What do AI Startups Want from Regulation?

July 2022

COADEC

Authors

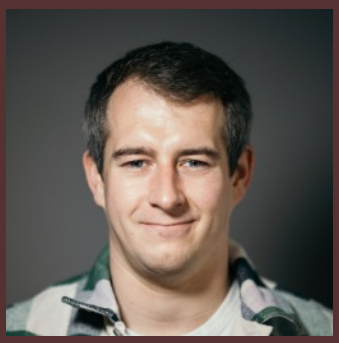

Dom Hallas
Executive Director
The Coalition for a Digital Economy

Camilla de Coverley Veale
Policy Director
The Coalition for a Digital Economy

Leah Friedlander
Policy Lead for Emerging Technologies
The Coalition for a Digital Economy

About Coadec

The Coalition for a Digital Economy (Coadec) is an independent advocacy group that serves as the policy voice for Britain's technology-led startups and scale ups.

Coadec was founded in 2010 by Mike Butcher, Editor-at-Large of technology news publisher TechCrunch, and Jeff Lynn, Executive Chairman and Co-Founder of online investment platform Seedrs.

We fight for a policy environment that enables early-stage British tech companies to grow, scale and compete globally. We have over 2500 startups in our network and have been instrumental in building proactive coalitions of businesses and investors on issues that are integral to the health of the UK's startup ecosystem. Our recent work has seen many successes, from the establishment of the Future Fund to the creation of the Scaleup visa.

Acknowledgements

We are very grateful to every startup, scale up, founder and investor who contributed to this report, particularly those who responded to our survey or provided us with their experiences and insight in interviews.

Executive Summary

AI is likely to be the definitive technology of the century, disrupting and reforming global economies.

At this early stage in the AI race, the UK can say it has many advantages - our AI ecosystem is envied around the world. The UK has an estimated 1400 high-growth AI startups, employing more than 20,000 and generating more than £1 bn in revenue.[1] In 2020 alone, the UK's AI ecosystem was worth more than £15 bn.[2] But the UK cannot afford to rest on its laurels: despite its strength, the AI ecosystem is not without challenges, nor is it immune to damage from regulatory overreach.

Startups by their very nature are created to do things differently and take the sorts of innovation risks other companies are unwilling or unable to. For policy makers this means startups will effectively be the litmus test for AI rules - if AI regulations don't work for startups, the entire AI sector will suffer.

Following the introduction of the EU's draft Artificial Intelligence Act (AIA), and the UK Government due to publish its own AI White Paper, we decided we needed to ask the AI ecosystem - what do they want from future regulation?

The startups we talked to were generally positive about the UK being able to cement its position as a world-leading hub for AI innovation - as long as it gets its regulation right.

Introduction

AI is likely to be the definitive technology of the century, disrupting and reforming global economies. As AI capabilities advance, the consequences of different types of automation and machine learning will impact different industries and the economic and political consequences likely to give rise to a new type of geopolitics. It is clear that if the UK wants to protect its interests as well as strengthen its economy and soft power, it is going to need to enable AI startups to succeed.

The UK is starting with an advantage in this global race: our AI ecosystem is already envied around the world.

The UK has an estimated 1,400 high-growth startups where AI is a critical part of their products or services, including Moneybox, Graphcore and Onfido, employing more than 20,000 people and generating more than £1bn in revenue.[3] In the decade between 2010 and 2021, AI startups raised more than £19bn, with 76% of AI startups securing equity funding. In 2020 alone, the UK's AI ecosystem was worth more than £15bn.[4]

AI startups in the UK also benefit from a thriving deep tech ecosystem, driven by world-recognised universities. While a small number of these university researchers and deep tech companies, such as DeepMind, are working to master the most advanced types of artificial intelligence, their development of machine learning and other software support techniques are filtering through to a much broader range of startups, the early adopters, who use this underlying tech to compete against incumbents.

AI startups do not exist in isolation, they compete with the entire innovation ecosystem, with capital and talent regularly recycled. A strong and vibrant ecosystem attracts new startups and here the UK is also well positioned. Tech Nation's 2022 Tech Week update, released in June, illustrated the strength of our startup ecosystem as a whole: in the first five months of 2022, startups in the UK trailed only those in the US, raising more investment than India and China.[5]

In fact, by June 2022, startups in the UK raised double the funding of any other European market. The update also showed we have built a strong unicorn pipeline - 258 startups are considered to have the potential to be valued at $1bn or over in the coming years - double any European neighbour. This mid-year update demonstrated a resilient startup ecosystem able to support our most innovative AI startups.

But despite this strength, and relative advantages, the UK's AI ecosystem is not without its challenges, nor is it immune to damage from both core challenges on talentand finance as well as regulatory overreach.

Startups by their very nature are created to do things differently and take the sorts of innovation risks corporates are much less likely to. For policy makers this means startups will effectively be the litmus test for AI rules - if AI regulations don't work for startups, the entire AI sector will suffer.

With the EU introducing a draft Artificial Intelligence Act, and the UK Government due to publish an AI White Paper, we decided we needed to ask the AI ecosystem - what are their challenges and what do they want from future regulation?

The startups we talked to were generally positive about the UK being able to cement its position as a world-leading hub for AI innovation - as long as it gets its regulation right. Polling the EU's key proposals on AI regulation, we found deep reservations for every one of them - this gives the UK a real Brexit opportunity to support its AI ecosystem to move faster and further.

UK / EU Regulatory State of Play

The UK's Path to AI Regulation

The UK's broad priorities for AI were first outlined in 2017 in an Industrial Strategy White Paper[6], setting out the Government's ambition to boost productivity and earning power throughout the UK. These included finding ways to use data and AI for the prevention, early diagnosis and treatment of chronic disease, as well as offering citizens and businesses more confidence and clarity on AI deployment, and helping people develop the skills needed for future jobs.

The AI Sector Deal[7], published in 2017 and updated in 2019, built upon this thinking, setting out a £1bn package of support to boost the UK's global position as a leader in developing AI. The Strategy called for the creation of the Centre for Data Ethics and Innovation (CDEI) to 'ensure safe, ethical and ground-breaking innovation in AI and data-driven technologies'.

In parallel, a Lords Select Committee was appointed to consider the economic, ethical and social implications of advances in AI. The Committee published its report titled 'AI in the UK: ready, willing, able?'[8] in 2018, setting out an ethical framework to guide the development and application of AI. These recommendations included a Government Office for AI to help coordinate and grow AI in the UK, an AI Council to create a system to inform people when decisions are being made by AI, and a cross-sector code to ensure AI development and implementation remains ethical. As a result, the AI Council, an independent-expert committee providing advice on the AI ecosystem and the implementation of the AI Sector Deal, was established in 2019.

In January 2021, the AI Council published a roadmap[9] making recommendations to the Government to double-down on investments, remain adaptable and develop a globally ambitious strategy.

Later in 2021, the Government released its new National AI Strategy as a follow-up to the 2018 AI Sector Deal. The Strategy outlined the Government's plan for turning the UK into a global AI powerhouse in the next 10 years and is centred around three key principles:

- Investing in, and planning for, the long-term needs of the AI ecosystem, with the objective of making the UK a computer science and AI superpower;
- Supporting the transition to an AI-based economy; and
- Taking an informed and balanced approach to regulating AI technologies, by encouraging innovation whilst protecting the public and fundamental values.

The Government is expected to publish proposals for sector-agnostic AI regulations in an AI White Paper in 2022. In tandem with this, the Government has announced its intention to introduce a Data Reform Bill[10], in which the Government indicated it will look at refining and clarifying ambiguous sections of GDPR, such as when personal data can be repurposed for research.[11]

The EU's Path to AI Regulation

The European Union's draft Artificial Intelligence Act (AIA) aims to establish the first comprehensive regulatory scheme for artificial intelligence.[12] In 2021 the EU proposed a common set of EU rules for AI systems that are intended to shape the way AI developers, providers, users and governments use AI. Alongside GDPR, the AIA is part of the EU's recent positioning of itself as a regulatory leader in tech.

One of the EU's central proposals is that it should be able to categorise AI systems on the basis of risk, from "unacceptable' through to 'minimal'. Developers and users of AI systems categorised as 'high-risk' would have mandatory legal requirements, including direction on how to develop the AI system, imposed upon them. In addition they will be required to demonstrate to a regulator that their systems comply with a series of strict obligations before they can be put on the market. Proposed obligations include mandatory risk assessments and strict data activity logs.[13]

The draft legislation's current plans for each category, but especially high-risk AI, are broad. In the current proposals the EU can also adjust what it considers to be high-risk AI without new legislation. For this reason it is proposed that limited and minimal-risk AI systems be encouraged to comply with high-risk AI compliance but legally will have far less onerous duties.

At the time of writing, negotiations on a revised draft of the AIA have begun with amendments estimated to be over three-thousand[14]. This included a submission by the French Presidency indicating it wanted the draft regulatory framework to include General Purpose AI systems within scope.[15] General Purpose AI systems are those able to be applied to different tasks without needing substantial modification. These systems are often characterised by their role as foundation models on which more specialised AI systems are built, so one General Purpose AI system for language processing might become the basis for hundreds of applied models such as chatbots, decision assistants and translations services. General Purpose AI is also often open source AI, used by startups to enable them to get to their innovation faster without deploying time and resources creating an underlying infrastructure. If adopted, this would represent a vast broadening of the scope of legislation - dramatically increasing regulatory burden by effectively pulling in all AI technology regardless of its risk classification.

Views from Our Ecosystem

Startups by their very nature are created to do things differently and take the sorts of innovation risks corporates are much less likely to. For policy makers this means startups will effectively be the litmus test for AI rules - if AI regulations don't work for startups, the entire AI sector will suffer. Coadec decided to delve into the AI ecosystem and survey what AI startups see as their greatest challenges and try to understand what they want from possible future regulation.

Categorisation

The EU's draft Artificial Intelligence Act (AIA) is considering introducing a 'high risk' classification based on what sectors companies operate in, such as education or healthcare, regardless of an AI system's features. Would you be concerned if the UK Government copied this approach?

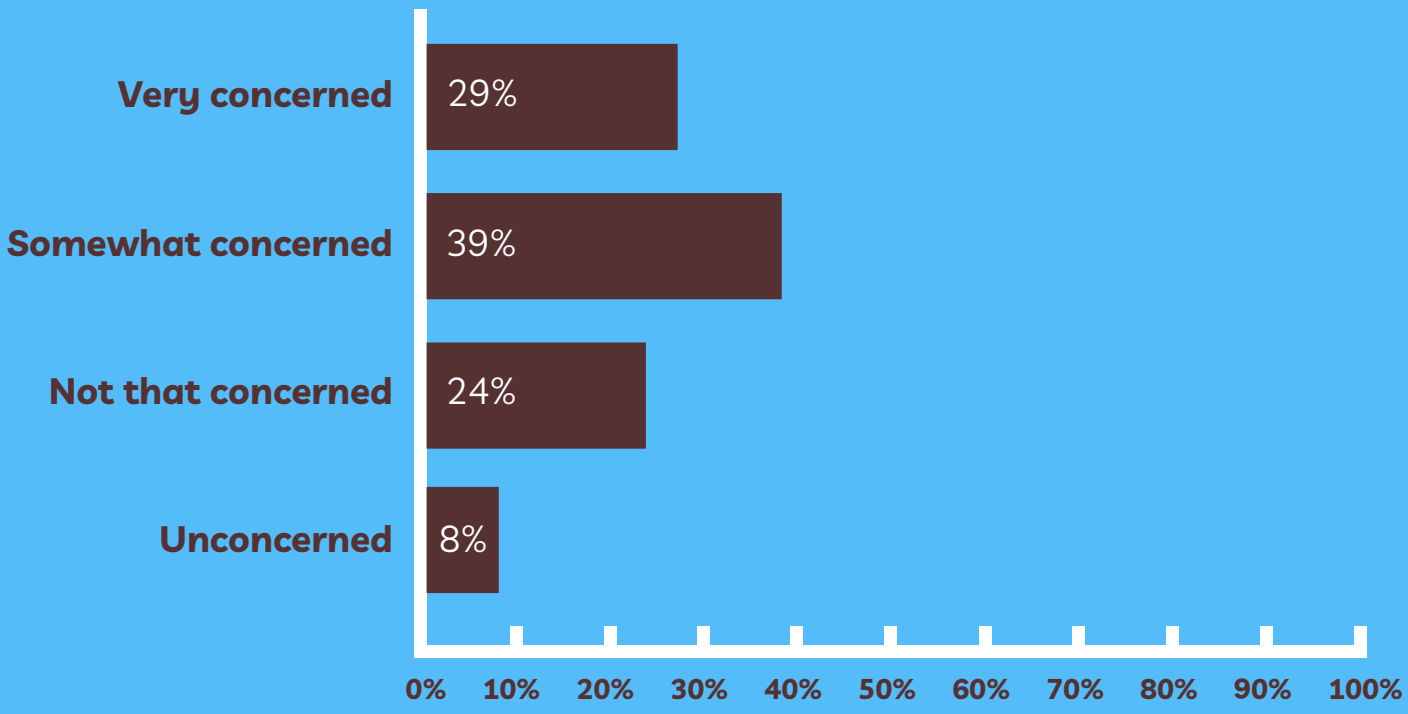

If the UK were to introduce a categorisation system for AI innovations, including Unacceptable AI, High Risk AI, and Low Risk AI, do you agree that there would need to be clear definitions for each category?

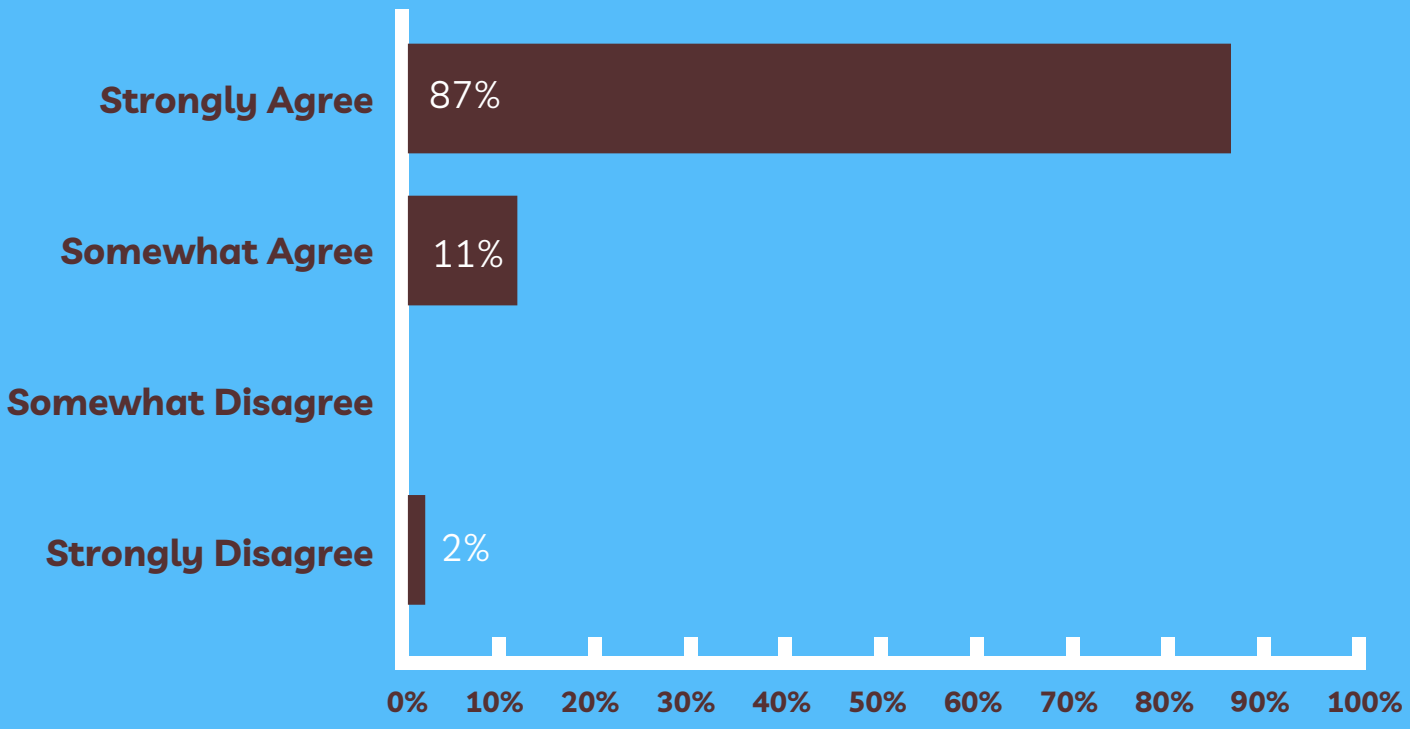

At Coadec we believe innovation in fields that will likely be categorised as high-risk AI, for example, facial recognition, should not be automatically equated with high-risk use. How strongly do you agree that any risk-classification regime the UK Government creates for AI should take care not to discourage innovation?

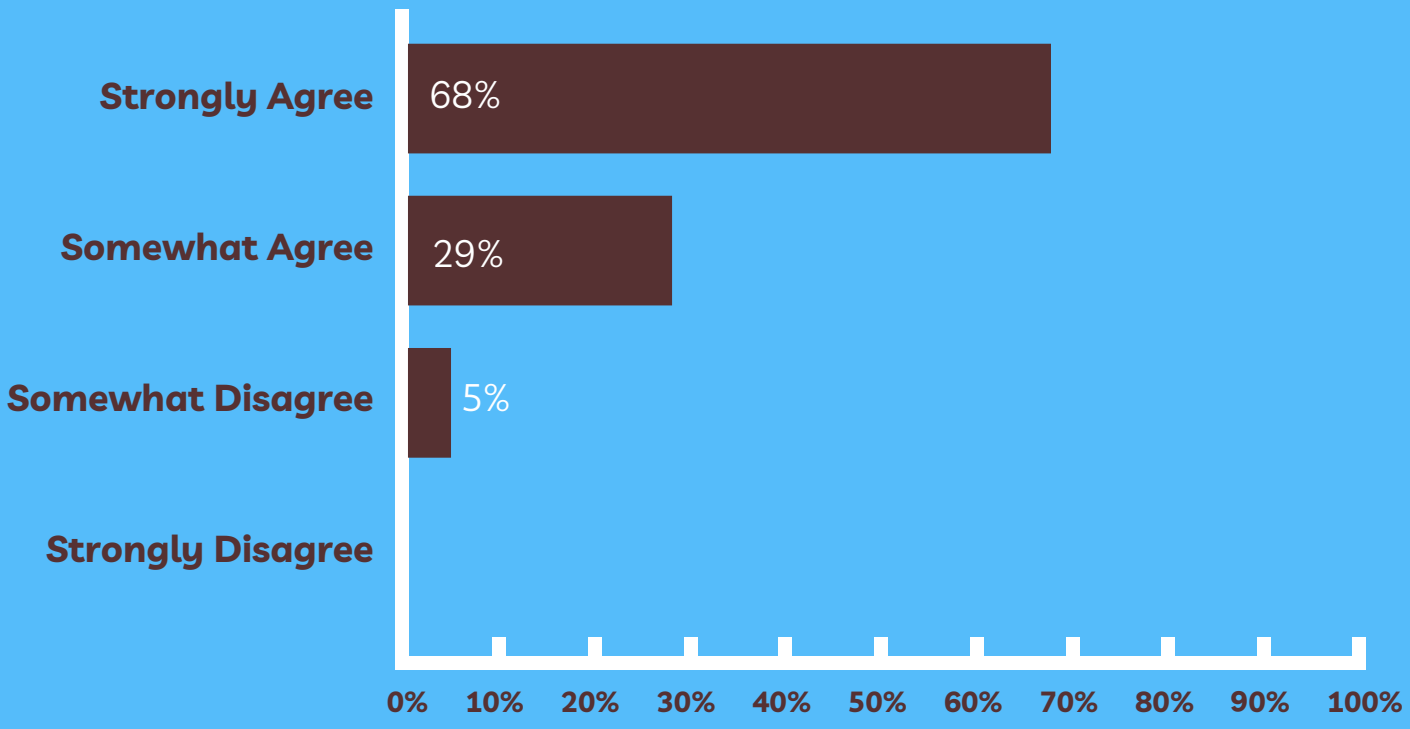

68% of respondents told us they would be either very or somewhat concerned if the UK Government were to introduce a similar categorisation system as the EU's proposed risk-based approach.

A common theme that came up in conversations with AI startups is innovation in a high-risk field does not automatically translate into high-risk use. In many areas where AI has the potential to revolutionise outcomes, such as in medical diagnosis, we need lots of different startups iterating to get to a usable product or service. The proposals for a categorisation system risks discouraging startups from innovating in fields where we desperately need them.

Liability

Rules around liability should be proportionate, clear and include only material or physical damage:

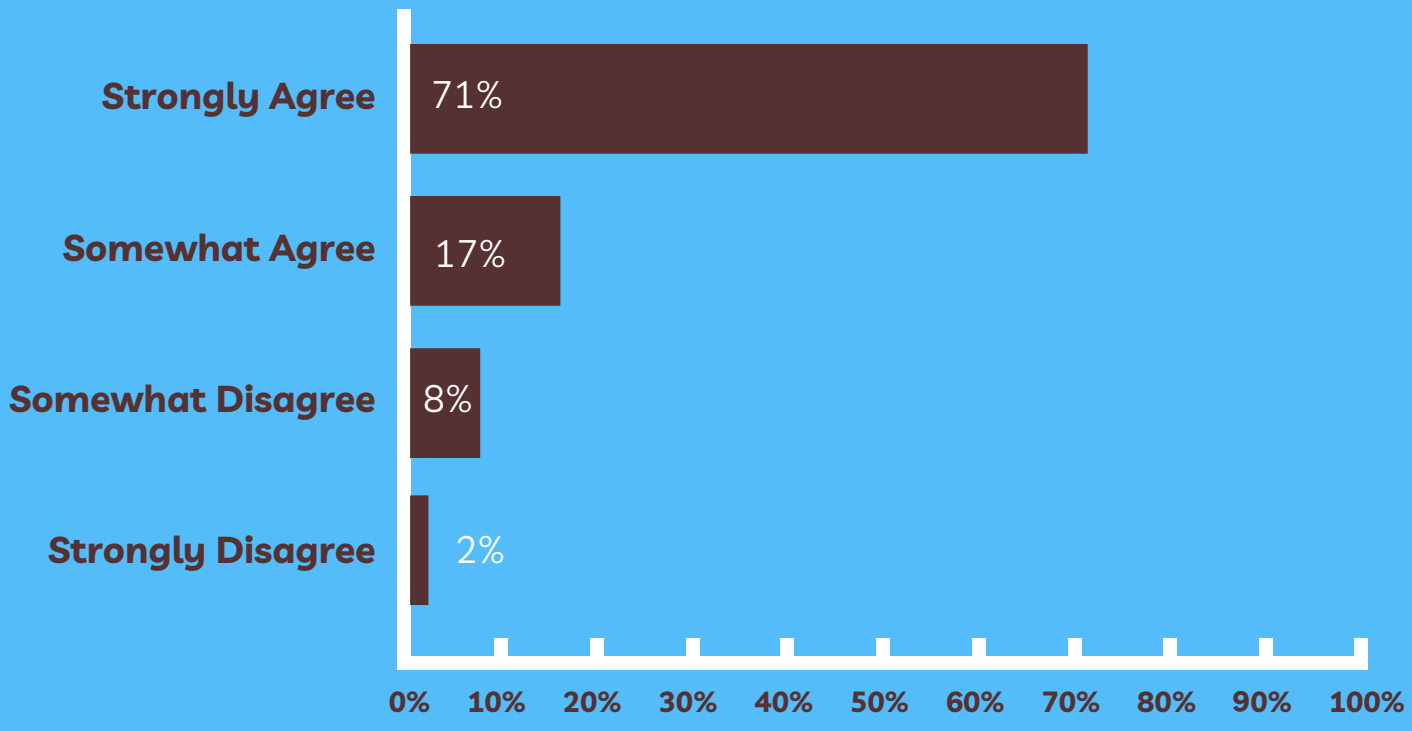

A lack of clarity regarding liability for AI developers will only discourage entrepreneurs:

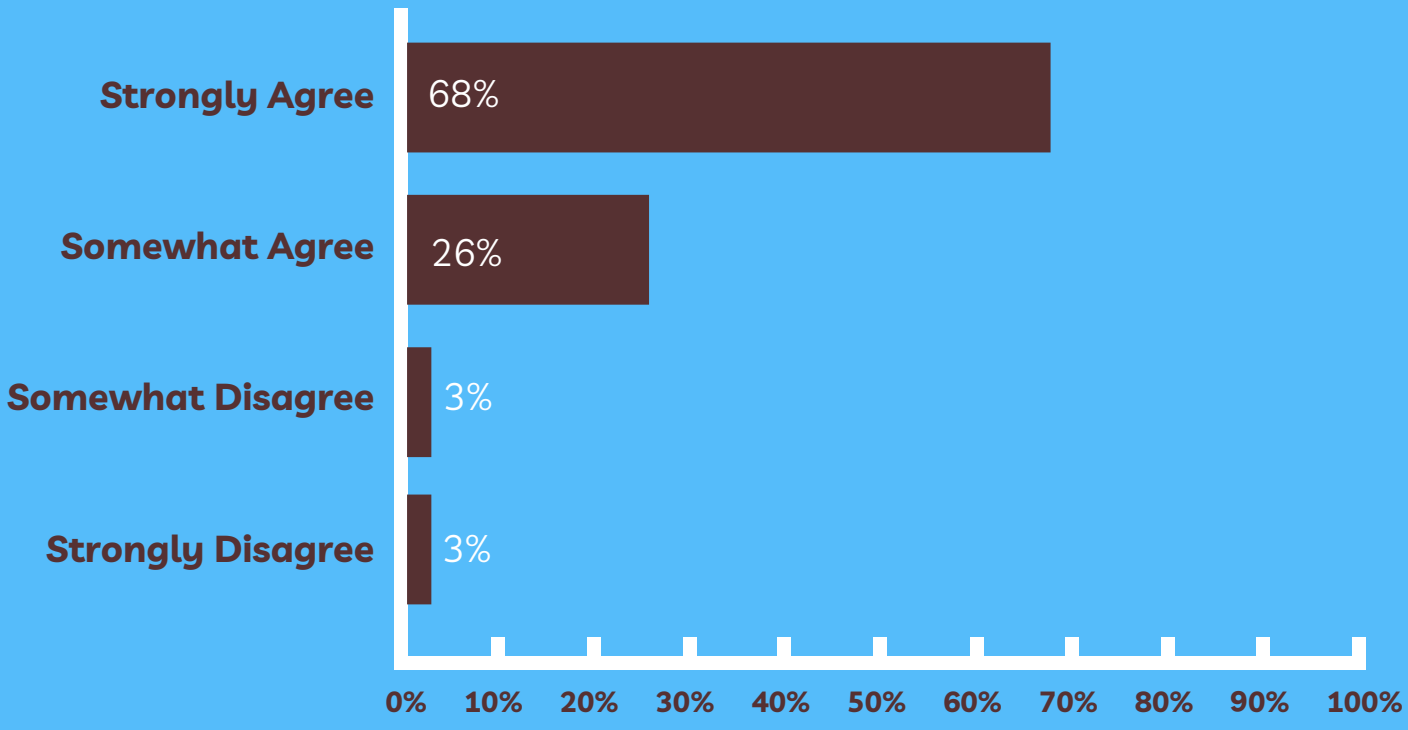

If the UK Government made AI developers liable, how do you anticipate this would affect your business?

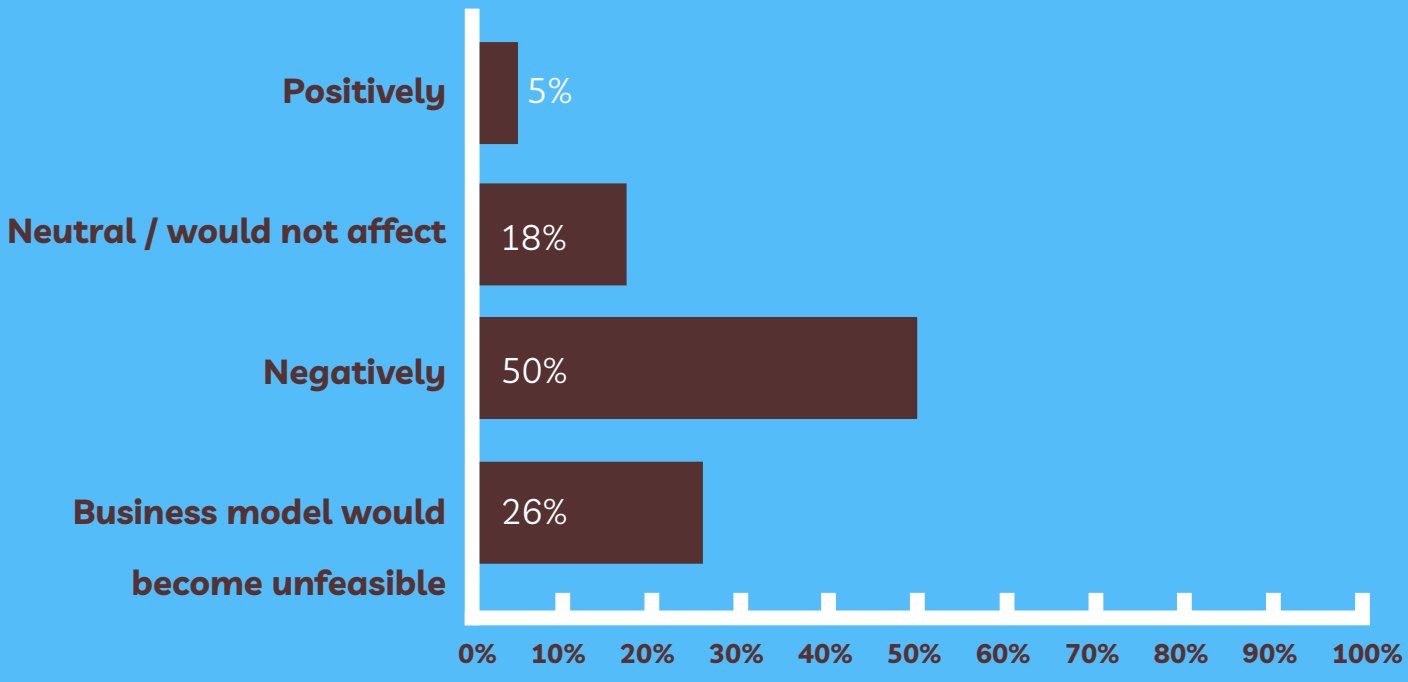

If the UK Government introduces new rules increasing the legal responsibility of AI providers, operators and users, do you agree that their obligations should be clearly defined?

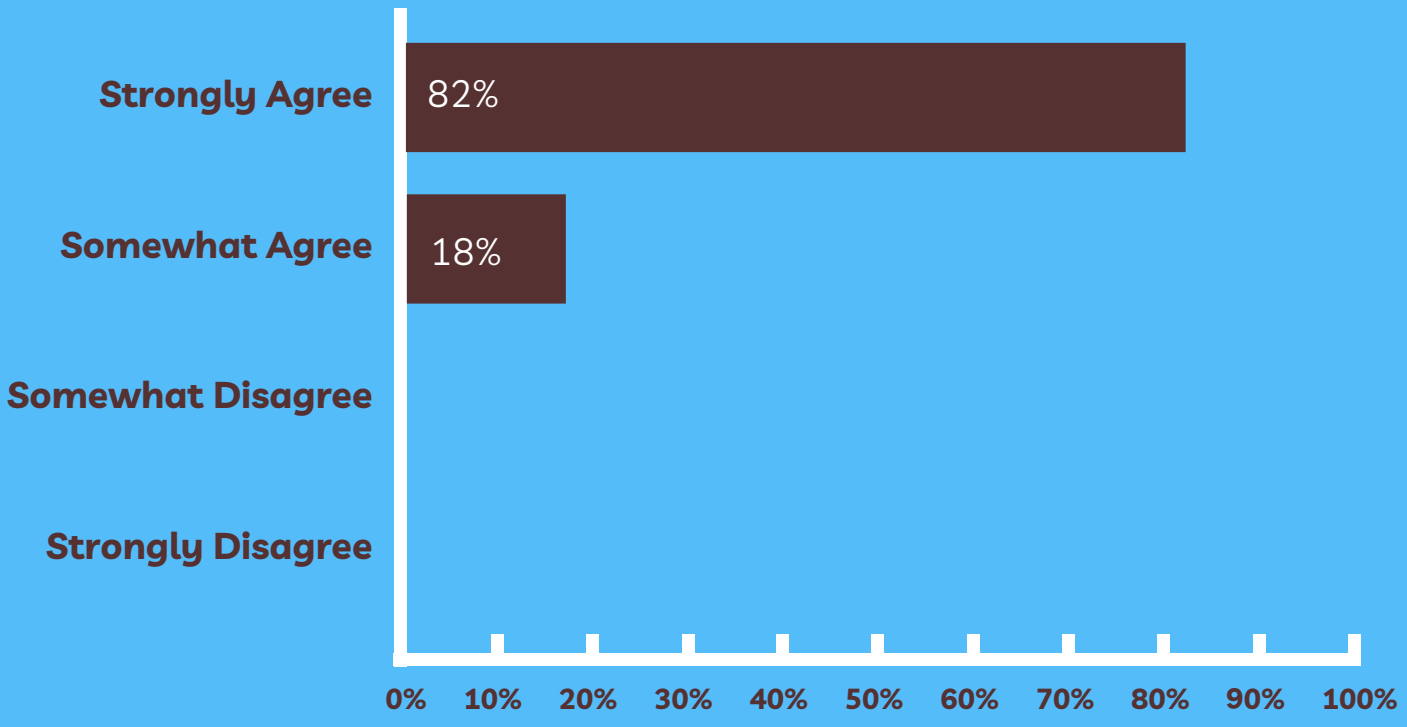

The EU is currently consulting on a potential update to its Product Liability Directive that would include digital services, as well as make AI developers and programmers liable for damage caused by their AI systems. The current Directive places the burden of proving the causal connection between product defect and damage on the injured party. Information asymmetry between the technical features of digital economy products and consumers prompted the consultation on an update.

The cause of an AI system's failure to perform is a key element to establishing breaches of duty of care or damage, but there are many parties involved in an AI system's creation. The EU's consultation has been an eye catching proposal to many AI startups as developers of AI systems are not always the ones who deploy it on the market, nor do providers typically have access to an AI system or its performance data after it has been purchased by a third party.

76% of startups told us their business model would be either negatively affected or become infeasible if the UK were to follow the EU in making AI developers liable.

If the UK Government were to introduce new rules increasing the legal responsibility of AI developers, operators and users, 100% of respondents agreed that any new obligations needed to be clearly defined.

89% told us the rules around liability would need to be clear, proportionate and cover only material or physical damage. If there were to be a lack of clarity when it comes to liability, 94% of startup founders were clear it would discourage entrepreneurship within the sector.

Algorithmic Transparency

The UK Government is considering developing a standard for algorithmic transparency, including potentially mandating commercially sensitive IP be shared with regulators. How concerned would you be with sharing commercially sensitive IP with regulators?

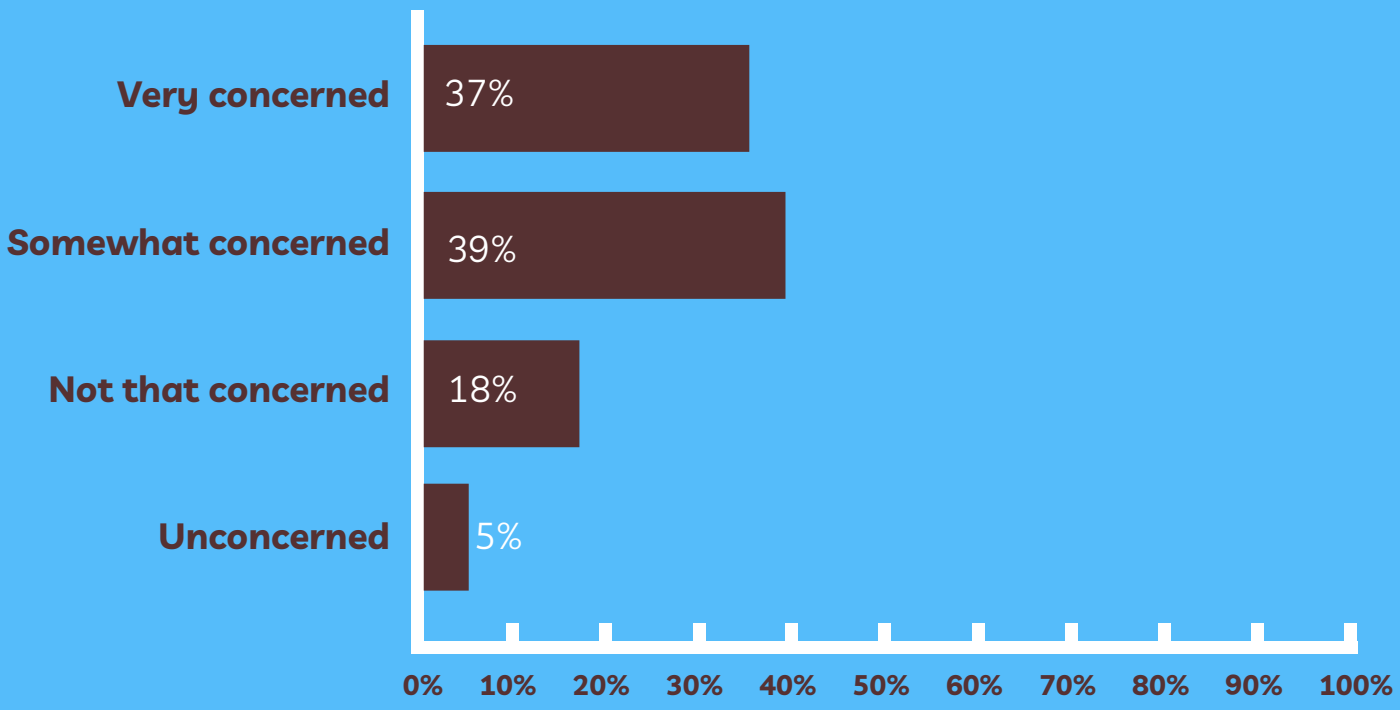

With companies and Governments increasingly using algorithms to make automated decisions about what services and products people can access, or what exam grades or diagnosis they receive, transparency standards can be a useful way to provide information about algorithmic decision-making. AI startups we talked to were very supportive of this aim but they expressed concerns about which transparency standards were effective, as well as what information they should have to make transparent.

78% of startups told us they would be either very or somewhat concerned if the UK were to mandate they share commercially sensitive IP with regulators.

For AI startups, their algorithm is critically commercially sensitive information. Being forced to share this with regulators makes startups fear that their IP will find its way into competitors, or that bad actors will be able to more easily discover how to game a system to their advantage.

Startups also pointed out to us that while their source coding can hold some of their most sensitive data - it's also not necessarily the most effective to examine for transparency. For many, their algorithm logic is learned from training data. These are often referred to as 'black box' models because of their opacity. One founder commented that we should not think of open algorithms like open operating systems - you cannot comb over them for bugs or vulnerabilities.

Startups are also concerned that transparency requirements may force them to use models that compromise on quality or efficiency because they are more transparent. Startups share the aim of building trust through transparency, but in a way that preserves the ability to innovate, compete and make the best products and services. Many startups are beginning to look at transparency vs plausible explainability. Accepting that there will be times when transparency will be necessary, such as medical diagnosis, many startups told us that they felt being able to plausibly explain how the system will have made its decision would be more useful to both regulators and consumers.

Getting to data that is 'free of errors'

The EU's draft AIA is considering mandating certain requirements of AI developers, including ensuring AI systems are trained on datasets that are 'relevant and free of errors'. If this were to happen in the UK, would you be:

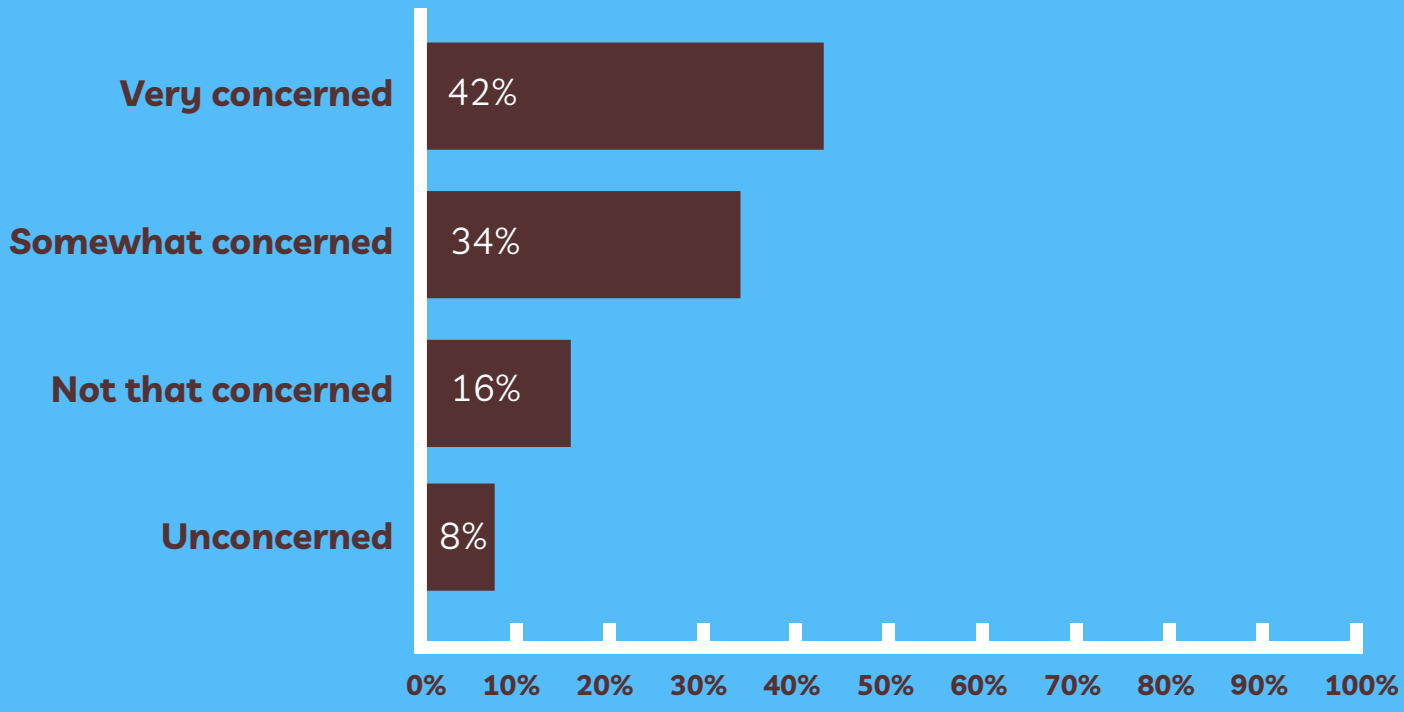

76% of startups told us they had varying concerns about whether the proposal was technically feasible.

The overwhelming majority of startups will want to use data sets that are free of errors - but it is unclear what the EU means by this - especially as data sets can always be improved. It is also unclear how the EU proposes to assess that a data set is free of errors.

Requiring startups to ensure their datasets have no errors places a heavier burden on startups who already struggle to find complete data sets and neglects the natural lifecycle of data - it can be hard to pinpoint the exact time in which an error can appear in a dataset.

Data Audits

Do you frequently conduct general audits of your data (training data, user data, etc.)?

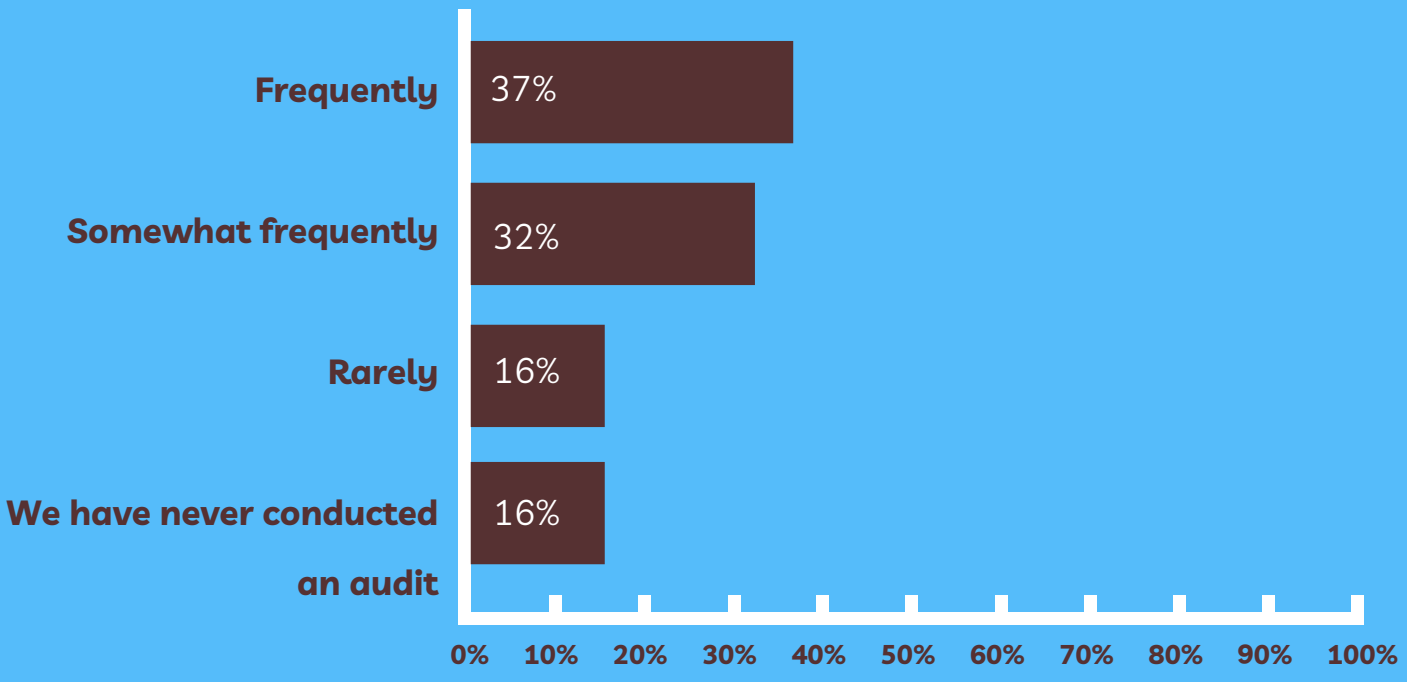

If the government were to propose mandating data audits, how concerned would you be handing over your datasets to regulators?

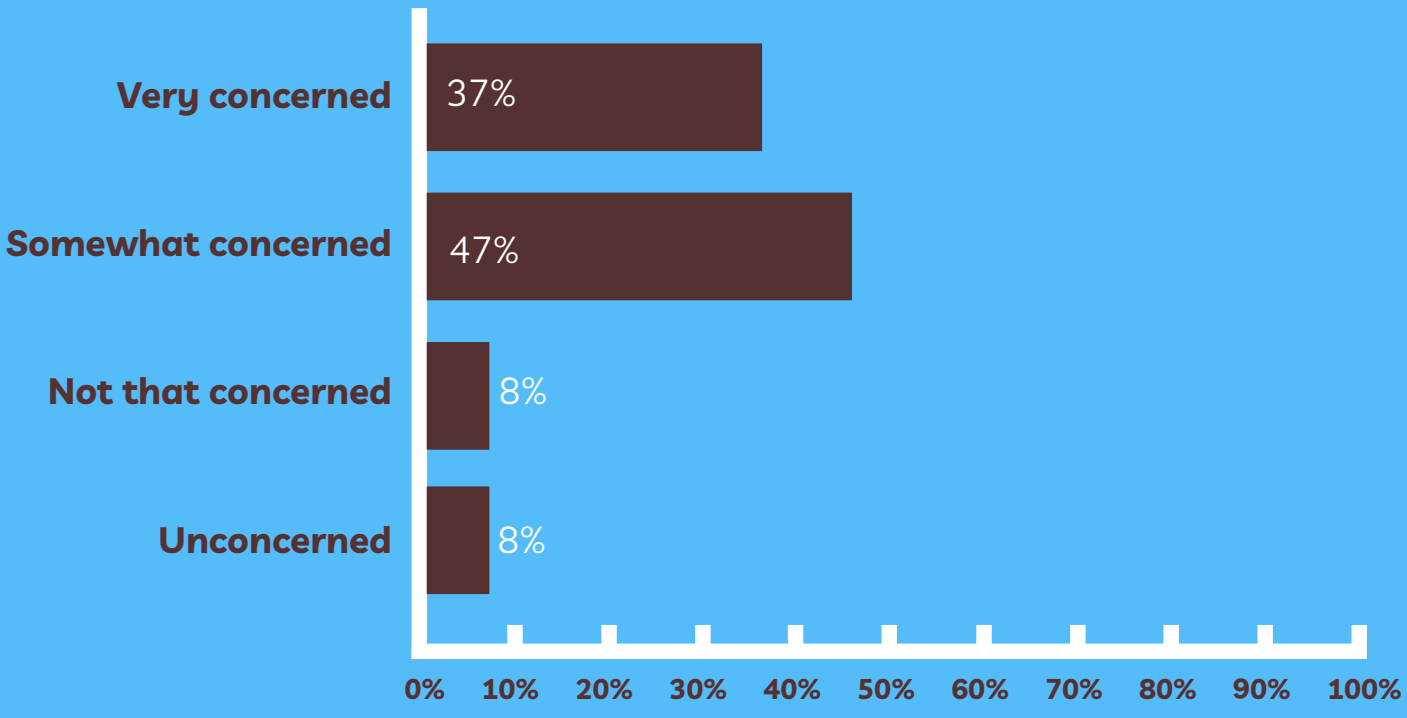

If the UK Government were to introduce regular data audits by regulators, how frequently do you feel would be appropriate?

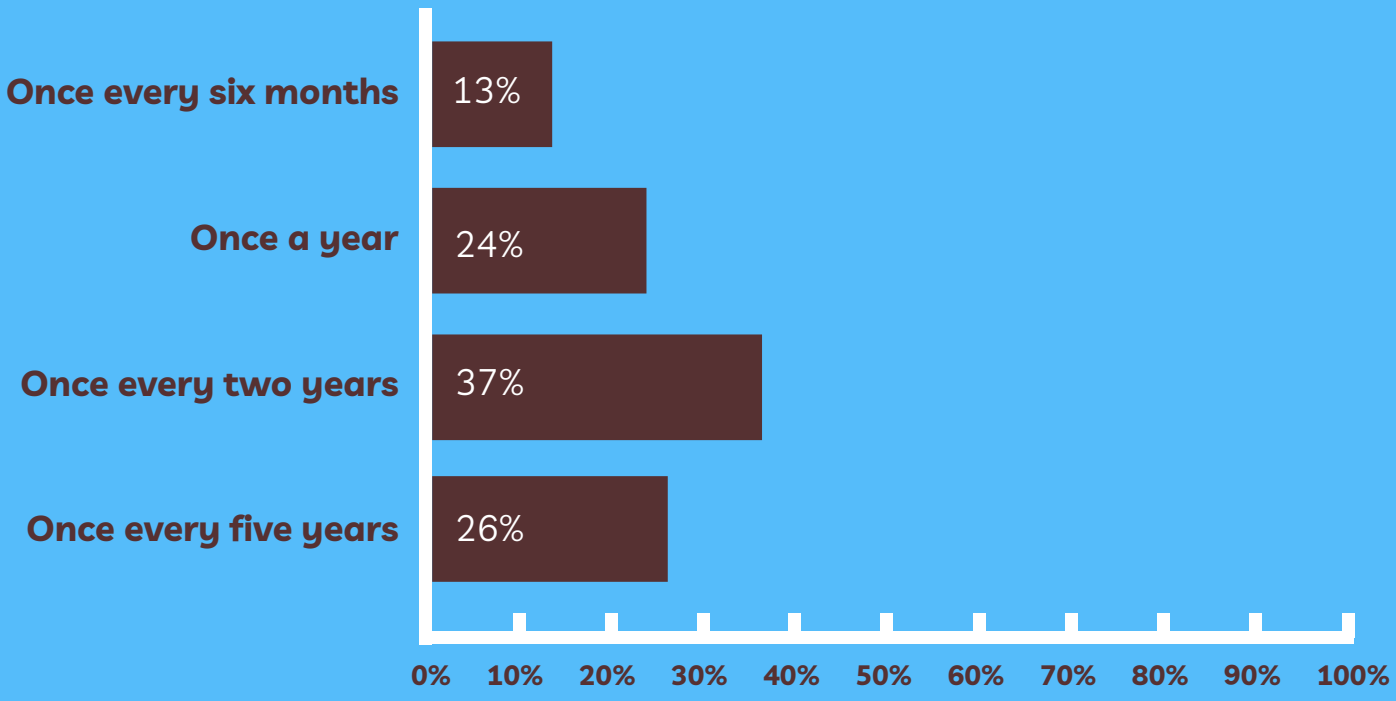

Data assurance helps to discover inconsistencies and anomalies in data, as well as help with data cleansing such as removing outliers and missing information in order to improve data quality. Best practice for data assurance in AI is still developing but conducting data audits is common.

Of the AI startups we spoke to, 68% told us they conducted general data audits frequently or fairly frequently, however 84% expressed concern at the prospect of these data audits being mandated, something which the EU is considering.

AI Ethics

When developing your AI product, which of the following approaches to ethics best describes your process today?

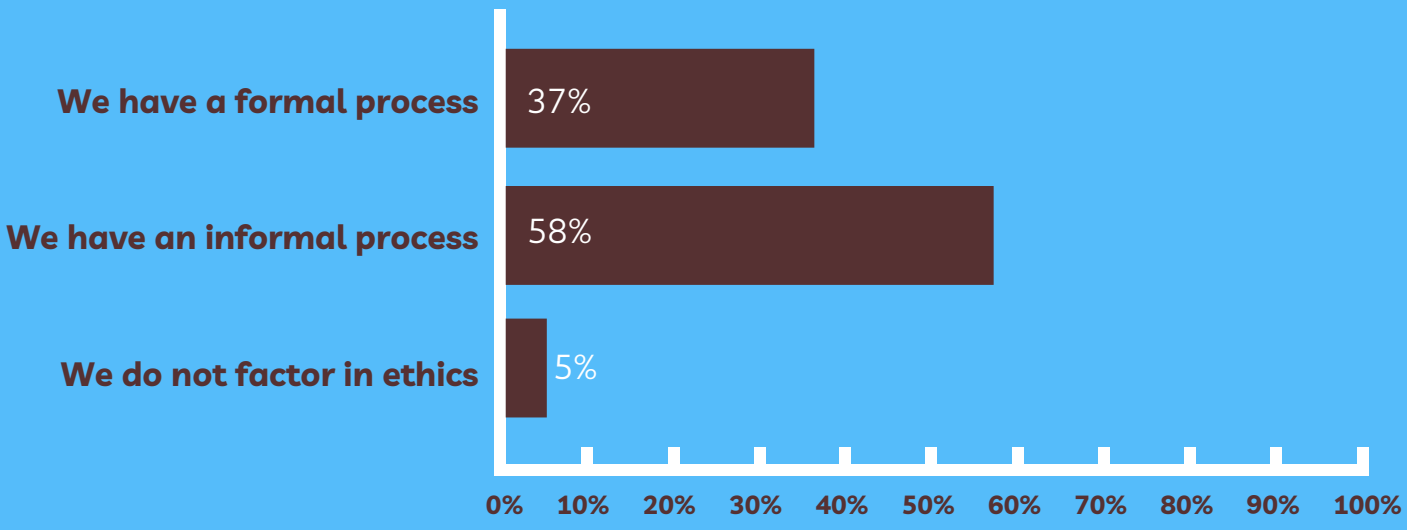

Startups care about making the best products and services possible, including innovating with ethical AI - only 6% of startups told us they did not factor in ethical considerations during their product or service development. In contrast 36% of startups told us formally factored in ethical considerations, a further 58% of startups said they factored in ethical considerations - they just didn't have a formal process for doing so.

Anecdotally, many startups told us that they considered their work too early for formal processes and that they had plans to factor in explicit processes for assessing ethics as their systems matured and developed. Many questioned whether regulators could give meaningful guidance on this in the immediate future. One startup also argued that most popular examples of AI ethics mistakes should more accurately be described as “bad uses of AI or bad AI methodology”.

Data Reform

Which of the following represents your biggest operational costs?

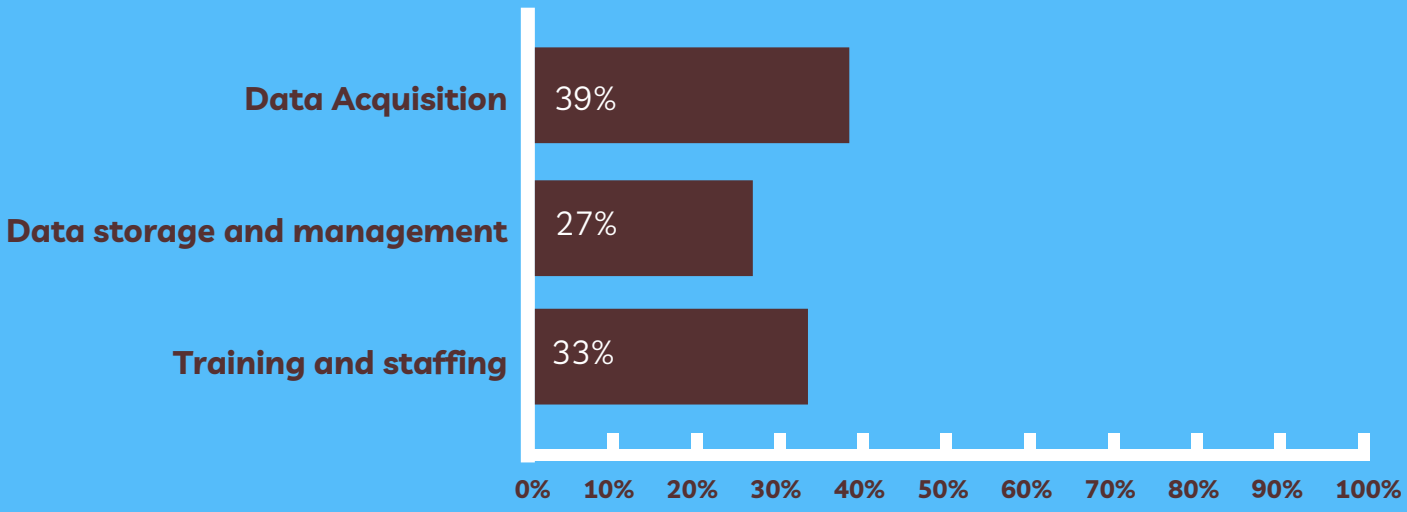

Access to good data is critical for AI startups. 40% of startups told us data acquisition was their biggest operational cost, 27% told it was data storage and management. When we asked startups what they thought would be the greatest barrier to their growth over the next five years, 52% told us they expected it to be the - lack of - availability of datasets.

Because accessing necessary data can be difficult to find and costly to procure, AI startups are often forced to expend disproportionate time and resources sourcing, collecting, purchasing, storing, and evaluating the data they need.

And once a startup has enough data, they need to clean and label it correctly. This process, effectively curating the data so that the machines can ingest it, is necessary to train a machine to perform a new task: in order to teach a machine to recognise planes, you first need to go through your data to tell the machine which of the photos feature planes and which do not.

Working with Regulators

Have you engaged with a regulator when developing your product?

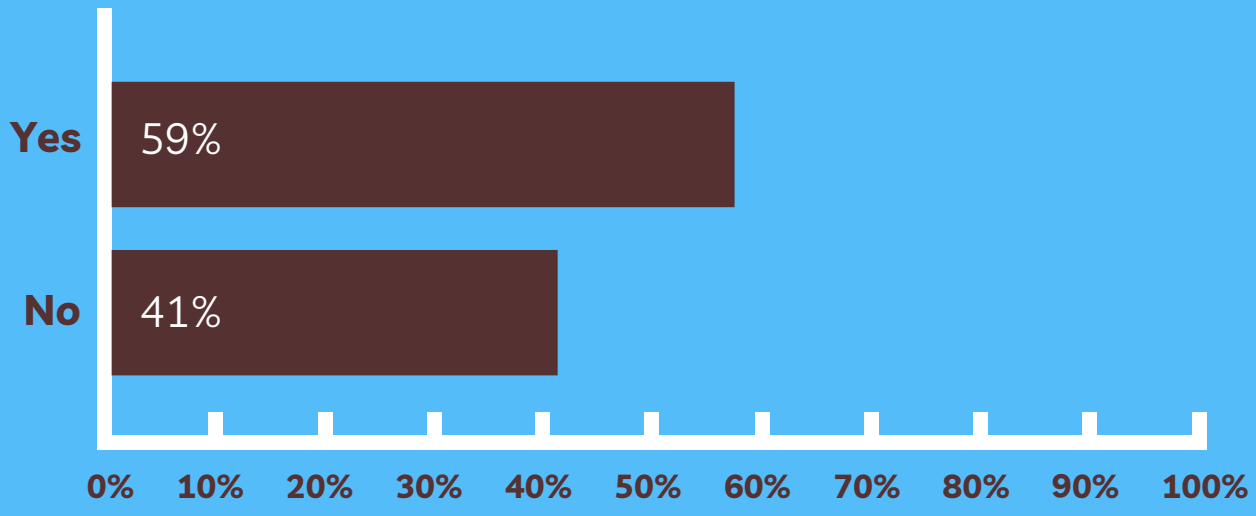

If you have engaged with a regulator in the past, how helpful were they?

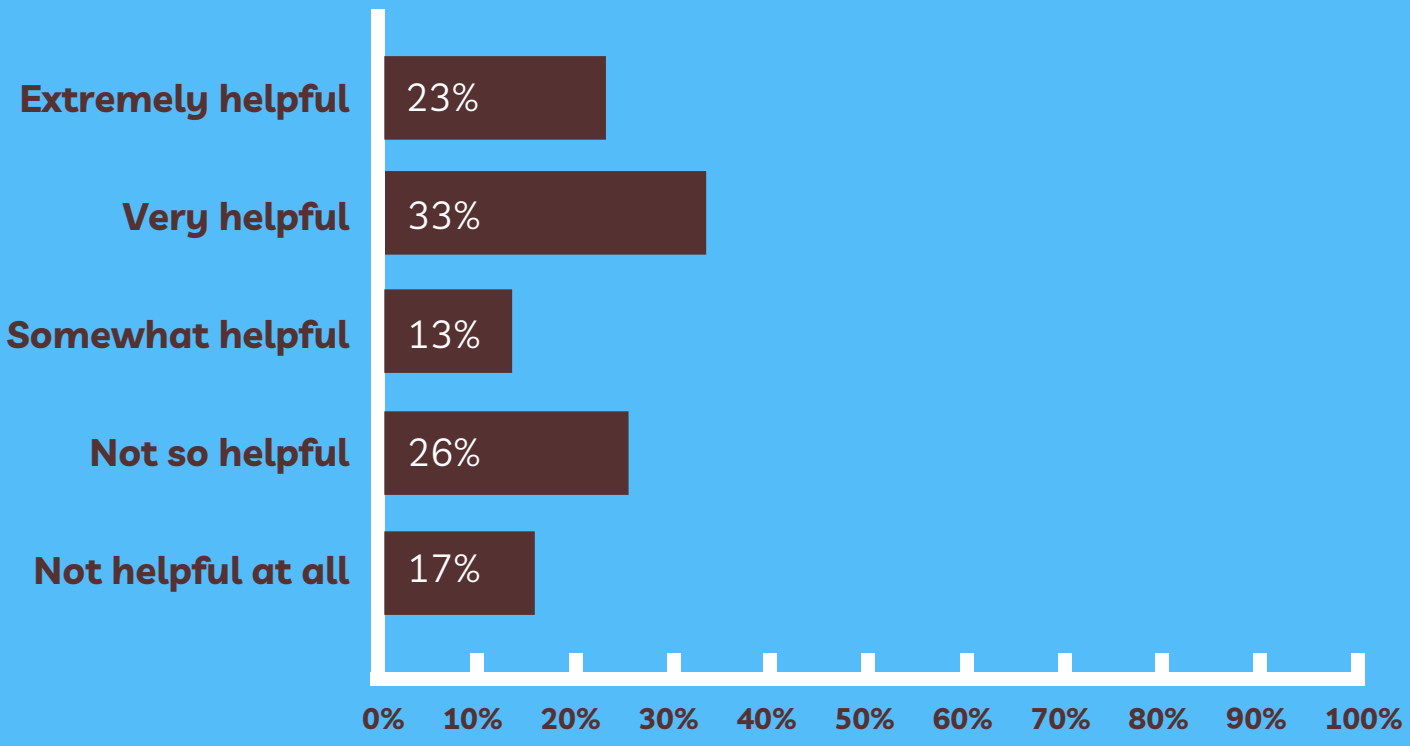

There are many regulatory sandboxes that allow AI businesses to test their innovative propositions in the market with real consumers. Under close supervision by the relevant regulator, sandboxes remove unnecessary regulatory barriers in order to promote innovation. This allows firms to test products and services in a controlled environment and reduce which can help to reduce time-to-market. Have you ever participated in a regulatory sandbox?

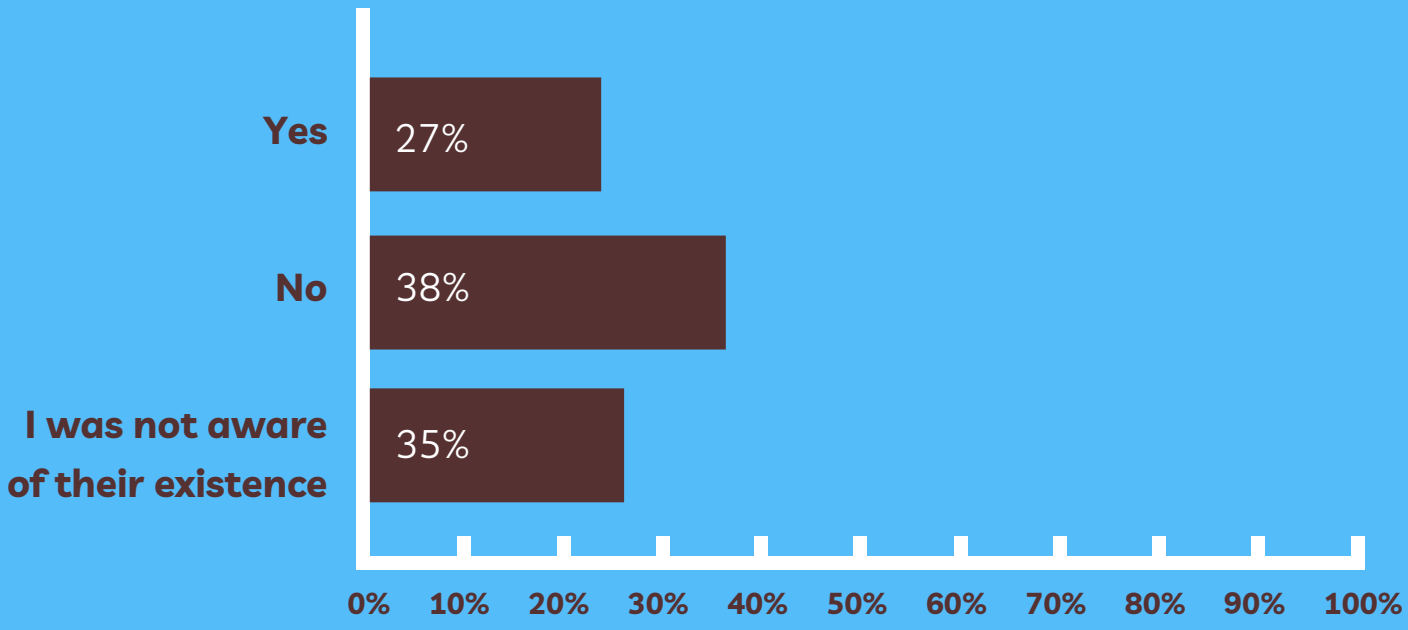

If you have participated, how important was it to developing your product?

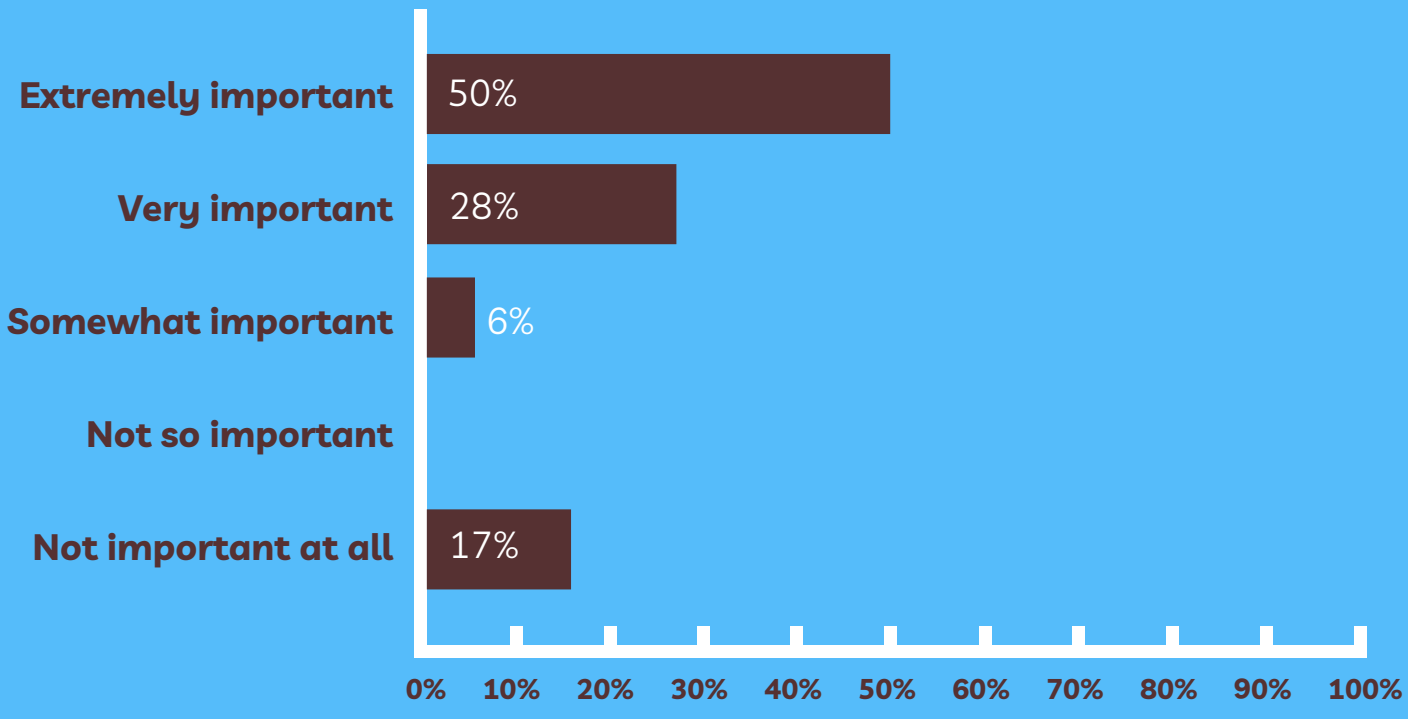

To your knowledge does your product/service cross-cut more than one regulator?

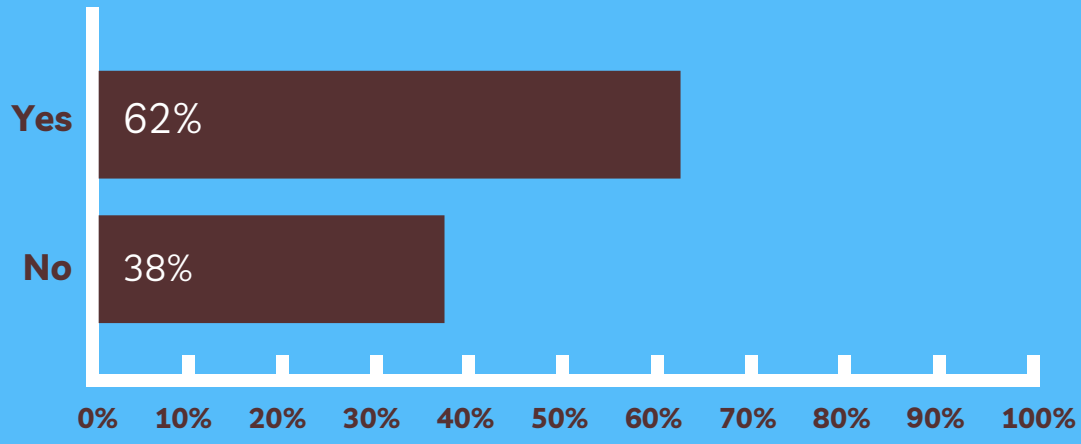

For many early stage startups, their algorithm is the product, and where and how it is deployed adapts to different markets. This makes them highly agile and ready to pivot away from a market if, for example, a regulator were to make it very difficult to iterate and bring a product to market.

62% of AI startups told us their product or service covers multiple sectors, meaning they would have to typically deal with more than one regulator. This can be confusing for a startup who can be unsure which regulator has oversight of a product or service. For emerging issues it can also mean startups get bounced between regulators. This means that for AI startups, regulators not only have to ensure they have the right technical expertise, they have to ensure they are identifying and working with colleagues across the regulatory landscape. The UK's new Digital Regulatory Cooperation Forum, established to help regulators collaborate on emerging issues, is well placed to help with these issues but so far we have found little awareness from startups about its work.

Many regulators are aware of this and are implementing different policies to help startups. One of the most popular are sandboxes. Regulatory sandboxes enable startups to test products in the market, with real consumers, in a controlled environment. At their best, sandboxes can remove burdensome regulatory barriers and help startups refine their products in a secure environment. Our survey showed a majority of AI startups, 38%, had never participated in a regulatory sandbox, more concerningly, a further 35% told us they were not aware of their existence. This is especially concerning given that of the 27% of startups who told us they had used a regulatory sandbox, 78% of them told us their use was at least either extremely or very important to the development of their product or service.

Operations and Funding

In order to fill requisite skills gaps, are the majority of your hires:

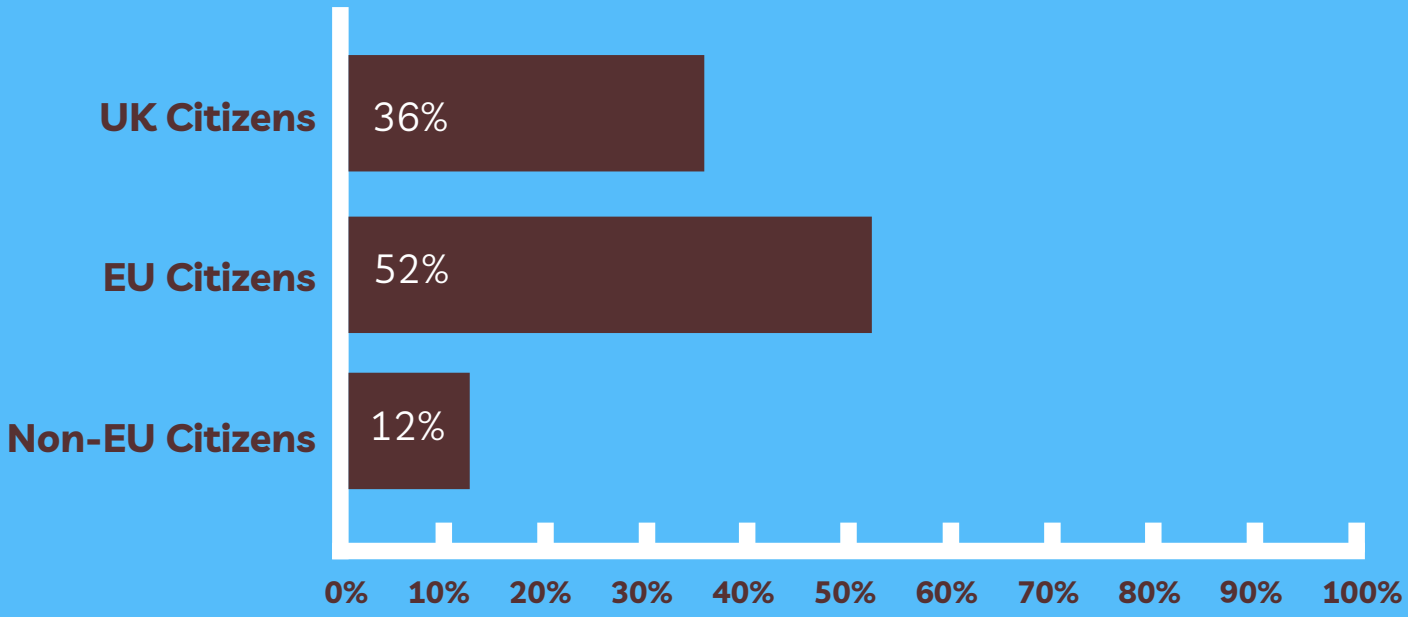

From your experience have you found the availability of private capital for AI ventures to be sufficient?

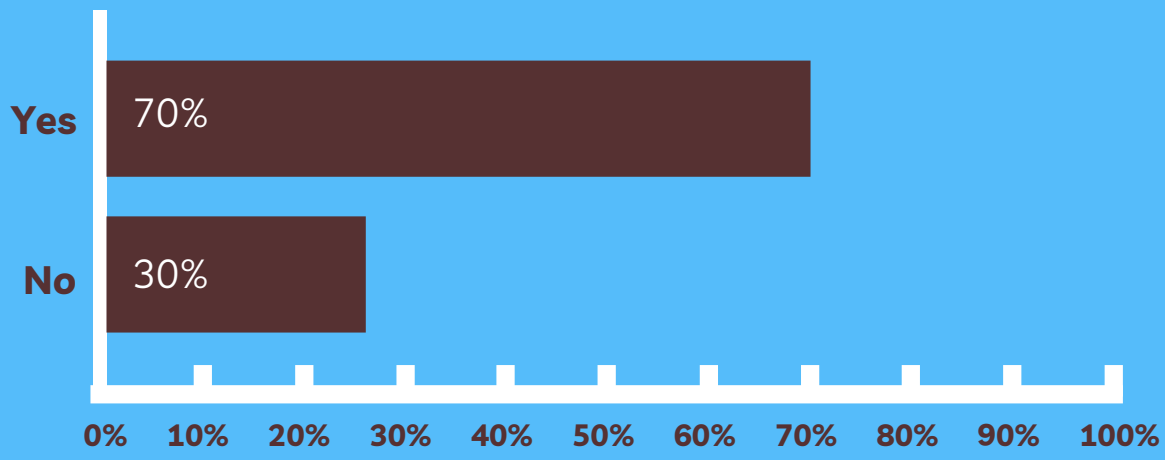

Where would you say there is a gap in funding for AI companies?

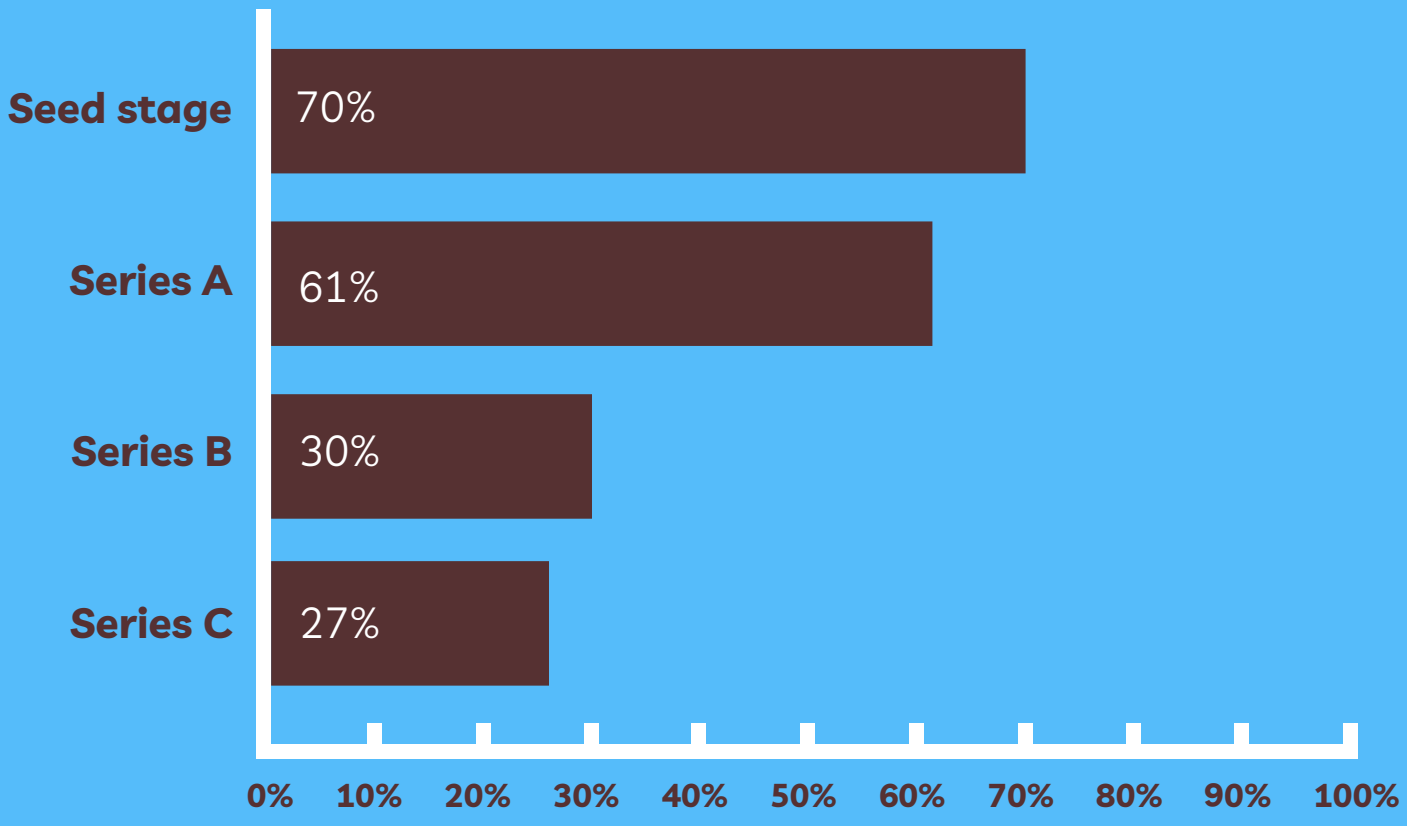

Which of the following factors would you cite as being the greatest barrier to your firms' growth over the next 5 years? (You can select more than one)

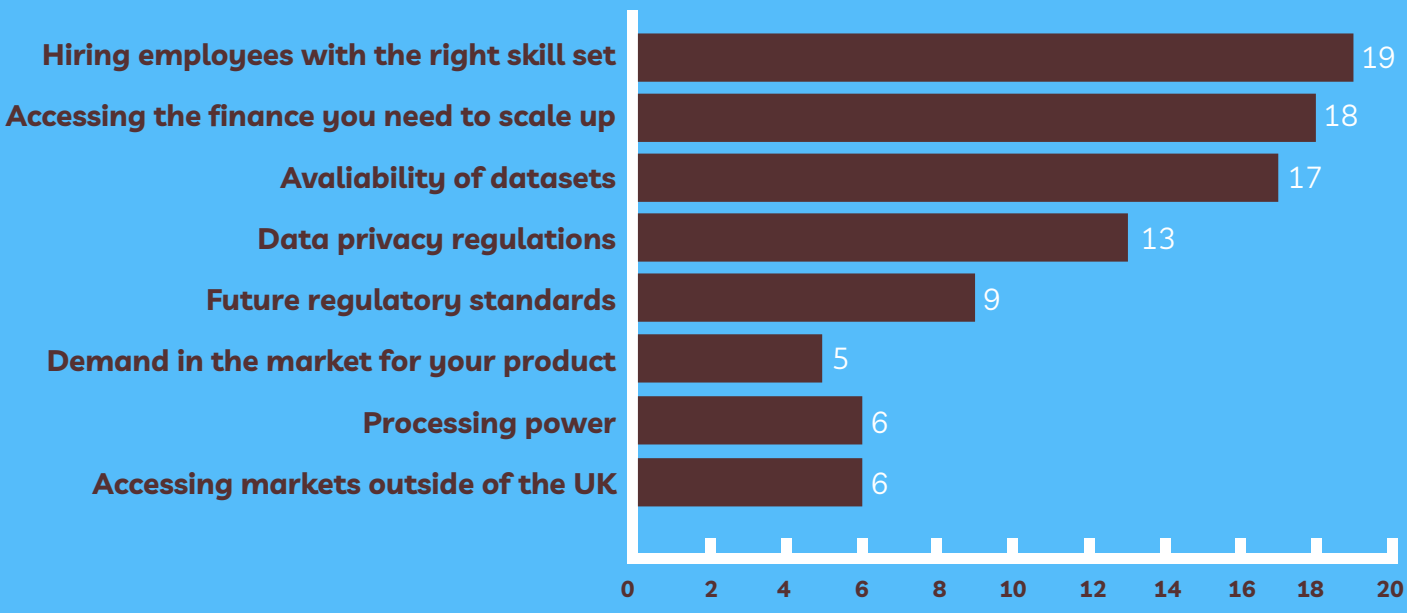

Case Studies

Startup A

Startup A is leading the way for automated transport technology. It uses vast amounts of training data, both data collected in real time and synthetic data produced in-house, to better train its deep learning models. Startup A feels broadly positive about the current direction of policy in the UK. Ensuring there is sufficient technical support for sector-specific regulators of applied AI is a priority. Startup A would like to see the UK Government focus support on sector-specific regulation and building understanding of AI for sectoral regulators such as in health, finance and transport. Describing the importance of the algorithm as 'everything' to an AI company, Startup A felt certain that algorithmic transparency in itself doesn't remove the risks of AI. We need to focus on assessing statistical outputs, reliability and validation. Startup A felt it would be better for the Government to design regulation with a focus on outcomes based assessments rather than interrogating specific decision making.

Startup B

Startup B is a revenue positive pre-seed platform that enables companies to stress test their AI systems to better identify vulnerabilities, helping the deployment of fair, trustworthy and responsible AI.

Startup B is positive about the opportunities for the UK to build on its well developed data ecosystem, though it commented that, for all the Government is clearly enthusiastic about embracing the AI ecosystem, it is unclear how this will translate into meaningful action.

Startup B agreed that many within the AI ecosystem would be following the debate on AI transparency closely. Startup B felt that mandating full transparency - including commercially sensitive information - could hurt innovation, commenting it should be possible to treat AI like a "human thinking...you cannot always see how decisions are made but you can explain the thought process".

More generally, Startup B is concerned that the policy conversation does not adequately appreciate how transparency can increase vulnerabilities. Startup B cites the malicious use of Adversarial AI to trick AI systems. Adversarial AI uses a system's own inputs against itself - publishing data sets and algorithms would enable bad actors to develop successful attacks.

More generally, Startup B felt innovative businesses too often struggle with regulators making it difficult for them to engage, as well as effectively being obstructive by moving too slowly.

Startup C

Startup C is an early stage bootstrapping startup helping people connect and organise sporting events.

They see their biggest challenge as legislative barriers, specifically mentioning that they were struggling with the unclear requirements of GDPR. The founder commented it was "very hard for bootstrappers…to get on top of everything… lawyers are expensive… Any extra barriers to how that data is used is going to make it tricker. We have a very exciting potential AI product but if the burden of legislation becomes too high… I don't want to see innovative products not get off the ground".

Startup D

Startup D is a seed stage B2B eCommerce startup that enables retailers to make videos of their products.

Startup D was broadly positive about the direction of the UK but felt that while the Government can have good ideas and schemes, these didn't always translate into benefits for startups. Startup D pointed to two recent experiences. First they were frustrated that they needed advice on the SEIS scheme but HMRC could find no one knowledgeable enough to answer the query. Startup D also made use of the global talent visa to hire a talented programmer but that the prerequisites were "incredibly hard" to fulfil for young people who have not had the chance to build up their CV. The founder commented that it sometimes felt like the Government effectively said: "Here's a great idea we had. Now go and do the idea but we can't help you do it."

Our Key Takeaways

1. **Regulations need to help, not hinder, innovation and should clearly outline the aspirational outcomes rather than prescribe the way to get there: startups have limited time, resources and expertise, and regulatory ambiguity leads to uncertainty.** The experience of GDPR demonstrated how unclear, complex and expensive regulations drove many startups out of business, and disproportionately impact startups that survived - GDPR compliance cost startups significantly more than it did the Tech Giants.[16] As we look at the future of AI regulation, it is critical we do not put AI startups through the same experience. Startups want clarity on what regulators want from them. They want to be told clearly what they can and cannot do - but they do not want to be told how they should achieve this end goal.

2. **Startup concerns are dominated by bread and butter issues:** the high cost of data are seen as critical to the immediate survivability of startups. Clearer regulations, such as making GDPR easier to understand and comply with, that help startups innovate will have the most immediate impact on their success.

3. **The EU's approach to AI regulation is not supported by the startup ecosystem:** Startups were overwhelmingly concerned that the EU's key proposals are too restrictive and will hurt innovation. This is an opportunity for the UK to cement its status as one of the most attractive places for AI innovation.

References

1. Top 50 AI companies in the UK
2. Curious Aritifical Intelligence Statistics UK
3. Top 50 AI Companies in the UK
4. Curious Artificial Intelligence Statistics UK
5. UK Tech Update - London Tech Week 2022
6. Industrial Strategy: Building a Britain fit for the future
7. Industrial Strategy- Artificial Intelligence Sector Deal
8. AI in the UK: ready, willing and able
9. AI Roadmap
10. Queen's Speech 2022
11. Data: a new direction
12. Proposal for a Regulation laying down harmonised rules on Artificial Intelligence
13. Regulatory Framework proposal on artificial intelligence
14. Europe's Artificial Intelligence Debate Heats Up
15. French Presidency: Proposal for a Regulation of the European Parliament
16. Privacy Regulation and Firm Performance: Estimating the GDPR Effect Globally