

Seizing Life, episode 134
Artificial Intelligence and Epilepsy: The Promise & Pitfalls of AI in Diagnosis and Treatment
Guest: Dr. Daniel Goldenholz
(Transcript)

Kelly Cervantes: Hi, I'm Kelly Cervantes, and this is Seizing Life, a monthly podcast produced by CURE Epilepsy. This month on Seizing Life, we take a look at something that's in the news and on people's minds a lot these days, artificial intelligence or AI. It seems like we are constantly hearing about the potential advancements and problems that AI promises for the future in many different areas. Today, we're going to explore AI as it relates to neurology and epilepsy, how it is currently being used in neurology, and what the future of AI may look like within the areas of epilepsy diagnosis, treatment, and research. My guest today is Dr. Daniel Goldenholz, who could not be a better person to help us understand the promise and pitfalls of artificial intelligence.

Dr. Goldenholz is an assistant professor of neurology at Beth Israel Deaconess Medical Center in Boston. He also leads the Goldenholz Epilepsy and Data Science Lab there, which combines epilepsy research with machine learning and data science statistics to improve the lives of epilepsy patients. Dr. Goldenholz, thank you so much for joining us today. I have been so excited for this conversation ever since I had the opportunity to hear you speak at the Partners Against Mortality in Epilepsy meeting ahead of the American Epilepsy Society Conference this past December, and was just blown away and immediately reached out to our producers to make sure we could get you on the show. So thank you for agreeing to come on.

To start off this episode, can you give us just a very basic definition of what AI is because, I think, a lot of us think we understand. My perception is a computer system that can process a whole bunch of information really fast and some days, going to take over the world a la Terminator. Am I wrong?

Dr. Daniel Goldenholz: First of all, thank you for inviting me. Thank you for having me, and secondly, no, I don't think that that's the definition that I'm going to go with. Intelligence is something that's difficult for us to understand and to define, but we kind of know it when we see it. Our cats and our dogs seem to be intelligent. Our colleagues, at least some of them seem to be intelligent, and we hope that we are intelligent, and those difficult-to-define ideas of intelligence are being translated into machines.

Long ago, we had very simple machines that made simple decisions based on simple rules, and nowadays, we have moved on to much more advanced systems that are able to make much more complicated decisions, predictions, and control systems that behave in ways that are like intelligence that we know and love, and that's the fuzzy definition that I would use for artificial intelligence. If we can make things that are smarter and smarter, I think that we will help humanity more and more, and not go to Terminator and kill us all, but actually, just make life better and better.

Kelly Cervantes: Awesome. I like your version much better than mine. Can you give us, I guess, a little bit of a history here as to what makes this recent boon in AI, because it feels like we've been on the cusp of this for a while, but now, all of a sudden, everyone's talking about it? It's everywhere. Why now? What is the difference?

Dr. Daniel Goldenholz: Yeah. AI has been around for a very long time, but in roughly the last decade or so, there's been some major developments both in, in terms of speed of computers, in terms of better algorithms, and in terms of new uses that people have demonstrated work really, really well, that we're beginning to get to the point now where I can take off the shelf tools and deploy AI on new problems that have never had AI deployed on them before and get really great results. The other big thing is that information, more than ever before in the humanity, is now in digital form. So I've got all these cool toys, I've got all these great algorithms, and I've got amazing sets of data, and when you combine all these three things, now, for the first time, we're able to do things that we could never do before. The biggest thing that happened in the past year or so is the explosion of AI in large language models. Many people have heard of ChatGPT, and that really brought an introduction to a lot of people for the first time, that AI is much more advanced than people realized.

But what's been happening behind the scenes before that introduction is there's been a huge, huge growth and explosion in these technologies that have gone way beyond just ChatGPT in a lot of different domains.

Kelly Cervantes: Thank you for that. That makes so much more sense. Now that we have a general understanding of AI and how it has become relevant to the layperson, I wonder if you can explain to us more specifically how AI is being used in neurology. And we'll sort of get to the epilepsy piece in a minute, but just sort of in the broader world of neurology and understanding the brain and how it works, how is AI being used there?

Dr. Daniel Goldenholz: So I'm going to answer you broadly, and then specifically. So broadly, I think that there's four kind of levels of this. The first is that years ago, we started to think AI might be useful, and we found various types of experiments that we could do, and we found that AI might be able to help us, but it's been kind of iffy and didn't work that well, and doctors weren't so sure that they wanted to really go down that road. Then, we move to level two, where we say, "No, AI really is useful," and that's where we are today, where we have some actual FDA approved neurological tools that are in our arsenal. Some examples of that is you can detect stroke using CT scans, using AI tools that are FDA cleared.

You can detect seizures with an FDA cleared seizure watch, and certain EEG and sleep studies can be read in part with AI assistance. So that's level two, and level three is when we get to the point where AI is required. We're not there yet. We are currently clearly in level two, but someday soon, we're getting to the point where AI is going to be necessary in order to do our jobs as doctors and as neurologists, and at that point, there's going to be so much information and so many different possibilities, that no one human will know them all and

understand them all and can keep it all in their head. At that point, we are going to rely on AI technologies in the same way that we, today, we rely on certain imaging technologies.

Finally, we'd move to level four, where not only do we need AI, but AI can take on some of the functions that we currently do as clinicians. At that point, we would move out to be more of manager of AI systems and help bring more humanity to the doctor-patient relationship, and do less of the simple day-to-day drudgery, if you will, that doctors do, and do more of the living with uncertainty and being there for patients, whereas figuring out what's the name of the drug and what's the name of the test that we need to do can be outsourced to a machine that does it better than us.

Kelly Cervantes: Okay. There is so much to unpack there, but all really, really exciting. I want to back up a little bit and go back to ... So you mentioned a few of the ways that AI is currently being used to treat patients with epilepsy. You mentioned the Empatica watch, reading EEGs.

What does that look like as that continues to progress in the near future? I know you mentioned that eventually, doctors are going to have to rely heavily on AI. It will be integrated into a clinical setting. What needs to happen, I guess, for us to get from a place where we're using it to assist with reading EEGs to a place where doctors are utilizing it in their daily practice?

Dr. Daniel Goldenholz: Well, I read EEGs on a regular basis here at Harvard Beth Israel, and we use AI in our EEG reading room on a regular basis in order to make decisions about how much EEG we need to continue reading, using a simple score-based system called 2HELPS2B, and that's available for free to any clinician reading EEG, but we also use AI in the Persyst software, which is an EEG reading and interpretation software that helps us identify patterns that are concerning. Right now, it's just kind of helpful. Both of those tools, they give us a clue and we can decide to pay attention to them or not. As things progress, we're going to get more and more expert assistance to the point where the AI will simply read the EEG and say, "Look, this is what I found. Here are the highlighted areas that I'm concerned about. This is what I recommend."

And we'll still be able to go and look and make sure that we agree or disagree with AI, but those days are coming soon. We've just saw a study that came out about six months ago now from Sándor Beniczky's group, where they were basically able to take in a complete 30-minute EEG and come up with a diagnosis for the whole thing. So we're seeing a lot of advances in automating the reading of complex tests. We're going to see that in EEG and imaging, and in wearable devices so that doctors don't have to dig through every little second of every little piece of information, and instead, they can step back, and like I said, focus a little bit more on the humanity. "What do you need, my patient, and what do I need, as your doctor, in order for us to reach some sort of consensus about how to best take care of you?"

Kelly Cervantes: Which sounds incredible. I mean, how many people go into the medical field because they want to help patients, and then realize that they're not actually in the clinic with the patients. They're sitting in their office, looking at paperwork and filing, and so, not to mention, the incredible benefit that that is for the patients to actually get FaceTime with a doctor, which is so incredibly challenging to do these days. The wait times to see an epileptologist specifically are atrocious across the country. So that is all incredibly exciting to hear.

Brandon: Hi, this is Brandon from CURE Epilepsy. Do you have questions about seizures, medications, treatments, or other areas of epilepsy? CURE Epilepsy's new video series *Epilepsy Explained* provides answers to help you better understand the basics of epilepsy. Each month a different expert offers short, easily understandable answers to questions from our community about a particular area of epilepsy. Doctors and researchers who are leaders in their field will cover questions about seizures, diagnosing epilepsy, medications, surgery, and many more topics. New episodes of *Epilepsy Explained* will be available on CURE Epilepsy's website and YouTube Channel on the third Wednesday of every month. Now back to Seizing Life.

Now back to Seizing Life.

Kelly Cervantes: We're talking a lot about AI in terms of analyzing data after a seizure. Is there a world in which we could use AI to prevent seizures, or to know or to predict seizures?

Dr. Daniel Goldenholz: So I think the answer is I hope so. My group and others are working furiously on this very question, seizure forecasting, and I've published some stuff on this, and some of my colleagues have published some amazing things as well. The consensus is that we probably can make some kind of forecast about the future that is better than nothing. In other words, we might be able to provide some kind of hints about the future. Are those hints good enough in order for patients and for doctors to make decisions about them?

That's the harder question. So we can mathematically do something, and it looks cool on paper, but will it help a patient take less medicine, or drive sometimes, or go do that activity that they wanted to do, or avoid certain days because there are certain days that are more dangerous? I hope so, and we're working furiously on that. There's a lot of barriers that we need to jump over, but I think that there is many, many signs that we're onto something because we are able to see patterns in this information about what the brain is doing differently right before a seizure, or even a long time before a seizure. So I hope so, but I don't know.

Kelly Cervantes: I mean, the ramifications for that would just be absolutely remarkable. I mean, you could see SUDEP numbers falling and all sorts of really remarkable life-saving ways if we had better guides on when that next seizure was coming. So now, you have me thinking about how AI could be used in epilepsy surgery, which ... I mean, epilepsy surgery, I mean, even 10 years ago was, it has come

just so far and how precise in the way that we are able to use lasers, and so that on its own is exciting. How could AI be used to help, I guess, maybe predict who would be a better candidate for surgery, or what the outcomes of surgery would be? What does that look like?

Dr. Daniel Goldenholz: Yeah. So we're seeing some really exciting work from our friends at Cincinnati Children's, who are actually trying to read the clinical notes from doctors, and based on those notes, refer patients for early evaluation for surgery. So that's coming. We're going to see, "Hey, this patient might be a good person to consider for evaluation," but then, once you get that, "Hey, let's check into it," then we're going to also see many, many, many studies, which have been saying, "We can take the clinical information that comes in the MRI, the EEG, the PET scan, et cetera, and come up with some kind of a prediction of, "Are you a good surgical candidate? Are you a bad surgical candidate?"

"If you're a good surgical candidate, what would be your chance of success? What would be your chance of bad side effects? And not only bad side effects, but what kind of a cognitive outcome can we expect? What kind of emotional outcome can we expect?," these kinds of things, but then, even more, "Where would be the best place to go?" For example, like you said, "What kind of surgery should we do?," because we now have many different types of surgeries, whereas it used to be you did this one removal, and that's it, and there's no other options.

Now, we have a bunch of different tools at our disposal, "Which tool is best?" I haven't seen a lot yet on that, but I think that we're going to see a lot more in the area of using AI to decide, "What's the best approach for a patient?," and then when the surgery is finished, "Did it work?" Right now, we have to wait and see if someone has a seizure, but could an AI, for example, look at an EEG or other behavioral measures or some other signals and say, "Yep, this person had a successful surgery. They can reduce their medicine. They can go about a normal life," or, "No, this surgery was not successful. We need further evaluation."?

So there's like on all of these different sides, and finally, on the pathology side, when the pathology from a removed part of the brain comes out, a pathologist needs to look at that, and sometimes they make mistakes. Could an AI help them and say, "This is the actual diagnosis," "These are the genes identified," "This is what the cells were doing," "This gives you some interpretation about, 'What can we expect in five years, in 10 years, in other parts of the brain and so on?'"? So the answer is yes, yes, and oh boy, yes.

Kelly Cervantes: I want to keep going down that pathology road that you're leading us down right now, specifically in terms of genetics. Just listening to you talk about all of these things, I'm hearing not just that we could potentially have more information or more accurate information, but that we are speeding up the process of receiving this information. And as someone who held her daughter through countless seizures, the shorter time that we can have to find out what's

going to work and what's not going to work, the fewer seizures, the less damage there is to the brain, the less risk there is to suit up, I mean, just that on its own, shortening the time of when you can get in front of an epileptologist, of when you can potentially find out if you're a surgical candidate, if you can know what your EEG read means, all of that is just going to increase the quality of life for the patient, but perhaps nothing more so than pathology, because there are so many different ways that someone can develop epilepsy, and two-thirds of people with epilepsy don't know. So I'm wondering, specifically in terms of genetics, how could we see AI help there in a place where the rare epilepsies, in particular, are just clawing for research and for information?

Dr. Daniel Goldenholz: I think the right moves are happening in that direction. We have more than 20,000 genes in the human genome, so there's no way that I'm going to know them all or understand them all, but we're seeing incredible advances in AI, now being able to take a gene sequence and saying, "That's going to turn into this shape," and there's also advances in understanding, "If I have this molecule that's this shape, will it fit with that molecule with that shape, and how can I modify this molecule to have this particular action?," and so on. So I think the day is going to come when we can scan someone's genes and say, "Okay, we think the reason you have seizures is because of this, this, this, this, and this. I don't even know what they are, but the AI said so. The AI is right 99.7% of the time, so let's go with that, and the AI recommends this medicine which was custom-designed for you and your particular situation, and it's also been checked to see if your other body physiology is going to have side effects, and this would be a good candidate treatment for you."

And I think that that kind of stuff is absolutely on the way. I think, exactly like you said, we need to speed up the amount of time from suffering to ending suffering so that we can get to that finish line faster, and these tools are exactly designed for that purpose.

Kelly Cervantes: So, I'm sure anyone who is listening, their brain is spinning right now in all of the different ways that AI could be applied in neurology, and epilepsy, and research, and it's exciting, but it's not coming tomorrow. So can you give us an idea of when we can expect to see some of these changes? Do you want to be your own fortune teller, and give us like a year or something that we can look forward to?

Dr. Daniel Goldenholz: Well, like I said, I think that we're in level two now. We are actually integrating small bits of AI into our clinical practice. I mentioned the 2HELPS2B clinical decision tool. There's another one called EpiPick to helping to choose a medication using AI. There is the Empatica seizure watch that uses AI to detect seizures.

There's a very large number of things that are coming online that are becoming approved very soon. So I think that in this year and next year, we're going to be using a lot more tools that have AI built into them, for starters, but if you want to ask, "When are we going to get the full potential of AI flowering and

blossoming, and making our life perfect?," I think that there's always a moving target, because until people are 100% seizure-free all the time, we're always going to want more. We're always going to want no side effects, no problem, no disease, nothing, and until we have a full cure, we're not going to be satisfied. So I think that when we reach stage four on my hierarchy over there, where the AI is doing a lot of the basic functions that clinicians are doing, and the clinicians are really spending much more humanity, face-to-face time, thinking about uncertainty and so forth, which I think is probably 20, 40-ish, but who knows? When we're there, unless everybody's cured, we're not going to be satisfied, and we're going to ask and demand our AI and our tools to do even better than that.

So it's not like when we reach stage four, we're going to say, "Ah, we're here. Everything's done." I don't think that we want to go to doctors when we're sick. I think we want to just not be sick, and that's the goal, right? We want to have people that keep us healthy so we never have to go to the doctor. I think that that's the eventual direction that we need to go, is health maintenance as opposed to disease management.

Kelly Cervantes: Yeah. Yes to all of that. I think that if we didn't have to manage doctor's appointments and specialists and all of that, there are many of us who would be rejoicing and celebrating. I want to turn a little bit to research, and specifically the work that you are doing in your lab, the Goldenholz Epilepsy and Data Science Lab. Can you talk about what you're studying and what you're working toward?

Dr. Daniel Goldenholz: Yeah, so the Epilepsy + Data Science Lab is basically trying to take all of the different tools that we have available in data science. That's mathematics, and that's computer science, and that's engineering, and it's statistics, and machine learning, and data visualization, basically any ingredients that we can and bring them to bear in the field of epilepsy because there's so many amazing things that are happening outside of epilepsy, and we just need to bring them in and use them for our patients. So the types of things that we worry about are, like we already discussed, seizure forecasting. We talk about ways to improve clinical trials, to speed up treatments in order to get them faster to patients, and we're also working on using machine learning for accelerating the interpretation and understanding of EEG, as well as using large language models like ChatGPT and others, to help us deepen our understanding and what's going on with the patient so that we can help the patient faster. So in a lot of different domains, really, but all of them are connected in the sense that we're using data tools in order to help patients in epilepsy.

Kelly Cervantes: That is incredible, and I'm so grateful to you and your team for pushing this forward because this research is so desperately needed and it is going to help so many, so many lives. We've mentioned ChatGPT a couple times. I think that's the AI that people are most familiar with. Is there a usage for AI that is accessible to the layperson, that they could be using to help themselves or to help them in meetings with their doctors?

Dr. Daniel Goldenholz: Definitely. So, we actually brought to the American Epilepsy Society, in this past December, a study where we asked three of the popular free tools for AI. One of them is Bing, which is now called Microsoft Copilot, and that's run by GPT-4, the thing that runs ChatGPT. So Bing is one, and the second one is Claude.ai, and the third one is Bard, which is google.bard.com or something like that. And those three tools, we asked them, "Hey, can you take this test, which is a practice test from the American Epilepsy Society, to prepare for the epilepsy exam?," and all three of them were able to achieve really great scores.

In other words, all three of them are showing near, I say near, expert-level performance on these graded tests, practice tests for understanding things about epilepsy. What does that tell us? Well, that means that you can ask a hard question to any of these free services today, and you can get a potentially very good, meaningful, thought-out answer. The problem is that it also makes up stuff that's complete gibberish. So what you get on the one side is someone who's really, really smart and capable, and on the other side, they're a complete fool.

So if you ask, for example, a medical student, "Hey, what should I do about my terrible epilepsy? My doctor gave me this drug and that drug, and I have these side effects," they might sound really good, but they're a medical student. They don't know practical neurology and epileptology quite yet, and that's sort of where these tools are at right now. You can go and get very good answers from them, sometimes better than your doctor, and sometimes totally ill-thought-out, poorly conceived ideas, but often, you can get good ones. And the reason I mentioned this is because you can, today, as a patient, or as a clinician, ask hard questions, get answers, and then use those answers as a conversation starter.

"Hey, doc, I was talking with Claude.ai, and it said that maybe I should consider Lamictal. What do you think about that?" And it might be a terrible idea for that patient, but it might be actually quite good, and it's worth asking.

Kelly Cervantes: That's incredible. It also makes me a little bit nervous that maybe we should be making the epilepsy exam a little bit harder, but-

Dr. Daniel Goldenholz: Maybe, or maybe, like I said before, that just memorizing stuff is what doctors are not all about.

Kelly Cervantes: Yeah. Okay.

Dr. Daniel Goldenholz: Yeah.

Kelly Cervantes: That makes sense. Okay, so I like this. We can use the AI to ask these questions, but we have to take the answers with a grain of salt and run them by a professional. They are not the end-all, be-all, and I think that's really, really important for people to remember if they are trying to use AI as an educational tool or as a clinical tool, at least where we stand today. I wonder, what are some

of the other ethical considerations that we need to be aware of as AI is being incorporated into epilepsy and the medical field in general?

Dr. Daniel Goldenholz: So I had the privilege to co-author a paper with Sharon Chiang. It's a fantastic paper, 2021 Neurology. I recommend people interested to take a look at that paper, and I will not do it justice because it was written better than I would say. I'm going to mention that, and then just move on to a couple of things that I think are interesting. One is this concept of bias.

So an AI is as smart as it's taught to be, not smarter. So if I bring an AI a bunch of pictures of white, old men and I ask it, "Show me a face," then it can become very excellent in finding the face of a white, old man, but if you bring, let's say, an African American female, and put her into the camera, the AI will make mistakes, because the AI was taught in a biased way. And the same would be true for, "Choose who's going to be a good surgical candidate, people with insurance versus people without insurance." Of course, the AI is going to focus on the insurance people because its training data is heavily biased in that direction. So we need to be very, very careful about what goes into these tools so that what comes out of them makes sense for our patients.

So that's one very important thing, is bias. A second ethical issue that I worry about is gaming the system. If my job is to make money, and I don't care about patients and I don't care about doing the right thing, I just want to make money, I can learn tricks to basically fool an AI to make a more expensive diagnosis. And there've been some interesting publications on this, and it's quite a frightening problem because we could potentially get AI making diagnoses that just cost more money and don't help anyone. So we just need to be aware of these problems and put protections around them. And then finally, a lot of people in ethical AI talk about the concept of explainability, or in our paper, we talked about transparency of method.

I think it's very important to understand what your AI is supposed to do and what it's not supposed to do, but somehow, the idea that an AI needs to be fully understandable has become a very important rallying cry in AI, and I don't understand it whatsoever, because I have a lot of brilliant colleagues, and I do not understand why they make the decisions, what they make. And I see a lot of clinicians that are like this, and when I ask crowded audiences, "Do you understand why doctor so-and-so is choosing what they're doing?," the answer is, "They're very good at what they do," and that's not an answer. And I think that that's the same level of scrutiny that we ask our physicians is the level we should be asking for our awesome AIs. They should do the job, they should do it safely and well, but if we don't understand how they do it, it's more important that we know the box around them, that they're going to do these things and they're never going to do those things, those things are dangerous, these things are helping and healthy, those are okay. So I think that there's a lot that goes into the ethics of AI, but these are few that I think about.

Kelly Cervantes: It's one thing for you to use AI in a research lab. It is quite another to use AI when treating patients. What are the barriers of moving from the lab to the clinic?

Dr. Daniel Goldenholz: Yeah, wow. When I take an AI tool and I say, "Look, this is mathematically good. It's accurate according to some method or whatever," and I want to give it to patients, I can't just do that. I need to make sure that there's no safety concerns. I wonder, "What happens if someone misuses this, and uses it in the wrong kind of patient?"

There's going to be regulatory concerns. There's going to be deployment concerns. "How much is it going to cost the hospital? Is the hospital responsible if the AI makes mistakes? Is the company that gave the AI responsible ... Who's responsible?"

Then, there's also this question of maintenance, because AI is really software. So whenever you make a piece of software, if you make it once and never maintain it, then after a while, it won't work as well because there's new operating systems, because there's new sensors, because there's new things, and unless you're constantly keeping it maintained, it's not going to work as well as it originally was intended. So all of these things are kind of barriers between, "I made something cool and flashy, and the patient wants to use it." We need to make sure that it works, that it's safe, that we understand it, but also, that we're maintaining it, that governments are allowing us to do it. And by the way, right now, our governments are not very quick to understand and to adopt AI technologies because they don't really know what we're doing, and they're not able to move at the speed that AI research is moving.

Kelly Cervantes: That's a whole lot.

Dr. Daniel Goldenholz: Yeah. It is.

Kelly Cervantes: So it's overwhelming to see how fast the technology is moving and also acknowledge the speed at which these things take to be approved, which is sluggish and slow, but it is important. You talked earlier, you were talking about certain technologies that are FDA approved, and I do think ... We need that approval. Do you agree that we need to make sure that all of these systems are approved in that way, or is there a way that we could speed some of this up safely?

Dr. Daniel Goldenholz: Yeah. I mean, I think that fundamentally, the FDA model is inappropriate for what we do with AI. I think that the FDA model is, "I found this compound, this chemical, and it's always going to be this chemical from now until eternity, and this is all that I'm talking about, and I want you to make sure that this will be safe in humans, and humans have been the same for 10,000 years. Please go check onto this." That system takes 10 years and several billion dollars, and that's terrible, but it makes sense.

When it comes to AI, when you have a piece of software that can be developed in three months and can be updated in one week, and the circumstances which it's using changes from week to week, month to month, year to year, it no longer makes sense to have a 10-year process where you're vetting and proving that everything's perfect, because 10 years later, everything is different. I mean, if you think about the world, before the introduction of ChatGPT to today, many people haven't even heard the term AI until ChatGPT came out. We're moving way faster than a system that's designed to move at this very, very slow and methodical speed. We need a much faster mechanism. So in my opinion, I think that we need to build a completely new regulatory system that is much faster and much more nimble.

Kelly Cervantes: Yeah. That makes perfect sense to me. As we're talking about all these systems and integrating them, and the regulatory, and I'm thinking of accreditations, and how do ... I mean, is this something that neurologists and epileptologists in medical school, is this something that is going to need to be taught, and how do clinicians stay on top of something that is so frequently changing?

Dr. Daniel Goldenholz: So medicine has always been changing, and it's been changing over the past few decades. The things that I learned in medical school, many of them have been already refuted and have been replaced with other things, so I don't think it's new that things are changing in medicine. What is new is that now, the technology, that really is used to be the domain of the computer scientists and the theoreticians, is now really becoming practical. So yeah, medical school has to upgrade itself. Here at Harvard, we're doing that and we're changing the curriculum in multiple ways.

With my students, I'm showing them these tools, these AI tools, because like I said, you can get great medical answers from them, which you can't trust yet, but we're getting there, and even if you can't trust it, you can say, "Ooh, that's a really cool thing I didn't know about. I'm going to go look that up using trusted sources now." So I think that the answer is absolutely, we need to upgrade our approach to AI. We need to understand some of the tools of AI and how to interact with them better, because if we don't, then the people that do are going to be doing our job.

Kelly Cervantes: Yeah. Dr. Goldenholz, this conversation has been absolutely fascinating. I am so grateful to you and your lab for pushing this conversation forward, and for you to speaking with us about this today, and especially in such a lay-friendly way that I hope everyone can understand. Thank you.

Dr. Daniel Goldenholz: Thank you very much.

Kelly Cervantes: Thank you, Dr. Goldenholz, for helping us understand the current and future impacts of artificial intelligence on epilepsy researchers, clinicians, and patients. CURE Epilepsy has been connecting those three communities for over 25 years, with the aim of accelerating science to discover new patient-focused therapies

and ultimately cures for epilepsy. If you would like to help us achieve our goal of a world without epilepsy, please visit cureepilepsy.org/donate. Thank you.

Legal Disclaimer:

The opinions expressed in this podcast do not necessarily reflect the views of CURE Epilepsy. The information contained herein is provided for general information only and does not offer medical advice or recommendations. Individuals should not rely on this information as a substitute for consultations with qualified healthcare professionals, who are familiar with individual medical conditions and needs. CURE Epilepsy strongly recommends that care and treatment decisions related to epilepsy and any other medical conditions be made in consultation with a patient's physician or other qualified healthcare professionals who are familiar with the individual's specific health situation.