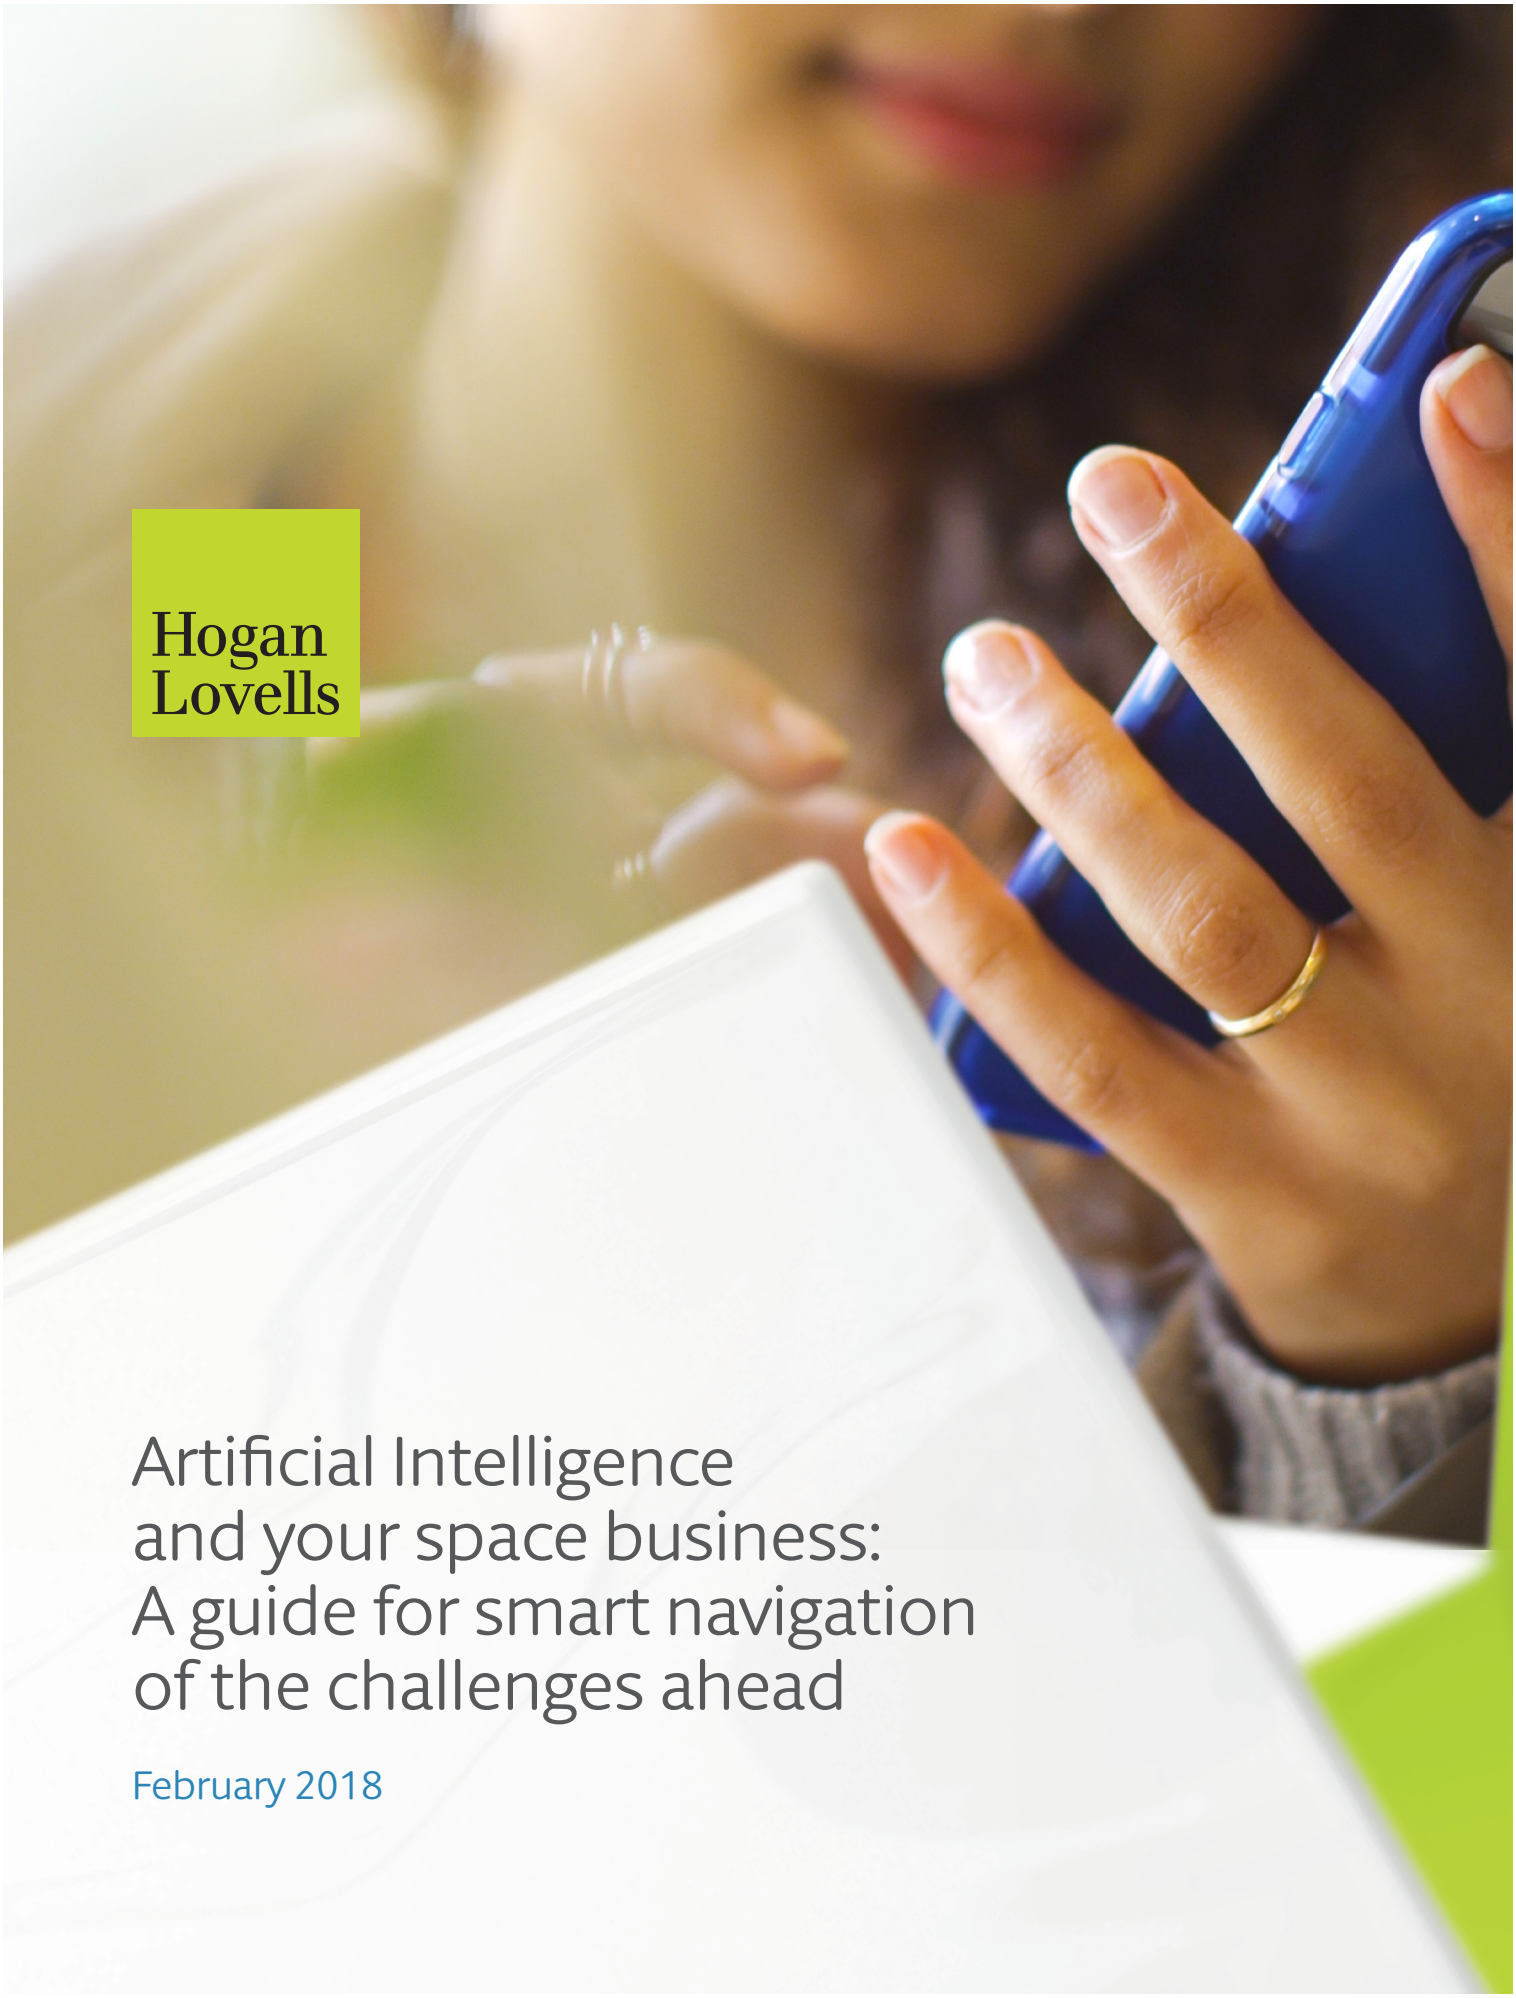

Hogan Lovells

Artificial Intelligence and your space business: A guide for smart navigation of the challenges ahead

February 2018

Global hotspots

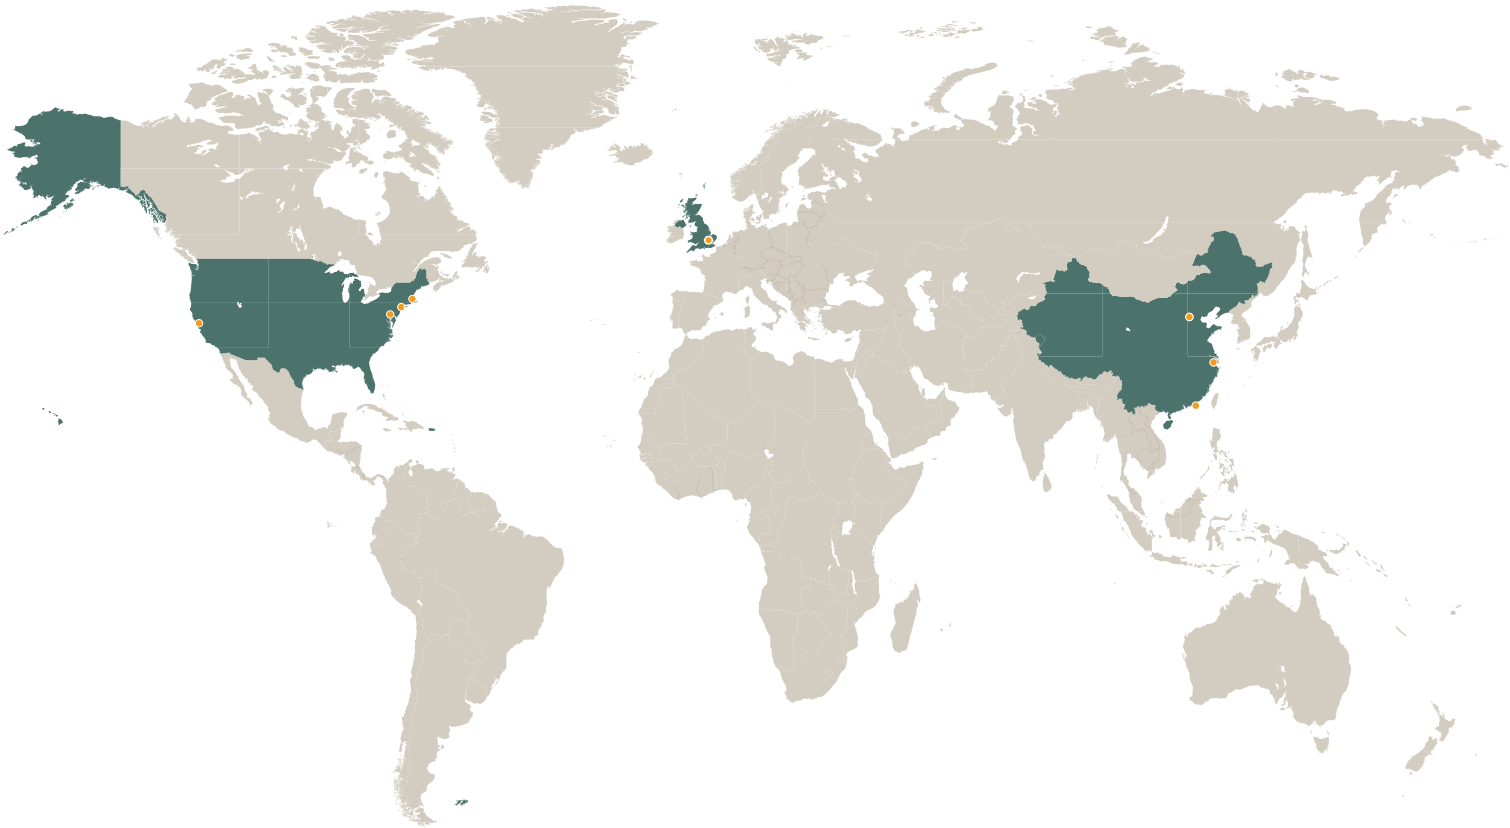

Silicon Valley

— Top global hub for startups with 12,000+ active startup businesses

— Global leader for venture capital investment

— Headquarters of many top high-tech companies

Washington, D.C.

— Leading center for U.S. policy and regulation of Artificial Intelligence (AI), including health, automotive, space, drones, and education

London

— Global finance center, supporting both investment and FinTech applications

— European leader of VC startup investments

New York

— Leading hub for financial and media industries

— Strong funding ecosystem, second in the world after Silicon Valley for absolute number of early stage investments

Boston

— Long history of cooperation between science and industry

— World-class universities such as MIT developing advanced technologies and providing a talent pipeline

China

— Leading in volume of academic research output in AI coming from universities

— AI identified as a strategically important technology by the Chinese government

Source: *Artificial Intelligence The Next Digital Frontier?*, McKinsey Global Institute, 2017

What is Artificial Intelligence?

Virtually every industry is being reshaped with the use of Artificial Intelligence (AI) and advanced machine-learning, ranging from space and drones, to healthtech to self-driving vehicles, to education and smart homes, social media, and everything in between and beyond. AI opens up new ways to accomplish existing missions, as well as generating huge databases of information to which new algorithms and data analytics can apply. These new technologies present a variety of commercial opportunities and the potential to change our daily lives and businesses in significant ways.

At the same time, new AI innovations bring legal, policy, commercial, and strategic challenges that need to be considered thoughtfully across jurisdictions and applications. In some instances, existing frameworks can be applied or adapted, but for others new paradigms and robust safeguards may be needed. And as machine-learning technologies continue to evolve, organizations will need dynamic and sophisticated compliance approaches.

In this guide, we highlight the key challenges and commercial opportunities for AI and advanced machine-learning, with particular focus on space-based business considerations. We also touch on AI in the areas of space, drones, and terrestrial convergence, particularly communications and imaging platforms or applications.

Space and Satellite

The nature of the space and satellite industry presents a quintessential use case for AI. Everything about the industry requires machine intelligence and assistance to launch, operate, maintain, control, repair, and ensure achievement of the business mission. Mission success heavily relies upon sophisticated computer-assisted models, algorithms, robotics, and communications across long distances (from geostationary earth orbit (GEO) to low earth orbit (LEO), and everything in-between (medium earth orbit (MEO)). Some examples of potential AI applications include:

- **Remote sensing and monitoring** of a broad array of potential targets, including environmental changes, dark ships and national security, fleet management, and aircraft and maritime tracking.
- **Communications** between ground and space, and from satellite-to-satellite (in the case of a multi-satellite constellation), using radio frequencies, optical-laser communications, radar and other technologies, along with growing complexity of satellite-to-satellite handoffs between satellites in different orbits.
- **Robotics** in space, including mission extension vehicles, space docking, satellite health monitoring, manned space vehicles support (including health, safety, medical, analytics, and repair), and spacecraft, such as automated transport vehicles, designed to make their own decisions to explore, learn, identify, and adapt during missions; and carry out repairs.
- **Data analytics** including the policy and regulatory issues inherent in gathering large amounts of information, and how that information can be used, from national security (and sovereignty), data privacy, and proprietary perspectives.
- **Reusable launch and manned vehicles** including sophisticated AI for return to Earth for completion of mission.
- **Asteroid mining** including analytics of substances discerned from asteroid samples, and remote mining of the same.
- **Remote missions** to Mars and beyond (and a broad variety of information transit, maneuvers, and return.)
- **Satellites as alternative to terrestrial-based systems** including cloud computing, cross-border broadband services, and other multi-jurisdictional data and information transfer.

The breadth of these space-based services requires consideration of a broad range of legal, policy, and commercial issues, including:

Regulatory jurisdiction

For traditional GEOs, single jurisdiction (plus applicable International Telecommunications Union) rules govern, with a more limited scope of cross-border questions raised based on landing rights. As systems expand to LEO and MEO orbits and operate in the area of more innovative technologies (such as remote sensing, high resolution data gathering, dark ship monitoring), the exercise of jurisdiction on a global basis becomes more complex. Additional complexities occur with new satellites with innovative missions, such as mission extension vehicles and satellite health monitoring, where the satellite missions require continued relocation amidst a field of other satellites.

Cross-border data collection and dissemination rules

A space-based business operating in multiple jurisdictions has to address the multiple jurisdictional rules on information gathering and collection. The space-based AI is subject to all privacy and data protection rules and any government national security restrictions (including those that may apply to their own airspace).

Product liability, cybersecurity, insurance, and litigation

Satellites and launches can give rise to large liabilities in the event of a satellite failure, collision, destruction (self or involving other satellites), or cybersecurity incidents. AI can be used to both protect the safety and security of operations, but can also be used as a tool for interference, hacking, and/or destruction.

Randy Segal
Partner, Northern Virginia, Silicon Valley, and Washington, D.C.
T +1 703 610 6237
randy.segal@hoganlovells.com

Steven Kaufman
Partner, Washington, D.C.
T +1 202 637 5736
steven.kaufman@hoganlovells.com

Stephen Propst
Partner, Washington, D.C.
T +1 202 637 5894
stephen.propst@hoganlovells.com

Jeffrey Epstein
Counsel, Northern Virginia
T +1 703 610 6144
jeffrey.epstein@hoganlovells.com

Tony Lin
Counsel, Washington, D.C.
T +1 202 637 5795
tony.lin@hoganlovells.com

Unmanned Aircraft Systems

Unmanned aircraft systems (UAS or drones) technology has moved forward rapidly in recent years, and what used to be considered toys are quickly becoming powerful commercial tools that can provide enormous benefits in terms of safety and efficiency. Estimates by consulting firms suggest that the global market for commercial UAS technology applications alone, which currently stands at about US$2bn, could increase to US$120bn by 2020.

Advances in AI and machine-learning technology are allowing UAS to see and act like human pilots, and to process huge amounts of data in real-time. Whether UAS are performing search and rescue missions, allowing farmers to be more efficient and environmentally friendly, inspecting power lines and cell towers, surveying and mapping large swaths of land, or performing package deliveries, AI is allowing drones to become more automated, safer, and efficient.

The applications for AI in the drone industry are limited only by our collective imagination. The use-cases range from data analytics for industrial infrastructure inspections, to navigating warehouses more efficiently, and everything in between.

Real-time data analytics

AI is allowing drones to collect and process huge volumes of data in real-time. Aerial imagery that used to take humans hours, days, or weeks to review and analyze is being streamlined and automated by AI that strategically determines what kind of data and images are important enough to collect, and can simulate a human looking at thousands or millions of images. For example, drones being used to perform railway inspections can use a variety of onboard sensors (cameras) to inspect track conditions and identify defects that are invisible to the naked eye. Once detected, AI software can be used to provide recommendations on what, if any, maintenance may be necessary.

Sense-and-avoid

A pilot manually flying a drone should be able to avoid obstacles like buildings or other aircraft. But what happens if a drone loses all connectivity? To fully enable many of tomorrow's most promising use-cases, drones will need to be capable of flying autonomously without human intervention, and this will require drones to be able to sense obstacles and react in time to avoid a collision. Computer vision plus machine learning is helping drones navigate more effectively by allowing drones to see the world the way humans do. AI software is enabling drones to fly autonomously, even in dark, obstacle-filled environments or beyond the reaches of GPS or other methods of connectivity.

Swarm technology

AI technology is enabling swarms of tens or hundreds of drones to operate entirely autonomously. The swarm collaborates by staying in constant communication with itself and by changing its configuration to complete the mission if any one drone is lost.

Situational awareness

AI is enabling better situational awareness and changing the way drones are able to interact with things in their environment. In the not-too-distant future, AI technology will enable fully autonomous drone operations. Civil airworthiness authorities around the world maintain air safety by placing ultimate responsibility for safe operation of aircraft on the entity operating the aircraft and on the human pilot. Since fully autonomous drones will not have a human pilot, countries around the world will need to put in place policies, laws, and regulations that fully address this profound change.

Lack of human judgment

Fully autonomous drones will raise important policy questions regarding the removal of human judgment from the equation. Human pilots make not only safety-related decisions, but in certain circumstances — especially emergencies — moral and ethical decisions, such as whether to crash land in a heavily populated area versus a less populated one. With the removal of the human pilot and human judgment, what level of AI will be needed for drones to learn from experience and use that learned knowledge to make appropriate moral and ethical judgments?

Security

Who will have the legal responsibility to maintain the security of a fully autonomous drone, and to ensure that its automation, navigation, and communications systems are not hacked into?

Regulatory and civil liability

What if something goes wrong with a fully autonomous drone? Who will be responsible from a regulatory compliance and civil liability perspective in the event of an incident involving personal injury or property damage, or a failure to comply with rules and regulations?

Lisa Ellman
Partner, Washington, D.C.
T +1 202 637 6934
lisa.ellman@hoganlovells.com

Matthew Clark
Senior Associate, Northern Virginia
T +1 703 610 6154
matt.clark@hoganlovells.com

Gretchen West
Senior Advisor, Silicon Valley
T +1 650 463 4062
gretchen.west@hoganlovells.com

Issues to consider with space-related AI

Drafting contracts with AI

As shown by our above examples, AI has or will transform virtually all industries including in ways not yet known. With this transformation will come uncertainties as to how existing and new legal frameworks will apply to the new technologies and the liabilities that may follow.

Innovations raise many regulatory questions, not just at the compliance level, but at the fundamental nature of the innovation itself. The issues range from whether and how the new technology is to be regulated to whether the regulatory scheme to be applied will support innovation or, conversely, create hurdles that will stand in the way of (or even block) its development. These innovations will also raise cross-border jurisdictional issues, application of multiple regulations, and how to navigate the rules in product development, deployment, and contracting matters.

So what do you do?

As in the case of all innovations and disruptive technologies, you should start with the basic premise that using contract boilerplate for the main terms and conditions will not achieve a good result. The contract for a new business model involving disruptive technologies must be built from the ground up, with a clean sheet of paper architecting the end-to-end system and service expectations, including technology development, technical capabilities, customer experience, financial model, risks, budgeting and handling cost increases, regulatory hurdles and changes, and termination strategy, to name a few. And you must build in provisions for systemic change, in other words have a contract that itself can evolve.

Care must be taken to consider:

— how will this new system operate,

— what flexibility is needed (or can be provided), and

— how are unknowns and possible (or probable) risks taken into account?

Next, you must consider what goes on the clean sheet of paper, as it forms the essence of the business arrangement between or among the parties. We have divided this analysis into three parts, which reflect three different goals in the contracting process.

First, develop a contract that contains the necessary terms and reflects the company's strategy

This includes taking an inventory of the knowns and unknowns of the technology to develop a contractual roadmap that will contain sufficient flexibility to change course based on technology, regulatory, and other developments. It also requires you to design an acquisition strategy, including:

— how to acquire the relevant rights for what exists today,

— how to acquire rights to the next stage of the technology (at least to the extent it is developed by the counterparty),

— how to price these acquisitions (including receiving credit for obsolete technology that has to be replaced),

— appropriate acceptance criteria,

— how much control and exclusivity is desired (considering exclusivities, rights of first refusal (ROFRs), and most favored nations (MFN) provisions, and
— appropriate decision mechanisms with off-ramps to protect the parties against situations too far from the envisioned business model.

Maintaining flexibility for change is valuable in a changing technology and regulatory landscape. There is no one-size-fits-all solution and only a careful consideration of your business situation, aligned with the legal and commercial toolkit of terms, will enable you to determine the likely optimal terms for your new technologies. In all cases, the ability to foresee the future will be imperfect, but careful planning and strategic thinking will help improve the clarity and certainty of reaching the best solution.

Second, anticipate third party events that need to be factored in, dealing with the changing regulatory landscape for the new technology and providing for its effects on the parties' deal

AI brings with it significant new issues involving product liability, data privacy, intellectual property, and almost every area of the law. When this is layered atop the global reach of products using AI, the cross-border issues in a developing legal and regulatory landscape are tremendous. In some cases, regulatory conditions should be considered to bound liability issues within your tolerance for risk and/or business model, since the changing regulation may well make achievement of certain contract goals impossible and these situations need to be handled by specifying some outcome. Allocation of responsibilities and costs for compliance with future laws should be considered since these costs could be substantial.

Third, overlay the standard allocation of known or anticipated risks between parties, with separate provisions for allocating unknown risks through contract adjustments or exit strategies

Thinking these issues through is critical, and it may be to your advantage to set cost and liabilities expectations, rather than leaving the implications of changes to later dispute resolution. All reasonable scenarios should be contemplated when drafting agreements to ensure that all compliance, approval, cost, indemnification, termination, insurance, and financing provisions support the desired business outcome.

Once there is agreement on the allocation of liabilities and risks, the parties need to support that agreement with appropriate indemnification provisions. These clauses, often considered boilerplate in more routine arrangements, take on greater importance because there are so many uncertainties with respect to which indemnification provisions may be called upon to address risk allocations. Insurance can play an important role to backstop indemnification provisions and the attendant risks, including the risk that the indemnifying party may not, as a practical matter, have the ability to step up to its contractual commitments.

In addition to this, there may need to be an overlay for unknown risks. When the parties' goals are frustrated, do you bring in an industry expert to reform the contract to best implement the parties' expectations? Or do you build in renegotiation points along the way, noting always the risks of an unenforceable agreement to agree? This can be done in bold ways, with agreements to supply based on technologies not yet developed, with prices to be set based upon future market conditions. There must also be a series of off-ramps, where the exposure gets too high (as specified in contract clauses) and the parties have the right to stop. This relies upon mutual assured injury to encourage a further negotiation at that time based upon new data.

There is no one best answer as to what goes on the clean sheet of paper

Indeed, it will vary based on the particular business plan, the nature of the contracting parties, the specific business plan risks, relative leverage, and many other factors. But one common theme is critical to all cases: taking the time to carefully consider a full range of outcomes and possibilities while structuring your contract. Even terms of early stage contracts can have long-lasting impacts on business flexibility, market positioning, and customer commitments. Therefore, getting it right from the start is imperative.

Export Controls

AI is a cutting-edge area that raises new and complex export control issues. Given that AI is a nascent technology that is rapidly evolving, export control rules do not yet impose express, specific restrictions on it. However, AI related software and technology may be caught under existing rules that were never intended to capture it, resulting in a potential mismatch between the regulatory regime and technology such as machine and deep learning. Accordingly, navigating the U.S. and non-U.S. export controls applicable to AI requires sound judgment and extensive experience with export control requirements.

Military applications

The International Traffic in Arms Regulations (ITAR) administered by the U.S. Department of State impose stringent restrictions on the export, re-export, temporary import, and brokering of defense articles, technical data, and defense services. As governments and defense companies apply AI to defense projects, including weapons platforms, such technology, software, and AI-enabled hardware may be subject to the strict controls of the ITAR, even where the underlying machine and deep learning technology is based on commercial techniques.

High performance computing

The rapid evolution and adoption of AI techniques is expected to drive the market for high performance computing in the coming years, with AI platforms consuming more and more computing power. Certain high performance computers, and related software and technology are subject to strict controls under the Export Administration Regulations (EAR) administered by the U.S. Department of Commerce. The export, re-export, and transfer of such hardware, software, and technology may be subject to licensing and other requirements under U.S. and non-U.S. law.

Space and satellite

Military and commercial space-based systems also are subject to significant export controls under the ITAR and EAR. As the space industry adopts machine and deep learning techniques to assist with launch, operation, maintenance, and other activities, related AI technology, software, and AI-enabled hardware also may be subject to significant control under export control regulations.

Drones

The drone industry is expected to adopt AI to enhance the operation of drones and other mission critical functions. Military drones are controlled under the ITAR, and certain commercial drones are subject to stringent controls under the EAR depending on their range and duration of flight. To the extent AI is incorporated into drones, such technology, software, and AI-enabled hardware may subject to the highly restrictive controls applicable to drones.

Ajay Kuntamukkala
Partner, Washington, D.C.
T +1 202 637 5552
ajay.kuntamukkala@hoganlovells.com

Stephen Propst
Partner, Washington, D.C.
T +1 202 637 5894
stephen.propst@hoganlovells.com

Media Regulation

Freedom of expression is one of the pillars of democratic society because without it, no other right could exist. AI can have adverse effects on freedom of expression because it can anticipate the kind of information that you like and simply feed you more of the same. This is called the filter bubble effect, which can lead to increasing polarization of society and the absence of democratic debate.

This bubble effect is a hard problem to solve, but the problem is broader than just a debate about AI. This issue instead relates to how the state should intervene to help make sure the marketplace of ideas functions properly. Generally, the state is the last person stakeholders would want to intervene in a marketplace of ideas because the state is, for many people, the most dangerous monopolist. In the age of analog television and radio, media regulators helped ensure that citizens received a diverse set of viewpoints on topics of interest to the public. In the digital age, providing viewpoint diversity is much more difficult given the diversity of content available. How do you encourage citizens to explore all areas of a vast public library? Many countries are looking at how public service broadcasters can fulfill their public service role in an online environment. The regulatory debate should center on the future of media regulation, not on the regulation of AI.

Communications Network Regulation

The data analytics and processing performed in most AI applications will require access to sophisticated computer hardware and software and high-capacity, low-latency communications network connections. For many AI applications, the most immediate communications link to the user device will be a wireless link, either terrestrial or satellite, because the user device collecting, processing, storing, and sending the data will be in motion. Because most AI users will not be in a position to build their own private communications networks, they will have to rely mainly on mobile connectivity provided by third-parties, including commercial terrestrial wireless and satellite operators. The communications networks of these providers will need to be high-capacity, ubiquitous, secure, and reliable.

Depending on the requirements or sensitivity of the particular AI application, AI users will need to establish redundancy measures to ensure that their AI applications will be maintained at a high quality and reliability level when a primary communications link is temporarily unavailable. The user device hardware that will be collecting, processing, storing, and transmitting the AI data (including on-board sensors and on-board radio communication devices) may themselves be subject to government radio frequency exposure, emissions, and other limits.

Care will need to be taken in procurement of the communications network capacity and equipment necessary to support these business objectives through service contracts with third-party providers and engagement with government regulators. Relatedly, issues will need to be borne in mind with respect to government regulation of communications law, spectrum policy, licensing, equipment, network construction, and service quality and reliability issues.

Michele Farquhar
Partner, Washington, D.C.
T +1 202 637 5663
michele.farquhar@hoganlovells.com

Ari Fitzgerald
Partner, Washington, D.C.
T +1 202 637 5423
ari.fitzgerald@hoganlovells.com

Mark Brennan
Partner, Washington, D.C.
T +1 202 637 6409
mark.brennan@hoganlovells.com

Trey Hanbury
Partner, Washington, D.C.
T +1 202 637 5534
trey.hanbury@hoganlovells.com

Alexander Maltas
Partner, Washington, D.C.
T +1 202 637 5651
alexander.maltas@hoganlovells.com

Tony Lin
Counsel, Washington, D.C.
T +1 202 637 5795
tony.lin@hoganlovells.com


“You may not realize it but AI is all around us.”
Judy Woodruff

Privacy and Cybersecurity 

The large volumes of data collected by AI systems, and the extensive and complex processing such data undergoes, may create challenges for compliance with laws focusing on individual privacy, and how such data is secured.

Many privacy regimes around the world are based on internationally recognized privacy principles known as the Fair Information Practice Principles (FIPPs), and several of the FIPPs may be challenging to implement in the context of AI. For example, the principle of data minimization, which calls for collecting only the minimum amount of data necessary to accomplish a specified purpose, is in tension with the need for AI systems to gather large amounts of data, not all of which may be able to be identified as relevant at the outset of collection.

In the U.S., there is no singular, comprehensive data privacy law, but rather a patchwork of sector-specific privacy protections. Although these laws were not drafted with AI systems in mind, companies will need to be mindful of the restrictions such laws may place on specific AI projects, which may need to track certain individual level activity or functions over time. For example, the healthcare industry is a ripe target for AI innovation, as AI products may help improve the speed and accuracy of diagnoses and refine and tailor treatment plans. Achieving these outcomes may require tracking individuals' treatment and response over time. However, obtaining the medical data necessary for training AI may be a challenge, as the federal Health Insurance Portability and Accountability Act (HIPAA) places restrictions on how health plans, healthcare clearinghouses, and healthcare providers can use and disclose protected health information. State medical privacy laws may similarly restrict access to health information. Thus, companies seeking to innovate in the healthcare space will need to thoughtfully consider how to lawfully obtain comprehensive data sets that can enhance learning and treatment.

Another example of a U.S. privacy law that may impact AI systems is the Fair Credit Reporting Act (FCRA), which governs the use of credit reports – essentially any information collected or compiled that will be shared with others for use in credit, insurance, or employment eligibility determinations. The FCRA provides consumers with broad rights of access and correction, and it imposes various requirements on consumer reporting agencies. Companies working on AI systems may inadvertently become swept into the purview of the FCRA depending on how recipients of the information developed make use of the information.

In the European Union, a new regulation coming into effect on May 25, 2018 will have far-reaching impacts on AI products. The General Data Protection Regulation (GDPR) defines personal data broadly, such that much of the data processed by AI systems arguably would likely be covered. The GDPR requires data controllers to provide individuals with privacy notices. For example, where the data processing involves "automated decision-making, including profiling," the privacy notice must include "meaningful information about the logic involved." The difficulties posed by AI in readily translating how algorithms function specifically may make such an explanation difficult to provide. Further, the GDPR requires appropriate precautions to avoid discriminatory effects from profiling. It may be challenging for companies to fully account for all unintended biases depending on how AI outputs are used, especially as the uses may not be controlled by the entity that developed a particular AI solution.

The GDPR also requires data controllers to provide individuals with rights of access, rectification, erasure, restriction of processing, data portability, and objection to certain types of processing. Companies will have to design AI products with these rights in mind and provide mechanisms for individuals to exercise such rights where AI outputs may include personal data. Similar issues may arise under other privacy law regimes globally. This likely will require creativity and careful construction throughout the design process.

From a cybersecurity perspective, the threats to AI data from attackers or negligent handling are many and varied. It is important to reasonably secure any personal data that AI outputs may analyze, especially where the information reveals sensitive characteristics, such as medical conditions or financial history. In addition to protecting the underlying information analyzed, companies may need to protect their algorithms and AI outputs, which in many cases will be confidential and proprietary as to the AI company itself or its customers. Companies will also need to develop and implement comprehensive cybersecurity programs to help protect information and implement, test, and adjust their programs and incident response plans as threats continually evolve.

Discrimination

AI can help make decisions based on historical data. However, the outcome of the data analysis may yield results that are socially unacceptable. The algorithm may predict that someone is a bad credit risk because they grew up in a certain part of Oregon, or because their parents were born in another country. In most instances, the algorithm itself is not the origin of the bias. The problem relates to the data that are analyzed. AI analyzes historical data from real life. Data from real life is messy, and reflects the biases and bigotry of human society — in other words, garbage in, garbage out.

The designers of AI systems are working on solutions to the problem. Ideally, we would analyze data from the world as we would like it to be, not the world as it actually exists. The answer may be to ensure that decisions that result from AI are checked by humans before they create effects for an individual. This kind of human review is precisely what Article 22 of the GDPR (EU Regulation 2016/679) attempts to do. The GDPR gives individuals an absolute right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects. In addition, existing laws prohibit all forms of discrimination based on sex, skin color, or religion whether in a work place or elsewhere. The existing legal mechanisms are not perfect, but they do exist.

AI applications will merit testing and risk assessments before they are deployed, to anticipate potential problems such as illegal discrimination. Article 35 of the GDPR requires data protection impact assessments to be conducted for any risky processing. These impact assessments should be expanded to cover other risks associated with AI, such as risks of discrimination.

Timothy Tobin
Partner, Washington, D.C.
T +1 202 637 6833
timothy.tobin@hoganlovells.com

Winston Maxwell
Partner, Paris
T +33 1 53 67 48 47
winston.maxwell@hoganlovells.com

Sam Choi
Associate, London
T +44 20 7296 2000
sam.choi@hoganlovells.com

Product Liability

Fast-paced development and innovation can raise interesting product liability challenges for manufacturers and those in the product supply chain, including ensuring that new products meet the requirements of relevant regulatory regimes, while also seeking to minimize future litigation risks. The latter can be especially difficult as regulatory and legislative regimes, and even the common law, often have not kept pace with product innovation. These considerations are especially relevant in light of the rapid advancements in AI in recent years.

AI's place in the product compliance and liability landscape

A number of jurisdictions including the U.S., EU, South Korea, and Japan have started to consider whether AI products need specific legislation, regulations, and standards. By way of example, from an EU perspective, there is currently no set of laws or regulations that apply to AI in particular. Instead, a manufacturer would need to look at the wider EU legislative landscape applicable to products. As for any product, that landscape will depend, among other things, on the product's features and functionality. The existing product laws and standards would need to be considered in much the same way as when any new product is being designed for market launch.

Similarly, existing U.S. legal requirements are likely to regulate AI, at least initially. Identifying pertinent legal standards, however, will not always be straightforward. For instance, courts will have to answer if, and under what circumstances, AI that is incorporated into a tangible object, such as an autonomous vehicle, qualifies as a product subject to strict liability.

When looking to launch a new AI product in a market there are likely to be additional complicating factors such as:

— the identification of appropriate safety and other product standards,

— determining the application of relevant product laws in circumstances where the laws could not possibly have envisioned the technology in question (and where relevant guidance or case law may be thin on the ground, or especially challenging to apply); and

— the appropriate testing of the product (this could include, for example, identifying a test house with the requisite expertise, experience, authority, and equipment).

The challenge of AI to existing product liability regimes

It has been argued that the most challenging legal issues arise when human intervention is taken out of the equation and AI begins to make its own independent decisions. For example, most defects traditionally exist at the time when the product was sold. But AI will increasingly be capable of learning on its own. If an AI product learns to become unsafe in response to its external environment, would the capacity to learn to become unsafe make it a defective product, bringing it within the scope of product liability regimes? What types of injuries would be the foreseeable consequences of AI continuing to learn? Who would be liable and under what theories (e.g, the product programmer/designer, the manufacturer who puts the "nuts and bolts" of the product together, or less traditional strict liability defendants like the owner of the AI's algorithm?) What about the consumer who home-programmed the product? These are the type of issues which manufacturers will need to grapple with assessing litigation risks associated with marketing new AI products.

There are also interesting practical and evidentiary issues to be considered. For example, a judge or jury may prefer the testimony or video recording of an AI product to a competing first-hand (human) witness of fact? Could an AI product perjure itself and if so, would the manufacturer be held liable for this offense? Perjured testimony by AI could be particularly damaging given the public's historical overreliance on the accuracy and reliability of new technologies.

To start addressing these issues, some commentators have argued that it would be sensible to assign legal personality or personhood to sophisticated AI products rather than placing the entire burden on the manufacturer (but it's important to note that this does not equate to giving machines legal rights). This would mean that a product/robot could be held liable for any damage it causes and could be sued in its own right. Of course, this approach would require that the product be covered by insurance. However, this approach is not without its own pitfalls. It remains to be seen whether the insurance market will offer affordable policies covering new AI products. In addition, consumer groups may question whether, if insurance largely replaces traditional product liability, manufacturers retain a sufficient incentive to ensure a high-level of safety for their products.

Conclusion

The fundamental question is how to ensure the safety and performance of AI products while not stifling their development and introduction to the market. Existing product compliance and liability regimes will be tasked with answering this question while AI-specific rules continue to develop. AI's fit within these legal regimes may at times be awkward, but is by no means impossible if past technological advances are any guide. Meanwhile, the AI-specific rules that emerge are an opportunity for creative, practical solutions and should be tailored to avoid a legal environment which becomes characterized by inefficiencies, stifled innovation, wasted opportunities, and the need for constant amendments as these emerging technologies present new challenges.

Christine Gateau
Partner, Paris
T +33 1 53 67 18 92
christine.gateau@hoganlovells.com

Valerie Kenyon
Partner, London
T +44 20 7296 5521
valerie.kenyon@hoganlovells.com

Michael Kidney
Partner, Washington, D.C.
T +1 202 637 5883
michael.kidney@hoganlovells.com

Intellectual Property

Ownership: Patents and copyrights

Patents and copyrights are forms of intellectual property (IP) in which the government grants protections to the creators of novel works – patents protect new and useful inventions, and copyrights protect original works of authorship fixed in any tangible medium of expression.

The twin questions of "Who is the inventor?" and "Who is the author?" bring up interesting and complex questions in the field of AI. For example, when an AI system creates visual images or audio compositions, are they copyrightable? To some extent, this is an extension of the monkey selfie case several years ago, in which it was argued that when a photographer set up his camera in the forest and a Celebes crested macaque managed to press the shutter-release button while looking into the lens, the monkey should be considered the author of the resulting photo. Similar questions arise when AI algorithms are able to develop new and useful objects (or even other algorithms). Is the AI the inventor or author? Can inventorship or authorship be attributed to a nonhuman?

Moreover, in the case of patentable inventions, if the solution to a technical problem is developed by the AI system, yet is obscured by the black box of the AI algorithm, how can the proprietors of the AI system even recognize or determine that the AI has devised a solution that is sufficiently novel to be potentially patentable? It may, for example, be entirely obscured how the solution is carried out.

Relatedly, can the human developers of the AI system be deemed to be the inventors or authors of the AI system's output? Would the answer be different when an AI system develops inventions or art or music that was not specifically foreseen by the human developers of the AI system?

Ownership: Trade secrets

Trade secret law, another traditional field of IP, raises a different, but equally challenging set of issues. To be protected as a trade secret, information must have independent economic value from not being generally known to the public, and must be subject to reasonable efforts to maintain its secrecy. In trade secrets litigation, it is common to require that the claimant specifically identify its trade secrets, and also explain the efforts to maintain secrecy.

Trade secret misappropriation generally involves taking, disclosing, or using trade secrets under circumstances where the taking, disclosure, or use is improper (such as stealing them or violating confidentiality agreements). Even where one party has trade secrets, it is generally not misappropriation to independently develop the same technology, or to reverse engineer publicly available aspects of the technology.

If an AI system comes up with a technical solution that happens to infringe on third parties' patent rights, or develops art or music that has too-uncanny similarities to known, existing works, who is the infringer?

Where information is the product of AI, it is possible to theorize that it could have independent economic value from being nonpublic, and that it would be subject to reasonable efforts to maintain its secrecy. The problems of inventorship or authorship may not arise in the same way they do with patents and copyrights. But, where the information is the product of an AI system – particularly where it is within the black box of how the AI system performs its analysis – there may be difficulties articulating specifically what the trade secrets are, and possibly also how their secrecy has been maintained.

Ownership: Data

A further area of proprietary rights also bears mentioning: ownership of data. Data is increasingly recognized as valuable in its own right. Yet it doesn't always fit easily within the traditional IP doctrines. With the increased processing complexity and speed of AI systems, data, particularly large data sets, are an ever-more important consideration.

Infringement

On the flip side, if an AI system comes up with a technical solution that happens to infringe on third parties' patent rights, or develops art or music that has too-uncanny similarities to known, existing works, who is the infringer? Can the AI system infringe a patent or a copyright?

Celine Crowson
Partner, Washington, D.C.
T +1 202 637 5703
celine.crowson@hoganlovells.com

Dr. Christian Mammen
Partner, San Francisco
T +1 415 374 2325
chris.mammen@hoganlovells.com

Antitrust

AI has consequences that go far beyond the direct purpose of the technical devices themselves. The same technology can have totally different outcomes when introduced into different contexts. Algorithms may facilitate perfect competition or they may facilitate collusion. For instance, some algorithms make markets more transparent and dynamic and thus have pro-competitive effects. On the other hand, AI using algorithms that implement collusive structures by monitoring and punishing deviation by any competitor without the need for explicit communication raise antitrust concerns. Detecting the difference between the pro-competitive and anticompetitive algorithms is, however, not an easy task.

Moreover, the DNA of AI is to take on a life of its own. This raises a difficult question regarding liability for antitrust liability. If there is no or only a weak link between the principal (the human) and the agent (the algorithm), who is on the hook for antitrust infringements? Some antitrust authorities already sent a clear warning message. With the words of the EU Competition Commissioner: “Companies can’t escape responsibility for collusion by hiding behind a computer program.”

One benchmark to hold someone liable under antitrust law for wrongdoing of AI could be whether the human could have been anticipated what the computer did. If it can be anticipated that an algorithm can lead to an anticompetitive action, such infringement by the algorithm will be attributed to the company. This is why businesses using AI (created in-house or by third parties) should be well aware of how their algorithms operate. Businesses should make sure that their algorithms comply with antitrust law by design. For compliance officers and legal counsel dealing with AI this means: talk to your technology departments to ensure that software is programmed to prevent any risks of collusion.

Dr. Falk Schöning
Partner, Brussels
T +32 2 505 0911
falk.schoening@hoganlovells.com

AI creates issues in many other market sectors, where it gives rise to new services, industries, data, and ethical issues. Some of these other sectors combine AI with space activity and yet others are experiencing similar issues that have been, or will be, experienced in the space industry. For all of these reasons, it is valuable to consider the use of AI and trends in other industries, such as connected cars, life sciences, robotics, smart homes, and FinTech.

Alicante
Amsterdam
Baltimore
Beijing
Birmingham
Boston
Brussels
Budapest
Caracas
Colorado Springs
Denver
Dubai
Dusseldorf
Frankfurt
Hamburg
Hanoi
Ho Chi Minh City
Hong Kong
Houston
Jakarta
Johannesburg
London
Los Angeles
Louisville
Luxembourg
Madrid
Mexico City
Miami
Milan
Minneapolis
Monterrey
Moscow
Munich
New York
Northern Virginia
Paris
Perth
Philadelphia
Rio de Janeiro
Rome
San Francisco
São Paulo
Shanghai
Shanghai FTZ
Silicon Valley
Singapore
Sydney
Tokyo
Warsaw
Washington, D.C.
Zagreb

Our offices
Associated offices

www.hoganlovells.com

"Hogan Lovells" or the "firm" is an international legal practice that includes Hogan Lovells International LLP, Hogan Lovells US LLP and their affiliated businesses.

The word "partner" is used to describe a partner or member of Hogan Lovells International LLP, Hogan Lovells US LLP or any of their affiliated entities or any employee or consultant with equivalent standing. Certain individuals, who are designated as partners, but who are not members of Hogan Lovells International LLP, do not hold qualifications equivalent to members.

For more information about Hogan Lovells, the partners and their qualifications, see www. hoganlovells.com.

Where case studies are included, results achieved do not guarantee similar outcomes for other clients. Attorney advertising. Images of people may feature current or former lawyers and employees at Hogan Lovells or models not connected with the firm.

© Hogan Lovells 2018. All rights reserved. 03887