

Artificial intelligence and civil liability—do we need a new regime?

Baris Soyer* and Andrew Tettenborn†

ABSTRACT

Artificial intelligence (AI) is almost ubiquitous, featuring innumerable facets of daily life. For all its advantages, however, it carries risks of harm. In this article, we discuss how the law of tort should deal with these risks. We take account of the need for any proposed scheme of liability to protect the existing values of tort law without acting as a barrier to innovation. To this end, we propose a strict liability regime in respect of personal injury and death, and a bespoke fault-based regime for dignitary or reputational injuries. For other losses, we take the view that there is no justification for introducing any new regime, on the basis that AI applications do not introduce substantial added degrees of risk that would justify departing from the existing scheme of liability arising under the current law of tort.

KEYWORDS: Artificial intelligence, tort liability, personal injury and death, dignitary injuries, damage to property, other losses

INTRODUCTION

If a job requiring creativity or intelligent decision-making can be carried out by a human, somebody somewhere has probably either compiled code to do it, or at least speculated on whether this might be possible. Thus computers today can not only draw, plan and design things, often better—and certainly more tidily—than we can. They can beat us at chess and Go (the latter being vastly harder than the former) and write passable literature and plausible Mozartian music. They can produce teaching materials, and use them to teach and examine students. Increasingly, they are becoming experts at predicting events and seeing patterns in them, thus allowing them to invade the human fastnesses of science.¹ And this is without smart homes, and the increasing omnipresence in them of Siri, Alexa and their competitors.

* Professor of Commercial and Maritime Law, Intitute of International Shipping and Trade Law, Swansea University, HRC Law School, Singleton Park, Swansea, UK. E-mail: B.Soyer@swansea.ac.uk.

† Professor of Law, Institute of International Shipping and Trade Law, Swansea University, HRC Law School, Singleton Park, Swansea, UK. E-mail: A.M.Tettenborn@swansea.ac.uk.

¹ For an excellent and readable account of these developments, especially as regards matters supposedly requiring creative efforts, see Marcus du Sautoy, *The Creativity Code: How AI Is Learning to Write, Paint and Think* (Harper Collins Publishers Ltd 2019).

What is important about these instances is that they involve the application not only of computer code but artificial intelligence (AI); that is, something done electronically that goes beyond a mere following of instructions. AI engages techniques enabling improvement in the accuracy, speed and scale of machine performance across complex or large data environments such as potentially to improve or supplant human performance.² Its applications typically, and increasingly, employ analytical methods such as machine learning,³ natural language processing, image recognition, neural network and deep learning⁴ to perform complex transactions or processes in a fairly rapid fashion.

One matter largely affected by this, even though the point may not be immediately obvious, is the ambit of civil legal liability in England and Wales under the law of tort⁵ for the operations of algorithms. It is that which forms the subject of this article.

Although the law has dealt with IT-based systems for some time,⁶ the ones it has encountered have hitherto normally been on a fairly low level of sophistication. As a result, they have required no particular adaptation or alteration of traditional legal reasoning.⁷ Stock recording, record-keeping, routine banking operations like the administration of BACS or SWIFT payments and direct debits, accounting tools, and so on may require voluminous code and enormous security precautions; but they still involve little more than applying complex flowcharts to complex situations. So too, for that matter, with operations such as automated trading, electronic presentation of international trade documents and electronic payment against them, and even most smart contracts. There is precious little of AI, as against simple electronic logic, about most of these applications. The computers that carry them out are in essence just super-efficient machines used by humans, and the traditional scheme of liability for negligence or breach of contract is generally adequate.

But this is changing. Electronic trading is now being aided by machine learning proper, under which algorithms update themselves in the light of based on experience. So too are a host of other matters: the electronic diagnosis of illnesses in humans and malfunctions in machines; the automated credit scoring of possible borrowers and the administration of their accounts; the electronic control of aircraft, ships, cars and other vehicles for transport and collision prevention; control of the environment to avoid dangers; the automated design of buildings and machines; and so on.

² More precisely, AI systems are systems designed by humans that, given a complex goal, act in the physical or digital dimension by perceiving their environment through data acquisition, interpreting the collected structured or unstructured data, reasoning on the knowledge, or processing the information, derived from this data and deciding on the best action to take to achieve the given goal. Such systems can either use symbolic rules or learn a numeric model, and they can also adapt their behaviour by analysing how the environment is affected by their previous actions. See generally a document from the European Commission's Independent High Level Expert Group on AI, *A Definition of AI: Main Capabilities and Disciplines* (2019), 6 onwards <https://ec.europa.eu/futurium/en/system/files/ged/ai_hleg_definition_of_ai_18_december_1.pdf> accessed on 25 January 2023. Compare the definition from Selmer Bringsjord and Naveen Sundar Govindarajulu, 'Artificial Intelligence', in *Stanford Encyclopedia of Philosophy* (2020) 'the field devoted to building artificial animals (or at least artificial creatures that – in suitable contexts – appear to be animals) and, for many, artificial persons (or at least artificial creatures that – in suitable contexts – appear to be persons)' <<https://plato.stanford.edu/entries/artificial-intelligence/>> accessed on 25 January 2023; and also Stuart Russell and Peter Norvig, *Artificial Intelligence: A Modern Approach* (4th edn, Pearson 2020), 2–5.

³ In simple terms, machine learning is the ability of a computer to modify its programming to account for new data and to modify its own operation accordingly. It involves computers running predictive models that can learn from existing data and use it to forecast future behaviour outcomes and trends. See, for a useful summary, D Schönberger, 'Artificial Intelligence in Healthcare: A Critical Analysis of the Legal and Ethical Implications' (2019) 27 *Int J Law Inf Technol* 171, 175.

⁴ Deep learning is in essence a branch of machine learning relying on complex statistical models and algorithms with multiple layers of parallel processing that loosely model the way the biological brain works. For a scientific analysis of the idea, see Jürgen Schmidhuber, 'Deep Learning in Neural Networks: An Overview' (2014) <<https://arxiv.org/pdf/1404.7828.pdf>> accessed on 25 January 2023.

⁵ This article will, for brevity, concentrate on the law of tort. But many of its conclusions will also be applicable to contractual obligations, particularly when the latter are implied obligations to take reasonable care.

⁶ According to LEXIS, Westlaw and the online version of Lloyd's Law Reports, the first reported English case mentioning the word 'computer' dates from mid-1965: *Burnett v Westminster Bank Ltd* [1966] 1 QB 742; [1965] 3 WLR 863.

⁷ Though not always. The tort of deceit traditionally requires a human being to be duped into acting in a particular way: it was not until 2007 that it was judicially confirmed in England that feeding falsehoods to a computer fell by analogy to be similarly treated (see *Renault UK Ltd v FleetPro Technical Services Ltd* [2007] EWHC 2541 (QB) [122]).

AI AND THE LAW OF OBLIGATIONS

The significance of these developments for civil legal liability is straightforward. The law of tort is currently predicated on human capability and moral agency. This is obviously inherent in the standard of the reasonable person which underlies the law of negligence.⁸ Furthermore, perhaps surprisingly much the same applies in practice even where liability is strict. This is because such liability normally only applies to a person deliberately carrying on some deliberate activity,⁹ or who takes no steps to stop a state of affairs which they had reason to know existed and had the opportunity to stop;¹⁰ in either case, issues of moral agency continue to arise.

With what can be called ‘dumb’ machines, this does not really matter. Where a machine, including a computer, intervenes in the causing of damage, the question is merely whether those handling or controlling it are personally at fault,¹¹ or occasionally (in cases of product liability) whether the machine they are using is itself defective.¹² Unfortunately, once we turn to AI, the process immediately becomes much more problematical. The workings of machines equipped with AI are not intended to be readily understandable, even on a general level, to people; indeed, many such devices are designed precisely to operate autonomously and independently of human supervision. In such a situation, the general rule in negligence, that there is no liability unless a human being ought to have spotted that something was wrong, gives seriously skewed results if applied unaltered. First, it means that in most cases of damage done by AI, a claimant will have either to go uncompensated or try to invoke some head of non-fault-based liability, such as product liability or liability under *Rylands v Fletcher*. And second, we now have the serious possibility that even if an error would give rise to damages when committed by a human, the self-same error having the self-same effect will create no liability if committed by a smart computer or an algorithm. This cannot be right.

THE WAY FORWARD: A SCHEME OF LIABILITY

When it comes to decisions taken, or actions initiated, by AI, there are at present no specific liability rules in place in England and Wales (save in one specialized case, automated vehicles),¹³ and therefore, the ordinary rules of tort will apply, with the results outlined above. Nor has much thought been given here to whether some different AI-specific regime, involving some specific combination of strict, fault and possibly hybrid liability, might be called for (though the government is said to be working on regulation in some areas¹⁴).

By contrast, there has been a little more development in the EU. Most recently, in April 2021, the European Commission proposed a fairly detailed EU regulation on the production

⁸ This is uncontroversial. See the classic authority of *Blyth v Birmingham Waterworks Co* (1856) 11 Ex 781.

⁹ Liability under the rule in *Rylands v Fletcher* (1868) LR 3 HL 330, for example, normally depends on some activity knowingly carried on land; so too strict product liability under Part I of the Consumer Protection Act 1987 assumes a deliberate act in manufacturing goods.

¹⁰ Nuisance liability, for instance, can arise where a defendant fails to take reasonable steps to deal with a state of affairs on their land: see for example *Goldman v Hargrave* [1967] 1 AC 645 and *Leakey v National Trust* [1980] QB 485.

¹¹ If the handler has taken reasonable care to employ an apparently sound machine, they escape.

¹² In the case of claims for product liability under the Consumer Protection Act 1987, Part I.

¹³ Under the Automated and Electric Vehicles Act 2018, ss.1–2, which came into force in April 2021.

¹⁴ For example, medical devices: see the Medicines & Healthcare Regulatory Agency’s report *Software and AI as a Medical Device Change Programme* (16 September 2021). The Information Commissioner’s Office is working on guidance in connection with AI and data protection <<https://ico.org.uk/for-organisations/guide-to-data-protection/key-dp-themes/guidance-on-artificial-intelligence-and-data-protection/>> accessed 25 January 2023.

and marketing of AI,¹⁵ but this largely concerned regulatory matters. Earlier, however, and more interestingly, the EU Parliament had in 2020 voted to approve an outline scheme specifically directed at civil liability.¹⁶ Though lacking in detail, this essentially provided for a series of different schemes of responsibility, dependent on the application in question and the risk attaching to it. There would be strict liability for AI mistakes arising from the use of a limited number of applications deemed ‘high-risk’ (a category to be listed in secondary EU legislation to be prepared and updated by a committee of experts, ethicists, anthropologists, sociologists and mental health experts). This would be limited to those applications whose autonomous operation involved ‘a significant potential to cause harm to one or more persons, in a manner that is random and goes beyond what can reasonably be expected’. Elsewhere, there would be fault liability, though subject to a reverse burden of proof;¹⁷ and as a coda, there would be a measure of strict liability for some untried applications that repeatedly caused serious damage. This looks at first sight convincing.¹⁸ However, as will appear below, it raises a number of its own problems—not least indeterminacy—and we do not ultimately find it attractive.¹⁹

In light of all this, what is the best way forward? One option, which should not necessarily be rejected out of hand, is to do nothing, simply leaving the law to sort out its relation with technology organically. Our view, however, is that at least in certain cases change is needed. For the reasons pointed out above, the current situation does not seem to reflect the values of a modern tort system, in that it not only skews the existing scheme of liability in a peculiar way, but draws what we see as arbitrary distinctions between equally deserving claimants.

THE STARTING POINT: UNDERLYING VALUES AND PRINCIPLES

With this in mind, our approach is for a number of amendments to the law, based on a number of values that deserve protection by a modern system of law of obligations and fundamental principles that we regard as essential when devising a novel liability regime.

First, as far as possible any scheme of fault liability should be harmonized across human and AI-based decisions. It seems anomalous to us that the same decision or determination should engender liability if human generated but not if AI based, or vice versa.

¹⁵ EU (2021/0116 (COD)). This proposed regulation is commonly known as the EU Artificial Intelligence (AI) Act and is currently deliberated at the EU level. Under these proposals, certain AI practices (such as systems with a significant potential to manipulate persons through subliminal techniques) would be prohibited (Title II, Article 5); and AI systems deemed high risk would only be permitted subject to compliance with certain mandatory requirements and an *ex ante* conformity assessment (Title III, Articles 6 and 7). Annex III listed a limited number of types of AI system that it said should be seen as high risk for these purposes, including those used for biometric identification or categorization of humans or for recruitment or selection for employment; those set up to decide eligibility for social security; and polygraphs and similar devices. See generally, ‘Commission Proposes New Rules and Actions for Excellence and Trust in Artificial Intelligence’ (2021) 403 EU Focus 27; in more detail, Chris Holder, ‘The Commission’s Proposed Artificial Intelligence Regulation’ (2021) 27 Comp Tel L Rev 130, and Vera Lúcia Rapso, ‘Ex Machina: Preliminary Critical Assessment of the European Draft Act on Artificial Intelligence’ (2022) Int J Law Inf Technol 88.

¹⁶ ‘Civil Liability Regime for Artificial Intelligence’ (2020/2014 (INL)) (4 May 2020, adopted 20 October 2020) (‘the EU Parliament report’).

¹⁷ See the draft Regulation attached to the EU Parliament report, Articles 4 and 8.

¹⁸ For a useful commentary, see Nikos Th Nikolinos, ‘Adapting the EU liability Regime to Artificial Intelligence (AI): The European Commission’s Proposed Policy Options’ (2022) 28 Comp Tel L Rev 66.

¹⁹ Largely because in our view, it is not so much the degree of abstract risk that should determine the liability regime, as the claimant’s interest at stake. In fact, it has been often emphasized by legal theorists that adopting a particular normative value is bound to influence the nature and type of regulation. See, for example, Timothy Caulfield and Roger Brownsword, ‘Human Dignity: A Guide to Making in the Biotechnological Era’ (2006) 7 Nat Rev Genet 72; Kieran Tranter, ‘The Law and Technology Enterprise: Uncovering the Template to Legal Scholarship on Technology’ (2011) 3 Law Innov Technol 31. Similar points are also made by various contributors in Roger Brownsword, Eloise Scotford and Karen Yeung (eds), *The Oxford Handbook of Law, Regulation and Technology*, (OUP 2017).

Second, we see it as important that there should also be a reasonable ‘fit’ between the existing law and any new law. This carries two corollaries. One is that any scheme should run in parallel to existing heads of liability which may affect AI, such as (for example) product liability under the 1985 EU Product Liability Directive, as reproduced in the Consumer Protection Act 1987.²⁰ We do not take the view that this would lead in practice to wasteful duplication of legal resources;²¹ indeed, elsewhere in the law alternative pleading of similar but technically distinct heads of liability is commonplace.²² The other is that in so far as liability for AI would operate in parallel to some existing analogous scheme of non-AI liability, there should be a presumption that the two schemes should be as near as possible similar in effect.

Thirdly, third-party victims who are unusually vulnerable or otherwise, particularly at risk from AI-type decisions ought to be assured of being compensated.

Fourthly, in marking off strict and fault liability (and for that matter liability and non-liability), we need to pay large attention to the interests we are seeking to protect, and where necessary recognize that different kinds of harm may deserve different treatment. After all, in the law of tort as a whole, we regularly distinguish personal injury, dignitary interests, property damage and pure financial loss;²³ and there is no reason to treat AI any differently.

Fifthly, we should not necessarily reject carve-outs. For example, while we may not want any general strict liability for property damage, we should not be closed to the idea of making a limited exception for, say, self-driving vehicles.

Sixthly, while detailed safety regulation is best left to relevant technical regulatory regimes rather than civil liability rules aimed at fair compensation rather than beneficial behaviour change,²⁴ we see a strong policy case for two prudential civil protections for defendants. For one thing, we must be careful lest any liability directly or otherwise²⁵ encourage wholly disproportionate expenditure on precautions. Further, we should avoid imposing liability which might make experimentation and development excessively risky,²⁶ since in a very new field without claims experience, this may well make developers essentially uninsurable.

Lastly, while we obviously wish to preserve the effect of consumer protection considerations and any relevant unfair contract terms legislation, we would otherwise preserve the freedom to contract out. In a field where we hope competition and efficiency considerations will drive innovation, we see no reason why the rules we suggest should be anything other than default rules applicable only in the absence of contrary agreement.

²⁰ On which, see Tiago Sérgio Cabral, ‘Liability and Artificial Intelligence in the EU: Assessing the Adequacy of the Current Product Liability Directive’ (2020) 27 *Maast J Eur Comp L* 615. This will often not be very relevant, AI not being a ‘product’ as such: but it might sporadically be. This might be where AI was incorporated into a tangible movable at the time of its production, or possibly where a product utilizes AI supplied over the Net in order to function, as with a satnav (as suggested in Tiago Sérgio Cabral, *ibid*, 619). But pure AI, of itself, is almost certainly not a ‘product’, since it is not tangible.

²¹ Notably, if the proposed new liability regime was to channel liability to a specific party, such as the manufacturer of the AI product, this would not necessarily reduce the need for others in the chain, such as operators, to purchase their own liability insurance. Given the lack of useful data on AI and non-AI risk, it seems unlikely that the extra expense would be compensated by a correspondingly lower premium. Further, even with channelling of liability to a manufacturer, there will still be instances where manufacturers bring contribution proceedings against others. In such a case, those others would still need to carry liability insurance.

²² For example, the common phenomenon of rolled-up complaints of negligence and nuisance; misuse of private information and GDPR liability; or product liability at common law and/or under the Consumer Protection Act 1987.

²³ In this connection, we can cite not only the obvious example of the economic loss rule in negligence, but the rules that nuisance only protects a person in their enjoyment of land, that the law of libel does not compensate personal injury, and that strict product liability ignores damage to profits or business property.

²⁴ See, in particular, Roger Bowles, *Law and Economy* (Marin Robertson 1982), Ch 7; Richard A Posner, *Economic Analysis of Law* (6th edn, Aspen Law & Business 2003), Ch 6. In fact, as mentioned above, the EU is already proposing a regime to that effect: EU (2021/0106 (COD)).

²⁵ For example, through insurance premiums. See, in particular, Gary T Schwartz, ‘Reality in Economic Analysis of Tort Law: Does Tort Law Really Deter?’ (1994) 42 *UCLA L Rev* 377.

²⁶ Compare Jonathan Morgan, ‘Torts and Technology’ in Roger Brownsword; Eloise Scotford; Karen Yeung (eds) *Oxford Handbook of Law, Regulation and Technology* (OUP 2017), 522–543.

THREE TECHNICAL QUESTIONS. AI AND FAULT: ALGORITHMS AND THE PROBLEM OF AGENCY: AND THE MEANING OF 'DEFECTIVE' AS REGARDS AI

Before we come to the details, we need also to discuss three important technical matters. One, arising in the context of fault liability, concerns the common law rules about negligence and how far they can be appropriately applied to decisions taken by AI. The second concerns the question of agency, and who should bear legal responsibility for decisions made by AI; and the third is a question arising in connection with product-liability-style strict liability, namely, what counts as defective AI.

AI and fault

Under the present law, the notion of fault in negligence law presupposes an act or omission, coupled with a human mental state failing to measure up to a particular standard (or at least its negative, blameworthy inadvertence).²⁷ Although an employer will be liable for the fault of its employee or someone else in an analogous position,²⁸ failure by a mere machine under its control, such as a computer set-up to supervise safety, detect risks or carry out some similar function, will not suffice. There must be a showing of personal fault by a person, such as an employee, controlling or programming it.²⁹

At least two factors suggest that this will become problematic in so far as decisions of a similar type come to be taken interchangeably by AI and humans. For one thing, the need to compensate equally deserving claimants consistently suggests that, assuming a decision is visibly wrong, it cannot be right to dismiss one claimant's complaint if the determination was made by algorithm but allow suit by a fellow claimant in the same position where it was made by a human. (And indeed conversely, horizontal equity between defendants suggests that an employer liable for the error of an employee should not escape responsibility for exactly the same mistake should it be taken through the medium of AI where no human can be shown to have been at fault.)

If this is right, then however incongruous it might seem, we need to find some means of ascribing blame to a machine. A short while ago, the European Parliament toyed with the idea of ascribing some kind of legal personality to certain things endowed with AI.³⁰ But this outré and over-complex solution makes little sense. This is not on the basis that a machine cannot be a moral agent (itself a very controversial proposition, but in any case largely irrelevant since even if it is true there is precious little of personal morality in negligence law anyway). Rather, it is for a more practical reason. A legal duty to compensate, and ownership of the money to do it with or the ability to contract for insurance cover, have in practice to be vested in either a human or a legal person acting through a human.³¹ Without this connection with humankind, the notion of fault liability becomes unworkable.³²

²⁷ See generally, Andrew Tettenborn (ed), *Clerk & Lindsell on Torts* (23rd edn, Sweet & Maxwell 2022), paras 7-157–7-236.

²⁸ *Ibid*, Ch 6. For who is equivalent to an employee despite not technically being one, see generally *Cox v Ministry of Justice* [2016] UKSC 10; [2016] AC 660.

²⁹ There is very little authority here, the point one suspects being regarded as so obvious as not to need stating. But see *Target Holdings Ltd v Redfern* [1996] 1 AC 421, 432 (Lord Browne-Wilkinson), where payment due to a computer failure is treated as a matter of course as being non-negligent. In the context of health care, the point is made in Daniel Schönberger, 'Artificial Intelligence in Healthcare: A Critical Analysis of the Legal and Ethical Implications' (2019) 27 *Int J Law Inf Technol* 171, 197.

³⁰ European Parliament Resolution with Recommendations to the Commission on Civil Law Rules on Robotics (2015/2103(INL)) (European Parliament, 16 February 2017), para 59(f).

³¹ True, one could say a corporation is just as artificial a seat of personality as an intelligent machine. But the property-holding powers of a corporation make limited sense except in the context of detailed rules as to capitalization, agency and powers of disbursement. All these are matters somewhat incongruous as regards an algorithm. A similar point has been made by Simon Chesterman, 'Artificial Intelligence and the Limits of Legal Personality' (2020) 69 *ICLQ* 822, 822–827.

³² This point is not new. See, for example, the interesting and perceptive article in Bartosz Brożek and Marek Jakubiec, 'On the Legal Responsibility of Autonomous Machines' (2017) 25 *Artif Intell Law* 293; also Roger Brown, 'Property Ownership and the Legal Personhood of Artificial Intelligence' (2021) 30 *1 C T Law* 208.

A simpler and more elegant legislative solution in the case of a decision taken by AI, which we would advocate here, is to concentrate not on the means by which the decision was reached, but on the decision itself. Faced with such a decision, we should simply require a court to ask whether that determination would have been found to have been negligent if taken by a person in comparable circumstances. If it would, then the legislation should provide for the automatic liability of the person responsible for the machine, notwithstanding a lack of personal fault in that person or what would otherwise be seen as reasonable reliance on an apparently faultless piece of machinery.

The reference just made to the ‘person responsible’ for the actions of the algorithm, something not yet defined, neatly leads to the next question, agency.

Algorithms and agency

To render workable the concept of the negligent machine referred to above, and also the suggestion appearing later that in at least some cases there ought to be liability without fault for harm caused by a machine operating an AI process, we have to fashion a legal rule of agency, or ascription. Assuming a machine is seen as at fault, or a decision taken by AI ought to give rise to strict liability, who pays? As mentioned above, it cannot be the machine itself, an entity with neither the propensity to own assets with which to satisfy an award of damages, nor the ability spontaneously to insure itself against the possibility of having to pay them. (Nor, unlike a human or a corporate entity, can it enter into a contract of employment or agency such as will make an employer liable for acts or omissions committed in the course of employment.)

How should we get around this? The best way, it is suggested, is by legislative ascription of liability, the underlying aim of which should be a rough correlation of benefit and burden (in other words, those who substantially benefit from the utilization of AI should be the ones to pay when it goes wrong). To achieve this aim, it is suggested that we should first single out parties who exploit AI with a view either to profit, or to the carrying on of some other substantial but non-commercial economic activity,³³ thus excluding purely private and personal uses.³⁴ Assuming this criterion is met, two categories of defendant, it is suggested, should be included.

The first, which we might call ‘users’, consists of parties who incidentally incorporate AI, whether created by themselves or by a third party, into some aspect of the running of their business or operations. Examples would include an employer using it to control on-site safety in a factory or mine; a broker or shipowner employing it to take trading decisions, or to monitor a vessel at sea; or a municipality making use of AI to control traffic through variable speed limits and road closures.

The second, perhaps best christened ‘suppliers’, comprises businesses that create, supply or control those AI facilities. We are thinking here of, for example, specialist software creators, or IT suppliers making AI-based solutions available to businesses, or being contracted to operate them once supplied.³⁵

³³ For example, the activity of a charity, hospital or municipality. They are not run for profit, but operate very much in the economic sphere. The conflation of such bodies with profit-making organizations is well established elsewhere. Compare, for example, in England and Wales s. 14 of the Unfair Contract Terms Act 1977, which is limited to business activity but then goes on to say that ‘business’ includes a profession and the activities of any government department or local or public authority’. Similarly, the defence under the EU Product Liability Directive, that the defendant never supplied the goods in the course of a business (see Article 7(c)) has been held not to encompass supply by a publicly funded hospital: see *Veedefald v Århus Amtskommune* [2003] 1 CMLR 41.

³⁴ Likely to be rare in any case, except in the instance of things like motor vehicles, which are subject to their own regime anyway.

³⁵ These categories very roughly correspond to the ‘frontend’ and ‘backend’ liability suggested by the European Parliament, though with some differences in detail. See the EU Parliament report at paras 12–13.

It should be noted that in this connection, we see no objection to more than one entity being seen as responsible for a given AI application. For example, if a serious underground mining accident was caused by a failure of AI-equipped surveillance equipment, we see no reason why there should not be a liability in both the mining company that made use of the equipment and also the specialist IT contractors who supervised its day-to-day operation. If liability arose, the claimant should be able to sue either or both, with if necessary a claim for contribution³⁶ between them.

Defective AI

Just as strict product liability under Part I of the Consumer Protection Act 1987³⁷ necessitates a definition of a defective product,³⁸ any parallel liability for harm done by AI demands at least some definition of defective AI.³⁹

The analogy of the position under the Consumer Protection Act 1987 is of some help here. Under this regime, a product is defective, and thus potentially liable to inculcate its producer, ‘if the safety of the product is not such as persons generally are entitled to expect.’⁴⁰ Such a definition has two good points as regards AI. For one, it rightly stresses the objective nature of safety and its complete independence from any idea of fault or foreseeability of harm. In addition, it also correctly plays up the strong social interest in favour of promoting people’s reasonable expectations of security, without reference to the position of those responsible for the use or production of that which endangers them.

However, the parallel is not exact. With AI, which differs from a paradigmatic product by being a purely intellectual creation without the familiarity or physical footprint of a car or a hairdryer, it also raises two particular difficulties.

First, AI is both very new and very unfamiliar. For the foreseeable future, it may therefore be difficult for a court to reckon what degree of ‘safety’ the public expects, or has a right to expect, from it.⁴¹ This being so, it might be better to define defectiveness by reference to function—does the algorithm do what it was being used to do?—rather than expectations of safety. Even if we do not as yet know what degree of safety we can legitimately expect from an AI application, we can at least have a good idea of what it is meant to do and whether it is doing it.

Secondly, there is a difficulty of timing. With defective products, the invariable cut-off is the time of production: problems appearing later, unless mere manifestations of some pre-existing defect, are out of account.⁴² But the utility of AI lies precisely in its adaptivity; that is, in the ability of the algorithms embodied in it to change and improve in the light of experience, without the original producer necessarily having the final word. It may, therefore, be appropriate to impose liability for miscalculations by AI even in respect of subsequent malfunctions in taking account of experience.⁴³ The point might well be significant. Take, for example, a case of AI used to preserve safety in a dangerous environment such as a mine or quarry. If after a number of unexpected incidents, it failed properly to take these new dangers into account and as a result,

³⁶ Under the Civil Liability (Contribution) Act 1978.

³⁷ Legislation in turn based on the EU Product Liability Directive (Council Directive 85/374/EEC of 25 July 1985), but remaining in force without change post-Brexit.

³⁸ See s 2(1) of the Consumer Protection Act 1987 (‘... where any damage is caused wholly or partly by a *defect* in a product ...’) (italics supplied). This can on occasion be far from straightforward: see, for example, the coverage *Clerk & Lindsell on Torts* (n 27), paras 10-55 to 10-66.

³⁹ For the obvious reason that a generalized liability for any harm caused by any AI decision whatever is entirely unworkable.

⁴⁰ See s 3(1) of the 1987 Act, anglicizing the original EU Product Liability Directive, Article 6.

⁴¹ Nor can it be said that there is an industry standard in the light of which defectiveness can be reckoned. These issues are doubtless at present occupying the minds of both theorists and regulators at national and EU level; but for the moment we are far away from imposing national, let alone international, standards.

⁴² See the Product Liability Directive, Article 7(b).

⁴³ A point neatly made in *Cabral* (n 20), 624.

an accident occurred which would not otherwise have done so, there would be a strong case for liability.

It is tentatively suggested that as regards defective AI, we need a definition on the following lines, which would capture both AI with an unexpected glitch and AI which was simply inadequate as designed. AI in a device should be regarded as defective if (i) it does not perform, or fails to perform properly, the function reasonably to be expected of it, and (ii) as a result of this failure, it is likely, or more likely, to cause damage to a claimant.

THE PROPOSED SCHEME

So far, we have been laying down the parameters that we think ought to underlie any scheme of AI liability. In this section, we draw these matters together, and in light of them propose a new liability regime.

It will immediately be noticed that we have not chosen to go down the EU route, which concentrates on the technical nature of the AI involved. This is for two reasons. First, we take the view that it is best to avoid the complex exercise in taxonomy which is inherent in the process of dividing applications into high risk, low risk and novel, and then treating each differently. Secondly, we are unattracted by features of the EU scheme such as presumptions of fault, which in our view would serve arbitrarily to differentiate AI liability from other liability without good reason. Our solution is, we suggest, less cumbersome in both respects.

Instead, we prefer to categorize our scheme by reference to the claimant's interest at stake, and as far as possible make this scheme harmonize with, and as far as possible operate parallel to, existing heads of tort liability. We take the view that four classes of potential claim are potentially relevant here. These are (i) personal injury and death; (ii) dignitary interests of natural persons; (iii) damage to property; and (iv) pure economic losses. As will appear, to some extent, these are already differentiated under the present law. This alignment is quite deliberate, in view of our opinion that AI liability should as far as possible be harmonized with the principles of liability in tort as a whole.

Liability for personal injury and death

AI already controls surprising numbers of matters relevant to personal safety, and this influence is likely to increase by leaps and bounds. Apart from vehicles (the subject of separate legislation which we do not cover in detail here),⁴⁴ one can cite its use in medical treatment and diagnosis; in access control and direction in dangerous environments such as mines or factories; in controlling aircraft, vessels and drones; in traffic control; in manufacturing; and in the design of large commercial and residential buildings.

We take the view that something more than mere negligence liability based on proof of fault is appropriate here. This is for several reasons. One is a general point about personal injury: for an individual claimant, the results are more likely to be catastrophic, and much less likely to be covered by insurance. In accordance with our scheme of values, the case for affording compensation is, therefore, strong.⁴⁵ The second point is that the risks arising from many of the uses of AI described above are big and serious, and claimants are likely to find it difficult, and increasingly so, to avoid running them. For example, if AI implanted in a human body (say to control a prosthetic limb or regular heart function, or to maintain a constant feed of medication) malfunctions, there may be consequences not only for the patient themselves, but for

⁴⁴ In the shape of the Automated and Electric Vehicles Act 2018 referred to above.

⁴⁵ This is what differentiates personal injury claims from those for compensation for property damage. The vast majority of the latter are brought by insurance companies, the owner of the property having already been compensated.

third parties injured in an ensuing accident. The third reason is that even though AI liability and product liability are not the same thing,⁴⁶ there is a strong parallel between the two. If we allow strict liability for personal injury in the case of the latter, which we do,⁴⁷ there is a strong case for doing so for the former as well, subject to many of the same conditions.⁴⁸

In the context of personal injury, therefore, we propose an area of AI strict liability essentially operating in parallel to product liability.⁴⁹ Indeed, to make the transition between the two smooth and prevent anomalies, it is suggested that the two heads of liability need to be harmonized as far as possible. Thus, in particular, at least some of the defences available in a claim under Part I of the Consumer Protection Act 1987 should be extended to an AI action too. The most obvious of these is the development risks defence under s.4(1)(e), exonerating a producer in respect of defects undiscoverable by the technology existing at the time of production. Although controversial and arguably inimical to the entire idea of strict liability,⁵⁰ so long as it remains in the case of product liability suits, there can be no justification for excluding it in AI claims. For similar reasons, other defences may also have to be carried over; for example, that in s.4(1)(a) (defects arising from compliance with legislative requirements).

There also should in our view be parallels to the defences in s.4(1)(b)(c) and (f), dealing with goods never supplied, or supplied only in a non-business context, or causing harm only because of the design of another product they are incorporated in. Although none can be directly applied to AI, a user or developer of AI should be able to escape liability by proving that it never made the AI available to anyone else, or did so only other than in the course of a business; or that the AI was harmful solely because of the design of some other product or AI in which it was integrated or with which it was associated.

Dignitary or reputational injuries

We are dealing here with interferences with a number of traditional legally protected interests of natural persons falling short of bodily integrity but nevertheless open to being vindicated under the law of tort: notably privacy, reputation and dignity.⁵¹ There are a number of ways these can be affected by the uses of AI. An AI-powered misdiagnosis of cancer or dementia, for example, could well affect dignity and well-being in a big way quite apart from any actual injury.⁵² So too with AI implanted in the body, as in the case of prosthetic limbs or software controlling the operation of the heart or other bodily functions. In addition, data mining and the manipulation of personal data by algorithm obviously have worrying implications for personal privacy; think of the case of personal information extracted and then misused, or stored by a data processor

⁴⁶ Because it seems fairly clear that AI, at least where not embodied in a physical object, is not a 'product': see Cabral (n 20). (The contrary argument in Gerhard Wagner, 'Robot Liability', (2018) <<https://ssrn.com/abstract=3198764>> (accessed 25 January 2023) is unconvincing).

⁴⁷ Under Part I of the Consumer Protection Act 1987, referred to above.

⁴⁸ It might be argued that we should go further and advocate compulsory liability insurance for all controllers of AI. Arguably compulsory insurance would also go a long way in affording third parties and users of AI applications an added degree of protection against the risk of insolvency. But it has not thus far been the practice to require companies to insure against product liability, or indeed public liability generally, and it is suggested that liability for AI is not sufficiently unique to justify such a requirement in that limited situation. Things are different for those operating in health sector as they might be required to obtain liability insurance (and in our view they should be required to do so). And in fact, one would expect an personal injury claim associated with an AI application to come from this sector so perhaps this is not a big problem as it might initially seem.

⁴⁹ It would be theoretically possible simply to do this by extending the Consumer Protection Act's definition of 'product' in s. 1(2)(c) to include AI, with a definition of the latter being provided. But that solution would be problematic when it came to concepts such as 'supply', and also because we see no reason to extend AI liability to cover damage to consumer property, to which at present the 1987 Act does, in our view anomalously, apply.

⁵⁰ See Duncan Fairgrieve and Richard S Goldberg, *Product Liability* (3rd edn, OUP 2020), 13-33–13-38.

⁵¹ Protected under the general law by a combination of torts: defamation, breach of confidence, misuse of private information, liability under the General Data Protection Regulation (Regulation (EU) 2016/679), and others.

⁵² Actual personal injury is of course also possible, for example, from unnecessary surgery induced by the false positive diagnosis. This is dealt within the previous section.

and then mistakenly publicized owing to malfunctions in the processor's computer, or even the malversations of external hackers.

Similarly, data controlled by AI can affect reputation. Falsehoods about individuals could be released to news outlets or a subject's personal acquaintances by AI contained in a company's press office or PR department; credit scores might be mishandled by financial institutions and end up in the wrong hands. Or for that matter, imagine the use by an employer of an algorithm that wrongly flags up an employee as guilty of misconduct or constant malingering, or an applicant for a job as having some undesirable characteristic.

Two questions arise in connection with such injuries. First, should there be a generalized liability at all where such injuries are caused by the actions of AI? Secondly, if the answer is Yes, at what level should that liability be set?

On the first question, the analogy of data protection law under both the EU's GDPR⁵³ and the UK Data Protection Act 2018 underlying it, where there is just such a generalized liability,⁵⁴ suggests that the answer should be Yes. The issues raised, and the interests threatened, by AI in this connection are similar to those raised by any other form of mass data storage, namely a novel and potentially potent source of harm that does not necessarily mesh neatly with the existing tort law. Although many cases of AI affecting reputation or dignity will engender some head of tort liability, not all will: there may, for example, be no liability in the case of failure to hire because of misinformation induced by an AI malfunction,⁵⁵ or possibly misdiagnosis leading to nothing more than distress.⁵⁶ Yet, consistently with our aim that deserving claimants should not be left in the cold, there is a case for ensuring that at least some compensation should be uniformly available for moral arms of this sort.

This brings us to the second issue: if there should be a generalized liability, should it be strict, or should it be in some way fault based? Although one might instinctively think the answer was the former (as did the European Parliament, at least in the case of high-risk AI⁵⁷), the case is actually rather weaker. If we take the analogy of the Data Protection Act 2018 and the GDPR, their liability depends not on the simple fact that information has reached the wrong hands (ie strict liability), but on fault, in the form of some kind of culpable breach of one of the many duties under the Regulation.⁵⁸ Furthermore, the imperative to compensate for physical personal injury applies with less force to distress or mere injury to dignity, as witnessed by the fact that in the general law, there is no overall duty to protect such interests.⁵⁹ It is therefore suggested that, consistently with the approach of the existing law, the best solution is to apply liability based on fault, though subject to the modifications suggested above ("AI and Fault") concerning decisions entrusted to AI applications.

Damage to property

In the nature of things, physical damage to property may well be less likely to occur than other kinds of injury as a result of AI problems. But it is perfectly possible. Apart from autonomous vehicles, subject to their own bespoke regime, consider the case where AI used in a construction project goes wrong, or where a factory process controlled by AI malfunctions, causing an explosion damaging nearby premises and vehicles.

⁵³ Regulation (EU) 2016/679 of 27 April 2016. Post-Brexit, an amended form of this remains in effect in the UK as retained EU law.

⁵⁴ See the detailed compensation provisions in ss. 168 and 169 of the Data Protection Act 2018.

⁵⁵ For example, where a computer failure causes a reference to become garbled. In such a case, the giver of the reference might well not be negligent so as to give rise to liability.

⁵⁶ This seems implicit in the decision in *Grieves v FT Everard & Sons Ltd* [2007] UKHL 39; [2008] 1 AC 281 that there is no liability if negligence which might cause disease causes not actual lesion but mere anxiety that disease might be there. The present case is a fortiori.

⁵⁷ See the EU Parliament report, referred to in Note [16] above.

⁵⁸ See the detailed compensation provisions in ss. 168 and 169 of the Data Protection Act 2018.

⁵⁹ The classic statement of this position appears in a dictum by Lord Denning MR in *Hinz v Berry* [1970] 2 QB 40, 42.

Traditionally, property damage is lumped together with personal injury, but this is arguably more from tradition than for any well-thought-out reason.⁶⁰ There are actually a number of reasons to consider a property damage claimant less deserving. For one thing, in contrast to the situation with bodily integrity, the immediate loss to the owner of damaged property is almost invariably covered by insurance.⁶¹ Furthermore, business losses consequential on property damage can also be if the owner so wishes. It follows that the immediate beneficiaries of any cause of action are likely to be not owners but insurers. (Theoretically, the availability of subrogation should cancel this factor out by reducing the premiums paid by owners seeking cover, but whether this actually happens is open to considerable doubt.)⁶² Third, a practice already recognizes at times that property damage might be different from bodily lesion. For example, the legislation based on the EU Product Liability Directive excludes from its protection all commercial property⁶³ (though it must be admitted that, somewhat illogically, it includes damage to domestic property subject to a franchise of €500).⁶⁴ Again, the scheme for state compensation for crimes of violence is limited to personal injuries,⁶⁵ with the making good of property damage claims left to the property insurer, if any.

As regards AI, in order to preserve horizontal equity between those suffering loss by reason of physical and AI causes, and also between those who do and do not use AI in their business, there should be a liability for fault on the lines suggested above. Moreover, this should include the extended definition of fault discussed above (“AI and Fault”). But it is suggested that it is difficult to see any justification for going further or extending the strict liability scheme suggested above so as to cover cases of this sort. Certainly, the imperative to compensate for personal injury—that the loss is likely to be both disastrous and uninsured—does not apply. On the contrary, most items of property, both domestic and commercial, are readily insurable and in the vast majority of cases actually insured. In these circumstances, it is suggested that for the moment at least, the best solution seems to be to say that losses of this sort are generally for property insurers to cover, and if they can to recoup by proving fault in a subrogation action. The authors believe that such an action is not only in line with the underlying values discussed earlier but also in accord with the established business and insurance practices.

Other losses

Traditionally, the attitude of the law of negligence towards compensating financial losses not resulting from physical injury has been highly restrictive. Although duties have been regularly imposed in a number of particular discrete situations, for example, negligent misrepresentations made in a business context,⁶⁶ or professional liabilities to clients,⁶⁷ there is no general duty of care in the absence of some positive reason to impose one, such as an undertaking of responsibility.⁶⁸ If there is no such reason, a claimant is expected to rely on what rights, if any, it has under the law of contract, whether against the defendant or a third party.

Should this be subject to review where damage is caused through the intervention of AI? Imagine, for example, that an AI-equipped government computer in a port malfunction and as

⁶⁰ The equiparation has, in fact, been questioned in academic circles, see Andrew Tettenborn, ‘Property Damage and Economic Loss: Should Claims by Property Owners Themselves Be Limited?’ (2005) 34 CLWR 128.

⁶¹ Admittedly subject to any excess: but the excess can often be adjusted by the assured according to how much they are willing to pay.

⁶² For a clear (though rather dated) suggestion that subrogation makes no practical difference, see Reuben Hasson, ‘Subrogation in Insurance Law - A Critical Evaluation’ (1985) 5 OJLS 416, 422–423. Those in insurance practice will readily testify that pursuing a subrogation action can be problematic for practical and evidential reasons in most cases.

⁶³ Council Directive 85/374/EEC, Art 9(b)(i), anglicized as the Consumer Protection Act 1987, s. 5(3).

⁶⁴ Council Directive 85/374/EEC, Art 9(b); Consumer Protection Act 1987, s. 5(4).

⁶⁵ See the Criminal Injuries Compensation Scheme 2012, paras 4–9 (subject to a very small exception for personal appliances such as walking sticks: para 52(a)).

⁶⁶ As in cases such as *Hedley Byrne & Co Ltd v Heller & Partners Ltd* [1964] AC 465 and *Smith v Eric S Bush* [1990] 1 AC 831.

⁶⁷ A point finally cemented by the House of Lords in *Henderson v Merrett Syndicates Ltd* [1995] 2 AC 145.

⁶⁸ See *Clerk & Lindsell* (n 27), para 7-103 onwards.

a result, the claimant is prohibited from landing goods which were in fact perfectly lawful; or AI in a farmer's computer misdetects an outbreak of some plague, causing movement restrictions to be put on the claimant's animals.⁶⁹ We suggest that no case for extending liability here is made out, since it is difficult to see any particular extra threats presented by AI in this context. We would, therefore, suggest that as regards liability for pure financial loss, there is no need for any provision to be made for AI, and that the existing law should be left as it is.

This is not to say that there would never be a liability in such cases. Where personal negligence was shown and the case fell within one of the existing areas where a duty is recognized, such as negligent misrepresentation, the fact that the representation happens to have been mediated through AI should make no difference to liability. Indeed, sometimes liability might in any case extend beyond its present limits. We suggested above that an act by a computer which if done by a human would have been negligent should engage the liability of the person in control of the computer, and this might well be significant. Imagine, for example, that without personal negligence on anyone's part, a stockbroker's malfunctioning algorithm caused a client to buy a wildly unsuitable security. Since a stockbroker undoubtedly owes a duty to his client in arranging the latter's purchases⁷⁰, and because the act of the algorithm would be assimilated to negligence and attributed to the broker, there would be a liability.

CONCLUSION

The AI applications are used in our daily lives and by large yield significant advantages for their users. However, just like anything else in life, things can go wrong and their users or third parties could suffer losses associated with the malfunctioning of the relevant AI application. The authors firmly believe that civil liability rules should not be used as an implicit regulatory proxy. No doubt, regulators will need to intervene (and this is currently debated in numerous jurisdictions) to impose manufacturing and operating standards with regard to AI applications to ensure the safety of these products and also to protect individuals. However, it should not be forgotten that the fundamental role of the law of obligations is to provide a satisfactory mechanism of compensation. The authors consider fundamental values that a tort regime should protect and various technical difficulties with regard to establishing liability for the use of AI that needs to be addressed, in this article advocate that there is a need for a new scheme of liability for AI which will fit neatly into the existing scheme of the law of obligations. Our suggestion is that any liability for harm done by AI should closely parallel other heads of liability, with changes made only where there are noticeable anomalies to be dealt with. It will have to be seen whether this scheme will stand the test of time, but it is perhaps better at this stage to advance slowly unless and until there is a noticeable need for more drastic reforms.

Also, such an approach is in line with the way these technologies are developed, that is, in an incremental fashion. Depending on the nature and speed of innovation, further changes in the new tort system might be needed but for now, we advocate the introduction of a strict liability system for personal injury and death claims emerging from the use of AI, a fault-based regime to be introduced for dignitary or reputational injuries caused by such applications. We do not think that there is a need to devise particular rules for damage to property or other losses emerging from the use of AI as we do not see how AI applications introduce added degrees of risks for their users or society as a whole that deserve additional protection afforded by a liability regime.

⁶⁹ Compare the analogue cases of *Weller & Co v Foot & Mouth Disease Research Institute* [1966] 1 QB 569 and *Pride & Partners (A Firm) v Institute for Animal Health* [2009] EWHC 685 (QB).

⁷⁰ *Clerk & Lindsell* (n 27), para 9-244.