

Data Augmentation for Conversational AI

The Web Conference 2024



Tutorial website

Presenters



Heydar Soudani

PhD Candidate
Radboud University
heydar.soudani@ru.nl



Evangelos Kanoulas

Full Professor
University of Amsterdam
e.kanoulas@uva.nl



Roxana Petcu

PhD Candidate
University of Amsterdam
r.m.petcu@uva.nl



Faegheh Hasibi

Assistant Professor
Radboud University
f.hasibi@cs.ru.nl

Part 3: Conversation Generation - Open Domain

Duration: 30 mins

Presenter: Heydar Soudani

Open Domain Dialogue (ODD) System

Definition

- Engage users in conversations across a wide variety of topics
 - without being confined to specific tasks or domains

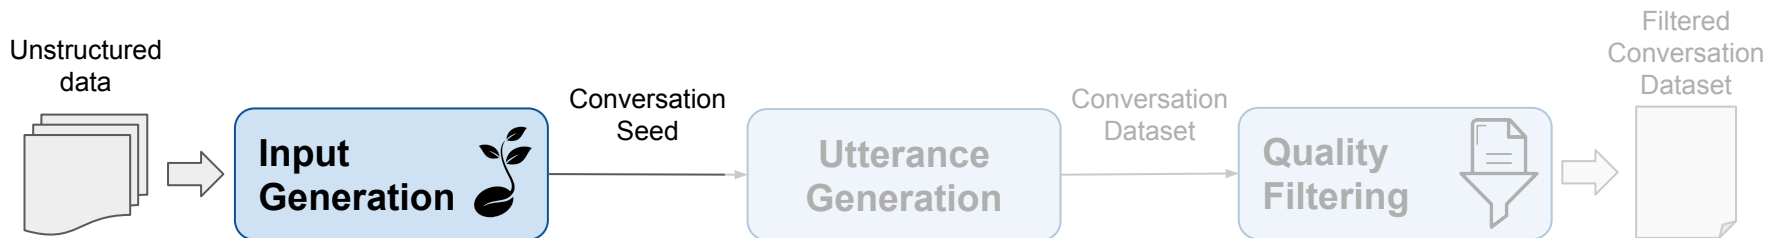
(Ni et al. 2023)

Key Features of ODD

- **Coherence:** Conversation's turns meaningfully connect to each other
- **Diversity:** Avoid bland and repetitive responses & encourage engaging interactions
- **Generality:** Encompass a broad spectrum of topics
- **Informativeness:** Elicit informative responses, knowledgeable and relevant conversations

(Mehri et al., 2020) (Hwang et al., 2022) (Hwang and Lee, 2022)

ODD Data Generation



Example

(Kim et al., 2023)



Conversation Seed

Participants:

Madeleine, Coach

Description:

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

ODD Data Generation



Example

(Kim et al., 2023)



Conversation Seed

Participants:

Madeleine, Coach

Description:

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

Prompt GPT-3.5 to generate a full, multi-turn conversation

Conversation Sample

Madeleine: Hey coach, I wanted to talk to you about my performance today. I was really pushing ...

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, ...

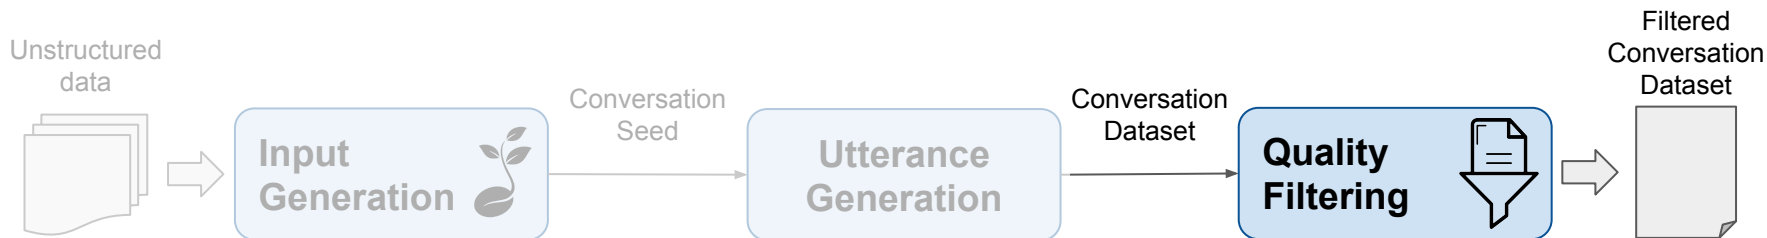
Madeleine: I know that. And I'm willing to put in the work. It's just that ...

Coach: It could be a number of things, Madeleine. But don't worry, we'll figure it ...

Madeleine: Alright, coach. Thanks for the talk.

Coach: No problem. See you at practice tomorrow

ODD Data Generation



Example

(Kim et al., 2023)



Conversation Seed

Participants:
Madeleine, Coach

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

Prompt GPT-3.5 to generate a full, multi-turn conversation

Conversation Sample

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.

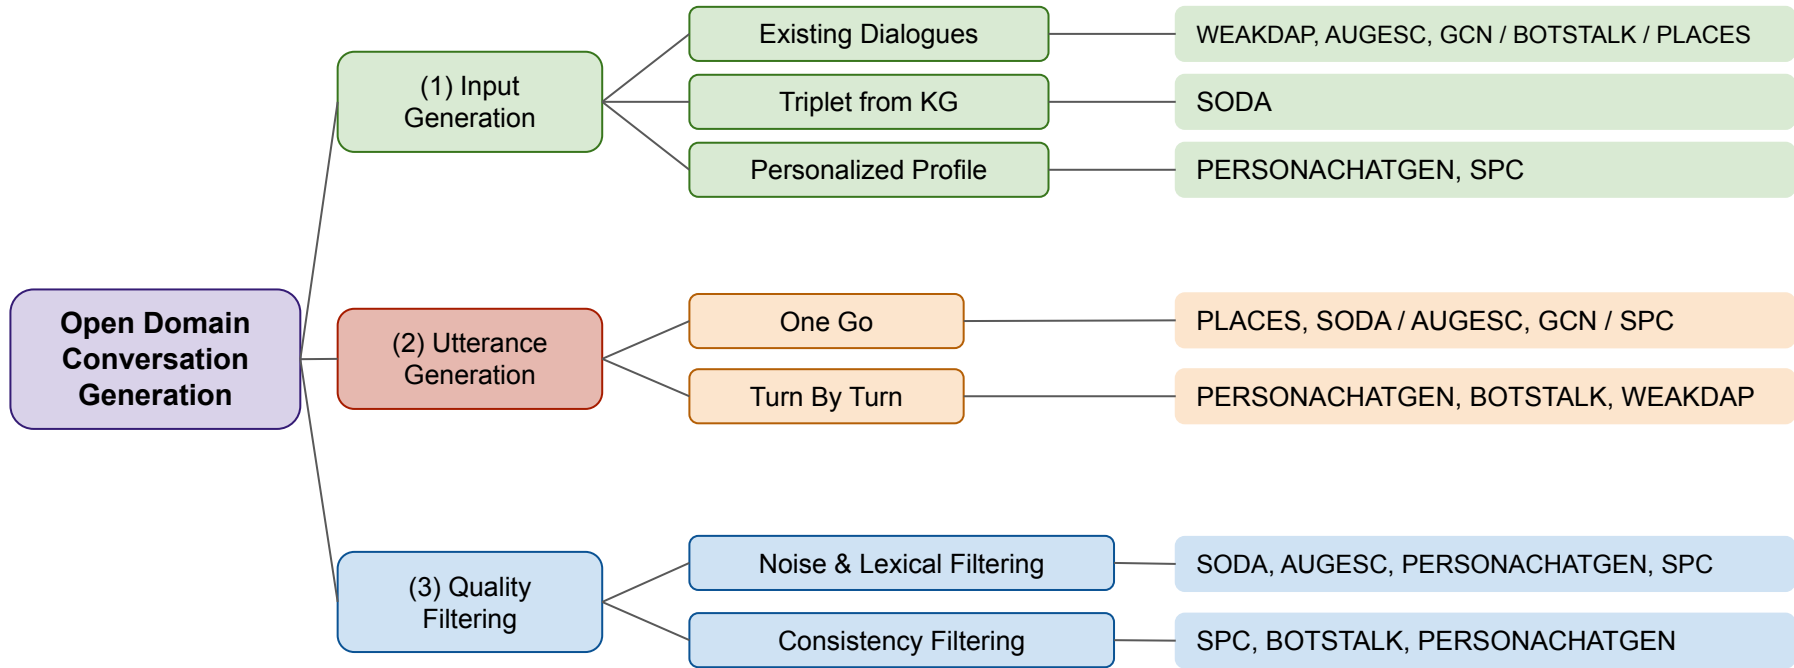
Coach: Well Madeleine, you're progressing nicely. You've come a long way since we first started working together. But if you want to reach your full potential, there's still some work to be done.

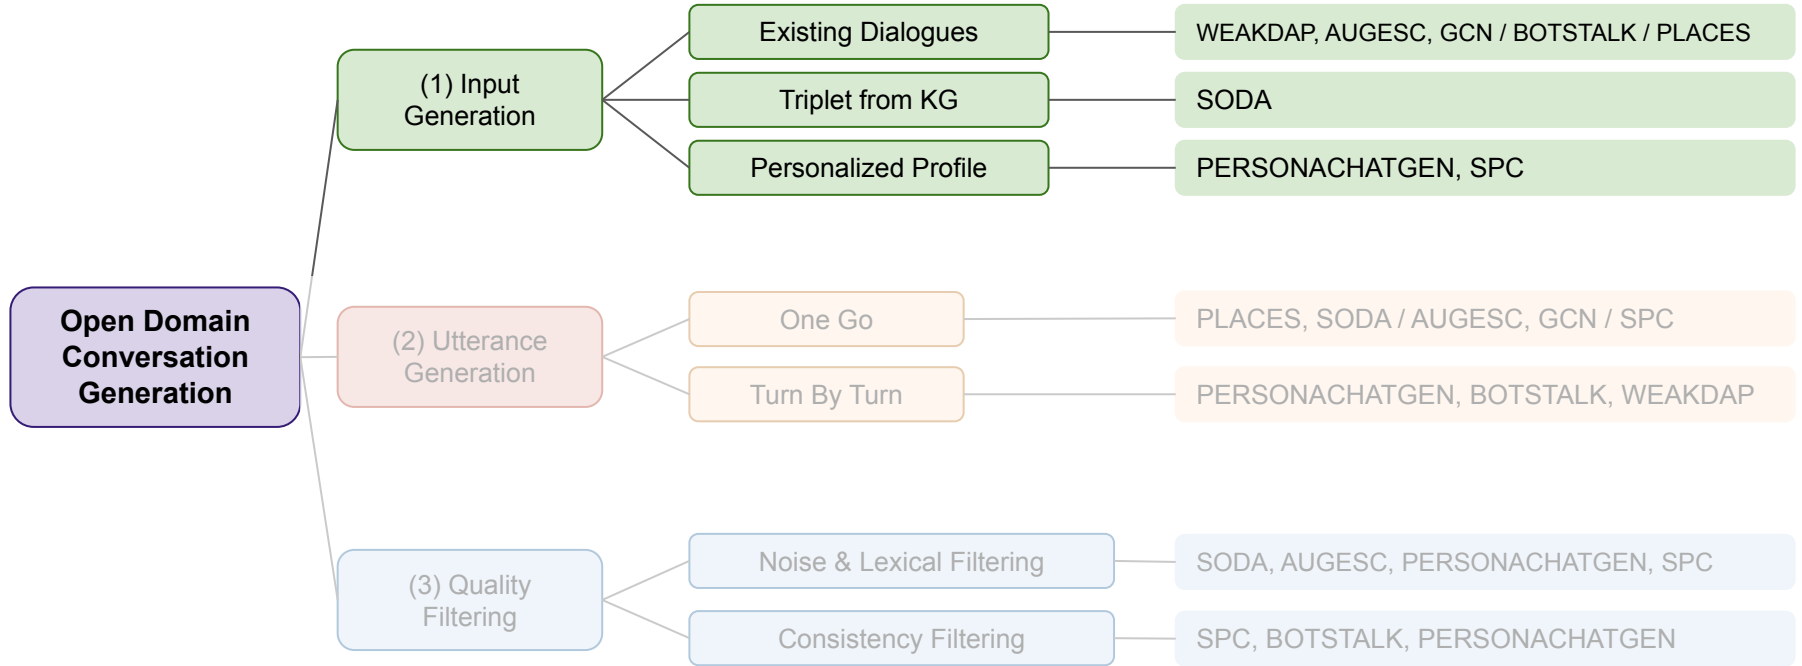
Madeleine: I know that. And I'm willing to put in the work. It's just that sometimes I feel like I'm not making as much progress as I should be. Maybe I'm not training hard enough? Or maybe my technique is off?

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





Input Generation

- What is the conversation seed and why has it been defined for the generation process
- **Conversation Seed:** An information card containing a main topic, subtopics, and key details about the topic
- The conversation about this topic is going to take place

(Kim et al., 2023), (Zheng et al., 2023)

Example from ESConv dataset

(Liu et al., 2021)

Conversation Seed

Experience type: Previous experience

Emotion type: Anxiety

Problem type: job crisis

Situation: I hate my job but I am scared to quit
and seek a new career



Conversation Sample

Supporter: Hello, what would you like to talk about?

Seeker: I am having a lot of *anxiety* about *quitting my current job*. It is *too stressful* but pays well.

Supporter: What makes your *job stressful* for you?

Seeker: I have to deal with many *people in hard financial situations* and it is *upsetting*.

Supporter: Do you help your clients to make it to a better financial situation?

...

Input Generation - Existing Dialogues

Using background information directly

The task description and starting utterance are selected from existing dataset

AugESC

(Zheng et al., 2023)

The following is a conversation with an AI assistant. The assistant is helpful, empathetic, clever, and very friendly. It can use various support skills to provide emotional support to human.

Human: I moved into a new state recently, and there's a lot to do, but I don't have any friends in the new place I stay at.

AI: What's it like being away from family?

Human: Family is all I have here. They aren't exactly close, so I haven't gotten to see them in a while.

AI: That must be difficult. How do you feel about where you live?

Human: It's OK. I'm learning to like it a little bit. At least now I have someone who is usually around when I wake up.

AI: If only you were that lucky with people in general. People move for so many different reasons. I've found that often when I move, I just need to adjust my social circle a little, and I find that I then end up liking where I am.

Human: That's true. Maybe I should just find some people to hang out with.

Human generates background information

Given a list of topics and tasks, humans are asked to generate background info and some dialogue samples

PLACES (Chen et al., 2022)

Topic: Relationships

Background info: Bob got engaged

<Conversation 0>

The following is a conversation between Alice and Bob about **relationships**. **Bob recently got engaged.**

Alice: Congrats on your engagement! When do you think you will have your wedding?

Bob: Thank you!! We're thinking of having it in November.

Alice: That's amazing! Will you pick a fancy destination?

...

Input Generation - Triplet from KG

Example



- Input: A Knowledge Graph
- **SODA:** social dialogue
- Generation Technique:
 - Sample a socially relevant triplet
 - Define the conversation participants

(Kim et al., 2023)

Triplet from Atomic 10x:

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

Name participants:

Speakers: Madeleine, Coach

Input Generation - Triplet from KG

- Input: A Knowledge Graph
- Generation Technique:
 - Sample a triplet
 - Define the conversation participants
 - Convert the triplet to a sentence
 - Expand the sentence

(Kim et al., 2023)

Example



Triplet to Sentence:

Madeleine took the first step. Madeleine moves a step closer to the goal

Sentence to Description:

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

Conversation Seed

Input Generation - Personalized Profile

Personalized Dialogue Systems

- User Profile (UP)
- Profile Sentences (PS)
 - Contain personalized information about the user

Fig. from PersonaChat dataset (Zhang et al., 2018)

Persona 1	Persona 2
I like to ski My wife does not like me anymore I have went to Mexico 4 times this year I hate Mexican food I like to eat cheetos	I am an artist I have four children I recently got a cat I enjoy walking for exercise I love watching Game of Thrones

[PERSON 1:] Hi
[PERSON 2:] Hello ! How are you today ?
[PERSON 1:] I am good thank you , how are you.
[PERSON 2:] Great, thanks ! My children and I were just about to watch Game of Thrones.
[PERSON 1:] Nice ! How old are your children?
[PERSON 2:] I have four that range in age from 10 to 21. You?
[PERSON 1:] I do not have children at the moment.
[PERSON 2:] That just means you get to keep all the popcorn for yourself.
[PERSON 1:] And Cheetos at the moment!
[PERSON 2:] Good choice. Do you watch Game of Thrones?
[PERSON 1:] No, I do not have much time for TV.
[PERSON 2:] I usually spend my time painting: but, I love the show.

Input Generation - Personalized Profile

Example



Generation steps

- 1) Collect/Generate a pool of PS
 - Given a list of topics, a LLM generates
 - Using PSs from PersonaChat dataset
 - Prompt a LLM to generate more PS
- 2) Group number of PS to create a UP
 - Contradiction score using NLI classifier
- 3) Filtering
 - Heuristic: Do not follow the template

(Lee et al., 2022), (Jandaghi et al., 2023)

User's persona: Want | Activity

Generate five profile sentences related to the given user's persona and the "activity" in each sentence:

1. I have 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)

Input Generation - Personalized Profile

Example



Generation steps

- 1) Collect/Generate a pool of PS
 - Given a list of topics, a LLM generates
 - Using PSs from PersonaChat dataset
 - Prompt a LLM to generate more PS
- 2) Group number of PS to create a UP
 - Contradiction score using NLI classifier
- 3) Filtering
 - Heuristic: Do not follow the template

(Lee et al., 2022), (Jandaghi et al., 2023)

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.

(a) An example of persona set containing contradiction between profile sentences

I am a very creative and imaginative person.

My older sister is a doctor.

I love to read books that are science fiction.

I enjoy watching suspenseful movies.

I have to be very careful in the springtime because of my allergies.

(b) An example of persona set containing no contradiction between profile sentences

Input Generation - Personalized Profile

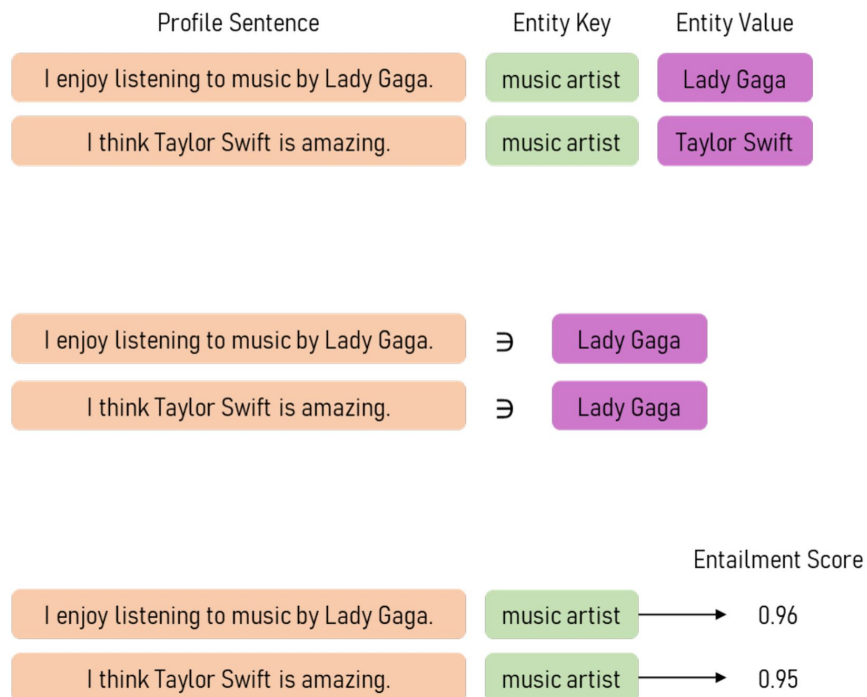
Example



Generation steps

- 1) Collect/Generate a pool of PS
 - Given a list of topics, a LLM generates
 - Using PSs from PersonaChat dataset
 - Prompt a LLM to generate more PS
- 2) Group number of PS to create a UP
 - Contradiction score using NLI classifier
- 3) Filtering
 - Heuristic: Do not follow the template

(Lee et al., 2022), (Jandaghi et al., 2023)



Component 2: Utterance Generation

Objective: Convert “Conversation Seed” to “Conversation Sample”

Example from PERSONACHATGEN

(Lee et al., 2022)

Conversation Seed



P1's Persona:

I love food and I love to eat.
I am a woman who loves fashion.
I love reality TV.
I prefer to watch comedies.
I have hay fever.

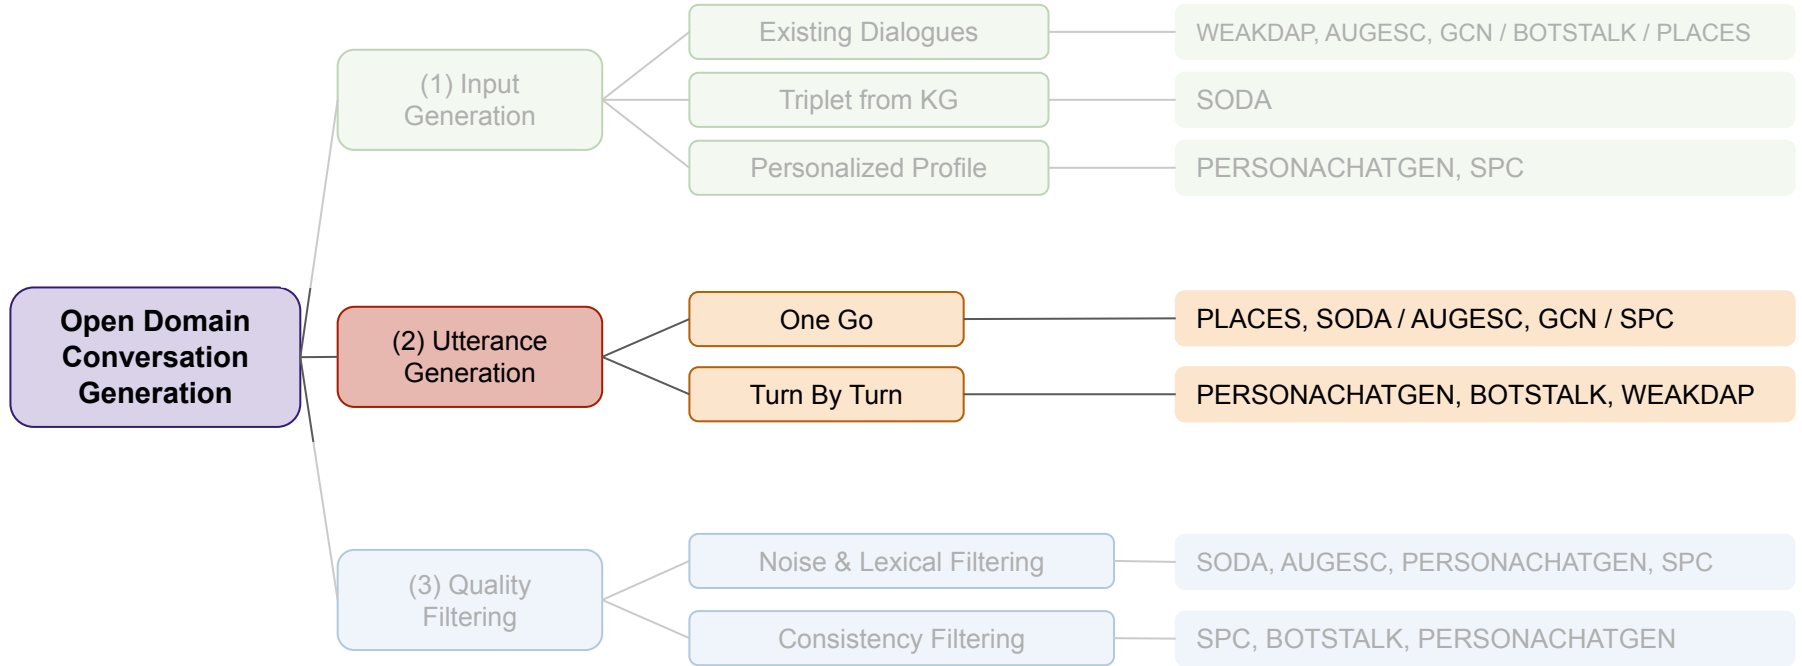
P2's Persona:

I've been to Italy three times.
I graduated from Yale.
I've read all of the books by Jodi Picoult.
I'm a big fan of books, and my favorite genre is fantasy.
I have asthma and it makes it hard to breathe sometimes.



Conversation Sample

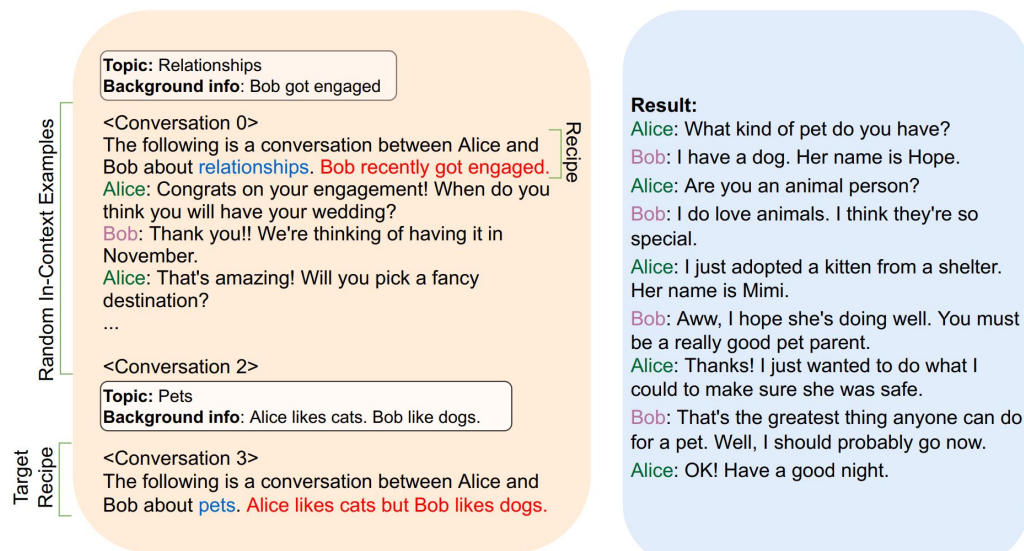
P1: Hey, how are you?
P2: I'm good, just got back from Italy.
P1: Oh, how was it?
P2: It was great, I love it there.
P1: What did you eat?
P2: Everything, I love Italian food.
P1: That sounds amazing, I need to go there sometime.
P2: Yeah, you should definitely go. It's a great place.
P1: I also love fashion, do you follow any trend?
P2: No, not really. I prefer classic styles.
P1: That's cool, I like to watch reality TV.
P2: Yeah, I don't really watch much TV. I prefer to read.
P1: That's great, what do you like to read?
P2: I like to read books, my favorite genre is fantasy.
P1: That's great, I have hay fever so I can't really read outside.
P2: Yeah, I have asthma so it's hard for me to breathe sometimes.



Utterance Generation - One Go

1) ICL: Only prompting, wo. Fine tuning

(Chen et al., 2022)

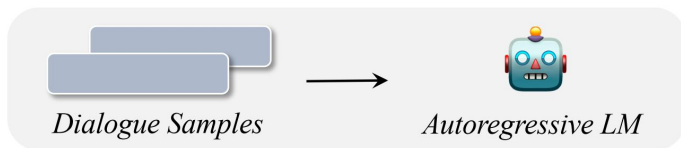


Utterance Generation - One Go

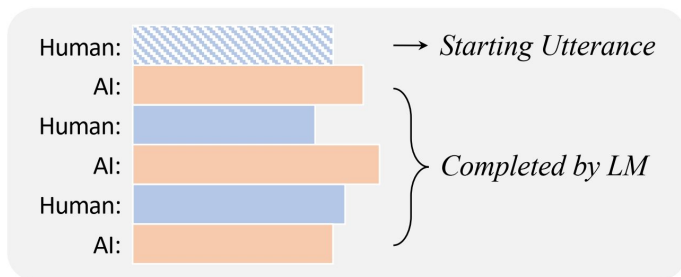
2) First Fine-tune on Dialogue Completion task, then prompt

(Zheng et al., 2023)

(1) Fine-tuning LM



(2) Dialogue Completion



The following is a conversation with an AI assistant. The assistant is helpful, empathetic, clever, and very friendly. It can use various support skills to provide emotional support to human.

Human: I moved into a new state recently, and there's a lot to do, but I don't have any friends in the new place I stay at.

AI: What's it like being away from family?

Human: Family is all I have here. They aren't exactly close, so I haven't gotten to see them in a while.

AI: That must be difficult. How do you feel about where you live?

Human: It's OK. I'm learning to like it a little bit. At least now I have someone who is usually around when I wake up.

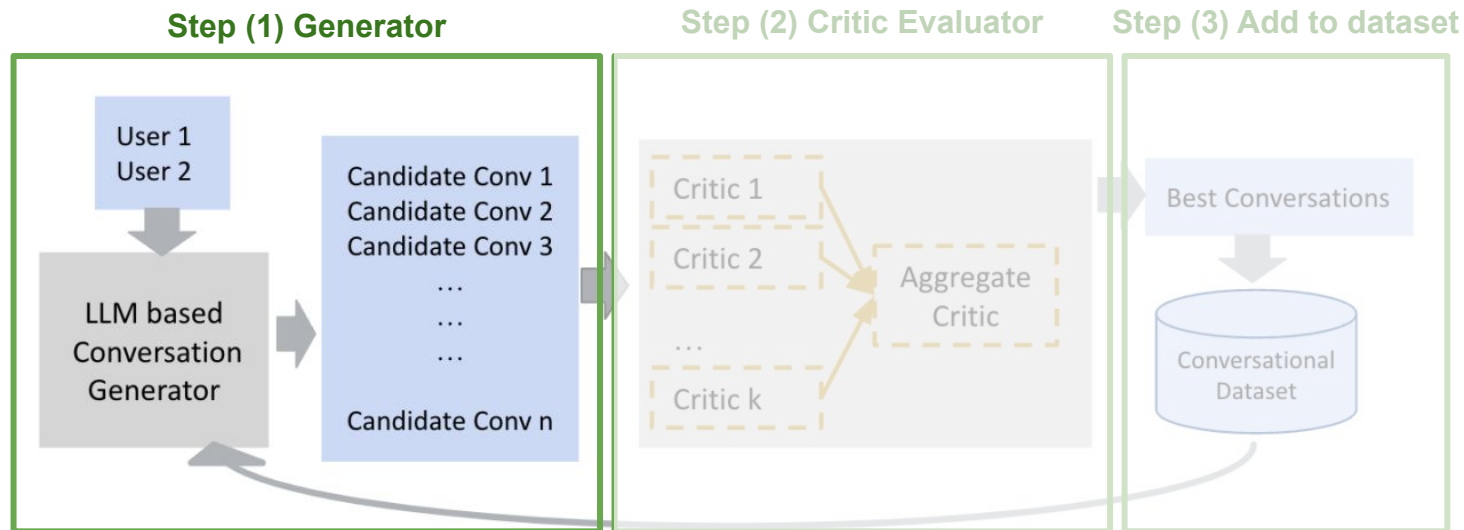
AI: If only you were that lucky with people in general. People move for so many different reasons. I've found that often when I move, I just need to adjust my social circle a little, and I find that I then end up liking where I am.

Human: That's true. Maybe I should just find some people to hang out with.

Utterance Generation - One Go

3) Prompting with Generator-Critic Architecture

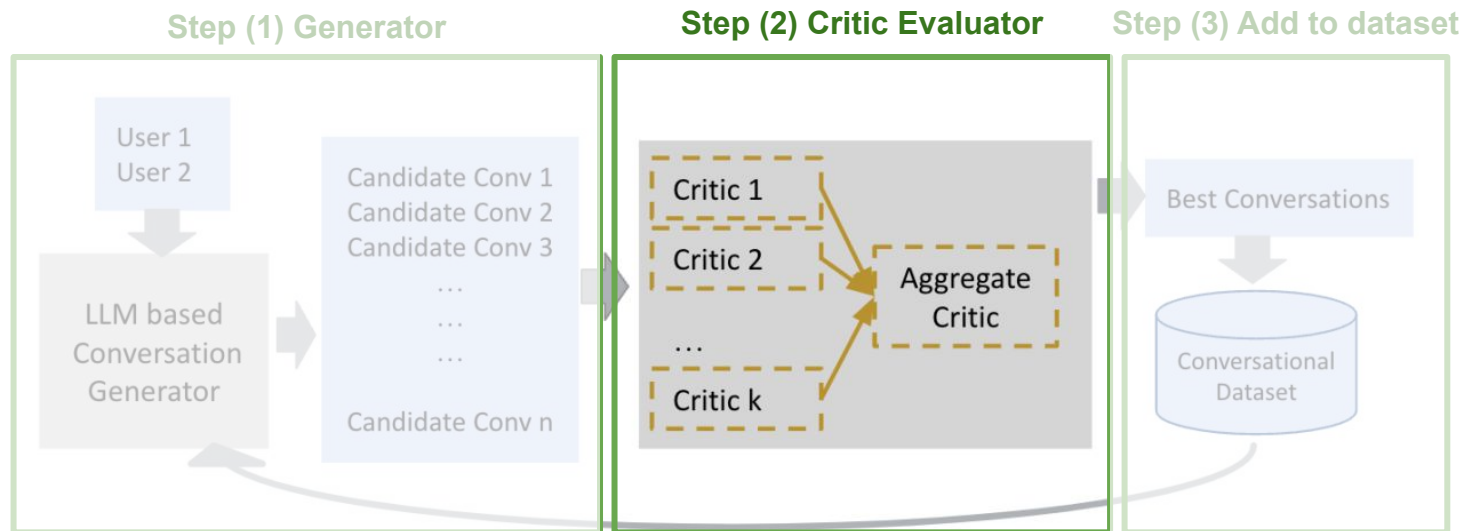
(Jandaghi et al., 2023)



Utterance Generation - One Go

3) Prompting with Generator-Critic Architecture

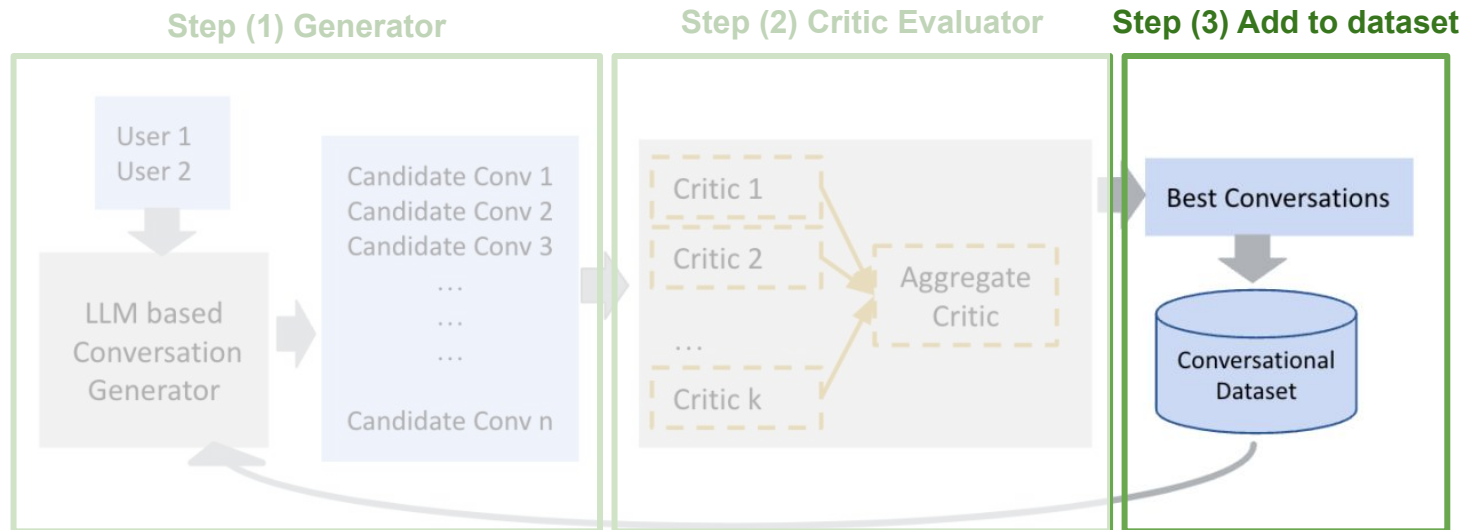
(Jandaghi et al., 2023)



Utterance Generation - One Go

3) Prompting with Generator-Critic Architecture

(Jandaghi et al., 2023)



Utterance Generation - Turn-by-Turn

(1) Two persona Profiles

- two LLMs, User simulation (Lee et al., 2022)

Reason 2: Merging multiple conversation

datasets (Kim et al., 2022)

Person A

Person B

Skill context from ConvAI2 I like to play soccer; I like to read; ...	Skill context from ConvAI2 I have 3 children; I am a karate black belt; ...
Skill context from WoW Nike Inc.	Skill context from WoW Nike Inc.; ... multinational corporation ... Air Jordan ...
Skill context from ED I really like this girl at my job, but I am ... ; Apprehensive	Skill context from ED None

Dialogue

A: Do you have much experience using the different types of cleats? Which do you like best? (P)

B: I have a little. I also know about Air jordans, a brand of footwear also popular with athletes. (K)

A: I enjoy Air jordans as well. I like to play soccer, and it's extremely hard to get good footwear. (P)

B: I agree. Air jordans are good for calves, and it's really easy to pull off. (K)

A: Air jordans are generally made with the most material so it makes sense they'd be easy to use. (K)

B: And now shoes can go together with clothes as well, like any other type of material. (K)

A: Do you know much about aeros then? They have excellent fit and beauty. (K)

B: I have a hard time finding it but they are great shoes. (P)

A: I hope you can find ones that are comfortable to you. (E)

B: Yes. I hope you can get those shoes too. (E)

Utterance Generation - Turn-by-Turn

Reason 3: More diversity and quantity

Trajectory Augmentation

All-turn Augmentation

Last-turn Augmentation

(Chen et al., 2022)

Original Conversation

Turn 1: Alice in a happy mood: Oh, man. I had the best supper last night. My wife made a stir-fry and it was amazing!

Turn 2: Bob in a happy mood: I love stir fry crispy bitesize vegetables covered in a mixture of soy sauce and oyster sauce. Wilted greens and fresh bean sprouts. Throw in some onion and garlic and ginger! Mmm! Mmm! It's almost lunchtime. I would die for a plate of stir fry right now!

Turn 3: Alice in a neutral mood: Well, you can keep the vegetables, I'll take the meat. The stir fry my wife made was really hearty, with chunks of beef and slivers of bell peppers and onion...

Turn 4: Bob in a surprised mood: What? You call that a stir fry? More meat than vegetables? That's the worst insult you could throw at a Chinese stir fry. What disgrace to the wok she fried it in! What you had is equivalent to a fajita without the wrap!

Output

Augmented Conversation

GT Turn 1: Alice in a happy mood: Oh, man. I had the best supper last night. My wife made a stir-fry and it was amazing!

GT Turn 2: Bob in a happy mood: I love stir fry crispy bitesize vegetables covered in a mixture of soy sauce and oyster sauce. Wilted greens and fresh bean sprouts. Throw in some onion and garlic and ginger! Mmm! Mmm! It's almost lunchtime. I would die for a plate of stir fry right now!

Generated Turn 3: Alice in a neutral mood: Stir fry is good. When my dad cooks, he makes stir fry for us.

Generated Turn 4: Bob in a surprised mood: Stir fry was was the only Asian food I had until I was ten. I've never known what it was called!

Utterance Generation - Turn-by-Turn

Reason 3: More diversity and quantity

Trajectory Augmentation

All-turn Augmentation

Last-turn Augmentation

(Chen et al., 2022)

Output



Original Conversation

Turn 1: Alice in a surprised mood: My goodness! She is thirty seconds faster than the world record in 5000 metres race.

Turn 2: Bob in a surprised mood: Excuse me, what did you say?

Turn 3: Alice in a happy mood: A Chinese girl has broken the world record in the Olympic Games.

Turn 4: Bob in a surprised mood: That's incredible. I can't believe it.

Turn 5: Alice in a happy mood: You have to. It is sure.

Turn 6: Bob in a surprised mood: So, what's her name?

Augmented Conversation 1

GT Turn 1: Alice in a surprised mood: My goodness! She is thirty seconds faster than the world record in 5000 metres race.

GT Turn 2: Bob in a surprised mood: Excuse me, what did you say?

Generated Turn 3: Alice in a happy mood: I said my goodness, that girl is fast.

Augmented Conversation 2

GT Turn 1: Alice in a surprised mood: ...

GT Turn 2: Bob in a surprised mood: ...

GT Turn 3: Alice in a happy mood: A Chinese girl has broken the world record in the Olympic Games.

Generated Turn 4: Bob in a surprised mood: I am dying to know who is the girl?

Augmented Conversation 3

GT Turn 1: Alice in a surprised mood: ...

GT Turn 2: Bob in a surprised mood: ...

GT Turn 3: Alice in a happy mood: A Chinese girl has broken the world record in the Olympic Games.

GT Turn 4: Bob in a surprised mood: That's incredible. I can't believe it.

Generated Turn 5: Alice in a happy mood: This is the first time in the history of the world

Augmented Conversation 4

...

Utterance Generation - Turn-by-Turn

Reason 3: More diversity and quantity

Trajectory Augmentation

All-turn Augmentation

Last-turn Augmentation

(Chen et al., 2022)

Output

Augmented Conversation

GT Turn 1: Alice informs Bob: Good morning.

GT Turn 2: Bob informs Alice: Er, good morning, yes, er...

GT Turn 3: Alice directs Bob: I'm phoning about the job that was in the paper last night.

GT Turn 4: Bob directs Alice: Oh, yes. Erm, well, could you tell me your name, please?

GT Turn 5: Alice informs Bob: Oh, Candida Fawcett.

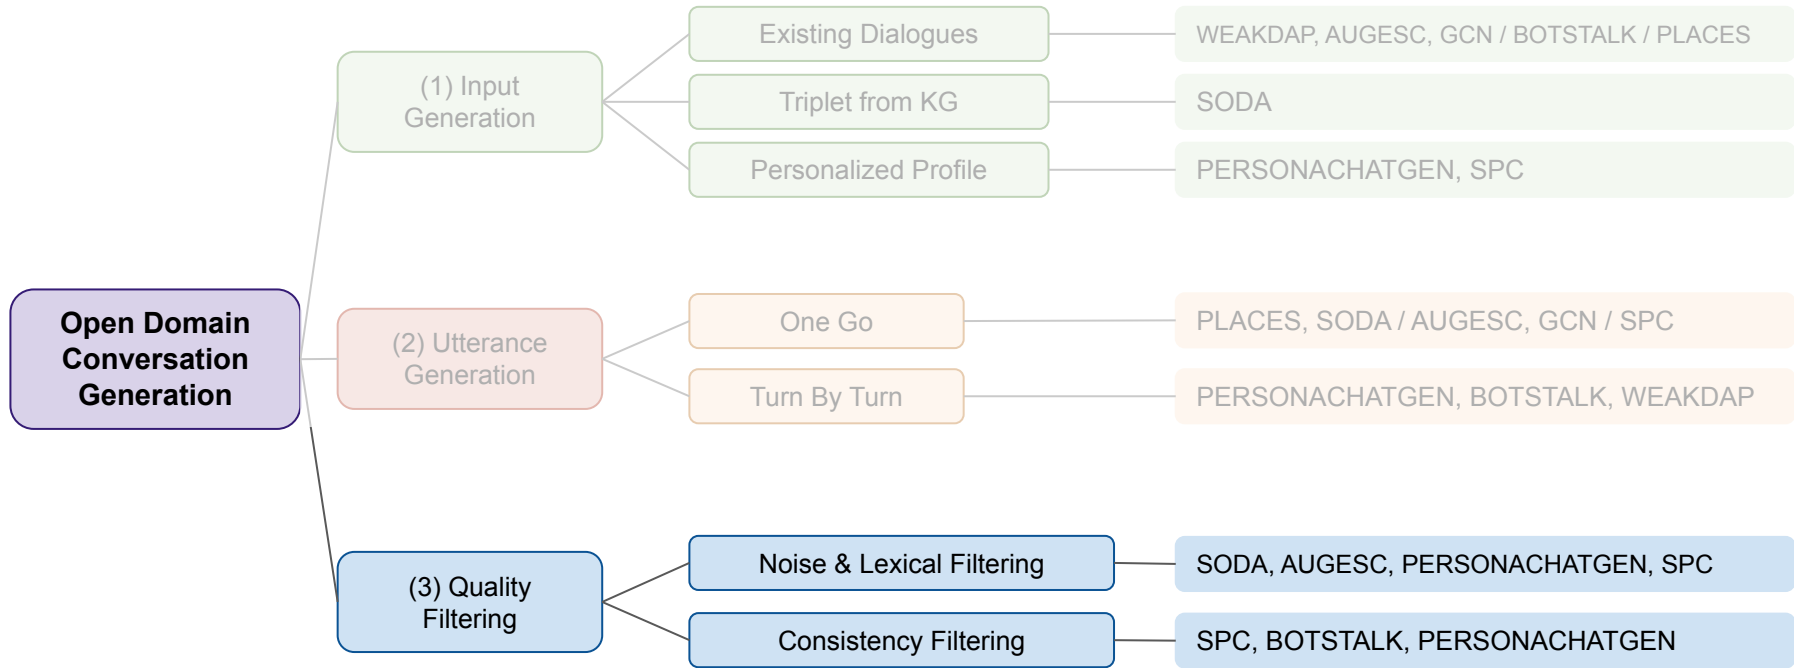
GT Turn 6: Bob questions Alice: Oh, yes. Erm, well, what exactly is it that interests you about the job?

GT Turn 7: Alice informs Bob: Well, I just thought that it was right up my street you know.

GT Turn 8: Bob directs Alice: Really, hmmm. Erm, well, could you perhaps tell me a little about yourself?

GT Turn 9: Alice informs Bob: Yes, erm, I'm 23. I've been working abroad, I'm um...

Generated Turn 10: Bob informs Alice: Well, perhaps you would like to think it over. Erm, and er...



Component 3: Quality Filtering

Objective

- Remove conversations do not contain ODD features
 - (e.i. Correctness, Diversity, Consistency, Informativeness)

Noise & Lexical Filtering:

Checking the Correctness, Diversity

Approach: Heuristic rules

- Unfinished conversations
- Do not follow the wanted patterns
- Contain repetitive pattern
- Dangerous or Toxic context with social bias

(Kim et al., 2023), (Zheng et al., 2023), (Lee et al., 2022)

Consistency Filtering:

Checking the consistency

- Between the turns
- Between persona sentences in one user profile

Approach: NLI Classifier

(Kim et al., 2022), (Lee et al., 2022), (Jandaghi et al., 2023)