

# Generative Adversarial Networks for Annotated Data Augmentation in Data Sparse NLU

**Olga Golovneva**

Amazon / Cambridge, MA  
olggol@amazon.com

**Charith Peris**

Amazon / Cambridge, MA  
perisc@amazon.com

## Abstract

Data sparsity is one of the key challenges associated with model development in Natural Language Understanding (NLU) for conversational agents. The challenge is made more complex by the demand for high quality annotated utterances commonly required for supervised learning, usually resulting in weeks of manual labor and high cost. In this paper, we present our results on boosting NLU model performance through training data augmentation using a sequential generative adversarial network (GAN). We explore data generation in the context of two tasks, the bootstrapping of a new language and the handling of low resource features. For both tasks we explore three sequential GAN architectures, one with a token-level reward function, another with our own implementation of a token-level Monte Carlo rollout reward, and a third with sentence-level reward. We evaluate the performance of these feedback models across several sampling methodologies and compare our results to up-sampling the original data to the same scale. We further improve the GAN model performance through the transfer learning of the pre-trained embeddings. Our experiments reveal synthetic data generated using the sequential generative adversarial network provides significant performance boosts across multiple metrics and can be a major benefit to the NLU tasks.

## 1 Introduction

Over recent years, various task-oriented conversational agents, such as Amazon Alexa, Apple’s Siri, Google Assistant, and Microsoft’s Cortana, have become more popular in people’s everyday life and are expected to be highly intelligent. For the NLU component, this means that we expect models to perform recognition of the actions and entities within a user’s request with high accuracy.

When first training an NLU model on a new language (a process referred to as bootstrapping a new language), there is a strong requirement for high quality annotated data that would support the most common user requests across a range of domains. As the modeling space expands to support new features and additional languages, NLU models are regularly re-trained on updated data sets to ensure support for these new functions. The major bottleneck in both of these processes is the labor and cost associated with collecting and annotating new training utterances for every new feature or language.

Recent advances in machine learning methods, including the use of techniques such as transfer learning (Lu et al., 2015) and active learning (Settles, 2009), can lead to more efficient data usage by NLU models and therefore decrease the need for annotated training data. Additionally, data augmentation models are being widely explored. The advantage of data augmentation is that once synthetic data is generated, it can be ingested into subsequent models without additional effort, allowing for faster experimentation.

NLU models in dialog systems can perform a variety of tasks (Ram et al., 2018; Gao et al., 2018). In this study, we will focus on three of them: **Domain classification (DC)** – identify the domain that the user request belongs to (music, reminders, alarm, etc.), **Intent classification (IC)** – extract actions requested by users (play music, find a restaurant, set an alarm, etc.), and **Named Entity Recognition (NER)** – identify and extract entities (names, values, dates, locations, etc.) from user requests.

For each utterance we expect our NLU model to output a domain, intent, and set of extracted entities with corresponding tags. For example, if a user requests “*play Bohemian Rhapsody by Queen*”, we expect the NLU model to return **{domain: music, intent: play\_song, named\_entities: [(bohemian**

rhapsody, *song\_name*), (queen, *artist\_name*)]}. We call this output *annotation*, and the utterance along with annotation is called an *annotated utterance*. Named entities with corresponding labels are called *slots*.

For our NLU model to perform well on real-time user requests, we need to train it on a large dataset of diverse annotated utterances. However, there could be some areas of functionality where large datasets for training are not available. To boost model performance in situations where training data is limited, we use synthetic data generated from a small set of unique utterances that cover the basic functionality of the user experience, called *Golden utterances*. We leverage a Sequence Generative Adversarial Networks (SeqGAN) introduced by Yu et al. (2017) to generate new utterances from this “seed” set, and use these generated utterances to augment training data and evaluate the performance of the classification and recognition tasks. We also investigate how the metrics that we use to evaluate the quality of the generated synthetic data links to the performance boost in the underlying tasks.

## 2 Related work

NLU model boosting through training data augmentation has been an active area of research over the last few years, with more sophisticated techniques and models being developed. Some of these techniques include data resampling, the use of Variational Autoencoders (VAEs) and GANs. Xie et al. (2017) generalize resampling methods by proposing noising schemes that are designed to smooth input data by randomly changing the word tokens in a sentence. First described by Kingma and Welling (2013), VAEs learn distributed representations of latent variables, and decode random samples to generate data that have similar characteristics to those that the network was trained on. GAN model proposed by Goodfellow et al. (2014) includes two competing neural networks: a generator that creates fake data, and a discriminator that is trained to distinguish between fake and real data. The generator is trained on the results of its success in fooling the discriminator and this contest results in synthetic data that is progressively more similar to real data.

Synthetic data have shown to be useful for IC model boosting. For example, Malandrakis et al. (2019) explored a set of encoder-decoder models

and proposed the use of conditional VAEs (CVAEs) to generate phrase templates, called carrier phrases. Authors used CVAEs to control the domain, intent, and slot types to generate desirable outputs that resulted in a higher F1 score on the intent classification task.

Kumar et al. (2019) focused on a few-shot IC problem where new categories with limited training data are introduced into an existing system with mature categories. They compared different techniques that were designed to augment training data, including upsampling, random perturbation, extrapolation, CVAEs, and delta-encoders, and combined feature space augmentation with popular BERT pre-training (Devlin et al., 2019) to provide better performance.

The use of GANs has been previously explored for text data augmentation in language modeling (Kusner and Hernández-Lobato, 2016; Yu et al., 2017; Che et al., 2017; Guo et al., 2018; Hu et al., 2017; Li et al., 2017; Lin et al., 2017; Zhang et al., 2017; Fedus et al., 2018) and sentiment classification (Gupta, 2019). However, discrete text sequence generation brings about several challenges: first, one needs to generate a set of discrete tokens from a random sample of real-valued continuous data, and second, GANs are designed to give feedback on entire sequences, whereas generators need guidance for each subsequent token. The SeqGAN model developed by Yu et al. (2017) attempts to resolve these issues by applying reinforcement algorithms for the GAN objective with a policy gradient that evaluates current state-action value using Monte Carlo (MC) search. In this work, we adopt a SeqGAN model to boost DC, IC, and NER tasks in NLU models that suffer from sparse data limitations.

## 3 Methods

### 3.1 Data

For our experiments, we used the English data<sup>1</sup> collected by Schuster et al. (2019). This data was consisted of three domains: *weather*, *alarm*, and *reminder*, and a total of 43000 utterances. It was collected in a three-step process: step 1 consisted of native English speakers producing utterances for each intent, step 2 consisted of two annotators labeling the intents and slots while any conflicts between these two annotators were resolved in step 3 by a third annotator. The data was processed

<sup>1</sup>[https://fb.me/multilingual\\_task\\_oriented\\_data](https://fb.me/multilingual_task_oriented_data)

further to match the format that was suitable for our models.

### 3.2 Models

Text data boosting in NLU was extensively used for classification tasks, so most previous research focused on generating sentences (Kumar et al., 2019), carrier phrases (Malandrakis et al., 2019), or embeddings (Guo et al., 2018). In our work we consider DC, IC, and NER problems, where both sequence and word tags are needed for model training. We leverage a GAN to synthesize training data as a sequence of intents and slots:  $X = \{x_0, x_1, \dots, x_n\}$ , where  $n$  varies between a length of 1 and the maximum allowed utterance length. Each slot  $x_i (i > 0)$  denotes the combination of the  $i^{\text{th}}$  word and its corresponding tag, and  $x_0$  is a concatenation of the utterance domain and intent. For example, for the utterance “*play Bohemian Rhapsody by Queen*”, the training sequence for text generation would be as follows: “*music/play\_song play:none bohemian:song\_name rhapsody:song\_name by:none queen:artist\_name*”.

#### 3.2.1 SeqGAN model

When applied to text data, traditional GANs have difficulty performing back-propagation due to the non-differentiable output of the generator model. The SeqGAN model addresses this issue by treating the generator as a reinforcement learning agent that optimizes the GAN objective. The discriminator itself is used within the reward function to evaluate output sequences and return feedback to guide the learning of the generative model. Traditional reward functions from classification models are also limited by the ability to only provide score/loss-based reward values for a complete sequence. The SeqGAN model enables evaluation of the action-value for an intermediate state of unfinished sequence for each initial state,  $s_0$ , by applying an MC search with a rollout policy to sample the unknown last tokens (Yu et al., 2017). MC search is a tree-search algorithm with a root node  $s_0$ . Each child node of the tree is drawn from a distribution parametrized by a stochastic parametrized policy. This policy can be any, for example current state of the generator  $G_\theta$ .

In our experiments, we compare different ways to compute the reward function (Figure 1). First, we use an implementation where the MC tree search strategy is replaced with token-level reward produced by the discriminator (Xu et al., 2018; Hu

et al., 2018). We use the output of the Long Short-Term Memory (LSTM)-based discriminator model  $D_\phi^L$ , cross-entropy, as the reward. For a synthetic sentence  $Y = \{y_0, y_1, \dots, y_n\}$ ,  $y_i \in \mathcal{Y}$ , where  $\mathcal{Y}$  is the vocabulary of candidate domain-intent and token-label pairs, the cross-entropy based reward for the  $i^{\text{th}}$  word is calculated as:

$$R(y_i) = -\log D_\phi^L(y_i|y_{<i}) \quad (1)$$

Next, we devise our own MC search-based method to produce a set of possible sequences to approximate the expectation of the token-level reward, rather than using a single evaluation of only one of the possible sequences. Using the current state of the generator model  $G_\theta$  as a stochastic rollout policy, for each incomplete sequence  $\{y_0, y_1, \dots, y_k\}$  we use Monte-Carlo search to produce  $N$  complete sequences. We evaluate the token-level reward for each of these sequences:

$$\{y_{k+1}, \dots, y_n\} \in MC^{G_\theta}(\{y_0, \dots, y_k\}; N), \quad (2)$$

for  $k < n$ . To approximate the expected value of the reward function we average the token-level reward (1) calculated for the  $N$  MC rollouts:

$$R(y_i) = \frac{1}{(n+1)N} \sum_{k=0}^n \sum_{j=1}^N (-\log D_\phi^L(y_i|y_{<i}))_{j,k}, \quad (3)$$

where  $(\dots)_{j,k}$  is the  $j^{\text{th}}$  MC rollout of the  $k^{\text{th}}$  sequence. Finally, we compare results with an expectation of a sentence-level reward calculated using MC tree rollout strategy for the convolutional neural network (CNN)-based discriminator feedback  $D_\phi^C$ , that provides feedback on the full sentence (Yu et al., 2017). Intermediate token-level reward for current token  $y_i$  is approximated using MC rollouts as follows:

$$R_{MC}(y_i) = \frac{1}{N} \sum_{j=1}^N (D_\phi^C(Y))_j, i < n, \quad (4)$$

where the sampling of the missing  $(n-i)$  tokens is governed by the equation (2). At each step  $i$ , generator reward is governed by the Monte-Carlo approximation (4), or discriminator feedback on the full sentence:

$$R(y_i) = \begin{cases} R_{MC}(y_i), i < n, \\ D_\phi^C(Y), i = n. \end{cases} \quad (5)$$

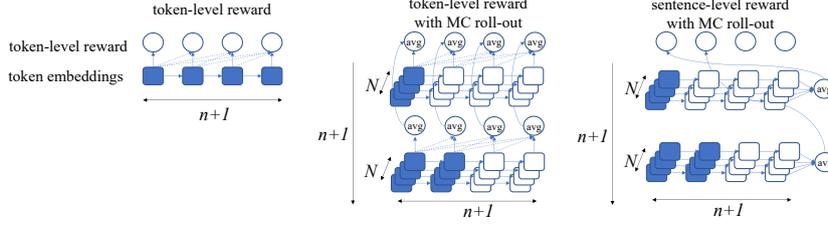


Figure 1: Schematic representation of the reward function produced by discriminator model for each of the  $n + 1$  input tokens and  $N$  Monte-Carlo rollouts ( $avg = average$ ). Tokens fixed at each iteration are represented by solid squares and tokens produced by the rollout policy are represented by blank squares. Sentence-level reward produces feedback on a complete sentence only, whereas token-level reward gives feedback on each token conditioned on previous set of tokens.

The generator is trained to maximize the reward from the discriminator model for the full sequence (Sutton et al., 2000):

$$J(\theta) = \mathbb{E}[R(Y|s_0, \theta)] = \sum_{y_0 \in \mathcal{Y}} G_\theta(y_0|s_0) Q_{D_\phi}^{G_\theta}(s_0, y_0) \quad (6)$$

where  $Q_{D_\phi}^{G_\theta}(s_0, y_0)$  can be estimated by  $R(y_0)$  using equations (1), (3) or (5).

The pre-training procedure for the generative decoder uses the maximum likelihood estimation metric, followed by the pre-training of the discriminative classifier on positive samples from the training data and negative samples produced by the generator. After pre-training, the generative and discriminative models are trained one at a time in the following loop: the last discriminator state is used to provide feedback for the generative model training, which then provides a new set of negative examples for the discriminator updates.

### 3.2.2 DC and IC/NER models

Our model architecture consists of two parts: a sentence classification model for the DC task, and a domain-specific joint model for the IC and NER tasks. The DC model is the same as the IC part of the joint IC-NER model, so we will describe only the latter in detail. The IC-NER model components are schematically shown in Figure 2, and are composed of the following:

- **Embedding:** concatenation of word embedding with 256 dimensions and character embedding with 16 dimensions trained on a 1-filter CNN with a  $\tanh$  activation function and dropout.
- **Encoder:** 2-layer bidirectional LSTM with 384 dimensions in each hidden layer with

dropout and layer normalization for token encoder, and a pooling stack for sequence encoder.

- **Decoders:** MLP classifier with 256 dimensions for IC task and CRF sequence labeler with 192 dimensions. For each block we also apply ELU activation function and dropout. All IC-NER models were trained for 500 epochs, and DC models for 100 epochs.

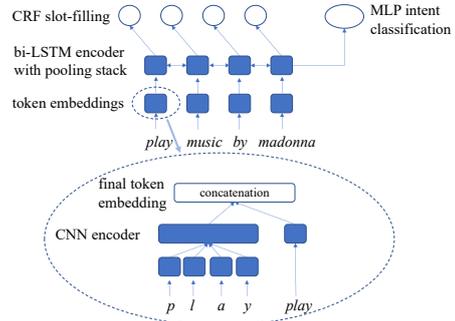


Figure 2: Multi-task model for IC and NER prediction.

### 3.3 Experimental setup

We build our experiments to mimic two major tasks that are of interest to us (Figure 3):

1. Bootstrapping of a new language
2. Handling low resource features

We manipulate our English data set to simulate modeling conditions in situations where training data is limited. For our tasks, we need to identify a set of utterances from within our data set that can be classified as Golden utterances, i.e., utterances that are found to be usually common

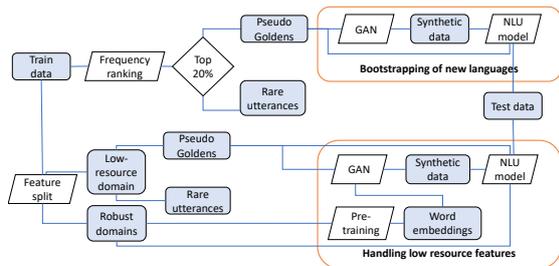


Figure 3: This schematic represents our data and experimental process. For our experiments we use annotated data collected by Schuster et al. (2019). Train and validation sets were combined together in a single train data set to extract pseudo Golden utterances. Pseudo Golden were further used to train and evaluate the GAN model. The test set was untouched and was used to evaluate model performance in full scale. To simulate the low resource feature task, we used the full train data set for the two robust domains, and the pseudo Golden to represent the third low-resource domain.

across languages and cover the basic functionality of the user experience. To select Golden from our data set we use the following process. We sort the utterances in our training data by the utterance frequency within each domain-intent combination. We then select the top 20% of utterances within each domain-intent combination, unique them, and call them *pseudo Golden*s. The unique utterances within the remaining set are called *Rare* utterances. Utterance counts are provided in Appendix A.

### 3.3.1 Bootstrapping NLU models for new languages

When bootstrapping an NLU model to support a new language for a conversational artificial intelligence (AI), there is usually a very limited data available. One major data source is Golden translated from another existing robust language model. In mimicking this task, we use our determined pseudo Golden as input for both data generation and NLU models. Synthetic data generated with the GAN model is added to the set of pseudo Golden and fed into our NLU model. The NLU models are tested using our fixed original test data set (Figure 3). Our baseline for comparison consists of an NLU model run on only the pseudo Golden.

### 3.3.2 Handling low resource features

Another common challenge faced by conversational AI NLU models is sparse data. When a feature is new, or not common, we do not have

enough data for the NLU model to generalize well on possible request variations. But, unlike the previous task, we typically do have a large amount of data collected in the same language, but within different intents and domains. Generative adversarial networks, paired with a transfer learning approach, can give us the opportunity to use other domains and intents which have robust data, to strengthen the performance of these low resource features. In this case we limit our experiments to exploring a case where we have a low resource domain. In mimicking this task, we pick one of the three domains we have as the low resource domain. We consider only the pseudo Golden as available data for this domain. The remaining two domains are considered as robust domains and we consider all data (i.e., both pseudo Golden and Rare utterances) as available data for the NLU model. We also use all training data from the two robust domains to pre-train word embedding using the fastText algorithm (Lample et al., 2017). We then run data generation for the low resource domain using our SeqGAN implementation while using the pre-trained word embeddings from the two robust domains to initialize word embeddings both in generator and discriminator. For the NLU model, we feed the pseudo Golden together with the synthesized data for the low resource domain and feed all available training data for the two robust domains. We measure the performance of the low resource domain against a baseline which is run using only pseudo Golden from the low resource domain together with all available data from the robust domains. As with task 1, performance is tested on the separate fixed test data set.

## 4 Data Generation

We explore the following different GAN frameworks and expansions to generate synthetic annotated utterances:

- **Original implementation with Monte Carlo rollout:** a selected set of experiments to benchmark the original implementation by Yu et al. (2017) on our open source data set.
- **Original implementation without rollouts:** a selected set of experiments to benchmark the original implementation by Yu et al. (2017), where the reward function is evaluated on a single MC rollout.

- **Generator with token-level reward:** training the generator using token-level reward as suggested by [Hu et al. \(2018\)](#).
- **Generator with token-level Monte Carlo rollout:** we expand the above implementation to include a token-level Monte Carlo rollout and test our task of bootstrapping new languages.
- **Generator with pre-trained embeddings:** we add fastText pre-trained embeddings to the generator and discriminator to selectively test our task of handling low-resource features).

The model architecture consists of three parts: generator pre-training, discriminator pre-training, and adversarial training. In the pre-training parts, we first train the generator, followed by the discriminator, for 80 epochs each. For each adversarial epoch, we update the generator once and then for 35 steps we generate negative examples using the current state of the generator, combined with the same number of positive examples from the training data, and re-train the discriminator. The total number of adversarial training epochs is set to 600. All hyper-parameters were chosen based on observations of when the loss functions and synthetic data quality evaluations either stabilize or clearly degrade.

## 5 Results and Discussion

### 5.1 Evaluation

We evaluate our models using domain accuracy, intent accuracy, slot F1, and frame accuracy. Domain and intent accuracy measure the accuracy of the domain and intent classification tasks, respectively. We use micro-averaging to calculate slot F1 to measure the performance of the NER task. Frame accuracy indicates the relative number of utterances for which the domain, intent, and all slots were correctly identified. In our case, we pay attention to individual metrics to understand which tasks are most affected by the synthesized data.

### 5.2 Language bootstrapping task

In this section we present the results of our experiment mimicking the task of bootstrapping a new language. See Section 3.3.1 for experimental setup. We compare three SeqGAN implementations each with a different reward policy. An open question when using generated data for NLU model training

is whether the improvements observed are due to the data enrichment gained by the new variations introduced in synthetic data or due to upsampling. To test this, we repeat each experiment on three subsamples of the generated synthetic data. First, we add a synthetic data set that is equal in size to the pseudo Golden data set. We call this *TopX* sampling. This enables us to explore the changes in performance obtained by adding a synthetic data set that reflects the original distribution produced by the trained GAN framework, but with limited effects of upsampling. In the second case, we produce a synthetic data set that is significantly larger than the pseudo Golden data set (9600 utterances per domain) and we take only the unique utterances within that set. We call this *Uniques* sampling. This enables us to explore the changes in performance obtained by adding a synthetic dataset that contains an exhaustive set of the different combinations of utterances that the GAN can create, given the input data. Finally, we add the full set of 9600 generated utterances per domain, that should be similar in distribution to the pseudo Golden set (*All* sampling). The results obtained for each SeqGAN implementation with TopX, Uniques, and All sampling strategies are summarized in Appendix B and Appendix C.

#### 5.2.1 Generator with token-level reward

When using the Generator with token-level reward, we observe that out of three synthetic data sampling strategies (i.e., TopX, Uniques, and All), the biggest overall gain is shown by the All sampling strategy with an overall intent accuracy improvement of 4%, and an overall frame accuracy improvement of 3% (Table 1). Domain accuracy and overall slot F1 do not show much change. In this setup, alarm and reminder domains' intent accuracies show increases of 12-14% with overall intent accuracy increasing by 4%. However, intent accuracy in the weather domain degrades by 4%. The generator with token-level reward outperforms the corresponding generator with sentence-level reward in domain accuracy, slot F1 and frame accuracy, but trails in intent metrics.

When the models are run using Golden's upsampled to the same counts, we observe that overall domain, intent, and frame accuracies are slightly better. However, individual intent accuracy of alarm and reminder domains do not perform as well as with synthetic data. No degradation of the weather domain is observed with upsampled data.

Table 1: Performance relative to the baseline for models with token-level and sentence-level reward using All sampling strategy. Baseline is a model trained on Goldens only.

Reward	Domain accuracy	Intent accuracy			Overall int. acc.	Slot F1			Overall slot F1	Frame accuracy
		alarm	reminder	weather		alarm	reminder	weather		
Token-level	0.29	11.84	14.62	-3.89	4.41	-1.04	0.8	1.25	0.39	3.26
Sent.-level	-0.26	14.45	17.4	0.27	7.81	-1.19	1.21	-2.77	-1.23	-4.43

Table 2: Performance relative to the baseline for models with Monte Carlo rollouts on token-level and sentence-level reward using Uniques sampling strategy. Baseline is a model trained on Goldens only.

Reward	Domain accuracy	Intent accuracy			Overall int. acc.	Slot F1			Overall slot F1	Frame accuracy
		alarm	reminder	weather		alarm	reminder	weather		
Token-level	0.41	14.83	14.43	-0.12	7.21	-0.43	2.82	0.25	0.67	4.3
Sent.-level	-1.37	9.07	9.19	0.23	4.71	-0.09	1.47	-0.52	-0.12	0.19

### 5.2.2 Generator with token-level Monte Carlo rollout

Using a Generator with token-level Monte Carlo rollout brings significant improvement, especially outperforming other models in the Uniques sampling strategy. Overall intent accuracy shows an improvement of 7%, and overall frame accuracy improves by 4% (Table 2). Alarm and reminder domains’ intent accuracy show improvements of 14-15%. The reminder domain’s slot F1 improves by 3%. This setup performs significantly better than the generator with sentence-level reward in all overall metrics and also in alarm and reminder domain intent accuracies.

When compared to the models run using Goldens upsampled to the same counts, we observe that the Generator with token-level Monte Carlo rollout policy performs better on domain accuracy, overall intent accuracy, and overall slot F1 while being slightly under on frame accuracy. Specifically, it performs much better on both the alarm and reminder domains which show approximately 50% the performance boost.

### 5.3 Handling low resource features

In this section we present the results of our experiment mimicking the task of handling low resource features. See Section 3.3.2 for the experimental setup. We conduct each experiment three times and present the mean results for Uniques sampling strategy in Appendix D.

#### 5.3.1 Generator with token-level reward and embeddings pre-trained on robust domains

Appendix D summarizes the results obtained using SeqGAN with a generator with token-level reward

and embeddings pre-trained on robust domains to synthesize data for the low-resource domain. We observe a 12% increase in intent accuracy in the alarm domain when compared to the baseline. We also observe a 13% increase in intent accuracy in the reminder domain when synthetic data is added. For the weather domain, we do not observe a significant change in intent accuracy. For alarm and reminder domains we see an increase in overall intent accuracy when synthetic data is added. For these same domains, overall frame accuracy shows small improvements while domain accuracy and overall slot F1 does not show any significant changes.

#### 5.3.2 Generator with token-level Monte Carlo rollout and embeddings pre-trained on robust domains

Appendix D shows the results obtained when using MC rollout policy in addition to using embeddings pre-trained on robust domains to synthesize data for the low-resource domain. For the reminder and alarm domains, we note that the performance boost in the IC task for the low-resource domain is larger by 2-3% than without the MC rollout. The weather domain shows a small but statistically significant improvement of 0.5% when compared to the baseline. These results suggest that the MC rollout policy provides additional guidance to the generator in all cases.

## 6 Synthetic data deep dive

### 6.1 Evaluating the quality of the data generated

To the best of our knowledge, there is no comprehensive metric that is commonly used for measuring the performance of a text generation model. To measure the performance of the SeqGAN, we

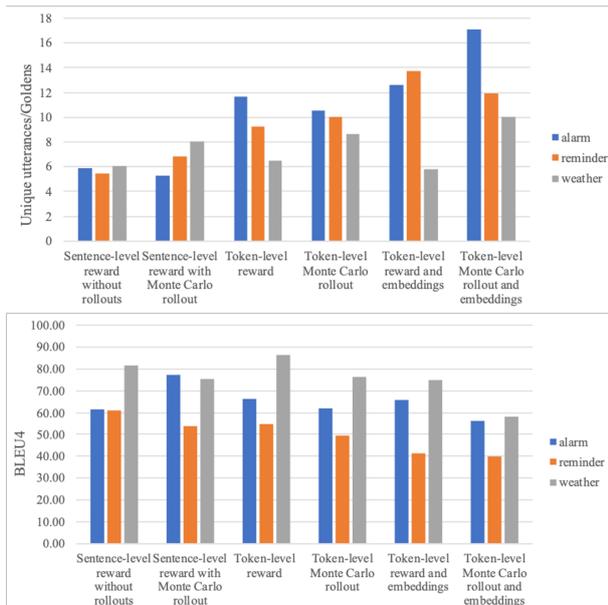


Figure 4: Data quality evaluation for synthetic data generated in different GAN models: number of unique utterances generated normalized by the number of Goldens used as an input for GAN model (left), and 4-gram BLEU score (right).

use the n-gram Bilingual Evaluation Understudy (BLEU) score (Papineni et al., 2002), calculated against a test set of Golden utterances. This metric measures the degree of similarity between the generated data and test set. Additionally, we calculate the diversity of generated data through the number of unique phrases we generated and number of unique words used from vocabulary. We also keep track on the mean utterance length to detect possible utterance collapse for token-level feedback, when the generator learns to generate shorter utterances. Detailed evaluations on bootstrapping experiments are provided in the Appendix E, and summarized in Figure 4. We see that compared to other domains, weather’s BLEU scores are higher, and the diversity of the generated data is lower, suggesting that the model potentially reproduces almost the same data as the input, and does not bring in much novelty for NLU model training.

## 6.2 Annotation review

Although, in general, we observed that the NLU model benefits from synthetic data, we have noted some degradation in NER models, especially in the weather domain. Deep diving into NER errors, we found the following major sources of errors: annotation errors and context-dependent annotations.

First, in seeding pseudo Goldens, we selected utterances based on their frequency, and then took all unique annotations to be a pseudo Goldens set. That process artificially increased weight of annotation errors in cases were for frequent phrases there were a few misannotated utterances. For example, for utterances “*remind me to ... tomorrow*” with and without label “*datetime*” frequently appeared in our goldens, and the NLU model fails to recognize “*tomorrow*” as an entity. In contrast, for the utterances “*remind me tomorrow to ...*” the model produces the correct “*datetime*” label. Additionally, small context-dependent words that have different annotation in the same domain, but appear to have a dominant annotation in pseudo Goldens and further in synthetic data, happen to be another cause of failures. One example of such a word is “*for*” in the weather domain, where in “*weather for London next week*” it has no label, while in “*forecast for next week please*” it is labeled as *datetime*.

## 7 Conclusions

In this paper, we evaluate the use of the SeqGAN model for synthetic annotated data generation to boost NLU model performance. We have shown that adding synthetic data to bolster our Goldens can significantly improve DNN model performance in intent classification and named entity recognition tasks. We propose a token-level reward with Monte Carlo search rollout to guide the generator model, that showed better performance when compared with a regular token-level reward implementation, sentence-level reward implementations both with and without Monte Carlo tree search, and with a pure upsampling strategy. We also show that using SeqGAN together with embeddings pre-trained on high-resource domains to generate synthetic data can significantly improve the performance of low-resource domains. Embeddings pre-trained on different tasks can carry over the information they have learned and that can be especially useful in low-resource model building scenarios.

## References

- Tong Che, Yanran Li, Ruixiang Zhang, R. Devon Hjelm, Wenjie Li, Yangqiu Song, and Yoshua Bengio. 2017. [Maximum-likelihood augmented discrete generative adversarial networks](#). *ArXiv*, abs/1702.07983.
- 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.
- William Fedus, Ian J. Goodfellow, and Andrew M. Dai. 2018. [Maskgan: Better text generation via filling in the ..](#) *ArXiv*, abs/1801.07736.
- Ge Gao, Eunsol Choi, Yejin Choi, and Luke Zettlemoyer. 2018. [Neural metaphor detection in context](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 607–613, Brussels, Belgium. Association for Computational Linguistics.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. [Generative adversarial nets](#). In *Advances in Neural Information Processing Systems*, pages 2672–2680. Curran Associates, Inc.
- Jiaxian Guo, Sidi Lu, Han Cai, Weinan Zhang, Yong Yu, and Jun Wang. 2018. [Long text generation via adversarial training with leaked information](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th Innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18)*, pages 5141–5148, New Orleans, Louisiana, USA. AAAI Press.
- Rahul Gupta. 2019. Data augmentation for low resource sentiment analysis using generative adversarial networks. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7380–7384. IEEE.
- Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, and Eric P. Xing. 2017. [Toward controlled generation of text](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1587–1596, International Convention Centre, Sydney, Australia. PMLR.
- Zhiting Hu, Zichao Yang, Tiancheng Zhao, Haoran Shi, Junxian He, Di Wang, Xuezhe Ma, Zhengzhong Liu, Xiaodan Liang, Lianhui Qin, Devendra Singh Chaplot, Bowen Tan, Xingjiang Yu, and Eric Xing. 2018. [Texar: A modularized, versatile, and extensible toolbox for text generation](#). In *Proceedings of Workshop for NLP Open Source Software (NLP-OSS)*, pages 13–22, Melbourne, Australia. Association for Computational Linguistics.
- Diederik P Kingma and Max Welling. 2013. [Auto-encoding variational bayes](#). *ArXiv*, abs/1312.6114.
- Varun Kumar, Hadrien Glaude, Cyprien de Lichy, and William Campbell. 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. Association for Computational Linguistics.
- Matt J. Kusner and José Miguel Hernández-Lobato. 2016. [Gans for sequences of discrete elements with the gumbel-softmax distribution](#). *ArXiv*, abs/1611.04051.
- Guillaume Lample, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2017. [Unsupervised machine translation using monolingual corpora only](#). *ArXiv*, abs/1711.00043.
- Jiwei Li, Will Monroe, Tianlin Shi, Sébastien Jean, Alan Ritter, and Dan Jurafsky. 2017. [Adversarial learning for neural dialogue generation](#). *ArXiv*, abs/1701.06547.
- Kevin Lin, Dianqi Li, Xiaodong He, Ming-Ting Sun, and Zhengyou Zhang. 2017. [Adversarial ranking for language generation](#). In *Advances in Neural Information Processing Systems 30*, pages 3155–3165. Curran Associates, Inc.
- Jie Lu, Vahid Behbood, Peng Hao, Hua Zuo, Shan Xue, and Guangquan Zhang. 2015. [Transfer learning using computational intelligence: A survey](#). *Knowledge-Based Systems*, 80:14–23.
- Nikolaos Malandrakis, Minmin Shen, Anuj Goyal, Shuyang Gao, Abhishek Sethi, and Angeliki Metallinou. 2019. [Controlled text generation for data augmentation in intelligent artificial agents](#). In *Proceedings of the 3rd Workshop on Neural Generation and Translation*, pages 90–98, Hong Kong. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [BLEU: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics, ACL ’02*, page 311–318, USA. Association for Computational Linguistics.
- Ashwin Ram, Rohit Prasad, Chandra Khatri, Anu Venkatesh, Raefer Gabriel, Qing Liu, Jeff Nunn, Behnam Hedayatnia, Ming Cheng, Ashish Nagar, Eric King, Kate Bland, Amanda Wartick, Yi Pan, Han Song, Sk Jayadevan, Gene Hwang, and Art Petigriue. 2018. [Conversational AI: The science behind the alexa prize](#). *ArXiv*, abs/1801.03604.
- Sebastian Schuster, Sonal Gupta, Rushin Shah, and Mike Lewis. 2019. [Cross-lingual transfer learning for multilingual task oriented dialog](#). 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 3795–3805, Minneapolis, Minnesota. Association for Computational Linguistics.

- Burr Settles. 2009. Active learning literature survey. Technical report, University of Wisconsin-Madison Department of Computer Sciences.
- Richard S Sutton, David A. McAllester, Satinder P. Singh, and Yishay Mansour. 2000. Policy gradient methods for reinforcement learning with function approximation. In *Advances in Neural Information Processing Systems 12*, pages 1057–1063. MIT Press.
- Ziang Xie, Sida I. Wang, Jiwei Li, Daniel Lévy, Aiming Nie, Dan Jurafsky, and Andrew Y. Ng. 2017. Data noising as smoothing in neural network language models. *ArXiv*, abs/1703.02573.
- Jingjing Xu, Xuancheng Ren, Junyang Lin, and Xu Sun. 2018. Diversity-promoting gan: A cross-entropy based generative adversarial network for diversified text generation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3940–3949, Brussels, Belgium. Association for Computational Linguistics.
- Lantao Yu, Weinan Zhang, Jun Wang, and Yong Yu. 2017. Seqgan: Sequence generative adversarial nets with policy gradient. In *Thirty-First AAAI Conference on Artificial Intelligence*, pages 2852–2858. AAAI Press.
- Yizhe Zhang, Zhe Gan, Kai Fan, Zhi Chen, Ricardo Henao, Dinghan Shen, and Lawrence Carin. 2017. Adversarial feature matching for text generation. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 4006–4015, International Convention Centre, Sydney, Australia. PMLR.