \times$ size phi-rtm	$n = 4, \theta = 4e-9$

Table 2: Error-rates using different models (Table 1) of media player queries, interpolated with the main LM at runtime. Letters between brackets denote sub-sections.

Method	Size (MB)	Head		Torso		Tail	
		WER	SER	WER	SER	WER	SER
Main LM only	-	10.7	38.2	12.1	47.2	13.2	51.6
+ ϕ -RTN only [2a]							
phi-rtm	86.1	5.3	18.2	5.7	20.8	6.3	23.2
+ small N-Gram/ ϕ -RTN combinations [2b]							
ngram ($<$)	42.9	5.2	18.8	6.0	22.6	7.0	27.2
phi-rtm + ngram ($<$)	129.1	5.0	17.6	5.6	20.7	6.1	23.0
+ medium N-Gram/ ϕ -RTN combinations [2c]							
ngram (\sim)	96.6	5.0	17.7	5.8	22.1	6.8	26.0
phi-rtm + ngram (\sim)	182.8	5.0	17.2	5.5	20.5	6.0	22.7
+ large N-Gram/ ϕ -RTN combinations [2d]							
ngram ($>$)	205.7	4.9	17.0	5.7	21.4	6.7	26.1
phi-rtm + ngram ($>$)	291.8	4.9	16.9	5.5	20.3	6.0	22.6

RQ2: We now investigate whether the increase in tail coverage provided by ϕ -RTN translates into improved ASR performance. Following the observations made from Fig. 2 on the dev. sets, we choose the **phi-rtm** model with $n = 3$ as it provides a good coverage/size trade-off. We choose **ngram** models relative to the size of the **phi-rtm** model by taking the **ngram** model closest to the target size (e.g., half size of **phi-rtm**) and pick the **ngram** within a $\epsilon = 5$ MB window with lowest perplexity on the dev. set (§3.2). Table 1 shows the selected models and selection criteria/hyperparameters. In Table 2, we evaluate the ϕ -RTN and back-off LMs interpolated individually with the main ASR LM, and combinations thereof (§3.2). Including **phi-rtm** results in a significant WER improvement across all test sets,

when compared to the main LM alone. This is unsurprising, since no media player-specific data was added to the main LM—except the queries occurring in transcribed data (end of §3.3)—as the models are complementary (§1). When comparing **phi-rtm** and same-sized **ngram**(\sim) (Table 2c), we note that while **phi-rtm** performs significantly better on the tail, **ngram**(\sim) wins on the head. Can a combination of both models result in overall better performance on all test sets? Comparing **ngram**($>$) (Table 2d) with **phi-rtm + ngram**(\sim) (Table 2c), the 11% smaller combination performs similarly on the head test set, but achieves a 10% relative WER reduction on the tail set. Finally, we did not observe a regression on the uniform sample of VA queries (§3.1) in terms of WER and end-to-end latency. We answer our 2nd RQ: improvements in model coverage translate to lower WER. In addition, the sum is greater than the parts: combining both methods yields an ASR quality increase at less space than would be obtained when using larger back-off LMs. **Discussion.** Since ϕ -RTN stores templates and entities as two separate sub-graphs, and $|E| \gg |T|$, the worst-case model size of ϕ -RTN is asymptotically bounded by the number of unique entities. The space-saving of ϕ -RTN over back-off LMs (§3.3) occurs since back-off LMs need to explicitly represent the relation between template contexts and entities (i.e., explicit probabilities). ϕ -RTN is suitable when the language used in the templates is sufficiently limited. Within our experiments, the ϕ -RTN model provides coverage for 99% of queries represented by the grammar (§3.1). However, in other applications where template language may be more complex, ϕ -RTN may not be the best fit. Even though ϕ -RTN has limitations, its inclusion in VA ASR systems allows the recognition of long tail entities that would otherwise not be recognized at all. Finally, collisions between the template/entity vocabulary that would lead to reduced coverage can be detected at training time and addressed by, e.g., rewriting the template/entity lists or including the affected query in the training data of a complimentary n-gram LM.

5. Conclusions

We introduced ϕ -RTN—a grammar-backed LM with low disk space cost that integrates with FST decoders and improves coverage for long tail entities in VA ASR scenarios. ϕ -RTN avoids the explicit expansion of non-terminals at creation time and stores a single sub-graph per non-terminal, leading to a low disk footprint. Despite limitations (end of §4), we show that ϕ -RTN is complementary to LMs trained on the expanded grammar—allowing us to reap the benefits of both models—and improves WER on long tail entity queries by 10%. Future work includes the expansion of ϕ -RTN to non-English languages where grammatical agreements and morphology may require specialized solutions. **Acknowledgements** We thank Amr Mousa, Barry Theobald, Man-Hung Siu, and the anonymous reviewers for their comments and feedback.

6. References

- [1] J. Research, “Digital voice assistants in use to triple to 8 billion by 2023, driven by smart home devices,” Press Release, Feb. 2019.
- [2] Y. Maarek, “Alexa, can you help me shop?” in *SIGIR*, 2019, pp. 1369–1370.
- [3] A. Broder, “A taxonomy of web search,” in *SIGIR forum*, vol. 36, no. 2, 2002, pp. 3–10.
- [4] C. Van Gysel, M. Tsagkias, E. Pusateri, and I. Oparin, “Predicting entity popularity to improve spoken entity recognition by virtual assistants,” in *SIGIR*, 2020, pp. 1613–1616.
- [5] S. Gondala, L. Verwimp, E. Pusateri, M. Tsagkias, and C. Van Gysel, “Error-driven pruning of language models for virtual assistants,” in *ICASSP*. IEEE, 2021, pp. 7413–7417.
- [6] M. Saebi, E. Pusateri, A. Meghawat, and C. Van Gysel, “A discriminative entity-aware language model for virtual assistants,” in *Interspeech*, 2021, pp. 2032–2036.
- [7] H. Hayashi, Z. Hu, C. Xiong, and G. Neubig, “Latent relation language models,” in *ACAI*, 2020, pp. 7911–7918.
- [8] L. Velikovich, I. Williams, J. Scheiner, P. Aleksic, P. Moreno, and M. Riley, “Semantic lattice processing in contextual automatic speech recognition for google assistant,” in *Interspeech*, 2018, pp. 2222–2226.
- [9] P. F. Brown, V. J. Della Pietra, P. V. Desouza, J. C. Lai, and R. L. Mercer, “Class-based n-gram models of natural language,” *Computational linguistics*, vol. 18, no. 4, pp. 467–480, 1992.
- [10] M. A. B. Shaik, D. Rybach, S. Hahn, R. Schlüter, and H. Ney, “Hierarchical hybrid language models for open vocabulary continuous speech recognition using wfst,” in *SAPA@ Interspeech*, 2012, pp. 46–51.
- [11] A. Stolcke and J. Segal, “Precise n-gram probabilities from stochastic context-free grammars,” in *ACL*, 1994, pp. 74–79.
- [12] D. Jurafsky, C. Wooters, J. Segal, A. Stolcke, E. Fosler, G. Tajchaman, and N. Morgan, “Using a stochastic context-free grammar as a language model for speech recognition,” in *ICASSP*, vol. 1. IEEE, 1995, pp. 189–192.
- [13] A. Gandhe, A. Rastrow, and B. Hoffmeister, “Scalable language model adaptation for spoken dialogue systems,” in *SLT*. IEEE, 2018, pp. 907–912.
- [14] X. L. Aubert, “A brief overview of decoding techniques for large vocabulary continuous speech recognition,” in *ASR2000 ITRW*, 2000, pp. 91–96.
- [15] H. J. Dolfing and I. L. Hetherington, “Incremental language models for speech recognition using finite-state transducers,” in *ASRU*, 2001, pp. 194–197.
- [16] W. A. Woods, “Transition network grammars for natural language analysis,” *Communications of the ACM*, vol. 13, no. 10, pp. 591–606, 1970.
- [17] C. Allauzen and M. Riley, “A pushdown transducer extension for the openfst library,” in *CIAA*, 2012, pp. 66–77.
- [18] OpenFst maintainers, “OpenFst advanced usage: Matchers,” <https://www.openfst.org/twiki/bin/view/FST/FstAdvancedUsage\#Matchers>, accessed: 2022-03-04.
- [19] S. M. Katz, “Estimation of probabilities from sparse data for the language model component of a speech recognizer,” *ASSP*, vol. 35, pp. 400–401, 1987.
- [20] S. Achanta, A. Antony, L. Golipour, J. Li, T. Raitio, R. Rasipuram, F. Rossi, J. Shi, J. Upadhyay, D. Winarsky *et al.*, “On-device neural speech synthesis,” in *ASRU*, 2021, pp. 1155–1161.
- [21] C. Zhu, R. H. Byrd, P. Lu, and J. Nocedal, “Algorithm 778: L-bfgs-b: Fortran subroutines for large-scale bound-constrained optimization,” *TOMS*, vol. 23, no. 4, pp. 550–560, 1997.
- [22] A. Stolcke, “Entropy-based pruning of backoff language models,” in *DARPA Broadcast News Transcription and Understanding Workshop*, 1998, pp. 270–274.
- [23] C. Allauzen, M. Riley, J. Schalkwyk, W. Skut, and M. Mohri, “OpenFst: A general and efficient weighted finite-state transducer library,” in *CIAA*. Springer, 2007, pp. 11–23.
- [24] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz *et al.*, “The Kaldi speech recognition toolkit,” in *ASRU*. IEEE, 2011.
- [25] J. Sorensen and C. Allauzen, “Unary data structures for language models,” in *Interspeech*, 2011, pp. 1425–1428.
- [26] Z. Huang, T. Ng, L. Liu, H. Mason, X. Zhuang, and D. Liu, “SNDCNN: Self-normalizing deep cnns with scaled exponential linear units for speech recognition,” in *ICASSP*. IEEE, 2020, pp. 6854–6858.
- [27] E. Pusateri, C. Van Gysel, R. Botros, S. Badaskar, M. Hannemann, Y. Oualil, and I. Oparin, “Connecting and comparing language model interpolation techniques,” in *Interspeech*, 2019, pp. 3500–3504.
- [28] M. Bacchiani and B. Roark, “Unsupervised language model adaptation,” in *ICASSP*, vol. 1. IEEE, 2003, pp. 224–227.