CIVIL LIABILITY FOR ARTIFICIAL INTELLIGENCE IN MEDICINE: IS THERE A NEED FOR A NEW PARADIGM?

Petr Šustek,* Martin Šolc**

Abstract: *The rapid integration of artificial intelligence (AI) into medical practice raises pressing questions for civil liability law. While AI systems promise enhanced diagnostic and therapeutic capabilities, their autonomous and opaque black box nature pose unique legal challenges. This paper investigates whether existing tort law frameworks, particularly within the Czech legal system, are sufficient to address harms caused by medical AI. It systematically analyses three core liability regimes: breach of the standard of care, strict liability based on a defect of a thing or product, and liability for damage caused by a thing in and of itself. The authors argue that, although each of these regimes presents interpretive and evidentiary difficulties in the context of AI, none is fundamentally unfit for addressing AI-related harms. Instead of advocating for a wholly new legal regime or the creation of a fictive electronic person, the paper supports a cautious evolution of existing doctrines, refined through case law, medical guidelines, and professional standards. Special attention is given to the distinction between certified medical AI devices and general-purpose systems like large language models. The authors also explore the implications of EU legislation, notably the AI Act. The paper concludes that legal conservatism, rooted in pragmatism, may serve the justice system best as the practical use of AI in medicine continues to unfold.*

Keywords: *Civil liability for artificial intelligence, medical artificial intelligence, tort law, medical law, health law*

IN PLACE OF INTRODUCTION: NEW PARADIGM IN MEDICINE, NEW PARADIGM IN LAW?[1]

Judging by the steadily growing body of academic studies and expert opinions, medicine appears to be on the brink of a paradigmatic shift. Artificial-intelligence (AI) systems are, on the one hand, merely one of many technological innovations encountered in medicine; on the other hand, they have a remarkable capacity to permeate the discipline in its full breadth and to reshape clinical practice profoundly. Scientific literature increasingly provides evidence that, in specific domains, AI enhances physicians' performance.[2] In 2024, *Nature Medicine* published the first study to conclusively demonstrate that an AI system can significantly reduce patient mortality: specifically, a system that predicts sudden clinical deterioration in hospitalised patients by interpreting electrocardiograms (ECGs).[3]

* Associate Professor, JUDr. Petr Šustek, Ph.D., is the head of the Department of Medical Law and a member of the Department of Civil Law at the Charles University, Faculty of Law, Prague, the Czech Republic. ORCID: 0000-0003-3411-6493.

** JUDr. Mgr. Martin Šolc, Ph.D., is a member of the Department of Civil Law and the Department of Medical Law at the Charles University, Faculty of Law, Prague, the Czech Republic. ORCID: 0000-0002-6274-8562.

[1] This article was written with the support of the research project Donatio Universitatis Carolinae "Legal Support for New Technologies and Innovations in Medicine", no. DUC/24/PF/2.

[2] See for example AL ZO'UBI, Mazen. Review of 2024 publications on the applications of artificial intelligence in rheumatology. *Clinical Rheumatology.* 2025, Vol. 44, Issue 4, pp. 1427–1438. doi: 10.1007/s10067-025-07382-3, or NORI, Harsha, DASWANI, Mayank, KELLY, Christopher (eds.). Sequential Diagnosis with Language Models. In: *ArXiv* [online]. 27.6.2025 [2025-06-30]. doi: 10.48550/arXiv.2506.22405.

[3] LIN, Chin-Sheng, LIU, Wei-Ting, TSAI, Dung-Jang (eds.). AI-enabled electrocardiography alert intervention and all-cause mortality: a pragmatic randomized clinical trial. *Nature Medicine.* 2024, Vol. 30, Issue 5, pp. 1461–1470. doi: /10.1038/s41591-024-02961-4.

Must a paradigmatic shift in an activity already subject to legal regulation entail a corresponding transformation of the legal framework itself? In this article we ask whether tort law, in its current form, can satisfactorily address the challenges posed by the rise of medical AI. Our aim is not to provide definitive answers, but rather to outline possible solutions that practice will test in the near future. Although our analysis is rooted in Czech law, we believe that our conclusions are, in principle, applicable across legal systems.

I. THE CONCEPT OF ARTIFICIAL INTELLIGENCE

The concept of artificial intelligence (AI) can be defined in a variety of ways, none of which has yet gained universal acceptance as definitive. Nevertheless, a useful starting point is the Regulation laying down harmonised rules on artificial intelligence – widely known as the Artificial Intelligence Act[4] (AI Act). – which in its Article 3(1) defines an AI system as

> *a machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.*

Despite the potential weaknesses of this definition, above all its vagueness and breadth, it can still be viewed as technologically neutral (that is, not tied to any particular AI technology) and, in principle, suitable.

The two core features that contribute to the definition of AI systems are autonomy and adaptiveness: they operate with a certain degree of independence from human intervention and can learn on their own, so the system evolves as it is used. AI is therefore not a set of simple *if-then* instructions; rather, it learns from large data sets. Both features are key to AI's effectiveness, yet they also reduce the transparency with which the system converts inputs into outputs. As a direct consequence, AI systems function as a black box, and therefore it is in most cases to describe precisely how the inputs were transformed into the output.[5]

While researchers have tried to devise technical solutions that would make AI systems *self-explaining,* or at least more transparent and explainable, progress has only been partial.[6] Greater transparency can impair performance,[7] which is often an unacceptable

[4] Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act).

[5] See for example HILDT, Elisabeth. What Is the Role of Explainability in Medical Artificial Intelligence? A Case-Based Approach. *Bioengineering*. 2025, Vol. 12, Issue 4, pp. 375 and following. doi: 10.3390/bioengineering12040375.

[6] See ibid.

[7] See for example LI, Bo, QI, Peng, LIU, Bo (eds.). Trustworthy AI: From Principles to Practices. *ACM Computing Surveys*. 2023. Vol. 55, Issue 9, pp. 1–46. doi:10.1145/3555803.

trade-off, particularly where AI helps save human lives. Moreover, full transparency may never be attainable, owing both to the models' internal architecture and to the natural limitations of the human mind in processing vast amounts of data.[8]

This characteristic is crucial for understanding how AI systems work and for assessing related legal liability. Because AI systems are often adaptive and change certain properties during use, several actors can influence their operation, not only the provider, but also the deployer who uses the system and potentially other parties who supply data. In addition, the low transparency of the input-to-output process substantially limits the ability to detect specific errors in an AI system. As we show below, these peculiarities are problematic under existing tort law**,** yet we maintain that arising issues are not, in principle, insurmountable.

II. EUROPEAN DIMENSION: AN OUTLINE

The **Artificial Intelligence Act** came into effect in August 2024 as the world's most comprehensive AI regulation. It distinguishes four categories of AI systems based on their risk. Many medical AI systems pertain to high-risk AI systems (the highest category that is still permitted by the AI Act). Perhaps most importantly, this category encompasses medical devices.[9] It also contains various other AI systems (unless they do not pose a significant risk of harm to the health, safety or fundamental rights, including by not materially influencing the outcome of decision making),[10] for example those intended to be used by public authorities or on their behalf to evaluate the eligibility of natural persons for essential public assistance benefits and services (including healthcare services), to be used for risk assessment and pricing in relation to health insurance, or to evaluate emergency calls or assist in emergency healthcare triage.[11]

High-risk AI systems entail stricter obligations for the parties involved. For example, the deployer must take appropriate technical and organisational measures to ensure they use such systems in accordance with the instructions for use, assign human oversight to competent natural persons, ensure that input data is relevant and sufficiently representative, or monitor the operation of the AI system. The provider must, among other duties, establish a risk management system, train the AI system on data that meet certain quality requirements, design the AI system so that it allows for the automatic recording of events (logs) over its lifetime, ensure that the operation of the AI system is sufficiently transparent to enable deployer to interpret and appropriately use its output, design the AI system in a way that it enables effective human oversight, or meet the criteria for accuracy, robustness, and cybersecurity.

The outlined and other obligations might be important for the establishment of liability in individual cases, the failure to comply with them potentially representing the

[8] For an interesting reflection on the role of transparency of healthcare AI systems, see KOHOUTOVÁ, Marie. Building Trust Through Transparency: Regulatory Suggestions for Healthcare AI. In: ŠTURMA, Pavel (ed.). *Czech Yearbook of Public & Private International Law*, 2025, Vol. 16.

[9] See Article 6(1) in connection with Annex I of the AI Act.

[10] See Article 6(3) of the AI Act.

[11] See Annex III of the AI Act.

breach of law that can be causally linked to harm. Nevertheless, the AI Act does not directly interfere with liability regulation, much less does it establish liability rules or a liability regime.

In 2022, the European Commission published a proposal for **Directive on Adapting Non-Contractual Civil Liability Rules to Artificial Intelligence** (**Artificial Intelligence Liability Directive)**. In early 2025, however, the proposal was withdrawn.[12] This might be seen as a certain setback, even if some aspects of the proposal were bound not to be universally welcome.

The basic assumption of the proposal was that the opacity (the black-box nature) of AI systems renders fault-based liability unsuitable; however, there is not yet enough information and experience to establish a specific strict liability regime. For these reasons, the AI Liability Directive would set several rules within a fault-based liability system, and after several years, the question of strict liability would be open for discussion.

The AI Liability Directive's most significant innovation would have been to shift the burden of proof where evidence relating to the operation of high-risk AI systems is not preserved or made available. In such instances a rebuttable presumption of causation would arise between the culpable breach of a legal duty and the AI output, though not between the AI output and the damage itself.

It should be stressed that the AI Liability Directive would not extend to specific cases of strict tort liability or to contractual liability. Its practical impact on healthcare would therefore be limited, as in many European jurisdictions, including the Czech Republic, the vast majority of health services are supplied under contract.

Even if outcome-based liability may be inappropriate, any enhancement of the injured party's position would still be welcome, given that victims of AI-related harm are often at a marked disadvantage. The solution chosen in the AI Liability Directive may thus be viewed positively. Only time will tell whether European or national legislators will ultimately take a similar course.

In 2024, the new **Directive on Liability for Defective Products**[13] came into force. It clarifies that intangible things such as software fit into the scope of product. It is also noticeable that a person who substantially modified a product created by another becomes a manufacturer of the modified product. The Directive on Liability for Defective Products also applies to harm caused by a medical device, or any other product used in provision of healthcare. Nevertheless, as its name suggests, it requires a defect in the product. Furthermore, the burden of proof is born by the claimant (with certain exceptions such as the manufacturer's failure to disclose required information or the product's failure to comply with the mandatory product safety requirements). In the realm of AI applications, defects are notoriously difficult to define and even more so to prove (we analyse this problem in more detail below). Practical applicability of product liability on medical AI applications is therefore seriously limited.

---

[12] See DE LUCA, Stefano. AI liability directive. In: *Legislative Train Schedule. European Parliament* [online]. 21.5.2025 [2025-06-27]. Available at: <https://www.europarl.europa.eu/legislative-train/theme-a-europe-fit-for-the-digital-age/file-ai-liability-directive>.

[13] Directive (EU) 2024/2853 of the European Parliament and of the Council of 23 October 2024 on liability for defective products and repealing Council Directive 85/374/EEC.

III. POSSIBLE LIABILITY REGIMES

Nearly four centuries ago, Lucius Cary, Viscount Falkland, stood in the British Parliament and voiced words that have since come to epitomise the conservative approach: *'When it is not necessary to change, it is necessary not to change.'*[14] Although this maxim can hardly be adopted uncritically, we believe that, in the legal sphere, it embodies a wisdom that should not be ignored. If the law is to provide a safe space for societal development and for the full realisation of the individual under conditions of legal certainty, it ought not introduce changes that are unnecessary, or worse, changes made merely for change's sake. The key question, therefore, is whether tort law can adequately address liability arising from the use of AI systems in medicine through the legal concepts and institutions already at its disposal.

III.1. Liability for Non-Compliance with the Standard of Care

A healthcare provider is generally not liable for the outcome of treatment; rather, liability is tied to the correctness of the procedure, i.e., compliance with the applicable standard of care in the given field.

Czech law defines that standard of care (commonly referred to as *lex artis*) in Section 4(5) of Act No. 372/2011 Sb., on Health Services and Conditions of Their Provision.[15] It consists of three cumulative elements:

1. Conformity with scientific knowledge and recognised medical practices;
2. Respect for the individuality of the patient; and
3. Consideration of the specific circumstances and objective possibilities under which the care is provided.

All three components must be met for the provider to have complied with the statutory standard of care.

III.1.1 Conformity with Scientific Knowledge and Recognised Medical Practices

Recognised procedures are, above all, those set out in guidelines, clinical algorithms, the scientific literature, and comparable professional sources. Even where no such precisely defined procedures exist, the standard of care may still be met if the course taken accords with more general scientific principles, that is, if it rests on sound scientific reasoning.[16]

Accordingly, the medical profession itself largely determines the content of the standard of care, and the same will hold true for medical applications of AI. At present, guidelines in this field are generally lacking. Yet as AI solutions become increasingly effective, their use can be expected to form part of the standard of care ever more widely and to

[14] See *Oxford Reference* [online]. [2025-06-30]. Available at: <https://www.oxfordreference.com/display/10.1093/acref/9780191826719.001.0001/q-oro-ed4-00004254>.

[15] This definition is also referenced in Section 2643 of the Civil Code that establishes the duty to act in accordance with the care of a diligent professional in the context of contract for health care.

[16] We address the question of proper scientific justification in more detail below in the context of evaluating unestablished medical methods.

appear in speciality-specific recommendations. In step with these developments, correct practice in the use of AI will acquire an increasingly concrete shape, and breaches of the standard of care will become easier to identify.

A crucial distinction must be drawn between AI systems that are certified as medical devices and those that are not. AI systems may obtain certification as medical devices, typically in the medium-risk classes IIa or IIb. Since 2024, several Czech AI systems have been certified under the EU Medical Device Regulation (MDR),[17] including Carebot, Kardi-AI, or Aireen.[18] The status of a medical device may attach both to embodied AI and to pure software (e.g., diagnostic support tools or clinical decision support systems). The manner in which a medical device is used will ordinarily follow, first and foremost, from its documentation, including the instructions for use.[19] It must then be assessed in the broader context of each indicated medical procedure. The MDR certification process appears fully adequate for AI systems; their specific characteristics have so far not required special legal rules.[20]

There are, however, AI systems that are not certified medical devices yet are nonetheless used in medical practice. Some are geared specifically towards healthcare, whereas others are entirely general-purpose.

General-purpose AI systems include widely available large language models and their platforms, such as ChatGPT, Claude, or Gemini. Their use can be described as **trivial** in the sense that it does not constitute a medical method. A *method* may be understood as a circumscribed procedure consisting of a sequence of defined steps.[21] If, for instance, a doctor thinking through a case consults ChatGPT on a mobile phone – alongside professional databases and an internet search engine – such use can hardly be regarded as part of any method, nor does the standard of care apply to it as such.

In essence, trivial use of AI does not differ from any other use of an information source employed in forming clinical judgement. It may be approached analogously to *consulting* the literature or an online search engine. While not required *by the standard of care*, it is neither prohibited nor inappropriate; it is permissible, often highly useful, yet potentially risky if the physician misjudges the reliability of the information obtained, its relevance to the specific case or other material factors. The full responsibility for handling such sources rests with the physician (or, where applicable, the healthcare provider as employer). The resulting clinical decision is the doctor's alone, and the law affords no

[17] Regulation (EU) 2017/745 of the European Parliament and of the Council of 5 April 2017 on medical devices, amending Directive 2001/83/EC, Regulation (EC) No 178/2002, and Regulation (EC) No 1223/2009 and repealing Council Directives 90/385/EEC and 93/42/EEC.

[18] See Nová expertní skupina pomáhá zavést AI do české medicíny [A new expert group helps introduce AI into Czech medicine]. In: *Medical Tribune* [online]. 22.1.2025 [2025.06-27]. Available at: <https://www.tribune.cz/medisekce/digitalni-medicina/ai-v-medicine/nova-expertni-skupina-pomaha-zavest-ai-do-ceske-mediciny-medi/#>.

[19] See the relevant regulation in the Medical Device Regulation and Czech Act No. 375/2022 Sb., on Medical Devices and In Vitro Diagnostic Medical Devices.

[20] Srov. ZIKMUNDOVÁ, Klára. Artificial Intelligence and Medical Devices: Do We Need New Regulation? *Časopis pro právní vědu a praxi*. 2023, Vol. 31, Issue 2, p. 390.

[21] We return to the definition of a method in more detail below in the context of liability for unestablished methods.

protection against the consequences of errors in the source used, much as in the case of off-label medication.

Other AI systems, though uncertified, are nevertheless designed for medical use. Examples include MedLM by Google, GatorTron (developed in collaboration between the University of Florida and NVIDIA), or Microsoft's BioGPT. Some models are even more narrowly tailored to particular fields, such as MGH Radiology Llama or Microsoft's LLaVA-Med, which converts visual inputs into textual outputs. Their use may no longer be trivial. It can be integrated into an existing procedure or method. In that event the standard of care may attach to such use, governing the very indication for employing the system, the way it is operated, and the evaluation of its outputs.

Within the standard of care, correctness may be assessed at both the input stage (e.g., supplying complete and accurate data, prompt engineering) and the output stage (verifying and working with the results). Defining the concrete standard of care for different specialties and scenarios is, for the future, a task for the medical profession through self-regulatory documents, guidelines, clinical algorithms, and the like.

III.1.2 Respect for the Individuality of the Patient

This requirement concerns the unique needs of the individual patient at the biological level and, in a broader sense, at the psychological, social and, where appropriate, spiritual levels. Every person's organism differs from the average and from that of other patients; likewise, each person has their own life values and goals. Blind adherence to guidelines without regard to an individual patient's needs may in fact constitute a breach of the standard of care.

Medical AI systems could, in this respect, help to achieve a more finely tuned individualisation of care and move medicine towards a genuinely personalised form. By processing extensive data on a particular patient – possibly in conjunction with telemedicine applications (for example, through the remote transmission of data from wearable devices) – it would be possible to obtain far more information about the patient's unique characteristics and illness, including their response to various therapeutic modalities.

At the same time, one must bear in mind the risk inherent in AI systems: a tendency to prioritise average results, which naturally constitute the largest part of the datasets on which the AI is trained. This is an inherent limitation of current AI systems. Future changes in system architecture may overcome it, but for now it remains essential that AI outputs in clinical practice serve only as a basis for the judgement of a human physician.

III.1.3 Consideration of the Specific Circumstances and Objective Possibilities under which the Care is Provided

According to a general maxim, *no one is bound to do the impossible.*[22] If a provider complies with all duties concerning, inter alia, the staffing, and the technical and material

[22] *'Ad impossibilia nemo tenetur.'* See for example ŠUSTEK, Petr. Concluding Remarks. In M. Foglia – T. Holčapek – M. Šolc – P. Šustek (eds.). *Crisis Patient Prioritisation and the Law: the Pandemic Experience.* Kluwer Law International, 2024, p. 252.

resources of its healthcare facility,[23] they cannot be held liable for failing to do something they objectively could not do – for example, for not using a diagnostic device that they neither possessed nor were required to possess.

The use of AI is not at present a legal obligation for healthcare providers. Where a provider does not have AI tools at its disposal, this cannot be regarded as a breach of the standard of care. As AI systems are likely to be incorporated into professional recommendations and thus come to form part of recognised medical practice, a point will approach at which, in certain circumstances, failure to employ AI will have to be judged as an incorrect professional procedure (save where objective factors make use of the system impossible, for instance, in an emergency outside a healthcare facility).

III.2 Specific Regimes of Strict Liability

Even where the physician's conduct has not breached the standard of care, certain special strict-liability regimes may still apply. Some of these regimes permit exoneration, typically on the ground that the defendant did not neglect proper supervision. In such cases the burden of proof is effectively reversed; it is not the claimant who must prove the defendant's unlawful conduct, but the defendant who must, if necessary, demonstrate that the required standard of behaviour was not breached.

Below we examine the strict-liability regimes most likely to be applied to AI-related harm under Czech civil law, while noting that broadly similar approaches are available in many other legal systems.

III.2.1 Liability for Damage Caused by Operational Activity

Liability for damage arising from operational activity is defined in Section 2924 of the Civil Code as follows:

> *Whoever operates an establishment serving a gainful purpose shall compensate the damage resulting from the operation, whether the damage was caused by the operational activity itself, by a thing used in that activity, or by the impact of the activity on the surroundings. The obligation is discharged if the operator proves that they exercised all care that can reasonably be required to prevent the damage.*

The extent to which this strict-liability regime applies in healthcare is not entirely settled. On the one hand, its application is not prevented by the fact that most health services in the Czech Republic are provided under the public healt insurance scheme. As the explanatory report to the Civil Code emphasises, liability under Section 2924 covers all establishments serving a gainful purpose, including those whose "profit" consists mainly of income from public budgets (for example, public hospitals).[24]

On the other hand, case law dating back to the 1990s holds that liability for damage arising from operational activity does not extend to the actual practice of medicine. The

[23] Decree of the Ministry of Health No. 92/2012 Sb., on the Requirements for the Minimum Technical and Material Equipment of Healthcare Facilities and Contact Facilities, and Decree of the Ministry of Health No. 99/2012 Sb., on the Requirements for Minimum Staffing of Healthcare Facilities.

[24] See Explanatory Report to the Civil Code, Special Part, To Sections 2924 2926.

reason lies chiefly in the special nature of medical science, which is inherently bound up with risk, technical difficulty, and a high complexity of the phenomena involved. These considerations, taken together with the physician's duty to perform interventions, led the courts to exclude strict liability from the domain of medical practice.[25] The use of an AI system in a physician's clinical decision-making process is unquestionably part of that practice.

The outlined judicial conclusion was reached under the former Civil Code. Commentary is not unanimous as to whether, and to what extent, it should also apply under the new Civil Code, which entered into force in 2014. Some authors argue that legislative changes to the relevant liability regimes, particularly the grounds for exoneration, render the earlier decisions obsolete.[26] A substantial part of legal scholarship, however, still supports excluding medical interventions from liability for damage arising from operational activity.[27] From the practical standpoint, the courts even today do not impose such broad strict liability in healthcare.

Importantly, liability for damage arising from operational activity covers not only damage caused by the operational activity itself (in healthcare, damage arising from running the facility as such, typically health injury due to a nosocomial infection),[28] but also damage caused by a thing used in that activity.

In this respect a careful distinction must be drawn between cases appropriately dealt with under liability for operational activity and cases better addressed through other strict-liability regimes, notably liability for damage caused by a defective thing (Section 2936) and liability for damage caused by a thing in and of itself (Section 2937(1)). From a teleological viewpoint these separate regimes are apt where the thing is used directly in providing the health service. A thing (whether a scalpel, a probe, or software) is employed in virtually every medical intervention. If liability for operational activity could be inferred merely because a thing was used, the case law outlined above would be undone, for the regime would then apply to practically all medical practice. Liability for operational activity should therefore be reserved for things used in running the healthcare facility (in the AI context, for example, a system

[25] See the Regional Court in Hradec Králové judgment of 17 September 1997, file no. 25 Co 167/97, and for example also the Supreme Court of the Czech Republic judgment of 31 August 2004, file no. 25 Cdo 2542/2003 or the Supreme Court of the Czech Republic judgment of 31 March 2010, file no. 25 Cdo 4758/2008. See also ŠUSTEK, Petr, HOLČAPEK Tomáš. Odpovědnost obecná a zvláštní. [General and Special Liability.] In: P. Šustek – T. Holčapek (eds.). *Zdravotnické právo. [Medical Law.]* Praha: Wolters Kluwer, 2016, p. 305.

[26] See BEZOUŠKA, Petr. Commentary to Section 2924. In: M. Hulmák (ed.). *Občanský zákoník VI. Závazkové právo. Zvláštní část (§ 2055-3014). Komentář. [Civil Code VI. Law of Obligations. Special Part (Sections 2055-3014). Commentary.]* 1st ed. Praha: C. H. Beck, 2014, pp. 1601–1602.

[27] See ŠUSTEK, Petr, HOLČAPEK Tomáš. Odpovědnost obecná a zvláštní. [General and Special Liability.], pp. 304–307; PAŠEK, Martin. Commentary to Section 2924. In: J. Petrov – M. Výtisk – V. Beran (eds.). *Občanský zákoník. Komentář. [Civil Code. Commentary.]* 2nd ed. Praha: C. H. Beck, 2019, p. 3053; VOJTEK, Petr. Commentary to Section 2924. In: J. Švestka – J. Dvořák – J. Fiala (eds.). *Občanský zákoník. Komentář. Svazek VI. [Civil Code. Commentary. Volume VI.]* 2nd ed. Praha: Wolters Kluwer, 2019, p. 973; ŠOLC, Martin. *Právo, etika a kmenové buňky. [Law, Ethics and Stem Cells.]* Praha: Wolters Kluwer, 2018, pp. 211–212.

[28] Srov. ŠUSTEK, Petr, HOLČAPEK Tomáš. Odpovědnost obecná a zvláštní. [General and Special Liability.], pp. 306–307.

performing logistics or purely administrative tasks).[29] Software employed directly to support diagnosis, therapeutic decision-making, or other clinical purposes should as a rule not fall under that regime.

The case law excluding medical practice from liability for damage arising from operational activity likewise applies to a similar strict-liability regime for damage caused by an **especially dangerous operation** (Section 2925), that is, an operation in which the possibility of serious damage cannot reasonably be ruled out even with proper care. Notably, this regime is conceived as absolute strict liability, without any possibility of exoneration. Even if one sought to characterise AI systems, by reason of their opacity, as an especially dangerous operation, the result would have little practical impact in healthcare. Moreover, classing AI systems as an especially dangerous operation would be unduly harsh and would not reflect their actual use. First, it must be recalled that in clinical processes the AI system does not make the final decision; it informs a human who then works with, or at least confirms, its output. This interpretation and application of AI outputs leaves room to correct potential errors, so it is difficult to maintain that serious damage cannot be avoided even with proper care when using AI. Secondly, the notion of "serious damage" in Section 2925 is tied more to large-scale harm, which typically is not at stake in the medical use of AI.

III.2.2 Liability for Damage Caused by a Defective Thing

Liability for damage caused by a thing is, under Section 2936 of the Civil Code, confined to harm resulting from the use of a defective thing in the performance of an obligation.[30] The tortfeasor is the person who employed the defective thing while discharging an obligation towards another. This is a form of strict liability without any possibility of exoneration: if its preconditions are met, the tortfeasor must provide compensation.

The term *defect* is not defined for this purpose. One might, by analogy, look to the definition used in the liability regime for damage caused by a defective product in Section 2941, while remembering that product-defect liability governs a different relationship, arising only where the product is used directly by a consumer.[31] In the context of harm caused by AI in medicine, the Section 2941 definition could become relevant chiefly for medical devices used directly by the consumer in telemedicine. Where the AI is employed by a healthcare provider in performing its obligations, that definition serves only as an interpretative guide.[32] Similar analogies could be drawn from the definitions of defects in the subject-matter of performance (Section 1916), in the

---

[29] See ŠOLC, Martin. *Občanskoprávní odpovědnost za umělou inteligenci v rámci klinického rozhodování v medicíně [Civil Liability for Artificial Intelligence in the Context of Clinical Decision-Making in Medicine]*. In: J. Suchoža – J. Husár – R. Hučková (eds.). *Právo, obchod, ekonomika IX*. Košice: Univerzita Pavla Jozefa Šafárika v Košiciach, 2019, p. 545.

[30] Section 2936 of the Civil Code: *'Whoever is obliged to perform something for another and uses a defective item in doing so shall compensate for the damage caused by the defect. This also applies in the case of the provision of healthcare, social, veterinary, and other biological services.'*

[31] See VOJTEK, Petr. Commentary to Section 2936. In: J. Švestka – J. Dvořák – J. Fiala (eds.). *Občanský zákoník. Komentář. Svazek VI. [Civil Code. Commentary. Volume VI.]*, 2019, p. 1009.

[32] See ŠOLC, Martin. *Občanskoprávní odpovědnost za umělou inteligenci v rámci klinického rozhodování v medicíně [Civil Liability for Artificial Intelligence in the Context of Clinical Decision-Making in Medicine]*, p. 546.

subject-matter of sale (Section 2099 in conjunction with Sections 2095–2096), and in the quality of the thing sold (Section 2161), always bearing in mind the distinct purpose of each provision.[33]

Legal doctrine attempts a general definition: a defect is a deficiency in the thing, manifested by malfunction, faulty operation, breakdown, or other difficulties compared with what is normally expected to be trouble-free operation.[34]

Whether the software in question has been certified as a medical device may influence the applicability of Section 2936. For a medical device, a defect may lie in a departure of its actual performance from the technical parameters recorded at registration. Clear malfunctions that demonstrably reduce the reliability or usability of the AI's outputs, however, tend to be rare. One can envisage, for instance, sudden changes in system behaviour that markedly diminishes output credibility; such temporal fluctuations may reflect the system's adaptiveness and complexity. Comparable failures, amounting to defects, may also occur in AI not certified as a medical device if they demonstrably undermine its trustworthiness. Proving a defect is sometimes easier with certified devices because there is a formal description of fault-free operation against which performance at the material time can be compared. In either case, the situations sketched above are unlikely to be the typical cause of patient harm.

For completeness, a defect might theoretically consist in erroneous or incomplete data within the database from which the AI system draws. Yet with the very large datasets used by major models such defects are improbable (not least because errors tend to be statistical outliers and are outweighed by less faulty data). Demonstrating such an error – and, above all, a causal link to the harm – is realistically feasible only in rare instances.

In practice, objectively inaccurate outputs arise far more often from minor errors during algorithmic processing of inputs. In particular, large language models should be viewed in light of their training: they are taught to produce outputs that humans regard as valuable and plausible. They optimise for persuasiveness rather than for maximal truthfulness. This can prove problematic, because errors in AI outputs are often harder for a professional to spot than errors made by a human.[35] Understanding these problematic aspects of AI will undoubtedly become part of the standard of care in their use and should therefore be included, inter alia, in medical education.[36]

A certain error rate – including the tendency to present some incorrect outputs with undue confidence (popularly called *hallucinations*) – is an inherent feature of AI

[33] See PAŠEK, Martin. Commentary to Section 2936. In: J. Petrov – M. Výtisk – V. Beran (eds.). *Občanský zákoník. Komentář. [Civil Code. Commentary.]*, p. 3072.

[34] See VOJTEK, Petr. Commentary to Section 2936, p. 1009.

[35] See explanation of this phenomenon in the context of mathematics by Professor Terence Tao. The outlined problems do not only apply to mathematics but to various fields where large language models can be used. FRIDMAN, Lex. Terence Tao: Hardest Problems in Mathematics, Physics & the Future of AI. In: *Lex Fridman Podcast #472* [online]. Time: 01:49:12. 14.6.2025 [2025-06-30]. Available at: <https://lexfridman.com/terence-tao-transcript>.

[36] See a prediction of the expansion of medical education to include, i.a., relevant areas of computer and data science in TOPOL, Eric. *Deep Medicine. How Artificial Intelligence Can Make Healthcare Human Again*. New York: Basic Books, p. 308.

systems. Such error rates are not departures from ordinary expectations and therefore do not constitute defects in the legal sense (just as a scalpel's ability to damage surrounding tissue, or its inability to cut bone, is not a defect; nor is a car's non-zero braking distance).

The possibility of objectively incorrect output is thus an internal risk of using an AI system. A rough analogy can be drawn with medicinal products (drugs) which, unlike most existing medical devices, act in the body independently of any direct human control; their efficacy and safety depend on multiple interacting factors, making them in a sense *autonomous* agents. A system's intrinsic error rate is, in this respect, akin to a drug's side-effects. The risks linked to pharmaceuticals cannot be eliminated entirely, and where harm occurs despite correct medical procedure, liability is generally considered only in cases of off-label use. It is noteworthy that one might regard the use of an AI system that is not a medical device as a kind of analogue to off-label drug administration, potentially attracting expanded liability.

It is important to remember that AI systems operate on probabilities, and that probabilistic reasoning is characteristic of every intellectual process, whether carried out by a machine or a human. When formulating a diagnosis or treatment plan, a physician never achieves one-hundred-per-cent certainty; rather, consciously or subconsciously, they reach a qualified estimate of the likelihood of success for various options. Hence liability in medicine is typically linked not to the outcome, but to the correctness of the procedure.[37] A diagnosis later shown to be objectively wrong does not in itself establish liability, unless the physician breached the standard of care when making it.[38]

The injured party's evidential position will be all the harder because defects in an AI system's operation can usually be inferred only from an incorrect output: the input-to-output process itself is opaque, a direct consequence of AI's black-box nature. Proving that a defect occurred in that process, as opposed to the materialisation of the inherent risk of a wrong output, will therefore be very difficult, even if one can define *defect* clearly enough in this context.

There might be an objection that, for certain hazardous activities, civil law creates special strict-liability regimes. If, earlier, we said that a normal braking distance is not a defect, its associated risk is nevertheless inherent in the operation of motor vehicles, and is addressed by strict liability for damage arising from the operation of transport vehicles under Sections 2927 ff., from which exoneration is impossible where the damage originated in the vehicle's operation (for example owing to its braking distance). In medical AI the closest analogue would be liability for damage arising from operational activity – but, as noted, medical practice itself (including the use of things during medical procedures) has been excluded from that liability regime by case law.

---

[37] See HOLČAPEK, Tomáš. Odpovědnost zpravidla za činnost, nikoli za výsledek. [Liability is Generally for the Activity, not for the Outcome.] In: P. Šustek – T. Holčapek (eds.). *Zdravotnické právo. [Medical Law.]*, pp. 299–300.

[38] See ŠUSTEK, Petr, JIRÁSKOVÁ, Kateřina. Náležitý odborný postup (lex artis). Obecně. [Appropriate Professional Conduct. In General.] In: P. Šustek – T. Holčapek (eds.). *Zdravotnické právo. [Medical Law.]*, pp. 266–267.

III.2.3 Liability for Damage Caused by a Defective Product

This liability regime is set out in Section 2939(1) of the Civil Code:

> *Damage caused by a defect in a movable thing intended to be placed on the market as a product for the purposes of sale, lease or other use shall be compensated by the person who manufactured, extracted, grew or otherwise acquired the product or its component, and, jointly and severally with that person, also by anyone who has marked the product or its component with their name, trade mark or other designation.*

In this context the new European Directive on Liability for Defective Products (see above) is highly relevant. However, as with liability under Section 2936, the practical application of liability for damage caused by a defective product will largely be constrained by the difficulty of proving a defect in the AI system employed.

III.2.4 Liability for Damage Caused by a Thing in and of Itself

Liability for damage caused by a thing *in and of itself* is governed by Section 2937(1) of the Civil Code: *'If a thing causes damage in and of itself, the person who ought to have supervised the thing shall compensate the damage; if no such person can otherwise be identified, the owner of the thing is deemed to be that person.'* The alleged tortfeasor may avoid liability by proving that proper supervision of the thing was not neglected.[39]

Determining who *ought to have supervised* an AI system is not always straightforward. In most situations that role rests with the healthcare provider. Where the system is cloud-based and the provider accesses servers it does not control, part of the necessary supervision may also lie with the manufacturer or provider of the software platform.

A thing caused damage in and of itself if the harm arose from the thing's internal characteristics: its shape, material, strength, structure, construction, and so forth.[40] Liability still applies even if an external influence acts on the thing (for physical objects, temperature, or weather, for example), provided that influence does not consist in a direct act of a fault-capable person,[41] nor an act that merely triggers a dangerous internal property.[42]

Damage can occur while the thing is in use, but only where the harm flows from the thing's internal property, not from the manner of its use. Section 2937(1) protects against hazards inherent in the thing itself, not against human action.[43] Liability is therefore excluded if the manner in which the thing was used led to the risk materialising.

---

[39] According to some authors, Section 2937 establishes fault-based liability (BEZOUŠKA, Petr. Commentary to Section 2937. In: M. Hulmák (ed.). *Občanský zákoník VI. Závazkové právo. Zvláštní část (§§ 2055-3014). Komentář. [Civil Code VI. Law of Obligations. Special Part (Sections 2055-3014). Commentary.]*, p. 1640). However, in line with the majority of scholarly literature (VOJTEK, Petr. Commentary to Section 2937. In: J. Švestka – J. Dvořák – J. Fiala (eds.). *Občanský zákoník. Komentář. Svazek VI. [Civil Code. Commentary. Volume VI.]*, p. 1011; PAŠEK, Martin. Commentary to Section 2937. In: J. Petrov – M. Výtisk – V. Beran (eds.). *Občanský zákoník. Komentář. [Civil Code. Commentary.]*, pp. 3073-3074), we interpret it as a case of strict liability.

[40] See PAŠEK, Martin. Commentary to Section 2937, p. 3074.

[41] See MELZER, F. Commentary to Section 2937. In: F. Melzer – P. Tégl (eds.). *Občanský zákoník – velký komentář. Svazek IX. [Civil Code – Large Commentary. Volume IX.]* Praha: Leges, 2018, p. 717.

[42] See PAŠEK, Martin. Commentary to Section 2937, p. 3074.

[43] See BEZOUŠKA, Petr. Commentary to Section 2937, p. 1640.

Czech scholarship shows two views. Some authors believe this liability regime applies even when the thing is used in treating the patient, though the tortfeasor may still exonerate themselves.[44] Others argue that it only applies where the thing was not used for patient care yet still affected the patient.[45] This latter view cannot be pressed too far. The regime clearly covers, for instance, an X-ray tube head that falls on a patient during an examination, while the device is in use. The decisive question is whether the harm arose from the manner of use; if so, Section 2937(1) does not apply.

A typical AI-related injury – an incorrect AI output – fits Section 2937(1) very well. The harm stems from the AI system's internal risk; the system was indeed being used, but the harm did not result from an erroneous manner of use. The essential distinction is between (i) cases where harm results from an unavoidable materialisation of the AI system's risk and (ii) cases where it flows from the physician's mishandling of the system or its outputs. A practical, even if in some cases non-trivial, rule is: If, given the physician's qualifications and role, the incorrect output was recognisably wrong, the matter concerns a breach of the standard of care; only where the error was reasonably undetectable within the physician's expertise does strict liability under Section 2937(1) arise.

As AI systems gain greater autonomy (particularly autonomous robotic systems or semi-autonomous procedures) this liability regime will likely be invoked more often.

What constitutes proper supervision upon which exoneration is conditional? Generally speaking, proper supervision denotes the level of care reasonably expected of a person acting with due professional diligence, taking into account the nature of the thing used and the circumstances.[46] For AI systems this typically includes:

- routine hardware maintenance;
- regular software updates;
- adequate cyber-security measures; and
- other safeguards proportionate to the system's purpose.[47]

An important concretisation is to be found in obligations imposed on a deployer of high-risk AI systems by the AI Act. For healthcare providers, obligations of high-risk AI systems deployers will be especially relevant (such as taking appropriate technical and organisational measures to ensure the use is compliant with the accompanying instructions, assigning human oversight, ensuring that the input data are relevant and sufficiently representative, monitoring the operation of the high-risk AI system, etc.).[48] Full adherence to these obligations may demonstrate due care and thus allow the deployer to exonerate themselves from strict liability.

---

[44] See VOJTEK, Petr. Dvě otázky medicínského práva, pro něž bude nový občanský zákoník přelomový. [Two Questions of Medical Law for Which the New Civil Code Will Be Crucial.] *Soudní rozhledy.* 2013, Vol. 19, Issue 4, pp. 122–126.

[45] See ŠUSTEK, Petr, HOLČAPEK Tomáš. Odpovědnost obecná a zvláštní. [General and Special Liability.], pp. 308–309.

[46] See PAŠEK, Martin. Commentary to Section 2937, pp. 3074–3075.

[47] See ŠOLC, Martin. *Občanskoprávní odpovědnost za umělou inteligenci v rámci klinického rozhodování v medicíně [Civil Liability for Artificial Intelligence in the Context of Clinical Decision-Making in Medicine]*, p. 548.

[48] See Article 26 of the AI Act and the sub-chapter on EU law above in this paper.

III.3 Summary: Three Typical Scenarios

On the basis of the foregoing, we can identify three typical scenarios in which, in principle, the application of three liability regimes may be considered:

Scenario	Liability regime	Pitfalls	Expected frequency
Harm arising from a defect in the AI system used	Liability for damage caused by a defective thing (Section 2936)	• Difficult to define a defect in the AI context (hallucinations are inherent features, not defects) • Very difficult to prove a defect because AI systems operate as a black box	Low frequency
Harm arising from internal causes of the AI system, unrelated to the way it was used	Liability for damage caused by a thing in and of itself (Section 2937 (1))	• In some cases, hard to distinguish from harm caused by incorrect use of the AI system • Possible arguments against applying the regime when the AI system is used directly in providing care • The precise scope of proper supervision as a ground for exoneration is unclear	Potentially medium frequency, especially for systems with a higher degree of autonomy
Harm arising from a breach of the standard of care in using the AI system (choice of indication, entering patient-specific inputs, handling outputs, etc.)	Liability for breach of the standard of care	• The standard of care for working with AI systems in healthcare will need to be defined clearly enough • The frequency of application may decline as AI systems achieve greater autonomy	Potentially high

III.4 Liability for New Medical Methods Involving AI

Conceptually, the introduction of AI into medicine can also be approached as the adoption of new methods. Framing it this way may help to overcome certain civil-liability

issues that arise when harm is caused primarily by *how* the AI is used (so none of the strict-liability regimes discussed above apply) yet an established standard of care has not yet emerged. The chief difficulties lie in distinguishing the AI system itself from the method in which it is embedded and in defining a concrete standard of care for that method.

Unlike the highly detailed statutory definitions of a medicinal product or a medical device, the term *method* is neither formally nor consensually defined. In general, it can be regarded as a particular act, or sequence of acts, underpinned by specific *know-how*.[49] It is a non-trivial, structured, concrete procedure comprising a series of systematically linked steps.

Accordingly, software alone is not a method; the method is the formalised manner in which the software is employed. Typical, though not exclusive, AI use-cases include:

- Interventional (e.g., surgical) procedures assisted by robotics that incorporate AI.
  - This represents an intermediate level of autonomy between surgical robots that are fully or almost fully operator-controlled and fully autonomous robotic systems, which are not yet routine in clinical practice (though successful cases exist, for example, the *Robodoc* system, which has been performing pre-programmed tasks in orthopaedic surgery since 1992,[50] the *Smart Tissue Autonomous Robot* (STAR), which has successfully carried out several experimental laparoscopic intestinal anastomoses on pigs,[51] and a robotic system SRT-H that performed a large part of a gallbladder removal on deceased pigs, the most complex autonomous surgery to date, in 2025[52]). Using a surgical robot is, like other surgical or interventional procedures, a medical method. The German Federal Court of Justice has said as much, although AI systems are medical devices in themselves, the operative deployment of Robodoc was judged a new method.[53]
- Imaging enhancement during interventional procedures in internal specialties.
  - Contemporary medicine increasingly replaces many surgical operations with less invasive, safer interventional procedures performed by specialists such as interventional cardiologists or gastroenterologists. Artificial intelligence, especially when embedded in existing software, can improve the physician's real-time visualisation of the procedure, thereby increasing efficacy and avoiding patient harm. In this context AI is one of the physician's tools, used within a specific, relatively well-defined method.

[49] See ŠUSTEK, Petr. Ověřování nezavedené metody. [Evaluation of an Unestablished Method.] In: P. Šustek – T. Holčapek (eds.). *Zdravotnické právo. [Medical Law.]*, p. 721.

[50] See FAN, Xuanze, WANG, Yan, ZHANG, Shouwei (eds.). Orthopedic surgical robotic systems in knee arthroplasty: a comprehensive review. *Frontiers in Bioengineering and Biotechnology*. 2025, Vol. 13. doi: 10.3389/fbioe.2025.1523631.

[51] See GRAHAM, Catherine. Robot Performs First Laparoscopic Surgery Without Human Help. In: *John Hopkins University. The Hub* [online]. 26.1.2022 [2025-06-28]. Available at: <https://hub.jhu.edu/2022/01/26/star-robot-performs-intestinal-surgery/>.

[52] See KIM, Ji Woong, CHEN, Juo-Tung, HANSEN, Pascal, KRIEGER, Alex (eds.). SRT-H: A Hierarchical Framework for Autonomous Surgery via Language-Conditioned Imitation Learning. *Science Robotics*. 2025, Vol. 10, Issue 104. doi: 10.1126/scirobotics.adt5254.

[53] Federal Court of Justice (BGH) decision of 13 June 2006, file no. VI ZR 323/04 (NJW 2006, 2477).

- Diagnostic methods.
  - At present AI is mainly employed to improve the accuracy of interpreting imaging outputs (e.g., CT or MRI). Interpretation, however, remains the task and responsibility of the physician, as does the prior indication of the diagnostic procedure of which the AI forms part. In the future, AI will probably play a larger role in the broader diagnostic reasoning process, offering suggested diagnoses. Even then it will still constitute an element of the diagnostic procedure, albeit one whose steps and sequence may not always be sharply delineated, so it cannot invariably be labelled a concrete method.

Methods that are already part of routine clinical practice are typically described and technically elaborated in guidelines and clinical algorithms. This is not invariably the case. A method may become part of practice even without being expressly set out in professional recommendations. In such cases the standard of care emerges organically – from practical experience and the prevailing views of specialists in the field – even though, strictly speaking, Czech law requires that standard to be reasonably clear from the outset, which in reality it seldom is.

Under Czech law, specifically Act No. 373/2011 Sb., on Specific Health Services, a method that has not yet been introduced into clinical practice may be carried out solely within the formal regime for the evaluation of unestablished methods, following approval by the Ministry of Health. If the evaluation succeeds, the Ministry subsequently designates the method as standard. From a civil liability perspective, this statutory regime has the advantage of a study protocol approved by the Ministry, which defines the standard of care for performing the method; conformity with the protocol can therefore serve as a valuable legal defence for the healthcare provider.

AI will not always render a procedure a new method. The Act exempts modifications of established methods from the evaluation requirement, without defining that term; it merely states that the modification must not have adverse effects on the patient's state of health. Where doubt arises as to whether a procedure is a new method or merely a modification, the Ministry of Health decides. Scholarly writing suggests that a method is *new* when two conditions are met cumulatively: (i) the risk introduced by the modification is sufficiently concrete and foreseeable (i.e., not purely hypothetical) and (ii) the risk has a non-negligible intensity, assessed case by case by combining its seriousness with the probability of occurrence.[54]

In many instances, incorporating AI into an existing procedure will not transform that procedure into a *new* method but merely modify it. An example might consist in embedding AI in diagnostic software to enhance performance without adding extra risk. By contrast, interventions that rely on robotics or on more intensive clinical-decision support may well amount to an unestablished method. Wherever the statutory hallmarks of an unestablished method are met, formal evaluation should take place.

In practice, formal evaluation is rare for a mixture of substantive and administrative reasons. Since Act No. 373/2011 Sb. entered into force in 2012, the Ministry of Health has

[54] See ŠOLC, Martin. *Nové metody v medicíně a právo. [New Methods in Medicine and the Law.]* Praha: Wolters Kluwer, 2022, pp. 65–67.

approved only four evaluation projects, none of which directly involved AI. Most new methods are therefore developed outside this statutory framework, typically under the approval of the local research-ethics committee.

Defining the standard of care is harder in such cases but remains indispensable in litigation. At a minimum, the method must be properly scientifically justified. Acceptable justification may include experience with a comparable method, even at another institution, or at least sound scientific knowledge of the pathological processes involved and of how they may be influenced. If the method represents the patient's only realistic hope of maintaining or improving health, this fact may be taken into account in assessing the scientific justification, but it cannot dispense with the requirement altogether.[55]

Civil liability for new methods is not the subject of a special statute, so suitable existing liability regimes must be applied. To derive strict liability by analogy would be questionable. Section 2895 of the Civil Code expressly provides that strict liability arises only where a statute so directs. The rationale is clear: strict liability places a heavy burden on the tortfeasor; even with a possibility of exoneration, it effectively reverses the burden of proof regarding compliance with the required standard of care and should be imposed only for good reason. Imposing outcome-based liability on those who pioneer new methods could chill innovation and would arguably be inequitable.

A more appropriate solution is to rely on fault-based liability for unlawful conduct, typically, a breach of the emerging standard of care. Such breaches might consist in deviating from the approved study protocol or in applying the new method without a sound scientific foundation or without following the necessary procedural safeguards.

IV. AN ELECTRONIC PERSON AS A POSSIBLE WAY FORWARD?

Some academic authors, and even certain official documents,[56] have suggested creating a new category of legally recognised person that would exist alongside natural and legal persons, that would possess legal capacity, and could itself bear legal liability. For convenience this hypothetical new category is referred to here as an *electronic person*.

The idea aligns more closely with the looser notion of *person* in the common-law world, where several types of juridical person are already recognised. Continental legal systems, by contrast, draw a sharper line between natural persons, legal persons, and things, although a few jurisdictions acknowledge sui generis entities that fall outside those traditional classes (for example, living animals or parts of the human body). Given the scale at which AI may eventually interact with society, one could imagine a material need for a new legal category even within a system as strongly classificatory as continental law.

---

[55] See ibid., pp. 201–202.

[56] See the European Parliament resolution of 16 February 2017 with recommendations to the Commission on Civil Law Rules on Robotics, which in its Article 59 calls the Commission *'to explore, analyse and consider the implications of all possible legal solutions, such as (…) creating a specific legal status for robots in the long run, so that at least the most sophisticated autonomous robots could be established as having the status of electronic persons responsible for making good any damage they may cause, and possibly applying electronic personality to cases where robots make autonomous decisions or otherwise interact with third parties independently.'*

Yet this raises numerous thorny questions. To begin with, it is not obvious how the legal personality of an electronic person should be constructed. Even with so venerable an institution as the legal person, jurists disagree on its true nature: under the *real entity theory*, a legal person is a community of people that truly exists and merely receives recognition from the law, whereas the *fiction theory* regards it simply as a statutory construct.[57] The distinction is not purely academic, under the fiction theory a legal person can never act in its own right and must always be represented (typically by a member of its governing body).

Like a legal person, an electronic person (at least given the current state of AI) could not form a will comparable in quality to human volition. That alone does not rule out its legal personhood: it might, likewise, be treated as a fiction. But here lies a serious paradox. If an electronic person is no more than a legal fiction, it lacks will and therefore cannot act without a representative; yet the autonomy of AI – inherent in the very definition found in the AI Act – is precisely what prompts discussion of granting it a separate status in the first place.

Even if incapable of acting itself, an AI system could in principle hold individual rights (most notably proprietary rights) and bear legal liability. This in turn prompts the question: out of what patrimony would compensation be paid to an injured party? Would an electronic person own assets? On what legal basis would it acquire them?[58] Would its assets be ring-fenced solely for liability (and perhaps other narrowly defined purposes), or could it dispose of them? And, if so, would that power be restricted? Might compensation instead be paid under liability insurance taken out for the electronic person? How would such insurance operate? Would the situation amount, in substance, to vicarious liability? If so, what would the status of an electronic person add to tort law, given that healthcare providers are presently liable for the conduct of their employees and carry mandatory liability insurance?

In the light of current AI capabilities, we do not consider the introduction of an electronic person appropriate.[59] Far from clarifying tort law's treatment of AI, it would probably make the landscape still more confusing, opening numerous difficult questions. It could also polarise the public debate on AI, which is already prone to oscillation between over-optimism and undue alarm. Judged against basic legal principles such as legal certainty, predictability, equity, and social cohesion, the notion of an electronic person seems ill-suited. Nevertheless, should AI autonomy and capabilities advance dramatically, the provisional conclusion set out here might require reconsideration.

---

[57] See HAVEL, B. Commentary to Section 20. In: F. Melzer – P. Tégl (eds.). *Občanský zákoník – velký komentář. Svazek I. [Civil Code – Large Commentary. Volume I.]* Praha: Leges, 2013, p. 270.

[58] See Expert Group on Liability and New Technologies. Liability for Artificial Intelligence and Other Emerging Digital Technologies. In: M. A. Geistfeld – E. Karner – B. A. Koch – C. Wendehorst (eds.). *Civil Liability for Artificial Intelligence and Software.* Berlin: De Gruyter, 2023, pp. 361–362.

[59] This is, of course, a relatively common statement. For example, the European Parliament resolution of 16 February 2017 with recommendations to the Commission on Civil Law Rules on Robotics states in its Article 56 that *'at least at the present stage, the liability must lie with a human and not a robot.'* This stance is also common in literature, see for example Expert Group on Liability and New Technologies. Liability for Artificial Intelligence and Other Emerging Digital Technologies, pp. 360–362.

V. CONCLUSION

The current social climate surrounding artificial intelligence calls not for further emotional agitation, but for realistic expectations and a more rational debate. One way to temper both exaggerated hopes and fears is simply to recognise that *not everything is new*.

The law can contribute to this aim, among other things by setting an appropriate civil-liability framework. At present, however, there is considerable uncertainty, both in legal theory and practice and at national and European level, as to what that framework should look like.

Legislation should also convey to the public that AI supports, rather than replaces, the doctor–patient relationship. In our view, granting AI systems the status of an electronic person is presently ill-advised: it would raise more problems than it solves and would stoke unwarranted hopes and anxieties. For similar reasons we do not consider it sound policy, at this juncture, to extend liability for auxiliaries to AI systems.

We have analysed several possible civil-liability regimes for medical AI against the backdrop of Czech law, though we believe our conclusions are broadly transferable to other jurisdictions: all face similar questions and practical challenges.

The liability of the healthcare professional (or their employer) for breach of the standard of care remains our starting point. This standard can be derived both for the indication of an AI system and for its use; at input and output alike. The situation would be clearer if the AI formed part of a new method being formally verified under the Specific Health Services Act, for there a Ministry-approved study protocol would define the correct procedure. In general, however, the medical profession itself must define the standard. As yet there is seldom a clear standard for medical AI, but guidelines, clinical algorithms, and other sources are likely to emerge in the near future.

The most problematic feature of AI for liability purposes is its black-box nature. Outputs can be erroneous, not because the system is *defective* but because of its inherent workings. Defining a defect is therefore difficult, and even if the system malfunctions, proving it will usually be hard. For that reason, the strict-liability regimes that presuppose a defect – liability for damage caused by a defective thing or by a defective product – will see limited use in practice, although they remain crucial where harm can clearly be traced to a defect.

More often, liability will rely on the regime of damage caused by a thing in and of itself. This strict-liability regime can apply even when the thing is used in performing an obligation; what matters is that the harm is *not* linked to the manner of use (for instance, harm caused by a large language model's *hallucination* or by a highly autonomous surgical robot). The key defence is proper supervision. Here the concrete duties laid down, inter alia, in the AI Act will be vital, as practice gradually fleshes out what proper supervision entails.

For these reasons we consider that, as long as fully autonomous AI systems are absent from clinical practice, the existing tools of tort law suffice for a fair allocation of loss.

We acknowledge that a part of the literature differs in its opinion, emphasising specifics of AI systems (such as their adaptiveness and shift of the risk control from the manufacturer to other actors such as machine trainers, data suppliers, or users). Based on this,

some authors call for the establishment of new specific rules for AI liability.[60] Existing civil-liability regimes nevertheless still provide adequate coverage (at least for the time being) for every typical situation in which compensable harm may arise from the use of artificial intelligence in healthcare.

A fresh tort-law regime might be required only if AI systems were to acquire a markedly higher degree of autonomy. For now, however, we advocate applying, and, where necessary, interpretatively adapting, the current tort-law institutes to medical AI, while recognising that empirical experience may in due course reveal the need for targeted legislative change in specific areas.

[60] See GEISTFELD, Mark A., KARNER, Ernst, KOCH, Bernhard A. Comparative Law Study on Civil Liability for Artificial Intelligence. In: M. A. Geistfeld – E. Karner – B. A. Koch – C. Wendehorst (eds.). *Civil Liability for Artificial Intelligence and Software,* p. 127, or ZIKMUNDOVÁ, Klára. Artificial Intelligence and Medical Devices: Do We Need New Regulation?, pp. 392–394.