

Data Augmentation with Paraphrase Generation and Entity Extraction for Multimodal Dialogue System

Eda Okur, Saurav Sahay, Lama Nachman

Intel Labs

USA

{eda.okur, saurav.sahay, lama.nachman}@intel.com

Abstract

Contextually aware intelligent agents are often required to understand the users and their surroundings in real-time. Our goal is to build Artificial Intelligence (AI) systems that can assist children in their learning process. Within such complex frameworks, Spoken Dialogue Systems (SDS) are crucial building blocks to handle efficient task-oriented communication with children in game-based learning settings. We are working towards a multimodal dialogue system for younger kids learning basic math concepts. Our focus is on improving the Natural Language Understanding (NLU) module of the task-oriented SDS pipeline with limited datasets. This work explores the potential benefits of data augmentation with paraphrase generation for the NLU models trained on small task-specific datasets. We also investigate the effects of extracting entities for conceivably further data expansion. We have shown that paraphrasing with model-in-the-loop (MITL) strategies using small seed data is a promising approach yielding improved performance results for the Intent Recognition task.

Keywords: Spoken Dialogue System, Natural Language Understanding, Intent Classification, Entity Recognition, Paraphrase Generation, Data Augmentation

1. Introduction

Building Artificial Intelligence (AI) systems that can assist students in the math learning process has been a challenging yet exciting area of research. Play-based learning systems can offer significant advantages in teaching fundamental mathematical concepts interactively, especially for younger kids (Skene et al., 2021). Such intelligent systems are often required to handle multimodal understanding of the children and their surroundings in real-time. Within these complex architectures, Spoken Dialogue Systems (SDS) are crucial building blocks for carrying out efficient task-oriented communication with kids in game-based learning settings. This work presents our multimodal dialogue system for children learning basic math concepts. This dialogue system technology needs to be constructed and evaluated carefully to handle goal-oriented interactions between the kids and a virtual conversational agent.

This study primarily focuses on creating and improving the Natural Language Understanding (NLU) module of a task-oriented SDS pipeline, especially with limited dataset resources. Building the NLU module of such goal-oriented SDS often involves the definition of intents (and entities if necessary), creation of domain-specific and task-relevant datasets, annotation of the data with intents and entities, iterative training and evaluation of NLU models, and repeating this process for each new or updated use-cases. This work explores the potential benefits of data augmentation with paraphrase generation for the NLU models trained on small-size task-specific datasets. The main NLU task we concentrate on improving is the Intent Recognition (IR) from possible user utterances. The ultimate goal of IR is to predict the user’s intent (i.e., what the user

wants to accomplish within a task-oriented SDS) given an input utterance. In addition to paraphrasing the possible user utterances for increased intent samples, we investigate the effects of extracting entities from these utterances for potentially further data expansion.

Our experiments show that paraphrasing with model-in-the-loop (MITL) strategies is a promising approach to boost performance results for the IR task on our small-scale task-specific datasets. To be more precise, we first increase the F1-score of our baseline NLU model by 5% (i.e., 90.6 to 95.6) for the intents by adopting a multi-task architecture. Then, we improve this further by 4% (i.e., 95.6 to 99.4) with MITL data augmentation. With a Transformer-based multi-task architecture for joint Intent and Entity Recognition, we investigate employing and auto-annotating the entities to improve the NLU performance. Our next goal is to obtain more variations in paraphrased samples with entity expansion to create semantically richer datasets.

2. Related Work

Investigating intelligent systems to assist children in their learning process has been an attractive area of research (Jia et al., 2020). Employing Natural Language Processing (NLP) for building educational applications has also gained popularity in the past decade (Meurers, 2012; Lende and Raghuvanshi, 2016; Taghipour and Ng, 2016; Cahill et al., 2020). Game-based learning environments can enhance significant benefits in teaching basic math concepts interactively, particularly for young learners (Skene et al., 2021). Since we aim to build conversational agents for early childhood education with scarce datasets, we summarize the existing SDS/NLU approaches and data augmentation methods.

2.1. Conversational AI Systems

Dialogue systems are often categorized as either task-oriented or open-ended, where the task-oriented dialogue systems are designed to fulfill specific tasks and handle goal-oriented conversations, whereas the open-ended systems or chat-bots allow more generic conversations such as chit-chat (Jurafsky and Martin, 2018). With the advancements of deep learning-based language technologies and increased availability of large datasets in the research community, the end-to-end trained dialogue systems are shown to produce promising results for both goal-oriented (Bordes et al., 2016) and open-ended (Dodge et al., 2015) applications. Dialogue Managers (DM) of goal-oriented systems are often sequential decision-making models where optimal policies are learned via reinforcement learning from a high number of user interactions (Shah et al., 2016; Dhingra et al., 2016; Liu et al., 2017; Su et al., 2017; Cuayáhuitl, 2017). However, building such systems with limited user interactions is remarkably challenging. Thus, supervised learning approaches with modular SDS pipelines are still widely preferred when initial training data is limited, basically to bootstrap the goal-oriented conversational agents for further data collection (Sahay et al., 2019). Statistical and neural network-based dialogue system toolkits and frameworks (Bocklisch et al., 2017; Ultes et al., 2017; Burtsev et al., 2018) are also heavily used in the academic and industrial research communities for implicit dialogue context management. A recent study named Conversation Learner (Shukla et al., 2020) describes an interactive DM via machine teaching with human-in-the-loop annotations. Although the majority of task-oriented dialogue systems require defining intents and entities, a recent work called SMCaFlow (Andreas et al., 2020) argues for a richer representation of dialogue state as a dataflow graph.

2.2. Natural Language Understanding

The NLU module within SDS processes user utterances as input and typically predicts the user intents and entities of interest. Long Short-Term Memory (LSTM) networks (Hochreiter and Schmidhuber, 1997) and Bidirectional LSTMs (BiLSTM) (Schuster and Paliwal, 1997) have been widely used for sequence learning tasks such as Intent Classification (Hakkani-Tur et al., 2016) and Slot Filling (Mesnil et al., 2015). Jointly training Intent Recognition and Entity Extraction models have been explored recently (Zhang and Wang, 2016; Liu and Lane, 2016; Goo et al., 2018; Varghese et al., 2020). Various hierarchical multi-task architectures are also proposed for these joint NLU tasks (Zhou et al., 2016; Wen et al., 2018; Okur et al., 2019; Vanzo et al., 2019), even some in multimodal context (Gu et al., 2017; Okur et al., 2020). Vaswani et al. (2017) proposed the Transformer, a novel network architecture based entirely on attention mechanisms (Bahdanau et al., 2014). Transformer-based models usually achieve

better results than RNN-based models for the NLU tasks (Sahay et al., 2019). Bidirectional Encoder Representations from Transformers (BERT) (Devlin et al., 2018) is one of the main breakthroughs in pre-trained language representations, showing strong performance in numerous NLP tasks, including the NLU. More recently, Bunk et al. (2020) introduced the Dual Intent and Entity Transformer (DIET), a lightweight multi-task architecture for joint Intent Classification and Entity Recognition. The authors showed that DIET outperforms fine-tuning BERT for predicting intents and entities on a complex multi-domain NLU-Benchmark dataset (Liu et al., 2021) and is much faster to train.

2.3. Data Augmentation

Earlier studies have shown that data augmentation can improve the classification performance for several NLP tasks (Barzilay and McKeown, 2001; Dolan and Brockett, 2005; Lan et al., 2017; Hu et al., 2019). Back-translation (Sennrich et al., 2016) is a popular method to construct paraphrase corpora by translating the samples to another language and then back to the original language. Wieting et al. (2017) employed the Neural Machine Translation (NMT) (Sutskever et al., 2014) approach for translating the non-English part of the parallel corpus to obtain English-to-English paraphrase pairs. Later, the authors massively scaled this approach to generate huge paraphrase corpora called ParaNMT-50M (Wieting and Gimpel, 2018). In order to learn how to generate meaningful paraphrases, some of the previous work have utilized the autoencoders (Socher et al., 2011; Bowman et al., 2016), encoder-decoder models such as BART (Lewis et al., 2019), and NMT (Sokolov and Filimonov, 2020). Recent studies explore data augmentation via Natural Language Generation (NLG) for few-shot intents (Xia et al., 2020) and paraphrase generation for intents and slots in task-oriented dialogue systems (Jolly et al., 2020). Another relevant recent work (Panda et al., 2021) is an extension of a transformer-based model by Jolly et al. (2020) that works for multilingual paraphrase generation for intents and slots, even in the zero-shot settings. Several other recent works have also been exploring data augmentation with fine-tuning large language models and few-shot learning for intent classification and slot-filling tasks (Kumar et al., 2019; Kumar et al., 2020; Lee et al., 2021).

3. Application Domain

The motivation behind this study is to build conversational agents as part of the Kid Space project for early childhood education. Kid Space aims to create smart spaces for children with traditional gaming motivations such as level achievements and virtually collecting objects. The space allows multiple children to interact, which can encourage social development. The agent should accurately comprehend inputs from children and provide feedback. The system needs to be



Figure 1: Learning basic math via play-based interactions.

physically grounded to allow children to bring meaningful objects into the play experience, such as physical toys and manipulatives as learning materials. Thus, Kid Space AI system would combine various sensing technologies that should interact with children, track each child, and monitor their progress (Anderson et al., 2018; Aslan et al., 2022).

The use-cases contain specific flows of interactive games facilitating elementary math learning designed for children 5-to-8 years old. The Flowerpot game builds on the math concepts of tens and ones, with the larger flowerpots representing ‘tens’ and smaller pots ‘ones’. The virtual character provides a number of flowers the children should plant, and when the children have placed the correct number of pots against the wall, digital flowers appear. In the NumberGrid game, children are shown math clues, and when the correct number is touched on the grid, water is virtually poured to water the flowers. These experiences require the space to be physically grounded. As an example, at certain point in the game, the virtual character needs help to get on the table, where children position physical boxes that the character can jump upon, showing an understanding of the physical space.

Figure 1 demonstrates our intelligent agent, Oscar the teddy bear, helping the kids with learning tens and ones concepts along with practicing simple counting, addition, and subtraction operations.

The technology behind Kid Space includes distinct recognition capabilities. In the Flowerpot game, the 2D computer vision algorithm based on AprilTags is used to associate the physical flowerpots with the virtual flowers using a standard RGB camera. Our dialogue system takes multimodal information to incorporate user identity, actions, gestures, and the audio context. During the Flowerpot experience, the virtual character asks the children if they are done placing pots, to which they respond ‘yes’. Our dialogue system needs to use the visual input for the agent to respond appropriately to the correct number of pots being detected. The visual, audio, and LiDAR-based gesture recognition enables physically situated interactions. User identification utilizing the body and face would allow to accurately recognize children across different cameras

and attribute their actions accordingly. 3D scene understanding algorithms are being used for the boxes experience to provide detection and 6D pose estimation of multiple concurrent boxes.

4. Multimodal SDS

In this section, we describe the overall architecture of our multimodal dialogue system that we build for the Kid Space project. The aim is to enable children to interact with the agent while performing various activities, including learning math concepts with physical manipulatives or objects. For that, the dialog system must accurately comprehend multimodal inputs from children and respond appropriately.

Note that the current use-cases are designed for two children playing and learning with the agent collaboratively, while an adult user (i.e., Facilitator) is also present in the room to interact with the agent for game progress and helping the kids whenever needed. Therefore, we need to build a robust multimodal and multi-party conversational system that incorporates number of learning modules for interacting with multiple users (i.e., two children and one adult). This goal-oriented dialogue system should provide game instructions, guide the kids, understand both the kids’ and Facilitator’s utterances and actions to respond appropriately.

4.1. Architecture

The SDS pipeline starts from recognizing user speech via Automatic Speech Recognition (ASR) module and feed the recognized text into our NLU component. We develop NLU models performing Intent Recognition and Entity Extraction to interpret user utterances. Then we pass these user intents and entities 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 entities) and separate external nodes processing visual and audio information (e.g., faces, poses, gestures, objects, events and actions). Note that we pass these multimodal inputs directly to the DM (bypassing the NLU) in the form of relevant multimodal intents for goal-oriented interactions. The Dialogue State Track-

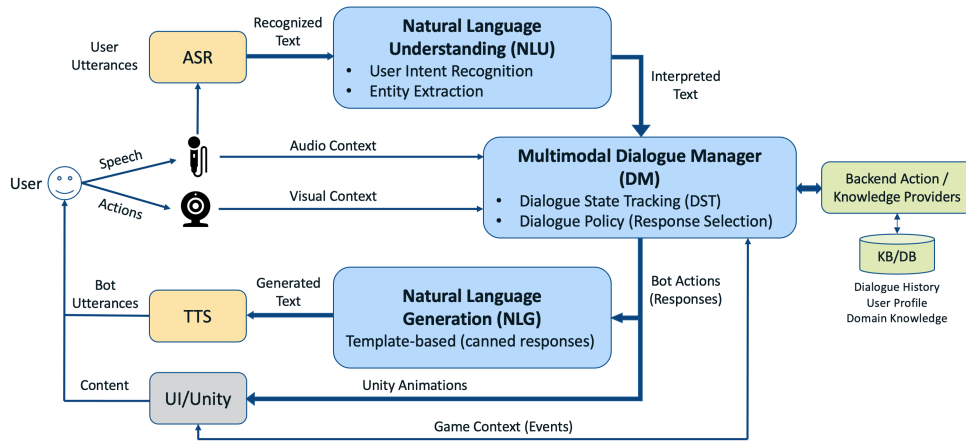


Figure 2: Multimodal SDS pipeline.

ing (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). That means the DM generates sequential actions for bot utterances and non-verbal events. When the verbal response types are predicted, based on the output classes of DM, the NLG module retrieves actual bot responses, which are template-based in our use-cases. We create a variety of responses by preparing multiple templates (i.e., 3-to-6 variations) for each response type. Among these, the final response template randomly assigned at runtime. Generating grammatically correct and semantically coherent responses is challenging with such scarce datasets. Hence, this approach is more reliable than training NLG models for our application. Finally, the generated text responses are sent to the Text-to-Speech (TTS) module to output agent utterances. Non-verbal agent actions such as animations and game events are sent to the Unity application, which serves as the end User Interface (UI) displaying the agent and learning game content. Figure 2 illustrates the schematic representation of our modular SDS pipeline.

5. Model Development

This section describes the models we develop for the NLU (i.e., Intent classification and Entity Recognition) and multimodal DM modules of our dialogue system pipeline. For data augmentation explorations, we also discuss the development of our paraphrase generation models here.

5.1. NLU and DM

Our NLU and DM models are built on top of the Rasa open-source framework (Bocklisch et al., 2017). The former baseline Intent Classifier in Rasa was based on

supervised embeddings provided within the Rasa NLU, which is an embedding-based text classifier that embeds user utterances and intents into the same vector space. This former baseline architecture is inspired by the StarSpace work (Wu et al., 2017), where the embeddings are trained by maximizing the similarity between intents and utterances. In previous work (Sahay et al., 2019), we enriched this former baseline NLU/Intent Recognition architecture available in Rasa by incorporating additional features and adapting alternative network architectures. To be more precise, we adapted the Transformer network (Vaswani et al., 2017) and incorporated pre-trained BERT embeddings (Devlin et al., 2018) to improve the Intent Recognition performance, as shown in Sahay et al. (2019). In this current study, our new baseline NLU model is the best-performing approach from our previous work (Sahay et al., 2019), which we would call TF+BERT in our experiments.

In this work, we explore potential improvements in Intent Classification performance by adapting the recent DIET architecture (Bunk et al., 2020). DIET is a transformer-based multi-task architecture for joint Intent Recognition and Entity Extraction. DIET can incorporate pre-trained word and sentence embeddings from language models as dense features, with the flexibility to combine these with token level one-hot encodings and multi-hot encodings of character n-grams as sparse features. Note that one can use any pre-trained embeddings as dense features in DIET, such as GloVe (Pennington et al., 2014), BERT (Devlin et al., 2018), and ConveRT (Henderson et al., 2020). Conversational Representations from Transformers (ConveRT) is another recent and promising architecture to obtain pre-trained representations that are well-suited for Conversational AI applications, especially for the Intent Classification task. Both DIET and ConveRT are lightweight architectures with faster training capabilities than their counterparts. For all the above reasons, we adapted the DIET architecture and incorpo-

rated pre-trained ConveRT embeddings to improve our Intent Classification performance (and later explore the Entity Recognition capabilities). We would call this approach DIET+ConveRT in our experiments¹.

Although the DM model development is beyond the scope of this work, we gradually migrated from the baseline Recurrent Embedding Domain Policy (REDP) (Vlasov et al., 2018) model, which is again inspired by the StarSpace algorithm (Wu et al., 2017) and used in our previous work (Sahay et al., 2019). We adopted a more recent and suitable Transformer Embedding Dialogue (TED) policy (Vlasov et al., 2019) architecture, where a transformer’s self-attention mechanism operates over the sequence of dialogue turns.

5.2. Paraphrase Generation

Data augmentation via paraphrase generation is an effective strategy that we explored in this study to improve the Intent Classification performance, particularly when we have limited original data to train our NLU models. With that motivation, we developed a data augmentation module via training a sequence-to-sequence paraphrasing model to generate paraphrased samples from the original seed utterances to augment the NLU training data. We propose several paraphrasing-based approaches and augmentation strategies to over-sample the intent classes and investigate their effects on the NLU performance. We examine the data augmentation via paraphrasing with a few simple heuristics (i.e., paraphrasing only for the low-sample intents or minority classes or excluding the intent types with samples having shorter utterance lengths). We also investigate model-in-the-loop data augmentation techniques (i.e., augmenting only the paraphrased utterances with successful predictions and checking the confidence level thresholds using the initial NLU models trained on original samples).

For effective paraphrase generation from the seed samples, we adapted the BART sequence-to-sequence model (Lewis et al., 2019) that we fine-tuned on the back-translated English sentences from the combination of following three datasets: the Microsoft Research Paraphrase (MSRP) corpus (Dolan and Brockett, 2005), ParaNMT corpora (Wieting and Gimpel, 2018; Wieting et al., 2017), and the PAWS dataset (Zhang et al., 2019; Yang et al., 2019). The MSRP (William B. Dolan and Chris Brockett, 2005) is a corpus containing 5800 pairs of sentences extracted from web news sources, along with human annotations indicating whether each pair captures a semantic equivalence/paraphrase relationship. The ParaNMT-50M (John Wieting and Kevin Gimpel, 2017) is a dataset of more than 50 million English-English sentential paraphrase pairs back-translated from the Czeng1.6 corpus (Bojar et al., 2016). The PAWS (Yuan Zhang,

¹Please refer to the Bunk et al. (2020) for hyperparameters, hardware specifications, and computational cost details.

Statistics/Dataset	Planting	Watering
# distinct intents	14	13
total # samples (utterances)	1927	2115
min # samples per intent	22	25
max # samples per intent	555	601
avg # samples per intent	137.6	162.7
# unique words (vocab)	1314	1267
total # words	10141	10469
min # words per sample	1	1
max # words per sample	74	65
avg # words per sample	5.26	4.95

Table 1: KidSpace NLU Dataset Statistics

Jason Baldrige, Luheng He, 2019) is a corpus containing 108,463 human-labeled and well-formed paraphrase pairs. Note that we also trained paraphrasing models using GPT-2 (Radford et al., 2019) fine-tuning to augment the training set. Finally, we decided to stick with the sequence-to-sequence paraphrasing model with BART fine-tuning as it performed slightly better than GPT-2 fine-tuning version, which is expected as the BART model can be seen as a generalized BERT (Devlin et al., 2018) encoder and GPT (Radford et al., 2018) decoder.

6. Experimental Results

6.1. Dataset

We conduct our experiments on the KidSpace NLU datasets having utterances from multimodal math learning experiences (i.e., Planting and Watering activities) designed for 5-to-8 years-old kids (Sahay et al., 2021; Aslan et al., 2022). These are the initial proof-of-concept (POC) datasets to bootstrap the agents to be deployed. These POC datasets are manually created based on User Experience (UX) design studies for training the SDS models and validated with UX sessions in the lab with multiple kids going through play-based learning activities. The NLU datasets have a limited number of utterances, which are manually annotated for intent types that we defined for each use-case or learning game/activity (see section 3). For the Flowerpot game, we have the Planting Flowers dataset with 1927 utterances, and for the NumberGrid activity, we have a separate Watering Flowers dataset with 2115 utterances. Some of our intents are highly generic across usages and activities (e.g., *affirm*, *deny*, *next_step*, *out_of_scope*, *goodbye*), whereas the rest are highly domain-dependent and task-specific (e.g., *intro_meadow*, *answer_flowers*, *answer_water*, *ask_number*, *counting*). Note that our dialogue system needs to process and interpret utterances received either from the kids or the adult (i.e., the Facilitator) present in the room. Therefore, almost half of our intent types are defined based on what the game flow expects from the Facilitator (e.g., to progress the games or guide the children). Table 1 shows the statistics of our NLU datasets.

Model/Dataset	Planting	Watering
TF+BERT (Baseline)	90.55	92.41
DIET+ConveRT	95.59	97.83
Performance Gain	+5.04	+5.42

Table 2: NLU/Intent Recognition F1-scores (%): Previous (TF+BERT) and Updated (DIET+ConveRT) Model Results (3 runs of 10-fold CV)

6.2. NLU Results

To evaluate the Intent Recognition performances, the baseline NLU model that we call TF+BERT is compared with the DIET+ConveRT model that we adapted recently (see section 5.1). We conduct these evaluations on both the Planting and Watering datasets. Table 2 summarizes the Intent Classification performance results in weighted average F1-scores. We perform 10-fold cross-validation (CV) over the dataset for each run, and we report the results based on the average of 3 runs. As we can observe from Table 2, adapting the lightweight DIET architecture (Bunk et al., 2020) with pre-trained ConveRT embeddings (Henderson et al., 2020) for our NLU models significantly improved the Intent Classification performance, which is consistent across different use-cases (i.e., Planting and Watering Flowers game datasets). With that observation, we have updated the NLU component in our multimodal SDS pipeline (see Figure 2) by replacing the TF+BERT model (i.e., previously best-performing Baseline + BERT + Transformer model in Sahay et al. (2019)) with this promising DIET+ConveRT model.

6.3. Data Augmentation

The goal here is to explore the potential benefits of paraphrasing-based data augmentation to further improve the Intent Recognition models. Our final BART-based Seq2Seq model² (see section 5.2) is utilized for paraphrase generation to conduct augmentation experiments. All paraphrasing experiments are conducted on the Planting Flowers dataset with a limited number of original seed utterances (i.e., less than 2K samples). We propose and investigate the following data augmentation methods with certain rule-based heuristics and model-in-the-loop (MITL) approaches:

- **baseline (aug3/aug5/aug10)**: Augment the original NLU dataset with paraphrased samples. We configured the number of samples to be generated as 3/5/10 (i.e., for each original utterance, x3/x5/x10 paraphrased samples are generated).
- **inc6low**: Augment data only for 6 low-sample intents (i.e., having less than 50 utterances). We simply used the original plus paraphrased samples for 6 intents with fewer number of ut-

terances (i.e., *intro_meadow*, *help_affirm*, *everyone_understand*, *oscar_understand*, *ask_number*, *next_step*), whereas we only used the original samples for rest of the 8 intents with higher number of utterances (no need for more variations for those).

- **exc5short**: Augment data except for 5 intents with seed samples having short utterance lengths. We only used the original samples for 5 intents with short utterances (i.e., *affirm*, *deny*, *answer_flowers*, *answer_valid*, *answer_invalid*), whereas we used the original plus paraphrased samples for rest of the 9 intents with longer utterances (as variations help for those).
- **success**: Augment only the paraphrased samples that are classified correctly into the same intent class as their seed samples (successful predictions). For this MITL approach, we first trained the NLU model (DIET+ConveRT) on the original/seed dataset, then classified the paraphrased samples using this initial NLU model to obtain successful predictions. We assume the paraphrased samples should belong to the same class as seed samples, and the idea is to filter out noisy synthetic samples that belong to other classes.
- **success_conf90**: Augment only the paraphrased samples that are classified correctly into the same intent class as their seed samples (successful predictions) with a confidence score of 0.9 or higher. Another MITL approach using the same initial NLU model as in *success*. The confidence check ensures better noise filtering, and the threshold of 0.9 is chosen empirically after checking the confidence histograms on paraphrased samples.
- **all_conf90**: Augment only the paraphrased samples that are classified into any intent type (regardless of their seed samples' intent class) with a confidence score of 0.9 or higher. Another MITL approach using the same initial NLU model as in *success*. We removed the assumption that paraphrased samples should be of the same class as their seed samples and still augment them into the predicted class samples if confidence is high.

Figure 3 depicts an example seed utterance and its paraphrase generation outputs for x5. In all data aug-

```

Enter text to paraphrase: Oscar has something he wants to say to you children
-----
Predictions >>>
Oscar has something to say to you, children.
Oscar has something he wants to tell you, kids.
He has something he wants to tell you, kids.
Oscar has something he wants to tell you, children.
Oscar wants to say something to you, kids.

```

Figure 3: Example seed and paraphrased utterances.

²<https://simpletransformers.ai/docs/seq2seq-model/>

Method	original	aug3	aug5	aug10
baseline	95.59	95.17	94.73	94.75
inc6low	-	95.84	96.06	96.41
exc5short	-	97.70	97.61	97.86
success	-	98.58	98.82	98.65
success_conf90	-	99.19	99.37	99.43
all_conf90	-	98.61	98.75	98.58
Perf. Gain	-	+3.60	+3.78	+3.84

Table 3: NLU/Intent Recognition F1-scores (%) with DIET+ConveRT Models Trained on Augmented Data via Paraphrasing

mentation experiments, we observed repetitions due to duplicate samples generated by the paraphraser, and we removed those duplicates from the final augmented datasets. Note that while augmenting the data, paraphrased samples are assumed to have the same intent labels as their seed samples they are generated from, except for the *all_conf90* case (for which we assign the labels predicted by the initial NLU model if the confidence is 0.9 or higher).

Table 3 summarizes the Intent Recognition performances in weighted-avg F1-scores after the data augmentation with paraphrased samples. We report the NLU results on an average of 3 runs and perform a 10-fold CV over original/augmented datasets for each run. In each such fold, the 10% test partition has the original samples only, whereas the 90% training partition is augmented with the paraphrased samples. This setup ensures we evaluate the models on the same original/seed samples only, while we can expand the training data with more variations. As shown in Table 3, the baseline approach of simply augmenting the data with paraphrased samples does not help but slightly hurt the NLU performances. That is due to the possible noises in synthetic data generated via paraphraser. However, with our proposed heuristics and MITL approaches, data augmentation helps improve our NLU results. When we augment for low-sample intents only (i.e., *inc6low*), we start observing slight improvements and increasing the number of paraphrased samples help. We get even better jumps when we exclude short-sample intents (i.e., *exc5short*). Augmenting only the paraphrased samples that are classified (with high confidence) into the same label as their seed samples (i.e., *success_conf90*) is the best performing approach, boosting the performance by nearly 4% (compared to training on original samples). Since we properly filter out noisy synthetic data in this case, generating more paraphrases beyond x3 and x5 also helps (i.e., $\text{aug10} > \text{aug5} > \text{aug3}$ for *success_conf90*).

6.3.1. Data Size versus NLU Performance

We analyzed the dataset sizes vs. Intent Recognition F1-scores for original and augmented datasets with paraphrased samples. Instead of performing a 10-fold




Figure 4: Paraphraser evaluation: data size vs. NLU performance for original and augmented datasets.

CV, we created a train/test split (i.e., 80/20%) from each dataset (i.e., original, aug3, aug5, aug10). Then we trained NLU models with 10, 20, ..., 100% of the training sets, each evaluated on the same test set, and the F1-scores compared. This process repeated for 3 runs (i.e., 3 test sets). In Figure 4, we show the plots for data size versus NLU performance with avg/std of F1-scores. Note that x-axis values (i.e., number of training samples) at each plot vary due to data augmentation. We observed that models trained on augmented datasets achieve a plateau of F1-scores faster with less original training data.

Next, we visualize the data size vs. performance with a superimposed chart for better comparison. We created a fixed train/test split (i.e., 80/20%) from the original data. Then, we used the same test set across all comparisons (models can be trained on original/augmented data but tested only on original samples). For training sets, we created 10, 20, ..., 100% of the original training set, and we augmented those with the paraphrased samples (aug3, aug5, aug10). In Figure 5, we plot the superimposed chart with common x-axis values (% of the original training set used) for comparison. We observed that we could reach a 0.8 F1-score with around 15% of the original samples with paraphrasing (aug10), whereas we need at least 35% of the data to achieve the same level of performance without paraphrasing (i.e.,




Figure 5: Paraphraser evaluation: data percentage vs. NLU performance for original and augmented datasets.

2.3x reduction in required original data). Similarly, we can reach a 0.9 F1-score with around 40% of the original set via paraphrasing (aug10), whereas we need at least 70% of the data to achieve the same level of performance without paraphrasing (i.e., 1.75x reduction in required data). We believe this paraphrasing approach would help us achieve better results with limited initial intent samples whenever there is a new use-case (e.g., future learning activities in Kid Space).

6.4. Entity Extraction

In addition to data augmentation, we aim to investigate the Entity Extraction and evaluate the potential improvements via entity expansion. The idea behind this is, if we can find a way to auto-extract the entities existing in the dataset, we can perform a lexical entity enrichment via ConceptNet (Liu and Singh, 2004; Speer et al., 2017) as an external Knowledge Graph (KG). Later on, we can also explore the use of lookup tables and synonyms in the dataset to create more variations via entity expansion on top of the original and paraphrased samples. With that motivation, we used a pre-trained SpaCy Entity Recognizer³ that performs Named Entity Recognition (NER) (McCallum and Li, 2003; Okur et al., 2016) to automatically extract the entities in our original NLU dataset (Planting Flowers). We then re-formatted the dataset with auto-tagged entities detected within utterances. We performed 3 runs of 10-fold CV on the original dataset with SpaCy NER tagged entities. We evaluated the joint Intent and Entity Recognition performances using DIET+ConveRT.

We observed that these auto-tagged entities generated via a pre-trained SpaCy NER model do not really help improve the NLU performances. We also realized that these generic named entity types are not very much relevant to our dataset. Hence, we had to define and extract more domain-specific entities, which requires a heavy task of word-level annotations. We completed these manual annotations for domain-specific entity types on the original dataset (Planting Flowers).

Table 4 summarizes our findings. Entity Recognition F1-score improved from 72.8% to 97.1% with these manual annotations, which is not surprising. Intent Classification F1-scores slightly drop when entities come into play, which aligns with the findings in Bunk et al. (2020). Although we can extract domain-specific entities pretty accurately, these token-level manual annotations are costly, even for small-size datasets. Next, we investigate auto-annotating the domain-specific entities using ConceptNet relatedness. We provide up to 6 sample values for each domain-specific entity types that we previously defined, then construct a synonym dictionary by returning the corresponding entities in the KG if the relatedness of the value is larger than an empirical threshold of 0.7. After extracting the tokens and noun chunks via SpaCy part-of-speech (POS) tagger⁴,

³<https://spacy.io/api/entityrecognizer>

⁴<https://spacy.io/api/tagger>

	Intent Classification	Entity Recognition
No entities	95.59	-
SpaCy NER tagged entities	94.70	72.82
Manually annotated entities	94.91	97.12
Auto-annotated entities	94.76	92.64

Table 4: NLU/Joint Intent Classification and Entity Recognition F1-scores (%): DIET+ConveRT Model Results (3 runs of 10-fold CV)

we automatically annotate the domain-specific entities. We do not expect this simplistic approach to work as accurately as manual annotations but want to see how much we can improve upon generic SpaCy NER tagged entities. Surprisingly, we achieved an Entity Recognition F1-score of 92.6% with the auto-annotated entities, which we believe is a very good compromise.

We believe these domain-specific entities would help us achieve lexical entity enrichment via ConceptNet as a knowledge graph on top of the original/paraphrased samples. We also believe these results are quite encouraging in our quest to make dialog systems more robust and generalizable to new intents with limited data.

7. Conclusion and Future Work

Constructing robust dialogue systems are critical to achieving efficient task-oriented communication with children in game-based learning settings. This study presents our multimodal dialogue system engaging with younger kids while learning basic math concepts. We focus on improving the NLU module of the task-oriented SDS pipeline with limited datasets. This exploration employs data augmentation with paraphrasing to increase NLU performances. Paraphrasing with model-in-the-loop strategies looks promising for achieving higher F1-scores for Intent Classification using small task-dependent datasets. Finally, we investigate the Entity Extraction to potentially further improve the NLU component of our multimodal SDS.

In future work, we plan to extend the Plug and Play Language Model (PPLM) (Dathathri et al., 2019) architecture applicable to the Decoder-only Unconditional Language Models (such as GPT-2) to Seq2Seq Encoder-Decoder based Conditional Language Models (CLM) (Keskar et al., 2019) (such as BART), where the text to be generated is constrained by a cross-attention to the Encoder input. We can further control this CLM with controllable attributes that require no training/fine-tuning of the model. During inference, control attributes directly update the latent activations to steer the model to generate fluent and attribute-specific text. We can explore the PPLM approach to get more paraphrased samples using entity expansion via ConceptNet and adapt this approach to Seq2Seq Encoder-Decoder models.

8. Acknowledgements

We show our gratitude to our current and former colleagues from the Intel Labs Kid Space team, especially Ankur Agrawal, Glen Anderson, Sinem Aslan, Arturo Bringas Garcia, Rebecca Chierichetti, Hector Cordourier Maruri, Pete Denman, Lenitra Durham, Roddy Fuentes Alba, David Gonzalez Aguirre, Sai Prasad, Giuseppe Raffa, Sangita Sharma, and John Sherry, for the conceptualization and the design of use-cases to support this research. The authors are also immensely grateful to the Intel Labs KAIU team, specifically Nagib Hakim, Ezequiel Lanza, and Gadi Singer, for their feedback and support on entity annotation tasks in collaboration with our team. Finally, we thankfully acknowledge the Rasa team for the open-source framework and the community developers for their contributions that enabled us to improve our research and build proof-of-concept models for our use-cases.

9. Bibliographical References

- Anderson, G. J., Panneer, S., Shi, M., Marshall, C. S., Agrawal, A., Chierichetti, R., Raffa, G., Sherry, J., Loi, D., and Durham, L. M. (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.
- Andreas, J., Bufe, J., Burkett, D., Chen, C., Clausman, J., Crawford, J., Crim, K., DeLoach, J., Dorner, L., Eisner, J., Fang, H., Guo, A., Hall, D., Hayes, K., Hill, K., Ho, D., Iwaszuk, W., Jha, S., Klein, D., Krishnamurthy, J., Lanman, T., Liang, P., Lin, C. H., Lintsbakh, I., McGovern, A., Nisnevich, A., Pauls, A., Petters, D., Read, B., Roth, D., Roy, S., Rusak, J., Short, B., Slomin, D., Snyder, B., Striplin, S., Su, Y., Tellman, Z., Thomson, S., Vorobev, A., Witoszko, I., Wolfe, J., Wray, A., Zhang, Y., and Zotov, A. (2020). Task-Oriented Dialogue as Dataflow Synthesis. *Transactions of the Association for Computational Linguistics*, 8:556–571, 09.
- Aslan, S., Agrawal, A., Alyuz, N., Chierichetti, R., Durham, L. M., Manuvinakurike, R., Okur, E., Sahay, S., Sharma, S., Sherry, J., 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.
- Bahdanau, D., Cho, K., and Bengio, Y. (2014). Neural machine translation by jointly learning to align and translate. *CoRR*, abs/1409.0473.
- Barzilay, R. and McKeown, K. R. (2001). Extracting paraphrases from a parallel corpus. In *Proceedings of the 39th Annual Meeting of the Association for Computational Linguistics*, pages 50–57, Toulouse, France, July. Association for Computational Linguistics.
- Bocklisch, T., Faulkner, J., Pawlowski, N., and Nichol, A. (2017). Rasa: Open source language understanding and dialogue management.
- Bojar, O., Dušek, O., Kocmi, T., Libovický, J., Novák, M., Popel, M., Sudarikov, R., and Variš, D. (2016). Czeng 1.6: enlarged czech-english parallel corpus with processing tools dockered. In *International Conference on Text, Speech, and Dialogue*, pages 231–238. Springer.
- Bordes, A., Boureau, Y.-L., and Weston, J. (2016). Learning end-to-end goal-oriented dialog. *arXiv preprint arXiv:1605.07683*.
- Bowman, S. R., Vilnis, L., Vinyals, O., Dai, A., Jozefowicz, R., and Bengio, S. (2016). Generating sentences from a continuous space. In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 10–21, Berlin, Germany, August. Association for Computational Linguistics.
- Bunk, T., Varshneya, D., Vlasov, V., and Nichol, A. (2020). DIET: lightweight language understanding for dialogue systems. *CoRR*, abs/2004.09936.
- Burtsev, M., Seliverstov, A., Airapetyan, R., Arkhipov, M., Baymurzina, D., Bushkov, N., Gureenkova, O., Khakhulin, T., Kuratov, Y., Kuznetsov, D., Litinsky, A., Logacheva, V., Lymar, A., Malykh, V., Petrov, M., Polulyakh, V., Pugachev, L., Sorokin, A., Vikhreva, M., and Zaynutdinov, M. (2018). DeepPavlov: Open-source library for dialogue systems. In *Proceedings of ACL 2018, System Demonstrations*, pages 122–127, Melbourne, Australia, July. Association for Computational Linguistics.
- Cahill, A., Fife, J. H., Riordan, B., Vajpayee, A., and Galochkin, D. (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, July. Association for Computational Linguistics.
- Cuayáhuitl, H. (2017). Simpleds: A simple deep reinforcement learning dialogue system. In *Dialogues with Social Robots*, pages 109–118. Springer.
- Dathathri, S., Madotto, A., Lan, J., Hung, J., Frank, E., Molino, P., Yosinski, J., and Liu, R. (2019). Plug and play language models: A simple approach to controlled text generation. *arXiv preprint arXiv:1912.02164*.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Dhingra, B., Li, L., Li, X., Gao, J., Chen, Y.-N., Ahmed, F., and Deng, L. (2016). Towards end-to-end reinforcement learning of dialogue agents for information access. *arXiv preprint arXiv:1609.00777*.
- Dodge, J., Gane, A., Zhang, X., Bordes, A., Chopra, S., Miller, A., Szlam, A., and Weston, J. (2015). Evaluating prerequisite qualities for learning end-to-end dialog systems. *arXiv preprint arXiv:1511.06931*.

- Dolan, W. B. and Brockett, C. (2005). Automatically constructing a corpus of sentential paraphrases. In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*.
- Goo, C.-W., Gao, G., Hsu, Y.-K., Huo, C.-L., Chen, T.-C., Hsu, K.-W., and Chen, Y.-N. (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, June. Association for Computational Linguistics.
- Gu, Y., Li, X., Chen, S., Zhang, J., and Marsic, I. (2017). Speech intention classification with multimodal deep learning. In *Canadian conference on artificial intelligence*, pages 260–271. Springer.
- Hakkani-Tur, D., Tur, G., Celikyilmaz, A., Chen, Y.-N. V., Gao, J., Deng, L., and Wang, Y.-Y. (2016). Multi-domain joint semantic frame parsing using bi-directional rnn-1stm. ISCA, June.
- Henderson, M., Casanueva, I., Mrkšić, N., Su, P.-H., Wen, T.-H., and Vulić, I. (2020). ConveRT: Efficient and accurate conversational representations from transformers. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2161–2174, Online, November. Association for Computational Linguistics.
- Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Hu, J. E., Singh, A., Holzenberger, N., Post, M., and Van Durme, B. (2019). Large-scale, diverse, paraphrastic bitexts via sampling and clustering. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 44–54, Hong Kong, China, November. Association for Computational Linguistics.
- Jia, J., He, Y., and Le, H. (2020). A multimodal human-computer interaction system and its application in smart learning environments. In Simon K. S. Cheung, et al., editors, *Blended Learning. Education in a Smart Learning Environment*, pages 3–14, Cham. Springer International Publishing.
- Jolly, S., Falke, T., Tirkaz, C., and Sorokin, D. (2020). Data-efficient paraphrase generation to bootstrap intent classification and slot labeling for new features in task-oriented dialog systems. In *Proceedings of the 28th International Conference on Computational Linguistics: Industry Track*, pages 10–20, Online, December. International Committee on Computational Linguistics.
- Jurafsky, D. and Martin, J. H. (2018). *Ch 24: Dialog Systems and Chatbots. Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. Prentice Hall PTR, Upper Saddle River, NJ, USA, 3rd (draft) edition.
- Keskar, N. S., McCann, B., Varshney, L. R., Xiong, C., and Socher, R. (2019). CTRL: A conditional transformer language model for controllable generation. *CoRR*, abs/1909.05858.
- Kumar, V., Glaude, H., de Lichy, C., and Campbell, W. (2019). A closer look at feature space data augmentation for few-shot intent classification. In *Proceedings of the 2nd Workshop on Deep Learning Approaches for Low-Resource NLP (DeepLo 2019)*, pages 1–10, Hong Kong, China, November. Association for Computational Linguistics.
- Kumar, V., Choudhary, A., and Cho, E. (2020). Data augmentation using pre-trained transformer models. In *Proceedings of the 2nd Workshop on Life-long Learning for Spoken Language Systems*, pages 18–26, Suzhou, China, December. Association for Computational Linguistics.
- Lan, W., Qiu, S., He, H., and Xu, W. (2017). A continuously growing dataset of sentential paraphrases. *CoRR*, abs/1708.00391.
- Lee, K., Guu, K., He, L., Dozat, T., and Chung, H. W. (2021). Neural data augmentation via example extrapolation. *CoRR*, abs/2102.01335.
- Lende, S. P. and Raghuvanshi, M. (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.
- Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., and Zettlemoyer, L. (2019). BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *CoRR*, abs/1910.13461.
- Liu, B. and Lane, I. R. (2016). Attention-based recurrent neural network models for joint intent detection and slot filling. *CoRR*, abs/1609.01454.
- Liu, H. and Singh, P. (2004). Conceptnet—a practical commonsense reasoning tool-kit. *BT technology journal*, 22(4):211–226.
- Liu, B., Tur, G., Hakkani-Tur, D., Shah, P., and Heck, L. (2017). End-to-end optimization of task-oriented dialogue model with deep reinforcement learning. In *NIPS Workshop on Conversational AI*.
- Liu, X., Eshghi, A., Swietojanski, P., and Rieser, V., (2021). *Benchmarking Natural Language Understanding Services for Building Conversational Agents*, pages 165–183. Springer Singapore, Singapore.
- McCallum, A. and Li, W. (2003). Early results for named entity recognition with conditional random fields, feature induction and web-enhanced lexicons.
- Mesnil, G., Dauphin, Y., Yao, K., Bengio, Y., Deng, L., Hakkani-Tur, D., He, X., Heck, L., Tur, G., Yu, D., and Zweig, G. (2015). Using recurrent neural networks for slot filling in spoken language understanding. *Trans. Audio, Speech and Lang. Proc.*, 23(3):530–539, March.
- Meurers, D. (2012). Natural language processing and

- language learning. *Encyclopedia of applied linguistics*, pages 4193–4205.
- Okur, E., Demir, H., and Özgür, A. (2016). Named entity recognition on twitter for turkish using semi-supervised learning with word embeddings. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 549–555.
- Okur, E., Kumar, S. H., Sahay, S., Esme, A. A., and Nachman, L. (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)*, April.
- Okur, E., H Kumar, S., Sahay, S., and Nachman, L. (2020). Audio-visual understanding of passenger intents for in-cabin conversational agents. In *Second Grand-Challenge and Workshop on Multimodal Language (Challenge-HML)*, pages 55–59, Seattle, USA, July. Association for Computational Linguistics.
- Panda, S., Tirkaz, C., Falke, T., and Lehnen, P. (2021). Multilingual paraphrase generation for bootstrapping new features in task-oriented dialog systems. In *Proceedings of the 3rd Workshop on Natural Language Processing for Conversational AI*, pages 30–39, Online, November. Association for Computational Linguistics.
- Pennington, J., Socher, R., and Manning, C. D. (2014). Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Radford, A., Narasimhan, K., Salimans, T., and Sutskever, I. (2018). Improving language understanding by generative pre-training.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al. (2019). Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Sahay, S., Kumar, S. H., Okur, E., Syed, H., and Nachman, L. (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, September. SEMDIAL.
- Sahay, S., Okur, E., Hakim, N., and Nachman, L. (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, June. Association for Computational Linguistics.
- Schuster, M. and Paliwal, K. (1997). Bidirectional recurrent neural networks. *Trans. Sig. Proc.*, 45(11):2673–2681, November.
- Sennrich, R., Haddow, B., and Birch, A. (2016). Improving neural machine translation models with monolingual data. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96, Berlin, Germany, August. Association for Computational Linguistics.
- Shah, P., Hakkani-Tur, D., and Heck, L. (2016). Interactive reinforcement learning for task-oriented dialogue management.
- Shukla, S., Liden, L., Shayandeh, S., Kamal, E., Li, J., Mazzola, M., Park, T., Peng, B., and Gao, J. (2020). Conversation Learner - a machine teaching tool for building dialog managers for task-oriented dialog systems. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 343–349, Online, July. Association for Computational Linguistics.
- Skene, K., O’Farrelly, C., Byrne, E., Kirby, N., Stevens, E., and Ramchandani, P. (2021). Can guidance during play enhance children’s learning and development in educational contexts? a systematic review and meta-analysis. *Child Development*.
- Socher, R., Huang, E. H., Pennington, J., Ng, A. Y., and Manning, C. D. (2011). Dynamic pooling and unfolding recursive autoencoders for paraphrase detection. In *Proceedings of the 24th International Conference on Neural Information Processing Systems, NIPS’11*, page 801–809, Red Hook, NY, USA. Curran Assoc. Inc.
- Sokolov, A. and Filimonov, D. (2020). Neural machine translation for paraphrase generation. *arXiv preprint arXiv:2006.14223*.
- Speer, R., Chin, J., and Havasi, C. (2017). Conceptnet 5.5: An open multilingual graph of general knowledge. In *Thirty-first AAAI conference on artificial intelligence*.
- Su, P.-H., Budzianowski, P., Ultes, S., Gasic, M., and Young, S. (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, August. Association for Computational Linguistics.
- Sutskever, I., Vinyals, O., and Le, Q. V. (2014). Sequence to sequence learning with neural networks. *CoRR*, abs/1409.3215.
- Taghipour, K. and Ng, H. T. (2016). A neural approach to automated essay scoring. In *Proceedings of the 2016 conference on empirical methods in natural language processing*, pages 1882–1891.
- Ultes, S., Rojas Barahona, L. M., Su, P.-H., Vandyke, D., Kim, D., Casanueva, I., Budzianowski, P., Mrkšić, N., Wen, T.-H., Gasic, M., and Young, S. (2017). PyDial: A multi-domain statistical dialogue system toolkit. In *Proceedings of ACL 2017, System Demonstrations*, pages 73–78, Vancouver, Canada, July. Association for Computational Linguistics.
- Vanzo, A., Bastianelli, E., and Lemon, O. (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, September. Association for Computational Linguistics.
- Varghese, A. S., Sarang, S., Yadav, V., Karotra, B., and Gandhi, N. (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.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. (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.
- Vlasov, V., Drissner-Schmid, A., and Nichol, A. (2018). Few-shot generalization across dialogue tasks. *arXiv preprint arXiv:1811.11707*.
- Vlasov, V., Mosig, J. E. M., and Nichol, A. (2019). Dialogue transformers. *CoRR*, abs/1910.00486.
- Wen, L., Wang, X., Dong, Z., and Chen, H. (2018). Jointly modeling intent identification and slot filling with contextual and hierarchical information. In Xuanjing Huang, et al., editors, *Natural Language Processing and Chinese Computing*, pages 3–15, Cham. Springer International Publishing.
- Wieting, J. and Gimpel, K. (2018). ParaNMT-50M: Pushing the limits of paraphrastic sentence embeddings with millions of machine translations. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 451–462, Melbourne, Australia, July. Association for Computational Linguistics.
- Wieting, J., Mallinson, J., and Gimpel, K. (2017). Learning paraphrastic sentence embeddings from back-translated bitext. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 274–285, Denmark, September. Association for Computational Linguistics.
- Wu, L., Fisch, A., Chopra, S., Adams, K., Bordes, A., and Weston, J. (2017). Starspace: Embed all the things! *arXiv preprint arXiv:1709.03856*.
- Xia, C., Xiong, C., Yu, P., and Socher, R. (2020). Composed variational natural language generation for few-shot intents. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3379–3388, Online, November. Association for Computational Linguistics.
- Yang, Y., Zhang, Y., Tar, C., and Baldrige, J. (2019). PAWS-X: A Cross-lingual Adversarial Dataset for Paraphrase Identification. In *Proc. of EMNLP*.
- Zhang, X. and Wang, H. (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.
- Zhang, Y., Baldrige, J., and He, L. (2019). PAWS: Paraphrase Adversaries from Word Scrambling. In *Proc. of NAACL*.
- Zhou, Q., Wen, L., Wang, X., Ma, L., and Wang, Y. (2016). A hierarchical lstm model for joint tasks. In Maosong Sun, et al., editors, *Chinese Computational Linguistics and Natural Language Processing Based on Naturally Annotated Big Data*, pages 324–335, Cham. Springer International Publishing.

10. Language Resource References

- John Wieting and Kevin Gimpel. (2017). *ParaNMT-50M: Pushing the Limits of Paraphrastic Sentence Embeddings with Millions of Machine Translations*.
- William B. Dolan and Chris Brockett. (2005). *The Microsoft Research Paraphrase Corpus (MSRP)*. Microsoft Shared Source License.
- Yuan Zhang, Jason Baldrige, Luheng He. (2019). *PAWS: Paraphrase Adversaries from Word Scrambling*. Google LLC.