



Morgan Lewis

SILICON VALLEY FIRST CUP OF COFFEE SEMINAR SERIES

UPCOMING SEMINARS:

Artificial Intelligence (AI) Boot Camp

- January 12 Computer-Implemented Inventions in Biotechnology and Healthcare, Patentability from European and US Perspective
- January 13 M&A and Investment into AI Companies
- January 19 Software As a Medical Device: US FDA Regulatory and Legal Framework
- January 20 Patent and Trade Secret Protection for Inventions That Use AI
- January 21 AI in Hiring and Recruiting
- January 28 AI and Copyright



Morgan Lewis

SILICON VALLEY **FIRST CUP OF COFFEE** SEMINAR SERIES

UPCOMING SEMINARS:

Artificial Intelligence (AI) Boot Camp


- February 2 The Ethics of Artificial Intelligence for the Legal Profession
- February 3 AI and Data Privacy
- February 4 Patents for MedTech AI: Opportunities and Pitfalls
- February 9 IP Landscape of AI Hardware Startups
- February 11 AI in Digital Advisory Offerings: Regulatory Considerations
- February 16 Bias Issues and AI

Morgan Lewis

BIAS ISSUES AND AI

**Mike Pierides, Oliver Bell
and Andrew J. Gray IV**

February 16, 2021



Presenters



Mike Pierides



Oliver Bell



Andrew J. Gray IV

Morgan Lewis

Agenda

1 The Problem

2 Risks and Liabilities

3 Mitigating and Removing Bias


4 The Future

The Problem



Morgan Lewis

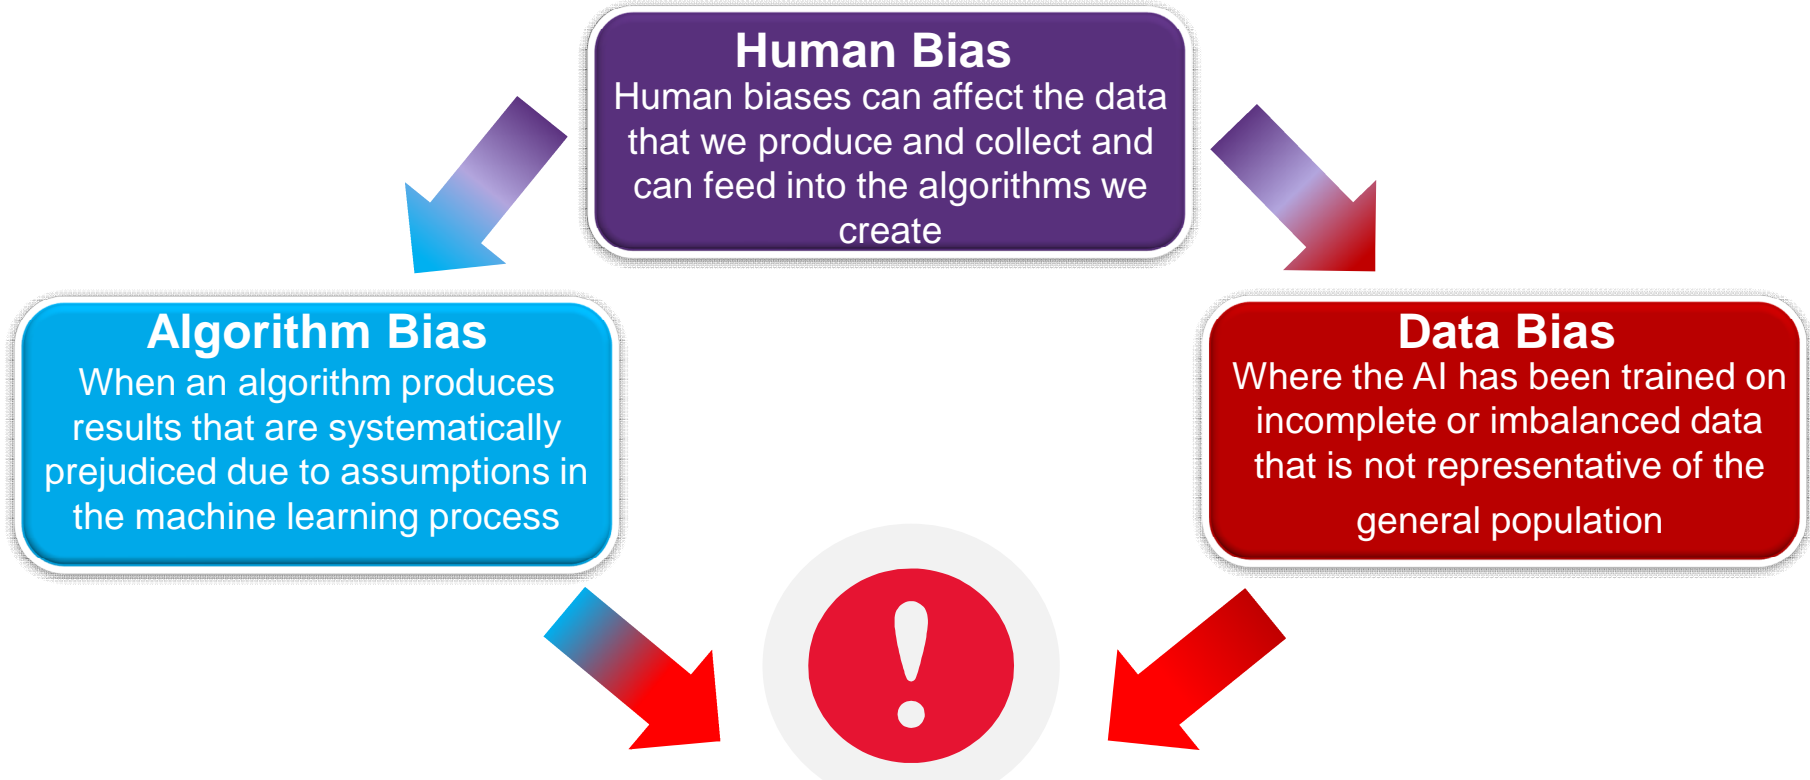
What is Bias in AI?



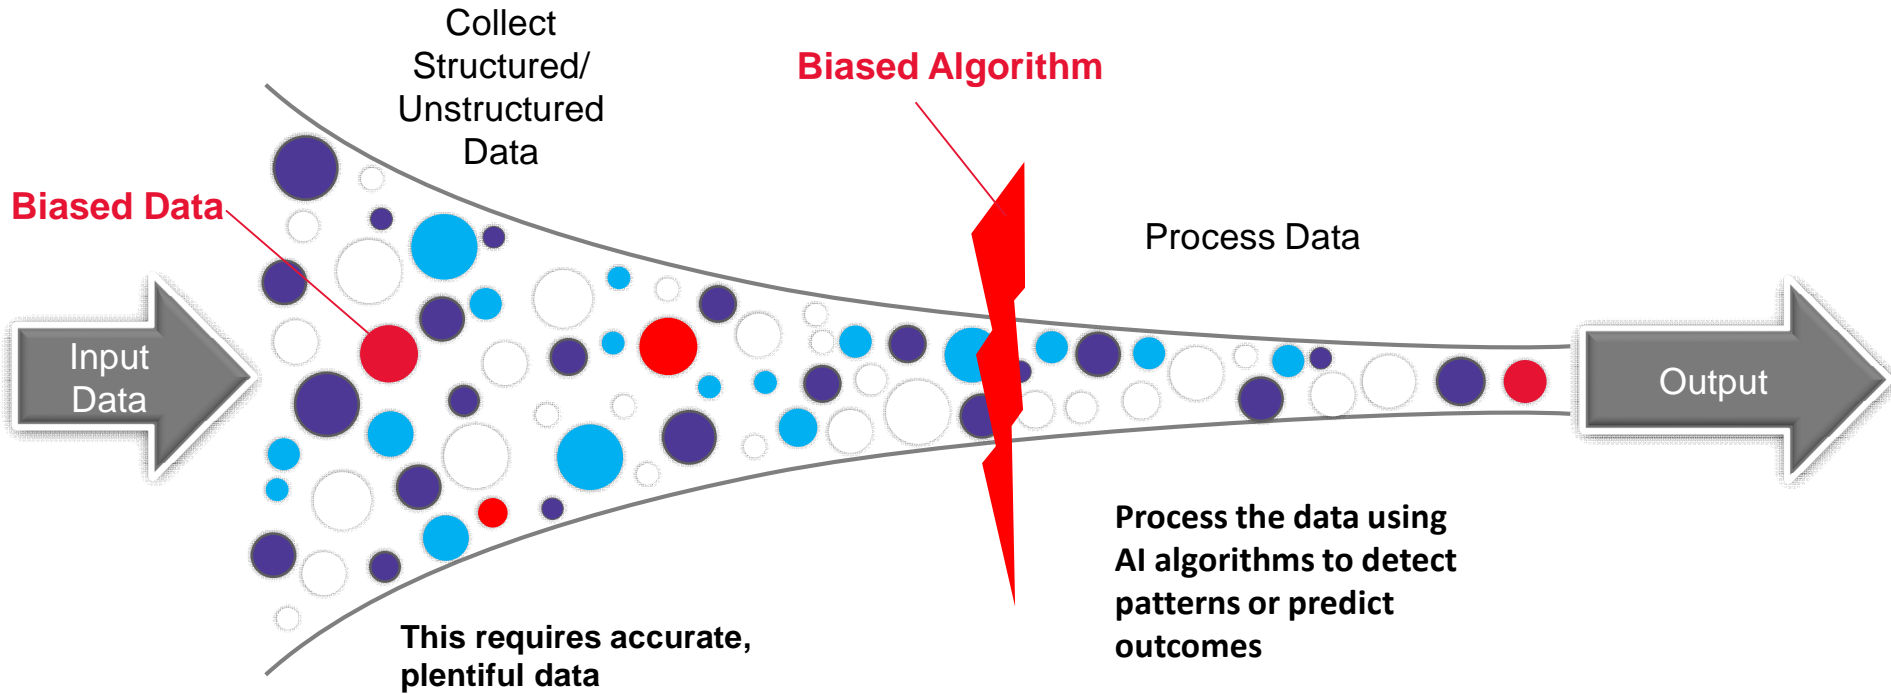
Bias in AI is a situation where an algorithm produces results that are biased/prejudiced due to assumptions in the machine learning process.

Examples include: gender, age, sexual orientation, racial.

Sources of Bias in AI



How AI works



Types of Bias

Sample Bias

The dataset does not reflect reality

Exclusion Bias

Excluding valuable data

Measurement Bias

A systematic or non-random error that occurs in the collection of data

Recall Bias

Inconsistent labelling of data resulting in lower accuracy

Confirmation Bias

The effect of seeing what you expect to see or want to see

Association Bias

The Data reinforces and/or multiplies a cultural bias

Popular Uses



**Recruitment
Processes**



**Insurance
Decision Making**



**Monitoring User
Behaviours**



**Credit
Referencing**



**Underwriting Loans,
Anti-money
Laundering and
Fraud Detection
Processes**

Popular Uses – Potential Bias

Recruitment

- Male candidates' CVs favoured
- Lower salaries offered to minorities

Insurance

- Higher premiums for the elderly
- Premiums for same policy differing due to name

Monitoring Behaviour

- Promotions offered to employees without children

Credit Referencing

- Disabled and ethnic minorities given lower credit scores

Loans, AML, Fraud

- More AML issues flagged for ethnic minorities
- Loans not extended to those living in certain areas

It's Not New!

Computer based bias is not a new issue – its potential existed for as long as computers have been programmed to make decisions

- In 1988, the UK Commission for Racial Equality found a British medical school guilty of discrimination
- The computer program it was using to determine which applicants would be invited for interviews was determined to be biased against women and those with non-European names
- The program had been developed to match human admissions decisions, doing so with 90 to 95 percent accuracy
- The issue was the algorithm, which perpetuated human bias

COMMISSION FOR
RACIAL EQUALITY



Data Bias Example

A Tech company used an AI tool to automate its recruiting process by rating applicants' resumes based on the resumes of past and current employees.



The input data used was biased - Used historical recruitment data from the last 10-years. Males made up the majority of applicants and hired employees.

The output was therefore biased - The recruiting system incorrectly learnt that male candidates were preferable. The system favored applicants based on words like “executed” or “captured” that were more commonly found on men’s resumes, and penalized resumes that included the word “women”.

Algorithm Bias Example


A healthcare provider used an algorithm to review data from over 200 million people to predict which patients would likely need extra medical care.



The algorithm was biased - The algorithm's designers used previous patients' health care spending as a proxy for medical needs. This was a bad interpretation of historical data because income and race are highly correlated metrics and making assumptions based on only one variable of correlated metrics led the algorithm to provide inaccurate results.

The output was therefore biased - The algorithm was producing faulty results that favored white patients over black patients, reducing the number of black patients identified for necessary additional care.

Why is AI Bias such a potential issue?



Some form of bias is likely to exist in a significant proportion of decision making

Is AI bias any worse than employees blindly following a corporate policy that has been based on bias?

Arguably not – But wait!

AI has the potential to increase bias issues as it can undertake decision making on a vast scale – extrapolating individual or minor bias issues into potentially significant issues with major consequences for organizations

Advantages

- ✓ It can improve on traditional human decision making
- ✓ If variables are not included, they are not taken into account – this can't be said for humans that even when told not to take something into account, may do so (intentionally or unintentionally)
- ✓ It can also be easier to probe algorithms for bias, potentially revealing human biases that had gone unnoticed or unproven
- ✓ Using AI to improve decision-making may benefit traditionally disadvantaged groups

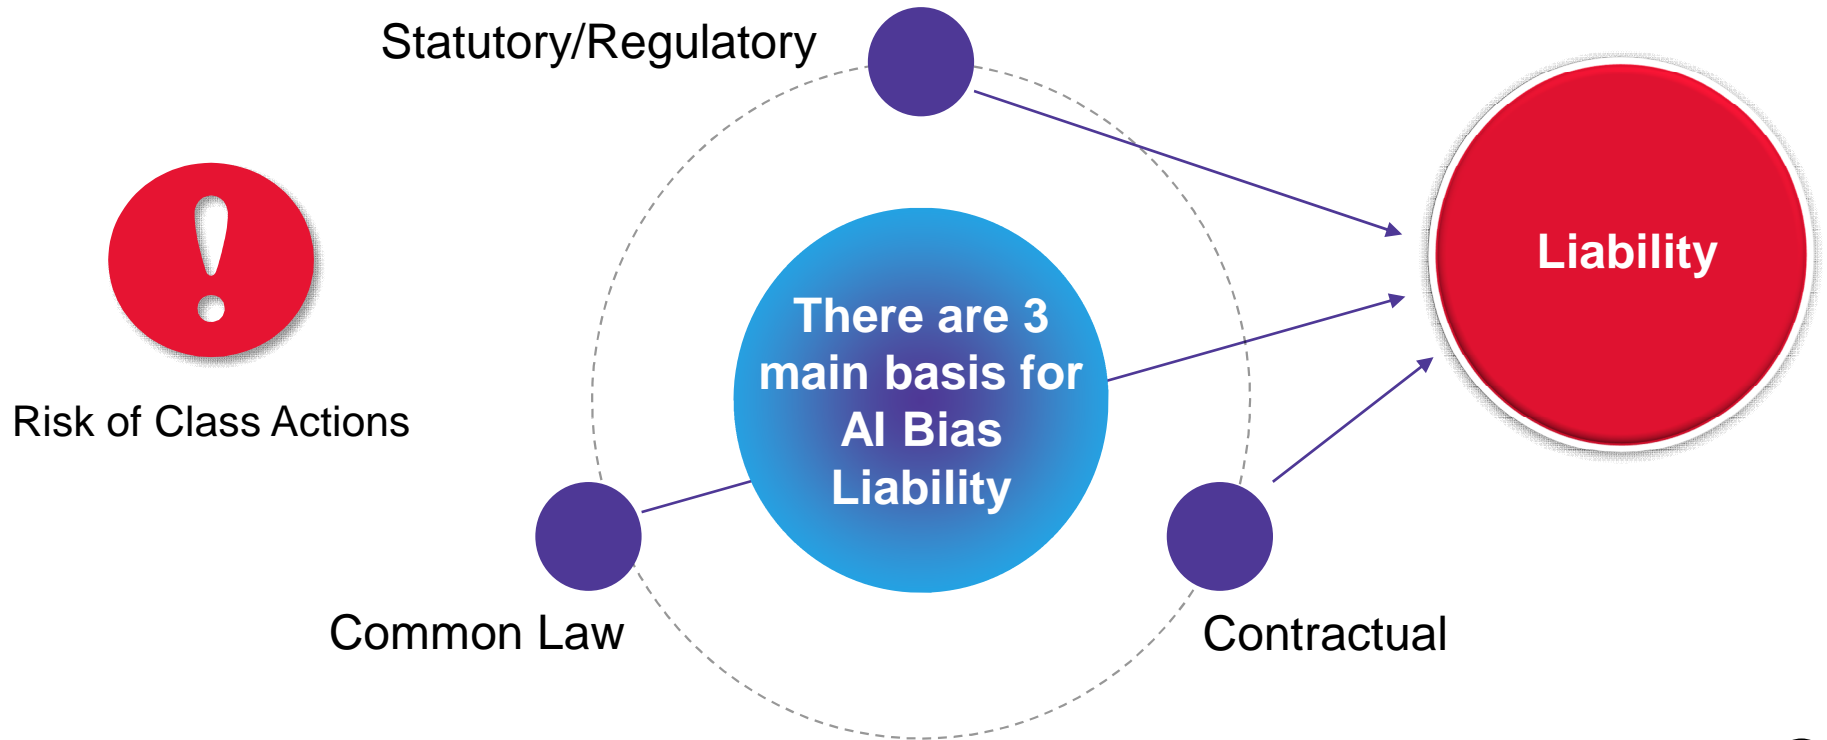
Risks & Liabilities



AI

Morgan Lewis


What is the basis for liability?



Statutory Liability - Examples

The Fair Housing Act prohibits housing-related discrimination on the basis of race, color, religion, sex, disability, familial status, and national origin.

Penalties – Compensation for discriminated persons and/or fine of up to \$65,000 for repeat breaches



Morgan Lewis

The Equality Act prohibits discrimination, in relation to nine protected characteristics: age, sex, disability, gender reassignment, marriage and civil partnership, pregnancy and maternity, race, religion or belief and sexual orientation.

Penalties – Compensation for discriminated persons

Statutory liability

Is there more
legislation /
regulation to
come?



An evolving
area – We will
come back to
this!

Common Law Liability

Tort of Negligence?


Duty of care to the claimant?

Breach of that duty?


Breach has caused harm?

Damage or loss has resulted from that harm?

Common Law Liability - Example




- We have seen that AI bias can cause issues with certain demographics receiving priority medical care



- What if this leads to serious harm or death based on decisions made by an AI tool?



- If a duty of care is owed to the patient, there could potentially be a claim for clinical negligence





Contractual Liability


Liability relating to biased AI outputs could arise under a number of contractual relationships



**Organization
and
AI Provider**



**Organization
and
Corporate Customer**



**Organization
and
Consumer**

Contractual Liability – Organization and AI Provider

Organization Breaches Contract



Warranty that input data provided by organization does not contain bias

The input data is biased and causes a biased output

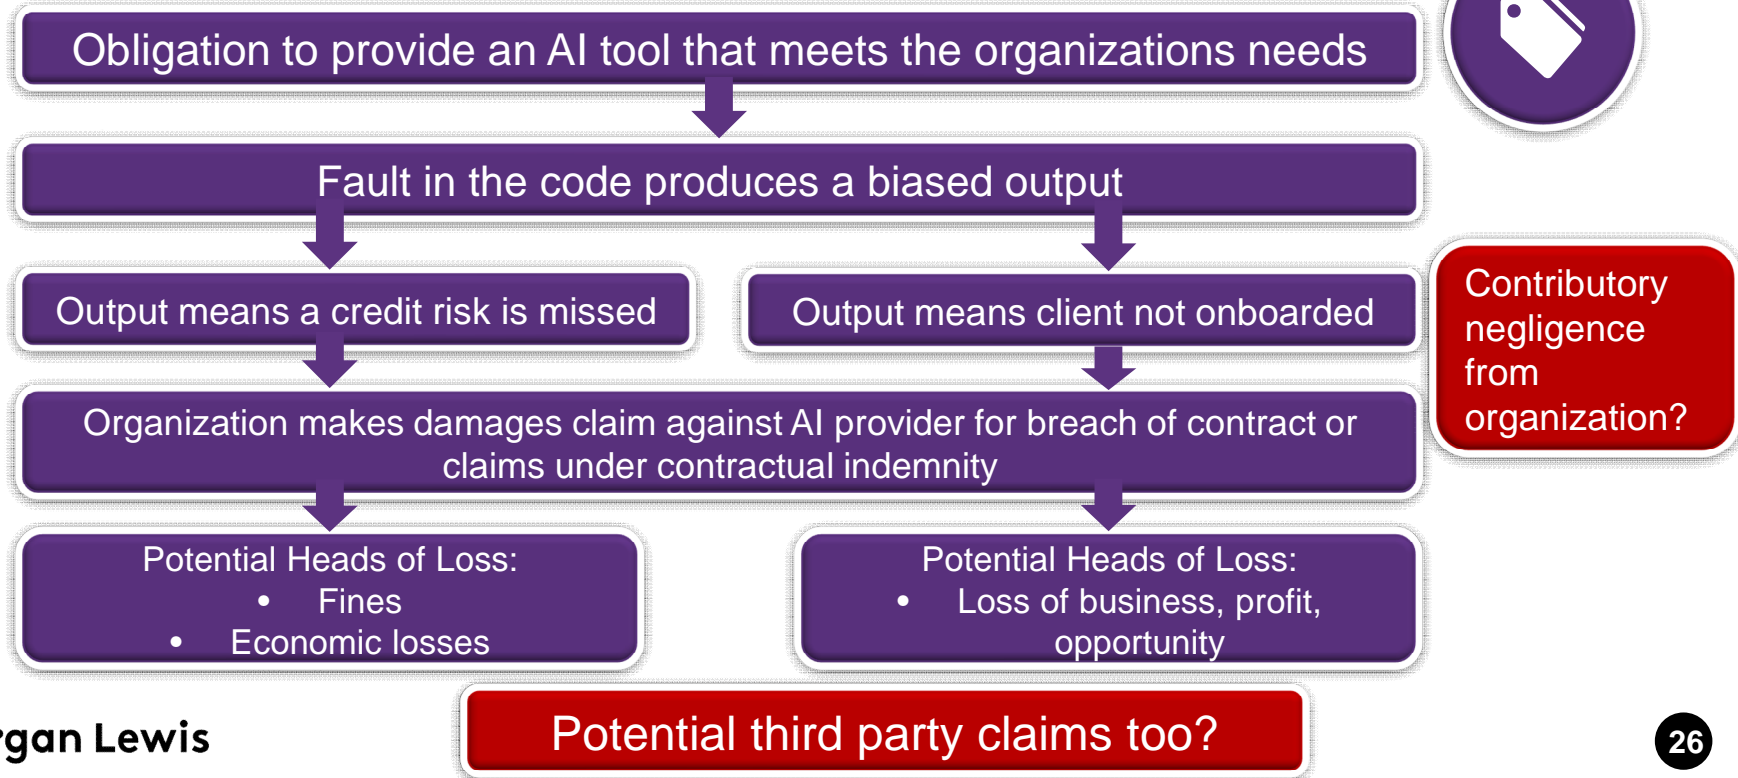
Decision made using the output discriminates against one or more persons

Third party claim brought against AI provider or AI provider suffers reputational damage


AI provider makes damages claim against organization for breach of contract or claims under contractual indemnity!!

Contractual Liability – Organization and AI Provider


AI Provider Breaches Contract



Contractual Liability – Organization and Corporate



Organization uses an AI tool as part of the provision of contracted services to a corporate client (e.g. the provision of suitable temporary staff)




The output of the AI tool is biased and results in unsuitable staff being provided

Corporate client suffers loss due to the unsuitable staff (e.g. costs of finding replacement staff and/or economic loss caused by the staff)

Corporate client makes damages claim against organization for breach of contract or claims under contractual indemnity!!

Liability
backed-off
against AI
provider?

Contractual Liability – Organization and Consumer



Organization provides insurance coverage to individual consumers



The organization uses an AI algorithm to decide on pay-outs under the policy

The AI algorithm produces bias outputs and certain consumers receive pay-outs whilst others don't for exactly the same event

Consumer makes damages claim against organization for breach of the insurance contract!!



Risk of Class Actions

Potential combined contractual and common law liabilities – for example, private medical

Contractual Liability – Exclusions and Limitations

Standard liability exclusions and limitations may be helpful




- Loss of profits
- Loss of business
- Loss of opportunity
- Indirect and consequential Loss
- Loss of goodwill
- Liability caps

Should AI specific exclusions and limitations be considered?



- No liability for decisions made based on outputs
- No liability for bias input data
- No liability for faults caused by organization's instructions/specification

Contractual Liability – Root Cause



Contractual Liability is likely
to be very fact specific

Potentially a significant
amount of litigation time
based on deciding what the
root cause of the AI bias was



Who is potentially liable?


Who has legal responsibility when an AI algorithm makes a decision that results in bias and ultimately harm?



The Organization?



The AI Provider?



The AI Itself?

Is a third party data provider involved?

Who is potentially liable? – The Organization



In the UK an All-Party Parliamentary Group on Artificial Intelligence concluded that organizations must be accountable for the decisions made by the algorithms they use

Primary Liability?

- Uses the AI tool
- May develop the algorithm itself or contract with a third party for development
- Responsible for the principles of how the AI tool works?
- May be responsible for the input data
- Makes decisions based on the outputs

Who is potentially liable? – The AI Provider



- Provides the AI tool (off-the-shelf or bespoke)
- Responsible for the code of the algorithm
- May also provide input data



Two key potential sources of liability



- Liable contractually to the customer it supplies the AI solution to?
- Vicarious liability to end users?
- The impact of decisions at the time of development may not be known – issues with liability for unknown issues!

Who is potentially liable? – The AI Itself



Can (and should) AI have a legal personality itself?

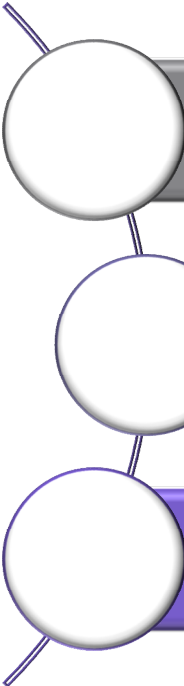
Recent case law, including the 'DABUS' decisions in the UK, EU and US, suggests not.



- AI is not a legal person and so cannot be held liable at law
- If there is harm then one or more legal persons connected to the AI must have liability – Fair?
- Some issues left open!

At present, only natural and legal persons can have liability.

Liability – Grey Areas




If an algorithm designed largely or completely by computers makes a mistake, whose fault is it?

True AI systems don't just implement human-designed algorithms, they create their *own* algorithms!

Do existing liability regimes provide for AI-related loss, or should new systems be created?

Mitigating and Removing AI Bias



AI

Morgan Lewis

Key Questions

In order to reduce or minimise bias, how do we define and measure “fairness”?

Can AI ever be completely unbiased while humans are pivotal to its development?

Can variables that potentially drive bias be removed from the start? – e.g. don't have male vs female!

Are businesses and industry leaders willing to accept the financial cost of minimizing risk?

Fairness



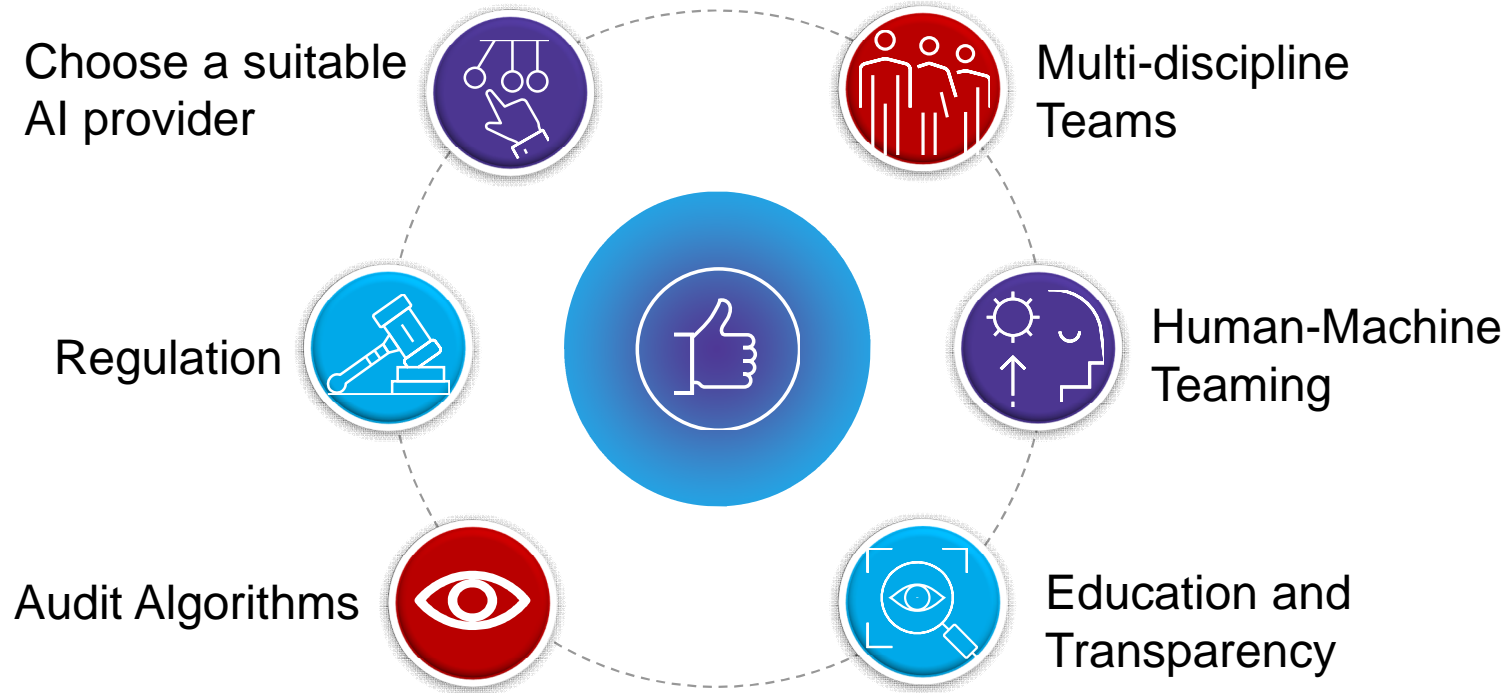
One of the most complex steps in considering AI bias is understanding and measuring “fairness”

Conflicting views can often arise – one person’s fair is another person’s unfair

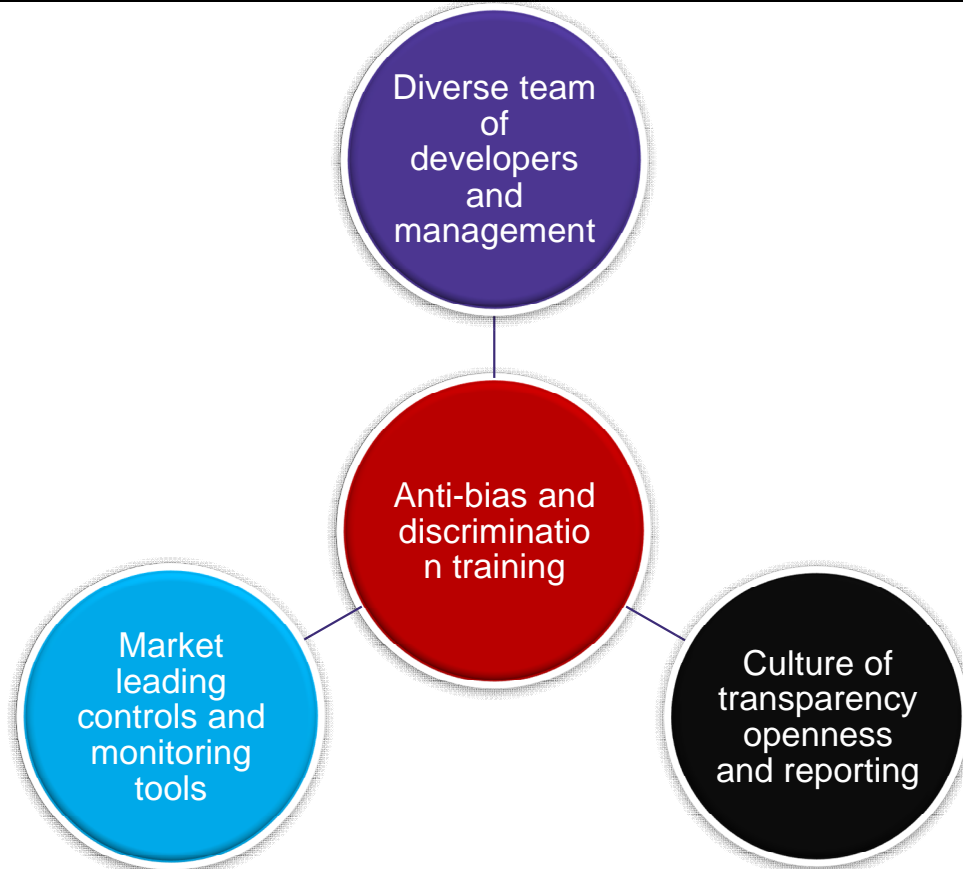
Different fairness definitions cannot be satisfied at the same time!

Even if we can agree what is fair, how should the issue be addressed?

Potential Ways to Mitigate or Remove AI Bias



Choose a suitable AI provider



Multi-Discipline Teams



Assemble a team of individuals covering multiple disciplines



The Office of AI in the UK recommends requiring AI providers to assemble teams that could include individuals that have domain expertise, commercial expertise, systems and data engineering capabilities, model development skills

Human-Machine Teaming



Though machines perform some tasks better than humans, there are traits where humans generally have the edge, such as leadership, judgment and common sense

Actively involve developers in testing the output of AI

Use AI as a tool to assist and not as a sole decision maker

Have human checks in place

Test against human outputs and analyse the results

Education and Transparency



Provide information about how the algorithms work – consider confidentiality and IP protection

Ensure employees understand how the tool works

Audit Algorithms



Examine the inputs, outputs and outcomes in a scientific way to ensure they are working as intended



Creation of new AI jobs



Trainers



Explainers




Sustainers

Regulation



- Artificial Intelligence Video Interview Act in Illinois
- More proposed legislation



No specific AI bias legislation in the UK

An emerging hot topic... Watch this Space!

Tools to Reduce Bias



'What-If' Tool

Using the What-If Tool, you can test performance in hypothetical situations, analyze the importance of different data features, and visualize model behavior across multiple models and subsets of input data, and for different ML fairness metrics



IBM OpenScale

IBM's Watson OpenScale performs bias checking and mitigation in real time when AI is making its decisions

IBM's AI Fairness 360

IBM released an open-source library to detect and mitigate biases in unsupervised learning algorithms that has currently 34 contributors (as of September 2020) on Github

Resources



Considerations for Organizations using AI

In-house development vs
third party development

Ensure that data
inputs are not biased
Own data vs third party data

Requirements to
include safeguards
in the code

Auto stop if issues identified to
stop escalating issues

Confidentiality requirements
to stop disclosure of data
inputs without consent

Keep control of AI use

Considerations for Organizations using AI

Only use AI where there are clear rules that can be followed as this ensures appropriate labels can be used

Heavy oversight of development and regular testing

Requirements to monitor outputs and override

Contractual commitments from AI providers

Considerations for AI Providers

Requirements on organization to ensure input data does not create bias results

Clear specification/requirements

Importance of testing procedures – working with the customer to ensure results are correct prior to live use

Don't ignore the issue – discuss AI bias and work with customer to come up with an AI tool that mitigates the risk – win-win for both parties

No liability for use of the results of the AI tool

Off-the-Shelf AI Products

- As the use of AI continues to increase, more off-the-shelf solutions will become available
- These solutions may be desirable for organizations, as the track record of the solution can be assessed
- Appropriate due diligence should be undertaken

Ensure the product is fit for purpose

Heavy scrutiny of legal terms – try to negotiate

Are any guarantees given about bias?

Contractual Protections

A significant amount of the risks presented by AI technologies cannot realistically be dealt with at a contractual level. However, some core issues can be addressed:

Responsibilities	<ul style="list-style-type: none">• Need to clearly set out who is responsible for issues with AI• Code issues – solely AI provider?• Raw Data Input – Organisation using the AI? Could be provided third party and/or the AI provider• Include data set parameters
Obligations	<ul style="list-style-type: none">• Consider obligations on each party and mutual obligations – need to work together to mitigate bias• Fairness to be taken into account when developing the algorithm• Monitoring of results and ability to override• Requirements for AI provider to evidence or undertake bias training for all personnel engaged and to have a diverse team

Contractual Protections


Specifications	<ul style="list-style-type: none">• Clear descriptions of the AI system's specifications, including non-discriminatory features and practices• Description of controls in place to mitigate bias
Service commitments	<ul style="list-style-type: none">• Any automated results of the AI system will be actively monitored by an employee of the AI provider• Real representative data will be made available and used to monitor the performance of the AI system
Representations and Warranties	<ul style="list-style-type: none">• Warranty that the datasets used are diverse• The AI provider represents and warrants that the AI tool is free of bias and discrimination, including as defined by any applicable law• AI tools will function and be maintained in accordance with industry standards

Contractual Protections


Indemnities	<ul style="list-style-type: none">• Indemnification obligations to cover third party claims that the AI system caused discrimination/damage/loss and any fines for breach of laws• May be one way or mutual – depending on specific circumstances and negotiating power
Liability Exclusions	<p>Consider excluding or limiting liability for certain events:</p> <ul style="list-style-type: none">• Use of the AI tool outside of a designated scope• Inputting information outside of specified fields or parameters
Rectification Plan Process	<ul style="list-style-type: none">• Include a clear process for rectifying any issues that arise and ensuring that such issues don't arise again• Agree who is responsible for the costs
Transparency and Reporting	<ul style="list-style-type: none">• Requirements to provide detailed data about how the AI tool works – this will be essential for good compliance and having transparent documentation• Obligations to ensure accurate recording keeping and reporting at all stages - a paper trail is key to show the right things were being done to avoid bias

Data Protection

GDPR requires data subjects to be informed of any automated decision making used in respect of their personal data – organizations will need to update their privacy policies to reflect their use of AI and may want to consider reputational issues when considering using AI tools for decision making!



Even where such transparency is not a legal requirement, organizations should be working to ensure transparency of data use as far as possible



Undertake Data Privacy Impact Assessments – this may be a legal requirement in certain jurisdictions (e.g. Europe and the UK)

The UK ICO has issued guidance on addressing the risks of bias in AI:

- Some of the protected characteristics outlined in the Equality Act are classified as special category data. These include race, religion or belief, and sexual orientation
- Before processing data of individuals with protected characteristics, ensure you have an appropriate lawful basis to process the data for such purposes
- Where you use biometric data for the **purpose** of uniquely identifying an individual, it is also special category data

Determine and document your approach to bias and discrimination mitigation from the beginning of any AI application lifecycle, so that you can take into account and put in place the appropriate safeguards and technical measures during the design and build phase

The Future

01010101 0101 010101 0101

0101



01010101 0101 0101

01010101 0101 010101 0101 0101 0101 01010101 0101

AI

0101

01010101 0101

01010101 0101 010101 0101

0101

01010101 0101

Morgan Lewis

Cultural Shift



- Global cultural shift – Black Lives Matter
- Focus on diversity and inclusion
- Companies having to follow suit and prioritise D&I

Reputational
Damage

Increases in Diversity and Inclusion Legislation

- Diversity and Inclusion is a current hot topic
- Proposals in the UK for changes to legislation to deal with carer's leave, workplace harassment, family leave and workplace modifications

New or updated
legislation

```
graph TD; A[New or updated legislation] --> B[Expressly covers AI Bias?]; A --> C[AI Bias caught?]
```

Expressly covers AI
Bias?

AI Bias caught?

Contracts



- Corporates starting to specifically address AI in their contracts
- Shifting risk onto AI developers
- Shifting risk onto supplier's using AI to provide services
- Updating and creating policies:
 - HR
 - IT
 - Specific AI Policies?

Currently advising on a Passenger Services System Agreement that specifically addresses AI bias

Specific AI Laws

Considering how AI
tech is being used

Wait and see?

Consider the
impact that will
have on society

US Regulation



- **Algorithmic Accountability Act of 2019**
 - Require companies to affirmatively evaluate and minimize the risks of algorithms that result in inaccurate, unfair, biased or discriminatory decisions
 - Large companies to audit their algorithms for potential bias and discrimination
- **Commercial Facial Recognition Act of 2019**
 - General ban the commercial use of facial recognition technology to "identify or track an end user" without obtaining their consent. Requirement for third-party testing
- **New York City Council - Local Law 49**
 - Address algorithmic bias and discrimination occurring as a result of algorithms used by city agencies
- **One Federal Law?** Companies starting to call for regulation – overarching federal approach potentially preferable

UK Review



Update Equality Act to reflect issues with AI Algorithms

Create national policing bodies

Information Commissioner's Office to update guidance

Mandatory transparency obligation on all public sector organizations using algorithms

Revise Product Liability Directive

- Expand the definition of “products” to include digital content
- Amend the scope of “damages”, “defects” and “producers”
- Potentially change to a Regulation, rather than a Directive

Maximum Compensation

- €2 million in case of death or harm to a person’s physical health or integrity resulting from an operation of a high-risk AI-system;
- €1 million in case of harm that results in economic loss or damage to property

Limitation Periods

- 30 years for claims concerning harm to life or health; and
- 10 years in cases of property damage or harm that results in economic loss

Strict Liability

- The operator of a “high-risk” AI-system “shall be strictly liable for any harm or damage that was caused by a physical or virtual activity, device or process driven by that AI-system”

Insurance

- Publicly funded compensation mechanisms are not an adequate answer to the rise of AI.
- Potentially consider mandatory liability insurance in the future

Next Steps

- The European Commission’s legislative proposal is expected to be issued during the first quarter of 2021

Increase in Class Actions?



The risks associated with AI bias lend themselves to potential class actions



Class action against technology company for racial discrimination

Coronavirus COVID-19 Resources

We have formed a multidisciplinary **Coronavirus/COVID-19 Task Force** to help guide clients through the broad scope of legal issues brought on by this public health challenge.

Morgan Lewis

To help keep you on top of developments as they unfold, we also have launched a resource page on our website at

[www.morganlewis.com/
topics/coronavirus-
covid-19](http://www.morganlewis.com/topics/coronavirus-covid-19)

If you would like to receive a daily digest of all new updates to the page, please visit the resource page to [subscribe](#) using the purple “Stay Up to Date” button.



Biography



Mike Pierides

London

+44.20.3201.5686

mike.pierides@morganlewis.com

Mike Pierides' practice encompasses a wide breadth of commercial and technology transactions. Mike advises on major outsourcings, strategic restructurings following divestments or acquisitions, and technology-specific transactions such as licensing and "as a service" arrangements. He is also active advising on new technologies such as blockchain and artificial intelligence.

Biography



Oliver Bell

London

+44.20.3201.5449

oliver.bell@morganlewis.com

Oliver Bell focuses his practice on large-scale IT and business process outsourcing arrangements. Oliver advises multinational clients on all aspects of their sourcing requirements from initial scoping of requirements through to negotiation, completion, and day to day contract management. He also advises clients on the disaggregation and exit of complex agreements.

In addition to his outsourcing services, Oliver advises clients across a number of industries, including financial services, leisure, retail, automotive, and the public sector. Oliver advises on and negotiates a wide range of commercial arrangements, such as supply of goods and services agreements, warehousing and distribution agreements, agency agreements, wholesale agreements, concession agreements, intellectual property licenses, and sponsorship arrangements.

Biography



Andrew J. Gray IV

Silicon Valley

+1.650.843.7575

andrew.gray@morganlewis.com

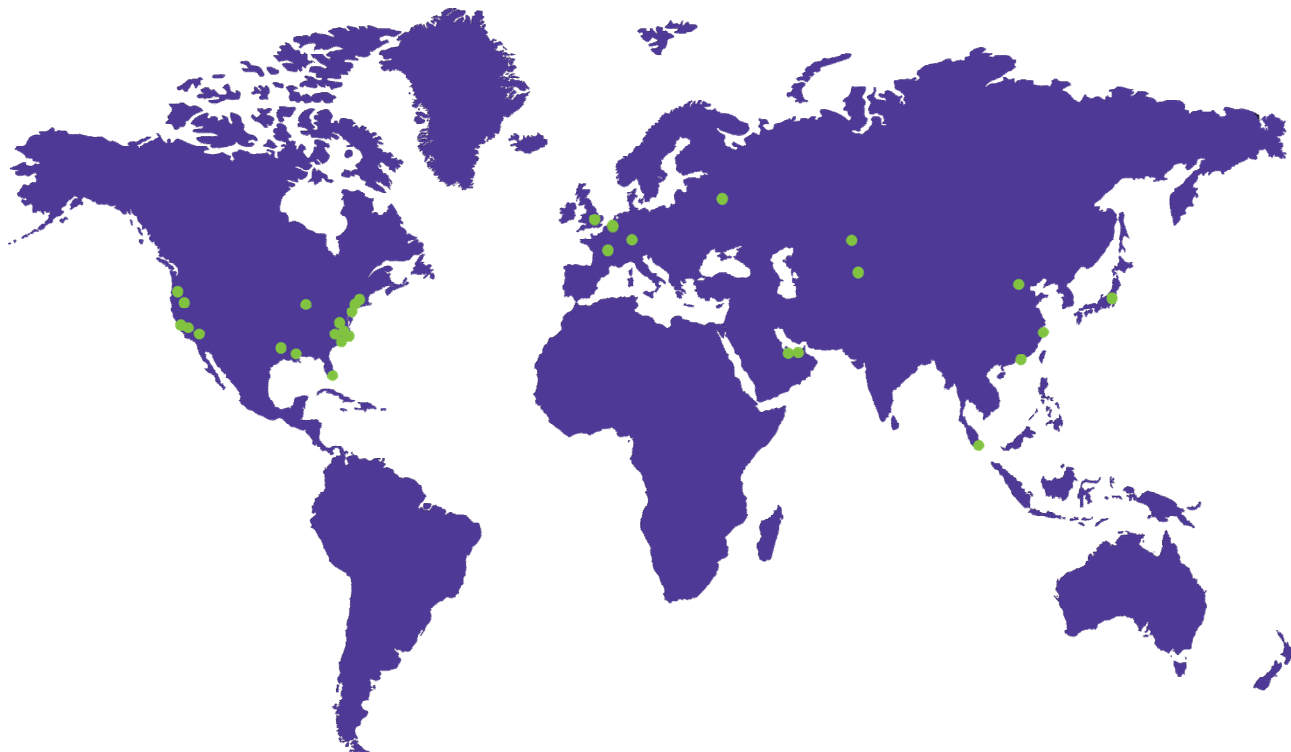
Serving as the leader of Morgan Lewis's semiconductor practice and as a member of the firm's fintech and technology practices, Andrew J. Gray IV concentrates his practice on intellectual property (IP) litigation and prosecution and on strategic IP counseling. Andrew advises both established companies and startups on Blockchain, cryptocurrency, computer, and Internet law issues, financing and transactional matters that involve technology firms, and the sale and licensing of technology. He represents clients in patent, trademark, copyright, and trade secret cases before state and federal trial and appellate courts throughout the United States, before the US Patent and Trademark Office's Patent Trial and Appeal Board, and before the US International Trade Commission.

Our Global Reach

Africa
Asia Pacific
Europe
Latin America
Middle East
North America

Our Locations

Abu Dhabi
Almaty
Beijing*
Boston
Brussels
Century City
Chicago
Dallas
Dubai
Frankfurt
Hartford
Hong Kong*
Houston
London
Los Angeles
Miami
Moscow
New York
Nur-Sultan
Orange County
Paris
Philadelphia
Pittsburgh
Princeton
San Francisco
Shanghai*
Silicon Valley
Singapore*
Tokyo
Washington, DC
Wilmington



Morgan Lewis

Our Beijing and Shanghai offices operate as representative offices of Morgan, Lewis & Bockius LLP. In Hong Kong, Morgan, Lewis & Bockius is a separate Hong Kong general partnership registered with The Law Society of Hong Kong. Morgan Lewis Stamford LLC is a Singapore law corporation affiliated with Morgan, Lewis & Bockius LLP.

THANK YOU

© 2021 Morgan, Lewis & Bockius LLP
© 2021 Morgan Lewis Stamford LLC
© 2021 Morgan, Lewis & Bockius UK LLP

Morgan, Lewis & Bockius UK LLP is a limited liability partnership registered in England and Wales under number OC378797 and is a law firm authorised and regulated by the Solicitors Regulation Authority. The SRA authorisation number is 615176.

Our Beijing and Shanghai offices operate as representative offices of Morgan, Lewis & Bockius LLP. In Hong Kong, Morgan, Lewis & Bockius is a separate Hong Kong general partnership registered with The Law Society of Hong Kong. Morgan Lewis Stamford LLC is a Singapore law corporation affiliated with Morgan, Lewis & Bockius LLP.

This material is provided for your convenience and does not constitute legal advice or create an attorney-client relationship. Prior results do not guarantee similar outcomes. Attorney Advertising.