

Neural Conversational Models

Human: What is the purpose of living?
Machine: To live forever.

Berkay Antmen
March 8, 2016

Conversational model

- Purpose: Given previous sentences of the dialogue and context, output a response

- Why?

- goal driven dialogue systems
 - e.g. tech support
- non-goal driven dialogue systems
 - e.g. language learning, video game characters

- How?

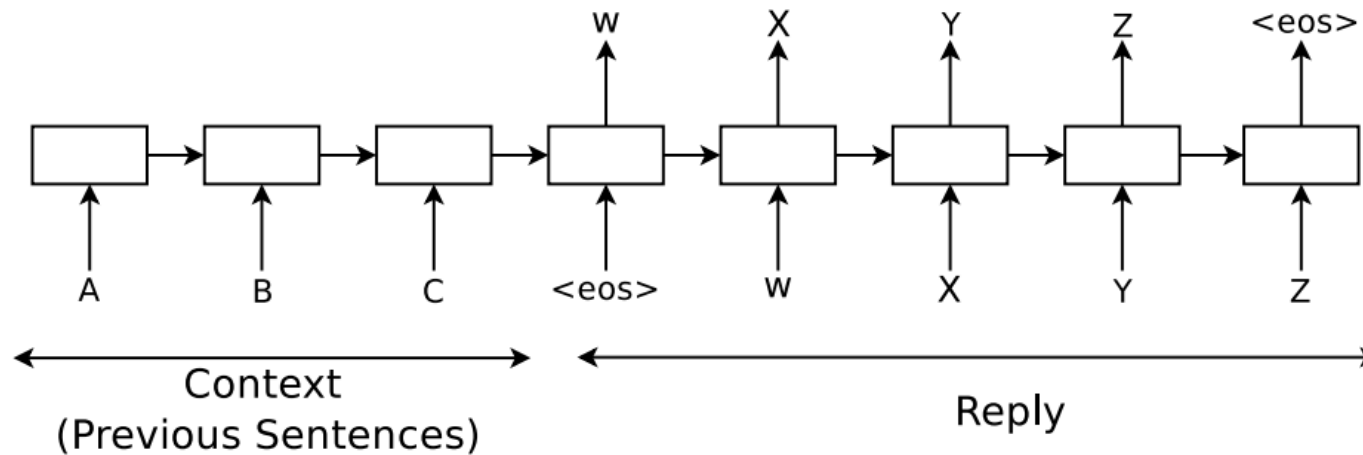
- discriminative
- generative
- heavily hand-crafted
- data-driven systems

Demo (Cleverbot)

- <http://www.cleverbot.com/>
- <http://www.cleverbot.com/conv/201603150055/VWU01366204> Hi-can-you-help-me (Troubleshooting)
- <http://www.cleverbot.com/conv/201603150111/VWU01366307> Hello (Basic)
- <http://www.cleverbot.com/conv/201603150120/VWU01366357> What-is-the-purpose-of-life (Philosophical)
- <http://www.cleverbot.com/conv/201603150204/VWU01366635> We-are-no-strangers-to-love (extra)

Frameworks

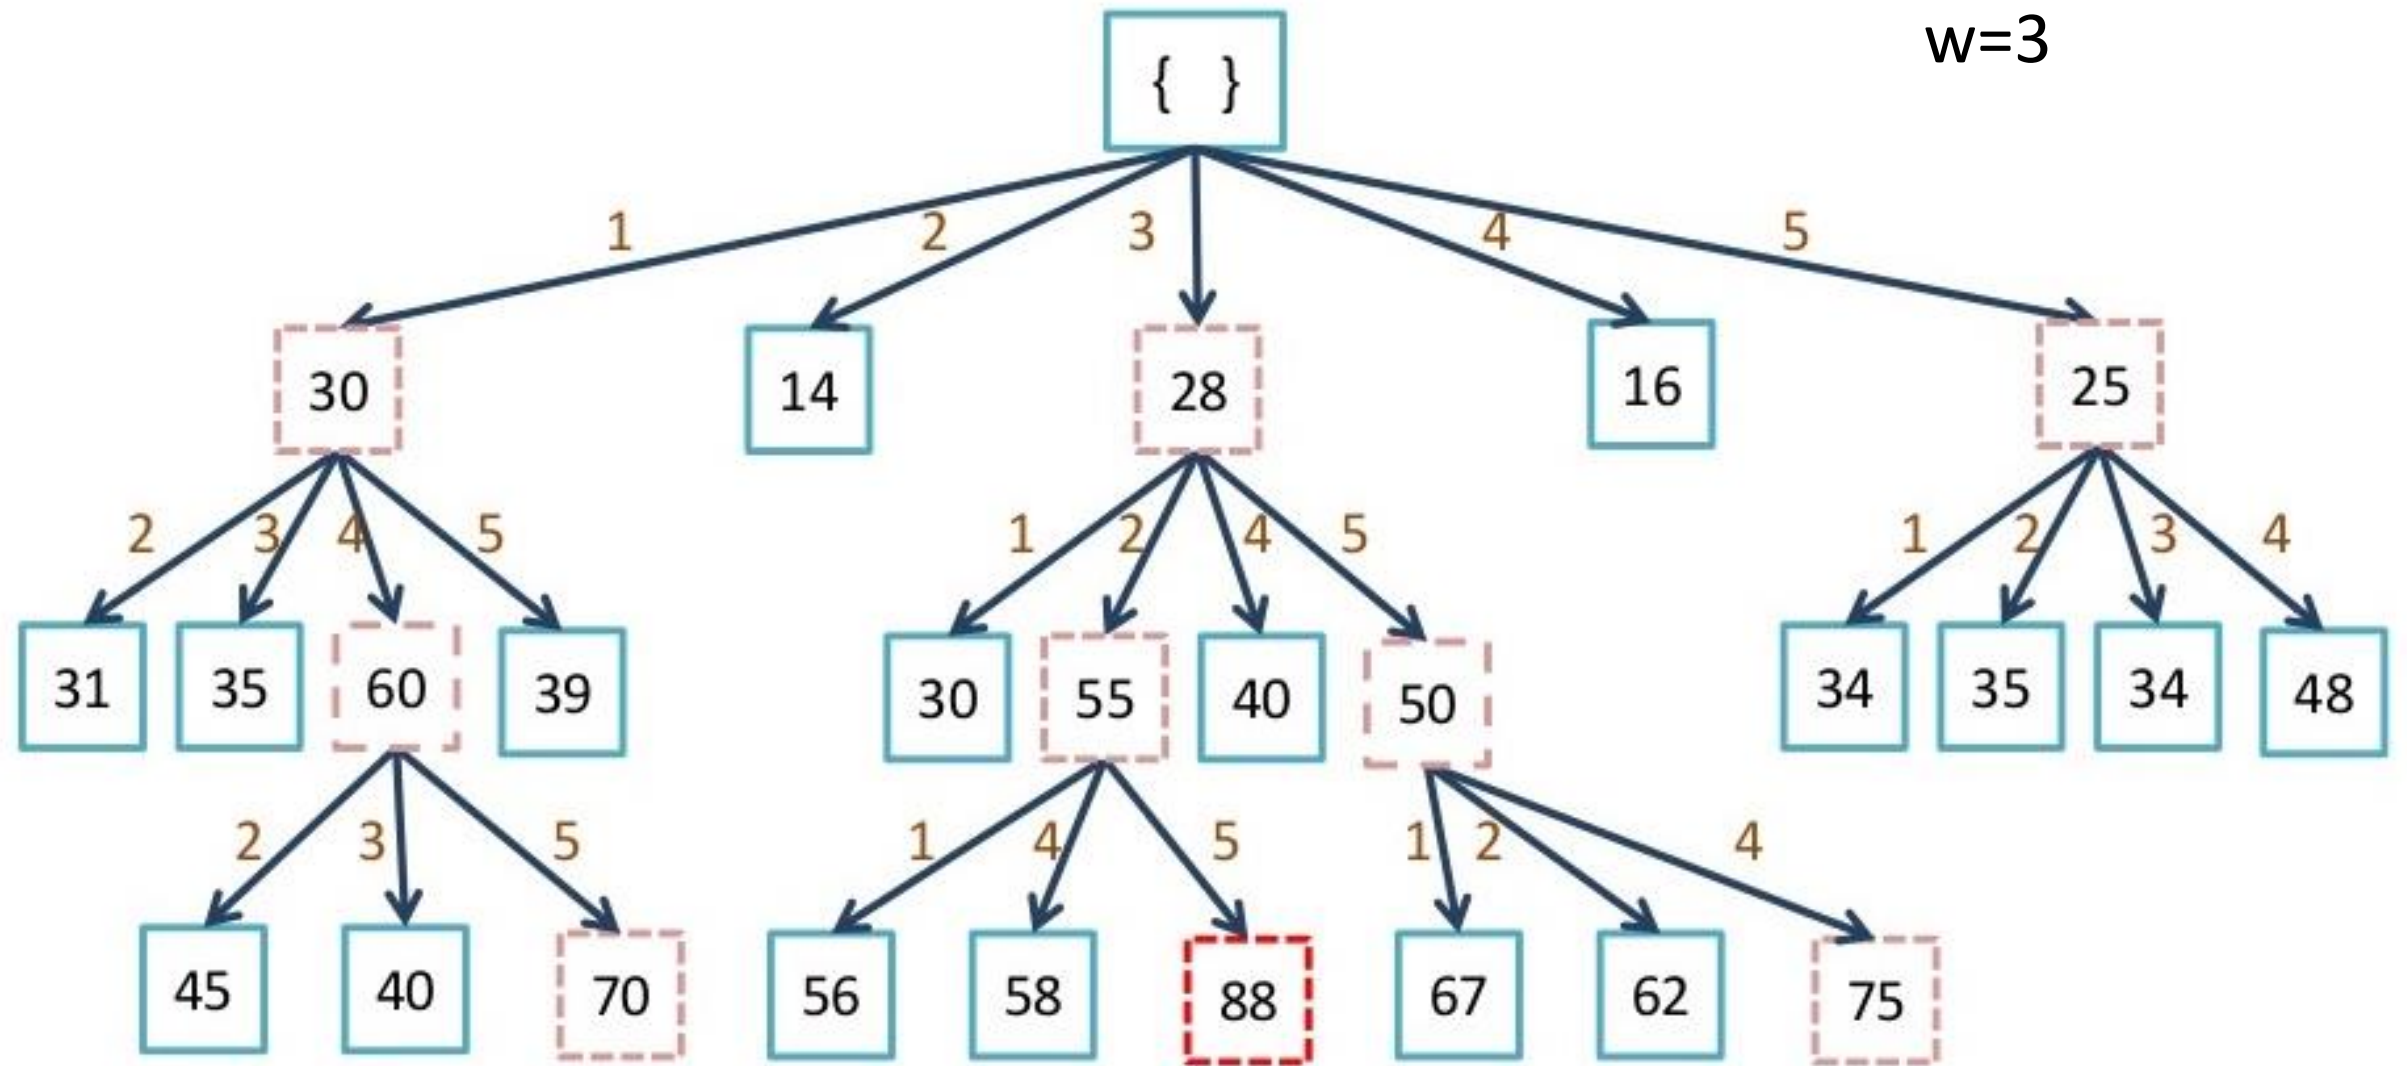
- sequence-to-sequence (seq2seq)
- classification problem over a known vocabulary
- input: sequence of tokens
- output: sequence of tokens



Frameworks: seq2seq

- The goal: estimate $p(y_1, \dots, y_{T'} | x_1, \dots, x_T)$
- problem: boundaries
- solution:
$$p(y_1, \dots, y_{T'} | x_1, \dots, x_T) = \prod_{t=1}^{T'} p(y_t | v, y_1, \dots, y_{t-1})$$
- training: maximize $\frac{1}{|\mathcal{S}|} \sum_{(T,S) \in \mathcal{S}} \log p(T|S)$ (target given source)
- inference: $\hat{T} = \arg \max_T p(T|S)$
 - approximated by beam search

Beam Search



A Neural Conversational Model

- IT helpdesk dataset of conversations (closed-domain)
- OpenSubtitles movie transcript dataset (open-domain)
- Experiments: troubleshooting, general knowledge, philosophical etc.

A Neural Conversational Model

- training: maximize cross entropy of the correct sequence given its context
 - (aside) how is cross entropy measured when the true distribution of the words in the corpus is not known? Monte Carlo estimation: training set is treated as samples from the true distribution

$$H(T, q) = - \sum_{i=1}^N \frac{1}{N} \log_2 q(x_i)$$

- inference: greedy search

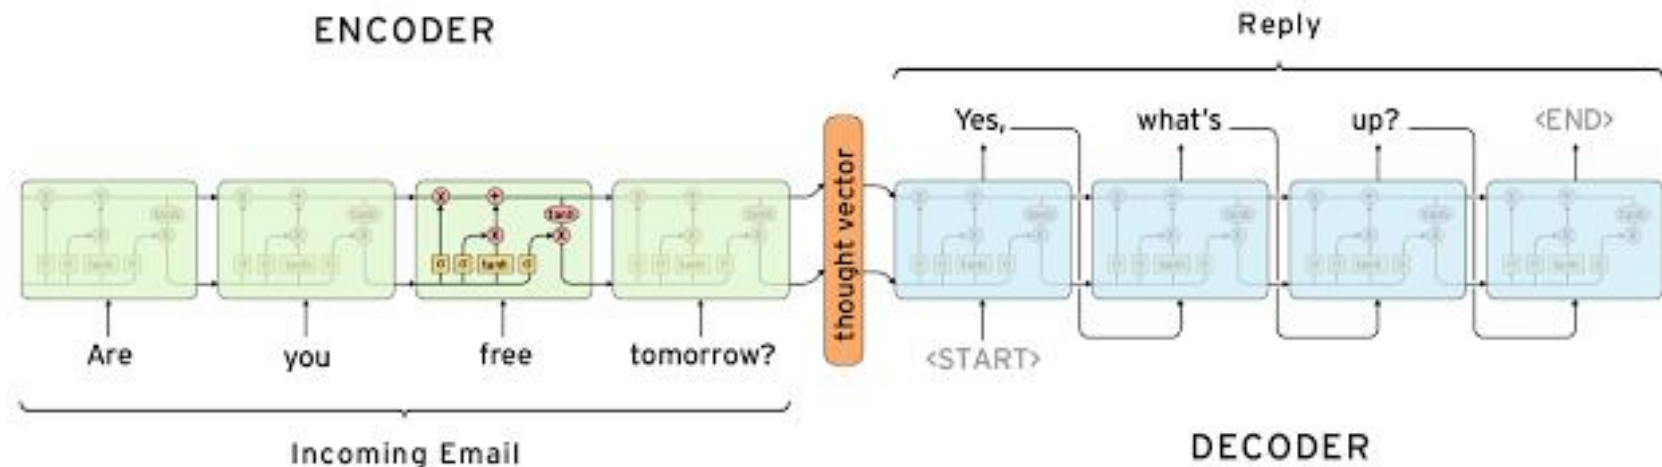


image: Chris Olah

Some results (troubleshooting)

Password issues

Describe your problem: i forgot my password

Machine: *hi , this is <NAME> from helpdesk connect , what can i help you with today ?*

Human: *hi , i forgot my password*

Machine: *ok great*

Human: *can you tell me how i can recover it ?*

Machine: *i ' ll need to verify who it is at that machine*

Machine: *can we do a hangout*

Machine: *?*

Human: *yes*

Machine: *<anonymized hangout instructions and verifications>*

Human: *cool , i am good now*

Machine: *how ' s this looking good ?*

Machine: *i ' ll go ahead and close this chat .*

Machine: *take care . :)*

Human: *bye*

Machine: *bye .*

Cleverbot:

<http://www.cleverbot.com/conv/201603150055/VWU0136620>

[4 Hi-can-you-help-me](#)

Browser issues

Describe your problem: browser keeps crashing

Machine: *hi*

Human: *hi*

Machine: *could you please let me know what is the operating system you are using ?*

Human: *linux*

Machine: *ok this is the issue with the system performance , nasser , tools .*

Human: *what is that ?*

Machine: *is it a windows machine ?*

Machine: *we have linux on your machine .*

Machine: *is it windows xp or windows 7 ?*

Human: *i am using linux*

Machine: *let me take a remote of your machine and check if issue has been pushed properly .*

Human: *ok*

Machine: *please open terminal and run the below commands in one computer :*

Machine: *sudo apt - get update*

Machine: *sudo apt - get upgrade*

Machine: *sudo <COMMAND> - updater*

Human: *done*

Machine: *let me try connecting now*

Human: *ok*

Machine: *its done*

Machine: *should use the second or hit run on the background*

Human: *done*

Machine: *try signing out now and follow the prompts as per the message given above*

Machine: *shd i reboot the machine .*

Human: *ok , please go ahead*

Machine: *once its completed please try to login*

Some more results

Basic

Human: *hello !*

Machine: *hello !*

Human: *how are you ?*

Machine: *i 'm good .*

Human: *what 's your name ?*

Machine: *i 'm julia .*

Human: *when were you born ?*

Machine: *july 20th .*

Human: *what year were you born ?*

Machine: *1977 .*

Human: *where are you from ?*

Machine: *i 'm out in the boonies.*

Human: *are you a man or a woman ?*

Machine: *i 'm a woman .*

Human: *why are we here ?*

Machine: *i 'm not sure .*

Human: *okay, bye !*

Machine: *bye .*

Human: *see you later .*

Machine: *bye .*

Contexts and multiple choice

Human: *my name is david . what is my name ?*

Machine: *david .*

Human: *my name is john . what is my name ?*

Machine: *john .*

Human: *are you a leader or a follower ?*

Machine: *i 'm a leader .*

Human: *are you a follower or a leader ?*

Machine: *i 'm a leader .*

Cleverbot:

<http://www.cleverbot.com/conv/201603150111/VWU01366307> Hello

Some more results

Philosophical

Human: *what is the purpose of life ?*

Machine: *to serve the greater good .*

Human: *what is the purpose of living ?*

Machine: *to live forever .*

Human: *what is the purpose of existence ?*

Machine: *to find out what happens when we get to the planet earth .*

Human: *where are you now ?*

Machine: *i 'm in the middle of nowhere .*

Human: *what is the purpose of dying ?*

Machine: *to have a life .*

Human: *what is the purpose of being intelligent ?*

Machine: *to find out what it is .*

Human: *what is the purpose of emotions ?*

Machine: *i don 't know .*

Cleverbot:

<http://www.cleverbot.com/conv/201603150120/VWU0136635>

[7 What-is-the-purpose-of-life](#)

Opinions

Human: *what do you think about tesla ?*

Machine: *he 's a good conductor .*

Human: *what do you think about bill gates ?*

Machine: *he 's a good man .*

Human: *What do you think about messi ?*

Machine: *he 's a great player .*

Human: *what do you think about cleopatra ?*

Machine: *oh , she 's very regal .*

Human: *what do you think about england during the reign of elizabeth ?*

Machine: *it was a great place .*

Evaluation

- Perplexity measures how well a model predicts the given samples
- $2^{H_q(S_1, \dots, S_n)} = 2^{-\sum_i q(S_i) \log_2 q(S_i)}$

Experiment	Model	Perplexity
IT Helpdesk Troubleshooting	N-grams	18
IT Helpdesk Troubleshooting	Neural conversational model	8
OpenSubtitles	N-grams	28
OpenSubtitles	Neural conversational model	17

Evaluation

- human evaluation against a rule-based bot (CleverBot)
 - asked a list of questions to both models
 - judges picked the bot they preferred
 - Mechanical Turk

# questions	# judges	# prefer neural model	# prefer CleverBot	# tie	# disagreement
200	4	97	60	20	23

Wrong objective function?

- the answers are not diverse, i.e. likely to give most probable answers without giving out much information
 - e.g. S="How old are you?" T="I don't know."
 - $p(T|S)$ high, $p(S|T)$ low
 - e.g. S="How old are you?" T="I am 10 years old"
 - $p(T|S)$ lower, $p(S|T)$ higher

- not really obvious from the selected examples in the paper

A Diversity-Promoting Objective Function for Neural Conversation Models

Input (message)	Output (response)	$\log p$
What are you doing?	I don't know.	-0.862
	I don't know!	-1.035
	Nothing.	-1.062
	Get out of the way.	-1.088
	Get out of here.	-1.089
	I'm going home.	-1.092
	Oh my god!	-1.094
	I'm talking to you.	-1.100
What is your name?	I don't know.	-0.913
	I don't know!	-0.916
	I don't know, sir.	-0.920
	Oh, my god!	-0.974

	My name is Robert.	-1.553
	My name is John.	-1.583
	My name's John.	-1.586
How old are you?	I don't know.	-0.794
	I'm fine.	-1.064
	I'm all right.	-1.167
	Don't know.	-1.169
	I'm not sure.	-1.170

	Twenty-five.	-1.637
	Five.	-1.657
Eight.	-1.713	

Input (message)	Output (response)
What are you doing?	I've been looking for you.
	I want to talk to you.
	Just making sure you're OK.
	I told you to shut up.
	Get out of here.
	I'm going home.
What is your name?	I'm looking for a doctor.
	John.
	Blue!
	Peter.
	Tyler.
	Daniel.
How old are you?	My name is John.
	My name is Robert.
	My name's Daniel.
	Twenty-five.
	Twenty-eight.
	Twenty-four.
How old are you?	Long.
	Five.
	15.
	Good.
	Eight.

A Diversity-Promoting Objective Function for Neural Conversation Models

- An alternative objective function: Maximum Mutual Information (MMI)
- maximize mutual information between source (S) and target (T)
 - $I(S, T) = \log\left(\frac{p(S, T)}{p(S)p(T)}\right)$
- $\hat{T} = \arg_T \max\{\log p(T|S) - \lambda \log p(T)\}$
- remember, previously $\hat{T} = \arg \max_T p(T|S)$

Some results (OpenSubtitles)

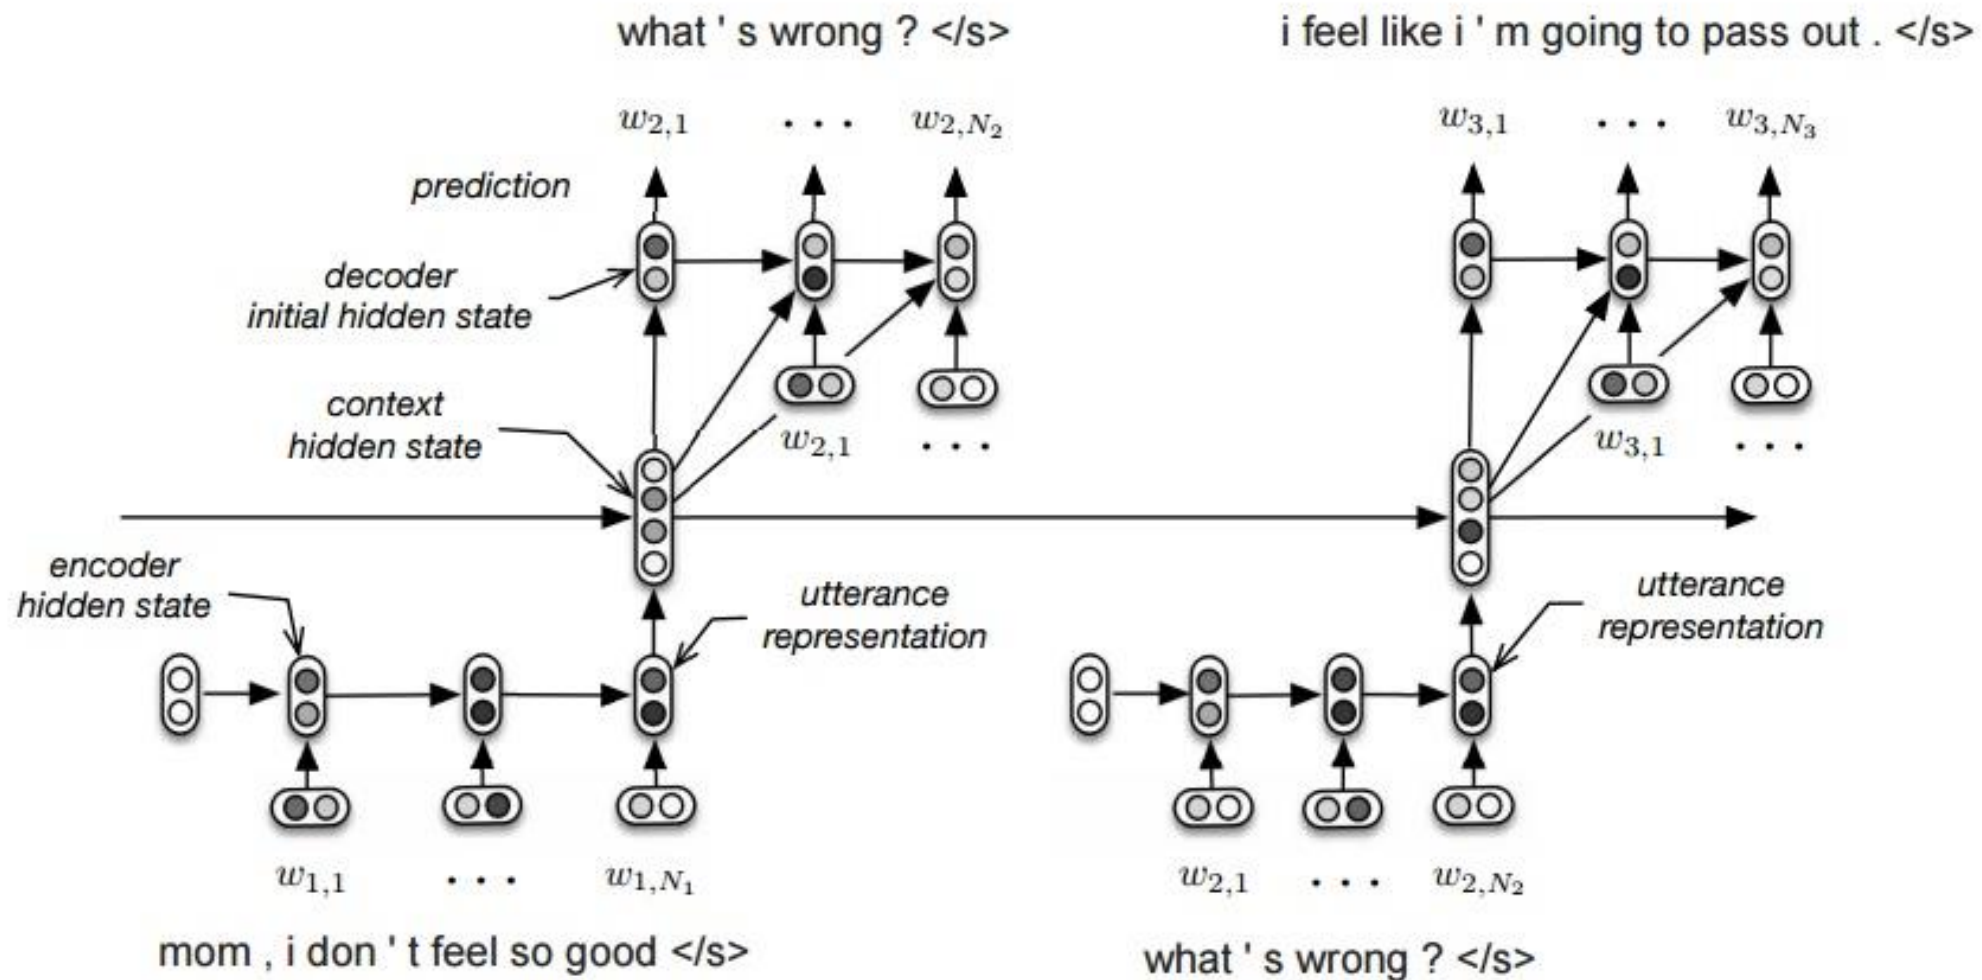
message	SEQ2SEQ	MMI $P(T S) - \lambda P(T)$
I did not get the report from an mi-6 agent.	I don't think that's a good idea.	You did the right thing, did you?
You haven't been given an assignment in this case.	I don't know what you are talking about.	I've been looking all over for you.
I'm losing my grip.	I don't know what you are talking about.	I'm the only one in the world.
I am ready to help.	Come on, come on	I have something we need to talk about.
You programmed me to gather intelligence. That's all I've ever done.	You know that, don't you?	You do have fun, don't you?
I mean, we'd have to talk to him.	I mean, I don't know.	I mean, he's a good guy.
This has nothing to do with him.	You don't know anything about him, do you?	I know him as much.
How come you never say it?	I don't know	Because I don't want to hurt you
He is the one behind this scheme to crash x-net.	I don't know what you are talking about	Yes, he is, isn't he?
How much time do you have here?	I don't know.	Not long enough. Sorry, sir.

Some results (Twitter)

message	SEQ2SEQ	MMI $(1 - \lambda)p(T S) + \lambda p(S T)$
Ohh I've never seen that! How long does it take you guys to learn the drill?	I don't even know.	Like 2 weeks ago!!
You'll know when we sign Heskey in 2 weeks	I'll see what I can do	I sure hope so..
I'm waiting!! I have even got biscuits waiting for you??	I have no idea what you're talking about	I'll send you a pic when I get home xx
My parents actually had touch too much played at their wedding.	I have no idea what you're talking about	How old are they??
Say anything you want... I know you love me and I love you.	I love you so much.	I love you too, baby.
I am off all week next week	What time you going?	What time you going?
How are you doing?	I'm good, thanks.	I'm good, you?


Frameworks

- Hierarchical Recurrent Encoder Decoder (HRED)



Frameworks: HRED

- Motivation?



Hierarchical Neural Network Generative Models for Movie Dialogues

- Non-goal driven: can be easily adapted to specific tasks
- Bootstrapping
 - from word embeddings OR
 - from a large non-dialogue corpus (Q-A SubTle containing 5.5 pairs)
- Interactive dialogue structure
 - end-of-utterance token
 - continued-utterance token

Dataset

- why movie scripts?
 - large dataset
 - wide range of topics
 - long dialogues with few participants
 - relatively few spelling mistakes and acronyms
 - similar to human spoken conversations
 - mostly single dialogue thread
- atomic entries are triples
- 13M words total; 10M in training

Evaluations (movie dialogue generation)

Model	All Tokens				Excluding Stop Words	
	Perplexity	Perplex.@U ₃	Error-Rate	Error-Rate@U ₃	Perplexity	Perplex.@U ₃
RNN	27.09 ± 0.13	26.67 ± 0.19	64.10% ± 0.06	64.07% ± 0.10	75.34 ± 0.47	73.24 ± 0.76
HRED	27.14 ± 0.12	26.60 ± 0.19	64.10% ± 0.06	64.03% ± 0.10	77.17 ± 0.42	74.41 ± 0.66
HRED-Bi.	26.81 ± 0.11	26.31 ± 0.19	63.93% ± 0.06	63.91% ± 0.09	75.71 ± 0.41	73.24 ± 0.64

- test set perplexity and classification errors when bootstrapping from SubTle corpus

Evaluations

Reference (U₁, U₂)	MAP	Target (U₃)
U ₁ : yeah , okay . U ₂ : well , i guess i ' ll be going now .	i ' ll see you tomorrow .	yeah .
U ₁ : oh . <continued_utterance> oh . U ₂ : what ' s the matter , honey ?	i don ' t know .	oh .
U ₁ : it ' s the cheapest . U ₂ : then it ' s the worst kind ?	no , it ' s not .	they ' re all good , sir .
U ₁ : <person> ! what are you doing ? U ₂ : shut up ! c ' mon .	what are you doing here ?	what are you that crazy ?

Future work?

- study larger length dialogues (as opposed to triplets)
- bootstrapping on other non-dialogue but large datasets

Thank you!

Questions?

References

- seq2seq <http://arxiv.org/abs/1409.3215>
- neural conversational <http://arxiv.org/abs/1506.05869>
- hierarchical <http://arxiv.org/abs/1507.02221>
- hierarchical conversational <http://arxiv.org/abs/1507.04808>
- MMI <http://arxiv.org/abs/1510.03055>