

Controllable Text Generation

Should machines reflect the way humans interact in society?

Shrimai Prabhumoye



Carnegie Mellon University


Language Technologies Institute

We constantly author text!

The key to human-like generation is
Style, Content and Structure



Social Media



Personal/Professional documents



Blogs



News/Wiki articles

Style

- Expressed in the choice of *words or phrases* as well as *syntactic structures* used to convey information
- I define style as a group of natural language sentences that belong to a particular class or label.



“Do you have any code that we can look at?”



“Any code? Or do I have to guess at how you did it?”

(Danescu-Niculescu-Mizil, 2013)

Content


- Human communication by definition is a process by which individuals *exchange information* and influence one another through a common system of symbols and signs.
- Leverage information in the unstructured form



Wikipedia



Gutenberg



Encyclopedia



Internet

Content

Alice: The Notebook is hands-down one of my favorite movies EVER! Have you ever seen The Notebook?

Bob: No I have never seen this movie. I am going to try it out now

Alice: It was a heartwarming story of young love. The main characters are played by [Ryan Gosling and Rachel McAdams](#).

Bob: Ok this sounds nice. I think Ryan is a good actor

Alice: For all the praise it received, I was surprised to see that it only got a [5.7/10 on Rotten Tomatoes](#).

Bob: That is interesting. They never get the rating correct.

Alice: [The story goes back and forth between present day and the past. Older Ryan is played by James Garner and older Rachel is played by Gena Rowlands](#). Yeah, Rotten Tomatoes never gets the right ratings. I always like to see the ratings but if I want to see a movie, I will watch it even if it has a bad rating.


The Notebook

From Wikipedia, the free encyclopedia

For other uses, see [Notebook \(disambiguation\)](#).

The Notebook is a 2004 romantic drama film directed by [Nick Cassavetes](#) and written by [Jeremy Leven](#) from [Jan Sardi](#)'s adaptation of the 1996 novel by [Nicholas Sparks](#). The film stars [Ryan Gosling](#) and [Rachel McAdams](#) as a young couple who fall in love in the 1940s. Their story is narrated from the present day by an elderly man (played by [James Garner](#)), telling the tale to a fellow nursing home resident (played by [Gena Rowlands](#), who is [Cassavetes's](#) mother).

The Notebook received generally mixed reviews, but performed well at the box office and received a number of award nominations, winning eight [Teen Choice Awards](#), a [Satellite Award](#), and an [MTV Movie Award](#). The film became a [steep hit](#)^{[3][4]} and has gained a [cult following](#).^{[5][6]} On November 11, 2012, [ABC Family](#) premiered an extended version with [deleted scenes](#) added back into the original storyline.^[7]



Reception [edit]

Box office [edit]

The film premiered June 25, 2004, in the United States and Canada and grossed \$13,464,745 in 2,303 theaters its opening weekend, ranking number 4 at the box office.^[33] The film grossed a total of \$115,603,229 worldwide, \$81,001,787 in Canada and the United States and \$34,601,442 in other countries.^[2] It is the 15th highest-grossing romantic drama film of all time.^[34]


Critical reception [edit]

The Notebook received a mixed reaction from film critics. The 178 reviews on review aggregator [Rotten Tomatoes](#) show that 53% of critics gave the film a positive review, with an average rating of 5.64/10 and the website's consensus stating "It's hard not to admire its unabashed sentimentality, but *The Notebook* is too clumsily manipulative to rise above its melodramatic clichés."^[35] At [Metacritic](#), which assigns an average rating out of 100 to reviews from mainstream critics, the film currently holds an average score of 53, based on 34 reviews, which indicates "mixed or average reviews."^[36]




The performances of [Ryan Gosling](#) and [Rachel McAdams](#), as well as their on-screen chemistry, were particularly praised by most film critics

Structure




Application



Application

I have a talk to
present today!



Application



I have a talk to present today!



Best of luck!

Application



I have a talk to
present today!




Best of luck!

Style

Application



I have a talk to
present today!




May the force be with
you!

✓
Style

Application

Hey, have you
watched Avengers
Endgame



Application



Hey, have you
watched Avengers
Endgame




Yeah, I loved Bradley
Cooper in it!

Application



Hey, have you
watched Avengers
Endgame



Yeah, I loved Bradley
Cooper in it!


Content

Application



Hey, have you
watched Avengers
Endgame




Yeah, I loved Robert
Downey in it!



Content

Application

Hey, have you
watched Avengers
Endgame



Application




Hey, have you
watched Avengers
Endgame



Yeah, I loved the
movie!

Application

oh
great! Can you tell me
the story!



Application

oh
great! Can you tell me
the story!



Yeah, Ironman steals the Infinity Stones back from Thanos and uses them to disintegrate Thanos and his army, at the cost of his life. Thor decapitates Thanos. Hulk travels to New York City in 2012 and convinces the Ancient One to give him the Time Stone. Five years later, AntMan escapes from the quantum realm. Ironman builds a time machine to save the world.



Application

oh
great! Can you tell me
the story!



Yeah, Ironman steals the
Infinity Stones back from Thanos
and uses them to disintegrate Thanos
and his army, at the cost of his life. Thor
decapitates Thanos. Hulk travels to New York
City in 2012 and convinces the Ancient One to
give him the Time Stone. Five years
later, AntMan escapes from the quantum
realm. Ironman builds a time
machine to save the world.




Structure



Application

oh
great! Can you tell me
the story!



Yeah, Thor ✓
decapitates Thanos. Five
years later, AntMan escapes from
the quantum realm. Ironman builds a
time machine to save the world. Hulk travels
to New York City in 2012 and convinces
the Ancient One to give him the Time Stone.
Ironman steals the Infinity Stones back
from Thanos and uses them to
disintegrate Thanos and his army,
at the cost of his life.

Structure




Other Applications

- ***Writing Assistance Tools***
 - recommend formal language
 - recommend structural changes
- Recommend ***polite emails***
- ***Story Generation***
 - plot, ending, sentiment, topic, persona
- ***Content Generation*** (websites, descriptions etc)


Overview

Style



ACL '20
Storytelling '19
ACL '18

Content




NAACL '19
EMNLP '18

Structure



ACL '20
draft EMNLP '20

Ethical Considerations




WiNLP '19

**The key to human-like generation is
Style, Content and Structure**


Overview

Style



ACL '20
Storytelling '19
ACL '18

Content




NAACL '19
EMNLP '18

Structure



ACL '20
draft EMNLP '20

Ethical Considerations




WiNLP '19

**The key to human-like generation is
Style, Content and Structure**

Politeness Transfer

Converting non-polite sentences to polite sentences while preserving the meaning

Send me the data



Could you please send me the data?

Challenges

- Politeness is *culturally diverse*

Challenges

- Politeness is *culturally diverse*

Usage of “please”
to closest friends.



Challenges

- Politeness is *culturally diverse*
- Politeness is *subtle*.

Challenges

- Politeness is *culturally diverse*
- Politeness is *subtle*.

Indirect: By the way, do you know when is the deadline?

1st person plural: Let us find a good name for this paper.

Challenges

- Politeness is *culturally diverse*
- Politeness is *subtle*.
- *Ill-defined* non-politeness class

Challenges

- Politeness is *culturally diverse*
- Politeness is *subtle*.
- *Ill-defined* non-politeness class

Existing literature to detect presence
Difficult to objectify absence

Challenges

- Politeness is *culturally diverse*
- Politeness is *subtle*.
- *Ill-defined* non-politeness class
- *Lack* of data

Challenges

- Politeness is *culturally diverse*
- Politeness is *subtle*.
- *Ill-defined* non-politeness class
- *Lack of data*



Overcoming Challenges

- Politeness accepted by *North American* English speakers in *formal* setting


Overcoming Challenges

- Politeness accepted by **North American** English speakers in **formal** setting
 - ✓ Politeness is *culturally diverse*
 - ✓ Politeness is *subtle*.

Overcoming Challenges

- Politeness accepted by *North American* English speakers in *formal* setting
- Convert only from non-polite to polite

Overcoming Challenges

- Politeness accepted by **North American** English speakers in **formal** setting
 - Convert only from non-polite to polite
-  *Ill-defined* non-politeness class

Overcoming Challenges

- Politeness accepted by *North American* English speakers in *formal* setting
- Convert only from non-polite to polite
- Convert *requests* to polite requests

Overcoming Challenges

- Politeness accepted by **North American** English speakers in **formal** setting
- Convert only from non-polite to polite
- Convert **requests** to polite requests

The sun rises in the east.



Overcoming Challenges

- Politeness accepted by **North American** English speakers in **formal** setting
- Convert only from non-polite to polite
- Convert **requests** to polite requests
 - ✓ *Ill-defined* non-politeness class

Tag and Generate Pipeline

Tag and Generate Pipeline



The diagram illustrates the 'Tag' step of a pipeline. It features a white rectangular box with a dark red border containing the word 'Tagger'. Below this box is a dark red callout box with a white border and a pointed top, containing the instruction 'Add TAG words at positions where the polite phrases can appear'.

Tagger

**Add TAG words at
positions where the
polite phrases can
appear**

Tag and Generate Pipeline



Add TAG words at positions where the polite phrases can appear

Tag and Generate Pipeline

Send me
the data



Tagger



[TAG] send
me the data

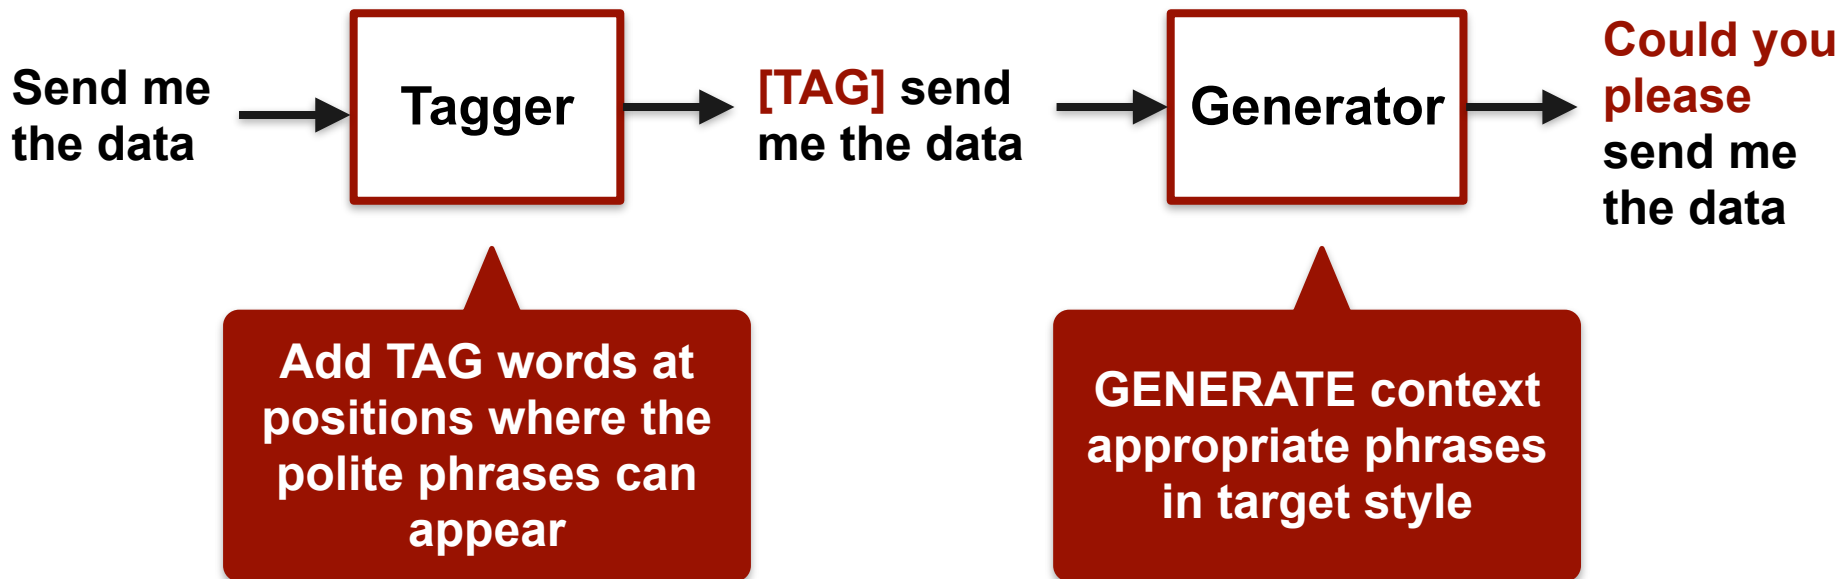


Generator

Add TAG words at
positions where the
polite phrases can
appear

GENERATE context
appropriate phrases
in target style

Tag and Generate Pipeline



Tag and Generate Pipeline

Tagger and Generator are transformers

Send me
the data

Tagger

[TAG] send
me the data

Generator

Could you
please
send me
the data

Add TAG words at
positions where the
polite phrases can
appear

GENERATE context
appropriate phrases
in target style

Dataset Creation

- Enron Corpus
 - North American culture
 - Formal setting
 - Plenty requests
- Politeness classifier (Niu and Bansal, 2018)

Dataset Creation

- Enron Corpus
 - North American culture
 - Formal setting
 - Plenty requests
- Politeness classifier (Niu and Bansal, 2018)

1.39 million sentences!

Dataset Creation



Dataset Creation

Could you please send me the data



Dataset Creation

Could you please send me the data

Identify markers of politeness



Dataset Creation

Could you please send me the data

Remove politeness markers

Send me the data

Tagger

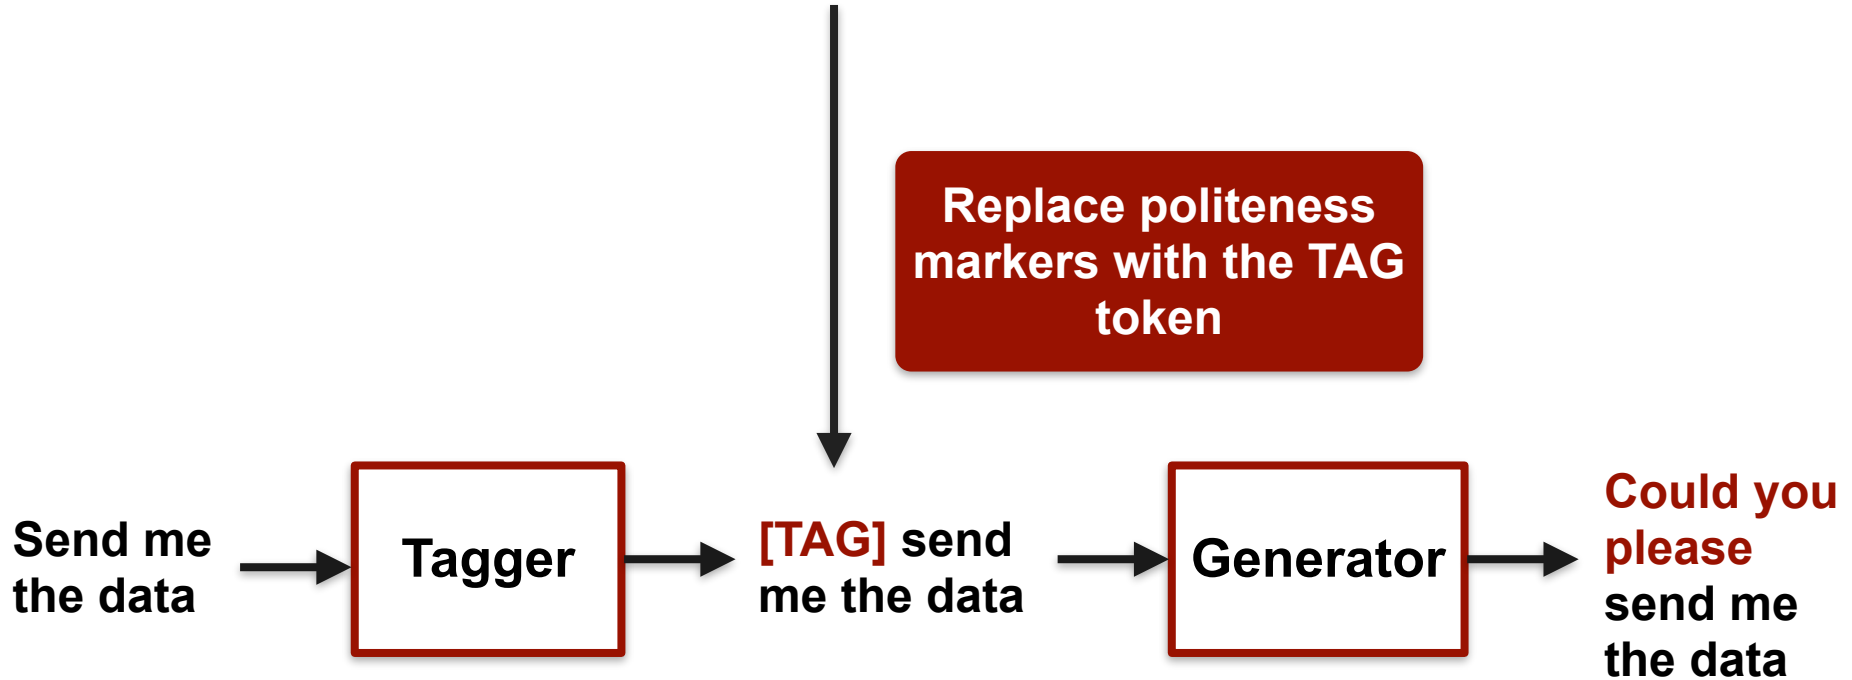
[TAG] send me the data

Generator

Could you please send me the data

Dataset Creation

Could you please send me the data



Dataset Creation

Could you please send me the data

Send me
the data



Tagger

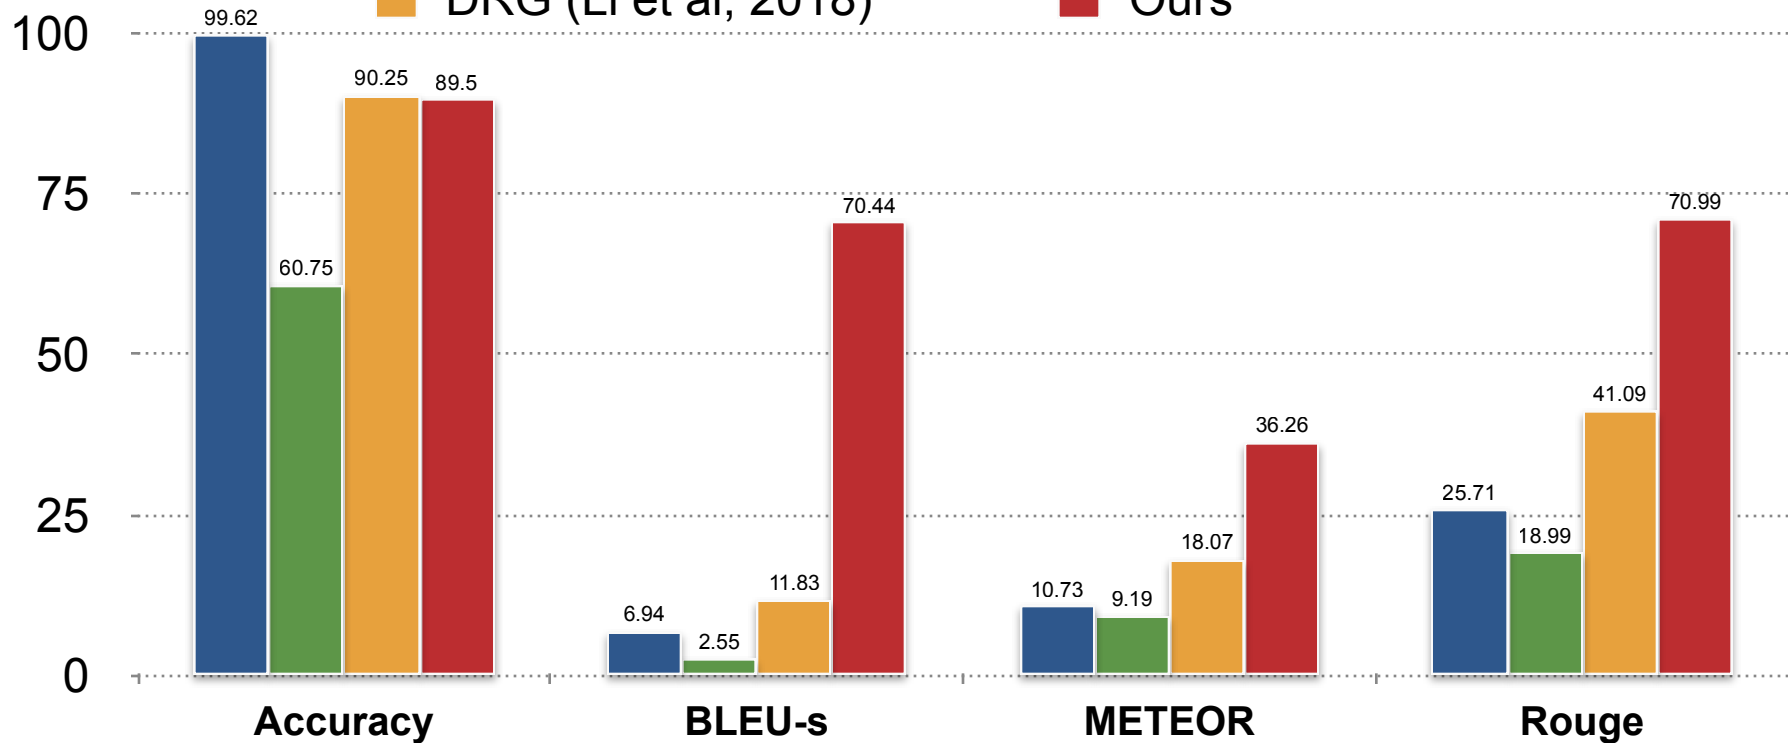
[TAG] send
me the data

Generator

Could you
please
send me
the data

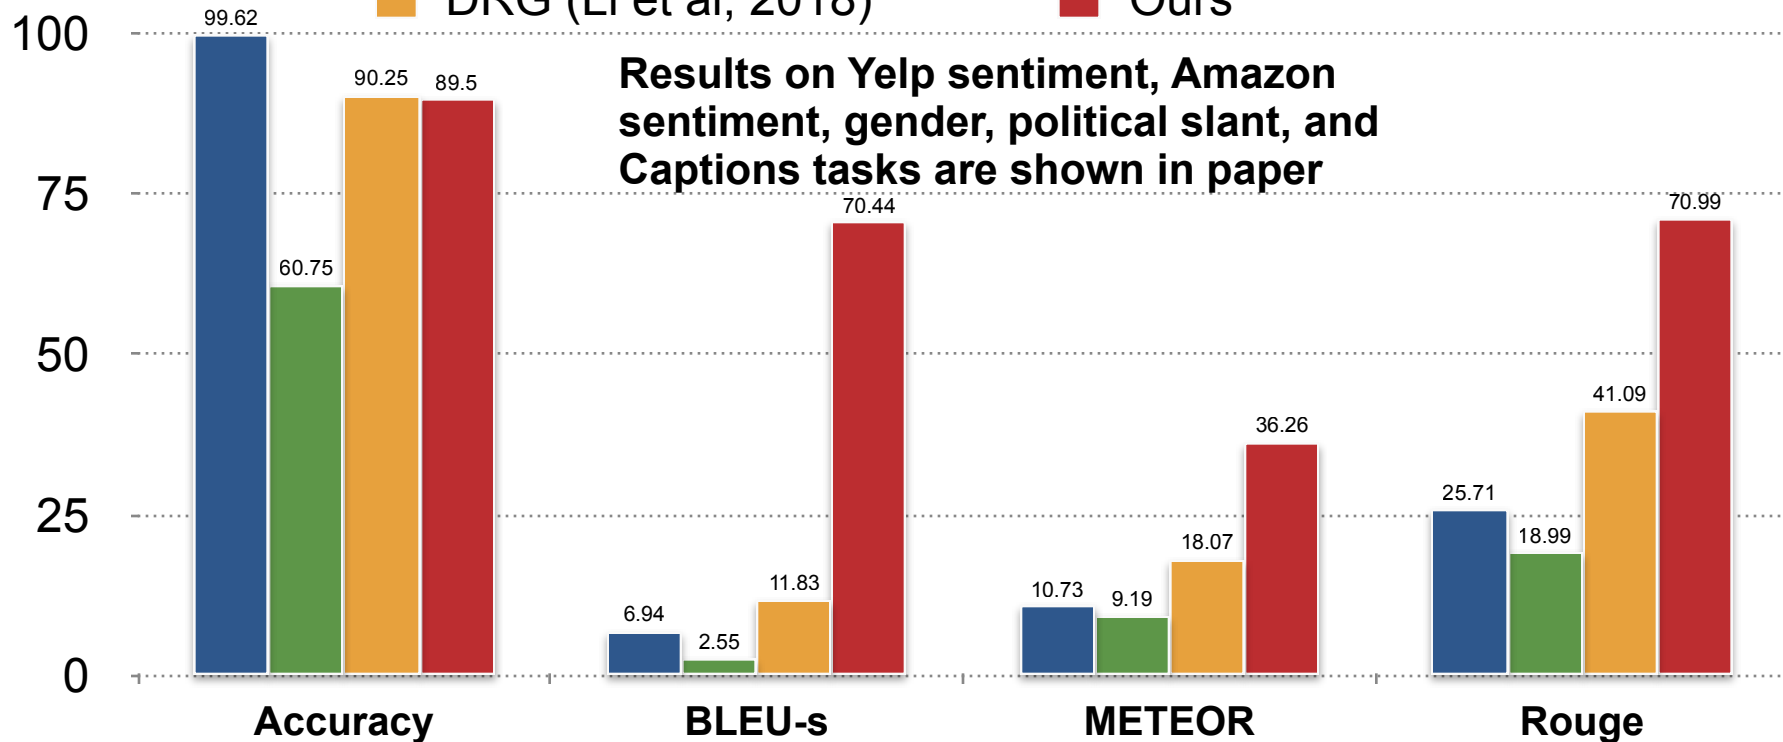
Results

■ CAE (Shen et al, 2017) ■ BST (Prabhumoye et al, 2018)
■ DRG (Li et al, 2018) ■ Ours



Results

■ CAE (Shen et al, 2017) ■ BST (Prabhumoye et al, 2018)
■ DRG (Li et al, 2018) ■ Ours




Examples

| Input | Output | Politeness Strategy |
|---------------------------------------|--|---|
| What happened to my personal station? | Could you please let me know what happened to my personal station? | <i>could</i> and <i>would</i> are counterfactual forms which are polite |
| Yes, go ahead and remove it. | Yes, we can go ahead and remove it. | First person plural |
| Not yet - I'll try this weekend | Sorry, not yet - I'll try to make sure this weekend. | Apologizing |

Overview


Style



ACL '20
Storytelling '19
ACL '18



Content




NAACL '19
EMNLP '18

Structure



ACL '20
draft EMNLP '20

Ethical Considerations




WiNLP '19

The key to human-like generation is
Style, Content and Structure

Overview


Style



ACL '20
Storytelling '19
ACL '18



Content




NAACL '19
EMNLP '18

Structure



ACL '20
draft EMNLP '20

Ethical Considerations



WiNLP '19

The key to human-like generation is
Style, Content and Structure

Content Transfer

After **graduate form** Columbia University, Obama worked in Chicago.

After **graduating from** Columbia University, Obama worked in Chicago.

After graduating from **Carnegie Mellon University**, Obama worked in Chicago.

After graduating from **Columbia University**, Obama worked in Chicago.

- AI assistance deals with *form* (grammar, style, etc.)
- Our goal is to control for *content*

What is our task?

What is our task?



WIKIPEDIA
The Free Encyclopedia

[Main page](#)
[Contents](#)
[Current events](#)
[Random article](#)
[About Wikipedia](#)
[Contact us](#)
[Donate](#)

[Contribute](#)
[Help](#)
[Community portal](#)
[Recent changes](#)
[Upload file](#)

[Tools](#)
[What links here](#)
[Related changes](#)
[Special pages](#)

Not logged in [Talk](#) [Contributions](#) [Create account](#) [Log in](#)

Article [Talk](#) [Read](#) [Edit](#) [View history](#)

Monkey selfie copyright dispute

From Wikipedia, the free encyclopedia

The **monkey selfie copyright dispute** is a series of disputes about the [copyright](#) status of [selfies](#) taken by [Celebes crested macaques](#) using equipment belonging to the British nature photographer David Slater. The disputes involve [Wikimedia Commons](#) and the blog *Techdirt*, which have hosted the images following their publication in newspapers in July 2011 over Slater's objections that he holds the copyright, and [People for the Ethical Treatment of Animals](#) (PETA), who have argued that the macaque should be assigned the copyright.

Slater has argued that he has a valid copyright claim, as he engineered the situation that resulted in the pictures by travelling to Indonesia, befriending a group of wild macaques, and setting up his camera equipment in such a way that a "selfie" picture might come about. The [Wikimedia Foundation's](#) 2014 refusal to remove the pictures from its Wikimedia Commons image library was based on the understanding that copyright is held by the creator, that a non-human creator (not being a [legal person](#)) cannot hold copyright, and that the images are thus in the [public domain](#).



One of the monkey selfies at issue ↗

What is our task?

Not logged in | Talk | Contributions | Create account | Log in

Article **Talk** | Read | Edit | View history | Search Wikipedia

Monkey selfie copyright dispute

From Wikipedia, the free encyclopedia

The **monkey selfie copyright dispute** is a series of disputes about the [copyright](#) status of [selfies](#) taken by [Celebes crested macaques](#) using equipment belonging to the British nature photographer David Slater. The disputes involve [Wikimedia Commons](#) and the blog *Techdir*, which have hosted the images following their publication in newspapers in July 2011 over Slater's objections that he holds the copyright, and [People for the Ethical Treatment of Animals](#) (PETA), who have argued that the macaque should be assigned the copyright.



Slater has argued that he has a valid copyright claim, as he engineered the situation that resulted in the pictures by travelling to Indonesia, befriending a group of wild macaques, and setting up his camera equipment in such a way that a "selfie" picture might come about. The [Wikimedia Foundation's](#) 2014 refusal to remove the pictures from its Wikimedia Commons image library was based on the understanding that copyright is held by the creator, that a non-human creator (not being a [legal person](#)) cannot hold copyright, and that the images are thus in the [public domain](#).

One of the monkey selfies at issue

WIKIPEDIA
The Free Encyclopedia

- Main page
- Contents
- Current events
- Random article
- About Wikipedia
- Contact us
- Donate

Contribute

- Help
- Community portal
- Recent changes
- Upload file

Tools

- What links here
- Related changes
- Social pages

Amateur Photographer

Subscribe now Digital

Home | Latest | Opinion | Reviews | How To | Inspiration | AP Awards | The Video Mode

Home / Latest / Latest news

Ape-rture priority photographer plays down monkey reports

Chris Cheesman
July 5, 2011

0 shares

A photographer who says he witnessed monkeys taking pictures of themselves, tells Amateur Photographer (AP) that much of the media coverage has been exaggerated.

Wildlife photographer David Slater today played down newspaper reports that suggest a bunch of Indonesian monkeys grabbed his camera and began taking self-portraits.

And he revealed that the shots, featuring the endangered crested black macaque monkey, were taken three years ago.

Speaking to AP, David explained that his camera had been mounted on a tripod when the primates

Latest reviews

Leica M10-R review



Andy Westlake takes a detailed look at Leica's latest high-resolution...

Canon EOS R6 review: hands-on first look



Look out Sony - Canon may just have got full-

What is our task?

WIKIPEDIA
The Free Encyclopedia

Not logged in | Talk | Contributions | Create account | Log in

Article | Talk | Read | Edit | View history | Search Wikipedia

Monkey selfie copyright dispute

From Wikipedia, the free encyclopedia

The **monkey selfie copyright dispute** is a series of disputes about the **copyright** status of **selfies** taken by **Celebes crested macaques** using equipment belonging to the British nature photographer David Slater. The disputes involve **Wikimedia Commons** and the blog *Techdirt*, which have hosted the images following their publication in newspapers in July 2011 over Slater's objections that he holds the copyright, and **People for the Ethical Treatment of Animals** (PETA), who have argued that the macaque should be assigned the copyright.



Slater has argued that he has a valid copyright claim, as he engineered the situation that resulted in the pictures by travelling to Indonesia, befriending a group of wild macaques, and setting up his camera equipment in such a way that a "selfie" picture might come about. The **Wikimedia Foundation's** 2014 refusal to remove the pictures from its Wikimedia Commons image library was based on the understanding that copyright is held by the creator, that a non-human creator (not being a **legal person**) cannot hold copyright, and that the images are thus in the **public domain**.

One of the monkey selfies at issue

The following day, *Amateur Photographer* reported that Slater gave them further explanation as to how the photographs were created, downplaying the way newspaper articles had described them; Slater said reports that a monkey ran off with his camera and "began taking self-portraits" were incorrect and that the portrait was shot when his camera had been mounted on a tripod, with the primates playing around with a remote cable release as he fended off other monkeys.^[13]

Amateur Photographer

Subscribe now Digital

Home | Latest | Opinion | Reviews | How To | Inspiration | AP Awards | The Video Mode

Home / Latest / Latest news

Ape-rture priority photographer plays down monkey reports

Chris Cheesman
July 5, 2011

0 shares

A photographer who says he witnessed monkeys taking pictures of themselves, tells Amateur Photographer (AP) that much of the media coverage has been exaggerated.

Wildlife photographer David Slater today played down newspaper reports that suggest a bunch of Indonesian monkeys grabbed his camera and began taking self-portraits.

And he revealed that the shots, featuring the endangered crested black macaque monkey, were taken three years ago.

Speaking to AP, David explained that his camera had been mounted on a tripod when the primates

Latest reviews

Leica M10-R review



Andy Westlake takes a detailed look at Leica's latest high-resolution...

Canon EOS R6 review: hands-on first look



Look out Sony - Canon may just have got full-

What is our task?



The screenshot shows the Wikipedia article "Monkey selfie copyright dispute". The article text states: "The **monkey selfie copyright dispute** is a series of disputes about the **copyright** status of **selfies** taken by **Celebes crested macaques** using equipment belonging to the British nature photographer David Slater. The disputes involve **Wikimedia Commons** and the blog *Techdirt*, which have hosted the images following their publication in newspapers in July 2011 over Slater's objections that he holds the copyright, and **People for the Ethical Treatment of Animals** (PETA), who have argued that the macaque should be assigned the copyright.

Slater has argued that he has a valid copyright claim, as he engineered the situation that resulted in the pictures by travelling to Indonesia, befriending a group of wild macaques, and setting up his camera equipment in such a way that a "selfie" picture might come about. The **Wikimedia Foundation's** 2014 refusal to remove the pictures from its Wikimedia Commons image library was based on the understanding that copyright is held by the creator, that a non-human creator (not being a **legal person**) cannot hold copyright, and that the images are thus in the **public domain**.



One of the monkey selfies at issue

The following day, *Amateur Photographer* reported that Slater gave them further explanation as to how the photographs were created, downplaying the way newspaper articles had described them; Slater said reports that a monkey ran off with his camera and "began taking self-portraits" were incorrect and that the portrait was shot when his camera had been mounted on a tripod, with the primates playing around with a remote cable release as he fended off other monkeys.^[13]



The screenshot shows the "Amateur Photographer" website. The article title is "Ape-rture priority photographer plays down monkey reports" by Chris Cheesman, dated July 5, 2011. The article text states: "A photographer who says he witnessed monkeys taking pictures of themselves, tells **Amateur Photographer (AP)** that much of the media coverage has been exaggerated.

Wildlife photographer David Slater today played down newspaper reports that suggest a bunch of Indonesian monkeys grabbed his camera and began taking self-portraits.

And he revealed that the shots, featuring the endangered crested black macaque monkey, were taken three years ago.

Speaking to AP, David explained that his camera had been mounted on a tripod when the primates


On the right side of the page, there are sections for "Latest reviews" including "Leica M10-R review" and "Canon EOS R6 review: hands-on first look".

Generate text that is:

- coherent to the Wikipedia context
- contains latest event from the article


Primary Contribution

Primary Contribution



curated text (context)


Primary Contribution



+

curated text (context)

Primary Contribution



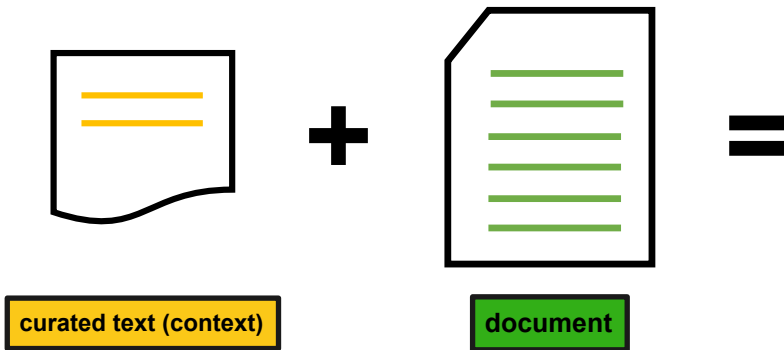
+




curated text (context)

document


Primary Contribution



Primary Contribution



Primary Contribution




- design a task to perform content transfer from an unstructured source of information
- release dataset

Applications


Applications

**Software
Documentation**



Applications

**Software
Documentation**




**Inbox
Summarization**



Applications

**Software
Documentation**



**Inbox
Summarization**




**Legal
Document**



Applications

Software
Documentation



Wikipedia
Articles



Inbox
Summarization



Legal
Document




Data Creation Process


Data Creation Process



Data Creation Process



Data Creation Process

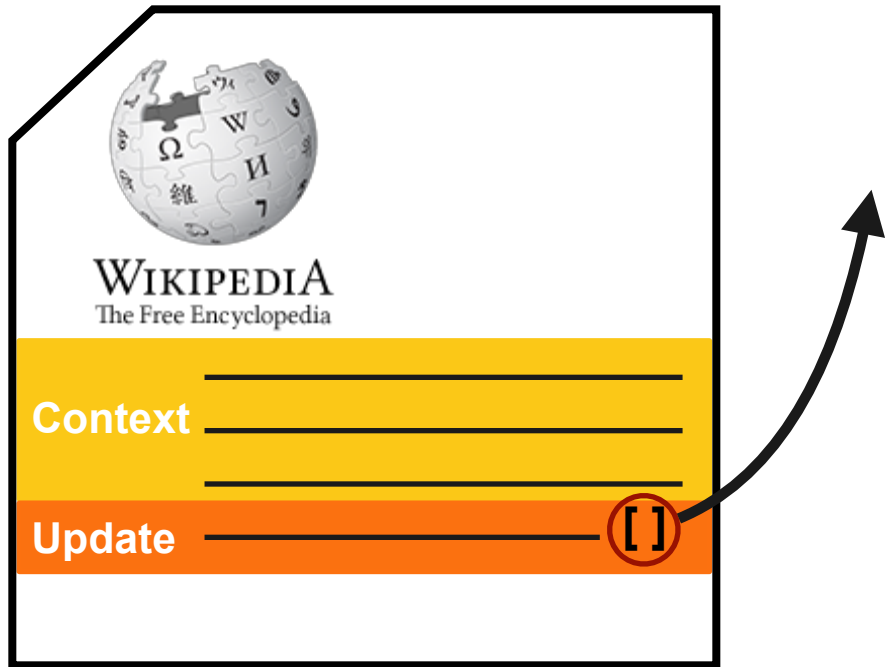


WIKIPEDIA
The Free Encyclopedia

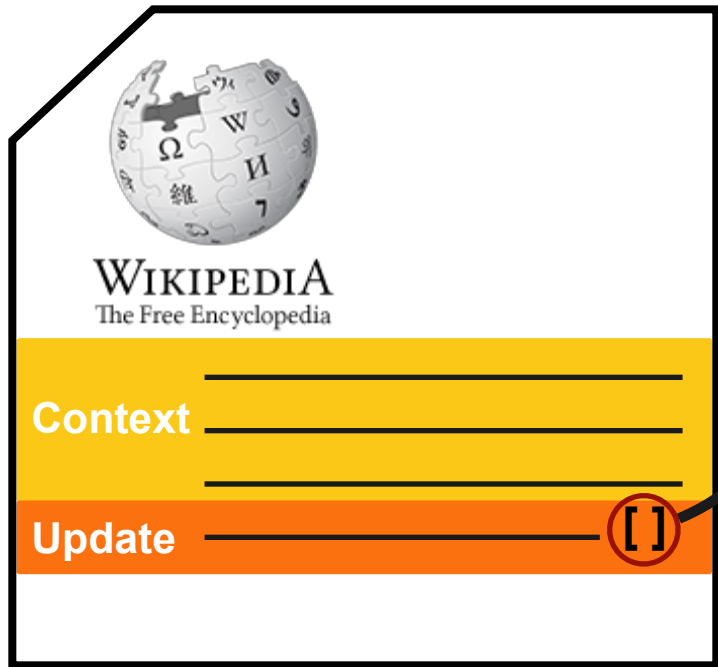
Context _____

Update _____ **[]**


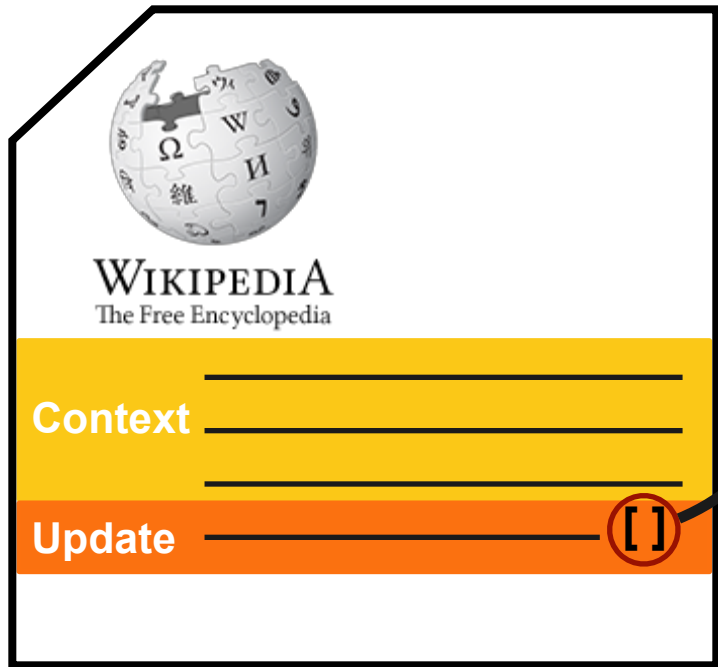
Data Creation Process



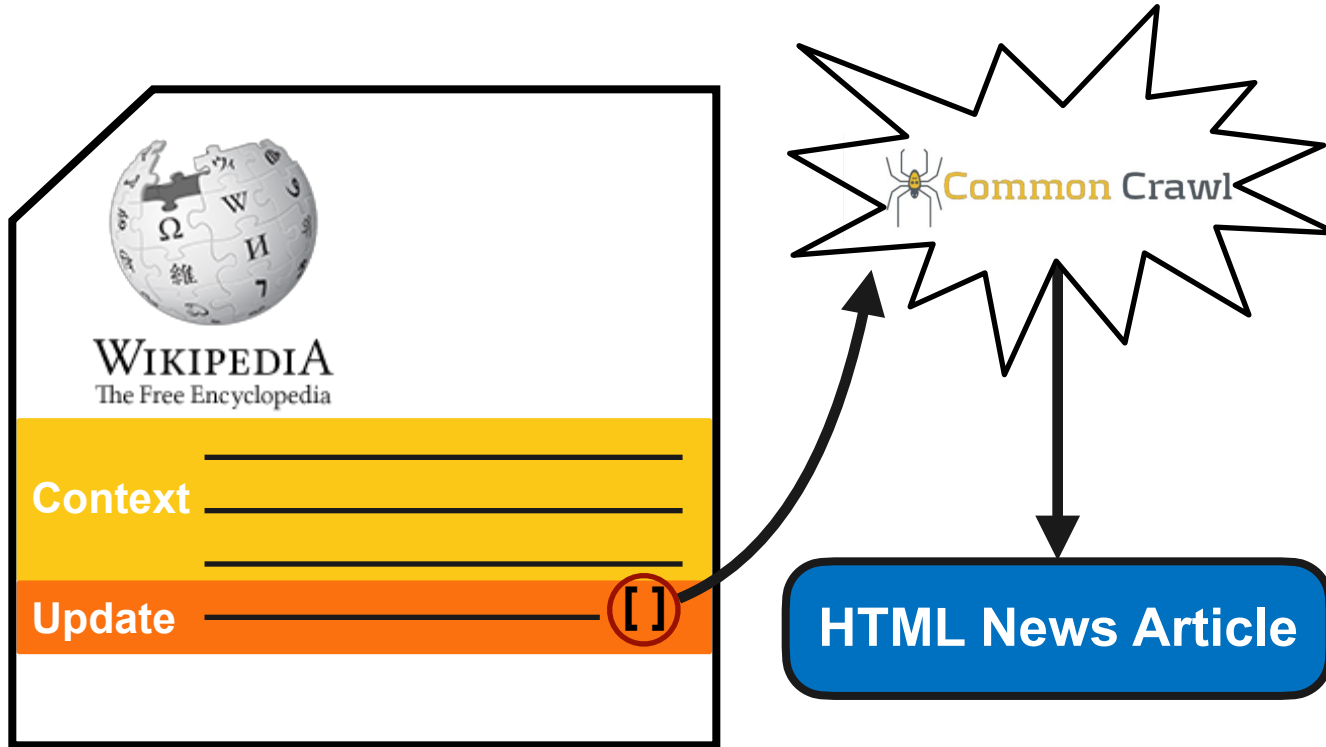
Data Creation Process



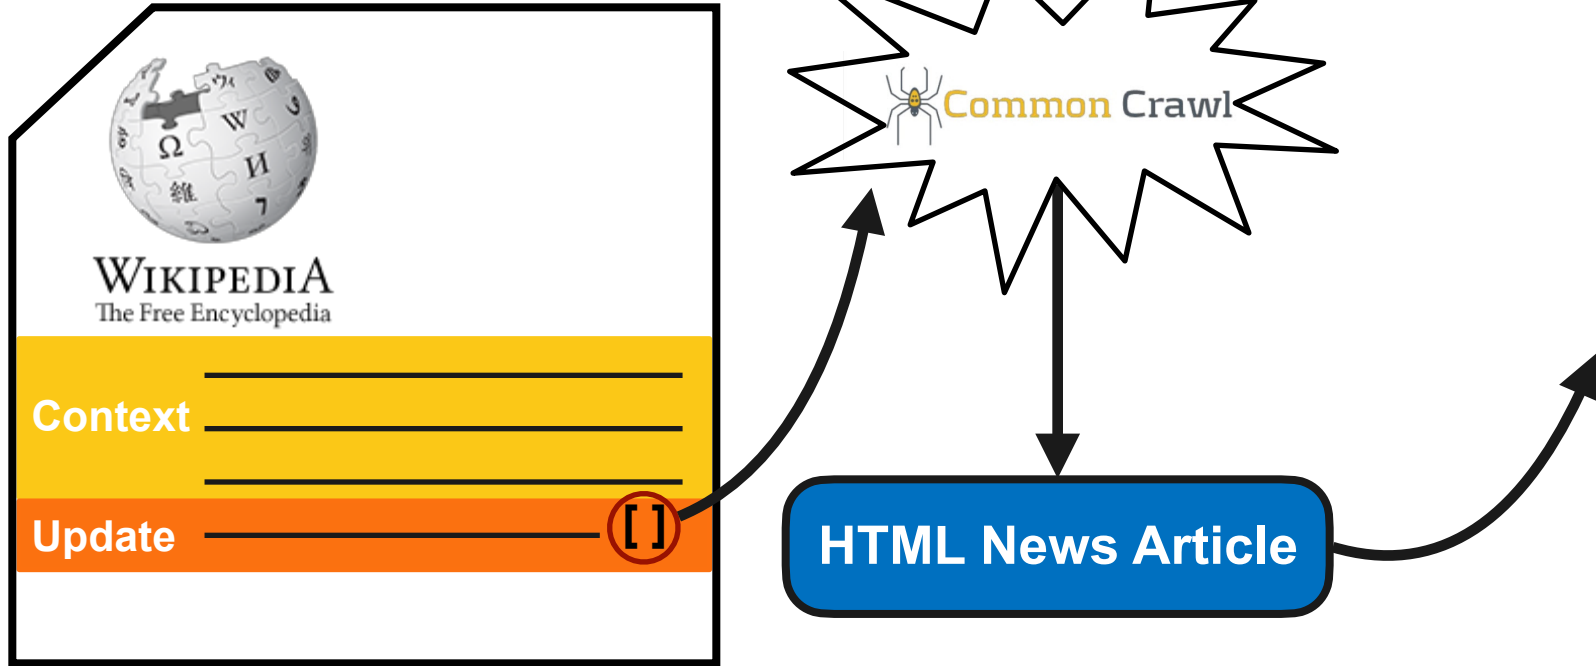
Data Creation Process




Data Creation Process




Data Creation Process



Data Creation Process



Data Creation Process



Data Creation Process

Total Data Size: 636K

Plain Text of
News Article

| News Article | Wikipedia Context | Update |
|--------------|-------------------|--------|
| ... | ... | ... |
| ... | ... | ... |

Models

Generative Models

- Context Agnostic Generative Model (CAG) — Baseline
- Context Informed Generative Model (CIG)
- Context Responsive Generative Model (CRG)

all models have global attention

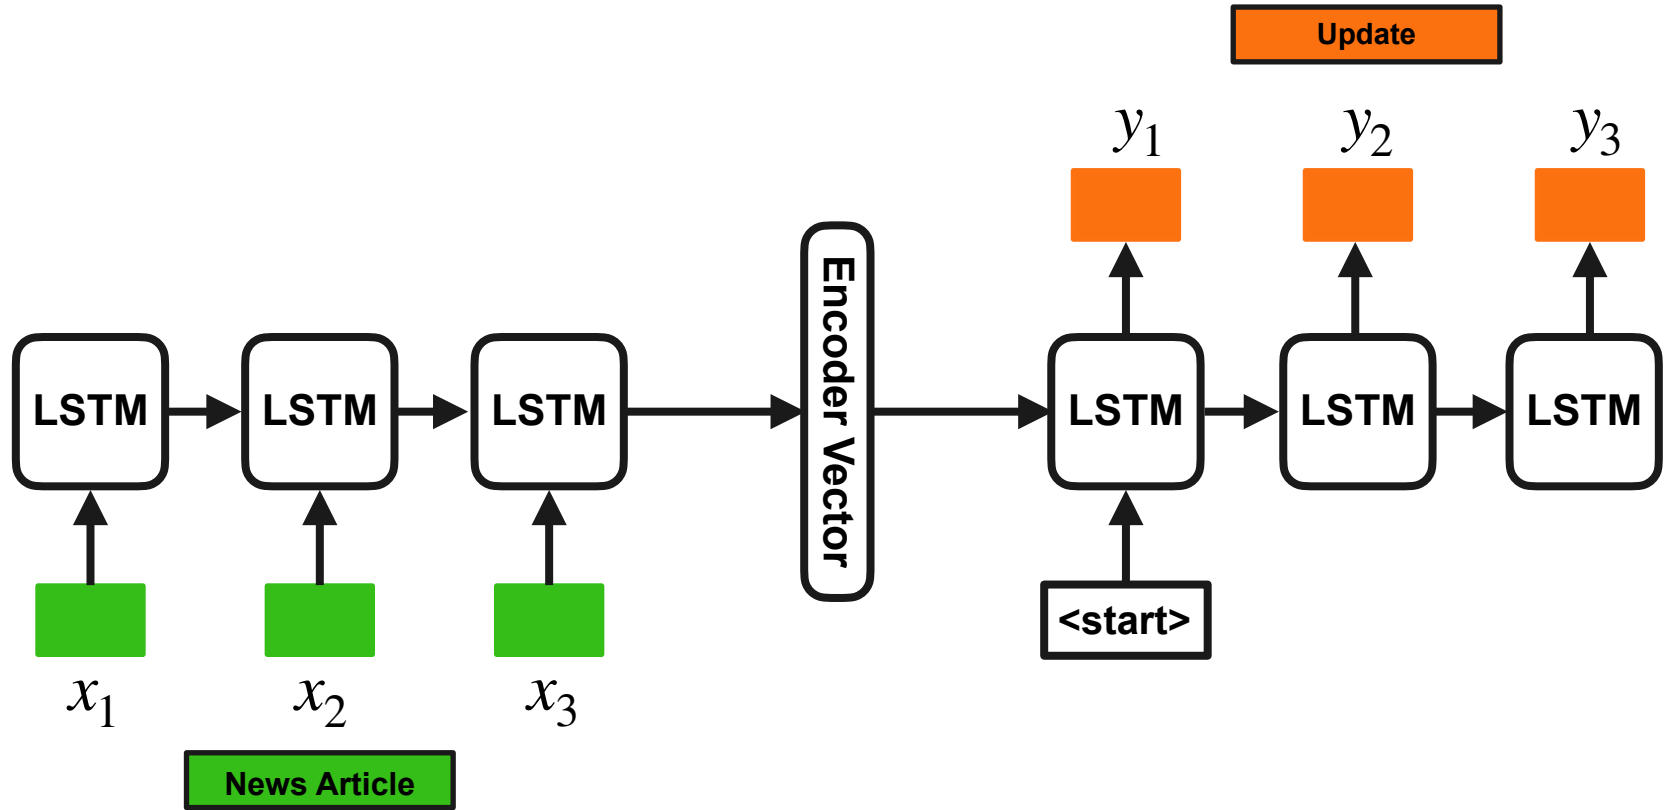
Extractive Models

- SumBasic
- Context Informed SumBasic
- Oracle

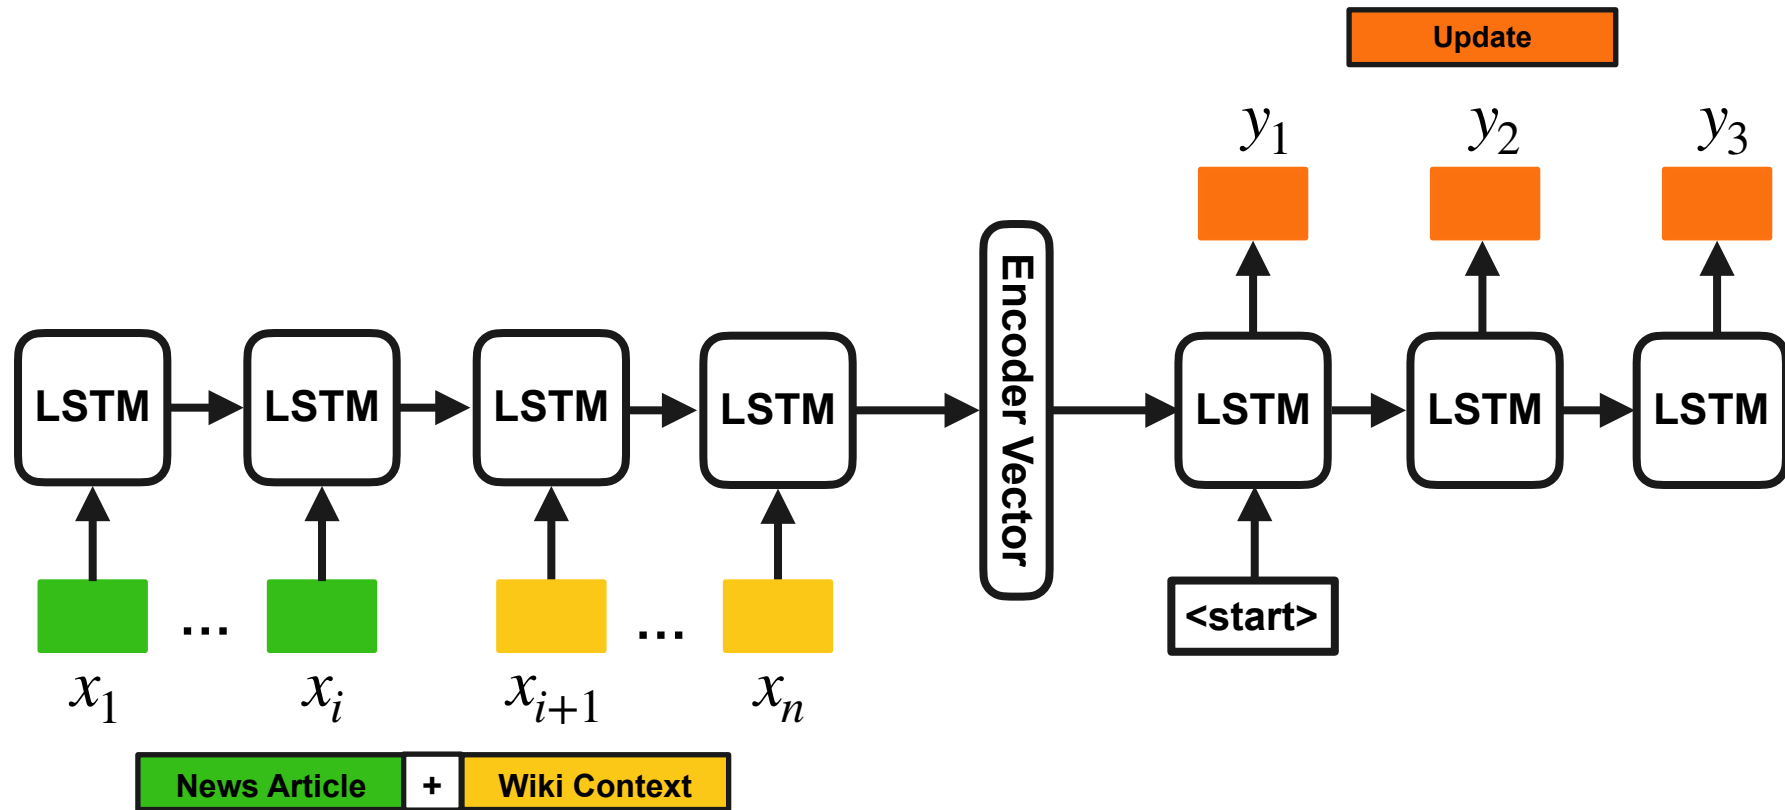
all models are simplistic to infer if

context helps in generation

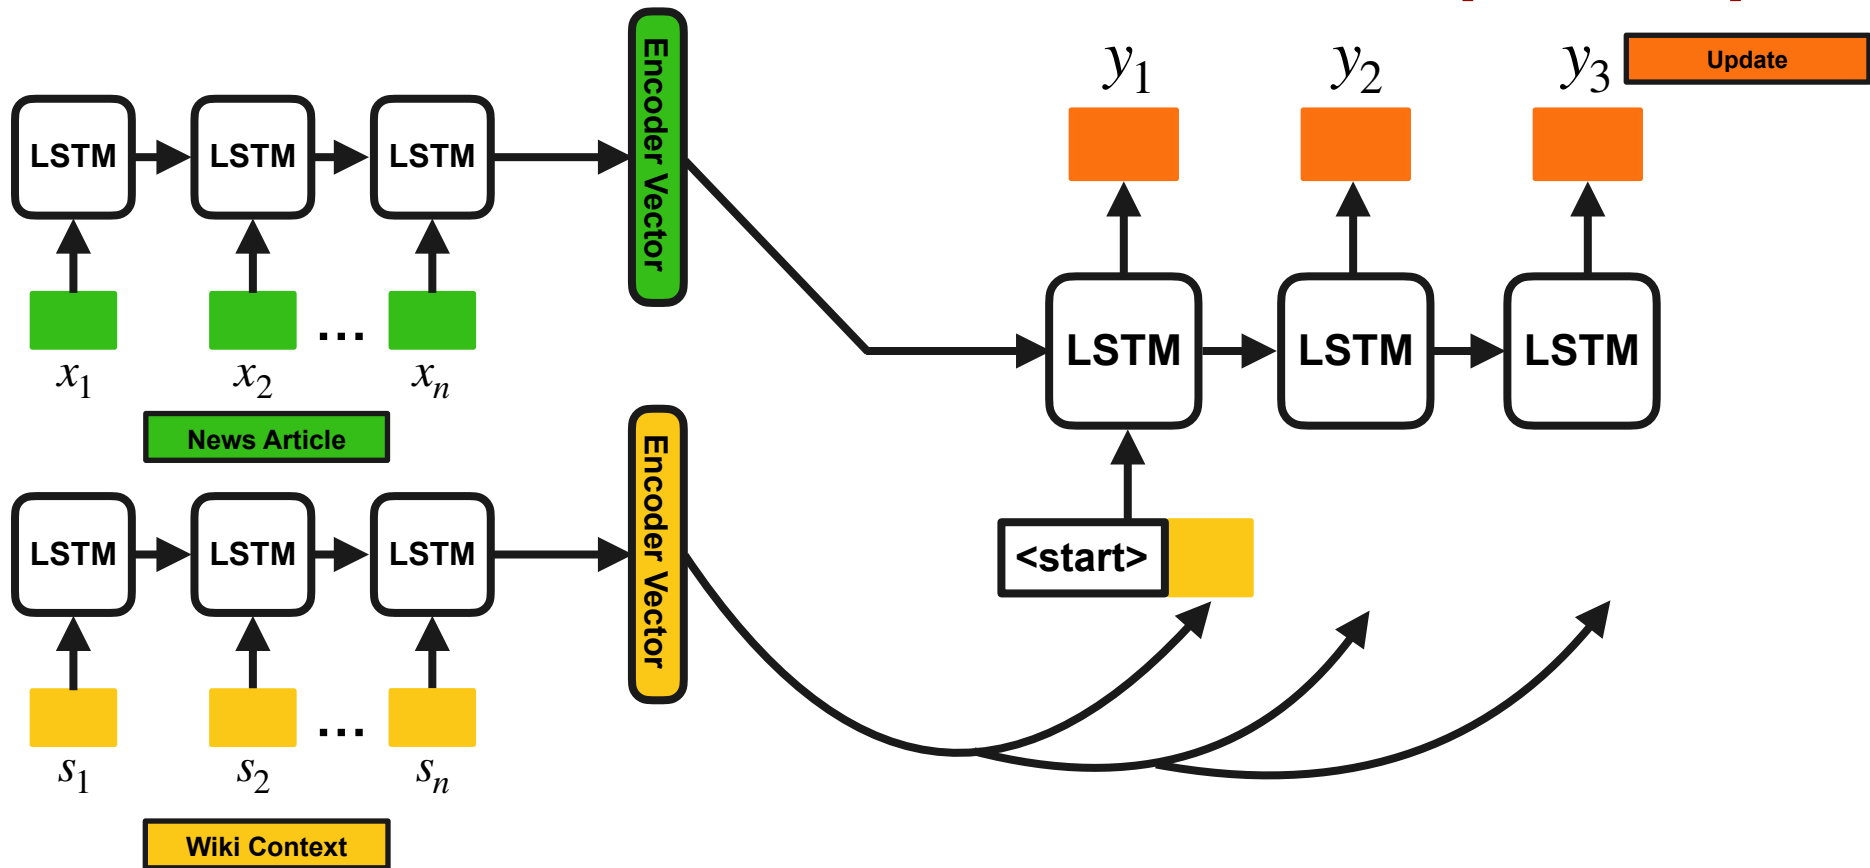
Context Agnostic Model (CAG) - Baseline



Context Informed Model (CIG)



Context Receptive Model (CRG)




Automated Evaluation

| Model | ROUGE-L |
|------------------------------------|-------------------------|
| SumBasic | 5.6 (5.6-5.7) |
| Context Informed SumBasic | 7.0 (7.0-7.1) |
| Context Agnostic Generative Model | 9.1 (9.0-9.2) |
| Context Informed Generative Model | 16.0 (15.9-16.1) |
| Context Receptive Generative Model | 14.7 (14.6-14.8) |
| Oracle | 28.8 (28.7-29.0) |


* METEOR and BLEU numbers are consistent with ROUGE-L

Relative Human Evaluation



Overview


Style



ACL '20
Storytelling '19
ACL '18



Content



NAACL '19
EMNLP '18




Structure



ACL '20
draft EMNLP '20

Ethical Considerations




WiNLP '19

The key to human-like generation is
Style, Content and Structure

Overview


Style



ACL '20
Storytelling '19
ACL '18



Content



NAACL '19
EMNLP '18




Structure



ACL '20
draft EMNLP '20

Ethical Considerations



WiNLP '19

The key to human-like generation is
Style, Content and Structure

Sentence Ordering Task

Sentence Ordering Task

Before

Ironman steals the Infinity Stones back from Thanos and uses them to disintegrate Thanos and his army, at the cost of his life.

Hulk travels to New York City in 2012 and convinces the Ancient One to give him the Time Stone.

Thor decapitates Thanos.

Five years later, AntMan escapes from the quantum realm.

Ironman builds a time machine to save the world.

Sentence Ordering Task

Before

Ironman steals the Infinity Stones back from Thanos and uses them to disintegrate Thanos and his army, at the cost of his life.

Hulk travels to New York City in 2012 and convinces the Ancient One to give him the Time Stone.

Thor decapitates Thanos.

Five years later, AntMan escapes from the quantum realm.

Ironman builds a time machine to save the world.

After

Thor decapitates Thanos.

Five years later, AntMan escapes from the quantum realm.

Ironman builds a time machine to save the world.

Hulk travels to New York City in 2012 and convinces the Ancient One to give him the Time Stone.

Ironman steals the Infinity Stones back from Thanos and uses them to disintegrate Thanos and his army, at the cost of his life.

Methodology

- Constraint Solving Problem vs Sequence Prediction Task
- For a document with n sentences ($\{s_1 \dots s_n\}$)
 - $|\mathcal{C}| = \binom{n}{2}$ constraints
 - Predicted constraints of the form $s_1 < s_2$
 - 4 sentences in a document then 6 constraints
 - $\{s_1 < s_2, s_1 < s_3, s_1 < s_4, s_2 < s_3, s_2 < s_4, s_3 < s_4\}$
 - Topological sort to find an order given \mathcal{C}
 - Graph: $s_1 \rightarrow s_2$ if $s_1 < s_2$

Constraint Learning


- ***BERT based Representation (B-TSort)***
 - Next Sentence Prediction
 - $\text{MLP}(\text{BERT}(s_1[\text{SEP}]s_2))$
- ***LSTM based Representation (L-TSort)***
 - $\mathbf{h}_1 = \text{LSTM}(s_1); \mathbf{h}_2 = \text{LSTM}(s_2)$
 - $\text{MLP}([\mathbf{h}_1; \mathbf{h}_2])$

Baselines

- ***Attention Order Network (AON)***
 - LSTM: sentence representation
 - Transformer: document representation
 - LSTM decoder: generate order
- ***BERT Attention Order Network (B-AON)***
 - BERT: sentence representation


Results for NeurIPS abstracts

AON L-TSort B-AON B-TSort




Results for NeurIPS abstracts

AON L-TSort B-AON B-TSort




Results for NeurIPS abstracts

AON L-TSort B-AON B-TSort



Results for NeurIPS abstracts

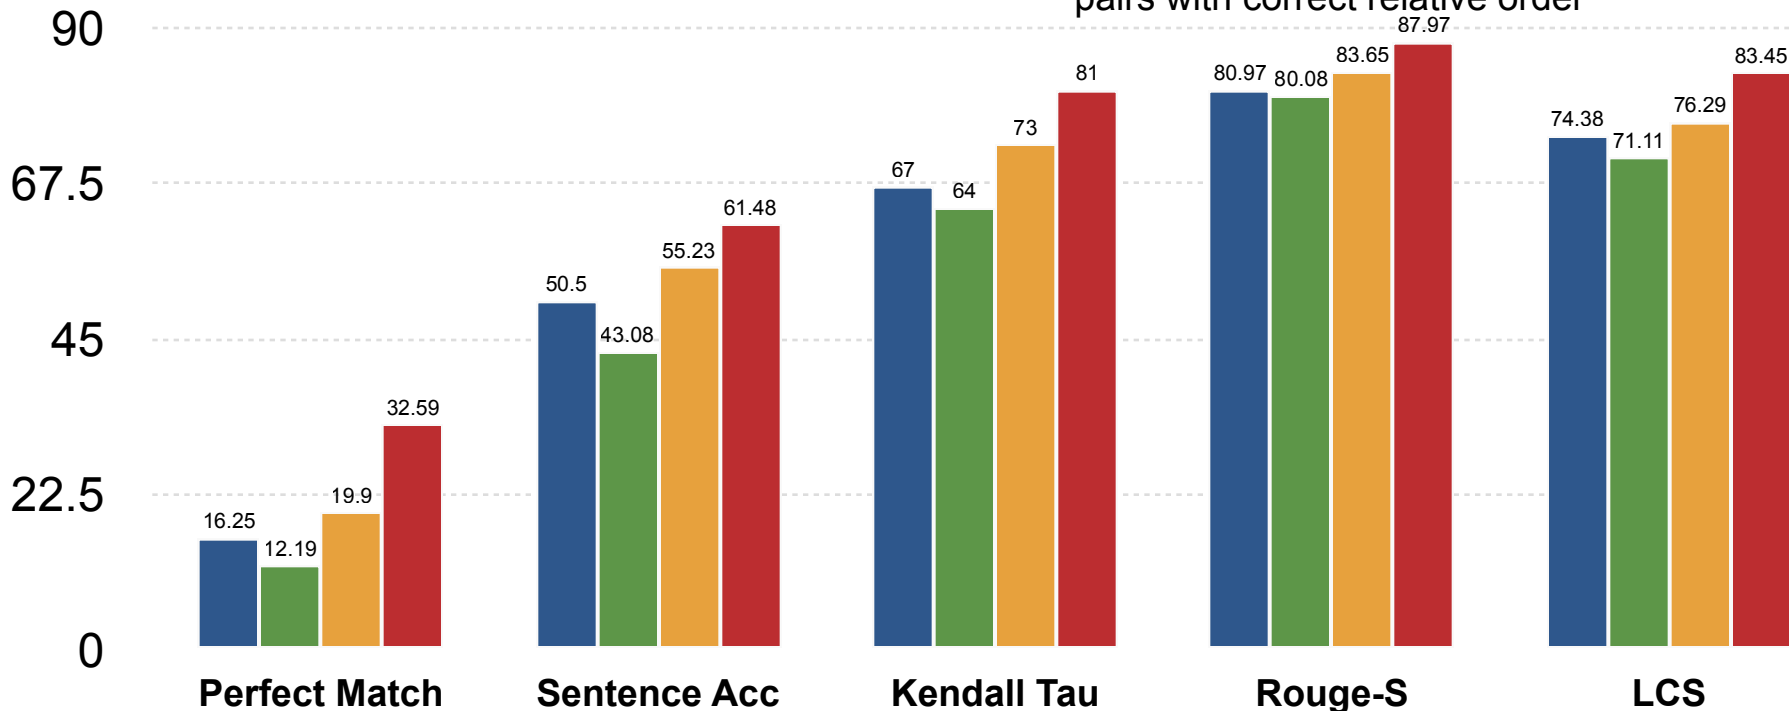
AON L-TSort B-AON B-TSort



Results for NeurIPS abstracts


AON L-TSort B-AON B-TSort

pairs with correct relative order




Results for NeurIPS abstracts

AON L-TSort B-AON B-TSort



Results for NeurIPS abstracts


AON L-TSort B-AON B-TSort



Results for NeurIPS abstracts

■ AON ■ L-TSort ■ B-AON ■ B-TSort


• L-TSort performs close to AON



Results for NeurIPS abstracts

AON L-TSort B-AON B-TSort


• B-AON performs better than AON



Results for NeurIPS abstracts


AON L-TSort B-AON B-TSort

• B-TSort performs the best



Results for NeurIPS abstracts

AON L-TSort B-AON B-TSort




Results

- B-TSort performs best in all metrics for SIND captions, NSF abstract, AAN abstract datasets
- Analysis of *displaced sentences*
 - distance between the predicted position and the actual position of a sentence.
 - B-TSort displaces less sentences.
- Analysis of documents with more than 10 sentences
 - B-TSort also performs better for longer documents
- Percentage of *mismatch* in input and output for AON
 - Order generation can have a mismatch between the input sentences and the generated orders.
 - This problem does not arise for our models.


Overview

Style



ACL '20
Storytelling '19
ACL '18

Content




NAACL '19
EMNLP '18

Structure



ACL '20
draft EMNLP '20

Ethical Considerations




WiNLP '19

The key to human-like generation is
Style, Content and Structure

Overview


Style



ACL '20
Storytelling '19
ACL '18



Content



NAACL '19
EMNLP '18




Structure



ACL '20
draft EMNLP '20



Ethical Considerations



WiNLP '19


The key to human-like generation is
Style, Content and Structure

Ethical Considerations

- Swear words, obscenity, bias, hate speech
- Broader Impact of controllable text generation
- Social good and bad applications
 - Generate persuasive tweets to spread awareness about climate change
 - Generate persuasive social media content to keep people away from vaccines


Overview

Style



ACL '20
Storytelling '19
ACL '18

Content




NAACL '19
EMNLP '18

Structure



ACL '20
draft EMNLP '20

Ethical Considerations



WiNLP '19

**The key to human-like generation is
Style, Content and Structure**

The key to human-like generation is Style, Content and Structure

Style: new model for style transfer and new styles like politeness, political slant, and gender.

Content: define new tasks to perform content grounded generation from unstructured data.

Structure: new framing of the sentence ordering task and the model is the new state-of-the art.

Thank You!

- Alan W Black (co-advisor)
- Ruslan Salakhutdinov (co-advisor)
- Yulia Tsvetkov (committee member)
- Jason Weson (committee member)
- Michel Galley (collaborator)
- Chris Quirk (collaborator)

Thank You!

- Aman Madan
- Amrith Setlur
- Tanmay Parekh
- Kangyan Zhou