DSC190: Machine Learning with Few Labels

Case study: text generation

Zhiting Hu

Lecture 17, November 23, 2021

UC San Diego
HALICIOĞLU DATA SCIENCE INSTITUTE

Outline

- Text generation (50mins)

- 2 Paper presentations (30 mins)
  - **Viswesh Uppalapati:**
  - **James Yu: UCPhrase: Unsupervised Context-aware Quality Phrase Tagging**

Text Generation Tasks

- Generates **natural language** from input **data or machine representations**

Text Generation Tasks

- Generates **natural language** from input **data or machine representations**
- Spans a broad set of natural language processing (NLP) tasks:

Task	Input *X*	Output *Y* (Text)
Chatbot / Dialog System	Utterance	Response
Machine Translation	English	Chinese
Summarization	Document	Short paragraph
Description Generation	Structured data	Description
Captioning	Image/video	Description
Speech Recognition	Speech	Transcript

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
  - 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
  - I.e., generating **natural** language

  - (1) Model
  - (2) Learning

Common Model for Text Generation: Left-to-Right Language Model

- Calculates the probability of a sentence:
  - Sentence:

$$\boldsymbol{y} = (y_1, y_2 \dots, y_T)$$

$$p_\theta(\boldsymbol{y}) = \prod_t p_\theta(y_t \mid \boldsymbol{y}_{1:t-1})$$

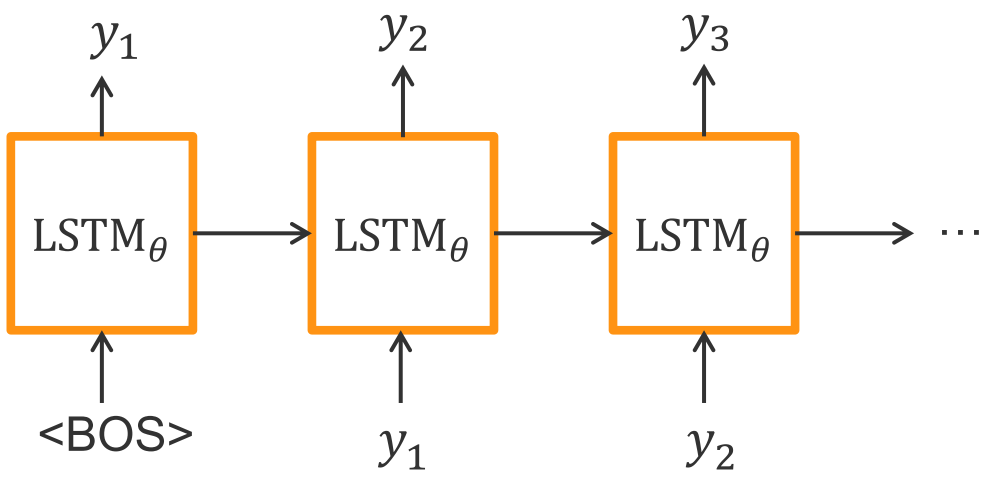

Common Model for Text Generation: Left-to-Right Language Model

- Calculates the probability of a sentence:
  - Sentence:

$$\boldsymbol{y} = (y_1, y_2 \dots, y_T)$$

$$p_\theta(\boldsymbol{y}) = \prod_t p_\theta(y_t \mid \boldsymbol{y}_{1:t-1})$$

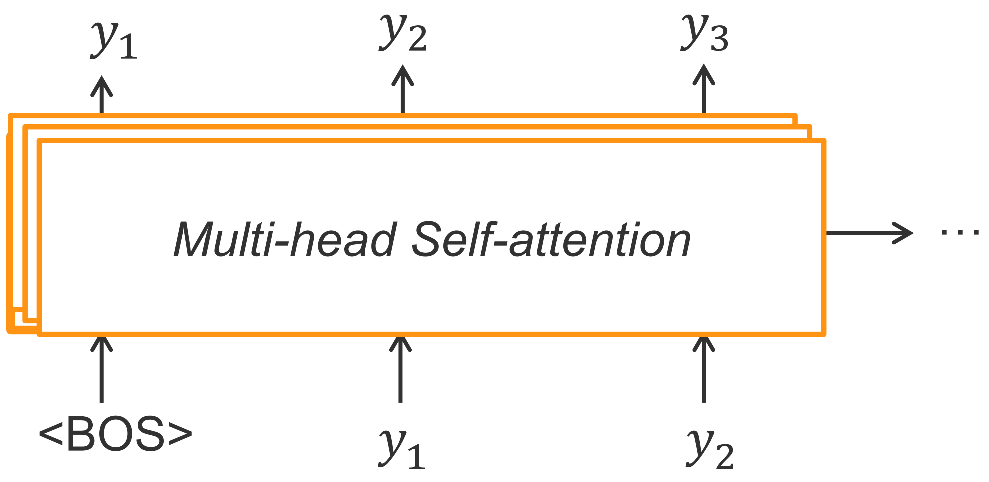

Common Model for Text Generation: Left-to-Right Language Model

- Difficulty of generating coherent long text (>1000 tokens)

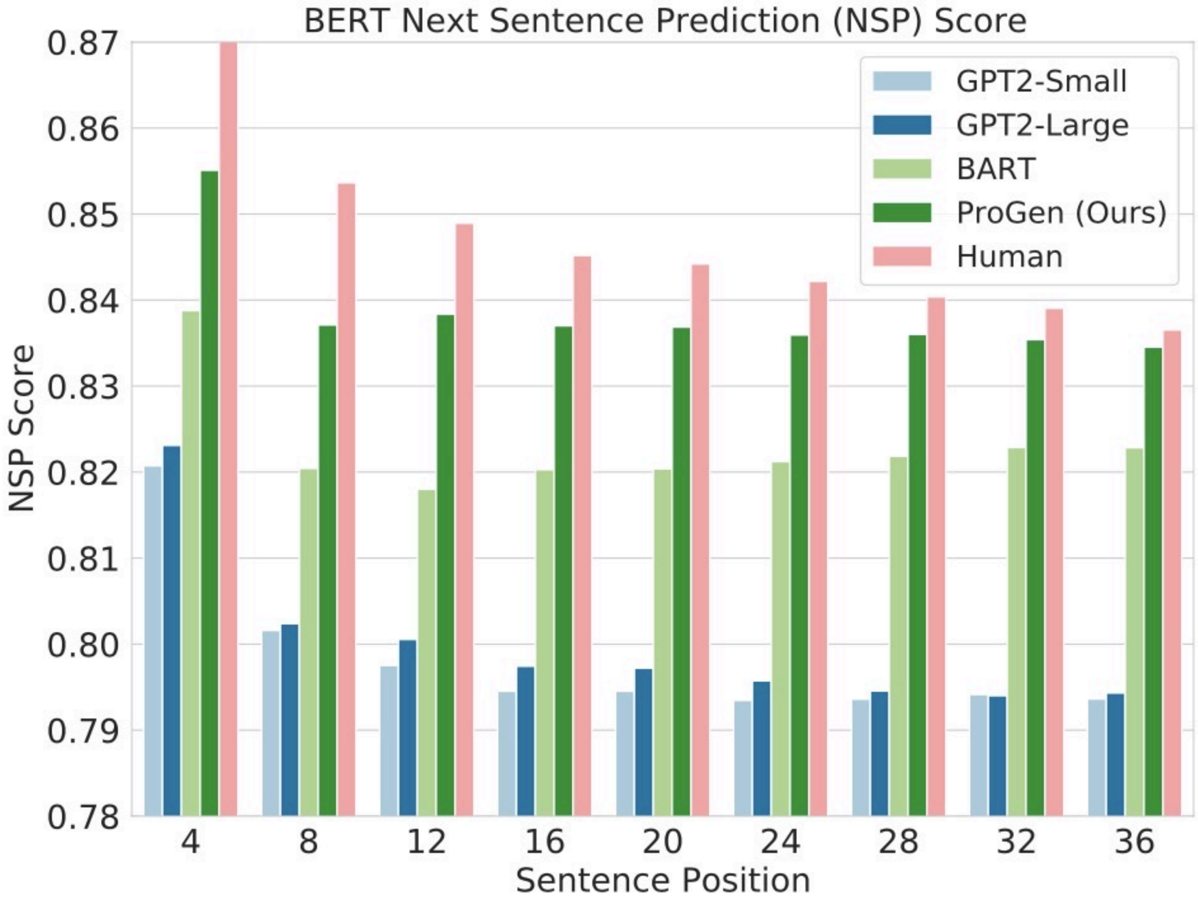

Common Model for Text Generation: Left-to-Right Language Model

- Difficulty of generating coherent long text (>1000 tokens)

AI MARKETING

September 18th 2020

GPT-3 Exposed: Behind the Smoke and Mirrors

The Guardian wrote an article in September with the title A robot wrote this entire article. Are you scared yet, human?

They had to piece together 8 different 500-word essays to come up with something that was fit to be published. Think about that for a minute. There's nothing efficient about that!

Joe Stanley-Smith @TaxStanley · Sep 8, 2020
Replying to @guardian
Editor's note: Actually, we wrote the standfirst and the rather misleading headline. Also, the robot wrote eight times this much and we organised it to make it better...

Jarno Duursma (Tech-Human-Future)
@JarnoDuursma

Exactly. GPT-3 created eight different essays. The Guardian journalists picked the best parts of each essay (!). After this manual selection they edited the article into a coherent article. That is not the same as: "this artificial intelligent system wrote this article".

12:45 PM · Sep 8, 2020

19 5 Share this Tweet

No human being could ever give an editor 4,000 words and expect them to edit it down to 500! What this reveals is that on average, each essay contained about 60 words (12%) of usable content.

Progressive Text Generation Model

- Informative words: decisive, have long-term impact on the whole content of the passage
- Non-informative words (e.g., stop words): do not require many contexts

- Intuition:
  - First generate most informative words
  - Then progressively refine the sequence by adding finer-grained details

Tan et al., NAACL 2021

Progressive Text Generation Model

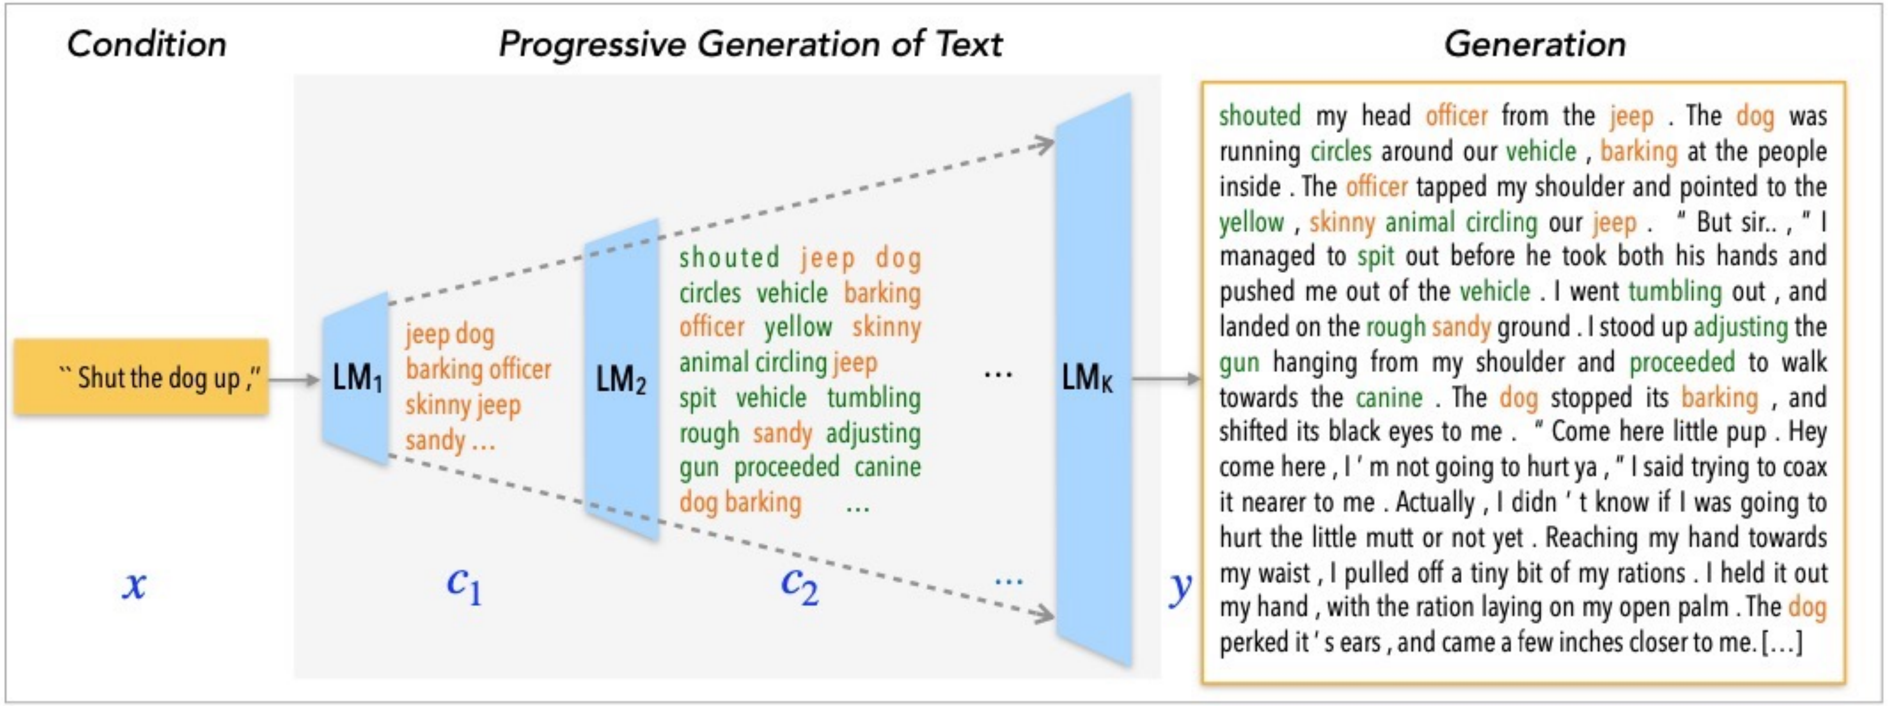

Figure 2: Progressive generation of long text $\mathbf{y}$ given any condition $\mathbf{x}$. Each stage refines the results from the previous stage by adding finer-grained details. Added content at each stage is highlighted in different colors.

Informativeness of words measured by TF-IDF

Tan et al., NAACL 2021

Progressive Text Generation Model

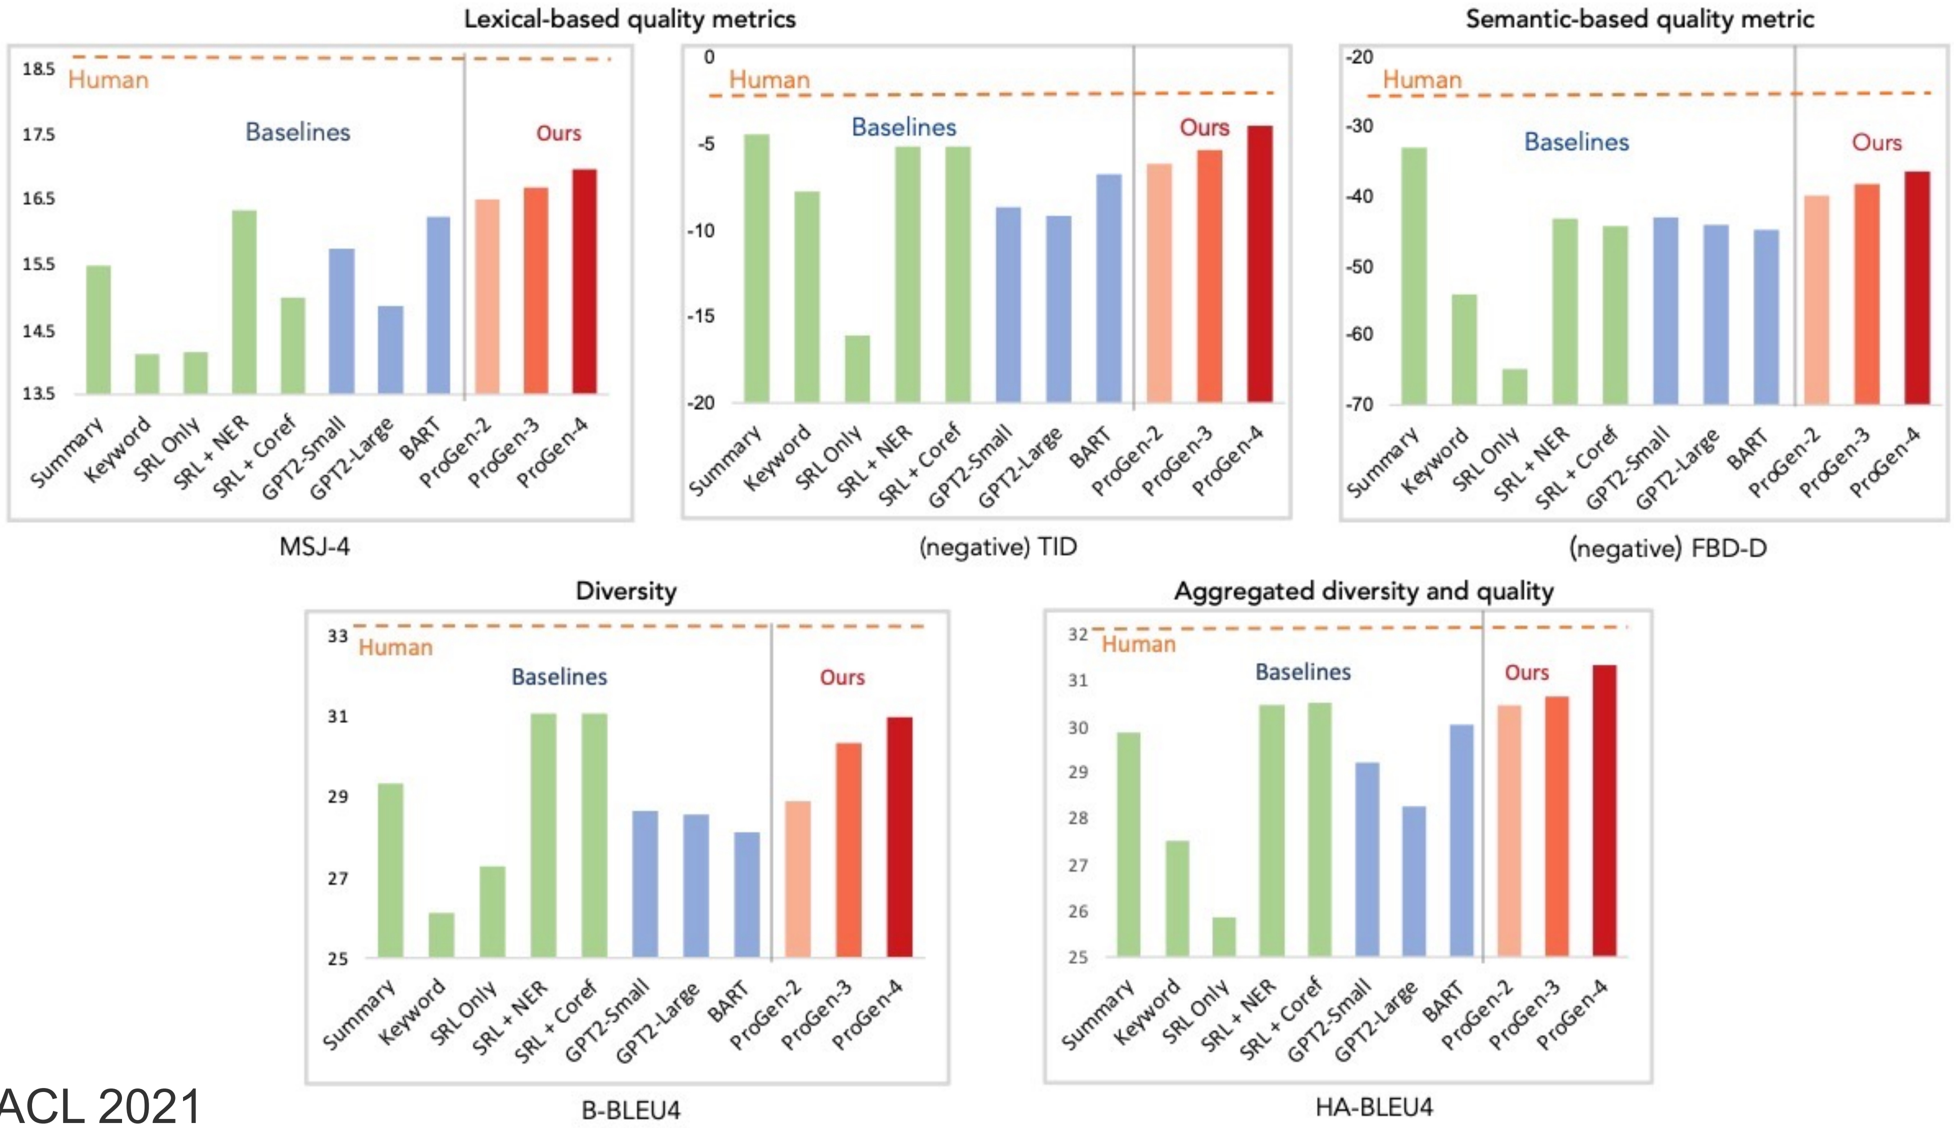

Tan et al., NAACL 2021

Progressive Text Generation Model

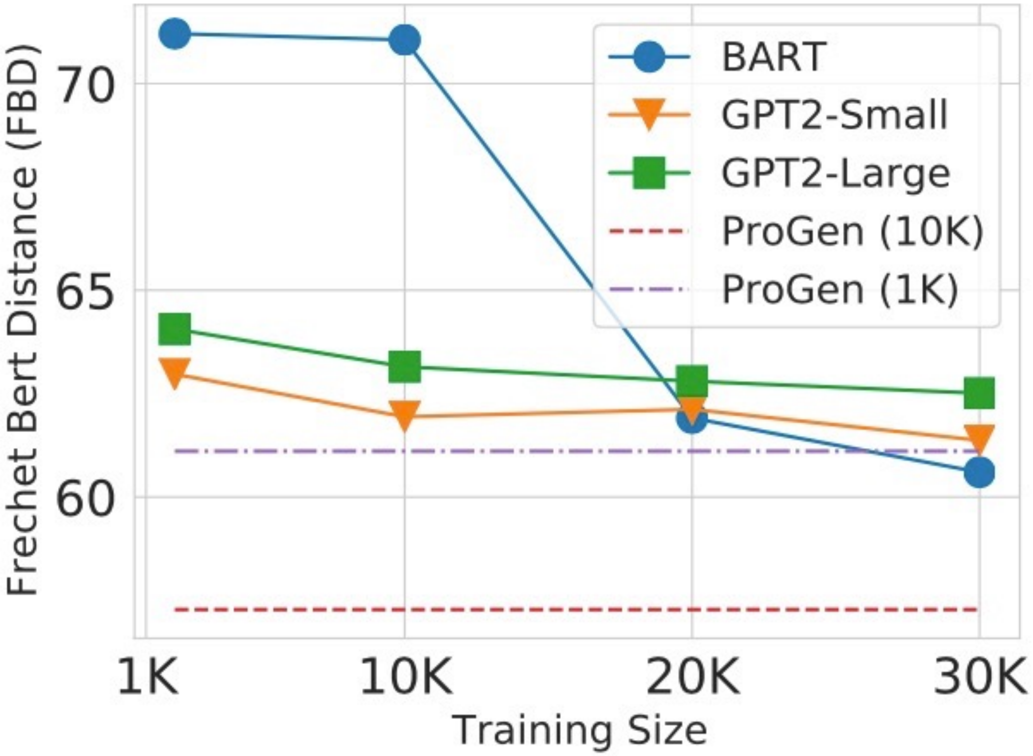

Figure 5: Sample efficiency on the story domain with the FBD metric (the lower, the better).

Tan et al., NAACL 2021

- Generating human-like, grammatical, and readable text
  - I.e., generating **natural** language
  - (1) Model
  - (2) Learning

Common Learning Algorithm: Maximum Likelihood Estimation (MLE)

- Training
  - Maximize data log-likelihood
  - Given ground-truth data

$$\boldsymbol{y}^* = (y_1^*, y_2^* \dots, y_{T^*}^*)$$

$$\mathcal{L}_{\mathrm{MLE}}(\boldsymbol{\theta}) = \log p_\theta(\boldsymbol{y}^* \mid \boldsymbol{x}) = \log \prod_t p_\theta(y_t^* \mid \boldsymbol{y}_{1:t-1}^*, \boldsymbol{x})$$

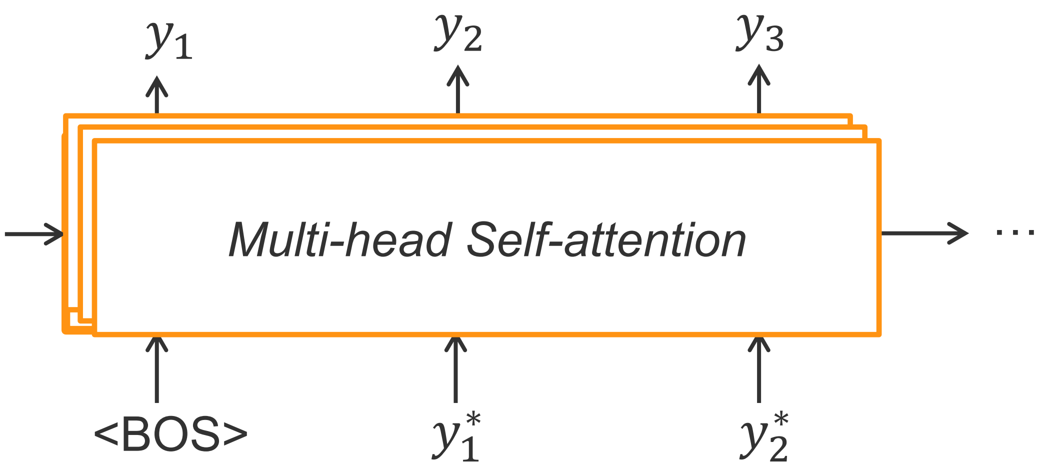

Common Learning Algorithm: Maximum Likelihood Estimation (MLE)

- Training
  - Maximize data log-likelihood
  - Given ground-truth data

$$\boldsymbol{y}^* = (y_1^*, y_2^* \dots, y_{T^*}^*)$$

$$\mathcal{L}_{\text{MLE}}(\boldsymbol{\theta}) = \log p_\theta(\boldsymbol{y}^* \mid \boldsymbol{x}) = \log \prod_t p_\theta(y_t^* \mid \boldsymbol{y}_{1:t-1}^*, \boldsymbol{x})$$

- Evaluation
  - Task-specific metrics
    - BLEU for machine translation
    - ROUGE for summarization
    - ….

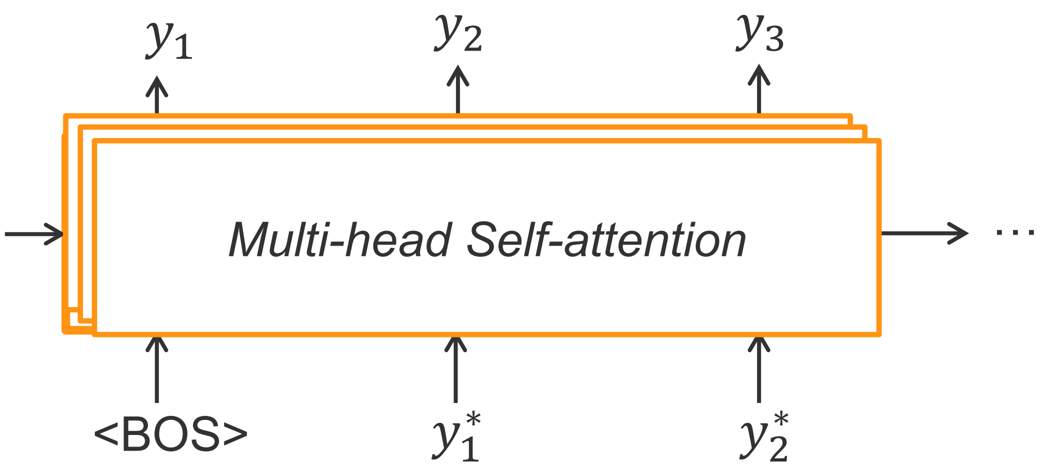

Two Issues of MLE

- Exposure bias [Ranzato et al., 2015]
  - **Training**: predict next token given the previous ground-truth sequence
  - **Evaluation**: predict next token given the previous sequence that are generated by the model itself

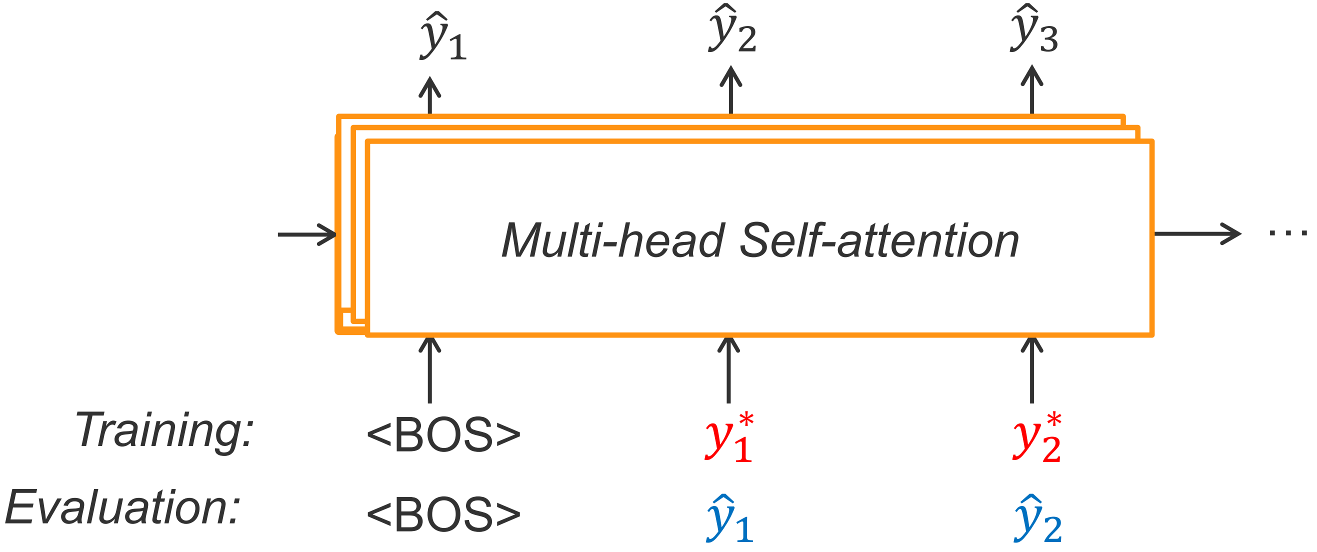

[Ranzato et al., 2015] Sequence Level Training with Recurrent Neural Networks

Two Issues of MLE

- Exposure bias [Ranzato et al., 2015]
  - **Training**: predict next token given the previous ground-truth sequence
  - **Evaluation**: predict next token given the previous sequence that are generated by the model itself

- Mismatch between training & evaluation criteria
  - Train to maximize **data log-likelihood**
  - Evaluate with, e.g., **BLEU**

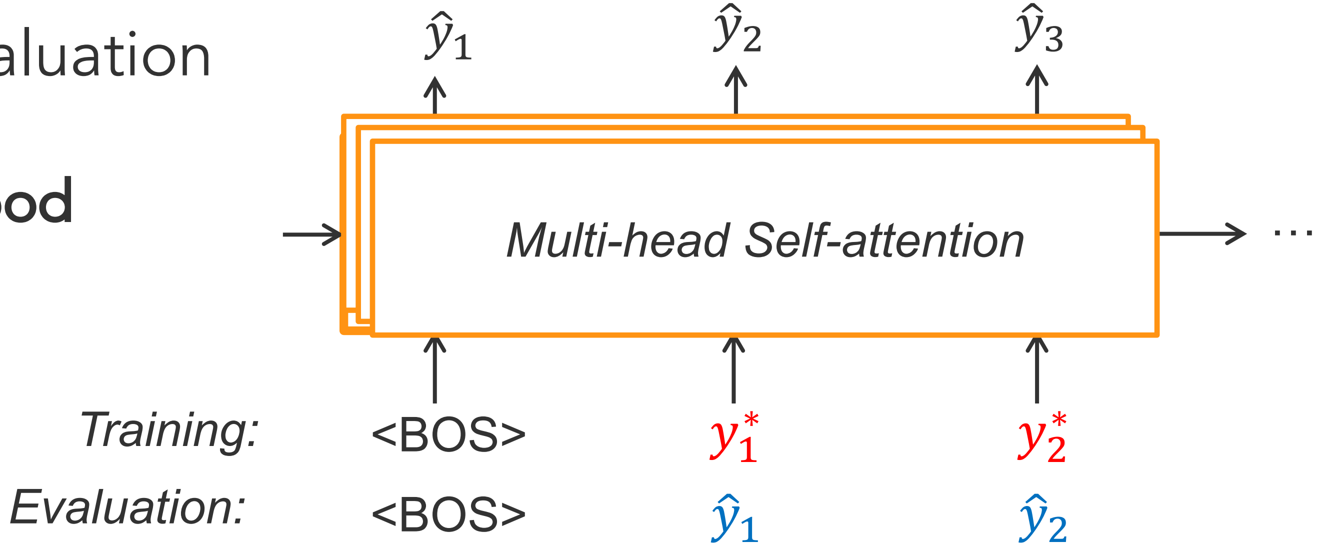

[Ranzato et al., 2015] Sequence Level Training with Recurrent Neural Networks

Two Issues of MLE

Solution: Reinforcement learning for text generation (next lecture)

- Exposure bias [Ranzato et al., 2015]
  - **Training**: predict next token given the previous ground-truth sequence
  - **Evaluation**: predict next token given the previous sequence that are generated by the model itself

- Mismatch between training & evaluation criteria
  - Train to maximize **data log-likelihood**
  - Evaluate with, e.g., **BLEU**

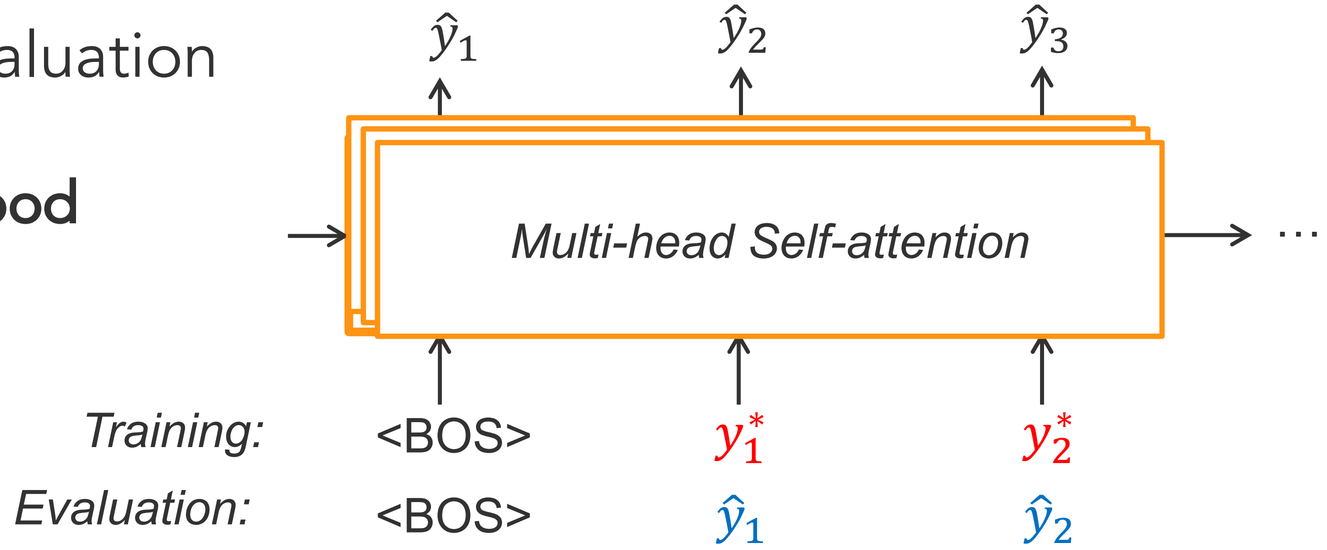

[Ranzato et al., 2015] Sequence Level Training with Recurrent Neural Networks

Two Central Goals

- Generating human-like, grammatical, and readable text
  - Progressive generation
  - 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
  - Progressive generation
  - 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
  - Progressive generation
  - 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

#supervision data

- Machine translation --> 10s of millions
- Data description --> 10s of 1000s
- Sentiment: positive --> ``I like this restaurant'' --> 10s of 1000s
- Modify sentiment from positive to negative --> 0
- Control conversation strategy and topic --> 0

Two Central Goals

Controlled generation in unsupervised settings

- Generating human-like, grammatical, and readable text
  - Progressive generation
  - 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

Task	#supervision data
Source sentence --> target sentence w/ the same meaning	10s of millions
Table --> data report describing the table	10s of 1000s
Sentiment: positive --> ``I like this restaurant''	10s of 1000s
Modify sentiment from positive to negative	0
Control conversation strategy and topic	0

Unsupervised Controlled Generation of Text

- Sentence-level control
  - Text attribute transfer (style transfer) [Hu et al., 2017; Yang et al., 2018]
  - Text content manipulation [Wang, Hu et al., 2019]
- Conversation-level control
  - Target-guided Open-domain Conversation

Unsupervised Controlled Generation of Text

- Sentence-level control
  - **Text attribute transfer** (style transfer) [Hu et al., 2017; Yang et al., 2018]
  - Text content manipulation [Wang, Hu et al., 2019]
- 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, conversation systems, authorship obfuscation

[Hu et al., 17] Toward Controlled Generation of Text

Text Attribute Transfer

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

Text Attribute Transfer: Solution

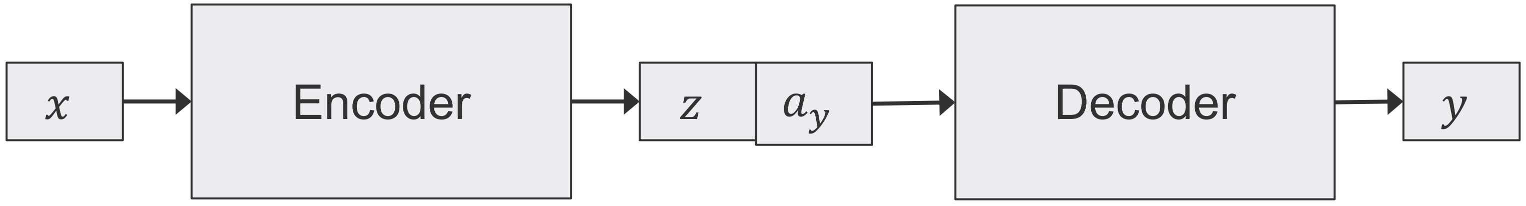

- Task: $(\boldsymbol{x}, \boldsymbol{a}_{\boldsymbol{y}}) \rightarrow \boldsymbol{y}$
  - $\boldsymbol{y}$ has the desired attribute $\boldsymbol{a}_y$
  - $\boldsymbol{y}$ keeps all attribute-independent properties of $\boldsymbol{x}$
- Model $p_\theta(\boldsymbol{y}|\boldsymbol{x}, \boldsymbol{a}_{\boldsymbol{y}})$

Text Attribute Transfer: Solution

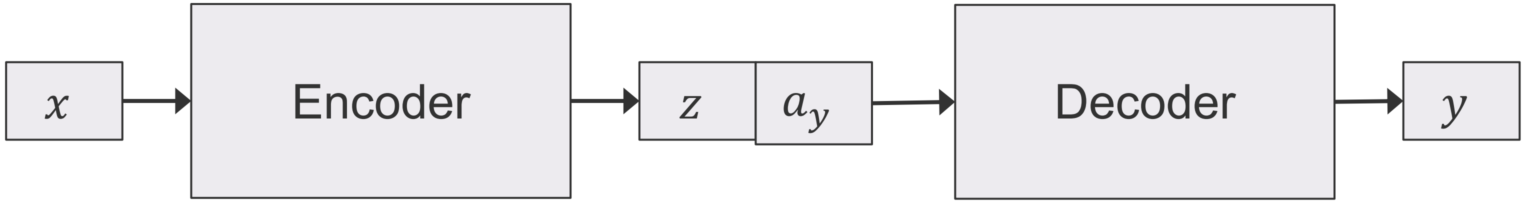

- Task: $(\boldsymbol{x}, \boldsymbol{a_y}) \rightarrow \boldsymbol{y}$
  - $\boldsymbol{y}$ has the desired attribute $\boldsymbol{a_y}$
  - $\boldsymbol{y}$ keeps all attribute-independent properties of $\boldsymbol{x}$
- Model $p_\theta(\boldsymbol{y}|\boldsymbol{x}, \boldsymbol{a_y})$
- Key intuition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

Text Attribute Transfer: Solution

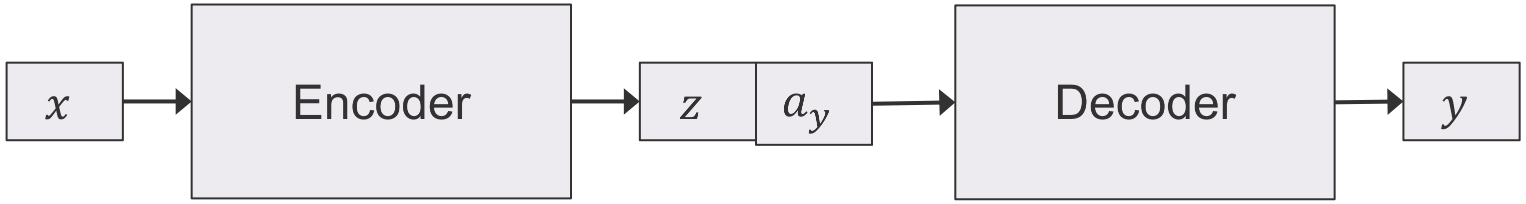

- Task: $(\boldsymbol{x}, \boldsymbol{a_y}) \rightarrow \boldsymbol{y}$
  - $\boldsymbol{y}$ has the desired attribute $\boldsymbol{a_y}$
  - $\boldsymbol{y}$ keeps all attribute-independent properties of $\boldsymbol{x}$
- Model $p_\theta(\boldsymbol{y}|\boldsymbol{x}, \boldsymbol{a_y})$
- Key intuition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

Text Attribute Transfer: Solution

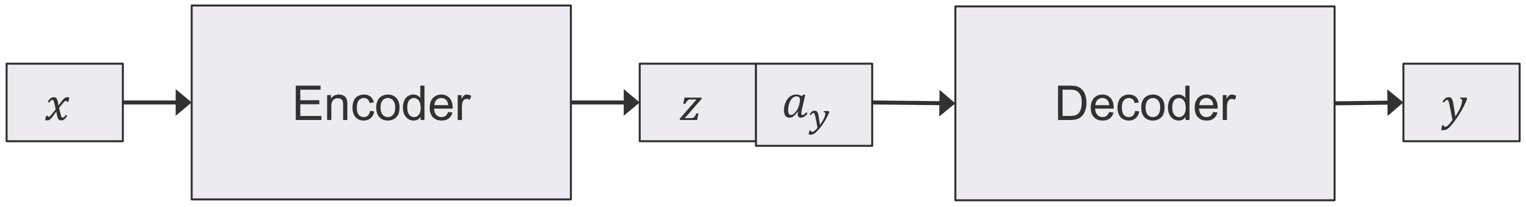

- Task: $(\boldsymbol{x}, \boldsymbol{a}_y) \rightarrow \boldsymbol{y}$
  - $\boldsymbol{y}$ has the desired attribute $\boldsymbol{a}_y$
  - $\boldsymbol{y}$ keeps all attribute-independent properties of $\boldsymbol{x}$
- Model $p_\theta(\boldsymbol{y}|\boldsymbol{x}, \boldsymbol{a}_y)$
- Key intuition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

Text Attribute Transfer: Solution

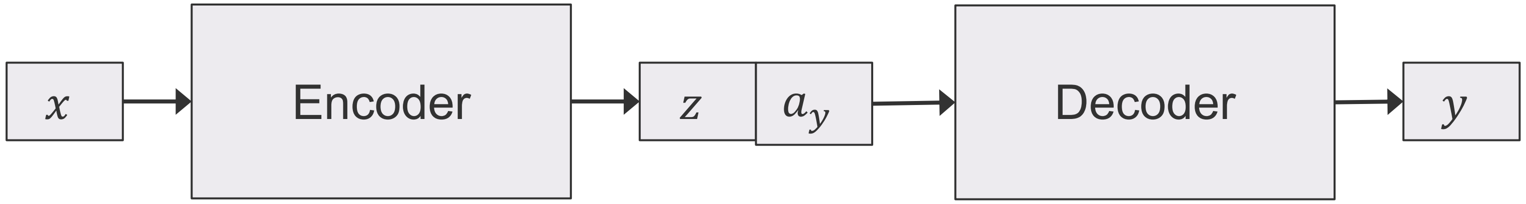

- Task: $(\boldsymbol{x}, \boldsymbol{a}_y) \to \boldsymbol{y}$
  - $\boldsymbol{y}$ has the desired attribute $\boldsymbol{a}_y$
  - $\boldsymbol{y}$ keeps all attribute-independent properties of $\boldsymbol{x}$
- Model $p_\theta(\boldsymbol{y}|\boldsymbol{x}, \boldsymbol{a}_y)$
- Key intuition for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

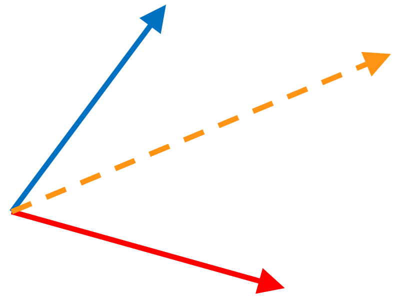

Text Attribute Transfer: Results & Improvement

- 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!
  - $\hat{\boldsymbol{y}} \sim p_\theta(\boldsymbol{y}|\boldsymbol{x}, \boldsymbol{a}_y)$, $\max_\theta \mathrm{LM}(\hat{\boldsymbol{y}})$
  - 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 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

[Yang et al., 18] Unsupervised text style transfer using language models as discriminators

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!
  - $\hat{\boldsymbol{y}} \sim p_\theta(\boldsymbol{y}|\boldsymbol{x}, \boldsymbol{a}_y)$, $\max_\theta \mathrm{LM}(\hat{\boldsymbol{y}})$
  - 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 lite 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 .

[Yang et al., 18] Unsupervised text style transfer using language models as discriminators

Unsupervised Controlled Generation of Text

- Sentence-level control
  - Text attribute transfer (style transfer) [Hu et al., 2017; Yang et al., 2018]
  - Text content manipulation [Wang, Hu et al., 2019]
- 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) [Hu et al., 2017; Yang et al., 2018]
  - **Text content manipulation** [Wang, Hu et al., 2019]
- 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
- But language is rich with variation -- there are diverse possible ways of saying the same content (writing style):
  - word choice, expressions, transitions, tones, …

Content Record	**PLAYER**	**PT**	**RB**	**AS**	**PLAYER**	**PT**
	LeBron_James	32	4	7	Kyrie_Irving	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 .					

[Wang, Hu et al., 18] Toward Unsupervised Text Content Manipulation

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

Content Record	**PLAYER**	**PT**	**RB**	**AS**	**PLAYER**	**PT**
	LeBron_James	32	4	7	Kyrie_Irving	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 .					

Text Content Manipulation - Results

Content x	PLAYER	PTS	FGM	FGA	FG3M	FG3A	FTM	FTA	AST
	Gerald_Henderson	17	6	13	1	2	4	4	5

Reference y′	Kawhi_Leonard also had a solid offensive game , scoring 16 points (7 - 13 FG , 0 - 1 3Pt , 2 - 5 FT) and adding 5 assists and 5 rebounds .
Rule-based	Gerald_Henderson also had a solid offensive game , scoring 17 points (6 - 13 FG , 1 - 2 3Pt , 4 - 4 FT) and adding 5 assists and 5 rebounds .
AdvST	Gerald_Henderson also had a solid offensive game , scoring 13 points (13 - 13 FG , 2 - 2 3Pt , 4 - 4 FT) and adding 5 assists and 5 rebounds .
Ours w/o Cover.	Gerald_Henderson also had a solid offensive game , scoring 17 points (6 - 13 FG , 1 - 2 3Pt , 4 - 4 FT) and adding 5 assists and 5 rebounds .
Ours	Gerald_Henderson also had a solid offensive game , scoring 17 points (6 - 13 FG , 1 - 2 3Pt , 4 - 4 FT) and adding 5 assists .

Text Content Manipulation - Results

	Model	Content Precision%	Content Recall%	Style BLEU
1	AttnCopy-S2S	88.71±2.45	60.64±1.31	39.15±5.48
	Rule-based	62.63	63.64	100
2	MAST	33.15±0.78	31.09±0.63	**95.29±2.53**
	AdvST	66.51±1.08	56.03±0.56	72.22±1.47
3	Ours w/o Cover.	75.61±1.03	62.93±0.53	75.09±2.17
	Ours	**78.31±0.94**	**65.64±0.47**	80.83±1.89

Unsupervised Controlled Generation of Text

- Sentence-level control
  - Text attribute transfer (style transfer) [Hu et al., 2017; Yang et al., 2018]
  - Text content manipulation [Wang, Hu et al., 2019]
- 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) [Hu et al., 2017; Yang et al., 2018]
  - Text content manipulation [Wang, Hu et al., 2019]

- 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: **e-books**

Agent : hi how are you doing
Human : not so good . i am really tired
Agent : oh i'm sorry to hear
Human : i have a lot of work to do
Agent : what kind was it ?
Human : i am writing a chatbot program
Agent : interesting . i read about coding from a book
Human: really
Agent : i prefer e-books over paperback book

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: **e-books**

	Agent : hi how are you doing
tired	Human : not so good . i am really tired
sorry	Agent : oh i'm sorry to hear
work	Human : i have a lot of work to do
work	Agent : what kind was it ?
program	Human : i am writing a chatbot program
coding, book	Agent : interesting . i read about coding from a book
	Human: really
e-books	Agent : i prefer e-books over paperback book

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: **e-books**

	Agent : hi how are you doing
tired	Human : not so good . i am really tired
sorry	Agent : oh i'm sorry to hear
work	Human : i have a lot of work to do
work	Agent : what kind was it ?
program	Human : i am writing a chatbot program
coding, book	Agent : interesting . i read about coding from a book
	Human: really
e-books	Agent : i prefer e-books over paperback book

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

Target-guided Open-domain Conversation

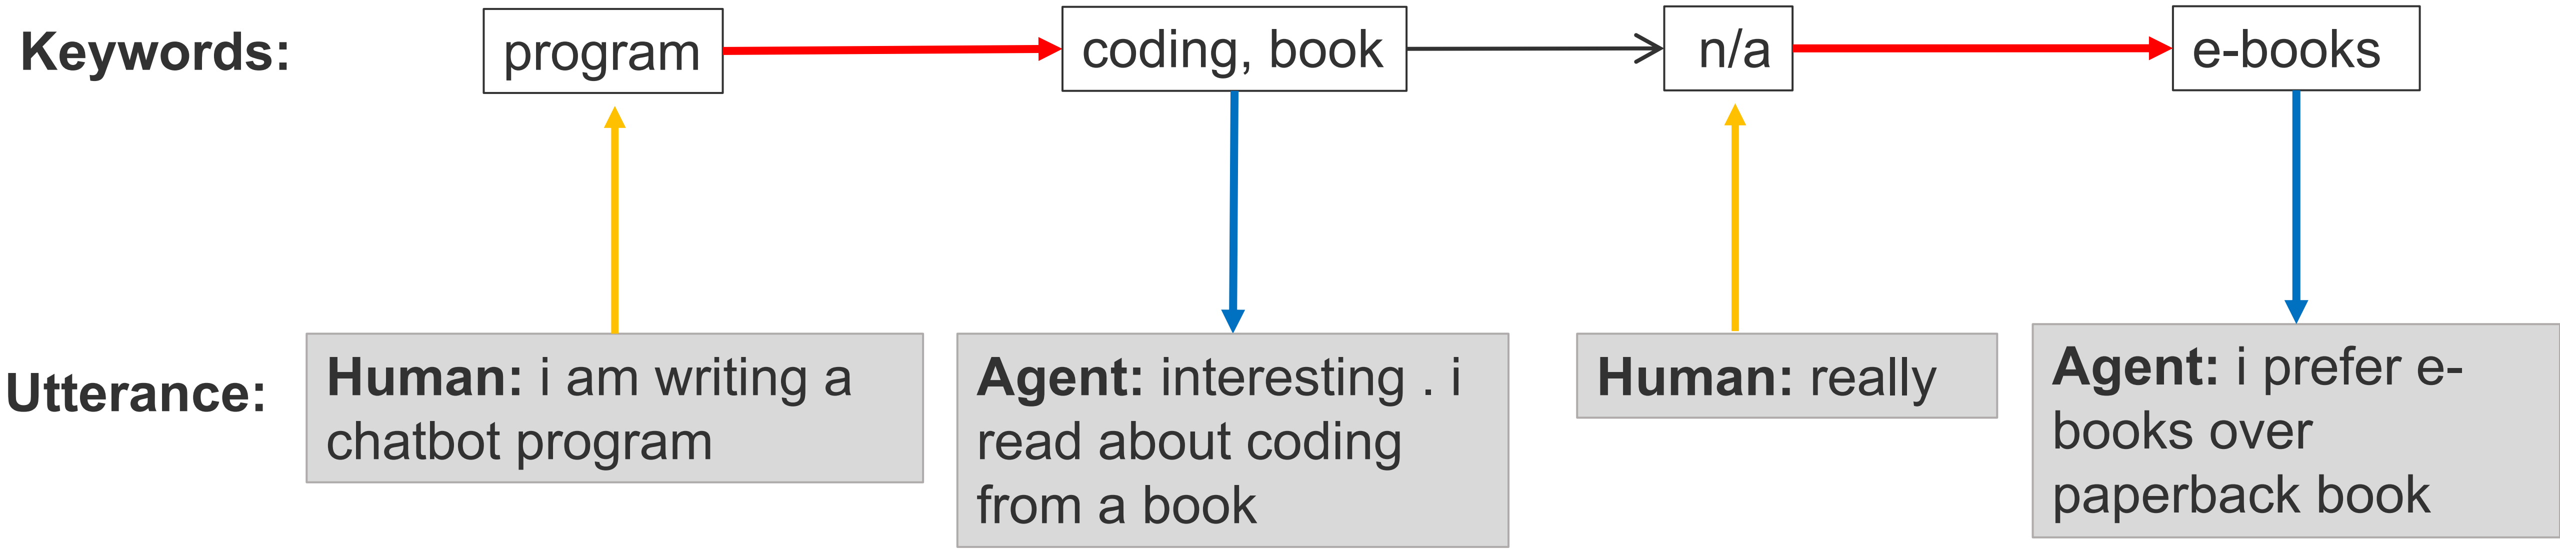

Target-guided Open-domain Conversation

- → keyword extraction

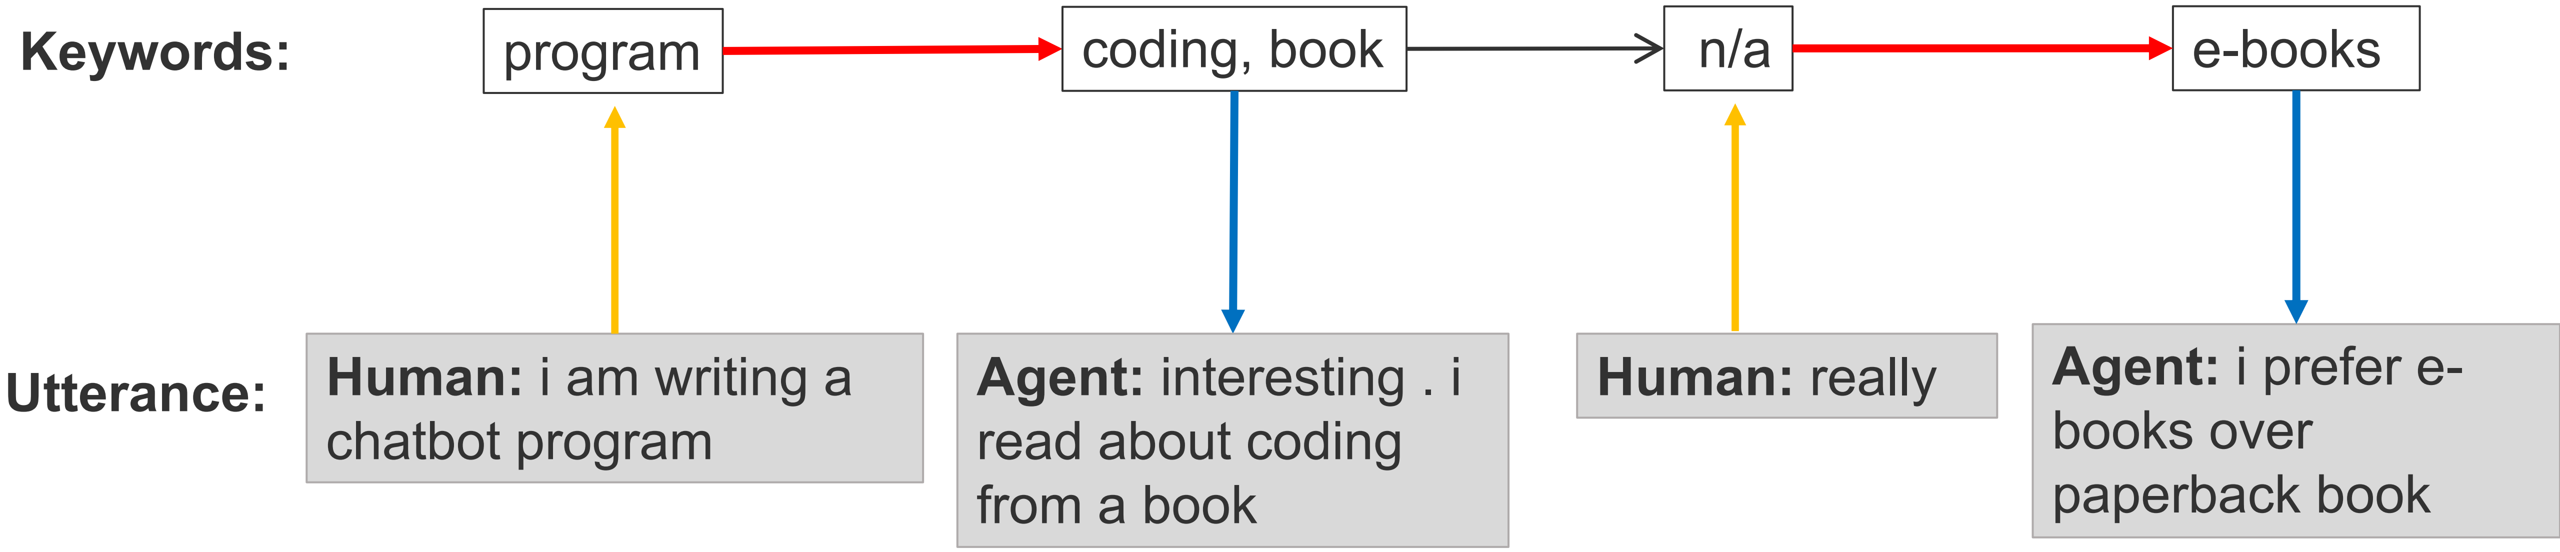

Target-guided Open-domain Conversation

- → keyword extraction
- → keyword conditional response retrieval

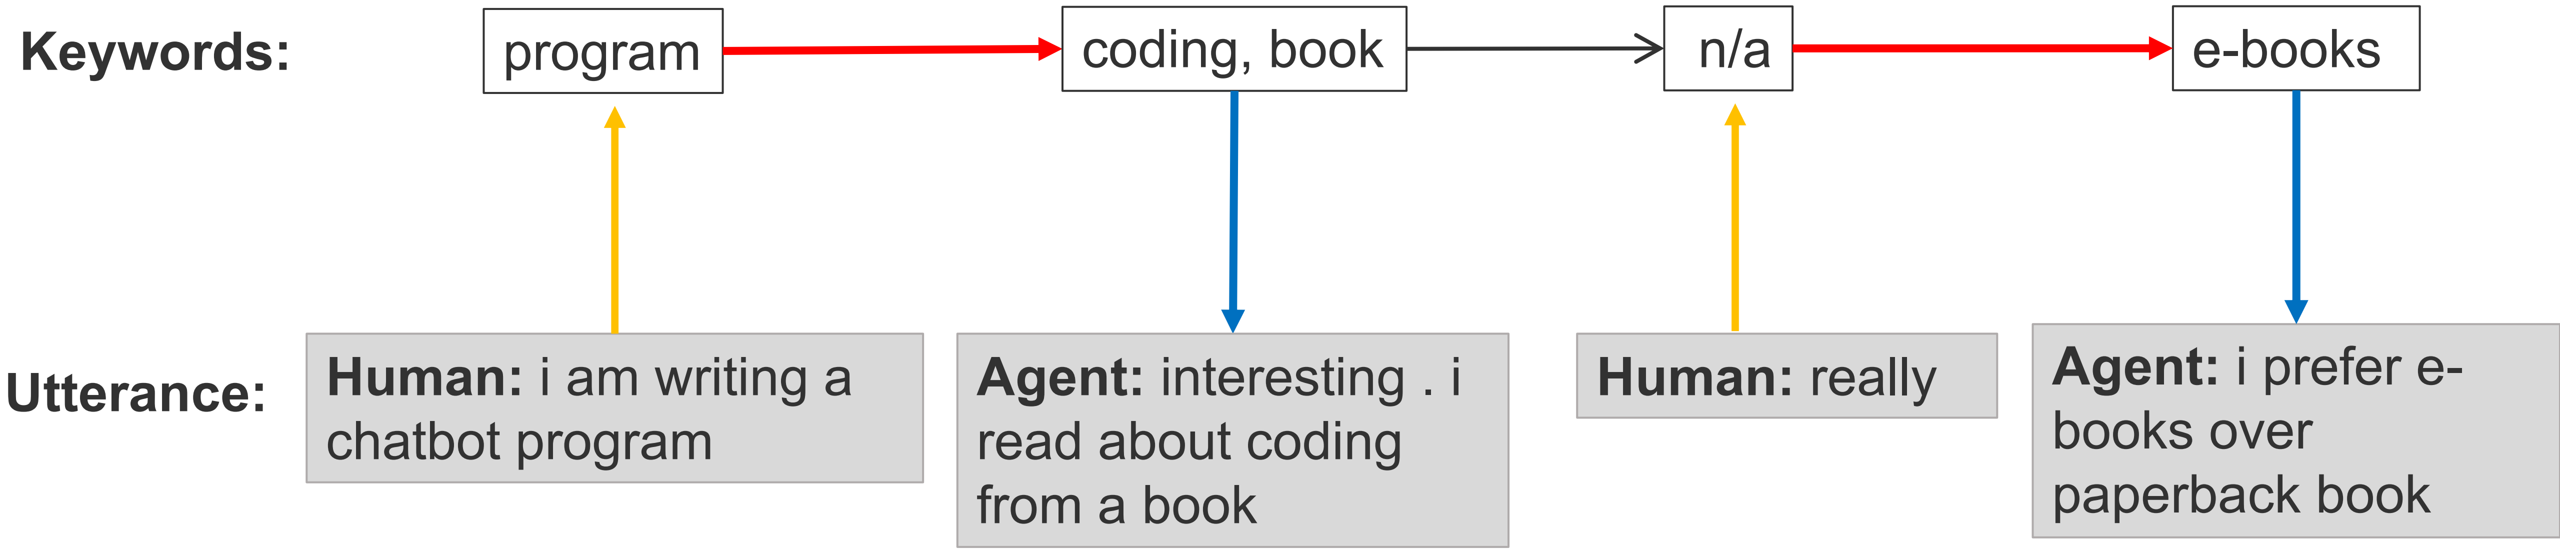

Target-guided Open-domain Conversation

- 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

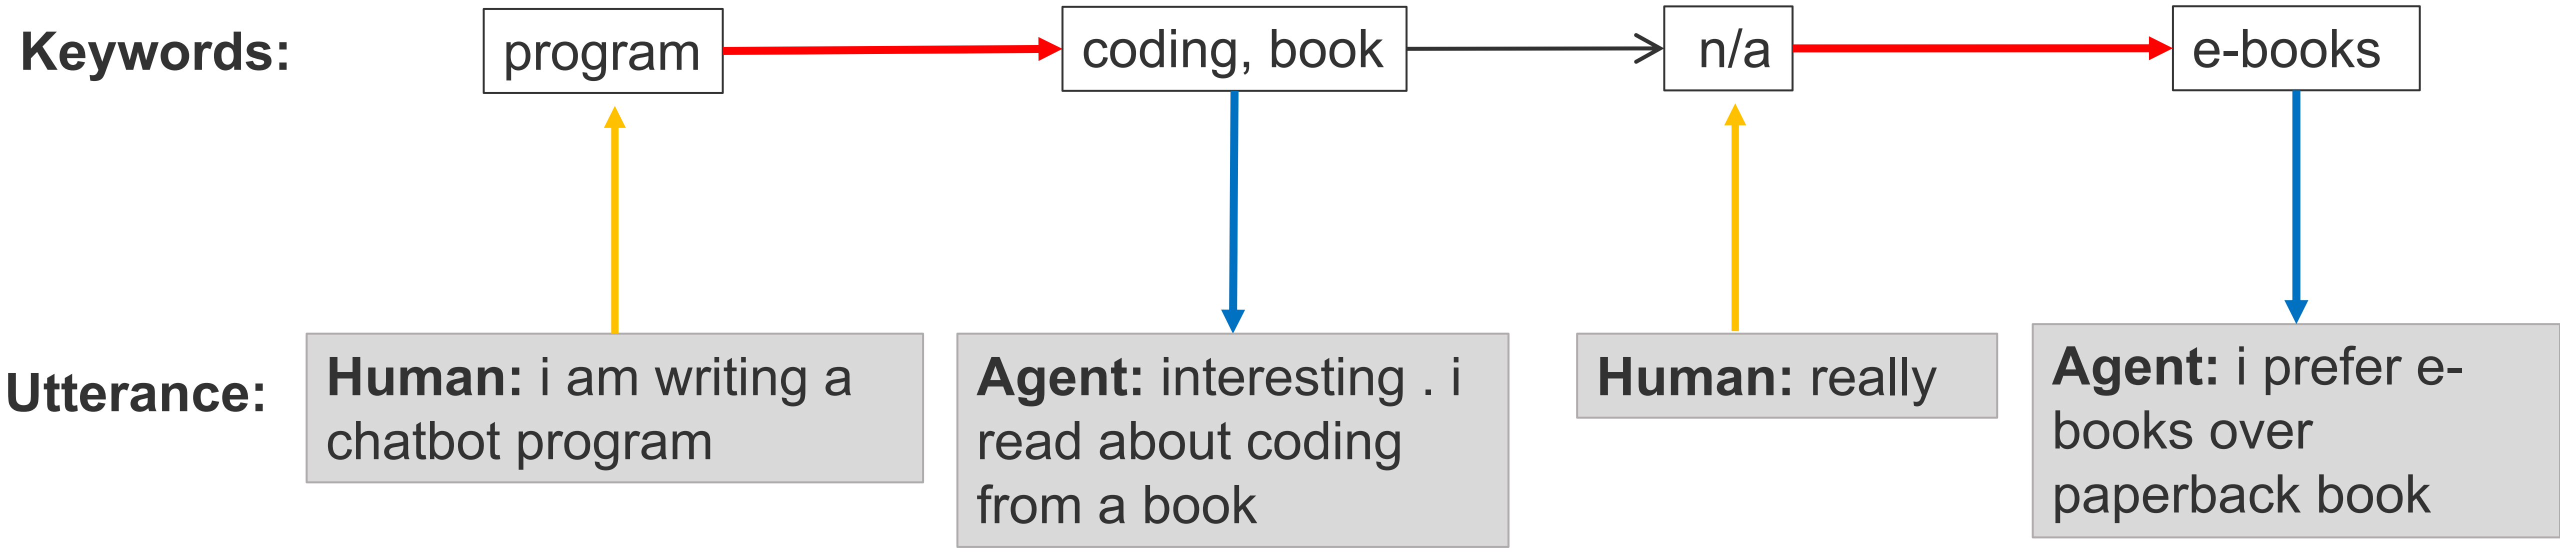

Example Results

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) [Hu et al., 2017; Yang et al., 2018]
  - Text content manipulation [Wang, Hu et al., 2019]
- 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
  - Progressive generation
  - 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?