

Algorithms and the Law

Jeremy Barnett, Adriano Soares Koshiyama and Philip Treleaven

ABSTRACT

In Law a company is treated as having the rights and obligations of a *person*. In this era of Artificial Intelligence (intelligent assistants, 'Robo'-advisors, robots, and autonomous vehicles) 'algorithms' are rapidly emerging as *artificial persons*: a legal entity that is not a human being but for certain purposes is considered by virtue of statute to be a natural person. Intelligent algorithms will increasingly require formal training, testing, verification, certification, regulation, insurance, and most importantly status in law.

For example, already in financial services Regulators require firms to demonstrate that trading algorithms have been thoroughly tested, demonstrate 'best execution' and are not engaged in market manipulation. Other interesting cases are healthcare algorithms; medical-assistant 'Chatbots' and patient screening systems which will increasingly dispense medical advice and treatments to patients. Regulators, who have traditionally regulated firms and individuals, are raising the status of 'algorithms' to 'persons'.

This paper discusses the emergence of 'Algorithms as *artificial persons*', with the need to formally verify, certify and regulate algorithms. Its aim is to start discussion in the Legal profession regarding the legal impact of algorithms on firms, software developers, insurers, and lawyers. This paper is written with the expectation that the reader is familiar with 'Law' but has a limited knowledge of algorithm technologies.

1. Introduction

The science fiction writer Isaac Asimov famously proposed "Three Laws of Robotics" (Asimov, 1950): 1) A robot may not injure a human being or, through inaction, allow a human being to come to harm; 2) A robot must obey orders given it by human beings except where such orders would conflict with the First Law; and 3) A robot must protect its own existence as long as such protection does not conflict with the First or Second Law. Later he added an additional law: 0) A robot may not injure humanity, or, by inaction, allow humanity to come to harm.

Fast forward, in 2007 the South Korean Government proposed a Robot Ethics Charter and in 2011 the UK Research Council EPSRC (Boden et al., 2011) published five ethical "principles for designers, builders and users of robots". More recently, the Association for Computing Machinery (ACM, 2017) in 2017 published seven principles for algorithmic transparency and accountability: 1) Awareness - Stakeholders of analytic systems should be aware of the potential harm that biases can cause to individuals and society; 2) Access and Redress - Regulators should encourage the adoption of mechanisms for individuals and groups affected by algorithmically informed decisions; 3) Accountability - Institutions should be held responsible for decisions made by the algorithms that they use; 4) Explanation - Institutions that use algorithmic decision-making are encouraged to produce explanations regarding both the procedures followed by the algorithm; 5) Data Provenance - A description of the training data and its potential biases should be maintained by the builders of the algorithms, with public scrutiny being a desirable feature as long as it not allows malicious actors to game the system; 6) Auditability - Models, algorithms, data, and decisions should be recorded so that they can be audited in cases where harm is suspected; and 7) Validation and Testing) Institutions should routinely perform use rigorous methods to assess their models and document those methods and results as well as encouraged to make the results of such tests public.

In this era of Artificial Intelligence (AI), 'algorithms' are rapidly emerging in Law as *artificial persons*. Already algorithmic trading systems (Treleaven et al., 2013) account for 70%-80% of US Equity trades. Apple, Google and Amazon provide 'intelligent' virtual assistants and *Chatbots* (Virtual assistant, 2017), such as Apple Siri, Google Assistant, and 'smart' devices such as Amazon Echo, Google Home and Apple HomePod that interact with speech. Numerous financial firms provide financial 'Robo' investment advisors (Robo-advisor, 2017). Baidu has a medical-assistant advisor currently running in China (Baidu Research, 2017). And Google, Ubur, Tesla and most car manufacturers are working on autonomous vehicles (Autonomous car, 2017). In response, governments and regulators are modifying national laws to encourage innovation, with lawyers and insurers scrambling to absorb the implications.

2. Algorithm Technology

The core underlying algorithm technologies are:

- **Artificial Intelligence (AI)** - AI and machine learning systems able to perform tasks normally requiring human intelligence.
- **Blockchain Technologies** - technology underpinning digital currencies and transactions, that secures, validates and processes transactional data.
- **Internet of Things (IoT)** - is the inter-networking of 'smart' physical devices, vehicles, buildings, etc. that enable these objects to collect and exchange data.
- **Behavioural and Predictive Analytics** – the analysis of large and varied data sets to uncover hidden patterns, unknown correlations, customer preferences etc. to help make informed decisions.

These four technologies are intimately linked: AI provides the algorithms, blockchain the data storage and processing infrastructure, IoT the data devices, and behavioural/predictive analytics are important for (human) behaviour analysis.

We will briefly review these technologies.

2.1 Artificial Intelligence Technologies

Artificial intelligence (AI) provides computers with the ability to make decisions and learn without explicit programming. There are two main branches:

- **Knowledge-based systems (KBS)** – are computer programs that reason, and knowledge is explicitly represented as ontologies or rules rather than implicitly via code. KBS can be subdivided into:
 - *Rule-based systems* – is one whose knowledge base contains the domain knowledge coded in the form of IF-THEN or IF-THEN-ELSE rules.
 - *Case-based Reasoning* - a form of so-called expert systems that bases decision-making on prior case experience, instead of on a pre-defined rule set.
- **Machine Learning** - is a type of AI program with the ability to learn without explicit programming, and can change when exposed to new data. Subdivisions include:
 - *Supervised learning* - is the task of inferring a function from labelled training data, where training data consist of a set of training examples.
 - *Unsupervised learning* - is the task of inferring a function to describe hidden structure from unlabelled data.

Other related 'intelligent' algorithms we should also mention include natural language processing (NLP) and sentiment analysis:

- **Natural language processing (NLP)** – the application of computational techniques to the analysis and synthesis of natural language and speech.
- **Sentiment analysis** - the process of computationally identifying and categorizing opinions expressed in a piece of text.

2.2 Blockchain Technologies

Blockchain technology originally conceived for Bitcoin and other cryptocurrencies are now recognized to have far-reaching potential in other areas. Blockchains are a way to order transactions in a distributed ledger, a record of consensus with a cryptographic audit trail maintained and validated by multiple nodes. Using this technology, many processes and third-party transactions are streamlined or collapsed entirely.

The core blockchain technologies are:

- **Distributed Ledger (DL)** – a decentralized database where transactions are kept in a shared, replicated, synchronized, distributed bookkeeping record, which is secured by cryptographic sealing. The key distinction between 'distributed ledgers' and 'distributed databases' is that nodes of the distributed ledger cannot/do not trust other nodes – and so must independently verify transactions before applying them. Distributed ledger technologies subdivide into two broad classes: those that seek to minimize the role of trusted and identifiable third parties; and those that explicitly rely on identifiable third parties for some subset of the system's properties.
- **Smart Contracts** - are simply the rules that participants have collectively signed up to that govern the evolution of the 'facts' in the distributed ledger. Possibly computer programs that attempt to codify transactions and contracts with the intent that the records managed by the distributed ledger are authoritative with respect to the existence, status and evolution of the underlying legal agreements they represent. When paired with a blockchain that records changes of asset ownership, they can act as a wrapper on a transaction that can automatically move value and can execute the terms of a contract; although some parts may require human input and control. Smart contract technology has the potential to automate laws and statutes.

The key attributes are: a) *Resilience* – blockchains operate as decentralized networks as opposed to a central server with a single point of failure; b) *Integrity* – blockchains operate using distributed open-source protocols removing the need to trust a third party for execution; c) *Transparency* – public blockchains have inherent transparency features, since all changes are visible by all parties; and d) *Unchangeable* – records in a distributed public blockchain are largely 'immutable', allowing applications and users to operate with a good degree of confidence.

2.3 Internet of Things (IoT)

IoT is becoming increasingly important as every device with an on and off switch will have a unique identity, a connection to the Internet, will communicate and be controlled by an *algorithm*. Devices range from individual lights in a smart building, your domestic appliances, to the national infrastructure. When a light fails an algorithm (Blockchain smart contract) runs, *UBERises* an electrician to come and fix the light; and then pays them.

2.4 Behavioural and Predictive Analytics

So called Big data analytics is the process of examining large and varied data sets -- i.e., big data -- to uncover hidden patterns, unknown correlations, market trends, customer preferences and other

useful information that can help organizations make more-informed business decisions. One of the most interesting areas for intelligent algorithms is behavioural and predictive analytics.

Behavioural analytics is a subset of Big data analytics which focus on finding out how and why people behave the way they do. Traditionally this has focussed on consumers, analysing eCommerce platforms, social media sites, online games, and any other web application. Now algorithms are expanding to all areas of data science from health to crime science. Closely related is predictive analytics the practice of extracting information from existing data sets in order to determine patterns and predict (possible) future outcomes and trends.

Fast forward, "Minority Report" is an action-detective thriller set in Washington D.C. in 2054, where police utilize algorithms to arrest and convict murderers before they commit their crime. An interesting concept for future regulation and for Law.

3. Virtual Assistant Algorithms

As discussed, a virtual assistant or a *Chatbot* are algorithms designed to simulate conversation with human users, especially over the Internet, where (machine learning) algorithms perform a particular task, like providing customer service or answering a question. Often the terms are used interchangeably.

Algorithms behind the technology have the capacity for learning, reasoning, and understanding. They range from search engines like Google, to increasingly sophisticated assistants such as Apple Siri or smart devices such as Amazon Echo/Alexa and Google Home. Regarding the Law and liability, if Google returns the wrong answer to a search, we try again. However, if Amazon's Alexa misinterprets a conversation or hears something on the television and make an expensive purchase, where does the Law stand?

Rogue Algorithms

Already rogue algorithms are emerging with advertisers abusing Amazon Alexa and Google Home. For example, Burger King 'hijacked' Google Home speakers by creating an ad that would trigger the devices to read its Wikipedia entry for the Whopper, which it had 'conveniently' edited beforehand to sound like marketing copy. Google quickly blocked the trigger, but not before the restaurant chain had gained a lot of free press; and Google consumer backlash.

Another 'entertaining' example is Microsoft's Tay, a Chatbot algorithm that was designed to learn from user interaction via Twitter; but had to be sent to bed early. Tay proved a smash hit with racists, trolls, and online troublemakers — who persuaded Tay to blithely use racial slurs, defend white-supremacist propaganda, and even outright call for genocide.

The Law

Although concern has been expressed about the urgent need to police and regulate these rogue algorithms, there exists through current criminal and civil laws, a considerable body of law that can be deployed where necessary. UK examples include the Data Protection Act 1998, The Consumer Protection against Unfair Trading Regulations 2008, and the Fraud Act 2006.

4. Financial Algorithms

The application of 'intelligent' algorithms has been driven to a large extent by the highly competitive financial services industry, starting with algorithmic trading (Treleven et al., 2013), and now the rise of financial Robo-advisors.

4.1 Algorithmic Trading

In electronic financial markets, algorithmic trading (AT) refers to the use of algorithms to automate one or more stages of the trading process: pre-trade analysis (data analysis), trading signal generation (buy and sell recommendations), and trade execution. Trade execution is further divided into agency/broker execution (when a system optimizes the execution of a trade on behalf of a client) and principal/proprietary trading (where an institution trades on its own account). Each stage of this trading process can be conducted by humans, by humans and algorithms, or fully by algorithms. AT usually involves learning, dynamic planning, reasoning and decision taking on the basis of logical inference from internal analytical, trading, risk and money management models.

Rogue Algorithms

Algorithmic trading due to its magnitude and proliferation has had a major impact on the financial markets, most notably the 2010 Flash Crash which wiped \$600 billion in market value off US corporate stocks in 20 minutes. However, the involvement and market impact of algorithmic trading to the flash crash is still the subject of much debate.

Another interesting example is that of market-making firm Knight Capital (Treleaven et al., 2013). On August 1, 2012, Knight Capital deployed untested software to a production environment which contained an obsolete function. The rogue algorithm started pushing erratic trades through on nearly 150 different stocks and lost \$440 Million in 30 minutes; destroying the company.

The Law

Trading algorithms are increasingly regulated. The fine legal line between algorithm-algorithm interaction that manipulate the trading of competitor algorithms and market manipulation has led to some industry practices being banned under legislation such as MiFID II (European Parliament, 2014).

4.2 Professional Robo-Advisors

Robo-advisors are a class of financial adviser software based on machine learning algorithms that provide professional financial advice or portfolio management with minimal human intervention, based on mathematical rules and AI algorithms. A typical Robo-advisor collects information from clients about their financial situation and future goals through an online survey, and then uses the data to offer advice and/or automatically invest client assets.

Rogue Algorithms

Warren Buffett when discussing the aftermath of the 2008 financial crisis, warned, "Wall Street's beautifully designed risk algorithms contributed to the mass murder of \$22 trillion." The potential problem is that today's Robo-advisor algorithms probably lack experience when it comes to managing assets during sustained periods of market turbulence and falling stocks prices.

The Law

Although Robo-advisors are registered with the Securities and Exchange Commission (SEC), they are not fiduciaries nor do they fit under the traditional standard applied to human registered investment advisors (RIA). And therein hides the problem for the investing public: Robos lack accountability for the investment decisions they make on your behalf.

4.3 Financial Regulation

Regulators are also exploring the use of algorithms to improve efficiency. Financial regulation is estimated to cost firms \$280b pa and involve 10% of the workforce (Treleaven and Batrinca, 2017). Regulators faces a myriad of pressures: increasing workload in monitoring small firms and individuals; cross-border cybercrime (e.g. AML, Binary Options); political pressure to curb excesses (e.g. Libor); escalating international and European Union regulations (e.g. MiFID II); governments relaxing regulations to increase competitiveness (e.g. US Dodd-Frank) etc. The monitoring challenges faced by regulators are illustrated by the UK Financial Conduct Authority (FCA). Previously, the FCA monitored

25,000 large and medium size firms. With essentially the same resources, the FCA now has to supervise an additional 30,000 small firms.

Hence Regulators are increasingly looking to automate compliance and regulation using AI algorithms.

Rogue Algorithms

As discussed, although the \$600 billion 2010 Flash Crash was initially blamed on algorithms it seems in retrospect more like due to a human trading error.

The Law

As discussed, financial services Regulators require firms to demonstrate that trading algorithms have been thoroughly tested, demonstrate 'best execution' and are not engaged in market manipulation. Regulators, who have traditionally regulated firms and individuals, are raising the status of 'algorithms' to 'persons'.

MiFID II introduces closer regulation and monitoring of algorithmic trading, imposing new and detailed requirements on algorithmic traders (in certain cases, even where they are exempt from authorisation under MiFID II) and the trading venues on which they trade; including regulated markets (RMs), multilateral trading facilities (MTFs) and organised trading facilities (OTFs).

The regime dictates that a firm engaging in algorithmic trading will have in place systems to ensure that the trading cannot create or contribute to disorderly trading on the market and to manage any such conditions that arise. The EU Authority ESMA has proposed regulatory and technical standards based on existing guidance such as the 2012 Guidelines on Systems and Controls in Automated Trading Environment (ESMA, 2012).

High Frequency Algorithmic Trading (HFAT) now have obligations to require time stored sequenced records and trading algorithms for at least five years. Further specific regulation has been introduced to deal with algorithms that are being used in market making, where protections for liquidity have been set down and technical standards are being considered to define a 'market making strategy'. It is thought that this will be a *principles* rather than *rules-based* approach.

5. Autonomous Vehicles and Robotics

Autonomous cars use a variety of techniques to detect their surroundings, such as radar, laser light, GPS, odometry, and computer vision. Advanced control systems interpret sensory information to identify appropriate navigation paths, as well as obstacles and relevant signage. Autonomous vehicles have control systems that are capable of analysing sensory data to distinguish between cars on the road, pedestrians and other potential hazards, which is clearly necessary for safe navigation.

Rogue Algorithms

An interesting recent video on the Internet shows how Tesla's autopilot algorithm attempts to predict a car accident before it happens. However, not everything is perfect for autopilot algorithms. Just about every autonomous vehicle company from Tesla and Google to Uber has had a number of car crashes, including the death of a Tesla driver. Some accidents are algorithm 'anomalies'; some caused by other driver's unpredictable behaviours.

The Law

Industrialised countries are adjusting their laws to accommodate autonomous vehicles. For example, the UK Vehicle Technology and Aviation Bill (UK Parliament, 2017) imposes liability on the owner of an uninsured automated vehicle when driving itself but makes provision for cases where the owner has made 'unauthorised alterations' to the vehicle or failed to update its software.

Difficult questions are already being posed as to what happens where there are collisions between two driverless cars, and both appear to have acted properly. Further ethical issues arise when for example a driverless car swerves to avoid a pedestrian but thereby causes a fatal accident. This type of question raises legal and philosophical issues that are beyond the scope of this paper.

6. Blockchain Smart Contracts

One of the most contentious algorithm technologies for lawyers is Blockchain smart contracts; frequently disparaged as “not smart, not contracts”.

But to quote Sean Murphy of Norton Rose Fulbright “Smart contracts in combination with distributed ledger technologies have the potential to automate an extensive array of transactions and services within the Service services sector. Legal compliance can be built into the program logic, providing a way of transacting that maximises operational efficiencies with the potential to reduce legal and regulatory cost and risk.”

Rogue Algorithms

Due to its embryonic state, smart contracts have yet to spawn rogue ‘contracts’.

The Law

With regard to Smart contract law, proponents fall into three ‘legal’ camps (Murphy and Cooper, 2016):

- **Code-is-contract** - those who promote the ‘code is contract’ approach (that is, that the entirety of a natural language contract can be encoded).
- **Hybrid Contract** - those using a hybrid smart contract model under which natural language contract terms are connected to computer code via parameters (for example, a smart contract template) that feed into computer systems for execution. This way a wet contract (code) is turned into a dry contract (paper based, enforceable in the courts).
- **Code as Business Logic** - those who see smart contracts as consisting of digitising performance of business logic (for example, payment), which may or may not be associated with a natural language contract.

7. Legal status of algorithms

Legal redress for algorithm failure seems straightforward. If something goes wrong with an algorithm, just sue the humans who deployed the algorithm. It may not be that simple: for example if an autonomous vehicle causes death does the lawsuit pursue the dealership, the manufacturer, or the third-party who developed the algorithm?

7.1 Black box Algorithms

Next we look at so-called black box algorithms; where its inputs and outputs are visible, without any knowledge of its internal workings. Many machine learning algorithms are referred to as ‘black box’ as they are unable to explain their decision making, which makes them often unusable for many financial consumer applications; where an unexplained rejection might be deemed discriminatory.

7.2 Algorithmic Dispute Resolution

An area already undergoing major innovation is alternative dispute resolution (ADR), especially automated online dispute resolution (ODR). Dispute resolution by algorithm can be broadly divided into: a) *Consumer ODR* – uses technology to facilitate the resolution of disputes between ecommerce parties, typically online suppliers and consumers; b) *Judicial ODR* - any means of settling ‘ordinary’ disputes where there is a hearing (using technology) but outside of the courtroom, such as divorce or personal injury cases; and what we refer to as c) *Corporate ODR* – the use of technology to manage

the resolution of any contractual disputes that may emerge from major multi-partner projects or financial transactions.

In consumer ODR, concerns have been expressed (Katsh and Rabinovich-Eny, 2017) about the use of algorithms by corporates against the interests of consumers. One early conclusion as to how this imbalance can be controlled is to investigate the software design and funding of the algorithm, to establish whose interest are being preferred during the operation of the code.

A more difficult issue arises where an algorithm fulfils a lawful instruction but does so in a manner which offends the law. A recent example is the decision of the Grand Chamber of the European Court of Human rights known as 'Sheilas Wheels' (Tobler, 2011). In this case, an algorithm calculated a premium that was favourable to the insurers' female clients, but was struck down as it contravened the prohibition on discrimination on the grounds of sex.

8. Algorithm Testing and Certification

In some sense, algorithms can be thought as ordinary software that before being deployed to the general public needs thorough testing. However, AI algorithms are different in that they adapt, learn and influence the environment without being explicitly programmed to do so. They will soon occupy an 'ecosystem' of interacting 'intelligent' algorithms making the testing and certification of algorithms an immensely greater challenge. As discussed, financial regulators are increasingly demanding the testing, verification and registration of trading algorithms and Robo-advisor algorithms.

The technologies required to fulfil these demands are broadly: algorithm testing/verification, algorithm certification and algorithm regulation.

Algorithm Testing

Algorithm testing is an investigation conducted to assess the risks of algorithm implementation and to provide a clear view for stakeholders on the quality of the product under test. Depending on the nature of the system, the techniques divides to Formal Verification and Cross-Validation:

- **Algorithm Formal Verification** - is the act of proving or disproving the correctness of algorithms underlying a system with respect to a certain formal specification or property. The verification is done by providing a formal proof on an abstract mathematical model of the system, which corresponds accurately to the nature of the system (usually known by construction).
- **Algorithm Cross-Validation** - is a set of techniques to assess how the results of an algorithm designed for prediction (AI algorithms, statistical methods, etc.) will generalize to an independent data set, that is, how well it will perform in practice. Generally, a predictive algorithm uses a training dataset with inputs (causes) and outputs (effects) and automatically learns a function (nexus) that maps their relationship. The aim of cross-validation is to run the same algorithm in an independent dataset (a.k.a. test set or holdout set), in order to evaluate potential risks (overfitting, sensibilities to noise, etc.) and measure its expected generalization accuracy.

Algorithm Certification

Algorithm certification is the act of ensuring that a particular algorithm conforms to one or more standards. Typical standards available for some algorithms are ISO 9000 (ISO/IEC 90003:2014, 2014), IEEE Standard for Software Quality Assurance Processes (IEEE Std 1012-2012, 2012) and FDA General Principles for Software Validation (FDA, 2002). Usually this process involves auditing whether the algorithm during the life cycle: (i) conforms to the protocoled requirements (e.g., for correctness, completeness, consistency, and accuracy); (ii) satisfy the standards, practices, and conventions; and

(iii) solve the right problem (e.g., correctly model physical laws) and satisfy intended use and user needs in the operational environment. However, for predictive algorithms such standards are still to be stipulated, although some institutions and research groups are trying to set some minimum standards (Amodei et al., 2016; Taylor et al., 2016; ACM, 2017).

Algorithmic Star Chamber

In the workplace, algorithms are rapidly becoming a judge & jury 'star chamber'. Uber only communicates with its drivers via algorithms that unilaterally decide on the level of revenue share, driver's rating, and whether to terminate employment; without the right of appeal. Online retailers live in fear of a drop in Google search-engine ranking, if they are judged by an algorithm to have done something fraudulently. And job applicants face 'CV-sorting' algorithms that may mysteriously rank them as unqualified for the job.

Gaps between the design and operation of algorithms, our understanding of their ethical implications, and the lack of redress can have severe consequences affecting individuals as well as groups and whole societies.

Ethical Challenges

Consider the following: you have been asked to code the algorithm for a self-driving car that is able to predict possible accidents. When a fatal accident appears unavoidable, does your algorithm: a) sacrifice the car; b) sacrifice the pedestrian; c) sacrifice passengers in other vehicles; or d) risk harming the occupants?

Less daunting ethical challenges in finance, might involve a Robo-advisor balancing the interests of the client and the firm, or an algorithmic trading system balancing the risk of making a large amount of money or potentially causing a 'flash crash' and damaging the markets.

Predictive Justice

In 1977, Anthony D'Amato asked the question, 'Can/should computers replace Judges?' (D'Amato, 1977). He asked what would be gained and what would be lost and quickly concluded that decisions on facts must be the province of humans. He pointed to aspects of human judgement that cannot be reduced to algorithms, and concluded that such systems will in all likelihood be used in 'pre-trial' stages only.

Predicting the outcome of court decisions through the use of algorithms is gaining traction, especially in France. Here software analyses of senior court decisions using natural language can predict for example that a lawyer is more likely to succeed with a dismissal application in Rennes (where the statistics are 35%) than in Versailles (12%). See also the UCL 'AI Judge Project' (Aletas et al., 2016).

Algorithms Regulator and new code

Behaviour can be regulated by the criminal and civil law, but in practice, where serious concerns have been identified about a certain market sector, a Regulator 'with teeth' is generally established. For example the Health and Safety Executive, backed up by a severe code of sanctions, has changed the behaviour of an entire industry. Although a case has to be made for specialist 'algorithm' regulator, as the use of algorithms is all pervasive, in not doing so, there would be risk that certain areas of the market would remain unregulated. The key issue is establishing the intention and behaviour of people but in certain cases, the law has to be strict. A good analogy is putting of unsafe products on to the market. The first step is to have a definition of safety: a safe product is defined.

The next step is to identify who in the supply chain has responsibility for breach. Producers and distributors responsibilities for corrective action vary depending on the circumstances. There are criminal offences which sometimes import the concept of 'strict liability', where the law mandates

that certain behaviour is criminal, regardless of the offenders' intention (such as driving through a red traffic light). However, in these cases, the law establishes certain statutory defences for example the offender took all reasonable steps in 'due diligence'.

In a recent paper calling for an FDA for algorithms, Tutt (Tutt, 2017) argues that US criminal and tort regulatory systems will prove no match for the difficult regulatory puzzles that algorithms pose. He concludes that algorithm regulation requires federal uniformity, expert judgment, political independence and pre market review to prevent the introduction of unacceptably dangerous algorithms into the market. He proposes an FDA for Algorithms to serve as an expert regulator that develops guidance, standards and expertise in partnership with industry to strike a balance between innovation and safety.

Algorithms as *artificial persons*

Another interesting concept is whether or not algorithms should be given a legal personality. As we know, a Legal person refers to a non – human entity that has a legal standing in the eyes of the law. A graphic example of a company having legal personality is the offence of corporate manslaughter, which is a criminal offence in UK law being an act of homicide committed by a company or organisation. Another important principle of law is that of Agency, where a relationship is created where a principal gives legal authority to an agent to act on the principal's behalf when dealing with a third party. An agency relationship is a fiduciary relationship. It is a complex area of law with concepts such as apparent authority, where a reasonable third party would understand that the agent had authority to act.

The current position is as follows: In *R. (on the application of Software Solutions Partners Ltd) v Revenue and Customs Commissioners* (2007) EWHC 971 (Admin) it was decided that a software program could not enter into a contract on behalf of an insurer. Similarly, a director of a company must be a 'legal person'. S156 A of the Companies Act is to be amended by the Small Business, Enterprise and Employment Act 2015 to provide that a person may not be appointed a director of a company unless that person is a 'natural person'. This followed press reports of a Venture Capital firm appointing an algorithm called Vital to vote on whether or not to invest in a specific company or not. (<http://www.bbc.co.uk/news/technology-27426942>).

As the combination of software and hardware is producing intelligent algorithms that learn from their environment and may become unpredictable, it is conceivable that, with the growth of multi algorithm systems, decisions will be made by algorithms that have far reaching consequences for humans. It is this potential of unpredictability that supports the argument that algorithms should have a separate legal identity, so that due process can occur in cases where unfairness occurs. The alternative to this approach would be to adopt a regime of strict liability for those who design or place dangerous algorithms on the market, so as to deter behaviours that appear or turn out to have been reckless. Is this a case of bolting the door after the horse has escaped?

9. Conclusions

This paper is written to stimulate discussion in the Legal profession concerning the emergence of 'Algorithms as artificial persons', with the increasing need on the one hand: to formally test, verify, certify and regulate algorithms; and the other to formalise their status in Law.

Jeremy Barnett is a regulatory Barrister and sits as a Recorder of the Crown and County Court, St. Paul's Chambers (www.stpaulschambers.com/jeremy-barnett-crime) and Gough Square Chambers London. With a background in advanced computing, he is currently involved in research and development of Blockchain and Smart Contracts.

Adriano Soares Koshiyama is a PhD Student in the UCL Department of Computer Science, funded by the Brazilian Government CNPq programme.

Philip Treleaven is Professor of Computing at UCL and Director, UK Centre for Financial Computing & Analytics (www.financialcomputing.org). UCL is a leading centre for AI and blockchain research (<http://blockchain.cs.ucl.ac.uk/>).

10. References

- Aletras, N., Tsarapatsanis, D., Preoțiu-Pietro, D., & Lampos, V. (2016). Predicting judicial decisions of the European Court of Human Rights: A natural language processing perspective. *PeerJ Computer Science*, 2, e93. Link for The Guardian article over it: <https://www.theguardian.com/technology/2016/oct/24/artificial-intelligence-judge-university-college-london-computer-scientists>
- Asimov, I. (1950). Three Laws of Robotics, Wikipedia, https://en.wikipedia.org/wiki/Laws_of_robotics.
- Association for Computing Machinery (ACM). (2017). Statement on Algorithmic Transparency and Accountability. https://www.acm.org/binaries/content/assets/public-policy/2017_usacm_statement_algorithms.pdf
- Autonomous car. (2017). Wikipedia, https://en.wikipedia.org/wiki/Autonomous_car
- Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., & Mané, D. (2016). Concrete problems in AI safety. <https://arxiv.org/abs/1606.06565>.
- Baker, T., Dellaert, T. (2017). Regulating Robo Advice Across the Financial Services Industry, *Iowa Law Review*, Vol. 103, U of Penn, Inst for Law & Econ Research Paper No. 17-11, May 2017.
- Baidu Research. (2017). Baidu's Melody: AI-Powered Conversational Bot for Doctors and Patients. <http://research.baidu.com/baidus-melody-ai-powered-conversational-bot-doctors-patients/>
- Boden, M., Bryson, J., Caldwell, D., Dautenhahn, K., Edwards, L., Kember, S., Newman, P., Parry, V., Pegman, G., Rodden, T. and Sorell, T. (2011). Principles of robotics. UK Engineering and Physical Sciences Research Council (EPSRC). www.epsrc.ac.uk/research/ourportfolio/themes/engineering/activities/principlesofrobotics/
- D'Amato, A. (1977). Can/should computers replace judges. *Georgia Law Review*, 11, 1277-1301, 1277.
- Katsh, E., Rabinovich-Einy, O. (2017). *Digital Justice: Technology and the Internet of Disputes*. Oxford University Press. 260p.
- European Parliament. (2014). Directive 2014/65/EU of the European Parliament and of the Council of 15 May 2014 on markets in financial instruments and amending Directive 2002/92/EC and Directive 2011/61/EU Text with EEA relevance. OJ L 173, 2014, p. 349–496. <http://data.europa.eu/eli/dir/2014/65/oj>
- European Securities and Markets Authority (ESMA). (2012). Guidelines on Systems and controls in an automated trading environment for trading platforms, investment firms and competent authorities. www.esma.europa.eu/document/guidelines-systems-and-controls-in-automated-trading-environment-trading-platforms
- Food and Drug Administration (FDA). (2002). General principles of software validation; final guidance for industry and FDA staff. Food and Drug Administration, 11.

- Fitzpatrick, K. K., Darcy, A., & Vierhile, M. (2017). Delivering Cognitive Behavior Therapy to Young Adults With Symptoms of Depression and Anxiety Using a Fully Automated Conversational Agent (Woebot): A Randomized Controlled Trial. *JMIR Mental Health*, 4(2), e19.
- IEEE Std 1012-2012. (2012). IEEE Standard for System and Software Verification and Validation, in Revision of IEEE Std 1012-2004, pp.1-223. doi: 10.1109/IEEESTD.2012.6204026
- ISO/IEC 90003:2014. (2014) Software engineering -- Guidelines for the application of ISO 9001:2008 to computer software. ISO/IEC JTC 1/SC 7 Software and systems engineering. 52p.
- Robo-advisor. (2017). Wikipedia, The Free Encyclopedia. Retrieved 11:30, June 23, 2017. <https://en.wikipedia.org/w/index.php?title=Robo-advisor&oldid=785832796>
- Murphy, S., Cooper, C. (2016) Can smart contracts be legally binding contracts? An R3 and Norton Rose Fulbright White Paper. www.nortonrosefulbright.com/files/norton-rose-fulbright--r3-smart-contracts-white-paper-key-findings-nov-2016-144554.pdf
- Taylor, J., Yudkowsky, E., LaVictoire, P., & Critch, A. (2016). Alignment for advanced machine learning systems. Technical Report 20161, MIRI.
- Tobler, C. (2011). Case C-236/09, Association belge des Consommateurs Test-Achats ASBL, Yann van Vurgt, Charles Basselier v. Conseil des ministres. *Common Market L. Rev.*, 48, 2041.
- Treleaven, P., Galas, M., & Lalchand, V. (2013). Algorithmic trading review. *Communications of the ACM*, 56(11), 76-85.
- Treleaven, P., Batrinca, B. (2017). Algorithmic Regulation: Automating Financial Compliance Monitoring and Regulation Using AI and Blockchain, *Journal of Financial Transformation*, 2017, vol. 45, 14-21
- Tutt, A. (2017). An FDA for Algorithms. *Administrative Law Review*, 69(83), 83-125. <https://ssrn.com/abstract=2747994>.
- UK Parliament. (2017). Vehicle Technology and Aviation Bill. Bill 143, www.publications.parliament.uk/pa/bills/cbill/2016-2017/0143/17143.pdf
- Virtual assistant. (2017). Wikipedia, [https://en.wikipedia.org/w/index.php?title=Virtual_assistant_\(artificial_intelligence\)&oldid=786918164](https://en.wikipedia.org/w/index.php?title=Virtual_assistant_(artificial_intelligence)&oldid=786918164)