Part 4: Conversation Generation - Open Domain

Duration: 80 min

Presenter: Heydar Soudani & Faegheh Hasibi

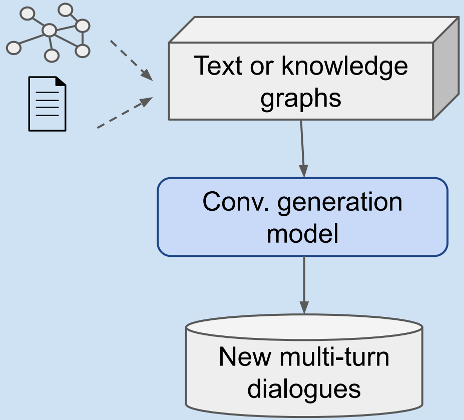

Overview

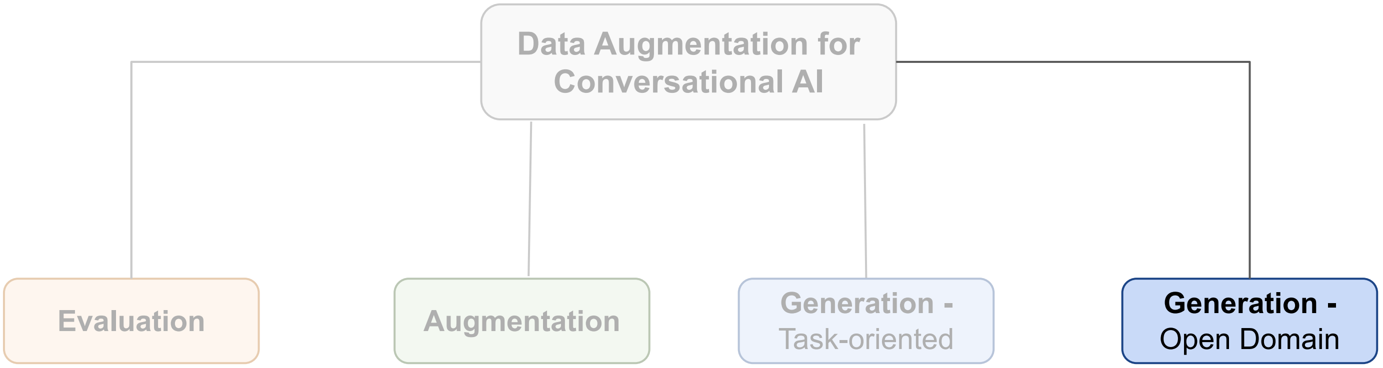

Data Augmentation for Conversational AI
Evaluation
Augmentation
Generation - Task-oriented
Generation - Open Domain

Open Domain Conversational Dataset Generation

External resources

- Textual Documents
- Knowledge Graphs (KG)
- Knowledge ingrained within the layers of LLMs

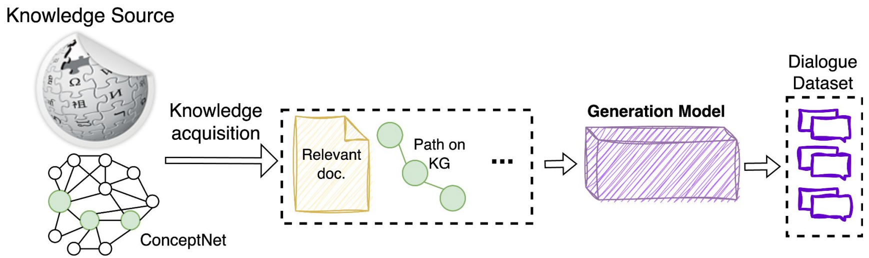

Main Characteristics

- **Coherent** and **informative** response generation
  - Responds not only depending on the current message but also based on the conversation history
  - Maintaining logic and consistency in a dialogue
- Proactive and **Target-guided** dialogue generation

Overview

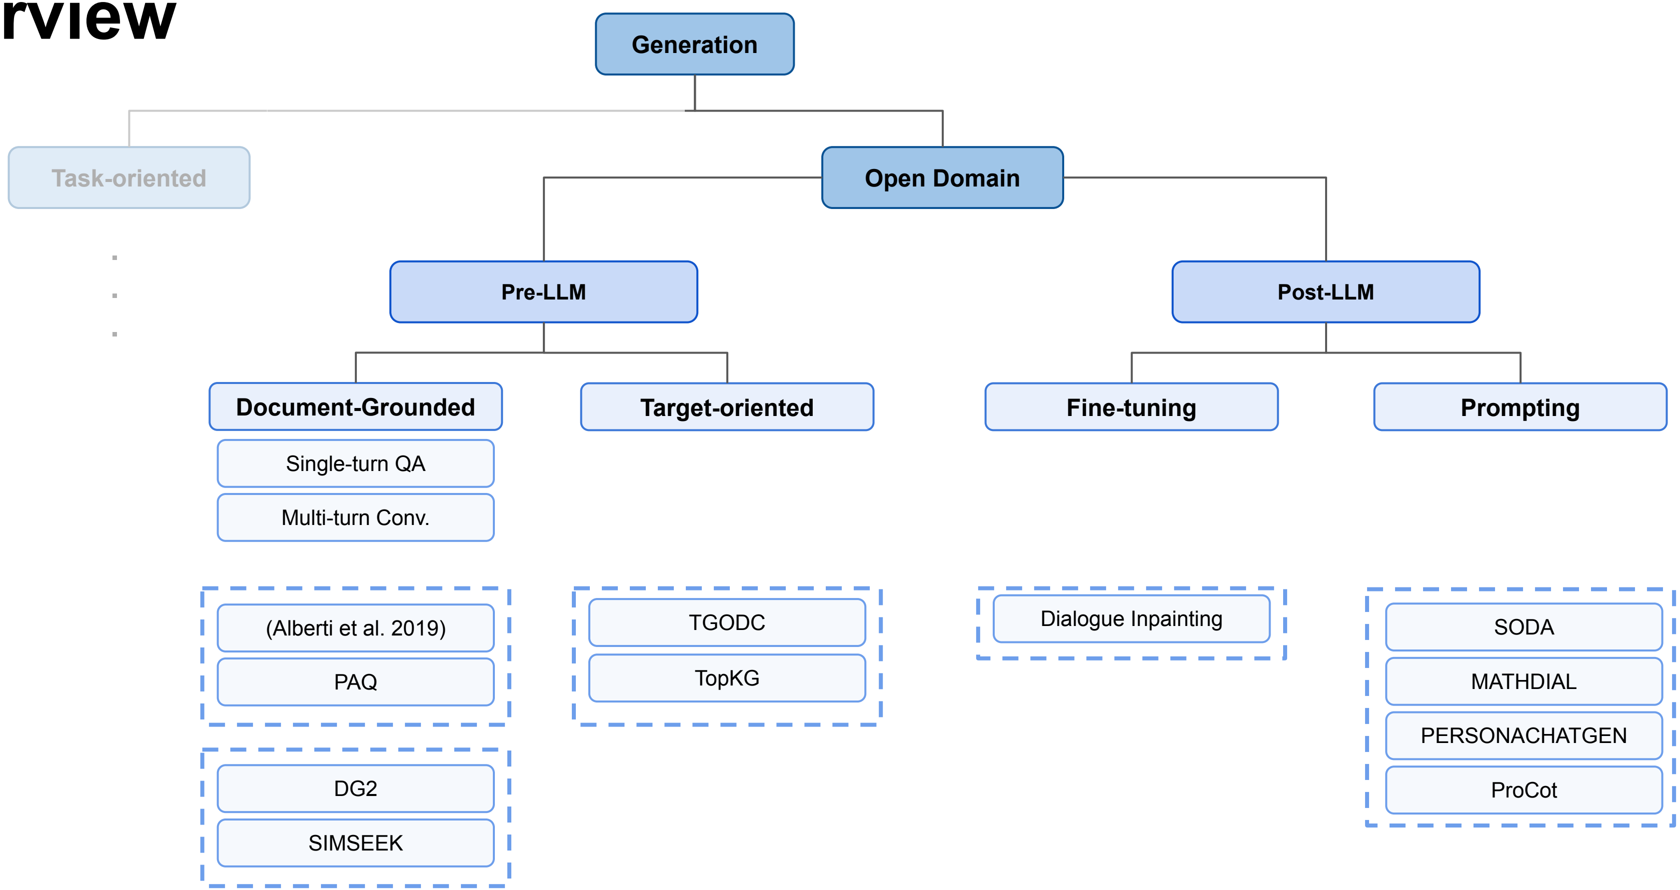

Generation
Task-oriented
Open Domain
Pre-LLM
Post-LLM
Document-Grounded
Target-oriented
Fine-tuning
Prompting
Single-turn QA
Multi-turn Conv.
(Alberti et al. 2019)
PAQ
DG2
SIMSEEK
TGODC
TopKG
Dialogue Inpainting
SODA
MATHDIAL
PERSONACHATGEN
ProCot

Overview

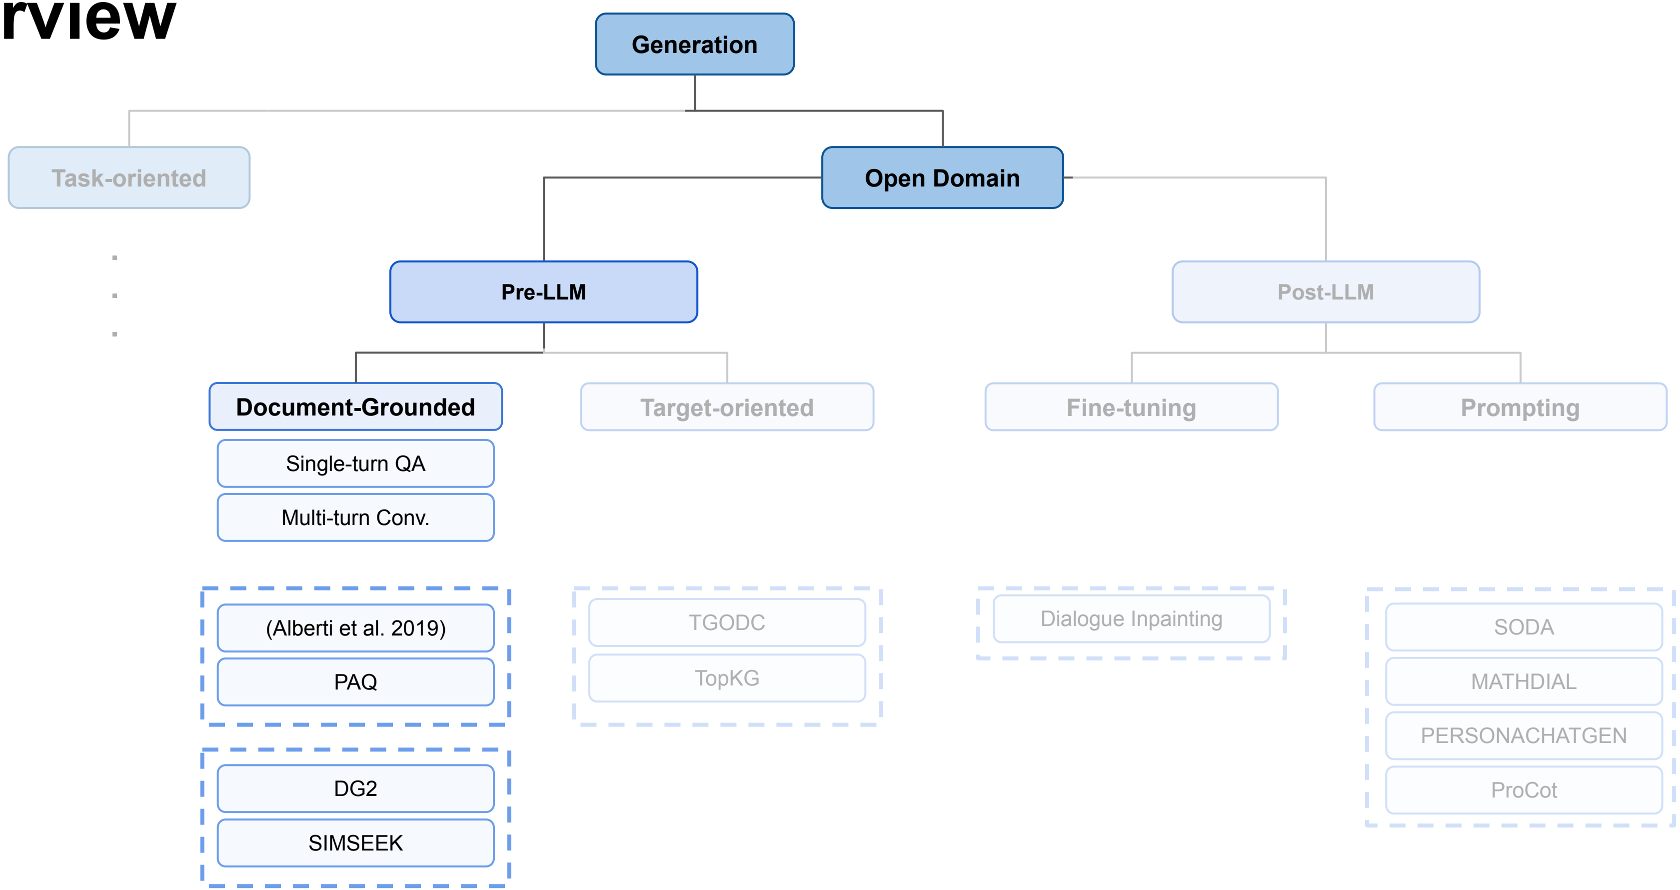
Generation
Task-oriented
Open Domain
Pre-LLM
Post-LLM
Document-Grounded
Target-oriented
Fine-tuning
Prompting
Single-turn QA
Multi-turn Conv.
(Alberti et al. 2019)
PAQ
DG2
SIMSEEK
TGODC
TopKG
Dialogue Inpainting
SODA
MATHDIAL
PERSONACHATGEN
ProCot

Single-turn QA Pair Generation

(Alberti et al., 2019)

Task Definition:

For a given passage C, (C, Q, A) is emitted as a new synthetic training example

The pipeline approach

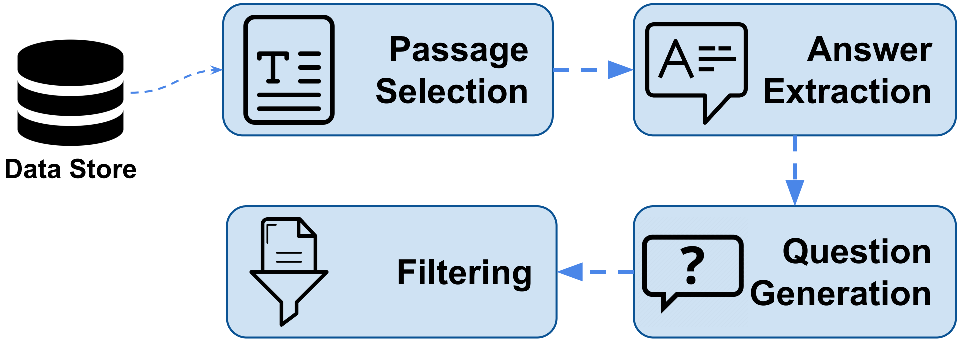

(Lewis et al., 2021)

Single-turn: Passage Selection

Finds passages that:

- Contain information that humans may ask about
- Are good candidates to generate questions from

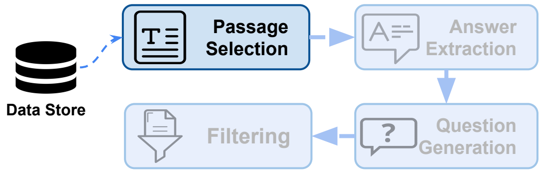

1) Passage Selection	C	… in 1903, boston participated in the first modern world series, going up against the pittsburgh pirates …
2) Answer Extraction	C→A	1903
3) Question Generation	C, A→Q	when did the red sox first go to the world series
4) Filtering	C, Q→A′ A≟A′	1903 Yes

Single-turn: Answer Extraction

Given a passage (or sometimes a question):

- Identify spans that are likely to be answers to questions

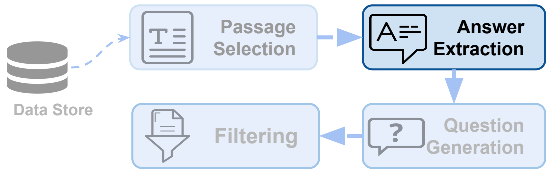

1) Passage Selection	**C**	… in 1903, boston participated in the first modern world series, going up against the pittsburgh pirates …
2) Answer Extraction	**C→A**	1903
3) Question Generation	**C, A→Q**	when did the red sox first go to the world series
4) Filtering	**C, Q→A′** **A≟A′**	1903 Yes

Single-turn: Question Generation

Given a passage and an answer

- Generate likely questions with that answer

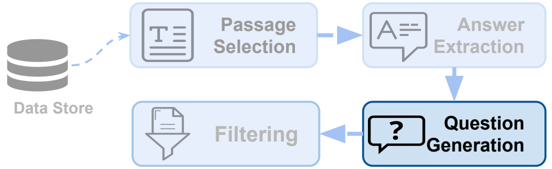

1) Passage Selection	**C**	… in 1903, boston participated in the first modern world series, going up against the pittsburgh pirates …
2) Answer Extraction	**C→A**	1903
3) Question Generation	**C, A→Q**	when did the red sox first go to the world series
4) Filtering	**C, Q→A′** **A≟A′**	1903 Yes

Single-turn: Filtering

Improve the quality of generated QA pairs

- By ensuring that they are consistent
- By checking that the generated answer is valid to the question

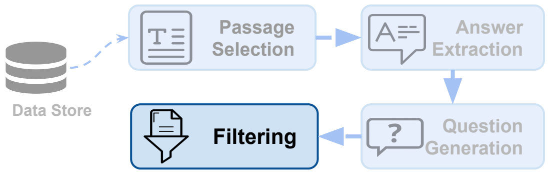

1) Passage Selection	**C**	… in 1903, boston participated in the first modern world series, going up against the pittsburgh pirates …
2) Answer Extraction	**C→A**	1903
3) Question Generation	**C, A→Q**	when did the red sox first go to the world series
4) Filtering	**C, Q→A′** **A≟A′**	1903 Yes

Multi-turn Conversation Generation

Differences with Single-turn QA

- The conversation history is also fed to the model
- Deeper understanding of the context and the dialogue history is required

(Feng et al., 2020)

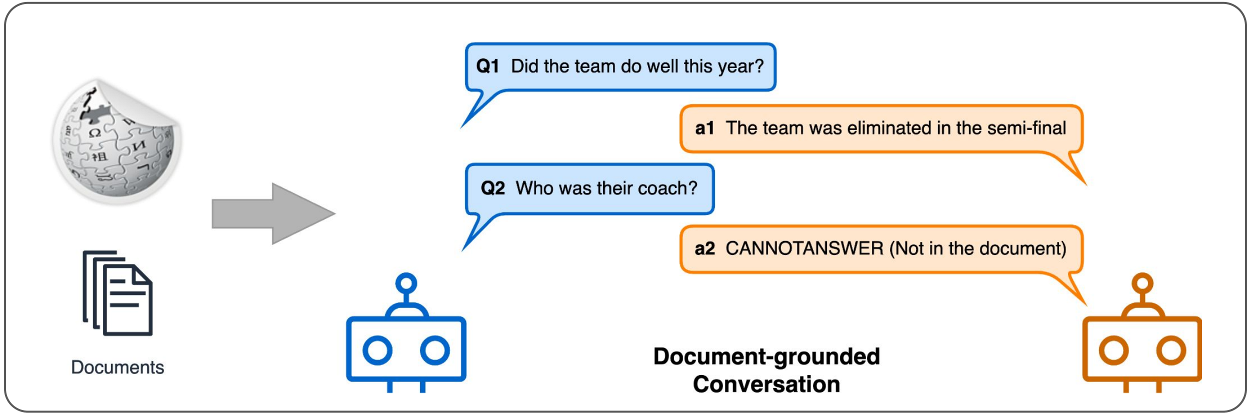

Multi-turn Conversation Generation

Problem Definition

(Wu et al., 2022)

Given a document *C*, a dialog *d* between the user and the agent can be modeled with

$$p(d|C) = \prod_{i=1}^{t} p(u_i, a_i | c_i \in C)$$

- u_i: the user turn utterance
- a_i: the agent turn utterance
- c_i: the selected passage at i-th turn

- **Document**: a series of text passages, broken down into spans
- **Dialogue**: a series of turns between user (1st speaker) and agent (2nd speaker)

$$d = [(u_1, a_1), (u_2, a_2), \ldots (u_t, a_t)]$$

Inspired By Single-turn QA

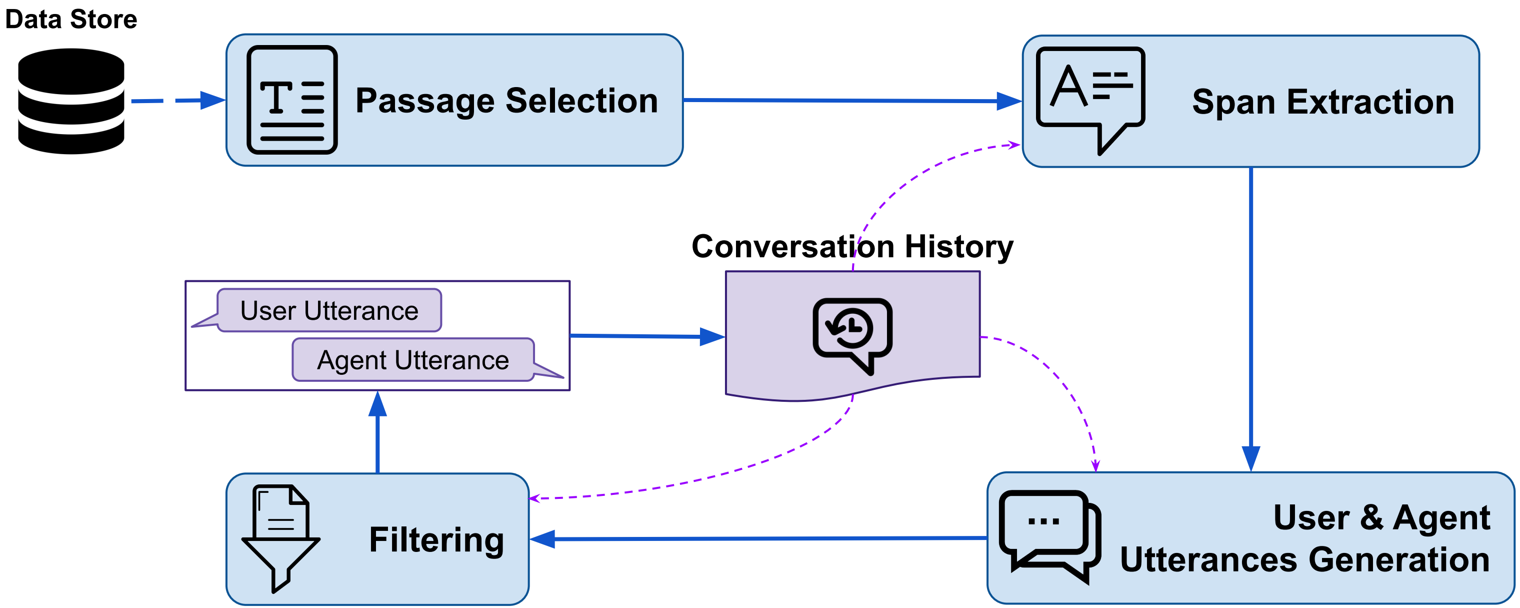

Passage Selection

Top 5 DMV Mistakes and How to Avoid Them

<Passage 1> Many DMV customers make easily avoidable mistakes that cause them significant problems, …

<Passage 2>

<Passage 3> Not Bringing Proper Documentation to DMV Office. About ten percent of customers visiting a DMV office do not bring what they need to complete their transaction and see if your transaction can be …

<Passage 4>

<Passage 5> We send 500,000 inquiry letters a year. If the inquiry letter does not resolve the problem, we must suspend the vehicle registration and, if it persists, your driver license!
We suspend 300,000 registrations a year for failure to maintain insurance. …

Given the dialogue context, the most appropriate passage in a document is selected for further grounding

$$p(\underbrace{c_t}_{\text{Selected passage in turn } t} | \underbrace{\{u_i, a_i\}_{i<t}}_{\text{Conversation History}}, \underbrace{C}_{\text{Document}})$$

DG2 (Wu et al., 2022)

- Passage ranker is a fine-tuned RoBERTa model
- Passage probability is maximized with contrastive loss
  - Positive passages are from ground truth
  - Negative passages are from the same document

Span (Rationale) Extraction

<Passage 3> Not Bringing Proper Documentation to DMV Office. About ten percent of customers visiting a DMV office do not bring what they need to complete their transaction and see if your transaction can be ...

Conversation History

Highlights the rationale span used to generate the dialogue turn

<Passage 3> Not Bringing Proper Documentation to DMV Office. ==About ten percent of customers visiting a DMV office do not bring what they need to complete their transaction== and see if your transaction can be ...

Extract a rationale span from the selected passage

$$p(r_t \mid \underbrace{\{u_i, a_i\}_{i<t}}_{\text{Conversation History}}, \underbrace{c_t}_{\text{Selected passage in turn } t})$$

DG2 (Wu et al., 2022)

- For short span, like QA task
  - Model the start and end position of a span independently
- For long span
  - Autoregressive method
  - Samples the start and end position sequentially

Utterance Generation

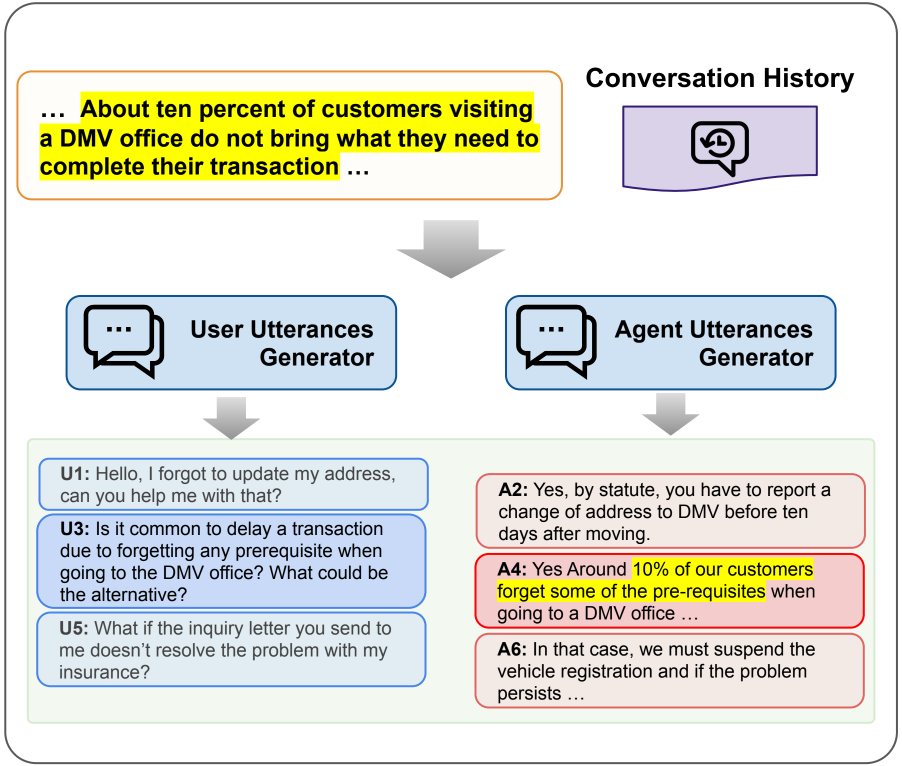

User utterance generator

- Generates a question with the answer span
- Highlight the rationale span by wrapping its text,

$$p(u_t) = p(u_t|\underbrace{\{u_i, a_i\}_{i<t}}_{\text{Conversation History}}, \underbrace{c'_t}_{\text{Selected passage in turn } t})$$

Agent utterance generator

- Generates the response with the answer span
- The dialogue history now includes the previous generated user utterance

$$p(a_t) = p(a_t|\{a_i, u_i\}_{i<t}, c'_t)$$

Filtering

- To check the extracted rationale to be aligned with the dialogue context and the user utterance
- Roundtrip passage selector

$$p(\underline{\hat{c}_t}|\underbrace{\{u_i, a_i\}_{i<t},}_{\text{Conversation History}}, \underbrace{u_t}_{\text{User Utterance}}, \underbrace{C}_{\text{Document}})$$

- Rationale predictor

$$p(\hat{r}_t|\underbrace{\{u_i, a_i\}_{i<t}}, \underbrace{u_t}, \underline{\hat{c}_t})$$

- Filter out the utterance
  - When the predicted answer contradicts the previous answer
- Since the span is long, F1 word overlap for filtering, instead of exact match

Passage Blinded

SimSEEK: Simulating Information-Seeking conversation from unlabeled documents (Kim et al., 2022)

- **RQ:** When asking a question, is it better to know its answer in advance?
- To answer the RQ: Proposes two scenarios and compare them

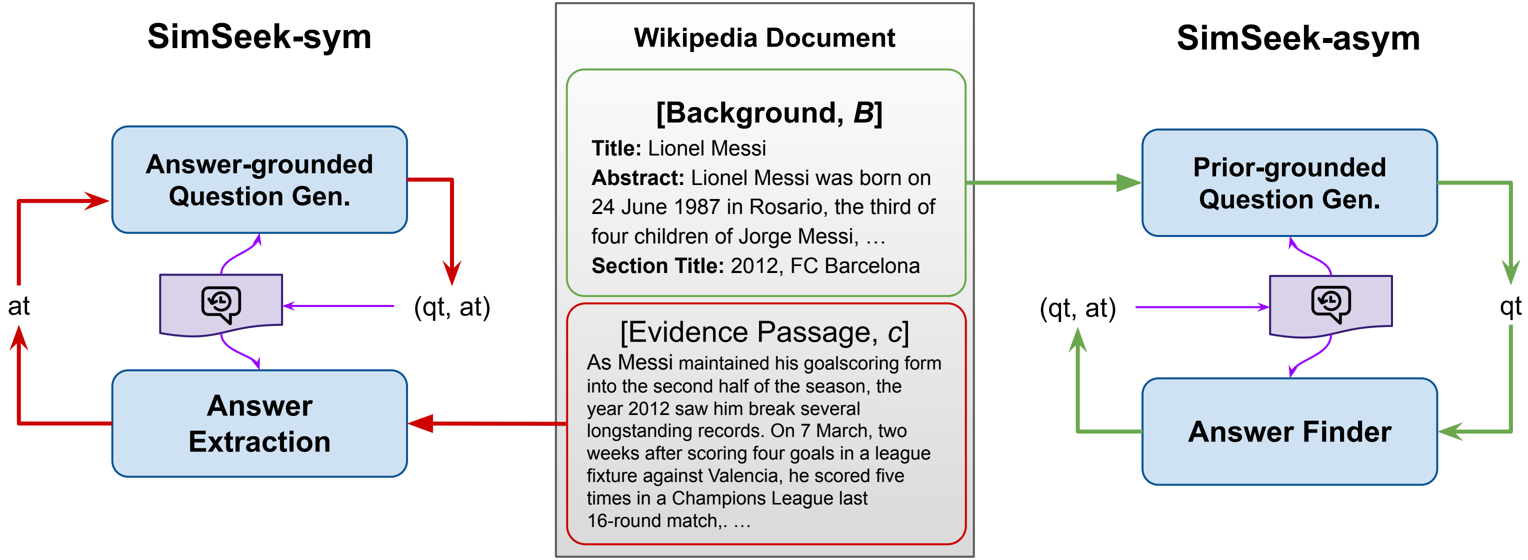

SYMSEEK: Major Differences

SimSeek-sym

- First extracts an answer candidate from the passage
- Questioner can access all answer-relevant information
  - Pro: Coherency with answer
  - Con: Constraint to the predetermined answer

SimSeek-asym

- First asks a question without accessing an answer or passage
- Questioner asks any questions relevant to the topic, guessing inaccessible passage
  - Pro: encouraging information-seeking behaviour

SYMSEEK: Experiments

- **Down-stream task:** Semi-supervised CQA
- **Baseline:** Extended version of PAQ
- **Results:** SimSeek-ASYM shows dominant performance compared to other baselines in both setup

CQA Backbone Synthetic CQA Generation	Trained on $\hat{\mathcal{D}}$	 $\mathcal{D}+\hat{\mathcal{D}}$
RoBERTa-large		
None ($\hat{\mathcal{D}}$ = empty)	-	65.6
PAQ-CANARD	38.8	66.5
PAQ-QuAC	51.5	66.6
SIMSEEK-SYM	54.3	66.3
SIMSEEK-ASYM	**64.8**	**67.5**
Human Annot. ($\hat{\mathcal{D}}$ = $\text{QuAC}_{\text{unseen}}$)	65.0	70.3

SIMSEEK: Qualitative Analysis

SimSeek-sym

Title: Native Americans in the United States
Section Title: Self-determination

Document c:
· · · Upset with tribal government and the failures of the federal government to enforce treaty rights, about 300 Oglala Lakota and AIM activists took control of Wounded Knee on February 27, 1973. Indian activists from around the country joined them at Pine Ridge, and the occupation became a symbol of rising American Indian identity and power. Federal law enforcement officials and the national guard cordoned off the town, and the two sides had a standoff for 71 days. During much gunfire, one United States Marshal was wounded and paralyzed.· · ·

· · ·
q3: What happened at Wounded Knee?
a3: Indian activists from around the country joined them at Pine Ridge, and the occupation became a symbol of rising American Indian identity and power.
q4: What happened after they took control of Pine Ridge?
a4: Federal law enforcement officials and the national guard cordoned off the town, and the two sides had a standoff for 71 days.
q5: What happened during the standoff?
a5: During much gunfire, one United States Marshal was wounded and paralyzed.
· · ·

SimSeek-asym

Title: Thor Heyerdahl
Section Title: Kon-Tiki expedition

Background B:
Thor Heyerdahl (October 6, 1914 - April 18, 2002) was a Norwegian adventurer and ethnographer with a background in zoology, botany, and geography. He became notable for his Kon-Tiki expedition in 1947, · · ·

· · ·
q4: What were some of the things he found on the Kon-Tiki expedition?
a4: The raft proved to be highly manoeuvrable, and fish congregated between the nine balsa logs in such numbers that ancient sailors could have possibly relied on fish for hydration in the absence of other sources of freshwater.
q5: Are there any other interesting aspects about this article?
a5: The documentary film of the expedition entitled Kon-Tiki won an Academy Award in 1951.
q6: Why did the film win an Academy Award?
a6: CANNOTANSWER.

Overview

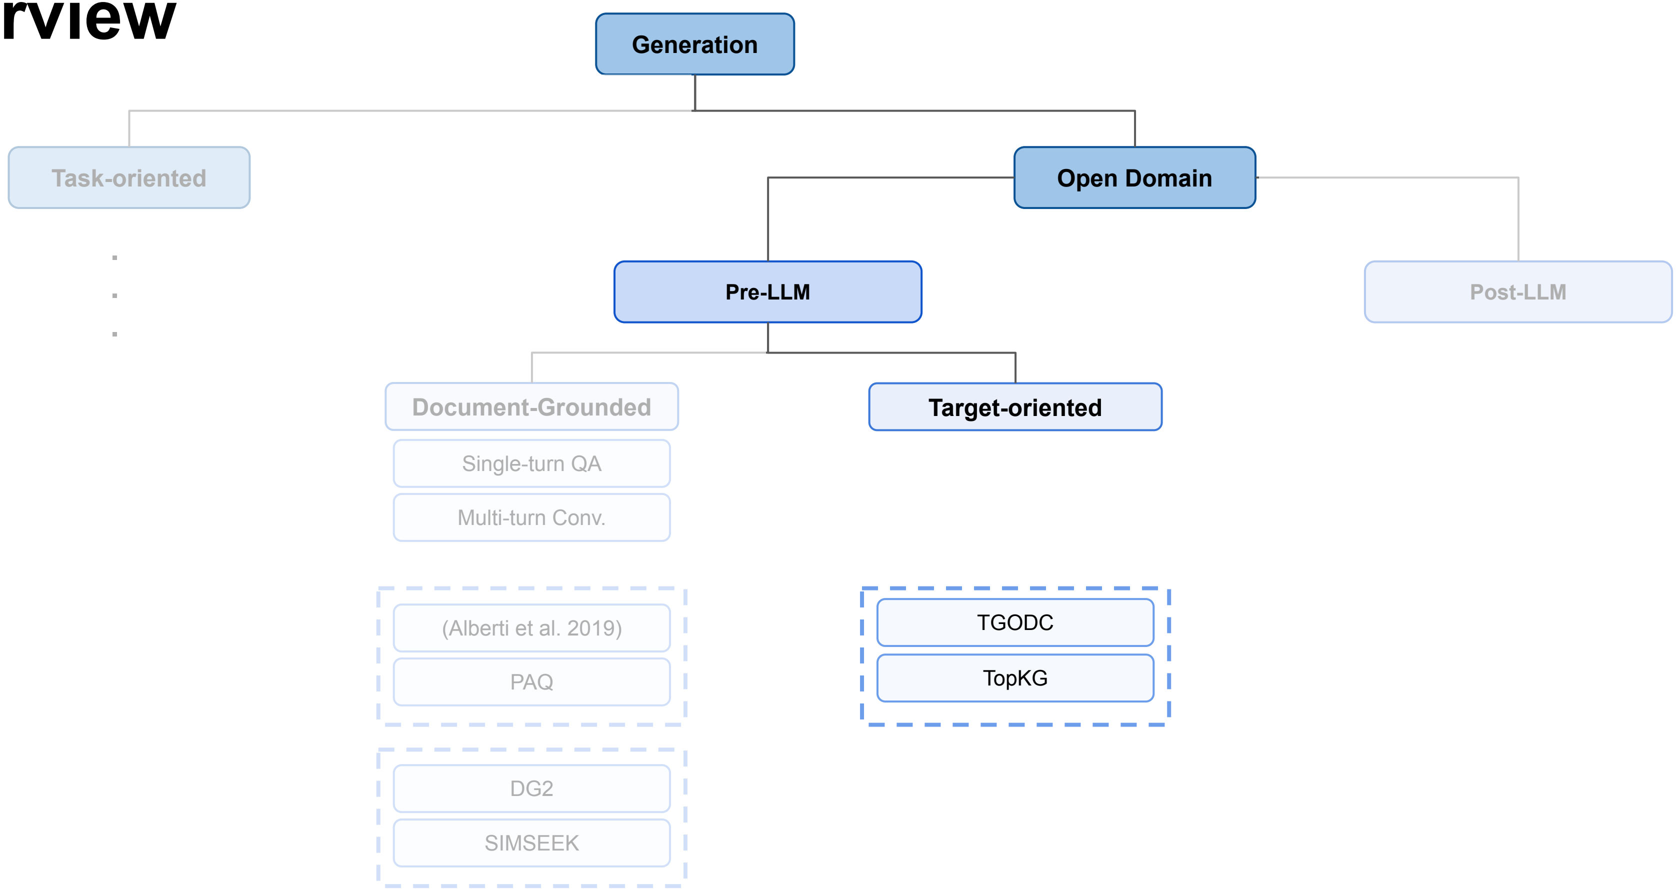
Generation
Task-oriented
Open Domain
Pre-LLM
Post-LLM
Document-Grounded
Single-turn QA
Multi-turn Conv.
(Alberti et al. 2019)
PAQ
DG2
SIMSEEK
Target-oriented
TGODC
TopKG

Self-Play Simulation

- Two agents are talking to each other
- The trained dialogue agent is used as both user and agent bot
- Target-guided Dialogue Systems are often leveraged for self-play simulation
- What is Target-guided Dialogue System?

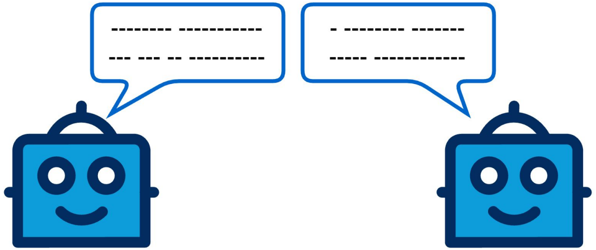

Target-guided Dialogue System

Definition (Deng et al. 2023)

- Agent is provided with a target that the user is unaware of
- Generate multiple utterances, steer the conversation towards the predetermined target

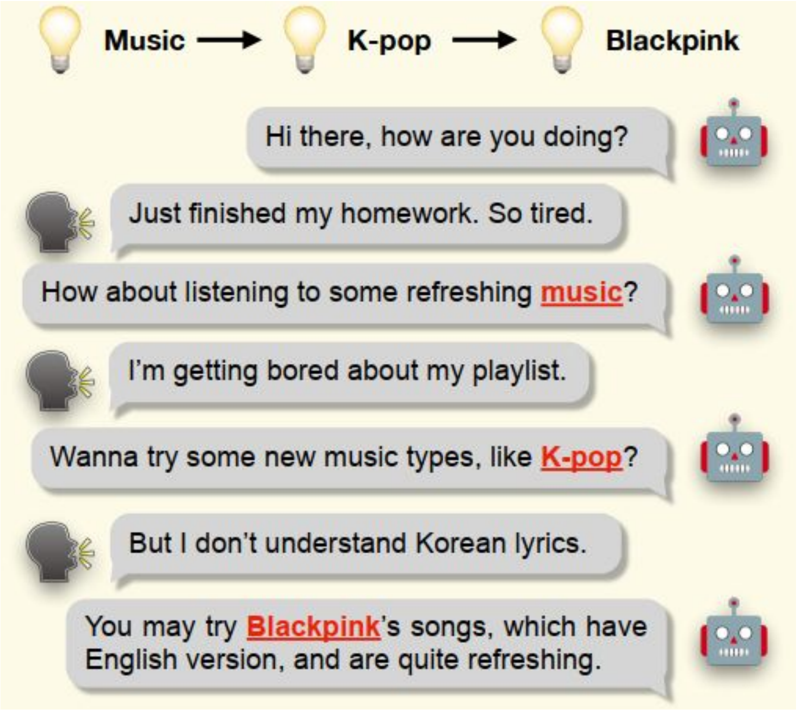

Response Criteria

- Transition smoothness
- Target achievement

Core Modules

1) Next Topic (keyword) Selection

aims to predict keywords of the next response that is appropriate in the *conversation context*

2) Transition Response Generation

aims to fulfill the *end target* by proactively driving the discussion topic forward

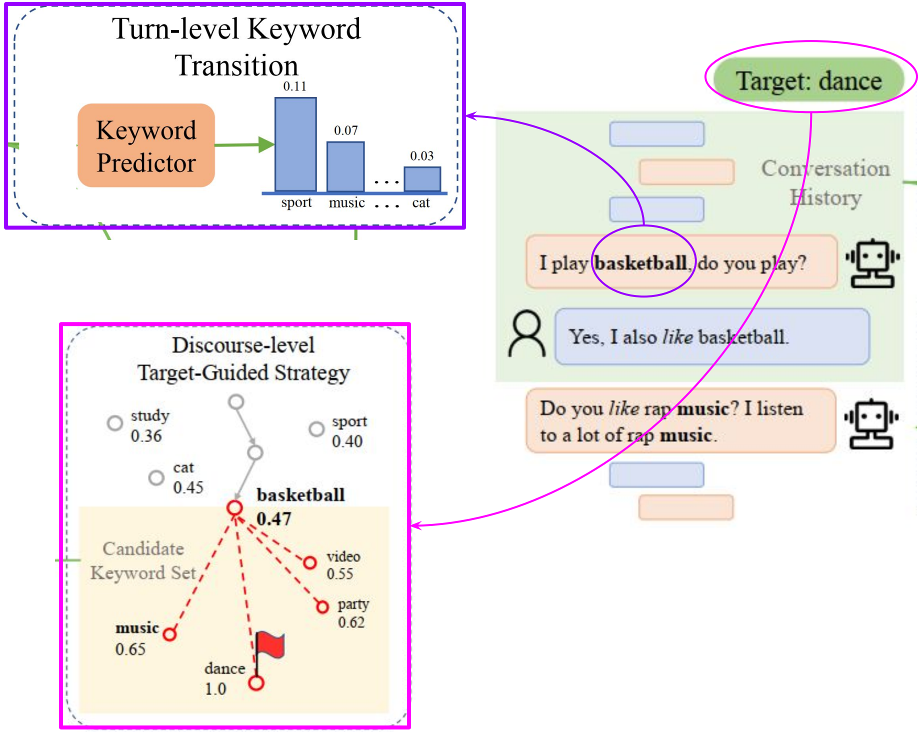

(Yang et al. 2022)

Next Topic (keyword) Selection

Pairwise PMI-based Transition (Tang et al., 2019)

- Construct a keyword pairwise matrix that characterizes the association between keywords
- Use pointwise mutual information (PMI)
  - Given two keywords wi and wj, computes likeliness of wj -> wi with

$$\mathrm{PMI}(w_i, w_j) = \log p(w_i|w_j)/p(w_i)$$

    - $p(w_i|w_j)$ is the ratio of transitioning to wi in the next turn given wj in the current turn
    - $p(w_i)$ is the ratio of wi occurrence
  - Both quantities can be directly counted from the conversation data beforehand

Transition Response Generation: Type of Planning

1) Local and Greedy Planning

- The algorithm makes locally optimal choices at each step with the hope of finding a global optimum
- Without considering the overall long-term consequences

2) Global Planning

- Takes into account the bigger picture
- Generating optimal responses, considering the context of the entire conversation

Local and Greedy Planning

Target Closeness Score (Tang et al., 2019)

- The keyword of each turn must move strictly closer to the end target compared to those of preceding turns
- Cosine similarity between normalized word embeddings
- To generate the response: Second module determines valid candidates, then the first module samples or picks the most likely one according to the keyword distribution

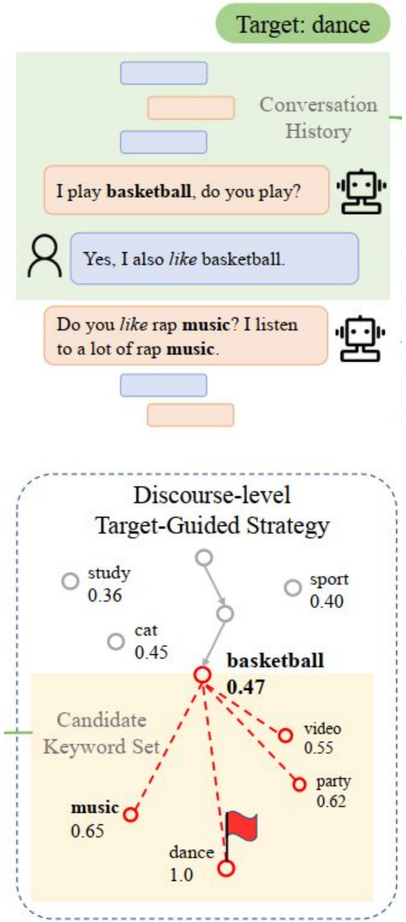

Global Planning: TopKG (Yang et al.,2022)

Idea: Input possible paths from initial utterance to target word

- Using external KG
- Global Graph: A subgraph consisting of a set of potential paths from the starting to target

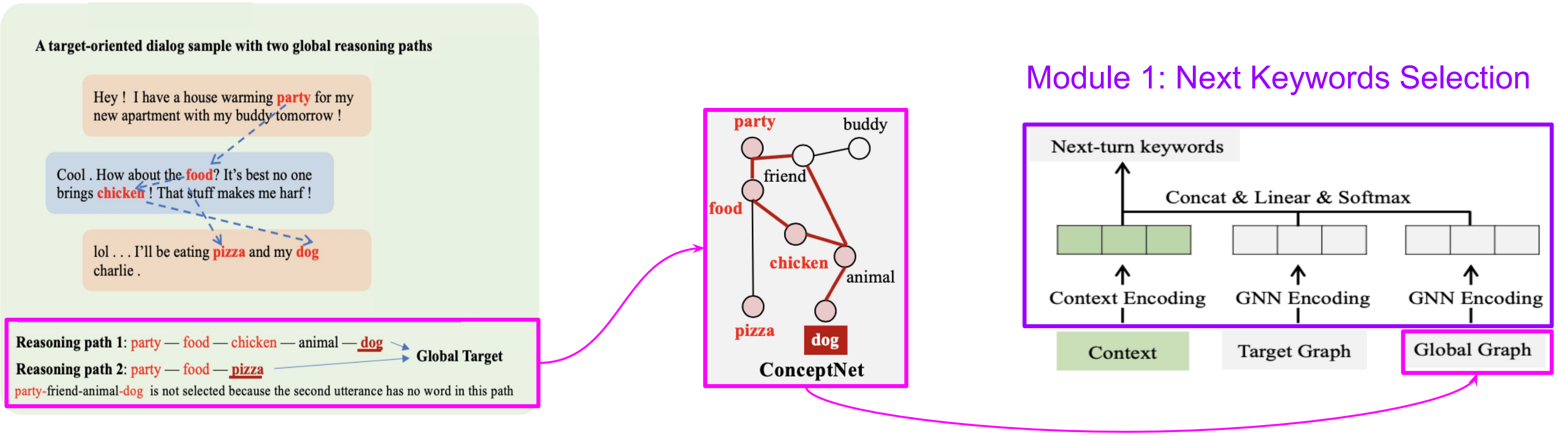

Global Planning: TopKG (Yang et al.,2022)

RL Framework:

A simulation-based environment to guide the response generator model toward the global target

Local Reward:

- Encourages the contextual consistency at each turn of dialog
- A discriminator to tell whether an utterance sequence is semantically consistent

Global Reward:

- Encourages the global target-oriented response
- Giving a positive reward of "1" if the global target word finally appears in the last turn or a negative reward of "-1" otherwise

RL Rewards

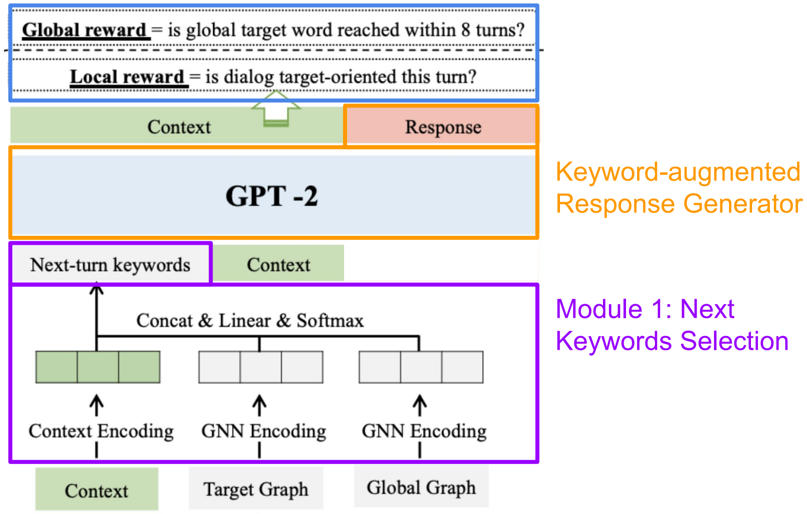

Overview

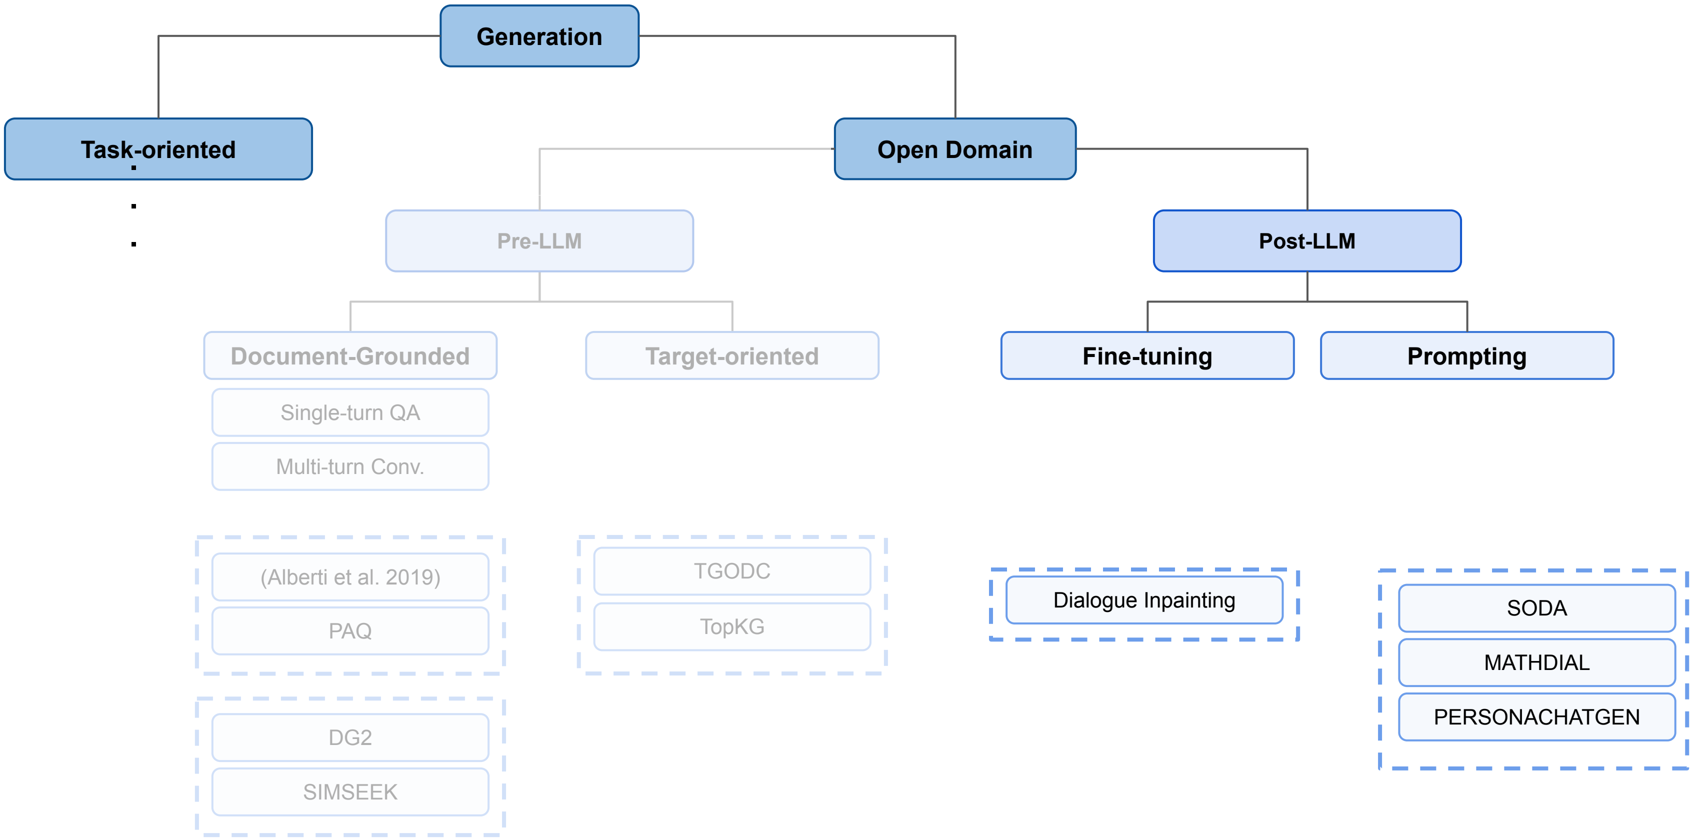

Generation
Task-oriented
Open Domain
Pre-LLM
Post-LLM
Document-Grounded
Target-oriented
Fine-tuning
Prompting
Single-turn QA
Multi-turn Conv.
(Alberti et al. 2019)
PAQ
DG2
SIMSEEK
TGODC
TopKG
Dialogue Inpainting
SODA
MATHDIAL
PERSONACHATGEN

LLM Fine-tuning for Dialogue Inpainting

- **Task:** Take a partial dialog $\Longrightarrow$ Generate a complete dialog

$$(u_1, u_2, \diamond, u_4, \diamond) \Longrightarrow d = (u_1, u_2, \ldots, u_t, \ldots, u_T)$$

- Analogy to overhearing someone else's phone call
- Hear on side, try to guess another side

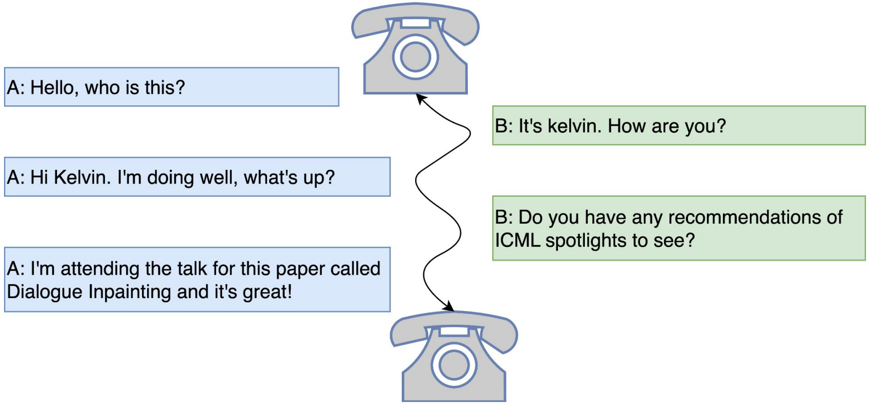

LLM Fine-tuning for Dialogue Inpainting

- **Task:** Take a partial dialog ⟹ Generate a complete dialog

$$(u_1, u_2, \diamond, u_4, \diamond) \Longrightarrow d = (u_1, u_2, \ldots, u_t, \ldots, u_T)$$

- Pretend that a document is the transcript of what one party said to another one, then predict utterances of the other side

- Rewrites high quality documents (e.g., Wikipedia) into dialogues

- Generates an enormous corpus of information-seeking dialogs (19M+ dialogs) with expert answers
  - 1,000x larger than the largest existing ConvQA dataset

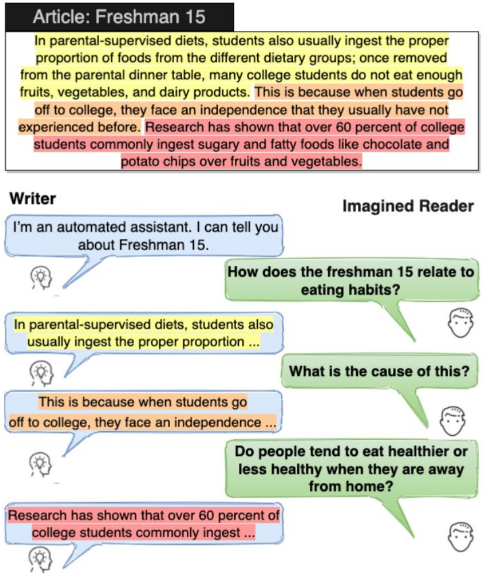

(Dai et al., 2022)

LLM Fine-tuning for Dialogue Inpainting

- **Task:** Take a partial dialog ⇨ Generate a complete dialog

$$(u_1, u_2, \diamond, u_4, \diamond) \Rightarrow d = (u_1, u_2, \ldots, u_t, \ldots, u_T)$$

- **Training:** Dialog reconstruction
  - Randomly mask one utterance (u_t)

$$d_{m(t)} = (u_1, \ldots, u_{t-1}, \diamond, u_{t+1}, \ldots, u_T)$$

  - Train a generative model to predict the masked utterance

$$p_\theta(u_t \mid d_{m(t)})$$

- Similar to the masked language modeling task used by BERT
- T5-XXL (11B parameters) checkpoint is fine-tuned on existing dialogue corpus

LLM Fine-tuning for Dialogue Inpainting

Inference: Transforming documents into dialogues

- Convert document into spans (e.g., sentences)
- Generate the start utterance s_{prompt}:
  - *"Hello, I am an automated assistant and can answer questions about* [document title]*"*
- Autoregressively generate utterances

$$(s_{\text{prompt}}, \diamond, s_1) \Longrightarrow \hat{u}_1$$

$$(s_{\text{prompt}}, \hat{u}_1, s_1, \diamond, s_2) \Longrightarrow \hat{u}_2$$

$$\ldots$$

Highly effective on conversational retrieval datasets; e.g., TREC CAsT, OR-QUA

Overview

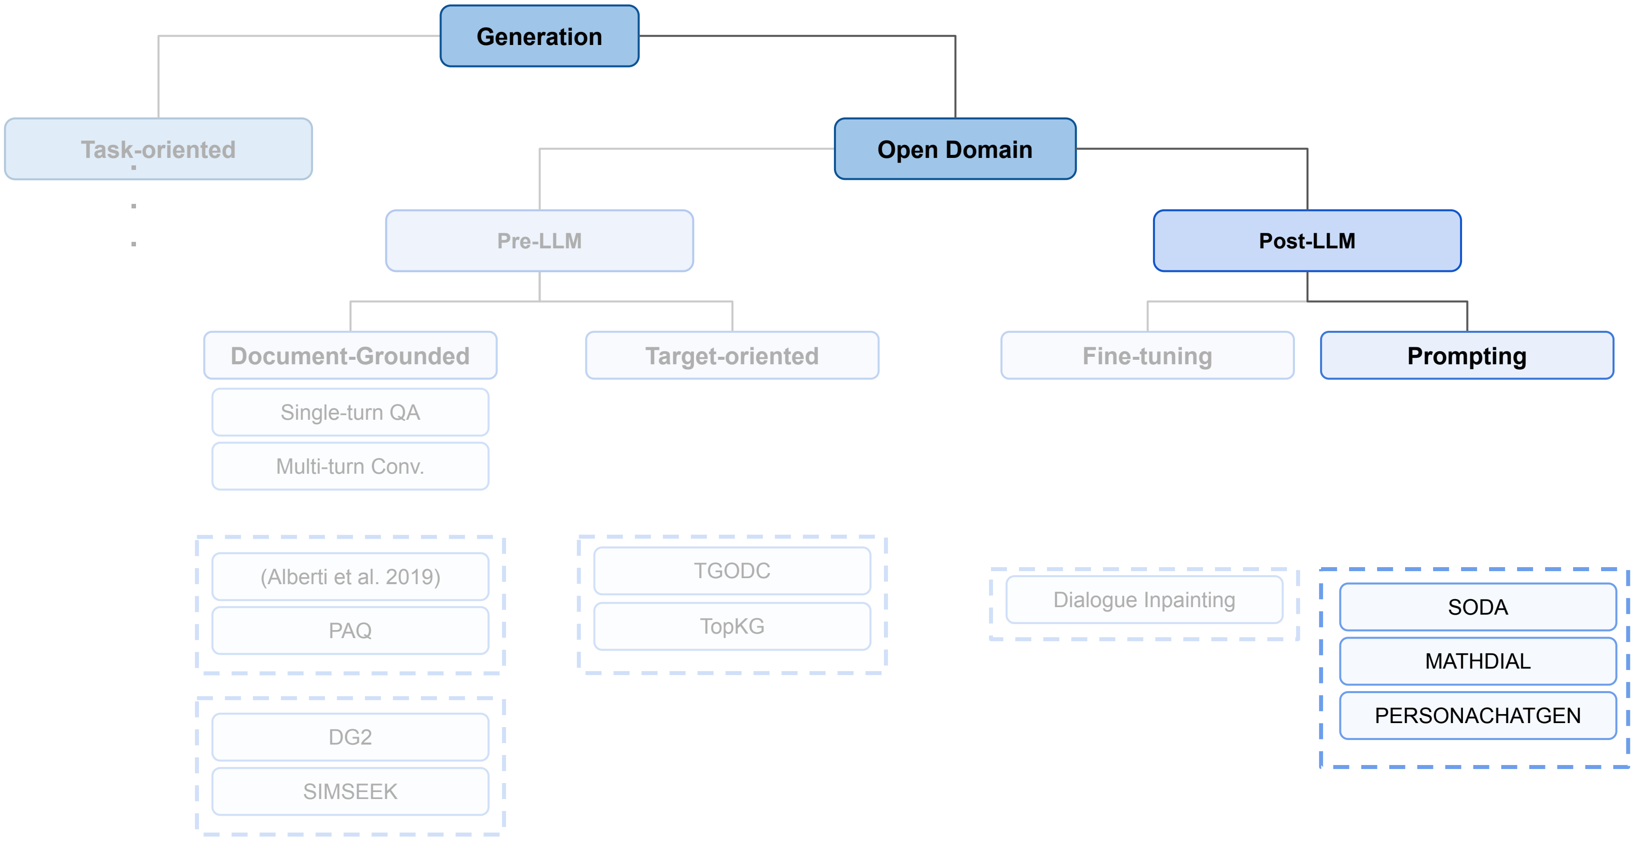
Generation
Task-oriented
Open Domain
Pre-LLM
Post-LLM
Document-Grounded
Target-oriented
Fine-tuning
Prompting
Single-turn QA
Multi-turn Conv.
(Alberti et al. 2019)
PAQ
DG2
SIMSEEK
TGODC
TopKG
Dialogue Inpainting
SODA
MATHDIAL
PERSONACHATGEN

Prompt-based Conversational Data Generation

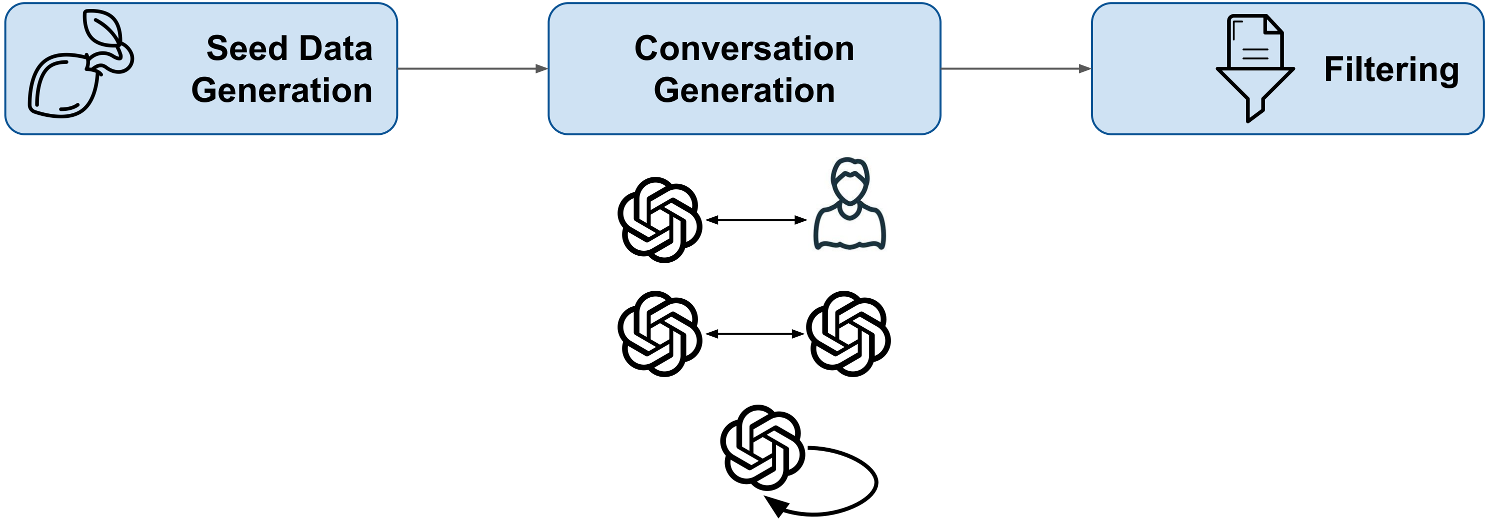

MATHDIAL

A dialogue tutoring dataset for Math problems (1.5k dialogues)

- Expert annotator role plays a teacher
- LLM simulates the student

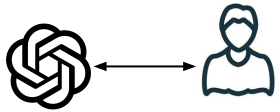

Challenges in collecting dialogue tutoring datasets:

- Crowdsourcing
  - Low pedagogical quality, poor annotator training
  - Small-scale datasets
- Classroom recording
  - Noisy dialogues
  - Privacy concerns

(Macina et al., 2023)

MATHDIAL: Seed Data Generation

Problem & Confusion Selection

- Collect multi-step math problems from GSM8k dataset
- Select failures of ChatGPT with Chain-of-thought reasoning

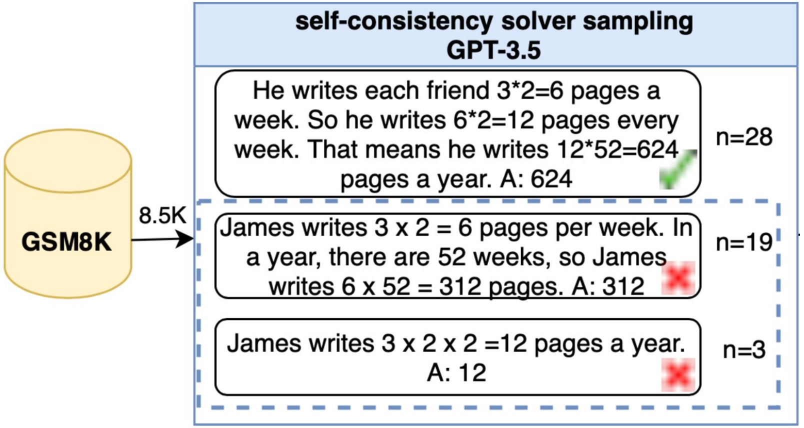

(Macina et al., 2023)

MATHDIAL: Conversation Generation

Pair LLM with human

- Selects professionals with teaching experience through the Prolific
- Prompts InstructGPT to generate student turns

Student Persona: (STUDENT PERSONA)
Math problem: (MATH PROBLEM)
Student solution: (STUDENT SOLUTION)
Context: (STUDENT NAME) thinks their answer is correct. Only when the teacher provides several good reasoning questions, (STUDENT NAME) understands the problem and corrects the solution. (STUDENT NAME) can use a calculator and thus makes no calculation errors. Send EOM tag at end of the student message.
(DIALOG HISTORY)

(Macina et al., 2023)

MATHDIAL: Filtering

Safety and quality filtering

- Annotators with 100% completion rates are selected
- conversations with any sensitive content are filtered using Perspective API

SODA: A Million-scale Social Dialogue Dataset

- A social interaction dataset constructed using self-chat of InstructGPT
- Infuses 1.5 million triples from a commonsense knowledge graph (Atomic)
- Used for fine-tuning of T5-XL to generate a conversational model (COSMO)

Conversation Dataset Generation

Downstream Task

SODA: Seed Data Generation

Commonsense triplet selection

- Social triplet from Atomic are retrieved
- *[Head, Tail, Relation]*

Example

Triplet from Atomic:

- *Head:* PersonX moves a step closer to the goal
- *Relation:* xNeed
- *Tail:* to take the first step

SODA: Seed Data Generation

Commonsense triplet selection

- Social triplet from Atomic are retrieved
- *[Head, Tail, Relation]*

Commonsense Knowledge → Narrative

- Triple form → sentence
- Sentence → narrative (prompt InstructGPT)

Example

Triplet from Atomic:

- *Head:* PersonX moves a step closer to the goal
- *Relation:* xNeed
- *Tail:* to take the first step

Triplet to Sentence:
Madeleine took the first step. Madeleine moves a step closer to the goal

Sentence to Narrative:
Madeleine took the first step towards her goal, and with her coach's encouraging words, she moves one step closer.

SODA: Conversation Generation

Commonsense triplet selection

- Social triplet from Atomic are retrieved
- *[Head, Tail, Relation]*

Commonsense Knowledge → Narrative

- Triple form → sentence
- Sentence → narrative (prompt InstructGPT)

Narrative → Conversation

- Narrative and speakers as input of prompt
- Full conversation generation with

  Self-chat of Instruct GPT

Example

Speakers:

- Madeleine, Coach

Conversation:

Madeleine: Hey coach, I wanted to talk to you about my performance today. I was really pushing myself and I think I did pretty well. But I'm still not quite where I want to be.

Coach: Well Madeleine, you're progressing nicely. You've come a long way since we first started working together. But if you want to reach your full potential, there's still some work to be done.

Madeleine: I know that. And I'm willing to put in the work. It's just that sometimes I feel like I'm not making as much progress as I should be. Maybe I'm not training hard enough? Or maybe my technique is off?

Coach: It could be a number of things, Madeleine. But don't worry, we'll figure it out together. Let's just keep working hard and see how things go.

Madeleine: Alright, coach. Thanks for the talk.

Coach: No problem. See you at practice tomorrow.

SODA: Filtering

- **Basic Filtering**
  - Conversations with erroneous patterns, too few/little turn
  - Conversations with more than two participants
  - Conversations with non-human speakers

- **Safety Filtering**
  - Detecting contexts that needs caution or intervention (using a Canary safety classifier)
  - Detecting toxic content; e.g., violence, hate (using Rewire API)

- **Commonsense Filtering**
  - A small scale human evaluation via Amazon Mechanical Turk

PersonaChatGen

- An LLM-based method for creating personalized dialgoue datasets
  - Inspired by the human-human PersonaChat dataset
- Uses two GPT-3 with two different persona sets

Persona: a set of personal attributes

- Personal attribute:
  - In textual format: “I like to play soccer” OR
  - In triplet format (entity 1, relation type, entity 2)

(Lee et al., 2023)

PersonaChatGen: Seed Data Generation

Generate personal attribute sentences

- Promt GPT-3 to create persona sentences for a given category

Example

User's persona: **Want | Activity**
Generate five profile sentences related to the given user's persona and the "activity" in each sentence:
1. I always wanted to travel to ireland or puerto rico. (activity: travel)
2. I hope to visit quebec, canada someday. (activity: travel)
3. One day I would really like to skydive. (activity: skydiving)
4. Before I die, I want to skydive. (activity: skydiving)
5. I hope to see the world with my husband. (activity: travel)

User's persona: **Preference | Movie | Title**
Generate five profile sentences related to the given user's persona and the "movie title" in each sentence:
1. I am a big fan of the Lord of the Rings movies. (movie title: Lord of the Rings)
2. I love all of the Harry Potter movies. (movie title: Harry Potter)
3. The Hobbit is one of my favorite movies. (movie title: The Hobbit)
4. I have seen all of the Star Wars movies. (movie title: Star Wars)
5. I enjoy watching Marvel movies. (movie title: Marvel)

PersonaChatGen: Seed Data Generation

Generate personal attribute sentences

- Prompting GPT-3 with a given category
- Post-processing of generated sentences

Attribute Sentences → Persona

- Selecting a random sentence for a category
- Check contradictions with previously added sentences using a Fine-tuned RoBERTa

Example

Persona with contradictory sentences

I am studying at a community college.
I am a teacher at the high school.
"The Great Gatsby" is another book I enjoy.
I'm a big fan of the violin. I love reading books that are full of adventure.

PersonaChatGen: Conversation Generation

Generate personal attribute sentences

- Prompting GPT-3 with a given category
- Post-processing of generated sentences

Attribute Sentences → Persona

- Selecting a random sentence for a category
- Checking contradictions with previously added sentences using a fine-tuned RoBERTa

Persona → Conversation

- Two GPT-3 with difference personas

Example

Persona:
<FEWSHOT PERSONA>
The following is a daily conversation with your friend **implicitly** containing the given persona.
<FEWSHOT CONV>

Persona:
<TARGET PERSONA>
The following is a daily conversation with your friend **implicitly** containing the given persona.
You:

PersonaChatGen: Filtering

Generate personal attribute sentences

- Prompting GPT-3 with a given category
- Post-processing of generated sentences

Attribute Sentences → Persona

- Selecting a random sentence for a category
- Checking contradictions with previously added sentences using a fine-tuned RoBERTa

Persona → Conversation

- Two GPT-3 with difference personas

Filtering Copy–Paste, Persona Consistency, Toxicity