Dialog Systems

Mausam

(Slides by Vivian Chen, Asli Celikylmaz, Dilek Hakkani-Tur, Graham Neubig, Manoj Chinnakotla and IITD students)

Dialog Agents

Apple Siri (2011)

Google Now (2012)
Google Assistant (2016)

Microsoft Cortana (2014)

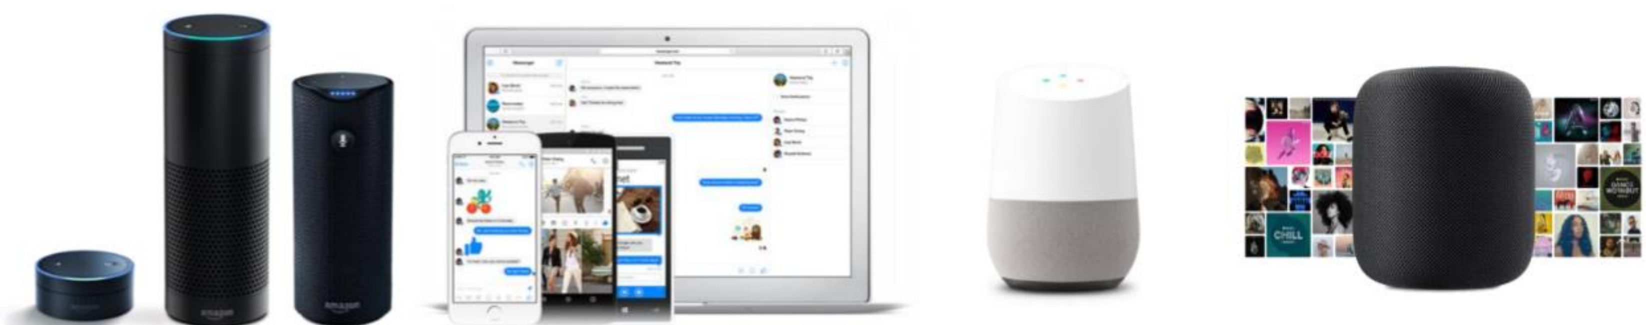

Amazon Alexa/Echo (2014)

Facebook M & Bot (2015)

Google Home (2016)

Apple HomePod (2017)

The Need for Dialog Agents

- Get things done
  - E.g. set up alarm/reminder, take note
- Easy access to structured data, services and apps
  - E.g. find docs/photos/restaurants
- Assist your daily schedule and routine
  - E.g. commute alerts to/from work

Why Natural Language?

Global Digital Statistics (2015 January)

Global Population	Active Internet Users	Active Social Media Accounts	Active Unique Mobile Users
7.21B	3.01B	2.08B	**3.65B**

App → Bot

- A bot is responsible for a “single” domain, similar to an app

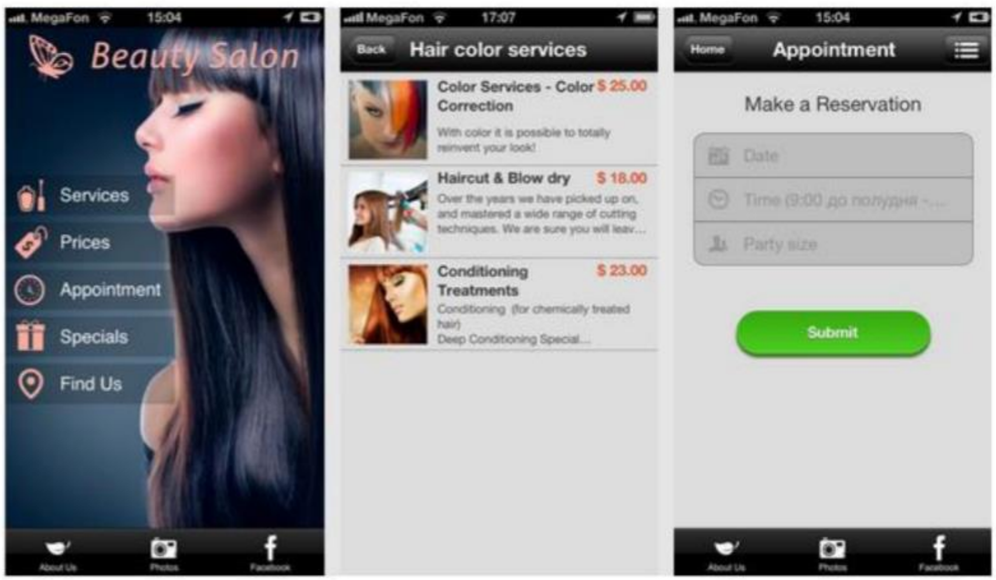

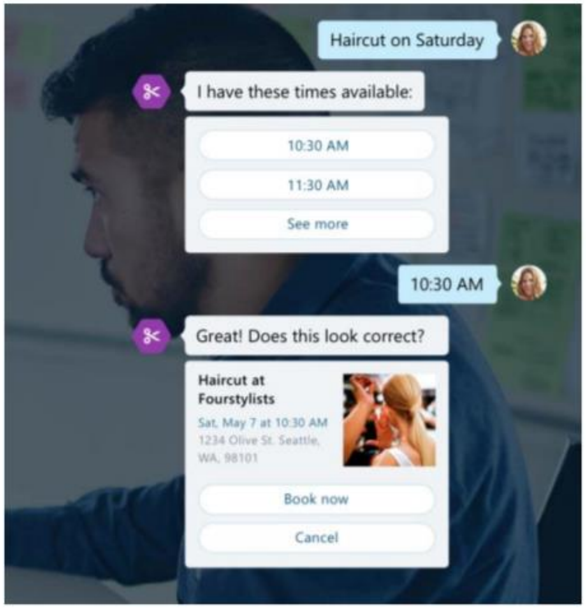

Users can initiate dialogues instead of following the GUI design

GUI vs Conversational UI

	Website/APP's GUI	Msg's CUI
Situation	Navigation, no specific goal	Searching, with specific goal
Information Quantity	More	Less
Information Precision	Low	High
Display	Structured	Non-structured
Interface	Graphics	Language
Manipulation	Click	mainly use texts or speech as input
Learning	Need time to learn and adapt	No need to learn
Entrance	App download	Incorporated in any msg-based interface
Flexibility	Low, like machine manipulation	High, like converse with a human

Two Branches of Bots

Task-Oriented Bot

- Personal assistant, helps users achieve a certain task
- Combination of rules and statistical components
  - POMDP for spoken dialog systems (Williams and Young, 2007)
  - End-to-end trainable task-oriented dialogue system (Wen et al., 2016)
  - End-to-end reinforcement learning dialogue system (Li et al., 2017; Zhao and Eskenazi, 2016)

Chit-Chat Bot

- No specific goal, focus on natural responses
- Using variants of seq2seq model
  - A neural conversation model (Vinyals and Le, 2015)
  - Reinforcement learning for dialogue generation (Li et al., 2016)
  - Conversational contextual cues for response ranking (Al-Rfou et al., 2016)

Dimensions

- Short Conversations
  - The expected number of turns in the dialogue is usually less
  - For Example: factoid question answering (is there a train from A to B?)
  - No need to maintain the entire context
- Long Conversations
  - Multi-turn conversations which may be long
  - For Example: Discussion about movie scene

Traditional Pipeline models

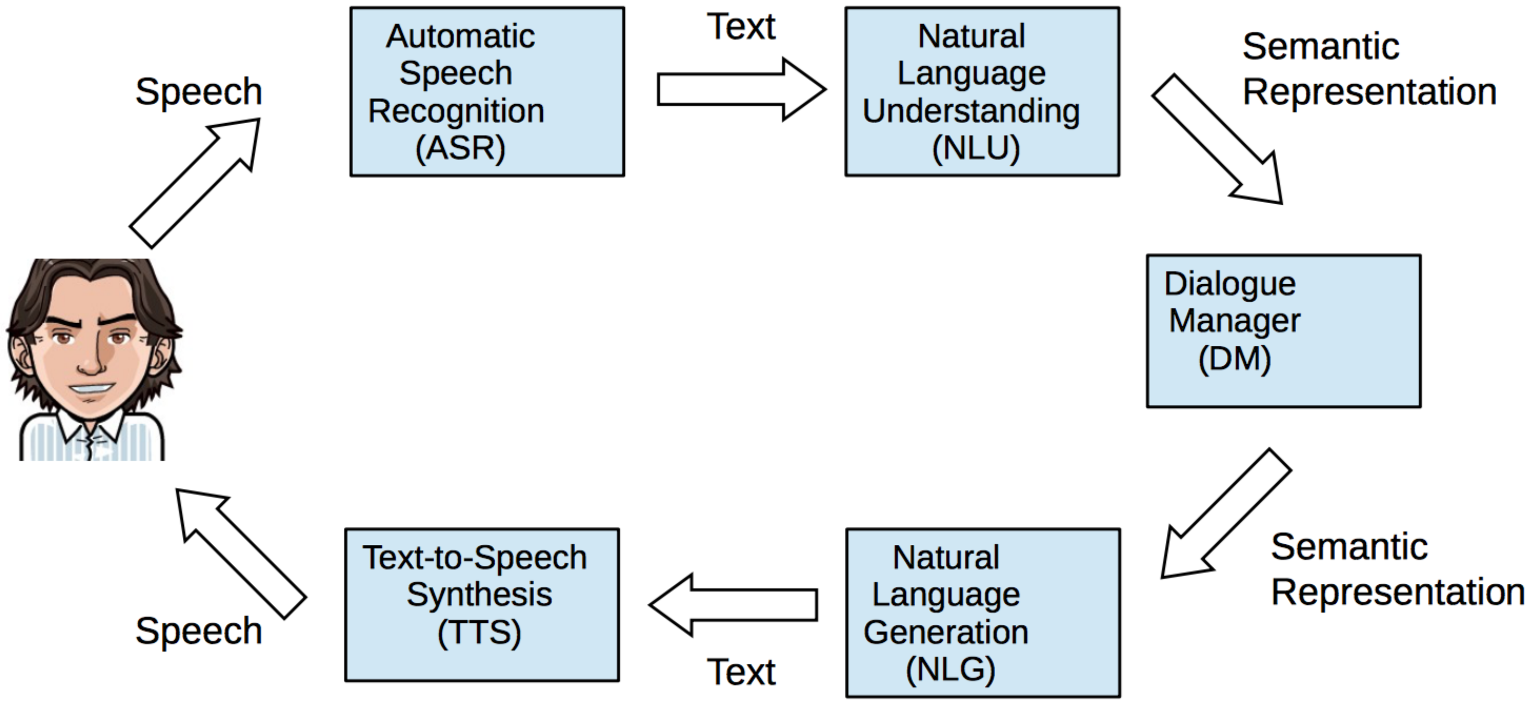

End-To-End models with DL

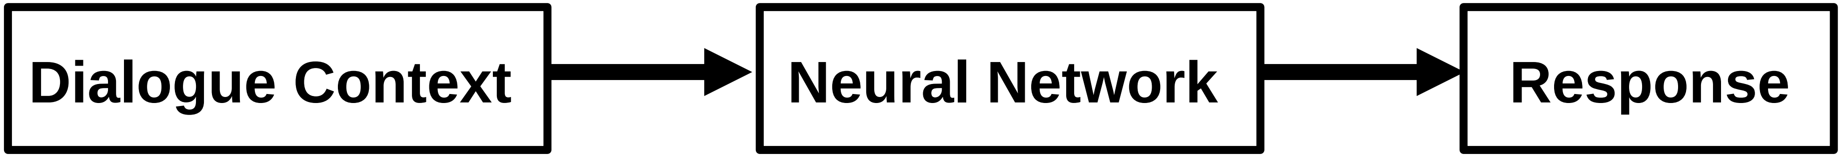

End-To-End models with DL

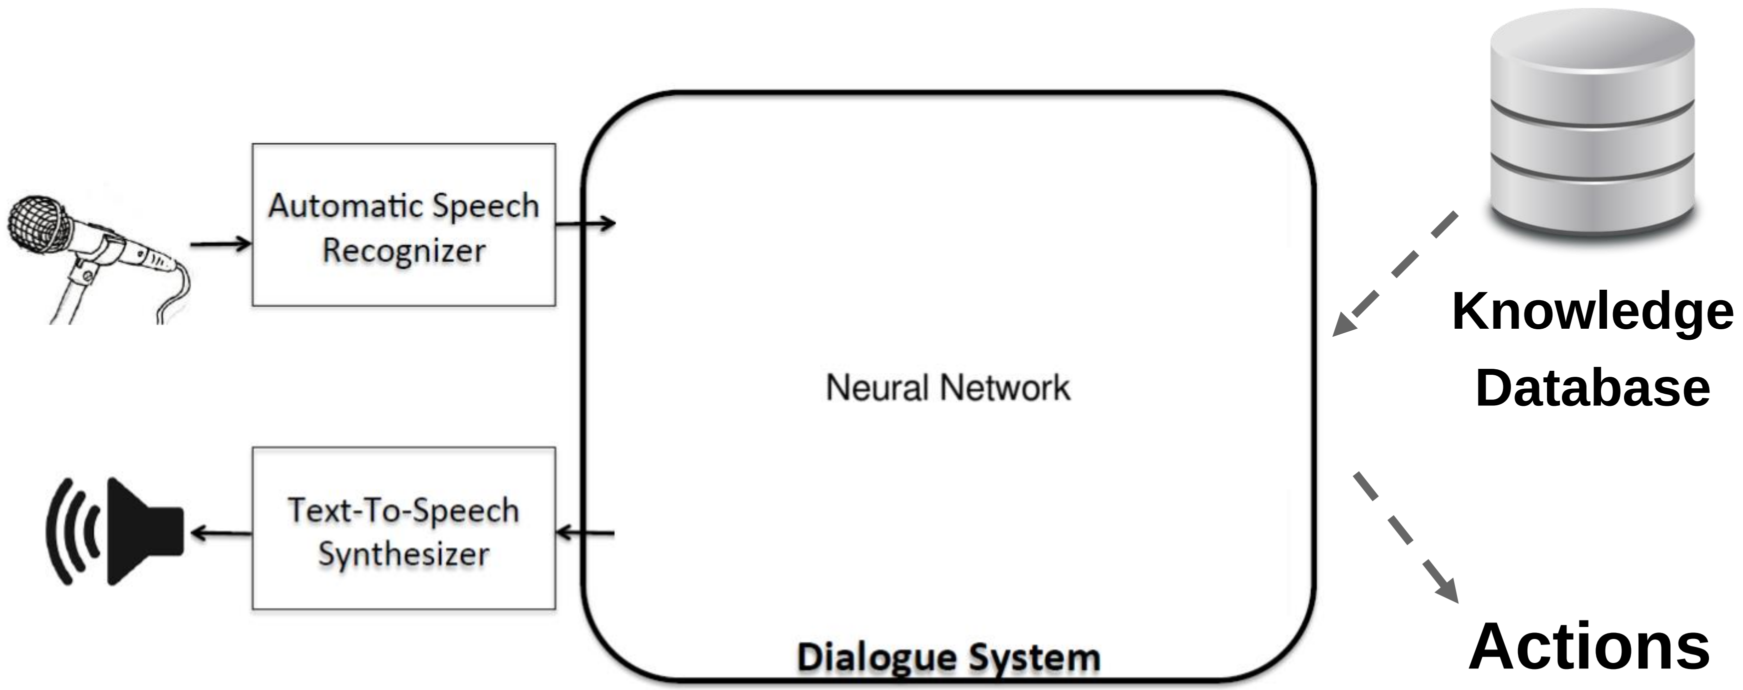

Neural Models

- Mine a large message-response conversation corpus

- Retrieval based Models
  - Given a user utterance, find the semantically most similar message and return it’s response
  - Can only output fixed-set of responses from the corpus

- Generative Models
  - Learn a Machine Learning (ML) based translation model to “translate” a given user message to an appropriate response in the same language
  - Can generate a totally new response not available in

What is a good Chatbot?

- The responses should be
  - Grammatical
  - Coherent
  - In Context
  - Ideally non-Generic responses

Challenges

- Variability in Natural Language
- Robustness
- Recall/Precision Trade-off
- Meaning Representation
- Common Sense, World Knowledge
- Ability to Learn
- Transparency

Encoder-Decoder Model

- Sheng et al 15

ID	Post	NRM-hyb
1	每年元旦都来高烧。 High fever attacks me every New Year's Day	祝早日康复，身体健康！ Get well soon and stay healthy!
2	又老了一岁，感谢我们组，我很高兴 I gain one more year. Grateful to my group, so happy	老了老了，岁月不饶人啊！ Getting old now. Time has no mercy.
3	吃素第一天，坚持住，崔朵拉。 First day of being a vegetarian. Hold on, CuiDuoLa	每天都是素食主义者 I have been being a vegetarian everyday
4	这种天气真应该带个帅锅出去踏青赏花 We should go out with some cute guys to enjoy a great outing in such a nice weather.	这种天气真不是盖的。 It is indeed a very nice weather.

Challenge: Global Coherence

- Incorporating Context
  - Keep tracking of what information has been exchanged, in which context etc.
  - Especially important in long conversations

User: Can you please tell me which flights from London are expected this evening after six

Bot: Sure, let me check

Bot: Here are the flights: EK 201, EK 522

User: and which ones tomorrow morning before twelve?

Soln: Use More Context

Sordoni et al. (2015) consider one additional previous context utterance concatenated together

Vinyals et al. (2015) just concatenate together all previous utterances and hope an RNN an learn

Describe your problem: i am having issues accessing vpn
Machine: *hi*
Human: *hello*
Machine: *could you please let me know what are the operating systems you are using ?*
Human: *i am using linux*
Machine: *ok it was in the ssl vpn on which you stated it was broken*
Machine: *please check if you are logged in to corp network*
Machine: *what is the error that you are running please*
Human: *i am seeing an error related to vpn*
Machine: *what is the error message that you are getting when connecting to vpn using network connect ?*
Human: *connection refused or something like that*
Machine: *may i know the version of network connect you connect ?*
Human: *i am not sure i know that*
Machine: *is the network connect prompting for any user certificate*
Machine: *?*
Human: *yes*
Machine: *are you at home or at work at the moment ?*
Machine: *office ?*
Human: *i am at home*
Machine: *try this*
Machine: *goto <URL> page and enter your name and password*
Human: *done*
Machine: *try logging in with <NAME> and check if its working fine now*
Human: *yes , now it works !*
Machine: *great*
Machine: *anything else that i can help*
Human: *no , that was very useful*
Machine: *you are other than <NAME> :)*
Machine: *have a nice night*

Hierarchical Recurrent Encoder-Decoder (HRED)

- *Encoder* RNN
  - For encoding each utterance independently into an utterance vector
- *Context* RNN
  - For encoding the topic/context of the dialogue up till the current utterance using utterance vectors
- *Decoder* RNN
  - For predicting the next utterance

HRED

- Also have utterance-level RNN track overall dialog state

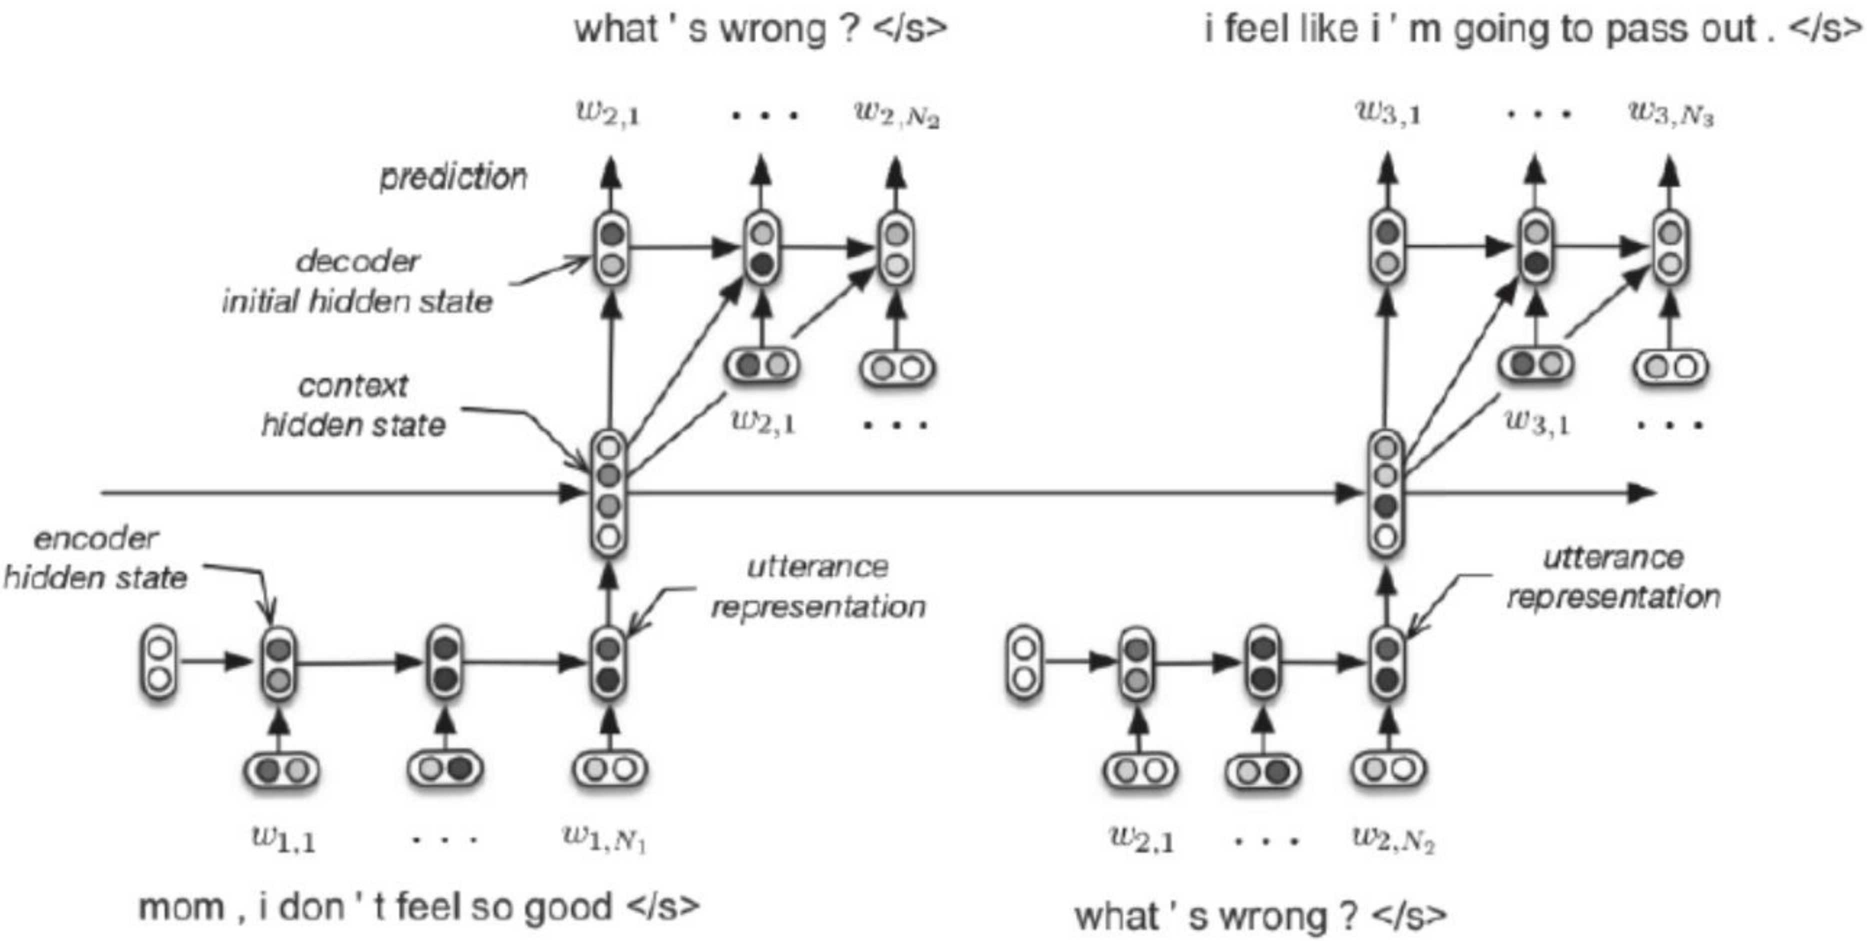

Bootstrapping in HRED

- Initialising with Word2Vec embeddings
  - Trained on Google News dataset

- Pre-training on *SubTle* Q-A dataset
  - 5.5M Q-A pairs
    - Converted to 2-turn dialogue
    - D = {U1 = Q, U2 = A}

Dataset - MovieTriples dataset

	Training	Validation	Test
Movies	484	65	65
Triples	196,308	24,717	24,271
Avg. tokens/triple	53	53	55
Avg. unk/triple	0.97	1.22	1.19

- Open Domain - Wide variety of topics covered
- Names and Numbers replaced with *<person>* and *<number>* tokens
- Vocab of 10K most popular tokens
- Special *<continued-utterance>* and *<end-of-utterance>* tokens to capture breaks

Model	Perplexity	Perplexity@U_3	Error-Rate	Error-Rate@U_3
Backoff N-Gram	64.89	65.05	-	-
Modified Kneser-Ney	60.11	54.75	-	-
Absolute Discounting N-Gram	56.98	57.06	-	-
Witten-Bell Discounting N-Gram	53.30	53.34	-	-
RNN	35.63 ± 0.16	35.30 ± 0.22	$66.34\% \pm 0.06$	$66.32\% \pm 0.08$
DCGM-I	36.10 ± 0.17	36.14 ± 0.26	$66.44\% \pm 0.06$	$66.57\% \pm 0.10$
HRED	36.59 ± 0.19	36.26 ± 0.29	$66.32\% \pm 0.06$	$66.32\% \pm 0.11$
HRED + Word2Vec	33.95 ± 0.16	33.62 ± 0.25	$66.06\% \pm 0.06$	$66.05\% \pm 0.09$
RNN + SubTle	27.09 ± 0.13	26.67 ± 0.19	$64.10\% \pm 0.06$	$64.07\% \pm 0.10$
HRED + SubTle	27.14 ± 0.12	26.60 ± 0.19	$64.10\% \pm 0.06$	$64.03\% \pm 0.10$
HRED-Bi. + SubTle	$\mathbf{26.81 \pm 0.11}$	$\mathbf{26.31 \pm 0.19}$	$\mathbf{63.93\% \pm 0.06}$	$\mathbf{63.91\% \pm 0.09}$

MAP Output

- - Most probable last utterance
- - Found using beam search for better approximation
-     - Generic responses observed
- - Stochastic sampling gives more diverse dialogues

Reference (U_1, U_2)	**MAP**	**Target** (U_3)
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 ?

Challenge: Varied/Interesting Responses

For translation, there is lexical variation but content remains the same

For dialog, content will also be different! (e.g. Li et al. 2016)

Input: What are you doing?			
−0.86	I don't know.	−1.09	Get out of here.
−1.03	I don't know!	−1.09	I'm going home.
−1.06	Nothing.	−1.09	Oh my god!
−1.09	Get out of the way.	−1.10	I'm talking to you.
Input: what is your name?			
−0.91	I don't know.	...	
−0.92	I don't know!	−1.55	My name is Robert.
−0.92	I don't know, sir.	−1.58	My name is John.
−0.97	Oh, my god!	−1.59	My name's John.
Input: How old are you?			
−0.79	I don't know.	...	
−1.06	I'm fine.	−1.64	Twenty-five.
−1.17	I'm all right.	−1.66	Five.
−1.17	I'm not sure.	−1.71	Eight.

Soln: Diversity Promoting Objective

- Basic idea: we want responses that are likely given the context, unlikely otherwise
- Method: subtract weighted unconditioned log probability from conditioned probability (calculated only on first few words)

$$\log p(T|S) - \lambda \log U(T)$$

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?

Diversity: A Problem for Evaluation

Translation uses BLEU score; while imperfect, not horrible

In dialog, BLEU shows very little correlation (Liu et al. 2016)

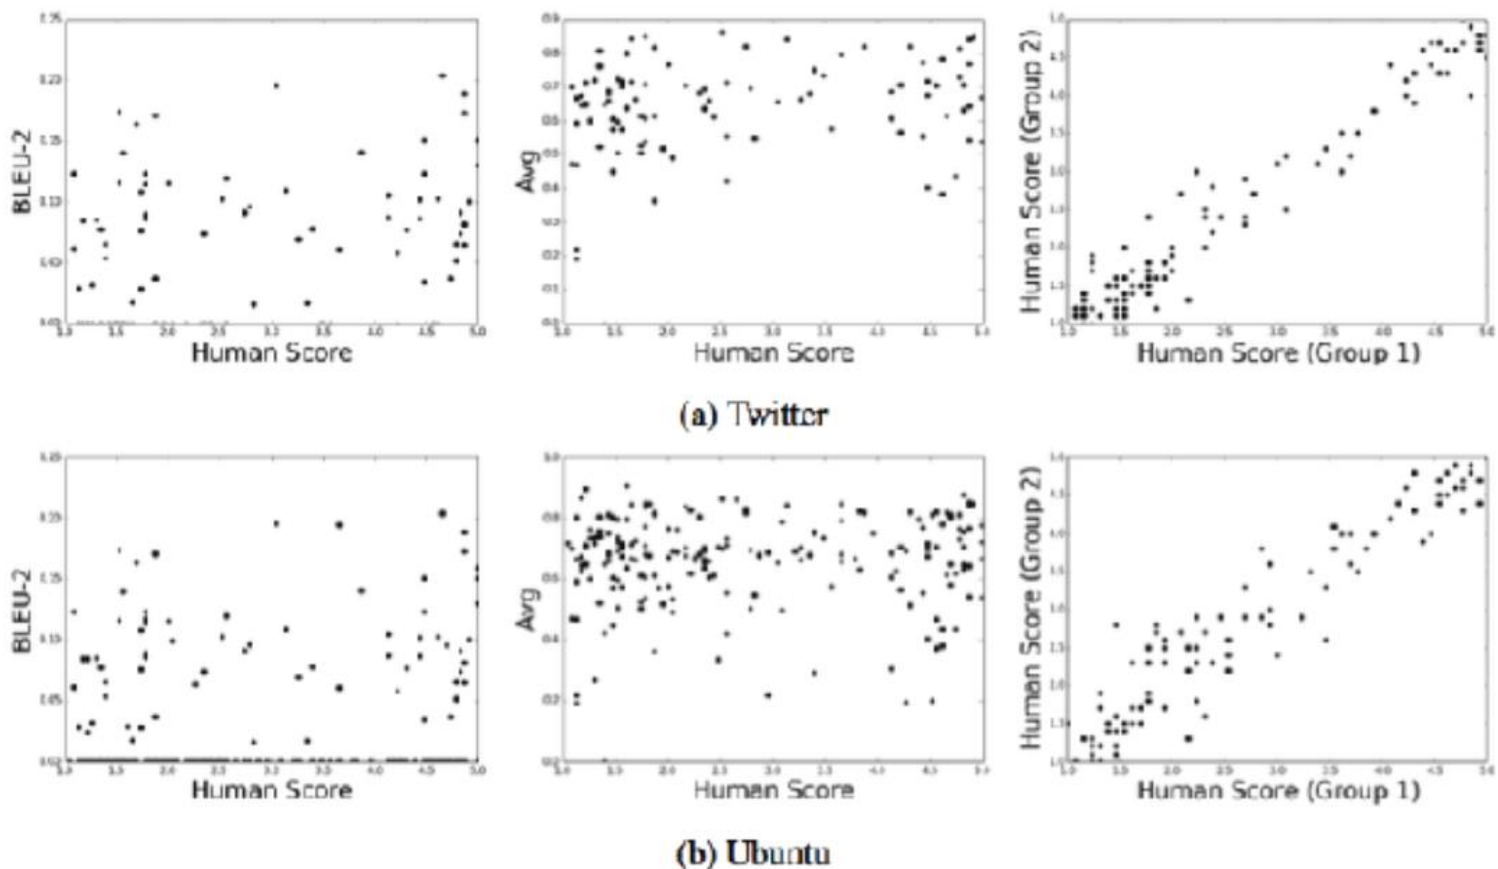

Soln 1: Multiple References

- Retrieve good-looking responses, perform human evaluation, up-weight good ones, down-weight bad ones

Context c	Message m	Response r	Score
i was about to text you and my two cousins got excited cause they thought you were "rihanna"	aww, i can imagine their disappointment	they were very disappointed!!!	0.6
yes. my ex-boyfriend, killed my cat. like i say, it was the start of a bad time...	*i can imagine!*	*yes. luckily, the whole thing feels very much of the past now.*	0.8
its good.. for some reason i can't name stand out tracks but i've been playing it since it dropped	*i can imagine, banks doesn't disappoint*	*na this is anything but a disappointment..*	0.6
at my lil cousins dancing to "dance for you". these kids are a mess.	*lmaoo i can imagine.*	*they were belly rolling, filarious.*	0.4
what's sick about it?? do you know how long it is?? no so how is it sick?	*i can imagine it*	*your imagination is wrong, very wrong at that.*	-0.1
it's saad oh yeah the snow is very beautiful	*yeah i can imagine*	*the weather in russia is very cool.*	-0.7

Soln 2: Learn to Evaluate

- Use context, true response, and actual response to learn a regressor that predicts goodness (Lowe et al. 2017)
- Important: similar to model, but has access to reference!

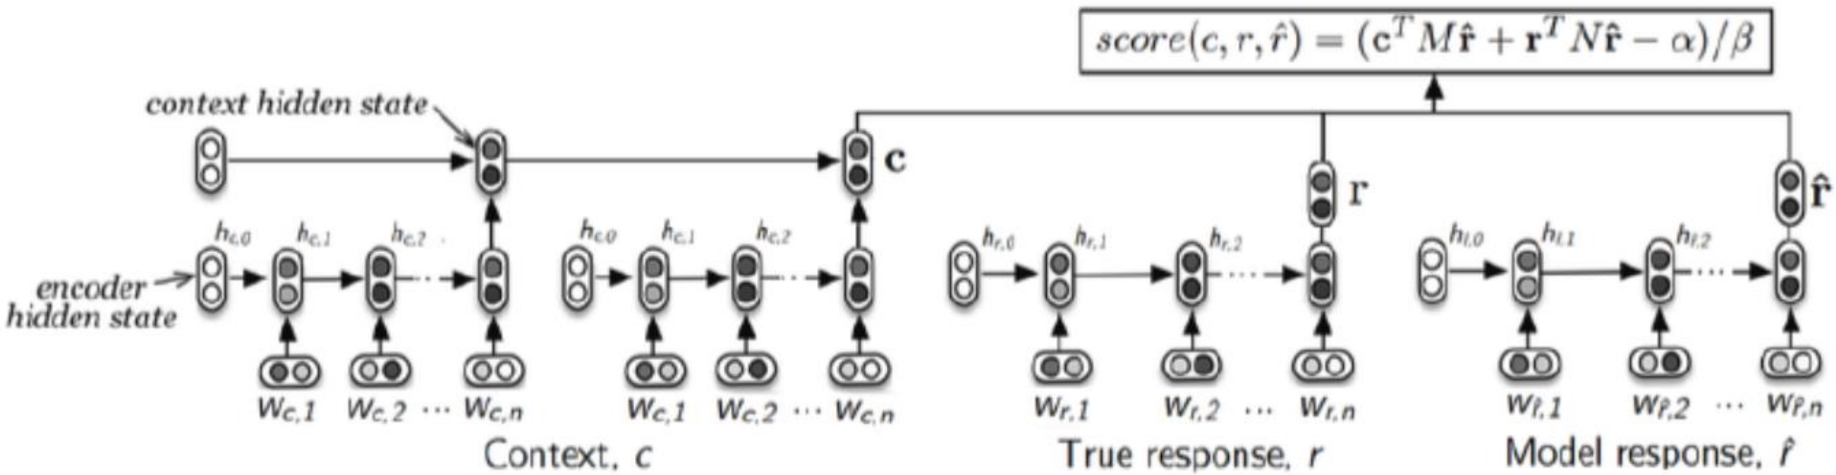

Challenge: Lack of Personality

If we train on all of our data, our agent will be a mish-mash of personalities (e.g. Li et al. 2016)

message	Where do you live now?
response	I live in Los Angeles.
message	In which city do you live now?
response	I live in Madrid.
message	In which country do you live now?
response	England, you?

We would like our agents to be consistent!

Personality Infused Dialog

(Mairesse et al. 2007)

Train a generation system with controllable "knobs" based on personality traits

e.g. Extraversion:

Non-neural, but well done and perhaps applicable

Alt	Realization	Extra
5	Err... it seems to me that Le Marais isn't as bad as the others.	1.83
4	Right, I mean, Le Marais is the only restaurant that is any good.	2.83
8	Ok, I mean, Le Marais is a quite french, kosher and steak house place, you know and the atmosphere isn't nasty, it has nice atmosphere. It has friendly service. It seems to me that the service is nice. It isn't as bad as the others, is it?	5.17
9	Well, it seems to me that I am sure you would like Le Marais. It has good food, the food is sort of rather tasty, the ambience is nice, the atmosphere isn't sort of nasty, it features rather friendly servers and its price is around 44 dollars.	5.83
3	I am sure you would like Le Marais, you know. The atmosphere is acceptable, the servers are nice and it's a french, kosher and steak house place. Actually, the food is good, even if its price is 44 dollars.	6.00
10	It seems to me that Le Marais isn't as bad as the others. It's a french, kosher and steak house place. It has friendly servers, you know but it's somewhat expensive, you know!	6.17
2	Basically, actually, I am sure you would like Le Marais. It features friendly service and acceptable atmosphere and it's a french, kosher and steak house place. Even if its price is 44 dollars, it just has really good food, nice food.	6.17

Persona based Dialog

What is a persona?

- Composite of various elements of identity
- Background facts, user profile, language behaviour and interaction style

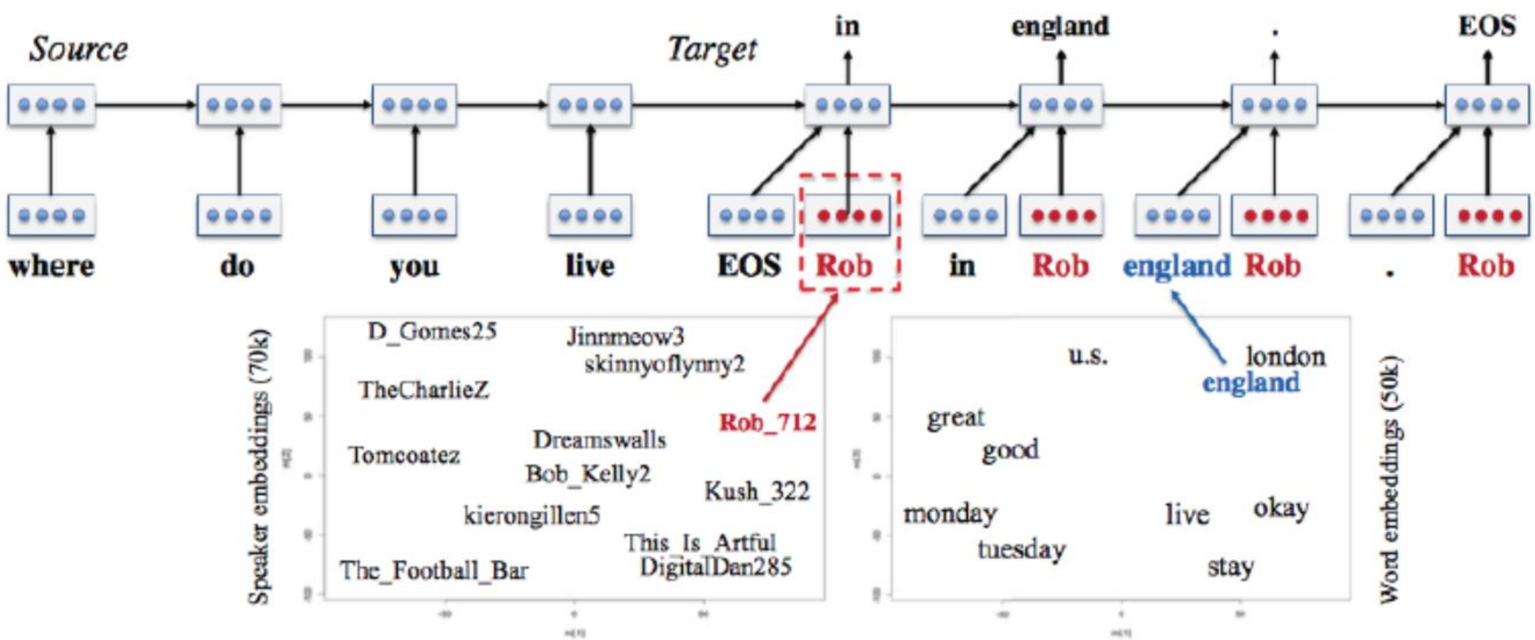

- Also model who the speaker is speaking to in speaker-addressee model

Loss Functions

- Language Model Log-likelihood

- BLEU
  - non differentiable
  - Use RL

Supervised learning

We are given the correct decisions

$$\ell_{\text{super}}(Y, X) = -\log P(Y \mid X)$$

In the context of reinforcement learning, this is also called "imitation learning," imitating a teacher (although imitation learning is more general)

Self-learning

Sample or argmax according to the current model

$$\hat{Y} \sim P(Y \mid X) \quad \text{or} \quad \hat{Y} = \mathrm{argmax}_Y P(Y \mid X)$$

Use this sample (or samples) to maximize likelihood

$$\ell_{\mathrm{self}}(X) = -\log P(\hat{Y} \mid X)$$

No correct answer needed! But is this a good idea?

Reinforcement Learning

Policy Gradient/Reinforce

Add a term that scales the loss by the reward

$$\ell_{\text{r.l.}}(X) = -R(\hat{Y}) \log P(\hat{Y} \mid X)$$

Outputs that get a bigger reward will get a higher weight

Quiz: Under what conditions is this equal to MLE?