

DYNAMIC SPEECH ENDPOINT DETECTION WITH REGRESSION TARGETS

Dawei Liang^{*}, Hang Su[†], Tarun Singh[†], Jay Mahadeokar[†], Shanil Puri[†], Jiedan Zhu[†],
Edison Thomaz^{*}, Mike Seltzer[†]

^{*} University of Texas at Austin, [†] Meta AI.

ABSTRACT

Interactive voice assistants have been widely used as input interfaces in various scenarios, e.g. on smart homes devices, wearables and on AR devices. Detecting the end of a speech query, i.e. speech end-pointing, is an important task for voice assistants to interact with users. Traditionally, speech end-pointing is based on pure classification methods along with arbitrary binary targets. In this paper, we propose a novel regression-based speech end-pointing model, which enables an end-pointer to adjust its detection behavior based on context of user queries. Specifically, we present a pause modeling method and show its effectiveness for dynamic end-pointing. Based on our experiments with vendor-collected smartphone and wearables speech queries, our strategy shows a better trade-off between endpointing latency and accuracy, compared to the traditional classification-based method. We further discuss the benefits of this model and generalization of the framework in the paper.

Index Terms— endpointing, end-of-query, interactive voice assistant.

1. INTRODUCTION

With rapid development of speech technologies in recent years, interactive voice assistants have been widely adopted as a mainstream for intelligent user interface [1, 2, 3, 4]. By taking users’ speech queries, these systems are able to perform a variety of tasks from basic question answering, music playing, calling and messaging, to device control. As an initial step, a voice assistant needs to determine the time point when a user finishes the query so that it knows when to close the microphone and continue downstream processing (e.g. language understanding and taking action). This process is often referred to as speech *endpoint detection* or *end-pointing*. In practice, the challenge for speech end-pointing lies in the conflict goals of fast response and endpoint accuracy (avoiding early cuts of user speech). Specifically, a rapid close of the microphone and response to user queries brings a better user experience if endpointed correctly, but this inevitably increases the risk of early-cut of user queries. The performance of an end-pointing system is thus evaluated on how this conflict is resolved in practical cases.

Canonically, voice activity detection (VAD) has been widely used for speech end-pointing [5, 6]. In the VAD setting, a model is developed to distinguish speech segments from non-speech segments, which usually includes silence, music, and background noises [7, 8, 9]. The end of the query can then be determined if a fixed duration of silence is observed by the VAD system. However, this approach is not reliable enough as pointed out by later work [10, 11]. The fact that a VAD system is typically not trained to distinguish pauses within and at the end of queries prevents the system from capturing enough acoustic cues related to end-of-query detection, such as speaking rhythm or filler sounds [11].

Recent researches on speech end-pointing mostly focus on classification methods, where a dedicated end-pointer is developed to classify audio frames into end-of-query frames and other frames [11, 12]. The success of the Long Short-Term Memory (LSTM) [13] architecture contributes to this method. Some other efforts include additional text decoding features [14] or user personalized i-vectors [15] to better fit the end-pointer for specific acoustic environments. In an end-to-end automated speech recognition (ASR) system, the end-pointer may also be jointly optimized with the recognition model [16]. Despite the promising results, all of the above works focus on binary detection of end-of-query with hard labels (e.g. 0 and 1). In real scenarios, however, an endpointer shall adjust its endpointing aggressiveness based on semantic, prosodic or other speaking patterns in the query. The traditional binary targets for classification can be less flexible in this respect.

In this paper we study a novel speech end-pointing strategy based on regression method. In this setup, an end-pointer is optimized to fit soft-coded targets during training, and the training targets are set by considering expected pause given semantic context of the queries. By testing on both 14.4M smartphone speech queries and 467K wearables user queries, we show that our proposed method effectively reduces the response delay of end-pointing while maintaining a comparable accuracy performance as the conventional classification-based method.




Fig. 1. Sample targets for end-pointer training based on hard-coded binary values (target in black) and soft-coded float values (target in red). EP: endpoint.

2. REGRESSION-BASED ENDPOINT MODELING

2.1. Principle

In a traditional classification setup, a speech end-pointing model is trained against targets of hard-coded values (0 and 1). At prediction stage, a decision threshold is picked and compared against model’s output end-pointing confidence. In a regression setting, the prediction stage remains the same, but the training targets are changed to be continuous float values from 0 to 1, representing the confidence of a frame being part of end-of-query. Specifically, we apply a transition curve to the area between user speech end and the expected endpoint position (i.e. position where we are confident to endpoint) (Figure 1). The motivation of our regression-based modeling includes: 1. the model should have an increased confidence on endpointing as more silence is observed; 2. the aggressiveness of the endpointing behavior shall be adjustable, which corresponds to the slope of the transition region. As an extreme case, a regression model with 0 transition region would have the same targets as a classification model. The top plot of Figure 1 is an example where the training targets of a regression model and classification model are the same. At the bottom, the regression target enforces the model to be less aggressive, with a smoother slope in the transition region.

Figure 2 presents the overall pipeline of end-pointing with regression. To enable a dynamic output behavior of the end-pointer, the transition slope of the training targets can be quantified based on unique patterns (e.g., expected pause) of individual input queries. In this paper, we propose a method to compute and leverage the pause statistics of queries for this purpose (Section 3). During inference, the regression model shares the same setup as the classification model, i.e., taking speech features as inputs, and applying a decision threshold to the model outputs for endpoint decision.

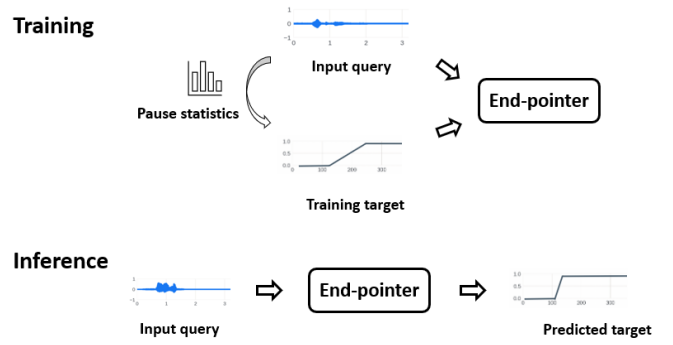


Fig. 2. The overall pipeline of our proposed method.

2.2. Loss Function

In our study, we use the mean-square-error (MSE) as loss function for end-pointer training. The loss function is defined as below:

$$l = \frac{\sum_{i=1}^n (\hat{y}_i - y_i)^2}{n} \quad (1)$$

where i is the i th frame of an utterance in the dataset, n is the total number of frames in the dataset, \hat{y}_i and y_i are the output of the end-pointer and training target, respectively. In our experiments, we found that regression with the binary targets yields the same performance as the traditional classification setup. Furthermore, the MSE loss tends to be more stable for model convergence than the L1 loss.

3. SPEAKER PAUSE MODELING

To enable dynamic adjustments of endpointing behavior, we explore the expected pause following a speech segment. Our intuition is that ideal end-pointing aggressiveness should be controlled by the expected pause following a query, and the expected pause is related to semantics. For example, the chance of a speaker continuing speaking following an utterance “*what’s the weather*” is bigger than that following the utterance “*what’s the weather in Seattle next week*”. By calculating the expected pause following a speech segment, we can adjust the slope of the transition region to be smoother for the first utterance, so that the end-pointer can learn to wait a bit longer for the first utterance than the second. We hope to estimate the expected duration of each training query and use it to adjust the targets for these queries.

We model the pause for a given text context as a Gaussian variable:

$$T_T|C \sim \mathcal{N}(\mu, \sigma^2) \quad (2)$$

Here T_T is the pause duration for a text query, C is the text context of the query, μ and σ are parameters for the Gaussian variable. The parameters of this model can be estimated by

aggregating pause statistics of utterances with a common context (prefix). For example, the query "what is the weather" is considered as a prefix for the queries "what is the weather", "what is the weather in Seattle", and "what is the weather today", and all these queries contribute to the statistics of pause duration given query "what is the weather". Note that utterances with transcription "what is the weather" have zero pause, indicating it is a complete utterance itself. In this work, the common prefix definition requires a strict match of transcriptions. Hence, the utterances "what is the weather" and "how is the weather" are considered to be two independent contexts, though they are semantically similar. Ideally, these utterances with similar semantics should be grouped together. Also, we used the 95th percentile statistics in our experiment for the expected pause duration estimate rather than mean μ to reduce the risk of early-cutting.

We further model the pause for a given speech query as follows:

$$T_S = T_T \cdot R \quad (3)$$

Here, T_S is the pause duration for a speech query, T_T is the duration for a text query, and R is the speaking rate. The intuition here is that different queries were spoken at different speaking rates, and the slower a speaker speaks, the longer the expected pause. The speaking rate factor can be estimated using the ratio between the duration of a speech query and the average duration of all those queries with the same prefix (duration of the prefix part only).

4. EXPERIMENTAL SETUP

4.1. Data

We used two datasets to evaluate our method - both are collected by 3rd party data vendors. The first one is collected using smartphones whereas the second one is collected using wearable devices (smart glasses). We refer to the two datasets as smartphone data and wearables data in this paper, respectively. The smartphone dataset contains clean user speech that are spoken in fluent manner, with little background noise. In total, 14.4M queries were collected, categorized into 55 domain contexts based on the text transcription. The most common contexts are *music* (1.4M), *weather* (1.3M), and *device handling* (1.2M). The average word count per query is 6. The wearables dataset is collected when users interact with a voice assistant on smart glasses, with real life noises in the background. A considerable amount of speech pauses and dis-fluency are identified in this data. It contains 447K queries, with an average word count of 11 per query. Both datasets contain speech and ground truth transcriptions. Ground truth speech endpoints are calculated by aligning speech to transcription and measure the last word endtime. From the alignments, we can also collect pause duration (if any) following each word in the queries.

For both datasets, we randomly split the queries with a ratio of 95:5 for model development and testing. A development set was further split out in random during training for learning rate adjustment. The training set and testing set share the same set of speakers.

4.2. Model Configuration

Our end-pointer is an LSTM neural network inspired by prior work [11, 14]. Specifically, the model consists of three uni-directional LSTM layers with a hidden unit size of 128 per layer. For the baseline classification model, the output is a fully-connected layer of two neurons, transformed by Log-Softmax activation. For the regression model, the output contains only a single neuron with sigmoid activation. During training, the baseline model leveraged the regular cross-entropy loss, while the regression model uses MSE loss. The learning rate was $2 \times 10^{-4} / 2 \times 10^{-3}$ for each model, and reduced by a factor of 0.5 if no improvement was observed on the validation set. The Adam optimizer [17] was used with a Beta of (0.9, 0.999). The mini batch size was 128. Besides, a sliding window of 10 ms was applied to the outputs of the baseline model at prediction stage for score smoothing. Both models were trained for 15 epochs when learning rate drops to a small enough value. All development was done using PyTorch toolkit[18].

We use 40-dim filter-bank features as inputs to the network. We extracted the features based on a 25 ms window and a stride of 10 ms. The sampling rate of the audio was 16 kHz.

5. RESULTS

5.1. Overall Results

In total, we obtained 106K and 18K unique prefixes respectively for the smartphone and the wearable training sets. Following the prior work [11, 14], we applied the early-cut rate as the model accuracy metric which reports the proportion of early-cut queries out of the entire test population. We do not use WER as a good endpointing metric, because we think it is not linearly correlated to endpointing accuracy, and thus not linearly correlated to user experience on endpointing. Specifically, an early cut in the beginning of an utterance and at the end of an utterance are both early cuts and both cause users to repeat their queries, but their WER differs greatly. For the latency measure, we used the 50th, 75th, 90th, and 99th percentile statistics of latency values out of the predicted results, noted as P50, P75, P90, and P99, respectively. For both metrics, a smaller value is more desirable. Tables 1 and 2 summarize the overall model performance of our study, based on sample decision thresholds from around 0.5. Given the closest early-cut rate for each pair of comparisons, we can see that the prediction latency measured by P50, P75, and P90 tends to be improved by the regression model on both

Threshold	Early-cut rate (%)	P50 (ms)	P75 (ms)	P90 (ms)	P99 (ms)
0.50 / 0.63	3.39 / 3.38	160 / 120	180 / 150	230 / 210	480 / 530
0.60 / 0.74	2.45 / 2.38	170 / 120	200 / 160	250 / 240	530 / 600
0.70 / 0.82	1.74 / 1.67	180 / 130	210 / 170	280 / 260	590 / 660

Table 1. End-pointing performance with classification / regression-based models for the clean speech dataset.

Threshold	Early-cut rate (%)	P50 (ms)	P75 (ms)	P90 (ms)	P99 (ms)
0.56 / 0.50	14.83 / 14.81	420 / 350	590 / 450	860 / 810	1990 / 2500
0.60 / 0.59	12.82 / 12.75	470 / 430	670 / 510	970 / 870	2050 / 2500
0.70 / 0.67	10.53 / 10.37	580 / 500	780 / 650	1100 / 1130	2120 / 2500

Table 2. End-pointing performance with classification / regression-based models for the noisy speech dataset.

datasets. This demonstrates a better balance of end-pointing behavior achieved by the regression model. For P99, the classification model was consistently better. This is in fact reasonable, because P99 shows the very tail latency of prediction, and the regression model was enforced to have a very smooth performance for queries of potentially a high expectation of succeeding pause. For the typical user experience of latency, regression-based end-pointing is still more desirable, as demonstrated by P50 and P75.

5.2. Discussions

To further study the potential of generalization for regression-based end-pointing, we visualized the model outputs based on a simpler setting. In this setting, the training target slope of the end-pointer was adjusted by hand-crafted criteria based on the max pause duration within a query (0: 10 ms slope duration; > 90 ms: 1,000 ms slope duration). Figure 3 shows predictions of the regression model for two sample queries. It can be seen that the regression model can achieve the same output behavior as classification-based end-pointing for a query with little within-sentence pauses (top plot). However, the output of the regression model is smoother for a query with significant within-sentence pauses (bottom plot). The plots help to demonstrate that regression-based end-pointing can be applied to different target adjustment strategies.

In addition, we noticed that regression-based end-pointing can help to maintain the model stability for end-pointers of lower complexity. In an extra study, we halved the hidden units of the LSTM layers for our neural end-pointer. By testing with the noisy data, the early-cut rate of the regression model increased to 18.77% at a threshold of 0.5. The P50, P75, P90 and P99 values also increased to 380 ms, 550 ms, 970 ms, and 2,370 ms, respectively. However, at the closest early-cut rate (27.95%, threshold = 0.81), the P50, P75, P90 and P99 results of the classification model significantly increased to 1,770 ms, 2,260 ms, 2,480 ms, and 2,690 ms, respectively. A possible explanation is that the abstraction of the unique prefixes from the original queries helps to general-




Fig. 3. Comparison of end-pointer outputs based on classification (black) and regression (red). The training targets of the regression model were adjusted by max pause within queries.

ize the end-pointer on the patterns of the queries. This benefit is meaningful for end-pointing on the edge, especially for devices with limited computational capabilities.

6. CONCLUSION AND FUTURE WORK

In this paper, we proposed a novel regression-based speech end-pointing model, which points to a new direction for solving the early cut issues in endpointing without sacrificing latency. We then describe how we utilize this model for dynamic end-pointing by adjusting training targets based on expected pause of speech queries. Based on experiments with speech data collected by smartphones and wearables, the proposed regression model shows a better performance on latency-accuracy trade-off, with an improved P50, P75 and P90, specifically, while maintaining a comparable early-cut performance to the traditional method. This regression-based endpointing approach is quite flexible and can be used to incorporate speaking behavior modeling in the future. A better pause model can be explored by grouping utterance into semantics rather than pure prefixes as well.

7. REFERENCES

- [1] Martin Porcheron, Joel E Fischer, Stuart Reeves, and Sarah Sharples, “Voice interfaces in everyday life,” in *proceedings of the 2018 CHI conference on human factors in computing systems*, 2018, pp. 1–12.
- [2] Alisha Pradhan, Kanika Mehta, and Leah Findlater, ““accessibility came by accident” use of voice-controlled intelligent personal assistants by people with disabilities,” in *Proceedings of the 2018 CHI Conference on human factors in computing systems*, 2018, pp. 1–13.
- [3] Radhika Garg and Subhasree Sengupta, “He is just like me: a study of the long-term use of smart speakers by parents and children,” *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 4, no. 1, pp. 1–24, 2020.
- [4] Frank Bentley, Chris Luvogt, Max Silverman, Rushani Wirasinghe, Brooke White, and Danielle Lottridge, “Understanding the long-term use of smart speaker assistants,” *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 2, no. 3, pp. 1–24, 2018.
- [5] Won-Ho Shin, Byoung-Soo Lee, Yun-Keun Lee, and Jong-Seok Lee, “Speech/non-speech classification using multiple features for robust endpoint detection,” in *2000 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (Cat. No. 00CH37100)*. IEEE, 2000, vol. 3, pp. 1399–1402.
- [6] Ramalingam Hariharan, Jula Hakkinen, and Kari Laurila, “Robust end-of-utterance detection for real-time speech recognition applications,” in *2001 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (Cat. No. 01CH37221)*. IEEE, 2001, vol. 1, pp. 249–252.
- [7] Rathinavelu Chengalvarayan, “Robust energy normalization using speech/nonspeech discriminator for german connected digit recognition,” in *Sixth European conference on speech communication and technology*, 1999.
- [8] Shuo-Yiin Chang, Bo Li, Gabor Simko, Tara N Sainath, Anshuman Tripathi, Aäron van den Oord, and Oriol Vinyals, “Temporal modeling using dilated convolution and gating for voice-activity-detection,” in *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2018, pp. 5549–5553.
- [9] Dawei Liang, Zifan Xu, Yinuo Chen, Rebecca Adaimi, David Harwath, and Edison Thomaz, “Automated detection of foreground speech with wearable sensing in everyday home environments: A transfer learning approach,” *arXiv preprint arXiv:2203.11294*, 2022.
- [10] Luciana Ferrer, Elizabeth Shriberg, and Andreas Stolcke, “A prosody-based approach to end-of-utterance detection that does not require speech recognition,” in *2003 IEEE International Conference on Acoustics, Speech, and Signal Processing, 2003. Proceedings.(ICASSP’03)*. IEEE, 2003, vol. 1, pp. I–I.
- [11] Matt Shannon, Gabor Simko, Shuo-Yiin Chang, and Carolina Parada, “Improved end-of-query detection for streaming speech recognition.,” in *Interspeech*, 2017, pp. 1909–1913.
- [12] Shuo-Yiin Chang, Bo Li, Tara N Sainath, Gabor Simko, and Carolina Parada, “Endpoint detection using grid long short-term memory networks for streaming speech recognition.,” in *Interspeech*, 2017, pp. 3812–3816.
- [13] Sepp Hochreiter and Jürgen Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [14] Roland Maas, Ariya Rastrow, Chengyuan Ma, Guitang Lan, Kyle Goehner, Gautam Tiwari, Shaun Joseph, and Björn Hoffmeister, “Combining acoustic embeddings and decoding features for end-of-utterance detection in real-time far-field speech recognition systems,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5544–5548.
- [15] Aditya Jayasimha and Periyasamy Paramasivam, “Personalizing speech start point and end point detection in asr systems from speaker embeddings,” in *2021 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2021, pp. 771–777.
- [16] Shuo-Yiin Chang, Rohit Prabhavalkar, Yanzhang He, Tara N Sainath, and Gabor Simko, “Joint endpointing and decoding with end-to-end models,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 5626–5630.
- [17] Diederik P Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [18] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al., “Pytorch: An imperative style, high-performance deep learning library,” *Advances in neural information processing systems*, vol. 32, 2019.