



Chapitre d'actes

2019

Published version

Open Access

This is the published version of the publication, made available in accordance with the publisher's policy.

Liability for damages caused by AI in medicine : Progress Needed

Junod, Valérie

How to cite

JUNOD, Valérie. Liability for damages caused by AI in medicine : Progress Needed. In: Responsabilité civile et nouvelles technologies. Journée de la responsabilité civile 2018. Christine Chappuis / Bénédicte Winiger (Ed.). Genève. Genève : Schulthess éditions romandes, 2019. p. 119–149. (Collection genevoise. Droit de la responsabilité)

This publication URL: <https://archive-ouverte.unige.ch/unige:146875>

Liability for damages caused by AI in medicine : Progress Needed

Valérie Junod*

*Intelligence is the ability to avoid doing work,
yet getting the work done.*

Linus Torvalds

Table des matières

| | | |
|------|--|-----|
| I. | Introduction..... | 120 |
| II. | Liability related to the use of AI | 122 |
| | A. Liability of health care professionals..... | 122 |
| | B. Liability of the AI provider-producer-developer | 125 |
| | 1) Liability under the Product Liability Act (PLA)..... | 126 |
| | 2) Liability under Article 55 CO | 129 |
| | C. Liability of researchers..... | 131 |
| III. | Recommendations | 133 |
| | A. Recommendations to improve the current liability regime . | 134 |
| | B. Other recommendations..... | 136 |
| | C. Further considerations..... | 138 |
| | 1) Some cost issues | 138 |
| | 2) Some antitrust issues | 139 |
| | 3) Some ethical issues..... | 140 |
| IV. | Conclusion | 142 |
| V. | Reference list..... | 144 |

* Professor at the University of Geneva and and the University of Lausanne.

I. Introduction

Medicine is the field where Artificial Intelligence (AI¹) has been claimed to have the greatest immediate potential². Some of these claims are almost certainly "hype"³, while others appear more tangible⁴. Several scientific studies have been published which demonstrate that AI can do as well as or better than human experts for some specific medical tasks⁵. For example, AI can diagnose certain cancers better than expert physicians⁶; some AI tools can read chest X-rays better than most doctors; AI has even been shown to predict patients' suicides with a remarkable accuracy⁷! The uses of AI can extend well beyond diagnosis: AI can suggest the best cancer treatment from an array of possible medications; it can help synthesize the relevant medical literature for a given case⁸. The areas of medicine where AI can be exploited are constantly growing; some suggest that patients' medical history could be captured more reliably by an AI robot; medical research may also be accelerated thanks to AI⁹. However, scientific evidence of efficacy remains limited to specific areas (i.e., narrow AI); hope for a broad general AI has yet to materialize¹⁰. Hence, AI will not replace doctors altogether- at least not in the near future.

¹ "There is no universally agreed definition of AI; the term broadly refers to computing technologies that resemble processes associated with human intelligence, such as reasoning, learning and adaptation, sensory understanding and interaction". Nuffield Council on Bioethics (2018), p. 1. See also GUIHOT ET AL. (2017), pp. 393-396; Académie des technologies (2018), pp. 15-20.

² TERRA NOVA, p. 5; BAMBAUER (2017), p. 383; GREENBERG (2017) ("AI, along with personalized big data, is projected to be the next disruptive technology that will transform medicine").

³ On the history of "hype" in the field of AI, see European Parliament Research Service (May 2018), p. 10; also SCHWARTZ ET AL. (1987), pp. 685-686.

⁴ On the development of IBM's Watson AI, see ALLAIN (2013); CHUNG (2017), p. 37; CHUNG/ZINK (2017); DELOITTE (2015). Explaining recent setbacks: DAVID H. FREEDMAN, *A Reality Check for IBM's AI Ambitions*, Business Report, Technology Review, 27 June 2017.

⁵ See LOH (2018); WISE (2018); CHOI ET AL. (2016); DIPROSE/BUIST (2016); GOLDEN (2017), p. 2184.

⁶ A correct diagnosis is of course essential, as this is the first necessary step for the subsequent correct treatment. Diagnostic errors (either a missed health condition or a wrongly identified health condition) can have dramatic consequences, leading to early deaths or decreased quality of life. According to GREENBERG (2017), "misdiagnoses are the leading cause of malpractice claims in both Canada and the United States". More generally, "the true number of premature deaths associated with preventable harm to patients [at hospitals in the United States] was estimated at more than 400,000 per year. Serious harm seems to be 10- to 20-fold more common than lethal harm." JAMES J. T., *A new, evidence-based estimate of patient harms associated with hospital care*, 9(3) Journal of Patient Safety. pp. 122-128, September 2013.

⁷ See GREENBERG (2017).

⁸ "Rien que dans la base de données PubMed, qui regroupe les articles de médecine et de biologie, environ 3'000 nouveaux articles sont indexés chaque jour. Un expert, s'il réussissait à lire trois articles par jour, ne lirait qu'un millième de ce qui est publié. Plus aucun humain ne peut donc lire, comprendre, inférer et résumer tout cela sans le secours d'un outil puissant." TERRA NOVA (2017), p. 7.

⁹ See GIL ET AL. (2014), p. 171.

¹⁰ On singularity as the ultimate step, see GUIHOT ET AL. (2017), pp. 397-399.

This is not to imply that AI will not become increasingly common in hospitals, clinics and perhaps even private medical practices. The economic interest in proposing AI products or AI services (hereafter : AI tools¹¹) is certainly present. Major companies are investing in this area and field testing their innovative solutions. The high cost of healthcare¹² and the high revenues generated by health care services certainly provide the necessary incentives.

Whereas Switzerland welcomes several established companies and many startups specializing in AI¹³, it has yet to introduce AI routinely into its medical establishments. For this reason, the present article on the liability resulting from the use of AI anticipates future and somewhat uncertain events¹⁴. Caution is “de rigueur”¹⁵. We will limit our analysis to what may occur now or in the next few years¹⁶. Fear of an all-powerful AI ruling over humans is best left to science fiction novels¹⁷.

The article is divided in three chapters. The first answers the question of the extent to which a health care professional, a health care institution or an AI producer may be found liable when a patient is injured following the use of an AI tool. The second chapter puts forward recommendations to improve the current liability system and analyzes the broader implications of AI use. The conclusion highlights the need for harmonization.

Beyond the scope of the paper are the damages that professionals themselves (e.g., hospitals, physicians) may incur because of AI tools ; also beyond its scope is deliberate hostile or criminal use of AI, regardless of whether it harms patients or institutions (e.g., AI being hacked to access patient data illegally)¹⁸. Finally, the paper only focuses on AI use in the context of patients' medical care, although one can imagine that AI could also be deployed more broadly in health care settings, for example in preparing hospital invoices or supervising medical staff.

¹¹ Of course, choosing the term "tool" is not entirely innocent, as it implies someone else, i.e., a person, using the tool.

¹² In Switzerland, health care expenses reached CHF 80.7 billion in 2016, a 3.8% increase over 2015. Overall, these expenses represent 12.2% of GDP. See Federal Office for Statistics, Press release of 19 April 2018.

¹³ ANOUCH SEYDTAGHIA, *La Suisse, leader en intelligence artificielle*, *Le Temps* du 31 janvier 2018 ; from the same author, *En Suisse, l'intelligence artificielle est devenue omniprésente*, *Le Temps* du 15 novembre 2017.

¹⁴ As would be expected, no lawsuit because of AI-caused damages has yet taken place in Switzerland – nor could I find reports of such judgments abroad.

¹⁵ See, e.g., TOM MORISSE (2017).

¹⁶ I leave aside the proposal advanced by some authors (e.g., CHUNG (2017) ; CHUNG/ZINK (2017)) that AI will become and behave so "human-like" that it should be granted "legal personhood". I am not convinced by the usefulness of such a proposal – at least for the mid-term future. For a more thorough discussion of the issue, see PETRELUZZI/MILIOTIS (2017) ; see already in 1981, LEHMAN-WILZIG, pp. 452-453.

¹⁷ See EUROPEAN PARLIAMENT RESEARCH SERVICE (May 2018).

¹⁸ On these issues, see, e.g., TSCHIDER (2018).

II. Liability related to the use of AI

A. Liability of health care professionals

Imagine the following hypotheses: a software using AI reaches the “conclusion” that patient X suffers from a certain cancer, although this actually not the case; an AI tool concludes that patient Y does not suffer from cancer, when in fact she does; an AI program recommends a certain course of drug treatment for patient Z, whereas another treatment would have cured the disease. In all three hypotheses, the patient incurs harm: added medical costs, loss of wages, possibly permanent body lesions or even death. Had the AI tool not been used, the outcome for the patient would have been different and, in all likelihood, better.

Should the physician handling the treatment of the patient in a private medical practice be held liable? Similarly, should the private clinic taking care of the patient face liability?

Under current Swiss law, a service provider – whether a doctor or a medical institution¹⁹ – is liable if the injured party can prove that the health care contract binding the two was breached²⁰; the claimant must also prove having incurred a damage, whether financial or moral (pain and suffering)²¹; this damage must have been caused by the above-mentioned breach of contract. Finally, depending on the fact pattern, a fault is presumed (the defendant can escape liability if she proves not to have been at fault²²) or not required (in case of vicarious liability of a private institution for the acts of its employees²³ or in case of an action against a State actor, e.g., most cantonal public hospitals)²⁴.

The first difficulty for the injured patient is to prove that the contract was indeed breached. When health care is delivered, the contract is typically one of mandate. A contract of mandate obliges the service provider to show the greatest diligence when delivering her services²⁵. Hence, the contract is breached if a reasonable health care provider, placed in the same

¹⁹ I will leave aside the liability of public hospitals, although ultimately it is not very different from the liability of private clinics.

²⁰ Articles 398, 97 or 101 of the Swiss Code of obligations (CO).

²¹ Articles 42, 47 and 49 CO; BAUME (2015), pp. 60-65, noting that they are no punitive damages under Swiss law.

²² Article 97 CO.

²³ Article 101 al. 1 CO – the exculpatory evidence to be provided by the defendant is nearly impossible to bring.

²⁴ On the legal regime of liability for public hospitals, see, e.g., Baume (2015), pp. 30-48.

²⁵ In the context of medical liability, see BAUME (2015), pp. 58-59.

circumstances, would have ordinarily exhibited a greater degree of diligence²⁶. In our hypothesis, it boils down to whether a diligent health care provider would have used the AI in the same manner. The questions to be answered are for example whether a reasonable provider would have used the AI tool at all, whether it would have cross-checked the AI's recommendations with other sources of information (e.g., a human medical expert), whether it would also have relied "blindly" on the AI recommendations, and whether it would have followed the implementation of the AI recommendations more closely or more regularly.

As this list of questions reveals, this requires a highly fact-intensive inquiry. For certain medical services, for example certain diagnostics, there may be scientific evidence that the AI performs so much better than human doctors that the diligent course of action is to rely on the AI recommendations, without confronting them with the perspective of human doctors²⁷. This would be the case if, for example, a relevant peer-reviewed study had established that the AI diagnostic tool reaches 99% sensitivity (true positive rate) and 99% specificity (true negative rate), while human experts in the field exhibit much lower scores (let's say 40%). In that case, even if the unfortunate patient falls in the 1% error rate due to AI, the actions of the health care provider are still diligent and no breach of contract can be shown.

The interesting questions may then turn to the cutoff point where "blind" reliance on AI becomes admissible. If AI has only 70% sensitivity and 70% sensibility, whereas human experts are at 60%, should the two combine their efforts? These questions are best answered by medical experts, usually in the course of a judicial expertise and/or by medical societies issuing recommendations to their members (e.g., "when doing a chest X-ray for breast cancer, we advise checking the recommendations of the AI against the reasoned opinions of two medical experts").

A problem that we are likely to encounter is how to quantify the performance of the AI tool. This problem stems from the fact that AI is designed to learn constantly by incorporating new data. Of course, it may be tested in a clinical setting when it is launched, but afterwards it keeps on "tweaking" its abilities. Thus, the AI may achieve 90% accuracy as compared to experts at time t_1 , but no one knows how it scores at time t_2 when the patient was injured. It is hoped that its performance will keep on improving,

²⁶ See WERRO (2012), p. 2397 et ss.

²⁷ See FROOMKIN ET AL., (2018). An additional issue is which kind of doctors should constitute the relevant baseline. Should it be primary care doctors? Doctors specialized in this particular field? Experts in the field? If, in a rural area, the primary doctor is also offering basic dermatologist services to the population, should her diligence when using an AI tool be compared to that of other rural doctors not using AI? Or to public hospital experts even if these experts are not readily available to her patients?

but there will be little or no scientific evidence for time t_2 ²⁸. Even worse, it will be very difficult and perhaps even impossible to acquire - after the facts - information about its t_2 performance²⁹. Similarly, it will be very difficult or impossible to know how the AI reached its recommendation in the case of the individual (injured) patient (the "black box" effect³⁰). Finally, conclusion reached for one AI tool (i.e., this one is defective) will not be transposable to a second AI tool, because the two, even when designed together, will have - over time - learned and evolved differently³¹. As a result, the legal decision regarding a possible lack of diligence will have to be taken based on possibly outdated and certainly incomplete information. This is not entirely unusual, as this also occurs in other medical malpractice cases, but it is an uncomfortable situation. All the more so since patients bear the burden of proof and failure to prove a breach of contract will lead to denial of indemnification.

Assuming that this first obstacle is overcome, the patient still has to prove causality³². This requirement is known to be difficult to satisfy in medical cases³³. It is hard to determine what would have happened to the patient, had a different medical course of action been decided upon. Would her life have been saved if her cancer had been diagnosed sooner? Would she have lived longer thanks to a more carefully chosen treatment? How much longer? These questions are even more difficult to answer if the patient is already very

²⁸ On the additional risk of data decay, see CHEN / ASH (2017), p. 2507.

²⁹ "[I]n the field of robotics, it could be difficult to distinguish between a defect which was present at the moment in which the robot was put into circulation from a defect that appeared after that moment, given the characteristics of self-learning systems." ERNST & YOUNG (2018), p. 85.

³⁰ Black box refers to the fact that the AI tool is designed in such a manner that it cannot explain which factors it used to reach its conclusions, nor can its human users. See WISE (2018). This concept is sometimes opposed to (the need for) "explainability" or "transparency". See U.K. SCIENCE AND TECHNOLOGY COMMITTEE (2016), p. 17. As the EUROPEAN COMMISSION (April 2018) explains; "in order to increase transparency and minimize the risk of bias or error, AI systems should be developed in a manner which allows humans to understand (the basis of) their actions". See also EUROPEAN PARLIAMENT (2017), point 12 ("it should always be possible to supply the rationale behind any decision taken with the aid of AI that can have a substantive impact on one or more persons' lives"); FROMKIN ET AL., (2018), p. 60 ;TERRA NOVA (2017), p. 22 ("l'affirmation d'un principe général de transparence des algorithmes risque d'être vaine ; elle serait d'ailleurs susceptible de nuire au respect du secret industriel et pourrait potentiellement freiner l'innovation."). See also PACKING/LEV-ARETZ.

³¹ The judgment of the EU Court of Justice in C-503/13 of 5 March 2015.

³² See BAUME (2015), pp. 65-66. In addition, because each AI tool is likely to be viewed as an individual and different product, since each will have "learned" differently, it will not be possible to infer that the present AI suffers from a defect, because another AI has already been proved to be defective. This adds another practical difficulty (compare with European Court Judgment of 5 March 2015, C-503/13, in case of classic products).

³³ When the contract breach resides in an omission – as opposed to an action – then causality must be determined hypothetically. According to the Federal Tribunal's case law, a strict level of proof (evidence) is not required. A preponderance of evidence suffices. The plaintiff must still demonstrate that if the proper medical action had been undertaken in her case, she would have been more likely than not to have avoided the damage. See, e.g., Judgment of July 4, 2006 in case 4C.125/2006. See also on the burden of proof, VIGNERON (2018), pp. 45-46.

sick and when the disease is known to have a high fatality rate³⁴. The problem is compounded by the fact that, once AI becomes broadly used, there will be few other benchmarks to make a comparison³⁵. If most medical practices use the same AI or even if they use different AI tools, how is the judge to decide what would have happened if no AI had been used. Moreover, the fact that AI tools will be put on the market based on (some) clinical evidence of their efficacy will make it hard to say that an alternative course would have led to a better outcome. If the AI has been shown to be, say, 10% better than experts, how could the judge rule that *not* using the AI in the specific situation of the injured patient would have avoided part or all of the harm? The availability of pre-launch clinical evidence for safety and efficacy of the tool is likely to tip the balance (regarding the causal pathway) in favor of AI use³⁶.

Finally, courts will find it increasingly difficult to understand the technical evidence presented to them³⁷. Expertise reports that they will have to assess will need to be prepared by interdisciplinary teams made of software and hardware engineers, doctors, biostatisticians.

B. Liability of the AI provider-producer-developer

While the exact business models of AI suppliers are yet to be precisely defined, it can safely be assumed that AI will be developed by companies which are not primarily health care providers. The AI tool will most likely be licensed as a program; it might also be incorporated in a hardware, which will be sold, leased or licensed to health care providers.

When thinking of the liability of companies providing medical goods, it is obviously the Federal Product Liability Act (PLA; SR 221.112.944) that comes to mind. Although this Act has not been frequently applied by Swiss courts³⁸,

³⁴ For cancers for which patients' life expectancy after diagnosis is low, it is difficult to prove that an earlier diagnosis would have led to an overall cost saving. It is only for diseases which respond well to early treatment that such proof is available, at least based on statistical data. Moreover, since the Federal Tribunal has rejected the loss of chance theory ("théorie de la perte d'une chance"), the patient is faced with an all-or-nothing dilemma: she succeeds in proving causality and obtains indemnification for the entire damage – or she fails and obtains nothing. See VIGNERON (2018), pp. 46-51; BAUME (2015), pp. 63-64.

³⁵ FROOMKIN ET AL., (2018) ("If a set of symptoms is consistently producing an erroneous ML [machine learning] diagnostic, and physicians act on that erroneous diagnostic, where will ML get the data to suggest a different diagnosis which lead to better treatment? If the answer is "nowhere" then we have a problem. Worse, it is not even clear that either the ML system or an outside observer necessarily would know that the results were sub-optimal.").

³⁶ See in a slightly different context, EUROPEAN PARLIAMENTARY RESEARCH SERVICE (February 2018), p. 33.

³⁷ "[E]ven experts may find it difficult to understand defects in new technological developments, because of the increasing number of chain events where possible causal connections are not understandable." ERNST & YOUNG (2018), p. 85.

³⁸ In Switzerland, only five cases decided by the Federal Tribunal applied directly the PLA (4A_365/2014 of 5 January 2015; 4A_16/2011 of 18 March 2011; 4A_319/2010 of 4 October

it is meant as a pro-consumer tool³⁹ to allocate liability for defective products. Additional objectives include encouraging safe product design⁴⁰ and ensuring the availability of sufficient information about proper use. Together these goals operate to guarantee people's trust in the products that they buy and use. Such trust is particularly important in the field of medicine.

The first legal difficulty that we encounter is that the PLA only applies to goods, and not to services⁴¹. When software is incorporated into some hardware, it qualifies as a good under the PLA. However, pure software, e.g., a website where data can be entered by the user to receive the result of computer calculations, is not deemed a good⁴². This good/service distinction has already been criticized as outdated⁴³. It does not make much sense to apply a completely different regime depending on whether the software is available on a USB device or purely on-line, when the key functions of the software remain the same. In the context of AI, several authors have already recommended considering all AI tools as goods under product liability regimes⁴⁴.

In the first part of this section, we will assume that AI qualifies as a "good" under the Swiss PLA. Then, we'll study the opposite hypothesis.

1) *Liability under the Product Liability Act (PLA)*

Under the PLA, a patient can be compensated for the economic and moral harm incurred if she proves that this harm was caused by a defective good put into circulation by a producer⁴⁵.

The producer is broadly defined⁴⁶. The company developing and selling/licensing the AI is to be viewed as a producer. The importer into Switzerland and (to some extent) even the final supplier are also producers⁴⁷. Depending on the circumstances, the health care institution using the AI could

2010 ; 4A_255/2010 of 29 June 2010 ; 4C.298/2006 of 19 December 2006). In the European Union, the equivalent Directive 85/374 has not been very much used either. "[M]ost product liability claims between 2000 and 2016 [a total of 798] were actually settled out of Court [46%]". EUROPEAN COMMISSION (May 2018), p. 4.

³⁹ Worth mentioning in passing is the fact that provisions of the PLA are mandatory and cannot be contracted away (Article 8 PLA). This can be a notable advantage compared to contractual or tort liability, where exonerative clauses are admissible within the limits of Articles 100 and 101 CO.

⁴⁰ See, e.g., ALLAIN (2013), p. 1067.

⁴¹ Articles 1 and 3 PLA.

⁴² See VOKINGER ET AL., (2017), p. 17.

⁴³ See, e.g. ERNST & YOUNG (2018), p. 71 ; ALLAIN (2013), p. 1070.

⁴⁴ See VOKINGER ET AL., (2017), p. 18.

⁴⁵ Article 1 PLA, with Article 2 defining the producer, Article 3 defining the product, Article 4 defining the defect.

⁴⁶ See MARCHAND (2018), pp. 24-25.

⁴⁷ Id. pp. 25-26.

also qualify as a producer, particularly if it improves on the program by feeding it data (see below)⁴⁸. However, should AI come to be developed in an open-source environment, identifying who added which lines of code might prove increasingly difficult. Thus, identifying all AI developers, or at least the most "guilty" coders among them, will be a challenge. A possible solution would be for the law to pinpoint (at least) one available "defendant"; thus, the plaintiff and injured party would always know who, in Switzerland, could and should be sued. This is only partly the case today since the local supplier/seller can only be sued under the PLA if it failed to indicate upon request the name of the producer or the importer⁴⁹.

The main issue for the injured patient is whether she can prove the AI to be defective. Defect is defined in an objective manner as the lack of the level of security which reasonable users can expect (of such a product)⁵⁰. One must first ascertain who the intended user is. Because the AI will generally be used by doctors, and not by patients directly⁵¹, the relevant user is the physician (or team of physicians, depending on the situation); the security expectations of the patient are not relevant⁵². As is the case for medicines, the expectations of the physician will be shaped by the information she is given. In the Federal Tribunal's judgment regarding Bayer's Yasmin contraceptive pill⁵³, the existence of a defect was negated because the professional notice of use clearly stated the medical risk which materialized in the case of the patient⁵⁴. Since physicians were explicitly warned of the (largely unavoidable) risk, they could not claim that the pill was defective. If this legal doctrine were to be transposed to AI, it would mean that the existence of a defect would be dismissed whenever users were clearly warned of the corresponding risk – at least so when the mentioned risks are unavoidable. For example, if the company offering the AI warns that the incidence of a false positive is 5% (the patient is told she has cancer, but has none) and the incidence of a false negative is 3% (the patient is told not she does not suffer from cancer, but does), then this should negate the defect if the patient is unlucky to fall within these percentages. A defect could still be admitted if the actual incidence of mistakes turns out to be 10% and 6%. However, in practice, clinical study

⁴⁸ In normal situations, simple users of products are not producers. See Judgment of 21 December 2011, C- 495/10.

⁴⁹ Article 2 paragraphs 2 and 3 PLA.

⁵⁰ Article 4 PLA. See MARCHAND (2018), p. 27 ("le critère est celui de la sécurité à laquelle le consommateur peut légitimement s'attendre").

⁵¹ In a more distant future, one can imagine each person having her or his own AI machines at home to provide medical services. But today, this seems more like the stuff of movies (e.g., the 2013 film *Elysium*) than impeding reality.

⁵² On the "learned intermediary doctrine" in the United States, see ALLAIN (2013), p. 1069.

⁵³ Federal Tribunal's judgment 4A_365/2014 = 4A_371/2015 of 5 January 2015.

⁵⁴ See also WERRO's comment of the Federal Tribunal judgment, WERRO (2018), pp. 95-102. See also EUROPEAN COMMISSION Staff Working Document (May 2018), p. 25.

results which rectify numbers initially provided rarely become available after launch. Therefore, the physician may suspect that the AI is not quite as reliable as claimed, but she may find it impossible to prove this to the required evidentiary standard⁵⁵. Moreover, relying on (human) expert evidence will be very difficult too; as mentioned above, the injured party may not find an expert that can credibly attest that the AI failed to work as claimed. Indeed, the skills to qualify as such an AI expert will be multiple and highly sophisticated. Identifying and retaining these experts for the purpose of court proceedings will not be easy.

More generally, AI producers will be inclined to include broad warnings. They may (truthfully) highlight that AI tools keep on learning (based on the data they are fed) and that, while it is hoped that this learning will constantly improve performance, this cannot be guaranteed. The AI producers may call for users to exercise caution and not to rely blindly on the AI results. Such warnings are likely to increase the liability of health care providers and to lower that of AI producers. Of course, Article 8 PLA prohibits liability waivers, but inevitable genuine risks inherently resulting from use of a good should not fall within the scope of this prohibition.

We have shown that injured patients may find it difficult to prove that the AI which was used in their medical care is defective⁵⁶. The next legal obstacle in their path toward indemnification is Article 5 paragraph 2 letter e of the PLA. Producers can escape liability (once the conditions of good, defect, damage⁵⁷, causality have been proved by the claimant) if they can establish that, at the time the AI was put into circulation, they were not aware and could not have been aware of the risk which materialized⁵⁸. This escape route has been labeled "the development risk"⁵⁹. Switzerland makes this broadly available to all producers, except for those offering transplants⁶⁰. In the case of AI, a producer could convincingly explain that it tested its AI before release, that it has clinical data to prove that the AI performed better than physicians at

⁵⁵ Generally on "consumers' difficulties in obtaining compensation in the pharmaceutical/medical sector", ERNST & YOUNG (2018), p. 28.

⁵⁶ According EUROPEAN COMMISSION (2018), "based on the desk research at national level, the most frequent reason for rejecting a claim for defective product [in any field, not specifically AI or medicine] is related to the burden of proof, and specifically to : (i) prove the defect (32% of cases) ; and (ii) prove the link between the defect and the damage (21% of cases)." (p. 27).

⁵⁷ As per Article 6 al. 1 PLA, only damages to things exceeding CHF 900 is covered by the Act. This limit has been criticized as unduly discouraging small claims, which – although small by themselves – would add up to high overall damages.

⁵⁸ On the scope of this exception, see, e.g., MARCHAND (2018), p. 19.

⁵⁹ For an analysis and a general background on this exception, see TRAN/ETIER (2018), pp. 106-128.

⁶⁰ In the European Union, Member States are free to retain or to reject this exception to liability. A minority of countries has chosen to make producers liable also for development risks. Some countries (Finland, Luxembourg, France) have a regime whereby this exculpatory factor is not available when the defect is due to pharmaceuticals or to products derived from the human body. See EUROPEAN COMMISSION (May 2018), p. 4. Also PICHONNAZ/WERRO (2018), p. 21.

the time of release, and that it has no reason to suspect that the way the AI was programmed could lead to a defect in the course of its use. The situation of the AI producer is similar to that of a pharmaceutical producer. A company like Bayer can persuasively argue that it tested its contraceptive pill carefully and that, to the best of its knowledge, the clinical trials it had to conduct gathered all possibly identifiable risks and that it could not have known of a risk that was not flagged during pre-launch clinical trials. In the pharmaceutical context, the development risk clause results in pharmaceutical companies rarely being found liable for adverse side effects not revealed in the Phase I, II and III clinical trials (i.e., the studies taking place before marketing authorization).

For AI, one added difficulty, that we already alluded to, is that the product is meant to evolve⁶¹. An AI used for diagnostic purposes is supposed to gain more accuracy the more diagnoses it performs ; an AI tasked with identifying the proper drug treatment is expected to improve with each increment of patient data it analyzes. This means that the tool continually changes as it is used. Data stemming from the original pre-launch clinical testing may no longer be relevant one year after the AI tool has first been put into circulation. However, this learning will probably not occur under the control of the AI producer-developer, but under the control of the health care provider (e.g., a university hospital). Can the initial developer still incur liability if further learning by the AI occurred under the control of third parties ? Should this situation be compared to that of a car maker which sells a car which is later modified (“tuned”) by its new owner ? Or should the situation be distinguished because the AI is inherently meant to be “fine-tuned” by its new owner, so that this further third-party learning should be expected and even anticipated by the product developer ? In practice, the answer may depend also on the concrete terms of the licensing agreement between the AI developer and the AI user (i.e., the hospital). Because the legal issue is undeniably thorny, it should be dealt with by the legislature, rather than by case-by-case interpretations.

2) *Liability under Article 55 CO*

Let us now turn to the second hypothesis : the software is not incorporated into any leased-licensed-sold hardware, but is, for example, a free-standing

⁶¹ Developers could also argue that they should no longer be liable for defects that arise after the release of the product. Indeed, under Articles 4 and 5 para 1 letter b PLA, the defect must be present at the time of release (“mise en circulation”) ; it is for the producer to establish that the defect must have arisen later. In my view, such exculpatory evidence should not be accepted in the case of a product which is designed specifically to evolve. The developer is selling a tool under the business proposition that this tool should get better and better at its tasks. In such circumstances, the developer should be held to its promises. It should be considered that the AI tool was designed so that, from the start, the origin of the later-materialized defect was present.

internet decision-aide, in other words more akin to a service. If the PLA is held not to apply, then the injured party must rely on the Code of obligations (CO ; SR 220) and its provisions on tort liability. Traditionally, before the entry into force of the PLA in 1993, Article 55 CO was the topical provision, although Article 41 CO also enters into consideration, since the patient is not bound by contract to the AI developer. Under Article 55, the patient must show an injury (economic loss or moral harm), an illicit act within the sphere of control of the defendant, and a causal link ; proof of fault is not required, but the defendant can escape liability if it can bring certain exculpatory evidence. The first step is to assess whether there is an illicit act or an illicit omission. Bodily harm is illicit because it infringes on the absolute right to personal integrity. An incorrect AI-led diagnostic or an ill-advised AI-recommended treatment may very well directly cause bodily harm, because the patient had to undergo treatment for nothing, because she did not undergo treatment that could have saved or helped her, or because a better chosen treatment would have been more effective. Once or if it can be shown that the AI was “wrong”, it then becomes clear that this mistake was the causal factor in the adverse outcome. We will assume here that the mistake was not so obvious that any reasonable physician should have noticed and corrected it. Because the patient has satisfied her burden of proof under Article 55 CO, the question turns to whether the AI producer can bring forth the exculpatory evidence allowed under this provision. If the AI was properly developed and tested before release, the producer can show that its employees and auxiliaries were well chosen, well instructed and well supervised. Since the AI was clinically tested before being put into the market, the producer can also show that it implemented a proper quality control. There is no reason to doubt that there should be a faulty system organization within the company producing the AI. Once more, the situation can be compared to that of a pharmaceutical company which releases a new drug or a new device that has been tested according to current regulations and guidelines before marketing. In conclusion, except in cases when the development work on the AI was faulty (e.g., hypothetical testing guidelines not followed), the AI developer will be able to provide the needed evidence to escape liability.

A last question under this section is whether the findings reached for the AI developer should be extended to the medical facility using the AI. Indeed, AI are likely to be released within hospitals, where they will continue to be trained chiefly or partly by the hospital staff ; AI will learn from patient data that the hospital will feed them. One can suppose that after a few years of use, an AI tool will function differently from when it was first released. Because the operation of the AI will have been altered by the hospital (its staff), should the hospital qualify as a producer ? We know that a pure user of a product is not a producer under the PLA ; for example, under these rules, a hospital using a MRI machine is not considered a producer. This changes, however, once the

hospital starts modifying a piece of equipment to further its needs. When data are fed into an AI program, the internal functioning of the AI changes. This is substantively different from using a software to fill cells in a form ; in case of an AI tool, with enough new data, the program itself is changed. For example, if erroneous data are systematically fed into the AI (e.g., by a dishonest employee), then the AI will start to give wrong recommendations or incorrect diagnoses even when, later on, correct information is being entered. In our view, the hospital faces a producer liability if it has the ability to change the functioning of the AI and has indeed done so⁶².

Whether the hospital-producer will actually be held liable for damages is more difficult to ascertain. Unlike the initial developer, the hospital will not have tested the AI while or after the hospital-led changes occurred to the program. On the one hand, this makes it harder for the hospital to explain that it did all it could to avoid the occurrence of defects. On the other hand, the hospital will have no reason to suspect that providing additional data to the AI could lead to defects, because AI is precisely meant to function this way, i.e., to improve based on additional data. Of course, if a corrupt hospital employee intentionally caused the AI to malfunction, the situation would be different and the question would then turn to whether the hospital properly chose, instructed and supervised its work force.

To sum up, we have shown that the action of an injured patient will face significant hurdles and is most likely to fail, even if she can show that the AI tool was in some sense "wrong". We should add that even showing that the AI was wrong may become increasingly difficult if the use of AI becomes more and more common. For example, if all AI tools recommend treatment 1 in case of cancer, and all hospitals implement this proposal, how will the patient be able to show that treatment 2 would have been better ? After a while, there may no longer be any established point of comparison. It will just be assumed that treatment 1 works well, since the alternatives are no longer being field-tested.

C. Liability of researchers

Before AI tools are released on the market, they will have to be tested. The purpose of this testing will be to ensure their safety and efficacy, the two notions being here closely intertwined⁶³. For instance, the AI developer will need to show that its AI diagnostic tool can diagnose conditions better than

⁶² "[T]he longer a robot's 'education' has lasted, the greater the responsibility of its 'teacher' should be". DIRECTORATE-GENERAL FOR INTERNAL POLICIES (2016), p. 17.

⁶³ Expressed in greater details by SHORTLIFFE/SEPULVEDA (2018), "[a] CDSS [clinical decision support systems] should have rigorous, peer-reviewed scientific evidence establishing its safety, validity, reproducibility, usability and reliability".

most doctors can, or that its treatment selection tool offers better recommendations. The data produced during medical testing will be submitted – along with proper analysis – to the authorities to decide whether the product can be put in circulation. We argue that AI tools are medical devices⁶⁴ and ought to be classified – at least at the beginning – as high-risk devices, making it necessary to obtain a State authorization before marketing⁶⁵. As explained by Tschider, the typology of risks that need to be examined for AI devices is much broader than for standard medical devices⁶⁶. Thus, in Switzerland, Swissmedic should be declared competent to decide whether the clinical data in the submitted file are enough to allow launch⁶⁷ and whether additional clinical trials after launch are necessary.

During the phase of clinical testing, the AI sponsor will have to set up insurance coverage to indemnify test subjects having incurred damages. According to the Therapeutic Products Act (TPA), clinical trials of uncertified medical devices can only take place if the liability of the sponsor is properly insured⁶⁸. Typically, the sponsor enters into an insurance contract with an insurance company. The insurance policy must cover all damages caused by the clinical trial irrespective of fault. In theory, neither the insurance company nor the sponsor can invoke any exculpatory evidence to avoid liability⁶⁹. Contrary to the Product Liability Act, ignorance of the defect cannot be claimed to escape sponsor's liability. Therefore, at least in principle, test subjects (during the clinical trials) should benefit from a broader liability coverage than patients (once the product has been authorized on the market).

⁶⁴ See the Federal Administrative Tribunal judgment of September 17, 2018 in the case C-669/2016. In that case, the Tribunal found that a basic diagnostic app was indeed a medical device if it was designed, based on the user's individual data, to give back medical information (here when the user was most fertile or not fertile so as to either favor or avoid a pregnancy). See also VOKINGER ET AL., (2017), p. 7 ; TSCHIDER (2018), part II ; ALLAIN (2013), pp. 1071-1072 ; BAMBAUER (2017), p. 388. For the more complex situation in the United States, see Food and Drug Administration (FDA), Software as a Medical Device (SAMD) : Clinical Evaluation, Guidance for Industry and FDA Staff (December 8, 2017). See also BARAK RICHMAN, *Health Regulation for the Digital Age – Correcting the Mismatch*, 379:18 *New England Journal of Medicine* 2018 pp.1694-1695.

⁶⁵ This is not how the classification works currently. Yet, given the number of patients who may be affected by the use of AI, we recommend implementation of higher safeguards, at least until it is established that the risks are not markedly higher. Of a more nuanced opinion, see VOKINGER ET AL. (2017).

⁶⁶ "New technology, however, introduces patient safety risks connected to backend systems and infrastructure, rather than devices interacting with the physical body. These new types of risks either necessitate a new formula for panel or expertise selection or require complementary reviews from experts who understand new technology functionality." TSCHIDER (2018), p. 200.

⁶⁷ The question of how long an AI tool should spend in clinical trial(s) could prove delicate. It seems that the more time the AI spends in a controlled learning environment, the better it performs. How will scientists then decide that the AI tool has reached its optimal learning point ?

⁶⁸ Article 20 of the Federal Act of 30 September 2011 on research involving human beings ("HRA" ; RS 810.30).

⁶⁹ See however Articles 10.1.c and 10.2 of the Ordinance on Clinical Trials in Human Research (ClinO ; SR 810.35).

This begs the question of whether it makes sense to distinguish between a product under testing and a product released for circulation in the case of AI. Indeed, one key characteristic of AI is that it keeps on learning, constantly incorporating the new data that it is provided with. Given that the AI tool is constantly evolving, how can one decide when it is leaving the testing phase? It could be argued that the constant fine-tuning of AI tools means that they remain in a testing phase for a long time, and possibly forever.

It might therefore be necessary to introduce a rule to decide when the AI is used in the context of research – with the ensuing application of the protective regime of the TPA – and when it has reached a sufficiently solid level of development where the TPA is no longer applicable and the PLA applies instead. However, if our proposed recommendation to transform the current liability regime into a true “no-fault” system (see chapter IV below), this would become moot.

III. Recommendations

This section is divided in two parts. First, we outline recommendations meant to improve the liability regime. Next, we lay forth other more general recommendations which go somewhat beyond narrowly-defined liability issues.

Before delving into this matter, a more general issue is determining what the overall objectives of a liability system should be⁷⁰. Many companies have claimed that AI development should be encouraged and that liability should therefore be limited, excluded or capped. They claim that actual lawsuits and even sometimes the threats of potential lawsuits will discourage research and then adoption of AI solutions. Because there is a race between scientists and between (most) countries⁷¹ to be the first to develop and then launch new AI tools, the countries with the most beneficial legal regime and its scientists are most likely to win it⁷². A comparison is sometimes made with other innovative medical techniques, which were first introduced without absolute guarantee of safety, but would never have been tried if such guarantees had been required from the start. For example, looking back in time, no one would have risked an

⁷⁰ On this issue with also the benefit of a historical overview, see PICHONNAZ/WERRO (2018), pp. 2-26 ; still on this issue, but under U.S. law and with regard to AI, KOWERT.

⁷¹ The Swiss government does not appear to be very interested by AI technology. Its September 2018 "Plan d'action Suisse numérique" mentions several initiatives, but AI is in no way highlighted (see its page 27 and point 7.5).

⁷² "The global AI developing ecosystem is already in a frenzy. [...] Funding more than tripled between 2016 and 2017 alone, to reach over 11 billion euro". EUROPEAN POLITICAL STRATEGY CENTRE (2018), p. 2. See also EUROPEAN COMMISSION (April 2018) ; European Commission's press release of 25 April 2018 titled "Artificial intelligence : Commission outlines a European approach to boost investment and set ethical guidelines".

organ transplant or in vitro fertilization if any mishaps – and there were many – would have resulted in lawsuits and victorious claims. A historical comparison can also be made to cars : the first automobiles were dangerous, at least judged by today's standards ; it is only because the law allowed their producers to progressively add security features that they quickly became such successful products.

Others have claimed the reverse⁷³. It is only if the patients and the physicians have full confidence in AI that AI tools will be adopted. Full confidence requires strong patient rights, which include strong rights to claim compensation. Moreover, the threat of lawsuits will push developers to create better and safer products. They are the parties best positioned to incorporate safety considerations into their design early on ; they are also best positioned to properly test their AI. Therefore, strong liability rules may create a virtuous circle, in which good AI tools are created and then exploited for the greater good, while actual lawsuits remain rare.

At this stage, there are no empirical ways to test the two viewpoints. For transparency's sake, I will admit my preference for the second perspective. Trust is essential in medicine. The potential for large-scale harm is so great that innovations ought to be carefully tested and accompanied by proper safeguards. The threat of lawsuits, at least to a certain extent, operates as such a safeguard. The history of pharmaceuticals' and medical devices' regulations shows that, without such a credible threat, massive harms can occur (e.g., elixir sulfanilamide, thalidomide, breast implants).

This preference also leads me to suggest improvements that mainly favor the interests of injured patients.

A. Recommendations to improve the current liability regime

In my view, based on the developments above, the following measures should be implemented :

- a) All AI tools should fall within the scope of the PLA, regardless of whether or not they are incorporated into hardware. The distinction between products and services makes less and less sense with respect to IT tools. Since the risks and benefits are the same, whether or not the product is physically incorporated, the legal regime should be the same. This corresponds to the reasonable expectations of users. A patient does not

⁷³ "[C]alls for innovation without any regulation must be viewed critically". GUIHOT ET AL. (2017), p. 419. The authors look carefully at the various regulatory schemes, including self-regulation, that could come into consideration. See also CHUNG (2017), p. 38.

know whether the AI recommendations was issued on-line or stemmed from a tangible machine – and it should make no difference.

- b) The development risk exception (Article 5, letter e PLA) should be removed and AI should follow the fate of transplants⁷⁴. There are inherent risks in AI that cannot be assessed with certainty by pre-launch testing, risks which will evolve as the AI continues to learn on its own⁷⁵. Because a development risk escape clause will almost always bar patients' lawsuits, the public interest in a fairer social regime calls for discarding this exception.
- c) Solidarity rules should be carefully considered⁷⁶. The present system of the PLA is that any producer is jointly and severally liable towards the injured patient for the entire amount of the damage⁷⁷; only after having paid what is owed to the patient, can this producer sue the other liable parties in the chain of production. This system could become increasingly unfair if the entity being sued has little to do with the defect that caused the harm. Ideally, a liability system should place liability on the person most capable of preventing or mitigating the harm⁷⁸. In this way, liability claims or threats of such claims can truly act as a deterrent for dangerous commercial practices. If the system operates so that Swiss suppliers or Swiss importers are the only parties being sued by patients in Switzerland, while the AI developers in China or in the United States remain de facto sheltered, this objective of product liability regime will hardly be attained.
- d) Since the development of AI is likely to involve multiple parties in multiple countries, the most efficient solution might be to introduce mandatory insurance for each AI tool⁷⁹. This is also the thinking behind vehicle insurance: people drive a variety of cars, including car which they do not own, across geographic frontiers. Victims should be indemnified, regardless of whether the car (or the AI) was built in one country, registered in a different one, owned by someone domiciled in a third country and (badly) used in a fourth one⁸⁰. With a system of insurance

⁷⁴ Article 5 al.1bis PLA ("organes, tissus ou cellules d'origine animale ainsi qu'aux transplants standardisés issus de ceux-ci destinés à être transplantés"). See also the regime for genetically modified organisms, explained by KRELL ZBINDEN/LANGHORST (2018), pp. 82-109.

⁷⁵ See also PICHONNAZ/WERRO (2018), p. 25.

⁷⁶ On this topic with regard to AI, but under U.S. law, see KOWERT, pp.202-203.

⁷⁷ Article 7 PLA.

⁷⁸ See EUROPEAN PARLIAMENT, point 56.

⁷⁹ See EUROPEAN PARLIAMENT, point 57. In its Resolution, the European Parliament recommends adding a public fund "to ensure that reparation can be made for damage in cases where no insurance cover exists". In my view, it is still too early to decide whether there would really be gaps in insurance coverage that make a fund necessary. Others, on the contrary, have suggested introducing liability caps, so that the overall damages that could be claimed would be known to be limited to a certain ceiling. ALLAIN (2013), p. 1076.

⁸⁰ Insurance may also somewhat simplify issues related to jurisdiction and applicable law.

coverage, victims would know that they would be properly indemnified in a relatively quick manner, without having to master the complete history of the AI product (e.g., who developed it where, who fed it data where, who imported it into Switzerland, who used it professionally). With the insurance company acting as the pivot point, it would then (i.e., after the patient received compensation) be easier to allocate responsibility among the involved parties, based either on their share of the guilt/fault or on other criteria (e.g., contract clauses).

We realize that these pro-consumer recommendations are likely to run into political obstacles. As Pichonnaz/Werro wrote : "Le futur de la responsabilité pour risque et sa généralisation sont des questions éminemment politiques"⁸¹.

B. Other recommendations

This subsection sets forth additional recommendations more indirectly connected with liability. We comment only briefly on them, since a detailed justification would go beyond the scope of this paper.

- a) The legislature should institute clear procedures for AI medical tools to be preauthorized, either by a medicine agency (e.g., Swissmedic in Switzerland) or by a specialized *ad hoc* agency. The authorization procedure should be strict, whether we apply the rules of high-risk medical devices or those for allopathic drugs. Self-certification or certification by notified bodies (as currently used for most medical devices) should not be held sufficient. The authorization process should be sufficiently transparent for professional users to know which level of security and efficacy to expect from the released AI tools.
- b) Producers of AI should be obliged to keep on supervising and studying their tools even after release. Because an AI tool learns continuously, it should be tested and retested periodically. Indeed, reliability results at time t_1 may no longer hold true at time t_2 . The law should specify how this ongoing testing is to be arranged, and whose responsibilities it should be (e.g, the initial producer only, the hospital now using the AI, or both of them together).
- c) Patients should be told whether AI is or was used in the course of their medical care⁸². Because the accuracy of the data fed to the AI is so crucial, patients should be made aware of the importance of the information that they themselves provide. As the fictional Dr. House says "all patients lie" ;

⁸¹ PICHONNAZ/WERRO (2018), p. 25.

⁸² See also VOKINGER ET AL., (2017), pp. 15-16.

whereas good doctors may be able to spot a lie, AI is unlikely to master this social skill – at least in the near future.

- d) Patients should also be given information about AI involvement in their medical care, as it will help them understand the “thought process” of their doctors. For example, if a patient is told that the doctor relied only on the diagnostic provided by the AI tool, she can challenge this conclusion and seek a different medical opinion elsewhere. Without this knowledge, the consent she would give to the suggested treatment course would not be truly informed⁸³.
- e) Based on the autonomy principle, patients should have the final say as to whether to allow AI to access their medical files⁸⁴. Whether consent is by opt-in or by opt-out remains open to discussion⁸⁵. However, it should be held that patients' data in medical files can never be viewed as fully anonymized and that AI's access to such files falls under the legal notion of data processing⁸⁶. Moreover, the data are sensitive data, since they contain health information. Therefore, in conformity with Data Protection Law, the individual is entitled to decide whether to confer or exclude access. The higher good hoped from AI data analysis and AI services should not bypass this requirement dictated by the fundamental interest in autonomy⁸⁷.
- f) AI tools should not be allowed to pool or merge the data from different databases without patients' approval, and for sensitive fields, without a prior authorization by a competent authority. One can imagine a future where the AI uses data beyond medical dossiers, and hunts for data on-line regarding the patient (e.g., Facebook's postings)⁸⁸. AI tools might also be programmed to allow the sharing of data among each other⁸⁹. Being able to use data from government databases (e.g., the tax file) might also be tempting, since socioeconomic data can be very informative, including for

⁸³ See NUFFIELD COUNCIL ON BIOETHICS (2018), p. 5.

⁸⁴ Because patients' files at least in Switzerland are not (yet) standardized, it will be difficult for AI to just access and use these data directly. However, with the introduction of the patient electronic dossier, standardization may improve and usefulness of these files for AI tools will improve too.

⁸⁵ On the difficulties related to patient's consent in relation with Big Data, see, e.g., DOCHERTY/LONE (2015), p. 470.

⁸⁶ Article 3 letter e of the Swiss Federal Act of 19 June 1992 on Data Protection (“FADP” ; RS 235.1). In the European Union, see MAISNIER-BOCHÉ (2017), p. 25.

⁸⁷ Withholding consent may become increasingly difficult, as CHAR ET AL. (2018) explain in their article at p. 983.

⁸⁸ See NUFFIELD COUNCIL ON BIOETHICS (2018), p. 5.

⁸⁹ See Directorate-General for Internal Policies (2016), p. 22 (“once robots have the skills to communicate and interact, not only will they exchange data [...], but this communication may be imperceptible to humans”).

medical purposes. Merging of data by AI tools should therefore not be allowed without strict legal requirements⁹⁰.

C. Further considerations

In this third subsection, we move away from liability issues to present a sampling of issues of significance related to the medical use of AI. Once more, our comments will be brief.

1) *Some cost issues*

A regular question cropping up in AI debates is the potential cost saving of AI medical tools. Some claims that AI has the potential to massively reduce health care costs. In our country where health care represents 12% of GDP, such a promise is enticing. Even slight reductions would be tempting. It is true that the main budget line is the wages of doctors, whether doctors in private practice or doctors in medical institutions. Thus, if AI could replace at least some of the doctors presently employed, significant savings could be achieved. For example, if far fewer dermatologists or radiologists were needed to diagnose patients⁹¹, this would have an impact on medical expenses ; likewise if some tasks could be delegated to an AI rather than being performed by nurses (e.g., checking on patients). One might also reasonably suppose that if the diagnoses are more reliable, unnecessary and costly care will be entirely avoided ; for example, patients who would have been needlessly treated for a falsely-diagnosed breast cancer will no longer incur the pain and the costs, if the AI tool can better differentiate between benign tumors and those requiring treatment.

Others claim that AI will not dent the current trend of higher medical expenses. It will "behave" as any other medical tool put on this unique market ; it will not be cheap⁹² and will generate its own additional costs. Possible savings will be offset by the possibility to offer even more or even better health care in another sector. For example, cancers will be diagnosed sooner in younger patients who will then have to undergo regular check-ups for the rest of their (long) lives.

The answer to this cost question has some (indirect) impact on liability. If an overall cost reduction is achieved, the impetus to use AI tools will be stronger. This may in turn cause more damages imputable to AI, simply

⁹⁰ A different yet interesting question is whether AI systems could be forced to share their data in the name of "open science".

⁹¹ See BAMBAUER (2017), p. 393.

⁹² "Watson generates fees between \$200 and \$1000 per patient." CHUNG (2017), p. 38.

because their use will have become common. As mentioned earlier, this may also make it harder to reach a conclusion with respect to causality : How can the judge decide that the damage would not have incurred if no AI tool had been used, if there is no longer a human benchmark to make the comparison ?

Another relevance of cost has to do with the feasibility of offering compensation. If significant savings can be obtained, there is a stronger case for redirecting (part of) these savings towards damage indemnification. If the savings benefit the State as payer, the State could then create an indemnification fund or co-finance a social insurance system.

Finally, cost issues may also be viewed in the light of “de-skilling” risks. As Froomkin et al. write : “overreliance on these machines could render obsolete the human cultivation of medical skills and knowhow developed over centuries”. Thus, certain medical professions could disappear or be profoundly altered. With European economies already suffering from the weighty financial burden of high unemployment rates and correspondingly low tax revenues, AI could make the situation even worse⁹³. Thus, the savings achieved by having less need for doctors could be – at least in the short-run – offset by the need to pay social benefits to those professionals now without work. Hiring the human doctors made redundant to oversee the continuous efficacy of AI tools could be a (partial) solution⁹⁴.

2) *Some antitrust issues*

AI are being developed both by giant companies and by start-ups. However, it appears likely that the IT giants will be first on the market. Given that medical AI tools may be seen as high-risk devices, hospitals may well favor the most-established players. To know that IBM is behind an AI diagnostic tool is more reassuring than if it is a start-up still relying on venture capital to make ends meet. However, this may imply an ever greater control of a few companies over our entire medical histories. The logic behind AI is that the more the data and the better the data⁹⁵, the more reliable the output⁹⁶. However, making a huge volume of data available to just a few companies raises concerns. Merging of data across various databases is likely to make re-identification of

⁹³ See GUIHOT ET AL. (2017), p. 411-413 ; GOUVERNEMENT FRANÇAIS, p. 2 ; LOH (2018) ; U.K. SCIENCE AND TECHNOLOGY COMMITTEE (2016).

⁹⁴ See FROOMKIN ET AL. (2018), p. 63.

⁹⁵ "The world is estimated to produce more than 2.5 quintillion bytes of data every day (a quintillion is 1 followed by 18 zeros)". DOUGLAS B. FRIDSMAN, *Health informatics : a required skill for 21st century clinicians*, 362 *BJM* Editorial (July 2018). According to DELOITTE (2015) citing IBM, "90 percent of the world's data was generated just in the past two years".

⁹⁶ See, e.g., EUROPEAN COMMISSION (April 2018) ("access to data is a key ingredient for a competitive AI landscape, which the EU should facilitate").

the individual patients possible ; harms caused by possible security breaches are then magnified.

Additionally, having a few companies exercising control over medical AI may trigger concerns about possible abuse of dominant positions⁹⁷. These companies may charge unfair prices ; they may impose unfair conditions, for example relating to data access and data re-use ; they may discriminate between clients favoring some over others (e.g., larger hospitals securing better deals). Their behavior on the market may be viewed as opaque. They may discourage innovation from smaller players, given the high likelihood for a possible network or platform effect.

For these reasons, the development of medical AI should go hand in hand with careful surveillance of the market to avoid the emergence of dominant players⁹⁸. Similarly, authorities should make sure that anticompetitive agreements are not entered into between AI developers and medical institutions. In particular, there should be no exclusivity clause blocking a health care facility with a certain AI developer. Access to medical data should be on fair and equal terms.

3) *Some ethical issues*

Use of AI will also drive us to question the patient-doctor relationship. AI is yet one more tool in the field of medicine that appears increasingly dehumanized⁹⁹. Of course, a better diagnosis is better for patient, as is a better treatment choice. However, when more and more tasks can be delegated to a non-human tool, the place that was ascribed to this central patient-doctor relationship evaporates. If or when AI robots can do everything better than doctors, why would a patient still see a doctor ?

Yet in many cases the patient is helped by the kind words or the generally benevolent attitude of the doctor or of the nurse. Knowing that someone really cares is an efficient yet subtle aspect of health care. Patients may fare worse just because they can no longer talk with a human being, even if that talk had little to do with medical decision-making¹⁰⁰. Social relationships matter, yet AI tools may lead us to discounting their effectiveness, as AI tools may exhibit

⁹⁷ See more generally, BENSAMOUN (2017), p. 31. Also LUTZMANN (2017), p. 386.

⁹⁸ "A proper enforcement of merger control, antitrust and state aid rules can prevent market distortions and the creation of bottlenecks in the digital value chain." See EUROPEAN POLITICAL STRATEGY CENTRE (2018), p. 11.

⁹⁹ See e.g., LOPER (2018) ; SHAPSHAY (2014), p. 191.

¹⁰⁰ As the EUROPEAN PARLIAMENT (2017) points out : "human contact is one of the fundamental aspects of human care ; [...] replacing the human factor with robots could dehumanize caring practices" (point 32). Some authors believe that patients should be entitled to refuse the involvement of robots in their care. See, e.g., DIRECTORATE-GENERAL FOR INTERNAL POLICIES (2016), p. 21 ("not accepting this refusal would violate their dignity") ; KIEFER (2017).

higher specificity (true negative rate) and sensitivity (true positive rate) when delivering a diagnosis or when choosing the right treatment (as compared to human action). Human interactions may also come to be perceived as just too expensive, when their costs are compared to that an AI robot¹⁰¹.

Thus, AI may force us to rethink the social nature of medicine. The new role of doctors may become less technical and more social. Of course, this is not – or not yet – what doctors have trained for. Such a radical change in the ascribed mission of doctors may feel uncomfortable to some. It may come with a loss in social status: the doctor may no longer be viewed as a respected medical expert, but more as a friend holding one's hand.

AI may also have direct adverse consequences on health care providers. Once AI tools become able to keep track of all human actions in a given medical setting, they may be used to point out the failure of human staff. For example, the AI may point out that the doctor did not wash her hands when entering the patient's room¹⁰², that the surgeon did not follow the hospital procedure when operating, that the nurse failed to check on certain patients. Working in surroundings where AI constantly checks on the employees' behavior can lead to a toxic work environment. It will therefore be necessary to impose certain constraints or employee's safeguards as AI is deployed for surveillance functions.

Finally, the issue of bias in medicine is often highlighted, as AI may build on and amplify such bias¹⁰³. Because AI learns directly from the data it is being fed with, often with very few human filters, AI is prone to reproduce the bias and mistakes of human doctors. It may even reinforce them¹⁰⁴. For example, if a AI tool is being given pictures of possible melanoma on white and brown skins and the labeling benign vs. cancerous is more often false in the case of patients with brown skins, AI will repeat the same mistake over and over. Or if women suffering from heart attacks are usually offered a less aggressive pharmaceutical treatment than men, the AI tool may assume this is indeed the "normal" standard of treatment and suggest the same when interrogated by health care professionals. Similarly, data samples from which AI tools would learn may not represent the entire population, because patients in the samples only come from certain groups (e.g., richer patients visiting this private clinic or poorer patients going to an emergency service); if the data exploited by AI

¹⁰¹ "[A]s it will be less expensive for people to have robots than a human helper at home, there is a risk that machines becoming the norm and people the exception". DIRECTORATE-GENERAL FOR INTERNAL POLICIES (2016), p. 24.

¹⁰² See YEUNG ET AL. (2018), p. 1272.

¹⁰³ NUFFIELD COUNCIL ON BIOETHICS (2018), p. 5; PACKIN/LEV-ARETZ; LOH (2018), p. 3; CHEN/ASH (2017), p. 2507; CHAR ET AL. (2018), p. 982. For examples of AI bias beyond medicine, see EUROPEAN POLITICAL STRATEGY CENTRE (2018), p. 7.

¹⁰⁴ See FROMKIN ET AL. (2018). According to certain authors, "AI-based predictions should only be relied on if someone is continuously checking predictions against reality". Id. p. 48.

are not representative of the population¹⁰⁵, relevant characteristics unique to certain groups may be overlooked¹⁰⁶. Because human beings are rarely aware of their biases or potential for bias, they are usually not able to educate AI tools to overcome them. In addition, broad scale use of AI tools may hide these biases, since AI often operates as a "black box"; it would then be even harder to identify and correct these recurrent mistakes. Identifying solutions to this problem will not be easy¹⁰⁷.

IV. Conclusion

Ideally, an adequate regime of AI should be agreed upon internationally so that the standard adopted is harmonized. Having a unique or at least roughly equivalent standard would create a level-playing field where companies could compete in a prosocial manner. This would avoid a race to the bottom where the cheaper products would come at the price of inadequate patient safety (whether physical harm or privacy harm). In the case of AI, patients are not in a position to choose the safer products. Even medical institutions may not be able to make the right choice, because they too may not fully understand how AI functions. Moreover, AI is likely to be licensed based on complex contract terms, which physicians will certainly not read. Finally, if AI tools become available only in a purely dematerialized way (i.e., purely on-line), users may not even be able to locate where the AI developer is based (or the AI developers are based). These difficulties are already occurring now with some internet services.

Hence, it is crucial to ensure that all AI tools meet basic safety standards, regardless of where they were produced. This uniform regime should guarantee that patients are not barred from opening action in their own jurisdiction under the law of their own country.

As a further step, an internationally agreed-upon regime could propose a basic insurance format, obliging all AI product to be licensed with such a standard insurance coverage. Thus, a hospital subscribing a license for an AI tool would know that a portion of the price paid would be pooled by a duly appointed insurance company to cover certain patient damages.

AI has the potential to improve human health. Among the most important questions we may need to ask ourselves is whether this is the type of

¹⁰⁵ See TERRA NOVA (2017), p. 30.

¹⁰⁶ For example, the AI Watson performed poorly when used in a Korean hospital, apparently because it had been trained using the data from US patients. CHUNG/ZINK (2017), p. 5.

¹⁰⁷ Proposing some solutions, GIANFRANCESCO ET AL. (2018), p. 1546. Most of these proposals, although relevant, are costly to implement (e.g., "ensure interdisciplinary approach and continuous human involvement"; "conduct follow-up studies to ensure results are meaningful").

improvement that we most need. AI will be deployed – at least at the beginning – in specialized medical settings to deliver specialty health care. However, there are other areas where the basic health services are still sorely lacking¹⁰⁸. Hence, excitement over such sophisticated tools seems a bit out of place when one realizes that millions of people still lack basic access to health care – let alone good and affordable health services. Even in Switzerland, many people encounter financial hardships when faced with medical bills¹⁰⁹ – because the medical services are not reimbursed (e.g., dental care), because of high deductibles, because of the high copay or because of the ever-increasing monthly premiums¹¹⁰.

Investing in broader access to health care and in better preventive services is likely to lead to greater social returns than investing in AI. Choosing the right priorities for our times is perhaps more important than dreaming of futuristic tools. Of course, one is not necessarily exclusive of the other. However, medicine tends to focus and rely on technical tools, sometimes excessively – forgetting about the simple measures that make the greatest difference.

As a conclusion, we will remind ourselves that health is primarily determined by social factors – first among them the wealth, the social and the educative attainment level of the individual. Statistically speaking, being rich, highly-educated and well-integrated in society leads to an overall higher life expectancy in good health – as compared to those who are not so fortunate. At least for now, AI is unlikely to change this grim but fundamental truth. Honesty and fairness call for us doing our best to improve the situations of those most in need. AI may contribute to this goal, but how exactly is still open for debate.

¹⁰⁸ For a short but poignant piece highlighting the gap between our medicine and the medicine of poorer countries, see LILIC (2018), p. 1202.

¹⁰⁹ "Lorsqu'un assuré ne paie pas sa prime, c'est au canton de combler le manque à gagner des caisses à hauteur de 85%. En 2016, les cantons romands ont ainsi versé 163,7 millions de francs aux assureurs, selon les chiffres recueillis par la RTS. Le montant total des primes impayées en Suisse est quant à lui estimé à 843 millions de francs par l'Office fédéral de la santé publique (OFSP)." THÉO ALLEGREZZA, *Les cantons règlent des millions de francs de primes maladie impayées*, 12 October 2017, RTS, at <https://www.rts.ch/info/suisse/8982179-les-cantons-reglent-des-millions-de-francs-de-primes-maladie-impayees.html> (consulté le 1er août 2019).

¹¹⁰ "La hausse de 1,2% est inférieure à la moyenne des années précédentes. Depuis 2008, la prime moyenne a augmenté de 3,5% par an et de 3,9% par an depuis l'entrée en vigueur de l'assurance obligatoire des soins en 1996." Federal Office for Public Health, press release of 24 September 2018.

V. Reference list

- Académie (française) des technologies, *Renouveau de l'Intelligence artificielle et de l'apprentissage automatique*, Rapport, March 2018, http://academie-technologies-prod.s3.amazonaws.com/2018/06/01/14/25/31/28/IA_Web.pdf (dernière consultation le 1er août 2019).
- ALLAIN Jessica, *From Jeopardy! to Jaundice: The Medical Liability Implications of Dr. Watson and Other Artificial Intelligence Systems*, 73:4 *Louisiana Law Review*, 2013, pp.1049-1078, <https://digitalcommons.law.lsu.edu/cgi/viewcontent.cgi?article=6423&context=lalrev> (dernière consultation le 1er août 2019).
- BAMBAUER Jane, *Dr. Robot*, 51:383 *University of California, Davis*, 2017, pp. 383-398, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3087983 (dernière consultation le 1er août 2019).
- BAUME Cédric, *Gestion des risques cliniques et responsabilité médco-hospitalière*, Collection IDS - Unine, 2015.
- BENSAMOUN Alexandra, *Intelligence artificielle et santé : l'intégration en droit de l'IA médicale*, 17 *Journal de droit de la Santé et de l'Assurance Maladie*, 2017.
- CHAR Danton S. et al., *Implementing Machine Learning in Health Care - Addressing Ethical Challenges*, 378:11 *New England Journal of Medicine* 2018, pp. 981-983.
- CHEN Jonathan/ASCH Steven, *Machine Learning and Prediction in Medicine - Beyond the Peak of Inflated Expectations*, 376:26 *New England Journal of Medicine*, 2017, pp. 2507-2509.
- CHOI Edward et al., *Doctor AI: Predicting clinical events via recurrent neural networks*, 56 *Proceedings of Machine Learning for healthcare*, pp. 301-318 2016, <http://proceedings.mlr.press/v56/Choi16.pdf> (dernière consultation le 1er août 2019).
- CHUNG Jason, *What should we do about artificial intelligence in health care*, 22:3 *Health Law Journal*, 2017, pp. 37-40.
- CHUNG Jason/ZINK Amanda, *Hey Watson, Can I sue you for malpractice? Examining the liability of Artificial Intelligence in medicine*, *Asia-Pacific Journal of Health Law, Policy and Ethics*, 2017, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3076576 (dernière consultation le 1er août 2019).

- DELOITTE, Disruption ahead, Deloitte's point of view on IBM Watson, March 2015.
- DIPROSE William/BUIST Nicholas, Artificial intelligence in medicine : humans need not apply? 129 NZMJ May 2016, pp. 73-76, https://www.nzma.org.nz/__data/assets/pdf_file/0007/49579/Diprose-2155FINAL1434.pdf (dernière consultation le 1er août 2019).
- Directorate-General for Internal Policies (European Parliament), Policy Department, Citizens' Rights and Constitutional Affairs, European Civil Law Rules in Robotics, Study for the JURI Committee, 2016, http://www.europarl.europa.eu/RegData/etudes/STUD/2016/571379/1/POL_STU%282016%29571379_EN.pdf (dernière consultation le 1er août 2019).
- Ernst & Young (EY) with the support of Technopolis Group and VVA Consulting, European Commission, Evaluation of Council Directive 85/374/EEC on the approximation of laws, regulations and administrative provisions of the Member States concerning liability for defective products, Final Report, January 2018, <https://publications.europa.eu/en/publication-detail/-/publication/d4e3e1f5-526c-11e8-be1d-01aa75ed71a1/language-en> (dernière consultation le 1er août 2019).
- European Commission, Artificial Intelligence for Europe, Communication, COM(2018) 237 final, April 2018, <https://ec.europa.eu/digital-single-market/en/news/communication-artificial-intelligence-europe>.
- European Commission, Report on the application of the Council Directive on the approximation of the laws, regulations, and administrative provisions of the Member States concerning liability for defective products (85/374/EEC), SWD(2018) 157 final, COM(2018) 246 final (May 2018), <https://ec.europa.eu/transparency/regdoc/rep/1/2018/EN/COM-2018-246-F1-EN-MAIN-PART-1.PDF>.
- European Commission Staff, Evaluation of Council Directive 85/374/EEC of 25 July 1985, Working Document SWD(2018) 157 final, May 2018, <https://ec.europa.eu/transparency/regdoc/rep/10102/2018/EN/SWD-2018-157-F1-EN-MAIN-PART-1.PDF>.
- European Parliament, European Parliament resolution of 16 February 2017 with recommendations to the Commission on Civil Law Rules on Robotics, P8_TA(2017)0051, 2017, <http://www.europarl.europa.eu/sides/getDoc.do?pubRef=-//EP//NONSGML+TA+P8-TA-2017-0051+0+DOC+PDF+V0//FR>.
- European Parliament Research Service (EPRS), Scientific Foresight Unit (STOA), Should we fear artificial intelligence ? In-Depth Analysis, Science and Technology Options Assessment, March 2018,

http://www.europarl.europa.eu/RegData/etudes/IDAN/2018/614547/EPRS_IDA%282018%29614547_EN.pdf.

European Parliamentary Research Service Scientific Foresight Unit (STOA), A Common EU approach to liability rules and insurance for connected and autonomous vehicles, European Added value Assessment, February 2018, [http://www.europarl.europa.eu/RegData/etudes/STUD/2018/615635/EPRS_STU\(2018\)615635_EN.pdf](http://www.europarl.europa.eu/RegData/etudes/STUD/2018/615635/EPRS_STU(2018)615635_EN.pdf).

European Political Strategy Centre, The Age of Artificial Intelligence, Towards a European Strategy for Human-Centric Machines, EPSC Strategic Notes, March 2018, https://ec.europa.eu/epsc/sites/epsc/files/epsc_strategicnote_ai.pdf.

FROOMKIN Michael et al., When AIs Outperform Doctors: The dangers of a tort-induced over-reliance on machine learning and what (not) to do about it, University of Miami Legal Studies Research Paper No. 18-3, available at SSRN: <https://ssrn.com/abstract=3114347> or <http://dx.doi.org/10.2139/ssrn.3114347>.

GIANFRANCESCO Milena et al., Potential Biases in Machine Learning Algorithms Using Electronic Health Record Data, 178 (11) JAMA Internal Medicine, 2018, <https://jamanetwork.com/journals/jamainternalmedicine/article-abstract/2697394>.

GIL Yolanda et al., Amplify scientific discovery with artificial intelligence, 346 Science Magazine 2014, pp. 171-172, <http://science.sciencemag.org/content/346/6206/171/tab-pdf>.

GOLDEN Jeffrey Alan, Deep learning Algorithms for Detection of Lymph Node Metastases from Breast Cancer, 318(22) JAMA 2017, pp.2184-2186, <https://jamanetwork.com/journals/jama/fullarticle/2665757>.

Gouvernement français, France Intelligence Artificielle, Rapport de synthèse, 2017, <https://www.ladocumentationfrancaise.fr/var/storage/rapports-publics/174000247.pdf>.

GREENBERG Anastasia, Artificial Intelligence in Health Care: Are the Legal Algorithms Ready for the Future? McGill Journal of Law and Health (2017), <https://mjlh.mcgill.ca/2017/10/06/artificial-intelligence-in-healthcare-are-the-legal-algorithms-ready-for-the-future/>.

GUIHOT Michael et al., Nudging Robots: Innovative Solutions to Regulate Artificial Intelligence, Vanderbilt Journal of Entertainment and Technology Law, 20(2), 2017, pp. 385-456, <https://eprints.qut.edu.au/109926/1/Guihot%20Matthew%20Suzor%2017%20Nudging%20Robots.pdf>.

- JACQUEMIN Hervé/de Streel Alexandre, *L'intelligence artificielle et le droit*, Editeur Larcier, Collection CRIDS, 2017.
- KIEFER Bertrand, *L'IA, l'autre et le soignant*, *Revue médicale Suisse*, 2017, p. 1816.
- KOWERT Weston, *The Foreseeability of Human-Artificial Intelligence Interactions*, 96:1 *Texas Law Review*, 2017, pp. 181-204, <https://texaslawreview.org/wp-content/uploads/2017/12/Kowert.pdf>.
- KRELL ZBINDEN Karola/Langhorst Maren, « Les organismes génétiquement modifiés - Une responsabilité de génie », in: *Les responsabilités fondées sur le risque*, Colloque du droit de la responsabilité civile 2017, Ed. F. Werro/P. Pichonnaz, Stämpfli, 2018.
- LEHMAN-WILZIG Sam, *Frankenstein Unbound, Towards a legal definition of Artificial Intelligence*, *Futures* 1981, pp. 442-457, https://www.researchgate.net/publication/229067941_Frankstein_unbound_Towards_a_legal_definition_of_artificial_intelligence (dernière consultation le 1er août 2019).
- LILIC Nikola, *Murmurs of Politics and Economics*, 379:13 *New England Journal of Medicine*, 2018, pp. 1202-1203, <https://www.nejm.org/doi/pdf/10.1056/NEJMp1806115> (dernière consultation le 1er août 2019).
- LOH Erwin, *Medicine and the rise of the robots : a qualitative review of recent advances of artificial intelligence in health*, *BMJ Leader*, June 2018, <https://bmjleader.bmj.com/content/leader/early/2018/06/01/leader-2018-000071.full.pdf>.
- LOPER Peter Louis, *The Electronic Health Record and Acquired Physician Autism*, 172(11) *JAMA Pediatrics* 2018, p. 1009, <https://jamanetwork.com/journals/jamapediatrics/fullarticle/2703482> (dernière consultation le 1er août 2019).
- LUTZMANN Steve, *Digitalisierung und Rechssicherheit*, *GesKR (Gesellschafts- und Kapitalmarktrecht)* 2017, pp. 383-386.
- MAISNIER-BOCHÉ Lorraine, *Intelligence artificielle et données de santé*, 17 *Journal de droit de la Santé et de l'Assurance Maladie*, 2017, pp. 25-29.
- MARCHAND Sylvain, « Les fondamentaux de la responsabilité du fait des produits - Exposé introductif : sources, for et droit applicable, producteur et défaut », in: *La responsabilité du fait des produits*, Journée de responsabilité civile 2016. Ed. Chappuis et Winiger, Schulthess 2018, pp. 11-30.
- MORISSE Tom, « Enjeux et histoire de l'Intelligence Artificielle », in: *Intelligence Artificielle, un nouvel horizon : Pourquoi la France a besoin d'une culture du numérique*, Les Cahiers Lysias, 2017.

- NUFFIELD COUNCIL ON BIOETHICS, Artificial Intelligence (AI) in healthcare and research, Bioethics Briefing Note May 2018, <http://nuffieldbioethics.org/wp-content/uploads/Artificial-Intelligence-AI-in-healthcare-and-research.pdf> (dernière consultation le 1er août 2019).
- PACKIN NIZAN Geslevich/Lev-Aretz Yafit, Research Handbook of Law and AI (Woodrow Barfield & Ugo Pagallo eds.) Baruch College Zicklin School of Business Research Paper No. 2018-04-03, 2018.
- PETRELLUZZI Julia/MILIOTIS Marie, L'intelligence artificielle : ses enjeux en droit civil européen, Editions Universitaires Européennes (EUE), Editeure : M. Clément-Fontaine, 2017.
- PICHONNAZ Pascal/WERRO Franz, « La responsabilité fondée sur le risqué : un état des lieux et quelques perspectives d'avenir », in Les responsabilités fondées sur le risque, Colloque du droit de la responsabilité civile 2017, Ed. F. Werro/P. Pichonnaz, Stämpfli, 2018, p.1-26.
- SCHWARTZ William et al., Artificial Intelligence in Medicine, Where we stand ?, 316:11 New England Journal of Medicine, 1987, pp. 685, <https://www.nejm.org/doi/pdf/10.1056/NEJM198703123161109> (dernière consultation le 1er août 2019).
- SHAPSHAY Stanley, Artificial Intelligence, The Future of Medicine ? 140:3 JAMA Otolaryngology, 2014, p. 191, <https://jamanetwork.com/journals/jamaotolaryngology/fullarticle/1812651> (dernière consultation le 1er août 2019).
- SHORTLIFFE Edward/SEPÜLVEDA Martin, Clinical Decision Support in the Era of Artificial Intelligence, JAMA e1 (on-line), 2018, https://jamanetwork.com/journals/jama/fullarticle/2713901?guestAccessKey=73555101-6acd-47b6-b1bf-b7ffd6265b37&utm_source=silverchair&utm_medium=email&utm_campaign=article_alert-jama&utm_content=olf&utm_term=110518 (dernière consultation le 1er août 2019).
- TERRA Nova, La santé à l'heure de l'intelligence artificielle, <http://tnova.fr/notes/la-sante-a-l-heure-de-l-intelligence-artificielle> (dernière consultation le 1er août 2019).
- TRAN Laurent/ETIER Guillaume, « Responsabilité du fait des produits : le risque de développement », in: La responsabilité du fait des produits, Journée de responsabilité civile 2016. Ed. Chappuis et Winiger, Schulthess 2018, pp. 105-130.
- TSCHIDER Charlotte, Deus ex Machina : Regulating Cybersecurity and Artificial Intelligence for Patients of the Future, 5 Savannah L. Rev. 177, 2018, available at SSRN : <https://ssrn.com/abstract=3070000> (dernière consultation le 1er août 2019)

- U.K (United Kingdom) Science and Technology Committee, House of Commons, Robotics and artificial intelligence, Fifth Report of Session 2016-17, 2016,
<https://publications.parliament.uk/pa/cm201617/cmselect/cmsstech/145/145.pdf> (dernière consultation le 1er août 2019).
- VIGNERON-MAGGIO-APRILE Sandra, « Les fondamentaux de la responsabilité du fait des produits – Entre dommage et causalité : la perte d'une chance au secours du lésé », in: La responsabilité du fait des produits, Journée de responsabilité civile 2016. Ed. Chappuis et Winiger, Schulthess 2018, pp. 31-58.
- VOKINGER, Kerstin Noëlle et al., Artificial Intelligence and Machine Learning in der Medizin, Jusletter du 28 août 2017.
- WERRO Franz, « La responsabilité pour produits pharmaceutiques défectueux dans la jurisprudence récente », in: La responsabilité du fait des produits, Journée de responsabilité civile 2016., Ed. Chappuis et Winiger, Schulthess 2018, pp. 77-104.
- WERRO Franz, Commentaire de l'art. 398 CO, in: Commentaire romand du Code des Obligations I, Ed. Thévenoz/Werro, Helbing Bâle 2012, p. 2397 et ss.
- WISE Jacqui, AI system interprets eye scans as accurately as top specialists, 362 BMJ, August 2018,
<https://www.bmj.com/content/bmj/362/bmj.k3484.full.pdf> (dernière consultation le 1er août 2019).
- YEUNG Serena et al., Bedside Computer Vision – Moving Artificial Intelligence from Driver Assistance to Patient Safety, 378:14 New England Journal of Medicine 2018, pp. 1271-1273,
<https://www.nejm.org/doi/pdf/10.1056/NEJMp1716891> (dernière consultation le 1er août 2019).