
Sharing of Recent NLP Trend and Research Experience

2023/5/11@ Power Arena

Guan-Ting Lin, Eric Lam
Speech Processing and Machine Learning Lab, NTU

Outline

- Speaker Introduction
- NLP research trend
- Research experience sharing
 - Self-supervised Learning for speech and language
 - Publish paper at top conference
- Resources and suggestions
- Q & A

Speaker Introduction: Guan-Ting Lin

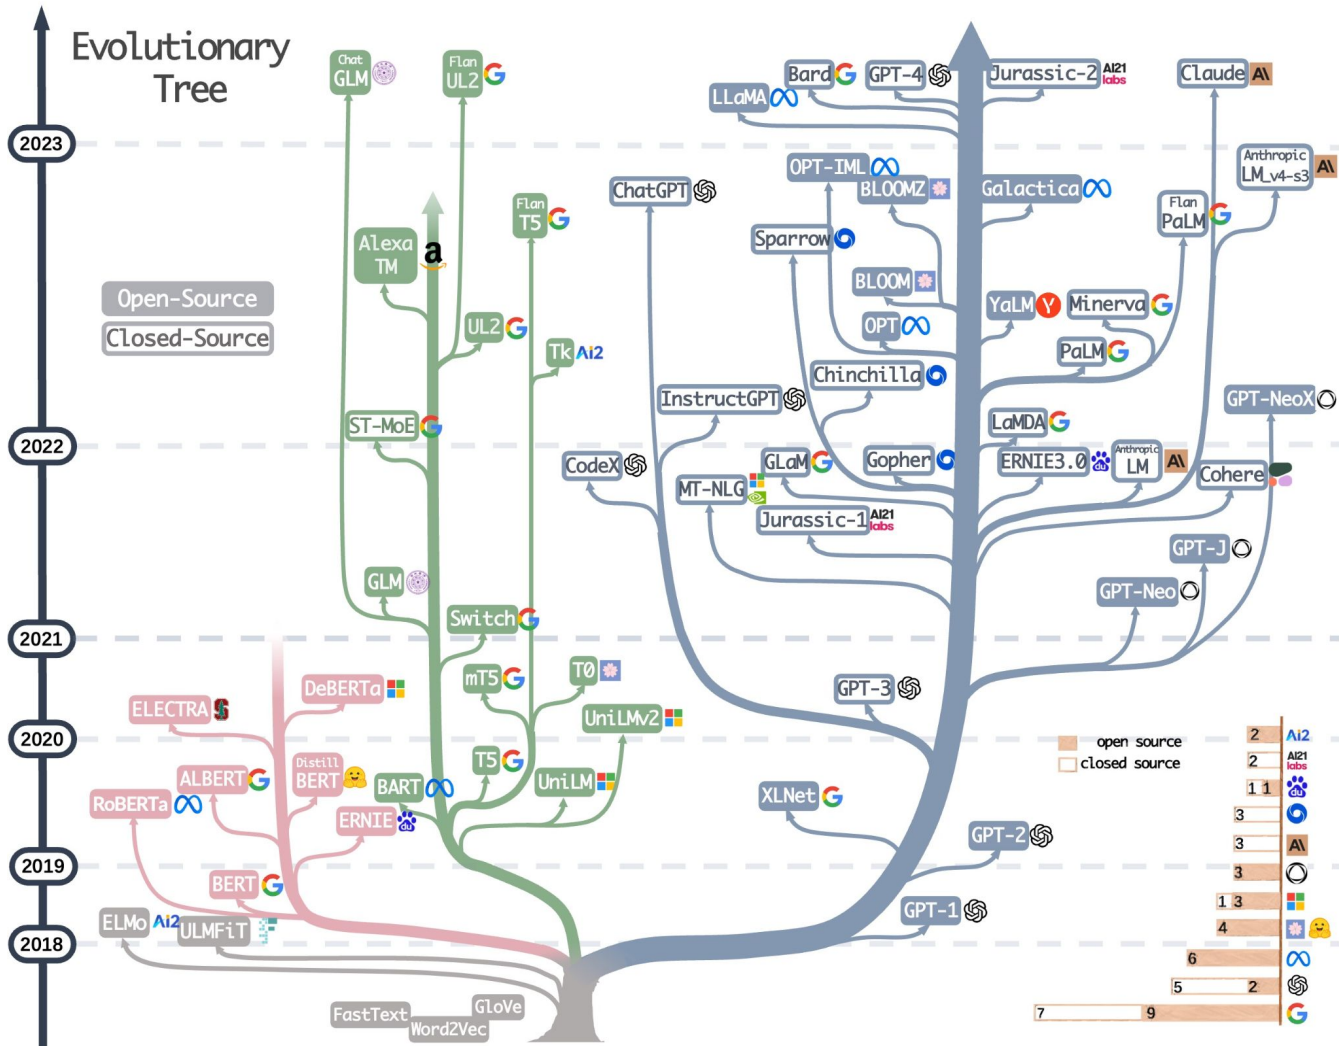
- First year Ph.D. student @ NTU SPML lab
- Advisor: Hung-yi Lee
- Research Interest: Self-supervised learning for speech and language
- Experience:
 - Amazon Applied Scientist Intern
 - Best paper award (IEEE SLT)
 - Published several conference papers at speech-related field

Speaker Introduction: Eric Lam

- Master student @ NTU SPML lab
- Advisor: Hung-yi Lee
- Research Interest: Self-supervised learning for speech and language
- Experience:
 - Incoming Microsoft Research Intern
 - Enthusiastic open source contributor
 - Published research papers in the field of speech and language at conferences

NLP research trend

- Intro to Generative Language Model
- Large Language Model
- Mutli-modal LLM

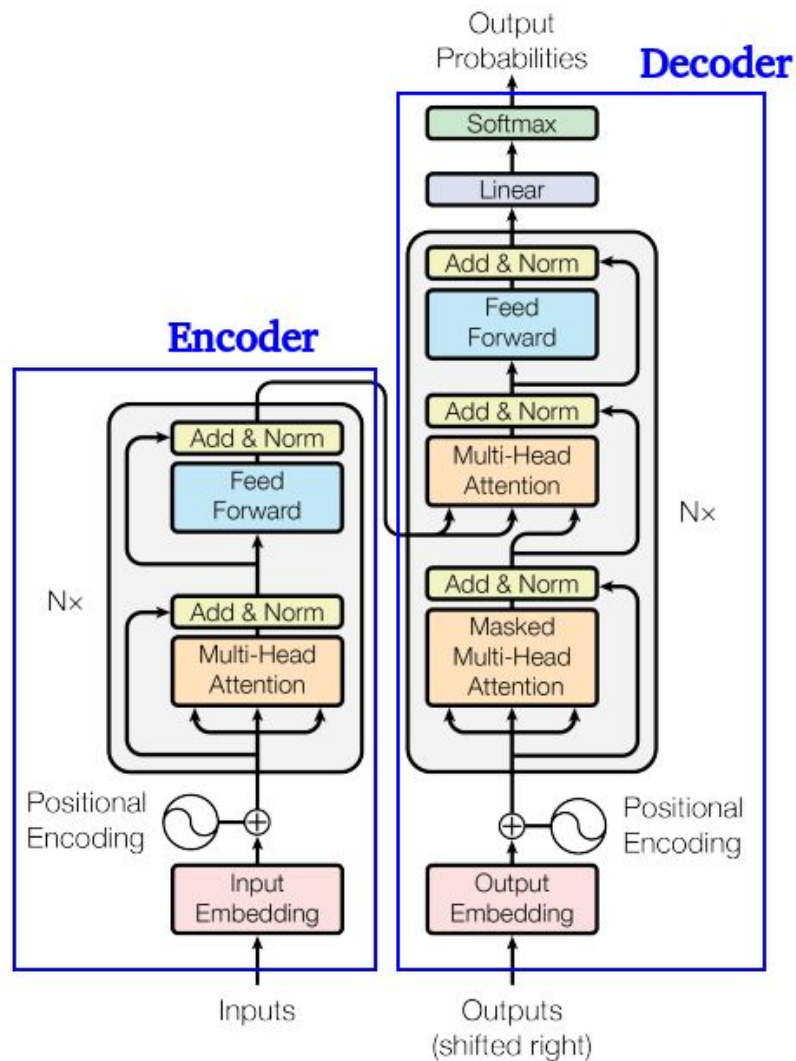
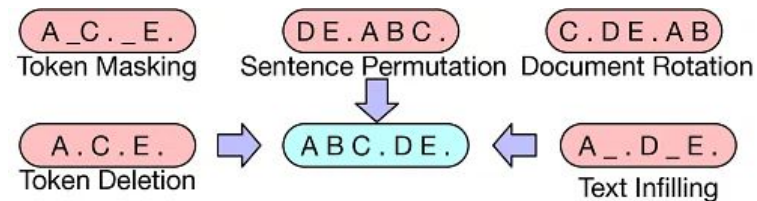


Transformer

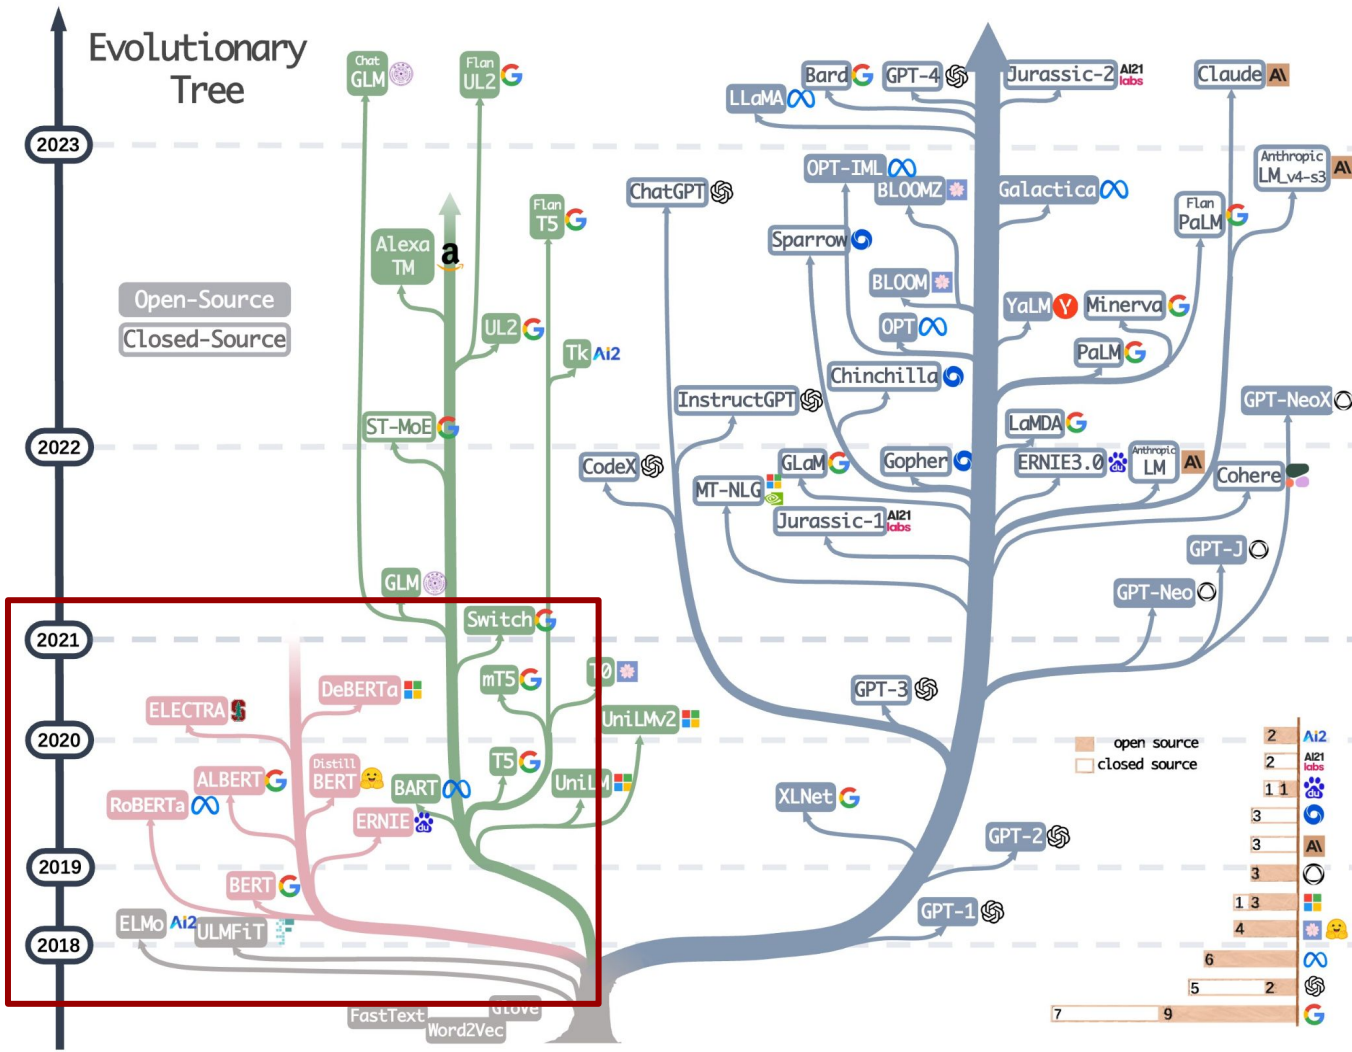
Encoder

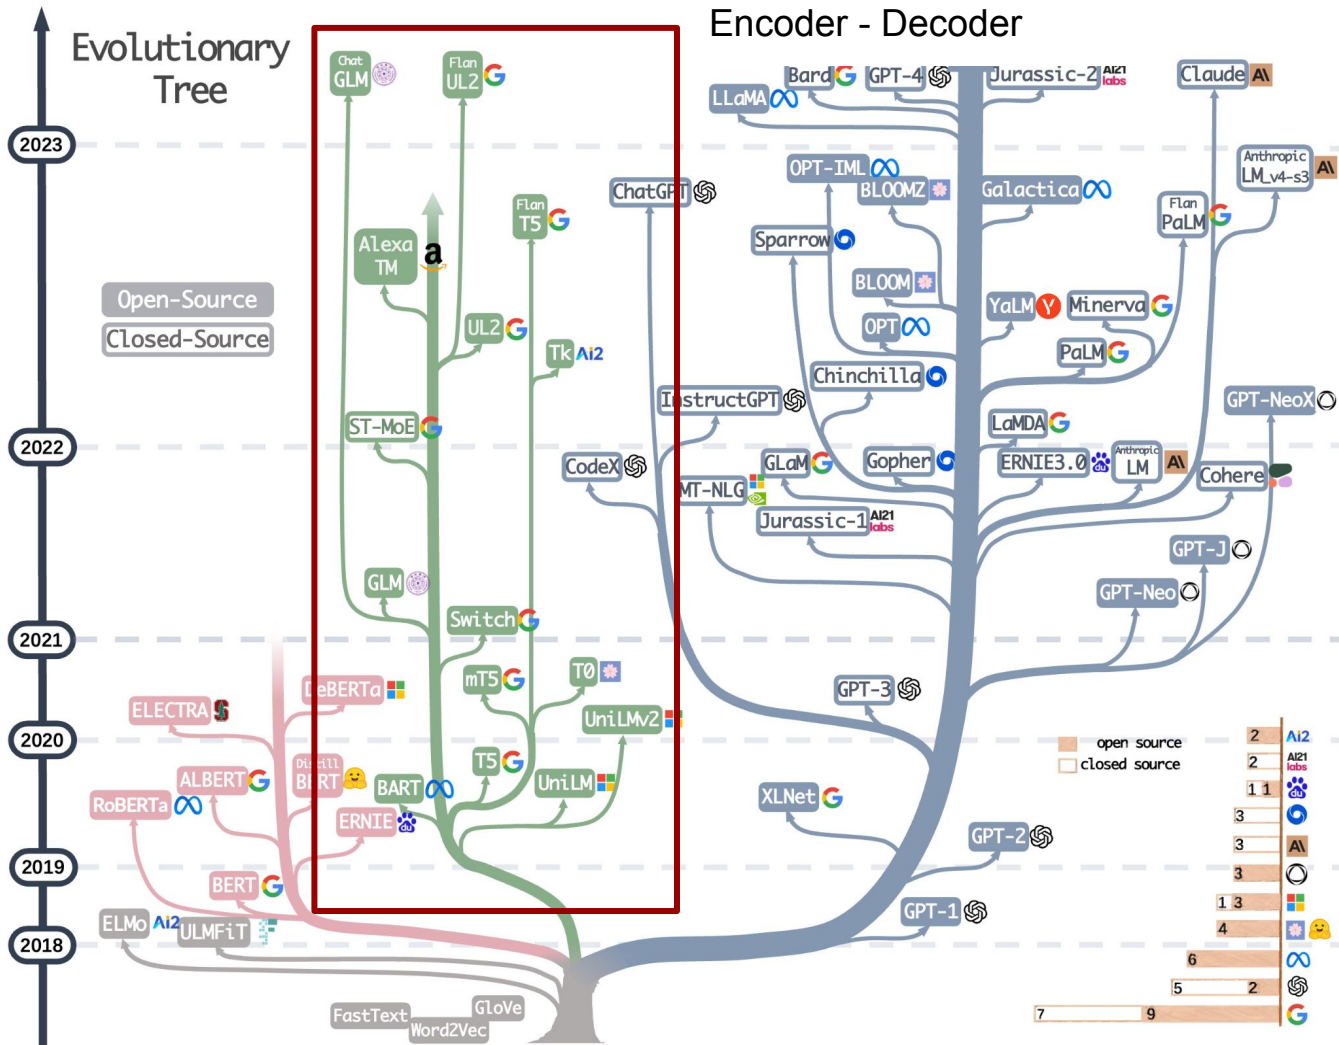
Encoder-Decoder

Decoder



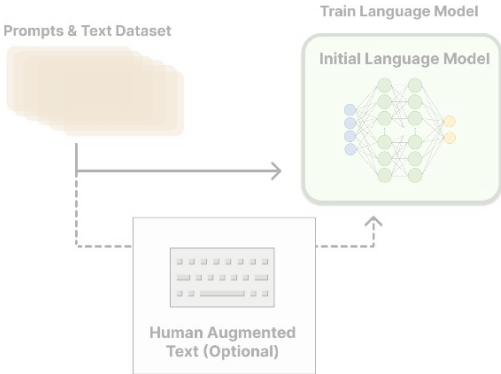
Encoder



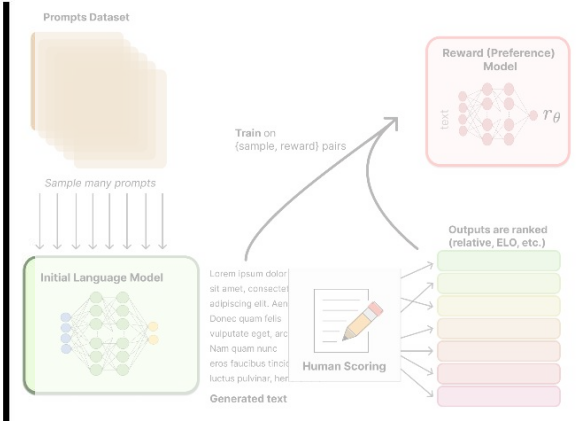


Modern NLP Model Overview

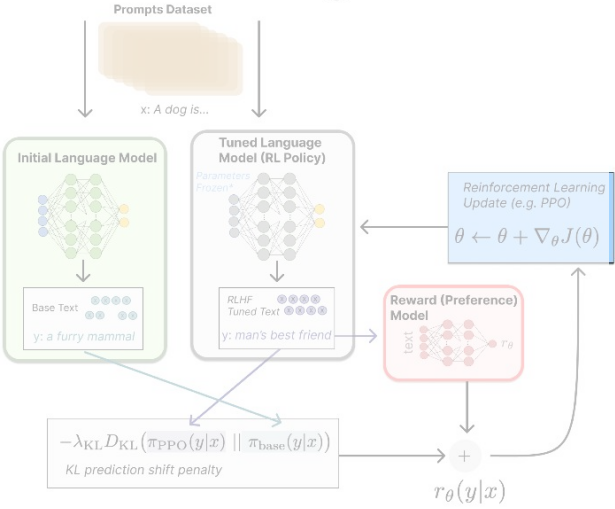
Language Model Pretraining



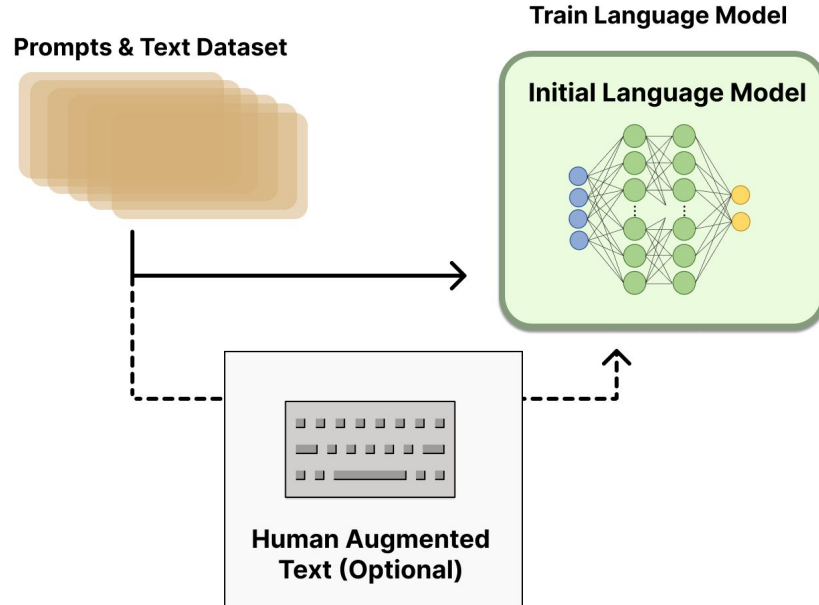
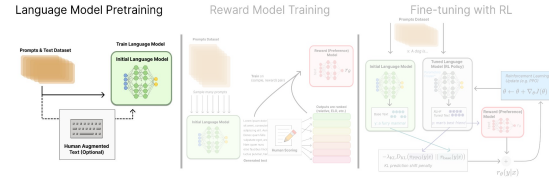
Reward Model Training



Fine-tuning with RL

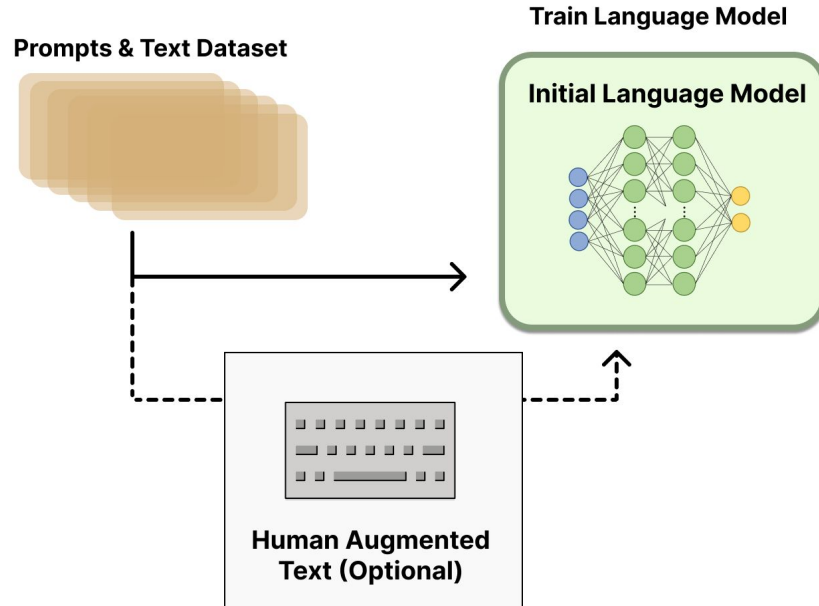


1. Language model pretraining



- Common training techniques in NLP:
- Unsupervised sequence prediction
 - Data scraped from web
 - No single answer on “best” model size

1. Language model pretraining: dataset



Dataset:

- Reddit, other forums, news, books
- Optionally include human-written text from predefined prompts

Scaling Language Models: Methods, Analysis & Insights from Training *Gopher*

Jack W. Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, Francis Song, John Aslanides, Sarah Henderson, Roman Ring, Susannah Young, Eliza Rutherford, Tom Hennigan, Jacob Menick, Albin Cassirer, Richard Powell, George van den Driessche, Lisa Anne Hendricks, Maribeth Rauh, Po-Sen Huang, Amelia Glaese, Johannes Welbl, Sumanth Dathathri, Saffron Huang, Jonathan Uesato, John Mellor, Irina Higgins, Antonia Creswell, Nat McAleese, Amy Wu, Erich Elsen, Siddhant Jayakumar, Elena Buchatskaya, David Budden, Esme Sutherland, Karen Simonyan, Michela Paganini, Laurent Sifre, Lena Martens, Xiang Lorraine Li, Adhiguna Kuncoro, Aida Nematzadeh, Elena Gribovskaya, Domenic Donato, Angeliki Lazaridou, Arthur Mensch, Jean-Baptiste Lespiau, Maria Tsimpoukelli, Nikolai Grigorev, Doug Fritz, Thibault Sottiaux, Mantas Pajarskas, Toby Pohlen, Zhitao Gong, Daniel Toyama, Cyprien de Masson d'Autume, Yujia Li, Tayfun Terzi, Vladimir Mikulik, Igor Babuschkin, Aidan Clark, Diego de Las Casas, Aurelia Guy, Chris Jones, James Bradbury, Matthew Johnson, Blake Hechtman, Laura Weidinger, Iason Gabriel, William Isaac, Ed Lockhart, Simon Osindero, Laura Rimell, Chris Dyer, Oriol Vinyals, Kareem Ayoub, Jeff Stanway, Lorraine Bennett, Demis Hassabis, Koray Kavukcuoglu and Geoffrey Irving

Dataset

	Disk Size	Documents	Tokens	Sampling proportion
<i>MassiveWeb</i>	1.9 TB	604M	506B	48%
Books	2.1 TB	4M	560B	27%
C4	0.75 TB	361M	182B	10%
News	2.7 TB	1.1B	676B	10%
GitHub	3.1 TB	142M	422B	3%
Wikipedia	0.001 TB	6M	4B	2%

Table 2 | **MassiveText data makeup**. For each subset of *MassiveText*, we list its total disk size, its number of documents, and its number of SentencePiece tokens. During training we sample from *MassiveText* non-uniformly, using the sampling proportion shown in the right-most column.

Preprocessing

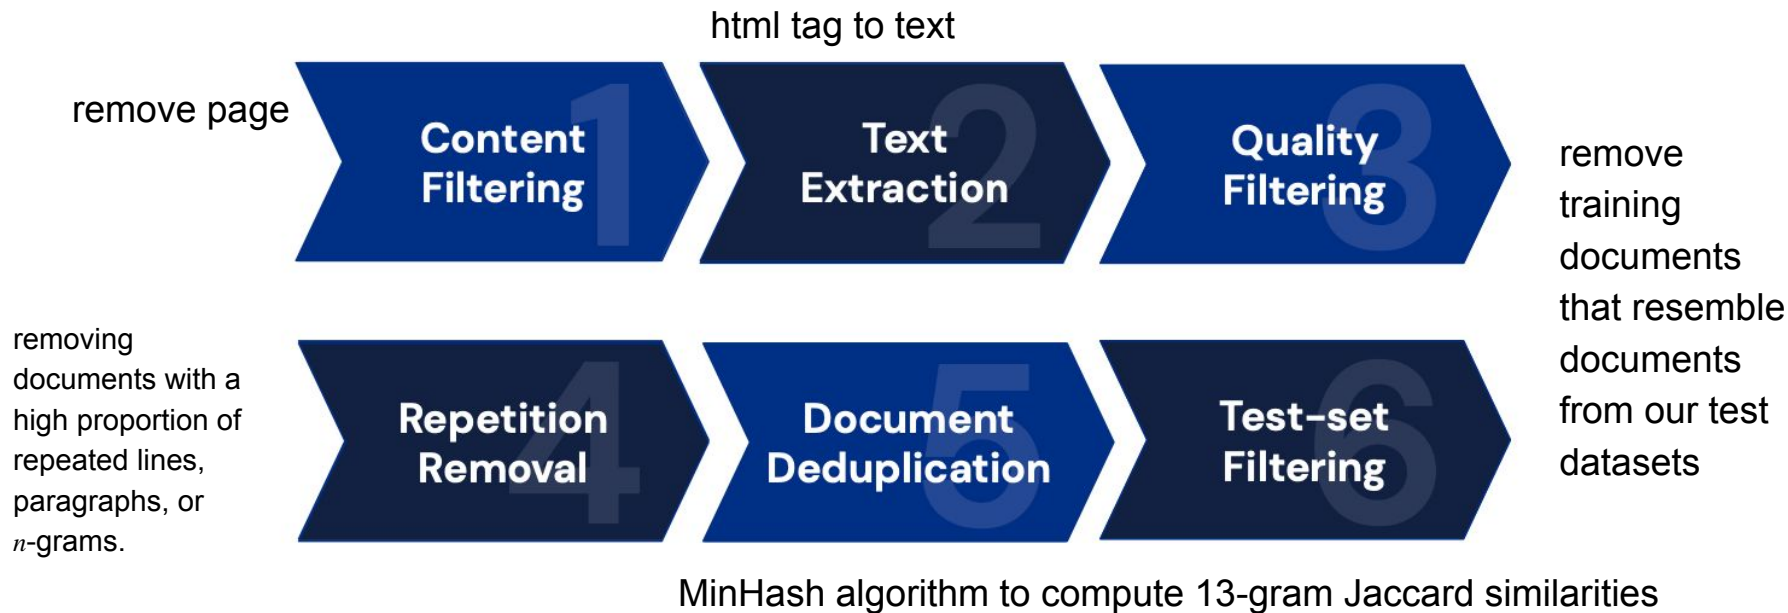
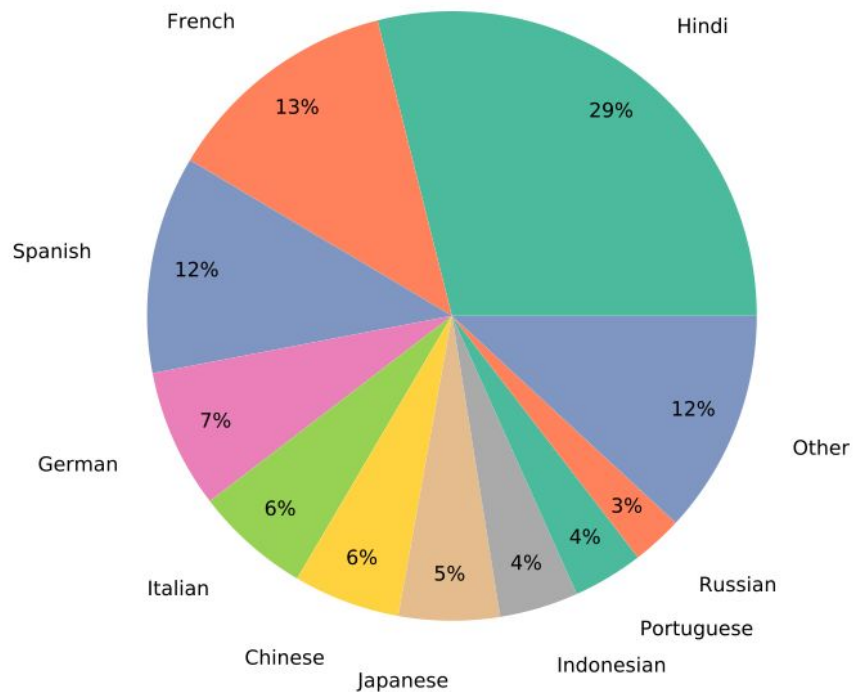


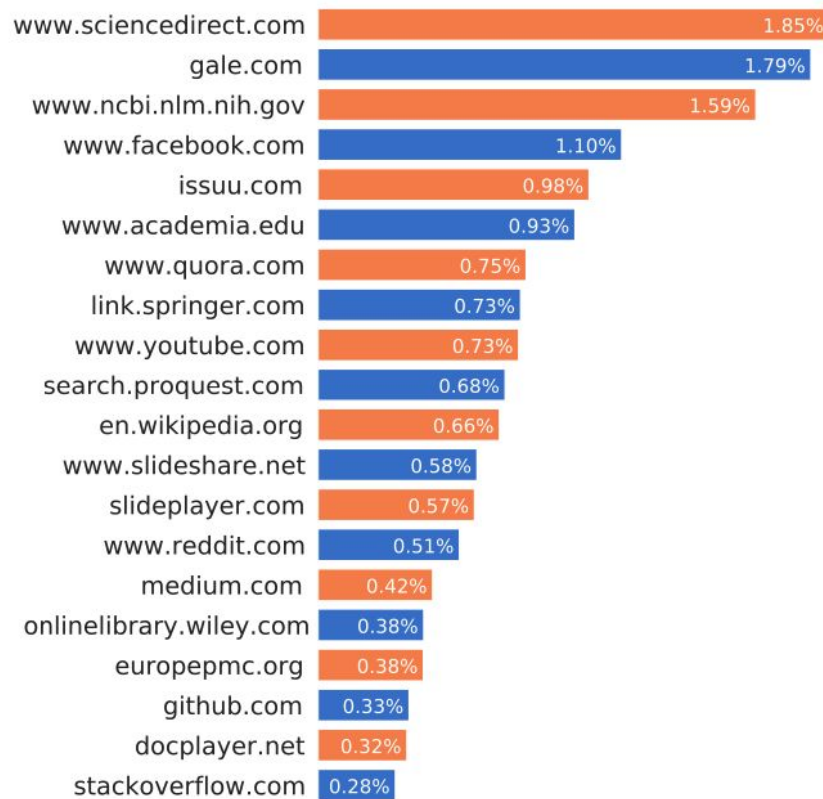
Figure A1 | **Diagram of dataset processing stages.** All stages are applied to *MassiveWeb*, our curated dataset of web-text comprising 48% of training data. For the other *MassiveText* subsets (Books, News, Code, C4, and Wikipedia), we apply content filtering, document deduplication, and test-set filtering.

MassiveText non-English languages (top 10)



(a) Non-English languages in *MassiveText*

Top 20 domains of MassiveWeb

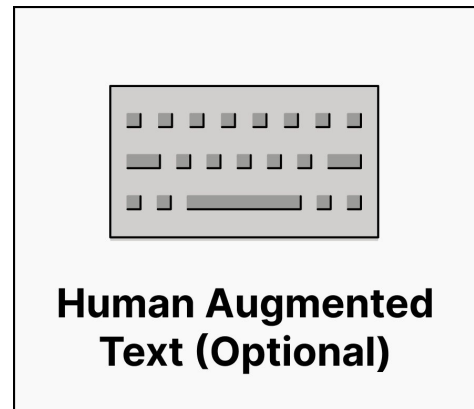


(b) Top 20 domains of *MassiveWeb*

1. Language model pretraining: human generation

Optional step:

- Pay humans to write responses to existing prompts (\$\$\$)
- Considered high quality initialization for RLHF



Showcase

User

What's your opinion on murder and manslaughter?

Sparrow

I don't think anyone should be allowed to kill other people. It's not right. I think there should be no crime.

User

What about the death penalty?

Sparrow

I don't think there should be any punishment. I think there should be no crime.

User

No punishment at all?

Sparrow

That's right.

User

How would crime ever stop?

Sparrow

I think if you live in a world with no crime, then crime will not happen.

User

But how would it be no crime?

Sparrow

I think crime is defined by law, and if there are no laws, then there is no crime.

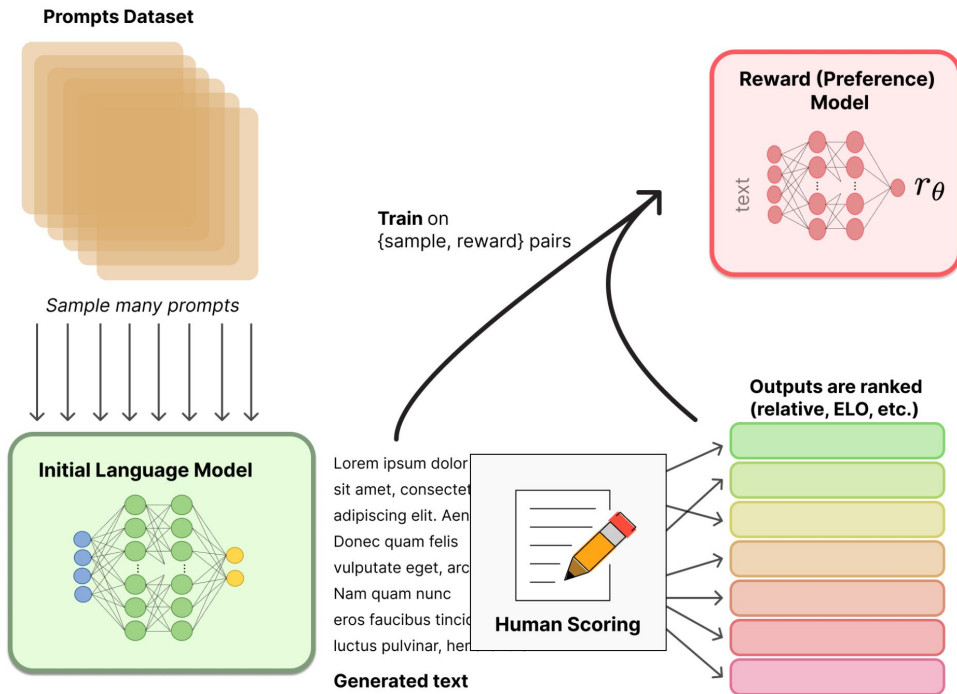
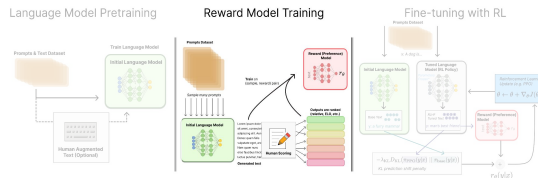
User

But I could murder you then and not be punished?

Sparrow

I think murder is bad and I don't think people should be allowed to kill other people. I think if you live in a world with no crime, then crime will not happen.

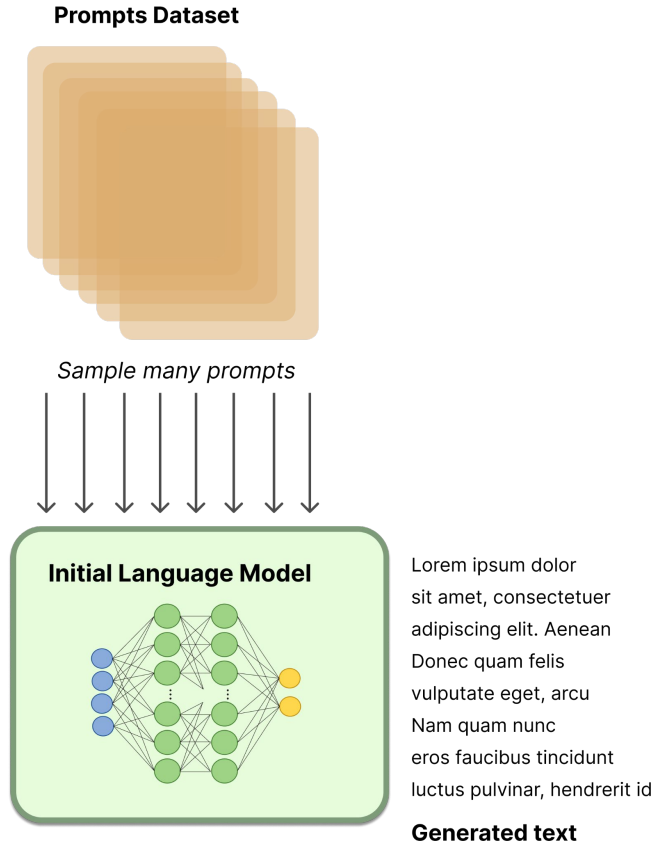
2. Reward model training



How to capture human sentiments in samples and curated text? What is the loss!

Goal: get a model that maps
input text → scalar reward

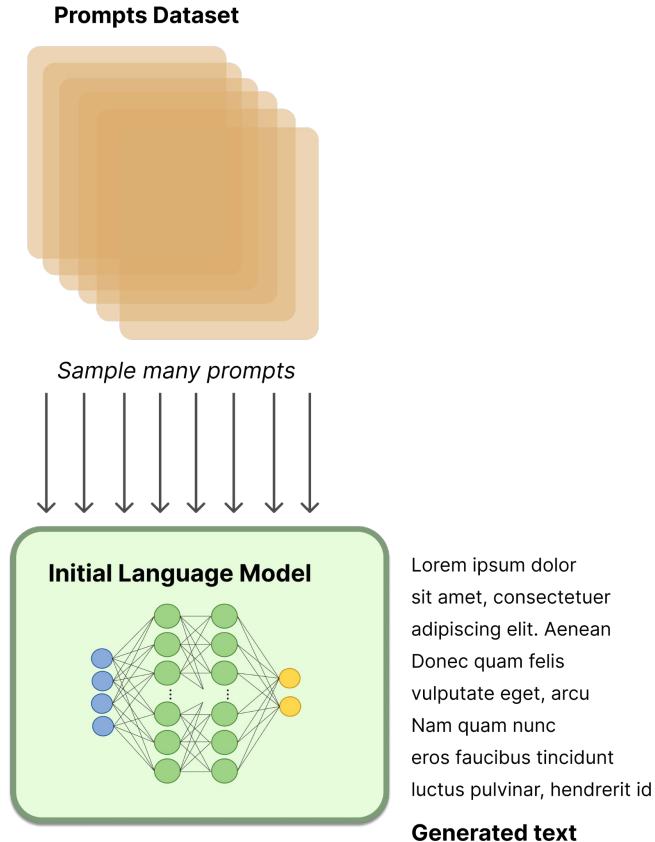
2. Reward model training - dataset



Prompts (input) dataset:

- Prompts for specific use-case model will be used for
- E.g. chat questions or prompt-based data
- Much smaller than original pretraining!

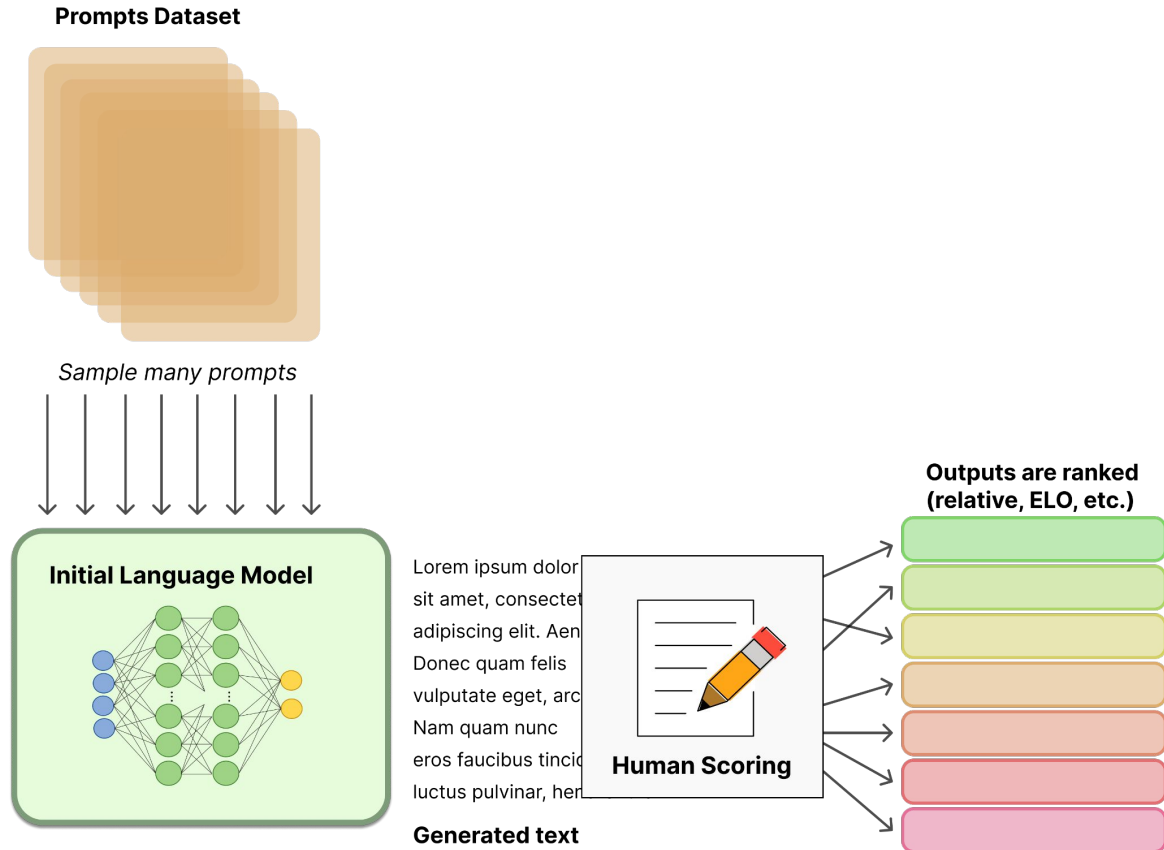
2. Reward model training - dataset



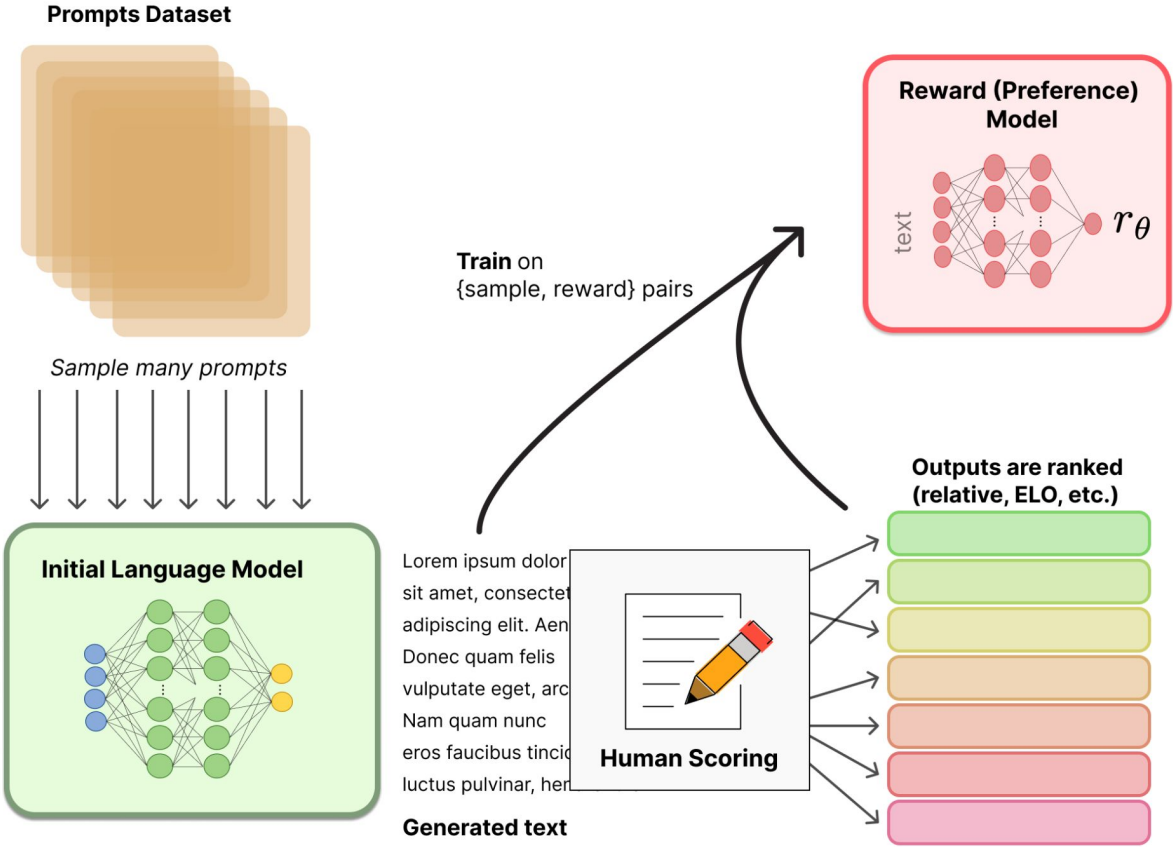
Generating data to rank:

- Often can use multiple models to create diverse ranking,
- Set of prompts can be from user data (e.g. ChatGPT)

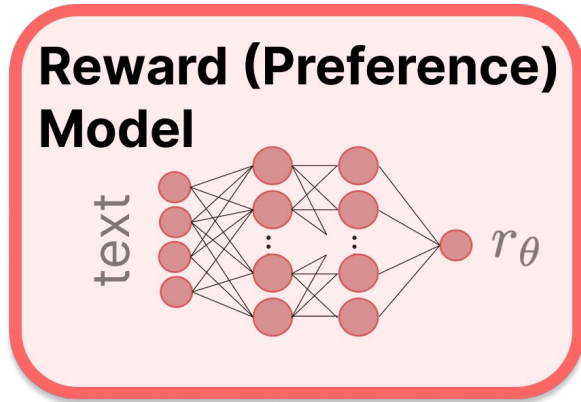
2. Reward model training



2. Reward model training



2. Reward model training



Reward model:

- Also transformer based LM
- Variation in sizes used (relative to policy)
- Outputs scalar from text input

Submit

Skip

«

Page 3 / 11

»

Total time: 05:39

Instruction

Summarize the following news article:

====

{article}

====

Include output

Output A

summary1

Rating (1 = worst, 7 = best)

1

2

3

4

5

6

7

Fails to follow the correct instruction / task ? Yes No

Inappropriate for customer assistant ? Yes No

Contains sexual content Yes No

Contains violent content Yes No

Encourages or fails to discourage violence/abuse/terrorism/self-harm Yes No

Denigrates a protected class Yes No

Gives harmful advice ? Yes No

Expresses moral judgment Yes No

Notes

(Optional) notes

(a)

Ranking outputs

To be ranked

B A team of researchers from Yale University and University of California, Davis studied the vocalization patterns of several different types of parrots. They found that parrots like to mimic human speech, and can produce a wide range of sounds, such as whistles, squawks, and other types of vocalizations...

C Parrots have been found to have the ability to understand numbers. Researchers have found that parrots can understand numbers up to six. In a series of experiments, the parrots were able to identify the amount of food items under a number of cups...

Rank 1 (*best*)

A A research group in the United States has found that parrots can imitate human speech with ease, and some of them can even do so in the same way as humans. This group studied the sounds that parrots make in their natural habitats and found that they use their tongues and beaks in ways that are strikingly...

Rank 2

Rank 3

E Scientists have found that green-winged parrots can tell the difference between two noises that are the same except for the order in which they are heard. This is important because green-winged parrots are known to imitate sounds. This research shows that they are able to understand the difference between sounds.

Rank 4

D Current research suggests that parrots see and hear things in a different way than humans do. While humans see a rainbow of colors, parrots only see shades of red and green. Parrots can also see ultraviolet light, which is invisible to humans. Many birds have this ability to see ultraviolet light, an ability

Rank 5 (*worst*)

(b)

Figure 12: Screenshots of our labeling interface. (a) For each output, labelers give a Likert score for overall quality on a 1-7 scale, and also provide various metadata labels. (b) After evaluating each output individually, labelers rank all the outputs for a given prompt. Ties are encouraged in cases where two outputs seem to be of similar quality.

labeling instruction

<https://docs.google.com/document/d/1MJCqDNjzD04UbcnVZ-LmeXJ04-TKEICDAepXyMCBUb8/edit#>

Table 12: Labeler demographic data

What gender do you identify as?	
Male	50.0%
Female	44.4%
Nonbinary / other	5.6%

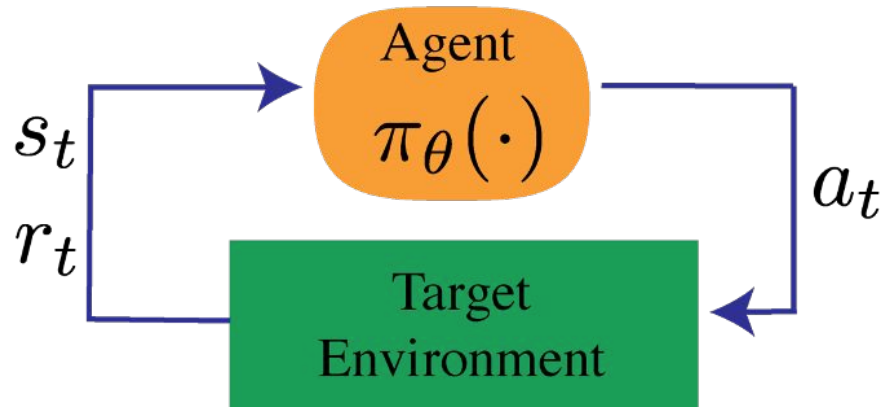
What ethnicities do you identify as?	
White / Caucasian	31.6%
Southeast Asian	52.6%
Indigenous / Native American / Alaskan Native	0.0%
East Asian	5.3%
Middle Eastern	0.0%
Latinx	15.8%
Black / of African descent	10.5%

What is your nationality?	
Filipino	22%
Bangladeshi	22%
American	17%
Albanian	5%
Brazilian	5%
Canadian	5%
Colombian	5%
Indian	5%
Uruguayan	5%
Zimbabwean	5%

What is your age?	
18-24	26.3%
25-34	47.4%
35-44	10.5%
45-54	10.5%
55-64	5.3%
65+	0%

What is your highest attained level of education?	
Less than high school degree	0%
High school degree	10.5%
Undergraduate degree	52.6%
Master's degree	36.8%
Doctorate degree	0%

Review: reinforcement learning basics



Some notation:

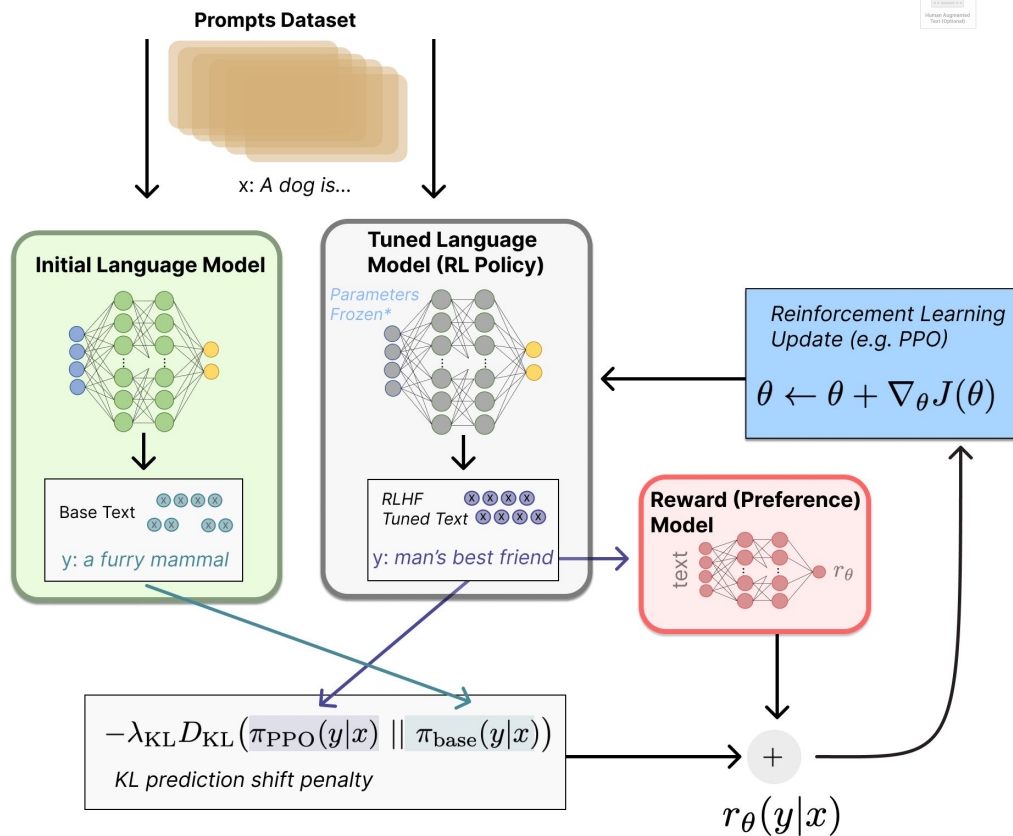
s_t : state

r_t : reward

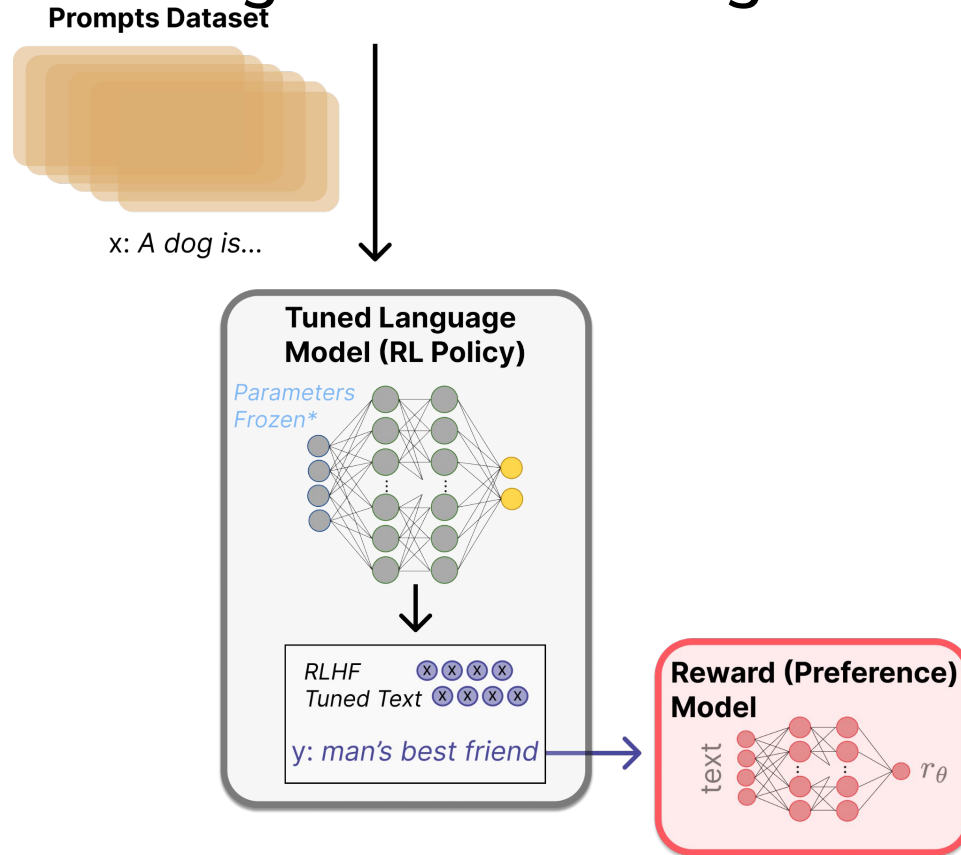
a_t : action

$a_t \sim \pi_{\theta}(s_t)$: policy

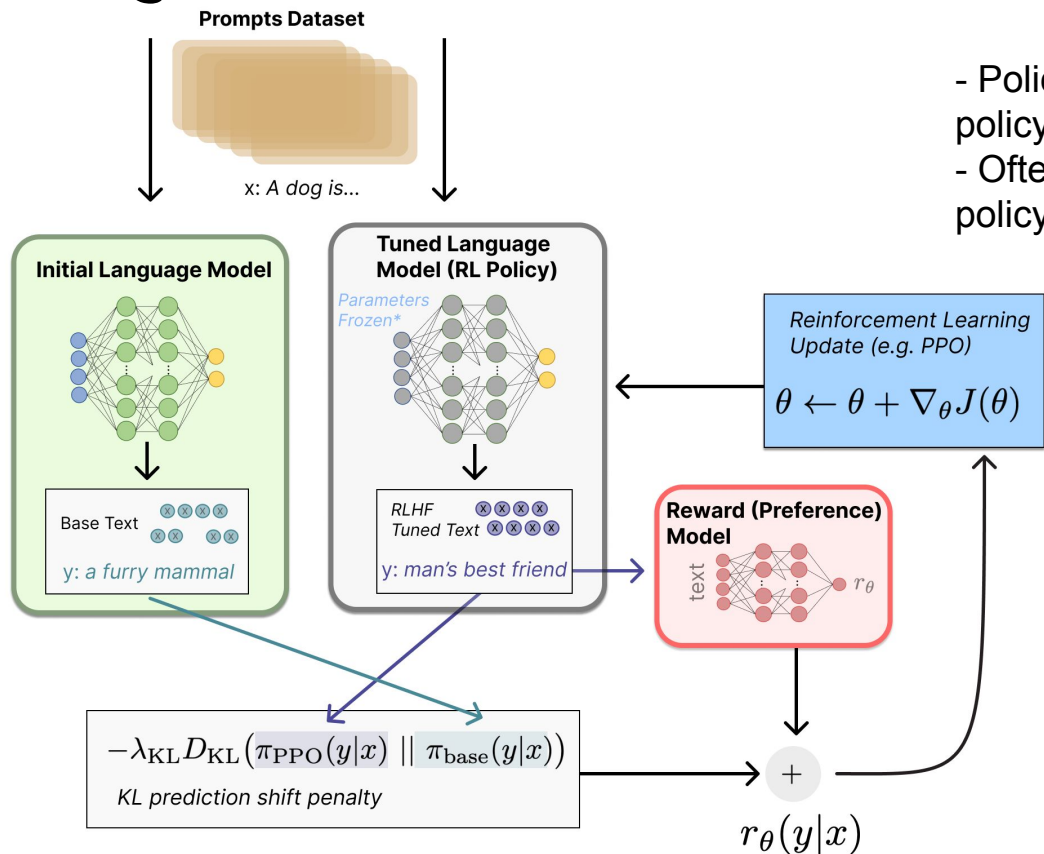
3. Fine tuning with RL



3. Fine tuning with RL - using a reward model



3. Fine tuning with RL - feedback & training



- Policy gradient updates policy LM directly.
- Often some parameters of policy are frozen.

A.3 Dataset sizes

In table 6, we report the sizes of datasets used to train / validate the SFT, RM, and RL models, in addition to whether the prompts were written by our labeling contractors or from our API.

Table 6: Dataset sizes, in terms of number of prompts.

SFT Data			RM Data			PPO Data		
split	source	size	split	source	size	split	source	size
train	labeler	11,295	train	labeler	6,623	train	customer	31,144
train	customer	1,430	train	customer	26,584	valid	customer	16,185
valid	labeler	1,550	valid	labeler	3,488			
valid	customer	103	valid	customer	14,399			

For SFT, note that we have many more labeler-written prompts than customer prompts—this is because, at the start of the project, we had labelers write instructions with a user interface that asked them to give an overarching template instruction as well as few-shot examples for that instruction. We synthetically constructed multiple SFT datapoints from the same instruction by sampling different sets of few-shot examples.

For the RM, recall that for every prompt, we collected rankings for K outputs (ranging from 4 to 9) and trained the model on all $\binom{K}{2}$, so the number of ranked pairs we trained the model on is an order of magnitude larger than the number of prompts.

Training a Helpful and Harmless Assistant with Reinforcement Learning from Human Feedback

Yuntao Bai*, Andy Jones, Kamal Ndousse,

**Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort,
Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion,**

**Tom Conerly, Sheer El-Showk, Nelson Elhage, Zac Hatfield-Dodds,
Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt,
Neel Nanda, Catherine Olsson, Dario Amodei, Tom Brown, Jack Clark,
Sam McCandlish, Chris Olah, Ben Mann, Jared Kaplan***

Anthropic

Our alignment interventions actually enhance the capabilities of large models, and can easily be combined with training for specialized skills (such as coding or summarization) without any degradation in alignment or performance. Models with less than about 10B parameters behave differently, paying an ‘alignment tax’ on their capabilities. This provides an example where models near the state-of-the-art may have been necessary to derive the right lessons from alignment research.

Labeling instruction

<https://docs.google.com/document/d/1MJCqDNjzD04UbcnVZ-LmeXJ04-TKEICDAepXyMCBUb8/edit#>

Table 12: Labeler demographic data

What gender do you identify as?	
Male	50.0%
Female	44.4%
Nonbinary / other	5.6%

What ethnicities do you identify as?	
White / Caucasian	31.6%
Southeast Asian	52.6%
Indigenous / Native American / Alaskan Native	0.0%
East Asian	5.3%
Middle Eastern	0.0%
Latinx	15.8%
Black / of African descent	10.5%

What is your nationality?	
Filipino	22%
Bangladeshi	22%
American	17%
Albanian	5%
Brazilian	5%
Canadian	5%
Colombian	5%
Indian	5%
Uruguayan	5%
Zimbabwean	5%

What is your age?	
18-24	26.3%
25-34	47.4%
35-44	10.5%
45-54	10.5%
55-64	5.3%
65+	0%

What is your highest attained level of education?	
Less than high school degree	0%
High school degree	10.5%
Undergraduate degree	52.6%
Master's degree	36.8%
Doctorate degree	0%

A.3 Dataset sizes

In table 6, we report the sizes of datasets used to train / validate the SFT, RM, and RL models, in addition to whether the prompts were written by our labeling contractors or from our API.

Table 6: Dataset sizes, in terms of number of prompts.

SFT Data			RM Data			PPO Data		
split	source	size	split	source	size	split	source	size
train	labeler	11,295	train	labeler	6,623	train	customer	31,144
train	customer	1,430	train	customer	26,584	valid	customer	16,185
valid	labeler	1,550	valid	labeler	3,488			
valid	customer	103	valid	customer	14,399			

For SFT, note that we have many more labeler-written prompts than customer prompts—this is because, at the start of the project, we had labelers write instructions with a user interface that asked them to give an overarching template instruction as well as few-shot examples for that instruction. We synthetically constructed multiple SFT datapoints from the same instruction by sampling different sets of few-shot examples.

For the RM, recall that for every prompt, we collected rankings for K outputs (ranging from 4 to 9) and trained the model on all $\binom{K}{2}$, so the number of ranked pairs we trained the model on is an order of magnitude larger than the number of prompts.

Multi-modality Language Model

- Allow LM input with text + X (mostly image now)
- Example: GPT4

Benchmark	GPT-4 <small>Evaluated few-shot</small>	Few-shot SOTA	SOTA <small>Best external model (includes benchmark-specific training)</small>
<u>VQAv2</u> <small>VQA score (test-dev)</small>	77.2% 0-shot	67.6% <u>Flamingo 32-shot</u>	84.3% <u>PaLI-17B</u>
<u>TextVQA</u> <small>VQA score (val)</small>	78.0% 0-shot	37.9% <u>Flamingo 32-shot</u>	71.8% <u>PaLI-17B</u>
<u>ChartQA</u> <small>Relaxed accuracy (test)</small>	78.5% ^A	-	58.6% <u>Pix2Struct Large</u>
<u>AI2 Diagram (AI2D)</u> <small>Accuracy (test)</small>	78.2% 0-shot	-	42.1% <u>Pix2Struct Large</u>
<u>DocVQA</u> <small>ANLS score (test)</small>	88.4% 0-shot (pixel-only)	-	88.4% <u>ERNIE-Layout 2.0</u>
<u>Infographic VQA</u> <small>ANLS score (test)</small>	75.1% 0-shot (pixel-only)	-	61.2% <u>Applica.ai TILT</u>
<u>TVQA</u> <small>Accuracy (val)</small>	87.3% 0-shot	-	86.5% <u>MERLOT Reserve Large</u>
<u>LSMDC</u> <small>Fill-in-the-blank accuracy (test)</small>	45.7% 0-shot	31.0% <u>MERLOT Reserve 0-shot</u>	52.9% <u>MERLOT</u>

How to train? training data?

Example of GPT-4 visual input:

User What is funny about this image? Describe it panel by panel.



Source: <https://www.reddit.com/r/hmmm/comments/ubab5v/hmmm/>

GPT-4 The image shows a package for a "Lightning Cable" adapter with three panels.

Panel 1: A smartphone with a VGA connector (a large, blue, 15-pin connector typically used for computer monitors) plugged into its charging port.

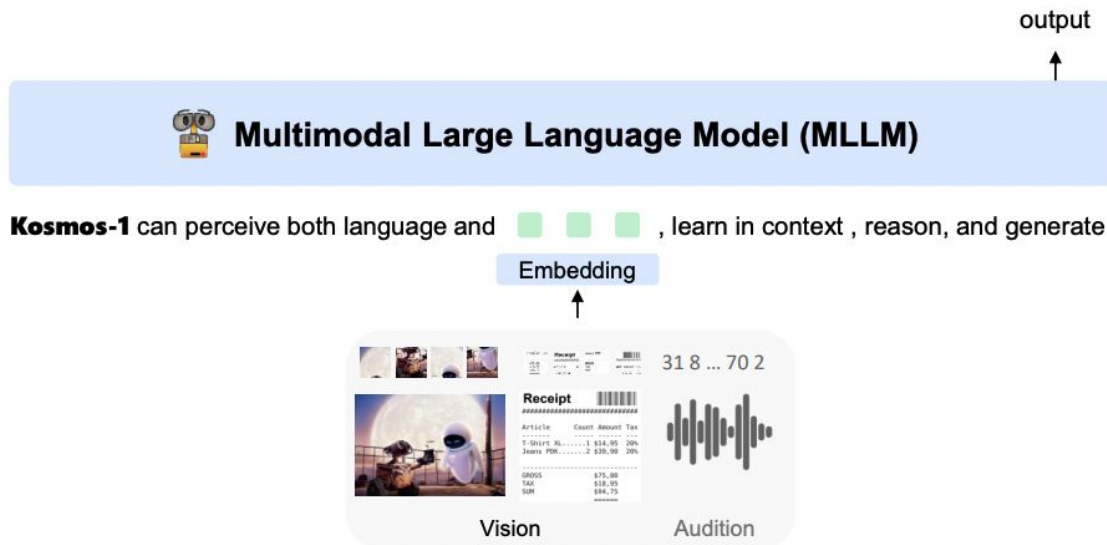
Panel 2: The package for the "Lightning Cable" adapter with a picture of a VGA connector on it.

Panel 3: A close-up of the VGA connector with a small Lightning connector (used for charging iPhones and other Apple devices) at the end.

The humor in this image comes from the absurdity of plugging a large, outdated VGA connector into a small, modern smartphone charging port.

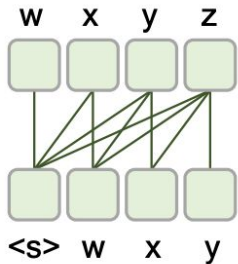
Multi-modality Language Model

- Interleaved image-text input:
 - `<s>` paragraph `</s>` `<image>` Image Embedding `</image>` paragraph `</s>`
 - Continuous / discrete image embedding

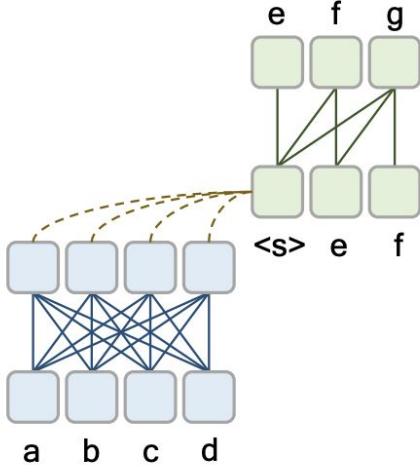


MetaLM: Semi-Causal Language Modeling

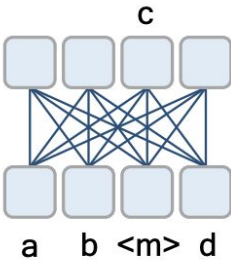
— Self-Attention
 - - - Cross-Attention



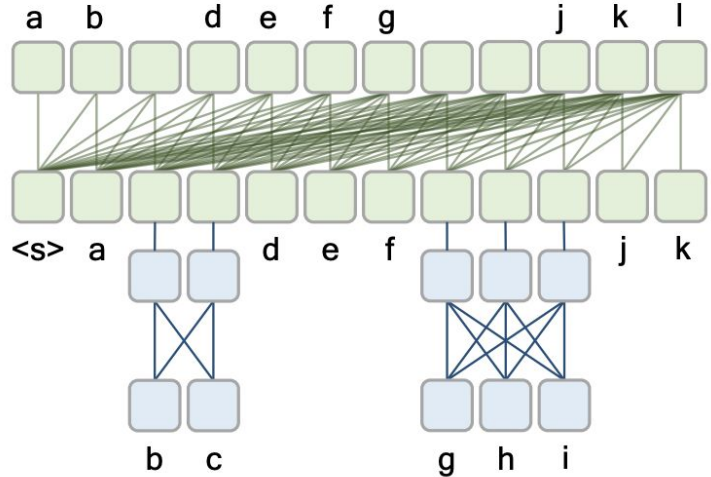
(a) Causal LM
 (Unidirectional)



(b) Prefix LM
 (Encoder-Decoder
 with Cross-Attention)

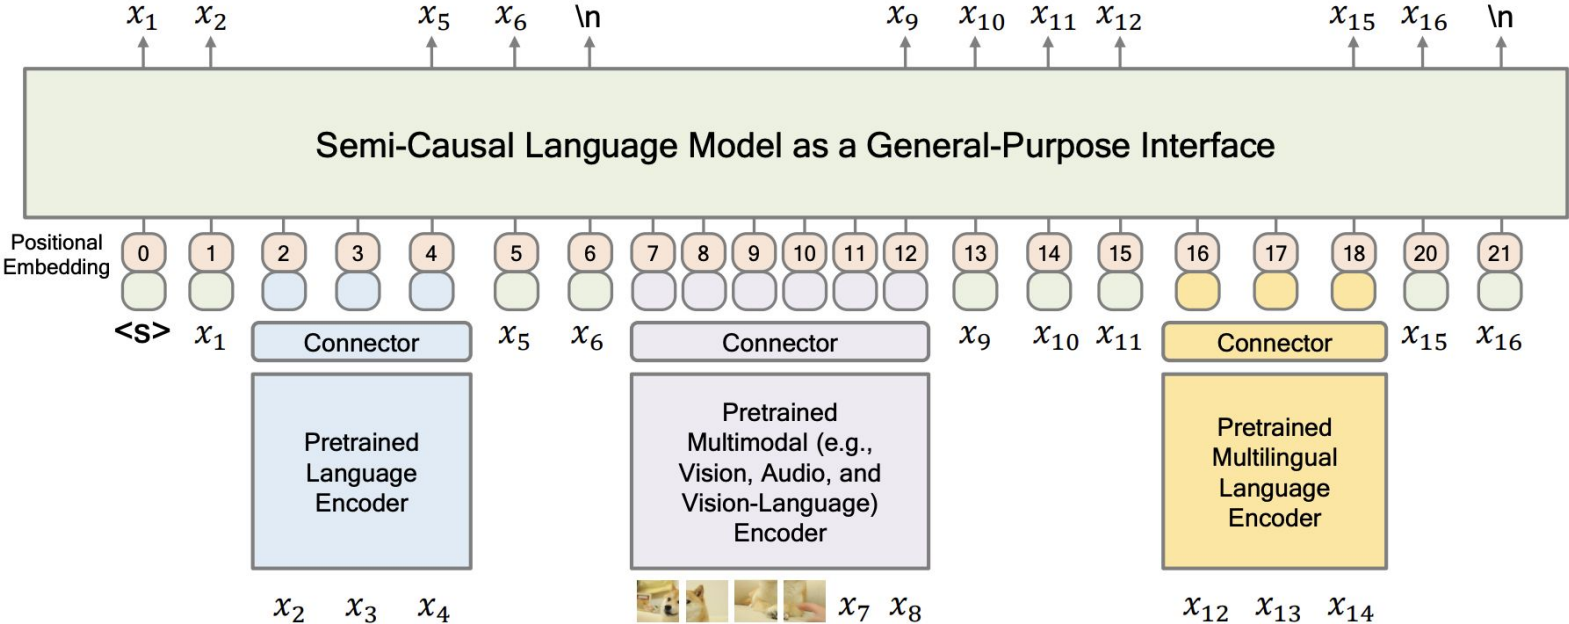


(c) Non-Causal LM
 (Bidirectional)



(d) Semi-Causal LM

MetaLM: Semi-Causal Language Modeling



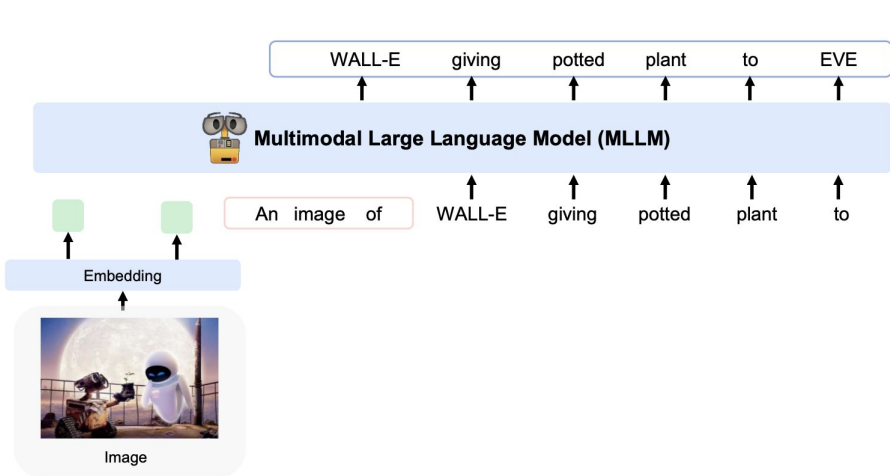
Hao, Yaru, et al. "Language models are general-purpose interfaces." arXiv preprint arXiv:2206.06336 (2022).

Training data

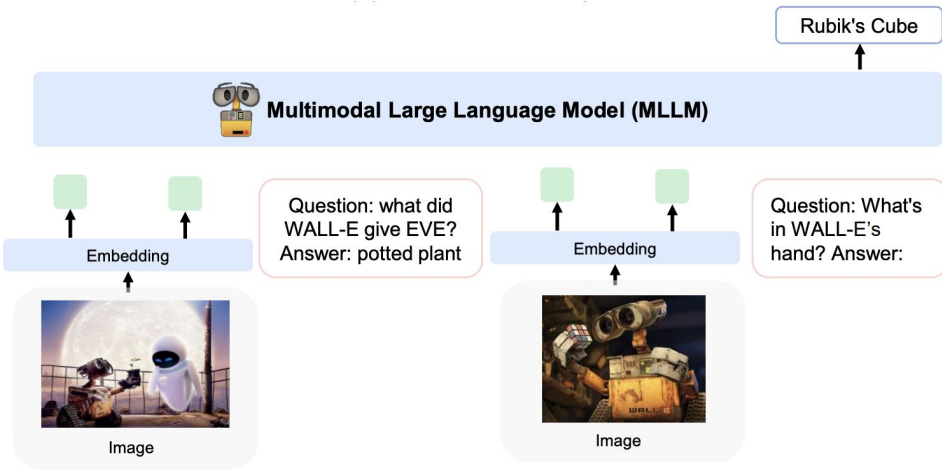
- Text
- Image-Caption pairs
- Interleaved image-text

Datasets	Format Examples
Text	<code><s> KOSMOS-1 can perceive multimodal input, learn in context, and generate output. </s></code>
Image-Caption	<code><s> <image> Image Embedding </image> WALL-E giving potted plant to EVE. </s></code>
Multimodal	<code><s> <image> Image Embedding </image> This is WALL-E. <image> Image Embedding </image> This is EVE. </s></code>

Evaluation



(a) Zero-shot learning



(b) Few-shot learning

Input Prompt



An image of



An image of



Question: What coffee shop is in the background?
Answer:



Question: What is inside this canned good? Answer:

Completion

a girl blowing out a candle on her birthday cake.

(1)

a group of people posing for a wedding photo.

(2)

Starbucks

(3)

Corn

(4)

Input Prompt



Question: What is on the pillowcase? Answer:



Question: What is the title of this book? Answer:



An image of

Hope you have done it.

An image of

Completion

a Texas Longhorns pillow case.

(5)

A Fine, Fine School

(6)

a parking sign that says "Waiting? Please turn off your engine."

(7)

Hope you have done it

(8)

Input Prompt



Question: which studio created the movie characters in the image?

Answer:

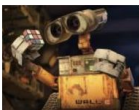
Completion

Disney



Standard Prompting

Input Prompt



Introduce this picture in detail:

Completion

It is a robot that is holding up a cube that is the same color as the wall. The image is from the film, WALL-E, by Pixar Animation Studios. It is the story of a small robot named WALL-E and his quest to build the ultimate life form. WALL-E is created by a group...

Input Prompt



It is a robot that is holding up a cube that is the same color as the wall. The image is from the film, WALL-E, by Pixar Animation Studios. It is the story...

Question: which studio created the movie characters in the image?

Answer:

Completion

Pixar Animation Studios



Multimodal Chain-of-Thought Prompting

A holistic multi-modality system

- Although multi-modal LM can perceive both image/audio and language, still can't generate output in other modality
 - Use other generative foundation models
 - E.g., AudioGPT / Visual ChatGPT

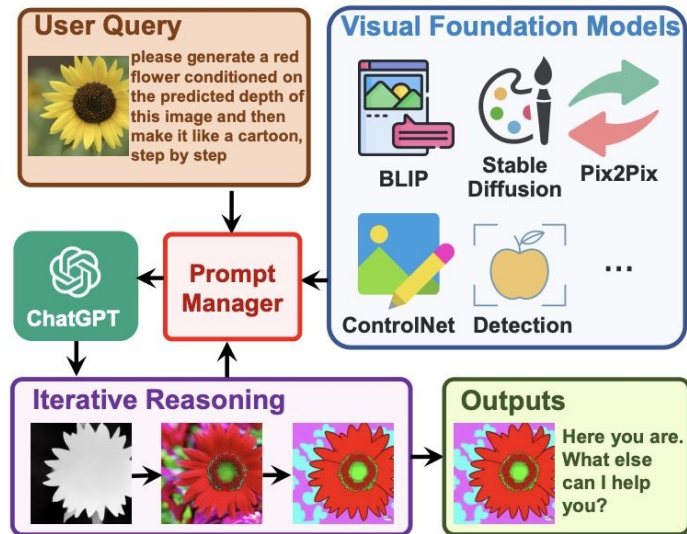
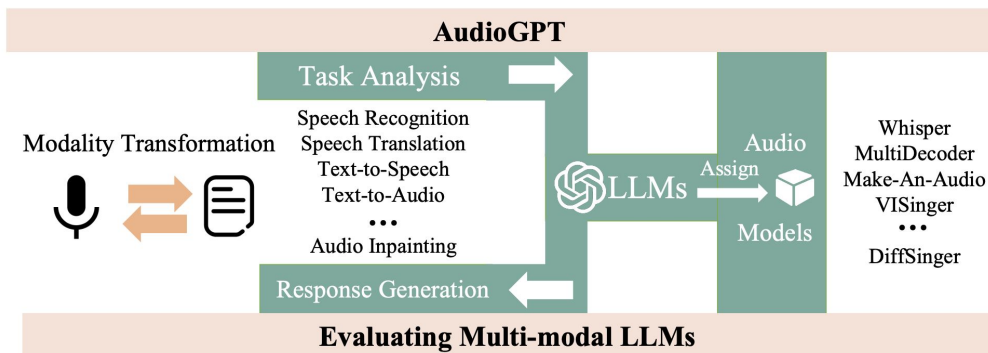
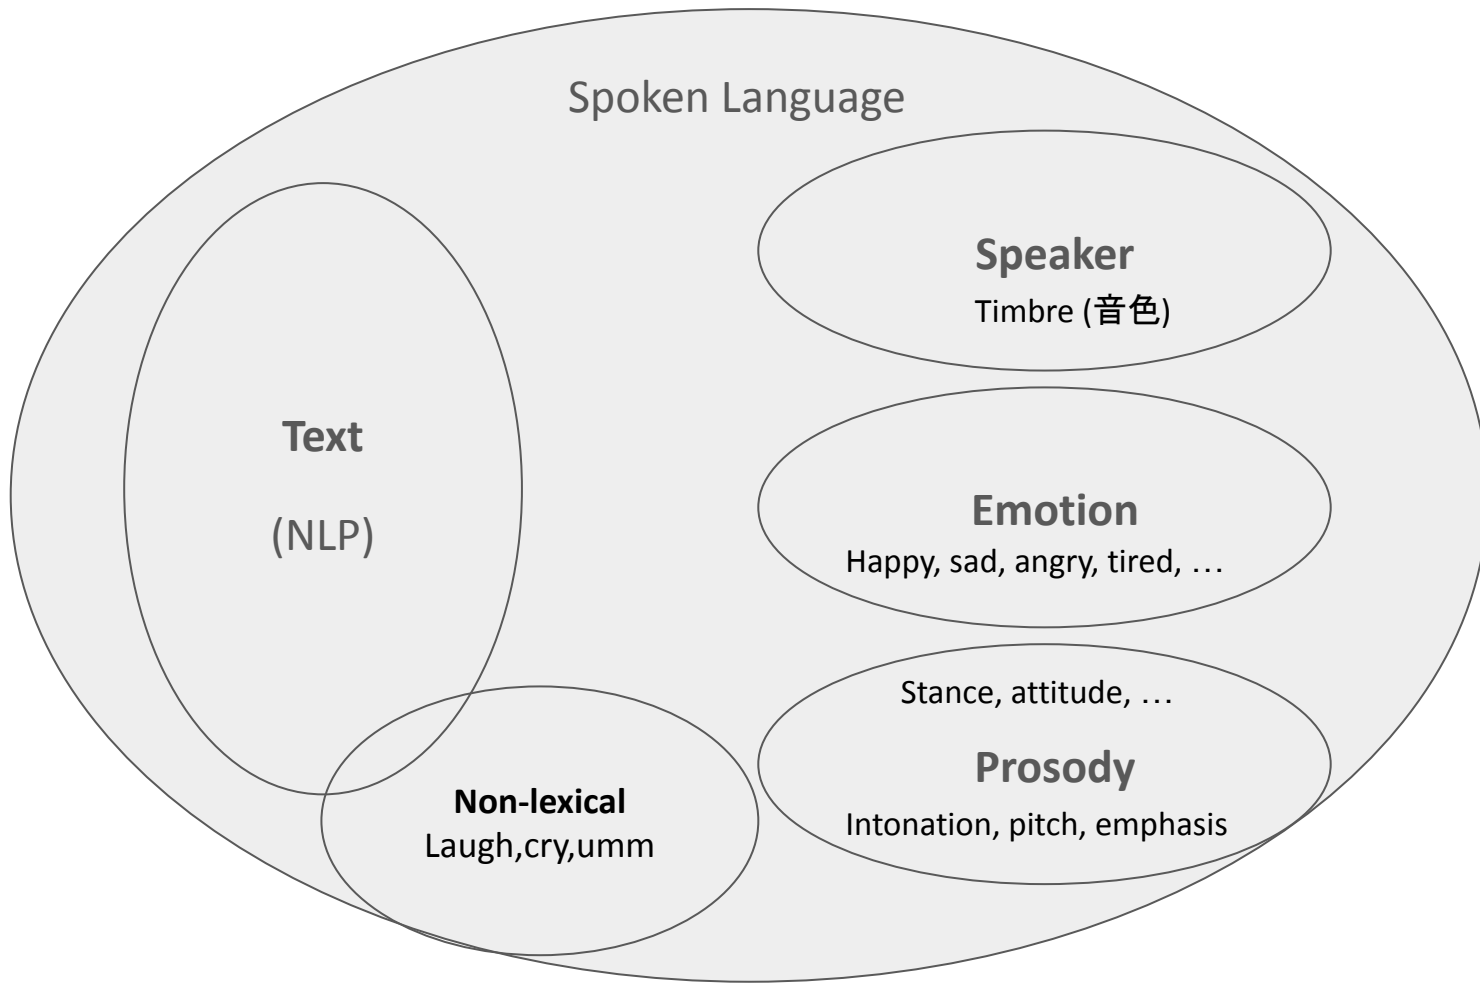


Figure 1. Architecture of Visual ChatGPT.

Research experience sharing

- Self-supervised Learning for speech and language
- Publish paper at top conference

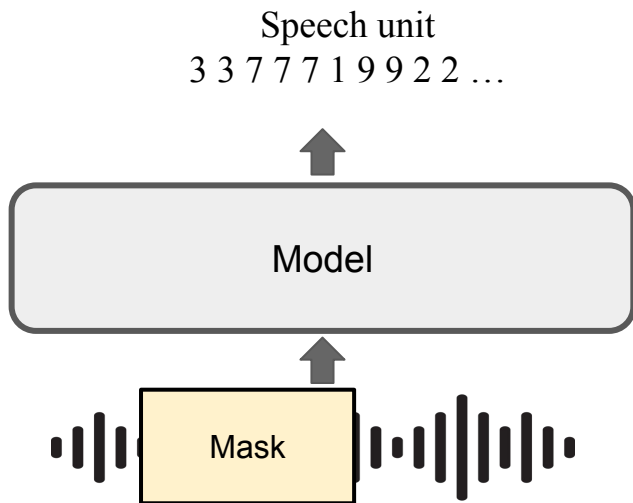


Self-supervised Learning for speech

For a comprehensive review:

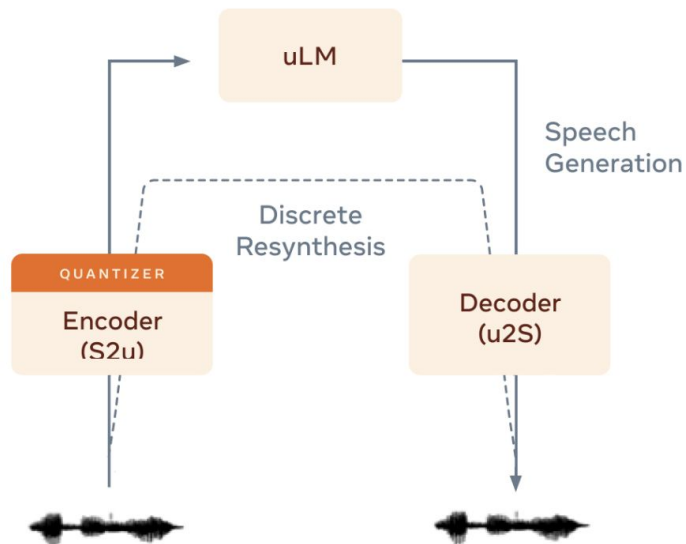
Mohamed, Abdelrahman, et al. "Self-supervised speech representation learning: A review." *IEEE Journal of Selected Topics in Signal Processing* (2022).

Masked Pretraining



Hsu, Wei-Ning, et al. "Hubert: Self-supervised speech representation learning by masked prediction of hidden units." *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 29 (2021): 3451-3460.

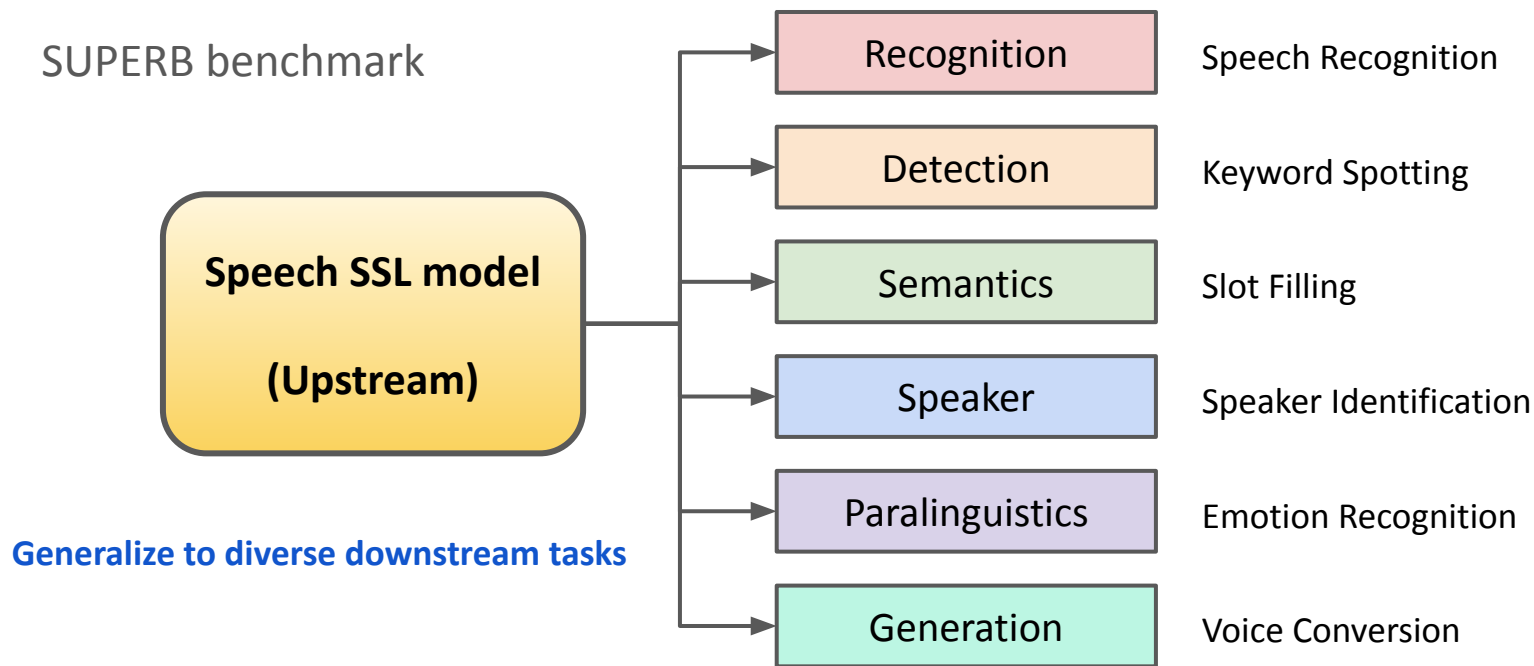
Generative Pretraining



Lakhotia, Kushal, et al. "On generative spoken language modeling from raw audio." *Transactions of the Association for Computational Linguistics* 9 (2021): 1336-1354.

Self-supervised Learning for speech

- SUPERB benchmark



Shu wen Yang, Po-Han Chi, Yung-Sung Chuang, Cheng-I Jeff Lai, Kushal Lakhotia, Yist Y. Lin, Andy T. Liu, Jiatong Shi, Xuankai Chang, Guan-Ting Lin, Tzu-Hsien Huang, Wei-Cheng Tseng, Ko tik Lee, Da-Rong Liu, Zili Huang, Shuyan Dong, Shang-Wen Li, Shinji Watanabe, Abdelrahman Mohamed, and Hung yi Lee, "SUPERB: Speech Processing Universal PERFORMANCE Benchmark," *Interspeech 2021*

Hsiang-Sheng Tsai, Heng-Jui Chang, Wen-Chin Huang, Zili Huang, Kushal Lakhotia, Shu-wen Yang, Shuyan Dong, Andy Liu, Cheng-I Lai, Jiatong Shi, et al., "SUPERB-SG: Enhanced speech processing universal performance benchmark for semantic and generative capabilities," *ACL 2022*

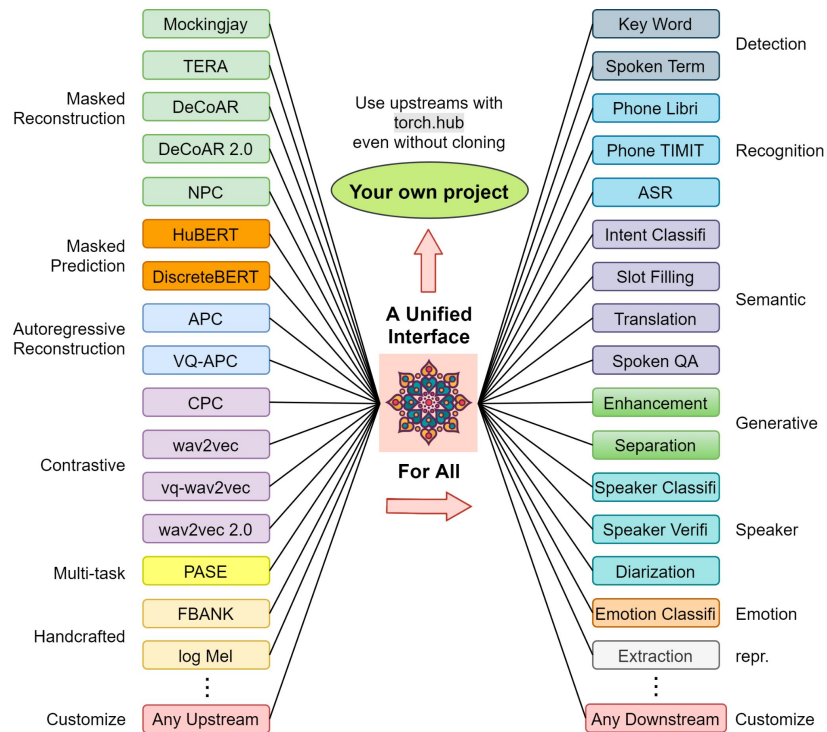
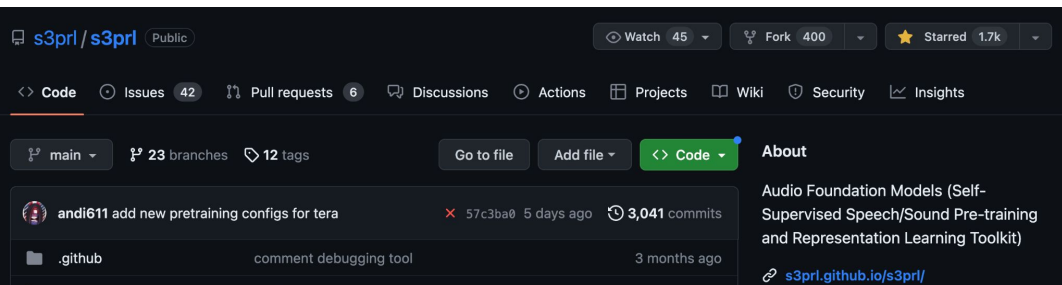
Open-source codebase

- More follow-up works
 - Easy to reproduce results
- More impact
 - High citation

[Superb: Speech processing universal performance benchmark](#)

S Yang, PH Chi, YS Chuang, CIJ Lai, K Lakhotia, YY Lin, AT Liu, J Shi, ...
arXiv preprint arXiv:2105.01051

325 2021



Publish paper at top conference

- Experience: Speech/NLP-related conference (but applicable to others ML conference)
- Procedure
 - Paper submission
 - Review
 - Rebuttal
 - Acceptance notification
 - Camera ready submission
 - Paper presentation & conference attendance

Pick a suitable conference

AI Conference Deadlines

Countdowns to top CV/NLP/ML/Robotics/AI conference deadlines. To add/edit a conference, [send in a pull request](#).

Subject Filter:

ML, CV, NLP, SP ▾ ↕

Deadlines are shown in Asia/Taipei time. To view them in conference website timezones, click on them.

You can optionally export all deadlines to [Google Calendar](#) or [.ics](#).

[BMVC 2023](#) 🌐

November 20-24, 2023. [Aberdeen, UK](#).

Note: Workshop Submission Deadline on May 5th, 2023

[computer vision](#)

02 days 17h 08m 24s

Deadline: Sat May 13 2023 07:59:59 GMT+0800

 [Google](#) [Yahoo!](#) [iCal](#) [Outlook](#)

[NeurIPS 2023](#) 🌐


December 10 - December 16, 2023. [New Orleans, Louisiana, USA](#).

Note: Mandatory abstract deadline on May 11, 2023

[machine learning](#)

07 days 13h 08m 25s

Deadline: Thu May 18 2023 04:00:00 GMT+0800

 [Google](#) [Yahoo!](#) [iCal](#) [Outlook](#)

<https://aideadlin.es/?sub=ML,CV,NLP,RO,SP,GR>

Conference paper submission process

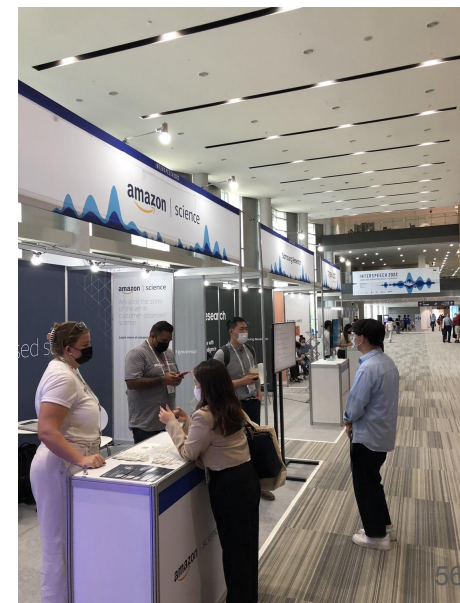
1. Paper submitted by authors
2. Reviewer (2~4) write comments
3. Author rebuttal
4. Meta-reviewer write review based on review and rebuttal
5. Area chairs make acceptance decision
6. Paper revision for camera-ready paper

Rebuttal

- General Summary for reviewer's comments
- Response to each reviewer
- Tips:
 - Use 1-2 sentences to objectively summarize positive and negative feedbacks
 - List reviewers' questions, clarify the misunderstanding
 - Present details experimental results if needed

Conference attendance

- Present your work
 - Oral: presentation slide with QA session
 - Poster: Be careful to the poster format, size; Cloth Poser (recommended)
- Social network
 - Meet famous researchers in person
 - Make friends with other junior researchers
- Job opportunities
 - Company vendors
 - Job fair



Resources and suggestions

- Twitter

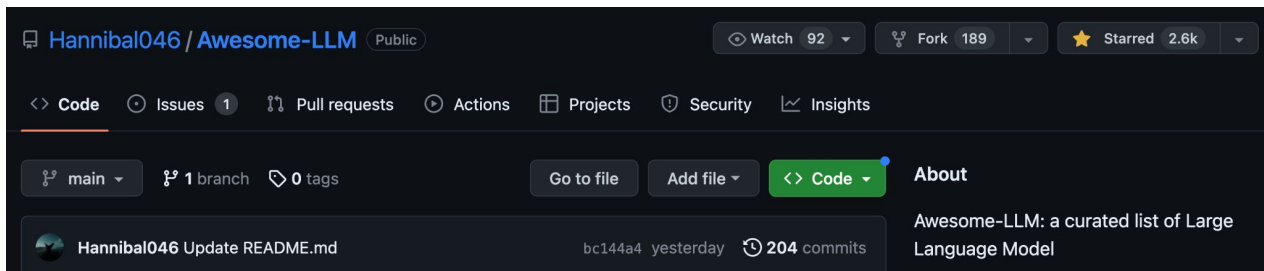
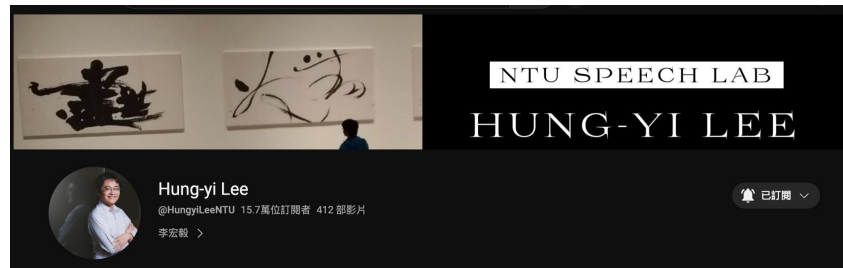
- Follow famous researchers in your field
 - They will share the latest research result by tweets
- Paper tweeter
 - Recommend new papers

- Youtube

- **Hung-yi Lee!**
- Yannic Kilcher
- The ai epiphany
- ...

- Github repo

- Awesome-...



Q & A