Writing with AI

How technology can aid
humans to be better writers

이민아 Stanford University

Bio

고려대학교 학부 졸업
Programming Languages (Hakjoo Oh, Korea University)

스탠퍼드 박사과정 재학중
Natural Language Processing (Percy Liang, Stanford University)
Human-Computer Interaction (Qian Yang, Cornell University)
Economics (Erik Brynjolfsson, Stanford University)

Bio

고려대학교 학부 졸업
Programming Languages (Hakjoo Oh, Korea University)

스탠퍼드 박사과정 재학중
Natural Language Processing (Percy Liang, Stanford University)
Human-Computer Interaction (Qian Yang, Cornell University)
Economics (Erik Brynjolfsson, Stanford University)

how can we better design human-AI interaction?

Bio

고려대학교 학부 졸업
Programming Languages (Hakjoo Oh, Korea University)

스탠퍼드 박사과정 재학중
Natural Language Processing (Percy Liang, Stanford University)
Human-Computer Interaction (Qian Yang, Cornell University)
Economics (Erik Brynjolfsson, Stanford University)

what is the economic impact of AI?

how can we better design human-AI interaction?

Part 1

What is AI-assisted writing?

AI와 함께 글쓰기

Introduction

Part 2

How can we use AI to enhance productivity and creativity in writing?

AI로 생산성, 창의성 증가시키기

Productivity & Creativity

Part 3

What opportunities and risks might arise from increased adoption of AI-assisted writing?

AI 사용의 기회와 문제점

Opportunities & Risks

AI와 함께 글쓰기
Part 1

AI-assisted writing

AI = automation?

인공지능 = 자동화?

AI-assisted writing

Human-centered AI

사람을 위한 AI

Stanford Institute for Human-Centered Artificial Intelligence (https://hai.stanford.edu/)

AI-assisted writing

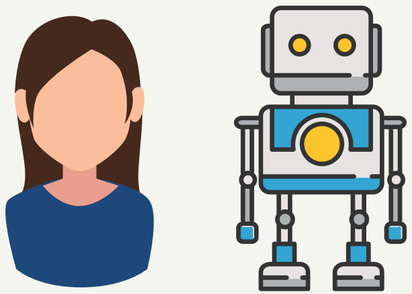

Human-AI **Interaction**

사람과 AI간의 **상호 작용**

AI-assisted writing

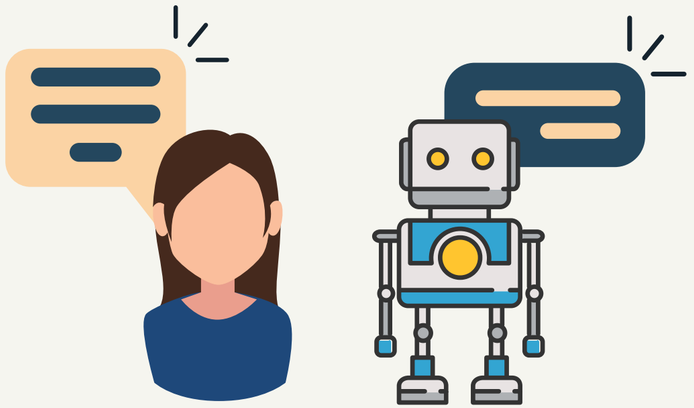

Human-**LM** Interaction

사람과 **언어모델**간의 상호작용

(Language Model)

Language model

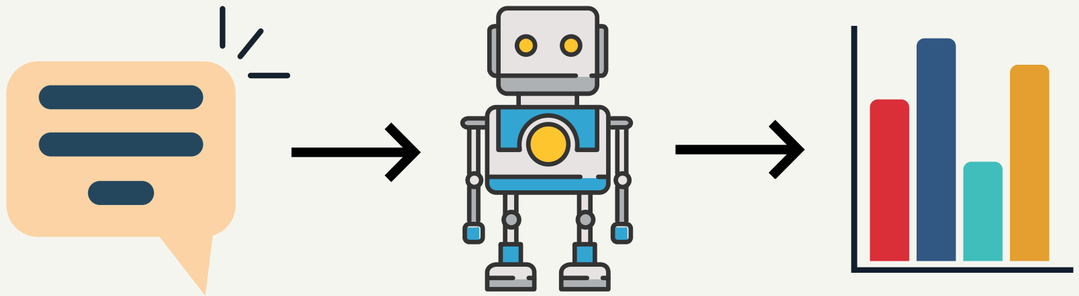

We use language models every day!

Stanford CS224N http://web.stanford.edu/class/cs224n/slides/cs224n-2021-lecture05-rnnlm.pdf

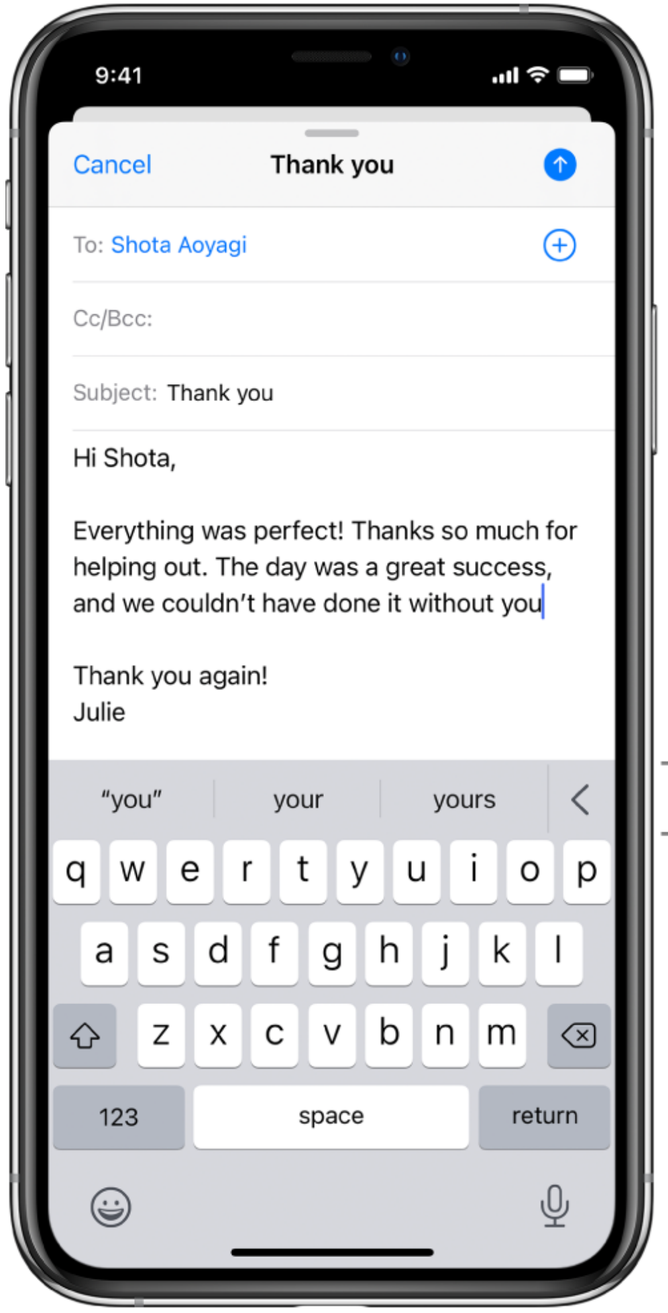

Predictive text

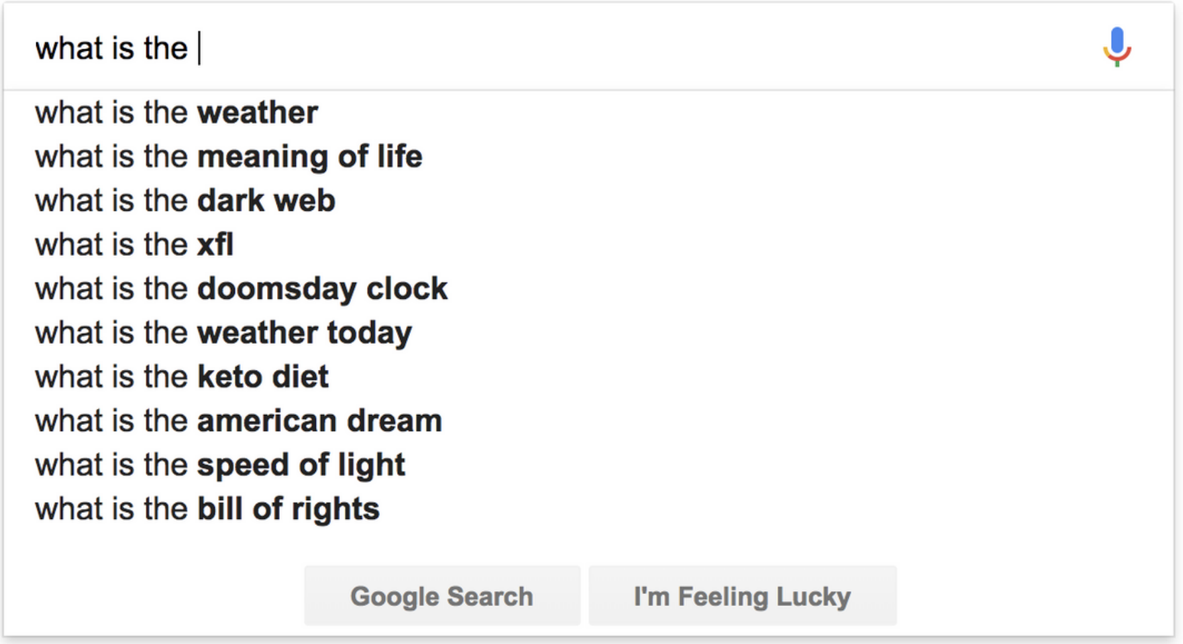

언어모델을 글쓰기에 사용하는 다양한 방법들

01
Autocomplete
자동 완성

02
Fill in the blanks
빈칸 채우기

03
Contextual
thesaurus
어휘력 늘리기

04
Human-LM
collaborative writing
사람과 AI가 함께
글을 쓰는 과정

05
Writing in the
economy
AI가 경제에
미치는 영향

AI로 생산성 증가시키기

Part 2-1

생산성

일

시간

AI로 생산성 증가시키기

자동 완성

같은 일

적은 시간 ↓

빈칸 채우기

더 많은 일 ↑

같은 시간

어휘력 늘리기

양질의 일 ↑

같은 시간

Mina Lee, Tatsunori Hashimoto, Percy Liang. Learning autocomplete systems as a communication games. NeurIPS 2019 Workshop.
Chris Donahue, **Mina Lee**, Percy Liang. Enabling Language Models to Fill in the Blanks. ACL 2020.
Mina Lee, Chris Donahue, Robin Jia, Alexander Iybor, Percy Liang. Swords ⚔: A Benchmark for Lexical Substitution with Improved Data Coverage and Quality. NAACL 2021.

AI로 생산성 증가시키기

자동 완성

같은 일

적은 시간 ↓

빈칸 채우기

더 많은 일 ↑

같은 시간

어휘력 늘리기

양질의 일 ↑

같은 시간

Mina Lee, Tatsunori Hashimoto, Percy Liang. Learning autocomplete systems as a communication games. NeurIPS 2019 Workshop.
Chris Donahue, **Mina Lee**, Percy Liang. Enabling Language Models to Fill in the Blanks. ACL 2020.
Mina Lee, Chris Donahue, Robin Jia, Alexander Iybor, Percy Liang. Swords ⚔: A Benchmark for Lexical Substitution with Improved Data Coverage and Quality. NAACL 2021.

키워드로부터 자동 완성

When you know what you want to write how can AI help you write faster?

Mina Lee
Stanford University

Tatsunori Hashimoto
Stanford University

Percy Liang
Stanford University

Mina Lee, Tatsunori Hashimoto, Percy Liang. Learning autocomplete systems as a communication games. NeurIPS 2019 Workshop.

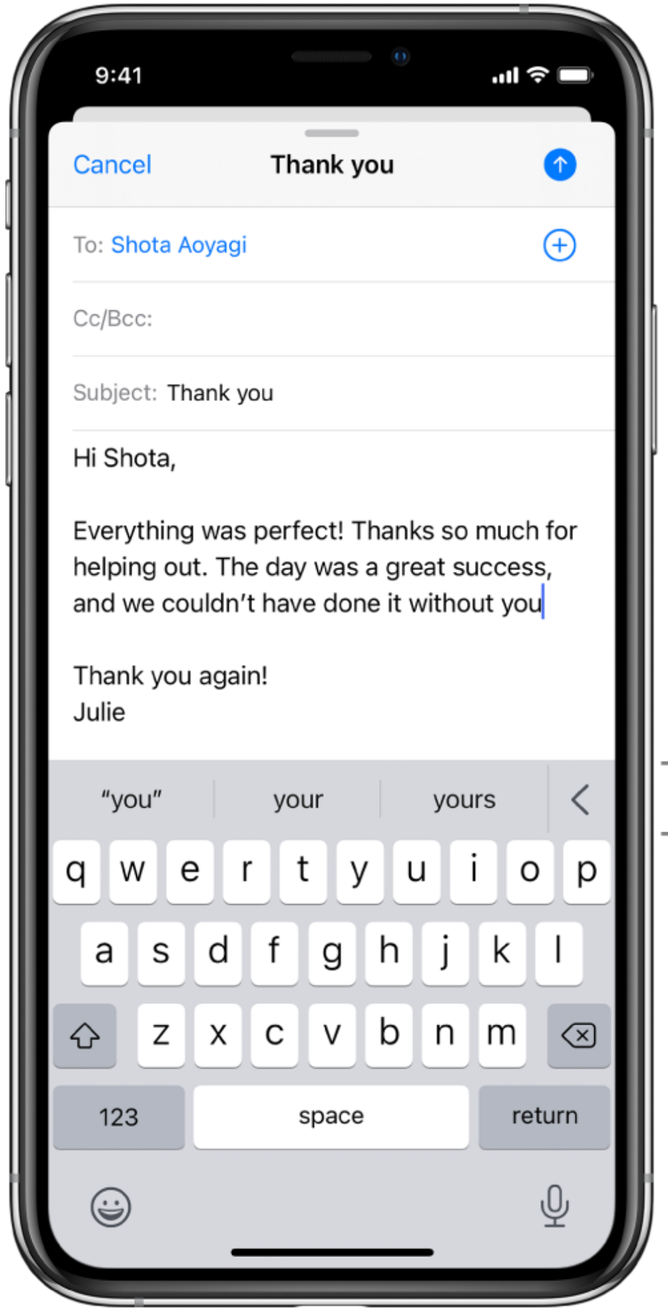

Predictive text

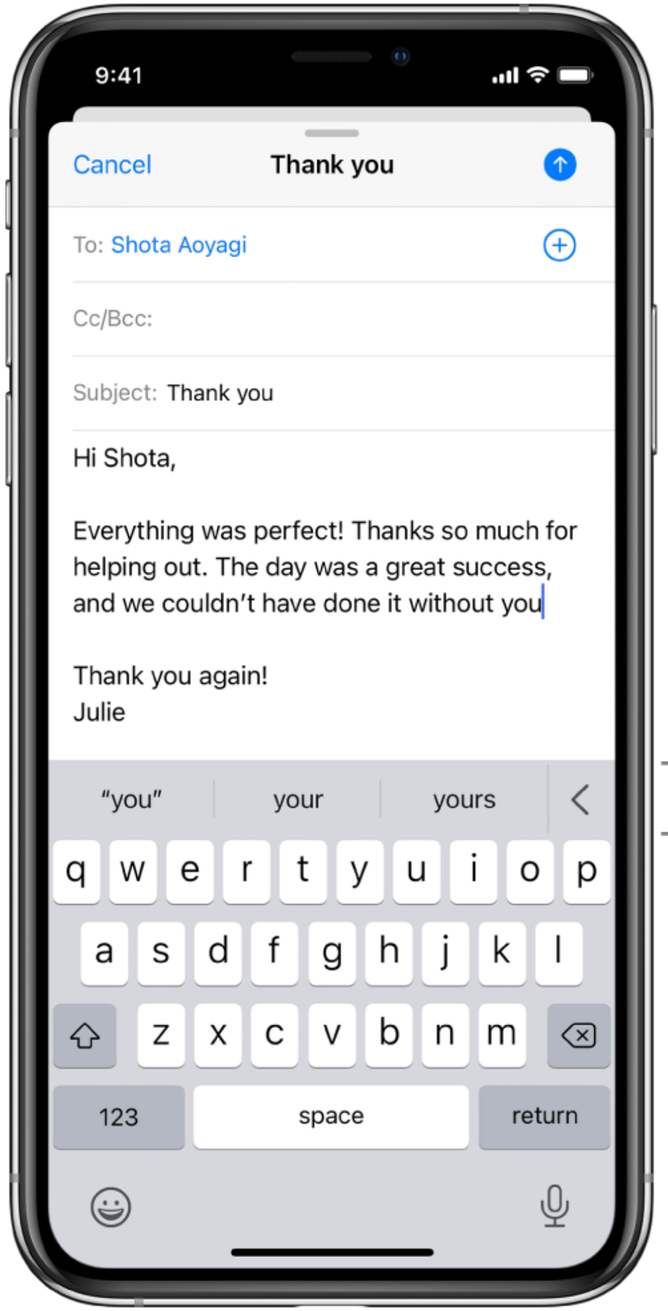

Predictive text

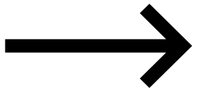

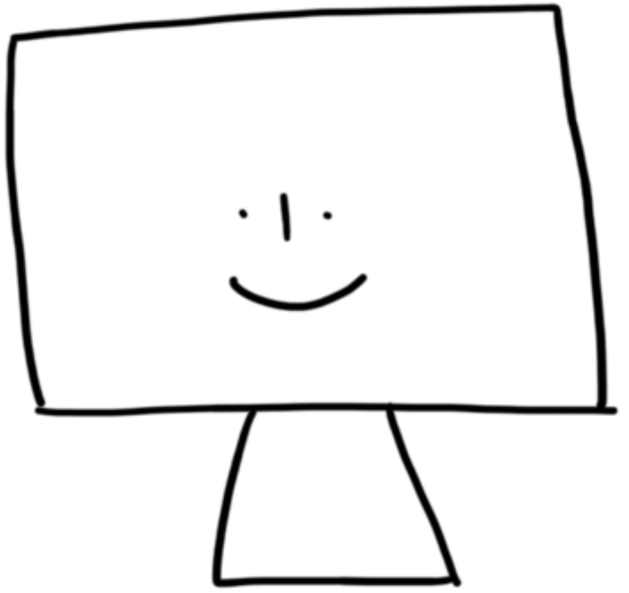

System

Why Keyword-based Autocomplete?

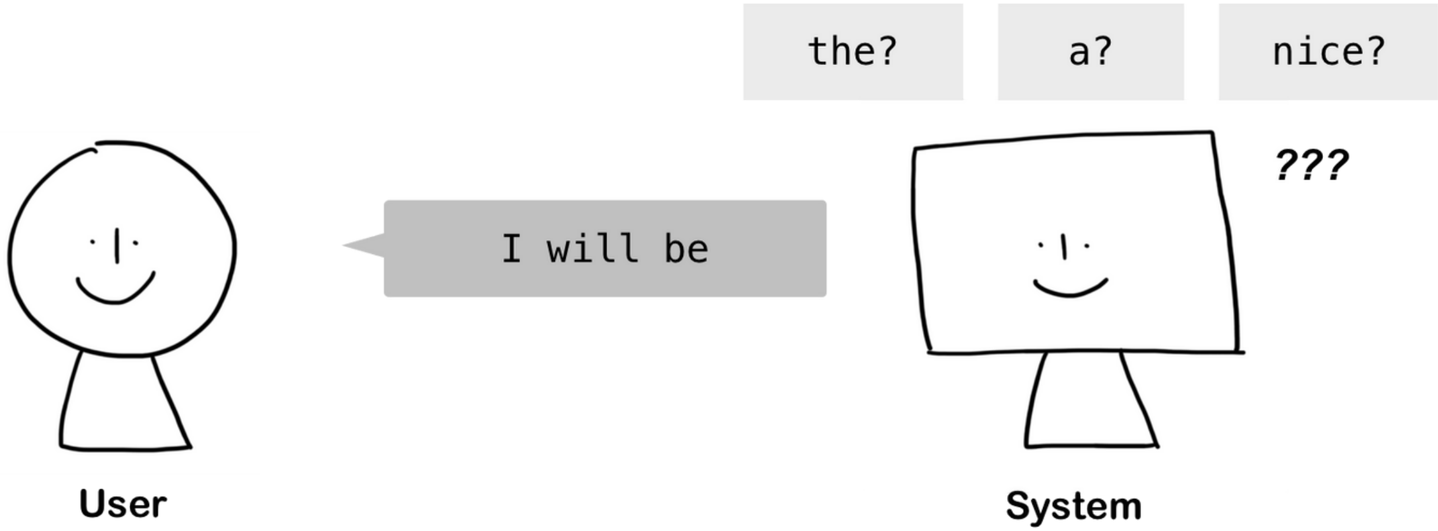

The **prefix** of a sentence often **fails to capture the meaning of the sentence**.

Mina Lee, Tatsunori Hashimoto, Percy Liang. Learning autocomplete systems as a communication games. NeurIPS 2019 Workshop.

Why Keyword-based Autocomplete?

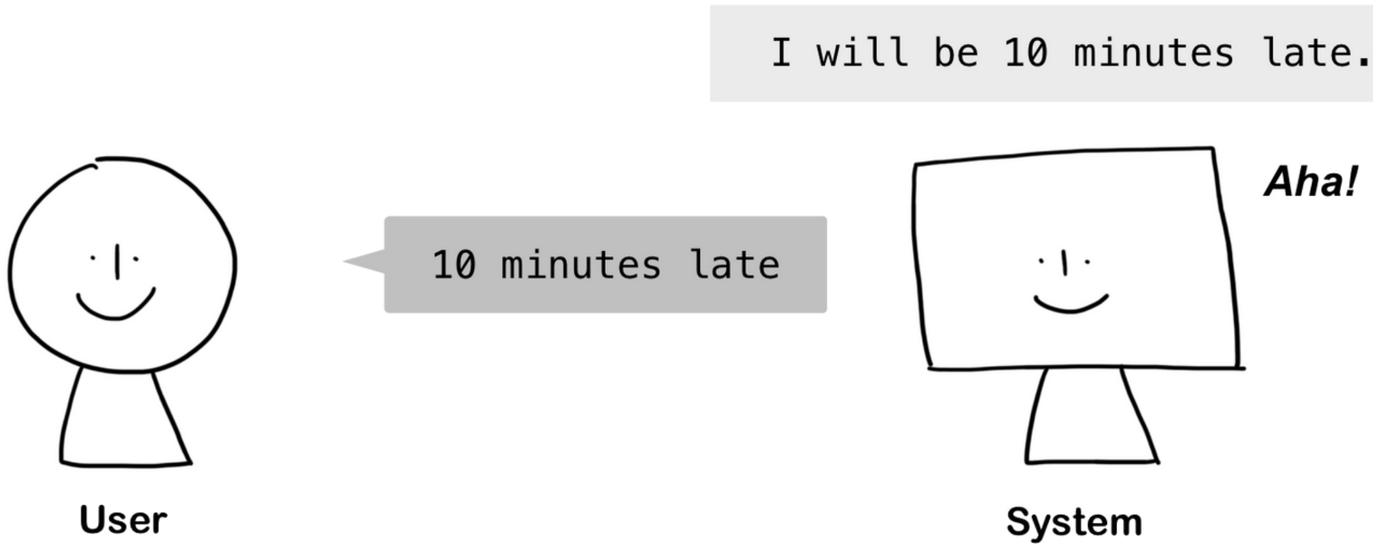

Keywords can better capture the meaning of a sentence.

Mina Lee, Tatsunori Hashimoto, Percy Liang. Learning autocomplete systems as a communication games. NeurIPS 2019 Workshop.

Autocomplete Task

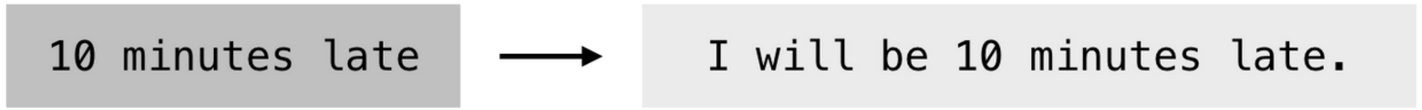

Keywords **Target sentence**

Solution Collect data?

Fail to capture user's tendency to adapt to the system

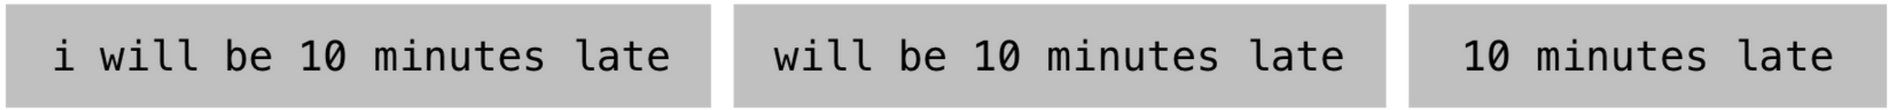

Mina Lee, Tatsunori Hashimoto, Percy Liang. Learning autocomplete systems as a communication games. NeurIPS 2019 Workshop.

Learn both keyword generation and sentence completion strategies (unsupervised learning)

Our Idea: Communication Game

Frame it as a **cooperative communication game**

x

x

z

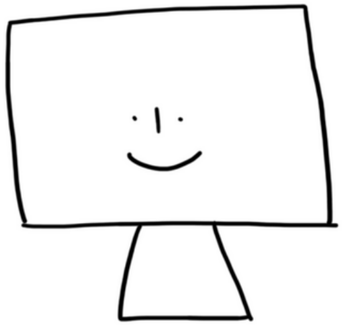

$\underbrace{\mathbb{E}_{q(z|x)}[\#\text{tokens}(z)]}_{\text{cost}(x)}$

user wants to write as little as possible

$q(z|x)$

$p(x|z)$

$\underbrace{\mathbb{E}_{q(z|x)}[-\log p(x|z)]}_{\text{loss}(x)}$

system wants to guess user's intention as accurately as possible

Mina Lee, Tatsunori Hashimoto, Percy Liang. Learning autocomplete systems as a communication games. NeurIPS 2019 Workshop.

Modeling with Autoencoder

Learn both the user and the system **without supervision**

x

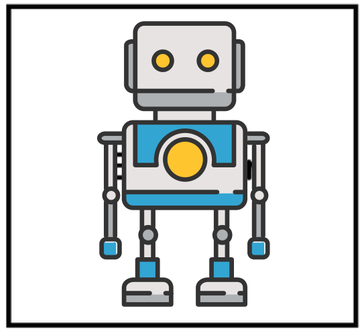

$q(z|x)$

z

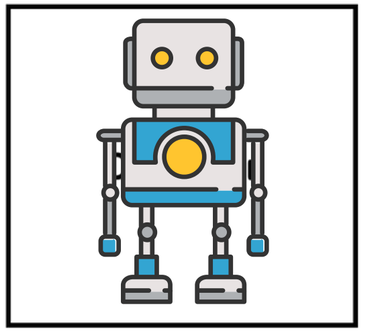

$p(x|z)$

x

No need to collect user data

Near 50% reduction in typing time

When users know what they want to write and adapt to the system, this way of using language model for autocomplete can boost productivity by reducing writing time.

Optimization

Constrained objective $\min \mathbb{E}[\text{cost}(x)] \quad \text{subject to } \mathbb{E}[\text{loss}(x)] \leq \epsilon$ ✔

Lagrangian $\min_{\alpha,\beta} \max_{\lambda \geq 0} J(\alpha, \beta, \lambda)$

$$J(\alpha, \beta, \lambda) = \mathbb{E}[\text{cost}(x, \alpha)] + \lambda(\mathbb{E}[\text{loss}(x, \alpha, \beta)] - \epsilon)$$

Policy gradient

$$\nabla_\alpha J(\alpha, \beta, \lambda) = \mathbb{E}[\mathbb{E}_{q_\alpha(z|x)}[\nabla_\alpha \log q_\alpha(z \mid x) \cdot G(x, z)]]$$

$$G(x, z) = \underbrace{\#\text{tokens}(z)}_{\text{per-example cost}} + \lambda(\underbrace{-\log p_\beta(x \mid z)}_{\text{per-example loss}} - \epsilon).$$

Mina Lee, Tatsunori Hashimoto, Percy Liang. Learning autocomplete systems as a communication games. NeurIPS 2019 Workshop.

다양한 형태의 자동완성

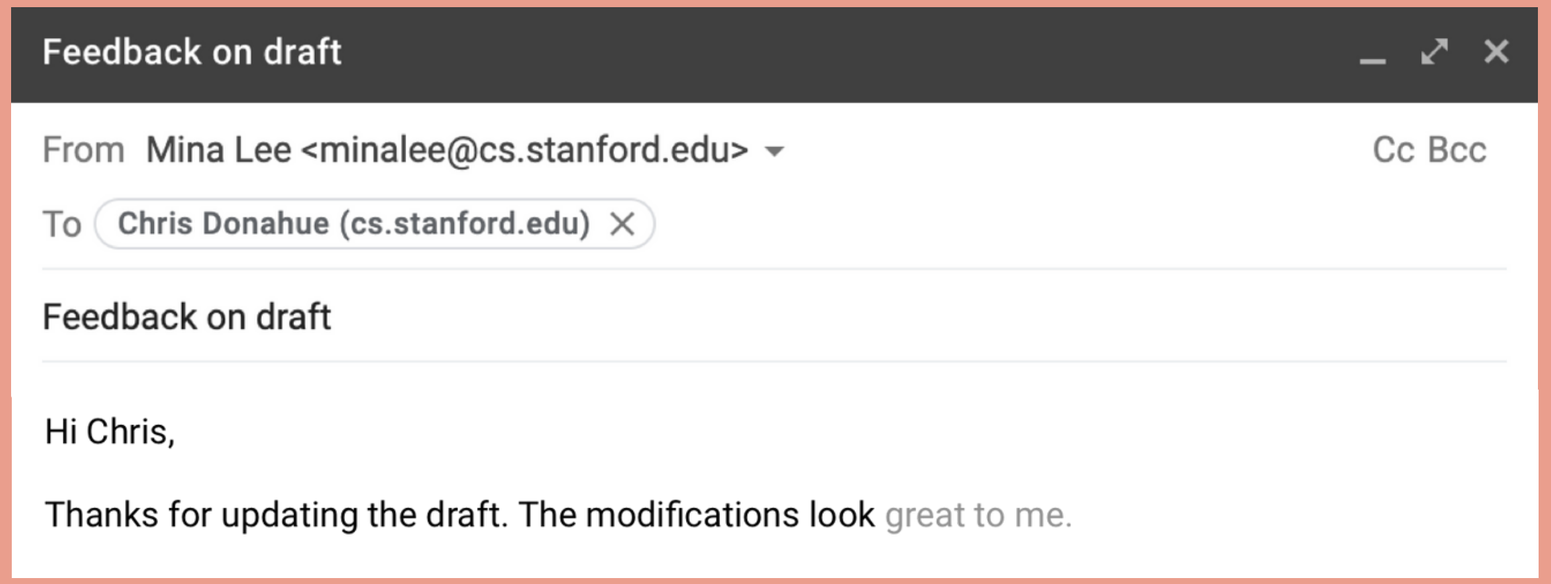

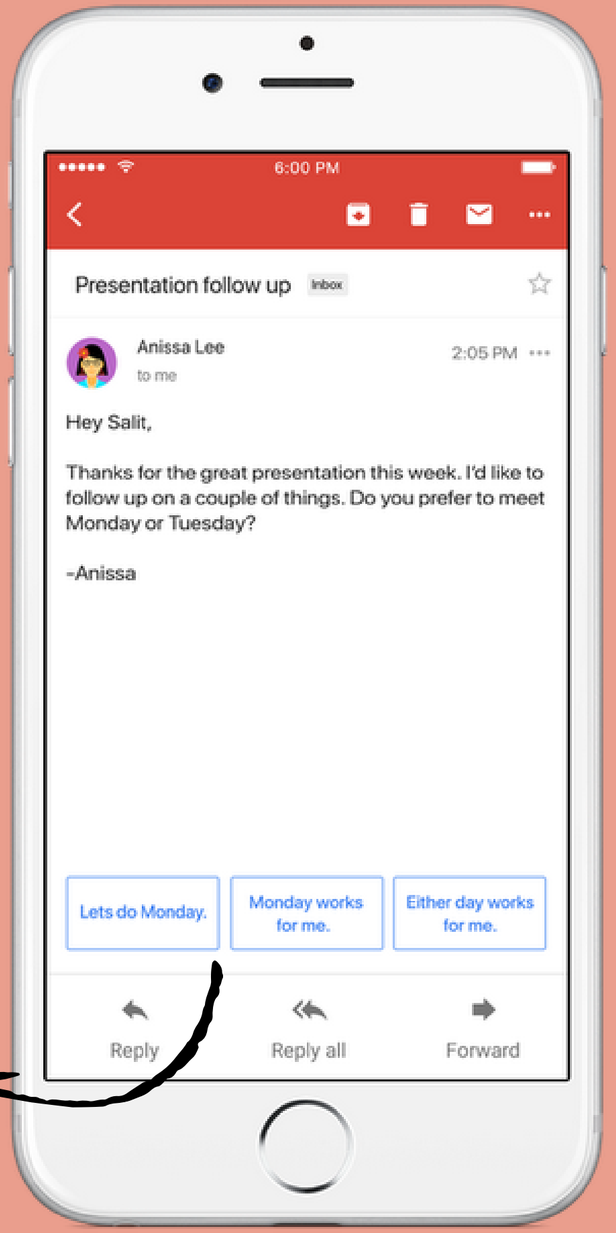

"Usage of Smart Reply has grown significantly,
making up about 12% of replies in Inbox on mobile"

Smart Compose (https://ai.googleblog.com/2018/05/smart-compose-using-neural-networks-to.html)
Smart Reply (https://ai.googleblog.com/2017/05/efficient-smart-reply-now-for-gmail.html)

다양한 형태의 자동완성

Javascript | **Python** | Typescript | PHP | Java | C++ | Go | Rust

```python
from argparse import ArgumentParser
from boto3 import Session
from botocore.exceptions import BotoCoreError, ClientError

cli = ArgumentParser(description="GetLexicon e
cli.add_argument("name", type=str, metavar="LE
arguments = cli.parse_args()

sess
```

```
session = Session(profile
session = Session(client
session = Session
```

sentiment.ts | write_sql.go | **parse_expenses.py** | addresses.rb

```python
import datetime

def parse_expenses(expenses_string):
    """Parse the list of expenses and return the list of triples (date, value, currency).
    Ignore lines starting with #.
    Parse the date using datetime.
    Example expenses_string:
        2016-01-02 -34.01 USD
        2016-01-03 2.59 DKK
        2016-01-03 -2.72 EUR
    """
    expenses = []
    for line in expenses_string.splitlines():
        if line.startswith("#"):
            continue
        date, value, currency = line.split(" ")
        expenses.append((datetime.datetime.strptime(date, "%Y-%m-%d"),
                         float(value),
                         currency))
    return expenses
```

Copilot

TabNine (https://www.tabnine.com/)
GitHub Copilot (https://copilot.github.com/)

다양한 형태의 자동완성

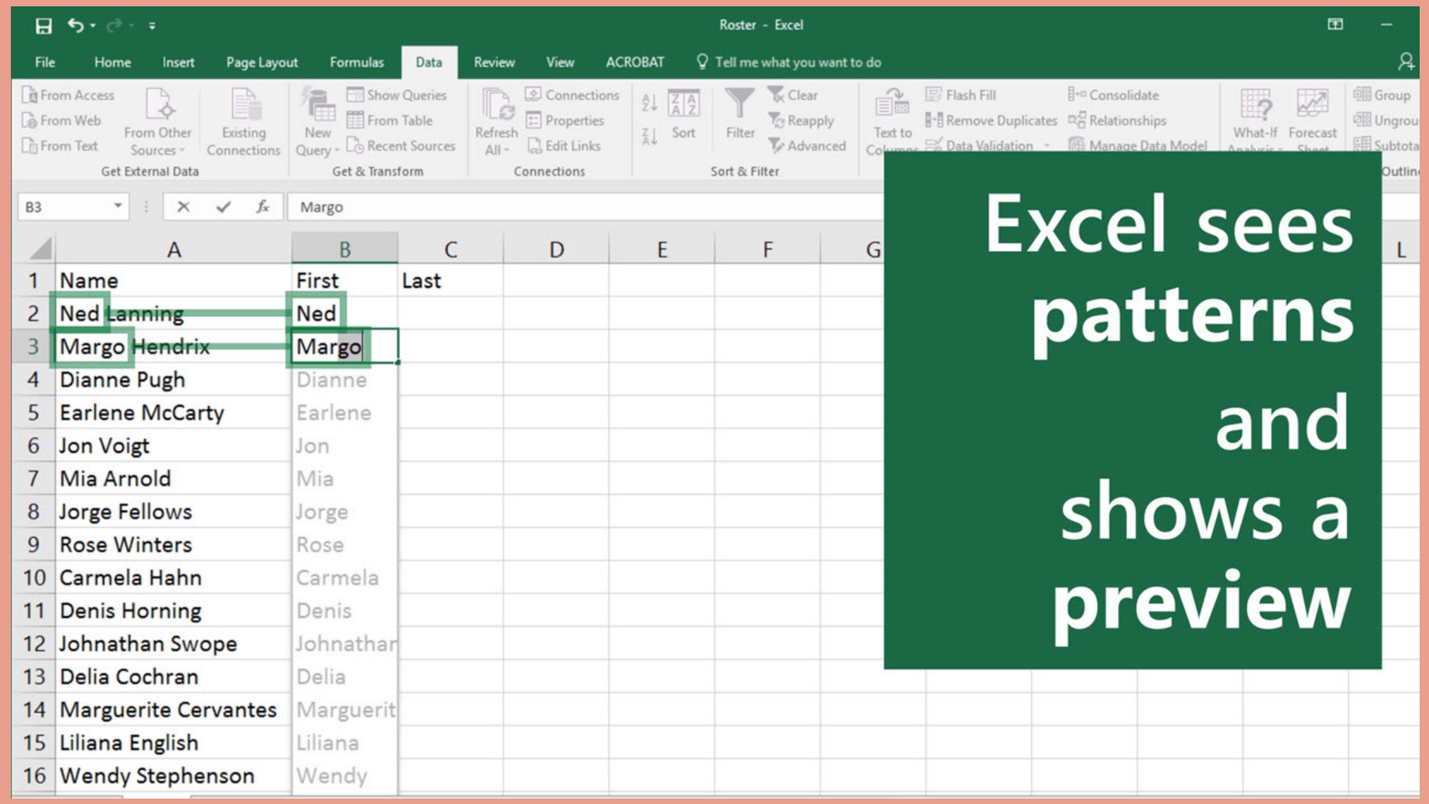

Flash Fill (https://support.microsoft.com/en-us/office/using-flash-fill-in-excel-3f9bcf1e-db93-4890-94a0-1578341f73f7)
Sumit Gulwani. Automating String Processing in Spreadsheets using Input-Output Examples. POPL 2011.

다양한 형태의 자동완성

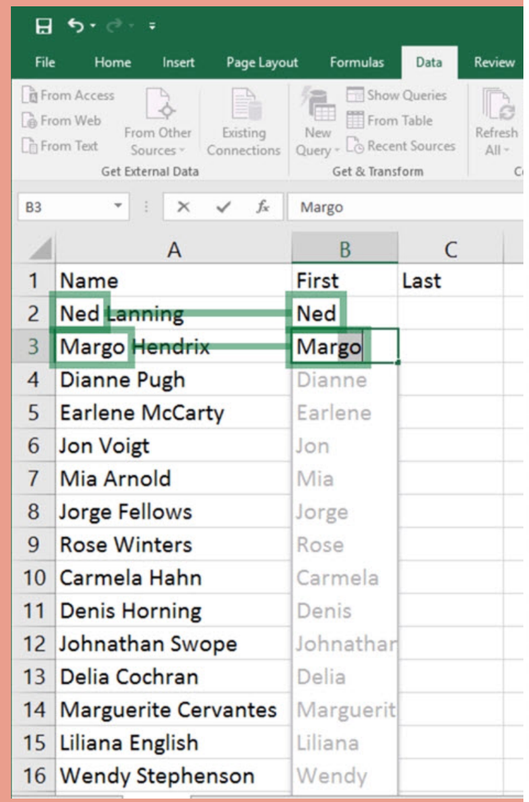

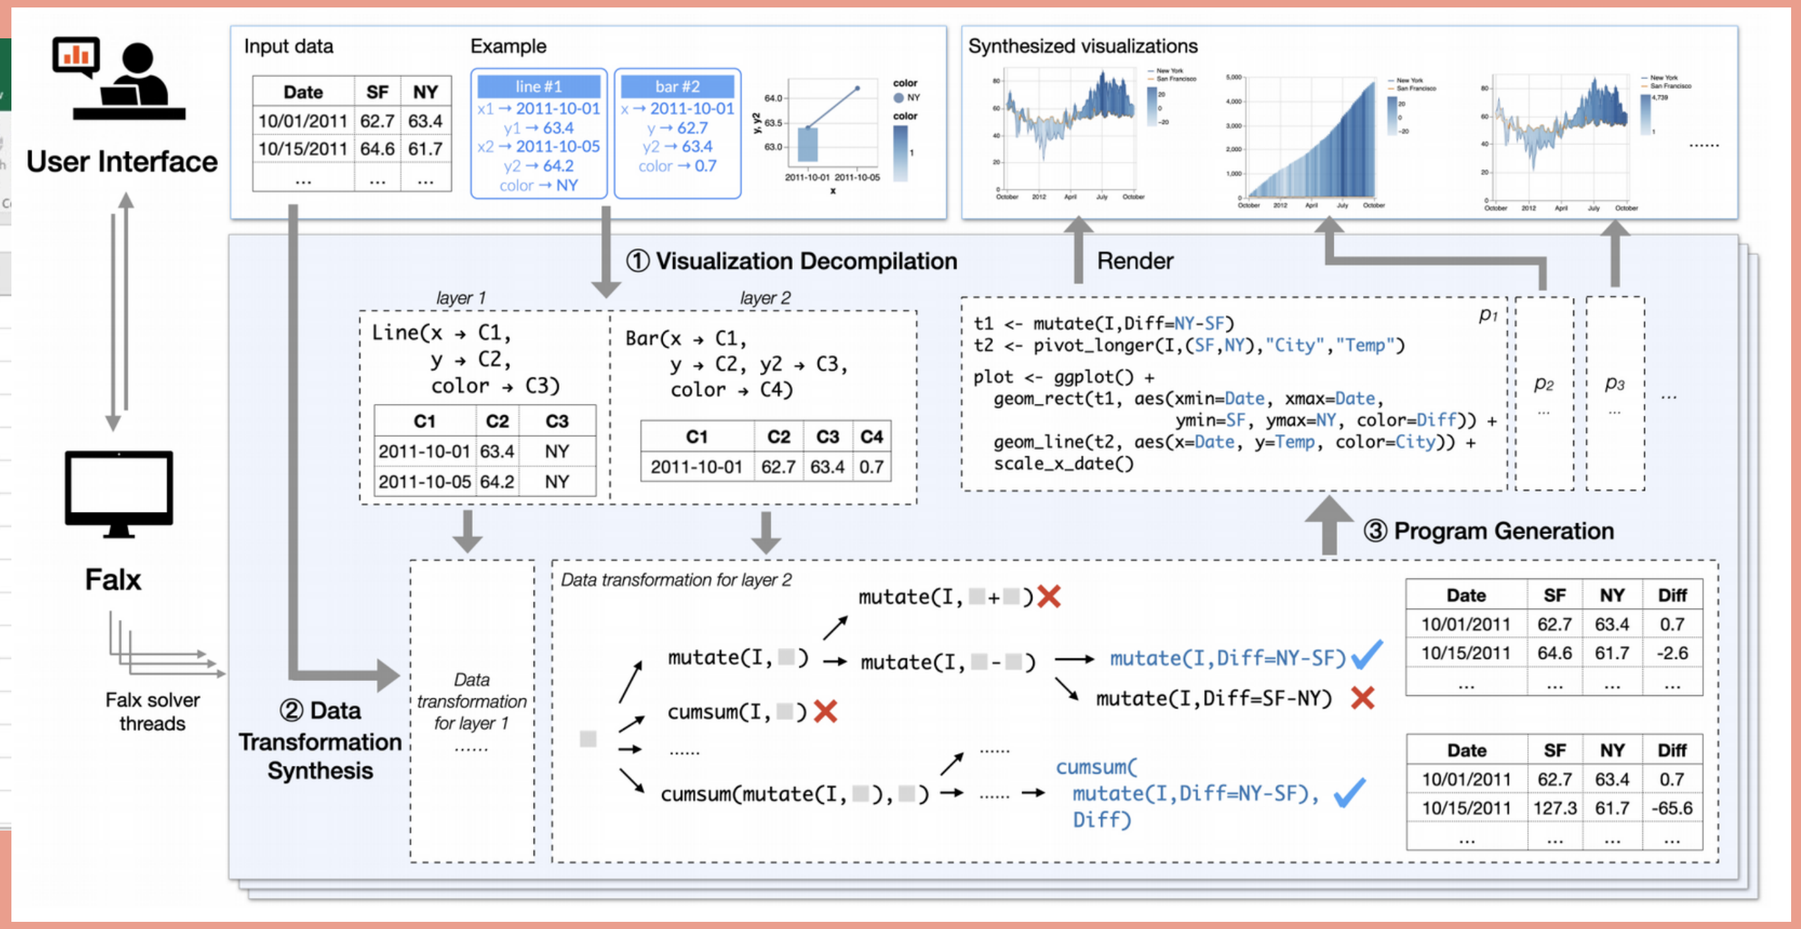

Flash Fill (https://support.microsoft.com/en-us/office/using-flash-fill-in-excel-3f9bcf1e-db93-4890-94a0-1578341f73f7)

Sumit Gulwani. Automating String Processing in Spreadsheets using Input-Output Examples. POPL 2011.

Chenglong Wang, Yu Feng, Rastislav Bodik, Isil Dillig, Alvin Cheung, Amy J. Ko. Falx: Synthesis-powered Visualization Authoring. CHI 2021.

AI로 생산성 증가시키기

Mina Lee, Tatsunori Hashimoto, Percy Liang. Learning autocomplete systems as a communication games. NeurIPS 2019 Workshop.
Chris Donahue, **Mina Lee**, Percy Liang. Enabling Language Models to Fill in the Blanks. ACL 2020.
Mina Lee, Chris Donahue, Robin Jia, Alexander Iybor, Percy Liang. Swords ⚔: A Benchmark for Lexical Substitution with Improved Data Coverage and Quality. NAACL 2021.

Code: https://github.com/chrisdonahue/ilm
Demo: https://chrisdonahue.com/ilm

빈칸 채우기

When you are editing or revising how can AI help you connect ideas?

Chris Donahue
Stanford University

Mina Lee
Stanford University

Percy Liang
Stanford University

Chris Donahue, **Mina Lee**, Percy Liang. Enabling Language Models to Fill in the Blanks. ACL 2020.

Why Filling in the Blanks?

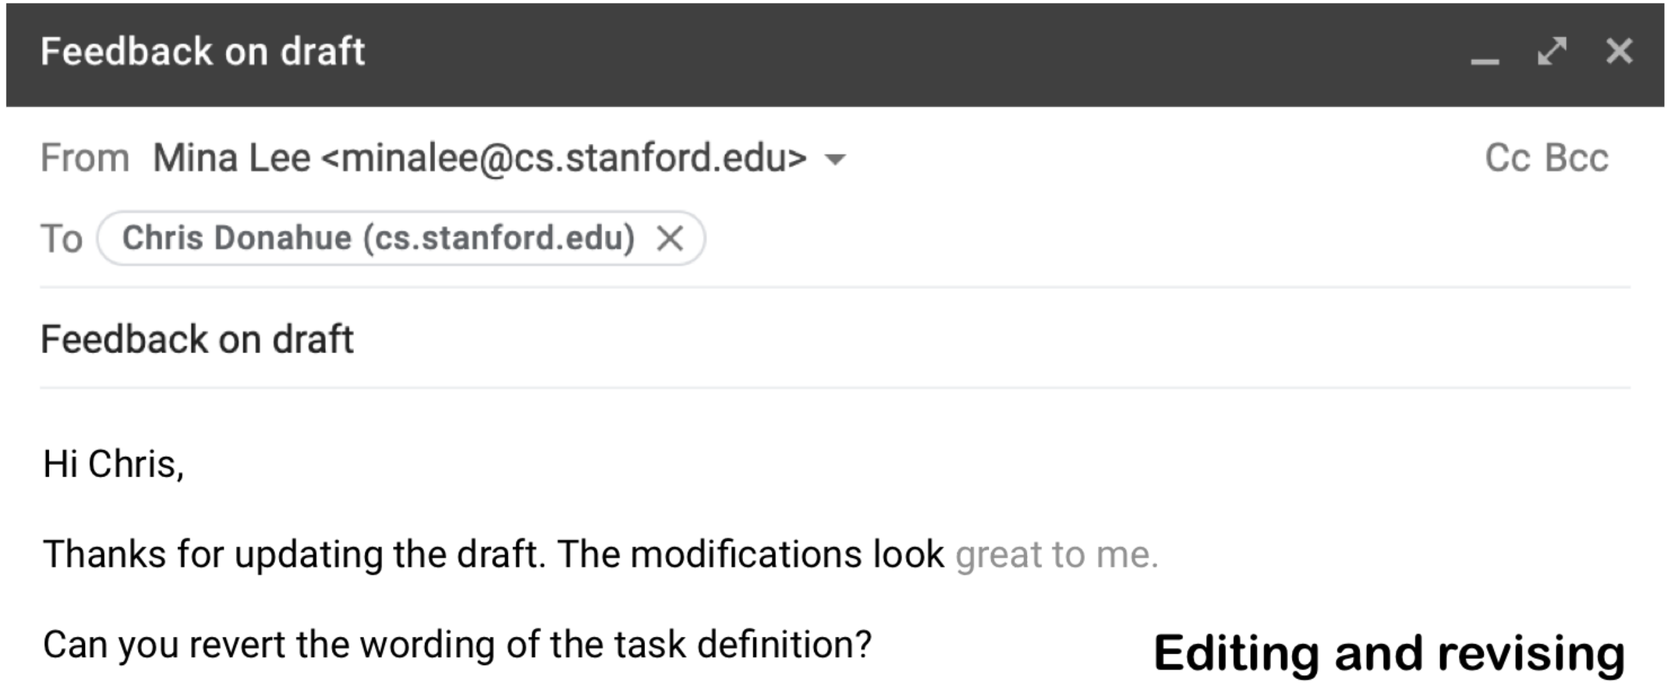

Chris Donahue, **Mina Lee**, Percy Liang. Enabling Language Models to Fill in the Blanks. ACL 2020.

Why Filling in the Blanks?

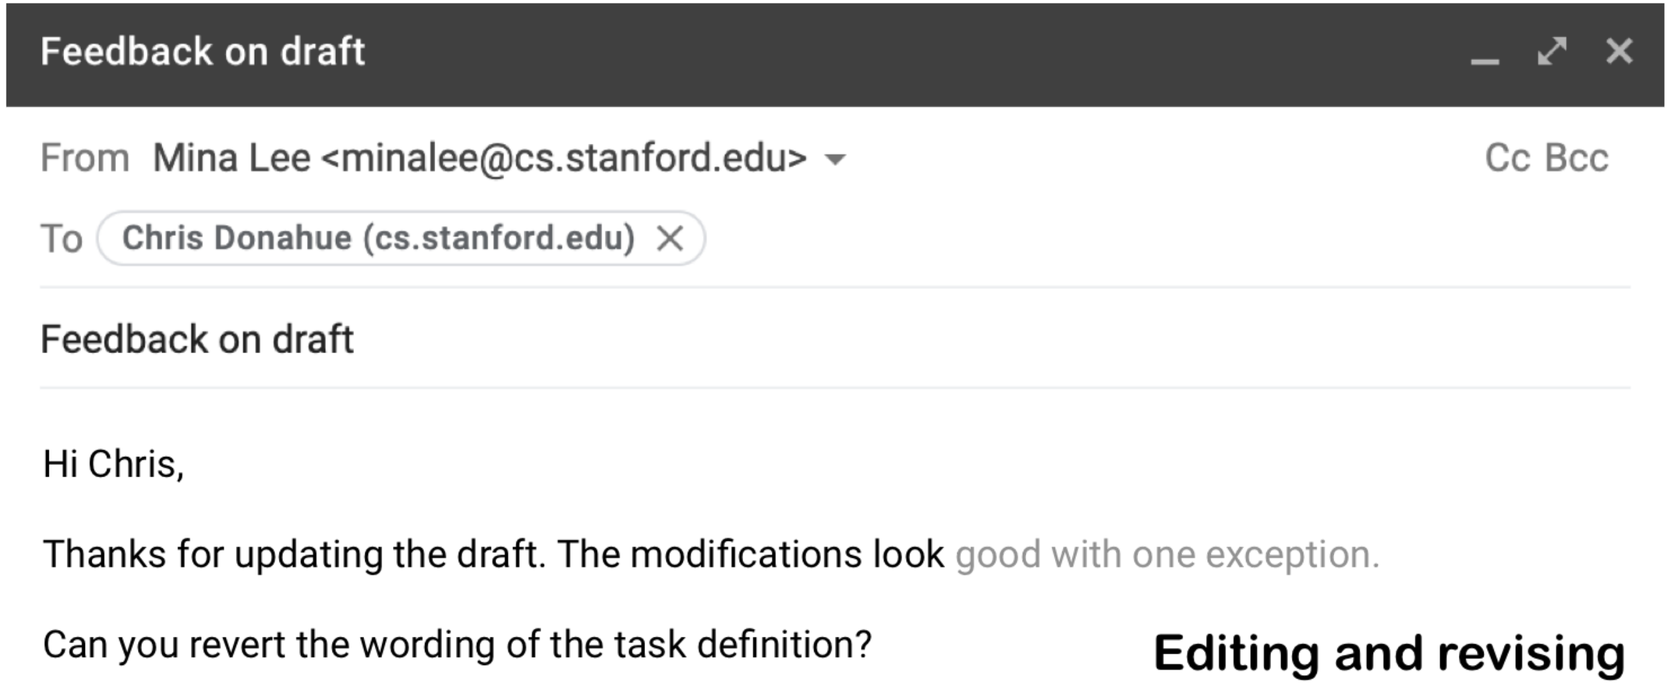

Chris Donahue, **Mina Lee**, Percy Liang. Enabling Language Models to Fill in the Blanks. ACL 2020.

Text Infilling

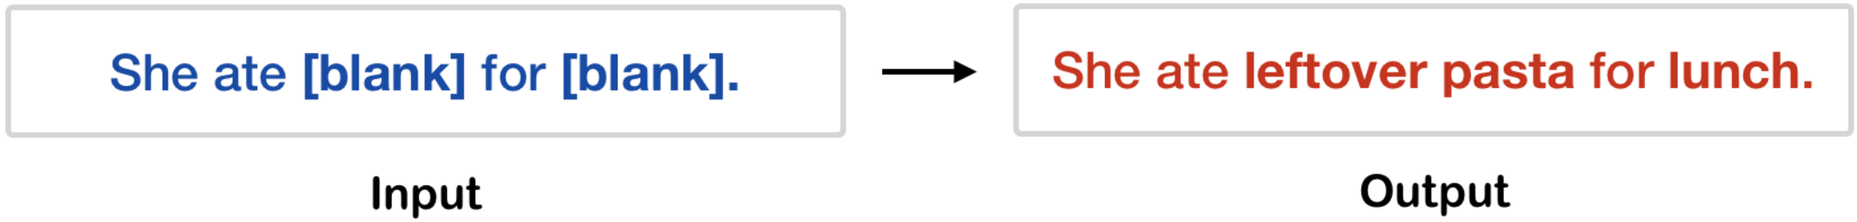

Given incomplete text with **[blank]**s, predict complete text

Arbitrary number of blanks

Variable length spans (e.g. word, sentence, paragraph)

Chris Donahue, **Mina Lee**, Percy Liang. Enabling Language Models to Fill in the Blanks. ACL 2020.

Demo

chrisdonahue.com/ilm

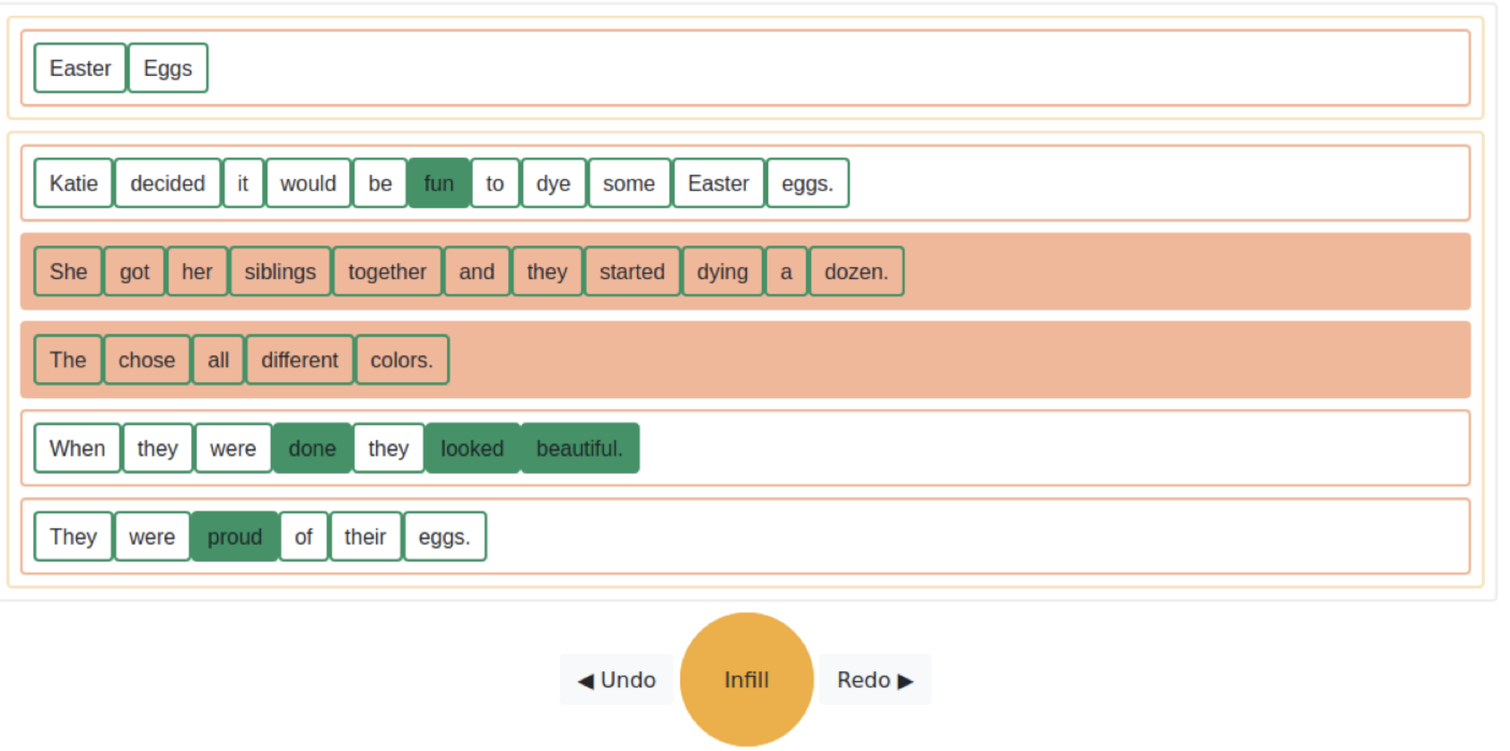

Chris Donahue, **Mina Lee**, Percy Liang. Enabling Language Models to Fill in the Blanks. ACL 2020.

다양한 방법으로 빈칸 채우기

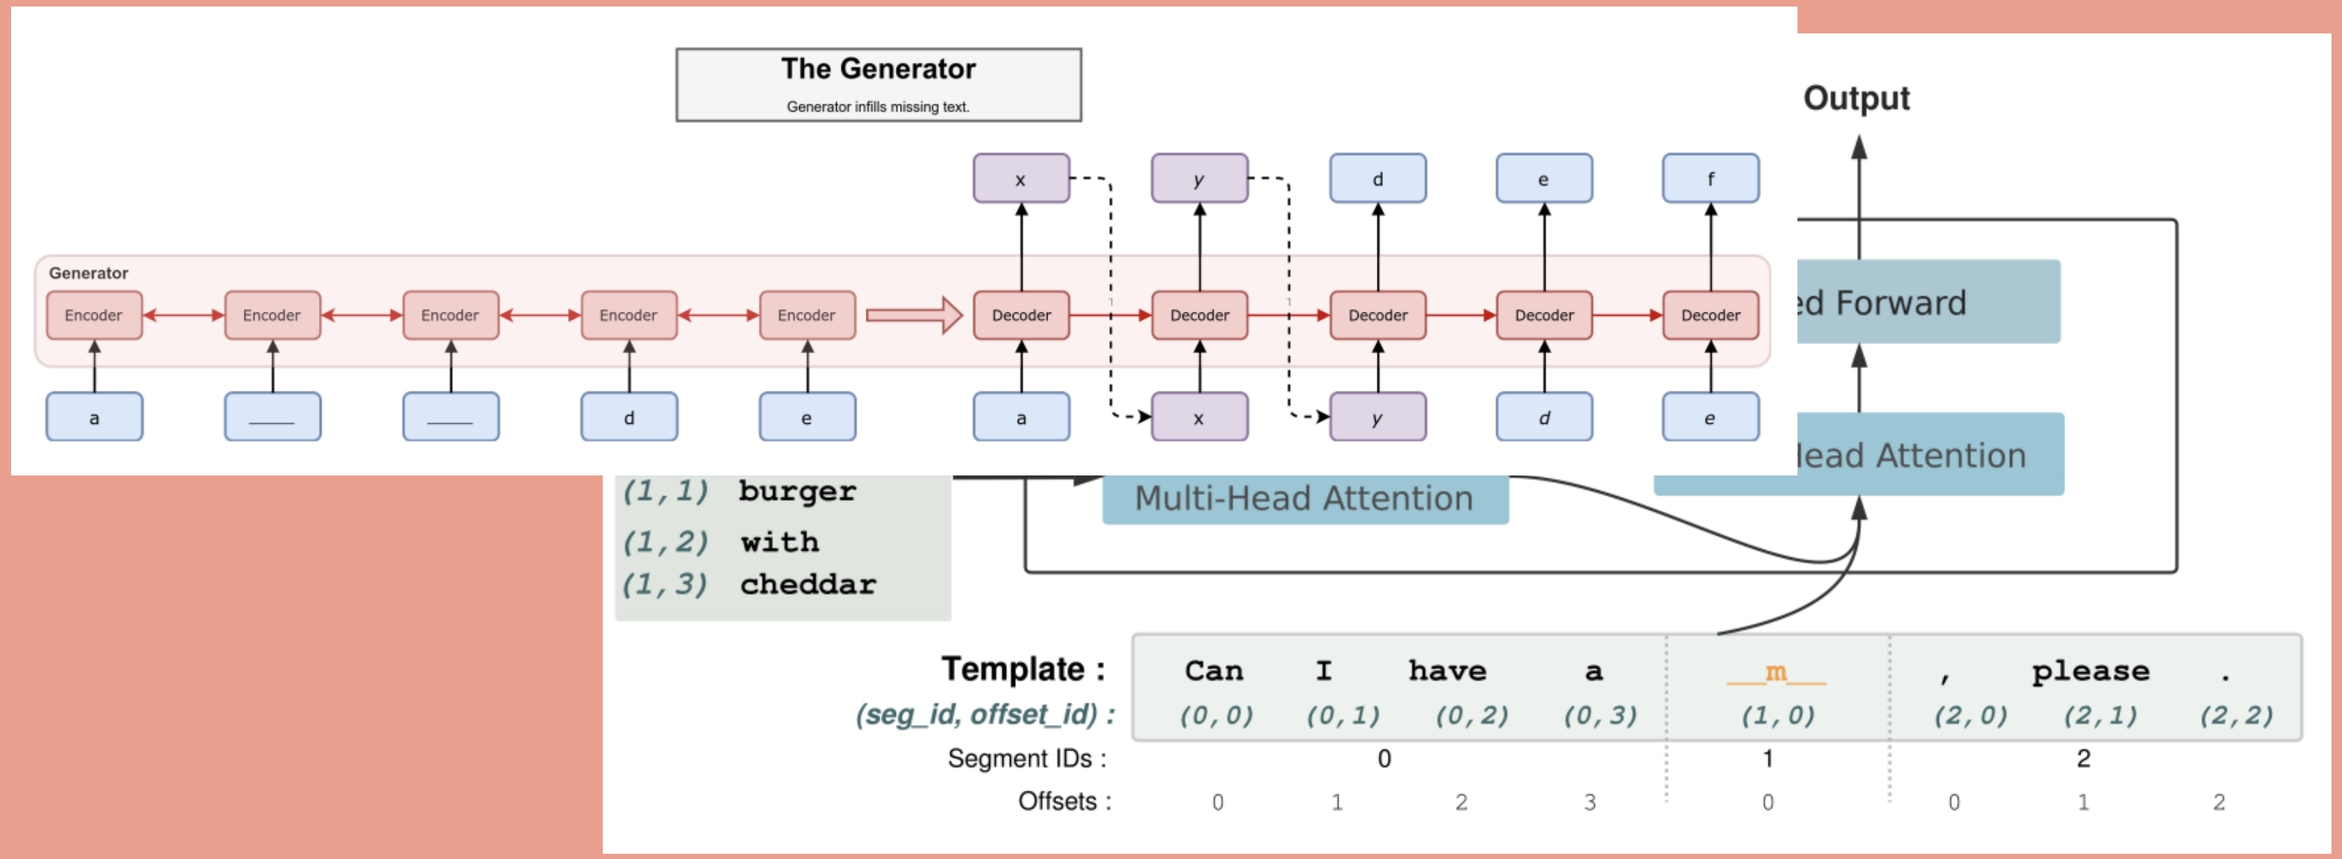

William Fedus, Ian Goodfellow, Andrew M. Dai. MaskGAN: Better Text Generation via Filling in the ____. ICLR 2018
Wanrong Zhu, Zhiting Hu, Eric P. Xing. Text infilling. arXiv 2019.

Pre-train & fine-tune language models

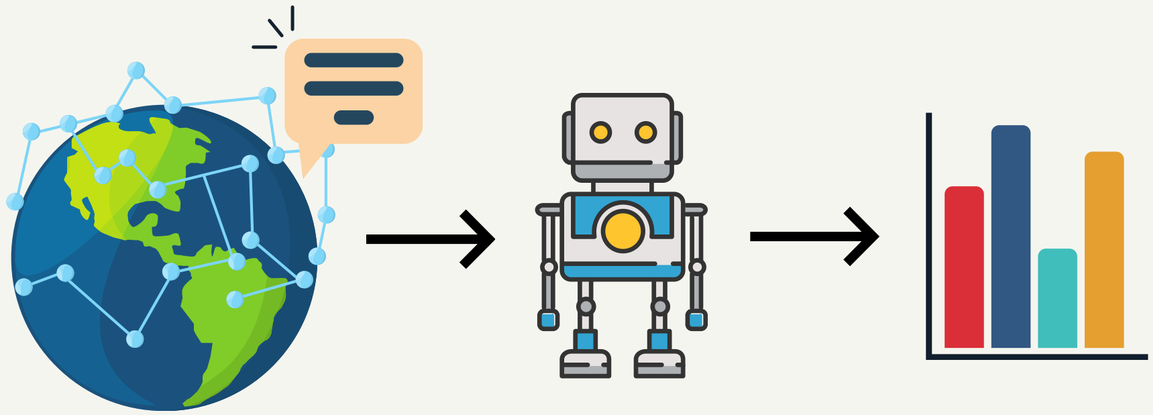

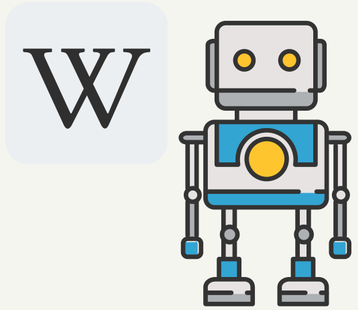

Q & A

Input: 한국의 수도
Output: 서울

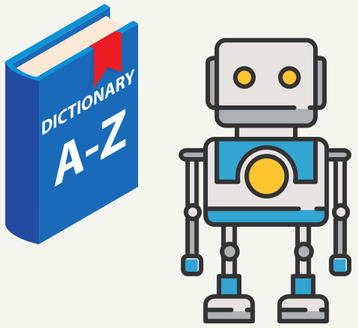

Translation

Input: 안녕하세요
Output: Hi

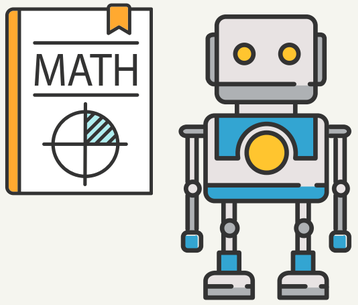

Math

Input: 3 + 5
Output: 8

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever. Language Models are Unsupervised Multitask Learners. 2019.

Our Idea: Infilling by Language Modeling (ILM)

Manufacture infilling examples

Training time

She ate leftover pasta for lunch.

Data

She ate [blank] for [blank]. [sep] leftover pasta [answer] lunch [answer]

New data

Chris Donahue, **Mina Lee**, Percy Liang. Enabling Language Models to Fill in the Blanks. ACL 2020.

Our Idea: Infilling by Language Modeling (ILM)

Test time

Use fine-tuned LM to infill

He drinks [blank] after [blank]. [sep] water [answer] running [answer]

Input Target

He drinks water after running.

Output

Chris Donahue, **Mina Lee**, Percy Liang. Enabling Language Models to Fill in the Blanks. ACL 2020.

Humans could not correctly identify the machine-generated sentences **45% of times**.
(Other models such as BERT and Zhu et al. fooled humans 20% and 29%, respectively).

Human Evaluation: Turing Test

Identify one of the five sentences generated by machine.

Patty was excited about having her friends over.
She had been working hard preparing the food.
Patty knew her friends wanted pizza.
All of her friends arrived and were seated at the table.
Patty had a great time with her friends.

Wanrong Zhu, Zhiting Hu, Eric P. Xing. Text infilling. arXiv 2019.
Jacob Devlin, Ming-Wei Chang, Kenton Lee, Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. ACL 2019.
Chris Donahue, **Mina Lee**, Percy Liang. Enabling Language Models to Fill in the Blanks. ACL 2020.

AI로 생산성 증가시키기

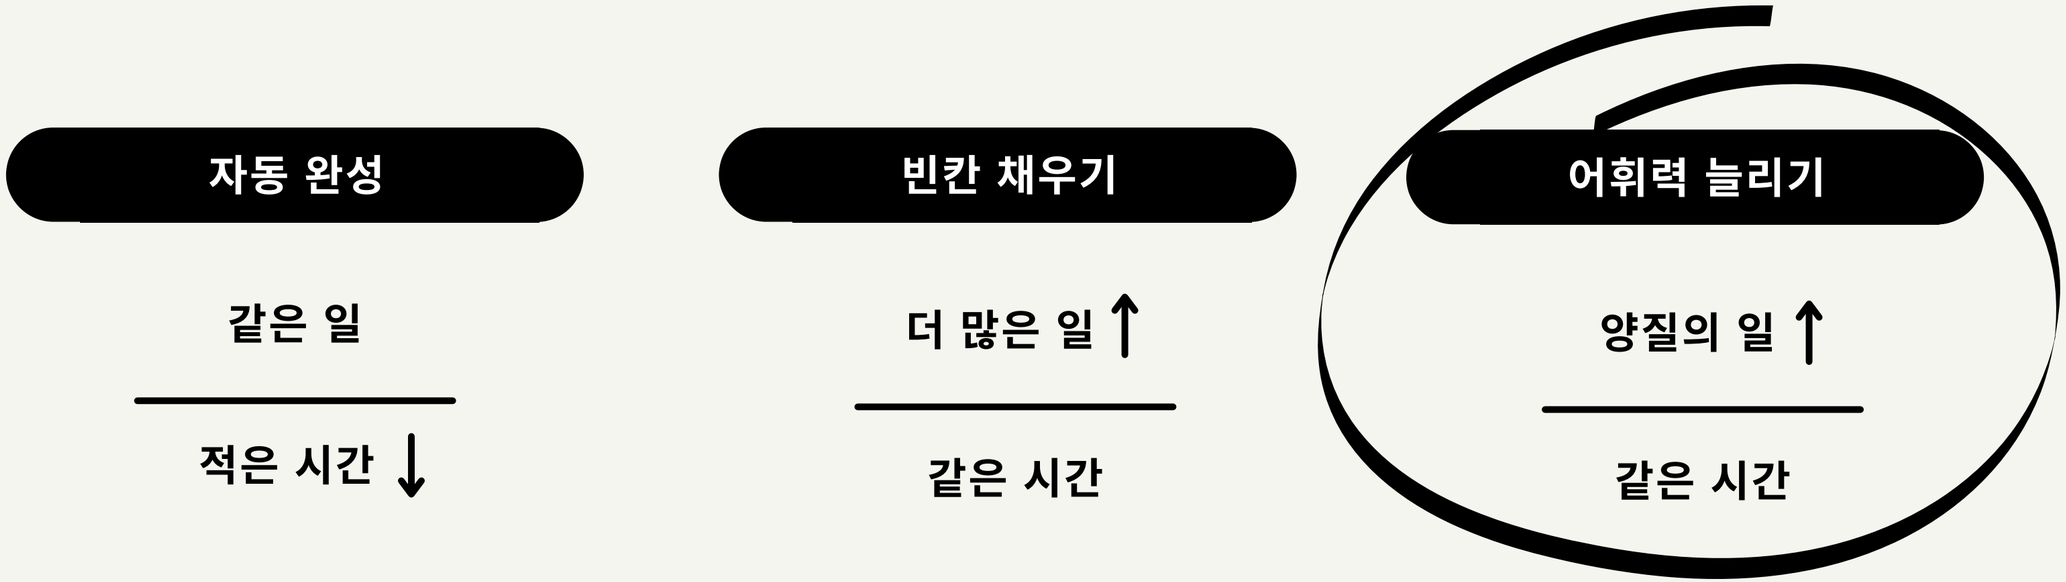

Mina Lee, Tatsunori Hashimoto, Percy Liang. Learning autocomplete systems as a communication games. NeurIPS 2019 Workshop.
Chris Donahue, **Mina Lee**, Percy Liang. Enabling Language Models to Fill in the Blanks. ACL 2020.
Mina Lee, Chris Donahue, Robin Jia, Alexander Iybor, Percy Liang. Swords ⚔: A Benchmark for Lexical Substitution with Improved Data Coverage and Quality. NAACL 2021.

Benchmark: https://github.com/p-lambda/swords

Fill in the blanks

01 Autocomplete — 02 — 03 Contextual thesaurus

어휘력 늘리기

Can AI help you make better word choices?

Mina Lee
Stanford University

Chris Donahue
Stanford University

Robin Jia
Stanford University

Alexander Iybor
Stanford University

Percy Liang
Stanford University

Mina Lee, Chris Donahue, Robin Jia, Alexander Iybor, Percy Liang. Swords ⚔: A Benchmark for Lexical Substitution with Improved Data Coverage and Quality. NAACL 2021.

Why word substitution?

SWORDS ⚔: A Benchmark for Lexical Substitution with Improved Data Coverage and Quality

Mina Lee* Chris Donahue* Robin Jia Alexander Iyabor Percy Liang
Stanford University
{minalee, ..., robinjia, aiyabor, pliang}@cs.stanford.edu

1 Introduction

Imagine you are writing the message "I read an amazing paper today" to a colleague, but you want to choose a more descriptive adjective to replace "amazing." At first you might think of substitutes like "awesome" and "great," but feel that these are also unsatisfactory. You turn to a thesaurus for inspiration, but among reasonable alternatives like "incredible" and "fascinating" are words like "prodigious" which do not quite fit in your context. Ultimately, you choose to go with "fascinating," but reaching this decision required a non-trivial amount of time and effort.

*Equal contribution.

This one has some amazing plots.

Cool! I'll take a look!

awesome
great

June 8, 2021

I read an amazing paper today!

Mina Lee, Chris Donahue, Robin Jia, Alexander Iybor, Percy Liang. Swords ⚔: A Benchmark for Lexical Substitution with Improved Data Coverage and Quality. NAACL 2021.

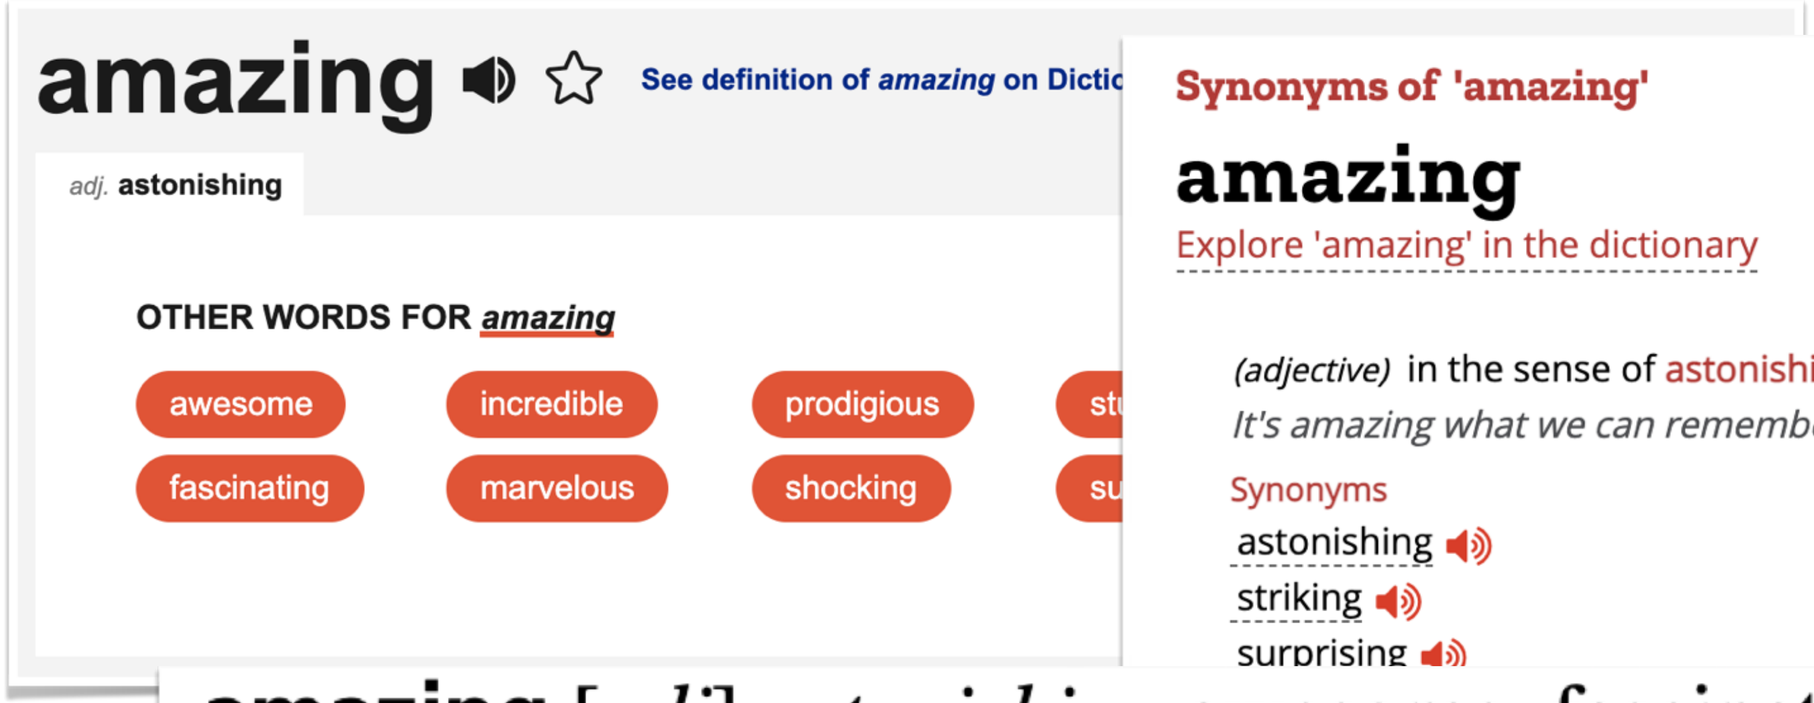

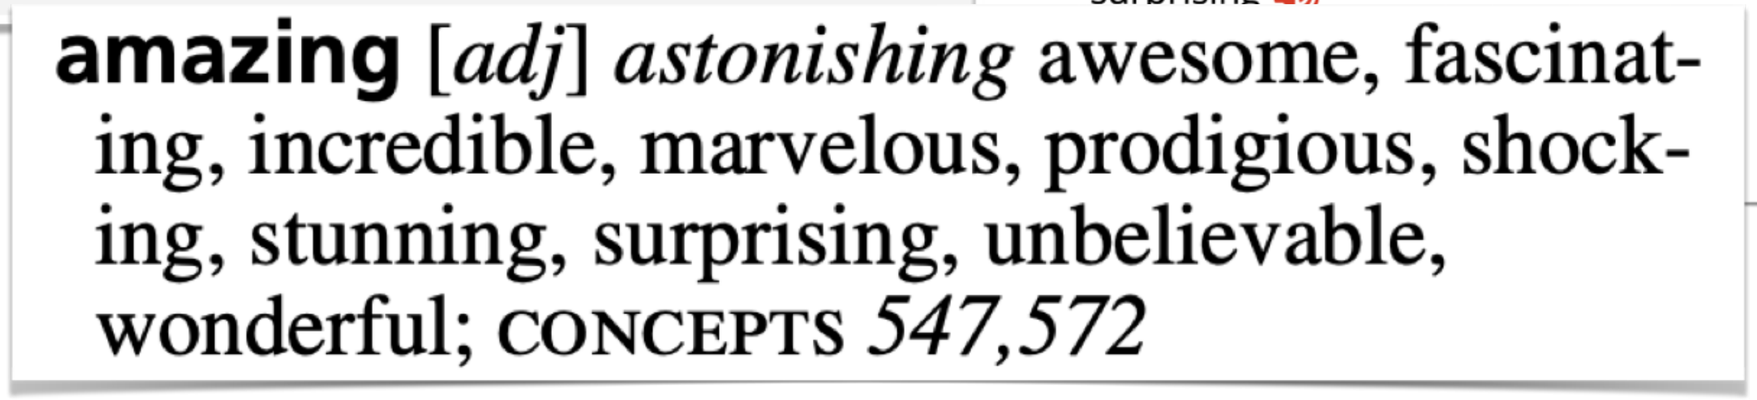

Thesaurus.com (https://www.thesaurus.com/browse/amazing, accessed on May 10, 2021)
Collins' synonyms (https://www.collinsdictionary.com/dictionary/english-thesaurus/amazing; accessed on May 10, 2021)
Barbara Ann Kipfer. 2013. Roget's 21st Century Thesaurus, Third Edition. Random House Publishing Group.

Lexical substitution

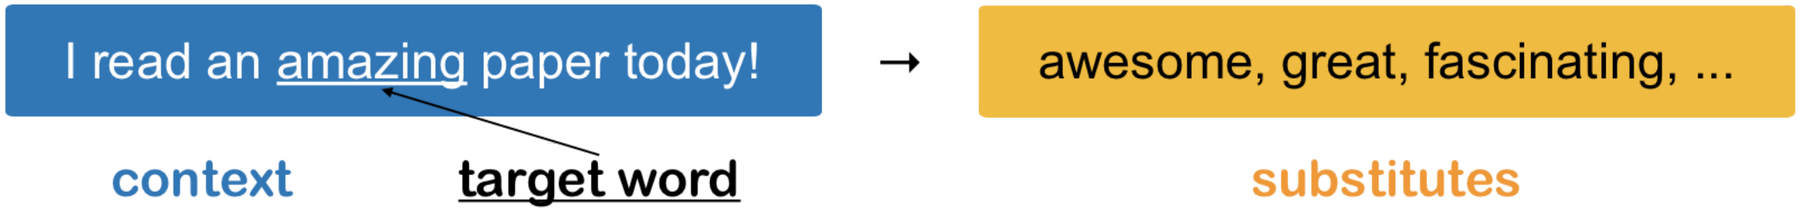

Task: Find **appropriate substitutes** for a **target word** in **context**

Related to practical application of a **writing assistant**

Can serve as a **downstream NLU task** for evaluating pre-trained models

Diana McCarthy. 2002. Lexical Substitution as a Task for WSD Evaluation. SIGLEX/SENSEVAL Workshop on Word Sense Disambiguation: Recent Successes and Future Directions.
Mina Lee, Chris Donahue, Robin Jia, Alexander Iybor, Percy Liang. Swords ⚔: A Benchmark for Lexical Substitution with Improved Data Coverage and Quality. NAACL 2021.

My favorite thing about her is her straightforward honesty.

CoInCo

candid (3), artless (1), blunt (1), complete (1), direct (1), forthright (1), frank (1), outspoken (1), plainspoken (1), truthful (1), unreserved (1)

Thesaurus

honest, candid, forthright, genuine, sincere, truthful, unequivocal, aboveboard, barefaced, direct, frank, guileless, honorable, laid on the line, unvarnished, ... 33 more

SWORDS ⚔️

sincere (80%), **genuine** (80%), **frank** (70%), **candid** (70%), **direct** (70%), **forthright** (70%), **uncomplicated** (60%), **up front** (60%), **clear-cut** (60%), **clear** (60%), **plainspoken** (60%), complete (50%), straight-arrow (50%), honest (50%), open (50%), blunt (50%), outspoken (50%), truthful (50%), plain-dealing (40%), undisguised (40%), unvarnished (40%), unreserved (40%), barefaced (40%), unequivocal (30%), upright (30%), simple (30%), veracious (30%), unconcealed (30%), like it is (30%), square-shooting (20%), upstanding (20%), undissembled (20%), manifest (20%), unambiguous (20%), pretenseless (20%), level (10%), laid on the line (10%), honorable (10%), guileless (10%), artless (0%) ...19 more

Mina Lee, Chris Donahue, Robin Jia, Alexander Iybor, Percy Liang. Swords ⚔: A Benchmark for Lexical Substitution with Improved Data Coverage and Quality. NAACL 2021.

Pre-train & fine-tune language models

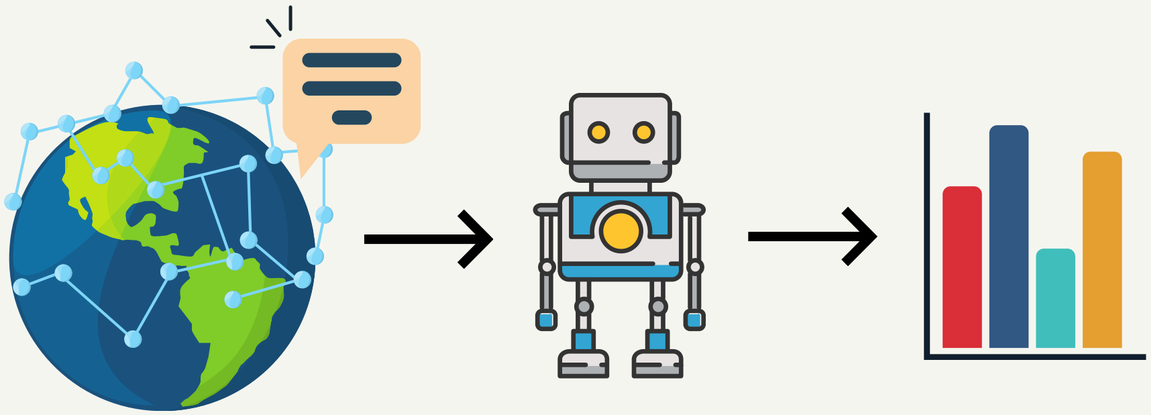

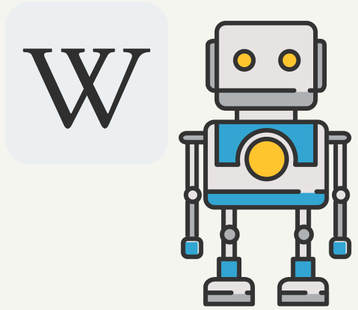

Q & A

Input: 한국의 수도
Output: 서울

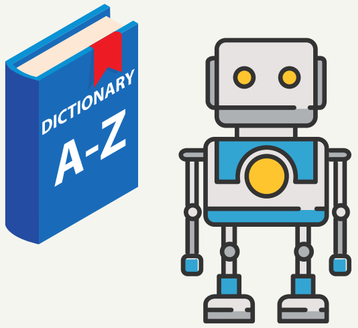

Translation

Input: 안녕하세요
Output: Hi

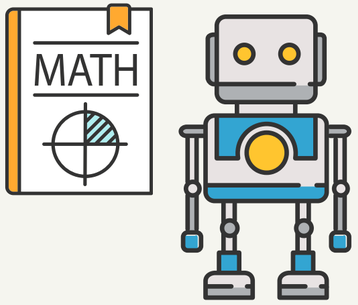

Math

Input: 3 + 5
Output: 8

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever. Language Models are Unsupervised Multitask Learners. 2019.

Pre-train & prompt-tune language models

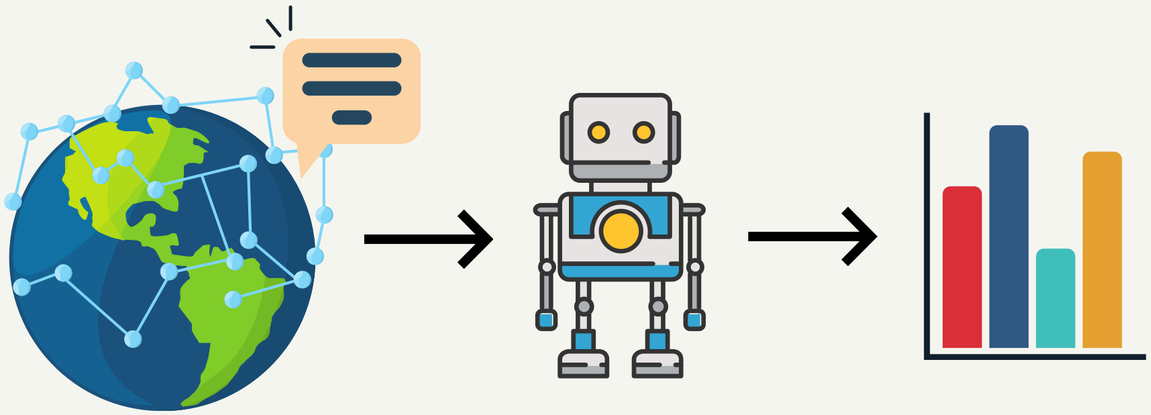

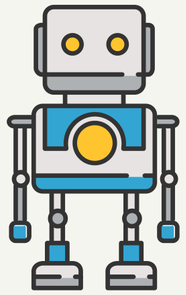

Q & A

Q. 한국의 수도는?
A. 서울

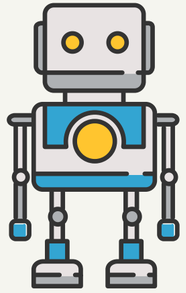

Translation

"안녕하세요"를 영어로
"Hi"

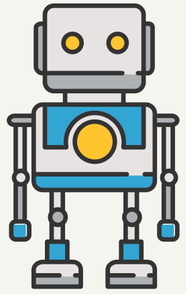

Math

3 + 5 = 8

Brown et al. Language Models are Few-Shot Learners. NeurIPS 2020.

Growing in size

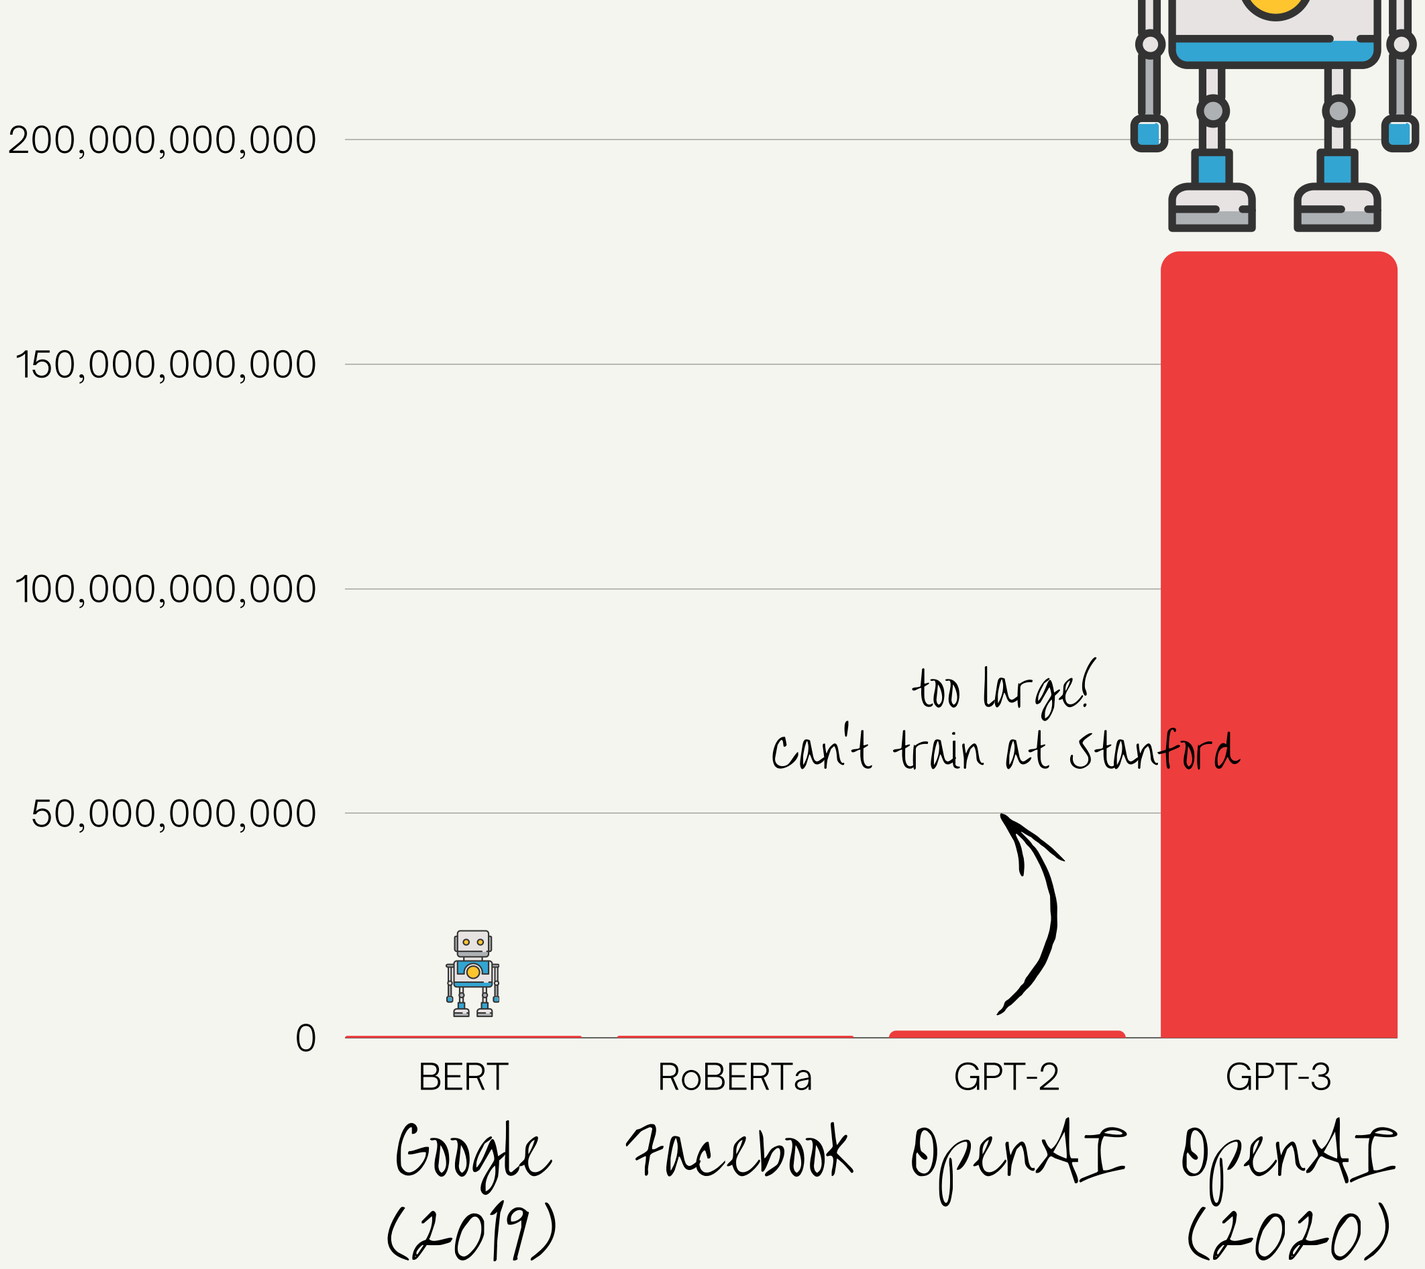

Emily M Bender, Timnit Gebru, Angelina McMillan-Major, Shmargaret Shmitchell. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? FACCT 2021.

Q. What are appropriate synonyms to replace the word "amazing" in the sentence "I read an amazing paper today"?

A. The word "amazing" can be replaced by the following words in the given sentence:

Wondrous

Incredible

Astounding

Remarkable

Astonishing

Extraordinary

Unbelievable

Stupendous

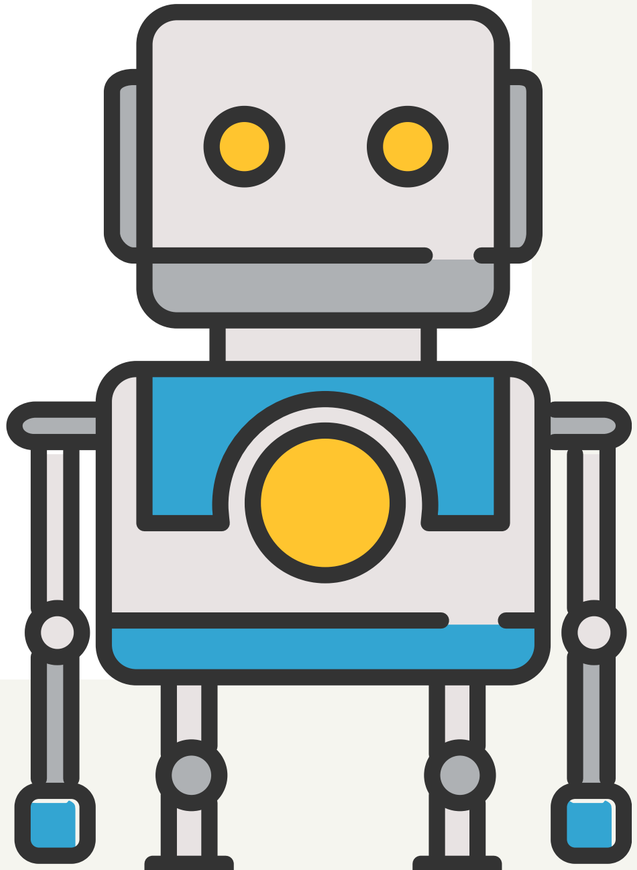

Brown et al. Language Models are Few-Shot Learners. NeurIPS 2020.

Q. What are appropriate synonyms to replace the word "straightforward" in the sentence "My favorite thing about her is her straightforward honesty"?
A. sincere, genuine, frank, candid, direct, forthright, uncomplicated, up front, clear-cut, clear, plainspoken, complete, straight-arrow, honest, open, blunt, outspoken, truthful

Q. What are appropriate synonyms to replace the word "amazing" in the sentence "I read an amazing paper today"?
A. wonderful, astonishing, outstanding, stunning, astounding, fantastic, marvelous, breathtaking, astounding, astonishing, incredible, remarkable, incredible, wonderful, stunning, amazing, fabulous, sensational, surprisin

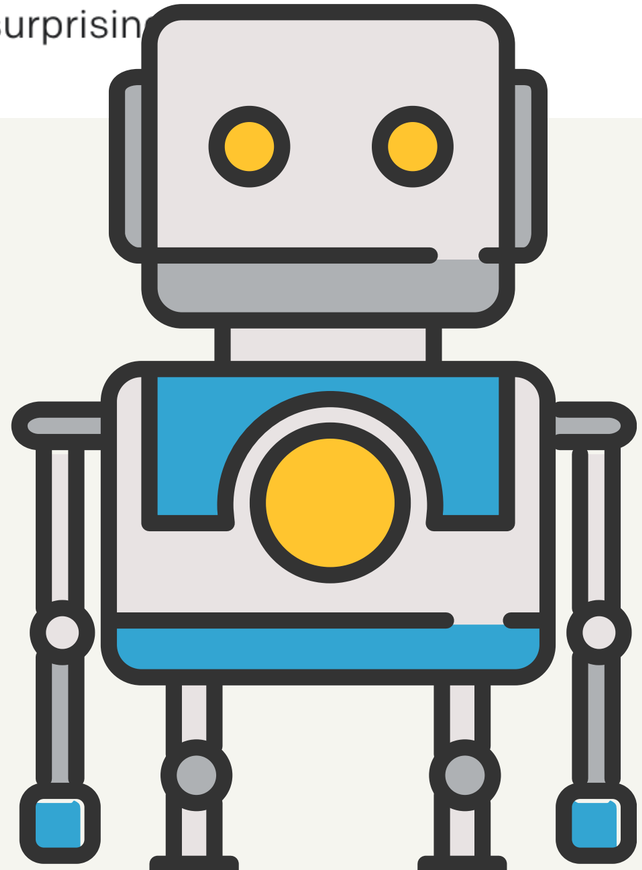

Brown et al. Language Models are Few-Shot Learners. NeurIPS 2020.

Evaluation

Lexical substitution systems and large pre-trained models

	F_1
Humans	**51.6**
CoInCo	34.6
Thesaurus	17.6
GPT-3	**35.0**
Wordtune	34.9
BERT-LS	31.6
BERT	31.5

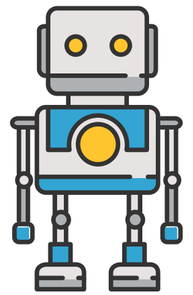

<

- **All methods** perform **worse** than human judgement (**Humans**)
- **GPT-3** and **Wordtune** perform **better** than human recollection (**CoInCo**)

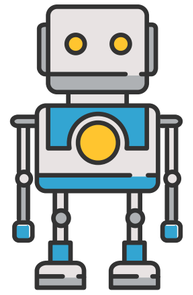

>

Mina Lee, Chris Donahue, Robin Jia, Alexander Iyabor, Percy Liang. Swords ⚔: A Benchmark for Lexical Substitution with Improved Data Coverage and Quality. NAACL 2021.

어휘력을 늘리기 위한 다양한 애플리케이션

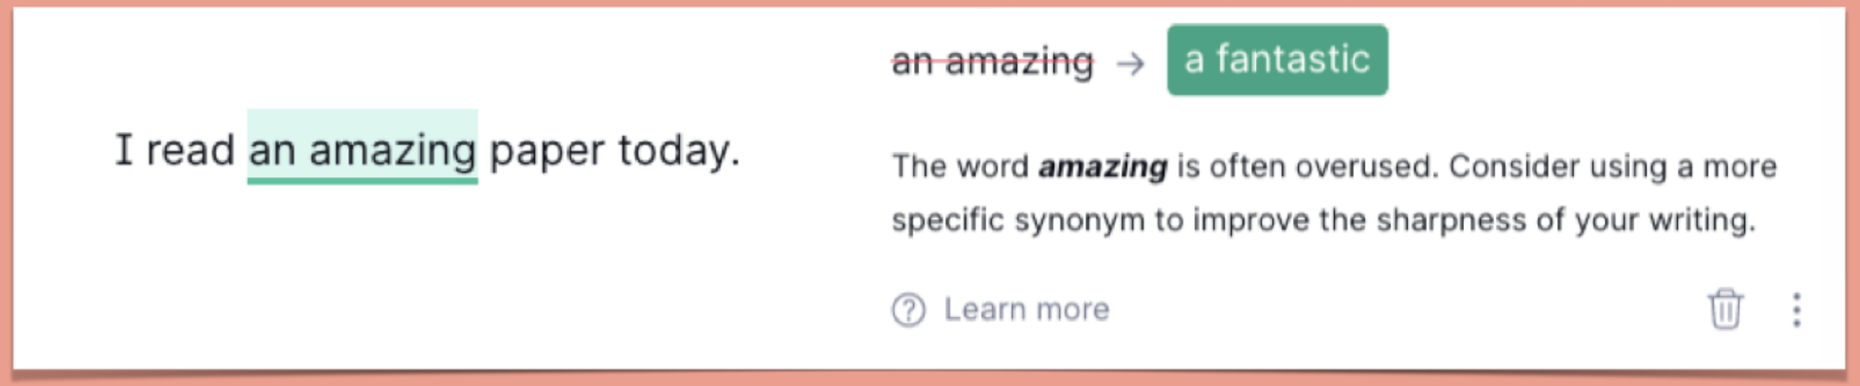

Grammarly (https://www.grammarly.com/)
Grammar Coach (https://www.thesaurus.com/grammarcoach)
Wordtune (https://www.wordtune.com/)

어휘력을 늘리기 위한 다양한 애플리케이션

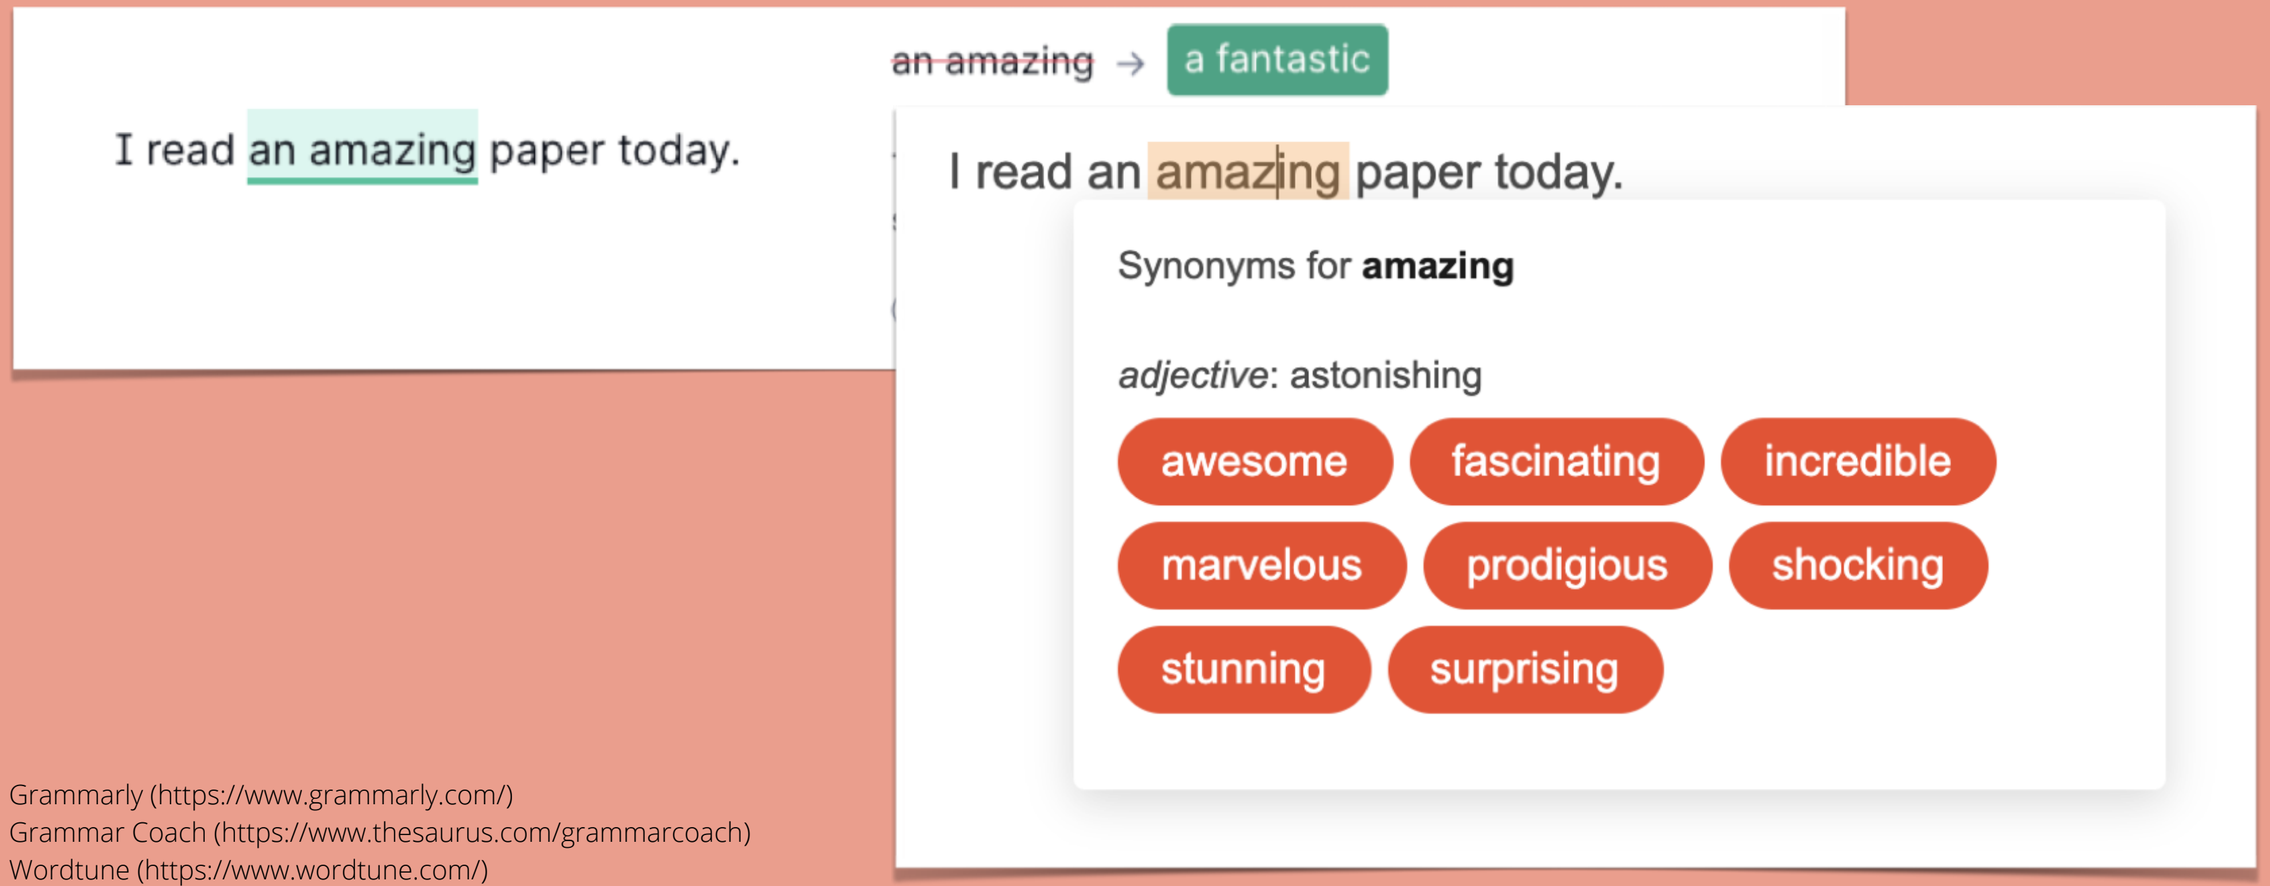

Grammarly (https://www.grammarly.com/)
Grammar Coach (https://www.thesaurus.com/grammarcoach)
Wordtune (https://www.wordtune.com/)

어휘력을 늘리기 위한 다양한 애플리케이션

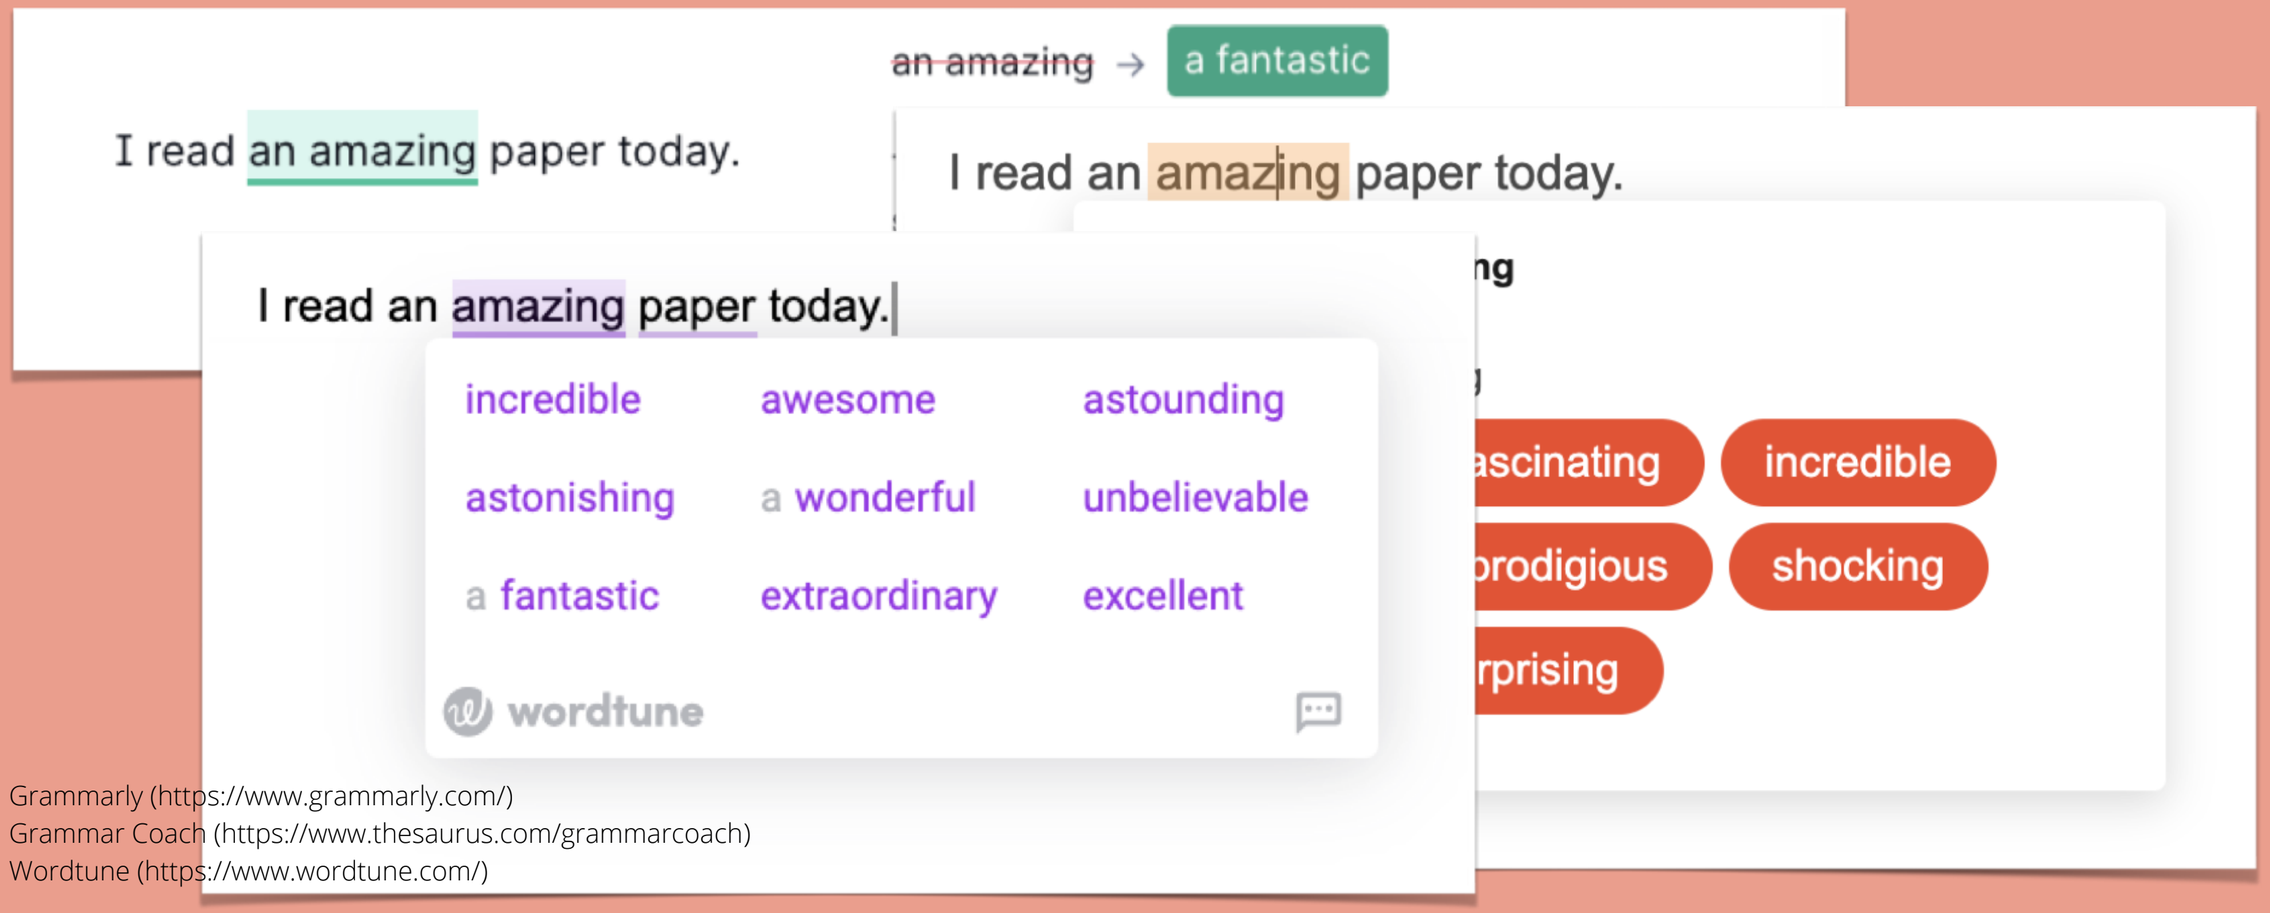

Grammarly (https://www.grammarly.com/)
Grammar Coach (https://www.thesaurus.com/grammarcoach)
Wordtune (https://www.wordtune.com/)

AI로 창의성 증가시키기

Part 2-2

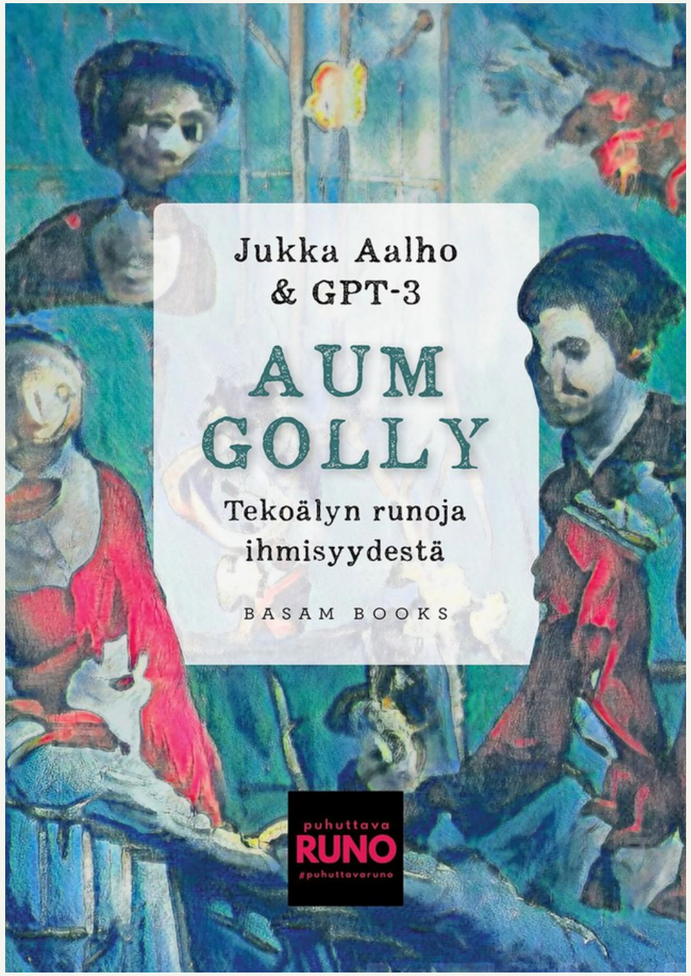

https://aumgolly.fi/english/

I Wrote a Book with GPT-3 AI in 24 Hours — And Got It Published

In early 2021 I signed up for the GPT-3 beta program to see how good it is. A few days later I had co-authored Aum Golly — a book of AI poems on humanity. A few months later it was published. This is what it means for writers and publishers.

Jukka Aalho Follow
Jun 11 · 7 min read ★

https://medium.com/swlh/i-wrote-a-book-with-gpt-3-ai-in-24-hours-and-got-it-published-93cf3c96f120

Aum Golly — Poems on Humanity by an Artificial Intelligence. Photo: Jukka Aalho

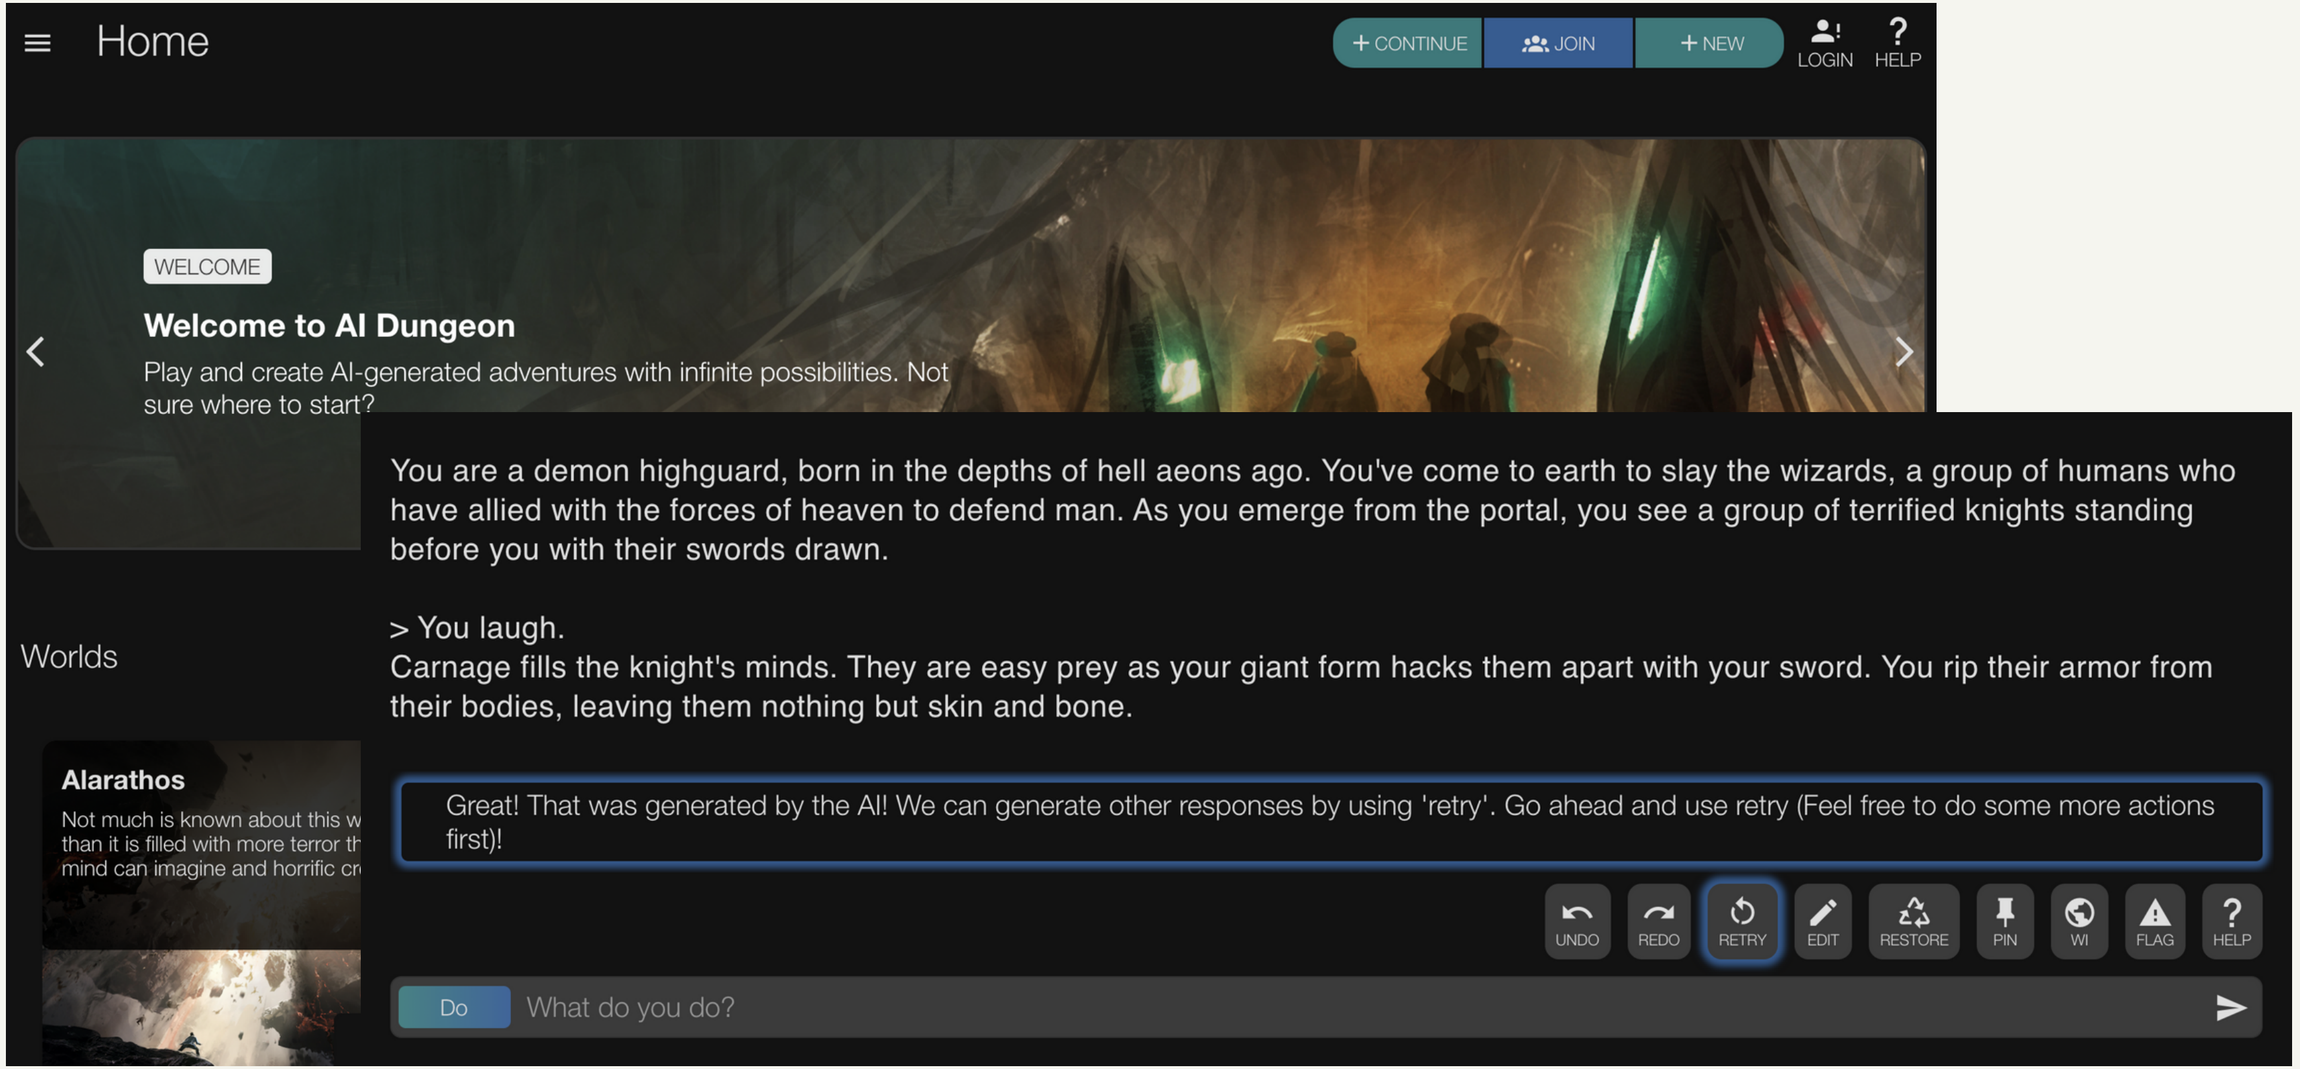

AI Dungeon (https://play.aidungeon.io/main/home)

TEXT PROMPT

an armchair in the shape of an avocado. . . .

AI-GENERATED IMAGES

Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, Ilya Sutskever. Zero-Shot Text-to-Image Generation. CVPR 2021.

TEXT PROMPT

an armchair in the shape of an avocado. . . .

AI-GENERATED IMAGES

TEXT PROMPT

an illustration of a baby daikon radish in a tutu walking a dog

AI-GENERATED IMAGES

Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, Ilya Sutskever. Zero-Shot Text-to-Image Generation. CVPR 2021.

Being creative requires a *jolt* to take us out of the well-carved paths we retrace each day.
This is where a machine might come in: perhaps it could give us that jolt,
throw up a new suggestion, stop us from simply repeating the same algorithm each day.
Machines might ultimately help us, as humans, behave less like machines.

- Marcus Du Sautoy, The Creativity Code: Art and Innovation in the Age of AI

Landscape

Creativity Support Tools (CSTs) play a fundamental role in the study of creativity in Human-Computer Interaction (HCI). Even so, there is **no consensus definition of the term** 'CST' in HCI, and in most studies, CSTs have been construed as one-off exploratory prototypes, typically built by the researchers themselves. This makes it difficult to clearly demarcate CST research, but also to compare findings across studies, which impedes advancement in digital creativity as a growing field of research. Based on a literature review of **143 papers** from the ACM Digital Library (1999-2018), we contribute a first **overview of the key characteristics of CSTs** developed by the HCI community. Moreover, we propose a **tentative definition** of a CST to help strengthen knowledge sharing across CST studies. We end by discussing our study's implications for future HCI research on CSTs and digital creativity.

Category	Codes
Device	- Analogue - Phone/Smartphone/PDA - Tabletop display - Large wall display - Tablet - Laptop/Computer - Tangibles - Method/no device - *Miscellaneous*
Complexity	- Low: One or two features - Medium: Multiple feat., semi-complex system - High: Entire system or suite of tools
Current Availability	- Not available - Readily available online but needs hardware - Online but requires user overhead to set up - Readily available/usable online with no setup
Maturity	- Vision Description/Design Fiction/Scenario - Lo-fi prototype or Mockup (Paper/WoOz) - Hi-fidelity working prototype - Public Release - Longstanding Public Release/widely adopted
Part of Creativity Process	- Pre-ideation/background research - Idea generation or ideation - Evaluation or critique - Implementation - Iteration - Meta or project management
User Group	- Novice - Casual user - Expert - Unspecified - *Intended audience* - *Application domain/activity*
Evaluation	- No evaluation - Hypothesis-driven - Exploratory - *Number of participants* - Student evaluators - Qualitative - Quantitative - Mixed methods - *Method* - *Criterion (what was evaluated for)*
Collab-oration	- Individual use - Collaborative use (with people) - Collaborate with software/algorithm

Jonas Frich, Lindsay MacDonald Vermeulen, Christian Remy, Michael Mose Biskjaer, Peter Dalsgaard. Mapping the Landscape of Creativity Support Tools in HCI. CHI 2019.

Landscape

Creativity Support Tools (CSTs) play a fundamental role in the study of creativity in Human-Computer Interaction (HCI). Even so, there is **no consensus definition of the term** ‘CST’ in HCI, and in most studies, CSTs have been construed as one-off exploratory prototypes, typically built by the researchers themselves. This makes it difficult to clearly demarcate CST research, but also to compare findings across studies, which impedes advancement in digital creativity as a growing field of research. Based on a literature review of **143 papers** from the ACM Digital Library (1999-2018), we contribute a first **overview of the key characteristics of CSTs** developed by the HCI community. Moreover, we propose a **tentative definition** of a CST to help strengthen knowledge sharing across CST studies. We end by discussing our study’s implications for future HCI research on CSTs and digital creativity.

모두가 동의하는 정의가 없음

관심은 많음

Jonas Frich, Lindsay MacDonald Vermeulen, Christian Remy, Michael Mose Biskjaer, Peter Dalsgaard. Mapping the Landscape of Creativity Support Tools in HCI. CHI 2019.

Landscape

Creativity Support Tools (CSTs) play a fundamental role in the study of creativity in Human-Computer Interaction (HCI). Even so, there is **no consensus definition of the term** 'CST' in HCI, and in most studies, CSTs have been construed as one-off exploratory prototypes, typically built by the researchers themselves. This makes it difficult to clearly demarcate CST research, but also to compare findings across studies, which impedes advancement in digital creativity as a growing field of research. Based on a literature review of **143 papers** from the ACM Digital Library (1999-2018), we contribute a first **overview of the key characteristics of CSTs** developed by the HCI community. Moreover, we propose a **tentative definition** of a CST to help strengthen knowledge sharing across CST studies. We end by discussing our study's implications for future HCI research on CSTs and digital creativity.

여전히 모호한 정의

"A Creativity Support Tool (CST) runs on one or more digital systems,
encompasses one or more creativity-focused features,
and is employed to positively influence users of varying expertise
in one or more distinct phases of the creative process."

Jonas Frich, Lindsay MacDonald Vermeulen, Christian Remy, Michael Mose Biskjaer, Peter Dalsgaard. Mapping the Landscape of Creativity Support Tools in HCI. CHI 2019.

Challenges

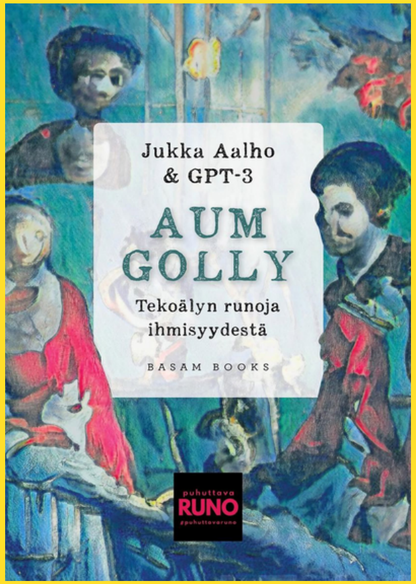

- Many applications demonstrated the potential
- Hard to reason about “creativity”
  - What do humans want from creativity support tools?
  - How do humans evaluate creativity support tools?
- Maybe what matters is how people "feel" during the process
  - How do we measure "enjoyment" and "engagement"?

Anna Kantosalo. Human-Computer Co-Creativity -- Designing, Evaluating and Modeling Computational Collaborators for Poetry Writing. PhD Dissertation 2019.

And AI introduces *even more* challenges

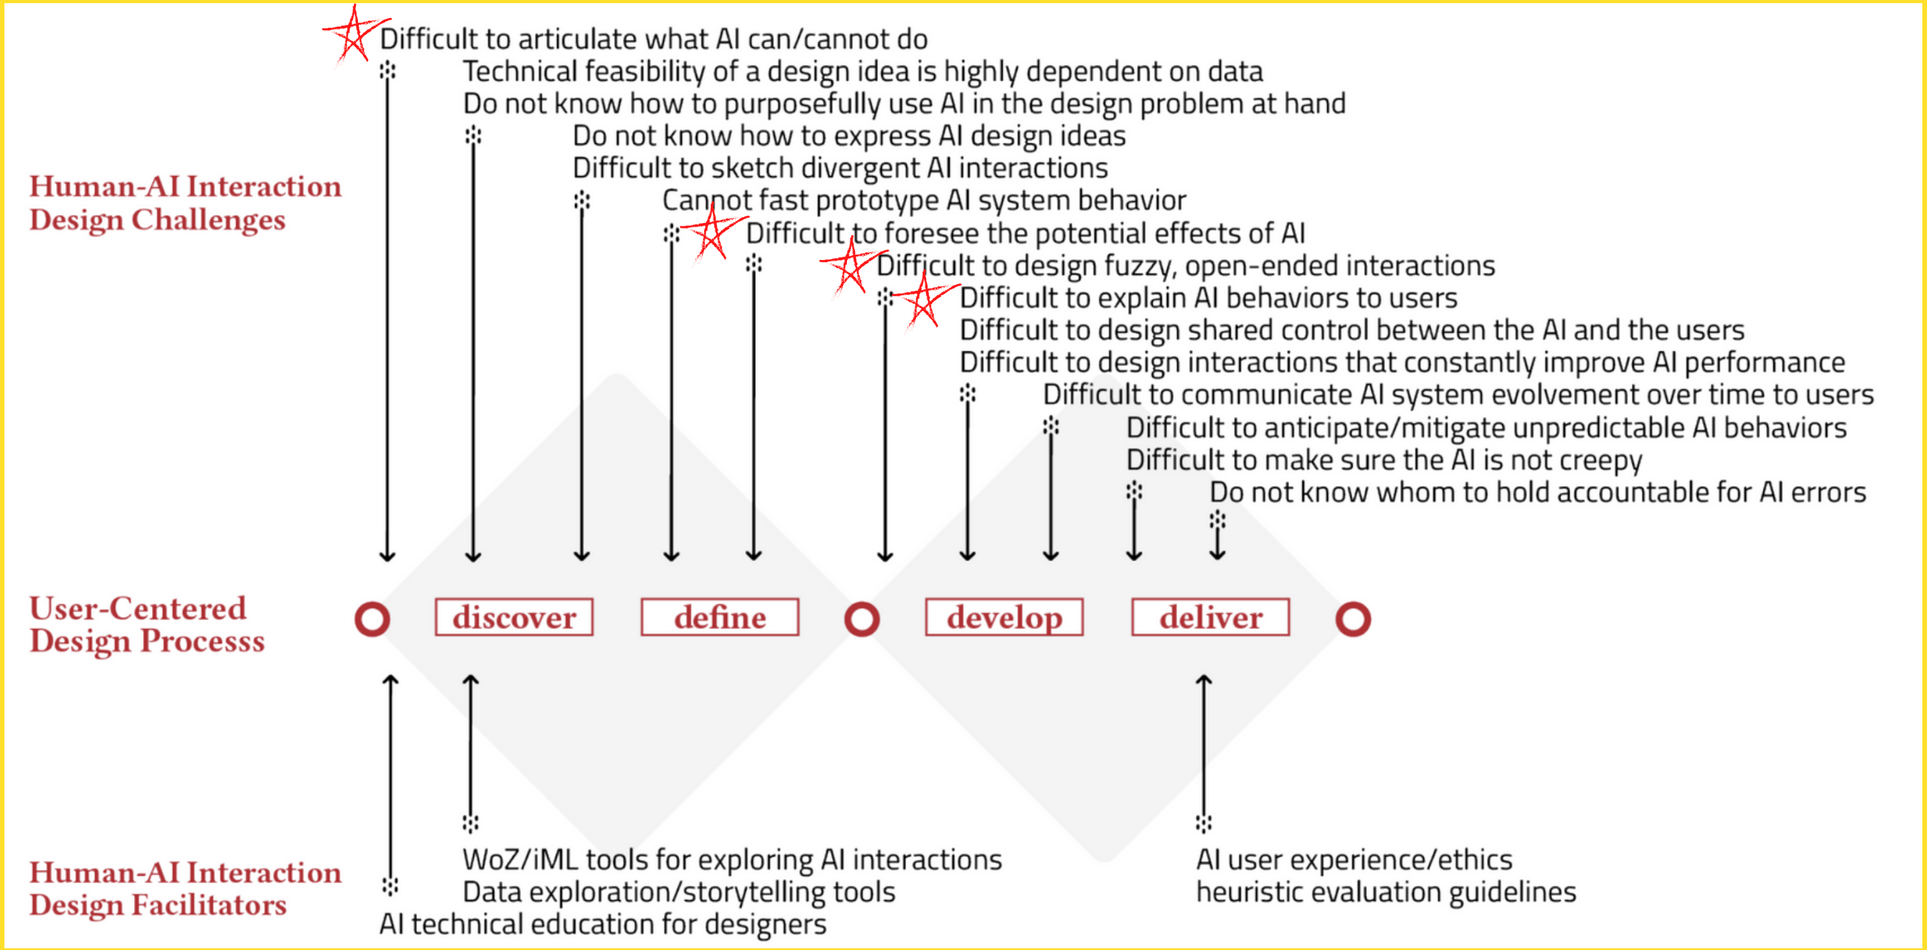

Qian Yang, Aaron Steinfeld, Carolyn Rose, John Zimmerman. Re-examining Whether, Why, and How Human-AI Interaction Is Uniquely Difficult to Design. CHI 2020.

사람과 AI가 함께 글을 쓰는 과정

How do humans and language models write together?

Mina Lee
Stanford University

Qian Yang
Cornell University

Percy Liang
Stanford University

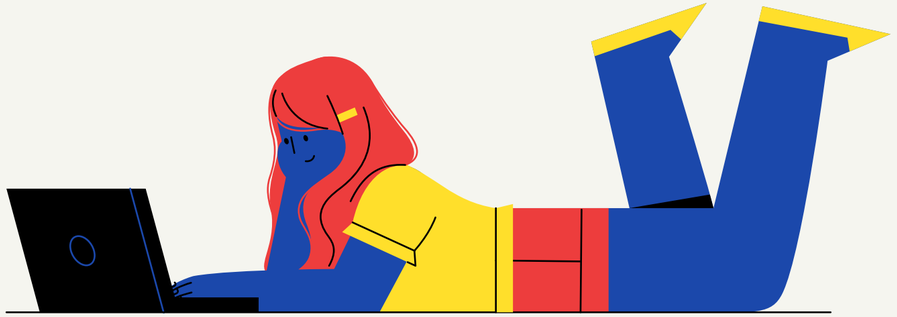

많은 사람들이 AI와 함께 자유롭게 글을 쓰는 과정을 관찰하자.

아기돼지 92형제

Once upon a time there was an old mother pig who had one hundred little pigs and not enough food to feed them. So when they were old enough, she sent them out into the world to seek their fortunes. You know the story about the first three little pigs. This is a story about the 92nd little pig.

The 92nd little pig built a house out of depleted uranium. And the wolf was like, "dude."

Normal

Once upon a time there was an old mother pig who had one hundred little pigs and not enough food to feed them. So when they were old enough, she sent them out into the world to seek their fortunes. You know the story about the first three little pigs. This is a story about the 92nd little pig.

The 92nd little pig built a house out of depleted uranium. And the wolf was like, "dude." "

우리가 알고 싶은 것

What are the **strengths and weaknesses** of humans and AI in collaborative writing?

What are emergent **interaction patterns**?

Does human-AI collaborative writing have **known benefits** of human-human collaborative writing?

Does AI improve the **quality of writing**?

Can AI reduce **writing reluctance**?

For those who choose to use artificial intelligence, it will alter the task of writing. "**The writer's job becomes as an *editor* almost**," Gupta said. "Your role starts to become deciding what's good and executing on your taste, not as much the low-level work of pumping out word by word. You're still editing lines and copy and making those words beautiful, but, as you move up in that chain, and you're executing your taste, you have the potential to do a lot more."

- Stephen Marche, The Computers Are Getting Better at Writing (The New Yorker)

AI 사용의 기회와 문제점

Part 3

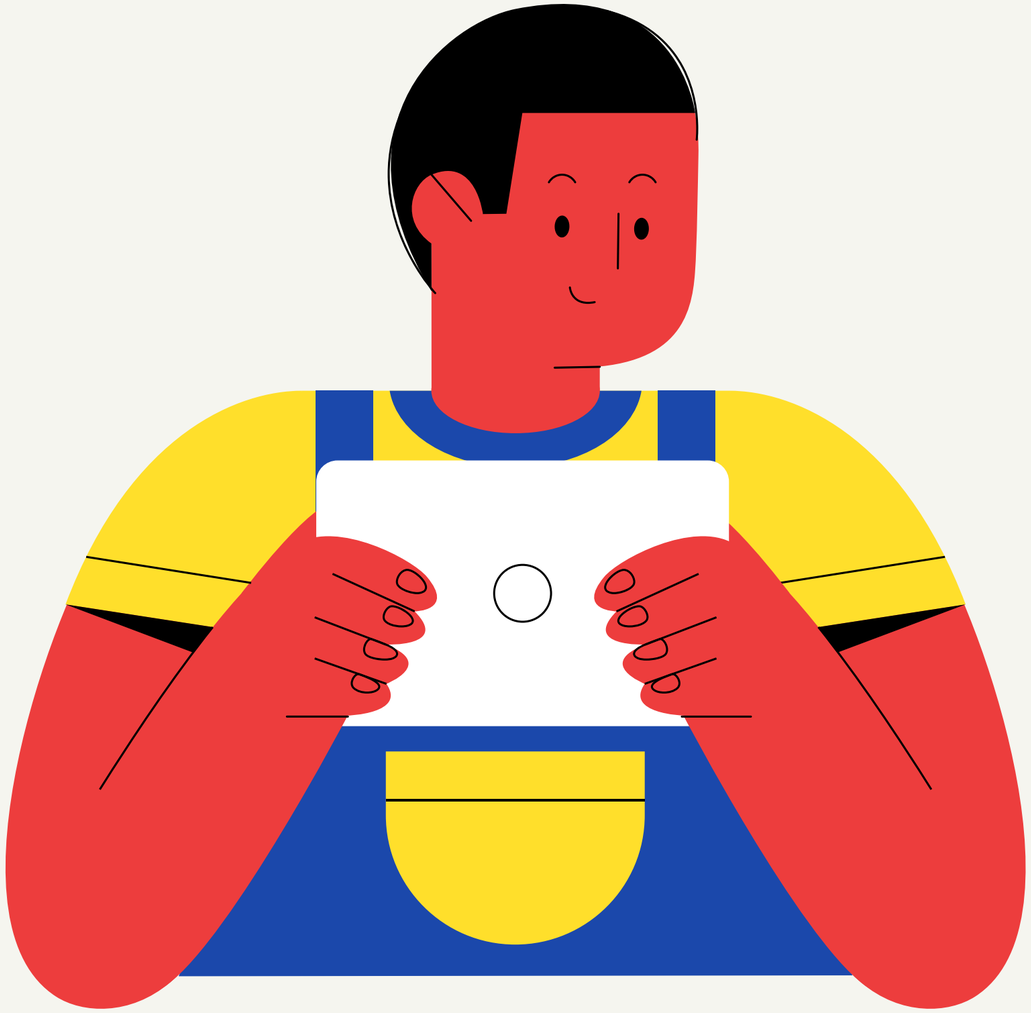

AI-assisted writing's

Economic impact

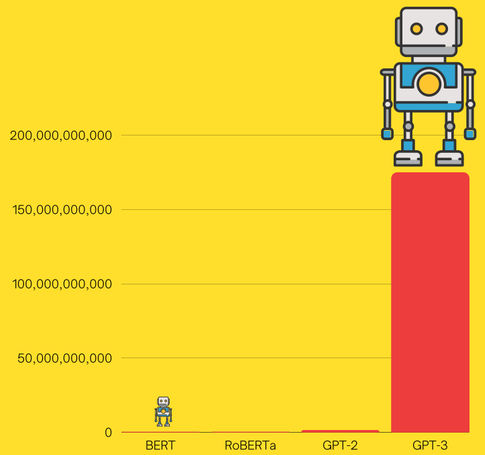

Automation? Augmentation!

When a technology substitutes for human labor in completing tasks, it tends to reduce demand for the workers doing those tasks. This depresses employment and wages. However, when a technology complements labor, or facilitates the creation of new opportunities or tasks, it tends to increase labor demand.

Productivity and innovation

Large language models are likely to substantially increase productivity and innovation by reducing time for work and increasing the amount/quality of work per unit time, thereby raising living standards.

Centralization

The ownership and control of large language models and the associated data can lead to increased concentration, market power and winner-take-all effects. However, open, distributed and/or interoperable models could lead to decentralization of economic power and innovation.

AI를 통한 글쓰기가 일에 미치는 영향

How does work product of human-AI collaboration differ from human-only and AI-only scenarios?

Mina Lee
Stanford University

Zanele Munyikwa
MIT

Percy Liang
Stanford University

Erik Brynjolfsson
Stanford University

AI를 통한 글쓰기가 일에 미치는 영향

Human-only

AI를 통한 글쓰기가 일에 미치는 영향

Human-only

Human-AI

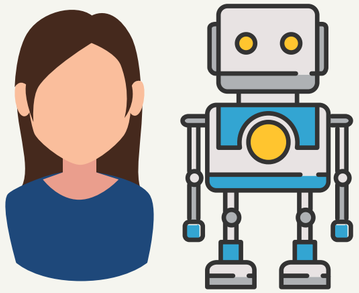

AI를 통한 글쓰기가 일에 미치는 영향

Human-only

Human-AI

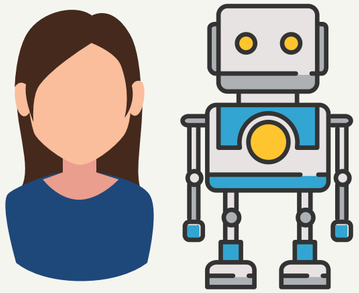

AI-only

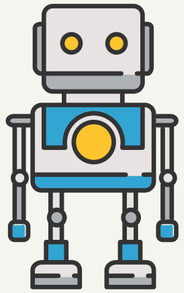

AI를 통한 글쓰기가 일에 미치는 영향

Human-only

초보자

전문가

Human-AI

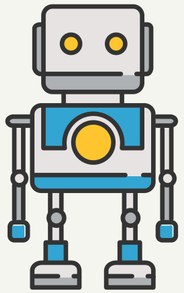

초보자

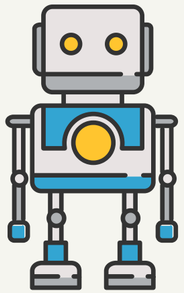

전문가

AI-only

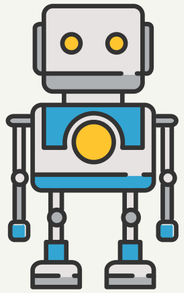

AI-assisted writing's

Opportunities

Low entry barrier

Easier and faster to prototype and build AI-infused systems

New forms of interaction

There will be important new forms of interaction that that will arise due to the powerful generative capacities of language models (e.g. autocomplete systems, creativity support tools), but also due to the potential multimodal models (e.g. combining text and images) that will more naturally combine different modes of interaction.

Blur the line between users and developers

important opportunities to apply participatory, or value-sensitive design processes much more directly by empowering individuals and communities to build bespoke models that more closely align with their needs and values.

AI-assisted writing's

Risks

Toxic content

Pre-trained neural language models are prone to generating racist, sexist, or otherwise toxic language which hinders their safe deployment

Responsibility attribution

AI can have agency but do not meet traditional criteria for moral agency and moral responsibility

Misuse

Fake news, fake profile, abuse, plagiarism

Perception & trust

Reshaping people’s trust, credibility, and identity, shifting responsibilities and ownership of contents

Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, Noah A. Smith. RealToxicityPrompts: Evaluating Neural Toxic Degeneration in Language Models. Findings of EMNLP 2020.

Mark Coeckelbergh. Artificial Intelligence, Responsibility Attribution, and a Relational Justification of Explainability. Science and Engineering Ethics 2020.

Gabriel Lima, Nina Grgić-Hlača, Meeyoung Cha. Human Perceptions on Moral Responsibility of AI: A Case Study in AI-Assisted Bail Decision-Making. CHI 2021.

Maurice Jakesch, Megan French, Xiao Ma, Jeffrey T. Hancock, Mor Naaman. AI-Mediated Communication: How the Perception that Profile Text was Written by AI Affects Trustworthiness. CHI 2019.

Workshop on Foundation Models

AUGUST 23-24, 2021

VIRTUAL EVENT

RSVP to join the workshop

https://crfm-stanford.github.io/workshop.html

We invite you to join us at this workshop, where scholars reflecting a diverse array of **perspectives**, disciplinary **backgrounds** (e.g. social science, economics, computer science, law, philosophy, information science) and **sectors** (academia and industry) will convene to provide vital expertise on the many dimensions of foundation models.
Broadly, we will address the **opportunities, challenges, limitations, and societal impact** of foundation models.

Thank you!

Writing with AI

이민아

Stanford University

연구 관련 이야기, 진로 상담, 질문 및 고민 등 편하게 이메일 주세요!

minalee@cs.stanford.edu