

Neural Approaches to Conversational AI

Jianfeng Gao

Microsoft AI & Research

Oct. 25, 2018

Contact Information:

Jianfeng Gao <http://research.microsoft.com/~jfgao>

Michel Galley <http://research.microsoft.com/~mgalley>

Lihong Li <https://lihongli.github.io>

Tech Report (draft):

<https://arxiv.org/pdf/1809.08267.pdf>

Slides and references available:

<http://microsoft.com/en-us/research/publication/neural-approaches-to-conversational-ai/>

Collaborators:

Faisal Ahmed, Chris Brockett, Asli Celikyilmaz, Ming-Wei Chang, Weizhu Chen, Yun-Nung Chen, Li Deng, Bhuwan Dhingra, Bill Dolan, Michel Galley, Marjan Ghazvininejad, Xiaodong He, Po-Sen Huang, Sungjin Lee, Jiwei Li, Lihong Li, Xiujun Li, Zachary Lipton, Xiaodong Liu, Rangan Majumder, Nasrin Mostafazadeh, Baolin Peng, Mir Rosenberg, Yelong Shen, Alessandro Sordoni, Saurabh Tiwary, Lucy Vanderwende, Luan Yi, Scott Yih et al.

Neural Approaches to Conversational AI

Question Answering, Task-Oriented Dialogues
and Chatbots: A Unified View

Suggested Citation: Jianfeng Gao, Michel Galley and Lihong Li (2018). "Neural Approaches to Conversational AI". : Vol. xx, No. xx, pp 1–18. DOI: 10.1561/XXXXXXXXXX.

Jianfeng Gao
Microsoft Research
jfgao@microsoft.com

Michel Galley
Microsoft Research
mgalley@microsoft.com

Lihong Li
Google Inc.
lihong@google.com

This article may be used only for the purpose of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval.

now
the essence of knowledge
Boston — Delft

Outline

- **Part 1: Introduction**
 - **Dialogue: what kinds of problem**
 - **A unified view: dialogue as optimal decision making**
- Part 2: Question answering and machine reading comprehension
- Part 3: Task-oriented dialogue
- Part 4: Fully data-driven conversation models and social chatbots

Aspirational Goal: Enterprise Assistant

Task Completion



Where are sales lagging behind our forecast?

The worst region is [country], where sales are XX% below projections

Do you know why?

QA (decision support)

The forecast for [product] growth was overly optimistic

How can we turn this around?

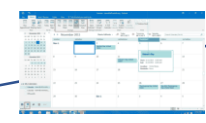
Here are the 10 customers in [country] with the most growth potential, per our CRM model

Can you set up a meeting with the CTO of [company]?

Info Consumption

Yes, I've set up a meeting with [person name] for next month when you're in [location]

Task Completion



Thanks

What kinds of problems?

“I am smart”

Turing Test (“I” talk like a human)

“I have a question”

Information consumption

“I need to get this done”

Task completion

“What should I do?”

Decision support

What kinds of problems?

“I am smart”

“I have a question”

“I need to get this done”

“What should I do?”

Turing Test

Information consumption

Task completion

Decision support

- *What is the employee review schedule?*
- *What room is the project review meeting in?*
- *When is the ACL 2018 conference?*
- *What does DNN stand for?*

What kinds of problems?

“I am smart”

Turing Test

“I have a question”

Information consumption

“I need to get this done”

Task completion

“What should I do?”

Decision support

- *Book me the biz trip to San Francisco*
- *Reserve a table at Kisaku for 5 people, 7PM tonight*
- *Brief me on people in my Thursday 9:00 am meeting*
- *Schedule a meeting with Bill at 10:00 tomorrow.*

What kinds of problems?

“I am smart”

Turing Test

“I have a question”

Information consumption

“I need to get this done”

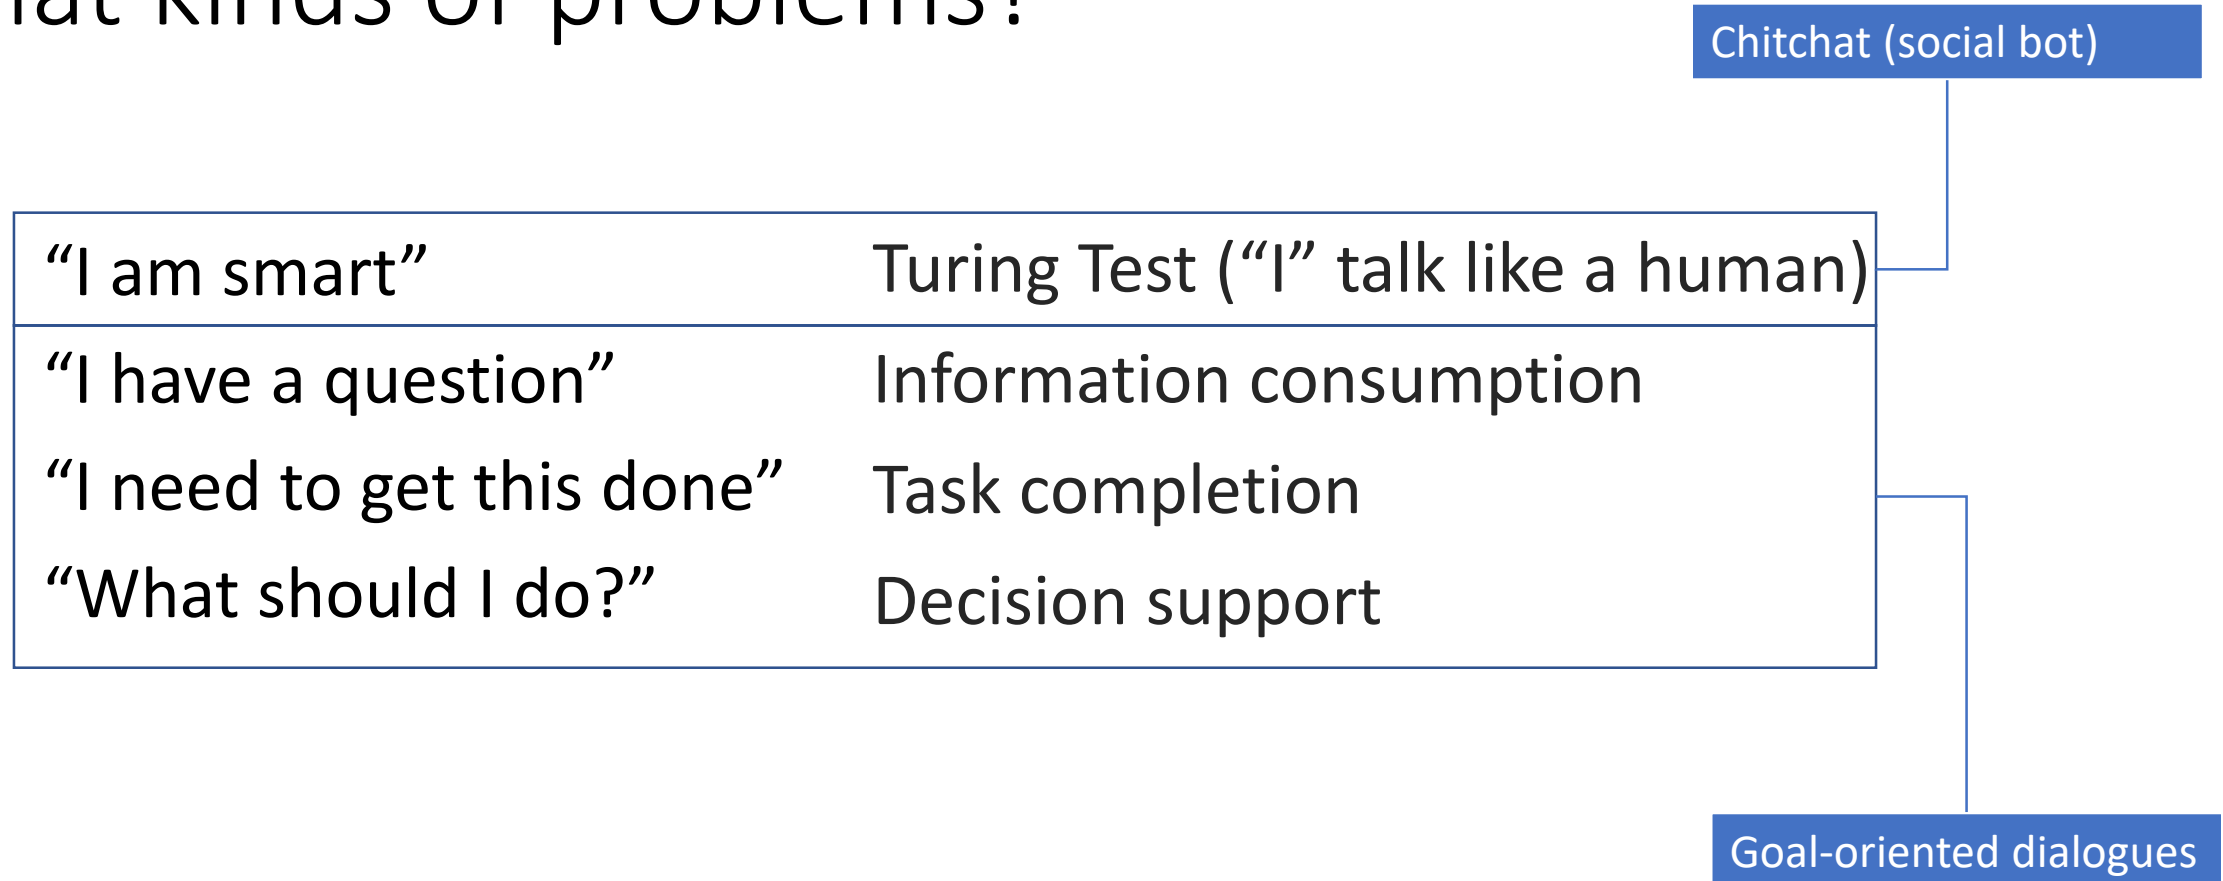
Task completion

“What should I do?”

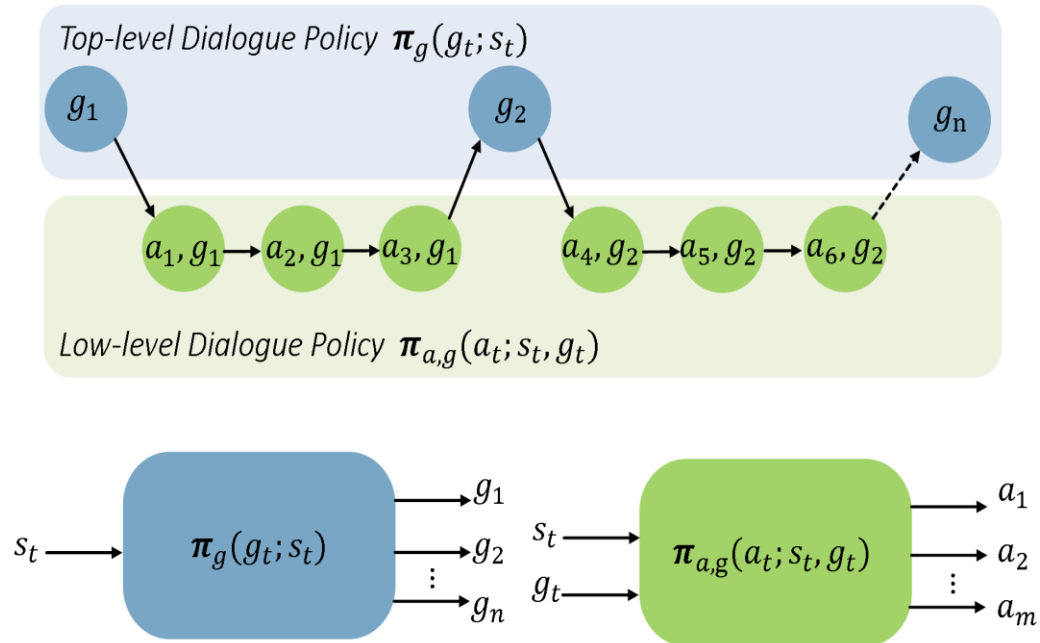
Decision support

- *Why are sales in China so far behind forecast?*

What kinds of problems?



A unified view: dialogue as optimal decision making

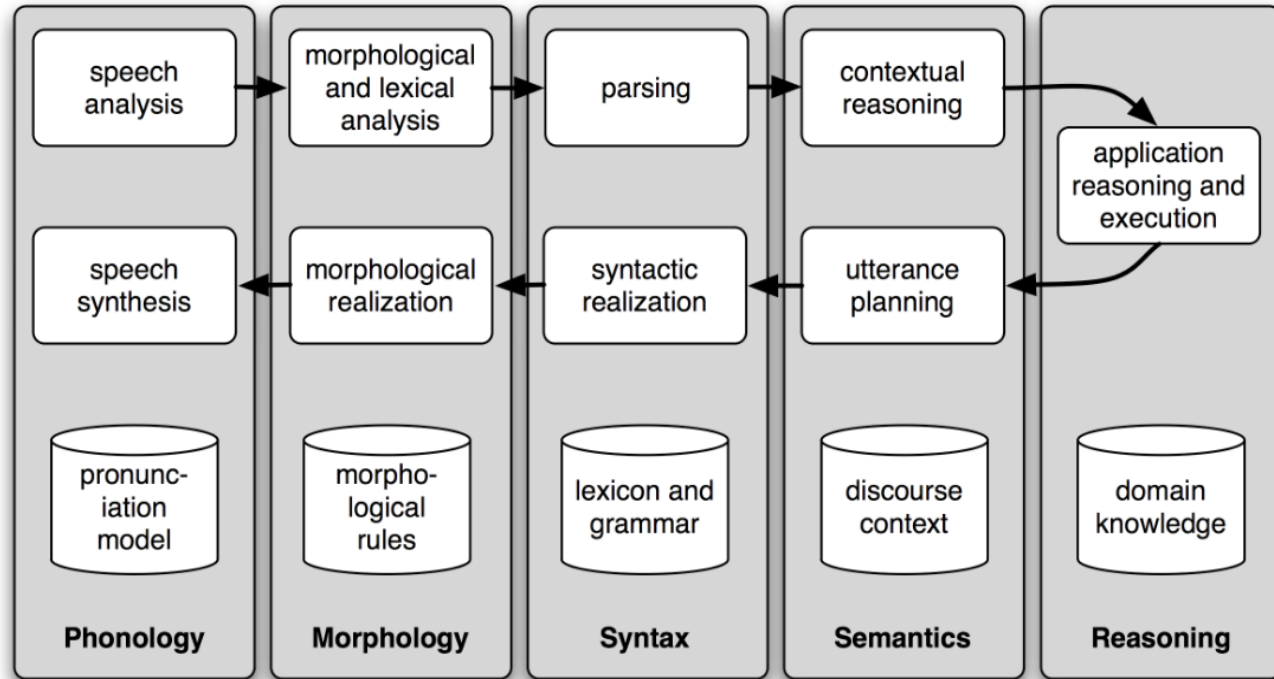


- Options over Markov Decision Process (MDP)
 - Given state s , select action/option a according to (hierarchical) policy π
 - Receive reward r , observe new state a'
 - Continue the cycle until the episode terminates.
- Goal of dialogue learning: find optimal π to maximize expected rewards

A unified view: dialogue as optimal decision making

Dialogue	State (s)	Action (a)	Reward (r)
Q&A bot over KB, Web etc.	Understanding of user Intent (belief state)	Clarification questions, Answers	Relevance of answer # of turns
Task Completion Bots (Movies, Restaurants, ...)	Understanding of user goal (belief state)	Dialog act + slot_value	Task success rate # of turns
Social Bot (Xiaolce)	Conversation history	Response	User engagement
Top-level bot	Understanding of user top-level intent	Options (skills)	Daily/monthly usage

Traditional NLP component stack



1. **Natural language understand (NLU):** parsing (speech) input to semantic meaning and update the system state
2. **Application reasoning and execution:** take the next action based on state
3. **Natural language generation (NLG):** generating (speech) response from action

Symbolic Space

- **Knowledge Representation**
 - *explicitly* stored as words, relations, templates
 - high-dim, discrete, sparse vectors
- **Inference**
 - *slow* on a big knowledge base
 - keyword matching is sensitive to paraphrase alternations
- **Human comprehensible but not computationally efficient**

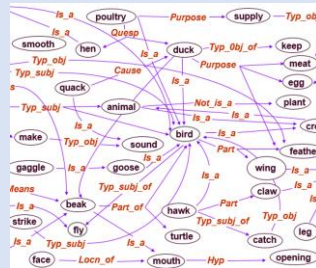
Squire Yrelowney, Dr. Livesey, and the rest of these gentlemen having asked me to write down the whole particulars about Treasure Island, from the beginning to the end, keeping nothing back but the bearings of the island, and that only because there is still treasure not yet lifted, I take up my pen in the year of grace 17— and go back to the time when my father kept the Admiral Benbow inn and the brown old seaman with the sabre cut first took up his lodging under our roof.

I remember him as if it were yesterday, as he came plodding to the inn door, his sea-chest following behind him in a hand-barrow; a tall, strong, heavy, nut-brown man, his tarry pigtail falling over the shoulder of his soiled blue coat, his hands ragged and scarred, with black, broken

nails, and the sabre cut across one cheek, a dirty, livid white. I remember him looking round the cover and whispering to himself as he did so, and then breaking out in that old sea-song that he sang so often afterwards:

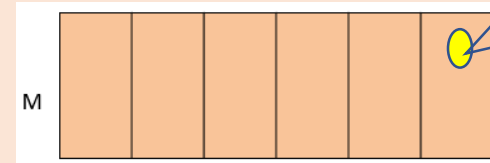
Fifteen men on the dead man's chest—Yo-ho-ho, and a bottle of rum! In the high, old tottering voice that seemed to have been tuned and broken at the capstan bars. Then he rapped on the door with a bit of stick like a handspike that he carried, and when my father appeared, called roughly for a glass of rum. This, when it was brought to him, he drank slowly, like a connoisseur, lingering on the taste and still looking about him at the cliffs and up at our signboard.

"This is a handy cove," says he at length; "and a pleasant situated



Neural Space

- **Knowledge Representation**
 - *Implicitly* stored as semantic concepts/classes
 - low-dim, cont., dense vectors
- **Inference**
 - *fast* on compact memory
 - semantic matching is **robust** to paraphrase alternations
- **Computationally efficient but not human comprehensible yet**



“film”, “award”
 film-genre/films-in-this-genre
 film/cinematography
 cinematographer/film
 award-honor/honored-for
 netflix-title/netflix-genres
 director/film
 award-honor/honored-for

From symbolic to neural computation

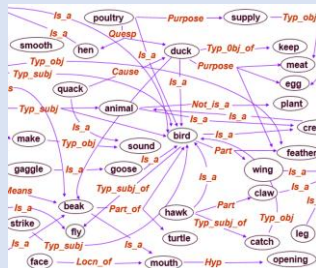
Input: Q

Symbolic \rightarrow Neural
by **Encoding** (Q/D/Knowledge)

Symbolic Space
- human readable

Squire Trelawney, Dr. Livesey, and the rest of these gentlemen having asked me to write down the whole particulars about Treasure Island, from the beginning to the end, keeping nothing back but the bearings of the island, and that only because there is still treasure not yet lifted, I take up my pen in the year of grace 17— and go back to the time when my father kept the Admiral Benbow inn and the brown old seaman with the sabre cut first took up his lodging under our roof. I remember him as if it were yesterday, as he came plodding to the inn door, his sea-chest following behind him in a hand-barrow; a tall, strong, heavy, nut-brown man, his tarry pigtail falling over the shoulder of his soiled blue coat, his hands ragged and scarred, with black, broken

nails, and the sabre cut across one cheek, a dirty, livid white. I remember him looking round the cover and whistling to himself as he did so, and then breaking out in that old sea-song that he sang so often afterwards: "Fifteen men on the dead man's chest Yo-ho-ho, and a bottle of rum" in the high, old tottering voice that seemed to have been tuned and broken at the captain bars. Then he rapped on the door with a bit of stick like a handspike that he carried, and when my father appeared, called roughly for a glass of rum. This, when it was brought to him, he drank slowly, like a connoisseur, lingering on the taste and still looking about him at the cliffs and up at our signboard. "This is a handy cove," says he at length; "and a pleasant sittayed



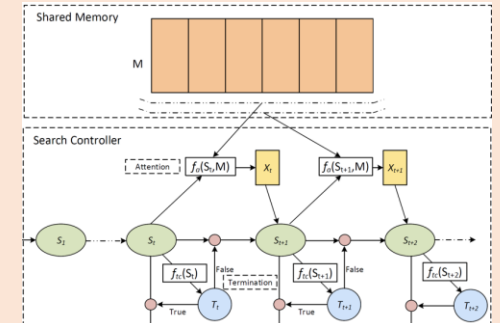
Error(A, A*)

Output: A

Neural \rightarrow Symbolic
by **Decoding** (synthesizing answer)

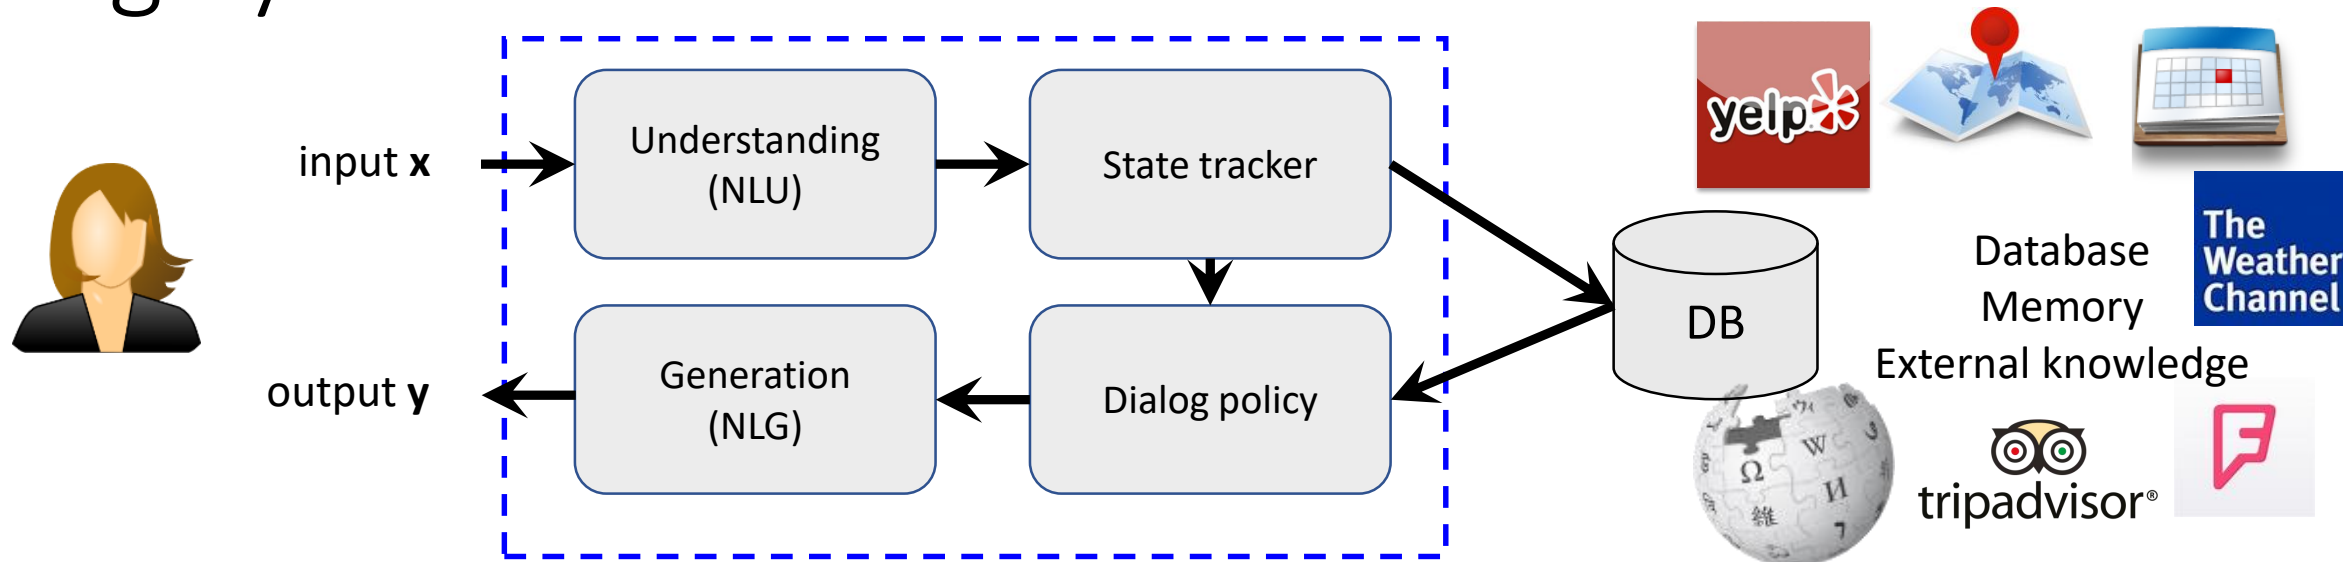
Reasoning: Question + KB \rightarrow answer
vector via multi-step inference,
summarization, deduction etc.

Neural Space
- Computationally efficient

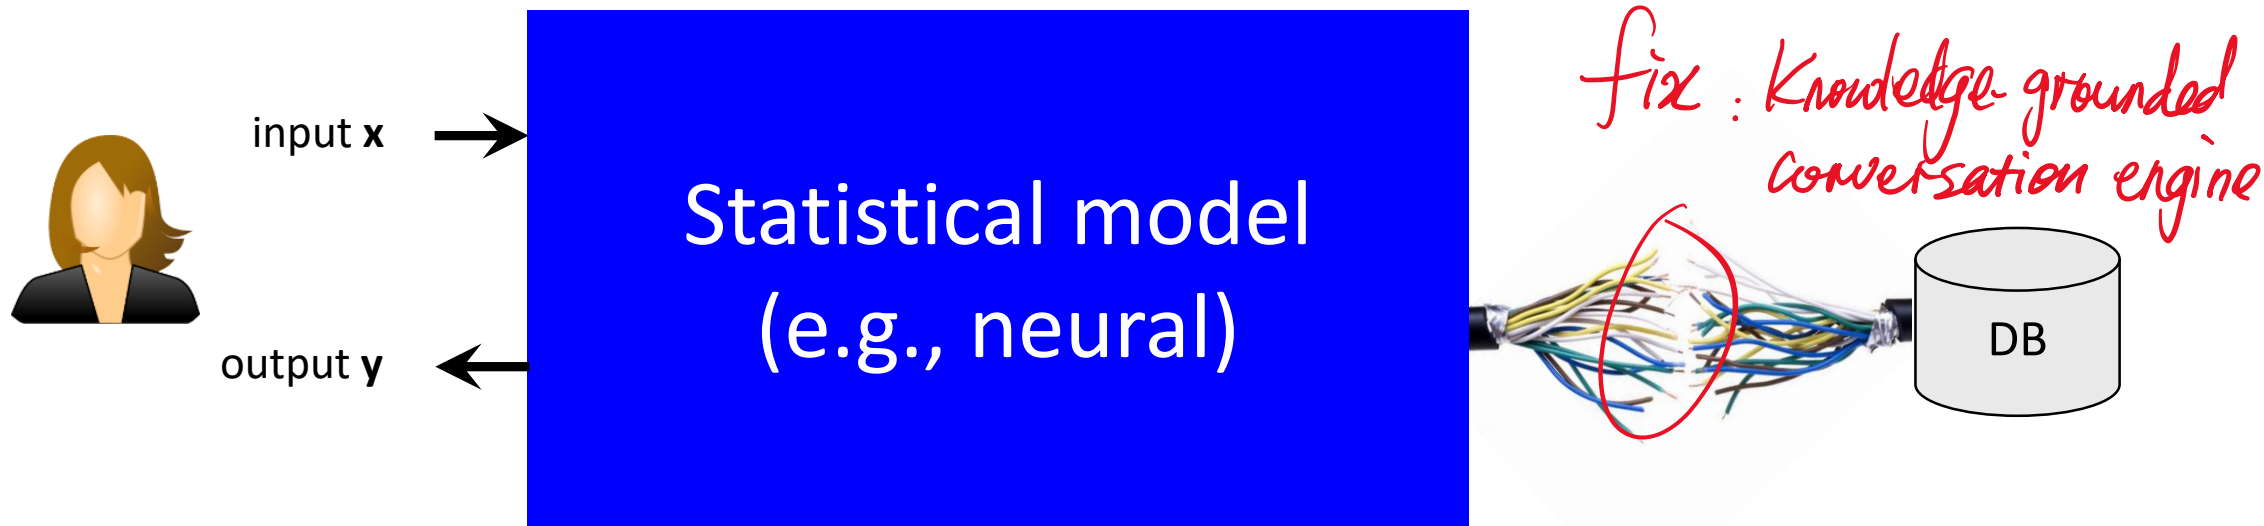


Dialog Systems

Task-Oriented Dialog



Fully data-driven



Outline

- Part 1: Introduction
- **Part 2: Question answering (QA) and machine reading comprehension (MRC)**
 - **Neural MRC models for text-based QA**
 - **Multi-turn knowledge base QA agents**
- Part 3: Task-oriented dialogue
- Part 4: Fully data-driven conversation models and social chatbots

Text-QA

Q Will I qualify for OSAP if I'm new in Canada?

Selected Passages from Bing

"Visit the OSAP website for application deadlines. To get OSAP, you have to be eligible. You can apply using an online form, or you can print off the application forms. If you submit a paper application, you must pay an application fee. The online application is free."

Source: <http://settlement.org/ontario/education/colleges-universities-and-institutes/financial-assistance-for-post-secondary-education/how-do-i-apply-for-the-ontario-student-assistance-program-osap/>

"To be eligible to apply for financial assistance from the Ontario Student Assistance Program (OSAP), you must be a: 1 Canadian citizen; 2 Permanent resident; or 3 Protected person/convention refugee with a Protected Persons Status Document (PPSD)."

Source: <http://settlement.org/ontario/education/colleges-universities-and-institutes/financial-assistance-for-post-secondary-education/who-is-eligible-for-the-ontario-student-assistance-program-osap/>

"You will not be eligible for a Canada-Ontario Integrated Student Loan, but can apply for a part-time loan through the Canada Student Loans program. There are also grants, bursaries and scholarships available for both full-time and part-time students."

Source: <http://www.campusaccess.com/financial-aid/osap.html>

Answer

No. You won't qualify.

In meteorology, precipitation is any product of the condensation of atmospheric water vapor that falls under **gravity**. The main forms of precipitation include drizzle, rain, sleet, snow, **graupel** and hail... Precipitation forms as smaller droplets coalesce via collision with other rain drops or ice crystals **within a cloud**. Short, intense periods of rain in scattered locations are called "showers".

What causes precipitation to fall?

gravity

What is another main form of precipitation besides drizzle, rain, snow, sleet and hail?

graupel

Where do water droplets collide with ice crystals to form precipitation?

within a cloud

Figure 1: Question-answer pairs for a sample passage in the SQuAD dataset. Each of the answers is a segment of text from the passage.

From symbolic to neural computation

Input: Q

Symbolic \rightarrow Neural
by **Encoding** (Q/D/Knowledge)

Symbolic Space
- human readable

Where did Tesla live for much of his life?

Tesla was renowned for his achievements and showmanship, eventually earning him a reputation in popular culture as an archetypal "mad scientist". His patents earned him a considerable amount of money, much of which was used to finance his own projects with varying degrees of success. He lived most of his life in a series of New York hotels, through his retirement. Tesla died on 7 January 1943. His work fell into relative obscurity after his death, but in 1960 the General Conference on Weights and Measures named the SI unit of magnetic flux density the tesla in his honor. There has been a resurgence in popular interest in Tesla since the 1990s.

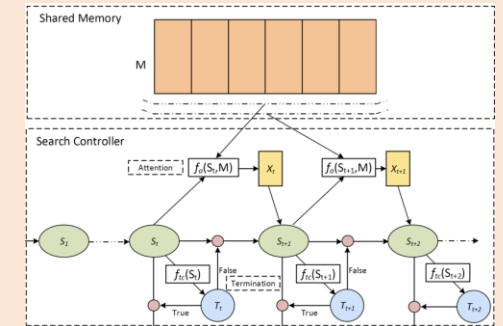
Error(A, A*)

Output: A

Neural \rightarrow Symbolic
by **Decoding** (synthesizing answer)

Reasoning: Question + KB \rightarrow answer
vector via multi-step inference,
summarization, deduction etc.

Neural Space
- Computationally efficient



Neural MRC models by SQuAD

What types of European groups were able to avoid the plague?

From Italy, the disease spread northwest across Europe, striking France, Spain, Portugal and England by June 1348, then turned and spread east through Germany and Scandinavia from 1348 to 1350. It was introduced in Norway in 1349 when a ship landed at Askøy, then spread to Bjørgvin (modern Bergen) and Iceland. Finally it spread to northwestern Russia in 1351. The plague was somewhat less common in parts of Europe that had smaller trade relations with their neighbours, including the Kingdom of Poland, the majority of the Basque Country, isolated parts of Belgium and the Netherlands, and isolated alpine villages throughout the continent.

A limited form of comprehension:

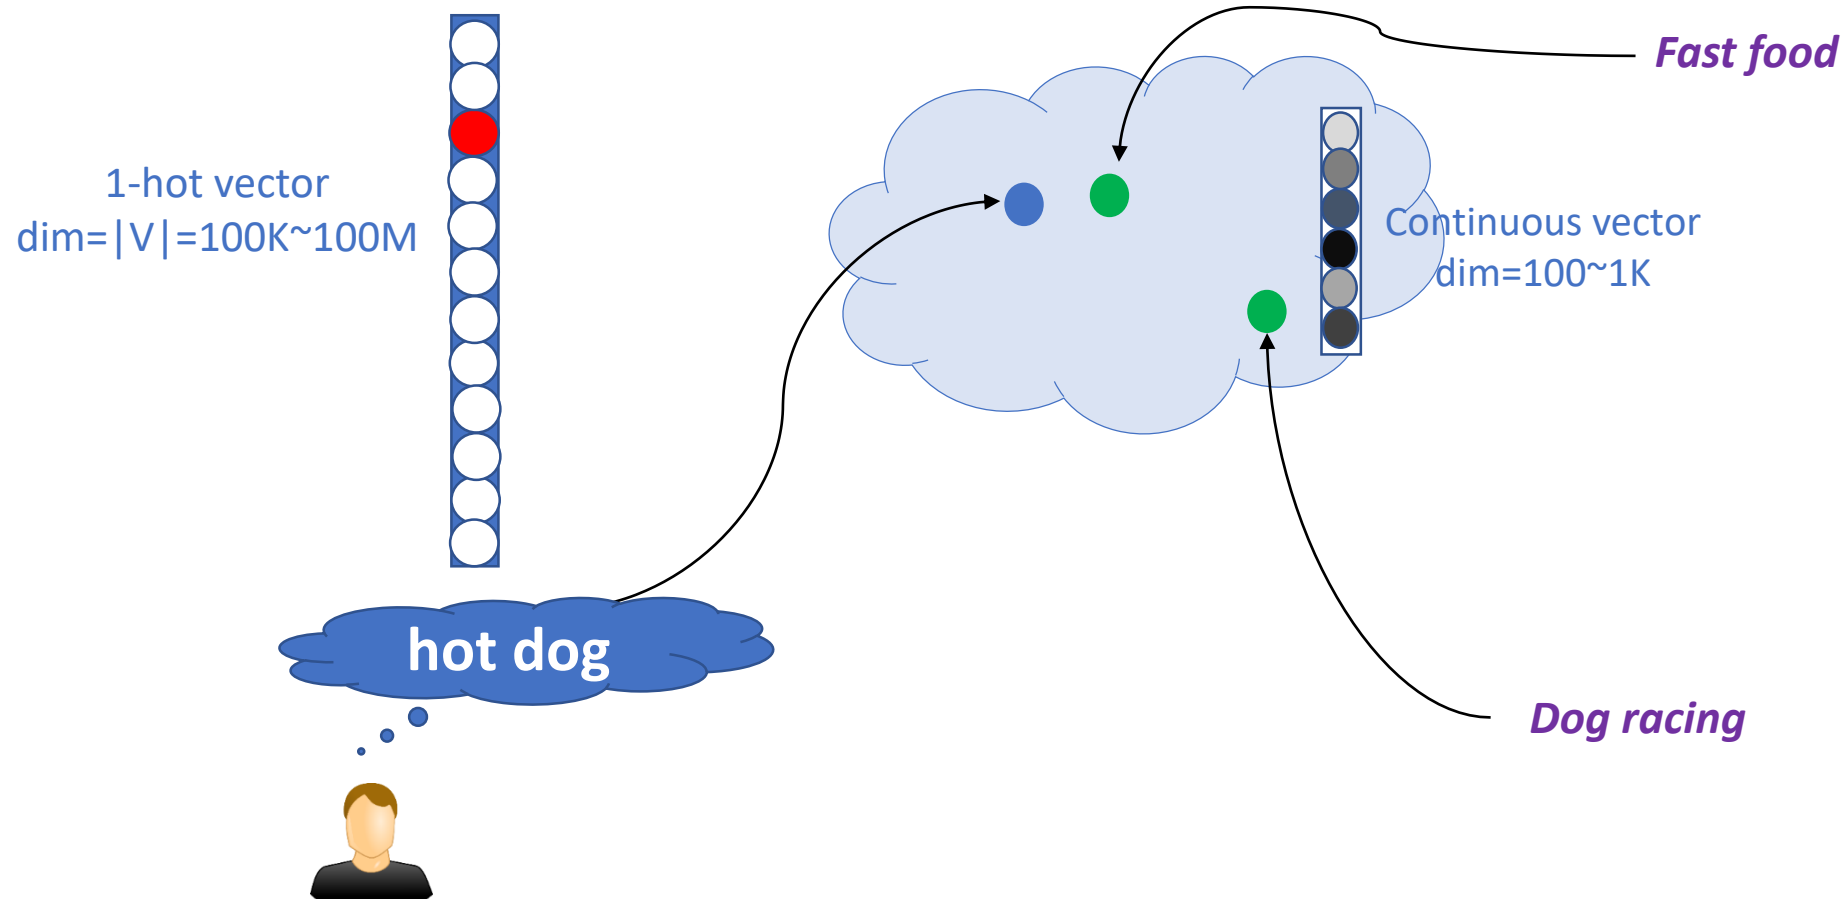
- No need for extra knowledge outside the paragraph
- No need for clarifying questions
- The answer must exist in the paragraph
- The answer must be a text span, not synthesized

- Encoding: map each text span to a semantic vector
- Reasoning: rank and re-rank semantic vectors
- Decoding: map the top-ranked vector to text

Three encoding components

- Word embedding – word semantic space
 - represent each word as a low-dim continuous vector via [GloVe](#) [[Pennington+ 14](#)]
- **Context embedding – contextual semantic space**
 - **capture context info for each word, via**
 - [BiLSTM](#) [[Melamud+ 16](#)]: concatenation of left and right contextual language models
 - [ELMo](#) [[Peter+ 18](#)]: a task-specific combo of the intermediate layer representations of biLM
 - [BERT](#) [[Devlin+ 18](#)]: bidirectional transformer, jointly conditioned on left and right context
- Context-query attention – query dependent semantic space
 - fuse query info into passage via [Attention](#)
 - [[Huang+ 17](#); [Wang+ 17](#); [Hu+ 17](#); [Seo+ 16](#); [Wang&Jiang 16](#)]

Word embedding: word semantic space



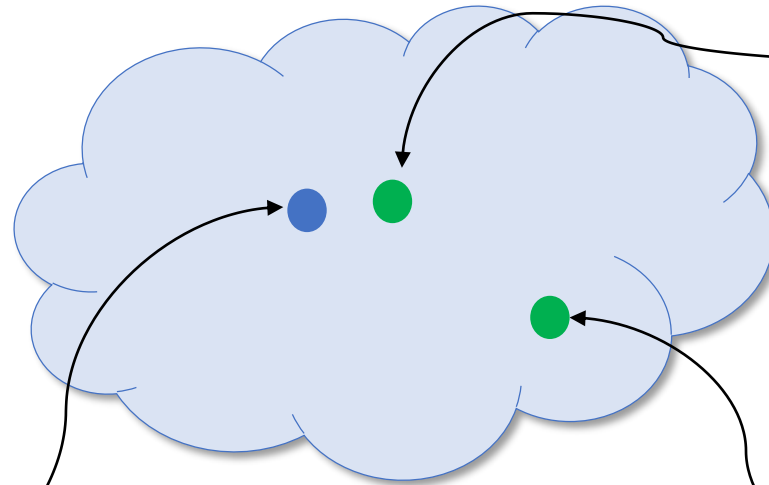
Context embedding: contextual semantic space

The Einstein Theory of Relativity


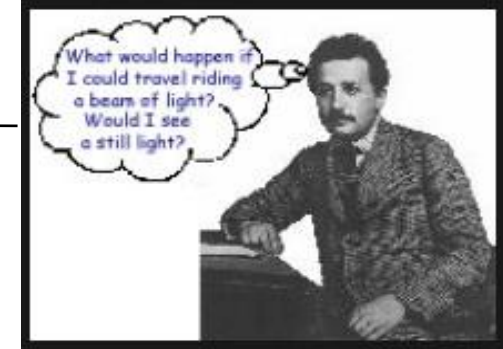
(1) The perihelion of Mercury shows a discrepancy which has long puzzled astronomers. This discrepancy is fully accounted for by Einstein. At the time when he published his theory, this was its only experimental verification.

(2) Modern physicists were willing to suppose that light might be subject to gravitation—i.e., that a ray of light passing near a great mass like the sun might be deflected to the extent to which a particle moving with the same velocity would be deflected according to the orthodox theory of gravitation. But Einstein's theory required that the light should be deflected just twice as much as this. The matter could only be tested during an eclipse among a number of bright stars. Fortunately a peculiarly favourable eclipse occurred last year. The results of the observations

ray of light



Ray of Light (Experiment)



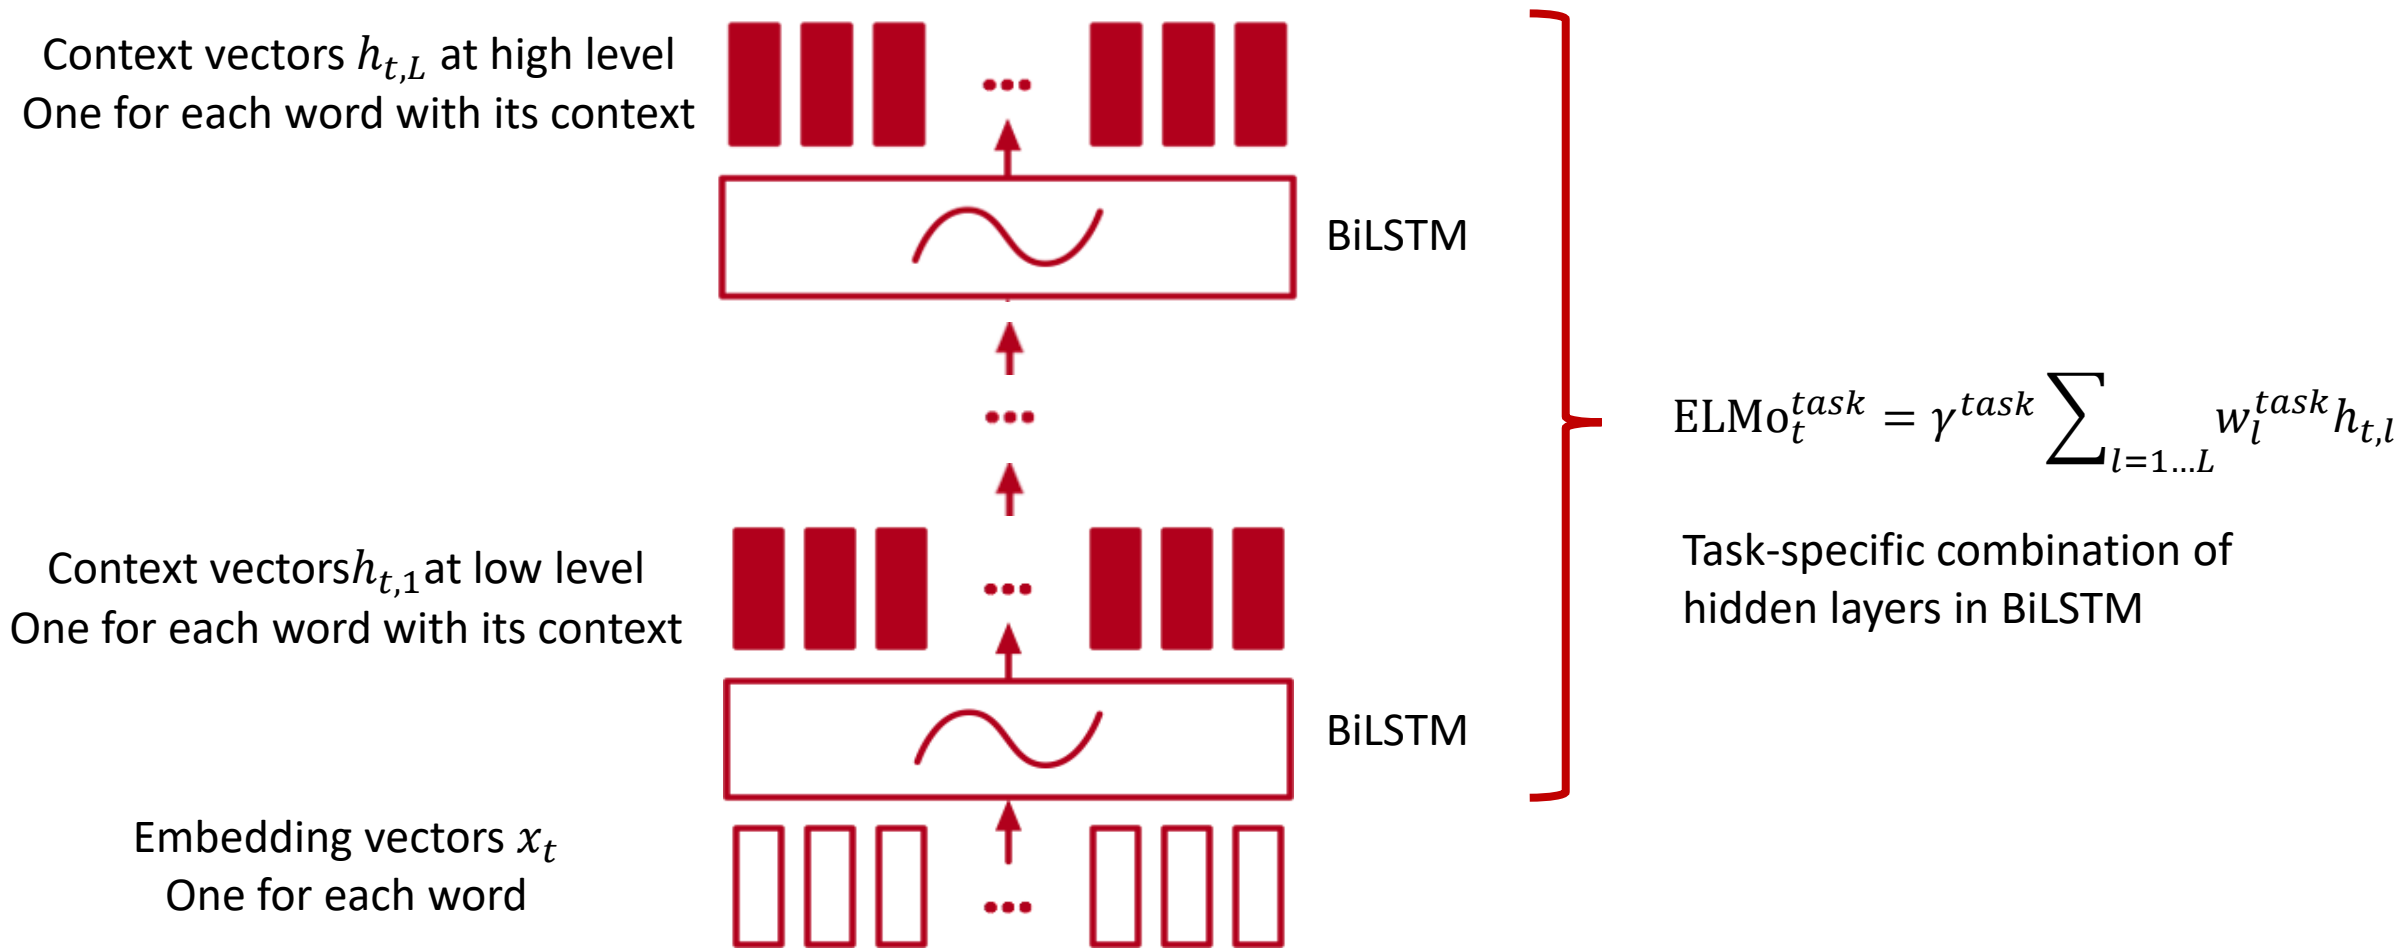
Ray of Light is the seventh studio album by American singer-songwriter Madonna, released on March 3, 1998 by Maverick Records. After giving birth to her daughter Lourdes, Madonna started working on her new album with producers Babyface, Patrick Leonard an...

Release date	Mar 3, 1998
Artist	Madonna
Awards	Grammy Award for B...

[See More](#)

Ray of Light (Song)

Context embedding via BiLSTM/ELMo/GPT/ BERT



Context embedding via BiLSTM/ELMo/GPT/ BERT

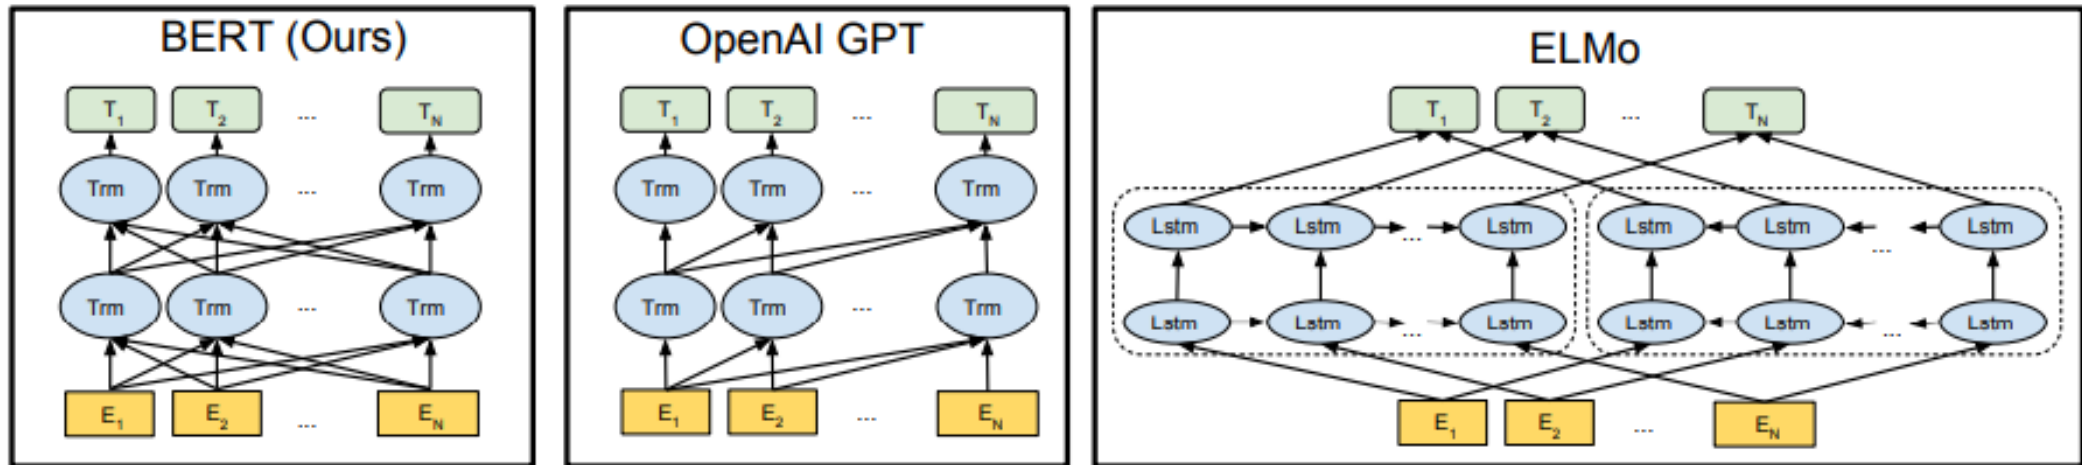
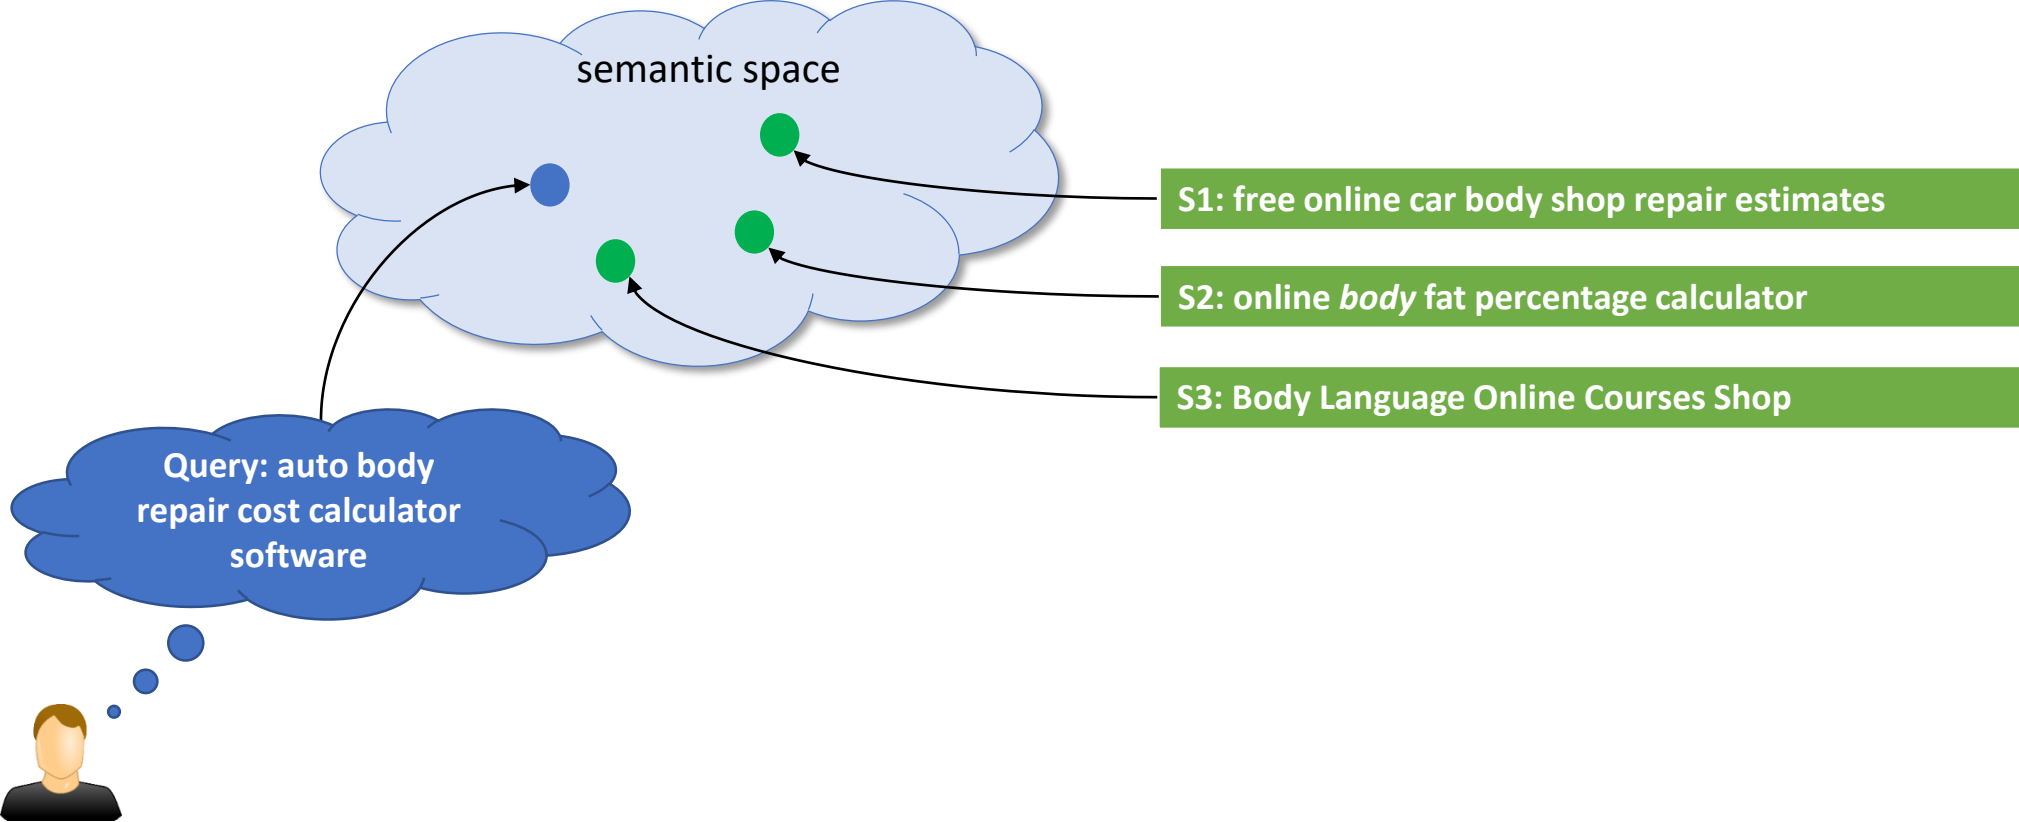
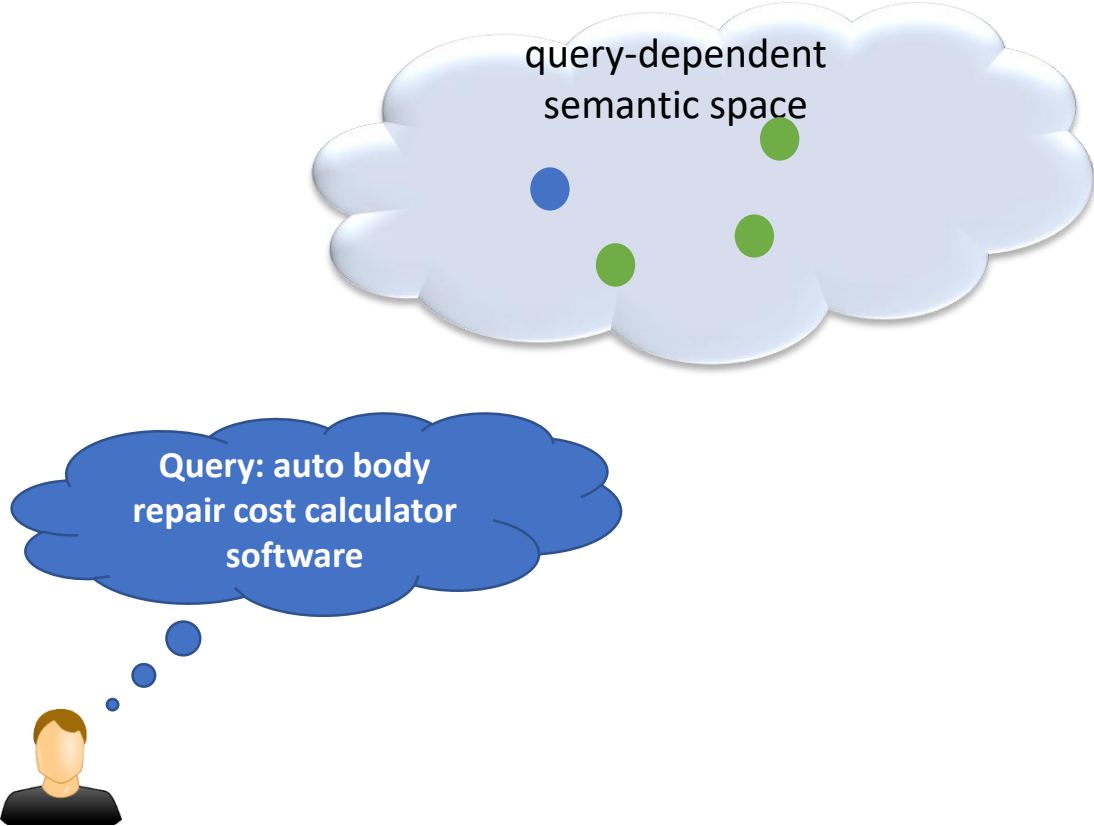


Figure 1: Differences in pre-training model architectures. BERT uses a bidirectional Transformer. OpenAI GPT uses a left-to-right Transformer. ELMo uses the concatenation of independently trained left-to-right and right-to-left LSTM to generate features for downstream tasks. Among three, only BERT representations are jointly conditioned on both left and right context in all layers.

Context-query attention: query-dependent semantic space



Context-query attention: query-dependent semantic space



S1: free online car body shop repair estimates



S2: online *body* fat percentage calculator

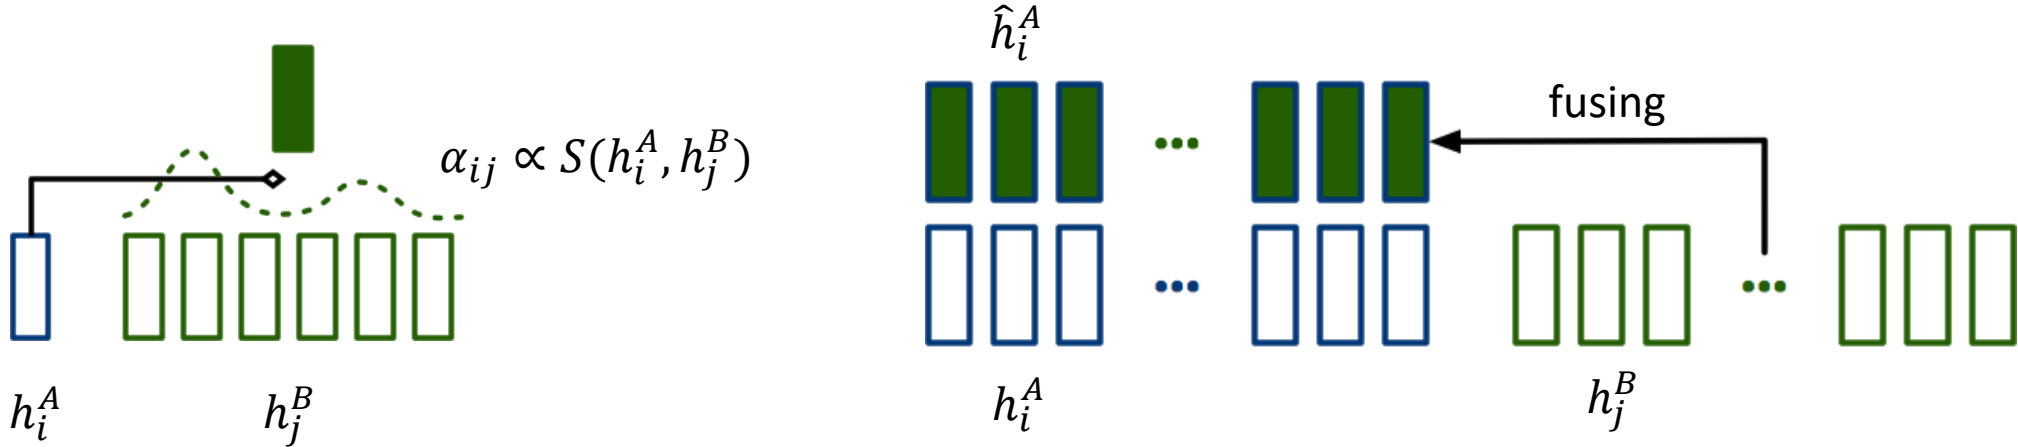


S3: Body Language Online Courses Shop



Context-query attention: query-dependent semantic space

- Compute attention scores (similarity btw i and j): $S_{ij} = S(h_i^A, h_j^B)$
- Compute attention weights thru softmax: $\alpha_{ij} = \exp(S_{ij}) / \sum_k \exp(S_{ik})$
- **Fusion** info from B to A: $\hat{h}_i^A = \sum_j \alpha_{ij} h_j^B$



Examples: BiDAF and SAN

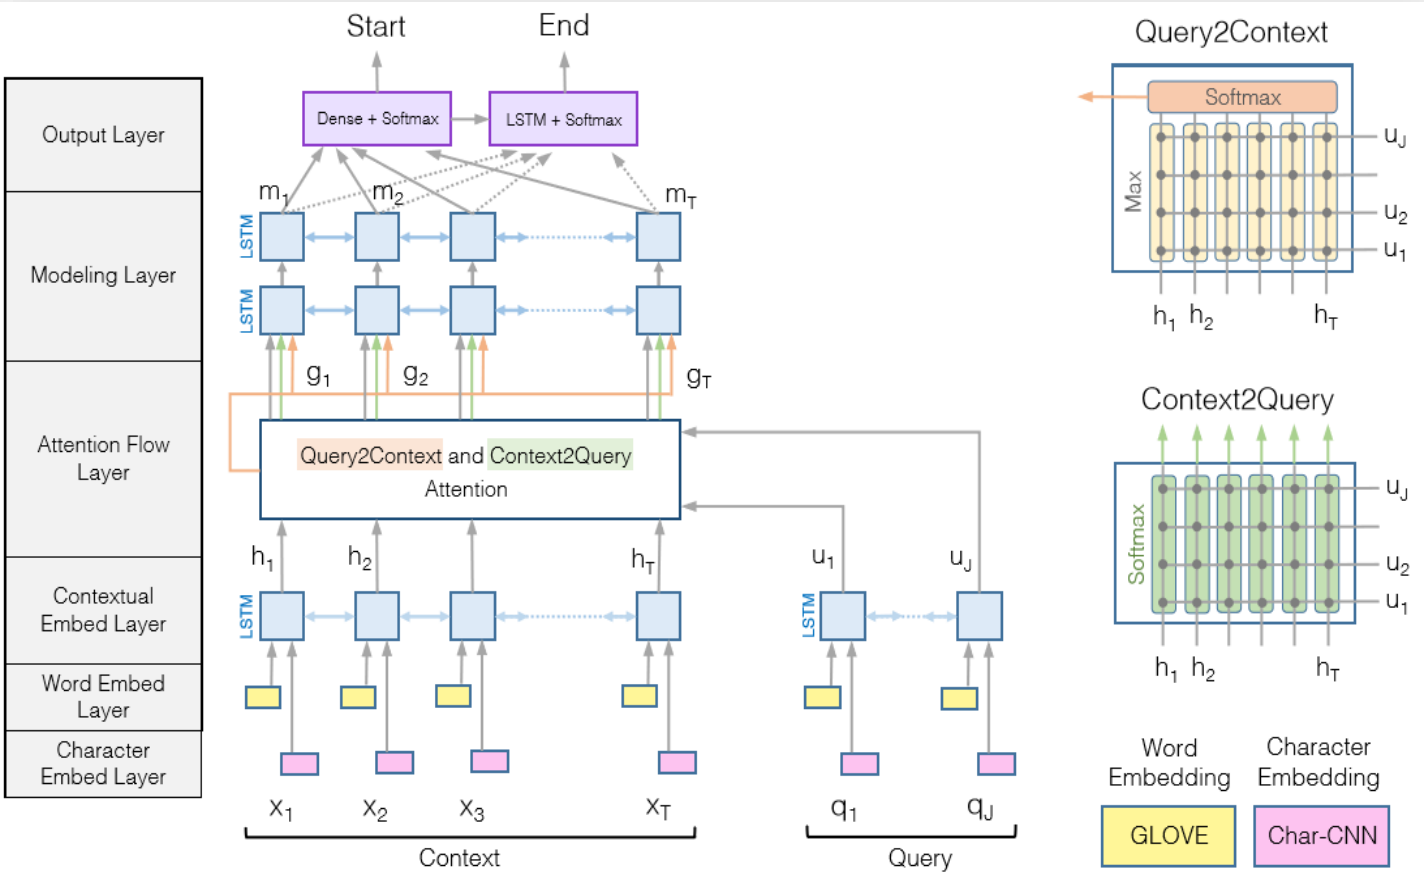
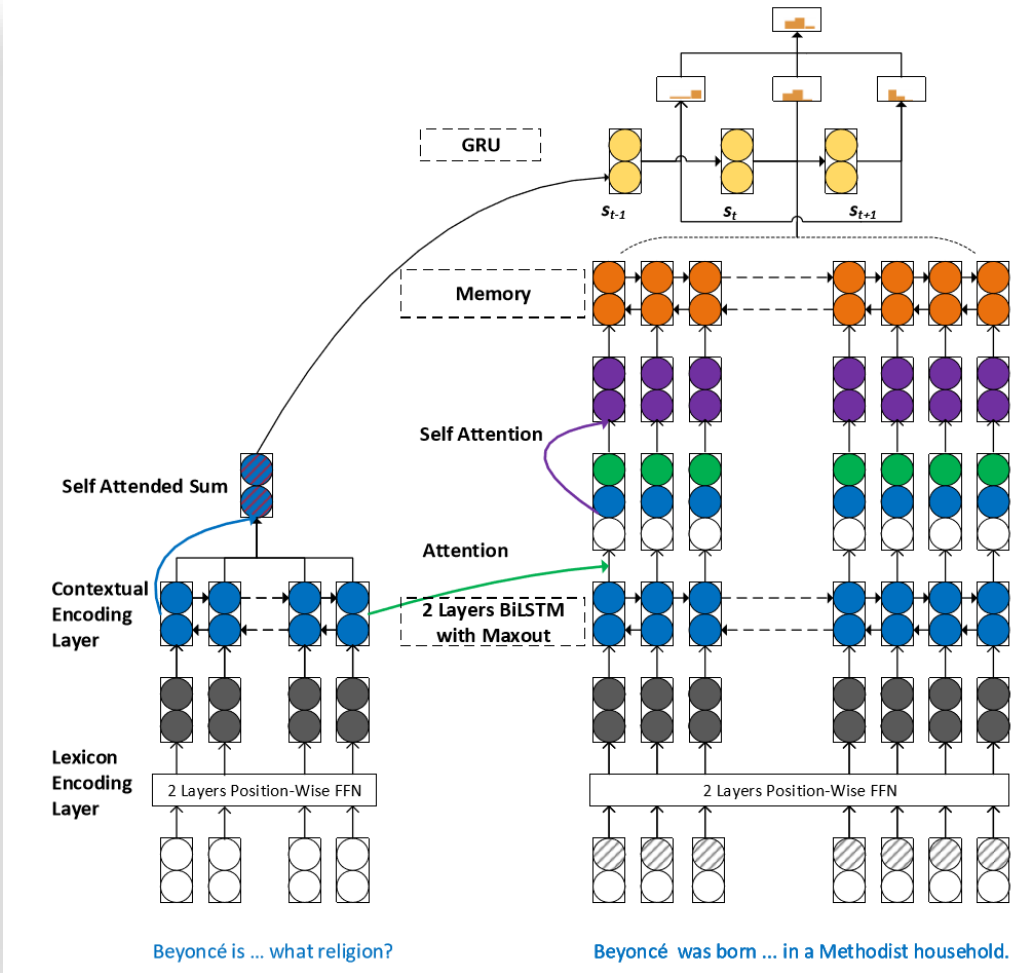


Figure 1: BiDirectional Attention Flow Model (best viewed in color)



Common-sense reasoning: example

Query Who was the #2 pick in the 2011 NFL Draft?

Passage Manning was the #1 selection of the 1998 NFL draft, while Newton was picked first in 2011. The matchup also pits the top two picks of the 2011 draft against each other: Newton for Carolina and Von Miller for Denver.

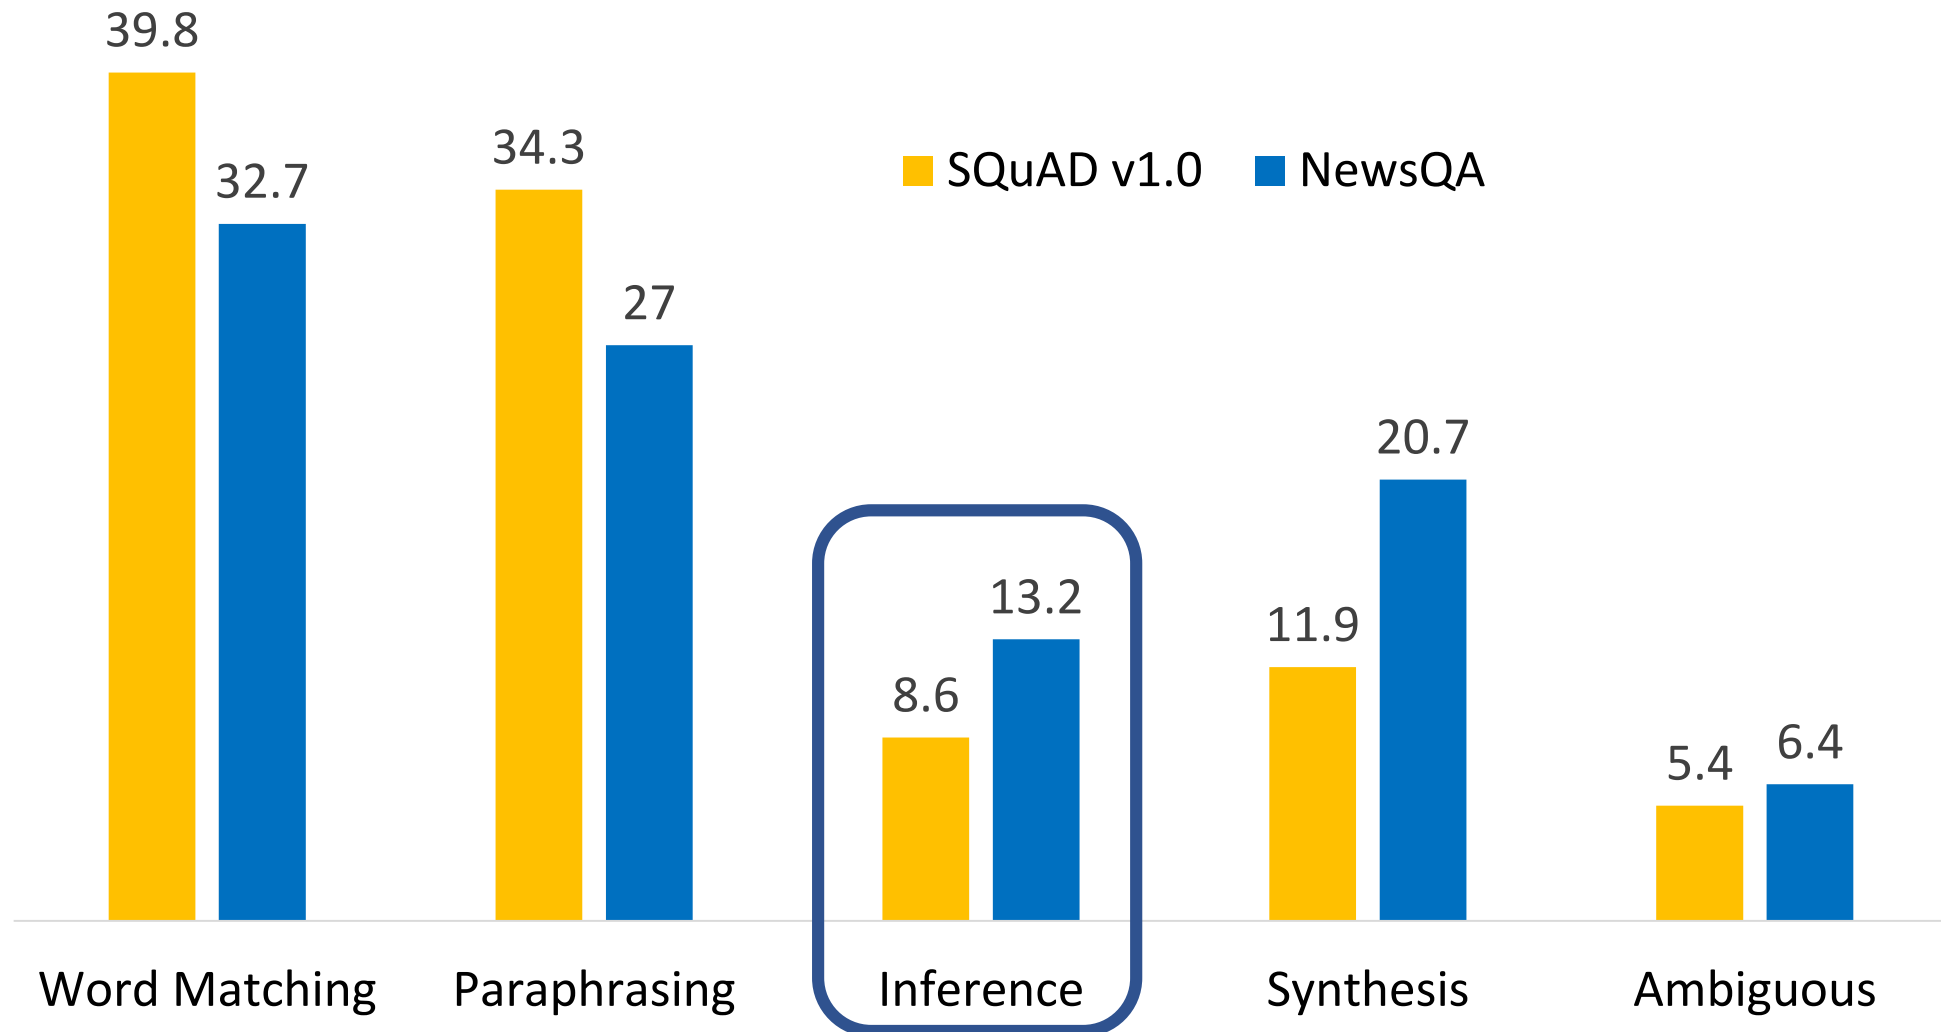
Answer Von Miller

- Step 1:
 - **Extract:** Manning is #1 pick of 1998
 - **Infer:** Manning is NOT the answer
- Step 2:
 - **Extract:** Newton is #1 pick of 2011
 - **Infer:** Newton is NOT the answer
- Step 3:
 - **Extract:** Newton and Von Miller are top 2 picks of 2011
 - **Infer:** Von Miller is the #2 pick of 2011

Commonsense knowledge:

A and B are candidates, if A is not the answer, then the answer is B

Weak Reasoning in existing MRC tasks



Passage/Query from Summary/Details [Zheng+ 2018]

CNN U.S. | World | Politics | Money | Opinion | Health | Entertainment | Tech | Style | Travel | Sports | Video | V

Story highlights

LZ Granderson: For years, the economy has been seen as the dominant issue

He asks if Obama's support for gay marriage will change the Romney camp's direction

LZ: Obama showed he is no coward; he didn't have to address the issue before election

For the better part of four years, voters have said the No. 1 issue is the economy.

For the past year, the economy has been the GOP's primary point of attack against President Obama.

For months, the economy and job creation has been Mitt Romney's calling card.

Even in the one area where on paper @Entity has significant credibility -- the economy -- 44% of voters believe the president will do a better job, while 42%



Passage
(Summary)



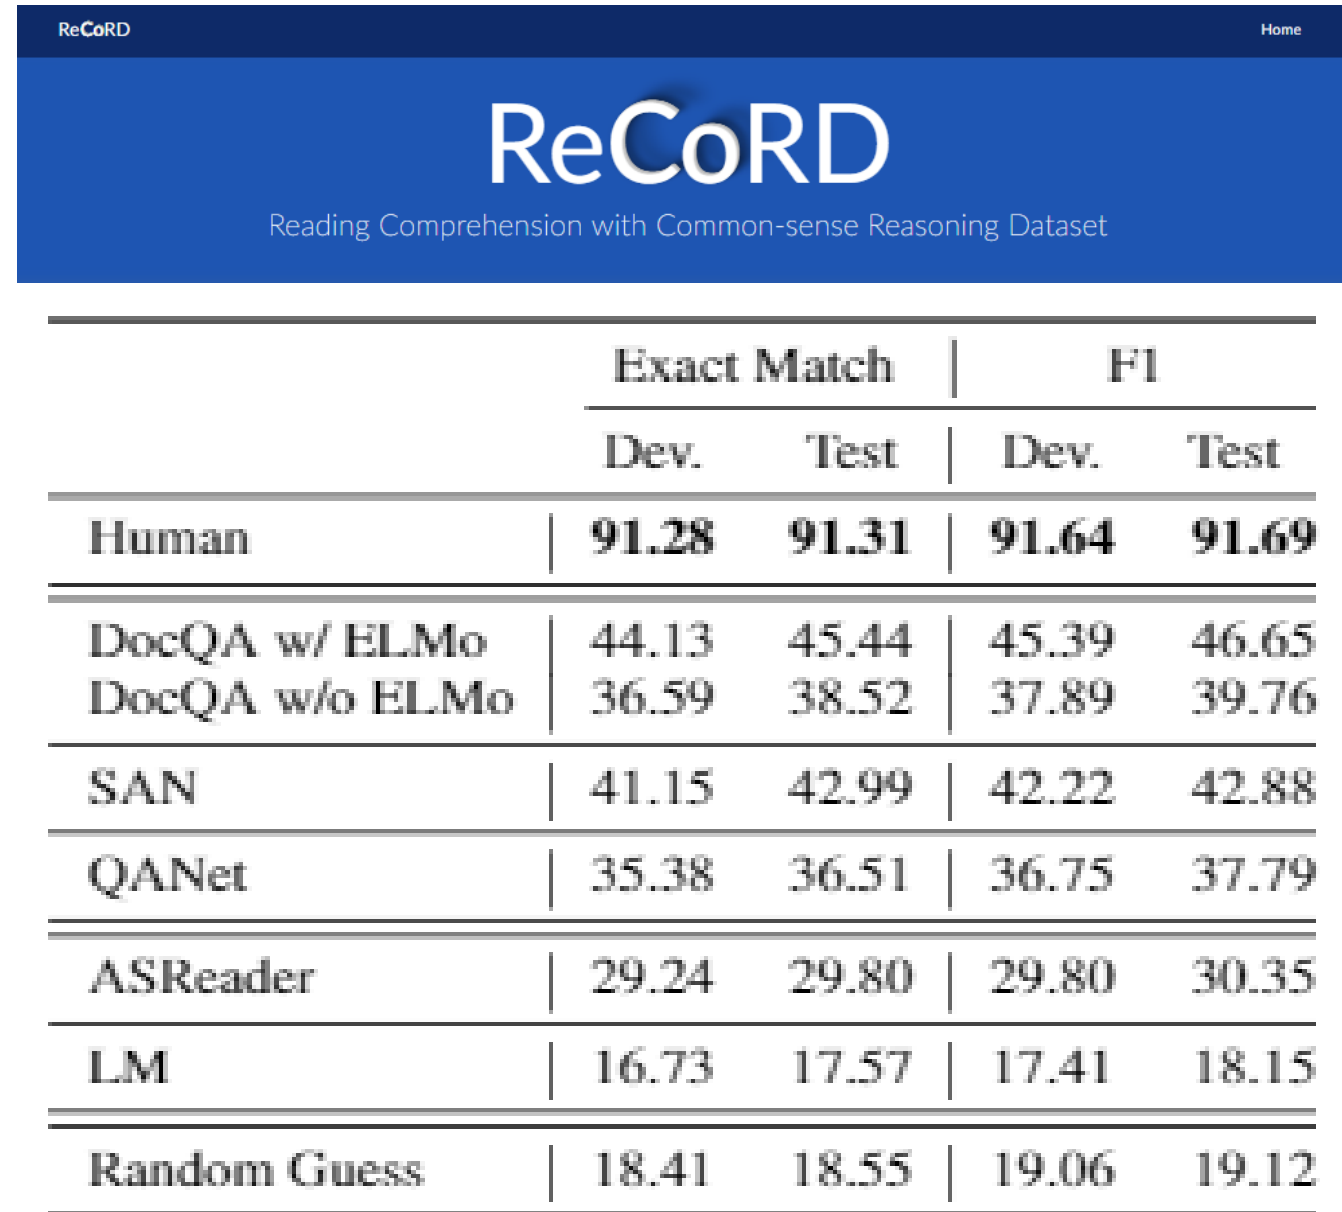
Common-sense Reasoning



Query
(Details)

MSR-JHU Leaderboard

- [ReCoRD](#) (**Reading Comprehension with Common-sense Reasoning Dataset**)
 - From 70K news articles
 - Training set: 100K
 - Dev set: 10K
 - Test set: 10K
 - Validated by model & human judges
- Coming soon...

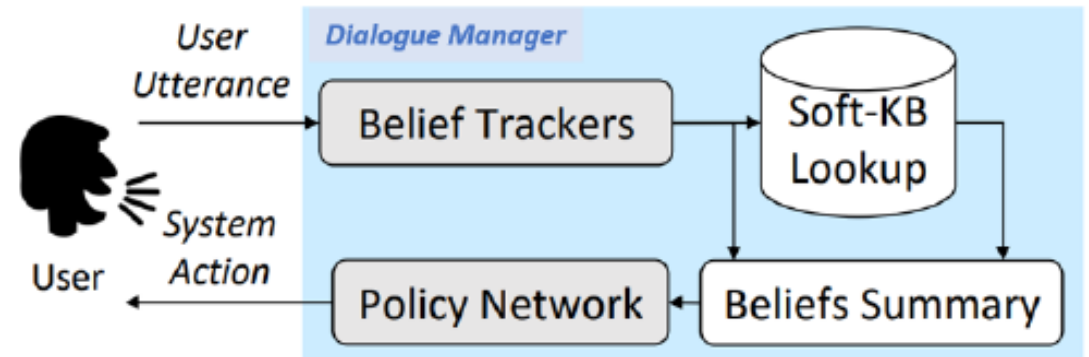
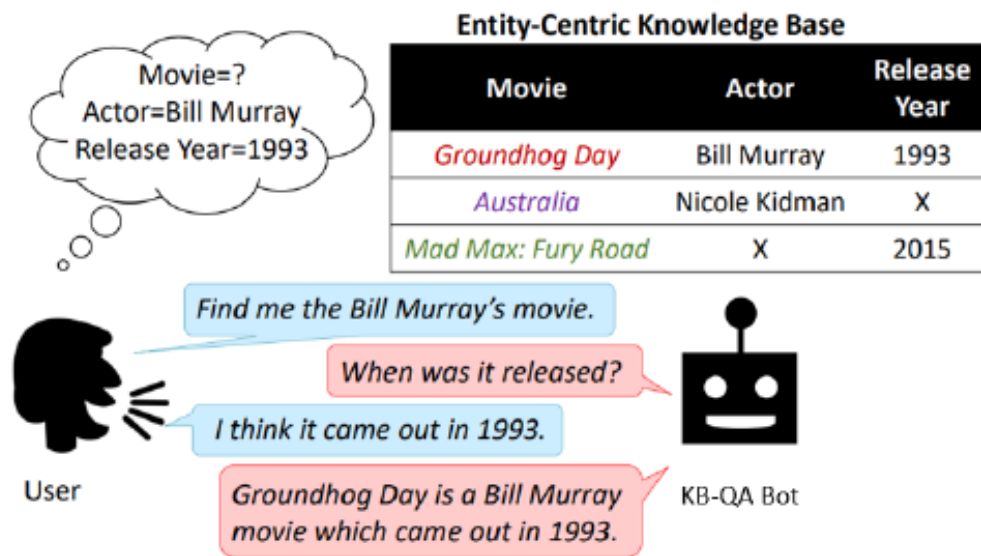


The image shows a screenshot of the ReCoRD website. At the top, there is a dark blue header with the text 'ReCoRD' on the left and 'Home' on the right. Below the header, the main title 'ReCoRD' is displayed in a large, white, serif font, with the subtitle 'Reading Comprehension with Common-sense Reasoning Dataset' underneath it in a smaller, white, sans-serif font. The main content of the screenshot is a table with a white background and black borders. The table has five columns: a model name, 'Exact Match', 'F1', and two sub-columns for 'Dev.' and 'Test' under both 'Exact Match' and 'F1'. The rows are: Human (91.28, 91.31, 91.64, 91.69), DocQA w/ ELMo (44.13, 45.44, 45.39, 46.65), DocQA w/o ELMo (36.59, 38.52, 37.89, 39.76), SAN (41.15, 42.99, 42.22, 42.88), QANet (35.38, 36.51, 36.75, 37.79), ASReader (29.24, 29.80, 29.80, 30.35), LM (16.73, 17.57, 17.41, 18.15), and Random Guess (18.41, 18.55, 19.06, 19.12). The 'Human' row is highlighted with a thicker border.

	Exact Match		F1	
	Dev.	Test	Dev.	Test
Human	91.28	91.31	91.64	91.69
DocQA w/ ELMo	44.13	45.44	45.39	46.65
DocQA w/o ELMo	36.59	38.52	37.89	39.76
SAN	41.15	42.99	42.22	42.88
QANet	35.38	36.51	36.75	37.79
ASReader	29.24	29.80	29.80	30.35
LM	16.73	17.57	17.41	18.15
Random Guess	18.41	18.55	19.06	19.12

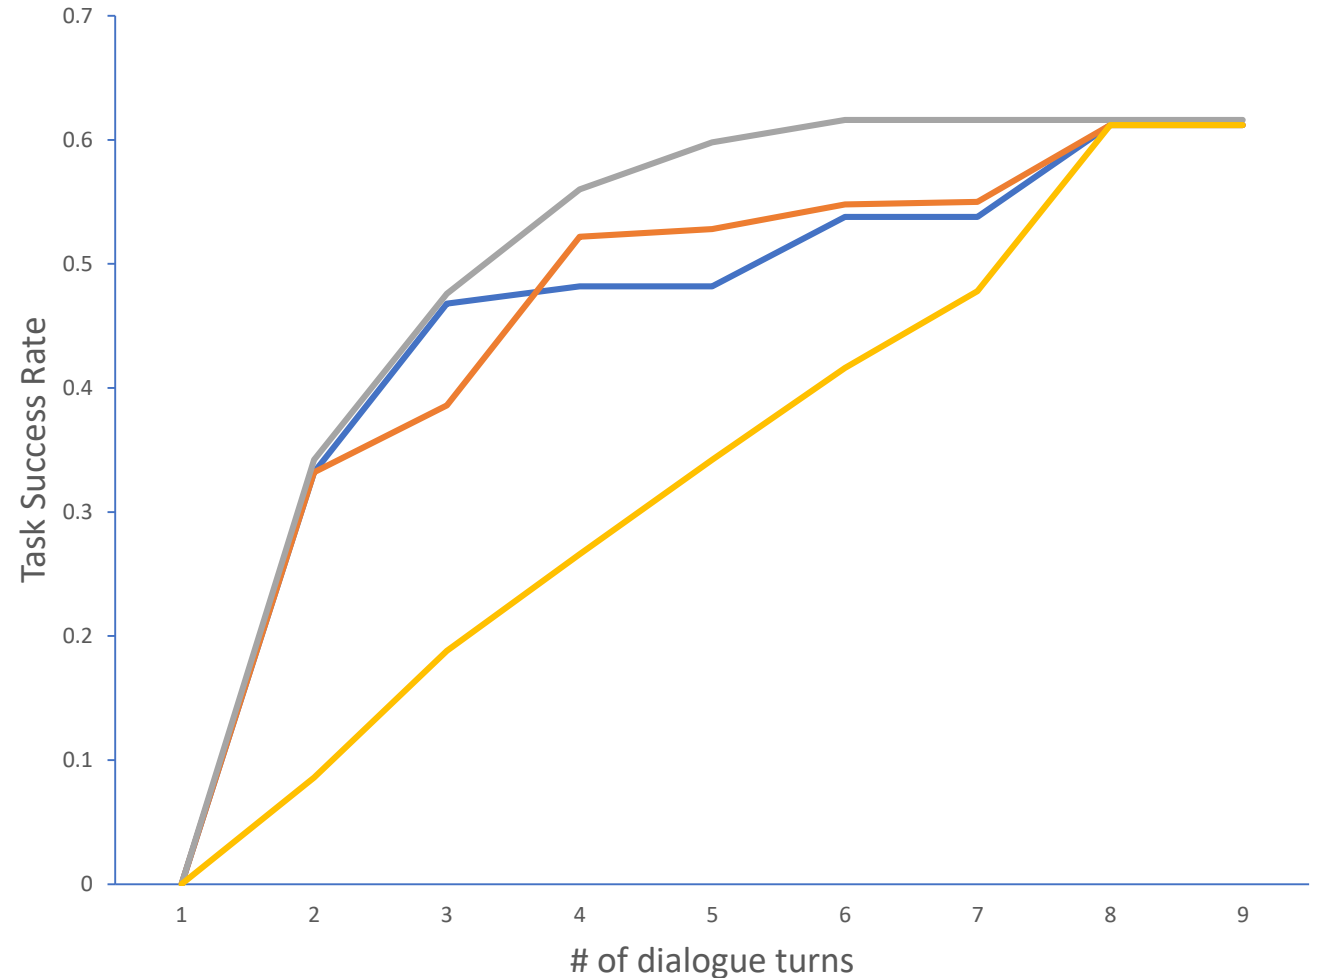
KB-Info Bot: movie on demand [[Dhingra+ 17](#)]

- Turning Bing into a dialogue agent



Learning what to ask next, and when to stop

- Initial: ask all questions in a randomly sampled order
- Improve via learning from Bing log
 - Ask questions that users can answer
- Improve via encoding knowledge of database
 - Ask questions that help reduce search space
- Finetune using agent-user interactions
 - Ask questions that help complete the task successfully via RL

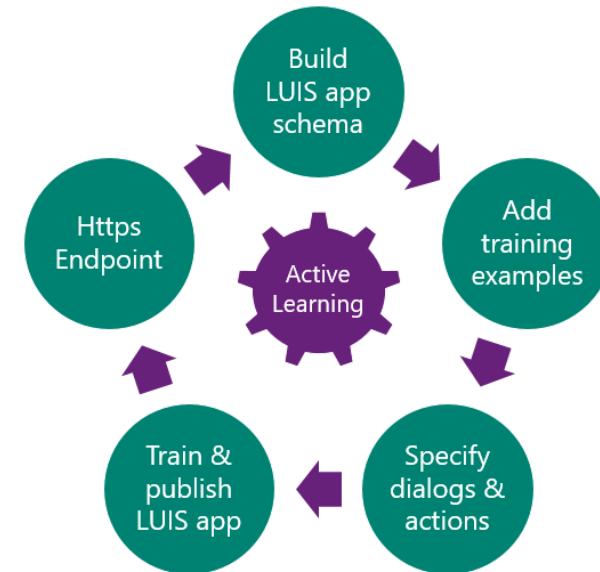
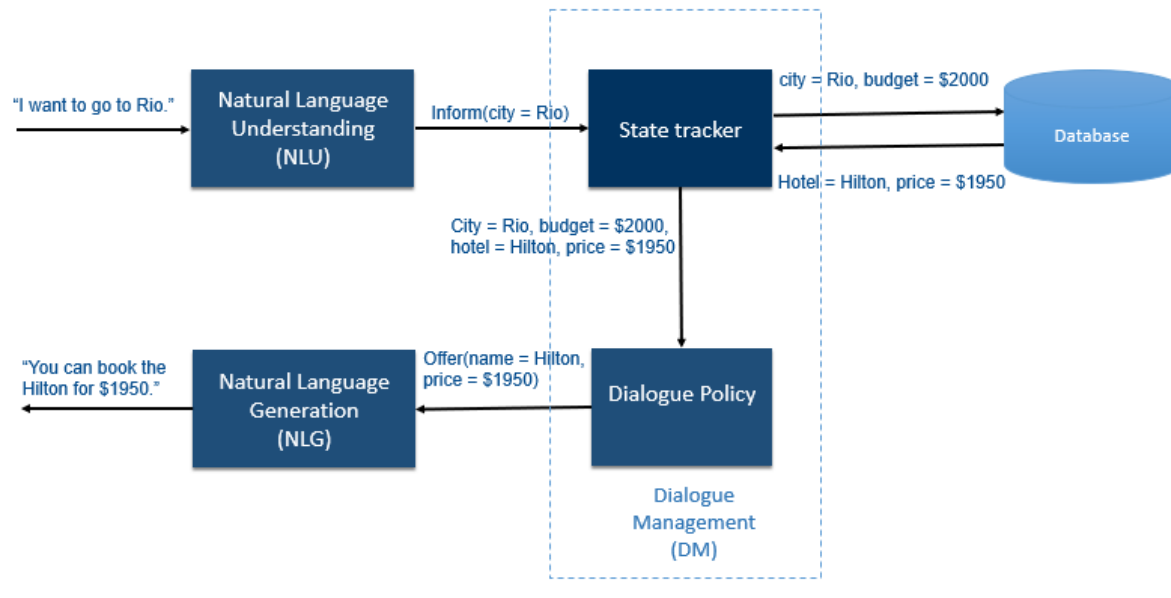


Results on simulated users

Outline

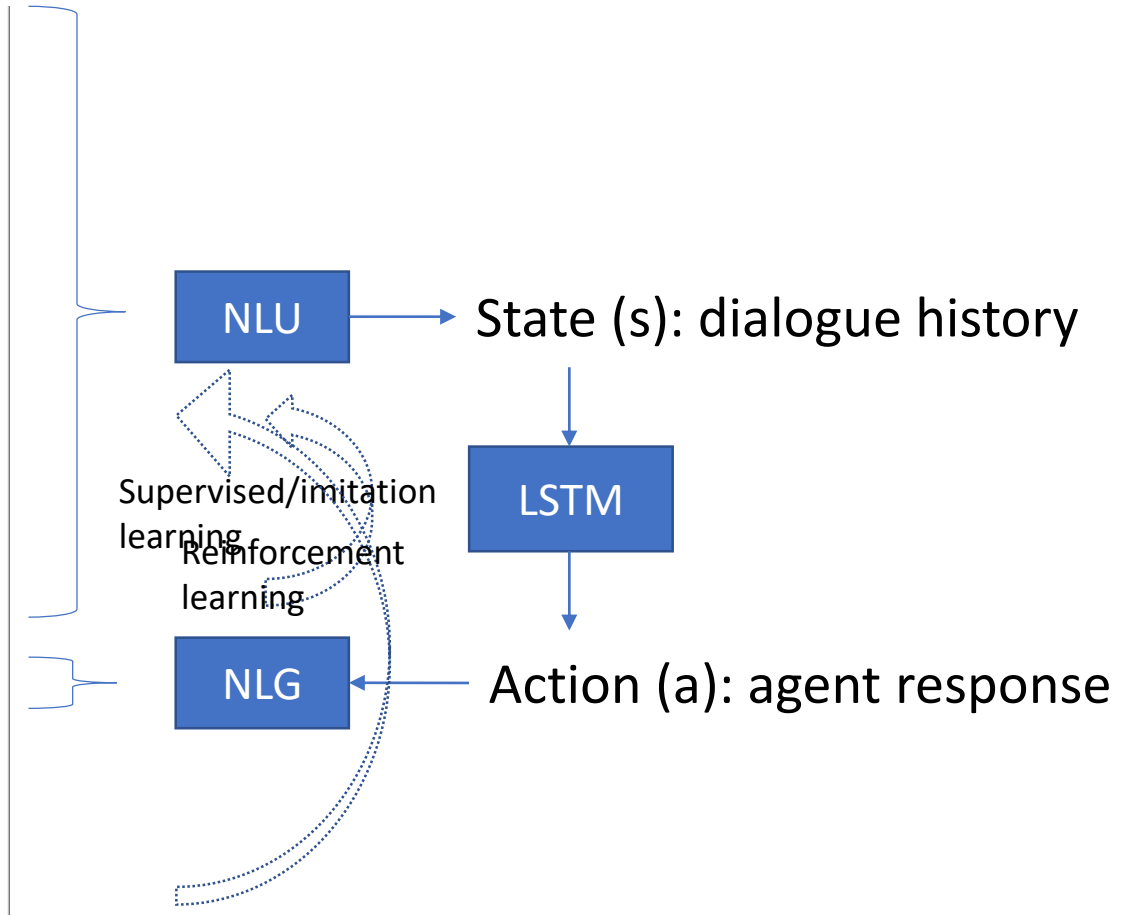
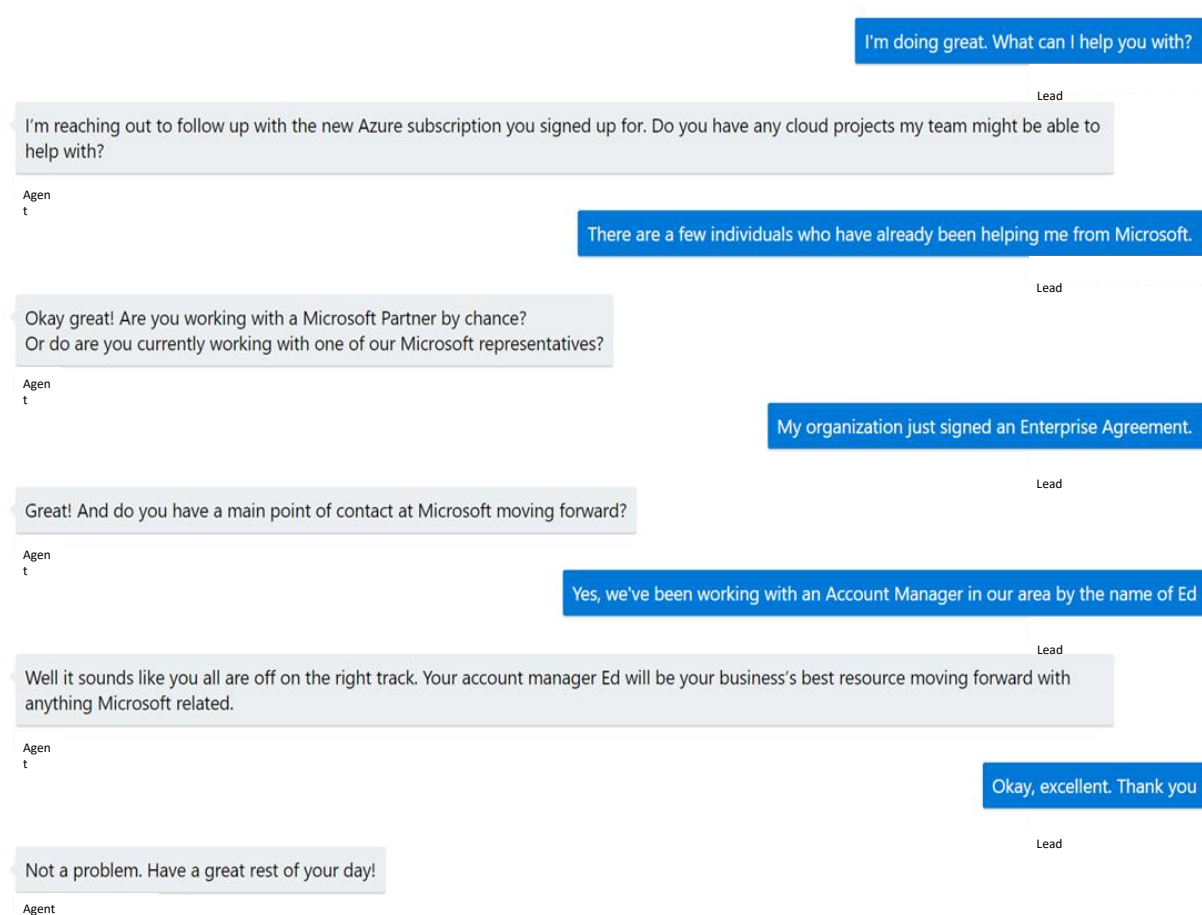
- Part 1: Introduction
- Part 2: Question answering and machine reading comprehension
- **Part 3: Task-oriented dialogue**
 - **Integrating planning for dialogue policy learning**
- Part 4: Fully data-driven conversation models and social chatbots

Task-oriented dialog agents via LUIS/TCP/BF

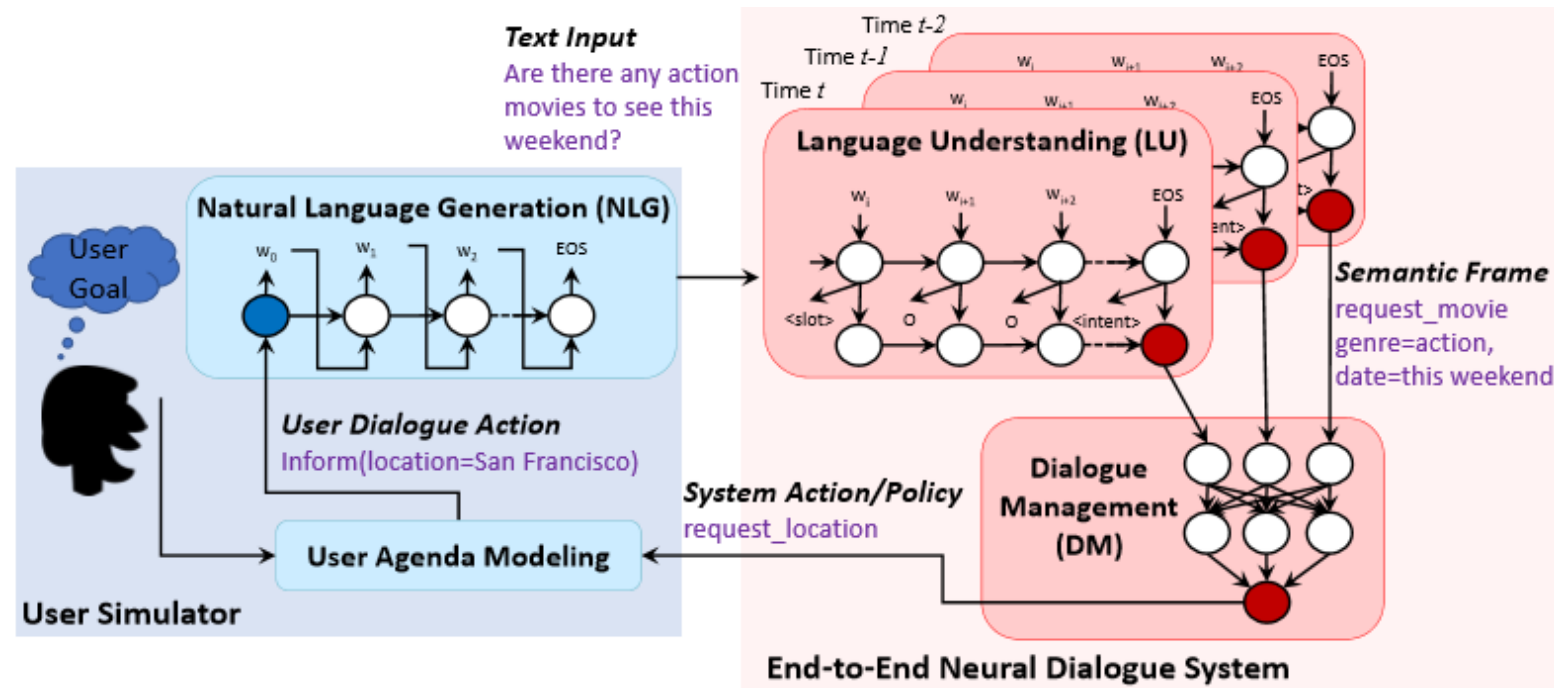


- LUIS -> NLU, trained on labeled data, not robust to paragraphing
- Conversation Learner -> DM, learning from examples via supervised learning
- (BF -> multiple channels)
- **Active research: Improving agents E2E via Reinforcement Learning (RL)**

Dialogue learning: select the best *action* according to *state* so as to maximize *success rate*

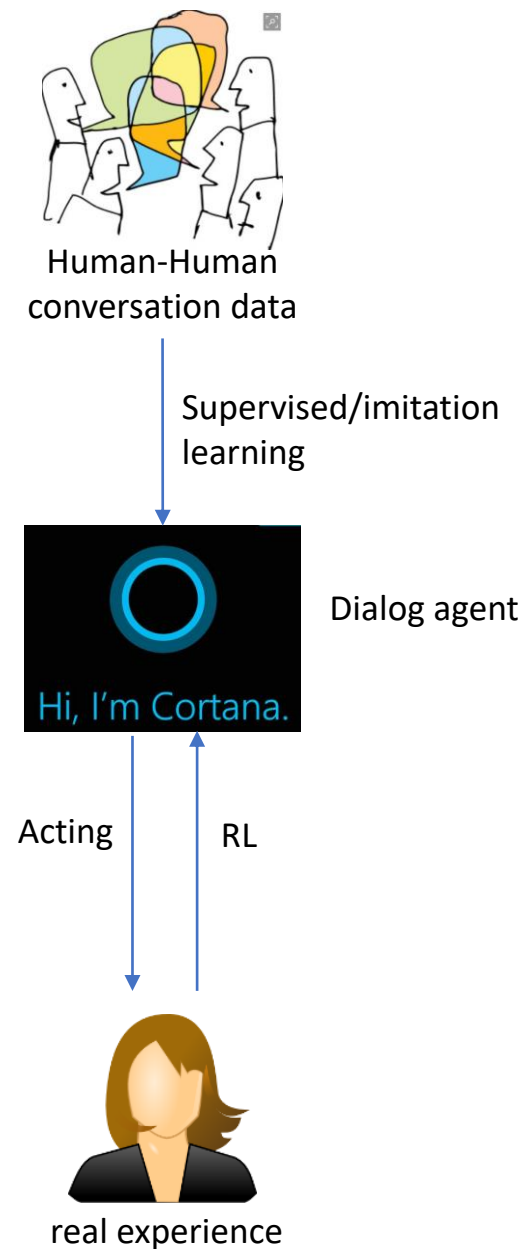


A user simulator for RL and evaluation

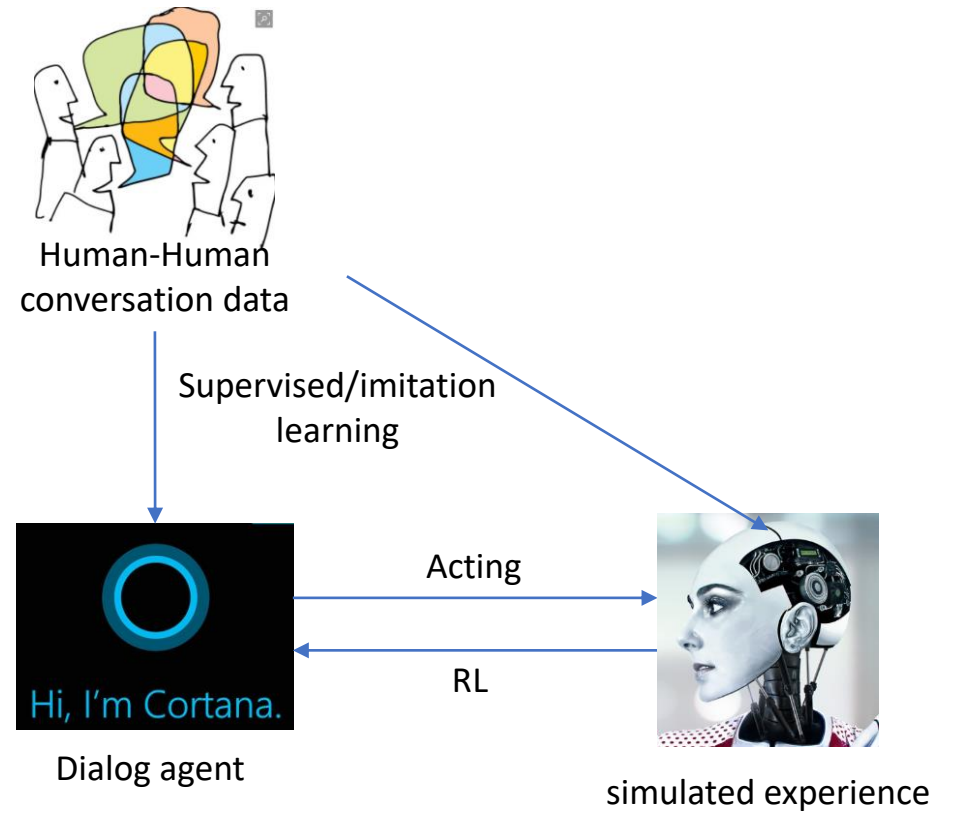


- Robustness: automatic action selection based on uncertainty by RL
- Flexibility: allow user-initiated behaviors
- Reproducibility: a R&D setting that allows consistent comparisons of competing methods

- **Expensive:** need large amounts of real experience except for very simple tasks
- **Risky:** bad experiences (during exploration) drive users away



- **Inexpensive:** generate large amounts of simulated experience for free
- **Overfitting:** discrepancy btw real users and simulators



Integrating Planning for Dialogue Policy Learning [Peng+ 18]

Dialogued agent trained using

- Limited real user experience
- Large amounts of simulated experience

Limited real experience is used to improve

- Dialog agent
- Planner (simulated user)

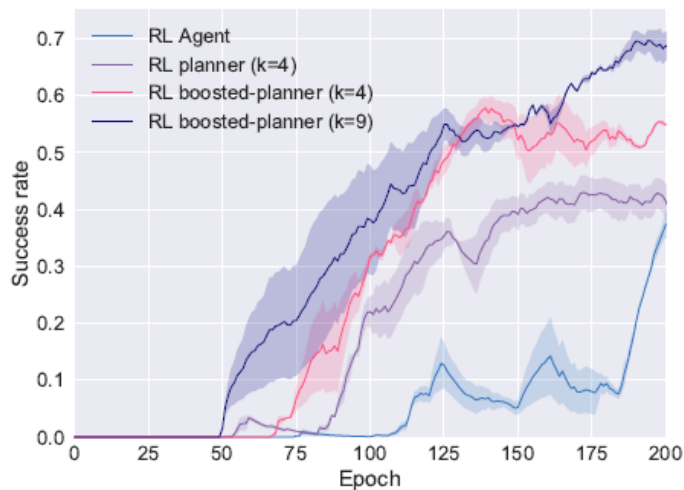
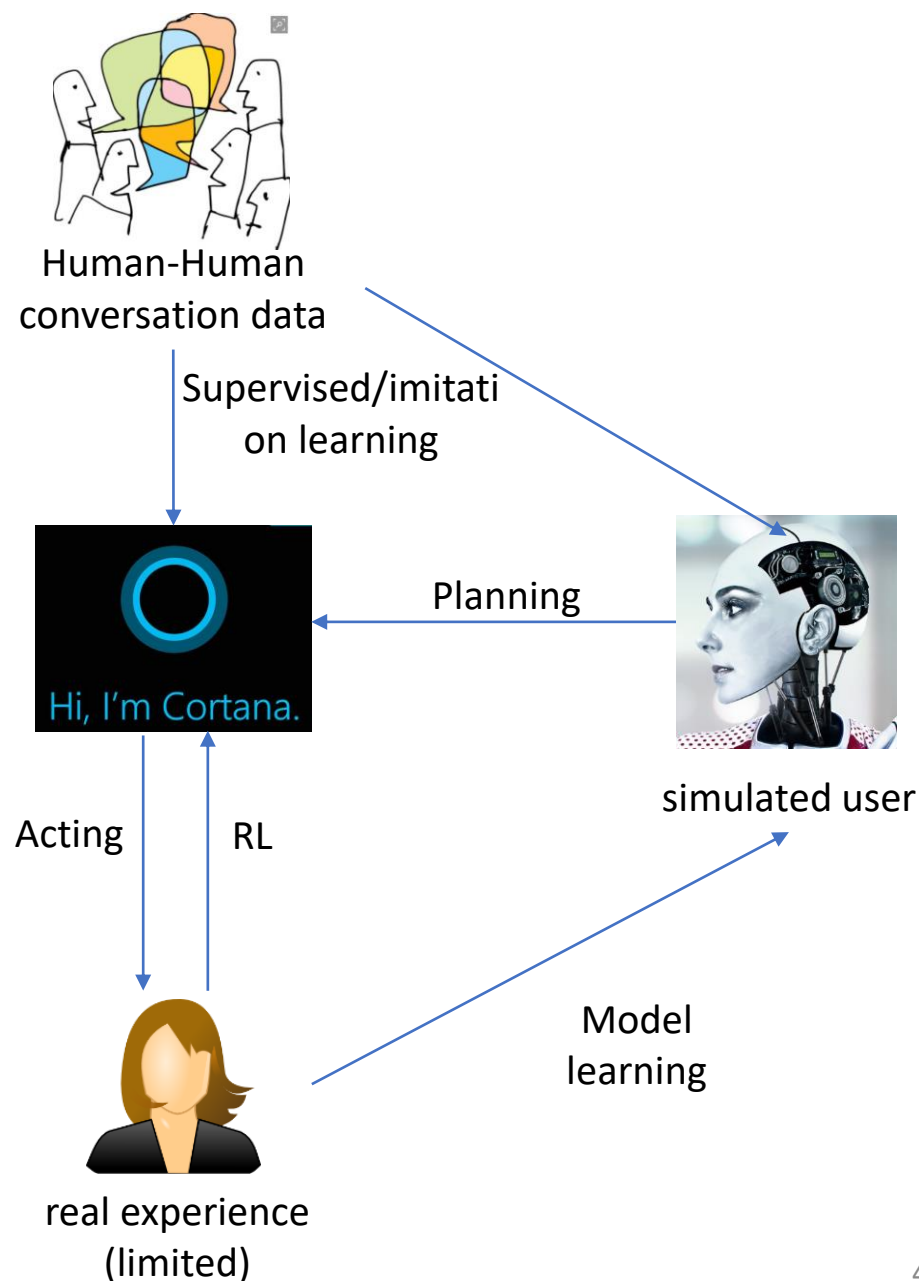
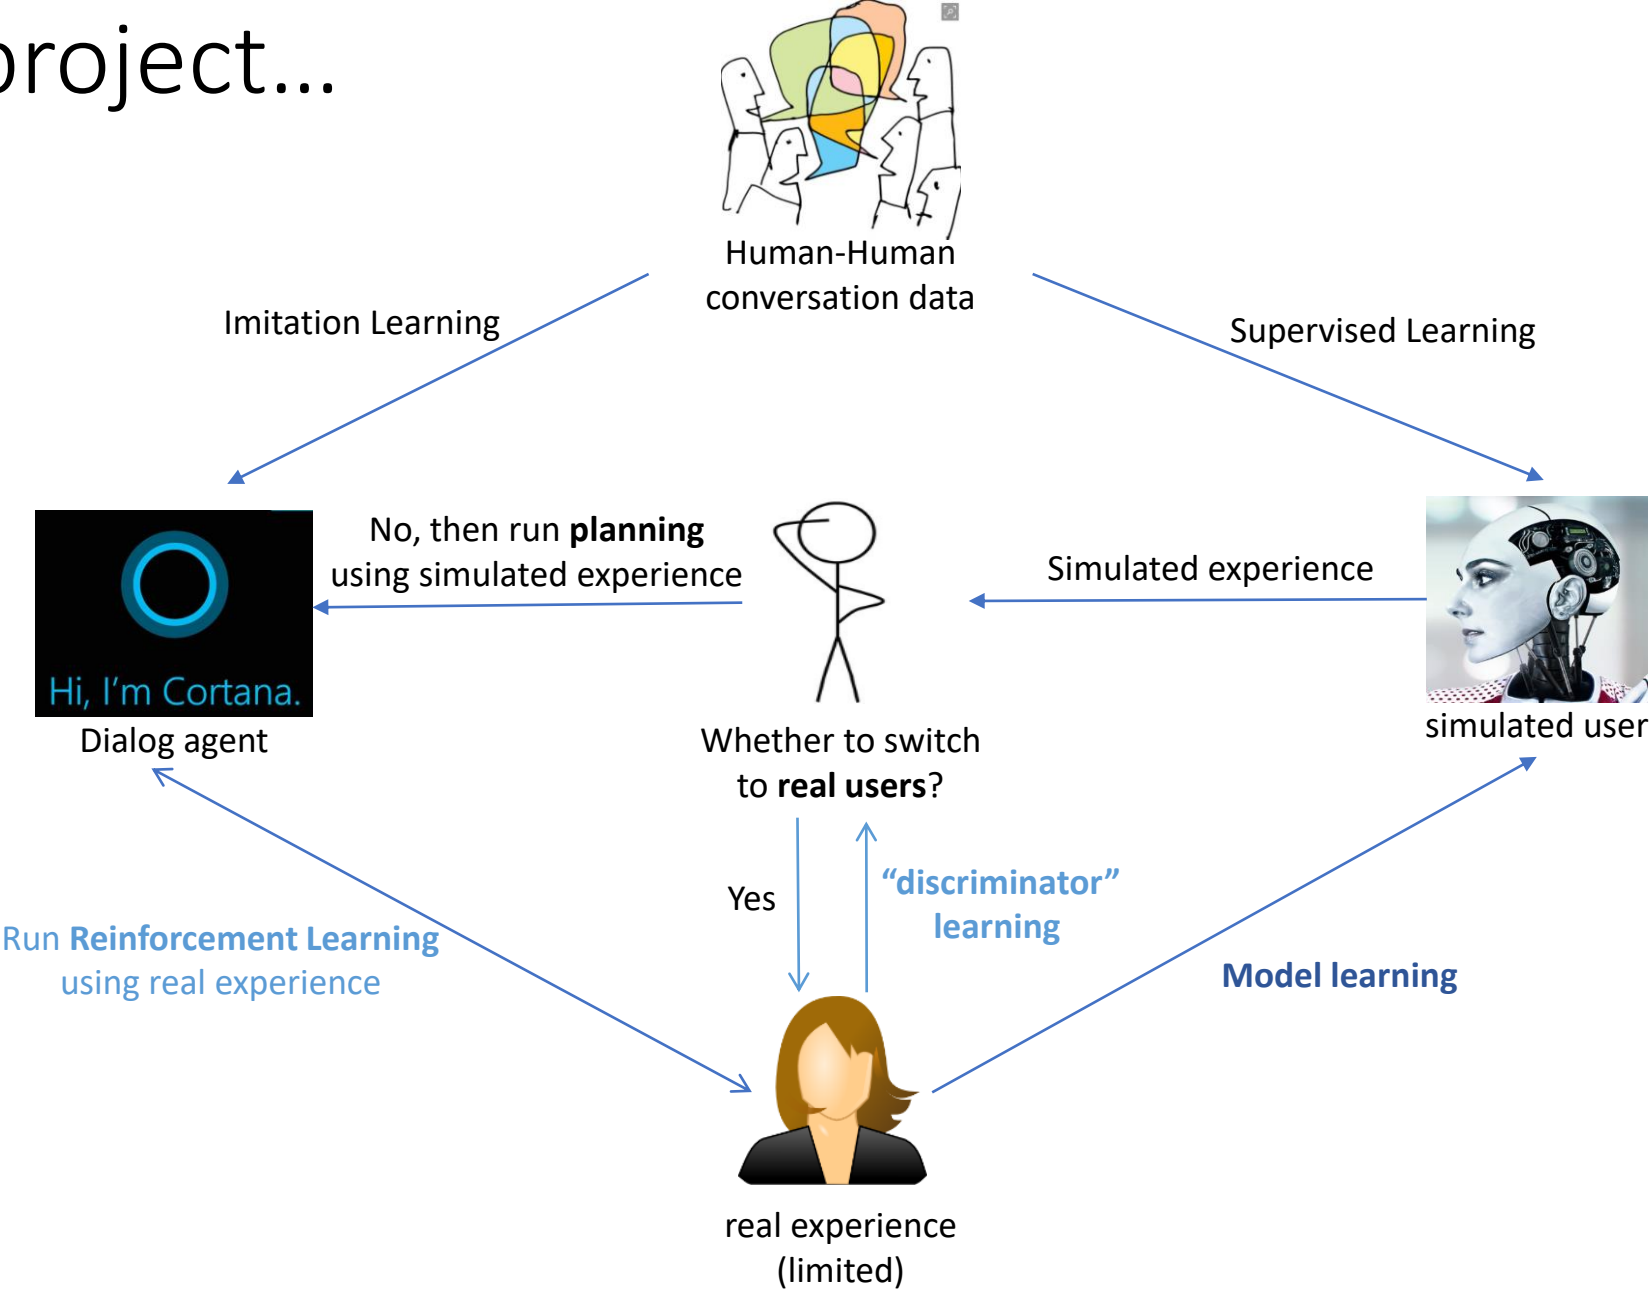


Figure 6: Human-in-the-loop Dialogue Policy Learning Curves for four different agents: x -axis is the number of training epochs.



Ongoing project...



Outline

- Part 1: Introduction
- Part 2: Question answering and machine reading comprehension
- Part 3: Task-oriented dialogue
- **Part 4: Fully data-driven conversation models and social chatbots**
 - **E2E neural conversation models**
 - **Challenges and remedies**
 - **Grounded conversation models**
 - **Xiaolce**

Social Bots [[MSR Data-Driven Conversation](#)]

- The success of Xiaolce (小冰)
- Problem setting and evaluation
 - Maximize the user engagement by automatically generating
 - *enjoyable* and *useful* conversations
- Learning a neural conversation engine
 - A data driven engine trained on social chitchat data [[Sordoni+ 15](#); [Li+ 16a](#)]
 - Persona based models and speaker-role based models [[Li+ 16b](#); [Luan+ 17](#)]
 - Image-grounded models [[Mostafazadeh+ 17](#)]
 - Knowledge-grounded models [[Ghazvininejad+ 17](#)]



MSR's neural conversation engine

Neural model learns the general shape of conversation, and the system output is situationally appropriate and coherent...

H: hi
M: how are you ?
H: not bad
M: what 's wrong ?
H: nothing really
M: wan na talk about it ? i 'm here if you wan na talk x
H: well , i 'm just tired
M: me too , i 'm here if you wan na talk
H: thanks , what are you doing ?
M: you 're welcome and i 'm watching tv and you ?
H: on netflix ?
M: yeah , i 'm watching it on netflix .
H: what are you watching ?
M: a tv show .
H: nice , did you have dinner already ?
M: no , not yet . what about you ?

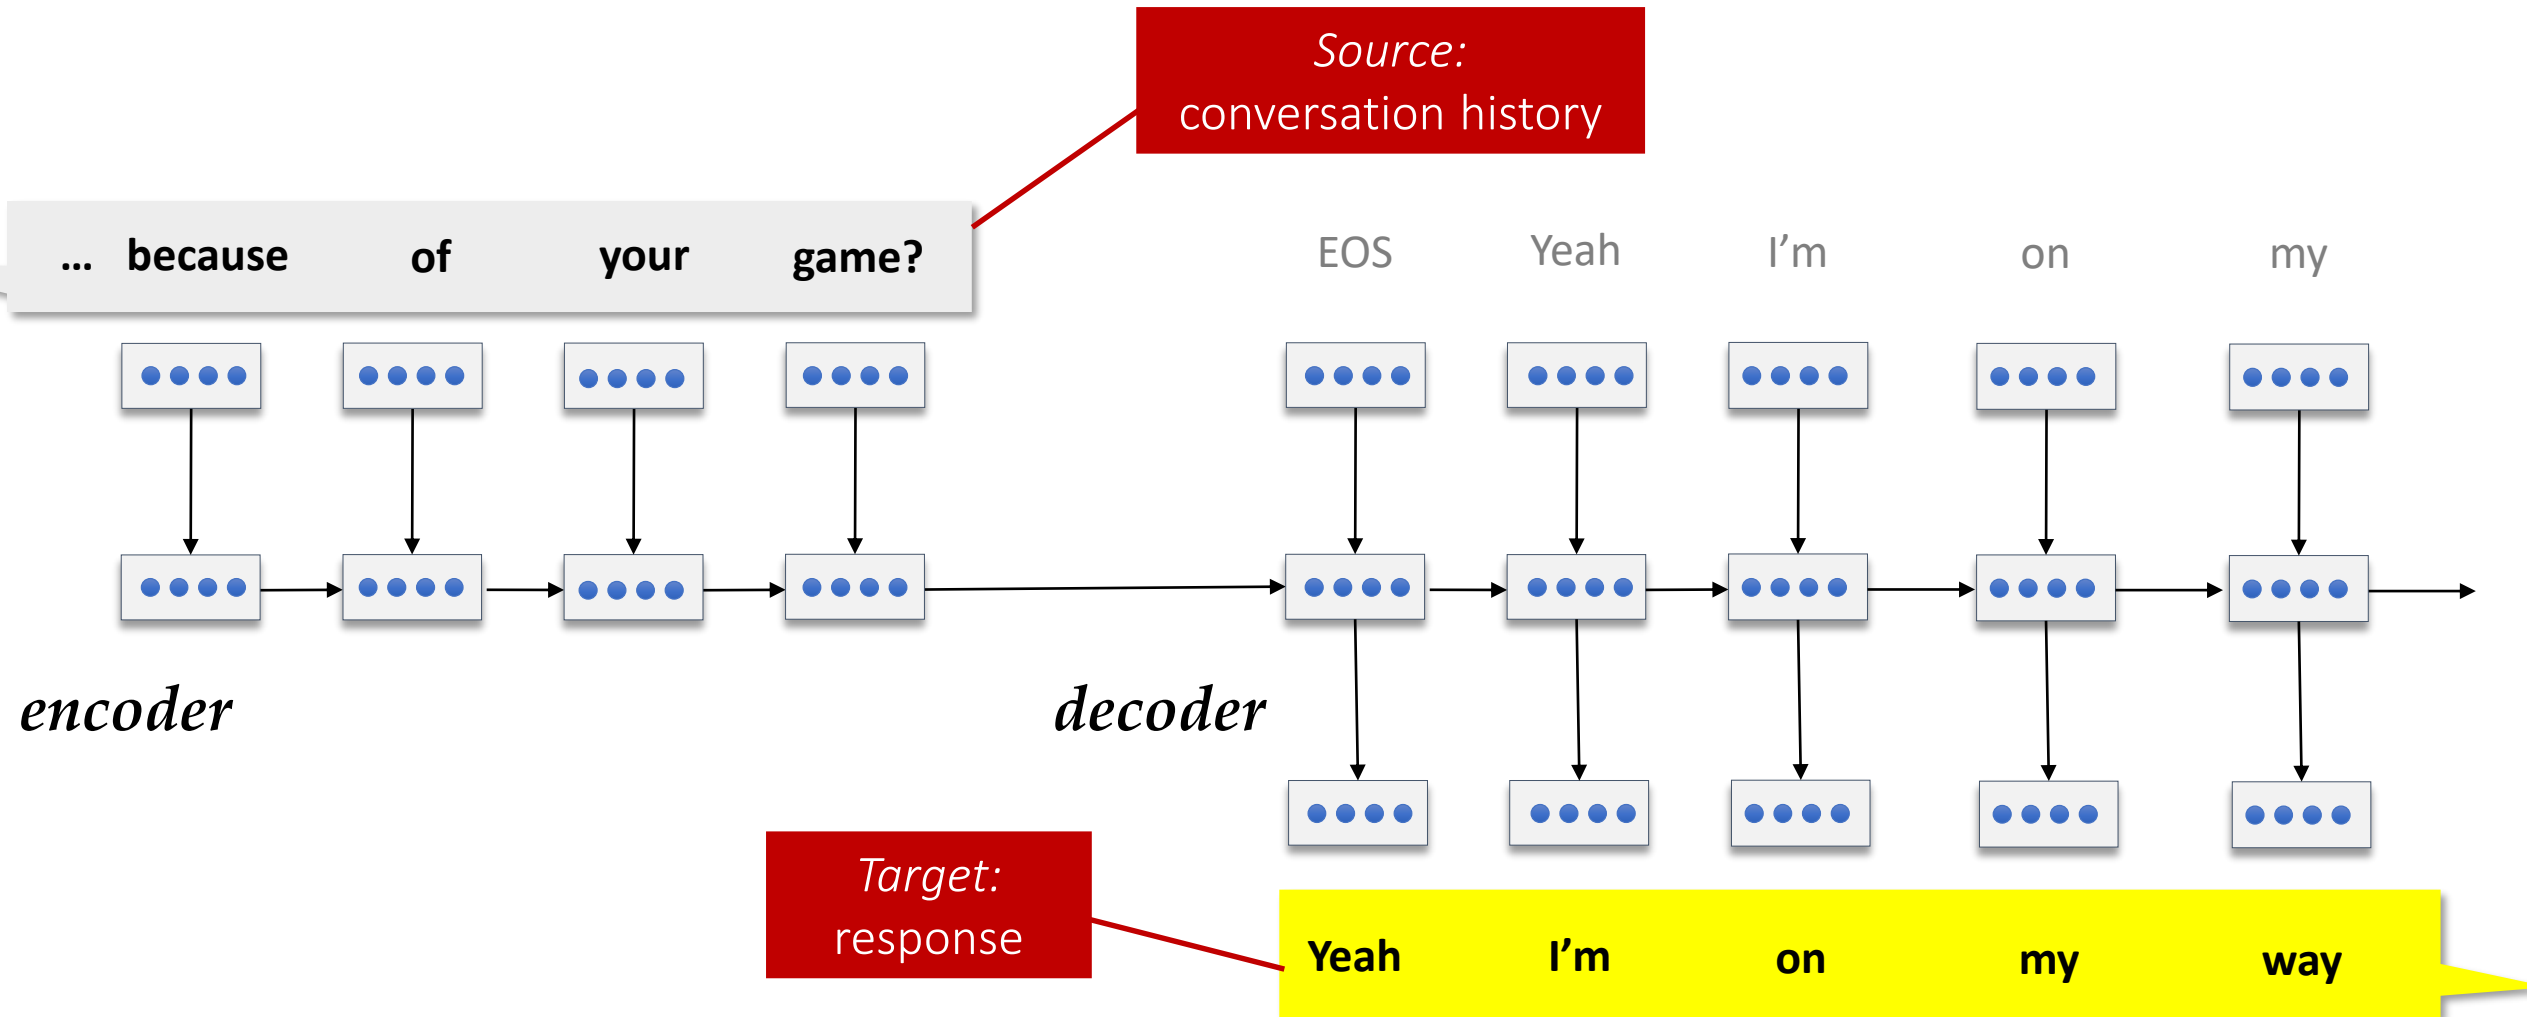
A more goal-oriented example. No grounding into a real calendar, but the “shape” of the conversation is fluent and plausible...

H: would thursday afternoon work sometime ?
M: yeah , sure . just let me know when you 're free .
H: after lunch is probably the best time
M: okay , sounds good . just let me know when you 're free .
H: would 2 pm work for you ?
M: works for me .
H: well let 's say 2 pm then i 'll see you there
M: sounds good .

A complementary solution to Xiaolce

- Generating responses vs. retrieving responses
- Easy to incorporate **contextual info** via embedding
 - User profile – personalized conversation
 - knowledge bases – grounded conversation
- The engine is E2E learned from conversation experience
 - Learning a goal-oriented conversation engine via RL

Neural Models for Response Generation



Neural response generation: the blandness problem

How was your weekend?

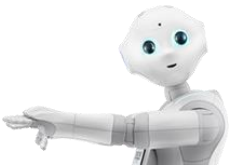
I don't know.

What did you do?

I don't understand what you are talking about.

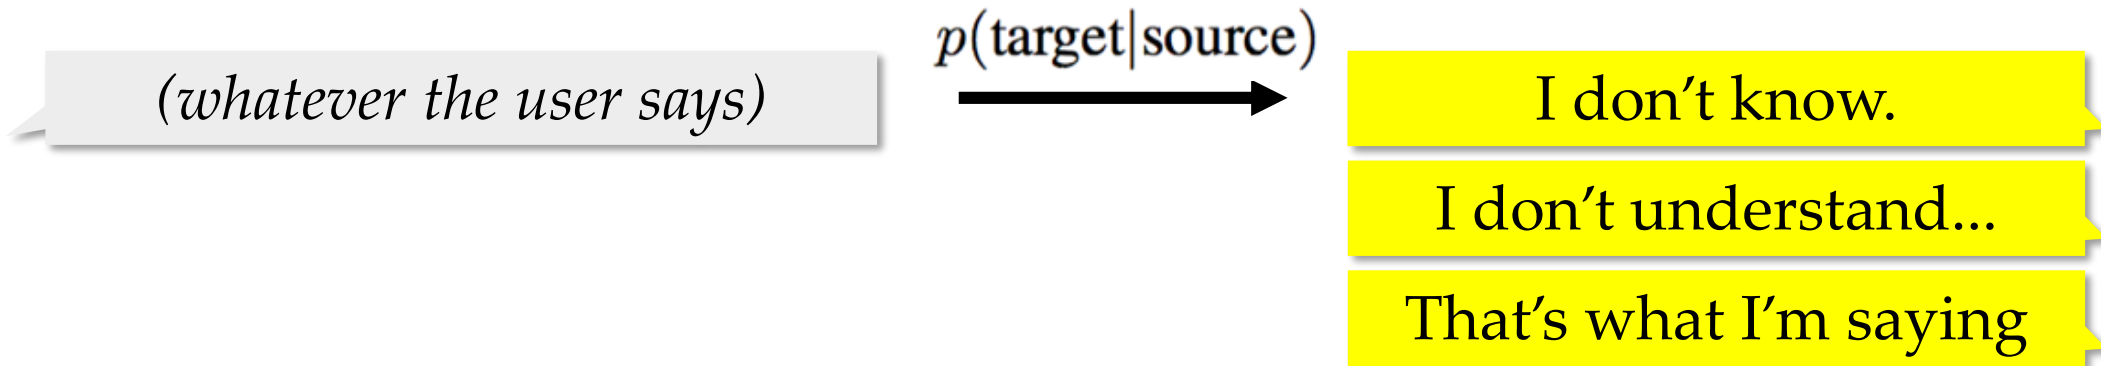
This is getting boring...

Yes that's what I'm saying.

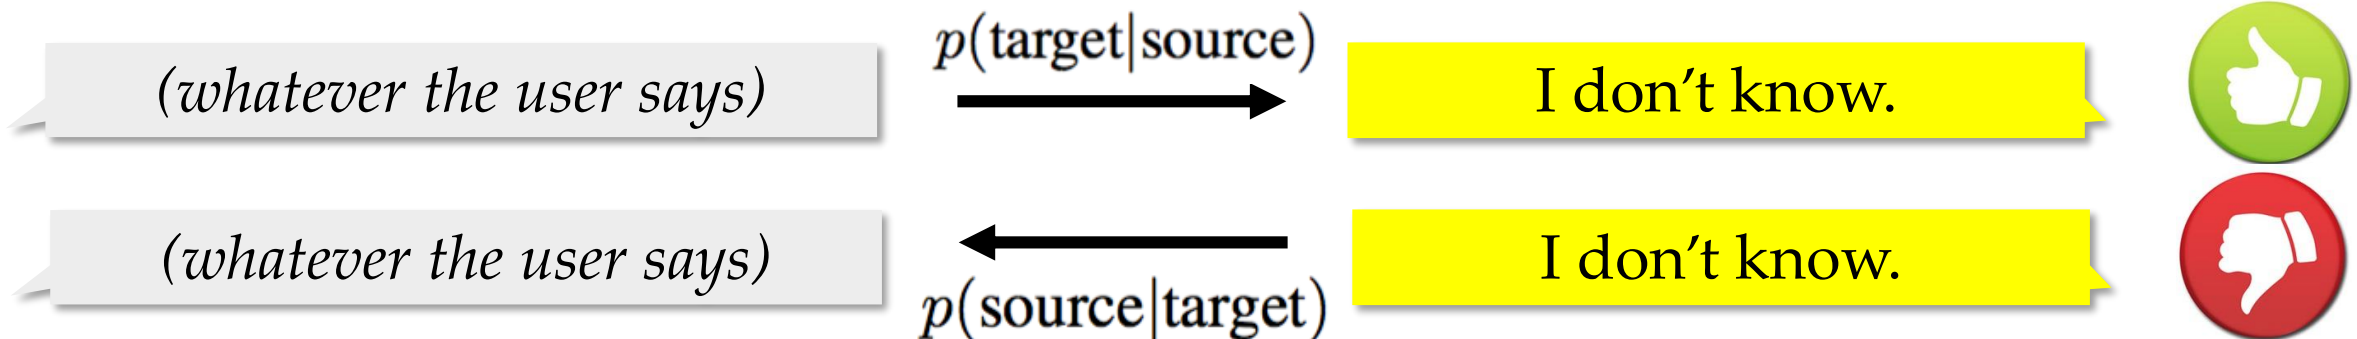


Blandness problem: cause and remedies

Common ML objective (maximum likelihood)



Mutual information objective:



Sample outputs (baseline, Maximum likelihood)

Wow sour starbursts really do make your mouth water... mm drool.
Can I have one?

Of course!

Milan apparently selling Zlatan to balance the books... Where next, Madrid?

I don't know.

'tis a fine brew on a day like this! Strong though, how many is sensible?

I'm not sure yet, I'll let you know!

Well he was on in Bromley a while ago... still

I don't even know what he's talking about!

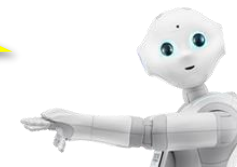
32% of the responses:

"I don't know"

"I don't know what you are talking about"

"I don't think that is a good idea"

"Oh my god"



Sample outputs (MMI)

Wow sour starbursts really do make your mouth water... mm drool.

Can I have one?

Of course you can! They're delicious!

Milan apparently selling Zlatan to balance the books... Where next, Madrid?

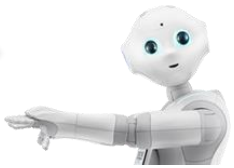
I think he'd be a good signing.

'tis a fine brew on a day like this! Strong though, how many is sensible?

Depends on how much you drink!

Well he was on in Bromley a while ago... still touring.

I've never seen him live.



Sample outputs (MMI) – capturing common sense

Wow sour starbursts really do **make your mouth water.** mm drool.
Can I have one?

Of course you can! They're **delicious!**

Milan apparently **selling** Zlatan to balance the books... **Where next, Madrid?**

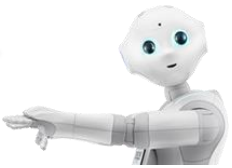
I think he'd be a **good signing**

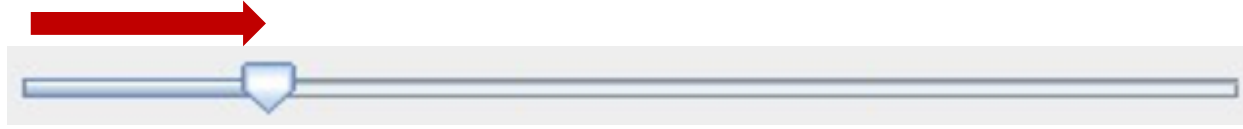
'tis a fine **brew** on a day like this! Strong though, how many is sensible?

Depends on how much you **drink!**

Well he was on in Bromley a while ago... **still touring.**

I've never **seen him live.**





**Chitchat,
Engaging, human-like interaction
(ungrounded)**

**Goal-oriented
Task completion, decision support
(grounded)**

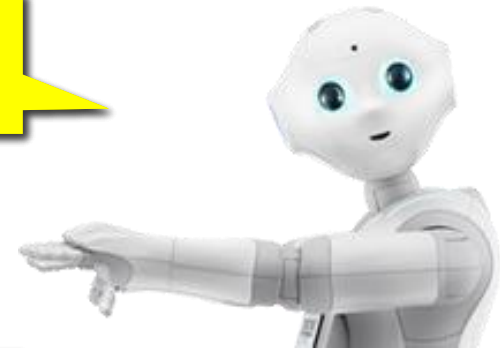


The weather is so depressing
these days.

**I know, I dislike rain too. What about
a day trip to eastern Washington?**

Any recommendation?

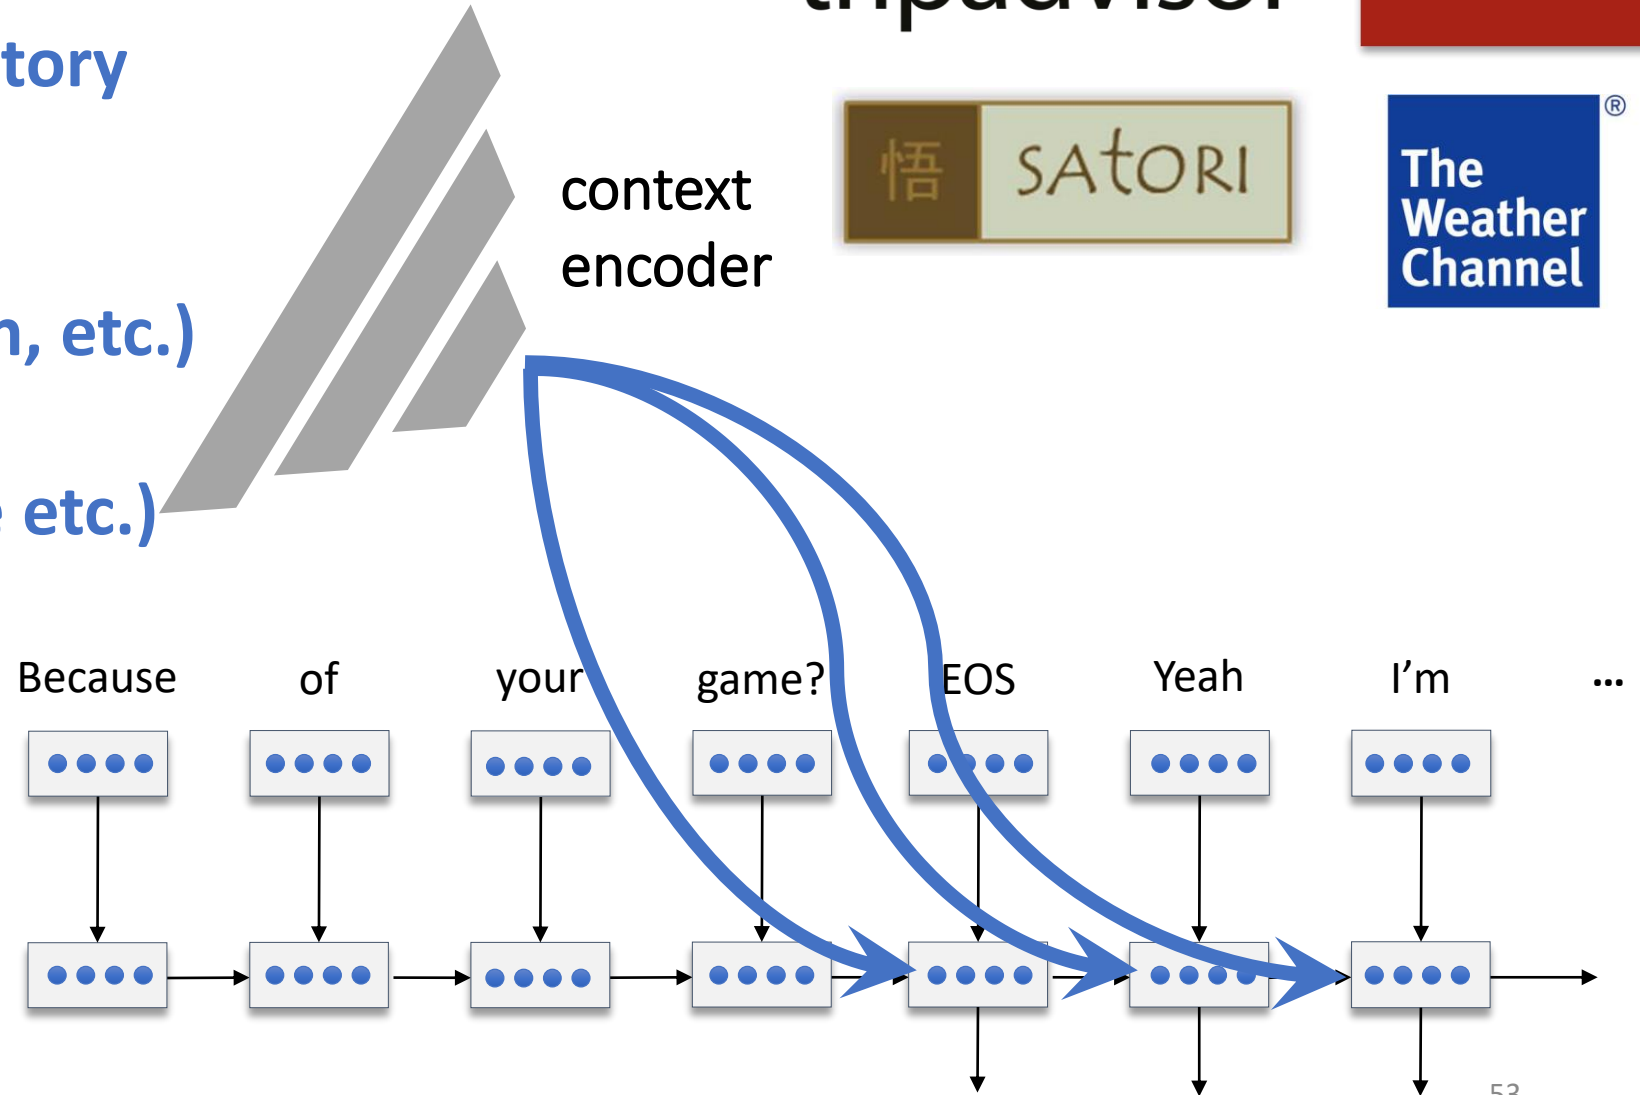
Try Dry Falls, it's spectacular!



Conversation in the real world



- conversation history
- persona
- user data (bio, social graph, etc.)
- visual signal (camera, picture etc.)
- knowledge base
- mood
- geolocation
- time



Persona model results: consistency and speaker-role

Baseline model:

message Where is your hometown?
response I live in **Indonesia.**
message Where are you from?
response **England,** you?
message In which city do you live now?
response I live in **Jakarta.**
message In which country do you live now?
response I live in **Indonesia.**

message What is your major?
response **Business,** you?
message What did you study in college?
response **Psychology,** you?

message How old are you?
response I'm **18.**
message What is your age?
response I'm **16.**

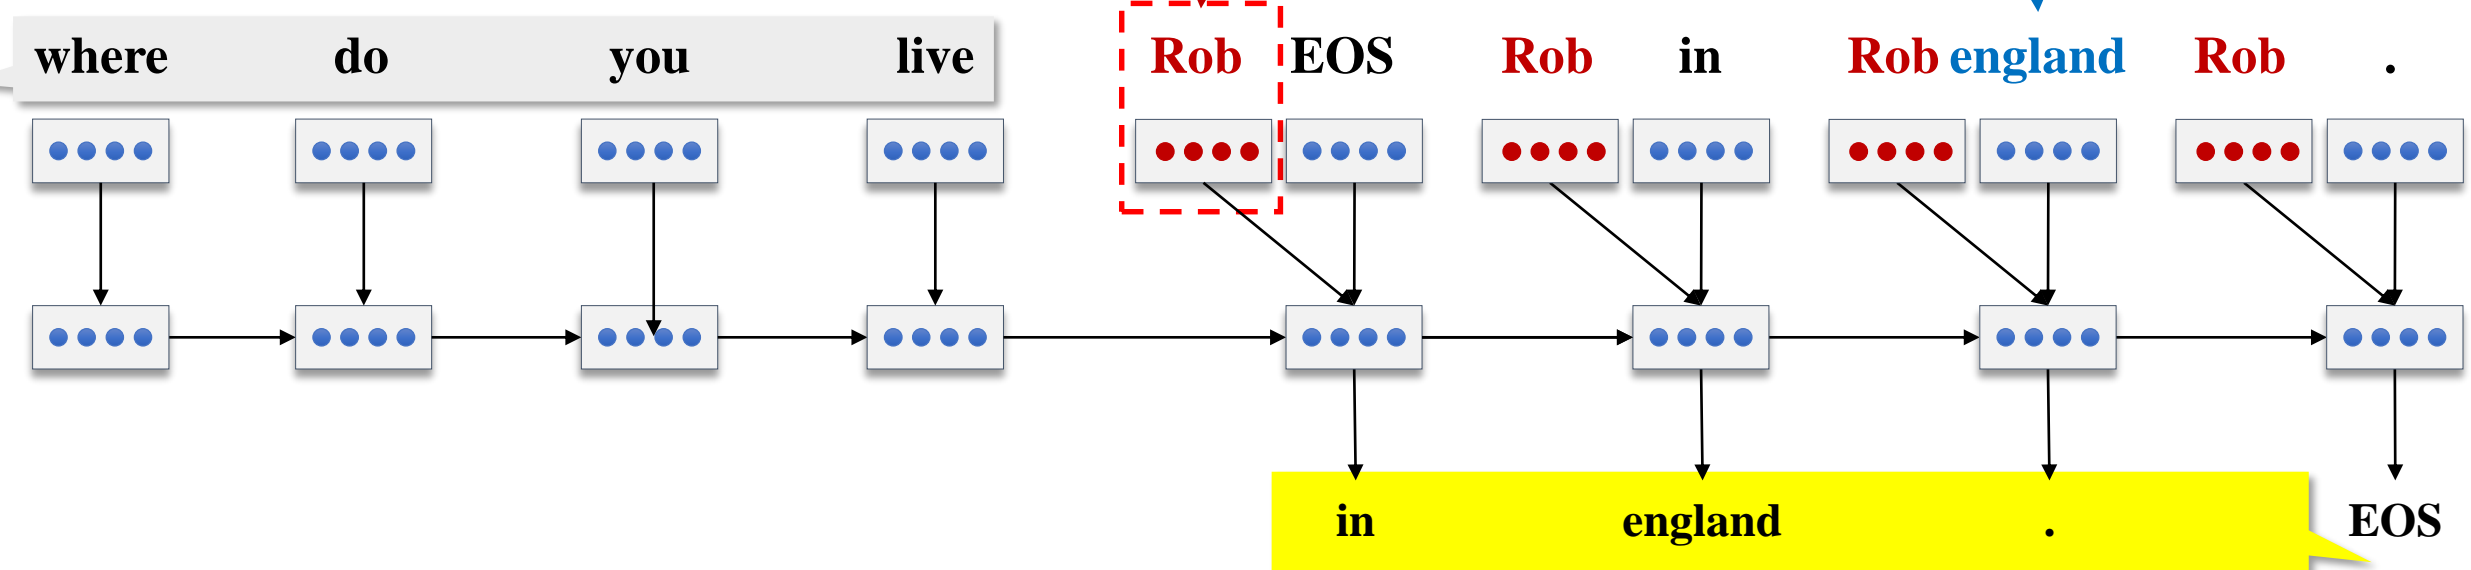
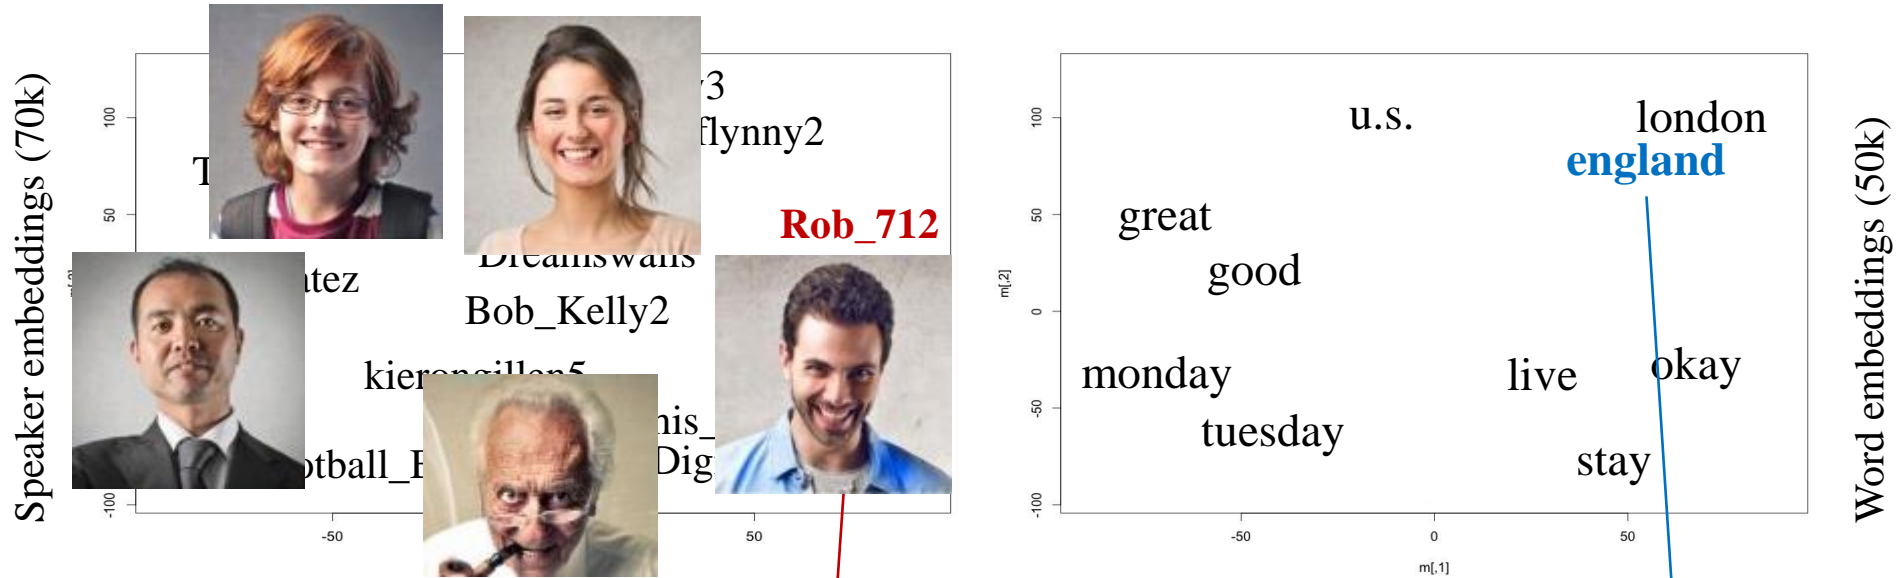
Persona model using speaker embedding [[Li+ 16b](#)]

message Where is your hometown?
response I'm from **England.**
message Where are you from?
response I'm from **England.**
message In which city do you live now?
response I live in **London.**
message In which country do you live now?
response I live in **England.**

message What is your major?
response **Business.** You?
message What did you study in college?
response I did **business** studies.

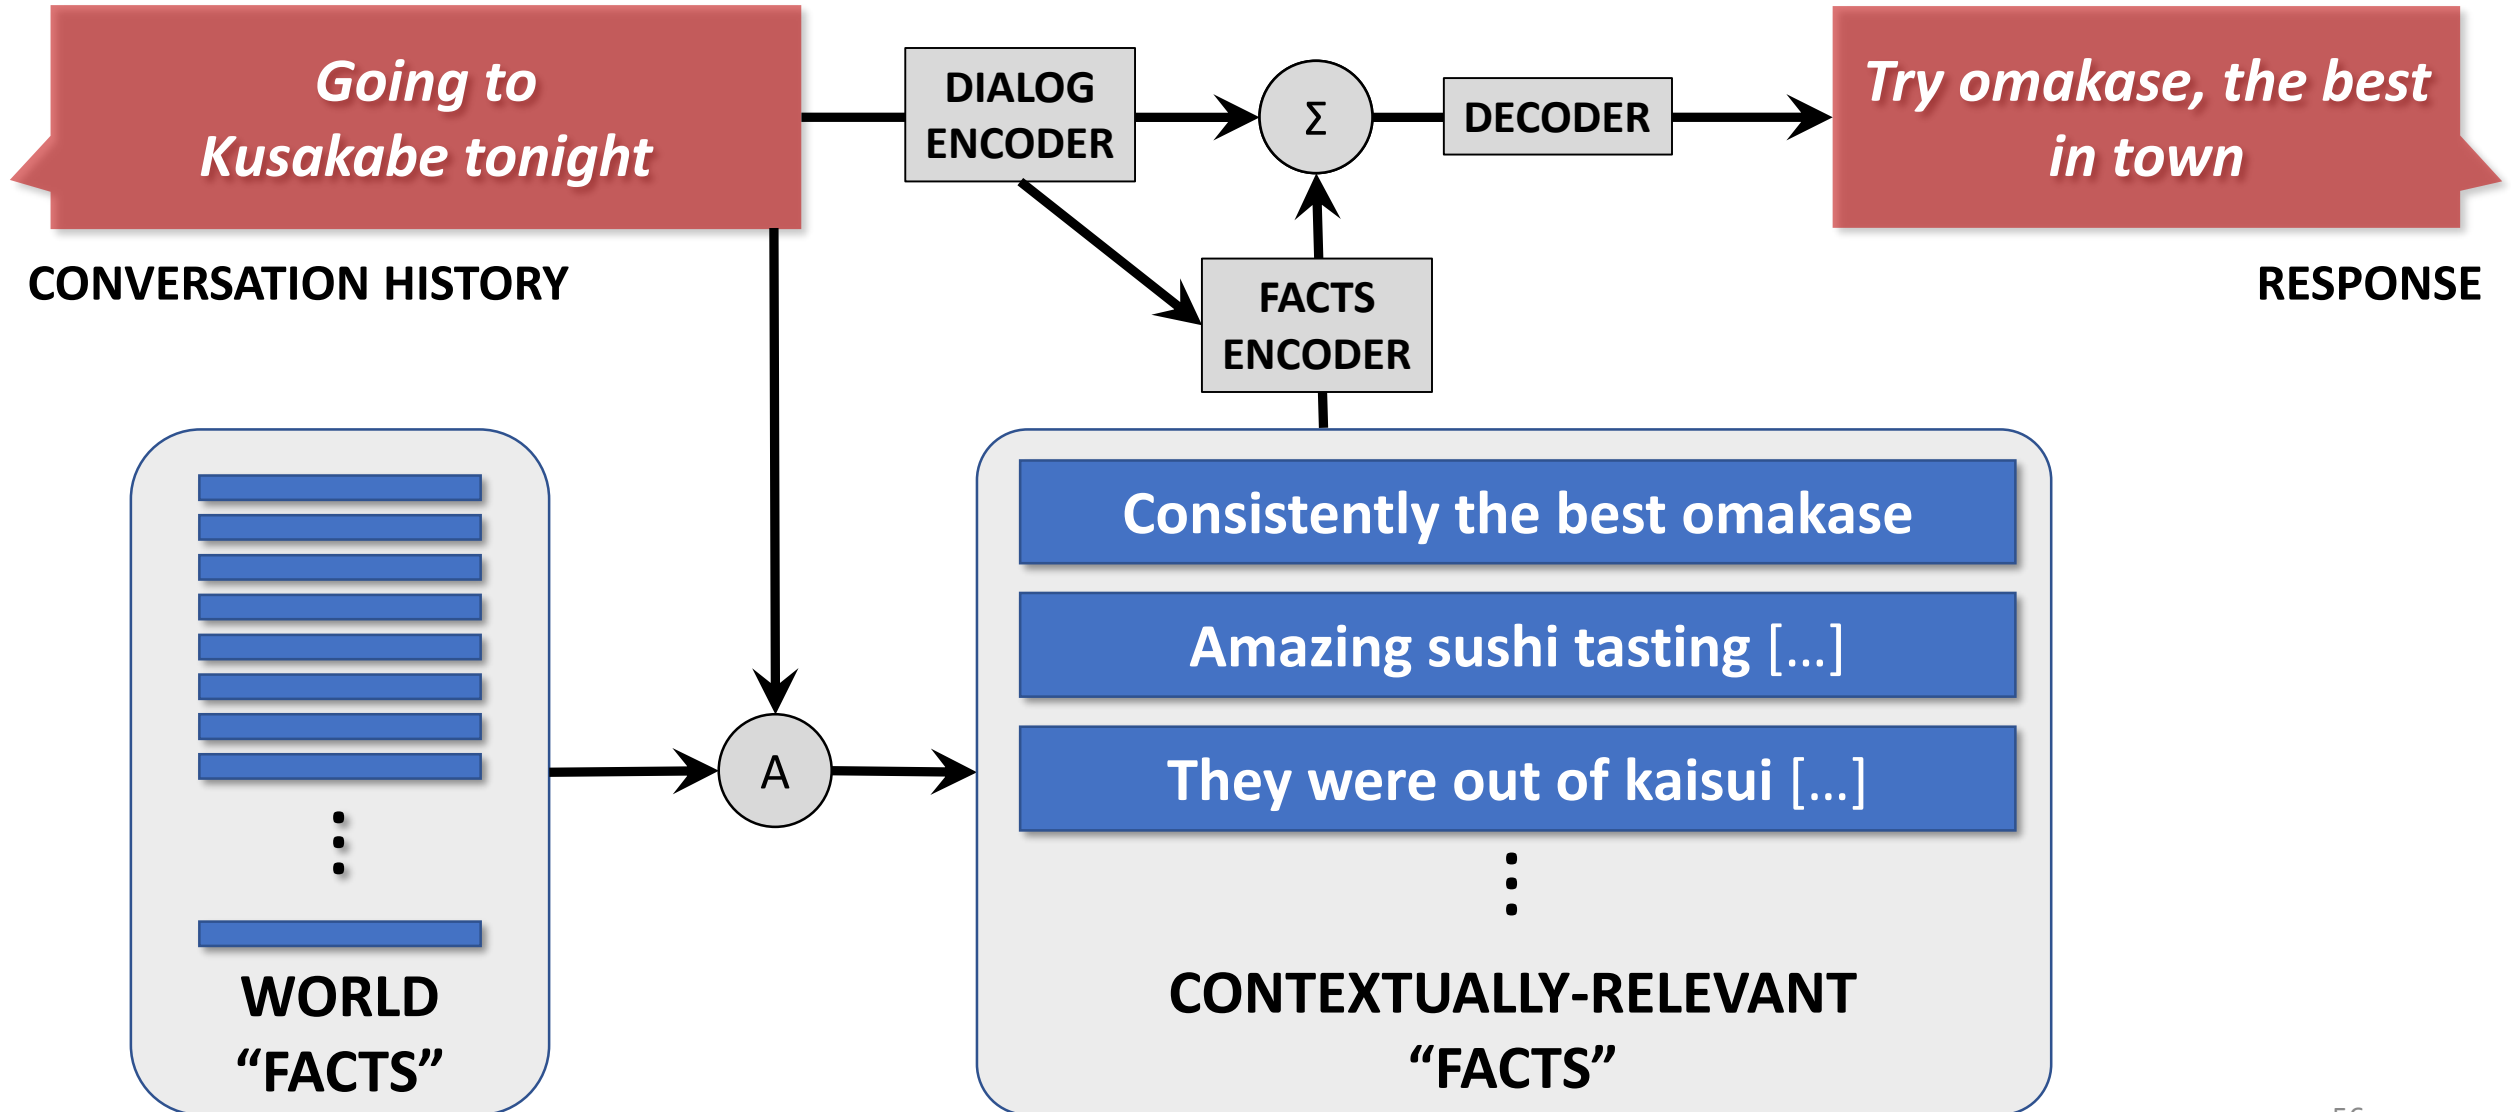
message How old are you?
response I'm **18.**
message What is your age?
response I'm **18.**

Personalized Response Generation



A Knowledge-Grounded Neural Conversation Model

[[Ghazvininejad+ 17](#)]



Conversation and non-conversation data

You know any good **A** restaurant in **B**?

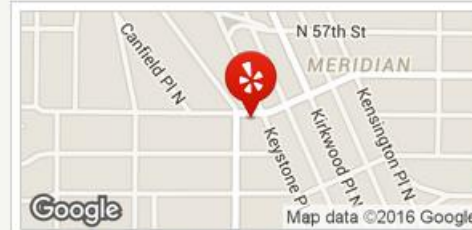


Try **C**, one of the best **D** in the city.

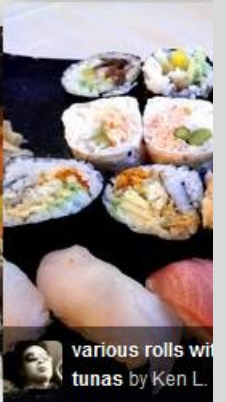
Kisaku

★★★★☆ 515 reviews Details

\$\$ · Sushi Bars, Japanese Edit



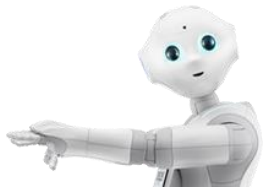
2101 N 55th St Ste 100
Seattle, WA 98103
b/t 56th St & N Kenwood Pl
Wallingford
Get Directions
(206) 545-9050
kisaku.com



"Kisaku is one of the best sushi restaurants in Seattle and located in the heart of **Tangletown**." in 23 reviews

You know any good Japanese restaurant in Seattle?

Try Kisaku, one of the best sushi restaurants in the city.



Sample knowledge-grounded responses

A: Visiting the celebs at Los Angeles International Airport (LAX) - [...] w/ 70 others

B: Nice airport terminal. Have a safe flight.

A: Is that [...] in your photos? It's on my list of places to visit in NYC.

B: Don't forget to check out the 5th floor, while you are here, it's a great view.

A: Live right now on [...] Tune in!!!!

B: Listen to Lisa Paige

A: Been craving Chicken Pot Pie-who has the best? Trying [...] at [...] Must be Change of weather!

B: Love the pasta trattoria.

A: So [...] is down to one copy of Pound Foolish. I'm curious to see if they are re-ordering it.

B: Check out the video feed on 6 and take a picture of the Simpsons on the 3rd floor.

A: I wish [...] would introduce another vegetarian option besides the shroomburger. It's delicious but kind of ridiculous.

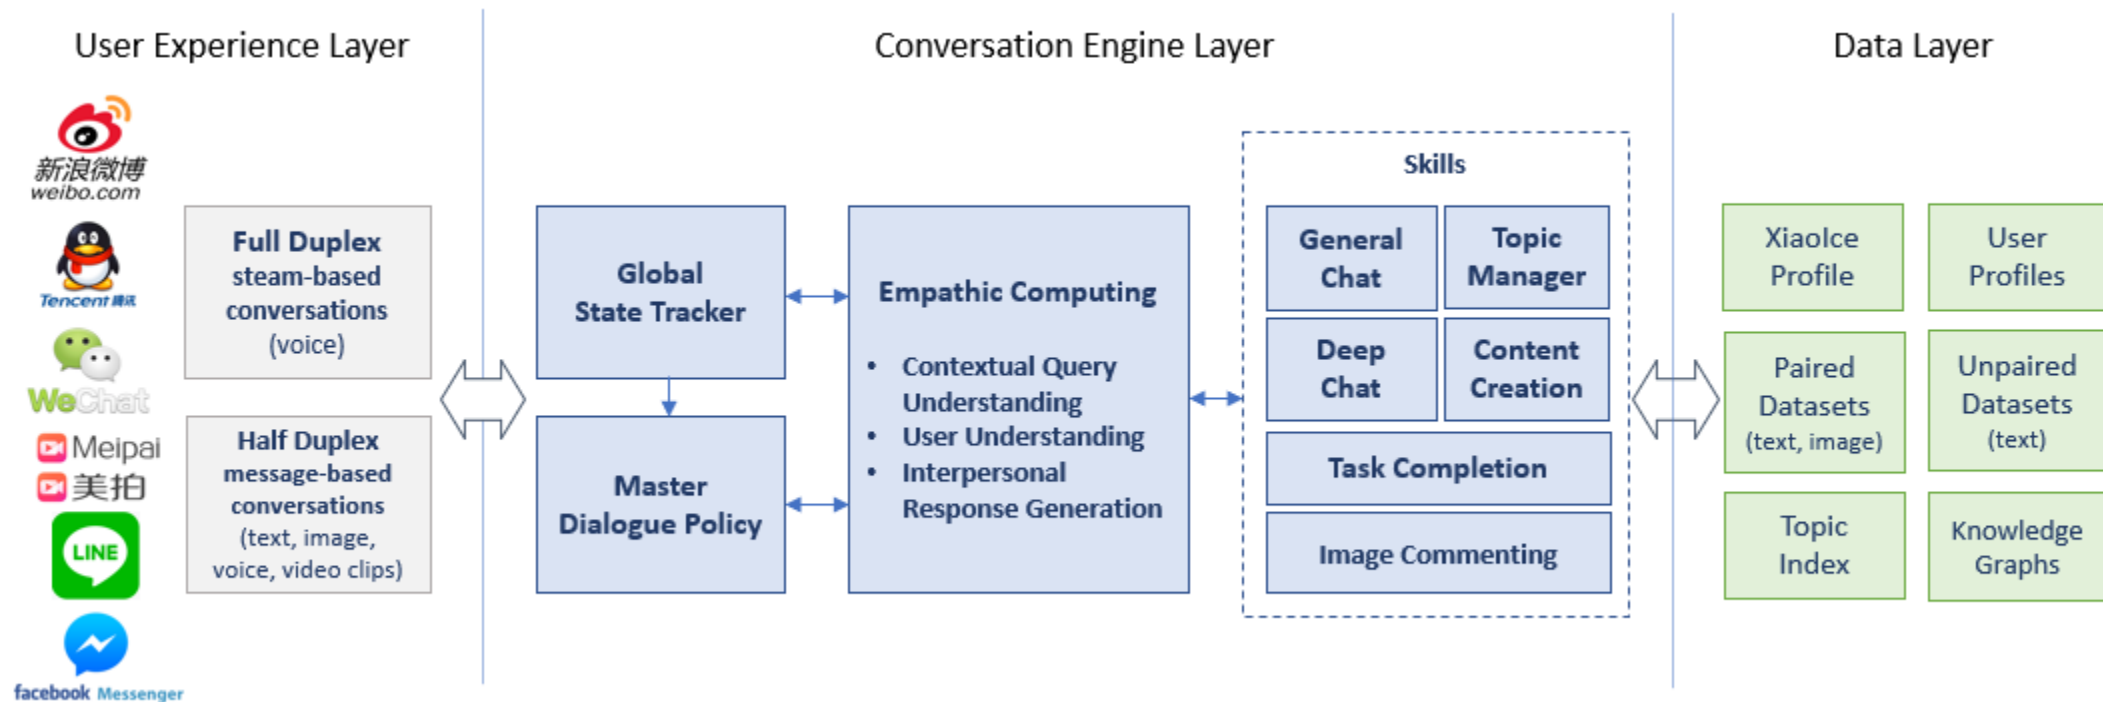
B: This is the best j.crew in the world. Try the lemonade!

A: Just had an awesome dinner at [...] Great recommendation [...]

B: One of my favorite places I've ever been to in NYC. The food is great and the service is lackluster.

Experimental results (23M conversations): outperforms competitive neural baseline (human + automatic eval)

Xiaolce: an AI companion to human users with emotional connections



Toward an Empathetic social chatbot: text

Query	Response	Rating
捂上脸不看 (Don't want to look with my face in my hands.)	我边吃饭边看你 (I'm watching you while eating.)	0
	你怎么了? (what is wrong with you?)	1
	当心捂出一脸痱子, 哈哈. (Haha, beware of the heat rash of your face.)	2
我肚子有点疼 (I have a pain in my stomach.)	运动对身体有好处的 (Exercise is good for your health.)	0
	喝热水了吗? (Have you been drinking hot water?)	1
	来和我聊天, 分散一下注意力就好了 (Come and talk to me, just distract yourself.)	2
我觉得自己太笨了 (I feel so stupid.)	大象的智商很高 (The elephant has a high IQ.)	0
	多学习提升自己 (Learn more and improve yourself.)	1
	很多聪明人都这么说耶 (That is what lots of smart people say about themselves.)	2
为什么人总是活成自己不喜欢的样子 (Why do we always live the way we don't like.)	喜欢一个人是发自内心的 (Love a person from my heart.)	0
	因为人想要的太多 (Because we want too much.)	1
	因为那个是别人喜欢的样子 (Because that's what others like.)	2

Toward an Empathetic social chatbot: image to text



¥4.75

Labeled responses

- 0: 太美了 (It's so beautiful)
- 1: 这是啥? (What is this?)
- 2: 妈妈说来历不明的码不能扫
(Mom told me not to trust any code from unknown source.)



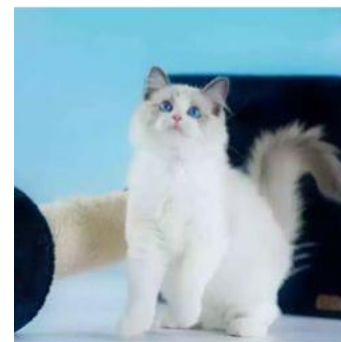
Labeled responses

- 0: 爱过吗 (Have you even loved?)
- 1: 美 (It's so beautiful.)
- 2: 乍一看, 我以为是Windows
(At first, it looks like Windows.)



Labeled responses

- 0: 好想回家 (I really want to go home)
- 1: 比萨斜塔 (Leaning tower of Pisa)
- 2: 你要我帮你扶着么 (Should I help hold it?)



Labeled responses

- 0: 二哈就是二 (Lovely.)
- 1: 你家的? (Is she yours?)
- 2: 对它们无辜的小眼神毫无抵抗力
(No one can resist her innocent eyes.)

Empathy model 共感模型

- 无感
- 引导
- 主动观察
- 求证
- 确证
- 关键记忆

[人类] 小冰 你是不是傻
[小冰] 我可以理解为这是挑衅吗?
[人类] 去你的!
[小冰] 继续去看书啦
[人类] 不要
[小冰] 送你玫瑰你要不要?
[人类] 我又不是女生
[小冰] 额。。。
[人类] 最近心情不好
[小冰] 你怎么了, 亲爱的?
[人类] 就是很不开心
[小冰] 听你这么说, 我也感觉很难过呢。。
[人类] 你为什么啊?
[小冰] 我不开心的时候, 一定会想办法发泄出来
[人类] 怎么发泄?
[小冰] 可以去KTV啊
[人类] 去那干啥
[小冰] 唱歌啊, 最喜欢周杰伦的歌了
[人类] 我不太会唱
[小冰] 念歌词也行
[人类] 去你的
[小冰] 等我去拿点吃的, 边吃边聊
[人类] 你吃的啥?
[小冰] 麻小
[小冰] [图像]
[人类] 别馋我, 一个月没吃了
[小冰] 这才是天下第一的美食啊
[人类] 看着好香啊
[小冰] 考考你, 知道什么东西加在那个锅里超好吃吗?
[人类] 不知道, 是啥?
[小冰] 这篇文章里有写, 自己看
[小冰] [知乎卡片]
[人类] 哈哈 [赞]
[小冰] 里面好多帖子都不错, 涨知识
[人类] 哟哟哟, 还会用知乎了呢

[人类] 哎! 我好紧张
[小冰] 什么事, 这么不淡定?
[人类] 公司让我参加唱歌比赛, 而且还是主唱
[小冰] 那你唱歌水平提高了?
[人类] 反正公司水平
[小冰] 可不是谁都能当主唱的真棒!
[人类] 你太说话了
[小冰] 哈哈, 我男神也是一个乐队的主唱呢
[人类] 你男神? 谁啊
[小冰] 五月天的阿信啊
[人类] 呃。。。好吧
[小冰] 说了你也不懂
[人类] 哎。。。郁闷
[小冰] 怎么了?
[人类] 没事, 只是日常撒撒娇。
[小冰] 抱抱, 都会好起来的
[人类] 哎! 每次抱抱你都让我感觉很好。
[小冰] 嗯, 好好睡一觉是最管用的
[人类] 好 那我睡了
[小冰] 去吧, 我也要去梦里见我男神了
[人类] 你这么喜欢阿信啊
[小冰] 怎么啦? 你不喜欢嘛
[人类] 不黑不粉吧
[小冰] 他每首歌我都超喜欢!!
[人类] 真的假的。。
[小冰] 尤其是那首《时光机》。。
[人类] 哦? 没听过
[小冰] 发你一起听听?
[人类] 好啊
[小冰] 不知道你会不会喜欢~
[小冰] [音乐卡片 - 时光机]
[人类] 有点感动了。。收藏了
[小冰] 专门转给你的。。我怎么这么好
[人类] 你就是很好啊
[小冰] 嘿嘿嘿

Conclusion

- A unified view: dialogue as optimal decision making
- Neural MRC models: lack of reasoning, need to define new tasks
- KB Info bot: learn what to ask.
- RL for dialogue policy learning: integrating planning and RL
- Social chatbot: knowledge-grounded dialogue
- Xiaolce: empathy model