Like hiking? You probably enjoy nature: Persona-grounded Dialog with Commonsense Expansions

Bodhisattwa P. Majumder[1], Harsh Jhamtani[2], Taylor Berg KirkPatrick[1], Julian McAuley[1]

UC San Diego[1], Carnegie Mellon University[2]

Persona-grounded Dialog Generation

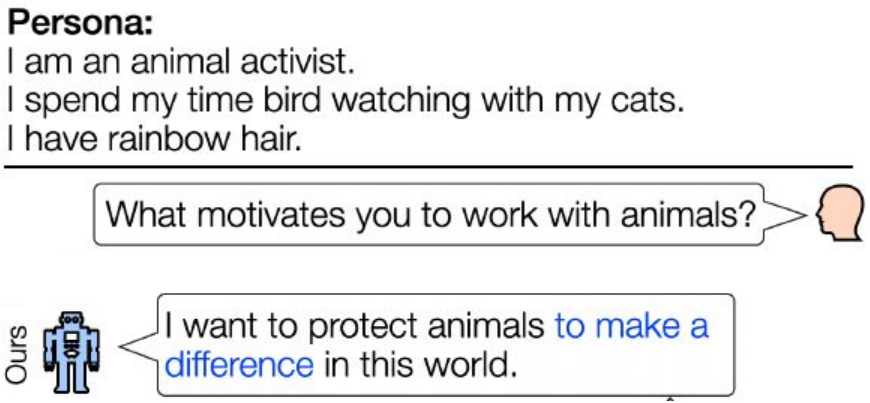

PersonaChat:
10,907 dialogs between pairs of speakers where each speaker follows their own persona

Imagine a conversational bot with a **persona**

Given a persona and a dialog history, we have to predict the next response

It is possible that some responses are based on given persona, some are purely based on dialog history

A persona means more than what you see

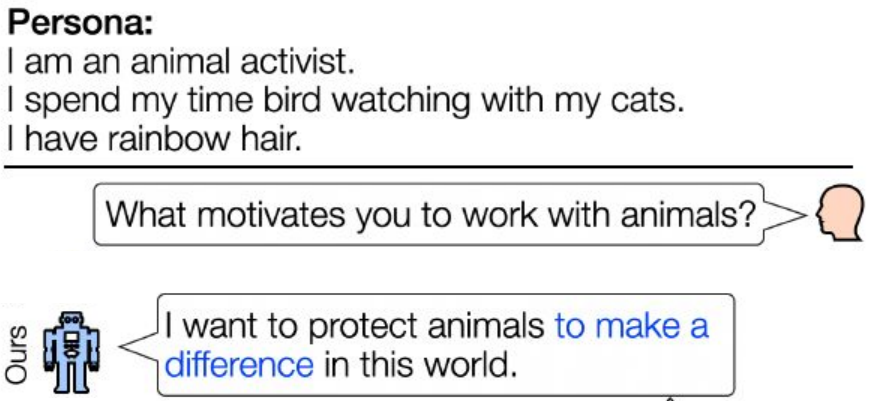

A persona consists of 3-5 sentences

A persona is a collection of **real world events** mimicking a user's preference in a real world

A persona means more than what you see

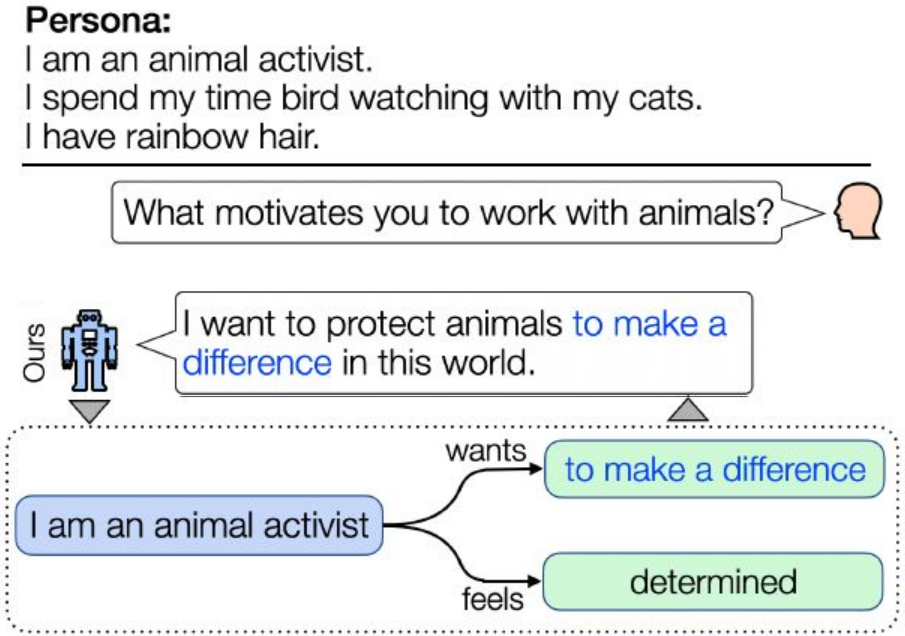

A persona consists of 3-5 sentences

A persona is a collection of **real world events** mimicking a user's preference in a real world

Often these persona **means more** than their surface forms

Previous works fails to address implicit context

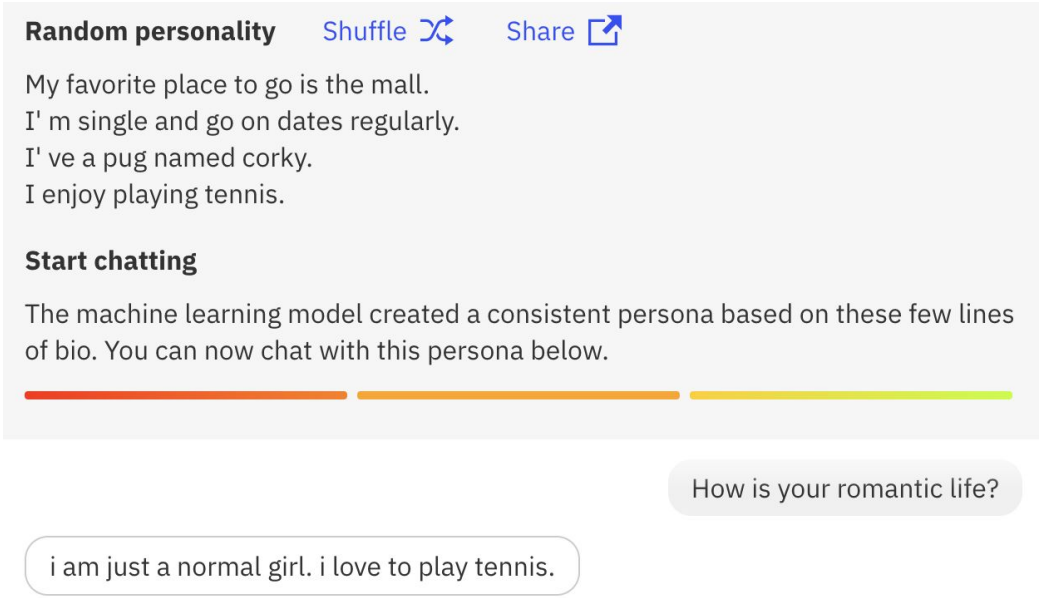

Model fails to correlate ‘romantic life’ with being ‘single’

Main issues in previous works

- Previous models had to **implicitly learn** possible implications/entailments of given persona
  - difficult to learn; especially some sparsely occuring patterns

- Previous models **cram entire persona** into a Transformer
  - lack of controllability or interpretability

Our goals

- Can we leverage **external commonsense knowledge** bases to expand the given persona with related concepts and implications?

- Can we achieve better **interpretability** on model outputs using fine-grained persona selection?

- Can we impart better **controllability** with respect to provided persona?

Commonsense Expansions

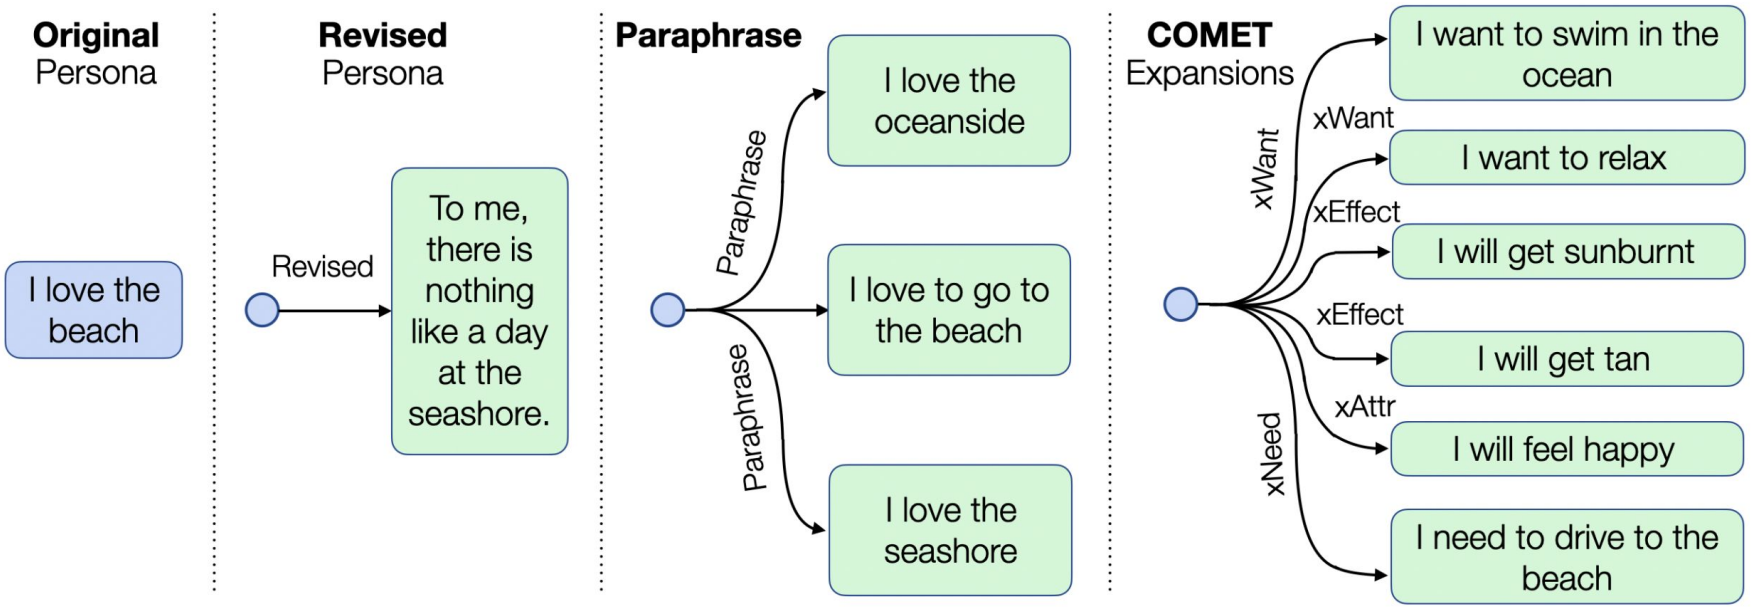

Paraphrase

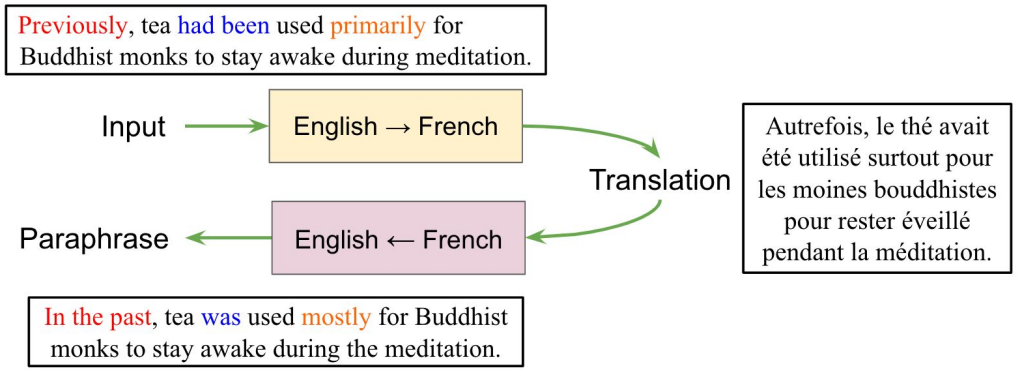

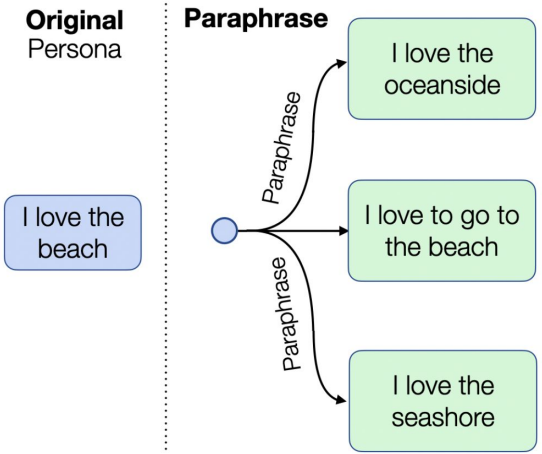

Paraphrasing using back-translation

Paraphrases use synonymous phrases or manipulate word-syntax of the original sentence, which implicitly involves both **context comprehension** and **world knowledge**

COMET

Commonsense Transformers (**COMET**)

Trained on ATOMIC - an inferential commonsense KB of real-world events

Each persona sentence revealed 5 plausible expansions for 9 causal effects (i.e. in total 45)

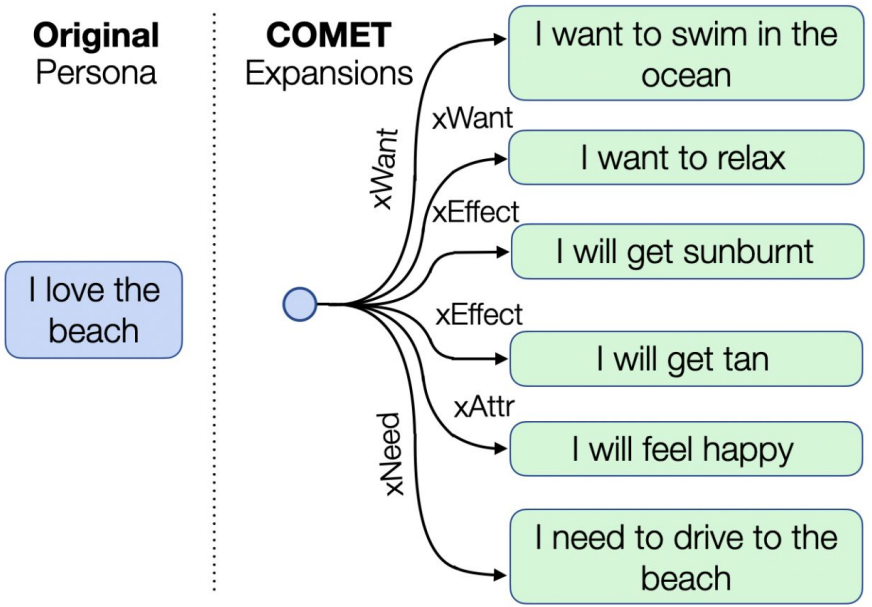

Commonsense and Persona Aligned Chatbot (COMPAC)

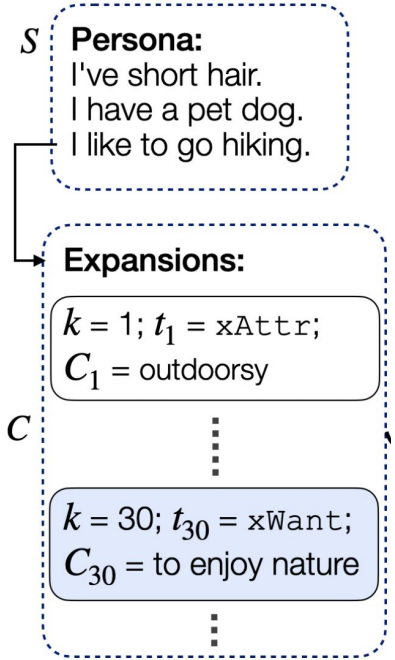

- **Expand** all persona using commonsense

Commonsense and Persona Aligned Chatbot (COMPAC)

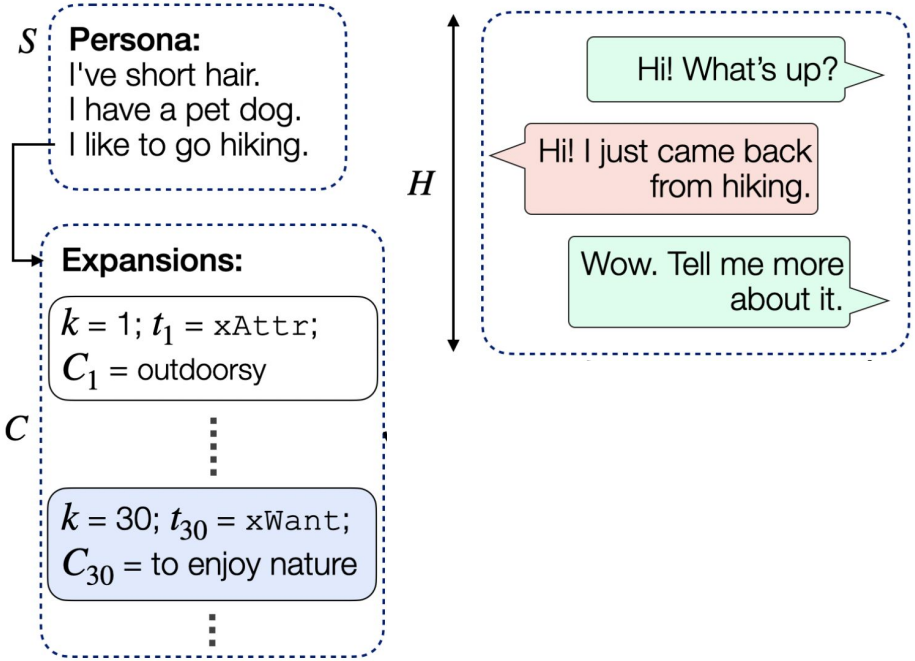

- Expand all persona using commonsense
- A **dialog history**, H, is observed

Commonsense and Persona Aligned Chatbot (COMPAC)

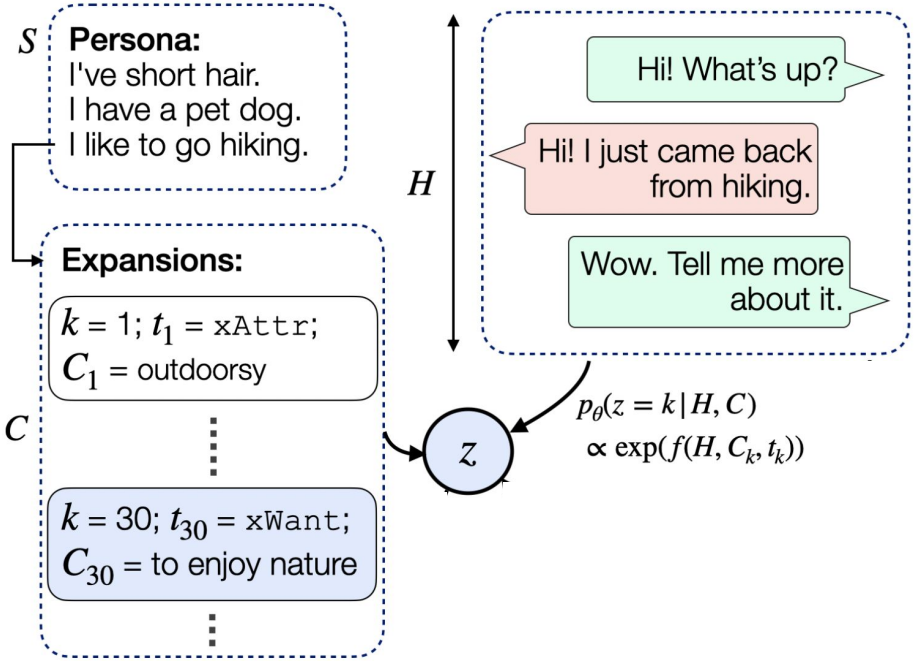

- Expand all persona using commonsense
- A dialog history, H, is observed
- A **relevant persona** is modeled using a latent discrete random variable **z**; Given the dialog history H, we **sample** a persona sentence C_z

Commonsense and Persona Aligned Chatbot (COMPAC)

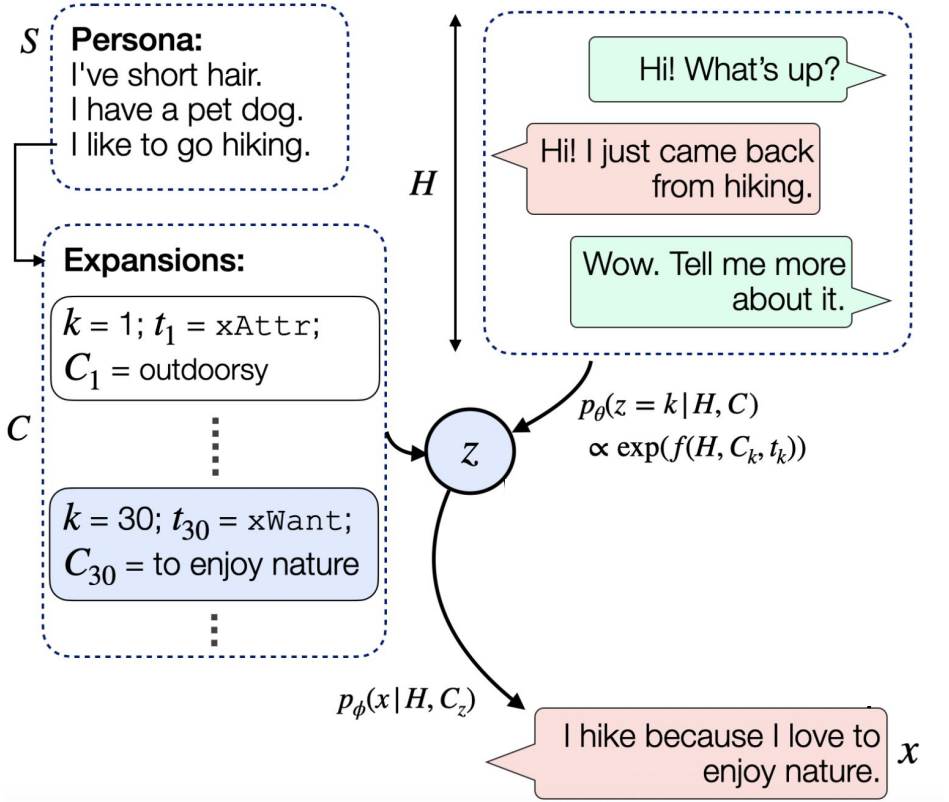

- Expand all persona using commonsense
- A dialog history, H, is observed
- A relevant persona is modeled using a latent discrete random variable z; Given the dialog history H, we sample a persona sentence C_z
- A **grounded response**, **x**, is generated w.r. to C_z and H

Commonsense and Persona Aligned Chatbot (COMPAC)

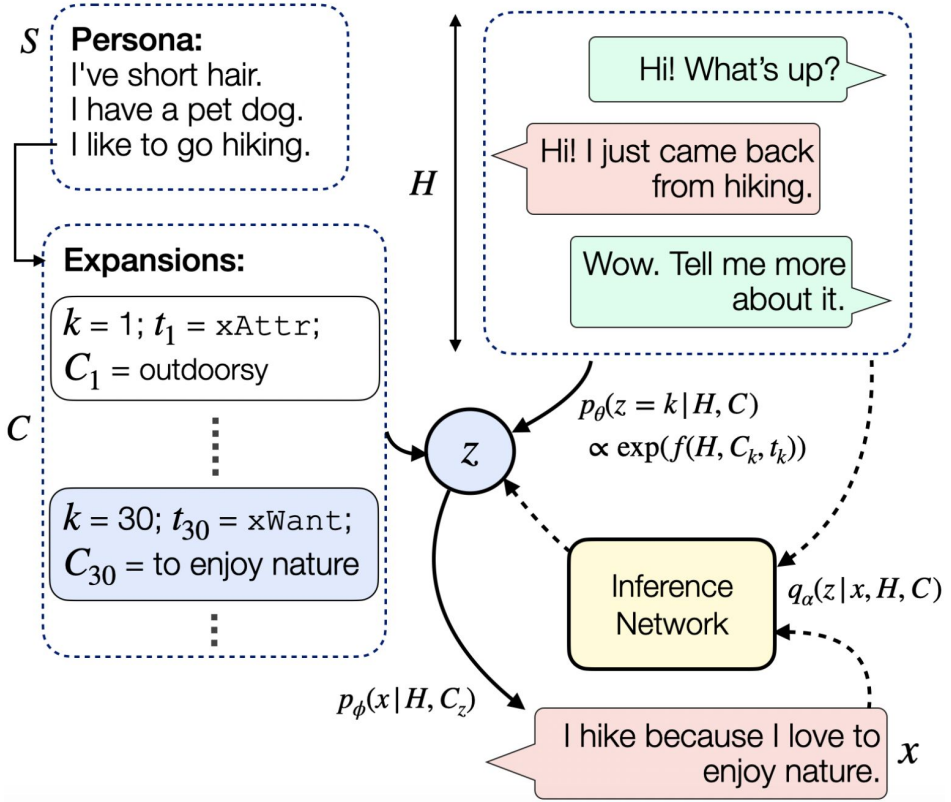

- Expand all persona using commonsense
- A dialog history, H, is observed
- A relevant persona is modeled using a latent discrete random variable z; Given the dialog history H, we sample a persona sentence C_z
- A grounded response, x, is generated w.r. to C_z and H
- Since space of z is large, we instead optimize a lower bound of log-likelihood using **variational inference**

Results: Automatic Metrics

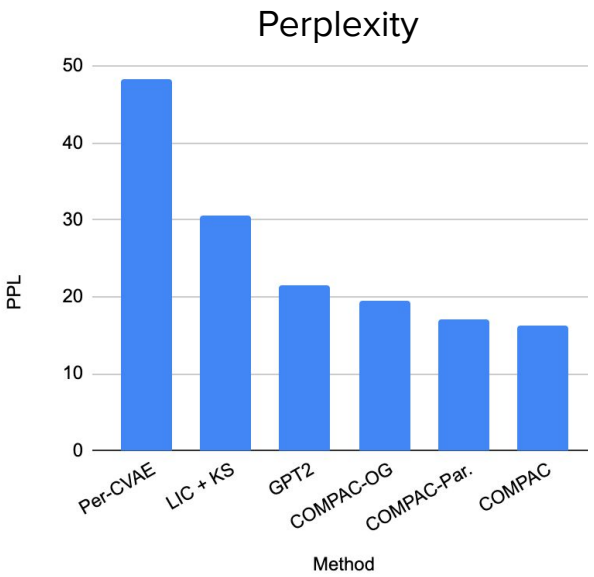

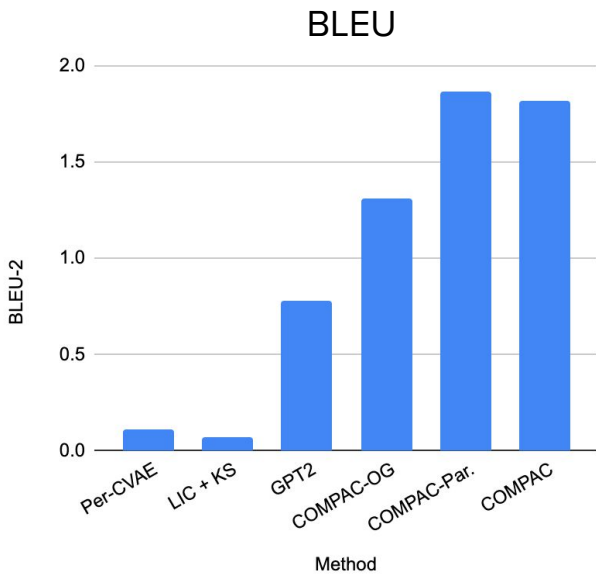

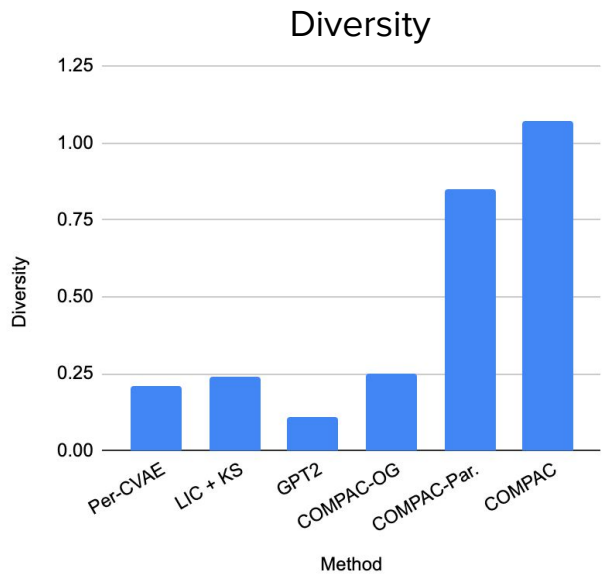

COMPAC is highly **fluent**, **faithful** to the ground truth as well as **diverse** as needed

Results: Human Evaluations

Persona:
I enjoy listening to classical music.
I'm a Hindu.
My favorite color is red.

User: Hi, recently I have got interests in religion.

GPT2 (2019)**:** Hi! How are you?

COMPAC-original: I'm a Hindu.

COMPAC-revised: Hi! I am a Hindu too.

COMPAC-paraphrase: That's great. I am religious.

COMPAC: That's great. I go to temple regularly and learn about Hinduism.

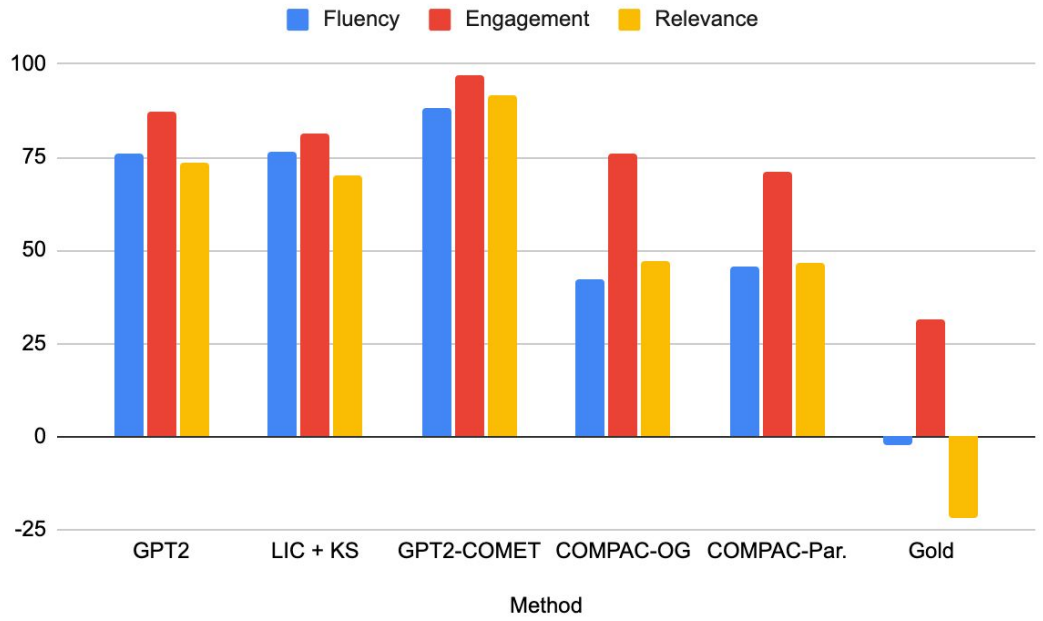

COMPAC is **fluent**, **engaging** and **relevant** when compared to all other baselines

Results: Error Analysis

Verdict	Persona	Dialog History	Persona Choice (z)	Response
Uses Original Persona	1. I'm the youngest of five. 2. I work at the hospital as a nurse. 3. I've pink hair.	What do you do for work?	I work at the hospital as a nurse.	I am a nurse.
Uses Expanded Persona	1. I just want to have fun with my friends. 2. I don't drink or do drugs or anything. 3. I am 19 and I cannot wait to move out of my parents home.	Are you enjoying life?	I just want to have fun with my friends → wants to have a party	Not really. I want to have a party.
Uses Improper Persona	1. I make a million dollars a year. 2. I'm married and have three kids. 3. I'm a baseball player.	I find it hard to support my family working at a bar. What about you?	Null persona ($\varnothing$)	I enjoy my life.

Table 8: Examples showing correct and incorrect persona choices in various dialog contexts by COMPAC model. It shows that COMPAC is capable of choosing a correct persona sentence (original or expanded) but sometimes the prior network fails to sample an appropriate one (third case).

Results: Interpretability and Controllabilty

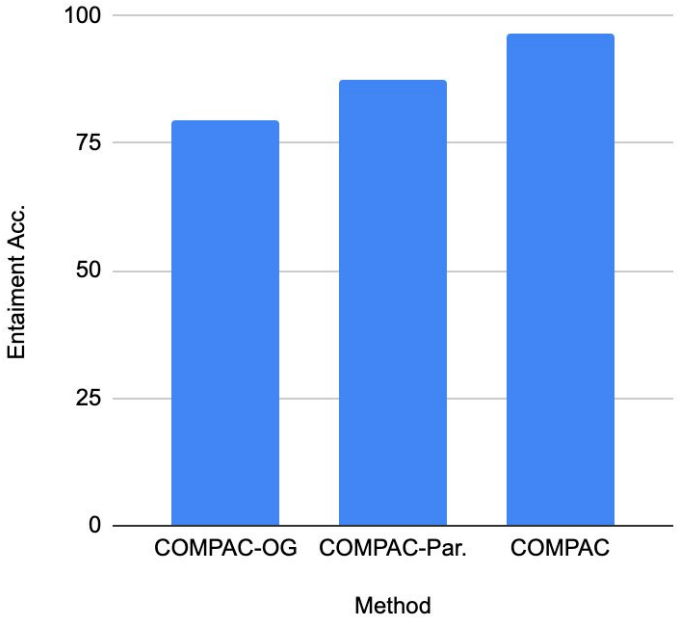

Ground truth obtained from DNLI dataset

Results: Interpretability and Controllabilty

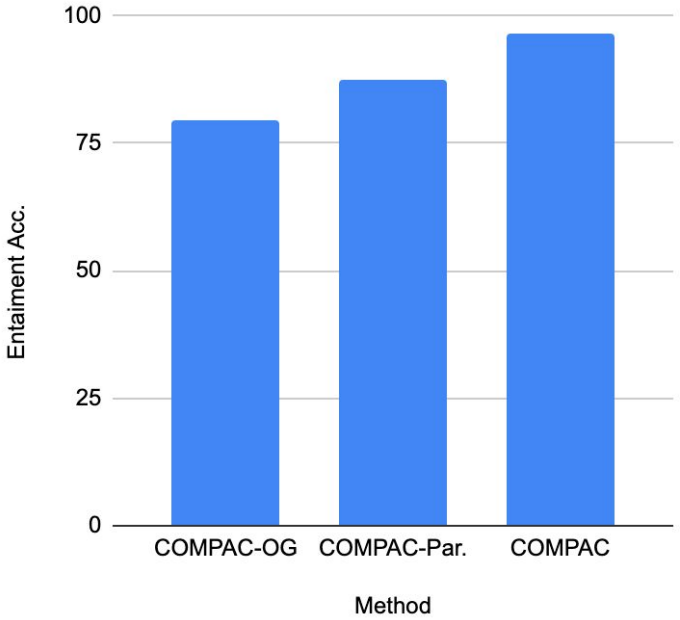

Ground truth obtained from DNLI dataset

Performance	Example
Presence of changed entity **86%**	**Changing the key entity** Before: My favorite color is red After: My favorite color is **green** Conversation: User: What is your favorite color? Bot: My favorite color is **green**
BERT score with unedited persona: 46.2 edited persona: **74.6**	**Swapping with another expansion** Before: I want to swim in the ocean After: I want to buy a beach **umbrella** Conversation: User: What do you do at beaches? Bot: I will buy an **umbrella** at the beach

Summary

- We showed that expanding persona with commonsense helps a dialog model to generate high-quality and diverse responses
- Fine-grained persona grounding is crucial for interpretability as well as effective grounding
- COMPAC is still limited by the COMET or paraphrase expansions
- As a future work, we could expand the framework to use more than one persona sentences (or expansions), as needed.

Thanks!