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/ai-powered-alzheimers-care/54784/

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

Earlier Detection, Smarter Prevention: AI-Powered Alzheimer's Care

Dr. McDonough:
Alzheimer's disease is usually recognized clinically after symptoms appear, but the biology may be changing years before a patient or clinician sees memory loss. That creates a difficult question for medicine: if we can detect risk earlier and more personally, how do we decide what to do with that information responsibly?

Welcome to *The Convergence* on ReachMD. I'm Dr. Brian McDonough, and joining me today are Dr. Adrian Noriega de la Colina and Arvid Gollwitzer. Adrian is an MIT Novo Nordisk AI Fellow working at the intersection of physician-scientist training, artificial intelligence, and Alzheimer's prevention. Arvid is a researcher affiliated with the Broad Institute of MIT and Harvard, MIT, and ETH Zurich. He is also a model design leader for FINGERPRINT, a platform using agentic AI and foundation models to reason across lifestyle, clinical, genomic, proteomic, biomarker, and population-level prevention data.

Adrian, Arvid, welcome to *The Convergence.*

Dr. Noriega de la Colina:
Thank you for having us, Brian.

Mr. Gollwitzer:
Thanks so much. Super excited.

Dr. McDonough:
Adrian and Arvid, can you tell me and our audience about the work you're doing, fundamentally?

Dr. Noriega de la Colina:
We're at the crossroads of artificial intelligence, biotechnology, and healthcare, and there's a number of chronic diseases that could benefit considerably at that intersection. We're focusing at the moment specifically on one of the hardest problems, which is Alzheimer's disease, which is a disease that is an enormous burden to so many families across the country and worldwide.

Mr. Gollwitzer:
We're currently in the process of building FINGERPRINT, which is a platform to ingest multimodal and multiomic data. As part of these efforts, we've already pre-trained our first seven billion-parameter foundation model called FINGERS-7B, and we're currently in the process of spinning this effort off the Broad Institute of Harvard and MIT together with our co-founders, Giovanni Traverso, Miia Kivipelto, and Li-Huei Tsai.

Dr. McDonough:
For a practicing physician, the question I have is not simply whether AI can predict risk, but what is a responsible way to use that information if and when we have it? And how do we provide personalized care? So I'll open it up to you, Adrian. When you hear the term “biological fingerprint,” it sounds fascinating, but what is that idea? How does that impact us biologically? And more importantly, how does that help us deal with the questions that we're raising here?

Dr. Noriega de la Colina:
Absolutely. We live in a very exciting time for Alzheimer's disease. For the first time, we have blood-based biomarkers. We have treatments available, and what we know is that those treatments, as effective as they are in the real world, are coming at a very late time in the disease pathology.

What we know from the biology is that there's, as you mentioned, a biological fingerprint. We can identify those biomarkers early, but

they don't explain the whole picture. But now we have tools; we have genomic data, proteomic data, and transcriptomic data. We have this abundance of data that allows us to identify those individuals 15 to 20 years earlier and know how the trajectory more or less is going to evolve over time.

So for a practicing physician, we can already not only identify the risk factors, but also, those risk factors can tell us where that individual more likely is going to go based not only on their family histories—we used to do that, traditionally—now, all of that genetic and proteomic material tells us a story of how that evolution is going to go and starts to provide actionable insight. We don't have treatments yet for those early stages, but we do have a modification of risk factors that it has enormous impact on how that individual is going to evolve as the pathology evolves.

Dr. McDonough:
Arvid, from the model design side, what does the AI see that a clinician or a researcher might miss when looking at these variables?

Mr. Gollwitzer:
So where the foundation model and what we're doing becomes really powerful is when there's a lot of different data involved that would be just really hard and way too slow and resource-intensive for a clinician or humans to analyze. So what we're doing is pulling all these different datasets—multiomic datasets that we can get from blood samples, stool samples, et cetera—together and looking at them with a pre-trained model, which is a pre-trained foundation model that's seen a lot of the data that exists from around the world and then looks at these massive amounts of complex multiomic data with this preexisting knowledge.

And that's how we can uncover hidden structures, patterns, and correlations between these complicated datasets and, for example, our understanding of early risk and an understanding of how we'd respond to different interventions. And what we found is that if we use some of these foundation models and AI to help with this analysis, we can identify multiomic precision signatures that can be quite complex at times and aren't intuitive if we just look at those directly.

Dr. McDonough:
So we've discussed some key issues here, but one of the things I see as a family doc—I look at dementia, Alzheimer's in particular, and we do have some disease-modifying therapies, but nothing, frankly, that makes a great big difference. On an individual level, I saw my father, in his later years, develop dementia, which we believe was Alzheimer's, and we watched the decline, and there wasn't much we could see. So I think of it personally or as a clinician. If I could find out early at this point, do I want to find out early? Do I want my patients to find out early? What's the bonus of knowing that, and where is this helpful at this stage where we are right now? And maybe Adrian can talk about now and going into the future, getting this diagnosed quickly—what do we think?

Dr. Noriega de la Colina:
Absolutely. So there's a couple things. First thing many families haven't gone through what you have with your father, and that means that you sometimes arrive at a diagnosis of dementia, but you're not exactly sure—was it Alzheimer's? Was it Parkinson's disease dementia? What kind of dementia was it? I had the same experience with my grandfather. We knew that he had some sort of mixed dementia. There was never a proper diagnosis. That tells us that there is a biological heterogeneity, right? A variability on what the trajectories mean, and not everything exactly goes and gets a label of Alzheimer's disease purely. That's usually a very small subset of the total population that develops dementia. The most vast numbers are individuals who have this mixed dementia, which have a component of vascular, and a component that might be based on, for example, amyloid and tau pathology, which is Alzheimer's disease pathology.

So all of this ends up being very complex, and because of that, treatments are not as effective. You were mentioning disease-modifying therapies; there's two medications in the market in the US, but those likely are going to be effective in one small piece of that puzzle, but they're not going to not be suited for all of that variability that comes in the dementia and the dementia family, let's say.

With identification of these biological fingerprints, we're able to dissect who goes into which bucket, and that translates into treatment because as a family physician, you're going to see a number of patients with all sort of chronic diseases overlapping. You might have a metabolic syndrome. You might have diabetes on top of a history of hypertension, and all of that gets mixed. No patient is a textbook definition of the case, or seldom they are. And that translates into, what can you do actionably?

The second thing is there's things that you can do before you get actually a formal diagnosis. The same with diabetes—you can intervene with lifestyle changes before you actually develop diabetes as a diagnosis. You're going to have a period of insulin resistance where there's potentially actionable items that you can do in lifestyle modification. We know that that is effective, and that explains over 40 percent of the risk, and there's things that you can do. If you manage that blood pressure and control it, and if you become a bit more active, all those things are going to reduce the risk of developing that mixed portion. So even if you have, for example, a genotype—there's an Alzheimer's disease APOE4, for example, which it carries a higher genetic risk for Alzheimer's disease pathology—still,

you're going to be able to modify the other portion. And that allows you to reduce the risk of eventually developing dementia.

And that's why I would like to know. I would like my parents to know. I would like people I know to know their risk before actually developing a diagnosis so you can actually postpone the onset, or eventually, you might not even ever develop the disease if the risk is managed properly.

Dr. McDonough:
Let me swing it back to you now, Arvid. I know the Gates Ventures ADDI Alzheimer's Insights Prize for AI focused specifically on AI for Alzheimer's research. What did that competition force you to do to clarify about the problem with FINGERPRINT and what you were trying to solve? Where does that all fit in? It's so complex.

Mr. Gollwitzer:
Where this becomes really powerful is if we don't only focus on the discovery of new diagnostic biomarkers, but also biomarkers and multiomic precision signatures that help us understand how we'd respond to certain interventions. And what that enables is we get to do both early diagnosis and risk quantification and then use this model to come up with complex multiomic precision signatures that couldn't be uncovered by just a human-led analysis alone to predict how we respond to different interventions or combinations of interventions.

And, that's what this competition and all these ongoing efforts have helped us accelerate a little bit—to really go ahead and use some of these datasets that are available but that I believe have been heavily underutilized. We never really looked at these datasets in this very specific, special new way that we have, where we've built a foundation model pre-trained really on all the data from around the world. And we've built this baseline knowledge into the model about how multiomic data works and how these different multiomic signatures relate to human health and disease. And then we went ahead and took that pre-trained model and fine-tuned it on datasets specific to Alzheimer's disease and from the World Wide FINGERS network. Those efforts have really proven the scaling loss we've observed. They've shown that we're able to come up with those multiomic precision signatures not only for early risk quantification and diagnosis, but also to accurately predict the outcome of interventions.

Dr. McDonough:
For a layman or for a physician who doesn't do a lot of work in the technical computer world, how does FINGERPRINT work?

Mr. Gollwitzer:
FINGERPRINT is a platform that ingests multiomic data from various sources. So this could be sequencing data, or metagenomic data from the gut microbiome, so the different microbes that live in our GI tract. It could be metabolomic data, so the different metabolomic molecules that exist in our digestive system as well, and proteomic and whole genome sequencing data from blood samples.

And this is a lot of complicated multiomic data and multimodal data together with clinical annotations that we ingest into this platform. And we can ingest this data from, for example, one specific patient as he moves through a clinical trial. Then what FINGERPRINT does is we see the first few data points of this specific patient as he moves through the trial, and based on inputting the first few data points of that patient, we ask the model to predict the trajectory of that patient and treatment outcomes.

Later, we go ahead and compare against the ground truth of what really happened in our retrospective analysis of a prior clinical trial, for example. And what we'll see is that in the first few attempts, our predictions are going to be a bit wrong, right? We didn't get it completely right. The ground truth is going to diverge a little bit from what actually happened for that specific patient. And we're going to do that a couple of times to teach the model to predict treatment outcomes really well. So we're going to run this learning loop a few times. It's called few-shot learning. And the output of this effort is a specific multiomic precision signature that is really predictive of treatment outcome for a specific intervention.

And so then what we can do is take any new multiomic data set that has a lot of information, we can throw it into the FINGERPRINT platform, and we'll get the prediction of treatment outcomes and the risk analysis as well.

Dr. McDonough:
Thank you for that explanation. Adrian, where do you see this having an impact short-term and long-term clinically?

Dr. Noriega de la Colina:
Long term, I think the vision, and this is from my early days in medical school and later, what I always wanted to see—I remember seeing this article, I think with the New York Times, saying, "Are we ever going to develop personalized treatments?" Meaning pharmacological treatments designed specifically for you and me. And that promise has been evolving, and I haven't seen this getting to the real world. But I do see a future in the near term now that we have the computing power to actually personalize interventions. We are doing this with lifestyle, where you can actually, based on your genome, your proteomics currently, and other clinical information, you can actually tailor an intervention devoted specifically to any of us. So that can be the same premise that can be applied eventually to

pharmacological treatments. I don't think that we're there now, but the long-term vision, I think, for the field—not only for Alzheimer's disease, but for most chronic diseases—is to create a personalized pharmacological treatment on top of personalized lifestyle interventions.

Short term, what this means is that we are able to select patients better, and I think eventually there's a number of initiatives currently ongoing that allow us to translate that and how that is being applied both into primary care and specialist care in a number of diseases. So I think it's a very exciting time for the field and for the broader precision medicine field, where we're able to provide better care and better recommendations for our patients.

Dr. McDonough:
Can either of you walk us through a real-world example from the platform, ideally from FINGERPRINT or FINGERS-7B work, where the system started with a research question, pulled together the heterogeneous data, ran analysis, and generated useful insights?

Dr. Noriega de la Colina:
The one that I love that is quite recent is when we're talking about lifestyle interventions, and there's always the assumption that when you, for example, as a family physician, present a patient and say, “Well, let's do some lifestyle changes because you might be at higher risk of a number of diseases,” let's say diabetes or higher risk, because of family risk factors, of Alzheimer's disease. So you're thinking, “Let's do a lifestyle intervention.” Let's propose you're going to do 150 minutes a week, and that's it. But what we learned from the analysis that went into our foundation model, FINGERS-7B, was that not every individual is going to receive benefit from that kind of lifestyle recommendation. Not because they're not achieving and the level that we're suggesting of activity in minutes per week; it's just because they might have already stayed in their disease. That is, for example, high inflammation that is going to prevent them from arriving to any benefit. For example, a proteomic marker called TREM2, which is based in your genetics, triggers an inflammation pathway. So we're using, for example, data from one of the lifestyle multi-domain interventions that was based in Finland, and we learned that a fourth of those patients who were doing the intervention and were doing everything right for two years didn't see a benefit cognitively just because they had this high inflammation burden. So even if they were adhering to the intervention, and even if we're following all the recommendations, one out of four didn't see a benefit, and it was triggered by these proteins that are triggered to an inflammatory profile.

So that I think supports a thesis that we shouldn't necessarily assume that every recommendation that we just give in general is going to be beneficial for every patient we’re going to see. There's going to be individuals who, because of the genetics or the state they are in in their patient journey, are going to be much more advanced, and they're going to require probably some sort of combination therapy in between lifestyle and pharmacological treatment.

The intervention was effective in three out of four. That’s wonderful. But I think the patients that come to see you and any physician at the primary care stage are going to be those who are already developing the symptoms. So they're probably going to be in that portion of one out of four who are at higher risk and have all of those comorbidities.

Dr. McDonough:
For those just tuning in, this is *The Convergence* on ReachMD. I'm Dr. Brian McDonough, and I'm speaking with Dr. Adrian Noriega de la Colina and Arvid Gollwitzer about FINGERPRINT, an AI-driven research platform designed to help advance Alzheimer's prevention by reasoning across lifestyle, clinical, genomic, proteomic biomarker, and population-level data.

So, the Davos Alzheimer's Collaborative and FINGERS Brain Health Institute collaboration emphasize globally representative data. “Global” is the key word there. What kinds of populations and environments need to be represented before a model like this can responsibly inform prevention science?

Dr. Noriega de la Colina:
I can mention the World Wide FINGERS. There's a number of studies that are targeting a multi-domain intervention worldwide. That includes studies in Sweden, Finland, and Spain. There's initiatives in Malaysia. There's the African FINGERS. There's initiatives in Latin America, the US, and Canada. A broad set of over 70 countries are thinking and coordinating a harmonized response-intervention base for lifestyle. That covers, I think, worldwide.

The newest of them all, I think, is the one based in Africa that covers, for example, Kenya and other countries. So it's really a global scope. And perhaps, Arvid, do you want to comment on the Davos group?

Mr. Gollwitzer:
What's really important for models like our foundation model, FINGERS-7B, is to see very diverse data, for example, from all over the world, where we see various biases being introduced into these datasets. For example, we get to see, as I've previously described,

longitudinal datasets where we get to see how a certain patient progresses through a clinical trial in Finland, and then we can look at similar or the same interventions from somewhere else around the world. And what we want is for our model to reliably come up with multiomic precision signatures to predict treatment response, both in the Finnish cohort and, for example, in the African one or from all over the world, including India.

And for this reason, we've set up a collaboration with the Davos Alzheimer's Collaborative, which has a lot of these cohorts available all around the world. They have a lot of multiomic data that would really benefit from this large-scale AI-enabled analysis, where we're able to take all these datasets and ingest these large amounts of multiomic data, coming up with non-trivial multiomic precision signatures.

And one of the hardest things in both pre-training and fine-tuning foundation models like these is dealing with the bias in those datasets. So for example, what could happen is that a model only performs well on a certain population or on a certain cohort, but our insights do not generalize to other cohorts. And we've come up with novel computational methods to harmonize these datasets, and then we pre-train our model on a few datasets from specific cohorts while holding out on another entire different cohort from a different clinical trial. And we took our frozen model that's only learned from a few specific cohorts, and we've asked it to apply what it's done in terms of discovery of multiomic precision signatures to the holdout set, and we'll see if we can still correctly predict treatment response and if our biomarkers and multiomic precision signatures still work in this different cohort.

And what we've seen is that this is possible. There are computational solutions that can be developed to build those generalizable models and multiomic precision signatures that work across cohorts, with bias introduced from, for example, environmental exposure or just different genomes and different diets, et cetera, around the world.

Dr. McDonough:
I like what I'm hearing when you talk about 70 different populations from around the world. That gives you a lot of that diversity. And to your point, diet, nutrition, and environment can be better evaluated. Without getting too deep, but at the same time thinking about issues, we talked a little bit about the potential for missing data or underrepresented populations, but what doesn't work yet? Where does your model, as you see it, struggle right now in its present state?

Mr. Gollwitzer:
There are still a lot of limitations attached to these foundation models. What we've seen is that if we have an abundance of data available, and sequencing data has become incredibly cheap, the cost of generating sequencing data has fallen faster than Moore's Law. And as a result, there is an abundance of sequencing data and generally multiomic data available. There are other data domains that are still more complicated and expensive to generate data, but I think we'll see the same exponential decline in cost for these domains as well. But what we don't have yet is the amount of multiomic data that we need for really rare diseases, for example. And often what we found as well is that for a specific disease, we're lacking the large amount of data that would need to be collected in combination with the specific indication to be fed into a model for pre-training. So there is data available, maybe for a very small population or a rare disease, but it doesn't work together with the tools we've built and with the foundation models we've built.

I think two things will address this in the future. The cost of generating those data sets will become much lower; it'll fall continuously. I think it'll decrease even faster than the cost of running computational analysis. And that'll give us, again, the situation where we have an abundance of data, where we've essentially collected data sets related to different indications and diseases that we didn't necessarily know at the time made sense. But then after some time, we'll have accumulated all these different data sets and say, "Okay, we can look at this now using this foundation model. We can come up with relationships that'll help us both for diagnosis and predicting the outcome of interventions."

And the other enabling component that we yet have to see is just better tools for building, training, and evaluating foundation models like we've built. What you want to ideally see is the ability to take large data sets from clinical trials and just, throw this into a platform like FINGERPRINT, and it does the work on its own. You check back in after a few days and you'll have really excellent results where you run counterfactual inference on these models and do few-shot learning and reinforcement fine-tuning on the models using these data sets. So for example, inputting a data set from a specific patient where you already know the diagnosis and you know how he responded to a specific treatment and asking the model to make a prediction and then compare against the ground truth. Tell the model how it got it wrong. Do this a couple of times. This is still something that requires a lot of human observation and human control. Ideally, and this has been some of our efforts at FINGERPRINT as well, you want to run this largely automatic, agentic way.

Dr. McDonough:
So, Adrian or Arvid, what would you see as a bad use of this technology? We want to focus on the good obviously, but what we see as a bad use?

Dr. Noriega de la Colina:

A bad use of the technology, I think, is deploying it perhaps too early in the clinic. There's a process here that requires, as I already mentioned, enormous amounts of data. I don't think that for deploying in the real world for a specific individual we're there yet because we're still mapping the world variability in terms of proteomics and genomics. There's so much that we're missing. And for example, in a population like the US, which has a very diverse population from people from across the world—I think those efforts are specifically beneficial for populations like the US that have that kind of representation, right? It's not as homogeneous a country as China or Japan, for example. So the benefit becomes very translational.

But the use of the technology in the clinic directly, for example, requires a little bit more work. I think the field is advancing so fast. Six years ago, you couldn't have thought of anything like this because you didn't have the tools. Some of this technology for proteomics didn't exist yet. You didn't have the computing power. I think now the models are becoming stronger but still require that input of data, specifically for proteomics. I think genomics is getting cheaper. Proteomics is still quite expensive in comparison, and specific technology for brain diseases is only getting there. With proteomics, you have different levels, and some of them allow you to capture more a central nervous system component, which is technology that goes to the atom level—very specific, very tiny. That is much more related to the brain, and that is still expensive in comparison.

Dr. McDonough:
I want practicing healthcare providers to get an understanding of where we are now and what remains investigational, so let me ask you a couple quick questions. What should clinicians understand now about Alzheimer's prevention that they may not have learned in training? That's a big one, but what would you say at this point—assuming you trained 15, 20 years ago—what should we be understanding now about how things have changed?

Dr. Noriega de la Colina:
I think the most important one is early detection is key. Using the right tools usually can be a simple cognitive test very. And many of you have heard of the Mini-Mental. There's also the MoCA test, the Montreal Cognitive Assessment. There's a number of other technologies and tools and digital assessments.

The important thing is to assess, specifically if there's a family history—you want to do the work of that early diagnosis. I think that's a message that is not always conveyed, especially in primary care or in other specialties outside neurology and geriatrics. So that's the first one.

Now, the second one is there's things that you can do. So once you identify certain risk factors, that's where you want to have that early diagnosis and start acting upon it. And there's a time component, which is essential. Those individuals who are diagnosed and are in a mild cognitive impairment stage—that's where you want to make sure they are captured and taken care of very quickly because there's a time window where they're eligible for current treatments. We might have discussions about whether those treatments are as effective as we want to, but we want to give them the option of having those.

And the third one is we're getting very close to, I think, making substantial progress on Alzheimer's, so bringing hope to those patients. They're probably going to feel anxiety receiving some kind of assessment or a diagnosis because there's no treatment in their mind, or they are not informed of treatments. So you want to reassure them, I think, that there is considerable hope in this space, and there's progress being made every single day. I think this is the closest that we've been in the field to achieving or arriving to a cure within our lifetime.

Dr. McDonough:
In a prior episode of *The Convergence*, one of our guests said to me, "Brian, if you think that you as a physician are controlling the narrative about how AI is being used, forget it. Patients are out there; they're already using the tools that are there, and they have expectations." So with this particular topic, how should we talk to patients who are already asking about whether AI can predict an Alzheimer's risk?

Dr. Noriega de la Colina:
So I'm sure with many of these large language models that are available to consumers and patients—you can have a subscription to Claude, ChatGPT, or whatever, and I'm sure there's people, even if you tell them, "Oh, you shouldn't put that information there," people are doing it, right?

I think we can fairly assume that some people are doing it. And you have to remember that while there might be a prediction that is imputed there, it might not be tailored to you specifically. So you have to take whatever information that you're taking from those models with a certain level of precaution and care.

I was telling you before that I don't think that this technology's ready necessarily to go to the individual level with confidence or a high

level of certainty because we're lacking that kind of representation. And I think the same goes for the models that are deployed for consumers for a number of applications. They are going to get very close, and you might get a hint of your individual risks, specifically if it's tied to input of, let's say, genomic sequencing to any of those tools—you're going to identify specific regions that might make you more prone to a specific disease, and in this case, Alzheimer's disease. But even in individuals who have a genetic test, for example, APOE4, which is a genetic marker for Alzheimer's disease pathology—not everybody who has that specific genotype is 100 percent going to develop that pathology. There's a much higher likelihood, but that doesn't mean that determines your destiny.

So there's a lot you can do on your side. If you're receiving any genetic counseling, you're going to know that it's not set in stone, and there's things that you can do to reduce your risk regardless of that genetic input. So I think the lesson to learn is the technology's advancing very quickly, but it's not necessarily ready to go to the individual level with the highest degree of certainty. And then you just have to be careful. If you see that information, always talk to your physician, consult with your physician, make sure that there is a communication channel going there.

Dr. McDonough:
Arvid, I've been waiting to ask you this: what is the one thing you wish every physician understood about AI models in biomedical research? I'm sure that always is going through your head.

Mr. Gollwitzer:
That's great. Thanks so much. I think what we've seen is that these tools are just heavily underutilized. That's also something the competition we've done with Gates Ventures helped us with—really bringing some of these advancements we've seen in AI to new problems that haven't been, addressed with these tools yet and are lagging behind significantly in terms of what we can do.

Now, I'd invite experimental use of tools like this. A big ethical challenge is basically that you might be tempted by having foundation models that are pre-trained on just population-level data where there is data available. But what happens is that you maybe just want to focus then on these populations where these datasets exist, and it's just maybe too much effort, too much work, and too slow to collect additional datasets. One way around this is really to take these models as early as they are—and there's still a lot wrong with this, and this is super experimental—and just try them out, right? We have a public demo for FINGERPRINT that could be used as a clinician-facing front end to really ingest multiomic, multimodal data that's available. And sometimes you can only have the simplest datasets—just clinical notes in natural language—and that'll still be really helpful to come up with feedback on early risk quantification.

And what I wish people did is just put all these tools to use—put both reasoning models and new foundation models to use—and maybe in return, while these things are still under development, contribute some of the data that'll help improve these models. And if that setup existed and these fields were interfaced like this, that'd be really powerful.

Dr. McDonough:
For both of you, my final question: five years from now, what will distinguish a clinician who understands AI-enabled prevention well from one who uses or interprets it poorly?

Mr. Gollwitzer:
I think what we'll see is in general that there are components that can't be replaced with AI. But using these different tools—agentic models, reasoning models, and new hypothesis-generating models like FINGERS-7B and FINGERPRINT—it will basically enable the generation of hypotheses that are too complex and too unintuitive basically for the clinician to come up with alone.

But once you do see these hypotheses, it becomes a tractable problem, again, coming up with strategies to validate this and to iterate through different options of what you can do that can be, for example, for certain suggestions of early risk or a response prediction. And I believe we'll see that these tools, just like any other tools that have been developed, will help move the field towards prevention in generally and also towards, just reduce the number of failed treatment attempts—just getting it right after much fewer iterations.

Dr. Noriega de la Colina:
I feel that because of the load we have in terms of healthcare spending and costs, eventually we're going to realize that we need to go earlier as healthcare systems, not necessarily as individuals. And that's going to change, eventually, physician behavior, which is extremely very difficult to change. As you probably know, changing processes within any healthcare system is very difficult, and specifically, physician behavior is a whole new space to modify. But we're not going to have, I think, a choice because of, first, the aging population that we face. That translates into a number of chronic diseases. We're already seeing family physicians especially overwhelmed by patients.

So the fact that we're able to predict and identify earlier, I think, is going to be forced upon the system as a solution to the high healthcare costs that we see, because it's much cheaper to treat a patient and prevent them from going into a high-cost state as a

chronic disease burden intensifies. So these tools allow us to do exactly that, not only in Alzheimer's disease, and not only neurodegenerative diseases, but across a number of chronic diseases overall.

So I think with those technologies, there's going to be an intersection of better treatments available. And then as I mentioned at the beginning when I did the example, there's going to be a portion of individuals, because of their genetics, who because of their baseline characteristics are going to need a combination treatment, not only a specific lifestyle modification. And those individuals are going to be able to benefit from not only early detection, but also a combination of both pharmacological and lifestyle intervention. And I think that goes back to my mention of when I was growing up and I saw this New York Times article saying that eventually we're going to get to a precision prevention era where we can actually provide treatments for you and me specifically, and not only for the general average.

Dr. McDonough:
That's a useful place to end—not with the idea that AI gives us a single Alzheimer's prediction test, but with a more careful view of how prevention science may be built. Better data, better models, better validation, and better ways to identify who may benefit from intervention.

I want to thank my guests, Dr. Adrian Noriega de la Colina and Arvid Gollwitzer, for joining me to discuss FINGERPRINT, FINGERS-7B, and the future of AI-enabled Alzheimer's prevention research.

Adrian, Arvid, thank you both. This was fascinating.

Dr. Noriega de la Colina:
Thank you, Brian.

Mr. Gollwitzer:
Thank you so much.

Dr. McDonough:
For ReachMD, I'm Dr. Brian McDonough. To access this and other episodes in our series, visit *The Convergence* on ReachMD.com, where you can Be Part of the Knowledge. Thanks for listening!