



A Discriminative Entity-Aware Language Model for Virtual Assistants

Mandana Saebi^{1*}, Ernest Pusateri², Aaksha Meghawat², Christophe Van Gysel²

¹University of Notre Dame, Notre Dame, IN, USA

²Apple, Cupertino, CA, USA

msaebi@nd.edu, {epusateri, ameghawat, cvangysel}@apple.com

Abstract

High-quality automatic speech recognition (ASR) is essential for virtual assistants (VAs) to work well. However, ASR often performs poorly on VA requests containing named entities. In this work, we start from the observation that many ASR errors on named entities are inconsistent with real-world knowledge. We extend previous discriminative n-gram language modeling approaches to incorporate real-world knowledge from a Knowledge Graph (KG), using features that capture entity type-entity and entity-entity relationships. We apply our model through an efficient lattice rescoring process, achieving relative sentence error rate reductions of more than 25% on some synthesized test sets covering less popular entities, with minimal degradation on a uniformly sampled VA test set.

Index Terms: Speech Recognition, Language Modeling, Semantic knowledge, Spoken Entities

1. Introduction

Digital virtual assistants (VAs) continue to gain popularity and reach [1], propelled by their ability to respond to a wide range of requests made with natural human speech. High-quality automatic speech recognition (ASR) is essential for these systems to work well. However, an important subset of requests, those containing named entities, still presents a significant challenge to ASR [2]. This is primarily because it is difficult to build a language model (LM) that accurately models entities that are less popular or only recently popular, as they show up rarely or not at all in training data. See Figure 1 for an example of the problem we describe. Given the VA request “Play Canyon Moon by Harry Styles”, ASR returns a top choice hypothesis which substitutes an acoustically similar phrase (“Can you Moon”) for the correct entity name (“Canyon Moon”).

In this work, we start from the observation that many ASR named entity errors are inconsistent with real-world knowledge. Returning to Figure 1, knowing the songs performed by Harry Styles would make it obvious that the fourth hypothesis was the correct one. Building on this intuition, our approach uses real-world knowledge in the form of a knowledge graph (KG) to improve ASR performance on named entities.

We bring a KG to bear through a discriminative entity-aware language model (DEAL) which builds on an n-gram discriminative language modeling approach introduced in [3]. In that work, a log-linear model consumes features including a base model score and word n-gram indicators. The model is trained on n-best lists or lattices generated from a base system to minimize a discriminative loss function. We extend that approach




Figure 1: Example of a synthesized VA utterance and ASR hypothesis. The ASR top choice can be improved using knowledge of the relationship between Canyon Moon and Harry Styles.

by defining n-gram indicator features that fire when a word sequence within an ASR hypothesis is consistent with a KG. We design our features to capture entity type – entity relationships (e.g. “Canyon Moon” is a song) as well as relationships between entities (e.g. Canyon Moon is a song by Harry Styles.) We apply the model using an efficient lattice rescoring process. By applying this approach, we are able to achieve large improvements on targeted test sets covering unpopular entities with minimal degradation on a uniformly sampled VA test set.

Our main contributions are: (1) a discriminative language modeling approach for incorporating knowledge from a KG into an ASR system, (2) a performance analysis of different feature sets, and (3) an efficient algorithm to apply the approach via lattice rescoring.

2. Approach

Following [3], the model underlying our approach takes the form of a set of features and weights. The first feature is defined as the score from a base system. All other features are specified by n-grams. N-gram features fire when the corresponding n-gram matches a word sequence within a hypothesis. At test time, the features and weights are used to rescore n-best lists or lattices generated from the base system. The score given to hypothesis h in utterance u is shown in Equation 1, where w_f is the learned weight for feature f , and $x_{u,h,f}$ is the value of feature f , for hypothesis h , utterance u . $x_{u,h,0}$ always takes the value of the combined LM and acoustic model score from the base system. At training time, n-best lists are generated for a training set using the base system, and the weights are learned to minimize a discriminative loss function.

$$s(u, h) = \sum_{f=0}^F w_f x_{u,h,f} \tag{1}$$

In the following subsections, we describe how we extended [3] to incorporate information from a KG via enhanced features, and how we were able to use our approach to efficiently do lattice

*work done during the internship at Apple

id	n-gram	weight
f1	play \$title by	1.2
f2	play \$artist	0.8
f3	to \$city \$state city	-0.4




Figure 2: An example of features, their learned weights and the R_{top} built from them using Algorithm 1. $\$state|city$ represents a non-terminal covering states with a contains relation to the realization of the preceding $\$city$ non-terminal. The σ symbol in R_{top} matches all input symbols and reproduces them on the output.

```

"12345": {
  "names": { "Canyon Moon": { "word count": 2 } },
  "types": { "music title": { "popularity": 0.0025 } },
  "relationships": { {
    "relation": "performed by",
    "entity id": "67890",
    "popularity": 0.0021 } }
}
"67890": {
  "names": {
    "Harry Styles": { "word count": 2 },
    "Harry Edward Styles": { "word count": 3 }
  },
  "types": { "music artist": { "popularity": 0.5 } },
  "relationships": { {
    "relation": "performed",
    "entity id": "12345",
    "popularity": 0.1 } }
}

```

Figure 3: Example of two related entities in our KG.

rescoring.

2.1. Features

Our features reference data stored in our KG, so we will start with a description of the KG’s structure. Each node in our KG is an *entity*, which represents a real-world object. An entity has *names*, which are word sequences used to denote it, *types*, which are broad semantic classes to which the entities belong and *relationships*, which represent connections to other entities. A name has a *word count*, which is the number of words in the entity. A relationship has an *entity id*, which specifies the related entity, a *relation* which defines the nature of the relationship, and a *popularity*, which indicates the frequency with which the entity and its related entity appear together. Fig. 3 shows an example of two related entities.

Our features are specified as n-grams comprised of words and at least one *non-terminal*. A non-terminal represents the union of the entity names sharing an entity type, possibly with conditions. Non-terminals can be conditioned on entity name word count or entity popularity, as well as entities whose names were previously seen in a hypothesis. A feature takes the value 1 for a word sequence in a hypothesis when the word sequence matches a valid realization of the n-gram. A realization is a word sequence generated by replacing each non-terminal with a corresponding entity name. If the feature does not match a valid realization, the value is set to 0.

To take into account entity popularity, for each entity type, we create three non-terminals, the first covering the most popular (head) entities, the second adding moderately popular (torso) entities, and the last adding the least popular (tail) entities. We also create three non-terminals for each entity type based on entity name word count, covering entity names with one or more words, two or more words, and three or more words. The entity lists corresponding to each conditional non-terminal intentionally overlap to reduce the number of features with few training examples. By using entity popularity and word count, we hoped to expand the space of trade-offs the model could make. For

instance, we hypothesized that the model can be more aggressive in boosting hypotheses for popular entities versus less popular ones as the risk of causing an error by doing so would be lower. Similarly, for entity name word count, we surmised that the model could more aggressively reward hypotheses with longer entity names, as it would be less likely for these to show up incorrectly in an ASR lattice.

To choose feature n-grams, we first derive text templates with frequency estimates. These templates consist of entity type non-terminals (without conditions) and surrounding word context and cover the most common requests where the targeted entity types or related entity type pairs appeared. For instance, the most common template for the city and state entity type pair was “Directions to \$city \$state”. From these templates, we extract all 3-grams containing a non-terminal. In order to better model entity relationships, we also include all 4-grams that begin and end with a non-terminal. We then factor the resulting n-grams into multiple n-grams with conditional non-terminals. Examples of feature n-grams can be found in Figure 2.

2.2. Lattice Rescoring

While it is straightforward to use our approach to do n-best rescoring, we also developed a way of using it to rescore lattices. We prefer rescoring lattices to n-best lists because lattices can represent a larger set of hypotheses more efficiently, and an efficient lattice rescoring scheme can avoid the redundant computations that occur with n-best rescoring due to the large overlap between nearby hypotheses.

In the approach described in [3], a lattice, represented as a weighted finite-state automaton (WFSA) [4], is rescored by composing it with a static, deterministic WFSA representing the learned model. That approach does not work for us here, as it would require expanding the non-terminals in all of the features for all of the corresponding entities in the KG. This would produce prohibitively large WSFAs for many use cases and also require that the resulting rescoring WFSA be rebuilt whenever the KG changed. To overcome these difficulties, we developed an approach that constructs a deterministic WFSA on-demand.

Our deterministic rescoring WFSA, R_{det} , is built by determining a weighted finite-state transducer (WFST), R_{kg} , on-demand. R_{kg} is constructed on-demand from another WFST, R_{top} , and the KG. R_{top} is built offline and can be thought of as a top-level feature tagging grammar. It is built from the model features and weights according to Algorithm 1. An example R_{top} is shown in Figure 2. R_{kg} is built by traversing R_{top} and replacing arcs labeled with non-terminals with WSFAs constructed from their corresponding entity names from the KG. In order to support non-terminals whose realizations depend on previously seen entities, non-terminal/realization pairs are stored as part of the dynamically created R_{kg} states. Non-terminal realizations are propagated until a feature end tag is reached. The result of composing a lattice with R_{kg} would be a WFST with two important properties: (1) exactly one sub-path would exist for

any lattice sub-path and n-gram feature realization spanning that sub-path, and (2) all arcs except those with feature end tags would have 0 weights.

Input: *features*, a list of *id*, *ngram*, *weight* triples
 Initialize R_{top} with start and final state, s_{top}
 Add arc from s_{top} to s_{top} with symbols $\sigma:\sigma$
foreach *id*, *ngram*, *weight* in *features* **do**
 Construct a WFSA that accepts *ngram*, with start state s_f , and end state e_f
 Add arc from s_{top} to s_f with symbols $\langle f \rangle:\langle f \rangle$
 Add arc from e_f to s_{top} with symbols $\epsilon:\langle id \rangle$ and weight *weight*
end

$R_{top} \leftarrow \text{Decode}(\text{Determinize}(\text{Encode}(R_{top})))$

Algorithm 1: Algorithm for constructing R_{top} from feature n-grams and weights. `Encode` encodes arc input symbols, output symbols and weights in the arc symbols. `Decode` reverses that operation.

R_{det} is built by determinizing R_{kg} on-demand. Note that this determinization operation is not equivalent to conventional determinization on the tropical or log semiring, as we want to apply the same operation, addition, to feature weights on consecutive arcs as well as to paths that converge to the same state. However, because we have ensured the properties above, the on-demand determinization process is straightforward. When the composition process requests an arc weight and end state from R_{det} for a start state, s and symbol i , the on-demand generation process will find all states in R_{kg} corresponding to s and find the set of arcs, A , with input symbol i exiting those states. A new R_{det} arc will be created with a weight equal to the sum of the weights on A . The end state will be a new R_{det} state associated with all of the R_{kg} end states on the arcs in A , and the input and output symbols will be i .

3. Related Work

Our approach was particularly inspired by [5], which presents a generative, neural network LM that incorporates a KG. A similar approach is presented in [6]. Roughly speaking, one can view our work as a non-neural analog to those approaches with a discriminative optimization criterion.

While they do not use a KG as their knowledge source, [7], [8], [9] and [10] share many similarities with our work. The models in all of those works can be viewed as log-linear models which consume n-gram features. Further, WFSTs are used to encode the models and WFST composition is used to apply them, and all of those works focus on improving ASR for virtual assistants. However, those approaches differ from ours in that model weights are determined by heuristics, sometimes in combination with context-specific n-gram LMs, rather than through discriminative training. The approaches in [7], [8], [9] additionally differ in that the models are applied early in the recognition process rather than during lattice rescoring (an advantage, all else equal.) The approach in [10] is especially similar to our work in that their analog to n-gram features can contain non-terminals. However, in [10], lattices are semantically tagged to determine the locations of the non-terminals before being rescored, and their n-gram features do not capture entity-entity relationships.

4. Experiments

4.1. Data

We tested our model using US English data. Our KG covered 5 entity types: city, state, music title, music artist, and soccer athlete. Cities and states were connected with *is in* and *contains* relations. Music titles and music artists were connected with *performed by* and *performed* relations. Our KG provided usage-based popularity estimates for each entity name and relationship as well as word counts for every entity name. We chose these entity types and relations for our experiments to cover a variety of domains, popularity distributions, and entity name word counts.

Our training and testing data consisted of a combination of anonymized general VA data and synthesized data covering the targeted entities and relations. We synthesized data to avoid the large expense of recording and manually transcribing data with adequate entity/relation coverage. From the general VA data, we used a manually transcribed uniform sample for training and testing.

To derive the templates and frequency estimates described in § 2.1, we started from a larger pool of automatically transcribed and named-entity-tagged VA requests. We first replaced regions tagged as named entities in each utterance with non-terminals, then extracted candidate templates sorted by count. We then hand-curated the template lists, removing erroneous templates and cutting off the template lists when the templates became obscure. Through inspection, we determined that 100 templates was a reasonable threshold for most entity types and entity type pairs. The exceptions were soccer athletes and cities where only 10 and 79 templates were chosen, respectively.

To synthesize targeted data, as well as to support popularity-conditioned non-terminals, we divided the entities of each entity type into three strata, head, torso, and tail, according to their popularity estimates. We determined the thresholds for the strata by inspecting the popularity distributions. All of the distributions were L-shaped, declining steeply at the start and flattening at the end, with a gradual curve in the middle. For all entity types and entity type pairs, the steep decline ended at about the 100th entity. This determined the threshold for the head. To set the lower bound on the torso, we chose the point at which the popularity distribution appeared to flatten out. We then included up to 60,000 entities in the tail. The number of entities in each stratum for each entity type and entity type pair is provided in Table 1. To synthesize targeted data for each stratum, we generated text data probabilistically according to the template frequency estimates and entity popularities in the KG. After generating text utterances, we synthesized audio for each text utterance using our previous generation speech synthesizer, a unit selection system described in [11].

To train the DEAL models, we used 7000 utterances from the general VA data and 2000 synthesized utterances for every entity type or entity type pair and stratum. To allay concerns that our approach would only work on the voice used for synthesizing our training data, we used a female voice to synthesize the training data and a male voice to synthesize test data.

4.2. Baseline Systems

Our ASR system used a deep convolutional neural network acoustic model [12], a 4-gram LM with Good-Turing smoothing in the first pass, and the same LM interpolated with a Feed-Forward Neural Network (FFNN) LM [13] in the second pass. For training and testing, we used the n-bests lists and lattices resulting from the second pass. We trained our models starting from two baseline systems, baseline-1 and baseline-2, which

Table 1: Experiment results. # entities indicates the number of unique entities in each dataset and # relations indicate the number of unique entity-entity relations. Bold numbers indicate the lowest sentence error rate (SER) w.r.t each baseline.

Data	Tier	# entities	# relations	Baseline-1	DEAL1	DEAL1-r	DEAL1-rc	DEAL1-rp	DEAL1-rpc	Oracle-1	Baseline-2	DEAL2	DEAL2-rpc	Oracle-2
City/State	Head	138	100	2.71	0.69	0.69	0.58	0.53	0.47	0.39	2.65	0.77	0.55	0.39
	Torso	1746	1900	5.13	2.83	2.49	2.42	2.24	2.21	1.62	4.08	2.02	1.96	1.42
	Tail	26946	40793	18.65	15.40	13.45	13.42	13.41	13.43	11.07	18.26	12.97	12.94	10.59
City	Head	100	—	3.94	2.12	2.31	2.28	2.02	1.79	1.99	3.58	2.05	1.89	1.69
	Torso	1900	—	9.50	7.02	7.06	7.01	6.20	5.91	2.30	8.63	6.57	5.70	1.87
	Tail	25437	—	29.70	22.84	22.69	22.72	23.51	23.44	13.58	29.15	22.86	22.57	13.12
Music Title/Artist	Head	178	100	12.70	9.97	9.20	9.04	8.78	8.99	8.27	6.87	5.97	5.89	5.46
	Torso	5405	3900	17.48	14.53	12.79	12.86	12.91	12.93	11.78	10.53	8.74	8.79	8.03
	Tail	82769	60000	57.70	50.19	45.01	45.27	45.26	45.28	43.78	51.04	40.17	40.05	38.50
Music Title	Head	100	—	11.42	8.51	8.37	8.78	6.38	7.00	3.02	10.53	7.49	6.68	2.67
	Torso	3900	—	12.66	9.97	9.90	9.57	9.59	9.15	4.12	11.14	8.40	8.20	3.25
	Tail	60000	—	48.35	48.18	47.92	47.77	48.21	47.81	34.42	45.89	45.68	45.68	32.73
Music Artist	Head	100	—	11.37	8.95	8.95	8.90	8.36	8.08	5.34	9.24	7.23	6.13	4.23
	Torso	900	—	14.30	11.76	11.81	12.00	11.53	11.41	8.65	11.27	9.80	9.45	6.92
	Tail	60000	—	59.59	100.00	57.74	57.37	57.66	57.66	45.19	52.05	51.27	51.02	40.04
Soccer Athletes	Head	100	—	6.25	5.44	5.44	5.44	5.44	5.44	5.44	4.84	3.83	3.83	3.83
	Torso	300	—	17.12	12.65	12.65	12.72	12.57	12.57	11.71	15.10	10.48	10.41	9.68
	Tail	60000	—	78.63	62.90	62.95	62.88	63.16	63.08	62.48	78.43	62.74	62.85	62.42
General VA	—	—	12.04	12.16	12.11	12.11	12.09	12.14	5.90	12.07	12.11	12.11	5.93	

Table 2: Overview of variations being compared.

Variation	Description	# features
DEAL	entity type – entity	513
DEAL-r	DEAL + entity-entity	560
DEAL-rc	DEAL-r + entity name word count	2278
DEAL-rp	DEAL-r + entity popularity	2278
DEAL-rpc	DEAL-rp \cup DEAL-rc	6535

Table 3: Examples of improvements on synthetic utterances resulting from DEAL.

Baseline-2	DEAL2-rpc
directions to Hammers Texas	directions to Amherst Texas
turn on the traffic	Toronto traffic
play Cinderella by young Jesus	play Cinderella by Yung Zeus
lyrics tablet on my jeans	lyrics to Blood on my jeans
play Brett Eldridge	play Brett Eldridge
cadence on cavity age	Edinson Cavani age

differed in their 4-gram LMs.

To build the 4-gram LM for baseline-1, we started by building component 4-gram models from data sources consisting of more than 10B manually and automatically transcribed anonymized VA requests. We then combined these components using count merging to minimize perplexity on held out manually transcribed VA data [14]. In baseline-2, we followed the same process but added LM components trained on synthetic text data. We generated the synthetic text in a way similar to how we generated the synthesized training data for our model but in much larger quantities. We generated 100 million utterances for each entity type and entity type pair and trained a separate 4-gram LM component on each data set. (This approach is similar to one of the approaches described in [15], where domain-specific LM components are built based on grammars and interpolated with more general components.)

4.3. Results

Table 1 show our experimental results in terms of Sentence Error Rate (SER) on test data using the model variations outlined in Table 2. We show results for all variations on top of baseline-1 (DEAL1) and selected variations on top of baseline-2 (DEAL2). We use SER (as opposed to WER) to measure performance, since getting an entity name only partially correct will often result in an incorrect VA response. We also include the oracle SER, which indicates the maximum possible improvement that can be achieved by rescored the baseline lattices. We can see

that SER of Oracle-2 is generally lower than Oracle-1, resulting from the inclusion of targeted synthetic data in the baseline LM. We also observe an increase in SER as we move from head to tail and torso across all models, since less popular entities will be less well represented in the general VA data used to train the baseline LMs.

We see that using DEAL1 results in significant improvement over baseline-1 on the synthetic datasets with minimal degradation on general VA queries. Comparing DEAL1-r and DEAL1, we see additional gains using DEAL1-r, mainly on test sets containing related entity pairs (Music Title/Artist and City/State).

We see improvements over DEAL1-r by adding word count (DEAL1-rc) and popularity (DEAL1-rp) features. Across most test sets DEAL1-rp outperforms DEAL1-r and DEAL1-rc. Combining both popularity and word count features (DEAL1-rpc) results in lower SER than either feature set on its own.

Before looking at the results for baseline-2 model variations, it is interesting to compare the performance of DEAL1-rpc and baseline-2. We can see that DEAL1-rpc outperforms baseline-2 across several datasets, despite lacking the additional hypotheses available in baseline-2 lattices.

Finally, comparing DEAL2 and baseline-2, we notice that our method can achieve SER reductions even over a stronger baseline. We see further improvements by adding the entity-entity and including both word count and popularity features (DEAL2-rpc). Table 3 shows utterances for which DEAL2-rpc was able to correct the error made by the baseline-2 model.

5. Conclusions

We introduced a discriminative LM that incorporates real-world knowledge from a KG via n-gram features and a method for efficiently applying the model with lattice rescoring. We evaluated the impact of using features that captured entity-entity relations, as well as features that took into account entity popularity and entity name word count. Our approach achieved relative SER reductions of more than 25% on some synthetic test sets covering less common entities while causing minimal degradation on a general VA test set. Including features that captured entity-entity relations led to large SER reductions on synthesized test sets covering related entity pairs, while including features conditioned on popularity and word count resulted in a modest SER reduction over all synthesized entity test sets.

6. References

- [1] Juniper Research, “Digital Voice Assistants in Use to Triple to 8 Billion by 2023, Driven by Smart Home Devices,” Press Release, Feb. 2019.
- [2] 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.
- [3] B. Roark, M. Saraclar, M. Collins, and M. Johnson, “Discriminative language modeling with conditional random fields and the perceptron algorithm,” in *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics*, 2004, pp. 47–54.
- [4] M. Mohri, F. Pereira, and M. Riley, “Weighted finite-state transducers in speech recognition,” *Computer Speech & Language*, vol. 16, no. 1, pp. 69–88, 2002.
- [5] R. L. Logan IV, N. F. Liu, M. E. Peters, M. Gardner, and S. Singh, “Barack’s wife hillary: Using knowledge-graphs for fact-aware language modeling,” in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2019, pp. 5962–5971.
- [6] H. Hayashi, Z. Hu, C. Xiong, and G. Neubig, “Latent relation language models,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 05, 2020, pp. 7911–7918.
- [7] K. Hall, E. Cho, C. Allauzen, F. Beaufays, N. Coccaro, K. Nakajima, M. Riley, B. Roark, D. Rybach, and L. Zhang, “Composition-based on-the-fly rescoring for salient n-gram biasing,” in *Interspeech*, 2015, pp. 1418–1422.
- [8] J. Scheiner, I. Williams, and P. Aleksic, “Voice search language model adaptation using contextual information,” in *Spoken Language Technology Workshop*. IEEE, 2016, pp. 253–257.
- [9] P. Aleksic, M. Ghodsi, A. Michaely, C. Allauzen, K. Hall, B. Roark, D. Rybach, and P. Moreno, “Bringing contextual information to google speech recognition,” in *Interspeech*, 2015, pp. 468–472.
- [10] L. Velikovich, I. Williams, J. Scheiner, P. S. Aleksic, P. J. Moreno, and M. Riley, “Semantic lattice processing in contextual automatic speech recognition for google assistant,” in *Interspeech*, 2018, pp. 2222–2226.
- [11] T. Capes, P. Coles, A. Conkie, L. Golipour, A. Hadjitarkhani, Q. Hu, N. Huddleston, M. Hunt, J. Li, M. Neeracher *et al.*, “Siri on-device deep learning-guided unit selection text-to-speech system,” in *Interspeech*, 2017, pp. 4011–4015.
- [12] Z. Huang, T. Ng, L. Liu, H. Mason, X. Zhuang, and D. Liu, “SND-CNN: Self-Normalizing Deep CNNs with Scaled Exponential Linear Units for Speech Recognition,” in *International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2020, pp. 6854–6858.
- [13] S. Zhang, H. Jiang, M. Xu, J. Hou, and L. Dai, “The Fixed-Size Ordinally-Forgetting Encoding Method for Neural Network Language Models,” in *IJCNLP*, 2015, pp. 495–500.
- [14] M. Bacchiani and B. Roark, “Unsupervised language model adaptation,” in *International Conference on Acoustics, Speech, and Signal Processing*, vol. 1. IEEE, 2003, pp. 224–227.
- [15] A. Gandhe, A. Rastrow, and B. Hoffmeister, “Scalable language model adaptation for spoken dialogue systems,” in *Spoken Language Technology Workshop*. IEEE, 2018, pp. 907–912.