NORTON ROSE FULBRIGHT

Key Issues and Considerations when Contracting for AI and AI-Enabled Services

SBOT Corporate Counsel Power Hour

Sean Christy & Chuck Hollis

September 12, 2023

Norton Rose Fulbright US LLP

Agenda and overview

1. Stage setting
2. State of regulation in the US
3. Areas of liability and risk
4. Contractual risk mitigation
5. Other mitigating steps

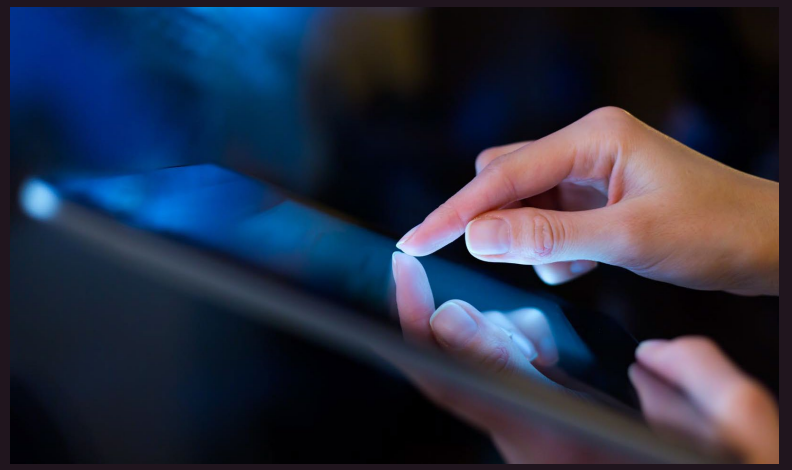

Stage settings

Key terms

Stage setting

What makes AI problematic?

Its complexity means AI system likely to be provided to companies that deploy it (users) by an expert third party service provider (provider)

Socio-technical systems:
AI technology + people + processes

Some AI can continue to learn (and change the model parameters) from the input data it is processing in the field, becoming autonomous unless monitored or constrained

Where the Model is fixed, the world can change and the outputs become inaccurate, so monitoring necessary

The mathematical relationships the AI uses to predict/create output range from easy to describe and follow, and/or reflective of our accepted real world causality, to so complex most humans cannot understand them and/or seemingly at odds with our accepted real world causality

Using complex and uninterpretable AI may not be suitable in a number of situations (e.g., where opportunities will be denied and decision subjects cannot alter their behavior to obtain such opportunities because the determinative factors are not known)

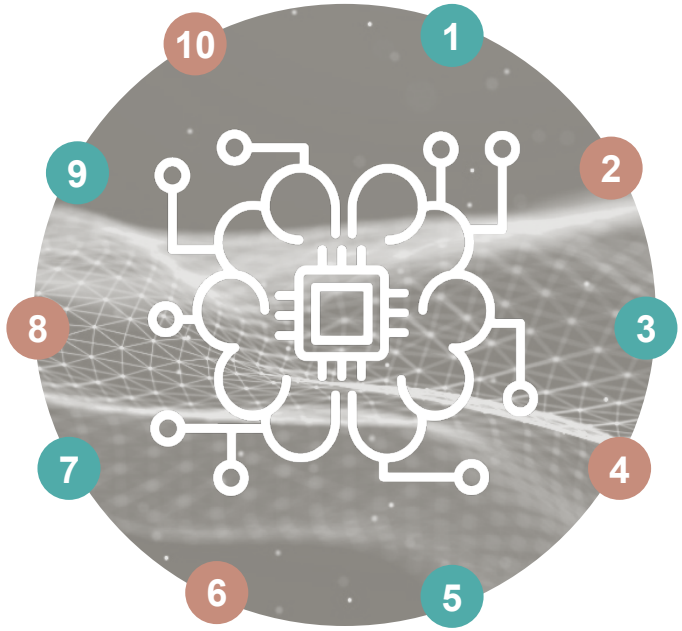

Digitization and automation of service delivery (with or without AI) is making ex ante human oversight and intervention unattractive as slow and expensive

Wide rapid roll out: impact many

Historic data (including relationships) – model (including relationships) – input data – output data (new content/prediction based on relationship)

Human may never have detected relationships before - may/may not be stronger/more reliable than the reasons humans have seen or believed caused the desired outcome

The model does not explain why the relationship occurs in the historic data: this needs to be interpreted and described by a human

If historic data was non-representative or biased the model will reflect this (without intervention)

Without human intervention the model will have no safety, legal or morality constraints:

- if AI system output used without further human checking/interpretation and without the ability to retake a flawed decision – all safety, legal and morality constraints will need to be coded in
- if good human oversight and understanding of how the AI works and how accurate it is, fewer constraints could be coded in as the human could introduce them (i.e. override/disregard AI prediction) when required (if alert and diligent)

State of regulation in the US

Federal action and legislation

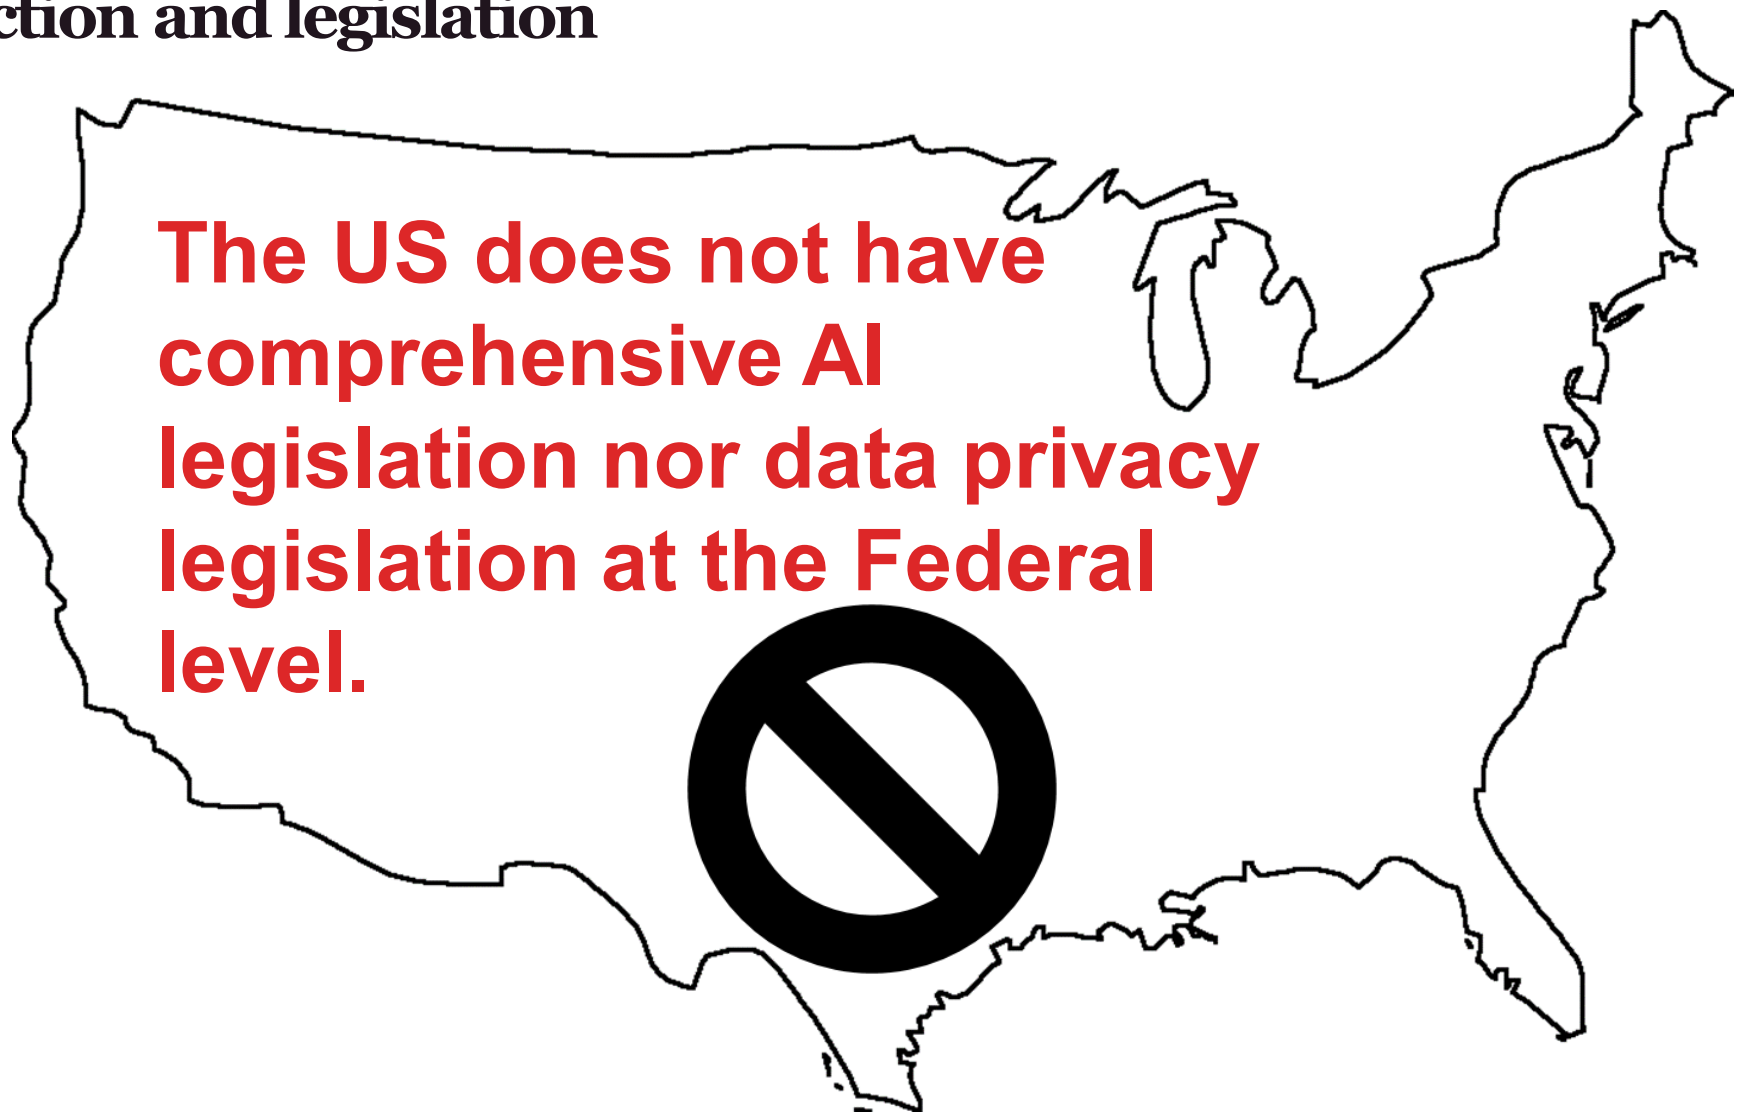

State of regulation in the US

- There have been attempts at legislation at the Federal Level:
  - Algorithmic Accountability Act of 2022
  - American Data Protection and Privacy Act – (ADPPA) (2022)
- Federal Level Guidance
  - Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People (October 2022)
  - Executive Branch Action (DOL, EEOC, DOJ, FTC, CFPB, DHHS, SEC)
    - NIST – AI Risk Management Framework
    - National AI Research Institutes

State of regulation in the US

Where does that leave us?

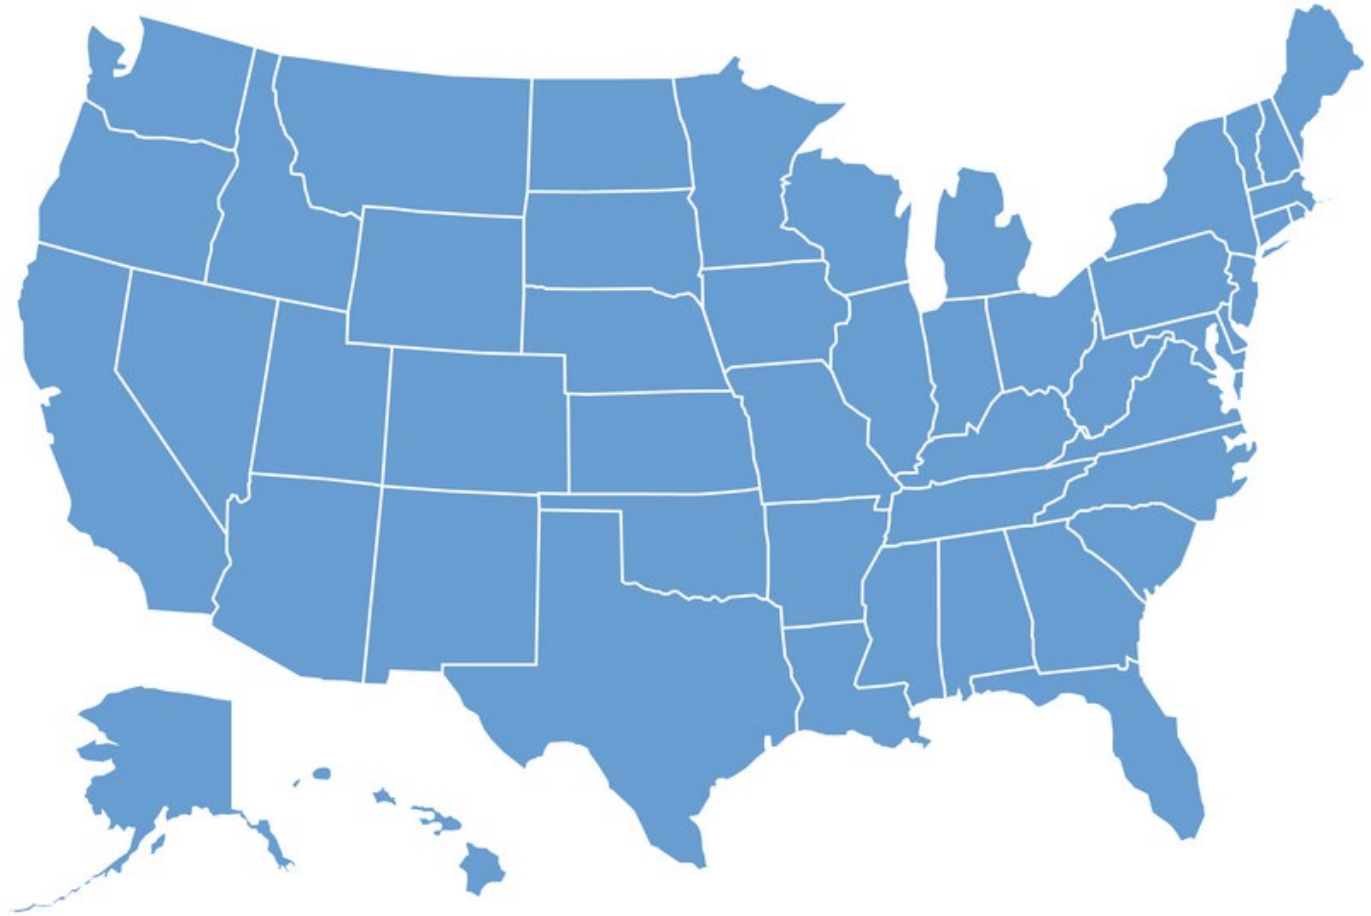

We are left with a patchwork of various State and local laws and actions.

State of regulation in the US

State laws and regulations

The State Legislatures have been active in 2023.

At least 25 states and territories have introduced bills or regulations related to AI (excluding facial recognition and autonomous vehicles).

State of regulation in the US

Texas legislation

- State legislation specifically focused on AI thus far is narrowly focused:
  - AI Advisory Council (Texas Session Laws, 2023 Tex. Gen. Laws H.B. 2060)
    - Establishes an AI advisory council to study the use and deployment of AI by Texas state agencies and to determine the need for ethics oversight and future regulation
  - True Source of Communication (Texas Statutes, Tex. Elec. Code § 255.004)
    - Prevents the use of deep fakes to influence elections
  - Sexually Explicit Deep Fakes (Texas Session Laws, 2023 Tex. Gen. Laws S.B. 1361)
    - Criminalizes the production and distribution of sexually explicit deep fakes
- New application of older laws:
  - Texas Capture or Use of Biometric Identifier Act (CUBI) is being used by the Texas AG to go after companies (e.g., Facebook and Google) who use biometric information to train AI models

State of regulation in the US

State data privacy laws

Some State Data Privacy Laws provide for the right to "opt out" of "profiling" or other "automated decision-making," and require business to conduct data privacy impact assessments. Others have AI relevant provisions.

- Texas Data Privacy and Security Act (TDPSA)
- California Consumer Privacy Act (CCPA) / California Privacy Rights Act (CPRA)
- Colorado Privacy Act
- Connecticut Data Privacy Act (also limited sourcing of AI by State)
- Virginia Consumer Data Protection Act
- Others (Indiana, Montana, and Tennessee)
- Other States have pending / failed legislation that follow these acts

State of regulation in the US

Be mindful of local laws and regulations

New York City (Local Law 144) – Prohibition on Automated Employment-Decision Tools – Requires the technology to undergo a bias audit each year. (Effective in 2023).

European Union

Where do we look for what the future may hold in the US?

- EU AI Act
- Classification system to determine level of risk
  - unacceptable / banned (generally prohibited except in rare circumstances)
  - high risk (significant governance obligations and assessments)
  - Generative AI (requires compliance with transparency requirements
  - limited risk and minimal risk (more limited requirements, e.g., monitoring / transparency)

Intellectual property

- IP ownership of AI-generated outputs:
  - Potential loss of copyrightability for AI-generated works
  - Potential loss of patent protection for AI-generated inventions
  - In all cases, “some level” of human authorship or invention is required
- Infringement exposure
  - Training data exposure
  - General patent infringement exposure in any emerging technology
- Inadvertent disclosure of proprietary information

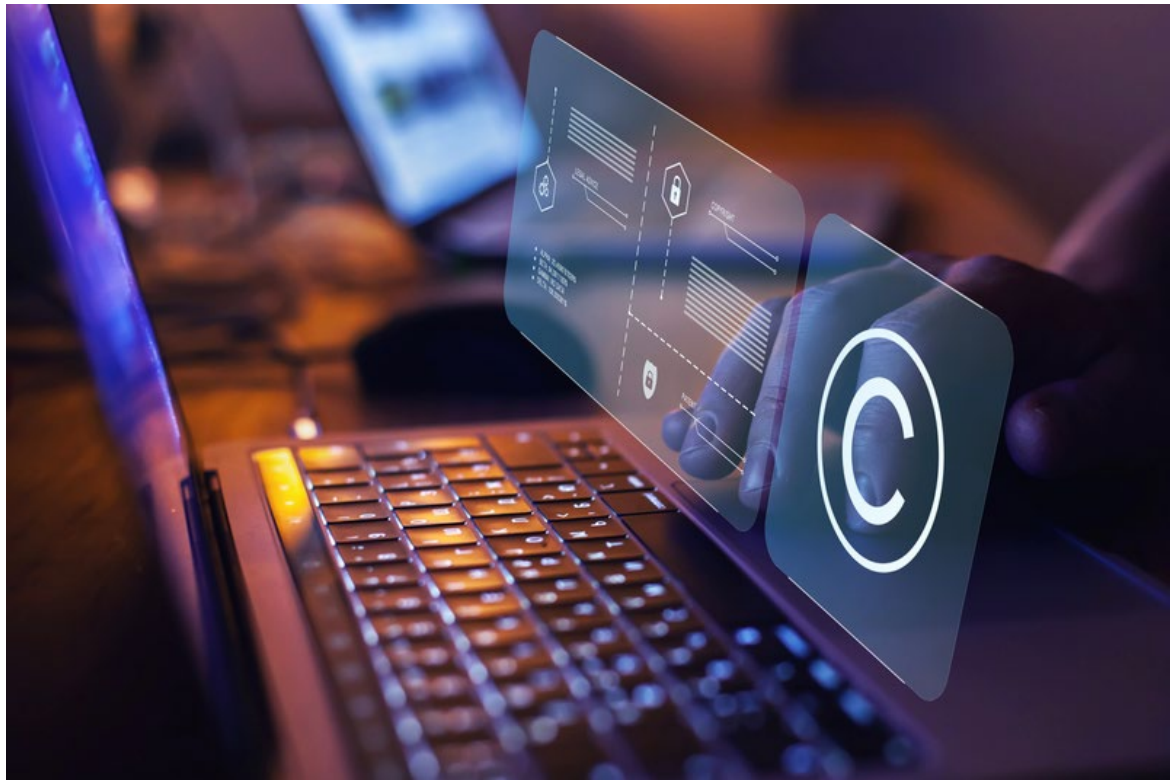

Discrimination and bias

- One of the primary focuses of new laws and regulations
- Existing laws and regulations can and are being applied
- Training data is often biased, resulting in biased outputs
- Algorithmic coding also often reflects inherent biases in coders
- Bias can occur in various use cases, for example:
  - Employment screening and decisioning
  - Creditworthiness determinations
  - Healthcare and clinical determinations
  - Tenant screening

Data privacy

- Inadvertent or unlawful disclosures
  - Use of personal data in generative AI prompts without appropriate consent
- Unlawful data collection
  - Use of personal data for algorithmic training without consent
- Management of opt-outs and deletion requests
  - For automated decision making
  - For algorithmic training
  - Practicality of “how” with AI as a black box
- Data breaches

Hallucinations

- Outputs may or may not achieve the intended or desired result
- Outputs may be inconsistently accurate or inaccurate
- Like all systems, there is a need for audit and verification

Contractual risk mitigation

Generally

- While the application of the technology is new (relatively), the contractual solutions to mitigate risk largely will seem familiar
- Guiding principle: Allocate responsibility and risk to the party in the best position to perform or mitigate
- Beware that the market trends provider friendly when it comes to contract terms
- Provider-friendly market bias can only be mitigated by effective sourcing strategy:
  - Competitive selection
  - Prioritization of contractual terms with solution and commercial terms
  - Proper sequencing of selection, negotiation, implementation and subscription/contracting
- The contract is not a substitute for due diligence and ongoing governance and oversight

Contractual risk mitigation

Representations, warranties and covenants

- The familiar:
  - Conformity with specifications and documentation
  - No material adverse changes
  - Compliance with laws
  - Non-infringement (may be more attainable in managed services vs. product/platform purchases or subscriptions)
  - Related termination, refund and indemnification/recovery remedies
- The more nuanced:
  - Transparency and explainability
  - Anti-discrimination
  - Historical and ongoing training (alignment with purpose of use)
  - Provider consents for ongoing training data
  - Use (or prohibition of use) of customer data for training
  - Responsible and ethical use by customer
  - Human oversight by customer

Contractual risk mitigation

Intellectual property

- Allocation of ownership and license rights in AI platform or services
- Allocation of ownership and license rights (if any) in AI-generated outputs
- Allocation of ownership and license rights in derivative works of the AI solution
  - Does training an AI solution create a derivative work?
- Rights of use of customer data for ongoing training and algorithmic improvement
  - Consider whether appropriate consents have been given for same
- Open source considerations
  - Limitations of open source licenses should be a red flag for corporate deployment of AI licensed on an open source basis
  - Increased training data risk
  - Less technical and operational support than commercially licensed solutions

Risk allocation

- Indemnities for:
  - IP infringement
    - Customary indemnity for infringement of patent, copyright, trade secret and any other IP or proprietary rights
    - Be mindful that scope is sufficiently broad to cover potential copyright claims related to historical training data
    - Protection for provider's use of customer data in ongoing training
  - Data breaches, data subject claims related to training data and breach of data privacy laws and regulations resulting from the provider's breach or other acts or omissions (or customer's failure to obtain appropriate consents)
  - Proper usage of the platform or services that results in a customer violation of applicable law
  - Improper use or decisioning / reliance by customer
- Exclusions from the limitations on types and amounts of damages recoverable for:
  - The indemnities referenced above
  - The customer's own damages arising in connection with the triggering events for the indemnified claims
- Data-related damages will usually be subject to a separate, higher damages cap and limited in type
  - Consider, though, that historical training data claims should be carved out and unlimited because those claims are more akin to IP infringement, solely within the provider's control and not subject to the same "data breaches are our new normal" line of reasoning
  - Similarly, data claims related to the customer's failure to obtain appropriate consents to use ongoing training data may also trigger claims
- Market practice dictates that claims for breaches of laws may also be subject to a higher cap (outcome here is driven by risk and negotiation leverage)

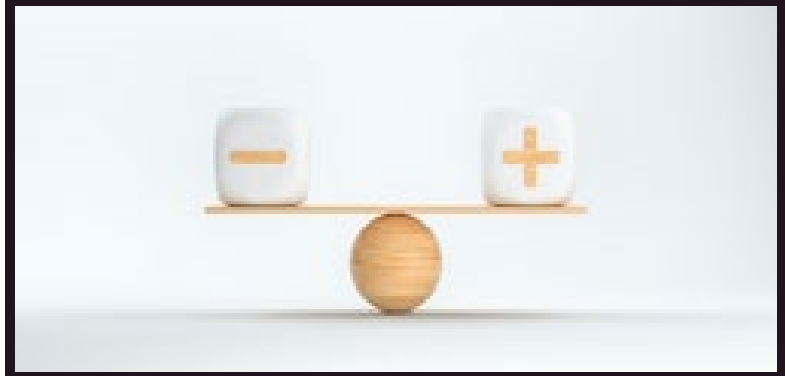

Audit and oversight

- The familiar:
  - Ability to audit compliance with the contract
  - Regulatory audit rights for regulated customers
- The more nuanced:
  - Regulatory audit rights must now extend to all customers in view of the regulatory focus and attention on AI generally
  - Transparency and explainability audits
    - Tying outputs to inputs and logic
    - Explaining automated decision making to data subjects as required by privacy and other laws
  - Providers may resist these audits on confidentiality grounds, but:
    - The provider must be able to demonstrate the input to output logic and use cases in a way that enables the customer to satisfy its explainability obligations; or
    - The parties can agree to heightened confidentiality protections (clean room, escrow, etc.)

Other mitigating steps

The non-contractual

- Provider due diligence:
  - Reputation, customer references, etc.
  - Financial (ability to back potential claims)
  - Litigation (e.g., claims against provider for infringement, training data issues, discrimination/bias, etc.)
  - Licensing models (commercial only or commercial and open source)
  - Training data assessments (lineage, quality, bias analysis)
- Solution choice:
  - Some providers offer different tiers of service (higher risk/lower cost vs. lower risk/higher cost)
    - Example: ChatGPT vs. ChatGPT Enterprise
  - Make sure to take advantage of what is marketed
    - For example: Microsoft to defend customers on AI copyright challenges | Reuters

Other mitigating steps

The non-contractual

- Corporate policies and procedures:
  - Approval of the use and deployment of AI
  - Limitations on data furnished to AI
  - Human governance and oversight of AI outputs and AI vendors
  - Board and executive level oversight and involvement
  - Written AI program
- Enterprise AI Governance Frameworks
  - Numerous examples and models in the industry
    - GSA – AI Guide for Government
    - Singapore – Model Artificial Intelligence Governance Framework
    - Microsoft – Responsible AI
    - Specific Industry Guidance (NAIC – National Association of Insurance Commissioners – Exposure Draft)

Questions

Appendix of references

- Society for Computers & Law (SCL) AL Group Artificial Intelligence Contractual Clauses
- GSA – AI Guide for Government
- Singapore – Model Artificial Intelligence Governance Framework
- Microsoft – Responsible AI Framework

nortonrosefulbright.com