

### Transcript Details

This is a transcript of an educational program. Details about the program and additional media formats for the program are accessible by visiting: <https://reachmd.com/programs/the-convergence/clinical-monitoring-inpatient-decline-ai-alerts/39924/>

### ReachMD

www.reachmd.com  
info@reachmd.com  
(866) 423-7849

---

### Next-Gen Clinical Monitoring: Predicting Inpatient Decline with AI-Driven Alerts

#### Dr. McDonough:

Welcome to *The Convergence* on ReachMD, where innovators and physicians explore the technology transforming medicine. I'm Dr. Brian McDonough, and today, we'll be examining a remarkable use of AI: a continuous clinical alert model that's able to predict adverse clinical outcomes up to 17 hours in advance.

And joining me to talk about this new tool is Dr. Theodoros Zanos. He's an Associate Professor at the Feinstein Institutes for Medical Research and the Zucker School of Medicine, and he leads the Division of Health Artificial Intelligence.

Dr. Zanos, it's great to have you here today.

#### Dr. Zanos:

Thank you, Dr. McDonough. Great to be here.

#### Dr. McDonough:

So let me start with just the first question. What does artificial intelligence actually see that doctors don't? I know that's open-ended, but there just seems to be so many options out there.

#### Dr. Zanos:

That's a great question. The AI actually sees the same things that doctors look at, but it can process them in a lot quicker ways and also at a scale that cannot be achieved by a human, no matter how capable they are. Humans are pattern-finding machines, and that's exactly what AI is also trying to do, but it does it in a scale and at a speed that cannot be matched by humans.

So all of a sudden, it can find extremely subtle patterns that might be related to a specific disease. So it becomes very good at diagnosing a disease, or it finds very subtle patterns that precede a specific outcome or event, and that enables that AI to become very good at predicting that event.

So it's not that it performs a completely different process of thinking, let's just say, than doctors in identifying, diagnosing or predicting something that will happen. But it does it at a scale and at a speed that is just not possible for a human—let alone an overworked, tired doctor.

It seems that lately, artificial intelligence has been mainly used for us to discuss chatbots and the latest instantiation of artificial intelligence. But it has been around for decades and way before chatbots, and it has also had different names, from systems identification to machine learning and deep learning models and all that. All that is part of the general field of artificial intelligence.

But all these different terms or subfields have had applications in medicine.

We see now more and more as the whole society's looking more and more to AI. Especially because of the chat bots, we see the medical establishment also focusing a lot more on this technology and seeing how well it works, but also where they can use it. So from our standpoint, we don't discriminate on what type of AI we use for a problem as long as it actually solves an actual real clinical problem.

#### Dr. McDonough:

If someone were to do a search for your name, there's all sorts of articles. You were kind enough to send me some, and of course, one of the articles, which was in *Nature*, talked about the use of AI and of course, a fascinating use: the 17-hour prediction.

Let's break that down. What is the 17-hour prediction? What does it mean, and where do you see its use clinically?

**Dr. Zanos:**

The main problem that we found—it's a well-known problem—is that patients that are admitted on medical surgical wards that are supposed to be relatively stable patients—some of them actually deteriorate while being inside the hospital and require elevation of care. And that percent can be anywhere from five to 10 to 15 depending on the site and the type of unit.

So it's not a trivial problem, and usually these deteriorations mean that the patient might require a longer hospital stay transfer to the ICU—unplanned—or even a death on the floor. So there is an absolute need to identify these high-risk patients and also give a heads up to the care team that somebody is at this high level of risk so they can do some preemptive treatments and take some preemptive actions.

17 hours was the average length of the window that our algorithm that we developed to try to predict these deteriorations came to be. So we trained it to flag deteriorations as early as it can, and it ends up being, on average, 17 hours before the actual event is happening.

**Dr. McDonough:**

And what are some of the things it's looking for? When you had designed the algorithm, what are some of the things that allow it to make this determination?

**Dr. Zanos:**

Yeah, so the main data that the algorithm is using is vitals. But it's vitals that are recorded continuously. So the whole basis of the algorithm is that it can be applied on clinical wearables—so devices much like your Apple Watch or the Fitbits of the world—but for clinical use. So there are specific clinical wearable devices that are cleared by the FDA that have a bit more accurate readings on these very specific important vitals such as heart rate, respiratory rate, and temperature. They can also record movement.

So these are the signals that we take from these clinical wearables, and we use them to develop this algorithm that essentially is finding early telltale signs of downstream deteriorations. And these are just subtle changes. Some of these vitals or combined changes in some of these vitals reliably predict these deteriorations down the line.

**Dr. McDonough:**

So let's say in an intensive care unit, you might have nurses or physicians or others watching monitors throughout the day and checking things. This is actually not just occasionally glancing over and making sure the rhythms trip hasn't changed. This is looking at every element and continuously looking at them?

**Dr. Zanos:**

Correct. And the other big difference is that this is in medical surgical worlds, where we don't have continuous monitors on these patients usually. These are more used in ICU settings. But in regular floors, we check these vitals every four to six hours. And the problem is that during these four to six hours, some things might change, and then all of a sudden we identify a patient deteriorating and we are asked to try to avert it. But maybe by then, it might be too late.

So my favorite quote that kind of inspired this whole line of research is that patients don't suddenly deteriorate—we suddenly notice. That's mainly because we have to check on them every four to six hours because we can't have thousands of doctors or nurses constantly checking on all these patients in these wards. But now we have a technology that could be monitoring every five seconds and take away a new measurement of the heart rate, respiratory rate, temperature, and heart rate variability. And using algorithms like the one we built, it can actually ingest these and flag a patient early in advance and identify these patterns of deterioration.

**Dr. McDonough:**

I'm thinking about my world of practice. I teach in a residency program, but we do inpatient and outpatient. So we might admit somebody and, on many occasions, I say, "We'd better admit them to the unit, stepdown, or whatever you want to say—telemetry—because I need the monitor." And also, some of the time, I'm not going to get there because I'm back in the office seeing patients and coming back at noon or five or whatever.

What you're saying is, it isn't just that—it's constant. So somebody could go to a regular floor bed, and they'd be in a regular floor bed just with a wearable device. And we could be alerted, and I'm assuming when you talk about the alert, basically that would then alert you through the computer almost like you would get an alert about anything else, right, that there's something to watch for?

**Dr. Zanos:**

Yeah, usually these alerts are surfaced in the nursing station. And you know that this patient at that specific room has an alert. And again, the point of this is not to alert that something is happening right now, but there's a high risk of this happening in some hours from now so you have enough time to intervene.

So it's not really a code red or something that all of a sudden people will need to start running immediately. It's more like a flag that you should be very attentive to this patient and maybe start calling your ICU to arrange an elevation of care. So there's a lot of different ways that this can be actionable. But it can be used as a risk stratification tool as well.

And then we've worked with similar ones and developed some other ones as well where you can kind of stratify and rank patients on a specific medical floor based on their acuity and how likely they are to experience these outcomes. So there are a couple of different ways of using this.

**Dr. McDonough:**

In case you're just listening, I'm Dr. Brian McDonough. I'm speaking with Dr. Theo Zano about a device basically predicting through clinical alerts potential adverse clinical outcomes up to 17 hours in advance.

So in my background, I was involved as a Chief Health Informatics Officer, so we were putting in these computers, and initially when we were doing it, as you know, it was somewhat annoying because it was more or less just telling doctors they had to do more typing. And this is what we were doing it for—it wasn't just for coding and billing and all the other metrics. You could make a difference in someone's care, and it wouldn't be possible if we couldn't get it to the machine.

So let me talk about you for a second. AI actually is, as you said, this term that's grown in popularity, but things have been done for quite some time. Tell me about your journey from neuroscience to bioelectronic medicine and what got you going in this direction, because without people like you, we wouldn't have these tools.

**Dr. Zanos:**

So my journey has been with quite a few turns. My first degree was in computer engineering. And while I was doing my undergrad in electrical and computer engineering, I constantly was talking to my family—my mother was a nurse, and my brother is an MD. I'm almost the black sheep of the family not pursuing medicine. So I had to kind of correct for that and try to figure out something worthwhile with my computer engineering degree.

So I decided to pursue a PhD that would utilize some of my computational background, and I thought the biomedical application would be worthwhile.

But I also wanted to try to answer a very hard question at the same time. And understanding how the brain works at that time—20 something years ago—seemed to be a worthwhile and hard question that maybe I could put a dent on.

So that's what I tried to do during my PhD. I joined the biomedical engineering department at USC, and my PhD was essentially developing what we now call

AI or machine learning models. Back then, we didn't use that term. It was systems identification for the hippocampus, so figuring out the inner circuitry of a very specific brain area. And through that, I started to really appreciate both the complexity of the brain and the power of advanced computational tools.

So I then moved on to do some more neurophysiology at the Montreal Neurological Institute. But when I came to New York, I always wanted to go back to my computational roots and utilize this expertise that I had in both machine learning, but also in how the brain organizes itself and you have the building blocks of the brain—the neurons—and how they communicate and how you have this incredible array of capabilities that the human brain possesses by building complexity from very rudimentary, let's just say, individual units that are the neurons.

But when these individual units connect in complicated ways and in the millions, we have the emergence of these really exciting properties. We have the human brain with all its glory, with consciousness and philosophy and all of these different amazing things that humans can do. So going back to that enabled me to work both in neuroscience, but also in trying to advance some more computational methods, kind of like combining them at the same time.

And when I joined Northwell in 2016, I also realized that there is a huge opportunity in utilizing the fact that I'm part of the largest health system in the Northeast. And as you said, because of work like you and other giants in the clinical informatics space, we had access to decades of data from millions of patients that we've served. So we could look back and start analyzing these with these advanced computational methods and try to develop tools that would solve actual clinical problems. That's why we built the Division of Health AI, and that's a lot of the work that we're doing right now.

**Dr. McDonough:**

It's so interesting talking to you, Theo, because in my career, certain things have crossed over. I used to be Chair of Pharmacy and Therapeutics for Trinity Health, and they're like 90 hospitals, and I learned how we could do studies based on data that we had at places

like Loyola in Chicago and Syracuse and Connecticut.

You had all these built-in patients' cases and things you could use with that existing information in a clinical way, but you're taking it to such a new and exciting level. What I was doing is with others—obviously pretty smart people—but conducting studies the way we have traditionally done them. You are more or less looking at, for lack of a better term, bioelectronic medicine. What is bioelectronic medicine in plain English?

**Dr. Zanos:**

So bioelectronic medicine is a new field of medicine that actually started here at Northwell and at the Feinstein Institute for Medical Research about 20 to 30 years ago. And it's essentially developing the science and the devices and tools to interface with our nervous system, trying to treat diseases that are not necessarily only linked to our brain and our nerves.

So, for instance, one of the breakthroughs that really drove a lot of research and now innovation in this field is the identification of the inflammatory reflex, which is a neural reflex. That governs inflammation in our body, and it controls how our body responds to injury or trauma and regulates inflammation.

And that has been work that has been done by multiple different labs, but was mainly spearheaded by Dr. Kevin Tracy, who is also the CEO of the Feinstein Institutes. And he described the inflammatory reflex about 25 years ago and identified the main conduit of it, which is the vagus nerve. So when the vagus nerve is stimulated, what they found is that it can actually put the brakes on the inflammatory response. And that helps a lot with autoimmune conditions or things like rheumatoid arthritis or Crohn's disease.

And now this year, there has been a company that got FDA approval for vagus nerve stimulation for rheumatoid arthritis treatment. So the whole field of bioelectronic medicine is not, of course, only the vagus nerve or only rheumatoid arthritis, but in general is harnessing through electricity and devices the ability of the body through its innate circuits to heal itself essentially.

**Dr. McDonough:**

So you're talking about, in many ways, two different things. Bioelectronics, for those of us in practice—if there's anything we've learned in the last 10 to 20 years, it's the role of inflammation. You can have poor dentition and that could lead to heart disease. And COVID, some have described as a disease of inflammation—that it just started attacking people's weak points because it was inflammation.

So all these things are tied in when you look at that and the work you're doing when you take that information—you're also using information from others' work and patients who have been studied and compiling it yourself as well.

**Dr. Zanos:**

So some of the work in electronic medicine that we did was, instead of just stimulating the nerve, we actually tried to record from it because nerves cut information either from the brain to the body or the other way. And what we found is that there is a very specific language that the nerve uses to encode inflammation, and there's also a specific language that he uses to encode metabolic signaling as well.

So now, where this is going is essentially devices that, by eavesdropping on this signaling, could provide the triggers for proper stimulation or optimization of the stimulation. But in parallel to that, the way that we decoded this language is by using very similar models like the ones that we used to identify the language that precedes deteriorations. And instead of looking at vagus nerve recordings, we used it to look at heart rate and breathing rate and temperature recordings in the same way.

**Dr. McDonough:**

Those of us who are practicing day-to-day medicine, we have this thing called alert fatigue where there's just so many alerts coming to us throughout the day. You almost have to get rid of some of them because people start to ignore them. With this 17-hour prediction, how do you make sure that doesn't get wound up with all their other alert fatigue? Because this could be crucial and somebody might say, "They just told me the potassium's a little lower," and they may not see it as that same urgency.

**Dr. Zanos:**

That's a good question, and it's actually a very important aspect and a barrier to deploying these tools and getting them to be used in the clinic. You've got to be extremely careful and listen to the main users, which are the nurses and the doctors, and make sure that you don't overburden them with more alerts, new buttons to click, or new popups to close.

So the way that we're trying to account for that is first by providing a high positive predictive value, or PPV for short. And what that value tells you is how many of the alerts will actually be true alerts on average based on my prior work and my prior test data

So usually, common alerting systems in the hospital are between five to 10 percent in this area. So out of the hundred alerts you're going to get, 10 are actually real, and 90 are false positives. But we are on the side of caution, so we're going to give you all 100 so that

we don't miss those 10. So we doubled that as our threshold of what would be useful.

But the reason why it's still one out of five, which is not great—it should ideally be that every single one is a real one, right? The problem is that AI, as powerful as it is, it's not magic. And there are things that our algorithm is not measuring necessarily. Or these, these telltale signs or subtle trends in the vitals might not necessarily be predictive of this. So sometimes it will give you a false positive. But we want to make sure that we have high accuracy, so we are really capturing the ones that do deteriorate. So, that's why our accuracy needs to be really high. We try to educate the clinician and also provide the hyper-positive performance value.

**Dr. McDonough:**

So, once again, knowing how it's worked, even from the days of telling you about CPOE for the first time with doctors, it's that communication, which is key. You're setting expectations and presenting it in real time. Now, has this been your experience at Northwell where you're getting this percentage? I would assume this is where you're using it clinically and getting to see it in real time.

**Dr. Zanos:**

So, we just developed it. We haven't deployed this specific model yet. That's one of our next steps. But absolutely, in any model that we've deployed, in any collaboration that we've had with our clinicians, our very first thing to do would be to involve them as early as possible in the development of these tools.

And that's the big value of not being just an academic lab but being a division that's embedded inside that health system. We get our ideas on what to build next based on our ongoing conversation with the frontline workers—the doctors and the nurses—and hearing their problems firsthand—what they're facing every day, and whether AI would necessarily be the right solution for some of these problems.

So, when we decide that, yes, this is a worthwhile project to work on, we always try to bring them in from the very first ideation phase. And of course, as part of our division, we have many MDs that have a clinical background as well as a computational background that can bring this clinical insight. That really helps a lot.

**Dr. McDonough:**

I'm Brian McDonough, and I'm speaking with Theo Zanos. We're talking about his work that he's doing using AI.

You know, when we started off our interview, we were chatting and you were talking about how family members and people were in healthcare. Clearly, a lot of that rubbed off on you as well because you're thinking in this environment like a clinician, because you're right. If we think it's going to help a patient, we'll do anything. But if we think it's tedious or just doing some other work for a machine, we tend to recoil and say, "We're not going to do it. I've got to move on." But if it's patients, they will.

So getting them involved is really important. And I can see that's where you're going to have your impact in clinical practice, because they're probably also going to feed back the good and the bad as when you start to test it out there. They'll say, "Oh, this has been a problem or not." You'll probably learn things from them as well, I would imagine.

**Dr. Zanos:**

Yeah, absolutely. And again, there is a whole process of how you deploy these types of tools in hospitals. There is a whole period where you're just making sure 100 percent that how you think the model works actually works in real life. And you do that by doing what we call a silent deployment, meaning that you turn on the tool for a certain amount of time—one, two, three months— but you don't surface the results to anyone. You're just running it in the background and making sure that it's running correctly. And you also check its results and you check, well, does it have this accuracy, and does it have this positive predictive value still? Or has something changed? Has something in my data changed?

So as soon as you do that, of course, parallel to that, you always talk to the stakeholders and you ask them, "How do you want this to be presented to you? Where do you want this to be presented to you? What part of your workflow? At what point? How do you want to acknowledge that you saw it? Or what do you want it to point to as a specific action?"

So all of these things are extremely important in building a tool and deploying it that people are going to actually use. No matter how advanced it is or accurate it is, at the end of the day, it needs to be usable and fit in the clinical workflow. And by doing that, you're also promoting what we call change management or a bit of training, right? Because they get involved in, what does it actually do? How do you expect it to work? And all of a sudden, you have your clinical champions actually using the tool, and that's a recipe for actually deploying, as far as we know and based on our experience, a tool like this appropriately.

**Dr. McDonough:**

So obviously, you're very realistic, and you're a champion of artificial intelligence. Where do you see AI get it wrong? What are our limitations as we sit here right now looking at it?



**Dr. Zanos:**

AI gets a lot of things wrong. And every single model will get things wrong no matter what. There's no such thing as a perfect model or a perfect AI algorithm. I think the, the bigger question is, first of all, how many times do humans get it wrong for a similar use case? And a lot of the times these algorithms are compared to human benchmarks, and a lot of the times they outperform it.

But I think right now, our focus is on building tools that will be usable by the doctors and the nurses because the actual human computer interface is the biggest limiting factor where I see the opportunity of this technology to really upend the way we're practicing medicine right now.

I think there is a huge opportunity to utilize these extremely capable pattern-finding machines in the clinic as a another data point and as another very useful tool for the doctors so that doctors can actually be a lot better in certain cases. But in order to do that, and that stands even for chatbots, we need to train the doctors to use them appropriately.

And we also need to ourselves to design them appropriately so that they understand what they're using and also use it the right way. And I don't want to open up Pandora's box, but there is this whole idea of human plus AI is always better than either AI or human alone. There have been actually studies lately that have shown that that's not true. Actually, AI beats both human and human plus AI. But I think the limiting factor there is that we haven't trained the humans yet to use AI properly, and we also haven't optimized the interface in the way that we are delivering the AI and the predictions and everything to the doctors.

**Dr. McDonough:**

One of the things which is clear is you're living in the real world. Any of us who have tried to have projects or get things through, whether it's through med exec, through the hospital, or budgeting, those things are always barriers. What access or implementation barriers, financial or otherwise, do you see?

**Dr. Zanos:**

There are a lot of barriers that are right to be there. And there are some that we could do without if we could to optimize the way that we deploy these tools.

So the first one is that you need to gain the trust of your stakeholders, of your users. So you need to gain the trust of the doctors and the nurses that are going to use these tools. So there is one barrier in how to deploy it and how to not only gain but also retain the trust of the user. And I think the main way that you're doing it is by being accurate. Because, as you mentioned, if I flag a patient for deterioration 20 times and they're completely fine and they remain completely fine for the next day, I messed up, right? And if this happens once or twice or three times, then the nurse or the doctor will say, "Well, this is useless" very quickly. And they'll most likely be right about it, right? So the very first barrier that you need to overcome is you need to have a really good and accurate model.

But on top of that, there are other barriers that relate to, how do you deploy this? How do you integrate it into existing electronic health records? Again, most of the work that our doctors or nurses are doing are inside specific EHR platforms. So if you now all of a sudden ask them to minimize this platform, open up this other window, and log in with your account there, they're never going to do it. Right? So you need to have a seamless integration into their clinical workflow. That's important. That is a bit of a barrier because many of these solutions do not necessarily exist or can exist easily inside these platforms. But that's something that we've seen be an important barrier in how you deploy this.

And there are, of course, just in general, barriers of getting the buy-in from leadership as well, which lately has been shifting about. I want to say three, four or five years ago, it was a bit harder to make a case to start using AI tools, but since ChatGPT came out and everybody started talking about it, the stance of a lot of both doctors and leadership changed in how open they are to using this type of tool and to explore possible use cases in medicine. I think that helped in a lot of ways to lower those barriers. But rightfully so, they exist for financial issues and for security reasons. You want to make sure that the tools you're using are not going to harm the patients.

**Dr. McDonough:**

So, I want to be respectful of your time, but I have a couple questions I have to ask you. First of all, you had a vision 20 years ago and you've been following along and working and growing. You're poised to be able to probably answer this question. Where do you see this technology going in, let's say, five years?

**Dr. Zanos:**

Five years is a century in AI times. So I really don't know where exactly it's going to go. I can only guess. But I see this technology getting more and more mature and providing more and more holistic approaches to the way that we're treating or diagnosing patients.

So I'm a huge proponent of the fact that you can only have a great model when you have great data, and those great data cannot be only from one specific small source, but it should be from multiple different sources—as many as you can that monitor different types of

signals that are related to somebody's health. That's why, for instance, for our deterioration model, we're using both some demographics and some electronic health record related data, but also those clinical, wearable, continuous signals as well.

In general, I would say that how AI is moving to the healthcare space—I would say that it's already here. It's just going to grow and change not only the technology itself to become better, but also the doctors. It's going to change the way that they're practicing because it's going to start touching a lot of different things that they do day to day, from an ambient AI that's transcribing and summarizing discussions happening right now to AI that's going to help them write notes and reimbursements to AI that's going to alert them if a certain patient might have a certain disease through opportunistic diagnosis.

You might be going to more your PCP for one specific reason, and they might alert you for something else just because there was an AI model that identified a pattern that is related to something different. So I think it's going to enable us to really move to a lot more preventive care rather than responsive care. And this composite aggregation of data about a patient is really going to enable us to realize this dream. The physician's going to have immense power in scanning across multiple different sources at the same time. While at the same time really focusing on the human connection and what to do next for that specific patient.

**Dr. McDonough:**

It's interesting. I teach residents and young physicians, and they're always coming up with better uses. I'll say, "Oh, well why don't you use OpenEvidence? This can help you." And they go, "Well, it depends what you enter in. Are you entering in certain journals with what your query is? Are you entering this?" And they're like, "You've got to give more information. You're not getting as much as me, Dr. McDonough, because we're doing this." And then we're all learning together.

And I'm not surprising you, but there's an age barrier. If you're under 35, you've grown up with it. You don't read directions. You get something and you just start, open it, and play. That's kind of what they do with this, whereas we're going, "What about this? What about that?" So there's doctors at all ends of the spectrum here.

You sent me a series of articles. They were fascinating. And these are journals I always read, yet I didn't read those articles. Because in my mind, I said, "Oh, that really won't have clinical relevance to me I think it should." But it does because what you're talking about is what's happening. How do you, I guess, help physicians start to change the way they view things? To read these articles to be open to things and to learn? Because they're going to need to do it.

**Dr. Zanos:**

I agree that they will need to do it. I think the advent of AI technologies in healthcare is not going to stop. It's here and it's going to stay and grow, and it's going to change the way that they practice medicine.

So I think as with any rapidly changing field, adaptability and trying to keep up with what's going on is the core competency to remain relevant. So I think future doctors will need to familiarize themselves as what now we call power AI users. I think future doctors will be just that in the same way that 30 years ago, most likely, my grandfather didn't know how to turn on a computer, and I'm pretty sure now someone his age would be online chatting to his chatbot.

So things change. Society changes, and we adapt. But I think doctors need to adapt to this new model of utilizing these tools. That's from the user's perspective, but also on doctors lies the opportunity and the responsibility to vet these models as well. We cannot just blindly adopt AI models and decide based on them. So they need to be able to check their accuracy, to read through the performance characteristics of these models, and be able to actually judge whether they are good enough for their practice. They should advocate, whether they're leaders or frontline physicians, for safe and effective deployment of these tools.

And when they are able, they should teach the other doctors but also their patients and talk to them about this technology and convey that when this technology is trustworthy, it's not that we're going to go out of business and people won't need us. It's that our lives will become a lot better, and our jobs will become a lot easier, and our patients' lives will also become a lot better.

**Dr. McDonough:**

Well, Dr. Zanos, I can't think of any way better to finish the conversation than to talk about those points because you hit so many of the areas where physicians and nurses and PAs and all everybody involved in care need to be a part of it. And I really like the fact that you're stressing that they be a part of the solution as well. As you know, in the early development of a lot of the tools we use, they didn't have clinicians involved early on and they learned over time. The more we get involved working with engineers and others, that collegiality and that teamwork can be good.

But I want to thank you so much for joining me on the program. This has been fascinating, and we only skimmed the surface of what you could share, but it has been wonderful, and I really appreciate your time.

**Dr. Zanos:**

Thank you so much, Brian. It was great to be here.

**Dr. McDonough:**

You've been listening to *The Convergence* on ReachMD, where innovators and physicians explore the technology transforming medicine to hear about other technological advances shaping the future of medicine. Visit *The Convergence* on ReachMD.com, where you can Be Part of the Knowledge. Thanks for listening, and we'll see you next time.