

Multi-turn Dialogue Response Generation in an Adversarial Learning Framework

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

We propose an adversarial learning approach for generating multi-turn dialogue responses. Our proposed framework, *hredGAN*, is based on conditional generative adversarial networks (GANs). The GAN’s generator is a modified hierarchical recurrent encoder-decoder network (HRED) and the discriminator is a word-level bidirectional RNN that shares context and word embeddings with the generator. During inference, noise samples conditioned on the dialogue history are used to perturb the generator’s latent space to generate several possible responses. The final response is the one ranked best by the discriminator. The *hredGAN* shows improved performance over existing methods: (1) it generalizes better than networks trained using only the log-likelihood criterion, and (2) it generates longer, more informative and more diverse responses with high utterance and topic relevance even with limited training data. This improvement is demonstrated on the Movie triples and Ubuntu dialogue datasets using both automatic and human evaluations.

1 Introduction

Recent advances in deep neural network architectures have enabled tremendous success on a number of difficult machine learning problems. While these results are impressive, producing a deployable neural network-based model that can engage in open domain conversation still remains elusive. A dialogue system needs to be able to generate meaningful and diverse responses that are simultaneously coherent with the input utterance and the overall dialogue topic. Unfortunately, earlier conversation models trained with naturalistic dialogue data suffered greatly from limited contextual information (Sutskever et al., 2014; Vinyals and Le, 2015) and lack of diversity (Li et al., 2016a).

These problems often lead to generic and safe responses to a variety of input utterances.

Serban et al. (2016) and Xing et al. (2017) proposed the Hierarchical Recurrent Encoder-Decoder (HRED) network to capture long temporal dependencies in multi-turn conversations to address the limited contextual information but the diversity problem remained. In contrast, some HRED variants such as variational (Serban et al., 2017b) and multi-resolution (Serban et al., 2017a) HREDs attempt to alleviate the diversity problem by injecting noise at the utterance level and by extracting additional context to condition the generator on. While these approaches achieve a certain measure of success over the basic HRED, the generated responses are still mostly generic since they do not control the generator’s output. This is because the output conditional distribution is not calibrated. Li et al. (2016a), on the other hand, consider a diversity promoting training objective but their model is for single turn conversations and cannot be trained end-to-end.

The generative adversarial network (GAN) (Goodfellow et al., 2014) seems to be an appropriate solution to the diversity problem. GAN matches data from two different distributions by introducing an adversarial game between a *generator* and a *discriminator*. We explore *hredGAN*: conditional GANs for multi-turn dialogue models with an HRED generator and discriminator. *hredGAN* combines ideas from both generative and retrieval-based multi-turn dialogue systems to improve their individual performances. This is achieved by sharing the context and word embeddings between the generator and the discriminator allowing for joint end-to-end training using back-propagation. To the best of our knowledge, no existing work has applied conditional GANs to multi-turn dialogue models and especially not

with HRED generators and discriminators. We demonstrate the effectiveness of hredGAN over the VHRED for dialogue modeling with evaluations on the Movie triples and Ubuntu technical support datasets.

2 Related Work

Our work is related to end-to-end neural network-based open domain dialogue models. Most neural dialogue models use transduction frameworks adapted from neural machine translation (Sutskever et al., 2014; Bahdanau et al., 2015). These Seq2Seq networks are trained end-to-end with MLE criteria using large corpora of human-to-human conversation data. Others use GAN’s discriminator as a reward function in a reinforcement learning framework (Yu et al., 2017) and in conjunction with MLE (Li et al., 2017; Che et al., 2017). Zhang et al. (2017) explored the idea of GAN with a feature matching criterion. Xu et al. (2017) and Zhang et al. (2018) employed GAN with an approximate embedding layer as well as with adversarial information maximization respectively to improve Seq2Seq’s diversity performance.

Still, Seq2Seq models are limited in their ability to capture long temporal dependencies in multi-turn conversation. Although Li et al. (2016b) attempted to optimize a pair of Seq2Seq models for multi-turn dialogue, the multi-turn objective is only applied at inference and not used for actual model training. Hence the introduction of HRED models (Serban et al., 2016, 2017a,b; Xing et al., 2017) for modeling dialogue response in multi-turn conversations. However, these HRED models suffer from lack of diversity since they are only trained with MLE criteria. On the other hand, adversarial system has been used for evaluating open domain dialogue models (Bruni and Fernandez, 2018; Kannan and Vinyals, 2017). Our work, hredGAN, is closest to the combination of HRED generation models (Serban et al., 2016) and adversarial evaluation (Kannan and Vinyals, 2017).

3 Model

3.1 Adversarial Learning of Dialogue Response

Consider a dialogue consisting of a sequence of N utterances, $\mathbf{x} = (x_1, x_2, \dots, x_N)$, 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 . At any time step i , the dialogue history is given by $\mathbf{x}_i = (x_1, x_2, \dots, x_i)$. The dialogue response generation task can be defined as follows: Given a dialogue history \mathbf{x}_i , generate a response $y_i = (y_i^1, y_i^2, \dots, y_i^{T_i})$, where T_i is the number of generated tokens. We also want the distribution of the generated response $P(y_i)$ to be indistinguishable from that of the ground truth $P(x_{i+1})$ and $T_i = M_{i+1}$. Conditional GAN learns a mapping from an observed dialogue history, \mathbf{x}_i , and a sequence of random noise vectors, z_i to a sequence of output tokens, y_i , $G : \{\mathbf{x}_i, z_i\} \rightarrow y_i$. The generator G is trained to produce output sequences that cannot be distinguished from the ground truth sequence by an adversarially trained discriminator D that is trained to do well at detecting the generator’s fakes. The distribution of the generator output sequence can be factored by the product rule:

$$P(y_i | \mathbf{x}_i) = P(y_i^1) \prod_{j=2}^{T_i} P(y_i^j | y_i^1, \dots, y_i^{j-1}, \mathbf{x}_i) \quad (1)$$

$$P(y_i^j | y_i^1, \dots, y_i^{j-1}, \mathbf{x}_i) = P_{\theta_G}(y_i^{1:j-1}, \mathbf{x}_i) \quad (2)$$

where $y_i^{1:j-1} = (y_i^1, \dots, y_i^{j-1})$ and θ_G are the parameters of the generator model. $P_{\theta_G}(y_i^{1:j-1}, \mathbf{x}_i)$ is an autoregressive generative model where the probability of the current token depends on the past generated sequence. Training the generator G with the log-likelihood criterion is unstable in practice, and therefore the past generated sequence is substituted with the ground truth, a method known as *teacher forcing* (Williams and Zipser, 1989), i.e.,

$$P(y_i^j | y_i^1, \dots, y_i^{j-1}, \mathbf{x}_i) \approx P_{\theta_G}(x_{i+1}^{1:j-1}, \mathbf{x}_i) \quad (3)$$

Using (3) in relation to GAN, we define our fake sample as the teacher forcing output with some input noise z_i

$$y_i^j \sim P_{\theta_G}(x_{i+1}^{1:j-1}, \mathbf{x}_i, z_i) \quad (4)$$

and the corresponding real sample as ground truth x_{i+1}^j .

With the GAN objective, we can match the noise distribution, $P(z_i)$, to the distribution of the ground truth response, $P(x_{i+1} | \mathbf{x}_i)$. Varying the

noise input then allows us to generate diverse responses to the same dialogue history. Furthermore, the discriminator, since it is calibrated, is used during inference to rank the generated responses, providing a means of controlling the generator output.

3.1.1 Objectives

The objective of a conditional GAN can be expressed as

$$\mathcal{L}_{cGAN}(G, D) = \mathbb{E}_{\mathbf{x}_i, x_{i+1}} [\log D(x_{i+1}, \mathbf{x}_i)] + \mathbb{E}_{\mathbf{x}_i, z_i} [1 - \log D(G(\mathbf{x}_i, z_i), \mathbf{x}_i)] \quad (5)$$

where G tries to minimize this objective against an adversarial D that tries to maximize it:

$$G^*, D^* = \arg \min_G \max_D \mathcal{L}_{cGAN}(G, D). \quad (6)$$

Previous approaches have shown that it is beneficial to mix the GAN objective with a more traditional loss such as cross-entropy loss (Lamb et al., 2016; Li et al., 2017). The discriminator’s job remains unchanged, but the generator is tasked not only to fool the discriminator but also to be near the ground truth x_{i+1} in the cross-entropy sense:

$$\mathcal{L}_{MLE}(G) = \mathbb{E}_{\mathbf{x}_i, x_{i+1}, z_i} [-\log P_{\theta_G}(x_{i+1}, \mathbf{x}_i, z_i)]. \quad (7)$$

Our final objective is,

$$G^*, D^* = \arg \min_G \max_D (\lambda_G \mathcal{L}_{cGAN}(G, D) + \lambda_M \mathcal{L}_{MLE}(G)). \quad (8)$$

It is worth mentioning that, without z_i , the net could still learn a mapping from \mathbf{x}_i to y_i , but it would produce deterministic outputs and fail to match any distribution other than a delta function (Isola et al., 2017). This is one key area where our work is different from Lamb et al.’s and Li et al.’s. The schematic of the proposed hredGAN is depicted at the right hand side of Figure 1.

3.1.2 Generator

We adopted an HRED dialogue generator similar to Serban et al. (2016, 2017a,b) and Xing et al. (2017). The HRED contains three recurrent structures, i.e. the encoder ($eRNN$), context ($cRNN$), and decoder ($dRNN$) RNN. The conditional probability modeled by the HRED per output word token is given by

$$P_{\theta_G}(y_i^j | x_{i+1}^{1:j-1}, \mathbf{x}_i) = dRNN(E(x_{i+1}^{j-1}), h_i^{j-1}, \mathbf{h}_i) \quad (9)$$

where $E(\cdot)$ is the embedding lookup, $\mathbf{h}_i = cRNN(eRNN(E(x_i), \mathbf{h}_{i-1}), eRNN(\cdot))$ maps a sequence of input symbols into fixed-length vector, and h and \mathbf{h} are the hidden states of the decoder and context RNN, respectively.

In the multi-resolution HRED, (Serban et al., 2017a), high-level tokens are extracted and processed by another RNN to improve performance. We circumvent the need for this extra processing by allowing the decoder to attend to different parts of the input utterance during response generation (Bahdanau et al., 2015; Luong et al., 2015). We introduce a local attention into (9) and encode the attention memory differently from the context through an attention encoder RNN ($aRNN$), yielding:

$$P_{\theta_G}(y_i^j | x_{i+1}^{1:j-1}, \mathbf{x}_i) = dRNN(E(x_{i+1}^{j-1}), h_i^{j-1}, a_i^j, \mathbf{h}_i) \quad (10)$$

where $a_i^j = \sum_{m=1}^{M_i} \frac{\exp(\alpha_m)}{\sum_{m=1}^{M_i} \exp(\alpha_m)} h_i^{\prime m}$, $h_i^{\prime m} = aRNN(E(x_i^m), h_i^{\prime m-1})$, h_i^{\prime} is the hidden state of the attention RNN, and α_k is either a logit projection of $(h_i^{j-1}, h_i^{\prime m})$ in the case of Bahdanau et al. (2015) or $(h_i^{j-1})^T \cdot h_i^{\prime m}$ in the case of Luong et al. (2015). The modified HRED architecture is shown in Figure 2.

Noise Injection: We inject Gaussian noise at the input of the decoder RNN. Noise samples could be injected at the utterance or word level. With noise injection, the conditional probability of the decoder output becomes

$$P_{\theta_G}(y_i^j | x_{i+1}^{1:j-1}, z_i^j, \mathbf{x}_i) = dRNN(E(x_{i+1}^{j-1}), h_i^{j-1}, a_i^j, z_i^j, \mathbf{h}_i) \quad (11)$$

where $z_i^j \sim \mathcal{N}_i(0, \mathbf{I})$, for utterance-level noise and $z_i^j \sim \mathcal{N}_i^j(0, \mathbf{I})$, for word-level noise.

3.1.3 Discriminator

The discriminator shares context and word embeddings with the generator and can discriminate at the word level (Lamb et al., 2016). The word-level discrimination is achieved through a bidirectional RNN and is able to capture both syntactic and conceptual differences between the generator output and the ground truth. The aggregate classification of an input sequence, χ can be factored over word-

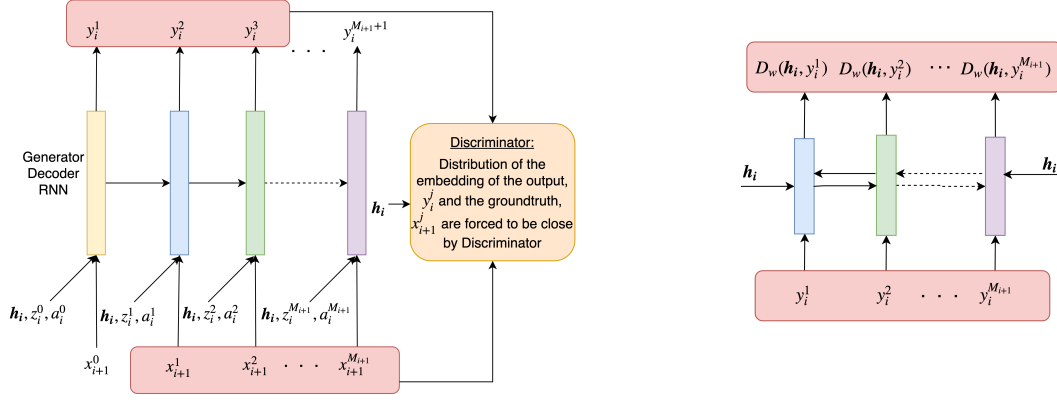


Figure 1: **Left: The hredGAN architecture** - The generator makes predictions conditioned on the dialogue history, \mathbf{h}_i , attention, a_i^j , noise sample, z_i^j , and ground truth, x_{i+1}^{j-1} . **Right: RNN-based discriminator** that discriminates bidirectionally at the word level.

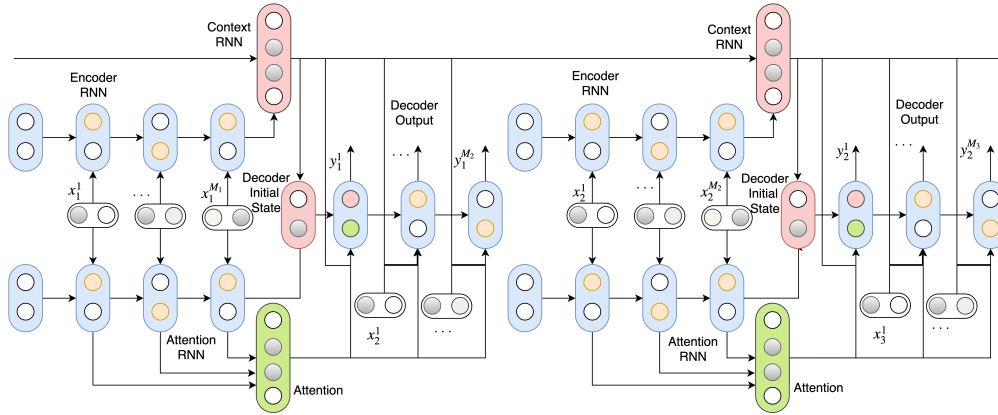


Figure 2: **The HRED generator with local attention** - The attention RNN ensures local relevance while the context RNN ensures global relevance. Their states are combined to initialize the decoder RNN and the discriminator BiRNN.

level discrimination and expressed as

$$D(\mathbf{x}_i, \chi) = D(\mathbf{h}_i, \chi) = \left[\prod_{j=1}^J D_{RNN}(\mathbf{h}_i, E(\chi^j)) \right]^{\frac{1}{J}} \quad (12)$$

where $D_{RNN}(\cdot)$ is the word discriminator RNN, \mathbf{h}_i is an encoded vector of the dialogue history \mathbf{x}_i obtained from the generator's $cRNN(\cdot)$ output, and χ^j is the j th word or token of the input sequence χ . $\chi = y_i$ and $J = T_i$ for the case of generator's decoder output, $\chi = x_{i+1}$ and $J = M_{i+1}$ for the case of ground truth. The discriminator architecture is depicted on the left hand side of Figure 1.

3.2 Adversarial Generation of Multi-turn Dialogue Response

In this section, we describe the generation process during inference. The generation objective can be

mathematically described as

$$y_i^* = \arg \max_{\mathbf{y}_i} \left\{ P(y_{i,l} | \mathbf{x}_i) + D^*(\mathbf{x}_i, y_{i,l}) \right\}_{l=1}^L \quad (13)$$

where $y_{i,l} = G^*(\mathbf{x}_i, z_{i,l})$, $z_{i,l}$ is the l th noise samples at dialogue step i , and L is the number of response samples. Equation 13 shows that our inference objective is the same as the training objective (8), combining both the MLE and adversarial criteria. This is in contrast to existing work where the discriminator is usually discarded during inference.

The inference described by (13) is intractable due to the enormous search space of $y_{i,l}$. Therefore, we turn to an approximate solution where we use greedy decoding (MLE) on the first part of the objective function to generate L lists of responses based on noise samples $\{z_{i,l}\}_{l=1}^L$. In order to facilitate the exploration of the generator's latent space, we sample a modified noise distribution, $z_{i,l}^j \sim \mathcal{N}_{i,l}(0, \alpha \mathbf{I})$, or $z_{i,l}^j \sim \mathcal{N}_{i,l}^j(0, \alpha \mathbf{I})$

Algorithm 1 Adversarial Learning of hredGAN

Require: A generator G with parameters θ_G .

Require: A discriminator D with parameters θ_D .

```
for number of training iterations do
  Initialize  $cRNN$  to zero_state,  $\mathbf{h}_0$ 
  Sample a mini-batch of conversations,  $\mathbf{x} = \{x_i\}_{i=1}^N$ ,  $\mathbf{x}_i = (x_{i1}, x_{i2}, \dots, x_{iM_i})$  with  $N$  utterances. Each utterance mini batch  $i$  contains  $M_i$  word tokens.
  for  $i = 1$  to  $N - 1$  do
    Update the context state.
     $\mathbf{h}_i = cRNN(eRNN(E(x_i)), \mathbf{h}_{i-1})$ 
    Compute the generator output using (11).
     $P_{\theta_G}(y_i | z_i, \mathbf{x}_i) = \{P_{\theta_G}(y_i^j | x_{i+1}^{1:j-1}, z_i^j, \mathbf{x}_i)\}_{j=1}^{M_{i+1}}$ 
    Sample a corresponding mini batch of utterance  $y_i$ .
     $y_i \sim P_{\theta_G}(y_i | z_i, \mathbf{x}_i)$ 
  end for
  Compute the discriminator accuracy  $D_{acc}$  over  $N - 1$  utterances  $\{y_i\}_{i=1}^{N-1}$  and  $\{x_{i+1}\}_{i=1}^{N-1}$ 
  if  $D_{acc} < acc_{D_{th}}$  then
    Update  $\theta_D$  with gradient of the discriminator loss.
     $\sum_i [\nabla_{\theta_D} \log D(\mathbf{h}_i, x_{i+1}) + \nabla_{\theta_D} \log(1 - D(\mathbf{h}_i, y_i))]$ 
  end if
  if  $D_{acc} < acc_{G_{th}}$  then
    Update  $\theta_G$  with the generator's MLE loss only.
     $\sum_i [\nabla_{\theta_G} \log P_{\theta_G}(y_i | z_i, \mathbf{x}_i)]$ 
  else
    Update  $\theta_G$  with both adversarial and MLE losses.
     $\sum_i [\lambda_G \nabla_{\theta_G} \log D(\mathbf{h}_i, y_i) + \lambda_M \nabla_{\theta_G} \log P_{\theta_G}(y_i | z_i, \mathbf{x}_i)]$ 
  end if
end for
```

where $\alpha > 1.0$, is the exploration factor that increases the noise variance. We then rank the L lists using the discriminator score, $\{D^*(\mathbf{x}_i, y_{i,l})\}_{l=1}^L$. The response with the highest discriminator ranking is the optimum response for the dialogue context.

4 Training of hredGAN

We trained both the generator and the discriminator simultaneously as highlighted in Algorithm ?? with $\lambda_G = \lambda_M = 1$. GAN training is prone to instability due to competition between the generator and the discriminator. Therefore, parameter updates are conditioned on the discriminator performance (Lamb et al., 2016).

The generator consists of four RNNs with different parameters, that is, $aRNN$, $eRNN$, $cRNN$, and $dRNN$. $aRNN$ and $eRNN$ are both bidirectional, while $cRNN$ and $dRNN$ are unidirectional. Each RNN has 3 layers, and the hidden state size is 512. The $dRNN$ and $aRNN$ are connected using an additive attention mechanism (Bahdanau et al., 2015).

The discriminator shares $aRNN$, $eRNN$, and $cRNN$ with the generator. D_{RNN} is a stacked bidirectional RNN with 3 layers and a hidden state size of 512. The $cRNN$ states are used to initialize the states of D_{RNN} . 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.

Others: All RNNs used are gated recurrent unit (GRU) cells (Cho et al., 2014). The word embedding size is 512 and shared between the generator and the discriminator. The initial learning rate is 0.5 with decay rate factor of 0.99, applied when the adversarial loss has increased over two iterations. We use a batch size of 64 and clip gradients around 5.0. As in Lamb et al. (2016), we find $acc_{D_{th}} = 0.99$ and $acc_{G_{th}} = 0.75$ to suffice. All parameters are initialized with Xavier uniform random initialization (Glorot and Bengio, 2010). The vocabulary size V is 50,000. Due to the large 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 evaluation. The model is trained end-to-end using the stochastic gradient descent algorithm.

5 Experiments and Results

We consider the task of generating dialogue responses conditioned on the dialogue history and the current input utterance. We compare the proposed hredGAN model against some alternatives on publicly available datasets.

5.1 Datasets

Movie Triples Corpus (MTC) dataset (Serban et al., 2016). This dataset was derived from the *Movie-DiC* dataset by Banchs (2012). Although this dataset spans a wide range of topics with few spelling mistakes, its small size of only about 240,000 dialogue triples makes it difficult to train a dialogue model, as pointed out by Serban et al. (2016). We thought that this scenario would really benefit from the proposed adversarial generation.

Ubuntu Dialogue Corpus (UDC) dataset (Serban et al., 2017b). This dataset was extracted from the Ubuntu Relay Chat Channel. Although the topics in the dataset are not as diverse as in the MTC, the dataset is very large, containing about 1.85 million conversations with an average of 5 utterances per conversation.

We split both MTC and UDC into training, validation, and test sets, using 90%, 5%, and 5% proportions, respectively. We performed minimal pre-processing of the datasets by replacing all words

except the top 50,000 most frequent words by an *UNK* symbol.

5.2 Evaluation Metrics

Accurate evaluation of dialogue models is still an open challenge. In this paper, we employ both automatic and human evaluations.

5.2.1 Automatic Evaluation

We employed some of the automatic evaluation metrics that are used in probabilistic language and dialogue models, and statistical machine translation. Although these metrics may not correlate well with human judgment of dialogue responses (Liu et al., 2016), they provide a good baseline for comparing dialogue model performance.

Perplexity - For a model with parameter θ , we define perplexity as:

$$\exp \left[-\frac{1}{N_W} \sum_{k=1}^K \log P_{\theta}(y_1, y_2, \dots, y_{N_k-1}) \right] \quad (14)$$

where K is the number of conversations in the dataset, N_k is the number of utterances in conversation k , and N_W is the total number of word tokens in the entire dataset. The lower the perplexity, the better. The perplexity measures the likelihood of generating the ground truth given the model parameters. While a generative model can generate a diversity of responses, it should still assign a high probability to the ground truth utterance.

BLEU - The BLEU score (Papineni et al., 2002) provides a measure of overlap between the generated response (candidate) and the ground truth (reference) using a modified n-gram precision. According to Liu et. al. (Liu et al., 2016), BLEU-2 score is fairly correlated with human judgment for non-technical dialogue (such as MTC).

ROUGE - The ROUGE score (Lin, 2014) is similar to BLEU but it is recall-oriented instead. It is used for automatic evaluation of text summarization and machine translation. To compliment the BLEU score, we use ROUGE-N with $N = 2$ for our evaluation.

Distinct n-gram - This is the fraction of unique n-grams in the generated responses and it provides a measure of diversity. Models with higher a number of distinct n-grams tend to produce more diverse responses (Li et al., 2016a). For our evaluation, we use 1- and 2- grams.

Normalized Average Sequence Length (NASL) - This measures the average number of

words in model-generated responses normalized by the average number of words in the ground truth.

5.2.2 Human Evaluation

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 presented 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 and informativeness. 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 estimated the per sample score variance between judges and then averaged 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 judgment for the model.

5.3 Baseline

We compare the performance of our model to (V)HRED (Serban et al., 2016, 2017b), since they are the closest to our approach in implementation and are the current state of the art in open-domain dialogue models. HRED is very similar to our proposed generator, but without the input utterance attention and noise samples. VHRED introduces a latent variable to the HRED between the *cRNN* and the *dRNN* and was trained using the variational lower bound on the log-likelihood. The VHRED can generate multiple responses per context like hredGAN, but it has no specific criteria for selecting the best response.

The HRED and VHRED models are both trained using the Theano-based implementation obtained from <https://github.com/julianser/hed-dlg-truncated>. The training and validation sets used for UDC and MTC dataset were obtained directly from the authors¹ of (V)HRED. For model comparison, we use a test set that is disjoint from the training and validation sets.

¹UDC was obtained from <http://www.iulianserban.com/Files/UbuntuDialogueCorpus.zip>, and the link to MTC was obtained privately.

Model	Teacher Forcing		Autoregression				Human Evaluation
	Perplexity	$-\log D(G(\cdot))$	BLEU-2	ROUGE-2	DISTINCT-1/2	NASL	
MTC							
HRED	31.92/36.00	NA	0.0474	0.0384	0.0026/0.0056	0.535	0.2560 ± 0.0977
VHRED	42.61/44.97	NA	0.0606	0.1181	0.0048/0.0163	0.831	0.3909 ± 0.0240
hredGAN_u	23.57/23.54	6.85/6.81	0.0493	0.2416	0.0167/0.1306	0.884	0.5582 ± 0.0118
hredGAN_w	24.20/24.14	13.35/13.40	0.0613	0.3244	0.0179/0.1720	1.540	0.7869 ± 0.1148
UDC							
HRED	69.39/86.40	NA	0.0177	0.0483	0.0203/0.0466	0.892	0.3475 ± 0.1062
VHRED	98.50/105.20	NA	0.0171	0.0855	0.0297/0.0890	0.873	0.4046 ± 0.0188
hredGAN_u	56.82/57.32	10.09/10.08	0.0137	0.0716	0.0260/0.0847	1.379	0.6133 ± 0.0361
hredGAN_w	47.73/48.18	8.37/8.36	0.0216	0.1168	0.0516/0.1821	1.098	0.6905 ± 0.0706

Table 1: Generator Performance Evaluation

5.4 Results

We have two variants of hredGAN based on the noise injection approach, i.e., hredGAN with utterance-level (*hredGAN_u*) and word-level (*hredGAN_w*) noise injections.

We compare the performance of these two variants with HRED and VHRED models.

Perplexity: The average perplexity per word performance of all the four models on MTC and UDC datasets (validation/test) are reported in the first column on Table 1. The table indicates that both variants of the hredGAN model perform better than the HRED and VHRED models in terms of the perplexity measure. However, using the adversarial loss criterion (Eq. (8)), the hredGAN_u model performs better on MTC and worse on UDC. Note that, for this experiment, we run all models in teacher forcing mode.

Generation Hyperparameter: For adversarial generation, we perform a linear search for α between 1 and 20 at an increment of 1 using Eq. (13), with sample size $L = 64$, on validation sets with models run in autoregression. The optimum values of α for hredGAN_u and hredGAN_w for UDC are 7.0 and 9.0 respectively. The values for MTC are not convex, probably due to small size of the dataset, so we use the same α values as UDC. We however note that for both datasets, any integer value between 3 and 10 (inclusive) works well in practice.

Quantitative Generator Performance: We run autoregressive inference for all the models (using optimum α values for hredGAN models and selecting the best of $L = 64$ responses using a discriminator) with dialogue contexts from a unique test set. Also, we compute the average BLEU-2, ROUGE-2(f1), Distinct(1/2), and normalized

Item	$D(G(\cdot))$	Utterance
MTC		
Context.0	NA	perhaps <person> had a word with the man upstairs .
Context.1	NA	a word ? i ' m sure by now he ' s engineered a hostile takeover .
Response.0	0.996	<person> , i know what you ' re saying , <person> , that ' s not what i ' m saying .
Response.1	0.991	<person> , i know . i was just about to help the guy .
Response.2	0.315	<person> , i ' m sorry .
Response.3	0.203	<person> , i ' m a little out .
<hr/>		
Context.0	NA	says he wanted food . <person> . he wanted the gold .
Context.1	NA	how ' s he going to want the gold ? he couldn ' t even know we had it .
Response.0	0.998	<person> , i know . but it ' s not him , it ' s the only way he ' s got it all figured
Response.1	0.981	<person> , i know . but i have to tell you . these things are really stupid and you think i was wrong ?
Response.2	0.690	<person> , i ' m sure he did .
Response.3	0.314	<person> , i ' m not sure .
<hr/>		
UDC		
Context.0	NA	The netboot one is suppose to download packages from the net.
Context.1	NA	like the ones to be installed? or the installed to be run?
Response.0	0.993	you don ' t need to install the whole system , just install the ubuntu installer
Response.1	0.952	you can install the ubuntu installer from the ubuntu menu
Response.2	0.749	I ' m not sure , I don ' t know .
Response.3	0.184	you can
Response.4	0.003	you a . 1
<hr/>		
Context.0	NA	DJones: update manager won't detect 12.04.1 as a new version if you are already running 12.04, because 12.04.1 = 12.04 + lots of package updates
Response.0	0.991	did you try a clean install of the latest version ?
Response.1	0.981	try installing the latest _UNK and see if it works
Response.2	0.615	I ' m not sure you have a problem . . .
Response.3	0.191	try sudo apt-get remove the package that is not installed
Response.4	0.002	try the _UNK . 1 . 1 . 1 . 1 . 1 . _UNK . deb

Table 2: Example of Discriminator Ranking with hredGAN_w

average sequence length (NASL) scores for each model and summarize the results in the middle of Table 1. Distinct(1/2) largely agrees with the perplexity score. Most scores, similar to the perplexity, indicate that hredGAN models perform better than (V)HRED on both datasets. However, on the UDC ROUGE and MTC BLEU, VHRED scores slightly better than hredGAN_u but still worse than hredGAN_w.

A good dialogue model should find the right balance between precision (BLEU) and diversity. We strongly believe that our adversarial approach is better suited to solving this problem.

Movie	
Context.0	perhaps <person> had a word with the man upstairs .
HRED	i ' m sorry , <person> .
VHRED	<person> ' t do it .
hredGAN_w	<person> , i ' m not gon na be okay . i just don ' t think it ' s a good idea .
hredGAN_u	<person> , i cause it was all the way to you .
Context.1	a word ? i ' m sure by now he ' s engineered a hostile takeover .
HRED	i ' m sorry , <person> .
VHRED	<person> ' t do it .
hredGAN_w	<person> , i know what you ' re saying , <person> , that ' s not what i ' m saying .
hredGAN_u	<person> , i have to tell you i ' m not the same .
Context.0	<person> ' s .
HRED	i ' m sorry , <person> .
VHRED	<person> ' t do it .
hredGAN_w	<person> , i ' m sure it ' s the only thing i ever imagined .
hredGAN_u	<person> , i ' ll be right there .
Context.1	can i ask you something about her ?
HRED	i don ' t know .
VHRED	you ' re not going to be here .
hredGAN_w	i don ' t know . i think i ' ve heard of it .
hredGAN_u	i ' m sorry , i can ' t .
Ubuntu	
Context.0	this is Dacs i am in i used 'sudo statx and it let me in
HRED	what ' s the problem ?
VHRED	I don't know what you want to do . if you want to do it , you need to do it to do it in the terminal ?
hredGAN_w	you don ' t have to do anything , just type in the terminal and it should be to find the right device
hredGAN_u	you can have a look at the output of the command . . .
Context.1	good deal... cat /etc/X11/default-display-manager
HRED	what ' s the problem ?
VHRED	do you know what you want to do ?
hredGAN_w	I ' m trying to figure a command that I can find to find out the file that I can find in the file"
hredGAN_u	I don ' t see the point , but I ' m not sure how to do that .
Context.2	/usr/sbin/lightdm http://paste.ubuntu.com/1286224/ <--- my
HRED	/etc/X11/xorg.conf
VHRED	what ' s the problem ?
hredGAN_w	is there a way to do that in the terminal ?
hredGAN_u	did you just type \" sudo mount -a \" ?
hredGAN_u	i have no idea , i just installed ubuntu and i have no idea how to do that

Table 3: Sample responses of HRED, VHRED and hredGAN.

As hredGAN generators explore diversity, the discriminator ranking gives hredGAN an edge over (V)HRED because it helps detect responses that are out of context and the natural language structure (Table 2). Also, the ROGUE(f1) performance indicates that hredGAN_w strikes a better balance between precision (BLEU) and diversity than the rest of the models. This is also obvious from the quality of generated responses.

Qualitative Generator Performance: The results of the human evaluation are reported in the last column of Table 1. The human evaluation agrees largely with the automatic evaluation. hredGAN_w performs best on both datasets although the gap is more on the MTC than on the UTC. This implies that the improvement of HRED with adversarial generation is better than with variational generation (VHRED). In addition, looking at the actual samples from the generator outputs in Table 6 shows that hredGAN, especially hredGAN_w, performs better than (V)HRED. While other models produce short and generic ut-

terances, hredGAN_w mostly yields informative responses. For example, in the first dialogue in Table 6, when the speaker is sarcastic about “the man upstairs”, hredGAN_w responds with the most coherent utterance with respect to the dialogue history. We see similar behavior across other samples. We also note that although hredGAN_u’s responses are the longest on Ubuntu (in line with the NASL score), the responses are less informative compared to hredGAN_w resulting in a lower human evaluation score. We reckon this might be due to a mismatch between utterance-level noise and word-level discrimination or lack of capacity to capture the data distribution using single noise distribution. We hope to investigate this further in the future.

Discriminator Performance: Although only hredGAN uses a discriminator, the observed discriminator behavior is interesting. We observe that the discriminator score is generally reasonable with longer, more informative and more persona-related responses receiving higher scores as shown in Table 2. It worth to note that this behavior, although similar to the behavior of a human judge is learned without supervision. Moreover, the discriminator seems to have learned to assign an average score to more frequent or generic responses such as “I don’t know,” “I’m not sure,” and so on, and high score to rarer answers. That’s why we sample a modified noise distribution during inference so that the generator can produce rarer utterances that will be scored high by the discriminator.

6 Conclusion and Future Work

In this paper, we have introduced an adversarial learning approach that addresses response diversity and control of generator outputs, using an HRED-derived generator and discriminator. The proposed system outperforms existing state-of-the-art (V)HRED models for generating responses in multi-turn dialogue with respect to automatic and human evaluations. The performance improvement of the adversarial generation (hredGAN) over the variational generation (VHRED) comes from the combination of adversarial training and inference which helps to address the lack of diversity and contextual relevance in maximum likelihood based generative dialogue models. Our analysis also concludes that the word-level noise injection seems to perform better in general.

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.
- E. Bruni and R. Fernandez. 2018. Adversarial evaluation for open-domain dialogue generation. In *Proceedings of the 18th Annual SIGdial Meeting*.
- T. Che, Y. Li, R. Zhang, R. D. Hjelm, W. Li, Y. Song, and Y. Bengio. 2017. Maximum-likelihood augmented discrete generative adversarial networks. In *arXiv preprint arXiv:1702.07983*.
- K. Cho, B. Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. In *Proceedings of International Conference of Learning Representation (ICLR 2015)*, pages 1724–1734.
- X. Glorot and Y. Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *International conference on artificial intelligence and statistics*.
- I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. 2014. Generative adversarial nets. In *Proceedings of Advances in Neural Information Processing Systems (NIPS 2014)*.
- P. Isola, J. Y. Zhu, T. Zhou, and A. A. Efros. 2017. Image-to-image translation with conditional adversarial networks. In *Conference on Computer Vision and Pattern Recognition (CVPR, 2017)*.
- 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*.
- A. Kannan and O. Vinyals. 2017. Adversarial evaluation of dialogue models. In *arXiv preprint arXiv:1701.08198v1*.
- A. Lamb, A. Goyal, Y. Zhang, S. Zhang, A. Courville, and Y. Bengio. 2016. Professor forcing: A new algorithm for training recurrent networks. In *Proceedings of Advances in Neural Information Processing Systems (NIPS 2016)*.
- 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, W. Monroe, A. Ritter, M. Galley, J. Gao, and D. Jurafsky. 2016b. Deep reinforcement learning for dialogue generation. In *arXiv preprint arXiv:arXiv:arXiv:1606.01541v4*.
- J. Li, W. Monroe, T. Shi, A. Ritter, and D. Jurafsky. 2017. Adversarial learning for neural dialogue generation. In *arXiv preprint arXiv:1701.06547*.
- C. Y. Lin. 2014. Rouge: a package for automatic evaluation of summaries. In *Proceedings of the Workshop on Text Summarization Branches Out*.
- C. Liu, R. Lowe, I. V. Serban, M. Noseworthy, L. Charlin, and J. Pineau. 2016. How not to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. In *Proceedings of EMNLP*, pages 2122–2132.
- M. T. Luong, I. Sutskever, Q. V. Le, O. Vinyals, and W. Zaremba. 2015. Addressing the rare word problem in neural machine translation. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics*.
- 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*.
- R. J. Williams and D. Zipser. 1989. A learning algorithm for continually running fully recurrent neural networks. *Neural computation*, 1(2):270–280.
- 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*.

- Z. Xu, B. Liu, B. Wang, S. Chengjie, X. Wang, Z. Wang, and C. Qi. 2017. Neural response generation via gan with an approximate embedding layer. In *EMNLP*.
- L. Yu, W. Zhang, J. Wang, and Y. Yu. 2017. Seqgan: sequence generative adversarial nets with policy gradient. In *Proceedings of The Thirty-first AAAI Conference on Artificial Intelligence (AAAI 2017)*.
- Y. Zhang, M. Galley, J. Gao, Z. Gan, X. Li, C. Brockett, and B. Dolan. 2018. Generating informative and diverse conversational responses via adversarial information maximization. In *arXiv preprint arXiv:arXiv:1809.05972v5*.
- Y. Zhang, Z. Gan, K. Fan, Z. Chen, R. Henao, D. Shen, and L. Carin. 2017. Adversarial feature matching for text generation. In *arXiv preprint arXiv:1706.03850*.

A Ablation Experiments

Before proposing the above adversarial learning framework for multi-turn dialogue, we carried out some experiments.

A.1 Generator:

We consider two main factors here, i.e., addition of an attention memory and injection of Gaussian noise into the generator input.

A.1.1 Addition of Attention Memory

First, we noted that by adding an additional attention memory to the HRED generator, we improved the test set perplexity score by more than 12 and 25 points on the MTC and UDC respectively as shown in Table 4. The addition of attention also shows strong performance at autoregressive inference across multiple metrics as well as an observed improvement in response quality. Hence the decision for the modified HRED generator.

A.1.2 Injection of Noise

Before injecting noise into the generator, we first train hredGAN without noise. The result is also reported in 4. We observe accelerated generator training but without an appreciable improvement in performance. It seems the discrimination task is very easy since there is no stochasticity in the generator output. Therefore, the adversarial feedback does not meaningfully impact the generator weight update.

Finally, we also notice that even with noise injection, there is no appreciable improvement in the auto-regressive performance if we sample with

$L = 1$ even though the perplexity is higher. However, as we increase L , producing L responses per turn, the discriminator’s adversarial selection gives a better performance as reported in Table 1.

Therefore, we conclude that the combination of adversarial training and adversarial inference helps to address the lack of diversity and contextual relevance observed in the generated responses.

A.2 Discriminator:

Before deciding on the word-level discrimination, we experimented with utterance-level discrimination. The utterance-level discriminator trains very quickly but it leads to mostly generic responses from the generator. We also note that utterance-level discriminator scores are mostly extreme (i.e., either low or high). Since we had used a convolutional neural network discriminator (Yu et al., 2017) in our experiments, we hope to investigate this further with other architectures.

A.3 Adversarial Training:

Lastly, we also tried a basic policy gradient approach (Li et al., 2017), where word-level discriminator score is used as a reward for each generated word token, but this leads to training instability. This is probably due to the instability of Monte Carlo sampling over a large vocabulary size. We believe this might improve with other sampling methods such as importance sampling and hope to investigate this further in the future.

Model	Teacher Forcing	Autoregression			
	Perplexity	BLEU-2	ROUGE-2	DISTINCT-1/2	NASL
MTC					
HRED	31.92/36.00	0.0474	0.0384	0.0026/0.0056	0.535
HRED+Attn	18.70/19.02	0.0425	0.2239	0.0397/0.1567	0.527
hredGAN_no_noise	18.93/19.19	0.0355	0.1839	0.0272/0.0978	0.471
UDC					
HRED	69.39/86.40	0.0177	0.0483	0.0203/0.0466	0.892
HRED+Attn	43.43/43.92	0.0140	0.0720	0.0473/0.1262	0.760
hredGAN_no_noise	43.48/44.04	0.0123	0.0827	0.0398/0.1147	0.908

Table 4: Generator Performance: HRED, HRED+Attn and hredGAN without noise

Model	Response
Context.0 HRED VHRED hredGAN_w hredGAN_u	<person> i ' ve raised my rate . i might even cut my sessions to forty minutes . i ' m sorry , <person> . <person> , <person> , <person> , <person> . <person> , i ' m sorry , but i ' m not sure if you ' re really a very good guy . i ' ll never see them again , <person> .
Context.0 Context.1 HRED VHRED hredGAN_w hredGAN_u	i ' m going to handle this exactly the way your father would have handled this . i ' ve got it solved . you ' ve got it solved ? where is she ? i don ' t know . <person> , i ' m sorry . to the end of a beautiful woman who will be very happy . i ' m sorry , i just didn ' t know -
Context.0 HRED VHRED hredGAN_w hredGAN_u	of course not ! <person> ' s like a daughter to me . her father and i taught at university together . a good man , but he never listened . i ' m sorry , <person> . <person> , mr . <person> . <person> , i know it ' s been all six years since i told him about you . i ' m sure you didn ' t , but i would have done a little more .
Context.0 Context.1 HRED VHRED hredGAN_w hredGAN_u	of course not ! <person> ' s like a daughter to me . her father and i taught at university together . a good man , but he never listened . neither did i . i ' m sorry , <person> . <person> , you ' re a <unk> . <person> , i know you ' re think about it . think it ' ll be time , right ? i ' ll be the first one . if it did you , you ' ll be it .
Context.0 Context.1 HRED VHRED hredGAN_w hredGAN_u	come on , <person> , we could have a great week up in vermont . i see what you mean . but wednesday we ' ve got to be back in new york . <person> . <person> , i don ' t know what you ' re going to do . <person> , i know . i ' m kind of busy here . i ' m sorry , i can ' t .
Context.0 Context.1 HRED VHRED hredGAN_w hredGAN_u	isabel ? hi . that ' s what i ' m talking about . you ' re not going to be a <unk> , <person> . <person> ? what do you want to know ?
Context.0 Context.1 HRED VHRED hredGAN_w hredGAN_u	you understand mr . <person> is constantly barraged with nuisance litigation and irs audits . of course . that ' s right , <person> . <person> , you ' re going to be a <unk> of the <unk> . <person> , i ' m sure , but you ' re probably still in to the new <person> ' s house , huh ? you ' re gon na be a little boy , and i ' m not gon na get out !
Context.0 Context.1 HRED VHRED hredGAN_w hredGAN_u	<person> ' t go into my room , she ' s asleep . <person> , but the nanny , mrs . <person> , she wants to see her . i ' m sorry , <person> . <person> , mr . <person> . <person> , i know . but she ' s never seen me anyway . i ' m sorry i didn ' t know what to say .
Context.0 Context.1 HRED VHRED hredGAN_w hredGAN_u	what was that for ? i ' d forgotten how exciting it is working together . i ' m sorry , <person> . <person> , <person> . i don ' t know . i ' m sorry , <person> . so many people are so interested , but you ' re a very nice guy . i ' ve always been so sure to <person> .
Context.0 Context.1 HRED VHRED hredGAN_w hredGAN_u	a teacher ' s . oh my god . i ' m sorry , <person> . <person> , <person> ' t worry about it . i don ' t know what you ' re doing . i ' m sorry . i didn ' t mean to . i didn ' t mean to hurt you . i just wanted you <person> , i don ' t want to be your friend .
Context.0 Context.1 HRED VHRED hredGAN_w hredGAN_u	i took the initiative . you haven ' t been given an assignment in this case . <person> . <person> , <person> , <person> , <person> . i ' m not going to get out of the <unk> . i ' m not sure . <continued.utterance> i ' m not gon na do you this shit anyway , <person> . <person> , i guess i didn ' t .

Table 5: Additional sample responses of HRED, VHRED and hredGAN on Movie.

Model	Response
Context.0 Context.1 HRED VHRED hredGAN_w hredGAN_u	what ' s the command you are running to chroot and the error message you get in return? i mounted the device sudo mount /dev/sda1 /mnt - that works - then sudo chroot /mnt - error is no such file or command what ' s the problem ? is there a way to do that ? what do you mean ? I have no idea , I just did a sudo fdisk and it said it was mounted I ' m trying to get the latest version of ubuntu , but I ' m getting a lot of problems .
Context.0 Context.1 HRED VHRED hredGAN_w hredGAN_u	Hmm, have you checked that ICMP is being allowed all the way through and that the server responds to other ICMP requests? now i can ping it (i had to enable the lzo compression on client) i cannot reach the remote subnet adding a rule in the firewall forward the lan to vpn what ' s the problem ? I don ' t know what you want to do I have no idea , I ' m just trying to get the ip address from the router to the router" I ' m not sure if that ' s the case , but I ' m not sure if it ' s a good idea to have a look at the output .
Context.0 HRED VHRED hredGAN_w hredGAN_u	The netboot one is suppose to download packages from the net. I don ' t know what you want to do . how do I do that ? I ' m not sure of a command to do so I can just install the ubuntu . iso and burn a new iso you can ' t have a look at the man page . . .
Context.0 Context.1 HRED VHRED hredGAN_w hredGAN_u	The netboot one is suppose to download packages from the net. like the ones to be installed? or the installed to be run? I don ' t know what the problem is . is there a way to install the package manager ? you don ' t need to install the whole system , just install the ubuntu installer I ' m trying to install a .UNK version of ubuntu , but I ' m trying to install it on my laptop .
Context.0 HRED VHRED hredGAN_w hredGAN_u	You trying to mount it open with the password? I have to do that , but I have a problem with the same problem . is there a way to do ? yes , I ' m not sure if it ' s mounted or not I ' m trying to install the latest version of Ubuntu , but I ' m running on a .UNK .
Context.0 HRED VHRED hredGAN_w hredGAN_u	are you using lvm for your root partition? yes , but I have a problem with a new user , but I don ' t want to do that . what is the best way to do that ? what is that ? yes , I ' m on a live system , I ' m trying to use the partition to boot from windows I ' m not sure if it ' s a good idea to have a backup of the system . . .
Context.0 Context.1 HRED VHRED hredGAN_w hredGAN_u	are you using lvm for your root partition? i have the same config working on a 10.04 box . It appears in 12.04 lvm is on upstart and multipath is not . i do use lvm for my root folder nifty , is there any way to work around it? have you heard of anyone posting a bug about this to have it looked at? darn... well i hope they can fix this eventually i hope they keep inited and move everything to upstart . Upstart is not easy to setup for programs that don ' t come packaged for upstart. HRED I don ' t know what the problem is , but I don ' t know what the problem is . ? VHRED is there a way to do that ? or **unknown** ? hredGAN_w you can also try the live cd and see if you can get it back up the live" hredGAN_u I ' m trying to get my wireless working on my laptop , but I ' m getting to the point that I can be able to use the wireless card to get the wireless working .
Context.0 Context.1 HRED VHRED hredGAN_w hredGAN_u	are you using lvm for your root partition? i have the same config working on a 10.04 box . It appears in 12.04 lvm is on upstart and multipath is not . i do use lvm for my root folder nifty , is there any way to work around it? have you heard of anyone posting a bug about this to have it looked at? darn... well i hope they can fix this eventually i hope they keep inited around and move everything to upstart . Upstart is not easy to setup for programs that don ' t come packaged for upstart. Context.2 yes, because i cant supply it at boot time, i want the truecrypt drive to come up by itself without manual intervention HRED what is the problem ? VHRED what do you mean ? hredGAN_w you can also mount a new one and put the mount command to the mount point" hredGAN_u I ' m trying to get my sound working , but I ' m trying to get my sound working .

Table 6: Additional sample responses of HRED, VHRED and hredGAN on Ubuntu.