

FAST INTENT CLASSIFICATION FOR SPOKEN LANGUAGE UNDERSTANDING SYSTEMS

Akshit Tyagi^o, Varun Sharma^o, Rahul Gupta⁺, Lynn Samson^o, Nan Zhuang^o, Zihang Wang^o, Bill Campbell⁺

^oUniversity of Massachusetts, Amherst, USA

⁺Amazon.com, Cambridge, USA

ABSTRACT

Spoken Language Understanding (SLU) systems consist of several machine learning components operating together (e.g. intent classification, named entity recognition and resolution). Deep learning models have obtained state of the art results on several of these tasks, largely attributed to their better modeling capacity. However, an increase in modeling capacity comes with added costs of higher latency and energy usage, particularly when operating on low complexity devices. To address the latency and computational complexity issues, we explore a BranchyNet scheme on an intent classification scheme within SLU systems. The BranchyNet scheme when applied to a high complexity model, adds exit points at various stages in the model allowing early decision making for a set of queries to the SLU model. We conduct experiments on the Facebook Semantic Parsing dataset with two candidate model architectures for intent classification. Our experiments show that the BranchyNet scheme provides gains in terms of computational complexity without compromising model accuracy. We also conduct analytical studies regarding the improvements in the computational cost, distribution of utterances that egress from various exit points and the impact of adding more complexity to models with the BranchyNet scheme.

Index Terms— Spoken Language Understanding, BranchyNet, Intent Classification

1. INTRODUCTION

Spoken Language Understanding systems are core components of voice agents such as Apple Siri, Amazon Alexa and Google Home and can be designed in one of the several ways including as an end to end modeling scheme [1], or a collection of task specific classifiers [2, 3]. For a complex SLU system, the machine learning architecture can be computationally expensive, posing a challenge for applications such as on device SLU. In this work, we explore a scheme that allows us to retain the complexity of the SLU system, while allowing an early decision making when possible. The intuition for such a modeling choice stems for the fact that different queries made to the SLU system may warrant different degrees of processing. For instance, a request such as “Stop” is arguably simpler to process for an SLU system in comparison to a more complex request such as “find me the closet open restaurant at 8 PM and reserve me a spot”. To capture this intuition, we use the BranchyNet modeling scheme for an intent classification model as is used in the SLU system described in [2]. The BranchyNet scheme allows for decision making at various *depths* in a deep learning model, thereby naturally fitting our use case. The overarching goal of this work is to assess the efficacy of the BranchyNet methodology as a tool to reduce computational complexity and latency, while maintaining model accuracy.

The SLU system used in [2] uses as many domain classifiers, intent classifiers, named entity recognizers and reranker as the number

of domains supported by an SLU system. The domain classification, intent classification and named entity recognizer models identify the domain, user intent and name entities within a request and collectively can take significant compute resources. Several previous works have suggested methods to reduce computational complexity of the models. Examples include regularization methods [4], model distillation [5] and compression [6, 7]. All the methods mentioned above attempt to modify the modeling architecture to reduce computational complexity. For instance, regularization methods (e.g. L1, dropout) help reduce the number of model parameters and prevent over-fitting. However, a further attempt to reduce modeling complexity typically results in an accuracy loss. Similarly, distillation and compression can prevent over-fitting, but their use for further reducing computational complexity typically leads to an accuracy loss. We consider the BranchyNet as an alternative in this work, as it does not significantly alter the modeling architecture and allows for an adaptive use of fewer parameters depending upon the input query. The BranchyNet scheme also allows us to fragment the model into multiple sections, allowing for models segments to reside at different locations (e.g. a few early exit points can be stored on device for on-the-edge computing and the rest on the cloud for more complex computation).

We use the Facebook Semantic Parsing Systems (FSPS) dataset [8] to evaluate the efficacy of BranchyNet modeling technique. Using a DNN and stacked-LSTM models as candidate architectures for intent classification, we observe that the introduction of BranchyNet does not lead to any degradation in accuracy. For the more complex stacked-LSTM model we observe a relative improvement of 10.4% in the computational complexity (measured in operations per second during inference). We conduct further analysis on utterances egressing at various exit points in a model and assess the impact of increased model complexity on the egress distribution. We provide a more detailed description of the BranchyNet scheme in the next section along with other related work on reducing SLU modeling complexity.

2. RELATED WORK

2.1. Fast/Lower complexity SLU

Previous work that addresses design of fast and lower complexity SLU include use of quantization and hashing [9] and *private by design* on the edge modeling [10]. Other compression techniques such as knowledge distillation [11], matrix factorization [12], random projection techniques [13] and model pruning [14]. These techniques require an initial model choice and often a trade-off is made between accuracy and efficiency. On the other hand, there exist a set of algorithms that explicitly focus on reducing run time inference complexity. Examples include code optimization [15] and fast convolution for CNNs [16, 17]. We experiment with the BranchyNet

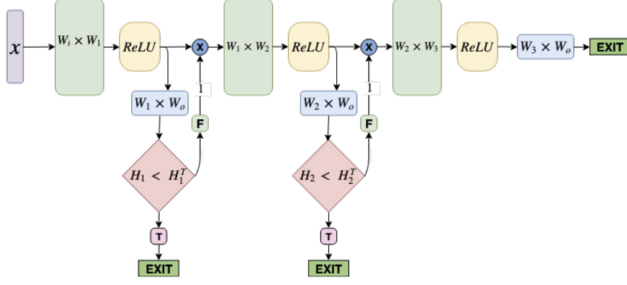


Fig. 1. NN Model with early exiting strategy. Exit can be done at each hidden layer based on model confidence.

network that does not alter the model architecture or inference algorithm. This allows us to retain the modeling architecture with the best accuracy, while allowing adaptive early inferences. We also note that the BranchyNet scheme can be used in combination with any of the aforementioned methods.

2.2. BranchyNet scheme

Given a model of choice with multiple potential exit points (e.g. each hidden layer in a DNN), the BranchyNet scheme aims at training a model that makes a decision as soon as it is confident in its prediction. Given N potential exit points with corresponding feature representations (e.g. outputs obtained at the hidden layer in a DNN), we optimize the following loss function L .

$$L = \sum_{n=1}^N \alpha_n L_n \quad (1)$$

Where L_n is the cross entropy loss, as determined by the output of the n^{th} exit point. We chose α_n to be a linearly decreasing function of n bounded by values r_l, r_u , as shown in the equation 2. Such a choice of α_n encourages the learning discriminative representations in earlier layers, thereby encouraging early exit. Note that during training, all exit points impact the model parameters and gradients computed using L_n influence the parameters contributing towards feature representation at the exit point n . Apart from learning the exit points for early exit, BranchyNet training also tends to have a regularizing effect on the model parameters as well as prevents vanishing gradients during training.

$$\alpha_n = r_l + \frac{r_u - r_l}{n}, \quad n = 1, \dots, N \quad (2)$$

Once trained, an entropy threshold H_n^T is defined for each exit point n . During inference, the entropy of the output class probabilities is computed at the exit point n , starting a scan from the exit point 1 through N . If the entropy at point n is less than the threshold H_n^T , a decision is made at that exit point. No further model computations are made for exit points beyond the chosen exit point n . For further details on the BranchyNet model, we refer the reader to [18].

2.3. Models used in our experimentation

We use two sets of models for the intent classification task in our experiments, as discussed below.

Deep Neural Network We use a DNN with ReLu activations with an exits at each layer. The hidden layer outputs from each are fed to an output layer as shown in Figure 1. The input features to the

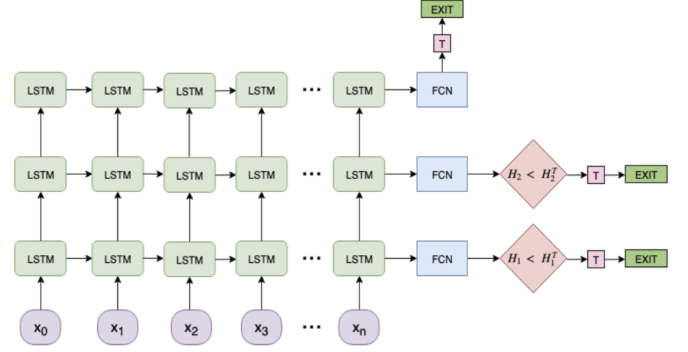


Fig. 2. Stacked LSTM model with early exiting strategy. The model can exit at each LSTM layer. FCN implies a fully connected layer.

NN are sentence embeddings computed as mean of the word embeddings constituting the sentence. We pre-train the word embedding on a larger Wikipedia corpus [19] and they are fine tuned during training for the task of our interest.

Stacked LSTM We also apply the BranchyNet technique to a Stacked LSTM network as shown in Figure 2. A uni-directional LSTM is applied to the word embeddings and sentence representation is obtained from the hidden layer at the last time step. The sentence representation at each LSTM layer is fed to an output layer for decision making.

We train both the above models using the training methodology defined in Section 2.2. We tune the number of nodes in the hidden layers in each DNN and Stacked LSTM layer for the best performance on the development set. The thresholds H_n^T for an exit point n is computed as the average value of the entropies obtained on the training data at each of the exit nodes. At inference time, if any node other than the last exit node obtains an entropy lesser than the average entropy, we make a decision at that exit point. Else, we make an egress at the last exit point. In the next section, we discuss the datasets used in our experimentation.

3. DATASET

Facebook Semantic Parsing System (FSPS) dataset: The FSPS dataset [8] consists of $\sim 44.7k$ utterances randomly split into $\sim 31.2k$ training, $\sim 4.4k$ validation and $\sim 9k$ test utterances. The utterances are annotated with a hierarchical representation; however for each given utterance we focus on predicting the overall utterance intent at the top node of the hierarchical annotation. The dataset consists of 25 intents including get-distance and get-directions intent. We tune the hyper-parameters of the model (e.g. number of hidden layers, number of stacked LSTM layers etc.) on the development set. In particular, we obtain best performance for three layered DNN and stacked-LSTM, leading to three exit points for the respective models.

4. RESULTS

We train the DNN and stacked LSTM models on the FSPS datasets with and without the BranchyNet mechanism built into it. The models without the BranchyNet mechanism serves as our baseline and we compare the change in model performance after adding the mechanism. Table 4 reports the macro F1 score and accuracy on the FSPS dataset.

Model	F1(Macro)	Acc.(%)
DNN	0.48	88.5
DNN + BranchyNet	0.55	89.6
Stacked LSTM	0.65	92.8
Stacked LSTM + BranchyNet	0.66	93.2

Table 1. Performance of DNN models on the FSPS dataset with and without the BranchyNet mechanism

Exit point	# Params($\times K$)	FLOPS($\times K$)
3-Layer Neural Network		
Without BranchyNet	36.4	36.2
Exit at first layer	32.6	32.5
Exit at second layer	38.9	38.7
Exit at final layer	40.8	40.6
Stacked LSTM		
Without BranchyNet	22.0	69.2
Exit at first layer	7.6	23.1
Exit at second layer	14.8	46.1
Exit at final layer	22.1	69.2

Table 2. FLOPS comparison for performing inference on one data-point using various models in our experiments

As expected, we observe that the Stacked LSTM outperforms the DNN counterparts. We also observe a boost in the respective model performances after addition of the BranchyNet mechanism. While this accuracy can be attributed to the regularization effect (as specified as an added benefit of the BranchyNet scheme in [18]), we also believe that with BranchyNet, each layer can learn to obtain tailored representations for classification at each layer. In order to quantify the increase in computational complexity, we present analysis in the next section along with an analysis on the utterances egressing from each exit node.

5. ANALYSIS

We conduct three sets of analysis on the trained BranchyNet network: (i) Evaluating the model size and computational complexity of the trained model, (ii) Understanding the lexical distribution of the utterances exit from each node, (iii) Exit proportions with increased complexity of the models.

5.1. Model size and computational complexity analysis

In this section, we compare the cumulative number of parameters before each exit node of a model as well as the number of floating point operations (FLOPS) to arrive at a decision. Table 2 presents these statistics for the baseline models as well as their versions with the BranchyNet mechanism. For the DNN model, we observe that the number of parameters of the overall model (equivalent to the number of parameters while exiting at the final layer) is greater than the baseline models. This happens due to the fact that each exit point contains parameters for entropy comparison as well as an output layer. The number of parameters for the Stacked LSTM is lower than three layered DNN due to repeated use of the same set of parameters at each time step. We also observe that the improvement in computational complexity is higher for stacked-LSTM. For instance, if inference is made at the first layer, it is done with a third of the computational cost when compared to the baseline model without BranchyNet. Additionally, exit at the final Stacked LSTM layer

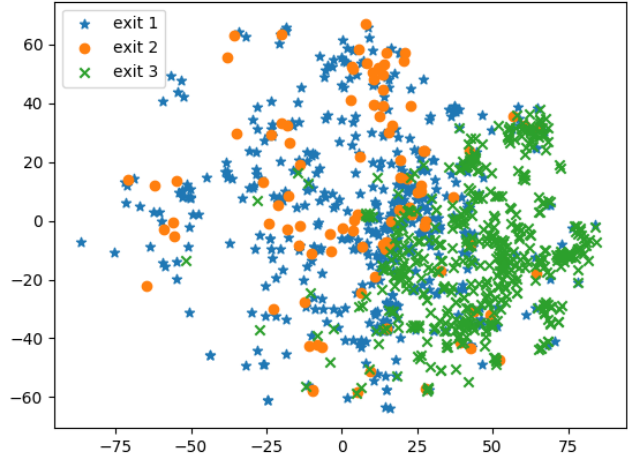


Fig. 3. Distribution of utterances as exit from each exit point as obtained after t-SNE projection. Axes are the two dimensions after the t-SNE projection

Model	Exit Point activation distribution	Expected complexity
DNN	1: 27.80	38.27
	2: 01.92	
	3: 70.20	
Stacked LSTM	1: 15.32	61.99
	2: 00.71	
	3: 83.97	

Table 3. Distribution of utterances egressing from each exit point. We also provide the expected complexity based off this observed distribution on the FSPS test set.

is almost as expensive as the baseline model, suggesting that the BranchyNet scheme does not add significant computational complexity even in the worst case. Overall, we observe that there is an opportunity for reduced computational complexity without taking a hit in the performance, particularly for more complex models. We look at the distribution of the exit point activation in the next section for the FSPS dataset and report expected reduction in computational complexity depending on this distribution.

5.2. Lexical distribution of utterance exit from each node

Given that we observe decrease in computational cost if exit early, we list at the distribution of exit point activation for the FSPS dataset in Table 3. We observe for the test portion of the FSPS dataset, majority of the test utterances are decoded either in the first layer or the last layer. Given the exit distribution of the test data, the table also lists the expected number of FLOPs (computed as the sum of FLOPs at each layer, weighted by the probability of exit at that layer). While we observe an increase in computational complexity of the BranchyNet based DNN model, for the more complex Stacked LSTM model, we observe a relative decrease of 10.4%. This observation advocates for the use of the BranchyNet scheme, particularly when using more complex models.

In figure 3, we also plot the distribution of the sentence representations as exit from each exit point. We extract the sentence embedding of the utterances from a Doc2Vec model pre-trained on a large

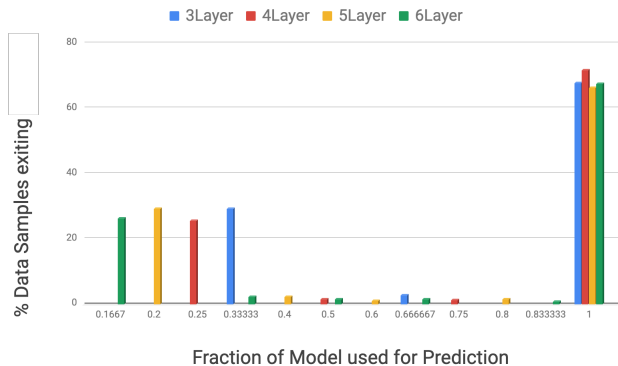


Fig. 4. Fraction of test data egressing from each exit point in a 3/4/5/6 layered DNN.

Wikipedia corpus [19]. We train a t-Distributed Stochastic Neighborhood Embedding (t-SNE) model to transform the higher dimensional sentence representations into a 2-dimensional representation for the purposes of visualization. In the obtained 2-dimensional representation, we observe that the test utterance egressing from each exit point cluster differently in the utterance representation space. This indicates that there is an inherent difference in the utterances. Eyeballing the data egressing from each point reflects the differences in them. For instance, all the sentences containing the word ‘restaurant’ egress from the first exit point, while all utterances containing the word ‘holiday’ egress from the last exit point. We also observed that the length of the utterances egressing from the first exit is generally higher than those egressing from other nodes. Arguably, the length of the utterance is an indicator of utterance complexity and this observation is against our expectation that more complex utterance will egress from latter exit nodes. The t-SNE plots along with these observation suggests that each exit node specializes in a specific semantic space and topic as opposed to utterance complexity.

5.3. Impact on increasing model complexity on model outputs

In this section, we observe if the exit distribution of the network changes as we add more exit points to the models. We add more hidden layers and LSTM layers to the models of our interest and observe the exit proportions at each layer. Figure 4 and 5 presents the exit proportions for the DNN and Stacked-LSTM models.

The figures indicate that the egress proportion of utterances from the first and the last exit point does not change significantly. We hoped that a further break down of the semantic space would be possible with added exit points; providing us flexibility in terms of fragmenting the model (e.g. a variable proportion of the model could be fragmented between device and the cloud). However, the results do not indicate this and we observe that majority of the decisions are made at two egress points. We recommend finding this optimal configuration based on expected accuracy maximization as done on the development set in our experiments.

6. CONCLUSION

SLU models often are an ensemble of complex models operating in unison to return machine interpretable hypotheses to downstream components. Their complexity makes them hard to run on low complexity devices and compression methods often lead to a loss in ac-

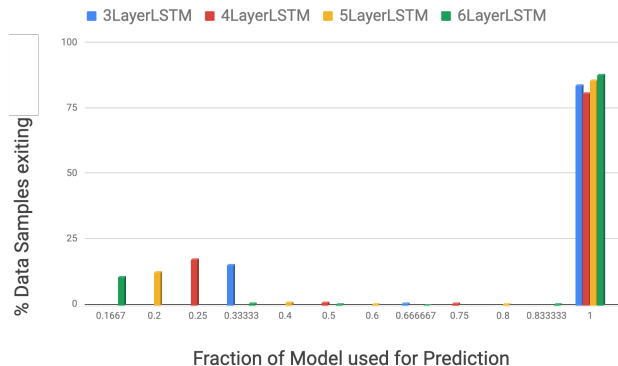


Fig. 5. Fraction of test data egressing from each exit point in a 3/4/5/6 layered Stacked-LSTM.

curacy. In this work, we experiment with a BranchyNet scheme that retains the architecture of an intent classification model, yet allowing for reduction in computational complexity. The scheme also allows for model fragmentation with a partial model storage on device. We also perform analyses on the model and observe distribution of utterance egressing from each exit point and the impact of increasing modeling capacity. Our results indicate the promise of using BranchyNet scheme for SLU systems as a way to reduce computational complexity, without the need to change modeling architecture or a drop in accuracy. Conducted analyses suggest an egress based on semantic clusters and no further splitting of clusters with addition of more exit points.

As the next steps, we aim at extending the BranchyNet scheme to the entire stack of SLU models [2]. The BranchyNet scheme can also be combined with other modeling schemes such as model distillation and compression. We also aim to test variants of BranchyNet scheme with exit points added at different points in the model (e.g. a partial set of neurons in a hidden layer) for a further reduction in the complexity. Exit criterion from each exit point is another parameter for experimentation and we will look for criterion beyond entropy. Finally, we observed that the egress is based on topic of the utterance. We aim to introduce other forcing functions that encourage egress based on utterance complexity or other factors of interest.

7. REFERENCES

- [1] Dmitriy Serdyuk, Yongqiang Wang, Christian Fuegen, Anuj Kumar, Baiyang Liu, and Yoshua Bengio, “Towards end-to-end spoken language understanding,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5754–5758.
- [2] Chengwei Su, Rahul Gupta, Shankar Ananthkrishnan, and Spyros Matsoukas, “A re-ranker scheme for integrating large scale nlu models,” in *2018 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2018, pp. 670–676.
- [3] Rahul Gupta, Aman Alok, and Shankar Ananthkrishnan, “One-vs-all models for asynchronous training: An empirical analysis,” *arXiv preprint arXiv:1906.08858*, 2019.
- [4] Christopher M Bishop, “Regularization and complexity control in feed-forward networks,” 1995.
- [5] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean, “Distill-

- ing the knowledge in a neural network,” *arXiv preprint arXiv:1503.02531*, 2015.
- [6] Yu Cheng, Duo Wang, Pan Zhou, and Tao Zhang, “A survey of model compression and acceleration for deep neural networks,” *arXiv preprint arXiv:1710.09282*, 2017.
 - [7] Edward WD Whittaker and Bhiksha Raj, “Quantization-based language model compression,” in *Seventh European Conference on Speech Communication and Technology*, 2001.
 - [8] Sonal Gupta, Rushin Shah, Mrinal Mohit, Anuj Kumar, and Mike Lewis, “Semantic parsing for task oriented dialog using hierarchical representations,” *arXiv preprint arXiv:1810.07942*, 2018.
 - [9] Grant P Strimel, Kanthashree Mysore Sathyendra, and Stanislav Peshterliev, “Statistical model compression for small-footprint natural language understanding,” *arXiv preprint arXiv:1807.07520*, 2018.
 - [10] Alaa Saade, Alice Coucke, Alexandre Caulier, Joseph Dureau, Adrien Ball, Théodore Bluche, David Leroy, Clément Doumouro, Thibault Gisselbrecht, Francesco Caltagirone, et al., “Spoken language understanding on the edge,” *arXiv preprint arXiv:1810.12735*, 2018.
 - [11] Liang Lu, Michelle Guo, and Steve Renals, “Knowledge distillation for small-footprint highway networks,” in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 4820–4824.
 - [12] Yunchao Gong, Liu Liu, Ming Yang, and Lubomir Bourdev, “Compressing deep convolutional networks using vector quantization,” *arXiv preprint arXiv:1412.6115*, 2014.
 - [13] Ella Bingham and Heikki Mannila, “Random projection in dimensionality reduction: applications to image and text data,” in *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2001, pp. 245–250.
 - [14] Song Han, Huizi Mao, and William J Dally, “Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding,” *arXiv preprint arXiv:1510.00149*, 2015.
 - [15] Vincent Vanhoucke, Andrew Senior, and Mark Z Mao, “Improving the speed of neural networks on cpus,” .
 - [16] Michael Mathieu, Mikael Henaff, and Yann LeCun, “Fast training of convolutional networks through ffts,” *arXiv preprint arXiv:1312.5851*, 2013.
 - [17] Andrew Lavin and Scott Gray, “Fast algorithms for convolutional neural networks,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 4013–4021.
 - [18] B. McDanel S. Teerapittayanon and H. Kung, *Branchynet: Fast inference via early exiting from deep neural networks*, 2016.
 - [19] Andrew M Dai, Christopher Olah, and Quoc V Le, “Document embedding with paragraph vectors,” *arXiv preprint arXiv:1507.07998*, 2015.