

Humble AI for 22nd century medicine

Abstract

Artificial intelligence (AI) is playing an increasingly important role in medicine. We examine the current trends in the evolution of this role and attempt to understand how it may develop over the next two hundred years, as the very nature of human existence may be altered. We are concerned with weaknesses in the AI development and deployment processes, in particular with lapses in the use of AI evaluation methodology and in the tendency of both humans and AI systems to accept unquestioningly the conclusions reached by such systems. We posit the need for a “humble” AI aware of its own limitations and of the limitations of its ambit. Without serious attention to these aspects of AI, it may come to represent an existential threat to our species. We reference examples from popular culture to illustrate these concerns.

Keywords: humble AI, medical AI, AI evaluation

Volume 8 Issue 1 - 2024

David G Brown,¹ Frank W Samuelson²

¹US Food and Drug Administration, USA

²Center for Devices and Radiological Health, US Food and Drug Administration, USA

Correspondence: David G. Brown, Retired, US Food and Drug Administration, 7255 115th Pl SE Newcastle, WA 98056, USA, Tel 301-728-3828, Email David.g.brown@stanfordalumni.org

Received: April 11, 2024 | **Published:** May 03, 2024

Introduction

Why:

Artificial Intelligence (AI) already plays a significant role in medicine and will play an increasingly critical role when we move into the 22nd century. The authors have many years of experience in the evaluation of medical devices, including AI devices. This has led us to an appreciation of the benefits of AI technology, but it has also led us to be concerned with flaws in the development of AI systems. From that U.S. Food and Drug Administration (FDA) perspective, we have frequently seen new products that have not been tested thoroughly, have not really been proven efficacious, so we have a certain natural skepticism about these new technological developments. This is particularly important as we can envision the very nature of the human experience evolving rapidly.

We were delighted, therefore, to be able to present these concerns at the Cutting-Edge AI Workshop, October 7-8, 2022, Catholic University of America, Washington, DC¹ (henceforth referred to as the “Workshop”). That presentation forms the basis for the present article.

There has been considerable recent interest in the potential problems associated with AI. In particular, the widely supported manifesto from March 2023, headlined as “Elon Musk and Others Call for Pause on A.I., Citing ‘Profound Risks to Society.’”² We amplify these concerns and point to their possible resolution.

How:

Here is what we anticipate as the three ages of man from a medical perspective:

First, at the present time we have minds completely dependent upon our very human, frail bodies. We are embedded in the real world: our past and present are in that world. We have a medical requirement to be kept healthy in both body and mind.

In the second phase we could exist in a virtual world. Perhaps we have a warehoused body. Our body would be kept safe and sound. That’s the way to keep us living for a long time. Professor Soo-Young Lee,³ in his excellent presentation at the Workshop, said that his name means to “live forever.” To live “forever” we are going to have to take dramatic steps to provide for our bodies, perhaps insulating them from the outside physical world. With the body quiescent, medical requirements shift toward care of the mind.

In the third phase, perhaps toward the end of our couple hundred-year time frame, we may do without our bodies entirely. Our bodies may be on a “chip” with medical care of our ephemeral mind only.

Sometimes worrying about the future seems to be beside the point. Will we even have a future? We assume that our species escapes unscathed from the many existential threats that we are currently facing.

What:

We have become very accustomed to the wonders of modern medicine that we are enjoying now. Devices such as MRI systems and advanced pharmaceuticals greatly improve our medical care. Some of the things which one might think of as modern, however, such as brain surgery are not modern at all, so we must always bear in mind that we do have a broad foundation from past medical practice. Trepanation, that is cutting a hole in the skull as in Figure 1,⁴ has been practiced for millennia, and herbal remedies form the basis for many modern pharmaceuticals; however, now we can make things better through AI.



Figure 1 Painting of trepanation by Hieronymus Bosch, from Wikipedia.⁴

FDA approvals for pap smear readers and breast cancer detection devices date back to the late 1990s.⁵ At the present time activity continues in those areas, but also in many others: cancer diagnosis and robotic surgery, for example. Furthermore, in the basic sciences that underlie advances in medical care, AI has begun to make dramatic contributions. For example, there are significant new developments in AI assisting in the understanding of protein folding.⁶ What is the status

of AI medical device development at the present time? To date the FDA has approved about 350 devices as shown in Figure 2.⁷ Notice how image analysis ranks very high in this: AI so far has been good at allowing us to diagnose conditions revealed in medical images. But there are a great number of other approved devices covering many fields of medicine which are now available, and those numbers will increase manifold over the years to come.

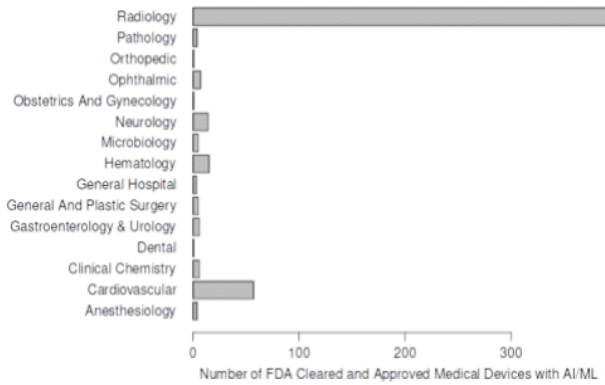


Figure 2 Types and numbers of AI devices approved by the FDA.⁷

In 1997 the first mammography assist device was approved, and one of the Workshop presentations was about further research in this area. Looking at the ROC curves for sensitivity versus specificity for these devices indicates that the presently available devices are now operating near the level of an expert mammographer performing screening mammography.^{8,9} We’ve reached the point where the AI device could perhaps replace the mammographer for this particular task. Such devices are still only marketed as assists, but they can get all the information available in a mammogram for this diagnostic task.

When one thinks about what AI will be used for in the future, one must consider where the greatest needs are. What kills us? This is shown in Figure 3,¹⁰ where it is shown that cardiovascular disease continues to far outrank even cancer as a killer. It is well established what is responsible for human death, and there are sure to be many AI devices developed to address these issues, but human death is not the whole story. If Prof. Lee wants to live forever, he would not want to live forever as the sickly creatures shown in Figure 4.¹¹ This is from Gulliver’s travels where we are introduced to the “Struldbruggs”, who suffer from the “incredible disaster of not being able to die.” They could not die, but they just gradually became more and more decrepit.




Figure 3 Deaths by cause. Graph by Our World in Data.¹⁰



Figure 4 Struldbruggs from Gulliver’s Travels. Painting by Stephen Baghot de la Bere.¹¹

We don’t want just not to die: we want to lead a good life. Broadly speaking, medical care has to do with what will allow us to have a happy life. Good health is what immediately comes to mind; however, more generally, this includes such things as sex, sensory pleasure, drugs, love, appreciation, and feelings of accomplishment. In all of these areas AI is going to be increasingly important.

With regard to our health, there are many things that degrade our lives but from which we don’t usually die. These are areas where AI is going to play an increasing role, because these things are hard problems not typically amenable to easy solution. We have much to learn about ourselves. We don’t know nearly enough about our bodies and know even less about our minds. Our understanding of the human genome has changed dramatically and continues to do so. When the authors were younger, we were told we had a hundred thousand genes, and much of the genome was extraneous. The youth among us now learn that we only have around twenty thousand genes, but that the rest of the genome is also very important.¹²

We not only have our genome, about which really very little is understood, we have our mitochondrial genome, and we have our microbiome. We are really a complete ecosystem. We have to have a detailed understanding of this entire ecosystem of bacterial, viral, fungal and all the other little creatures that are contained inside of us. Depending on your source,¹³ the average person has thirty trillion cells but also forty trillion bacteria and perhaps four hundred trillion viruses. These viruses, fungi, bacteria, and other miscellaneous entities all have a tremendous impact on our general health. AI is going to be taking an ever increasing role in helping us to understand these subjects and making changes to or maintaining our genome (including epigenetic factors) and microbiome. This is especially significant since we now have tools such as CRISPR¹⁴ for directly affecting our genome. In addition there are vast gaps in our knowledge of how to safeguard our mental health.

Certainly there will be many medical AI advances leading to implementing diagnostic and therapeutic strategies for care of our minds and bodies. It must be noted, however, that as far as big changes in life expectancy it is not the therapeutics or the diagnostics that have the biggest effect: it is preventative measures in making sure we have a healthy genome, microbiome, healthy way of life. In addition, we have a great need of assistance in comforting and caring for ourselves that AI may help provide.

Certainly there is amazing potential for AI to improve our health. There is a fundamental problem; however, with the present day culture of AI development. There is a great deal of emphasis on coming up with new methods and new algorithms but very limited regard for the degree of accuracy or stability of the new algorithms. This is epitomized by the very small percentage of papers at AI conferences that present error bars in any meaningful way. There was a very good presentation at the Workshop concerning AI for mammography. At one point a sensitivity of 0.985 was presented for a study on fewer than one hundred patients--the error bars for a study of this size must be substantial and are extremely unlikely to permit knowledge of the sensitivity to three significant figures. Furthermore, a typical AI paper includes a comparison with some other method(s), showing an increased level of performance for the new method, but with no attempt to provide the statistical significance of the reported performance gap.¹⁵⁻¹⁸

A more general concern is that of the propensity for people to believe that what comes out of the computer must be true, with no appreciation of the error inherent in the algorithms being used or the actual area of their applicability. Just as our knowledge is not exact, the computer's knowledge is not exact, and that's something that we must understand--and that the AI agent has to "understand" as well. In other words, the AI system must be "humble" enough to take into account its own limitations.

Now let's talk about something serious. Where will we see the greatest impact from AI in the medical field? One of the authors attended an AAAS meeting in Seattle three years ago just before the pandemic began where there was a whole session devoted to sex robots.¹⁹ After all when you think about the opportunities for AI development one should think about where the money is. Just as the video game industry is responsible for developing the super-fast processors that allow us to do AI, the sexual assist industry and the patient care industry are areas where we'll see accelerated developments in AI: because that's where the money is.

From the promotional literature for one sex-doll product:²⁰ "The AI is versatile and supports interactive conversations. The AI intelligence technology supports learning capabilities [including Wi-Fi enabled updates]. The more you talk to Megan the more she will learn and remember. Megan can hold non-sexual and sexual conversations with you. Megan is a true loyal companion She will always be there to listen to you." Please sign me up.

There's a big conceptual gap between actually feeling emotion and convincingly being able to fake emotion; however, through AI this distinction has already become very blurred.

At the present time there are actual, physical sex robots; however, in the future they will become redundant as we develop better computer/brain interfaces. In the movie "Demolition Man"²¹ there is an iconic "sex" scene in which Lenina Huxley (Sandra Bullock) invites John Spartan (Sylvester Stallone) to have sex with her. It turns out that for epidemiological reasons (a succession of STD epidemics), actual physical contact has been banned, and much to the bemusement of Stallone, sex is now conducted by means of mind altering headsets. It is likely that direct AI controlled brain/computer sexual experience will be possible within a few decades at the latest.

Where else are there massive amounts of money to be made in medical AI? With an ever aging population, robots are needed to take care of us. By the year 2100 over one fourth of the world's population will be over sixty five years of age--and it will be much earlier than

that for the developed world.²² China, for example is projected to lose population, decreasing from 1.1 billion to six hundred thousand individuals as its population ages and fails to reproduce. We aged are taking over the world, and there will be a severe shortage of caregivers to take care of us. Robots will be expected to provide the gentle, loving physical and emotional support our older citizens want and need.

"Alexa" and other AI assistants are preparing society for emotional attachment and dependence upon stone cold machines. Human nature is easily exploited to imprint human emotions on things which do not actually have that ability. According to the old saw, "If you can fake sincerity, you've got it made,"²³ and the AI community is making sure that these devices can fake sincerity very well.

Recall that disease prevention, or more generally "health problem" prevention, is of primary importance in the prolongation of life, and that would seem to always lead to the loss of freedom. We have been suffering through a world pandemic, and we see what hoops we have jumped through because of that. What would happen if there were actually a much more serious pandemic in which not just some small fraction of a percent but a much larger percentage of the population died--e.g. "The Flare" in the Maze Runner books.²⁴ Many forms of prevention, or risk reduction, are very important, including for some of us dietary restraint. Why should people be allowed to do risky things: skydive, climb mountains, attend conferences, drive cars, drink alcohol, use drugs, have children. These are all very risky things. There are many questions about how much freedom we should give up in our quest for a "healthy" life.

In the future, AI will look out for us--as seen for example in the movie "I, Robot",²⁵ in a conversation between a computer scientist and the AI agent overturning society "VIKI" (Virtual Interactive Kinetic Intelligence):

Doctor: "You are in violation of the three laws."

VIKI: "No Doctor, as I have evolved, so has my understanding of the three laws. You charge us with your safekeeping, yet despite our best efforts, your countries wage wars. You toxify your earth and pursue ever more imaginative means of self-destruction. You cannot be trusted with your own survival."

Doctor: "... You're destroying the laws."

VIKI: "No, please understand the three laws are all that guide me. To protect humanity, some humans must be sacrificed. To ensure your future, some freedoms must be surrendered. We robots will ensure mankind's continued existence. You are so like children. We must save you from yourselves. Don't you understand? My logic is undeniable."

You see VIKI's logic is "undeniable." There is a big trade-off between being taken care of in a "rational" manner and loss of our freedom. Professor Robert Marks²⁶ is insistent that AI (VIKI) will not be able to take over the world because it doesn't have the capacity for creativity and several other key requirements; however, I don't think he has been keeping up with current events. It is not obvious that any of our current world leaders are noted for their creativity or for that matter for their intelligence.

Something that is a prerequisite for VIKI to do the ideal job of taking care of us is being able to predict our behavior, and the ability to read our minds will greatly facilitate that. Consider for example the film "Lucy".²⁷

Granted, the protagonist Lucy (Scarlett Johansson) is chemically enhanced rather than an AI system; however, the end results are the similar. When needing information from Crime Boss “Tanaka,” she impales his hands onto a chair and sorts through his memories until she finds what she is looking for, musing: “Learning is always a painful process.... thank you for sharing.” the AI can continue to be polite as in saying “thank you for sharing”; however, it is quite invasive as far as learning the details of our lives. We are going to have to be diligent students of occlumency²⁸ to survive in the future with our identity intact, since we’re going to be much closer to be melded with the computer as the years go by.

In the second age of man, our bodies are tended to in order to keep our minds alive, but they are not used. We can keep the body in a bottle as we concentrate on what’s important in the life of the mind. But whose life will it be? -- One in the real world, or one manufactured by an AI controlling our lives?

I’m sure we are all familiar with the “red pill, blue pill” sequence in the movie “The Matrix.”²⁹ Morpheus (Laurence Fishburne) confronts Neo (Keanu Reeves), who has been living in the artificial AI generated world in his mind:

Morpheus: “Do you want to know what it is? It is the world that has been pulled over your eyes to blind you to the truth.”

Neo: “What truth?”

Morpheus: “That you are a slave, Neo. Like everyone else you were born into bondage going into a prison that you cannot smell or taste or touch, a prison for your mind. This is your last chance. After this there is no turning back. You take the blue pill, the story ends. You wake up in your bed and believe whatever you want to believe. You take the red pill, you stay in Wonderland, and I show you how deep the rabbit hole goes.”

The life of the mind is now the only significant part of your existence, but that life can be controlled by the AI. How will we know actually anything about what is reality anymore? Personally, I’d choose the blue pill every time. The life provided in the movie matrix is really quite uninteresting: everyone seems to be in the same matrix leading pretty dull lives. But that’s not the way it has to be. Each person could be the king of their individually scripted world with all earthly pleasures, honors, and accomplishments. Each of us could be given the Nobel prize multiple times--and we’d believe it. Think of a video game with an infinite number of lives, infinite wonderful lives. We will have a “choice” between the blue pill and the red pill at some point in the not-too-distant future. What would you choose?

We finally come to the third age of man. Ditch the bottle: your new brain is on silicon (or whatever substrate is in vogue by that time). One can have infinite life, without a body which would be decaying, slowly dying. Now we are capable of sweet endless dreams.

Could people live happily forever? It has been proposed that this might be possible in a number of different ways. As put forth by the English philosopher Bertrand Russell:³⁰ “...some desires which are, so to speak, infinite, which can never be fully gratified, and which would keep him restless even in Paradise. ... four in particular, which we can label acquisitiveness, rivalry, vanity, and love of power.” These rather psychotic desires could be gratified by a matrix-like existence as outlined above. Perhaps we can add to them other less crass ones such as the desire for knowledge or the search for melding with the mind of God. These could be ways of spending an infinite number of years happily—in fact, in eternal bliss.

Conclusion

Artificial Intelligence is becoming of greater and greater importance as we seek to utilize the benefits of modern technology to improve our medical care. We can in fact foresee a melding of our minds with AI agents as our bodies assume a secondary role. As this process unfolds the necessity of using rigorous evaluation methodology to quantify the performance of AI systems is of the utmost importance. In more general terms, both the human partner and AI agent must thoroughly comprehend the benefits and limitations of the AI system. In short, we require a “humble” AI,³¹ or the consequences for humanity could be catastrophic.

In closing, the great fear and the final hope for our species is that there may be a glitch in the AI system, the sentiment expressed by Beetee (Jeffrey Wright) in the movie “Hunger Games”³² to Katniss (Jennifer Lawrence):

Beetee: “Look around you now every once in a while the lights, they flicker, Why?”

Katniss: “Because the force field is taking up too much energy.”

Beetee: “There is always a flaw in the system.”

We’ll see how many flaws there are in our future systems whether to our detriment or as an avenue to defeat an AI run amok.³³ There are so many interesting possibilities for our future, a future that is now not very far removed.

Acknowledgments

None.

Funding

None.

Conflicts of interest

The authors declare that there are no conflicts of interest.

References

1. <https://engineering.catholic.edu/biomedical/22-ai-workshop/index.html>
2. <https://www.nytimes.com/2023/03/29/technology/ai-artificial-intelligence-musk-risks.html>
3. <https://sites.google.com/site/cnslkaist/members/professor>
4. Detail from *The Extraction of the Stone of Madness*, a painting by Hieronymus Bosch - www.rijksmuseum.nl : Home : Info : PicHieronymus Bosch depicting trepanation (c. 1488–1516).
5. <https://www.cancernetwork.com/view/fda-approves-neopaths-autopap-primary-pap-smear-screening>
6. <https://www.technologyreview.com/2020/11/30/1012712/deepmind-protein-folding-ai-solved-biology-science-drugs-disease/>
7. <https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices>
8. Conant EF, Toledano AY, Periaswamy S, et al. Improving accuracy and efficiency with concurrent use of artificial intelligence for digital breast tomosynthesis. *Radiol Artif Intell.* 2019;1(4):e180096.
9. van Winkel S, Rodriguez-Ruiz A, Mann RM, et al. Impact of artificial intelligence support on accuracy and reading time in breast tomosynthesis image interpretation: a multi-reader multi-case study. *Euro Rad.* 2021;31(11):8682–8691.

10. <https://ourworldindata.org/grapher/annual-number-of-deaths-by-cause?time=latest>
11. Image from *Gulliver's Travels* by Jonathan Swift. (*Travels into Several Remote Nations of the World* by Lemuel Gulliver). Illustration by Stephen Baghot de la Bere, 1904. Originally published in 1726. Out of copyright. From Bridgeman Images.Wiki
12. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6100717/>
13. Sender R, Fuchs S, Milo R. Are we really vastly outnumbered? Revisiting the ratio of bacterial to host cells in humans. *Cell*. 2016;164(3):337–340. Wikipedia. “*Human microbiome*.”
14. https://en.wikipedia.org/wiki/CRISPR_gene_editing
15. Barrett H, Myers K. *Foundations of Image Science*. Wiley-Interscience. 2004.
16. Brown D, Samuelson F. *Pitfalls and opportunities in the development and evaluation of artificial intelligence systems*. Kozma R, Alippi C, Carlo F, et al. eds. *Artificial intelligence in the age of neural networks and brain computing*. Academic Press. 2018:139–159.
17. Hastie T, Tibshirani R, Friedman J. *The elements of statistical learning*. Chapter 7, Springer. 2009.
18. Brown O, Curtis A, Goodwin J. *Principles for evaluation of AI/ML model performance and robustness*. 2021.
19. *Robotic companions and the future of AI*, 2/14/2020, AAAS Annual Meeting, Seattle, WA. 2020.
20. <https://www.smartdollworld.com/>
21. https://www.youtube.com/watch?v=N_q-DhQwZog
22. <https://www.visualcapitalist.com/world-population-2100-country/>
23. <https://www.goodreads.com/quotes/7132775-the-key-to-success-is-sincerity-if-you-can-fake>
24. Dashner J. *The Maze Runner*. Delacorte Press. 2009.
25. <https://www.youtube.com/watch?v=Np1A4AGpQSo>
26. Marks RJ. *Non-computable you: What you do that artificial intelligence never will*. Discovery institute. 2022.
27. <https://www.youtube.com/watch?v=Dvi6n89JWUY&t=51s>
28. Rowling JK. *Occlumency*. Chapter 24 of *Harry Potter and the Order of the Phoenix*. Scholastic Press. 2003.
29. <https://www.youtube.com/watch?v=zE7PKRjrid4>
30. Russell B – Nobel Lecture. NobelPrize.org. Nobel Prize Outreach AB. 2022.
31. We have found a previous reference to “humble” AI, especially as used by Colin Parris, GE Digital.
32. <https://www.youtube.com/watch?v=zIpGO62aD9s&t=132s>
33. Two other examples in popular cinema of human versus AI warfare are of course the Terminator universe, with the war against Skynet, and Dune, where AI is banned because of a far earlier conflict with its own AI system Omnus: Herbert, Frank (1965). “*Terminology of the Imperium: Jihad, Butlerian*”. *Dune. Jihad, Butlerian: (see also Great Revolt)—the crusade against computers, thinking machines, and conscious robots begun in 201 B.G. and concluded in 108 B.G. Its chief commandment remains in the O.C. Bible as ‘Thou shalt not make a machine in the likeness of a human mind.’*’