

Large Language Models in Belgian Primary Care: Clinical Applications, Ethical Challenges, and Epistemological Reflections

Marc Jamouille, MD, PhD^{1,2}, Valentin Angenot, Eng.¹, and Ashwin Ittoo,
PhD¹

¹Business Management System, HEC Liège

²Department of Family Medicine, CAMG-UCL

October 7, 2025

Abstract

Background: Large language models (LLMs) are increasingly integrated into healthcare, with potential applications spanning clinical documentation, patient communication, decision support, and biomedical research. Their use, however, raises significant challenges in terms of ethics, epistemology, and health system organization.

Objective: This narrative review critically examines the opportunities and limitations of LLMs in medicine, with a particular focus on their implications for Belgian primary care and family medicine.

Methods: We conducted a structured narrative review of the recent literature (2023–2025) on LLMs in healthcare, complemented by clinical case illustrations and analysis of the Belgian information system. The review emphasizes practical applications, ethical issues, and epistemological reflections relevant to clinicians and policymakers.

Results: LLMs can enhance efficiency in clinical documentation, support patient-centered communication across languages, and facilitate evidence retrieval and research through standardized terminologies such as the Human Phenotype Ontology. However, their integration into Belgian healthcare is hindered by fragmented information systems and rigid data governance. Ethical concerns include reliability, bias, confidentiality, transparency, patient trust, and environmental impact. Epistemological biases in scientific production and the invisibilization of field knowledge further limit the reliability of LLM outputs. Importantly, LLMs cannot substitute for the therapeutic relationship, which remains central to care.

Conclusion: LLMs represent promising auxiliary tools in Belgian and international medicine, capable of reducing administrative burden and supporting clinical practice and research. Their safe and effective integration requires ethical oversight, systemic reform of fragmented health information structures, and renewed attention to the physician–patient relationship. For Belgian primary care, these technologies should be adopted with caution, transparency, and critical vigilance to ensure they serve patients and reinforce, rather than undermine, the human dimension of medicine.

Keywords

Large Language Models (LLMs); Artificial Intelligence in Medicine; Primary Care; Family Medicine; Clinical Documentation; Human Phenotype Ontology (HPO); Medical Ethics; Epistemology of Medicine; Health Information Systems; Belgium

1 Introduction

The integration of Artificial Intelligence (AI), and particularly Large Language Model (LLM), into contemporary medicine is expanding (see [1, 2]). These models, trained on massive text corpora, are capable of generating, summarizing, translating, and structuring information. Their potential applies both to everyday care and to biomedical research.

However, caution is needed when using the term ‘artificial intelligence. An LLM should above all be considered as a methodological tool, capable of aggregating an impressive quantity of information. As the adage ‘garbage in, garbage out’ reminds us, these models can only reproduce the information they have absorbed.

Of human intelligence, it only reflects certain facets:

- **Linguistic:** mastery of words, language, and formulations.
- **Logical-mathematical:** reasoning, calculation, solving structured problems.
- **Spatial:** visualization, orientation, graphic representation.

On the other hand, it is difficult to imagine that an LLM could reproduce the emotional intelligence characteristic of human beings — at the heart of the therapeutic relationship and social interactions.

It should be emphasized that LLMs work with numerically identified concepts: natural languages are merely an expression intended for human users. As such, an LLM is, by definition, intrinsically multilingual.

The potential applications of LLMs are vast, and only a few can be highlighted here. Each example could easily warrant a full-length discussion, turning this article into a book. The intentionally concise bibliography offers guidance for readers seeking further detail. In what follows, we focus on selected aspects, with particular attention to the ethical and epistemological dimensions of these technologies.

This article offers a narrative critical review of current applications, ethical considerations, and epistemological implications of LLMs in medicine, drawing on recent literature (2023–2025) and examples from primary care practice

2 Use of LLMs in Daily Care

2.1 Clinical Documentation and Workflow Support

- Automatic generation of consultation notes from clinical interviews.
- Coding assistance International Classification of Diseases, 10th Revision (ICD-10), International Classification of Primary Care (ICPC), Systematized Nomenclature of Medicine — Clinical Terms (SNOMED-CT), procedures, prescriptions.
- Drafting forms, certificates, referral letters, and clinical summaries (see [3, 4]).

2.2 Communication with Patients

- Translation of medical jargon into accessible language.
- Multilingual support for non-native patients.
- Production of discharge or self-care instructions (see [5]).

2.3 Medical Decision Support

- Summaries of clinical guidelines tailored to a specific case.
- Verification of dosages, interactions, and contraindications.
- Reminders for prevention and follow-up (vaccinations, screenings).

A crucial limitation: LLMs should never be used as autonomous diagnostic tools. They can complement but not replace the clinician’s expertise [6]) and important limitations in clinical reasoning capabilities [7] .

2.4 Education and Training

- For professionals: rapid access to medical knowledge, updates, preparation of teaching materials.(see [8]).
- For patients: personalized documents explaining chronic diseases, preventive measures, or therapeutic adherence (see [9]).

2.5 Administrative Efficiency

- Assistance with appointment management and initial triage (see [10]).
- Support in responding to patients’ emails.
- Automatic extraction of structured data from free text (e.g., BMI, HbA1c, blood pressure) (See [11]).
- Clinical risk management perspective (See [12]).

These clinical applications highlight the usefulness of LLMs at the bedside. But their potential does not stop there: beyond care, LLMs also open up vast opportunities in biomedical research.

3 Use of LLMs in Biomedical Research

This is an area of potentially considerable expansion. Large language models indeed offer new possibilities for querying scientific literature, integrating heterogeneous data (clinical, genetic, omics), generating hypotheses, and facilitating translational research. Their ability to rapidly analyze vast biomedical corpora could transform the way clinicians and researchers identify new biomarkers, establish phenotype–genotype correlations, and design research protocols (see [13]). The fragmentation of the medical information system in Belgium and the “end-of-the-world” position of family medicine make the applicability of LLMs challenging but certainly attractive.

3.1 Fragmentation of the Medical Information System in Belgium

The general model of medical information processing in Belgium is modeled on the pre-existing paper-based system. Information is generated in hospitals and by specialists, then used to operate the hospital, with each institution remaining an isolated unit. It is then partially made available to the general practitioner, and after some times to the patient, with the patient serving as the connecting thread between the different actors.

This information often remains locked in family doctors' records, with little opportunity for dynamic circulation. The available IT systems reproduce this highly fragmented vertical model, which does not allow for any genuine quality assurance approach (see Figure 1).

Circulating information is essentially in the form of PDF files, without consensus on metadata ([14]), and electronic health records are locked by rigid technical and legal rules. In such informational cacophony, it is hard to see how to effectively implement the immense possibilities offered by artificial intelligence and LLMs.

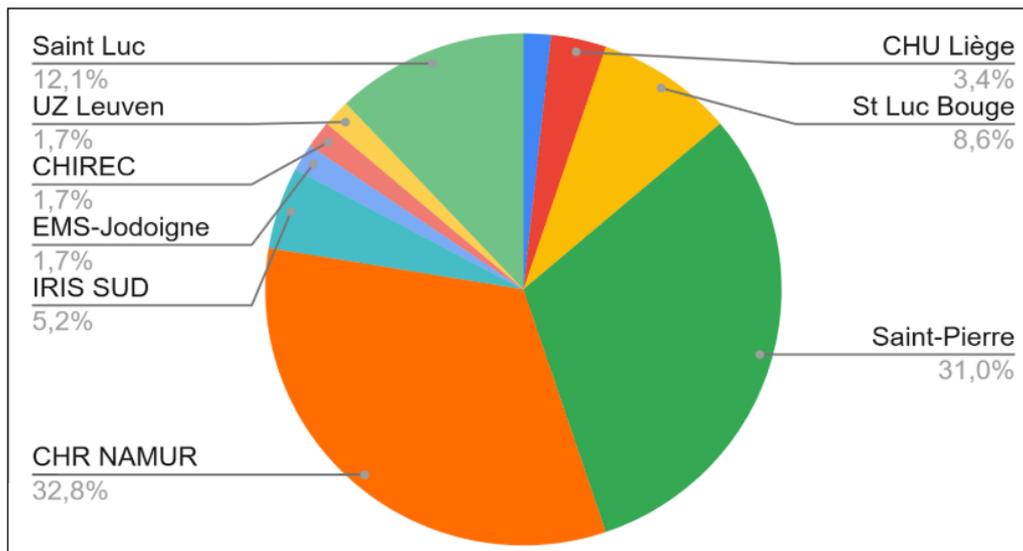


Figure 1: Male, 29 years old, master's student in humanities. Acute COVID in 06/2020, two additional episodes in 2021 and 2023. Outpatient follow-up. Diagnosis of Long COVID in 02/2025. Technetium scintigraphy 03/2025: heterogeneous cerebral perfusion with multiple nonsystematic focal uptake zones, suggesting SARS-CoV-2 encephalitis. Patient consulted in nine care institutions across seven cities between June 2020 and January 2025. Excluding laboratory medicine, 38 physicians from 15 specialties involved (source: Walloon Health Network Hub, Medispring software). (in [15])

3.2 Multi-azimuth Research in Family Medicine

Yet despite these structural barriers, family medicine itself provides fertile ground for innovative approaches. One promising direction lies in multi-azimuth research, where the family doctor's record constitutes an inexhaustible source of data (see [16]). Used locally, it enables genuine quality assurance of practices, either by the physician for personal use, or by comparison in *benchmarking* with a peer group.

The persistent question — “does what is done correspond to what was decided?” (see [17]) — can find a possible answer through thoughtful and structured information

management.

The family physician’s perspective on the health system can be formidable, and is probably feared. Indeed, a simple analysis of the care episodes, as they appear for each patient in the Walloon Health Network HUB, immediately highlights the structural disorganization of the Belgian health system (see Fig. 1).

3.3 Terminological Research via HPO

In this text, we address only one research angle: that of discriminating patient information by means of a specific ontology, as an example of what is possible.

The Human Phenotype Ontology (HPO) ([18, 19]) belongs to the field of ontologies: it is a structured set of concepts, along with their synonyms and hierarchical relations, designed to be used directly by computer systems.

Automated coding by LLMs is generally accurate, since HPO terms align closely with patients’ own formulations. The use of HPO serves a dual purpose:

- **Structuring clinical diversity**, by capturing verbatims through a controlled language that reduces the ambiguity of patients’ spontaneous descriptions.
- **Ensuring semantic interoperability**, so that the collected data can be consistently used by different clinicians, researchers, and information systems.

Table 1: Illustrative case: 11-year-old boy with post-COVID symptoms since October 2021. Withdrawn from school, previously athletic and skilled in strategy video games. Consultation by M. Jamouille, June 2025. HPO codes are clickable links. (in [15])

Symptom	HPO Code (PURL)	Child’s Formulation
Chest pain	HP:0100749	“I really hurt in my sternum.”
Palpitations	HP:0001962	“My heart beats faster and faster.”
Syncope	HP:0001279	“I realize I fainted.”
Cognitive fluctuation	HP:0033630	“Sometimes I say $1+1=3$; the next day it’s fine.”
Fatigability	HP:0003388	“I get too tired so I stop.”
Post-exertional fatigue	HP:0009020	“After playing, I have to lie down.”
...

The use of HPO enables not only standardized clinical descriptions but also advanced computational applications such as automatic indexing, comparative analysis, and integration with other biomedical databases. By structuring symptoms in this way, HPO functions as a **terminological biomarker**: a reproducible linguistic–semantic signature derived from patients’ narratives that ensures consistent mapping to controlled vocabularies while capturing the clinical specificity of syndromes that remain poorly defined.

This example illustrates how terminology can transform clinical narratives into computable data. Building on such foundations, LLMs can further enhance phenotyping and cohort construction, where uniform patient descriptions are essential for research.

3.4 Phenotyping and Cohort Building

In genetic research, patients must be described in a homogeneous manner. LLMs can suggest relevant HPO terms from narrative descriptions, reducing inter-observer variability and accelerating the creation of registries (see Table 1).

3.5 Phenotype–Genotype Reasoning

Matching algorithms such as Exomiser, PhenIX (software that exploits HPO to link observed phenotypes with candidate genes in diagnostics, see [20]) rely on HPO annotations. LLMs can enrich these annotations, suggest links with genes or diseases such as Online Mendelian Inheritance in Man, database of human genetic diseases (OMIM) or database on rare diseases and orphan drugs (Orphanet), and summarize the associated literature.

3.6 Literature Search and Scientific Monitoring

An LLM can query biomedical literature to identify recent associations between an HPO phenotype and genetic variants, thereby facilitating scientific monitoring and hypothesis generation (see [21]).

3.7 Interdisciplinary Harmonization

LLMs facilitate interoperability across disciplines by standardizing clinical descriptions, notably through the use of terminologies and ontologies such as HPO (Human Phenotype Ontology) or SNOMED CT. They thus enable better communication between clinicians, researchers, and bioinformaticians, by translating free-text medical records into structured, shareable data. This harmonization process helps reduce knowledge silos and accelerate translational research (see [22]).

In short, LLMs hold the potential to harmonize research across disciplines. Yet such technical advances cannot be considered in isolation. Their integration into healthcare raises major ethical questions, which we now turn to.

4 Ethics

The use of LLMs in healthcare raises major ethical questions that go beyond purely technical aspects. First and foremost arises the issue of responsibility: who is accountable for an error produced by a model — the developer, the institution, or the healthcare professional who used it? This section examines the immediate ethical issues related to their integration into clinical practice (see [23]).

4.1 Reliability and Bias

Large Language Models (LLMs) have remarkable power to analyze and generate text, but they can also produce factual errors, fallacious reasoning, or invented information (hallucinations). Moreover, these models inevitably reflect the biases present in their training data, whether cultural, linguistic, socio-economic, or related to the underrepresentation of certain populations in medical corpora.

In healthcare, such biases can have significant clinical consequences, affecting the quality of recommendations, reinforcing existing inequalities, or inappropriately guiding decision-making. Human validation by healthcare professionals therefore remains essential to ensure the safety, relevance, and fairness of the results produced by these systems (see [24]).

4.2 Confidentiality and Data Protection

The processing of medical information must comply with applicable regulations as General Data Protection Regulation (GDPR) or Health Insurance Portability and Accountability Act (HIPAA)).

- Health data entrusted to an LLM must be protected against any unauthorized use.
- Models must be deployed in secure environments, with appropriate encryption and restricted access.
- Exchanges with LLMs must not expose information that can directly identify a patient, except within a strictly controlled and documented framework (see [25]).

4.3 Transparency and Acceptability

Patients must be informed whenever an artificial intelligence tool, such as an LLM, is used in their care. Such transparency is essential to build and maintain trust, while enabling responsible and ethical use of these technologies. It involves not only disclosing the use of AI but also explaining in accessible terms its purposes, limitations, and the modalities of medical supervision. The social and clinical acceptability of LLMs depends on this trust contract, where technological innovation remains at the service of the patient and under the control of the healthcare professional (see [26]).

4.4 Changing Access to Knowledge and Physician–Patient Communication

Patients now have direct access to LLMs and can freely question them on any medical topic. This profoundly alters the traditional balance between professional and lay knowledge:

- The physician lives and works in uncertainty, reasons in terms of probabilities, and accepts the complexity of clinical situations.
- The patient, questioning an LLM, receives answers often phrased as certainties. Where the physician thinks “there is a probability,” the patient understands “there is a certainty.”
- This discrepancy often creates misunderstandings: the practitioner offers nuance, while the patient cites the machine as absolute truth.
- The physician thus loses part of their traditional prerogative: being the mediator and guarantor of access to medical knowledge. This transformation weakens the therapeutic relationship, which must be rethought in the era of LLMs (see [27]).

4.5 The Therapeutic Relationship, an Irreplaceable Dimension

We must not underestimate the power of the therapeutic relationship, which remains to this day irreplaceable by any machine. This relationship, grounded in the encounter between two human beings, engages the mirror neurons highlighted by [28] and activates still-uncertain mechanisms such as the placebo effect.

The mirror-neuron metaphor illustrates this phenomenon: those of the therapist and the patient begin to “reflect” together, in the literal sense (shared neuronal activation) as well as the figurative sense (a joint construction of meaning). This subtle synchronization fosters trust and amplifies the effectiveness of care (see [29]).

The placebo effect, for its part, remains a scientific enigma: only briefly touched upon today by the neurosciences, it testifies to the power of relationship and expectation in the healing process.

No technology can reproduce the impact of this profound human interaction, where trust, listening, and emotional resonance contribute to recovery. However advanced they may be, LLMs cannot substitute for the power of the clinical relationship, which remains the heart of care.

4.6 Environmental Impact

The use of LLMs also raises ecological issues. Their training and operation require considerable consumption of electricity and water, linked to data centers and cooling processes.

Yet the health system already contributes significantly to global warming, at a level comparable to the automotive industry. Introducing new energy-intensive technologies without prior reflection risks further increasing this environmental footprint.

The ethics of using LLMs must therefore integrate an ecological dimension:

- systematically assess the energy and water costs of these technologies;
- limit their use to applications that are relevant and genuinely beneficial for patients and caregivers;
- promote sober and responsible digital solutions (see [30] and [31]).

Thus, ethical reflection extends from issues of bias and confidentiality to the ecological footprint of LLMs. But ethics alone does not capture the whole picture: LLMs also reveal deeper epistemological biases in how medical knowledge is produced and transmitted.

5 Epistemology and Biases of Knowledge

Beyond ethical questions, LLMs reflect the limits and biases of scientific production itself. This section addresses the conditions under which medical knowledge is produced and the tensions between technique, culture, and field practices.

5.1 Limits Related to Scientific Content

An LLM relies primarily on the existing scientific literature. However, the latter is marked by several biases:

1. Emerging fields: if the question concerns a new disease (for example, *Long Covid*), the model’s responses reflect researchers’ beliefs or hypotheses more than established knowledge.
2. Never say no : A key limitation of LLMs is their tendency to provide an answer even when no reliable information exists. We need models and interfaces that can explicitly recognize uncertainty and refrain from answering when evidence is lacking.
3. Academic biases: publications predominantly come from the academic world subject to the rule of “publish or perish,” which can orient scientific production toward quantity rather than clinical relevance.
4. Lack of access: frontline physicians, rich in direct experience with patients, often have neither the time nor the means to publish.
5. Publication costs: many journals require substantial publication fees so-called Article Processing Charges (APC), which limits the ability of non-academic practitioners to share their clinical experience.
6. Invisibilization of field knowledge: the physician–patient relationship, sometimes established over two or three generations, gives the practitioner a global and *anthropological* understanding of the patient’s life and family. This knowledge is captured neither in scientific literature nor in medical training. Invoking the “biopsychosocial approach” is not enough: no published article can, by itself, solve systemic problems rooted in patients’ histories and living contexts.
7. Health as a market: producers of technological and pharmaceutical goods finance certain publications, remunerate ghost authors, and sometimes influence medical services. This structural bias feeds LLM databases, once again distorting the information returned. (see [32]).
8. Cultural prejudice: another bias lies in the widely shared idea that medicine is always beneficial to health. This conviction creates a favorable prejudice toward procedures and medications. Side effects are often poorly documented and rarely emphasized; this imbalance is reflected in the corpora used by LLMs. (see [33]).
9. Prevention paradigm: as taught, it derives from the stages of syphilis elaborated after the Second World War. Prevention is conceived as an intervention within a disease course, without a quality-assurance approach. Quaternary prevention, which aims to limit overmedicalization, has no place in a productivist world (see [34, 35, 36]).

5.2 Gap Between Technique and Culture

Another fundamental epistemological bias lies in the widening gap between technique and culture. Traditionally, patients and physicians aged together within the same community, sharing a cultural heritage and medical knowledge patiently transmitted from generation to generation. This common framework fostered mutual understanding and continuity in the meaning given to health and illness.

The rapid evolution of technology disrupts this balance: it provides universal, multiple, and constantly expanding knowledge, yet inscribed in a culture of production and consumption. This knowledge—impossible to absorb in its entirety by the actors of care—is moreover susceptible to all kinds of malfeasance or manipulation.

The result is a risk of collective disorientation: a “health bouillabaisse” in which patients, physicians, and institutions struggle to discern the reliable from the uncertain, and where mutual trust tends to fray. As stated by [37] *"Successful deployment requires addressing critical challenges, including data privacy, algorithmic bias, model interpretability, regulatory oversight, and maintaining human clinical oversight"*

These epistemological tensions highlight how technical progress, cultural shifts, and systemic biases intertwine in shaping medicine’s future. Against this backdrop, we can now draw some conclusions about the role of LLMs in care, research, and ethics.

6 Conclusion

LLMs today represent powerful assistants both in daily medical practice and in biomedical research. Their combined use can reduce administrative burden, improve communication, standardize data, and accelerate scientific discovery.

However, their integration must be accompanied by:

- thorough ethical reflection;
- critical vigilance toward the structural biases of scientific production;
- a redefinition of the physician–patient relationship in the digital era;
- and consideration of environmental impacts.

LLMs cannot replace the human dimension of care, but they can serve as powerful auxiliaries if used judiciously and in line with the core values of medicine.

List of Acronyms

AI Artificial Intelligence

APC Article Processing Charges

ICD-10 International Classification of Diseases, 10th Revision

ICPC International Classification of Primary Care

GDPR General Data Protection Regulation

HIPAA Health Insurance Portability and Accountability Act

LLM Large Language Model

Medispring Belgian family medicine software

OMIM Online Mendelian Inheritance in Man, database of human genetic diseases

Orphanet database on rare diseases and orphan drugs

SARS-CoV-2 Severe Acute Respiratory Syndrome Coronavirus 2

SNOMED-CT Systematized Nomenclature of Medicine — Clinical Terms

Conflict of Interest Statement

The author declares no conflicts of interest.

Note on the Use of Artificial Intelligence

This manuscript was drafted with the occasional assistance of a generative artificial intelligence tool (ChatGPT, OpenAI), which was used exclusively for writing support, idea structuring, and reformulation. All scientific content was reviewed, validated, and revised by the human author(s). This process illustrates a responsible and transparent example of AI use in health care publishing.

References

- [1] A. A. J. Thirunavukarasu et al. “Large language models in medicine”. In: *Nature Medicine* 29.8 (2023), pp. 1930–1940. DOI: 10.1038/s41591-023-02448-8 (cit. on p. 2).
- [2] S. Tian et al. “Opportunities and challenges for ChatGPT and large language models in biomedicine and health”. In: *Briefings in Bioinformatics* 25.1 (2023) (cit. on p. 2).
- [3] S. B. Patel and K. Lam. “ChatGPT: the future of discharge summaries?” In: *The Lancet Digital Health* 5.3 (2023), e107–e108. DOI: 10.1016/S2589-7500(23)00021-3. URL: [https://www.thelancet.com/journals/landig/article/PIIS2589-7500\(23\)00021-3/fulltext](https://www.thelancet.com/journals/landig/article/PIIS2589-7500(23)00021-3/fulltext) (cit. on p. 2).
- [4] M. Karino et al. “Exploring the Potential of ChatGPT for the Summarization of Patient Medical Histories: A Pilot Study”. In: *Cureus* 17.5 (2025). DOI: 10.7759/cureus.84133 (cit. on p. 2).
- [5] C. Park, M. H. An, G. Hwang, R. W. Park, and J. An. “Clinical Performance and Communication Skills of ChatGPT Versus Physicians in Emergency Medicine: Simulated Patient Study”. In: *JMIR Medical Informatics* 13.1 (2025), e68409. DOI: 10.2196/68409 (cit. on p. 3).
- [6] F. Shah-Mohammadi and J. Finkelstein. “Accuracy evaluation of GPT-assisted differential diagnosis in emergency department”. In: *Diagnostics* 14.16 (2024), p. 1779. DOI: 10.3390/diagnostics14161779 (cit. on p. 3).
- [7] L. G. McCoy et al. “Assessment of Large Language Models in Clinical Reasoning: A Novel Benchmarking Study”. In: *NEJM AI* 2.10 (2025), AIdbp2500120. DOI: 10.1056/AIdbp2500120 (cit. on p. 3).

- [8] T. H. Kung, M. Cheatham, A. Medenilla, et al. “Performance of ChatGPT on USMLE: Potential for AI-assisted medical education using large language models”. In: *PLOS Digital Health* 2.2 (2023), e0000198. DOI: 10.1371/journal.pdig.0000198 (cit. on p. 3).
- [9] Z. Yan et al. “Ability of ChatGPT to Replace Doctors in Patient Education: Cross-Sectional Comparative Analysis of Inflammatory Bowel Disease”. In: *Journal of Medical Internet Research* 27 (2025), e62857. DOI: 10.2196/62857 (cit. on p. 3).
- [10] J. Li, W. Zhang, X. Chen, H. Wang, M. Zhou, and F. Liu. “Patient Triage and Guidance in Emergency Departments Using Large Language Models: Multimetric Study”. In: *Journal of Medical Internet Research* 27 (2025), e71613. DOI: 10.2196/71613. URL: <https://www.jmir.org/2025/1/e71613> (cit. on p. 3).
- [11] B. Gu et al. “Scalable information extraction from free text electronic health records using large language models”. In: *BMC Medical Research Methodology* 25.1 (2025), p. 23. DOI: 10.1186/s12874-025-02470-z (cit. on p. 3).
- [12] M. Ferrara et al. “Risk management and patient safety in the artificial intelligence era: a systematic review”. In: *Healthcare* 12.5 (2024), p. 549. DOI: 10.3390/healthcare12050549 (cit. on p. 3).
- [13] H. Rao, W. Liu, H. Wang, I. Huang, Z. He, X. Huang, et al. “A Scoping Review of Synthetic Data Generation for Biomedical Research and Applications”. In: *arXiv preprint* (2025). arXiv: 2506.16594 [cs.LG]. URL: <https://arxiv.org/abs/2506.16594> (cit. on p. 3).
- [14] M. Jamouille. “Réseau Santé Wallon en Médecine de Famille, essai d’utilisation en temps réel”. In: 2019. eprint: <https://orbi.uliege.be/handle/2268/239840> (cit. on p. 4).
- [15] M. Jamouille. *Long COVID, invisible illness*. Tech. rep. 2025. eprint: <https://orbi.uliege.be/handle/2268/335655> (cit. on pp. 4, 5).
- [16] L. A. Green and J. Hickner. “A short history of primary care practice-based research networks: from concept to essential research laboratories”. In: *Journal of the American Board of Family Medicine* 19.1 (2006), pp. 1–10. DOI: 10.3122/jabfm.19.1.1 (cit. on p. 4).
- [17] R. Grol and J. Grimshaw. “From best evidence to best practice: effective implementation of change in patients’ care”. In: *The Lancet* 362.9391 (2003), pp. 1225–1230. DOI: 10.1016/S0140-6736(03)14546-1 (cit. on p. 4).
- [18] P. N. Robinson, S. Köhler, S. Bauer, D. Seelow, D. Horn, and S. Mundlos. “The Human Phenotype Ontology: A Tool for Annotating and Analyzing Human Hereditary Disease”. In: *American Journal of Human Genetics* 83.5 (2008), pp. 610–615. DOI: 10.1016/j.ajhg.2008.09.017. URL: <https://www.sciencedirect.com/science/article/pii/S0002929708005351> (cit. on p. 5).
- [19] M. A. Gargano et al. “The Human Phenotype Ontology in 2024: Phenotypes around the World”. In: *Nucleic Acids Research* 52.D1 (2024), pp. D1333–D1346. DOI: 10.1093/nar/gkad1005 (cit. on p. 5).
- [20] P. N. Robinson et al. “Improved exome prioritization of disease genes through cross-species phenotype comparison”. In: *Genome Research* 24.2 (2014), pp. 340–348. DOI: 10.1101/gr.160325.113 (cit. on p. 6).

- [21] I. Taleb, A. N. Navaz, and M. A. Serhani. “Leveraging large language models for enhancing literature-based discovery”. In: *Big Data and Cognitive Computing* 8.11 (2024), p. 146 (cit. on p. 6).
- [22] E. Chang and S. Sung. “Use of SNOMED CT in Large Language Models: Scoping Review”. In: *JMIR Medical Informatics* 12.1 (2024), e62924. DOI: 10.2196/62924 (cit. on p. 6).
- [23] J. C. L. Ong et al. “Ethical and regulatory challenges of large language models in medicine”. In: *The Lancet Digital Health* 6.6 (2024), e428–e432. DOI: 10.1016/S2589-7500(24)00061-X (cit. on p. 6).
- [24] J. Morley et al. “The ethics of AI in health care: A mapping review”. In: *Social Science & Medicine* 260 (2020), p. 113172. DOI: 10.1016/j.socscimed.2020.113172 (cit. on p. 7).
- [25] S. Benjamens, P. Dhunoo, and B. Meskó. “The state of artificial intelligence-based FDA-approved medical devices and algorithms: an online database”. In: *NPJ Digital Medicine* 3.118 (2020). DOI: 10.1038/s41746-020-00324-0 (cit. on p. 7).
- [26] A. Jobin, M. Ienca, and E. Vayena. “The global landscape of AI ethics guidelines”. In: *Nature Machine Intelligence* 1.9 (2019), pp. 389–399. DOI: 10.1038/s42256-019-0088-2 (cit. on p. 7).
- [27] A. Choudhury and Z. Chaudhry. “Large language models and user trust: consequence of self-referential learning loop and the deskilling of health care professionals”. In: *Journal of Medical Internet Research* 26 (2024), e56764 (cit. on p. 7).
- [28] V. S. Ramachandran. *Le cerveau fait de l’esprit*. Édition française de *The Tell-Tale Brain*. Paris: Dunod, 2011 (cit. on p. 8).
- [29] T. D. Wager and L. Y. Atlas. “The neuroscience of placebo effects: Connecting context, learning and health”. In: *Nature Reviews Neuroscience* 16.7 (2015), pp. 403–418. DOI: 10.1038/nrn3976 (cit. on p. 8).
- [30] N. Jegham, M. Abdelatti, L. Elmoubarki, and A. Hendawi. “How hungry is ai? benchmarking energy, water, and carbon footprint of llm inference”. In: *arXiv preprint arXiv:2505.09598* (2025) (cit. on p. 8).
- [31] M. Jamoulle. “Quaternary Prevention and Climate Global Warming”. In: *1st WONCA Conference on Planetary Health, Brazil*. Opening lecture. 2021. URL: <http://orbi.uliege.be/handle/2268/262531> (cit. on p. 8).
- [32] R. Moynihan, I. Heath, and D. Henry. “Selling sickness: the pharmaceutical industry and disease mongering”. In: *BMJ* 324.7342 (2002), pp. 886–891. DOI: 10.1136/bmj.324.7342.886. URL: <https://www.bmj.com/content/324/7342/886> (cit. on p. 9).
- [33] R. Moynihan and D. Henry. “The fight against disease mongering: generating knowledge for action”. In: *PLoS Medicine* 3.4 (2006), e191. DOI: 10.1371/journal.pmed.0030191. URL: <https://doi.org/10.1371/journal.pmed.0030191> (cit. on p. 9).
- [34] J. P. A. Ioannidis. “Why most published research findings are false”. In: *PLoS Medicine* 2.8 (2005), e124. DOI: 10.1371/journal.pmed.0020124 (cit. on p. 9).
- [35] T. Greenhalgh, J. Howick, and N. Maskrey. “Evidence based medicine: a movement in crisis?” In: *BMJ* 348 (2014), g3725. DOI: 10.1136/bmj.g3725. URL: <https://www.bmj.com/content/348/bmj.g3725> (cit. on p. 9).

- [36] M. Jamouille et al. “Implications éthiques, pédagogiques, sociopolitiques et anthropologiques de la prévention quaternaire / Ethical, pedagogical, socio-political and anthropological implications of quaternary prevention”. In: *Revue Médicale de Bruxelles* 39.4 (2018), pp. 383–393. URL: <https://tinyurl.com/7ysd7wwb> (cit. on p. 9).
- [37] Y. A. Fahim, I. W. Hasani, S. Kabba, . . . , and W. M. Ragab. “Artificial intelligence in healthcare and medicine: clinical applications, therapeutic advances, and future perspectives”. In: *European Journal of Medical Research* 30.1 (2025), p. 848. DOI: 10.1186/s40001-025-03196-w. URL: <https://eurjmedres.biomedcentral.com/articles/10.1186/s40001-025-03196-w> (cit. on p. 10).