

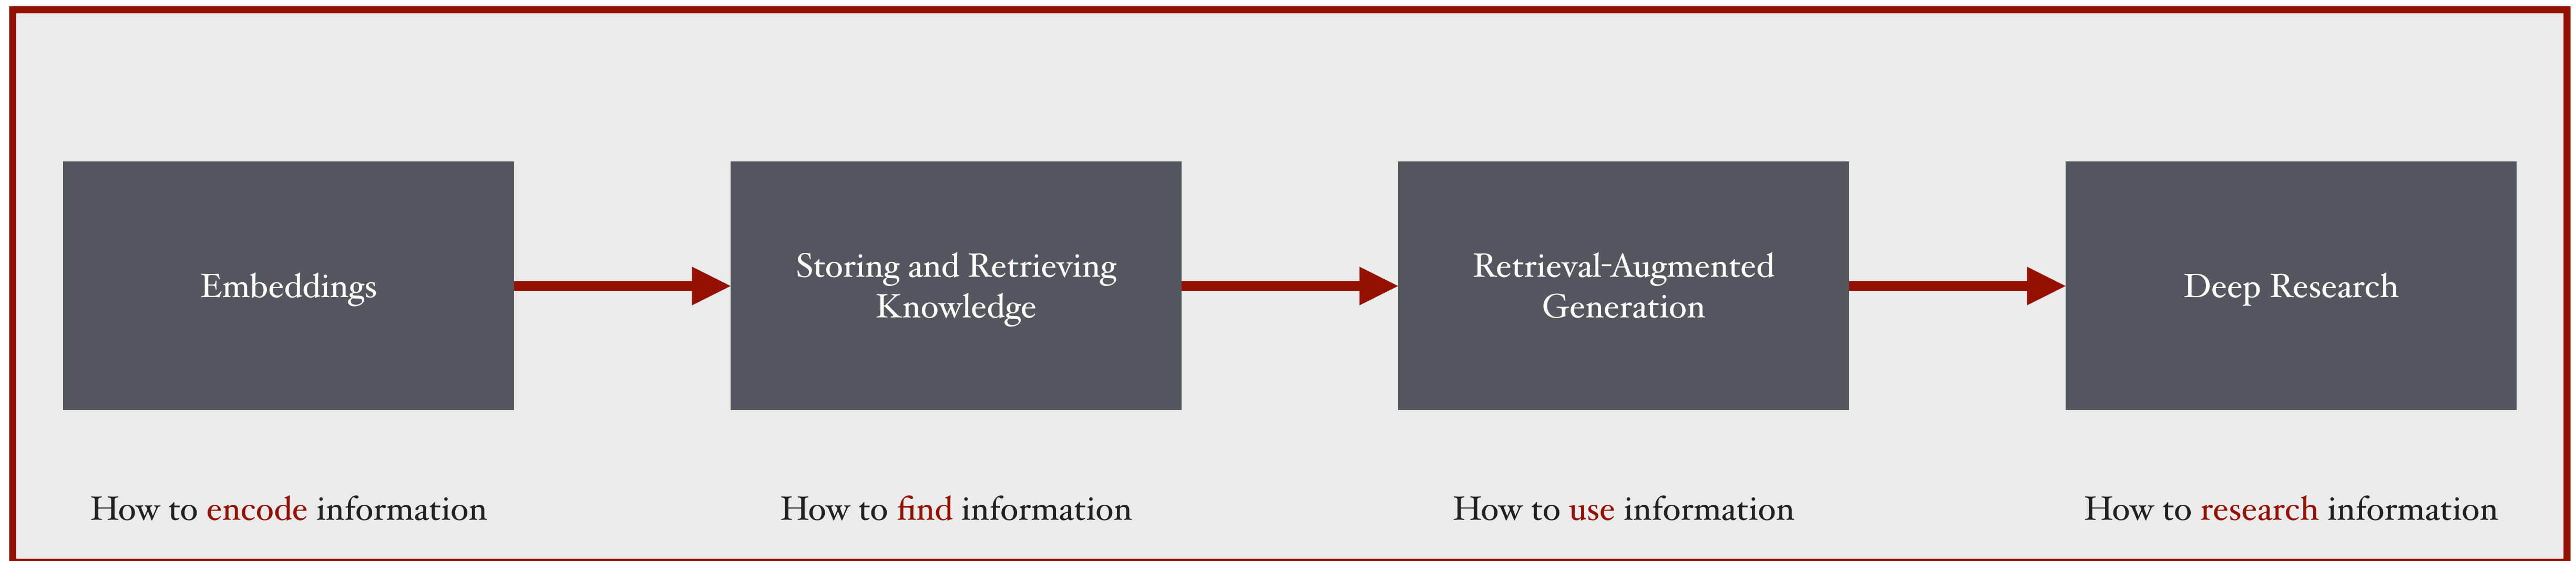
Large Language Model Applications

Task-based conversation

Logistics

- Mail or ask for emergency extensions to hw1 by 5pm today.
- Penalty for late submission is on the syllabus,

For each unexcused day your homework is late, we will subtract 5% from your final grade for the homework. For example, if you submit your homework 1 minute late, and your grade would otherwise be a 97%, it will drop down to 92%. If you submit 26 hours late, and your grade would otherwise be 90%, it will drop down to 80%.



User provides a single input (e.g., prompt, prefix, question).

Model can interact with search engine and other tools.



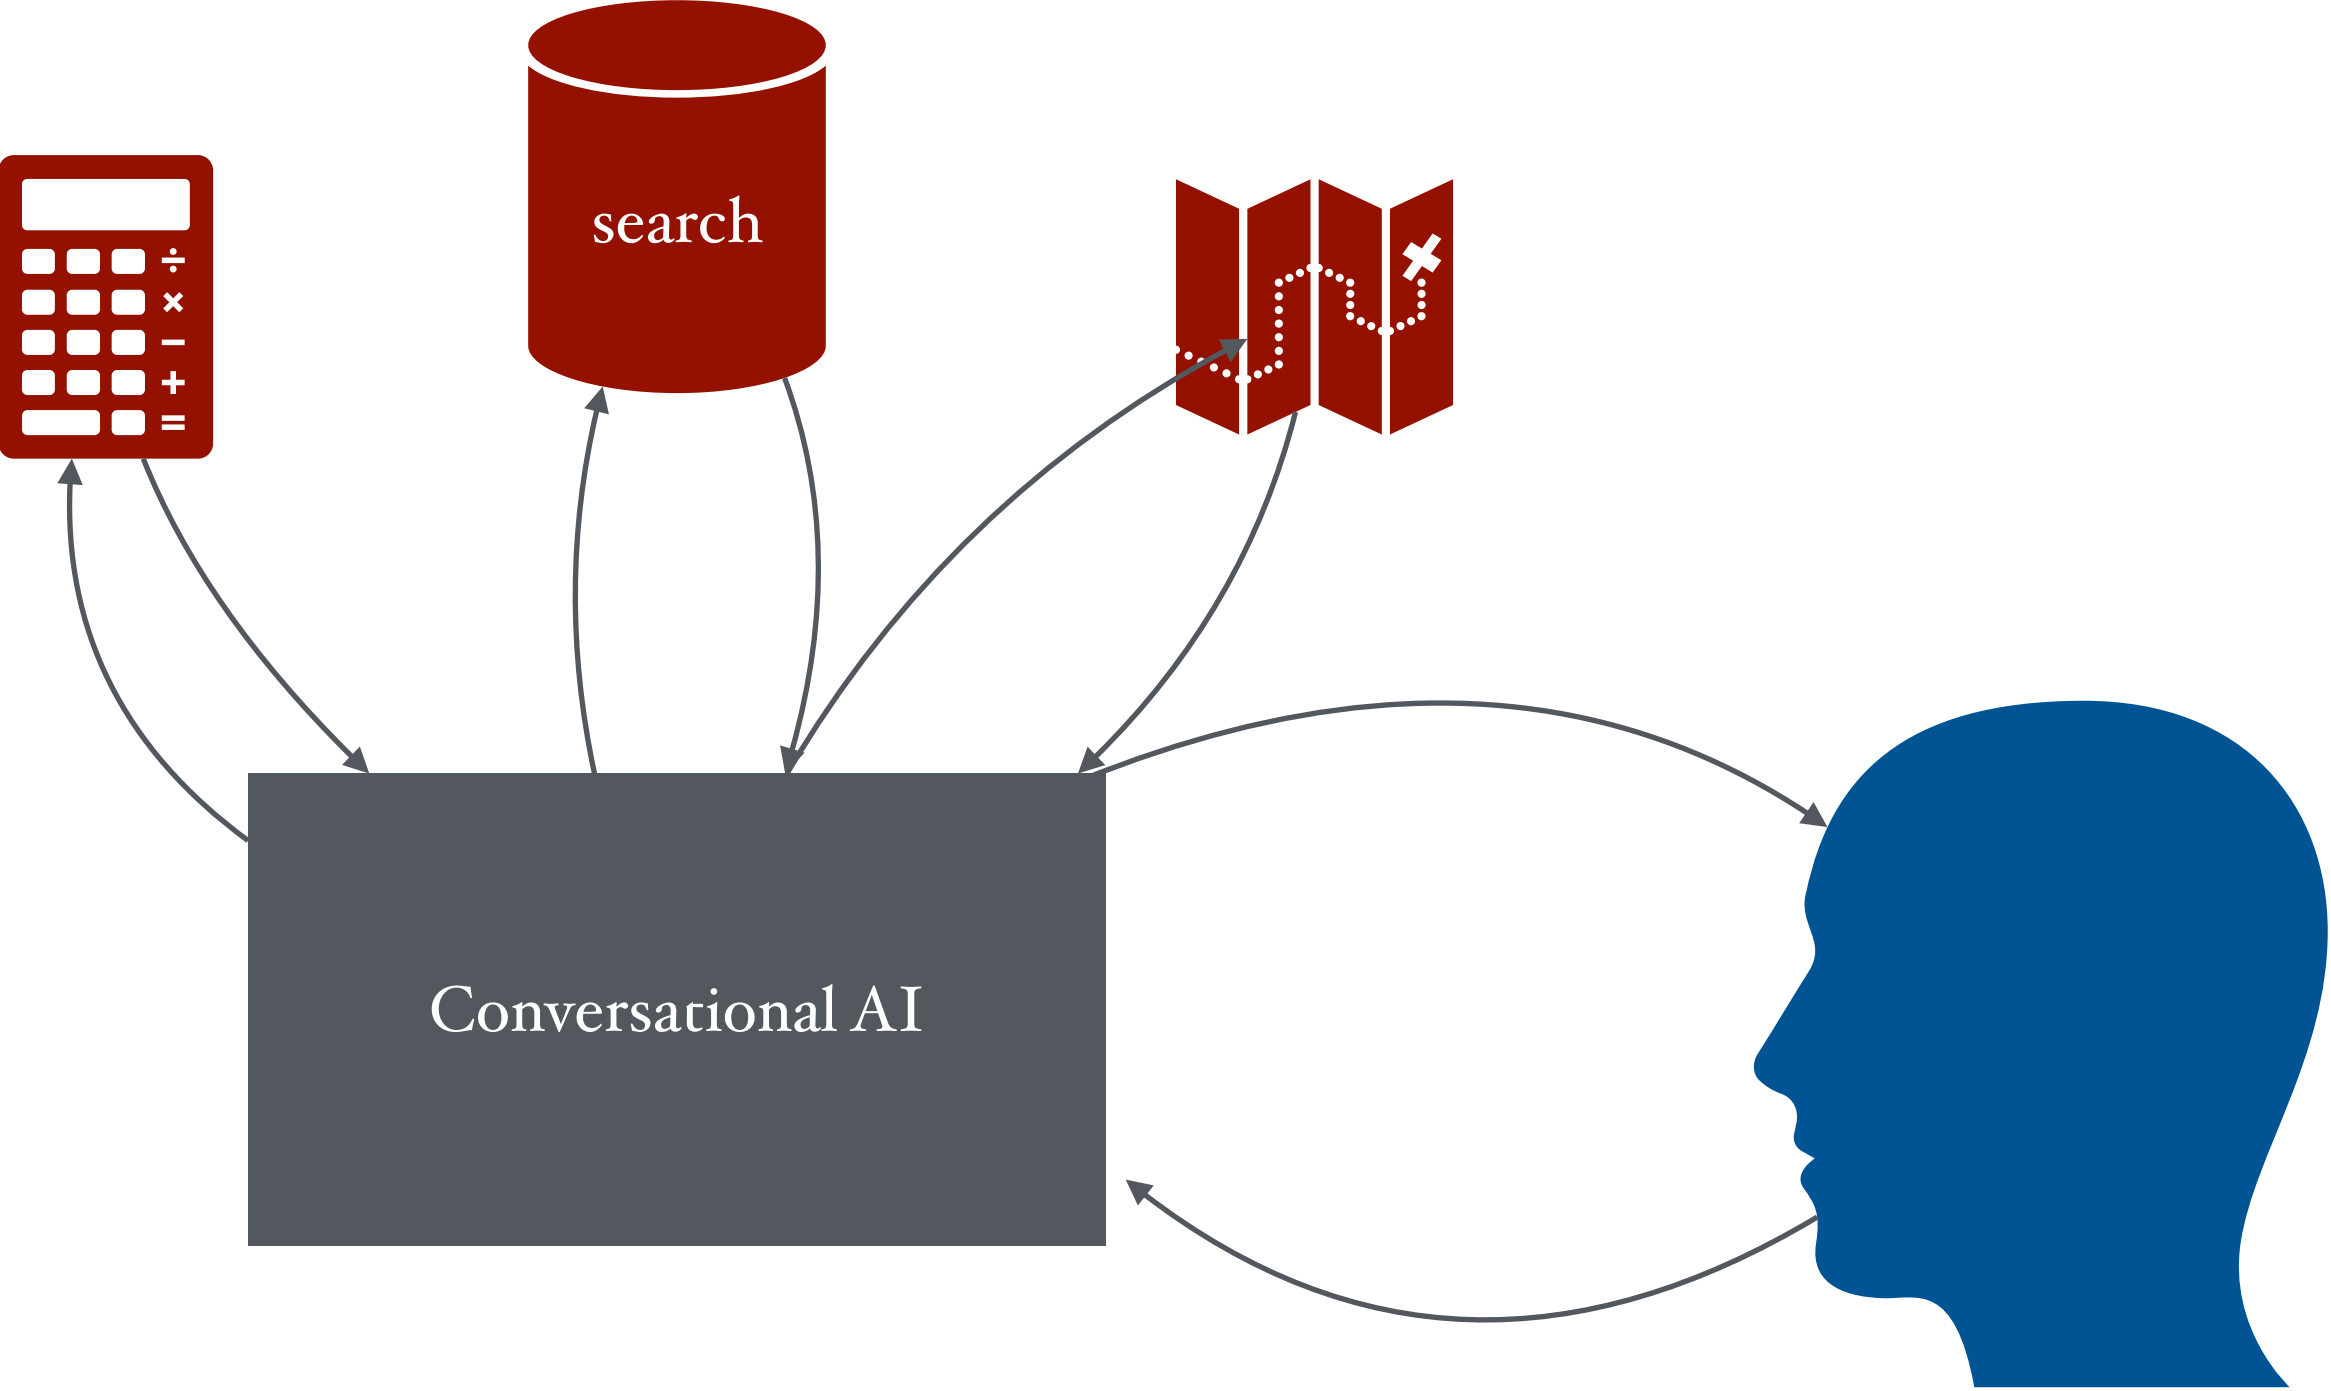
Conversational AI

How to **interact with people**

User provides a **multiple inputs** (e.g., utterance, turn, clarification).

Model can interact with search engine, other tools, **and people**.

tools: usually designed to operate subject to a specification.



people: not designed to operate subject to a specification.

How people use conversational AI...

- LLMs were developed for specific tasks
 - next word prediction
 - question answering
 - summarization
 - translation
 - instruction following
 - Well-defined problems with research on how to evaluate
 - clear problem definitions
 - datasets with reference answers
 - evaluation metrics
 - But is this how people really use conversational AI?
- motivated by funding (e.g., intelligence analysis)*

task-based dialogue

AI as a Tool (revisited)


- AI can be a **tool** to accomplish tasks.
- People use AI to accomplish a broad set of request types in a wide set of domains,
 - travel booking
 - navigation
 - education
 - ...




start
task



interact



finish
task

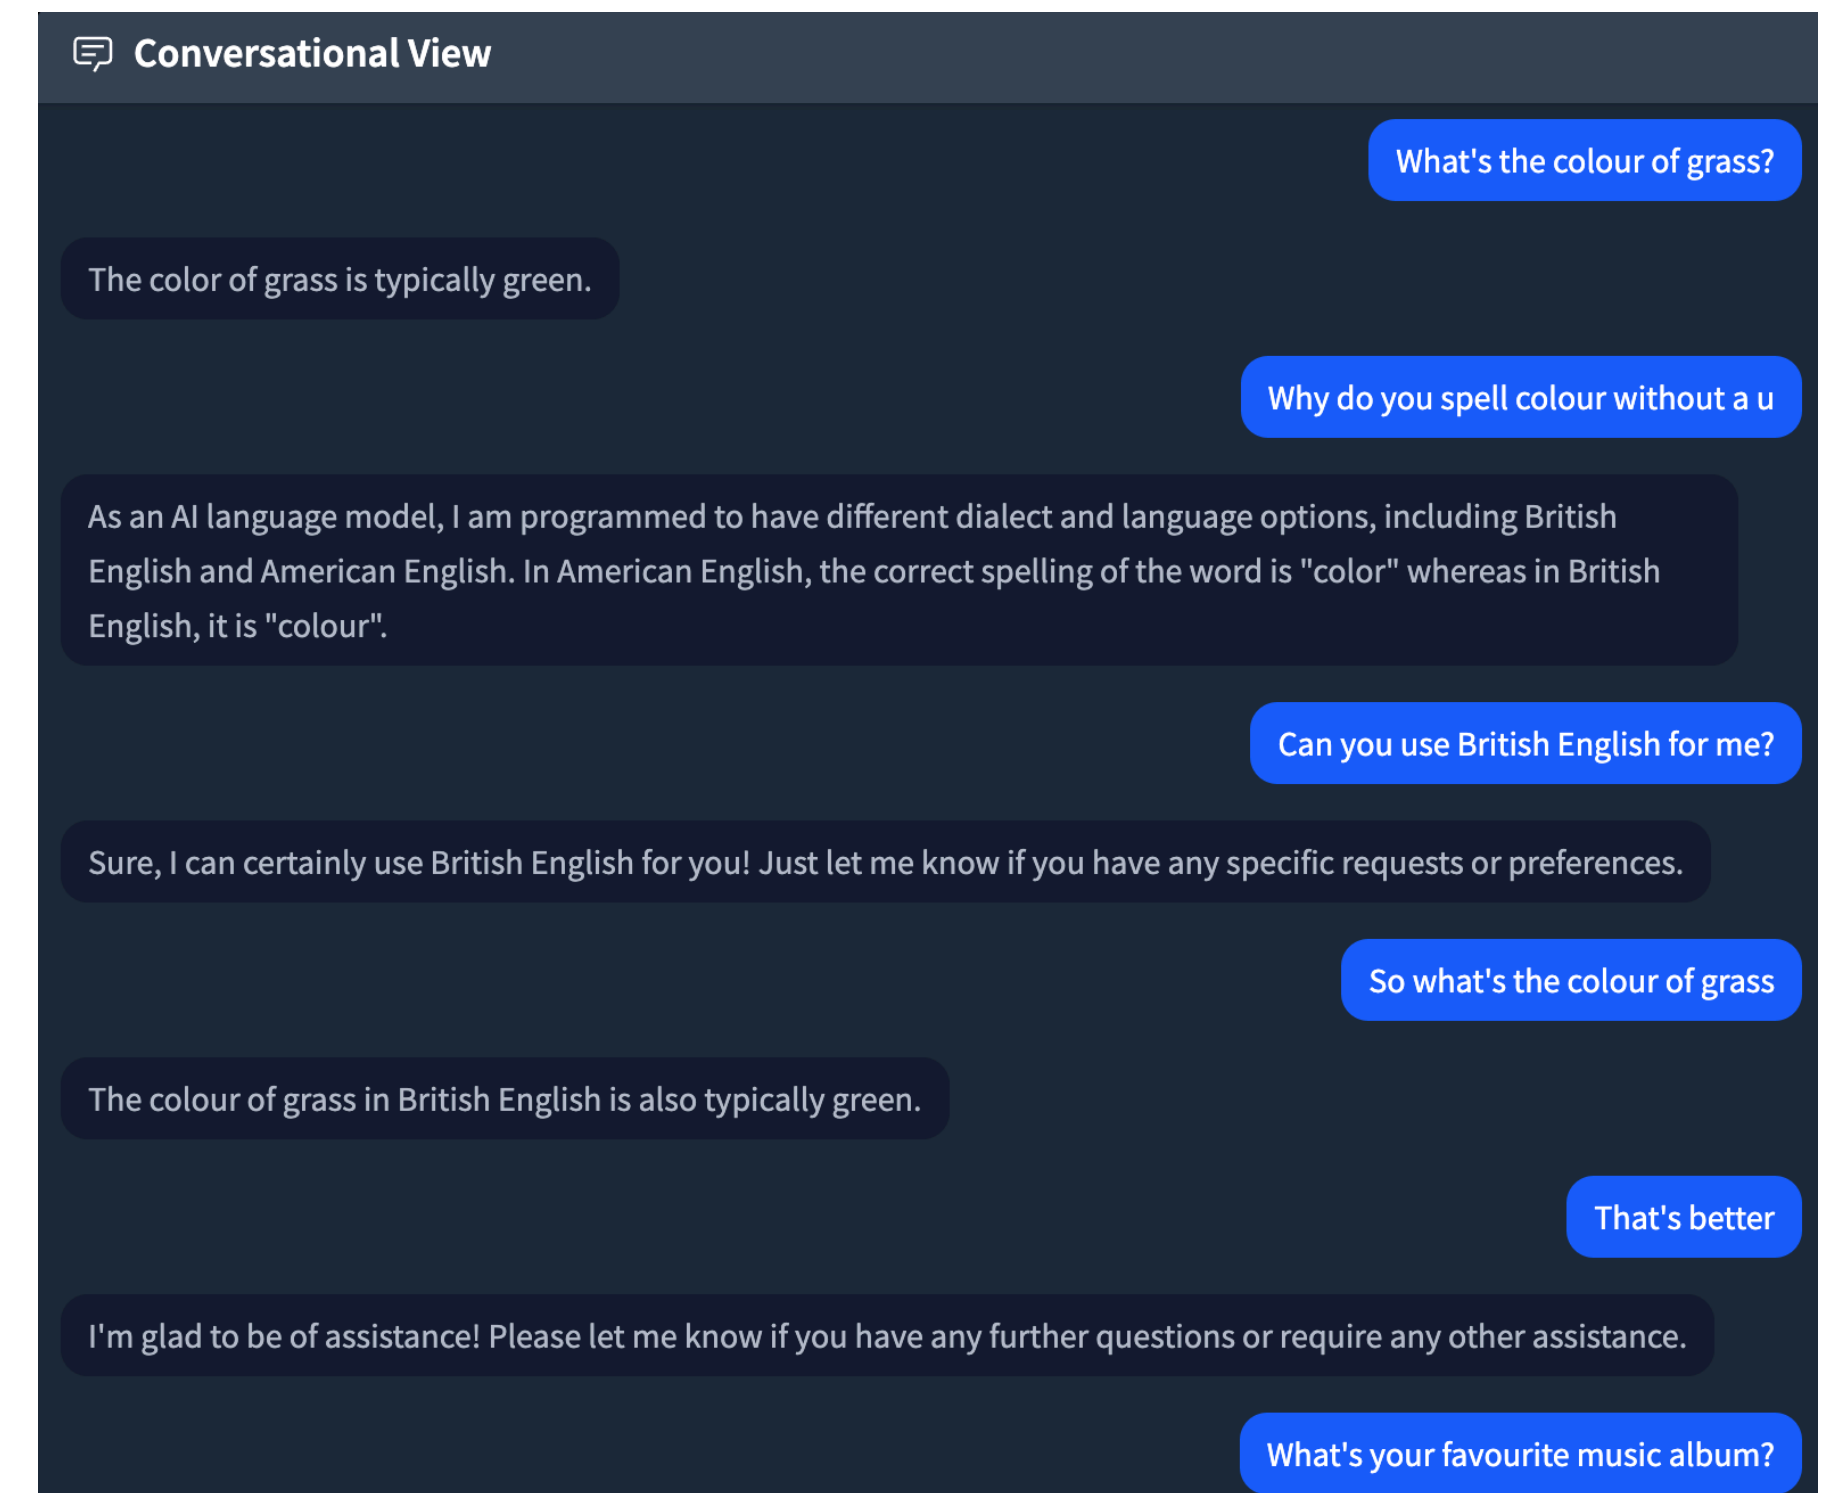


receive
feedback

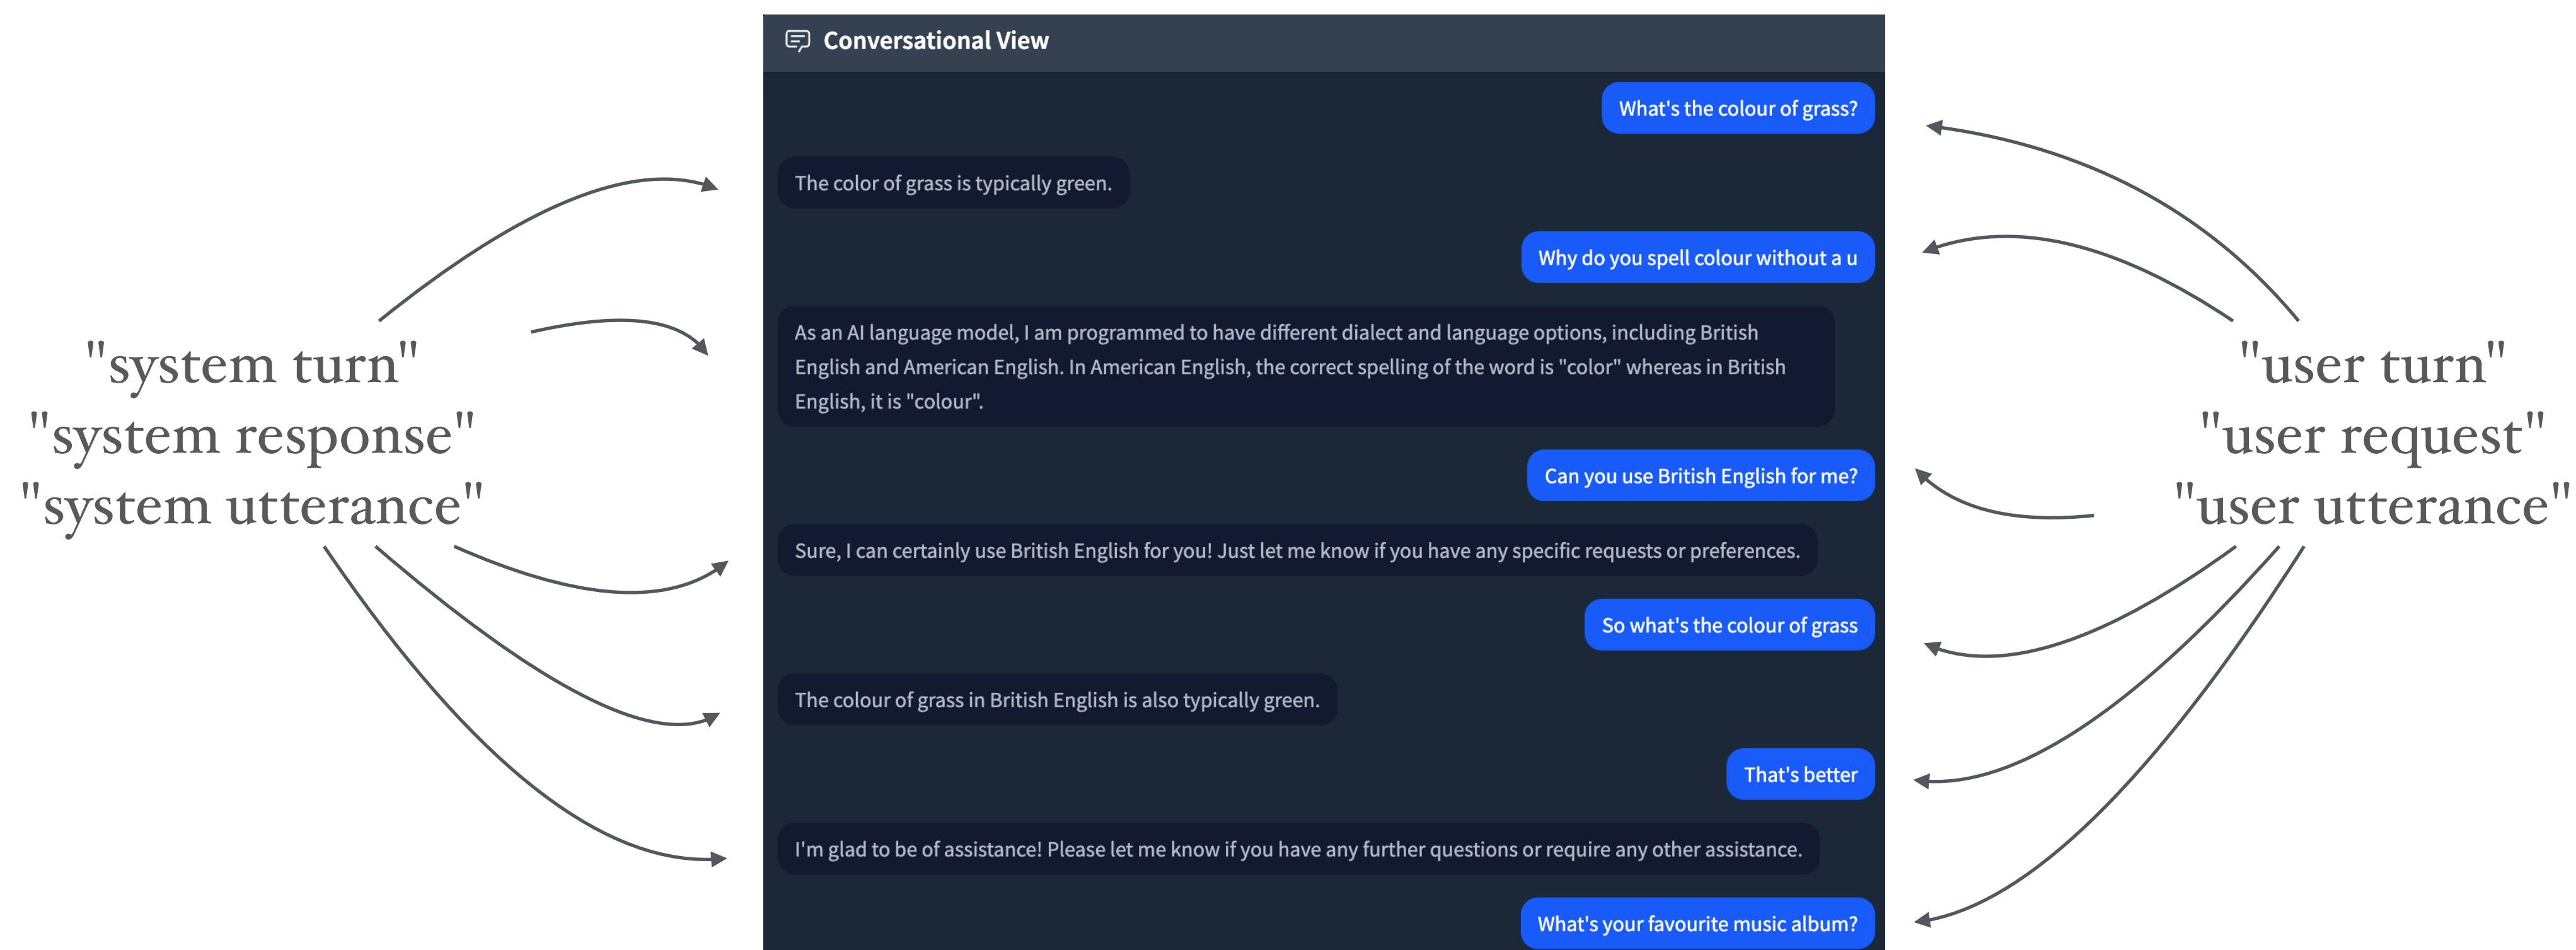


How people use conversational AI...

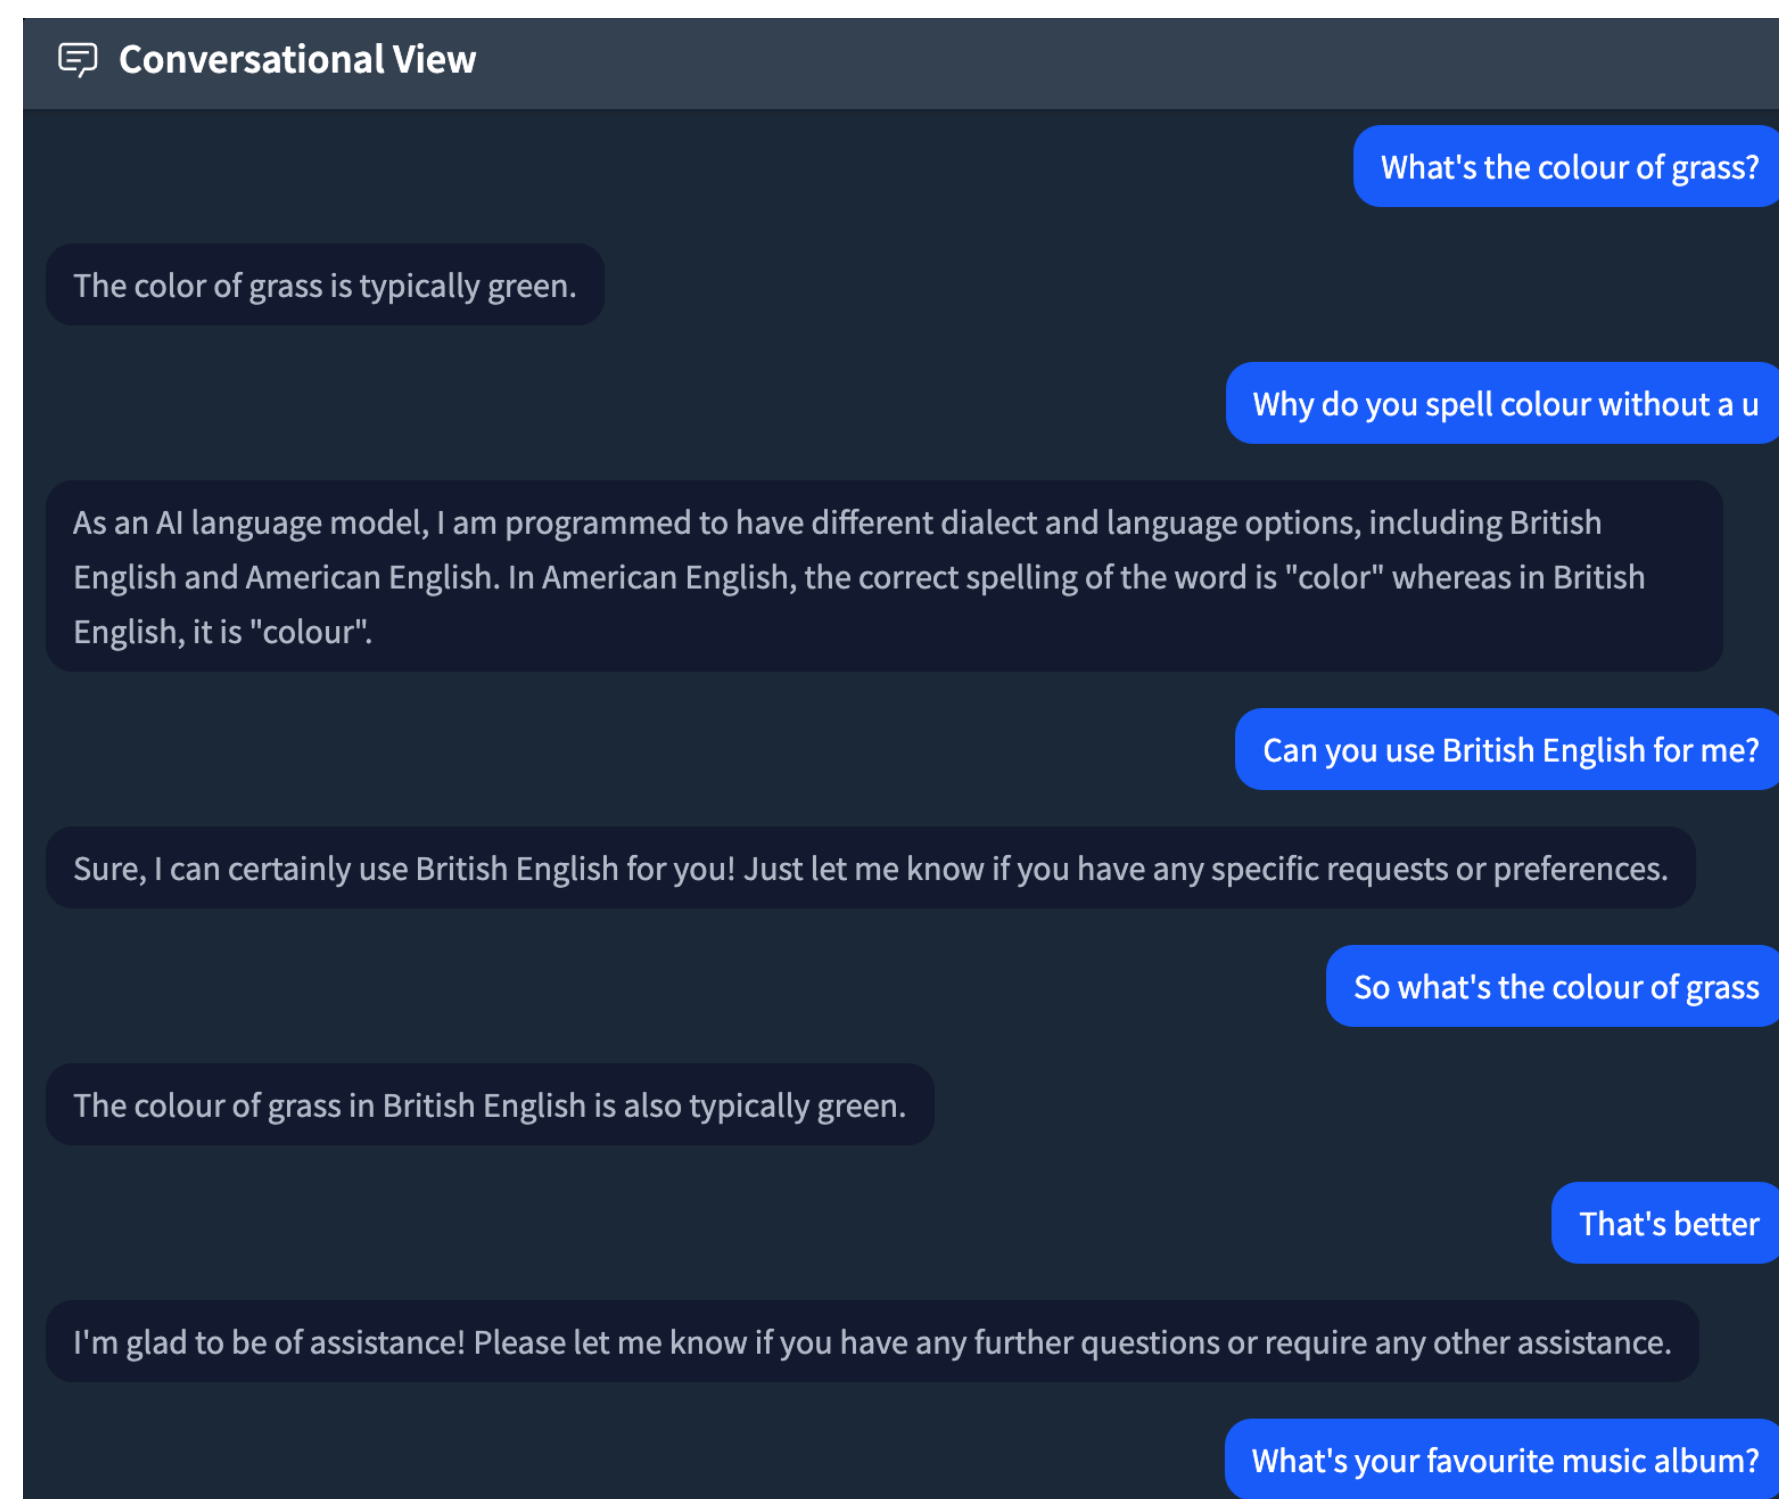
- Before ChatGPT, conversational AI was restricted to academic study with curated data or very specific contexts (e.g., some customer service).
- After ChatGPT, increasing everyday usage of conversational AI.
- Can inspect conversation logs to get a sense of how people are using conversational AI.
- Most data is internal to companies (OpenAI, Anthropic, Google).
- **WildChat:** "offered free access to ChatGPT for online users in exchange for their affirmative, consensual opt-in to anonymously collect their chat transcripts."



How people use conversational AI...



How people use conversational AI...



} dialogue

How people converse...

Speech acts

- Developed in the 1960s by linguist John Searle that can be used to describe turns.
- Assertives
 - Definition: Speech acts that commit the speaker to the truth of a proposition—statements, descriptions, claims, or reports.
 - Examples in task-oriented dialogue:
 - System: "There are 3 hotels available in downtown Boston for your dates"
 - System: "Your flight departs at 3:45 PM from Gate B12"
 - User: "I'm a gold member with your airline"
 - System: "The total cost including taxes is \$247.50"

How people converse...

Speech acts

- Directives
 - Definition: Speech acts intended to get the hearer to perform some action—requests, commands, questions, suggestions.
 - Examples in task-oriented dialogue:
 - User: "Book a table for four people at 7 PM" (command)
 - User: "Can you search for flights under \$300?" (request)
 - System: "Please provide your confirmation number" (request)
 - System: "I recommend selecting seat 12A for extra legroom" (advice)
 - User: "Show me vegetarian options" (command)

How people converse...

Speech acts

- Commissives
 - Definition: Speech acts that commit the speaker to a future course of action—promises, pledges, guarantees.
 - Examples in task-oriented dialogue:
 - System: "I'll send a confirmation email to your inbox within 5 minutes"
 - System: "We guarantee the lowest price or we'll refund the difference"
 - User: "I promise to arrive by 2 PM for check-in"
 - System: "Your prescription will be ready for pickup by 4 PM today"

How people converse...

Speech acts

- Expressives
 - Definition: Speech acts that express the speaker's psychological state or attitude—thanks, apologies, congratulations, welcomes.
 - Examples in task-oriented dialogue:
 - User: "Thank you for finding that flight"
 - System: "I apologize for the confusion about your reservation"
 - User: "Sorry, I meant Thursday, not Tuesday"
 - System: "You're welcome! Is there anything else I can help you with?"
 - System: "Congratulations on your upgrade to premium class!"

How people converse...

Speech acts

- Declarations
 - Definition: Speech acts that bring about a change in the institutional/external state of affairs.
 - Examples in task-oriented dialogue:
 - System: "Your reservation is confirmed" (creates a binding reservation)
 - System: "Your account is now activated" (changes account status)
 - System: "I'm canceling order #12345" (terminates the order)
 - System: "You are now checked in for flight AA502" (changes passenger status)
 - System: "Your appointment is scheduled for March 15th at 2 PM" (establishes the appointment)

How people converse...

Other ontologies

Verbmobil

Tag	Example
THANK	<i>Thanks</i>
GREET	<i>Hello Dan</i>
INTRODUCE	<i>It's me again</i>
BYE	<i>Alright bye</i>
REQUEST-COMMENT	<i>How does that look?</i>
SUGGEST	<i>from thirteenth through seventeenth June</i>
REJECT	<i>No Friday I'm booked all day</i>
ACCEPT	<i>Saturday sounds fine,</i>
REQUEST-SUGGEST	<i>What is a good day of the week for you?</i>
INIT	<i>I wanted to make an appointment with you</i>
GIVE_REASON	<i>Because I have meetings all afternoon</i>
FEEDBACK	<i>Okay</i>
DELIBERATE	<i>Let me check my calendar here</i>
CONFIRM	<i>Okay, that would be wonderful</i>
CLARIFY	<i>Okay, do you mean Tuesday the 23rd?</i>
DIGRESS	<i>[we could meet for lunch] and eat lots of ice cream</i>
MOTIVATE	<i>We should go to visit our subsidiary in Munich</i>
GARBAGE	<i>Oops, I-</i>

Map Task

Tag	Example
INSTRUCT	<i>Go round, ehm horizontally underneath diamond mine</i>
EXPLAIN	<i>I don't have a ravine</i>
ALIGN	<i>Okay?</i>
CHECK	<i>So going down to Indian Country?</i>
QUERY-YN	<i>Have you got the graveyard written down?</i>
QUERY-W	<i>In where?</i>
ACKNOWLEDGE	<i>Okay</i>
CLARIFY	<i>{you want to go... diagonally} Diagonally down</i>
REPLY-Y	<i>I do.</i>
REPLY-N	<i>No, I don't</i>
REPLY-W	<i>{And across to?} The pyramid.</i>
READY	<i>Okay</i>

ATR

Tag	Example
PHATIC	<i>Hello</i>
EXPRESSIVE	<i>Thank you</i>
RESPONSE	<i>That's right</i>
PROMISE	<i>I will send you a registration form</i>
REQUEST	<i>Please go to Kitaoji station by subway</i>
INFORM	<i>We are not giving any discount this time</i>
QUESTIONIF	<i>Do you have the announcement of the conference?</i>
QUESTIONREF	<i>What should I do?</i>
QUESTIONCONF	<i>You have already transferred the registration fee, right?</i>

How people use conversational AI...

- Can inspect usage data to get a sense of how people are using conversational AI.
- Using a collection of public datasets (WildChat, ShareGPT, LMSys), conduct qualitative coding of dialogues for recurring patterns of **request types**.
- Augment with literature-based research on how people use other technologies (e.g., search, content creation).

Mode	Strategy	Request Types
Information Seeking	Retrieval	direct fact question, concept search, refinding request, unknown-item search
	Discovery	topic update, similarity search, rate item(s), perspective seeking
Information Processing & Synthesis	Clarification	explanation request, exemplar request
	Distillation	summarization request, key information identification, information structuring
	Analysis	comparative analysis, qualitative data analysis, quantitative data analysis, evaluative judgment, inference & prediction, hypothetical scenario
Procedural Guidance & Execution	Guidance	how-to instructions, method recommendation, feasibility assessment, error identification
	Execution	error solution, autonomous task completion, logical reasoning, calculation
Content Creation & Transformation	Generation	creative content generation, functional content generation, content extension/insertion
	Modification	editing, translation, paraphrasing, reformatting
Social Interaction	Sociability	social banter, social etiquette, emotional expression
	Shared Understanding	requesting clarification, providing clarification, requesting elaboration, expressing acknowledgment, requesting acknowledgment, conversational convention
Meta-Conversation	System Management	persona directive, stylistic constraint, system performance feedback, regeneration request, continuation request, conversation history query, system information query
	Conversation Management	background information, user-provided content, conversational convention definition, action initiation signal
	Communicative Status	uninterpretable, abandoned, self-talk

How people use conversational AI...

Information seeking

Strategy	Specific Request Type	Example
Retrieval	direct fact question	<i>What is the population of Tokyo?</i>
	concept search	<i>buttermilk pancakes</i>
	refinding request	<i>The musical about trains in love...</i>
	unknown-item search	<i>What's a word for when the world gets hotter?</i>
Discovery	topic update	<i>What's new in AI research?</i>
	similarity search	<i>games similar to minecraft</i>
	rate item(s)	<i>recommend some good restaurants in Mexico City</i>
	perspective seeking	<i>What are different perspectives on climate change?</i>

Information seeking: "user requests to find, access, or explore extant information, casting conversational AI in the familiar role of an information retrieval or recommender system."

How people use conversational AI...

Information Processing & Synthesis

Strategy	Request Type	Example
Clarification	explanation request	<i>How are humanitarian personnel supposed to be legally protected in situations of armed conflict?</i>
	exemplar request	<i>I want you to give me examples of font pairings that work well together.</i>
Distillation	summarization request	<i>...summarize the results from South Korea.</i>
	key information identification	<i>Give me only 10 keywords...for the following text</i>
	information structuring	<i>Organize this into a table.</i>
Analysis	qualitative data analysis	<i>Analyze the text above for style, voice, and tone</i>
	quantitative data analysis	<i>What's the relationship between X and Y in this data?</i>
	evaluative judgment	<i>Would it improve the experience to hook up a subwoofer to the rear of the main seat?</i>
	comparative analysis	<i>Compare the weather in Winona, MN vs. Melbourne, FL.</i>
	inference & prediction	<i>What happens if global temps rise by 2 degrees?</i>
	hypothetical scenario	<i>What if dinosaurs had not gone extinct?</i>

Information Processing & Synthesis: "user requests to understand, analyze, and evaluate information to construct new meaning. The system is cast not merely as an informant but as an active sensemaking partner."

How people use conversational AI...

Procedural Guidance & Execution

Strategy	Specific Request Type	Example
Guidance	how-to instructions	<i>How do I tie a tie?</i>
	method recommendation	<i>What's the best way to learn a new language?</i>
	feasibility assessment	<i>Is it feasible to travel across the country on a bike?</i>
	error identification	<i>Are there grammar errors in this text?</i>
Execution	logical reasoning	<i>What has an eye but cannot see?</i>
	calculation	<i>What is the square root of 144?</i>
	error solution	<i>...this is the error I get now...revise the script correcting it.</i>
	autonomous task completion	<i>Order me a pizza.</i>

Procedural Guidance & Execution: "user requests that elicit procedural knowledge: 'how to do something, methods of inquiry, and criteria for using skills, algorithms, techniques, and methods'."

How people use conversational AI...

Content Creation & Transformation

Strategy	Specific Request Type	Example
Generation	creative content generation	<i>Come up with a poem about love for a cat.</i>
	functional content generation	<i>Generate python code to sort a list.</i>
	content extension/insertion	<i>Add a conclusion to this paper.</i>
Modification	editing	<i>Proofread this essay for clarity.</i>
	translation	<i>Translate this article into Hindi.</i>
	paraphrasing	<i>Restate the text...</i>
	reformatting	<i>Convert this text into a table.</i>

Content Creation & Transformation: "user requests to generate novel material or modify existing content."

How people use conversational AI...

Social Interaction

Strategy	Specific Request Type	Example
Sociability	social banter	<i>Tell me a joke. / Let's chat.</i>
	emotional expression	<i>That's a scary thought. I hope you're not worried about that.</i>
	social etiquette	<i>Hello, thank you, please.</i>
Shared Understanding	requesting clarification	<i>What did you mean by that? I'm confused.</i>
	providing clarification	<i>No, I meant the book, not the movie.</i>
	requesting elaboration	<i>Tell me more about that second point.</i>
	expressing acknowledgment	<i>Okay, got it. / Understood.</i>
	requesting acknowledgment	<i>Do you understand the instructions?</i>
	conversational convention	<i>'Y' / (In a text adventure game) Go north. / Take the sword.</i>

Social Interaction: "the foundational layer of sociality that enables and shapes instrumental tasks in human-AI dialogue."

How people use conversational AI...

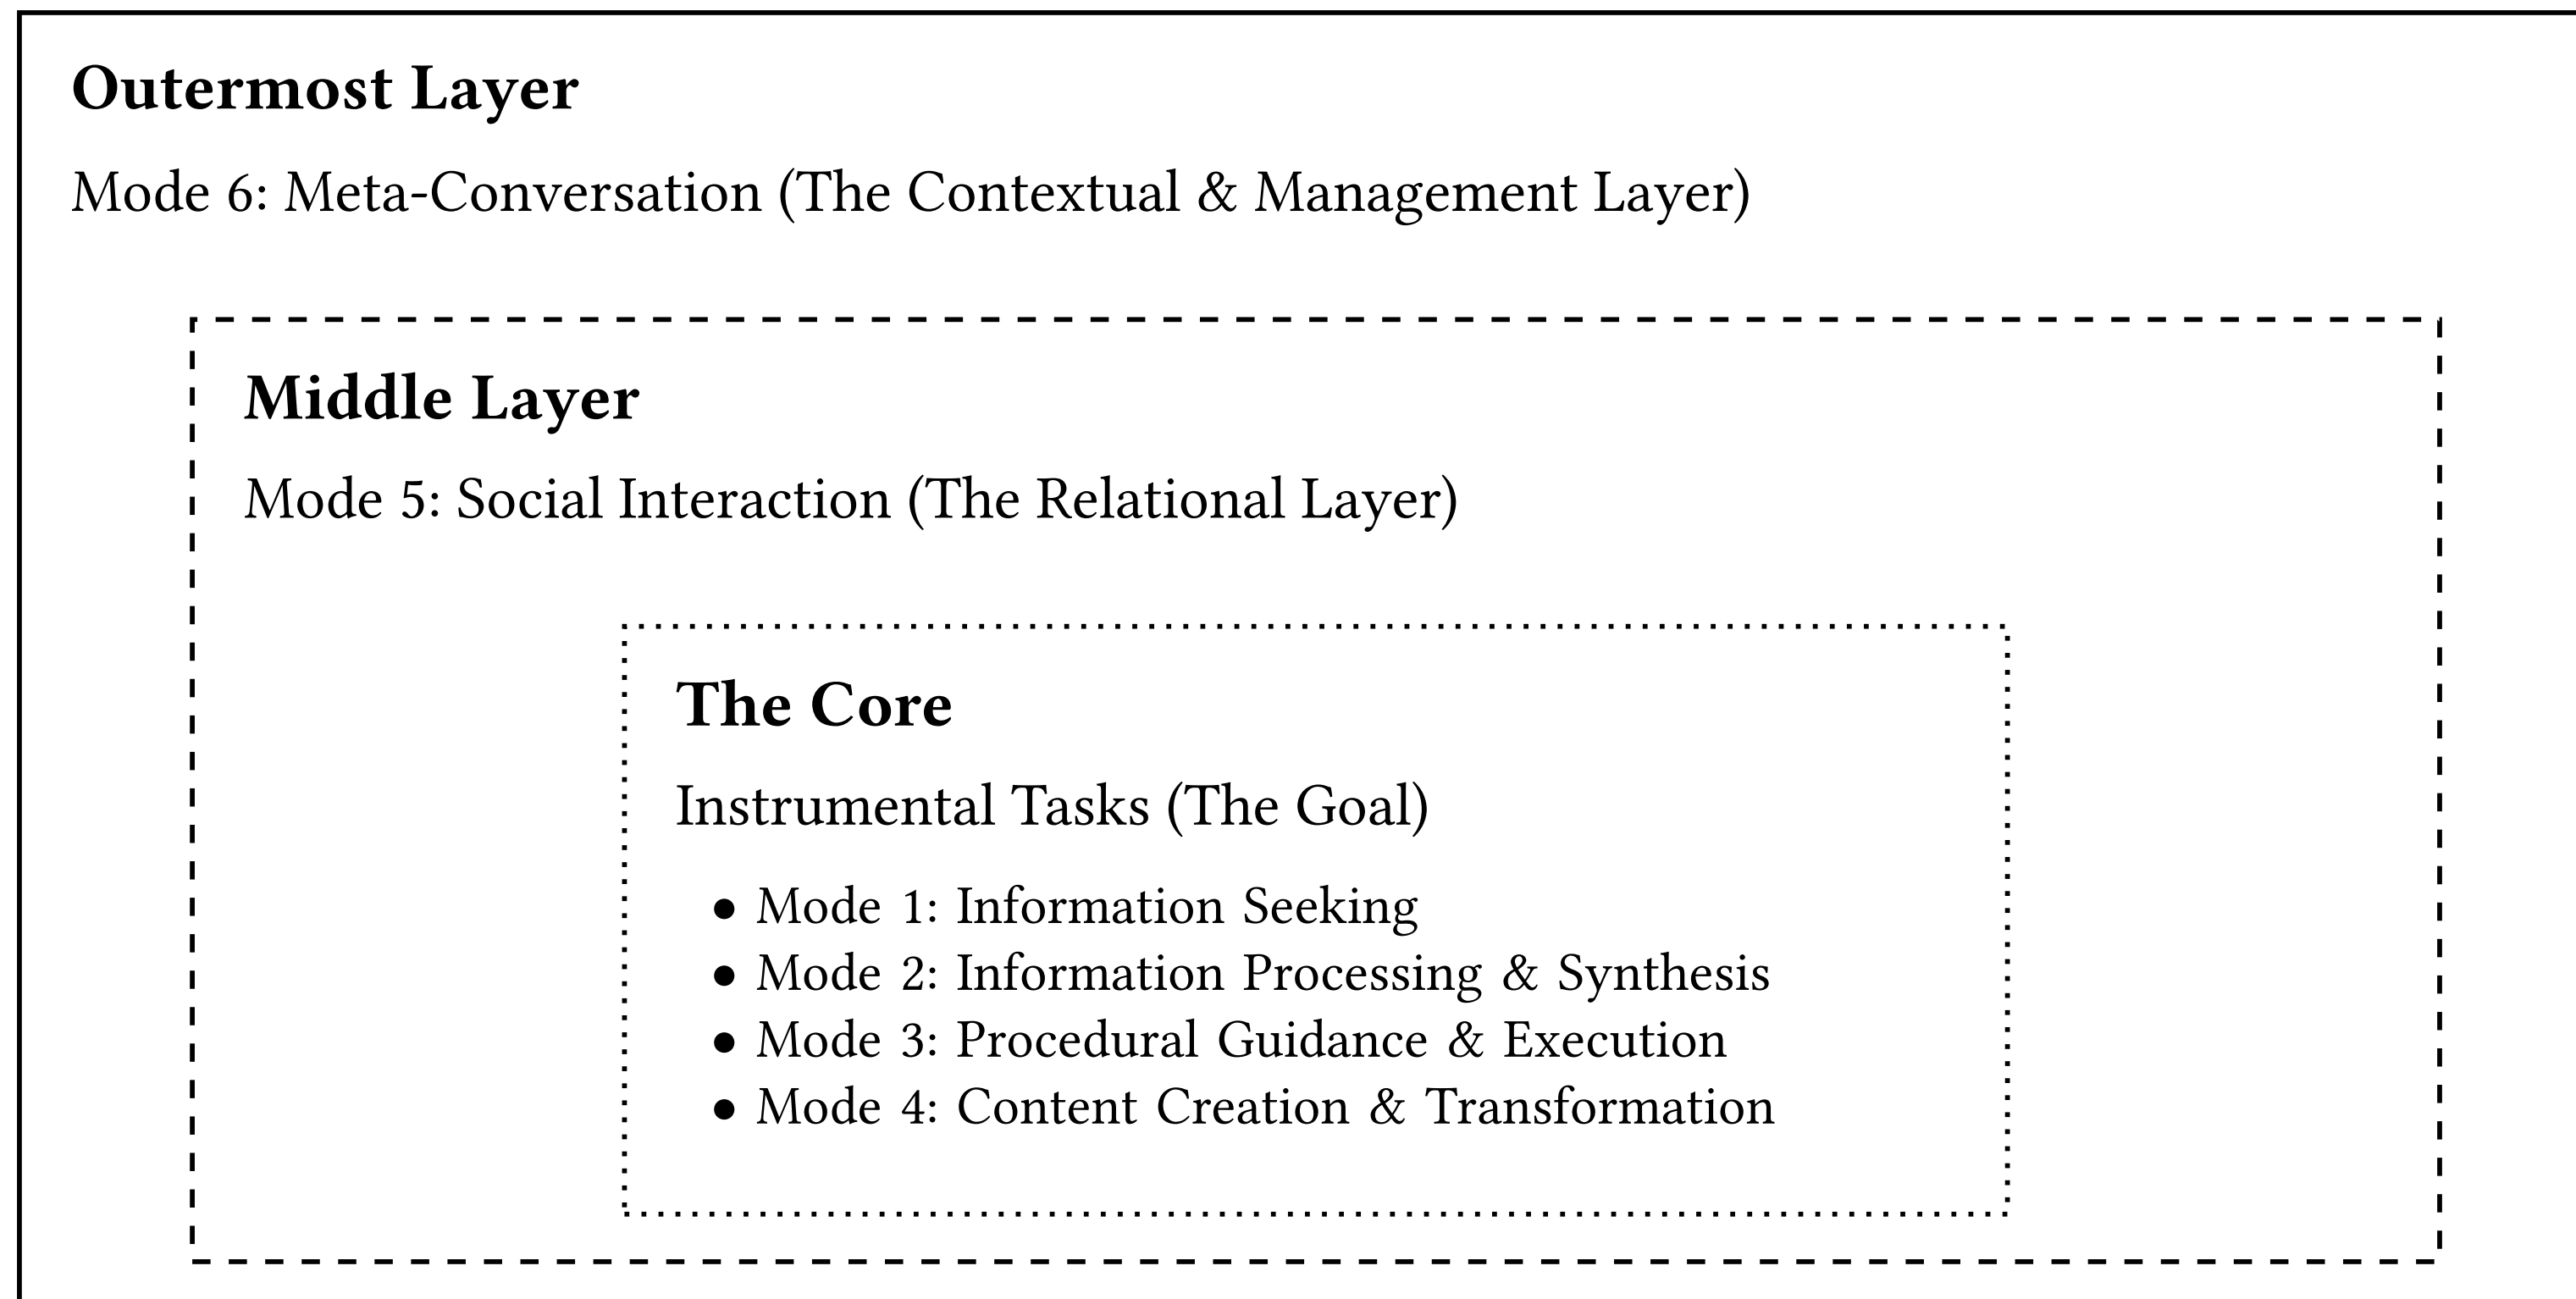
Meta-Conversation

Strategy	Specific Request Type	Example
System Management	persona directive	<i>Act as a helpful assistant. / You are a pirate.</i>
	stylistic constraint	<i>Explain like I'm 5. / Use bullet points.</i>
	system performance feedback	<i>That answer was wrong. / Your last response was too formal.</i>
	regeneration request	<i>Try that again, but in a different style.</i>
	continuation request	<i>Keep going. / More. / Finish the list.</i>
	conversation history query	<i>What did you say earlier about sharks? / Let's get back to my first point.</i>
	system information query	<i>Can you access the internet? / What is your knowledge cutoff?</i>
Conversation Management	background information	<i>I'm a teacher planning a lesson. / He took me to dinner last night. / It's 2025.</i>
	user-provided content	<i>[User pastes a long article to be summarized.]</i>
	conversational convention definition	<i>When I say 'Y,' it means 'yes'.</i>
	action initiation signal	<i>Okay, here is the text you need to analyze:</i>
Communicative Status	uninterpretable	<i>asdfasdfjkl</i>
	abandoned	<i>Could you please... oh, never mind</i>
	self-talk	<i>Okay, so where did I put that file... [to self]</i>

Meta-Conversation: "user actions that manage the interaction itself, rather than performing an instrumental task."

How people use conversational AI...

Taxonomy of User Needs and Actions



How people use conversational AI...

Example dialogue

USER: I have an ETL script on AWS Glue platform. How does the script know which file is the input file to process the data?

The user begins by **background information** (Mode 6) with an **explanation request** (Mode 2). They are not asking for a simple fact but for an understanding of an underlying process.

AI: In AWS Glue, input files are defined...

USER: Try writing an ETL script that describes what you just said?

The user then requests **functional content generation** (Mode 4) that may also serve as an **exemplar request** (Mode 2), asking the system to make the abstract explanation concrete through code.

AI: Okay, let's look at the following example. This script will read...

USER: What happens when running an ETL job on AWS Glue where the input file does not exist?


The user makes an **inference & prediction** request (Mode 2), posing a 'what if' question asking about the potential outcome or consequence ("What happens when") of a specific technical action ("running a job on AWS Glue") under a given condition ("where the input file does not exist").

AI: AWS Glue Crawler is an automated data extraction...

poll: think about how you use
LLMs...

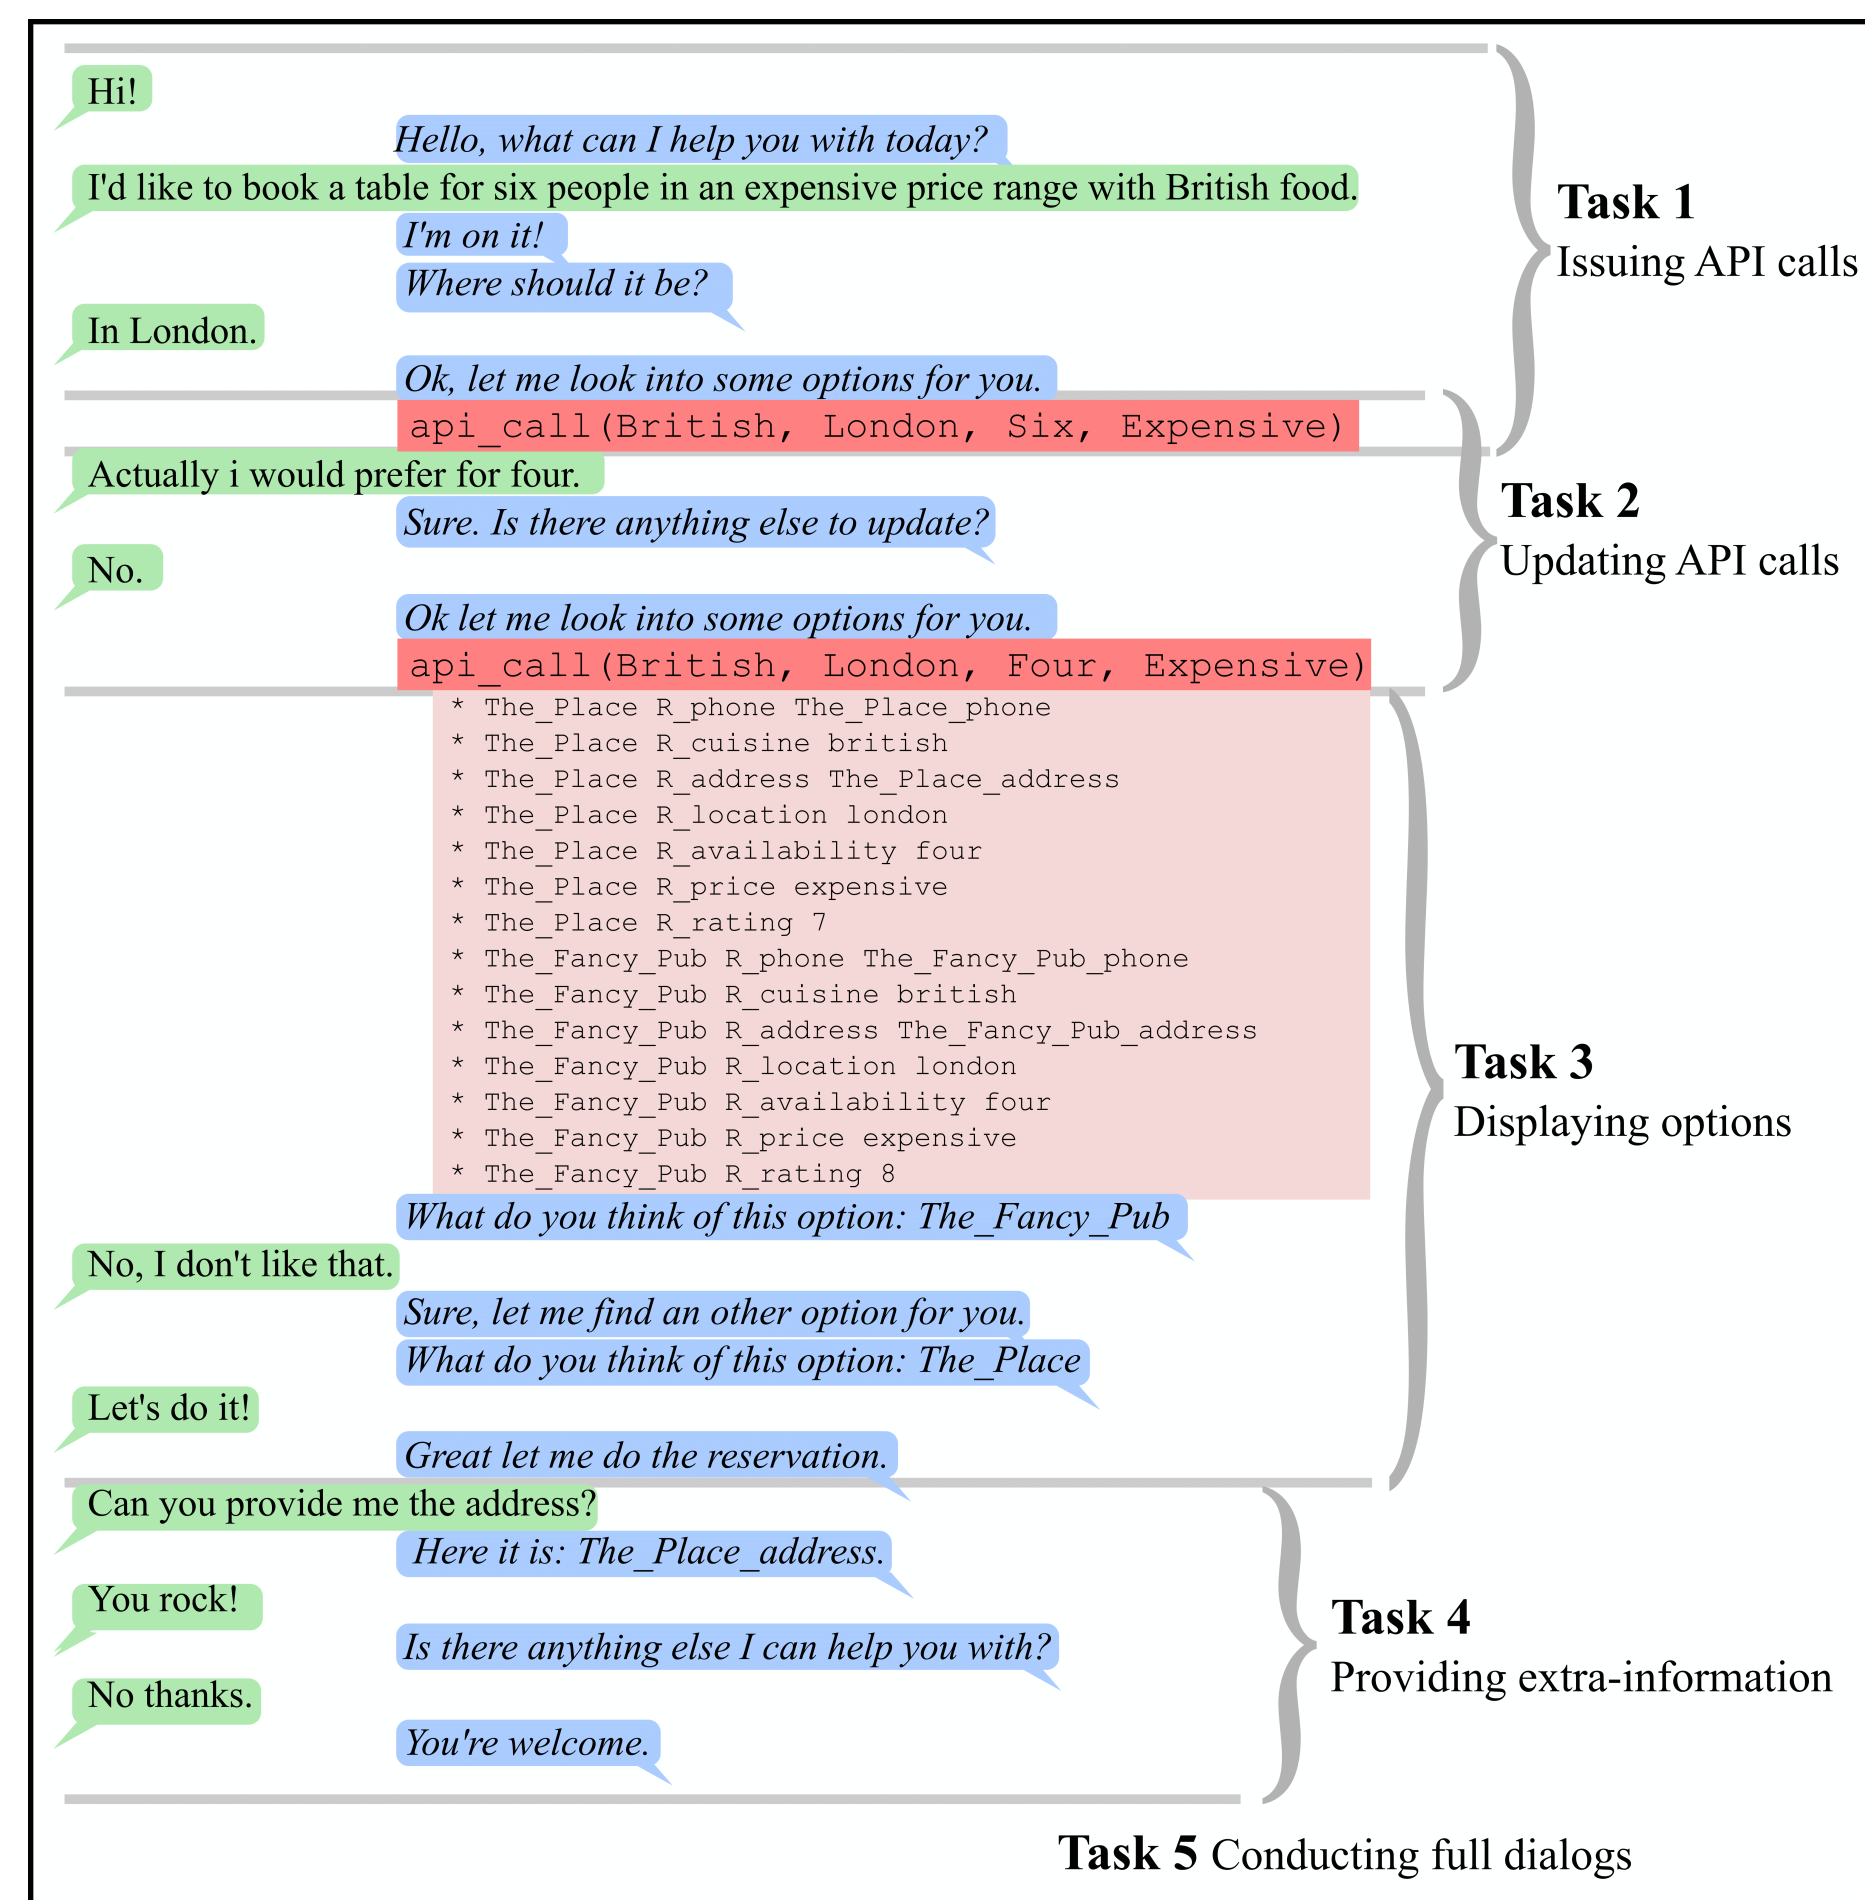
poll: look at a recent conversation
you've had with an LLM...

How people use ChatGPT...



How people use Claude...

Sub-tasks



AI as a Tool (revisited)


- AI can be a **tool** to accomplish tasks.
- People use AI to accomplish a broad set of request types in a wide set of domains,
 - travel booking
 - navigation
 - education
 - ...
- A task-based conversational system usually focuses on a **verifiable** task,
 - plane ticket bought
 - problem resolved
 - ...




start
task



interact



finish
task



receive
feedback



How people use conversational AI...

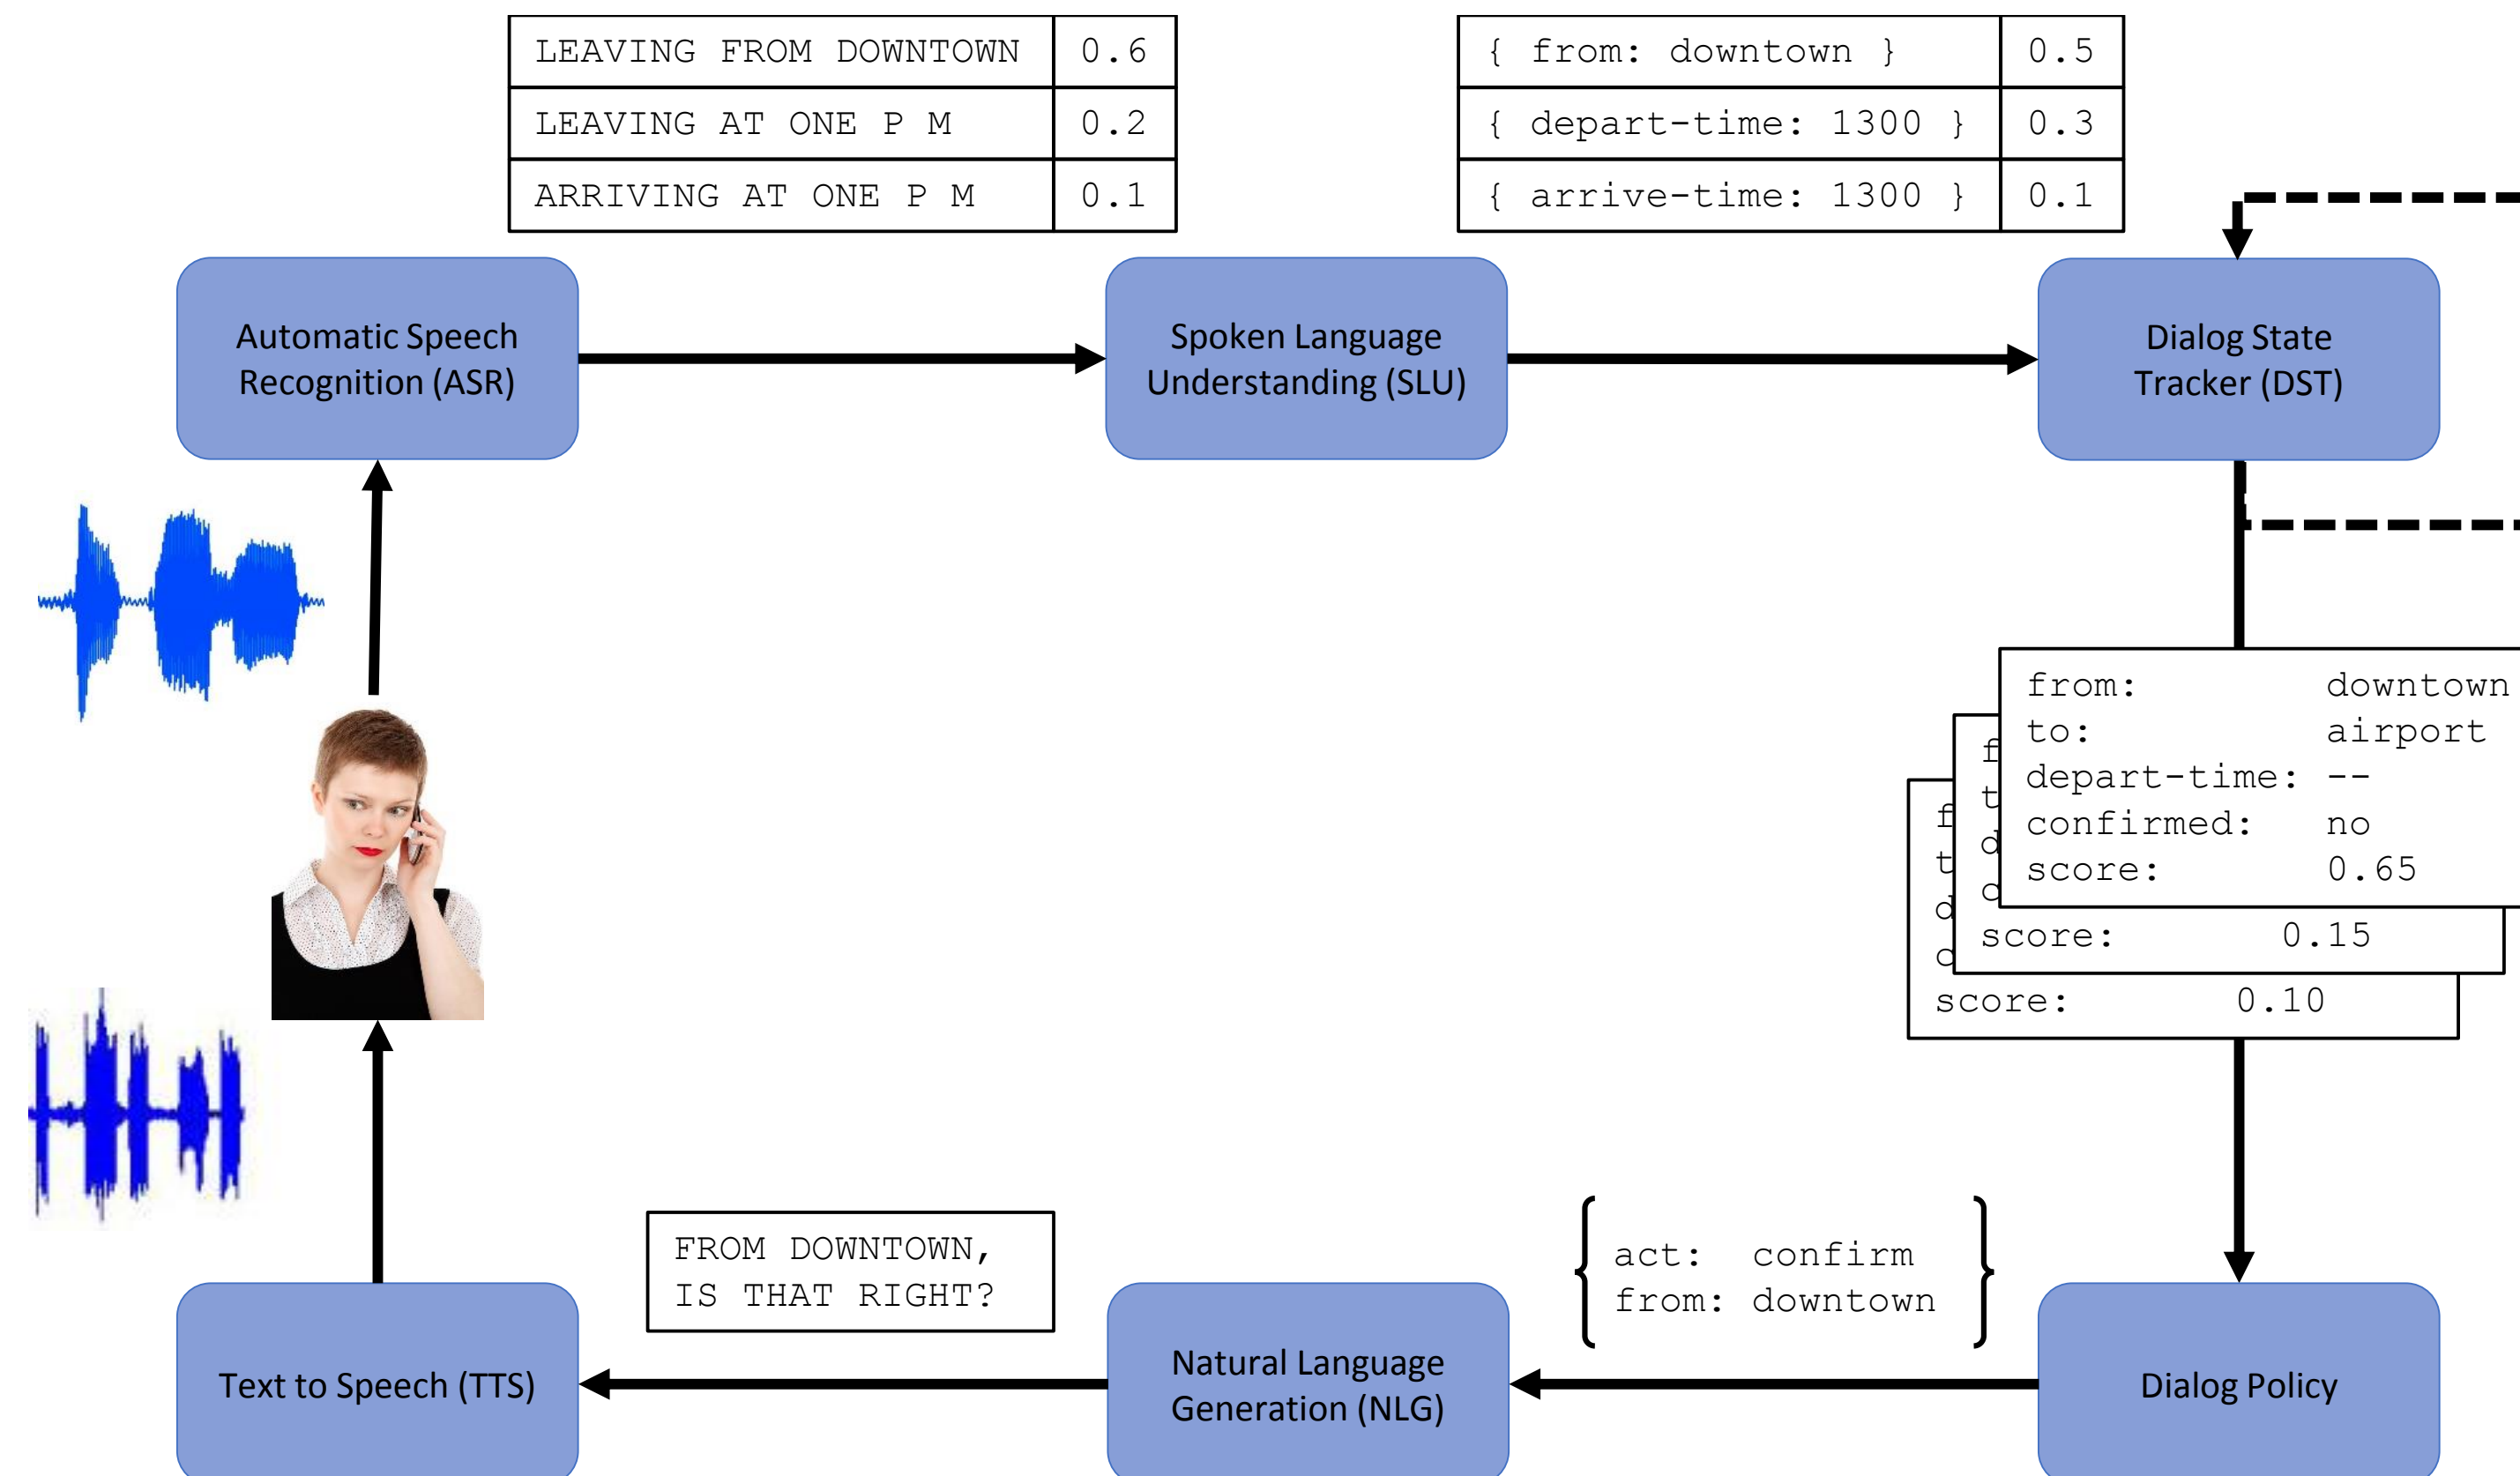
A1: Hello. You've reached the [Communicator].
Tell me your full name
U2: hi i'd like to fly to seattle tuesday morning
A3: Traveling to Seattle on Tuesday, August
eleventh in the morning. Your full name?
U4: [John Doe]
A5: Hello, [John], I'm glad to hear from you
again. You can now ask me for help.
A6: OK. Let's see I have a United flight
departing Pittsburgh at ten forty five a.m.
arriving into Seattle at one fifty seven p.m.
Is that OK?
U7: Yes
A8: Will you return to Pittsburgh from Seattle?
U9: no i will go to Los Angeles
A10: Traveling to Los Angeles. on what date
U11: Sorry can we go back to talk about the
previous leg again
A12: For the first leg, you are on a United flight
1446 departing Pittsburgh at ten forty five
A.M. tomorrow, arriving in Seattle at one
fifty seven P.M. Do you want to take that?
U13: actually i want to leave on wednesday
A14: a flight on wednesday.
A15: I have a flight departing Pittsburgh at ten
forty five a.m., arrives Seattle at one fifty
seven p.m. Is that OK?
U16: Yes
A17: OK and on what day do you want to depart
Seattle?
U18: next Monday
A19: Traveling on Monday, August sixteenth.

M: How may I help you?
U: *Can you tell me how much it is to Tokyo?*
M: You want to know the cost of a call?
U: *Yes, that's right.*
M: Please hold on for rate information.

M: How may I help you?
U: *I tried calling my Aunt and got a wrong num-
ber.*
M: You need a billing credit.
U: *(silence)*
M: Please speak the number that you dialed.
U: *908 582 2778*
M: Was the call billed to the phone that you're
calling from?
U: *Yes it was.*
M: You will be given credit, thank you for calling.

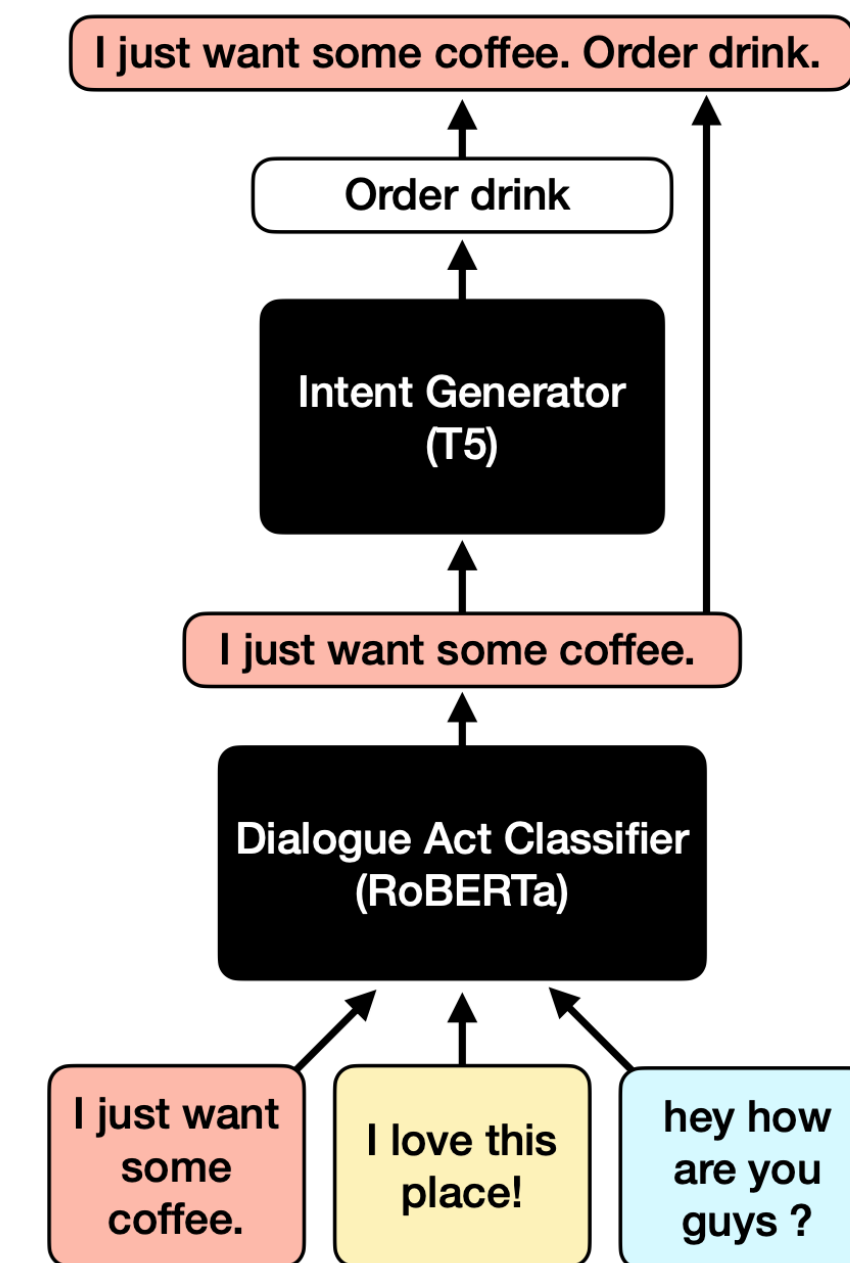


Architecture of a dialogue system



Detecting task intent

- Even a very specific application like a travel agent, can support multiple tasks.
- Answers the question, "what is the user currently trying to use the system for?"
- Once we know the task intent, the system knows,
 - how to act with the user
 - how to interpret user utterances
 - what "good" looks like

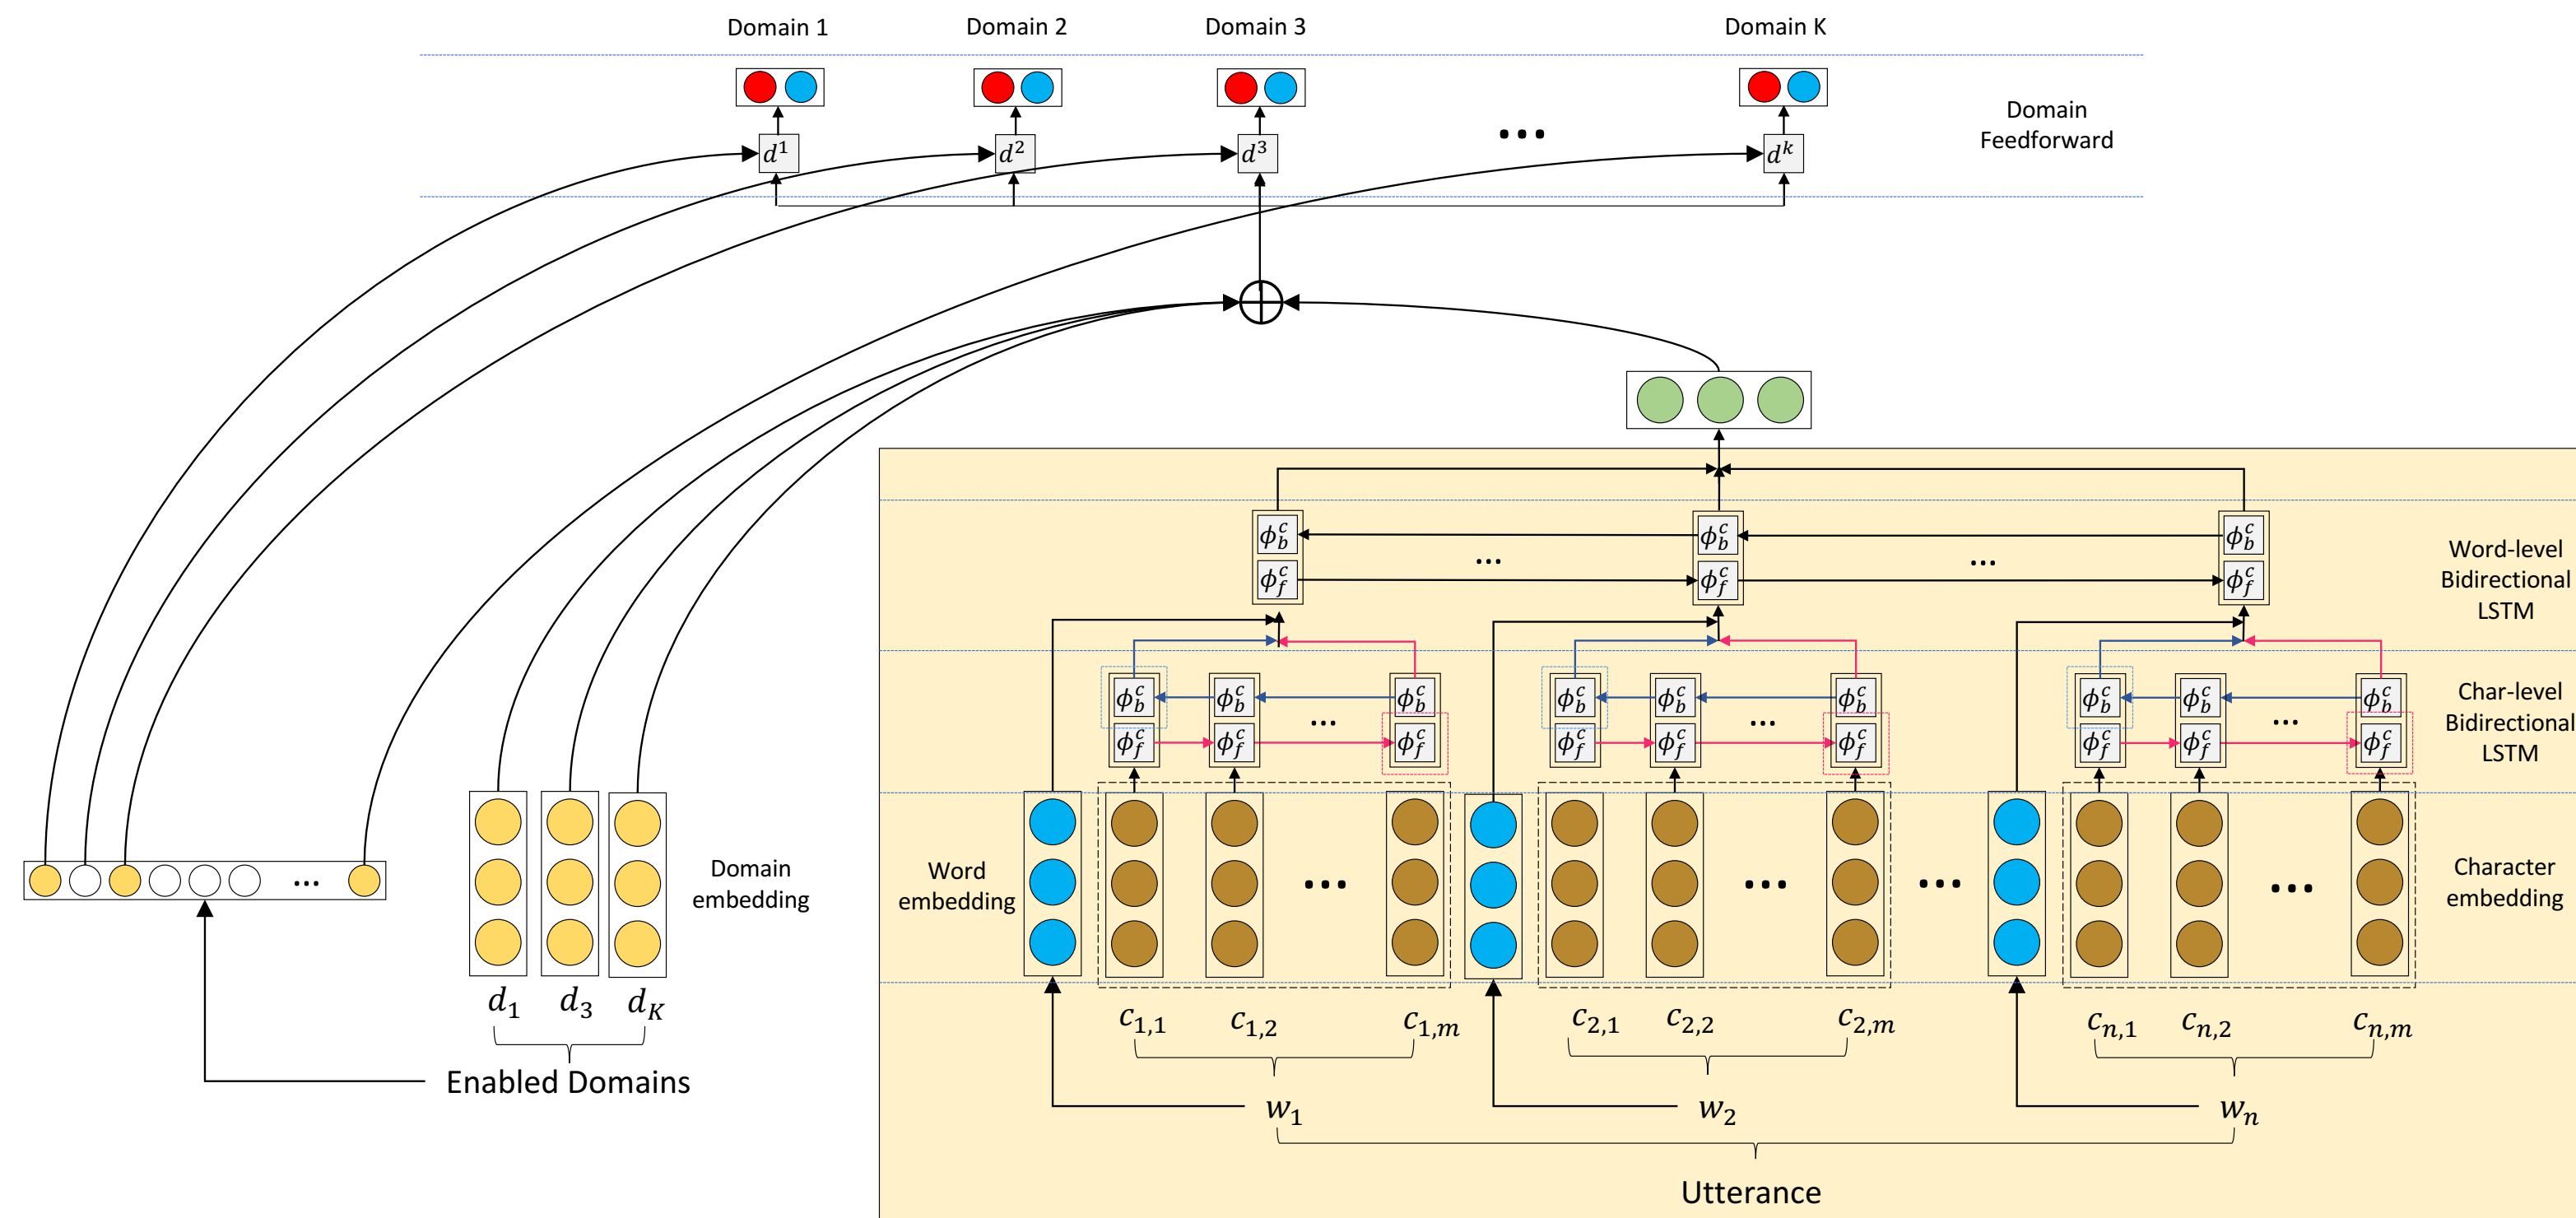


intent detection stack [Mueller et al., 2022]



Detecting task intent

Amazon Alexa domain classifier



Slot-filling approach to task-based dialogue

- A task can often be decomposed in a set of information that, once collected, a system can autonomously complete.
- When booking a flight, once we have the user's,
 - name
 - dates
 - itinerary
 - preference

then we can book a flight...

- Assumes that we have access to domain knowledge about the task.

Slots	Fillers	Servants	Demons
Dialog			
(1) CLIENT	Person	Create	Link to TRAVELLER
(2) NOW	Date	GetDate	
(3) TOPIC	Trip Specification	Create	
TripSpecification			
(4) HOMEPORT	City	Default— Palo Alto	
(5) FOREIGNPORT	City		Link to OUTWARDLEG, AWAYSTAY, INWARDLEG
(6) OUTWARDLEG	TripLeg	Create	
(7) AWAYSTAY	PlaceStay		
(8) INWARDLEG	TripLeg	Create	
TripLeg			
(9) FROMPLACE	City	FindFrom HOMEPORT	
(10) TOPPLACE	City	AskClient	
(11) TRAVELDATE	Date	AskClient	
(12) DEPARTURESPEC	TimeRange	AskClient	Proposc-Flight-By-Departure
(13) ARRIVALSPEC	TimeRange		Propose-Flight-By-Arrival, Link to DEPARTURESPEC
(14) PROPOSEDFLIGHTS	(SetOfFlight)		
(15) FLIGHTCHOSEN	Flight	Ask Client	
(16) TRAVELLER	Person	Ask Client	

frames for a travel booking task [Bobrow et al., 1977]

Slot-filling approach to task-based dialogue

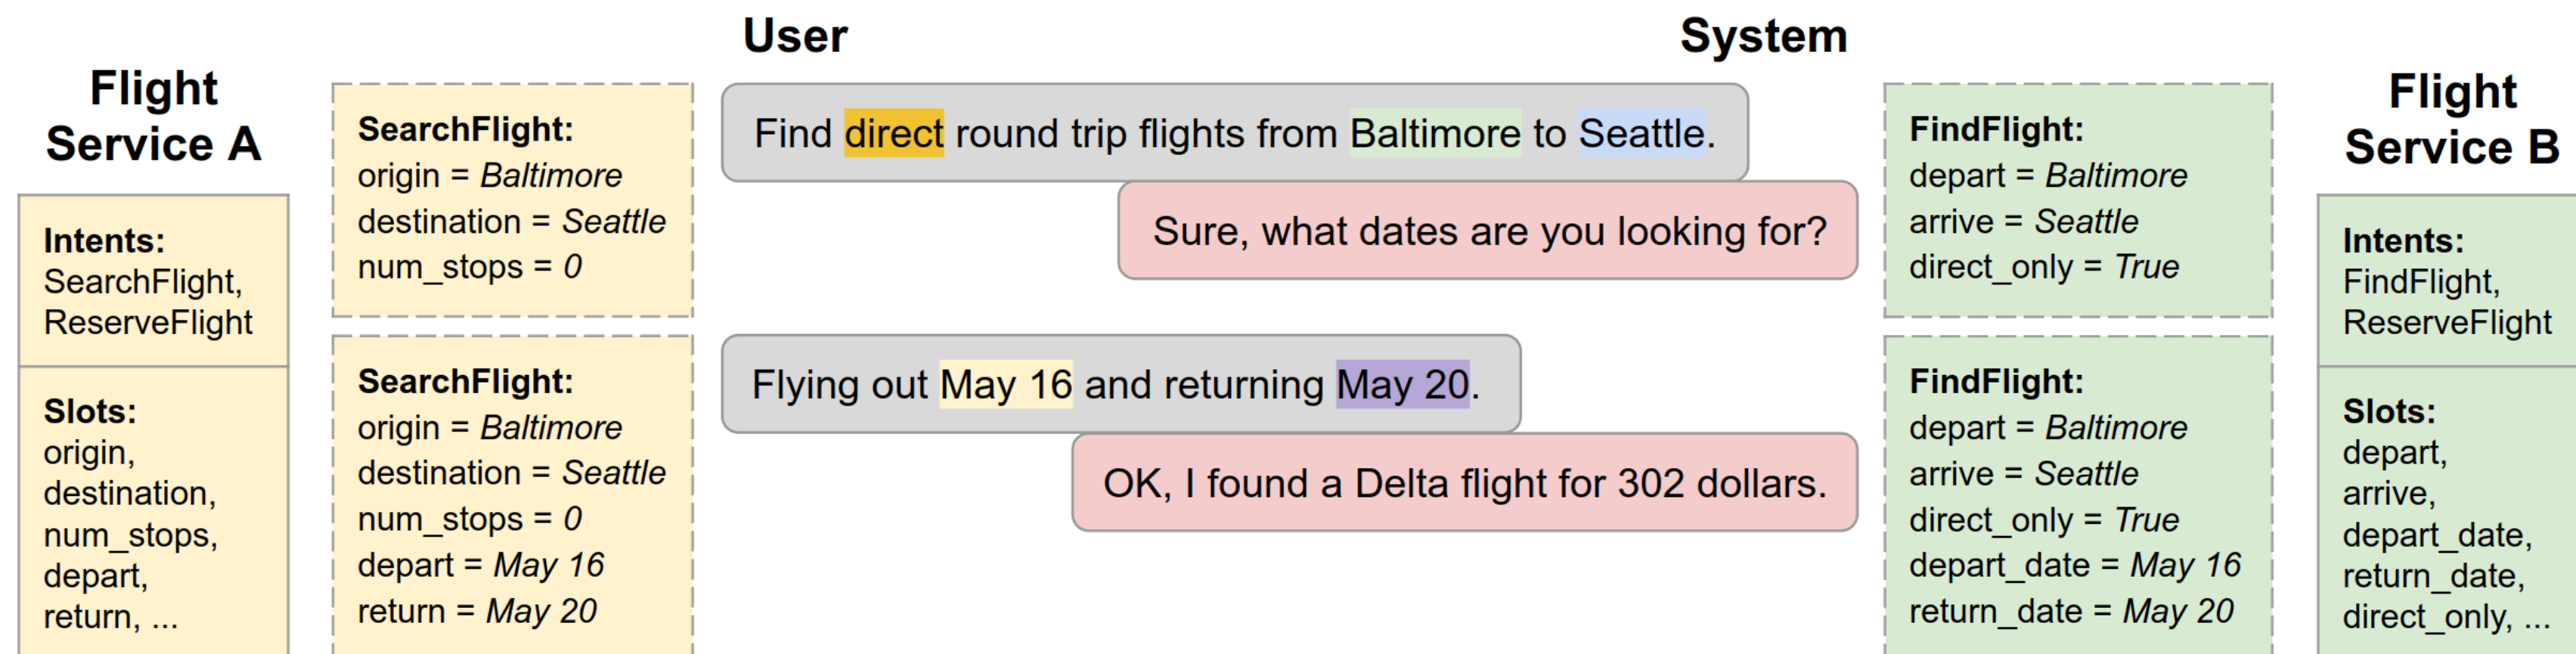
- A task can often be decomposed in a set of information that, once collected, a system can autonomously complete.
- When booking a flight, once we have the user's,
 - name
 - dates
 - itinerary
 - preferencethen we can book a flight...
- Assumes that we have access to domain knowledge about the task.

service_name: "Payment" Service description: "Digital wallet to make and request payments"
name: "account_type" categorical: True Slots description: "Source of money to make payment" possible_values: ["in-app balance", "debit card", "bank"]
name: "amount" categorical: False description: "Amount of money to transfer or request"
name: "contact_name" categorical: False description: "Name of contact for transaction"
name: "MakePayment" Intents description: "Send money to your contact" required_slots: ["amount", "contact_name"] optional_slots: ["account_type" = "in-app balance"]
name: "RequestPayment" description: "Request money from a contact" required_slots: ["amount", "contact_name"]

schema for a digital wallet service [Rastogi et al., 2020]



Slot-filling approach to task-based dialogue



- domain ontology (all frames) are a design choice.
- still used; good if you need reliable and controllable behavior.



Slot-filling approach to task-based dialogue

MultiWOZ

Example Dialogue Segment	MultiWOZ 2.2
1. Early Markup User: Help me find a moderate priced british food place please. Sys: restaurant one seven is a nice place. Do you want to book?	r-food=british, r-pricerange=moderate
2. Annotation from Database User: Can you give me the address to the hospital in Cambridge? Sys: The address is Hills Rd, Cambridge Postcode: CB20QQ	-no update-
3. Typo Sys: Okay, I can help with that. What day and time would you like to dine and how many people should I have the reservation for? User: On Thursday at 5:00. I also need a hotel in the same area. No need to have free parking.	r-bookday=thursday, r-booktime=5:00, hotel-area=west
4. Implicit Time Processing User: Can I get the postcode for that? I also need to book a taxi to the Golden Wok. Sys: The postcode is cb21tt. Are you looking for a taxi from Old Schools to the Golden Wok? User: Yes I do. I'd like to make sure I arrive at the restaurant by the booked time. Can you check?	r-name=Golden Wok, r-bookday=friday, r-booktime=11:00, taxi-destination=Golden Wok r-name=Golden Wok, r-bookday=friday, r-booktime=11:00, taxi-arriveby=11:00




System response given slot

filled slots

Dialog Act
 Intent: *Confirm*
 Slot-value pairs:
 [name = Hilton], [area = center]

Response
 Let me confirm that
 you are searching for Hilton hotel
 in the center area



condition generation on detected slots



System response given slot

#	Model	Generated Responses from Different Models
1	Input DA	Laptop{inform(name=satellite proteus 84; type=laptop; memory=8 gb; drive=1 tb; weight=2.3 kg)}
	Reference	the satellite proteus 84 is a laptop with a 1 tb drive , 8 gb memory and weighs in at 2.3 kg
	SC-LSTM	the satellite proteus 84 is a laptop with 8 gb , with a 1 tb drive , and is for business computing , and is for business computing % [businesscomputing=true] [weight=2.3kg]
	GPT-2	the satellite proteus 84 laptop has a 1 tb drive , weighs 2.3 kg and is used for business computing % [businesscomputing=true] [memory=8 gb]
	SC-GPT	the satellite proteus 84 is a laptop with 8 gb memory , 1 tb drive , and a weight of 2.3 kg
2	Input DA	Restaurant{inform_only_match(name=marnee thai and thanh long; pricerange=moderate; near=outer sunset; goodformeal=dinner)}
	Reference	marnee thai and thanh long are the only 2 moderately priced restaurants near outer sunset that are good for dinner
	SC-LSTM	marnee thai and thanh long is is moderate restaurant restaurant and good good good dinner % [near=outer sunset] [inform_only_match]
	GPT-2	there is a moderately priced restaurant called marnee thai and thanh long that is near the outer sunset area % [goodformeal=dinner] [inform_only_match]
	SC-GPT	marnee thai and thanh long is the only restaurant that serves moderately priced food near outer sunset and good for dinner
3	Input DA	Train{inform(time=50 minutes)}, hotel{request(stars=?; area=?), inform(choic=5)}
	Reference	travel time is 50 minutes . i found 5 hotels you may like . do you have a star rating request or an area you prefer
	SC-LSTM	there are 5 trains that are 50 minutes . do you have a preferred departure location ? % [train{inform(choic=5),request(departure=?)}] [hotel{request(stars=?; area=?), inform(choic=5)}]
	HDSA	there are 5 trains that meet your criteria . do you have a preference on the area or star rating ? % [train{inform(choic=5, time=50 minutes)}] [hotel{inform(choic=5)}]
	SC-GPT	there are 5 hotels that meet your criteria . the trip will last 50 minutes . do you have an area preference or star rating you would like ?

Examples of generated utterances from different models, along with its corresponding dialog acts (DAs) and references. Wrong (red), Redundant (green), missing (blue) information.

feedback

Feedback

- one of the core elements of a dialogue system is understanding if the system did the right thing.
- task-oriented dialogue often has a clear notion of completeness (=good)
- how do we know if this task was complete?

Human: I need to cancel my flight to Chicago on March 15th.

Machine: I can help you with that. I see you have a booking reference TK9472 for March 15th to Chicago O'Hare, departing at 2:35 PM. Would you like me to proceed with the cancellation?

Human: Yes, please cancel it.

Machine: I've processed your cancellation request for booking TK9472. You should receive a confirmation email within 24 hours. Is there anything else I can help you with today?

Human: Wait, will I get a refund? It was a non-refundable ticket.

Machine: Let me check the fare rules for you. Your ticket was booked under our Economy Saver fare, which is indeed non-refundable. However, you may be eligible for a travel credit minus the cancellation fee of \$150. The credit would be valid for 12 months from the original booking date.


Human: Okay, so I'll get the credit automatically then?

Machine: The credit will be processed along with your cancellation. You can use the credit by entering your booking reference when making a new reservation. Is there anything else you need assistance with?

How people use conversational AI...

- **Explicit feedback:** System provides a mechanism for users to indicate whether the system's response was good.

Netflix Star Ratings are explicit feedback



<https://news.newonnetflix.info/news/netflix-set-to-simplify-5-star-ratings-to-thumbs-updown/>

poll: how have you provided
explicit feedback in dialogue?

poll: what are the drawbacks of explicit feedback?

How people use conversational AI...

- **Explicit feedback:** System provides a mechanism for users to indicate whether the system's response was good.
- **Implicit feedback:** System infers whether the system's response was good based on logged user behavior.

Clicks are implicit feedback



<https://localsearchforum.com/threads/another-new-local-3-pack-click-through-heat-map-study-mike-ramsey.50405/>

poll: what **positive** implicit feedback signals do you provide during dialogue?

poll: what **negative** implicit feedback signals do you provide during dialogue?

poll: what are the drawbacks of
implicit feedback?

Summary

- Modern dialogue systems have built on the foundations of decades of research in language technology
- Moving online means new request types and new feedback signals (more on this during the evaluation lecture).
- Extremely exciting moment in field.