



Carnegie
Mellon
University

Unsupervised Enrichment of Persona-grounded Dialog with Background Stories

Bodhisattwa Prasad Majumder, Taylor Berg-kirkpatrick, Julian McAuley, Harsh Jhamtani
UC San Diego, Carnegie Mellon University
ACL 2021



Persona-grounded Dialog

Persona


I have two children and a dog
I like outdoor activities with my kids

Persona-grounded Dialog


Persona

I have two children and a dog
I like outdoor activities with my kids

I went camping last
weekend with my family



Oh great! How was
your experience?




Persona-grounded Dialog


Persona

I have two children and a dog
I like outdoor activities with my kids


I went camping last weekend with my family



Oh great! How was your experience?



It was eventful. Do you like camping?




Persona-grounded Dialog


Persona

I have two children and a dog
I like outdoor activities with my kids


I went camping last weekend with my family



Oh great! How was your experience?



It was scary. A howl pierced the night. The children huddled closer to the campfire. We were terrified.




Persona-grounded Dialog

Persona

I have two children and a dog
I like outdoor activities with my kids


I went camping last weekend with my family




Background Story

A howl pierced the black night. The kids huddled closer to the campfire. Everyone was terrified, even the camp counselor. The howl came again, right on top of them!

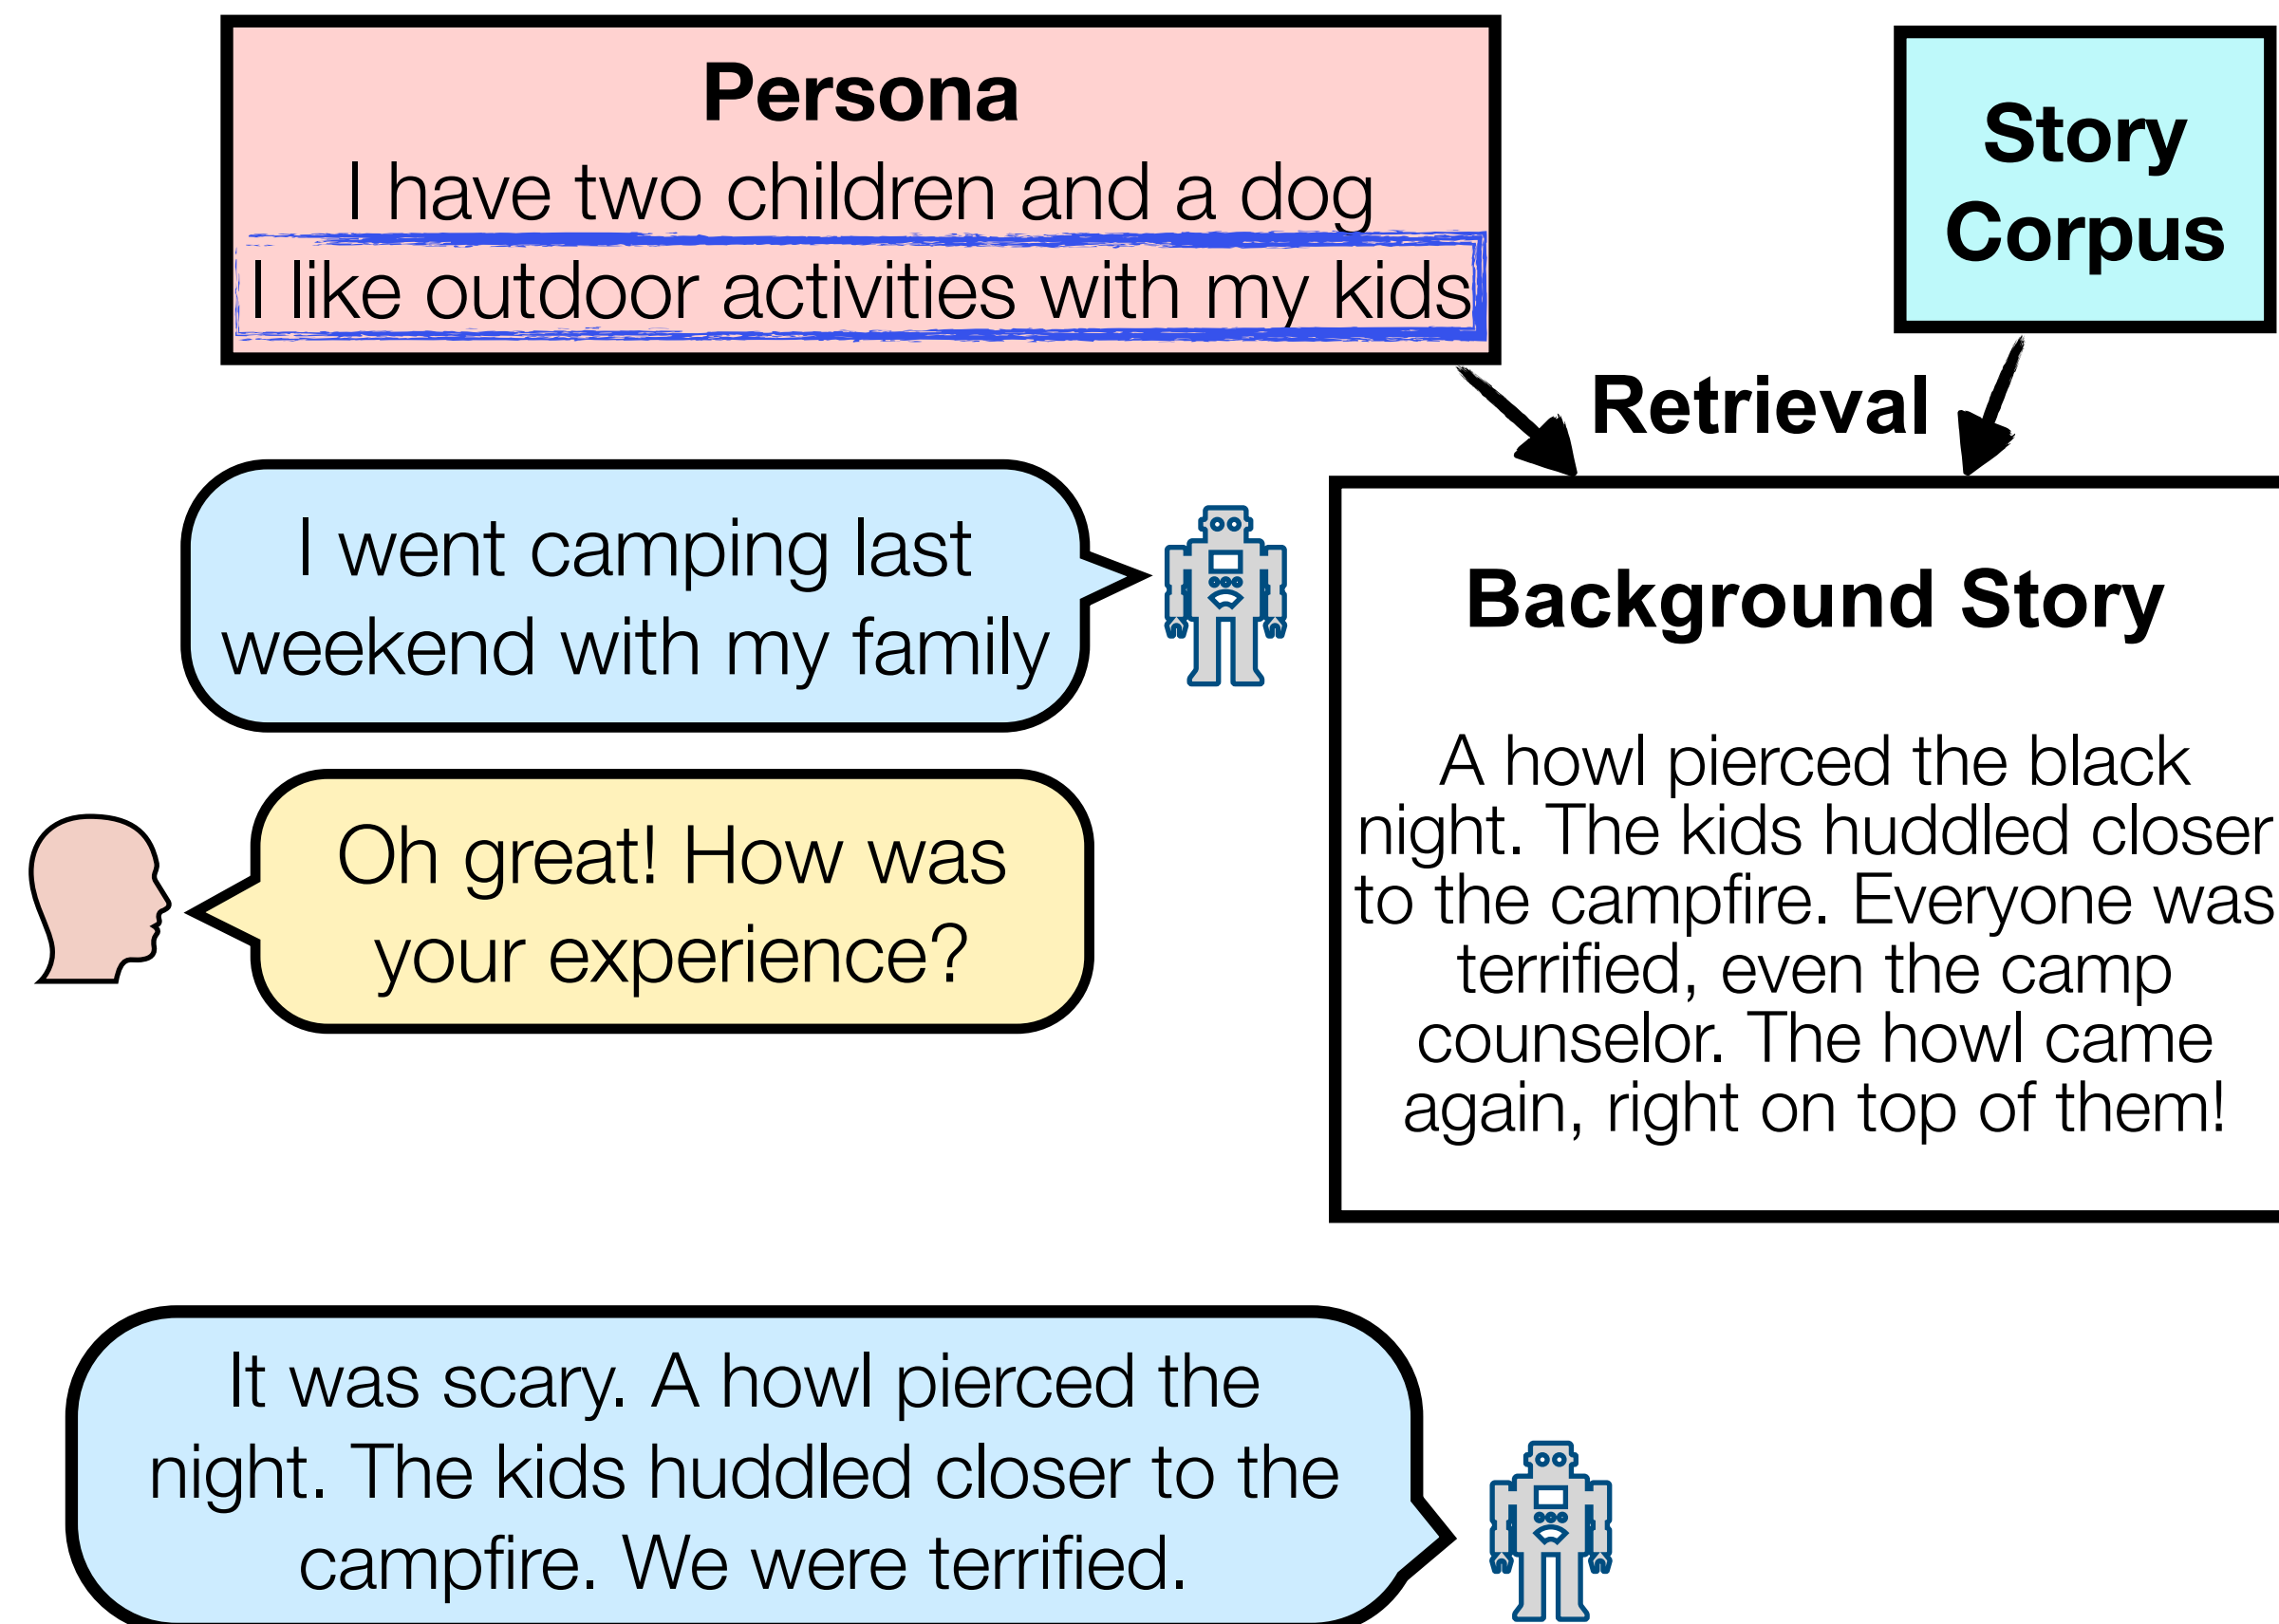
Oh great! How was your experience?



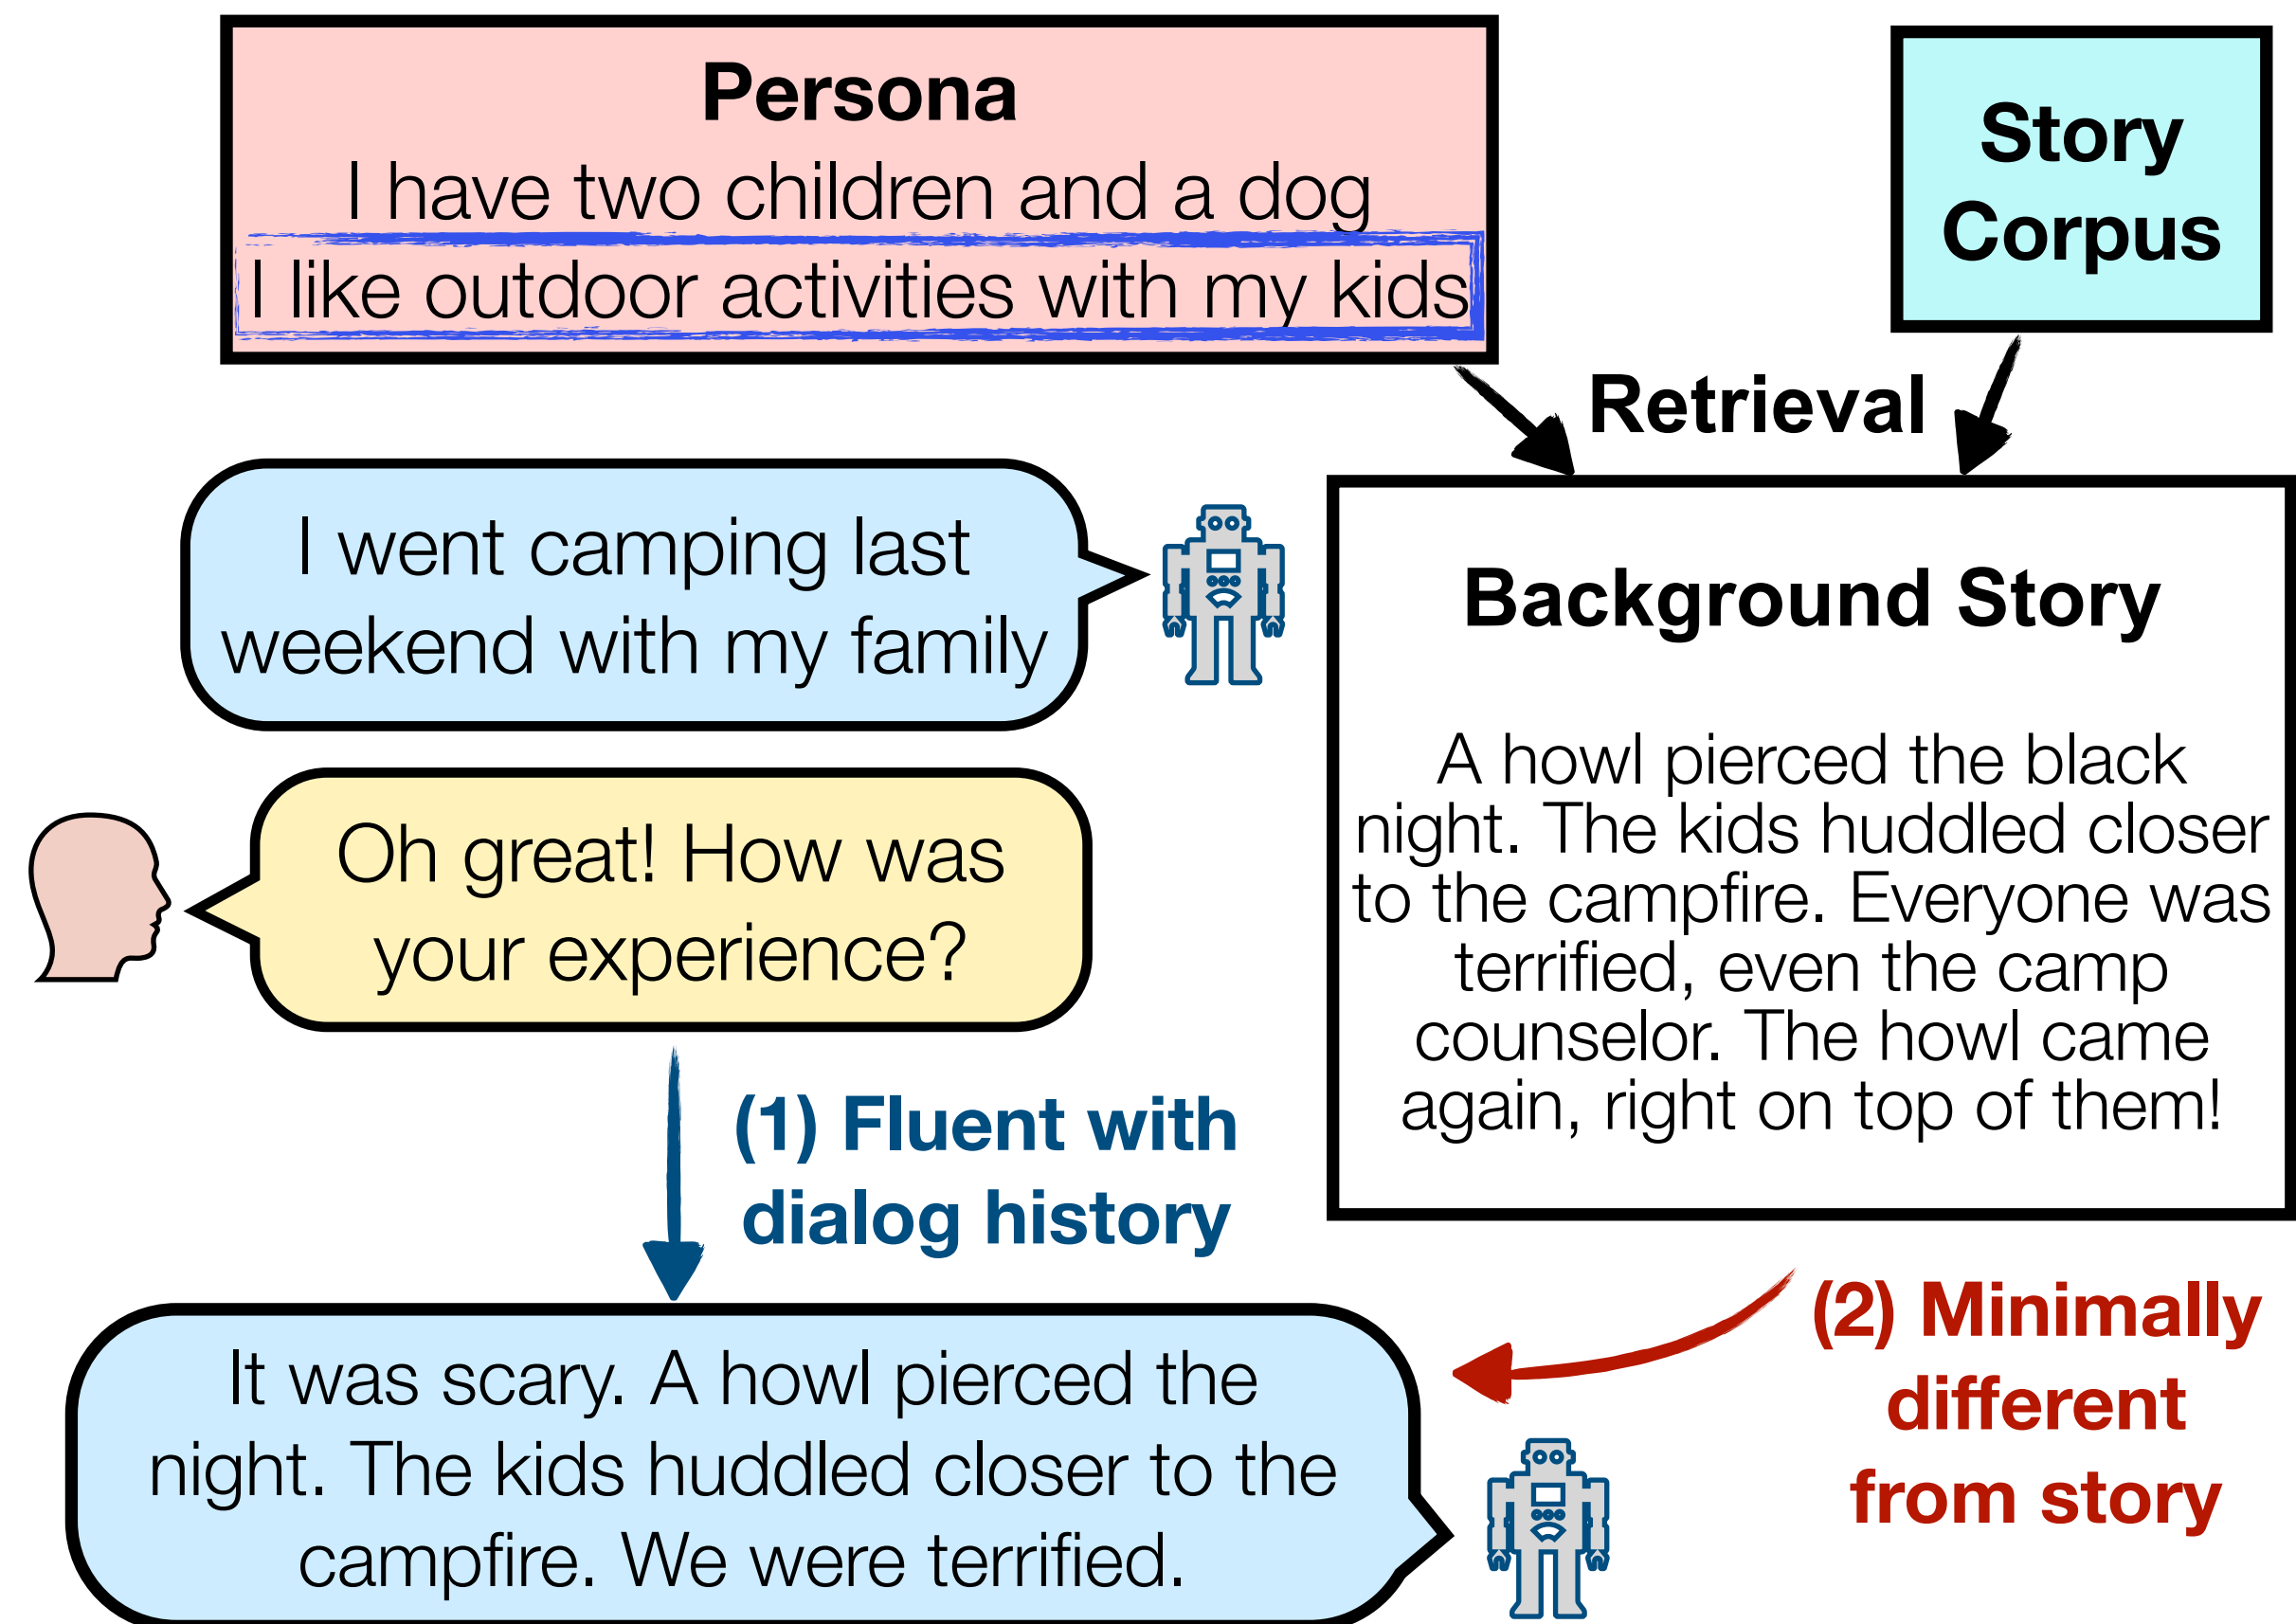
It was scary. A howl pierced the night. The kids huddled closer to the campfire. We were terrified.



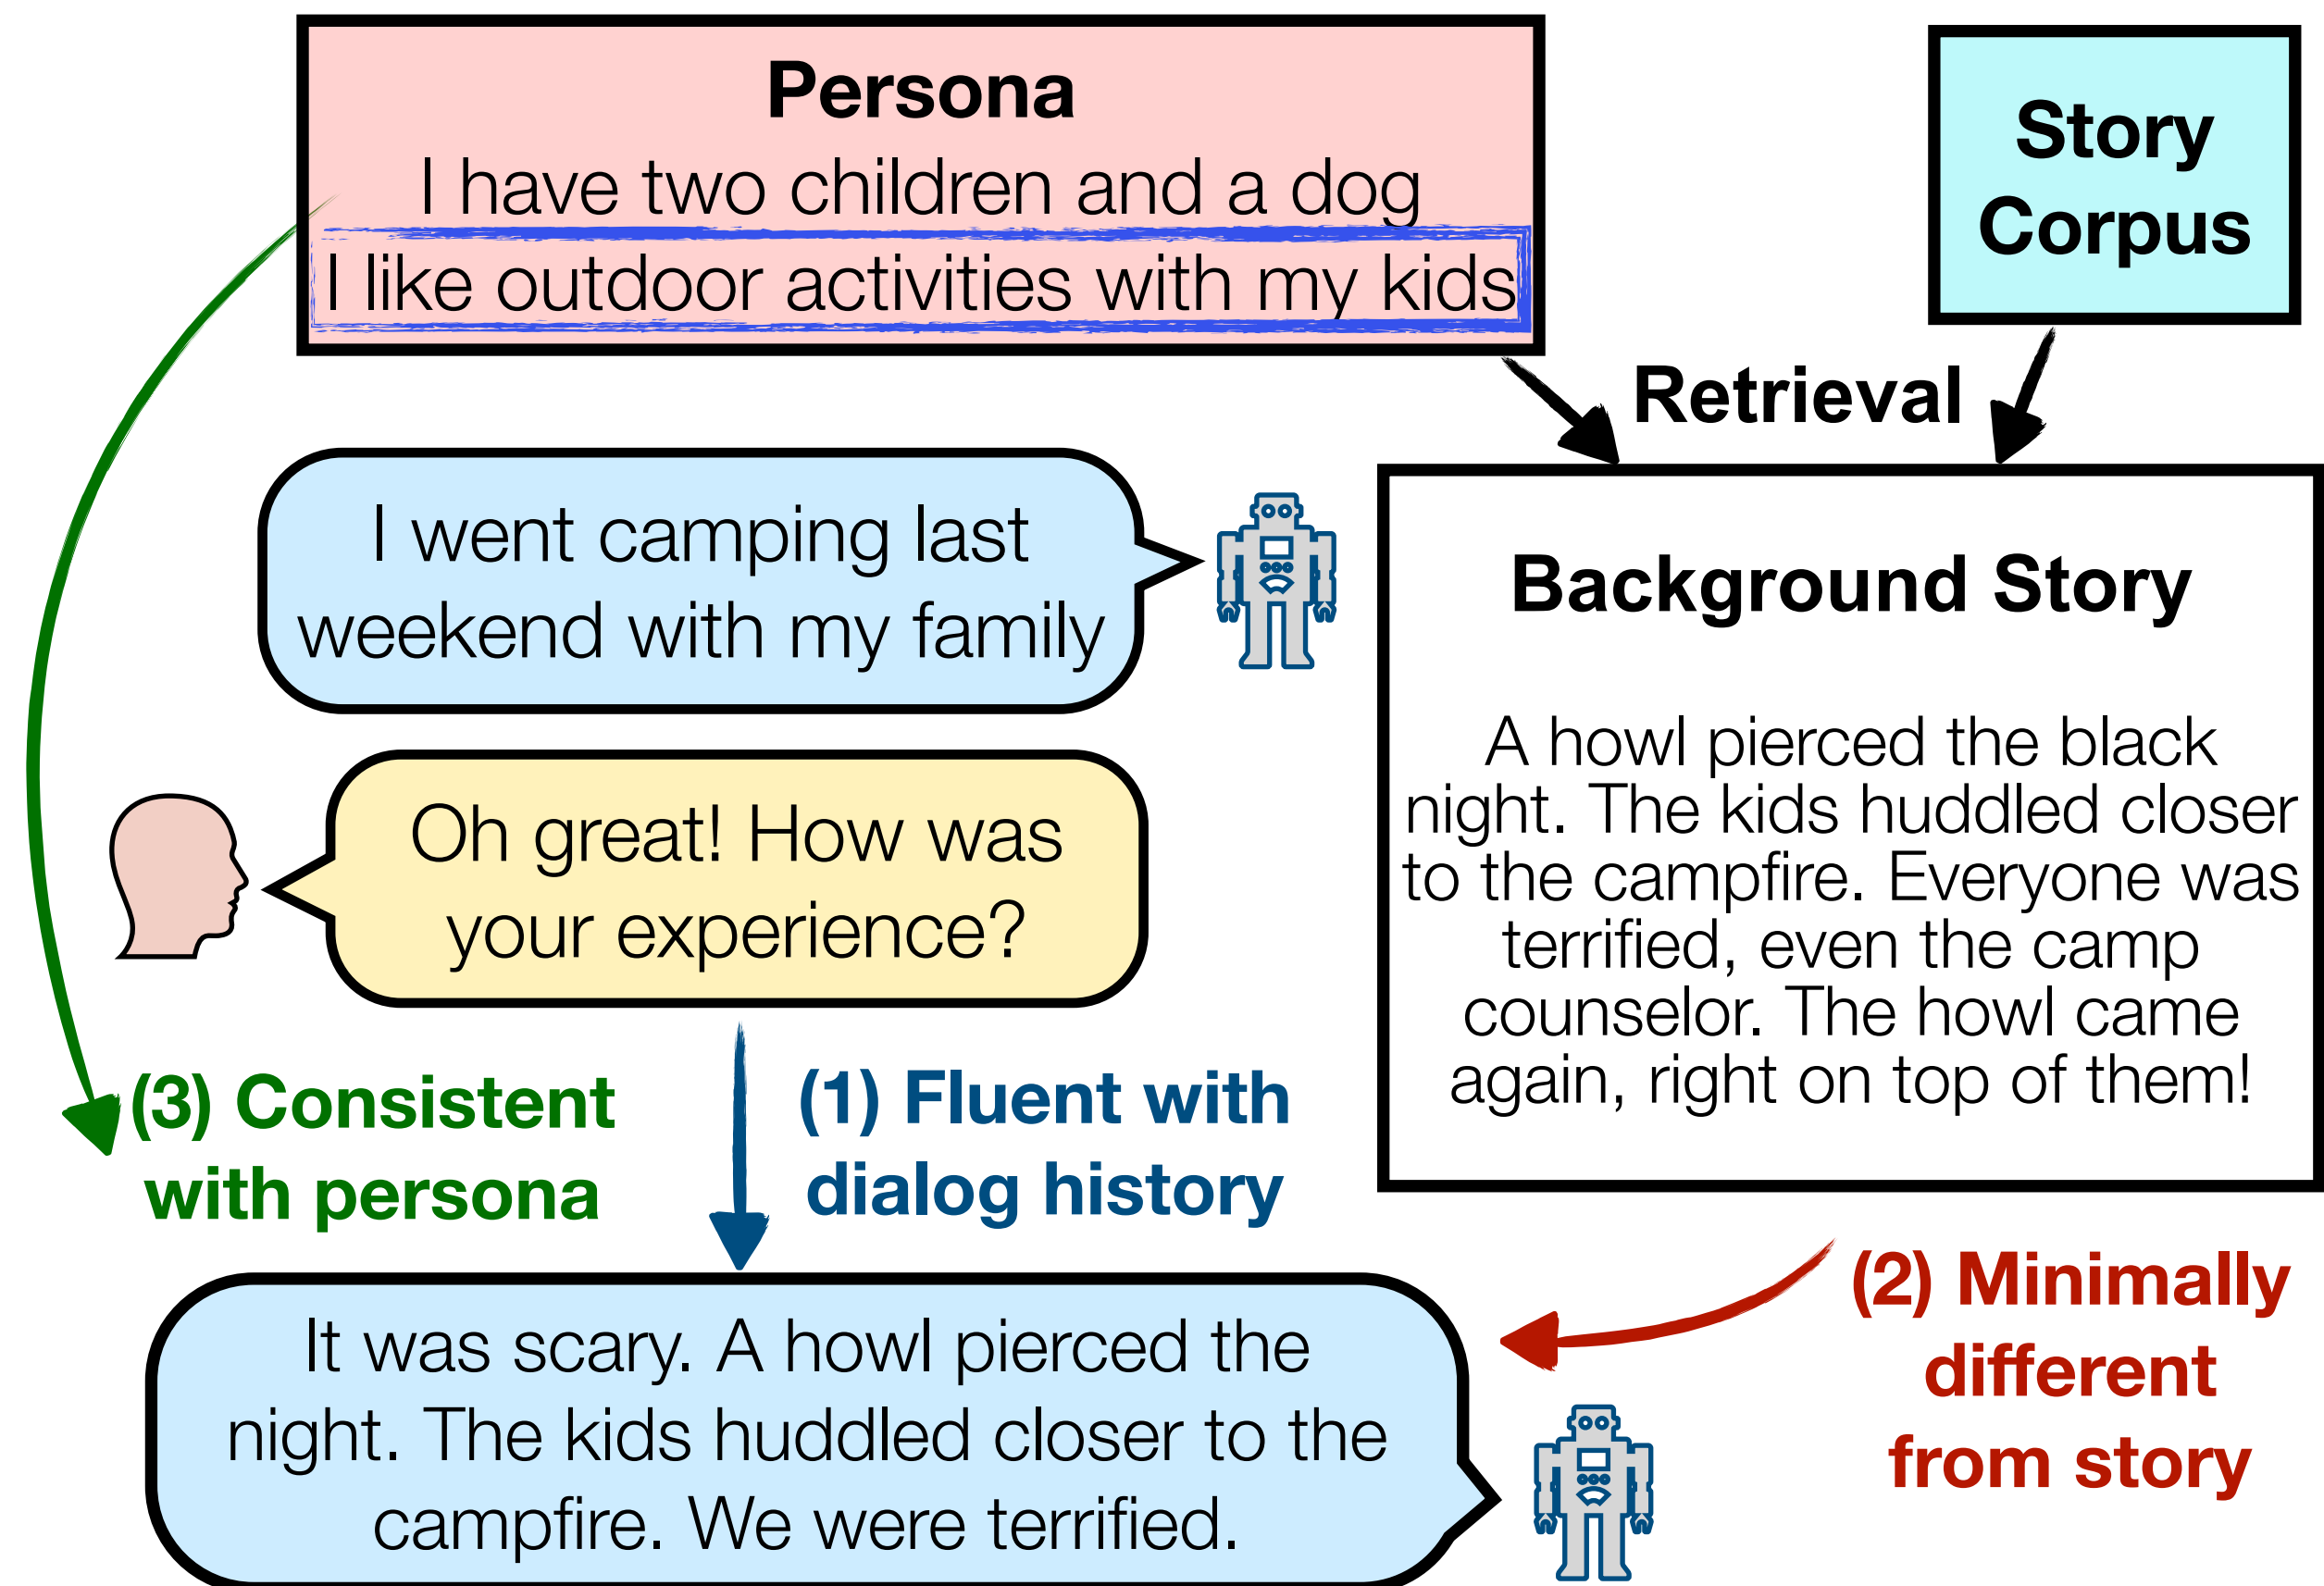
Persona-grounded Dialog



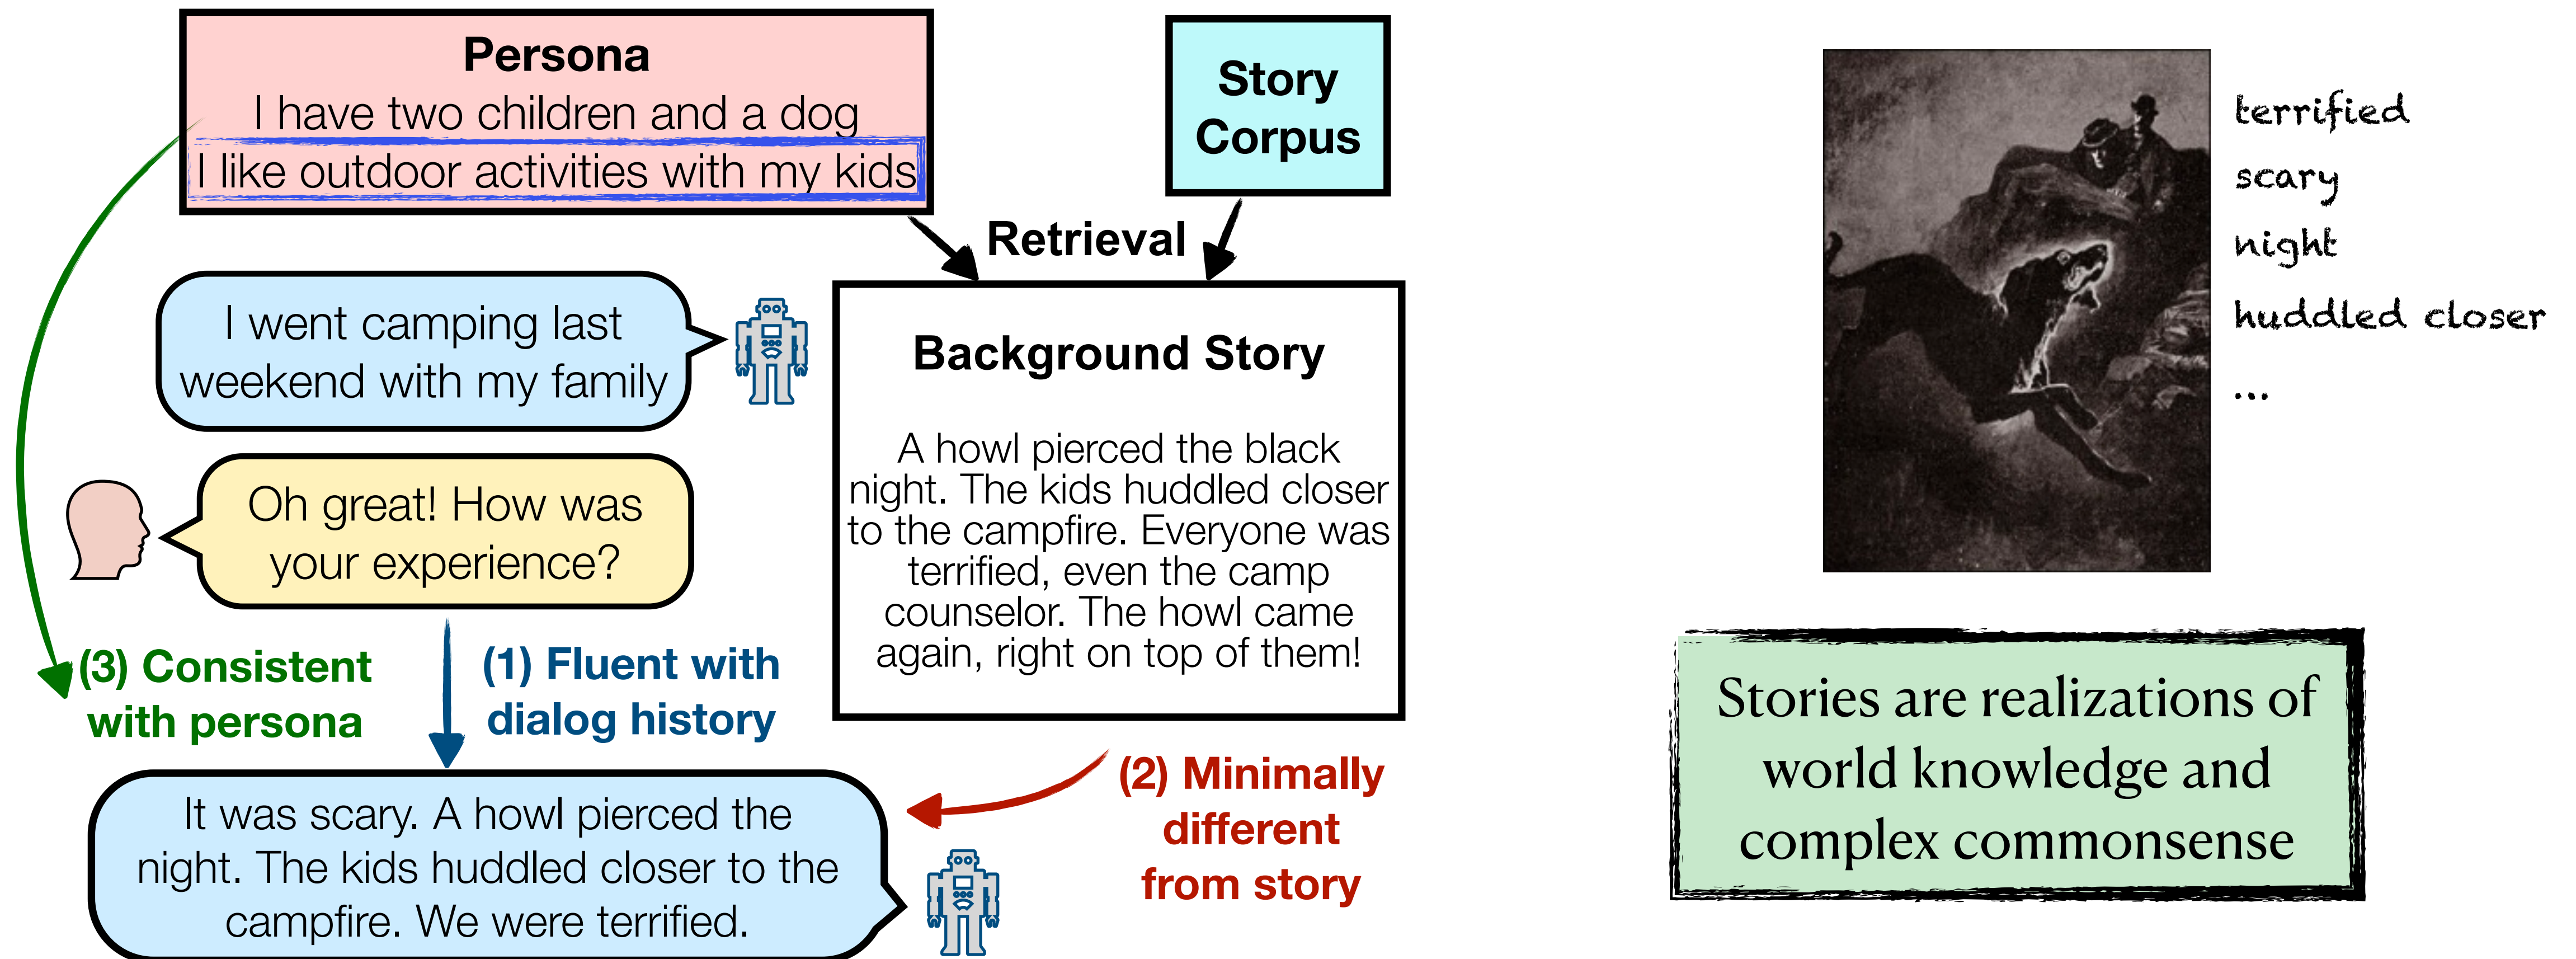
Persona-grounded Dialog



Persona-grounded Dialog



Persona-grounded Dialog



Our Goals

- Can we **retrieve relevant stories** based on persona facts?
- Can we synthesize a response that is **fluent** with dialog history as well as **sticks** to both the **retrieved story** and the **persona facts**

What is different from previous works

- Previous work (Su et al., 2020) attempts to generate dialog responses using retrieved non-conversational texts (posts from social forums). They use retrieved text as **pseudo-labels** and **train** a model from scratch.
- Can we bypass this noisy distant supervision and extra training burden? We aim to synthesize the response completely at the **inference time** using a **pre-existing** dialog model (i.e. without any **additional training**)

Unsupervised Persona Enrichment with Background Stories (PABST)



PABST



(Majumder et al., 2020)


Suggests the most relevant persona to be used for the next response based on dialog history

PABST




(Majumder et al., 2020)

Suggests the most relevant persona to be used for the next response based on dialog history



Story that has highest Bert-Score with sampled persona

Persona
+
Story



PABST



(Majumder et al., 2020)


Suggests the most relevant persona to be used for the next response based on dialog history

Story that has highest Bert-Score with sampled persona


- 1) Fluent with Dialog History
- 2) Minimally different from retrieved story
- 3) Consistent with Persona

Gradient-based decoding in PABST

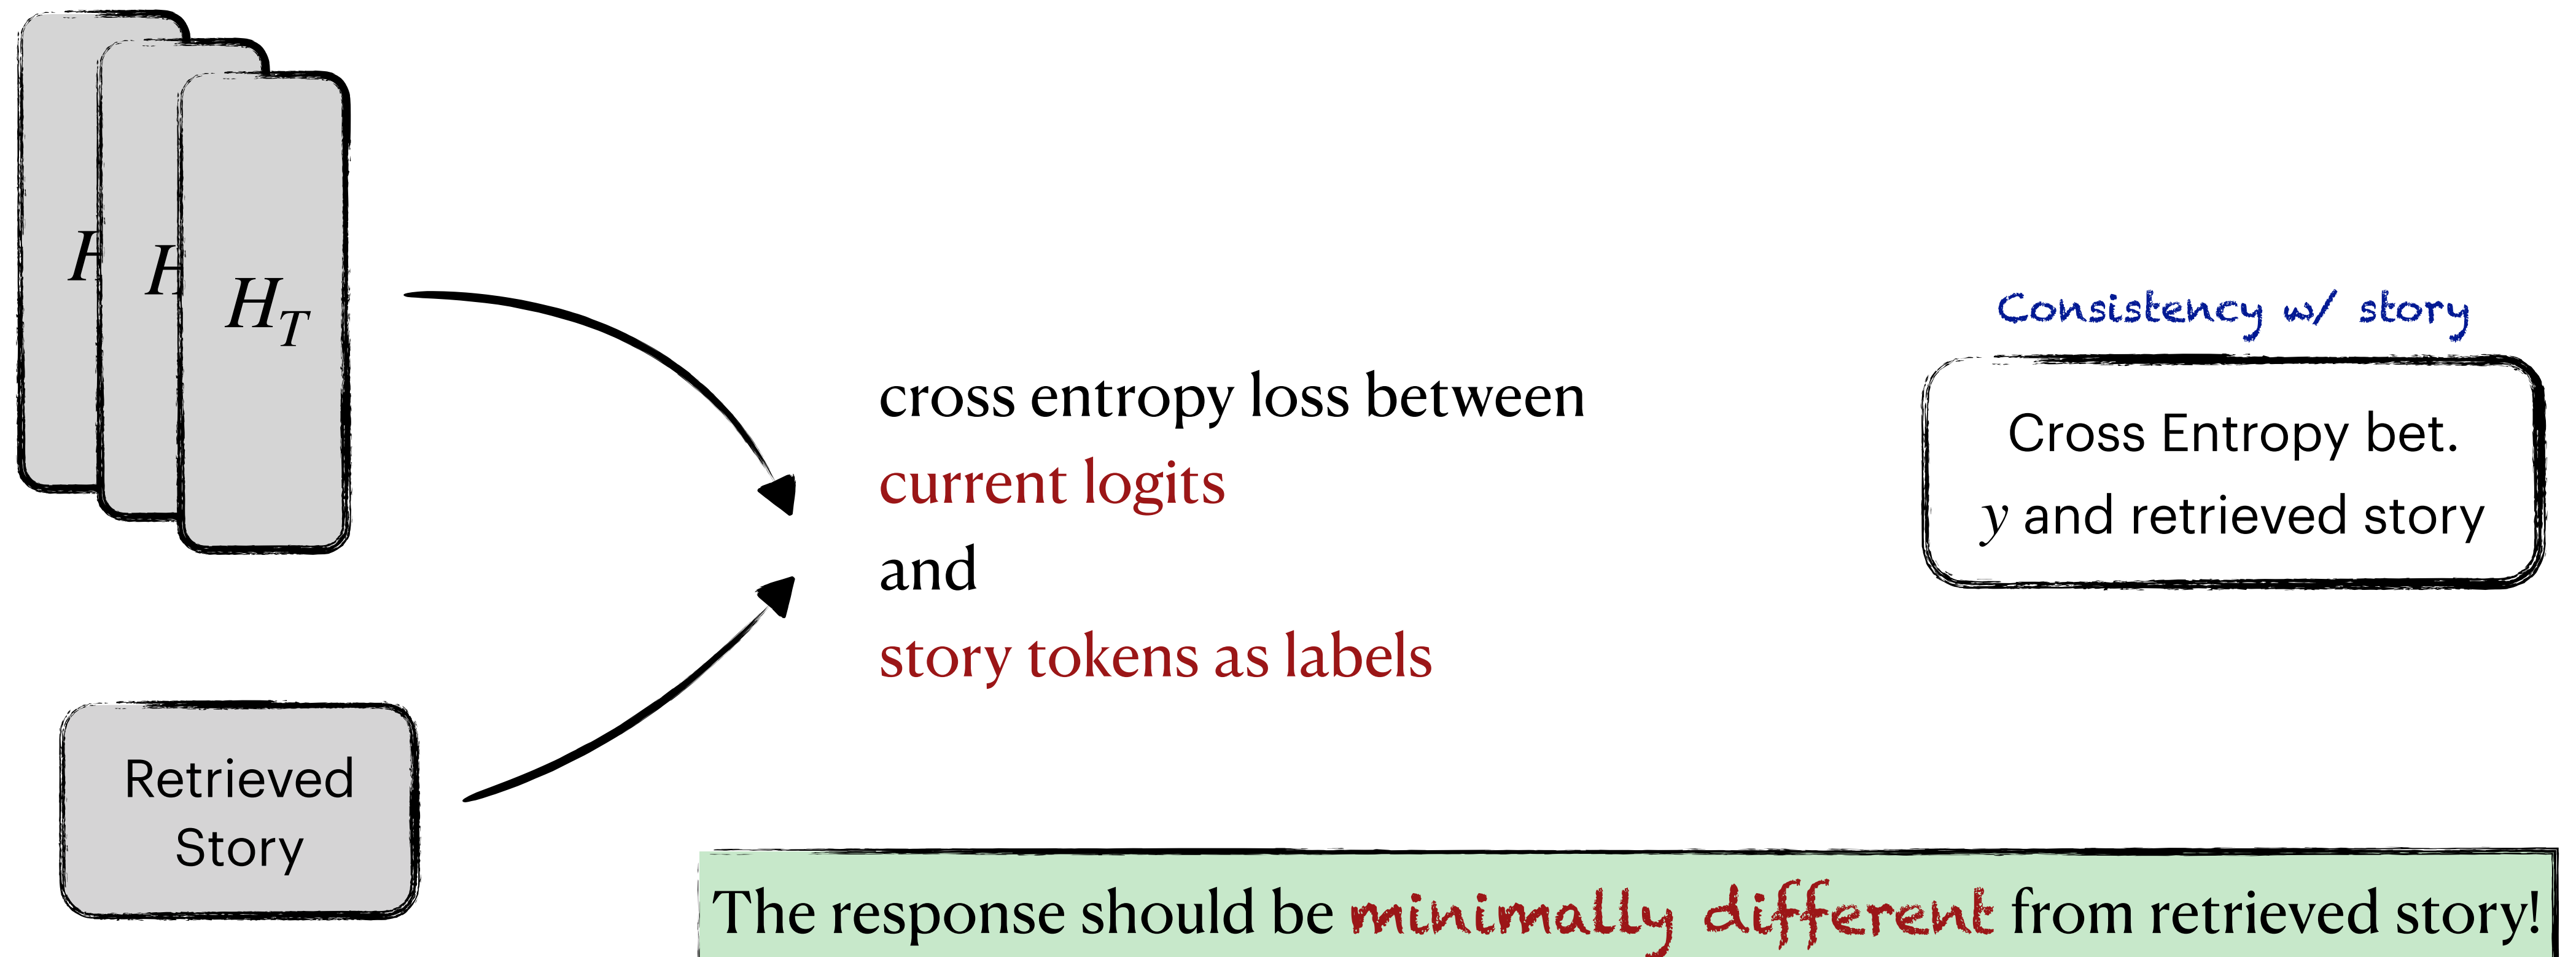
Gradient-based decoding in PABST




Gradient-based decoding in PABST



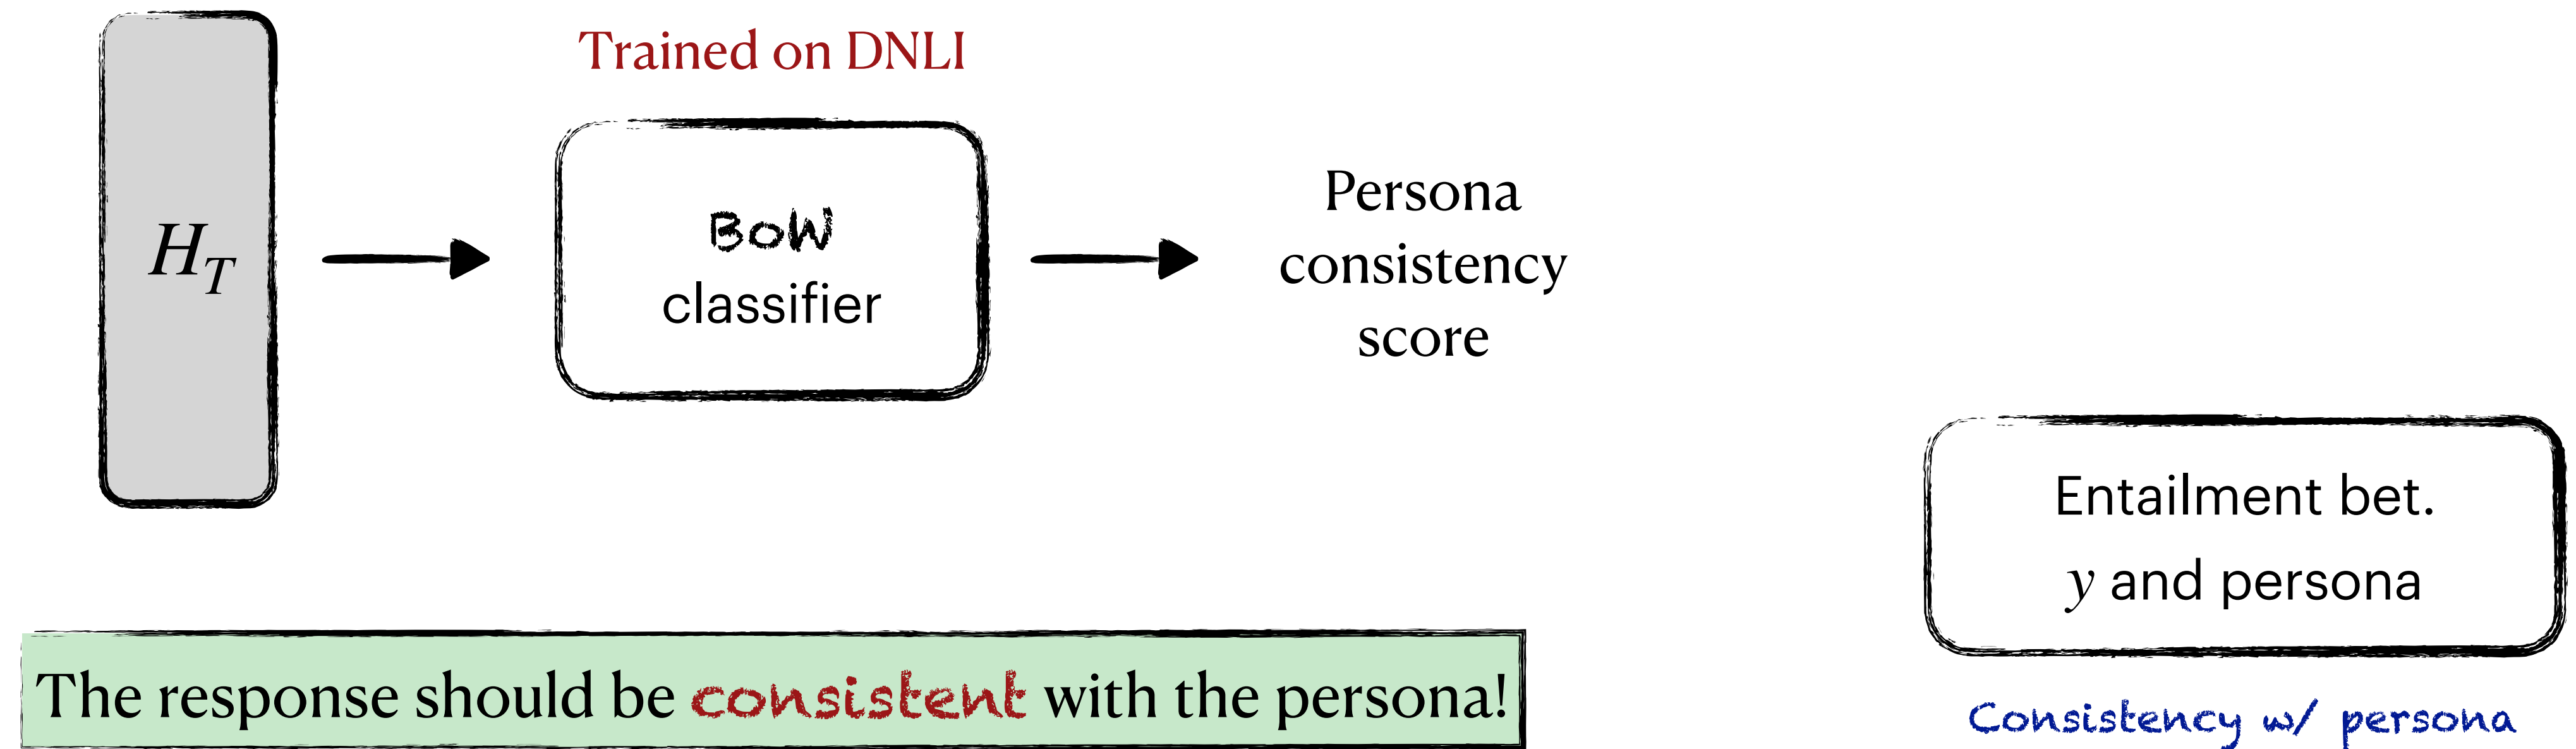
Gradient-based decoding in PABST



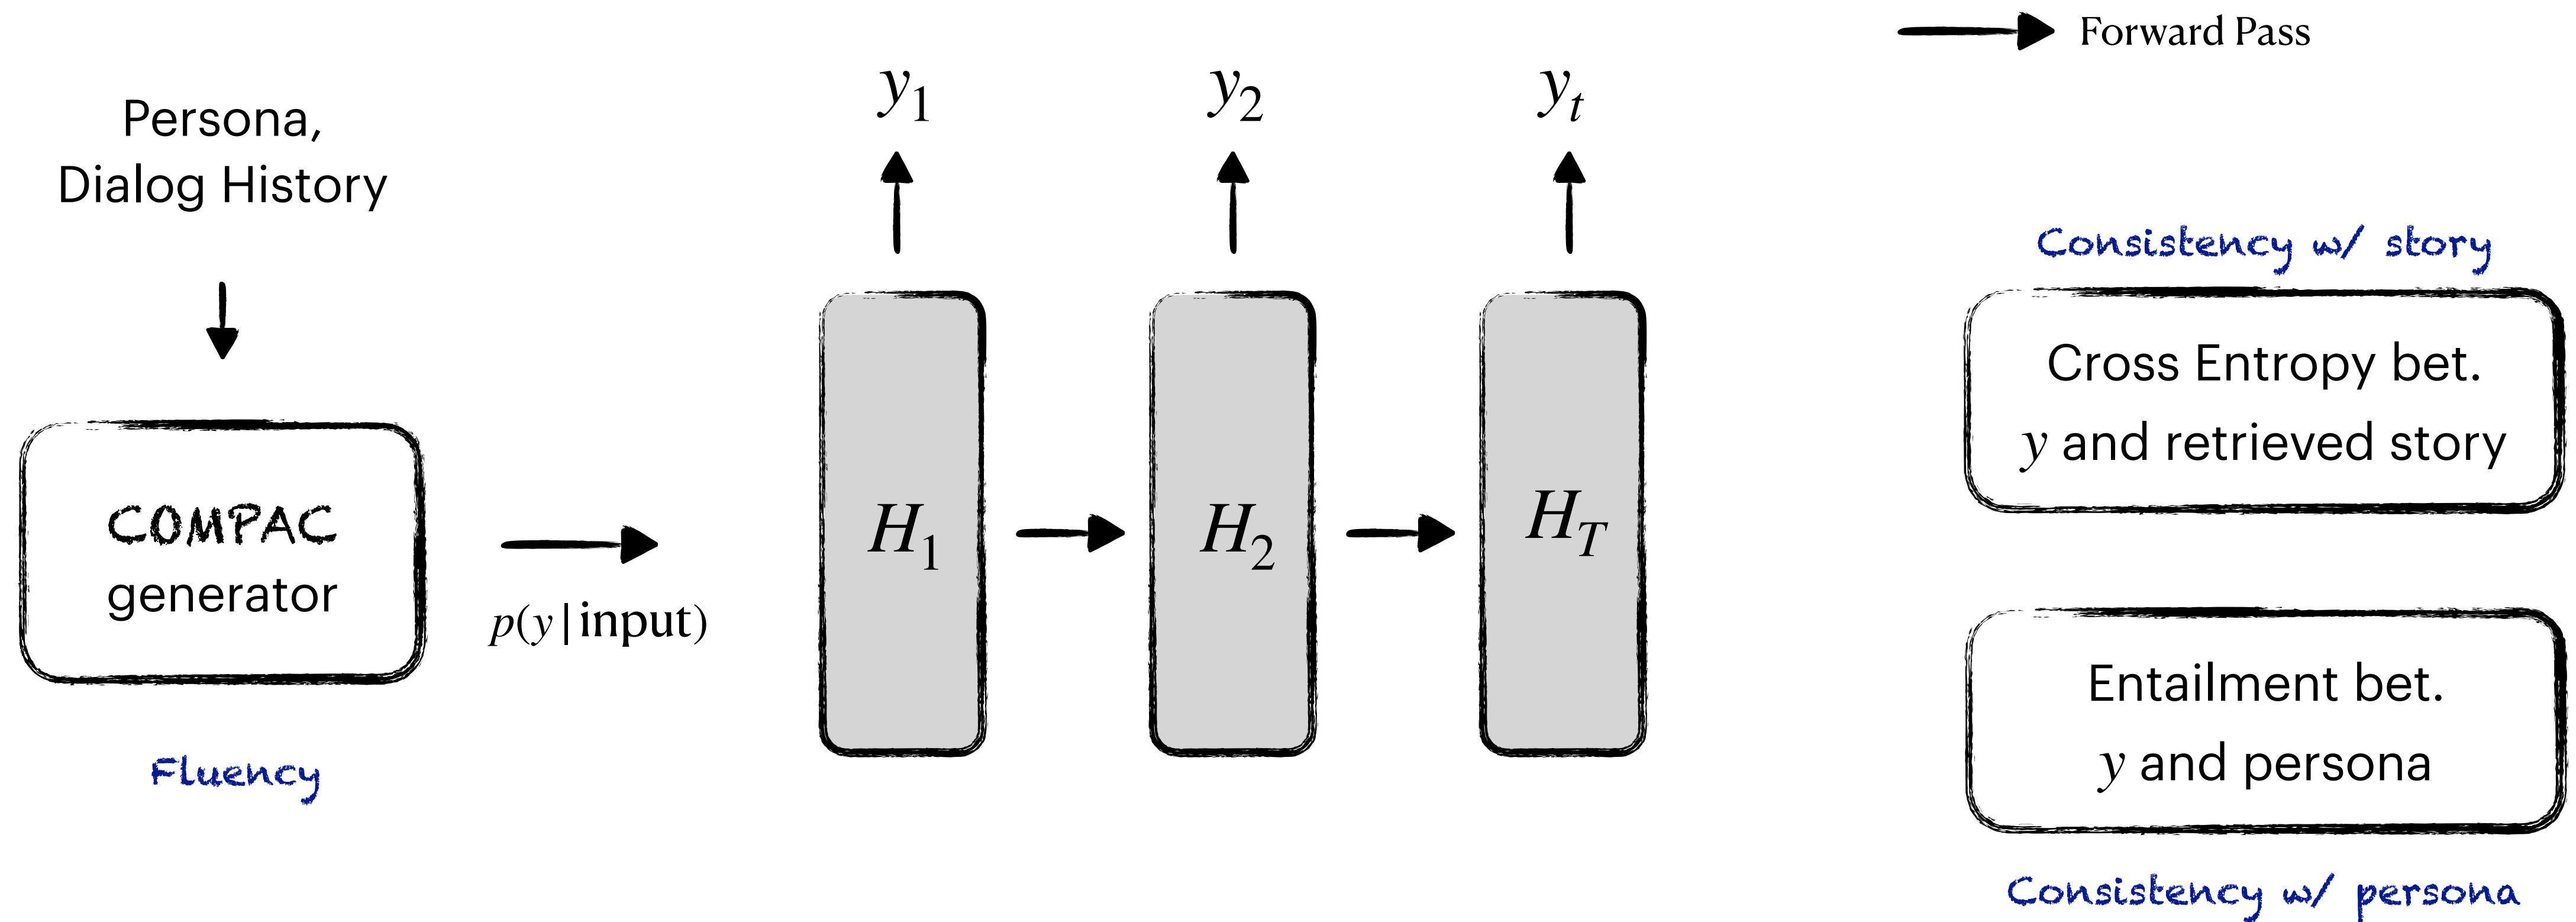
Gradient-based decoding in PABST




Gradient-based decoding in PABST



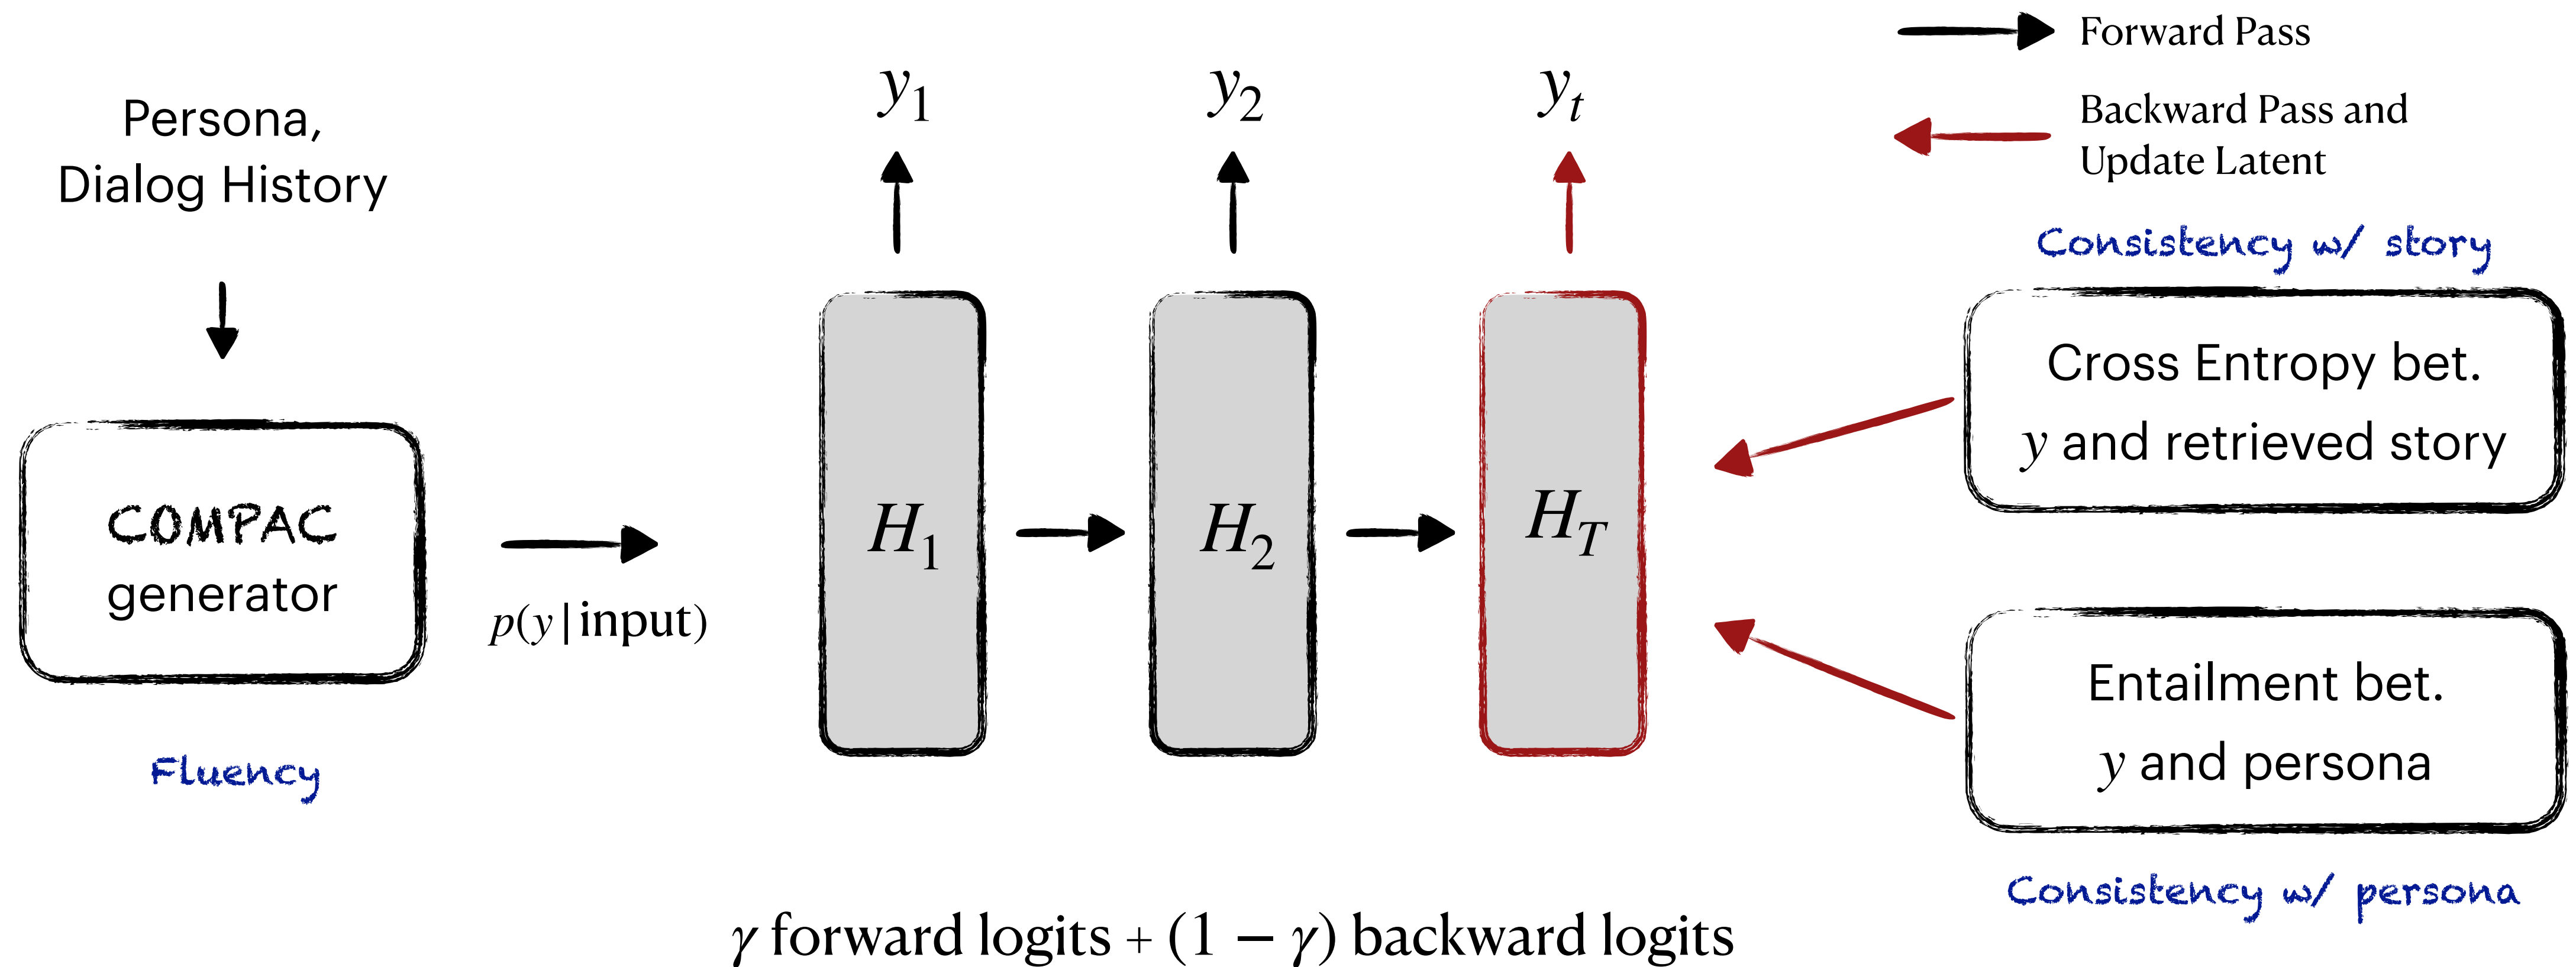
Gradient-based decoding in PABST



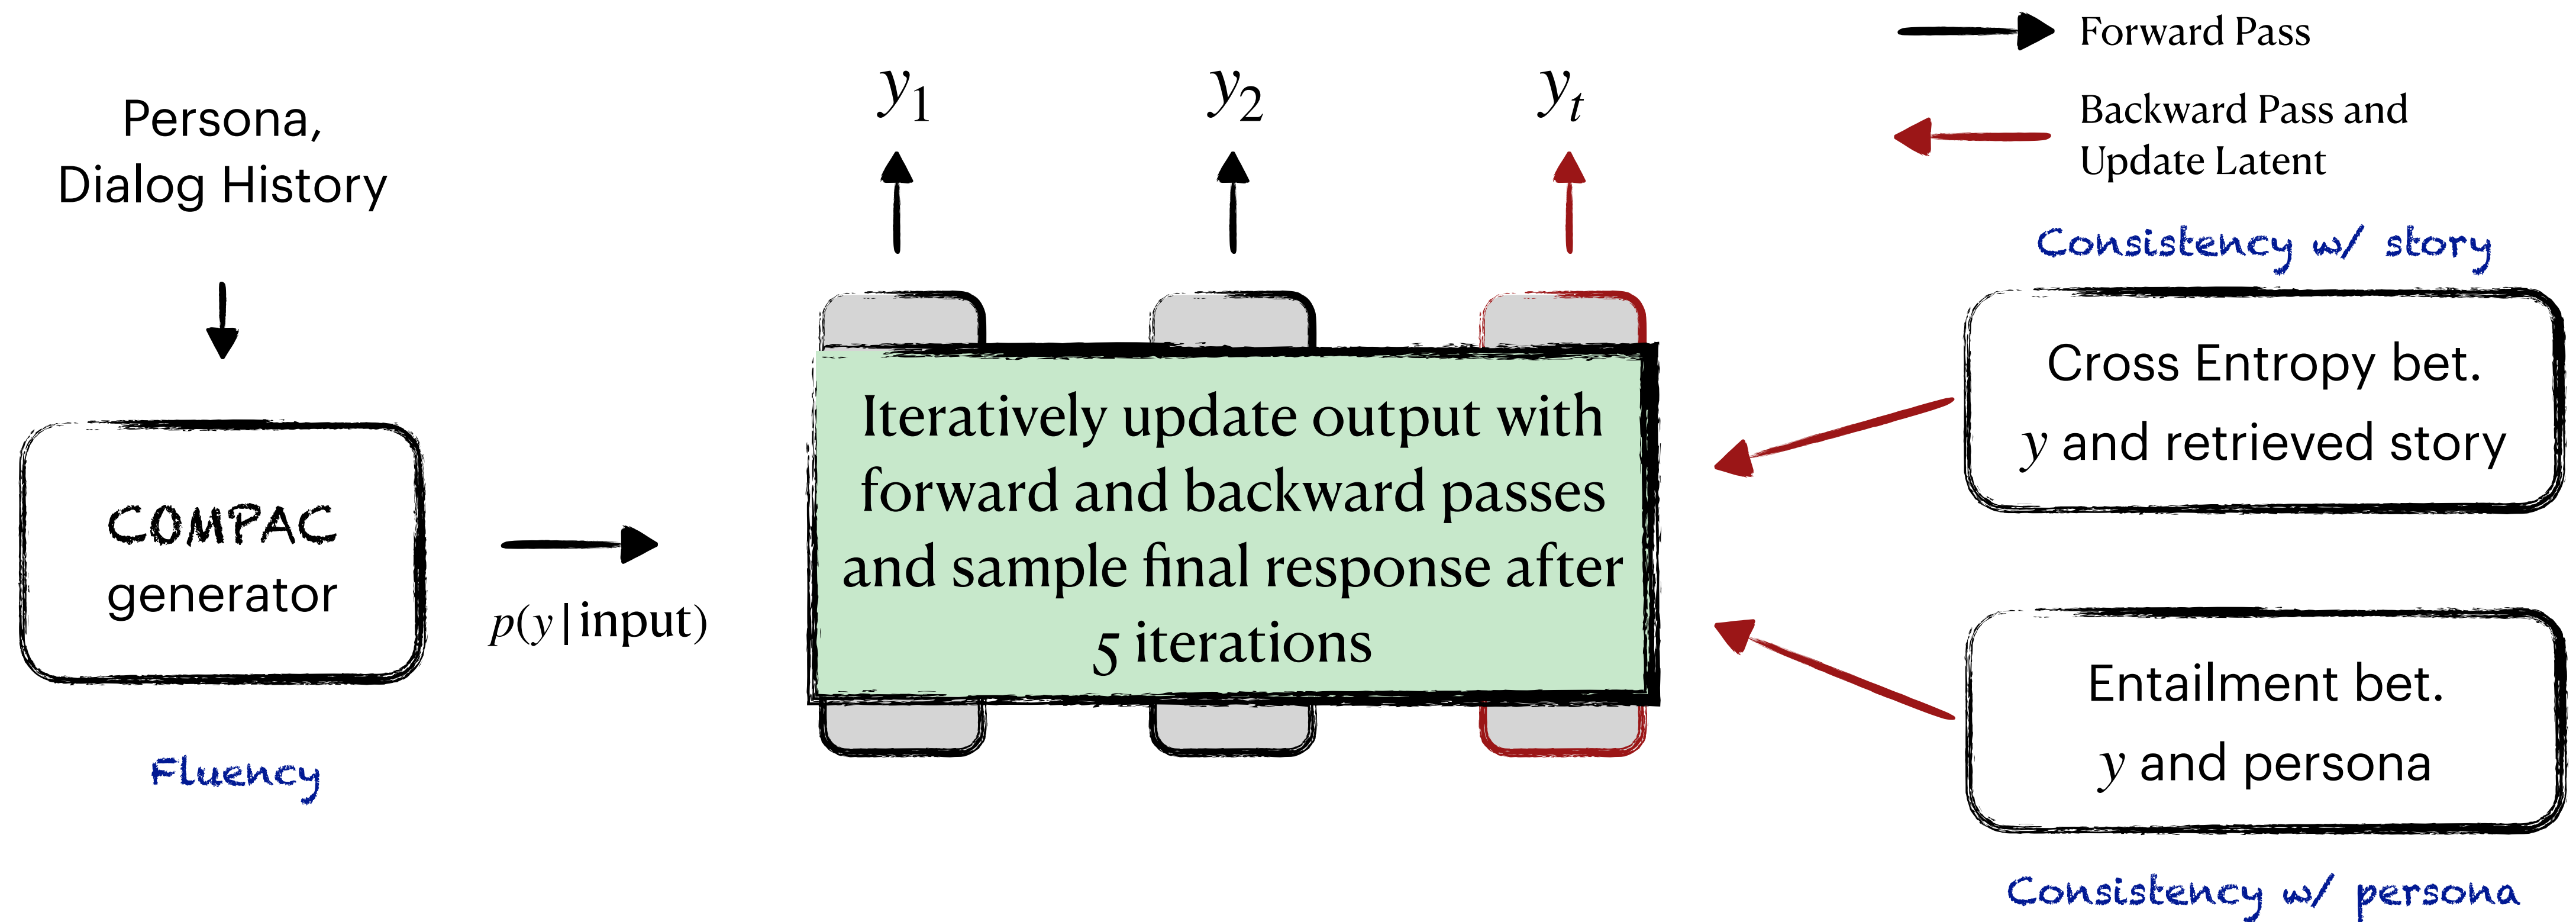
Gradient-based decoding in PABST



Gradient-based decoding in PABST



Gradient-based decoding in PABST



Inspired from PPLM (2020) and DELOREAN (2020)

Experiments

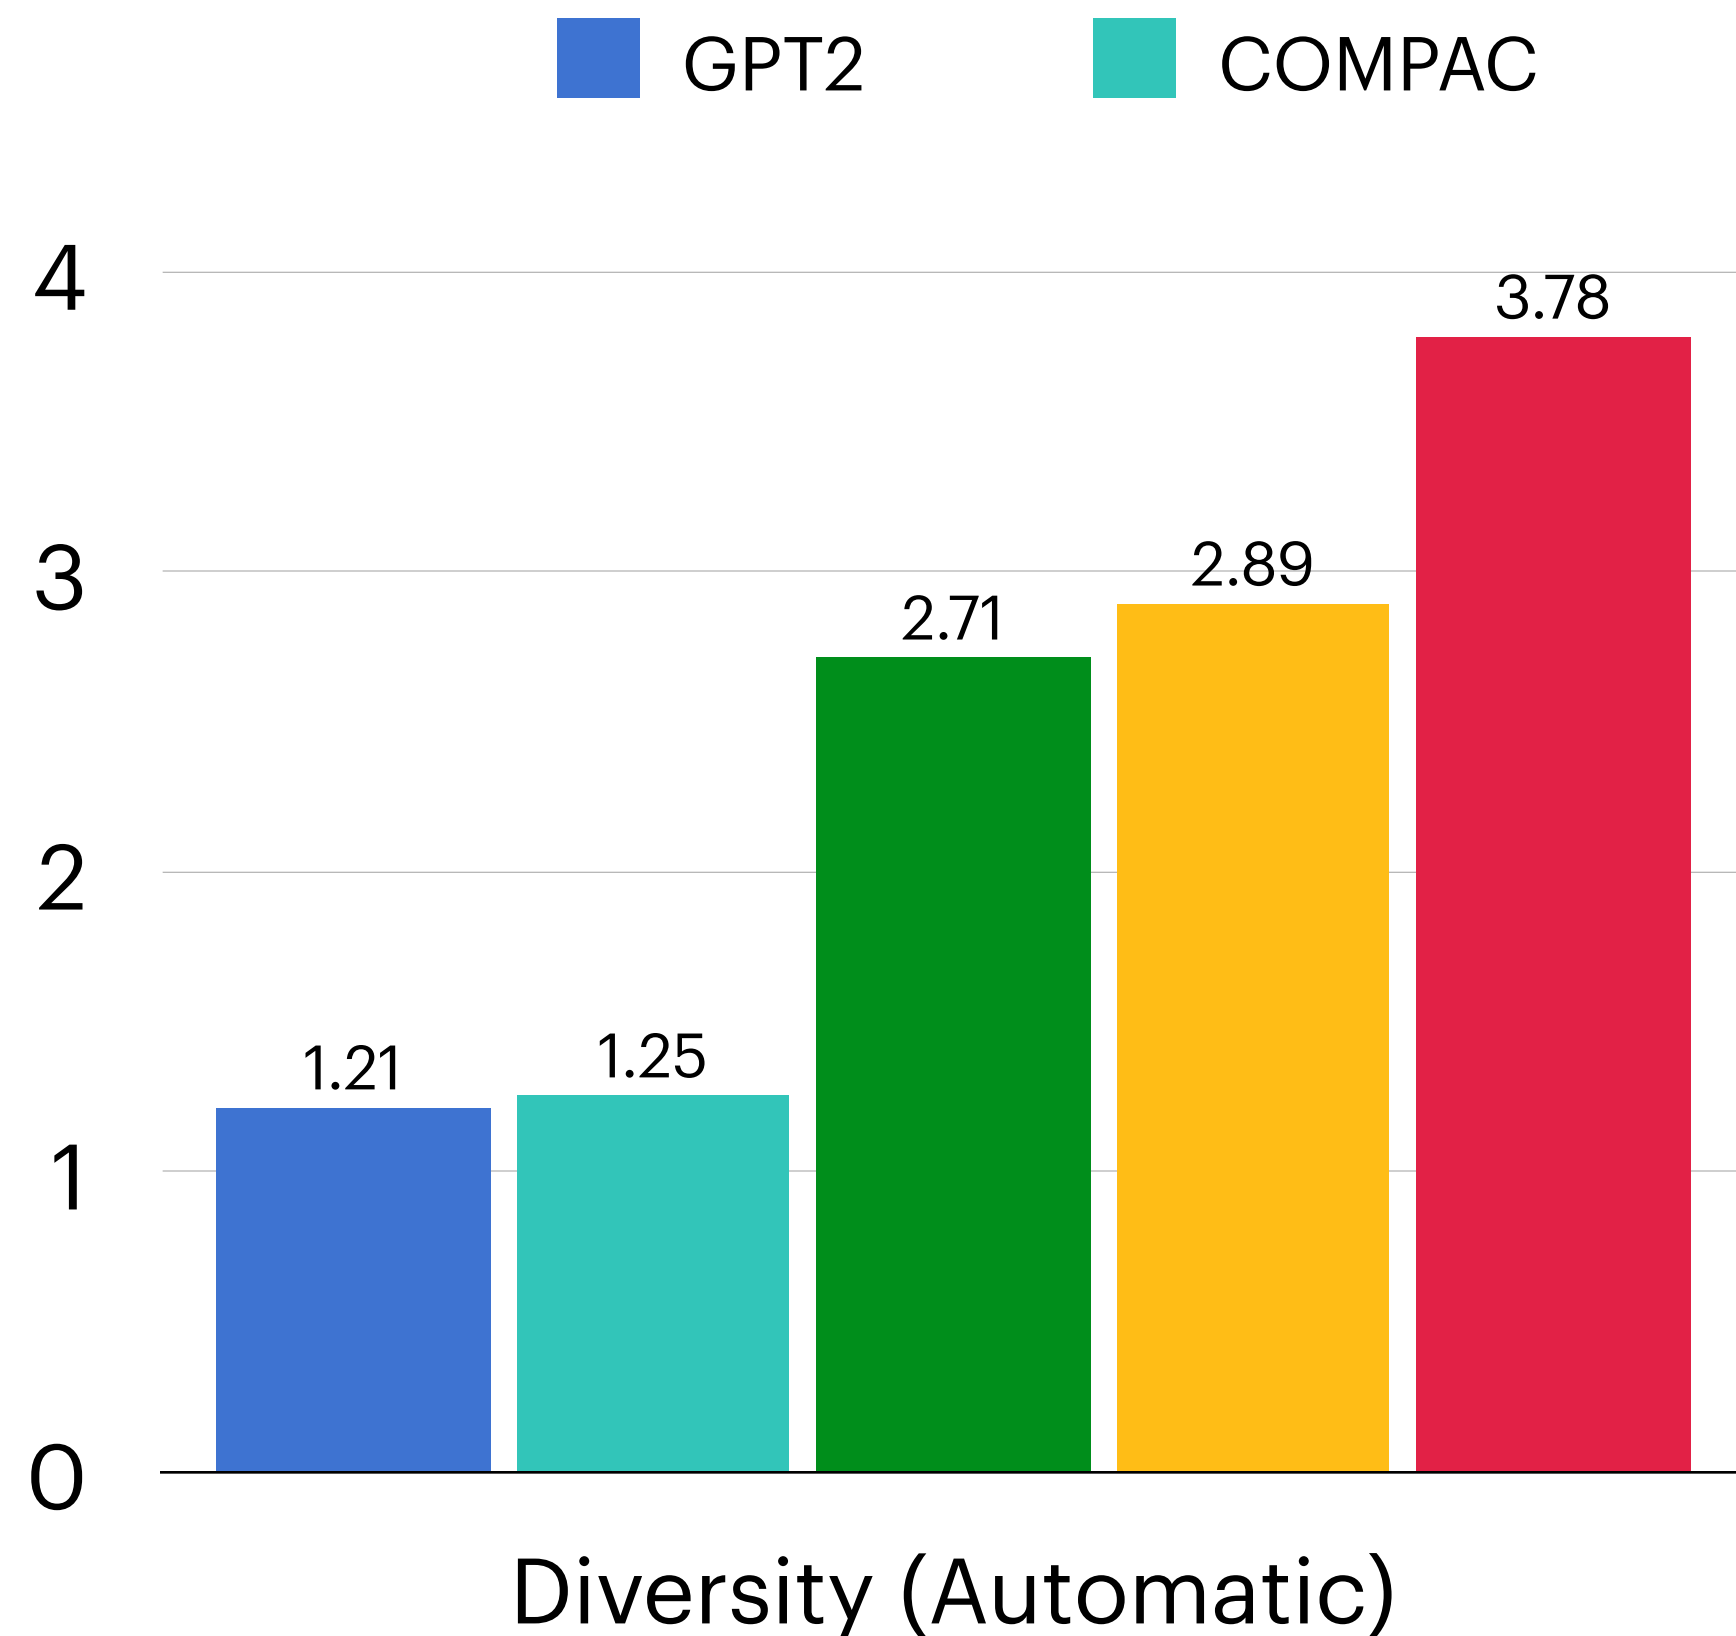
Research Questions

- Do decoding-level constraints help?
- Do pretrained LMs help generate sensible and engaging responses?

Baselines

- GPT2
- **COMPAC**
- Multi-task
- Pseudo-labeling
- **PABST**

Evaluation

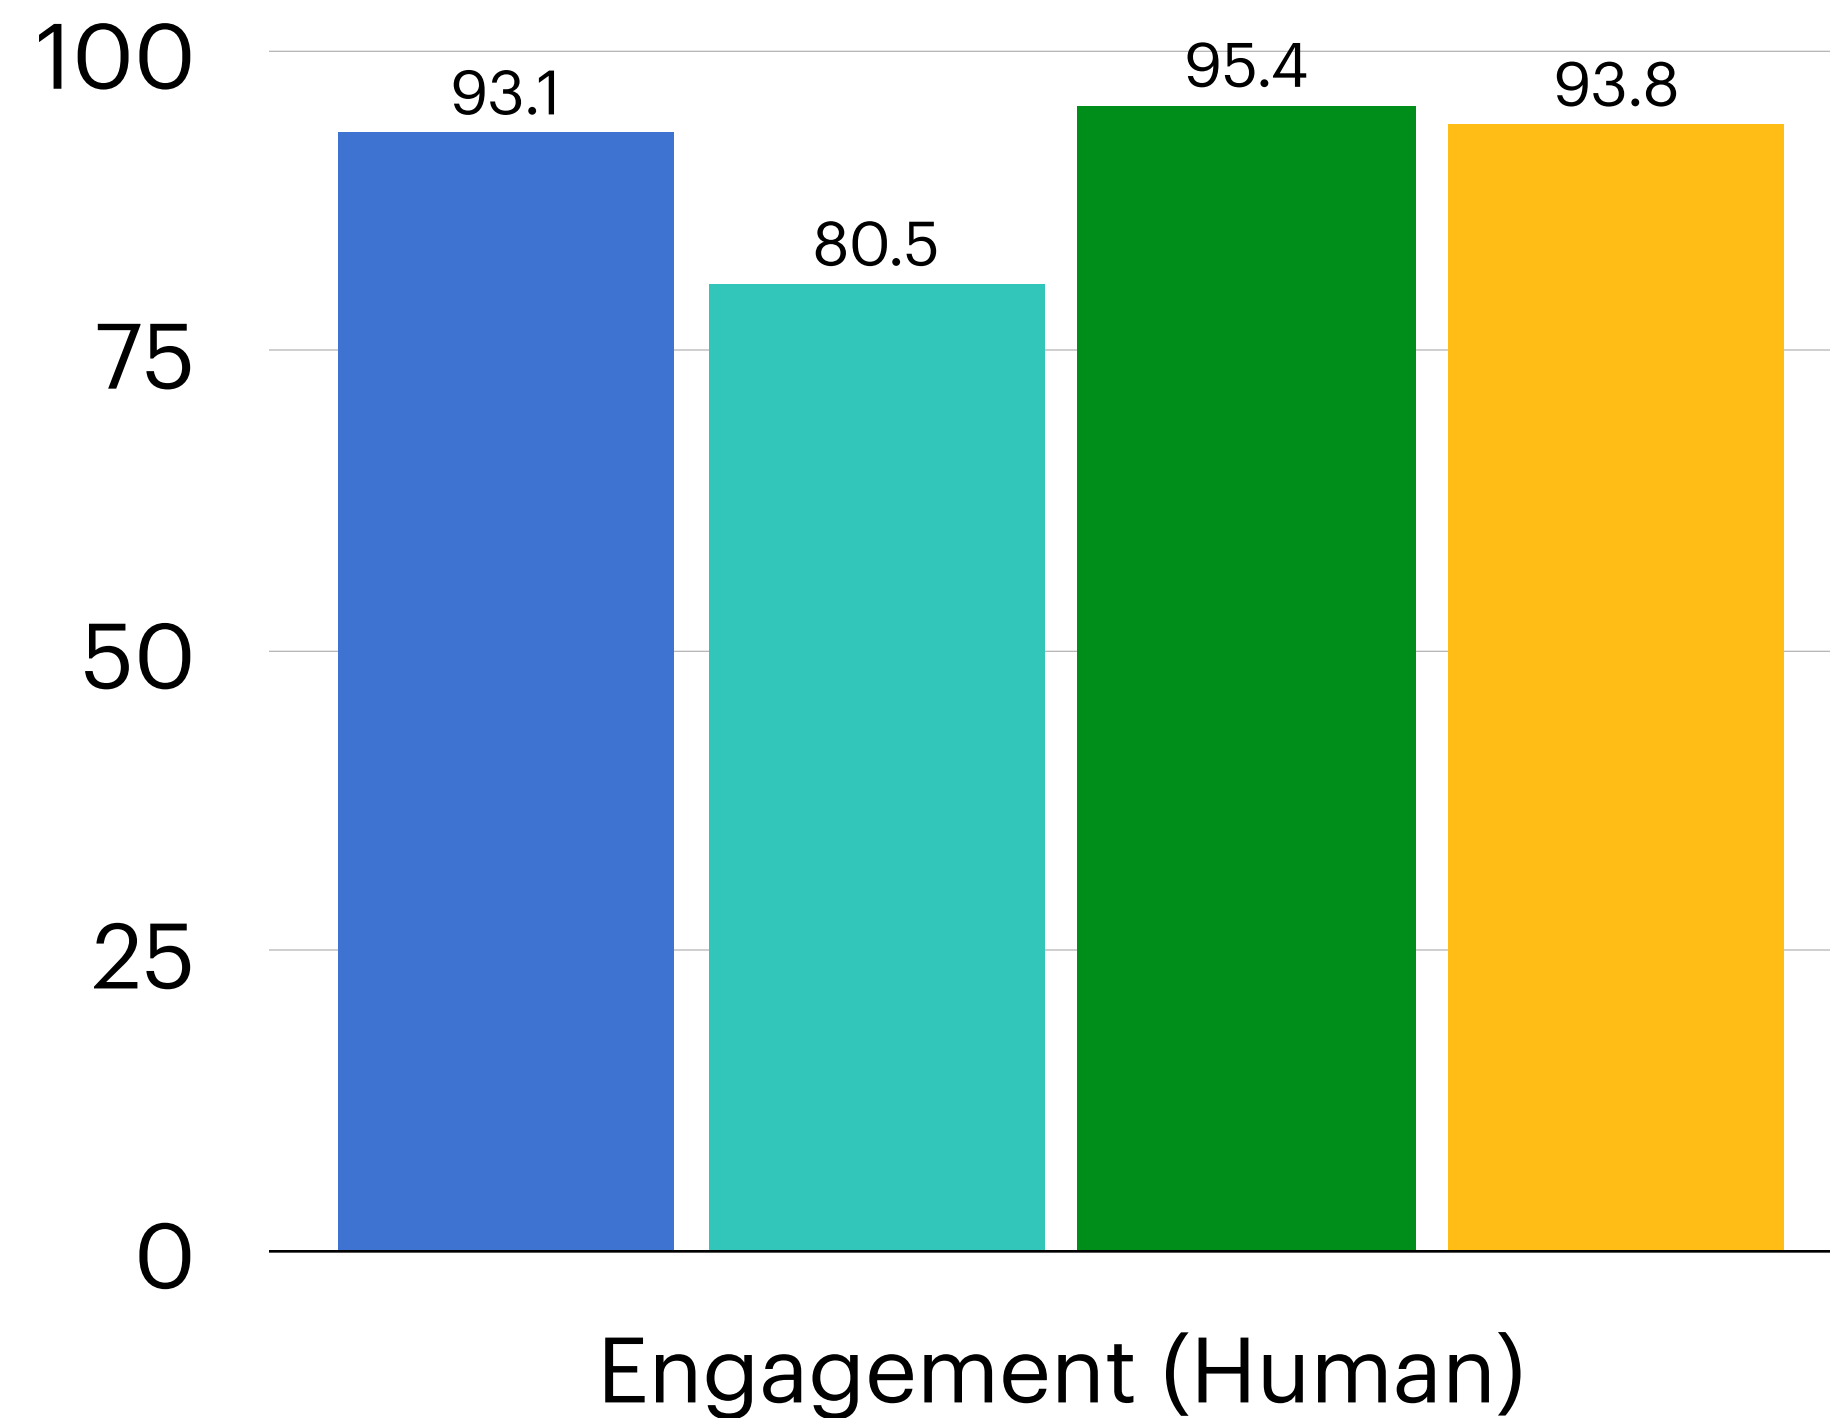


Diversity in dialog response correlates with user **engagingness**

PABST is more diverse than previous supervised approaches

Evaluation

■ GPT2 ■ COMPAC ■ Multi-task ■ Pseudo



Pairwise comparison with PABST, (*difference between win and loss %*, *higher is better*) shows our model is unanimously rated as **more engaging**

Qualitative Examples

Dialog History

Agent: How was your last trip?

User: It was great. I went to Utah and saw the Canyonlands.

DISCCHOICE: That's great. You must be adventurous.

RETRIEVAL: I went on a trip. I went to Florida. When we got there, there was no air conditioning. I was irritated. It was too hot.

PSEUDO: I have a house in Florida.

MULTITASK: I like Florida. I went to the beaches.

PABST ($\lambda_d = 5$): I went on a trip. I went to Florida. I have a house there. We do not have air condition. It was hot.

PABST ($\lambda_d = 1$): That's great. I frequently go to Florida. I have a house there but it gets hot in summer.

DISCCHOICE (aka **COMPAC**)

augments **no story**

Previous work (PSEUDO or MULTITASK) cannot adhere to the story completely, generating **less engaging** responses

RETRIEVAL **copies** the whole story, hence not fluent

Controlling verbosity

Dialog History

Agent: How was your last trip?

User: It was great. I went to Utah and saw the Canyonlands.

DISCHOICE: That's great. You must be adventurous.

RETRIEVAL: I went on a trip. I went to Florida. When we got there, there was no air conditioning. I was irritated. It was too hot.

PSEUDO: I have a house in Florida.

MULTITASK: I like Florida. I went to the beaches.

PABST ($\lambda_d = 5$): I went on a trip. I went to Florida. I have a house there. We do not have air condition. It was hot.

PABST ($\lambda_d = 1$): That's great. I frequently go to Florida. I have a house there but it gets hot in summer.

λ_d is a control for **verbosity**

Higher value of λ_d yields the response to be **verbatim** similar to the story


Human evaluation suggest a **moderate** λ_d is most engaging

Summary

Inference-time synthesis of non-conversational text yields **diverse** and **engaging** responses

Decoding-level constraints are useful and controllable

Thanks!



Code here!