Tutorial: Sunday, August 11th
09:00 - 12:30, at Lotus 5-7 (Level 22)

ACL 2024

Automatic & Human-AI Interactive Text Generation

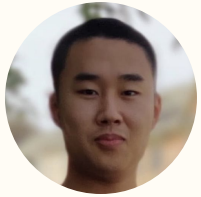

Yao Dou
(Georgia Tech)

Philippe Laban
(Salesforce)

Claire Gardent
(CNRS / Université de Lorraine)

Wei Xu
(Georgia Tech)

https://acl2024-text-generation-tutorial.github.io/

Text Editing Universe

Definition: A task involving the modification of a text to enhance its readability, accuracy, or suitability for a specific purpose. *Huh?*

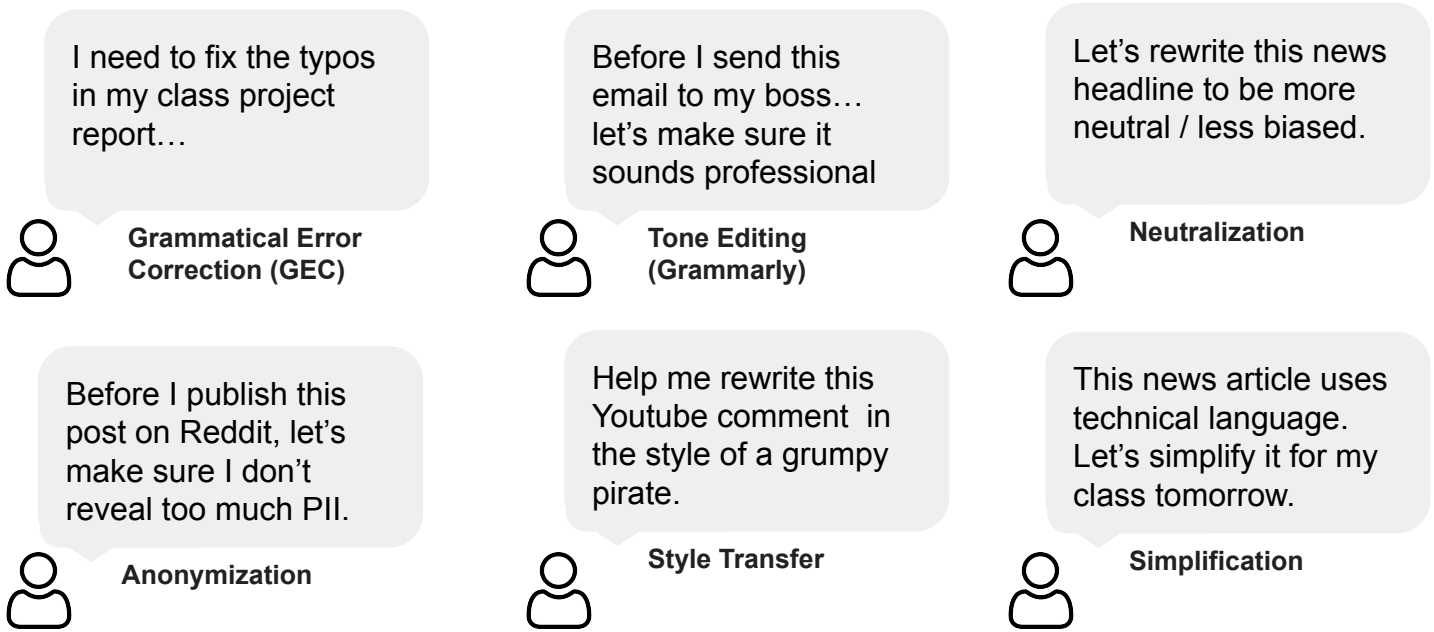

Text Editing Universe

Definition: A task involving the modification of a text to enhance its readability, accuracy, or suitability for a specific purpose. *Huh?*

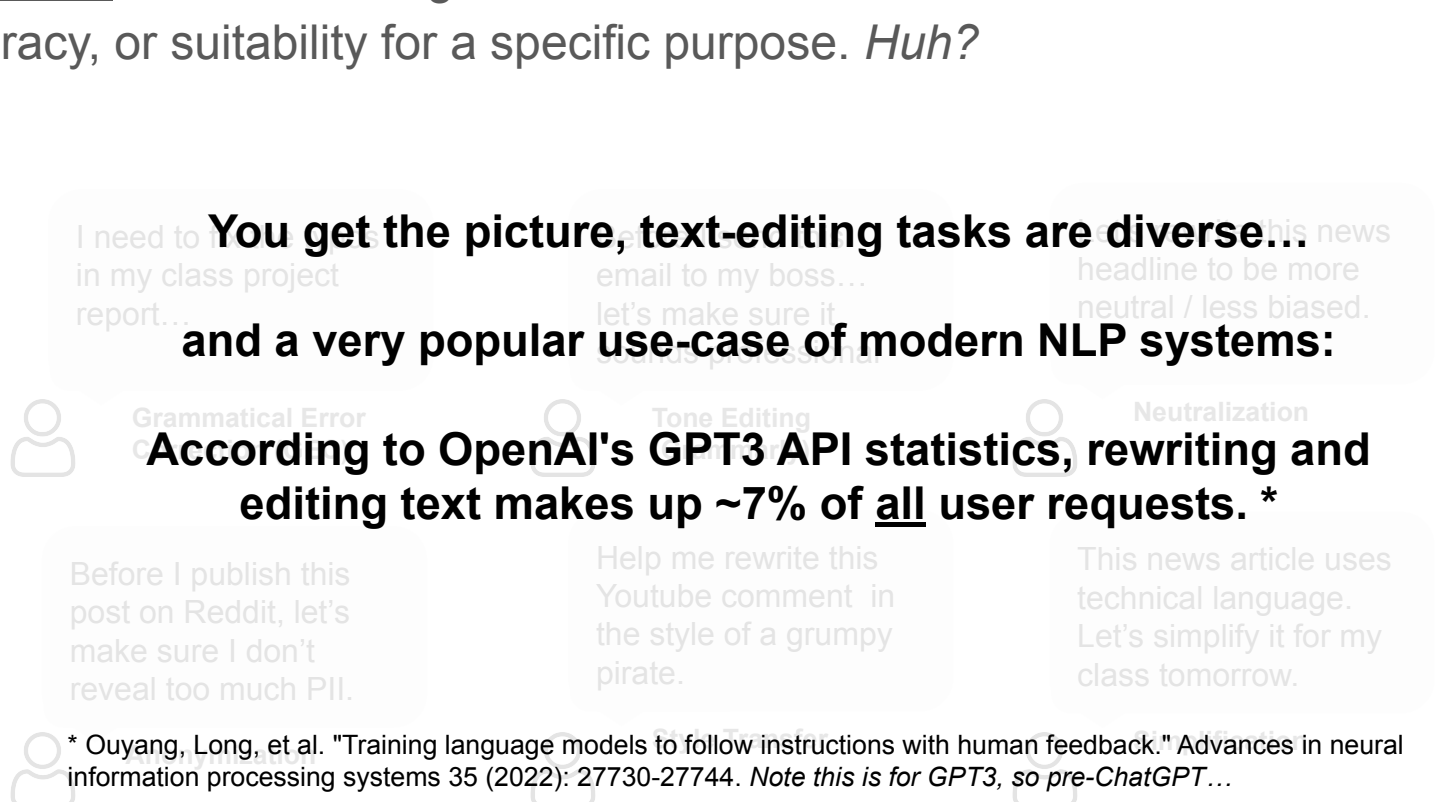

You get the picture, text-editing tasks are diverse…

and a very popular use-case of modern NLP systems:

According to OpenAI's GPT3 API statistics, rewriting and editing text makes up ~7% of all user requests. *

* Ouyang, Long, et al. "Training language models to follow instructions with human feedback." Advances in neural information processing systems 35 (2022): 27730-27744. *Note this is for GPT3, so pre-ChatGPT…*

TC

Login

Search

Startups

Venture

Apple

Security

AI

Apps

Events

Startup Battlefield

More

AI

Automattic launches AI writing tool that aims to make WordPress blogs more readable and succinct

Paul Sawers / 2:32 AM PDT • August 7, 2024

Comment

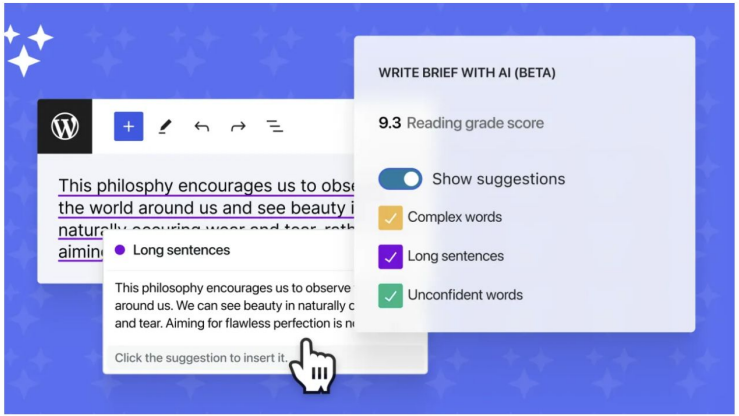

Image Credits: Automattic

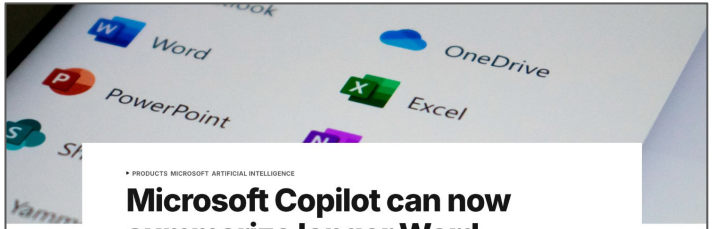

▸ PRODUCTS MICROSOFT ARTIFICIAL INTELLIGENCE

Microsoft Copilot can now summarize longer Word documents

You can now review the main idea of an 80,000 words document without having to read the entire document.

by Oyinebiladou Omemu Updated July 23, 2024

Microsoft is rolling out a new update toMicrosoft Word's AI assistant, Copilot, that will let you summarize documents up to a staggering 80,000 words – about four times its previous limit. That's like getting a sneak peek into entire novels or complex textbooks in a few minutes.

This update lets you quickly summarize lengthy documents, articles, or reports, saving you time. The information will be condensed while retaining key points, making it easier to get the essential content.

For those unfamiliar, using this feature is simple. Open your document, hit the Copilot button on the Home tab, and type "Summarize this doc" into the chat box. Within moments, you have a neat summary ready to go.

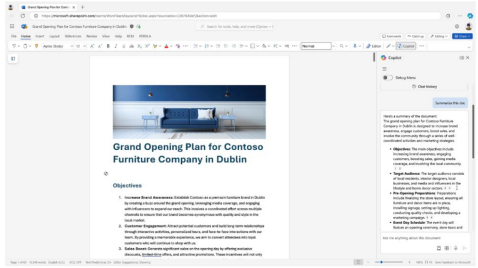

This feature is rolling out across all platforms; Windows, Mac, iOS, and Android. Specifically, you'll need Word version 2310 (Build 16919.20000) or later on Windows, version 16.78 (Build 23082801) or later on Mac, Build 16.0.16827100000

AI

Automattic launches AI writing tool that aims to make WordPress blogs more readable and succinct

Paul Sawers / 2:32 AM PDT • August 7, 2024

Comment

Login

Search

Startups

Venture

Apple

Security

AI

Apps

Events

Startup Battlefield

More

Image Credits:

Artificial Intelligence › | Replacing Meaningless Jobs | Rapid Weather Forecasts | A.I.'s Math Problem | Microsoft's Risk-Taker | Qui

Google Tests A.I. Tool That Is Able to Write News Articles

The product, pitched as a helpmate for journalists, has been demonstrated for executives at The New York Times, The Washington Post and News Corp, which owns The Wall Street Journal.

Share full article | 181

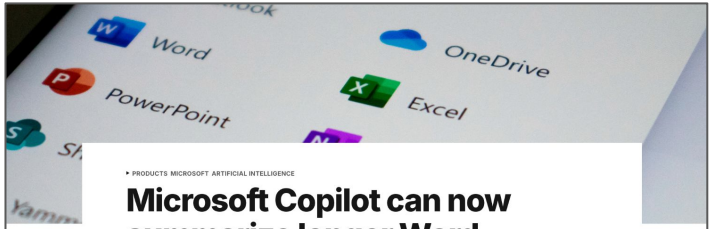

PRODUCTS MICROSOFT ARTIFICIAL INTELLIGENCE

Microsoft Copilot can now summarize longer Word documents

You can now review the main idea of an 80,000 words document without having to read the entire document.

by Oyinebiladou Omemu Updated July 23, 2024

Microsoft is rolling out a new update toMicrosoft Word's AI assistant, Copilot, that will let you summarize documents up to a staggering 80,000 words – about four times its previous limit. That's like getting a sneak peek into entire novels or complex textbooks in a few minutes.

This update lets you quickly summarize lengthy documents, articles, or reports, saving you time. The information will be condensed while retaining key points, making it easier to get the essential content.

For those unfamiliar, using this feature is simple. Open your document, hit the Copilot button on the Home tab, and type "Summarize this doc" into the chat box. Within moments, you have a neat summary ready to go.

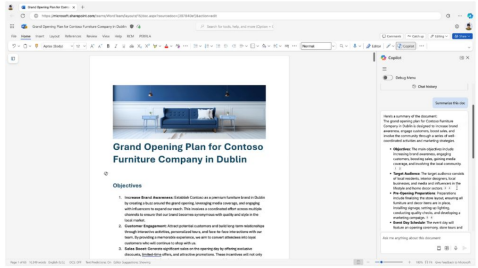

This feature is rolling out across all platforms; Windows, Mac, iOS, and Android. Specifically, you'll need Word version 2310 (Build 16919.20000) or later on Windows, version 16.78 (Build 23082801) or later on Mac, Build 16.0.16827100000

TC

Login

Search

Startups

Venture

Apple

Security

AI

Apps

Events

Startup Battlefield

More

AI

Automattic launches AI writing tool that aims to make WordPress blogs more readable and succinct

Paul Sawers / 2:32 AM PDT • August 7, 2024

Comment

Image Credits:

Artificial Intelligence › | Replacing Meaningless Jobs | Rapid Weather Forecasts | A.I.'s Math Problem | Microsoft's Risk-Taker | Qui

Google Tests A.I. Tool That Is Able to Write News Articles

The product, pitched as a helpmate for journalists, has been demonstrated for executives at The New York Times, The

engadget

Sections

Proton Mail now has a privacy-focused AI writing assistant

Proton says Scribe can't train itself on your inbox data.

Kris Holt
Contributing Reporter
Thu, Jul 18, 2024 · 2 min read

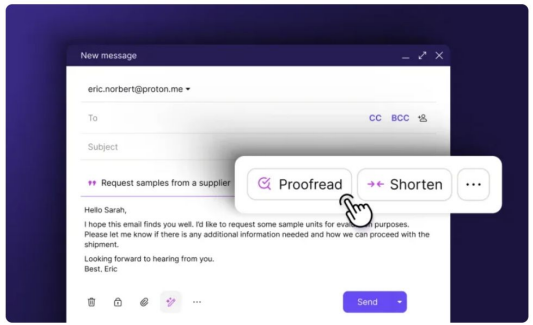

Proton Mail

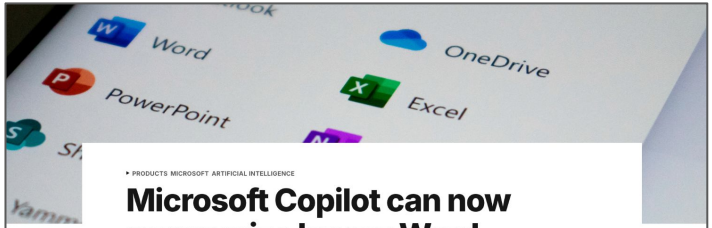

PRODUCTS MICROSOFT ARTIFICIAL INTELLIGENCE

Microsoft Copilot can now summarize longer Word documents

You can now review the main idea of an 80,000 words document without having to read the entire document.

by Oyinebiladou Omemu Updated July 23, 2024

Microsoft is rolling out a new update to Microsoft Word's AI assistant, Copilot, that will let you summarize documents up to a staggering 80,000 words – about four times its previous limit. That's like getting a sneak peek into entire novels or complex textbooks in a few minutes.

This update lets you quickly summarize lengthy documents, articles, or reports, saving you time. The information will be condensed while retaining key points, making it easier to get the essential content.

For those unfamiliar, using this feature is simple. Open your document, hit the Copilot button on the Home tab, and type "Summarize this doc" into the chat box. Within moments, you have a neat summary ready to go.

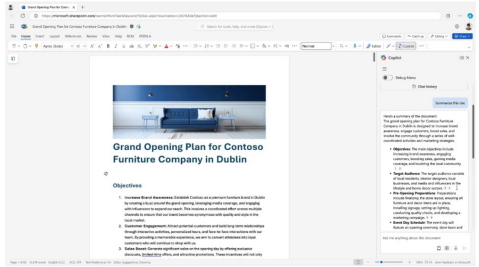

This feature is rolling out across all platforms; Windows, Mac, iOS, and Android. Specifically, you'll need Word version 2310 (Build 16919.20000) or later on Windows, version 16.78 (Build 23082801) or later on Mac, Build 16.0.16827100000

Today's Tutorial Schedule

9:00-9:15	Introduction (Yao, Claire, Wei, Philippe)
9:15-10:00	Evaluation of LLM-generated Text (Yao)
10:00-10:45	Models: Text Simplification and Text Rewriting (Claire)
10:45-11:00	☕ Coffee Break ☕
11:00-11:45	More Models: Decoding, Distillation, and Diffusion (Wei)
11:45-12:30	From method -> usable system: HCI & NLP (Philippe)

Q&A - please feel free to ask questions during the tutorial anytime

Yao Dou

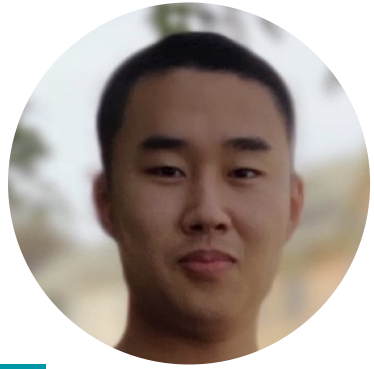

PhD student at Georgia Tech

Research: Evaluation, LLM Privacy

Today's tutorial

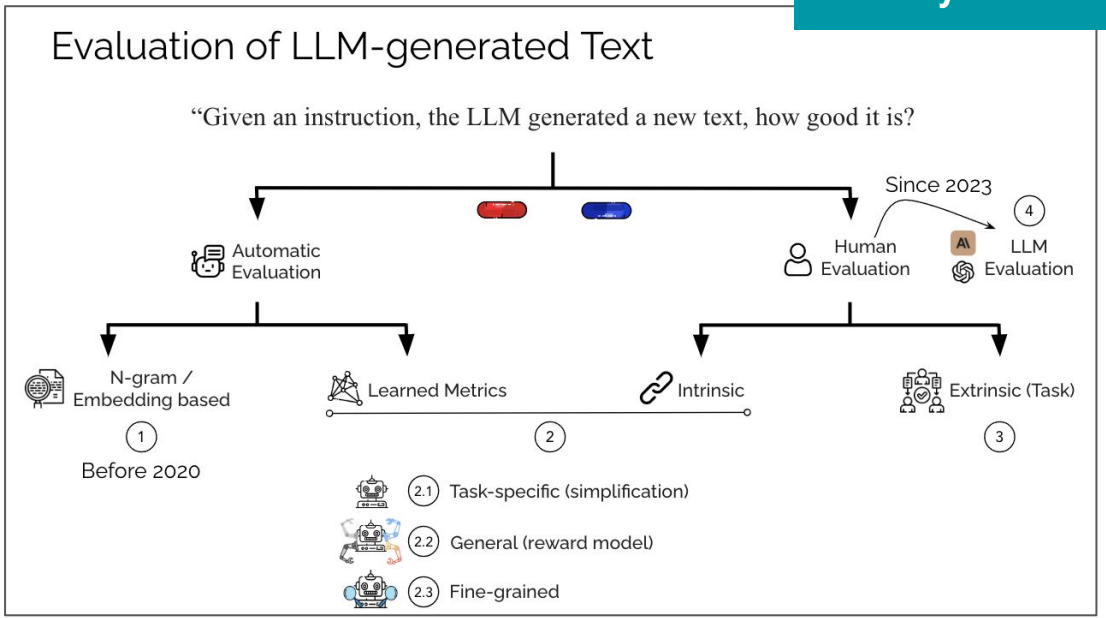

Claire Gardent

Research Scientist at the French National Center for Scientific Research (CNRS)

Research: Generating Text from Multiple Sources into Multiple Languages; Human-Machine Dialog

Today's tutorial

Text Rewriting Models

1. **Text Simplification**
   Simple MLM, Planning,
   Using Summarisation data / models
2. **Generic Text Rewriting**
   Mimicking the human writing process,
   Using Synthetic Data and RL

Wei Xu

Associate Professor in CS

Director of **NLP X Lab** @ Georgia Tech

Research: Multilingual/Multicultural LLMs, Generative AI, NLP+X interdisciplinary research

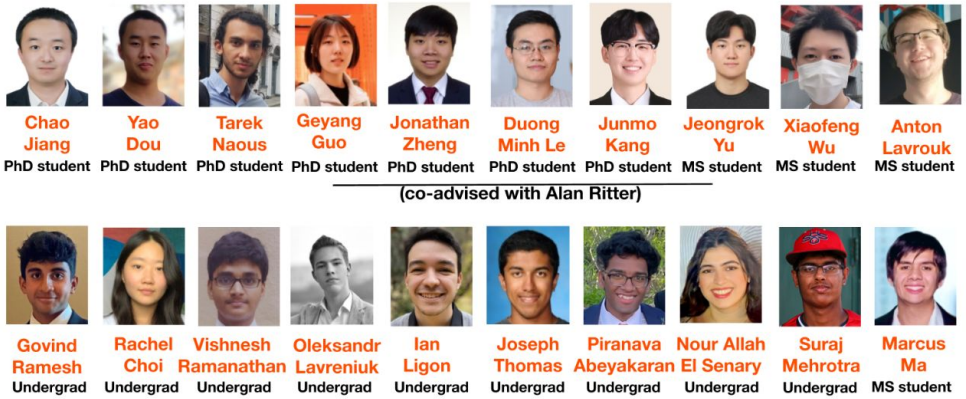

Today's tutorial

Three Popular Methods for Generation

1. **Decoding**:
an inference-time solution to optimize LLM outputs
(Survey by Welleck+ 2024 & Bertsch+, 2023; MBR with Multi-Prompt by Heineman+, 2024)

2. **Diffusion**:
an alternative to Transformer-based LLM
(Diffusion-LM by Li+ 2022; DiffuSeq by Gong+, 2022; SeqDiffuSeq by Yuan+, 2024)

3. **Distillation**:
reproduce GPT-4 performance by small open-source LLMs
(Edit-based generation by Dou+ 2024; Feedback to refine LLM outputs by Wadhwa+ 2024)

Georgia Tech College of Computing

Philippe Laban

Research Scientist @ Salesforce Research
Previously: UC Berkeley (Ph.D.), Georgia Tech (undergrad)
Research: Summarization, Simplification, NLP+HCI, Factuality

Today's tutorial

An Intro to Human-Centered NLP

PART 1
Beyond the ML Method
Situation: I have built a cool NLP model and I want to people to use it!

What should I consider when building the system that uses my NLP method?

PART 2
Usability Testing Crash Course
Situation: I have built a system that uses my NLP model. How do if "it works" for users?

A deep-dive into two usability studies of an LLM-based text-editing system.