Supplementary Material for the Paper:
"A Hierarchical Latent Structure for Variational Conversation Modeling"

Yookoon Park, Jaemin Cho & Gunhee Kim
Department of Computer Science and Engineering & Center for Superintelligence
Seoul National University, Korea
yookoonpark@vision.snu.ac.kr, {jaemin895, gunhee}@snu.ac.kr

what?
→ **you're a good man, you're a liar. you're a good man.**
→ **what's wrong with you**?
what's wrong?
→ it's a good idea.
→ you're a good man, you don't have to do anything like that.
isn't there a problem with you?
→ well, maybe ...
→ then, you'll be fine.
isn't that a little joke, isn't it?
→ yes.
→ you don't have to be a hero, do you?
isn't that a fact that you're a <unk>?
→ **no, i'm fine.**
→ **then why are you here?**

Table 1: An example of interpolated 3-turn responses over $\mathbf{z}^{\text{conv}}$ on Cornell Movie Dialog.

you're a liar?
→ i don't know.
→ i'll tell you what. i'll be there. i'll be back in a minute.
you don't know?
→ no, i mean, i don't know, i just don't think i'd be in love with you.
→ well?
you want a drink?
→ sure.
→ good. good night, mary.
you're sure?
→ we're going to have to do something, don't we?
→ well, i don't know what to say.
you know that?
→ sure.
→ you're a good man, are n't you?

Table 2: An example of 3-turn responses conditioned on sampled $\mathbf{z}^{\text{utt}}$ for a single fixed $\mathbf{z}^{\text{conv}}$.

1 Data Processing

In both datasets, we truncate utterances longer than 30 words. Tokenization and text preprocessing is carried out using Spacy [1].

As Cornell Movie Dialog does not provide a separate test set, we randomly choose 80% of the conversations in Cornell Movie Dialog as training set. The remaining 20% is evenly split into validation set and test set.

2 Implementation Details

We use Pytorch Framework [2] for our implementations. We plan to release our code public.

We build a dictionary with the vocabulary size of 20,000, and further remove words with frequency less than five. We set the word embedding dimension to 500. We adopt Gated Recurrent Unit (GRU) (Cho et al., 2014) in our model and all baseline models, as we observe no improvement of LSTMs (Hochreiter and Schmidhuber, 1997) over GRUs in our experiments. We use one-layer GRU with the hidden dimension of 1,000 (2,000 for bidirectional GRU) for our RNN decoders. Two-layer MLPs with hidden layer size 1000 parameterizes the distribution of latent variables. All latent variables have a dimension of 100. We apply dropout ratio of 0.2 during training. Batch size is 80 for Cornell Movie Dialog, and 40 for Ubuntu Dialog. For optimization, we use Adam (Kingma and Ba, 2014) with a learning rate of 0.0001 with gradient clipping. We adopt early stopping by monitoring the performance on the validation set. We apply the KL annealing to all variational models, where the KL multiplier λ gradually increases from 0 to 1 over 15,000 steps on Cornell Movie Dialog and over 250,000 steps on Ubuntu Dialog. For both the word drop and the utterance drop, we use drop probability of 0.25.

3 Experimental Results

Table 1 – 4 shows additional sample generation results.

4 Human Evaluation

We perform human evaluation study on Amazon Mechanical Turk (AMT). We first filter out con-

---

[1] https://spacy.io/
[2] http://pytorch.org/

Context	HRED	VHRED	VHRED+bow	VHCR
thanks, jack. i'll be in touch. → do you know what you're doing. → trust me.	i'm not going to be a good man.	you're not going to be a good time.	you ever seen a <unk> or something? you never seen a <unk>?	how much do you know about the <unk>?
i saw desert sun, i wanna tell you ... → no, i was, i was, i was just u learning u on, it's a ... → how'd you like working with richard hill?	oh, i don't know ...	what did you say?	i don't know ... i mean, i know ... i mean, you look like this ... it's ... you know, i've been	i've been thinking about it, but i ...
i don't think the people who live around here would like that very much ... → do you think they like the fact that <unk> has dropped ninety eight percent in the last ten years? → no, but ...	you're a good man, you know.	... and then it's not a <unk>. it's a <unk> of <unk>. it's a long time. it's a long time.	i know. it's a beautiful woman to do with your own. so, you're not married.	they're not going to kill me.
every one of these women thinks she should be the next angel of the month. it's your job, ted, to decide which of them is → let me get this straight ... → if it isn't straight by now, son, you probably should see a doctor.	no, no, no, no, no, no, no, no, no, no, no, no, no, no, no,	... i'm not a <unk>, i'm a <unk> ...	you look dead.	it's not the same thing. it's not the truth.
like following me around and <unk> the people i work with wasn't bad enough, but breaking into my apartment → it was open. → you got a lotta nerve.	i'm not a <unk>.	i was a <unk> <unk>.	you're not a cop. i didn't tell you. it was just a little <unk>, but it wasn't me. you're just trying to	why don't you just go to the police?

Table 3: Comparison of generated responses on Cornell Movie Dialog.

Context	HRED	VHRED	VHRED+bow	VHCR
hi, everyone. i download the <unk>, and want to install it with sudo dpkg i google earth stable current i386.deb, it tells me google earth → the error suggests running sudo apt get f install → i tried that, it comes the same error	then you should be able to get the .deb from the deb file	that's a good idea, sorry	what is it, please.	it says error dependency is not satisfiable <unk> that's the error i get is that
click the database you <unk> to import it into, and look for the import button. import as sql. → noticed. this will take me a while because i am not used to this subject and not fluent in it. → you may need to destroy the tables in the database before importing.	i am not sure what you mean. i am using the default ubuntu install.	it is a <unk> of the <unk>. i have a <unk> <unk> and i want to be able to use the database on the <unk>. i am	i dont think that s what you are saying is it isnt an option	then how to do that?
i'm trying to setup a shortcut dir of this path home user downloads in this path jail <unk> site how do i do this? i tried using → ln s home user downloads jail <unk> site downloads if you're trying to do something with a chroot jail guessing from the path you can't symlink outside → thank you for the answer. but making a vlan for each user and then limit the vlan is not possible when you only have ssh root access right	i'm not sure if that's possible, but i'm not sure if it's possible to do that.	<unk> is the user that owns the home directory	i'd say that's the case ... i've never even used a shell to draw the data with the destination folder	look at the permissions of the folder and the other is <unk>, yes
hi i am facing a problem! while upgrading my ubuntu to 12.04 lts unfortunately shutdown the system and after that when i reboot my system it shows black → did the shutdown occur in the middle of the upgrade? → yes, one of my friend just shut down, he was not aware that upgrade is going on now the system freezes at seemingly random points and leaves	i 'm not sure what you mean by 'shutdown'.	not sure what the problem is, try sudo apt get update sudo apt get upgrade sudo apt get upgrade sudo apt get upgrade	it has been interrupted.	so you are saying that it is not able to boot the cd or something?

Table 4: Comparison of generated responses on Ubuntu Dialog.

texts that contain generic unknown word (unk) token from the test set. Using these contexts, we generate model response samples. Samples that contain less than 4 tokens are removed. The order of the samples and the order of model responses are randomly shuffled.

Evaluation procedure is as follows: given a context and two model responses, a Turker decides which response is more appropriate in the given context. In the case where the Turker thinks that two responses are about equally good or bad or does not understand the context, we ask the Turker to choose "tie". We randomly select 100 samples to build a batch for a human intelligence test (HIT). For each pair of models, we perform 3 HITs on AMT and each HIT is evaluated by 5 unique humans. In total we obtain 9000 preferences in 90 HITs.

References

Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078* .

Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9(8):1735–1780.

Diederik P Kingma and Jimmy Ba. 2014. Adam: A

method for stochastic optimization. *arXiv preprint arXiv:1412.6980* .