

End-to-End Evaluation of a Spoken Dialogue System for Learning Basic Mathematics

Eda Okur

Intel Labs, USA
eda.okur@intel.com

Saurav Sahay

Intel Labs, USA
saurav.sahay@intel.com

Roddy Fuentes Alba

Intel Labs, Mexico
roddy.fuentes.alba@intel.com

Lama Nachman

Intel Labs, USA
lama.nachman@intel.com

Abstract

The advances in language-based Artificial Intelligence (AI) technologies applied to build educational applications can present AI for social-good opportunities with a broader positive impact. Across many disciplines, enhancing the quality of mathematics education is crucial in building critical thinking and problem-solving skills at younger ages. Conversational AI systems have started maturing to a point where they could play a significant role in helping students learn fundamental math concepts. This work presents a task-oriented Spoken Dialogue System (SDS) built to support play-based learning of basic math concepts for early childhood education. The system has been evaluated via real-world deployments at school while the students are practicing early math concepts with multimodal interactions. We discuss our efforts to improve the SDS pipeline built for math learning, for which we explore utilizing MathBERT (Shen et al., 2021b) representations for potential enhancement to the Natural Language Understanding (NLU) module. We perform an end-to-end evaluation using real-world deployment outputs from the Automatic Speech Recognition (ASR), Intent Recognition, and Dialogue Manager (DM) components to understand how error propagation affects the overall performance in real-world scenarios.

1 Introduction

Following the advances in Artificial Intelligence (AI) research, building innovative applications to support education can present exciting opportunities with a positive and broader social impact. The United Nations (UN) Sustainable Development Goals¹ (SDGs) (Desa et al., 2016) represent an urgent call for action to help address critical global problems, where education is among the top five of these development areas (i.e., poverty, hunger, health, education, gender equality). Across many

¹<https://sdgs.un.org/goals>

disciplines, improving the quality of mathematics education is crucial in building critical thinking and problem-solving skills at younger ages, which is a fundamental component of comprehensive and successful STEM education (i.e., science, technology, engineering, and mathematics). Language-based AI systems are starting to mature to a point where they could play a significant role in helping students learn and practice mathematical concepts. Despite its importance, applied Natural Language Processing (NLP) technologies for enhancing mathematics education still remain a highly under-explored area of research.

This study presents a task-oriented Spoken Dialogue System (SDS) developed to facilitate play-based learning of basic mathematical concepts for early childhood education. The system has been developed in the lab and evaluated via real-world deployments at school while the students are learning and practicing basic math concepts with multimodal interactions. These fundamental early math concepts and basic operations include constructing numbers using ones and tens, counting, addition, subtraction, measurement of length and size, etc. The multimodal interactions involve speech-based interactions to answer early math-related or game-related questions, counting and placing learning-specific tangible objects (i.e., manipulatives) in a visually observed space, touch-based interactions with the 1-to-100 number grid projected on the wall, to name a few.

This work discusses our efforts to improve the modular SDS pipeline built for game-based math learning and perform an end-to-end evaluation with various SDS components using real-world deployment outputs. The main SDS module we focus on investigating is Natural Language Understanding (NLU). The NLU arguably is the most critical component of goal-oriented dialogue systems that enables efficient communication between humans and intelligent conversational agents via application-

specific comprehensive sub-tasks. Intent Recognition (IR) is at the heart of these NLU tasks, where the goal is to identify users’ objectives from the input text and determine their intentions.

The representation of human language is a crucial factor determining the success of conversational agents, especially in real-world applications. In the math learning domain, this language representation gains further significance due to the quite peculiar nature of mathematical language. With that motivation, we explored employing MathBERT (Shen et al., 2021b) representations for potential enhancement to the NLU module of our play-based math-learning system. MathBERT is a recently proposed language model created by pre-training the well-known BERT-base model (Devlin et al., 2019) on a large mathematical corpus. We compared the NLU results obtained by simply using BERT (Devlin et al., 2019) and ConveRT (Henderson et al., 2020) representations versus the new MathBERT representations. We further investigated the two variations of MathBERT models, one pre-trained with mathematics-specific vocabulary and the other with BERT-base vocabulary, to see their effects on our domain-specific math-learning NLU task.

As most of the application-specific and modular SDS pipelines do, our task-oriented SDS contains particular building blocks or modules for Automatic Speech Recognition (ASR), Natural Language Understanding (NLU), multimodal Dialogue Management (DM), Natural Language Generation (NLG), and Text-to-Speech (TTS). At its current stage of research & development, for practical reasons, we employ template-based responses at the NLG module with off-the-shelf TTS. Thus, this study emphasizes more on speech recognition (ASR), intent understanding (NLU), and response selection (DM) tasks for a conversational agent that supports elementary math learning. For a complete end-to-end evaluation of our task-oriented SDS, we evaluated the ASR, NLU, and DM components to understand how error propagation affects the overall performance in real-world scenarios.

2 Related Work

2.1 NLP for Mathematics Education

Exploring the advancements in AI systems for social good and positive impact in the education domain, specifically to amplify students’ learning experiences, has recently gained increasing

interest from the research community (D’Mello and Graesser, 2013; Chassignol et al., 2018; Jia et al., 2020; Baker, 2021; Zhai et al., 2021). Intelligent and interactive play-based learning systems have shown remarkable benefits for teaching mathematical concepts in smart and collaborative spaces (Lester et al., 2013; Pires et al., 2019; Richey et al., 2021; Sun et al., 2021). For early childhood education, a recent study by Skene et al. (2022) has showcased that game-based learning environments can offer significant advantages over traditional learning approaches while practicing fundamental math concepts, especially for younger learners.

Harnessing NLP technologies to construct innovative applications for education is gaining popularity as an emerging area of research having various examples in the last decade (Blanchard et al., 2015; Lende and Raghuvanshi, 2016; Taghipour and Ng, 2016; Raamadhurai et al., 2019; Cahill et al., 2020; Chan et al., 2021; Dutta et al., 2022). Within those efforts, exploring conversational agents for intelligent tutoring systems and smart education applications is a glaring sub-field of NLP in education, having several research studies tackling the problem from different angles (Graesser et al., 2004; Kerry et al., 2009; Winkler and Söllner, 2018; Palasundram et al., 2019; Winkler et al., 2020; Datta et al., 2020; Zhang et al., 2021).

Furthermore, there have been recent attempts to bridge the gap between general AI research and mathematics education (Davila and Zanibbi, 2017; Jiang et al., 2018; Mansouri et al., 2019; Yuan et al., 2020; Huang et al., 2021; Kumar and Rajagopal, 2021). To further narrow our attention to language-based technologies applied to mathematics education, relatively few recent studies exist (Shen et al., 2021a; Suresh et al., 2022; Logionova and Benoit, 2022) which explore transfer learning to improve language representation for math-related tasks (Peng et al., 2021; Shen et al., 2021b). Among these, MathBERT (Shen et al., 2021b) has been built specifically for challenging downstream NLP tasks in math education (e.g., knowledge component prediction, auto-grading open-ended question-answering, and knowledge tracing). It is indeed a mathematics-customized BERT model (Devlin et al., 2019). MathBERT representations are created by pre-training the BERT-base model on large mathematical corpora, including pre-kindergarten, to high-school and college graduate-level mathematical text.

2.2 Dialogue Systems and NLU

For interactive early math learning applications, as we aim to build a spoken dialogue system for kids, we will briefly discuss the existing dialogue system technologies and language understanding approaches in a more generic context here.

Conversational agents or dialogue systems are mainly categorized as either open-ended or task-oriented (Chen et al., 2017). The open-ended dialogue systems or chatbots allow generic conversations such as chit-chat (Serban et al., 2016; Jurafsky and Martin, 2018). On the other hand, task-oriented conversational AI systems are designed to accomplish specific tasks and handle goal-oriented conversations (Serban et al., 2018; Mehri et al., 2020). With the advances of deep learning-based language technologies, improved access to high computing power, and increased availability of large datasets; the end-to-end trained dialogue systems can achieve encouraging results for both open-ended (Serban et al., 2016; Dodge et al., 2016) and task-oriented (Wen et al., 2017; Bordes et al., 2017; Ham et al., 2020) applications. Dialogue Managers (DM) of task-oriented conversational AI systems are mostly sequential decision-makers. At that step, learning the optimal dialogue policies can be achieved via reinforcement learning (RL) from an excessive number of user interactions (Zhao and Eskenazi, 2016; Shah et al., 2016; Cuayáhuil, 2017; Dhingra et al., 2017; Liu et al., 2017; Su et al., 2017). However, building RL-based dialogue systems with highly limited user interaction data is immensely challenging. For this reason, supervised learning methods with traditional pipeline-based modular dialogue systems are still widely accepted when training data is initially limited to bootstrap the task-oriented SDS for further data collection (Budzianowski et al., 2018). For implicit dialogue context management, statistical and neural network-based dialogue system frameworks and toolkits (Bocklisch et al., 2017; Ultes et al., 2017; Burtsev et al., 2018) are employed popularly in the research communities and industrial applications.

The NLU module of a dialogue system pipeline processes the user utterances as input text and usually predicts the user’s intents or dialogue acts. For sequence learning tasks of Intent Recognition and Slot Filling (Mesnil et al., 2015; Hakkani-Tür et al., 2016), joint training of Intent Classification and Entity Recognition models have been explored widely (Zhang and Wang, 2016; Liu and

Lane, 2016; Goo et al., 2018; Varghese et al., 2020). Many hierarchical multi-task learning architectures have been proposed for these joint NLU methods (Zhou et al., 2016; Gu et al., 2017; Wen et al., 2018; Okur et al., 2019; Vanzo et al., 2019). Vaswani et al. (2017) proposed the Transformer as a game-changing neural network architecture based entirely on attention mechanisms (Bahdanau et al., 2015). Right after the Transformers, Bidirectional Encoder Representations from Transformers (BERT) (Devlin et al., 2019) has been proposed. BERT has become one of the most significant breakthroughs in language representations research and has shown strong performance in several NLP tasks, including NLU. Lately, Bunk et al. (2020) introduced the Dual Intent and Entity Transformer (DIET) model as a lightweight multi-task architecture. DIET has been shown to outperform fine-tuning the BERT model for predicting intents and entities on a complex multi-domain NLU-Benchmark dataset (Liu et al., 2021). For efficient language representation learning, Henderson et al. (2020) recently proposed the Conversational Representations from Transformers (ConveRT) model, which is another lightweight approach to obtain pre-trained embeddings as sentence representations successfully used in several conversational AI tasks.

3 Language Understanding Methods

For the NLU module within our early math-learning dialogue system pipeline, we have examined numerous options for Intent Recognition and built our NLU models on top of the Rasa open-source framework (Bocklisch et al., 2017).

3.1 Baseline: TF+BERT

The previous baseline Intent Recognition architecture available in the Rasa platform was based on supervised embeddings as part of the Rasa NLU (Bocklisch et al., 2017). It was an embedding-based text classifier that embedded user utterances and intent labels into the same vector space. This former baseline architecture was inspired by the StarSpace algorithm (Wu et al., 2018), where the supervised embeddings were trained by maximizing the similarity between intents and utterances. Sahay et al. (2019) enriched this embedding-based former baseline Rasa Intent Classifier by incorporating additional features and adapting alternative network architectures. To be more specific, they adapted the Transformer ar-

chitecture (Vaswani et al., 2017) and employed pre-trained BERT embeddings (Devlin et al., 2019) using the bert-base-uncased² model for Intent Recognition. We treat this simple and initial approach as our baseline NLU model in this study. We will refer to this method as TF+BERT in our experiments.

3.2 DIET+BERT

Next, we explored the potential benefits of adopting the recent DIET architecture (Bunk et al., 2020) for the Intent Recognition task in the basic math-learning domain. DIET is a transformer-based multi-task architecture for joint Intent Classification and Entity Recognition. The architecture includes a two-layer transformer shared for both NLU tasks. A sequence of entity labels is predicted with a Conditional Random Field (CRF) (Lafferty et al., 2001) tagging layer on top of the transformer output sequence corresponding to the input sentences treated as a sequence of tokens. For the intent labels, the transformer output for the classification token and the intent labels are embedded into the same semantic vector space. The dot-product loss is employed to maximize the similarity with the target label and minimize similarities with the negative samples. DIET architecture enables the incorporation of the pre-trained word and sentence embeddings from language models as dense features, with the flexibility to combine these with token-level one-hot and multi-hot encodings of character n-grams as sparse features. These sparse features are passed through a fully-connected layer with shared weights across all sequence steps. The output of the fully-connected layer is concatenated with the dense features from the pre-trained models. The high flexibility of this architecture allowed us to use any pre-trained embeddings as dense features in DIET, such as BERT (Devlin et al., 2019), ConveRT (Henderson et al., 2020), and MathBERT (Shen et al., 2021b). To investigate the net benefits of DIET architecture, we adopted DIET with off-the-shelf pre-trained BERT embeddings using the bert-base-uncased model (Devlin et al., 2019) and compared that against our baseline TF+BERT model. This approach (i.e., combining out-of-the-box BERT representations with the DIET classifier) will be referred to as DIET+BERT in the experiments³.

²<https://huggingface.co/bert-base-uncased>

³Refer to Bunk et al. (2020) for hyper-parameters, computational costs, and hardware specifications.

3.3 DIET+ConveRT

Conversational Representations from Transformers (ConveRT) (Henderson et al., 2020) is a promising architecture recently proposed to learn pre-trained representations that are well-suited for conversational AI applications, especially for the real-world Intent Classification tasks. ConveRT is a transformer-based dual-encoder network leveraging quantization and sub-word level parameterization, where the pre-trained representation outputs from its sentence encoder can be utilized for the conversational Intent Classification tasks. DIET and ConveRT are both lightweight architectures with faster and more efficient training capabilities than their counterparts. When incorporating the ConveRT embeddings within the DIET classifier, the initial embeddings for the classification tokens are set as the input sentence encoding obtained from the ConveRT model. That enables exploiting extra contextual information from the complete sentence on top of the word embeddings. For these reasons, we adopted the DIET architecture and utilized pre-trained ConveRT embeddings to potentially improve the Intent Recognition performances on our domain-specific early math-learning datasets. We will call this approach DIET+ConveRT in our experiments.

3.4 DIET+MathBERT

Finally, a pre-trained language model named MathBERT (Shen et al., 2021b) has been presented lately for downstream NLP tasks in the domain of mathematics education. MathBERT is a BERT-like language representation model further pre-trained from the bert-base-uncased model with dedicated mathematical corpora. Note that BERT is a general-purpose language model trained on a vast amount of unlabeled text corpus (i.e., Wikipedia and BookCorpus) with 3.3 billion words, which can be further pre-trained to obtain a new set of model weights for transfer learning. On the contrary, MathBERT is pre-trained on 100 million tokens of mathematical corpora, including instructional texts from books, curriculum, Massive Open Online Courses (MOOCs), and arXiv.org paper abstracts, covering all possible grade levels from pre-k to college graduate-level content. Although the scale of training corpora is much smaller than the BERT-base model, MathBERT can still have the potential to be more effective in math-related NLP tasks. That is hypothesized because mathe-

Statistics	Math-Game Data	
	Planting	Watering
# Distinct Intents	14	13
Total # Samples (Utterances)	1927	2115
# Math-related Samples	452	599
Min # Samples per Intent	22	25
Max # Samples per Intent	555	601
Avg # Samples per Intent	137.6	162.7
Min # Words per Sample	1	1
Max # Words per Sample	74	65
Avg # Words per Sample	5.26	4.95
# Unique Words (Vocab)	1314	1267
Total # Words	10141	10469

Table 1: KidSpace-POC NLU Dataset Statistics

mathematical language frequently uses domain-specific vocabulary and concepts that require better word representations within the math context. With that motivation, we utilized the publicly available MathBERT release ⁴, which includes the PyTorch and TensorFlow versions of MathBERT models and the tokenizers. We explored adopting the DIET architecture with pre-trained MathBERT embeddings to empower the NLU/Intent Recognition task in our dialogue system designed for teaching basic math concepts. Furthermore, we investigated the representations from the MathBERT-base model that is trained with the BERT-base vocabulary (i.e., *origVocab*) and compared that against the representations from MathBERT-custom model pre-trained with math-specific vocabulary (i.e., *mathVocab*). This custom *mathVocab* set is also released to reflect the specific nature of mathematical jargon and concepts used in math corpora. In our NLU experiments, we will refer to these two distinct approaches as DIET+MathBERT-base (*origVocab*) and DIET+MathBERT-custom (*mathVocab*).

4 Experimental Results

4.1 Math-Game Datasets

Our experiments are conducted on the NLU datasets of Kid Space Planting and Watering use cases (Anderson et al., 2018; Aslan et al., 2022), having utterances from gamified math learning experiences designed for early childhood education (i.e., 5-to-8 years old kids). The intelligent conversational agent should accurately understand these children’s utterances and provide appropriate feedback. The use cases include a specific flow of inter-

⁴<https://github.com/tbs17/MathBERT>

Statistics	Math-Game Data	
	Planting	Watering
# Distinct Intents	12	11
Total # Samples (Utterances)	2173	2122
# Math-related Samples	549	602
Min # Samples per Intent	4	6
Max # Samples per Intent	1005	1005
Avg # Samples per Intent	181.1	192.9
Min # Words per Sample	1	1
Max # Words per Sample	45	44
Avg # Words per Sample	4.80	4.48
# Unique Words (Vocab)	772	743
Total # Words	10433	9508

Table 2: KidSpace-Deployment NLU Dataset Statistics

active games facilitating elementary math learning. The FlowerPot (i.e., Planting) game builds on the math concepts of tens and ones, with the larger flower pots representing tens and smaller pots representing ones. The virtual character provides the number of flowers the children should plant, and when the children have placed the correct number of large and small pots against the wall, digital flowers appear. In the NumberGrid (i.e., Watering) game, the children are presented with basic math questions with clues. When the correct number is touched on the 1-to-100 number grid projected on the wall, water is virtually poured to water the flowers. The virtual character supports the kids with learning to construct numbers using the ‘tens’ and ‘ones’ digits, practicing simple counting, addition, and subtraction operations. These math-game datasets have a limited number of user utterances, which are annotated manually for intent types defined for each learning activity. Some of the intents are highly generic across learning activities (e.g., *affirm*, *deny*, *next-step*, *out-of-scope*, *good-bye*), whereas others are highly domain-dependent and game-specific (e.g., *intro-meadow*, *answer-flowers*, *answer-water*, *answer-valid*, *answer-invalid*) or math-learning/task-specific (e.g., *ask-number*, *counting*).

The NLU models are trained and validated on the initial proof-of-concept (POC) datasets (Sahay et al., 2021) to bootstrap the agents for real-world deployments. These POC datasets were curated manually to train the initial SDS models based on User Experience (UX) studies. The models are then validated on the UX sessions in the lab with five kids going through these early math-learning games. Table 1 shows the statistics

Model	Math-Game Data	
	Planting	Watering
TF+BERT (baseline)	90.50±0.25	92.43±0.32
DIET+BERT	94.00±0.38	96.39±0.14
DIET+ConveRT	95.88±0.42	97.69±0.11
DIET+MathBERT-base	90.40±0.16	93.56±0.26
DIET+MathBERT-custom	90.82±0.10	93.67±0.10

Table 3: NLU/Intent Recognition micro-avg F1-scores (%): TF+BERT (baseline), DIET+BERT, DIET+ConveRT, DIET+MathBERT-base (origVocab), and DIET+MathBERT-custom (mathVocab) models trained and validated on KidSpace-POC datasets.

of these KidSpace-POC NLU datasets. Planting and Watering game POC datasets have 1927 and 2115 user utterances, respectively. The deployment datasets were collected later from twelve kids, where the math-learning system was deployed in a classroom at school (Okur et al., 2022b). Table 2 shows the statistics of KidSpace-Deployment NLU datasets, where Planting and Watering deployment datasets have 2173 and 2122 user utterances, respectively. These deployment datasets are used only for testing purposes, where we train our NLU models on the POC datasets. For both in-the-lab and real-world datasets, the spoken user utterances and agent responses are transcribed manually at first. These transcriptions are annotated for the user intent and agent response types we defined for each math learning activity. These transcribed and annotated final utterances are analyzed and used in our experiments.

4.2 DIET Classifier with MathBERT Representations for NLU

Tables 3 and 4 present a summary of our NLU/Intent Recognition experimental results for the Planting and Watering math-game use cases at school, covering a series of task-oriented interactions for early math education. On the proof-of-concept (POC) datasets that we created for each math-game, we achieved above 95% F1-scores for Intent Recognition performances with our best-performing DIET+ConveRT NLU models (see Table 3). We trained and cross-validated these NLU models on math-game or activity-specific datasets having around 2K POC samples. When we later tested these models on real-world deployment data collected at school, we observed F1-score performance drops of around 7% with our best-performing DIET+ConveRT models (see Ta-

Model	Math-Game Data	
	Planting	Watering
TF+BERT (baseline)	85.08±0.49	90.06±0.56
DIET+BERT	87.03±0.30	89.63±0.62
DIET+ConveRT	89.00±0.29	90.57±0.86
DIET+MathBERT-base	84.52±1.19	86.85±1.15
DIET+MathBERT-custom	85.22±0.78	87.80±1.45

Table 4: NLU/Intent Recognition micro-avg F1-scores (%): TF+BERT (baseline), DIET+BERT, DIET+ConveRT, DIET+MathBERT-base (origVocab), DIET+MathBERT-custom (mathVocab) models trained on KidSpace-POC & tested on KidSpace-Deployment datasets.

ble 4). This performance drop is anticipated and explainable as we operate on the noisier real-world utterances collected from younger children in dynamic play-based environments.

To the best of our knowledge, this study presents the first attempt to adopt the lightweight multi-task DIET architecture and incorporate pre-trained MathBERT embeddings as dense features. We combine these MathBERT representations with sparse word and character-level n-gram features in a plug-and-play fashion. The motivation behind this was empowering math domain-specific embeddings for the NLU task targeted at kids playing basic math-learning games. However, we could not observe any benefits of employing MathBERT-base (i.e., pre-trained using origVocab of BERT-base) or MathBERT-custom (i.e., pre-trained using customized mathVocab) representations for the Intent Recognition task. Although MathBERT-custom seems to perform slightly better than MathBERT-base as expected, the gain is insignificant and still way lower than ConveRT and even BERT. There could be many reasons for this, such as the possible mismatch between our domain and advanced mathematical corpora (e.g., mathematical equations and symbols) with graduate-level textbooks that MathBERT trained on. Compared to this, our early childhood math education domain involves more basic concepts and simple operations (e.g., ones and tens, counting, adding, and subtracting). In addition, MathBERT is pre-trained on a relatively small set (i.e., 100M tokens) compared to the massive general-purpose corpora that the BERT models are trained on (3.3B words). ConveRT embeddings have already been shown to perform well on conversational tasks such as Intent Classification, partly because these are pre-trained on large cor-

NLU Model	Math-Game Activity	# Test Samples (Speech/Utterances)	NLU/Intent Recognition F1 (%)	ASR+NLU F1 (%)
DIET+ConveRT	Planting	588	91.8	75.8
DIET+ConveRT	Watering	664	97.4	84.7

Table 5: End-to-End (ASR + NLU/Intent Recognition) Evaluation Results on Planting and Watering Math-game activity datasets: DIET+ConveRT models trained on KidSpace-POC datasets and tested on KidSpace-Deployment datasets.

Math-Game Activity	#test-NLU (utterances)	NLU/Intent F1 (%)	ASR+NLU F1 (%)	#test-DM (responses)	DM/Response F1	NLU+DM F1	ASR+NLU+DM F1
Planting	184	90.5	73.3	209	0.89	0.87	0.82
Watering	346	95.2	84.4	403	0.93	0.91	0.89

Table 6: End-to-End (ASR + NLU/Intent Recognition + DM/Response Selection) Evaluation Results on Planting and Watering Math-game activity datasets: DIET+ConveRT NLU models and TED DM models trained on KidSpace-POC datasets and tested on KidSpace-Deployment datasets.

pora of natural conversational datasets (e.g., Reddit conversational threads). The success of ConveRT representations over MathBERT could also indicate that our educational game datasets involve numerous utterances around play-based conversations tailored towards planting and watering flower use cases. Compared to those, our datasets include limited interactions that directly involve numbers or counting/addition/subtraction operations. Our intent class distributions also support this observation, where we have around 450-600 samples within approximately 2K utterances in POC datasets annotated with directly math-related intents (e.g., *ask-number*, *counting*). In the deployment datasets, we observed around 550-600 math-related utterances within a total of 2.1K samples. Nevertheless, instead of pre-training the BERT models to create MathBERT, one can also explore pre-training the ConveRT model on large and more elementary math-related corpora as the next step.

4.3 End-to-End Evaluation

Our SDS pipeline starts by recognizing user speech via the ASR module and feeds the recognized text into our NLU component. We developed NLU models performing Intent Recognition to interpret user utterances. Then we pass these user intents together with multimodal inputs, such as user actions and objects, into the DM component. The multimodal dialogue manager handles verbal and non-verbal communication inputs from the NLU (e.g., intents) and external nodes processing audiovisual information (e.g., poses, gestures, objects, game events, and actions). We pass these multi-

modal inputs directly to the DM in the form of relevant multimodal intents. The Dialogue State Tracking (DST) model tracks what has happened (i.e., the dialogue state) within the DM. Then, the output of DST is used by the Dialogue Policy to decide which action the system should take next. Our DM models predict the appropriate agent actions and responses based on all the available contextual information (i.e., language-audio-visual inputs, game events, and dialogue history/context from previous turns). When the DM predicts verbal response types, the NLG module retrieves actual bot responses that are template-based in our use cases. We create a variety of response text by preparing multiple templates for each response type, where the final response template is randomly assigned at run-time. Finally, the generated text responses are sent to the TTS module to output agent utterances. Please refer to [Okur et al. \(2022a\)](#) for our multimodal SDS pipeline diagram.

We were assuming perfect (or human-level) speech recognition performances for the NLU results obtained on manual transcriptions (by human transcribers) until now. However, we observed around 30% word-error-rate (WER) in ASR transcriptions for kids’ speech. These ASR outputs are obtained via the top hypothesis given by Google Cloud Speech-to-Text API ⁵. Considering these ASR errors propagating into the ASR+NLU pipeline, the Intent Recognition F1-score performances drop around 11-to-17% (see Tables 5 and 6) when evaluated directly on noisy ASR outputs acquired using Google ASR engine. We are cur-

⁵<https://cloud.google.com/speech-to-text/>

Math-Game	Human Transcription	ASR Output	Intent	Prediction
Planting	Sunflowers	7 flour	<i>answer-flowers</i>	<i>counting</i>
	We need just one more	jasmine ranvir	<i>counting</i>	<i>answer-flowers</i>
	Twenty four	trailer for	<i>counting</i>	<i>intro-game</i>
	We just counted nineteen	he doesn't canton my team	<i>counting</i>	<i>out-of-scope</i>
	Twenty two	tell me too	<i>counting</i>	<i>out-of-scope</i>
	Six	snakes	<i>counting</i>	<i>answer-valid</i>
	Twelve	towel	<i>counting</i>	<i>answer-invalid</i>
	Nine	nah	<i>counting</i>	<i>deny</i>
Watering	Water could help them bloom	why don't can count them you	<i>answer-water</i>	<i>counting</i>
	Thirteen flowers	sure thing flowers	<i>counting</i>	<i>affirm</i>
	Seven	I haven't	<i>counting</i>	<i>deny</i>
	Twenty eight	try it	<i>counting</i>	<i>out-of-scope</i>
	Three	tree	<i>counting</i>	<i>answer-valid</i>
	Five	bye	<i>counting</i>	<i>goodbye</i>
	Bye bye Oscar	buy a hamster	<i>goodbye</i>	<i>answer-invalid</i>

Table 7: ASR + NLU/Intent Recognition Error Samples from Kid Space Planting and Watering Math-game datasets.

rently working towards improving the Automatic Speech + Intent Recognition (i.e., ASR+NLU) performances by exploring the N-best ASR hypotheses (Ganesan et al., 2021) instead of using only the top ASR hypothesis.

For the DM model development, we adopted a recently proposed Transformer Embedding Dialogue (TED) policy architecture (Vlasov et al., 2019), which is highly suitable to our multimodal math-learning use cases. In TED architecture, a transformer’s self-attention mechanism operates over the sequence of dialogue turns to select the appropriate agent response. Despite the noisy ASR outputs with relatively higher WER in kids’ speech compared to adults, when we performed end-to-end evaluations with the ASR+NLU+DM pipeline, we observed only 4-to-7% drops in response prediction F1-score performances (see Table 6). That means error propagation from ASR and NLU has much less effect on the Dialogue Manager (DM) outputs, which are the agent’s final actions and selected responses.

4.4 Error Analysis

Table 7 presents several concrete examples to compare the utterance text obtained by manual human transcriptions (i.e., ground truth) versus problematic ASR outputs (i.e., speech transcriptions). The ground truth intent classes based on gold data annotations on human transcriptions are shown along with the predicted intent classes on ASR outputs obtained by our best-performing DIET+ConveRT NLU models. These ASR errors, especially on domain-specific math-related intents, could explain how errors propagate into the NLU module of

our SDS pipeline and significantly degrade the performance of the Intent Recognition task. Although such ASR errors are expected in noisy real-world application data, especially with kids of age 5-to-8 (Dutta et al., 2022), this analysis also points to a significant room for improvements in the ASR+NLU pipeline. It also encourages us to explore mitigation strategies such as utilizing N-best ASR outputs (Ganesan et al., 2021) and employing phonetic-aware representations (Sundararaman et al., 2021) that can be more robust to ASR errors.

5 Conclusion and Future Work

Improving the quality of mathematics education is vital in developing problem-solving and critical thinking skills for younger learners, which are fundamental for comprehensive STEM education. This study showcased a task-oriented SDS built to promote play-based math learning for early childhood education. The conversational AI system is implemented and evaluated on real-world deployment data collected in classrooms while the kids are practicing basic math concepts. We presented our attempts to enrich the modular SDS pipeline for gamified math learning and prosecuted an end-to-end evaluation using several SDS components tested on real-world deployment data. For NLP applied to math education, language representations can play a significant role due to the exceptional nature of math language. We investigated employing the language representations from the MathBERT model created by pre-training the BERT-base on mathematical corpora. We compared the Intent Recognition results obtained using BERT and ConveRT representations versus the recently proposed

MathBERT embeddings on top of the DIET architecture for NLU. To perform an end-to-end evaluation of our SDS pipeline, we evaluated the ASR, NLU, and DM modules to investigate how error propagation affects the overall SDS performance in real-world math-learning scenarios.

In future work, we aim to explore adopting the N-Best-ASR-Transformer architecture (Ganesan et al., 2021) to utilize multiple ASR hypotheses. This approach can improve the Intent Recognition performances and mitigate recognition errors propagated into the ASR+NLU+DM pipeline due to using only the top ASR hypothesis. Another future direction is to enhance math-specific language representation learning by pre-training the ConveRT model on large math corpora, especially tailored towards early math education (e.g., pre-k to 2nd-grade math curriculum), and then fine-tuning them on our NLU tasks for game-based math learning.

Limitations

Before discussing the limitations of our study, note that the goal of this multimodal dialogue system that we built is to improve the quality of mathematics education for younger learners. To begin with, the cost of the overall setup currently deployed at school (e.g., projector, RGB-D/3D cameras, LiDAR sensor, lapel microphones) can be a limitation, especially for public schools in disadvantaged regions. That can potentially prevent us from having a broader positive impact with our AI for social-good efforts.

Another limitation of this work is the size of the collected datasets. Since the multimodal data collection from authentic classrooms and their labor-intensive annotation process is costly, we need to be innovative to work on such low-data regimes and benefit from the transfer learning paradigm whenever possible. Unfortunately, this data scarcity also affects the generalizability and reliability of our experimental results and end-to-end evaluations, which affects the overall robustness of such real-world applications.

In addition to the deployment costs and data-size concerns, we are bound to use lapel microphones to capture the speech from subjects. That affects the overall unobtrusiveness of the system. Although our ultimate goal is to use the microphone arrays in the classroom, the high WER observed in ASR with kids' speech, even with lapel mics, prevents us from using these far-field mic-array recordings.

Finally, instead of pre-training the BERT model on large math corpora (as performed to create MathBERT), we aimed to pre-train the ConveRT model on a more early-math-related subset of the corresponding corpora. However, although the code and models are publicly available, the authors of MathBERT (Shen et al., 2021b) have not released the fine-tuning math dataset per the data owner's request. Hence, we cannot perform these DIET+MathConveRT experiments before we collect our early-math-related corpora for transfer learning, which is another limitation of this work.

Ethics Statement

It is worth noting that before the proof-of-concept research deployments, a rigorous Privacy Impact Assessment process has been followed. Legal approvals have been sought and received to conduct research with parents, their children, and educators. All research participants signed consent forms before the studies, informing them of all vital details about the studies, including the research goals and procedures, along with how their data would be collected and used to support our research. As with all research projects, our collaborators abide by strict data privacy policies and adhere to ongoing oversight.

Acknowledgements

We gratefully acknowledge our current and former colleagues from the Intel Labs Kid Space team, especially Ankur Agrawal, Glen Anderson, Sinem Aslan, Benjamin Bair, Arturo Bringas Garcia, Rebecca Chierichetti, Hector Cordourier Maruri, Pete Denman, Lenitra Durham, David Gonzalez Aguirre, Sai Prasad, Giuseppe Raffa, Sangita Sharma, and John Sherry, for the conceptualization and the design of school use-cases to support this research. We also would like to show our gratitude to the Rasa team for the open-source framework and the community developers for their contributions that enabled us to conduct our research and evaluate proof-of-concept models for our usages.

References

Glen J. Anderson, Selvakumar Panneer, Meng Shi, Carl S. Marshall, Ankur Agrawal, Rebecca Chierichetti, Giuseppe Raffa, John Sherry, Daria Loi, and Lenitra Megail Durham. 2018. *Kid space: Interactive learning in a smart environment*. In *Proceedings of the Group Interaction Frontiers in Technology*,

- GIFT'18, New York, NY, USA. Association for Computing Machinery.
- Sinem Aslan, Ankur Agrawal, Nese Alyuz, Rebecca Chierichetti, Lenitra M Durham, Ramesh Manuvinakurike, Eda Okur, Saurav Sahay, Sangita Sharma, John Sherry, et al. 2022. Exploring kid space in the wild: a preliminary study of multimodal and immersive collaborative play-based learning experiences. *Educational Technology Research and Development*, pages 1–26.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). In *International Conference on Learning Representations (ICLR 2015)*.
- Ryan S Baker. 2021. Artificial intelligence in education: Bringing it all together. *OECD Digital Education Outlook 2021: Pushing the Frontiers with Artificial Intelligence, Blockchain and Robots*, pages 43–51.
- Nathaniel Blanchard, Michael Brady, Andrew M. Olney, Marci Glaus, Xiaoyi Sun, Martin Nystrand, Borhan Samei, Sean Kelly, and Sidney D’Mello. 2015. A study of automatic speech recognition in noisy classroom environments for automated dialog analysis. In *Artificial Intelligence in Education*, pages 23–33, Cham. Springer International Publishing.
- Tom Bocklisch, Joey Faulkner, Nick Pawlowski, and Alan Nichol. 2017. [Rasa: Open source language understanding and dialogue management](#). In *Conversational AI Workshop, NIPS 2017*.
- Antoine Bordes, Y-Lan Boureau, and Jason Weston. 2017. [Learning end-to-end goal-oriented dialog](#). In *International Conference on Learning Representations (ICLR 2017)*.
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gašić. 2018. [MultiWOZ - a large-scale multi-domain Wizard-of-Oz dataset for task-oriented dialogue modelling](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5016–5026, Brussels, Belgium. Association for Computational Linguistics.
- Tanja Bunk, Daksh Varshneya, Vladimir Vlasov, and Alan Nichol. 2020. [DIET: lightweight language understanding for dialogue systems](#). *CoRR*, abs/2004.09936.
- Mikhail Burtsev, Alexander Seliverstov, Rafael Airapetyan, Mikhail Arkhipov, Dilyara Baymurzina, Nikolay Bushkov, Olga Gureenkova, Taras Khakhulin, Yuri Kuratov, Denis Kuznetsov, Alexey Litinsky, Varvara Logacheva, Alexey Lyman, Valentin Malykh, Maxim Petrov, Vadim Polulyakh, Leonid Pugachev, Alexey Sorokin, Maria Vikhрева, and Marat Zaynutdinov. 2018. [DeepPavlov: Open-source library for dialogue systems](#). In *Proceedings of ACL 2018, System Demonstrations*, pages 122–127, Melbourne, Australia. Association for Computational Linguistics.
- Aoife Cahill, James H Fife, Brian Riordan, Avijit Vajpayee, and Dmytro Galochkin. 2020. [Context-based automated scoring of complex mathematical responses](#). In *Proceedings of the Fifteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 186–192, Seattle, WA, USA → Online. Association for Computational Linguistics.
- Ying-Hong Chan, Ho-Lam Chung, and Yao-Chung Fan. 2021. [Improving controllability of educational question generation by keyword provision](#). *CoRR*, abs/2112.01012.
- Maud Chassignol, Aleksandr Khoroshavin, Alexandra Klimova, and Anna Bilyatdinova. 2018. [Artificial intelligence trends in education: a narrative overview](#). *Procedia Computer Science*, 136:16–24. 7th International Young Scientists Conference on Computational Science, YSC2018, 02-06 July 2018, Heraklion, Greece.
- Hongshen Chen, Xiaorui Liu, Dawei Yin, and Jiliang Tang. 2017. [A survey on dialogue systems: Recent advances and new frontiers](#). *SIGKDD Explor. Newsl.*, 19(2):25–35.
- Heriberto Cuayáhuitl. 2017. [Simpleds: A simple deep reinforcement learning dialogue system](#). In *Dialogues with Social Robots*, pages 109–118. Springer.
- Debajyoti Datta, Maria Phillips, Jennifer L. Chiu, Ginger S. Watson, James P. Bywater, Laura E. Barnes, and Donald E. Brown. 2020. [Improving classification through weak supervision in context-specific conversational agent development for teacher education](#). *CoRR*, abs/2010.12710.
- Kenny Davila and Richard Zanibbi. 2017. [Layout and semantics: Combining representations for mathematical formula search](#). In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1165–1168.
- UN Desa et al. 2016. [Transforming our world: The 2030 agenda for sustainable development](#).
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [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, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Bhuwan Dhingra, Lihong Li, Xiujun Li, Jianfeng Gao, Yun-Nung Chen, Faisal Ahmed, and Li Deng. 2017. [Towards end-to-end reinforcement learning of dialogue agents for information access](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 484–495, Vancouver, Canada. Association for Computational Linguistics.

- Sidney D’Mello and Art Graesser. 2013. [Autotutor and affective autotutor: Learning by talking with cognitively and emotionally intelligent computers that talk back](#). *ACM Trans. Interact. Intell. Syst.*, 2(4).
- Jesse Dodge, Andreea Gane, Xiang Zhang, Antoine Bordes, Sumit Chopra, Alexander H. Miller, Arthur Szlam, and Jason Weston. 2016. [Evaluating prerequisite qualities for learning end-to-end dialog systems](#). In *International Conference on Learning Representations (ICLR 2016)*.
- Satwik Dutta, Dwight Irvin, Jay Buzhardt, and John H.L. Hansen. 2022. [Activity focused speech recognition of preschool children in early childhood classrooms](#). In *Proceedings of the 17th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2022)*, pages 92–100, Seattle, Washington. Association for Computational Linguistics.
- Karthik Ganesan, Pakhi Bamdev, Jaivarsan B, Amresh Venugopal, and Abhinav Tushar. 2021. [N-best ASR transformer: Enhancing SLU performance using multiple ASR hypotheses](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 93–98, Online. Association for Computational Linguistics.
- Chih-Wen Goo, Guang Gao, Yun-Kai Hsu, Chih-Li Huo, Tsung-Chieh Chen, Keng-Wei Hsu, and Yun-Nung Chen. 2018. [Slot-gated modeling for joint slot filling and intent prediction](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 753–757, New Orleans, Louisiana. Association for Computational Linguistics.
- Arthur C Graesser, Shulan Lu, George Tanner Jackson, Heather Hite Mitchell, Mathew Ventura, Andrew Olney, and Max M Louwerse. 2004. Autotutor: A tutor with dialogue in natural language. *Behavior Research Methods, Instruments, & Computers*, 36(2):180–192.
- Yue Gu, Xinyu Li, Shuhong Chen, Jianyu Zhang, and Ivan Marsic. 2017. [Speech intention classification with multimodal deep learning](#). In *Canadian conference on artificial intelligence*, pages 260–271. Springer.
- Dilek Hakkani-Tür, Gokhan Tur, Asli Celikyilmaz, Yun-Nung Chen, Jianfeng Gao, Li Deng, and Ye-Yi Wang. 2016. [Multi-domain joint semantic frame parsing using bi-directional rnn-lstm](#). *Interspeech 2016*, pages 715–719.
- Donghoon Ham, Jeong-Gwan Lee, Youngsoo Jang, and Kee-Eung Kim. 2020. [End-to-end neural pipeline for goal-oriented dialogue systems using GPT-2](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 583–592, Online. Association for Computational Linguistics.
- Matthew Henderson, Iñigo Casanueva, Nikola Mrkšić, Pei-Hao Su, Tsung-Hsien Wen, and Ivan Vulić. 2020. [ConveRT: Efficient and accurate conversational representations from transformers](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2161–2174, Online. Association for Computational Linguistics.
- Shifeng Huang, Jiawei Wang, Jiao Xu, Da Cao, and Ming Yang. 2021. [Real2: An end-to-end memory-augmented solver for math word problems](#). In *Workshop on Math AI for Education (MATHAI4ED), 35th Conference on Neural Information Processing Systems (NeurIPS 2021)*.
- Jiyoun Jia, Yunfan He, and Huixiao Le. 2020. [A multi-modal human-computer interaction system and its application in smart learning environments](#). In *Blended Learning. Education in a Smart Learning Environment*, pages 3–14, Cham. Springer International Publishing.
- Zhuoren Jiang, Liangcai Gao, Ke Yuan, Zheng Gao, Zhi Tang, and Xiaozhong Liu. 2018. [Mathematics content understanding for cyberlearning via formula evolution map](#). In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 37–46.
- Daniel Jurafsky and James H. Martin. 2018. *Ch 24: Dialog Systems and Chatbots. Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*, 3rd (draft) edition. Prentice Hall PTR, Upper Saddle River, NJ, USA.
- Alice Kerry, Richard Ellis, and Susan Bull. 2009. [Conversational agents in e-learning](#). In *Applications and Innovations in Intelligent Systems XVI*, pages 169–182, London. Springer London.
- Aashish Kumar and Anoop Rajagopal. 2021. [Phygital math learning with handwriting for kids](#). In *Workshop on Math AI for Education (MATHAI4ED), 35th Conference on Neural Information Processing Systems (NeurIPS 2021)*.
- John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. [Conditional random fields: Probabilistic models for segmenting and labeling sequence data](#). In *International Conference on Machine Learning, ICML*, pages 282–289.
- Sweta P Lende and MM Raghuwanshi. 2016. [Question answering system on education acts using nlp techniques](#). In *2016 world conference on futuristic trends in research and innovation for social welfare (Startup Conclave)*, pages 1–6. IEEE.
- James C. Lester, Eun Y. Ha, Seung Y. Lee, Bradford W. Mott, Jonathan P. Rowe, and Jennifer L. Sabourin. 2013. [Serious games get smart: Intelligent game-based learning environments](#). *AI Magazine*, 34(4):31–45.

- Bing Liu and Ian Lane. 2016. [Attention-based recurrent neural network models for joint intent detection and slot filling](#). In *Interspeech 2016*, pages 685–689.
- Bing Liu, Gokhan Tur, Dilek Hakkani-Tur, Pararth Shah, and Larry Heck. 2017. End-to-end optimization of task-oriented dialogue model with deep reinforcement learning. In *Conversational AI Workshop, NIPS 2017*.
- Xingkun Liu, Arash Eshghi, Pawel Swietojanski, and Verena Rieser. 2021. [Benchmarking Natural Language Understanding Services for Building Conversational Agents](#), pages 165–183. Springer Singapore, Singapore.
- Ekaterina Loginova and Dries Benoit. 2022. Structural information in mathematical formulas for exercise difficulty prediction: a comparison of nlp representations. In *Proceedings of the 17th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2022)*, pages 101–106.
- Behrooz Mansouri, Shaurya Rohatgi, Douglas W Oard, Jian Wu, C Lee Giles, and Richard Zanibbi. 2019. Tangent-cft: An embedding model for mathematical formulas. In *Proceedings of the 2019 ACM SIGIR international conference on theory of information retrieval*, pages 11–18.
- Shikib Mehri, Mihail Eric, and Dilek Hakkani-Tür. 2020. [Dialoglue: A natural language understanding benchmark for task-oriented dialogue](#). *CoRR*, abs/2009.13570.
- Grégoire Mesnil, Yann Dauphin, Kaisheng Yao, Yoshua Bengio, Li Deng, Dilek Hakkani-Tur, Xiaodong He, Larry Heck, Gokhan Tur, Dong Yu, and Geoffrey Zweig. 2015. [Using recurrent neural networks for slot filling in spoken language understanding](#). *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 23(3):530–539.
- Eda Okur, Shachi H Kumar, Saurav Sahay, Asli Arslan Esme, and Lama Nachman. 2019. [Natural language interactions in autonomous vehicles: Intent detection and slot filling from passenger utterances](#). *20th International Conference on Computational Linguistics and Intelligent Text Processing (CICLing 2019)*.
- Eda Okur, Saurav Sahay, and Lama Nachman. 2022a. [Data augmentation with paraphrase generation and entity extraction for multimodal dialogue system](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 4114–4125, Marseille, France. European Language Resources Association.
- Eda Okur, Saurav Sahay, and Lama Nachman. 2022b. [NLU for game-based learning in real: Initial evaluations](#). In *Proceedings of the 9th Workshop on Games and Natural Language Processing within the 13th Language Resources and Evaluation Conference*, pages 28–39, Marseille, France. European Language Resources Association.
- Kulothunkan Palasundram, Nurfadhline Mohd Sharef, Nurul Nasharuddin, Khairul Kasmiran, and Azreen Azman. 2019. [Sequence to sequence model performance for education chatbot](#). *International Journal of Emerging Technologies in Learning (iJET)*, 14(24):56–68.
- Shuai Peng, Ke Yuan, Liangcai Gao, and Zhi Tang. 2021. [Mathbert: A pre-trained model for mathematical formula understanding](#). *CoRR*, abs/2105.00377.
- Ana Cristina Pires, Fernando González Perilli, Ewelina Bakafa, Bruno Fleisher, Gustavo Sansone, and Sebastián Marichal. 2019. [Building blocks of mathematical learning: Virtual and tangible manipulatives lead to different strategies in number composition](#). *Frontiers in Education*, 4.
- Srikrishna Raamadhurai, Ryan Baker, and Vikraman Poduval. 2019. [Curio SmartChat : A system for natural language question answering for self-paced k-12 learning](#). In *Proceedings of the Fourteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 336–342, Florence, Italy. Association for Computational Linguistics.
- J. Elizabeth Richey, Jiayi Zhang, Rohini Das, Juan Miguel Andres-Bray, Richard Scruggs, Michael Mogessie, Ryan S. Baker, and Bruce M. McLaren. 2021. Gaming and frustration explain learning advantages for a math digital learning game. In *Artificial Intelligence in Education*, pages 342–355, Cham. Springer International Publishing.
- Saurav Sahay, Shachi H. Kumar, Eda Okur, Haroon Syed, and Lama Nachman. 2019. [Modeling intent, dialog policies and response adaptation for goal-oriented interactions](#). In *Proceedings of the 23rd Workshop on the Semantics and Pragmatics of Dialogue*, London, United Kingdom. SEMDIAL.
- Saurav Sahay, Eda Okur, Nagib Hakim, and Lama Nachman. 2021. [Semi-supervised interactive intent labeling](#). In *Proceedings of the Second Workshop on Data Science with Human in the Loop: Language Advances, NAACL 2021*, pages 31–40, Online. Association for Computational Linguistics.
- Iulian Serban, Alessandro Sordani, Yoshua Bengio, Aaron Courville, and Joelle Pineau. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30.
- Iulian Vlad Serban, Ryan Lowe, Peter Henderson, Laurent Charlin, and Joelle Pineau. 2018. A survey of available corpora for building data-driven dialogue systems: The journal version. *Dialogue & Discourse*, 9(1):1–49.
- Pararth Shah, Dilek Hakkani-Tur, and Larry Heck. 2016. [Interactive reinforcement learning for task-oriented dialogue management](#). In *Deep Learning for Action and Interaction Workshop, NIPS 2016*.

- Jia Tracy Shen, Michiharu Yamashita, Ethan Prihar, Neil Heffernan, Xintao Wu, Sean McGrew, and Dongwon Lee. 2021a. Classifying math knowledge components via task-adaptive pre-trained bert. In *International Conference on Artificial Intelligence in Education*, pages 408–419. Springer.
- Jia Tracy Shen, Michiharu Yamashita, Ethan Prihar, Neil T. Heffernan, Xintao Wu, and Dongwon Lee. 2021b. [Mathbert: A pre-trained language model for general NLP tasks in mathematics education](#). *CoRR*, abs/2106.07340.
- Kayleigh Skene, Christine M O’Farrelly, Elizabeth M Byrne, Natalie Kirby, Eloise C Stevens, and Paul G Ramchandani. 2022. Can guidance during play enhance children’s learning and development in educational contexts? a systematic review and meta-analysis. *Child Development*.
- Pei-Hao Su, Paweł Budzianowski, Stefan Ultes, Milica Gasic, and Steve Young. 2017. [Sample-efficient actor-critic reinforcement learning with supervised data for dialogue management](#). In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 147–157, Saarbrücken, Germany. Association for Computational Linguistics.
- Yueqiu Sun, Tangible Play, Rohitkrishna Nambiar, and Vivek Vidyasagan. 2021. [Gamifying math education using object detection](#). In *Workshop on Math AI for Education (MATHAI4ED), 35th Conference on Neural Information Processing Systems (NeurIPS 2021)*.
- Mukuntha Narayanan Sundararaman, Ayush Kumar, and Jithendra Vepa. 2021. Phoneme-bert: Joint language modelling of phoneme sequence and asr transcript. *arXiv preprint arXiv:2102.00804*.
- Abhijit Suresh, Jennifer Jacobs, Margaret Perkoff, James H Martin, and Tamara Sumner. 2022. Fine-tuning transformers with additional context to classify discursive moves in mathematics classrooms. In *Proceedings of the 17th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2022)*, pages 71–81.
- Kaveh Taghipour and Hwee Tou Ng. 2016. [A neural approach to automated essay scoring](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1882–1891, Austin, Texas. Association for Computational Linguistics.
- Stefan Ultes, Lina M. Rojas Barahona, Pei-Hao Su, David Vandyke, Dongho Kim, Iñigo Casanueva, Paweł Budzianowski, Nikola Mrkšić, Tsung-Hsien Wen, Milica Gasic, and Steve Young. 2017. [PyDial: A multi-domain statistical dialogue system toolkit](#). In *Proceedings of ACL 2017, System Demonstrations*, pages 73–78, Vancouver, Canada. Association for Computational Linguistics.
- Andrea Vanzo, Emanuele Bastianelli, and Oliver Lemon. 2019. [Hierarchical multi-task natural language understanding for cross-domain conversational AI: HERMIT NLU](#). In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 254–263, Stockholm, Sweden. Association for Computational Linguistics.
- Akson Sam Varghese, Saleha Sarang, Vipul Yadav, Bharat Karotra, and Niketa Gandhi. 2020. Bidirectional lstm joint model for intent classification and named entity recognition in natural language understanding. *International Journal of Hybrid Intelligent Systems*, 16(1):13–23.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 5998–6008.
- Vladimir Vlasov, Johannes E. M. Mosig, and Alan Nichol. 2019. [Dialogue transformers](#). *CoRR*, abs/1910.00486.
- Liyun Wen, Xiaojie Wang, Zhenjiang Dong, and Hong Chen. 2018. Jointly modeling intent identification and slot filling with contextual and hierarchical information. In *Natural Language Processing and Chinese Computing*, pages 3–15, Cham. Springer International Publishing.
- Tsung-Hsien Wen, David Vandyke, Nikola Mrkšić, Milica Gašić, Lina M. Rojas-Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. 2017. [A network-based end-to-end trainable task-oriented dialogue system](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 438–449, Valencia, Spain. Association for Computational Linguistics.
- Rainer Winkler, Sebastian Hobert, Antti Salovaara, Matthias Söllner, and Jan Marco Leimeister. 2020. [Sara, the Lecturer: Improving Learning in Online Education with a Scaffolding-Based Conversational Agent](#), page 1–14. Association for Computing Machinery, New York, NY, USA.
- Rainer Winkler and Matthias Söllner. 2018. [Unleashing the potential of chatbots in education: A state-of-the-art analysis](#). In *Academy of Management Annual Meeting (AOM)*.
- Ledell Wu, Adam Fisch, Sumit Chopra, Keith Adams, Antoine Bordes, and Jason Weston. 2018. Starspace: Embed all the things! In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, volume 32.
- Ke Yuan, Dafang He, Zhuoren Jiang, Liangcai Gao, Zhi Tang, and C Lee Giles. 2020. Automatic generation of headlines for online math questions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 9490–9497.

- Xuesong Zhai, Xiaoyan Chu, Ching Sing Chai, Morris Siu Yung Jong, Andreja Istenic, Michael Spector, Jia-Bao Liu, Jing Yuan, and Yan Li. 2021. A review of artificial intelligence (AI) in education from 2010 to 2020. *Complexity*, 2021.
- Xiaodong Zhang and Houfeng Wang. 2016. A joint model of intent determination and slot filling for spoken language understanding. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI'16*, pages 2993–2999. AAAI Press.
- Zheng Zhang, Ying Xu, Yanhao Wang, Tongshuang Wu, Bingsheng Yao, Daniel Ritchie, Mo Yu, Dakuo Wang, and Toby Jia-Jun Li. 2021. Building a storytelling conversational agent through parent-ai collaboration.
- Tiancheng Zhao and Maxine Eskenazi. 2016. Towards end-to-end learning for dialog state tracking and management using deep reinforcement learning. In *Proceedings of the 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 1–10, Los Angeles. Association for Computational Linguistics.
- Qianrong Zhou, Liyun Wen, Xiaojie Wang, Long Ma, and Yue Wang. 2016. A hierarchical lstm model for joint tasks. In *Chinese Computational Linguistics and Natural Language Processing Based on Naturally Annotated Big Data*, pages 324–335, Cham. Springer International Publishing.