LLOYD TEVIS
INVESTMENTS, LLC

Securing your future through Precision Investing™

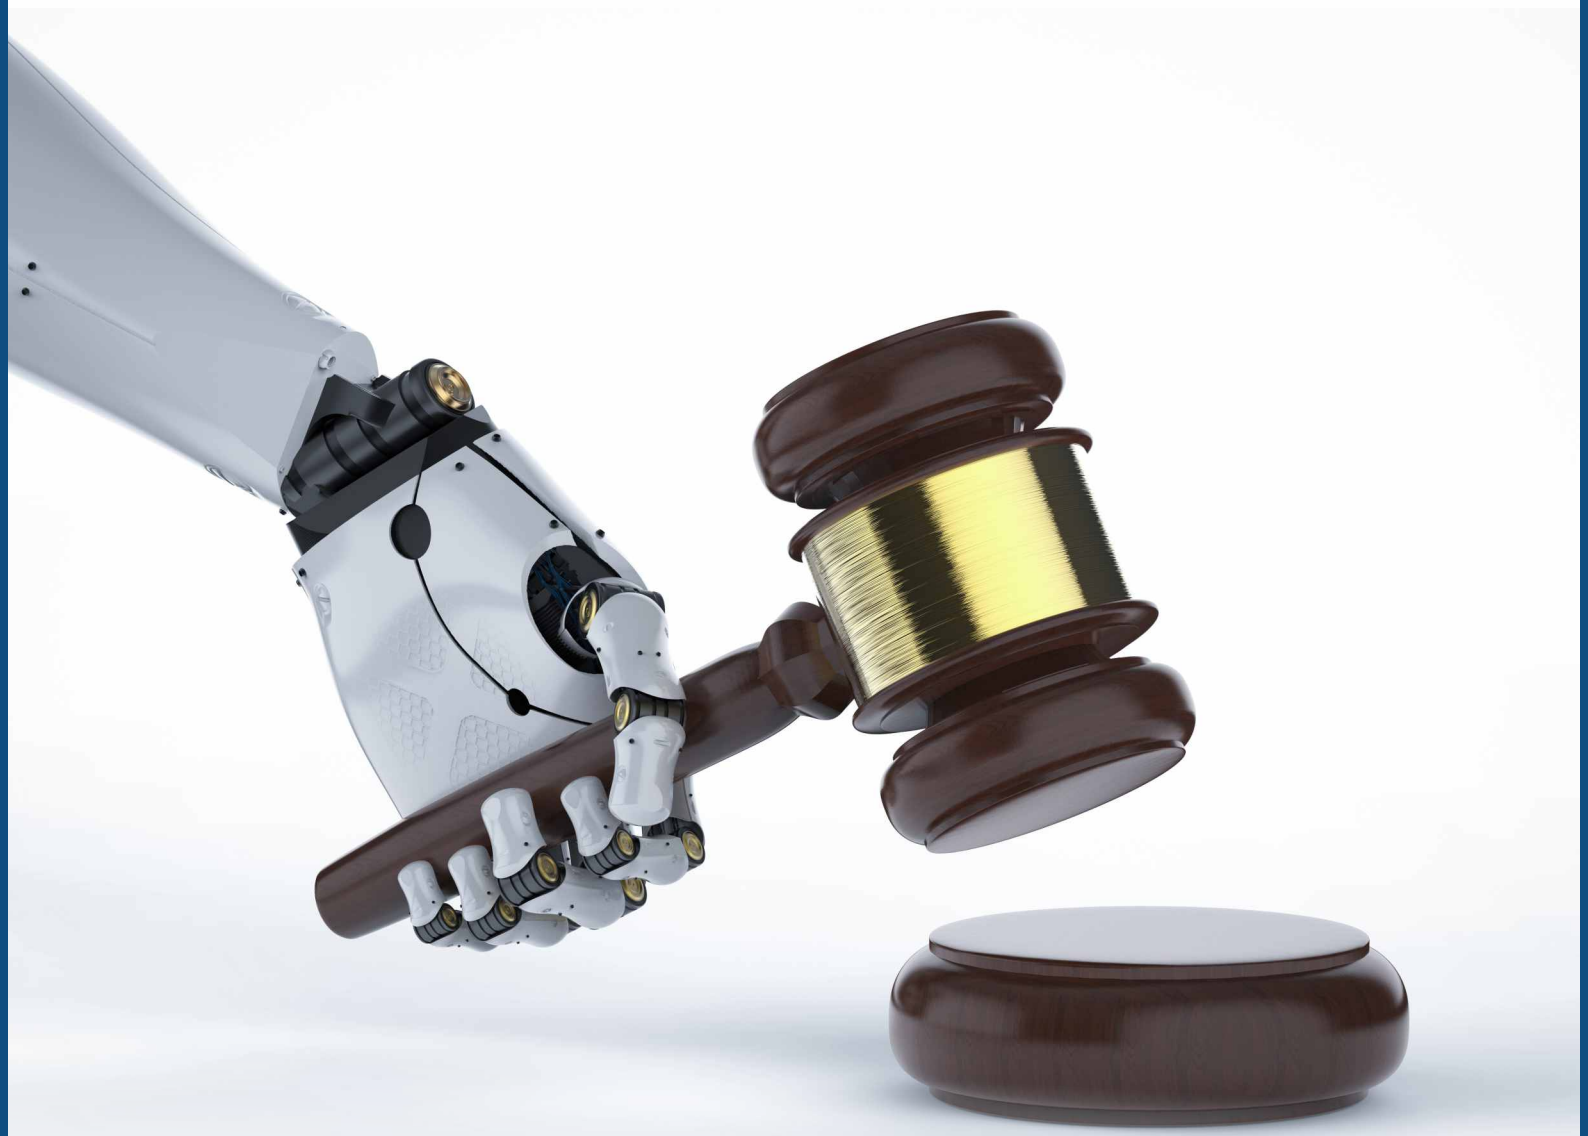

Regulatory Regimes for Artificial Intelligence

Legal Regulation of AI

Introduction

Recently there have been calls from many quarters to regulate AI technology. Here by regulate is mostly meant restrict or restrain as opposed to license. This is a rapidly emerging field in which understanding is still evolving rapidly. In this note we attempt to pull together the current state of play. Doing so requires us first to look at the law, next to consider AI and third to consider how the two might interact.

Existing Legal Frameworks

Let us begin by noting well settled law. If a machine is used properly and it damages persons or property then civil liability normally results. This liability will potentially attach to the manufacturer, the owner or the operator depending on the circumstances of the case. In general the particular technology involved is not an issue in assessing who is liable. Rather various duties of care and warrants of reliability are what is considered.

However, exceptions to this basic rule do exist. Certain technologies are considered inherently dangerous and extra burdens of care are imposed. For instance, impounding water behind a dam is considered an obvious danger. Absolute liability rather than negligence liability attaches to damages resulting from dam breaks.

Often with these inherently dangerous technologies a special regulatory regime is created. These regimes establish best practices, oversight and inspections to mitigate the danger of the technology. Compliance with the regulatory regime mitigates the liability exposure otherwise attaching to the manufacturer / owner / operator. Flouting of the regulatory regime, however, will normally enhance liability or remove legal defenses to liability. Generally it is understood that the dangerous technology confers benefits, but the default standards of liability create too much risk for the manufacturer/owner/operator and thus discourage putting in use of the technology. Alternately, it is recognized that the damage created by the technology may be beyond financial redress and that avoidance of an incident is of paramount importance. Overall the purpose of regulation is to create a defined space in which the dangerous activity may be permitted to proceed. There are many examples of such regulatory regimes: nuclear power plants, underground mines, pharmaceutical manufacture and sales, steam boiler manufacture and operation, civil aircraft manufacture and operation are just a few examples.

If a machine is used with deliberate purpose to harm persons or property then criminal liability may attach. Here the critical element is malicious intent and the liability falls directly on the holder of the intent. Normally this is the operator of the machine. For instance, when a gun is used to kill someone the homicide is charged to the trigger man rather than the owner or the manufacturer. Liability may, however, attach to the owner as an accessory. In most cases the manufacturer will be shielded from liability as long as the machine was produced and sold in a lawful manner. The basic framework of the law does not consider the particular technology involved. However, again there are limited exceptions. For instance, killing people with bombs is considered to exhibit an extra level of reckless disregard for life given the frequent death of uninvolved bystanders. Thus, misuse of a particular technology may act as an intensifier of an act already judged criminal.

Certain situations are both common and deleterious and thus have evolved somewhat specialized legal treatment. Perhaps the most obvious example is vehicular manslaughter.

Society adapting to new technologies has been going on for centuries. Usually the new technology is greeted with fear and suspicion, but as its benefits and detriments become better understood a regulatory regime is worked out. For instance, in the nineteenth century steam railroads were a new thing and sparks from the firebox might set standing crops on fire. Naturally this created hostility between farmers and railroad operators. However a compromise was worked out which involved requiring spark arresters on the smoke stack and cutting back flammable materials from right of ways. As a result railroads could be built and farmers benefitted by getting their crops to market. The duties of care imposed created costs for the railroads, but as these costs were common to all operators it was possible to recoup the costs in freight fees and thus spread the costs over the economic base directly benefitted by the technology.

We often tend to think of regulations as something of a band aid which gets applied after a problem emerges. As such the attitude is a bit of "thank god the government finally stepped in to save us from that problem." However, there are plenty of instances of premature regulations that have led to problems. A famous example is English regulation of mechanical vehicles on public highways. The regulations were framed around heavy drayage machinery drawing freight wagons at slow speed – a sort of road train. The regulations required a man with a signal lamp to walk ahead of the vehicle to give warning of the oncoming machine. The unintended consequence was to stifle development of passenger cars and thus to give Germany a head start in automotive technology. This lead persisted for decades after the English modernized their road regulations. The Germans lead in automotive engineering became an issue of national survival when cars evolved into tanks.

In the case of most existing machines there is no question of autonomous action on the part of the machine. With existing AI technologies also there is no real autonomy. However, autonomous robots or computer systems have been the subject of much fictional speculation. If we ever got to the point of having such autonomous machines (and it is generally assumed to be a question of when and not if) then seemingly new legal issues might arise. However, that is not really so. Already parents are in a sense the manufacturers / owners / operators of children. The law recognizes that children can act autonomously and with greater recklessness or stupidity than adults would. The law imposes parental responsibility upon the parents to supervise children and to not empower their reckless tendencies. One can easily see how such legal approaches could evolve to address autonomous machines. Similarly live stock and cell lines are autonomous beings. The law recognizes responsibility on owners to not permit straying of live stock. Operators of laboratories are responsible for keeping pathogenic cell lines under tight control and to not allow them to become a nuisance to the community. Autonomy is not a fundamentally new issue for the law.

In this context, we should take note of legal trials to decide whether capital punishment should be applied to dogs who attack people. Such trials often turn on questions of provocation by the victim and the character of the animal. One can foresee similar issues arising in future trials concerning lethally armed autonomous vehicles employed to guard large facilities. An important difference is that dog behavior is considered common knowledge and a responsibility to avoid provocation can be laid on victims. What an algorithm considers provocative may not be so much part of common knowledge.

Finally property and persons is not the total of what may be damaged. Reputations also may be damaged. Here speech and publication, being favored activities, the law generally narrows the definition of what is considered a cause for liability. Historically, marital relations were considered as something which the law should protect from third party damage. This concept appears to slowly be going out as emphasis shifts to the autonomy of spouses. However, loss of companionship remains a cognizable loss in certain civil liability assessments. Trespass over land is a long established legal tort. Building on this foundation, there is a recognized right to privacy. However, its boundaries and limitations remain an unsettled side of the law.

Generally law needs to evolve with society. There are two very different philosophies as to how this should occur. In the common law tradition, the law evolves through judges settling cases. As new situations arise judges mold and adapt the prior existing law and practice to address the new situation to conform it to community standards of justice and equity. General law is gradually built up from consideration of numerous particular cases. Eventually legislators may codify these evolved understandings into statute law to create uniformity and predictability in the application of the law. In contrast to this bottom up approach, the Roman (also known as Civil) law tradition offers a top down approach. Here the legislator articulates a very general, often rather abstract, legal principle. Legal scholars engage in academic study and debate to clarify the meaning and application of the principle in concrete settings and to resolve clashes of principles. Judges draw on this evolved legal understanding to settle the particular matter before them. Both approaches work in the sense that advanced economies have been successfully created on both foundations. Neither is perfect. Judge made law emerges slowly. Civil law can emerge quickly but can be out of touch with reality. American jurisprudence is basically a mix of both. American tort law has mostly emerged from a common law tradition. Constitutional rights and the law around them partakes much more of the abstractions of the Roman law tradition.

Law is inextricably bound to ideas of authority. Historically, law derived its authority from divine or royal decree. Even in this setting, however, there were areas of detail which did not interest the divinity or king and

which were left to somewhat self governing communities. Thus, the customs of merchants were recognized to govern contractual relations and they evolved into the current commercial code. Democracy has grafted on to this tradition by emphasizing a view of legislative authority even more extreme than traditional royal authority. However the exploding complexity of the modern economy has retained a role for self governing communities in the area of technical standards. Similarly, legislators have often recognized the need to delegate decision making power in complex areas to expert regulatory agencies. This delegation, incidentally, is by no means new – the patent office dates back to the inception of the Republic. Even in the middle ages rights in coats of arms (a forerunner of modern trademarks) were left to the decision of the heralds and the Marshall's court. Also long standing is the recognition that regulatory agencies, whether community or state sponsored, can act to protect established interests from competition and change, rather than to foster their prime missions of safety, reliability or efficiency.

What is AI

A number of issues have been raised about AI. To come to grips with them one must have some understanding of what AI is. Here it should be immediately noted that AI is a somewhat misleading convenience term. Strictly AI means artificial intelligence. As such it describes a strand of academic/commercial research and not a technology. As a research program AI has three different objectives

1. to endow computers with capacities resembling higher cognitive functions in humans as for instance vision.

2. by creating a mechanical model of intelligence to advance our understanding of what human intelligence is.

3. to create computer systems which perform tasks currently considered to require human intelligence.

The research program is an exciting one, but for a long time its results were sufficiently unimpressive that the field was described as “an attempt to get to the moon by climbing trees.” At the same time, the excitement of the program means that substantive forward steps can be over hyped. This field is exceptionally unlikely to achieve its goals as the result of one breakthrough discovery, but that circumstance is not generally understood.

Results achieved to date have been impressive along narrow fronts. AI softwares have been successful at recognizing faces, playing strategy games such as chess, taking dictation, scanning x-rays for tumors, diagnosing diseases and managing personal finances. A whole range of different technologies have been used, often in combination, to achieve these results.

Each result, however, has been achieved by a focused effort. What has not been achieved is so called general artificial intelligence. A general artificial intelligence would somewhat resemble the human brain in having a general capacity to learn any particular task. In principle a GAI might be applied to the problem of creating a more capable GAI and in this way rapidly evolve itself to have a greater than human intelligence. This, still fanciful, idea is known as The Singularity - the moment when machine intelligence surpasses human intelligence.

We would like to note that emergence of a GAI is not a given. Nature did not endow birds with powers of celestial navigation by evolving a general intelligence and letting it reason out how to do celestial navigation. Indeed a general intelligence capable of proceeding that way would be too heavy to get airborne on muscle power. Instead nature evolved a highly focused navigational capacity in birds which does not impose a significant load on their airframe. Nature proceeds by Darwinian survival of the fittest. However the economy also operates by survival of the most competitive technology. To the extent that commercial objectives drive the development of AI, that effort may prove to be focused in specific areas of application so that high performance may be achieved by with minimal resources. To be sure, the effort to create GAI will remain part of the research agenda. But even if successful, GAI may not be the main road forward with commercial applications of AI.

When people express concerns about AI they may be addressing the research program as a whole, but more typically they are expressing concerns about two technologies in particular: machine learning and large language models.

Machine learning is more an academic re-branding exercise than something new. Originally the field which focused on distilling data into information was known as statistics. The field got going at a time when computing was limited to tabulating machines. As a result the field built analytic power by creating refined mathematical models of data which could be executed with only simple computations. Once substantial computing power came along, some researchers wanted to attack data analysis with brute force computing rather than with refined mathematical tools. This group started calling themselves data scientists about ten years ago and currently data science is one of the college majors which is most rapidly growing in popularity with students. Within data science the process which statisticians refer to as model fitting is renamed as machine learning. In both cases, the idea is simply that raw data is reduced by some computational process into a more distilled form which takes meaning from the model the analyst selected for the data. In statistics this process is known as fitting a model whereas in data science it is known as training the model. The models used by data scientists are often identical to those used by statisticians, but in some cases they are newer models whose properties are not yet so well studied and delineated. The flavor of statistics is influenced by notions of mathematical rigor and the field tends to emphasize careful data analysis. Data science has a bit more of a pragmatic spirit – lets just turn the computer loose on the data and see if we get anything useful. However, these differences are not important for skeptics outside the field. It has long been recognized that statistics can be used not in a search for truth but as a polemical tool to put a scientific veneer on a set position. Data science can be abused in the same way, but the rhetorical benefit of so doing is less valuable as data science has already abandoned the mantle of mathematical rigor which gives statistics a certain authority.

Large language models are a model of language created by data scientists. Again the analysis of language had been begun by statisticians. In fact cryptanalysis, a centuries old field, is simply an application of statistically determined language structure to the practical problem of extracting meaning from coded messages. It is thus no surprise that language contains significant structure which can be exploited by a computational analysis. Large language models take the typical data science approach of using relatively simple models and applying massive computational power. The interesting discovery is how far that gets you. Large language models have endowed computers with a certain glibness in processing natural language texts and a facility for recognizing the emotional color present in the text. What large language models do not do is allow computers to understand text. Language is about words, whereas understanding is about ideas. For computers to possess understanding they would need to have an internal representation of ideas and be able to map from word sentences to ideas and back again. Naturally this is a field of AI research (known as semantic analysis) but it is not what large language models do. Without an understanding of text, computers can still achieve some moderately impressive results. They can mine text off the internet and assemble it into mostly accurate and relevant reports. They can translate with some facility between two human languages or between a human language and a computer language. But because they do not understand the texts they are working with they cannot detect errors or internal contradictions in the reports they create nor can they recognize mistranslations. This makes their current work product not especially reliable.

It is useful to compare the computer's performance to that of a human's to gain some sense of where the technology is currently at. By age three children have learned to prattle with some fluency but still limited understanding. Like dogs they have a reliable understanding of emotional coloration in what they hear and can produce colored utterances themselves albeit not yet with the level of nuance of adult humans. By age ten or eleven children can use a scissors, paste pot and encyclopedia to turn out reports on any number of subjects. But those reports may jumble together important and miscellaneous facts. And we would not expect such reports to contain careful discussions of contradictory opinions leading to a persuasive judgments as to correctness.

From the viewpoint of AI research large language models represent exciting progress over prior approaches to natural language processing. From the viewpoint of commerce the achieved results are useful. Every year vastly more material is created than the small number of human translators can keep up with. As a result foreigners's understanding of what another language group is talking about is limited to the small stream of material which makes its way through the translation bottleneck. Having machines create a rough draft

translation which a human translator can clean up expands the flow of material through the bottleneck. For certain text streams, the level of reliability computers can achieve may be adequate and in this case the bottleneck is removed entirely. Similarly each year commercial enterprises churn out vast volumes of reports, most of which have limited significance. Again computers can produce useful first drafts in an assistant role and in specific applications may achieve adequate reliability to produce final drafts.

These considerations imply that this area of research is going to attract considerable funding from commercial sources as well as academic attention. Rapid progress may result from this combination. As we have seen a major limitation of current systems is their reliability. Improving reliability will undoubtedly be a major focus of effort. One could seek to improve reliability by working on sematic analysis and driving towards GAI. Such an approach would represent progress at a fundamental level. Alternately, it may prove that faster progress can be achieved in specific application domains with simpler techniques. Commercially oriented research will likely orient around such less fundamental but more competitive approaches.

Along the road to GAI we can see a number of way points which must be achieved. A representation of ideas must be found, language must be translated to and from ideas, a standardized representation of ideas must be developed so researchers can collaborate, a capacity to reason with ideas is required, the ability to detect and define differences in ideas must be developed, an ability to weigh ideas and select the better one must be formed, the ability to fashion new ideas from old is required and finally the whole chain of processing must be explained to a human such a way that it can be verified. To be sure there is already work along all these lines. Some way points appear easier to achieve than others. But putting the entire pipeline together is still a considerable challenge. It is not going to happen in a few months or without first delivering intermediate results. By implication, we are going to see GAI coming and not suddenly be surprised by its arrival. Which is not to say that progress might be more rapid than we might guess at the present moment. But it might also prove to be as challenging as sustainable fusion – a century long research project where each breakthrough allows one to see just how challenging the next step is. Reliance on brute force methods is straightforward, but limits progress to a few well funded teams and paces the entire field at the overall rate of technical progress in delivering computational power. Clever ideas can be force multipliers by allowing more teams to participate and freeing the field from the limitations of hardware. However, clever ideas are few and hard to find.

Calls For Regulation

The eye catching progress large language models have made with natural language processing after decades of unremarkable incremental work has stirred speculation about where the field may be going and has generated calls for deployment of regulatory power to shape that speculative future.

One of the wilder concerns is that AI could represent an existential threat to the human species. In this nightmare GAI is achieved, rapidly progresses through the singularity, escapes from human control and decides that humans are a buggy experiment in natural evolution which is now best terminated. Clearly this is science fiction and not a proximate reality. Nuclear weapons are, by contrast, a real and current risk of something approaching human extinction. They are also a data point that suggests we can recognize and to some extent control existential technological threats. We have better things to worry about than the existential threat from AI.

Another fairly wild idea is that robots endowed with AI should be considered as sentient beings and acquire legal rights. Here I would like to call attention to my dishwasher. It is a robot (i.e. self directed machine) albeit one single mindedly focused on cleaning dishes. It communicates with me (not very well) through a weird set of boops and flashing lights. It also senses its internal state and notifies me that it is suffering from some form of mechanical flu today (error condition 12!) such that it has to shirk its work until its doctor pays a house call. Quite clearly it meets technical requirements for being a sentient being, but I do not consider it such. I am certainly not going to extend it legal rights. Sentient being is in fact a religiously inspired term which can be applied to meal bugs, a functionally less advanced being than dishwashers, but not to mechanical devices. The terminology in fact rests on a recognition of a common substrate of cellular life and a world experience shaped by pains and rewards. Recognition of this commonality of experience is what inspires the empathy for other living creatures and the hope that we could tread more lightly through our shared world. It would seem then

that meal bugs are farther along the road to claiming legal rights than any robot we are likely to see in our lifetime.

Worries less in the realm of science fiction are that AI is going to displace humans from jobs and leave them unemployable sitting by the side of the road. This concern has a nugget of reality to it.

Dishwashers have undoubtedly displaced many scullery maids from employment in domestic situations. But the roads are not ornamented with listless crowds of scullery maids. Humans are quite adaptable and can retrain themselves to do whatever needs to be done next.

AI will undoubtedly have an impact on the work place. Machines reading x-rays will reduce the number of radiologists which must be trained each year. This will displace a small number of extremely trainable humans into other activities. It will not create an economic crisis. Machines translating texts will remove a bottle neck, potentially promote translators into a more lucrative role of supervising machine translations and undoubtedly open up new jobs in reading and understanding the flood of newly translated material. AI also shows near term promise in fulfilling client support roles. If AI can effectively dialog with humans and relay useful information, it will be a welcome improvement on the current technology of traipsing through phone trees to an aggravatingly unrewarding destination. With a higher level of reliability it could credibly take over some remote support tasks from humans. The number of persons employed in call center support is certainly greater than radiologists and they may not be as retrain-able as radiologists. But they are still very retrain-able individuals and there are many jobs requiring similar skill sets of learning a particular body of knowledge and combining it with good people skills. Again we do not see an economic crisis in the making. In fact, we have empirical evidence on this point. In recent years many call center support jobs have been transferred overseas. As a result, some call center employees in the US became unemployed and some foreigners acquired obscure skills in making American small talk. However, massive employment problems did not develop in the US. The main impact of AI taking over support jobs may be to reduce viewer-ship of popular prime time American TV shows in Bangalore and Manila.

To create a shock in the labor market, a technology must rapidly displace a large population of fairly specialized workers and the skills they possess must be relatively nontransferable to other jobs. Weavers displaced by automated looms in the English midlands c. 1811 are the prime example of such a shock. Even in this case, weavers were only locally concentrated in the midlands and they did not represent a large component the total laboring class. AI may prove to have large, most likely transient, impact on the labor market. Currently that is a speculative hypothesis and as a creator of long term unemployment it is a hypothesis contrary to prior experience. A true GAI might be a more credible threat of permanently reduced employment. On the other hand, we must consider that around the industrial world workers are so miserable in their current jobs that they are not reproducing at a rate sufficient to maintain the current labor force. It may just as well turn out that economic deployment of GAI ends up being controlled by the need to sustain output in the face of human reluctance to do so.

Another concern about AI revolves around machine learning. Here the concern is that training sets will mirror the majority population and that products incorporating such technologies may serve minority subsections of the population poorly. This concern is reasonable but hardly new. Commerce always orients itself to the main market and serves the niche market as an after thought. Most industrial products are, for instance, designed for right handed people or for persons with normal color vision. Persons who are left handed or color blind are less well served and certain products designed for the majority population may even be unusable or frankly dangerous for such minorities. Machine learning does not appear to present a unique new danger in this regard. However, understanding and recognizing the limitations of new products and methods may take some time.

Motions Towards Regulation

Governments have been relatively eager to join the AI party and have bestirred themselves to study the matter.

The Europeans

Fastest to move has been the EU. The EU is a creature of the Civil law approach to regulation, which as we have noted lends itself to early sometimes unrealistic regulation. The EU is also motivated by what it views as its recent success with the Global Data Protection Regime. Frustrated by what it saw as an American failure to regulate data and privacy rights coupled to American companies voraciously acquiring and applying data about users and other members of the public, the EU passed a data regulatory regime which is intended to have global effect. Charged up by that action, the EU has been quick to move on AI and on June 14 2023 it passed the first reading of the AI Act. This law, extending to some 350 pages and including 771 Amendments, is testament to the rapidity of the EU legislative machinery.

The law moves, in good Civil law tradition, from the abstract to the concrete. It begins by positing a progressive ladder of harms and then legislating appropriate restraints. Every AI product must be assessed in terms of this ladder and its proper restraints determined. The topmost rung of the ladder is products considered to pose a threat to people and which therefore are banned. Examples given are

1. voice activated toys which encourage dangerous behavior in children

2. statistical models which classify people on the basis of behavior, socio-economic status or personal characteristics

3. real-time and remote personal recognition systems such as facial recognition systems.

We note that human intelligence is routinely deployed to recognize faces, to draw conclusions about people based on personal characteristics (“I'd like to date her”) and even to use voice to egg children on to activities which may prove risky for them (“you will learn to drive the car.”) As the EU is not banning these uses of intelligence, the legislative perception of danger, and a danger requiring in fact a total ban, seems to revolve around machines possessing these capabilities. Note that it is outright possession which is banned and not how the capability is used which draws legislative attention.

Moving down the ladder one finds dangerous but permitted activities. These include products already subject to safety regulation (a huge category) as well as biometric identification (other than facial recognition), facilities management, education and vocational training, law enforcement, representation and interpretation, and systems impacting employability, movement across borders or access to public services. It should be noted that some of these areas (e.g. movement across borders) is a perennial subject of EU legislation with a general motive of preventing member states from doing things contrary to the EU's fundamental treaties. So some of this legislation should be seen as an effort to restrict suspect entities (the member states) from applying AI to get around current regulations. Still it is interesting to see education classified as a potentially dangerous activity. Potentially dangerous AI products (e.g. a software to tutor one n a foreign language) must be assessed prior to being put in use and through out their life cycle. What exactly this assessment will amount to remains to be seen. Perhaps tutoring software which taught slang or omitted to teach regional dialects could find itself denied a sales license. Probably a software teaching French would have to conform to the Academic Francaise's standards for the language, which for instance disfavor loan words from other languages (and especially english.) Such words, for instance, often fail to fit into the French system of genders and conjugations and thus mess up the euphony of the tongue if they become prevalent.

Next down the ladder we come to products of more limited risk. Here AI products which generate text from public sources meet their regulator. Their products must be properly labeled (“produced by AI”) and footnoted (“based on Wikipedia”.) Here the goals are to protect the public through proper disclosure and to plug the AI product into the existing regime of intellectual property control.

Finally systems of limited risk which do not interact with intellectual property rights need simply be labeled appropriately. For instance, a photo retouched by a software program will probably require labeling.

Applicability of this regulatory system requires a definition of an AI system. The EU's approach is to be expansive and abstract so as to remain applicable even as technology evolves:

"Artificial intelligence system (AI system) means a system that is designed to operate with a certain level of autonomy and that, based on machine and/or human-provided data and inputs, infers how to achieve a given set of human-defined objectives using machine learning and/or logic- and knowledge based approaches, …"

Motivations for the AI are act are thus seen to be several. First, to protect the public from harm. Second, to protect existing regulatory regimes from being undermined by change. Third, by establishing a uniform set of rules to preempt national legislation. Fourth, to carve out a space within which this technology may be developed.

The AI Act is a classic piece of EU legislation in the best tradition of the Civil law. It will now be discussed for generations of legal academicians who will evolve many fine distinctions to guide the licensing authorities which will actually implement the law. All of this will take some considerable time and energy and likely stifle development of AI products by EU firms. It should be understood that is as much a goal as a feature of the EU system. Rather like the elves of Lothlorien the EU elite is pursuing their vision of the beautiful society and they do not much mind falling behind the world outside their borders – a world perceived as ugly and dangerous rather than varied and interesting.

The United Kingdom

In the UK they have had enough of EU regulation and they recently decided to escape it at the cost of shrinking their economy by at least a decades worth of growth and losing substantial influence over their major trading partner. The UK is of course a common law tradition rather than a civil law tradition, and the two traditions are as nearly incompatible as AC and DC circuits. Finely attuned to the ambition and menace of EU legislation creating global standards, the UK is keen to offer an alternative.

The UK's answer is guidance rather than legislation. It proposes to state five fundamental values and to urge firms to voluntarily comply with them. The UK intends a significant learning period to take place before any statutes are written to enforce compliance with the standards. By complying voluntarily with the guidance, firms are reasonably protected in their investment from any statutes which might get written. The UK does not envision new regulatory authorities to enforce standards or laws. Rather it expects existing regulators and courts to go about their usual business and when an AI issue presents itself the authority in question may, if it chooses, refer to the guidance to pick its way forward. Of course different regulators and courts might go in different directions and ultimately that could undermine the promise of protecting investment which is being held out to businesses to motivate their voluntary compliance. Accordingly, the UK contemplates a central coordinating body which will inform the various legal authorities of each other's decisions and seek to evolve the framework in a harmonious manner.

The values the UK identifies as central to its guidance share much in common with the EU:

1) safety, security and robustness

2) appropriate transparency and explain-ability

3) fairness

4) accountability and governance

5) contest-ability and redress

Well surely no one is against fairness, but reasonable people might disagree on what is fair. In true common law tradition, the UK intends the meaning off this value, in applications of AI, to be worked out in the context of specific cases and questions by the ordinary arbiters of fairness.

The core underlying fact of the UK's approach to regulation is that it is a constitutional monarchy with a fairly compact and homogeneous governing elite. Senior judges are few in number and of high quality. His lordship disposes of large and somewhat nebulous powers. The UK system is comfortable charging a single senior judge with a mission to determine if the Prime Minister has acted in an unbecoming manner – a noncriminal behavior nowhere defined in statute. A Prime Minister determined to have so acted would be expected to promptly resign and retire to the Outer Thule appointed as the home of disgraced politicians. In such a system a common understanding of values is broadly shared across the community and guidance is taken seriously.

The United States

The United States takes a different approach again. We genuinely love liberty and our first thought to any proposed restriction on an activity is to ask "where is the harm in permitting that." Even when potential harms can be identified, we will ask if perhaps the victim should have been more careful to protect themselves. Where the public can protect themselves, we are more comfortable asking them to do so than deploying state power to restrict the liberty of those who create the harm. We are a big society and we seek to keep participation in decision making broadly open. We are under no illusions as to our society sharing a common set of values or perspective. We know we do not now and never have. Nor do we have great respect for our judges. Even our supreme court judges are so routinely political placemen that a true legal intellect on the top bench is a noteworthy event. Lower court judges are often the second grade intellects who could not hold their own in private practice. The true lights of American jurisprudence are working advocates and not academicians or judges. Our preferred way to manage society is to leave it to its own devices until the necessity of government action is clear. Once necessity is granted a rather voluminous law is written as the result of much political horse trading. Inevitably this messes up whatever governing concepts the law sought to embody and results in a text where much detail remains to be settled. That translation from text to practice then takes place through a long extended and expensive judicial process. A goal and not just a feature of this system is to keep government as weak as possible while still meeting the core needs of society.

As would be expected, the American system is not rushing to deal with AI. President Biden floated the idea that perhaps text and pictures produced by AI should bear some sort of watermark. There was no particular enthusiasm for this idea. The right is against anything Biden proposes out of political tactics, while the would be Roman law dictators of the left thought it too modest. But the idea is a simple practical one which addresses a real issue and we would not be surprised if it ultimately is legislated as an amendments to an unrelated must pass bill.

Showing a livelier sense of political opportunity, Senate Majority Leader Schumer has put forward a plan to form nine (!) blue ribbon committees, each tasked with considering some aspect of AI. Members of these panels are to be drawn broadly from law, academia, government and industry. It would then be left to senior lawmakers and to lawmakers strongly interested in the subject to propose possible legislation. This legislation might bear some casual connection to whatever the blue ribbon committees had to say, but it well might not. Schumer sees the necessity of this approach not in AI itself but in the competition of regulatory regimes. He wants the US to get out ahead of other jurisdictions in proposing AI regulations. Schumer's aspiration is to set the new "rules of the road" both at home and hopefully for emulation abroad. Finally there are the American regulatory agencies. We may take the Security and Exchange commission as representative. The SEC's mission is keeping the capital markets safe and efficient for investors. The Commission has cogitated, concluded that AI might have something to do with this, and so it would like its budget expanded so it can hire staff to think about the matter.

Some American politicians have a taste for competitive regulation, but unless industry persuades itself that it is necessary to head the EU off from "grabbing the regulatory high ground," not much is likely to come of this train of thought. The likelihood is modest amounts of money will be spent spinning the government's wheels but doing nothing substantive until a palpable harm emerges and a constituency for doing something about it forms.

The Entrepreneurs

Finally there are the governments more prone to see AI as an opportunity than a threat. Israel is undoubtedly eager to deploy remote facial recognition systems to combat terrorism. China has undoubtedly identified AI

as a critical military and economic technology. On past form it has organized a whole of society effort. Its businessmen will be out to license whatever technology they can. What its businessmen cannot buy its spies will try to steal. Its academies and businesses will be active on joint R&D programs. Its promising students will be directed into the field. Its politicians and diplomats will be eager to lull the rest of the world into thinking its not doing very much, even as it seeks to race ahead as fast as possible. Its policemen and secret policemen will throw a mantle of security and secrecy over its efforts. That secrecy will lead to much waste of effort and perhaps slow China down from achieving its goals.

Assessments

Different governments; different approaches. Perhaps the most obvious take away is that the only thing weirder than ones own government is the other guy's government.

There is currently a gold rush going on in AI technology. The US has the initial advantage and looks likely to keep it. The Europeans are more concerned about domesticating AI than advancing it. The English have their own approach which will work well for them but perhaps not export as well as they hope. The Chinese will spend a lot of money and might overtake the US. Where the technology is going and how quickly it gets there is anyone's guess. Successes in fields we do not yet identify will probably come faster than we would expect. Challenges in areas we now identify may prove much harder to surmount than we currently understand.

Many of the fears about AI are overblown. Existing legal frameworks should prove adequate to deal with the new questions raised by AI with only modest tweaks. The UK is probably right that its best to let this field evolve for a while and see how far we can get with existing tools before trying to legislate prematurely.