

Narrative Summarization

Candidacy Exam

Melanie Subbiah - March 10, 2023

1. Narrative Summarization

What is narrative summarization and why is it important?

2. Long Documents

How can we summarize very long narratives?

3. Controllability

How do we target summaries to specific tasks/users?

4. Evaluation

Can we automatically evaluate how well a system works?

1. Narrative Summarization

What is narrative summarization and why is it important?

2. Long Documents

How can we summarize very long narratives?

3. Controllability

How do we target summaries to specific tasks/users?

4. Evaluation

Can we automatically evaluate how well a system works?

What is narrative?

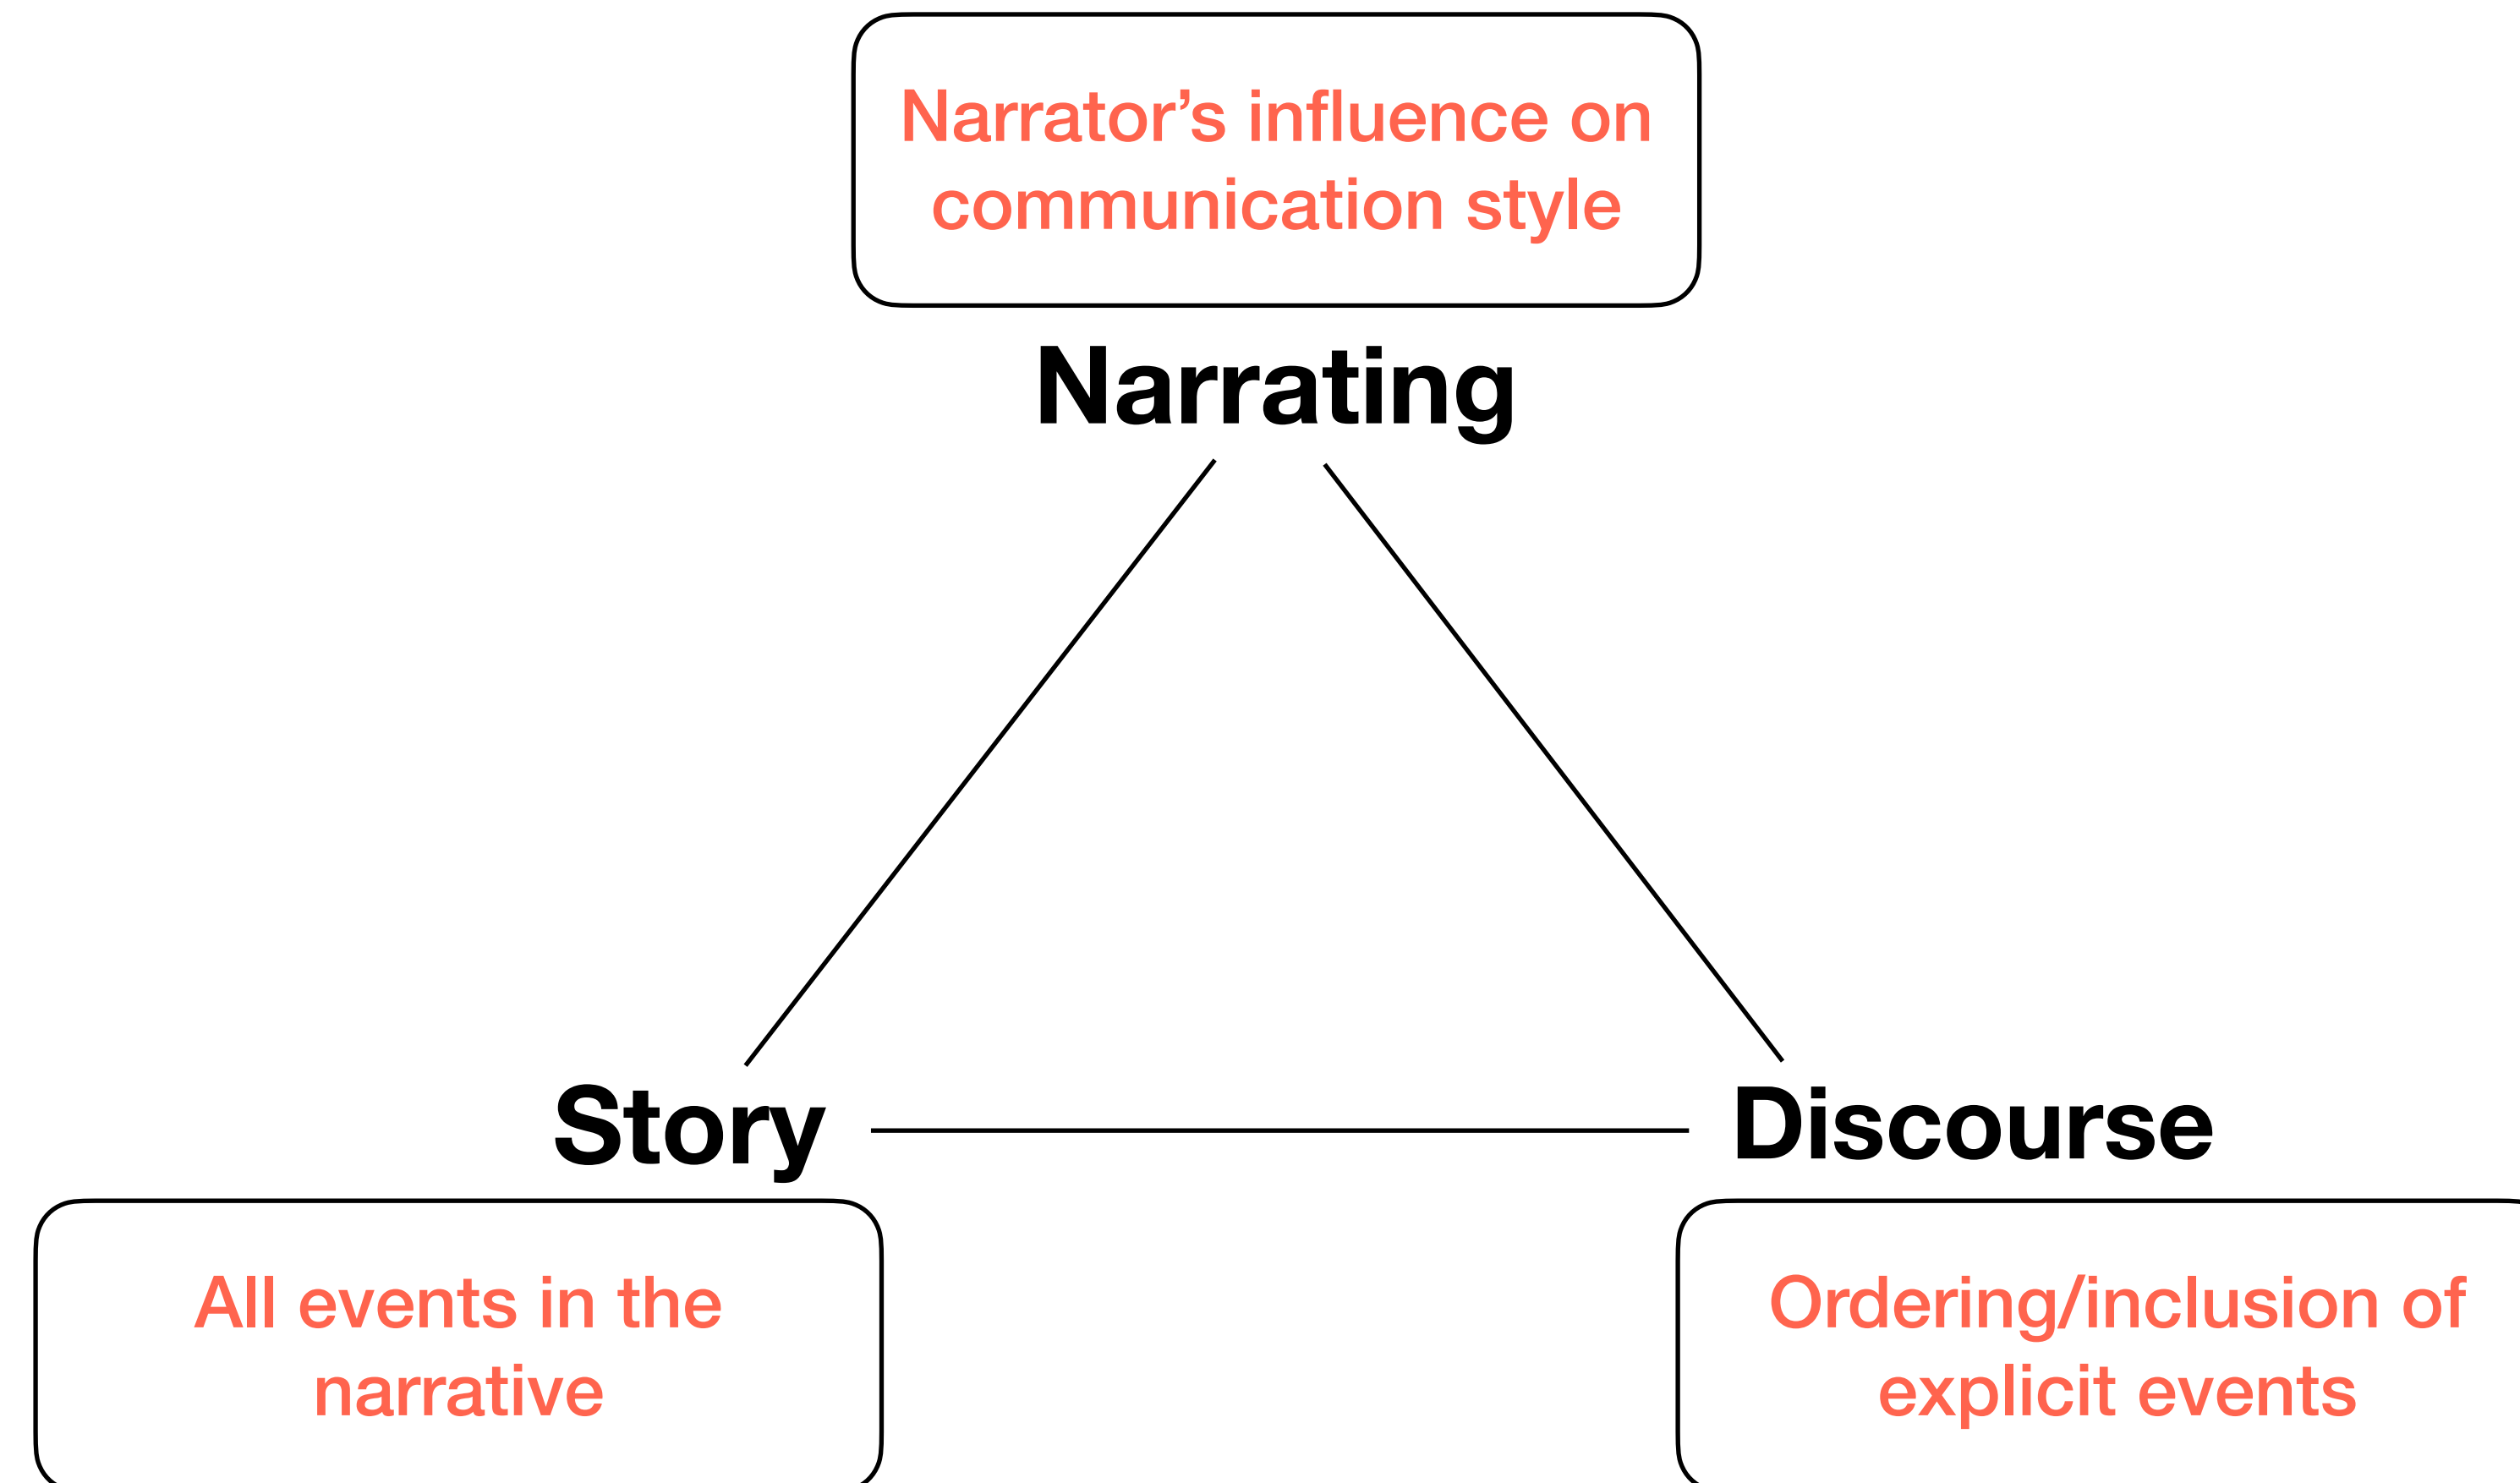
“Narrative roots itself in the **lived, felt experience** of human or human-like agents interacting in an ongoing way with their surrounding environment... [Narrative] is a **basic human strategy** for coming to terms with **time, process, and change.**”

What is narrative?

Different definitions across cultures/traditions with common elements:

- **State change**
- **Sequence of events**
- **Communication context**
- **Human-like experienter**

What is narrative?



Narrative Discourse: An Essay in Method, Genette (1983) - shown in *Narrative Theory for Computational Narrative Understanding*, Piper et al. (2021)

What is narrative?

Huge variation:

- books 
- screenplays 
- online birth stories 
- ... (investigative journalism, interviews, etc.)

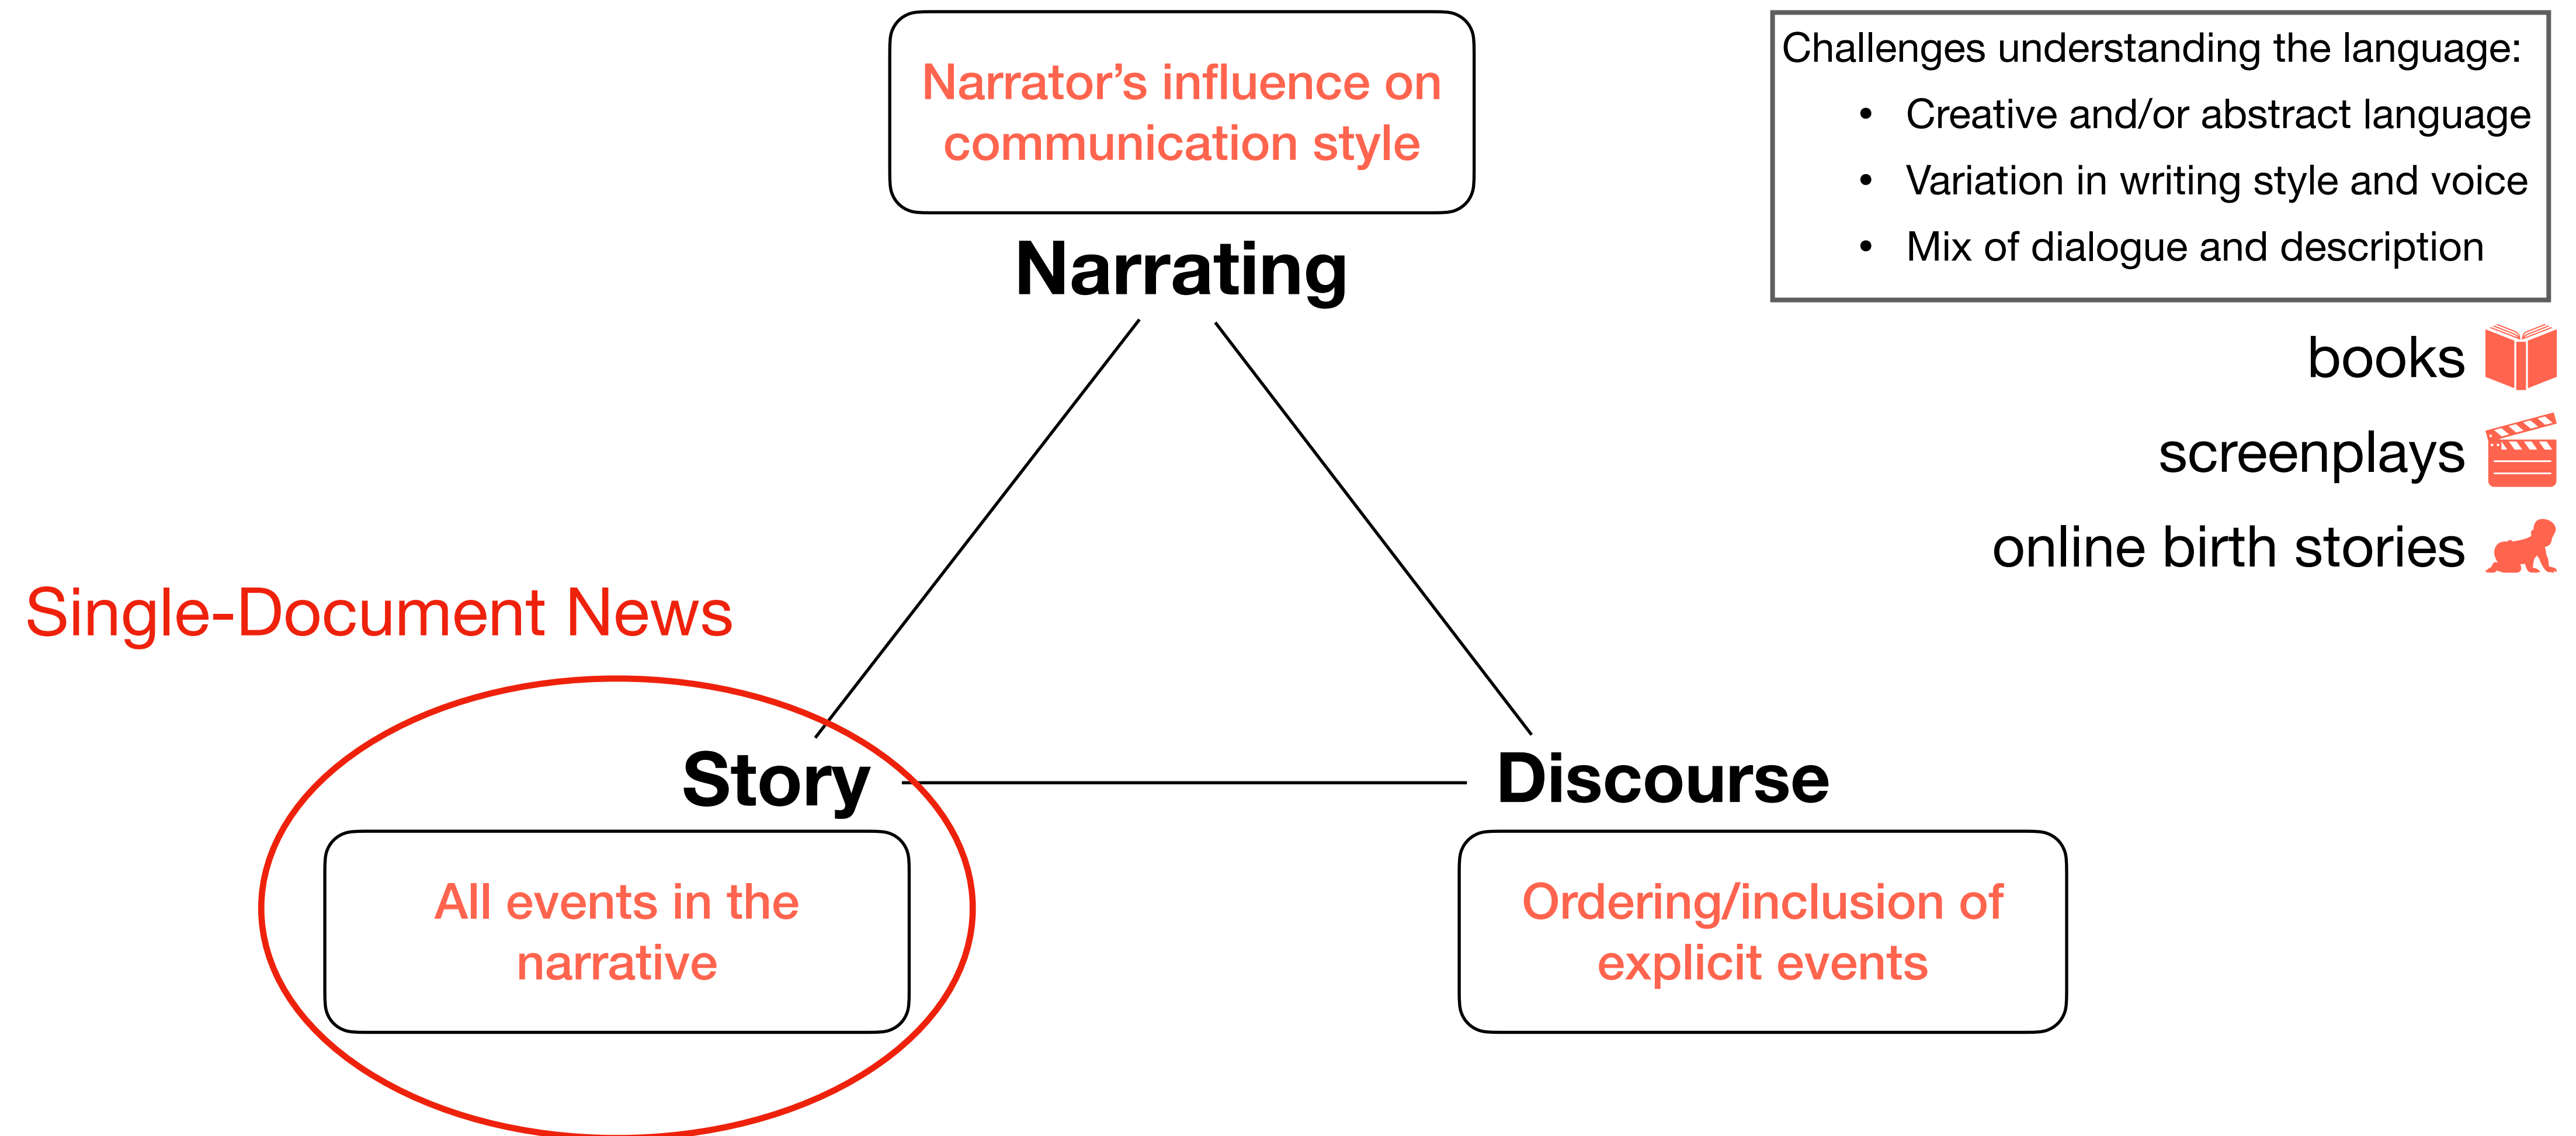
Why is summarizing narrative important?

Human - universal part of communication, teaching, and understanding the world

Practical - narratives can be long and complex → we need summaries

Technical - automatic narrative summarization demonstrates advanced summarization ability

Why is summarizing narrative difficult?



Narrative Discourse: An Essay in Method, Genette (1983) - shown in *Narrative Theory for Computational Narrative Understanding*, Piper et al. (2021)

Why is summarizing narrative difficult?

Unique technical challenges:

- Working with very long narratives **Long Documents**
- Generating context-dependent summaries **Controllability**
- Evaluating summary quality **Evaluation**

Narrative/Summary Variants

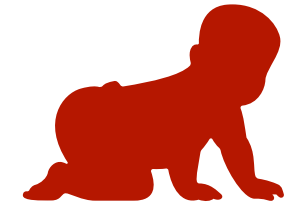
 Online birth stories

 Screenplays

 Books

Narrative/Summary Variants

1. ***Narrative Paths and Negotiation of Power in Birth Stories***, Antoniak et al. (2019)



I finally had my gorgeous baby girl at 41 weeks and 3 days on 3/3/2017!

So I heard from 37 weeks that because of the size of the baby, I probably wouldn't be able to get all the way to 40 weeks and induction might be necessary. Well 39 weeks came and my doctor said that I shouldn't go past 41 weeks.

...

8:30 AM: AT LAST at 41+1 I went to my appointment and was sent to hospital. I go as fast as I can to labor and delivery knowing that they'll have to schedule an induction. Baby was fine in his current spot but the on call OB decided to start an induction.

...

My partner was really relieved that I decided to get an epidural...I think he was getting nervous! Since he was witnessing me in pain. Honestly the epidural wasn't bad, especially comparing the short term pain to the endless contractions.

...

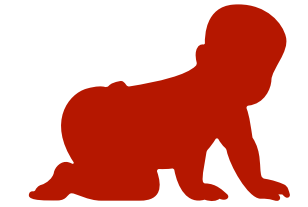
I kept telling the nurse that I felt some pressure, and she was so surprised when she checked me. The nurse said that I was 10cm and ready to push! She went to get the midwife and I began pushing.

...

He scored 9/10 and immediately latched. Breastfeeding wasn't as strange as I expected and it actually came really naturally. All my fears were unnecessary after all.

Narrative/Summary Variants

1. *Narrative Paths and Negotiation of Power in Birth Stories*, Antoniak et al. (2019)



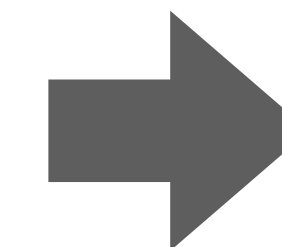
*I finally had my gorgeous baby girl at 41 weeks and 3 days on 3/3/2017!
So I heard from 37 weeks that because of the size of the baby, I probably wouldn't be able to get all the way to 40 weeks and induction might be necessary. Well 39 weeks came and my doctor said that I shouldn't go past 41 weeks.*

...
8:30 AM: AT LAST at 41+1 I went to my appointment and was sent to hospital. I go as fast as I can to labor and delivery knowing that they'll have to schedule an induction. Baby was fine in his current spot but the on call OB decided to start an induction.

...
My partner was really relieved that I decided to get an epidural...I think he was getting nervous! Since he was witnessing me in pain. Honestly the epidural wasn't bad, especially comparing the short term pain to the endless contractions.

...
I kept telling the nurse that I felt some pressure, and she was so surprised when she checked me. The nurse said that I was 10cm and ready to push! She went to get the midwife and I began pushing.

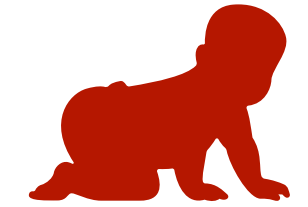
...
He scored 9/10 and immediately latched. Breastfeeding wasn't as strange as I expected and it actually came really naturally. All my fears were unnecessary after all.



Persona	N-Grams
AUTHOR	I, me, myself
We	we, us, ourselves
BABY	baby, son, daughter
DOCTOR	doctor, dr, doc, ob, obgyn, gynecologist, physician
PARTNER	partner, husband, wife
NURSE	nurse
MIDWIFE	midwife
FAMILY	mom, dad, mother, father, brother, sister
ANESTHESIOLOGIST	anesthesiologist
DOULA	doula

Narrative/Summary Variants

1. *Narrative Paths and Negotiation of Power in Birth Stories*, Antoniak et al. (2019)



I finally had my gorgeous baby girl at 41 weeks and 3 days on 3/3/2017!

So I heard from 37 weeks that because of the size of the baby, I probably wouldn't be able to get all the way to 40 weeks and induction might be necessary. Well 39 weeks came and my doctor said that I shouldn't go past 41 weeks.

...

8:30 AM: AT LAST at 41+1 I went to my appointment and was sent to hospital. I go as fast as I can to labor and delivery knowing that they'll have to schedule an induction. Baby was fine in his current spot but the on call OB decided to start an induction.

...

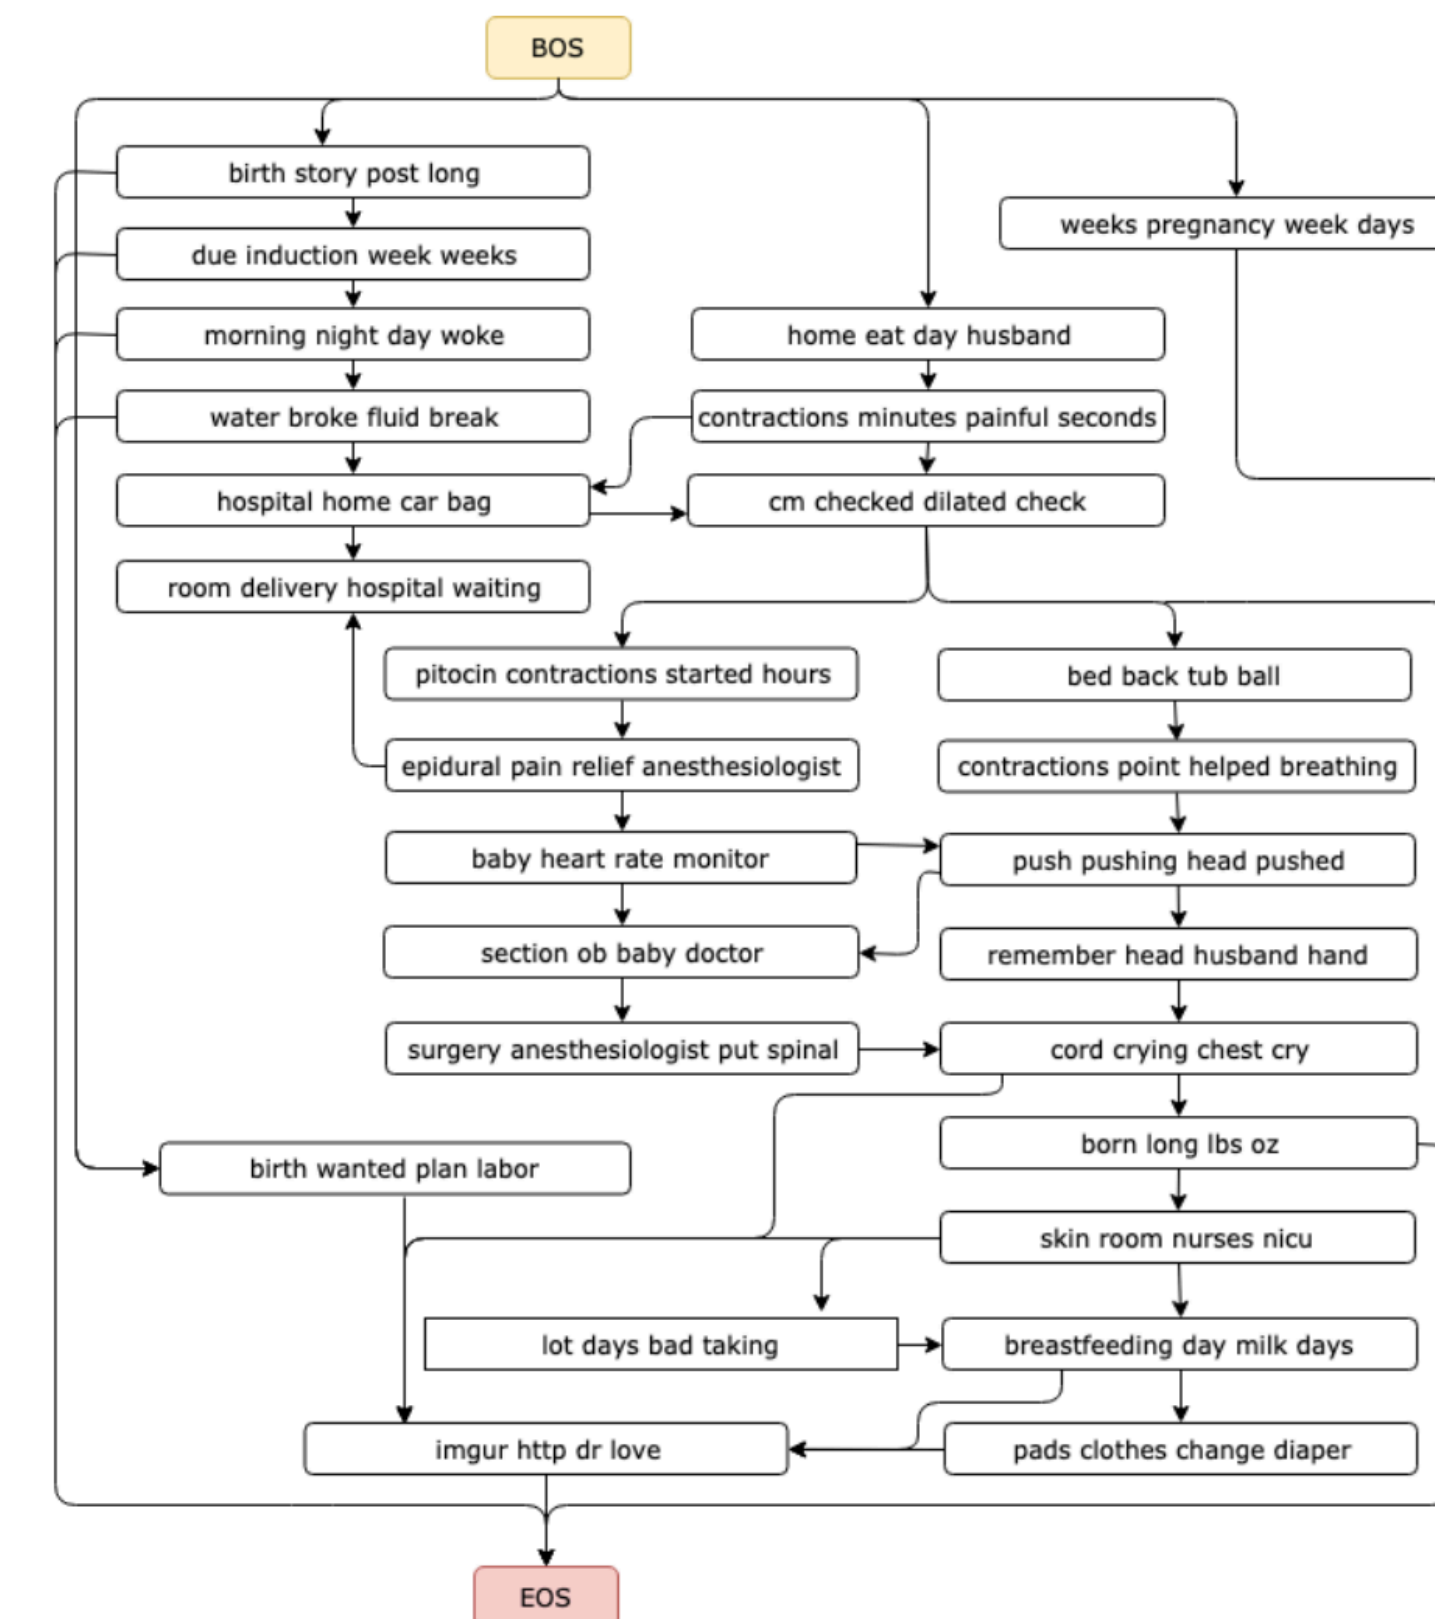
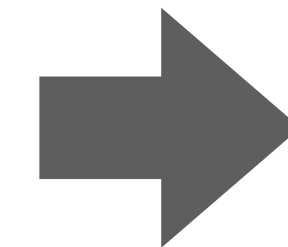
My partner was really relieved that I decided to get an epidural...I think he was getting nervous! Since he was witnessing me in pain. Honestly the epidural wasn't bad, especially comparing the short term pain to the endless contractions.

...

I kept telling the nurse that I felt some pressure, and she was so surprised when she checked me. The nurse said that I was 10cm and ready to push! She went to get the midwife and I began pushing.

...

He scored 9/10 and immediately latched. Breastfeeding wasn't as strange as I expected and it actually came really naturally. All my fears were unnecessary after all.

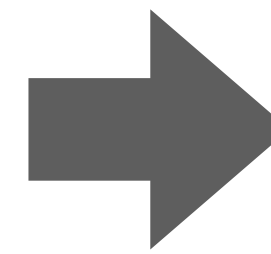
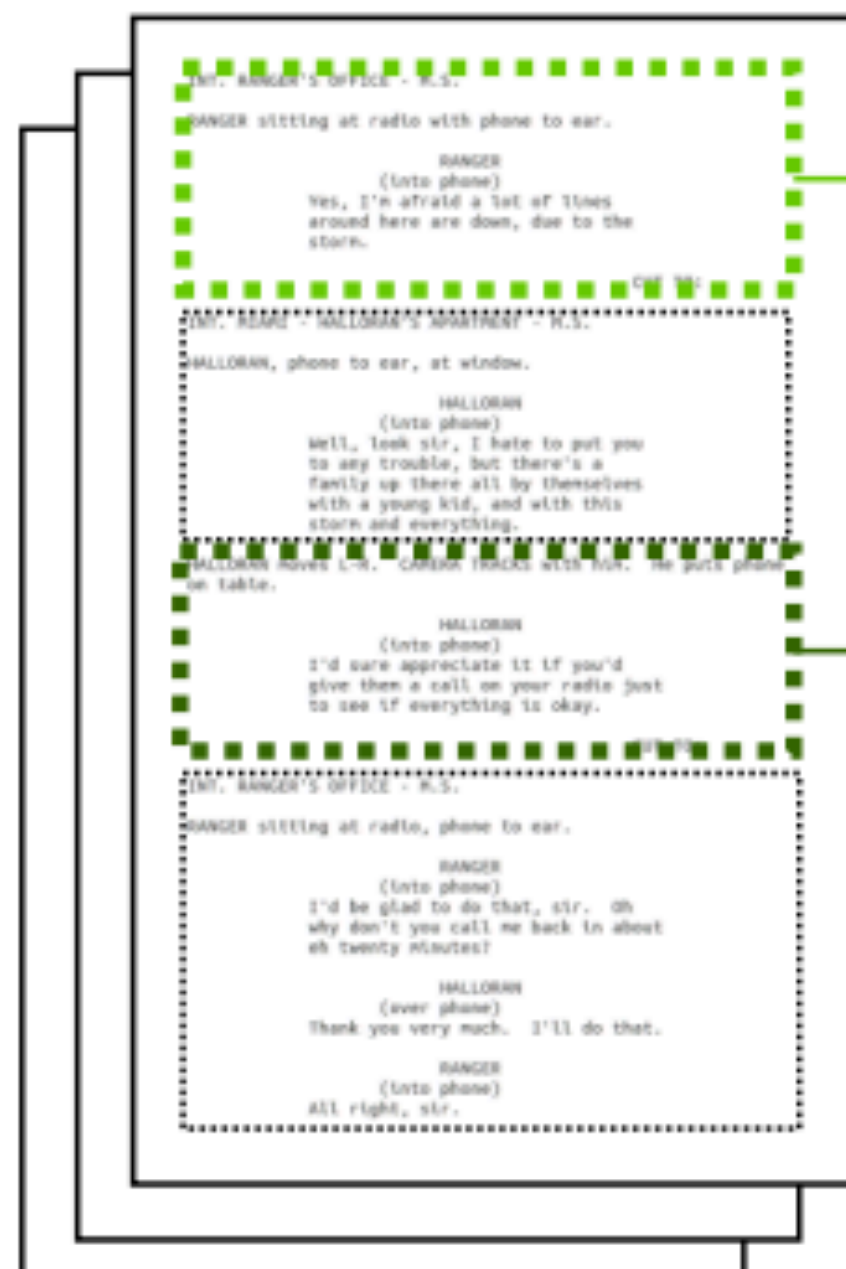


Narrative/Summary Variants

1. *Narrative Paths and Negotiation of Power in Birth Stories*, Antoniak et al. (2019) 

2. **Screenplay Summarization Using Latent Narrative Structure**, Papalampidi et al. (2020) 


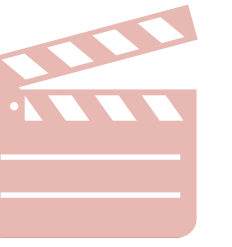

Screenplay



Video summary



Narrative/Summary Variants

1. *Narrative Paths and Negotiation of Power in Birth Stories*, Antoniak et al. (2019) 
2. *Screenplay Summarization Using Latent Narrative Structure*, Papalampidi et al. (2020) 
3. ***SummScreen: A Dataset for Abstractive Screenplay Summarization***, Chen et al. (2022) 

Transcript:

[The apartment]

Sheldon : What color would you like to be ?

Leonard : Well , I 'd like to be green , but you know you always take it .

Sheldon : That 's not true . Any color 's fine with me . Yeah , I could be a -
a combination of blue and yellow .

Leonard : Blue and yellow make green .

Sheldon : Well , then it 's settled .

Penny : Hi . Ready to go ?

Sheldon : Oh , good news , we ordered lunch , so we can all stay here and
play **Lord of the Rings Risk** .

Amy : Sheldon , we said that we would play games with you tonight .

Sheldon : Oh , no , we 'll still be playing it tonight , **this game** can easily
take eight hours .

Penny : Sweetie , you really thought I 'd want to do this ?

Leonard : No .

Penny : Well , did you tell him that ?

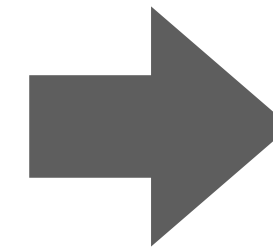
Leonard : Yes .

Penny : Did you say it out loud with words ?

Leonard : No .

Penny : I do n't want to spend the whole day playing **a board game** .

...



Recap:

Sheldon and Leonard are happy playing **a board game** until Amy and Penny say they are tired of doing what the guys want ...

Narrative/Summary Variants - Books

Datasets

Paper	Input Data	Reference Summaries	Input Length
Exploring Content Selection in Summarization of Novel Chapters , Ladhak et al. (2020)	Project Gutenberg novels	Study guide websites	Chapter
BookSum: A Collection of Datasets for Long-form Narrative Summarization , Kryściński et al. (2021)	Project Gutenberg stories, plays and novels	Study guide websites	Paragraph, chapter, and book
Recursively Summarizing Books with Human Feedback , Wu et al. (2021)	GPT-3 Books1 and Books2	Human-written	Book

Narrative/Summary Variants - Books

Datasets

Paper	Input Data	Reference Summaries	Input Length
Exploring Content Selection in Summarization of Novel Chapters , Ladhak et al. (2020)	Project Gutenberg novels	Study guide websites	Chapter
BookSum: A Collection of Datasets for Long-form Narrative Summarization , Kryściński et al. (2021)	Project Gutenberg stories, plays and novels	Study guide websites	Paragraph, chapter, and book
Recursively Summarizing Books with Human Feedback , Wu et al. (2021)	GPT-3 Books1 and Books2	Human-written	Book

Narrative/Summary Variants - Books

Datasets

Paper	Input Data	Reference Summaries	Input Length
Exploring Content Selection in Summarization of Novel Chapters , Ladhak et al. (2020)	Project Gutenberg novels	Study guide websites	Chapter
BookSum: A Collection of Datasets for Long-form Narrative Summarization , Kryściński et al. (2021)	Project Gutenberg stories, plays and novels	Study guide websites	Paragraph, chapter, and book
Recursively Summarizing Books with Human Feedback , Wu et al. (2021)	GPT-3 Books1 and Books2	Human-written	Book

Narrative/Summary Variants - Books

Note:

Free unpublished or past-copyright books

Datasets

Paper	Input Data	Reference Summaries	Input Length
Exploring Content Selection in Summarization of Novel Chapters , Ladhak et al. (2020)	Project Gutenberg novels	Study guide websites	Chapter
BookSum: A Collection of Datasets for Long-form Narrative Summarization , Kryściński et al. (2021)	Project Gutenberg stories, plays and novels	Study guide websites	Paragraph, chapter, and book
Recursively Summarizing Books with Human Feedback , Wu et al. (2021)	GPT-3 Books1 and Books2	Human-written	Book

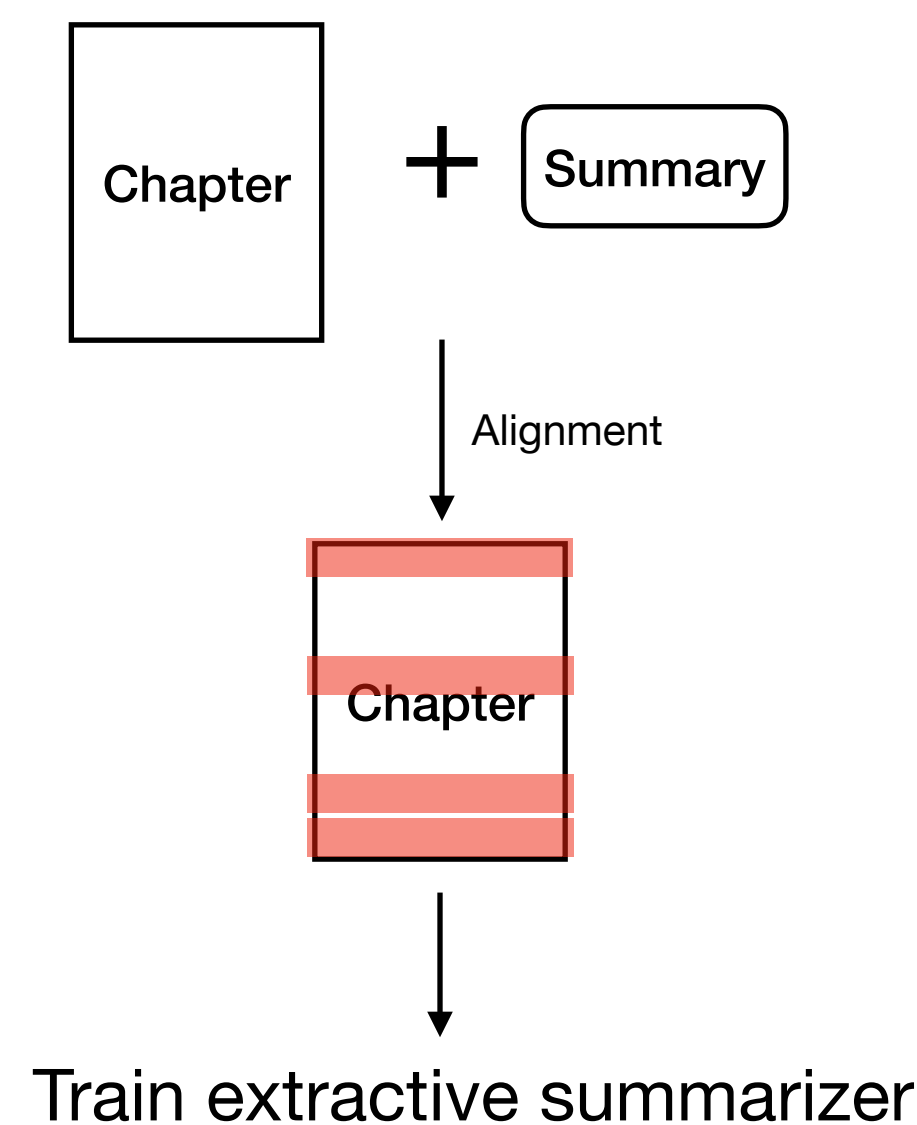
Narrative/Summary Variants - Books

Methods

Exploring Content Selection in Summarization of Novel Chapters,
Ladhak et al. (2020)

BookSum: A Collection of Datasets for Long-form Narrative Summarization,
Kryściński et al. (2021)

Recursively Summarizing Books with Human Feedback,
Wu et al. (2021)

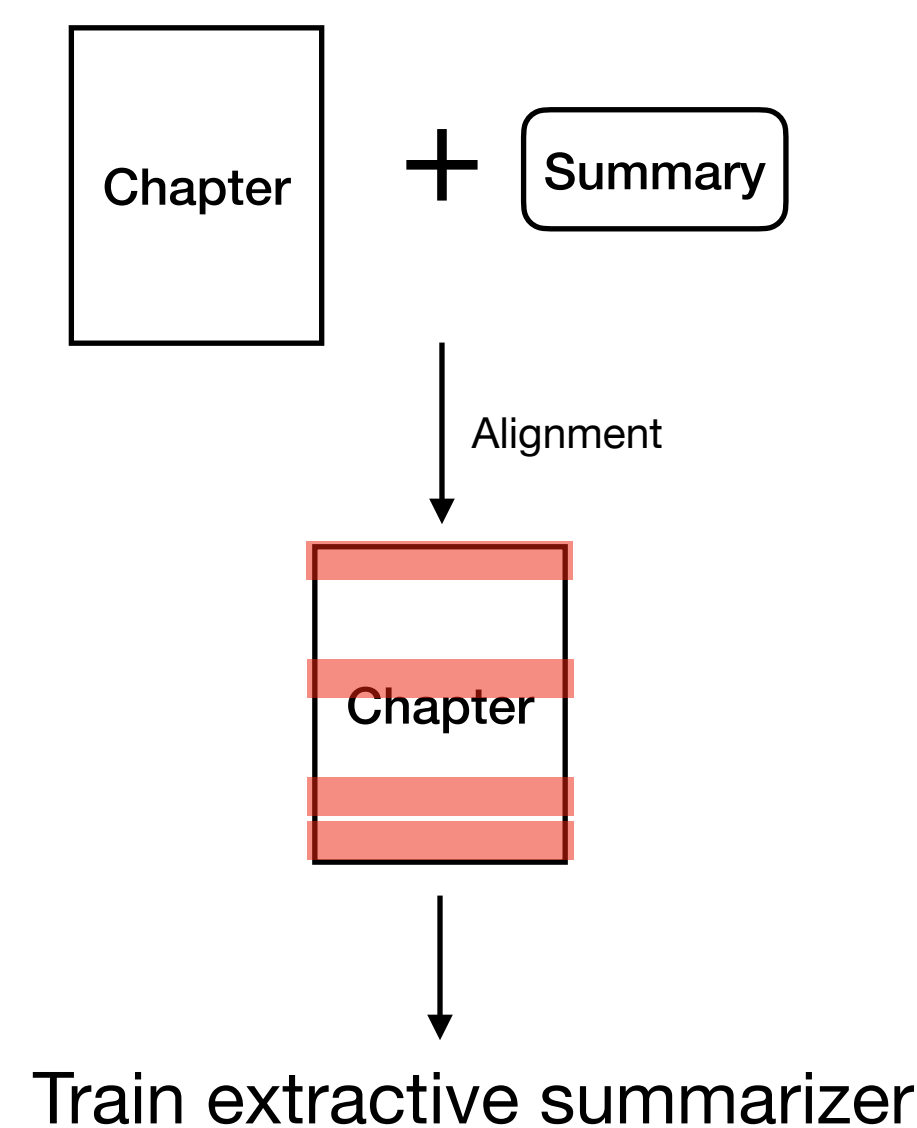


Extractive

Narrative/Summary Variants - Books

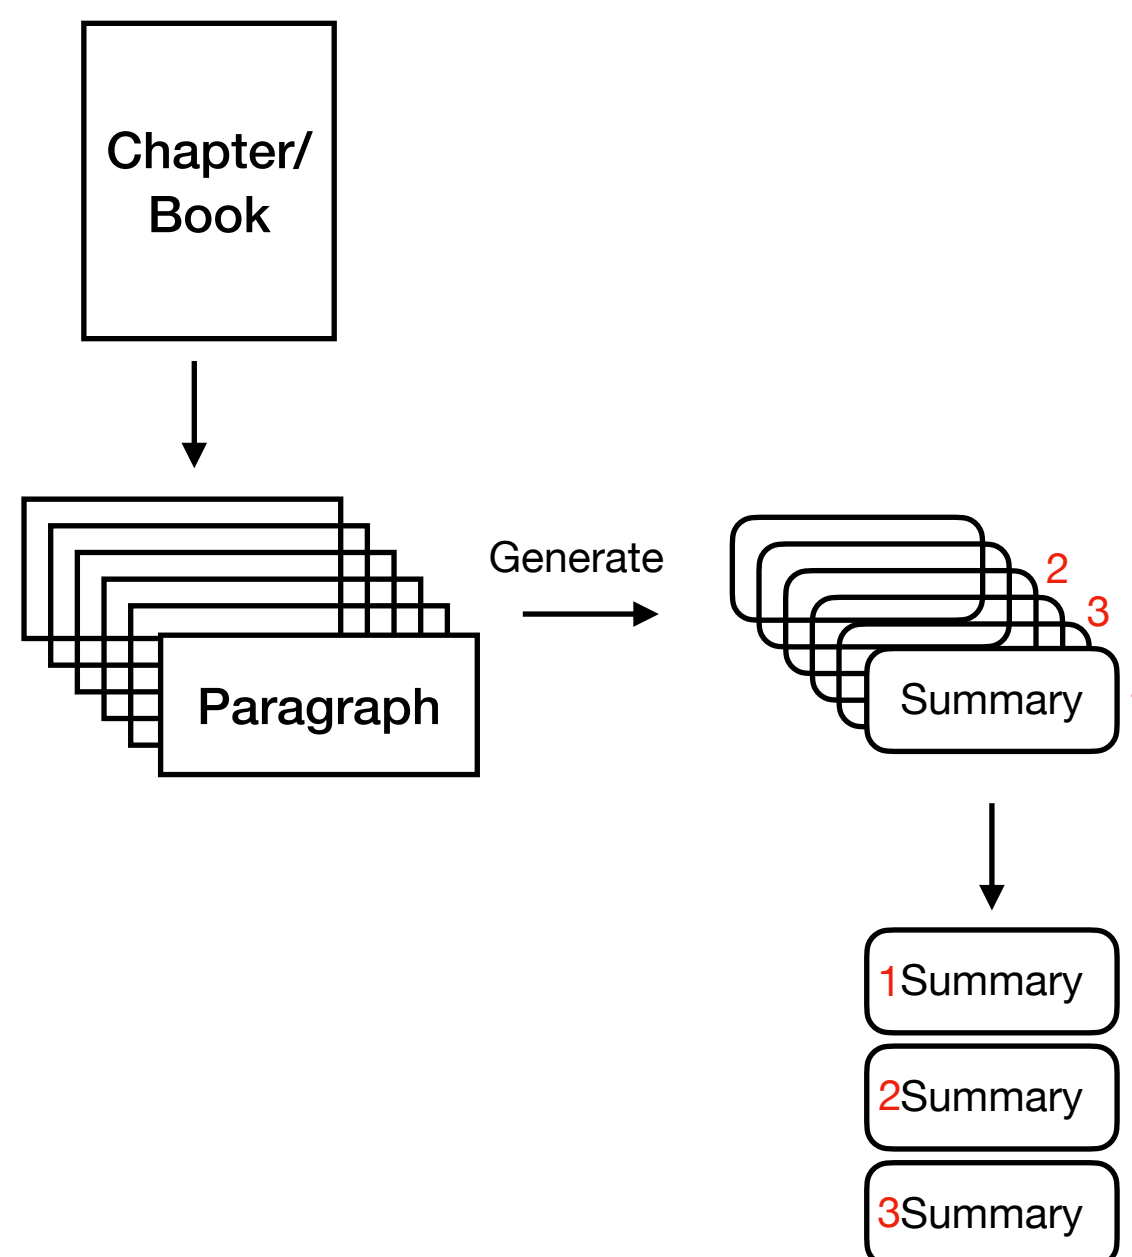
Methods

Exploring Content Selection in Summarization of Novel Chapters,
Ladhak et al. (2020)



Extractive

BookSum: A Collection of Datasets for Long-form Narrative Summarization,
Kryściński et al. (2021)



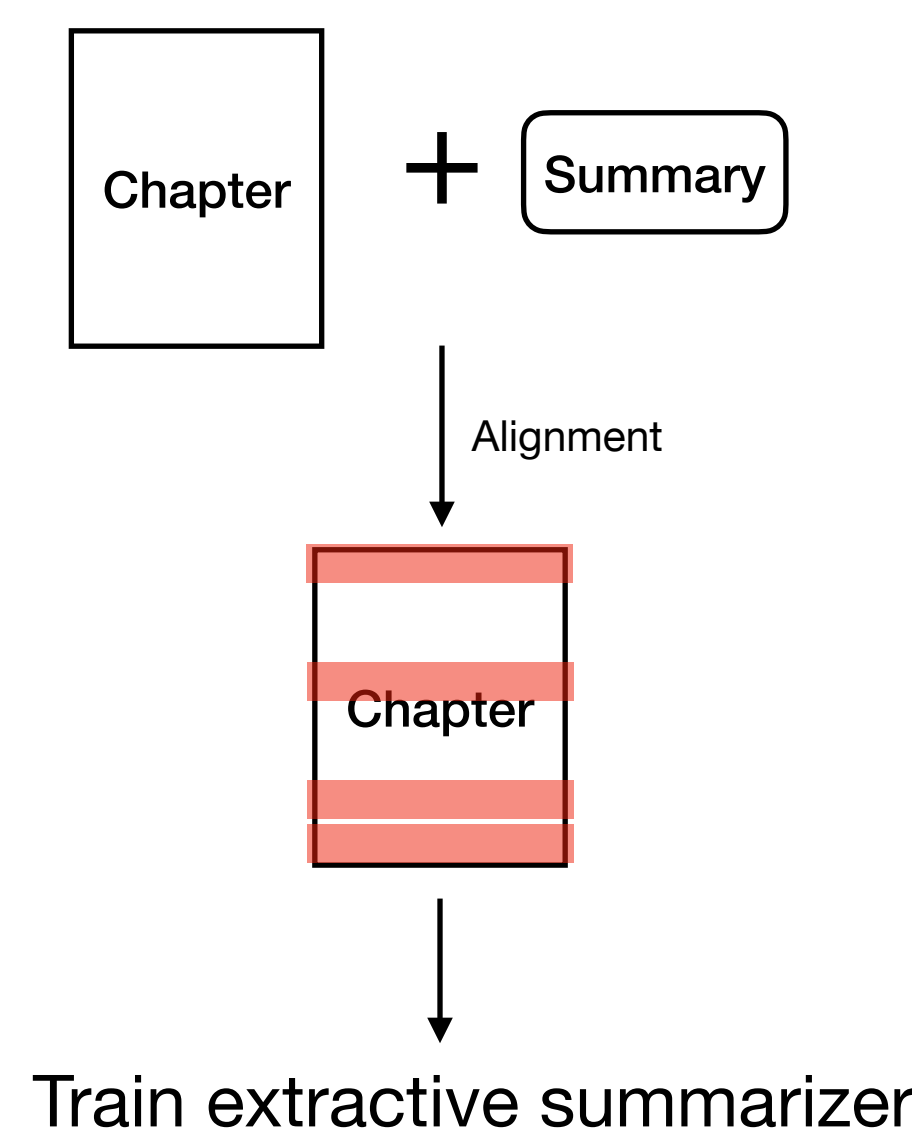
Extractive + Abstractive

Recursively Summarizing Books with Human Feedback,
Wu et al. (2021)

Narrative/Summary Variants - Books

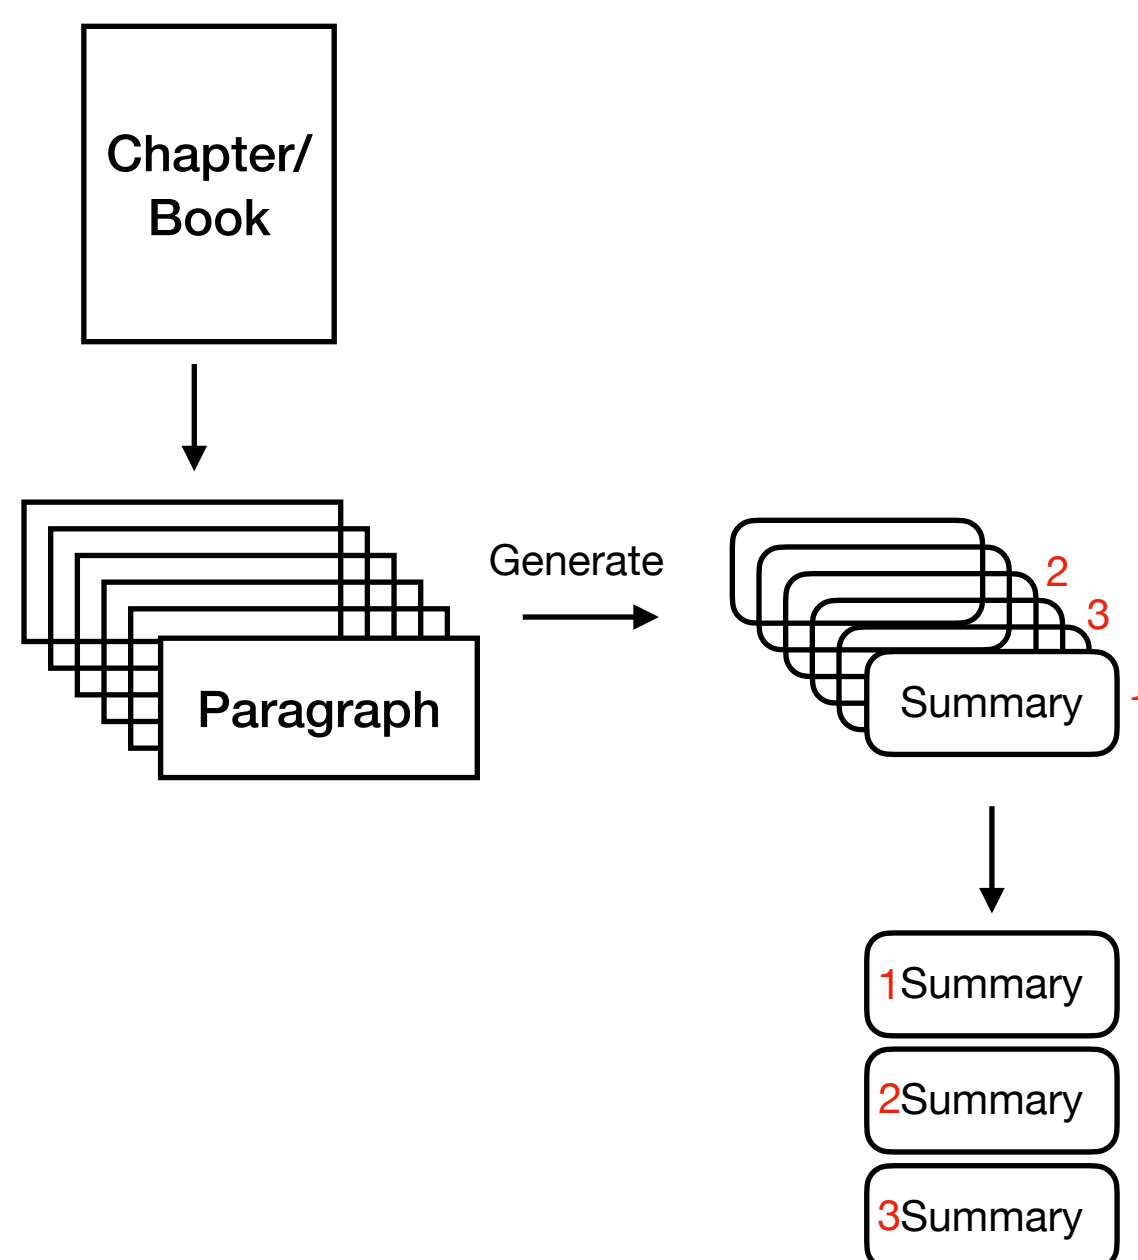
Methods

Exploring Content Selection in Summarization of Novel Chapters,
Ladhak et al. (2020)



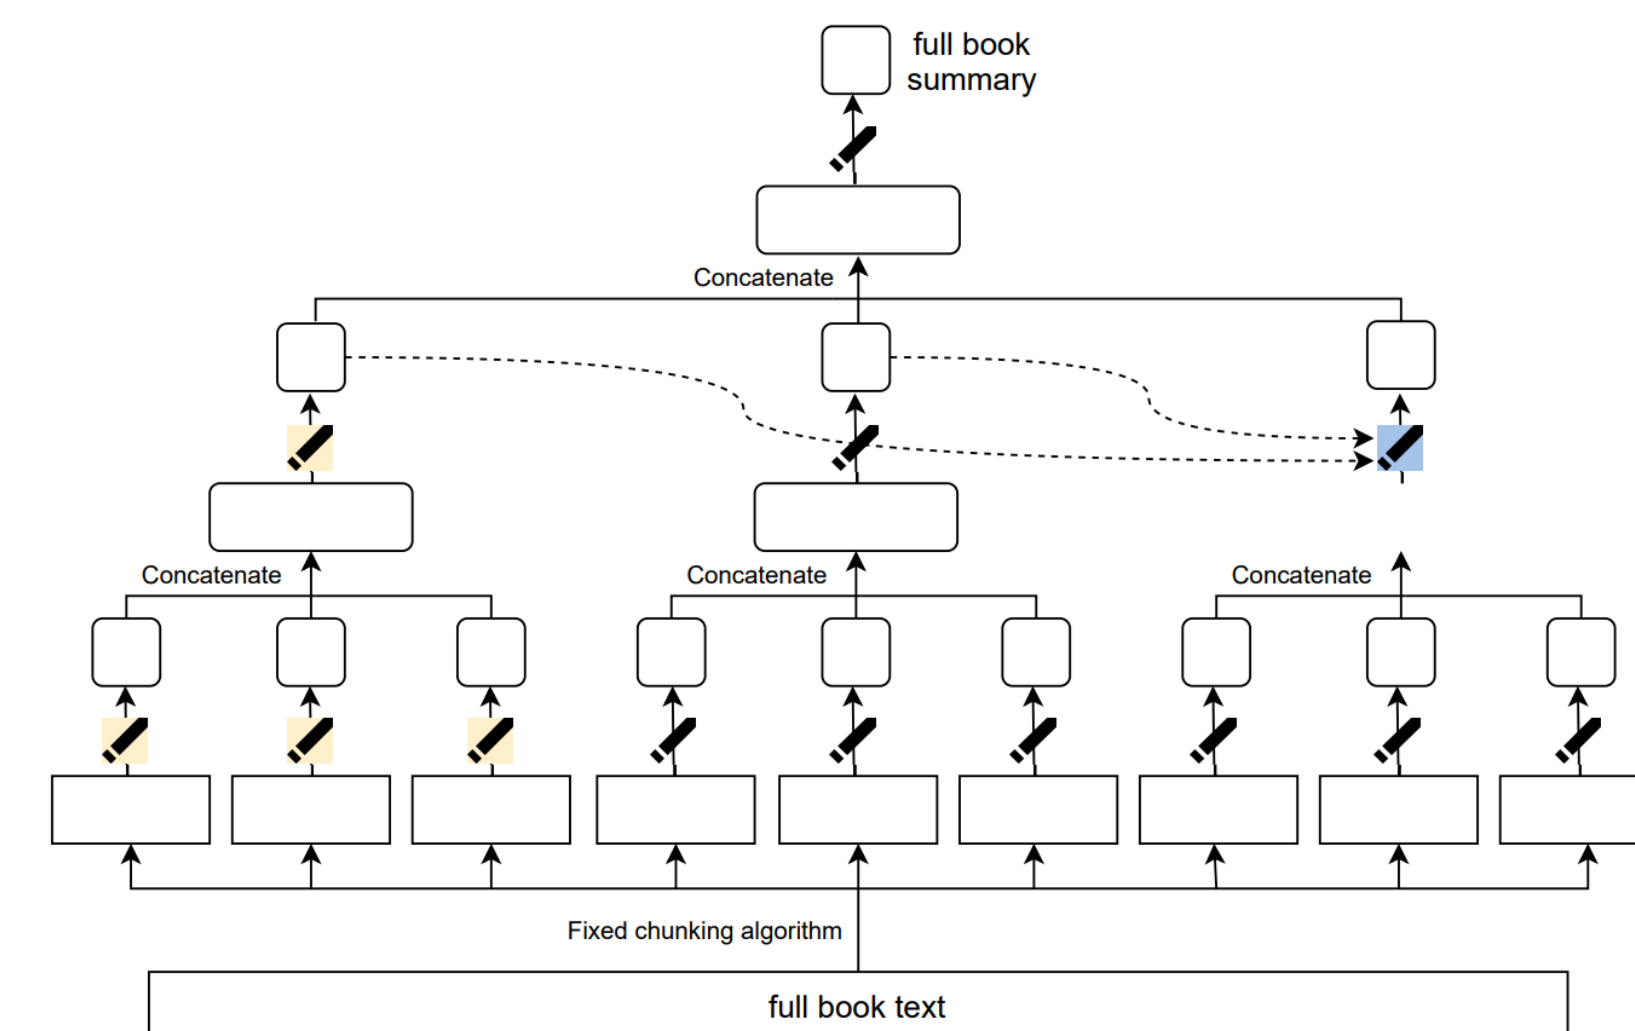
Extractive

BookSum: A Collection of Datasets for Long-form Narrative Summarization,
Kryściński et al. (2021)



Extractive + Abstractive

Recursively Summarizing Books with Human Feedback,
Wu et al. (2021)



Abstractive

Narrative/Summary Variants - Books

Results

	Abstractive	ROUGE-1	ROUGE-2	ROUGE-L	BERTScore
Extractive Oracle		46.62	9.17	18.31	0.082
BertExt		36.71	6.16	13.40	0.028
BookSum , Kryściński et al. (2021) T5 zero-shot	✓	35.43	5.62	12.02	0.011
T5 fine-tuned	✓	39.46	7.69	13.77	0.060
Recursively Summarizing Books with Human Feedback , Wu et al. (2021) 175b full tree RL	✓	41.51	10.46	16.88	0.1821
175b first subtree RL	✓	43.19	10.63	17.10	0.1778
6b full tree RL	✓	36.79	7.22	14.84	0.1246

Narrative/Summary Variants - Books

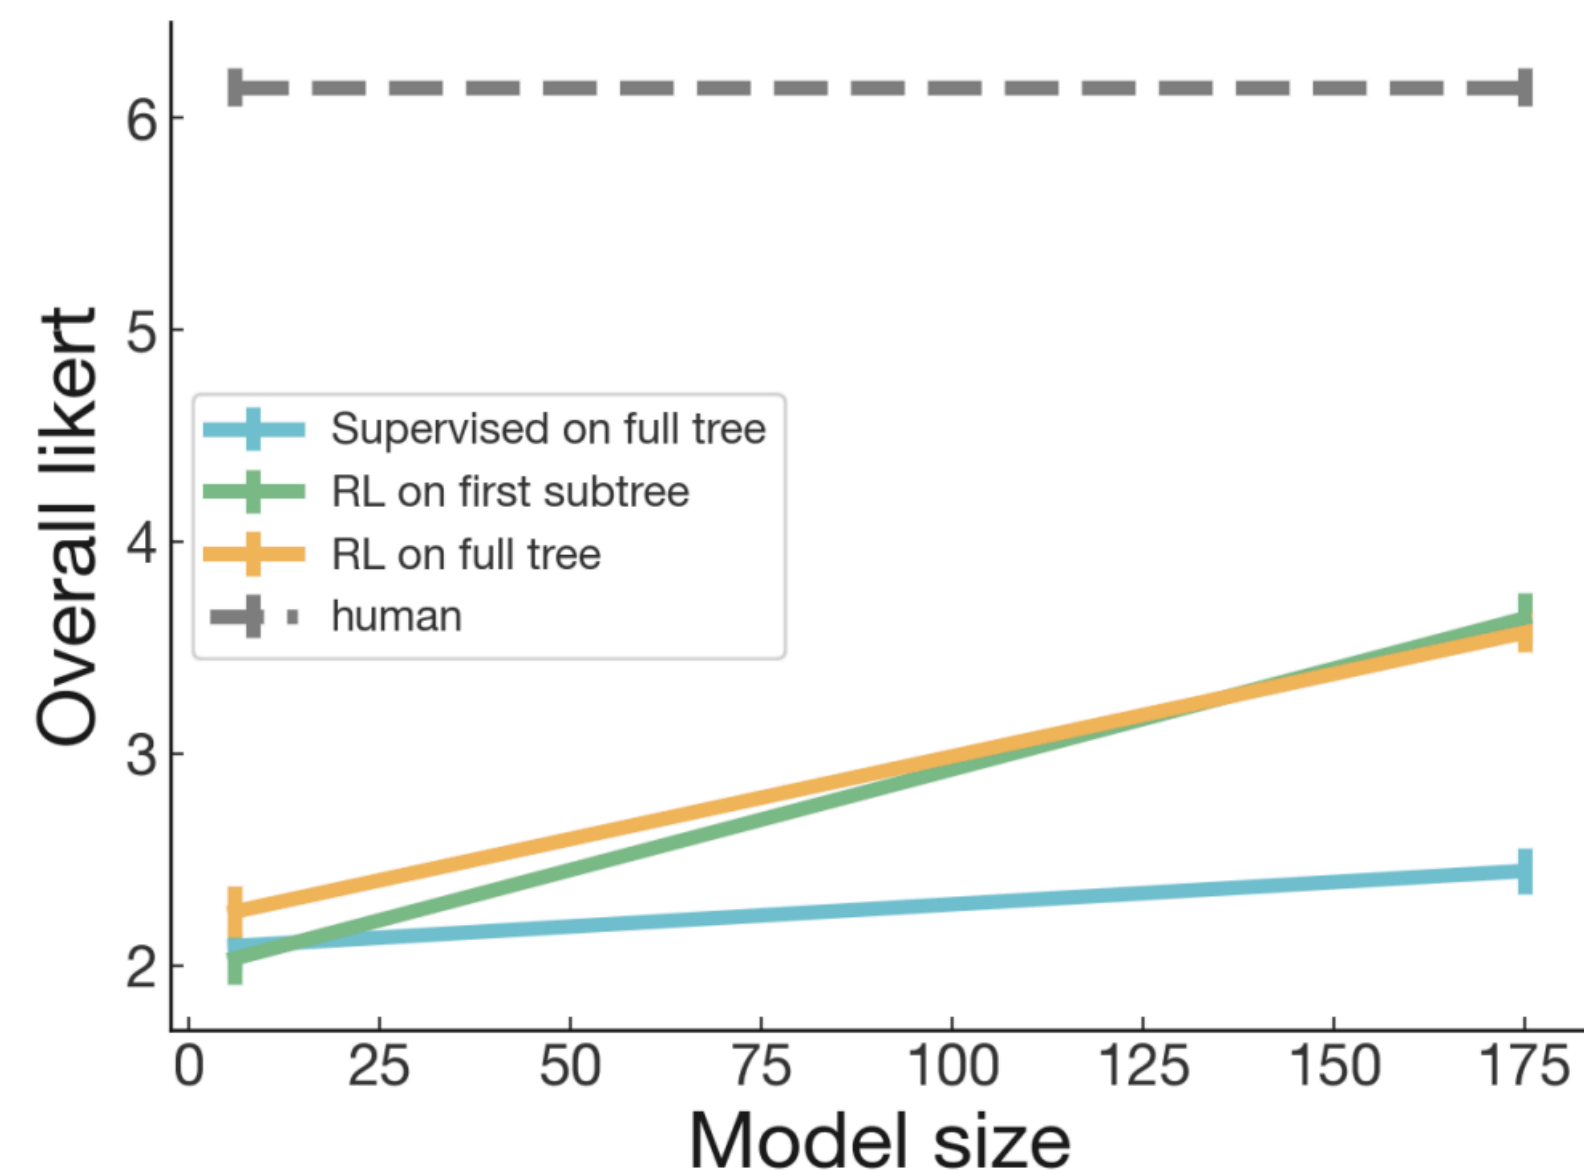
Results

	Abstractive	ROUGE-1	ROUGE-2	ROUGE-L	BERTScore
Extractive Oracle		46.62	9.17	18.31	0.082
BertExt		36.71	6.16	13.40	0.028
BookSum , Kryściński et al. (2021)	T5 zero-shot ✓	35.43	5.62	12.02	0.011
	T5 fine-tuned ✓	39.46	7.69	13.77	0.060
Recursively Summarizing Books with Human Feedback , Wu et al. (2021)	175b full tree RL ✓	41.51	10.46	16.88	0.1821
	175b first subtree RL ✓	43.19	10.63	17.10	0.1778
	6b full tree RL ✓	36.79	7.22	14.84	0.1246

Narrative/Summary Variants - Books

Results

Recursively Summarizing Books with Human Feedback, Wu et al. (2021)



Summary of “A Promised Land” by Barack Obama

Best 175B RL summary

Barack Obama wins the presidency in 2008, and his first term is marked by economic turmoil and political polarization. His domestic policy includes the Affordable Care Act, which reformed healthcare in the United States.

Obama’s presidency includes terrorism, the Arab-Israeli conflict, and the Deepwater Horizon oil spill. His foreign policy includes finding and eliminating Osama bin Laden. After the successful Abbottabad raid, Obama reflects on the unity of effort displayed by all the individuals involved in the operation. He wonders if that same unity of effort could be applied to other areas of government.

Need for: **Controllability** **Evaluation**

Summary - It's a good time to work on narrative summarization!

Interesting narrative datasets across different domains, data formats, and summarization styles:

Birth stories, screenplays, books, etc.



Summary - It's a good time to work on narrative summarization!

Interesting narrative datasets across different domains, data formats, and summarization styles:

Birth stories, screenplays, books, etc.



Early **attempts at abstractive summarization of long narratives** (e.g. books)...

- LLMs make this more possible but lots of room for improvement!
- Limited by challenges with processing long documents, unfaithful summarization, and identifying what details to include/exclude

1. Narrative Summarization

What is narrative summarization and why is it important?

2. Long Documents

How can we summarize very long narratives?

3. Controllability

How do we target summaries to specific tasks/users?

4. Evaluation

Can we automatically evaluate how well a system works?

Long Document Summarization

1. Efficient attention for Transformer-based models

Address computational bottleneck to scale to longer inputs

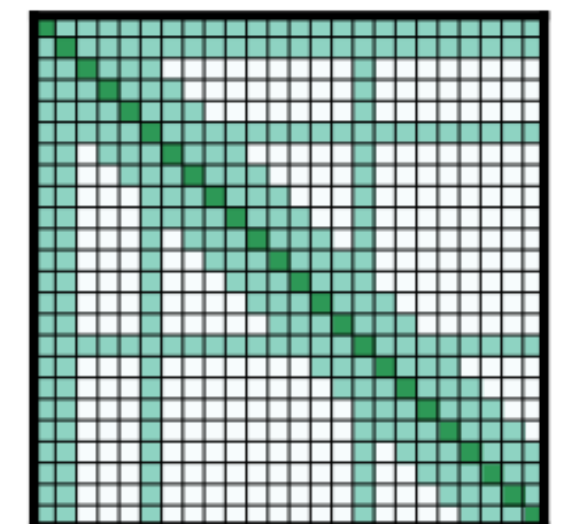
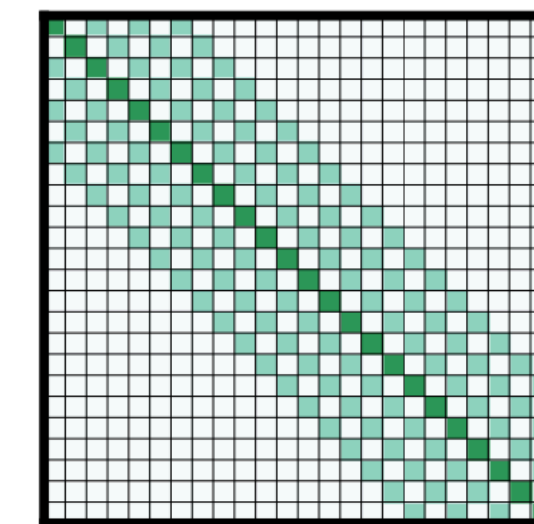
2. Summarization-specific approaches

Computational Bottleneck in Transformers

Full self-attention has **$O(n^2)$** memory/compute complexity in terms of input length

Efficient Attention Operations

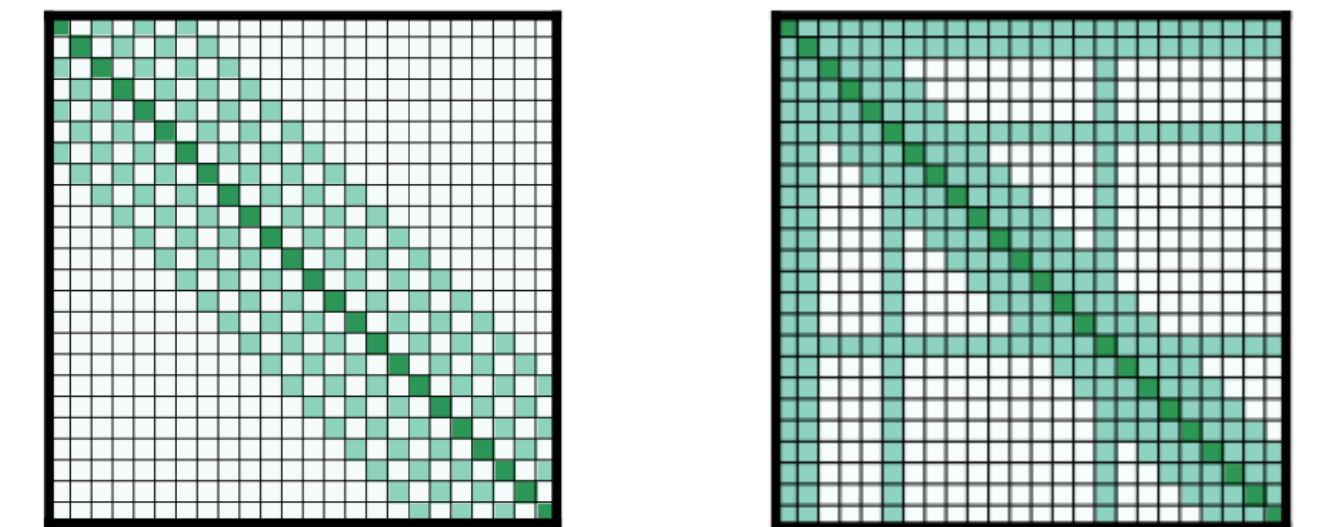
Paper	Model Context Length	Complexity	Strategy
Longformer: The Long-Document Transformer , Beltagy et al. (2020)	23,040	$O(nw)$ <small>w = window size</small>	Sliding window attention
Sparse Sinkhorn Attention , Tay et al. (2020)	2048	$O(2nb)$ <small>b = number of tokens to attend in each block</small>	Learnable sparse attention pattern
Efficient Attentions for Long Document Summarization , Huang et al. (2021)	10,000	$O(mn/s)$ <small>m = output length s = stride size</small>	Head-wise positional strides



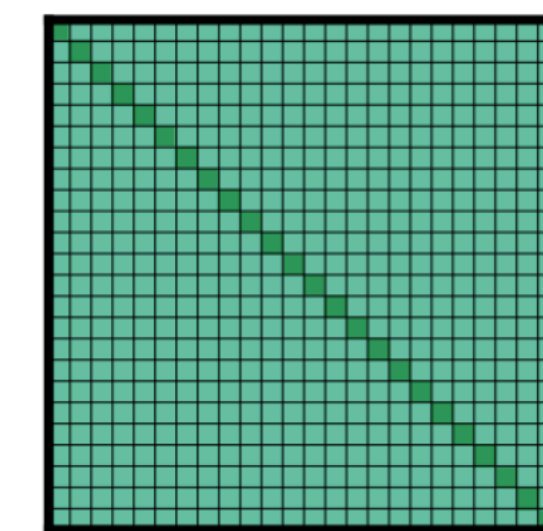
Dilated sliding window

Efficient Attention Operations

Paper	Model Context Length	Complexity	Strategy
Longformer: The Long-Document Transformer , Beltagy et al. (2020)	23,040	$O(nw)$ <small>w = window size</small>	Sliding window attention
Sparse Sinkhorn Attention , Tay et al. (2020)	2048	$O(2nb)$ <small>b = number of tokens to attend in each block</small>	Learnable sparse attention pattern
Efficient Attentions for Long Document Summarization , Huang et al. (2021)	10,000	$O(mn/s)$ <small>m = output length s = stride size</small>	Head-wise positional strides

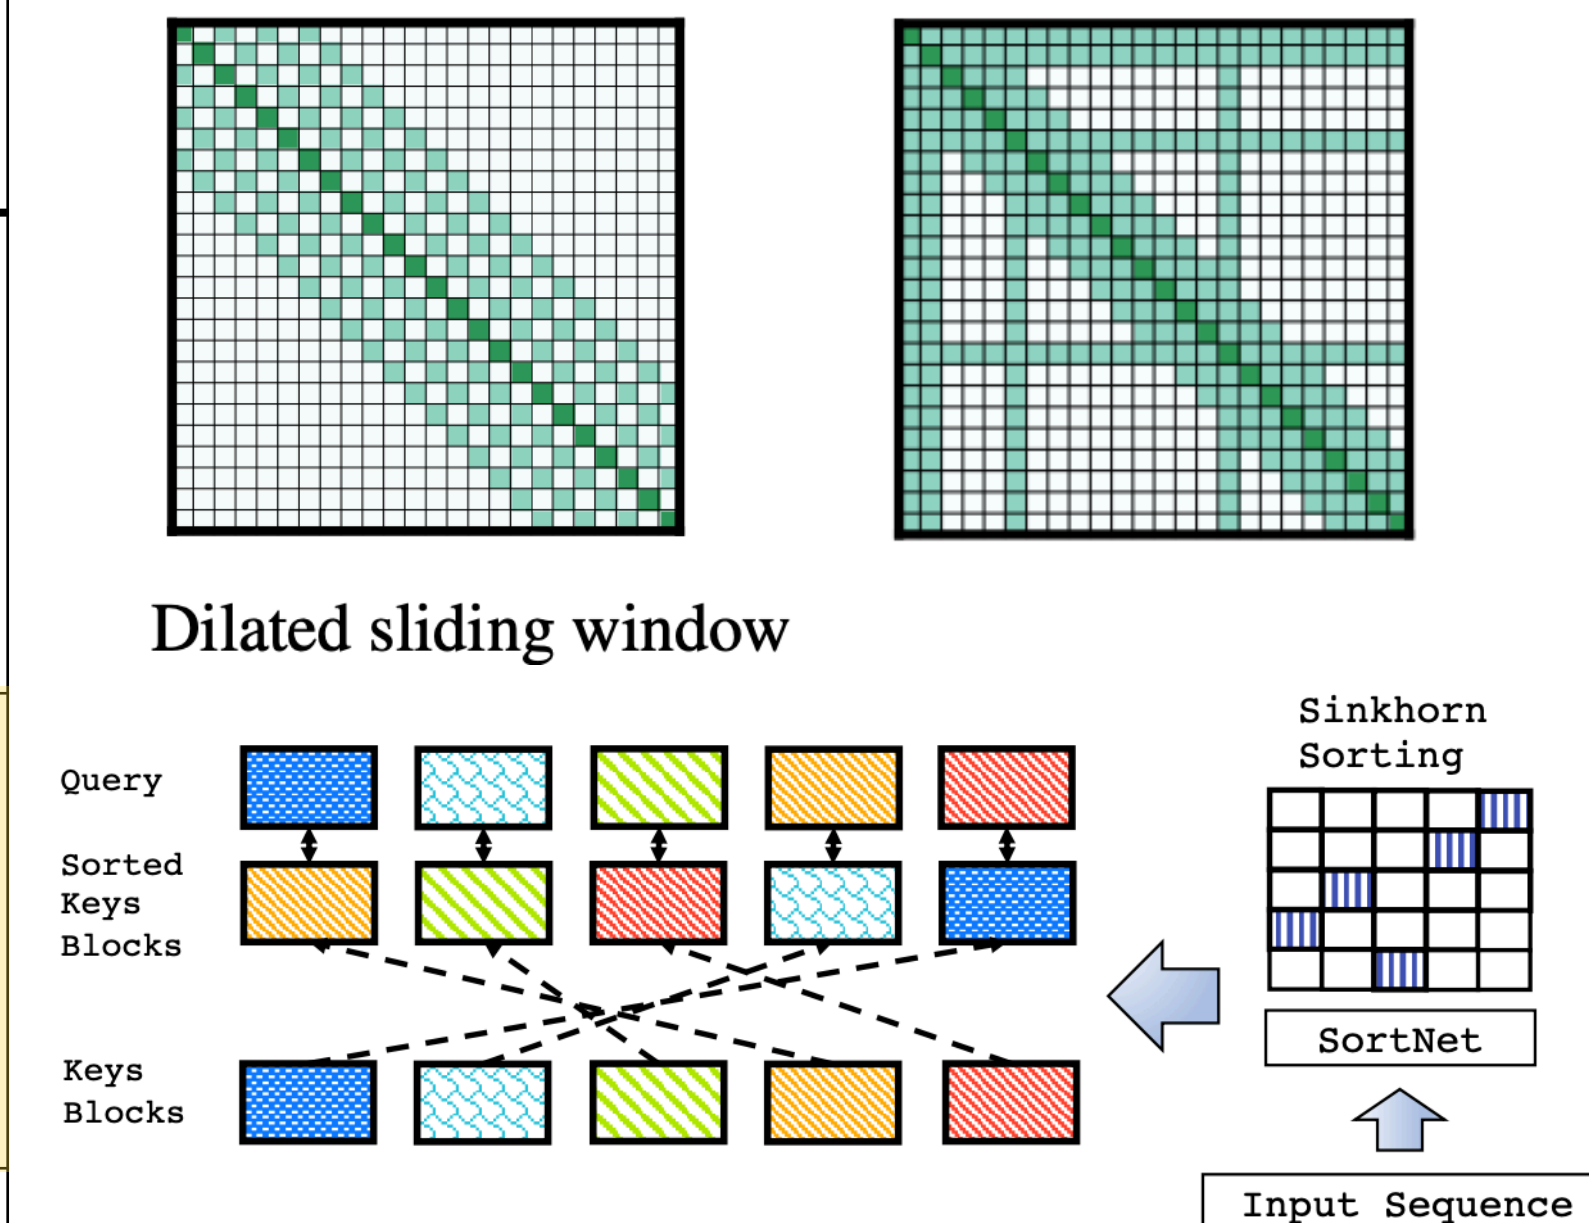


Dilated sliding window

(a) Full n^2 attention

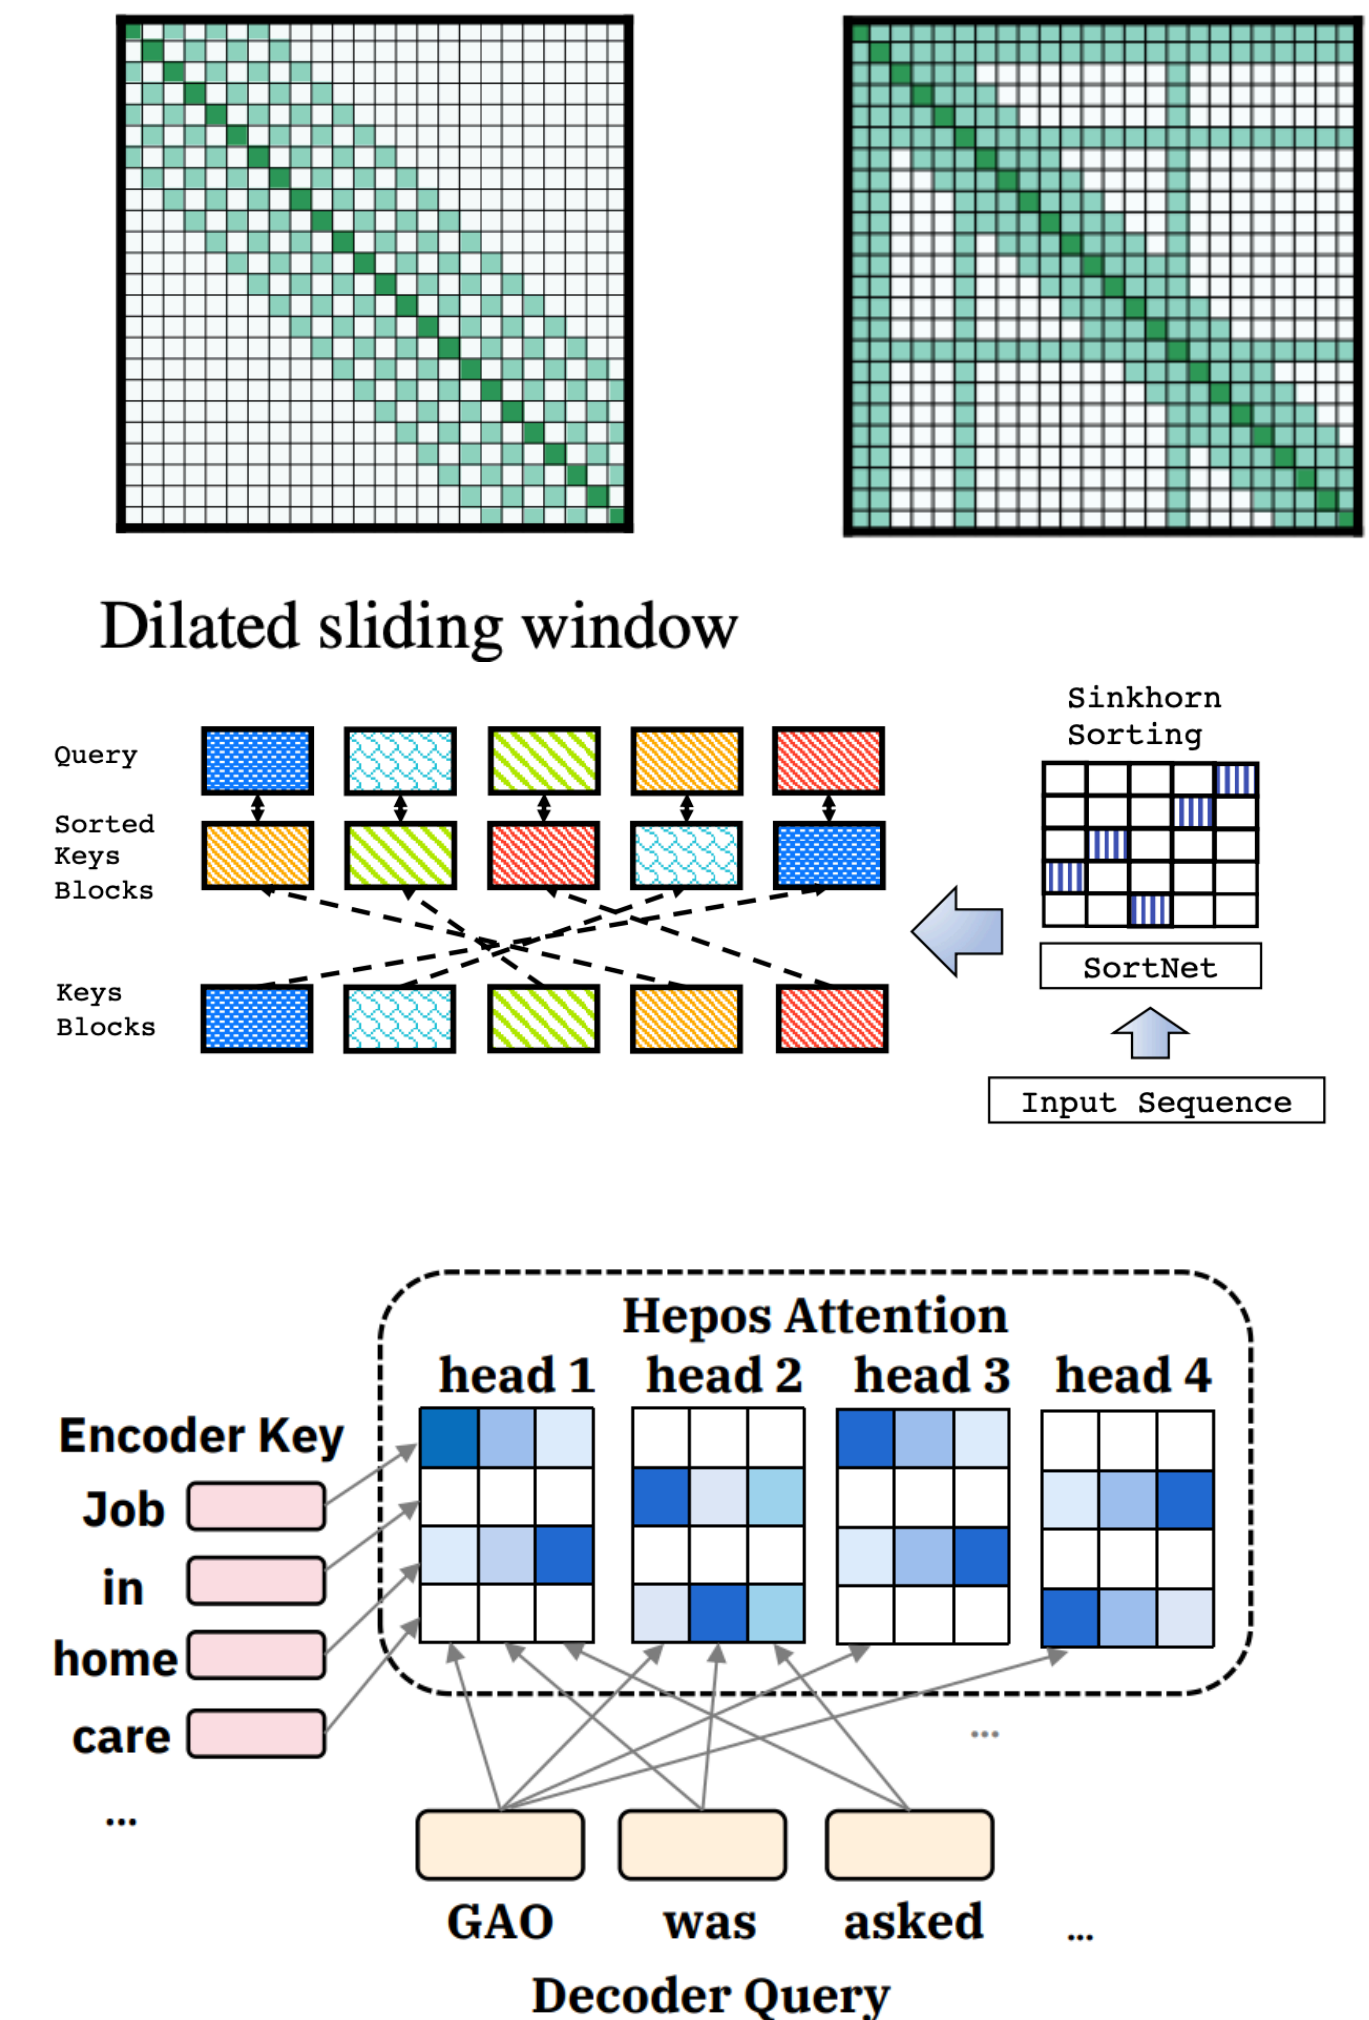
Efficient Attention Operations

Paper	Model Context Length	Complexity	Strategy
Longformer: The Long-Document Transformer , Beltagy et al. (2020)	23,040	$O(nw)$ <small>w = window size</small>	Sliding window attention
Sparse Sinkhorn Attention , Tay et al. (2020)	2048	$O(2nb)$ <small>b = number of tokens to attend in each block</small>	Learnable sparse attention pattern
Efficient Attentions for Long Document Summarization , Huang et al. (2021)	10,000	$O(mn/s)$ <small>m = output length s = stride size</small>	Head-wise positional strides



Efficient Attention Operations

Paper	Model Context Length	Complexity	Strategy
Longformer: The Long-Document Transformer , Beltagy et al. (2020)	23,040	$O(nw)$ w = window size	Sliding window attention
Sparse Sinkhorn Attention , Tay et al. (2020)	2048	$O(2nb)$ b = number of tokens to attend in each block	Learnable sparse attention pattern
Efficient Attentions for Long Document Summarization , Huang et al. (2021)	10,000	$O(mn/s)$ m = output length s = stride size	Head-wise positional strides



Efficient Attention Operations

Different strategies perform similarly across 2 summarization tasks

Longer documents

GovReport

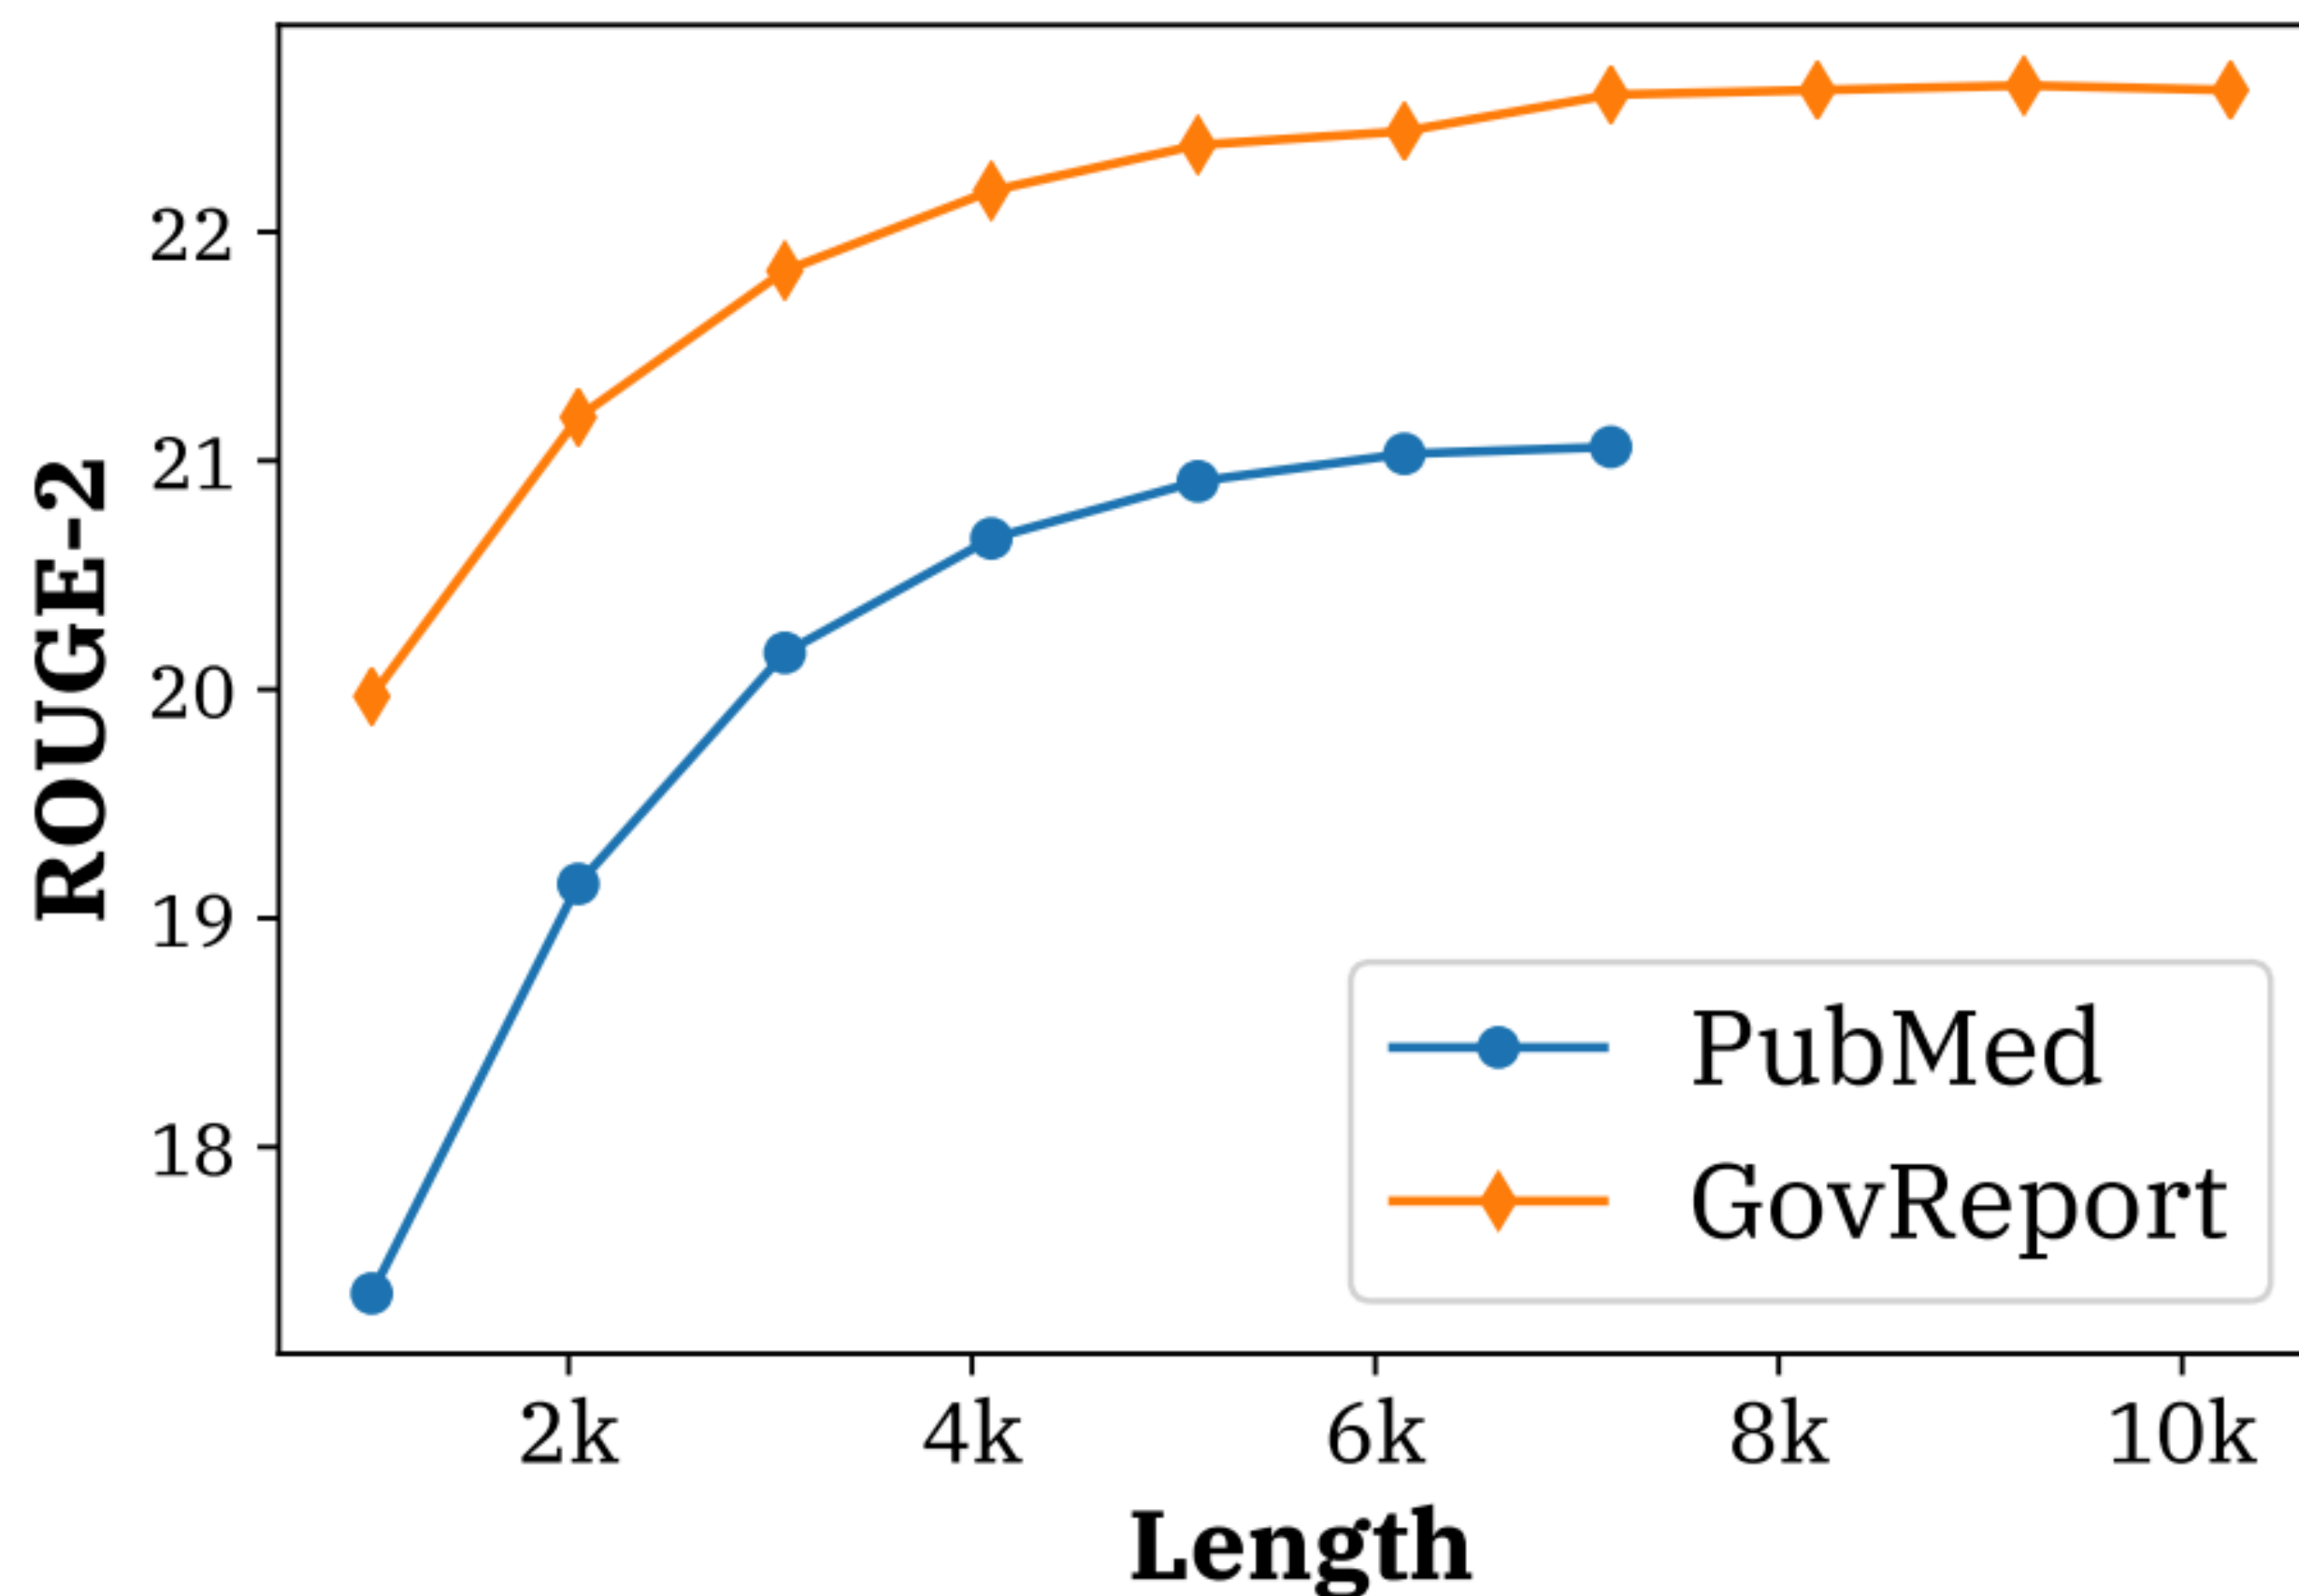
PubMed

	R-1	R-2	R-L	R-1	R-2	R-L
GPT-series	51.5	19.1	48.7	43.7	17.3	38.8
Longformer	51.2	19.0	48.6	43.4	17.1	38.6
Sinkhorn	53.0	20.1	50.3	45.1	18.4	40.1
Hepos	51.3	19.1	48.7	44.9	18.2	39.9

Efficient Attentions for Long Document Summarization, Huang et al. (2021)

Context Window Length Matters

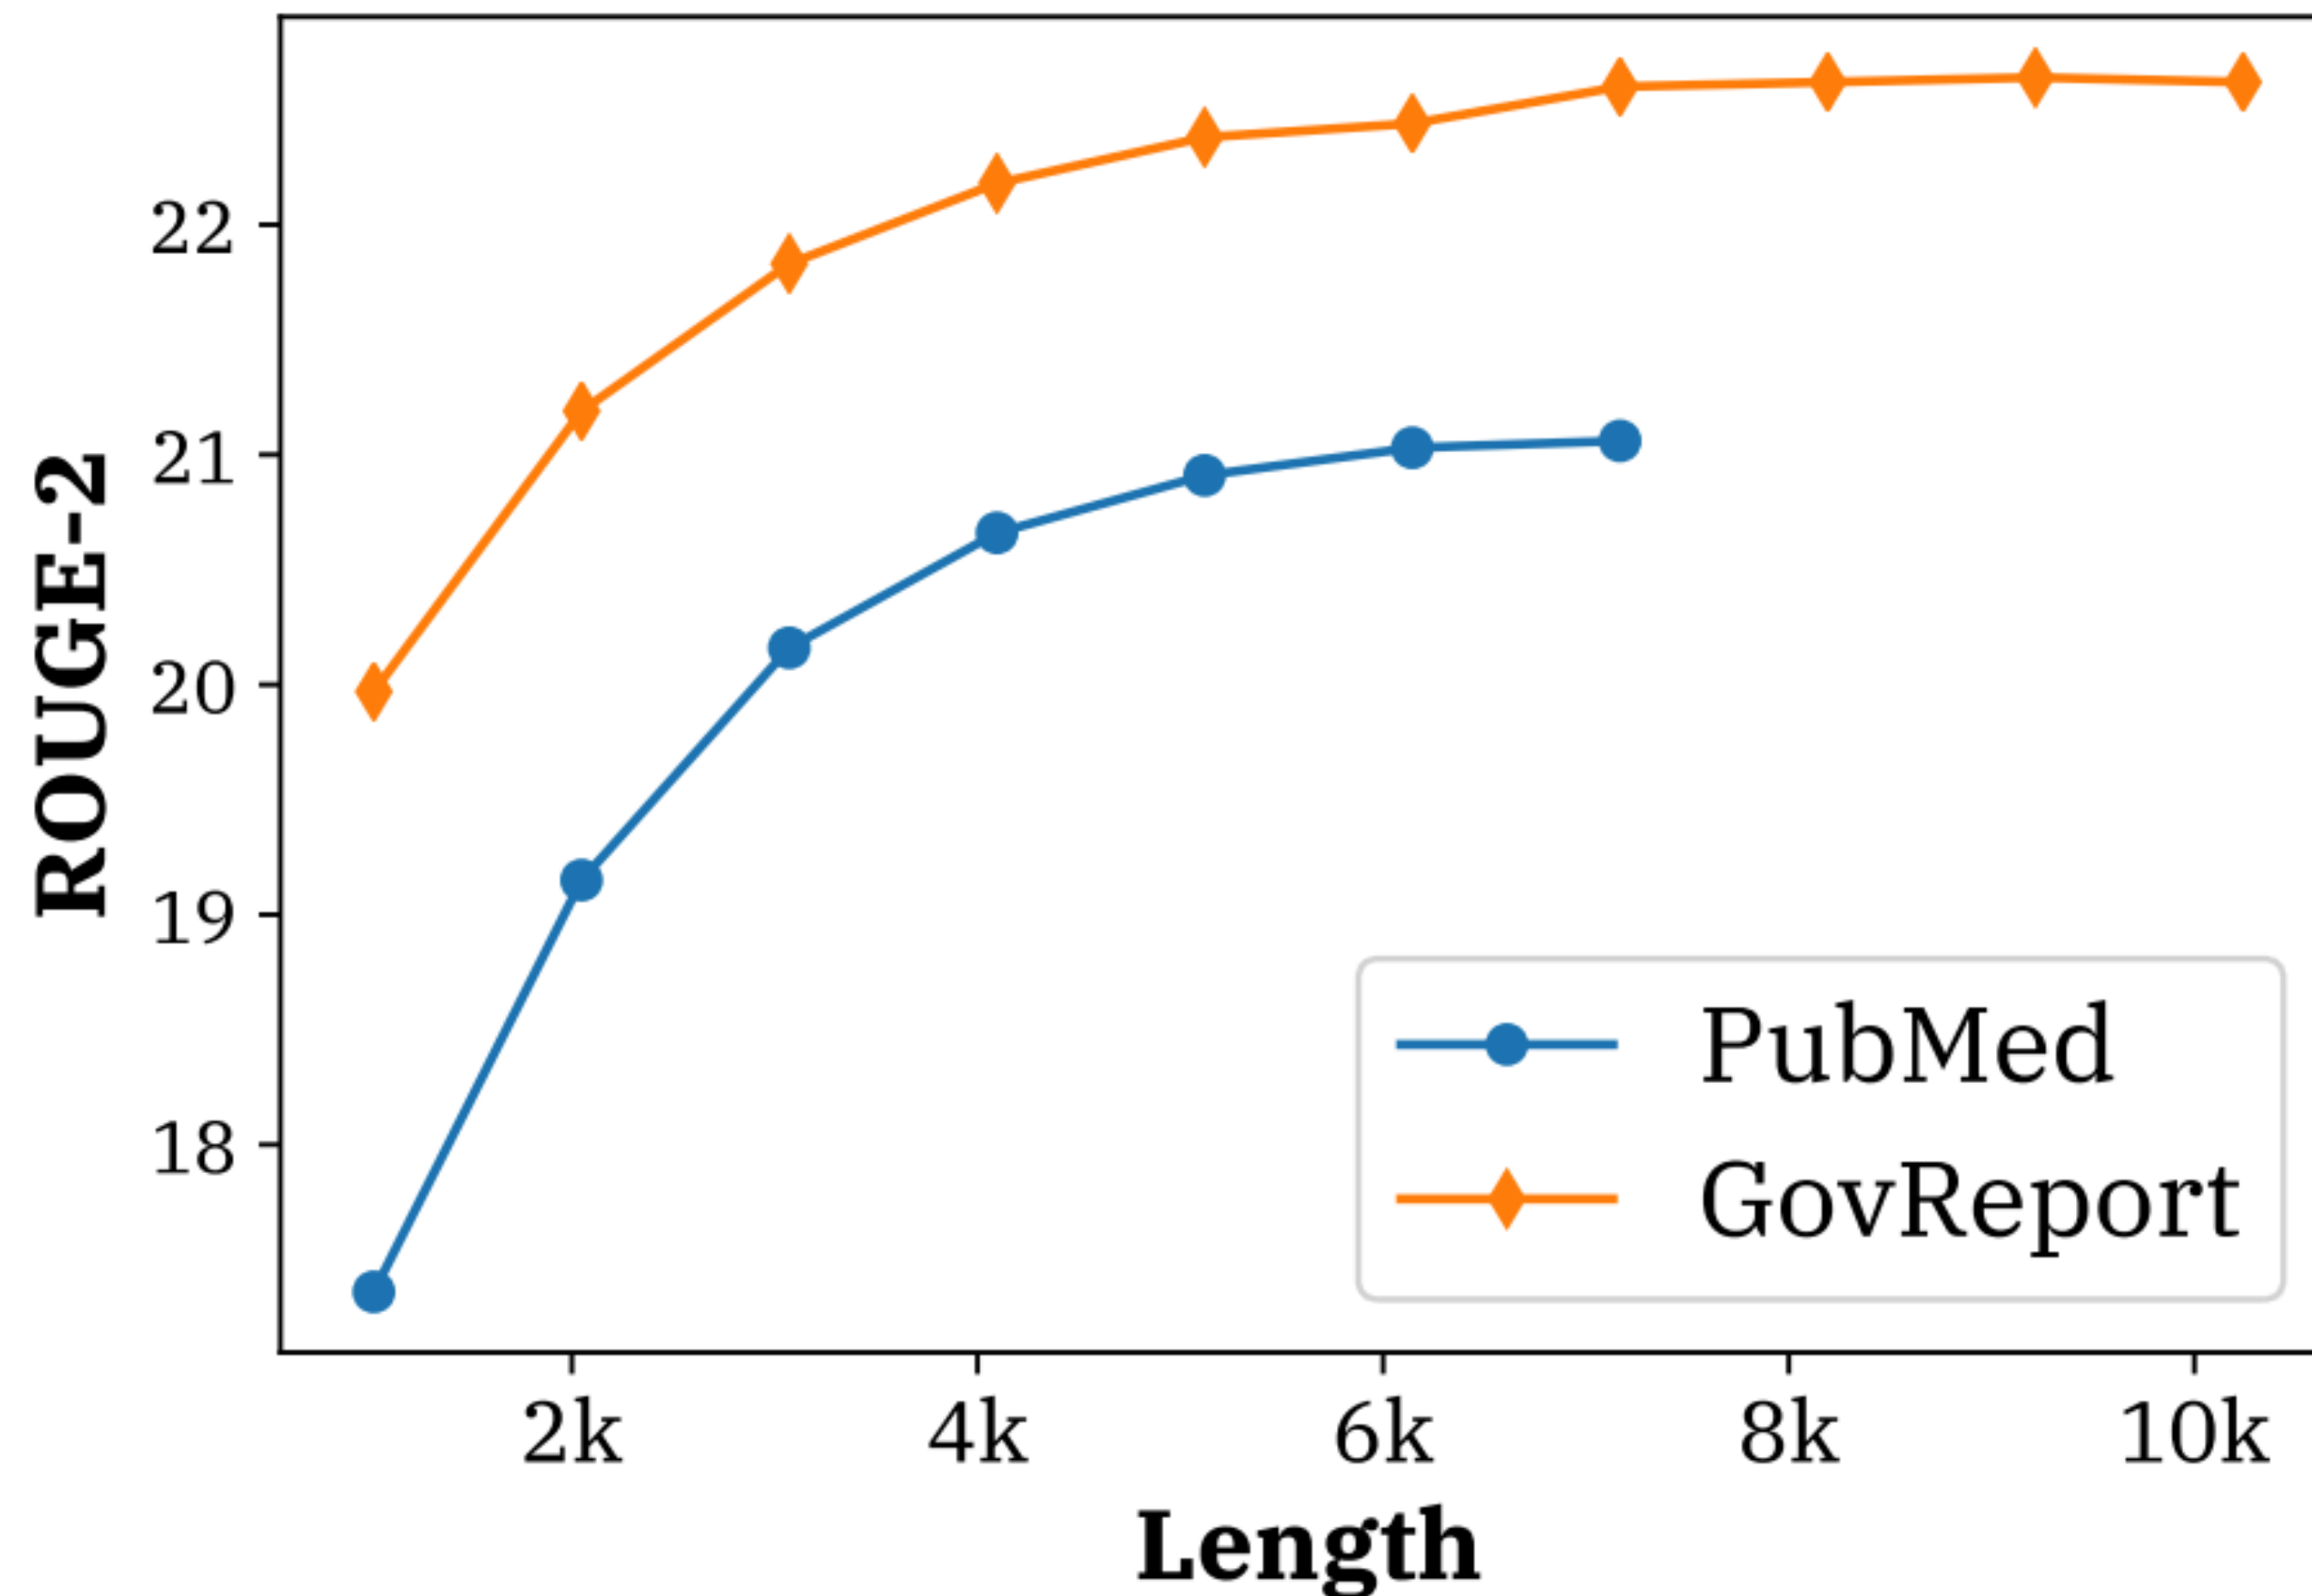
Summarization improves with a longer input context



Efficient Attentions for Long Document Summarization, Huang et al. (2021)

Context Window Length Matters

Summarization improves with a longer input context

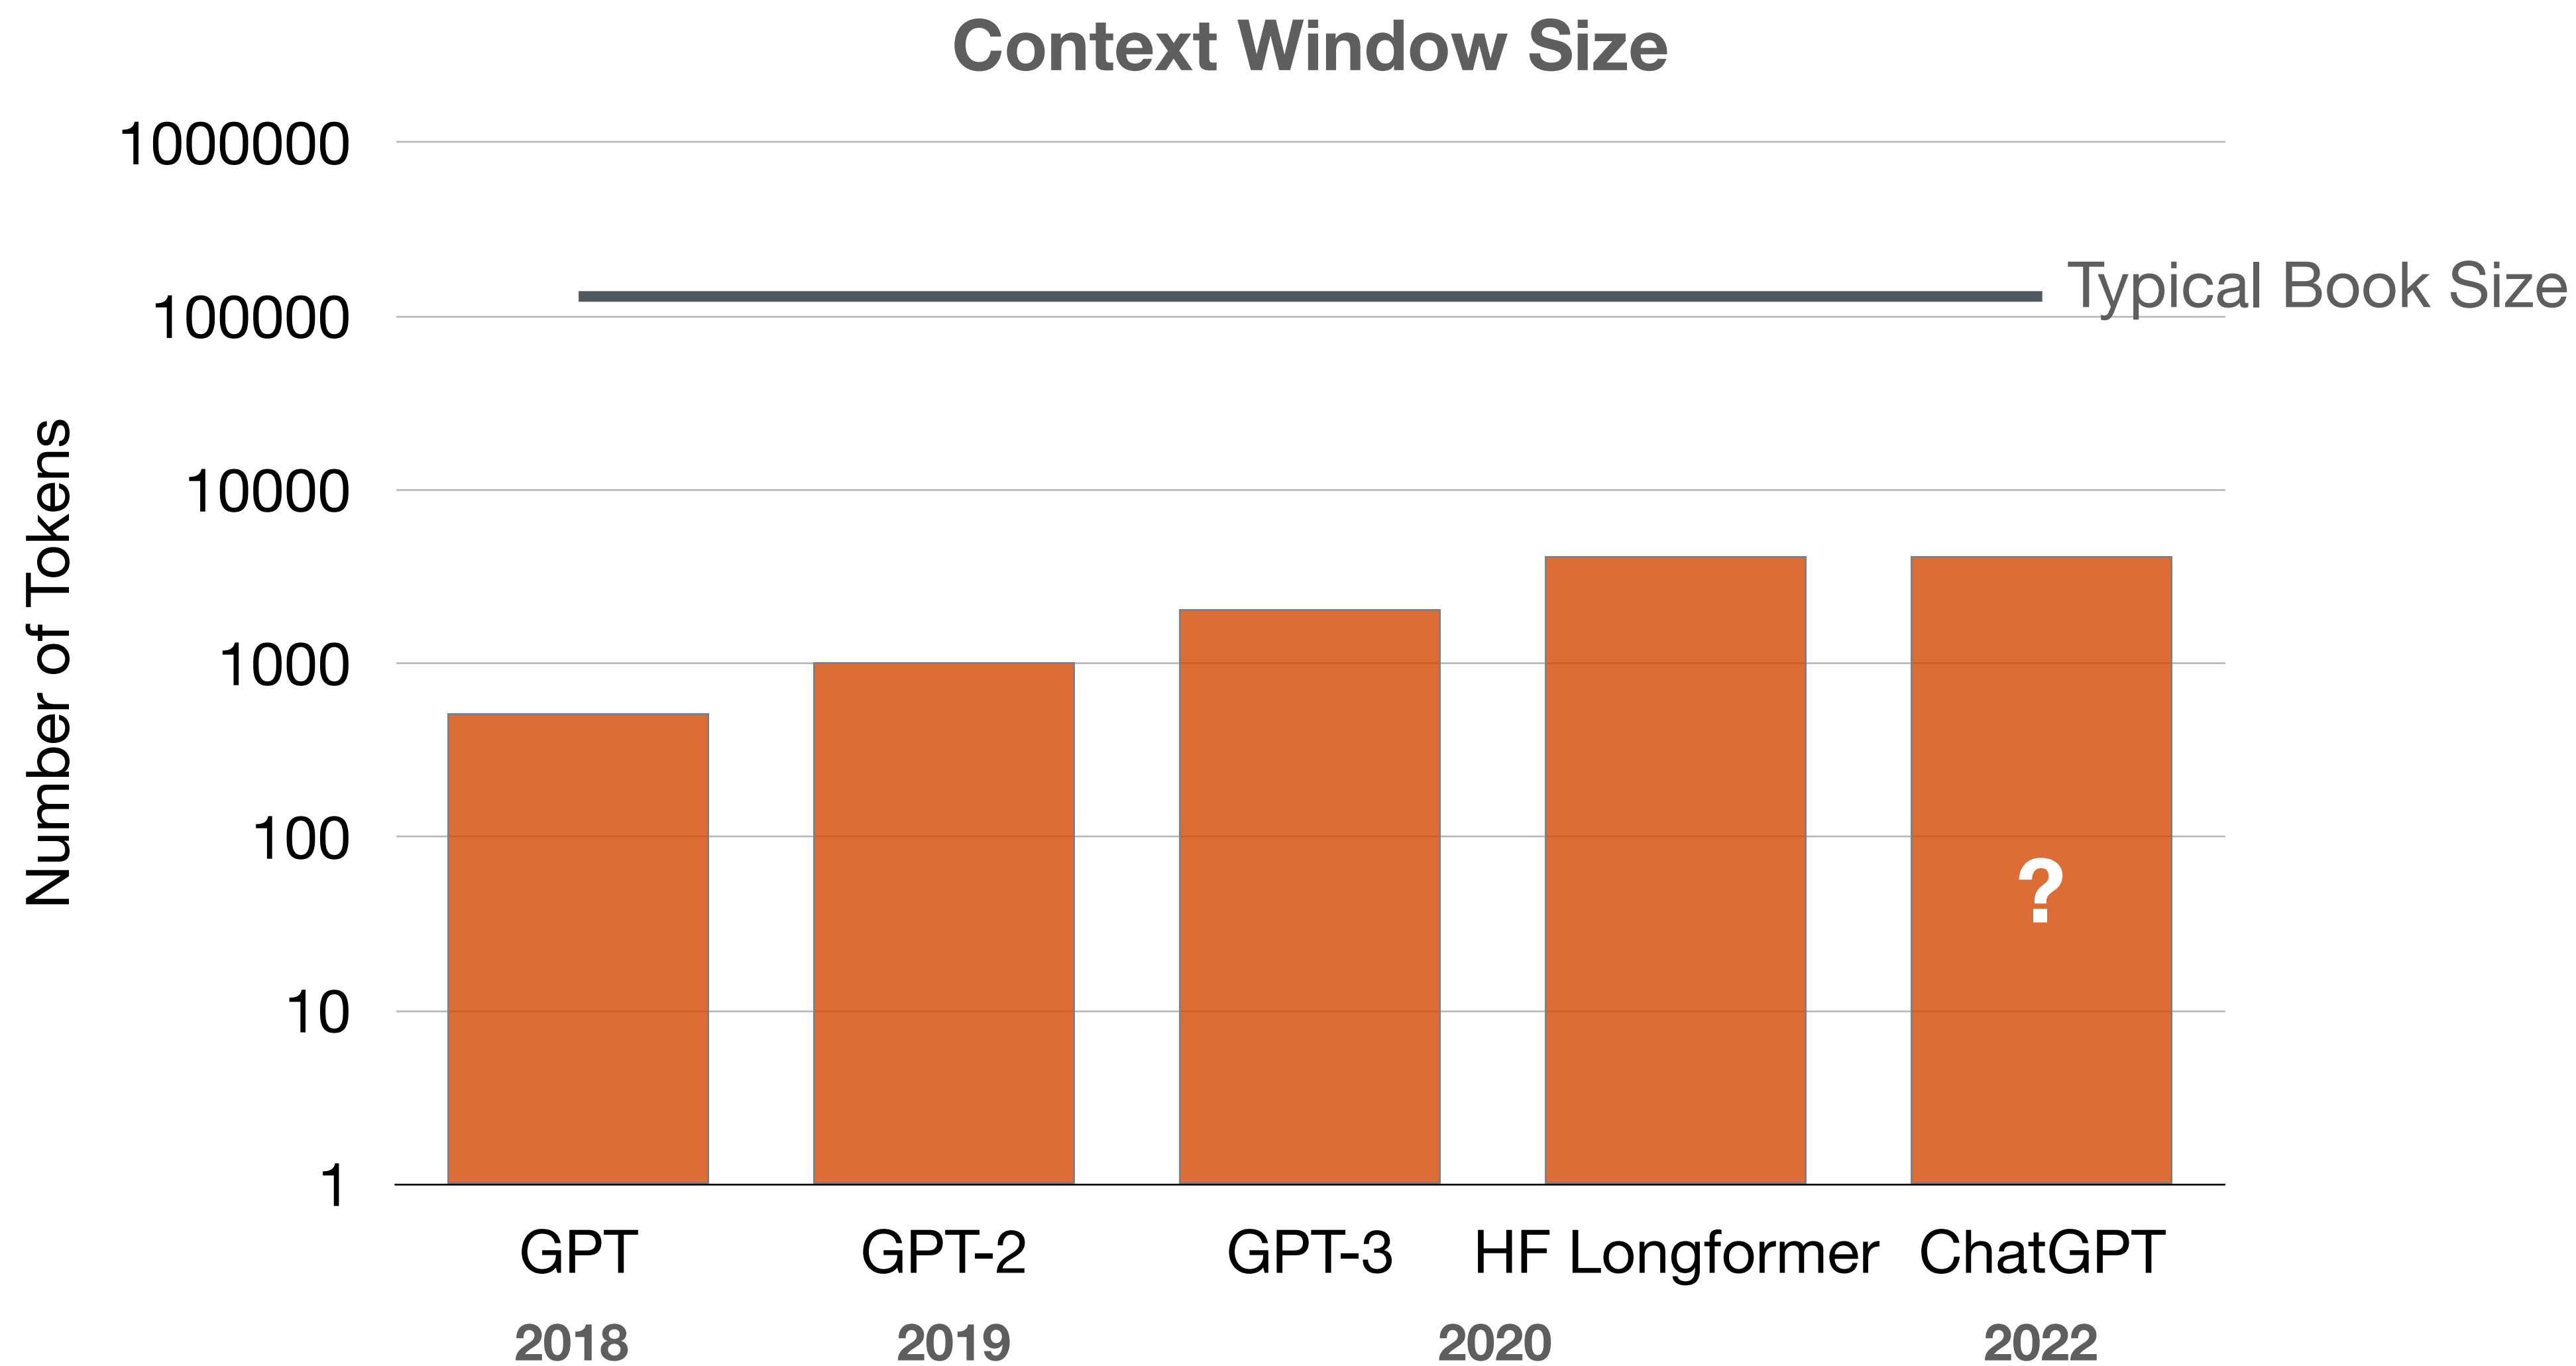


Longer documents

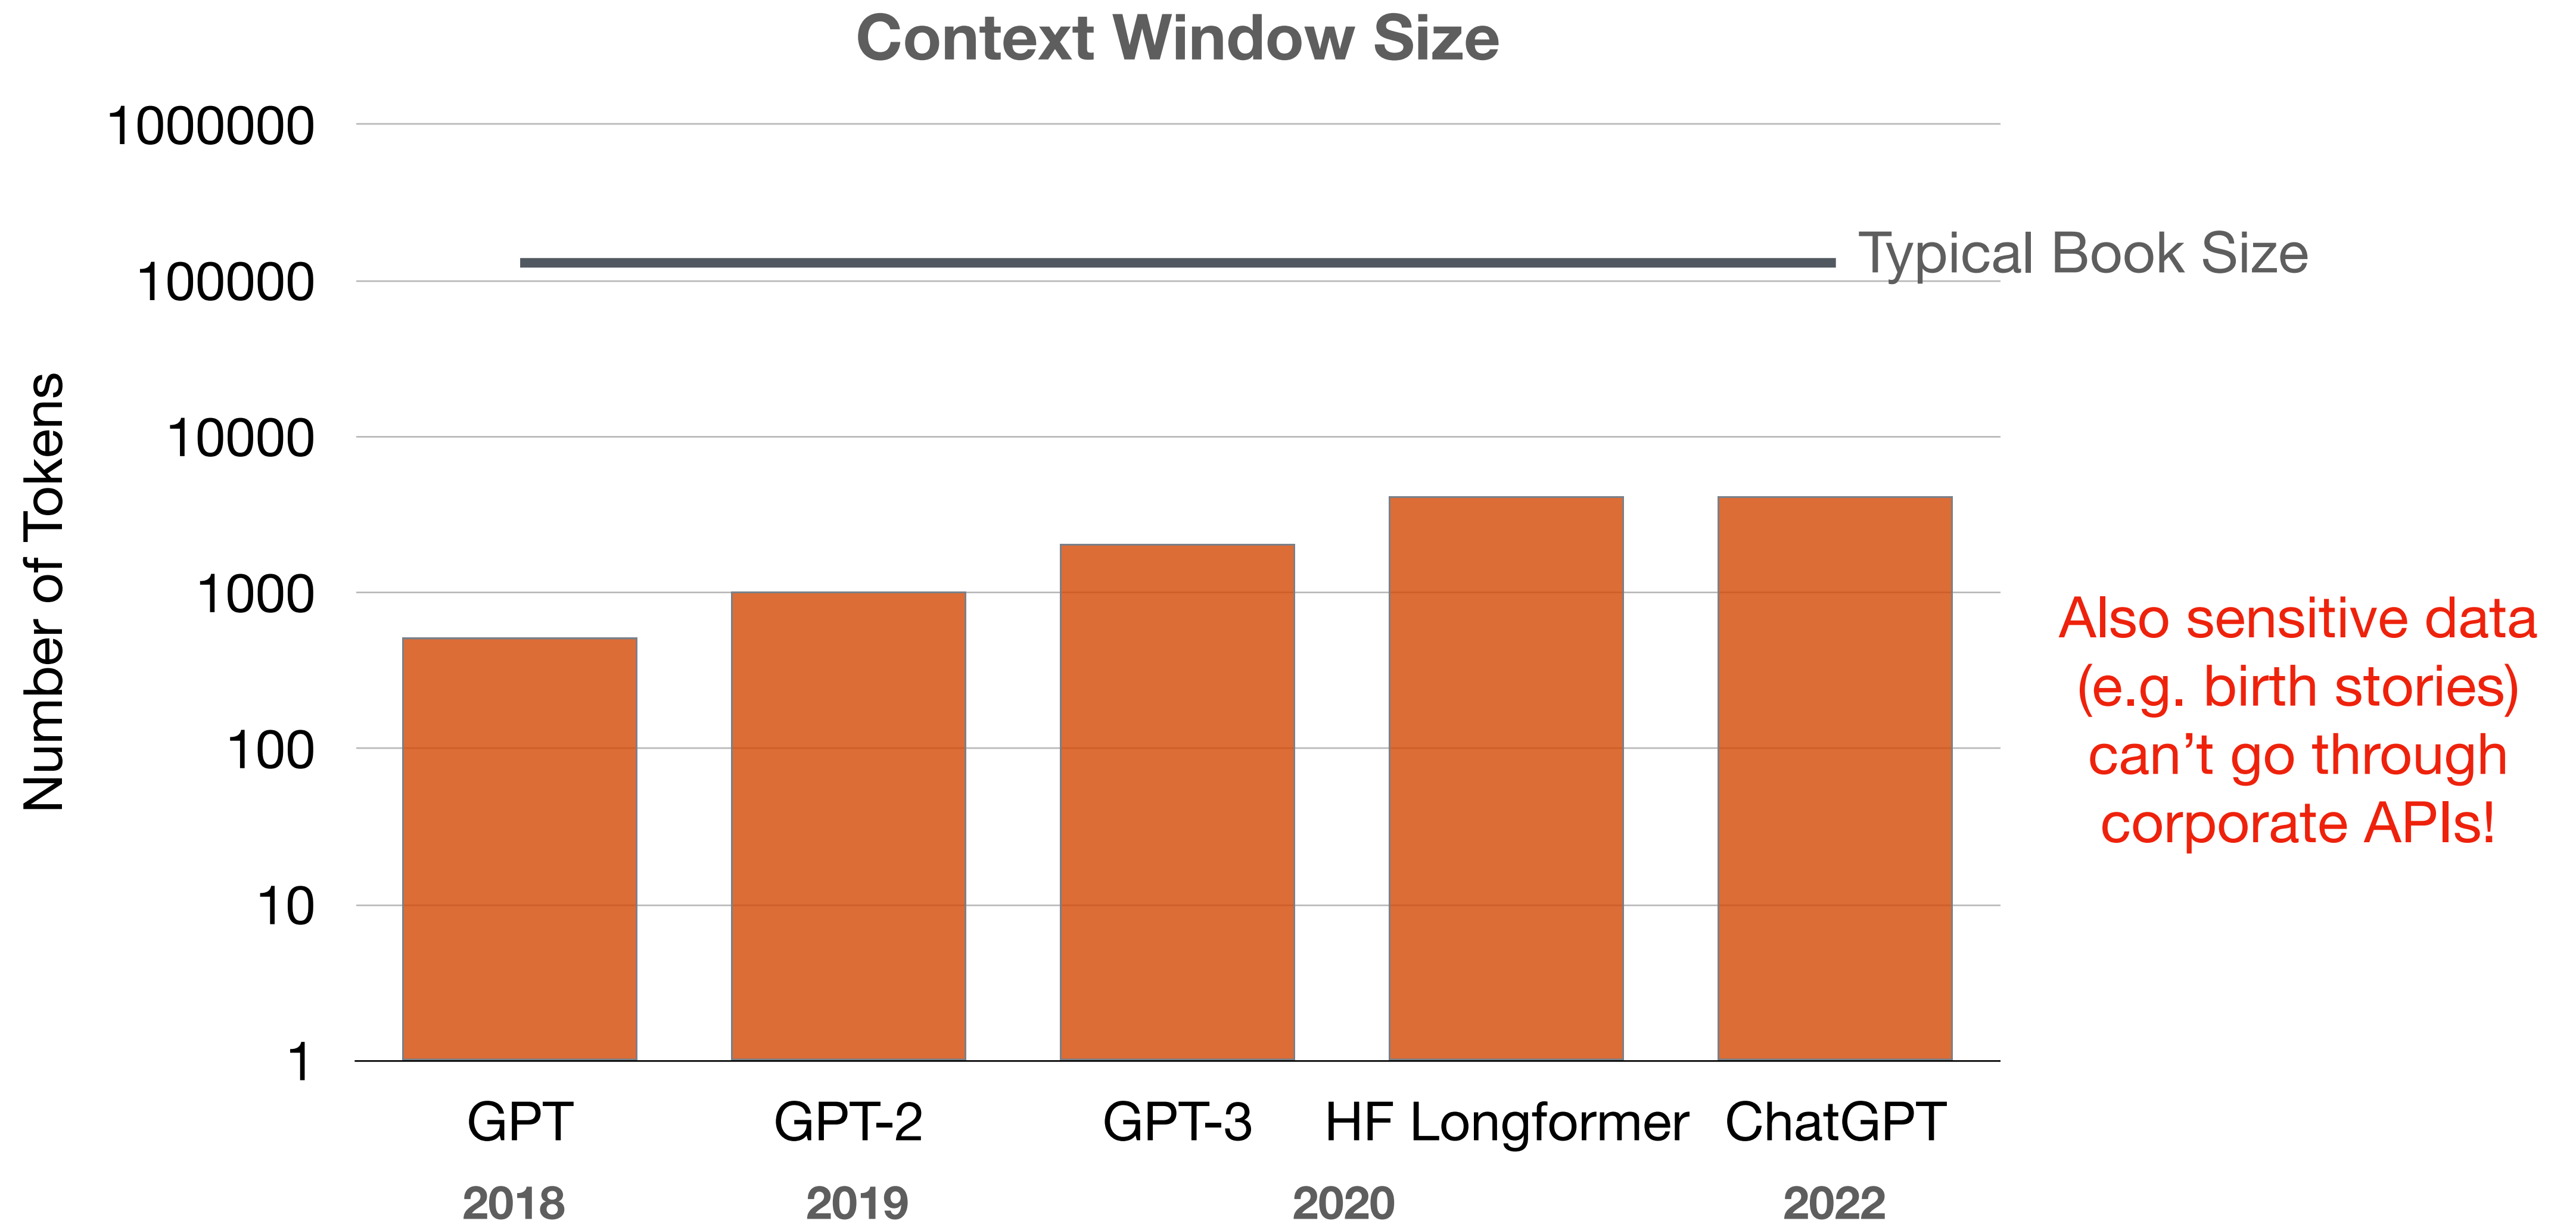
Narrative like books
would likely flatten out at
a much higher length

Efficient Attentions for Long Document Summarization, Huang et al. (2021)

Are efficient attention operations enough?



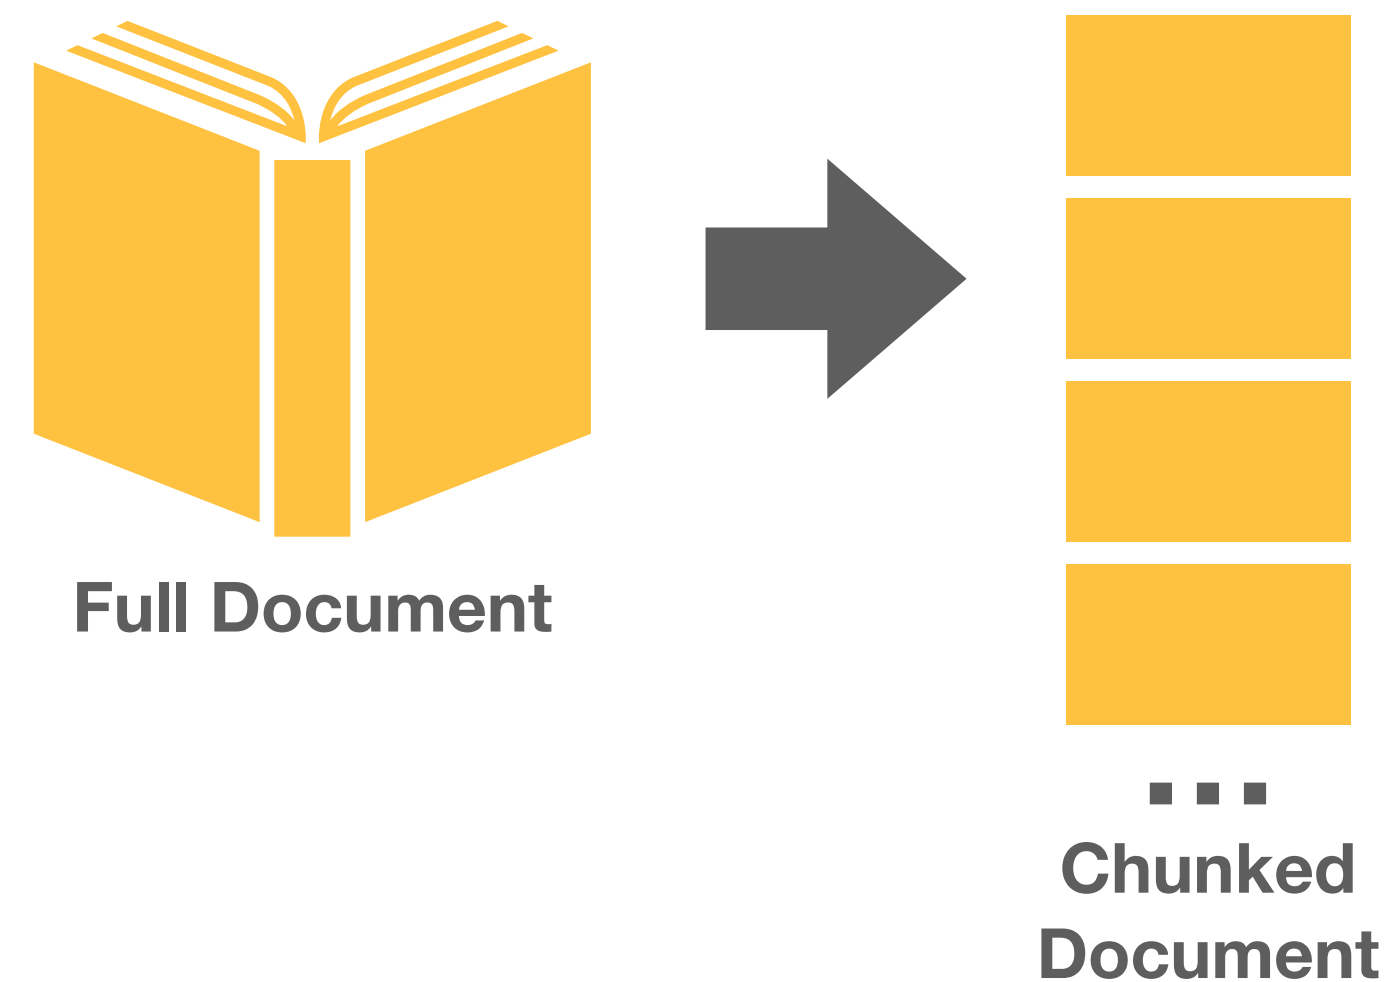
Are efficient attention operations enough?



Summarization-Specific Approaches

1. Extract (-then-generate)

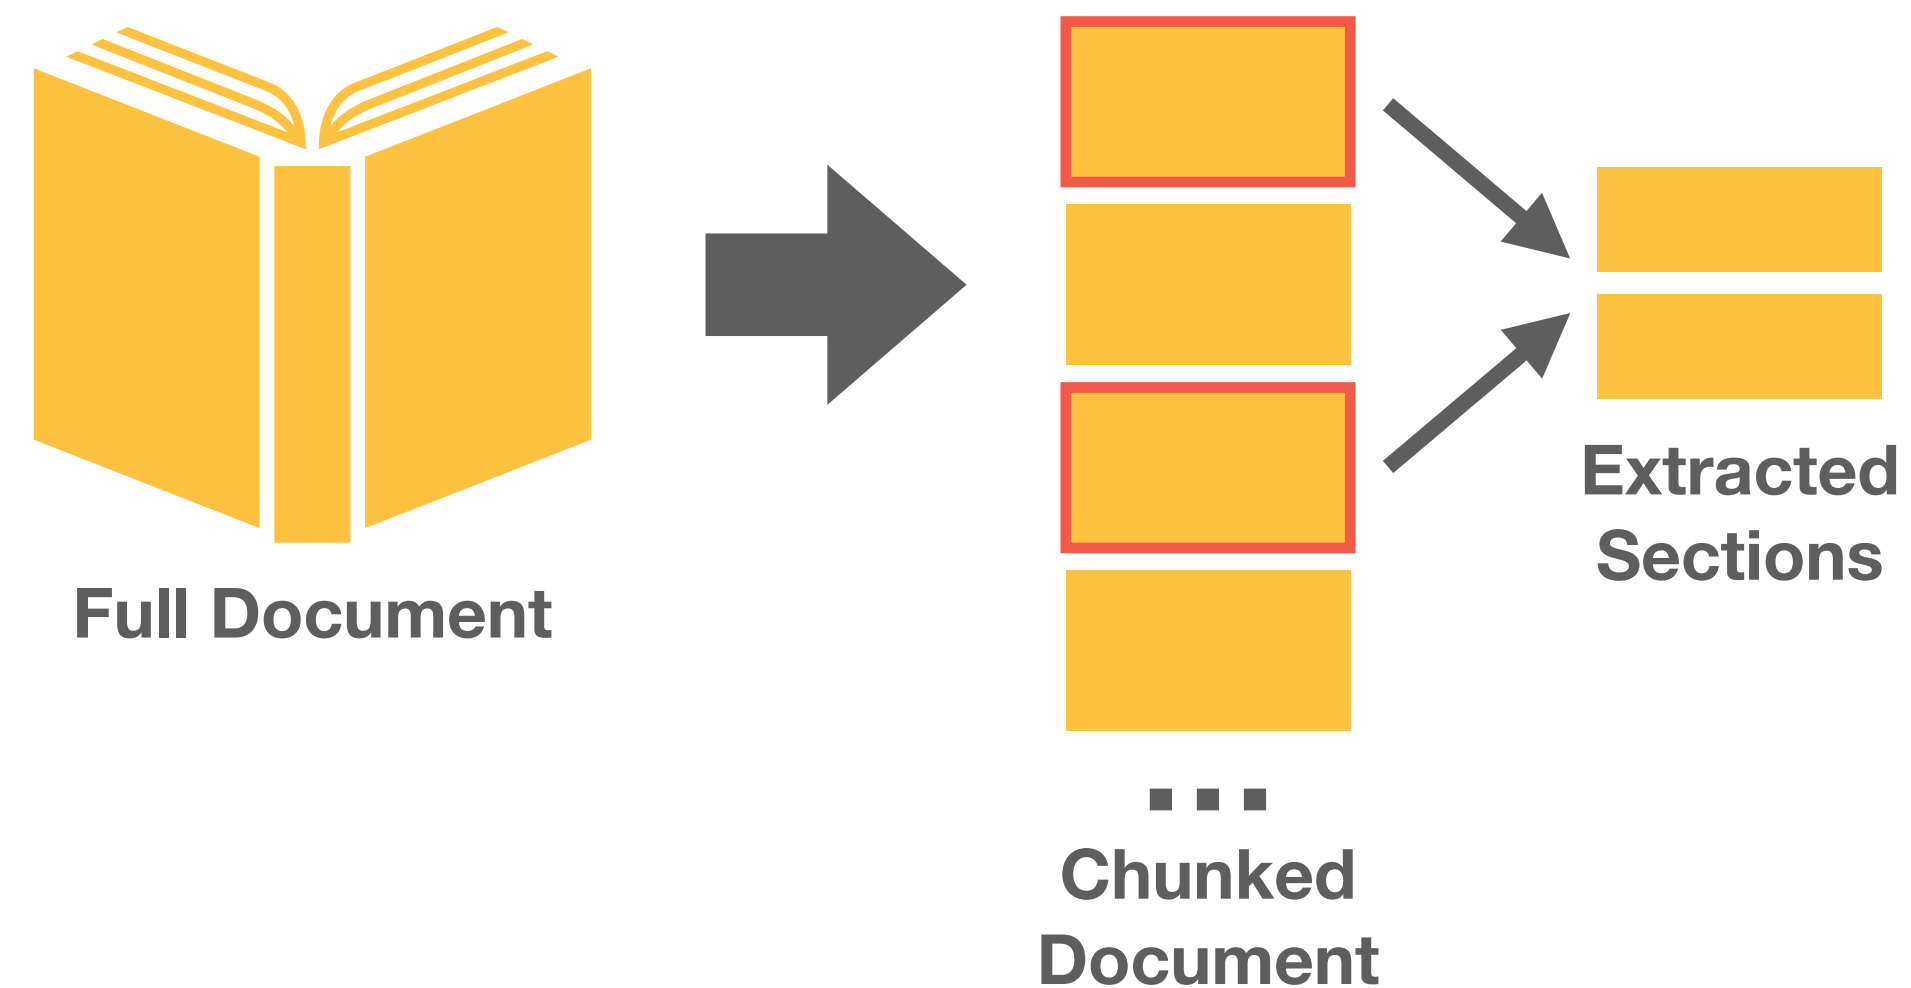
- *DYLE: Dynamic Latent Extraction for Abstractive Long-Input Summarization*, Mao et al. (2021)
- *Discourse-Aware Unsupervised Summarization of Long Scientific Documents*, Dong et al. (2020)



Summarization-Specific Approaches

1. Extract (-then-generate)

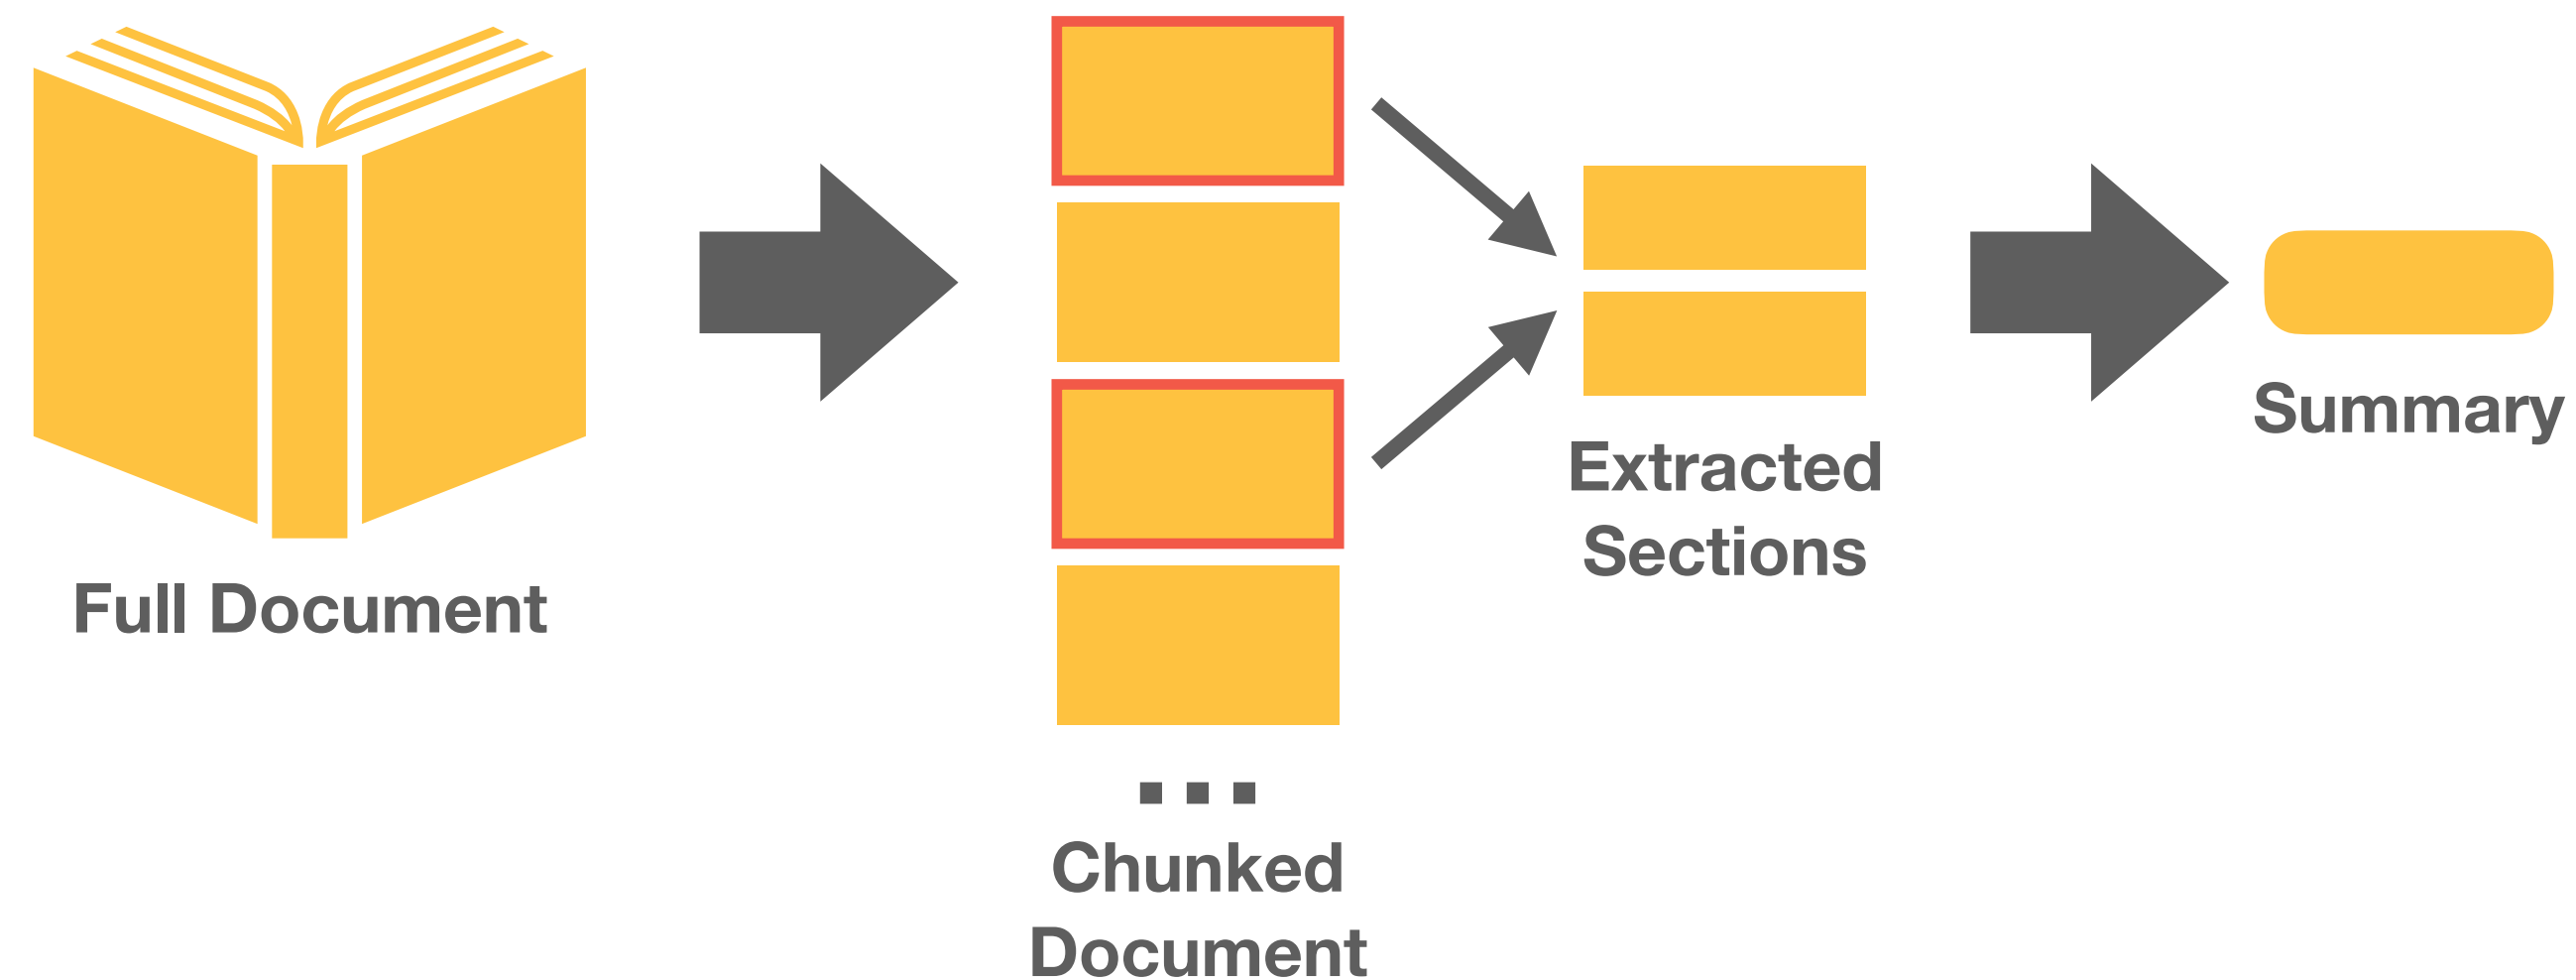
- *DYLE: Dynamic Latent Extraction for Abstractive Long-Input Summarization*, Mao et al. (2021)
- *Discourse-Aware Unsupervised Summarization of Long Scientific Documents*, Dong et al. (2020)



Summarization-Specific Approaches

1. Extract (-then-generate)

- *DYLE: Dynamic Latent Extraction for Abstractive Long-Input Summarization*, Mao et al. (2021)
- *Discourse-Aware Unsupervised Summarization of Long Scientific Documents*, Dong et al. (2020)



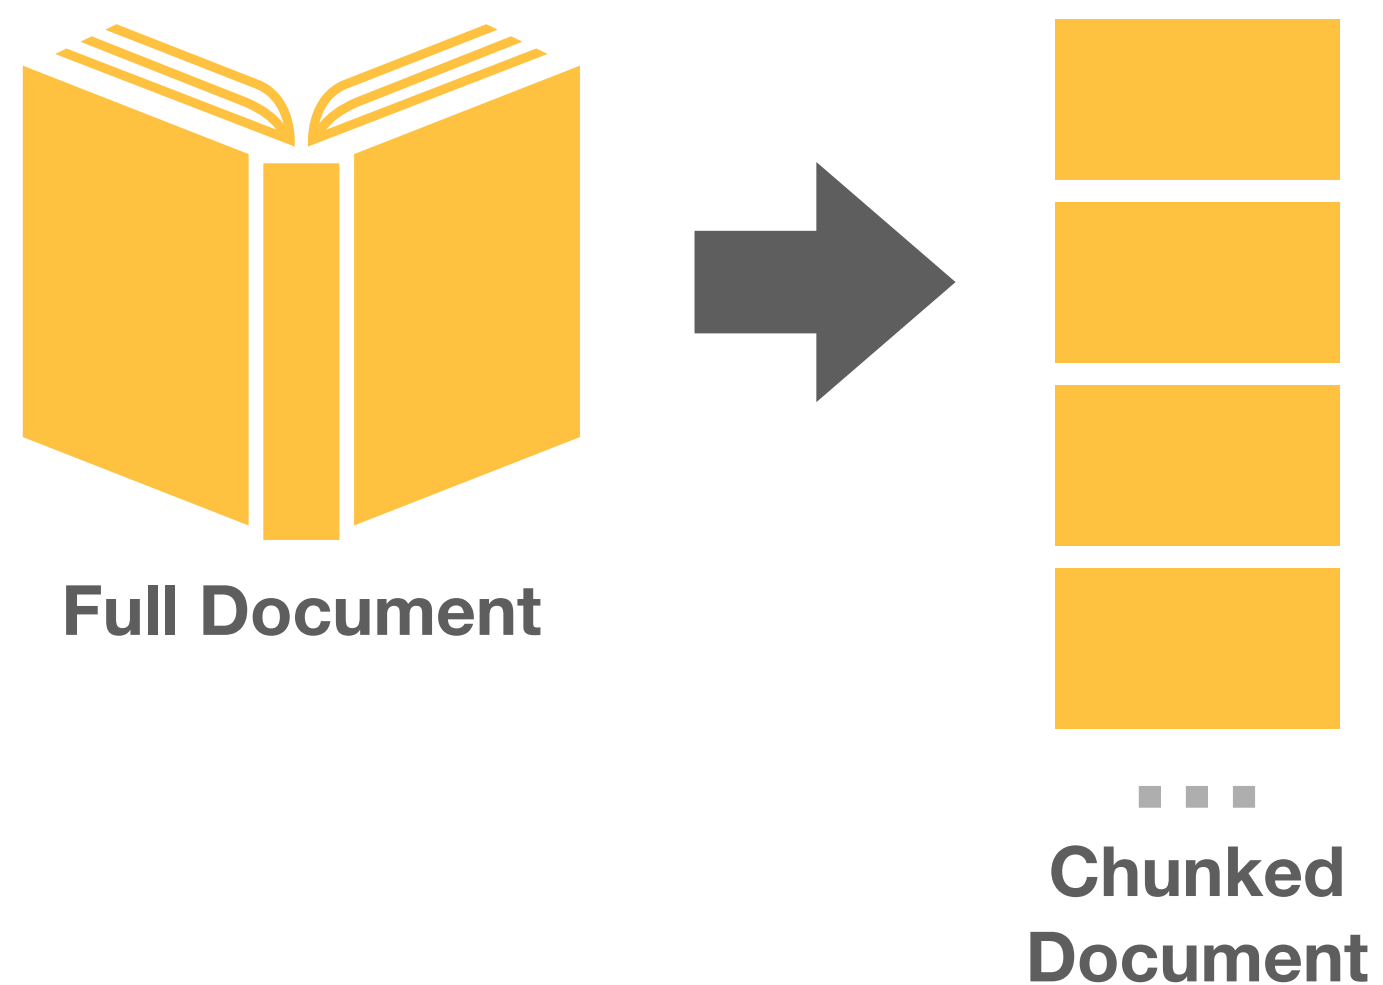
Summarization-Specific Approaches

1. Extract (-then-generate)

- *DYLE: Dynamic Latent Extraction for Abstractive Long-Input Summarization*, Mao et al. (2021)
- *Discourse-Aware Unsupervised Summarization of Long Scientific Documents*, Dong et al. (2020)

2. Divide-and-conquer

- *A Divide-and-Conquer Approach to the Summarization of Long Documents*, Gidiotis et al. (2020)



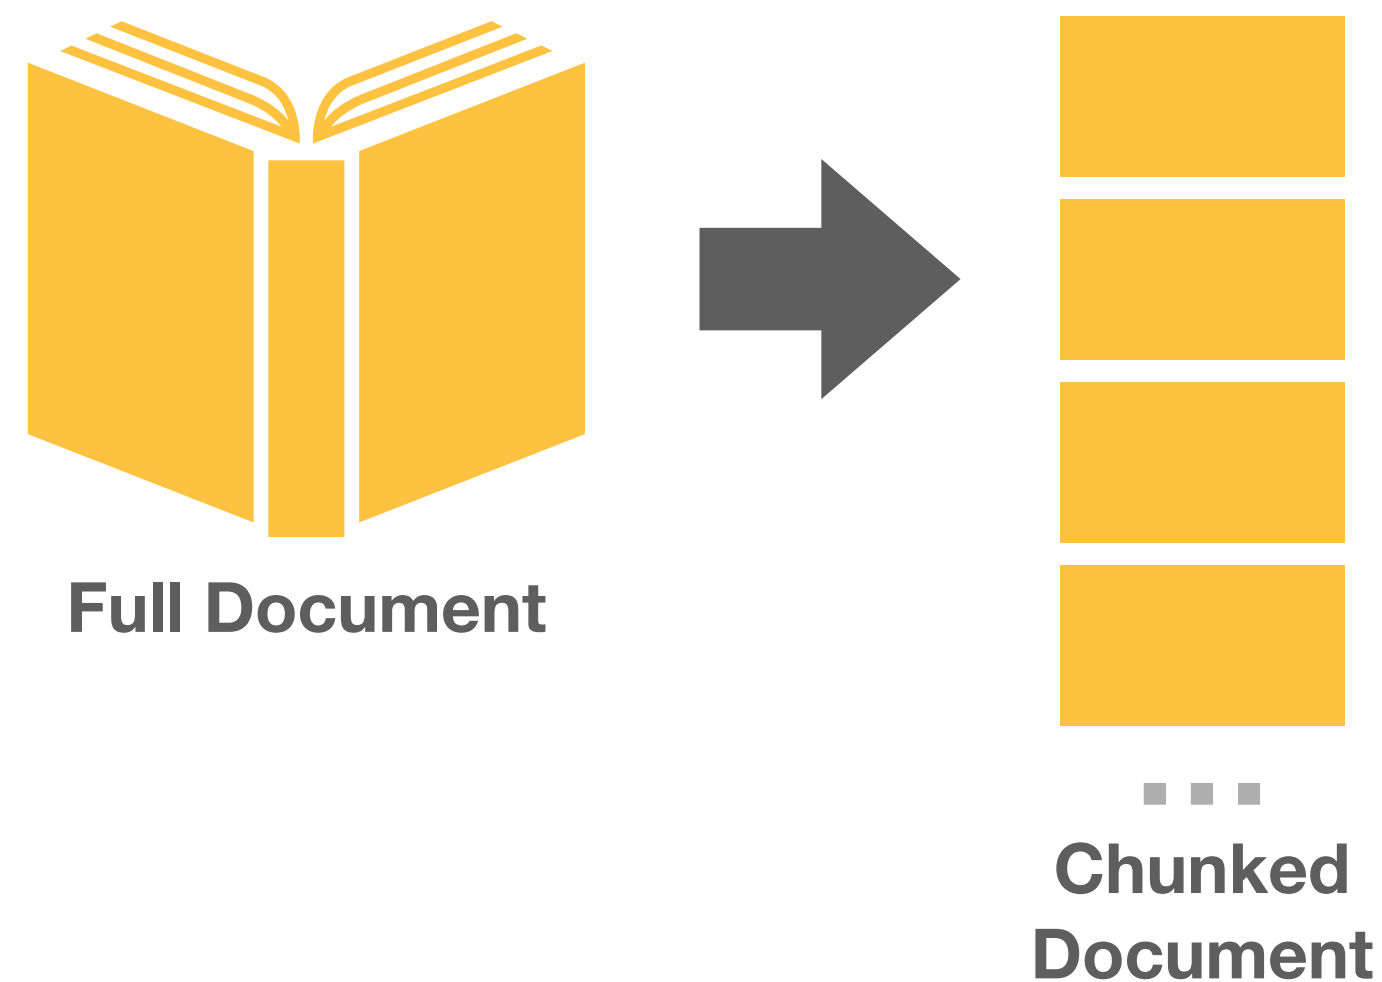
Summarization-Specific Approaches

1. Extract (-then-generate)

- *DYLE: Dynamic Latent Extraction for Abstractive Long-Input Summarization*, Mao et al. (2021)
- *Discourse-Aware Unsupervised Summarization of Long Scientific Documents*, Dong et al. (2020)

2. Divide-and-conquer

- *A Divide-and-Conquer Approach to the Summarization of Long Documents*, Gidiotis et al. (2020)



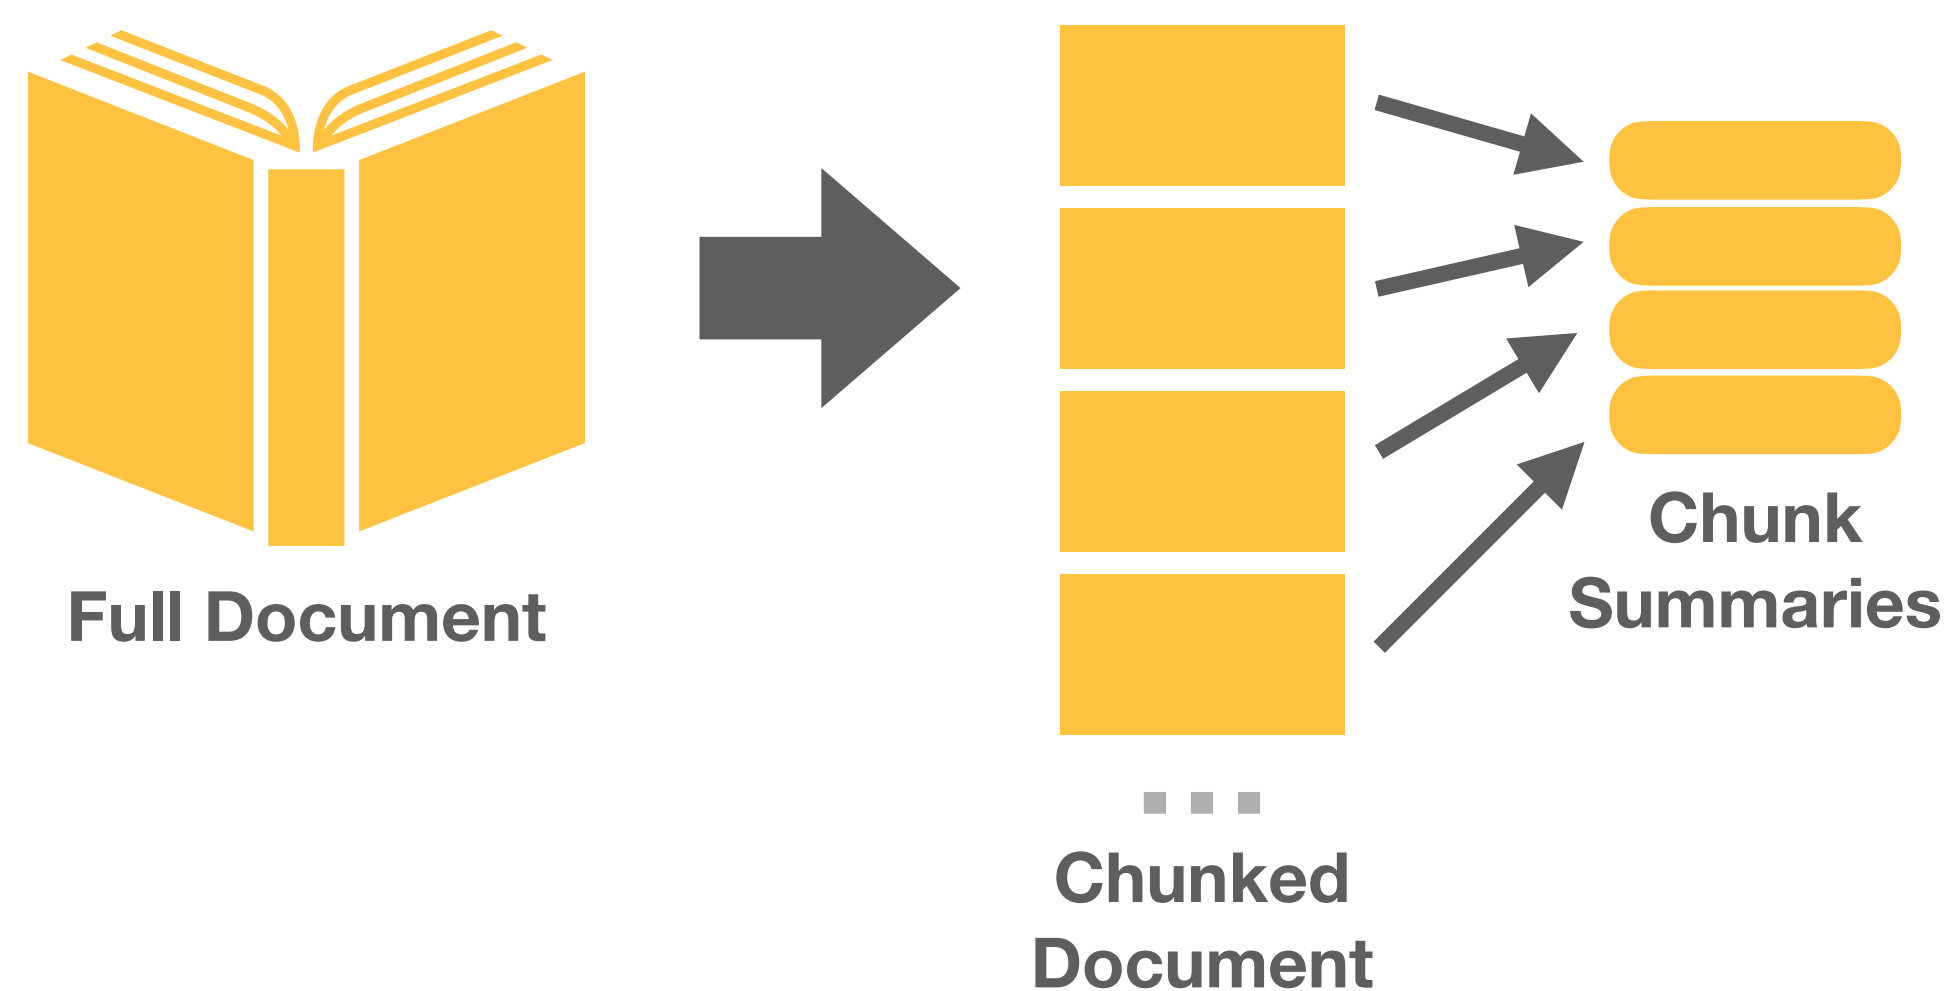
Summarization-Specific Approaches

1. Extract (-then-generate)

- *DYLE: Dynamic Latent Extraction for Abstractive Long-Input Summarization*, Mao et al. (2021)
- *Discourse-Aware Unsupervised Summarization of Long Scientific Documents*, Dong et al. (2020)

2. Divide-and-conquer

- *A Divide-and-Conquer Approach to the Summarization of Long Documents*, Gidiotis et al. (2020)



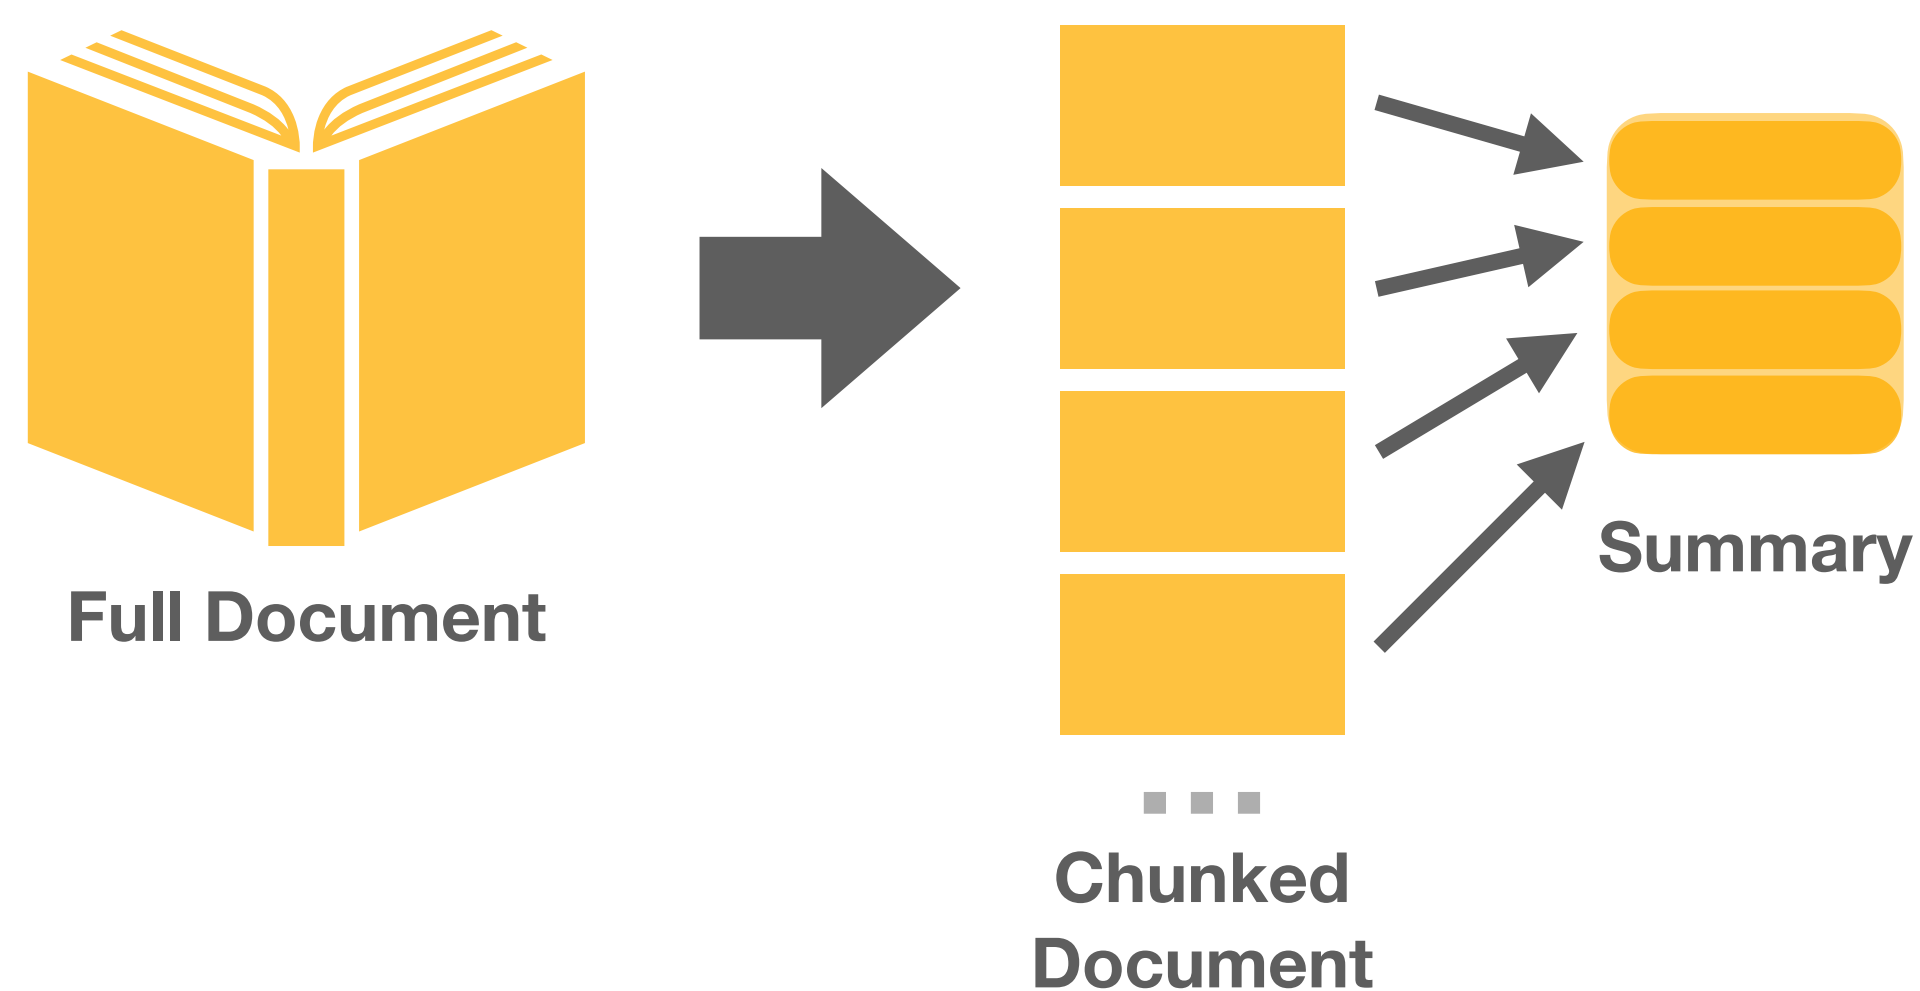
Summarization-Specific Approaches

1. Extract (-then-generate)

- *DYLE: Dynamic Latent Extraction for Abstractive Long-Input Summarization*, Mao et al. (2021)
- *Discourse-Aware Unsupervised Summarization of Long Scientific Documents*, Dong et al. (2020)

2. Divide-and-conquer

- *A Divide-and-Conquer Approach to the Summarization of Long Documents*, Gidiotis et al. (2020)



Summarization-Specific Approaches

1. Extract (-then-generate)

- *DYLE: Dynamic Latent Extraction for Abstractive Long-Input Summarization*, Mao et al. (2021)
- *Discourse-Aware Unsupervised Summarization of Long Scientific Documents*, Dong et al. (2020)

2. Divide-and-conquer

- *A Divide-and-Conquer Approach to the Summarization of Long Documents*, Gidiotis et al. (2020)

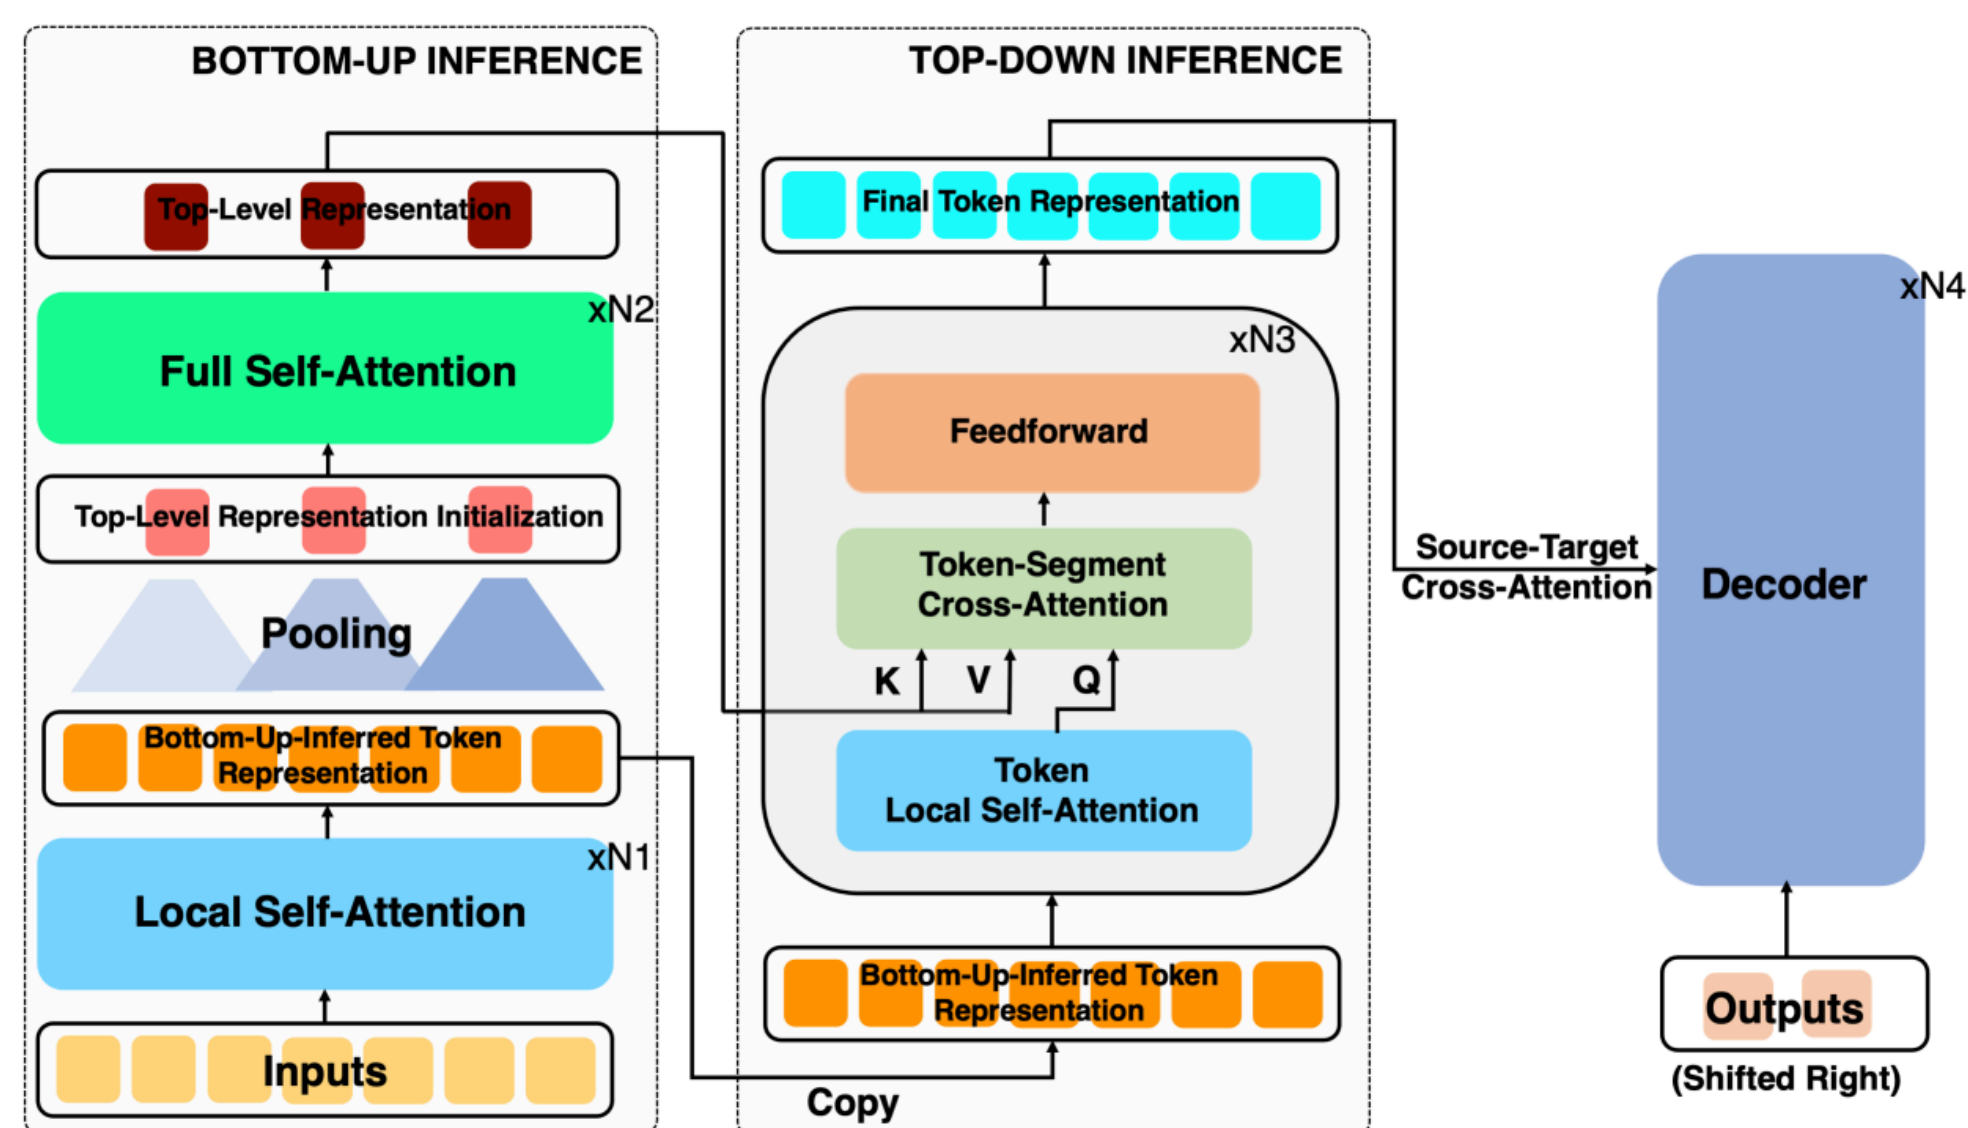
3. Hierarchical models

- *Long Document Summarization with Top-Down and Bottom-Up Inference*, Pang et al. (2022)
- *A Hierarchical Network for Abstractive Meeting Summarization with Cross-Domain Pretraining*, Zhu et al. (2020)

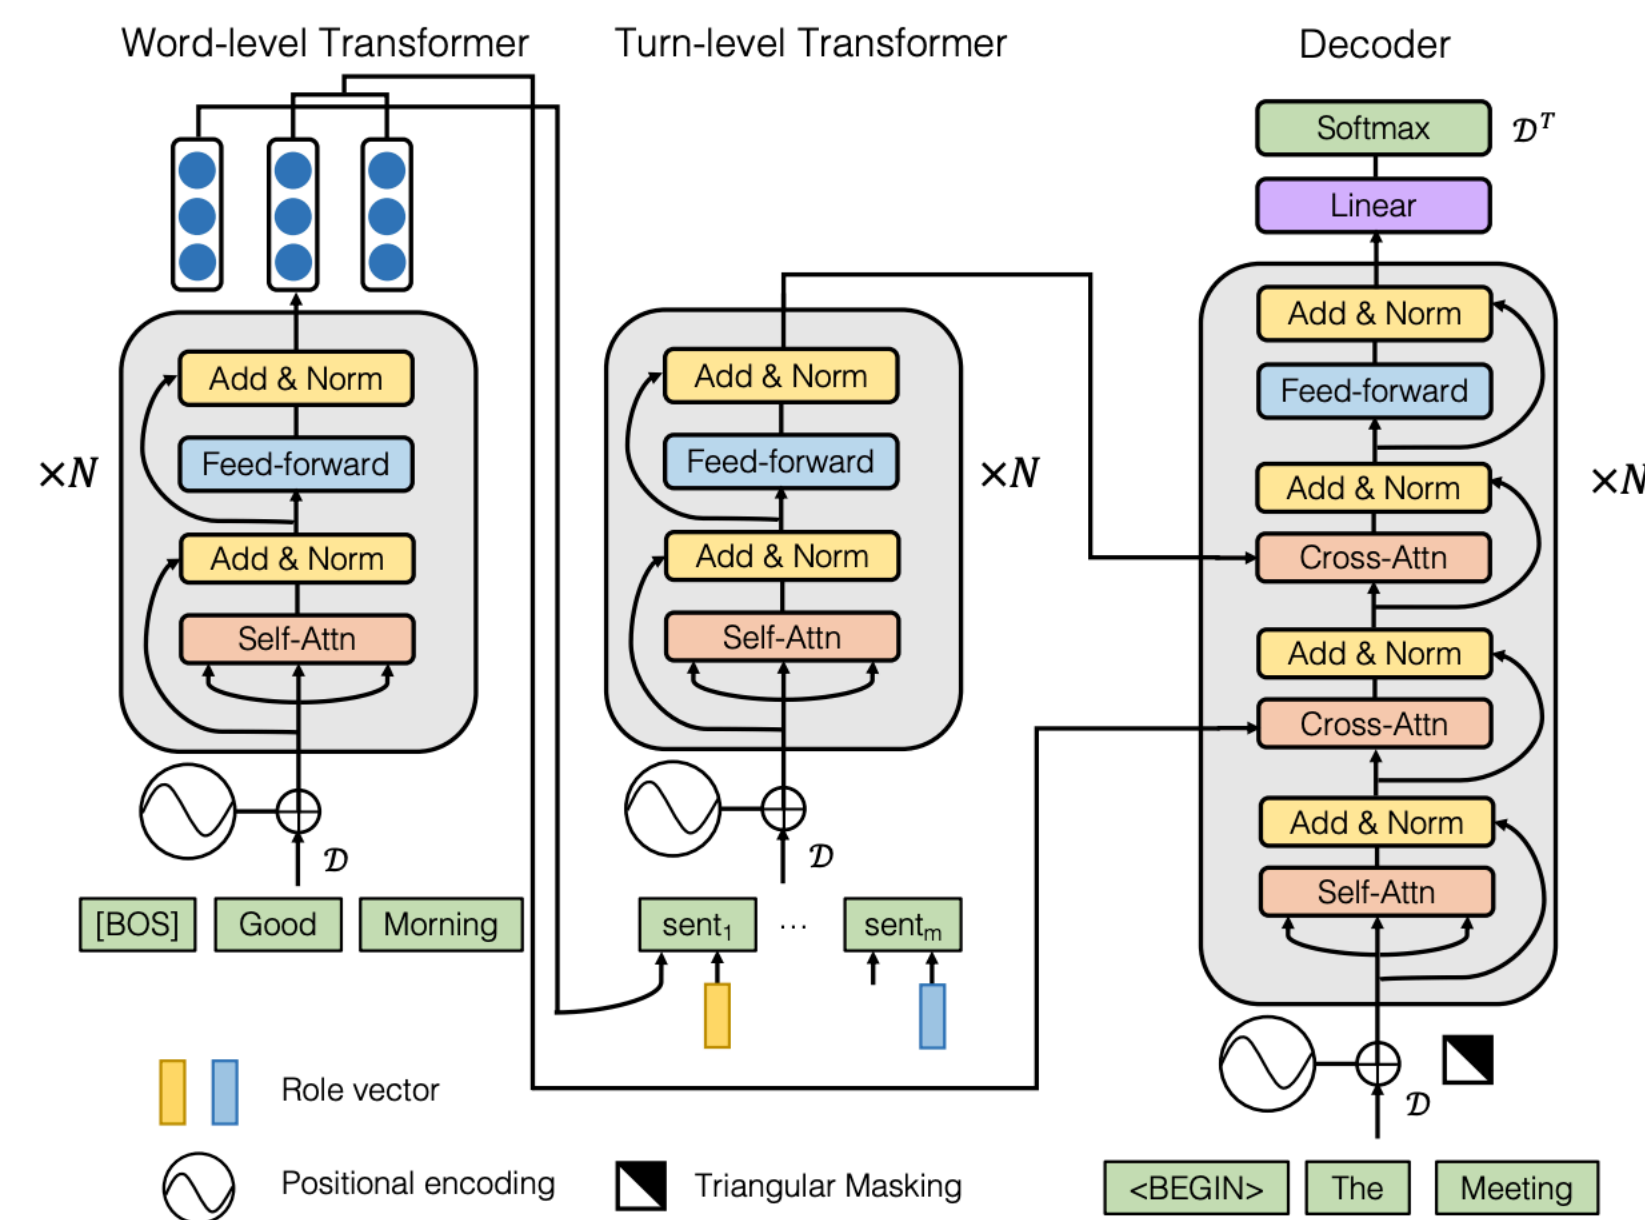
Summarization-Specific Approaches

Hierarchical Models

Long Document Summarization with Top-Down and Bottom-Up Inference, Pang et al. (2022)



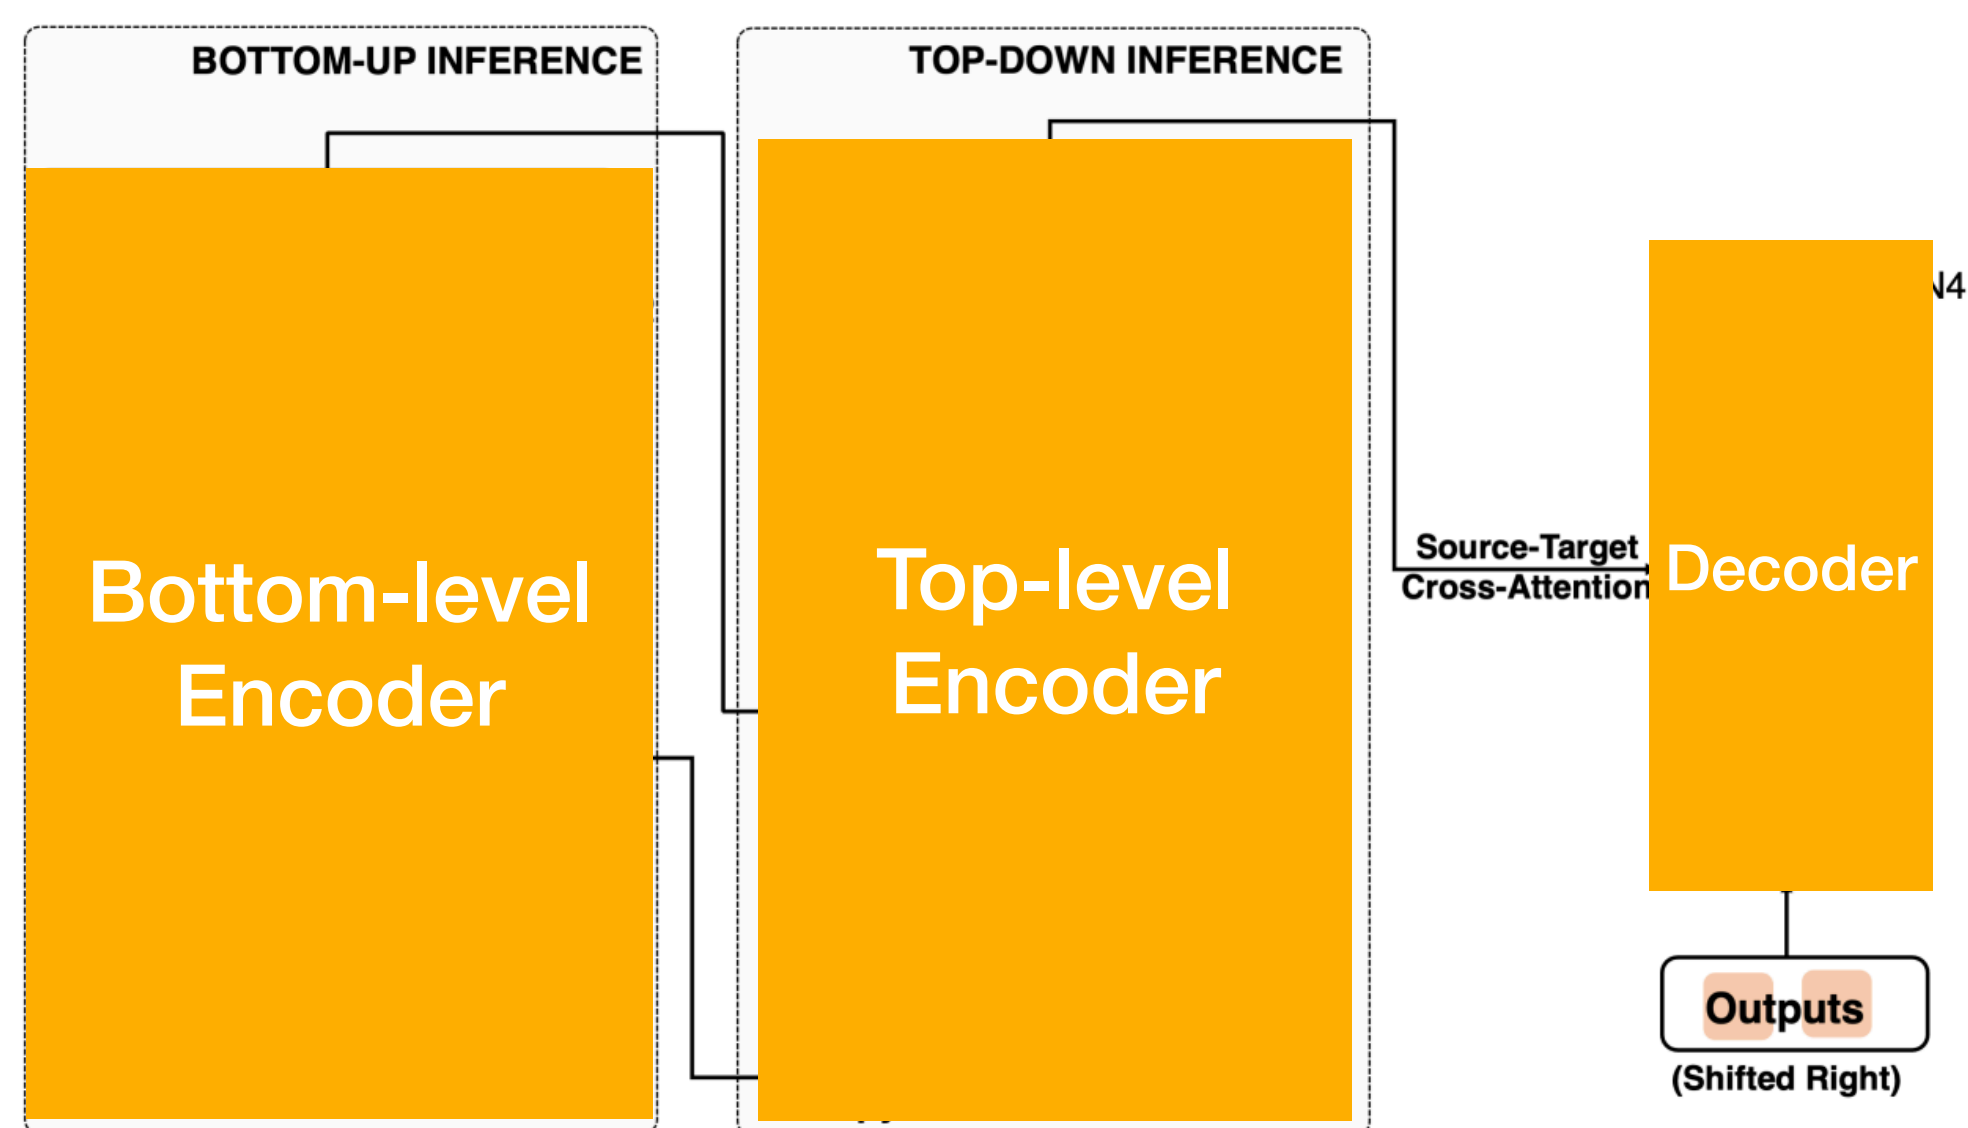
A Hierarchical Network for Abstractive Meeting Summarization with Cross-Domain Pretraining, Zhu et al. (2020)



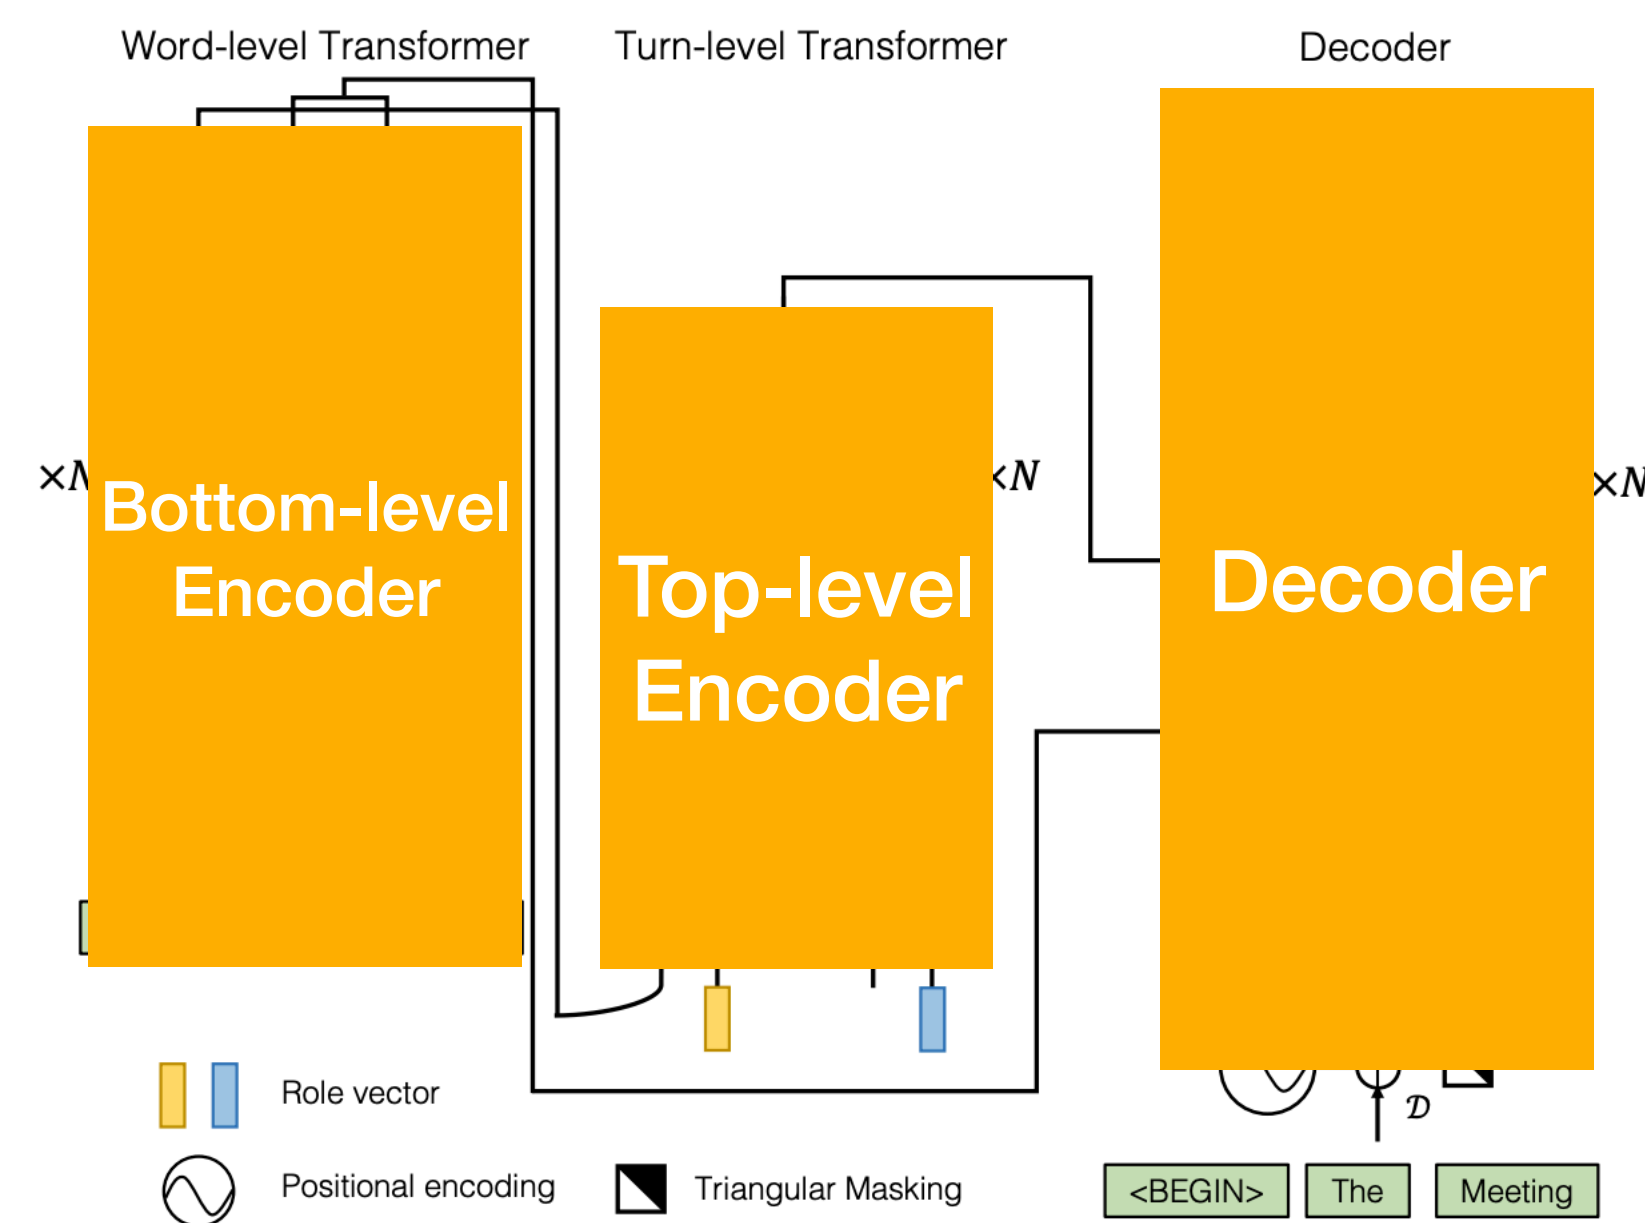
Summarization-Specific Approaches

Hierarchical Models

Long Document Summarization with Top-Down and Bottom-Up Inference, Pang et al. (2022)



A Hierarchical Network for Abstractive Meeting Summarization with Cross-Domain Pretraining, Zhu et al. (2020)



Summarization-Specific Approaches

1. Extract (-then-generate)

- *DYLE: Dynamic Latent Extraction for Abstractive Long-Input Summarization*, Mao et al. (2022)
- *Discourse-Aware Unsupervised Summarization of Long Scientific Documents*, Dong et al. (2021)

2. Divide-and-conquer

- *A Divide-and-Conquer Approach to the Summarization of Long Documents*, Gidiotis et al. (2020)

3. Hierarchical models

- *Long Document Summarization with Top-Down and Bottom-Up Inference*, Pang et al. (2022)
- *A Hierarchical Network for Abstractive Meeting Summarization with Cross-Domain Pretraining*, Zhu et al. (2020)

Summarization-Specific Approaches

BookSum Book Dataset

	R-1	R-2	R-L
Extractive Oracle	46.62	9.17	18.31
BART	29.97	6.02	10.97
BookSum , Kryściński et al. (2021) T5	39.46	7.69	13.77
Pegasus	35.29	6.79	12.71
175B full tree RL	41.51	10.46	16.88
Recursively Summarizing Books with Human Feedback , Wu et al. (2021) 175B first subtree RL	43.19	10.63	17.10
6B full tree RL	36.79	7.22	14.84
Top-Down and Bottom-Up Inference , Pang et al. (2022) Top Down Transformer (464M)	44.19	10.89	16.13

Long Document Summarization with Top-Down and Bottom-Up Inference, Pang et al. (2022)

Summarization-Specific Approaches

BookSum Book Dataset

	R-1	R-2	R-L	
	46.62	9.17	18.31	
	29.97	6.02	10.97	
BookSum , Kryściński et al. (2021)	39.46	7.69	13.77	
	35.29	6.79	12.71	
	41.51	10.46	16.88	
Recursively Summarizing Books with Human Feedback , Wu et al. (2021)	43.19	10.63	17.10	
	36.79	7.22	14.84	
Top-Down and Bottom-Up Inference , Pang et al. (2022)	44.19	10.89	16.13	<1% of the parameters

Long Document Summarization with Top-Down and Bottom-Up Inference, Pang et al. (2022)

Summarization-Specific Approaches

Dorian Gray is sitting in the garden of his uncle's house in London, playing the piano. In the center of the room is a portrait of a young man of extraordinary beauty, and in front of it is Basil Hallward's portrait of Dorian Gray. The portrait belongs to Hallward, and he refuses to show it at the Grosvenor Academy, because it contains too much of himself in it. He tells Dorian's uncle, Lord Henry Wotton, that he is in love with Dorian and wants to marry him. Dorian is sitting at the piano and finishes a painting that he has been working on. He looks at the portrait and realizes that it is of him. He realizes that he will never be able to look at it again, but he cannot stop looking at it because he wants to remember the beauty of his own face. He asks his uncle for some information about Dorian, and his uncle tells him that Dorian has a beautiful mother who was married to a poor man who was killed in a duel. She left him a son, who is very good-looking and who has inherited all of her property. Lord Henry tells him to write to him and ask for some advice, and Dorian agrees. One day, Dorian meets Sibyl Vane, a beautiful young woman who works as a governess for a rich family in the East End of London. She is in the employ of Lord Henry's friend, Mr. Erskine of Treadley, and Lord Henry wants to see her. He also wants to get her out of the hands of the Jew who has her bound to him for three years and eight months. He proposes to her, but she refuses him. She says that she does not think he is good enough for her, and she will never love anyone of his rank. He is disappointed, but does not say anything to his mother about it. The next day, he meets the Duchess of Monmouth, who tells him he should find a wife and marry her. She wants him to have a future and not to spend his money frivolously. He agrees, but when he tells her that he does not love her, she laughs at him and refuses to call him by his new name, Prince Charming. He goes to see the play, and is horrified to see that the face on the canvas is that of the portrait of Romeo and Juliet. He cannot believe that he could have done such a terrible thing to Juliet and that she could still be his wife. He leaves the theater and wanders the streets of London until he finds himself in Covent Garden. He finds some women waiting for him, and one of them laughs when he calls her by his nickname, "Prince Charming." She curses him and runs away. He runs into a dark alley and is suddenly grabbed by a man with a gun pointed at his head. It is James Vane. Vane threatens to kill Dorian if he doesn't make peace with God. He gives Dorian one minute to make his peace before he kills him. When Dorian gets to the street, he finds that the man he was trying to kill is not the same man he thought he was. It turns out that Vane is twenty-eight years younger than Dorian. The woman who took his money tells him not to talk to her again. She runs off, and when Dorian looks back, the woman has disappeared. When he wakes up the next morning, he has not had a nightmare. He writes two letters to his assistant, Alan Campbell, telling him that there is a dead man sitting on a table in his house, and that he must destroy the body so that no one will ever know who he is. He then goes to his bedroom and finds a small box of lacquer, which he takes out and puts inside. He puts the box back, gets into a horse-drawn carriage, and gives the driver an address. The driver takes him to the address, and as he is leaving the house, he sees the dead body of a man on the table. When Campbell returns, he tells Alan not to disturb the body, but to come back at seven o'clock in the evening. When the man arrives, he throws the picture over the table, but Dorian does not believe that it has been disturbed. He returns home and finds that Campbell has brought back the chemicals and the irons, and the other things that he needs to do the job. He opens the cabinet where he had hidden Basil's coat and bag, and finds the green paste. At midnight, he gets a hansom and leaves the house with the instructions to meet him at 7 o'clock the next day. He sits in the back of the carriage as the driver drives him through the streets. He wonders if it is possible to cure the soul by means of the senses and the body by way of the soul. He wakes up in the middle of the night to find that the portrait has not changed.

Long Document Summarization with Top-Down and Bottom-Up Inference, Pang et al. (2022)

Summarization-Specific Approaches

Dorian Gray is sitting in the garden of his uncle's house in London, playing the piano. In the center of the room is a portrait of a young man of extraordinary beauty,

and in front of it is Basil Hallward's portrait, which is too much of himself in it. He tells Dorian that he has finished a painting that he has been working on for years, and cannot stop looking at it because he was so beautiful. Dorian has a beautiful mother who was the property of Lord Henry. Lord Henry tells him to write as a governess for a rich family in the country. He also wants to get her out of the hands of the Duke. She does not think he is good enough for her.

Best 175B RL summary

Barack Obama wins the presidency in 2008, and his first term is marked by economic turmoil and political polarization. His domestic policy includes the Affordable Care Act, which reformed healthcare in the United States.

Obama's presidency includes terrorism, the Arab-Israeli conflict, and the Deepwater Horizon oil spill. His foreign policy includes finding and eliminating Osama bin Laden. After the successful Abbottabad raid, Obama reflects on the unity of effort displayed by all the individuals involved in the operation. He wonders if that same unity of effort could be applied to other areas of government.

One day, he meets the Duchess of Monmouth, who tells him he should find a wife and marry her. She wants him to have a future and not to spend his money frivolously. He agrees, but when he tells her that he does not love her, she laughs at him and refuses to call him by his new name, Prince Charming. He goes to see the play, and is horrified to see that the face on the canvas is that of the portrait of Romeo and Juliet. He cannot believe that he could have done such a terrible thing to Juliet and that she could still be his wife. He leaves the theater and wanders the streets of London until he finds himself in Covent Garden. He finds some women waiting for him, and one of them laughs when he calls her by his nickname, "Prince Charming." She curses him and runs away. He runs into a dark alley and is suddenly grabbed by a man with a gun pointed at his head. It is James Vane. Vane threatens to kill Dorian if he doesn't make peace with God. He gives Dorian one minute to make his peace before he kills him. When Dorian gets to the street, he finds that the man he was trying to kill is not the same man he thought he was. It turns out that Vane is twenty-eight years younger than Dorian. The woman who took his money tells him not to talk to her again. She runs off, and when Dorian looks back, the woman has disappeared. When he wakes up the next morning, he has not had a nightmare. He writes two letters to his assistant, Alan Campbell, telling him that there is a dead man sitting on a table in his house, and that he must destroy the body so that no one will ever know who he is. He then goes to his bedroom and finds a small box of lacquer, which he takes out and puts inside. He puts the box back, gets into a horse-drawn carriage, and gives the driver an address. The driver takes him to the address, and as he is leaving the house, he sees the dead body of a man on the table. When Campbell returns, he tells Alan not to disturb the body, but to come back at seven o'clock in the evening. When the man arrives, he throws the picture over the table, but Dorian does not believe that it has been disturbed. He returns home and finds that Campbell has brought back the chemicals and the irons, and the other things that he needs to do the job. He opens the cabinet where he had hidden Basil's coat and bag, and finds the green paste. At midnight, he gets a hansom and leaves the house with the instructions to meet him at 7 o'clock the next day. He sits in the back of the carriage as the driver drives him through the streets. He wonders if it is possible to cure the soul by means of the senses and the body by way of the soul. He wakes up in the middle of the night to find that the portrait has not changed.

Long Document Summarization with Top-Down and Bottom-Up Inference, Pang et al. (2022)

Summarization-Specific Approaches

Some good details, correct characters

Dorian Gray is sitting in the garden of his uncle's house in London, playing the piano. In the center of the room is a portrait of a young man of extraordinary beauty, and in front of it is Basil Hallward's portrait of Dorian Gray. The portrait belongs to Hallward, and he refuses to show it at the Grosvenor Academy, because it contains too much of himself in it. He tells Dorian's uncle, Lord Henry Wotton, that he is in love with Dorian and wants to marry him. Dorian is sitting at the piano and finishes a painting that he has been working on. He looks at the portrait and realizes that it is of him. He realizes that he will never be able to look at it again, but he cannot stop looking at it because he wants to remember the beauty of his own face. He asks his uncle for some information about Dorian, and his uncle tells him that Dorian has a beautiful mother who was married to a poor man who was killed in a duel. She left him a son, who is very good-looking and who has inherited all of her property. Lord Henry tells him to write to him and ask for some advice, and Dorian agrees. One day, Dorian meets Sibyl Vane, a beautiful young woman who works as a governess for a rich family in the East End of London. She is in the employ of Lord Henry's friend, Mr. Erskine of Treadley, and Lord Henry wants to see her. He also wants to get her out of the hands of the Jew who has her bound to him for three years and eight months. He proposes to her, but she refuses him. She says that she does not think he is good enough for her, and she will never love anyone of his rank. He is disappointed, but does not say anything to his mother about it. The next day, he meets the Duchess of Monmouth, who tells him he should find a wife and marry her. She wants him to have a future and not to spend his money frivolously. He agrees, but when he tells her that he does not love her, she laughs at him and refuses to call him by his new name, Prince Charming. He goes to see the play, and is horrified to see that the face on the canvas is that of the portrait of Romeo and Juliet. He cannot believe that he could have done such a terrible thing to Juliet and that she could still be his wife. He leaves the theater and wanders the streets of London until he finds himself in Covent Garden. He finds some women waiting for him, and one of them laughs when he calls her by his nickname, "Prince Charming." She curses him and runs away. He runs into a dark alley and is suddenly grabbed by a man with a gun pointed at his head. It is James Vane. Vane threatens to kill Dorian if he doesn't make peace with God. He gives Dorian one minute to make his peace before he kills him. When Dorian gets to the street, he finds that the man he was trying to kill is not the same man he thought he was. It turns out that Vane is twenty-eight years younger than Dorian. The woman who took his money tells him not to talk to her again. She runs off, and when Dorian looks back, the woman has disappeared. When he wakes up the next morning, he has not had a nightmare. He writes two letters to his assistant, Alan Campbell, telling him that there is a dead man sitting on a table in his house, and that he must destroy the body so that no one will ever know who he is. He then goes to his bedroom and finds a small box of lacquer, which he takes out and puts inside. He puts the box back, gets into a horse-drawn carriage, and gives the driver an address. The driver takes him to the address, and as he is leaving the house, he sees the dead body of a man on the table. When Campbell returns, he tells Alan not to disturb the body, but to come back at seven o'clock in the evening. When the man arrives, he throws the picture over the table, but Dorian does not believe that it has been disturbed. He returns home and finds that Campbell has brought back the chemicals and the irons, and the other things that he needs to do the job. He opens the cabinet where he had hidden Basil's coat and bag, and finds the green paste. At midnight, he gets a hansom and leaves the house with the instructions to meet him at 7 o'clock the next day. He sits in the back of the carriage as the driver drives him through the streets. He wonders if it is possible to cure the soul by means of the senses and the body by way of the soul. He wakes up in the middle of the night to find that the portrait has not changed.

Long Document Summarization with Top-Down and Bottom-Up Inference, Pang et al. (2022)

Summarization-Specific Approaches

Dorian Gray is sitting in the garden of his uncle's house in London, playing the piano. In the center of the room is a portrait of a young man of extraordinary beauty, and in front of it is Basil Hallward's portrait of Dorian Gray. The portrait belongs to Hallward, and he refuses to show it at the Grosvenor Academy, because it contains too much of himself in it. He tells Dorian's uncle, Lord Henry Wotton, that he is in love with Dorian and wants to marry him. Dorian is sitting at the piano and finishes a painting that he has been working on. He looks at the portrait and realizes that it is of him. He realizes that he will never be able to look at it again, but he cannot stop looking at it because he wants to remember the beauty of his own face. He asks his uncle for some information about Dorian, and his uncle tells him that Dorian has a beautiful mother who was married to a poor man who was killed in a duel. She left him a son, who is very good-looking and who has inherited all of her property. Lord Henry tells him to write to him and ask for some advice, and Dorian agrees. One day, Dorian meets Sibyl Vane, a beautiful young woman who works as a governess for a rich family in the East End of London. She is in the employ of Lord Henry's friend, Mr. Erskine of Treadley, and Lord Henry wants to see her. He also wants to get her out of the hands of the Jew who has her bound to him for three years and eight months. He proposes to her, but she refuses him. She says that she does not think he is good enough for her, and she will never love anyone of his rank. He is disappointed, but does not say anything to his mother about it. The next day, he meets the Duchess of Monmouth, who tells him he should find a wife and marry her. She wants him to have a future and not to spend his money frivolously. He agrees, but when he tells her that he does not love her, she laughs at him and refuses to call him by his new name, Prince Charming. He goes to see the play, and is horrified to see that the face on the canvas is that of the portrait of Romeo and Juliet. He cannot believe that he could have done such a terrible thing to Juliet and that she could still be his wife. He leaves the theater and wanders the streets of London until he finds himself in Covent Garden. He finds some women waiting for him, and one of them laughs when he calls her by his nickname, "Prince Charming." She curses him and runs away. He runs into a dark alley and is suddenly grabbed by a man with a gun pointed at his head. It is James Vane. Vane threatens to kill Dorian if he doesn't make peace with God. He gives Dorian one minute to make his peace before he kills him. When Dorian gets to the street, he finds that the man he was trying to kill is not the same man he thought he was. It turns out that Vane is twenty-eight years younger than Dorian. The woman who took his money tells him not to talk to her again. She runs off, and when Dorian looks back, the woman has disappeared. When he wakes up the next morning, he has not had a nightmare. He writes two letters to his assistant, Alan Campbell, telling him that there is a dead man sitting on a table in his house, and that he must destroy the body so that no one will ever know who he is. He then goes to his bedroom and finds a small box of lacquer, which he takes out and puts inside. He puts the box back, gets into a horse-drawn carriage, and gives the driver an address. The driver takes him to the address, and as he is leaving the house, he sees the dead body of a man on the table. When Campbell returns, he tells Alan not to disturb the body, but to come back at seven o'clock in the evening. When the man arrives, he throws the picture over the table, but Dorian does not believe that it has been disturbed. He returns home and finds that Campbell has brought back the chemicals and the irons, and the other things that he needs to do the job. He opens the cabinet where he had hidden Basil's coat and bag, and finds the green paste. At midnight, he gets a hansom. **Inaccuracies, Dorian should be dead and portrait changed** back of the carriage as the driver drives him through the streets. He wonders if it is possible to cure the soul by means of the senses and the body by way of the soul. He wakes up in the middle of the night to find that the portrait has not changed.

Long Document Summarization with Top-Down and Bottom-Up Inference, Pang et al. (2022)

Summarization-Specific Approaches

Dorian Gray is sitting in the garden of his uncle's house in London, playing the piano. In the center of the room is a portrait of a young man of extraordinary beauty, and in front of it is Basil Hallward's portrait of Dorian Gray. The portrait belongs to Hallward, and he refuses to show it at the Grosvenor Academy, because it contains too much of himself in it. He tells Dorian's uncle, Lord Henry Wotton, that he is in love with Dorian and wants to marry him. Dorian is sitting at the piano and finishes a painting that he has been working on. He looks at the portrait and realizes that it is of him. He realizes that he will never be able to look at it again, but he cannot stop looking at it because he wants to remember the **Incoherence**. He asks his uncle for some information about Dorian, and his uncle tells him that Dorian has a beautiful mother who was married to a poor man who was killed in a duel. She left him a son, who is very good-looking and who has inherited all of her property. Lord Henry tells him to write to him and ask for some advice, and Dorian agrees. One day, Dorian meets Sibyl Vane, a beautiful young woman who works as a governess for a rich family in the East End of London. She is in the employ of Lord Henry's friend, Mr. Erskine of Treadley, and Lord Henry wants to see her. He also wants to get her out of the hands of the Jew who has her bound to him for three years and eight months. He proposes to her, but she refuses him. She says that she does not think he is good enough for her, and she will never love anyone of his rank. He is disappointed, but does not say anything to his mother about it. The next day, he meets the Duchess of Monmouth, who tells him he should find a wife and marry her. She wants him to have a future and not to spend his money frivolously. He agrees, but when he tells her that he does not love her, she laughs at him and refuses to call him by his new name, Prince Charming. He goes to see the play, and is horrified to see that the face on the canvas is that of the portrait of Romeo and Juliet. He cannot believe that he could have done such a terrible thing to Juliet and that she could still be his wife. He leaves the theater and wanders the streets of London until he finds himself in Covent Garden. He finds some women waiting for him, and one of them laughs when he calls her by his nickname, "Prince Charming." She curses him and runs away. He runs into a dark alley and is suddenly grabbed by a man with a gun pointed at his head. It is James Vane. Vane threatens to kill Dorian if he doesn't make peace with God. He gives Dorian one minute to make his peace before he kills him. When Dorian gets to the street, he finds that the man he was trying to kill is not the same man he thought he was. It turns out that Vane is twenty-eight years younger than Dorian. The woman who took his money tells him not to talk to her again. She runs off, and when Dorian looks back, the woman has disappeared. When he wakes up the next morning, he has not had a nightmare. He writes two letters to his assistant, Alan Campbell, telling him that there is a dead man sitting on a table in his house, and that he must destroy the body so that no one will ever know who he is. He then goes to his bedroom and finds a small box of lacquer, which he takes out and puts inside. He puts the box back, gets into a horse-drawn carriage, and gives the driver an address. The driver takes him to the address, and as he is leaving the house, he sees the dead body of a man on the table. When Campbell returns, he tells Alan not to disturb the body, but to come back at seven o'clock in the evening. When the man arrives, he throws the picture over the table, but Dorian does not believe that it has been disturbed. He returns home and finds that Campbell has brought back the chemicals and the irons, and the other things that he needs to do the job. He opens the cabinet where he had hidden Basil's coat and bag, and finds the green paste. At midnight, he gets a hansom and leaves the house with the instructions to meet him at 7 o'clock the next day. He sits in the back of the carriage as the driver drives him through the streets. He wonders if it is possible to cure the soul by means of the senses and the body by way of the soul. He wakes up in the middle of the night to find that the portrait has not changed.

Long Document Summarization with Top-Down and Bottom-Up Inference, Pang et al. (2022)

Summary - We can't handle long documents well yet

Efficient attention approaches...

- More computationally intense but allow for **benefits of scale**
- Input **more of the entire document** which improves summarization performance
- Still far **short of book-length**

Summary - We can't handle long documents well yet

Efficient attention approaches...

- More computationally intense but allow for **benefits of scale**
- Input **more of the entire document** which improves summarization performance
- Still far **short of book-length**

Summarization-specific approaches...

- **Build on attention-based models**
- **Incorporate helpful structure** which encourages inclusion of important details across narrative
- Work flexibly with **any length text**

Summary - We can't handle long documents well yet

Efficient attention approaches...

- More computationally intense but allow for **benefits of scale**
- Input **more of the entire document** which improves summarization performance
- Still far **short of book-length**

Summarization-specific approaches...

- **Build on attention-based** models
- **Incorporate helpful structure** which encourages inclusion of important details across narrative
- Work flexibly with **any length text**

For the foreseeable future, **we need both** attention-efficient and summarization-specific approaches

1. Narrative Summarization

What is narrative summarization and why is it important?

2. Long Documents

How can we summarize very long narratives?

3. Controllability

How do we target summaries to specific tasks/users?

4. Evaluation

Can we automatically evaluate how well a system works?

Why do we need targeted summaries?

“A Promised Land” by Barack Obama

Barack Obama wins the presidency in 2008, and his first term is marked by economic turmoil and political polarization. His domestic policy includes the Affordable Care Act, which reformed healthcare in the United States.

Obama’s presidency includes terrorism, the Arab-Israeli conflict, and the Deepwater Horizon oil spill. His foreign policy includes finding and eliminating Osama bin Laden. After the successful Abbottabad raid, Obama reflects on the unity of effort displayed by all the individuals involved in the operation. He wonders if that same unity of effort could be applied to other areas of government.

Why do we need targeted summaries?

“A Promised Land” by Barack Obama

Assumes ACA = reform

Barack Obama wins the presidency in 2008, and his first term is marked by economic turmoil and political polarization. His domestic policy includes the Affordable Care Act, which reformed healthcare in the United States.

Obama’s presidency includes terrorism, the Arab-Israeli conflict, and the Deepwater Horizon oil spill. His foreign policy includes finding and eliminating Osama bin Laden. After the successful Abbottabad raid, Obama reflects on the unity of effort displayed by all the individuals involved in the operation. He wonders if that same unity of effort could be applied to other areas of government.

Why do we need targeted summaries?

“A Promised Land” by Barack Obama

Barack Obama wins the presidency in 2008, and his first term is marked by economic turmoil and political polarization. His domestic policy includes the Affordable Care Act, which reformed healthcare in the United States.

Obama’s presidency includes terrorism, the Arab-Israeli conflict, and the Deepwater Horizon oil spill. His foreign policy includes finding and eliminating Osama bin Laden. After the successful Abbottabad raid, Obama reflects on the unity of effort displayed by all the individuals involved in the operation. He wonders if that same unity of effort could be applied to other areas of government.

No further details on these issues

Why do we need targeted summaries?

“A Promised Land” by Barack Obama

Barack Obama wins the presidency in 2008, and his first term is marked by economic turmoil and political polarization. His domestic policy includes the Affordable Care Act, which reformed healthcare in the United States.

Obama’s presidency includes terrorism, the Arab-Israeli conflict, and the Deepwater Horizon oil spill. His foreign policy includes finding and eliminating Osama bin Laden. After the successful Abbottabad raid, Obama reflects on the unity of effort displayed by all the individuals involved in the operation. He wonders if that same unity of effort could be applied to other areas of government.

Randomly focuses on Abbottabad raid

Why do we need targeted summaries?

Unlike news, narrative is personal and complex.

Depending on the situation, we may want to...

- Speak to a particular audience (age, beliefs, depth of knowledge, etc.)
- Tradeoff breadth and depth
- Provide all details relevant to a specific question

Types of Control

1. Attribute-based

- *MACSum: Controllable Summarization with Mixed Attributes*, Zhang et al. (2022)
- *HydraSum: Disentangling Style Features in Text Summarization with Multi-Decoder Models*, Goyal et al. (2021)

2. Query-focused

- *Text Summarization with Latent Queries*, Xu et al. (2021)
- *Educational Question Generation of Children Storybooks via Question Type Distribution Learning and Event-Centric Summarization*, Zhao et al. (2022)

3. Reward function

- *Controllable Neural Story Plot Generation via Reward Shaping*, Tambwekar et al. (2019)
- *Controllable Summarization with Constrained Markov Decision Processes*, Chan et al. (2021)
- *Learning to Summarize from Human Feedback*, Stiennon et al. (2020)

Types of Control

1. Attribute-based **Style**

- *MACSum: Controllable Summarization with Mixed Attributes*, Zhang et al. (2022)
- *HydraSum: Disentangling Style Features in Text Summarization with Multi-Decoder Models*, Goyal et al. (2021)

2. Query-focused **Content**

- *Text Summarization with Latent Queries*, Xu et al. (2021)
- *Educational Question Generation of Children Storybooks via Question Type Distribution Learning and Event-Centric Summarization*, Zhao et al. (2022)

3. Reward function **General Quality**

- *Controllable Neural Story Plot Generation via Reward Shaping*, Tambwekar et al. (2019)
- *Controllable Summarization with Constrained Markov Decision Processes*, Chan et al. (2021)
- *Learning to Summarize from Human Feedback*, Stiennon et al. (2020)

Attribute-Based Control

MACSum, Zhang et al. (2022)

Control through input strategy

HydraSum, Goyal et al. (2021)

Control through output strategy

Attribute-Based Control

MACSum, Zhang et al. (2022)

Control through input strategy

Attributes:

- Specificity
- Length
- Extractiveness
- Topic
- Speaker

HydraSum, Goyal et al. (2021)

Control through output strategy

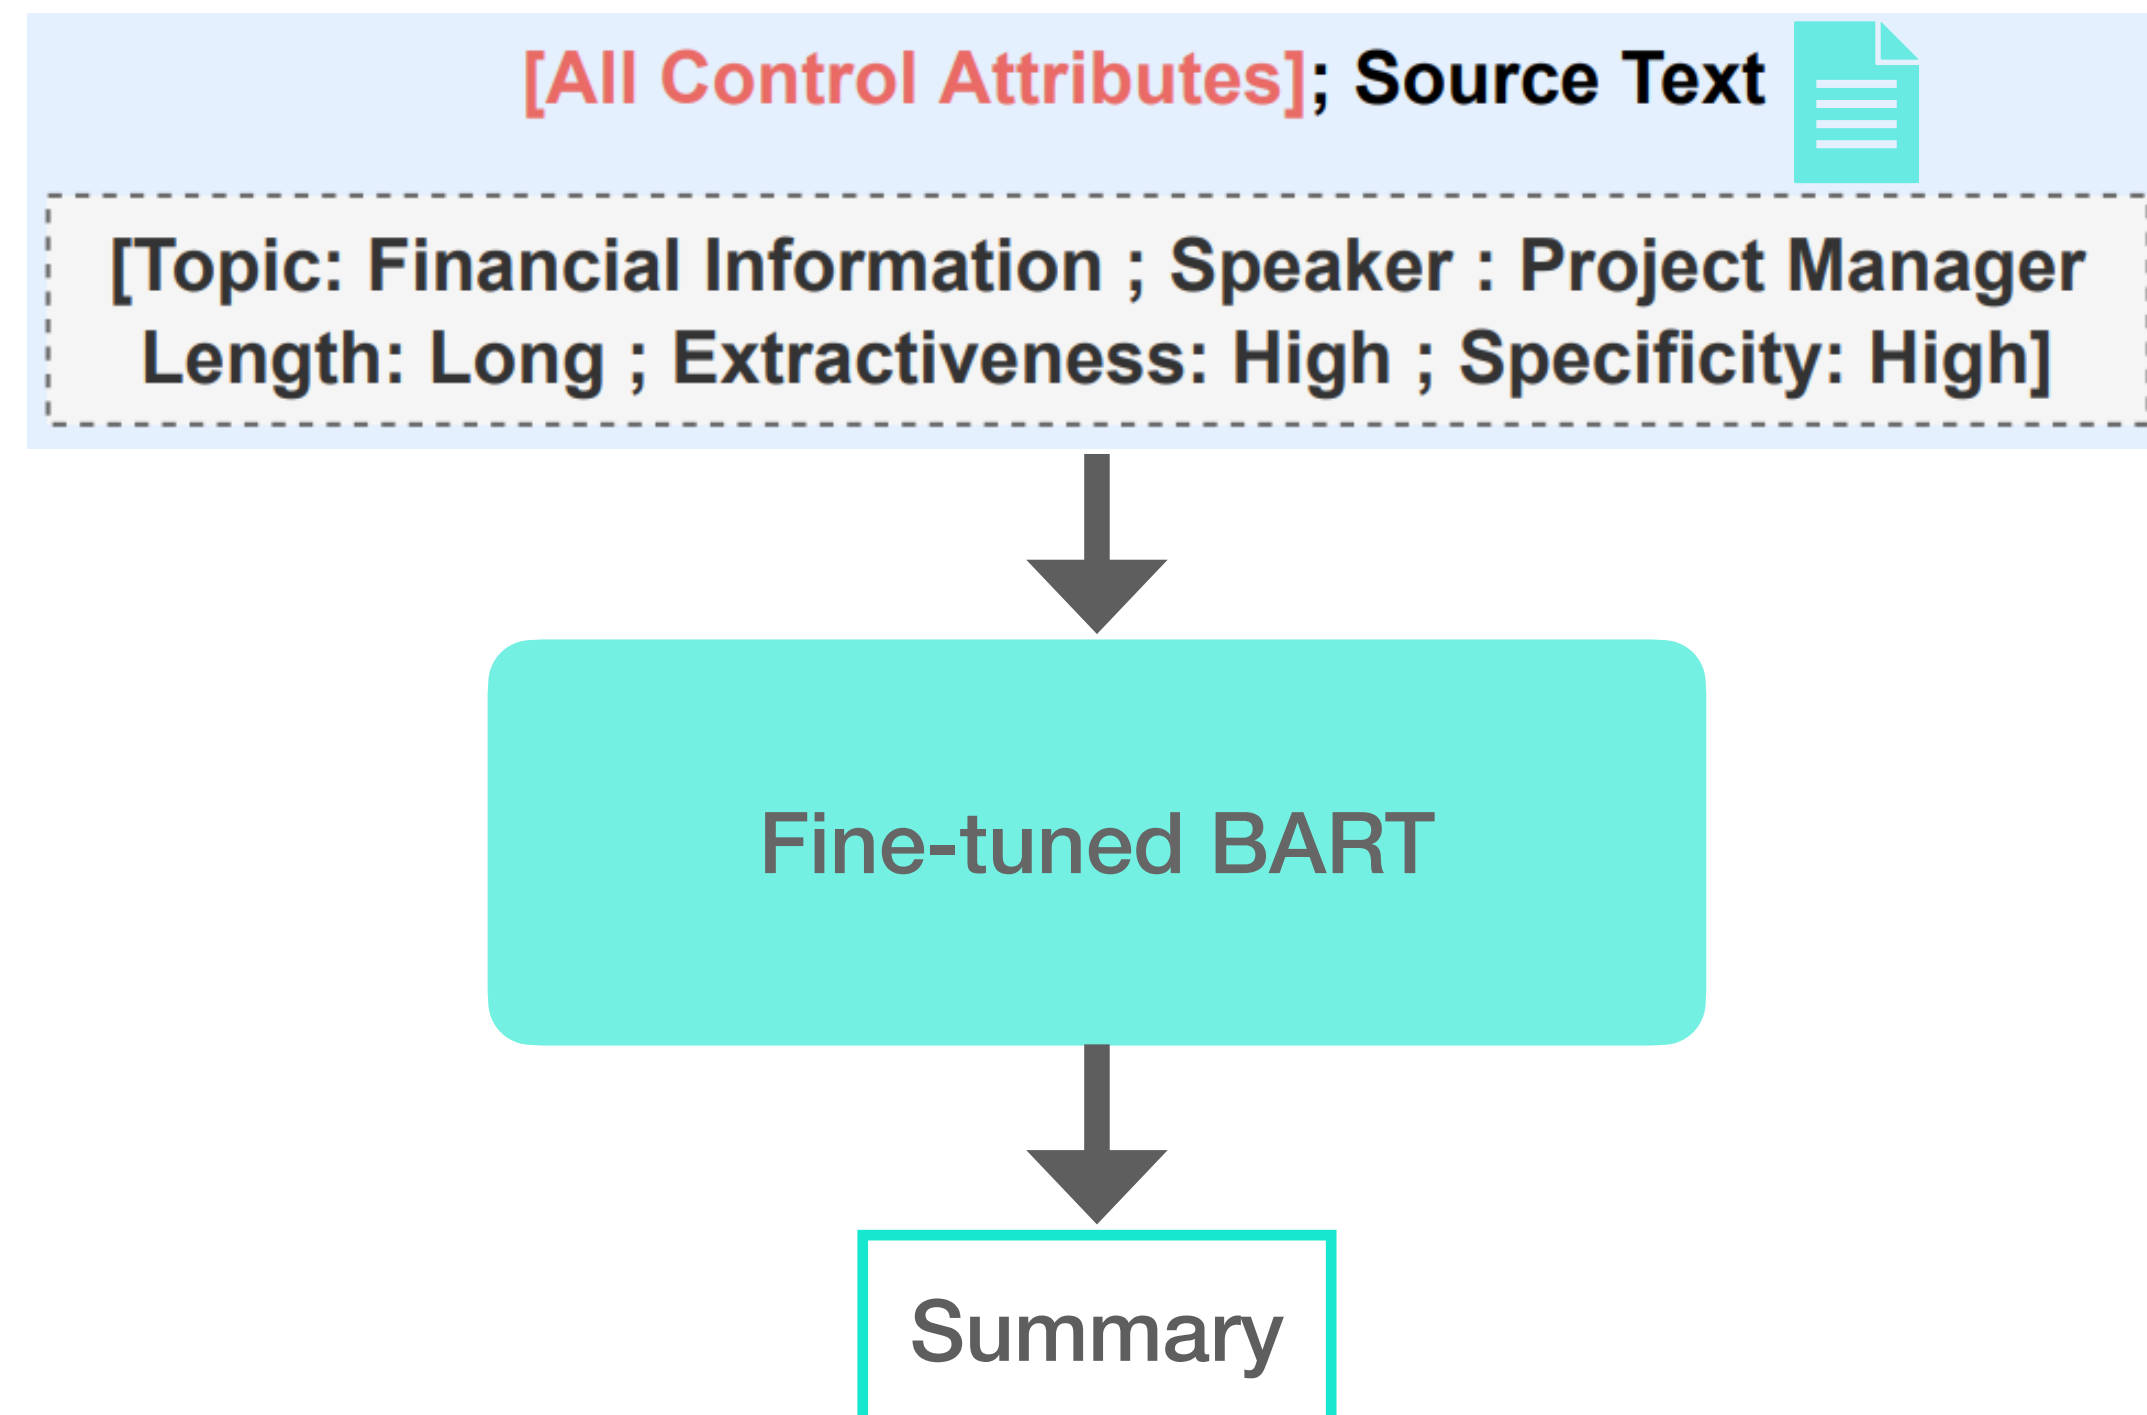
Attributes:

- Specificity
- Length
- Extractiveness
- Readability

Attribute-Based Control

MACSum, Zhang et al. (2022)

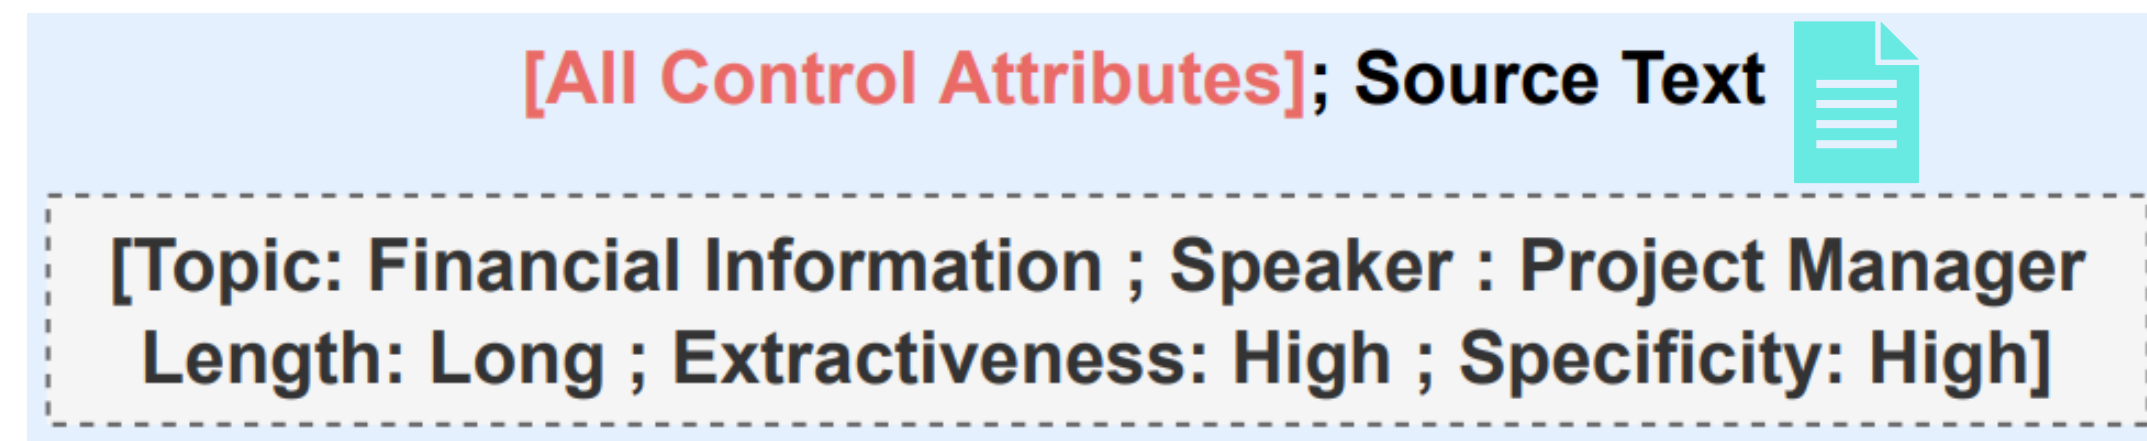
Control through input strategy



Attribute-Based Control

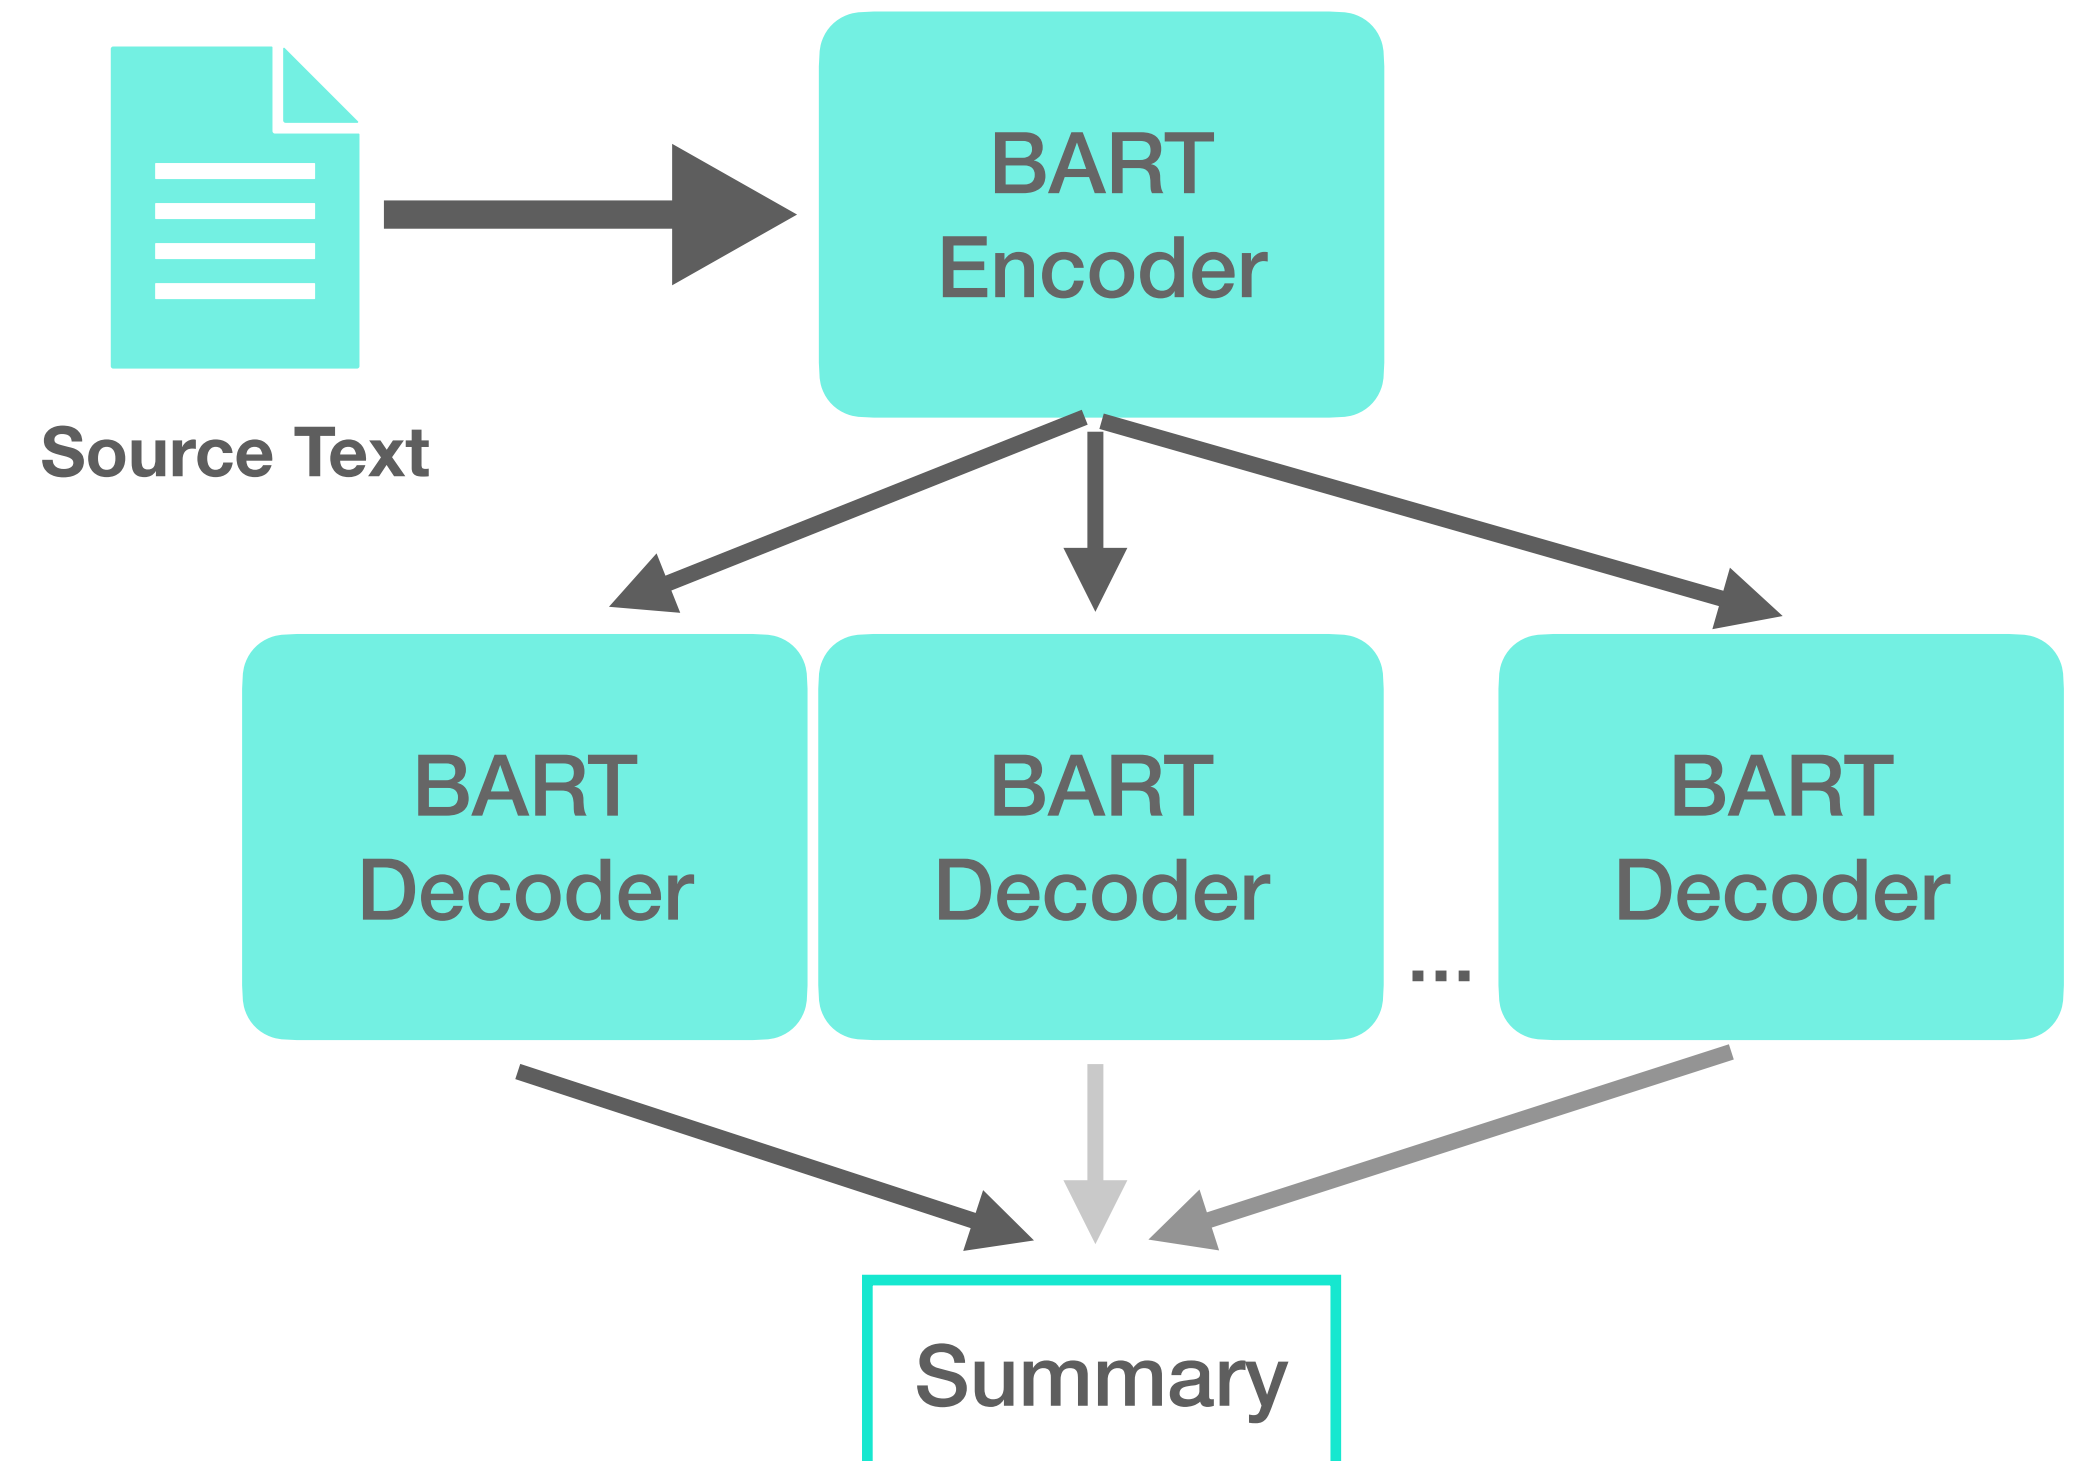
MACSum, Zhang et al. (2022)

Control through input strategy



HydraSum, Goyal et al. (2021)

Control through output strategy



Types of Control

1. Attribute-based

- *MACSum: Controllable Summarization with Mixed Attributes*, Zhang et al. (2022)
- *HydraSum: Disentangling Style Features in Text Summarization with Multi-Decoder Models*, Goyal et al. (2021)

2. Query-focused

- *Text Summarization with Latent Queries*, Xu et al. (2021)
- *Educational Question Generation of Children Storybooks via Question Type Distribution Learning and Event-Centric Summarization*, Zhao et al. (2022)

3. Reward function

- *Controllable Neural Story Plot Generation via Reward Shaping*, Tambwekar et al. (2019)
- *Controllable Summarization with Constrained Markov Decision Processes*, Chan et al. (2021)
- *Learning to Summarize from Human Feedback*, Stiennon et al. (2020)

Query-Focused Approaches

Text Summarization with Latent Queries, Xu et al. (2021)

Query Type	Query Example
Empty	\emptyset
Keywords	<i>Marina Beach, Incidents</i>
Question	<i>Is euthanasia better than withdrawing life support?</i>
Composite	<i>Amnesty International - What is the scope of operations of Amnesty International and what are the international reactions to its activities?</i>
Title	<i>Alzheimer's Disease</i>

Educational Question Generation of Children Storybooks via Question Type Distribution Learning and Event-Centric Summarization, Zhao et al. (2022)

Query-Focused Approaches

Text Summarization with Latent Queries, Xu et al. (2021)

Query Type	Query Example
Empty	\emptyset
Keywords	<i>Marina Beach, Incidents</i>
Question	<i>Is euthanasia better than withdrawing life support?</i>
Composite	<i>Amnesty International - What is the scope of operations of Amnesty International and what are the international reactions to its activities?</i>
Title	<i>Alzheimer's Disease</i>

Educational Question Generation of Children Storybooks via Question Type Distribution Learning and Event-Centric Summarization, Zhao et al. (2022)

P1: Why did the bonze want to get a good price for the pears? (causal relationship) What did the bonze ask for? (action)

P2: What did the Islanders want to express when they were married? (action) Why did the Islanders hold to the belief that Snorro was spirited away? (causal relationship)

Types of Control

1. Attribute-based

- *MACSum: Controllable Summarization with Mixed Attributes*, Zhang et al. (2022)
- *HydraSum: Disentangling Style Features in Text Summarization with Multi-Decoder Models*, Goyal et al. (2021)

2. Query-focused

- *Text Summarization with Latent Queries*, Xu et al. (2021)
- *Educational Question Generation of Children Storybooks via Question Type Distribution Learning and Event-Centric Summarization*, Zhao et al. (2022)

3. Reward function

- *Controllable Neural Story Plot Generation via Reward Shaping*, Tambwekar et al. (2019)
- *Controllable Summarization with Constrained Markov Decision Processes*, Chan et al. (2021)
- *Learning to Summarize from Human Feedback*, Stiennon et al. (2020)

Reward Function Approaches

Step-by-step reward

Controllable Neural Story Plot Generation via Reward Shaping, Tambwekar et al. (2019)

- Reward: **each step, shaped from verb relationships in corpus**

Reward Function Approaches

Step-by-step reward

Controllable Neural Story Plot Generation via Reward Shaping, Tambwekar et al. (2019)

- Reward: **each step, shaped from verb relationships in corpus**

Full summary reward

Controllable Summarization with Constrained Markov Decision Processes, Chan et al. (2021)

- Reward: **BERTScore against reference with attribute constraints**

Reward Function Approaches

Step-by-step reward

Controllable Neural Story Plot Generation via Reward Shaping, Tambwekar et al. (2019)

- Reward: **each step, shaped from verb relationships in corpus**

Full summary reward

Controllable Summarization with Constrained Markov Decision Processes, Chan et al. (2021)

- Reward: **BERTScore against reference with attribute constraints**

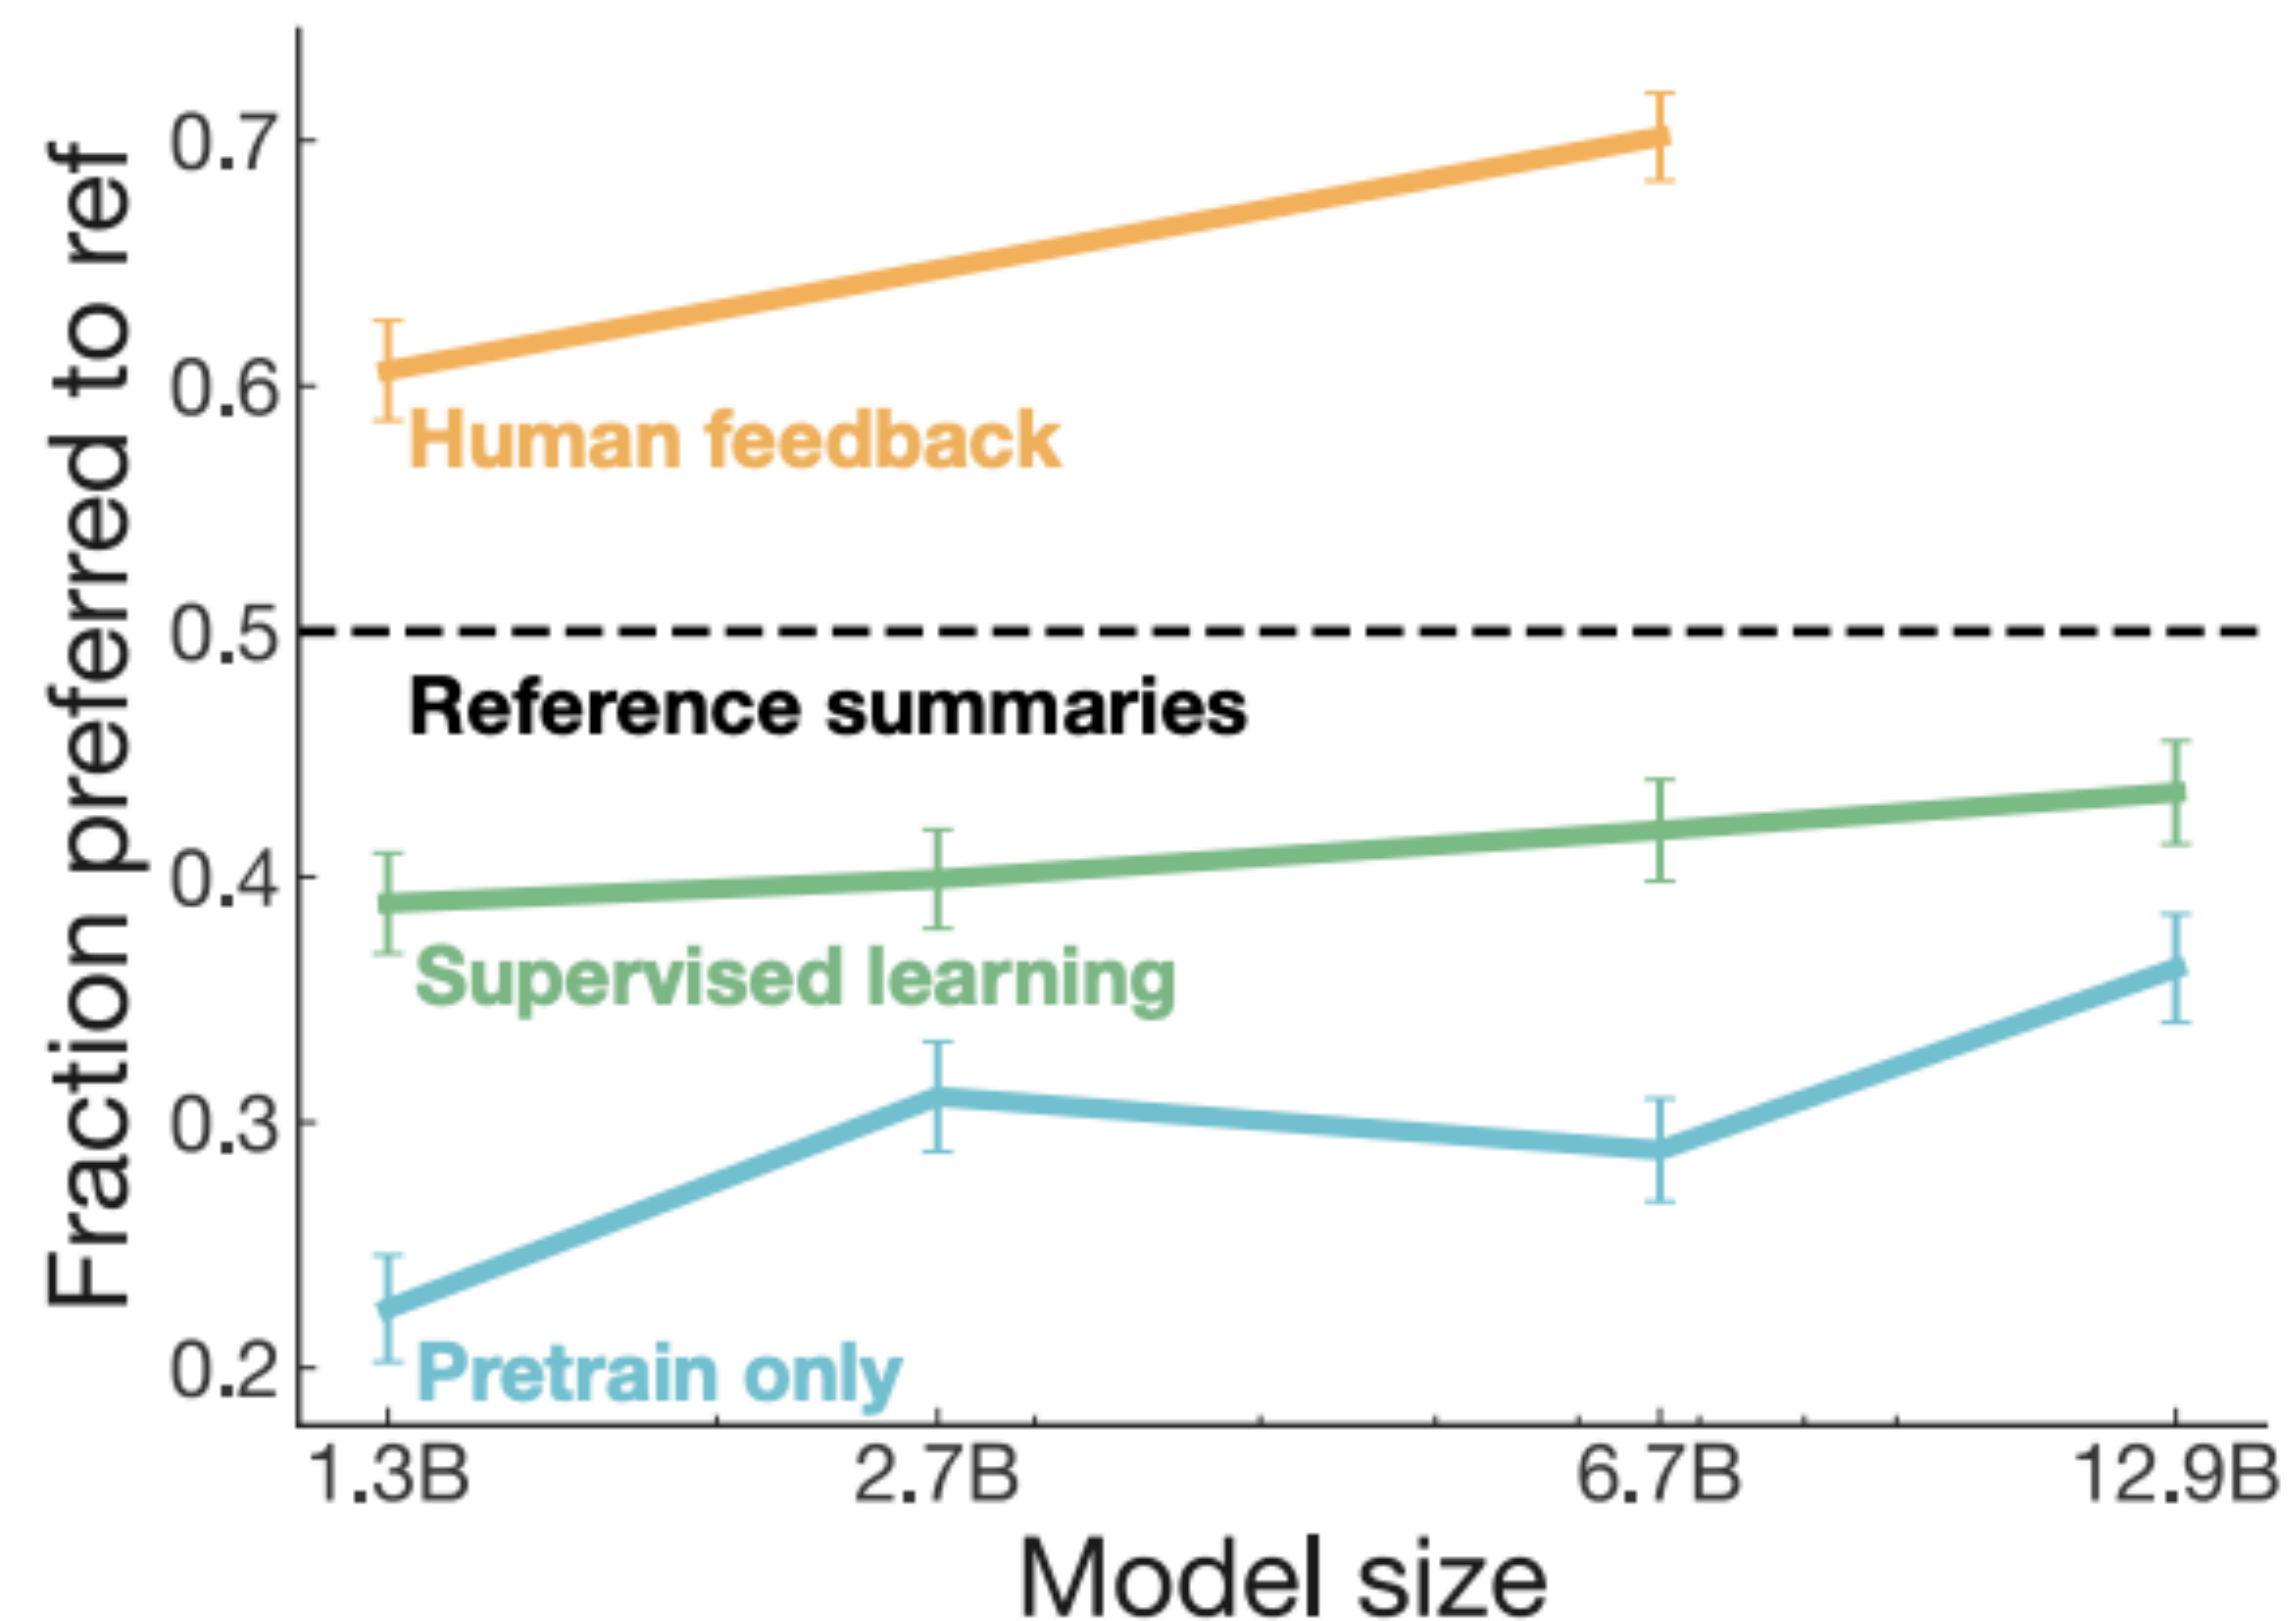
Full summary reward

Learning to Summarize from Human Feedback, Stiennon et al. (2020)

- Reward: **learned from human rankings of summary generations**

Reward Function Approaches

Learning to Summarize from Human Feedback, Stiennon et al. (2020) - RLHF



Does RLHF solve the problem?



Summarize "The Picture of Dorian Gray" for a young child with a focus on what happens to the painting throughout the book.



"The Picture of Dorian Gray" is a story about a young man named Dorian who is very handsome and likes to have fun. One day, he has his portrait painted by an artist named Basil, and he becomes very interested in the painting.


As Dorian goes on living his life, he starts to do some bad things, like lying and hurting other people's feelings. But instead of Dorian getting older and uglier like everyone else, his portrait gets older and uglier, showing all the bad things he's done.

Dorian realizes that the painting is like a mirror of his soul, and he becomes very afraid of it. He tries to destroy it, but it keeps getting worse and worse. In the end, Dorian confesses to all his bad deeds and dies, while the portrait goes back to looking young and beautiful.

So, the painting in "The Picture of Dorian Gray" is like a magical mirror that shows how a person's inner beauty or ugliness can change over time.

Does RLHF solve the problem?

M Summarize "The Picture of Dorian Gray" for a young child with a focus on what happens to the painting throughout the book.

 "The Picture of Dorian Gray" is a story about a young man named Dorian who is very handsome and likes to have fun. One day, he has his portrait painted by an artist named Basil, and he becomes very interested in the painting.

As Dorian goes on living his life, he starts to do some bad things, like lying and hurting other people's feelings. But instead of Dorian getting older and uglier like everyone else, his portrait gets older and uglier, showing all the bad things he's done.

Dorian realizes that the painting is like a mirror of his soul, and he becomes very afraid of it. He tries to destroy it, but it keeps getting worse and worse. In the end, Dorian confesses to all his bad deeds and dies, while the portrait goes back to looking young and beautiful.

So, the painting in "The Picture of Dorian Gray" is like a magical mirror that shows how a person's inner beauty or ugliness can change over time.

Does RLHF solve the problem?

M Summarize "The Picture of Dorian Gray" for a young child with a focus on what happens to the painting throughout the book.

 "The Picture of Dorian Gray" is a story about a young man named Dorian who is very handsome and likes to have fun. One day, he has his portrait painted by an artist named Basil, and he becomes very interested in the painting.


As Dorian goes on living his life, he starts to do some bad things, like lying and hurting other people's feelings. But instead of Dorian getting older and uglier like everyone else, his portrait gets older and uglier, showing all the bad things he's done.

Dorian realizes that the painting is like a mirror of his soul, and he becomes very afraid of it. He tries to destroy it, but it keeps getting worse and worse. In the end, Dorian confesses to all his bad deeds and dies, while the portrait goes back to looking young and beautiful.

So, the painting in "The Picture of Dorian Gray" is like a magical mirror that shows how a person's inner beauty or ugliness can change over time.

Does RLHF solve the problem?

M Summarize "The Picture of Dorian Gray" for a young child with a focus on what happens to the painting throughout the book.

 "The Picture of Dorian Gray" is a story about a young man named Dorian who is very handsome and likes to have fun. One day, he has his portrait painted by an artist named Basil, and he becomes very interested in the painting.

As Dorian goes on living his life, he starts to do some bad things, like lying and hurting other people's feelings. But instead of Dorian getting older and uglier like everyone else, his portrait gets older and uglier, showing all the bad things he's done.

Dorian realizes that the painting is like a mirror of his soul, and he becomes very afraid of it. He tries to destroy it, but it keeps getting worse and worse. In the end, Dorian confesses to all his bad deeds and dies, while the portrait goes back to looking young and beautiful.

So, the painting in "The Picture of Dorian Gray" is like a magical mirror that shows how a person's inner beauty or ugliness can change over time.

Dorian does not confess!
But the full story was not actually input to the model

Summary - Reward function approaches seem promising

Attribute-based...

- **Direct control** over the attributes but **constrained** by what they are

Summary - Reward function approaches seem promising

Attribute-based...

- **Direct control** over the attributes but **constrained** by what they are

Query-focused...

- Allows for **nuanced topic control** but **not style** of communication

Summary - Reward function approaches seem promising

Attribute-based...

- **Direct control** over the attributes but **constrained** by what they are

Query-focused...

- Allows for **nuanced topic control** but **not style** of communication

Reward function...

- Enables **flexible and effective** control across **both topic and style** but:
 - Expensive/challenging to define if using LLMs or humans for feedback
 - Is this enough? How can it combine with long document strategies?

1. Narrative Summarization

What is narrative summarization and why is it important?

2. Long Documents

How can we summarize very long narratives?

3. Controllability

How do we target summaries to specific tasks/users?

4. Evaluation

Can we automatically evaluate how well a system works?

Common Narrative Summarization Errors

Summaries may be long and complex, creating challenges for:

Coherence - Organized, consistent communication at a linguistic and semantic level *within the summary*

Faithfulness - Factual consistency *between source document and summary*

Common Narrative Summarization Errors

Summaries may be long and complex, creating challenges for:

Coherence - Organized, consistent communication at a linguistic and semantic level *within the summary*

Faithfulness - Factual consistency *between source document and summary*

Detecting Incoherence

How coherent are neural models of coherence?, Pishdad et al. (2020)

- **Linguistically based coherence tasks - *connective substitution, sentence cloze, topic switching, and story cloze***

Detecting Incoherence

How coherent are neural models of coherence?, Pishdad et al. (2020)

- **Linguistically based coherence tasks - *connective substitution, sentence cloze, topic switching, and story cloze***

Entity-based Neural Local Coherence Modeling, Jeon et al. (2022)

- **Operate over noun phrases and proper nouns instead of words/tokens**

Detecting Incoherence

How coherent are neural models of coherence?, Pishdad et al. (2020)

- **Linguistically based coherence tasks - *connective substitution, sentence cloze, topic switching, and story cloze***

Entity-based Neural Local Coherence Modeling, Jeon et al. (2022)

- **Operate over noun phrases and proper nouns instead of words/tokens**

Towards Coherent and Consistent Use of Entities in Narrative Generations, Papalampidi et al. (2022)

- **Metrics for long-range entity coherence and consistency**

Detecting Incoherence

How coherent are neural models of coherence?, Pishdad et al. (2020)

- **Linguistically based coherence tasks - *connective substitution, sentence cloze, topic switching, and story cloze***

Entity-based Neural Local Coherence Modeling, Jeon et al. (2022)

- **Operate over noun phrases and proper nouns instead of words/tokens**

Towards Coherent and Consistent Use of Entities in Narrative Generations, Papalampidi et al. (2022)

- **Metrics for long-range entity coherence and consistency**

SNaC: Coherence Error Detection for Narrative Summarization, Goyal et al. (2022)

- **Error taxonomy and annotation across generated summaries**

Detecting Incoherence

How coherent are neural models of coherence?, Pishdad et al. (2020)

- **Linguistically based coherence tasks - *connective substitution, sentence cloze, topic switching, and story cloze***

Entity-based Neural Local Coherence Modeling, Jeon et al. (2022)

- **Operate over noun phrases and proper nouns instead of words/tokens**

Towards Coherent and Consistent Use of Entities in Narrative Generations, Papalampidi et al. (2022)

- **Metrics for long-range entity coherence and consistency**

SNaC: Coherence Error Detection for Narrative Summarization, Goyal et al. (2022)

- **Error taxonomy and annotation across generated summaries**

Detecting Incoherence

SNaC: Coherence Detection for Narrative summarization, Goyal et al. (2022)

Examples of semantic errors:

- New character not introduced
- Missing reference to event/object
- Abrupt scene transition

Examples of language/fluency errors:

- Unnecessary repetition
- Ungrammatical/nonsensical
- Unclear coreference

Detecting Incoherence

SNaC: Coherence Detection for Narrative summarization, Goyal et al. (2022)

CharE, v = 3
 Miss Manette receives a letter from the bank informing her that information about her father's small property has been discovered. She wants to travel to France to identify him and restore him to life.

CharE, v = 3
 Mr. Lorry explains that her father has been found under another name and is being held in a house in Paris.

CharE, v = 3 *SceneE, v = 2* *InconE, v = 2*
 In court, Mr. Darnay is accused of treason. However, he is acquitted after his patriot friend, Roger Cly, testifies against him.

CharE, v = 3 *SceneE, v = 3*
 Mr. Lorry visits the Doctor's house on a Sunday afternoon as he often does. Miss Pross, the housekeeper, worries that many people will come to the house to look for Ladybird.

CharE, v = 3
 Suddenly, the Doctor starts to feel ill and says they should go inside.

SceneE, v = 2 *CharE, v = 3* *RefE, v = 2*
 Charles Darnay, the Marquis' nephew, returns to France to pursue the sacred object that took him away. He tells the Marquis that he renounces his French property as it is full of misery.

RefE, v = 2
 Charles has been in love with Lucie Manette for a long time but has never told her about his feelings.

CharE, v = 3
 Stryver tells Lorry that he intends to marry Lucie for pragmatic reasons.

New character without introduction (CharE)

Missing reference to object/event (RefE)

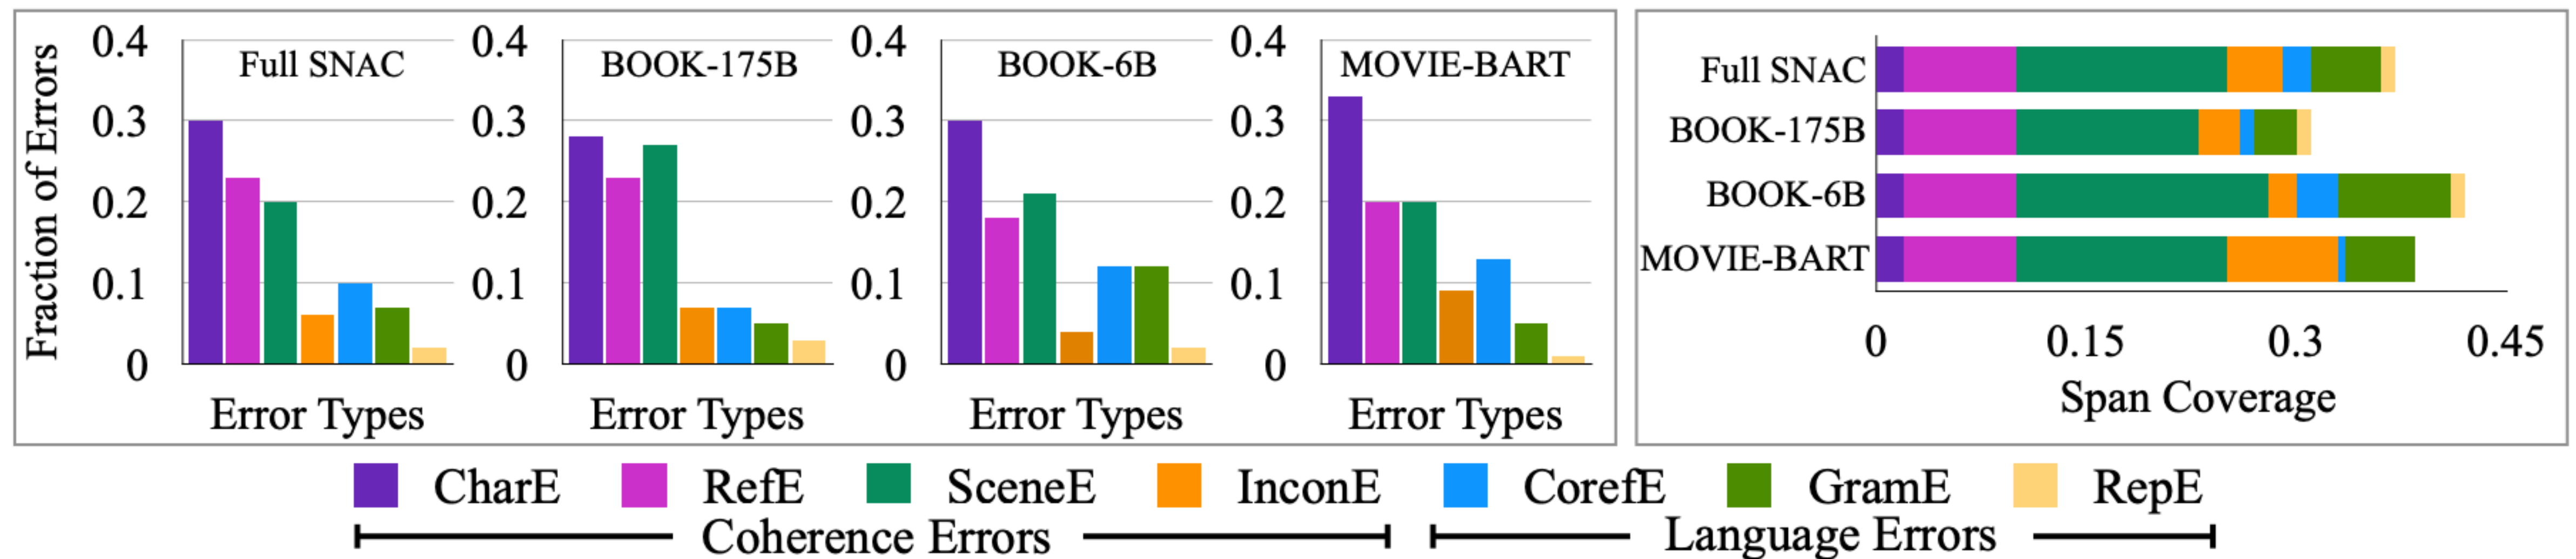
Abrupt scene transition (SceneE)

Inconsistent (InconE)

Figure 3: An example of expert annotations for a BOOK-175B summary (we only show coherence errors). The

Detecting Incoherence

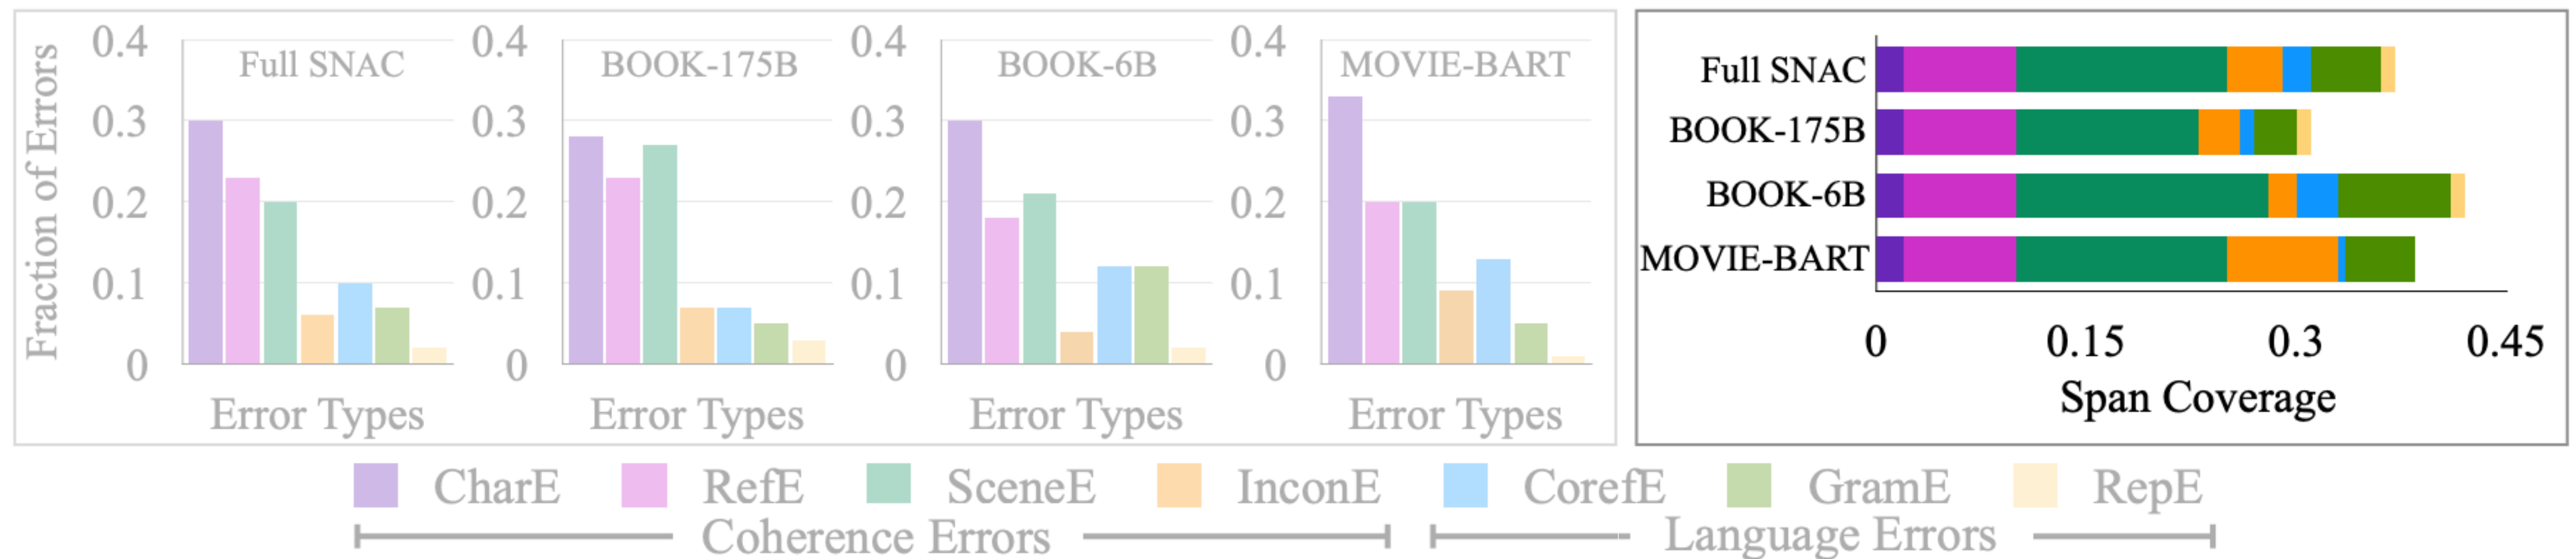
SNaC: Coherence Detection for Narrative summarization, Goyal et al. (2022)



Detecting Incoherence

On average, ~30 errors per summary!

SNaC: Coherence Detection for Narrative summarization, Goyal et al. (2022)



Challenges with Narrative Summarization Evaluation

Summaries may be long and complex, creating challenges for:

Coherence - Organized, consistent communication at a linguistic and semantic level *within the summary*.

Faithfulness - Factual consistency *between source document and summary*.

Detecting Factual Inconsistency

1. Natural Language Inference (NLI) *Is the summary entailed by the document?*

- *FalseSum: Generating Document-level NLI Examples for Recognizing Factual Inconsistency in Summarization*, Utama et al. (2022)
- *SummaC: Re-Visiting NLI-based Models for Inconsistency Detection in Summarization*, Laban et al. (2022)

Detecting Factual Inconsistency

1. Natural Language Inference (NLI)

- *FalseSum: Generating Document-level NLI Examples for Recognizing Factual Inconsistency in Summarization*, Utama et al. (2022)
- *SummaC: Re-Visiting NLI-based Models for Inconsistency Detection in Summarization*, Laban et al. (2022)

2. Questing Answering (QA) **Can you answer questions about the summary using information from the document?**

- *FEQA: A Question Answering Evaluation Framework for Faithfulness Assessment in Abstractive Summarization*, Durmus et al. (2020)
- *QAFactEval: Improved QA-Based Factual Consistency Evaluation for Summarization*, Fabbri et al. (2022)

Detecting Factual Inconsistency

1. Natural Language Inference (NLI)

- *FalseSum: Generating Document-level NLI Examples for Recognizing Factual Inconsistency in Summarization*, Utama et al. (2022)
- *SummaC: Re-Visiting NLI-based Models for Inconsistency Detection in Summarization*, Laban et al. (2022)

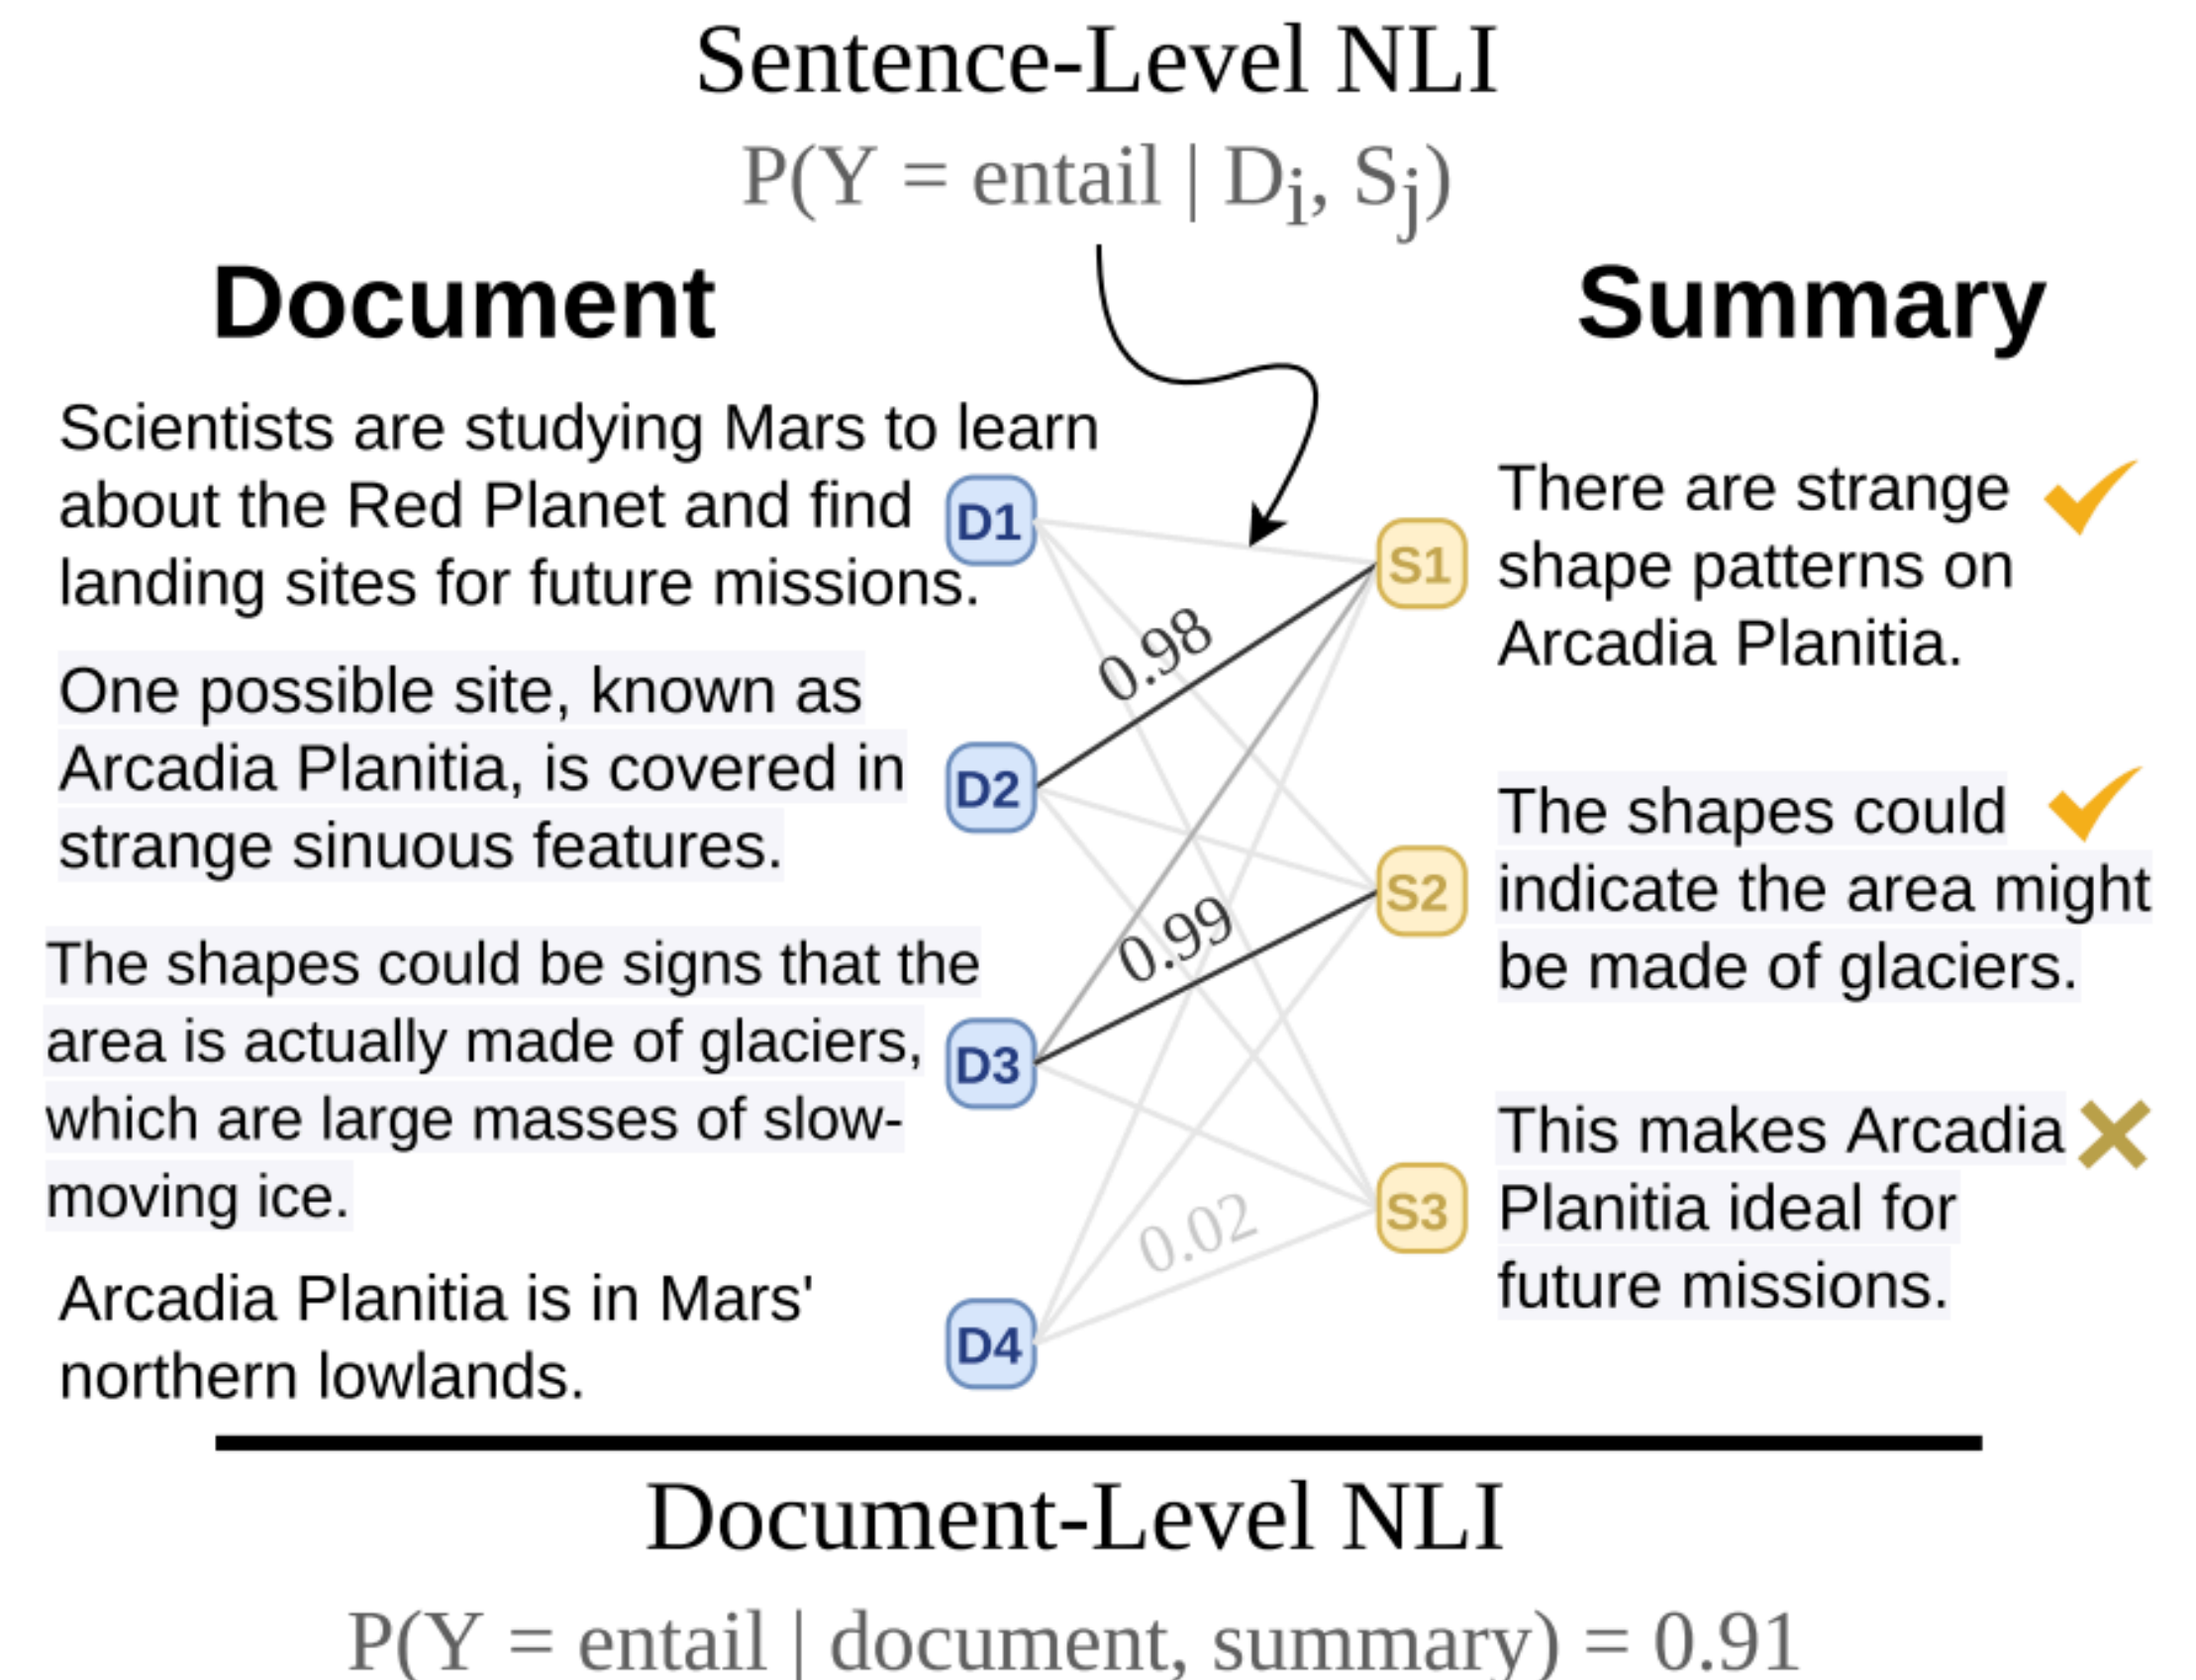
2. Question Answering (QA)

- *FEQA: A Question Answering Evaluation Framework for Faithfulness Assessment in Abstractive Summarization*, Durmus et al. (2020)
- *QAFactEval: Improved QA-Based Factual Consistency Evaluation for Summarization*, Fabbri et al. (2022)

Detecting Factual Inconsistency w/ NLI

SummaC: Re-Visiting NLI-based Models for Inconsistency Detection in Summarization, Laban et al. (2022)

- Important benchmark dataset
- NLI methods can be effective when applied at the right level of granularity



Detecting Factual Inconsistency w/ QA

QAFactEval: Improved QA-Based Factual Consistency Evaluation for Summarization, Fabbri et al. (2022)

- Extensive comparison of entailment and QA
- Improves QA with good question generation and answerability filtering
- Shows entailment and QA are complementary and can be combined

Detecting Factual Inconsistency w/ QA

SummaC Benchmark Datasets

	FactCC	SummEval	CGS	XSF	Polytope	FRANK
FEQA	53.6	53.8	61.0	56.0	57.8	69.9
FalseSum	83.52	65.18				
SummaC	89.5	81.7	64.7	66.4	62.7	81.6
QAFactEval	89.3	80.5	78.1	60.9	83.7	84.3

QAFactEval is current SOTA but both NLI and QA are continually improving

Summary - Some good strategies but not consistently used

For Coherence... designing detection strategies with a **basis in linguistics** helps

Summary - Some good strategies but not consistently used

For Coherence... designing detection strategies with a **basis in linguistics** helps

For Faithfulness... both NLI and QA metrics can be effective with the right settings

- NLI methods **directly compare the summary against the source** but require a similar pretraining setting
- QA methods can use an **out-of-the-box QA** system but require question/answer generation

Summary - Some good strategies but not consistently used

For Coherence... designing detection strategies with a **basis in linguistics** helps

For Faithfulness... both NLI and QA metrics can be effective with the right settings

- NLI methods **directly compare the summary against the source** but require a similar pretraining setting
- QA methods can use an **out-of-the-box QA** system but require question/answer generation

Rouge is still broadly used by the leading systems but, as we've seen, further analysis shows **coherence and faithfulness errors**

Overall Patterns

Educational Question Generation of Children Storybooks Via Question Type Distribution Learning and Event-Centric Summarization
 Recursively Summarizing Books with Human Feedback
 BookSum: A Collection of Datasets for Long-form Narrative Summarization
 Exploring Content Selection in Summarization of Novel Chapters
 Screenplay Summarization Using Latent Narrative Structure
 SummScreen: A Dataset for Abstractive Screenplay Summarization
 Narrative Paths and Negotiation of Power in Birth Stories
 Narrative Theory for Computational Narrative Understanding
 HydraSum: Disentangling Style Features in Text Summarization with Multi-Decoder Models
 Long Document Summarization with Top-Down and Bottom-Up Inference
 A Divide-and-Conquer Approach to the Summarization of Long Documents
 Longformer: The Long-Document Transformer
 Sparse Sinkhorn Attention
 Efficient Attentions for Long Document Summarization
 Discourse-Aware Unsupervised Summarization of Long Scientific Documents
 DYLE: Dynamic Latent Extraction for Abstractive Long-Input Summarization
 A Hierarchical Network for Abstractive Meeting Summarization with Cross-Domain Pretraining
 MacSum: Controllable Summarization with Mixed Attributes
 Text Summarization with Latent Queries
 Controllable Summarization with Constrained Markov Decision Processes
 Learning to Summarize from Human Feedback
 Controllable Neural Story Plot Generation Via Reward Shaping
 SNaC: Coherence Error Detection for Narrative Summarization
 Entity-based Neural Local Coherence Modeling
 SummaC: Re-Visiting NLI-based Models for Inconsistency Detection in Summarization
 FalseSum: Generating Document-level NLI Examples for Recognizing Factual Inconsistency in Summarization
 Towards Coherent and Consistent Use of Entities in Narrative Generation
 How Coherent are Neural Models of Coherence?
 QAFactEval: Improved QA-based Factual Consistency Evaluation for Summarization
 FEQA: A Question Answering Evaluation Framework for Faithfulness Assessment in Abstractive Summarization

Overall Patterns - Evaluates on narrative

Educational Question Generation of Children Storybooks Via Question Type Distribution Learning and Event-Centric Summarization

Recursively Summarizing Books with Human Feedback

MacSum: Controllable Summarization with Mixed Attributes

BookSum: A Collection of Datasets for Long-form Narrative Summarization

Text Summarization with Latent Queries

Exploring Content Selection in Summarization of Novel Chapters

Screenplay Summarization Using Latent Narrative Structure

Controllable Summarization with Constrained Markov Decision Processes

SummScreen: A Dataset for Abstractive Screenplay Summarization

Learning to Summarize from Human Feedback

Narrative Paths and Negotiation of Power in Birth Stories

Controllable Neural Story Plot Generation Via Reward Shaping

Narrative Theory for Computational Narrative Understanding

HydraSum: Disentangling Style Features in Text Summarization with Multi-Decoder Models

Long Document Summarization with Top-Down and Bottom-Up Inference

SNaC: Coherence Error Detection for Narrative Summarization

A Divide-and-Conquer Approach to the Summarization of Long Documents

Entity-based Neural Local Coherence Modeling

Longformer: The Long-Document Transformer

SummaC: Re-Visiting NLI-based Models for Inconsistency Detection in Summarization

Sparse Sinkhorn Attention

FalseSum: Generating Document-level NLI Examples for Recognizing Factual Inconsistency in Summarization

Efficient Attentions for Long Document Summarization

Towards Coherent and Consistent Use of Entities in Narrative Generation

Discourse-Aware Unsupervised Summarization of Long Scientific Documents

DYLE: Dynamic Latent Extraction for Abstractive Long-Input Summarization

How Coherent are Neural Models of Coherence?

A Hierarchical Network for Abstractive Meeting Summarization with Cross-Domain Pretraining

QAFactEval: Improved QA-based Factual Consistency Evaluation for Summarization

FEQA: A Question Answering Evaluation Framework for Faithfulness Assessment in Abstractive Summarization

Overall Patterns - Evaluates on books

Educational Question Generation of Children Storybooks Via Question Type Distribution Learning and Event-Centric Summarization

Recursively Summarizing Books with Human Feedback

MacSum: Controllable Summarization with Mixed Attributes

BookSum: A Collection of Datasets for Long-form Narrative Summarization

Text Summarization with Latent Queries

Exploring Content Selection in Summarization of Novel Chapters

Screenplay Summarization Using Latent Narrative Structure

Controllable Summarization with Constrained Markov Decision Processes

SummScreen: A Dataset for Abstractive Screenplay Summarization

Learning to Summarize from Human Feedback

Narrative Paths and Negotiation of Power in Birth Stories

Controllable Neural Story Plot Generation Via Reward Shaping

Narrative Theory for Computational Narrative Understanding

HydraSum: Disentangling Style Features in Text Summarization with Multi-Decoder Models

Long Document Summarization with Top-Down and Bottom-Up Inference

SNaC: Coherence Error Detection for Narrative Summarization

A Divide-and-Conquer Approach to the Summarization of Long Documents

Entity-based Neural Local Coherence Modeling

Longformer: The Long-Document Transformer

SummaC: Re-Visiting NLI-based Models for Inconsistency Detection in Summarization

Sparse Sinkhorn Attention

FalseSum: Generating Document-level NLI Examples for Recognizing Factual Inconsistency in Summarization

Efficient Attentions for Long Document Summarization

Towards Coherent and Consistent Use of Entities in Narrative Generation

Discourse-Aware Unsupervised Summarization of Long Scientific Documents

DYLE: Dynamic Latent Extraction for Abstractive Long-Input Summarization

How Coherent are Neural Models of Coherence?

A Hierarchical Network for Abstractive Meeting Summarization with Cross-Domain Pretraining

QAFactEval: Improved QA-based Factual Consistency Evaluation for Summarization

FEQA: A Question Answering Evaluation Framework for Faithfulness Assessment in Abstractive Summarization

Overall Patterns - Input length > 50,000 tokens

Educational Question Generation of Children Storybooks Via Question Type Distribution Learning and Event-Centric Summarization

Recursively Summarizing Books with Human Feedback

MacSum: Controllable Summarization with Mixed Attributes

BookSum: A Collection of Datasets for Long-form Narrative Summarization

Text Summarization with Latent Queries

Exploring Content Selection in Summarization of Novel Chapters

Screenplay Summarization Using Latent Narrative Structure

Controllable Summarization with Constrained Markov Decision Processes

SummScreen: A Dataset for Abstractive Screenplay Summarization

Learning to Summarize from Human Feedback

Narrative Paths and Negotiation of Power in Birth Stories

Controllable Neural Story Plot Generation Via Reward Shaping

Narrative Theory for Computational Narrative Understanding

HydraSum: Disentangling Style Features in Text Summarization with Multi-Decoder Models

Long Document Summarization with Top-Down and Bottom-Up Inference

SNaC: Coherence Error Detection for Narrative Summarization

A Divide-and-Conquer Approach to the Summarization of Long Documents

Entity-based Neural Local Coherence Modeling

Longformer: The Long-Document Transformer

SummaC: Re-Visiting NLI-based Models for Inconsistency Detection in Summarization

Sparse Sinkhorn Attention

FalseSum: Generating Document-level NLI Examples for Recognizing Factual Inconsistency in Summarization

Efficient Attentions for Long Document Summarization

Towards Coherent and Consistent Use of Entities in Narrative Generation

Discourse-Aware Unsupervised Summarization of Long Scientific Documents

DYLE: Dynamic Latent Extraction for Abstractive Long-Input Summarization

How Coherent are Neural Models of Coherence?

A Hierarchical Network for Abstractive Meeting Summarization with Cross-Domain Pretraining

QAFactEval: Improved QA-based Factual Consistency Evaluation for Summarization

FEQA: A Question Answering Evaluation Framework for Faithfulness Assessment in Abstractive Summarization

Overall Patterns - Coherence/Faithfulness Benchmark

Narrative Summarization	Long Documents	Controllability	Evaluation
Educational Question Generation of Children Storybooks Via Question Type Distribution Learning and Event-Centric Summarization			
Recursively Summarizing Books with Human Feedback			MacSum: Controllable Summarization with Mixed Attributes
BookSum: A Collection of Datasets for Long-form Narrative Summarization			Text Summarization with Latent Queries
Exploring Content Selection in Summarization of Novel Chapters			
Screenplay Summarization Using Latent Narrative Structure		Controllable Summarization with Constrained Markov Decision Processes	
SummScreen: A Dataset for Abstractive Screenplay Summarization			Learning to Summarize from Human Feedback
Narrative Paths and Negotiation of Power in Birth Stories		Controllable Neural Story Plot Generation Via Reward Shaping	
Narrative Theory for Computational Narrative Understanding			
	HydraSum: Disentangling Style Features in Text Summarization with Multi-Decoder Models		
Long Document Summarization with Top-Down and Bottom-Up Inference			
A Divide-and-Conquer Approach to the Summarization of Long Documents			
Longformer: The Long-Document Transformer			
Sparse Sinkhorn Attention			
Efficient Attentions for Long Document Summarization			
Discourse-Aware Unsupervised Summarization of Long Scientific Documents			
DYLE: Dynamic Latent Extraction for Abstractive Long-Input Summarization			
A Hierarchical Network for Abstractive Meeting Summarization with Cross-Domain Pretraining			

Compared performance against another model using a measure of coherence or faithfulness that was not ROUGE

Overall Patterns - Used a Non-ROUGE Benchmark

Educational Question Generation of Children Storybooks Via Question Type Distribution Learning and Event-Centric Summarization

Recursively Summarizing Books with Human Feedback

MacSum: Controllable Summarization with Mixed Attributes

BookSum: A Collection of Datasets for Long-form Narrative Summarization

Text Summarization with Latent Queries

Exploring Content Selection in Summarization of Novel Chapters

Screenplay Summarization Using Latent Narrative Structure

Controllable Summarization with Constrained Markov Decision Processes

SummScreen: A Dataset for Abstractive Screenplay Summarization

Learning to Summarize from Human Feedback

Narrative Paths and Negotiation of Power in Birth Stories

Controllable Neural Story Plot Generation Via Reward Shaping

Narrative Theory for Computational Narrative Understanding

HydraSum: Disentangling Style Features in Text Summarization with Multi-Decoder Models

Long Document Summarization with Top-Down and Bottom-Up Inference

A Divide-and-Conquer Approach to the Summarization of Long Documents

Longformer: The Long-Document Transformer

Sparse Sinkhorn Attention

Efficient Attentions for Long Document Summarization

Discourse-Aware Unsupervised Summarization of Long Scientific Documents

DYLE: Dynamic Latent Extraction for Abstractive Long-Input Summarization

A Hierarchical Network for Abstractive Meeting Summarization with Cross-Domain Pretraining

Compared performance against another model using a measure of coherence or faithfulness that was not ROUGE

Where we are now...

- Useful narrative summarization benchmarks exist
- LLMs are capable of long document processing and abstractive summarization
- Incorporating RL enables flexible control of summaries
- Variety of useful automatic metrics for coherence and faithfulness

Where we are now...

- Useful narrative summarization benchmarks exist
- LLMs are capable of long document processing and abstractive summarization
- Incorporating RL enables flexible control of summaries
- Variety of useful automatic metrics for coherence and faithfulness

Where we hope to get to...

1. Thorough benchmarking of existing summarization strategies across narrative tasks using rich evaluation metrics

Where we are now...

- Useful narrative summarization benchmarks exist
- LLMs are capable of long document processing and abstractive summarization
- Incorporating RL enables flexible control of summaries
- Variety of useful automatic metrics for coherence and faithfulness

Where we hope to get to...

1. Thorough benchmarking of existing summarization strategies across narrative tasks using rich evaluation metrics
2. Faithful, coherent and controllable summaries of long documents

Where we are now...

- Useful narrative summarization benchmarks exist
- LLMs are capable of long document processing and abstractive summarization
- Incorporating RL enables flexible control of summaries
- Variety of useful automatic metrics for coherence and faithfulness

Where we hope to get to...

1. Thorough benchmarking of existing summarization strategies across narrative tasks using rich evaluation metrics
2. Faithful, coherent and controllable summaries of long documents
3. Effective approaches that are accessible outside of resource-rich industry labs