Concetta Press
Journal of Biomedical Informatics and AI

From Leibniz' Quest to Dirac Notation

Toward a Digital Epistemology for Reasoning under Uncertainty in Patient-Centric Learning Healthcare Systems

Michael Rebhan[1], Laura Ferrarello[2], Moritz Freidank[3], Maïté Moniatte[4], Rukshan Ranatunge[5], Christian Seebode[6], and Barry Robson[7]

[1]Novartis Pharma AG, R&D, Basel, Switzerland; co-founder of the elevateHealth Open Science community, and the Institute for Human-Centric Health Innovation, Basel

[2]Laboratory for Intelligent Global Health Technologies, EPFL, Lausanne, Switzerland

[3]Visium SA, Prilly, Switzerland

[4]Novartis Pharma AG, R&D, Basel, Switzerland

[5]Swiss Tropical and Public Health Institute, Basel, Switzerland; World Health Organization; Ministry of Health, Sri Lanka; member of elevateHealth

[6]Patient Centered IT, Switzerland

[7]Ingine Inc., Cleveland, Ohio, USA, and The Dirac Foundation, Oxfordshire, UK

Article Type	Original Research
Volume / Issue	1 / 1:5 (2026)
DOI	https://doi.org/10.5281/zenodo.20719208
Correspondence	Barry Robson, PhD DSc Email: barryrobson@ingine.com Ingine Inc., 11000 Cedar Ave Ste 100, Cleveland, Ohio
Received / First Decision / Accepted	May 07, 2026 / May 20, 2026 / June 13, 2026
Disclosure	Funding: None. Conflicts of Interest: None.

Abstract

At the dawn of the scientific age, before scientific disciplines had clearly marked boundaries, Leibniz was a polymath on the quest for a language that would enable reasoned discourse to be resolved by calculation (or formal inference), using mathematically precise notation, across different intellectual traditions. Today, enabled by scientific progress, there is interplay between human and machine intelligence, new ways of encoding certain and uncertain knowledge, and new ways of organizing work based on knowledge. We propose a contemporary interpretation of Leibniz' vision as a '*digital epistemology'* for AI in medicine, focusing on difficult medical cases where evidence-based decision making by patients and healthcare professionals is non-trivial. In this approach, reasoning is based on Dirac notation and algebra based on hyperbolic-complex (split-complex) vector spaces as a mathematical foundation. Dirac's system has served science for almost a century, and as Dirac noted, it should be applicable to all aspects of human thought where numbers such as probabilities are involved (Dirac, 1930). To allow such learning systems to discover and validate the most informative patterns in patient data, combining human and machine intelligence, Open-Science-style transparency is used, rather than ‘black box AI’, including its encoding of uncertainty, contradictions, and conflicting recommendations. With increasing “FAIRification” (i.e. the process of making health-related data Findable, Accessible, Interoperable, and Reusable), and new health data spaces (federations across health data silos), we propose that we should soon have improved frame conditions for building such learning systems using the proposed digital epistemology. By “learning system” is usually meant a system that seeks to minimize the disagreement between the predictions the system makes and the actual situations; the minimization may involve local minima. However, any system that extracts knowledge, such as data analytics and data mining, can be considered as learning from data. We discuss systems designed for the translation of P4 science, i.e. systems that aim to make care more Personalized, Preventive, Participatory and Predictive, digital twins, and the modeling and simulation of patient trajectories, including trajectories that do not map well to the medical knowledge and evidence landscape. We refer to these as *unguided patients*, who typically have little recorded data. Importantly, however, they are patients who are clearly different from the populations studied in large clinical studies.

Keywords: healthcare; Leibniz; universal exchange language; unguided patients

1 Introduction

In the late 17th century, *Gottfried Wilhelm Leibniz* (1646-1716) dreamed of a *characteristica universalis,* a symbolic logic through which all (or most) human disputes could be resolved by calculation (or formal inference). As an epistemological[1] dream that has inspired and puzzled many, it remains largely elusive, apart from systems deployed in narrowly scoped areas that can be modeled well with mathematical calculations, e.g. as in many engineered systems (Milkov, 2006).

The various paradigm shifts[2] experienced by the physics community since the 19th century reinforced the key role of mathematics as a solid, trusted scientific foundation of the field, for its epistemology (Weigert, 2023). While the human mind often struggles to comprehend theory related to quantum mechanics and relativity, it can learn to ‘trust the math’ if it is scientifically robust and useful (Gibney, 2025; Dine, 2022; Robson & St. Clair, 2022). Note that, while quantum mechanics is often defined as the fundamental theory in physics that describes the behavior of nature at and below the atomic scale, where classical mechanics no longer applies, many physicists, including some of its founders (e.g. Dirac, 1933), believed that it is extensible to the everyday world of human experience. It is an idea that becomes important in the present paper. More than three centuries after Leibniz’ quest, considerable progress has been made in many areas of science, often involving many paradigm shifts, that make it timely to revisit his dream. Leibniz’ main motivation was to use the power of mathematics as a way of thinking and communicating that would be more precise than natural language. In contrast, artificial human languages like Esperanto have focused on ease of mastery because of their regular grammar and logical and flexible word-building, rather than precision per-se, by adapting word-roots from natural languages. They are essentially subject to symbolic manipulation rather than involving quantitative notions of scope, probability, and uncertainty. Mathematics has indeed become a robust foundation for discourse in a great variety of sciences requiring such quantification such as data science, decision theory, statistics, probabilistic reasoning, and causal inference. Not least, there is artificial intelligence (AI). Ironically, however, we currently also face problems of understanding the ‘inner workings’ of modern ‘black box AI’. The learning is ultimately a series of arbitrarily distributed weights that are hard to interpret (compared, for example, with probabilities), so that the information is in effect invisible. This is especially worrying in mission-critical

[1] We refer to ‘epistemology’ (and ‘epistemological’ concepts) throughout this text, using a broadened meaning that is aligned with the Oxford Dictionary’s definition of ‘epistemology’ as “The theory of knowledge, especially with regard to its methods, validity, and scope, and the distinction between justified belief and opinion.” Our use goes beyond the philosophical definition, toward applications in the context of an Open Science toolkit we aim to develop. Other definitions of epistemology exist.” See also our glossary.

areas such as medicine, where robust, explainable decision making and related scientific learning is crucial, requiring a visible reasoning approach that both humans and machines can relate to (Raposo, 2025). This increases the tension between the complex solutions generated by AI and human interpretability.

Despite great progress, with increasing specialization and demarcation of disciplinary boundaries in science, there arises the challenge of navigating those transdisciplinary boundaries, which manifest as specialized natural languages and distinctive mathematical styles. Crossing these boundaries implies entering what we may call 'epistemological fields', meaning clusters of similar epistemologies in science that agree on many norms, beliefs and rules related to what counts as valid (high-quality) knowledge, as a contrast to subjective, personal opinion. Each epistemological field uses different shared notions of encoding, using and judging knowledge that enables an epistemological community to act in concert, i.e. performing innovation within the boundaries of a particular paradigm (or 'Disciplinary Matrix')[2]. However, as a community develops its own modifications and extensions of language for effective coordination among domain experts, the increasing specialization at work in the professional communities[3] may not always result in users agreeing on the best way to judge the quality of knowledge, i.e. for a particular use case and related decision making where diverse inputs from multiple fields are desired.

The medical domain is notably a diverse field that is particularly rich in such epistemology across its own scientific subdisciplines and medical specialties, as its decision-making logic often relies on their aggregated knowledge (Tonelli and Bluhm, 2021). Also, no medical discipline seems untouched, or fails to be called upon, by "Translational Medicine" or "Translational Research" that brings latest discoveries into medicine from physics, chemistry, biology, mathematics and computer science. It is a recognized discipline, with organizations such as the European Society for Translational Medicine (EUSTM) offering qualifications in it.

2 Paradigm Shifts

In the present paper we wish to promote a paradigm shift, both regarding personalized patient-centric medicine including rare cases, with the patient having more control (Section 3), and by seeking to find the best digital means of achieving that (Section 4 onwards). It is of course impossible to avoid discussion of current AI. Many may argue against our concern that current AI, including Large Language Models, is a risky consideration for medicine. But future events are, of course, not always to be firmly specified for the patient of immediate interest. In the foreseeable future, predictive models seem unsuited to rare or unique situations where relevant scientific evidence from other patients is weak or entirely absent from an uncommon patient's pathway. However, given relevant information, future pathways can be modeled and expected to hold true with a certain probability: see e.g. Maurer *et al*. (2024). By rare or unique patient cases we do not, of course, mean collectively, since patients can be rare and unique in many ways, as discussed in Section 3. It is well known that AI runs into difficulties with sparse data and rare events, and much research is underway to ameliorate the problem (Carou, Sartal, and Davim, 2023) although, arguably, attempted solutions seem a little fragmented (case-dependent). Be that as it may,

[2] The term 'paradigm shift' is based on Thomas Kuhn's work on the history of scientific revolutions. Note that in his later work he preferred to use the term 'Disciplinary Matrix' (roughly matching his original notion of 'paradigm'), with emphasis on the power of 'exemplars' in paradigm construction.

[3] From a patient perspective this may look quite different, so that notion is under-represented in such a perspective focusing on expert communities. At system level, in healthcare organizations, since the time of Leibniz, there is a trend to increasing specialization that affects how organizations work, and make decisions. In preventive care, for example, knowledge outside healthcare organizations is likely a key factor for success, but often not sufficiently involved in system design compared to professional communities with their epistemologies and biases (note the 'participatory' aspect of P4, see Appendix). Being serious about the patient-centric approach then also means to tap into non-scientific non-expert, individual or experiential knowledge where it helps us learn.

efforts should be compared with an example of a general approach, the Theory of Expected Information, developed early by Robson and as refined by Robson (2005) and applied by Robson and Boray (2015), which also forms an intrinsic part of the proposed approach to a healthcare language discussed in the present paper. Although discussion is beyond present scope, it is worth noting that this approach to rareness forms a natural mathematical continuity with the mathematics of “Black Swan” events that are rare, unexpected and expensive, such as earthquakes, market crashes, and pandemics.

It is also of course good for all patients that western medicine has developed greatly since the less scientific days of the witch doctor and village shaman, and cults like the Asclepius cult, but not always in a smooth continuum of progress. Alvin Toffler (1928 - 2016) is well-known for analyzing and even predicting sudden bursts of social and technological changes. Medicine has often reflected sudden social transitions and the paradigm shifts that they imply. All the major periods of medical advancement seem to coincide with paradigm shifts in the sense of changes in worldview on a large social scale, perhaps most notably as in the Age of Enlightenment (1685 to 1815), but in any event when scholars and practitioners felt that they “get it!”. “The Demon-Haunted World: Science as a Candle in the Dark” was a 1995 book by Carl Sagan emphasizing the transitioning concept jumps from unscientific to scientific thinking, individually and collectively. Beginning with the revelations of Hippocrates, but particularly since the days of Leibniz whose life (1646–1716) overlapped with the dawn of the Age of Enlightenment, medicine rapidly became scientific in the more modern mechanistic sense, and increasingly ‘precise’ in its diagnostic and therapeutic logic (from visible symptoms down to the molecular biology of disease). Though Leibniz was not a physician, he engaged deeply with biological and medical questions, debated mechanism vs. vitalism (e.g., with Stahl), and speculated about preformation, organisms, and life as structured systems.

It may be an apocryphal story that around 1899-1902 many thought that physics was complete and that “everything that can be invented has been invented”. That seems an exception to the notion of a paradigm shift, but in fact it was a new paradigm with the sense that a great deal had been accomplished in the Age of Enlightenment and two major industrial revolutions, and for many there was sadness that the most exciting days seemed to have passed. If so, that misguided perception was short-lived. It was already running into the rise of early quantum theory 1900–1924, and Canadian physician and educator Sir William Osler (1849–1919) was shaping modern clinical medicine with focus on compassionate, personalized, patient-centered care, lifelong learning in the spirit of modern Evidence Based Medicine, and scientific "equanimity" (calmness) in practice, highlighting that “medicine is a science of uncertainty and an art of probability”. It might be argued that this more the personal, compassionate side of Osler’s medicine did benefit from reflection in an earlier period of more scientific, technical, and industrial calm, but, be that as it may, medicine has since moved forward dramatically in terms of probability-based decision aspects, making increasingly precise scientific predictions and increasingly sophisticated medical reasoning that should in the end enable better care decisions leading to better care outcomes.

Consequently, while in many disciplines we want to take full advantage of all the impressive scientific and technical progress for improving medicine and patient outcomes further, we must reflect upon the need to develop new kinds of solutions, i.e. *to create more room for new paradigms to emerge and develop*. Only then can we learn how to best navigate all these epistemological fields and sciences (including those that are not integrated well yet into medical epistemology[4]) in a way that enables the best possible decision making, and the effective dissemination of what has been learned. This is essentially

[4] While physicians and scientists, in collaborations, often experience considerable overlap in their epistemologies, a strong participation by patients can create particular epistemological challenges. The term bricolage, in this context, is a way of developing a shared language and understanding in such complex communities, e.g. when comparing clinical research endpoints and what really matters to patients and their families. See Levi-Strauss (1962) for a seminal text that contrasts the notion of bricolage with the mindset of an engineer or scientist.

the same concern behind the December 2010 PCAST report, "Realizing the Full Potential of Health Information Technology to Improve Healthcare for Americans: The Path Forward". See the reference President's Council of Advisors on Science and Technology (2010). It urged a shift from simple electronic health records (EHRs) to a nationwide, patient-centric health information technology (HIT) infrastructure. It emphasized using data for a learning health system and for improving care quality and safety. Specifically, it was concerned with interoperability and sharing across epistemological niches and facilitating personalized medicine by development of some kind of a "universal exchange language" for health data, for making medical record information securely available at any time an any place, and enhancing patient access to their own information.

Our focus here follows the spirit of William Osler mentioned above, with his more personalized medicine as "a science of uncertainty and an art of probability". In this paper, it is specifically in the most difficult cases for care decisions when current science-based medical knowledge is still too weak for guiding decisions with high confidence. Such difficult cases provide valuable material for learning about the limitations of currently established paradigms used in science-based medical decision making, and how we learn from the outcomes of past decisions. While biostatistics, epidemiology, and Evidence Based Medicine must continue to collect and analyze data from millions of patients to form the basis of rational, objective interpretations and decisions, each individual patient is an important data point. Even a modest patient record might minimally have three distinct descriptions or measurements as demographic and clinical data (e.g. low, normal, high clinical values) for each of a mere, say, 200 entries, giving an astronomically large combination of possible different records. The number is more than the estimated number of electrons in the observable universe, and so certainly far more than the human population of Earth, making every patient ultimately unique even into the far future. On top of that, new observations and a flood of genomic and proteomic and other data lead to an explosion of health state biomarkers, add valuable knowledge.

A difficulty for data-analytical methods is that there cannot be meaningfully said to be much of a pattern in a single unique case. So far, the medical domain has already gathered considerable experience with the application of science (and its mathematics-based toolkit) for enabling the discovery and validation of *informative patterns* in patient datasets that capture important information related to what happened in such difficult cases. Pioneering work, for example, was performed in collaborative, transdisciplinary[5] communities consisting of different kinds of physicians and scientists, working closely together with patients on a specific set of well-defined problems, i.e. in a context that requires *epistemic bricolage*[4] to create new kinds of solutions across epistemological fields, for understanding such difficult cases. In collaborative diverse communities that share the goal of learning how to inform better decisions in healthcare in that specific area where such *epistemic bricolage* is approached by intensive collaboration. That is, selecting, together, what seems most relevant for such bricolage, and developing a shared language that enables effective bricolage.

Stated in the language of the biostatistician, it remains crucially important to apply data analysis of many patients for interpretation, prediction, and decision-making in each patient case, and not only to provide knowledge to guide interpretation of the rarer and unique cases but also, more fundamentally, to reveal exactly *what is* rare or unique. Conversely, those physicians placing too much emphasis on their personal experience might form the erroneous conclusion that a seemingly appropriate interpretation and treatment is appropriate for most patients. Here, understanding the epistemological impact by the field called *Evidence-Based Medicine* (EBM) is of particular importance. EBM has contributed considerably to

[5] We use the term 'transdisciplinary' rather than 'interdisciplinary' as it is not only about an interface between two different disciplines but also about knowledge from several disciplines and across epistemological fields (not only 2). Such transdisciplinary work may need to be more encouraged to better deal with some of the problems outlined in this paper. See also footnote 4 (bricolage).

make medicine more scientific, enabled by a math-based 'evidence paradigm' [6] which is used increasingly to inform evidence-based decision making in many modern health care systems (Sackett, 1996; Robson, 2016).

Within the EBM paradigm, RCTs (*randomized controlled trials*, e.g., a particular type of clinical trial) play a crucial role. They are considered the most robust, generally applicable method for gathering reliable evidence. Nonetheless, the RCT is interventional and so not always appropriate: epidemiology is considered the mother of EBM and John Snow is considered the father of epidemiology, by investigating the cholera outbreak in Soho, London, in 1854. The initial view was that *miasma* was the cause, i.e. that it was due to bad gases in the air. Snow was the founder of anesthesia and experienced in studies of gas toxicity. In his hands it was essentially the first use of the four pillars of non-interventional observational evidence and the basis of estimating probabilities of the possible causes through counting. In that case, these four pillars were as follows. How many drank from the local water pump and got cholera, how many drank from the pump and did not get cholera, how many did not drink from the pump and got cholera, and how many did not drink from the pump and did not get cholera? The intervention was ultimately in removing the handle from the pump. But whether a study is interventional or not, an important aim is to separate clear causation from correlation[7], e.g. about a care decision involving a new drug or diagnostic which may provide extra value compared to an already widely used treatment. Not least, it is important for a physician to find and remove the etiology (cause, if known and treatable) of a disease, not just treat the symptoms.

In the last decades causal analysis has become a dominant paradigm for furthering scientific progress in medicine, no less guiding the generation of trusted evidence, especially in the case of some rare diseases that were once considered difficult including to diagnosis (Triposkiadis & Brutsaert, 2025). Nonetheless, there are a range of unresolved issues under debate, e.g. the "frequentist" methods taught from High School days that are commonly used for the generation of evidence in scientific publications. There is, for example, the over-emphasis of the notion of *p-values* and *statistical significance* in many biomedical publications (Ioannidis, 2005). Researchers are pressed to disprove the likelihood of obtaining the data that they have obtained when assuming a null hypothesis that a patient will, for example, not regain health. This is in large part because "classical" frequentist statistics is forbidden from going beyond that likelihood and computing the far more interesting probability of a hypothesis being true given their data. It requires the famous Bayes' Equation, and the notion of a prior belief in something being so before obtaining any data. This has revived the interest in Bayesian statistics that goes back to Reverend Thomas Bayes (1701–1761). Bayes was a teenager when Leibniz passed away. Leibniz and Bayes were both were intensely interested in God as causal (Bayes work was published posthumously, and Bayes was educated at Nonconformist academies known for having a mathematics and philosophy curriculum that must certainly have included the great and often controversial Leibniz). There is a case for arguing that Bayes' work ultimately put some of Leibniz ideas on causality on a more useful and quantitative basis, and this evolution continues today with the impact of uncertainty management from information theory and the study of the statistics of rare and severe ("Black Swan") events in other disciplines. We can expect that EBM is not a static system but also prone to undergo paradigm shifts over time.

[6] A simplified version is captured in the 'evidence hierarchy' with randomized controlled trials (RCTs) at the top of this pyramid, as the most reliable type of evidence for causation, and real-world evidence (RWE) derived from healthcare data as a less reliable type of evidence within that paradigm, for making medicine more scientific.

[7] We only touch briefly on this very important problem here. Note also many new emerging paradigms related to causal modeling in the field, e.g. to extract patterns from 'messy' real-world health data.

While traditional, more ‘bureaucratic’[8] healthcare models often struggle to integrate the patient’s experience and knowledge into shared decision making, cutting-edge digital approaches often aim to find a better balance between the needs of the system and the individual, aiming at a more complete picture around the case (beyond expert epistemologies) that combines information from different sources about the patient history, past decisions, current state and trajectory (as argued by the Agency for Healthcare Research and Quality, CDS Innovation Collaborative, 2025). Note that this may include information in narrative form describing a patient’s lived experience. Such approaches are indicative of *qualitative research*, formally defined as a subjective, non-numerical method used to explore human experiences, perceptions, and behaviors to understand the "how" and "why" behind phenomena. It focuses on in-depth, context-rich, and textual data (interviews, observations) rather than statistical, large-scale data, but still allowing researchers to identify patterns, build theories, and gain deep insights. Key Characteristics include emphasis on non-numerical data, e.g. words, images, and audio, but not statistics. Iterative data analysis occurs simultaneously with data collection, with very small sample sizes focusing on depth over breadth, often reaching data saturation (where no new information *seems to be* gained) rather than needing large, representative samples. However, it is non-trivial to integrate all this expert and patient knowledge to make it available for a patient to reason and navigate all kinds of accessible options that could impact the patient’s health (including but not limited to services offered by healthcare systems).

3 The Uncommon and the Unguided Patient

It is widely acknowledged among health innovation experts that such science-based advances in medicine (in the last 300 years since the days of Leibniz) have strongly improved health outcomes and life expectancy of populations, especially in the last 100 years or so (Wang *et al*., 2020; Deaton, 2013; Oeppen & Vaupel, 2002), in many countries that have been able to build strong healthcare systems for their populations. However, despite all the great science, physicians today still regularly face difficult decisions. Such difficult cases are often about an incoming case not fitting well into current medical categories, reasoning and its science-based evidence landscape[9], reminding us of the challenge traditional bureaucratic systems face as they struggle with incoming cases that do not fit their category-to-process “organizational cognition” (Rebhan, 2025)[8]. In particular, EBM may help to guide real-world cases that are highly similar to the patient populations investigated in those high-quality clinical studies (considering their inclusion and exclusion criteria), but it can’t reliably guide cases that are too different from the characteristic of those studied populations[10]. Determining when a case is similar enough to apply EBM-based clinical care guidelines can still be a difficult decision. A rare disease is generally defined as a condition affecting fewer than 1 in 2,000 people in Europe or fewer than 200,000 people in the US, but while individually rare, they are collectively common, affecting 1 in 17 people (approximately 3.5 million) in the UK and 300 million globally, with over 7,000 distinct diseases identified. Even common descriptions rapidly become combinatorially complex: with a variety of impacting demographic details and diverse therapies involving several drugs, some 20 common morbidities of which 20.7% of adults have

[8] A key property of bureaucratic systems that is relevant here is the tendency to quickly force-fit an incoming case to a set of fixed categories that inform the selection of an appropriate process (Rebhan, 2025). While many incoming cases may fit those categories reasonably well, the atypical cases who do not fit well provide a challenge to the system that it may not be able to resolve well, too often resorting to simple force-fitting of such cases to its categorical system. During system design, it may have been assumed that any incoming case will fit a limited set of categories in this category-to-process bureaucratic logic. Also, incentive systems can hinder other responses.

[9] Imagine a landscape of EBM-based evidence as hills and valleys, with the highest hills are peaks of strong (RCT-based) evidence that can guide care decisions well, while further down, in the valleys, this is not the case. See Fig. 1 also.

[10] A cohort of similar patients that is being studied, here, depends a lot on how similarity is defined, and what we consider ‘similar enough’ to include a case in the cohort. Like in a swarm of fish, some are closer to the center of the swarm, some more on the edge. This may relate to a case being considered guided or unguided. If a patient matches particular RCTs, their inclusion and exclusion criteria may be the easiest starting point.

combinations of three or more conditions and 2.3% of adults have four or more conditions. Added to that, what makes a patient rare or unique is not of course necessarily the disease or diseases.

In the Summary to this paper, we referred to unguided patients as those with medical history trajectories that do not map well to the medical knowledge and evidence landscape and who typically have little recorded data, and who are clearly different from the populations studied in large clinical studies. They do not fit into a "case" in the sense of "case" taught in medical school, i.e. a common pattern of disease and treatment. In the word "unguided" there is a sense of following pathways of treatment as a personal choice or at the initiative of the physician that are not dictated by medicine via approved guidelines, perhaps because no such appropriate best practice exists (or is at least recorded). From time to time in this paper, this definition of unguided patient may seem to vary. Of course, that is in large part because in the broadest perception it not yet seen as an accepted technical term but at very least describes a patient who do not fit in with what might be considered the norm. There are a myriad ways to be unusual. But it also depends on context. In clinical procedures and therapeutics, it often means a physician or patient taking an action (an intended judicious action) without some kind of clinical laboratory result or imaging that would normally be requested. That is, there were guidelines but they were not followed, perhaps for good reason. In mental health, the patient may turn to help for Apps or well-meaning humans who are not medically qualified. In nursing, "unguided" is when a nurse must decide what physical assessments and subjective data to gather without step-by-step prompts or checklists: that usage of the term appears particularly prevalent in nursing education, so "unguided patient" might well be seen by a nurse as a patient with initiative. In addition, at national healthcare level, the word "guidelines" may often have connotations that give the word "unguided" a special force. In many countries there is an organization like the UK's NICE (National Institute for Health and Care Excellence). The NICE is an executive non-departmental public body responsible for producing evidence-based guidelines to improve health and social care (in Switzerland, that role is embedded directly within the government via the Federal Office of Public Health). By a system of study and publicly open opinion by experts or involved persons called "stakeholders" in various medical domains, the NICE provides guidance and advice to the National Health Service to ensure that care is of high quality based on EBM and RCT principles. However, there is also the consideration (often with much controversy and much debate) that treatments be affordable to the national system, i.e. represent good return for the tax-payer's money. It seems a consideration that seems the antithesis of organized care for our typical "unguided patient", who now appears more like a "neglected patient". Interestingly, an initial NICE decision on certain Alzheimer's drugs as too costly was finally out-debated because it did not properly account for the heavy burden on *unpaid carers*, i.e. family members, partners, or friends outside of the formal healthcare system.

In the commonest use by the present authors with their diverse backgrounds, the term most generally refers to a real or simulated clinical scenario where a patient follows a path not structured by standard care protocols, and since the state of unhealth may not fit well into current medical categories, that can include self-guided interventions. The new point being made here is that this may involve feasible herbal remedies or even off-label medical interventions that take the patient very far from Sagan's "Demon-Haunted World" but may draw on the power of the Internet to guide more conformational guidance. By "off-label" (a US term) is usually meant using a licensed medicine for a different disease, patient group (age), dosage, or method of administration than what is listed on its license. It is very common, especially in pediatrics, oncology, and psychiatry, and quite legal. The responsibility is the physician's, but it usually involves discussion with, and recognition of the wishes of, the patient. In practice, in other countries "off label" often simply means "other than a therapy not approved as efficacious for a specific purpose", which arguably overlaps with herbal and home remedies. It is such *unguided patients* (representing many of the above-described complex cases) that are discussed here as a focus application for our digital epistemology. They are cases in which evidence-based decision making in healthcare is difficult, because a) the case does not match well any EBM-based evidence that can guide care decisions[11] (i.e. by not being similar enough to the populations EBM studied), or b) they match multiple guidelines for different populations but then result in conflicting recommendations (Damarell, 2020). For instance, such cases include situations where patients do not match well the genetic, ethnic, or gender composition of the populations studied in EBM RCTs. These foundational studies often exhibit a systemic bias toward certain Caucasian subpopulations that are logistically easier to recruit[12].

Examples for such *unguided patients* can include ethnicities and genders that are under-represented in EBM-based evidence, less-studied rare disorders, complex multi-morbidity cases that are often excluded in, and well-known diseases for which insufficient EBM-based evidence exists to guide care decisions (e.g. what to do once first-line therapy has failed). Also, EBM-based evidence is often biased towards later stages of disease (with strong symptoms) that are easier to study with current clinical research approaches, resulting in major gaps in the EBM evidence landscape when we want to use EBM-based evidence *proactively* to inform decisions about the personalization of *primary prevention,* with the aim of keeping populations healthy (in organizations emphasizing the *preventative* principle of P4, see below)[13].

For such different kinds of *unguided patients* who unfortunately fall into such gaps (or grey areas) in the EBM landscape (Fig. 1), it may not be obvious how to approach science-based, robust decision making in care due to the lack of robust evidence.

Figure 1. The Granular Nature of Healthcare in which Unguided Patients fall as Small Grains between the Main Populations.

[11] If a case (patient with a particular profile) matches EBM-based evidence or not may not be a binary decision (it matches, or it does not), but we would expect many grey area cases in which there is some similarity to the populations EBM focused on but also a concern that known differences to those populations could affect confidence in a decision to follow EBM-guided care (without any adjustment considering such differences).

[12] While EBM does not solely rely on RCTs and there is increasing interest in considering 'real-world evidence' e.g. in situations where it is difficult to study populations with RCTs (Sheldrick, 2023), the EBM community has a clear preference for well-designed RCTs that decrease the influence of confounders, as the most reliable type of evidence. If there is a lack of RCT-based evidence to guide medical decisions then the *unguided patients* challenge present itself.

[13] One important factor here is the ability to recruit participants into studies, if they do not yet experience symptoms strong enough to motivate them to participate in a study. Also, such preventive studies can be difficult to finance.

William Osler's focus on personalized care and oft-cited quote that "medicine is a science of uncertainty and an art of probability" warned of such uncertainty problems ahead but, for patients that do not fall between the gaps, the subsequent development and hardening of many kinds of statistical test might better make it expressed "a science of probability and an art of uncertainty". Nonetheless, "classical" frequentist tests and even current Bayesian methods are not well suited to sparse and sometimes exotic data. We argue that, with increasing digitalization of patient trajectories and healthcare processes, we can record the decisions that were made in such difficult cases (by physicians, patients and others) and, importantly, the outcomes, and then find ways to learn what works best for which trajectory, even if they are unguided patients. As cutting-edge *Learning Health System* approaches are being developed to facilitate such learning based on past decisions, it implies a search for *informative patterns* in these patient trajectories that correlate with better outcomes[14]. That includes complex combinations of big and small decisions made by different actors that correlate with good outcomes. Creating a new innovation area, including a role for Open Science communities, could help healthcare systems learn how to deal with such difficult cases, and show how to enable such learning in a way that fits the special situation of these cases. See Fig. 1 above for a visual metaphor of the landscape we discussed above. We can see areas with dark colors that are rich in evidence that can confidently guide care, and areas with much lighter colors that are less rich in such trusted evidence. An individual patient trajectory (current state and past decisions, history) would fall somewhere in that landscape, e.g. a) within a dark color area, b) nearby, or c) far away from it. Highly similar patient trajectories would be close to each other, considering the sequence of events in the trajectory (for an example, see Kapur et al. (2022)). Outside those dark spots we can see various colors, with the different shades of reddish brown indicating that there is some potentially informative science that may be relevant in healthcare decisions, but it's more patchy, less robust and not at the level of high-quality RCT-based evidence from large studies. Darker brown spots may have evidence from smaller studies, for example. The part of this landscape that is relevant for our discussion on unguided patients is therefore heterogeneous as well, with many colors found in small and large spots with different degrees of uncertainty in potentially relevant scientific knowledge.

How large such *unguided patient* populations are in a typical healthcare population is non-trivial to quantify at present[15], given the lack of attention to this way of framing the problem in previously designed healthcare systems' digitalization efforts, and a lack of consensus on how to count such cases reliably and reproducibly. Is it somewhere between 20-50% of a typical all-comers population entering a hospital for a variety of reasons? Therefore, we try to shine the light of attention on the problem of evidence-based decision making for unguided patients, based on their digitalized trajectories, hoping that they will receive more attention by different innovators in the near future, to further develop the science they need. To help us learn about the extent of the challenge, what common situations look like that we may be able to generalize and disseminate, how new roles can help organizations deal with them before they are force-fitted, and how learning loops can help us improve outcomes and patient / provider experiences. We would expect that with the increased use of EHR (electronic healthcare record) systems, improved digital capture of the patient experience, more sophisticated decision support systems, health data spaces using federated learning and analytics (such as PHEMS[16]), better patient participation in

[14] Basically, the main point of EBM is to use good science to properly distinguish causation from correlation in its evidence base. To optimize care processes for unguided patients based on correlations alone may be problematic.

[15] A baseline approach could map if a patient would have fit the inclusion / exclusion criteria defined for a particular large RCT that provided much of the trusted evidence, for example. However, this does not account for all biases in the composition of the study population in the RCT, e.g. ethnicity, genetics, socio-economic status. Also, often it is a more complex mix of studies that leads to care recommendations e.g. using the Cochrane process https://www.cochrane.org/evidence/why-our-evidence-trusted.

[16] A European health data space for pediatric hospitals that see a lot of unguided patients, see https://phems.eu/

decision making, better digital twins (Emmert-Streib, 2025), and AI-enabled learning from trajectory data, and a combination of those (e.g. using epistemic bricolage), will help us develop new kinds of solutions, together with those who are willing to collaborate on this using Open Science principles (UNESCO, 2021).

We propose that the *digital epistemology* described in this whitepaper (see below for specifics) may help us design an improved *learning healthcare system* (LHS), that aim at finding new, more patient-centric ways of improving care outcomes and experience for *unguided patients*. This would help us find the most *informative patterns* in the patient trajectory data to enable good care decisions for such difficult cases, even if modeling across health data silos is required to enable such learning. Assuming that patient outcomes and insights are tracked well enough to enable well-designed learning loops in such LHS, as originally outlined in the paradigm of *value-based health care* (Porter, 2010)[17], knowing that such learning will be difficult or impossible if appropriate definitions of success and meaningful progress are not well tracked.

In addition to the types of *unguided patients* described above, there is a trend towards more proactive and preventive care models that will add further fuel, increasingly emphasizing 'health maintenance'[18] as an ambition, implying a paradigm shift in organizational and digital design related to healthcare. As mentioned above, the problem of unguided patients is about the lack of evidence within the EBM paradigm, as a basis for science-enabled robust decision making in care. When healthcare systems try to shift to more proactive and preventative care models, they will face a lack of EBM-grade evidence in such relatively healthy populations and early stages of disease with no or only mild symptoms. In principle, it is possible to design an EBM-based strategy to fill such evidence gaps, but at present it is unclear how it would be resourced and implemented within a reasonable time frame (to enable organizations that have prioritized health maintenance and prevention, as in section 3.1).

As this *unguided patients* situation unfolds in front of our eyes, we can see new forms of collaborative transdisciplinary community-based learning emerge, facilitated by new health data spaces similar to the EHDS (European Health Data Space), European Commission (2022) and the Swiss HDS (Swiss Health Data Space19[19]) that allow to do science on patient trajectories. This is across the parts of the patient trajectory that were previously fragmented, hidden in health data silos. Over time, as those data spaces become a reality that is adopted by different users, we may accumulate a variety of (potentially) *informative patterns* found using sophisticated analytics and AI[20] in those complex data landscapes, e.g. combinations of very different 'data points' such as quantified blood-based circulating biomarkers, narrative patient insights and outcomes. That includes SNOMED-coded clinical events and health state models (Rebhan, 2017) in patient trajectories. Such complex *informative patterns* could in turn help to identify and describe particular types of *unguided patient* populations in ways that differ from

[17] VBHC (value-based health care) can only work if outcomes are tracked and used in feedback loops that help improve how care works, and how care decisions are made. Patient-reported outcomes (PRO) are clearly an enabling tool here. However, many healthcare organizations struggle to track outcomes and build solid feedback loops to enable VBHC, for a variety of reasons. Again, this can only work if the organization is serious about it, and has sufficient slack to invest in such innovation (and bringing in relevant expertise to make it work).

[18] Health maintenance is about maintaining physical and mental health in populations, e.g. in many P4 ecosystems (see Appendix). Example: the health of working-age populations has effects relevant for employers.

[19] https://www.digisante.admin.ch/de/swiss-health-dataspace-de

[20] AI is good at pattern recognition and connecting multiple sources of data, which can be difficult for human brains to recognize and connect, making it possible to see informative patterns in unguided patients that human doctors may not see.

traditional approaches for describing populations in medical publications[21]. That is even if such populations do not map well to current diagnostic procedures and diagnostic coding practices[22]. With AI increasingly adopted in that context to aid in the discovery, study and annotation of such informative patterns is plausible (see a few emerging trends described in Allam et al., 2021).

Considering the importance of gaining mechanistic insights into the disease pathobiology that could help to guide care decisions, a promising subset of *informative patterns* would help us model the elusive 'endotype'[23] in clusters of similar *unguided patient* trajectories, and how that endotype model relates to the more easily visible phenotype / symptoms observed in the clinic (or 'at home' or 'at work', in the patient experience). Beyond what healthcare systems record on their side, with their digitalized medical logic about a particular case here, e.g. in EHR systems[24], the integration of patient knowledge, experience and outcomes is likely to be a key factor in building a more complete picture of a patient trajectory; however, considering this may provide a interdisciplinary challenge in epistemology which takes us beyond medical epistemology[25]. Here, we are proposing a patient-centric approach that considers the patient's knowledge as it taps into expert medical knowledge (as well as other knowledge that is relevant, depending on the patient's health goals and aspirations).

Such efforts to record multiple perspectives about a patient trajectory (in a balanced way, from a system and individual's perspective[26]) are then likely to help us find out what the most *informative patterns* are, as this increases the chance of discovering new *informative patterns* that so far have not been considered by those used to focus their attention on medical epistemology only. Aiming to find a healthy balance between system-level optimization (driven mostly by medical professionals and operational optimization experts in healthcare systems, using medical epistemology) and improved participation and engagement of patients and their knowledge, insights, and data[26]. In LHS that emphasize the discovery of such *informative patterns* in *unguided patient* trajectories using a P4 approach aimed at improved health maintenance and prevention, such a balance between system and individual focus will not only inform personalization strategies that learn from what works best, but also how to build trust and solid relationships to enable collaborative progress.

To help us design such a new generation of *LHS* that can help us figure out such *unguided patient* trajectory challenges in a new way, this paper situates Barry Robson's "digital epistemology", grounded in Dirac notation and algebra (Dirac, 1939; Robson, 2007; Robson, 2014; Robson, 2016; Robson & St. Clair, 2022; Deckelman & Robson, 2014), as a design frame, and a philosophical, scientific and technical

[21] Exemplified by the classic Table 1 in medical publications that describes the characteristics of the population that was studied, including demographics, selected quantified biomarkers, and clinical code frequencies.

[22] An example would be the SNOMED approach for diagnostic coding, one of the most comprehensive toolkits. https://www.snomed.org

[23] *Endotype* is the mechanism that is driving the development of the visible part of disease (phenotype) as it shows up as symptoms considered in diagnostic processes. In other words, the biology of disease that is less visible ('hidden from sight') but important for modern medical decision making, e.g. related to disease subtype.

[24] Electronic Health Record digital systems can, in principle, contain patient trajectory data that reflect what happens in healthcare systems, including past decisions and hopefully also outcomes.

[25] As patients are usually not heavily trained in medical knowledge and logic, so when capturing their experience and insights we enter different epistemology that is less standardized, more influences by social context and culture etc. Therefore, the use of natural language is often preferred here to capture insights that are complementary to medical epistemologies and clinical coding systems (e.g. SNOMED).

[26] Cultures that are less individualistic may need a more comprehensive perspective of the social context around the individual. This may also be a factor to consider even in Europe, e.g. with migrant populations. See also Seebode (2013) for a patient-centric perspective on this.

successor of Leibniz's dream[27]. A number of papers by Robson and collaborators are cited here and below because they represent the development and application of the language and address a number of points relevant to the course of the discussion. In the medical domain, with its traditions of generating and using scientific evidence for robust decision making, focusing on *unguided patients* and related decision making (i.e. cases where EBM-grade evidence is lacking) new paradigms are needed to build outcome-based learning loops in LHS.

Note that, while Barry Robson's *digital epistemology* may also be useful outside the medical domain (in non-medical domains that can benefit from improved 'epistemic humility' in the design of similar kinds of learning digital systems[28]), our discussion here is restricted to applications in the medical domain. Specifically, our scope here is the discussion of applications of this *digital epistemology* in patient-centric *learning healthcare systems (LHS)*, which aim to a) improve care for patients that fall into gaps in the EBM evidence landscape, i.e. *unguided patients*, b) shift to more proactive and preventive care and health maintenance for their populations by learning how to personalize preventive care, using the P4 approach. Such applications of our *digital epistemology* need to consider how systems and people learn in a digitally connected care organization[29], increasingly supported by federated *health data spaces* that enable the discovery of *informative patterns* across traditionally separated health data silos. This brings new possibilities for more complex cross-silo patterns to emerge that are relevant for improving the care for *unguided patients*[30]. Note that both a) and b) come with the requirement to deal with different kinds of uncertainty including *epistemic uncertainty* (imperfect knowledge in the context of evidence-based decision making; see Appendix).

As healthcare systems can be very different from each other, e.g. in different countries and regions, we use the example of Switzerland to illustrate this further. Considering that Switzerland has a particular healthcare system, population, decentralized culture, federated tech scene, silo'ed digitalization of patient trajectories to overcome, and pioneering community culture in medicine. Relevant aspects in that Swiss landscape include *LHS* efforts that aim to track patient trajectories with outcomes for unguided patients, including related national and regional digitalization programs, AI-enabled digital decision support in healthcare, and a variety of relevant technical / scientific efforts. *P4 ecosystems*[31] that combine all those aspects to enable a shift towards more proactive and personalized care will be given special attention (see the Appendix for a definition of *P4 ecosystems*, beyond classic notions of P4 medicine in the Systems Biomedicine literature[32]).

---

[27] We approximate the 'essence' of his dream here as a form of reasoning across 'knowledge atoms' understandable to both humans and machines, enabling an epistemology that has a high degree of 'epistemic humility' in the sense that it knows its own imperfection well and then learns how to improve by updating its beliefs (in the sense Bayes' once described it).

[28] e.g. to better treat unguided cases who do not fit "organizational cognition" well, i.e. cases that do not fit the set of categories used by an organization for classifying incoming cases: see Rebha (2025).

[29] People here include healthcare professionals and non-experts (patients, their families, others).

[30] Healthcare systems are often fragmented, and so are their data landscapes. By enabling cross-silo modeling of a patient trajectory, we have a more complete picture of that trajectory and a better chance of finding patterns that we missed before.

[31] P4 ecosystems in this context are basically a subtype of learning healthcare systems, focused on making healthcare more P4 i.e. more preventative, personalized, predictive and participatory. Aiming for improved health maintenance at population level. See the Appendix, section 3.1 and https://www.linkedin.com/pulse/p4-ecosystems-recent-progress-2024-michael-rebhan-x6mqf/

[32] Original definitions of P4 medicine emerged from the Systems Biology and Systems Biomedicine academic community around Leroy Hood, in the early days of genomics, proteomics and related technologies that promised to provide insights into human patho-biology. Here, our perspective on it reflects our view on progress made in 20 years, considering translation of such P4 science into healthcare systems, organizational aspects of P4 ecosystems, startup-led ecosystems and consumer-centric P4-related services.

4 From Leibniz' Vision to Hyperbolic Reasoning over Imperfect Health Knowledge

Leibniz's goal was twofold: to create a universal symbolic language (*characteristica universalis*) and a formal calculus (*calculus ratiocinator*) that could encode reasoning across domains. He envisioned modular, recursive symbols representing *knowledge atoms* (Leibniz's indivisible units of understanding)[33], logical operations performed via manipulation of these atoms, and the resolution of complex reasoning problems through symbolic computation. As a mathematically grounded polymath, he tried to generalize and connect across the emerging scientific epistemologies that were accessible to hi at the time (Milkov, 2006). Leibniz was further encouraged in later phases of his effort by letters from Bouvet about 'ancient reasoning over symbols'[34] that may be considered a previous polymathic attempt at such 'computation', quite distinct in the 'way of thinking' from the epistemologies and mathematics accessible to Leibniz in Europe (Berkowitz & Cook, 2015).

Interpreting Leibniz' epistemological vision in its historic and cultural context, we can note that many intellectuals in his age in Europe had a fascination for sophisticated designs behind seemingly complex machines, the role of precise computation in it, enabled by progress in mathematical calculation and scientific reasoning, as engineering developed as a field that could create increasingly complex machines with such paradigms. In other words, we can imagine his generation of European intellectuals as 'owning' those scientific paradigms, being inspired by all the good things they will create for humanity. In that historic context, we can imagine his vision as a kind of generalization of the available calculus that would work beyond these engineered systems, to help humanity solve many other important problems related to the use of knowledge and reasoning. Such a more general calculus should then allow those who master it to cross many different areas of human knowledge, including intersectional problems where different *epistemological fields* and their languages and symbols overlap (as discussed above). In the context of our paper here, this is the interpretation of his vision that we propose, knowing that there is considerable debate about the meaning of his vision (Milkov, 2006).

In medicine, knowledge and reasoning often taken an 'Aristotelian approach to truth', by forcing a decision about a statement to be either 'true' or 'false', i.e. they use Boolean logic (Alencar, 2024). For example, a clinical *lab value* measuring a blood-based circulating biomarker (informing about a physiological health state) is either below or above a threshold that separates 'normal' from 'abnormal' values, a simplification designed to streamline decision making in hectic work situations and minimize required training for professionals. Is a particular symptom 'present' or 'absent'? Due to its simplicity such Boolean logic is popular, easy to learn, disseminate and implement, and therefore omnipresent, including in digitalized decision support systems that are widely used. An Open Science example for such Boolean logic being implemented in a clinical decision context and related decision support systems is CQL (*Clinical Quality Language*), a project developing a standard that medical professionals and digital

---

[33] His notion of 'knowledge atoms' and 'logical operations' has to be interpreted in its historic context, considering that in his time, such language was used more 'playfully' from a modern perspective, while meanings have narrowed in our time. In his age, there was a widespread fascination to new possibilities enabled by mathematics and, in particular, 'algebraic thinking' with symbols and logic in many domains of life.

[34] Bouvet, at the time, had just gained access to precious documents held by scholars near the Chinese emperor, as he proposed a cross-cultural intellectual exchange between European science culture (math-based) enabling precise engineering, while the Chinese scholars helped him understand what they cherished in Asian ancient logic e.g. in the form of hexagram-based cosmological logic. Leibniz was excited about the exchange with Bouvet on this, as it seemed to fit well with his new 'binary logic' that later enabled digitalization of knowledge (encoding information as 0 and 1 sequences, or, in Chinese, as 'yin' and 'yang' encoding). From a contemporary perspective his understanding, through Bouvet, of that hexagram logic may have been superficial in some parts, as Leibniz had limited information and was focused on finding something that could make his binary logic more universal. One interpretation of this exchange with Bouvet is that it further encouraged Leibniz to pursue his epistemological vision, even if it was difficult.

systems in medicine can understand[35], including those working in low-resource environments which may not have access to the latest technology for clinical decision support.

When facing health-related decisions that go beyond such Boolean logic, probabilities are often used to express a degree of uncertainty about a situation, e.g. a probability of developing severe complications after a surgery. Those probabilities typically are somewhere between 0 and 1 (or expressed in %, i.e. between 0 and 100%). In terms of probabilistic reasoning over medical knowledge, in a common interpretation, *Bayes Nets* can express conditional probabilities where the probability of an *event A* occurring can change if *event B* occurred first. See e.g. Ni *et al*. (2010). There is a conditional probability, regarding those two events, written as P(A|B). For example, if a patient is diagnosed with type 2 diabetes (event B), the probability of a diagnosis of diabetic renal complications (event A) after x years in that patient may have increased. In practice, a causal process is not necessarily observed and may or may not have occurred (ignoring any other evidence). Rather, one has examined the population of type 2 diabetic patients and counted the fraction of that population that has renal complications. A sense of what occurred first and common sense understanding of mechanism may, however, justify a causal interpretation. In many situations such *Bayes Nets* composed of such conditional probabilities can express medical reasoning relatively well, so they have found their way into many decision support systems, to complement what can be expressed with simple Boolean logic (e.g. in the form of *decision trees)* (Ni et al., 2010; Lucas, 2001). However, there are fundamental problems, related to the previously mentioned important distinction between causation and correlation (how to properly capture it), and the difficulties when using such probabilistic approaches to encode knowledge about describing bidirectional influences between event A and event B (event A influencing the probability of B and vice versa). For example, sampling gives the probability of being type 2 diabetic in overweight patients as about 30%, and the probability of being overweight in type 2 diabetic patients as about 83% (depending on demographics), suggesting, but not alone proving, that causality works in both directions, but not to the same degree, and not predominantly in the direction as is traditionally thought. The distinctions of causality and correlation are easier to see in chemical kinetics: the fact that compound A, B, and C finally occur in certain stable proportions is a matter of relative free energy and so *thermodynamics*, but says nothing about the *kinetics* in each direction (matter of crossing energy barriers), nor what formed first. Such considerations occur throughout metabolism and molecular biology, and these can change in health and disease, so it is to see that an examination of the interplay of conditions such as type 2 diabetes and being overweight requires detailed laboratory studies, not just sampling.

Attempts have recently been made to develop new knowledge and reasoning frameworks for health that would overcome the well-known limitations of Boolean logic and Bayes Nets for robust decision making in increasingly digitalized healthcare contexts (Lucas et al., 2004; Hunter & Williams, 2012; Friedman et al., 2014). Here, we argue that the problem of improving evidence-based decision making for *unguided patients* in *learning healthcare systems* is especially relevant, in the context of such efforts. This is to enable faster system-level (organizational) learning around the above-described *informative patterns* in patient trajectories, including those enabling endotype modeling. Such cases require special attention, in terms of our ability to combine many different and distributed weak signals that are hard to interpret using simpler medical logic.

The "*digital epistemology*" (digital health informatics) that we propose here for such purposes is based on the work of *Barry* Robson[36] and collaborators, building on earlier work by *Paul Adrien Maurice*

---

[35] https://github.com/cqlabio/cqlab-typescript, https://cql.hl7.org/

[36] As the founder CEO of *The Dirac Foundation*, Barry Robson was inspired by Dirac's earlier work, and worked on its extension into the *digital epistemology* described here, including Q-UEL and the Hyperbolic Dirac Net (HDN) as key pillars of it.

Dirac (1902-1984) on what is often referred to as *Dirac notation*. Inspired by the ability of Dirac's approach to encode more complex but real knowledge situations such as the representation of probabilistic and sometimes seemingly contradictory states in quantum mechanics, which would not fit the above simple logic (restricted to Boolean logic and Bayes Net style probabilistic reasoning), Robson et al. developed a language for encoding complex situations in medical decision making that could not be represented well with a combination of Boolean logic and Bayes Nets. The result of those efforts, the knowledge-encoding language Q-UEL *(Quantum Universal Exchange Language)*, encodes medical 'knowledge atoms' in a particular *hyperbolic vector space* (or *split-complex* vector space), based on split-complex vectors and their scalar products.

While any proposed standard can be expected to undergo considerable change, the basis of the approach has already been tested and refined over the years, and mathematically there is even some case for arguing that it simply represents well-established and well-proven quantum mechanics with something "legally removed". Hyperbolic-complex (or split-complex) values contain an imaginary ***h*** number rediscovered by Dirac (initially as a linear operator σ, later as γ-matrices) that squares to *plus* one. Contrast the more familiar imaginary number ***i*** that squares to *minus* one. This ***h*** in some form represents the basis of the Quantum Field Theory (of particle physics), but in the absence of ***i*** it makes probabilities behave as in the everyday world of human experience. "Removing" ***i*** is legal because quantum mechanical probability amplitudes can be purely real, or equivalently one may consider operators that return probabilities as expected values of measurements, or the ***h***-complex-valued spinor projectors of Quantum Field Theory as operators that return real values as eigensolutions (Robson & St. Clair, 2022). Initially, this was to facilitate the construction of *bidirectional general graphs* of probabilistic knowledge representation for use in inference and prediction, called *Hyperbolic Dirac Nets* (HDNs) (Robson, 2014; Deckelman & Robson, 2014). However, Q-UEL has been developed as a full intuitive system for managing probabilistic knowledge that conveniently maps to natural language (e.g., Robson, Caruso, and Balis, 2013; Robson, 2015; Robson. 2016), as well as to mathematical computing including quantum computing (Robson & St. Clair, 2022). While the language is naturally capable of managing certain knowledge, special attention was given to situations in which knowledge is incomplete, weak or conflicting, and hence difficult to use for medical decision making, based on probabilistic knowledge networks as bidirectional general graphs.

This paper is not primarily about learning and prediction and certainly not about mathematics in any detail; hence, readers from an AI background who think in those terms and are not familiar with the literature of Q-UEL and/or its associated HDN may well note that the present paper lacks formal definition of an inference engine or any discussion of a training algorithm with objective/loss functions. They may therefore think it is assumed that, because the mathematics works for quantum physics, it will naturally solve uncertain clinical reasoning. They may feel that it lacks mathematical proof beyond the brief introduction in the Appendix, so that a bridge connecting an elegant mathematical-physics representation to actual predictive or diagnostic utility in healthcare is missing. It may even be concluded that processes analogous to learning, inference, and prediction are not the forte of Q-UEL and not even the purpose of it. It should therefore be emphasized that these are some of the important applications of Q-UEL and that there is a substantial body of study on those issues. There are some 33 peer reviewed papers specifying and describing Q-UEL and the HDN, with collaborators from many universities and organization including the Linux Foundation (Robson & St Clair, 2022), with 1,190-1,200 citations at the time of writing. They include several ML and DL comparative studies including blind studies, many diverse use-case studies, experiments, baseline comparisons against standard probabilistic and graph-based models. They use real world digital patient health, genomic, health insurance, socioeconomic health, environmental toxicology, cell type, adverse drug reaction, and drug discovery databases. A sample of these papers given as a bibliography in references, cited and described in this text, and these include further references to Q-UEL and the HDN for automated probabilistic reasoning. Algorithms such as POPPER combine symbolic and arithmetic reasoning processes (Robson, 2015). Some algorithms such as XTRACTOR autosurf and mine the Internet for knowledge, and module MARPLE may be the first algorithm, in 2014-2015, to pass with high marks a medical licensing exam: see e.g. Robson and Boray (2016). Medical data analysis for Big Data runs on Amazon Web Services. In addition, Q-UEL and the HDN also make extensive use of a Theory of Expected Information first developed by Robson (and now integrated with Dirac algebra) that has and even older and more highly cited heritage.

A key design feature of Q-UEL's associated HDN, which is expressed in Q-UEL, *is* to avoid the constraint of Bayes Nets being unidirectional in their conditional probabilities, in a causal interpretation, changing from event B occurring first to event A occurring first (Polotskaya, 2024)[37]. The ability to encode *bidirectional* influences, from event A influencing B and *vice versa*, is considered important here, in the design of Robson's digital epistemology as Q-UEL (Robson, 2014). Following the well-known adage attributed to Sir Ronald Fisher that correlation does not alone imply causation, it is even more important to separate out both directions of conditionality to explore which one (or both) is causal of the other. Moreover, it ensures that the probabilities in the inference net are *coherent*, in large part meaning that the equation known as Bayes' Rule is obeyed (Robson, 2014). Applications of this *digital epistemology* will thus enable us to go beyond unidirectional probabilistic reasoning in medicine as enabled by Bayes Nets. For example, in chronic disease management, e.g. in cardio-renal-metabolic diseases (see 3.2), feedback cycles in the form of patho-physiological loops involved in what makes co-morbidities difficult to manage in healthcare cannot be properly represented with Bayes Nets. Consequently, we propose this *digital epistemology* as the logical next step in improving medical reasoning and epistemology, after the shift from Boolean logic to Bayes Nets style probabilistic reasoning, enabling the construction of learning loops in LHS architectures that learn from the trajectories of unguided patients. This serves to find ways of improving outcomes and experience for such patients, even if EBM-grade evidence to guide their care is relatively weak.

Our *digital epistemology* thus consists of the following core elements:

• **Mathematical foundation**: Barry Robson's extension of Dirac notation into hyperbolic-complex (split-complex) vector space enables us to go beyond combinations of Boolean logic and Bayes Nets (Robson & St Clair, 2022), as a foundation of Q-UEL and HDN

• **Q-UEL**: The language for encoding medical knowledge, e.g. a patient's state and potential hypotheses, decisions or actions related to that patient's state (Robson, 2007, Robson, 2014)

• **Hyperbolic Dirac Nets (HDN)** enable reasoning over knowledge encoded in Q-UEL (Robson, 2014)

• **Automation**: Processes for automatically converting existing medical data into Q-UEL and HDN, including automated quality control (Robson & Boray, 2018)

• **Digital systems** providing user-friendly interfaces to such integrated, complex knowledge and reasoning, to facilitate the discovery and validation of *informative patterns* related to unguided patient care, using the Q-UEL + HDN digital core (including patient-centric agentic AI)

• **System designs** that link all the above to human knowledge, behavior, decisions, education and other human activities that can affect a person's health state

Recall Leibniz' vision about 'knowledge atoms' (that he sometimes called "monads"), and logical operations and reasoning using symbolic computation over these knowledge atoms. Q-UEL basically provides the language for encoding and symbolically manipulating medical knowledge atoms and their relationships, while HDNs perform the arithmetic 'calculation' (reasoning) over the knowledge atoms.

[37] Cp. Joshua Pearl's do-calculus, as it sets probabilities to 1 (assuming certainty), 'what-if studies' in medicine.

Leibniz strongly emphasized that his atoms of language, and alphabet of human thought, had fundamental, *structural relationships with one another*. These atoms were not viewed as isolated entities but as basic, irreducible concepts that could be combined through relationships to form complex thoughts, to some extent mirroring the way numbers are formed by multiplying prime numbers. The commonest entity in Q-UEL is essentially Dirac's important "bra-operator-ket". In the present context, it emphasizes relationships between "atoms" in expressions in a form that looks like a simple sentence in a subject-verb-object language, and so easily readable by the human eye.

< *subject expression* | ***relationship expression*** | *object expression* >

The expressions are typically logical expressions. However, the important thing is that they represent states or situations often involving several attributes (logical constants or variables) that are observable and countable states or situations from which we can compute probabilities, or about which we can hold comparable degrees of belief. For example, there may be a situation in which A and B or C are seen, and not D, is both observable and countable. Each expression is formally represented by at least one attribute, exemplified by patient ethnicity or blood pressure measurement, separated by logical or other relationship operators.

Leibniz' vision was not focused on medicine as a domain, as our *digital epistemology* is, but it is quite possible that with some adaptation, similar approaches may also develop outside the medical domain. For example, there have been some use cases for Q-UEL shown in financial market analysis and regarding socioeconomic issues of equality. It is evident that Q-UEL comes very close to Leibniz if the variables and constants in the (typically logical) expressions (and separated by operators) are considered as Leibniz' atoms or small clusters of them, in which case it is interesting to note that they relate to noun or verb clauses, and to nodes or clusters of nodes on a graph. Moreover, Leibniz appears to have considered his atoms or monads as the most basic elements of human understanding, such as *being*, *same*, *different*, *true*, and *false*, and it is certainly evident that these represent words that are not meaningfully divisible. This comes very close to the elements of Q-UEL's *Attribute Metadata language* (AML) with e.g. nouns, verbs or prepositions alongside article, qualifier, and quantifier words that Q-UEL uses to form its entities analogous to noun and verb clauses. To strengthen that link, and recall the analogy of Leibniz' atoms to prime numbers, note that what are now considered as Q-UEL attributes can be usefully encoded as prime numbers, with many algorithmic benefits (Robson, 2015). Leibniz also explored that idea and considered hierarchical order or dominance as a relation where one monad "dominates" another if it contains the framework or reasons for the actions of others, in effect a role played by Q-UEL's *metadata operator* := in which the structure to the right is a more specific example (subset, element or value) of what lies to the left. However, we do not explore these concepts from Leibniz' perspective any further in this paper.

A Learning Health System (LHS) in healthcare is typically considered as a model in which data, technology, and culture are aligned for continuous improvement, allowing providers to learn from routine care and enhance patient outcomes in real-time. In principle, this is not necessarily learning in the sense used in current AI, but techniques of data analytics such as data mining and even more traditional kinds of statistics, though any approach must be able to handle sparse and data and much uncertainty in a way in which uncertainty is meaningful and ideally quantifiable. One needs to transform daily, routine care into reusable knowledge to improve quality, safety, and efficiency, rather than waiting for traditional, slow research cycles. While our *digital epistemology* (described in this whitepaper) is a science-based paradigm for building patient-centric LHS, we noted that in the humanities the notion of a digital epistemology can have a different flavor, that may be complementary – e.g., design, epistemology is grounded on process; this is an active and practice-based form of enquiry producing different forms of knowledge, tacit, experiential, and analytical (Page, et al, 2025). For example, Ingvarsson's *digital epistemology* notion that views knowledge as a complex living web with many different kinds of relations, including causal effects that are a focal point for EBM, but also many other kinds of relations that help human cultures adapt to environmental changes, shift "organizational cognition" and respond to technological disruptions (Ingvarsson, 2021). Such '*epistemic bricolage'* across epistemologies in science and the humanities may help to find ways to balance the medical system's perspective and the individual's (patient) perspective in healthcare, as outlined above.

The way knowledge is digitally encoded, to better capture the real complexity of difficult cases and their trajectories, is important here. To avoid forcing reality into an over-simplified logical frame that does not work well for robust scientific decision making, for unguided patients. Q-UEL (Robson, 2014; Robson, 2016; Robson & St. Clair, 2022; Deckelman & Robson, 2014) is a language that was designed to enable such progress. Q-UEL manages uncertainty, by associating each statement of knowledge with a pair of bidirectional probabilities encoded in a single *split-complex scalar value* [38], thus overcoming the unidirectional and acyclic constraints of Bayes Nets[39] as discussed briefly above.

With this, our *digital epistemology* allows the integration of health-related knowledge, symbolic logic and probabilistic reasoning, in a vector space representation of belief systems[40] that can evolve through further observation. The architecture is designed to learn, very much in the spirit of Bayes' original thoughts regarding the ability to 'update beliefs' with new data. It accommodates different kinds of uncertainty and contradictions in the way knowledge is captured and utilized, thereby enabling increasingly robust decisions in Learning Healthcare Systems (LS) without forcing over-simplification or arbitrary constraints when modeling complex clinical realities. In terms of the different sources of knowledge our language can handle, it can also use knowledge that was automatically extracted from natural language, which can provide a much larger amount of (potentially lower-quality or mixed-quality) knowledge than is typically available in highly structured form, e.g. as knowledge triples encoded in RDF[41]. Learning, here, in this digital epistemology, is modeled as a transformation on vectorized knowledge atoms, e.g. of a patient's current state and a potentially fitting diagnosis, risk or other care decision. Note that in its design it does not carry a bias towards certain types of knowledge but allows for considerable flexibility in the incorporation of the patient's experience and outcome expectations, in addition to the medical epistemologies used by healthcare organizations in their operations. Being able to handle all these different sources of knowledge is important, as it can increase our chance of finding *informative patterns* to inform the care of unguided patients, e.g. by considering distributed weak signals that may otherwise be overlooked in classic frameworks.

For a deeper discussion of underlying mathematics in our *digital epistemology*, the role of the split-complex (hyperbolic) vector space, and practical additions since Dirac's original work, see Deckelman & Robson (2014) and Robson & St. Clair (2022). For a gently narrative introduction without any mathematical notation, see our Appendix (those who are very comfortable with mathematics may want to

---

[38] As a reminder, recall the complex number based on ***i*** that squares to -1. Emphasis is now on the split-complex number ***h*** (rediscovered by Dirac) that squares to +1. It was rediscovered by Dirac and is fundamental to Quantum Field Theory, and when used in the absence of ***i*** it allows Dirac notation to be applied to the everyday world of human experience avoiding ***i*** which leads to the notorious predictions of quantum mechanics as wave mechanics. An important point now is that the mathematics has been found useful by other authors in certain AI-related areas, such as ***h***-complex neural networks that can solve the XOR problem in a single neuron, and recommender algorithms (e.g. on dating sites). It is no longer "weird".

[39] As with Dirac's approach, probabilistic elements of knowledge expressed with bidirectional probabilities can also be built up from vectors and matrices in which the elements have, in this case, ***h***-complex values.

[40] The system 'holds itself up' because the *representation* of states (basis vectors) is defined by the very scalar overlaps they compute, and the operators acting on them are defined by their action on that same basis.

[41] https://www.w3.org/TR/rdf11-concepts/ - It also allows the distinction of assertions extracted from natural language from other types of encoded medical knowledge, treating it in a consistent probability-theoretic way according to Popper's principle (Robson, 2015)

skip that part in the Appendix). Note that related mathematical approaches such as *Clifford Algebras* (Breuils et al., 2022) and other hyperbolic vector spaces may provide alternative paths to solutions. In recent AI research it was found that such hyperbolic (complex) spaces have clear advantages over classic Euclidean spaces, e.g. regarding the learning of hierarchical representations in AI[42] (Nickel & Kiela, 2017).

5 The N-of-1 Trial

First, recall the EBM paradigm and its preferred RCT-based evidence generation approach to new knowledge, noting that it helps to improve current decision logic in healthcare by updating Boolean logic (e.g. thresholds or new options) or Bayes Nets in decision support systems. Such EBM evidence is often summarized by medical expert communities that develop clinical guidelines for specified populations to facilitate the use of EBM evidence in clinical practice (Higgins et al., 2024). Developers of decision support systems then often try to encode such logic in their systems to make EBM evidence and derived guidelines (including algorithms) easily available in healthcare contexts when it is most appropriate and fits a particular case and decision. But as *unguided patients* do not benefit much from EBM evidence to inform care, the question of how to design an LHS and its corresponding digital infrastructure comes to the forefront.

In this context, in the last 10 years or so, the scientific paradigm of *n-of-1 trials*[43] is emerging as a potential future extension of well-validated EBM methodology, with potential relevance for at least some unguided patients (Kim-McManus et al., 2024). This methodology toolkit for personalizing care for atypical cases was designed with patients in mind who do not fit EBM-guided care logic well (Lillie, 2011). Here, several treatment switches can be combined with biomarkers (or other observations) that indicate positive or negative treatment response to find the best treatment approach for that particular patient, considering their complex trajectory, history, and special needs. For example, if the patient is quite different from the populations studied by EBM, the n-of-1 trial approach may be considered appropriate to achieve better personalization of care. From a statistical perspective, it includes an ambition to find ways to move from the discovery of *informative patterns* as pure correlations (and related endotype modeling) towards insights into causation, e.g. using Bayesian modeling (Samuel et al., 2023). This could be an area of emerging synergies with the digital epistemology described here, as such n-of-1 trials accumulate precious longitudinal data (i.e. trajectories) about a patient's response to different treatments, if biomarkers are available to quickly assess treatment response before the next treatment switch happens. This may facilitate the discovery of *informative patterns* with a potential re-use in other patient populations, e.g. those with similar endotypes that influence treatment response and biomarker profile.

N-of-1 trials do, of course, imply very sparse data, notably single or very few observations. We must face the fact that 1 out 5 occurrences carries less information than 100 out of 500, but also recognize that a great deal of weak and circumstantial evidence in a criminal trial can add up to outweigh the judgement that would be made without it. Conversely, we should also note that even zero observations are informative if we expected to see several occurrences. Research into approaches for the automatic

---

[42] Their exponential volume expansion naturally matches the exponential branching of hierarchical structures — which is why hyperbolic embeddings are effective for hierarchies.

[43] N-of-1 trials are less about average effects in populations as in many EBM studies, but more about the personalization aspect of P4, i.e. that particular patient with that particular profile even if it does not fit well into classic EBM methodology, guidelines or clinical research traditions. When using multiple treatment switches and biomarkers that help to understand efficacy and safety trade-offs quickly, Bayesian-style probabilistic modeling can be used to increasingly tease apart causation and correlation. How to best extract informative patterns from a collection of n-of-1 studies is an active area of clinical research. In the context of the proposed digital epistemology discussed here, the importance of epistemic humility being properly captured in the learning healthcare system, as well as the belief updating aspect may be focal points for synergy verification as data from n-of-1 studies accumulate.

conversion of existing medical data in common formats into Q-UEL and HDN have shown that a method based on *information theory* and particularly *Expected Information Theory* (Robson, 2005; Robson, 2015) promises to be able to transform complex data landscapes with both plentiful and sparse data components into the bidirectional relations encoded by Q-UEL in a robust way, e.g. using the k-method (Robson, 2014). This approach could also help enable the discovery of new *informative patterns* that combine weak signals, while imposing fewer assumptions and beliefs than classic methods, thereby facilitating the discovery of informative patterns to improve unguided patient care and the interpretation of N-of-1 trial data.

To help a patient navigate a great variety of options, knowledge and potential decisions, we are currently in an early stage. However, as new LHS are entering the phase of system design (e.g. see section 3.1), this may provide an opportunity to upgrade their digital foundations as well, once it is clear what level of emphasis the new system will have on *unguided patient* trajectory problems outlined here. For example, such a system may find a balance between a) areas where classic approaches work well and can be deployed also in the new organization, e.g. where EBM-grade evidence is strong, b) areas where new approaches such as the one outlined here are needed to improve the care of unguided patients, allowing safe prototyping while monitoring the decisions made by physicians and patients, and their outcomes. For safety reasons it would therefore be advised to first build further confidence in the approach before influencing any medical decisions using the approach described in this whitepaper. On the other hand, patients are already using a variety of tools available to them to navigate health-related decisions, so the question poses itself if what we propose here could provide added value, e.g. by helping them make better decisions, build more solid health-related knowledge and improve health outcomes using an approach similar to n-of-1 trials.

6 The Swiss Landscape

Switzerland has been selected as an example context here for several reasons: a) the emergence of the new *Visana VIVA* preventive healthcare model and the synergistic 8P innovation framework (Bernier, 2024) fits the above-described problem of *unguided patient trajectories* in a preventive care context (see 3.1 below), as it provides an example of a new system design for P4-based healthcare[44], b) the SPHN (*Swiss Personalized Health Network*) collaborative communities of physicians, scientists and patients, and their approach to the collaborative, inclusive design of prototype *LHS* for unguided patients through patient-centric research initiatives, c) nationally-coordinated efforts in the next 10 years coordinated by *digiSanté*, including the planned Swiss Health Data Space (SwissHSD), as they may increase chances of finding new kinds of *informative patterns* across health care data silos to inform the care of *unguided patients*[45], d) progress made in Oncology centers in Switzerland in improving care for unguided patients, using a multi-omics rapid endotyping approach[46], e) the emergence of a highly collaborative and transdisciplinary pediatric community in Switzerland, as it develops a more holistic, science-based approach to the personalization of care including diagnostic challenges related to rare

[44] It is a LHS in the sense described here, matching the 4 principles of P4 (make medicine more preventive, personalized, predictive and participatory) to maintain the health of the population. This entails a high level of attention on unguided patient trajectories, as most EBM-grade evidence is focused on classic healthcare models when patients are more sick with more severe symptoms.

[45] https://www.digisante.admin.ch/de/swiss-health-dataspace-de

[46] https://eth-nexus.github.io/tu-pro_website/. Rapid endotyping, here, is performed through the fast analysis of a tumor biosample from that patient, using a combination of different lab-based and computational methods, including genomics, transcriptomics, proteomics and other omics technologies that help us understand the endotype that is driving disease progression. See also the personalized combination therapy design approach tested in the Rapid-01 clinical trial in Switzerland, using *pharmacoscopy*.

disorders and n-of-1 trials[47]. Also note that many of the authors on this whitepaper are Swiss residents, and through their work are familiar with different aspects of the Swiss situation that is relevant here.

Even beyond unguided patients, Switzerland represents a need and an ultimately appropriate challenge. Traditionally, the casual observer outside Switzerland expects the Swiss healthcare system to be an integrated nationalized healthcare system somewhat like the UK's NHS. This is not the case. The Swiss healthcare system is decentralized rather than fully integrated. While the entire system is regulated nationally by the Federal Office of Public Health (FOPH), the 26 cantons have primary responsibility for organizing and delivering care. Instead of a single, centralized network, it relies on a highly competitive, regulated private siloed market with private insurance mandates rather than government-run clinics. There are difficulties sharing records in different natural languages, and while HL7 FHIR is prominent as a computer language for medical record and patient data representation, it does not extend much beyond that, is not universally accepted, and not always equivalently implemented. Switzerland is thus not in a dissimilar state to that of concern in the US; see reference President's Council of Advisors on Science and Technology (2010) and note that the interoperability problem in the US is not fully resolved today. However, in Switzerland there is progress, as follows.

7 P4 medicine and unguided patients, a new frontier in Switzerland and beyond

The above is not to say that Switzerland lacks any notions or efforts at greater interoperability. The VIVA care and insurance model, designed by Swiss Medical Network (SMN) and Visana (a health insurance organization), implemented in three different rural regions in Switzerland, organizes integrated care delivery using a health system design that emphasizes preventive care, early detection of disease and health risks, and early intervention strategies, to keep people healthy, and avoid bad health outcomes where possible. In other words, this model emphasizes health maintenance (including preventative medicine) at population level for the population that signed up to the VIVA model. It uses a bespoke health economic model (using capitation) to improve the alignment of provider incentives to this organizational goal. In this context, Bernier et al. (2024) have developed the 8P frame, to extend the classic P4 principles (to make medicine more *Predictive, Preventive, Personalized, Participatory*) by adding 4 additional system design principles (*Purpose, Platform, Performance Sharing and Pooling*). This provides an example for a paradigm shift in organizational design in health, which entails the challenge of finding smarter ways to personalize and deploy preventive care at population level, a new competence for Swiss healthcare systems, at that scale[48]. Note that organizations outside Switzerland have a longer history of implementing such models, e.g. Kaiser Permanente in the US[49].

As outlined above, such efforts require a bespoke digital foundation (custom digital infrastructure) to enable robust decision making and system-level learning, even in cases that are not well-informed by EBM-grade evidence (and EBM-derived guidelines that encode logic and algorithms derived from that evidence). In this context, a patient-centric *LHS* has to figure out how to implement evidence-based algorithms that guide care decisions in a way that work in that particular population and setting, in addition to all those cases that would be considered atypical or not fitting into classic diagnostic categories at all. When shifting towards more proactive care, health maintenance and prevention, this organizational shift

[47] https://swisspedhealth.com/

[48] A 'smarter way' in this context is essentially to build an 'organizational cognition' (and ability to learn at system level) beyond classic category-to-process bureaucracy systems (Rebhan, 2025), which tend to force-fit incoming cases (patients) into a set of categories that match the more frequent cases. In medicine, finding this 'smarter way' is important not only for patients and their outcomes, but also for the sustainability of the healthcare system itself with its limited resources.

[49] https://healthy.kaiserpermanente.org/washington/health-wellness/preventive-care

beyond the ‘safe zone’ of EBM-guided care captured well in high-confidence clinical guidelines becomes even more acute as EBM-level evidence tends to be weak in early stages of disease when symptoms are absent or quite weak.

Stated alternatively, such complex, atypical patient trajectories (Allam et al., 2021) from unguided patients (as defined above) and P4-style preventive care organizational transformations (personalized preventive care, such as the above examples) push us beyond established EBM paradigms for robust evidence-based decision making in health. See also e.g. Froehlich*et al*. (2018). This is pushing us to explore new paradigms that are based on science, transparent mathematical foundations, e.g. using Open Science and co-design collaborative principles that facilitate cross-organizational and transdisciplinary human and technological learning on what works well in which setting. In that context, it is reassuring that the proposed digital epistemology has been shown to fit the scientific foundational principles of EBM very well (Robson, 2016), while it also enables opportunities to explore new scientific paradigms that give special attention to complex knowledge and reasoning situations, including epistemic uncertainties and conflicting evidence, that are likely to be especially impactful for unguided patients.

Considering the difficulty in finding such new paradigms when similarities between patient trajectories (and informative patterns they may contain) are less obvious, we may find inspiration in cutting-edge approaches to the use of multi-omics tumor biosample profiling for endotype modeling in the *Tumor Profiler* project in Switzerland[50] as well as efforts related to earlier and improved diagnosis of rare disorders in pediatric care in Switzerland[51]. In those transdisciplinary innovation programs we can observe a trend towards the increased use of multi-omics profiling for unguided patients to find patterns related to endotype (disease-driving biology / physiology) that may help to find the most appropriate treatment[52]. We would expect that as such efforts mature, they will learn where the most informative patterns are likely to show up, to make such efforts economically sustainable. However, the Swiss healthcare and health data digitalization landscape also provides a number of challenges in this context: a) digitalized health data being available for research at the appropriate level of quality, with the definition of quality depending on intended data use, b) fragmentation in the landscape of healthcare organizations reflected in digital connectivity gaps in a particular patient trajectory[53] (data being stuck in health data silos), c) healthcare providers can be overwhelmed with administrative burden in their workday, and may be reluctant to engage in such innovation as proposed here, d) expecting patients to take responsibility to consent on personal data treatment, e) data being ‘AI-ready’ in the context of FAIRification[54], apart from some data subsets that have already reached that level (overlaps with a), f) a bias towards hospital episodes[55] in the Swiss digitalized health data landscape, with many gaps in other areas of healthcare important for more preventive healthcare (covering earlier stages of disease and health risk monitoring).

[50] https://tumorprofilercenter.ch/

[51] https://swisspedhealth.com/

[52] This is conceptually similar to the challenge about an organization noticing that an incoming case does not fit the categories it uses to assign a process to that case.

[53] Once the Swiss health data space (SwissHDS) is operational this situation is likely to improve.

[54] FAIRification of data as outlined above also relates to the data being ready for use in AI, to make it interpretable for humans and machines. Note that Q-UEL is one of the approaches to enable human-AI collaboration.

[55] Hospital data are heavily biased towards urgent, acute care and later stages of disease progression, when symptoms are severe, while progress towards better preventive care requires more data about earlier stages of disease progression, when symptoms tend to be much weaker.

8 Cardio-Renal-Metabolic Patient Populations and Population Health

Cardiovascular diseases (CVD), chronic kidney diseases (CKD) and chronic metabolic diseases (CMD), and a widespread more elaborate “bundle” called Cardiovascular Kidney and Metabolic (CKM) syndrome that has gained attention in the US, are highly prevalent in many geographies, causing considerable burden for patients and their families, healthcare systems, payors and other stakeholders. These diseases share risk factors, epidemiological patterns and pathophysiological connections. Therefore, it is not surprising that they show a tendency to overlap, in patient trajectories, resulting in considerable co-morbidity. It has been suggested to increasingly coordinate care decisions in this space, across different healthcare providers and medical specialties, and to shift care towards a more proactive model focused on those shared risk factors, by personalizing preventive care and focusing more on early risk detection. For example, if a patient gets a CKD diagnosis to also assess (and then monitor, over time) CVD and CMD risks (Ndumele et al., 2023). New drugs like GLP-1 and SGLT2 modulators have shown considerable value in such populations, contributing to an increasing interest in smart prevention models. However, healthcare systems often struggle to operationalize such shifts towards more proactive and connected care, i.e. if their original system design was all about a focus on acute care, including related incentives and organizational structures.

In that cardio-renal-metabolic population, we can find subpopulations with elevated genetic risk that tend to experience a faster progression to bad outcomes. An example is a Lp(a)-defined CVD subpopulation with its increased risk of fast-progressing atherosclerosis (Reyes-Soffer, 2022). One wonders whether lifestyle interventions will be sufficient here, also in later stages of the subclinical phase of ASCVD, and asks when it is too early, for a non-lifestyle intervention, in those with high genetic risk. It all starts for an affected individual when the results from Lp(a) diagnostic testing are available to patients and medical professionals who may then face such questions, e.g. about timing. In Switzerland, under the umbrella of SPHN, CVD patient trajectories have been simulated using a digital twin that models an individual’s personal anatomy in the cardiovascular system (Laumer et al., 2025), in collaboration with the UK Biobank[56] and their imaging effort[57].

When simulating trajectories and outcomes in such cardio-renal-metabolic populations, with their fast-progressing subpopulations, it can be helpful to use tools to simulate various secondary effects of improved health outcomes, i.e. the impact of changes in how we make care decisions. For example, at the level of social impact, a shift in the health status of this population can affect many stakeholders in society. An example is PALY (Productivity-Adjusted Life Years) as a tool for such estimation. It quantifies health effects related to workforce productivity. That is with the basic assumption that a healthier population is more productive in many ways, including both paid and unpaid work in PALY estimates. Unpaid work can include informal care for children, the elderly and the disabled, and even volunteering work as well. Related reference work has been performed in Finland, by Janne Martikainen et al. using Finnish population data (Lavikainen et al., 2025; Ademi et al., 2025). With tools similar to PALY we can then learn which measures have maximal effect at population level, where effects show up and how they can be stabilized. See also *elevateHealth* examples below for such modeling and simulation using the IOOI (input, output, outcomes, impacts) approach (Section 9).

[56] The UK Biobank generated a large health dataset for an adult population that includes much information on such genetic risks, in addition to other medical data.

[57] An effort to generate imaging data at population level in the UK Biobank enabling the development of digital twins.

9 Agentic AI, Digital Twins, Elevatehealth and Related Digital Trends

In the last years we have seen much technical progress that may help to address the challenge of unguided patient trajectories and more patient-centric LHS designs in new ways. For example:

- Improved reasoning abilities of AI systems including medical and mathematical knowledge
- Improved understandability of expert knowledge and language for patients, also in medicine
- AI systems performing well on medical exams and similar question-answer tasks
- Progress related to agentic AI, multi-agent collaboration, and AI helping to code software

These kind of approaches have undoubtedly been impressive, and most observers say unexpectedly so. Considering such AI-related trends and their speed, in the next years, we may increasingly see more sophisticated patient-centric LHS emerge that empower patients to navigate a large, complex knowledge landscape related to their health and wellbeing in a better way. On the clinical side, Large Language Models (LLMs) are fine-tuned with clinicians' knowledge to increase accuracy, reliability, and, ultimately, safety of AI use for care plans formulation and delivery (MOOVE, nd). Furthermore, agentic AI systems may be developed that achieve a high degree of personalization, considering the knowledge about health and disease that a patient currently has, the language they prefer, the style that works well to keep them engaged care. This is to tap into an integrated knowledge landscape, using a reasoning approach that builds on EBM-style scientific reasoning about evidence and confidence, to then explore what works well outside of that space, where it matters to that patient, using informative patterns derived from unguided patient trajectories as inputs. Depending on the level of interest in learning that patient has, such an AI could help that patient to build up health-related knowledge based on increasingly robust reasoning. Across areas of knowledge, epistemological fields (see above), medical specialties and areas of knowledge that are not well understood by healthcare systems yet e.g. regarding the early detection of weak signals that can be used as inputs for such learning.

Agentic AI systems require a language that coordinates what different AI and non-AI agents do collaboratively, and to enable agent outputs to be evaluated and combined in that context. For example, natural language as produced and 'understood' by LLMs and humans can be used for such coordination, typically with humans in the loop on key decisions and the language used for keeping that human in the loop adjusted to that particular person (Zhou et al., 2025). In that context, humans can be professionals working in healthcare systems, patients and their families, or other actors that collaborate in patient-centric LHS on improving decisions, outcomes, based on digitalized collections of patient trajectories[58]. Such agentic AI may then contain a number of different agents specialized on certain tasks, areas of knowledge and decisions. For example, an agent may focus on modeling the patient trajectory, and how it compares to a reference dataset of patient trajectories. Another agent may use that patient trajectory model to check if it's an unguided patient trajectory, and then interact with relevant agents to deal with the case accordingly. Yet another agent may try to model endotype-related patterns that could help to improve the prediction of therapy response. And so on.

For the foreseeable future, recall that the difficulty with LLMs, Machine Learning and Deep learning is that they are not data analytic methods and they do not produce "real" probabilities and probabilistic measures of the kind that biostatistics, Evidence Based Medicine, Epidemiology, Comparative Effectiveness Research, and organizations like the FDA and WHO would approve of; nor are they truly simulations. Q-UEL-based approaches do provide these in a natural way, but there are

[58] In Europe such efforts would typically operate on a federated technology stack that keeps health data in their data silos while enabling learning and analytics. Synthetic data that represent particular populations may also play a key role esp. in prototyping.

certainly many other efforts. The *elevateHealth* Open Science community performs modeling and simulation of patient trajectories using a variety of modeling approaches, including Markov state probabilistic modeling (a particular kind of digital twin that simplifies access to complex longitudinal health datasets; Laubenbacher et al., 2024). Their framework is specifically designed to simulate patient journeys across individual, organizational, and societal levels, which inherently requires handling uncertainty and variation through probabilistic methods. These models allow stakeholders to test intervention scenarios and quantify outcomes before committing resources. This approach works well, but the limitation is missing interactions and hence causing loss of potential accuracy. As is usually so in the traditional Markov one, future states depend *only* on the current state, the chance of moving between states usually stays fixed over time. In reality, risks often increase as patients age, the model assumes risk does not change the longer you stay in a single state, models often assume a constant risk of death that poorly reflects real-world human aging and disease progression, and adding history requires creating many new states that makes the model too large and difficult to manage. The problem of lack of memory of history is not just in neglecting that "the past can come back to haunt you". Medicine well recognizes the effect of triggers on events that happened many years ago (e.g. excessive production of active oxygen species triggering pancreatitis for many years after being stung by a scorpion). Compare Bayes' Rule, which written in the symmetric way $P(A|B)P(B) = P(B|A)P(A)$ is equally about deducing probability of cause in the past as it is effect in the future. In the context of the patient-centric agentic AI described above, digital twins (Emmert-Streib,2025), derived from the available knowledge that is relevant for that particular patient can be used to visualize the likely effects of potential health-related decisions considering what we learned from a larger set of patient trajectories. Done up to personally relevant aspects of social impact as captured by the PALY approach (see above), this enables patients to simulate what will happen, in which time frame, if they do X or Y. This can enhance patients' understanding of personal health data usage in AI systems to improve health through preventive modeling. A digital twin is a virtual, real-time replica of a physical object, process, or system, in this context a patient. In contexts other than medicine, it continuously updates using sensor data and artificial intelligence, allowing engineers to simulate behavior, predict failures, and optimize performance without risking the real-world asset. That level of simulation remains an ideal for patients and there is much research, because it implies anatomic, physiological, metabolic, and neural simulations, but much simpler representations such as the digital patient record combined with some form of reasoning and analytical activity can be very powerful. Note, however, that Q-UEL has been used to run metabolic and neural simulations using elements of form such as < substrates | **enzyme** | products > and < input neuron | **synapses** | output neuron >. The form also seems naturally well suited to simulating biological processes, as it is for < patient state A | **event** | patient state B >.

Learning at system level, here, also means to learn when updates based on new data from that patient help to decrease uncertainties captured in the digital epistemology, e.g. as Q-UEL. There may be situations where the opposite can happen, e.g. new data actually increase uncertainties. Note that different actors as mentioned above, e.g. different healthcare professionals, patients and scientists, may have different, complementary perspectives and understanding of uncertainties, which can be captured in Q-UEL. Additional interactivity between *elevateHealth*-style simplistic Markov models and the proposed digital epistemology may include a data-driven discovery of health states (Markov states)[59], and for checking consistency between elevateHealth-style models and the full complexity captured in the digital epistemology. In case of conflict between information, Q-UEL naturally treats it as a process of balance

[59] Typically, health states in Markov models are decided by a team of modelers who make decisions about the simplification of complex realities in the model design. A data-driven approach is possible as well, where the health states are learned from data, and then checked by the modeling team.

of evidence as in a court of law, using essentially the same underlying approach as computing EBM measures such as Relative Risk.

Specialized LLM-style AIs to model and generate patient trajectories using UK Biobank and Danish data have recently been published, as LLMs are increasingly used in related research (Shmatko et al., 2025). Note that such health LLMs could be used to generate *synthetic datasets* composed of a large number of patient trajectories for prototyping agentic AI work. However, note that any AI work as outlined above needs to consider what we know currently about different biases that affect AI e.g. due to the biases we can describe in the data that were used to train the AI (Nonori et al., 2021).

10 Human-Centric Design, Patient Empowerment and Knowledge

In Switzerland, many of the above-mentioned SPHN or digiSanté programs emphasize the use of patient-centric, participatory approaches, including increased patient empowerment (Varela et al., 2025), human-centric design of AI-based digital systems, and digitally enabled personalization to better fit digital systems into people's lives (Seebode, 2013). For example, the increased use of patient-reported outcome (PRO) tools to better estimate changes in delivered tangible value for patients, in addition to expert clinical definitions of value, complement the understanding of a patient trajectory, and helps to achieve a better balance between system and individual in system design (see above). Here empowerment is also about enabling the personalization of preventive care, tapping into a rich knowledge landscape to help a patient make better decisions related to their health.

With increasing use of digital systems, including agentic AI, as outlined above we can foresee an improved consideration of the patient experience of care in a form that can help us with the discovery and validation of informative patterns in patient trajectories, even if they are combinations of weaker signals that are distributed across data silos currently. Particularly challenging are high dimensional cases, meaning that many demographic and clinical factors may be involved. High dimensionality means sparse counts, and methods of managing and including evidence from sparse data, and detection of weak signals, becomes important.

11 The Outlook for a “Digital Epistemology”

On the one hand, the advantage of epistemological approaches over current mainstream AI, as the learning of arbitrary weights from massive amounts of human knowledge on the Internet, is that current AI takes little account of epistemology but learns primarily by rote, based on what humans have done. It often appears to reason, and in a sense, it is plausible that it can learn abstractions as metarules and so do a level of reasoning. For example, it might do so using syllogisms such as “All X are Y and all Y are Z, so all X are Z”, and even fill in gaps that general way in mathematical proof, using mathematical steps. The problem is that on the use of learning arbitrarily distributed weights it is extremely difficult to find out exactly how current AI is reasoning, and it seems woefully easy to make mistakes neglecting the vast arena of human experience that gives rise to common sense, wisdom, intuition and experience. There are many areas in which we can tolerate this to enormous advantage, using human common sense, wisdom, intuition and experience by a human user as a check. But in medicine, which as Evidence Based Medicine has argued, is sensitive to human bias and personal experience, this seems a dangerous enterprise. The main case for EBM is to maximize patient safety and treatment effectiveness by moving away from intuition, tradition, and blind reliance on unproven expert authority, and turn to direct evidence based on statistical evidence from RCTs and epidemiological analysis.

One the other hand, it could be said that there is lack of empirical evidence supporting the proposed digital solution. In part that is, of course, because it has yet to be implemented: rather, the present paper provides the background for the need and for the project. The substantial published Q-UEL literature discussed in Section 4, including proven capabilities in a variety of medical areas, and comparison with other methods, only addresses some of that concern, because the unguided patient has not yet been a use case. But it would certainly seem foolish (and indeed unhumanitarian) not to seek and explore some kind of digital solution, or at least digital support to ameliorate the situation. The critics may well clarify their concern regarding the feasibility of any digital approach as being extremely slim because of the diversity, uncertainty, and limited knowledge about the unguided patient. That is, it is "an impossible task" problem, unsuited as a significant initiative when there are so many other challenges in medicine and healthcare to meet. However, the unguided patents represent a large fraction of the population, perhaps in part at some time every patient, and managing uncertainty and sparse data is a major purpose of Q-UEL. In any case, in our opinion, diversity and uncertainty and limited knowledge are precisely the argument for some kind of effort that addresses the power and limitations of knowledge, and arguably in the fundamental and elegant way in which theoretical physics has managed uncertainty and probability. This is what Dirac's approach conveniently brings to the vision of Leibniz.

In terms of an outlook, it may help us to understand key differences between the proposed "digital epistemology" (and its relevance for learning based on unguided patient trajectories) and the current medical epistemology used in healthcare decisions. Our "digital epistemology" is not meant to replace the existing one, but to extend it, focused on providing additional value related to the problem of unguided patients, and the design of patient-centric systems, where existing epistemology does not provide sufficient guidance yet. If such an unguided patient is lucky to meet a highly knowledgeable healthcare professional who has all the relevant knowledge[60] to make a good decision on a difficult case, and enough time and attention (including the patient's knowledge, experience and insights), not being incentivized to force-fit the case into a diagnostic code that does not really fit. If all these factors are optimal in the moment when a key decision is made, then our case as presented here may be deemed weaker. However, how likely is it, in healthcare system realities with limited resources, limited time, imperfect attention etc. that all factors are optimal at that time? How much time would the patient have to communicate what is missing in the expert's perspective on the problem, and what would be the most effective and safe way to integrate patients' feedback?

As healthcare shifts increasingly towards smarter personalization of preventive care, we will need digital systems that help us learn how we can improve, build knowledge and make decisions. These will hopefully not be using an 'over-dose' of black-box AI[61], but something more scientific and robust that can enhance our shared knowledge and reasoning. That is even if that patient is not highly educated in relevant scientific paradigms. Current AI itself is well known to be now acknowledging the need for some dynamical process of reasoning that goes beyond rote learning and somehow emulates the reasoning processes that occur in the higher parts of the human brain. With Health Data Spaces then improving our ability to see informative patterns across health data silos, agentic AI improving the patient's ability to navigate all this, we may soon be able to witness a new phase of health innovation.

Once we have been able to co-create an Open Science catalog of such informative patterns in unguided patient trajectories, and a better understanding of the types of unguided patient trajectories that can be found, such information, if well-organized, could help to better connect the efforts of a) those working to improve decision making as outlined above, with b) those who need information about underlying endotypes that drive disease progression in unguided patients, to develop new interventions

[60] In other words, while a highly knowledgeable individual can make a big difference, it does not mean that this knowledge is accessible to others e.g. within that learning healthcare system.

[61] There may be helpful applications of black-box AI in medicine even if we don't really understand the decision the AI made, but in general, in medicine, we want to use a more scientific approach to decision making, where we can, which can extend our knowledge and improve our reasoning for the benefit of patients and other stakeholders who are interested in good healthcare and healthy, productive populations.

for these difficult cases. Considering that our knowledge about disease biology is quite patchy in many areas, there may be different ways to describe 'tentative' endotypes, from different angles, as we capture different hypotheses and the evidence that supports them. Without over-simplifying what we know so far about the less studied endotypes, including those active in earlier stages of disease (where we tend to have less EBM-grade evidence but may have animal models to study). To accelerate efforts to improve care decisions and the options available to such unguided patients, through endotype modeling.

While currently widespread, simpler logic such as Boolean reasoning in decision trees and Bayes Nets for probabilistic reasoning (Owens & Sox, 2021), can solve many common cases quite well, based on currently available evidence, we argue that unguided patients provide a clear opportunity to apply more sophisticated digital epistemologies such as the one outlined here, to overcome the known limitations of such simpler reasoning and knowledge representation, to bring patient-centric LHS to the next level. Making expert knowledge more accessible to more people as they make health-related decisions.

In that context we plan to build an Open Science initiative with a Python code and synthetic data repository on GitHub, and appreciate any constructive inputs and contributions to that community effort. Therefore, this whitepaper is meant to provide a reference and scope for this Open Science initiative.

References

1. Ademi, Z. et al. (2025), Scoping Review of Productivity-Adjusted Life Years (PALYs): Methods, Applications and Policy Implications. PharmacoEconomics 43:1367-88.
2. Alencar, J. N. et al. (2024), A Journey Through Philosophy and Medicine: From Aristotle to Evidence-Based Decisions. Philosophies 9:189.
3. Allam, A., Feuerriegel. S., Rebhan, M. and Krauthammer, M. (2021), Analyzing Patient Trajectories With Artificial Intelligence. J. Med Internet. Res. 3;23(12):e29812.
4. Berkowitz, A. and Cook, D. J. (2015) Leibniz-Bouvet correspondence. Translation of letters. https://leibniz-bouvet.swarthmore.edu/ (last access May 19, (2026).
5. Bernier, J. (2024), Revolutionizing Health Care Management: The 8Ps Value Proposition. Personal communication and circulated document. See also: https://www.linkedin.com/pulse/revolutionizing-health-care-management-8ps-value-jacques-bernierzw2ue/
6. Breuils, S. et al. (2022), New applications of Clifford's Geometric Algebra.. Advances in Applied Clifford Algebras, Nature, 32(17), .
7. Buergel, T. et al. (2022), Metabolomic profiles predict individual multi-disease outcomes. Nature Medicine, 28: 2309-2320.
8. Carou, D., Sartal, A.,and Davim, J. P. (Eds), (2022), Machine Learning and Artificial Intelligence with Industrial Applications: From Big Data to Small Data, Springer.
9. Agency for Healthcare Research and Quality, CDS Innovation Collaborative (2025), US Gov. How to empower and engage patients with patient-centered clinical decision support. Featuring products from the AHRQ Clinical Decision Support Innovation Collaborativehttps://digital.ahrq.gov/sites/default/files/Patient%20Empowerment%20How%20to%20Guide_Final_8.27.25.pdf
10. Damarell, R. A. (2020), General practitioner strategies for managing patients with multimorbidity: a systematic review and thematic synthesis of qualitative research. BMC Family Practice 21.
11. Deaton, A. (2015), The Great Escape: Health, Wealth, and the Origins of Inequality. Princeton University Press.
12. Deckelman, S, and Robson, B. (2014), Split-complex numbers and Dirac bra-kets. Communications in Information and Systems 14:135-159.
13. Dine, M. (2022), This Way to the Universe: A Theoretical Physicist's Journey to the Edge of Reality, Dutton.
14. Dirac, P. A. M. (1930), Nobel Prize Banquet Speech, NobelPrize.org, https://www.nobelprize.org/prizes/physics/1933/dirac/speech/ (last accessed May 7 2026).
15. Dirac, P. A. M. (1939), A new notation for quantum mechanics. Mathematical Proceedings of the Cambridge Philosophical Society 35: 416-418.
16. Dirac, P. A. M. (1958), The principles of quantum mechanics. Oxford Science Publications. 4thedition (revised).
17. Emmert-Streib, F. (2025), The role of digital twins in P4 medicine: A paradigm for modern healthcare. Nature, Npj Digital Medicine, 8:735. https://doi.org/10.1038/s41746-025-02115-x
18. European Commission (2022), Proposal for a Regulation on the European Health Data Space. COM(2022) 197 final. https://eur-lex.europa.eu/legal-content/EN/TXT/ (last accessed May 20 2026)
19. Friedman, C. et al. (2014), Toward a science of learning systems: a research agenda for the high-functioning Learning Health System. J. Am. Med. Inform. Assoc. 22:43-50.

20. Froehlich, H. et al. (2018), From hype to reality: data science enabling personalized medicine. BMC Medicine, 16(1):150.
21. Wang et al. (2020), Global age-sex-specific fertility, mortality, healthy life expectancy (HALE), and population estimates in 204 countries and territories, 1950–2019: a comprehensive demographic analysis for the Global Burden of Disease Study 2019. Lancet 396:1160-203.
22. Gibney, E. (2025), Physicists disagree wildly on what quantum mechanics says about reality, Nature survey shows. Nature News, 30 July 2025.
23. Higgins, J. P. T. et al. (2024), Cochrane Handbook for Systematic Reviews of Interventions, version 6.5 (updated August 2024). Cochrane. https://www.cochrane.org/authors/handbooks-and-manuals/handbook
24. Hüllermeier, E. and Waegeman, W. (2021), Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods, Machine Learning 110: 457-506.
25. Hunter, A. and Williams, M. (2012), Aggregating evidence about the positive and negative effects of treatments. Artificial intelligence in medicine, 56(3): 173-190.
26. Ingvarsson, J, (2021), Towards a Digital Epistemology: Aesthetics and Modes of Thought in Early Modernity and the Present Age. Palgrave Macmillan, Springer. Open access book (PDF).
27. Ioannidis, J. P. A. (2005), Why most published research findings are false. PLoS Medicine 2:e124.
28. Karsala, T. et al. (2025), Balanced hyperbolic embeddings are natural out-of-distribution detectors. arXiv:2506.10146v1.
29. Kim-McManus, O. et al. (2024), A framework for N-of-1 trials of individualized gene-targeted therapies for genetic diseases. Nature Communications 12;15(1):9802.
30. Laubenbacher, R., Mehrad, B., Shmulevich, I. and Trayanova, N. (2024), Digital twins in medicine. Nature Computer Science 4(3):184-191.
31. Laumer, F. et al. (2025), 2D echocardiography video to 3D heart shape reconstruction for clinical application. Medical Image Analysis 101:103434
32. Lavikainen, P. T. et al. (2025), The Impact of Chronic Conditions on Productivity-Adjusted Life-Years in Both the Workplace and Household Settings in the General Adult Population in Finland, Value Health, 28(3):379-388
33. Lemos, N. (2020), An introduction to the Theory of Knowledge, Cambridge University Press
34. Levi-Strauss, C. (1962), The savage mind. University of Chicago Press.
35. Lucas, P. J. F. (2001), Bayesian networks in medicine: A model-based approach to medical decision making. Proceedings of the EUNITE Workshop on Intelligent Systems in Patient Care, Vienna, 73-97.
36. Lucas, P. J. F. , van der Graag, L. C., and Abu-Hanna, A. (2004), Bayesian networks in biomedicine and health-care. Artificial intelligence in medicine. 30:201-214.
37. Maurer, M. et al. (2024), Modelling of patient journey in chronic spontaneous urticaria: Increasing awareness and education by shorten patients' disease journey in Germany. J. European Academy of Dermatology and Venereology, 38:2093-2101
38. Milkov, N. (2006), A new interpretation of Leibniz' project for Characteristica Universalis. In Hans Poser, Einheit in der Vielheit, Proceedings of the 8th International Leibniz-Congress. pp. 606–614
39. Massive Open Online Validation and Evaluation (MOOVE). (n.d.). MOOVE. https://jointhemoove.org (last accessed May 24, 2026).
40. Ndumele, C. E. et al. (2023), Cardiovascular-Kidney-Metabolic Health: A Presidential Advisory From the American Heart Association. Circulation 148.
41. Ni, Z. et al. (2010), The Use of Bayesian Networks in Decision-Making, in Athanasiou, T., Debas, H., Darzi, A. (Eds), Key Topics in Surgical Research and Methodology. Springer.
42. Nonori, N. et al. (2021), Addressing bias in big data and AI for health care: A call for open science. Patterns (NY) 2(10):100347.
43. Oeppen, J. and Vaupel, J. W. (2002), Broken limits to life expectancy. Science 296:1029-31.
44. Owens, D. K. and Sox, H. (2021), Biomedical decision making: Probabilistic clinical reasoning. In: E. H. Shortliffe, J. J. Cimino (eds.), Biomedical Informatics, Springer.
45. Page, R., Noches, P., & Gomez-Hernandez, M. (2025). Primary Perspectives on Design Research: Making and Action-Oriented Epistemology in Design. https://dl.designresearchsociety.org/iasdr/iasdr2025/fullpapers/45/ (last accessed May 24, 2026).
46. Polotskaya, K. et al. (2024), Bayesian Networks for the Diagnosis and Prognosis of Diseases: A Scoping Review. Mach. Learn. Knowl. Extr. 2024, 6(2), 1243-1262
47. Porter, M. E. (2010), What is value in health care? New England Journal of Medicine 363:2477–2481
48. President's Council of Advisors on Science and Technology (2010), Realizing the Full Potential of Health Information Technology to Improve Healthcare for Americans: The Path Forward. https://obamawhitehouse.archives.gov/sites/default/files/microsites/ostp/pcast-health-it-report.pdf (last accessed May 7 2026).
49. Raposo, V. L. (2025), The fifty shades of black: about black box AI and explainability in healthcare. Medical Law Review 33: fwaf005.
50. Rebhan, M. (2017),Towards a systems approach for chronic diseases, based on health state modeling. F1000Research. https://f1000research.com/articles/6-309/v1

51. Reyes-Soffer G. et al. (2022), Lipoprotein(a): A Genetically Determined, Causal, and Prevalent Risk Factor for Atherosclerotic Cardiovascular Disease: A Scientific Statement From the American Heart Association. Arteriosclerois and Thrombosis Vascular Biology, 42(1):e48-e60
52. Robson, B. (2005), Clinical and Pharmacogenomic Data Mining: 3. Zeta Theory As a General Tactic for Clinical Bioinformatics” Journal of Proteome Research. (Am. Che. Soc.), 4, 2, 445-455
53. Robson, B. (2007), The new physician as unwitting quantum mechanic: is adapting Dirac’s inference system best practice for personalized medicine, genomics and proteomics?, J. Proteome Research, 6:3114.
54. Robson, B. (2014), Hyperbolic Dirac Nets for medical decision support. Theory, methods, and comparison with Bayes Nets. Computers in Biology and Medicine 51:183-197
55. Robson, B (2015), POPPER, a simple programming language for probabilistic semantic inference in medicine. Computers in Biology and Medicine 56:107-123
56. Robson, B. (2016). Studies in using a universal exchange and inference language for evidence based medicine. Semi-automated learning and reasoning for PICO methodology, systematic review, and environmental epidemiology. Computers in Biology and Medicine 79:299-323
57. B. Robson, B. and S. Boray (2015)., Implementation of a web based universal exchange and inference language for medicine. Sparse data, probabilities and inference in data mining of clinical data repositories, Computers in Biology and Medicine, 66, 82-102
58. Robson, B. and Boray, S. (2016), Data-mining to build a knowledge representation store for clinical decision support. Studies on curation and validation based on machine performance in multiple choice medical licensing examinations, B Robson, S Boray, Computers in Biology and Medicine, 73, 71-93
59. Robson, B. and Boray, S. (2018), Studies in the extensively automatic construction of large odds-based inference networks from structured data. Examples from medical, bioinformatics, and health insurance claims data. Computers in Biology and Medicine 95:147-66.
60. Robson, B., Caruso, T. and U. G. J. Balis, (2013), Suggestions for a Web Based Universal Exchange and Inference Language for Medicine, Computers in Biology and Medicine, 43(12):2297-310
61. Robson, B. and St Clair, J. (2022), Principles of quantum mechanics for artificial intelligence in medicine. Discussion with reference to the Quantum Universal Exchange Language (Q-UEL). Computers in Biology and Medicine, 143, 105323.
62. Samuel, J. P. , Wooton, S. H and Tyson, J. E. (2023), N-of-1 trials: The epitome of personalized medicine? J. Clinical Translational Science, 7(1):e161.
63. Seebode C (2013), A patient centered infrastructure. IMECS 2013, Hong Kong.
64. Sheldrick, R. C. (2023), Randomized trials vs real-world evidence: how can both inform decision making?, J. Am. Chem. Soc. 329(16):1352-1353
65. Tonelli, M.R. and Bluhm, R, (2021), Teaching Medical Epistemology within an Evidence-Based Medicine Curriculum, Teaching and learning in medicine, 33(1):98-105.
66. Triposkiadis, F. and Brutsaert, D. L. (2025), Evidence-Based Medicine: Past, Present, Future, J. Clinical Medicine, 14(14):5094
67. UNESCO (2021), UNESCO recommendation on Open Science, https://unesdoc.unesco.org/ark:/48223/pf0000379949 (last accessed May 20, 2026).
68. Varela, A. J. Gallamore, M. J.Hansen, N. R., and Martin, D. C. (2025), Patient empowerment: a critical evaluation and prescription for foundational definition, Frontiers in Psychology, 15, 1473345.
69. Weigert, S. (2023), Relativity and Quantum Theory: Under the Spell of Today’s Paradigms. arXiv:2312.17693v1.
70. Zhou. P. et al. (2025), Why do AI agents communicate in human language? arXiv:2506.02739v1

Acknowledgements

We would like to acknowledge comments by Jacques Bernier, MD (Swiss Medical Network, Switzerland) and Prof. Joachim Buhmann (ETH Zurich, Switzerland) that helped us improve the text.

Ethics and Conflicts of Interest

The authors have no conflicts of interest. The Q-UEL language suggested in this paper and associated software applications required for its basic use in healthcare will be made available as Open Source and supports ethical principles such as privacy and fine-grained consent. In accordance with the policy of The Dirac Foundation and Ingine Inc., Q-UEL has been promoted for general use in human and veterinary medicine, and potentially more generally, and its specifications and examples of use have been described in publications without patenting. It is a transport language for input and/or output in the sense of XML, JSON, CSV files, etc., albeit with the advantages described in this paper. Ingine Inc. makes use of Q-UEL as an architectural principle and coauthor B. Robson is a cofounder of, and holds stock, in that company. Ingine's commercial interests are not regarding Q-UEL *per se* but relate to techniques required for Very Big Data, High Dimensional Data, Artificial Intelligence, and Quantum Computing. The Dirac Foundation is a UK Company Limited by Guarantee but is non-financial in operations and performs charitable and *pro bono* works in public health and epidemiology.

APPENDIX

A.1 Future Epistemological Concerns

Epistemologies are frameworks for knowing, with a set of rules and norms about what counts as a valid question, a legitimate method, a credible source, an acceptable proof, a justified belief and a contradiction (and how to handle it). They are socially created and maintained (Lemos, 2020). In the context of Thomas Kuhn's original view on paradigm shifts in the history of science (Kuhn, 1962), such shifts are enabled by a change in epistemology. Different historical periods and cultures have built different epistemologies that facilitate human coordination. As scientific disciplines establish themselves and their boundaries, they create their own epistemology. The proposed *digital epistemology* above involves choosing how we want to know and learn, including complex, uncertain, and evolving domains in medicine, and other areas of knowledge related to human health. In such efforts we may therefore be informed by past epistemologies, evaluating their strengths and blind spots. Epistemology-related questions for building patient-centric LHS may then include a) do we want to preserve epistemic uncertainty rather than de-emphasize it?, b) does it need a transparent handling of assumptions and versioning of knowledge?, c) how to handle edge cases and contradictions?, d) how much can meaning and validity vary by context and perspective?, e) how do algorithmic approaches that work well with machines connect with human sensemaking (patient, experts)? f) can we hold hypotheses in tension, waiting for resolution? How does resolution happen? Some of those questions relate to the overarching notion of *epistemic humility* and transparency of epistemic uncertainties that may not be considered well in past epistemologies. The current wave of innovation related to AI-supported patient-centric decision making may provide an opportunity for re-thinking along such lines to avoid over-reliance on intransparent black-box AI in decision making. Therefore, unguided patients who don't fit well into clinical care guidelines and the EBM evidence landscape help us see the limits of our current medical epistemologies and what we need to improve.

A.2 Different Types of Uncertainty

Hüllermeier & Waegeman (2021) propose that a trustworthy representation of uncertainty is desirable, especially in safety-critical application domains such as medicine and socio-technical systems. Not only as an on-average uncertainty for a particular type of prediction, but for a particular prediction on a particular instance for a particular patient (considering the patient's trajectory in the context of a larger collection of other trajectories). In the AI literature, there is increasing recognition of the need to distinguish different types of uncertainty, with effects in how we model and do machine learning. Basically, in addition to uncertainty caused by randomness (aleatoric uncertainty), a different type of uncertainty is 'epistemic uncertainty', i.e. uncertainty due to weak knowledge. The latter has the advantage that it can be reduced with additional information. In other words, we can reduce 'ignorance' by learning, i.e. by reducing epistemic uncertainty. Therefore, it is also considered the 'reducible' part of uncertainty, in contrast to irreducible parts considered in previous theories about uncertainty. In patient-centric LHS, their system-level learning goals can often be described in terms of reducing such ignorance, through a learning process the aims to reduce epistemic uncertainty, where such reduction is likely to result in improved patient outcomes and other process-level improvements considering the principles of value-based health care (Porter, 2010).

A.3 P4 Ecosystems and unguided patient trajectories

Using the meanings behind the fundamental improvement principles of P4 medicine (i.e. to make medicine more *personalized*, *participatory*, *preventive* and *predictive*) we can derive criteria for answering the question of whether a LHS composed of complementary organizations is indeed a P4 ecosystem, or not. If ecosystem-level goals that guide the collaborative effort (i.e. they are used for defining the meaning of success and related KPIs[62]) include making the practice of medicine more personalized, participatory, preventive and predictive - often in addition to other goals, and areas of focus, that are specific to that ecosystem. In addition to criteria derived from those 4 principles, we can define other ecosystem properties as criteria: a) ecosystem orchestration processes that help to manage conflicts (potential negative outcomes) and maximizes synergies (potential positive outcomes) between the participating organizations, b) a way of plugging in for a new organization that wants to join, and c) a shared learning process at ecosystem level that matches at least some Open Science principles (UNESCO, 2021) in the sense that sufficient transparency is designed in to keep the ecosystem engaged.

Therefore, different archetypes of P4 ecosystems would fulfill all those criteria, but in different ways. Traditionally, ecosystem orchestration (Valkokari, 2015; Valkokari, 2023) is centralized, driven by the organization that invited others into ecosystem co-design and risk sharing negotiation. Exemplified by the Integrated Care archetype of P4 ecosystems (section 3.1), the Nightingale archetype for metabolomics-enabled health risk detection and health maintenance for working populations in Finland (Buergel, 2022) and the HiNouNou connected health ecosystem (a patient-centric LHS with digital twins evolving towards agentic AI as described above)[63]. However, new orchestration models that increasingly deviate from classic, top-down ways of organizing may emerge over time, e.g. through decomposing what orchestration is about in ecosystem design. Our whitepaper here relates to P4 ecosystems in terms of a) ecosystems figuring out how to make progress with P4 translation into health care, b) enabling aspects of that shared learning process that require a paradigm shift in its scientific foundations, beyond classic probabilistic modeling and reasoning, and over-reliance on black-box AI or even Boolean logic, c) enable a data-driven refinement of decentralized orchestration models to adjust the level of decentralization where it helps to improve at ecosystem level. In that sense it will be exciting to watch how this space will develop in the next years, and which P4 ecosystems will learn faster and become smarter, using which approach.

A.4 The Q-UEL Language

We attempt here only the simplest account, though it captures the important principles. The mathematization of thought and language in the manner of Leibniz plus rendering it probabilistically quantitative for disciplines such as Evidence Based Medicine seems no easy task, yet Dirac's quantum mechanics does it for us, and arguably to an extent and in such a natural way that one may still consider our language as quantum mechanics. We merely demote the use of the more familiar imaginary component (the ***i***-complex part) responsible for the notorious predictions of quantum mechanics in the everyday world of human experience. Patients are too big to be seen as waves,. The primary elements of Q-UEL follow the Dirac braket notation in which "bra" row vectors of from <…| and "ket" vectors of form |…> are multiplied together to form <…|…> which is both an expression and the scalar product of the two vectors. In quantum physics one may see, in Dirac notation,

< momentum(eV-sec):=0.2 | position(Angstrom):=6.3 >

[62] KPI = key performance indications, used to track an organization's performance at system level.

[63] https://www.hinounou.com/

Textbooks can be a bit “sloppier” and varied in how they represent attributes such as momentum p and position x, e.g. writing < p | x > and explaining p and x in the text, but the braket form <…|…> is key. In medicine, one might see an analogous conditional probability…

P(systolicBP(mmHg):=145 | glucose(mg/dl):=180)

In Q-UEL for medicine, the above would translate to…

< systolicBP(mmHg):=145 | glucose(mg/dl):=180 >

The difference between the probability P(…|…) and braket <…|…> form is that the scalar complex *probability amplitude* of form <A|B> encodes a *probability dual*, a *pair* of probabilities deducible by a process known as the Dirac recipe (P. A. M. Dirac 1958). It may briefly be represented as follows. By A and B we mean any attribute, thing, state, event, observation, measurement, property, description of logical or similar expressions involving them, which can be a situation that may be observed and counted to compute probabilities P, or analogous Bayesian degrees of belief.

$$\langle A|B\rangle \rightarrow \{P(A|B), P(B|A)\} \qquad (1)$$

Note that whatever is put in a dual, the two forms must be *adjoint forms*, such that they are both meaningful in an inverse sense. Eqn. 1. is not in general reversible, because it is ***i*** that leads to wave behavior, and waves can cancel as well as augment each other. Hence information is lost in the process described by Eqn. 1. In Q-UEL we usually wish things to behave classically, we wish to assign empirical probabilities, and we wish to do familiar probability calculations using the full machinery of quantum mechanics. That also means we wish to use the Dirac notation and algebra (Dirac,1939) with empirical probabilities as input. That requires

$$\langle A|B\rangle = \{P(A|B), P(B|A)\} \qquad (2)$$

We also wish that <A|B> can be used in equations, inference nets, and inference engines in a way that follows classical probability laws. To do this, use is made of the hyperbolic complex number ***h*** such that ***hh*** = +1 rediscovered in other guises by Dirac. First introduced early in his book as the linear operator σ such that σσ =+1, it is part of the overall probability amplitude, and an innovation beyond Schrödinger’s wave mechanics as an ***i*** and ***h*** complex algebra. For most purposes we are not calculating regarding waves, so we drop the ***i*** part. To compute what {P(A|B), P(B|A)} means as an ***h***-complex number we make use of a *Hermitian Commutator* as follows.

$$\langle A|B\rangle = \tfrac{1}{2}\,[P(A|B) + P(B|A)] + \boldsymbol{h}\,\tfrac{1}{2}\,[P(A|B) - P(B|A)] \qquad (3)$$

Note that (<A|B>)* = <B|A> with * meaning “*take the complex conjugate”*, i.e. simply change the sign of the imaginary part. Change ***h*** to ***-h*** in Eqn. 1. We have a super-probability (probability amplitude) in one complex number, because it is not true that P(A|B) = P(B|A)*…it has no imaginary part. But Eqn. 1 is hard to manipulate. Easier to manipulate is the algebraically equivalent *Spinor Projector form* of Quantum Field Theory.

$$\langle A|B\rangle = \mathbf{Z}P(A|B) + \mathbf{Z}^*P(B|A) \qquad (4)$$

Here the spinor projectors are **Z** = ½ (1+***h***) and **Z*** = ½(1-***h***), and they are easy to manipulate because **ZZ** = **Z**, and **Z*Z***=**Z*** and ***ZZ**** = **Z*Z** = 0 and **Z**+**Z*** = 1. But even more simply for most purposes that means we can manipulate probability duals directly, e.g.

$$\{0.3, 0.6\} + \{0.2, 0.1\} = \{0.3+0.2, 0.6+0.1\} = \{0.5, 0.7\}$$

$$\{0.3, 0.6\} \times \{0.2, 0.1\} = \{0.3 \times 0.2, 0.6 \times 0.1\} = \{0.06, 0.06\} = 0.06$$

$$\ln(\{0.3, 0.6\} = \{\ln(0.3), \ln(0.6)\} = \{-1.2040, -0.5108\}$$

Because the braket is really an algebraic expression, we can introduce an operator, e.g. a matrix, between them, say giving <A| times **R** times |B> giving <A|**R**|B>. In Q-UEL, as is often in quantum mechanics, the **R** is a Hermitian operator and a projection matrix that again gives probabilities in both directions.so taking the complex conjugate again, we have

< ‘type 2 diabetes’ | **causes** | obesity > = (< obesity | **causes** | ‘type 2 diabetes’ >)*

= < obesity | **causes*** | ‘type 2 diabetes’ > = < obesity | **‘is caused by’** | ‘type 2 diabetes’ >

= {0.83, 0.30}

…at least, in many populations. Note that the relationship operator **causes** used in Q-UEL is potentially distinct from simple conditionality <A|B> that implies <A| **if** | B>. When Q-UEL tags are generated, transmitted, or stored the probabilities must be carried up front, as follows. Notice the following includes in *one entity* all the above forms (and the probabilities).

< ‘type 2 diabetes’ Pfwd:=0.83 | **causes** | obesity Pbwd:=0.30 >

Beyond this level, there appear to be complications and limitations in mapping to language, but in the footsteps of Leibniz, problems resolve. Other operators such as articles, qualifiers, quantifiers, **the**, **a**, **other**, **many, ten** and categorical operators **some**, **all**, **no**, are important “atoms”, and Q-UEL has intuitive defaults if they are omitted, such as probabilities of 1 and the logical operator **and**. Also, by < adult beavers | **build** | dams > we mean < **all** adult beavers | **build** | **some** dams >, essentially as in many natural languages, because the above operations would otherwise suggest that all dams are built by beavers (the Hoover Dam is not). Also note that when the above vectors and matrices are manipulated separately algebraically, they have ***h***-complex elements, but they are most easily represented and manipulated as arrays of real numbers multiplied by **Z** or **Z***. Finally note that there are nested forms of <…|…|…> that correspond to parsed forms of a sentence.

A.5 Programming Aspects

A few comments on how to ‘code this in Python’ or other programming languages. For example, in Python, split-complex vectors and relevant vector algebra are not part of the standard distribution of Python or any commonly used libraries. However, as shown above, they are not required for routine use. Some open source code development effort may be needed to facilitate experimentation using code. Note that Q-UEL itself can be considered a programming language. See especially POPPER language, a simplified form of Q-UEL to teach medical students how medical decision support tools are built, based on Q-UEL principles (Robson, 2015). Other fuller expressions of Q-UEL as a complete programming language involved writing a preprocessor in the PERL language that was essentially standard PERL with brakets <A|**R**|B> and vectors <A| and |B> introduced as variables to which constants, variables including these, or expressions including these, could be assigned. The preprocessor then converted this to code acceptable to the PERL interpreter. It is plausible that a similar idea might be used to combine Q-UEL and Python.

Glossary of Terms

Bureaucracy: in this paper, an organization that at its core is a category-to-process system, mapping incoming cases to a fixed set of categories, with processes attached to those categories. During the organizational ossification process (also called organizational aging) the fixed set of categories is increasingly out-of-touch with a changing world surrounding the organization. As a consequence, the ossified organization struggles to learn and adapt by using cases that do not fit those categories well as inputs to the learning process.

Digital epistemology: an epistemology that enables the transformation of classic *bureaucracies* into more patient-centric, learning systems that can learn from cases that do not fit well their inherited system of fixed categories in organizational cognition, exemplified here by *unguided patients* and their trajectories. Note the more detailed definition in section 2.

Endotype: a subtype of a disease defined by a distinct functional or pathological mechanism, rather than just its outward symptoms. While a **phenotype** describes observable characteristics (symptoms, lab findings), an **endotype** reveals the underlying molecular or biological pathway causing that disease, crucial for precision medicine and targeted treatments

Epistemological fields: Clusters of similar epistemologies in science that agree on many norms, beliefs and rules related to what counts as valid (high-quality) knowledge, as a contrast to subjective, personal opinion.

Epistemology: the theory of knowledge, especially regarding its methods, validity, and scope, and the distinction between justified belief and opinion. See footnote 1 for details, regarding the broader semantic field relevant here.

Hyperbolic Dirac Net (HDN): a probabilistic inference net, i.e. a probabilistic knowledge net for automated reasoning, built with Q-UEL components.

Learning healthcare system (LHS): a digital system that tracks patient trajectories and learns using outcomes, including what happens in healthcare systems but also outside of them. Learning is about improving care decisions.

LLM: a large language model, an AI trained on natural language patterns, and knowledge encoded in natural language.

Patient-centric LHS: an LHS that integrates medical expert knowledge and patient knowledge as it tracks patient trajectories comprehensively to optimize a patient's ability to make good health-related decisions, balancing the healthcare system and patient perspective in the way how it encodes knowledge and reasoning.

Patient trajectory: a dataset of time-stamped health(care)-related events for a single patient, including events related to diagnosis, treatment, outcomes and decisions made. Including events that are recorded by healthcare systems, and those recorded by patients. A set of patient trajectories may consist of patients with the same diagnosis but different treatment approaches and outcomes. See e.g. Allam et al. (2021).

Unguided patient: a special case of an incoming case that does not fit the organization's fixed set of categories well. Here, it mostly means a patient who does not fit the evidence landscape. In other words, the patient is not similar enough to those for which a lot of science exists to guide care decisions. See Section 1 and 2, but particularly Section 3 which discusses diversity of definition.