

An Adversarial Learning Framework For A Persona-Based Multi-Turn Dialogue Model

Oluwatobi Olabiyi

Capital One Conversation Research
Vienna VA
oluwatobi.olabiyi@capitalone.com

Anish Khazane

Capital One Conversation Research
San Francisco CA
anish.khazane@capitalone.com

Alan Salimov

Capital One Conversation Research
San Francisco CA
alan.salimov@capitalone.com

Erik T. Mueller

Capital One Conversation Research
Vienna VA
erik.mueller@capitalone.com

Abstract

In this paper, we extend the persona-based sequence-to-sequence (Seq2Seq) neural network conversation model to a multi-turn dialogue scenario by modifying the state-of-the-art hredGAN architecture to simultaneously capture utterance attributes such as speaker identity, dialogue topic, speaker sentiments and so on. The proposed system, phredGAN has a persona-based HRED generator (PHRED) and a conditional discriminator. We also explore two approaches to accomplish the conditional discriminator: (1) *phredGAN_a*, a system that passes the attribute representation as an additional input into a traditional adversarial discriminator, and (2) *phredGAN_d*, a dual discriminator system which in addition to the adversarial discriminator, collaboratively predicts the attribute(s) that generated the input utterance. To demonstrate the superior performance of phredGAN over the persona Seq2Seq model, we experiment with two conversational datasets, the Ubuntu Dialogue Corpus (UDC) and TV series transcripts from the Big Bang Theory and Friends. Performance comparison is made with respect to a variety of quantitative measures as well as crowd-sourced human evaluation. We also explore the trade-offs from using either variant of *phredGAN* on datasets with many but weak attribute modalities (such as with Big Bang Theory and Friends) and ones with few but strong attribute modalities (customer-agent interactions in Ubuntu dataset).

1 Introduction

Recent advances in machine learning especially with deep neural networks has lead to tremendous progress in natural language processing and dialogue modeling research (Sutskever et al., 2014; Vinyals and Le, 2015; Serban et al., 2016). Nevertheless, developing a good conversation model

capable of fluent interaction between a human and a machine is still in its infancy stage. Most existing work relies on limited dialogue history to produce response with the assumption that the model parameters will capture all the modalities within a dataset. However, this is not true as dialogue corpora tend to be strongly multi-modal and practical neural network models find it difficult to disambiguate characteristics such as speaker personality, location and sub-topic in the data.

Most work in this domain has primarily focused on optimizing dialogue consistency. For example, Serban et al. (2016, 2017b,a) and Xing et al. (2017) introduced a Hierarchical Recurrent Encoder-Decoder (HRED) network architecture that combines a series of recurrent neural networks to capture long-term context state within a dialogue. However, the HRED system suffers from lack of diversity and does not have any guarantee on the generator output since the output conditional probability is not calibrated. Olabiyi et al. (2018) tackles these problems by training a modified HRED generator alongside an adversarial discriminator in order to increase diversity and provide a strong and calibrated guarantee to the generator’s output. While the hredGAN system improves upon response quality, it does not capture speaker and other attributes modality within a dataset and fails to generate persona specific responses in datasets with multiple modalities.

On the other hand, there has been some recent work on introducing persona into dialogue models. For example, Li et al. (2016b) integrates attribute embeddings into a single turn (Seq2Seq) generative dialogue model. In this work, Li et al. consider persona models, one with Speaker-only representation and the other with Speaker and Addressee representations (Speaker-Addressee model), both of which capture certain

speaker identity and interactions. Nguyen et al. (2018) continue along the same line of thought by considering a Seq2Seq dialogue model with Responder-only representation. In both of these cases, the attribute representation is learned during the system training. Zhang et al. (2018) proposed a slightly different approach. Here, the attributes are a set of sentences describing the profile of the speaker. In this case, the attributes representation is not learned. The system however learns how to attend to different parts of the attributes during training. Still, the above persona-based models have limited dialogue history (single turn); suffer from exposure bias worsening the trade-off between personalization and conversation quality and cannot generate multiple responses given a dialogue context. This is evident in the relatively short and generic responses produced by these systems, even though they generally capture the persona of the speaker.

In order to overcome these limitations, we propose two variants of an adversarially trained persona conversational generative system, *phredGAN*, namely *phredGAN_a* and *phredGAN_d*. Both systems aim to maintain the response quality of *hredGAN* and still capture speaker and other attribute modalities within the conversation. In fact, both systems use the same generator architecture (PHRED generator), i.e., an *hredGAN* generator (Olabiya et al., 2018) with additional utterance attribute representation at its encoder and decoder inputs as depicted in Figure 1. Conditioning on external attributes can be seen as another input modality as is the utterance into the underlying system. The attribute representation is an embedding that is learned together with the rest of model parameters similar to Li et al. (2016b). Injecting attributes into a multi-turn dialogue system allows the model to generate responses conditioned on particular attribute(s) across conversation turns. Since the attributes are discrete, it also allows for exploring different what-if scenarios of model responses. The difference between the two systems is in the discriminator architecture based on how the attribute is treated.

We train and sample both variants of *phredGAN* similar to the procedure for *hredGAN* (Olabiya et al., 2018). To demonstrate model capability, we train on a customer service related data such as the Ubuntu Dialogue

Corpus (UDC) that is strongly bimodal between question poser and answerer, and transcripts from a multi-modal TV series *The Big Bang Theory* and *Friends* with quantitative and qualitative analysis. We examine the trade-offs between using either system in bi-modal or multi-modal datasets, and demonstrate system superiority over state-of-the-art persona conversational models in terms of human evaluation of dialogue response quality as well as automatic evaluations with perplexity, BLEU, ROUGE and distinct n-gram scores.

2 Model Architecture

In this section, we briefly introduce the state-of-the-art *hredGAN* model and subsequently show how we derive the two persona versions by combining it with the distributed representation of the dialogue speaker and utterance attributes, or with an attribute discrimination layer at the end of the model pipeline.

2.1 *hredGAN*: Adversarial Learning Framework

Problem Formulation: The *hredGAN* (Olabiya et al., 2018) formulates multi-turn dialogue response generation as: given a dialogue history of sequence of utterances, $\mathbf{x}_i = (x_1, x_2, \dots, x_i)$, where each utterance $x_i = (x_i^1, x_i^2, \dots, x_i^{M_i})$ contains a variable-length sequence of M_i word tokens such that $x_i^j \in V$ for vocabulary V , the dialogue model produces an output $y_i = (y_i^1, y_i^2, \dots, y_i^{T_i})$, where T_i is the number of generated tokens. The framework uses conditional GAN structure to learn a mapping from an observed dialogue history to a sequence of output tokens. The generator, G , is trained to produce sequences that cannot be distinguished from the ground truth by an adversarially trained discriminator, D akin to a two-player min-max optimization problem. The generator is also trained to minimize the cross-entropy loss $\mathcal{L}_{MLE}(G)$ between the ground truth x_{i+1} , and the generator output y_i . The following objective summarizes both goals:

$$G^*, D^* = \arg \min_G \max_D (\lambda_G \mathcal{L}_{cGAN}(G, D) + \lambda_M \mathcal{L}_{MLE}(G)). \quad (1)$$

where λ_G and λ_M are training hyperparameters and $\mathcal{L}_{cGAN}(G, D)$ and $\mathcal{L}_{MLE}(G)$ are defined in Eqs.

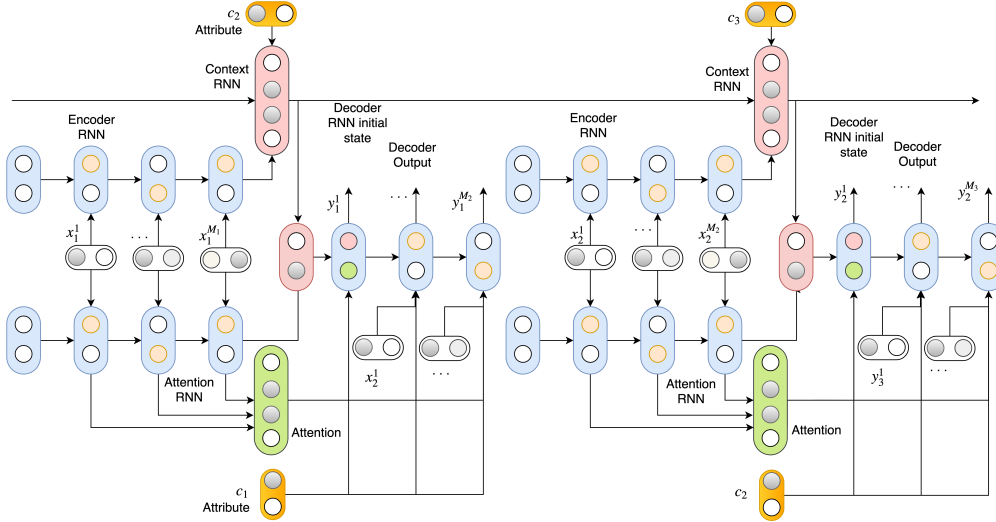


Figure 1: **The PHRED generator with local attention** - The attributes c_i , allows the generator to condition its response on the utterance attributes such as speaker identity, subtopics and so on.

(5) and (7) of Olabiyi et al. (2018) respectively. Please note that the generator G and discriminator D share the same encoder and embedding representation of the word tokens.

2.2 phredGAN: Persona Adversarial Learning Framework

The proposed architecture of *phredGAN* is very similar to that of *hredGAN* (Olabiyi et al., 2018). The only difference is that the dialogue history is now $\mathbf{x}_i = ((x_1, c_1), (x_2, c_2), \dots, (x_i, c_i))$ where c_i is additional input that represents the speaker and/or utterance attributes. Please note that c_i can either be a sequence of tokens or single token such that $c_i^j \in Vc$ for vocabulary Vc . Also, at the i th turn, c_i and c_{i+1} are the source/input attribute and target/output attribute to the generator respectively. The embedding for attribute tokens is also learned similar to that of word tokens.

Both versions of *phredGAN* shares the same generator architecture (PHRED) but different discriminators. Below is the highlight of how they are derived from the *hredGAN* architecture.

Encoder: The context RNN, $cRNN$ takes the source attribute c_i as an additional input by concatenating its representation with the output of $eRNN$ as in Figure 1. If the attribute c_i is a sequence of tokens, then an attention (using the output of $eRNN$) over the source attribute representations is concatenated with the output of $eRNN$. This output is used by the generator to create a context state for a turn i .

Generator: The generator decoder RNN,

$dRNN$ takes the target attribute c_{i+1} as an additional input as in Fig. 1. If the attribute c_{i+1} is a sequence of tokens, then an attention (using the output of $dRNN$) over the attribute representations is concatenated with the rest of the decoder inputs. This forces the generator to draw a connection between the generated responses and the utterance attributes such as speaker identity.

Noise Injection: As in Olabiyi et al. (2018), we also explore different noise injection methods.

Objective: For *phredGAN*, the optimization objective in eq. (1) can be updated as:

$$G^*, D_{adv}^*, D_{att}^* = \arg \min_G \left(\begin{aligned} & \max_{D_{adv}} \lambda_{G_{adv}} \mathcal{L}_{cGAN}^{adv}(G, D_{adv}) \\ & + \min_{D_{att}} \lambda_{G_{att}} \mathcal{L}_c^{att}(G, D_{att}) \\ & + \lambda_M \mathcal{L}_{MLE}(G) \end{aligned} \right). \quad (2)$$

where $\mathcal{L}_{cGAN}^{adv}(G, D_{adv})$ and $\mathcal{L}_c^{att}(G, D_{att})$ are the traditional adversarial and attribute prediction loss respectively and dependent on the architectural variation. It is worth to point out that while the former is adversarial, the later is collaborative in nature. The MLE loss is common and can be expressed as:

$$\mathcal{L}_{MLE}(G) = \mathbb{E}_{x_{i+1}} [-\log P_G(x_{i+1} | \mathbf{x}_i, c_{i+1}, z_i)]. \quad (3)$$

where z_i the noise sample and depends on the choice of either utterance-level or word-level noise input into the generator (Olabiyi et al., 2018).

2.3 $phredGAN_a$: Attributes as a Discriminator Input

$phredGAN_a$ shares the same discriminator architecture as the $hredGAN$ but with additional input, c_{i+1} . Since it does not use attribute prediction, $\lambda_{G_{att}} = 0$.

The adversarial loss, $\mathcal{L}_{cGAN}^{adv}(G, D)$ can then be expressed as:

$$\begin{aligned} \mathcal{L}_{cGAN}^{adv}(G, D_{adv}) = & \mathbb{E}_{\mathbf{x}_i, c_{i+1}, x_{i+1}} [\log D_{adv}(\mathbf{x}_i, c_{i+1}, x_{i+1})] + \\ & \mathbb{E}_{\mathbf{x}_i, c_{i+1}, z_i} [1 - \log D_{adv}(\mathbf{x}_i, c_{i+1}, G(\mathbf{x}_i, c_{i+1}, z_i))] \end{aligned} \quad (4)$$

The addition of speaker or utterance attributes allows the dialogue model to exhibit personality traits given consistent responses across style, gender, location, and so on.

2.4 $phredGAN_d$: Attributes as a Discriminator Target

$phredGAN_d$ does not take the attribute representation at its input but rather uses the attributes as the target of an additional discriminator D_{att} . The adversarial and the attribute prediction losses can be respectively expressed as:

$$\begin{aligned} \mathcal{L}_{cGAN}^{adv}(G, D_{adv}) = & \mathbb{E}_{\mathbf{x}_i, x_{i+1}} [\log D_{adv}(\mathbf{x}_i, x_{i+1})] \\ & + \mathbb{E}_{\mathbf{x}_i, z_i} [1 - \log D_{adv}(\mathbf{x}_i, G(\mathbf{x}_i, c_{i+1}, z_i))] \end{aligned} \quad (5)$$

$$\begin{aligned} \mathcal{L}_c^{att}(G, D_{att}) = & \mathbb{E}_{c_{i+1}} [-\log D_{att}(c_{i+1} | \mathbf{x}_i, x_{i+1})] \\ & + \mathbb{E}_{c_{i+1}} [-\log D_{att}(c_{i+1} | \mathbf{x}_i, G(\mathbf{x}_i, c_{i+1}, z_i))] \end{aligned} \quad (6)$$

Attribute Discriminator: In addition to the existing word-level adversarial discriminator D_{adv} from $hredGAN$, we add an attribute discriminator, D_{att} , that discriminates on an utterance level to capture attribute modalities since attributes are assigned at utterance level. The discriminator uses a unidirectional RNN (D_{attRNN}) that maps the input utterance to the particular attribute(s) that generated it. The attributes can be seen as hidden states that inform or shape the generator outputs. The attribute discriminator can be expressed as:

$$D_{att}(c_{i+1} | \mathbf{x}_i, \chi) = D_{attRNN}(\mathbf{h}_i, E(\chi)) \quad (7)$$

where $E(\cdot)$ is the word embedding lookup (Olabi et al., 2018), $\chi = x_{i+1}$ for groundtruth and $\chi = y_i$ for the generator output.




Figure 2: **The $phredGAN_d$ dual discriminator** - **Left:** D_{adv} is a word-level discriminator used by both $phredGAN_a$ and $phredGAN_d$ to judge normal dialogue coherency as in $hredGAN$. **Right:** D_{att} , an utterance-level attribute discriminator is used only in $phredGAN_d$ to predict the likelihood a given utterance was generated from a particular attribute.

3 Model Training and Inference

3.1 Model Training

We train both the generator and the discriminator (with shared encoder) of both variants of $phredGAN$ using the training procedure in Algorithm 1 (Olabiyi et al., 2018). For both variants, $\lambda_{G_{adv}} = \lambda_M = 1$, and for $phredGAN_a$ and $phredGAN_d$, $\lambda_{G_{att}} = 0$ and $\lambda_{G_{att}} = 1$ respectively. Since the encoder, word embedding and attribute embedding are shared, we are able to train the system end-to-end with back-propagation.

Encoder: The encoder RNN, $eRNN$, is bidirectional while $cRNN$ is unidirectional. All RNN units are 3-layer GRU cell with hidden state size of 512. We use word vocabulary size, $V = 50,000$ with word embedding size of 512. The number of attributes, V_c is dataset dependent but we use an attribute embedding size of 512. In this study, we only use one attribute per utterance so there is no need to use an attention mechanism to combine the attribute embeddings.

Generator: The generator decoder RNN, $dRNN$ is also a 3-layer GRU cell with hidden state size of 512. The $aRNN$ outputs are connected to the $dRNN$ input using an additive attention mechanism (Bahdanau et al., 2015).

Adversarial Discriminator: The word-level discriminator RNN, D_{RNN} is a bidirectional RNN, each 3-layer GRU cell with hidden state size of 512. The output of both the forward and the backward cells for each word are concatenated and passed to a fully-connected layer with binary output. The output is the probability that the word is from the ground truth given the past and future words of the sequence, and in the case of $phredGAN_a$, the responding speaker’s embed-

Algorithm 1 Adversarial Learning of phredGAN

Require: A generator G with parameters θ_G .
Require: An adversarial discriminator D_{adv} with parameters $\theta_{D_{adv}}$.
Require: An attribute discriminator D_{att} with parameters $\theta_{D_{att}}$.
Require: Training hyperparameters, $isTarget$, $\lambda_{G_{att}}$, $\lambda_{G_{adv}}$, and λ_M .

for number of training iterations **do**
 Initialize $cRNN$ to zero state, \mathbf{h}_0 .
 Sample a mini-batch of conversations, $\mathbf{x} = \{x_i, c_i\}_{i=1}^N$, $\mathbf{x}_i = ((x_1, c_1), (x_2, c_2), \dots, (x_i, c_i))$ with N utterances. Each utterance mini batch i contains M_i word tokens.
 for $i = 1$ **to** $N - 1$ **do**
 Update the context state.
 $\mathbf{h}_i = cRNN(eRNN(E(x_i)), \mathbf{h}_{i-1}, c_i)$
 Compute the generator output similar to Eq. (11) in (Olabiya et al., 2018).
 $P_{\theta_G}(y_i | z_i, \mathbf{x}_i, c_{i+1}) = \{P_{\theta_G}(y_i^j | x_{i+1}^{1:j-1}, z_i^j, \mathbf{x}_i, c_{i+1})\}_{j=1}^{M_{i+1}}$
 Sample a corresponding mini batch of utterance y_i .
 $y_i \sim P_{\theta_G}(y_i | z_i, \mathbf{x}_i, c_{i+1})$
 end for
 Compute the adversarial discriminator accuracy D_{adv}^{acc} over $N - 1$ utterances $\{y_i\}_{i=1}^{N-1}$ and $\{x_{i+1}\}_{i=1}^{N-1}$
 if $D_{adv}^{acc} < acc_{D_{adv}^{th}}$ **then**
 if $isTarget$ **then**
 Update $phredGAN_d$'s $\theta_{D_{adv}}$ and $\theta_{D_{att}}$.
 $\sum_i [\nabla_{\theta_{D_{adv}}} \log D_{adv}(\mathbf{h}_i, x_{i+1}) + \nabla_{\theta_{D_{adv}}} \log(1 - D_{adv}(\mathbf{h}_i, y_i)) + \nabla_{\theta_{D_{att}}} -\log D_{att}(c_{i+1} | \mathbf{h}_i, x_{i+1})]$
 else
 Update $phredGAN_a$'s $\theta_{D_{adv}}$ with gradient of the discriminator loss.
 $\sum_i [\nabla_{\theta_{D_{adv}}} \log D_{adv}(\mathbf{h}_i, c_{i+1}, x_{i+1}) + \nabla_{\theta_{D_{adv}}} \log(1 - D_{adv}(\mathbf{h}_i, c_{i+1}, y_i))]$
 end if
 end if
 if $D_{adv}^{acc} < acc_{G_{th}}$ **then**
 Update θ_G with the generator's MLE loss only.
 $\sum_i [\nabla_{\theta_G} -\log P_{\theta_G}(y_i | z_i, \mathbf{x}_i, c_{i+1})]$
 else
 Update θ_G with attribute, adversarial and MLE losses.
 $\sum_i [\lambda_{G_{att}} \nabla_{\theta_G} -\log D_{att}(c_{i+1} | \mathbf{h}_i, y_i) + \lambda_{G_{adv}} \nabla_{\theta_G} \log D_{adv}(\mathbf{h}_i, c_{i+1}, y_i) + \lambda_M \nabla_{\theta_G} -\log P_{\theta_G}(y_i | z_i, \mathbf{x}_i, c_{i+1})]$
 end if
end for

ding.

Attribute Discriminator: The attribute discriminator RNN, D_{attRNN} is a unidirectional RNN with a 3-layer GRU cell, each of hidden state size 512. A softmax layer is then applied to project the final hidden state to a prespecified number of attributes, V_c . The output is the probability distribution over the attributes.

Others: All parameters are initialized with Xavier uniform random initialization (Glorot and Bengio, 2010). Due to the large word vocabulary size, we use sampled softmax loss (Jean et al., 2015) for MLE loss to expedite the training process. However, we use full softmax for model evaluation. For both systems, parameters updates are conditioned on the word-level discriminator accuracy performance as in Olabiya et al. (2018) with $acc_{D_{adv}^{th}} = 0.99$ and $acc_{G_{th}} = 0.75$. The model is trained end-to-end using the stochastic

gradient descent algorithm. Finally, the model is implemented, trained, and evaluated using the TensorFlow deep learning framework.

3.2 Model Inference

We use an inference strategy similar to the approach in Olabiya et al. (2018).

For the modified noise sample, we perform a linear search for α with sample size $L = 1$ based on the average word-level discriminator loss, $-\log D_{adv}(G(\cdot))$ (Olabiya et al., 2018) using trained models run in autoregressive mode to reflect performance in actual deployment. The optimum α value is then used for all inferences and evaluations. During inference, we condition the dialogue response generation on the encoder outputs, noise samples, word embedding and the attribute embedding of the intended responder. With multiple noise samples, $L = 64$, we rank the generator outputs by the discriminator which is also conditioned on encoder outputs, and the intended responder's attribute embedding. The final response is the response ranked highest by the discriminator. For $phredGAN_d$, we average the confidences produced by D_{adv} and D_{att} .

4 Experiments and Results

In this section, we explore the performance of PHRED, $phredGAN_a$ and $phredGAN_d$ on two conversational datasets and compare their performances to non-adversarial persona Seq2seq models (Li et al., 2016b) as well as to the adversarial $hredGAN$ (Olabiya et al., 2018) with no explicit persona.

4.1 Datasets

TV Series Transcripts dataset (Serban et al., 2016). We train all models on transcripts from two popular TV drama series, Big Bang Theory and Friends. Following a similar preprocessing setup in Li et al. (2016b), we collect utterances from the top 12 speakers from both series to construct a corpus of 5,008 lines of multi-turn dialogue. We split the corpus into training, development, and test set with a 94%, 3%, and 3% proportions, respectively, and pair each set with a corresponding attribute file that maps speaker IDs to utterances in the combined dataset.

Due to the small size of the combined transcripts dataset, we first train the models on the larger Movie Triplets Corpus (MTC) by Banchs

(2012) which consists of 240,000 dialogue triples. We pre-train the models on this dataset to initialize the model parameters to avoid overfitting on a relatively small persona TV series dataset. After pre-training on MTC, we reinitialize the attribute embeddings in the generator from a uniform distribution following a Xavier initialization (Glorot and Bengio, 2010) for training on the combined person TV series dataset.

Ubuntu Dialogue Corpus (UDC) dataset (Serban et al., 2017b). We train the models on 1.85 million conversations of multi-turn dialogue from the Ubuntu community hub, with an average of 5 utterances per conversation. We assign two types of speaker IDs to utterances in this dataset: questioner and helper. We follow a similar training, development, and test split as the UDC dataset in Olabiyi et al. (2018), with 90%, 5%, and 5% proportions, respectively, and pair each set with a corresponding attribute file that maps speaker IDs to utterances in the combined dataset

While the overwhelming majority of utterances in UDC follow two speaker types, the dataset does include utterances that do not classify under either a questioner or helper speaker type. In order to remain consistent, we assume that there are only two speaker types within this dataset and that the first utterance of every dialogue is from a questioner. This simplifying assumption does introduce a degree of noise into each persona model’s ability to construct attribute embeddings. However, our experiment results demonstrate that both $phredGAN_a$ and $phredGAN_d$ are still able to differentiate between the larger two speaker types in the dataset.

4.2 Evaluation Metrics

We use similar evaluation metrics as in Olabiyi et al. (2018) including perplexity, BLEU (Papineni et al., 2002), ROUGE (Lin, 2014), distinct n-gram (Li et al., 2016a) and normalized average sequence length (NASL) scores. For human evaluation, we follow a similar setup as Li et al. (2016a), employing crowd-sourced judges to evaluate a random selection of 200 samples. We present both the multi-turn context and the generated responses from the models to 3 judges and asked them to rank the general response quality in terms of relevance, informativeness, and persona. For N models, the model with the lowest quality is assigned a score 0 and the highest is assigned a score N .

1. Ties are not allowed. The scores are normalized between 0 and 1 and averaged over the total number of samples and judges. For each model, we also estimate the per sample score variance between judges and then average over the number of samples, i.e., sum of variances divided by the square of number of samples (assuming sample independence). The square root of result is reported as the standard error of the human judgement for the model.

4.3 Baseline

We compare the non-adversarial persona HRED model, PHRED with the adversarially trained ones, i.e. $hredGAN$, $phredGAN_a$ and $phredGAN_d$, to demonstrate the impact of adversarial training. Please note that no noise was added to the PHRED model.

We also compare the persona models to Li et al.’s work (Li et al., 2016b) which uses a Seq2Seq framework in conjunction with learnable persona embeddings. Their work explores two persona models in order to incorporate vector representations of speaker interaction and speaker attributes into the decoder of their Seq2Seq models i.e., Speaker model (SM) and Speaker-Addressee model (SAM). All reported results are based on our implementation of their models in Li et al. (2016b).

4.4 Hyperparameter Search

For both $phredGAN_a$ and $phredGAN_d$, we determine the noise injection method and the optimum noise variance α that allows for the best performance on both datasets. We find that $phredGAN_d$ performs optimally with word-level noise injection on both Ubuntu and TV transcripts, while $phredGAN_a$ performs the best with utterance-level noise injection on TV transcripts and word-level injection on UDC. For all $phredGAN$ models, we perform a linear search for optimal noise variance values between 1 and 30 at an increment of 1, with a sample size of $L = 1$. For $phredGAN_d$, we obtain an optimal α of 4 and 6 for the UDC and TV Transcripts respectively. For $phredGAN_a$, we obtain an optimal value of 2 and 5 for the combined TV series dataset and the much larger UDC respectively.

4.5 Results

We will now present our assessment of performance comparisons of $phredGAN$ against the

Table 1: *phredGAN* vs. Li et al. (2016b) on BBT Friends TV Transcripts.

Model	Teacher Forcing Perplexity	Autoregression				Human Evaluation
		BLEU	ROUGE-2	DISTINCT-1/2	NASL	
TV Series						
SM	22.13	1.76 %	22.4 %	2.50%/18.95%	0.786	0.5566 ± 0.0328
SAM	23.06	1.86 %	20.52 %	2.56%/18.91%	0.689	0.5375 ± 0.0464
<i>hredGAN</i>	28.15	2.14 %	6.81 %	1.85 %/6.93 %	1.135	0.5078 ± 0.0382
<i>phred</i>	30.94	2.41 %	14.03 %	0.66 %/2.54 %	1.216	0.3663 ± 0.0883
<i>phredGAN_a</i>	25.10	3.07 %	30.47 %	2.19 %/19.02 %	1.218	0.6127 ± 0.0498
<i>phredGAN_d</i>	28.19	2.76 %	14.68 %	0.70 %/4.76 %	1.163	0.4284 ± 0.0337

Table 2: *phredGAN* vs. Li et al. (2016b) on UDC.

Model	Teacher Forcing Perplexity	Autoregression				Human Evaluation
		BLEU-2/4	ROUGE-2	DISTINCT-1/2	NASL	
UDC						
SM	28.32	0.437%/~ 0%	9.19 %	1.61%/5.79%	0.506	0.4170 ± 0.0396
SAM	26.12	0.490%/~ 0%	10.23 %	1.85%/6.85%	0.512	0.4629 ± 0.0171
<i>hredGAN</i>	48.18	2.16%/~ 0%	11.68 %	5.16%/18.21 %	1.098	0.5876 ± 0.0532
<i>phred</i>	34.67	0.16%/~ 0%	7.41 %	0.56%/1.44%	0.397	0.4399 ± 0.0445
<i>phredGAN_a</i>	31.25	1.94%/~ 0%	19.15%	1.05%/5.28%	1.520	0.4920 ± 0.0167
<i>phredGAN_d</i>	28.74	2.02%/ 0.10%	16.82%	1.38%/5.77%	1.387	0.5817 ± 0.0615

baselines, PHRED, *hredGAN* and Li et al.’s persona Seq2Seq models.

4.6 Quantitative Analysis

We first report the performance on TV series transcripts in table 1. The performance of both SM and SAM models in Li et al. (2016b) compared to the *hredGAN* shows a strong baseline and indicates that the effect of persona is more important than that of multi-turn and adversarial training for datasets with weak multiple persona. However, once the persona information is added to the *hredGAN*, the resulting *phredGAN* shows a significant improvement over the SM and SAM baselines with *phredGAN_a* performing best. We also observe that PHRED performs worse than the baseline S(A)M models on a number of metrics but we attribute this to the effect of persona on a limited dataset that results into less informative responses. This behavior was also reported in Li et al. (2016b) where the persona models produce less informative responses than the non-personal Seq2seq models but it seems to be even worse in multi-turn context. However, unlike the Speaker-Addressee and PHRED models that suffer from lower response quality due to persona conditioning, we note that conditioning the generator and discriminator of *phredGAN* on speaker embeddings does not compromise the systems ability to

produce diverse responses. This problem might have been alleviated by the adversarial training that encourages the generator model to produce longer, more informative, and diverse responses that have high persona relevance even with a limited dataset.

We also compare the models performances on the UDC. The evaluation result is summarized in table 2. While the deleterious effect of persona conditioning on response diversity is still worse with PHRED than with S(A)M models, we note that *hredGAN* performs much better than the S(A)M models. This is because, the external persona only provides just a little more information than is already available from the UDC utterances. Therefore, performance on UDC is mostly driven by longer dialogue context and adversarial training. We also note an improvement of *phredGAN* variants over the *hredGAN* in a variety of evaluation metrics including perplexity, ROUGE with the exception of distinct n-grams. This is expected as *phredGAN* should be generally less diverse than *hredGAN* since each persona attribute of *phredGAN* covers only a limited region of the data distribution. This, however, leads to better response quality with persona, something not achievable with *hredGAN*. Also, the much better ROUGE(F1) score indicates that *phredGAN* is able to strike a better balance be-

tween diversity and precision while still capturing the characteristics of the speaker attribute modality in the UDC dataset. Within the *phredGAN* variants, *phredGAN_d* seems to perform better. This is not surprising as speaker classification is much easier on UDC than on TV series. The attribute discriminator, D_{att} is able to provide more informative feedback on UDC than on TV series where it is more difficult to accurately predict the speaker. Therefore, we recommend *phredGAN_a* for datasets with weak attribute distinction and *phredGAN_d* for strong attribute distinction.

4.7 Qualitative Analysis¹

In addition to the quantitative analysis above, we report the results of the human evaluation in the last column of Tables 1 and 2 for the TV Series and UDC datasets respectively. The human evaluation scores largely agrees with the automatic evaluations on the TV Series with *phredGAN_a* clearly giving the best performance. However, on the UDC, both *hredGAN* and *phredGAN_d* performs similarly which indicates that there is a trade off between diversity and persona by each model. We believe this is due to the strong persona information that already exists in the UDC utterances.

An additional qualitative assessment of these results are in Table 3 with responses from several characters in the TV series dataset and the two characters in UDC.

We see that for TV drama series, *phredGAN* responses are comparatively more informative than that of the Speaker-Addressee model of Li et al. (2016b). For example, all the characters in the TV series respond the same to the dialogue context. Similar behavior is reported in Li et al. (2016b) where for the Speaker-Addressee model, nearly all the characters in the TV series respond with “Of course I love you.” to the dialogue context, “Do you love me?” despite the fact that some of the responders sometimes have unfriendly relationship with the addressee. Many of the novel situations explored by *phredGAN* are unachievable with the Speaker-Addressee model due to lack of informative responses. For example, by conditioning as Sheldon from The Big Bang Theory and asking “Do you like me?”, our model responds with annoyance if conditioned as Penny (“No, you don’t understand. You’re an idiot”), brevity with

Leonard (“Yes?”) and sarcasm with Raj (“Well , you know , we could be a little more than my friend’s friends.”) The wide range of responses indicate our model’s ability to construct distinct attribute embeddings for each character even from a limited dataset. The other interesting responses in Table 3 indicate *phredGAN*’s ability to infer not only the context of the conversation but important character information about the addressee.

We also see similar results with our model’s output on UDC in Table 4. We demonstrate that by conditioning as either a helper or questioner from the UDC dataset, *phredGAN* models are able to respond differently to input utterances as well as stay close to the context of the conversation. For the purpose of completeness, we also show some samples from PHRED generator on both UDC and TV series dataset in Table 5.

5 Conclusion and Future Work

In this paper, we improve upon state-of-the-art persona-based response generation models by exploring two persona conversational models: *phredGAN_a* which passes the attribute representation as an additional input into a traditional adversarial discriminator, and *phredGAN_d* a dual discriminator system which in addition to the adversarial discriminator from *hredGAN*, collaboratively predicts the attribute(s) that are intrinsic to the input utterance. Both systems demonstrate quantitative improvements upon state-of-the-art persona conversational systems such as the work from Li et al. (2016b) with respect to both quantitative automatic and qualitative human measures.

Our analysis also demonstrates how both variants of *phredGAN* perform differently on datasets with weak and strong modality. One of our future direction is to take advantage of *phredGAN_d*’s ability to predict utterance attribute such as speaker identity from just the utterance. We believe its performance can be improved even with weak modality by further conditioning adversarial updates on both the attribute and adversarial discriminator accuracies. Overall, this paper demonstrates clear benefits from adversarial training of persona generative dialogue system and leaves the door open for more interesting work in this domain.

¹Tables 3, 4 and 5 referenced in this section are in the appendix.

References

- D. Bahdanau, K. Cho, and Y. Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proceedings of International Conference of Learning Representation (ICLR 2015)*.
- R. E. Banchs. 2012. Movie-dic: A movie dialogue corpus for research and development. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics*, pages 203–207.
- X. Glorot and Y. Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *International conference on artificial intelligence and statistics*.
- S. Jean, K. Cho, R. Memisevic, and Y. Bengio. 2015. On using very large target vocabulary for neural machine translation. In *arXiv preprint arXiv:1412.2007*.
- J. Li, M. Galley, C. Brockett, J. Gao, and B. Dolan. 2016a. A diversity-promoting objective function for neural conversation models. In *Proceedings of NAACL-HLT*.
- J. Li, M. Galley, C. Brockett, G. Spithourakis, J. Gao, and B. Dolan. 2016b. A persona-based neural conversation model. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 994–1003.
- C. Y. Lin. 2014. Rouge: a package for automatic evaluation of summaries. In *Proceedings of the Workshop on Text Summarization Branches Out*.
- H. Nguyen, D. Morales, and T. Chin. 2018. A neural chatbot with personality. In *Stanford NLP Course website: <https://web.stanford.edu/class/cs224n/reports/2761115.pdf>*.
- O. Olabiyi, A. Salimov, A. Khazane, and E. Mueller. 2018. Multi-turn dialogue response generation in an adversarial learning framework. In *arXiv preprint arXiv:1805.11752*.
- K. Papineni, S. Roukos, T. Ward, and W. Zhu. 2002. Bleu: A method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318.
- I. Serban, A. Sordoni, Y. Bengio, A. Courville, and J. Pineau. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *Proceedings of The Thirtieth AAAI Conference on Artificial Intelligence (AAAI 2016)*, pages 3776–3784.
- I. V. Serban, T. Klinger, G. Tesauro, K. Talamadupula, B. Zhou, Y. Bengio, and A. Courville. 2017a. Multiresolution recurrent neural networks: An application to dialogue response generation. In *Proceedings of The Thirty-first AAAI Conference on Artificial Intelligence (AAAI 2017)*.
- I. V. Serban, A. Sordoni, R. Lowe, L. Charlin, J. Pineau, A. Courville, and Y. Bengio. 2017b. A hierarchical latent variable encoder-decoder model for generating dialogue. In *Proceedings of The Thirty-first AAAI Conference on Artificial Intelligence (AAAI 2017)*.
- I. Sutskever, O. Vinyals, and Q. Le. 2014. Sequence to sequence learning with neural networks. In *Proceedings of Advances in Neural Information Processing Systems (NIPS)*, pages 3104–3112.
- O. Vinyals and Q. Le. 2015. A neural conversational model. In *Proceedings of ICML Deep Learning Workshop*.
- C. Xing, W. Wu, Y. Wu, M. Zhou, Y. Huang, and W. Ma. 2017. Hierarchical recurrent attention network for response generation. In *arXiv preprint arXiv:1701.07149*.
- S. Zhang, E. Dinan, J. Urbanek, A. Szlam, D. Kiela, and J. Weston. 2018. Personalizing dialogue agents: I have a dog, do you have pets too? In *arXiv preprint arXiv:1801.07243v3*.

A Appendix

Table 3: *hredGAN* vs *phredGAN* vs Li et al. (2016b) on TV Transcripts

Model	Speaker	Response
Big Bang Theory TV Series		
Context.0	Sheldon	Do you like me?
<i>phredGAN_a</i>	Leonard	Yes ?
<i>phredGAN_d</i>	Leonard	You can ' t do that ? !
<i>phredGAN_a</i>	Penny	No , you don ' t understand . You ' re an idiot .
<i>phredGAN_d</i>	Penny	You ' re right . . .
<i>phredGAN_a</i>	Howard	Well , I don ' t know what to say . This is the part I ' ve ever seen .
<i>phredGAN_d</i>	Howard	You ' re right , you ' re just gonna go with this guy .
<i>phredGAN_a</i>	Raj	Well , you know , we could be a little more than my friends ' s friends .
<i>phredGAN_d</i>	Raj	You know , I can ' t can ' t do I can ' go , I can ' t do that . . .
SM	L/P/H/R ²	Yeah , I ' m sure . Oh , but I can ' t believe it . I ' m not gonna be in here
SAM	L/P/H/R	I ' m not going to be a good time .
<i>hredGAN</i>	NA	Okay , I ' m gonna go see my friend .
Context.0	Howard	We can ' t go home.
Context.1	Leonard	Alright, I'll be right back.
<i>phredGAN_a</i>	Howard	Okay , well , before you and Penny get out and get married with Penny just get married.
<i>phredGAN_d</i>	Howard	You ' ll see , I ' ll just go .
SM	Howard	Okay .
SAM	Howard	Okay , well , you ' re really good at this .
<i>hredGAN</i>	NA	Okay , I ' m gonna go see my old .
Context.0	Amy	Spent a lot of my childhood throwing coins into wishing wells hoping for friends.
Context.1	Penny	Whoa, whoa, give me a second to catch up here. Hold on.
Context.2	Amy	That'd be my boyfriend. Happier playing his dopey Star Trek game with his friends than hanging out with me.
<i>phredGAN_a</i>	Penny	You know , when you ' re being so nice about Sheldon and I want to take him to like his own girlfriend
<i>phredGAN_d</i>	Penny	You , I ' m sorry , I didn ' t have it . d be back in the bathroom . .
SM	Penny	I ' m sorry , I ' m sorry . I ' m just saying that you ' re going to have to answer to
SAM	Penny	It ' s not gonna happen .
<i>hredGAN</i>	NA	Okay , I ' m gonna go see my new .
Context.0	Sheldon	I recently read that during World War Two, Joseph Stalin had a research program to create supersoldiers by having women impregnated by gorillas.
<i>phredGAN_a</i>	Howard	You ' re right . I don ' t know what you ' re talking about , Sheldon .
<i>phredGAN_d</i>	Howard	You ' re kidding . .
SM	Howard	Well , my last was a little . I was just trying to be a good time .
SAM	Howard	That would be great .
<i>hredGAN</i>	NA	Okay , I ' m gonna go see my new friend , and then we ' ll go .
Friends TV Series		
Context.0	Rachel	Do you like me?
<i>phredGAN_a</i>	Joey	Yeah .
<i>phredGAN_d</i>	Joey	You ' re right
<i>phredGAN_a</i>	Ross	Well , no , no , no . I mean , I just , I don ' t know
<i>phredGAN_d</i>	Ross	You ' want to know sex . .
<i>phredGAN_a</i>	Monica	Well , I do . And I ' m not gonna tell you another guy .
<i>phredGAN_d</i>	Monica	You ' re a , you ' re really good
<i>phredGAN_a</i>	Chandler	No , I do not . You are a mean little man . .
<i>phredGAN_d</i>	Chandler	You ' re right , I can ' t do that .
SM	J/R/M/C ³	Yeah , I ' m sure . Oh , but I can ' t believe it . I ' m not gonna be in here
SAM	J/R/M/C	I ' m not going to be a good time .
<i>hredGAN</i>	NA	Okay , I ' m gonna go see my friend .

²Actors: Leonard/Penny/Howard/Raj

³Actors: Joey/Ross/Monica/Chandler

Table 4: *hredGAN* vs *phredGAN* vs Li et al. (2016b) on UDC

Model	Speaker	Response
UDC		
Context.0	asker	all i need some quick amarok help. i cant hear my music
Context.1	helper	is amarok muted?
Context.2	asker	no
<i>phredGAN_a</i>	helper	use the UNK drivers , and then run the UNK command to get the UNK
<i>phredGAN_d</i>	helper	ok , so you re not using the right driver for the network card , you re using the UNK ?
<i>phred</i>	helper	you can try to install the _UNK package
SM	helper	ok , thanks
SAM	helper	ok , thanks
<i>hredGAN</i>	helper	I have no idea why it would be a bit of a bit.
Context.0	asker	anyone had problems with the kernel update from today? giving me a kernel panic
Context.1	helper	you can select previous kernels at the bootloader (grub) menu on boot.
<i>phredGAN_a</i>	asker	it says that it is not installed . . .
<i>phredGAN_d</i>	asker	ok , so i ll have to reinstall the new kernel , i ll try that.
<i>phred</i>	asker	you can try to install the drivers from the live cd
SM	asker	I ' m not sure what you mean . . .
SAM	asker	I ' m not sure how to do that . . .
<i>hredGAN</i>	asker	I ' m not sure how to do that , but I can ' t boot from a CD . . .
Context.0	asker	how do I install Ubuntu?
<i>phredGAN_a</i>	helper	use the alternate cd , it should be in the repos , its a good place to get the source of the kernel
<i>phredGAN_d</i>	helper	ok , so you have to reinstall the kernel from the CD , and you can t install the iso to the CD
<i>phred</i>	helper	yes
SM	helper	you can use the command line
SAM	helper	what is your question ?
<i>hredGAN</i>	helper	you can use the _UNK to install the _UNK

Table 5: Sample of PHRED outputs on UDC and TV Series

Source	Speaker	Utterance
Big Bang Theory and Friends TV Series		
Context.0	Howard	Okay, you have to understand something, we're in a hospital right now.
Response.0	Penny	Oh , I ' m sorry . I was just trying to be a girl .
Context.0	Sheldon	Did I? Well, that can't be right. No one's ever done that before. Except me, because I just did it. Sheldon and his brain, yeah
Response.0	Leonard	I ' m sorry . I just don ' t want to be a man who ' s going to be my girlfriend .
Context.0	Leonard	She's gotten really hooked on Age of Conan, she's playing non-stop.
Response.0	Raj	Oh , I ' m sorry . I ' m sorry . I ' m sorry . I ' m sorry . I ' m
UDC		
Context.0	asker	all i need some quick amarok help. i cant hear my music
Context.1	helper	is amarok muted?
Context.2	asker	no
Response.0	helper	you can try to install the _UNK package
Context.0	asker	anyone had problems with the kernel update from today? giving me a kernel panic
Context.1	helper	you can select previous kernels at the bootloader (grub) menu on boot.
Response.0	asker	you can try to install the drivers from the live cd
Context.0	asker	how do I install Ubuntu?
Response.0	helper	yes