



Neural Multi-task Learning for Teacher Question Detection in Online Classrooms

Gale Yan Huang^{1,3}, Jiahao Chen¹, Haochen Liu², Weiping Fu¹,
Wenbiao Ding¹, Jiliang Tang², Songfan Yang¹, Guoliang Li³, and Zitao Liu¹(✉)

¹ TAL Education Group, Beijing, China
{galehuang, chenjiahao, fuweiping, dingwenbiao,
yangsongfan, liuzitao}@100tal.com

² Data Science and Engineering Lab, Michigan State University, East Lansing, USA
{liuhaoc1, tangjili}@msu.edu

³ Department of Computer Science, Tsinghua University, Beijing, China
liguoliang@tsinghua.edu.cn

Abstract. Asking questions is one of the most crucial pedagogical techniques used by teachers in class. It not only offers open-ended discussions between teachers and students to exchange ideas but also provokes deeper student thought and critical analysis. Providing teachers with such pedagogical feedback will remarkably help teachers improve their overall teaching quality over time in classrooms. Therefore, in this work, we build an end-to-end neural framework that automatically detects questions from teachers' audio recordings. Compared with traditional methods, our approach not only avoids cumbersome feature engineering, but also adapts to the task of multi-class question detection in real education scenarios. By incorporating multi-task learning techniques, we are able to strengthen the understanding of semantic relations among different types of questions. We conducted extensive experiments on the question detection tasks in a real-world online classroom dataset and the results demonstrate the superiority of our model in terms of various evaluation metrics.

Keywords: Question detection · Multi-task learning · Natural language understanding · Online classroom

1 Introduction

Teachers utilize various pedagogical techniques in their classrooms to inspire students' thought and inquiry at deeper levels of students' comprehension. These techniques may include lectures, asking questions, assigning small-group work, etc. [6, 9, 25]. A large body of research has demonstrated that asking certain types of questions can increase student engagement and it is an important factor of student achievement [1, 2, 16, 20, 27, 32]. Asking questions has become a central

H. Liu—Work was done when the author did internship in TAL Education Group.

© Springer Nature Switzerland AG 2020

I. I. Bittencourt et al. (Eds.): AIED 2020, LNAI 12163, pp. 269–281, 2020.

https://doi.org/10.1007/978-3-030-52237-7_22

component of teachers’ dialogic instructions and often serves as a catalyst for in-depth classroom discussions [21, 26, 28].

A large spectrum of approaches have been developed and successfully applied in generating classroom feedback to evaluate teachers’ performances and help them improve their pedagogical techniques [15, 19, 26, 28, 31]. For example, the Nystrand and Gamoran coding scheme provides a general template for recording teachers’ activities, which are used by trained human judges to manually assess teachers’ classroom practices [15, 26]. However, manually analyzing teacher questions is very subjective, expensive, time-consuming and not scalable. Thus, it is crucial to develop computational methods that can automatically detect teacher questions in live classrooms. By automatically analyzing when teachers ask questions and the corresponding question types, we are able to evaluate the question impact on teaching achievements and help teachers make adjustments to improve their pedagogical techniques. Previous endeavors have been conducted to tackle this problem using traditional machine learning (ML) algorithms [4–6, 12, 29]. However, the majority of these methods are not sufficient for teacher question detection due to the following challenges:

- *Question type variation.* Different from questions in daily chatting, routine conversation or other scenarios, teacher questions in classrooms are very diverse and open-ended. There are different types of classroom questions, such as knowledge-solicitation questions (e.g., “*What’s the definition of quadrangle?*”), open questions (e.g., “*Could you tell me your thought?*”), procedural questions (e.g., “*Can everyone hear me?*”), and discourse-management questions (e.g., “*What?*”, “*Excuse me?*”) [7, 28]. Traditional methods fail to perform a deep semantic understanding on natural languages, which is necessary for detecting questions of various types.
- *Subject and speaker variability.* Teaching materials and styles vary dramatically for different subjects and teachers, which leads to significantly distinguished classroom question sentences. Traditional methods show poor adaptability. When new subjects or teachers appear, most existing approaches have to be redesigned and retrained with the newly arrived data.
- *Tedious feature engineering.* Traditional ML-based methods detect questions based on complex acoustic and language features. It’s time-consuming to construct manually-engineered patterns.

In this work, we aim to investigate accurate teacher question detection in online classrooms. In particular, we study two variants of the teacher question detection problem. One is a two-way detection task that aims to distinguish questions from non-questions. The other is a multi-way detection task that aims to classify different types of questions. Please note that the formal definitions of the above two tasks are introduced in Sect. 3.2. We design a neural natural language understanding (NLU) model to automatically extract semantic features from teachers’ sentences for both the two-way task and the multi-way task. Our approach shows a powerful generalization capability for detecting questions of various types from different teaching scenarios. With the neural model as a core

component, we build an end-to-end framework that directly takes teacher audio tracks as input and outputs the detection results. Experiments conducted on a real-world online education dataset demonstrate the superiority of our proposed framework compared with competitive baseline models.

2 Related Work

2.1 Teacher Question Detection

Blanchard et al. explore classifying Q&A discourse segments based on audio inputs [4]. A simple amplitude envelope thresholding-based method is developed to detect teachers’ utterances. Then the authors extract 11 speech-silence features from detected utterances and train supervised classifiers to differentiate Q&A segments from other segments. Following this work, Blanchard et al. further introduce an automatic speech recognition (ASR) system to convert audio features into domain-general language features for teacher question detection [5,6]. They extract 37 NLP features from ASR transcriptions and train different classical ML models to distinguish questions from non-questions. Besides linguistic features, Donnelly et al. try both prosodic and linguistic features for supervised question classification and conclude that ML classifiers can achieve better performance with linguistic features [12].

The line of research presented above focuses on detecting questions from non-questions, which is a binary classification problem. Besides, we are interested in classifying questions into specific categories. Samei et al. build ML models to predict two properties “uptake” and “authenticity” of questions in live classrooms [29]. They extract 30 linguistic features related to part-of-speech and pre-defined keywords from each individual question. Samei et al. show that ML models are able to achieve comparable question detection performance as human experts.

Different from previous works of building question detection ML models based on manually selected linguistic and acoustic features, our approach eliminates the feature engineering efforts and directly learns effective representations from the ASR transcriptions. Furthermore, we introduce multi-task learning techniques into our model to classify different types of questions.

2.2 Multi-task Learning

Multi-task learning is a promising learning paradigm that aims at taking advantage of information shared in multiple related tasks to help improve the generalization performance of all tasks [8]. In multi-task learning, a model is trained with multiple objectives towards different tasks simultaneously, where all or some of the tasks are related. Researches have shown that learning multiple tasks jointly can achieve better performance than learning each task separately. Yang et al. propose a novel multi-task representation learning model that learns cross-task sharing structures at each layer of a neural network [37]. Hashimoto et al. propose a joint multi-task model for multiple NLP tasks [17]. The authors point out that

training a single network to model the hierarchical linguistic information from morphology, syntax to semantics can improve its generalization ability. Kendall et al. observe that the performance of the multi-task learning framework heavily depends on the weights of the objectives for different tasks [22]. They develop a novel method to learn the multi-task weightings by taking the homoscedastic uncertainty of each task into consideration.

3 Problem Statement

The teacher question detection task in live classrooms identifies questions from teachers’ speech and classify those questions into correct categories. In this section, we first introduce the method for coding questions and then formulate the problem of teacher question detection.

3.1 Question Coding

By analyzing thousands of classroom recordings and surveying hundreds of instructors and educators, we categorize teacher questions into the following four categories:

- **Knowledge-solicitation Question (KQ):** Questions that ask for a knowledge point or a factual answer. Some examples include: “What’s the solution to this problem?”, “What’s the distance between A and B?”, and “What is the area of this quadrilateral?”.
- **Open Question (OQ):** Questions to which no deterministic answer is expected. Open questions usually provoke a cognitive process of students such as explaining a problem and talking about knowledge points. Some Examples are: “How to solve this problem?”, “Can you share your ideas?”, and “Why did you do this problem wrong?”.
- **Procedural Question (PQ):** Questions that teachers use to manage the teaching procedure, such as testing teaching equipment, greeting students, and asking them something unrelated to course content. Examples are: “Can you hear me?”, “How are you doing?”, and “Have I told you about it?”.
- **Discourse-management Question (DQ):** Questions that teachers use to manage the discourses, such as making transitions or drawing students’ attention. Examples include: “Right?”, “Isn’t it?”, and “Excuse me?”.

We ask crowdsourcing annotators to code each utterance segment as non-question or one of the above four types of questions. The annotators code utterance segments by listening to the corresponding audio tracks. To ensure the coding quality, we first test the annotators on a set of 400 gold-standard examples. The 400 gold-standard examples are randomly sampled from the dataset and annotated by two experienced specialists in education. We only keep the top five annotators who achieve precision scores over 95% and 85% on the two-way and multi-way tasks on the gold-standard set to code the entire dataset.

3.2 Problem Formulation

We define the **two-way task** and the **multi-way task** for the teacher question detection problem as follows. Let $X = (x_1, \dots, x_n)$ be a transcribed utterance where x_i is the i -th word and n is the length of the utterance. In the two-way task, each utterance X is assigned a binary label $Y \in \{Q, NQ\}$ where Q indicates that X is a question and NQ indicates it is not a question. In the multi-way task, each utterance X is assigned a label $Y \in \{KQ, OQ, PQ, DQ, NQ\}$ where KQ , OQ , PQ , DQ indicate that X is a knowledge-solicitation question, open question, procedural question or discourse-management question, respectively, and NQ denotes that X is not a question. Both the two-way task and the multi-way task are treated as classification problems where we seek for predicting the label Y of a given utterance X .

4 The Proposed Framework

In this section, we present our framework for teacher question detection in both two-way and multi-way prediction settings. We first introduce the overview of the proposed framework. After that, we discuss the details of our neural natural language understanding module, which is a key component in our question detection framework.

4.1 The Framework Overview

The overall workflow of our end-to-end approach is shown in Fig. 1. Similar to [24], we efficiently process the large-volume classroom recordings by utilizing a well-studied voice activity detection (VAD) system to cut an audio recording into small pieces of utterances [30, 40]. The VAD algorithm is able to segment the audio stream into segments of utterances and filter out the noisy and silent ones. Then, each utterance segment is fed into an ASR system for transcription. After that, we build a neural NLU model to extract the semantically meaningful information within each sentence and make the final question detection prediction. Please note that as an end-to-end framework, our model can be integrated seamlessly into a run-time environment in the practical usage.




Fig. 1. The overall workflow of our end-to-end question detection framework.

4.2 The Neural Natural Language Understanding Model

In the task of text classification, traditional ML models only use simple word-level features, a.k.a., word embeddings. Due to the fact that such models are not able to extract contextual information, they fail to understand the sentence-level semantics and yield satisfactory detection performance. Therefore, in the work, we propose a neural NLU model to address above issues.

In our NLU module, given a sentence $X = (x_1, \dots, x_N)$ that contains N tokens, similar to Devlin et al. [11], we first insert a special token $[CLS]$ in front of the token sequence X . Then the sequence of the corresponding token embeddings $E = (E_{[CLS]}, E_1, \dots, E_N)$ is passed through multiple Transformer encoder layers [35]. Within each Transformer block, each token is repeatedly enriched by the combination of all the words in the sentence so that the contextualized information is captured. At last, we obtain the final hidden states $H = (H_{[CLS]}, H_1, \dots, H_N)$. We treat the final hidden state $H_{[CLS]}$ of the special token $[CLS]$ as the aggregated representation of the entire sentence and use $H_{[CLS]}$ for our two-way and multi-way prediction tasks. Our neural NLU module is shown in Fig. 2.

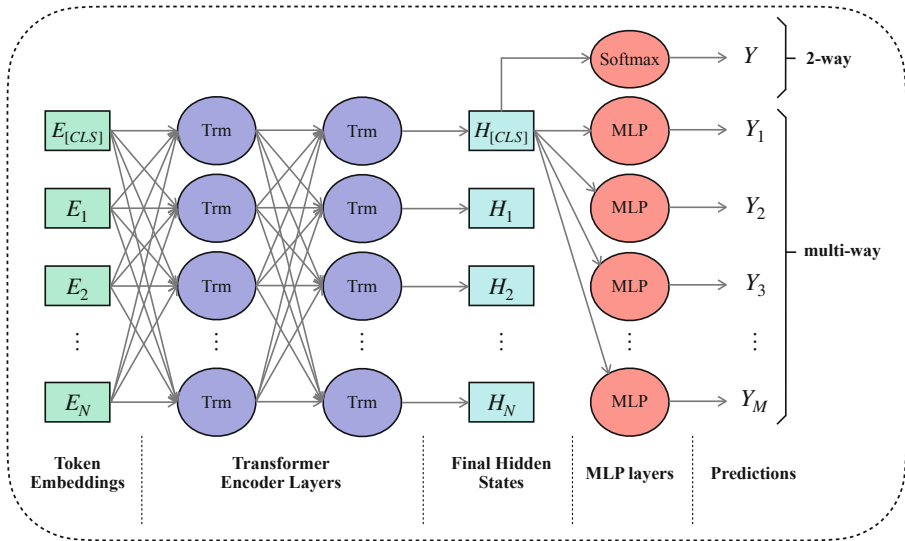


Fig. 2. An overview of our neural NLU module for question detection.

The NLU structures are different for the two-way and the multi-way tasks. For the two-way task, we feed the final hidden state of the special token $[CLS]$ into a Softmax layer for binary classification. While for the multi-way task, we convert the multi-class classification problem into multiple binary classification problems and train the model in a multi-task learning manner. Suppose that the number of classes is M . The final hidden state of $[CLS]$ is fed into M different multi-layer perceptron (MLP) layers to calculate the probabilities of

class memberships for each utterance segment. For class c_i , the cross-entropy loss function is

$$L_i = -(\mathbb{I}\{c_i = 0\} \log(1 - p_i) + \mathbb{I}\{c_i = 1\} \log p_i)$$

where $\mathbb{I}\{\cdot\}$ is an indicator function. c_i is 1 if the utterance segment belongs to the i -th class and is 0 otherwise. p_i is the predicted probability that the utterance segment belongs to the i -th class. We minimize the sum of cross-entropy loss functions of the M tasks, which is defined as $L_{multi} = \sum_{i=1}^M L_i$. In the inference phase, we make predictions for utterance segments by picking question types with the highest estimated probability.

In this multi-task learning model, multiple binary question classification tasks are learned simultaneously. This method provides several benefits. First, for different tasks, lower layers of the model are shared while the upper layers are different. The shared layers learn to extract the deep semantic features of the input utterance and the upper layers are responsible for making accurate question type predictions. This design yields more modeling capabilities. Second, in teacher question detection, different types of questions share some common patterns, such as interrogative words. But they typically have vastly different contents. When learning a task, the unrelated parts of other tasks can be viewed as auxiliary information, which prevents the model from overfitting and further improves the generalization ability of the model.

5 Experiment

To verify the effectiveness and superiority of our proposed model, we conduct extensive experiments on a real-world dataset. In this section, we first introduce our dataset that is collected from a real-world online learning platform. Then we describe the details of our experimental setup and the competitive baselines. Finally, we present and discuss the experimental results.

5.1 Dataset

We collect 548 classroom recordings of different subjects and grades from a third party K-12 online education platform¹. Recordings from both teacher and the students are stored separately in each online classroom. Here, we only focus on the teacher’s audio recording. The audio recordings are cut into utterance segments by a self-trained VAD system and each audio segment is transcribed by an ASR service (see Sect. 5.2). As a result, we obtain 39313 segments in total that are made up of 5314, 16934, and 17065 segments from classes in elementary school, middle school, and high school respectively. The average length of the segments is 3.5 seconds. The detailed segment-level per school-age and per subject question distribution is shown in Fig. 3(a).

¹ <https://www.xeslv1.com/>.




Fig. 3. Question distributions of our real-world education dataset.

As described in Sect. 3.1, each segment is labeled by five qualified annotators. The average pairwise Cohen’s Kappa agreement score is 0.696, which indicates a strong annotator agreement. Therefore, we choose to use majority votes as the final labels. The detailed distribution of questions with different types is shown in Fig. 3(b). We split the whole dataset into a training set, a validation set, and a test set with the proportion of 8:1:1, and the details of data statistics are shown in Table 1.

Table 1. Data statistics of the training set, the validation set, and the test set.

	KQ	OQ	PQ	DQ	NQ	Total
Training	6450	3551	2786	8514	10149	31450
Validation	861	431	328	1104	1207	3931
Test	778	458	372	1086	1238	3932

5.2 Implementation Details

In the work, we train our VAD model by using a four-layer DNN neural network to distinguish normal human utterances from background noises and silences [33]. Similar to Blanchard et al. [3], we find that publicly available ASR service may yield inferior performance in the noisy and dynamic classroom environments. Therefore, we train our own ASR models on the classroom specific datasets based on a deep feed-forward sequential memory network proposed by Zhang et al. [38]. Our ASR has a word error rate of 28.08% in our classroom settings.

Language model pre-training techniques have achieved great improvements on various NLU tasks [11, 36]. In the implementation of our neural NLU model,

we first pre-train the model with a large-scale language corpus and then use question specific classroom data to conduct the model fine-tuning. Here, we adopt the pre-trained NEZHA-base model released by Wei et al. [36]. In the multi-task setting, we apply a two-layer MLP with hidden sizes 256, 64 for each class. The output is passed through a sigmoid function to calculate the predictive probability. An optimal set of hyper-parameters is picked according to the model’s performance on the validation set and we report its performance on the test set.

5.3 Baselines

We compare our approach with the following representative baseline methods: (1) Logistic Regression (LR) [23], (2) K-Nearest Neighbor (KNN) [14], (3) Random Forest (RF) [18], (4) Support Vector Machine (SVM) [10], (5) Gradient Boosted Decision Tree (GBDT) [13] and (6) Bidirectional Long Short Term Memory Network (Bi-LSTM) [39]. For the first five baselines, we use the sentence embedding of a given transcribed utterance as the feature vector for classification. The sentence embedding is computed by taking the average of the pre-trained word embeddings within each sentence. For Bi-LSTM, the word embeddings are fed into an LSTM network sequentially and the concatenation vector of the final hidden states in two directions is fed into a Softmax layer for classification.

5.4 Experimental Results

We show the results of the two-way and the multi-way tasks in Table 2 and Table 3, respectively. In the two-way task, we report the classification results of different models in terms of accuracy, precision, recall, F1 score, and AUC score, respectively. In the multi-way task, we report the classification results on each question type in terms of F1 score, as well as the overall results in terms of precision, recall and F1 scores from both micro and macro perspectives [34]. From Table 2 and Table 3, we make the following observations:

Table 2. Performance comparison of the **two-way task**.

	Accuracy	Precision	Recall	F1 score	AUC
LR	0.724	0.863	0.711	0.779	0.811
KNN	0.740	0.745	0.943	0.832	0.763
RF	0.766	0.758	0.968	0.850	0.824
SVM	0.798	0.874	0.824	0.848	0.854
GBDT	0.817	0.826	0.929	0.874	0.837
Bi-LSTM	0.873	0.882	0.940	0.910	0.915
Our Model	0.885	0.888	0.952	0.919	0.933

Table 3. Performance comparison of the **multi-way task**. ma-Pre., ma-Rec., mi-F1 and ma-F1 represent the macro precision, macro recall, micro F1 score and macro F1 score respectively.

Type	KQ	OQ	PQ	DQ	NQ	Overall			
	F1	F1	F1	F1	F1	ma-Pre.	ma-Rec.	mi-F1	ma-F1
LR	0.621	0.584	0.532	0.734	0.620	0.611	0.634	0.637	0.618
KNN	0.450	0.461	0.450	0.616	0.540	0.580	0.490	0.540	0.503
RF	0.564	0.454	0.483	0.699	0.632	0.661	0.537	0.612	0.566
SVM	0.644	0.629	0.561	0.791	0.694	0.655	0.681	0.688	0.664
GBDT	0.629	0.583	0.516	0.758	0.676	0.662	0.616	0.668	0.632
Bi-LSTM	0.743	0.751	0.654	0.914	0.778	0.769	0.769	0.794	0.768
Our Model	0.767	0.768	0.686	0.912	0.793	0.781	0.794	0.808	0.785

- First, in terms of both the two-way and multi-way tasks and most of the evaluation metrics, our model outperforms all the baseline methods. Due to the fact that our dataset consists of different subjects, school-ages, teachers and question types, we believe that the performance improvements achieved by our approach show the adaptability and robustness towards the real challenging educational scenarios.
- Second, by comparing the performances of the models on different types of questions, we find that procedural questions are relatively harder to identify compared to discourse-management questions. We believe the reason is that procedural questions typically involve a wide range of topics and appear in diverse forms. While discourse-management questions are short and succinct, and their forms are relatively fixed.
- Third, the baselines LR, KNN, RF, SVM, and GBDT achieve unsatisfactory performance in both tasks. Because they simply average the word embeddings as the features for classification, which fail to capture any contextualized information. Bi-LSTM performs better by learning better contextualized representations. The proposed framework outperforms Bi-LSTM because of its powerful ability of deep semantic understanding learned through the Transformer layers, the pre-training procedure, and the multi-task learning technique.

6 Conclusion

In this paper, we present a novel framework for the automatic detection of teacher questions in online classrooms. We propose a neural NLU model, which is able to automatically extract semantic features from teachers’ utterances and adaptively generalize across recordings of different subjects and speakers. Experiments conducted on a real-world education dataset validate the effectiveness of our model in both two-way and multi-way tasks. As a future research direction,

we are going to explore the relationship between the use of teacher questions and student achievement in live classrooms, thus we can make corresponding suggestions to teachers to improve their teaching efficiency.

Acknowledgements. Haochen Liu and Jiliang Tang are supported by the National Science Foundation of United States under IIS1714741, IIS1715940, IIS1715940, IIS1845081 and IIS1907704.

References

1. Applebee, A.N., Langer, J.A., Nystrand, M., Gamoran, A.: Discussion-based approaches to developing understanding: classroom instruction and student performance in middle and high school English. *Am. Educ. Res. J.* **40**(3), 685–730 (2003)
2. Beck, I.L., McKeown, M.G., Sandora, C., Kucan, L., Worthy, J.: Questioning the author: a yearlong classroom implementation to engage students with text. *Elementary Sch. J.* **96**(4), 385–414 (1996)
3. Blanchard, N., et al.: A study of automatic speech recognition in noisy classroom environments for automated dialog analysis. In: Conati, C., Heffernan, N., Mitrovic, A., Verdejo, M.F. (eds.) *AIED 2015. LNCS (LNAI)*, vol. 9112, pp. 23–33. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-19773-9_3
4. Blanchard, N., D’Mello, S., Olney, A.M., Nystrand, M.: Automatic classification of question & answer discourse segments from teacher’s speech in classrooms. *Int. Educ. Data Mining Soc.* (2015)
5. Blanchard, N., et al.: Identifying teacher questions using automatic speech recognition in classrooms. In: *Proceedings of the 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pp. 191–201 (2016)
6. Blanchard, N., et al.: Semi-automatic detection of teacher questions from human-transcripts of audio in live classrooms. *Int. Educ. Data Mining Soc.* (2016)
7. Blosser, P.E.: *How to Ask the Right Questions*. NSTA Press (1991)
8. Caruana, R.: Multitask learning. In: Thrun, S., Pratt, L. (eds.) *Learning to Learn*, pp. 95–133. Springer, Heidelberg (1998). https://doi.org/10.1007/978-1-4615-5529-2_5
9. Chen, J., Li, H., Wang, W., Ding, W., Huang, G.Y., Liu, Z.: A multimodal alerting system for online class quality assurance. In: Isotani, S., Millán, E., Ogan, A., Hastings, P., McLaren, B., Luckin, R. (eds.) *AIED 2019. LNCS (LNAI)*, vol. 11626, pp. 381–385. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-23207-8_70
10. Cortes, C., Vapnik, V.: Support-vector networks. *Mach. Learn.* **20**(3), 273–297 (1995)
11. Devlin, J., Chang, M., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. In: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, 2–7 June 2019*, pp. 4171–4186 (2019)
12. Donnelly, P.J., Blanchard, N., Olney, A.M., Kelly, S., Nystrand, M., D’Mello, S.K.: Words matter: automatic detection of teacher questions in live classroom discourse using linguistics, acoustics, and context. In: *Proceedings of the Seventh International Learning Analytics & Knowledge Conference, Vancouver, BC, Canada, 13–17 March 2017*, pp. 218–227 (2017)

13. Drucker, H., Cortes, C.: Boosting decision trees. In: *Advances in Neural Information Processing Systems*, pp. 479–485 (1996)
14. Fix, E.: *Discriminatory analysis: nonparametric discrimination, consistency properties*. USAF school of Aviation Medicine (1951)
15. Gamoran, A., Kelly, S.: Tracking, instruction, and unequal literacy in secondary school English. In: *Stability and Change in American Education: Structure, Process, and Outcomes*, pp. 109–126 (2003)
16. Graesser, A.C., Person, N.K.: Question asking during tutoring. *Am. Educ. Res. J.* **31**(1), 104–137 (1994)
17. Hashimoto, K., Xiong, C., Tsuruoka, Y., Socher, R.: A joint many-task model: growing a neural network for multiple NLP tasks. *CoRR abs/1611.01587* (2016). <http://arxiv.org/abs/1611.01587>
18. Ho, T.K.: Random decision forests. In: *Proceedings of 3rd International Conference on Document Analysis and Recognition*, vol. 1, pp. 278–282. IEEE (1995)
19. Kane, T.J., Staiger, D.O.: Gathering feedback for teaching: combining high-quality observations with student surveys and achievement gains. Research paper. met project. Bill & Melinda Gates Foundation (2012)
20. Kelly, S.: Classroom discourse and the distribution of student engagement. *Soc. Psychol. Educ.* **10**(3), 331–352 (2007)
21. Kelly, S.: Race, social class, and student engagement in middle school English classrooms. *Soc. Sci. Res.* **37**(2), 434–448 (2008)
22. Kendall, A., Gal, Y., Cipolla, R.: Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. *CoRR abs/1705.07115* (2017). <http://arxiv.org/abs/1705.07115>
23. Kleinbaum, D.G., Dietz, K., Gail, M., Klein, M., Klein, M.: *Logistic Regression*. Springer, Heidelberg (2002)
24. Li, H., et al.: Multimodal learning for classroom activity detection. In: *2020 IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 9234–9238. IEEE (2020)
25. Liu, Z., et al.: Dolphin: a spoken language proficiency assessment system for elementary education. In: *Proceedings of the Web Conference 2020*, pp. 2641–2647. ACM (2020)
26. MacNeilley, L.H.: Opening dialogue: understanding the dynamics of language and learning in the English classroom by Martin Nystrand with Adam Gamoran, Robert Kachur, and Catherine Prendergast. *Language* **74**(2), 444–445 (1998)
27. Nystrand, M., Gamoran, A.: Instructional discourse, student engagement, and literature achievement. In: *Research in the Teaching of English*, pp. 261–290 (1991)
28. Nystrand, M., Wu, L.L., Gamoran, A., Zeiser, S., Long, D.A.: Questions in time: investigating the structure and dynamics of unfolding classroom discourse. *Discourse Process.* **35**(2), 135–198 (2003)
29. Samei, B., et al.: Domain independent assessment of dialogic properties of classroom discourse. In: *Proceedings of the 7th International Conference on Educational Data Mining*, London, UK, 4–7 July 2014, pp. 233–236 (2014)
30. Sohn, J., Kim, N.S., Sung, W.: A statistical model-based voice activity detection. *IEEE Sig. Process. Lett.* **6**(1), 1–3 (1999)
31. Stivers, T., Enfield, N.J.: A coding scheme for question-response sequences in conversation. *J. Prag.* **42**(10), 2620–2626 (2010)
32. Sweigart, W.: Classroom talk, knowledge development, and writing. In: *Research in the Teaching of English*, pp. 469–496 (1991)
33. Tashev, I., Mirsamadi, S.: DNN-based causal voice activity detector. In: *Information Theory and Applications Workshop* (2016)

34. Van Asch, V.: Macro-and micro-averaged evaluation measures. Belgium: CLiPS49 (2013)
35. Vaswani, A., et al.: Attention is all you need. In: Advances in Neural Information Processing Systems, pp. 5998–6008 (2017)
36. Wei, J., et al.: NEZHA: neural contextualized representation for Chinese language understanding. CoRR abs/1909.00204 (2019). <http://arxiv.org/abs/1909.00204>
37. Yang, Y., Hospedales, T.: Deep multi-task representation learning: A tensor factorisation approach. arXiv preprint [arXiv:1605.06391](https://arxiv.org/abs/1605.06391) (2016)
38. Zhang, S., Lei, M., Yan, Z., Dai, L.: Deep-FSMN for large vocabulary continuous speech recognition. In: 2018 IEEE International Conference on Acoustics, Speech and Signal Processing, pp. 5869–5873. IEEE (2018)
39. Zhang, S., Zheng, D., Hu, X., Yang, M.: Bidirectional long short-term memory networks for relation classification. In: Proceedings of the 29th Pacific Asia Conference on Language, Information and Computation, pp. 73–78 (2015)
40. Zhang, X.L., Wu, J.: Deep belief networks based voice activity detection. IEEE Trans. Audio Speech Lang. Process. **21**(4), 697–710 (2012)