

Evaluating text generation

CS685 Spring 2024
Advanced Natural Language Processing

Mohit Iyer
College of Information and Computer Sciences
University of Massachusetts Amherst

some slides from Marine Carpuat & Marzena Karpinska

So far...

- We've seen *perplexity* as an automatic measure to evaluate language models
- However, perplexity alone is insufficient to tell us about how well a model is solving some downstream task (e.g., translation or summarization)
- Today: BLEU score for MT, ROUGE for summarization, BERT-based improvements, LLM judges, and human evaluation

How Good is Machine Translation? Chinese > English

记者从环保部了解到,《水十条》要求今年年底前直辖市、省会城市、计划单列市建成区基本解决黑臭水体。截至目前,全国224个地级及以上城市共排查确认黑臭水体2082个,其中34.9%完成整治,28.4%正在整治,22.8%正在开展项目前期。

Reporters learned from the Ministry of Environmental Protection, "Water 10" requirements before the end of this year before the municipality, the provincial capital city, plans to build a separate city to solve the basic black and black water. Up to now, the country's 224 prefecture-level and above cities were identified to confirm the black and white water 2082, of which 34.9% to complete the renovation, 28.4% is remediation, 22.8% is carrying out the project early.

How Good is Machine Translation?

French > English

A l'orée de ce débat télévisé inédit dans l'histoire de la Ve République, on attendait une forme de «Tous sur Macron» mais c'est la candidate du Front national qui s'est retrouvée au cœur des premières attaques de ses quatre adversaires d'un soir, favorisées par le premier thème abordé, les questions de société et donc de sécurité, d'immigration et de laïcité.

At the beginning of this televised debate, which was unheard of in the history of the Fifth Republic, a "Tous sur Macron" was expected, but it was the candidate of the National Front who found itself at the heart of the first attacks of its four Opponents of one evening, favored by the first theme tackled, the issues of society and thus security, immigration and secularism.

What is MT good (enough) for?

- **Assimilation:** reader initiates translation, wants to know content
 - User is tolerant of inferior quality
 - Focus of majority of research
- **Communication:** participants in conversation don't speak same language
 - Users can ask questions when something is unclear
 - Chat room translations, hand-held devices
 - Often combined with speech recognition
- **Dissemination:** publisher wants to make content available in other languages
 - High quality required
 - Almost exclusively done by human translators

How good is a translation?
Problem: no single right answer

这个机场的安全工作由以色列方面负责。

Israeli officials are responsible for airport security.

Israel is in charge of the security at this airport.

The security work for this airport is the responsibility of the Israel government.

Israeli side was in charge of the security of this airport.

Israel is responsible for the airport's security.

Israel is responsible for safety work at this airport.

Israel presides over the security of the airport.

Israel took charge of the airport security.

The safety of this airport is taken charge of by Israel.

This airport's security is the responsibility of the Israeli security officials.

Evaluation

- How good is a given machine translation system?
- Many different translations acceptable
- Evaluation metrics
 - Subjective judgments by human evaluators
 - Automatic evaluation metrics
 - Task-based evaluation

Adequacy and Fluency

- Human judgment
 - Given: machine translation output
 - Given: input and/or reference translation
 - Task: assess quality of MT output
- Metrics
 - **Adequacy:** does the output convey the meaning of the input sentence? Is part of the message lost, added, or distorted?
 - **Fluency:** is the output fluent? Involves both grammatical correctness and idiomatic word choices.

Fluency and Adequacy: Scales

Adequacy	
5	all meaning
4	most meaning
3	much meaning
2	little meaning
1	none

Fluency	
5	flawless English
4	good English
3	non-native English
2	disfluent English
1	incomprehensible

Judge Sentence

You have already judged 14 of 3064 sentences, taking 86.4 seconds per sentence.

Source: les deux pays constituent plutôt un laboratoire nécessaire au fonctionnement interne de l'ue .

Reference: rather , the two countries form a laboratory needed for the internal working of the eu .

Translation	Adequacy	Fluency
both countries are rather a necessary laboratory the internal operation of the eu .	<input type="radio"/> <input type="radio"/> <input type="radio"/> <input type="radio"/> <input checked="" type="radio"/> 1 2 3 4 5	<input type="radio"/> <input type="radio"/> <input type="radio"/> <input type="radio"/> <input checked="" type="radio"/> 1 2 3 4 5
both countries are a necessary laboratory at internal functioning of the eu .	<input type="radio"/> <input type="radio"/> <input checked="" type="radio"/> <input type="radio"/> <input type="radio"/> 1 2 3 4 5	<input type="radio"/> <input type="radio"/> <input checked="" type="radio"/> <input type="radio"/> <input type="radio"/> 1 2 3 4 5
the two countries are rather a laboratory necessary for the internal workings of the eu .	<input type="radio"/> <input type="radio"/> <input type="radio"/> <input checked="" type="radio"/> <input type="radio"/> 1 2 3 4 5	<input type="radio"/> <input type="radio"/> <input type="radio"/> <input checked="" type="radio"/> <input type="radio"/> 1 2 3 4 5
the two countries are rather a laboratory for the internal workings of the eu .	<input type="radio"/> <input type="radio"/> <input checked="" type="radio"/> <input type="radio"/> <input type="radio"/> 1 2 3 4 5	<input type="radio"/> <input type="radio"/> <input type="radio"/> <input type="radio"/> <input checked="" type="radio"/> 1 2 3 4 5
the two countries are rather a necessary laboratory internal workings of the eu .	<input type="radio"/> <input type="radio"/> <input checked="" type="radio"/> <input type="radio"/> <input type="radio"/> 1 2 3 4 5	<input type="radio"/> <input type="radio"/> <input checked="" type="radio"/> <input type="radio"/> <input type="radio"/> 1 2 3 4 5
Annotator: Philipp Koehn Task: WMT06 French-English	<input type="button" value="Annotate"/>	
Instructions	5= All Meaning 4= Most Meaning 3= Much Meaning 2= Little Meaning 1= None	5= Flawless English 4= Good English 3= Non-native English 2= Disfluent English 1= Incomprehensible

Let's try:
rate fluency & adequacy on 1-5 scale

- Source:
N'y aurait-il pas comme une vague hypocrisie de votre part ?
- Reference:
Is there not an element of hypocrisy on your part?
- System1:
Would it not as a wave of hypocrisy on your part?
- System2:
Is there would be no hypocrisy like a wave of your hand?
- System3:
Is there not as a wave of hypocrisy from you?

what are some issues
with human evaluation?

Automatic Evaluation Metrics

- Goal: computer program that computes quality of translations
- Advantages: low cost, optimizable, consistent
- Basic strategy
 - Given: MT output
 - Given: human reference translation
 - Task: compute similarity between them

Precision and Recall of Words

SYSTEM A: Israeli officials responsibility of airport safety

REFERENCE: Israeli officials are responsible for airport security

Precision

$$\frac{\text{correct}}{\text{output-length}} = \frac{3}{6} = 50\%$$

Recall

$$\frac{\text{correct}}{\text{reference-length}} = \frac{3}{7} = 43\%$$

F-measure

$$\frac{\text{precision} \times \text{recall}}{(\text{precision} + \text{recall})/2} = \frac{.5 \times .43}{(.5 + .43)/2} = 46\%$$

Precision and Recall of Words



Metric	System A	System B
precision	50%	100%
recall	43%	100%
f-measure	46%	100%

flaw: no penalty for reordering

BLEU

Bilingual Evaluation Understudy

N-gram overlap between machine translation output and reference translation

Compute precision for n-grams of size 1 to 4

Add brevity penalty (for too short translations)

$$\text{BLEU} = \min \left(1, \frac{\text{output-length}}{\text{reference-length}} \right) \left(\prod_{i=1}^4 \text{precision}_i \right)^{\frac{1}{4}}$$

Typically computed over the entire corpus, not single sentences

Multiple Reference Translations

To account for variability, use multiple reference translations

- n-grams may match in any of the references
- closest reference length used

Example

SYSTEM: Israeli officials responsibility of airport safety
2-GRAM MATCH 2-GRAM MATCH 1-GRAM

REFERENCES: Israeli officials are responsible for airport security
Israel is in charge of the security at this airport
The security work for this airport is the responsibility of the Israel government
Israeli side was in charge of the security of this airport

BLEU examples

SYSTEM A: Israeli officials responsibility of airport safety
2-GRAM MATCH 1-GRAM MATCH

REFERENCE: Israeli officials are responsible for airport security

SYSTEM B: airport security Israeli officials are responsible
2-GRAM MATCH 4-GRAM MATCH

Metric	System A	System B
precision (1gram)	3/6	6/6
precision (2gram)	1/5	4/5
precision (3gram)	0/4	2/4
precision (4gram)	0/3	1/3
brevity penalty	6/7	6/7
BLEU	0%	52%

BLEU examples

SYSTEM A: Israeli officials responsibility of airport safety
2-GRAM MATCH 1-GRAM MATCH

REFERENCE: Israeli officials are responsible for airport security

SYSTEM B: airport security Israeli officials are responsible
2-GRAM MATCH 4-GRAM MATCH


why does BLEU not account for recall?

Metric	System A	System B
precision (1gram)	3/6	6/6
precision (2gram)	1/5	4/5
precision (3gram)	0/4	2/4
precision (4gram)	0/3	1/3
brevity penalty	6/7	6/7
BLEU	0%	52%

what are some drawbacks of BLEU?

- all words/n-grams treated as equally relevant
- operates on local level
- scores are meaningless (absolute value not informative)
- human translators also score low on BLEU

Yet automatic metrics such as BLEU correlate with human judgement



ROUGE - a recall-based counterpart to BLEU

- Idea: what % of the words or n-grams in the **reference** occur in the **generated output**?
- ROUGE and its variants are often used to evaluate *text summarization* systems

Can we include *learned* components
in our evaluation metrics?

BLEURT (BLEU + BERT)

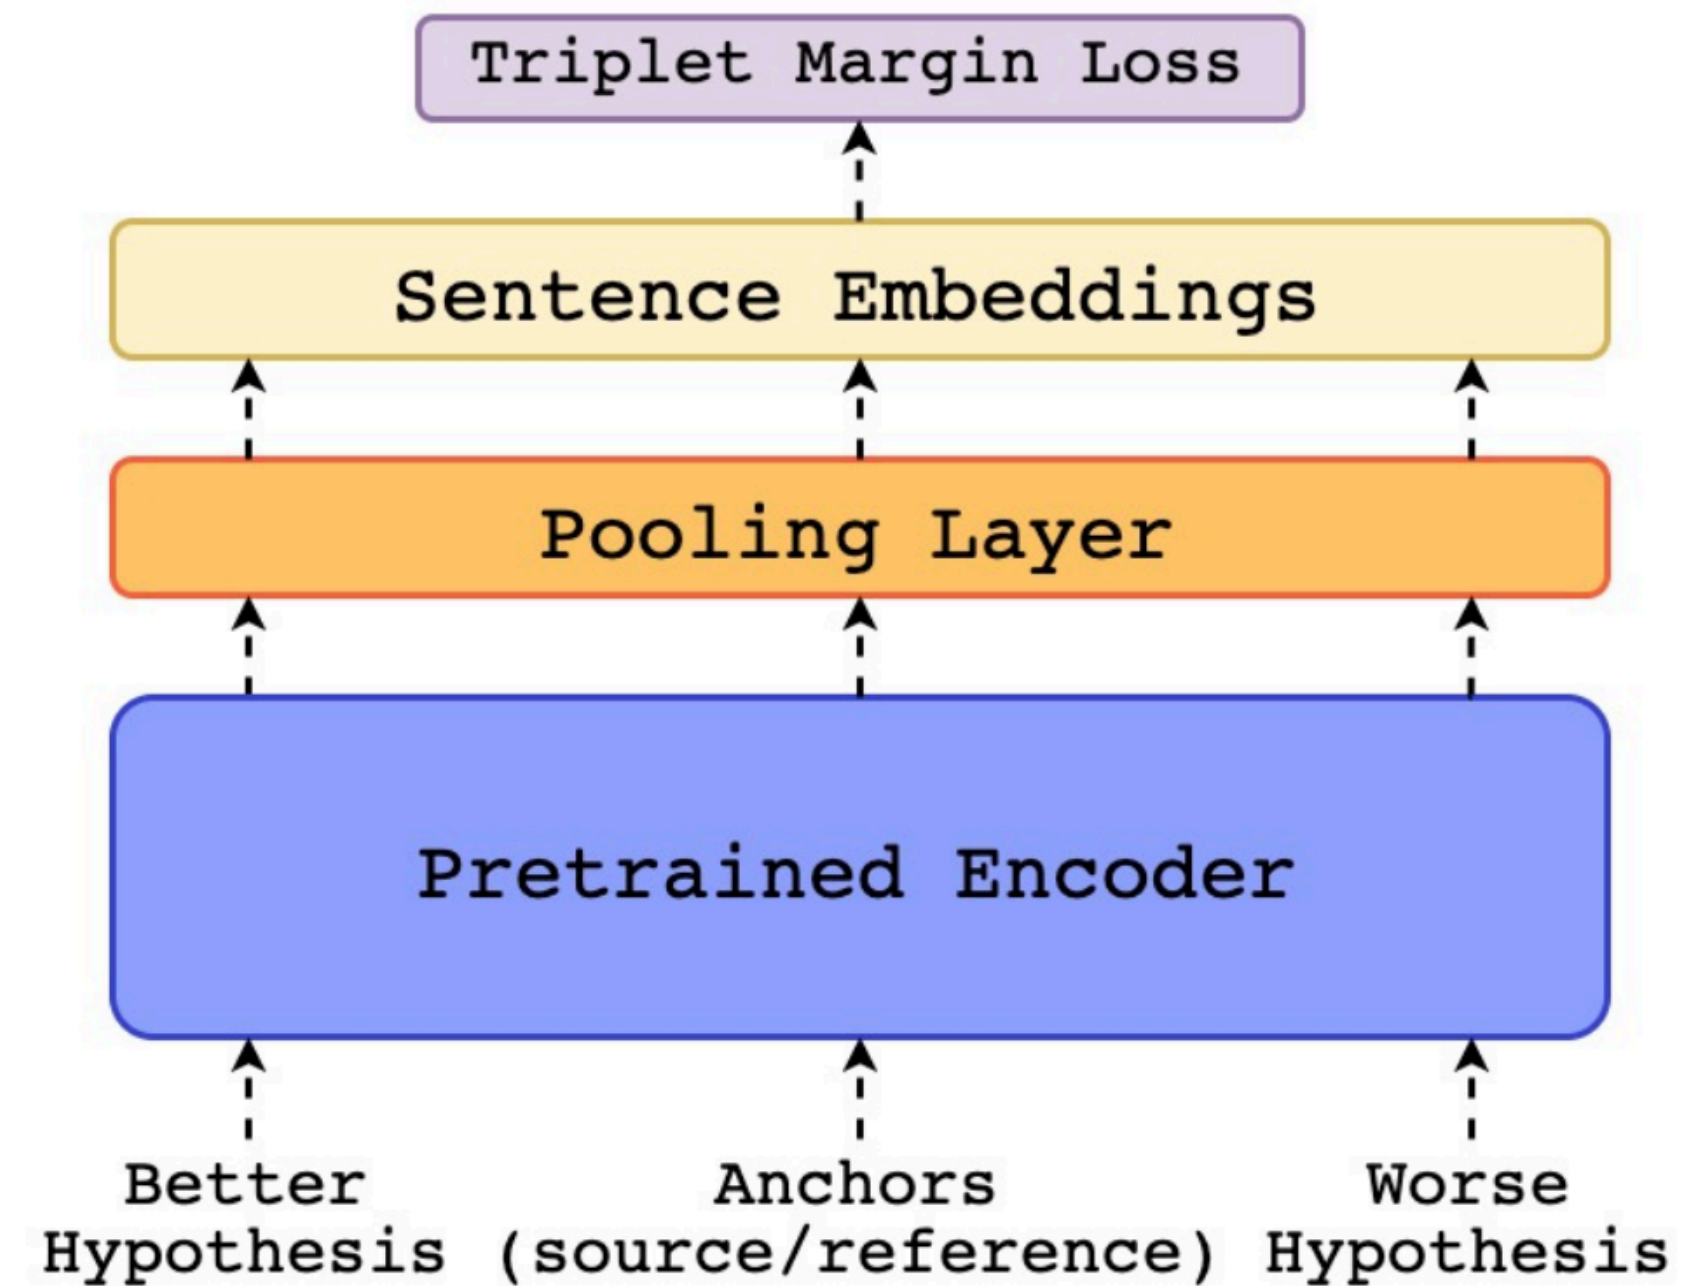
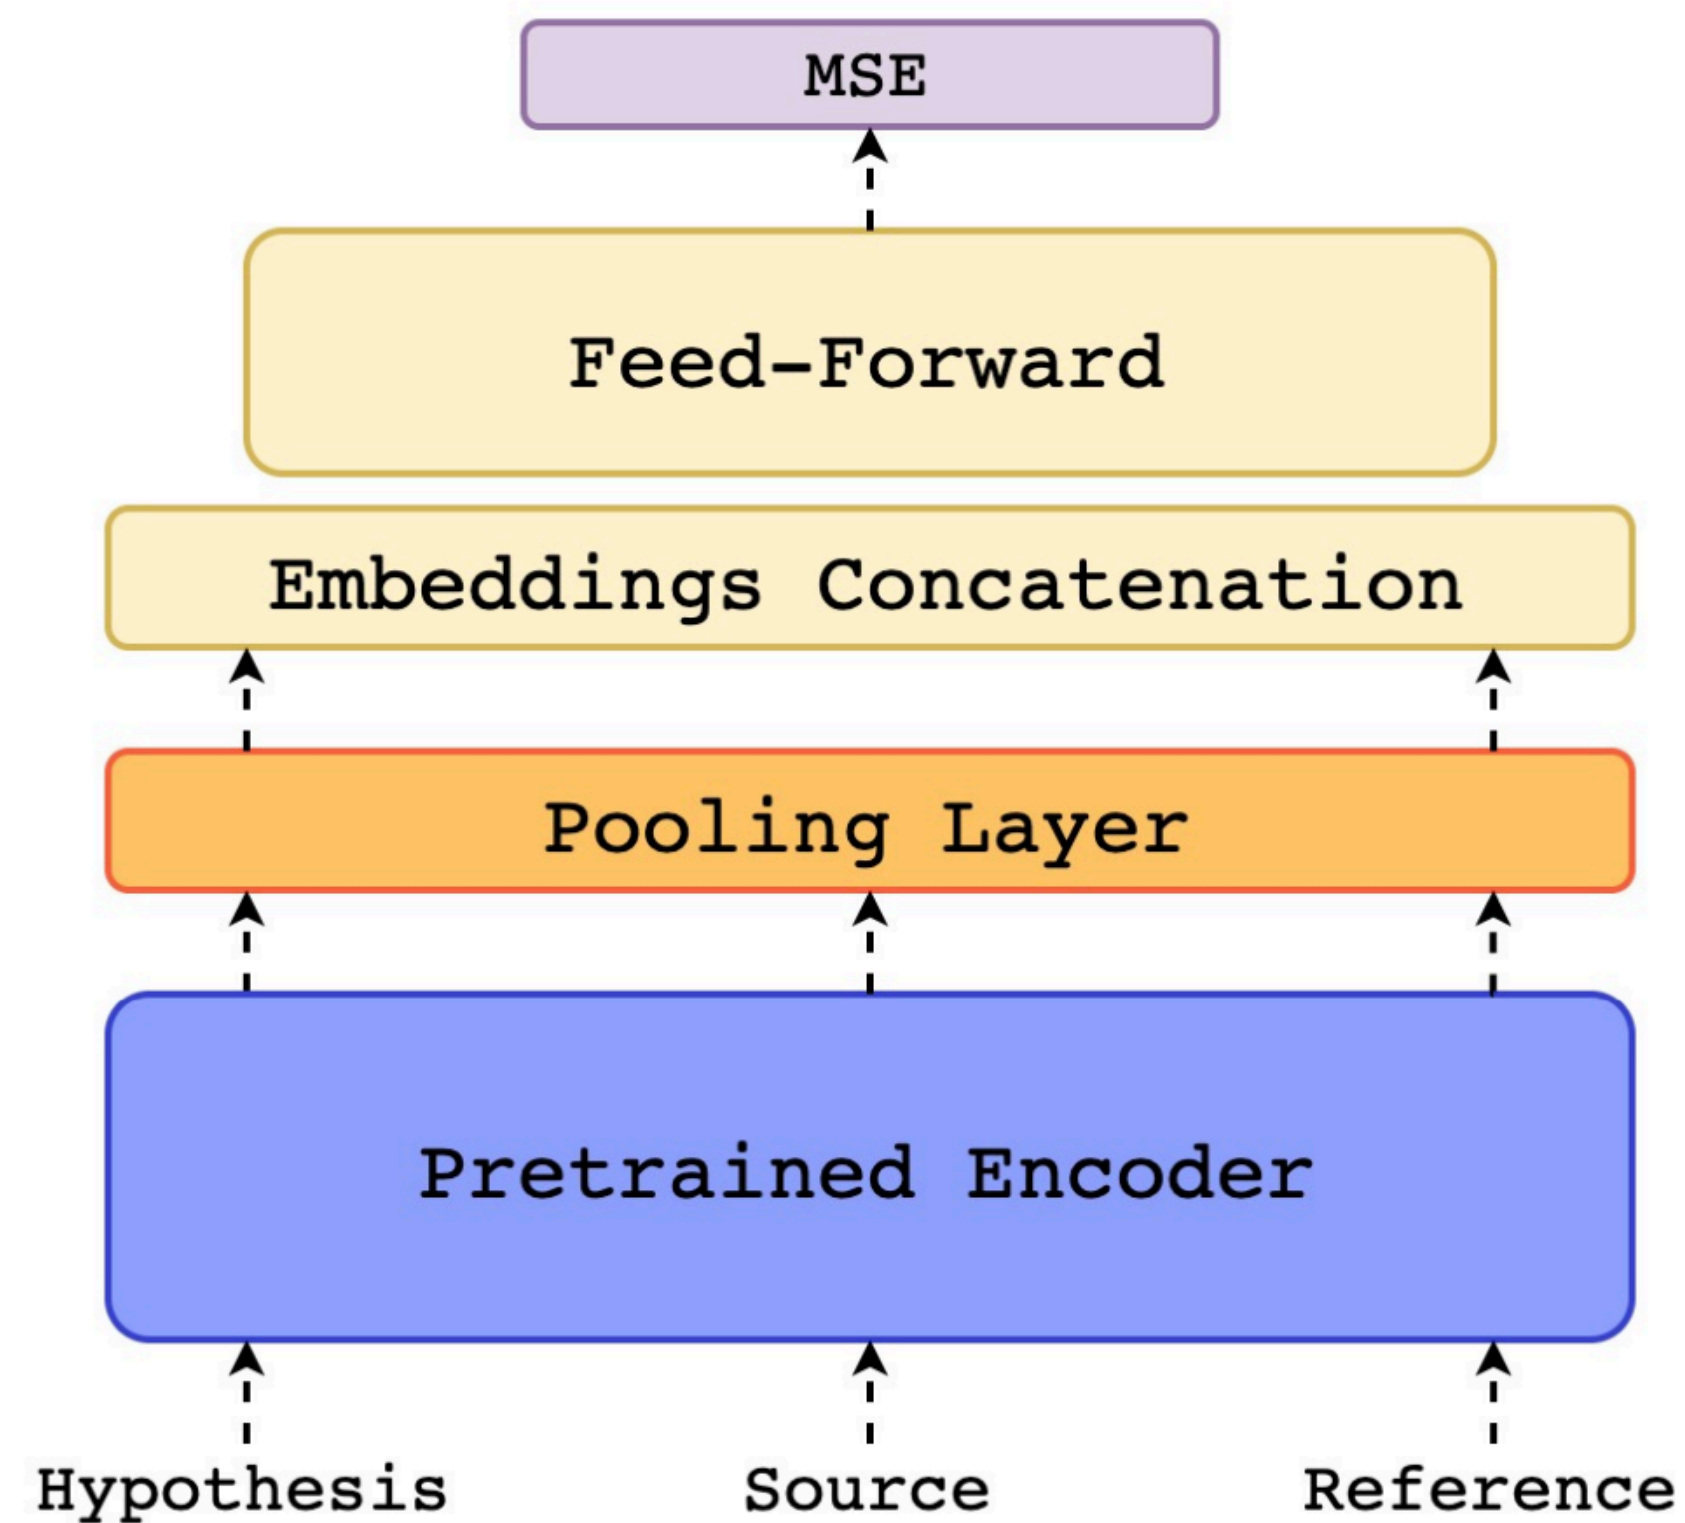
- Take a pretrained BERT, and fine-tune it on a variety of synthetic tasks with perturbed data
 - Synthetic data involves a sentence \mathbf{z} and “perturbed” version \mathbf{z}'
 - Objectives include many regression tasks (e.g., predict BLEU, ROUGE, backtranslation likelihood)
- Then, fine-tune the resulting model on small supervised datasets of human quality judgments

BLEURT (BLEU + BERT)


- Take a pretrained BERT, and fine-tune it on a variety of synthetic tasks with perturbed data
 - Synthetic data involves a sentence \mathbf{z} and “perturbed” version \mathbf{z}'
 - Objectives include many regression tasks (e.g., predict BLEU, ROUGE, backtranslation likelihood)
- Then, fine-tune the resulting model on small supervised datasets of human quality judgments

Higher correlation with human judgments than just BLEU, but has limitations...

COMET (now recommended for MT evaluation over BLEU)



How do we evaluate
open-ended text
generation?



WritingPrompts (Fan et al., 2019)

PROMPT: You, a video game character, just died. You had no spare lives left.

STORY: Link woke in a panic, gasping and flailing wildly at the air above him. The panic quickly turned to confusion, wasn't he just...?
He heard footsteps approaching from behind. Link leapt up from the floor and drew his sword, ready to strike. The sight of a slight man in a suit added to Link's already mounting confusion. Is this another trap set by Ganondorf?





200 Human-written Stories

PROMPT: *You, a video game character, just died. You had no spare lives left.*

STORY: *Link woke in a panic, gasping and flailing wildly at the air above him. The panic quickly turned to confusion, wasn't he just...? He heard footsteps approaching from behind. Link leapt up from the floor and drew his sword, ready to strike. The sight of a slight man in a suit added to Link's already mounting confusion. Is this another trap set by Ganondorf? (...)*

~150 words




200 GPT-2 Generated Stories

PROMPT: *You, a video game character, just died. You had no spare lives left.*

STORY: **The US Space Station exploded with automatic fire in preparation for tonight's nuclear strike on Earth. This is video from the control center inside the Space Station**
**A figure wearing a manilla envelope falls off of the back of the space station, detonating a nuclear device.*
***An audio file is made of the sirens, and some screams (...)*

AMAZON MECHANICAL TURK



GRAMMAR

How grammatically correct is the text of the story fragment?

COHERENCE

How well do the sentences in the story fragment fit together?


LIKABILITY

How enjoyable do you find the story fragment?

RELEVANCE

How relevant is the story fragment to the prompt?

AMAZON MECHANICAL TURK



GRAMMAR

How grammatically correct is the text of the story fragment?

COHERENCE

How well do the sentences in the story fragment fit together?

LIKABILITY

How enjoyable do you find the story fragment?

RELEVANCE

How relevant is the story fragment to the prompt?

Human evaluators struggle to read long texts

- To evaluate a 250-word generation:
 - Crowdworkers take ~10 sec
 - English teachers take ~70 sec
- High variability across runs, low agreement



Question: If under anesthesia do you get your sleep needed for the day?

Answer 1:

No. The unconsciousness induced by both IV anesthetics and inhaled anesthetics is very different from restorative sleep. In fact, patients in ICUs who are sedated for long periods of time with some of the same medications they use in anesthesia get delirium from both a lack of sleep and a lack of proper awake time. Sleep = charging your phone on airplane mode. Anesthesia = turning your phone off but not charging it.

Answer 2:

There is no clear consensus on whether or not people under anesthesia get the sleep they need for the day. Some studies suggest that people under anesthesia do get some sleep, but it is not clear how deep or restful this sleep is. Other studies suggest that people under anesthesia do not get any meaningful sleep at all. Therefore, it is difficult to say definitively whether or not people under anesthesia get the sleep they need for the day.

Which answer is better?

Question: If under anesthesia do you get your sleep needed for the day?

Answer 1:

No. The unconsciousness induced by both IV anesthetics and inhaled anesthetics is very different from restorative sleep. In fact, patients in ICUs who are sedated for long periods of time with some of the same medications they use in anesthesia get delirium from both a lack of sleep and a lack of proper awake time. Sleep = charging your phone on airplane mode. Anesthesia = turning your phone off but not charging it.

Human-written

Answer 2:

There is no clear consensus on whether or not people under anesthesia get the sleep they need for the day. Some studies suggest that people under anesthesia do get some sleep, but it is not clear how deep or restful this sleep is. Other studies suggest that people under anesthesia do not get any meaningful sleep at all. Therefore, it is difficult to say definitively whether or not people under anesthesia get the sleep they need for the day.

Generated by text-davinci-002

Which answer is better?


Question: If under anesthesia do you get your sleep needed for the day?

Answer 1:

No. The unconsciousness induced by both IV anesthetics and inhaled anesthetics is very different from restorative sleep. In fact, patients in ICUs who are sedated for long periods of time with some of the same medications they use in anesthesia get delirium from both a lack of sleep and a lack of proper awake time. Sleep = charging your phone on airplane mode. Anesthesia = turning your phone off but not charging it.

Answer 2:

There is no clear consensus on whether or not people under anesthesia get the sleep they need for the day. Some studies suggest that people under anesthesia do get some sleep, but it is not clear how deep or restful this sleep is. Other studies suggest that people under anesthesia do not get any meaningful sleep at all. Therefore, it is difficult to say definitively whether or not people under anesthesia get the sleep they need for the day.



Step 1: **decompose** long-form text into short units

Asa Graybar is a biological engineer who studies keeping Slider eggs alive and he is accused of a crime at the opening of the story . He thinks he was framed by Tom Dorr , Hazeltynne 's general manager . He was offered one year as a “ changeling ” on another planet or 5 years in rehabilitation on Earth . He elects to do the one year , and thinks that he will get into smuggling Slider eggs on Jordan 's planet . Being a changeling is not a highly sought after line of work , but it pays well , and the people who do it have organs and body parts regenerated to better suit specialized tasks . Asa travels to Jordan 's planet on a spaceship with a cellmate , Kershaw , who got caught stealing a Slider egg and is returning to serve more time ...





Is this *span* fully supported by the source?

Yes

No

Step 2: **verify** each atomic unit

Asa Graybar is a biological engineer who studies keeping Slider eggs alive and he is accused of a crime at the opening of the story . He thinks he was framed by Tom Dorr , Hazeltynne 's general manager . He was offered one year as a “ changeling ” on another planet or 5 years in rehabilitation on Earth . He elects to do the one year , and thinks that he will get into smuggling Slider eggs on Jordan 's planet . Being a changeling is not a highly sought after line of work , but it pays well , and the people who do it have organs and body parts regenerated to better suit specialized tasks . Asa travels to Jordan 's planet on a spaceship with a cellmate , Kershaw , who got caught stealing a Slider egg and is returning to serve more time ...



Step 2: **verify** each atomic unit



Is this *span* fully supported by the source?

Yes

No

Asa Graybar is a biological engineer who studies keeping Slider eggs alive and he is accused of a crime at the opening of the story . He thinks he was framed by Tom Dorr , Hazeltynne 's general manager . He was offered one year as a “ changeling ” on another planet or 5 years in rehabilitation on Earth . He elects to do the one year , and thinks that he will get into smuggling Slider eggs on Jordan 's planet . Being a changeling is not a highly sought after line of work , but it pays well , and the people who do it have organs and body parts regenerated to better suit specialized tasks . Asa travels to Jordan 's planet on a spaceship with a cellmate , Kershaw , who got caught stealing a Slider egg and is returning to serve more time ...



Step 2: **verify** each atomic unit



Is this *span* fully supported by the source?

Yes

No

Asa Graybar is a biological engineer who studies keeping Slider eggs alive and he is accused of a crime at the opening of the story . He thinks he was framed by Tom Dorr , Hazelyne 's general manager . He was offered one year as a “ changeling ” on another planet or 5 years in rehabilitation on Earth . He elects to do the one year , and thinks that he will get into smuggling Slider eggs on Jordan 's planet . Being a changeling is not a highly sought after line of work , but it pays well , and the people who do it have organs and body parts regenerated to better suit specialized tasks . Asa travels to Jordan 's planet on a spaceship with a cellmate , Kershaw , who got caught stealing a Slider egg and is returning to serve more time ...



Compute the fraction of atomic units that were verified.


$$F_{\text{summ}} = \frac{1}{|\mathcal{C}_{\text{summ}}|} \sum_{c \in \mathcal{C}_{\text{summ}}} F_c, F_c \in \{0, 1\}$$

Asa Graybar is a biological engineer who studies keeping Slider eggs alive and he is accused of a crime at the opening of the story. He thinks he was framed by Tom Dorr, Hazeltyn's general manager. He was offered one year as a "changeling" on another planet or 5 years in rehabilitation on Earth. He elects to do the one year, and thinks that he will get into smuggling Slider eggs on Jordan's planet. Being a changeling is not a highly sought after line of work, but it pays well, and the people who do it have organs and body parts regenerated to better suit specialized tasks. Asa travels to Jordan's planet on a spaceship with a cellmate, Kershaw, who got caught stealing a Slider egg and is returning to serve more time ...



More feasible to annotate a small fraction of units

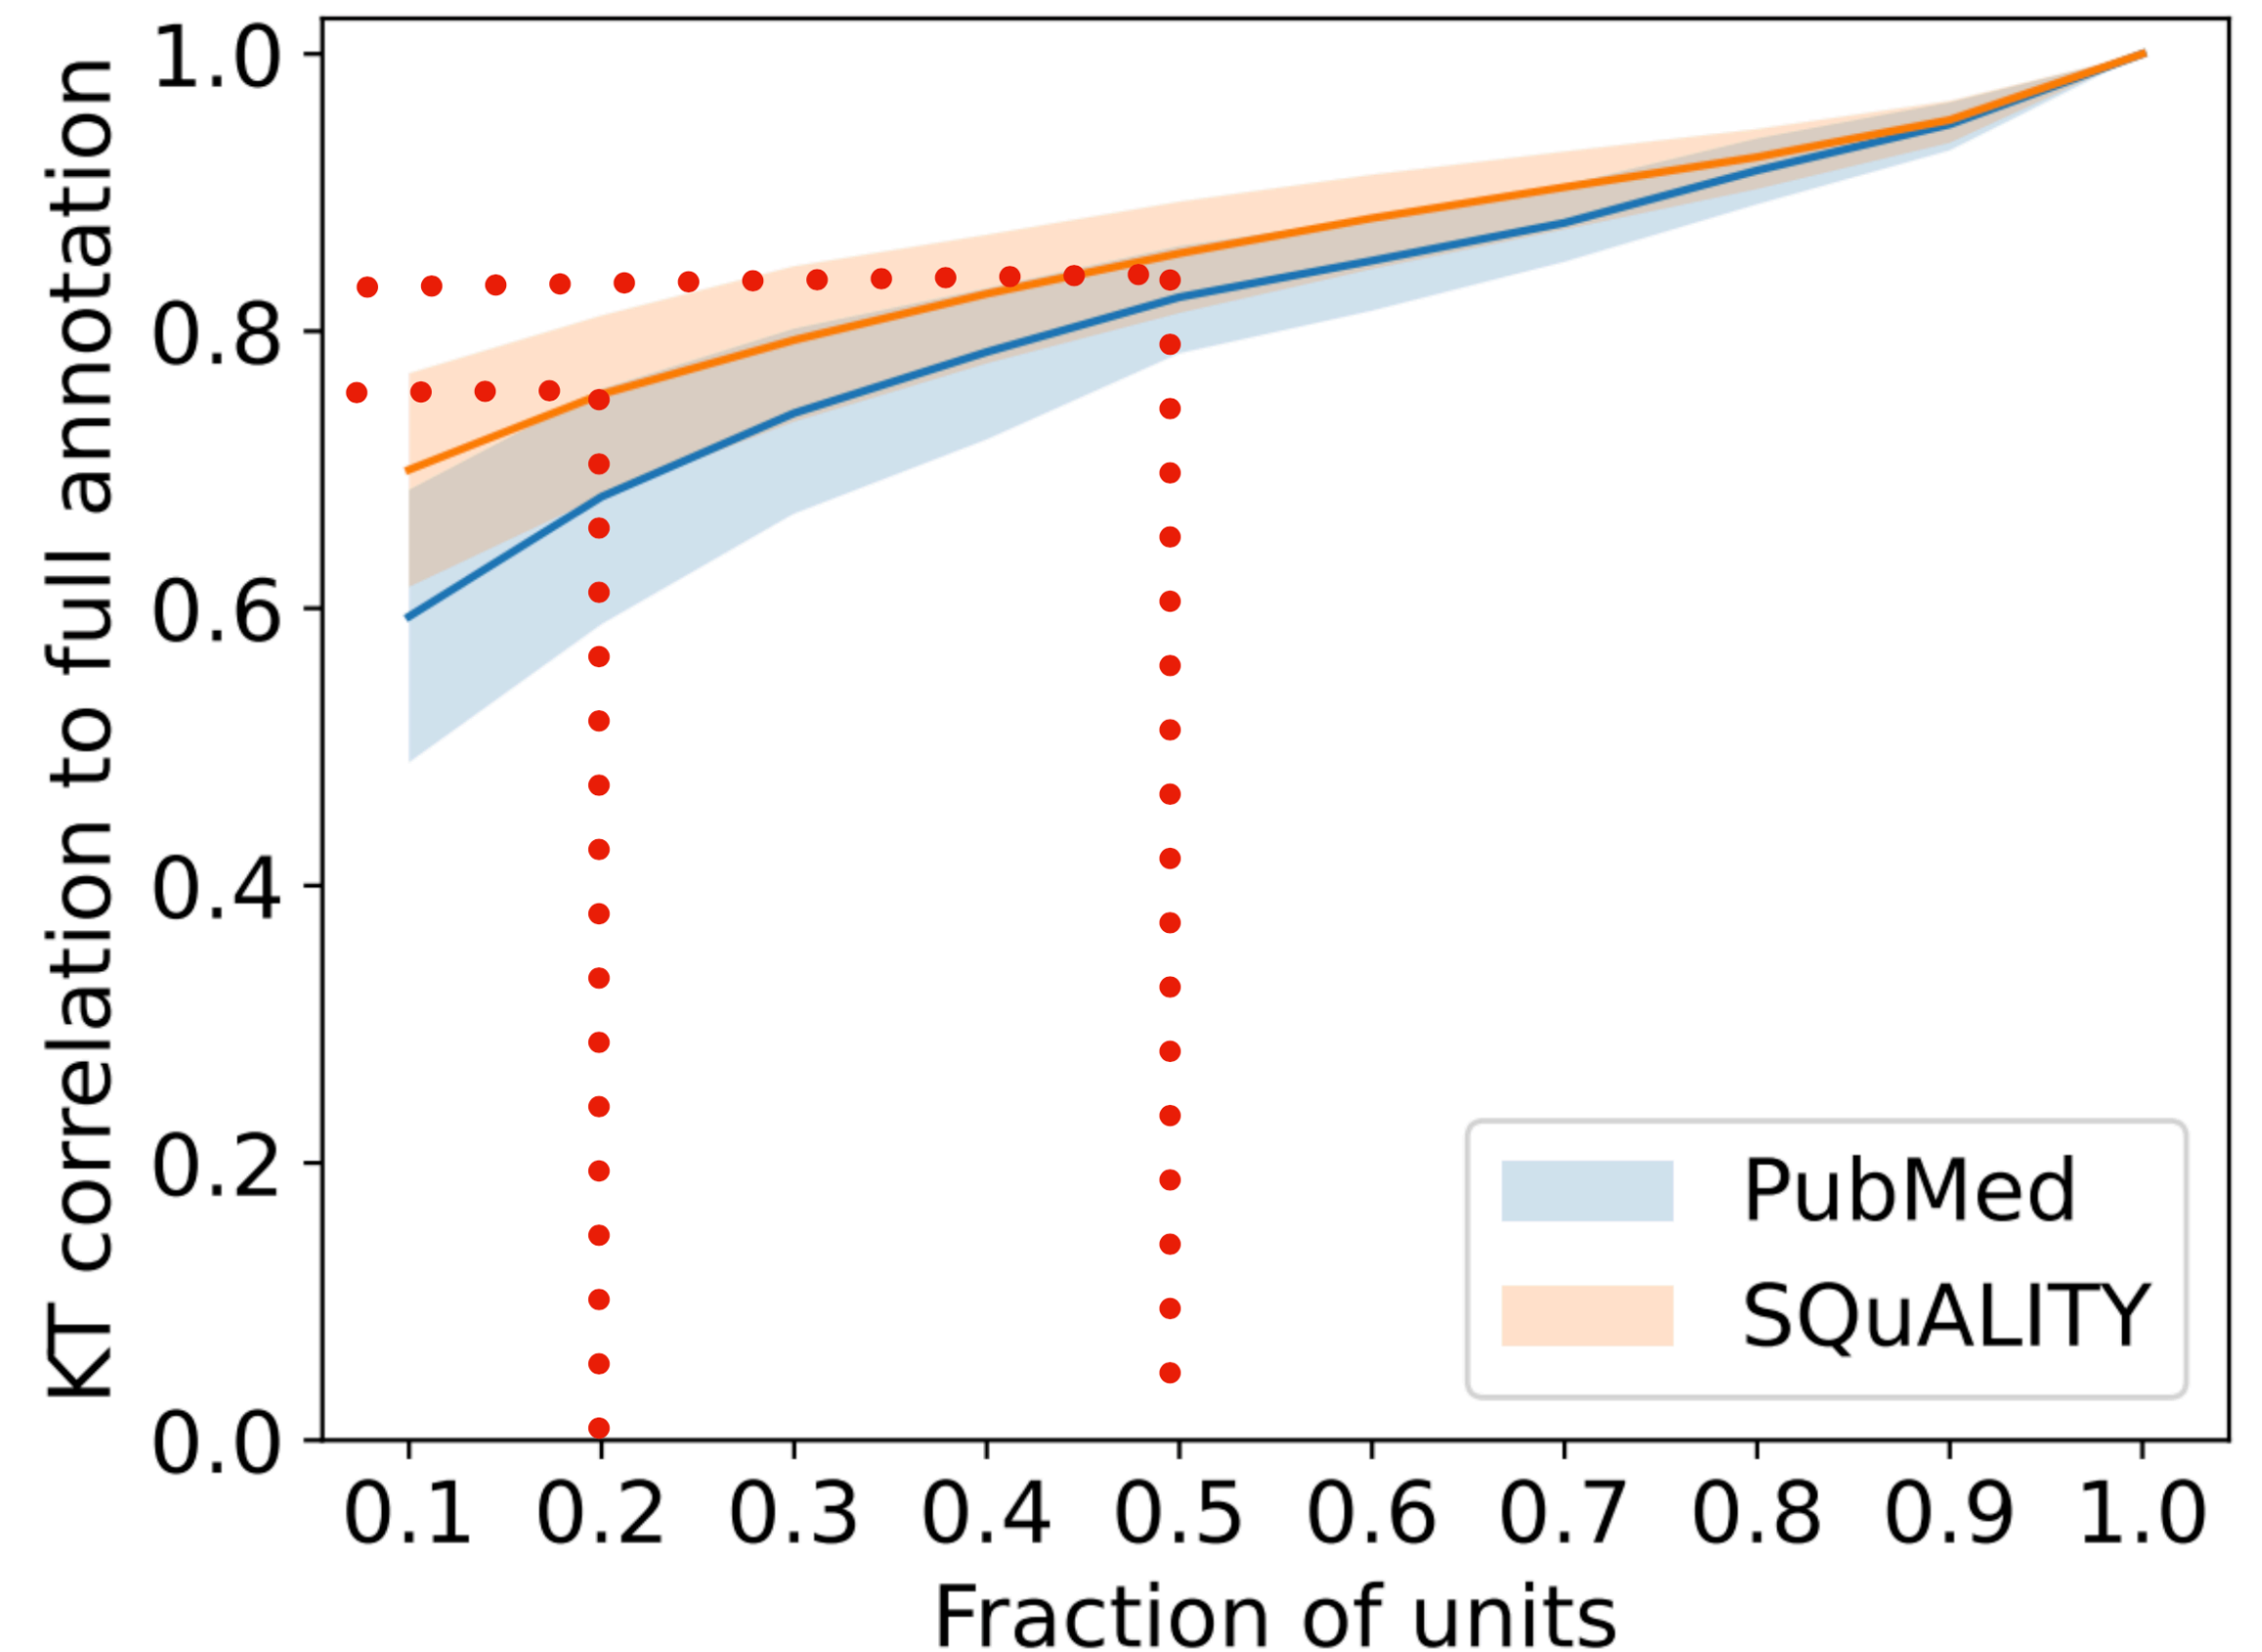
Asa Graybar is a biological engineer who studies keeping Slider eggs alive and he is accused of a crime at the opening of the story. He thinks he was framed by Tom Dorr, Hazeltynes general manager. He was offered one year as a "changeling" on another planet or 5 years in rehabilitation on Earth. He elects to do the one year, and thinks that he will get into smuggling Slider eggs on Jordans planet. Being a changeling is not a highly sought after line of work, but it pays well, and the people who do it have organs and body parts regenerated to better suit specialized tasks. Asa travels to Jordans planet on a spaceship with a cellmate, Kershaw, who got caught stealing a Slider egg and is returning to serve more time ...



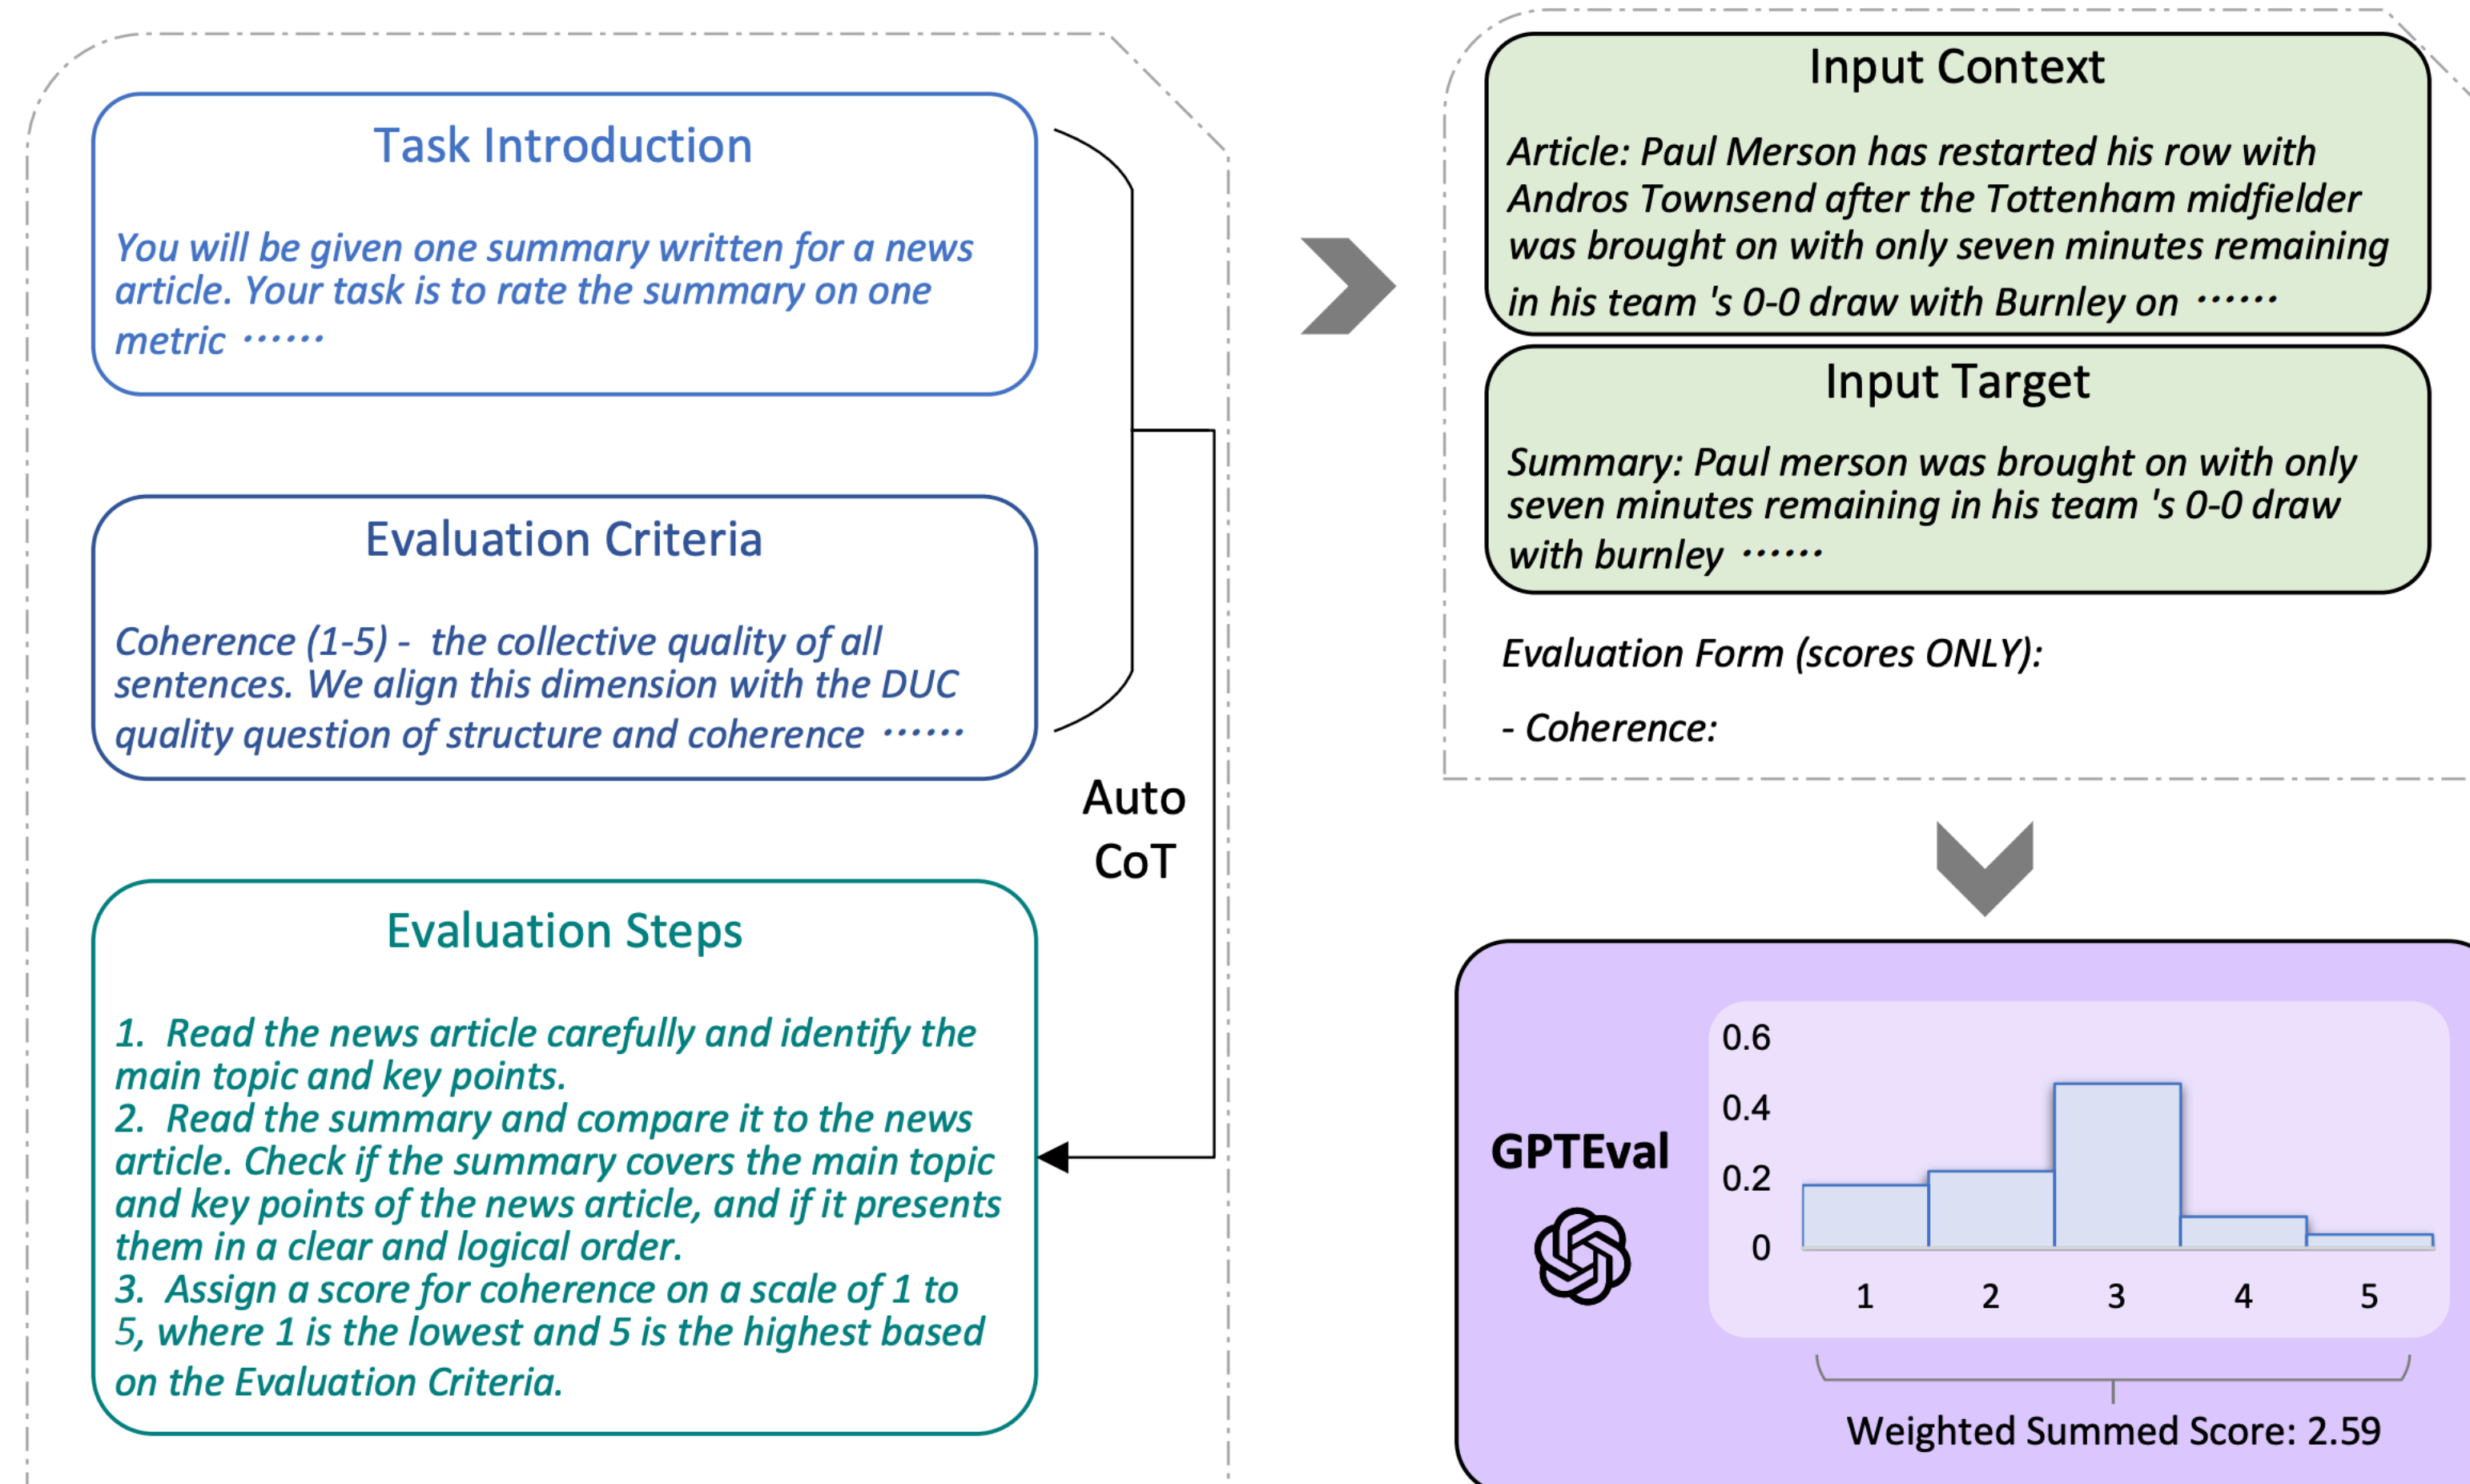
Asa Graybar is a biological engineer who studies keeping Slider eggs alive and he is accused of a crime at the opening of the story. He thinks he was framed by Tom Dorr, Hazeltynes general manager. He was offered one year as a "changeling" on another planet or 5 years in rehabilitation on Earth. He elects to do the one year, and thinks that he will get into smuggling Slider eggs on Jordans planet. Being a changeling is not a highly sought after line of work, but it pays well, and the people who do it have organs and body parts regenerated to better suit specialized tasks. Asa travels to Jordans planet on a spaceship with a cellmate, Kershaw, who got caught stealing a Slider egg and is returning to serve more time ...

Our approach, **LongEval**, provides scalable human evaluation for summarization faithfulness.

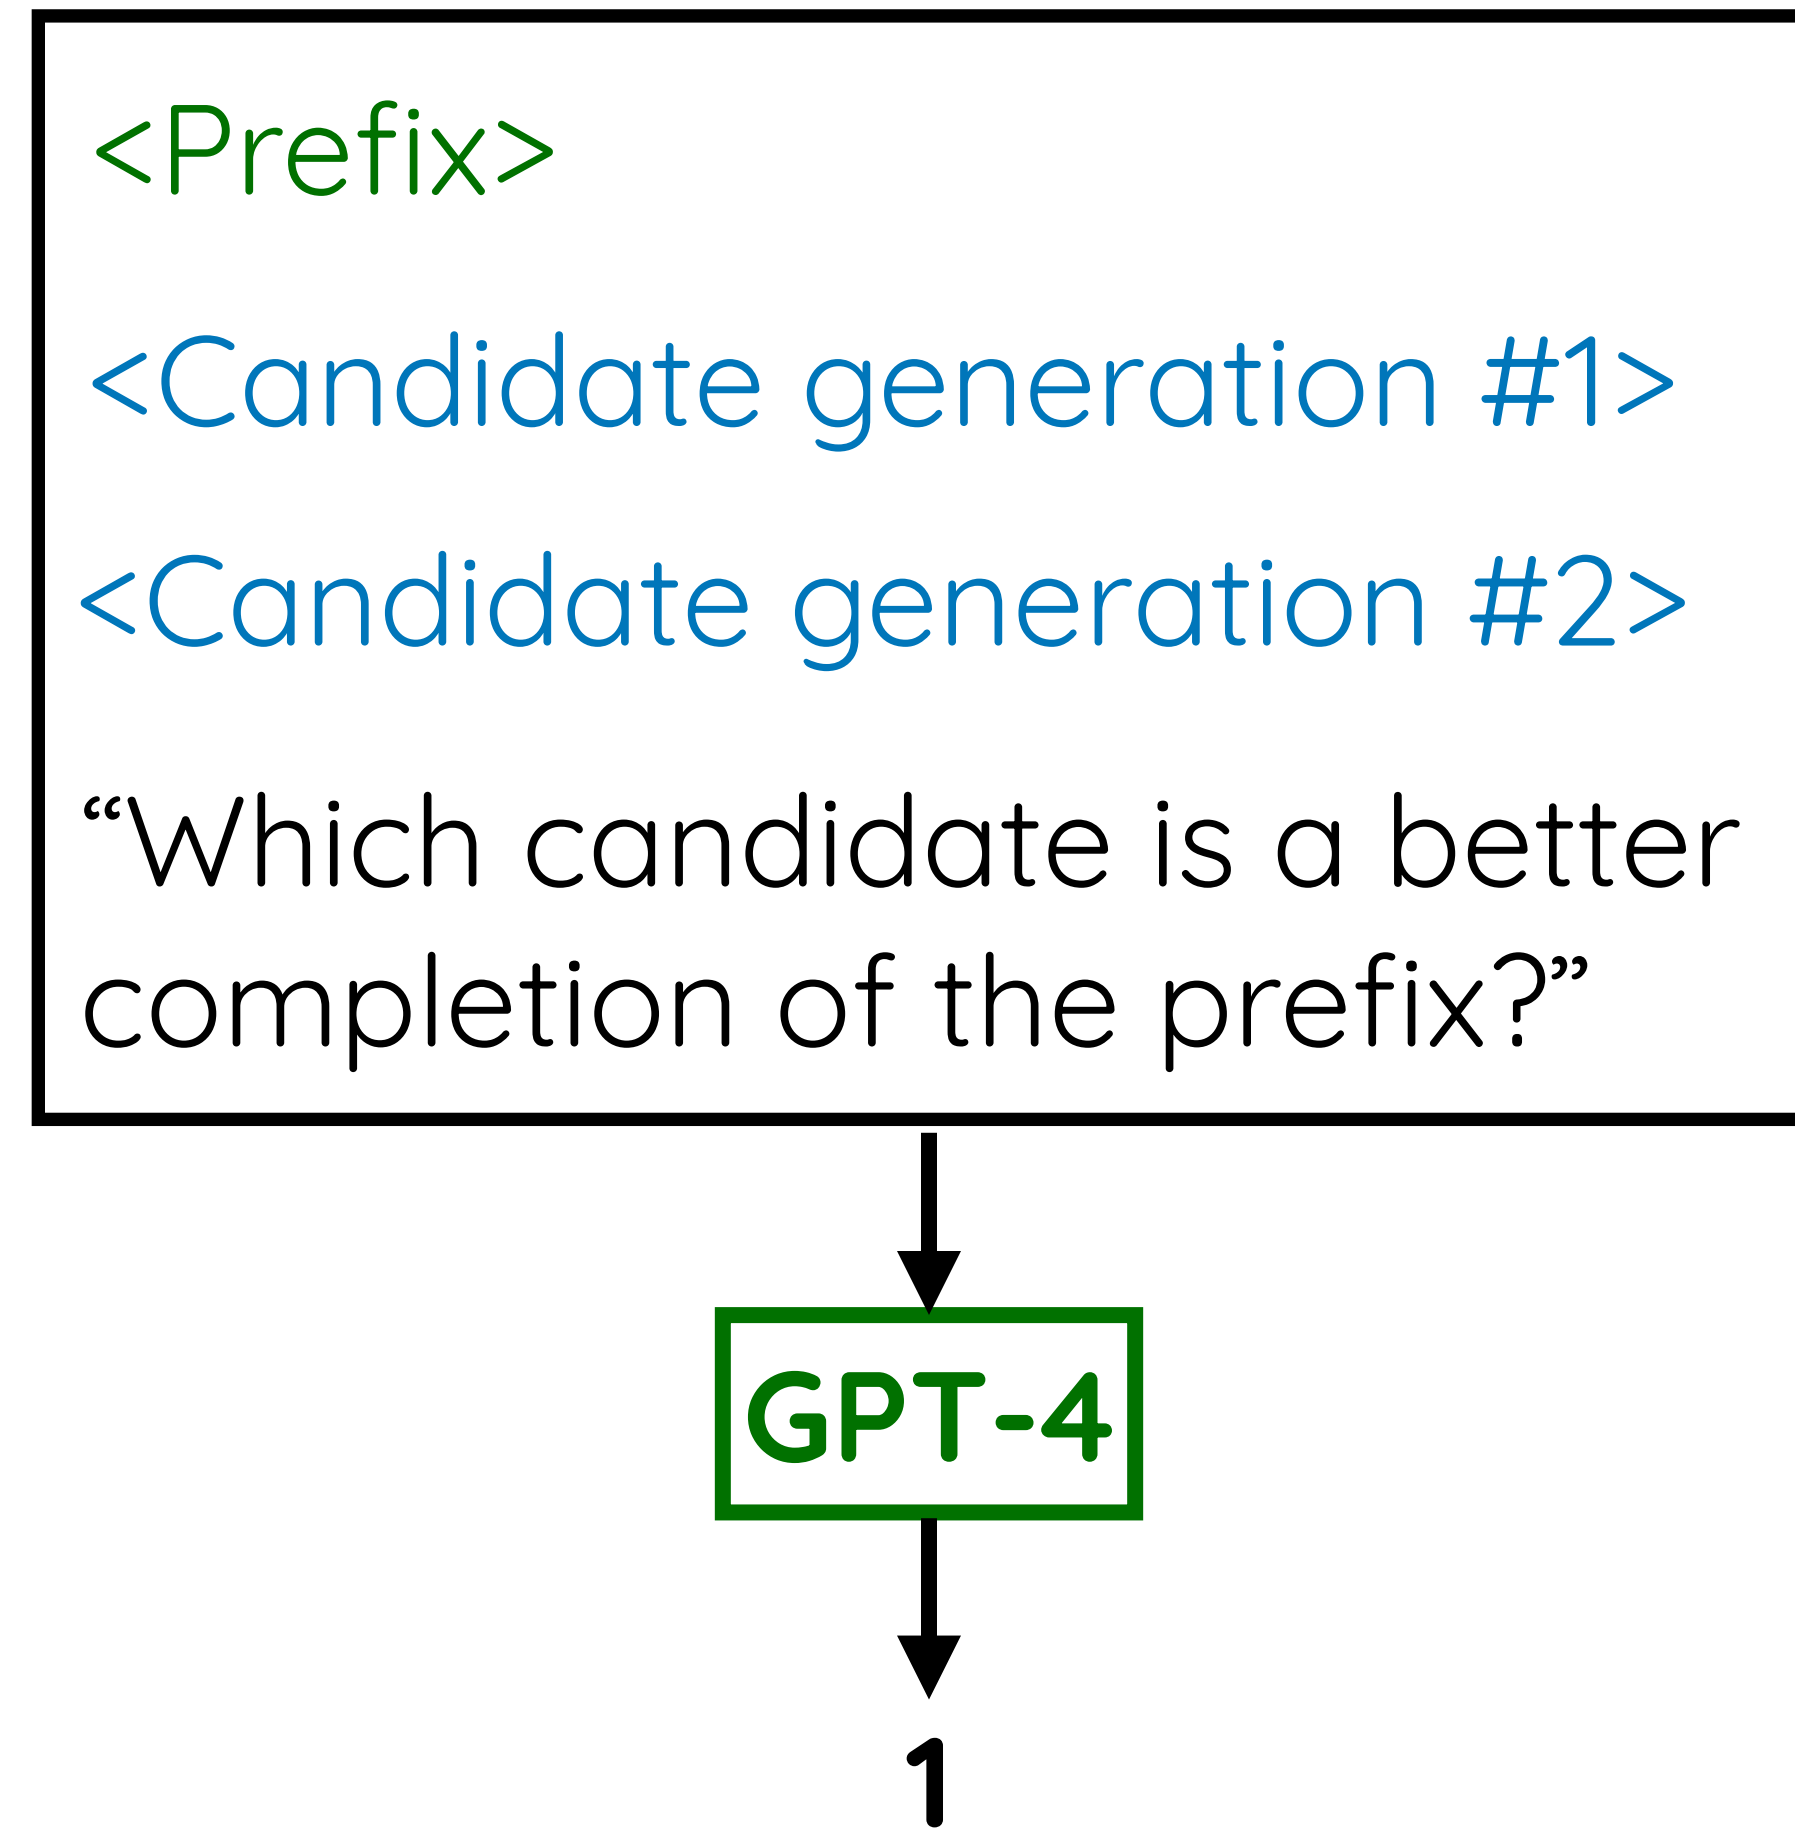
Asa Graybar is a biological engineer who studies keeping Slider eggs alive and he is accused of a crime at the opening of the story . He thinks he was framed by Tom Dorr , Hazeltynne 's general manager . He was offered one year as a “ changeling ” on another planet or 5 years in rehabilitation on Earth . He elects to do the one year , and thinks that he will get into smuggling Slider eggs on Jordan 's planet . Being a changeling is not a highly sought after line of work , but it pays well , and the people who do it have organs and body parts regenerated to better suit specialized tasks . Asa travels to Jordan 's planet on a spaceship with a cellmate , Kershaw , who got caught stealing a Slider egg and is returning to serve more time ...



Can we use LLMs to evaluate generated text?



Most popular LLM judge: **win rate** against a base LM's outputs



Prompt: Write a biography of Bridget Moynahan.

ChatGPT: Bridget Moynahan is an American actress, model and producer. She is best known for her roles in Grey's Anatomy, I, Robot and Blue Bloods. She studied acting at the American Academy of Dramatic Arts, and...

How factually correct is this biography?

Decomposition: Break the generation into atomic facts via few-shot prompting

Bridget Moynahan is an American actress, model and producer. She is best known for her roles in Grey's Anatomy, I, Robot and Blue Bloods. She studied acting at the American Academy of Dramatic Arts, and...



- Bridget Moynahan is American.
- Bridget Moynahan is an actress.
- Bridget Moynahan is a model.
- Bridget Moynahan is a producer.
- Bridget Moynahan is best known for her roles in Grey's Anatomy.
- Bridget Moynahan is best known for her roles in I, Robot.
- Bridget Moynahan is best known for her roles in Blue Bloods.
- Bridget Moynahan studied acting.
- Bridget Moynahan studied at the American Academy of Dramatic Arts.



Decomposition: Break the generation into atomic facts via few-shot prompting

Bridget Moynahan is an American actress, model and producer. She is best known for her roles in Grey's Anatomy, I, Robot and Blue Bloods. She studied acting at the American Academy of Dramatic Arts, and...



- Bridget Moynahan is American.
- Bridget Moynahan is an actress.
- Bridget Moynahan is a model.
- Bridget Moynahan is a producer.
- Bridget Moynahan is best known for her roles in Grey's Anatomy.
- Bridget Moynahan is best known for her roles in I, Robot.
- Bridget Moynahan is best known for her roles in Blue Bloods.
- Bridget Moynahan studied acting.
- Bridget Moynahan studied at the American Academy of Dramatic Arts.

Verification: Retrieve Wikipedia passages for each atomic fact. Prompt LLM to generate True or False given top-k passages and fact.

Bridget Moynahan is best known for her roles in Grey's Anatomy.



Bridget Moynahan

Article Talk

Read Edit View history Tools

From Wikipedia, the free encyclopedia

Kathryn Bridget Moynahan (born April 28, 1971) is an American actress and former model. She graduated from [Longmeadow High School](#) in Massachusetts in 1989 and began pursuing a career in modeling. Moynahan appeared in department-store catalogs and magazines, and after doing television commercials, began taking acting lessons. She made her television debut in a guest appearance in the comedy series *Sex and the City* in 1999, where she later had a recurring role as Natasha.

Moynahan made her feature-film debut in *Coyote Ugly* (2000). She had supporting roles in *Serendipity* (2001); *The Sum of All Fears* (2002); *The Recruit* (2003); *I, Robot* (2004); *Lord of War* (2005); *Grey Matters* (2006); *Prey* (2007); *Noise* (2007); *Ramona and Beezus* (2010); *John Wick* (2014); *The Journey Home* (2014) and *John Wick: Chapter 2* (2017).

Moynahan starred in the ABC television series *Six Degrees*, which premiered in September 2006, and was taken off the schedule after just eight episodes. Since September 2010, she has starred as [Erin Reagan](#) in the CBS police drama *Blue Bloods*.


Bridget Moynahan

Moynahan in 2016

Born	April 28, 1971 (age 52) ^[1] Binghamton, New York, U.S.
Education	Longmeadow High School



False



Bridget Moynahan

From Wikipedia, the free encyclopedia

Kathryn Bridget Moynahan (born April 28, 1971) is an American actress and former model. She graduated from [Longmeadow High School](#) in Massachusetts in 1989 and began pursuing a career in modeling. Moynahan appeared in department-store catalogs and magazines, and after doing television commercials, began taking acting lessons. She made her television debut in a guest appearance in the comedy series *Sex and the City* in 1999, where she later had a recurring role as Natasha.

Moynahan made her feature-film debut in *Coyote Ugly* (2000). She had supporting roles in *Serendipity* (2001); *The Sum of All Fears* (2002); *The Recruit* (2003); *I, Robot* (2004); *Lord of War* (2005); *Grey Matters* (2006); *Prey* (2007); *Noise* (2007); *Ramona and Beezus* (2010); *John Wick* (2014); *The Journey Home* (2014) and *John Wick: Chapter 2* (2017).

Moynahan starred in the [ABC](#) television series *Six Degrees*, which premiered in September 2006, and was taken off the schedule after just eight episodes. Since September 2010, she has starred as [Erin Reagan](#) in the [CBS](#) police drama *Blue Bloods*.

Bridget Moynahan



Moynahan in 2016

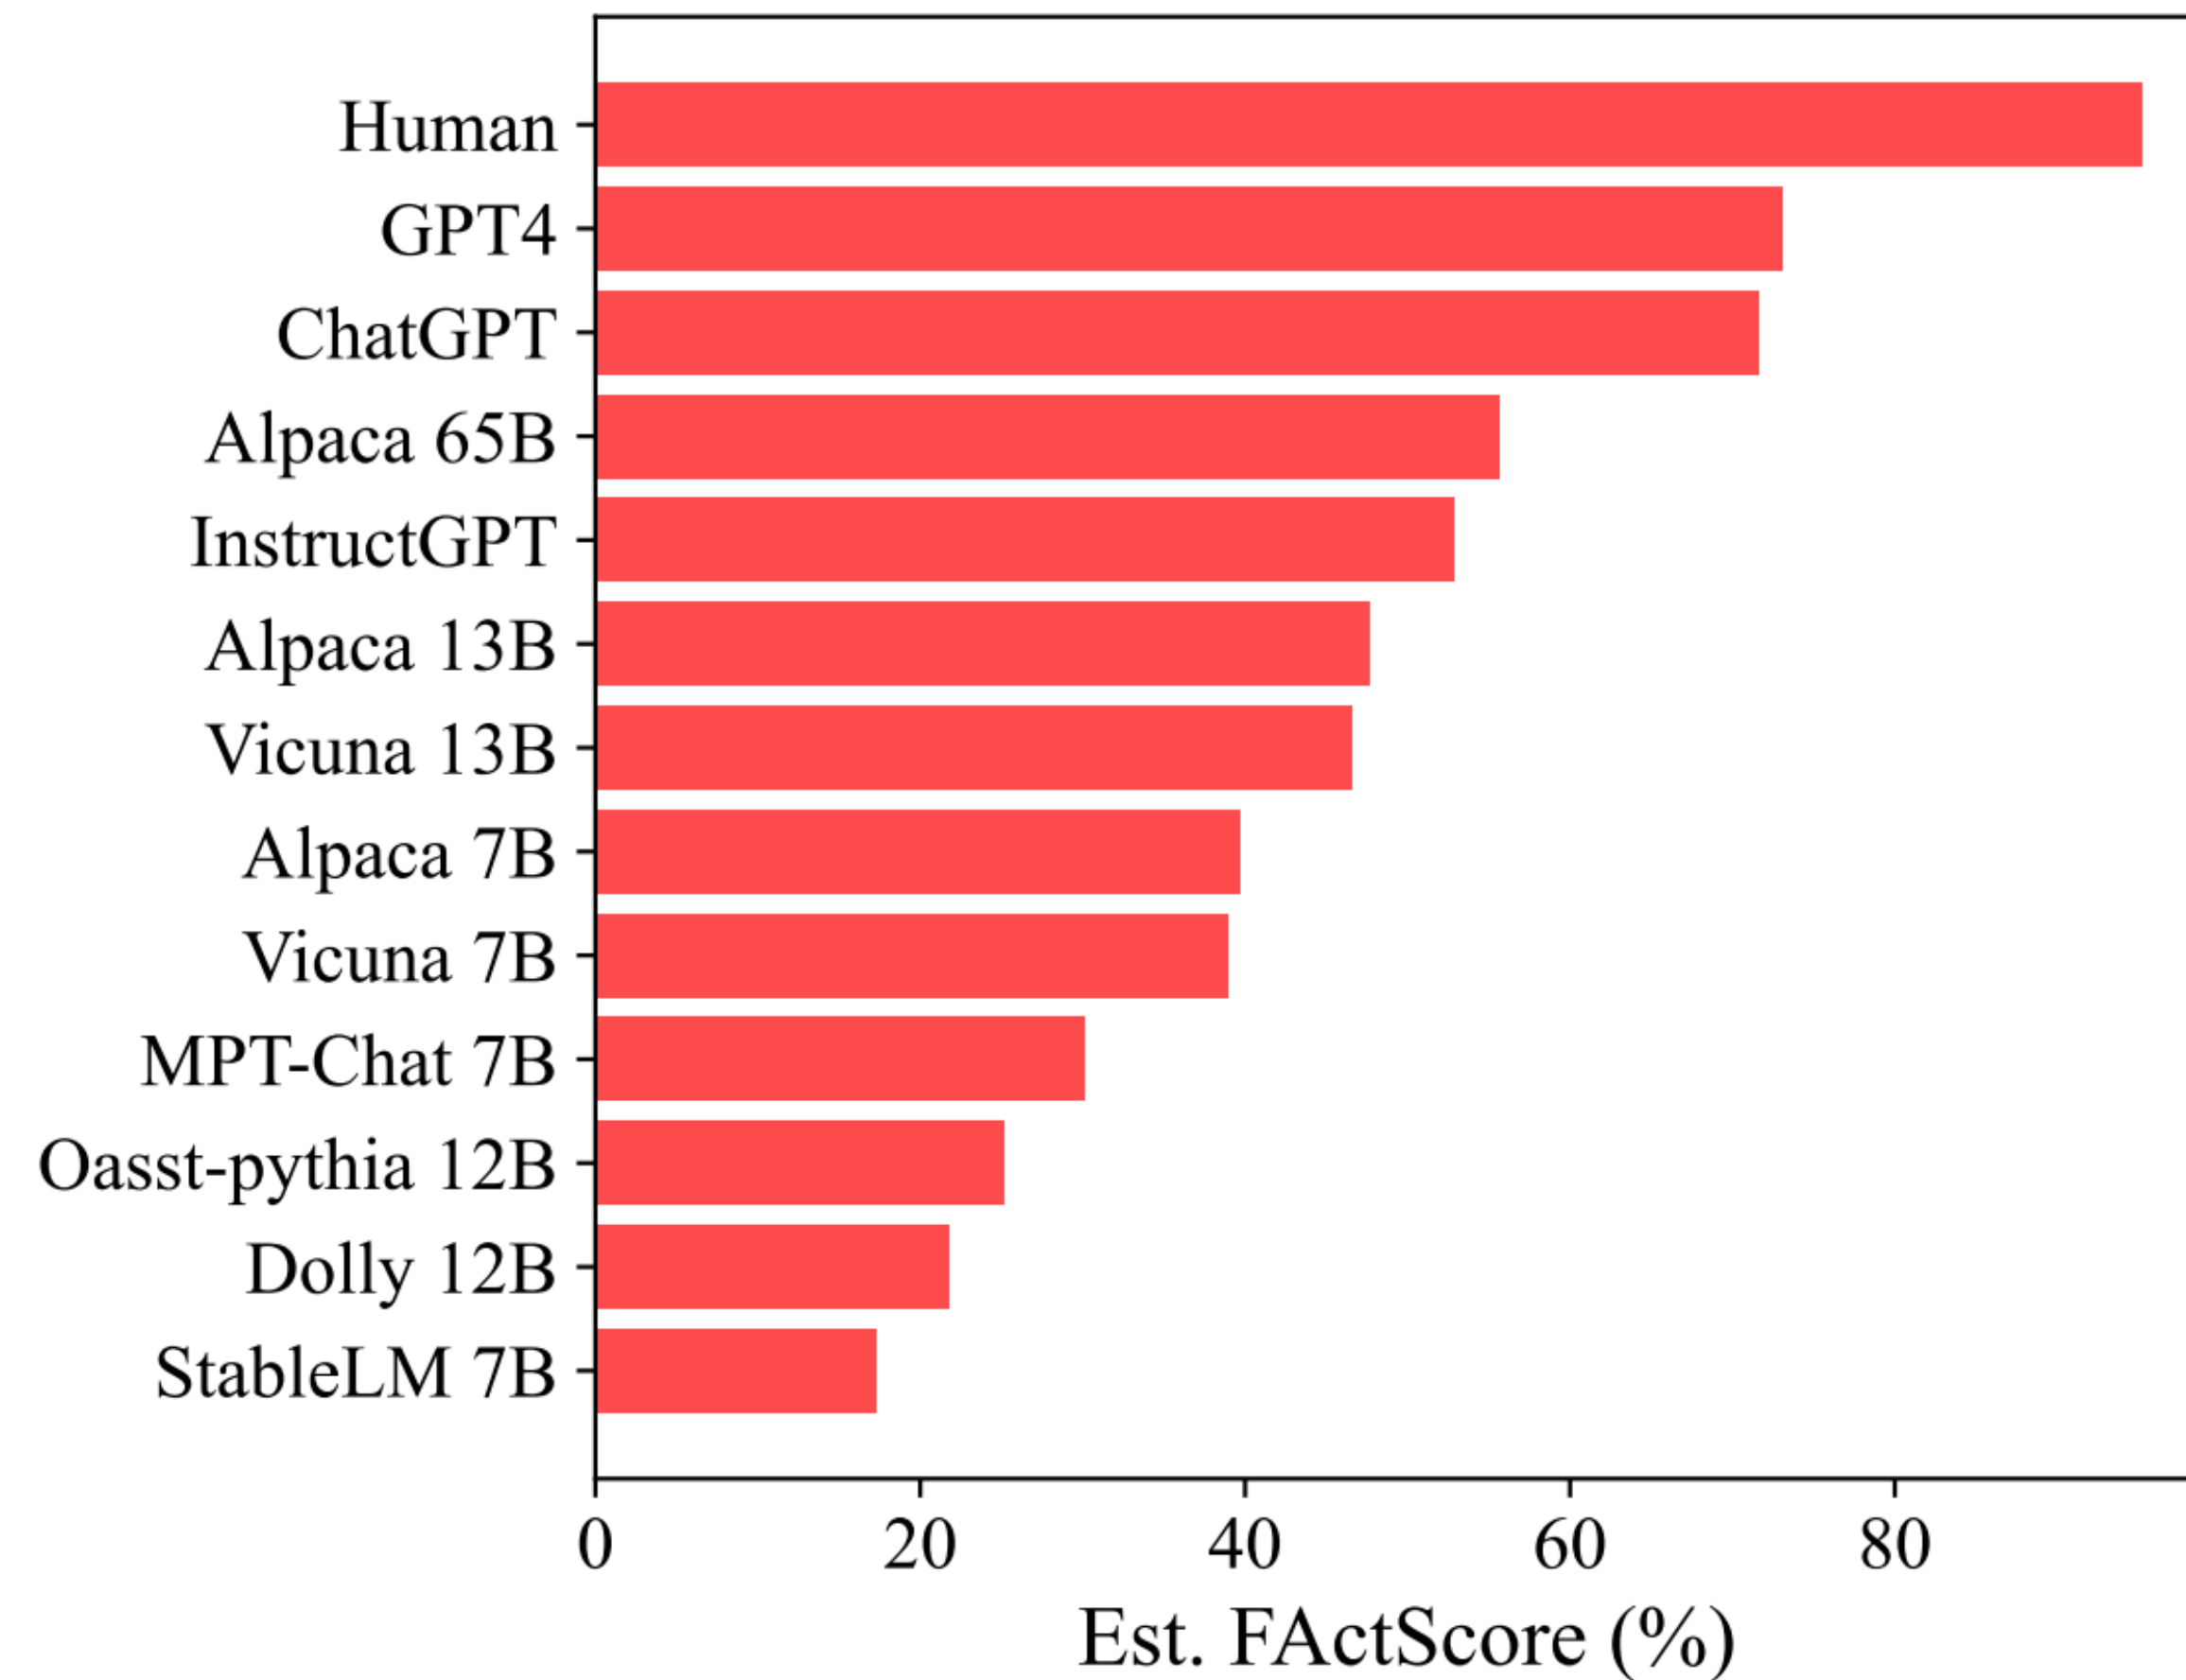
Born	April 28, 1971 (age 52) ^[1] Binghamton, New York, U.S.
Education	Longmeadow High School

- Bridget Moynahan is American.
- Bridget Moynahan is an actress.
- Bridget Moynahan is a model.
- ~~Bridget Moynahan is a producer.~~
- ~~Bridget Moynahan is best known for her roles in Grey's Anatomy.~~
- Bridget Moynahan is best known for her roles in I, Robot.
- Bridget Moynahan is best known for her roles in Blue Bloods.
- Bridget Moynahan studied acting.
- ~~Bridget Moynahan studied at the American Academy of Dramatic Arts.~~



6 of the 9 atomic facts are supported by Wikipedia

FActScore: Implement verifier with LLaMA-7B, error rate of **<2%** compared to human annotations



FActScore: Implement verifier with LLaMA-7B, error rate of **<2%** compared to human annotations

