NORTON ROSE FULBRIGHT

Navigating the metaverse: A global legal and regulatory guide

Part 7: AI and the metaverse

Introduction

In the space of a very short time, businesses are focusing on what the metaverse means for them. In addition to commercialising the opportunities available to them, such as new channels to market and enhanced customer engagement, businesses will need to understand and address the associated risks.

Such matters are extremely important for businesses, consumers, law-makers and lawyers alike. In this seven-part guide we consider the following key legal and regulatory issues in relation to the metaverse:

Part 1

What is the metaverse?

Who are the current big players building it?

What will the metaverse mean for business?

What are key technical, operational and governance considerations?

Part 2

Intellectual property and the metaverse

What are virtual reality worlds and virtual items?

Non-fungible tokens

How do traditional IP concepts sit with non-fungible tokens and other works in the metaverse?

Part 3

Anti-trust/competition law issues

Developer and participant conduct

Will the EU Digital Markets Act apply to the metaverse?

Competitors communicating and co-operating with each other in relation to metaverse offerings

Overview of the legal and regulatory issues

The diagram shows the key legal issues and subject areas this guide covers. The breadth of issues means that mitigating risk associated with the metaverse is going to be a significant challenge for any business, but particularly so for a regulated business.

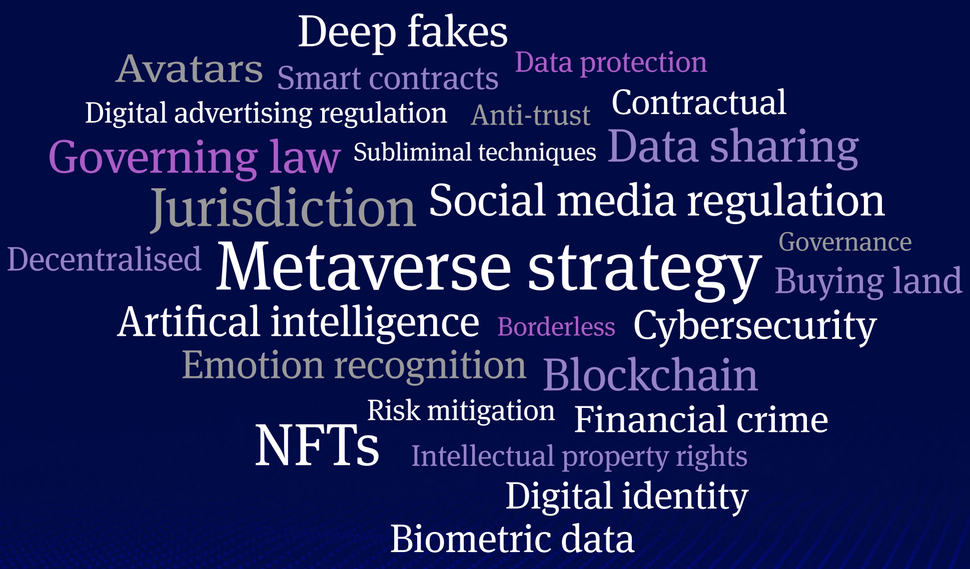

Deep fakes
Avatars
Smart contracts
Data protection
Digital advertising regulation
Anti-trust
Contractual
Governing law
Subliminal techniques
Data sharing
Jurisdiction
Social media regulation
Decentralised
Metaverse strategy
Governance
Buying land
Artifical intelligence
Borderless
Cybersecurity
Emotion recognition
Blockchain
Risk mitigation
Financial crime
NFTs
Intellectual property rights
Digital identity
Biometric data

AI and the metaverse

Why is AI relevant to the Metaverse?

The gaming world has long been an early adopter of AI. Virtual characters, interactions between characters and players, narrative developments, and even scenic backgrounds depend on AI to create and maintain realism in games. As the metaverse develops out of the gaming world (among other foundations), we see the same trend. To create human-like chatbots and participant avatars, engaging in realistic interactions in cinema-quality environments, requires many different AI techniques. For example, virtual characters presenting believable emotions in response to interactions with humans is an area where AI is applied.

Moreover, as economic and social interactions move into the metaverse, increasing quantities of data will be generated, leading to new applications of AI and predictive analytics.

How might AI regulation impact upon the metaverse?

The EU is at the forefront globally in proposals to regulate AI at the moment. As with the EU GDPR in relation to data protection regulation globally, it is expected that the proposed EU AI Regulation (now known as the EU "AI Act") will have an influence on how other jurisdictions regulate AI, so it is useful to understand the EU position.

The EU AI Act

The European Commission unveiled a proposal for the EU AI Act in April 2021. Due to the technical and political complexity of the legislative proposal, and the major impact it is expected to have on business and society, the negotiations through the European Parliament and the Council (made up of the 27 EU member states) have been lengthy, and are not near conclusion yet.

With a fair wind blowing, it is just about possible that the AI Act could be passed into law in the second half of 2023, with a 24 month transitional period (that is, requiring full compliance in 2025). However, it is all quite possible this could take longer.

Even though the text of the EU AI Act is far from agreed, as with the GDPR, many of the concepts in the AI Act will get picked up and applied by data protection and financial services regulators before the AI Act becomes law, so it is important businesses understand and respect the concepts as soon as possible.

As presently articulated in the Council's December 6, 2022, Common Position draft, the AI Act will apply to systems designed to "operate with elements of autonomy" which "infers how to achieve given objectives" and "generates outputs influencing the environment within which it interacts".

To be caught by the Act, such systems must generate content predictions, recommendations, or make decisions that influence the environment (that is, in a physical or digital dimension) with which they interact. As the AI Act specifies that impact on the "digital dimension" falls within scope, the metaverse is within scope.

Robo-advisor

How might the AI Act's current provisions affect possible AI use cases in the Metaverse? To answer that question in what follows, we will use the example of a bank's robo-advisor.

Let's assume that it analyses certain key data provided by an avatar/user, in circumstances where the robo-advisor then makes a decision or nudges the avatar/user towards an investment type, where the decision or nudge is based on a machine-learnt evaluation of the avatar's/user's:

- Financial sophistication.
- Stated or known assets.
- Indicators of risk appetite.

The European Commission undoubtedly recognises the potential and accompanying benefits of AI, and has declared its intent to support a metaverse-based economy in a recent letter of intent. Given its vision to protect and strengthen fundamental rights of people and businesses, while simultaneously encouraging AI innovation across the EU, the Commission proposes a risk-based approach to triaging categories of AI systems. It categorises AI systems into four risk types:

- Unacceptable.
- High.
- Limited.
- Minimal risk.

Unacceptable risks

What sort of AI-enabled metaverse use cases could land a business in the unacceptable risk category?

Unacceptable risks are those considered to be a clear threat to the safety, livelihoods, and rights of people, and are prohibited. AI systems that use "dark or deceptive patterns," and/or deploy manipulative "subliminal techniques" that, in either case, materially distort human behaviour in a manner that causes (or is likely to cause) physical or psychological harms are prohibited.

The tipping point for material distortion and psychological harm is undefined at present. Using our robo-advisor use case, it could conceivably cover instances where:

- An AI robo-advisor, or a series of programmatic personalised ads, nudges an individual towards investing in a newly launched crypto/NFT fund with the lure of high returns (notwithstanding the known volatility of digital asset markets compared to safer, more traditional investment avenues).
- Such nudge is the outcome of covert techniques without the associated risks having been clearly articulated.

Whether a consequential harm is psychological or not is key and remains to be seen.

The other category of prohibited unacceptable risk that could be tripped arises where AI technologies are used to exploit specific vulnerable groups (due to their age, imbalance of power, knowledge, economic or social circumstances), again in a way that materially distorts their behaviour and causes them physical or psychological harm.

For example, a bank may be dealing with an avatar operated by a child or a user with impaired mental capacity. Against this backdrop, it will be essential for such a bank to very quickly establish the true identity of the user of an avatar to ensure the interactions it is having are appropriate (that is, so as to ensure that its “real world” procedures can be applied).

Are businesses always aware of the involvement of AI in their operations?

Many organisations are not aware of all the AI in their systems, particularly for software procured from external vendors. Often, the internal champions for adoption of AI methods are people in the IT department, and other people in a business may not be aware of the use of these methods.

Auditing systems to identify those which use any of the methods defined as AI by the EU AI Act is important for good governance of these systems. The skills required for such audits are a mix of technical, legal and regulatory, and commercial. Compared with traditional IT systems, there are many new challenges in assessing risk – for example, assessing whether systems operate without bias or discrimination, and whether they can explain their decisions.

AI audits are ultimately about surfacing AI risk, particularly in the areas covered by the AI Act but not exclusively so.

High risk systems

High risk systems are not prohibited, but rather, are subject to a lot of operational controls and conformity assessments under the AI Act before they can be offered to the EU market.

According the draft AI Act, high risk systems are generally those:

- Covered by EU product safety certification rules.
- Falling within a specific list of AI areas of application, known as the "Annex III list".

If any business starts using AI in the metaverse to make decisions about whom to recruit, that would be a high risk use case under the AI Act.

For, say, a bank in the metaverse, a high risk use case under the AI Act will also be in relation to the use of AI to make consumer credit decisions. Credit decisions today are already dependent on more and more complex models and subject to financial services "treating customers fairly" regulation and data protection rules on fairness, transparency and automated decision-making.

High risk use case: an on-chain credit scoring model

The model could be based on blockchain data, and conducts comprehensive processing and evaluation of data in various dimensions, such as:

- Credit history.
- On-chain behaviour preference.
- Address activity level.
- Asset holdings and portfolio.
- Address correlation.

The AI Act adds a requirement for a robust quality management system for high risk AI to avoid ethical blindspots and technical malfunctioning. The control framework under the AI Act for the on-chain credit scoring model use case outlined above will include a strategy for regulatory compliance, including:

- Compliance with conformity assessment procedures.
- Procedures for the management of modifications to the high-risk AI system.
- Systems and procedures for data management.
- A post-market monitoring system.
- An accountability framework.

Commercial considerations

Before putting such an on-chain credit scoring model into a metaverse with EU users, a bank will have had to go through all the steps outlined above and self-certify that it has conformed to the requirements of the AI Act.

The main commercial issue for a bank in relation to high risk systems will be the time to market to go through this process (at least until all players are operating similar processes).

On the positive side, if done properly a bank should avoid a PR disaster – for example, where the system uses racial proxies to allocate credit etc.

How to operationalise AI risk mitigation in the metaverse

How can a business operationalise AI risk mitigation in relation to the AI Act?

High risk AI

We consider this issue by looking at high risk AI first, using our use case of using AI to assess consumer creditworthiness in the metaverse. In our case study:

- The algorithmic system awards users with ranked badges, which are unique NFTs, based on their credit score.
- The badges can serve as a kind of metaverse bank ID verification system, assuring other users and platforms that the bearer of a badge is who they claim to be, own what they claim to own, and is being truthful about their past on-chain behaviour.

High risk AI: steps to operationalising risk mitigation

A business can operationalise risk mitigation under the AI Act through:

- An assessment framework.
- An identification/triage process.
- The use of standards.
- Addressing vendor dependency.
- Providing for contingencies for malfunctions.

Assessment framework, identification and triage

- The AI Act requires businesses to have in place robust and effective governance, including a risk management framework, to identify, reduce and control any of the ethical, technical and legal risks associated with the use of high risk AI applications.
- Due to the fast-evolving nature of AI, and the increasing levels of adoption of AI solutions within many businesses, in both physical and virtual dimensions, the first job of the framework is to actually identify where AI is being used, and to triage it so that riskier applications (that is, ones that are unacceptable, high or limited risk) get more scrutiny than lower risk/better understood ones. AI audits are relevant here (see *Are Businesses Always Aware of the Involvement of AI in their Operations?*).
- The triaging and assessment process requires checklists of ethical, legal and technical points, requiring a lot of information and dependable ways of analysing it. For example, bias detection requires datasheets providing information about the training data, modelling and testing. There are also thresholds to be established (such as at what level of bias can actually be tolerated). The dialogue between technical experts (often who are heavily invested in seeing the AI application go live) and less technical compliance or ethical stakeholders is essential. The outcome of such dialogue is both the gating factor for AI to go live, and a defence should a discounted risk manifest itself later. This is time-consuming and should be factored into the commercial launch time line.
- Levels of explainability and human override functionality are determined through this process.

Standards

- The AI Act's requirements are conceptual, so tangible standards will come later via standards-setting bodies. In the meantime, data protection and financial services regulators are publishing more granular, actionable guidance which will start to set thresholds until more specific standards emerge.

Vendor dependency

- The length of time it takes to complete the evaluation, and the fact that autonomous AI applications keep evolving after going into production, means that: (1) the oversight and evaluation must be continuous; and (2) the ongoing evolution affects contracts with vendors providing such AI.
- Will a business want to pay for a system that fails the assessment after a 12 month evaluation? How much work and how much exposure of its IP will a vendor put at risk without commitment from the business?
- It is likely contracts with such vendors will be more like managed service agreements, with payments spread over the anticipated use time of the application in order to keep the vendor engaged.
- Accordingly vendor dependency needs to be factored in as both a risk factor and an element impacting timing in the AI roll-out.

Malfunction reporting

The AI Act introduces breach reporting and conformity assessments for high risk AI, so the operational framework needs to anticipate:

- How malfunctions will be communicated to regulators and users.
- Fall-back plans should the application need to be overridden or withdrawn.

Operationalising risk mitigation – high risk AI: a computer scientist's perspective

Computer scientist and AI specialist, Professor Peter McBurney, has this to say:

- Whether or not a system operates fairly and without bias may depend on the data used to train (or calibrate) the system and on the data input to it for each particular case it considers.
- Identifying systems having a bias, and rectifying any problems found, may therefore require access to all the datasets used as training or input data, along with all the output data generated.
- Best practice in this area is therefore to hold copies of all these datasets in repositories, for ease of possible future access, just as best practice in software development has long kept repository copies of all versions of the software created as it is developed.
- Unlike most other software systems, AI systems require assessment of aspects such as bias or fairness, both before and after deployment of the system into production. Undertaking such an assessment only before deployment is not sufficient, as there can be particular operational factors in production environments that can influence performance against these aspects.
- Waiting until a system is in production before undertaking an assessment of these aspects risks possible damage to users of the system, damage to the business's reputation, and the displeasure (or worse) of regulators.

Operationalising risk mitigation: limited risk AI

What type of AI is caught by the rules on limited risk AI under the EU AI Act and what requirements apply in relation to them? AI systems with limited risks are permitted, but have to fulfil specific transparency obligations:

- The most basic transparency obligation is that, when interacting with an AI system, users must be informed that they are acting with an AI and not a human. In the metaverse, where avatars could be completely automated, partially automated with certain interactions prompting a human to take over the control of the avatar, or fully human controlled, some sort of universal system for denoting when a human is in control would seem the way forwards.
- The other relevant area of transparency is that operators must inform users when they are using emotion recognition systems.

Use case: emotion-recognition avatar

A bank proposes to use an emotion-recognition AI avatar:

- Let us assume that the AI operates a so-called "Fear and Greed" index algorithm to evaluate users on various indicators, and to alert investors to their own emotions and biases.
- In this way the avatar AI could play a cautionary role in influencing user behaviour in relation to the stock market.

In such a use case:

- The user will have to be informed about the functioning of the algorithm.
- However, the use of such an emotion-recognition system to encourage users to plunge into risky investments will have the effect of converting the "limited risk" to the "unacceptable risk" category.

Similarly, metaverse users must be protected from, and made expressly aware of, the use of "deep fakes", which are images, audio or video content that appear to be deceptively genuine or real.

Operationalising risk mitigation: minimal risk AI

Applications of minimal risk AI under the EU AI Act include the vast majority of AI systems, such as video games or spam filters. The AI Act has a provision that encourages member states to facilitate the drawing up of codes of conduct to apply the same operational safeguards to minimal risk systems as those that apply to high risk systems.

Although applying such requirements to minimal risk AI might appear onerous (particularly given what is required to operationalise these safeguards), it could be sensible, and may eventually reflect good practice, particularly given a seemingly innocuous use of technology (such as a video game) could in fact have a harmful effect on users by psychologically manipulating or coercing them into taking excessive risk or exploiting others for their own gain.

Psychological manipulation or coercion

Examples of this kind of impact include:

- The use of NFTs in video games for injecting an artificial sense of scarcity into digital worlds for the benefit of an investor class, to the clear detriment of the gamer.
- Crypto or stocks trading AI simulators that perversely incentivise and encourage aggressive trading manoeuvres and deception tactics, if a user's portfolio worth is the sole determinant of their performance in that game.

Data protection and AI

We have already dealt with data protection generally in relation to the metaverse (see *Decentralised Models and Data Issues*). Here we focus on the narrower question of data protection issues that arise in relation to AI operating in the metaverse.

From a practical perspective, a business's presence in the metaverse will require it to ensure an individual it is interacting with is actually who they say they are. For example, a bank might:

- Seek to detect identity fraud through AI-led document verification that reaches outside the Metaverse.
- Add a virtual ID verification lounge where avatars can reveal themselves on camera to a bank AI agent that verifies the individual against the documents.

AI chatbots

A key concern from a data protection perspective is that AI chatbots could be programmed to never forget content that has been disclosed to them, which includes assessments that such bots have made on humans as part of their interactions. They might also draw inferences from such data. In the absence of controls, AI chatbots or avatars could breach data protection principles, especially those of lawfulness, fairness and transparency, purpose limitation, data minimisation, and accountability.

Use case: ID verification avatar

A bank wishes to use an ID verification avatar in the bank's virtual lounge that can review an individual's documents presented for loan approval.

The ID verification avatar is fitted with computer vision, and can "see" that the applicant is pregnant, for instance, and can "flag" such applicant as presenting a higher repayment risk (given the avatar's biased assumption that the applicant is likely to take time off work to care for their new-born).

Having an avatar in the metaverse that is able to auto-flag applicants, without meaningful human involvement embedded in such automated decision-making, may contravene Article 22 of the EU GDPR (in addition to other data protection principles).

NEON CITY
STAR

Key contacts

Australia

Nick Abrahams
Global Co-leader, Digital Transformation Practice
Tel +61 2 9330 8312
nick.abrahams@nortonrosefulbright.com

Ross Phillipson
Senior Advisor
Tel +61 8 6212 3449
ross.phillipson@nortonrosefulbright.com

Belgium

Jay Modrall
Senior Counsel
Tel +32 2 237 61 47
jay.modrall@nortonrosefulbright.com

Canada

Maya Medeiros
Partner
Tel +1 604 641 4846
maya.medeiros@nortonrosefulbright.com

France

Nadège Martin
Partner
Tel +33 1 56 59 53 74
nadege.martin@nortonrosefulbright.com

Clement Monnet
Counsel
Tel +33 1 56 59 53 91
clement.monnet@nortonrosefulbright.com

Sébastien Praicheux
Partner
Tel +33 1 56 59 54 25
sebastien.praicheux@nortonrosefulbright.com

Geoffroy Coulouvrat
Senior Associate
Tel +33 1 56 59 52 98
geoffroy.coulouvrat@nortonrosefulbright.com

Germany

Daniel Marschollek
Partner
Tel +49 69 505096 215
daniel.marschollek@nortonrosefulbright.com

Christoph Ritzer
Partner
Tel +49 69 505096 241
christoph.ritzer@nortonrosefulbright.com

Hong Kong

Justin Davidson
Partner
Tel +852 3405 2426
justin.davidson@nortonrosefulbright.com

Japan

Sam Inohara
Partner
Tel +813 4545 3213
sam.inohara@nortonrosefulbright.com

The Netherlands

Nikolai de Koning
Counsel
Tel +31 20 462 9407
nikolai.dekoning@nortonrosefulbright.com

United Arab Emirates

Adjou Ait Ben Idir
Partner
Tel +971 4 369 6393
adjou.aitbenidir@nortonrosefulbright.com

United States

Felicia J. Boyd
Head of IP Brands, United States
Tel +1 612 321 2206
felicia.boyd@nortonrosefulbright.com

Sean Christy
Partner
Tel +1 404 443 2146
sean.christy@nortonrosefulbright.com

Chuck Hollis
Partner
Tel +1 404 443 2147
chuck.hollis@nortonrosefulbright.com

Andrew Lom
Global Head of Private Wealth
Tel +1 212 318 3119
andrew.lom@nortonrosefulbright.com

Daniel Farris
Partner-in-Charge, Chicago
Tel +1 312 964 7730
daniel.farris@nortonrosefulbright.com

Susan Ross
Counsel
Tel +1 212 318 3280
susan.ross@nortonrosefulbright.com

Robert A. Schwinger
Partner
Tel +1 212 408 5364
robert.schwinger@nortonrosefulbright.com

Rachael Browndorf
Senior Associate
Tel +1 303 801 2763
rachael.browndorf@nortonrosefulbright.com

United Kingdom

James Russell
Partner
Tel +44 20 7444 3902
james.russell@nortonrosefulbright.com

Marcus Evans
EMEA Head of Information Governance, Privacy and Cybersecurity
Tel +44 20 7444 3959
marcus.evans@nortonrosefulbright.com

Lara White
Partner
Tel +44 20 7444 5158
lara.white@nortonrosefulbright.com

Sean Murphy
Global Head of FinTech
Tel +44 20 7444 5039
sean.murphy@nortonrosefulbright.com

Mike Knapper
Head of Intellectual Property, EMEA
Tel +44 20 7444 3998
mike.knapper@nortonrosefulbright.com

Harriet Jones-Fenleigh
Partner
Tel +44 20 7444 2867
harriet.jones-fenleigh@nortonrosefulbright.com

Michael Sinclair
Knowledge Director, Campaigns
Tel +44 20 7444 2344
michael.sinclair@nortonrosefulbright.com

NORTON ROSE FULBRIGHT

Norton Rose Fulbright is a global law firm. We provide the world's preeminent corporations and financial institutions with a full business law service. We have more than 3500 lawyers and other legal staff based in Europe, the United States, Canada, Latin America, Asia, Australia, Africa and the Middle East.

Law around the world

nortonrosefulbright.com

Norton Rose Fulbright Verein, a Swiss verein, helps coordinate the activities of Norton Rose Fulbright members but does not itself provide legal services to clients. Norton Rose Fulbright has offices in more than 50 cities worldwide, including London, Houston, New York, Toronto, Mexico City, Hong Kong, Sydney and Johannesburg. For more information, see nortonrosefulbright.com/legal-notices. The purpose of this communication is to provide information as to developments in the law. It does not contain a full analysis of the law nor does it constitute an opinion of any Norton Rose Fulbright entity on the points of law discussed. You must take specific legal advice on any particular matter which concerns you. If you require any advice or further information, please speak to your usual contact at Norton Rose Fulbright.

50962_EMEA – 09/23