

Mask & Focus: Conversation Modelling by Learning Concepts

Gaurav Pandey, Dinesh Raghu, Sachindra Joshi

IBM Research, New Delhi, India

{gpandey1, diraghu1, jsachind}@in.ibm.com

Abstract

Sequence to sequence models attempt to capture the correlation between all the words in the input and output sequences. While this is quite useful for machine translation where the correlation among the words is indeed quite strong, it becomes problematic for conversation modelling where the correlation is often at a much abstract level. In contrast, humans tend to focus on the essential concepts discussed in the conversation context and generate responses accordingly. In this paper, we attempt to mimic this response generating mechanism by learning the essential concepts in the context and response in an unsupervised manner. The proposed model, referred to as Mask & Focus maps the input context to a sequence of concepts which are then used to generate the response concepts. Together, the context and the response concepts generate the final response. In order to learn context concepts from the training data automatically, we *mask* words in the input and observe the effect of masking on response generation. We train our model to learn those response concepts that have high mutual information with respect to the context concepts, thereby guiding the model to *focus* on the context concepts. Mask & Focus achieves significant improvement over the existing baselines in several established metrics for dialogues.

Introduction

Standard neural conversation modelling (Vinyals and Le 2015; Serban et al. 2016) follows a sequence-to-sequence framework for learning the relationship between sequence of words in the conversation-context and the next response. Attempts have been made to incorporate abstract information in the form of latent variables to introduce variability in response generation (Zhao, Zhao, and Eskenazi 2017; Serban et al. 2017b). In these architectures, all words in the conversation-context and the next response are treated to be equally important. Such an assumption works well for machine translation as most words/phrases in the source language have to be encoded and translated to the target language. However, for modelling conversations, the encoder is expected to understand the conversation context with a higher level of abstraction to generate the next response.

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

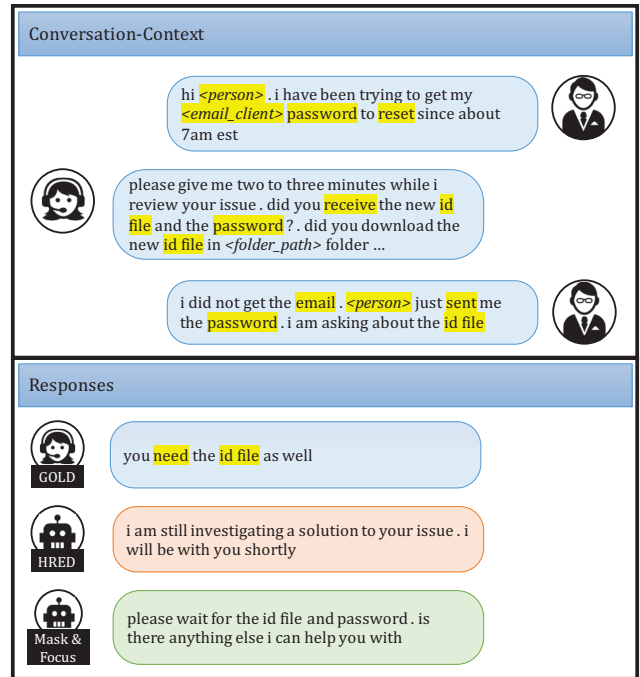


Figure 1: An example conversation from a tech support dataset along with responses generated responses by HRED and Mask & Focus. The highlight indicates concepts identified by Mask & Focus.

In contrast to these models, humans focus primarily on those parts of the context that are relevant to generate a response. Moreover, we form a rough sketch of the important words/phrases to be mentioned in the response, before letting the language model take over. This process is illustrated using an example conversation in Figure 1. The spans in the conversation context necessary to generate the response are highlighted. The spans that convey the crux of the response are also highlighted in the gold response. We refer to these important spans that captures the semantics of the conversations as concepts.

Past work has shown (Serban et al. 2017a) that repre-

senting contexts and responses using concepts improves the performance of conversation modeling but they require a tedious process of manual identification of these concepts for each domain. In this paper, we propose Mask & Focus, an approach which mimics human response generation process by discovering concepts in an unsupervised fashion. The model maps the context to the concept space. The concepts and the context are then used to predict the response concepts. Finally, the response concepts, along with the context and context concepts are used to generate the response. Learning to converse through concept has three main advantages: (1) It disentangles the learning of language and semantics in conversation modelling. The context in natural language is mapped to concept or semantic space. The model then translates the context concepts to response concepts. Finally the response concepts are mapped back to natural language. (2) Using words/phrases to represent concepts in context and response makes them interpretable (3) When the amount of available data is limited, learning in the concept space can be much more data-efficient. Figure 1 shows an example response generated by Mask & Focus and HRED on a tech-support dataset. The example illustrates the ability of Mask & Focus to leverage the identified concepts to generate a topically coherent response, while HRED resorts to a generic response.

In order to discover concepts in an unsupervised fashion, we systematically probe a trained conversation model by *masking* words in the context and computing the pointwise mutual information (PMI) of the masked words with respect to the response. Then, the decoder is trained to predict the response concepts by optimizing a variational lower bound (Jordan et al. 1999) to the marginal log-likelihood of the response. To ensure that the model focuses on the context concepts, we choose the variational distribution in the lower bound as a function of the mutual information between the concepts in the context and response. We show that Mask & Focus outperforms existing approaches on two datasets on automatic evaluations and human evaluation study.

To summarize, we make the following contributions:

1. *Mask*: We propose a novel technique for learning essential concepts in the conversation context by using PMI-guided systematic probing.
2. *Focus*: We propose a novel response generation technique that focuses on the context concepts by optimizing a variational lower bound to the marginal log-likelihood of the response.
3. We show that Mask & Focus achieves significant improvement in performance over baselines in several established metrics on Ubuntu dialog corpus and an internal technical support dataset.

Approach

A conversation is a sequence of utterances (u_1, \dots, u_m) between two or more individuals. Moreover, an utterance u_i is a sequence of words (w_{i1}, \dots, w_{il}) . At any given time-step $t < m$, the utterances till time t are referred to as the context \mathbf{c} , whereas the $(t + 1)^{th}$ utterance is referred to as the response \mathbf{r} . The set of all context response pairs in the training

data is denoted by \mathcal{D} . The set of all words used for training a model is referred to as the vocabulary and is denoted by \mathcal{V} .

Generative Conversation Modelling

Generative conversation models maximize the probability of the groundtruth response given the input context. That is, for a given context response pair (\mathbf{c}, \mathbf{r}) from the training data, they maximize the probability of the response \mathbf{r} given the context \mathbf{c} . Towards that end, the context is fed to an encoder to get a context embedding. The encoder can either be a flat RNN or a hierarchical RNN for processing words and utterances at different levels. The context embedding is fed to a decoder and the probability of the groundtruth response is maximized. In other words, generative conversation models maximize the following quantity:

$$\mathcal{L}(\theta) = \sum_{(\mathbf{c}, \mathbf{r}) \in \mathcal{D}} \log p(\mathbf{r}|\mathbf{c}; \theta), \quad (1)$$

where θ is the set of parameters that parametrize the above distribution. For a trained generative conversation model \mathcal{G} , the corresponding response distribution will be denoted as $p_{\mathcal{G}}(\mathbf{r}|\mathbf{c})$.

Mask & Focus

Given the context \mathbf{c} , we predict the context concepts \mathbf{z}_c . The context and the context concepts are used together to predict the response concepts \mathbf{z}_r . Finally, the context \mathbf{c} , and the concepts \mathbf{z}_c and \mathbf{z}_r are used for predicting the response. An overview of the model is shown in Figure 2.

Mask-Context Concepts: While generative conversation models capture the correlation between the words in the context and response, they tend to generate generic and non-informative responses for most contexts. This is due to the fact that these models fail to identify the words in the context that need to be focused while generating a response. In contrast, humans selectively focus on the essential concepts discussed in the context and frame a response accordingly. We wish to mimic this mechanism by identifying the essential concepts in the context. We restrict our context concepts to be subsets of words present in the context.

Given a context-response pair (\mathbf{c}, \mathbf{r}) , the context concepts \mathbf{z}_c are those words in the context that are ‘vital’ for generating the response \mathbf{r} . The concepts are unknown and hence, they need to be learnt from the data. We learn these concepts by maximizing the pointwise mutual information between these concepts and the response.

For any given word w in the vocabulary, the context obtained after masking the word is referred to as the masked context: $\mathbf{c}^{-w} = \mathbf{c} - w$. Conditioned on the masked context, the PMI between the word w and the response \mathbf{r} for a pretrained conversation model $p_{\mathcal{G}}$ is given by

$$\begin{aligned} \text{PMI}(\mathbf{r}; w) &= \log \frac{p_{\mathcal{G}}(\mathbf{r}, w|\mathbf{c}^{-w})}{p_{\mathcal{G}}(\mathbf{r}|\mathbf{c}^{-w})p_{\mathcal{G}}(w|\mathbf{c}^{-w})} \\ &= \log p_{\mathcal{G}}(\mathbf{r}, w|\mathbf{c}^{-w}) - \log p_{\mathcal{G}}(\mathbf{r}|\mathbf{c}^{-w}) \\ &= \log p_{\mathcal{G}}(\mathbf{r}|\mathbf{c}, w) - \log p_{\mathcal{G}}(\mathbf{r}|\mathbf{c}^{-w}) \end{aligned} \quad (2)$$

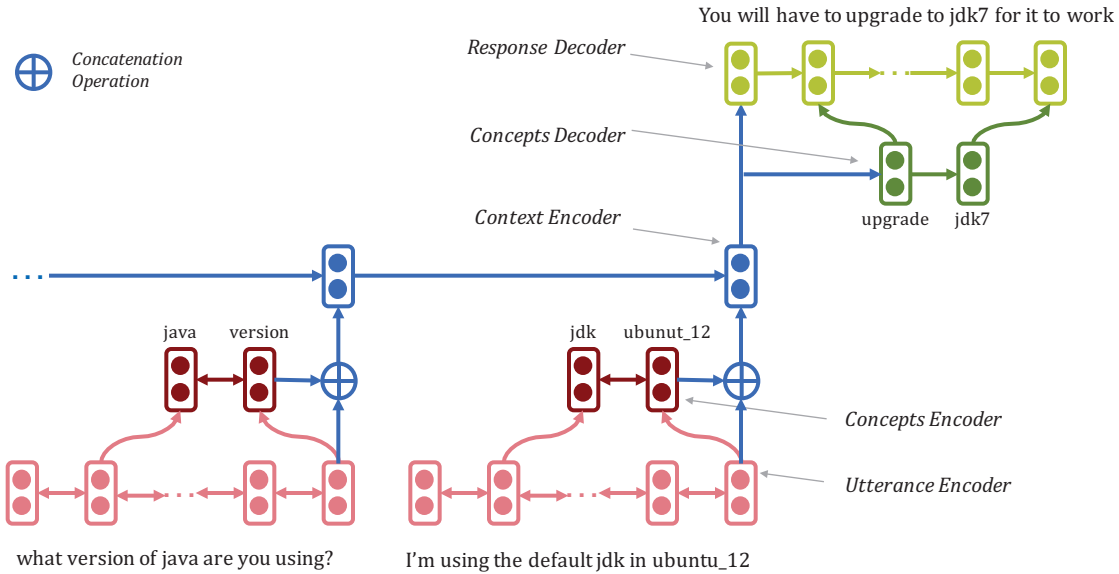


Figure 2: An overview of the proposed model. The bidirectional arrows indicate a bidirectional RNN while the unidirectional arrows indicate a unidirectional RNN. The curved arrows are used to indicate the words that are copied (either from utterances to concepts or vice-versa.)

Here, the last line follows since $\{w\} \cup \mathbf{c}^{-w} = \mathbf{c}$. For the generative conversation model $p_G(\mathbf{r}|w, \mathbf{c}) = p_G(\mathbf{r}|\mathbf{c})$. Hence, the pointwise mutual information between a word and the response is the difference between the log-probability of the response when the word w is present and the log-probability of the response when w is absent from the context. Intuitively, this quantity captures the importance of word w for predicting the response. This definition can easily be extended from a word to any set of words.

Hence, we propose the following strategy for learning the context concepts: A context-response pair is sampled from the data. The context is probed systematically to obtain a masked context. The probe randomly masks the words in the context. The context as well as the masked context are passed through a generative conversation model to evaluate the log-probability of the response before and after masking. The difference in log-probabilities are used for updating the scores of the masked words. The words with scores above a certain threshold are pushed to the concept bank. The context concepts \mathbf{z}_c are all those words in the context that are also available in the concept bank.

Having obtained the context concepts \mathbf{z}_c , we encode them as follow: Each utterance in the context is fed to an RNN encoder to obtain the utterance embedding. The corresponding concepts are also fed as input to an RNN encoder and the resultant embedding is concatenated with the utterance embedding. The embeddings of all utterances and corresponding concepts are fed to a second-level RNN to generate the context vector. This context vector will be denoted as \mathbf{v} . This corresponds to the lower-left portion of Figure 2.

Focus-Response Concepts: Since neural conversation models are black boxes, the concepts learnt by probing a

conversation model can be useful for obtaining a better understanding of the trained model. However, our aim here is to utilize the concepts learnt so far for generating better responses. Specifically, we want the model to generate responses that take into account the context concepts.

An obvious strategy is to provide the context concepts \mathbf{z}_c learnt in the previous section along with the corresponding utterances as input to the encoder as discussed in the end of the previous section. Unfortunately, a model can still choose to ignore \mathbf{z}_c while generating a response. Hence, instead of predicting the response directly from the context and context concepts, we break the process of response generation by predicting the response concepts \mathbf{z}_r at an intermediate step. By choosing \mathbf{z}_r to have high mutual information with \mathbf{z}_c , the model will be forced to focus on \mathbf{z}_c and hence generate more informative responses.

Given the context and its concepts, the joint distribution of the response and its concepts can be written as

$$p(\mathbf{r}, \mathbf{z}_r | \mathbf{c}, \mathbf{z}_c) = p(\mathbf{z}_r | \mathbf{c}, \mathbf{z}_c) p(\mathbf{r} | \mathbf{z}_r, \mathbf{c}, \mathbf{z}_c) \quad (3)$$

Since the response concepts are unknown, it is not possible to directly optimize the above probability. Instead, we propose to optimize a lower bound to the marginal likelihood of the response:

$$\mathbb{E}_{q(\mathbf{z}_r)} \log \frac{p(\mathbf{r}, \mathbf{z}_r | \mathbf{c}, \mathbf{z}_c)}{q(\mathbf{z}_r)} \leq \log \sum_{\mathbf{z}_r} p(\mathbf{r}, \mathbf{z}_r | \mathbf{c}, \mathbf{z}_c) \quad (4)$$

The above inequality is a straightforward consequence of Jensen's inequality. The term to the left of the inequality is referred to as the evidence lower bound objective (ELBO). Interestingly, the inequality holds true irrespective of the choice of the distribution q , also known as the variational distribution.

Standard variational training proceeds by optimizing ELBO with respect to the variational distribution q and the model distribution p alternatively. However, we want to capture the property that the response concepts have high mutual information with respect to the context concepts. Thus, to ensure this property, we optimize the above objective with respect to the model distribution only. The distribution q is chosen to give high probability to those words in the response that have high mutual information with respect to the context concepts.

Specifically, let w_ℓ be the ℓ^{th} word in the response. The pointwise mutual information (PMI) of w_ℓ with the context concepts \mathbf{z}_c conditioned on the masked context and the previous words for the model p is given by

$$\begin{aligned} \text{PMI}(w_\ell, \mathbf{z}_c) &= \log \frac{p(w_\ell, \mathbf{z}_c | \mathbf{c}^{-z})}{p(w_\ell | \mathbf{c}^{-z})p(\mathbf{z}_c | \mathbf{c}^{-z})} \\ &= \log p(w_\ell | \mathbf{c}^{-z}, \mathbf{z}_c) - \log p(w_\ell | \mathbf{c}^{-z}) \\ &= \log p(w_\ell | \mathbf{c}, \mathbf{z}_c) - \log p(w_\ell | \mathbf{c}^{-z}) \end{aligned} \quad (5)$$

Here, the last line follows because $\mathbf{c} = \mathbf{c}^{-z} \cup \mathbf{z}_c$. Note that unlike the pretrained conversation model p_G , the proposed model p takes the context concepts explicitly as input as shown in Figure 2. For the sake of brevity, we have not explicitly stated the dependence of w_ℓ on $w_{1:\ell-1}$ explicitly.

To capture this PMI in the ELBO in (4), we define the variational distribution q as follows:

$$q(w_\ell) = \begin{cases} \frac{\text{PMI}(w_\ell; \mathbf{z}_c)}{1 + \text{PMI}(w_\ell; \mathbf{z}_c)}, & \text{if } \text{PMI}(w_\ell) > 0 \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

Conditioned on the response, we assume that the response concepts are selected independently. Hence, the distribution of response concepts $q(\mathbf{z}_r)$ factorizes over the words in \mathbf{z}_r .

Next, we maximize the ELBO with respect to the model distribution p . Since q is fixed for this step, the ELBO objective can be simplified as given below:

$$\mathbb{E}_{q(\mathbf{z}_r)} \log p(\mathbf{z}_r | \mathbf{c}, \mathbf{z}_c) + \mathbb{E}_{q(\mathbf{z}_r)} \log p(\mathbf{r} | \mathbf{z}_r, \mathbf{c}, \mathbf{z}_c) \quad (7)$$

The two terms in the above equation correspond to the training objective for the concept and response decoder respectively. We assume that the distribution of the response concepts is modelled autoregressively, that is:

$$\log p(\mathbf{z}_r | \mathbf{c}, \mathbf{z}_c) = \sum_{w_\ell \in \mathbf{z}_r} \log p(w_\ell | w_{1:\ell-1}, \mathbf{c}, \mathbf{z}_c) \quad (8)$$

Hence, the first term in (7) can be approximated by expanding it over the words in the response as follows:

$$\mathbb{E}_{q(\mathbf{z}_r)} \log p(\mathbf{z}_r | \mathbf{c}, \mathbf{z}_c) \approx \sum_{w_\ell \in \mathbf{r}} q(w_\ell) \log p(w_\ell | w_{1:\ell-1}, \mathbf{c}, \mathbf{z}_c)$$

The distribution $p(w_\ell | w_{1:\ell-1}, \mathbf{c}, \mathbf{z}_c)$ is implemented using a decoder RNN. The variational distribution q weighs the words of the response for training the concept decoder.

The second term in (7) corresponds to the training objective for the response decoder. To optimize this term, we generate samples from the variational distribution $\mathbf{z}_r \sim q(\mathbf{z}_r)$. A decoder with copy mechanism (Gu et al. 2016) is trained

to predict the words of the response by copying the response concepts.

Note that changing the model distribution p for training the concept and response decoder causes the variational distribution q to change. Hence, the computation of q needs to alternate with training of concept and response decoders.

Experimental Setup

To demonstrate the utility of Mask & Focus, we evaluate it on two datasets.

Ubuntu Dataset: The proposed model is evaluated for generating responses on the Ubuntu Dialogue Corpus (Lowe et al. 2015). This dataset has been extracted from Ubuntu chatlogs, where users seek support for technical problems related to Ubuntu. The conversations in this dataset deal with a wide variety of technical problems related to Ubuntu OS and hence the total number of concepts discussed in these conversations is indeed quite large. This makes this dataset an ideal candidate for evaluating the proposed dialogue model. Table 1 depicts some statistics for this dataset:

	Ubuntu	Tech Support
Training Pairs	499,873	20,000
Validation Pairs	19,560	10,000
Test Pairs	18,920	10,000
Average No. of Turns	8	11.9
Vocabulary Size	538,328	293,494

Table 1: Dataset statistics for Ubuntu Dialogue Corpus v2.0 (Lowe et al. 2015) and Tech Support dataset

Tech Support Dataset: We also conduct our experiments on an internal technical support dataset with $\sim 20\text{K}$ conversations. We will refer to this dataset as Tech Support dataset in the rest of the paper. Tech Support dataset contains conversations pertaining to an employee seeking assistance from a technical support agent to resolve problems such as password reset, software installation/licensing, and wireless access. In contrast to Ubuntu dataset, this dataset has clearly two distinct users — employee and agent. In our experiments we model the *agent* responses only. For evaluating the efficacy of Mask & Focus for low resource settings, we create subsets of this dataset with 5,000, 10,000 and 20,000 conversations respectively.

Baselines

We compare the performance of the proposed model against several models that have been proposed for modelling conversations in the past few years: (i) HRED: An extension of sequence to sequence models for dealing with sequence of sequences (Serban et al. 2016) (ii) VHRED: An extension of HRED to incorporate continuous-valued latent variables (Serban et al. 2017b). This model is trained using variational inference with reparametrization (Diederik, Welling, and others 2014). (iii) HRED-attend: An extension of HRED that attend on the words of the context.

Model	Entity			Activity			Tense	Cmd
	F1	P	R	F1	P	R	Acc.	Acc.
VHRED	2.53	3.28	2.41	4.63	6.43	4.31	20.2	92.02
HRED	2.22	2.81	2.16	4.34	5.93	4.05	22.2	92.58
HRED-attend	2.44	3.13	2.08	4.86	6.04	4.65	29.75	99.89
MrRNN-Noun	6.31	8.68	5.55	4.04	5.81	3.56	24.03	90.66
Mask & Focus	7.82	8.78	8.69	5.43	6.42	5.90	24.47	93.19
MrRNN-ActEnt	3.72	4.91	3.36	11.43	16.84	9.72	29.01	95.04

Table 2: Activity & Entity metrics for the Ubuntu corpus. HRED, VHRED, MrRNN-Noun and MrRNN-ActEnt as reported by Serban et al. (2017a). Note that MrRNN-ActEnt is supervised since the activities and entities used in the evaluation metrics were also used during training.

(iii) MrRNN-Noun: A multi-resolution conversation model where the nouns from the context and response are identified and modelled in conjunction with the actual context and response (Serban et al. 2017a). (iv) MrRNN-ActEnt: A multi-resolution conversation model where the activities and entities are predefined for the conversations. MrRNN-ActEnt models the relationship between these predefined activities and entities via an RNN in conjunction with the actual context and response.

Implementation Details

We implemented Mask & Focus using the Pytorch library (Paszke et al. 2017). The word embeddings as well as weights of various neural networks are initialized using a standard normal distribution with mean 0 and variance 0.01. The biases are initialized with 0. We use 500-dimensional word embeddings for all our experiments.

For the Ubuntu dataset, the utterance and utterance concept encoders are single-layer bidirectional encoders, where each direction has a hidden size of 1,000. The context encoder is a single-layer unidirectional encoder with a hidden size of 2,000. We use a single-layer decoder with a hidden size of 2000 for decoding the response concepts. The response decoder with copy attention also has a hidden size of 2,000. We use a fixed vocabulary size of 20,000. The word embeddings are shared by the utterance and utterance concept encoders. Similarly, the response concept decoder as well as the response decoder also share the word embeddings. A similar model is also used for Tech Support dataset, with the exception that the dimensions of hidden layers are halved for both encoders and decoders.

We used the Adam optimizer with a learning rate of 1.5×10^{-4} . A batch size of 10 conversations is used for training. Note that all the responses in the conversation are trained simultaneously by treating all utterances preceding the response as the context. As discussed before, an HRED is pretrained and is used for guiding the training of context concepts. The dimensions of utterance and context encoder as well as the response decoder are exactly the same as their corresponding dimensions in the proposed model. To prevent the model from overfitting, we use early stopping with log-likelihood on validation set as evaluation criteria.

Results

Quantitative Evaluation: The metrics that are often used for comparing machine translation models have been shown to be uncorrelated with human judgements for conversation modelling (Liu et al. 2016). This is understandable since a single context can be associated with multiple correct responses that have little word-overlap with each other.

Hence, a new set of metrics were proposed in (Serban et al. 2017a), which compare responses based on a list of predefined entities and activities. The coarse-level representation of the predicted response is obtained in terms of activities and entities and compared with the corresponding representation of the ground truth response. These metrics were designed specifically for the Ubuntu Dialogue Corpus. A brief description of each metric is given below:

- **Activity:** Activity metric compares the activities present in a predicted response with the ground truth response. Activity can be thought of as a verb.
- **Entity:** This compares the technical entities that overlap with the ground truth response. A total of 3115 technical entities were identified using public resources such as Debian package manager APT.
- **Tense:** This measure compares the tense of ground truth with predicted response.
- **Cmd:** This metric computes accuracy by comparing commands identified in ground truth utterance with a predicted response.

We use the code provided by (Serban et al. 2017a) for computing precision, recall and F1-score for entities and activities. Table 2 compares the proposed model against several baselines based on each of the above metrics. MrRNN-ActEnt learns the activities and entities in a supervised manner, that is, the information about activities and entities is available during training. In contrast, all the other models used for comparison are trained in an unsupervised manner. As can be observed from the table, the proposed model significantly outperforms the unsupervised baselines on each metric. This suggests that Mask & Focus can predict the correct activities and entities with higher accuracy and hence, can aid in resolving the technical support problems of Ubuntu dataset.

trash, glxgears, DNS, fonts, ndiswrapper, clamscan, vmware, torrent, inittab, gmake, libmp3lame, sound, Beryl, cube, firewall, bogomips, ipv6, dragon, updates, smbfs, envy, raid, antivirus, suspend, madwifi, fakeraid, runlevel, workspaces, tar, mono, Define, burning, netstat, proxy, automatix, port, NickServ, RAID, font, chkconfig, truecrypt, fsck, apache, realplayer, beryl, burn, initrd, evolution, dns, xampp, SD, skydome, brightness, photoshop, muine, patch, mbox, locale, logs, Firewall, mic, 64-bit, flash64, game, dd, spanish, dock, lsb_release, webmin, xvid, transparency, JDK, /boot, filesharing, motd, dist-upgrade, netflix, vi, encrypt, /etc/environment, wubi, pcpkr, IA64, OpenOffice, serpentine, i686, mkinitrd, upgrade, theme, uname, Apache, ltspp,

Table 3: The top scoring concepts in Ubuntu that weren't manually annotated by (Serban et al. 2017a) for training MrRNN-ActEnt.

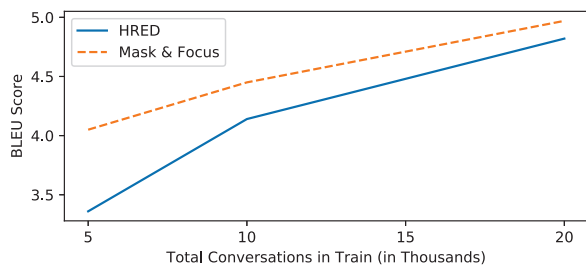


Figure 3: Comparison of Mask & Focus against HRED for Technical Support dataset.

Even more surprisingly, we were able to significantly outperform the supervised baseline MrRNN-ActEnt for the entity metrics. This indicates that the focus-based approach for learning response concepts is better than directly identifying the entities and training a supervised model to predict them. Moreover, we were able to discover several interesting concepts that were not identified manually in MrRNN-ActEnt. Some of these concepts are shown in Table 3.

Low Resource Setting: Next, we evaluate the performance of Mask & Focus on the Technical Support Dataset. As the Technical Support dataset is from a narrower domain than the Ubuntu corpus, we use this dataset to understand if Mask & Focus is able to capture these concepts even when the amount of resources available is small. In order to show this, we create subsets with randomly sampled 5,000, 10,000 and 20,000 context response pairs from the dataset and train HRED and Mask & Focus on these subsets. Unlike Ubuntu, we do not have a predefined list of activities and entities for this dataset. Hence, we use BLEU score (Papineni 2002) to compare the performance. The results are shown in Figure 3. As can be observed from the Figure, Mask & Focus outperforms HRED considerably when the dataset size is small. However, as the dataset grows in terms of training samples the gap reduces.

Human Evaluation Liu et al. (2016) have shown that metric based evaluations of response prediction in conversations are not strongly correlated with human judgements. So, we conduct a human study to compare the quality of the responses generated by Mask & Focus and other existing approaches. Annotators were shown a conversation-context along with the responses generated by unsupervised models (i.e) Mask & Focus, HRED, VHRED and MrRNN-Noun. The name of the algorithms were anonymized and we randomly shuffled the order as well. All annotators were asked to rate the relevance of the response on a scale of 0-4. We noticed that existing approaches often generated generic responses such as "what are you talking about" or "i see". So, we also requested the annotators to mark such generic responses. We removed these generic responses before computing the average relevance score reported in Table 4. We randomly sampled 100 context response pairs from the dataset, and collected a total of 500 relevance annotations.

Model	Relevance	GR Ratio
Gold	2.66	0.05
HRED	1.25	0.53
VHRED	1.48	0.39
MrRNN-Noun	1.56	0.32
Mask & Focus	1.79	0.18

Table 4: Human judgements on model responses on 100 random samples from Ubuntu corpus. Relevance scores are on a scale of (0-4). GR stands for Generic Response

Table 4 shows the relevance scores of gold and various algorithms. GR Ratio indicate the ratio of responses that were generic. Mask & Focus's ability to produce topically coherent responses is reflected by the significantly low generic response ratio. The relevance scored of Mask & Focus is significantly better than the existing approaches. A low relevance score on the gold responses captures the inherent difficulty in modelling the Ubuntu Corpus.

Discussion

We inspect the response concepts associated with the concepts in the context in Table 5. To obtain this table, we evaluate Mask & Focus on all conversations containing the specific context concept. The response concepts are predicted by the concept decoder. We list the most frequent response concepts in Table 5. As can be observed from the Table, the response concepts learnt by the model are indeed quite meaningful. Moreover, masked probing of generative conversation models, discovers concepts that weren't even extracted by manual inspection of Ubuntu chatlogs or by extracting terms from the package managers (the approach used in (Serban et al. 2017a)).

Responses generated by various algorithms are shown in Table 6. The contexts are highlighted with the discovered context concepts and the Mask & Focus responses are highlighted with the copied response concepts. We see that Mask

ipod	unrar	webcam	skype	eclipse	pdf	java	raid
ipod gtkpod	unrar rar	webcam cheese	skype mic	eclipse ide	pdf evince	java sun	raid software
itunes	unrar-nonfree	cam	deb	jdk	adobe	jre	mdadm
ipods	multiverse	webcams	voice	netbeans	print	sun-java6-plugin	array
rockbox	unrar-free	camera	medibuntu	sdk	xpdf	eclipse	controller
/media/ipod	.rar	camorama	call	sun	reader	sun-java6-jre	hardware
amarok	file-roller	lsusb	ekiga	pycharm	document	plugin	raid1
banshee	sources.list	logitech	libqt3c102-mt	jre	pdfs	openjdk	fakeraid
rhythmbox	non-free	motion	sip	java6	acrobat	update-alternatives	dmraid
plug	nonfree	amsn	uvcvideo	javac	openoffice	gcj	fake

Table 5: Context concepts (top row) and their corresponding response concepts.

Conversation-Context	Responses
can someone tell me how to install the downloads i got off the internet ? → ... → if you need help otherwise you need to name the package type you have like is it a tar.gz etc → its a .rar and a .exe	mask & focus: you need to unrar the package manager to unrar rar archives , then it doesnt work with the archive manager hred: What version of Ubuntu are you using ? vhred: I think you need to install the . tar . gz . mrrnn-noun: unk
my java has been faulty lately and i want to reinstall ... could someone please help me remove all my java ? → and tell us which of the javas you have → i also have some downloaded from oracle	mask & focus: java -version? what does java -version tell you? hred: what is the error? vhred: and what is the output of : dpkg -l grep unk mrrnn-noun: how do i uninstall oracle ?
you can use gparted to resize the win7 partition to make room for a ubuntu partition → thanks	mask & focus: you can resize ntfs partitions , but gparted wont let you resize partitions hred: what is the error? vhred: what version of ubuntu are you running ? mrrnn-noun: is it possible to resize a partition ?

Table 6: Model Responses from Ubuntu Dialog Corpus. The context concepts and the response concepts are in bold. Arrow (→) indicates end of turn.

& Focus is able to capture important concepts and generate a topically coherent response. For example, it was able to learn that if the context contains "download" and "tar", then the response should have "unrar".

Related Work

The idea of generating an abstract representation of the response before generating the response has been explored by Serban et al. (2017b) and Zhao et al. (2017). Pandey et al. (2018) proposed to ground the generated response in a set of retrieved responses. Liu et al. (2018) incorporated external knowledge to guide response generation in a dialog. Serben et al. (2017a) proposed Multi resolution RNN (Mr-RNN) to augment sequence-to-sequence architecture to support multiple higher levels of abstractions of input and output sequences. However, unlike Mask & Focus, the model requires the higher level abstractions to be pre-defined and available during training.

Systematic probing of trained models using manually defined masks has been explored by Hendricks et al. (2018) and Caglayan et al. (2019) to overcome bias in captioning and multi-modal machine translation. Unlike these ap-

proaches, Mask & Focus discovers masks required for probing in an unsupervised manner. Pointwise mutual information (PMI) has proven useful in conversation modelling for addressing the problem of generic responses (Li et al. 2016). While Li et al. used mutual information for reranking the candidate responses only, we incorporate it in our training mechanism for discovering context and response concepts.

Conclusion

In this paper, we presented Mask & Focus, a conversation model that learns the concepts present in the context and utilizes them for predicting the response concepts and the final response. The concepts are learnt by a strategy that involves masking the concepts in the context and focusing on the masked concepts. We show that the proposed Mask & Focus model discovers interesting concepts from challenging datasets that fail to be discovered by manual annotation. Mask & Focus achieves significant improvement in performance over existing baselines for conversation modelling with respect to several metrics. Moreover, we also observe that Mask & Focus generates interesting responses specific to the input context.

References

- Caglayan, O.; Madhyastha, P. S.; Specia, L.; and Barrault, L. 2019. Probing the need for visual context in multimodal machine translation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4159–4170.
- Diederik, P. K.; Welling, M.; et al. 2014. Auto-encoding variational bayes. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Gu, J.; Lu, Z.; Li, H.; and Li, V. O. 2016. Incorporating copying mechanism in sequence-to-sequence learning. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1631–1640.
- Hendricks, L. A.; Burns, K.; Saenko, K.; Darrell, T.; and Rohrbach, A. 2018. Women also snowboard: Overcoming bias in captioning models. In *European Conference on Computer Vision*, 793–811. Springer.
- Jordan, M. I.; Ghahramani, Z.; Jaakkola, T. S.; and Saul, L. K. 1999. An introduction to variational methods for graphical models. *Machine learning* 37(2):183–233.
- Li, J.; Galley, M.; Brockett, C.; Gao, J.; and Dolan, B. 2016. A diversity-promoting objective function for neural conversation models. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 110–119.
- Liu, C.-W.; Lowe, R.; Serban, I.; Noseworthy, M.; Charlin, L.; and Pineau, J. 2016. How not to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 2122–2132.
- Liu, S.; Chen, H.; Ren, Z.; Feng, Y.; Liu, Q.; and Yin, D. 2018. Knowledge diffusion for neural dialogue generation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1489–1498.
- Lowe, R.; Pow, N.; Serban, I.; and Pineau, J. 2015. The ubuntu dialogue corpus: A large dataset for research in unstructured multi-turn dialogue systems. In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, 285–294.
- Pandey, G.; Contractor, D.; Kumar, V.; and Joshi, S. 2018. Exemplar encoder-decoder for neural conversation generation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1329–1338.
- Papineni, K. 2002. Bleu: a method for automatic evaluation of machine translation. In *40th Annual meeting of the Association for Computational Linguistics, 2002*, 311–318.
- Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; and Lerer, A. 2017. Automatic differentiation in pytorch.
- Serban, I. V.; Sordoni, A.; Bengio, Y.; Courville, A.; and Pineau, J. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *Thirtieth AAAI Conference on Artificial Intelligence*.
- Serban, I. V.; Klinger, T.; Tesauro, G.; Talamadupula, K.; Zhou, B.; Bengio, Y.; and Courville, A. 2017a. Multiresolution recurrent neural networks: An application to dialogue response generation. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Serban, I. V.; Sordoni, A.; Lowe, R.; Charlin, L.; Pineau, J.; Courville, A.; and Bengio, Y. 2017b. A hierarchical latent variable encoder-decoder model for generating dialogues. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Vinyals, O., and Le, Q. 2015. A neural conversational model. *arXiv preprint arXiv:1506.05869*.
- Zhao, T.; Zhao, R.; and Eskenazi, M. 2017. Learning discourse-level diversity for neural dialog models using conditional variational autoencoders. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 654–664.