

# DSC291: Advanced Statistical Natural Language Processing

## Text Generation

**Zhiting Hu**

Lecture 16, May 19, 2022

**UC San Diego**

**HALICIOĞLU DATA SCIENCE INSTITUTE**

# Outline

- Controllable text generation
- 2 Paper presentations (15 x 2 mins)
  - 
  -

# Two Central Goals

- Generating human-like, grammatical, and readable text
  - I.e., generating **natural** language
- Generating text that contains desired information inferred from inputs
  - Machine translation
    - Source sentence --> target sentence w/ the same meaning
  - Data description
    - Table --> data report describing the table
  - Attribute control
    - Sentiment: positive --> "I like this restaurant"
  - Conversation control
    - Control conversation strategy and topic

# Two Central Goals

- Generating human-like, grammatical, and readable text
  - Exposure bias, criteria mismatch: reinforcement learning (next lecture)
- Generating text that contains desired information inferred from inputs
  - Machine translation
    - Source sentence --> target sentence w/ the same meaning
  - Data description
    - Table --> data report describing the table
  - Attribute control
    - Sentiment: positive --> "I like this restaurant"
    - Modify sentiment from positive to negative
  - Conversation control
    - Control conversation strategy and topic

# Two Central Goals

- Generating human-like, grammatical, and readable text
  - Exposure bias, criteria mismatch: reinforcement learning (next lecture)

- Generating text that contains desired information inferred from inputs

#supervision data

- Machine translation
  - Source sentence --> target sentence w/ the same meaning -----> 10s of millions
- Data description
  - Table --> data report describing the table -----> 10s of 1000s
- Attribute control
  - Sentiment: positive --> "I like this restaurant" -----> 10s of 1000s
  - Modify sentiment from positive to negative -----> 0
- Conversation control
  - Control conversation strategy and topic -----> 0

# Two Central Goals

*Controlled generation in unsupervised settings*

- Generating human-like, grammatical, and readable text
  - Exposure bias, criteria mismatch: reinforcement learning (next lecture)

- Generating text that contains desired information inferred from inputs

#supervision data

- Machine translation
  - Source sentence --> target sentence w/ the same meaning -----> 10s of millions
- Data description
  - Table --> data report describing the table -----> 10s of 1000s
- Attribute control
  - Sentiment: positive --> "I like this restaurant" -----> 10s of 1000s
  - Modify sentiment from positive to negative -----> 0
- Conversation control
  - Control conversation strategy and topic -----> 0

# Unsupervised Controlled Generation of Text

- Sentence-level control
  - Text attribute transfer (style transfer)
  - Text content manipulation
- Conversation-level control
  - Target-guided open-domain conversation

# Unsupervised Controlled Generation of Text

- Sentence-level control
  - Text attribute transfer (style transfer)
  - Text content manipulation
- Conversation-level control
  - Target-guided open-domain conversation



# Text Attribute Transfer

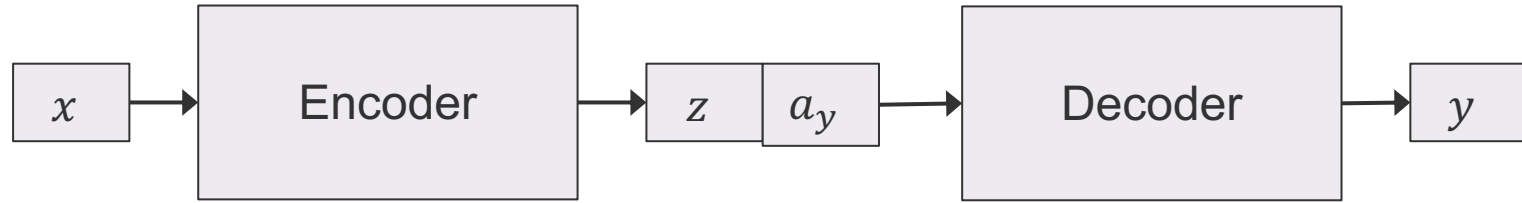
- Modify a given sentence to
  - Have desired attribute values
  - While keeping all other aspects unchanged
- Attribute: sentiment, tense, voice, gender, ...
  
- E.g., transfer sentiment from **negative** to **positive**:
  - “It was super **dry** and had a **weird** taste to the entire slice .”
  - “It was super **fresh** and had a **delicious** taste to the entire slice .”
- Applications:
  - Personalized article writing, emotional conversation systems, ...

# Text Attribute Transfer

- Original sentence  $\mathbf{x}$ , original attribute  $\mathbf{a}_x$
- Target sentence  $\mathbf{y}$ , target attribute  $\mathbf{a}_y$
- Task:  $(\mathbf{x}, \mathbf{a}_y) \rightarrow \mathbf{y}$ 
  - $\mathbf{y}$  has the desired attribute  $\mathbf{a}_y$
  - $\mathbf{y}$  keeps all attribute-independent properties of  $\mathbf{x}$
- Usually, only have pairs of  $(\mathbf{x}, \mathbf{a}_x)$ , but no  $((\mathbf{x}, \mathbf{a}_x), (\mathbf{y}, \mathbf{a}_y))$  for training
  - E.g., two sets of sentences: one with positive sentiment, the other with negative

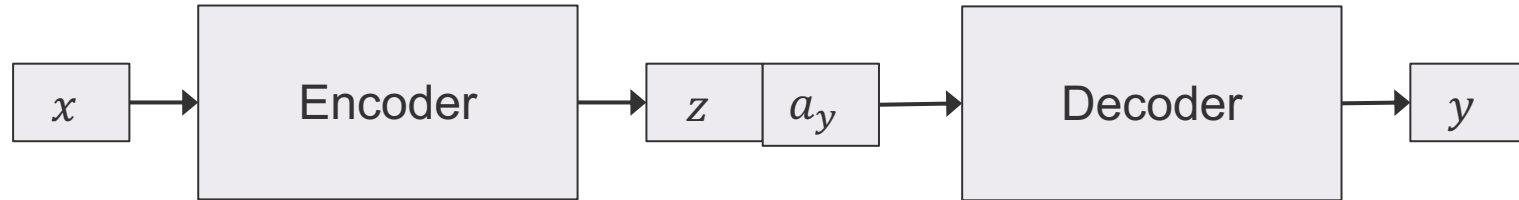
# Text Attribute Transfer: Solution

- Task:  $(x, a_y) \rightarrow y$ 
  - $y$  has the desired attribute  $a_y$
  - $y$  keeps all attribute-independent properties of  $x$
- Model  $p_\theta(y|x, a_y)$

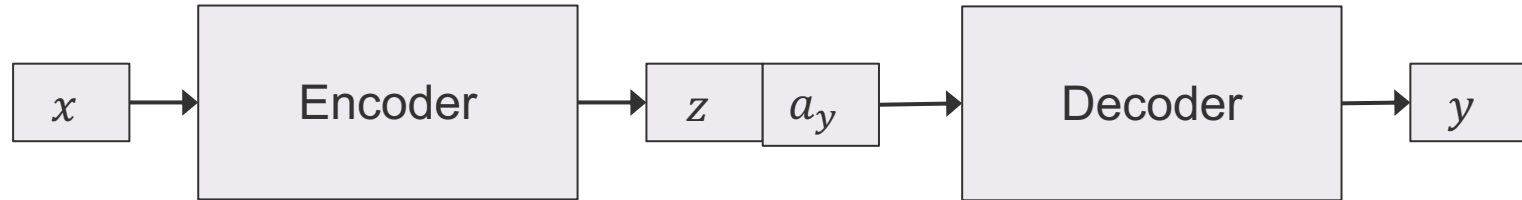


# Text Attribute Transfer: Solution

- Task:  $(x, a_y) \rightarrow y$ 
  - $y$  has the desired attribute  $a_y$
  - $y$  keeps all attribute-independent properties of  $x$
- Model  $p_\theta(y|x, a_y)$
- Key intuition for learning:
  - Decompose the task into competitive sub-objectives
  - Use direct supervision for each of the sub-objectives

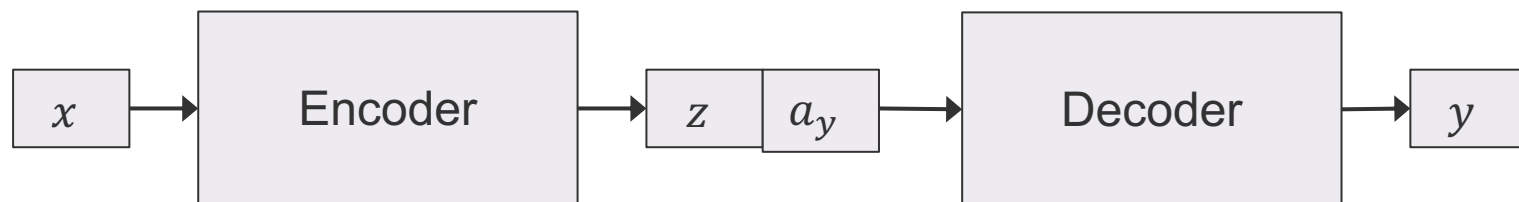


# Text Attribute Transfer: Solution



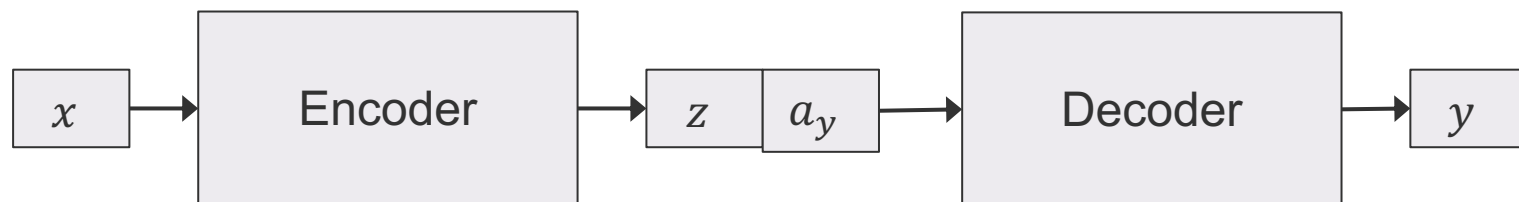
- Task:  $(x, a_y) \rightarrow y$ 
  - $y$  has the desired attribute  $a_y$
  - $y$  keeps all attribute-independent properties of  $x$
- Model  $p_\theta(y|x, a_y)$
- Key intuition for learning:
  - Decompose the task into competitive sub-objectives
  - Use direct supervision for each of the sub-objectives
- Auto-encoding loss:  $(x, a_x) \rightarrow x$

# Text Attribute Transfer: Solution

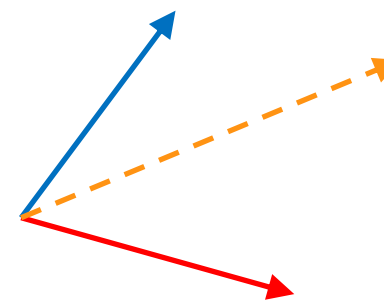


- Task:  $(x, a_y) \rightarrow y$ 
  - $y$  has the desired attribute  $a_y$
  - $y$  keeps all attribute-independent properties of  $x$
- Model  $p_\theta(y|x, a_y)$
- Key intuition for learning:
  - Decompose the task into competitive sub-objectives
  - Use direct supervision for each of the sub-objectives
- Auto-encoding loss:  $(x, a_x) \rightarrow x$
- Classification loss:  $\hat{y} \sim p_\theta(y|x, a_y), f(\hat{y}) \rightarrow a_y$ 
  - where  $f$  is a pre-trained attribute classifier

# Text Attribute Transfer: Solution



- Task:  $(x, a_y) \rightarrow y$ 
  - $y$  has the desired attribute  $a_y$
  - $y$  keeps all attribute-independent properties of  $x$
- Model  $p_\theta(y|x, a_y)$
- Key intuition for learning:
  - Decompose the task into competitive sub-objectives
  - Use direct supervision for each of the sub-objectives
- Auto-encoding loss:  $(x, a_x) \rightarrow x$
- Classification loss:  $\hat{y} \sim p_\theta(y|x, a_y), f(\hat{y}) \rightarrow a_y$ 
  - where  $f$  is a pre-trained attribute classifier
- The above two losses are competitive; minimize jointly to avoid collapse



# Text Attribute Transfer: Results & Improvement

- Performance on sentiment:
  - Accuracy: 92%
  - BLEU against input sentence: 54



# Text Attribute Transfer: Results & Improvement

- Performance on sentiment:
  - Accuracy: 92%
  - BLEU against input sentence: 54
- Problem:
  - Language quality is often not good
  - LM perplexity: 239.8

**Original:** if i could give them a zero star review i would !

**Output:** if i **lite** give them a **sweetheart** star review i would !

**Original:** uncle george is very friendly to each guest

**Output:** uncle george is very **lackluster** to each guest

**Original:** the food is fresh and the environment is good

**Output:** the food is **atrocious** and the environment is **atrocious**

# Text Attribute Transfer: Results & Improvement

- Performance on sentiment:
  - Accuracy: 92%
  - BLEU against input sentence: 54
- Problem:
  - Language quality is often not good
  - LM perplexity: 239.8
- Improvement:
  - Use an LM as a direct supervision!
  - $\hat{\mathbf{y}} \sim p_{\theta}(\mathbf{y}|\mathbf{x}, \mathbf{a}_y)$ ,  $\max_{\theta} \text{LM}(\hat{\mathbf{y}})$
  - Accuracy: 91%
  - BLEU against input sentence: 57
  - LM perplexity: 60.9

**Original:** if i could give them a zero star review i would !

**Output:** if i **like** give them a **sweetheart** star review i would !

**Original:** uncle george is very friendly to each guest

**Output:** uncle george is very **lackluster** to each guest

**Original:** the food is fresh and the environment is good

**Output:** the food is **atrocious** and the environment is **atrocious**

# Text Attribute Transfer: Results & Improvement

- Performance on sentiment:
  - Accuracy: 92%
  - BLEU against input sentence: 54
- Problem:
  - Language quality is often not good
  - LM perplexity: 239.8
- Improvement:
  - Use an LM as a direct supervision!
  - $\hat{y} \sim p_{\theta}(y|x, a_y), \max_{\theta} \text{LM}(\hat{y})$
  - Accuracy: 91%
  - BLEU against input sentence: 57
  - LM perplexity: 60.9

**Original:** if i could give them a zero star review i would !

**Output:** if i **like** give them a **sweetheart** star review i would !

**+ LM:** if i can give them a great star review i would !

**Original:** uncle george is very friendly to each guest

**Output:** uncle george is very **lackluster** to each guest

**+ LM:** uncle george is very rude to each guest

**Original:** the food is fresh and the environment is good

**Output:** the food is **atrocious** and the environment is **atrocious**

**+ LM:** the food is bland and the environment is bad .

# Unsupervised Controlled Generation of Text

- Sentence-level control
  - Text attribute transfer (style transfer)
  - Text content manipulation
- Conversation-level control
  - Target-guided open-domain conversation

Key idea:

- Decompose the task into **competitive** sub-objectives
- Use **direct supervision** for each of the sub-objectives

# Unsupervised Controlled Generation of Text

- Sentence-level control
  - Text attribute transfer (style transfer)
  - Text content manipulation
- Conversation-level control
  - Target-guided open-domain conversation

Key idea:

- Decompose the task into **competitive** sub-objectives
- Use **direct supervision** for each of the sub-objectives

# Text Content Manipulation

- Generate a sentence to describe content in a given data record

*Data Record*

Name	Food	Area	Price	Near
Loch Fyne	Italian	Riverside	£20-25	Strada

# Text Content Manipulation

- Generate a sentence to describe content in a given data record
- But language is rich with variation -- there are diverse possible ways of saying the same content (writing style):
  - word choice, expressions, transitions, tones, ...

*Data Record*

Name	Food	Area	Price	Near
Loch Fyne	Italian	Riverside	£20-25	Strada

# Text Content Manipulation

- Generate a sentence to describe content in a given data record
- But language is rich with variation -- there are diverse possible ways of saying the same content (writing style):
  - word choice, expressions, transitions, tones, ...
- We want to control the **writing style**: use the writing style of a reference sentence

*Data Record*

Name	Food	Area	Price	Near
Loch Fyne	Italian	Riverside	£20-25	Strada



# Text Content Manipulation

- Generate a sentence to describe content in a given data record

*Data Record*

Name	Food	Area	Price	Near
Loch Fyne	Italian	Riverside	£20-25	Strada

*Exemplar 1*

Zizzi is a pub providing fine French dining but with an expensive price, located near Cocum in the city center.

*Generation 1*

Loch Fyne provides fine Italian dining with a £20-25 price, located near Strada at the riverside.

*Exemplar 2*

Located near the Blue Spice, there is a highly-rated place, the Mill, as a choice that frugally priced.

*Generation 2*

Located near Strada by the river, there is a place with Italian foods, Loch Fyne, as a choice that priced £20-25.

*Exemplar 3*

With a family-friendly atmosphere and a 5-star rating, Aromi is a pub in the city center.

*Generation 3*

With Italian foods and a moderate price range, Loch Fyne is near Strada at the riverside.

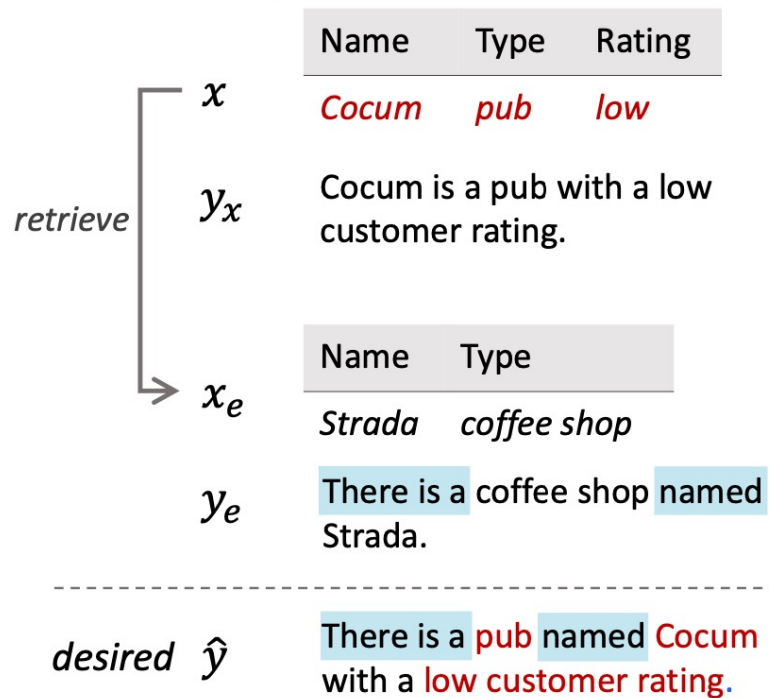
# Text Content Manipulation

- Generate a sentence to describe content in a given data record

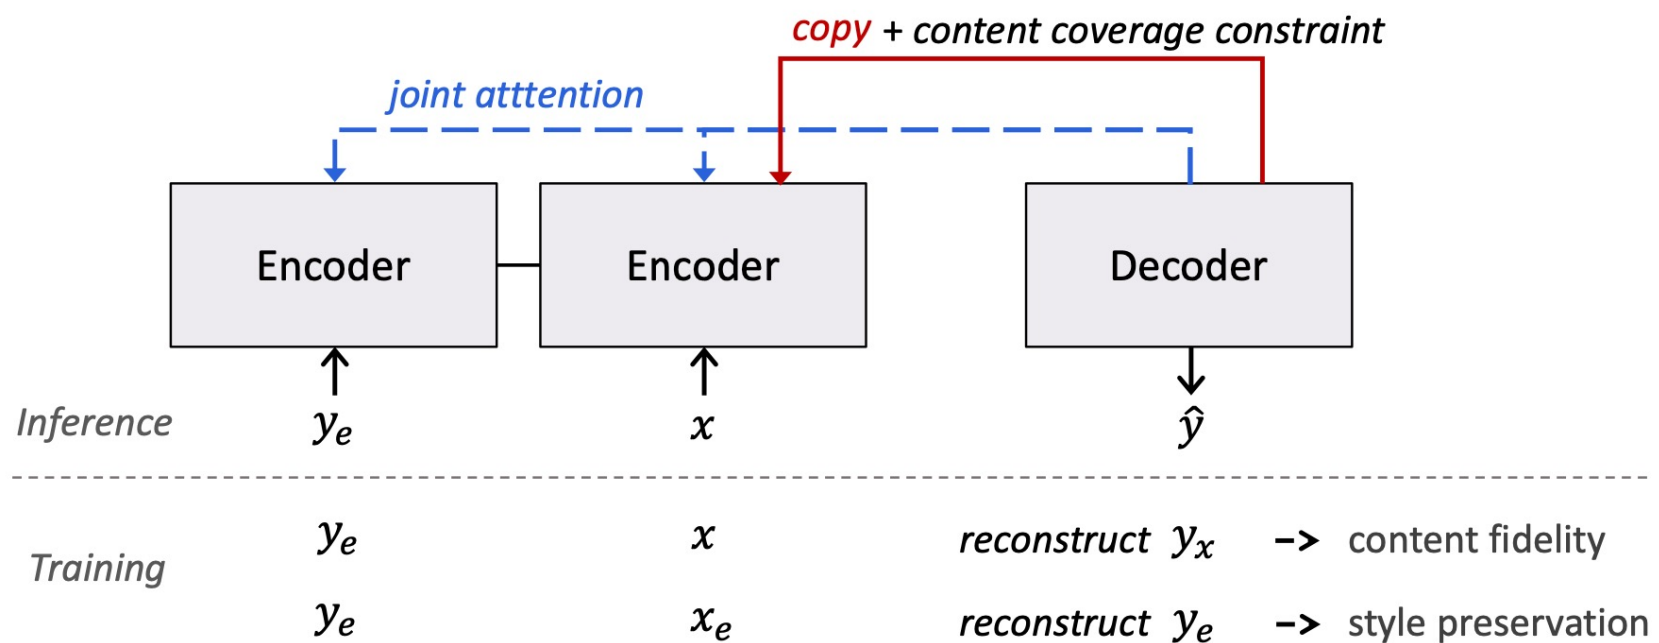
Content Record	<b>PLAYER</b> LeBron_James	<b>PT</b> 32	<b>RB</b> 4	<b>AS</b> 7	<b>PLAYER</b> Kyrie_Irving	<b>PT</b> 20
Reference Sentence	Jrue_Holiday led the way with 26 points and 6 assists , while Goran_Dragic scored 23 points and pulled down 8 rebounds .					
Output	LeBron_James led the way with 32 points , 7 assists and 4 rebounds , while Kyrie_Irving scored 20 points .					

# Method

Record and exemplar:



Model:



# Results

Content Record	Name	EatType	Food	PriceRange	CustomRating	FamilyFriendly
	Cocum	coffee shop	Italian	£20-25	high	family friendly
<b>Exemplar 1</b>	Looking for French food near Zizzi? Come try Strada, which has a 3-star customer rating and priced lowly.					
Slot filling	Looking for Italian [...] food near Zizzi? Come try [...] Cocum, which has a high customer rating and priced £20-25.					
AdvST	For Italian [...] place near Zizzi? Come try [...] Cocum, which has a high customer rating with priced £20-25.					
Ours	Looking for an Italian coffee shop? Come try family-friendly Cocum, which has a high customer rating and priced £20-25.					
<b>Exemplar 2</b>	Along the riverside near Cafe Rouge, there is a Japanese food place called The Golden Curry. It has an average customer rating since it is not a family-friendly environment.					
Slot-filling	Along the riverside near Cafe Rouge [...], there is a Italian food [...] place called Cocum. It has an high customer rating since it is not a family-friendly environment.					
AdvST	Along the riverside near the Ranch [...], there is a Italian food [...] place called Cocum. It has [...] high customer rating since it is not a family-friendly environment.					
Ours	Priced £20-25, there is an Italian food coffee shop called Cocum. It has a high customer rating since it is a family-friendly environment.					



# Results

		Restaurant Recommendations			NBA Reports		
Method		Content		Style	Content		Style
		% Incl.-new	% Excl.-old	m-BLEU	Precision	Recall	m-BLEU
<b>Reference</b>	AttnCopy-S2S	78.88 $\pm$ 2.08	99.71 $\pm$ 0.06	13.95 $\pm$ 0.52	81.62 $\pm$ 3.25	75.65 $\pm$ 7.42	45.5 $\pm$ 0.71
	Slot-filling	61.23	66.2	100	56.69	71.34	100
<b>Baselines</b>	MAST	36.28 $\pm$ 0.25	37.06 $\pm$ 0.16	<b>91.76<math>\pm</math>0.28</b>	23.06 $\pm$ 3.90	27.37 $\pm$ 3.88	<b>95.43<math>\pm</math>2.71</b>
	AdvST	51.64 $\pm$ 4.45	57.06 $\pm$ 4.44	76.02 $\pm$ 5.27	67.37 $\pm$ 0.66	66.79 $\pm$ 1.43	64.67 $\pm$ 4.81
<b>Ours</b>	Transformer w/o Coverage	60.03 $\pm$ 2.16	74.65 $\pm$ 2.69	77.81 $\pm$ 3.83	62.58 $\pm$ 2.88	70.22 $\pm$ 3.58	81.75 $\pm$ 2.32
	+ Coverage	61.84 $\pm$ 1.31	81.14 $\pm$ 2.73	80.29 $\pm$ 0.35	67.74 $\pm$ 0.79	<b>74.35<math>\pm</math>1.22</b>	81.97 $\pm$ 2.87
	LSTM w/o Coverage	60.83 $\pm$ 1.29	81.45 $\pm$ 1.10	78.91 $\pm$ 1.05	68.74 $\pm$ 3.07	69.35 $\pm$ 3.30	79.88 $\pm$ 2.44
	+ Coverage	<b>65.02<math>\pm</math>4.16</b>	<b>82.53<math>\pm</math>0.70</b>	82.92 $\pm$ 3.18	<b>69.54<math>\pm</math>1.16</b>	73.27 $\pm$ 1.18	80.66 $\pm$ 1.89

# Unsupervised Controlled Generation of Text

- Sentence-level control
  - Text attribute transfer (style transfer)
  - Text content manipulation
- Conversation-level control
  - Target-guided open-domain conversation

Key idea:

- Decompose the task into **competitive** sub-objectives
- Use **direct supervision** for each of the sub-objectives

# Unsupervised Controlled Generation of Text

- Sentence-level control
  - Text attribute transfer (style transfer)
  - Text content manipulation
- Conversation-level control
  - Target-guided open-domain conversation

Key idea:

- Decompose the task into **competitive** sub-objectives
- Use **direct supervision** for each of the sub-objectives

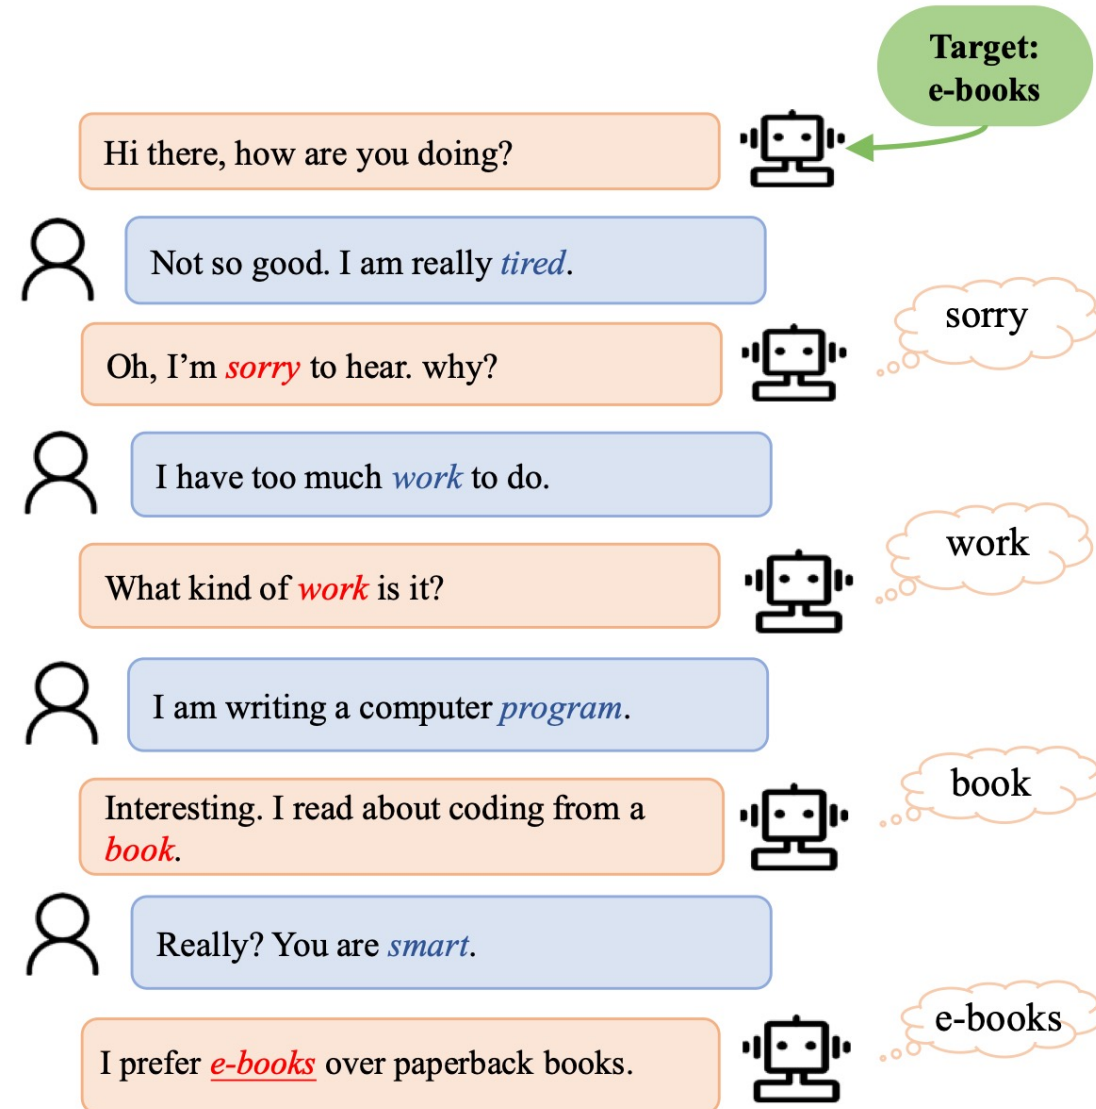
# Target-guided Open-domain Conversation

- Task-oriented dialog:
  - Address a specific task, e.g., booking a flight
  - Close domain
- Open-domain chit-chat:
  - Improve user engagement
  - Random conversation, hard to control
- Target-guided conversation:
  - Open-domain conversation
  - Controlled conversation strategy to reach a *desired topic* in the end of conversation
  - Applications:
    - Bridges task-oriented dialog and open-domain chit-chat
    - Conversational recommender system, education, psychotherapy



# Target-guided Open-domain Conversation

- Two goals:
  - Starting from any topic, reach a desired topic in the end of conversation
  - Natural conversation: smooth transition



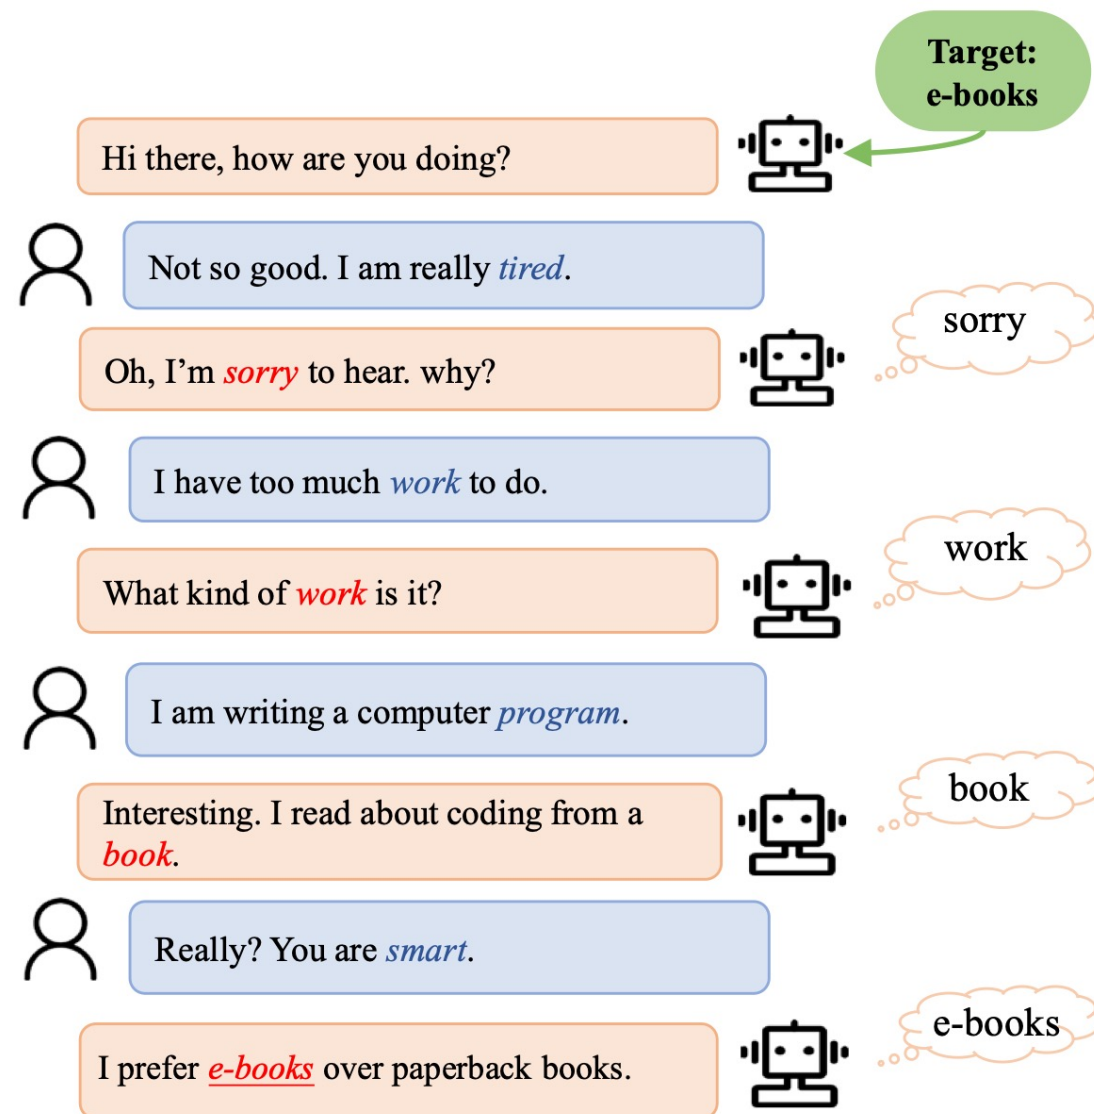
# Target-guided Open-domain Conversation

- Two goals:
  - Starting from any topic, reach a desired topic in the end of conversation
  - Natural conversation: smooth transition

**Challenge:** No supervised data for the task

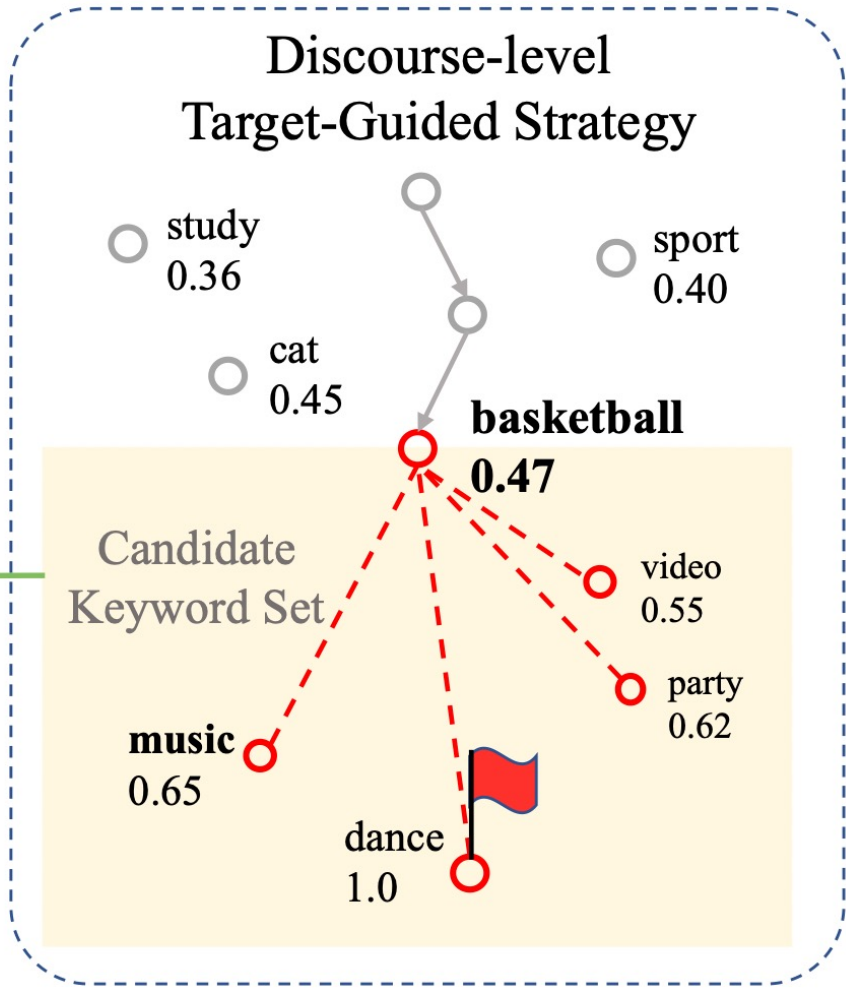
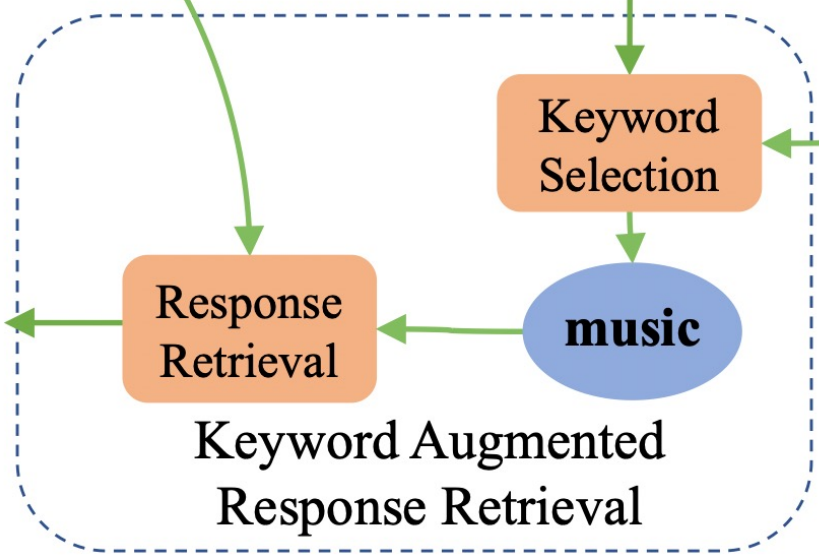
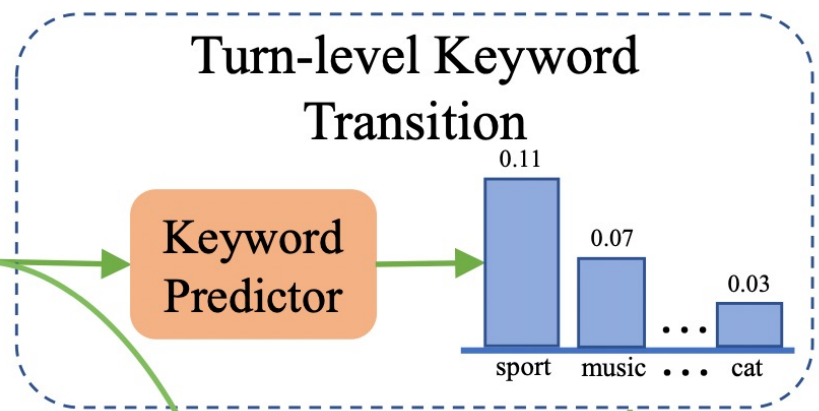
**Solution:** Use competitive sub-objectives and partial supervision

- **Natural conversation:** rich chit-chat data to learn smooth **single-turn** transition
- **Reaching desired target:** rule-based **multi-turn** planning

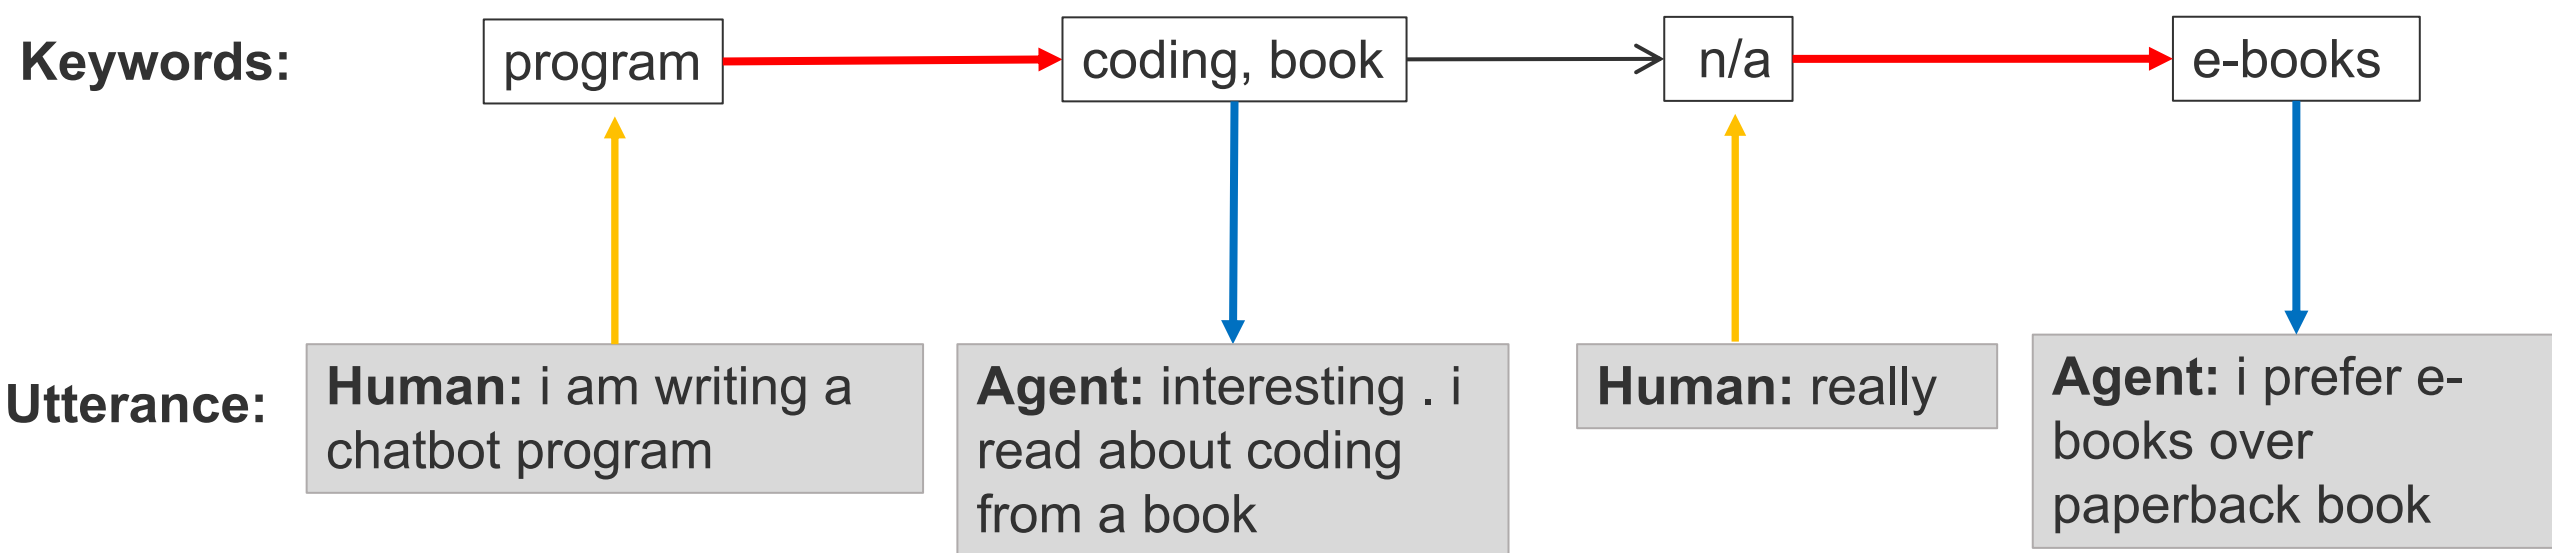


# Method

Target: dance

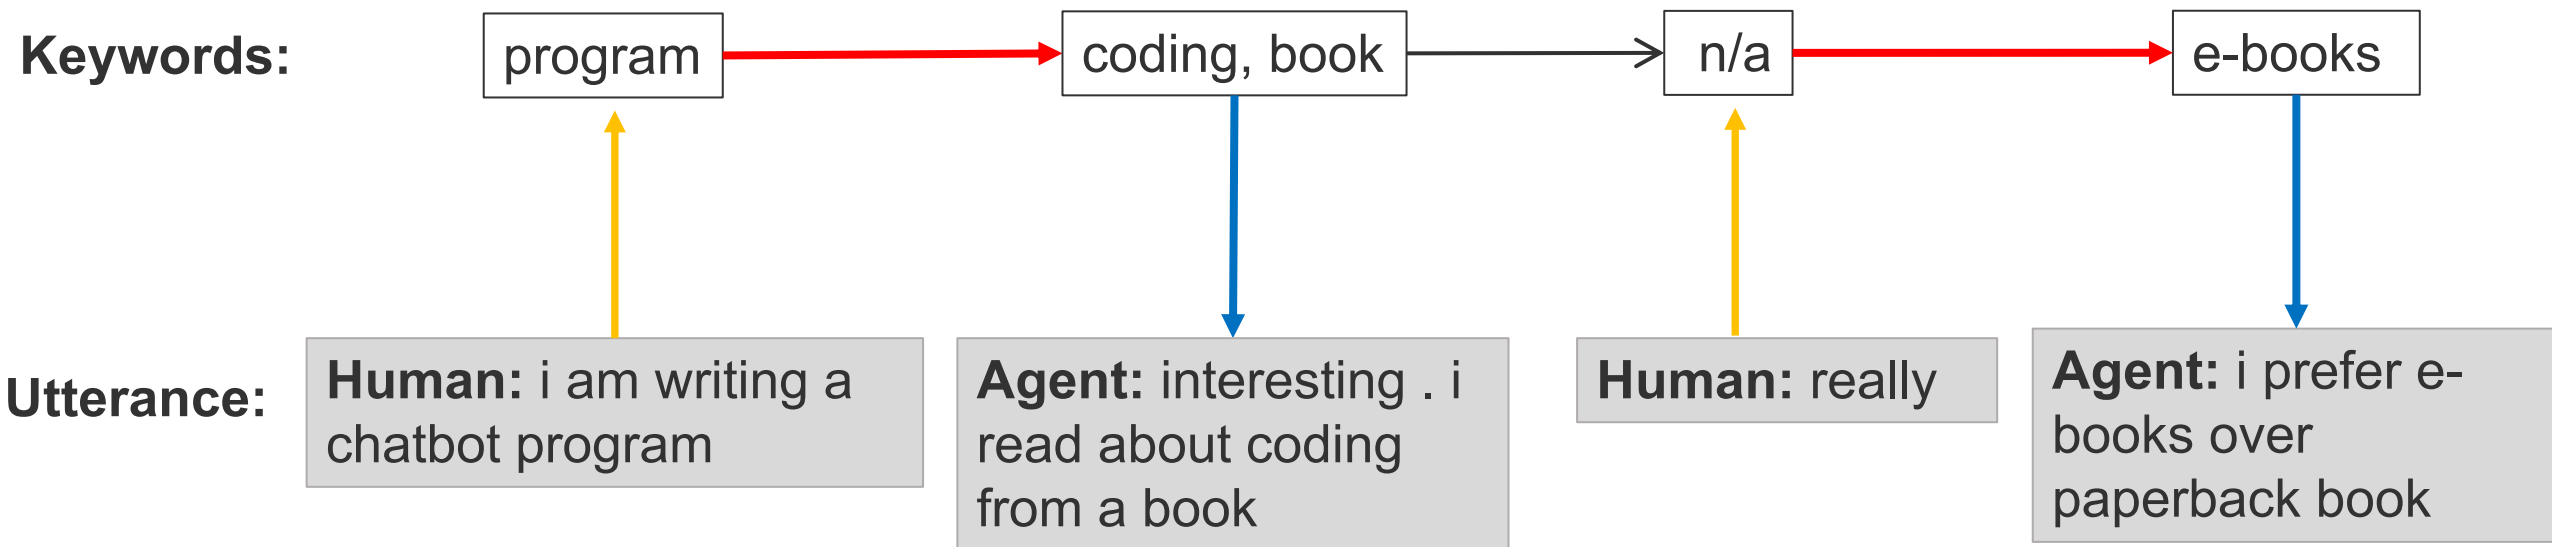


# Method





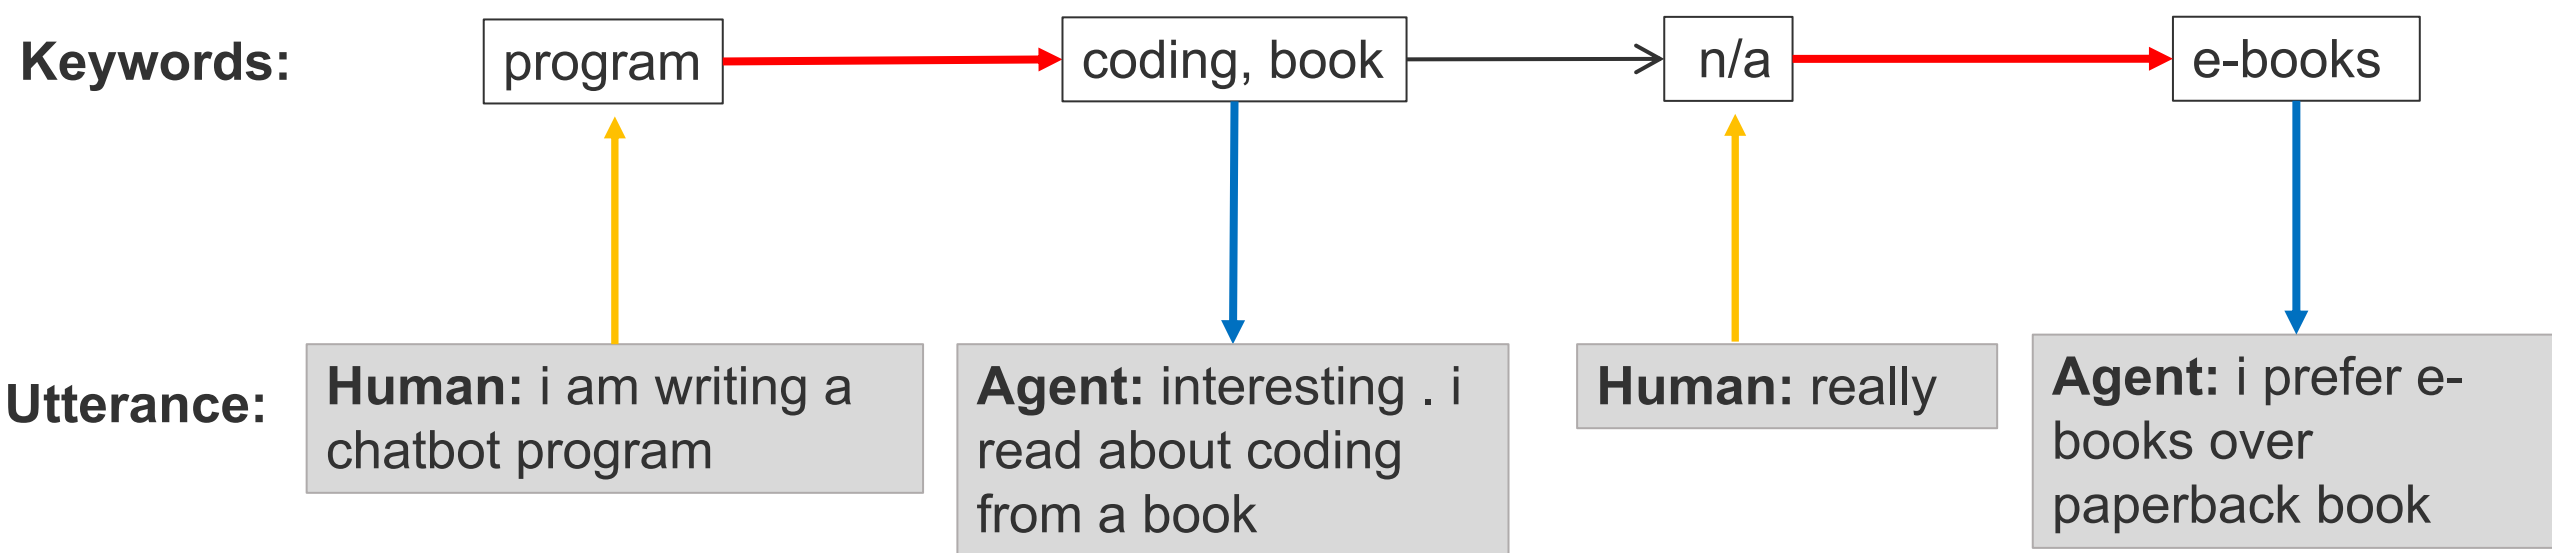
# Method

- → keyword extraction






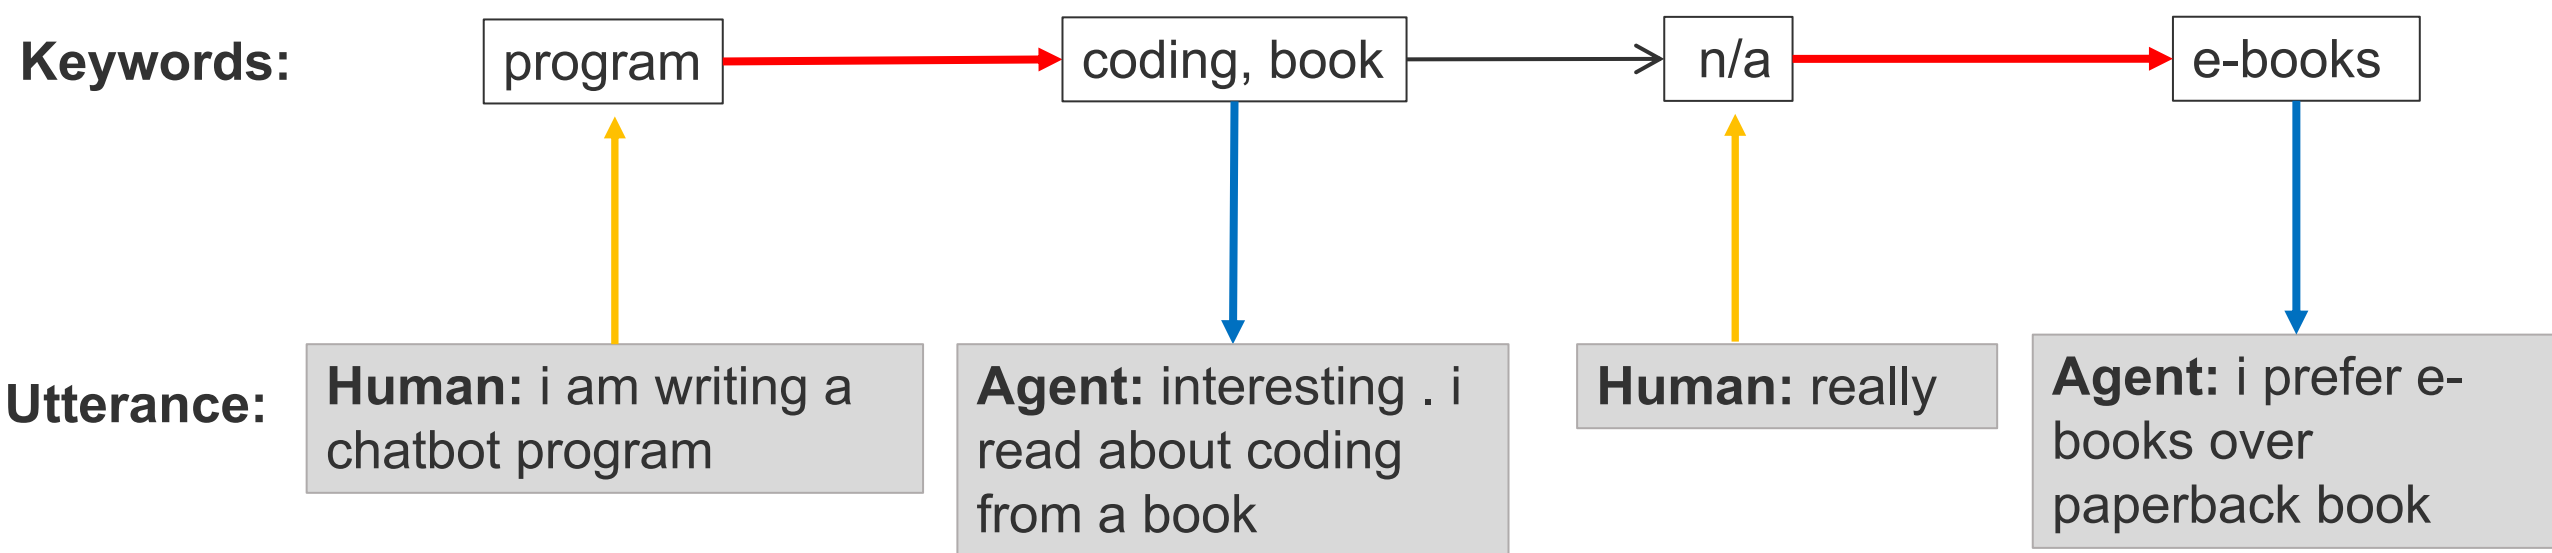
# Method

-  keyword extraction
-  keyword conditional response retrieval



# Method

-  keyword extraction
-  keyword conditional response retrieval
-  learned kernel-based topic transition + target-guided rule:
  - (1) transit to keywords that are close in the word embedding space – smooth
  - (2) next keywords must get closer to the target keyword – target-guided
  - Tune relative weight of (1) and (2) to control aggressiveness



# Example Results

**Target: dance**

Agent : hi how is it going with you

Human : fine . I'm playing sports

Agent : i play basketball . u play ?

Human : yeah i also like basketball .

Agent : you like rap music ? i listen to a lot of rap music

Human : no i like classic music

Agent : its a taylor swift song

Human : taylor swift is a great singer

Agent : sing to me , what are you going to sing

Human : i know nothing about singing

Agent : oh i think i can dance lol



# Example Results – Failure Case

**Target: listen**

Agent : hi ! how are you ? i am very exciting about my new job !

Human : fine . what is your new job ?

Agent : that sounds exciting .

Human : i like read books

Agent : what kind of books do you read ?

Human : i read comics

Agent : i do not have tv

Human : watch anime is fun

Agent : what position d you play ?

Human : i play a lot of sports

Agent : oh really ? what kind of music . i listen to classical

# Unsupervised Controlled Generation of Text

- Sentence-level control
  - Text attribute transfer (style transfer)
  - Text content manipulation
- Conversation-level control
  - Target-guided open-domain conversation

Key idea:

- Decompose the task into **competitive** sub-objectives
- Use **direct supervision** for each of the sub-objectives

# Key Takeaways: Two Central Goals

- Generating human-like, grammatical, and readable text
  - Exposure bias, criteria mismatch: reinforcement learning (next lecture)
- Generating text that contains desired information inferred from inputs
  - Machine translation
    - Source sentence --> target sentence w/ the same meaning
  - Data description
    - Table --> data report describing the table
  - Attribute control
    - Sentiment: positive --> "I like this restaurant"
  - Conversation control
    - Control conversation strategy and topic

Questions?