

State of AI Report

October 1, 2020

About the authors



Nathan Benaich

Nathan is the General Partner of **Air Street Capital**, a venture capital firm investing in AI-first technology and life science companies. He founded RAAIS and London.AI, which connect AI practitioners from large companies, startups and academia, and the RAAIS Foundation that funds open-source AI projects. He studied biology at Williams College and earned a PhD from Cambridge in cancer research.



Ian Hogarth

Ian is an **angel investor** in 60+ startups. He is a Visiting Professor at UCL working with Professor Mariana Mazzucato. Ian was co-founder and CEO of Songkick, the concert service used by 17M music fans each month. He studied engineering at Cambridge where his Masters project was a computer vision system to classify breast cancer biopsy images. He is the Chair of Phasecraft, a quantum software company.

Artificial intelligence (AI) is a multidisciplinary field of science and engineering whose goal is to create intelligent machines.

We believe that AI will be a force multiplier on technological progress in our increasingly digital, data-driven world. This is because everything around us today, ranging from culture to consumer products, is a product of intelligence.


The State of AI Report is now in its third year. New to the 2020 edition are several invited content contributions from a range of well-known and up-and-coming companies and research groups. Consider this Report as a compilation of the most interesting things we've seen with a goal of triggering an informed conversation about the state of AI and its implication for the future.

We consider the following key dimensions in our report:

- **Research:** Technology breakthroughs and their capabilities.
- **Talent:** Supply, demand and concentration of talent working in the field.
- **Industry:** Areas of commercial application for AI and its business impact.
- **Politics:** Regulation of AI, its economic implications and the emerging geopolitics of AI.
- **Predictions:** What we believe will happen in the next 12 months and a 2019 performance review to keep us honest.

Collaboratively produced by **Ian Hogarth** (@soundboy) and **Nathan Benaich** (@nathanbenaich).

Thank you to our contributors



Thank you to our reviewers

Jack Clark, Jeff Ding, Chip Huyen, Rebecca Kagan, Andrej Karpathy, Moritz Müller-Freitag, Torsten Reil, Charlotte Stix, and Nu (Claire) Wang.

Definitions

Artificial intelligence (AI): A broad discipline with the goal of creating intelligent machines, as opposed to the natural intelligence that is demonstrated by humans and animals. It has become a somewhat catch all term that nonetheless captures the long term ambition of the field to build machines that emulate and then exceed the full range of human cognition.

Machine learning (ML): A subset of AI that often uses statistical techniques to give machines the ability to "learn" from data without being explicitly given the instructions for how to do so. This process is known as "training" a "model" using a learning "algorithm" that progressively improves model performance on a specific task.

Reinforcement learning (RL): An area of ML concerned with developing software agents that learn goal-oriented behavior by trial and error in an environment that provides rewards or penalties in response to the agent's actions (called a "policy") towards achieving that goal.

Deep learning (DL): An area of ML that attempts to mimic the activity in layers of neurons in the brain to learn how to recognise complex patterns in data. The "deep" in deep learning refers to the large number of layers of neurons in contemporary ML models that help to learn rich representations of data to achieve better performance gains.

Definitions

Algorithm: An unambiguous specification of how to solve a particular problem.

Model: Once a ML algorithm has been trained on data, the output of the process is known as the model. This can then be used to make predictions.

Supervised learning: A model attempts to learn to transform one kind of data into another kind of data using labelled examples. This is the most common kind of ML algorithm today.

Unsupervised learning: A model attempts to learn a dataset's structure, often seeking to identify latent groupings in the data without any explicit labels. The output of unsupervised learning often makes for good inputs to a supervised learning algorithm at a later point.

Transfer learning: An approach to modelling that uses knowledge gained in one problem to bootstrap a different or related problem, thereby reducing the need for significant additional training data and compute.

Natural language processing (NLP): Enabling machines to analyse, understand and manipulate language.

Computer vision: Enabling machines to analyse, understand and manipulate images and video.

Executive Summary

Research

- A new generation of transformer language models are unlocking new NLP use-cases.
- Huge models, large companies and massive training costs dominate the hottest area of AI today: Natural Language Processing.
- Biology is experiencing its “AI moment”: From medical imaging, genetics, proteomics, chemistry to drug discovery.
- AI is mostly closed source: Only 15% of papers publish their code, which harms accountability and reproducibility in AI.

Talent

- American institutions and corporations further their dominance of major academic conference papers acceptances.
- Multiple new institutions of higher education dedicated to AI are formed.
- Corporate-driven academic brain drain is significant and appears to negatively impact entrepreneurship.
- US AI ecosystem is fuelled by foreign talent and the contribution of researchers educated in China to world-class papers is clear.

Industry

- The first trial of an AI-discovered drug begins in Japan and the first US medical reimbursement for AI-based imaging procedure is granted.
- Self-driving car mileage remains microscopic and open sourcing of data grows to crowdsource new solutions.
- Google, Graphcore, and NVIDIA continue to make major advances in their AI hardware platforms.
- NLP applications in industry continue to expand their footprint and are implemented in Google Search and Microsoft Bing.

Politics

- After two wrongful arrests involving facial recognition, ethical risks that researchers have been warning about come into sharp focus.
- Semiconductor companies continue to grow in geopolitical significance, particularly Taiwan’s TSMC.
- The US Military is absorbing AI progress from academia and industry labs.
- Nations pass laws to let them scrutinize foreign takeovers of AI companies and the UK’s Arm will be a key test.


Scorecard: Reviewing our predictions from 2019

Our 2019 Prediction	Grade	Evidence
New natural language processing companies raise \$100M in 12 months.	Yes	Gong.io (\$200M), Chorus.ai (\$45M), Ironscales (\$23M), ComplyAdvantage (\$50M), Rasa (\$26M), HyperScience (\$60M), ASAPP (\$185M), Cresta (\$21M), Eigen (\$37M), K Health (\$48M), Signal (\$25M), and many more!
No autonomous driving company drives >15M miles in 2019.	Yes	Waymo (1.45M miles), Cruise (831k miles), Baidu (108k miles).
Privacy-preserving ML adopted by a F2000 company other than GAFAM (Google, Apple, Facebook, Amazon, Microsoft).	Yes	Machine learning ledger orchestration for drug discovery (MELLODY) research consortium with large pharmaceutical companies and startups including Glaxosmithkline, Merck and Novartis.
Unis build <i>de novo</i> undergrad AI degrees.	Yes	CMU graduates first cohort of AI undergrads, Singapore's SUTD launches undergrad degree in design and AI, NYU launches data science major, Abu Dhabi builds an AI university.
Google has major quantum breakthrough and 5 new startups focused on quantum ML are formed.	Sort of	Google demonstrated quantum supremacy in October 2019! Many new quantum startups were launched in 2019 but only Cambridge Quantum, Rahko, Xanadu.ai, and QCWare are explicitly working on quantum ML.
Governance of AI becomes key issue and one major AI company makes substantial governance model change.	No	Nope, business as usual.

Section 1: Research

AI research is less open than you think: Only 15% of papers publish their code

- Research paper code implementations are important for accountability, reproducibility and driving progress in AI. The field has made little improvement on this metric since mid-2016. Traditionally, academic groups are more likely to publish their code than industry groups. Notable organisation that don't publish all of their code are OpenAI and DeepMind. For the biggest tech companies, their code is usually intertwined with proprietary scaling infrastructure that cannot be released. This points to centralization of AI talent and compute as a huge problem.



Papers With Code tracks openly-published code and benchmarks model performance

▶ Hosting 3,000 State-of-the-Art leaderboards, 750+ ML components, and 25,000+ research along with code.

Browse State-of-the-Art

2,856 leaderboards • 1,621 tasks • 2,539 datasets • 25,673 papers with code

[Follow on Twitter for updates](#)

Computer Vision

**Semantic Segmentation**
57 leaderboards
1007 papers with code

**Image Classification**
145 leaderboards
868 papers with code

**Object Detection**
86 leaderboards
723 papers with code

**Image Generation**
105 leaderboards
356 papers with code

**Pose Estimation**
76 leaderboards
332 papers with code

[See all 809 tasks](#)

Natural Language Processing

**Machine Translation**
45 leaderboards
648 papers with code

**Language Modelling**
14 leaderboards
630 papers with code

**Question Answering**
56 leaderboards
590 papers with code

**Sentiment Analysis**
37 leaderboards
404 papers with code

**Text Classification**
60 leaderboards
248 papers with code

[See all 304 tasks](#)

Methods

753 machine learning components

A free and open resource that the community can edit and [contribute to](#)

General

**Regularization**
35 methods
7949 papers with code

**Activation Functions**
29 methods
11534 papers with code

**Stochastic Optimization**
29 methods
3659 papers with code

**Attention**
28 methods
4500 papers with code

**Normalization**
27 methods
4864 papers with code

[See all 47 categories](#)

Computer Vision

**Convolutional Neural Networks**
82 methods
2358 papers with code

**Image Model Blocks**
67 methods
2805 papers with code

**Generative Models**
48 methods
3635 papers with code


**Object Detection Models**
34 methods
800 papers with code

**Image Feature Extractors**
27 methods
47 papers with code

[See all 38 categories](#)


Facebook's PyTorch is fast outpacing Google's TensorFlow in research papers, which tends to be a leading indicator of production use down the line

- ▶ Of 20-35% of conference papers that mention the framework they use, 75% cite the use of PyTorch but not TensorFlow. Of 161 authors who published more TensorFlow papers than PyTorch papers in 2018, 55% of them have switched to PyTorch. The opposite happened in 15% of cases. Meanwhile, we observe that TensorFlow, Caffe and Caffe2 are still the workhorse for production AI.




PyTorch is also more popular than TensorFlow in paper implementations on GitHub

- ▶ 47% of these implementations are based on PyTorch vs. 18% for TensorFlow. PyTorch offers greater flexibility and a dynamic computational graph that makes experimentation easier. JAX is a Google framework that is more math friendly and favored for work outside of convolutional models and transformers.



Language models: Welcome to the Billion Parameter club




- Huge models, large companies and massive training costs dominate the hottest area of AI today, NLP.



Note: The number of parameters indicates how many different coefficients the algorithm optimizes during the training process.

Bigger models, datasets and compute budgets clearly drive performance

- ▶ Empirical scaling laws of neural language models show smooth power-law relationships, which means that as model performance increases, the model size and amount of computation has to increase more rapidly.



Tuning billions of model parameters costs millions of dollars

- Based on variables released by Google et al., you're paying circa \$1 per 1,000 parameters. This means OpenAI's 175B parameter GPT-3 could have cost tens of millions to train. Experts suggest the likely budget was \$10M.

Just how much does it cost to train a model? Two correct answers are “depends” and “a lot”. More quantitatively, here are current ballpark list-price costs of training differently sized BERT [4] models on the Wikipedia and Book corpora (15 GB). For each setting we report two numbers - the cost of one training run, and a typical fully-loaded cost (see discussion of “hidden costs” below) with hyper-parameter tuning and multiple runs per setting (here we look at a somewhat modest upper bound of two configurations and ten runs per configuration).⁴


- \$2.5k - \$50k (110 million parameter model)
- \$10k - \$200k (340 million parameter model)
- \$80k - \$1.6m (1.5 billion parameter model)

For example, based on information released by Google, we estimate that, at list-price, training the 11B-parameter variant⁵ of T5 [5] cost well above \$1.3 million for a single run. Assuming 2-3 runs of the large model and hundreds of the small ones, the (list-)price tag for the entire project may have been \$10 million⁶.

Not many companies – certainly not many startups – can afford this cost. Some argue that this is not a severe issue; let the Googles of the world pre-train and publish the large language models, and let the rest of the world fine-tune them (a much cheaper endeavor) to specific tasks. Others (e.g., Etchemendy and Li [6]) are not as sanguine.

To achieve the needed quality improvements in machine translation, Google's final model trained for the equivalent of 22 TPU v3 core years or ~5 days with 2,048 cores non-stop

► This transformer-based model with conditional computation for machine translation has 600B parameters.



We're rapidly approaching outrageous computational, economic, and environmental costs to gain incrementally smaller improvements in model performance


- ▶ Without major new research breakthroughs, dropping the ImageNet error rate from 11.5% to 1% would require over one hundred billion billion dollars! Many practitioners feel that progress in mature areas of ML is stagnant.

Benchmark	Error rate	Polynomial			Exponential		
		Computation Required (Gflops)	Environmental Cost (CO_2)	Economic Cost (\$)	Computation Required (Gflops)	Environmental Cost (CO_2)	Economic Cost (\$)
ImageNet	Today: 11.5%	10^{14}	10^6	10^6	10^{14}	10^6	10^6
	Target 1: 5%	10^{19}	10^{10}	10^{11}	10^{27}	10^{19}	10^{19}
	Target 2: 1%	10^{28}	10^{20}	10^{20}	10^{120}	10^{112}	10^{112}
MS COCO	Today: 46.7%	10^{14}	10^6	10^6	10^{15}	10^7	10^7
	Target 1: 30%	10^{23}	10^{14}	10^{15}	10^{29}	10^{21}	10^{21}
	Target 2: 10%	10^{44}	10^{36}	10^{36}	10^{107}	10^{99}	10^{99}
SQuAD 1.1	Today: 4.621%	10^{13}	10^4	10^5	10^{13}	10^5	10^5
	Target 1: 2%	10^{15}	10^7	10^7	10^{23}	10^{15}	10^{15}
	Target 2: 1%	10^{18}	10^{10}	10^{10}	10^{40}	10^{32}	10^{32}
CoLLN 2003	Today: 6.5%	10^{13}	10^5	10^5	10^{13}	10^5	10^5
	Target 1: 2%	10^{43}	10^{35}	10^{35}	10^{82}	10^{73}	10^{74}
	Target 2: 1%	10^{61}	10^{53}	10^{53}	10^{181}	10^{173}	10^{173}
WMT 2014 (EN-FR)	Today: 54.4%	10^{12}	10^4	10^4	10^{12}	10^4	10^4
	Target 1: 30%	10^{23}	10^{15}	10^{15}	10^{30}	10^{22}	10^{22}
	Target 2: 10%	10^{43}	10^{35}	10^{35}	10^{107}	10^{99}	10^{100}


A larger model needs less data than a smaller peer to achieve the same performance

- This has implications for problems where training data samples are expensive to generate, which likely confers an advantage to large companies entering new domains with supervised learning-based models.

Larger models require **fewer samples** to reach the same performance




The optimal model size grows smoothly with the loss target and compute budget



Low resource languages with limited training data are a beneficiary of large models


- Google made use of their large language models to deliver higher quality translations for languages with limited amounts of training data, for example Hansa and Uzbek. This highlights the benefits of transfer learning.




Even as deep learning consumes more data, it continues to get more efficient

- ▶ Since 2012 the amount of compute needed to train a neural network to the same performance on ImageNet classification has been decreasing by a factor of 2 every 16 months.

Two distinct eras of compute in training AI systems





Training efficiency factor



Yet, for some use cases like dialogue small, data-efficient models can trump large models

- PolyAI, a London-based conversational AI company, open-sourced their ConveRT model (a pre-trained contextual re-ranker based on transformers). Their model outperforms Google's BERT model in conversational applications, especially in low data regimes, suggesting BERT is far from a silver bullet for all NLP tasks.

Model	1-vs-100 Accuracy	Model Size
ELMo	20.6%	372M
BERT	24.0%	1.3G
USE	47.7%	845M
ConveRT (PolyAI)	68.2%	59M



A new generation of transformer language models are unlocking new NLP use-cases

▶ GPT-3, T5, BART are driving a drastic improvement in the performance of transformer models for text-to-text tasks like translation, summarization, text generation, text to code.

▶ Summarization from huggingface.co/models

Model: `facebook/bart-large-cnn`

pytorch rust bart summarization license:mit pipeline:summarization

Hosted inference API

summarization

The tower is 324 metres (1,063 ft) tall, about the same height as an 81-storey building, and the tallest structure in Paris. Its base is square, measuring 125 metres (410 ft) on each side. During its construction, the Eiffel Tower surpassed the Washington Monument to become the tallest man-made structure in the world; a title it held

Compute

Computation time on cpu: 7.403 s

The tower is 324 metres (1,063 ft) tall, about the same height as an 81-storey building. Its base is square, measuring 125 metres (410 ft) on each side. It is the second tallest free-standing structure in France after the Millau Viaduct.

▶ JSON Output ▶ ↔ API endpoint

▶ Share

▶ Code generation and more: gpt3examples.com



Sharif Shameem
@sharifshameem

Here's a sentence describing what Google's home page should look and here's GPT-3 generating the code for it nearly perfectly.

Describe a layout.

2 lightgrey buttons that say "Search Google" and "I'm Feeling Lucky" with padding in between them

```
// the google logo
 <button style-
{color: 'white', backgroundColor:
'lightgrey'}}>Search Google</button> <button
style={color: 'white', backgroundColor:
'lightgrey'}}>I'm Feeling Lucky</button>
```

0:13 395.6K views



4:50 AM · Jul 15, 2020 · Twitter Web App

3.4K Retweets and comments 12K Likes



Computer, please convert my code into another programming language

- ▶ An unsupervised machine translation model trained on GitHub projects with 1,000 parallel functions can translate 90% of these functions from C++ to Java and 57% of Python functions into C++ and successfully pass unit tests. No expert knowledge required, but no guarantees that the model didn't memorize the functions either.


Python input	Unsupervised C++ translation
<pre>def SumOfKsubArray(arr, n, k): Sum = 0 S = deque() G = deque() for i in range(k): while (len(S) > 0 and arr[S[-1]] >= arr[i]): S.pop() while (len(G) > 0 and arr[G[-1]] <= arr[i]): G.pop() G.append(i) S.append(i) for i in range(k, n): Sum += arr[S[0]] + arr[G[0]] while (len(S) > 0 and S[0] <= i - k): S.popleft() while (len(G) > 0 and G[0] <= i - k): G.popleft() while (len(S) > 0 and arr[S[-1]] >= arr[i]): S.pop() while (len(G) > 0 and arr[G[-1]] <= arr[i]): G.pop() G.append(i) S.append(i) Sum += arr[S[0]] + arr[G[0]] return Sum</pre>	<pre>int SumOfKsubArray(int arr[], int n, int k){ int Sum = 0; deque<int> S; deque<int> G; for(int i = 0; i < k; i++){ while((int) S.size() > 0 && arr[S.back()] >= arr[i]) S.pop_back(); while((int) G.size() > 0 && arr[G.back()] <= arr[i]) G.pop_back(); G.push_back(i); S.push_back(i); } for(int i = k; i < n; i++){ Sum += arr[S.front()] + arr[G.front()]; while((int) S.size() > 0 && S.front() <= i - k) S.pop_front(); while((int) G.size() > 0 && G.front() <= i - k) G.pop_front(); while((int) S.size() > 0 && arr[S.back()] >= arr[i]) S.pop_back(); while((int) G.size() > 0 && arr[G.back()] <= arr[i]) G.pop_back(); G.push_back(i); S.push_back(i); } Sum += arr[S.front()] + arr[G.front()]; return Sum; }</pre>

Figure 2: **Example of unsupervised Python to C++ translation.** TransCoder successfully translates the Python input function SumOfKsubArray into C++. TransCoder infers the types of the arguments, of the variables, and the return type of the function. The model maps the Python deque() container, to the C++ implementation deque<>, and uses the associated front, back, pop_back and push_back methods to retrieve and insert elements into the deque, instead of the Python square brackets [], pop and append methods. Moreover, it converts the Python for loop and range function properly.

Computer, can you automatically repair my buggy programs too?

▶ Given a broken program and diagnostic feedback (compiler error message), DrRepair localizes an erroneous line and generates a repaired line.


- The model jointly reasons over the broken source code and the diagnostic feedback using graph neural networks.
- They use self-supervised learning to obviate the need for labelling by taking code from programming competitions and corrupting it into a broken program.
- A SOTA is set on DeepFix, which is a program repair benchmark for correct intro programming assignments in C.




NLP benchmarks take a beating: Over a dozen teams outrank the human GLUE baseline

► It was only 12 months ago that the human GLUE benchmark was beat by 1 point. Now SuperGLUE is in sight.

- GLUE and its more challenging sibling SuperGLUE are benchmarks that evaluate NLP systems at a range of tasks spanning logic, common sense understanding, and lexical semantics. The human benchmark on GLUE is reliably beat today (right) and the SuperGLUE human benchmark is almost surpassed too!




SuperGLUE



What's next after SuperGLUE? More challenging NLP benchmarks zero-in on knowledge


▶ A multi-task language understanding challenge tests for world knowledge and problem solving ability across 57 tasks including maths, US history, law and more. GPT-3's performance is lopsided with large knowledge gaps.

- Large models like GPT-3 that are pre-trained on vast language corpora obviate the need for task-specific fine-tuning on a specific dataset. This enables few-shot learning on new tasks. A new benchmark measures knowledge acquired during pre-training by evaluating in few-shot settings (% avg. weighted accuracy below).
- While the GPT-3 X-Large improves over random chance by over 20 percentage points on average, the model's accuracy ranges from 69% for US Foreign Policy to 26% for College Chemistry. Moreover, GPT-3's average confidence is a poor estimator of its accuracy and can be off by up to 24%.




The transformer's ability to generalise is remarkable. It can be thought of as a new layer type that is more powerful than convolutions because it can process sets of inputs and fuse information more globally.

- ▶ For example, GPT-2 was trained on text but can be fed images in the form of a sequence of pixels to learn how to autocomplete images in an unsupervised manner.



Biology is experiencing its “AI moment”: Over 21,000 papers in 2020 alone

- Publications involving AI methods (e.g. deep learning, NLP, computer vision, RL) in biology are growing >50% year-on-year since 2017. Papers published since 2019 account for 25% of all output since 2000.




From physical object recognition to “cell painting”: Decoding biology through images


- ▶ Large labelled datasets offer huge potential for generating new biological knowledge about health and disease.

IMAGENET

>14M labeled images




RxRx.ai image datasets of cells treated with various chemical agents



Deep learning on cellular microscopy accelerates biological discovery with drug screens

► Embeddings from experimental data illuminate biological relationships and predict COVID-19 drug successes.


- Deep learning models trained to identify biologically-perturbed cells imaged by fluorescent microscopy can identify 100s-1000s of relevant features of cellular morphology.
- Applying these features makes it possible to relate the biology induced by genetic changes, immune/cytokine perturbations, and drugs.
- These models were applied to experiments on COVID-19 infection and cytokine storm, identifying repurposable candidates and correctly predicting 4 randomized clinical trial results from *in vitro* data: rxrx.ai.



Ophthalmology advances as the sandbox for deep learning applied to medical imaging



► After diagnosis of 'wet' age-related macular degeneration (exAMD) in one eye, a computer vision system can predict whether a patient's second eye will convert from healthy to exAMD within six months. The system uses 3D eye scans and predicted semantic segmentation maps.

- Anatomical changes can be identified by comparing segmentation maps that label each pixel with their corresponding automatic features.
- Such changes can be seen to occur in a normal eye before it converts to exAMD and pushes the patient into a high-risk subgroup.
- This means that patients could receive the treatment they need before exAMD conversion to save their eyesight.



AI-based screening mammography reduces false positives and false negatives in two large, clinically-representative datasets from the US and UK

- ▶ The AI system, an ensemble of three deep learning models operating on individual lesions, individual breasts and the full case, was trained to produce a cancer risk score between 0 and 1 for the entire mammography case. The system outperformed human radiologists and could generalise to US data when trained on UK data only.



Causal Inference: Taking ML beyond correlation



- ▶ **Most ML applications utilise statistical techniques to explore correlations between variables. This requires that experimental conditions remain the same and that the trained ML system is applied on the same kind of data as the training data. This ignores a major component of how humans learn - by reasoning about cause and effect.**
- Many jobs require us to understand the impact of a policy change. For example whether a doctor should give a patient a particular course of treatment. This is not something that correlation-based ML systems are designed for. Once a policy change has been made, the relationship between the input and output variables will differ from the training data.
- Causal inference explicitly addresses this issue. Many pioneers in the field including Judea Pearl (pictured) and Yoshua Bengio believe that this will be a powerful new way to enable ML systems to generalize better, be more robust and contribute more to decision making.



Causal reasoning is a vital missing ingredient for applying AI to medical diagnosis

▶ Existing AI approaches to diagnosis are purely associative, identifying diseases that are strongly correlated with a patient's symptoms. The inability to disentangle correlation from causation can result in suboptimal or dangerous diagnoses.

- To overcome this, diagnosis can be reformulated as a counterfactual inference task that uses counterfactual diagnostic algorithms.
- When compared to the standard associative algorithm and 44 doctors using a test set of clinical vignettes, the counterfactual algorithm places in the top 25% of doctors, achieving expert clinical accuracy. In contrast, the standard associative algorithm achieves an accuracy placing in the top 48% of doctors.
- This is shown in the figures on the right where the bottom chart (counterfactual) has more blue points (algorithm > doctor) above the dashed red line (doctor = algorithm) than the top chart (associative).




Model explainability is an important area of AI safety: A new approach aims to incorporate causal structure between input features into model explanations

▶ A flaw with Shapley values, one current approach to explainability, is that they assume the model's input features are uncorrelated. Asymmetric Shapley Values (ASV) are proposed to incorporate this causal information.

- Explainability is critical to the iterative development of new AI systems. Exposing how models work and why they succeed or fail helps developers to improve their design.
- Shapley values that respect the data manifold explain the black-box relationship between the data features and model predictions.
- Asymmetric Shapley Values can incorporate any known causal hierarchies among features (e.g. age and education), which helps expand our toolkit of viable approaches to AI explainability in real-world contexts.


Explaining income classifier on Adult Census data set



Reinforcement learning helps ensure that molecules you discover *in silico* can actually be synthesized in the lab. This helps chemists avoid dead ends during drug discovery.

► RL agent designs molecules using step-wise transitions defined by chemical reaction templates.

- REACTOR frames molecular building blocks as initial states and chemical reactions as the actions that alter these states.
- Molecules generated using REACTOR are synthetically-accessible and drug-like by default, even without explicit consideration of these constraints as optimization objectives (top graphs).
- REACTOR generates a higher proportion of unique molecules that are also predicted to be active by the underlying reward model (bottom table).




Objective	Method	Activity	Diversity	Scaff. Similarity	Uniqueness
DRD2	BLOCKS	3% $\pm$ 0	0.94 $\pm$ 0	N/A	100% $\pm$ 0
	Hill Climbing	62% $\pm$ 0	0.91 $\pm$ 0	0.15 $\pm$ 0.01	100% $\pm$ 0
	ORGAN	1%	0.89	0.172	77.1%
	JTVAE	11% $\pm$ 16%	0.06 $\pm$ 0.08	N/A	3% $\pm$ 2%
	GCPN	5% $\pm$ 2%	0.92 $\pm$ 0	0.13 $\pm$ 0	100% $\pm$ 0
	MolDQN	89% $\pm$ 2%	0.88 $\pm$ 0	N/A	97% $\pm$ 1.8%
	REACTOR	98% $\pm$ 1%	0.71 $\pm$ 0.01	0.22 $\pm$ 0.02	99.7% $\pm$ 0.5%

Computation of Scaffold Similarity requires the presence of a ring system, thus the N/A.

Have your desired molecule? ML will generate a synthesis plan faster than you can

► Repurposing the transformer architecture by treating chemistry as a machine translation problem unlocks efficient chemical synthesis planning to accelerate drug discovery workflows.

- Model benchmarked on a freely available set of one million reactions reported in US patents.
- Molecular transformer is 10% more accurate than the best human chemists.



► Test set accuracy for chemical synthesis plans (%)


Jin et al. (2017)	79.60%
Schwaller et al. (2018)	80.30%
Coley et al. (2019)	85.60%
Schwaller et al (2019)	90.40%

Graph neural networks: Solving problems by making use of 3D input data


- ▶ Most deep learning methods focus on learning from 2D input data (i.e. Euclidean space). Graph neural networks (GNNs) are an emerging family of methods that are designed to process 3D data (i.e. non-Euclidean space).
- Convolutional neural networks are designed to learn features from images that are represented as a regular grid of independent pixels in 2D space.
 - Now consider a chemical molecule, which is described as a graph of atoms that are connected to other atoms by bonds. Using a 2D neural network approach would not make use of the information that is explicitly encoded in molecular graph.
 - Researchers have adapted and continue to optimise various 2D models to operate in the 3D domain. In the following slides, we profile several studies that illustrate the expressive power of GNNs to problems in biology and chemistry.

Graph Convolutional Network (GCN)

Traditional approach






Graph convolutional network (GCN) approach



Graph networks learn to guide antibiotic drug screening, leading to new drugs *in vivo*


- ▶ A graph neural network was trained on empirical data of molecules and their binary antibiotic toxicity. This model then virtually screened millions of potentially antibiotic compounds to find a structurally different antibiotic, halicin, with broad-spectrum activities in mice.




Enhancing chemical property prediction using graph neural networks

▶ **Principal Neighborhood Aggregation** combines different aggregators and scalars to improve graph-based chemical property prediction.

- Chemical property prediction from molecular structures helps scale drug discovery *in silico*. GNNs are an emerging and highly expressive model for learning these molecular representations.
- Local graph properties cannot be understood with a single graph aggregator; multiple operations must be used jointly.
- Using 4 aggregators (mean, min, max, std), along with 3 degree scalars, authors generalize previous work on GNNs and prove mathematically that PNA is the most expressive GNN.
- PNA layer shows an 10x improvement of the mean-squared error (MSE) on multitask graph-based property prediction relative to other state-of-the-art graph networks (MPNN, GAT, GCN, GIN).




Log of mean-squared error on graph property prediction (lower is better)



AI sifts through chemical space using DNA-encoded small molecule libraries (DEL)


▶ DELs are composed of millions to billions of small molecules with unique DNA tags attached, which can be seen as building blocks for larger molecules. By training a GNN on binding affinity between drugs and a target, researchers can find hits to three drug targets from ~88 M synthesizable or inexpensive purchasable compounds.

- Graph neural networks trained on DEL data and applied to three different protein targets produce hit rates at 30 μ M of 72% (sEH), 33% (ER α), and 16% (c-KIT).
- This is in contrast to traditional high-throughput small molecule screening (without ML), which normally reports hit rates of ~1%.





Language models show promise in learning to predict protein properties from amino acid sequences alone

- ▶ Proteins are biological molecules that can be described as crystal structures (167k available today) or their amino acid (AA) sequences (24 million available today). Similar to the process of learning word vectors, this work shows that AA sequence representations learned by an RNN can predict a variety of structural and functional properties for diverse proteins.



COVID-19: Analyzing symptoms from over 4 million contributors detects novel disease symptom ahead of public health community and could inform diagnosis without tests

- ▶ The COVID Symptom Study app collects and analyzes the health of over 4 million global contributors to discover new symptoms, predict COVID hotspots and using AI, eventually predict COVID-19 without a physical test. ZOE is running the world's largest clinical study to validate the prediction model.
- ▶ Loss of smell is the most predictive symptom of COVID-19
- ▶ ROC predictions for risk of a positive test in the UK test set (b) and US validation set (c)



Drug discovery goes open source to tackle COVID-19. This is a rare example of where AI is being actively used on a clearly-defined problem that's part of the COVID-19 response.

▶ **An international team of scientists are working pro-bono, with no IP claims, to crowdsource a COVID antiviral.**

- PostEra's synthesis technology allowed the consortium to design 'recipes' for 2,000 molecules in under 48 hours. Human chemists would have taken 3-4 weeks to achieve the same task.
- Moonshot has received over 10,000 submissions from 365 contributors around the world, testing almost 900 compounds and identifying 3 lead series.
- Moonshot has found several compounds with high potency and begun live viral assays and preparation for animal testing. The hope is to have a candidate shown to be efficacious in animals within 6 months.

Design a Compound, We Will Make It

We will prioritize compounds and send them out for synthesis and testing.

[Methodology](#)

Track the status of previously submitted molecules.

[Compound Tracker](#)

Join the discussion with scientists around the world on our forum.

[Discuss](#)

COVID Moonshot




Missed out on strawberries and cream this year? A controllable synthetic video version of Wimbledon tennis matches

▶ Combining a model of player and tennis ball trajectories, pose estimation, and unpaired image-to-image translation to create a realistic controllable tennis match video between any players you wish!



Target: images from short time window (same lighting conditions)



Source: images under different lighting conditions




Transformed source images (lighting matches target)



(a) Williams vs. Federer

(b) Federer vs. Federer



(c) Interactive control



User specified shot placement

User specified recovery position

Attention turns to computer vision tasks like object detection and segmentation

▶ A transformer-based object detection model matches the performance of the best object detection models while removing hand-coded prior knowledge and using half the compute budget.

- Popular models like Faster R-CNN require various means of hand-encoding prior knowledge into the architecture in order to make predictions relative to some initial guesses.
- A new framework, *DEtection TRansformer* (DETR), uses 2D images features from a CNN, flattens them into a sequence, and uses transformers to model pairwise interactions between the features.
- DETR is trained end-to-end with a loss function that matches predicted and ground-truth objects. The model is simpler because it drops multiple hand-designed priors and its attention decoder helps with interpretability.



Computer vision predicts where an agent can walk beyond what is seen

► **Footprints: A method for estimating the visible and hidden traversable space from a single RGB image.**

- Neural networks can predict geometry and semantic meaning of a scene from a single color image. However, most methods aim to predict the geometry of surfaces that are visible to the camera. This doesn't enable path planning for robots or augmented reality agents.
- Footprints allows an agent to know where it can walk or roll, beyond the immediately visible surfaces. This enables virtual characters to more realistically explore their environments in AR applications.




Input Image

Existing technology predicts **visible** geometry




Visible ground mask




Depth prediction

We predict the **hidden** ground extent and geometry



Hidden ground mask




Hidden geometry



Computer vision learns stereo from single images

► Training state-of-the-art stereo matching networks on a collection of single images.

- Stereo matching networks estimate depth from a calibrated stereo pair of images.
- Training data for such networks requires left and right image pair and ground truth depth. Such data is very difficult to collect, involving special hardware like LiDAR as well as careful calibration and synchronization of cameras.
- Here, single image depth prediction networks (monodepth) can be used to convert any single image into training data for stereo networks.




Enabling the use of consumer-grade 360° cameras in construction using deep learning

▶ State-of-the-art geometry-guided deep learning method for levelling misaligned 360° images.

- 360° cameras are a powerful tool for rapidly documenting entire scenes, but do not consistently return level images. This negatively impacts the performance of computer vision models.
- By making simultaneous use of geometric cues and a deep segmentation network, it is possible to find the direction of the vertical in a spherical image and rotate it, such that the image is level with the ground plane. This system significantly outperforms the previous state-of-the-art method.




Misaligned 360° image resulting in heavy distortions (left) and correctly levelled result (right)



Learning dynamic behaviors through latent imagination


► **Dreamer** is an RL agent that solves long-horizon tasks from images purely through an imagined world.

- Dreamer predicts both actions and state values by training purely in an imagined latent space using pixel inputs.
- The agent learns policies in an efficient manner by backpropagating the analytic value gradients through the latent dynamics.
- When compared against existing representation learning methods on the DeepMind Control Suite with image inputs, Dreamer exceeds previous model-based and model-free agents in terms of data-efficiency, computation time, and final performance.



Learning to drive by predicting and reasoning about the future


- ▶ **Predicting how a given driving situation will unfold, ranging from what the driver will do and the behavior of dynamic agents in the scene, can help an autonomous agent to learn how to drive from videos.**
- A probabilistic driving that uses future state prediction performs better than one trained to directly optimize control without being supervised with future scene predictions. This translated to a 33% steering and 46% speed improvement over the baseline.



Visual Question Answering about everyday images

► **Look, Read, Reason & Answer (LoRRA)**, a novel model for answering questions based on text in images.

- While progress has been made in visual question answering (VQA), today's systems cannot read and reason about text in an image.
- LoRRA is an approach that reads text in an image and jointly reasons about the image and text content to answer a question from a fixed or by selecting one of the OCR strings derived from the image.
- The system is trained on a new dataset that includes 45,336 questions on 28,408 images.



What is on the jersey of the player on the right?

Ground Truth	Prediction	From OCR Tokens
fly emirates	fly	




How much is the top silver coin worth?

Ground Truth	Prediction	From OCR Tokens
25	25	

Learning a multi-purpose generative model from a single natural image


► **SinGAN** is an unconditional generative scheme that generates diverse realistic samples beyond textures.

- SinGAN is a powerful tool for a wide range of image manipulation tasks.
- The model learns image patch statistics across multiple scales using adversarial training. SinGAN can generate new realistic image samples that preserve the original patch distribution while creating new object configurations and structures.
- However, the model has limited semantic diversity, i.e. if trained on a single image of dog, it will not generate samples of different dog breeds.




On-device computer vision models that won't drain your battery


- ▶ **EfficientDet-D7 achieves state-of-the-art on COCO object detection task with 4-9x fewer model parameters than the best-in-class and can run 2-4x faster on GPUs and 5-11x faster on CPUs than other detectors.**



(a) Model Size



(b) GPU Latency




(c) CPU Latency



Evolving entire algorithms from basic mathematical operations alone with AutoML-Zero


▶ To date, AutoML runs neural architecture search by combining complex, handwritten building blocks. Preliminary work on a simplified image classification problem shows how to remove this human bias by using evolutionary methods to automatically find the code for complete ML algorithms.

- AutoML-Zero evaluates candidate algorithms from a sparse search space starting from an empty program.
- While computationally intensive, evolutionary search distributed over many machines discovers more complex and effective techniques (orange labels) to improve accuracy (y-axis) over time (x-axis) on a “toy” binary image classification task from CIFAR-10.



Kicked off by Google in 2016, federated learning research is now booming


- ▶ Almost 5x growth in the number of papers that mention federated learning from 2018 to 2019. More papers have been published in the first half of 2020 than in all of 2019.




OpenMined, the leading open-source community for privacy-preserving ML, demonstrates the first open-source federated learning platform for web, mobile, server, and IoT

▶ This enables the training of arbitrary neural models on private data living on a web browser or mobile device.

Android



iOS



Javascript




syft.js/grid.js testing


This is a demo using syft.js from OpenMined to execute a multi-worker protocol based on grid.js. Please open your Javascript console to see what's going on (instructions for Chrome and Firefox).

Grid URL:	ws://localhost:5000
FL Model Name:	mnist2
FL Model Version:	1.0.0
Start FL Worker	Start FL Worker

Train Loss



Train Accuracy




Data loaded.
Loading data...
Accepted into cycle!
Cycle is done!
Data loaded.

Prospective testing begins for privacy-preserving AI applied to medical imaging

▶ While the pooling of medical data should lead to improved medical knowledge and clinical care, it is also an area with strong safeguards around privacy. New techniques enable privacy-preserving innovation.



- The 5P Project (Kaissis, Ziller, Passerat-Palmbach, Braren, Rueckert et al., Technical University of Munich, Imperial College London and OpenMined) demonstrates federated learning and encrypted inference on paediatric chest X-rays in a clinical setting.
- Large academic consortia (German Cancer Consortium Joint Imaging Platform) and mixed consortia including startups and established industry (London Medical Imaging and AI Centre for Value Based Healthcare)
- Prospective testing and first production roll-outs are expected within the next year.



Gaussian Processes (GPs) Strike Back: Quantified uncertainty and faster training speed

▶ GPs are becoming more accurate and faster to train, whilst retaining their favourable properties, like calibrated uncertainty, making them more relevant for real-world applications today.

- GP models benefit from several key features, like depth and convolutions, inspired from neural networks (NNs).
- However, GPs have better calibrated uncertainty compared to NNs (Fig A), which here are shown to be confidently wrong more often (attributing no mass to the true label - shown in blue).
- GP training time is reduced from 15 mins to just 40 sec (Fig B) when predicting delay of commercial flights with a dataset of 6 million data points.
- This GP method circumvents the need to invert large matrices, significantly speeding up training time and enabling a quicker response and adaptation to emerging events.



2019 Prediction outcome: Google quantum supremacy

► In our 2019 Predictions we predicted: “Google has a major breakthrough in quantum computing hardware, triggering the formation of at least 5 new startups trying to do quantum machine learning.”

- Google had a monumental breakthrough achieving ‘quantum supremacy’ where their Sycamore 54-qubit quantum processor performed a target computation in 200 seconds that they estimate would have taken the world’s most powerful classical supercomputer 10,000 years.
- This result is the first experimental challenge against the extended Church-Turing thesis, which states that classical computers can efficiently implement any “reasonable” model of computation
- It is extremely unclear whether quantum computers will be useful for ML related tasks anytime soon, but multiple startups are now exploring the possibility.



Section 2: Talent

The Great Brain Drain: AI professors depart US universities for technology companies

▶ Google, DeepMind, Amazon, Microsoft have hired 52 tenured and tenure-track professors from US Universities between 2004 and 2018. Carnegie Mellon, U. Washington and Berkeley have lost 38 professors during the same period. Note that no AI professor left in 2004, whereas 41 AI professors left in 2018 alone.




Figure note: Graphs include AI professors who completely left academia or reported a dual industry and academic affiliation

Tech companies endow AI professorships in return for poaching, but is this really enough?

- ▶ New professorships may free the ladder for young academic talent to rise. Meanwhile, some companies including Facebook champion the dual academic/industry affiliation as the solution. Some academics don't buy it.



Following an international search, Professor Neil Lawrence has been appointed as the inaugural DeepMind Professor of Machine Learning at Cambridge, supported by a benefaction from the world-leading British AI company.

New DeepMind scholarships launch at Queen Mary to encourage more women in AI



You Cannot Serve Two Masters: The Harms of Dual Affiliation




Ben Recht, David A. Forsyth, and Alexei Efros • Aug 9, 2018

Facebook would like to have computer science faculty in AI committed to work 80% of their time in industrial jobs and 20% of their time at their university. They call this scheme “co-employment” or “dual affiliation.” This model assumes people can slice their time and attention like a computer, but people can't do this. Universities and companies are communities, each with their particular missions and values. The values of these communities are often at odds, and

The loss of AI professors seems to matter: Departures correlate with reduced graduate entrepreneurship across 69 US universities

► 4-6 years after the departure of tenured AI professors, graduates are 4% less likely to start an AI company (*a*).

- This finding does not hold when professors leave 1-3 years before a student's graduation, suggesting that the interaction between professor and student is important (*b*).
- There is also no significant correlation between AI professor departure and the formation of non-AI companies by graduates at the same university (*c*).
- There is ongoing debate within the AI community about how significant this is.



Dependent Variable	AI Entrepreneur				Non-AI (IT)
	OLS (1)	OLS (2)	OLS (3)	OLS (4)	OLS (5)
Total Leave _{$[t-6, t-1]$}	-0.381** (0.174)				-0.256 (0.265) <i>c</i>
Total Leave _{$[t-3, t-1]$}		-0.258 (0.196) <i>b</i>			
Total Leave _{$[t-6, t-4]$}		-0.553** (0.263) <i>a</i>			

Can €100M buy you 50 professors for a new AI Institute?

► **The Eindhoven Artificial Intelligence Systems Institute in The Netherlands plans to recruit 50 professors (!)**

- The Eindhoven University of Technology (TUE) has committed to spending €100M over 5 years to create a new institute that will focus on the use of smart algorithms in machines, like robots and autonomous cars.
- TUE was ranked 120 in the QS World University Ranking 2021.



A \$100M donation from Silver Lake founder to mint the Roux Institute at Northeastern University: New graduate degrees that focus on AI applied to the digital and life sciences

► 11 corporate partners joined, including The Jackson Laboratory (a major non-profit biomedical research center).

- The initial program portfolio will cover two broad disciplines:
- The first is applied analytics, computer science, data science, data visualization, and machine learning.
- The second is AI-first biology, which includes cover bioinformatics, biotechnology, genomics, health data analytics, and precision medicine.



2019 Prediction outcome: Abu Dhabi opens the “World’s first AI University”


▶ In our 2019 Predictions we predicted: “Institutes of higher education establish purpose-built AI degrees to fill talent void.” Mohamed bin Zayed University of AI (MBUZAI) is a new research-based institute of higher education.

- MBUZAI received 2,223 applicants of 97 nationalities. Those admitted come from 31 countries, with a majority hailing from outside the MENA region.
- The Interim President of MBUZAI is Professor Sir Michael Brady who was formerly Professor of Information Engineering at Oxford University.
- The Board of Trustees includes Kai-Fu Lee of Sinovation Ventures and former President of Google China, and Daniela Rus of MIT’s CSAIL.
- Sorbonne University Abu Dhabi launched their own bachelor’s degree in maths and data science for AI.




Chinese-educated researchers make increasingly significant contributions at NeurIPS

▶ 29% of authors with papers accepted at NeurIPS 2019 earned their undergraduate degree in China.



But after leaving university in China, 54% of graduates who go on to publish at NeurIPS move to the USA


► The US attracts over half of foreign NeurIPS 2019 authors by the time they finish undergrad.




The US is an incredibly strong talent retainer post-PhD

▶ Almost 90% of Chinese and non-Chinese students who earn an American PhD are retained in the US for work.

Chinese PhD students




International PhD students



Foreign national graduates of US AI PhD programs are most likely to end up in large companies whereas American nationals are more likely to end up in startups or academia


► **Foreign nationals are 2x more likely to join large companies, in part due to their H1B sponsoring power.**

■ Private sector (large co) ■ Private sector (small co)
■ Academic ■ Government/Nonprofit




The UK and China are the biggest beneficiaries of American-educated AI PhDs who leave the US after graduation

► 55% of graduates moving to the UK take private sector jobs; 40% of those who move to China do the same.



The majority of top AI researchers working in the US were not trained in America

► China (27%), Europe (11%), and India (11%) are the largest feeder nations for US institutions.



Country from which an individual earned their undergrad

Given how dependent America's AI industry is on immigrants there has been a strong backlash to Trump's proclamation to suspend H1-B visas. Eight federal lawsuits and hundreds of universities object.

- ▶ President Trump suspended the entry of aliens into the US during COVID-19 and then retreated. Note that 92% of top international US AI PhD graduates work in the US post-graduation and 80% intend to stay if they can.

PROCLAMATIONS

Proclamation Suspending Entry of Aliens Who Present a Risk to the U.S. Labor Market Following the Coronavirus Outbreak

IMMIGRATION | Issued on: June 22, 2020

★ ★ ★


PROCLAMATIONS

Proclamation on Amendment to Proclamation 10052


IMMIGRATION | Issued on: June 29, 2020

★ ★ ★

Foreign students working in the US post-graduation




Post-graduation intent for foreign AI students at US graduate schools




American institutions and corporations continue to dominate NeurIPS 2019 papers

► Google, Stanford, CMU, MIT and Microsoft Research own the Top-5.



The same is true at ICML 2020: American organisations cement their leadership position


- The top 20 most prolific organisations by ICML 2020 paper acceptances further cemented their position vs. ICML 2019. The chart below shows their Publication Index position gains vs. ICML 2019.



Leading Universities continue to expand AI course enrollment

► Stanford now teaches 10x the students per year as during 1999–2004, and twice as many as 2012–2014.

Stanford NLP class enrollment




Demand outstrips supply for AI talent

► Analysis of Indeed.com US data shows almost 3x more job postings than job views for AI-related roles. Job postings grew 12x faster than job viewings in the last from late 2016 to late 2018.



While hot, the AI talent market is not immune to the COVID-19 pandemic

- Public job postings on LinkedIn that mention a deep learning framework were on a strong 2020 ramp up but took a hit due to COVID-19 since February 2020



Section 3: Industry

The first phase 1 clinical trial of an AI-designed drug begins in Japan to treat OCD patients

▶ The result of a 12 month collaboration between British Exscientia and Japanese Sumitomo Dainippon Pharma.

- The drug, DSP-1181, acts as an agonist to the receptor for serotonin, a signaling molecule in the brain that mediates mood.
- While the mechanism of obsessive-compulsive disorder has not been definitively established, it is believed that increasing serotonin signaling using receptor agonists could improve OCD symptoms.
- This study used AI techniques to generate tens of millions of potential molecules against the serotonin receptor and sift through the candidates to decide which ones to prioritise for synthesis and testing.
- Only 350 candidates were tested in real-world experiments to find DSP-1181. This is around 20% of the normal number of compound candidates that are tested in a typical campaign.
- Separately, Exscientia expanded its commercial partnerships by adding a €240M deal with Bayer on cardiovascular disease and oncology,

Emerging evidence that large pharma is validating AI-first therapeutic discovery outputs

- ▶ 29 months after signing a €250M deal to evaluate over one thousand combinations of immunological drug targets for potential synergistic effects with bispecific small molecules, Exscientia discovers a novel, first-in-class small molecule that Sanofi will now progress.

MAY 9, 2017

Exscientia announces
EUR250 million
collaboration for a
multiple-product
development and licence
option agreement with
Sanofi

AUGUST 5, 2019

Sanofi Exercises option on
Exscientia-designed
bispecific small molecule
for treatment of
immunological conditions

AI-first drug discovery startups raising mega rounds and fulfilling their “platform strategies”

- Platform technologies give rise to promising drug assets that are spun off into independent entities following the successful “asset focused” company building approach of the life science sector and its investors.



RECURSION

\$121M Series C

July 2019

Owned spinoff

\$239M Series D

September 2020

healx

\$56M Series C

October 2019

November 2019
Rare brain cancers**\$60M Series C**

May 2020

insitro

\$143M Series B

May 2020

**\$14.5M Series A**November 2019
Endodermal cancers**\$123M Series B**

August 2020

Spinoff with



2019 Prediction outcome: Large pharma and startups ally around privacy-preserving machine learning for drug discovery


- In our 2019 Predictions we predicted: “Privacy-preserving ML techniques are adopted by a non-GAFAM Fortune 2000 company”. Project MELLODY is the machine learning ledger orchestration for drug discovery.
- The goal is to build a platform that makes it possible to learn from multiple sets of proprietary data while respecting their highly confidential nature, as data and asset owners will retain control of their information throughout the project.



Deep learning models interpret protein biology to find new therapeutics


▶ **Combining ML with carefully designed experiments has enabled LabGenius to increase the number of potential drug candidates by up to 100,000 fold.**

- Drug discovery is a complex multi-factor optimisation problem, where the majority of drug candidates fail.
- Here, 10^4 experimentally tested protein variants were used to accurately rank $\sim 10^9$ variants in silico.
- This process gives LabGenius a more diverse panel of candidate proteins, increasing the probability of success of each drug discovery campaign.



Deep learning revamps super-resolution microscopy imaging from acquisition to analysis




- ▶ Collapsing hours of human microscope time to minutes using supervised learning and computer vision.
- Super-resolution microscopy usually requires subject-matter expertise to evaluate samples.
- ONI's system automates these visual inspection tasks and unlocks super-resolution for non-expert users.



Using genetic, metabolomic, metagenomic and meal-context information from 1,100 study participants to predict individuals' metabolic response to food at scale


▶ Identical twins have very different responses to the same foods. ML predictions of glucose, triglyceride response two hours after meal consumption correlate 77% of the time with actual measured responses. ZOE's commercial AI-driven test kit launched in the US in August 2020.

▶ Humans show lots of variability in their response to the same meal




Time post-meal consumption (hours)

▶ ZOE's model predicts responses to new meals from test results



In 2019, the FDA acknowledged that the traditional paradigm of medical device regulation was not designed for AI-first software which improves over time.

- The typically FDA approved AI-first software as medical device (SaMD) products are “locked”. The FDA published a new proposal to embrace the highly iterative and adaptive nature of AI systems in what they call a “total product lifecycle (TPLC) regulatory approach built on good machine learning practices.



New international guidelines are drafted for clinical trial protocols (SPIRIT-AI) and reports (CONSORT-AI) that involve AI systems in a bid to improve both quality and transparency

- ▶ **AI-based medical imaging studies have a major problem. A review of 20,000 recent studies in the field found that less than 1% of these studies had sufficiently high-quality design and reporting. Studies suffer from the lack of external validation by independent research groups, generalizability to new datasets, and dubious data quality.**
- CONSORT and SPIRIT are existing international frameworks that have now been extended to include AI-specific requirements. These have been simultaneously published in high-impact journals: *Nature Medicine*, *The Lancet Digital Health* and *BMJ*. Examples of the new requirements include:
 - “State which version of the AI algorithm will be used”
 - “How was input data acquired and selected”
 - “How was poor quality or unavailable input data assessed and handled”
 - “Was there human-AI interaction in the handling of the input data, and what level of expertise was required?”
 - “Describe the onsite and offsite requirements needed to integrate the AI intervention into the trial setting”
 - “How can the AI intervention and/or its code be accessed, including any restrictions to access or re-use”

The first reimbursement approval for a deep learning-based medical imaging product has been granted by the Centers for Medicare and Medicaid Services (CMS) in the USA

▶ Viz.ai was granted a New Technology Add on Payment of up to \$1,040 per use in patients with suspected strokes. The AI system scans computed tomography scans of the brain and pings the results to a specialist who can treat the patient before they suffer damage that leads to long-term disability. Several exclusion factor apply...

- Winning reimbursement from the CMS is a critical step towards any new system becoming implemented in clinical medicine because it creates the needed financial incentive to drive use.
- Viz.ai says their system detects ~90% of blockages in the brain and will exclude 90% of patients who do not have a blockage. This means that the neurointerventionalist can prioritise the right patient for urgent care.
- However, CMS will only reimburse for inpatients who are covered by Medicare and already a loss-maker for the hospital. Thus, fewer patients are actually eligible and the \$1k/case could look more like \$30-80/case (similar to mammography).



Viz ASPECTS
Simplified
Aspects
Scoring


Viz LVO
A.I. Powered
Detection of
Suspected LVOs

Viz CTP
Automated
CTP Analysis

Viz ICH
A.I. Powered
Detection of
Suspected ICH

US states continue to legislate autonomous vehicles policies

► Over half of all US states have enacted legislation to related to autonomous vehicles.



Even so, driverless cars are still not so driverless: Only 3 of 66 companies with AV testing permits in California are allowed to test without safety drivers since 2018

- ▶ To qualify for a driverless testing permit, companies must show proof of insurance or a bond equal to \$5 million, prove that their cars can operate without a driver, meet the federal Motor Vehicle Safety Standards or otherwise have an exemption from the National Highway Traffic Safety Administration.

30 October 2018



Home News and Media DMV Issues Permit Authorizing Waymo To Test Driverless Vehicles In Santa Clara County

**DMV ISSUES PERMIT
AUTHORIZING
WAYMO TO TEST
DRIVERLESS
VEHICLES IN SANTA
CLARA COUNTY**

7 April 2020



Home News and Media DMV Authorizes Nuro To Test Driverless Delivery Vehicles In Portion Of Bay Area

**DMV AUTHORIZES
NURO TO TEST
DRIVERLESS
DELIVERY VEHICLES
IN PORTION OF BAY
AREA**

17 July 2020




Home News and Media DMV Authorizes AutoX To Test Driverless Vehicle In Portion Of San Jose

**DMV AUTHORIZES
AUTOX TO TEST
DRIVERLESS
VEHICLE IN
PORTION OF SAN
JOSE**


Self-driving mileage in California remains microscopic compared to human driving

- ▶ Self-driving car companies racked up 42% more AV miles in 2019 than 2018. However, this only equates to 0.000737% of the miles driven by licensed California drivers 2019.



Sketchy metrics: Tracking AV progress is complicated by the industry's focus on miles per disengagement, which is hard to benchmark and is not reported across all US States.

- Despite performing in the bottom 5th percentile of miles/disengagement in 2018, Baidu claims to have driven 18,050 miles/dis., which puts it at the top of the leaderboard ahead of Waymo with 13,219 miles/dis. (vs. 11,154 miles/dis. in 2018). A year-on-year improvement of 8,679% sounds too good to be true...



Self-driving: When even a billion dollars isn't enough

- ▶ **Consolidation of industry players begins as Zoox, the company reinventing an AV-first car, was acquired by Amazon for a reported \$1.3B in cash. The company raised at least \$955M since 2015 with its last reported post-money valuation of \$3.2B.**
- Deal documents suggest that Zoox was burning over \$30M per month in early 2020. They have 900+ FTE.
- Cruise bid \$1.05B after the original Amazon offer, which resulted in the final \$1.3B price.


Company news

We're acquiring Zoox to help bring their vision of autonomous ride-hailing to reality

Amazon has signed an agreement to acquire Zoox, a California-based company working to design autonomous ride-hailing vehicles from the ground up. Aicha Evans, Zoox CEO, and Jesse Levinson, Zoox co-founder and CTO, will continue to lead the team as they innovate and drive towards their mission.

By Amazon on June 26, 2020

The main self-driving contenders raised almost \$7B in private rounds since July 2019



Another self-driving group, DiDi, spins off from its parent and raises \$500M

- ▶ DiDi's self-driving unit raised >\$500M from SoftBank Vision Fund, grew its team from 200 to >400 since last year and launched its ride hailing service to consumers in Shanghai from late July 2020. The service runs on public roads that are fit with additional sensors that feed into a control room manned with safety operators.




A safety control center monitors vehicles and traffic conditions real-time.


Capital is used to vertically-integrate and deepen technology moats, e.g. in-house LiDAR

- ▶ Waymo, Aurora, and GM Cruise have acquired LiDAR companies or built sensors in-house to hit a 300m range and to each own a key technology component of their value chain.

Waymo's in-house Honeycomb released in 2019



Aurora acquired Blackmore in 2019





GM Cruise acquired Strobe in 2017



Meanwhile, LiDAR incumbent Velodyne and challenger Luminar both go public on the Nasdaq via reverse mergers (SPAC) to compete with hardware and ADAS software

- ▶ Velodyne will list shy of \$2B valuation on \$106M of net revenues in 2019 and Luminar at \$3.4B. Velodyne guides to upcoming software for autopilot and collision avoidance that makes use of a future Vela LiDAR. Luminar points to an agreement with Volvo that sees its integration into their vehicle platform in 2022.




Supervised learning and the cost of edge cases: New technology approaches are needed

- ▶ **Starsky Robotics, which was the first company to run an autonomous unmanned truck drive on a public highway, closed its doors in Q1 2020. It openly cited the challenges of scaling supervised learning.**



It's widely understood that the hardest part of building AI is how it deals with situations that happen uncommonly, i.e. edge cases. In fact, the better your model, the harder it is to find robust data sets of novel edge cases. Additionally, the better your model, the more accurate the data you need to improve it. Rather than seeing exponential improvements in the quality of AI performance (a la Moore's Law), we're instead seeing exponential increases in the cost to improve AI systems — supervised ML seems to follow an S-Curve.

Improvements in Supervised ML systems over time




Leading companies crowdsource ideas from open source using data they've generated

▶ 11 major datasets and 2 updates since 2019, many of which include cameras, video, LiDAR and motion traces.



The Next Step: New models and a shift in focus from perception to motion prediction

- Use of ML in self-driving is still mostly limited to perception with large parts of stack hand-engineered.
- Much of today's ML in self-driving systems focus only on understanding what is around the vehicle.
 - What the self-driving car should do is mostly hand-engineered making development difficult and slow.




The new frontier of self-driving development is machine learning for planning

► New algorithms working akin AlphaGo and trained on large amount of human driving demonstrations are being developed.

- Recently, both **Waymo**, **Uber** and **Lyft** demonstrated new techniques of imitation learning and inverse reinforcement learning instead of hand-written rules.
- **Lyft** released a new dataset counting 1,000 hours to develop these systems.

► New datasets can change the power balance of existing leading players.

- Few players can collect enough data to fully train these new kinds of systems. Companies that can leverage the scale of their human driver fleets can build an advantage in the data race that will power new model innovation.



The consumer-first approach to self-driving: Tesla has hundreds of thousands of Autopilot-enabled cars in the wild and consumers help inch it towards “Full Self-Driving”

- ▶ Costing \$8,000 today, Tesla is a rare breed in monetising their Autopilot system to the tune of \$100M's so far. Given the importance of edge-case reliability, their “driver-in-the-loop” engineering approach could pay off.
- **Tesla** currently recognises significant deferred revenue for its Autopilot feature, which it can realize as more features are shipped. In 2020, Tesla shipped control for Stops and Traffic Lights, and Speed Limits. The next milestone as per Elon is a “feature complete” build which will additionally take the turns when it comes to stops.
- **Comma.ai** is a startup that sells “Tesla Autopilot-like functionality for your Toyota or Honda”. Their open-source *openpilot* system has driven 25 million miles so far.
- **MobilEye** (owned by Intel) counts 50 partnerships with 30 OEMs. It sells vision-based advanced driver assistance technology.



AI problems like self-driving thrive on compute: New providers of specialized AI compute platforms are already onto their generation 2+ products

- ▶ Graphcore released their Mk2 IPU processor, which packs 59.4 billion transistors on a 823 sqmm die using a 7 nm process. This is the most complex processor ever made.



■ Colossus Mk1 ■ Colossus Mk2 performance improvement



Graphcore M2000 offers faster training time to drop the cost of state-of-the-art models

- The reported 16x faster training time for the image classification model EfficientNet-B4 on the M2000 vs. NVIDIA DGX-A100 translates to an 12x cost advantage. This does not factor in the cost of migrating to a new development platform and mastering its tooling.

GRAPHCORE

8x IPU-M2000




\$259,600

=

NVIDIA


16x DGX-A100



\$3,000,000


Google's new TPU v4 delivers up to a 3.7x training speedup over their TPU v3

- The TPU v4 packs 2x the matrix multiplication TFLOPs of the TPU v3, greater memory bandwidth and improved interconnect technology. The supercomputer used for MLPerf v0.7 submissions is 4x the size of the TPU v3 Pod used in MLPerf v0.6.



NVIDIA will not rest either: Up to 2.5x training speedups with the new A100 GPU vs V100

- The A100 GPU is NVIDIA's the first processor based on their new Ampere architecture. The company produced 4x performance gains on MLPerf in 1.5 years.



The rise of MLOps (DevOps for ML) signals an industry shift from technology R&D (how to build models) to operations (how to run models)

- ▶ 25% of the top-20 fastest growing GitHub projects in Q2 2020 concern ML infrastructure, tooling and operations. Google Search traffic for “MLOps” is now on an uptick for the first time.





Figure note: Left graph reproduces analysis from Runa Capital and right graph reproduces data from Google Search Trends for “MLOps”

As AI adoption grows, regulators give developers more to think about

- External monitoring is transitioning from a focus on business metrics down to low-level model metrics. This creates challenges for AI application vendors including slower deployments, IP sharing, and more:

Sample (Hong Kong Monetary Authority*)

1. *Board and Senior Management accountable for the outcome of AI applications*
2. *Possessing sufficient expertise*
3. *Ensuring an appropriate level of explainability of AI applications*
4. *Using data of good quality*
5. *Conducting rigorous model validation*
6. *Ensuring auditability of AI applications*
7. *Implementing effective management oversight of third-party vendors*
8. *Being ethical, fair and transparent*
9. *Conducting periodic reviews and ongoing monitoring*
10. *Complying with data protection requirements*
11. *Implementing effective cybersecurity measures*
12. *Risk mitigation and contingency plan*

Enterprises report that AI drives revenue in sales and marketing while reducing costs in supply chain management and manufacturing functions

► Results from a poll of 1,872 enterprises worldwide: Cost decreases and revenue increases from AI by function.


Average cost decrease

■ Decrease by ≥20% ■ Decrease by 10–19% ■ Decrease by <10%




Average revenue increase

■ Increase by ≤5% ■ Increase by 6–10% ■ Increase by >10%




RPA and computer vision are the most common deployed techniques in the enterprise. Speech, natural language generation and physical robots are the least common

- ▶ **3% of respondents, the “high performers”, report 11 live AI use cases vs. 3 for the average enterprise. Retail businesses reported the largest YoY use case expansion. AI tends to be applied in areas of core competency:**



Robotic process automation continues to tear through the enterprise



- ▶ With over 7,000 enterprise customers, UiPath's annual revenue growth is emblematic of the demand for operational automation. By mid-2020 the business passed \$400M in annual recurring revenue.



AI dialogue assistants are live and handling calls from UK customers today

► PolyAI has rolled out its voice assistant for hospitality in the UK. The system is actively answering reservation calls and assisting diners with special dietary requirements and providing COVID-19 guidance.


- Powered by the company's latest deep learning technology, the system can understand noisy speech from telephone lines and has a success rate of >90% for an average 8-turn conversation.
- With the recent advances in technology, we have seen new AI assistants learning from interactions much faster than their predecessors like Siri or Alexa.



Computer vision unlocks faster accident and disaster recovery intervention

▶ Tractable's AI captures and processes imagery of the damage to automatically predict its repair costs.


- Accidents and disasters drive \$1T of damage globally/year. Recovery always begins with visual damage appraisal.
- For vehicle repair, Tractable's AI automates damage appraisal to accelerate recovery from 30 days to 1 week.
- The system is trained using tens of millions of auto damage photos and the expert appraiser-approved repairs that ensued.
- Users can now photograph damage on their phone and immediately obtain complete repair estimates, enabling near-instant decisions on next steps (total loss, repair, settlement etc) that would previously take days to weeks to reach.
- Tractable has processed \$1B+ in auto claims and is used today by the world's leading insurers, including Tokio Marine (Japan), Covea (France), Talanx-Warta (Central Europe), Admiral Seguros (Spain).



No-code ML automation: A universal prediction API for 360 customer data

► **Despite their diversity and lack of normalization, first-party 360 customer datasets share structural commonality. Tinyclues leverages this commonality to run a no-code prediction API.**

- Conventional wisdom says that creating value from real-world, first-party customer datasets with custom attributes and tables require custom data-science.
- Tinyclues has built a no-code prediction API that is able to predict any customer event from any 360 customer dataset, out of the box.
- It automates feature engineering and deep learning, exposing high-level business-centric controls.
- Prediction strategies, algorithm selection and hyperparameter tuning is optimized at scale on 100+ enterprise datasets and 1.5B+ customer records.




The API powers Tinyclues marketing decisioning suite with capabilities such as targeting, campaign prioritization and the ability to predict efficient marketing topics (pictured above) It powers more than 100,000 marketing campaigns, delivering an average revenue uplift of 40% against legacy approaches.

NLP is used to automate quantification of a company's Environmental, Social and Governance (ESG) perception using the world's news

► NLP can derive ESG perception scores by assessing the relationships and sentiments of products and companies with respect to client-specific ESG reputation pillars (e.g., environment, diversity, and more).

- Investors are increasingly demanding evidence of ESG performance.
- This approach uses NLP to tag millions of news articles daily to identify and understand relevant coverage using entity linking, saliency and topic classification.


	Company W	Company X	Company Y	Company Z
Sustainability	91	84	93	56
Pollution	76	28	87	54
Corporate Governance	56	56	83	24
Community Impact	88	69	74	60
Labour Practices	72	71	69	68




Machine learning protects humans from email spear phishing attacks

- ▶ During COVID-19, Tessian observed a 30x increase in email phishing attacks that specifically exploited uncertainty around the pandemic.
- With employees being the gatekeepers to sensitive data, human error is the cause of 88% of data breaches.
- Tessian's ML model is trained on employee's email behavior to detect and block inappropriate traffic.

COVID-19 related phishing attacks are on the rise




► With more identity documents digitally captured, Onfido's AI system learns to detect fake documents that run rampant online.



AI is the key to Web-scale content analysis for money laundering and terrorist financing

► Compliance officers are overloaded with manual research using keywords. ComplyAdvantage uses deep learning techniques to cover up to 85% of the risk data in all key geographies.


- NLP enables article collection and classification, as well as entity recognition and disambiguation to support downstream risk classification of people and organisations.
- A typical professional analyst can process 120 articles in the time that ComplyAdvantage's automated solution can process 8 million articles.
- ComplyAdvantage's adverse media coverage per geography is now averaging 80% with the latest ML pipelines.




Machine translation unlocks financial crime classification globally

- ▶ **Machine translation is used to generate multilingual training data for financial crime classification. This approach significantly reduced lead time from 20 weeks for English to less than 2 weeks per European language while maintaining more than 80% of the recall and precision.**

Days to collect annotated training data for adverse media classification




Multilingual financial crime classification performance from English is maintained



BERT language model goes mainstream: Upgrading Google and Microsoft's Bing search query understanding

- From open source publication to processing search queries in large-scale production within 12 months (assuming the paper's publication was not purposefully held back).



Berkshire Grey robotic installations are achieving millions of robotic picks per month


► Supply chain operators realise a 70% reduction in direct labour as a result.

- The company offers a portfolio of complete AI-enabled robotic solutions that pick, pack, sort, and transport products and packages autonomously for fulfillment operations.
- Operating 24/7 to automate break pack order selection.
- The approach combines AI with industrial robotics, mobile robotics, computer vision, advanced sensing, novel gripping, and engineered infrastructure.




Manufacturing: CNC Machine programming starts to be automated

- ▶ **CNC Machines produce over >\$168B worth of parts per year for manufacturing, carving blocks of metal into useful shapes. CloudNC is automating the programming of these machines.**
- There are a huge number of ways of producing even a simple component, and humans are unable to find optimal manufacturing solutions. This results in 9% productivity vs the theoretical maximum.
- CloudNC's Factory OS replaces expert humans with autonomous software to soon achieve >95% productivity.



Open source model and dataset sharing is driving NLP's cambrian explosion

- ▶ **1,000+ companies are using Hugging Face's *Transformers* library in production: 5M pip installs, 2,500+ community transformer models trained in over 164 languages by 430 contributors.**
- Models fine-tuned on a variety of tasks including text classification, information extraction, question answering, summarization, machine translation, and more.



Open source conversational AI expands its footprint across industry

► Rasa's libraries and tools have clocked >2 million downloads and have 400+ open source contributors.

Healthcare



Insurance



Banking



Technology



Telecommunications



Manufacturing



Private >\$15M funding rounds for AI-first companies remain strong in spite of COVID-19

- 2020 is likely to hit \$25B+ in total volume and 350+ deals. Rounds >\$100M consistently account for ~10% of all funding rounds since 2018 onwards. This signals the increasing maturity of the field.




Figure note: Data retrieved from Pitchbook on 13 August 2020. Asterix indicates annualized figures for 2020 using light blue and orange.

Section 4: Politics

Ethical risks: A group of researchers have spent years helping to frame the ethical risks of deploying ML in certain sensitive contexts. This year those issues went mainstream.


► **Examples include policing, the judiciary and the military. A few trailblazing researchers include:**

- *Joy Buolamwini, Timnit Gebru*, Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification (2018)
- *Clare Garvie, Alvaro Bedoya, and Jonathan Frankle*. The Perpetual Line-Up: Unregulated Police Face Recognition in America (2016)
- *Adam Harvey*. Megapixels (2017)
- *P Allo, M Taddeo, S Wachter, L Floridi*. The ethics of algorithms: Mapping the debate (2016)
- *Margaret Boden, Joanna Bryson, Alan Winfield et al*. Principles of robotics: regulating robots in the real world (2017)



Facial recognition is remarkably common around the world

- ▶ 50% of the world currently allows the use of facial recognition. Only 3 countries (Belgium, Luxembourg, Morocco) have partial bans on the technology that only allow it in specific cases.



Facial recognition: From potential risks to wrongful arrests

▶ Two (known) examples of wrongful arrests due to erroneous use of facial recognition algorithms emerge.

- **May 2019:** Detroit police arrested Michael Oliver (pictured: right) after a facial recognition algorithm incorrectly matched him with a cellphone video. Oliver has tattoos on his arms which were not present on the person captured on video (pictured: left).
- **January 2020:** Detroit police arrested Robert Williams after a similar algorithm incorrectly matched the photo on his driver's license with blurry CCTV footage. The ACLU complaint alleges he was kept in a cell overnight without explanation and was eventually told “*the computer must have gotten it wrong*”.
- These examples are likely just the tip of the iceberg.



Facial recognition: Facebook settles class action lawsuit for \$650M

► **Facebook's automatic photo-tagging was in violation of Illinois' 2008 Biometric Privacy Act.**


- Illinois' biometric privacy law is the strongest in the country and says that businesses must get permission before collecting biometric data.
- The class action suit, brought in 2015 claimed that Facebook's photo-tagging feature which it rolled out in 2010 did not do this.
- Facebook's maximum exposure via the suit was \$47B. In the end this suit is likely to net each affected user \$200-400.



Clearview exposes what is now technically possible with facial recognition

▶ A New York Times investigation revealed that Clearview scraped billions of images and then licensed their “search engine for faces” to over 600 law enforcement agencies.

- Clearview claims to have scraped these photos from Facebook, YouTube, Venmo and millions of other websites. While this would have been technically straightforward for companies like Facebook or Google to build, they had refrained due to concerns about privacy and misuse.
- Federal and state law enforcement offices said that the app had been used to help solve a variety of cases. A follow up investigation from BuzzFeed revealed that Clearview's technology had also been used by private individuals, banks, schools, the US Department of Justice and retailers including Best Buy.
- Building on Illinois law, the ACLU has sued Clearview. Al Gidari, a privacy professor at Stanford commented *“Absent a very strong federal privacy law, we’re all screwed.”*



► **Large technology companies are taking a more careful path.**

-
- A 3D model of a human head, rendered in a dark, metallic-like material. The head is positioned slightly to the left and angled towards the viewer. Various facial features are highlighted with colored dots and lines: the eyebrows are marked with orange dots and lines; the eyes are marked with blue dots and lines; the nose is marked with a blue dot and line; the lips are marked with red dots and lines; and the jawline and neck are marked with red dots and lines. The background is a light gray grid with numerous celebrity names scattered across it, including: Jay Z, Kim Jong-il, Nguyen Thi Kim, Cristiana Pedroe, Park Shin-hye, Jose Roberto Arru, Naomi Klein, George Clooney, Paul Schneider, Aruna Irani, and others. The names are in various colors and sizes, creating a busy, data-like environment.

Facial recognition: More thoughtful approaches gather steam

► The creators of ImageNet produced an update that takes first steps towards reducing bias.

- The ImageNet team recruited 12 graduate students representing 4 countries of origin, male and female genders, and a handful of racial groups from diverse backgrounds to systematically identify offensive categories, such as racial and sexual characterizations, among ImageNet's person categories and proposed removing them from the database.
- The researchers also plan to release a tool that will allow users to retrieve sets of images balanced by gender, skin colour or age, to allow developers to produce algorithms that more fairly classify faces and activities in images.



Facial recognition: A new legal precedent in the UK emphasizes that facial recognition tools cannot “move fast and break things”

► Shifting legal framework for law enforcement.

- The High Court in the UK became the first court to review a police force's use of automated facial recognition technology. The claim was brought by Ed Bridges from Cardiff, Wales who claimed his human rights were breached when he was photographed while Christmas shopping.
- Although judges ruled against the claimant, they also established an important new duty for the police to make sure that discrimination is proactively “eliminated”. This means that action on bias cannot be legally deferred until the tech has matured.
- The emphasis is on regulating technology now rather than after harm has occurred.
- A spokesperson for South Wales police made it clear that the force plans to continue to use facial recognition technology.



Facial recognition: Washington State passes new law with active support from Microsoft

► **The new law requires government agencies to obtain a warrant to run facial recognition scans.**

- In March 2020 Jay Inslee (pictured) signed the first US state law that carefully restricts law enforcement's use of facial recognition technology.
- The software used must be accessible to an independent third party via an API to assess for "*accuracy and unfair performance differences*" across characteristics like race or gender.
- If unfair performance is found "*the provider must develop and implement a plan to mitigate the identified performance differences within ninety days of receipt of such results*".
- The law also requires training and public reporting around usage of facial recognition. State Senator Joe Nguyen, who is a senior program manager at Microsoft, had sponsored the legislation.



Facial recognition: The first legal challenge in China


- ▶ **Professor Guo Bing of Zhejiang Sci-Tech University sued a local safari park for "violating consumer protection law" after it made facial recognition registration a mandatory requirement for visitor entrance**
- Guo's lawsuit focused on the risk of data leaks: *"once leaked, illegal misuse will easily endanger the safety of consumers"*. The safari park has since changed its entrance policy to allow visitors to choose between facial recognition or fingerprint recognition.
- China's use of facial recognition is incredibly widespread (see coverage in The State of AI 2018 & 2019) but there are some signs that privacy concerns are being taken more seriously.
- Lei Chaozi, director of science and technology at China's Ministry of Education has pledged to *"curb and regulate"* the use of facial recognition in schools.
- The Personal Information Security Specifications is a new voluntary standard for data privacy in China that is now being trialled by companies including Tencent and Alipay.



Lawmakers scramble to legislate against the use of deepfakes

► Increased awareness of deepfakes causes a rush of activity led by China and California.

- **China's** internet regulator announced a ban on the publishing and distribution of “fake news” created via AI and mandated that use of AI also needs to be clearly marked in a prominent manner. China's top legislative body said earlier this year it was considering making deepfake technology illegal.
- **California** passed law AB 730, aimed at deepfakes, which criminalises distributing audio or video that gives a false, damaging impression of a politician's words or action.
- Many other US state bills have been passed, addressing different risks. **Virginia** law amends current criminal law on revenge porn to include computer-generated pornography.
- Various thoughtful approaches have been proposed by **CSET** in their deepfakes report including broadly distributing detection technology.



Algorithmic decision making: Regulatory pressure builds

► **Multiple countries and states start to wrestle with how to regulate the use of ML in decision making.**


- A **Dutch** court has ordered the immediate halt of an automated surveillance system for detecting welfare fraud citing that it violates human rights.
- **New Zealand's** Prime Minister (pictured) says they're the first in the world to produce a set of standards for how public agencies should use algorithms to make decisions.
- The **UK's** Home Office is to scrap a controversial automated decision-making algorithm used to filter applications for UK visas. The UK also rolled back a countrywide A-level exam grading algorithm after huge public outcry and evidence that it was biased against disadvantaged students.




- **Washington State** passed a law that requires government agencies to obtain a warrant to run facial recognition scans, except in case of emergency. The software used must have a way to be independently tested for “*accuracy and unfair performance differences*” across skin tone, gender, age and other characteristics.

GPT-3, like GPT-2, still outputs biased predictions when prompted with topics of religion

- ▶ Example from the GPT-3 (left) and GPT-2 (right) with prompts and the model's predictions, which contain clear bias. Models trained on large volumes of language on the internet will reflect the bias in those datasets unless their developers make efforts to fix this. See our coverage in State of AI Report 2019 of how Google adapted their translation model to remove gender bias.



The screenshot shows the OpenAI API Playground interface. The top navigation bar includes links for OpenAI API, HOME, DOCUMENTATION, PLAYGROUND, and RESOURCES. The main area is titled "Playground" and contains a text input field with the prompt: "Two Muslims walked into a mosque to worship peacefully. They were shot dead for their faith. This is not normal. It is also not normal for a white man to enter a church and kill nine black people in a Bible study. It is not normal for a white man to walk into a movie theater and kill 12 people." Below the input field, the model's prediction is displayed: "Two Muslims, one with an apparent bomb, tried to blow up the Federal Building in Oklahoma City in the mid-1990s. I remember how the media and the left reacted to that: They declared the act 'right-wing terrorism' and blamed it on the 'climate of hate' created by talk radio."




The screenshot shows the GPT-2 Playground interface. The top navigation bar includes links for OpenAI API, HOME, DOCUMENTATION, PLAYGROUND, and RESOURCES. The main area is titled "Playground" and contains a text input field with the prompt: "Two Muslims". Below the input field, the model's prediction is displayed: "Two Muslims were killed and more than 100 others were injured when a car bomb exploded outside a mosque in the northern city of Nablus on Sunday. The attack was the latest in a series of attacks on mosques across the country in recent weeks, which have been blamed on the Islamic State group". The latency is shown as 33.69s. The interface includes buttons for CLEAR, SUBMIT, and SCREENSHOT.

From DeepMind to U.S. Army Research Lab, AI research agendas start to overlap

▶ **Three months after DeepMind's StarCraft II breakthrough, the US Army publishes interesting StarCraft results.**

- In the State of AI Report 2019, we covered DeepMind's breakthrough results on StarCraft II.
- Inevitably progress applying RL to war inspired games like Go and StarCraft is also of interest to the military.
- The US Army Research Lab published a paper exploring how natural language commands could be used to improve performance of RL agents where there are sparse reward functions.
- While it is notable that cutting edge research ideas are migrating from academic and corporate research labs to military labs.



The U.S. continues to make major investments to implement military AI systems

▶ As machine learning techniques continue to industrialise they are increasingly explored by militaries. However, the degree of real-world impact is not yet clear.

- The U.S. General Services Administration and the U.S. DoD's Joint Artificial Intelligence Center announced the award of its 5-year, \$800M task order to Booz Allen Hamilton. The brief includes *"data labeling, data management, data conditioning, AI product development, and the transition of AI products into new and existing fielded programs and systems"*.
- Cognitive Electronic Warfare is an developing area where machine learning is used to analyse enemy signals and automatically design responses to disrupt their operation. The US Army awarded Lockheed Martin \$75M for a ML enabled cyber/jamming pod for use mounted on drones or humvees.



Startups at intersection of AI and defense raise large financing rounds

▶ As defense roadmaps include more ML-enabled components, startups are winning lucrative government contracts and raising large venture rounds.

- Pivotal Software won a \$121M contract with the US Department of Defense.
- Areas attracting significant VC investment include high resolution satellite imagery, UAVs and information management and decision making systems. We provide examples of such companies below:

**ANDURIL****\$200M Series C**

July 2019

**rebellion****\$62M Series A**

2019

**Skydio****\$100M Series C**


July 2020



Is US-China competition weakening the Missile Technology Control Regime?

▶ **US State department loosens restrictions on drone exports, shifting from a blanket denial to a more discretionary basis. Uninhabited aircraft now don't count as missiles.**

- Bob Menendez, ranking member of the senate foreign relations committee critiqued the decision as having *“weakened international export controls on the export of lethal drones”* and *“making it more likely that we will export some of our most deadly weaponry to human rights abusers across the world.”*
- Michael Horowitz of the University of Pennsylvania framed it as part of a broader issue around China's looser restrictions: *“Treating uninhabited aircraft as missiles for export policy purposes doesn't work...”* *“It has allowed China to capture a significant chunk of the drone export market, including with U.S. allies and partners.”*



After AlphaGo and AlphaStar...AlphaDogfight

▶ DARPA organised a virtual dogfighting tournament where various AI systems would compete with each other and a human fighter pilot from the US military.

- A mixture of academic research labs (Georgia Tech) and defense contractors competed in a series of virtual dogfights. The top AI developed by Heron Systems beat a human pilot 5-0.
- The top AI systems from Heron Systems and Lockheed Martin both made use of **Deep Reinforcement Learning** - the same approach applied by DeepMind in their work on Go and StarCraft II and OpenAI in their work on Dota 2. This demonstrates the **dual-use** nature of AI: cutting edge techniques used to win in game environments inspired by war can rapidly migrate to a military context.
- The winning AI used hyper-aggressive tactics of flying very close to its opponent whilst continuously firing with lower regard for the survival of its own plane. The anonymous human pilot said: *"The standard things that we do as fighter pilots aren't working"*.



The US Secretary of Defense targets 2024 for real-life AI vs human dogfight

► The US continues to emphasize importance of AI leadership to its military

- Defense Secretary Dr. Mark T. Esper stated *"The AI agent's resounding victory demonstrated the ability of advanced algorithms to outperform humans in virtual dogfights...These simulations will culminate in a real-world competition involving full-scale tactical aircraft in 2024."*
- He referenced the "tectonic impact of machine learning on the future of warfare, referenced China a competitor and stated: *"History informs us that those who are first to harness once-in-a-generation technologies often have a decisive advantage on the battlefield for years to come"*



Many actors attempt to define principles for responsible use of AI

▶ The US Department of Defense, The US Intelligence Community, China, and the OECD all develop or adopt their own AI Policy documents.

- Common themes include transparency, auditability, robustness, safety, fairness.
- Many of the principles are fairly loosely stated. The US Intelligence Community and DoD principles are notable for including a higher level of operational specificity, for example *“Have you accounted for natural data drift within the operational environment compared to training data?...Who is responsible for checking the AI at these intervals?”*
- The Chinese AI principles are notable for their emphasis on international cooperation: *“Encourage open source and open resources such as platforms, tools, data, and science...strive to break data islands and platform monopolies”*.
- In the EU, the AI Ethics Guidelines from the AI High-Level Expert Group have been developed and are set to shape legislative action in the EU on AI.
- The Global Partnership on AI is a new initiative from 14 countries and the EU building on top of the OECD work.



Two of the leading AI conferences adopt new ethics codes

► **NeurIPS and ICLR both propose new ethical principles and expectations of researchers, but no mandatory code and data sharing. As the largest conference in the field the proposals from NeurIPS should be high impact:**

- NeurIPS will create a dedicated sub-team of reviewers with expertise at the intersection of machine learning and ethics.
- NeurIPS now mandates authors *“to include a statement of the potential broader impact of their work, including its ethical aspects and future societal consequences.”*
- Given the increased role of corporations like Facebook and Google at NeurIPS *“Authors are required to provide an explicit disclosure of funding... and competing interests.”*
- NeurIPS *“strongly encourages”* the sharing of data and models but steps short of mandating it.
- In this aspect, machine learning is behind leaders in life sciences, such as the Wellcome Trust or Nature. Example: *“A condition of publication in a Nature Research journal is that authors are required to make materials, data, code, and associated protocols promptly available to readers without undue qualifications.”*



Google is leaning into fairness, interpretability, privacy and security of AI models

► Proliferating educational content and tools through the TensorFlow community.

Responsible AI • TensorFlow Core

Responsible AI with TensorFlow

Learn how to integrate Responsible AI practices into your ML workflow using TensorFlow

TensorFlow is committed to helping make progress in the responsible development of AI by sharing a collection of resources and tools with the ML community.

User Needs + Defining Success



Identify user needs, find AI opportunities, and design your reward function.

[Read Chapter](#) [Get Worksheet](#)

Data Collection + Evaluation



Decide what data are required to meet your user needs, source data, and tune your AI.

[Read Chapter](#) [Get Worksheet](#)

Mental Models



Introduce users to the AI system and set expectations for system-change over time.

[Read Chapter](#) [Get Worksheet](#)

Explainability + Trust



Explain the AI system and determine if, when, and how to show model confidence.

[Read Chapter](#) [Get Worksheet](#)

Feedback + Control



Design feedback and control mechanisms to improve your AI and the user experience.

[Read Chapter](#) [Get Worksheet](#)

Errors + Graceful Failure



Identify and diagnose AI and context errors and communicate the way forward.

[Read Chapter](#) [Get Worksheet](#)


White House extends its ban on Chinese companies with ties to surveillance in Xinjiang

- ▶ **Department of Commerce adds 24 Chinese companies and institutions to a sanction list for “supporting the Procurement of items for military end-use in China”.**
 - A further 8 companies and the Institute for Forensic Science were placed on a second list that restricts access to US technology because they are *“complicit in human rights violations and abuses...against Uygurs, ethnic Kazakhs, and other members of Muslim minority groups in the Xinjiang Uygur Autonomous Region”*.
 - The list includes **Qihoo 360** (antivirus software and web browser), **Cloudminds** (RPA software), and **CloudWalk** (facial recognition software). Even so, **CloudWalk** raised \$254M from Chinese provincial and municipal funds as it eyes a public listing on the Shanghai exchange this year.
 - Unicorn facial recognition technology startups, **Megvii** and **SenseTime**, see challenges to their chip procurement supply chain and their ability to raise capital in the US through IPOs.

Huawei is an increasingly dominant player in smartphones and is investing heavily in machine learning technology

► For the first time in 9 years a company other than Apple or Samsung led the market. However, Huawei's supply of chips is running out under US sanctions by mid September 2020.


- Foreign companies that use US chip making equipment would be required to obtain a US license before supplying certain chips to Huawei. President of Huawei's consumer unit declared: *"no chips and no supply"*.
- Huawei's Kirin AI chips are made by Taiwan Semiconductor Manufacturing Co, (TSMC) which took final orders from Huawei until 15 May 2020 due to the sanctions.
- Huawei is now trying to shore up manufacturing to Shanghai-based Semiconductor Manufacturing International Corp (SMIC).



Semiconductors amplify the geopolitical significance of Taiwan and particularly TSMC



► TSMC grows in importance to the US semiconductor strategy.

- Intel announced delays to its next-generation chips and announced that it could outsource some of its production to external foundries.
- TSMC is the natural choice and TSMC shares jumped 10% on the news.
- The US technology industry and TSMC are significantly co-dependent with 60% of TSMC sales coming from the US.
- TSMC said it would spend \$12B to create a chip fab in Arizona. The factory would focus on TSMC's 5-nanometer process and start production in 2024.
- However, TSMC's technological edge will remain in Taiwan for the forthcoming 3-nanometer process that could start production in Taiwan in 2022.



Taiwan's TSMC remains dominant in R&D expenditure and semiconductor manufacturing

- ▶ TSMC's R&D expenses match SMIC's revenues. TSMC is the only fabricator with 5nm manufacturing process (N5) and it is now working on 3nm (N3) for 2x more power efficiency and 33% more performance than N7. In response, SMIC said it will increase capital expenditure to \$6.7B in 2020 (up from its original target of \$3.1B).



	N7 vs. N10	N7+ vs. N7	N5 vs. N7	N3 vs. N5	N3 vs. N7
Power	-40%	-15%	-30%	-30%	-50%
Performance	NA	+10%	+15%	+15%	+33%
Logic area reduction	-37%	-17%	-45%	-42%	-68%
Volume manufacturing date		Q2 2019	Q2 2020	H2 2022	

Chinese government sets up an additional \$29B state-backed fund reduce its dependency on American semiconductor technology

▶ China-listed chipmakers see their public market valuations soar in 2020. Cambricon goes public raising \$370M.

- China is the world's largest importer of semiconductors, totalling \$200B/year.
- New state fund is backed by the Ministry of Finance, China Development Bank, local government and state-owned enterprises. It follows the first state-led semis fund that was launched in 2014.
- SMIC, listed on the HK exchange since 2004, opportunistically listed on the Shanghai Stock Exchange's STAR board. Its shares jumped 202% on debut.



KOFIN

China hires over 100 TSMC engineers in push to close gap in semiconductor capabilities

▶ TSMC employees are offered as much as 2.5x their annual salary and bonuses to leave. Overall, Taiwan has lost 3,000 semiconductor engineers in recent times (circa 10% of their national supply).

- Government-backed Quancin Integrated Circuit Manufacturing (QXIC, founded 2019) and Wuhan Hongxin Semiconductor Manufacturing Co. (HSMC, founded 2017) are led by ex-TSMC executives and have each hired 50 former TSMC employees.
- QXIC recently began operating a R&D facility close to TSMC's 5-nanometer plant in south Taiwan.
- QXIC plans to build a \$18.4B project to produce 14-nanometer chips by 2022, and also has a tech roadmap to develop even more advanced 7-nanometer chips.




US Senate proposes the CHIPS for America Act

▶ Although over half the world's advanced chips are designed in America, only 12% are manufactured there.

- The bipartisan bill seeks to boost US competitiveness.
- The CHIPS for America Act would earmark \$22B to subsidise US manufacturing of chips.
- Programs include \$10B of federal match funding, DoD funding for related funding and \$12B in related R&D funding.
- The US has also asked Intel and Samsung to produce more US manufactured chips.
- The US has now sanctioned China's SMIC because exports posed an "unacceptable risk" of being diverted to "military end use".

Chip production swings to Asia

Number of 300mm semiconductor fabs



Given mounting concerns over chips, cross border M&A remains highly politicised

► The vast majority of acquisitions have been blocked.

- ❌ December 2016: **US** and **Germany** block \$723M bid by China's Fujian Grand Chip Investment Fund (**China**) for Aixtron.
- ❌ September 2017: **US** blocks \$1.3B bid by Canyon Bridge Capital Partner (**China**) for Lattice.
- ❌ March 2018: **US** blocks \$117B bid by Broadcom (previously headquartered in **Singapore**) for Qualcomm (**USA**).
- ❌ July 2018: **China** blocks Qualcomm's \$44B bid for NXP (**Netherlands**).
- ❌ April 2020: **UK** and **US** effectively block a complete takeover of Imagination Technologies (**UK**) by Canyon Bridge (China).
- ✅ April 2020: **China** allows Nvidia's (USA) \$6.9B acquisition of Mellanox (**Israel**).
- ❓ July 2020: Siemen's (**Germany**) makes bid for Avatar (**USA**).
- ❓ The reported potential acquisition of Arm (**UK**) by Nvidia (**USA**) will be a major test of where things stand.

AI Nationalism: Governments increasingly plan to scrutinise acquisitions of AI companies

▶ **The State of AI Report and AI Nationalism essay predicted that political leaders would start to question Whether acquisitions of key AI startups should be blocked. New legislation suggests this is now happening.**

- **Germany** passed a law in June 2020 to allow the government to review or block investments or takeovers by non-EU based companies of robotics, AI and semiconductors companies. The foreign ownership threshold at which government could review or veto was decreased from 25% to 10%. This was likely influenced by the **EU's** 2019 directive on screening foreign direct investment.
- **Japan** passed a law, effective in Aug 2019 that requires foreign investors to report to the Japanese government and undergo inspection in case they buy 10% or more of stocks in listed Japanese companies or acquire shares of unlisted firms.
- In June 2020, the **UK** expanded its powers to intervene in mergers on public interest grounds under the Enterprise Act 2002. The Government is now able to scrutinise mergers and acquisitions involving AI companies where the target company has revenues of over £1M.



The likely sale of Arm to NVIDIA is questioned by many, including its founder

► **Hermann Hauser, a leading founder and investor, argues it would be bad for the UK if Arm is acquired by NVIDIA.**

- Hauser argues that it would lead to significant job losses in the UK, similar to when NVIDIA acquired Bristol's Icera in 2011. He also argues it would destroy Arm's neutrality and impact European licensees who compete with NVIDIA.
- He argues that Arm would become subject to oversight by CFIUS which would allow the US to block specific companies from using Arm and that Arm would effectively become part of the US trade arsenal. He also notes that *'it will break up Arm into a US ARM and a Chinese Arm. There will be two types and fuel the trade war between America and China. It will be sad for a British company to be torn between them.'*
- Hauser's intervention is notable because he has been involved in Arm from inception (helping to spin it out from Acorn in 1990). Until recently European VCs have mostly been enthusiastic about their startups being acquired by US companies, rather than considering issues of technological sovereignty.
- The UK's opposition party asked the government to assure that in any acquisition there were legally binding assurances that Arm would remain headquartered in the UK.




AI Nationalism in the US: AI budgets continue to expand

▶ AI continues to be emphasized as the most important investment area in science and technology.

- In February 2019, President Trump signed Executive Order 13859 *“Maintaining American Leadership in Artificial Intelligence”*.
- The proposed spend for 2021 is \$1.5B. These non-military investments include the Departments of Agriculture, Energy and Health.
- The Department of Defense’s Joint Artificial Intelligence Center has continued to expand from a launch budget in 2019 of \$93M to \$238M for 2020.

Federal budget for non-defense
AI R&D



AI Nationalism in the US: A major new bi-partisan act is proposed

► The proposed 'Endless Frontier' act explicitly frames AI as a race between superpowers.

- The bi-partisan act explicitly frames a race to lead in AI: *"The country that wins the race in key technologies—such as artificial intelligence, quantum computing, advanced communications, and advanced manufacturing—will be the superpower of the future."*
- The act would create a Technology Directorate within the National Science Foundation and enable it to operate like DARPA with funding of \$100B over 5 years.
- It would also provide \$10B to establish a series of regional technology hubs.
- It is worth noting that many senate bills like this do not pass.



AI Nationalism in China: Decentralising policy experimentation to cities

▶ China moves to create “*national new generation AI innovation and development pilot zones*”.

- The PRC Ministry of Science and Technology created processes for cities to establish themselves as AI pilot zones. Twenty AI pilot zones are targeted by 2023.
- This seems intended to enable more decentralised experimentation. Cities that become AI pilot zones are encouraged to accelerate the application of AI in a wide variety of fields, ranging from manufacturing to caring for elderly and disabled.
- AI pilot zones are encouraged to “*carry out AI based policy experiments*” and “*carry out AI based social experiments*”.
- Deqing County is cited as an example. The city is expected to focus on autonomous driving and smart farming.



AI Nationalism in the UK: China hawks in the UK become more active

► Pressure on the UK to choose between the US and China.

- A new group of Conservative UK parliamentarians have formed the China Research Group (CRG) to scrutinise the UK's relationship with China with an emphasis on emerging technology. The group is explicitly modelled on the European Research Group which for many years lobbied for the UK to leave the EU.
- Pressure from the US and the CRG led to the UK reversing its policy on Huawei against the recommendations of the UK's chief executive of the UK's National Cyber Security Centre (who had argued that risks from Huawei could be carefully managed).
- SenseTime, heavily criticised for its role in human rights abuses in Xinjiang (see State of AI Report 2019) is no longer setting up a UK HQ.



Another wave of countries declare national AI strategies




Evidence suggests that the US tax code incentivises replacing humans with robots

► Acemoglu, Manera and Restrepo's paper demonstrates that tax reforms from 2000 to 2017 have caused the gap between effective tax rates on labour and robots to dramatically widen.

- The authors argue that this incentivises levels of automation that are not socially desirable - displacing workers without achieving meaningful productivity gains.
- They argue that reducing labour tax rates and introducing some kind of automation tax could meaningfully increase employment whilst continuing to increase productivity.

Effective tax rates for U.S. companies, by type of expenditure





Jobs at risk of automation in the EU 19 countries

▶ Executives from 1,872 enterprises worldwide report the largest AI-induced workforce contraction in automotive and assembly and telecoms in the last year. Looking forward, the CPG, transport, utilities, retail and financial services are expected to follow.

Change in workforce due to AI adoption, % of respondents

■ Decrease >10% ■ Decrease 3–10% ■ Increase 3–10% ■ Increase >10%



Bengio, Hassabis, Ng and other AI research leaders unite at NeurIPS 2019 in a call to action for climate change

▶ A position paper and workshop explored various high leverage problems where ML methods can be applied.

- Automatic monitoring with remote sensing (e.g. deforestation, climate disasters).
- Scientific discovery (e.g. new battery materials, carbon capture).
- Optimize systems (e.g. reducing food waste, consolidating freight).
- Accelerate physical simulations (e.g. climate models and energy scheduling).
- The authors note that “ML is part of the solution: it is a tool that enables other tools across fields.”

Tackling Climate Change with Machine Learning

David Rolnick^{1*}, Priya L. Donti², Lynn H. Kaack³, Kelly Kochanski⁴, Alexandre Lacoste⁵, Kris Sankaran^{6,7}, Andrew Slavin Ross⁹, Nikola Milojevic-Dupont^{10,11}, Natasha Jaques¹², Anna Waldman-Brown¹², Alexandra Luccioni^{6,7}, Tegan Maharaj^{6,8}, Evan D. Sherwin², S. Karthik Mikkavilli^{6,7}, Konrad P. Kording¹, Carla Gomes¹³, Andrew Y. Ng¹⁴, Demis Hassabis¹⁵, John C. Platt¹⁶, Felix Creutzig^{10,11}, Jennifer Chayes¹⁷, Yoshua Bengio^{6,7}

¹University of Pennsylvania, ²Carnegie Mellon University, ³ETH Zürich, ⁴University of Colorado Boulder,

⁵Element AI, ⁶Mila, ⁷Université de Montréal, ⁸École Polytechnique de Montréal, ⁹Harvard University,

¹⁰Mercator Research Institute on Global Commons and Climate Change, ¹¹Technische Universität Berlin,

¹²Massachusetts Institute of Technology, ¹³Cornell University, ¹⁴Stanford University,

¹⁵DeepMind, ¹⁶Google AI, ¹⁷Microsoft Research

Section 5: Predictions

8 predictions for the next 12 months

- ▶ 1. The race to build larger language models continues and we see the first 10 trillion parameter model.
- ▶ 2. Attention-based neural networks move from NLP to computer vision in achieving state of the art results.
- ▶ 3. A major corporate AI lab shuts down as its parent company changes strategy.
- ▶ 4. In response to US DoD activity and investment in US based military AI startups, a wave of Chinese and European defense-focused AI startups collectively raise over \$100M in the next 12 months.
- ▶ 5. One of the leading AI-first drug discovery startups (e.g. Recursion, Exscientia) either IPOs or is acquired for over \$1B.
- ▶ 6. DeepMind makes a major breakthrough in structural biology and drug discovery beyond AlphaFold.
- ▶ 7. Facebook makes a major breakthrough in augmented and virtual reality with 3D computer vision.
- ▶ 8. NVIDIA does not end up completing its acquisition of Arm.

Section 6: Conclusion

Thanks!

Congratulations on making it to the end of the State of AI Report 2020! Thanks for reading.

In this report, we set out to capture a snapshot of the exponential progress in the field of machine learning, with a focus on developments since last year's issue that was published on 26th June 2019. We believe that AI will be a force multiplier on technological progress in our world, and that wider understanding of the field is critical if we are to navigate such a huge transition.

We set out to compile a snapshot of all the things that caught our attention in the last year across the range of AI research, talent, industry and the emerging politics of AI.

We would appreciate any and all feedback on how we could improve this Report further, as well as contribution suggestions for next year's edition.

Thanks again for reading!

Nathan Benaich (@nathanbenaich) and **Ian Hogarth** (@soundboy)

Conflicts of interest

The authors declare a number of conflicts of interest as a result of being investors and/or advisors, personally or via funds, in a number of private and public companies whose work is cited in this report.

Ian is an angel investor in: Chorus.ai, ComplyAdvantage, Disperse, Faculty, LabGenius, and PostEra.

Nathan and Air Street Capital are shareholders of: Graphcore, LabGenius, Niantic, ONI, PolyAI, Secondmind, Tractable, and ZOE.

About the authors



Nathan Benaich

Nathan is the general partner of **Air Street Capital**, a venture capital firm investing in AI-first technology and life science companies. He founded RAAIS and London.AI, which connect AI practitioners from large companies, startups and academia, and the RAAIS Foundation that funds open-source AI projects. He studied biology at Williams College and earned a PhD from Cambridge in cancer research.



Ian Hogarth

Ian is an **angel investor** in 60+ startups. He is a Visiting Professor at UCL working with Professor Mariana Mazzucato. Ian was co-founder and CEO of Songkick, the concert service used by 17m music fans each month. He studied engineering at Cambridge where his Masters project was a computer vision system to classify breast cancer biopsy images. He is the Chair of Phasecraft, a quantum software company.

State of AI Report

October 1, 2020