

Supervised Domain Enablement Attention for Personalized Domain Classification

Joo-Kyung Kim Young-Bum Kim

Amazon Alexa

{jookyk, youngbum}@amazon.com

Abstract

In large-scale domain classification for natural language understanding, leveraging each user’s domain enablement information, which refers to the preferred or authenticated domains by the user, with attention mechanism has been shown to improve the overall domain classification performance. In this paper, we propose a supervised enablement attention mechanism, which utilizes sigmoid activation for the attention weighting so that the attention can be computed with more expressive power without the weight sum constraint of softmax attention. The attention weights are explicitly encouraged to be similar to the corresponding elements of the ground-truth’s one-hot vector by supervised attention, and the attention information of the other enabled domains is leveraged through self-distillation. By evaluating on the actual utterances from a large-scale IPDA, we show that our approach significantly improves domain classification performance.

1 Introduction

Due to recent advances in deep learning techniques, intelligent personal digital assistants (IPDAs) such as Amazon Alexa, Google Assistant, Microsoft Cortana, and Apple Siri have been widely used as real-life applications of natural language understanding (Sarıkaya et al., 2016; Sarıkaya, 2017).

In natural language understanding, domain classification is a task that finds the most relevant domain given an input utterance (Tur and de Mori, 2011). For example, “make a lion sound” and “find me an apple pie recipe” should be classified as *ZooKeeper* and *AllRecipe*, respectively. Recent IPDAs cover more than several thousands of diverse domains by including third-party developed domains such as Alexa Skills (Kumar et al., 2017; Kim et al., 2018a; Kim and Kim, 2018),

Google Actions, and Cortana Skills, which makes domain classification to be a more challenging task.

Given a large number of domains, leveraging user’s enabled domain information¹ has been shown to improve the domain classification performance since enabled domains reflect the user’s context in terms of domain usage (Kim et al., 2018b). For an input utterance, Kim et al. (2018b) use attention mechanism so that a weighted sum of the enabled domain vectors are used as an input signal as well as the utterance vector. The enabled domain vectors and the attention weights are automatically trained in an end-to-end fashion to be helpful for the domain classification.

In this paper, we propose a supervised enablement attention mechanism for more effective attention on the enabled domains. First, we use logistic sigmoid instead of softmax as the attention activation function to relax the constraint that the weight sum over all the enabled domains is 1 to the constraint that each attention weight is between 0 and 1 regardless of the other weights (Martins and Astudillo, 2016; Kim et al., 2017). Therefore, all the attention weights can be very low if there are no enabled domains relevant to a ground-truth so that we can disregard the irrelevant enabled domains, and multiple attention weights can have high values when multiple enabled domains are helpful for disambiguating an input utterance. Second, we encourage each attention weight to be high if the corresponding enabled domain is a ground-truth domain and if otherwise, to be low, by a supervised attention method (Mi et al., 2016) so that the attention weights can be directly tuned for the downstream classification task. Third, we

¹Enabled domain information specifically refers to preferred or authenticated domains in Amazon Alexa, but it can be extended to other information such as the list of recently used domains.




Figure 1: Model architecture: the input utterance is represented as a dense vector through word embedding and BiLSTM. Domain enablement vector is computed as a weighted sum of enabled domain vectors through the proposed attention mechanism. The two vectors are concatenated for the final domain prediction through a feed-forward neural network.

apply self-distillation (Furlanello et al., 2018) on top of the enablement attention weights so that we can better utilize the enabled domains that are not ground-truth domains but still relevant.

Evaluating on datasets obtained from real usage in a large-scale IPDA, we show that our approach significantly improves domain classification performance by utilizing the domain enablement information effectively.

2 Model

Figure 1 shows the overall architecture of the proposed model.

Given an input utterance, each word of the utterance is represented as a dense vector through word embedding followed by bidirectional long short-term memory (BiLSTM) (Graves and Schmidhuber, 2005). Then, an utterance vector is composed by concatenating the last outputs of the forward LSTM and the backward LSTM.²

To represent the domain enablement information, we obtain a weighted sum of domain enablement vector where the weights are calculated by logistic sigmoid function on top of the multiplicative attention (Luong et al., 2015) for the utterance vector and the domain enablement vectors. The attention weight of an enabled domain e is formu-

lated as follows:

$$a_e = \sigma(u \cdot v_e),$$

where u is the utterance vector, v_e is the enablement vector of enabled domain e , and σ is sigmoid function. Compared to conventional attention mechanism using softmax function, which constraints the sum of the attention weights to be 1, sigmoid attention has more expressive power, where each attention weight can be between 0 and 1 regardless of the other weights. We show that using sigmoid attention is actually more effective for improving prediction performance in Section 3.

The utterance vector and the weighted sum of the domain enablement vectors are concatenated to represent the utterance and the domain enablement as a single vector. Given the concatenated vector, a feed-forward neural network with a single hidden layer³ is used to predict the confidence score by logistic sigmoid function for each domain.

One issue of the proposed architecture is that the domain enablement can be trained to be a very strong signal, where one of the enabled domains would be the predicted domains regardless of the relevancy of the utterances to the predicted domains in many cases. To reduce this prediction bias, we use randomly sampled enabled domains

²We have also evaluated word vector summation, CNN (Kim, 2014), BiLSTM mean-pooling, and BiLSTM max-pooling (Conneau et al., 2017) as alternative utterance representation methods, but they did not show better performance on our task.

³We utilize scaled exponential linear units (SeLU) as the activation function for the hidden layer (Klambauer et al., 2017).

instead of the correct enabled domains of an input utterance with 50% probability during training so that the domain enablement is used as an auxiliary signal rather than determining signal. During inference, we always use the correct domain enablements of the given utterances.

The main loss function of our model is formulated as binary log loss between the confidence score and the ground-truth vector as follows:

$$\mathcal{L}_m = - \sum_{i=1}^n y_i \log o_i + (1 - y_i) \log (1 - o_i),$$

where n is the number of all domains, o is an n -dimensional confidence score vector from the model, and y is an n -dimensional one-hot vector whose element corresponding to the position of the ground-truth domain is set to 1.

2.1 Supervised Enablement Attention

Attention weights are originally intended to be automatically trained in an end-to-end fashion (Bahdanau et al., 2015), but it has been shown that applying proper explicit supervision to the attention improves the downstream tasks such as machine translation given the word alignment and constituent parsing given annotations between surface words and nonterminals (Mi et al., 2016; Liu et al., 2016; Kamigaito et al., 2017).

We hypothesize that if the ground-truth domain is one of the enabled domains, the attention weight for the ground-truth domain should be high and vice versa. To apply this hypothesis in the model training as a supervised attention method, we formulate an auxiliary loss function as follows:

$$\mathcal{L}_a = - \sum_{e \in E} y_e \log a_e + (1 - y_e) \log (1 - a_e),$$

where E is a set of enabled domains and a_e is the attention weight for the enabled domain e .

2.2 Self-Distilled Attention

One issue of supervised attention in Section 2.1 is that enabled domains that are not ground-truth domains are encouraged to have lower attention weights regardless of their relevancies to the input utterances and the ground-truth domains. Distillation methods utilize not only the ground-truth but also all the output activations of a source model so that all the prediction information from the source model can be utilized for more effective knowledge transfer between the source model

and the target model (Hinton et al., 2014). Self-distillation, which trains a model leveraging the outputs of the source model with the same architecture or capacity, has been shown to improve the target model’s performance with a distillation method (Furlanello et al., 2018).

We use a variant of self-distillation methods, where the model outputs at the previous epoch with the best dev set performance are used as the soft targets for the distillation,⁴ so that the enabled domains that are not ground-truths can also be used for the supervised attention. While conventional distillation methods utilize softmax activations as the target values, we show that distillation on top of sigmoid activations is also effective without loss of generality. The loss function for the self-distillation on the attention weights is formulated as follows:

$$\mathcal{L}_d = - \sum_{e \in E} \tilde{a}_e \log a_e + (1 - \tilde{a}_e) \log (1 - a_e),$$

where \tilde{a}_e is the attention weight of the model showing the dev set performance in the previous epochs. It is formulated as:

$$\tilde{a}_e = \sigma \left(\frac{u \cdot v_e}{T} \right),$$

where T is the temperature for sufficient usage of all the attention weights as the soft target. In this work, we set T to be 16, which shows the best dev set performance.

We have also evaluated soft-target regularization (Aghajanyan, 2017), where a weighted sum of the hard ground-truth target vector and the soft target vector is used as a single target vector, but it did not show better performance than self-distillation.

All the described loss functions are added to compose a single loss function as follows:

$$\mathcal{L} = \mathcal{L}_m + \alpha \{ (1 - \beta) \mathcal{L}_a + \beta^t \mathcal{L}_d \},$$

where α is a coefficient representing the degree of supervised enablement attention and β^t denotes the degree of the self-distillation. We set α to be 0.01 in this work. Following Hu et al. (2016), $\beta^t = 1 - 0.95^t$, where t denotes the current training epoch starting from 0 so that the hard ground-truth targets are more influential in the early epochs and the self-distillation is more utilized in the late epochs.

⁴This approach is closely related to Temporal Ensembling

Model no	Attention method	Biased ground-truth inclusion			Unbiased ground-truth inclusion		
		Top1	MRR	Top3	Top1	MRR	Top3
(1)	sfm	95.81	97.27	99.08	90.65	93.60	97.31
(2)	sgmd	95.98	97.43	99.19	91.03	93.92	97.49
(3)	sgmd, spvs	96.10	97.50	99.21	91.11	93.98	97.53
(4)	sgmd, spvs, sdst	96.29	97.65	99.32	91.33	94.14	97.62
(5)	sfm, bias	97.01	98.26	99.75	90.07	93.03	96.84
(6)	sgmd, spvs, sdst, bias	97.48	98.51	99.76	90.58	93.30	96.73

Table 1: Accuracies (%) on a test set with biased ground-truth inclusion in the enabled domains (90%) (left) and a test set with unbiased inclusion (70%) (right) with various enablement attention methods. **sfm**, **sgmd**, **spvs**, **sdst**, and **bias** denote softmax, sigmoid, supervised, self-distilled, and domain enablement bias, respectively.

Utterance	Ground-truth	Enabled domain: [attention weights for model (1), (2), and (4)], ...
what is the price of bitcoin	Crypto Price	Sleep and Relaxation Sounds: [0.9998, 0.0004, 0.2029], Crypto Price: [0.0001, 9.21e-0.6, 0.9977]
find me a round trip ticket flight	Expedia	Expedia: [0.0048, 5.37e-08, 0.6205], KAYAK: [0.9952, 0.0004, 0.461]
find my phone	Find My Phone	The Name Game: [1.0, 0.0001, 0.1677]

Table 2: Sample utterances correctly predicted with model (4) but not with model (1) and (2).

3 Experiments

We evaluate our proposed model on domain classification leveraging enabled domains. The enabled domains can be a crucial disambiguating signal especially when there are multiple similar domains. For example, assume that the input utterance is “*what’s the weather*” and there are multiple weather-related domains such as `NewYorkWeather`, `AccuWeather`, and `WeatherChannel`. In this case, if `WeatherChannel` is included as an enabled domain of the current user, it is likely to be the most relevant domain to the user.

3.1 Datasets

Following the data collection methods used in [Kim et al. \(2018b\)](#), our models are trained using utterances with explicit invocation patterns. For example, given a user’s utterance, “*Ask {ZooKeeper} to {play peacock sound}*,” “*play peacock sound*” and `ZooKeeper` are extracted to compose a pair of the utterance and the ground-truth, respectively. In this way, we have generated train, development, and test sets containing 4.4M, 500K, and 500K utterances, respectively. All the utterances are from the usage log of Amazon Alexa and the ground-truth of each utterance is one of 1K frequently used domains. The average number of enabled domains per utterance in the test sets is 8.47.

One issue of this collected data sets is that the

[\(Laine and Aila, 2017\)](#), but we just leverage the model outputs at the previous epoch rather than accumulating the outputs over multiple epochs.

ground-truth is included in the enabled domains for more than 90% of the utterances, where the ground-truths are biased to enabled domains.⁵ For more correct and unbiased evaluation of the models on the input utterances from real live traffic, we also evaluate the models on the same sized train, development, and test sets where the utterances are sampled to set the ratio of ground-truth inclusion in enabled domains to be 70%, which is closer to the ratio for actual input traffic.

3.2 Results

Table 1 shows the accuracies of our proposed models on the two test sets. We also show mean reciprocal rank (MRR) and top-3, accuracy⁶ which is meaningful when utilizing post reranker, but we do not cover reranking issues in this paper ([Robichaud et al., 2014](#); [Kim et al., 2018a](#)).

From Table 1, we can first see that changing softmax attention to sigmoid attention significantly improves the performance. This means that having more expressive power for the domain enablement information by relaxing the softmax constraint is effective in terms of leveraging the domain enablement information for domain classification. Along with sigmoid attention, supervised attention leveraging ground-truth slightly improves the performance, and supervised attention combined with self-distillation shows significant performance improvement. It demon-

⁵Since the data collection method leverages utterances where users already know the exact domain names, such domains are likely to be the enabled domains of the users.

⁶Top-3 accuracy is calculated as # (utterances one of whose top three predictions is a ground-truth) / # (total utterances).

strates that supervised domain enablement attention leveraging ground-truth enabled domains is helpful, and utilizing attention information from other enabled domains is synergistic.

Kim et al. (2018b)’s model also adds a domain enablement bias vector to the final output, which is helpful when the ground-truth domain is one of the enabled domains. Such models (5) and (6) also show good performance for the test set where the ground-truth is one of the enabled domains with more than 90% probability. However, for the unbiased test set where the ground-truth is included in the enabled domains with a smaller probability, not adding the bias vector is shown to be better overall.

Table 2 shows sample utterances correctly predicted with model (4) but not with model (1) and (2). For the first two utterances, the ground-truths are included in the enabled domains, but there were only hundreds or fewer training instances whose ground-truths are `CryptoPrice` or `Expedia`. In these cases, we can see that model (1) attends to unrelated domains, model (2) attends to none of the enabled domains, but model (4), which uses supervised attention, is shown to attend to the ground-truth even without many training examples. “*find my phone*” has a single enabled domain which is not a ground-truth. In this case, model (1) still fully attends to the unrelated domain because of softmax attention while model (2) and (4) do not highly attend to it so that the unrelated enabled domain is not impactful.

3.3 Implementation Details

The word vectors are initialized with off-the-shelf GloVe vectors (Pennington et al., 2014), and all the other model parameters are initialized with Xavier initialization (Glorot and Bengio, 2010). Each model is trained for 25 epochs and the parameters showing the best performance on the development set are chosen as the model parameters. We use ADAM (Kingma and Ba, 2015) for the optimization with the initial learning rate 0.0002 and the mini-batch size 128. We use gradient clipping, where the threshold is set to 5. We use a variant of LSTM, where the input gate and the forget gate are coupled and peephole connections are used (Gers and Schmidhuber, 2000; Greff et al., 2017). We also use variational dropout for the LSTM regularization (Gal and Ghahramani, 2016). All the models are implemented with DyNet (Neubig et al.,

2017).

4 Conclusion

We have introduced a novel domain enablement attention mechanism improving domain classification performance utilizing domain enablement information more effectively. The proposed attention mechanism uses sigmoid attentions for more expressive power of the attention weights, supervised attention leveraging ground-truth information for explicit guidance of the attention weight training, and self-distillation for the attention supervision leveraging enabled domains that are not ground truth domains. Evaluating on utterances from real usage in a large-scale IPDA, we have demonstrated that our proposed model significantly improves domain classification performance by better utilizing domain enablement information.

References

- Armen Aghajanyan. 2017. Softtarget regularization: An effective technique to reduce over-fitting in neural networks. In *IEEE Conference on Cybernetics (CYBCONF)*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *International Conference on Learning Representations (ICLR)*.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. Supervised Learning of Universal Sentence Representations from Natural Language Inference Data. In *EMNLP*, pages 670–680.
- Tommaso Furlanello, Zachary C. Lipton, Michael Tschannen, Laurent Itti, and Anima Anandkumar. 2018. Born Again Neural Networks. In *International Conference on Machine Learning (ICML)*, pages 1602–1611.
- Yarin Gal and Zoubin Ghahramani. 2016. A theoretically grounded application of dropout in recurrent neural networks. In *Advances in Neural Information Processing Systems 29 (NIPS)*, pages 1019–1027.
- Felix A. Gers and Jürgen Schmidhuber. 2000. Recurrent Nets that Time and Count. In *IJCNN*, volume 3, pages 189–194.
- Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 249–256.

- Alex Graves and Jürgen Schmidhuber. 2005. Frame-wise phoneme classification with bidirectional LSTM and other neural network architectures. *Neural Networks*, 18(5):602–610.
- Klaus Greff, Rupesh Kumar Srivastava, Jan Koutník, Bas R. Steunebrink, and Jürgen Schmidhuber. 2017. LSTM: A search space odyssey. *Transactions on Neural Network Learning and Systems (TNNLS)*, 28(10):2222–2232.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2014. Distilling the knowledge in a neural network. In *NIPS 2014 Deep Learning Workshop*.
- Zhiting Hu, Xuezhe Ma, Zhengzhong Liu, Eduard Hovy, and Eric Xing. 2016. Harnessing Deep Neural Networks with Logic Rules. In *ACL*, pages 2410–2420.
- Hidetaka Kamigaito, Katsuhiko Hayashi, Tsutomu Hirao, and Masaaki Nagata. 2017. Supervised attention for sequence-to-sequence constituency parsing. In *International Joint Conference on Natural Language Processing (IJCNLP)*.
- Joo-Kyung Kim and Young-Bum Kim. 2018. Joint learning of domain classification and out-of-domain detection with dynamic class weighting for satisficing false acceptance rates. In *Interspeech*.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1292–1302.
- Yoon Kim, Carl Denton, Luong Hoang, and Alexander M. Rush. 2017. Structured Attention Networks. In *International Conference on Learning Representations (ICLR)*.
- Young-Bum Kim, Dongchan Kim, Joo-Kyung Kim, and Ruhi Sarikaya. 2018a. A scalable neural shortlisting-reranking approach for large-scale domain classification in natural language understanding. In *NAACL*, pages 16–24.
- Young-Bum Kim, Dongchan Kim, Anjishnu Kumar, and Ruhi Sarikaya. 2018b. Efficient Large-Scale Neural Domain Classification with Personalized Attention. In *ACL*, pages 2214–2224.
- Diederik P. Kingma and Jimmy Lei Ba. 2015. ADAM: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*.
- Günter Klambauer, Thomas Unterthiner, Andreas Mayr, and Sepp Hochreiter. 2017. Self-normalizing neural networks. In *Advances in Neural Information Processing Systems 30 (NIPS)*, pages 972–981.
- Anjishnu Kumar, Arpit Gupta, Julian Chan, Sam Tucker, Bjorn Hoffmeister, Markus Dreyer, Stanislav Peshterliev, Ankur Gandhe, Denis Filiminov, Ariya Rastrow, Christian Monson, and Agnika Kumar. 2017. Just ASK: Building an Architecture for Extensible Self-Service Spoken Language Understanding. In *NIPS Workshop on Conversational AI*.
- Samuli Laine and Timo Aila. 2017. Temporal ensembling for semi-supervised learning. In *International Conference on Learning Representations (ICLR)*.
- Lemao Liu, Masao Utiyama, Andrew Finch, and Ei-ichiro Sumita. 2016. Neural machine translation with supervised attention. In *COLING*, pages 3093–3102.
- Minh-Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Effective Approaches to Attention-based Neural Machine Translation. In *EMNLP*, pages 1412–1421.
- André F. T. Martins and Ramón Fernandez Astudillo. 2016. From Softmax to Sparsemax: A Sparse Model of Attention and Multi-Label Classification. In *International Conference on Machine Learning (ICML)*, pages 1614–1623.
- Haitao Mi, Zhiguo Wang, and Abe Ittycheriah. 2016. Supervised Attentions for Neural Machine Translation. In *EMNLP*, pages 2283–2288.
- Graham Neubig, Chris Dyer, Yoav Goldberg, Austin Matthews, Waleed Ammar, Antonios Anastasopoulos, Miguel Ballesteros, David Chiang, Daniel Clothiaux, Trevor Cohn, et al. 2017. DyNet: The Dynamic Neural Network Toolkit. *arXiv preprint arXiv:1701.03980*.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global Vectors for Word Representation. In *EMNLP*, pages 1532–1543.
- Jean-Philippe Robichaud, Paul A. Crook, Puyang Xu, Omar Zia Khan, and Ruhi Sarikaya. 2014. Hypotheses ranking for robust domain classification and tracking in dialogue systems. In *Interspeech*, pages 145–149.
- Ruhi Sarikaya. 2017. The technology behind personal digital assistants: An overview of the system architecture and key components. *IEEE Signal Processing Magazine*, 34(1):67–81.
- Ruhi Sarikaya, Paul A Crook, Alex Marin, Minwoo Jeong, Jean-Philippe Robichaud, Asli Celikyilmaz, Young-Bum Kim, Alexandre Rochette, Omar Zia Khan, and Xiaohu Liu. 2016. An overview of end-to-end language understanding and dialog management for personal digital assistants. In *Spoken Language Technology Workshop (SLT)*, page 391–397.
- Gokhan Tur and Renato de Mori. 2011. *Spoken Language Understanding: Systems for Extracting Semantic Information from Speech*. New York, NY: John Wiley and Sons.