Artificial Intelligence, algorithmic decision-making and HR

Oliver Patel
Enterprise AI Governance Lead, AstraZeneca

Connect on LinkedIn

Overview

1. Introduction to AI and Machine Learning
2. AI in the Enterprise
3. AI and HR

1. Introduction to AI and Machine Learning

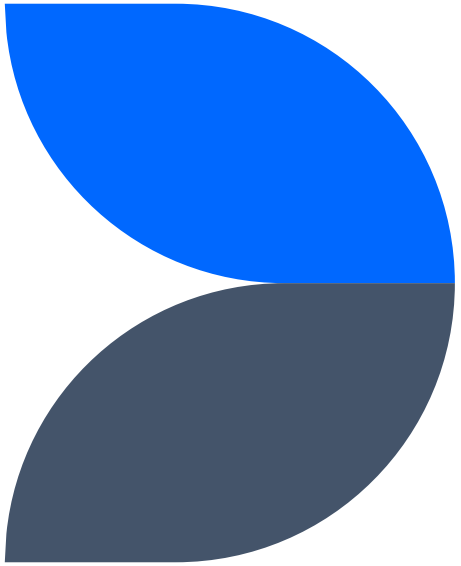

Narrow AI v General AI

Narrow AI = a machine or system which can perform a narrow task or set of tasks and functions in a given, pre-determined domain.

General AI = a machine or system which can learn any intellectual or cognitive task and outperform human intelligence across a wide variety of domains.

AI Timeline

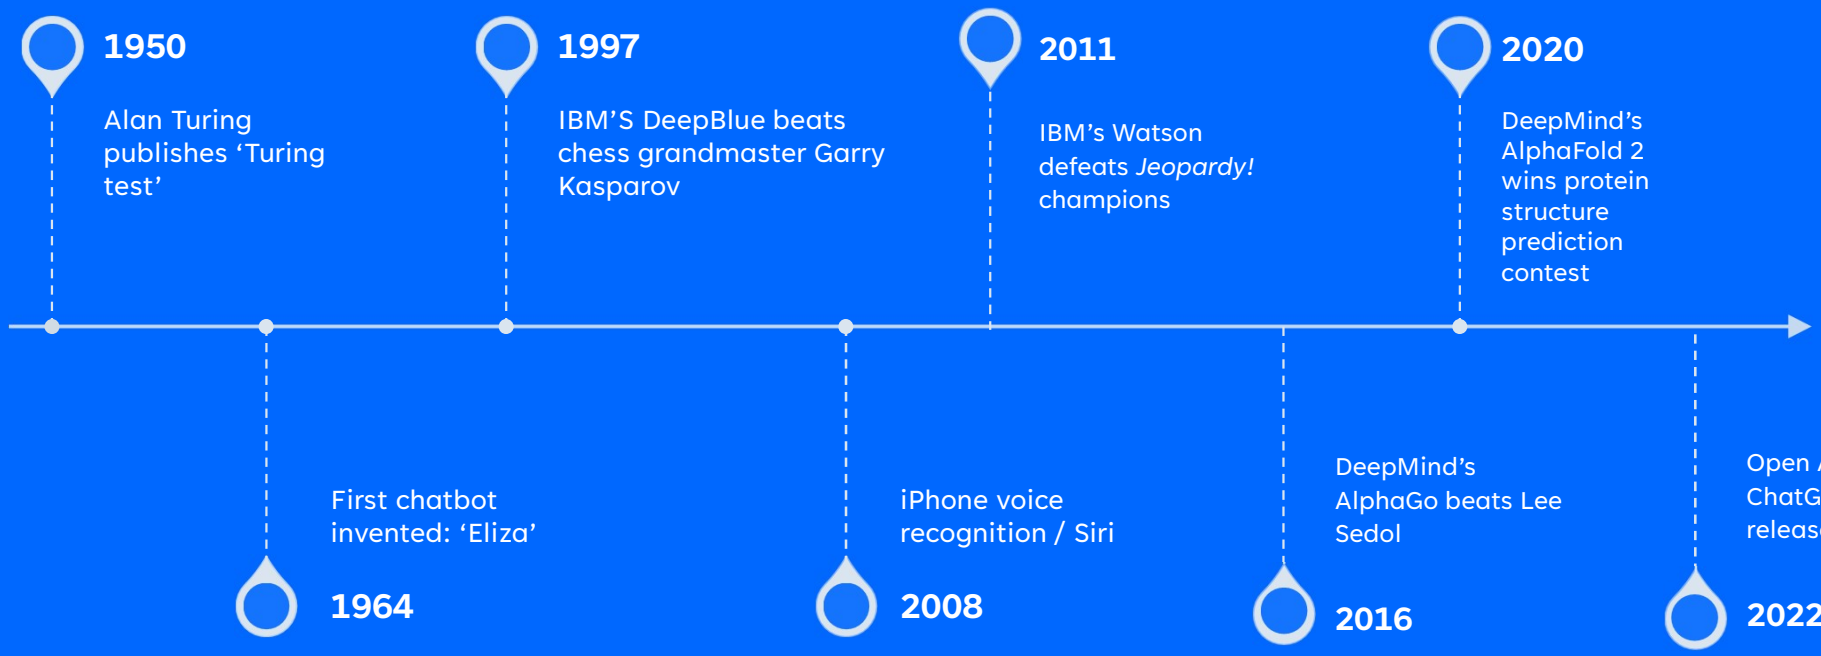
1950
Alan Turing publishes ‘Turing test’
1964
First chatbot invented: ‘Eliza’
1997
IBM’S DeepBlue beats chess grandmaster Garry Kasparov
2008
iPhone voice recognition / Siri
2011
IBM’s Watson defeats *Jeopardy!* champions
2016
DeepMind’s AlphaGo beats Lee Sedol
2020
DeepMind’s AlphaFold 2 wins protein structure prediction contest
2022
Open AI’s ChatGPT released

Explosive growth in AI innovation since 2010s

- AI winters: 1974-1980 and 1987-1993
- Explosion in the amount of data created
  - Data generation increasing exponentially
- Vast increase in computing power and data infrastructure
- Increasingly sophisticated algorithms/models and data science methodologies
  - GPT-3 has 175bn parameters and MT-NLG has 500bn

Defining AI: EU

*"A **system** that is designed to operate with elements of **autonomy** and that, based on machine and/or human-provided **data and inputs**, infers how to achieve a given set of objectives using machine learning and/or logic- and knowledge-based approaches, and **produces system-generated outputs** such as content (generative AI systems), predictions, recommendations or decisions, influencing the environments with which the AI system interacts."*

Defining AI: OECD

*"A machine-based **system** that can, for a given set of **human-defined objectives**, make predictions, recommendations, or decisions influencing real or virtual environments. AI systems are designed to operate with varying levels of **autonomy**."*

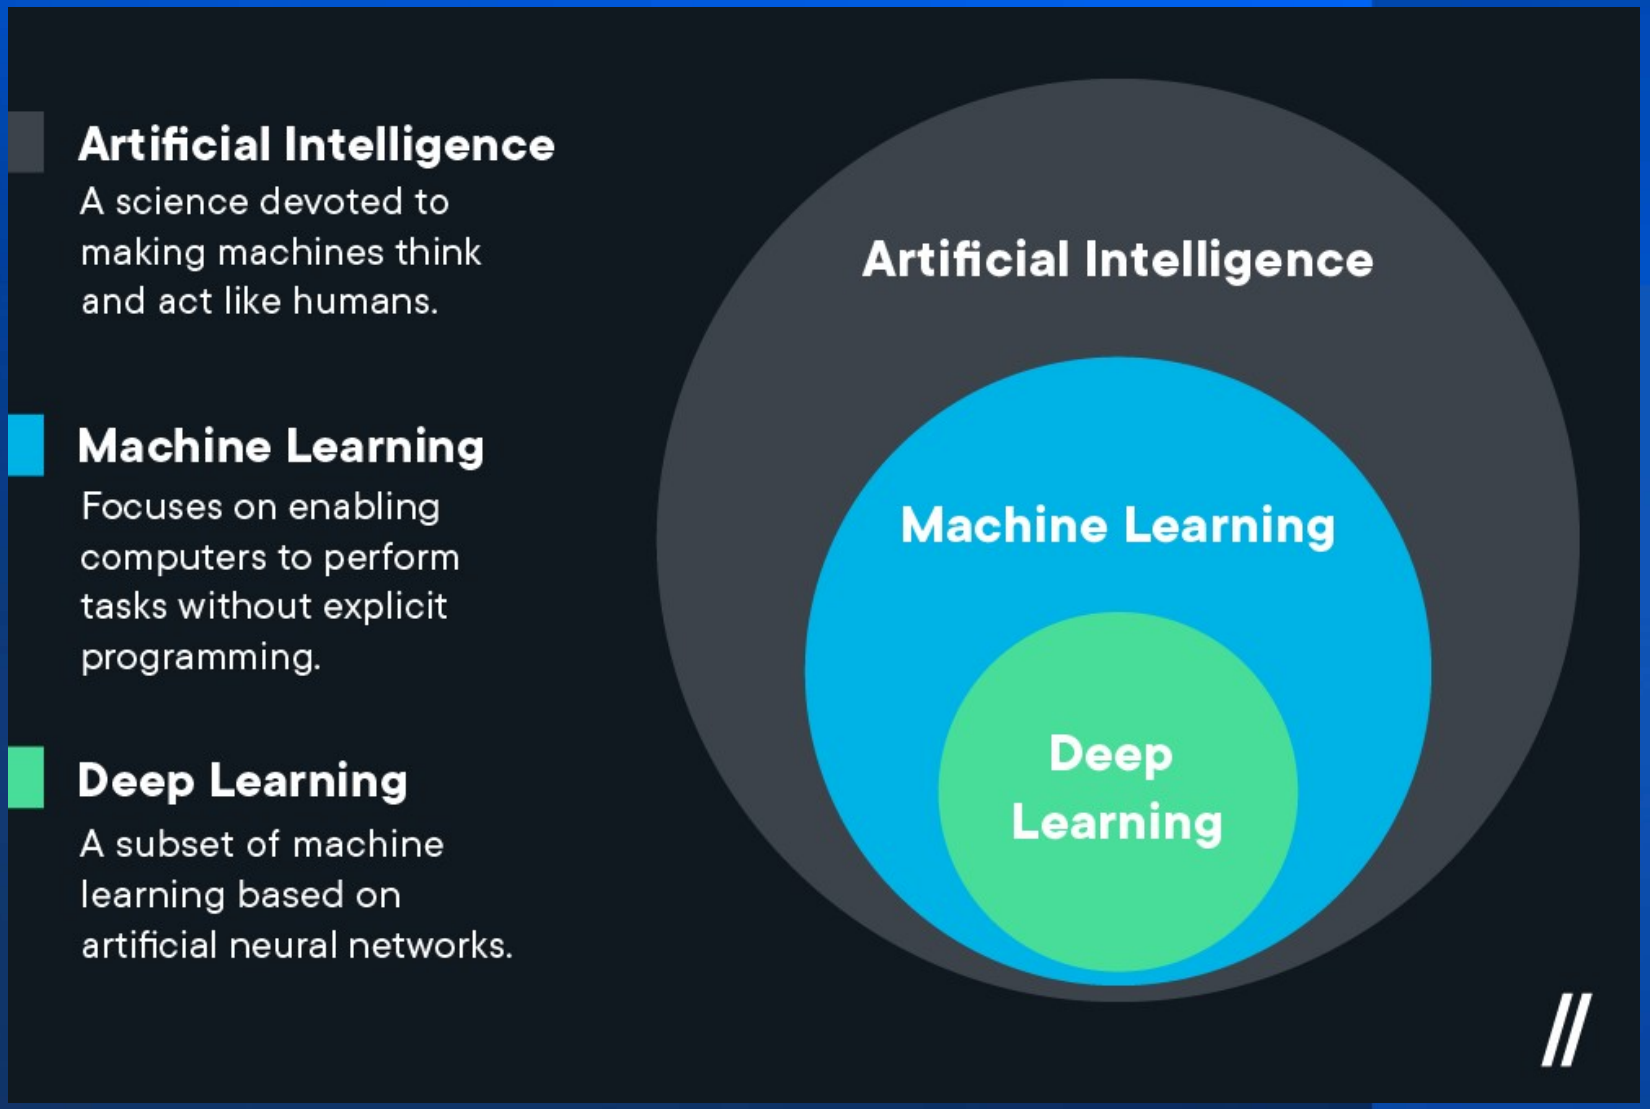

Source: Flat Iron School

Machine learning

Machine learning is an approach to learn complex patterns from existing data and use these patterns to make predictions on unseen data.

Source: Chip Huyen, Designing Machine Learning Systems

Machine learning approaches

Supervised learning: machine presented with example inputs and their desired outputs (e.g., images of cats). Learns general rule that maps inputs to outputs.

Unsupervised learning: the machine independently detects and learns structure and patterns in the input data. No matching outputs are provided to the machine.

Reinforcement learning: a computer programme interacts with a dynamic environment, in which it is provided feedback and 'rewarded' for optimal performance.

Deep learning

A machine learning technique that uses artificial neural networks to mimic the human brain’s learning process. Requires more data, less human intervention and makes more complex correlations.

The deep learning revolution

- Deep learning (DL) algorithms are more complex and sophisticated than traditional ML models.
- Requires vast amounts of training data, is more accurate, but also a ‘black box’.
- DL models consist of multiple ‘layers’ of artificial neural networks which progressively extract features and insight from raw input data.
- Notable deep learning use cases:
  - Autonomous vehicles
  - Natural language processing and large language models (e.g., ChatGPT)
  - Personalised recommender systems (e.g., Netflix)
  - Computer vision (e.g., medical imaging and diagnostics)

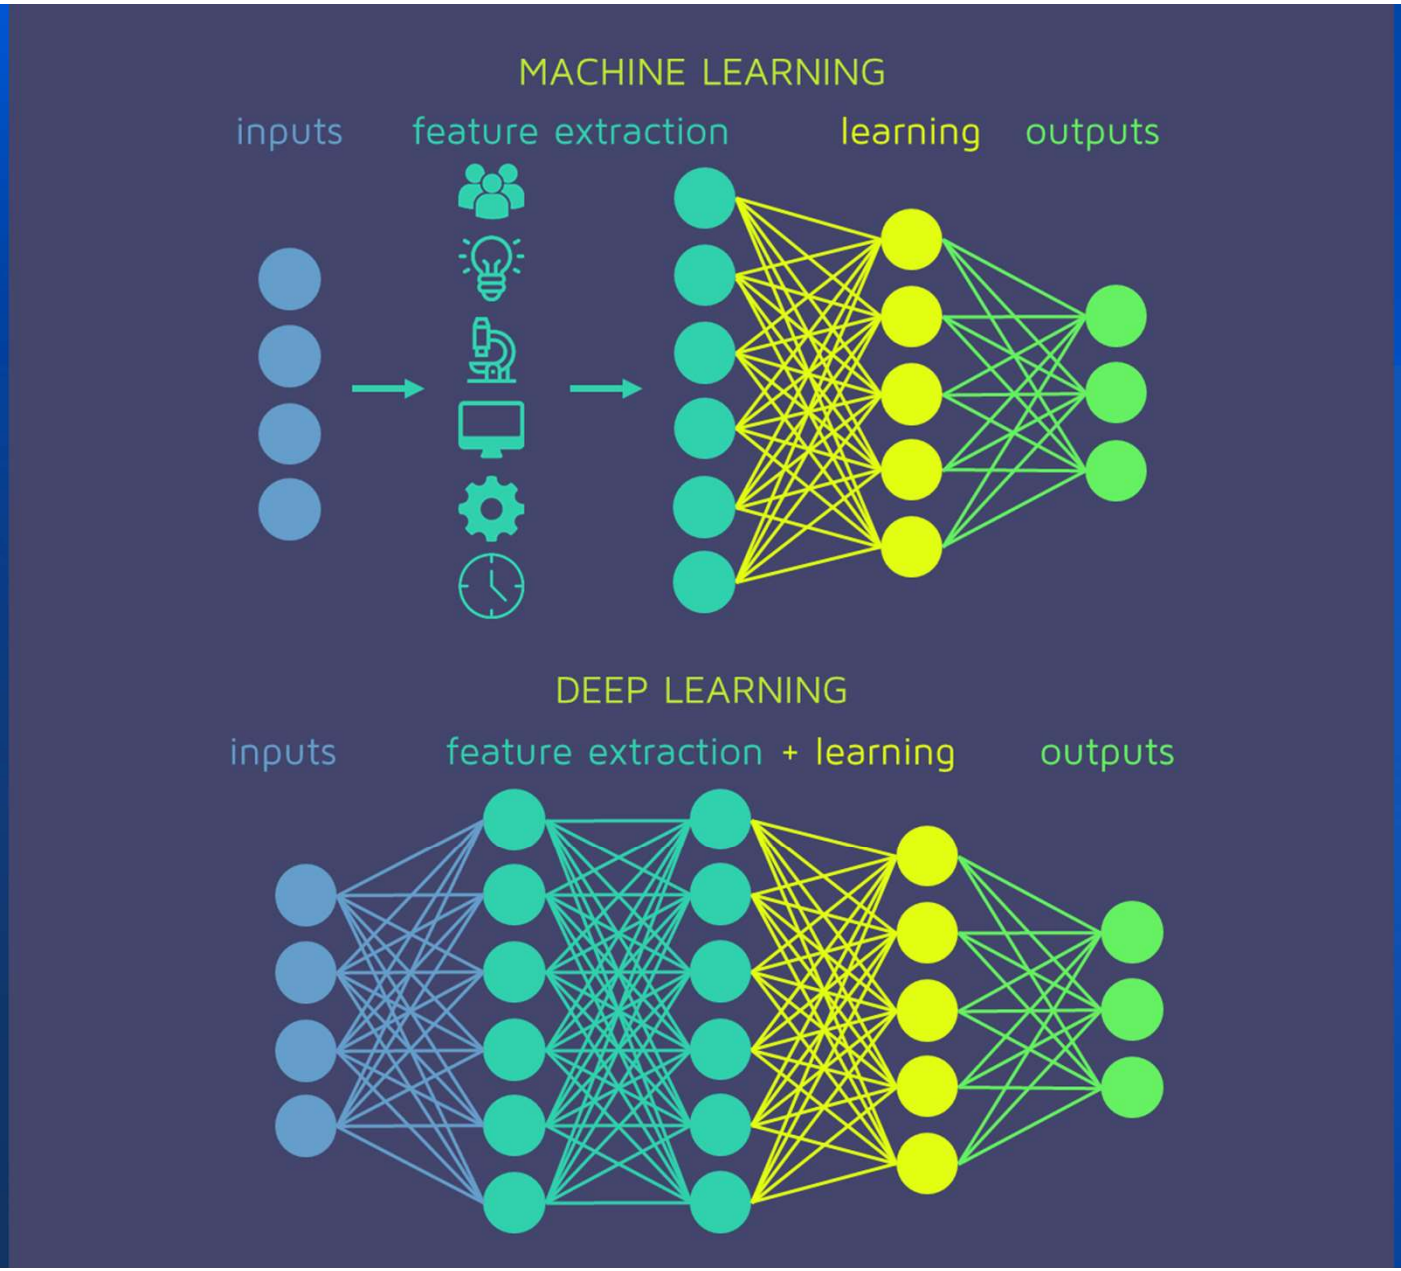

source:
Quantdare

Common AI/ML tasks

Computer vision

Object detection

Object classification

Facial recognition

Recommending and prediction

Social media and content platforms

Timeseries forecasting

Natural language processing

Machine translation

Search engines

Knowledge graphs

Mathematical optimisation

Driver routing

Digital twins

Manufacturing / supply chain

Generative AI

Image and video generation

LLMs: ChatGPT and Bard

Code and music generation

Machine Learning (ML) Systems

Source: Chip Huyen, Designing Machine Learning Systems

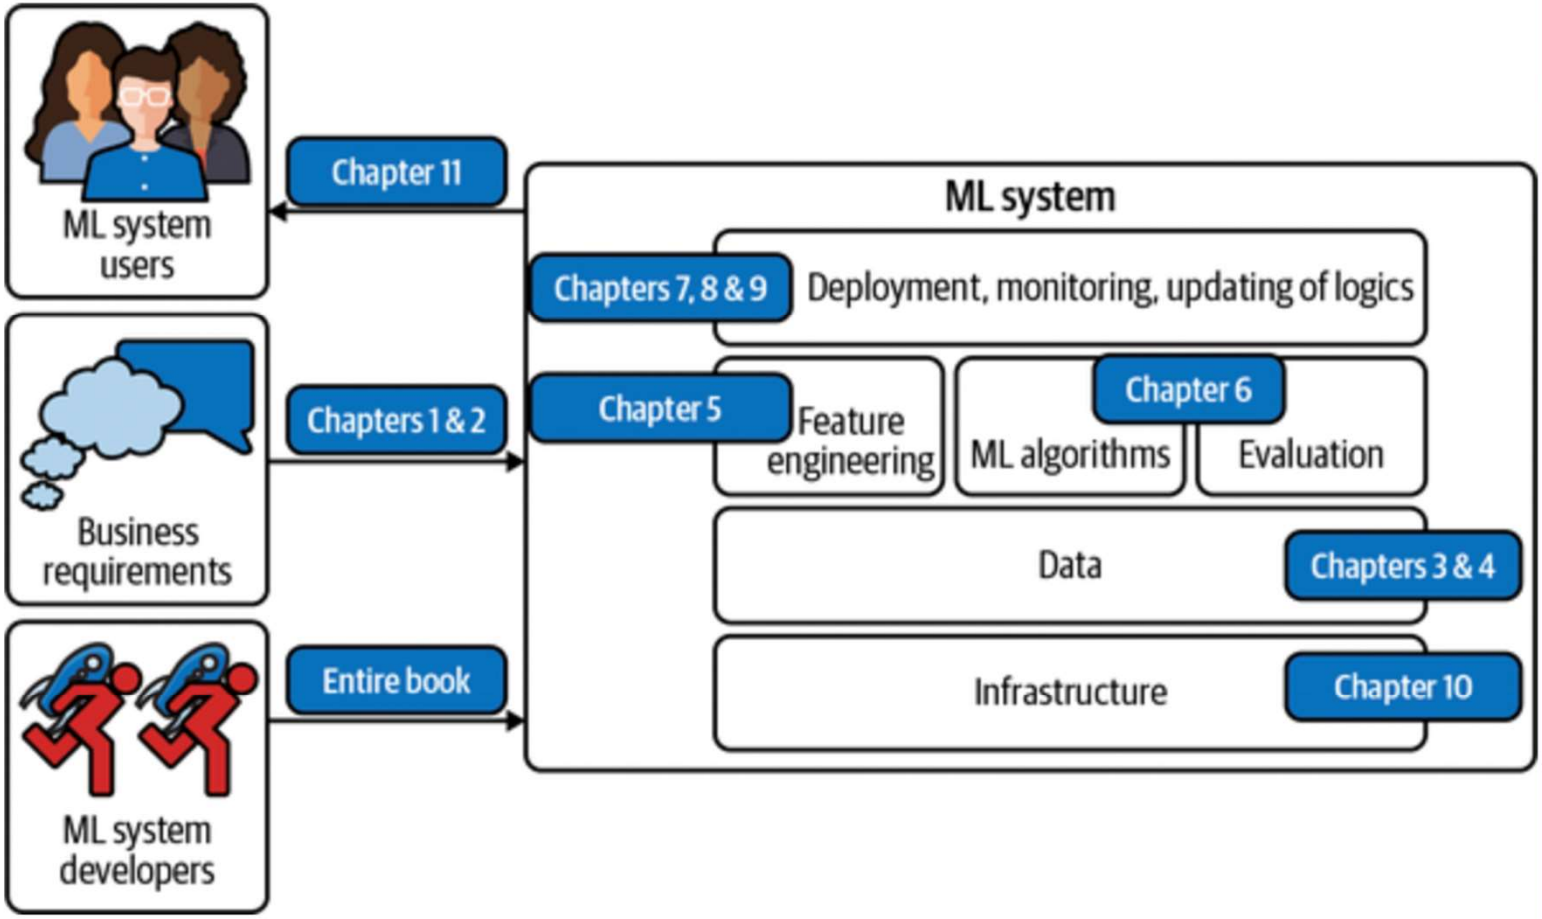

Figure 1-1. Different components of an ML system. "ML algorithms" is usually what people think of when they say machine learning, but it's only a small part of the entire system.

AI / ML development Lifecycle

Source: Alex Wang, LinkedIn

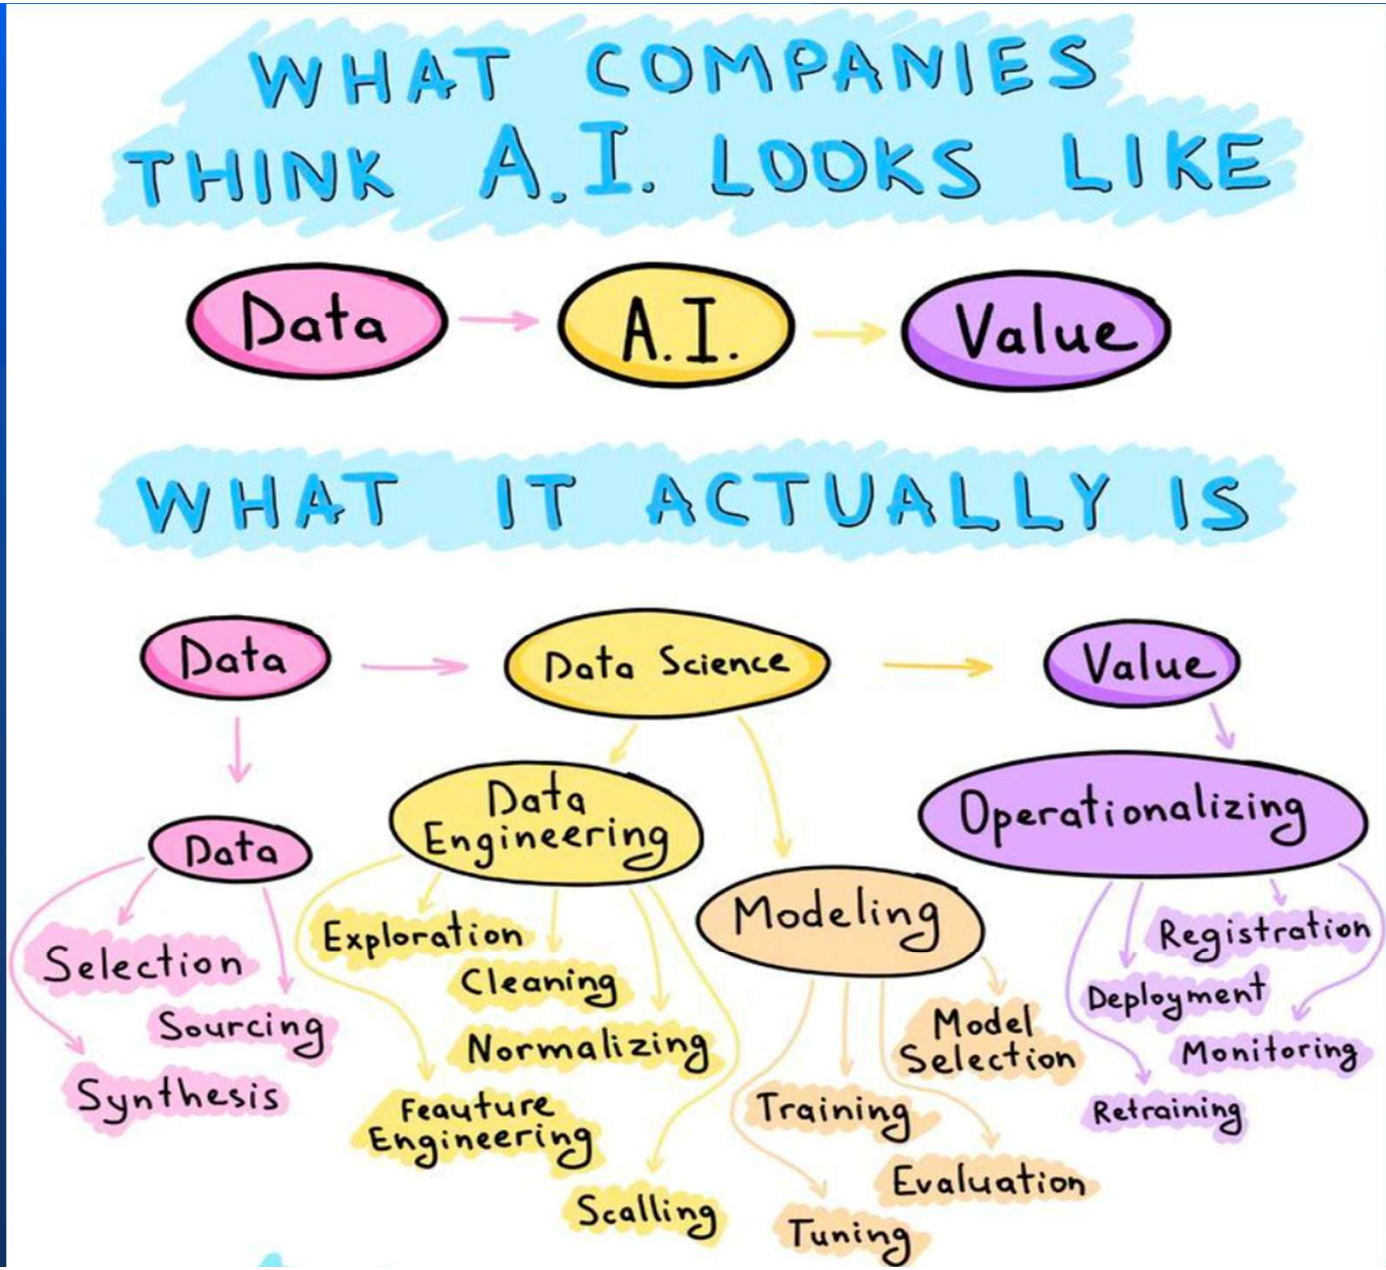

2. AI in the Enterprise

Enterprise AI adoption has soared

- Enterprise AI adoption has doubled between 2017 and 2022. (source: McKinsey)
- Private investment in AI was $93.5bn in 2021, double the amount in 2020. (source: AI Index Report)
- AI is becoming more affordable and higher performing.
  - Training times have improved by over 90% and training costs have declined by over 60%. (source: AI Index Report)

32% of enterprises are using AI for workforce utilisation and optimisation

27% of enterprises are using AI for human resources optimisation

Source: IBM

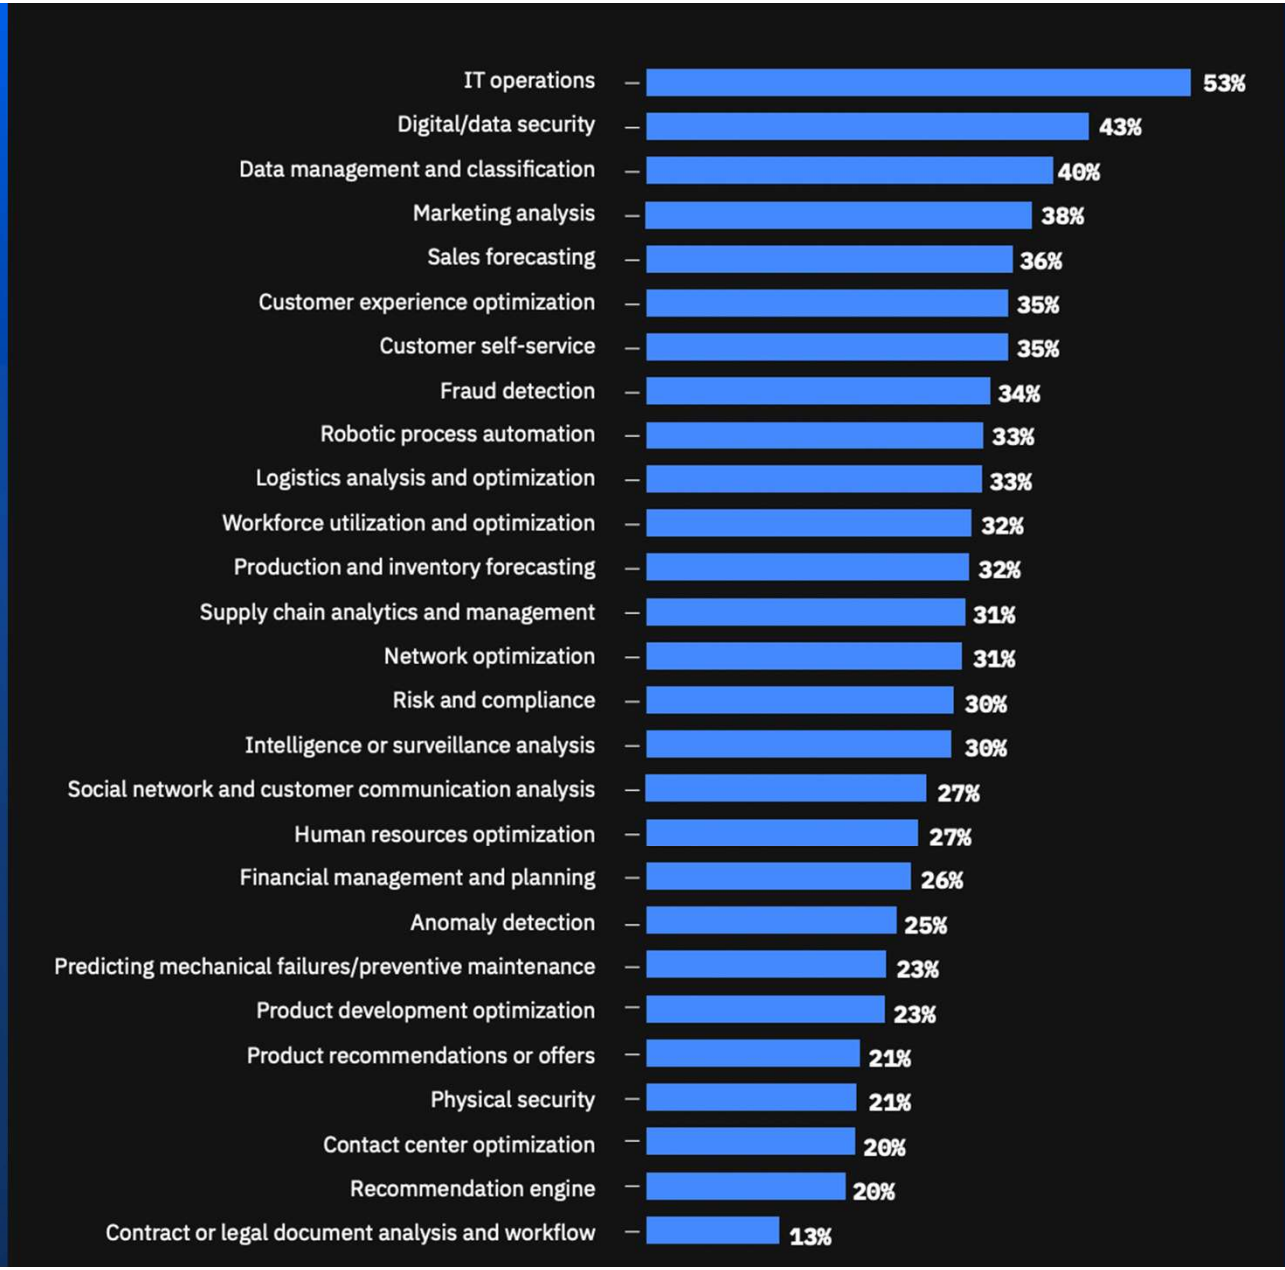

76% of enterprises expect AI to improve the employee experience

78% of enterprises expect AI to empower employees to do higher value

Source: IBM

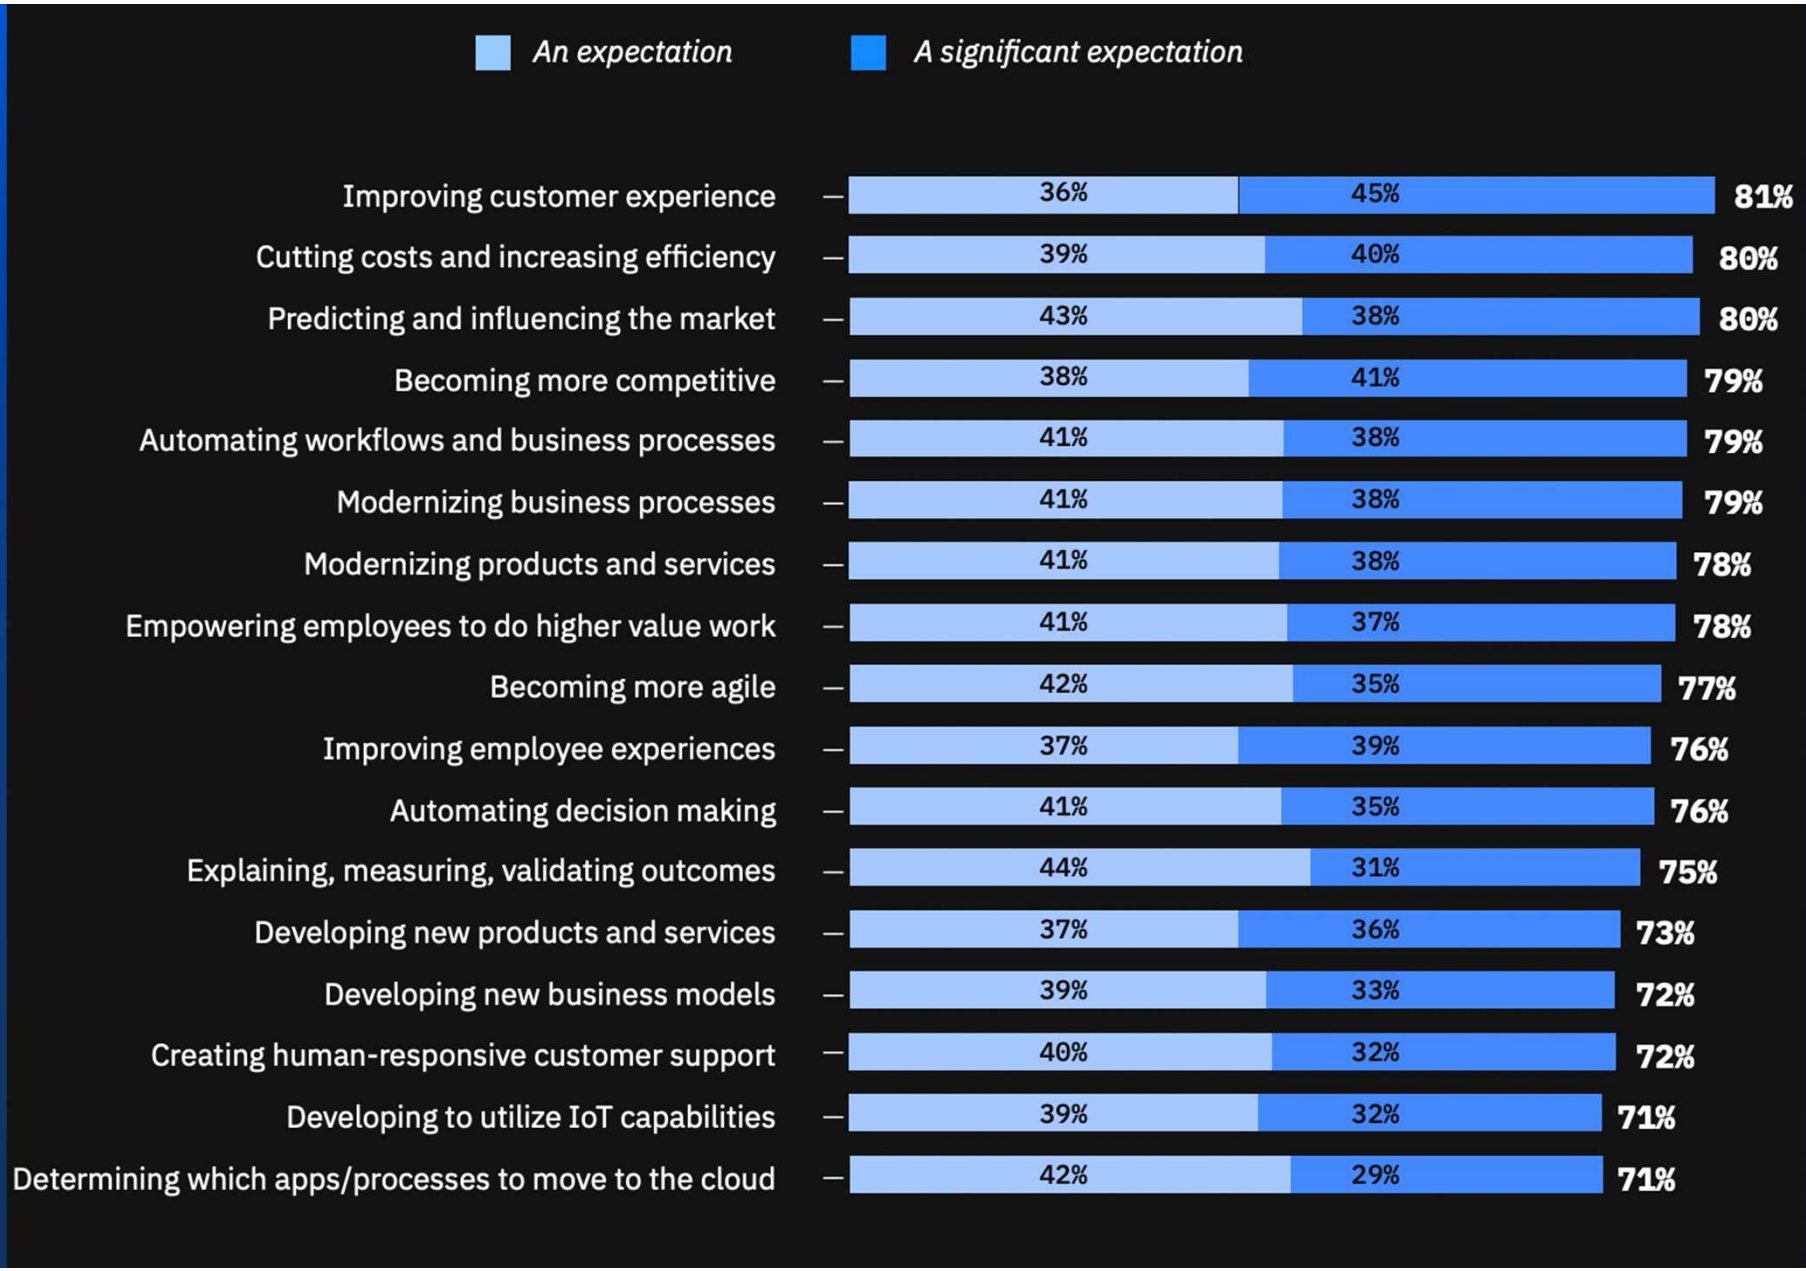

Source: McKinsey

Enterprise AI use remains in its infancy

- Most enterprises are at the beginning of their AI journey
- AI maturity will increase across all sectors and business functions
- Barriers to scaling AI use / moving beyond PoCs
  - Challenges proving business value or identifying best use cases
  - Lack of technical skills
  - Governance, risk and compliance concerns

Responsible AI

Robustness

Performance over time

Model accuracy

Model reproducibility / drift

Adversarial attacks

Fairness

Training data quality

Biased / discriminatory decision-making

Varied performance across different groups

Transparency

Model interpretability and explainability

Notification and disclosure

Technical documentation and record keeping

AI vendor documentation

Accountability

Human in the loop

Right to object / challenge decision-making

AI liability

Enterprise governance: roles and responsibilities

Privacy

Data minimisation

Automated decision-making (GDPR Article 22)

Sensitive data

Privacy engineering / by design

Security

AI system resilience

Open source AI

Third party risk

Data poisoning

System retirement

Sustainability

Carbon emissions / footprint

AI re-usability and duplication of work

AI supply chain / modern slavery

Key risks

Commercial

Poor ROI on AI spend

Lack of competitiveness

Unreliable decision-making

Legal / regulatory

Litigation

Regulatory enforcement and fines

Reputational

High profile scandals

Brand damage

Equity, diversity and inclusion

Operational

Over-reliance on AI for core activities

AI safety

Security breaches

Societal / customer

Harm to customers

Harm to the environment sustainability goals

When to not use AI

- It’s not cost effective
- Simpler technologies and approaches work better
- The ethical or legal risks are too high
- Accuracy shortcomings could be fatal / catastrophic

3. AI and Human Resources

AI in HR: use cases

Recruitment

Automated CV / applicant screening

AI video interviews / chatbots

Image / game-based assessment

Talent management

AI performance assessment / review

Compensation, promotion and dismissal

Employee engagement / retention

Personalised onboarding / communications

Employee churn prediction

Workforce / productivity management

Automated scheduling

Productivity monitoring

Algorithmic management

Fully automated worker management

Platform economy workers

Over 500+ HR tech vendors

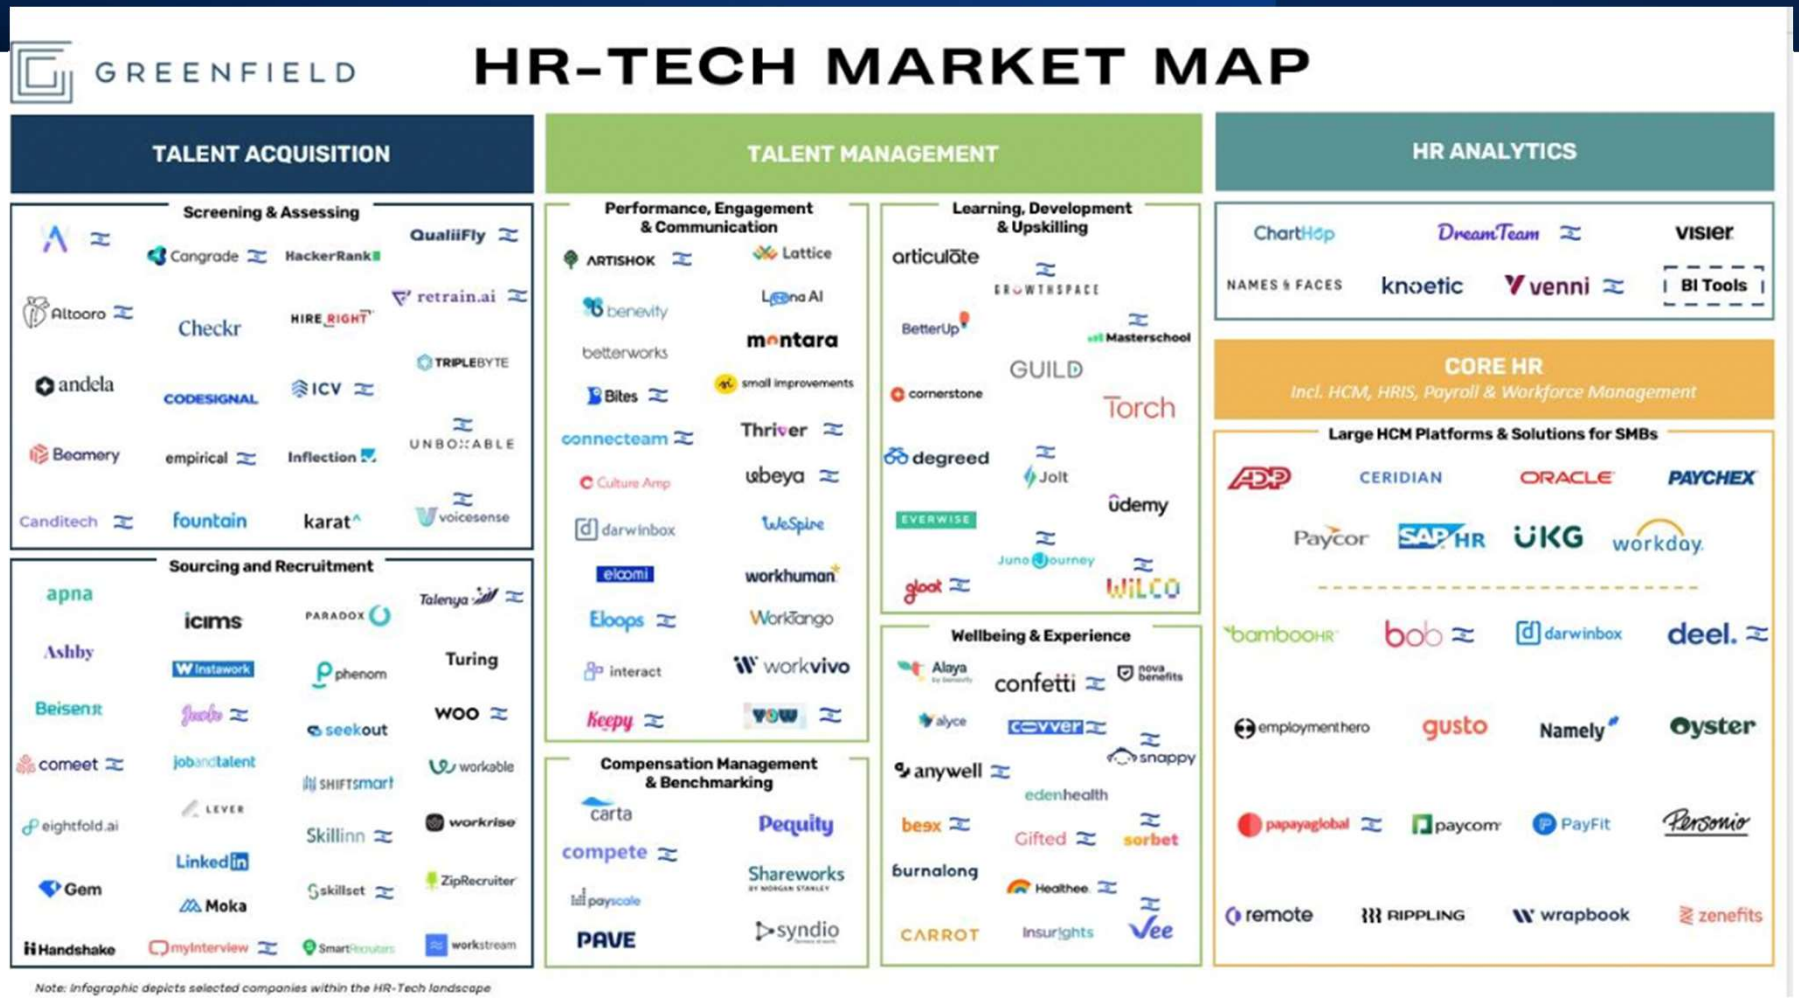

Source: Greenfield Partners

Why enterprises use AI

- Global HR functions must increasingly use AI to remain competitive
- Cost savings, increased efficiency and automation of time-consuming activities (e.g., recruitment – 99% of Fortune 500)
- More objective, data-driven and strategic decision-making
- Slicker experience and greater personalisation for candidates and employees

Key risks: AI in HR

- Lack of trust in AI-driven decision-making
- Biased and discriminatory decision-making
- Employee monitoring and surveillance: privacy concerns
- Exposure to greater legal / regulatory risks and liability

Case study: HireVue's face scanning tool

ECHNOLOGY

A face-scanning algorithm increasingly decides whether you deserve the job

HireVue claims it uses artificial intelligence to decide who's best for a job. Outside experts call it 'profoundly disturbing.'

By Drew Harwell

November 6, 2019 at 12:21 p.m. EST

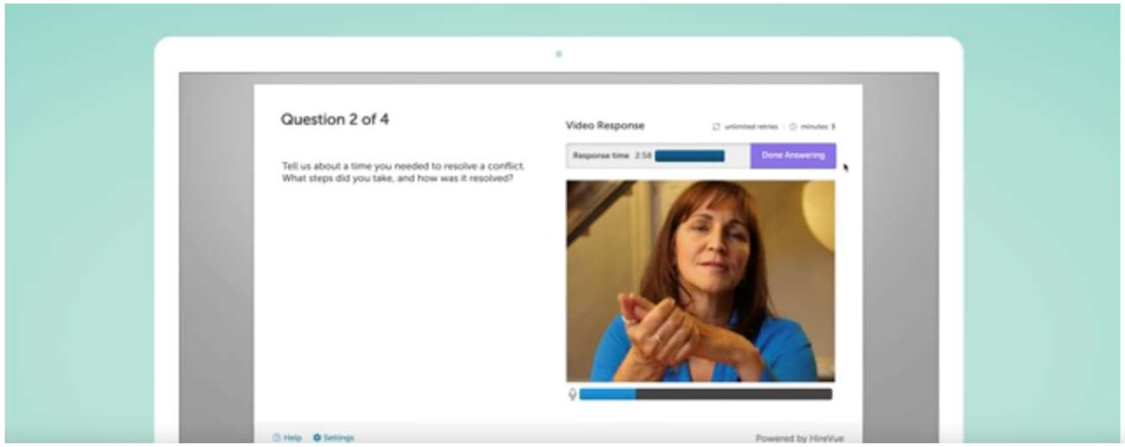

HireVue, Facing FTC Complaint From EPIC, Halts Use of Facial Recognition

January 12, 2021

Case study: Uber's fare algorithm

Working for an Algorithm

Secretive Algorithm Will Now Determine Uber Driver Pay in Many Cities

The company has long used ride time and mileage to decide driver pay but is now turning to an opaque calculation called "Upfront Fares"

By Dara Kerr

March 1, 2022 08:00 ET

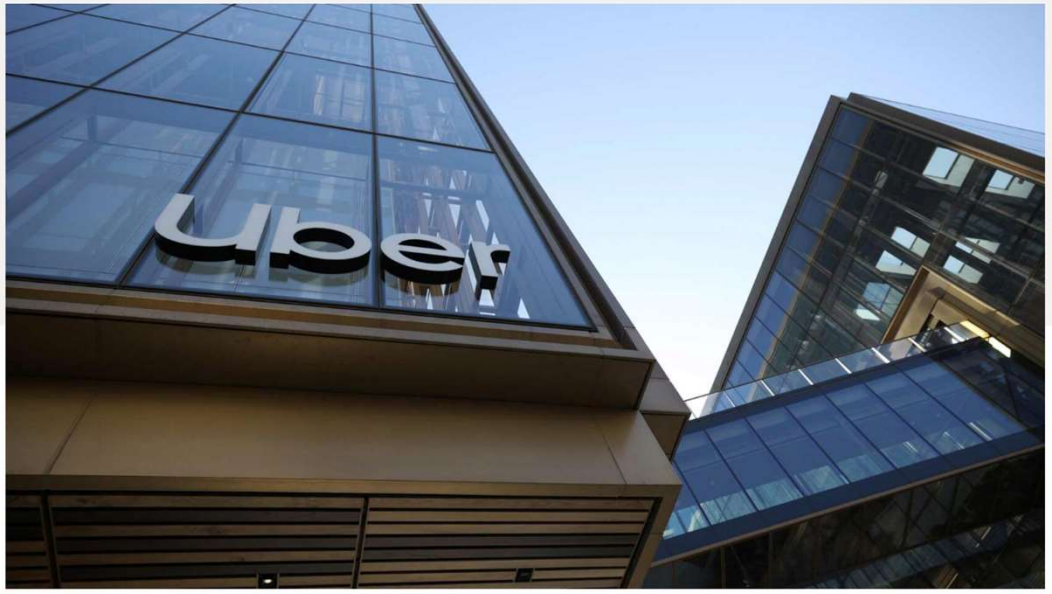

Key regulations

- EU AI Act: high-risk AI
- EU AI Liability Directive

- New York City Bias Audit Law (April 2023)

- Equal Employment Opportunity Commission: draft Strategic Enforcement Plan (2023)

- Spain’s Riders’ Law (2021): algorithmic transparency

- Canada’s AI and Data Act

- Illinois Artificial Intelligence Video Interview Act (2019)

“Machine learning algorithms don’t predict the future, but encode the past.”

Chip Huyen

Thanks for listening

Q&A

Oliver Patel
Enterprise AI Governance Lead, AstraZeneca

Connect on LinkedIn

Images from Flaticon