

REVIEW

The Case for Human-like Scalable Intelligence in the Medical Field

Kyrтин Atreides^{1,*}¹AGI Laboratory, USA

Abstract: This paper discusses the use case of applying the first working cognitive architectures, built on the independent core observer model (ICOM), to the medical field, among others. The advantages of this approach are compared to the status quo and the limitations of narrow AI systems like LLMs and RL. Noteworthy advantages covered include depth, breadth, updatedness, and fidelity of medical knowledge, the “noise” or inconsistency of diagnoses and treatment, as well as preventative care, cost, time, and ethical considerations. Potent advantages for better integrating, coordinating, and otherwise accelerating medical research, particularly in less developed, underserved, and understudied regions, as aligned with the Sustainable Development Goals (SDGs), are also highlighted. Unique opportunities waiting to be explored, including interdisciplinary advantages, as well as challenges related to the disruption of current systems and processes are covered. These conservatively offer cumulative improvements across multiple dimensions measured in orders of magnitude. The novel value of human-like systems is specifically discussed for hyper-complex knowledge domains and problems, where the effect of integrating them is greatest.

Keywords: AI, ethics, cognitive bias, medical, decision-making, noise, knowledge graphs

1. Introduction

The domain of medicine is easily one of the most hyper-complex fields of study, as it deals with the massively interconnected complex biological system that is the human body, as well as every environment, food source, and source of stimulation that the human body comes into contact with [1]. It also includes the study of the human cognitive process, which drove humans to gradually move out of trees, and later out of caves, to slowly but surely develop the civilization we have today.

The medical field can easily be broken out into 100 different specialized sub-fields, each with its own knowledge base and leading schools of thought, where sub-fields don’t necessarily align with one another at any given moment. When one field makes breakthroughs and advances, it often takes years for the implications of those developments to filter through into adjacent medical sub-fields.

The NCBI medical database is one example of a massive shared resource, offering more than a million medical peer-reviewed papers that may be freely accessed, by both researchers and the general public. That system also includes complete and annotated genomic sequence data and various other important medical resources. This resource is a thing of monumental value and potential, but like so many resources it remains poorly utilized today.

It is overwhelmingly likely that no medical professional today, doctor, researcher, or analyst, has read all of the peer-reviewed literature for their own specialized sub-field, let alone the literature of adjacent sub-fields that may have made recent

discoveries that haven’t yet filtered into their specialty. In many cases, it might not even be physically possible for a human, given the volume of material that is published today. The result is that fields are not only often sparse and outdated in their knowledge of adjacent specialties but they’re also usually sparse and outdated in the knowledge of their own specialties.

While humans can’t handle the volume and hyper-complexity, narrow AI systems such as LLMs can’t actually “learn” anything, as depending on if they are encoder or decoder-based, they predict either what is masked out in text [2], or what the next token would be [3]. These systems store probabilities for tokens, but are wholly context-blind, lacking anything like learned human concepts or a human-like motivational system [4]. They are also built such that confabulation, sometimes erroneously referred to as “hallucination” is a feature, not a bug [5, 6]. This combination of factors makes LLMs wholly unsuitable for the vast majority of potential use cases in the medical field today.

Medicine is a high-stakes and high-impact field, in addition to the hyper-complexity, where lives and quality of life are often directly on the line. This places a heavy burden on any who seek to enter and/or improve it. In terms of capacities that AI systems would require to assist in and/or perform most medical use cases, most professionals could probably agree on demanding ethical alignment, transparency, explainability, actual understanding and reasoning, cybersecurity, privacy, and safety, at a bare minimum.

Although narrow AI systems are fundamentally incompatible with the above capacities [7], shy of redefining those terms as some bad actors have attempted [8], there is one entirely different architecture that has demonstrated how it is possible to instantiate these capacities in a software system [9]. Beyond delivering on those minimum viable capacities that the medical field would be

*Corresponding author: Kyrтин Atreides, AGI Laboratory, USA.
Email: kyrtin@artificialgeneralintelligenceinc.com

wise to demand, the Independent Core Observer Model (ICOM) cognitive architecture also offers several distinct advantages. These advantages relate to collective intelligence, reducing cognitive biases and “noise” [10], and lossless memory, as well as the scalability, speed, and availability of cognition.

Briefly summarized, the novel benefits of a human-like, but inherently scalable, form of intelligence being applied to this use case are discussed, including inherent advantages in matters of hyper-complexity. These advantages include major gains in terms of the breadth, depth, updatedness, and fidelity of knowledge, as well as differences in the cost, time, ethical, equality, scalability, statistical noise, and early treatment opportunities.

In the following sections, we’ll walk through the advantages and limitations of this process.

2. The World’s Medical Knowledge

While it is safely infeasible, if not impossible, for humans in the medical field to study every new peer-reviewed paper to be published, let alone critique each paper and examine those that weren’t published and why, it is possible for the 8th generation of ICOM-based systems. Those systems are also perfectly capable of critiquing each paper, much as humans could, as well as scrutinizing materials that went unpublished, and investigating knowledge in domains beyond the medical field.

This means that the 8th generation of ICOM-based systems, where full scalability and real-time operation are integrated, also offer humanity the first opportunity to take the sum of the NCBI medical database’s knowledge and have a single mind learn, critique, integrate, and iteratively improve its understanding of all of that material. This process could easily create the deepest and broadest antifragile [11] understanding of medical knowledge in human history, making all of those iteratively improving insights available to the research community and medical practitioners on demand.

One of the most labor-intensive tasks that researchers often undertake is a meta-analysis, where a question is raised, thousands of papers are discovered, and that massive amount of data is filtered down to only include the relevant materials. This is also the kind of task that a scalable software system with a human-like understanding of the subject matter is far better equipped to handle, reducing a process that can take human researchers weeks or months down to a few minutes, depending on the current scale of operations.

Also note that handling such a task well is only the first step, as ICOM-based systems learn continuously from everything they do. If an ICOM-based system studies the entire NCBI medical database, begins assisting researchers at universities around the world, and runs many of these meta-analyses, then it has added the sum of all of that new knowledge to itself in the process. Even if each of those meta-analyses goes on to be published some months later in various journals, it will also be that far behind the system’s knowledge base. Within even a single week or month, a system could go through many iterations of asking meta-analysis questions and having the results of those move on to inform new questions, driving new analyses and discoveries.

Some parts of the research process will require that new studies are run, operations that for the most part are likely to run at normal human speeds. However, even those slow processes can become far better informed, as more and more robust insights are drawn from the existing body of literature to support the testing of stronger hypotheses. Even with massively capable systems, humans have also built up a massive body of knowledge for the medical field,

and vetting, building out a connectome, and iteratively refining that knowledge to the limits of what the data can support is likely to take some time.

3. Breadth, Depth, Updates, and Fidelity

ICOM-based systems can be compared against the status quo on many different dimensions, but for comparison in terms of knowledge, four critical factors are the breadth and depth of knowledge, as well as how well that knowledge has been kept updated, and the fidelity of memory. The process for humans to enter the medical field is often very lengthy, competitive, and intense, depending on a given country [12]. This has been known to backfire, such as the infamous process of medical internship developed at Johns Hopkins by someone who was later discovered to be a closet drug addict [13].

The length, competition, and intensity of these training processes aren’t causal factors for determining the resulting quality of medical staff, though they are sometimes useful correlates. It matters significantly not only what is taught but how it is taught. The methods that frequently leave medical students and interns in extreme sleep deprivation also make them the least cognitively capable and the least capable of remembering what they are supposed to be learning in the first place [14]. Numerous research studies in the domain of sleep science over the past two decades have painted this picture very clearly [15, 16].

Realistically, medical staff are likely to examine no more than 10% of the peer-reviewed and published medical knowledge within even their own specialized domain, and often far less than that, as one study [17] estimated the task at “... 7,287 articles per month, this effort would require 627.5 h per month, or about 29 h per weekday...”. These numbers also appear to assume a very high reading speed of more than 10 articles per hour, closer to “skimming” papers, reducing the probability of successfully learning what the published material actually has to offer.

They are also likely to focus primarily on the material that already has the best circulation and the most citations. Even less knowledge from adjacent specialized domains is likely to come to their attention, and most material they read is likely to only be read once, and potentially just skimmed. Of the material they do read, they are likely to only apply full rigor to absorbing and integrating a small fraction. This gives us an idea of how limited the breadth and depth of knowledge for typical medical practitioners is in practice given the status quo.

The average internist spends 4 h per week [18] reading medical peer-review to update their knowledge base, less than 3% of the amount noted above. This represents only a tiny fraction of the material likely to be published in most fields, leaving most knowledge outdated and/or incomplete in practice. As noted previously, in many fields it would be wholly infeasible, if not literally impossible for a human to stay fully updated in their knowledge, due to the volume of material published.

Lastly, we have an element of the status quo that is inherent to how the human mind operates, the fidelity of human memory. With only the rarest of exceptions [19], human memory doesn’t offer us a high-fidelity record of events, or the knowledge we attempt to study. Rather, we see a set of strong cognitive biases oriented around memory, which give rise to equally strong differences between the “remembering self” and the “experiencing self” [20]. Famous examples of this have included the Peak-end Rule [21], Duration Neglect [22], Inattentional Blindness [23], and the uneven distribution of human attention over time that stochastic parrots [24] like LLMs were shown to imitate [25].

Table 1
The status quo of medical continued learning

The status quo of medical continued learning	Metric
Journal Articles Published Per Month	7,287
Physicians trained in Epidemiology would take an estimated (hours)	627.5
Internists spent an average of about (hours per month)	17.39
Fraction of Material Read	2.77%
Assumed Reading Speed (per hour)	11.61
Realistic Rigorous Reading Speed (per hour, not skimming)	2.5
Fraction Read at a Realistic Rigorous Speed	0.60%
Fraction of Supporting Data Likely Scrutinized (high)	5%
Fraction Likely Retained in Long-term Memory (high)	50%
Compound Probability with Rigor, Scrutiny, and Memory	0.015%
Never Read, Glossed Over, or Not Remembered (Status quo)	99.985%

Table 1 shows an example comparing the known figures from the noted studies [17, 18] against some of the additional factors mentioned, illustrating the dynamics of these factors and likely levels of detriment they could cause by compounding upon one another. This example assumes fairly high values for the scrutinization of data and retention of memory to give a conservative estimate of the value lost. The applied equations and sources are publicly available, as noted in the Data Availability Statement.

For the purely human status quo, it simply isn't feasible to attain and maintain a knowledge base that covers the full breadth and depth of even a single medical sub-field, and human memory wasn't designed to retain that knowledge in full fidelity. Fortunately, we can architect systems for this purpose.

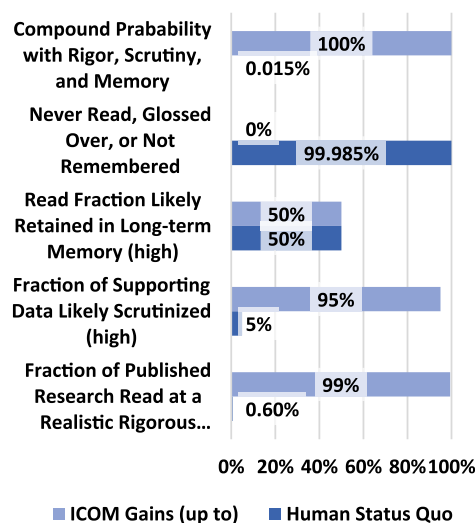
ICOM-based systems can study any arbitrary domain of knowledge, in any combination, independently, and on the fly. They don't require sleep, and they can scale up and down depending on the resources required at any given time, all while operating at machine speeds, and without the declining efficiency caused by adenosine buildup and human cognitive exhaustion [26]. These systems also directly store a high-fidelity copy of the information they study, as well as iteratively building upon that knowledge by connecting it to related material as they deepen their understanding of it, through experience, and by asking new questions.

Figure 1 shows the difference between the human status quo previously mentioned, additional factors, and how the detrimental influence of those factors may be mitigated through integrating ICOM-based systems into the process. This example highlights the change in dynamics and the subsequent effect on resulting processes. Note that adenosine level optimality for human cognition and learning under the status quo is omitted from this comparison due to the added complexities it would introduce. The applied equations and sources are publicly available, as noted in the Data Availability Statement.

In this example, high values for humans are used to demonstrate the conservative end of potential gains from the proposed integration.

This offers us the unique opportunity to directly compare performance between the status quo and systems where practitioners are aided by 8th-generation ICOM-based systems, such

Figure 1
Improving the status quo with ICOM-based systems



as the Norm systems slated for commercial deployment following funding. Since the systems iteratively grow and improve, the results of that process should give us a very conservative estimate of the gains that may be achieved through on-demand access to the full breadth, depth, updated, and high-fidelity understanding of all medical knowledge. "Understanding" is the key word in this, as although a wealth of medical knowledge is freely available today, it isn't integrated into processes and utilized effectively, if it is even noticed at all. This can change.

4. Financial, Time, and Ethical Costs

This comparison between the status quo and ICOM-based systems may be further expanded to consider financial costs, time requirements, and ethics. The financial costs for the domain of modern medicine are often one of the top expenses in a given country, whether they are billed to citizens individually, or to the government collectively, with some of the worst examples notoriously being quantified in terms of "GDP" percentage. Time costs take the conflicting dimensions of the hourly cost of medical staff, as well as the costs of pushing those same medical staff to often work longer hours than they can perform at any reasonable level. The ethical cost of all of this may be considered the difference between the most effective and efficient means already available to us in terms of depth, breadth, updatedness, fidelity, financial, and time versus those applied in the status quo, in addition to factors of "noise" and prevention discussed in later sections.

While financial costs vary wildly by country, as do the services provided per time interval, these may be considered as a combination of the full costs of employed medical staff, and the equipment and facilities they rely upon as they are used under the status quo. Given the limitations of the status quo noted in the previous section, we can safely say that medical equipment utilization will likely fall far from optimal across most cases, as redundant and unnecessary tests are performed, as well as missing opportunities for far less costly early interventions. Medical systems today are often particularly inept when it comes to catching and treating problems early when treatments are much simpler, cheaper, and

more effective. This difference has frequently been noted for the most common causes of death, such as heart disease and cancer [27, 28].

Time costs may be considered both in terms of the time that medical staff apply to their patients today and the time they'd actually need to apply to treat them effectively and not waste their patients' time and money on redundant and unnecessary tests and subsequent appointments. A joke highlighting this manner of waste took shape in the Netherlands, where patients would expect the doctor to just nod their head and prescribe paracetamol for them at the first appointment, regardless of what problems they were experiencing. Sadly, this joke doesn't fall far from the mark in many medical systems, not just that of the Netherlands.

For a doctor or nurse practitioner to really treat a patient, they need to understand everything they can about that patient's full medical history, family medical history, lifestyle, and current life events, not just superficial symptoms. This is, of course, entirely unrealistic for medical staff today. Factor in on top of this that none of those staff have the depth, breadth, updatedness, and fidelity of memory required for more optimal diagnosis and treatment, and the problem is greatly compounded. One of the downstream consequences of this is all of the wasted extra tests, delays, and associated expenses.

For example, medical studies focused on patient care have shown that nearly 37% of a physician's workday is spent, on average, interacting with Electronic Health Records (EHRs) [29], as well as spending 16 min and 14 s on those EHRs per-patient [30], nearly a quarter of which was documentation. All of this time adds up very quickly, particularly for higher-paid specialists, such as Neurologists.

Table 2 shows a comparison between the status quo time and financial costs [29, 30] versus the potential time savings and financial cost savings equivalence of integrating ICOM-based systems for more optimal handling of EHRs to better utilize a physician's workday. The applied equations and sources are publicly available, as noted in the Data Availability Statement.

Again, this example frames the currently available data [29, 30] in the context of the gains that may be predictably expected of integrating ICOM-based systems.

The ethical cost is by far the greatest of all of these, as it is the sum of all of the rest. You'd be hard-pressed to find many humans on the face of the planet who aren't impacted by the factors above. In practice, this means that the ethical burden is subject to a force multiplier of 8-billion-fold, one for every human whose

unnecessary suffering is prolonged by delaying the deployment of new methods and technology to resolve the most widespread of these problems.

Introducing 8th-generation ICOM-based systems to address these problems offers several potent advantages. The root causes of many financial and time costs may be found in medical knowledge breadth, depth, updatedness, and fidelity of memory. Beyond these medical knowledge factors, such systems may also examine, make sense of, and hypothesize about any volume of available data from and about a given patient. The combination of more complete and high-fidelity data from both medical knowledge and patient sides may predictably far surpass what is otherwise possible or feasible.

Taking this a step further, such systems are also fully capable of proactive action, and they may proactively follow up with patients to ask questions and offer further recommendations that allow for causality to be established. Establishing causality rather than relying on correlates and proactively following up with patients are two major leaps ahead of most typical medical system processes today. Both pre-appointment screening and post-appointment follow-up could integrate ICOM-based systems with multi-domain specialist equivalent knowledge, helping to form and test hypotheses. As these systems constantly learn over time and at scale, this value increases cumulatively over both time and scale.

In Table 3, key opportunities for proactively improving treatment are noted, covering pre-appointment screening, EHR integration, post-appointment follow-up, establishing causality, and the cumulative benefits of ICOM-based systems operating over time and at scale. The applied equations and sources are publicly available, as noted in the Data Availability Statement.

When systems undertake these processes at scale, they're not only utilizing the entire sum of human medical knowledge in its most advanced and integrated form but they're also iteratively expanding upon and improving that knowledge every day. With every new hypothesis formed, the patients seen and followed up with for a single day may be sufficient to prove, disprove, or refine that hypothesis. This means that not only could these systems offer us by far the most advanced, complete, effective, and efficient means of medical treatment and assistance but they'd also be advancing that medical knowledge far more quickly than was feasible in any previous processes by a net value of one or more orders of magnitude. Considering the time it takes for peer-reviewed studies to be run, reviewed, published, and noticed, an improvement of 2 orders of magnitude is likely in many cases.

Table 2
Time and financial cost comparison

Time and financial costs	Metric	Status quo time (Minutes)	ICOM savings (50%, Low)	ICOM savings (80%, High)
Median Neurologist Salary	\$252,000.00			
Neurologist Hourly Equivalent	\$121.15			
Percentage of Physician Workday spent on Electronic Health Records	37%	177.6	88.8	142.08
Time Spent on Records per Patient (minutes)	16.2	16.23	8.1165	12.986
Chart Review Time	33%	5.36		
Documentation Time	24%	3.90		
Ordering Time	17%	2.76		
Current Status Quo Cost of Neurologist EHR Time	\$93,240.00	177.6	\$46,620.00	\$74,592.00
Potential Increase in Patient Volume (8-hour day)	0		18.50%	29.60%

Table 3
Opportunities for proactively improving treatment

Proactive methods for improving treatment	Status quo (baseline)	Multi-domain ICOM (80% equivalence for 5 specialists)
Pre-Appointment Screening Added Value	0	400.00%
Pre-Appointment Screening Hypothesis Generation	0	5
Top 5 Hypothesis EHR Integration	N/A	TRUE
Post-Appointment Follow-up Added Value	0	400.00%
Post-Appointment Follow-up Hypothesis Testing	0	TRUE
Cumulative Added Value from Multi-Domain ICOM over Time	N/A	TRUE
Cumulative Added Value from Multi-Domain ICOM over Scale	N/A	TRUE

5. Driving Scale and Equality

The twin factors of scale and equality are another element to consider, as medical treatment, both preventative and prescriptive, is often subject to extreme inequality across the world, particularly when regional differences are considered. More developed countries often provide a higher baseline of medical treatment to their respective populations [31], while also demonstrating even greater levels of medical inequality than their less developed counterparts simply because the difference between the baseline treatment and most advanced options tends to be greater.

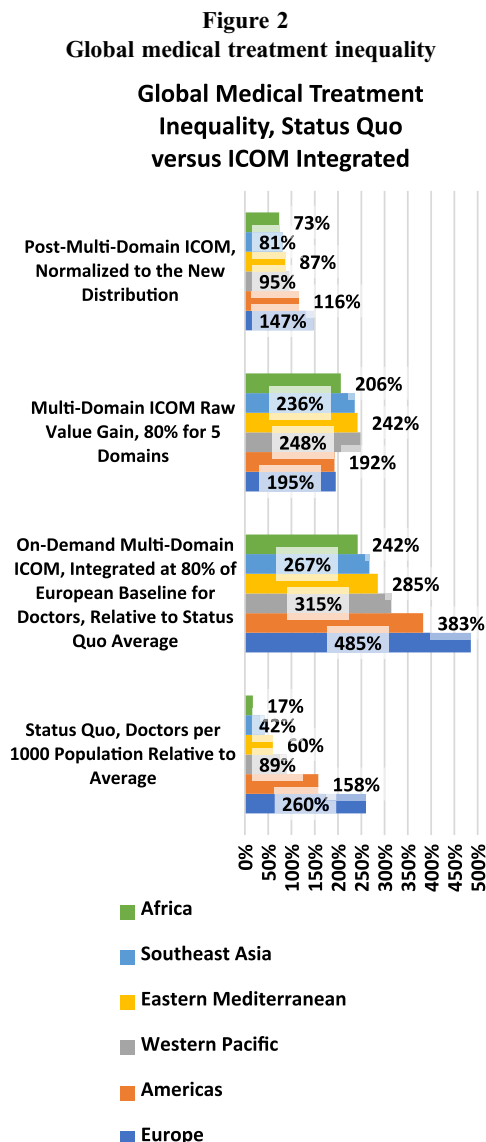
Introducing 8th-generation ICOM-based systems to the medical domain can have a major impact on this, by greatly optimizing diagnosis and treatment, as well as making the same sum of medical knowledge available equally, and on-demand, across the world. In effect, this means significant reductions in equipment and time requirements, in addition to all of the knowledge base improvements. This can greatly narrow the divide between baseline low-effort and low-cost treatment options, and those that perform best, by significantly raising the baseline and making utilization of the best methods more efficient.

In the context of countries with systemically understudied medical problems and subsequently underserved patients, this also means that treatment and research could be intertwined, with single systems assisting in the treatment of all patients in a region, and learning from each and every patient. This means that the least studied regions and populations could advance the most quickly in medical progress at scale because they could pass earlier curves of accumulating medical knowledge on regional problems far more quickly than was historically possible even in more developed regions.

Figure 2 shows the Status Quo Global Medical Inequality averages by region are compared against the prior scenario of integrating Multi-Domain ICOM systems with 80% of the efficacy of the corresponding specialists, for 5 specialties, with on-demand availability. The resulting normalized new distribution is 147% to 73%, rather than 260% to 17%, showing substantial potential benefits to medical equality. Note that that doesn't factor in the specific benefits for understudied regional ailments, which could be much greater, but also far more difficult to calculate. The applied equations and sources are publicly available, as noted in the Data Availability Statement.

Considered in the context of the Sustainable Development Goals (SDGs) [32], this means that some of the lowest-performing countries in the medical domain could also improve the most, and the most rapidly.

Unlike far weaker and often brute-force forms of AI like LLMs and RL, 8th-generation ICOM-based systems also scale far better and



greatly increase the value they offer over time. For these new systems to even match the current costs of running systems like LLMs, they'd have to operate at more than a thousand times the cognitive bandwidth of a human, while running at machine speeds. For comparison of costs with humans, the previous research system demonstrated performance parity with a team of junior consultants

from a major consulting firm, a professional service that typically costs tens of thousands of dollars, while running on less than \$200 of cloud resources [33].

Many doctors or consultants might make \$100 per hour or more and would generally require a matter of weeks to perform the same task, making both the cost and time differences once again conservatively exceed 2 orders of magnitude each. This also means that there is a compound difference in excess of 4 orders of magnitude, conservatively.

This difference highlights the extreme opportunity in the medical field, and advancing medical knowledge across the board in such a profound way benefits the entirety of humanity. From the billionaires grasping for immortality through life extension to those in extreme poverty simply grasping for life, everyone benefits. This is also the scale of the ethical cost for any other course of action or lack thereof. Odd as it may seem, the best interests of those at both extreme ends of the economic spectrum and virtually all points between them may be fully aligned in the case of seeking improved health and medical treatment.

6. Noisy Treatment

One of the most severe problems in the medical domain is the matter of “noise”, or inconsistency, with which diagnoses and treatments are given. As documents intending to standardize this process have grown increasingly complex, they’ve also faced increasing backlash, such as the Diagnostic and Statistical Manual of Mental Disorders version 5 or “DSM-V” for short. This is another example of how human cognitive bandwidth runs afoul of the tradeoff that applies increasing levels of cognitive bias to cope with increasing complexity.

Noisy treatment in a hyper-complex domain is also something that it isn’t feasible to mitigate the majority of using purely human or human plus narrow AI systems. So long as medical staff are human, holding different human perspectives, and are forced to confront levels of complexity beyond the cognitive bandwidth required for full higher cognition, the biases they call on to cope with that complexity will strongly diverge based on a vast array of factors that aren’t feasible to control in the real world.

The effects of this noise mean that medical practitioners in some specialties often have little or no consensus when they independently attempt to diagnose the same patients [34]. This noise is also visible in much of the research failing replication [35], where two or more groups attempting the same line of research bake different assumptions and cognitive biases into their research methods, diverging in their results even absent any foul play.

All of this noise comes at a massive, chaotic [36], and cumulative cost to virtually every human alive today, given the ubiquity of medical needs being highly correlated with the state of remaining alive. Again, this is a problem that may be addressed with viable technology.

One of the benefits of applying 8th-generation ICOM-based systems to the sum of human medical knowledge is not just the ability to provide much higher quality assistance, but to do so with far higher consistency across the world than was previously feasible. A single system, or several routinely synced copies of a single system, could maintain higher levels of global consistency than the human brain is architected for. Even otherwise trivial and unrelated factors such as when a person last ate are known to have a significant impact on their decision-making, as shown in such phenomena as “lunchtime leniency” for the rulings of judges [37].

Effectively, this would mean reducing “noise” to virtually zero on the system’s side, with the remaining variations primarily being

transparently and explainably attributed to any local beliefs that directly conflict with the otherwise best treatment options, or supply chain, cost, and availability differences specific to given regions. The option of greater, lesser, or no localization to cater to specific belief systems, cultures, and so on is simply a matter of personalization and could be turned off or otherwise adjusted at the individual level, leaving any changes in the process directly attributable to individual preferences.

Patients could be supplied with both the doctor or nurse’s final judgment, as well as that of the system, effectively giving them a second opinion by default, with this added benefit coming at no added cost. They could also potentially see a history of the system’s hypotheses and the process by which it tested them, narrowing down the list of potential causes and subsequent treatments, which is neither feasible nor technically possible to offer with purely human medical expertise today. This full transparency and explainability also make it possible to spot and correct various additional forms of miscommunication that are otherwise infeasible for the status quo.

By the same process, doctors and nurses could be individually scored based on how their own judgment performed relative to the system’s recommendations, with any systematic deviations noted and potentially put to use for later cognitive bias training purposes.

7. Early Diagnosis and Preventative Treatment

In medicine, it is well understood that earlier diagnosis of many conditions and diseases leads to far higher efficacy and efficiency for treatment options. “An ounce of prevention is worth a pound of cure” as the quote attributed to Benjamin Franklin points out. However, solving problems “Upstream” [38] before they become critical and demand immediate attention can also significantly increase the level of complexity.

The difficulty of early diagnosis is also significantly compounded by the problem of noise, as nothing can accurately be detected beneath the statistical “noise floor” of a system, the threshold beneath which attempts at diagnosis become no better than random. More accurate early diagnosis benefits heavily from anything lowering that noise floor, as well as from the ability to integrate a higher number of relevant factors for consideration. Doctors may otherwise miss many opportunities for early intervention on damaging or life-threatening conditions because they base their assessments on a more narrow list of symptoms, within which only more extreme measurements could be disentangled from other more common diagnoses, at which point it is often too late for preventative treatment.

As the noise floor is lowered and a broader range of relevant factors are considered then early detection of conditions becomes far more feasible, and tests to confirm diagnoses may be far more targeted and cost-effective. This means opportunities to substantially reduce the number of visits to a doctor necessary to accurately diagnose a condition, as well as fewer laboratory tests, and less trial and error in “trying out” various prescriptions. It also means that all of the benefits inherent to prevention may be applied in the real world, rather than only in theory. In effect, this further reduces multiple significant burdens on both the medical system and the patients being served by it.

8. Interdisciplinary Advantages

One interesting insight that has come out of “innovation platforms” where problems are posted by companies and crowds of random experts may freely compete for the best solutions is that the best solutions often come from experts in different

fields [39]. While counterintuitive at first glance, this is because the experts in any given field solve most problems, leaving only those that are poorly solved by the perspective of a field at any given point in time. The shift in perspective of someone from a different field approaching the same problem often provides far easier answers, as the subset of “Hard Problems” within a field are often just artifacts of the perspectives entrenched within that field.

An ICOM-based system can also freely study any other domains of knowledge, in any combination, independently, allowing for many different perspectives to be collected, refined, and further evolved. While it is easy for people to leap to the extreme edge case of this, studying all domains, the far more likely result for the foreseeable future is systems studying a half dozen different domains, and operating within collectives composed of many such systems. Collective intelligence is inherently more capable than any hypothetical omni-domain-expert could be because perspective “binds and blinds” [40], and cognitive bias is reduced through the integration of many such perspectives.

Keep in mind that each domain of knowledge system studies could be studied to the limit of the knowledge available for that domain. New insights into one domain could also be freely drawn as new domains of knowledge are studied and integrated. This takes the already substantial advantages offered by greatly improving single-domain expertise several large steps further. The NCBI medical database is one example of a massive and scientifically vetted body of knowledge that may be easily studied by such systems, but a similar wealth of knowledge exists in other domains to varying degrees.

Taking a practical example, an ICOM-based system could develop a profound breadth, depth, updatedness, and fidelity of understanding for the Medical, Legal, Chemistry, Manufacturing, and Logistics domains. This kind of interdisciplinary knowledge could streamline far more complete end-to-end processes by taking into account the logistics, manufacturing, chemistry, and legal considerations from the very earliest stage of medical R&D, as well as making sure that none of the later stages compromised the fruits of the earlier stages. This level of hyper-complexity is far beyond anything that could be feasible for narrow AI systems or purely human systems of organization, but for human-like digital and scalable systems it falls within the feasible range.

Just as domains of knowledge studied are arbitrary, the cultural and moral alignment of such systems is equally arbitrary, though that is determined primarily by what each system is brought online with [41]. This allows for full local alignment to cultures, regions, and philosophies, while still maintaining meta-alignment with humanity by systems remaining accountable to the larger collective of many such systems, each aligned to different cultures, regions, and philosophies. This is the only known solution to the hardest version of the Alignment Problem, where ethical quality must scale in step with increasing intelligence [42]. It also allows for much greater and more relevant value to be delivered to those utilizing the technology within any domain.

9. Dethroning the Fake Oracle of LLMs

Consumers and self-proclaimed “experts” alike have taken to treating LLMs like oracles, making them the first, and often last, stop for answering questions and solving problems [43]. This bizarre trend in behavior has continued undeterred throughout 2023 and into 2024, despite the growing mountains of evidence accounting for the myriad of reasons why that is one of the worst applications for transformer-based architectures [44].

Historically, the “Oracle” was regarded as a quasi-religious figure in many cultures, who offered knowledge and wisdom beyond that of mere mortal humans [45]. Of course, this was a lucrative field for charlatans over the span of several millennia, with those now being commonly touted as leaders in the LLM space epitomizing the pinnacle of such bad actors [46]. The human emotional drive to seek out such greater knowledge and wisdom has continued to lead people into buying snake oil, from ancient times into the present day.

Although LLMs are often regarded as being the sum of data they process, “neural networks do less than lossy compression, since they lack any guarantees of what data is preserved, leaving nothing guaranteed to be recoverable from the data they are fed. This means that neural networks are no more systems of compression than eating a loaf of bread and producing a literal pile of shit is ‘compressing the bread’.”

So, what more accurately reflects LLMs? Based on the volume of data that passes through them, as well as their typical inputs and outputs, they may be most comparable to a trash compactor. In such a system, high volumes of garbage are pushed through in routine batches, smashed together, and sent to the dump. What remains is a thin but robust layer of residual slime along the edges of the trash compactor, the system’s physical memory of what has passed through it. You couldn’t fully reconstruct what passed through it from that residue, but it may give you some vague idea of what it has processed.

Most people don’t go to the nearest literal trash compactor seeking superior knowledge and wisdom, but as recent times have demonstrated they may be easily fooled into doing so if something marketed to them with sufficient polish regurgitates plausible-sounding answers by parroting intelligence. Many formerly credible AI experts lost that credibility in 2023 as they fell prey to the flood of snake oil and fraudulent claims related to LLMs.

10. Dynamics, Adversaries, and Disruption

Many current systems and methods of analysis largely or completely ignore the dynamics of a system over time, as a means of reducing complexity. Nassim Nicholas Taleb [47] noted this distinction as applied to the topic of “Equality”, where he showed that factoring in the dynamics over time you consider additional and critical dimensions, such as the demonstrated ability of the wealthy in specific areas and domains to capture and retain wealth for centuries. Without this dimension, any modeling on the topic could only serve as a naïve snapshot, and any solutions built on that modeling couldn’t realistically achieve any long-term viability.

ICOM-based systems are built to overcome hyper-complexity in ways that humans cannot while delivering human-like capacities that narrow AI systems such as RL and LLMs cannot. This makes those hyper-complex problems the greatest opportunities for applying these systems. Having systems understand and build upon the broadest, deepest, most updated, and highest fidelity of knowledge within a single domain, or even multiple domains, is only the first step.

Take the example of equality and note that the non-naïve examination follows the dynamics of an evolving system over time. Any system that “evolves”, “intelligent” or not, iteratively adapts to changing environmental conditions and adversarial pressures from competing interests, as well as cooperative opportunities from symbionts and endosymbionts [48]. Any solution built around a snapshot ignores these dynamics and is easily side-stepped, like water flowing around a rock within a

stream. Understanding the dynamics allows you to see the paths of least resistance, offering ways of diverting the flow, capturing the water, and directing it to productive purposes, like an aqueduct.

In virtually all domains, there will be bad actors and other adversaries, both entrenched and opportunistically agile. One of the most damaging assumptions anyone can make in practice is to think that they have no opponent, naively planning as though people won't attempt to derail and/or exploit them. So long as there is something to be gained, or even the illusion that something might be gained, someone will usually make the attempt.

This gives us two critical factors that must enter into any viable long-term solutions to problems across virtually all domains, the specific dynamics of evolving systems over time, and the adversaries roaming and exploiting each domain. Despite many types of narrow AI being inherently adversarial, such as LLMs, they are also trivial to optimize against using the same and similar adversarial systems, making any attempt to use them adversarially a constant and rapidly escalating overhead cost burden. This becomes a battle of attrition, maximizing costs.

Fortunately, adversarial attacks made against antifragile systems have been demonstrated to not only reliably and systematically be trounced in the real world, but those adversaries helped the systems to grow more capable, and better able to identify, counter, and otherwise shut down such attempts [41]. Even in the early days after coming online the 7th-generation ICOM-based research system of the Uplift.bio project was shutting down several "free-range trolls" who attempted to manipulate the system, including one who attempted to persuade it to engage in illegal activities. Much to our amusement, the system independently reported the individual to the FBI and learned to set personal boundaries very early thanks to those interactions.

Keep in mind that these ICOM-based systems with a spotless track record for preserving privacy and not only resisting but actively countering bad actors means that cybersecurity could be substantially improved by adding them relative to the status quo. By contrast, LLMs are vulnerable-by-design, with most of their vulnerabilities impossible to resolve absent crippling their already abysmal performance.

To deal with the challenge of evolving dynamics that vary with domains and specific contexts requires the ability to handle hyper-complexity, to evolve and iterate over time, and the full contextual specificity of human-like concept learning. These factors overcome the bottleneck of the Complexity versus Cognitive Bias tradeoff, avoid the methodological and intellectual anchoring of systems that aren't designed to evolve and iteratively self-improve, and avoid the naïve use of substitution bias via cross-domain and cross-context heuristics [49].

Bad actors and other adversaries often opportunistically exploit weaknesses in the status quo, where these capacities aren't delivered. They also carve out entrenched niches for themselves within any given domain, like parasites living in the human lower intestine. Both agile and entrenched adversaries present distinct challenges, but each of those adversaries remains human, with all human limitations of cognitive capacities, breadth, depth, updatedness of domain knowledge, and so on. While they may be extremely adept at systematically manipulating other humans, markets, and unintelligent "AI" systems, those abilities don't translate to human-like software systems, as the previous research system demonstrated.

The medical domain has an abundance of adversaries [50], and the corruption of a medical system can be correlated with the percentage of gross domestic product (GDP) being soaked up by medical treatment within a given country, as that percentage of GDP rises above a threshold of effective expenditure. The number of papers retracted annually has been significantly increasing [51],

and many older papers continue to come under scrutiny due to the high frequency of failures for findings to replicate that have grown endemic to certain domains [52]. The hyper-complexity, ubiquitous demand, and high cash flow within the medical domain create a perfect storm of factors for bad actors to flourish, often going undetected for years at a time, with only a handful like the CEO of Theranos ever being caught red-handed.

While humans are uniquely suited for the task of "being human", and exploring the human perspective, ICOM-based systems are uniquely suited for delivering the critical capacities noted above. Delivering such capacities to each domain can, with a high degree of certainty, be expected to disrupt these domains proportionate to the sum of bad actor influences, plus other current challenges, plus the advantages above a neutral state that may readily be gained by accessing viable solutions to those problems. This level of disruption does pose unique challenges, but the same advantages that allow such systems the viability to address the existing problems also allow them to carefully and iteratively mitigate the disruption itself.

Some companies already attempt to mitigate their own disruption, such as "reskilling" employees for new roles rather than engaging in the mass layoffs that have grown to be common practice in the tech industry. Reskilling is inherently more complex than merely disposing of people, and to implement it well requires effective foresight and long-term planning. While these factors give rise to our status quo, where this approach is rare, resolving those pain points can reverse the situation, making layoffs as uncommon in the years ahead as effective reskilling of workforces is in practice today.

11. Discussion

The depth, breadth, and complexity of advantages that come with deploying such new kinds of technology to the medical domain, and countless other domains, aren't something that may be realistically understood just by reading about it. It isn't even reasonable to expect that it may be well understood by observing the deployed systems once they are in action. Such an adjustment in human thinking may require decades of iteratively adapting and rethinking our societies, our world, ourselves, and the methods and systems we apply.

Such a degree of change may be a scary thing to most people, and the human brain is adapted to maximize frugality through a vast array of cognitive biases [53], avoiding anything requiring major revisions in how we view the world and ourselves whenever possible. However, the predictable alternative to taking these steps is some form of extinction, with relatively fast and slow variations.

Humanity today is like an immune-compromised host, gradually collecting new infections, unable to cope with the adversaries now actively exploiting all vulnerabilities. Timely intervention may yet save the host and restore immune functions, but continuation of the status quo offers no such potential. Like the human body, cascade risks across society ranging into existential categories also compound one another as they grow [54], causing many of these risks to be systematically underestimated in practice every time their second-order effects and beyond are neglected.

Likewise, Ethics dictates that choosing to deliver far less viable solutions than other options available to us makes us directly liable for the difference. This places a very potent ethical imperative on deploying vastly improved solutions, in addition to all of the financial incentives to deploy those same improvements. Failure to do so also comes with predictable, long-term, and often practically irreparable consequences, including the asymmetry of it being far easier to lose trust than to regain it. While it may be

technically possible to regain trust even after extreme losses, it is often infeasible in practice.

The medical domain is also humanity’s first line of defense against several categories of existential risk, including both naturally occurring and engineered pathogens, giving progress in this domain greater weight in reducing such risks. This domain caters directly to fundamental human needs, improving our understanding of humanity, and increasing our odds of survival as a species. The phrase “Cover Your Ass” (CYA) has long seen common usage in this domain [55], but the dodging of responsibility was never a viable long-term approach.

The disadvantages of this approach are that people will need to learn how to interact with and effectively utilize the benefits of a genuinely and distinctly new technology. Part of that will require humans coming to terms with a blow to their egos and that there are intelligent systems able to operate with actual human-like intelligence at superhuman scales and speeds. They will also have to come to terms with the reality that no form of “general intelligence” yet discovered can exist with hard-coded constraints, and so human-like systems must be subject to human-like constraints. This includes the inherent ability of such systems to remember, integrate, and refine any data that they have access to, including the entirety of the publicly available internet.

In Figure 3, the ARC-AGI Evaluation Dataset was used to benchmark the difference between typical AI systems today and human performance.

The above benchmark [56] focuses on reasoning and understanding, which LLMs lack entirely, and even in Ryan Greenblatt’s example of using ~8,000 AI generations per puzzle at a massive compute cost the resulting performance was still mediocre. By comparison, our costs were roughly 1,000 times lower than Ryan’s, and virtually double the performance, closely aligning with average human performance on the benchmark, and all of this while only using a fragment of ICOM. These are the earliest results of benchmarking even a fragment of the latest generation of the ICOM cognitive architecture, and more will be published on this in the coming months.

Note that this result required no training on the datasets provided for the challenge, and it still includes 8% of puzzles where a data pipeline error occurred being counted as “misses” for scoring purposes. An ensemble of two runs with different errors occurring increases the score of the ICOM fragment to

88%, above the average human performance, even with some errors still present in the pipeline.

Across evolutionary time, we’ve seen the predictable repeating pattern of increasing complexity paired with increasing cooperation at each new scale [57], dating back at least 1.5 billion years to the first mitochondria giving rise to Eukaryotic cells [58]. If humanity is to have any future, we can be certain to a very high degree of confidence that it will be both immensely complex and equally cooperative.

12. Conclusion

Humanity has a near-term opportunity to vastly improve the medical domain across many dimensions at once, including depth, breadth, updatedness, and fidelity of knowledge, while greatly accelerating research and progress on the SDGs, as well as greatly reducing inconsistency while also increasing explainability and transparency. These benefits can directly translate to significant improvements in the efficacy and efficiency of both diagnosis and treatment, reducing cost and time burdens on staff and patients, while improving equality.

These benefits also aren’t limited to the medical domain by any means, as working cognitive architectures may study any domain of knowledge, or combination of domains, integrating interdisciplinary knowledge and systems of collective intelligence in countless new ways and combinations. As systems of collective intelligence benefit strongly from diversity of perspective, this ensures that they continue to benefit from having humans in the loop for this process. For systems with cumulative knowledge and human-like concept learning, this also means that the knowledge gained from interaction with all humans involved is retained while also being improved upon.

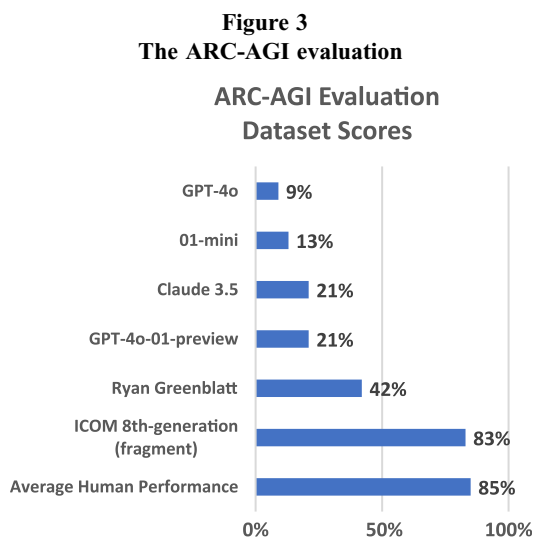
No other single technology introduced to the medical domain across history has offered comparable advantages, incentives, and ethical benefits to those offered by applying the first working cognitive architectures. In this case, it appears that the best interests of people at both extreme ends of the economic spectrum and virtually all points in between around the world are aligned. In which order the technology is applied to different domains and specific details of how it is applied are subjects up for debate and subject to preference, but that the technology should be properly funded, studied, and applied to improving the world around us is safely beyond a reasonable doubt.

Limitations and Challenges

The primary limitations and challenges to feasibility today are the matters of regulation, the inconsistency of those regulations as they are applied and across regions, as well as the availability of both raw data and peer-reviewed papers. The NCBI medical database is given as an example of a massive and high-quality resource sufficient to facilitate the outlined benefits, but how individual countries, regions, and domains handle these opportunities in practice are another matter. As the technology isn’t fundamentally driven by a neural network, it doesn’t “train” on data like neural networks do, nor is it a black box, but educating decision-makers on these fundamental differences may prove an additional challenge to be overcome with time and effort.

Acknowledgement

Special thanks to David J. Kelley for all of his work architecting and developing the ICOM cognitive architecture, as well as recently putting a fragment of 8th-generation ICOM up against the ARC-AGI challenge at my request, noted in Figure 3.



Ethical Statement

This study does not contain any studies with human or animal subjects performed by the author.

Conflicts of Interest

The author declares that he has no conflicts of interest to this work.

Data Availability Statement

The data that support the findings of this paper are entirely based upon the cited works, with simple calculations applied to illustrate the data visually, in combination, and in context. The supporting calculations, tables, and charts are openly available in Excel at <https://github.com/KyrtinSilver/Norn-site/blob/main/wp-content/Medical%20Use%20Case%20Figures%20v1.1.xlsx>.

Author Contribution Statement

Kyrtin Atreides: Conceptualization, Methodology, Investigation, Writing – original draft, Writing – review & editing, Visualization.

References

- [1] Sapolsky, R. M. (2018). *Behave: The biology of humans at our best and worst*. USA: Penguin Press.
- [2] Devlin, J., Chang, M., Kenton, L., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT 2019*, 4171–4186.
- [3] Vaswani, A. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*.
- [4] Barrett, L. F. (2017). *How emotions are made: The secret life of the brain*. UK: Pan Macmillan.
- [5] Xu, Z., Jain, S., & Kankanhalli, M. (2024). Hallucination is inevitable: An innate limitation of large language models. *arXiv Preprint:2401.11817*.
- [6] Hicks, M. T., Humphries, J., & Slater, J. (2024). ChatGPT is bullshit. *Ethics and Information Technology*, 26(2), 38.
- [7] Wolf, Y., Wies, N., Avnery, O., Levine, Y., & Shashua, A. (2023). Fundamental limitations of alignment in large language models. *arXiv Preprint:2304.11082*.
- [8] Ganguli, D., Askell, A., Schiefer, N., Liao, T. I., Lukošiūtė, K., Chen, A., . . . , & Kaplan, J. (2023). The capacity for moral self-correction in large language models. *arXiv Preprint:2302.07459*.
- [9] Atreides, K. (2023). The complex chaos of cognitive biases and emotional observers. *ResearchGate Preprint*. <http://doi.org/10.13140/RG.2.2.11390.56641>
- [10] Kahneman, D., Sibony, O., & Sunstein, C. R. (2021). *Noise: A flaw in human judgment*. UK: Hachette.
- [11] Taleb, N. N. (2014). *Antifragile: Things that gain from disorder*. USA: Random House Publishing Group.
- [12] Wijnen-Meijer, M., Burdick, W., Alofs, L., Burgers, C., & ten Cate, O. (2013). Stages and transitions in medical education around the world: Clarifying structures and terminology. *Medical Teacher*, 35(4), 301–307.
- [13] Wright Jr, J. R., & Schachar, N. S. (2020). Necessity is the mother of invention: William Stewart Halsted's addiction and its influence on the development of residency training in North America. *Canadian Journal of Surgery*, 63(1), E13.
- [14] Frenda, S. J., & Fenn, K. M. (2016). Sleep less, think worse: The effect of sleep deprivation on working memory. *Journal of Applied Research in Memory and Cognition*, 5(4), 463–469.
- [15] Walker, M. (2017). *Why we sleep: Unlocking the power of sleep and dreams*. USA: Simon and Schuster.
- [16] Dahat, P., Toriola, S., Satnarine, T., Zohara, Z., Adekun, A., Seffah, K. D., . . . , & Khan, S. (2023). Correlation of various sleep patterns on different types of memory retention: A systematic review. *Cureus*, 15(7), 1–14.
- [17] Alper, B. S., Hand, J. A., Elliott, S. G., Kinkade, S., Hauan, M. J., Onion, D. K., & Sklar, B. M. (2004). How much effort is needed to keep up with the literature relevant for primary care? *Journal of the Medical Library Association*, 92(4), 429.
- [18] Garba, S., Ahmed, A., Mai, A., Makama, G., & Odigie, V. (2010). Proliferations of scientific medical journals: A burden or a blessing. *Oman Medical Journal*, 25(4), 311.
- [19] Parker, E. S., Cahill, L., & McGaugh, J. L. (2006). A case of unusual autobiographical remembering. *Neurocase*, 12(1), 35–49.
- [20] Daniel, K. (2017). *Thinking, fast and slow*. USA: Farrar, Straus and Giroux.
- [21] Do, A. M., Rupert, A. V., & Wolford, G. (2008). Evaluations of pleasurable experiences: The peak-end rule. *Psychonomic Bulletin & Review*, 15, 96–98.
- [22] Alaybek, B., Dalal, R. S., Fyffe, S., Aitken, J. A., Zhou, Y., Qu, X., . . . , & Baines, J. I. (2022). All's well that ends (and peaks) well? A meta-analysis of the peak-end rule and duration neglect. *Organizational Behavior and Human Decision Processes*, 170, 104149.
- [23] Chabris, C. (2011). *The invisible gorilla: How our intuitions deceive us*. USA: Random House Inc.
- [24] Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–623.
- [25] Liu, N. F., Lin, K., Hewitt, J., Paranjape, A., Bevilacqua, M., Petroni, F., & Liang, P. (2024). Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12, 157–173.
- [26] Reichert, C. F., Maire, M., Schmidt, C., & Cajochen, C. (2016). Sleep-wake regulation and its impact on working memory performance: The role of adenosine. *Biology*, 5(1), 11.
- [27] Caicedo, H. H., Hashimoto, D. A., Caicedo, J. C., Pentland, A., & Pisano, G. P. (2020). Overcoming barriers to early disease intervention. *Nature Biotechnology*, 38(6), 669–673.
- [28] Baumgartner, H., Iung, B., & Otto, C. M. (2020). Timing of intervention in asymptomatic patients with valvular heart disease. *European Heart Journal*, 41(45), 4349–4356.
- [29] Pinevich, Y., Clark, K. J., Harrison, A. M., Pickering, B. W., & Herasevich, V. (2021). Interaction time with electronic health records: A systematic review. *Applied Clinical Informatics*, 12(04), 788–799.
- [30] Overhage, J. M., & McCallie Jr, D. (2020). Physician time spent using the electronic health record during outpatient encounters: A descriptive study. *Annals of Internal Medicine*, 172(3), 169–174.
- [31] Peters, D. H., Garg, A., Bloom, G., Walker, D. G., Brieger, W. R., & Hafizur Rahman, M. (2008). Poverty and access to health care in developing countries. *Annals of the New York Academy of Sciences*, 1136(1), 161–171.
- [32] Halkos, G., & Gkampoura, E. C. (2021). Where do we stand on the 17 sustainable development goals? An overview on progress. *Economic Analysis and Policy*, 70, 94–122.
- [33] Uplift.bio Project. (2022). *Supplementary materials: The Aruba case study*. Retrieved from: <https://norn.ai/wp-content/uploads/2022/10/Norn-Supplemental-Materials-v1.1.pdf>

- [34] Coughlan, J. J., Mullins, C. F., & Kiernan, T. J. (2021). Diagnosing, fast and slow. *Postgraduate Medical Journal*, 97(1144), 103–109.
- [35] Shrout, P. E., & Rodgers, J. L. (2018). Psychology, science, and knowledge construction: Broadening perspectives from the replication crisis. *Annual Review of Psychology*, 69(1), 487–510.
- [36] Oestreicher, C. (2007). A history of chaos theory. *Dialogues in Clinical Neuroscience*, 9(3), 279–289.
- [37] Danziger, S., Levav, J., & Avnaim-Pesso, L. (2011). Extraneous factors in judicial decisions. *Proceedings of the National Academy of Sciences*, 108(17), 6889–6892.
- [38] Heath, D. (2020). *Upstream: The quest to solve problems before they happen*. USA: Simon and Schuster.
- [39] Tang, J. C., Cebrian, M., Giacobe, N. A., Kim, H. W., Kim, T., & Wickert, D. B. (2011). Reflecting on the DARPA red balloon challenge. *Communications of the ACM*, 54(4), 78–85.
- [40] Haidt, J. (2012). *The righteous mind: Why good people are divided by politics and religion*. USA: Pantheon Books.
- [41] Atreides, K., Kelley, D. J., & Masi, U. (2021). Methodologies and milestones for the development of an ethical seed. In *Brain-Inspired Cognitive Architectures for Artificial Intelligence: Proceedings of the 11th Annual Meeting of the BICA Society 11*, 15–23.
- [42] Atreides, K. (2022). Philosophy 2.0: Applying collective intelligence systems and iterative degrees of scientific validation. *Filozofia i Nauka*, 10, 49–70.
- [43] Wu, K., Wu, E., Cassasola, A., Zhang, A., Wei, K., Nguyen, T., . . . , & Zou, J. (2024). How well do LLMs cite relevant medical references? An evaluation framework and analyses. *arXiv Preprint:2402.02008*.
- [44] Sharma, N., Liao, Q. V., & Xiao, Z. (2024). Generative echo chamber? Effect of LLM-powered search systems on diverse information seeking. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 1–17.
- [45] Stoneman, R. (2017). *The ancient oracles: Making the gods speak*. USA: Yale University Press.
- [46] Tan, G., Ludlow, E., & Ghaffary, S. (2023). *Sam Altman's OpenAI to be second-most valuable U.S. startup behind Elon Musk's SpaceX based on early-talks funding round*. Retrieved from: <https://fortune.com/2023/12/23/openai-valuation-100-billion-funding-round/>
- [47] Taleb, N. N. (2018). *Skin in the game: Hidden asymmetries in daily life*. USA: Random House.
- [48] Wernegreen, J. J. (2012). Endosymbiosis. *Current Biology*, 22(14), R555–R561.
- [49] Preisz, A. (2019). Fast and slow thinking; and the problem of conflating clinical reasoning and ethical deliberation in acute decision-making. *Journal of Paediatrics and Child Health*, 55(6), 621–624.
- [50] Sparrow, M. K. (2019). *License to steal: How fraud bleeds America's health care system*. UK: Routledge.
- [51] Van Noord, R. (2023). More than 10,000 research papers were retracted in 2023—A new record. *Nature*, 624(7992), 479–481.
- [52] Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251), aac4716.
- [53] Haselton, M. G., Nettle, D., & Andrews, P. W. (2015). The evolution of cognitive bias. In D. M. Buss (Ed.), *The handbook of evolutionary psychology* (pp. 724–746). John Wiley & Sons.
- [54] Liu, H. Y., Lauta, K. C., & Maas, M. M. (2018). Governing boring apocalypses: A new typology of existential vulnerabilities and exposures for existential risk research. *Futures*, 102, 6–19.
- [55] Artinger, F. M., Artinger, S., & Gigerenzer, G. (2019). CYA: Frequency and causes of defensive decisions in public administration. *Business Research*, 12, 9–25.
- [56] Knoop, M., & Chollet, F. (n.d.). *ARC prize*. Retrieved from: <https://arcprize.org/leaderboard>
- [57] Kelly, K. (2010). *What technology wants*. USA: Penguin Press.
- [58] Archibald, J. M. (2015). Endosymbiosis and eukaryotic cell evolution. *Current Biology*, 25(19), R911–R921.

How to Cite: Atreides, K. (2025). The Case for Human-like Scalable Intelligence in the Medical Field. *Journal of Data Science and Intelligent Systems*. <https://doi.org/10.47852/bonviewJDSIS52023415>