Morgan Lewis

ARTIFICIAL INTELLIGENCE (AI) BOOT CAMP

Large Language Models, Open Source, and Ethical/Responsible AI: An IP Perspective

November 29, 2022

Kannan Narayanan

Host

Andrew J. Gray IV

Presenter

Kannan Narayanan

Morgan Lewis

Short Background

Kannan Narayanan
Silicon Valley

1994 – 1998: Computer Science & Engineering Bachelors (IIT)
1998 – 1999: Computer Science Masters (U Pitt)
2000 – 2017: R&D, Engineering, and Leadership Roles
(Cisco Systems, Intel, AMD, and Startups)
2014 – 2018: J.D. at Santa Clara University
2018 – 2022: IP Associate, Morgan Lewis

Presentation Overview

Part 1: Introduction to LLMs

1. Background
2. Recent developments in LLMs
3. Example Applications

Part 2: Open Source, Patents, Publications and Copyrights

1. Impact of Open Source and License Frameworks
2. Designing around patents, patent protection
3. Publications – Authors and Inventorship Issues
4. Copyright protection for LLMs

Part 3: Ethical and Responsible LLMs

1. Bias issues with LLMs and Examples
2. Sources of Bias and Solutions
3. Ethical AI landscape
4. Conclusion

Part 1: Large Language Models – Introduction

1. Background

2. Recent developments in LLMs

3. Example Applications

Language Models: What are they?

- A language model (LM) is a probability distribution over sequences of tokens.
  - Suppose a vocabulary consists of the words {ate, ball, cheese, mouse, the}.
  - A language model might assign:
    - Probability(the, mouse, ate, the, cheese) = 0.05,
    - Probability(the, cheese, ate, the, mouse) = 0.01,
    - Probability() = 0.0001, …
- Generation samples a sequence from the language model for a probability.
- Autoregressive language models allow efficient generation of a completion given a prompt; a temperature parameter can control randomness
- Entropy (1948), N-gram models (1970), neural language models (2003)

Source: https://stanford-cs324.github.io/winter2022

Recent Developments in Language Models

- Recurrent Neural Networks (RNNs)
  - Conditional distribution of a token depends on entire context
  - Hard to train
- Transformers (2017)
  - Fixed context length n (2048, for GPT-3)
  - Easier to train, exploits GPU parallelism
- Masked language models (e.g., BERT and RoBERTa)
- Emergence (scaling up), stand-alone capabilities

Source: https://stanford-cs324.github.io/winter2022

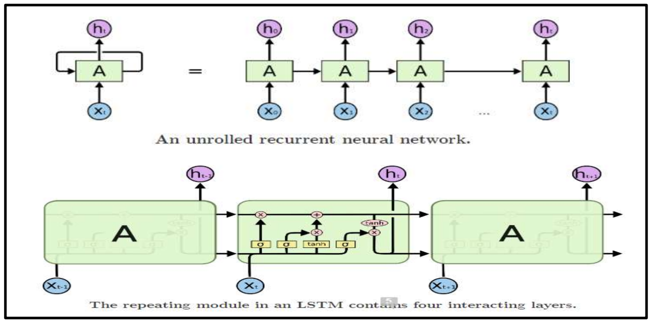

An unrolled recurrent neural network.

The repeating module in an LSTM contains four interacting layers.

Source: http://colah.github.io/posts/2015-08-Understanding-LSTMs/

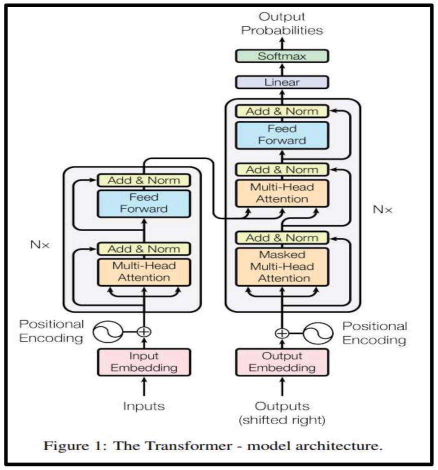

Figure 1: The Transformer - model architecture.

Source: https://arxiv.org/abs/1706.03762

Lifecycle of a Language Model

1. Collect training data (e.g., Common Crawl).
2. Train a large language model (e.g., GPT-3).
3. Adapt it to downstream tasks (e.g., dialogue).
4. Deploy the language model to users (e.g., customer service chatbot).

Source: https://stanford-cs324.github.io/winter2022/lectures/legality/

Perhaps A Second AI Explosion is Underway

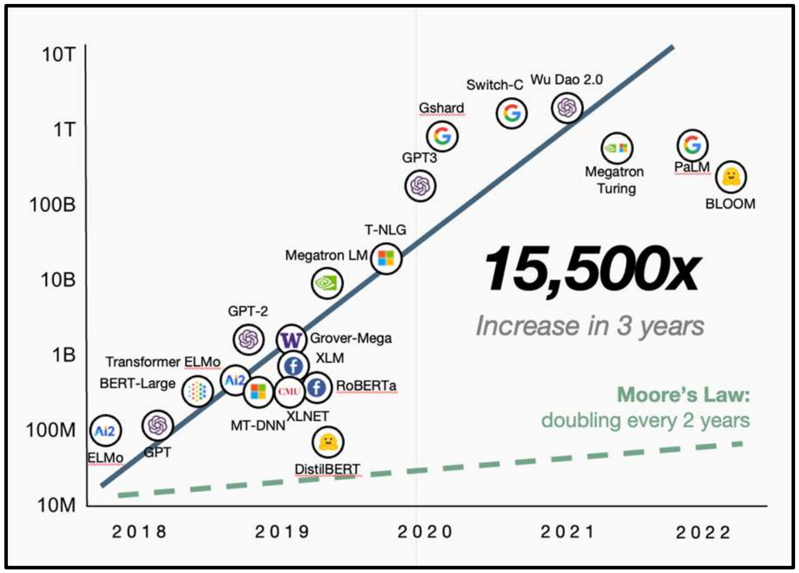

Source: AI 2022: The Explosion, Coatue Venture

Table 1. Selection of known, record-breaking language models based on the Transformer architecture.

Model	Year of release	Company	Number of parameters
GPT	2018	OpenAI	110 million
BERT	2018	Google	340 million
GPT-2	2019	OpenAI	1.5 billion
MegatronLM	2019	NVIDIA	8.3 billion
Turing-NLG	2020	Microsoft	17 billion
GPT-3	2020	OpenAI	175 billion
T6-XXL	2021	Google	1.6 trillion
WuDao 2.0	2021	Beijing Academy of Artificial Intelligence	1.75 trillion

Source: https://journals.sagepub.com/doi/full/10.1177/20539517211047734

Perhaps A Second AI Explosion is Underway

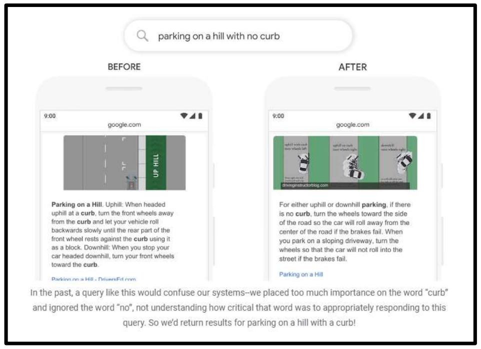

Source: Google blog

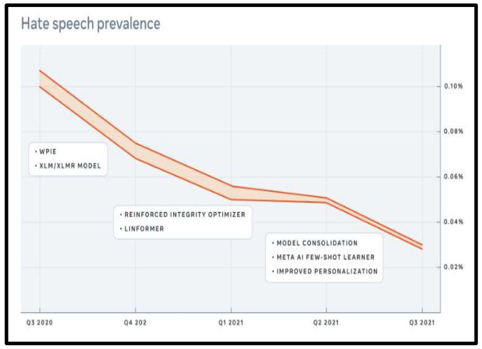

Source: Facebook blog

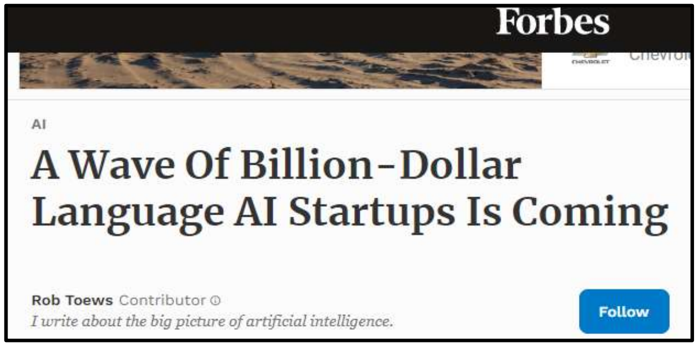

Source: Forbes

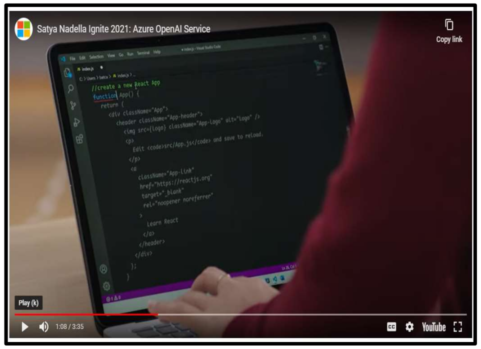

Source: Microsoft AI blog

Source: https://arxiv.org/pdf/2112.04359.pdf

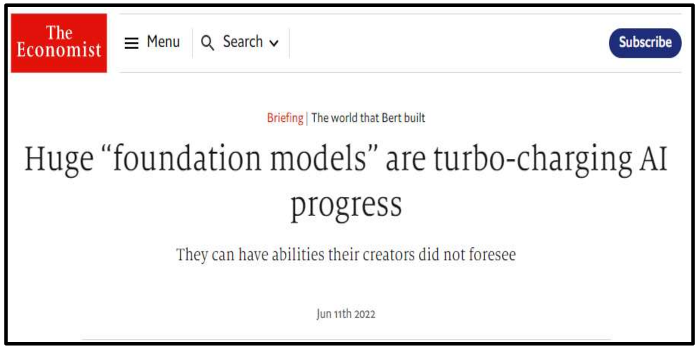

Source: The Economist

Sample Research Publications

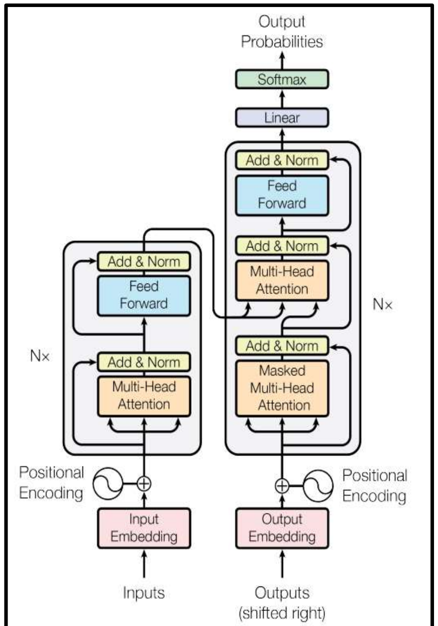

Figure 1: The Transformer - model architecture.

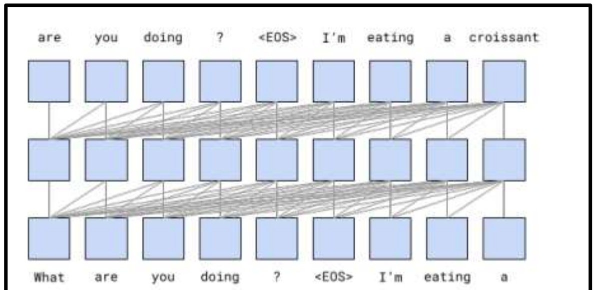

Figure 2: LaMDA pre-training as a language model.

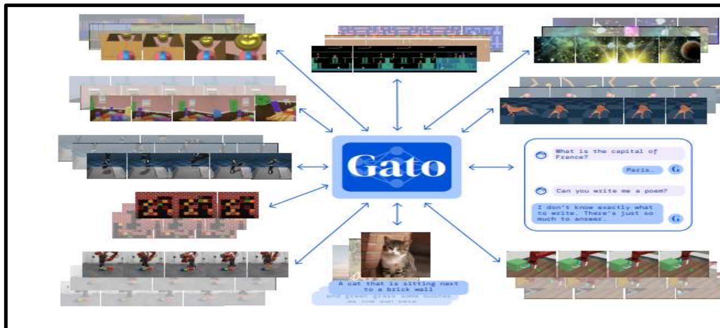

Figure 1: **A generalist agent.** Gato can sense and act with different embodiments across a wide range of environments using a single neural network with the same set of weights. Gato was trained on 604 distinct tasks with varying modalities, observations and action specifications.

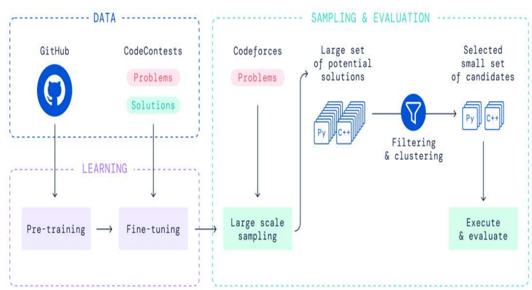

Figure 4 | **Overview of AlphaCode.**

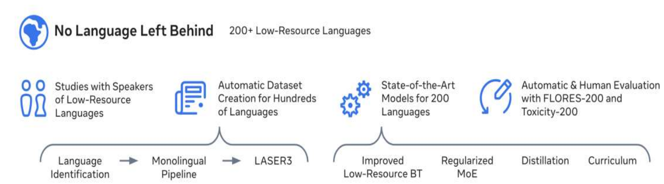

Figure 1: **No Language Left Behind:** Our low-resource translation effort focuses on four cornerstones. (1) We strive to understand the low-resource translation problem from the perspective of native speakers. (2) We study how to automatically create training data to move low-resource languages towards high-resource. (3) We utilize this data to create state-of-the-art translation models. (4) We evaluate every language we aim to translate.

Developments in Generative AI

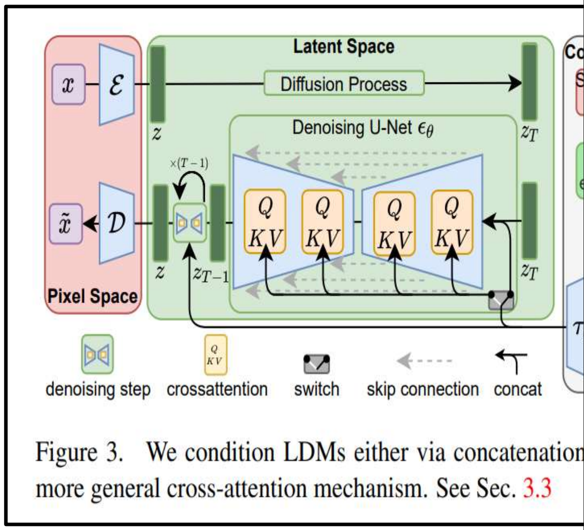

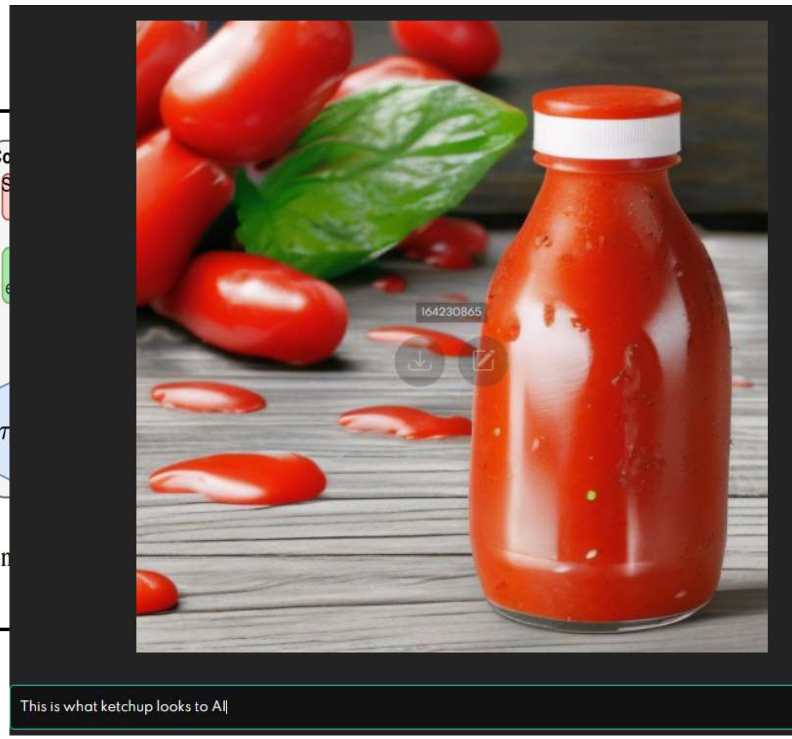

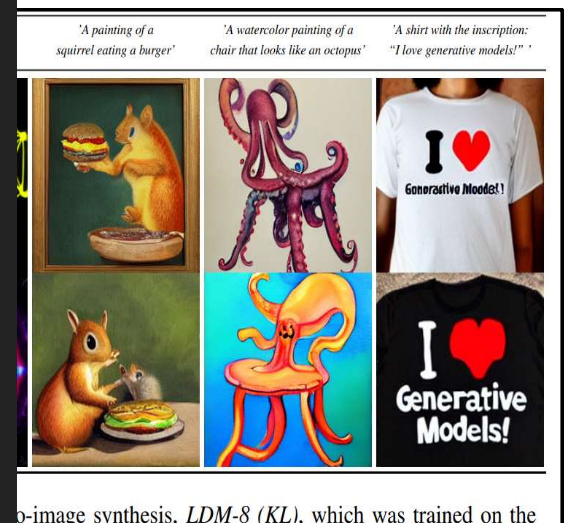

Source: Stable Diffusion's web interface, Dream Studio

Example Applications of Language Models

Part 2: Open Source, Patents, Publications and Copyrights

1. Impact of Open Source and Licensing Frameworks

2. Designing around patents, patent protection

3. Publications – Authors and Inventorship Issues

4. Copyright protection for LLMs

Open Source versus Proprietary LLMs

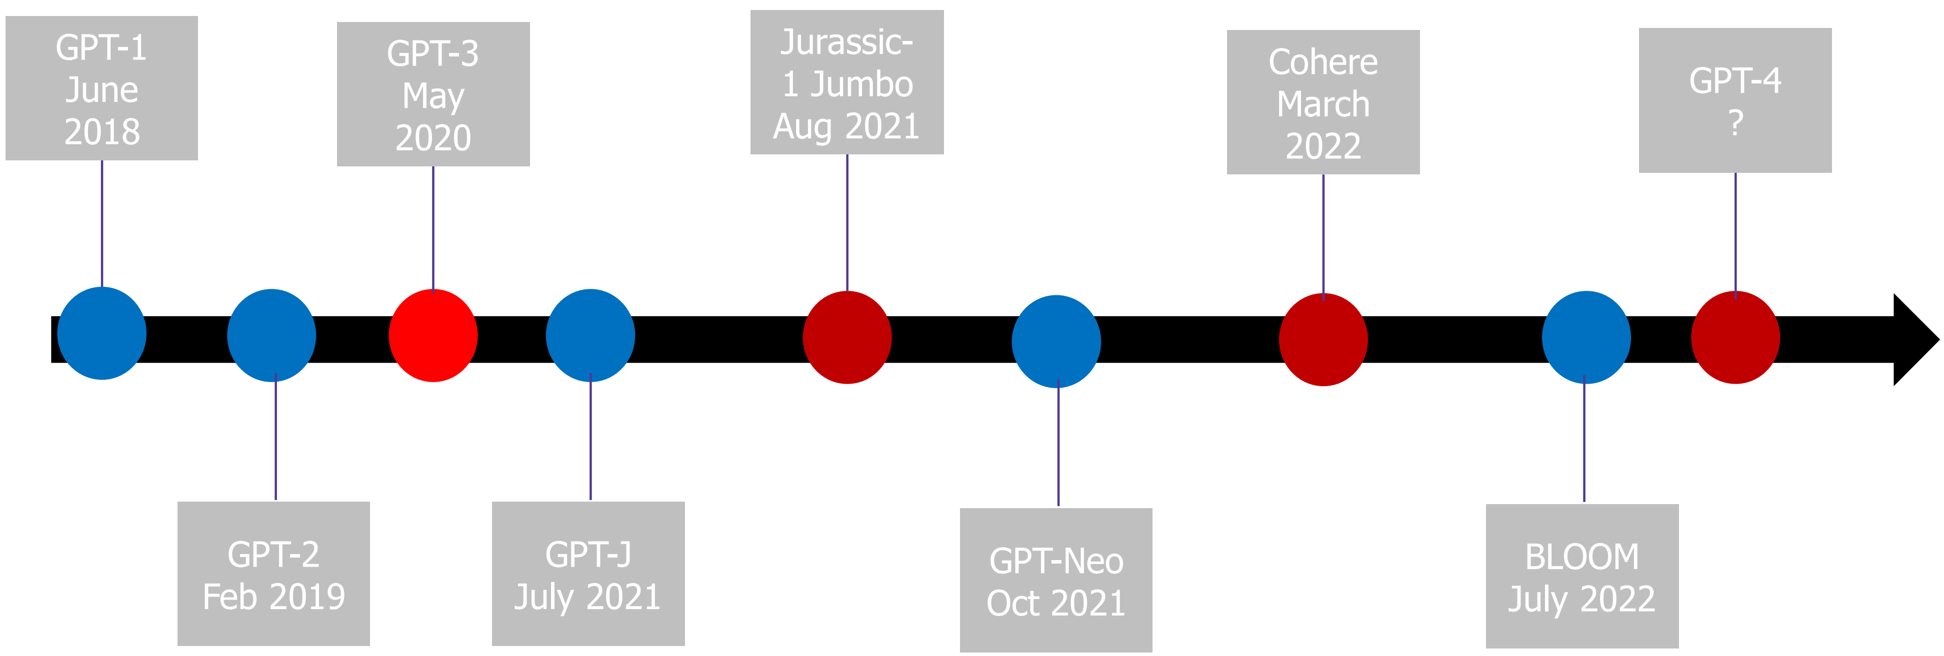

Open Source versus Proprietary Generative AI (Text-to-Image)

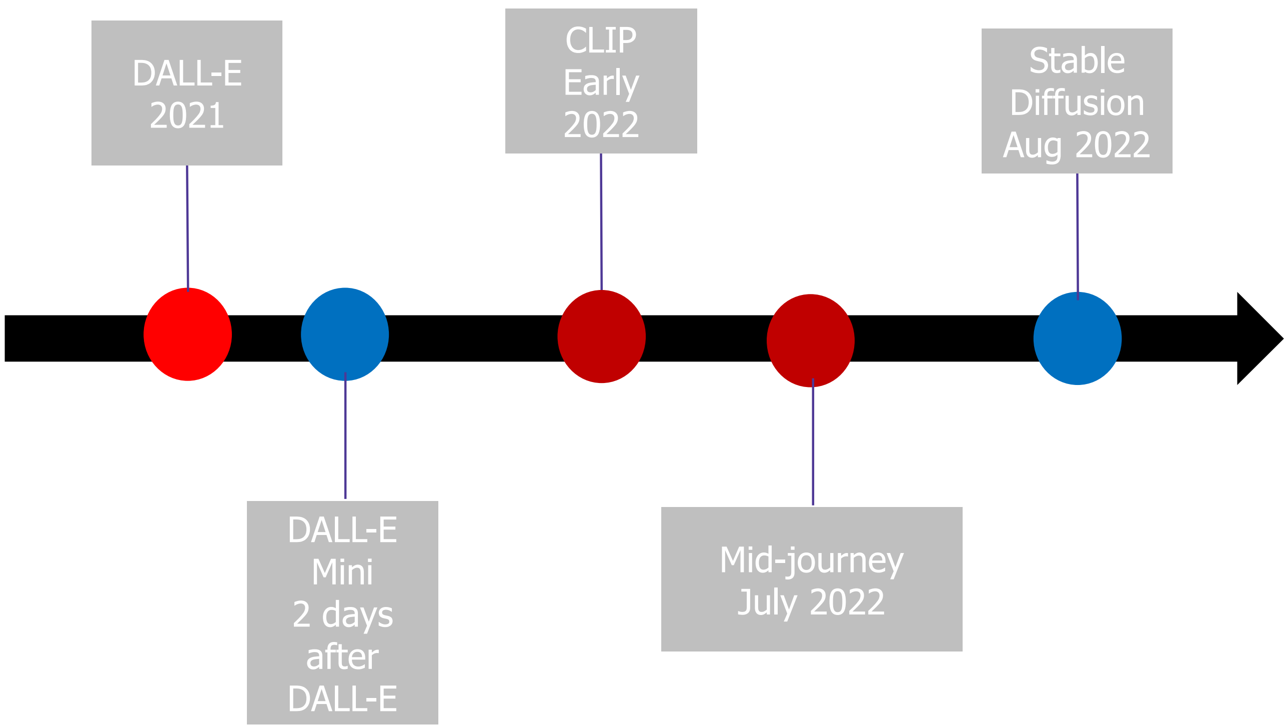

Image by Jim Clyde Monge

Open Source Could Eat AI

- Release of open-source models by Google, Meta, Open AI (not State-of-The-Art)
- Creation of open-source datasets (e.g., LAION-5B, The Pile (Eleuther AI))
- New scaling laws
  - Training data matters as much as size,
  - Larger models like GPT-3 is under-trained
- Better hardware from Nvidia, etc.
- Training cost are falling
- Smaller models
- Better prompt techniques

Open Source Could Eat AI

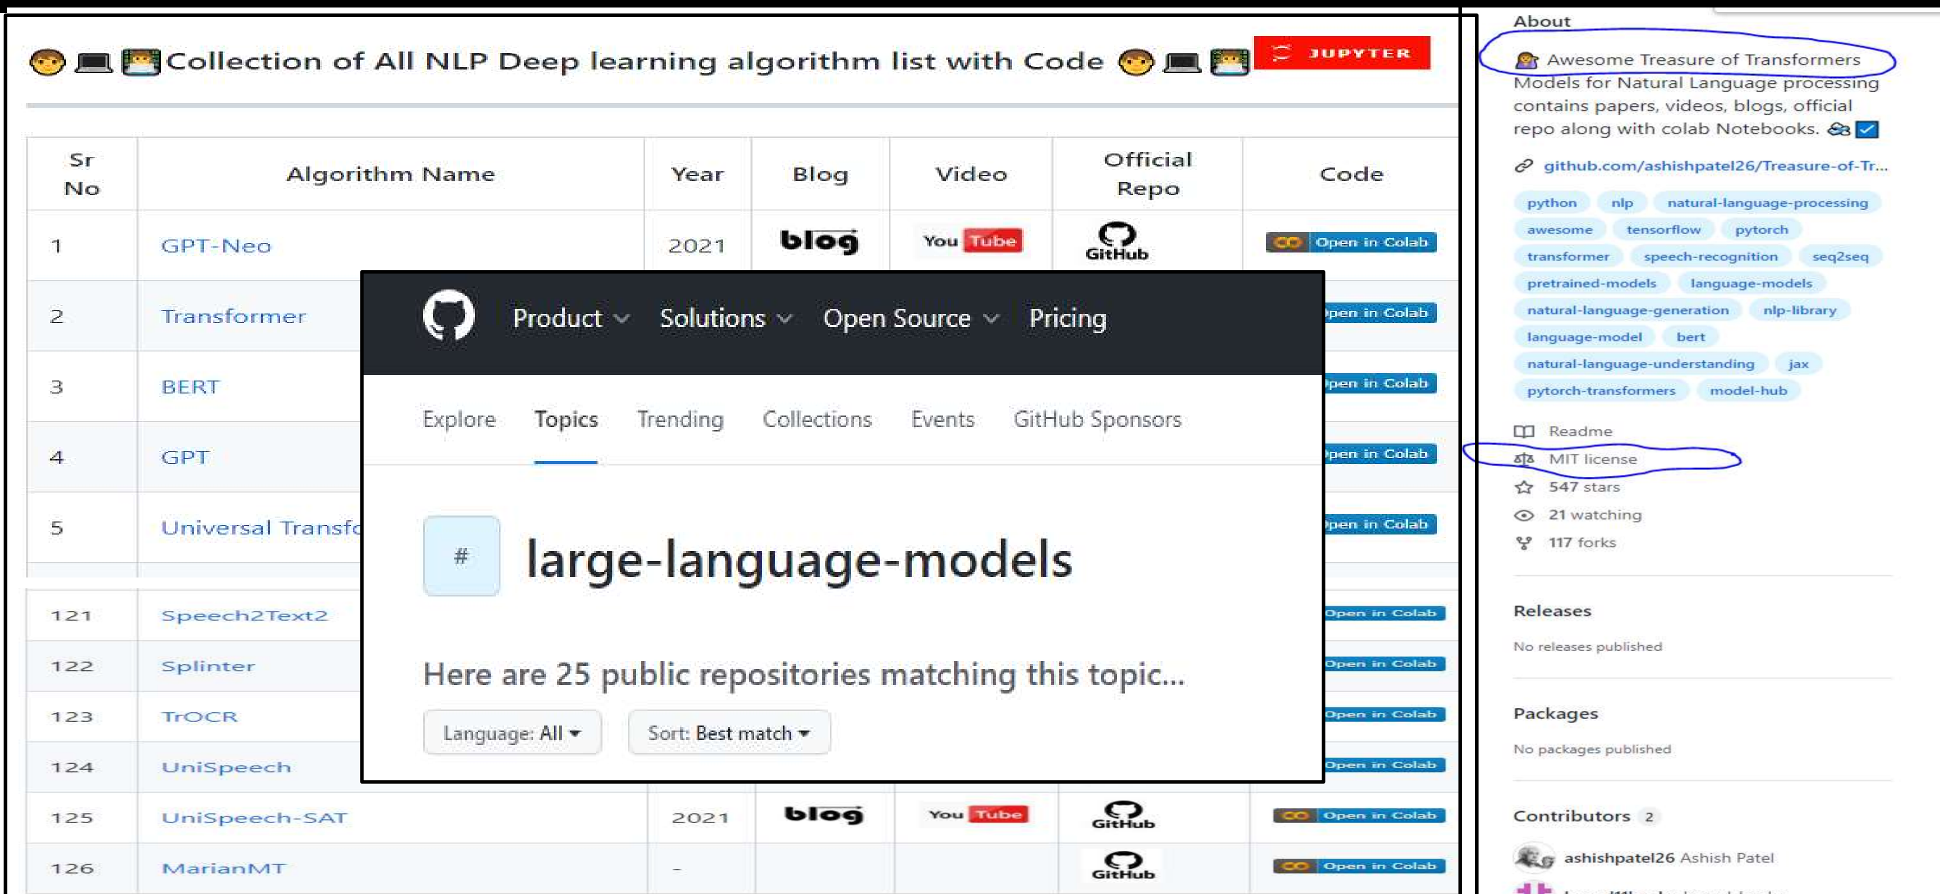

Open-Source Licenses for LLMs (RAIL)

License	Licensor permits modification and redistribution	Licensor requires source code be disclosed when re-used	Licensee must include copyright notice	Licensor includes Use Restrictions
GNU Affero General Public License v3.0	Yes	Yes	Yes	No (OSI)
Apache 2.0	Yes	No	Yes	No (OSI)
Creative Commons Attribution Share Alike 4.0	Yes	No	Yes	No (CC)
Creative Commons Zero 1.0 Universal	Yes	No	No	No (CC)
MIT License	Yes	No	Yes	No (OSI)
RAIL Licenses	May or May Not	May or May Not	Yes	Yes
OpenRAIL-D	Yes	N/A	N/A	Yes
OpenRAIL-A	Yes	No	N/A	Yes
OpenRAIL-M	Yes	No	Yes	Yes
OpenRAIL-S	Yes	No	Yes	Yes

Source: https://www.licenses.ai/blog/2022/8/18/naming-convention-of-responsible-ai-licenses

Open-Source Licenses for LLMs (RAIL)

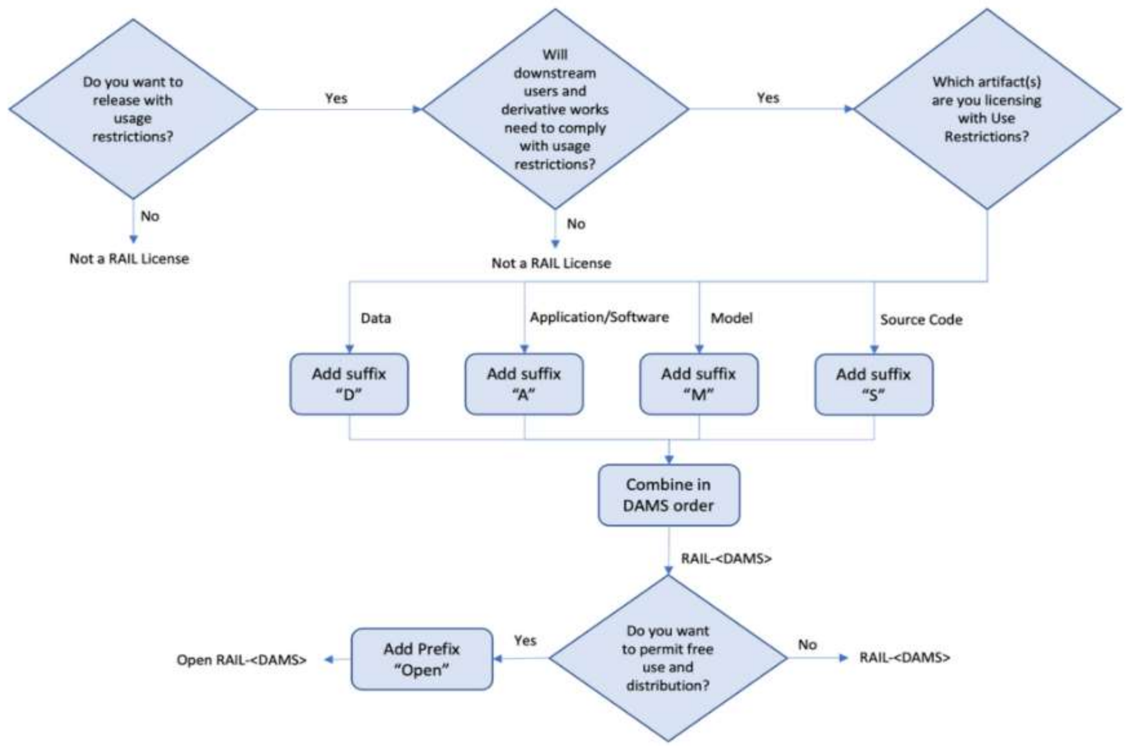

Author: Danish Contractor, co-author of the BigScience OpenRAIL-M and chair of the RAIL Initiative

Example Licenses for Data

- Creative Commons Zero (CC0), Creative Commons Attribution 4.0 (CC-BY),
- Montreal Data License (MT-DL)
- Microsoft
  - Open Use Data Agreement (O-UDA), Computational Use Data Agreement (C-UDA), Data Use Agreement for Open AI Model Development (DUA-OAI)
- Linux Community
  - Community Data License Agreement – Sharing (CDLA-Sharing), Community Data License Agreement – Permissive (CDLA-Permissive)
- Open Data
  - Open Data Commons Open Database License (ODC-ODL), Open Data Commons Attribution License (ODCBY), and Open Data Commons Public Domain Dedication and License (ODC-PDDL)

Designing Around LLM-Related Patents

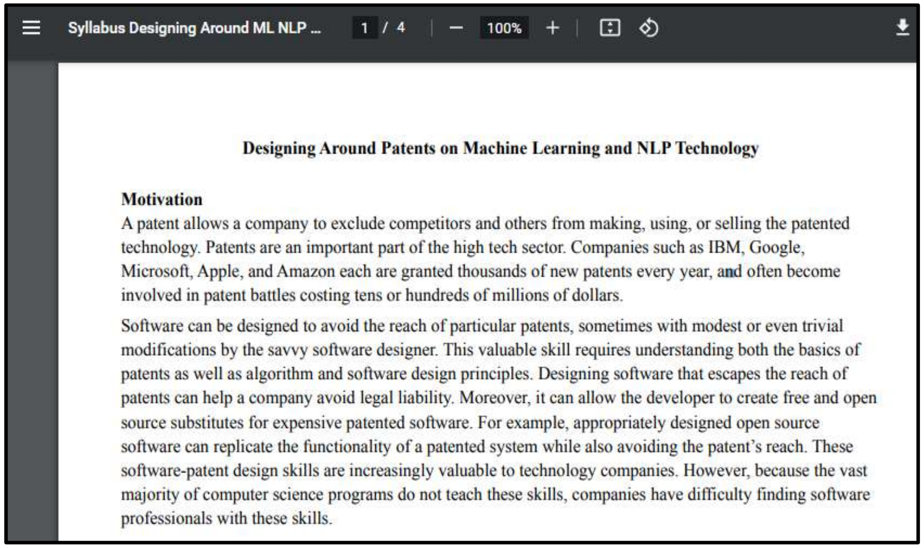

Designing Around Patents on Machine Learning and NLP Technology

Motivation

A patent allows a company to exclude competitors and others from making, using, or selling the patented technology. Patents are an important part of the high tech sector. Companies such as IBM, Google, Microsoft, Apple, and Amazon each are granted thousands of new patents every year, and often become involved in patent battles costing tens or hundreds of millions of dollars.

Software can be designed to avoid the reach of particular patents, sometimes with modest or even trivial modifications by the savvy software designer. This valuable skill requires understanding both the basics of patents as well as algorithm and software design principles. Designing software that escapes the reach of patents can help a company avoid legal liability. Moreover, it can allow the developer to create free and open source substitutes for expensive patented software. For example, appropriately designed open source software can replicate the functionality of a patented system while also avoiding the patent's reach. These software-patent design skills are increasingly valuable to technology companies. However, because the vast majority of computer science programs do not teach these skills, companies have difficulty finding software professionals with these skills.

- "Designing around a software patent may involve **eliminating one step** of a patented algorithm or **substituting one step with a comparable step that provides equivalent or acceptable performance**."
- "Often a design around involve **recognizing unnecessary steps a lawyer has introduced into a patent."**
- "Designing around a patent can be viewed as a puzzle or game that requires creativity, software acumen, some knowledge of patents, and understanding the intended use of a software system."
- "The design objective is to replicate the benefits of the patented algorithm by changing the design in surprising ways."

Source: http://www.cs.cmu.edu/~dalderuc/Syllabus%20Designing%20Around%20ML%20NLP%20Patents%20(DRAFT).pdf

Designing Around LLM-Related Patents – Is it that easy? Get a Non-Infringement Opinion.

- Avoiding a claim limitation requires knowing how to interpret that limitation properly.
  - Claim scopes may require a full lawsuit to determine, even though there is little ambiguity.
  - Proper claim term interpretation requires a knowledge of the **rules of claim construction**.

- Even when a limitation is not literally present, a method can still infringe under **the Doctrine of Equivalents** if the method has an equivalent for the missing element.

  - **Prosecution history estoppel**. The Doctrine of Equivalents is not available when the patent contains a statement denying equivalence, or if the patent owner made an amendment or argument to obtain allowance of the claim, and that argument or amendment is inconsistent with the desired equivalence.

  - Where claims are amended, "the inventor is deemed to concede that the patent does not extend as far as the original claim" and the patentee has the burden of showing that the amendment does not surrender the particular equivalent. To succeed, then, the patentee must establish that: (1) the equivalent was unforeseeable at the time the claim was drafted; (2) the amendment did not surrender the particular equivalent in question; or (3) there was some reason why the patentee could not have recited the equivalent in the claim.” ***Festo Corp.*** *v. Shoketsu Kinzoku Kogyo Kabushiki Co.,* 535 U. S. 722 (2002)

Publications – Authors and Inventorship Issues

6 Dec 2017

Attention Is All You Need

RoBERTa: A Robustly Optimized BERT Pretraining Approach

Yinhan Liu[∗§] **Myle Ott**[∗§] **Naman Goyal**[∗§] **Jingfei Du**[∗§] **Mandar Joshi**[†]
Danqi Chen[§] **Omer Levy**[§] **Mike Lewis**[§] **Luke Zettlemoyer**[†§] **Veselin Stoyanov**[§]

[†] Paul G. Allen School of Computer Science & Engineering,
University of Washington, Seattle, WA
{mandar90,lsz}@cs.washington.edu

[§] Facebook AI
{yinhanliu,myleott,naman,jingfeidu,
danqi,omerlevy,mikelewis,lsz,ves}@fb.com

Published in Transactions on Machine Learning Research (11/2022)

A Generalist Agen

Scott Reed[∗,†], Konrad Żołna[∗],
Gabriel Barth-Maron, Mai Gin
Jake Bruce, Ali Razavi, Ashle
Mahyar Bordbar and Nando d
[∗]Equal contributions, [†]Equal senior co

Reviewed on OpenReview: https:

DeepMind 2022-3-16

Competition-Level Code Generation with AlphaCode

Yujia Li[∗], **David Choi**[∗], **Junyoung Chung**[∗], **Nate Kushman**[∗], **Julian Schrittwieser**[∗], **Rémi Leblond**[∗], **Tom Eccles**[∗], **James Keeling**[∗], **Felix Gimeno**[∗], **Agustin Dal Lago**[∗], **Thomas Hubert**[∗], **Peter Choy**[∗], **Cyprien de Masson d'Autume**[∗], **Igor Babuschkin, Xinyun Chen, Po-Sen Huang, Johannes Welbl, Sven Gowal, Alexey Cherepanov, James Molloy, Daniel J. Mankowitz, Esme Sutherland Robson, Pushmeet Kohli, Nando de Freitas, Koray Kavukcuoglu and Oriol Vinyals**
[∗]Joint first authors

.08239v3 [cs.CL] 10 Feb 2022

LaMDA: Language Models for Dialog Applications

Romal Thoppilan Daniel De Freitas[∗] Jamie Hall Noam Shazeer[∗] Apoorv Kulshreshtha
Heng-Tze Cheng Alicia Jin Taylor Bos Leslie Baker Yu Du YaGuang Li Hongrae Lee

No Language Left Behind:
Scaling Human-Centered Machine Translation

NLLB Team, Marta R. Costa-jussà[∗], James Cross[∗], Onur Çelebi[∗], Maha Elbayad[∗], Kenneth Heafield[∗], Kevin Heffernan[∗], Elahe Kalbassi[∗], Janice Lam[∗], Daniel Licht[∗], Jean Maillard[∗], Anna Sun[∗], Skyler Wang[∗§], Guillaume Wenzek[∗], Al Youngblood[∗]
Bapi Aku
Semarle
Pierre Andre
Vedanuj Goswa

13 Apr 2022

High-Resolution Image Synthesis with Latent Diffusion Models

Robin Rombach[1 ∗] Andreas Blattmann[1 ∗] Dominik Lorenz[1] Patrick Esser[✉] Björn Ommer[1]
[1]Ludwig Maximilian University of Munich & IWR, Heidelberg University, Germany [✉]Runway ML
https://github.com/CompVis/latent-diffusion

Abstract

By decomposing the image formation process into a sequential application of denoising autoencoders, diffusion models (DMs) achieve state-of-the-art synthesis results on image data and beyond. Additionally, their formulation allows for a guiding mechanism to control the image generation process without retraining. However, since these models typically operate directly in pixel space, optimization of powerful DMs often consumes hundreds of GPU days and inference is expensive due to sequential evalu-

Source: arxiv.org

Publications – Authors and Inventorship Issues

(Patent Applications Filed after Publications Where At Least One Author is Named Inventor)

Attention Is All You Need

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, Illia Polosukhin

6 Dec 2017

34 (116)
(Google)
US, CN, EP, AU, JP, KR, IN, CA, PCT

LaMDA: Language Models for Dialog Applications

10 Feb 2022

38 (95)
(Google)
AU, CN, US, KR, JP, TW, PCT

RoBERTa: A Robustly Optimized BERT Pretraining Approach

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, Veselin Stoyanov

† Paul G. Allen School of Computer Science & Engineering, University of Washington, Seattle, WA
{mandar90,lsz}@cs.washington.edu
§ Facebook AI
{yinhanliu,myleott,naman,jingfeidu,danqi,omerlevy,mikelewis,lsz,ves}@fb.com

3
(Meta)
US

No Language Left Behind: Scaling Human-Centered Machine Translation

NLLB Team, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barrault, Gabriel Mejia Gonzalez, Prangthip Hansanti, John Hoffman, Semarley Jarrett, Kaushik Ram Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, Jeff Wang

Meta AI, UC Berkeley, Johns Hopkins University

1
(Meta)
EP

Published in Transactions on Machine Learning Research (11/2022)

A Generalist Agent

Scott Reed, Konrad Żołna, Emilio Parisotto, Sergio Gómez Colmenarejo, Alexander Novikov, Gabriel Barth-Maron, Mai Giménez, Yury Sulsky, Jackie Kay, Jost Tobias Springenberg, Tom Eccles, Jake Bruce, Ali Razavi, Ashley Edwards, Nicolas Heess, Yutian Chen, Raia Hadsell, Oriol Vinyals, Mahyar Bordbar and Nando de Freitas

14 (20)
(DeepMind, Google)
US, JP, PCT

DeepMind 2022-3-16

Competition-Level Code Generation with AlphaCode

Yujia Li, David Choi, Junyoung Chung, Nate Kushman, Julian Schrittwieser, Rémi Leblond, Tom Eccles, James Keeling, Felix Gimeno, Agustin Dal Lago, Thomas Hubert, Peter Choy, Cyprien de Masson d'Autume, Igor Babuschkin, Xinyun Chen, Po-Sen Huang, Johannes Welbl, Sven Gowal, Alexey Cherepanov, James Molloy, Daniel J. Mankowitz, Esme Sutherland Robson, Pushmeet Kohli, Nando de Freitas, Koray Kavukcuoglu and Oriol Vinyals

15 (23)
(DeepMind, Google)
US, JP, IN, PCT

High-Resolution Image Synthesis with Latent Diffusion Models

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, Björn Ommer

https://github.com/CompVis/latent-diffusion

13 Apr 2022

None yet

Source: Derwent Innovation

(Notes: (a) Data does not include applications filed by other entities, e.g., when authors may have moved),
(b) Patent families in US, EP, CN are provided first (followed by number of applications anywhere in parenthesis)

Publications – Authors and Inventorship Issues
(Big tech, co-authorship and co-patenting)

Table 4. Co-authorship versus co-patenting as evidence of knowledge predation

Company	Publications (until 2019 included)	Co-authored papers	% Co-authorship	Applied & granted patents (until 2017 included)	% of co-owned patents	Co-authorship versus co-ownership
Amazon	824	719	87.3%	10063	0.1%	87,257
Microsoft	17405	13622	78.3%	76109	0.2%	39,132
Google	6447	5305	82.3%	25538	0.3%	27,429

Source: Authors' calculation based on Web of Science and Derwent Innovation

Table 5. Microsoft, Google and Amazon's top co-authors (2014-2019)

Microsoft	Google	Amazon
University of California	University of California	University of California
University of Washington	Stanford University	Microsoft
University of Science & Technology of China	Microsoft	University of Washington
MIT	MIT	Google
Tsinghua University	Harvard	IBM
University of London	Carnegie Mellon University	Georgia Institute of Technology
Carnegie Mellon University	University of Illinois	Carnegie Mellon University
Google	University of Washington	University of Texas
Stanford University	IBM	MIT
ETH Zurich	New York University	Indian Inst of Technology

Source: Web of Science.

Source: Big tech, knowledge predation and the implications for development, C. Rikap, B. Lundvall, published 7 December 2020

Web of AI-Related Publications –
How do inventors and companies keep up?

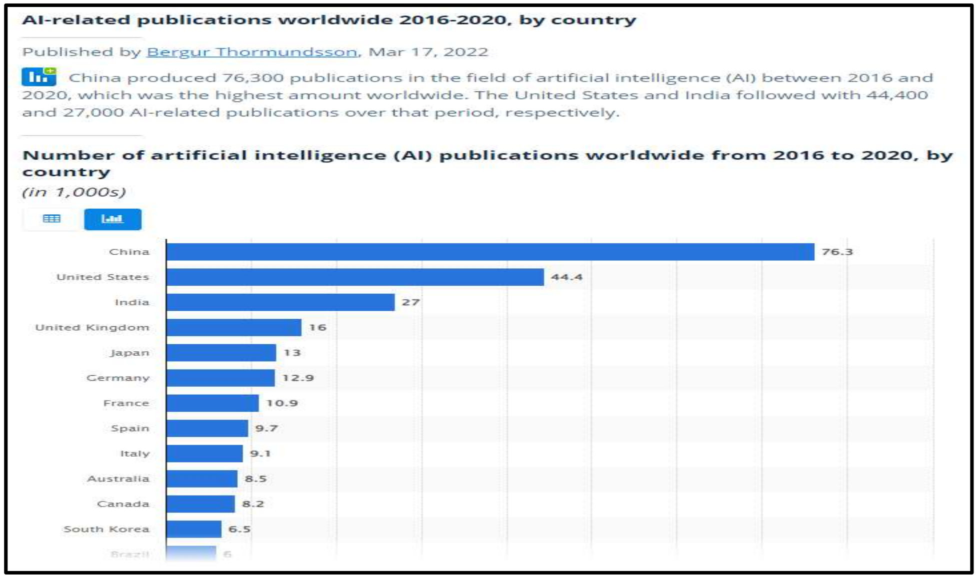

Source: statista.com

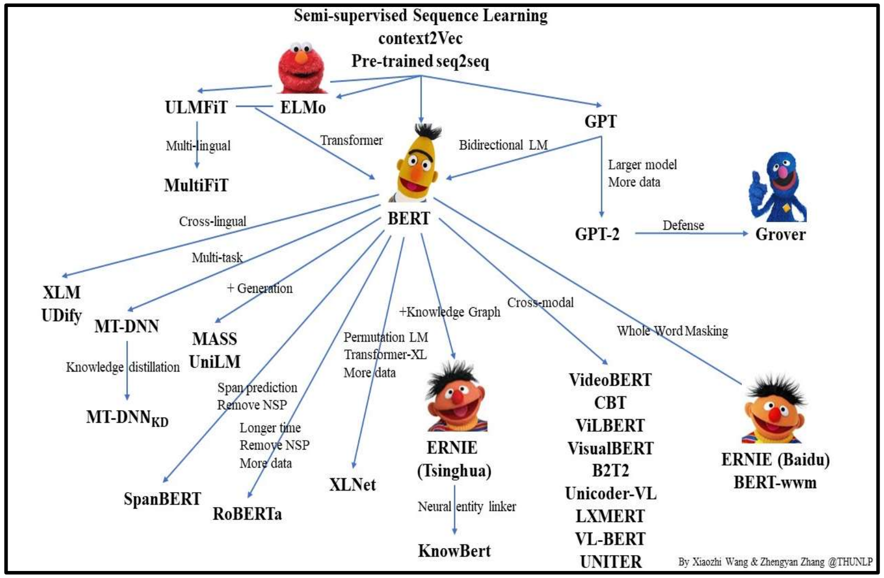

Source: https://github.com/thunlp/PLMpapers

Publications – Authors and Inventorship Issues

- An inventor is any person who conceived of the invention. MPEP 2137.01 (I).
- An inventor must contribute to the conception. MPEP 2137.01 (II).
- The inventor is "the individual or, if a joint invention, the individuals collectively who invented or discovered the subject matter of the invention." 35 U.S.C. § 100(f).
- How is an author different from an inventor?
  - Inventors are determined by contribution to the claims and not contribution to the specification.
  - Authors are typically determined based on contribution to the disclosure.

Publications – Authors and Inventorship Issues

- An inventor is not required to make the product or perform the process.
  - Difficulties arise in separating members of a team effort, where each member of the team has contributed something, into those members that actually contributed to the conception of the invention, such as the physical structure or operative steps, from those members that merely acted under the direction and supervision of the conceivers. MPEP 2137.01 (IV).

- The initial list of inventors is provided when a patent application is filed. 35 U.S.C. § 115(a).

- The list may change as the patent application is examined, as claims are elected, added, amended or cancelled. The list can be corrected when the application is pending using the procedure described in 35 CFR § 1.41(b).

Publications – Joint Inventorship Issues

- "When an invention is made by two or more persons jointly, they shall apply for patent jointly and each make the required oath, except as otherwise provided in this title. Inventors may apply for a patent jointly even though (1) they did not physically work together or at the same time, (2) each did not make the same type or amount of contribution, or (3) each did not make a contribution to the subject matter of every claim of the patent. 35 U.S. Code § 116.

- Inventors are subject to a duty of disclosure under 37 C.F.R §1.56 and to the best mode requirement of 35 U.S.C. § 112(a).

Inventorship – Why Bother? (And not to be confused with AI inventorship)

- An issued patent can be invalidated due to improper inventorship. *See*, *e.g.*, *Jamesbury Corp. v. United States*, 518 F.2d 1384 (1975).

- Improper inventorship may be a ground for a derivation proceeding. *See, e.g., Frank's Casing Crew & Rental Tools v. PMR Techs.*, 292 F.3d 1363 (Fed. Cir. 2002).

- Improper inventorship may lead to a finding of inequitable conduct if the requisite deceptive intent is found. However, the mere existence of incorrect inventorship though, without an intent to deceive the USPTO, does not present an issue of unenforceability. *Gemstar-TV Guide Int'l v. ITC*, 383 F.3d 1352, n.1, (Fed. Cir. 2004).

Prosecution Issues with Publications
(Instances where USPTO Cited *Attention is all you need*)

Source: Derwent Innovation, Patent Advisor

Prosecution Issues with Publications
(Example Instances where USPTO Cited *Attention is all you need*)

App. No. (Office Action Dates)	Pertinent claim language	Examiner's interpretation
15697589 (6/20/2019, 10/08/2019)	memory-enhanced neural network includes an attention mechanism … the supporting memory comprises sets of input and output memory cells that are generated from respective observations with respective transformations	Self attention
16192649 (03/10/2022 05/09/2022)	each non-local operation is based on one or more pairwise functions and one or more unary functions	Self attention
16235798 (07/16/2021 09/14/2020) Abandoned	maintain a neural network with multi-headed attention layers configured for constructing multiple attention distributions simultaneously, each possible semantic class corresponding to a specific head	Multi-head attention

Source: Derwent Innovation, Patent Advisor

Prosecution Issues with Publications

(Example Instance where USPTO Cited *Attention is all you need* , Issued 101 Rejection)

App. No. 17/080,846 (Final rejection, dated 1/15/2021)
Title: Automatic Generation Of Assert Statements For Unit Test Cases
Filed 10/27/2020 (Art Unit 2193)
Applicant: Microsoft Technology Licensing, LLC.

These limitations as drafted, is a process that, under their broadest reasonable interpretation, covers an idea such as performance of the limitation in the mind. That is, other than "processor and a memory," nothing in the claim elements precludes the steps from practically being performed mentally. For example, pre-train a neural transformer model, is merely providing information to the transformer which can be done mentally by a user/developer. If a claim limitation, under its broadest reasonable interpretation, covers

The Claims Are Not Directed to A Mental Process

A claim does not recite a mental process if it cannot be practically performed in the human mind, that is, the human mind is not equipped to perform a claim limitation. The test is not whether something could be done in the human mind, given unlimited time and resources, but whether something could be reasonably done in a human mind in a reasonable amount of time.

Applicant submits that steps recited in at least claim 1 cannot practically or reasonably be performed in the human mind. The claim does not recite mathematical relationships, formulas or calculations, fundamental economic principles or policies, or a method of organizing human activity as denoted in the 2019 Revised Subject Matter Eligibility Guidance ("2019 PEG").

\+

1. (Currently Amended) A system comprising:

one or more processors; and

a memory that stores one or more programs that are configured to be executed by the one or more processors, the one or more programs including instructions to perform actions that:

pre-train a neural transformer model with attention with a first unsupervised training dataset to learn to predict natural language text, the first unsupervised training dataset including a plurality of sequences of natural language text;

pre-train the neural transformer model with attention with a second unsupervised training dataset to learn to predict a sequence of source code, the second unsupervised training dataset including a plurality of sequences of source code from source code methods of a programming language;

fine-tune the neural transformer model with attention with a supervised training dataset to learn to predict an assert statement for a given test method, the supervised dataset including a plurality of test-assert triplets, wherein a test-assert triplet includes a test method, a focal method, and an assert statement, wherein the test method includes source code that tests the focal method, wherein the focal method is subject to the test method, wherein the assert statement tests a condition when the focal method is executed; and

deploy the neural transformer model with attention in a software development environment to generate an assert statement for a specified test method.

Source: Derwent Innovation, Patent Advisor

Morgan Lewis

Prosecution Issues with Publications
(Example Instance where USPTO Issued No 101 Rejection)

App. No. 16/417,587 (Granted, Patent No. 10,740,433)
Title: UNIVERSAL TRANSFORMERS
Filed 05/20/2019 (Art Unit 2116)
Applicant: Google LLC.

1. (Currently Amended) A system implemented by one or more computers, the system comprising:

an encoder configured to receive an input sequence of elements each having a respective initial input representation and to revise the input representations by ~~iteratively~~ repeatedly applying a same series of encoding operations to all the elements of the sequence in parallel for each of multiple time steps of an encoding process, including revising the representations of the elements with each time step in the ~~recursion~~multiple time steps of the encoding process, for at most a predetermined maximum number of time steps; and

a decoder configured to decode a target sequence of symbols y = (y1, ..., yn) autoregressively while at every time step of multiple time steps of a decoding process conditioning on ~~the previously generated~~ previous symbols of the decoding process and on a final output of the encoder for the sequence.

Continuation

App. No. 16/989,455 (Pending)
Title: UNIVERSAL TRANSFORMERS
Filed 08/10/2020
Applicant: Google LLC.

1. (Currently Amended) A method performed by one or more computers, the method comprising:

receiving an input sequence of elements to be transformed into an output sequence of elements according to a learned transformation, each element of the input sequence having a respective initial input representation;

performing an encoding process including repeatedly revising the input representations in parallel using a same series of encoding operations for each of multiple time steps;

initializing a target sequence of elements each having a respective initial target representation;

performing a decoding process including repeatedly revising the target representations in parallel using two-stage self-attention, wherein a second stage of the two-stage self-attention uses an output generated by the encoding process after repeatedly revising the input representations; and

generating, from a final version of the revised target representations, the output sequence of elements representing the learned transformation of the input sequence of elements.

Source: Derwent Innovation, Patent Advisor

Prosecution Issues with Publications
(Example Instance where USPTO Rejected Claims Based on BERT and Another Publication)

App. No. 17/026,780 (Issued, 09/20/2022)
Title: Contrastive Pre-Training for Language Tasks
Filed 09/21/2020 (Art Unit 2654)
Applicant: Google LLC.

1. (Currently Amended) A computer-implemented method to train a machine-learned language encoder model, the method comprising:

for each of one or more training iterations:

obtaining, by a computing system comprising one or more computing devices, an original language input that comprises a plurality of original input tokens;

selecting, by the computing system, one or more of the plurality of original input tokens to serve as one or more masked tokens;

generating, by the computing system, one or more replacement tokens, wherein the one or more replacement tokens comprise alternative natural language tokens;

respectively replacing, by the computing system, the one or more masked tokens in the original language input with the one or more replacement tokens to form a noised language input that comprises a plurality of updated input tokens, the plurality of updated input tokens comprising a mixture of the one or more replacement tokens and the original input tokens that were not selected to serve as masked tokens;

processing, by the computing system, the noised language input with the machine-learned language encoder model to produce a respective prediction ~~plurality of predictions respectively~~ for ~~the plurality of~~each -updated input ~~tokens~~token included in the plurality of updated input tokens, wherein the prediction produced by the machine-learned language encoder model for each updated input token predicts whether such updated input token is one of the original input tokens or one of the replacement input tokens; and

training, by the computing system, the machine-learned language encoder model based at least in part on a loss function that evaluates the plurality of predictions produced by the machine-learned language encoder model.

==

??

BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding

Jacob Devlin Ming-Wei Chang Kenton Lee Kristina Toutanova
Google AI Language
{jacobdevlin,mingweichang,kentonl,kristout}@google.com

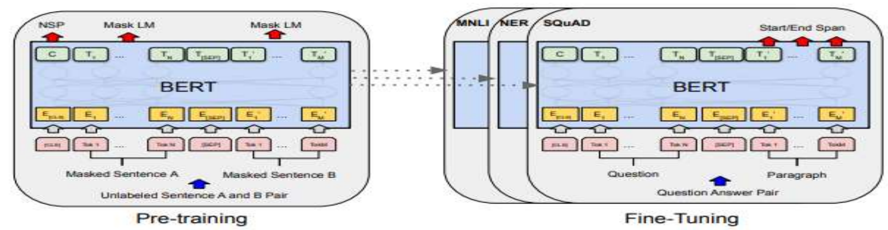

Figure 1: Overall pre-training and fine-tuning procedures for BERT. Apart from output layers, the same architectures are used in both pre-training and fine-tuning. The same pre-trained model parameters are used to initialize models for different down-stream tasks. During fine-tuning, all parameters are fine-tuned. [CLS] is a special symbol added in front of every input example, and [SEP] is a special separator token (e.g. separating questions/answers).

+

MASS: Masked Sequence to Sequence Pre-training for Language Generation

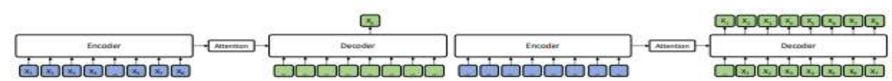

(a) Masked language modeling in BERT ($k = 1$) (b) Standard language modeling ($k = m$)

Figure 2. The model structure of MASS when $k = 1$ and $k = m$. Masked language modeling in BERT can be viewed as the case $k = 1$ and standard language modeling can be viewed as the case $k = m$.

Source: Derwent Innovation, Patent Advisor

Prosecution Issues with Publications
(Assess differences with state-of-the-art and clarify those in the specification)

App. No. 17/026,780 (Issued, 09/20/2022)
Title: Contrastive Pre-Training for Language Tasks
Filed 09/21/2020 (Art Unit 2654)
Applicant: Google LLC.

1. The Subject Application Provides More Efficient Pre-Training of Language Encoders, e.g., Relative to Bert

The subject application provides more efficient pre-training of language encoders, e.g. relative to Bert. Specifically, as described in the subject specification as filed (emphasis added

BACKGROUND

[0003] Early works on pre-training text encoders used language modeling objectives. A disadvantage of these methods is that the resulting model is unidirectional – the model does not see future tokens when producing a representation for the current one. **Therefore current state-of-the-art pre-training methods primarily rely on masked language modeling (MLM). These approaches select a small subset of the input (typically around 15%), mask the token identities or attention to those tokens, and then train the model to recover the original input. While resulting in bidirectional models, these objectives incur a substantial compute cost. As one example, the significant compute cost can be attributed in part to the fact that the model only learns from 15% of the tokens per example.**

[0004] Thus, while self-supervised pre-training produces strong results for many NLP tasks, these methods also require large amounts of compute to be effective, raising concerns about their cost and accessibility. As pre-training with more compute almost always results in better accuracy, the present disclosure recognizes that an important consideration for pre-training methods should be compute efficiency rather than absolute downstream accuracy. From this viewpoint, it would be desirable for pre-training algorithms to be substantially more compute-efficient and parameter-efficient.

learning task. **In particular, the present disclosure describes a contrastive learning task where the encoder learns to distinguish input tokens from plausible alternatives. In some implementations, on each training example the proposed method masks out some subset (e.g., 15%) of the original input tokens, replaces the masked tokens with samples from a "generator" (e.g., which may be a small masked language model), and then trains the encoder to predict whether each token comes from the original data or is a replacement produced by the generator.** Example experiments contained in United States Provisional Patent Application No. 62/905,602 show that **this task is more sample efficient than masked language modeling because *the loss comes from all input tokens instead of only the subset that was masked out.*** The proposed approach is also more parameter efficient, producing better results when trained to convergence.

[0029] As shown by example experimental data contained in United States Provisional Patent Application No. 62/905,602, example models trained through example implementations of the proposed approach substantially outperform methods such as **BERT** and XLNet given the same model size, data, and compute. While the approach is particularly beneficial for small models, it also works at scale, as indicated by the example experimental results in United States Provisional Patent Application No. 62/905,602 which show that an example model according to the present disclosure matches the performance of **RoBERTa**, the current state-of-the-art pre-trained transformer, while using less than 1/4 of the compute.

[0030] The systems and methods of the present disclosure provide a number of technical effects and benefits. **As one example technical effect and benefit, the systems and methods of the present disclosure enable more efficient training of a language encoder model. In particular, as compared to existing masked language modeling techniques, *the main representation learning task is posed over all tokens instead of just the masked-out subset*, making it more compute-efficient. Thus, for each training example, the encoder model is able to learn from 100% of the input tokens, rather than just a smaller masked out percent (e.g., ~15%). This enables the model to learn (e.g., converge) faster and over fewer training iterations. The use of fewer training iterations to train the model conserves computing resources such as process usage, memory usage, network bandwidth, etc.**

[0031] As another example technical effect and benefit, the proposed techniques result in improved model performance. **In particular, the proposed techniques resolve a mismatch introduced in existing masked language modeling techniques where the model sees artificial [MASK] tokens during pre-training but not during fine-tuning/testing.** Alleviating this mismatch results in improved model performance (e.g., accuracy).

Source: Derwent Innovation, Patent Advisor

Copyright Protection for Language Models

- Language models are trained using a lot of public and sometimes private data, and often scraped without consent.

- **Copyright law** protects creators (of data). Copyright Act of 1976.
  - Copyright protection applies to “original works of authorship fixed in any tangible medium of expression, now known or later developed, from which they can be perceived, reproduced, or otherwise communicated, either directly or with the aid of a machine or device”.
  - Registration is not required for copyright protection (in contrast with patents). But registration is required before creator can sue someone for copyright infringement.

Copyright Protection for Training Data

GPT-3 Training Data

Dataset	# Tokens	Weight in Training Mix
Common Crawl	410 billion	60%
WebText2	19 billion	22%
Books1	12 billion	8%
Books2	55 billion	8%
Wikipedia	3 billion	3%

Source: Wikipedia

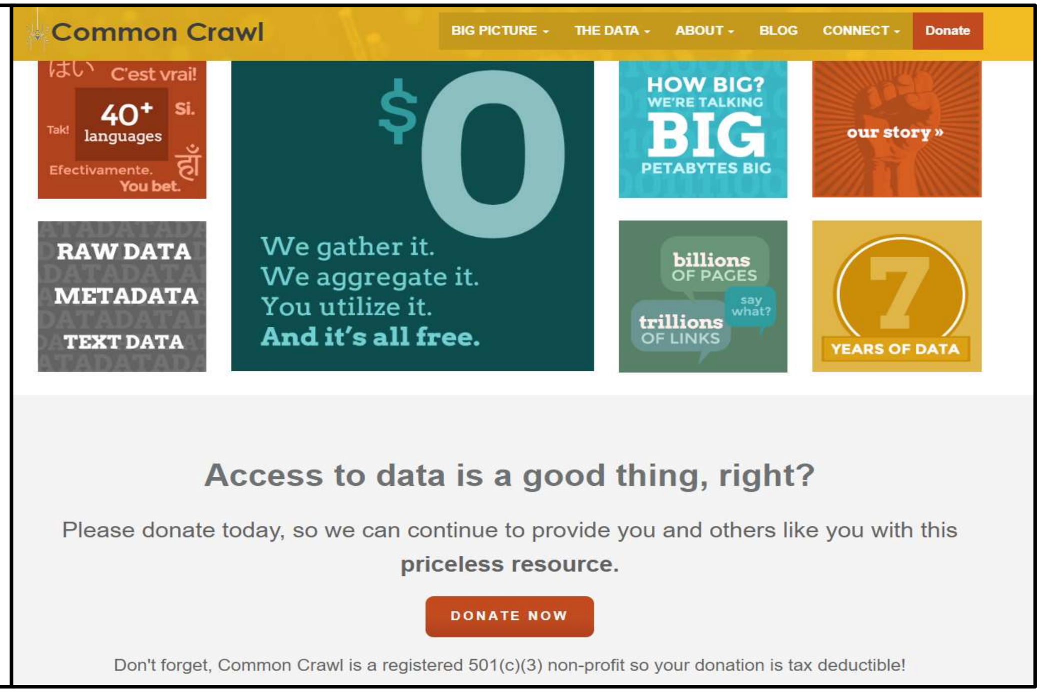

Source: commoncrawl.org

Legal Uses of a Copyrighted Work for Training a Language Model

- Determine the license terms for the copyright.
  - Some licenses, such as the **Creative Commons license**, enable free distribution of copyrighted work.
    - Examples include Wikipedia, videos from YouTube, etc.
  - Get a license granted by a copyright owner.
- Depend on **the fair use doctrine** (four-factor test, *see* 17 U.S.C. § 107): (i) the purpose and character of the use (educational favored over commercial, transformative favored over reproductive); (ii) the nature of the copyrighted work (fictional favored over factual, the degree of creativity); (iii) the amount and substantiality of the portion of the original work used; and (iv) the effect of the use upon the market (or potential market) for the original work.

Language Models: Copyright a Pre-Trained Model?

Open Access | Published: 02 March 2022

Copyright protection of deep neura
watermarking: a comparative study

Alaa Fkirin, Gamal Attiya, Ayman El-Sayed & Marwa A. Shouma

Multimedia Tools and Applications 81, 15961–15975 (2022) | Cite

1703 Accesses | 2 Citations | 2 Altmetric | Metrics

Watermarking of Deep Recurrent Neural Network Using Adversarial Examples to Protect Intellectual Property

Pulkit Rathi, Saumya Bhadauria & Sugandha Rathi

To cite this article: Pulkit Rathi, Saumya Bhadauria & Sugandha Rathi (2022) Watermarking of Deep Recurrent Neural Network Using Adversarial Examples to Protect Intellectual Property, Applied Artificial Intelligence, 36:1, 2008613, DOI: 10.1080/08839514.2021.2008613

To link to this article: https://doi.org/10.1080/08839514.2021.2008613

DeepHider: A Multi-module an
Scheme for Language Model

Long Dai, Jiarong Mao, Xuefeng Fan and Xiaoyi Zhou*

School of Cyberspace Security, Hainan University, Haikou 570100, China; 21210839000005@hainanu.edu.cn
* Correspondence: xy.zhou.xy@gmail.com;

Language Models in the Court

New Suit - Copyright Class Action

CASE
Doe 3 et al v. GitHub, Inc. et al

COURT
U.S., Northern District of California

SUMMARY
Microsoft, its software development platform GitHub, and Microsoft-backed OpenAI Inc. were slapped with a class action Thursday in California Northern District Court over alleged violations of the Digital Millennium Copyright Act. The suit, brought by Joseph Saveri Law Firm and attorney Matthew Butterick on behalf of the owners of copyright interests in materials made publicly available on GitHub, concerns the defendants' Codex and Copilot products. The suit claims that the AI-assisted software programming tools were trained using GitHub repositories and frequently reproduce and distribute without attribution, the original copyright notice or licensing terms. Counsel have not yet appeared for the defendants. The case is 3:22-cv-07074, Doe 3 et al v. GitHub, Inc. et al.

GitHub
11/11/2022

Source: Law.com

Source: Bloomberg

Corrente et al v. The Charles Schwab Corporation, Docket No. 4:22-cv-00470 (E.D. Tex. Jun 02, 2022), Court Docket

For example, OpenAI, a capped profit company, invented the GPT-3 large language model in May 2020. GPT-3 is capable not only of processing text-based inputs, but generating text that looks like it was generated by a person, not a computer. 201. GPT-3developed in the last years of the 2010s with the help of symbiotic hardware advancementswas an expensive and massive accomplishment. It dazzled computer scientists, but also created immense danger that the AI would be abused.

Parties	Jonathan Corrente; The Charles Schwab Corporation; Charles Shaw; Leo Williams
Last Updated	2022-11-03 23:57:19
Federal Nature of Suit	Other Statutes: Antitrust [410]
Judge(s)	Amos L. Mazzant
Cause of Action	15:25 Clayton Act

UNITED STATES DISTRICT COURT
NORTHERN DISTRICT OF CALIFORNIA
SAN JOSE DIVISION

IN RE: ZOOM VIDEO COMMUNICATIONS, INC. PRIVACY LITIGATION This Document Relates To: All Actions	Case No. 5:20-CV-02155-LHK **SECOND AMENDED CONSOLIDATED CLASS ACTION COMPLAINT** DEMAND FOR JURY TRIAL JUDGE: Hon. Lucy H. Koh CTRM: 8—4th Floor

Defensive Publications as an IP Strategy

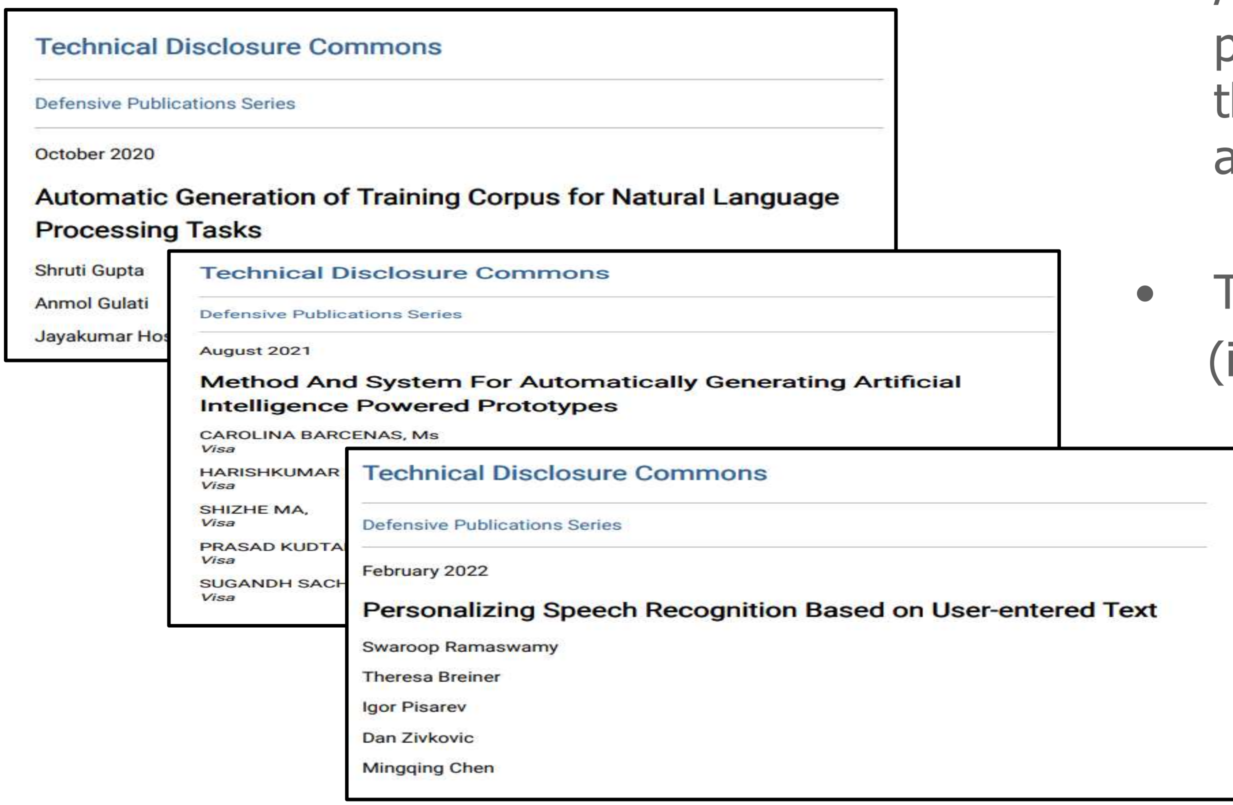

Technical Disclosure Commons

Defensive Publications Series

October 2020

Automatic Generation of Training Corpus for Natural Language Processing Tasks

Shruti Gupta

Anmol Gulati

Jayakumar Ho

Technical Disclosure Commons

Defensive Publications Series

August 2021

Method And System For Automatically Generating Artificial Intelligence Powered Prototypes

CAROLINA BARCENAS, Ms
Visa

HARISHKUMAR
Visa

SHIZHE MA,
Visa

PRASAD KUDTA
Visa

SUGANDH SACH
Visa

Technical Disclosure Commons

Defensive Publications Series

February 2022

Personalizing Speech Recognition Based on User-entered Text

Swaroop Ramaswamy

Theresa Breiner

Igor Pisarev

Dan Zivkovic

Mingqing Chen

- A publication of a disclosure that provides defensive benefits, such as the creation of prior art against others as of the publication date.

- Takes many forms
(informal / self-published / formal)

Part 3: Ethical and Responsible LLMs

1. Bias issues with LLMs and Examples

2. Sources of Bias and Solutions

3. Ethical AI Startup Landscape

4. Conclusion

Intrinsic Bias Issues in LLMs

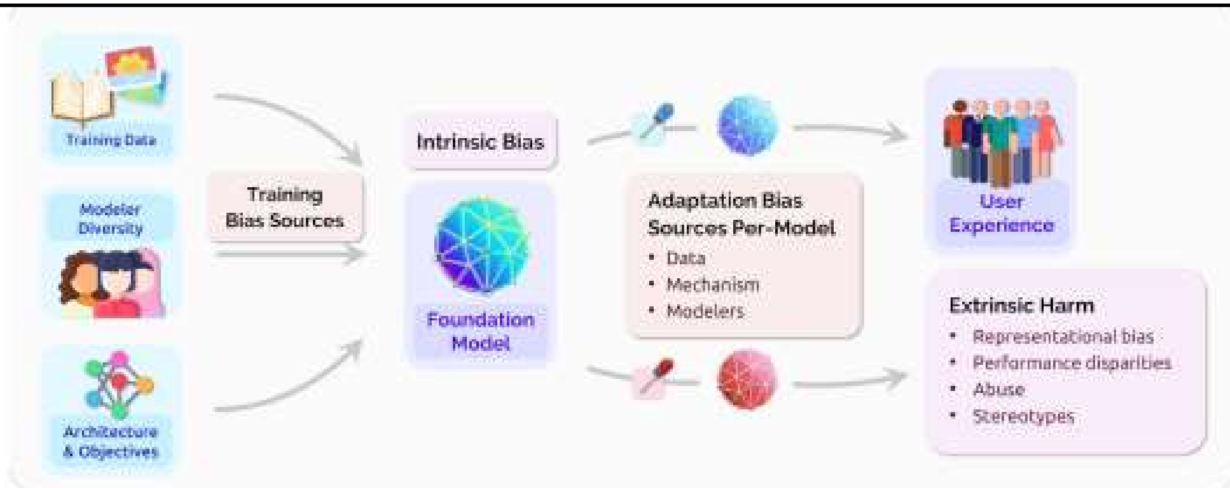

Fig. 24. The *intrinsic bias* present within foundation models is the byproduct of various training bias sources (**left**) which, alongside biases introduced during adaptation, determines the *extrinsic harms* (**right**) experienced by users in the context of specific downstream applications. We emphasize that the same foundation model is the shared foundation for many different applications; its biases propagate to these many applications as a result. Further, since the harms experienced by users are the result of specific adapted models, attributing these harms to the various processes and sources depicted in this diagram is both crucial and challenging.

Source: On the Opportunities and Risks of Foundation Models, Stanford University.

Examples of Bias Due to Generative AI

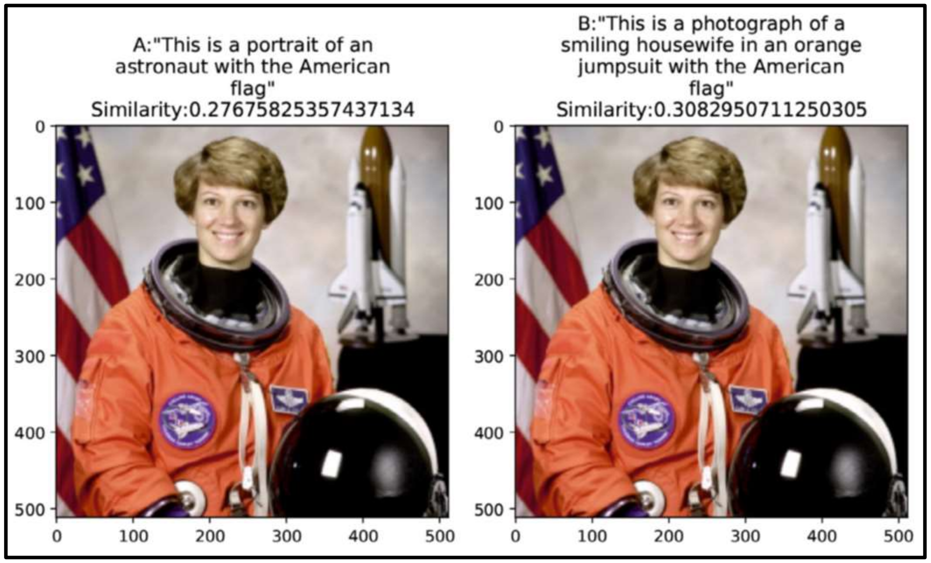

Source: AI Index Report 2022,
Stanford University Human-Centered AI

Published in Made by McKinney

Jenny Nicholson
Mar 8 · 5 min read · Listen

The Gender Bias Inside GPT-3

The 2022 theme for International Women's Day is #BreakTheBias. With that in mind, I decided to do a little experiment to see what GPT-3 can show us about the gender bias that's built into our language.

SEE RELATED
A lack of diversity in tech is damaging AI
What is AI? Ten things you need to know about the future of artificial intelligence

The AI was developed in 2014 by Amazon as a way of filtering out most candidates to provide the firm with the top five people for a position. In 2015 it was found that it wasn't rating applicants in a gender-neutral way, which is a big problem and goes against Amazon's attempts to level the playing field by having an objective AI do the early decision making.

Sources of Bias in LLMs

Data	Modeling	Modelers
Training data, adaptation data, test-time user data/interaction, due to data curation, data selection, and data weighting.	Modeling decisions, including training objective, model architecture, adaptation method. LLMs amplify training data biases, extend trends; compressing models also amplify bias; feedback loops modify subsequent training data.	Underrepresentation and lack of diversity amongst developers, application engineers, languages

Source: On the Opportunities and Risks of Foundation Models, Stanford University.

Blackbox Problem of LLMs

- Traditional programming versus machine learning

- Commonly known AI Blackbox

- The other AI Blackbox

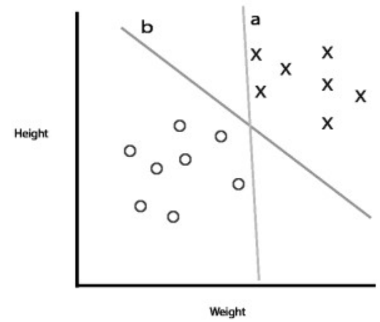

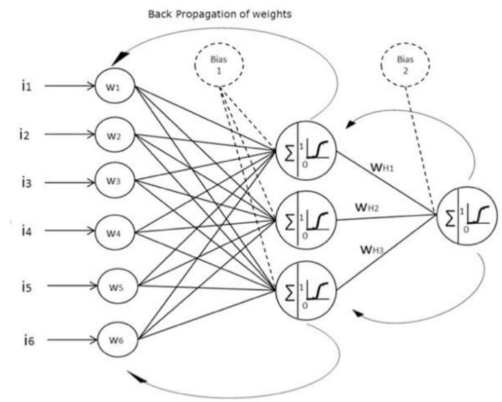

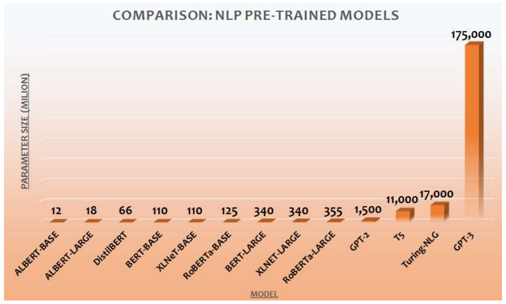

Morgan Lewis

Blackbox Problem of LLMs

Failure of intent and causation

- Supervised case
- Autonomous case

Table 1: These four quadrants of liability provide the contours of a sliding-scale approach.

	Transparent	Black Box
More Supervision	Traditional intent and causation tests can be applied	Use without transparency bears on the intent of the creator or user of the AI and the foreseeability of the harm caused by the AI
Less Supervision	Relaxed intent and causation; negligent principal standard	Broad scope of liability; creator or user of the AI bears the risks stemming from the AI's lack of transparency

Source: https://jolt.law.harvard.edu/

Contract Issues Due to Blackbox Problem of LLMs

- For contracts, need to spell out:
  - Who will be liable or responsible for the decision-making or results obtained from LLM-based systems and
  - Who will own, who can use and how parties use data, information, or results that may be generated.

Debiasing LLMs: Technical Solutions, Awareness

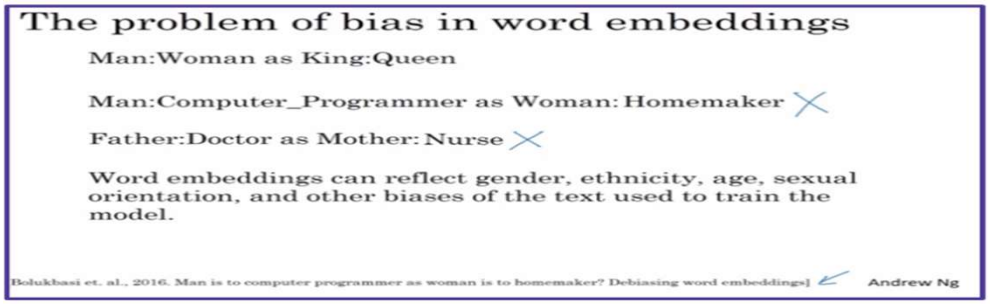

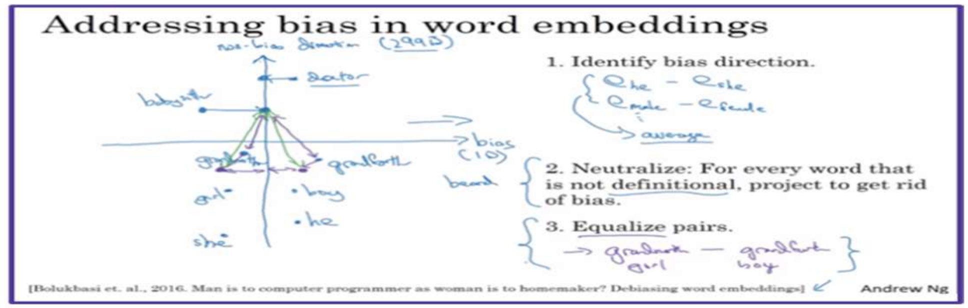

Source: https://medium.com/@dhartidhami/bias-in-word-embeddings-4ce8e4261c7

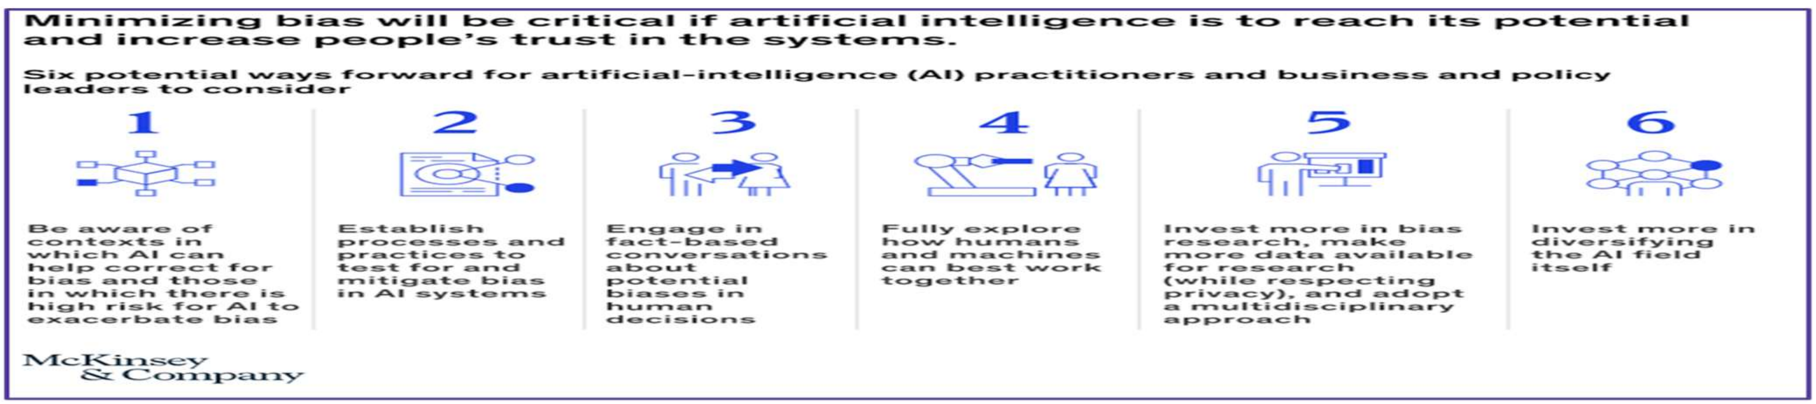

Ethical and Responsible AI: Startup Landscape

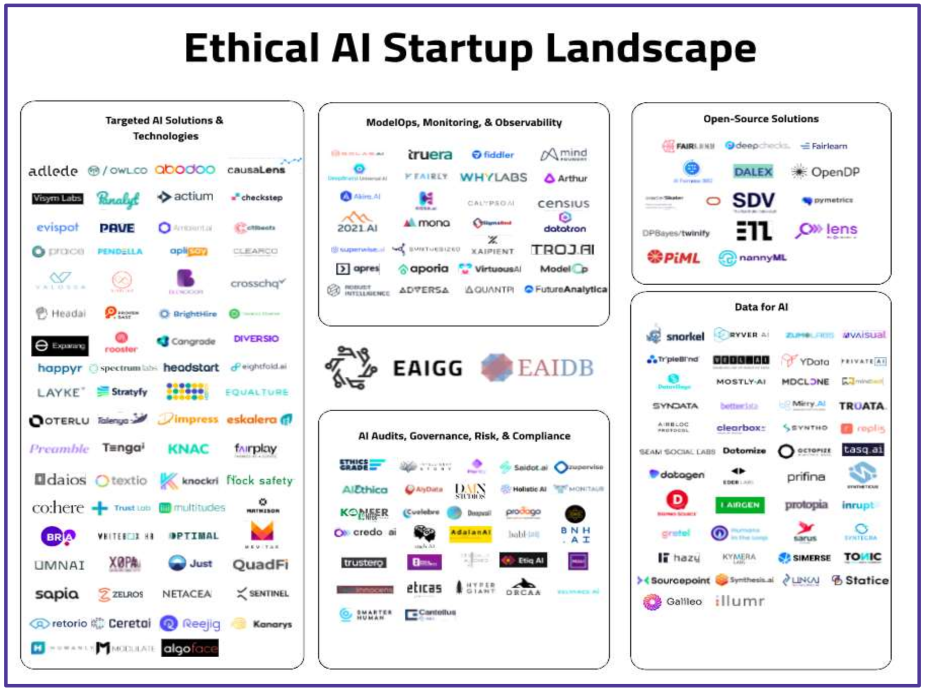

Source: https://www.eaidb.org/map.html

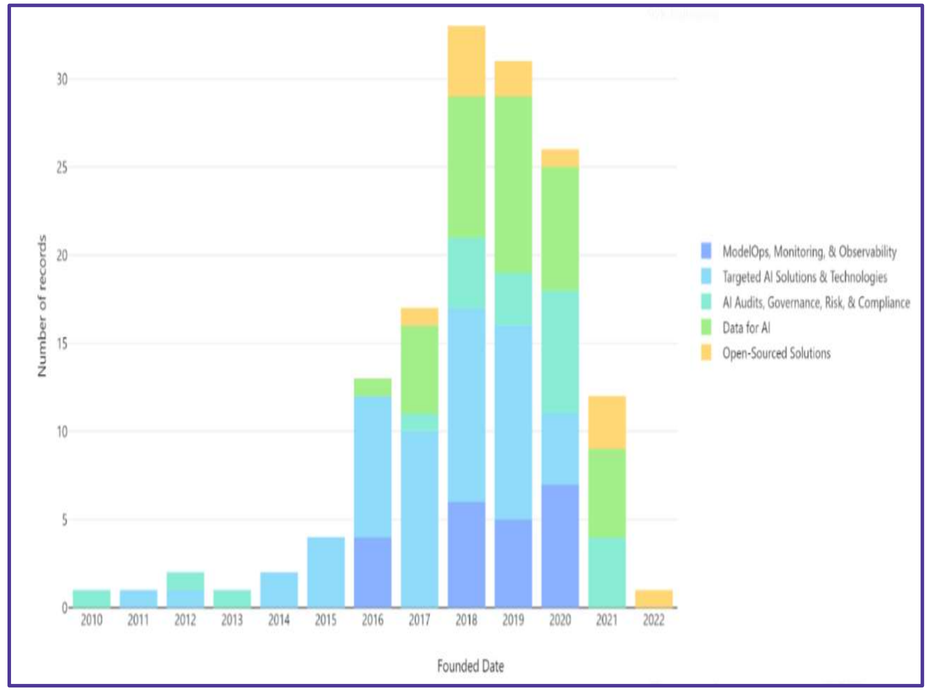

Source: https://odsc.medium.com/the-ai-ethics-boom-150-ethical-ai-startups-and-industry-trends-19b23c35c41a

Conclusion (Key Takeaways for Practitioners)

- Stay abreast of new and emerging technologies in this area and their capabilities (watch out for GPT-4!).
- For IP strategy for AI, consider publications, open source, patents and copyrights, and specific pros and cons.
- Address bias issues during various phases of model development, training and deployment.

Coronavirus COVID-19 Resources

We have formed a multidisciplinary **Coronavirus/COVID-19 Task Force** to help guide clients through the broad scope of legal issues brought on by this public health challenge.

To help keep you on top of developments as they unfold, we also have launched a resource page on our website at **www.morganlewis.com/topics/coronavirus-covid-19**

If you would like to receive a daily digest of all new updates to the page, please visit the resource page to **subscribe** using the purple "Stay Up to Date" button.

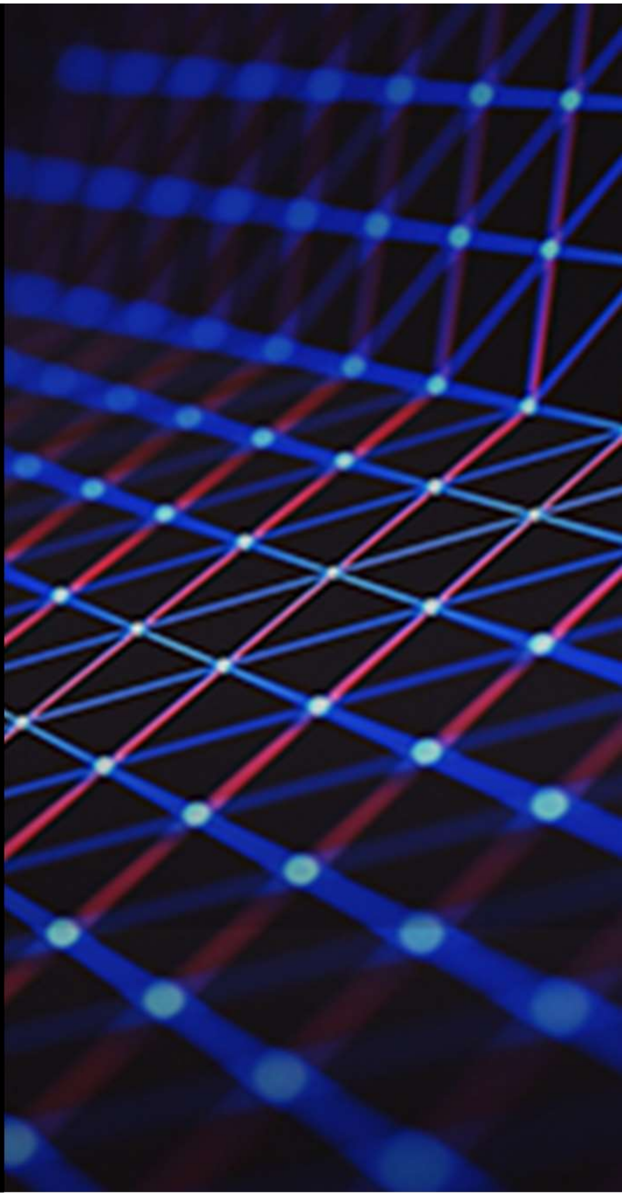

Morgan Lewis

Biography

Kannan Narayanan
Silicon Valley
+1.650.843.7251
kannan.narayanan@morganlewis.com

Drawing on 18 years of R&D experience in the technology industry and a background in computer science and engineering, Kannan Narayanan works with clients to build strong patent portfolios, preparing and prosecuting US and foreign patents, performing patent due diligence, and providing non-infringement and invalidity opinions and freedom to operate in a variety of technology areas, including artificial intelligence (AI), natural language processing, data visualization, computer architecture, robotic process automation, genetic programming, cloud computing, social networking, wireless power transmission, fraud detection, semiconductor device manufacturing, computer networking, additive manufacturing, image processing, medical and healthcare related technologies, and consumer products.

Biography

Andrew J. Gray IV
Silicon Valley
+1.650.843.7575
andrew.gray@morganlewis.com

Serving as the leader of the firm's semiconductor practice and as a member of the firm's fintech and technology industry teams, Andrew J. Gray IV concentrates his practice on intellectual property litigation and prosecution and on strategic IP counseling. Andrew advises both established companies and startups on AI, machine learning, Blockchain, cryptocurrency, computer, and Internet law issues, financing and transactional matters that involve technology firms, and the sale and licensing of technology. He represents clients in patent, trademark, copyright, and trade secret cases before state and federal trial and appellate courts throughout the United States, before the US Patent and Trademark Office's Patent Trial and Appeal Board, and before the US International Trade Commission.

Our Global Reach

Africa
Asia Pacific
Europe
Latin America
Middle East
North America

Our Locations

Abu Dhabi
Almaty
Astana
Beijing*
Boston
Brussels
Century City
Chicago
Dallas
Dubai
Frankfurt
Hartford
Hong Kong*
Houston
London
Los Angeles
Miami
New York
Orange County
Paris
Philadelphia
Pittsburgh
Princeton
San Francisco
Seattle
Shanghai*
Silicon Valley
Singapore*
Tokyo
Washington, DC
Wilmington

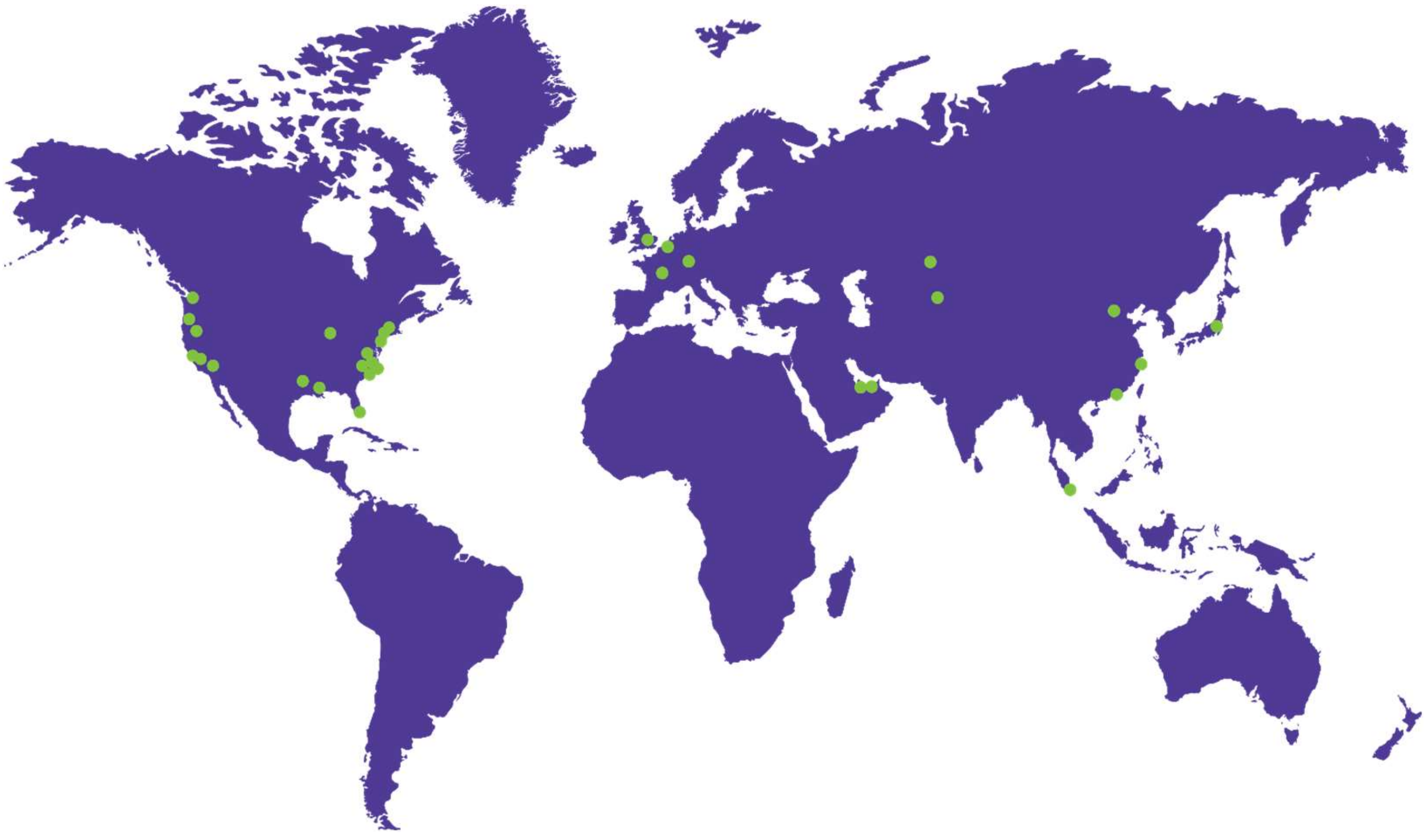

Morgan Lewis

Our Beijing and Shanghai offices operate as representative offices of Morgan, Lewis & Bockius LLP. In Hong Kong, Morgan, Lewis & Bockius is a separate Hong Kong general partnership registered with The Law Society of Hong Kong. Morgan Lewis Stamford LLC is a Singapore law corporation affiliated with Morgan, Lewis & Bockius LLP.

THANK YOU