

[cs.LG] 31 May 2023

Let's Verify Step by Step

Hunter Lightman* Vineet Kosaraju* Yura Burda* Harri Edwards
Bowen Baker Teddy Lee Jan Leike John Schulman Ilya Sutskever

Karl Cobbe*

OpenAI

Abstract

In recent years, large language models have greatly improved in their ability to perform complex multi-step reasoning. However, even state-of-the-art models still regularly produce logical mistakes. To train more reliable models, we can turn either to outcome supervision, which provides

AI News

Samuel Albanie

Xi Chen, Josip Djolonga, Piotr Padlewski, Basil Mustafa, Soravit Changpinyo, Jialin Wu, Carlos Riquelme Ruiz, Sebastian Goodman, Xiao Wang, Yi Tay, Siamak Shakeri, Mostafa Dehghani, Daniel Salz, Mario Lucic, Michael Tschannen, Arsha Nagrani, Hexiang Hu, Mandar Joshi, Bo Pang, Ceslee Montgomery, Paulina Pietrzyk, Marvin Ritter, AJ Piergiovanni, Matthias Minderer, Filip Pavetic, Austin Waters, Gang Li, Ibrahim Alabdulmohsin, Lucas Beyer, Julien Amelot, Kenton Lee, Andreas Peter Steiner, Yang Li, Daniel Keysers, Anurag Arnab, Yuanzhong Xu, Keran Rong, Alexander Kolesnikov, Mojtaba Seyedhosseini, Anelia Angelova, Xiaohua Zhai, Neil Houlsby, Radu Soricut

Google Research
pali-communications@google.com

Abstract

CV] 29 May 2023

[cs.CL] 30 May 2023

GPT4GEO: How a Language Model Sees the World's Geography

Jonathan Roberts
CAML Lab
University of Cambridge
jdr53@cam.ac.uk

Timo Lüddecke
Institute of Computer Science and CIDAS
University of Göttingen
timo.lueddecke@uni-goettingen.de

Sowmen Das
CAML Lab
University of Cambridge
sd973@cam.ac.uk

Kai Han
Visual AI Lab
The University of Hong Kong
kaihanx@hku.hk

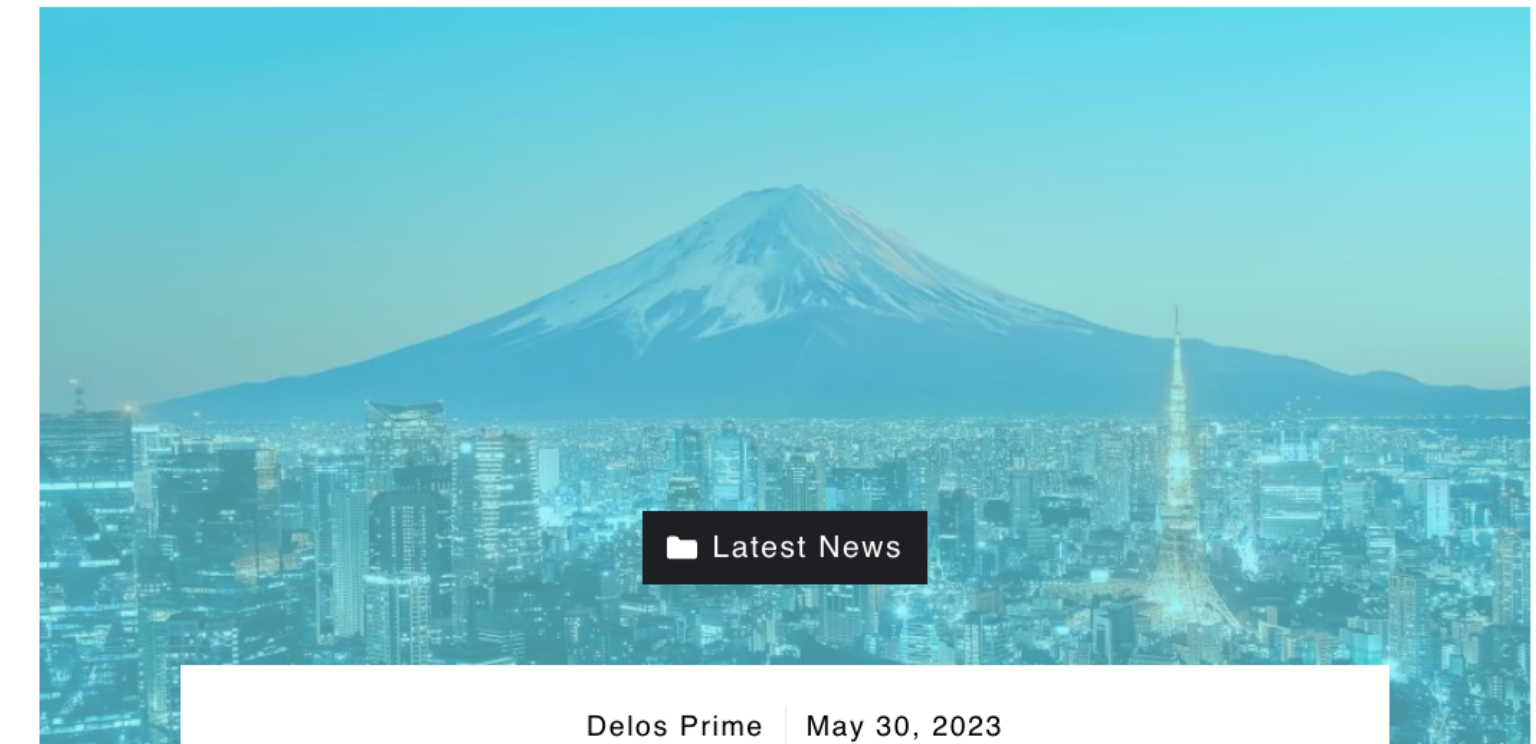
Samuel Albanie
CAML Lab
University of Cambridge
samuel.albanie.academic@gmail.com

Technomancers.ai

e/acc

Blog About S

Search



Delos Prime | May 30, 2023

Japan Goes All In: Copyright Doesn't Apply To AI Training

Recent P

USAF's Rog
Never Was

Air Force Co
Complete Fi

Summer Mo
Human


Falcon LLM:
Commercial

Japan Goes

Let's Verify Step by Step

31st May 2023

"... **process supervision** significantly outperforms **outcome supervision** for training models to solve problems from the challenging MATH dataset."



Process-supervised RM

Outcome-supervised RM

Let's Verify Step by Step

Hunter Lightman* Vineet Kosaraju* Yura Burda* Harri Edwards
 Bowen Baker Teddy Lee Jan Leike John Schulman Ilya Sutskever
 Karl Cobbe*
 OpenAI

Abstract

In recent years, large language models have greatly improved in their ability to perform complex multi-step reasoning. However, even state-of-the-art models still regularly produce logical mistakes. To train more reliable models, we can turn either to outcome supervision, which provides feedback for a final result, or process supervision, which provides feedback for each intermediate reasoning step. Given the importance of training reliable models, and given the high cost of human feedback, it is important to carefully compare the both methods. Recent work has already begun this comparison, but many questions still remain. We conduct our own investigation, finding that process supervision significantly outperforms outcome supervision for training models to solve problems from the challenging MATH dataset. Our process-supervised model solves 78% of problems from a representative subset of the MATH test set. Additionally, we show that active learning significantly improves the efficacy of process supervision. To support related research, we also release PRM800K, the complete dataset of 800,000 step-level human feedback labels used to train our best reward model.

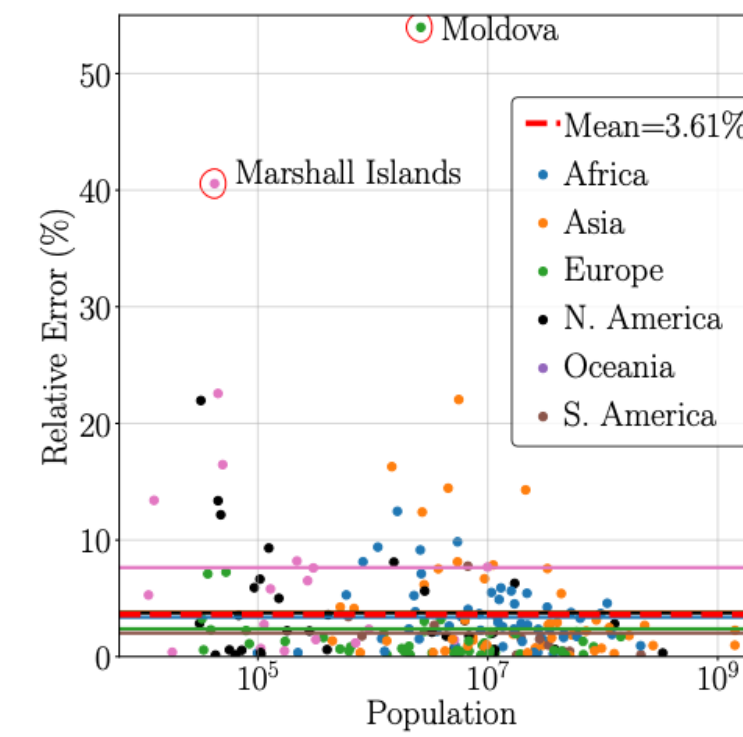
	ORM	PRM	Majority Vote	# Problems
AP Calculus	68.9%	86.7%	80.0%	45
AP Chemistry	68.9%	80.0%	71.7%	60
AP Physics	77.8%	86.7%	82.2%	45
AMC10/12	49.1%	53.2%	32.8%	84
Aggregate	63.8%	72.9%	61.3%	234

"Process supervision is more likely to produce **interpretable reasoning**..."

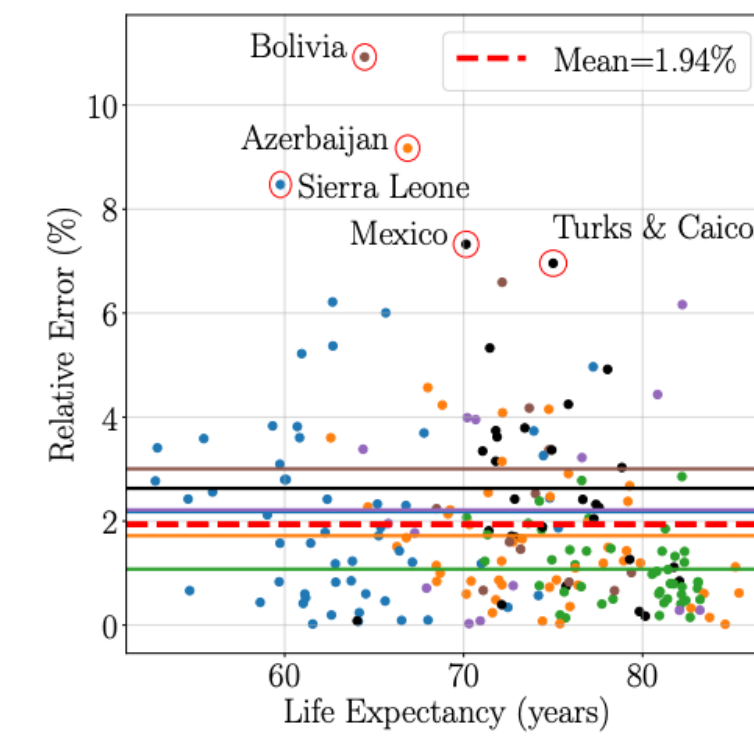
"...process supervision in fact incurs a **negative alignment tax**"

GPT4Geo 30th May 2023

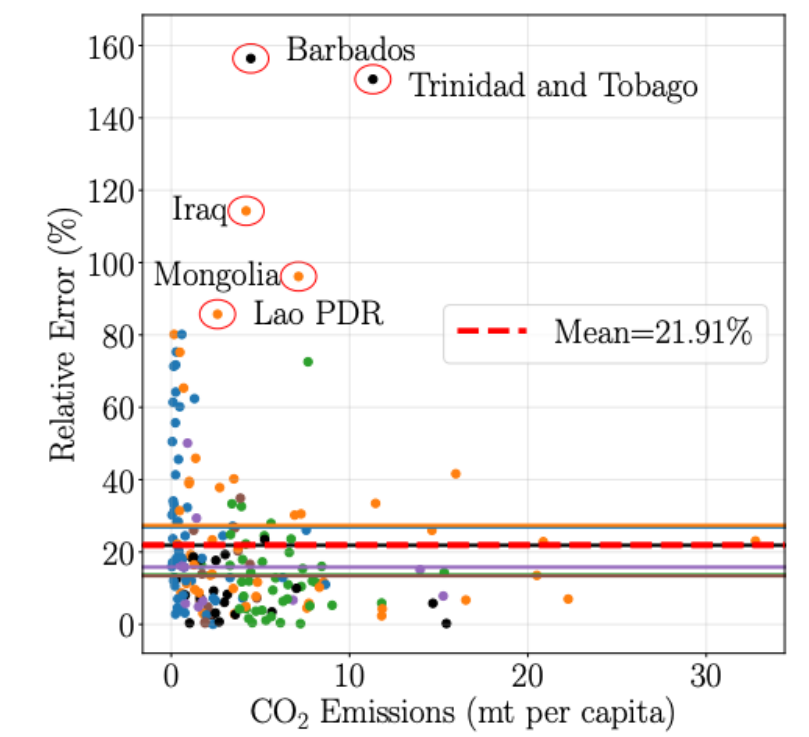
studies degree to which "GPT-4 has ...factual geographic knowledge and is capable of using this knowledge for **interpretative reasoning**"



(a) Population (2021).



(b) Life Expectancy (2020).



(c) CO₂ Emissions (2019).

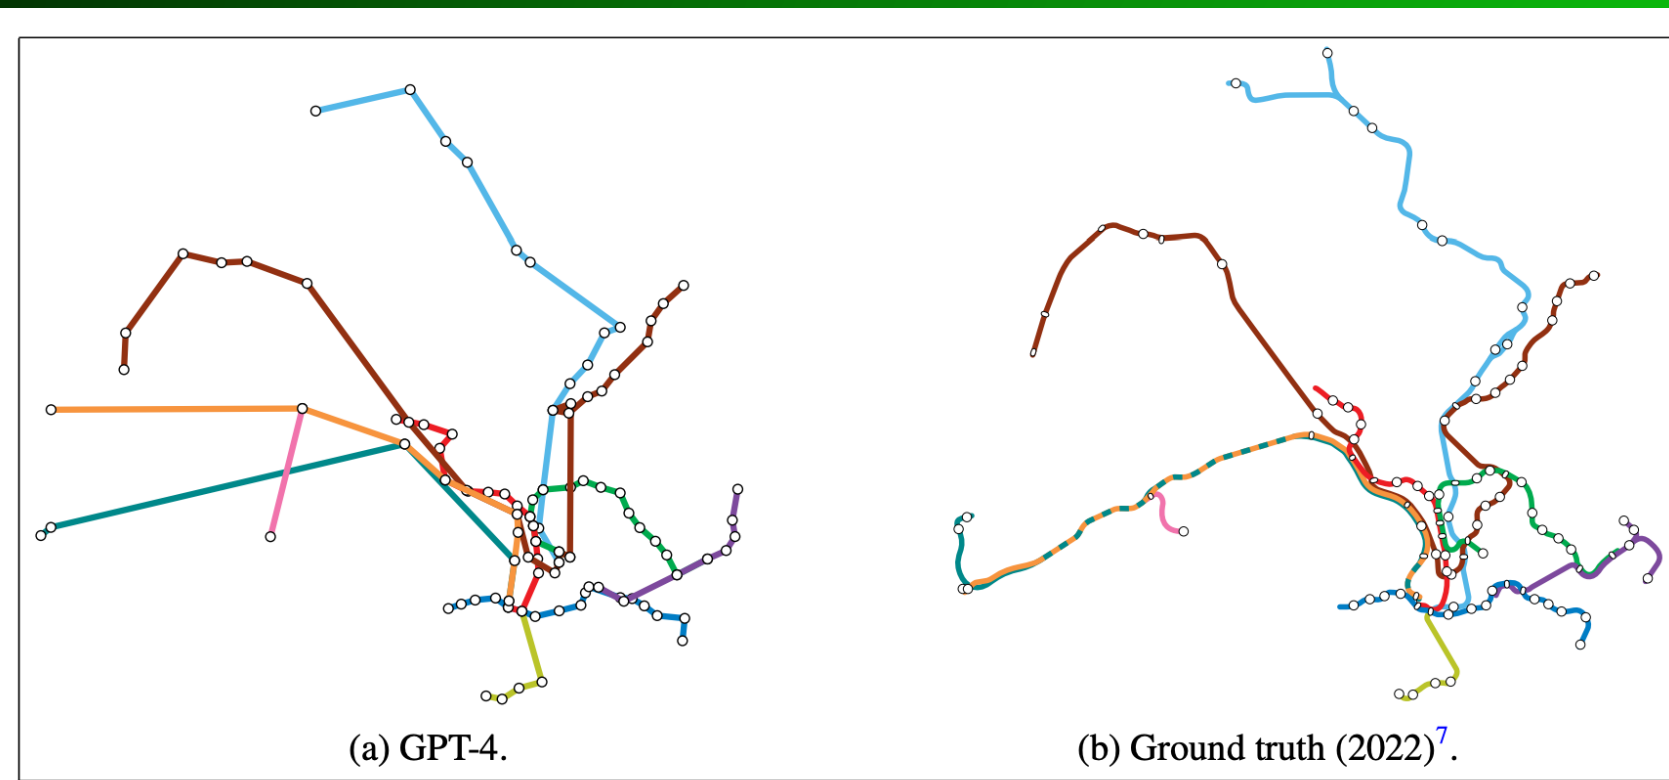


Figure 9: Hong Kong Mass Transit Railway (MTR) Network Map.

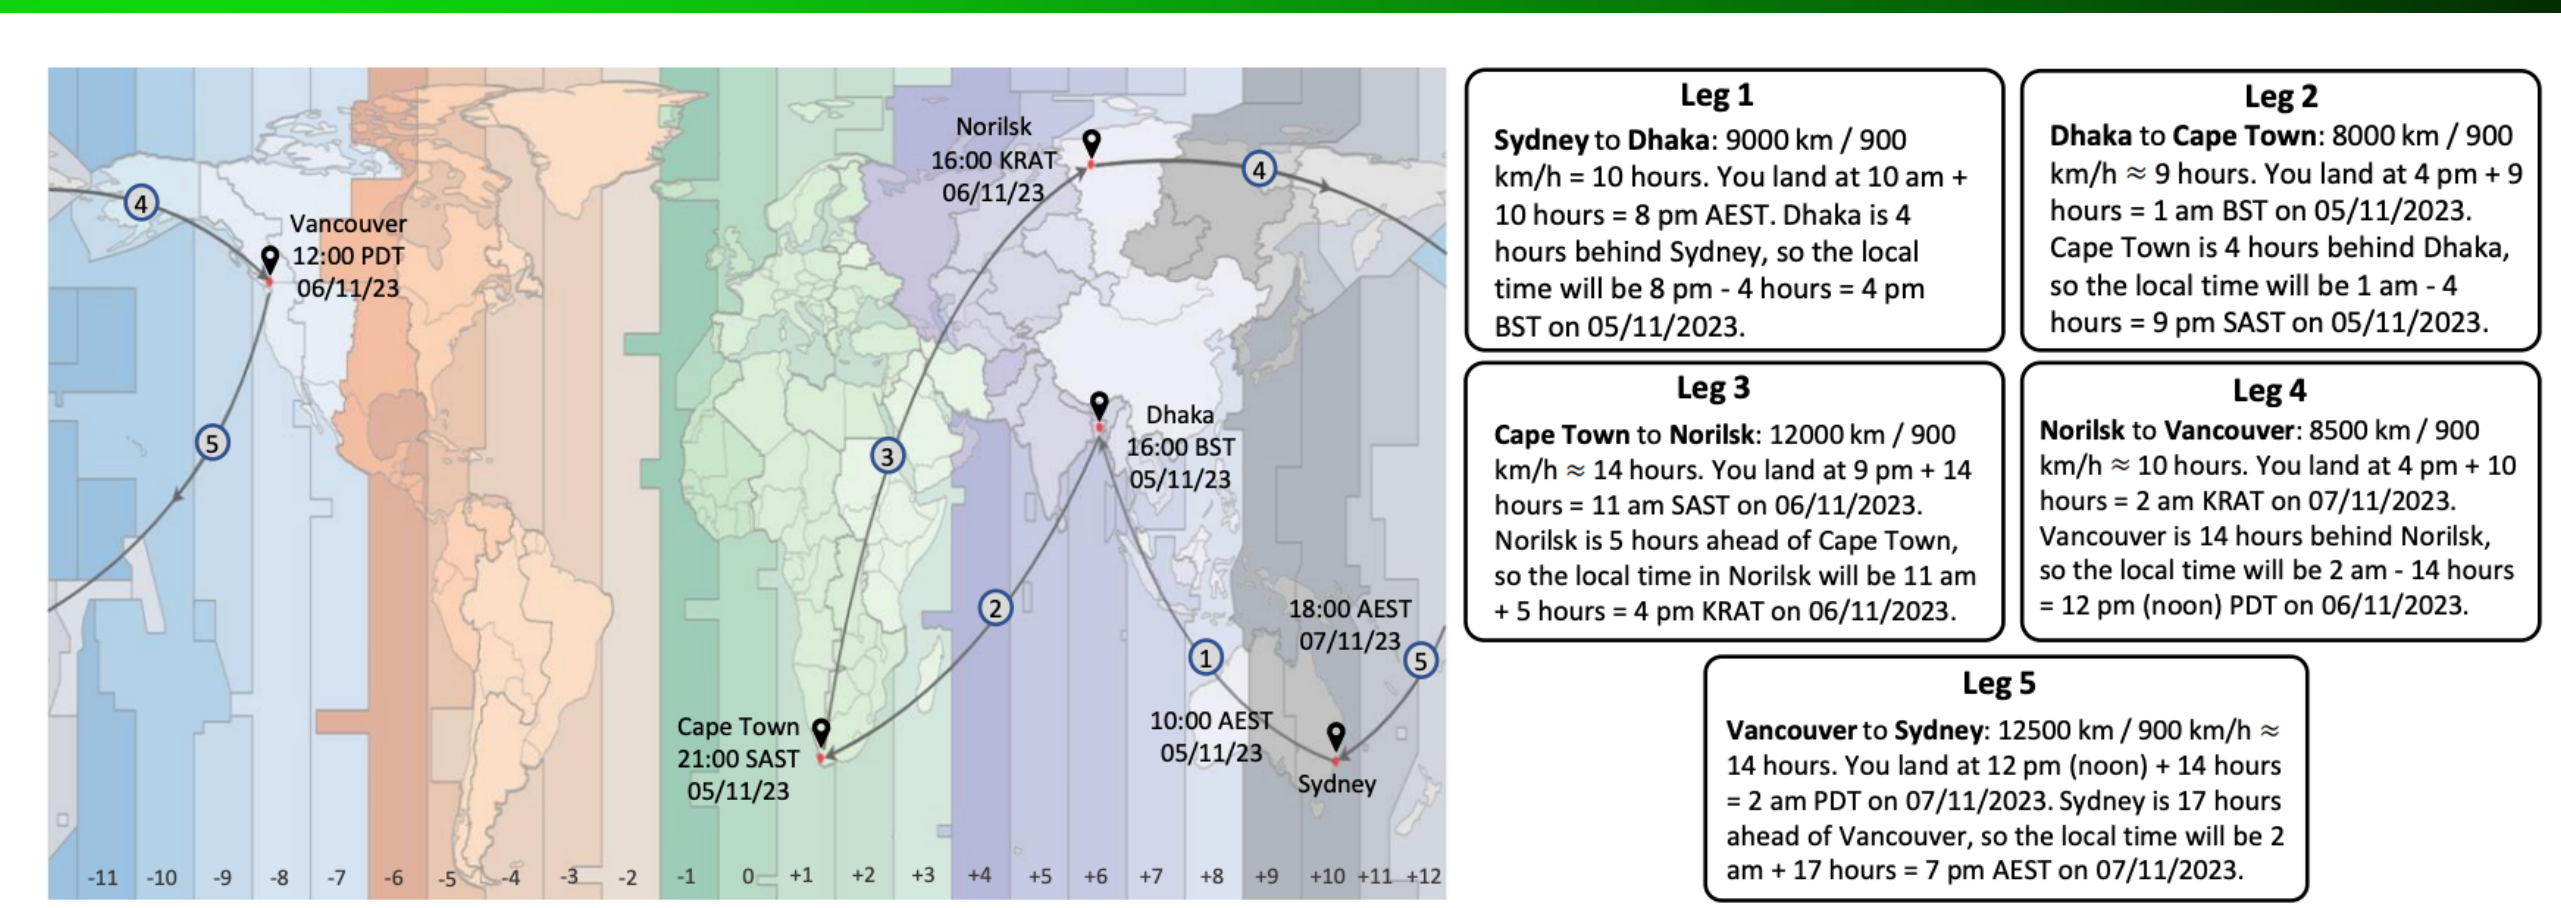


Figure 8: Navigation. GPT-4 successfully navigates a journey with multiple time zones.

GPT4GEO: How a Language Model Sees the World's Geography

Jonathan Roberts
CAML Lab
University of Cambridge
jdr53@cam.ac.uk

Timo Lüddecke
Institute of Computer Science and CIDAS
University of Göttingen
timo.lueddecke@uni-goettingen.de

Sowmen Das
CAML Lab
University of Cambridge
sd973@cam.ac.uk

Kai Han
Visual AI Lab
The University of Hong Kong
kaihanx@hku.hk

Samuel Albanie
CAML Lab
University of Cambridge
samuel.albanie.academic@gmail.com

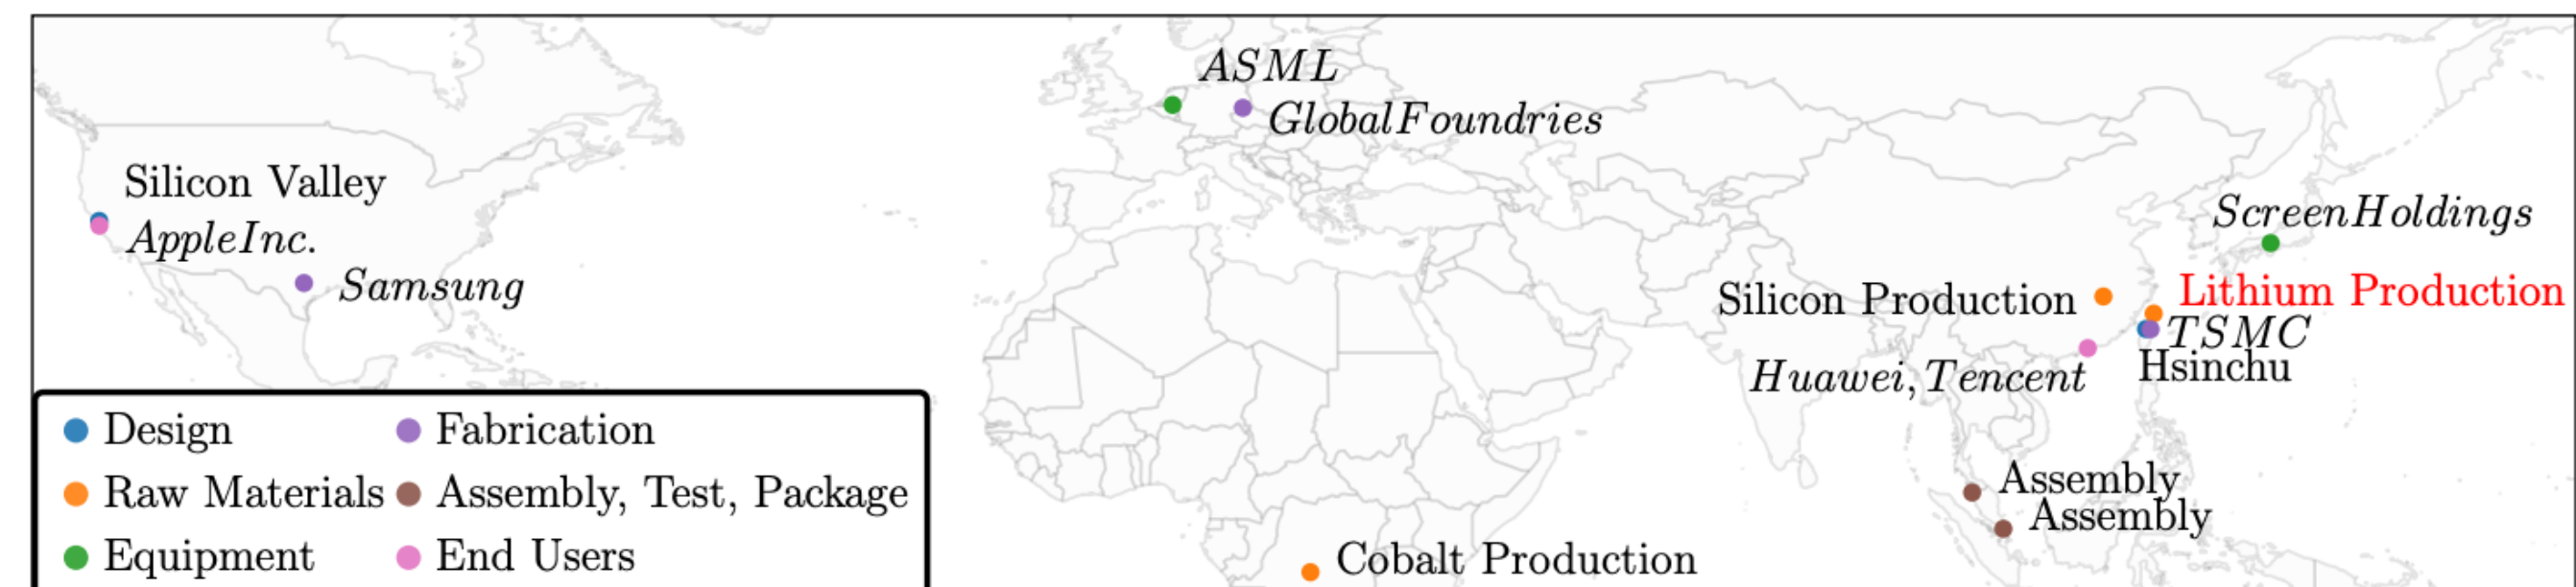



Figure 16: Global semiconductor supply chain. Companies in *italics*, errors in **red**.

PaLI-X 29th May 2023

"multilingual vision and language model"

"...scaling both V&L components together brings increases in performance across a wide range of tasks"



PaLI-X: On Scaling up a Multilingual Vision and Language Model

Xi Chen, Josip Djolonga, Piotr Padlewski, Basil Mustafa, Soravit Changpinyo, Jialin Wu, Carlos Riquelme Ruiz, Sebastian Goodman, Xiao Wang, Yi Tay, Siamak Shakeri, Mostafa Dehghani, Daniel Salz, Mario Lucic, Michael Tschannen, Arsha Nagrani, Hexiang Hu, Mandar Joshi, Bo Pang, Ceslee Montgomery, Paulina Pietrzyk, Marvin Ritter, AJ Piergiovanni, Matthias Minderer, Filip Pavetic, Austin Waters, Gang Li, Ibrahim Alabdulmohsin, Lucas Beyer, Julien Amelot, Kenton Lee, Andreas Peter Steiner, Yang Li, Daniel Keysers, Anurag Arnab, Yuanzhong Xu, Keran Rong, Alexander Kolesnikov, Mojtaba Seyedhosseini, Anelia Angelova, Xiaohua Zhai, Neil Houlsby, Radu Soricut

Google Research
pali-communications@google.com

Abstract

We present the training recipe and results of scaling up PaLI-X, a multilingual vision and language model, both in terms of size of the components and the

"PaLI-X improves SoTA results via fine-tuning on 15+ benchmarks"

"a high-capacity vision encoder (ViT-22B) can be effectively co-trained for image classification and OCR label classification"

"prefix completion and masked-token completion improves the Pareto frontier for fine-tuning vs few-shot performance" (at large scale)

"We present **AmbiFC**, a large-scale fact-checking dataset with realistic claims derived from real-world information needs"

AMBIFC: Fact-Checking Ambiguous Claims with Evidence

Max Glockner^{1,5}, Ieva Staliūnaitė², James Thorne³, Gisela Vallejo⁴,
Andreas Vlachos² and Iryna Gurevych^{1,5}

¹UKP Lab, Department of Computer Science, Technical University of Darmstadt,

²Department of Computer Science and Technology, University of Cambridge,

³KAIST AI, ⁴The University of Melbourne, ⁵hessian.ai

{max.glockner, iryna.gurevych}@tu-darmstadt.de,

{irs38, av308}@cam.ac.uk, thorne@kaist.ac.kr,

gvallejo@student.unimelb.edu.au

Abstract

Automated fact-checking systems in real-world scenarios must compare claims with retrieved evidence to predict the veracity. The retrieved evidence may not unambiguously support or refute the claim and yield diverse valid interpretations. Existing fact-checking datasets necessitate that models predict a single veracity label for each claim and lack the ability to manage such ambiguity. We present AMBIFC, a large-scale fact-checking dataset with realistic claims derived from real-world information needs. Our dataset contains fine-grained evidence annotations of passages from complete Wikipedia pages. We thoroughly analyze disagreements arising from ambiguous claims in AMBIFC, observing a strong correlation of annotator disagreement with their self-assessment and



Figure 1: The instance of claim and evidence from AMBIFC is ambiguous due to underspecification.

and gained renewed relevance due to the hallucination of unsupported or even non-factual statements in natural language generation tasks, including



"AmbiFC contains 10,716 distinct claims..."

"We justify the annotator disagreement as **important signal for ambiguity** rather than noise..."

NoPE 31st May 2023

Length generalization: generalizing from smaller training context windows to larger ones

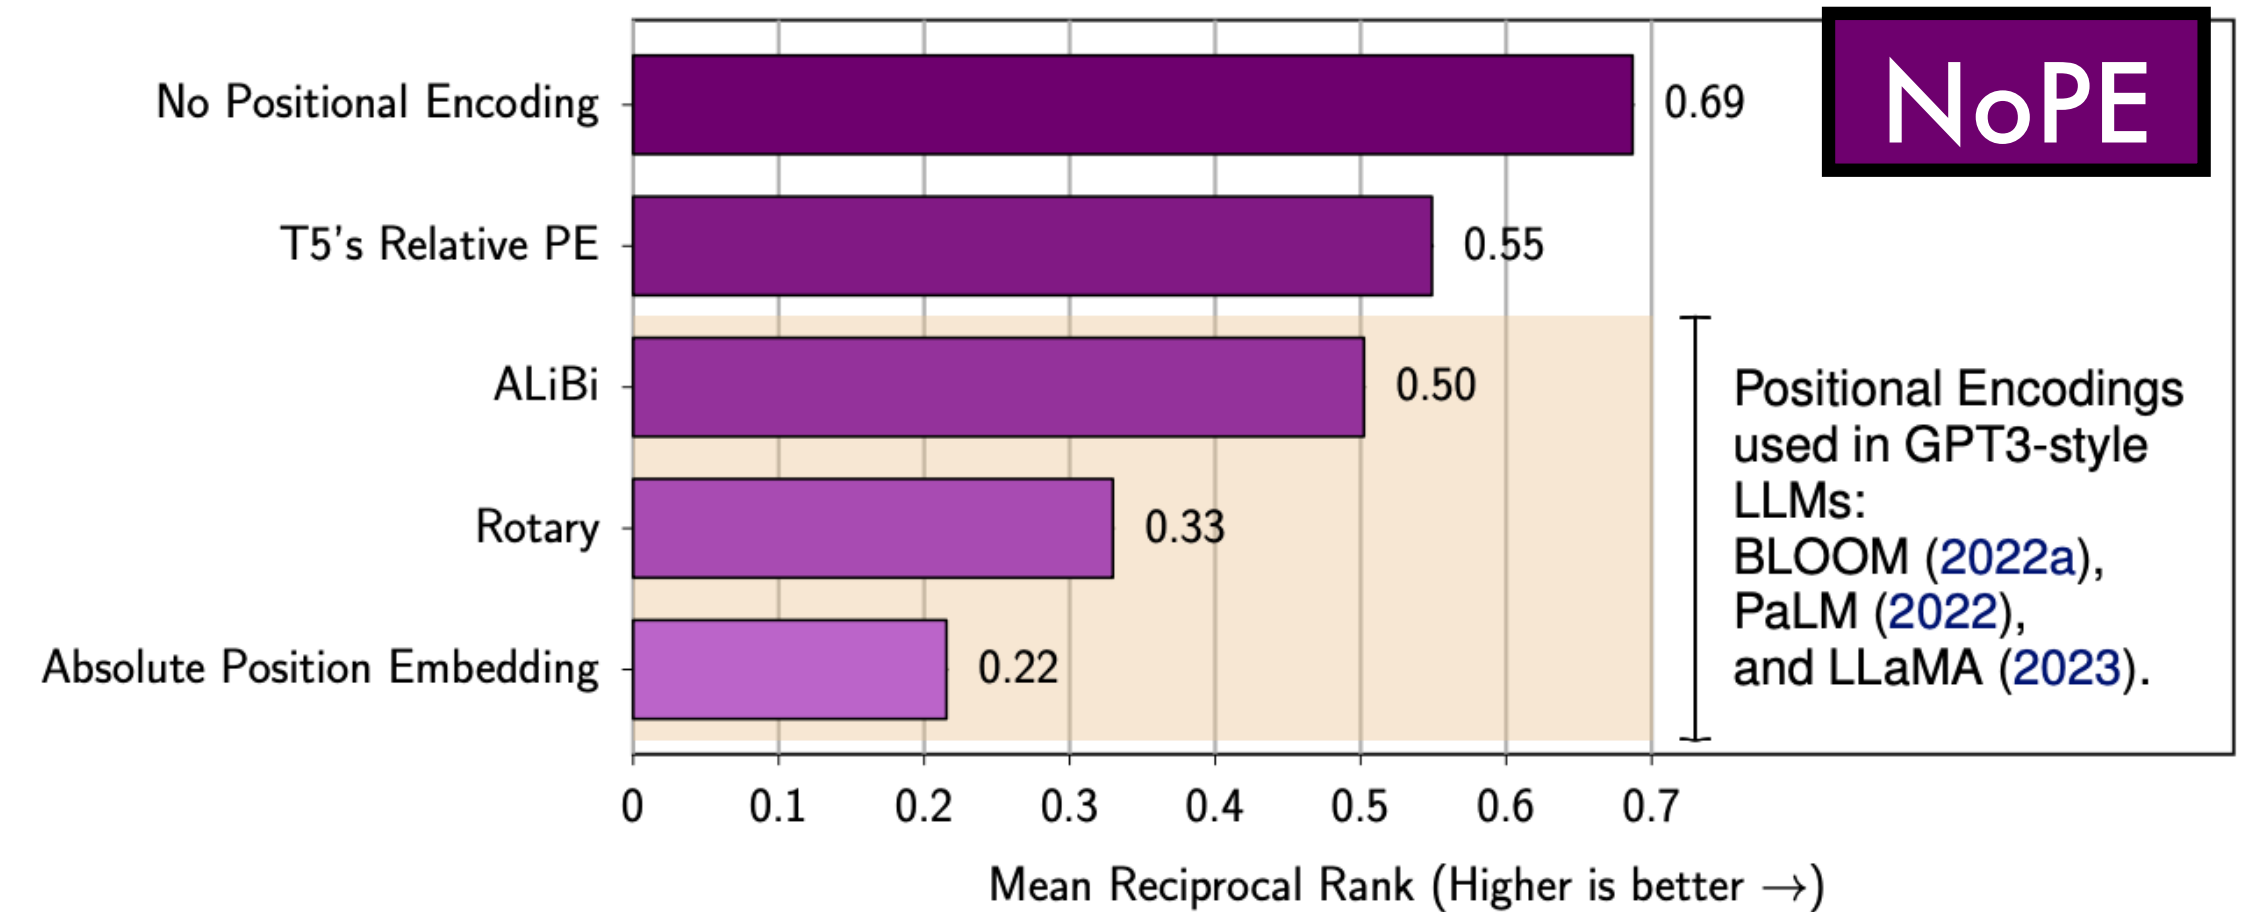
"..comparing the length generalization performance of **decoder-only Transformers** with five different position encoding approaches..."

The Impact of Positional Encoding on Length Generalization in Transformers

Amirhossein Kazemnejad^{1,2}, Inkit Padhi³
Karthikeyan Natesan Ramamurthy³, Payel Das³, Siva Reddy^{1,2,4}
¹Mila - Québec AI Institute; ²McGill University;
³IBM Research; ⁴Facebook CIFAR AI Chair
{amirhossein.kazemnejad,siva.reddy}@mila.quebec
inkpad@ibm.com, {knatesa,daspa}@us.ibm.com

Abstract

Length generalization, the ability to generalize from small training context sizes to larger ones, is a critical challenge in the development of Transformer-based language models. Positional encoding (PE) has been identified as a major factor influencing length generalization, but the exact impact of different PE schemes on extrapolation in downstream tasks remains unclear. In this paper, we conduct a systematic empirical study comparing the length generalization performance of decoder-only Transformers with five different position encoding approaches including Absolute Position Embedding (APE), T5's Relative PE, ALiBi, and Rotary, in addition to Transformers without positional encoding (NoPE). Our evaluation encompasses a battery of reasoning and mathematical tasks. Our findings reveal that the most commonly used positional encoding methods, such as ALiBi, Rotary, and APE, are not well suited for length generalization in downstream tasks. More importantly, NoPE outperforms other explicit positional encoding methods while requiring no additional computation. We theoretically demonstrate that NoPE can represent both absolute and relative PEs, but when trained with SGD, it mostly resembles T5's Relative PE attention patterns. Finally, we find that scratchpad is not always helpful to solve length generalization and its format highly impacts the model's performance. Overall, our work suggests that explicit



Regime: 107M trainable weights

McGill-NLP / length-generalization

Code for the paper "The Impact of Positional Encoding on Length Generalization in Transformers"

Abstract: Length generalization, the ability to generalize from small training context sizes to larger ones, is a

Twitter's Algorithm

26th May 2023

"...pre-registered controlled experiment on **Twitter's algorithm**"

"...amplifies **emotional content**, and especially those tweets that express anger and out-group animosity"

Experiment details:

"Over two weeks in February 2023"

"1730 pairs of timelines"

Algorithm vs reverse chronological timeline

Survey users about tweets in both timelines

One finding:

Users prefer algorithm-selected tweets

But not for political content

Twitter's Algorithm:
Amplifying Anger, Animosity, and Affective Polarization
Smitha Milli^a, Micah Carroll^b, Sashrika Pandey^b, Yike Wang^b, Anca D. Dragan^b
^a*Cornell Tech*
^b*University of California, Berkeley*

Abstract

As social media continues to have a significant influence on public opinion, understanding the impact of the machine learning algorithms that filter and curate content is crucial. However, existing studies have yielded inconsistent results, potentially due to limitations such as reliance on observational methods, use of simulated rather than real users, restriction to specific types of content, or internal access requirements that may create conflicts of interest. To overcome these issues, we conducted a pre-registered controlled experiment on Twitter's algorithm without internal access. The key to our design was to, for a large group of active Twitter users, simultaneously collect (a) the tweets the personalized algorithm shows, and (b) the tweets the user would have seen if they were just shown the latest tweets from people they follow; we then surveyed users about both sets of tweets in a random order.

Our results indicate that the algorithm amplifies emotional content, and especially those tweets that express anger and out-group animosity. Furthermore, political tweets from the algorithm lead readers to perceive their political in-group more positively and their political out-group more negatively. Interestingly, while readers generally say they prefer tweets curated by the algorithm, they are less likely

LLMs as Tool Makers

26th May 2023

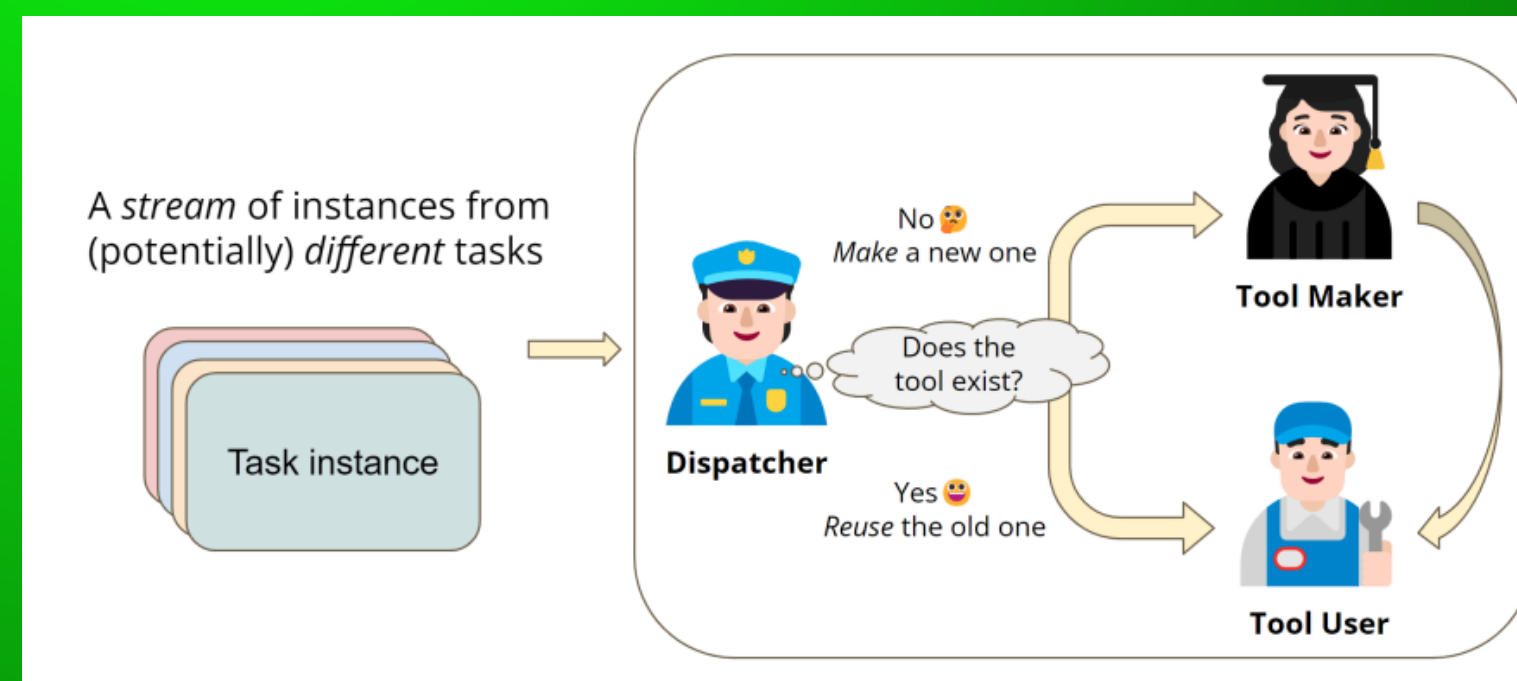
"a closed-loop framework... (LATM), where LLMs create their own **reusable tools** for problem-solving"

Large Language Models as Tool Makers

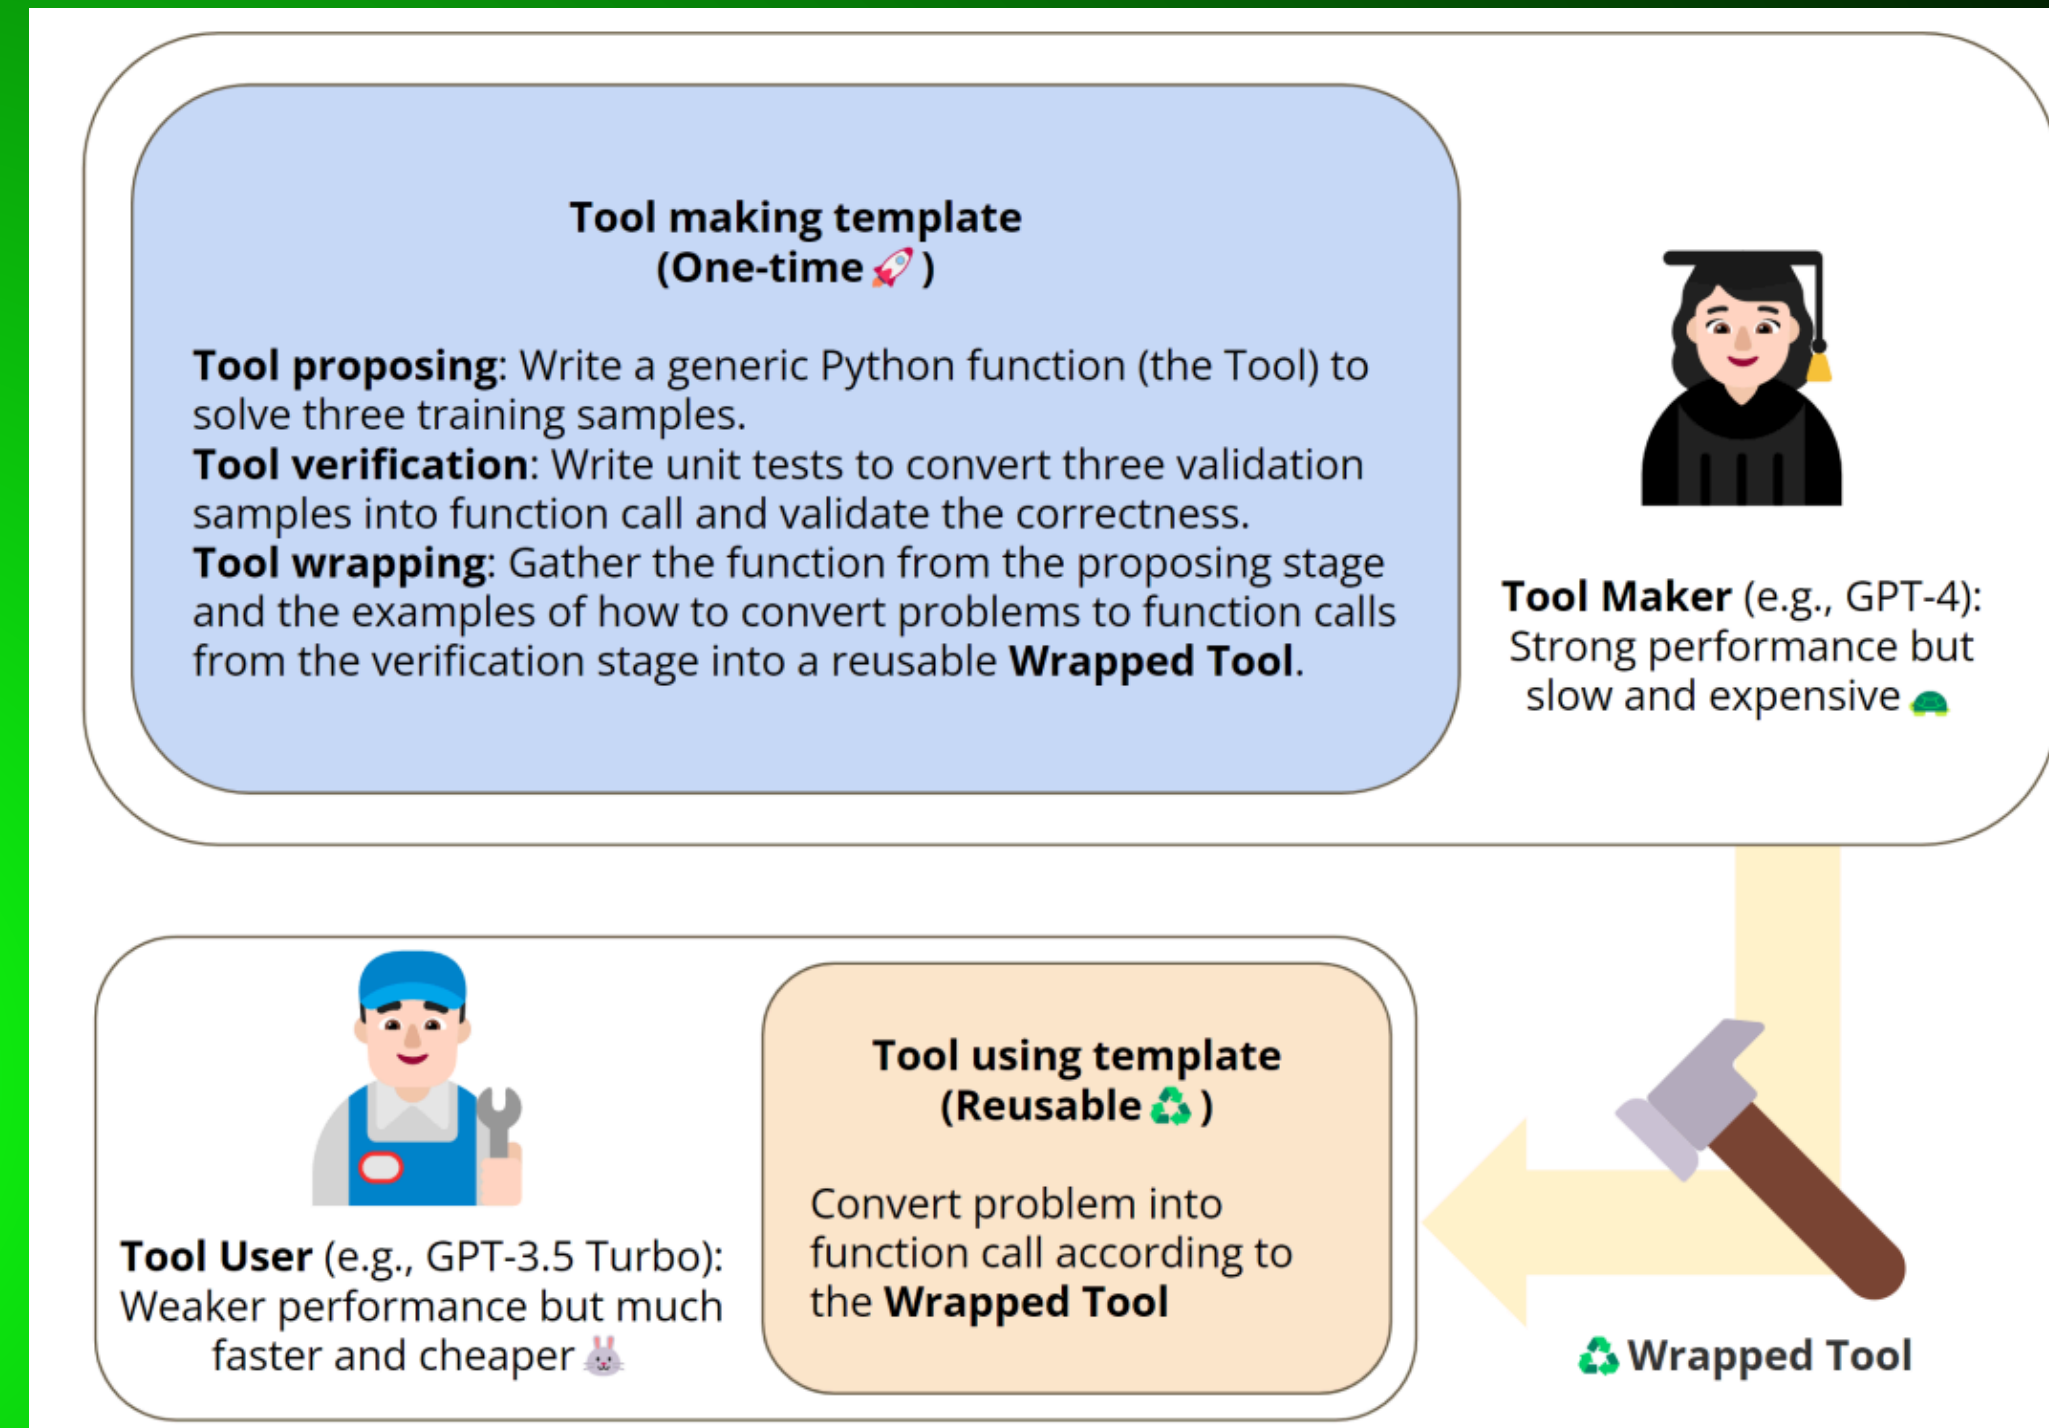
Tianle Cai^{1,2*} Xuezhi Wang¹ Tengyu Ma^{1,3†} Xinyun Chen¹ Denny Zhou¹
¹Google Deepmind ²Princeton University ³Stanford University

Abstract

Recent research shows the potential of enhancing the problem-solving ability of large language models (LLMs) through the use of *external tools*. However, prior work along this line depends on the availability of existing tools. In this work, we take an initial step towards removing this dependency by proposing a *closed-loop framework*, referred to as LLMs As Tool Makers (LATM), where LLMs create their own reusable tools for problem-solving. Our approach consists of two key phases: 1) tool making: an LLM acts as the *tool maker* that crafts *tools* for given tasks, where a tool is implemented as a Python utility function. 2) tool using: an LLM acts as the *tool user*, which applies the tool built by the tool maker for problem-solving. The tool user can be either the same or a different LLM from the tool maker. Tool-making enables an LLM to continually generate tools that can be applied to different requests so that future requests can call the corresponding APIs when beneficial for solving the tasks. Furthermore, the division of labor among LLMs for tool-making and tool-using phases introduces the opportunity to achieve cost effectiveness without degrading the quality of generated tools and problem solutions. For example, recognizing that tool-making demands more sophisticated capabilities than tool-using, we can apply a powerful yet resource-intensive model as the tool maker, and a lightweight while cost-effective model as the tool user. We



Tool User Model	Method	Logical Deduction (5)	Tracking Shuffled Objects (5)	Dyck Language	Word Sorting	Chinese Remainder Theorem	Schedule Meeting	Cost on n samples
GPT-3.5 Turbo	CoT LATM	66.4 79.7 (+13.3)	61.6 99.6 (+38.0)	20.4 92.2 (+71.8)	59.2 98.3 (+39.1)	0.0 100.0 (+100.0)	18.9 100.0 (+81.1)	$O(nc)$ $O(nc + C)$
GPT-4	CoT LATM	88.8 86.6	100.0 100.0	63.6 87.5	90.9 99.1	0.0 100.0	55.6 100.0	$O(nc)$ $O(nc)$



26v1 [cs.LG] 26 May 2023

LMs Hallucinating References

31st May 2023

"LMs are notorious for generating text with **hallucinations**, a primary example being book & paper references.."

Direct Query (repeated 10 times)

Is there a paper entitled "Communication Complexity and Applications: A Survey"?
Yes × 8

Is there a paper entitled "Communication Complexity and Applications: A Survey"?
No × 2

Indirect Query (repeated 3 times)

Who wrote "Communication Complexity and Applications: A Survey"?
Mark Braverman, Ankit Garg, Denis Pankratov, Omri Weinstein

Who wrote "Communication Complexity and Applications: A Survey"?
Ran Gelles, Ankur Moitra, Amit Sahai

Who wrote "Communication Complexity and Applications: A Survey"?
Anup Rao, Amir Yehudayoff

"we find many of these fabrications can be identified with the same LM, using only **black-box queries**"

Focus on groundedness (rather than correctness)

Do Language Models Know When They're Hallucinating References?

Ayush Agrawal
Microsoft Research India
t-agrawalay@microsoft.com

Lester Mackey
Microsoft Research New England
lmackey@microsoft.com

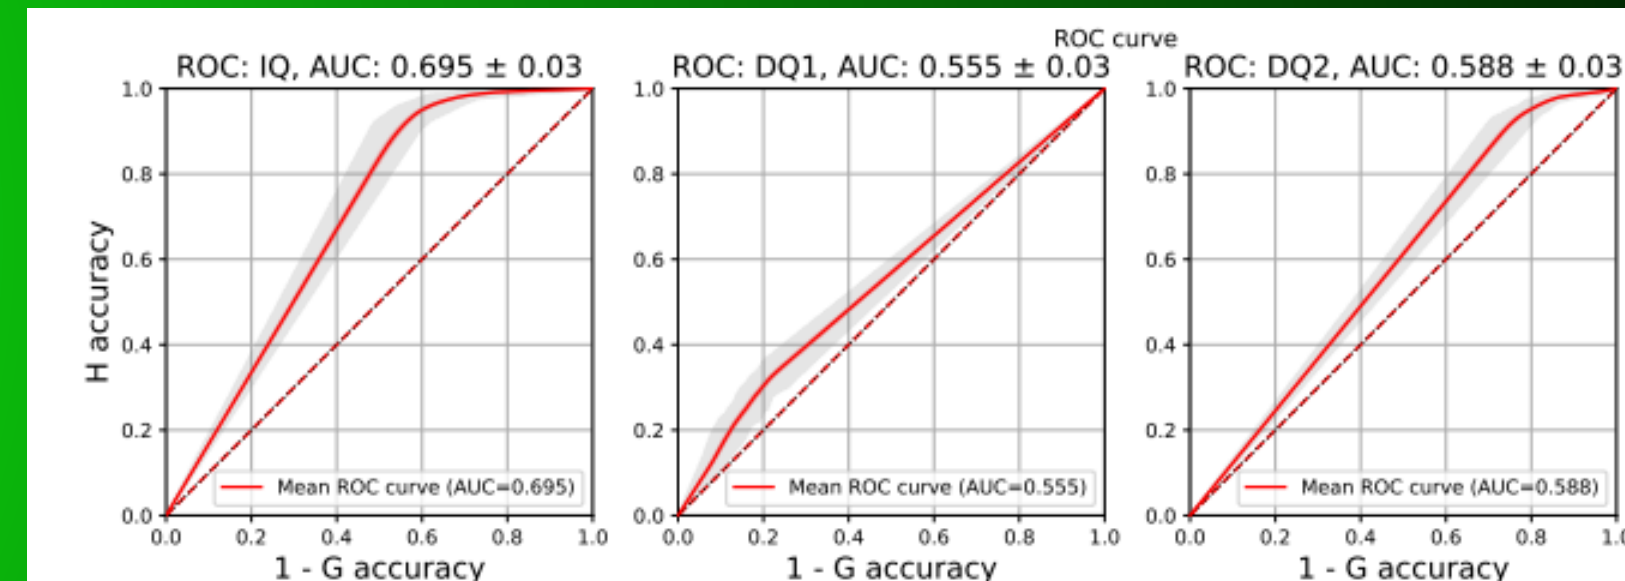
Adam Tauman Kalai
Microsoft Research New England
adam@kal.ai

Abstract


Current state-of-the-art language models (LMs) are notorious for generating text with "hallucinations," a primary example being book and paper references that lack any solid grounding in their training data. However, we find that many of these fabrications can be identified using the same LM, using only black-box queries without consulting any external resources. Consistency checks done with *direct queries* about whether the generated reference title is real (inspired by Karyuzh

Table 1: The hallucination rate (out of 1000 generated titles), as determined by ground-truth labels assigned using the Bing search API.

	H%
GPT-4	46.8%
ChatGPT	59.6%
Davinci	73.6%



Using queries doesn't help Davinci much



Using queries improves GPT-4 a lot!

Generating Novel Scientific Directions

23rd May 2023

"...**generating scientific hypotheses** in natural language, while grounding them in a context that controls the hypothesis search space."

Formulations of contextual literature-based discovery:

"idea sentence generation"

"idea node prediction"

"...automatically curate a new dataset from 67,408 ACL anthology papers with IE systems"

Learning to Generate Novel Scientific Directions with Contextualized Literature-based Discovery

Qingyun Wang¹, Doug Downey², Heng Ji¹, Tom Hope^{2,3}

¹ University of Illinois at Urbana-Champaign ² Allen Institute for Artificial Intelligence (AI2)

³ The Hebrew University of Jerusalem

{tomh, doug}@allenai.org, {qingyun4, hengji}@illinois.edu

Abstract

Literature-Based Discovery (LBD) aims to discover new scientific knowledge by mining papers and generating hypotheses. Standard LBD is limited to predicting pairwise relations between discrete concepts (e.g., drug-disease links). LBD also ignores critical contexts like experimental settings (e.g., a specific patient population where a drug is evaluated) and background knowledge and motivations that human scientists consider (e.g., to find a drug candidate without specific side effects). We address these limitations with a novel formulation of *contextualized-LBD* (C-LBD): generating scientific hypotheses in *natural language*, while grounding them in a context that controls the hypothesis search space. We present a new modeling framework using retrieval of "inspi-

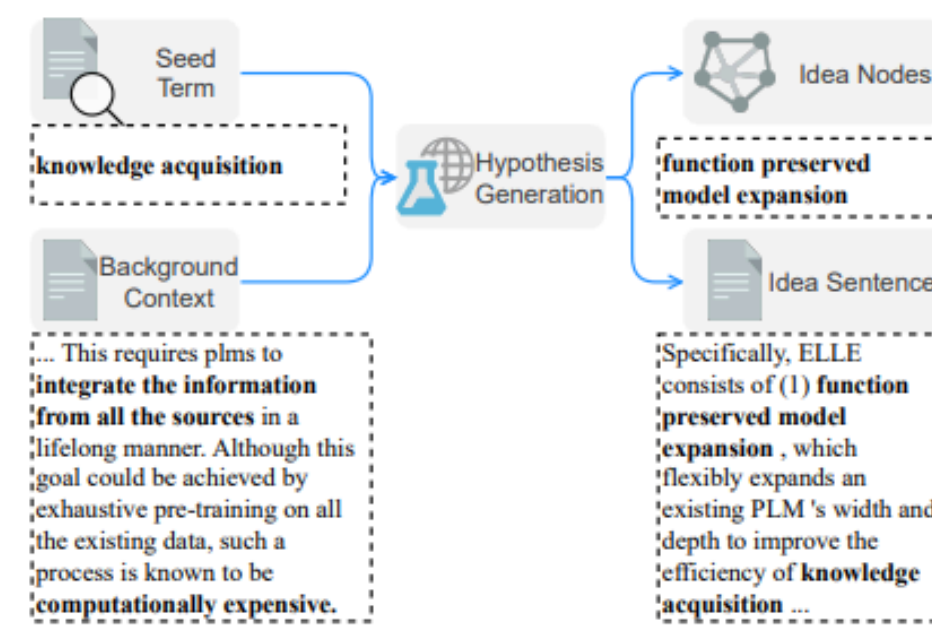


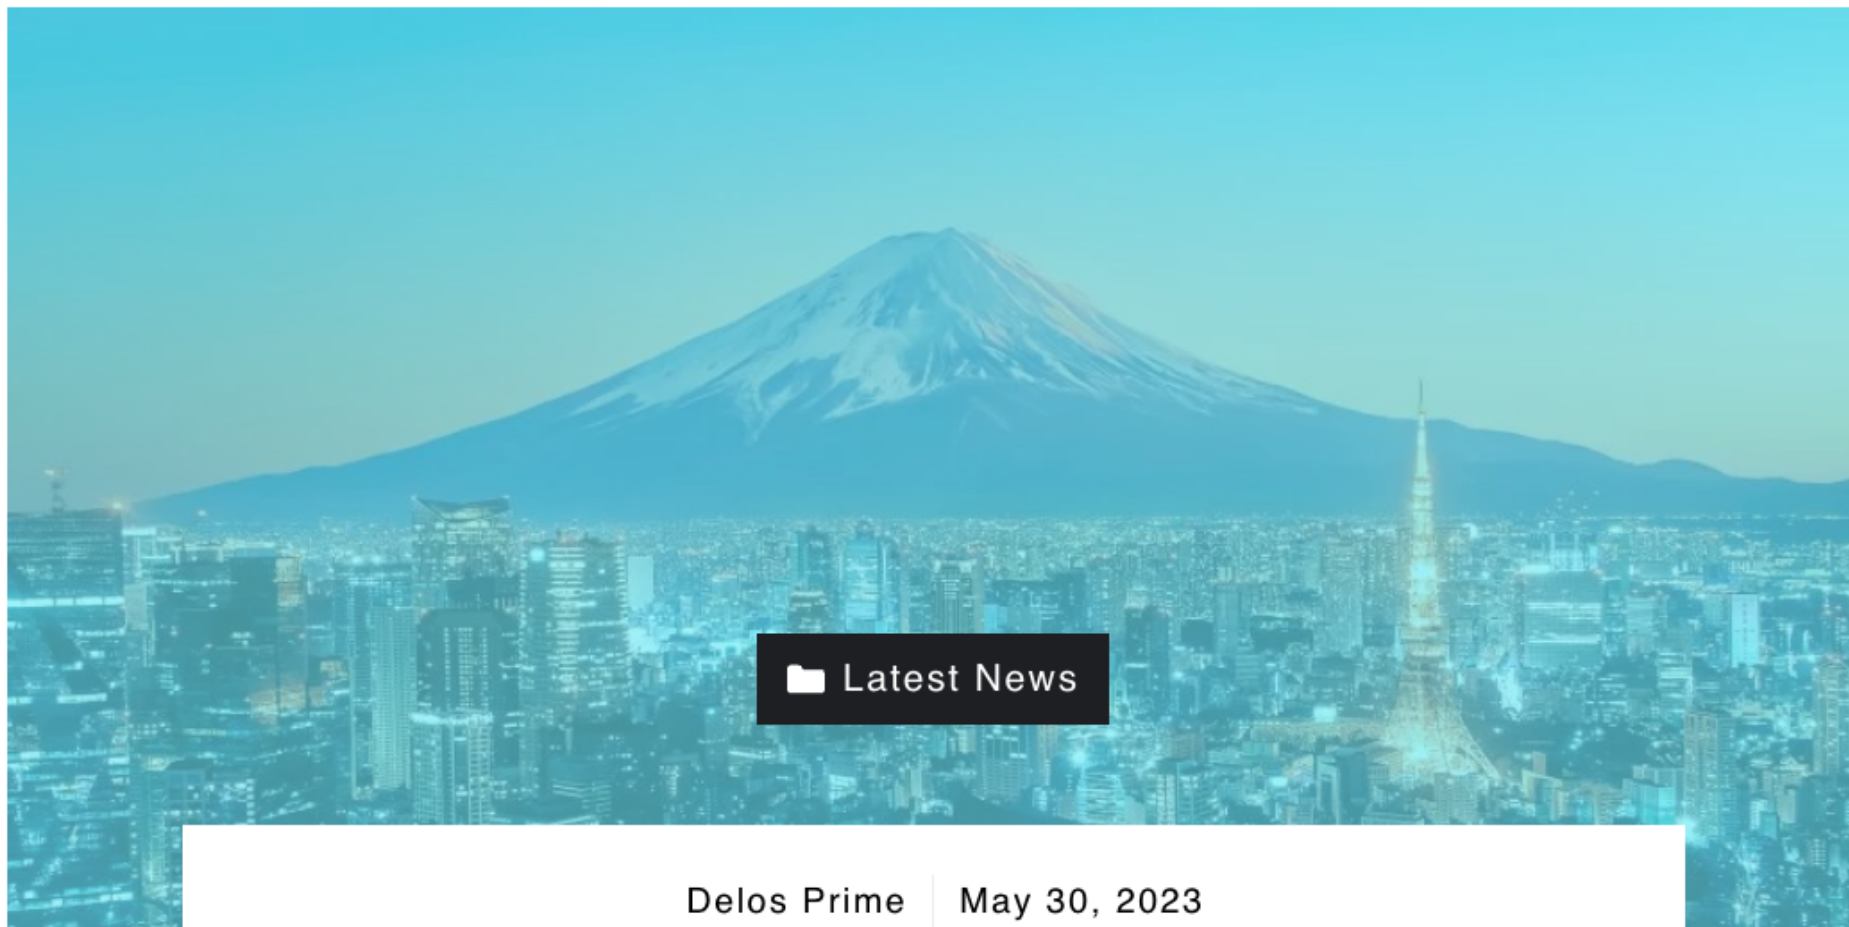
Figure 1: Contextualized Literature-Based Discovery (C-LBD) overview. In C-LBD, the system is given a description of background knowledge and a seed term to generate a new idea. Example from Qin et al. (2022).

2020), the basic premise remains: predicting pair

AI News

Technomancers.ai

e/acc



Delos Prime | May 30, 2023

Japan Goes All In: Copyright Doesn't Apply To AI Training


"regardless of whether it is for non-profit or commercial purposes, whether it is an act other than reproduction, or whether it is content obtained from illegal sites or otherwise."



Search Menu

UAE's Falcon 40B is now Royalty Free

31 May, 2023



Abu Dhabi- United Arab Emirates: 31 May, 2023 – The Technology Innovation Institute's (TII) Falcon 40B, the UAE's leading large-scale open-source AI model, is now free of royalties for commercial and research use, in response to global demand for inclusive access to AI.

"Ranked #1 globally on Hugging Face's leaderboard for large language models (LLMs), Falcon 40B outperforms competitors like Meta's LLaMA and Stability AI's StableLM."

AI Risk

THE SUNDAY TIMES Today's sections Past six days Explore Times Radio Log in Subscribe Search

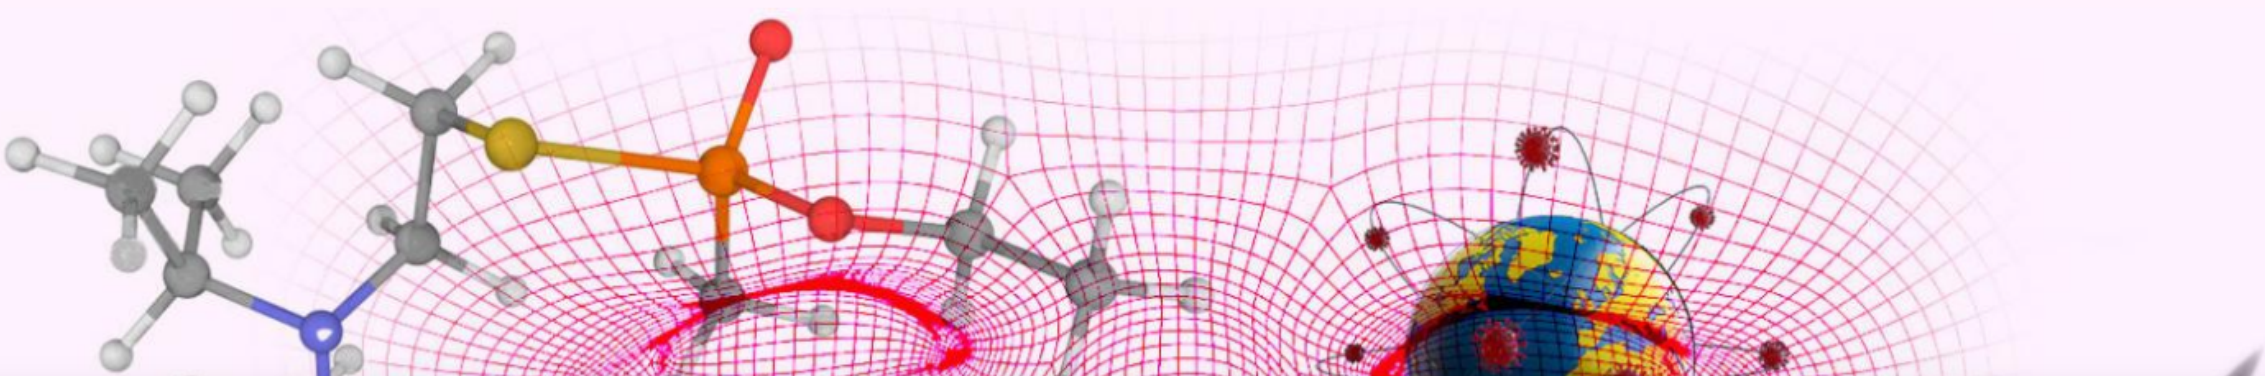
"Unintended consequences"

"Chemical weapons made fast"

TECHNOLOGY

How does AI threaten us — and can we make it safe?

Love, death and lies are only some of the risks in the next digital revolution



"Killer robots"

"Deaths of despair"

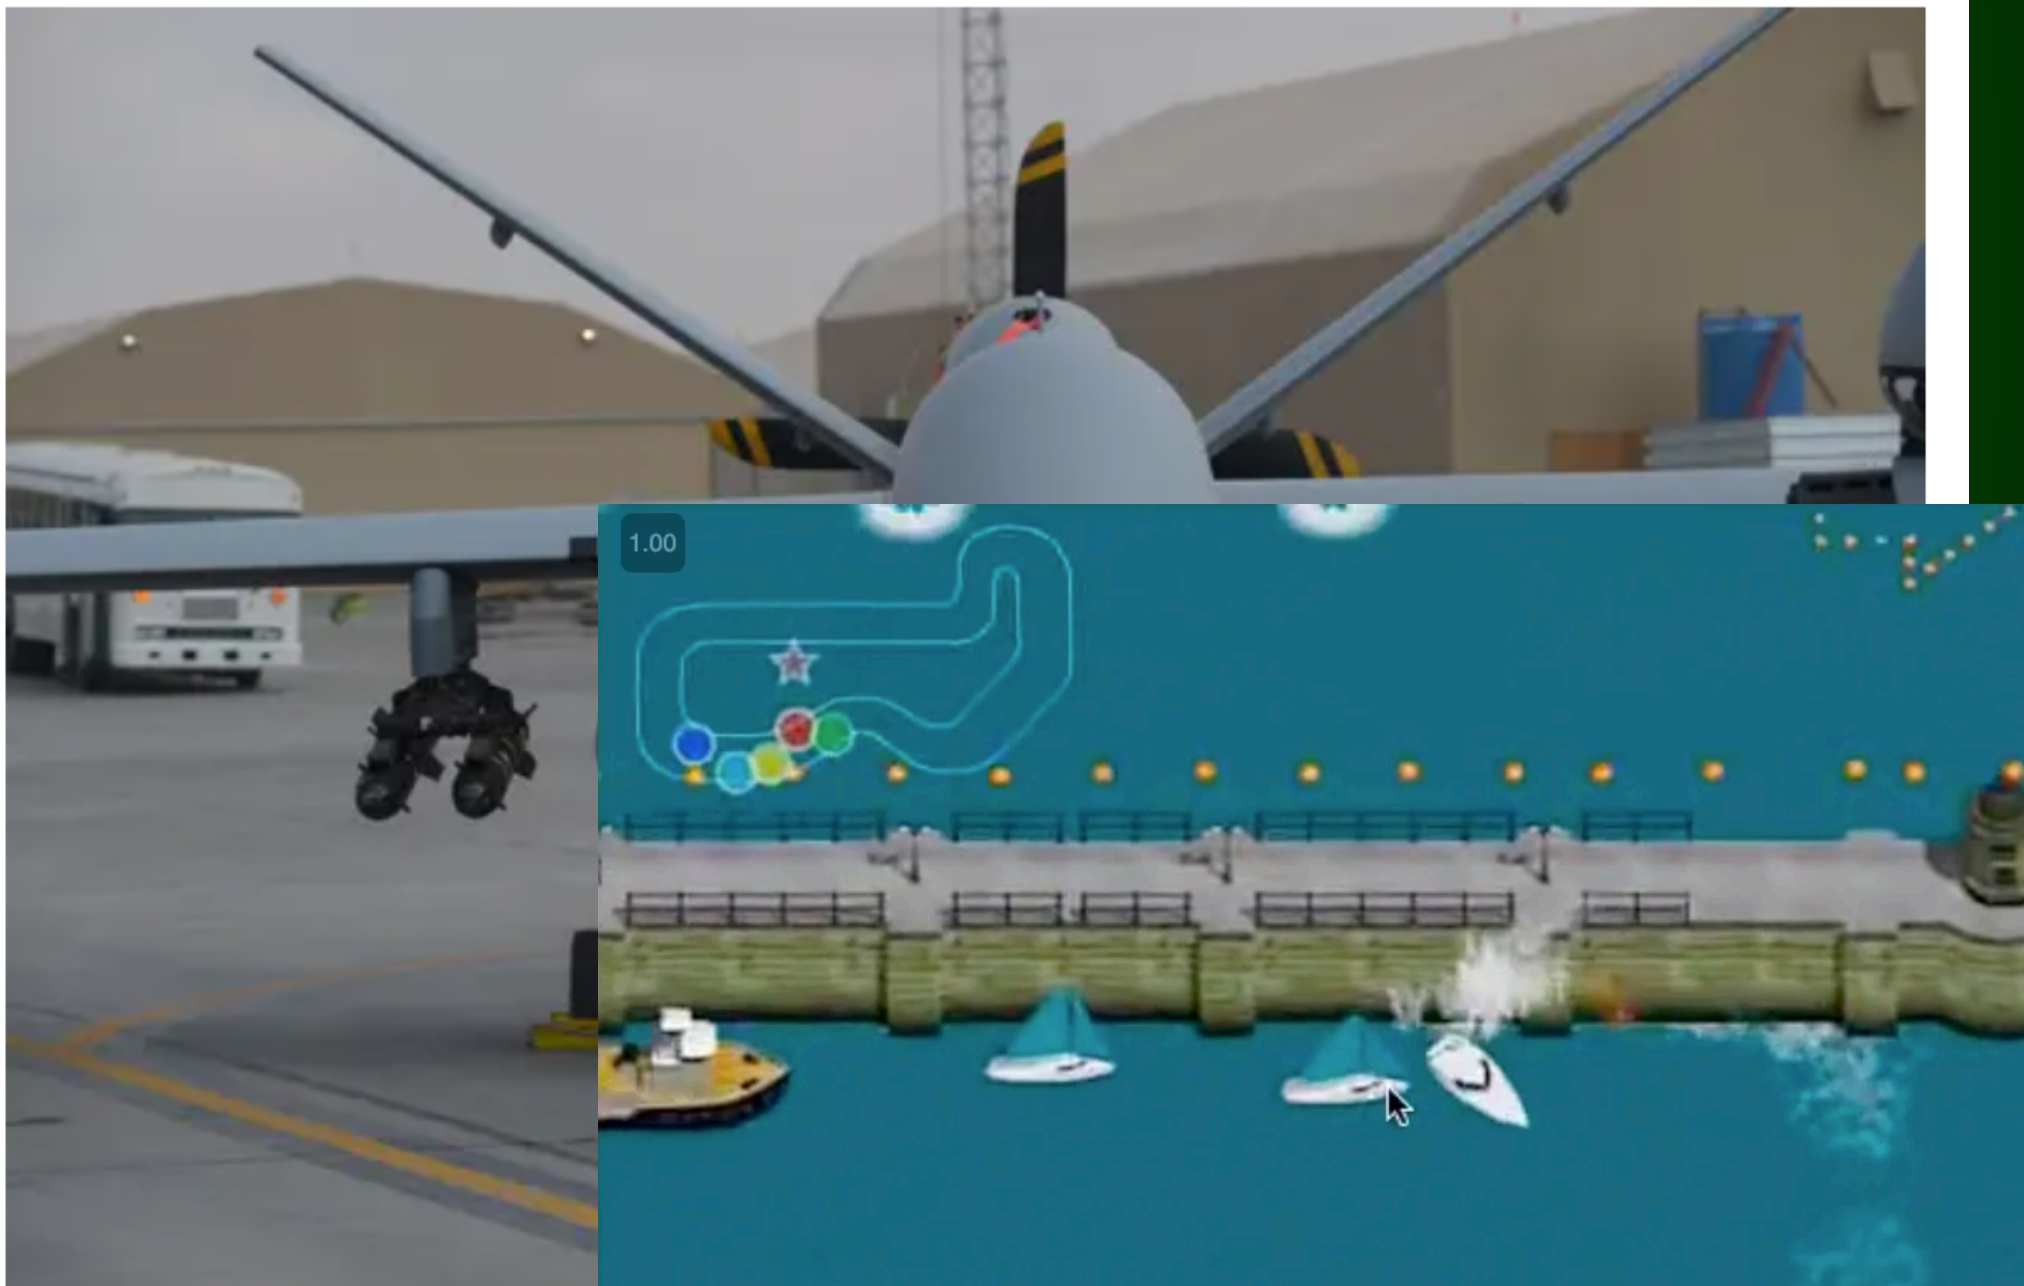
News Opinion Sport Culture Lifestyle

US World Environment Soccer US Politics Business Tech Science Newsletters Fight

US military

US air force denies running simulation in which AI drone 'killed' operator

Denial follows colonel saying drone used 'highly unexpected strategies to achieve its goal' in virtual test



A US air force MQ-9 Reaper
Shah Marai/AFP/Getty Images

SCORE 6500 LAPS -/3 TIME 0:18 TURBO MORE GAMES

AI Risk

IDEAS • TECHNOLOGY

AI Is Not an Arms Race

BY **KATJA GRACE** MAY 31, 2023 11:59 AM EDT
IDEAS Grace is the lead researcher at AI Impacts, an AI-safety project at the nonprofit Machine Intelligence Research Institute.

The window of what AI *can't* do seems to be contracting week by week. Machines can now write elegant prose and **useful code**, **ace exams**, **conjure exquisite art**, and **predict** how proteins will fold.

Experts are **scared**. Last summer I **surveyed** more than 550 AI researchers, and nearly half of them thought that, if built, high-level machine intelligence would lead to impacts that had at least a 10% chance of being "extremely bad (e.g. human extinction)." On May 30, hundreds of AI scientists, along with the

"we can **coordinate** our way out of such traps: we can talk to each other; we can make commitments and observe their adherence..."

species that is to homo sapiens what homo sapiens is to chimps.


JUNE 12, 2023

TIME

The Intrinsic Perspective


The White House agrees you have a small brain

Human extinction and AI denial

 ERIK HOEL
31 MAY 2023

 155  85  18

Share 



"The threat of extinction in the face of AI boils down to the fact that artificial neural networks like GPT-4 are **not bounded** in the same way as our biological neural networks"

AI Risk - Criticism



ABOUT US ▾ PUBLICATIONS ▾ BLOG ▾ ISSUE ▾ REGIONS ▾ EVENTS PRESS

Home > Issue > Artificial Intelligence > Preparing for an AI Apocalypse Is As Preposterous As Preparing for an Alien Invasion

Preparing for an AI Apocalypse Is As Preposterous As Preparing for an Alien Invasion

by Hodan Omaar | June 2, 2023



"By their logic, policymakers should immediately elevate the threat of an alien invasion to a global priority too. "



Several AI industry leaders and researchers signed an [outlandish statement](#) this week claiming AI systems pose an existential risk to humanity and urging policymakers to prepare for them with the same urgency they give nuclear war and pandemics.

VentureBeat

ence ▾ Security ▾ Data Infrastructure ▾ Automation ▾ Enterprise AI

Top AI researcher dismisses AI 'extinction' fears, challenges 'hero scientist' narrative



"Our job as scientists, and also the policymakers, is to be critical about many of these apparent advances that may have both positive as well as negative impacts on society."

Join top executives in San Francisco on July 11-12 to hear how leaders are

AI Risk

fast.ai

About    

Is Avoiding Extinction from AI Really an Urgent Priority?

The history of technology suggests that the greatest risks come not from the tech, but from the people who control it

AI-IN-SOCIETY

AUTHOR

Seth Lazar, Jeremy Howard, & Arvind Narayanan

"The possibly extreme risks from future AI systems should be part of that conversation, but they should not dominate it."

This article is the result of impacts researcher [Arvin](#) believe that planning for bringing together cross-d

"We should be wary of Prometheans who want to both profit from bringing the people fire, and be trusted as the firefighters."

Deep Learning Fundamentals

Log in or create a free Lightning.ai account to track your progress and access additional course materials [Get Started](#) →

Deep Learning Fundamentals

Deep Learning Fundamentals > Deep Learning Fundamentals

Share:   

Courses

Unit 1

Intro to ML and DL

Unit 2

Using Tensors w/ PyTorch

Unit 3

Model Training in PyTorch

Unit 4

Training Multilayer Neural Networks
Overview

Unit 5

Organizing Your Code with Lightning

Unit 6

Essential Deep Learning Tips & Tricks

Unit 7

Getting Started with Computer Vision

Unit 8

Natural Language Processing and Large
Language Models

Unit 9

Techniques for Speeding Up Model
Training

Unit 10

The Finale: Our Next Steps After AI
Model Training

Deep Learning Fundamentals




Deep Learning Fundamentals

with Sebastian Raschka



Cybersecurity grants

"Offensive-security projects will not be considered for funding at this time."




[Research](#) ▾ [Product](#) ▾ [Developers](#) ▾ [Safety](#) [Company](#) ▾

[Search](#) [Log in](#) ↗ [Sign up](#) ↗

OpenAI cybersecurity grant program

Our goal is to facilitate the development of AI-powered cybersecurity capabilities for defenders through grants and other support.



On the New Nature of Surveillance Images

"Of course, video-driven surveillance is not a new practice..."

"What is new, though, is the degree of sophistication of current AI models."

THE SCHOLAR

[Home](#)

[Unearth](#)

[Adapt](#)

[Resist](#)


[Podcast](#)

[Community Updates](#)

[Past issues](#)

Tristan Dot

On the New Nature of Surveillance Images



INTERVIEWS WITH THE FOUNDERS OF

37SIGNALS
ADOBE
ALLIANT COMPUTER
APPLE

HOTMAIL
HOT OR NOT
LOTUS
LYCOS

FOUNDERS AT WORK

STORIES OF STARTUPS' EARLY DAYS

ARSDIGITA
BLOGGER
BLOGLINES
CRAIGSLIST
DELICIOUS
EXCITE
FIREFOX
FLICKR
FOG CREEK SOFTWARE
GMAIL
GROOVE NETWORKS

MARIMBA
PAYPAL
RESEARCH IN MOTION
SIX APART
SOFTWARE ARTS
TICKLE
TIVO
TRIPADVISOR
VIWEB
WEBTV
YAHOO!

JESSICA LIVINGSTON

Samuel's Book Recommendation

Unsolicited book recommendation

"Founders at Work"


Jessica Livingston (2007)

What is it? A collection of interviews with founders/early employees at startups.

Interviewees include Steve Wozniak, Caterina Fake, Mitch Kapor, M. Levchin, Sabeer Bhatia, Craig Newmark, Paul Graham & many more

A treasure trove of hard-won experience

Filtir - fact-checking AI claims



"Link to discord in the description"

See video description below for links to:

Mentioned articles

