

Welcome to the 4sight Health Roundup podcast. 4sight Health's podcast series for healthcare revolutionaries. Outcomes matter. Customers count and value rules. Hello again, everyone. This is Dave Burda, news editor at 4sight Health. It is Thursday, August 10th, hard to believe. A wannabe fascist dictator who tried to overthrow our government is still walking around free despite being charged with multiple crimes, even harder to believe smart people in healthcare voted for him and will vote for him again if given the chance, you're the problem. Another problem we're trying to get a rope around is runaway innovation and artificial intelligence, especially in healthcare. Late last month, the Biden administration took the latest stab at it by getting seven big AI tech companies to voluntarily commit to using AI for good, not evil.

Those companies include open AI, which develop ChatGPT. The companies agreed to eight best practices in three domains. One, safety, two security, and three, trust of particular interest to healthcare is their commitment under the trust domain to use AI to address society's greatest challenges. Those challenges, according to the agreement include climate change, mitigation and adaption, early cancer detection and prevention, and combating cyber threats. There's not a healthcare organization in the country right now that's not dealing with at least one of those three, probably all three. Anyway, we talked about how to best regulate healthcare AI on the June 1st episode of the 4sight Health Roundup Podcast. Rather than reinventing the wheel, which AI is not supposed to do, we're gonna rebroadcast that June 1st podcast. Should we regulate healthcare ai, can we regulate healthcare ai, please enjoy the rebroadcast.

A day hasn't gone by since last November when OpenAI released ChatGPT without some announcement of yet another healthcare AI breakthrough. It's starting to freak people out so much so that people are calling for healthcare AI standards or rules or guardrails to make sure that when we ask Hal to open the pod bay doors, he'll open them to tell us what kind of rules we need to govern healthcare AI. Our Dave Johnson, Founder and CEO of 4sight Health, and Julie Murchinson, partner at Transformation Capital.

Let me tee up the conversation with a few things. Uh, way back in 2019, you remember, 2019, the National Academy of Medicine came out with its famous artificial intelligence and healthcare, the hope, the hype, the promise, the peril report, the report called for a Graduated Approach to regulation. The more risk the patients, the more of regulation needed. Now, fast forward to April of this year, and the Coalition for Health AI released its blueprint for trustworthy AI implementation, guidance and assurance for healthcare. And then in May, the World Health Organization issued a call for the safe and ethical use of AI for health. Things are moving fast, and many well-respected people are calling for something to stop AI from killing us. So I'm gonna ask each of you the same three questions.

Should we regulate healthcare ai? If so, who should regulate it?

And if we do regulate it, where is the line between, between promoting and stifling market innovation? In other words, how would you do it? Dave, you go first.

Well, Dave, first of all, in 2019, um, I don't know how you didn't bring this up, but the customer revolution in healthcare published in 2019, you know, the greatest book on healthcare ever, ever written. So just wanna make sure our readers check in on that fact. We got, um, but I'm <laugh>, I'm gonna go to your last, um, last question first about drawing the line and give you, uh, DJ's three rules for design regulation, the most important of which I call the Goldilocks rule. Um, not too hot, not too cold,

just, just right, you know, getting that balance that you talk about, uh, between, uh, protection and ensuring good behavior and not stifling, um, market innovation.

[\(20:46\)](#):

An example of regulation that's too cold, uh, would be the hospital pricing transparency guidelines. The penalty wasn't big enough, so most hospitals have just chosen to pay the penalty rather than go along with the reg. That's not a good outcome, too cold. Um, an example of one that too hot was in the early two thousands when Congress imposed a moratorium on specialty hospitals and completely shut down that business, uh, sent med cath into bankruptcy among others, um, be too hot. So balance not too hot, not too cold. Um, the second is, uh, the need for regulation, uh, depends on the potential for either harm or benefit. So the greater the harm or the greater the benefit, uh, the higher the need for governmental regulation. Um, as opposed to industry regulation, uh, by and large, um, industries haven't been terribly good at regulating themselves. So at the end of the day, I think we default to government for most regulation, and it's, it's an art more than a science.

[\(22:00\)](#):

Um, and when designing regulations, this is point number three. They should be as simple as possible, but not simpler. So clarity is essential. Needs to eliminate ambiguity. There may need to be a common lexicon. In fact, there almost always needs to be a common lexicon and standardization whenever and wherever possible. Um, so as simple as possible, but not simpler. So it has to get the job done, but it should do it in, um, most efficient way possible. Again, that's as art, as much as a science. There are always unintended consequences, uh, with regulation. Uh, so the regulations themselves have to adapt to market circumstances, just like they've increased the prices on hospital transparency violations. Uh, but just think of all of the unintended consequences when we, we put in, uh, regulations to stop behavior, abortion restrictions, minimum wage hikes, uh, endangered species act. There's always a, um, a counter reaction.

[\(23:06\)](#):

So getting that balance right, trying that line in the right way, not too hot, not too cold. More important, when, when consequences are are, are big. And then, um, as simple as possible, but not simpler. Those are DJ's rules for regulation. You heard 'em here first. Um, when it comes to AI regulation, that's an awfully broad topic for a two to three minute allocution. So I'm gonna focus on what I think is AI's biggest potential benefit in healthcare. And that's earlier in better diagnosis. Uh, I've just written a piece that we'll publish, uh, next month in the H F M magazine called Diagnostic Determinism. And I'm looking at the combination of big data with our much advanced understanding of the human body's mechanics, so all the omics, uh, genomics, proteomics, uh, and so on. Um, and when you do longitudinal studies in different mediums, so blood scans, signal devices, um, you ma you build up these massive data sets.

[\(24:18\)](#):

We then unleash, uh, analytics machine learning, looking for patterns, uh, finding correlations, ultimately finding very, very early indicators of disease correlated to, um, largely chronic conditions. And oftentimes very serious chronic conditions. You know, heart disease, cancers, Alzheimer's, that type of thing. And imagine a not too distant future day, um, when we will get these indications three, four years in advance with more than enough time to change lifestyle behaviors, um, take drugs, uh, do whatever is necessary to either slow the course of disease progression or reverse it entirely. Um, and there'll be continuous, um, learning. Another great thing about this is the ability to do retrospective analysis. So looking for patterns in retrospect, um, and honing, um, the algorithms in that way. Um, so the question is, you know, this means I think we're going to have a lot more preventive healthcare 10 years from now than we do today.

[\(25:30\)](#):

Prevention is only 3% of the budget. What's it in 10 years, 10, 15, 20, 25, whatever it is, it's gonna be very disruptive to the current healthcare, um, business models, but ultimately very good for humanity. So, should we regulate that? Huh? Really good question. Um, my initial instinct is we shouldn't regulate it any more than we do now when developing treatment protocols. Unless, and this is a big caveat, there are serious unintended consequences that I don't see right now. Um, false positives are less concerning than false negatives. We also need to avoid over-testing, you know, uh, on that. So, uh, that's, that's again, the art and science of this.

Let me wrap up with, um, one example of generative AI that really has me on edge, actually scares the crap outta me, and that's hallucinations what we're discovering, uh, with chat G p t four now, which makes chat G p t three look like a toy, is that it can produce in seconds a very thorough report on any topic of your liking in any language, uh, with footnotes. But what we've discovered, and the, um, technologists don't know why is some of the footnotes, some of the references, uh, are either wrong or they go to dead ends. That's what they call hallucinations, and they don't know why they occur. And I'm thinking in a not too future date when the machines are talking to the machines and we don't really understand how they communicate with one another. I mean, that's the whole beauty of machine learning, is they, they see things that human beings can't, that the potential, um, for disinformation errors, maybe even malicious manipulation, um, increases.

[\(27:51\)](#):

So when some of the technologists called for a six month delay, so we can understand this a little bit better. I was, I'm in that camp, obviously that hasn't happened, but I worry about, uh, hallucinations. I don't think it gets us all the way to Hal Dave. I mean, your, uh, your fear and my fear because we're both named Dave. Um, but, uh, you know, it, it's definitely something to worry about and pay attention to <laugh>. Yeah. It's like kind of a, a G P s to nowhere, right? That is pretty scary. Thanks Dave. Uh, Julie, any questions for Dave?

Dave, when I really sat back and looked at our Congress in particular, and their experience based and their age and their ability to think about regulating ai, now of course, we're good at drawing on experts. I get that. What is the w h o doing in this whole game? Did you see the fact that the W h O has identified six core principles that it's now saying should be part of the regulation of ai? What is, has a W H O been involved in really driving regulation in the past? It's crazy.

the WHO has done incredibly well in terms of regulation is coming up with surgical checklists that they've implemented broadly all over the world. Um, and with dramatic improvement in, in outcomes, um, standardizing lexicon like I talked about earlier, um, making sure that everyone is on the same page, um, making sure that all the preventive steps, um, and necessary checks have been in place. Uh, so just standard checklist stuff like happens on an airline or anywhere else. So if what the W H O is promulgating here, um, with, with regard to ai, that's worth paying attention to with all of the caveats I did in, uh, talked about before in, in DJ's, rules of regulations.

[\(30:16\)](#):

So, uh, get the temperature right. Um, you know, surgical interventions are pretty serious. So regulation of those, uh, through checklists is a really smart thing to do. Low cost. Um, and then really understanding, uh, the trade offs, making them as simple as possible, but not simpler. So hopefully the A W H O is acting in that regard.

Yeah. Interesting. Thanks Dave. Uh, Julie, it's your turn. The, uh, same three questions. Should we regulate healthcare AI? And if so, who should regulate it? And if we regulate it, where is the line between promoting and stifling market innovation? Or what would you do?

Well, I mean, this is so much bigger than a healthcare issue, right? I mean, I don't know if you guys saw yesterday the Center for AI Safety, which is some no name, you know, organization, from what I can tell from friends I know in the tech history, put out a 22 word statement that was signed by over a hundred researchers, academics, C-suite executives, like the c e O of DeepMind with Google, the c e o Open AI that said, and I quote, mitigating the risk of extinction from AI should be a global priority alongside other societal scale risks, such as pandemics and nuclear war <laugh>. Just let that sink in

Speaker 1 ([31:57](#)):

A little. Yeah. I don't, I don't like the word extinction, but, uh, extinction, please continue. <laugh>, that's so, no, that's a hell of a unintended consequence, isn't it? Uh, <laugh> the end of life as we know it.

Speaker 1 ([32:11](#)):

And I just want somebody to, to open the door so I can get in, right,

Speaker 4 ([32:17](#)):

<laugh>, right.

It's all a time game.

Speaker 1 ([32:23](#)):

Right?

Speaker 4 ([32:24](#)):

Um, so, you know, this is serious stuff. I mean, just issued yesterday, signed by hundreds of leaders, um, who I think Dave, are some of the same people calling for kind of a pause at the same time, regulations happening. China, the eu, Brazil among others, have already drafted, you know, completely unique pieces of legislation to regulate AI and their countries. So let's just stop for a minute on that one and ask ourselves, could you actually have a fragmented approach to this that wa works in our global economy? Like, it just, it doesn't make a lot of sense to me, <laugh>. And when I looked at what some of the regulation is, the EU is most extensive. They're looking at regulating both the inputs and the outputs. So think about in healthcare are inputs, what, here's this week's quiz. All right, get ready. Ready.

Speaker 1 ([33:17](#)):

All right. Boy, I've got a winning streak going here, Dave. I don't wanna blow it. Okay.

Speaker 4 ([33:24](#)):

Let's see if I can ask the question well enough that you might have a chance of getting it. So, what's one of the most overused phrases in healthcare about the data that sits in claim systems or EHRs that, you know, potentially poses one of the biggest issues or risks with AI?

Speaker 1 ([33:46](#)):

All right, Dave, you want to go first? Or you want me to take a stab? I'm gonna let you go first. Uh, I've got an idea what I'm gonna say, but, uh, I'm gonna say bias. That's a pretty good guess.

I'm gonna say, Julie, prior authorization.

Speaker 4 ([34:29](#)):

You know, what I love about you guys and us really, is that we're so inside baseball that, um, sometimes you don't think simply enough. So wrong on both counts, but very good answers, <laugh>, <laugh>. And the answer is, I, I have heard this time and time again, garbage in, garbage out. We've said this, Dave, right? Um,

Speaker 1 ([34:49](#)):

Yeah, yeah.

Speaker 4 ([34:51](#)):

in the AI world, you know, they talk about that as value and value out, but in healthcare, it's garbage in. We all know that. If you just put this in context, we have a decades of upcoding and frankly, at fraudulent levels, right? We have patient information edited from the beginning really to address liability concerns.

We have patient information that's actually sometimes exaggerated for coverage. So regardless of whether our data is representative from a health equity perspective or bias verree answer, is it actually representative of the real care provided or is it crafted for other motivations? So while it's complex for Congress to consider regulating inputs like that, there are inputs in general, like we all know that healthcare is a mess when it comes to our inputs, <laugh>.

([35:58](#)):

So how about the outputs on our side? Should we, you know, I guess, um, limit certain uses of output. Should we require disclaimers about the accuracy of work that's based on ai? I mean, there could be things that could be useful. I don't know. So the more I looked into what's really happening around regulation here, the scarier it gets when you start to really break it down in healthcare. But the fun discoveries I had, first of all, Dave, I agree with you. Just the fact that we have a term called AI hallucinations makes me laugh and it's terrifying. Um, you know, some of, some tech companies talk about AI as really helping them, um, be able to do what they can do with analytics, but just at warp speed and be able to see customer preferences shifting on almost like a realtime basis.

([36:56](#)):

And customer preferences are not something that healthcare really prioritizes or even is built around knowing how to harness, right? We, we don't think about really how customer preferences might shift our products. So, um, some of the most traditional uses of AI aren't even really applicable to us in current day. Uh, and there's also, um, you know, I guess on the upside, I also, uh, see how there are some, um, I guess applications of AI that are really looking at how do you actually choose data sets that allow you to control, um, you know, almost like more of a micro dataset way of controlling what you can act on based on maybe the insights of, of other types of data sets in historically. And Dave, this kind of gets to one of the points you made about, um, could it be used in a retrospective way, in a really powerful way?

([37:53](#)):

So mm-hmm. <affirmative>, there's downside and upside for sure. But one thing that scared me to death was, uh, the fact that, I don't know if you know this, but, um, some of the data, or I guess the data collected in, uh, I think it was opening Eyes program hasn't been updated in over 18 months.

And there have been bugs that have shut chat gpt down, for instance, one recently where users were able to see the titles of other users' conversation threads, which, you know, took them offline for a while until they could fix it. And I, I just don't think with all the hallucinations with, um, the dirty data with what is happening in the program, that if there are bugs in there, which of course are, cause every technology has bugs, like we, we do need to slow down and figure out what we're doing here. But, you know, large language models are being used all over the place to, in healthcare, around patient engagement, communication, um, informing hospital, a DT choices, making waves across healthcare workforce. Um, so there, there are a lot of ways it's being used today, I would say, I would argue kind of in back office ways that are probably highly beneficial to healthcare deficiency.

Speaker 1 ([39:14](#)):

Wow. So Dr. R g creep could get me a heart transplant when all I needed was an angioplasty <laugh>.

Speaker 4 ([39:23](#)):

That's right.

Speaker 1 ([39:23](#)):

Right. Yeah. Watch out.

Speaker 1 ([39:26](#)):

Wow. Okay. Dave, any questions for Julie? Well, as long as I don't get cide when I want an aspirin, I think <laugh>, <laugh>, we all have our issues. <laugh>.

Julie, you know, kind of based on everything you were throwing at us, um, and given the fast-paced development and distributed character of AI applications, can we, even if we really want to actually regulated use, uh, aren't there enough end runs around potential regulations and the old guys to put 'em in place, um, to render them essentially ineffective? Isn't, isn't the market almost moving too fast? What do you think?

Speaker 4 ([40:41](#)):

Well, I do feel like this topic has got me a little worked up this week. So <laugh> thanks beta for the, uh, the heart conversations.

Speaker 1 ([40:49](#)):

Sure. You're welcome. Thrown

Speaker 4 ([40:51](#)):

That at you. I, I've actually been quite interested in trying to learn more and more and more, um, you know, uh, HBS has had a lot of traffic back and forth on this, on a couple things that they've written. And I found one comment really interesting as it relates to your question, Dave. Uh, someone said that, you know, AI can only do what human intelligence can do because it's all based on our work over time, right? The broader, our, so AI just does it faster, and you can't regulate human intelligence. You can really only regulate the ex expression of that intelligence. So likewise, this person would argue that you can't regulate ai, you can only regulate its expression and use. And of course, governments like China are trying to regulate ai, which is kind of imposing the same kind of censorship that they impose on

This transcript was exported on Aug 07, 2023 - view latest version [here](#).

human expressions. So we're gonna have to ask ourselves these questions that, um, are very nuanced and probably difficult parse to really do something meaningful here if it can be done. Dave, to your point

Speaker 1 ([42:02](#)):

That that's interesting. Now, uh, when I wrote this script, I had myself coming down in the middle, and then after Dave's answer, I'm like, no way we should regulate this. After Julie's answer, I'm like, we have to regulate the hell out of this. So I think the only thing I can do is to try to stay healthy, right?

Wow. What a great discussion. Thank, thank you both.

Outro:

I'm back and it's still August 10th. We hope you enjoyed the re-broadcast of our June 1st podcast. "Should We Regulate Healthcare AI? Can We Regulate Healthcare AI? How would you answer those two questions? If you'd like to learn more about the topics we discussed on today's show, please visit our website at [4sighthealth.com](http://4sighthealth.com). And don't forget to tell a friend about the 4sight Health Roundup podcast. Subscribe now and don't miss another segment of the best 20 minutes in healthcare. Thanks for listening. I'm Dave Burda for 4sight Health.