

# QuesNet: A Unified Representation for Heterogeneous Test Questions

Yu Yin<sup>1</sup>, Qi Liu<sup>1</sup>, Zhenya Huang<sup>1</sup>, Enhong Chen<sup>1,\*</sup>, Wei Tong<sup>1</sup>, Shijin Wang<sup>2,3</sup>, Yu Su<sup>2,3</sup>

<sup>1</sup>Anhui Province Key Laboratory of Big Data Analysis and Application, School of Computer Science and Technology & School of Data Science, University of Science and Technology of China,

{yxonic,huangzhy}@mail.ustc.edu.cn, {qiliuql,cheneh}@ustc.edu.cn, tongw@mail.neea.edu.cn

<sup>2</sup>FLYTEK Research, iFLYTEK CO., LTD., <sup>3</sup>State Key Laboratory of Cognitive Intelligence, {sjwang3,yusu}@iflytek.com

## ABSTRACT

Understanding learning materials (e.g. test questions) is a crucial issue in online learning systems, which can promote many applications in education domain. Unfortunately, many supervised approaches suffer from the problem of scarce human labeled data, whereas abundant unlabeled resources are highly underutilized. To alleviate this problem, an effective solution is to use pre-trained representations for question understanding. However, existing pre-training methods in NLP area are infeasible to learn test question representations due to several domain-specific characteristics in education. First, questions usually comprise of heterogeneous data including content text, images and side information. Second, there exists both basic linguistic information as well as domain logic and knowledge. To this end, in this paper, we propose a novel pre-training method, namely QuesNet, for comprehensively learning question representations. Specifically, we first design a unified framework to aggregate question information with its heterogeneous inputs into a comprehensive vector. Then we propose a two-level hierarchical pre-training algorithm to learn better understanding of test questions in an unsupervised way. Here, a novel holed language model objective is developed to extract low-level linguistic features, and a domain-oriented objective is proposed to learn high-level logic and knowledge. Moreover, we show that QuesNet has good capability of being fine-tuned in many question-based tasks. We conduct extensive experiments on large-scale real-world question data, where the experimental results clearly demonstrate the effectiveness of QuesNet for question understanding as well as its superior applicability.

## CCS CONCEPTS

• **Computing methodologies** → **Information extraction**; *Neural networks*; • **Applied computing** → **Education**.

## KEYWORDS

question representation, heterogeneous data, pre-training

## ACM Reference Format:

Yu Yin, Qi Liu, Zhenya Huang, Enhong Chen, Wei Tong, Shijin Wang, Yu Su. 2019. QuesNet: A Unified Representation for Heterogeneous Test Questions. In The 25th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '19), August 4-8, 2019, Anchorage, AK, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3292500.3330900>

## 1 INTRODUCTION

In recent years, many online learning systems, such as Khan Academy and LeetCode, have gained more and more popularities among learners of all ages from K12, to college, and even adult due to its convenience and autonomy [1, 20]. In these platforms, it is necessary to well organize such abundant questions in advance [18]. For example, we need to sort them by difficulty attributes or create curricula designs with their knowledge concepts. Such managements are vitally necessary since they could help students save effort to locate required questions for targeted training and efficient learning [6]. Therefore, it is essential to find an effective way for systems to understand test questions. In fact, since it is the fundamental issue promoting many question-based applications, such as difficulty estimation [13], knowledge mapping [10, 33] and score prediction [27], much attention has been attracted from both system creators and researchers.

In the literature, many efforts have been developed for understanding question content by taking advantage of natural language processing (NLP) techniques [10, 13]. In general, existing solutions usually design end-to-end frameworks, where the questions are represented as syntactic patterns or semantic encodings, and furthermore directly optimized in specific downstream tasks by supervised learning [13, 29]. However, these task-specific approaches mostly require substantial amounts of manually labeled data (e.g., labeled difficulty), which restricts their performance in many learning systems that suffer from the sparsity problem of limited label annotations [13]. Comparatively, in this paper, we aim to explore an unsupervised way by taking full advantage of large-scale unlabeled question corpus available for question representation.

Unfortunately, it is a highly challenging task. Although several pre-training methods have shown their superiority in NLP on tasks such as question answering [5, 23], they just exploit linguistic contexts with homogeneous text. They are infeasible for understanding question materials due to following domain-specific characteristics. First, test questions contain heterogeneous data. For example, typical math questions in Figure 1 comprise of multiple parts with different forms including text (red), image (green) and side information such as knowledge concept (yellow). All these kinds information are crucial for question understanding, which requires

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

KDD '19, August 4–8, 2019, Anchorage, AK, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6201-6/19/08...\$15.00

<https://doi.org/10.1145/3292500.3330900>

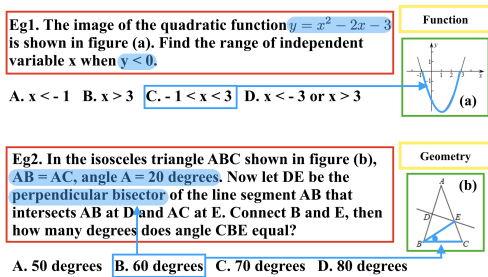


Figure 1: Two examples of heterogeneous questions.

us to find an appropriate way to aggregate them to a comprehensive representation. Second, for a certain question, not only should we extract its basic linguistic context, but we also need to carefully consider the advanced logic information, which is a nontrivial problem. As shown in Figure 1, in addition to linguistic context and relations from its content, a test question also contains high-level logic, taking the information of four options into consideration. The right answers are more related to the question meaning compared with the wrong ones, reflecting the unique mathematical logic and knowledge. E.g., to find the right answer (B) in question example 2, we need to focus more on the expression (“ $AB = AC, \dots$ ”) in text and the related  $\angle CBE$  in the image. Third, in practice, the learned question representations should have great accessibility and be easy to apply to downstream tasks such as difficulty estimation. In actual educational tasks, question representations are often used as part of a complex model, which requires the method to have simple yet powerful structure and easy to mix-in into task-specific models.

To address the above challenges, in this paper, we provide a unified domain-specific method, namely QuesNet, for comprehensive learning test question representations. Generally, QuesNet is able to aggregate heterogeneous data of a certain question into an integrated form and gain a deeper understanding with the benefits of both low-level linguistic information and high-level domain logic knowledge. It can also be naturally applied to many downstream methods, which effectively enhances the performance on-the-fly. Specifically, we first design a unified model based on Bi-LSTM and self-attention structures to aggregate question information with its heterogeneous inputs into a vector representation. Then we propose a two-level hierarchical pre-training algorithm to learn better understandings of test questions. On the lower level, we develop a novel holed language model (HLM) objective to help QuesNet extract linguistic context relations from basic inputs (i.e., words, images, etc.). Comparatively, on higher level pre-training, we propose a domain-specific objective to learn advanced understanding for each question, which preserves its domain logic and knowledge. With objectives, QuesNet could learn the integrated representations of questions in an unsupervised way. Furthermore, we demonstrate how to apply QuesNet to various typical question-based tasks with fine-tuning in education domain including difficulty estimation, knowledge mapping and score prediction. We conduct extensive experiments on large-scale real world question data, with three domain-specific tasks. The experimental results clearly demonstrate that QuesNet has good capability of understanding test questions and also superior applicability of fine-tuning.

## 2 RELATED WORK

We briefly summarize our related works as follows.

### 2.1 Question Understanding

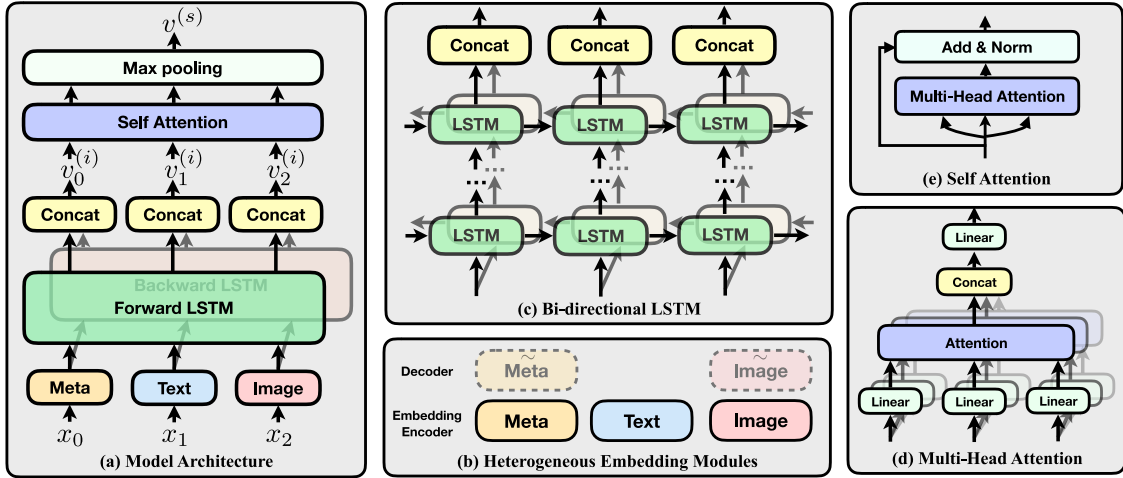
Question understanding is a fundamental task in education [25]. Existing approaches could be roughly divided into two categories: rule-based representation and vector-based representation. For rule-based representation, scholars devote efforts to designing many fine-grained rules or grammars and learn to understand questions by parsing the question text into semantic trees or pre-defined features [7, 9]. However, these initiative works heavily rely on expertise for designing effective rule patterns, which is obviously labor intensive. Comparatively, in vector-based representation methods, each question could be learned as a semantic vector in latent space automatically through many natural language processing (NLP) techniques [28, 30]. Recently, as an extension and combination of previous studies, deep learning techniques have become state-of-the-art models due to their superiority of learning complex semantics [13, 35]. For example, Tan et al. [29] used Long Short-Term Memory (LSTM) model to capture the long-term dependency of question sentences. Huang et al. [13] utilized convolutional neural network for question content understanding, targeting at the difficulty estimation task. Although great success have been achieved, all these supervised methods suffer from the problem of scarce labeled data. That is, with labels only in specific task supervising both question understanding and task modeling parts, the understanding of the question is quite limited, while large volume of unlabeled question data bank is not leveraged. Moreover, none of the work have considered different question input forms, which causes an information loss for heterogeneous question understanding.

### 2.2 Text Pre-training

Recent years have witnessed the development of pre-training methods, which is a good way to make the most of unlabeled corpus in NLP field [5]. These methods can be divided into two categories: feature-based methods, where text is represented by some sort of feature extractors as fixed vectors [22, 23], and pre-training based methods, where parameters of model are pre-trained on corpus and then fine-tuned to specific tasks [5, 12]. Among them, the most successful model would be BERT [5]. It utilizes Transformer [30] together with language related pre-training goals, solving many NLP tasks with impressive performance. Although these solutions have been fully examined in a range of NLP tasks, they could hardly be directly applied to the understanding of test questions mainly due to the following three reasons. First, test questions are heterogeneous, where much information in other input forms would be ignored. Second, test questions contain much domain logic and knowledge to be understood and represented other than just linguistic features, which makes it hard for models to capture. Third, many of these approaches are difficult to be applied due to the need of model modification, which is inconvenient under many education setup.

### 2.3 Question-based Applications

There are many question-based applications in education domain, which play important roles in traditional classroom setting or online learning system [1]. The representative tasks include difficulty



**Figure 2: QuesNet model architecture.** (a) shows the whole structure, which can be divided into three layers: 1. *Embedding Layer*, with heterogeneous embedding modules in (b); 2. *Content Layer*, with bi-directional LSTM detailed in (c); and 3. *Sentence Layer*, which is a global self attention layer demonstrated in (e), the implementation shown in (d).

estimation [3, 13], knowledge mapping [4] and score prediction [24]. Specifically, difficulty estimation requires us to evaluate how difficult a question is from its content without preparing a group of students to test on it. Knowledge mapping aims at automatically mapping a question to its corresponding knowledge points. Score prediction is a task of predicting how well a student performs on a specific question with their exercising history. All these applications benefit the system management and services, such as personalized recommendation [16].

Our work provides a unified representation for heterogeneous test questions compared with previous studies, and provides a solid backbone for applications in use of questions. We take the heterogeneous nature of test questions and the difficulty of understanding domain information into consideration and design a more powerful yet accessible pre-training algorithm. With heterogeneous question representation model and two-level pre-training, QuesNet captures much more information from test questions.

### 3 QUESNET: MODELING AND PRE-TRAINING

In this section, we introduce QuesNet modeling and pre-training in detail. First we give a formal definition of the question representation problem. Then, we describe the QuesNet architecture for heterogeneous question representation. Afterwards we describe the pre-training process of QuesNet, i.e. the two-level pre-training algorithm. Finally, in Section 3.4, we discuss how to apply QuesNet to downstream tasks and do fine-tuning.

#### 3.1 Problem Definition

In this subsection, we formally introduce the question representation problem and clarify mathematical symbols in this paper.

In our setup, each test question  $q$  is given as input in a heterogeneous form, which contains one or all kinds of content including text, images and side information (meta data such as knowledge).

Formally, we could define it as a sequence  $x = \{x_t\}, t \in \{1, \dots, T\}$ , where  $T$  is the length of the input question sequence, together with side information as a one-hot encoded vector  $m \in \mathbb{R}^K$ , where  $K$  is the number of categories in side information. Each input item  $x_t$  is either a word from a vocabulary (including formula piece), or an image in size  $W \times H$ .

For better usability, the final representation of each question  $q$ , i.e. the output, should contain both individual content representation as a sequence of vectors and the whole question representation as a single vector. We will see in Section 3.4 why all these representations are necessary. With the setup stated above, more formally, we define the question representation problem as follows:

*Definition 3.1. (Question Representation Problem).* Given a question  $q$  with heterogeneous input, as a sequence  $\mathbf{x} = \{x_0, x_1, \dots, x_T\}$ , with side information as  $x_0 = m$ , sequence length as  $T$ , each input content as  $x_i$  (either a word or an  $W \times H$  image), our goal is to represent  $q$  as a sequence of content representation vectors  $\mathbf{v}^{(i)} = \{v_0^{(i)}, \dots, v_T^{(i)}\}$  and one sentence representation vector  $v^{(s)}$ , each of which is of dimension  $N$ . The representation should capture as much information as possible.

In the following sections, we will address the main three challenges: (1) how QuesNet generates question representation; (2) how the representation is pre-trained; (3) how the representation is applied to downstream tasks.

#### 3.2 QuesNet Model Architecture

QuesNet model maps a heterogeneous question  $q$  to a unified final representation  $(\mathbf{v}^{(i)}, v^{(s)})$ . The architecture is shown in Figure 2 (a), which can be seen as three layers: Embedding Layer, Content Layer and Sentence Layer. Specifically, given a question  $q$ , in Embedding Layer, its heterogeneous input embedding is performed. Then in Content Layer, Bi-LSTM is used to model different input content

and generate each content representation  $\mathbf{v}^{(i)}$ . Finally, in Sentence Layer, we use self-attention to combine vectors in an effective way.

**3.2.1 Embedding Layer.** We first introduce the building blocks of the Embedding Layer. The aim of this layer is to project heterogeneous input content to a unified space, which enables our model for different input forms. In order to do so, in first layer, we setup embedding modules to map each kind of inputs into fixed length vectors. Embedding module for words is a mapping  $\text{emb}_w$  with parameters as  $\theta_{we}$ , which directly maps each word in the vocabulary to a vector of size  $N_e$ . Image embedding module  $\text{emb}_i$ , with parameters denoted as  $\theta_{ie}$ , as shown in the upper part of Figure 2 (b), consists of three convolutional layers followed by activations [32]. Features are also max-pooled into a vector of size  $N_e$ . Meta data embedding module  $\text{emb}_m$ , with parameters denoted as  $\theta_{me}$ , as shown in Figure 2 (b), uses two layers of fully-connect neural network to embed input meta data as a vector of size  $N_e$ .

With these basic embedding modules, for each input item  $x_t$  in  $q$ , we generate an embedded vector  $e_t$  in the first layer, so that we can get an embedding sequence  $\mathbf{e} = \{e_0, e_1, \dots, e_T\}$  from input  $\mathbf{x} = \{x_0, x_1, \dots, x_T\}$ . Formally:

$$e_t = \begin{cases} \text{emb}_w(x_t; \theta_{we}), & \text{if } x_t \text{ is word,} \\ \text{emb}_i(x_t; \theta_{ie}), & \text{if } x_t \text{ is image,} \\ \text{emb}_m(x_t; \theta_{me}), & \text{if } x_t \text{ is meta data.} \end{cases}$$

**3.2.2 Content Layer.** In this layer, we aim at modeling relation and context for each input item. Existing methods like LSTM [11] only cares about context on one side, while in Transformer [30], context and relation modeling relies on position embedding, which loses some locality. Therefore, with embedded vector sequence  $\mathbf{e}$  described above as input, we incorporate a multi-layer bi-directional LSTM structure [14], which is more capable of gaining context information. Here, we choose Bi-LSTM because it can make the most of contextual content information of question sentence from both forward and backward directions [14, 17]. Specifically, given question embedding sequence  $\mathbf{e} = \{e_0, e_1, \dots, e_T\}$ , we set the input of the first layer of LSTM as  $\vec{\mathbf{h}}^{(0)} = \overleftarrow{\mathbf{h}}^{(0)} = \{e_0, e_1, \dots, e_T\}$ . At each position  $t$ , forward hidden states  $(\vec{h}_t^{(l)}, \vec{c}_t^{(l)})$  and backward hidden states  $(\overleftarrow{h}_t^{(l)}, \overleftarrow{c}_t^{(l)})$  at each layer  $l$  are updated with input from previous layer  $\overleftarrow{h}_t^{(l-1)}$  or  $\overleftarrow{h}_t^{(l-1)}$  for each direction, and previous hidden states  $(\vec{h}_{t-1}^{(l)}, \vec{c}_{t-1}^{(l)})$  for forward direction or  $(\overleftarrow{h}_{t+1}^{(l)}, \overleftarrow{c}_{t+1}^{(l)})$  for backward direction in a recurrent formula as:

$$\begin{aligned} \vec{h}_t^{(l)}, \vec{c}_t^{(l)} &= \text{LSTM}(\vec{h}_t^{(l-1)}, \vec{h}_{t-1}^{(l)}, \vec{c}_{t-1}^{(l)}; \vec{\theta}_{\text{LSTM}}), \\ \overleftarrow{h}_t^{(l)}, \overleftarrow{c}_t^{(l)} &= \text{LSTM}(\overleftarrow{h}_t^{(l-1)}, \overleftarrow{h}_{t+1}^{(l)}, \overleftarrow{c}_{t+1}^{(l)}; \overleftarrow{\theta}_{\text{LSTM}}), \end{aligned}$$

where recurrent formula follows Hochreiter et al. [11]

More layers are needed for modeling deeper relations and context. With  $L$  Bi-LSTM layers, deep linguistic information is able to be captured in hidden states. As hidden state at each direction only contains one-side context, it is beneficial to combine the hidden state of both directions into one vector at each step. Therefore, we obtain content representation at each time step  $t$  as:

$$\mathbf{v}_t^{(i)} = \text{concatenate}(\vec{h}_t^{(L)}, \overleftarrow{h}_t^{(L)}).$$

**3.2.3 Sentence Layer.** After we model lower-level linguistic features, we still have to aggregate this information in a way that focus on long-term and global complex relations, so that domain logic and knowledge can be captured. To this end, the third Sentence Layer consists of a self-attention module for aggregating item representation vector sequence  $\mathbf{v}^{(i)}$  into a sentence representation  $\mathbf{v}^{(s)}$ . Compared with LSTM that focuses on context, attention mechanism is more capable of modeling long-term logic and global information [30]. Following Vaswani et al. [30], we use Multi-Head Attention module to perform global self attention. Given a set of queries of dimension  $N$  (as matrix)  $Q$ , keys (as matrix)  $K$ , values (as matrix)  $V$ , Multi-Head Attention computes output matrix as:

$$\begin{aligned} \text{MultiHead}(Q, K, V) &= \text{concatenate}(\text{head}_1, \dots, \text{head}_H)W^O, \\ \text{head}_j &= \text{Attention}(QW_j^Q, KW_j^K, VW_j^V), \\ \text{Attention}(Q, K, V) &= \text{softmax}(QK^T/\sqrt{N})V, \end{aligned}$$

where  $H$  is the number of attention heads,  $W^O, W_j^Q, W_j^K, W_j^V$  are some projection matrices. Intuitively, Multi-Head Attention module performs several different attentions in parallel, which helps it to aggregate high level logic and knowledge for lower layers.

Within our setup, we use self-attention to aggregate content vectors  $\mathbf{v}^{(i)}$  together with position embedding  $pe^{(i)}$ , by setting  $Q, K, V$  in Multi-Head Attention all as  $\mathbf{v}_+^{(i)} = \text{concatenate}(\mathbf{v}^{(i)}, pe^{(i)})$ , and then attended values in all time step into one single vector with max-pooling. More formally:

$$\mathbf{v}^{(s)} = \max \left\{ \text{MultiHead}(\text{LayerNorm}(\mathbf{v}_+^{(i)}, \mathbf{v}_+^{(i)}, \mathbf{v}_+^{(i)}) + \mathbf{v}_+^{(i)}) \right\},$$

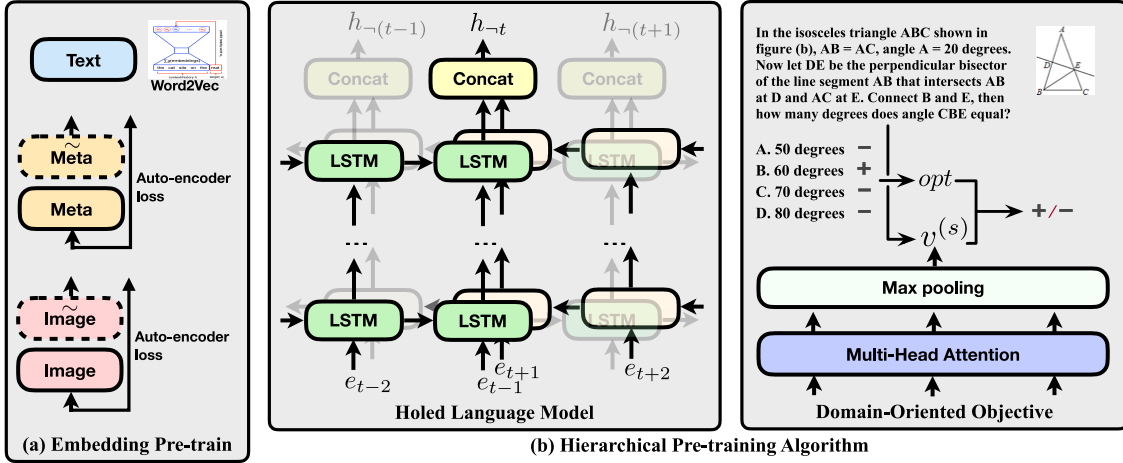
where LayerNorm refers to layer normalization technique proposed by Ba et al. [2], and position embedding follows Vaswani et al. [30]

Till now we already generated a unified representation of a question. To summarize, with embedding layer, we embed heterogeneous content into a unified form. Then with multi-layer Bi-LSTM in content layer, we capture deep linguistic relation and context. Finally in sentence layer, we aggregate the information into a single vector with high level logic and knowledge.

### 3.3 Pre-training

However, we still need a way to learn all the linguistic features and domain logic on large unlabeled question corpus, which we will describe in this subsection. Specifically, we fully describe how to pre-train QuesNet to capture both linguistic features and domain logic and knowledge from question corpus. For this purpose, as shown in Figure 3, we design a novel hierarchical pre-training algorithm. We first separately pre-train each embedding modules. Then in the main pre-training process, we propose two level hierarchical objectives. At low level pre-training, we proposed a novel holed language model as the objective for learning low-level linguistic features. At high-level learning, a domain-oriented objective is added for learning high-level domain logic and knowledge. The objectives of both levels are learned together within one pre-training process.

**3.3.1 Pre-training of Embedding.** We first separately pre-train each embedding module to set up better initial weights for them. For word embedding, we incorporate *Word2Vec* [19] on the whole corpus to get an initial word to vector mapping. For image and side



**Figure 3: Pre-training of QuesNet.** For pre-training, first, as shown in (a), *Embedding Pre-training* is done in advance. Then in (b), two hierarchical objectives are defined: the low-level objective, Holed Language Model (HLM, middle), and the high-level domain-oriented objective (right).

information embedding, we first construct decoders for each embedding that decodes the vector given by embedding module. Then we train these embedding modules using auto-encoder losses [21]. If we take image embedding module  $emb_i$  with parameters  $\theta_{ie}$  as an example, we first construct image decoder  $\tilde{emb}_i$  also with trainable parameters  $\theta_{id}$ . Then on all images in the corpus  $I$ , we can construct auto-encoder loss as:

$$\mathcal{L}_I = \sum_{x \in I} l(\tilde{emb}_i(emb_i(x)), x),$$

where  $l(y, x)$  is a loss function that measures distance between  $y$  and  $x$ , such as mean-squared-error (MSE) loss. Then initial weights of image embedding module would be:

$$\hat{\theta}_{ie} = \arg \max_{\theta_{ie}} \mathcal{L}_I.$$

We can train initial values for side information embedding similarly. With all side information as  $\mathcal{L}_M$ , side information decoder  $\tilde{emb}(\cdot; \theta_{md})$  is implemented as a multi-layer fully-connected neural network, and the initial weights of it would be:

$$\hat{\theta}_{me} = \arg \max_{\theta_{me}} \mathcal{L}_M.$$

**3.3.2 Holed Language Model.** The pre-training objective at low level aims at learning linguistic features from large corpus. Language model (LM) being the most used unsupervised linguistic feature learning objective, is limited by its one-directional nature. In this paper, we proposed a novel holed language model (HLM) that jointly combines context from both sides. Intuitively, the objective of HLM is to fill up every word with both left and right side context of it. It is different from the bi-directional LM implementation in ELMo [23] where context from both sides are trained separately without any interaction. And it does not rely on random masking of tokens as BERT does, which is much more sample efficient.

In HLM, comparative to traditional language model, the probability of input content at each position  $t$  is conditioned by its context at

both sides, and our objective is to maximize conditioned probability at each position. Formally speaking, for an input sequence  $\mathbf{x}$ , the objective of HLM calculates:

$$\mathcal{L}_{\text{HLM}} = - \sum_{t=0}^T \log P(x_t | x_{-t}),$$

where we use  $x_{-t}$  to stand for all other inputs that are not  $x_t$ , and the goal is to minimize  $\mathcal{L}_{\text{HLM}}$  (sum of negative log likelihood). As described in Section 3.2, inputs on the left of position  $t$  are modeled in the adjacent left hidden vector  $\vec{h}_{t-1}$ , on the right are modeled in the adjacent right hidden vector  $\overleftarrow{h}_{t+1}$ . Therefore, the conditional probability of each input item in HLM can be modeled using these two vectors combined:

$$h_{-t} = \text{concatenate}(\vec{h}_{t-1}, \overleftarrow{h}_{t+1})$$

along with a succeeding output module, and a specific loss function compatible with negative log likelihood in the original.

Due to heterogeneous input of a test question, we have to model each input kind separately. For words, the output module would be a fully-connected layer  $\text{out}_w$  with parameters  $\theta_{ow}$ , and the loss function would be cross entropy. The output module takes  $h_{-t}$  as input, and generates a vector of vocabulary size, which, after Softmax, models the probability of each word at position  $t$ . For images, the output module is a fully-connected layer  $\text{out}_i(\cdot; \theta_{oi})$ , followed by the image decoder  $\tilde{emb}(\cdot; \theta_{id})$  described before, and the loss function is mean-squared-error (MSE) loss. Next, the output module of side information is also a fully-connected layer  $\text{out}_m(\cdot; \theta_{om})$ , followed by the image decoder  $\tilde{emb}(\cdot; \theta_{md})$ . The loss function of this kind of input is also cross entropy. Therefore, applying these output modules and loss functions above, the HLM loss at each position  $t$  would become:

$$l_t = \begin{cases} \text{CrossEntropy}(\text{out}_w(h_{-t}), x_t), & \text{if } x_t \text{ is a word;} \\ \text{MSE}(\tilde{emb}_i(\text{out}_i(h_{-t})), x_t), & \text{if } x_t \text{ is an image;} \\ \text{MSE}(\tilde{emb}_m(\text{out}_m(h_{-t})), x_t), & \text{if } x_t \text{ is side information.} \end{cases}$$

Therefore, the objective of holed language model in the low level for question  $q$  is

$$\mathcal{L}_{\text{low}} = \sum_{t=0}^T l_t.$$

**3.3.3 Domain-Oriented Objective.** The low level HLM loss only helps the model to learn linguistic features such as relation and context. However, domain logic and knowledge is still ignored. Take a look back at Figure 1, we can see that the relation between content and options contains much domain specific logic and knowledge. In order to also include such information in final representation, in this section, we designed a domain-oriented objective for high-level pre-training. We use the natural domain guidance of a test question: its answer and options, to help train QuesNet representation. For questions with one correct answer and other false options, we setup a pre-training task for QuesNet that, given an option, the model should output whether it is the correct answer. More specifically, we encode the option with a typical text encoder  $\text{enc}(\cdot; \theta_{\text{opt}})$  and get the answer representation as  $v_{\text{opt}} = \text{enc}(\text{opt})$  where  $\text{opt}$  represents the option. Then we model the probability of the option being the correct answer as:

$$P(\text{opt}|q) = \text{sigmoid}\left(D(v_q^{(s)}, v_{\text{opt}}; \theta_D)\right),$$

where  $v_q^{(s)}$  is the sentence representation for  $q$  generated by QuesNet, and  $D(\cdot, \theta_D)$  is a fully-connected neural network with output of 1 dimension. Therefore, the domain-oriented objective in the high level for question  $q$  is

$$\mathcal{L}_{\text{high}} = -\log P(\text{opt}|q).$$

With both low and high level objective, the pre-training process can be conducted on a large question bank with heterogeneous questions. With embedding modules' weights initialized and pre-trained separately, we can now apply stochastic gradient descent algorithm to optimize our hierarchical pre-training objective:

$$\mathcal{L} = \mathcal{L}_{\text{low}} + \mathcal{L}_{\text{high}}.$$

After pre-training, QuesNet question representation should be able to capture both low-level linguistic features and high-level domain logic and knowledge, and transfer the understanding of questions to downstream tasks in the area of education.

### 3.4 Fine-tuning

Downstream tasks in the area of education are often rather complicated. Taking knowledge mapping for example, as in the research by Yang et al. [31], the authors use a fine-grained model on this multi-label problem, which requires representation for each input content. Another example is score prediction. In the paper of Su et al. [27], each exercise (test question) is represented as a single vector and then serves as the input of a sequence model.

As we can see, different tasks require different question representations. To apply QuesNet representation to a specific task, we just have to provide the required representation to replace the equivalent part of the downstream model, which minimizes the cost of model modification. Moreover, on each downstream task, only some fine-tuning of QuesNet is needed, which leads to faster training speed and better results.

In summary, QuesNet has the following advantages for question understanding. First, it provides a unified and universally applicable representation for heterogeneous questions. Second, it is able to learn not only low-level linguistic features such as relation and context, but also high-level domain logic and knowledge. Third, it is easy to apply to downstream tasks and do fine-tuning. In next section, we will conduct extensive experiments to further demonstrate these advantages.

## 4 EXPERIMENTS

In this section, we conduct extensive experiments with QuesNet representation on three typical tasks in the area of education to demonstrate the effectiveness of our representation method.

### 4.1 Experimental Setup

**4.1.1 Dataset.** The dataset we used, along with the large question corpus, are supplied by iFLYTEK Co., Ltd., from their online education system called Zhixue<sup>1</sup>. All the data are collected from high school math tests and exams. Some important statistics are shown in Table 1 and Figure 4. The dataset is clearly heterogeneous, as shown in the table. About 25% of questions contain image content, and about 72% of questions contain side information. To clarify, side information used in knowledge mapping task is some other question meta data (its grade, the amount of which is shown in Table 1). In all other tasks, knowledge concepts are used as side information. Ignoring the heterogeneous information would definitely cause a downgrade in question understanding. We also observed that questions contain about 60 words in average, but the information contained is much more, as there are plenty of formulas represented as  $\LaTeX$  in question text.

**4.1.2 Evaluation Tasks.** We pick three typical tasks related with test questions in the area of education, namely: knowledge mapping, difficulty estimation, and student performance prediction. The unlabeled question corpus contains around 0.6 million questions. All of the questions in corpus are later used for pre-training each comparison models. For traditional text representation models, image and side information inputs are omitted.

The main objective for knowledge mapping task is to map a given question to its corresponding knowledge [24]. This is a multi-label task, where about 13,000 questions are labeled (only 1.98% of the whole unlabeled question dataset). To show how a representation method alleviate this scarce label problem and how it performs on this task, we choose a state-of-the-art knowledge mapping model, and replace its question representation part with each representation model we want to compare. After fine-tuning, we use some of the mostly used evaluation metrics for multi-label problem including accuracy (ACC), precision, recall, and F-1 score. Details of these metrics can be found in Piech et al. [24], Yang et al. [31].

The second task, namely difficulty estimation, is a high-level regression task to estimation the difficulty of a question. The scarce problem is even worse, in that merely 0.37% of the questions have been labeled. Meanwhile, the task needs more domain logic and knowledge as a guidance to get better performance, as estimation of the difficulty of an exercise requires a deeper understanding of the

<sup>1</sup><http://www.zhixue.com>






Figure 4: Distribution of question inputs and labels.

Table 1: Statistics of datasets.

	All	KM	DE	SP
#Questions	675,264	13,372	2,465	15,045
#Questions with image	165,859	3,318	299	2,952
#Questions with meta	488,352	8,030	1,896	5,948
#Questions with option	242,960	4840	1,389	4,364
Avg. Words per question	59.10	58.43	60.25	51.94
#Students	-	-	-	50,945
#Records	-	-	-	3,358,111
Label sparsity	-	1.98%	0.37%	2.22%

question. The dataset on this task consists of only 2400 questions. The evaluation metrics, following Huang et al. [13], includes Mean-Absolute-Error (MAE), Root-Mean-Squared-Error (RMSE), Degree of Agreement (DOA), and Pearson Correlation Coefficient (PCC).

The score prediction task, on the other hand, is a much more complicated domain task, where the main goal is to sequentially predict how well a student performs on each test question they exercises on [27]. While student record is of large scale, the amount of questions for this task is still quite limited, only about 2.22%. For better modeling student exercising sequence, some of the state-of-the-art model incorporate question content combining a question representation into the modeling. We replace this module with our comparison methods, and evaluate the performance using MAE, RMSE (mentioned earlier), accuracy (ACC) and Area Under the Curve (AUC), as used in many studies [27, 34].

**4.1.3 QuesNet Setup.**<sup>2</sup> The embedding modules all output vectors of size 128. The image embedding module and related decoder are implemented as 4 layer convolutional neural networks and transposed convolutional neural networks, respectively. The size of feature maps at each layer is 16, 32, 32, 64. The side information embedding module and related decoder are setup as two-layer fully-connected neural networks. The size of the hidden layer is set to 256. For the main part of QuesNet, we use  $L = 4$  layer of bi-LSTM and 1 layer of self-attention. Sizes of hidden states in these modules are set to 256. To prevent overfitting, we also introduce dropout layers [26] between each layer, with dropout probability as 0.2.

Before any pre-training, the layers are first initialized with Xavier initialization strategy [8]. Then at pre-training process, parameters are updated by Adam optimization algorithm [15]. We pre-train our model on question corpus long enough so that the pre-training

<sup>2</sup>The code is available at <https://github.com/yxonic/QuesNet>.

Table 2: Comparison methods.

Method	Text	Image	Meta	Low level	High level
Original	✓			-	-
ELMo	✓			-	-
BERT	✓			-	-
H-BERT	✓	✓	✓	-	-
QN-T	✓			✓	✓
QN-I		✓		✓	✓
QN-M			✓	✓	✓
QN-TI	✓	✓		✓	✓
QN-TM	✓		✓	✓	✓
QN-IM		✓	✓	✓	✓
QN (no pre)	✓	✓	✓	✓	✓
QN-L	✓	✓	✓	✓	
QN-H	✓	✓	✓		✓
QuesNet	✓	✓	✓	✓	✓

loss converges. For optimizers in each task, we follow the setups described in the corresponding paper.

**4.1.4 Comparison Methods.** We compare QuesNet with several representation methods. All these methods are able to generate question representation, and then be applied to the three evaluation tasks mentioned above. Specifically, these methods are:

- **Original** refers to original supervised learning methods on each of the evaluation task. We choose state-of-the-art methods for difficulty estimation [13], knowledge mapping [31] and score prediction [27], without any forms of pre-training.
- **ELMo** is a LSTM based feature extraction method with bidirectional language model as pre-training strategy [23]. It is only capable of text representation, so we omit other types of input when using this method.
- **BERT** is a state-of-the-art pre-training method featuring Transformer structure and masked language model [5]. It is only capable of text representation, so we also omit other types of input.
- **H-BERT** is a modified version of BERT, which allows it to process heterogeneous input. We use the same input embedding modules as QuesNet, and set the embedded vectors as the input of text-oriented BERT.

All comparison methods are listed in Table 2. For a fair comparison, all these methods are adjusted to contain approximately same amount of layers and parameters, and all of them are tuned to have the best performance. All models are implemented by PyTorch, and trained on a cluster of Linux servers with Tesla K20m GPUs.

## 4.2 Experimental Results

The comparison results on each of three tasks with four different models including QuesNet are shown in Table 3. We can easily see that pre-trained methods is able to boost the performance on each task, among which QuesNet has the best performance on almost all metrics, no matter what the size of each task is. This proves that QuesNet gains a better understanding of questions and is transferred more efficiently from large unlabeled corpus to small

**Table 3: Performance of comparison methods on different tasks.**

Methods	Knowledge mapping				Difficulty estimation				Student performance prediction			
	ACC	Precision	Recall	F-1	MAE	RMSE	DOA	PCC	MAE	RMSE	ACC	AUC
Original	0.5744	0.4147	0.7872	0.5432	0.2200	0.2665	0.6064	0.3050	0.4245	0.4589	0.7459	0.5400
ELMo	0.6942	0.7960	0.7685	0.7820	0.2250	0.2655	0.5561	0.4299	0.3569	<b>0.4361</b>	0.7866	0.5773
BERT	0.6224	0.7326	0.6711	0.7005	0.2265	0.2975	0.6258	0.3600	0.4009	0.4630	0.7390	0.5279
H-BERT	0.6261	0.7608	0.6911	0.7243	0.2097	0.2698	<b>0.6597</b>	0.3713	0.3925	0.4528	0.7784	0.5838
QuesNet	<b>0.7749</b>	<b>0.8659</b>	<b>0.8075</b>	<b>0.8357</b>	<b>0.2029</b>	<b>0.2530</b>	0.6137	<b>0.4499</b>	<b>0.3445</b>	0.4403	<b>0.7999</b>	<b>0.6354</b>

**Table 4: Ablation experiments.**

Methods	Knowledge mapping				Difficulty estimation				Student performance prediction			
	ACC	Precision	Recall	F-1	MAE	RMSE	DOA	PCC	MAE	RMSE	ACC	AUC
QN-T	0.7050	0.8264	0.7436	0.7829	0.2166	0.2733	0.6123	0.3040	0.4488	0.4713	0.7454	0.6052
QN-I	0.1136	0.2232	0.3195	0.2628	0.2265	0.2713	0.5961	0.2178	0.4711	0.4899	0.7400	0.5921
QN-M	0.0355	0.1396	0.2853	0.1875	0.2251	0.2737	0.5549	0.2205	0.4719	0.4908	0.7410	0.5502
QN-TI	0.7207	0.8307	0.7595	0.7935	0.2110	0.2647	0.6029	0.3333	0.4279	0.4678	0.7523	0.6221
QN-TM	0.7196	0.8428	0.7523	0.7950	0.2114	0.2664	0.6151	0.3315	0.4353	0.4803	0.7456	0.6156
QN-IM	0.1428	0.2323	0.2818	0.2547	0.2277	0.2707	0.5766	0.2279	0.4710	0.4906	0.7411	0.5513
QN (no pre)	0.5659	0.6816	0.7091	0.6951	0.2225	0.2657	0.5750	0.3087	0.4349	0.4759	0.7488	0.5891
QN-L	0.7185	0.8352	0.7457	0.7879	0.2193	0.2630	0.5721	0.3359	0.3843	0.4561	0.7747	0.6237
QN-H	0.6807	0.8052	0.7271	0.7642	0.2161	0.2665	<b>0.6291</b>	0.3328	0.3946	0.4475	0.7814	0.6058
QuesNet	<b>0.7749</b>	<b>0.8659</b>	<b>0.8075</b>	<b>0.8357</b>	<b>0.2029</b>	<b>0.2530</b>	0.6137	<b>0.4499</b>	<b>0.3445</b>	<b>0.4403</b>	<b>0.7999</b>	<b>0.6354</b>

labeled datasets. However, there are more to be explained in this table. First, models that support heterogeneous input have superior results over similar structures without heterogeneous input, which proves that it is crucial to handle heterogeneous inputs when understanding questions. Second, as all methods are adjusted to contain similar amount of parameters, QuesNet turns out to be the most efficient one. Third, the result of the Transformer based methods is slightly lower than other LSTM based methods. This is probably because of the low sample efficiency of masked language model pre-training strategy used in BERT, while bi-directional language model in ELMo performs better, and our novel holed language model used in QuesNet outperforms both of them.

### 4.3 Ablation Experiments

In this section, we conduct some ablation experiments to further show how each part of our method affect final results. In Table 4, there are eight variations of QuesNet, each of which takes out one or more opponents from the full QuesNet. Specifically: *QN-T*, *QN-I*, *QN-M* refer to QuesNet with only text, image or side information is used in pre-training process, respectively. *QN-TI*, *QN-TM*, *QN-IM* refer to combinations of different kinds of input, i.e. text and image, text and side information, and image and side information, respectively. And finally, *QN-L* refers to QuesNet that only includes low-lever pre-training objective (holed language model), and *QN-H* refers to QuesNet that only includes high-level domain objective.

The result in Table 4 indeed shows many interesting conclusions. First, the more information a model incorporates, the better the performance, which agrees with the intuition. Second, if we focus on comparison between *QN-H* and *QN-L*, we will notice that they gain

different effects on different tasks. On tasks more focusing on lower level features like knowledge mapping, *QN-L* outperforms *QN-H* slightly, while on other more domain-oriented high-level tasks (difficulty estimation and student performance prediction), *QN-H* performs a little better. This clearly demonstrates different aspects the two objectives focus on, and with full QuesNet outperforms both *QN-L* and *QN-H*. we know that QuesNet is able to take account of both aspects of these objectives, and build an understanding at both low linguistic level and high logic level. Third, we notice that *QN-T* has better performance than *QN-I*, which is better than *QN-M*, indicating that text carries most information in a question, then images, then side information. But as omitting either of the input kind will cause a performance loss, all of the information in heterogeneous inputs is essential to a good question understanding.

### 4.4 Discussion

From the above experiments, it is clear that QuesNet can effectively gain an understanding of test questions. First, it can well aggregate information from heterogeneous inputs. The model is able to generate unified representation for each question with different forms, and leverage information in all kinds of input. Second, with low-level objective capturing linguistic features, and high-level objective learning domain logic and knowledge, QuesNet is also able to gain both low-level and high-level understanding of test questions. Impressive performance of QuesNet can be seen among all three typical educational applications, which highlights the usability and superiority of QuesNet in the area of education.

There are still some directions for future studies. First, we may work on some domain specific model architectures to model logic



among questions in a more fine-grained way. Second, the understanding of QuesNet on test questions is not comprehensible, and in the future, we would like to work on comprehension and explanation, to generate a more convincing representation. Second, the general idea of our method is applicable in more heterogeneous scenarios, and we would like to further explore the possibilities of our work on other heterogeneous data and tasks.

## 5 CONCLUSION

In this paper, we presented a unified representation for heterogeneous test questions, namely QuesNet. Specifically, we first designed a heterogeneous modeling architecture to represent heterogeneous input as a unified form. Then we proposed a novel hierarchical pre-training framework, with holed language model (HLM) for pre-training low-level linguistic features, and a domain-oriented objective for learning high-level domain logic and knowledge. With extensive experiments on three typical downstream tasks in education from the low-level knowledge mapping task, to the domain-related difficulty estimation task, then to complex high-level student performance prediction task, we proved that QuesNet is more capable of question understanding and transferring, capturing both low-level linguistic features and high-level domain logic and knowledge. We hope this work builds a solid basis for question related tasks in the area of education, and help boost more applications in this field.

## ACKNOWLEDGEMENTS

This research was partially supported by grants from the National Key Research and Development Program of China (No. 2016YFB1000904) and the National Natural Science Foundation of China (Grants No. 61727809, U1605251, 61672483). Qi Liu gratefully acknowledges the support of the Young Elite Scientist Sponsorship Program of CAST and the Youth Innovation Promotion Association of CAS (No. 2014299).

## REFERENCES

- Ashton Anderson, Daniel Huttenlocher, Jon Kleinberg, and Jure Leskovec. 2014. Engaging with massive online courses. In *Proceedings of the 23rd international conference on World wide web*. ACM, 687–698.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450* (2016).
- C Boopathiraj and K Chellamani. 2013. Analysis of test items on difficulty level and discrimination index in the test for research in education. *International journal of social science & interdisciplinary research* 2, 2 (2013), 189–193.
- Michel C Desmarais, Behzad Beheshti, and Rhouma Naceur. 2012. Item to skills mapping: deriving a conjunctive q-matrix from data. In *International Conference on Intelligent Tutoring Systems*. Springer, 454–463.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- David E Douglas and Glen Van Der Vyver. 2004. Effectiveness of e-learning course materials for learning database management systems: An experimental investigation. *Journal of Computer Information Systems* 44, 4 (2004), 41–48.
- Huizhong Duan, Yunbo Cao, Chin-Yew Lin, and Yong Yu. 2008. Searching questions by identifying question topic and question focus. *Proceedings of ACL-08: HLT* (2008), 156–164.
- Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*. 249–256.
- Arthur C Graesser, Zhiqiang Cai, Max M Louwerse, and Frances Daniel. 2006. Question Understanding Aid (QUAID) a web facility that tests question comprehensibility. *Public Opinion Quarterly* 70, 1 (2006), 3–22.
- Ulf Hermjakob. 2001. Parsing and question classification for question answering. In *Proceedings of the workshop on Open-domain question answering-Volume 12*. Association for Computational Linguistics, 1–6.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- Jeremy Howard and Sebastian Ruder. 2018. Universal language model fine-tuning for text classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Vol. 1. 328–339.
- Zhenya Huang, Qi Liu, Enhong Chen, Hongke Zhao, Mingyong Gao, Si Wei, Yu Su, and Guoping Hu. 2017. Question Difficulty Prediction for READING Problems in Standard Tests. In *AAAI*. 1352–1359.
- Zhiheng Huang, Wei Xu, and Kai Yu. 2015. Bidirectional LSTM-CRF models for sequence tagging. *arXiv preprint arXiv:1508.01991* (2015).
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- George D Kuh, Jillian Kinzie, Jennifer A Buckley, Brian K Bridges, and John C Hayek. 2011. *Piecing together the student success puzzle: research, propositions, and recommendations: ASHE Higher Education Report*. Vol. 116. John Wiley & Sons.
- Xuezhe Ma and Eduard Hovy. 2016. End-to-end sequence labeling via bidirectional lstm-cnns-crf. *arXiv preprint arXiv:1603.01354* (2016).
- Md Anwar Hossain Masud and Xiaodi Huang. 2012. An e-learning system architecture based on cloud computing. *system* 10, 11 (2012), 255–259.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781* (2013).
- Michael G Moore and Greg Kearsley. 2011. *Distance education: A systems view of online learning*. Cengage Learning.
- Jiquan Ngiam, Aditya Khosla, Mingyu Kim, Juhan Nam, Honglak Lee, and Andrew Y Ng. 2011. Multimodal deep learning. In *Proceedings of the 28th international conference on machine learning (ICML-11)*. 689–696.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*. 1532–1543.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365* (2018).
- Chris Piech, Jonathan Bassen, Jonathan Huang, Surya Ganguli, Mehran Sahami, Leonidas J Guibas, and Jascha Sohl-Dickstein. 2015. Deep knowledge tracing. In *Advances in Neural Information Processing Systems*. 505–513.
- Norbert Ed Schwarz and Seymour Ed Sudman. 1996. *Answering questions: Methodology for determining cognitive and communicative processes in survey research*. Jossey-Bass.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research* 15, 1 (2014), 1929–1958.
- Yu Su, Qingwen Liu, Qi Liu, Zhenya Huang, Yu Yin, Enhong Chen, Chris Ding, Si Wei, and Guoping Hu. 2018. Exercise-Enhanced Sequential Modeling for Student Performance Prediction. (2018).
- Martin Sundermeyer, Ralf Schlüter, and Hermann Ney. 2012. LSTM neural networks for language modeling. In *Thirteenth annual conference of the international speech communication association*.
- Ming Tan, Cicero dos Santos, Bing Xiang, and Bowen Zhou. 2015. LSTM-based deep learning models for non-factoid answer selection. *arXiv preprint arXiv:1511.04108* (2015).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*. 5998–6008.
- Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. Hierarchical attention networks for document classification. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 1480–1489.
- Yu Yin, Zhenya Huang, Enhong Chen, Qi Liu, Fuzheng Zhang, Xing Xie, and Guoping Hu. 2018. Transcribing Content from Structural Images with Spotlight Mechanism. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. ACM, 2643–2652.
- Dell Zhang and Wee Sun Lee. 2003. Question classification using support vector machines. In *Proceedings of the 26th annual international ACM SIGIR conference on Research and development in information retrieval*. ACM, 26–32.
- Jiani Zhang, Xingjian Shi, Irwin King, and Dit-Yan Yeung. 2017. Dynamic key-value memory networks for knowledge tracing. In *Proceedings of the 26th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 765–774.
- Liang Zhang, Keli Xiao, Hengshu Zhu, Chuanren Liu, Jingyuan Yang, and Bo Jin. 2018. CADEN: A Context-Aware Deep Embedding Network for Financial Opinions Mining. In *2018 IEEE International Conference on Data Mining (ICDM)*. IEEE, 757–766.