

Inspecting Spoken Language Understanding from Kids for Basic Math Learning at Home

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

Enriching the quality of early childhood education with interactive math learning at home systems, empowered by recent advances in conversational AI technologies, is slowly becoming a reality. With this motivation, we implement a multimodal dialogue system to support play-based learning experiences at home, guiding kids to master basic math concepts. This work explores Spoken Language Understanding (SLU) pipeline within a task-oriented dialogue system developed for Kid Space, with cascading Automatic Speech Recognition (ASR) and Natural Language Understanding (NLU) components evaluated on our home deployment data with kids going through gamified math learning activities. We validate the advantages of a multi-task architecture for NLU and experiment with a diverse set of pretrained language representations for Intent Recognition and Entity Extraction tasks in the math learning domain. To recognize kids' speech in realistic home environments, we investigate several ASR systems, including the commercial Google Cloud and the latest open-source Whisper solutions with varying model sizes. We evaluate the SLU pipeline by testing our best-performing NLU models on noisy ASR output to inspect the challenges of understanding children for math learning in authentic homes.

1 Introduction and Background

The ongoing progress in Artificial Intelligence (AI) based advanced technologies can assist humanity in reducing the most critical inequities around the globe. The recent widespread interest in conversational AI applications presents exciting opportunities to showcase the positive societal impact of these technologies. The language-based AI systems have already started to mature to a level where we may soon observe their influences in mitigating the most pressing global challenges. Education is among the top priority improvement areas identified by the United Nations (UN) (i.e., poverty,

hunger, healthcare, and education). In particular, increasing the inclusiveness and quality of education is within the UN development goals¹ with utmost urgency. One of the preeminent ways to diminish societal inequity is promoting STEM (i.e., Science, Technology, Engineering, Math) education, specifically ensuring that children succeed in mathematics. It is well-known that acquiring basic math skills at younger ages builds students up for success, regardless of their future career choices (Cesarone, 2008; Torpey, 2012). For math education, interactive learning environments through gamification present substantial leverages over more traditional learning settings for studying elementary math subjects, particularly with younger learners (Skene et al., 2022). With that goal, conversational AI technologies can facilitate this interactive learning environment where students can master fundamental math concepts. Despite these motivations, studying spoken language technologies for younger kids to learn basic math is a vastly uncharted area of AI.

This work discusses a modular goal-oriented Spoken Dialogue System (SDS) specifically targeted for kids to learn and practice basic math concepts at home setup. Initially, a multimodal dialogue system (Sahay et al., 2019) is implemented for Kid Space (Anderson et al., 2018), a gamified math learning application for deployment in authentic classrooms. During this preliminary real-world deployment at an elementary school, the COVID-19 pandemic impacted the globe, and school closures forced students to switch to online learning options at home. To support this sudden paradigm shift to at-home learning, previous school use cases are redesigned for new home usages, and our dialogue system is recreated to deal with interactive math games at home. While the play-based learning activities are adjusted for home usages with a much simpler setup, the multimodal aspects of these games are partially preserved along with the

¹<https://sdgs.un.org/goals>




Figure 1: Multimodal Dialogue System Pipeline

fundamental math concepts for early childhood education. These math skills cover using ones and tens to construct numbers and foundational arithmetic concepts and operations such as counting, addition, and subtraction. The multimodal aspects of these learning games include kids’ spoken interactions with the system while answering math questions and carrying out game-related conversations, physical interactions with the objects (i.e., placing cubes and sticks as manipulatives) on a visually observed playmat, performing specific pose and gesture-based actions as part of these interactive games (e.g., jumping, standing still, air high-five).

Our domain-specific SDS pipeline (see Figure 1) consists of multiple cascaded components, namely Automatic Speech Recognition (ASR), Natural Language Understanding (NLU), Multimodal Dialogue Manager (DM), Natural Language Generation (NLG), and Text-to-Speech (TTS) synchronizing the agent utterances with virtual character animations on Student User Interface (UI). Here we concentrate on the Spoken Language Understanding (SLU) task on kids’ speech at home environments while playing basic math games. Such application-dependent SLU approaches commonly involve two main modules applied sequentially: (i) Speech-to-Text (STT) or ASR module that recognizes speech and transcribes the spoken utterances into text, and (ii) NLU module that interprets the semantics of those utterances by processing the transcribed text. NLU is one of the most integral components of these goal-oriented dialogue systems. It empowers user-agent interactions by understanding the meaning of user utterances via performing domain-specific sub-tasks. Intent Recognition (IR) and Named Entity Recognition (NER) are essential sub-tasks within the NLU module to resolve the complexities of human language and extract meaningful information for the application at hand. Given a user utterance as input, the Intent Classification aims to identify the user’s intention (i.e.,

what the user desires to achieve with that interaction) and categorize the user’s objective at that conversational turn. The Entity Extraction targets locating and classifying entities (i.e., specific terms representing existing things such as person names, locations, and organizations) mentioned in user utterances into predefined task-specific categories.

In this study, we present our efforts to convert the task-oriented SDS (Okur et al., 2022b) designed for school use cases (Aslan et al., 2022) to home usages after COVID-19 and inspect the performance of individual SDS modules evaluated on the home deployment data we recently collected from 12 kids individually at their homes. The current work focuses on assessing and improving the SLU task performance on kids’ utterances at home by utilizing this real-world deployment data. We first investigate the ASR and NLU module evaluations independently. Then, we inspect the overall SLU pipeline (ASR+NLU) performance on kids’ speech by evaluating our NLU tasks on ASR output (i.e., recognized text) at home environments. As the erroneous and noisy speech recognition output would lead to incorrect intent and entity predictions, we aim to understand these error propagation consequences with SLU for children in the math learning domain. We experiment with various recent ASR solutions and diverse model sizes to gain more insights into their capabilities to recognize kids’ speech at home. We then analyze the effects of these ASR engines on understanding intents and extracting entities from children’s utterances. We discuss our findings and observations for potential enhancements in future deployments of this multimodal dialogue system for math learning at home.

2 Related Work

2.1 Conversational AI for Math Learning

With the ultimate goal of improving the quality of education, there has been a growing enthusiasm for exploiting AI-based intelligent systems to boost students’ learning experiences (Chassignol et al., 2018; Aslan et al., 2019; Jia et al., 2020; Zhai et al., 2021; Baker, 2021). Among these, interactive frameworks that support guided play-based learning spaces revealed significant advantages for math learning (Pires et al., 2019; Sun et al., 2021; Richey et al., 2021), especially for building foundational math skills in early childhood education (Nrupatunga et al., 2021; Skene et al., 2022). To attain this level of interactivity within smart

learning spaces, developing innovative educational applications by utilizing language-based AI technologies is in growing demand (Taghipour and Ng, 2016; Lende and Raghuvanshi, 2016; Raamadurai et al., 2019; Cahill et al., 2020; Chan et al., 2021; Rathod et al., 2022). In particular, designing conversational agents for intelligent tutoring is a compelling yet challenging area of research, with several attempts presented so far (Winkler and Söllner, 2018; Wambsganss et al., 2020; Winkler et al., 2020; Datta et al., 2020; Okonkwo and Ade-Ibijola, 2021; Wollny et al., 2021), most of them focusing on language learning (Bibauw et al., 2022; Tyen et al., 2022; Zhang et al., 2022).

In the math education context, earlier conversational math tutoring applications exist, such as SKOPE-IT (Nye et al., 2018), which is based on AutoTutor (Graesser et al., 2005) and ALEKS (Falmagne et al., 2013), and MathBot (Grossman et al., 2019). These are often text-based online systems following strict rules in conversational graphs. Later, various studies emerged at the intersection of cutting-edge AI techniques and math learning (Mansouri et al., 2019; Huang et al., 2021; Azerbayev et al., 2022; Uesato et al., 2022; Yang et al., 2022). Among those, employing advanced language understanding methods to assist math learning is relatively new (Peng et al., 2021; Shen et al., 2021; Loginova and Benoit, 2022; Reusch et al., 2022). The majority of those recent work leans on exploring language representations for math-related tasks such as mathematical reasoning, formula understanding, math word problem-solving, knowledge tracing, and auto-grading, to name a few. Recently, TalkMoves dataset (Suresh et al., 2022a) was released with K-12 math lesson transcripts annotated for discursive moves and dialogue acts to classify teacher talk moves in math classrooms (Suresh et al., 2022b).

For the conversational AI tasks, the latest large language models (LLMs) based chatbots, such as BlenderBot (Shuster et al., 2022) and ChatGPT (OpenAI, 2022), gained a lot of traction in the education community (Tack and Piech, 2022; Kasneci et al., 2023), along with some concerns about using generative models in tutoring (Macina et al., 2023; Cotton et al., 2023). ChatGPT is a general-purpose open-ended interaction agent trained on internet-scale data. It is an end-to-end dialogue model without explicit NLU/Intent Recognizer or DM, which currently cannot fully comprehend the

multimodal context and proactively generate responses to nudge children in a guided manner without distractions. Using these recent chatbots for math learning is still in the early stages because they are known to miss basic mathematical abilities and carry reasoning flaws (Frieder et al., 2023), revealing a lack of common sense. Moreover, they are known to be susceptible to triggering inappropriate or harmful responses and potentially perpetuate human biases since they are trained on internet-scale data and require carefully-thought guardrails.

On the contrary, our unique application is a task-oriented math learning spoken dialogue system designed to perform learning activities, following structured educational games to assist kids in practicing basic math concepts at home. Our SDS does not require massive amounts of data to understand kids and generate appropriate adaptive responses, and the lightweight models can run locally on client machines. In addition, our solution is multimodal, intermixing the physical and digital hybrid learning experience with audio-visual understanding, object recognition, segmentation, tracking, and pose and gesture recognition.

2.2 Spoken Language Understanding

Conventional pipeline-based dialogue systems with supervised learning are broadly favored when initial domain-specific training data is scarce to bootstrap the task-oriented SDS for future data collection (Serban et al., 2018; Budzianowski et al., 2018; Mehri et al., 2020). Deep learning-based modular dialogue frameworks and practical toolkits are prominent in academic and industrial settings (Bocklisch et al., 2017; Burtsev et al., 2018; Reyes et al., 2019). For task-specific applications with limited in-domain data, current SLU systems often use a cascade of two neural modules: (i) ASR maps the input audio to text (i.e., transcript), and (ii) NLU predicts intent and slots/entities from this transcript. Since our main focus in this work is investigating the SLU pipeline, we briefly summarize the existing NLU and ASR solutions.

2.2.1 Language Representations for NLU

The NLU component processes input text, often detects intents, and extracts referred entities from user utterances. For the mainstream NLU tasks of Intent Classification and Entity Recognition, jointly trained multi-task models are proposed (Liu and Lane, 2016; Zhang and Wang, 2016; Goo et al., 2018) with hierarchical learning approaches (Wen

et al., 2018; Okur et al., 2019; Vanzo et al., 2019). Transformer architecture (Vaswani et al., 2017) is a game-changer for several downstream language tasks. With Transformers, BERT (Devlin et al., 2019) is presented, which became one of the most pivotal breakthroughs in language representations, achieving high performance in various tasks, including NLU. Later, Dual Intent and Entity Transformer (DIET) architecture (Bunk et al., 2020) is invented as a lightweight multi-task NLU model. On multi-domain NLU-Benchmark data (Liu et al., 2021b), the DIET model outperformed fine-tuning BERT for joint Intent and Entity Recognition.

For BERT-based autoencoding approaches, RoBERTa (Liu et al., 2019) is presented as a robustly optimized BERT model for sequence and token classification. The Hugging Face introduced a smaller, lighter general-purpose language representation model called DistilBERT (Sanh et al., 2019) as the knowledge-distilled version of BERT. ConveRT (Henderson et al., 2020) is proposed as an efficiently compact model to obtain pretrained sentence embeddings as conversational representations for dialogue-specific tasks. LaBSE (Feng et al., 2022) is a pretrained multilingual model producing language-agnostic BERT sentence embeddings that achieve promising results in text classification.

The GPT family of autoregressive LLMs, such as GPT-2 (Radford et al., 2019) and GPT-3 (Brown et al., 2020), perform well at what they are pretrained for, i.e., text generation. GPT models can also be adopted for NLU, supporting few-shot learning capabilities, and NLG in task-oriented dialogue systems (Madotto et al., 2020; Liu et al., 2021a). XLNet (Yang et al., 2019) applies autoregressive pretraining for representation learning that adopts Transformer-XL (Dai et al., 2019) as a backbone model and works well for language tasks with lengthy contexts. DialoGPT (Zhang et al., 2020) extends GPT-2 as a large-scale neural response generation model for multi-turn conversations trained on Reddit discussions, whose representations can be exploited in dialogue tasks.

For language representations to be utilized in math-related tasks, MathBERT (Shen et al., 2021) is introduced as a math-specific BERT model pretrained on large math corpora. Later, Math-aware-BERT and Math-aware-RoBERTa models (Reusch et al., 2022) are proposed based on BERT and RoBERTa, pretrained on Math Stack Exchange².

²<https://math.stackexchange.com/>

2.2.2 Speech Recognition with Kids

Speech recognition technology has been around for some time, and numerous ASR solutions are available today, both commercial and open-source. Rockhopper ASR (Stemmer et al., 2017) is an earlier low-power speech recognition engine with LSTM-based language models, where its acoustic models are trained using an open-source Kaldi speech recognition toolkit (Povey et al., 2011). Google Cloud Speech-to-Text³ is a prominent commercial ASR service powered by advanced neural models and designed for speech-dependant applications. Until recently, Google STT API was arguably the leader in ASR services for recognition performance and language coverage. Franck Deroncourt (2018) reported that Google ASR could reach a word error rate (WER) of 12.1% on LibriSpeech clean dataset (28.8% on LibriSpeech other) (Panayotov et al., 2015) at that time, which is improved drastically over time. Recently, Open AI released Whisper ASR (Radford et al., 2022) as a game-changer speech recognizer. Whisper models are pretrained on a vast amount of labeled audio-transcription data (i.e., 680k hours), unlike its predecessors (e.g., Wav2Vec 2.0 (Baevski et al., 2020) is trained on 60k hours of unlabeled audio). 117k hours of this data are multilingual, which makes Whisper applicable to over 96 languages, including low-resourced ones. Whisper architecture follows a standard Transformer-based encoder-decoder as many speech-related models (Latif et al., 2023). The Whisper-base model is reported to achieve 5.0% & 12.4% WER on LibriSpeech clean & other.

Although speech recognition systems are substantially improving to achieve human recognition levels, problems still occur, especially in noisy environments, with users having accents and dialects or underrepresented groups like kids. Child speech brings distinct challenges to ASR (Stemmer et al., 2003; Gerosa et al., 2007; Yeung and Alwan, 2018), such as data scarcity and highly varied acoustic, linguistic, physiological, developmental, and articulatory characteristics compared to adult speech (Claus et al., 2013; Shivakumar and Georgiou, 2020; Bhardwaj et al., 2022). Thus, WER for children's voices is reported two-to-five times worse than for adults (Wu et al., 2019), as the younger the child, the poorer ASR performs. There exist efforts to mitigate these difficulties of speech recognition with kids (Shivakumar et al.,

³<https://cloud.google.com/speech-to-text/>

2014; Duan and Chen, 2020; Booth et al., 2020; Kelly et al., 2020; Rumberg et al., 2021; Yeung et al., 2021). Few studies also focus on speech technologies in educational settings (Reeder et al., 2015; Blanchard et al., 2015; Bai et al., 2021, 2022; Dutta et al., 2022), often for language acquisition, reading comprehension, and story-telling activities.

3 Methods

3.1 Home Learning Data and Use Cases

We utilize two datasets for gamified basic math learning at home usages. The first set is a proof-of-concept (POC) data manually constructed based on User Experience (UX) studies (e.g., detailed scripts for new home use cases) and partially adopted from our previous school data (Okur et al., 2022a). This POC data is used to train and cross-validate various NLU models to develop the best practices in later home deployments. The second set is our recent home deployment data collected from 12 kids (ages 7-8) experiencing our multimodal math learning system at authentic homes. The audio-visual data is transcribed manually, and user utterances in these reference transcripts are annotated for intent and entity types we identified for each learning activity at home. Table 1 compares the NLU statistics for Kid Space Home POC and Deployment datasets. Manually transcribed children’s utterances in deployment data are employed to test our best NLU models trained on POC data. We run multiple ASR engines on audio recordings from home deployment data, where automatic transcripts (i.e., ASR output) are utilized to compute WER to assess ASR model performances on kids’ speech. We also evaluate the SLU pipeline (ASR+NLU) by testing NLU models on ASR output from deployment data.

The simplified home deployment setup includes a playmat with physical manipulatives, a laptop with a built-in camera, a wireless lavalier mic, and a depth camera on a tripod. Home use cases follow a particular flow of activities designed for play-based learning in early childhood education. These activities are Introduction (Meet & Greet), Warm-up Game (Red Light Green Light), Training Game, Learning Game, and Closure (Dance Party). After meeting with the virtual character and playing jumping games, the child starts the training game, where the agent asks for help planting flowers. The agent presents tangible manipulatives, cubes representing ones and sticks representing tens, and instructs the kid to answer ba-

NLU Data Statistics	POC	Deployment
# Intents Types	13	12
Total # Utterances	4091	733
# Entity Types	3	3
Total # Entities	2244	497
Min # Utterances per Intent	105	1
Max # Utterances per Intent	830	270
Avg # Utterances per Intent	314.7	61.1
Min # Tokens per Utterance	1	1
Max # Tokens per Utterance	40	33
Avg # Tokens per Utterance	4.49	2.30
# Unique Tokens (Vocab Size)	702	149
Total # Tokens	18364	1689

Table 1: Kid Space Home POC and Deployment Data

sic math questions and construct numbers using these objects, going through multiple rounds of practice questions where flowers in child-selected colors bloom as rewards. In the actual learning game, the agent presents clusters of questions involving ones & tens, and the child provides verbal (e.g., stating the numbers) and visual answers (e.g., placing the cubes and sticks on the playmat, detected by the overhead camera). The agent provides scaffolding utterances and performs animations to show and tell how to solve basic math questions. The interaction ends with a dance party to celebrate achievements and say goodbyes in closure. Some of our intents can be considered generic (e.g., *state-name*, *affirm*, *deny*, *repeat*, *out-of-scope*), but some are highly domain-specific (e.g., *answer-flowers*, *answer-valid*, *answer-others*, *state-color*, *had-fun-a-lot*, *end-game*) or math-related (e.g., *state-number*, *still-counting*). The entities we extract are activity-specific (i.e., *name*, *color*) and math-related (i.e., *number*).

3.2 NLU and ASR Models

Customizing open-source Rasa framework (Bocklisch et al., 2017) as a backbone, we investigate several NLU models for Intent Recognition and Entity Extraction tasks to implement our math learning conversational AI system for home usage. Our baseline approach is inspired by the StarSpace (Wu et al., 2018) method, a supervised embedding-based model maximizing the similarity between utterances and intents in shared vector space. We enrich this simple text classifier by incorporating SpaCy (Honnibal et al., 2020) pre-

trained language models⁴ for word embeddings as additional features in the NLU pipeline. CRF Entity Extractor (Lafferty et al., 2001) with BILOU tagging is also part of this baseline NLU. For home usages, we explore the advantages of switching to a more recent DIET model⁵ for joint Intent and Entity Recognition, a multi-task architecture with two-layer Transformers shared for NLU tasks. DIET leverages combining dense features (e.g., any given pretrained embeddings) with sparse features (e.g., token-level encodings of char n-grams). To observe the net benefits of DIET, we first pass the identical SpaCy embeddings used in our baseline (StarSpace) as dense features to DIET. Then, we adopt DIET with pretrained BERT⁶, RoBERTa⁷, and DistilBERT⁸ word embeddings, as well as ConVeRT⁹ and LaBSE¹⁰ sentence embeddings to inspect the effects of these autoencoding-based language representations on NLU performance (see 2.2.1 for more details). We also evaluate pretrained embeddings from models using autoregressive training such as XLNet¹¹, GPT-2^{12,13}, and DialoGPT¹⁴ on top of DIET. Next, we explore recently-proposed math-language representations pretrained on math data for our basic math learning dialogue system. MathBERT (Shen et al., 2021) is pretrained on large math corpora (e.g., curriculum, textbooks, MOOCs, arXiv papers) covering pre-k to college-graduate materials. We enhance DIET by incorporating embeddings from MathBERT-base¹⁵ and MathBERT-custom¹⁶ models, pretrained with BERT-base original and math-customized vocabularies, respectively. Math-aware-BERT¹⁷ and Math-aware-RoBERTa¹⁸ mod-

⁴https://github.com/explosion/spacy-models/releases/tag/en_core_web_md-3.5.0

⁵Please check Bunk et al. (2020) for hyper-parameter tuning, hardware specs, and computational costs.

⁶<https://huggingface.co/bert-base-uncased>

⁷<https://huggingface.co/roberta-base>

⁸<https://huggingface.co/distilbert-base-uncased>

⁹<https://github.com/connorbrinton/polyai-models/releases>

¹⁰<https://huggingface.co/rasa/LaBSE>

¹¹<https://huggingface.co/xlnet-base-cased>

¹²<https://huggingface.co/gpt2>

¹³Excluded GPT-3 and beyond that are not open-source.

¹⁴<https://huggingface.co/microsoft/DialoGPT-medium>

¹⁵<https://huggingface.co/tbs17/MathBERT>

¹⁶<https://huggingface.co/tbs17/MathBERT-custom>

¹⁷https://huggingface.co/AnReu/math_pretrained_bert

¹⁸https://huggingface.co/AnReu/math_pretrained_roberta

els (Reusch et al., 2022) are initialized from BERT-base and RoBERTa-base, and further pretrained on Math StackExchange¹⁹ with extra LaTeX tokens to better tokenize math formulas for ARQMath-3 tasks (Mansouri et al., 2022). We exploit these representations with DIET to investigate their effects on our NLU tasks in the basic math domain.

For the ASR module, we explore three main speech recognizers for our math learning application at home, which are explained further in 2.2.2. Rockhopper ASR²⁰ is the baseline local approach previously inspected, which can be adjusted slightly for kids. Its acoustic models rely on Kaldi²¹ generated resources and are trained on default adult speech data. In the past explorations, when Rockhopper’s language models fine-tuned with limited in-domain kids’ utterances (Sahay et al., 2021) from previous school usages, WER decreased by 40% for kids but remained 50% higher than adult WER. Although this small-scale baseline solution is unexpected to reach Google Cloud ASR performance, Rockhopper has a few other advantages for our application since it can run offline locally on low-power devices, which could be better for security, privacy, latency, and cost (relative to cloud-based ASR services). Google ASR is a commercial cloud solution providing high-quality speech recognition service but requiring connectivity and payment, which cannot be adapted or fine-tuned as Rockhopper. The third ASR approach we investigate is Whisper²², which combines the best of both worlds as it is an open-source adjustable solution that can run locally, achieving new state-of-the-art (SOTA) results. We inspect three configurations of varying model sizes (i.e., base, small, and medium) to evaluate the Whisper ASR for our home math learning usage with kids.

4 Experimental Results

To build the NLU module of our SLU pipeline, we train Intent and Entity Classification models and cross-validate them over the Kid Space Home POC dataset to decide upon the best-performing NLU architectures moving forward for home. Table 2 summarizes the results of model selection experiments with various NLU models. We report the average of 5 runs, and each run involves a 10-fold

¹⁹<https://archive.org/download/stackexchange>

²⁰https://docs.openvino.ai/2018_R5/_samples_speech_sample_README.html

²¹<https://github.com/kaldi-asr/kaldi>

²²<https://github.com/openai/whisper>

NLU Model	Intent Detection	Entity Extraction
StarSpace+SpaCy	92.71±0.25	97.08±0.21
DIET+SpaCy	94.29±0.05	98.38±0.12
DIET+BERT	97.25±0.23	99.23±0.02
DIET+RoBERTa	95.50±0.18	99.11±0.12
DIET+DistilBERT	97.41±0.20	99.49±0.12
DIET+ConveRT	98.80±0.25	99.61±0.03
DIET+LaBSE	98.19±0.18	99.72±0.04
DIET+XLNet	94.99±0.19	98.38±0.14
DIET+GPT-2	95.35±0.27	99.01±0.27
DIET+DialogPT	96.00±0.49	98.94±0.12
DIET+MathBERT-base	94.55±0.22	98.10±0.21
DIET+MathBERT-custom	94.61±0.34	97.48±0.29
DIET+Math-aware-BERT	95.95±0.15	98.94±0.19
DIET+Math-aware-RoBERTa	94.20±0.16	98.75±0.21

Table 2: NLU Model Selection Results in F1-scores (%) Evaluated on Kid Space Home POC Data (10-fold CV)

cross-validation (CV) on POC data. Compared to the baseline StarSpace algorithm, we gain almost 2% F1 score for intents and more than 1% F1 for entities with multi-task DIET architecture. For language representations, we observe that incorporating DIET with the BERT family of embeddings from autoencoders achieves higher F1 scores relative to the GPT family of embeddings from autoregressive models. We cannot reveal any benefits of employing math-specific representations with DIET, as all such models achieve worse than DIET+BERT results. One reason we identify is the mismatch between our early math domain and advanced math corpora, including college-level math symbols and equations, that these models trained on. Another reason could be that such embeddings are pretrained on smaller math corpora (e.g., 100 million tokens) compared to massive-scale generic corpora (e.g., 3.3 billion words) that BERT models use for training. DIET+ConveRT is the clear winner for intents and achieves second-best but very close results for entities compared to DIET+LaBSE. ConveRT and LaBSE are both sentence-level embeddings, but ConveRT performs well on dialogue tasks as it is pretrained on large conversational corpora, including Reddit discussions. Based on these results, we select DIET+ConveRT as the final multi-task architecture for our NLU tasks at home.

Next, we evaluate our NLU module on Kid Space Home Deployment data collected at authentic homes over 12 sessions with 12 kids. Each child goes through 5 activities within a session, as described in 3.1. In Table 3, we observe overall F1% drops (Δ) of 4.6 for intents and 0.3 for entities when our best-performing DIET+ConveRT models are tested on home deployment data. These findings are expected and relatively lower than

Activity	Intent Detection			Entity Extraction		
	POC	Deploy	Δ	POC	Deploy	Δ
Intro (Meet & Greet)	99.9	97.3	-2.6	99.2	97.4	-1.8
Warm-up Game	98.8	93.4	-5.4	-	-	-
Training Game	98.4	94.2	-4.2	99.9	99.8	-0.1
Learning Game	98.9	94.3	-4.6	99.8	99.4	-0.4
Closure (Dance)	98.8	98.7	-0.1	-	-	-
All Activities	98.8	94.2	-4.6	99.6	99.3	-0.3

Table 3: NLU Evaluation Results in F1-scores (%) for DIET+ConveRT Models Trained on Kid Space Home POC Data & Tested on Home Deployment Data

the performance drops we previously observed at school (Okur et al., 2022c). We witness distributional and utterance-length differences between POC/training and deployment/test datasets. Real-world data would always be noisier than anticipated as these utterances come from younger kids playing math games in dynamic conditions.

To further improve the performance of our Kid Space Home NLU models (trained on POC data) by leveraging this recent deployment data, we experiment with merging the two datasets for training and evaluating the performance on individual deployment sessions via leave-one-out (LOO) CV. At each of the 12 runs (for 12 sessions/kids), we merge the POC data with 11 sessions of deployment data for model training and use the remaining session as a test set, then take the average performance of these runs. That would simulate how combining POC with real-world deployment data would help us train more robust NLU models that perform better on unseen data in future deployment sessions. The overall F1-scores reach 96.5% for intents (2.3% gain from 94.2%) and 99.4% for entities (0.1% gain) with LOOCV, which are promising for our future deployments.

To inspect the ASR module of our SLU pipeline, we experiment with Rockhopper, Google, and Whisper-base/small/medium ASR models evaluated on the same audio data collected during home deployments. Using the manual session transcripts as a reference, we compute the average WER for kids with each ASR engine to investigate the most feasible solution. Table 4 summarizes WER results before and after standard pre-processing steps (e.g., lower casing and punctuation removal) as well as application-specific filters (e.g., num2word and cleaning). The numbers are transcribed inconsistently within reference transcripts plus ASR output (e.g., 35 vs. thirty-five), and we need to standardize them all in word forms. The cleaning

ASR Model	Raw Output	Lowercase (LC)	Remove Punct (RP)	Num2Word (NW)	LC & RP	LC & RP & NW	NW & Clean	LC & RP & NW & Clean
Rockhopper	0.939	0.919	0.924	0.937	0.886	0.884	0.937	0.884
Google Cloud	0.829	0.798	0.775	0.763	0.695	0.602	0.763	0.602
Whisper-base	1.042	1.020	0.971	0.985	0.946	0.856	0.622	0.500
Whisper-small	0.834	0.804	0.760	0.756	0.720	0.621	0.537	0.405
Whisper-medium	0.905	0.870	0.824	0.814	0.785	0.675	0.522	0.384

Table 4: ASR Model Results: Avg Word Error Rates (WER) for Child Speech at Kid Space Home Deployment Data

step is applied to Whisper ASR output only due to known issues such as getting stuck in repeat loops and hallucinations (Radford et al., 2022). We seldom observe trash output from Whisper (4-to-7%) having very long transcriptions with non-sense repetitions/symbols, which hugely affect WER due to their length, yet these samples can be easily auto-filtered. Even after these steps, the relatively high error rates can be attributed to many factors related to the characteristics of these recordings (e.g., incidental voice and phrases), very short utterances to be recognized (e.g., binary yes/no answers or stating numbers with one-or-two words), and recognizing kids’ speech in ordinary home environments. Still, the comparative results indicate that Whisper ASR solutions perform better on kids, and we can benefit from increasing the model size from base to small, while small to medium is close.

For SLU pipeline evaluation, we test our highest-performing NLU models on noisy ASR output. Table 5 presents the Intent and Entity Classification results achieved on home deployment data where the DIET+ConveRT models run on varying ASR models output. Note that Voice Activity Detection (VAD) is an integral part of ASR that decides the presence/absence of human speech. We realize that the VAD stage is filtering out a lot of audio chunks with actual kid speech with Rockhopper and Google. Thus, our VAD-ASR nodes can ignore a lot of audio segments with reference transcripts (57.9% for Rokchopper, 49.1% for Google). That is less of an issue with Whisper-base/small/medium, missing 7.1%/5.7%/4.4% of transcribed utterances (often due to filtering very long and repetitive trash Whisper output). When we treat these entirely missed utterances with no ASR output as classification errors for NLU tasks (i.e., missing to predict intent/entities when no speech is detected), we can adjust the F1-scores accordingly to evaluate the VAD-ASR+NLU pipeline. These VAD-adjusted F1-scores are compared in Table 5, aligned with the WER results, where NLU on Whisper ASR

ASR Model	Intent Detection		Entity Extraction	
	F1	Adjusted-F1	F1	Adjusted-F1
Rockhopper	36.7	15.5	82.9	35.0
Google Cloud	78.0	39.7	96.2	49.0
Whisper-base	64.7	60.0	95.4	88.5
Whisper-small	72.2	68.1	96.6	91.1
Whisper-medium	76.5	73.1	98.5	94.1

Table 5: SLU Pipeline Evaluation Results in F1-scores (%) for ASR+NLU and VAD-Adjusted ASR+NLU on Kid Space Home Deployment Data

performs relatively higher than Google and Rockhopper. For enhanced Intent Recognition in real-world deployments with kids, increasing the ASR model size from small to medium could be worth the trouble for Whisper. Yet, the F1 drop is still huge, from 94.2% with NLU to 73.1% with VAD-ASR+NLU, when VAD-ASR errors propagate into the SLU pipeline.

5 Error Analysis

For NLU error analysis, Table 6 reveals utterance samples from our Kid Space Home Deployment data with misclassified intents obtained by the DIET+ConveRT models on manual/human transcripts. These language understanding errors illustrate the potential pain points solely related to the NLU model performances, as we are assuming perfect or human-level ASR here by feeding the manually transcribed utterances into the NLU. Such intent prediction errors occur in real-world deployments for many reasons. For example, authentic user utterances can have multiple intents (e.g., “Yeah. Can we have some carrots?” starts with *affirm* and continues with *out-of-scope*). Some utterances can be challenging due to subtle differences between intent classes (e.g., “Ah this is 70, 7.” is submitting a verbal answer with *state-number* but can easily be mixed with *still-counting* too). Moreover, we observe utterances having *colors* and “flowers” within *out-of-scope* (e.g., “Wow, that’s a lot of red flowers.”), which can be confusing for the NLU models trained on cleaner POC datasets.

Sample Kid Utterance	Intent	Prediction
Pepper.	<i>state-name</i>	<i>answer-valid</i>
Wow, that’s a lot of red flowers.	<i>out-of-scope</i>	<i>answer-flowers</i>
None.	<i>state-number</i>	<i>deny</i>
Nothing.	<i>state-number</i>	<i>deny</i>
Yeah. Can we have some carrots?	<i>affirm</i>	<i>out-of-scope</i>
Okay, Do your magic.	<i>affirm</i>	<i>out-of-scope</i>
Maybe tomorrow.	<i>affirm</i>	<i>out-of-scope</i>
He’s a bear.	<i>out-of-scope</i>	<i>answer-valid</i>
I like the idea of a bear	<i>out-of-scope</i>	<i>answer-valid</i>
Oh, 46? Okay.	<i>still-counting</i>	<i>state-number</i>
94. Okay.	<i>still-counting</i>	<i>state-number</i>
Now we have mountains.	<i>out-of-scope</i>	<i>answer-valid</i>
A pond?	<i>out-of-scope</i>	<i>answer-valid</i>
Sorry, I didn’t understand it. Uh, five tens.	<i>state-number</i>	<i>still-counting</i>
Ah this is 70, 7.	<i>state-number</i>	<i>still-counting</i>

Table 6: NLU Error Analysis: Intent Recognition Error Samples from Kid Space Home Deployment Data

Human Transcript	ASR Output	ASR Model	Intent	Prediction
Six.	thanks	Rockhopper	<i>state-number</i>	<i>thank</i>
fifteen	if he	Rockhopper	<i>state-number</i>	<i>out-of-scope</i>
fifteen	Mickey	Google Cloud	<i>state-number</i>	<i>state-name</i>
Five.	bye	Google Cloud	<i>state-number</i>	<i>goodbye</i>
Blue.	Blair.	Whisper-base	<i>state-color</i>	<i>state-name</i>
twenty	Plenty.	Whisper-base	<i>state-number</i>	<i>had-fun-a-lot</i>
A lot.	Oh, la.	Whisper-base	<i>had-fun-a-lot</i>	<i>out-of-scope</i>
A lot.	Oh, wow.	Whisper-small	<i>had-fun-a-lot</i>	<i>out-of-scope</i>
Two.	you	Whisper-small	<i>state-number</i>	<i>out-of-scope</i>
Four.	I’m going to see this floor.	Whisper-small	<i>state-number</i>	<i>out-of-scope</i>
twenty	Swamy?	Whisper-medium	<i>state-number</i>	<i>state-name</i>
Eight.	E.	Whisper-medium	<i>state-number</i>	<i>out-of-scope</i>

Table 7: SLU Pipeline (ASR+NLU): Intent Recognition Error Samples from Kid Space Home Deployment Data

For further error analysis on the SLU pipeline (ASR+NLU), Table 7 demonstrates Intent Recognition error samples from Kid Space Home Deployment data obtained on ASR output with several speech recognition models we explored. These samples depict anticipated error propagation from speech recognition to language understanding modules in the cascaded SLU approach. Please check Appendix A for a more detailed ASR error analysis.

6 Conclusion

To increase the quality of math learning experiences at home for early childhood education, we develop a multimodal dialogue system with play-based learning activities, helping the kids gain basic math skills. This study investigates a modular SLU pipeline for kids with cascading ASR and NLU modules, evaluated on our first home deployment data with 12 kids at individual homes. For NLU, we examine the advantages of a multi-task architecture and experiment with numerous pre-

trained language representations for Intent Recognition and Entity Extraction tasks in our application domain. For ASR, we inspect the WER with several solutions that are either low-power and local (e.g., Rockhopper), commercial (e.g., Google Cloud), or open-source (e.g., Whisper) with varying model sizes and conclude that Whisper-medium outperforms the rest on kids’ speech at authentic home environments. Finally, we evaluate the SLU pipeline by running our best-performing NLU models, DIET+ConveRT, on VAD-ASR output to observe the significant effects of cascaded errors due to noisy voice detection and speech recognition performance with kids in realistic home deployment settings. In the future, we aim to fine-tune the Whisper ASR acoustic models on kids’ speech and language models on domain-specific math content. Moreover, we consider exploring N-Best-ASR-Transformers (Ganesan et al., 2021) to leverage multiple Whisper ASR hypotheses and mitigate errors propagated into cascading SLU.

Limitations

By building this task-specific dialogue system for kids, we aim to increase the overall quality of basic math education and learning at-home experiences for younger children. In our previous school deployments, the overall cost of the whole school/classroom setup, including the wall/ceiling-mounted projector, 3D/RGB-D cameras, LiDAR sensor, wireless lavalier microphones, servers, etc., can be considered as a limitation for public schools and disadvantaged populations. When we shifted our focus to home learning usages after the COVID-19 pandemic, we simplified the overall setup for 1:1 learning with a PC laptop with a built-in camera, a depth camera on a tripod, a lapel mic, and a playmat with cubes and sticks. However, even this minimal instrumentation suitable for home setup can be a limitation for kids with lower socioeconomic status. Moreover, the dataset size of our initial home deployment data collected from 12 kids in 12 sessions is relatively small, with around 12 hours of audio data manually transcribed and annotated. Collecting multimodal data at authentic homes of individual kids within our target age group (e.g., 5-to-8 years old) and labor-intensive labeling process is challenging and costly. To overcome these data scarcity limitations and develop dialogue systems for kids with such small-data regimes, we had to rely on transfer learning approaches as much as possible. However, the dataset sizes affect the generalizability of our explorations, the reliability of some results, and ultimately the robustness of our multimodal dialogue system for deployments with kids in the real world.

Ethics Statement

Prior to our initial research deployments at home, a meticulous process of Privacy Impact Assessment is pursued. The legal approval processes are completed to operate our research with educators, parents, and the kids. Individual participants and parties involved have signed the relevant consent forms in advance, which inform essential details about our research studies. The intentions and procedures and how the participant data will be collected and utilized to facilitate our research are explained in writing in these required consent forms. Our collaborators comply with stricter data privacy policies as well.

Acknowledgements

We aspire to share our gratitude and acknowledge our former and current colleagues in the Kid Space team at Intel Labs. Particularly: (i) Hector Coudrier Maruri, Juan Del Hoyo Ontiveros, and Georg Stemmer for developing the VAD-ASR node to obtain the ASR output that we use in our SLU pipeline; (ii) Benjamin Bair, Lenitra Durham, Sai Prasad, Giuseppe Raffa, Celal Savur, and Sangita Sharma for designing the HW/SW architectural setup, developing the Wizard UI and game logic nodes, and supporting data collection; (iii) Ankur Agrawal, Arturo Bringas Garcia, Vishwajeet Narwal, and Guillermo Rivas Aguilar for developing the Student UI via the Unity game engine; (iv) Glen Anderson, Sinem Aslan, Rebecca Chierichetti, Pete Denman, John Sherry, and Meng Shi for conducting UX studies and performing interaction design to conceptualize school and home usages; (v) David Gonzalez Aguirre, Gesem Gudino Mejia, and Julio Zamora Esquivel for developing the visual understanding nodes to support this research. We would also want to gratefully acknowledge our field team members from Summa Linguae Technologies, especially Rick Lin and Brenda Tumbalobos Cubas, for their exceptional support in executing the data collection and transcription/annotation tasks in collaboration with our Intel Labs Kid Space team. Finally, we should thank the Rasa team and community developers for their open-source framework and contributions that empowered us to conduct our research.

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, Giuseppe Raffa, and Lama Nachman. 2022. *Exploring kid space in the wild: a preliminary study of multimodal and immersive collaborative play-based learning experiences*. *Educational Technology Research and Development*, 70:205–230.
- Sinem Aslan, Nese Alyuz, Cagri Tanriover, Sinem E.

- Mete, Eda Okur, Sidney K. D’Mello, and Asli Arslan Esme. 2019. [Investigating the impact of a real-time, multimodal student engagement analytics technology in authentic classrooms](#). In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI ’19, page 1–12, New York, NY, USA. Association for Computing Machinery.
- Zhangir Azerbayev, Bartosz Piotrowski, and Jeremy Avigad. 2022. [Proofnet: A benchmark for autoformalizing and formally proving undergraduate-level mathematics problems](#). In *Workshop MATH-AI: Toward Human-Level Mathematical Reasoning, 36th Conference on Neural Information Processing Systems (NeurIPS 2022)*, New Orleans, Louisiana, USA.
- Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. 2020. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in neural information processing systems*, 33:12449–12460.
- Yu Bai, Ferdy Hubers, Catia Cucchiari, and Helmer Strik. 2021. [An ASR-based Reading Tutor for Practicing Reading Skills in the First Grade: Improving Performance through Threshold Adjustment](#). In *Proc. IberSPEECH 2021*, pages 11–15.
- Yu Bai, Ferdy Hubers, Catia Cucchiari, Roeland van Hout, and Helmer Strik. 2022. [The Effects of Implicit and Explicit Feedback in an ASR-based Reading Tutor for Dutch First-graders](#). In *Proc. Interspeech 2022*, pages 4476–4480.
- 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.
- Vivek Bhardwaj, Mohamed Tahar Ben Othman, Vinay Kukreja, Youcef Belkhier, Mohit Bajaj, B Srikanth Goud, Ateeq Ur Rehman, Muhammad Shafiq, and Habib Hamam. 2022. Automatic speech recognition (asr) systems for children: A systematic literature review. *Applied Sciences*, 12(9):4419.
- Serge Bibauw, Wim Van den Noortgate, Thomas François, and Piet Desmet. 2022. Dialogue systems for language learning: a meta-analysis. *Language Learning & Technology*, 26(1).
- 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*.
- Eric Booth, Jake Carns, Casey Kennington, and Nader Rafla. 2020. [Evaluating and improving child-directed automatic speech recognition](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 6340–6345, Marseille, France. European Language Resources Association.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- 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, Nickolay Bushkov, Olga Gureenkova, Taras Khakhulin, Yuri Kuratov, Denis Kuznetsov, Alexey Litinsky, Varvara Logacheva, Alexey Lymar, 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.
- Bernard Cesarone. 2008. Early childhood mathematics: Promoting good beginnings. *Childhood Education*, 84(3):189.
- 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.

- Felix Claus, Hamurabi Gamboa Rosales, Rico Petrick, Horst-Udo Hain, and Rüdiger Hoffmann. 2013. A survey about databases of children’s speech. In *INTERSPEECH*, pages 2410–2414.
- Debby RE Cotton, Peter A Cotton, and J Reuben Shipway. 2023. Chatting and cheating: Ensuring academic integrity in the era of chatgpt. *Innovations in Education and Teaching International*, pages 1–12.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc Le, and Ruslan Salakhutdinov. 2019. [Transformer-XL: Attentive language models beyond a fixed-length context](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2978–2988, Florence, Italy. Association for Computational Linguistics.
- 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.
- 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.
- Richeng Duan and Nancy F Chen. 2020. Unsupervised feature adaptation using adversarial multi-task training for automatic evaluation of children’s speech. In *INTERSPEECH*, pages 3037–3041.
- 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.
- Jean-Claude Falmagne, Dietrich Albert, Christopher Doble, David Eppstein, and Xiangen Hu. 2013. *Knowledge spaces: Applications in education*. Springer Science & Business Media.
- Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Arivazhagan, and Wei Wang. 2022. [Language-agnostic BERT sentence embedding](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 878–891, Dublin, Ireland. Association for Computational Linguistics.
- Walter Chang Franck Dernoncourt, Trung Bui. 2018. A framework for speech recognition benchmarking. In *Interspeech*.
- Simon Frieder, Luca Pinchetti, Ryan-Rhys Griffiths, Tommaso Salvatori, Thomas Lukasiewicz, Philipp Christian Petersen, Alexis Chevalier, and Julius Berner. 2023. Mathematical capabilities of chatgpt. *arXiv preprint arXiv:2301.13867*.
- 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.
- Matteo Gerosa, Diego Giuliani, and Fabio Brugnara. 2007. Acoustic variability and automatic recognition of children’s speech. *Speech Communication*, 49(10-11):847–860.
- 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.
- A.C. Graesser, P. Chipman, B.C. Haynes, and A. Olney. 2005. [Autotutor: an intelligent tutoring system with mixed-initiative dialogue](#). *IEEE Transactions on Education*, 48(4):612–618.
- Joshua Grossman, Zhiyuan Lin, Hao Sheng, Johnny Tian-Zheng Wei, Joseph J Williams, and Sharad Goel. 2019. Mathbot: Transforming online resources for learning math into conversational interactions. *AAAI 2019 Story-Enabled Intelligence*.
- 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.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. [spaCy: Industrial-strength natural language processing in python](#).
- 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)*.
- Jiyou 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.
- Enkelejda Kasneci, Kathrin Seßler, Stefan Küchemann, Maria Bannert, Daryna Dementieva, Frank Fischer, Urs Gasser, Georg Groh, Stephan Günemann, Eyke Hüllermeier, et al. 2023. Chatgpt for good? on opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103:102274.
- Amelia C Kelly, Eleni Karamichali, Armin Saeb, Karel Veselý, Nicholas Parslow, Agape Deng, Arnaud Letondor, Robert O’Regan, and Qiru Zhou. 2020. Soapbox labs verification platform for child speech. In *INTERSPEECH*, pages 486–487.
- 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.
- Siddique Latif, Aun Zaidi, Heriberto Cuayahuitl, Fahad Shamshad, Moazzam Shoukat, and Junaid Qadir. 2023. Transformers in speech processing: A survey. *arXiv preprint arXiv:2303.11607*.
- Sweta P Lende and MM Raghuvanshi. 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.
- 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.
- Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. 2021a. Gpt understands, too. *arXiv preprint arXiv:2103.10385*.
- Xingkun Liu, Arash Eshghi, Pawel Swietojanski, and Verena Rieser. 2021b. Benchmarking natural language understanding services for building conversational agents. In *Increasing Naturalness and Flexibility in Spoken Dialogue Interaction: 10th International Workshop on Spoken Dialogue Systems*, pages 165–183. Springer.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized BERT pretraining approach](#). *CoRR*, abs/1907.11692.
- Ekaterina Logina 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.
- Jakub Macina, Nico Daheim, Lingzhi Wang, Tanmay Sinha, Manu Kapur, Iryna Gurevych, and Mrinmaya Sachan. 2023. Opportunities and challenges in neural dialog tutoring. *arXiv preprint arXiv:2301.09919*.
- Andrea Madotto, Zihan Liu, Zhaojiang Lin, and Pascale Fung. 2020. Language models as few-shot learner for task-oriented dialogue systems. *arXiv preprint arXiv:2008.06239*.
- Behrooz Mansouri, Vít Novotný, Anurag Agarwal, Douglas W. Oard, and Richard Zanibbi. 2022. Overview of arqmath-3 (2022): Third clef lab on answer retrieval for questions on math. In *Experimental IR Meets Multilinguality, Multimodality, and Interaction*, pages 286–310, Cham. Springer International Publishing.
- 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.
- Nrupatunga, 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)*.
- Benjamin D Nye, Philip I Pavlik, Alistair Windsor, Andrew M Olney, Mustafa Hajeer, and Xiangen Hu. 2018. Skope-it (shareable knowledge objects as portable intelligent tutors): overlaying natural language tutoring on an adaptive learning system for mathematics. *International journal of STEM education*, 5:1–20.
- Chinedu Wilfred Okonkwo and Abejide Ade-Ibijola. 2021. Chatbots applications in education: A systematic review. *Computers and Education: Artificial Intelligence*, 2:100033.
- 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](#). In *Computational Linguistics and Intelligent Text Processing*, pages 334–350, Cham. Springer Nature Switzerland.
- Eda Okur, Saurav Sahay, Roddy Fuentes Alba, and Lama Nachman. 2022a. [End-to-end evaluation of a spoken dialogue system for learning basic mathematics](#). In *Proceedings of the 1st Workshop on Mathematical Natural Language Processing (MathNLP)*, pages 51–64, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.

- Eda Okur, Saurav Sahay, and Lama Nachman. 2022b. [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. 2022c. [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.
- OpenAI. 2022. Chatgpt: Optimizing language models for dialogue.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. Librispeech: an asr corpus based on public domain audio books. In *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 5206–5210. IEEE.
- 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 Bakala, 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.
- Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Lukas Burget, Ondrej Glembek, Nagendra Goel, Mirko Hannemann, Petr Motlicek, Yanmin Qian, Petr Schwarz, et al. 2011. The kaldi speech recognition toolkit. In *IEEE 2011 workshop on automatic speech recognition and understanding*, CONF. IEEE Signal Processing Society.
- 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.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2022. Robust speech recognition via large-scale weak supervision. *arXiv preprint arXiv:2212.04356*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Manav Rathod, Tony Tu, and Katherine Stasaski. 2022. [Educational multi-question generation for reading comprehension](#). In *Proceedings of the 17th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2022)*, pages 216–223, Seattle, Washington. Association for Computational Linguistics.
- Kenneth Reeder, Jon Shapiro, Jane Wakefield, and Reg D’Silva. 2015. Speech recognition software contributes to reading development for young learners of english. *International Journal of Computer-Assisted Language Learning and Teaching (IJ-CALLT)*, 5(3):60–74.
- Anja Reusch, Maik Thiele, and Wolfgang Lehner. 2022. Transformer-encoder and decoder models for questions on math. *Proceedings of the Working Notes of CLEF 2022*, pages 5–8.
- Roberto Reyes, David Garza, Leonardo Garrido, Víctor De la Cueva, and Jorge Ramirez. 2019. Methodology for the implementation of virtual assistants for education using google dialogflow. In *Advances in Soft Computing: 18th Mexican International Conference on Artificial Intelligence, MICAI 2019, Xalapa, Mexico, October 27–November 2, 2019, Proceedings 18*, pages 440–451. Springer.
- 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 confrustion explain learning advantages for a math digital learning game. In *Artificial Intelligence in Education*, pages 342–355, Cham. Springer International Publishing.
- Lars Rumberg, Hanna Ehlert, Ulrike Lüdtke, and Jörn Ostermann. 2021. Age-invariant training for end-to-end child speech recognition using adversarial multi-task learning. In *Interspeech*, pages 3850–3854.
- 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 - Full Papers*, London, United Kingdom. SEM-DIAL.
- 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*, pages 31–40, Online. Association for Computational Linguistics.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. [Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter](#). In *5th EMC2 Workshop - Energy Efficient Training and Inference of Transformer Based Models, 33rd Conference on Neural Information Processing Systems (NeurIPS 2019)*.
- 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.
- Jia Tracy Shen, Michiharu Yamashita, Ethan Prihar, Neil T. Heffernan, Xintao Wu, and Dongwon Lee. 2021. [Mathbert: A pre-trained language model for general NLP tasks in mathematics education](#). *CoRR*, abs/2106.07340.
- Prashanth Gurunath Shivakumar and Panayiotis Georgiou. 2020. Transfer learning from adult to children for speech recognition: Evaluation, analysis and recommendations. *Computer speech & language*, 63:101077.
- Prashanth Gurunath Shivakumar, Alexandros Potamianos, Sungbok Lee, and Shrikanth Narayanan. 2014. [Improving speech recognition for children using acoustic adaptation and pronunciation modeling](#). In *Fourth Workshop on Child Computer Interaction (WOCCI 2014)*.
- Kurt Shuster, Jing Xu, Mojtaba Komeili, Da Ju, Eric Michael Smith, Stephen Roller, Megan Ung, Moya Chen, Kushal Arora, Joshua Lane, et al. 2022. Blenderbot 3: a deployed conversational agent that continually learns to responsibly engage. *arXiv preprint arXiv:2208.03188*.
- 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*.
- Georg Stemmer, Munir Georges, Joachim Hofer, Piotr Rozen, Josef G Bauer, Jakub Nowicki, Tobias Bocklet, Hannah R Colett, Ohad Falik, Michael Deisher, et al. 2017. Speech recognition and understanding on hardware-accelerated dsp. In *Interspeech*, pages 2036–2037.
- Georg Stemmer, Christian Hacker, Stefan Steidl, and Elmar Nöth. 2003. Acoustic normalization of children’s speech. In *Eighth European Conference on Speech Communication and Technology*.
- 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).
- Abhijit Suresh, Jennifer Jacobs, Charis Harty, Margaret Perkoff, James H. Martin, and Tamara Sumner. 2022a. [The TalkMoves dataset: K-12 mathematics lesson transcripts annotated for teacher and student discursive moves](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 4654–4662, Marseille, France. European Language Resources Association.
- Abhijit Suresh, Jennifer Jacobs, Margaret Perkoff, James H. Martin, and Tamara Sumner. 2022b. [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, Seattle, Washington. Association for Computational Linguistics.
- Anais Tack and Chris Piech. 2022. The ai teacher test: Measuring the pedagogical ability of blender and gpt-3 in educational dialogues. In *Proceedings of the 15th International Conference on Educational Data Mining*, page 522.
- 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.
- Elka Torpey. 2012. Math at work: Using numbers on the job. *Occupational Outlook Quarterly*, 56(3):2–13.
- Gladys Tyen, Mark Brenchley, Andrew Caines, and Paula Buttery. 2022. [Towards an open-domain chatbot for language practice](#). In *Proceedings of the 17th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2022)*, pages 234–249, Seattle, Washington. Association for Computational Linguistics.
- Jonathan Uesato, Nate Kushman, Ramana Kumar, H Francis Song, Noah Yamamoto Siegel, Lisa Wang, Antonia Creswell, Geoffrey Irving, and Irina Higgins. 2022. [Solving math word problems with process-based and outcome-based feedback](#). In *Workshop MATH-AI: Toward Human-Level Mathematical Reasoning, 36th Conference on Neural Information Processing Systems (NeurIPS 2022)*, New Orleans, Louisiana, USA.
- 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.
- 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.
- Thiemo Wambsganss, Rainer Winkler, Matthias Söllner, and Jan Marco Leimeister. 2020. A conversational agent to improve response quality in course evaluations. In *Extended Abstracts of the 2020 CHI conference on human factors in computing systems*, pages 1–9.

- 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.
- 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. In *Proceedings of the 2020 CHI conference on human factors in computing systems*, pages 1–14.
- 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)*.
- Sebastian Wollny, Jan Schneider, Daniele Di Mitri, Joshua Weidlich, Marc Rittberger, and Hendrik Drachler. 2021. Are we there yet?-a systematic literature review on chatbots in education. *Frontiers in artificial intelligence*, 4:654924.
- Fei Wu, Leibny Paola García-Perera, Daniel Povey, and Sanjeev Khudanpur. 2019. Advances in automatic speech recognition for child speech using factored time delay neural network. In *Interspeech*, pages 1–5.
- Ledell Wu, Adam Fisch, Sumit Chopra, Keith Adams, and Antoine Bordes Jason Weston. 2018. Starspace: Embed all the things! In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence and Thirtieth Innovative Applications of Artificial Intelligence Conference and Eighth AAAI Symposium on Educational Advances in Artificial Intelligence*, AAAI’18/IAAI’18/EAAI’18. AAAI Press.
- Zhicheng Yang, Jinghui Qin, Jiaqi Chen, Liang Lin, and Xiaodan Liang. 2022. [LogicSolver: Towards interpretable math word problem solving with logical prompt-enhanced learning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 1–13, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. [Xlnet: Generalized autoregressive pretraining for language understanding](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Gary Yeung and Abeer Alwan. 2018. On the difficulties of automatic speech recognition for kindergarten-aged children. *Interspeech 2018*.
- Gary Yeung, Ruchao Fan, and Abeer Alwan. 2021. Fundamental frequency feature normalization and data augmentation for child speech recognition. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6993–6997. IEEE.
- 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.
- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020. [DIALOGPT : Large-scale generative pre-training for conversational response generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 270–278, Online. Association for Computational Linguistics.
- Zheng Zhang, Ying Xu, Yanhao Wang, Bingsheng Yao, Daniel Ritchie, Tongshuang Wu, Mo Yu, Dakuo Wang, and Toby Jia-Jun Li. 2022. Storybuddy: A human-ai collaborative chatbot for parent-child interactive storytelling with flexible parental involvement. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pages 1–21.

A Appendix: Additional Error Analysis

Please refer to Table 8 for additional error analysis on ASR output from our home deployment data. Here, we compare manually transcribed utterances (i.e., human transcripts) with the speech recognition output (i.e., raw ASR transcripts) using five different ASR models that we investigated in this study. These ASR errors demonstrate the challenges faced in the speech recognition model performances on kids’ speech, which potentially would be propagated into the remaining modules in the conventional task-oriented dialogue pipeline.

We may attribute various factors to these speech recognition errors, often related to our deployment data characteristics. Incidental voices and phrases constitute a good chunk of the overall home deployment data, along with very short utterances to be recognized (e.g., stating names, colors, types of flowers, numbers, and binary answers with one-or-two words), plus the remaining known challenges present with recognizing kids’ speech in noisy real-world environments.

Human Transcript	Rockhopper	Google Cloud	Whisper-base	Whisper-small	Whisper-medium
Atticus.	-	-	Yeah, that's cute.	I have a kiss.	Now I have to kiss.
I am Genevieve.	i'm twenty-two	I'm going to be	I'm Kennedy.	I'm Genevieve.	I'm Genevieve.
Red.	rab	-	Ralph.	Red.	Red.
Blue.	lil	blue	Blair.	Blue.	Blue.
Yes,	laughs	yes	Yes?	Yes?	Yes?
Roses.	it is	roses	Okay.	Okay	focus
Zero.	you know	no	No.	No, no.	No.
four.	you swore	-	forward.	Over.	Over.
five.	-	bye	Bye.	Bye.	Bye.
eight	all	-	Thank you.	Bye.	Oh
forty eight	wall e	48	48	48	48
forty nine	already	49	49	49	49
fifty one	if you want	51	51	51	51
seventy four	stopping before	74	74	74	74
Maybe tomorrow.	novarro	tomorrow	I need some water, though.	I'm going to leave it tomorrow.	I'm leaving tomorrow.
Flowers, flowers in the greenhouse?	lean forward phelps hours than we	Greenhouse	In forward, in forward, in the green house.	I think forward, both flowers and the greenhouse.	In the green house.
There are seventeen, and seventeen minus ten equals seven.	seventeen seventeen rooms	17 + 17 - 27	There are 17 and 17 minus 10 equals 7.	There are 17 and 17 minus 10 equals 7.	What is the maximum number of children in the world? Um... There are 17 and 17 minus 10 equals 7.

Table 8: ASR Error Samples from Kid Space Home Deployment Data