OpenAI | faculty

Making better decisions with Generative AI

An executive briefing

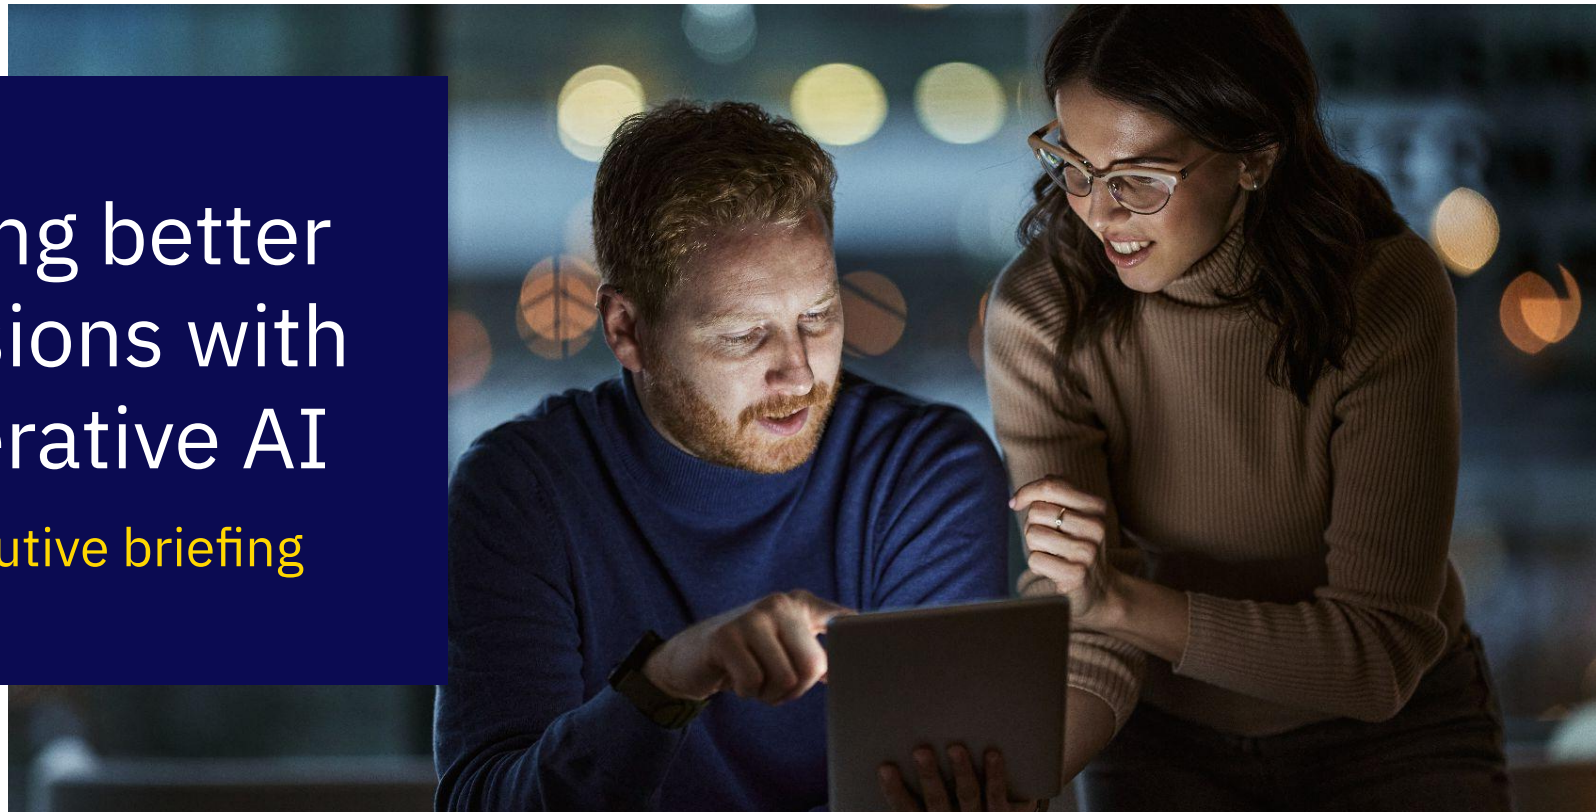

Today's agenda

01 An introduction to Faculty + OpenAI

02 Large Language Models (LLMs) - an overview

03 Putting LLMs to work in your organisation

04 How to get started

05 Q&A

OpenAI | faculty

Leading today's call

Marc Warner

Co-founder and CEO

Kat James

Technical Director
Retail & Consumer

Josh Muncke

Commercial Director
Retail & Consumer

Andrew Perry

Commercial Director
Energy Transition & Environment

Faculty design, build and implement intelligent decision systems

Faculty is Europe's leading applied AI firm, with a high density of PhD+ Data Scientists and ML Engineers

We have built and deployed over 450+ high-impact AI systems for over 250 organisations across both public and private sectors

Faculty's AI solutions underpin critical operations for healthcare providers, leading retail and consumer businesses, energy companies and governments.

Our technology is trusted to operate safely in some of the world's most demanding contexts. We are pioneers and global leaders in safe and ethical deployments of AI.

Our partnership with OpenAI helps you unlock the potential of Generative AI safely and responsibly in the real world.

AI Technology Experts

We have deep expertise in AI technology including NLP, LLMs, computer vision and transformer architecture.

Applied Experience

We are Europe's largest applied AI firm. We have designed, built and deployed AI extensively in both public and private sector.

Commitment to AI Safety

We are world leaders in AI safety. We are proud that customers trust us to mitigate risks in the real world.

OpenAI | faculty

Leverage the power of OpenAI to make better decisions - learn more

OpenAI

"We chose to work with Faculty because they understand our technology and its potential but also have the experience to ensure it is implemented successfully and safely.

Generative AI is going to change the world, and we look forward to Faculty helping us on our journey to extend the benefits of Generative AI to businesses and their customers."

Zack Kass, Chief Customer Officer, OpenAI

Large Language Models - an overview

We have reached an inflection point when it comes to the performance of AI against a range of human tasks

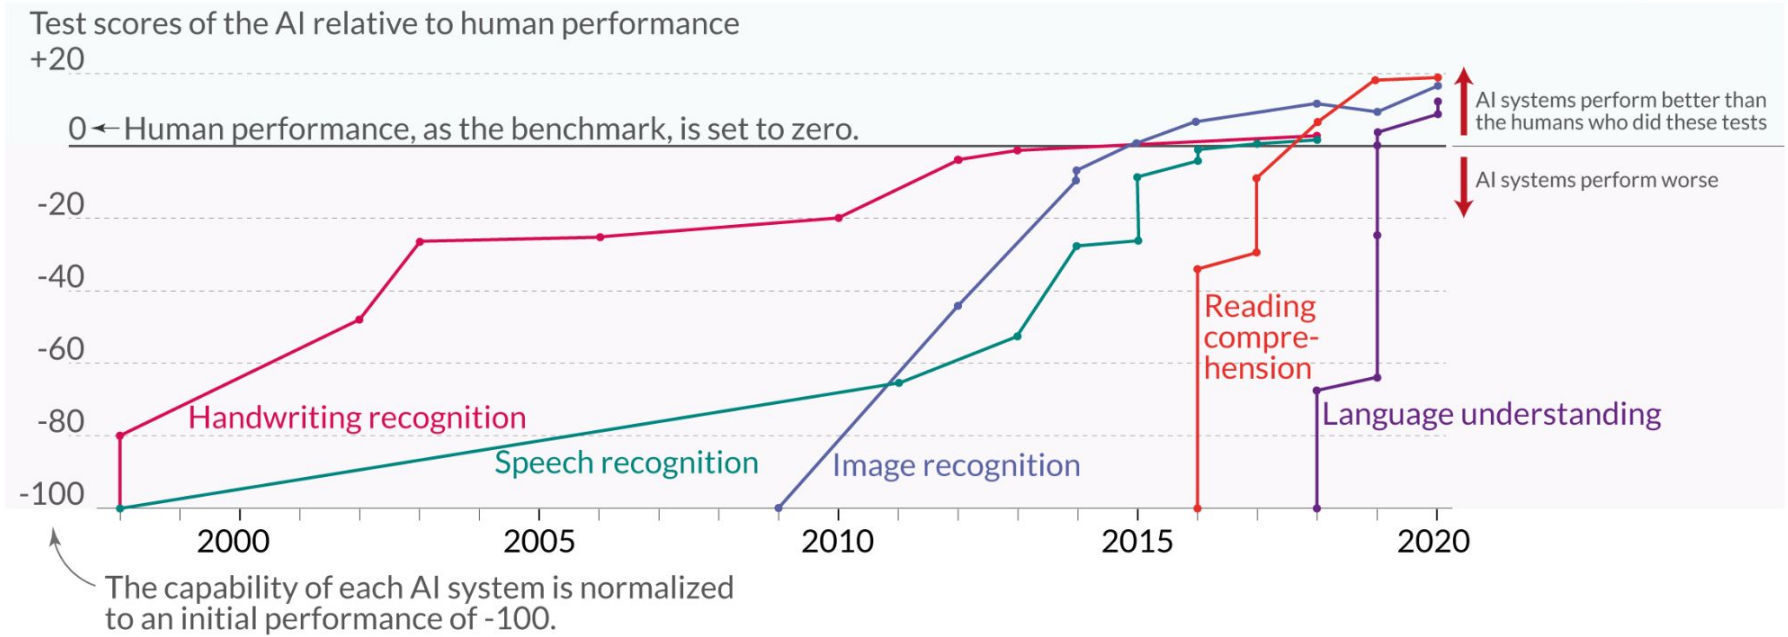

Data source: Kiela et al. (2021) – Dynabench: Rethinking Benchmarking in NLP

Generative AI is the latest advance in the field - but it represents a massive step forward in capability

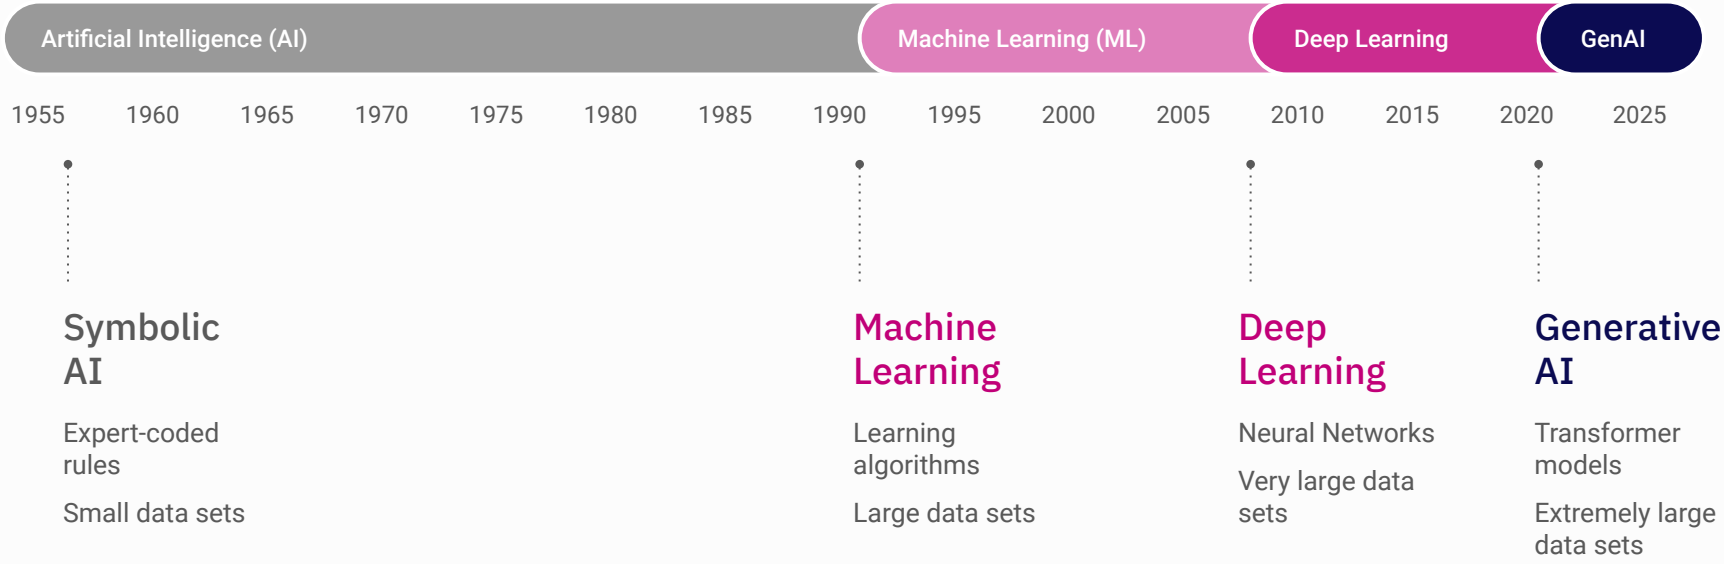

Generative AI moves us beyond classification and prediction tasks, and into creative capabilities (traditionally things machines were bad at)

Unlike “traditional” AI models - Generative AI delivers novel, generalisable outputs akin to what a human might produce

Traditional “discriminative” AI	Generative AI
Outputs are aggregations, summarisations or extrapolations of the inputs	Outputs are synthetic and entirely novel
Models with a narrow focus - built for a specific purpose and context	Large, general models that are multi-purpose and often exhibit “zero-shot” capabilities
Can often be trained and deployed / operated by in-house teams	Extremely expensive to train. Currently exclusively served via API
Input is generally data generated by some sort of process or behaviour (e.g. sales transactions)	Input is typically in the form of a text “prompt”
Depending on the model, outputs can be easy or difficult to explain	Outputs are in text/image format and are easily consumed and understood by humans

3 key breakthroughs have enabled the acceleration of Large Language Model (LLM) capability

1. Massive data

This unlocked the ability to synthesize an extremely large amount of human knowledge across multiple domains

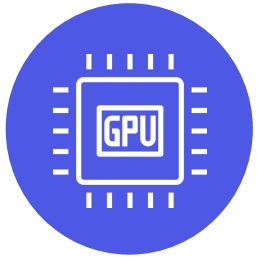

2. Massive compute

Advances in computing hardware (Moore's Law) have made parallel computations extremely fast and more cost effective than ever.

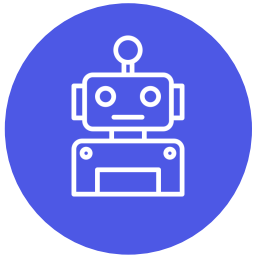

3. Algorithm advances

Transformer models have unlocked huge improvements in language models by allowing sequential operations to be trained in parallel.

How GPT-*x* is built

01

Train on massive amounts of data

Objective: base model is trained over a massive corpus of text, building a probabilistic distribution used to predict the next token

02

Fine tune with some humans in the loop

Objective: base model is fine-tuned on a reward signal from human feedback. This biases the model to the domain where human preferences have assigned a higher reward

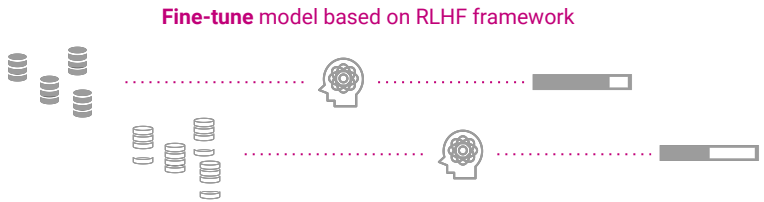

03

Add safety features

Objective: as well as including safety in RLHF, rule-based reward models (RBRMs) are applied to reduce likelihood that GPT will return harmful content

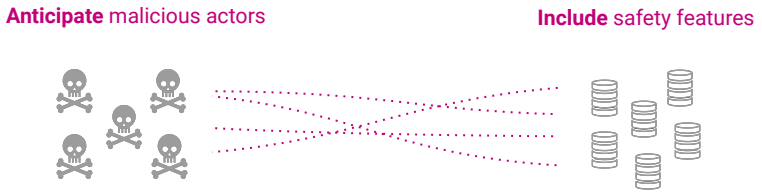

GPT-4 changes the game again and shows significant performance improvement compared with GPT-3

GPT-3.*x*	GPT-4
Work with text only	Multimodal
Cannot (probably) reason	Shows evidence of reasoning
Known problems with factual accuracy	40% more likely to produce factually correct answers than GPT-3
Can have performance issues	Shows better understanding in tasks in at least 25 languages than GPT-3 does in English
Safety concerns around the ability to generate harmful content	82% reduction in tendency to reply to requests for prohibited content

General purpose capabilities associated with Large Language Models

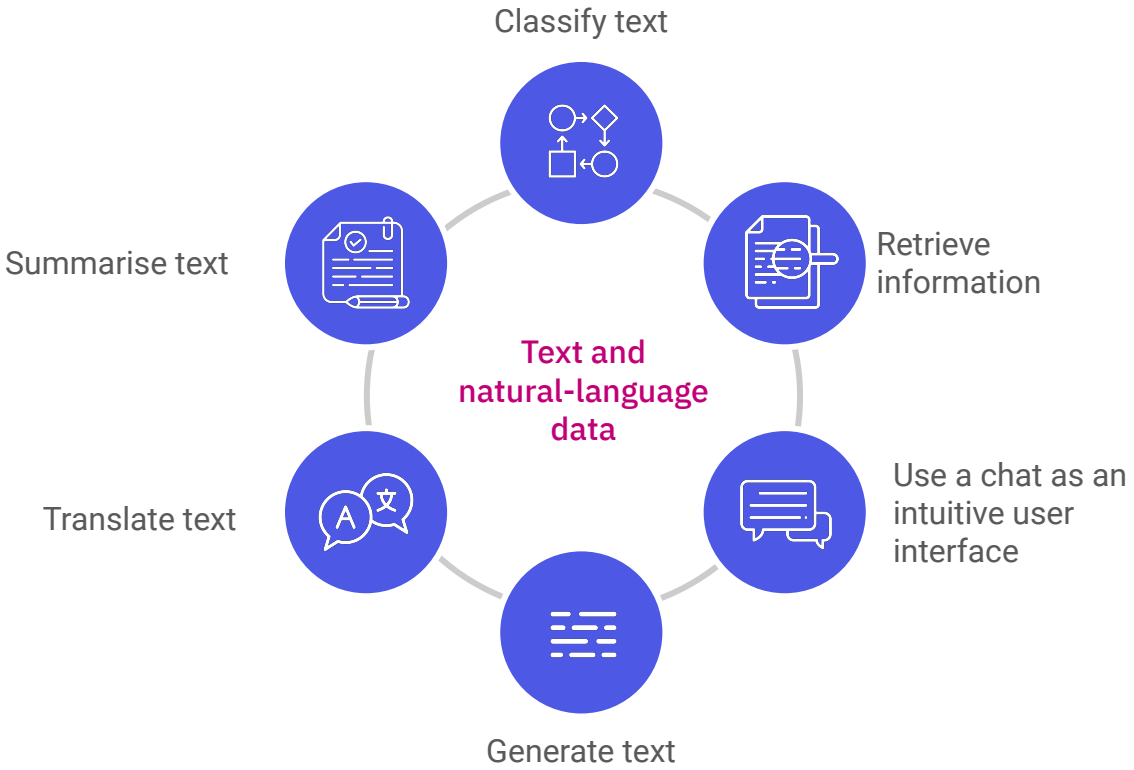

There are known risks and issues with the current generative systems that need to be taken into account

Hallucination

Frozen in time

Bias

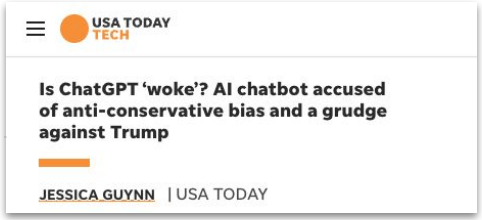

Unexpected behaviour

Legal Uncertainty

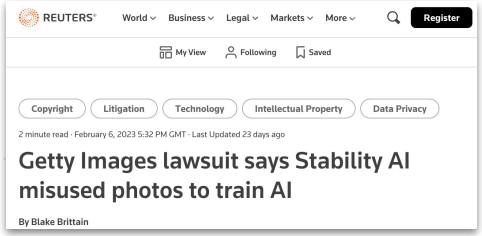

Sensitive data considerations

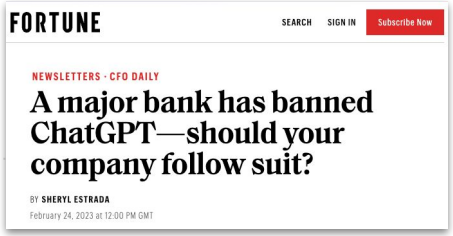

The right way to think of the models that we create is a reasoning engine, not a fact database,"

"They can act as a fact database, but that's not really what's special about them – what we want them to do is something closer to the ability to reason, not to memorize."

Sam Altman, CEO OpenAI

Putting LLMs to work

3 essential rules for using LLMs in an enterprise context

Find the right type of problem

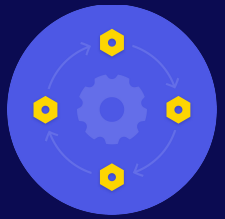

Integrate into an end-to-end system

Build in safety from the outset

Finding the right type of problem

A framework for Large Language Model use cases (illustrative)

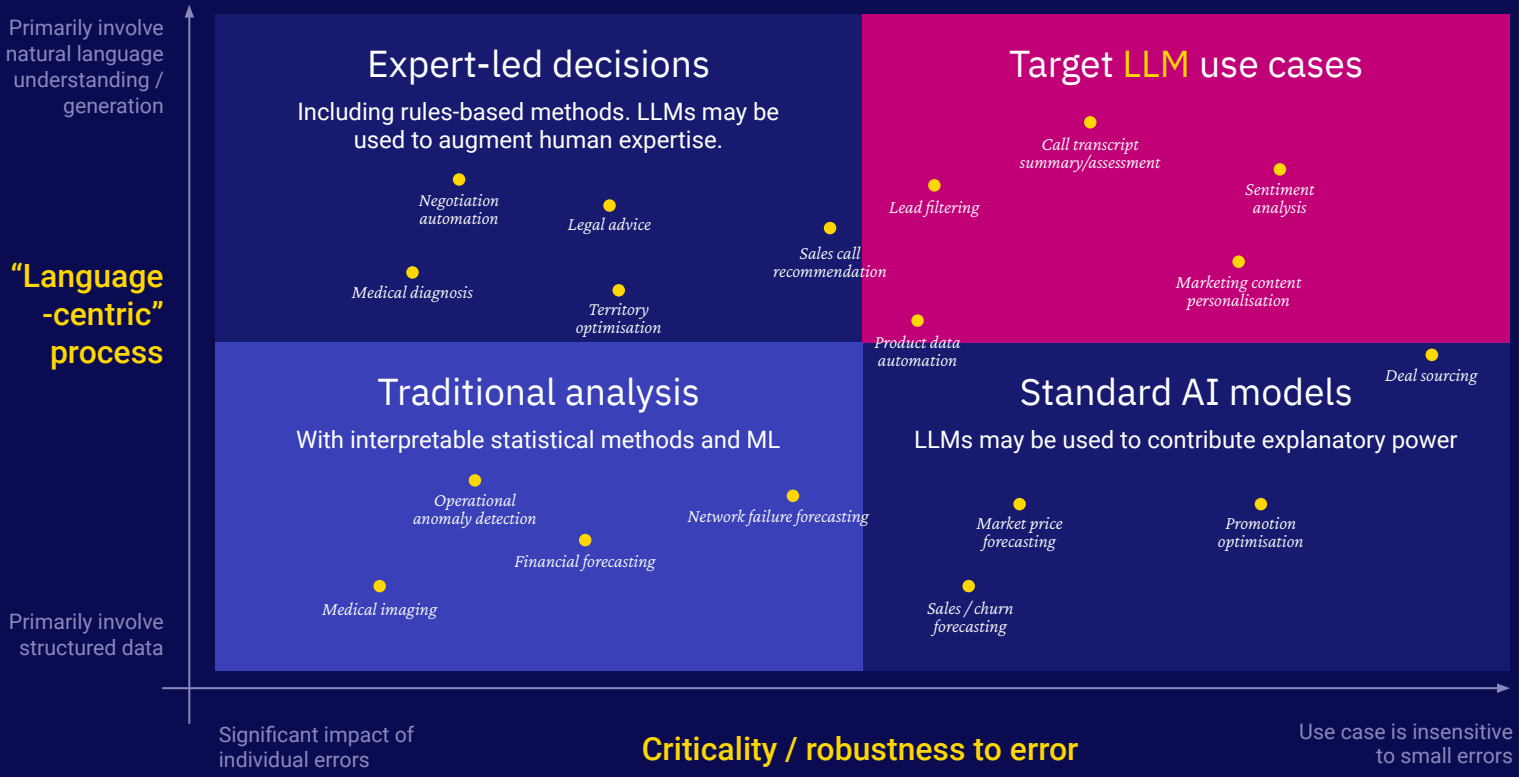

Integrate into an end-to-end system

Customer acquisition

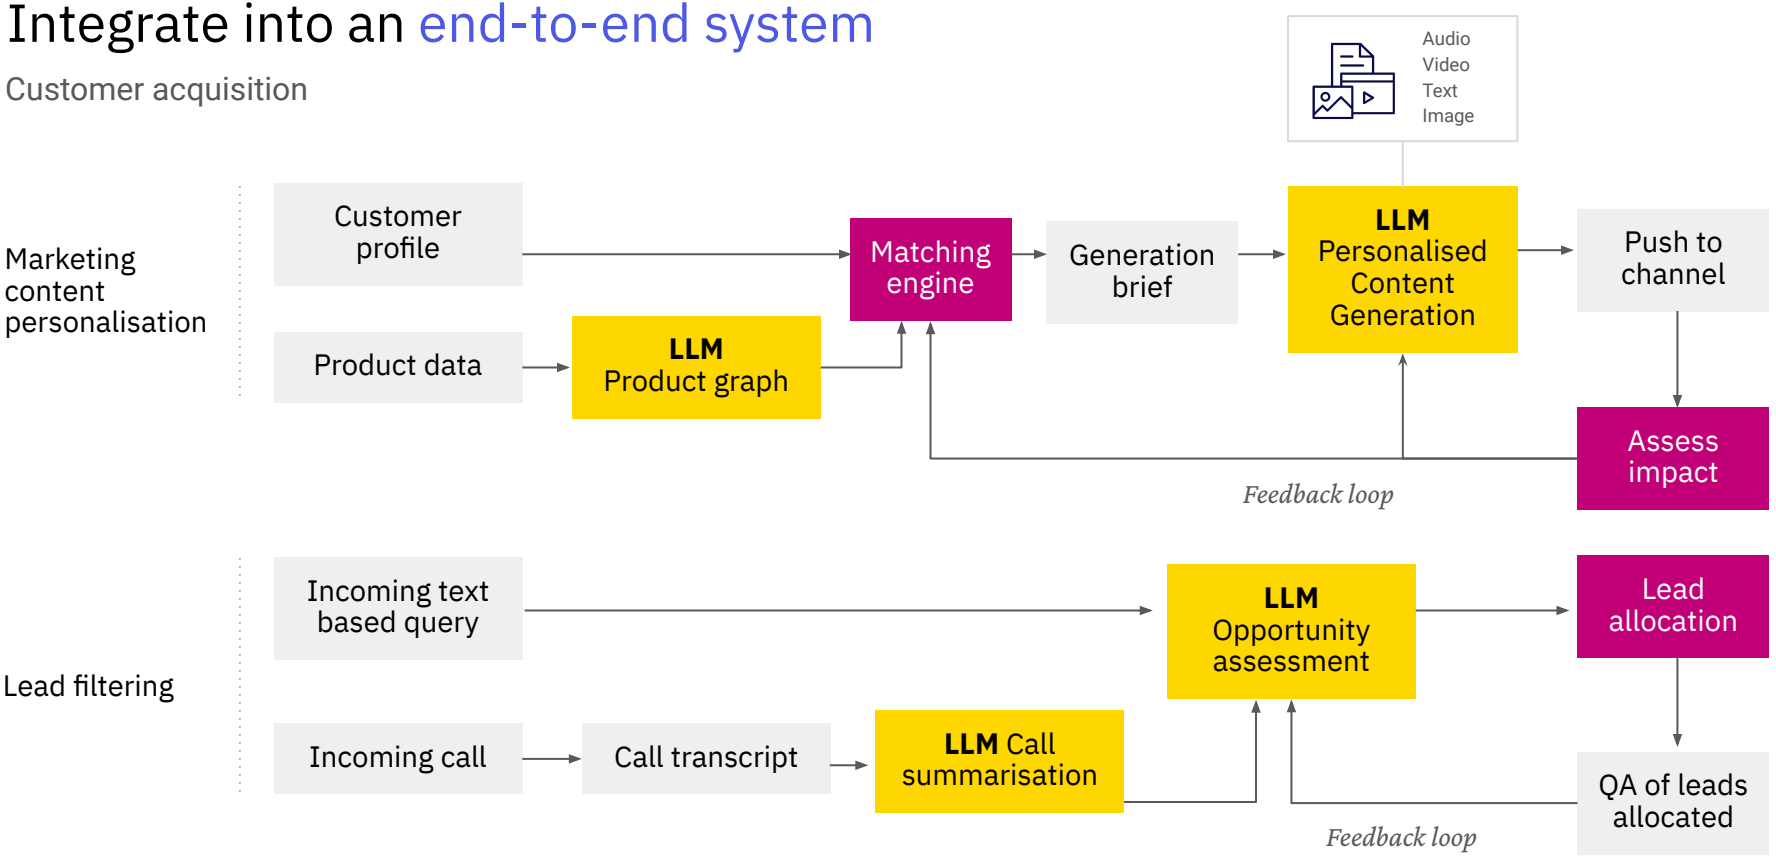

Integrate into an end-to-end system

Workforce management

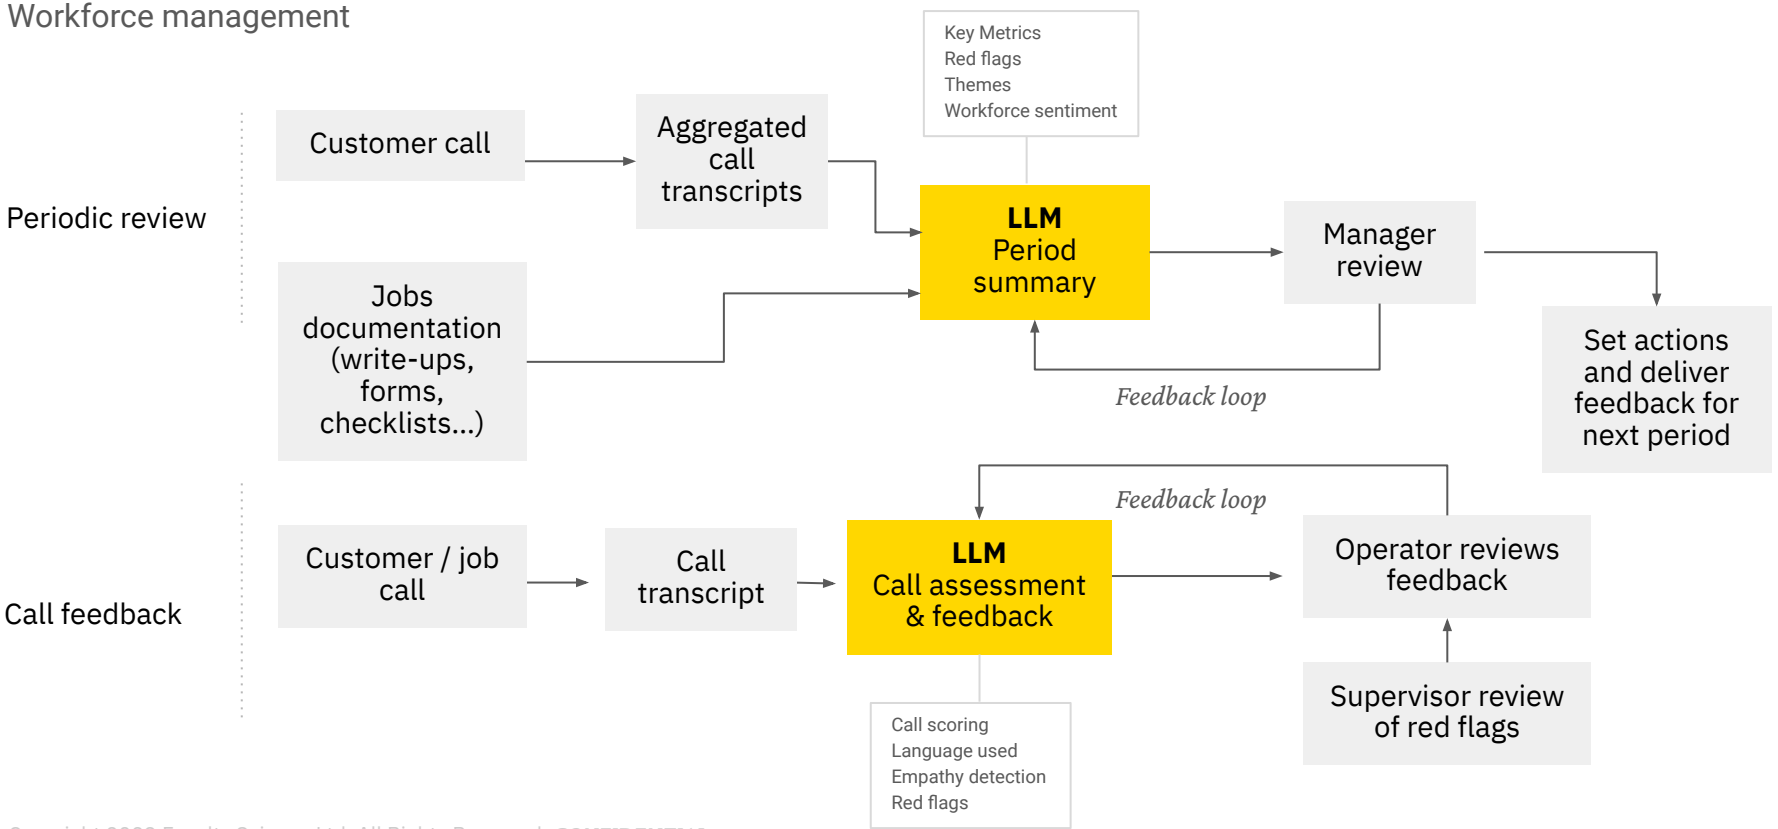

Build in AI safety from the outset

Faculty Responsible AI Framework

Explainability

How can I understand it?

Can I trust the output?

How do I monitor my model over time?

Who is accountable?

Is there a QA loop to capture internal feedback?

Fairness

Does it do the right thing?

Has it been legitimately and statistically validated?

Is it unfairly biased?

Privacy

Does it preserve privacy?

Are measures in place to minimise unnecessary exposure?

Can sensitive information be extracted from the model?

Robustness

Is it robust in the real world?

Is the impact of errors effectively mitigated within the E2E process?

Can I estimate uncertainty?

Does it generalise well?

Is it robust to attacks?

Is a layer needed between the model and consumers?

How to get started

Getting started with LLMs - key considerations

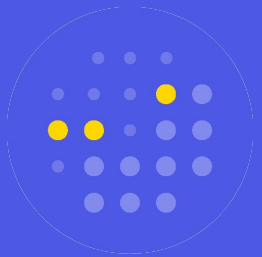

How you interact with them

Cost and commercial model

Privacy and legal considerations

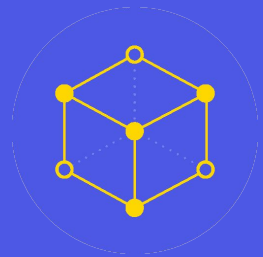

Internal technical architecture

There are numerous ways for organisations to interact with LLMs - different modes are appropriate for different use cases

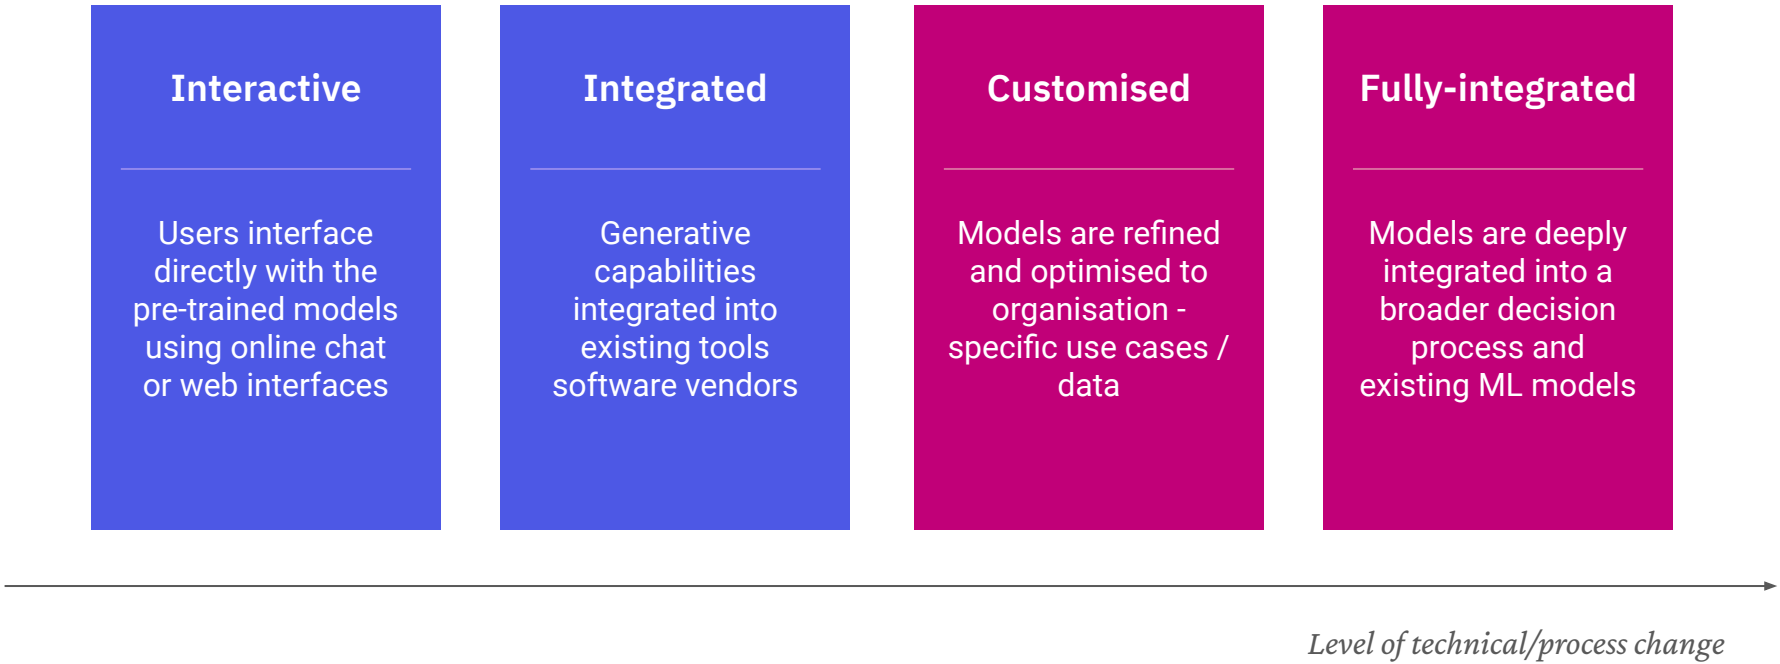

Better understanding LLM cost and commercial models

01

There are **three components** to LLM costs

- "Fine tuning"
- "Prompting"
- "Completion"

02

They are typically charged on the **volume of data** input/output

- 100 "tokens" ~ 75 words
- 30 min discussion = 6.5k - 9.5k tokens

03

Costs are **highly variable** across different models (and with time)

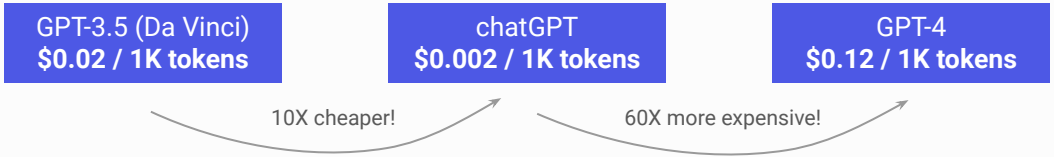

Key takeaways

- You won't always need the best performing model for all use cases
- Estimate and optimise your use case up front
- Remain flexible in your deployment approach as things change

The key legal, security and privacy considerations

Transfer and storage of your data

Data that the foundation models are trained on

Ownership of prompts and generated outputs

The impact of hallucination and accuracy

A new system design for using Generative AI in an enterprise context

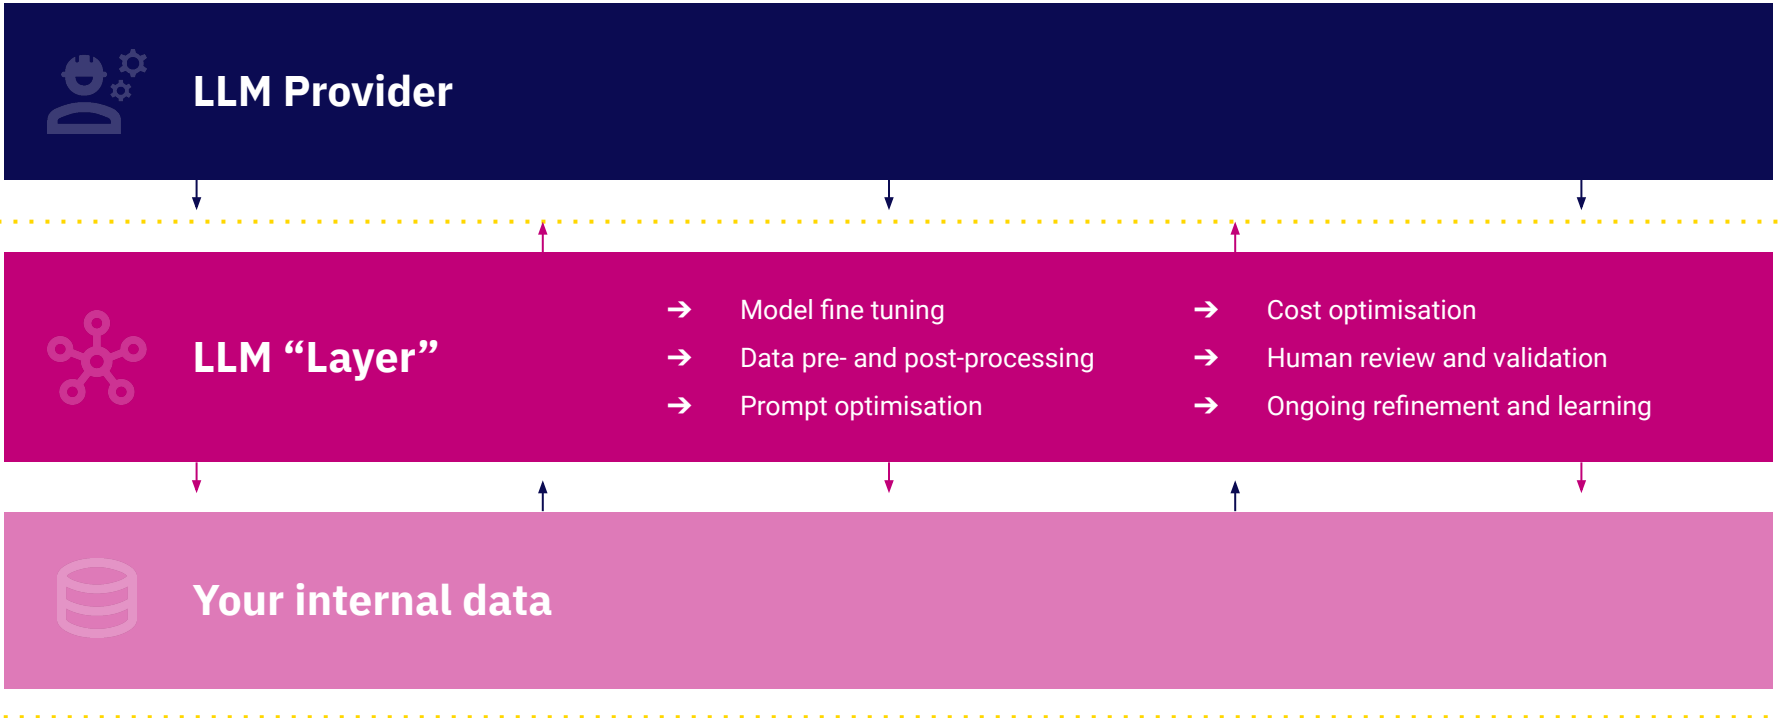

How we help our customers chart a course with Generative AI

1

Educate and ideate

1-2 workshops (½ day)

- Interactive Generative AI workshops and demos of latest models
- Outside-in perspective of relevant applications
- Develop a customised list of use cases for your business

2

Prioritise and scope

2-3 weeks

- Assess each of use case on feasibility and benefits
- Define the best interaction mode for each use case
- Evaluate cost, safety and performance implications

3

Plan for deployment

1 week

- Define the target technical architecture
- Develop a build plan for the top use case(s)
- Knowledge transfer and handover

Summary / Q&A

Summary

01	Generative AI is here and will be one of the most profoundly transformative technologies of our time
02	There are some immediate opportunities for businesses to use this technology today
03	Picking the right problems, thinking systematically and starting with a safety-first mindset are key

Q&A session

Iain Mackay
Senior Manager
Moderator

Marc Warner

Co-founder and CEO

marc@faculty.ai

Kat James

Technical Director
Retail & Consumer

katherine.james@faculty.ai

Josh Muncke

Commercial Director
Retail & Consumer

josh@faculty.ai

Andrew Perry

Commercial Director
Energy Transition and Environment

andrew.perry@faculty.ai

OpenAI | faculty

Thank you

If you're interested in finding out more about how Faculty can help you leverage OpenAI services, **get in touch with us!**

160 Old Street, London, EC1V 9BW, UK