Dataminr®

AI in Real-Time:
From Emergency Response to Corporate Security

Alejandro (Alex) Jaimes, Ph.D.

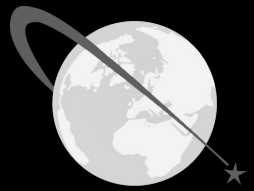

Center of Excellence
WIRELESS AND INFORMATION TECHNOLOGY
AT STONY BROOK UNIVERSITY

November 16th 2023


LLMs
Emergency Response
AI
Research
AI in Practice

King’s Cross Fire, 31 casualties (1987)

Kings Cross: What Failed?

- Various point to point communications
- Little/no inter-departmental communication
- Speed


To New York
DAVID M. CALABRESE
917.244.5111
DUNKIN' DONUTS

Centralized Model

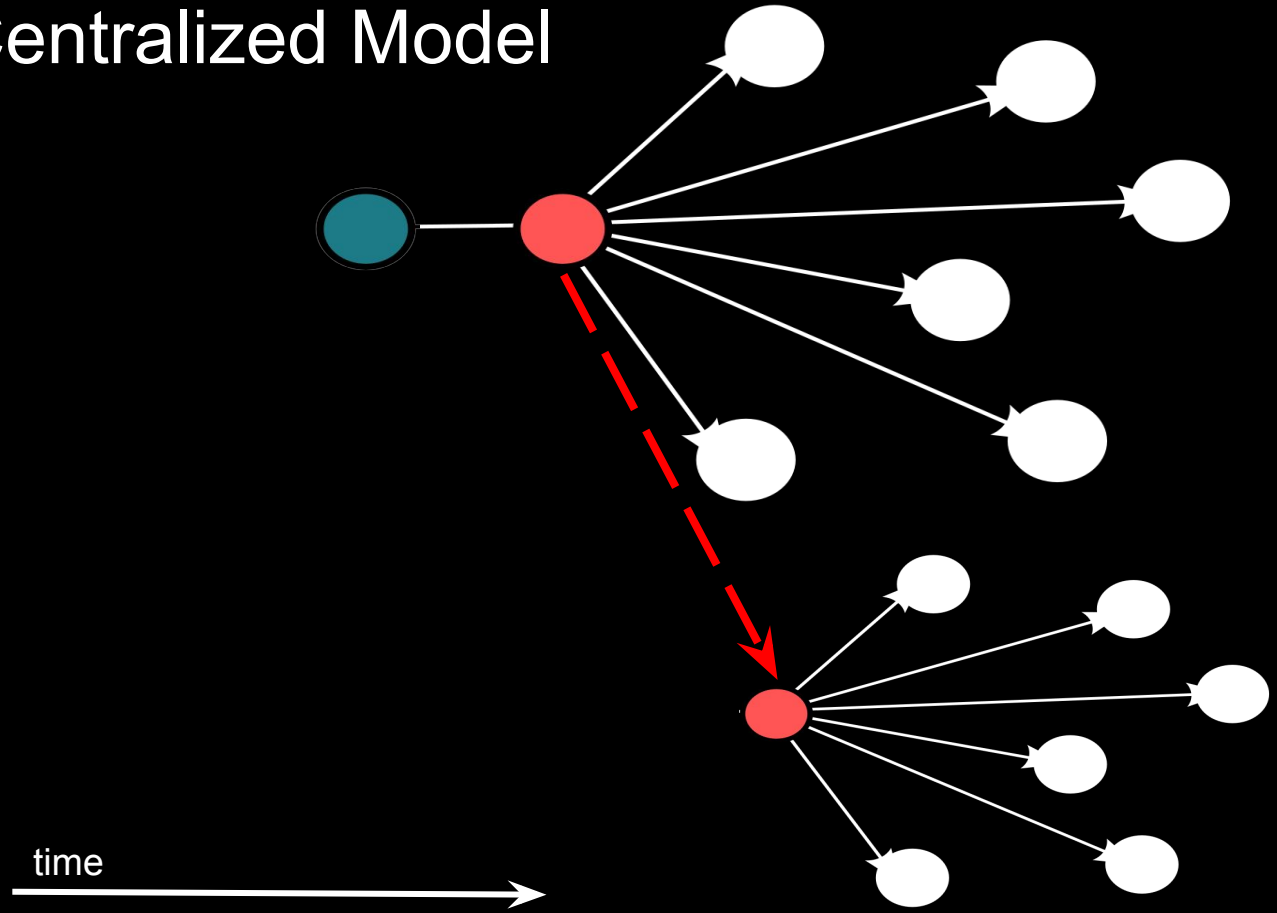

Broadcast Model (Public data)

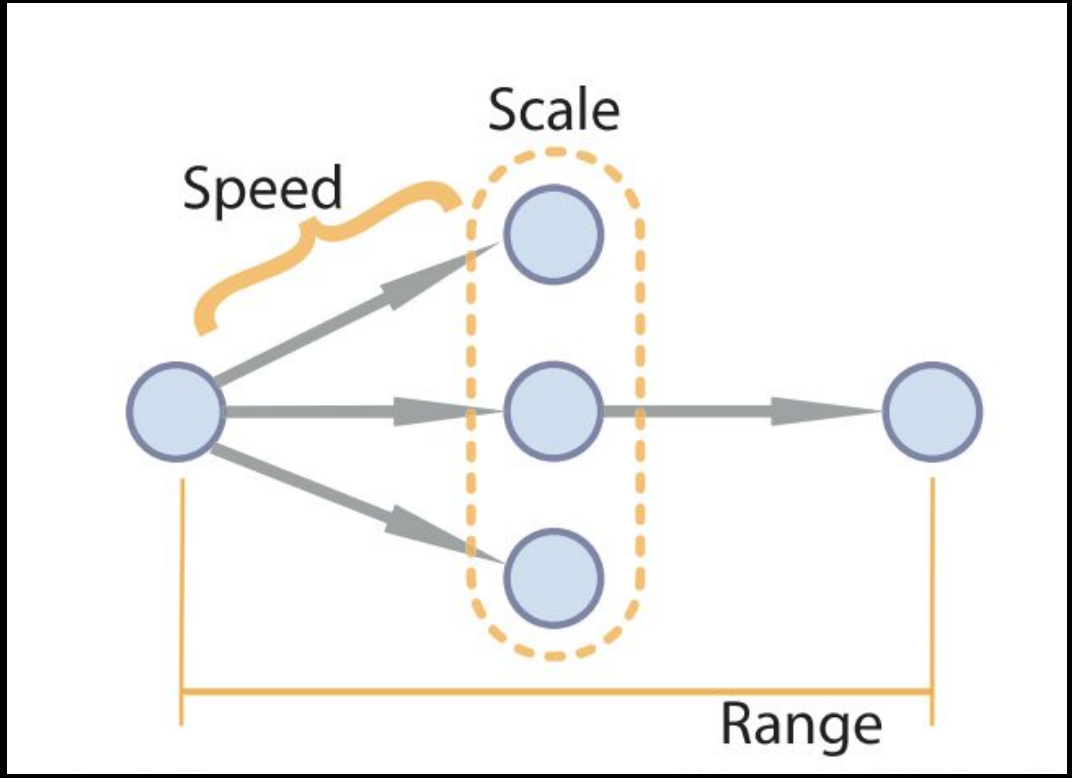

https://www.aaai.org/ocs/index.php/ICWSM/ICWSM10/paper/viewFile/1468/1896

Social Media in Emergency Response

The convergence of social networks and mobile has thrown the old response playbook out the window.

Michael Beckerman, president and CEO of the Internet Association.

During and immediately following Hurricane Sandy, “users sent more than 20 million Sandy-related tweets, despite the loss of cell phone service during the peak of the storm.”

FEMA National Preparedness report (2013)

https://www.ustelecom.org/emergency-911-service-and-the-role-of-social-media/

What’s an Event?

a thing that happens,
especially one of importance..

Event: 5W1H
What? Who? Where? When?
Why? How? (Xie et. al. ‘08)

Types of Emergency Events

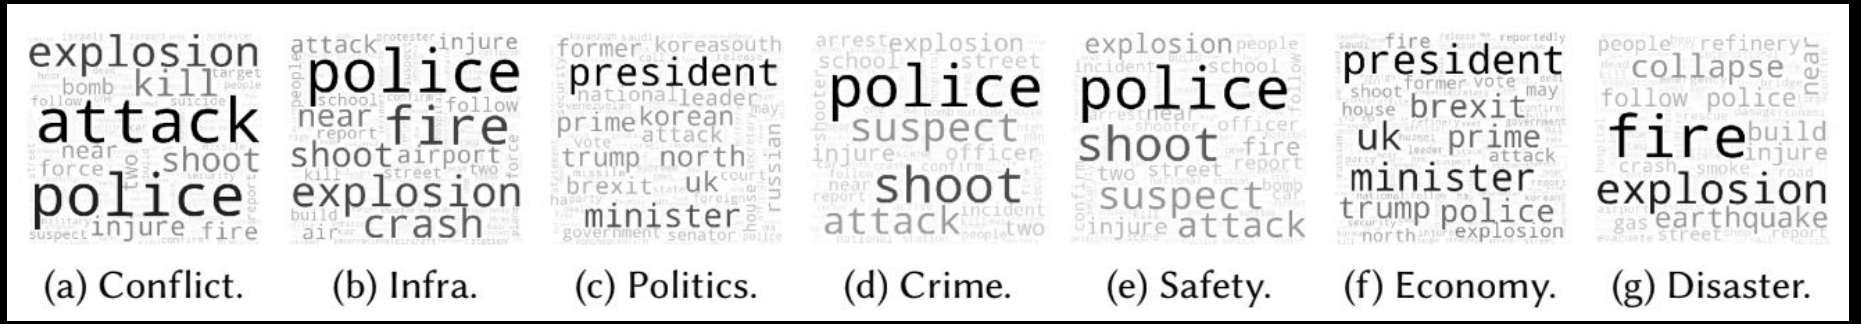

- Chemical
- Drought
- Earthquake
- Fire
- Flu
- Food Safety
- Heat Wave
- Highway Safety
- Hurricane
- Landslide
- Nuclear explosion
- Poisoning
- Power outage
- Terrorism
- Thunderstorm
- Tornado
- Tsunami
- Volcano
- Water Safety
- Wildfire
- Winter Storm

https://www.redcross.org/get-help/how-to-prepare-for-emergencies/types-of-emergencies.html

Why are emergencies so complicated?

- Unexpected events, speed
- Geographic spread
- So MUCH information (what is informative and what is not?)
- Fragmented information
- Logistics: efficient deployment of resources
- What’s important, to whom?

Important Events

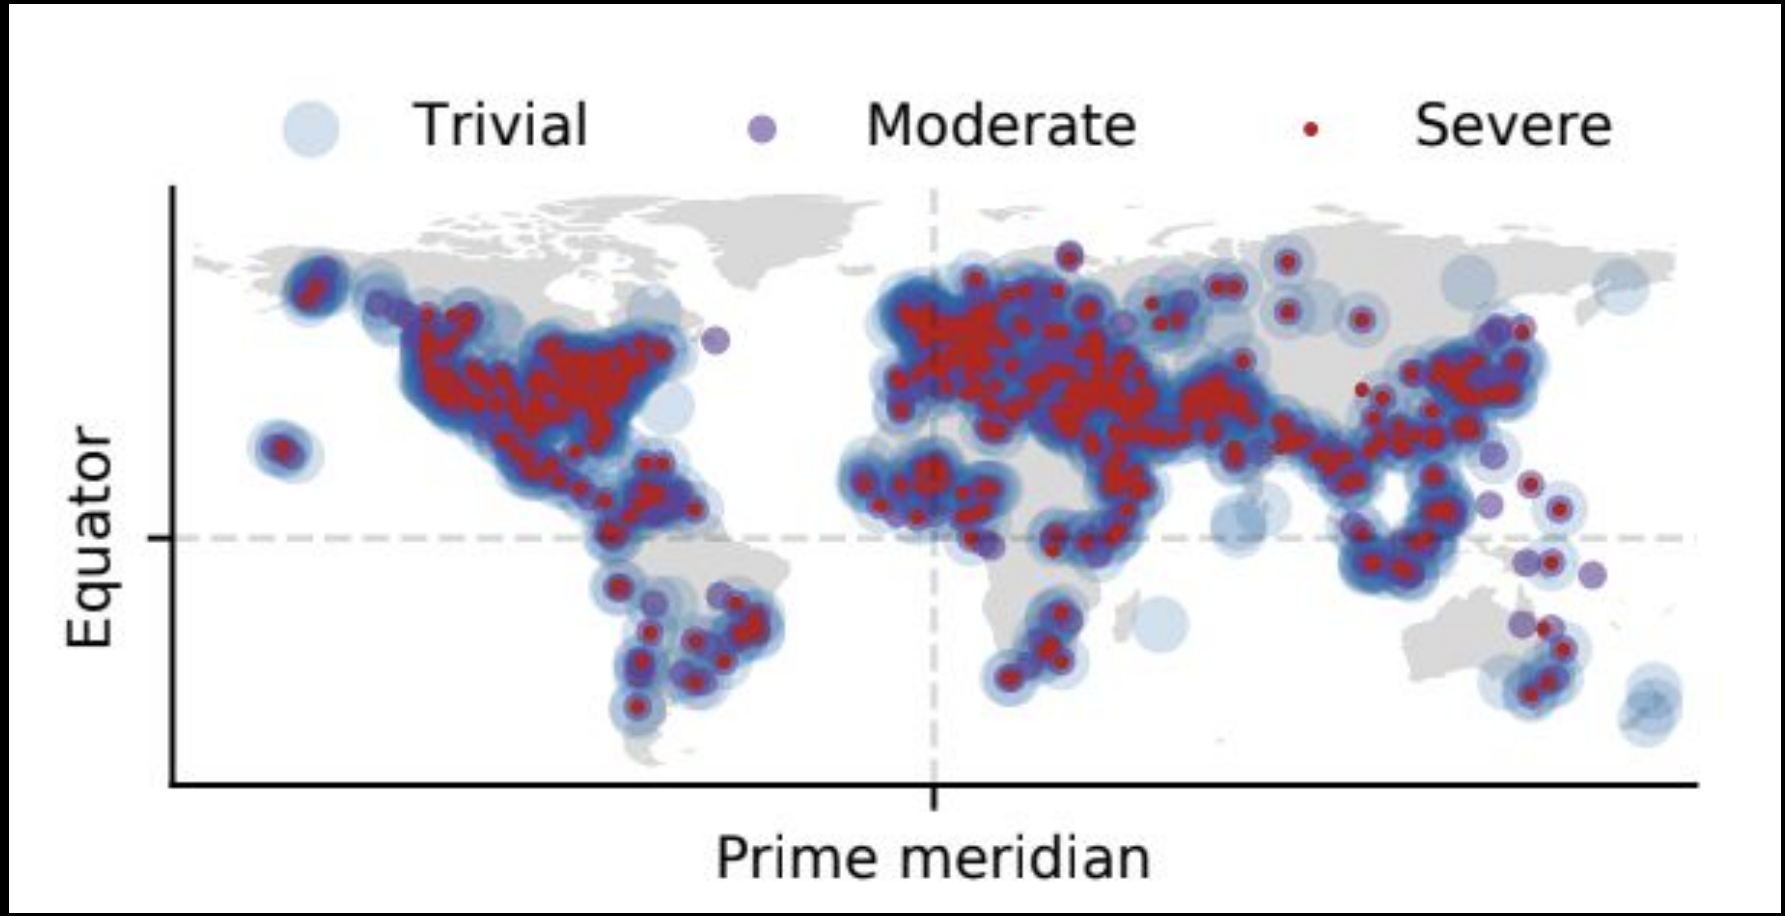

Temporal vs Spatial Distribution

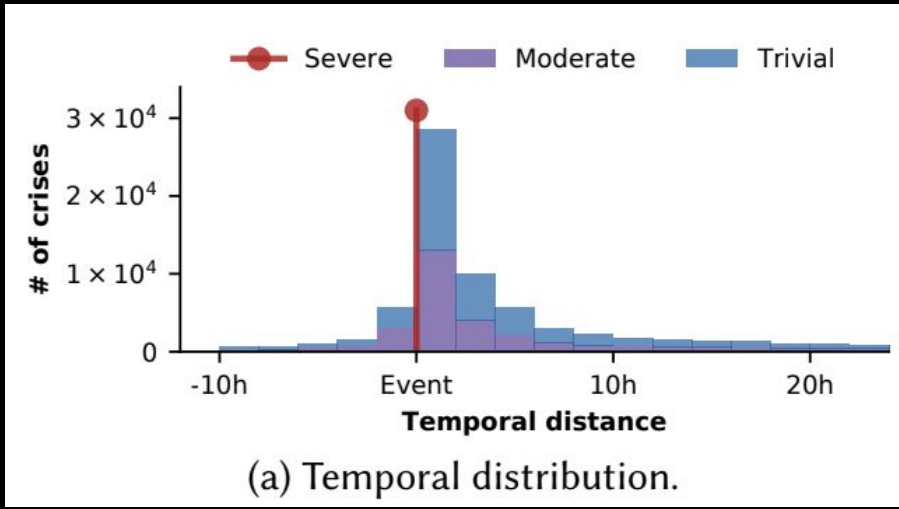

(a) Temporal distribution.

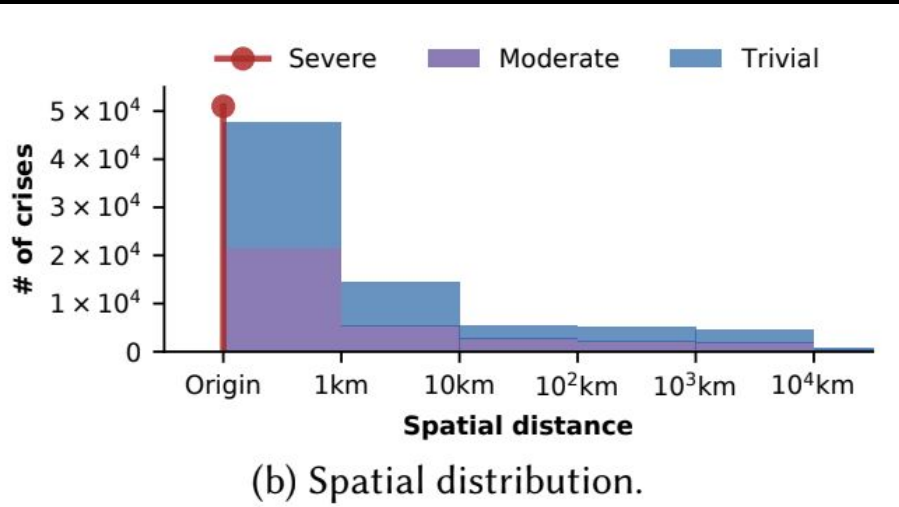

(b) Spatial distribution.

Frequency of Emergencies

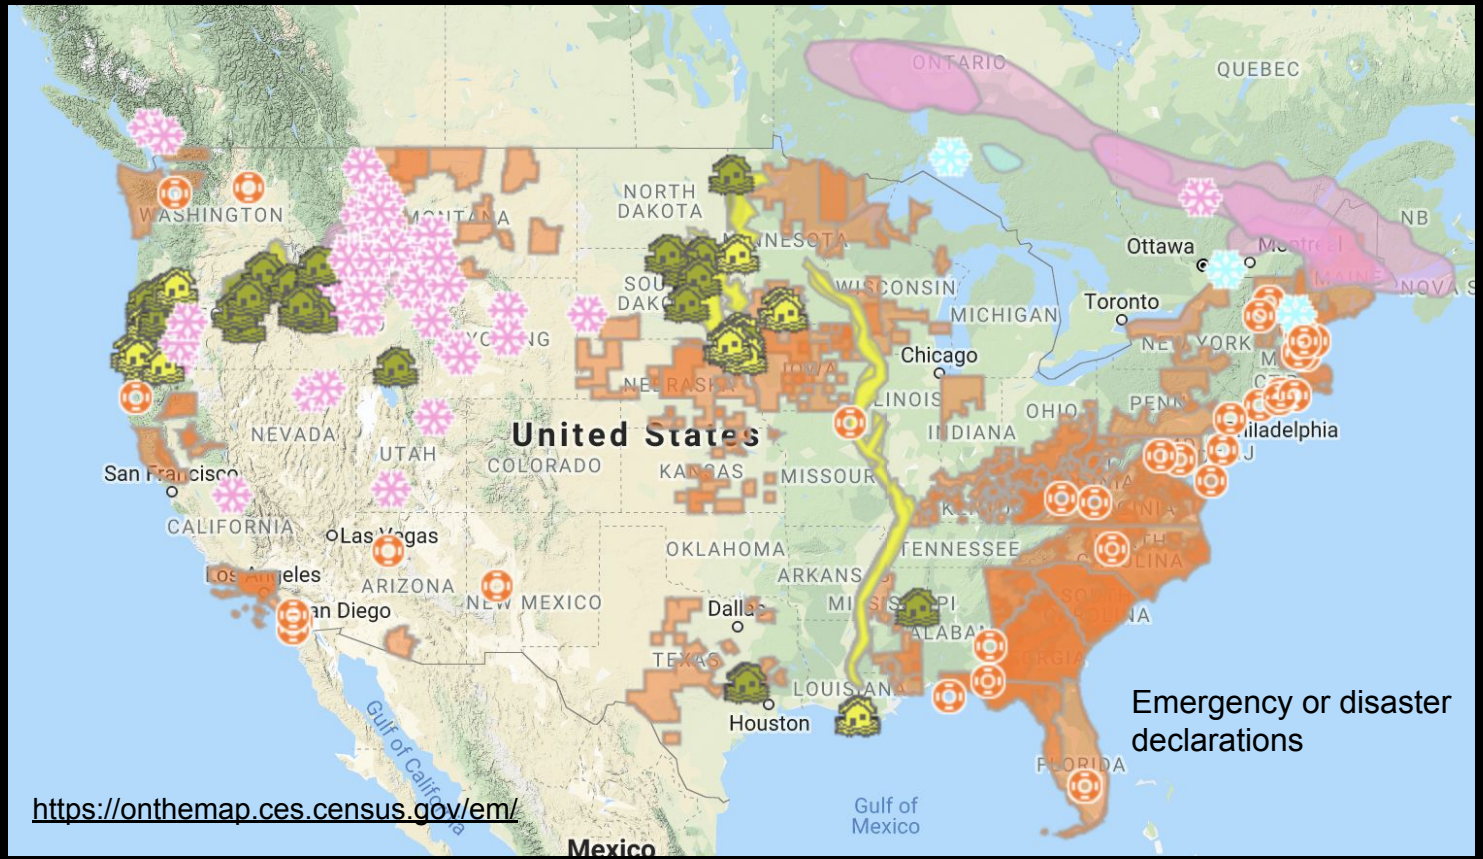

Number of wildfires in the United States

The number of wildfire events in the United States. Data is shown from 1983 onwards, when comparable data records began.

Our World in Data

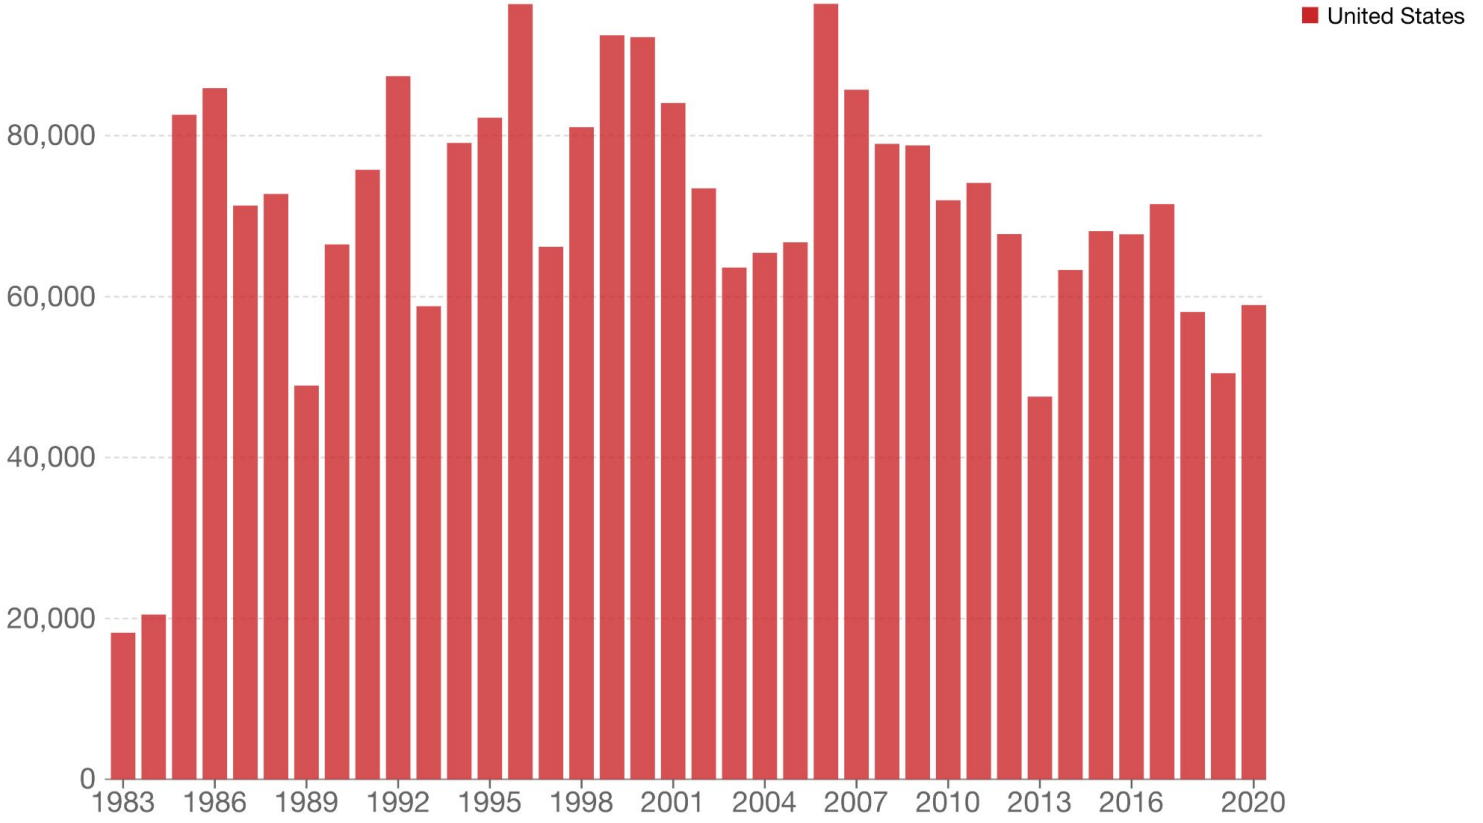

Source: National Interagency Fire Center (NIFC)

Number of recorded natural disaster events, 1900 to 2023

Our World in Data

The number of global reported natural disaster events in any given year. Note that this largely reflects increases in data reporting, and should not be used to assess the total number of events.

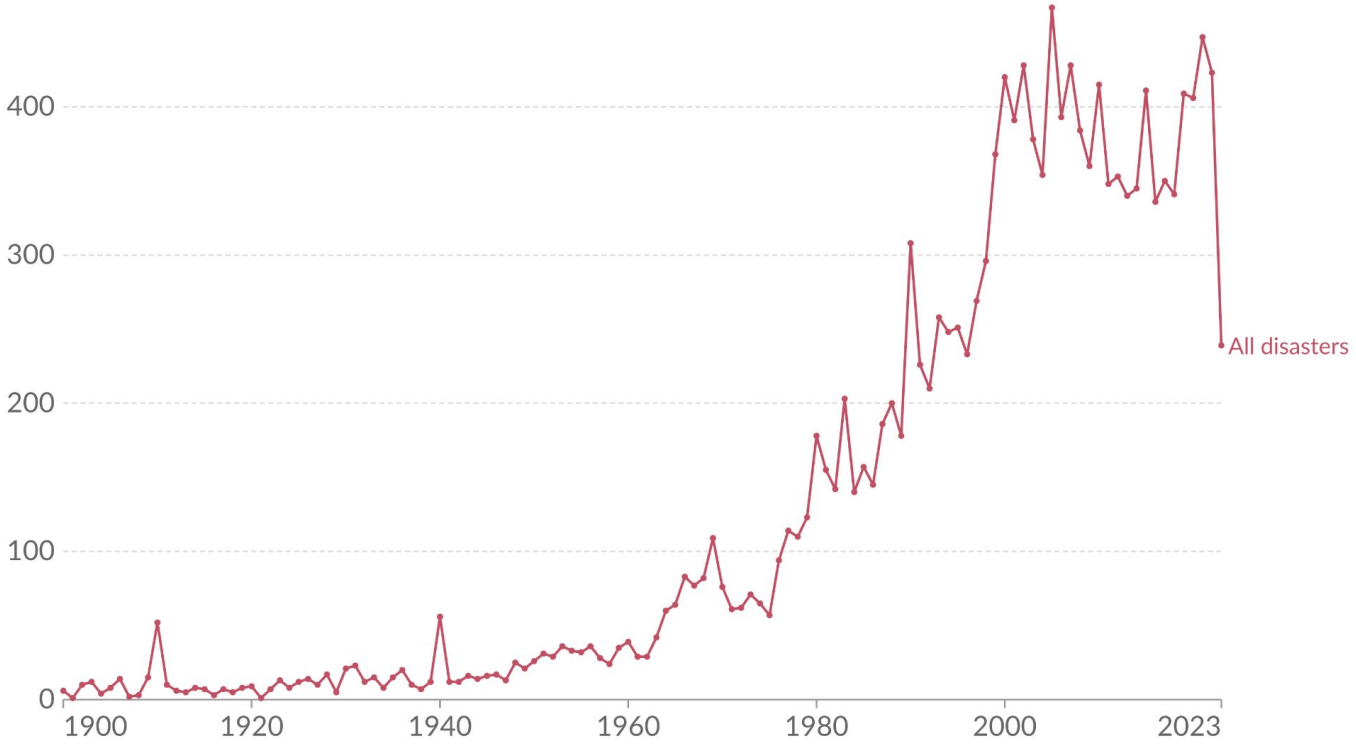

Data source: EM-DAT, CRED / UCLouvain (2023)
Note: Data includes disasters recorded up to September 2023.

OurWorldInData.org/natural-disasters | CC BY

https://ourworldindata.org/grapher/number-of-natural-disaster-events?time=earliest..latest

The Largest Peacekeeping Operations in 2023

Personnel involved in the largest active multilateral peacekeeping operations as of May 31, 2023*

* Includes troops, experts on mission, formed police units, police and staff officers

Source: UN Peacekeeping

statista

Jason Patinkin/The Associated Press

Global deaths in conflicts since the year 1400 – by Max Roser

Our World in Data

Each circle represents one conflict. [Data from the *Conflict Catalog* (1400-2000)]
The **size** represents the absolute number of fatalities (military + civilian fatalities
The **position** on the y-axis represents the fatality rate* (military + civilian fatalities)

Military + civilian death rate* for 1400-2000 [Data from *Conflict Catalog*] – 15 year moving-average

Military death rate* for 1946-2013 [Data from the PRIO Institute]

* All death rates are calculated as the share of fatalities relative to the world population at the time (rate of deaths per 100,000 people).

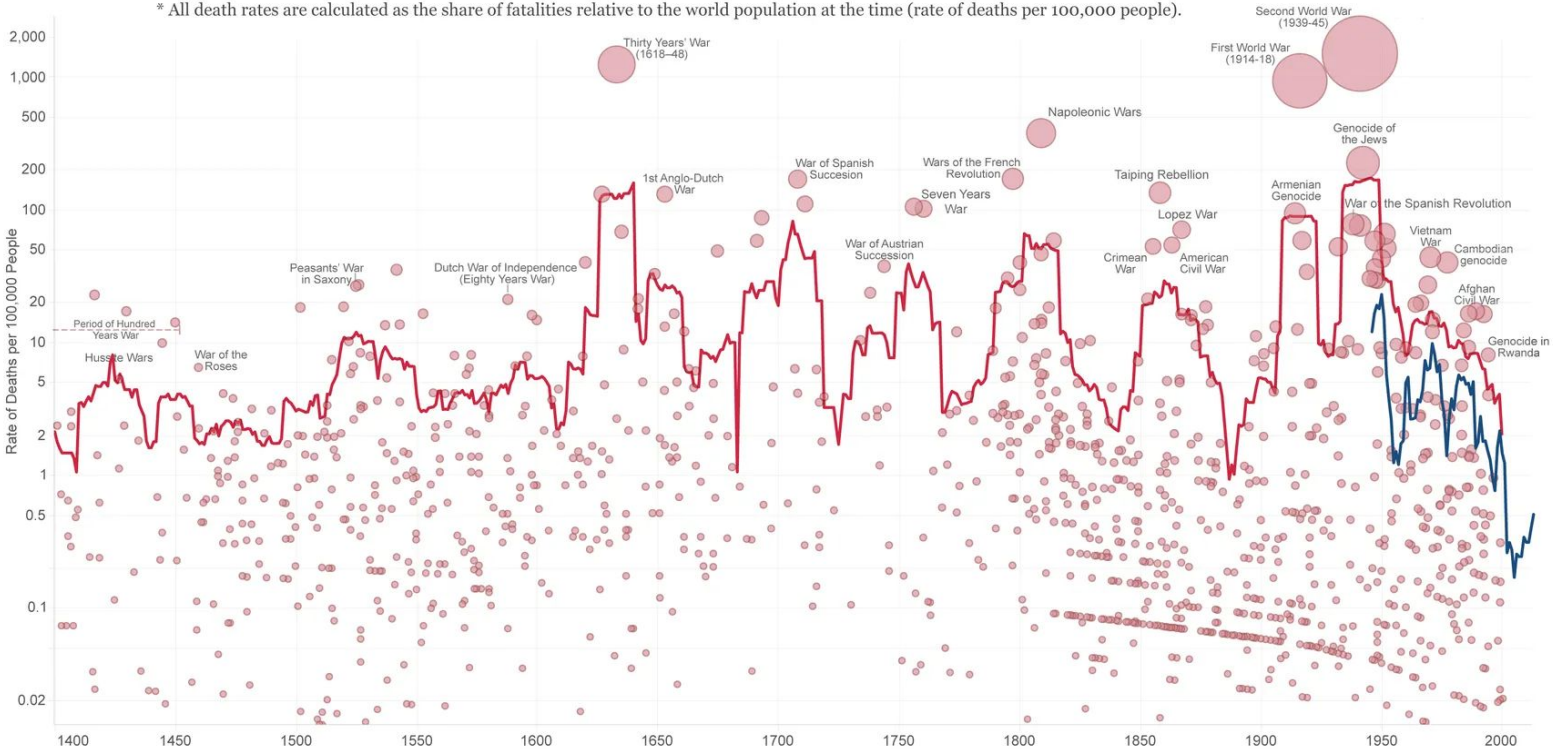

Data sources: Battle Deaths Dataset v.3.0. published by the PRIO Institute and Conflict Catalog by Peter Brecke for data on battle deaths. And world population data from HYDE and UN.

This is a data visualisation from OurWorldinData.org. There you find more visualisations on this topic.

https://www.vox.com/2015/6/23/8832311/war-casualties-600-years

World-wide battle deaths per 100,000 people

Colonial | Interstate | Civil | Civil with foreign interaction

Human Security Report Project, Uppsala Conflict Data Project, Peace Resarch Institute of Oslo, via WSJ

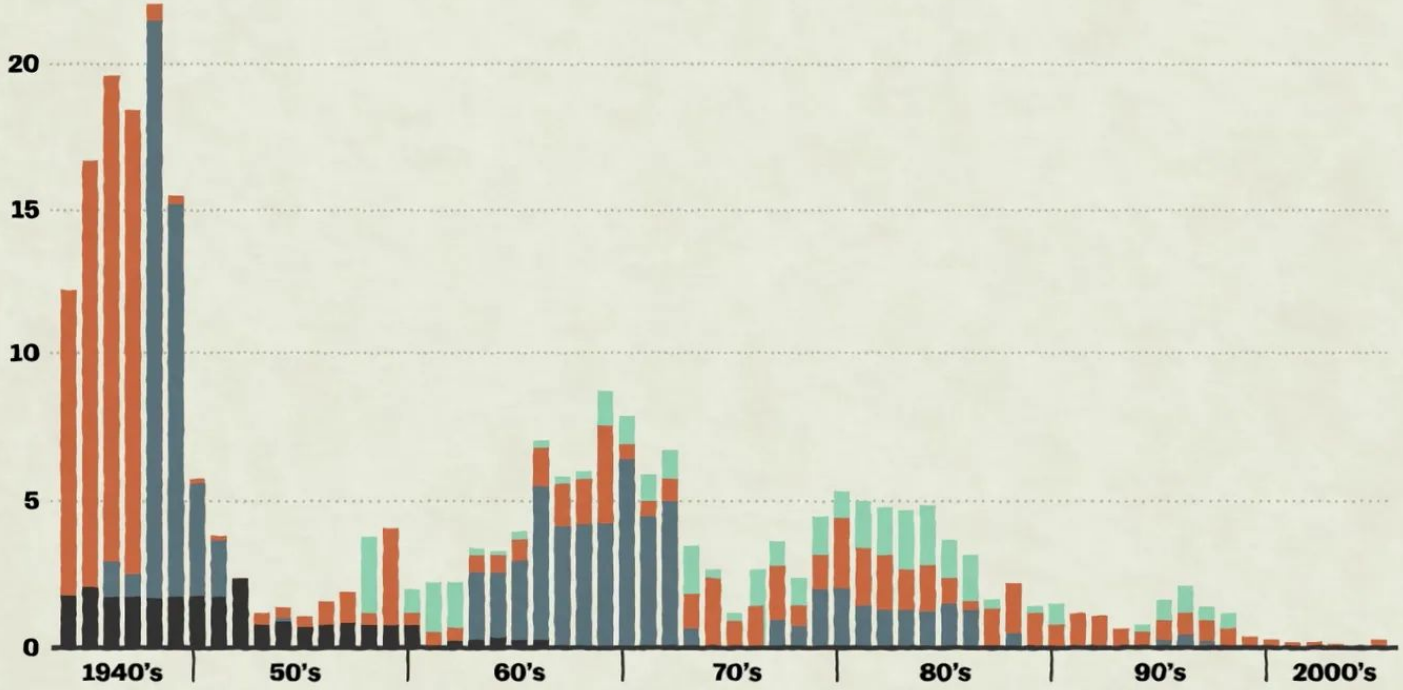

https://www.vox.com/2015/6/23/8832311/war-casualties-600-years

Use Cases

- Public sector, NGOs
- Corporate risk (all industries)
- News & journalism
    - Physical
    - Cyber
    - Executive protection

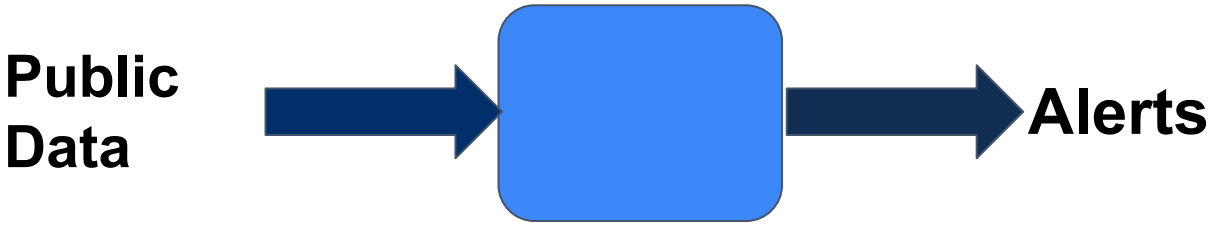


AI

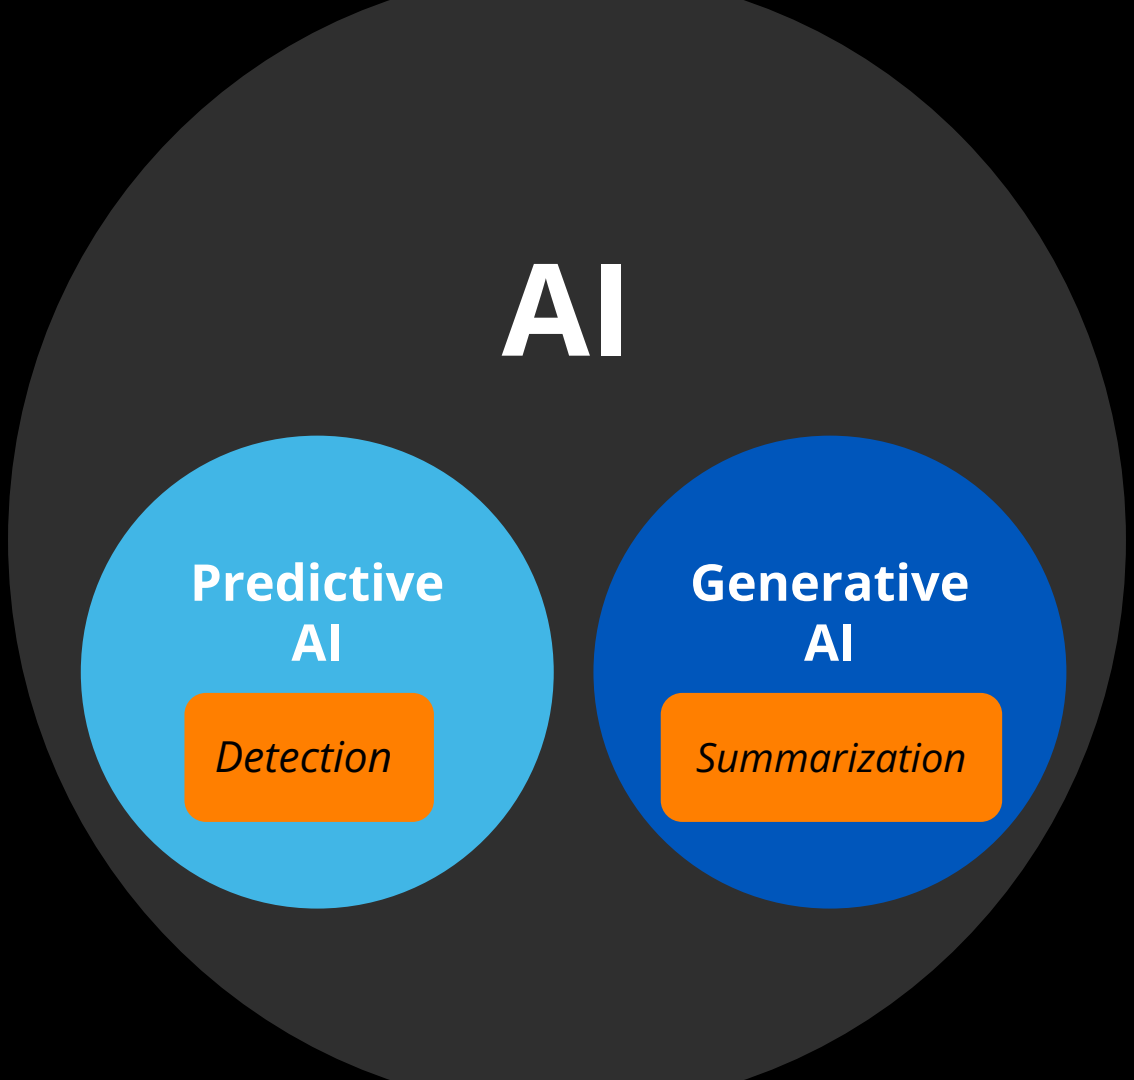
AI
Predictive AI
Detection
Generative AI
Summarization

AI

Dataminr®

https://www.vox.com/science-and-health/2019/1/4/17989224/intellectual-humility-explained-psychology-replication?fbclid=IwAR0sCKV_r8cAqVkX02AXYyeI1LJvMEBxMubcFfxR3vivi7Ab0Iqm2UnACnM

AI from 380 BC to 1900

AI from 1900-1950

AI in the 1950s

AI in the 1960s

AI in the 1970s

AI in the 1980s

AI in the 1990s

AI from 2000-2010

AI 2010 to present day

https://www.forbes.com/sites/gilpress/2016/12/30/a-very-short-history-of-artificial-intelligence-ai/?sh=1f1171fb6fba

AI Model

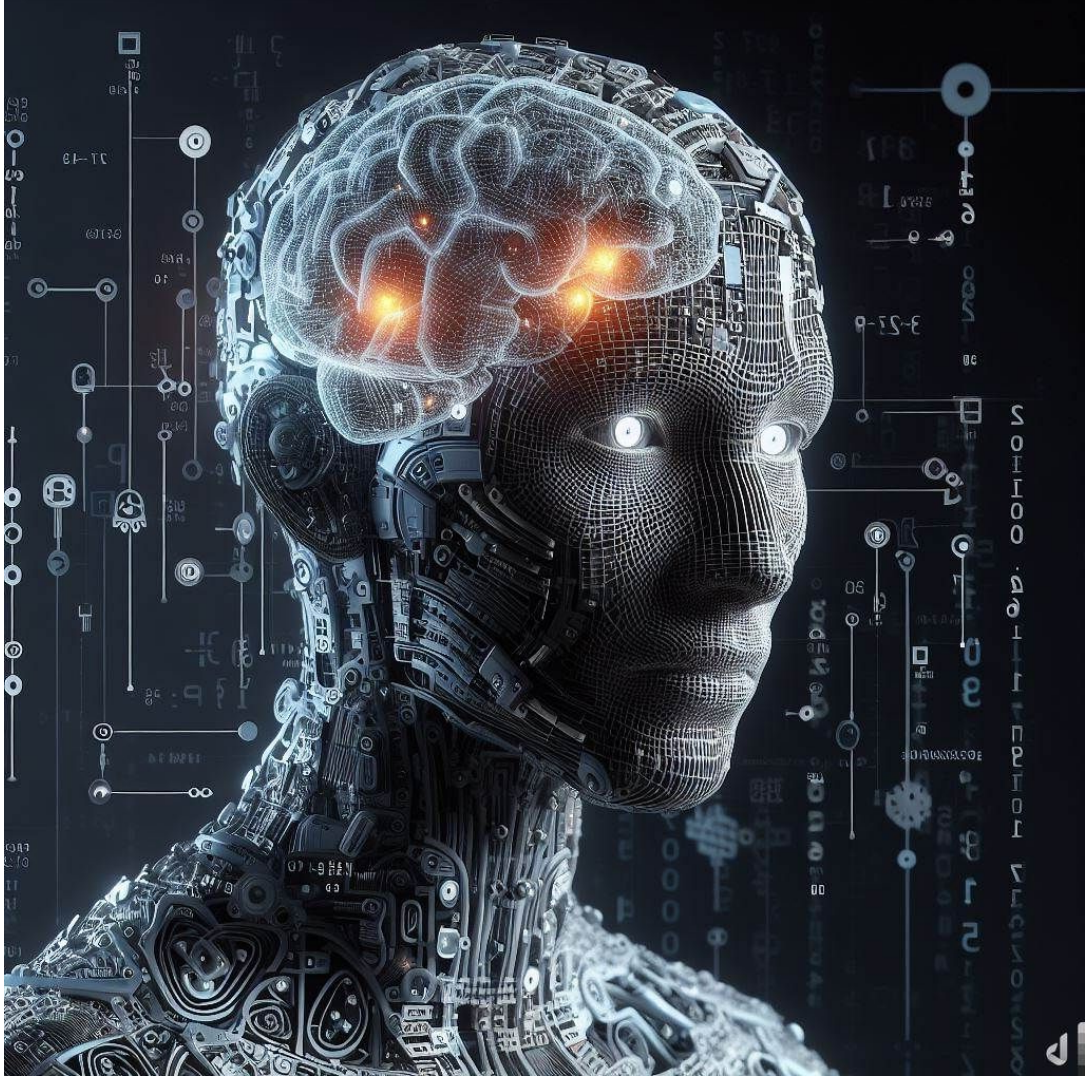

Dataminr

AI Model

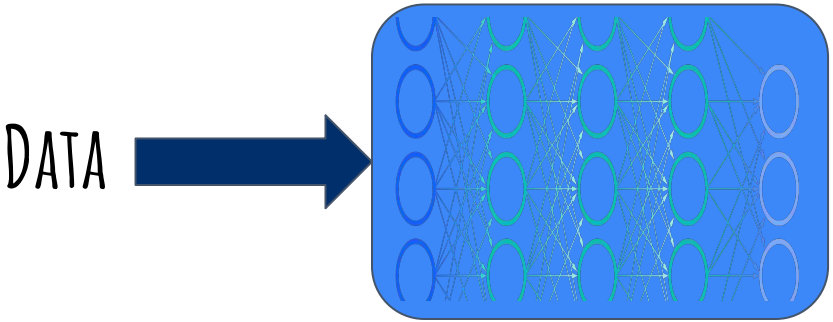

AI Model

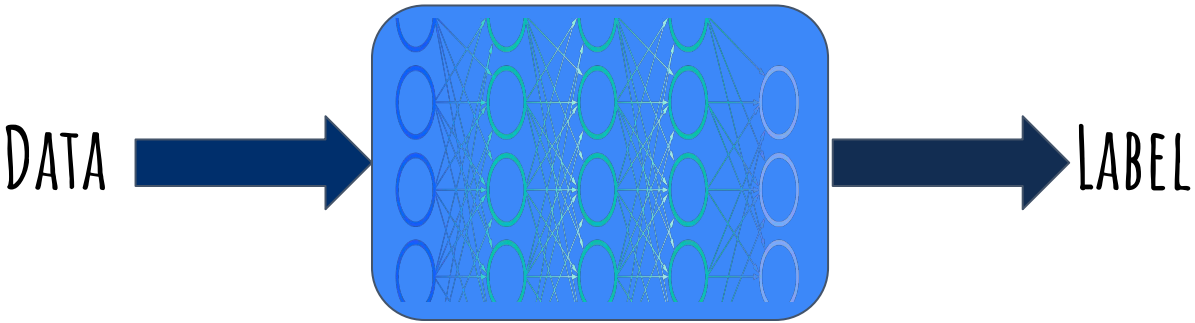

AI Model

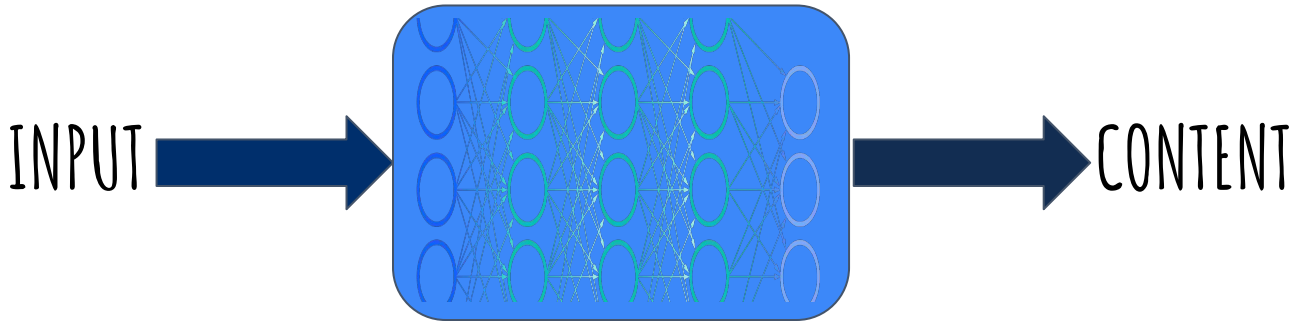

Dataminr®

Eliza
1966

Men are all alike. (human)

IN WHAT WAY (machine)

They're always bugging us about something or other.

CAN YOU THINK OF A SPECIFIC EXAMPLE

Well, my boyfriend made me come here.

YOUR BOYFRIEND MADE YOU COME HERE

He says I'm depressed most of the time.

I AM SORRY TO HEAR YOU ARE DEPRESSED

It's true. I am unhappy.

DO YOU THINK COMING HERE WILL HELP YOU NOT TO BE UNHAPPY

Data

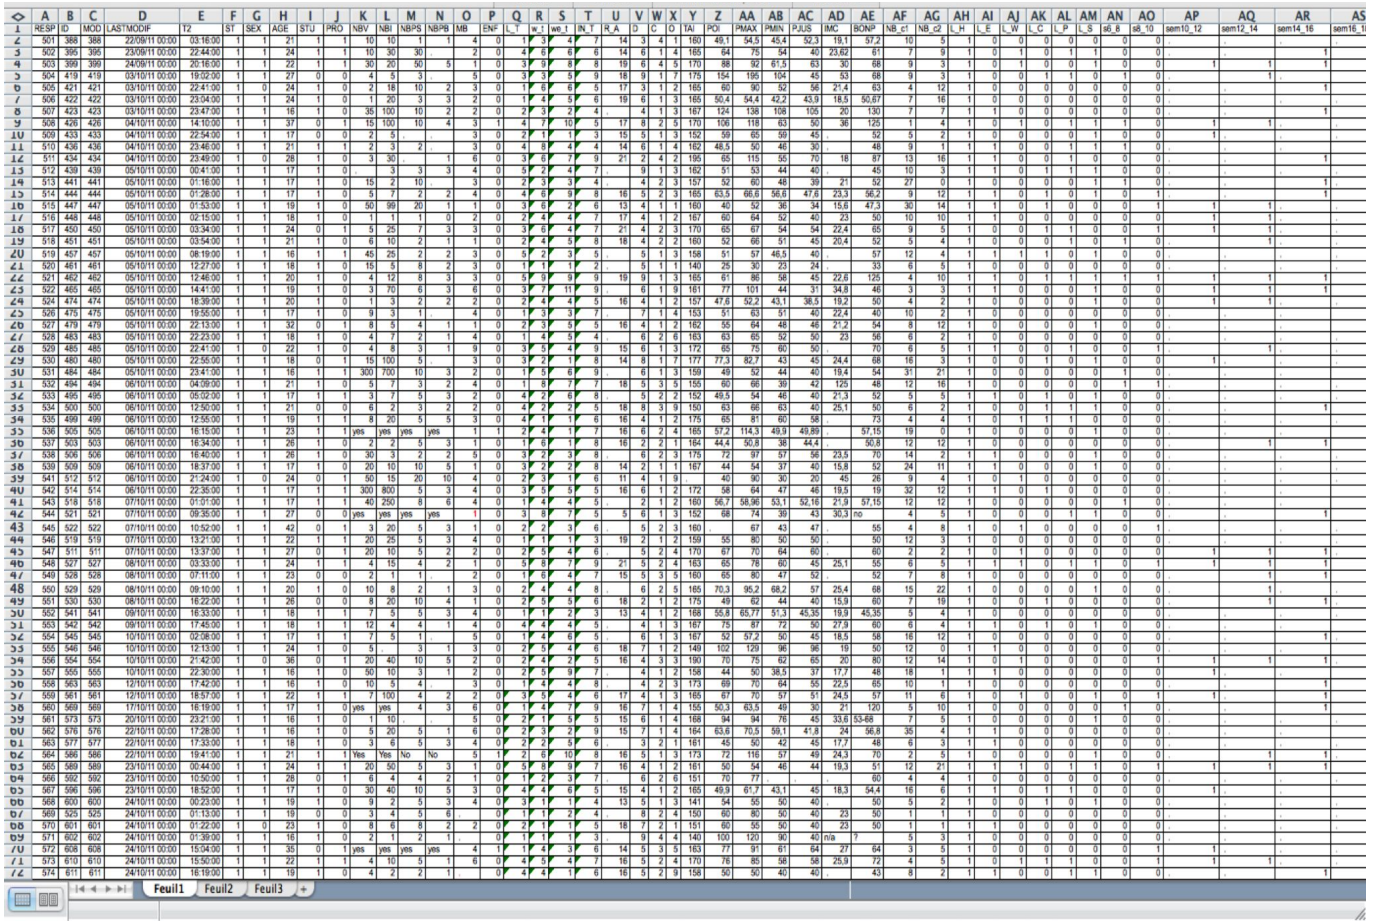

Dataminr

AI Model

Dataminr®

AI

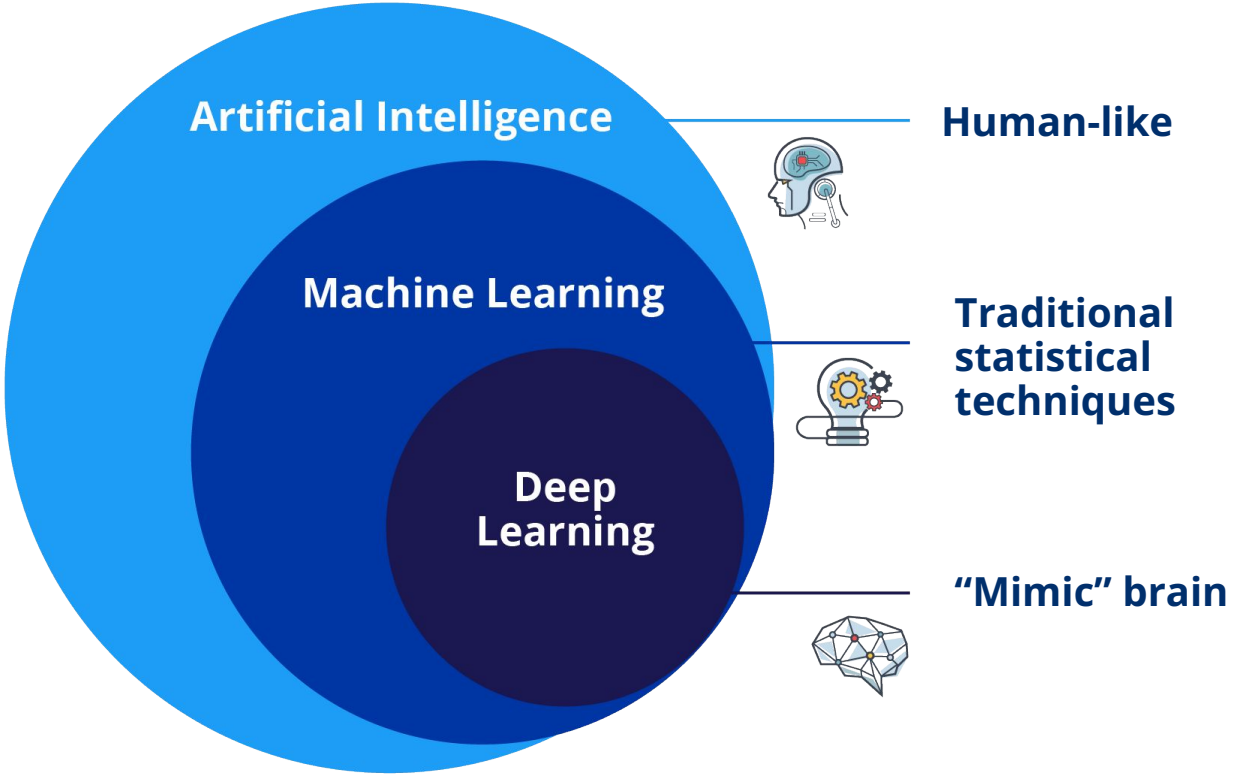
Artificial Intelligence
Machine Learning
Deep Learning
Human-like
Traditional statistical techniques
“Mimic” brain

Dataminr®

Deep Learning

Dataminr

https://towardsdatascience.com/everything-you-need-to-know-about-neural-networks-and-backpropagation-machine-learning-made-easy-e5285bc2be3a

AI

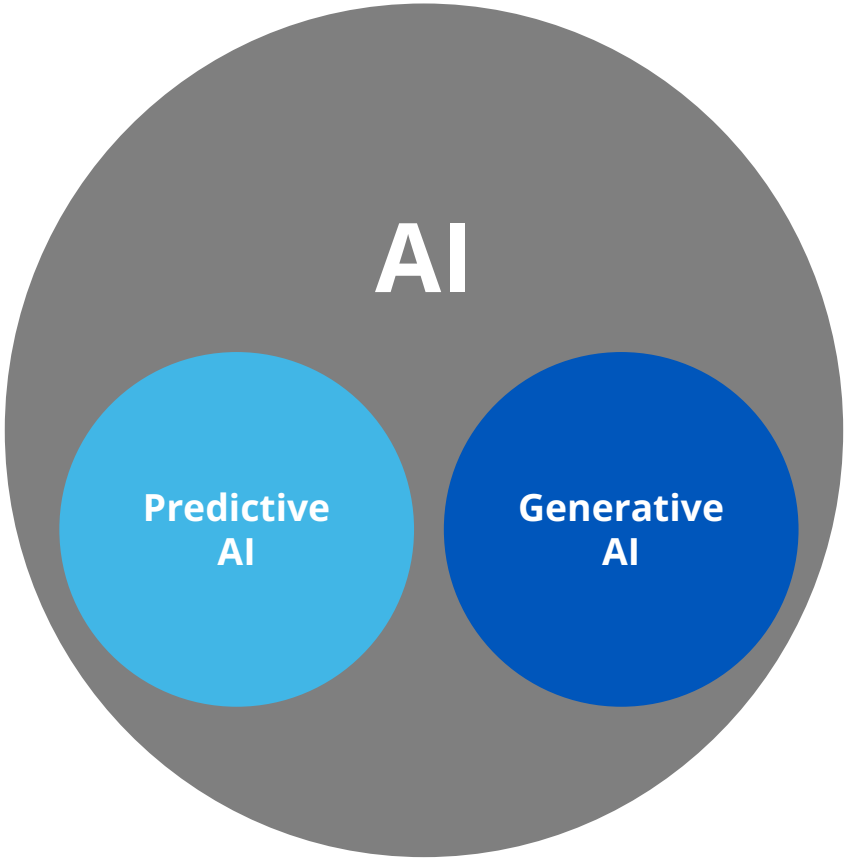

Applications of

AI

Advertising

Search

Spam

News

Manufacturing

Customer Support

Social Media

Transportation

Everywhere there are repetitive patterns

AI

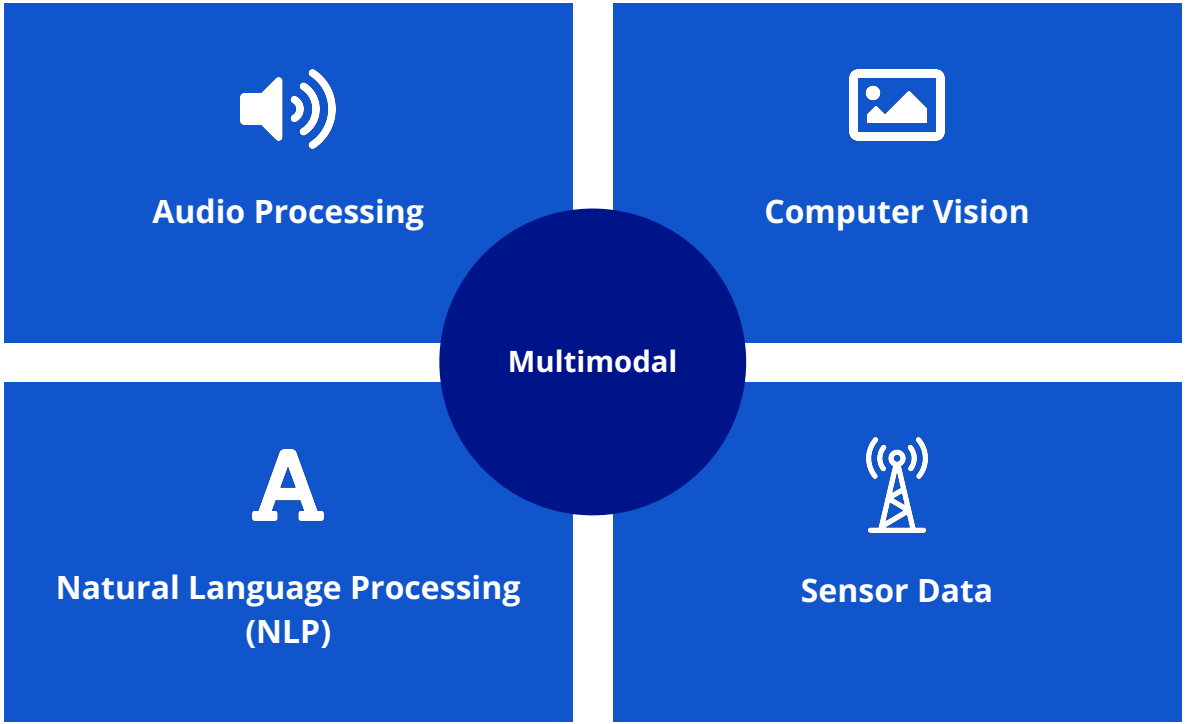
Audio Processing
Computer Vision
Multimodal
Natural Language Processing (NLP)
Sensor Data

Computer Vision

Computer Vision

Given an image or video…

identify and label scene, objects, actions, etc.

Computer Vision

Dataminr®

https://medium.com/@ageitgey/machine-learning-is-fun-part-3-deep-learning-and-convolutional-neural-networks-f40359318721#.wukv4xnp4

Audio


What do you hear!?!
YANNY
LAUREL

Brief Intro to LLMs

(Large Language Models)

What is an LLM?

Nothing is _____

Dataminr®

LLMs

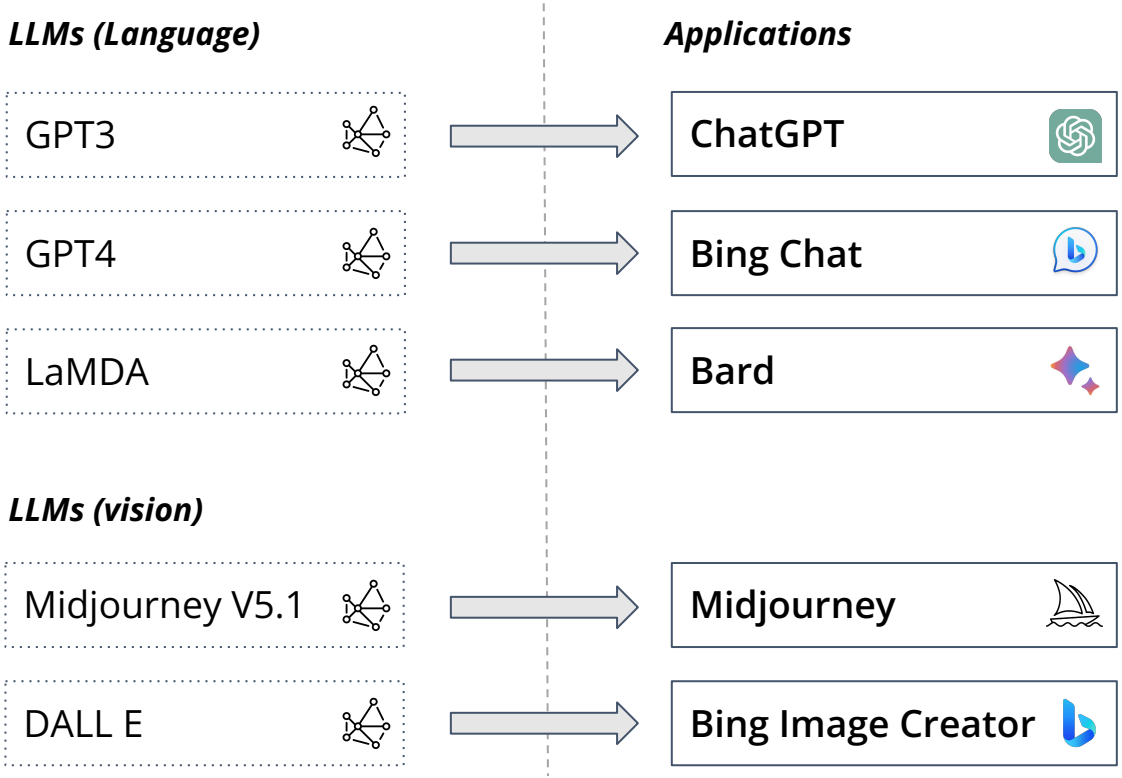
LLMs (Language)
Applications
GPT3
ChatGPT
GPT4
Bing Chat
LaMDA
Bard
LLMs (vision)
Midjourney V5.1
Midjourney
DALL E
Bing Image Creator

Building LLMs

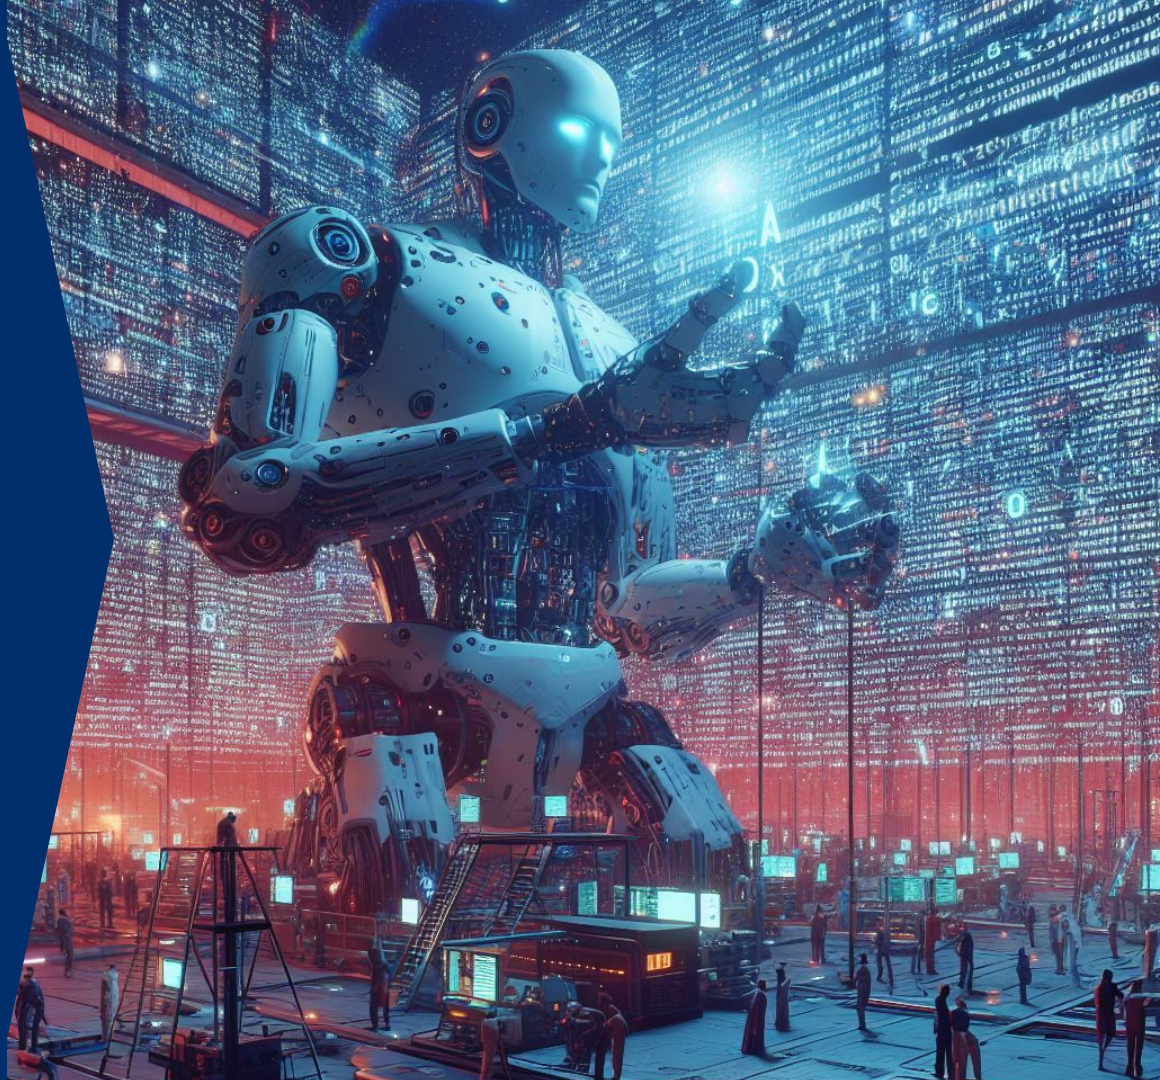

Foundation Models

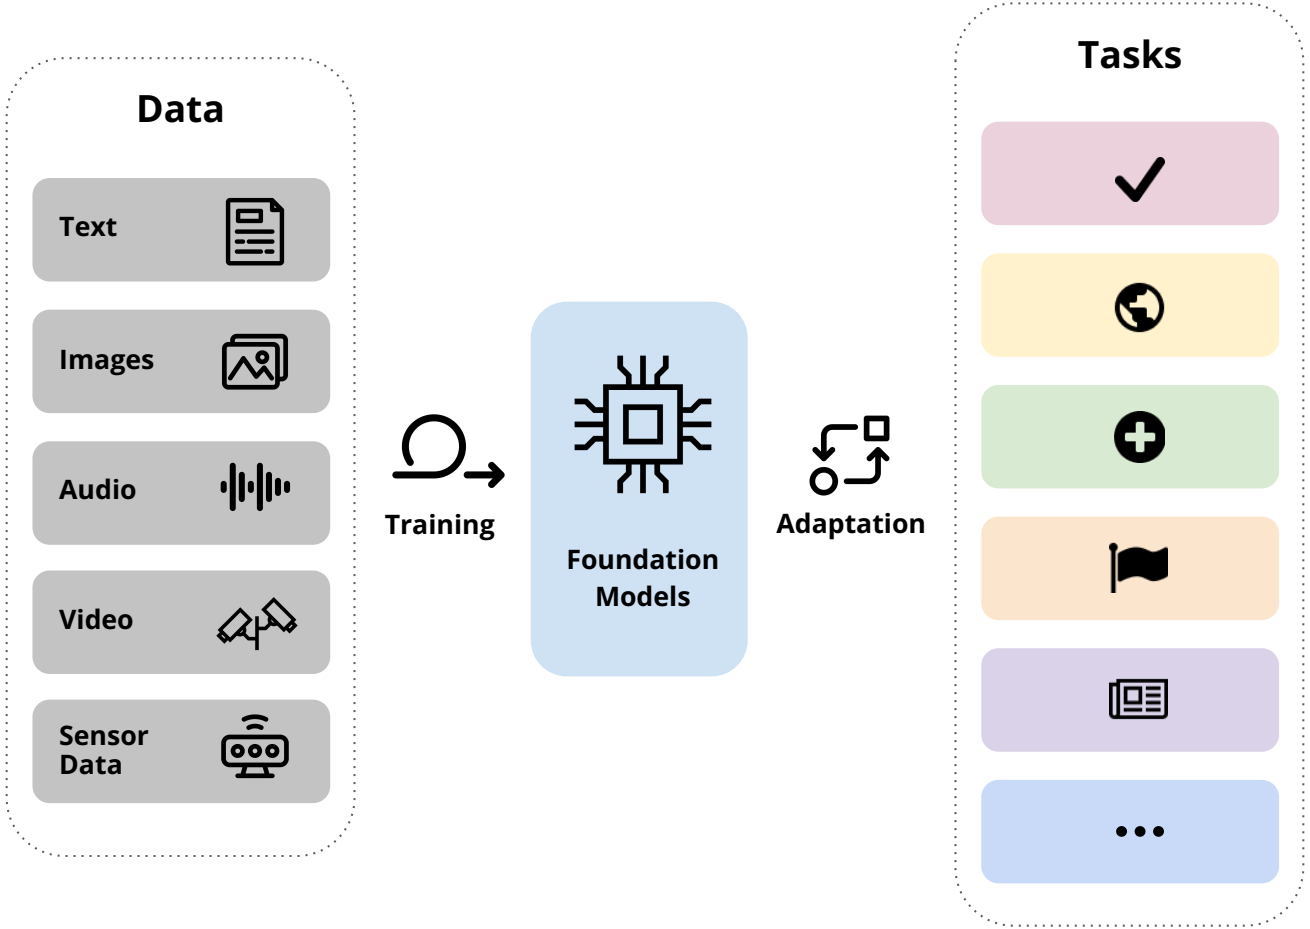

BERT & Friends

Dataminr

LLMs

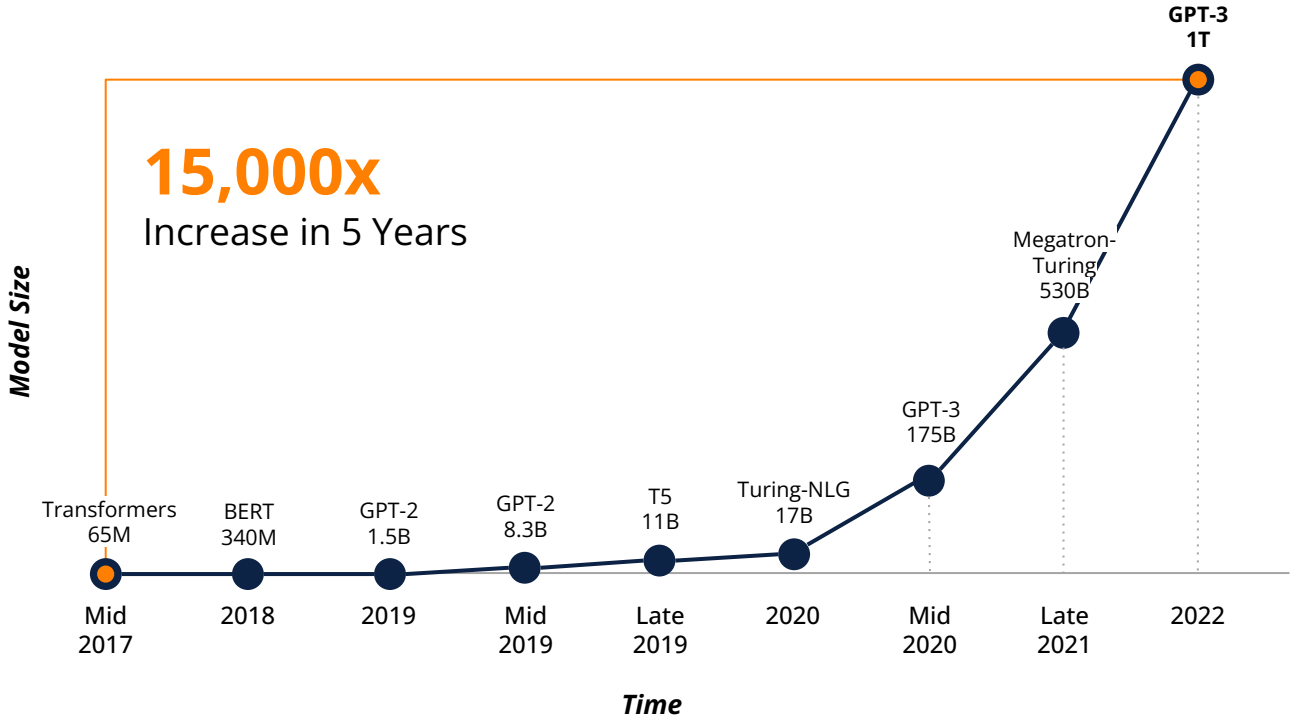
15,000x
Increase in 5 Years
Model Size
Time
Transformers 65M
BERT 340M
GPT-2 1.5B
GPT-2 8.3B
T5 11B
Turing-NLG 17B
GPT-3 175B
Megatron-Turing 530B
GPT-3 1T
Mid 2017
2018
2019
Mid 2019
Late 2019
2020
Mid 2020
Late 2021
2022

LLMs

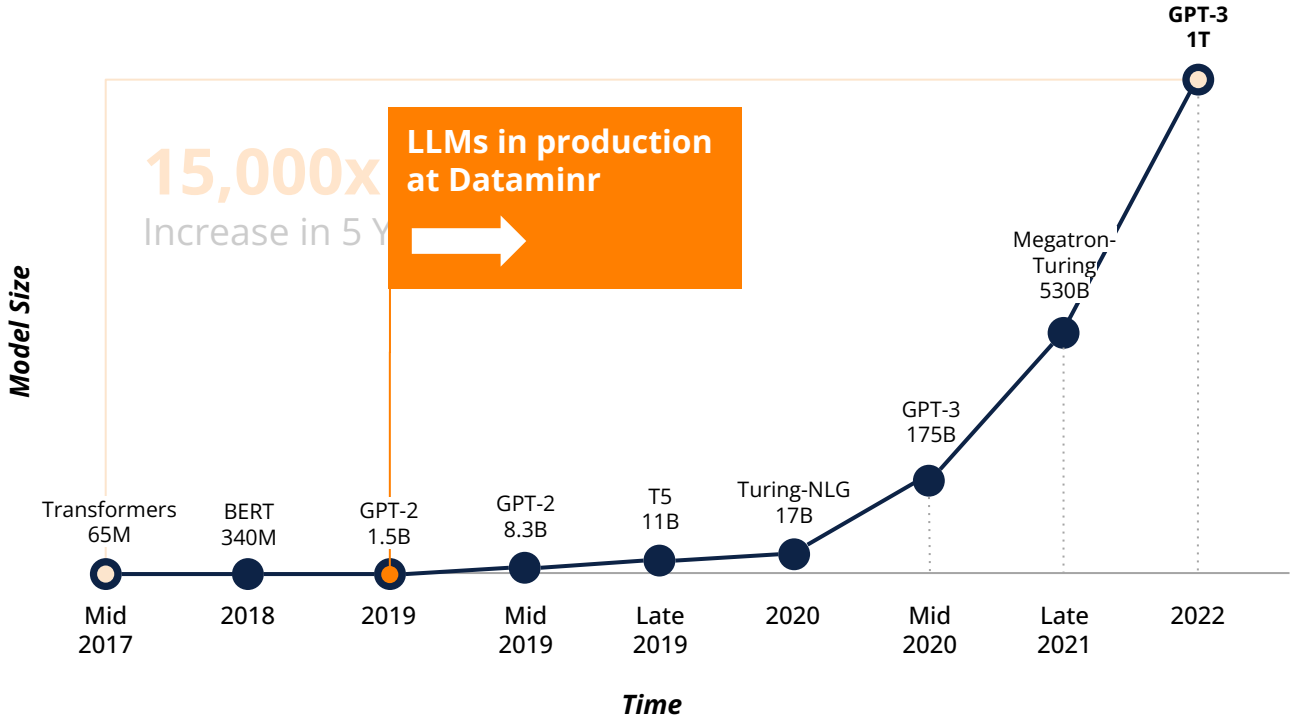
15,000x
Increase in 5 Y
LLMs in production at Dataminr
Model Size
Transformers 65M
BERT 340M
GPT-2 1.5B
GPT-2 8.3B
T5 11B
Turing-NLG 17B
GPT-3 175B
Megatron-Turing 530B
GPT-3 1T
Mid 2017
2018
2019
Mid 2019
Late 2019
2020
Mid 2020
Late 2021
2022
Time

LLM Training

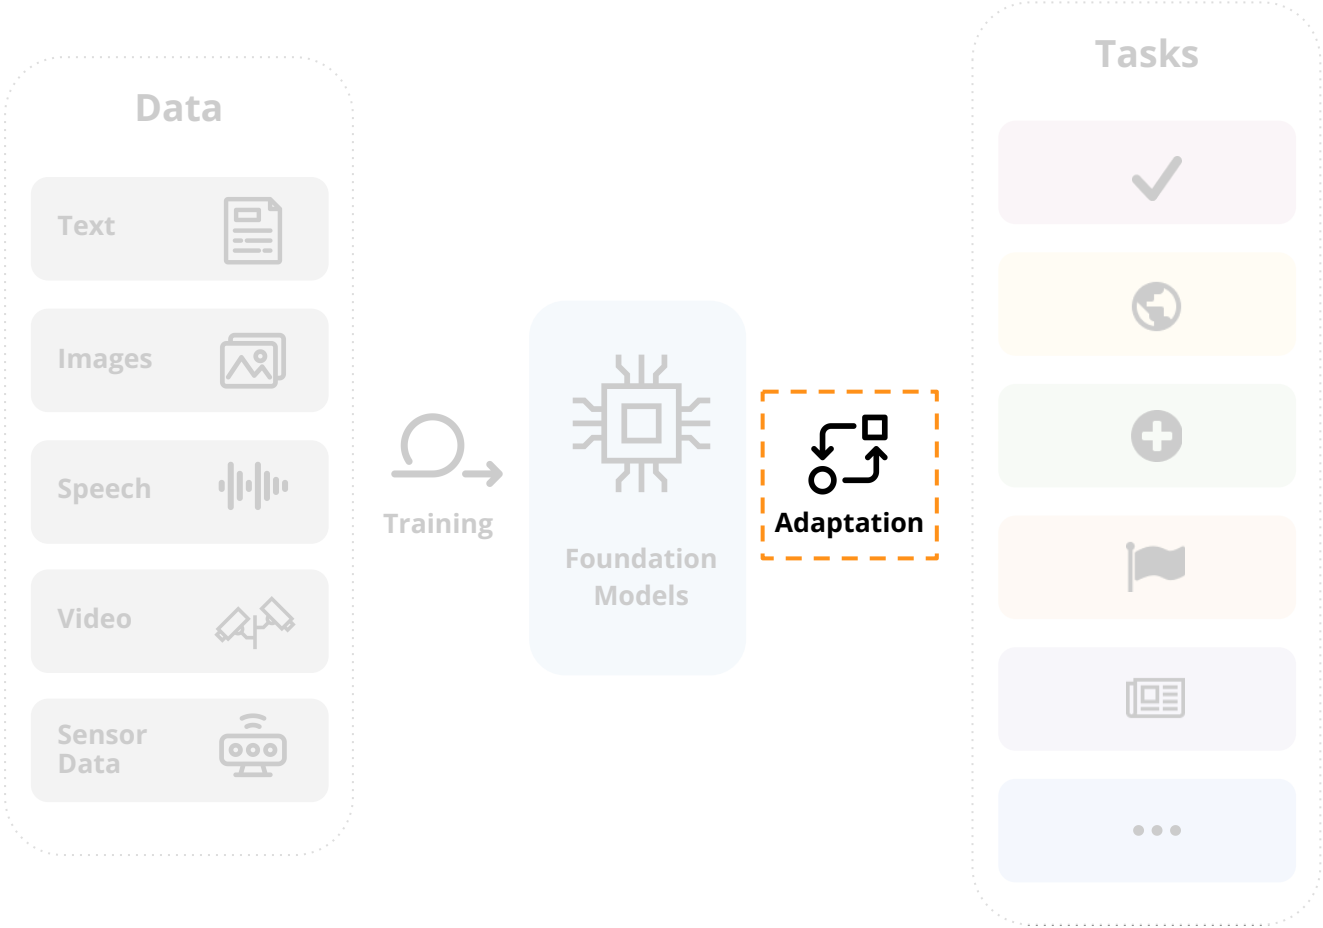

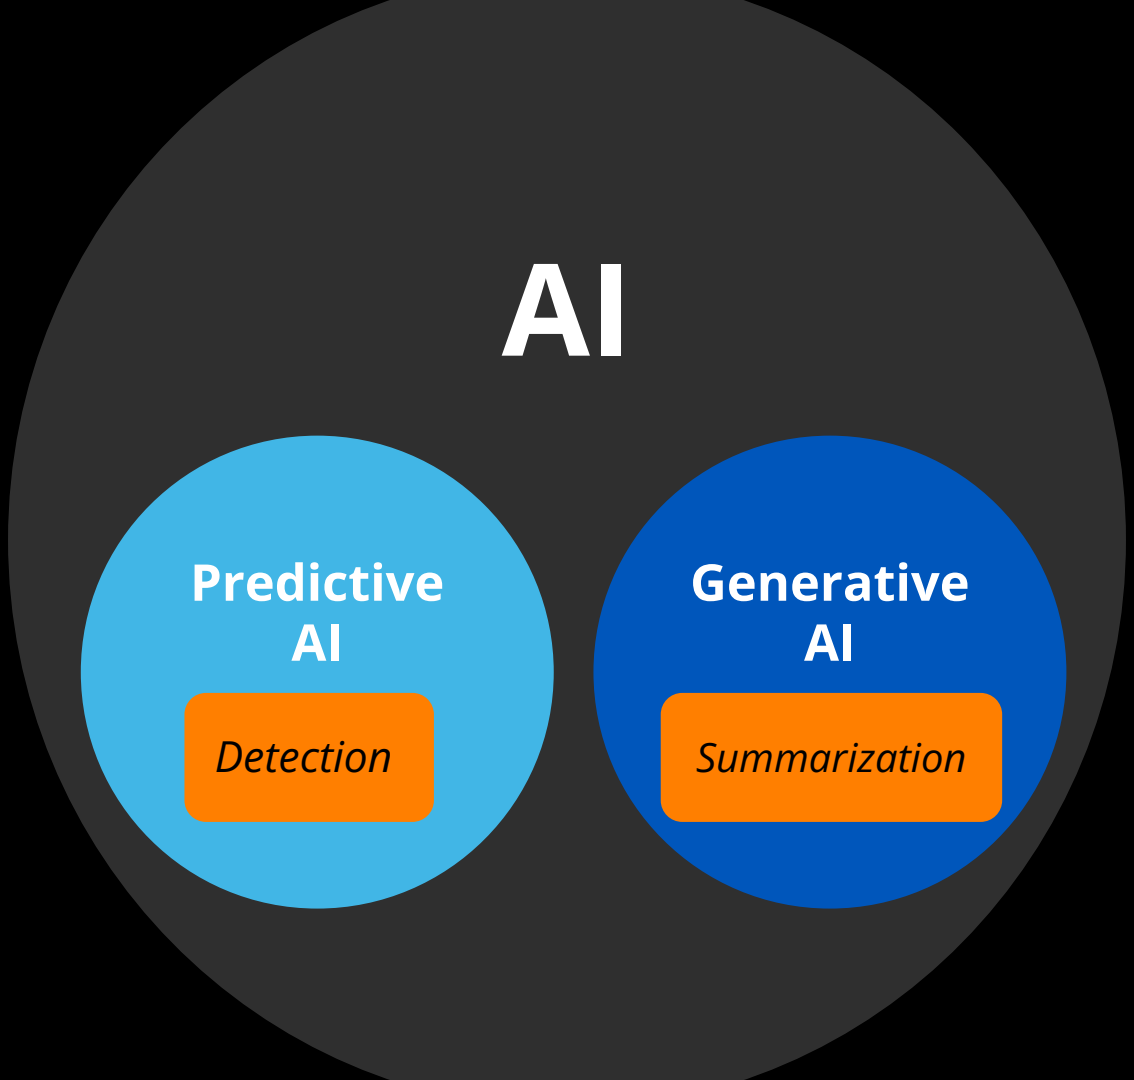
AI
Predictive AI
Detection
Generative AI
Summarization

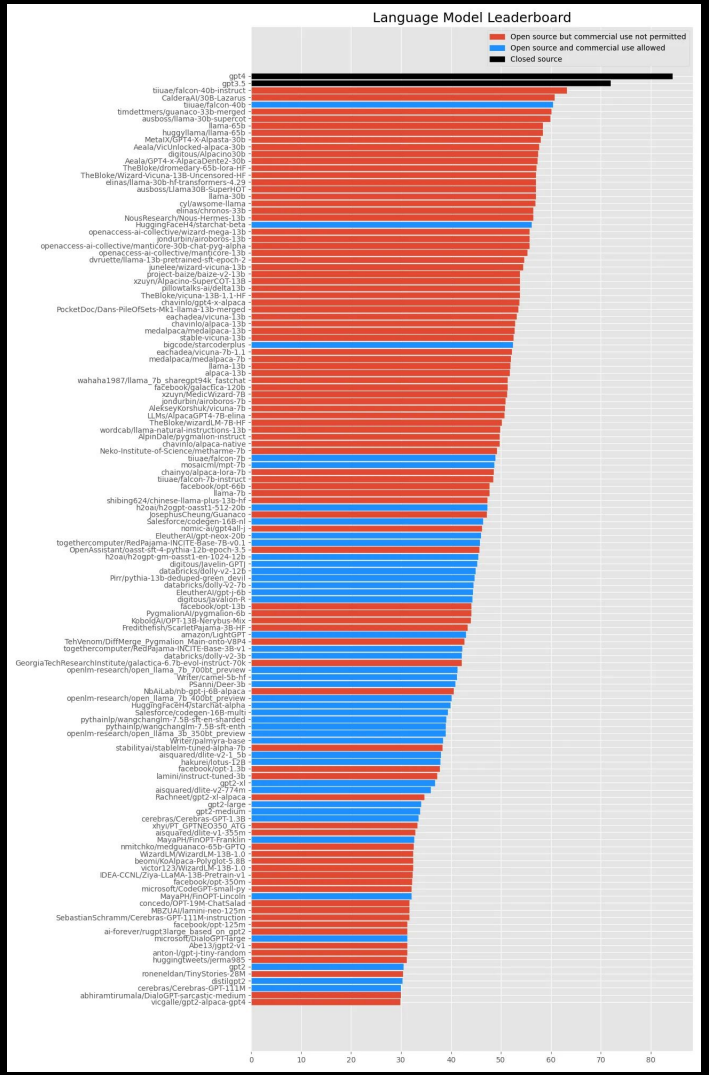

https://www.saattrupdan.com/posts/2023-04-16-open-source-chatgpt-alternatives

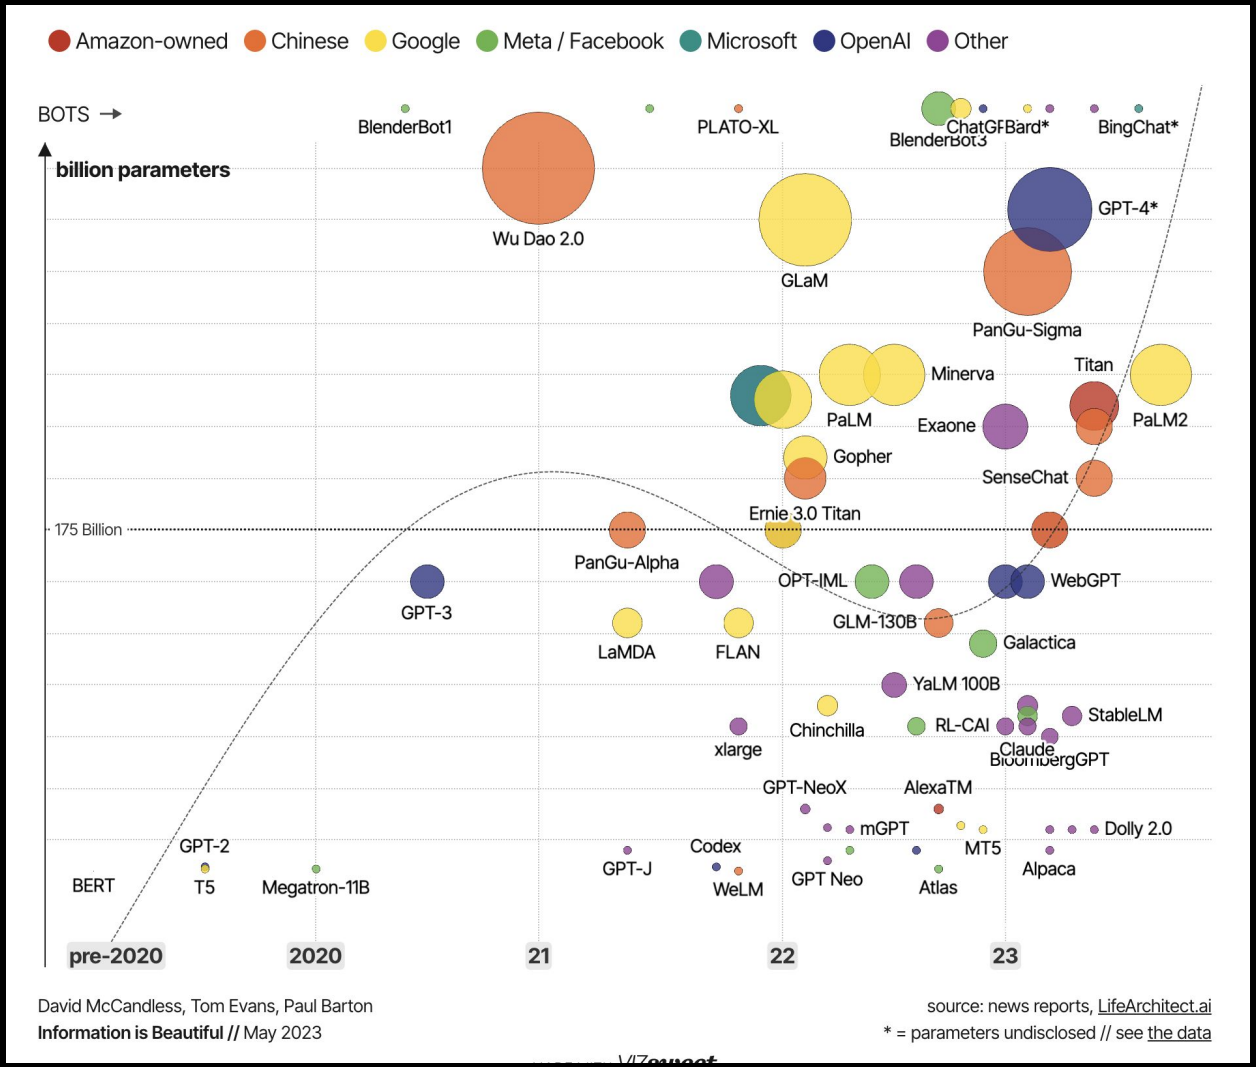

David McCandless, Tom Evans, Paul Barton
Information is Beautiful // May 2023

source: news reports, LifeArchitect.ai
* = parameters undisclosed // see the data

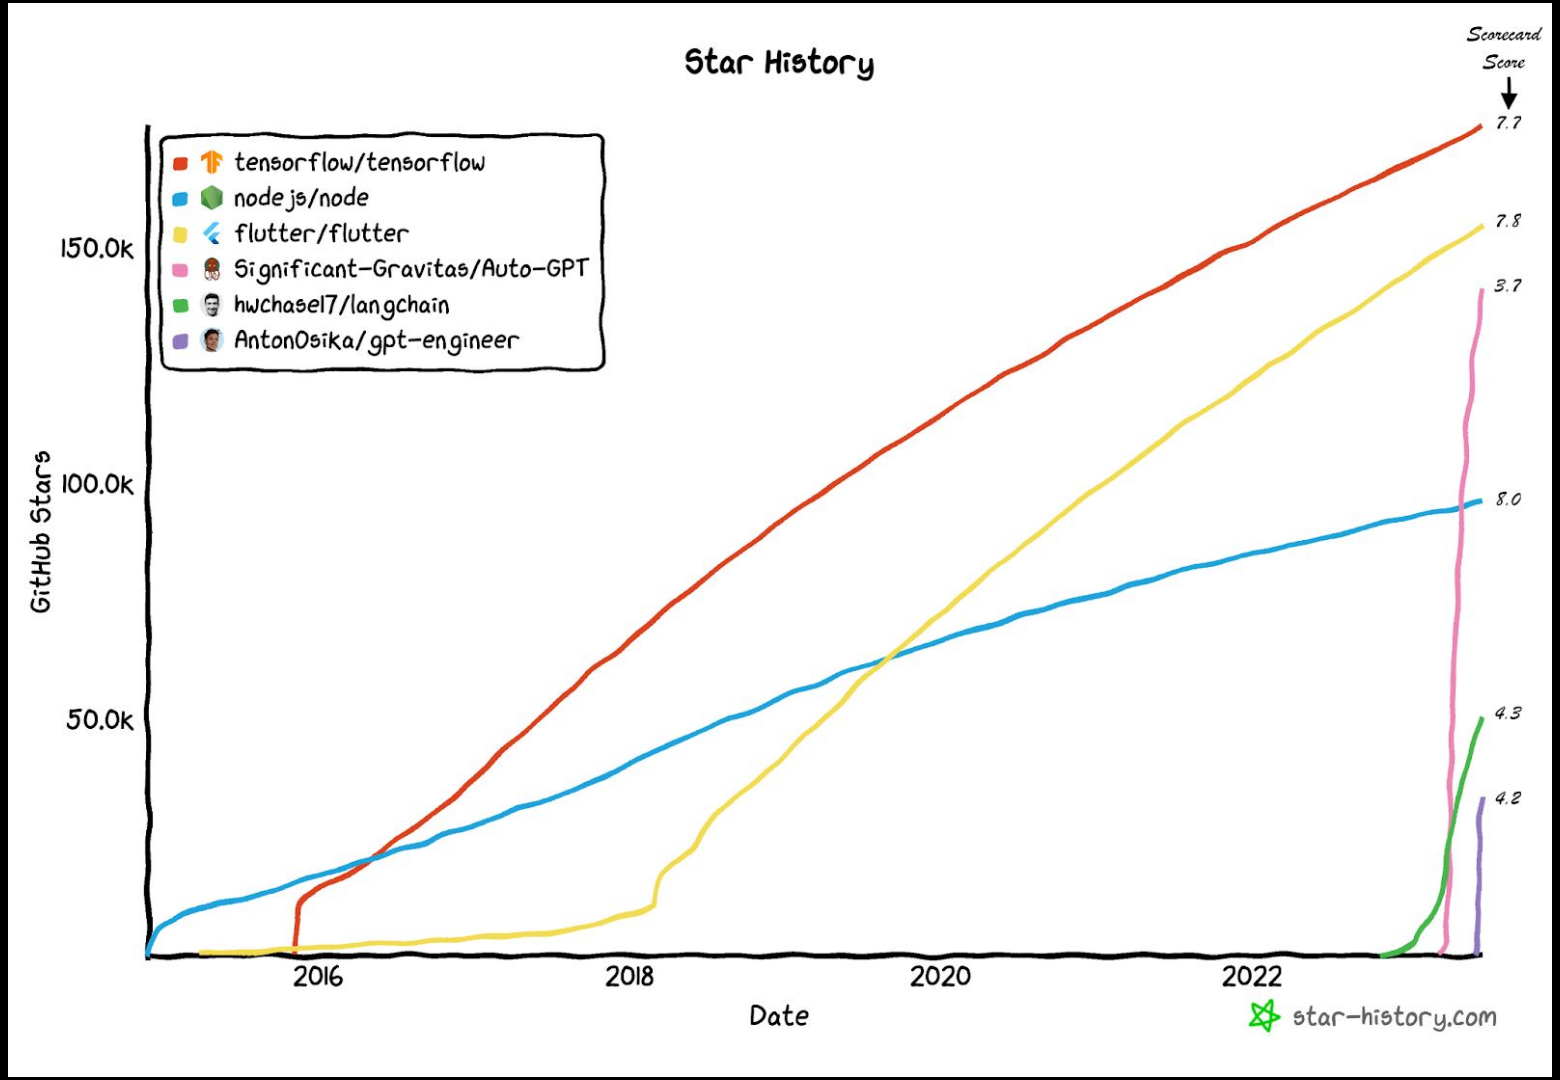
Star History
Scorecard Score
tensorflow/tensorflow
nodejs/node
flutter/flutter
Significant-Gravitas/Auto-GPT
hwchase17/langchain
AntonOsika/gpt-engineer
7.7
7.8
3.7
8.0
4.3
4.2
150.0k
100.0k
50.0k
GitHub Stars
2016
2018
2020
2022
Date
star-history.com

Challenges

Bad Actors Use AI to Rapidly Scale

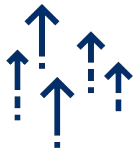

Speed

Rapid acceleration in open source LLM availability (>100K)

Scale

LLM availability lowers technical barriers, making it easier to scale attacks

Disruption

Attacks on LLMs will increase, amplifying the disruption

AI Advantages

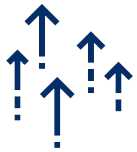

Speed

Only AI can detect in real-time.

Scope

Only AI can address the volume and variety.

Relevance

Only AI delivers & prioritizes at scale.

Which LLM?

https://towardsdatascience.com/choosing-the-right-language-model-for-your-nlp-use-case-1288ef3c4929

Model	Core differentiator	Pre-training objective	Para-meters	Access	Information Extraction	Text Classification	Conversa-tional AI	Summari-zation	Machine Translation	Content generation
BERT	First transformer-based LLM	AE	370M	Source code						
RoBERTa	More robust training procedure	AE	354M	Source code						
GPT-3	Parameter size	AR	175B	API						
BART	Novel combination of pre-training objectives	AR and AE	147M	Source code						
GPT-2	Parameter size	AR	1.5B	Source code						
T5	Multi-task transfer learning	AR	11B	Source code						
LaMDA	Dialogue; safety and factual grounding	AR	137B	No access						
XLNet	Joint AE and AR	AE and AR	110M	Source code						
DistilBERT	Reduced model size via knowledge distillation	AE	82M	Source code						
ELECTRA	Computational efficiency	AE	335M	Source code						
PaLM	Training infrastructure	AR	540B	No access						
MT-NLG	Training infrastructure	AR and AE	530B	API						
UniLM	Optimised both for NLU and NLG	Seq2seq, AE and AR	340M	Source code						
BLOOM	Multilingual (46 languages)	AR	176B	Source code						

AR = Autoregression

AE = Autoencoding

Seq2seq = Sequence-to-sequence

Highly appropriate

Appropriate

Somewhat appropriate

Options on Leveraging AI

Applications

- Consumer-facing
- Many free/low cost options (Bard, Bing Chat, ChatGPT, Midjourney, etc.)

Public APIs

- Used to build applications
- Low cost (for low volume)

Open Source Models

- Foundation models

Customized Models

- Custom built

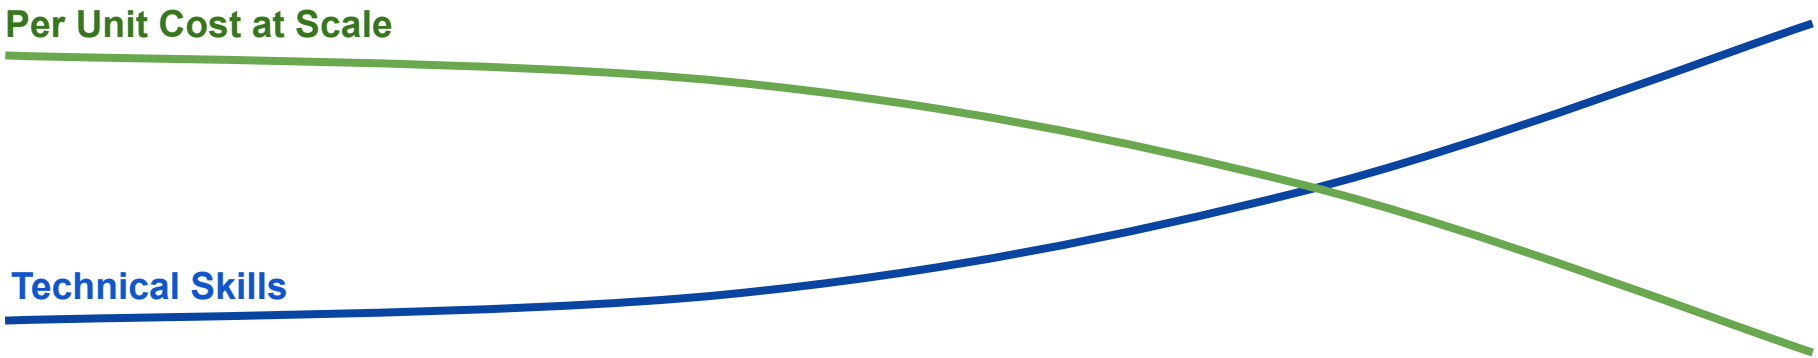

LLM Challenges

- Cost
- Performance (latency, accuracy)
- Evaluation
- Prompt “art” (not really “Engineering”)
- Stability
- Guardrails
- Hallucinations
- Benchmarks

Can’t rely “just” on LLMs (just yet)

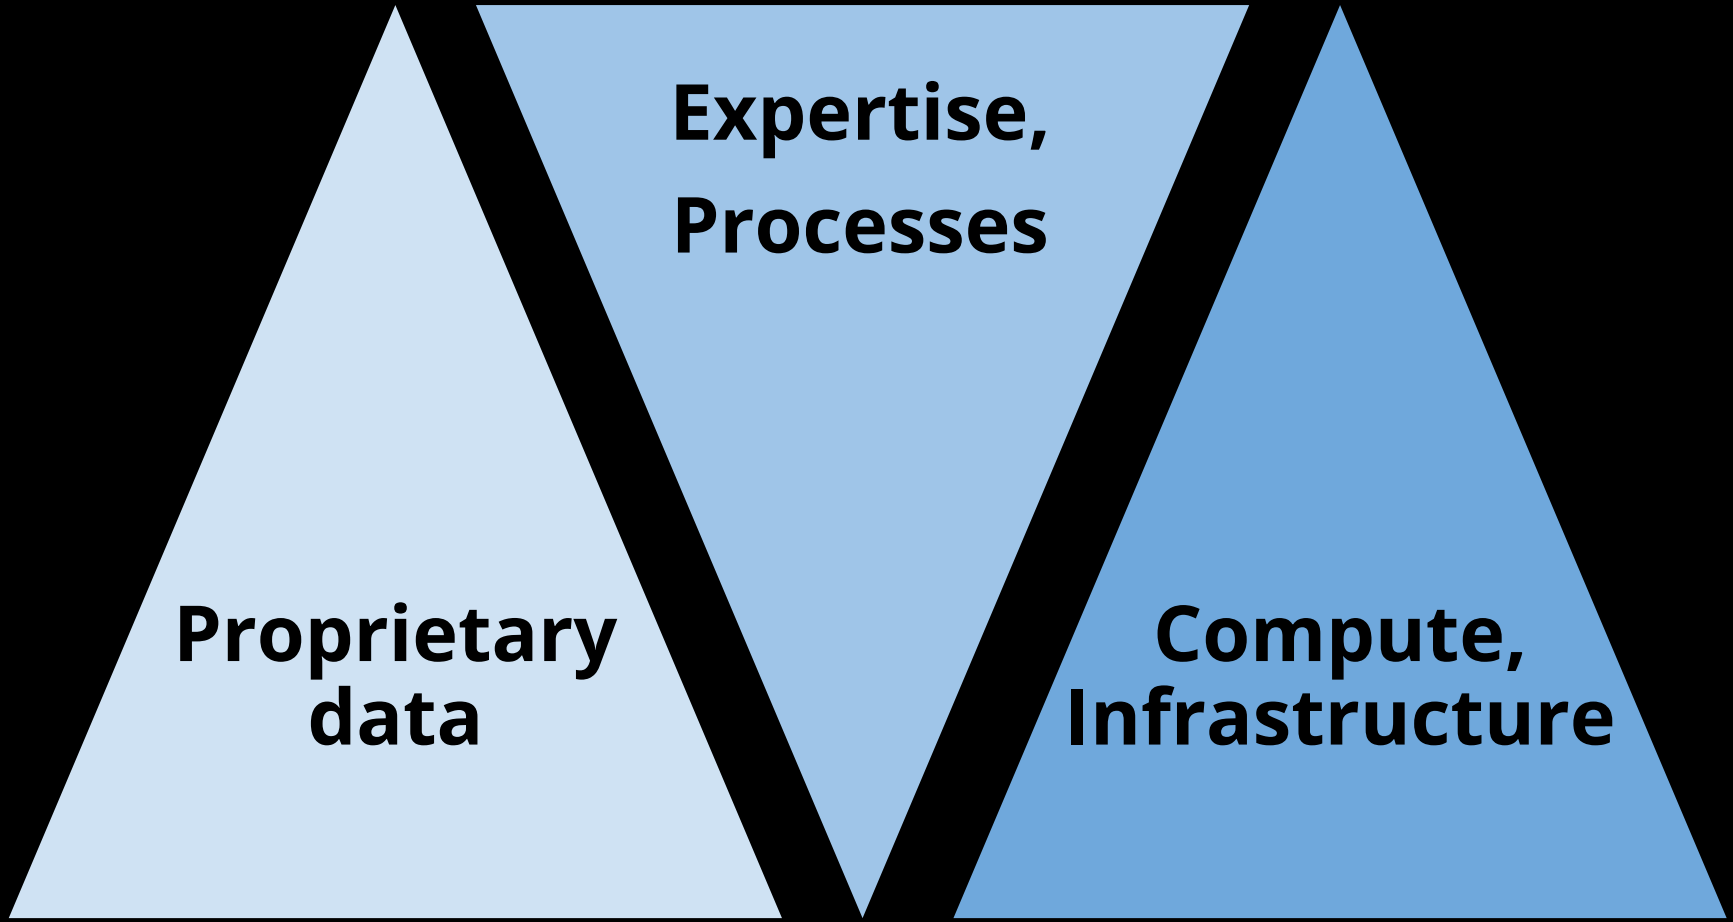
Expertise,
Processes
Proprietary
data
Compute,
Infrastructure

What's Needed?

Data

- For training and/or evaluation

Engineering Skills

- Integrate into workflows/applications

Problem Definitions

- Clearly define metrics and expected outcomes

Processes

- Achieve desired goals

Responsible Application of AI

Publicly Available Information

Applications	Public APIs	Data Repositories	Custom
• Consumer-facing/B2B • Different pricing models	• Often subsets of public data	• Academic and commercial	• Collected internally from public sources

Cost of LLM APIs

As a thought exercise, in 2021, DoorDash ML models made 10 billion predictions a day. If each prediction costs $0.004, that'd be $40 million a day!

https://huyenchip.com/2023/04/11/llm-engineering.html

Application Areas
(VC Perspective)

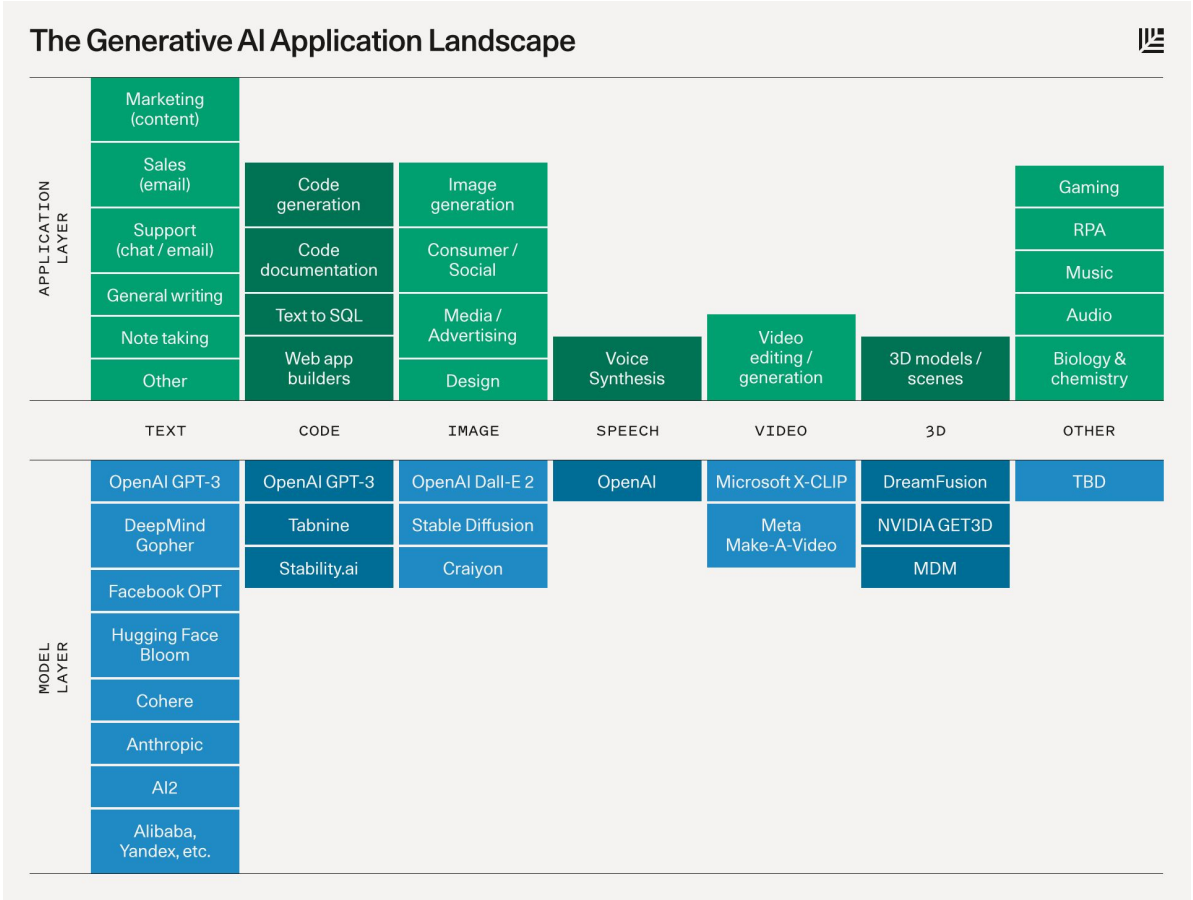

https://www.sequoiacap.com/article/generative-ai-a-creative-new-world/

Where We Are
(VC Perspective)

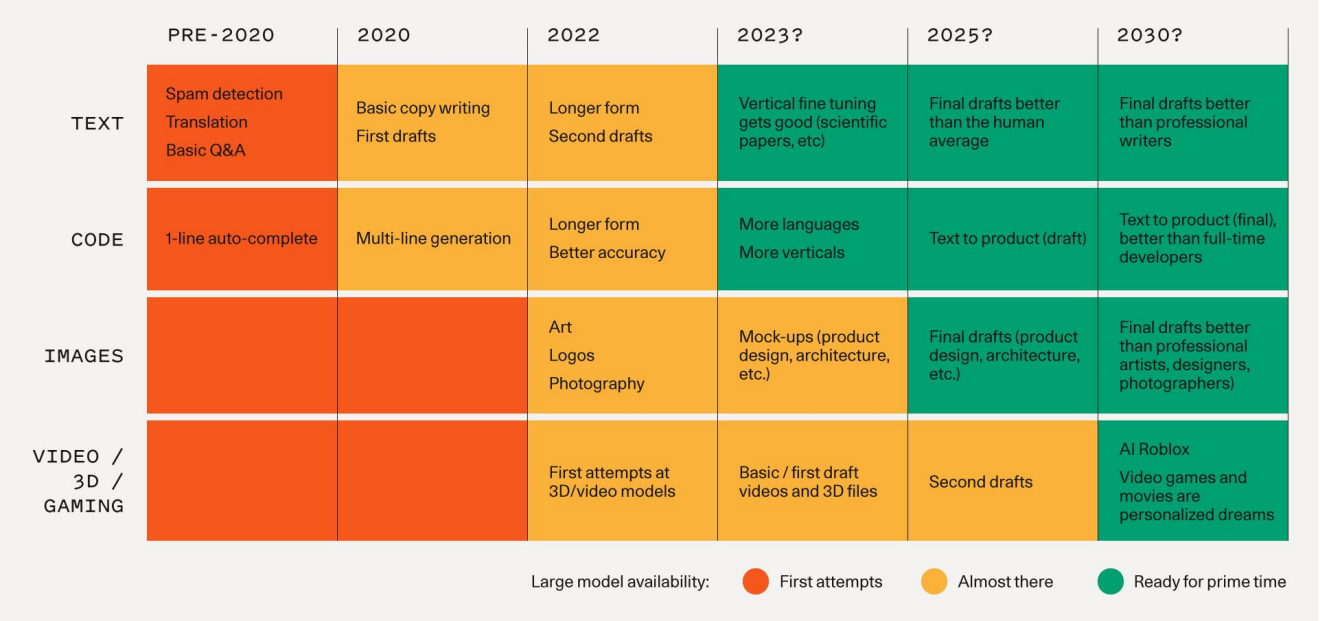

https://www.sequoiacap.com/article/generative-ai-a-creative-new-world/

Speeding up at Scale

- Inference matters more than training
- Shifts in distribution of input and output data over time put pressure on sampling methods and fast (continuous) training (and labeling).
- Every millisecond counts (saves lives)
- Optimization everywhere (including cost)

Training & Inference

- Hardware
- Software: model compression, pruning
- Processes: stacking, ensembles ++
- Complex and simple models

Research Directions

Three Phases

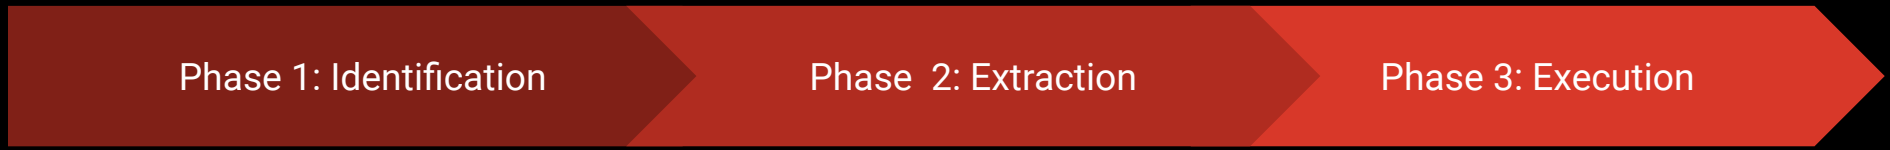

Three Phases

Phase 1: Identification

- Is the data relevant?
- Informative?
- What is the intent of the post?

Phase 2: Extraction

- Who, what, where?

Phase 3: Execution

- Allocate resources quickly and effectively

Popular Tasks

- Informativeness
- Aid Detection
- Humanitarian Response
- Eyewitness Account Detection

Predictive Challenges

- Context (sometimes!)
- Reasoning (connecting some dots)
- “Extraction” of basic semantic information is STILL VERY HARD (disambiguating, …)

Summarization Challenges

- Many (thousands, millions) of events
- Clustering
- Causality- critical
- Granularity (how much detail is enough? too much?)
- What to summarize?

Multimodal (our work at CVPR ‘20)

- Informativeness, Event, Severity
- Outperforms
  - Text-only and image-only approaches & strong multimodal baselines

Cross-Attention Module

- A fusion method for integrating complementary information from noisy modalities.
- Prevents negative transfer of information between modalities on a case by case basis.
- Cross-attention layers combined with a self-attention layer.

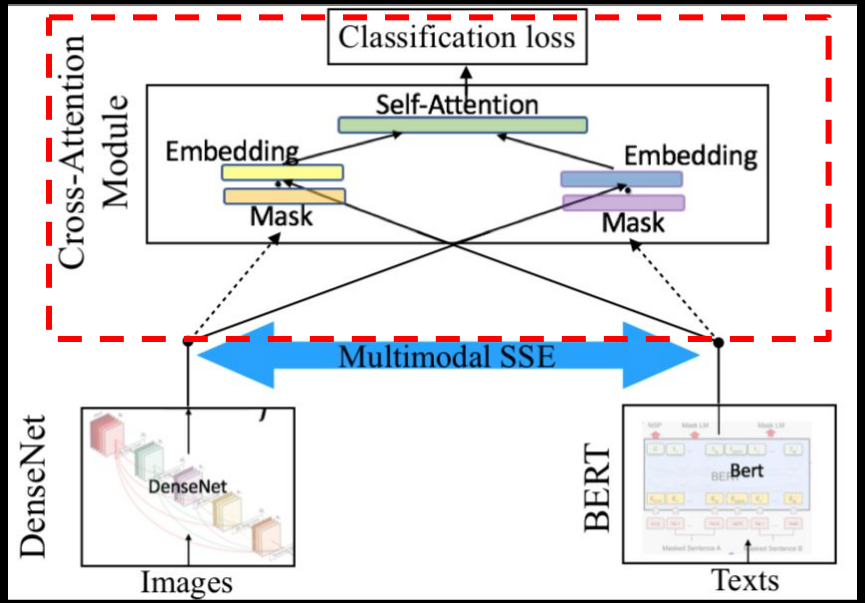

Cascading Networks

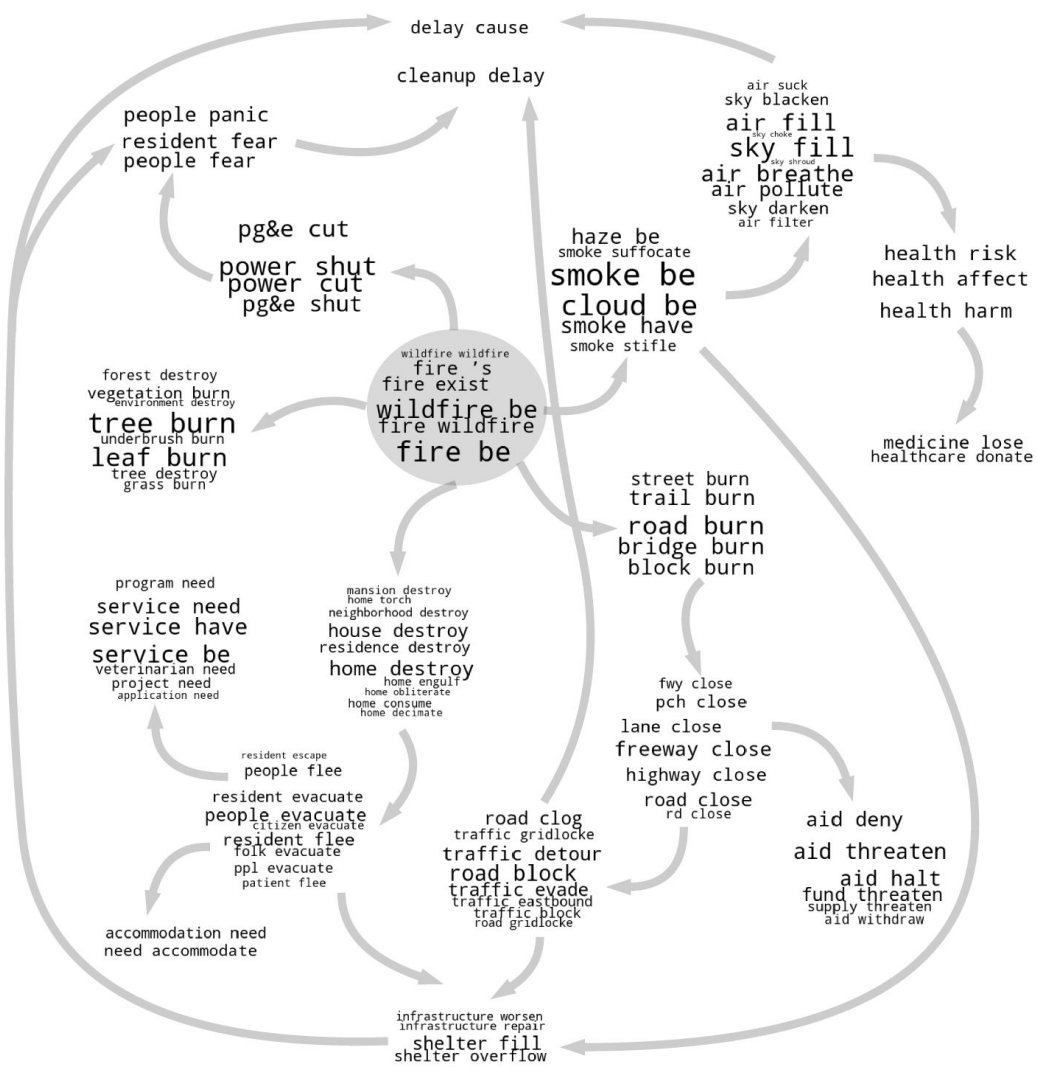

Dataminr®

Cascading Networks

Our work
at
ICML AI for Social Good ‘19

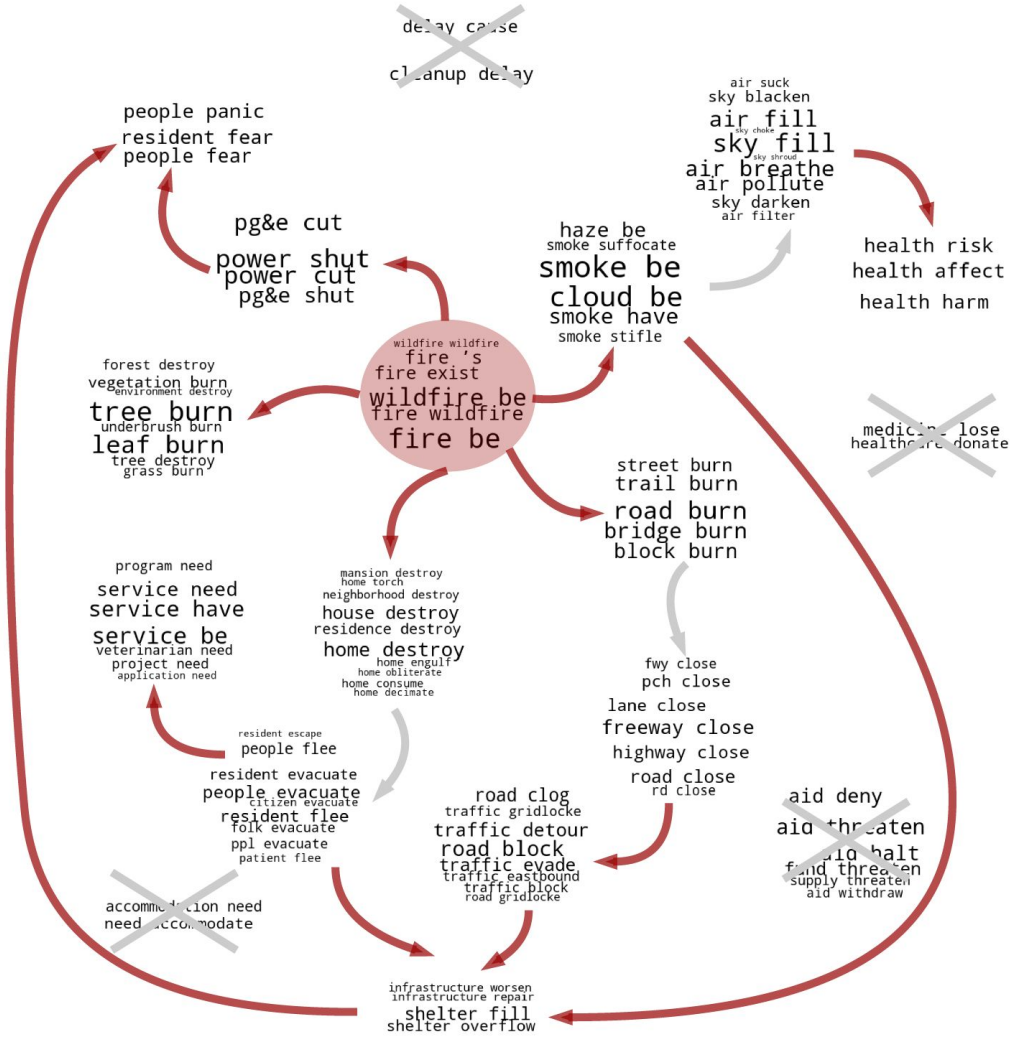

Dataminr

LLMs

CLIP

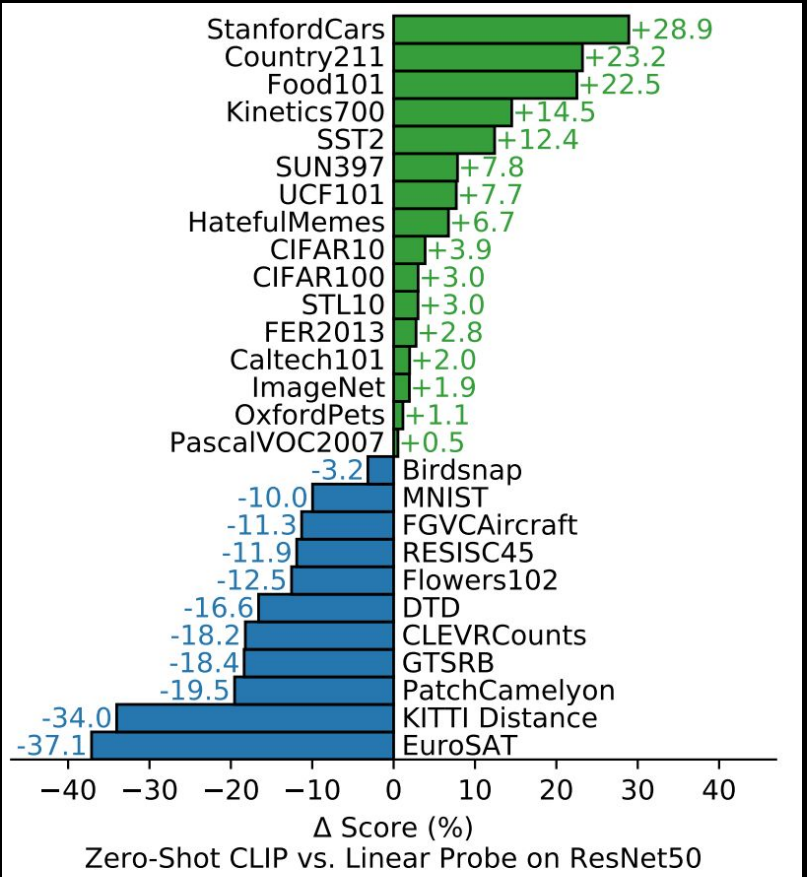

Figure 5. **Zero-shot CLIP is competitive with a fully supervised baseline.** Across a 27 dataset eval suite, a zero-shot CLIP classifier outperforms a fully supervised linear classifier fitted on ResNet-50 features on 16 datasets, including ImageNet.

Flamingo

Which city is this?

This is a street sign in New York City.

Which street corner is it?

It is the intersection of Madison Avenue and East 42nd Street.

Which subway station is nearby?

The closest subway station is Grand Central Station.

How about the nearest park?

The nearest park is Bryant Park.

There is another sign. What does it say?

It says "One way".

https://huyenchip.com/2023/10/10/multimodal.html

Summary

Research Topics

- Duplicate detection
- Integrating information from multiple sources in real-time
- Real-time corroboration
- Detecting sub-events in real-time
- Disinformation
- DeepFake & AI generated content detection
- Optimization of LLMs
- KGs
- Hallucination mitigation
- Attribution

State of the Art vs. Real World

- Data challenges
  - Bias, sampling, unseen data, metrics, labeling
- Computational challenges
  - Cost/performance/speed tradeoff *at scale*
- Knowledge challenges
  - Different stakeholders at different levels. What's relevant, to whom?
  - DIfferent ontologies and Knowledge *Graphs*
- Situational awareness
  - Context, context, context (who, how, etc.)
  - Human behavior

Value of Real Time AI

Emergency Response

- Humanitarian assistance

Cyber Risk

- Vulnerabilities & digital risk

Corporate Risk

- Physical asset protection

News

- Breaking news

Thank you!

Alex Jaimes
ajaimes@dataminr.com

https://www.linkedin.com/in/alexjaimes/

dataminr.com

END