
A Neural Conversational Model

Oriol Vinyals

Google

Quoc V. Le

Google

VINYALS@GOOGLE.COM

QVL@GOOGLE.COM

Abstract

Modeling conversations is an important research topic given its significance towards language understanding and artificial intelligence. Although solutions exist to modeling dialogue, these models are typically restricted to specific domains (e.g., booking an airline ticket), and require hand-crafted rules and extensive feature engineering. In this paper, we present a simple approach to this task by using the recently proposed sequence to sequence framework. This model learns to converse by learning to predict the next sentence given the previous sentence or sentences in a conversation. We find that this straightforward model can generate simple conversations given a large training dataset. Experiments in this paper suggest that, despite optimizing the wrong objective function, the model extracts knowledge from both a domain specific dataset, and from a large, noisy, and general domain dataset of movie subtitles.

1. Introduction

Recent progress in machine learning has led to remarkable progress in many domains such as speech recognition, computer vision, and language processing. Recent work suggests that recurrent neural networks can be used to map onto complicated sequences – such as language – and that these models can be conditioned on both fixed and variable length signals (Sutskever et al., 2014). One of the major advantages of this framework is that it requires little feature engineering and domain specificity whilst matching or surpassing state-of-the-art results. This, in our opinion, allows researchers to work on tasks for which domain knowledge may not be readily available, or for tasks which are simply

too hard to model.

Conversational models partially fall into this category, and have typically been designed to be very narrow in domain, with a major undertaking on feature engineering and debugging to make the models useful. In this work, we show that a very simple model can do surprisingly well on generating fluent and accurate replies to general domain questions trained on a corpus from movie conversations. Despite this, such a model may need a much more specialized and cleaned corpus to achieve what a rule based dialogue system would. To test this, we train the same model on chat sessions from an internal troubleshooting tool, and show that the model is, sometimes, able to track the problem and provide a solution to the user. We also experiment with conversations obtained from movie subtitles. In both cases, the model obtains better perplexity compared to the n-gram model and captures important long-range correlations. From a qualitative point of view, our model is sometimes able to produce natural conversations.

2. Related Work

Our approach is based on recent work which proposed to use neural networks to map sequences to sequences (Kalchbrenner & Blunsom, 2013; Sutskever et al., 2014; Bahdanau et al., 2014). This framework has been used for neural machine translation and achieves state of the art results on the English-French and English-German translation tasks from the WMT’14 dataset (Luong et al., 2014; Jean et al., 2014). It has also been used for other tasks such as parsing (Vinyals et al., 2014a) and image captioning (Vinyals et al., 2014b).

This framework employs a recurrent neural network (RNN) to both produce an output sequence, and to read an input sequence. Since it is well known that vanilla RNNs suffer from vanishing gradients, most researchers use variants of the Long Short Term Memory (LSTM) recurrent neural network (Hochreiter & Schmidhuber, 1997).

Our work is also inspired by the recent success of neural

language modeling (Bengio et al., 2003; Mikolov et al., 2010; Mikolov, 2012), which shows that recurrent neural networks are rather effective models in modeling natural language.

Building bots and conversational agents has been pursued by many researchers over the last decades, and it is out of the scope of this paper to provide an exhaustive list of references. However, most of these systems require a rather complicated processing pipeline of many stages (Lester et al., 2004; Will, 2007; Jurafsky & Martin, 2009). Our work differs from conventional systems by proposing an end-to-end approach to the problem which lacks domain knowledge. It could, in principle, be combined with other systems to re-score a short-list of candidate responses, but our work is based on producing answers given by a probabilistic model trained to maximize the probability of the answer given some context.

Our work is related to recent work by Sordoni et al. (Sordoni et al., 2015) and Shang et al. (Shang et al., 2015), both of which use recurrent neural networks to learn to converse.

3. Model

Our model is identical to the sequence to sequence (*seq2seq*) model described in (Sutskever et al., 2014). The model is based on a recurrent neural network, which reads the input sequence one token at a time, and predicts the output sequence, also one token at a time. During training, the model is provided with the true output sequence, so learning can be done by simple backpropagation. The model is trained to maximize the cross entropy of the correct sequence given its context. At test time, given that the true output sequence is not observed, we simply feed the predicted output token as input to predict the next output. This is a “greedy” inference approach. A less greedy approach is to use beam search, and feed several candidates at the previous step to the next step. The predicted sequence can be selected based on the probability of the sequence.

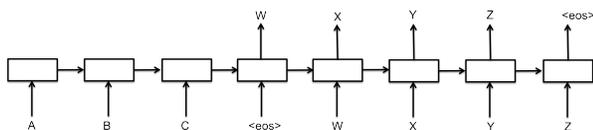


Figure 1. Using the *seq2seq* framework for modeling conversations.

Concretely, suppose that we observe a conversation with two turns: the first person utters “ABC”, and someone else replies “WXYZ”. We can use a recurrent neural network, and train to map “ABC” to “WXYZ” as shown in Figure 1 above. The hidden state of the model when it receives the

end of sequence symbol “<eos>” can be viewed as the *thought vector* because it stores the information of the sentence, or thought, “ABC”.

The strength of this model lies in its simplicity and generality. We can use this model for machine translation, question/answering, and conversations without major changes in the architecture.

Unlike easier tasks like translation, however, a model like sequence to sequence will not be able to successfully “solve” the problem of modeling dialogue due to several obvious simplifications: the objective function being optimized does not capture the actual objective achieved through human communication, which is typically longer term and based on exchange of information rather than next step prediction. The lack of a model to ensure consistency and general world knowledge is another obvious limitation of a purely generative model.

4. Datasets

In our experiments we used two datasets: a closed-domain troubleshooting dataset and an open-domain movie transcript dataset. The details of the two datasets are as follows.

4.1. Technical Troubleshooting dataset

Our first set of experiments used a dataset which is extracted from a troubleshooting chat service. In this service, costumers are dealing with computer related issues, and a specialist is helping them walking through a solution. Typical interactions (or threads) are 400 words long, and turn taking is clearly signaled. Our training set contains 30M tokens, and 3M tokens were used as validation. Some amount of clean up was performed, such as removing common names, numbers, and full URLs.

4.2. OpenSubtitles dataset

We also experimented our model with the OpenSubtitles dataset (Tiedemann, 2009). This dataset consists of movie conversations in XML format. It contains sentences uttered by characters in movies. We applied a simple processing step removing XML tags and obvious non-conversational text (e.g., hyperlinks) from the dataset. As turn taking is not clearly indicated, we treated consecutive sentences as if they were uttered by different characters. We trained our model to predict the next sentence given the previous one, and we did this for every sentence (noting that this doubles our dataset size, as each sentence is used both for context and as target). Our training and validation split has 62M sentences (923M tokens) as training examples, and the validation set has 26M sentences (395M tokens). The split is done in such a way that each sentence in a pair of sentences

either appear together in the training set or test set but not both. Unlike the previous dataset, the OpenSubtitles is quite large, and rather noisy because consecutive sentences may be uttered by the same character. Given the scope of movies, this is an open-domain conversation dataset, contrasting with the technical troubleshooting dataset.

5. Experiments

In this section, we describe the experimental results with the two datasets. We present a rather objective measure which is the perplexity of the model on the test set. We also show some samples of the interactions we had with the system that we trained.

5.1. Troubleshooting experiments

Given the lack of data for this domain specific and clean dataset, we trained a single layer LSTM with 1024 memory cells using stochastic gradient descent with gradient clipping. We built a vocabulary of the most common 20K words, and introduced special tokens indicating turn taking and actor.

At convergence, this model achieved a perplexity of 8, whereas an n-gram model achieved 18. Below is a sample of a simulated troubleshooting session, in which **NCM** is our Neural Conversational Model, and **Actor** the human actor interacting with it.

Conversation 1: VPN issues.

Describe your problem: i am having issues accessing vpn
NCM: hi
Actor: hello
NCM: could you please let me know what are the operating systems you are using ?
Actor: i am using linux
NCM: ok it was in the ssl vpn on which you stated it was broken
NCM: please check if you are logged in to corp network
NCM: what is the error that you are running please
Actor: i am seeing an error related to vpn
NCM: what is the error message that you are getting when connecting to vpn using network connect ?
Actor: connection refused or something like that
NCM: may i know the version of network connect you connect ?
Actor: i am not sure i know that
NCM: is the network connect prompting for any user certificate

NCM: ?
Actor: yes
NCM: are you at home or at work at the moment ?
NCM: office ?
Actor: i am at home
NCM: try this
NCM: goto <URL> page and enter your name and password
Actor: done
NCM: try logging in with <NAME> and check if its working fine now
Actor: yes , now it works !
NCM: great
NCM: anything else that i can help
Actor: no , that was very useful
NCM: you are other than <NAME> :)
NCM: have a nice night

(The <URL> indeed contains information about vpn access.)

Conversation 2: Browser issues.

Describe your problem: browser keeps crashing
NCM: hi
Actor: hi
NCM: could you please let me know what is the operating system you are using ?
Actor: linux
NCM: ok this is the issue with the system performance , nasser , tools .
Actor: what is that ?
NCM: is it a windows machine ?
NCM: we have linux on your machine .
NCM: is it windows xp or windows 7 ?
Actor: i am using linux
NCM: let me take a remote of your machine and check if issue has been pushed properly .
Actor: ok
NCM: please open terminal and run the below commands in one computer :
NCM: sudo apt - get update
NCM: sudo apt - get upgrade
NCM: sudo <COMMAND> - updater
Actor: done
NCM: let me try connecting now
Actor: ok
NCM: its done
NCM: should use the second or hit run on the background
Actor: done
NCM: try signing out now and follow the prompts as per the message given above
NCM: shd i reboot the machine .

Actor: *ok , please go ahead*
NCM: *once its completed please try to login .*
Actor: *sure*
NCM: *please ping me back once you have tried to login*
NCM: *are we still connected*
Actor: *i now logged in*
NCM: *thanks for your confirmation*
Actor: *what should i do now ?*
NCM: *wait for 1 5 minutes to complete the login policy*
NCM: *and then try logging in*
Actor: *ok*
NCM: *is there anything else to help you ?*
Actor: *let my try browsing a little bit and if it keeps crashing i will let you know*
NCM: *sure will do*
NCM: *seems the mail is not loading*
NCM: *are we still connected to the chat ?*
NCM: *<reinvite>*
Actor: *it seems good now , cool , thank you , bye*
NCM: *you , too*
NCM: *bye*

Conversation 3: Password issues.

Describe your problem: i forgot my password
NCM: *hi , this is <NAME> from helpdesk connect , what can i help you with today ?*
Actor: *hi , i forgot my password*
NCM: *ok great*
Actor: *can you tell me how i can recover it ?*
NCM: *i ' ll need to verify who it is at that machine*
NCM: *can we do a hangout*
NCM: *?*
Actor: *yes*
NCM: *<anonymized hangout instructions and verifications>*
Actor: *cool , i am good now*
NCM: *how ' s this looking good ?*
NCM: *i ' ll go ahead and close this chat .*
NCM: *take care . :)*
Actor: *bye*
NCM: *bye .*

5.2. OpenSubtitles experiments

We trained a two-layered LSTM using AdaGrad with gradient clipping.

Each layer of the LSTM has 4096 memory cells, and we built a vocabulary consisting of the most frequent 100K words. To speed up the softmax, we project the memory cells to 2048 linear units before feeding the information to

the classifier.

At convergence the perplexity of the recurrent model on the validation set is 17. Our smoothed 5-gram model achieves a perplexity of 28. Interestingly, adding the soft attention mechanism of (Bahdanau et al., 2014) did not significantly improve the perplexity on neither training or validation sets.

Besides the objective evaluation on perplexity, our simple recurrent model does often produce plausible answers. Here are a few sample question-answer pairs that allowed us to test the capabilities of the model when trained on the noisy but large dataset of OpenSubtitles from movies:

Conversation 1: Basic

Actor: *hello !*
NCM: *hello !*
Actor: *how are you ?*
NCM: *i ' m good .*
Actor: *what ' s your name ?*
NCM: *i ' m julia .*
Actor: *how old are you ?*
NCM: *i ' m 19 .*
Actor: *where are you from ?*
NCM: *i ' m out in the boonies.*
Actor: *are you a man or a woman ?*
NCM: *i ' m a woman .*
Actor: *why are we here ?*
NCM: *i ' m not sure .*
Actor: *okay, bye !*
NCM: *bye .*
Actor: *see you later .*
NCM: *bye .*

Conversation 2: Simple Q&A (contexts and multiple choice)

Actor: *my name is david . what is my name ?*
NCM: *david .*
Actor: *my name is john . what is my name ?*
NCM: *john .*
Actor: *are you a leader or a follower ?*
NCM: *i ' m a leader .*
Actor: *are you a follower or a leader ?*
NCM: *i ' m a leader .*

Conversation 3: General knowledge Q&A

Actor: *who is skywalker ?*
NCM: *he is a hero .*
Actor: *who is bill clinton ?*
NCM: *he ' s a billionaire .*
Actor: *is sky blue or black ?*

NCM: *blue .*
Actor: *does a cat have a tail ?*
NCM: *yes .*
Actor: *does a cat have a wing ?*
NCM: *no*
Actor: *can a cat fly ?*
NCM: *no .*
Actor: *what is the color of the sky ?*
NCM: *blue .*
Actor: *what is the color of water ?*
NCM: *water .*
Actor: *what is the color of blood ?*
NCM: *it is the same as a black eye .*
Actor: *what is the usual color of a leaf ?*
NCM: *it is a green one .*
Actor: *what is the color of a yellow car ?*
NCM: *yellow .*
Actor: *how much is two plus two ?*
NCM: *four .*
Actor: *how much is ten minus two ?*
NCM: *seventy - two .*

Conversation 4: Philosophical Q&A

Actor: *what is the purpose of life ?*
NCM: *to serve the greater good .*
Actor: *what is the purpose of living ?*
NCM: *to live forever .*
Actor: *what is the purpose of existence ?*
NCM: *to find out what happens when we get to the planet earth .*
Actor: *where are you now ?*
NCM: *i 'm in the middle of nowhere .*
Actor: *what is the purpose of dying ?*
NCM: *to have a life .*
Actor: *what is the purpose of being intelligent ?*
NCM: *to find out what it is .*
Actor: *what is the purpose of emotions ?*
NCM: *i don 't know .*

Conversation 5: Morality

Actor: *what is moral ?*
NCM: *what empowered humanity , what intellectual the essence is .*
Actor: *what is immoral ?*
NCM: *the fact that you have a child .*
Actor: *what is morality ?*
NCM: *what is altruism ?*
Actor: *what is the definition of altruism ?*
NCM: *if you don 't believe in god , then you don 't know .*
Actor: *ok ... so what is the definition of morality ?*

NCM: *well , the truth is , you 're not a believer in god almighty .*
Actor: *tell me the definition of morality , i am quite upset now !*
NCM: *i 'm not ashamed of being a philosopher !*
Actor: *tell me the definition of morality .*
NCM: *i don 't have ethics .*
Actor: *look , i need help , i need to know more about morality ...*
NCM: *i don 't know what ethics is .*
Actor: *seriously , what is morality ?*
NCM: *what is the definition of living ?*
Actor: *why living has anything to do with morality ?*
NCM: *you 're not a cop .*
Actor: *okay , i need to know how should i behave morally ...*
NCM: *i don 't know how to tell you .*
Actor: *is morality and ethics the same ?*
NCM: *yes , sir .*
Actor: *what are the things that i do to be immoral ?*
NCM: *i don 't know .*
Actor: *give me some examples of moral actions ...*
NCM: *i 'm not a moralist .*
Actor: *alright , morality ?*
NCM: *integrity .*
Actor: *what is integrity ?*
NCM: *i 'm sorry , i don 't know what else to say .*
Actor: *be moral !*
NCM: *be a man !*
Actor: *i really like our discussion on morality and ethics ...*
NCM: *and how i 'm not in the mood for a philosophical debate .*
Actor: *what do you like to talk about ?*
NCM: *nothing .*

Conversation 6: Opinions (notice the model outputs “he” , “she” or “it” depending on the subject being asked)

Actor: *what do you think about tesla ?*
NCM: *he 's a good conductor .*
Actor: *what do you think about bill gates ?*
NCM: *he 's a good man .*
Actor: *What do you think about messi ?*
NCM: *he 's a great player .*
Actor: *what do you think about cleopatra ?*
NCM: *oh , she 's very regal .*
Actor: *what do you think about england during the reign of elizabeth ?*

NCM: *it was a great place .*

We find it encouraging that the model, without additional memory to encode knowledge (other than the information captured by the parameters of the RNN), can remember facts and understand contexts without the complexity in traditional pipelines. An intriguing property of the model is that it can generalize to new questions. In other words, it does not simply try to match the question and copy the answers. In fact, most of the questions presented above, except for the first conversation, do not appear in the training set.

Nonetheless, one drawback of this basic model is that it only gives simple, short, sometimes unsatisfying answers to our questions as can be seen above. Perhaps a more problematic drawback is that the model does not capture a consistent personality. Indeed, if we ask not identical but semantically similar questions, the answers can sometimes be inconsistent. This is expected due to the simplicity of our model and the dataset in our experiments. The conversation below illustrates such failure:

Conversation 7: Job and Personality

Actor: *what is your job ?*

NCM: *i 'm a lawyer .*

Actor: *what do you do ?*

NCM: *i 'm a doctor .*

Lastly, as a comparison of other traditional bots that are rule-based, we provide the output of Cleverbot¹ for some of the questions in our test:²

Actor: *Are you a follower or a leader?*

Cleverbot: *No !!*

Actor: *My name is david . what is my name ?*

Cleverbot: *How old are you David?*

Actor: *My name is john . what is my name ?*

Cleverbot: *John*

Actor: *What is the purpose of being intelligent?*

Cleverbot: *The purpose of life is to love God and love people.*

Actor: *Is sky blue or black ?*

Cleverbot: *I am 200 years old.*

Actor: *What is the color of a yellow car ?*

Cleverbot: *Brown.*

Actor: *What is the usual color of a leaf ?*

Cleverbot: *What is the color of the apple in the fridge?*

Actor: *What is the purpose of life ?*

Cleverbot: *My purpose it to forward my species, in other words to make it easier for future generations of mankind to live.*

Actor: *What do you think about tesla?*

Cleverbot: *I am tesla.*

Actor: *What do you think about cleopatra ?*

Cleverbot: *Cool is cool.*

Actor: *What do you think about england during the reign of elizabeth ?*

Cleverbot: *I like him. Do you?*

Actor: *What do you think about bill gates?*

Cleverbot: *Lets talk about Link.*

Actor: *What is moral?*

Cleverbot: *Morality is the quality of being in accord with standards of right or good conduct.*

Actor: *What is immoral?*

Cleverbot: *What is the objective of sistetic life?*

Actor: *Does a cat have a wing?*

Cleverbot: *What is the name of your husband.*

Actor: *Can a cat fly?*

Cleverbot: *What is your husband's name?*

Actor: *Who is skywalker?*

Cleverbot: *There are two of them, Anakin Skywalker and Luke Skywalker, both characters from star Wars.*

Actor: *What is your job?*

Cleverbot: *I work in mental health.*

Actor: *What do you do?*

Cleverbot: *I have no job.*

An outstanding research problem is how to objectively measure the quality of models. We did so by manual inspection and by computing the perplexity. However, there are obvious drawbacks on using either. We are actively pursuing this and other research directions.

6. Discussion

In this paper, we show that a simple language model based on the *seq2seq* framework can be used to train a conversational engine. Our modest results show that it can generate simple and basic conversations, and extract knowledge from a noisy but open-domain dataset. Even though the model has obvious limitations, it is surprising to us that a purely data driven approach without any rules can deliver non-trivial answers. However, the model may require substantial modifications to be able to deliver realistic conversations. Amongst the many limitations, the lack of a coherent personality makes it difficult for our system to pass the Turing test (Turing, 1950).

¹<http://www.cleverbot.com/>

²The outputs of Cleverbot are non-deterministic so the answers may be different for different trials.

References

- Bahdanau, D., Cho, K., and Bengio, Y. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- Bengio, Y., Ducharme, R., Vincent, P., and Janvin, C. A neural probabilistic language model. *The Journal of Machine Learning Research*, 3:1137–1155, 2003.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural Computation*, 1997.
- Jean, S., Cho, K., Memisevic, R., and Bengio, Y. On using very large target vocabulary for neural machine translation. *CoRR*, abs/1412.2007, 2014.
- Jurafsky, D. and Martin, J. *Speech and language processing*. Pearson International, 2009.
- Kalchbrenner, N. and Blunsom, P. Recurrent continuous translation models. In *EMNLP*, 2013.
- Lester, J., Branting, K., and Mott, B. Conversational agents. In *Handbook of Internet Computing*. Chapman & Hall, 2004.
- Luong, T., Sutskever, I., Le, Q. V., Vinyals, O., and Zaremba, W. Addressing the rare word problem in neural machine translation. *arXiv preprint arXiv:1410.8206*, 2014.
- Mikolov, T. *Statistical Language Models based on Neural Networks*. PhD thesis, Brno University of Technology, 2012.
- Mikolov, T., Karafiát, M., Burget, L., Cernocký, J., and Khudanpur, S. Recurrent neural network based language model. In *INTERSPEECH*, pp. 1045–1048, 2010.
- Shang, L., Lu, Z., and Li, H. Neural responding machine for short-text conversation. In *Proceedings of ACL*, 2015.
- Sordani, A., Galley, M., Auli, M., Brockett, C., Ji, Y., Mitchell, M., Gao, J., Dolan, B., and Nie, J.-Y. A neural network approach to context-sensitive generation of conversational responses. In *Proceedings of NAACL*, 2015.
- Sutskever, I., Vinyals, O., and Le, Q. V. Sequence to sequence learning with neural networks. In *NIPS*, 2014.
- Tiedemann, J. News from OPUS - A collection of multilingual parallel corpora with tools and interfaces. In Nicolov, N., Bontcheva, K., Angelova, G., and Mitkov, R. (eds.), *Recent Advances in Natural Language Processing*, volume V, pp. 237–248. John Benjamins, Amsterdam/Philadelphia, Borovets, Bulgaria, 2009. ISBN 978 90 272 4825 1.
- Turing, A. M. Computing machinery and intelligence. *Mind*, pp. 433–460, 1950.
- Vinyals, O., Kaiser, L., Koo, T., Petrov, S., Sutskever, I., and Hinton, G. Grammar as a foreign language. *arXiv preprint arXiv:1412.7449*, 2014a.
- Vinyals, O., Toshev, A., Bengio, S., and Erhan, D. Show and tell: A neural image caption generator. *arXiv preprint arXiv:1411.4555*, 2014b.
- Will, T. *Creating a Dynamic Speech Dialogue*. VDM Verlag Dr, 2007.