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/missing-link-clinical-ai-institutional-knowledge/67129/

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

The Missing Link in Clinical AI: Unlocking Institutional Knowledge

Dr. McDonough:
Clinical expertise can only get us so far when we don't understand how care delivery works in our own institutions. What does AI need to understand about individual institutions to be genuinely useful to clinicians?

My guest today, Dr. Declan Kelly, encountered this problem while working as an NHS physician and went on to build Eolas Medical, a platform that aims to make clinical knowledge easier to find and use.

Dr. Kelly, welcome to *The Convergence*.

Dr. Kelly:
Thanks very much, Dr. McDonough. Thank you for having me.

Dr. McDonough:
Tell me a little about yourself, your clinical experience, and your interest in doing what you're doing, and then we'll get into questions.

Dr. Kelly:
So I'm a medical doctor by background. I worked clinically mostly in emergency medicine, and I very much loved that career, and I was hurtling down the path of that. I wanted to do all training in critical care and emergency medicine and had done all my postgraduate exams. And I started coding again. So I did a lot of software engineering during school and really enjoyed it, and I never really thought I was ever going to really get back at it in any sort of serious form. But I started coding again and just fell in love with it. And the ability to create something from nothing that then actually solves a problem in the real world can become quite addictive, actually, and quite enjoyable.

So we developed a solution for our own health system, and then within about six months, that was used by about 20 or 30 health systems across the world. So it became evident this was becoming a problem. That's kind of the path that I went on. And it seems to be a path that's shared actually by a lot of physicians. So there's seven physicians in our team, and six of us are actually software engineers. We have colleagues in lots of different other companies or health systems who are all doing a lot of engineering on the side or as their full-time job. So it's funny. It's a strange overlap.

Dr. McDonough:
I think it's really interesting, Dr. Kelly. I shared the one concern that you had, that overriding concern about just understanding healthcare delivery, even in my own hospital—the protocols, all these rules. I knew the medical side but struggled with the other things as well as the resources. And I think that's what you went after. Take me back to the time when no guidance existed, when you recognized the problem working in your institutions. What was it that you were looking for that caused you to get back into this?

Dr. Kelly:
Yeah, great question. So I suppose exactly what we do and what we've built was based upon the problem that we'd experienced ourselves. When treating patients, we found it very difficult to find the information that we needed that was based in our own health system. So someone's having a major heart attack, what's the referral pathway for a STEMI? How do I refer them to cardiology? What antiplatelet do we use in this hospital? Or someone's failed their cancer treatment; what medicines do we actually have in formulary? What's our next path? How do we get them referred to the social worker, and what should we do from a follow-up point of view?

So it allowed us to look at this problem holistically. Doctors make about 100 decisions a day, and we started looking into this to try to

© 2026 ReachMD

measure this. The majority of those decisions actually aren't purely clinical decisions. Those decisions don't happen in a vacuum. They happen within their own health system, which is different from the health system in the next state or the next country and different in the health system that exists across the road. So that context needs to be taken in.

So we then came up with this framework of doctors either need information from external sources or internal sources. And with the external sources, that problem has now largely been solved with external sources such as UpToDate or the latest journals or OpenEvidence, et cetera, and that's now becoming democratized. You can turn to OpenAI or ChatGPT or OpenEvidence and you get all the information at your fingertips. But the ability to search for that content within your own health system is still a bit of a mess and a bit of a nightmare. Usually, this is either a policy system or SharePoint or stuck to the wall or in someone's drawer or in someone's head or all of the above, and that's very difficult. And I think what we're now excited about in healthcare is the ability to start to move towards full AI adoption that allows us to start doing autonomous tasks in healthcare, and that's only possible if you really unlock all of that knowledge that exists within your own health system.

So that's the system that we ultimately built. We built that about four years ago and launched it, and it's now live in about 500 sites across the world. About 90 percent of all acute hospitals in the UK use it, and then in the US, there's about 50 centers using it. So Stanford, Mass General, Boston Children's, and a number of other centers are all now using it. That's really, I suppose, the path and a bit of background about the journey as well.

Dr. McDonough:
So for those who don't really know, you started a company called Eolas, and you built this more or less while you were an NHS physician. Tell me a little bit about Eolas and also how being a software engineer in your background helped you look at these problems in a different way.

Dr. Kelly:
So Eolas, if you're interested, is the Irish word for “knowledge” or “information.” And it’s quite difficult to pronounce, actually, because it's not pronounced how it's spelled. So we probably should have chosen a name that was a wee bit easier to say, but we're four or five years down the line now. And it was chosen because of the problem that it solves.

So to your question of how having an engineering background helps solve problems in healthcare, I think it's the combination of, is this a technical problem, or is this a process and people problem? And in healthcare it's usually both, and the fact that you're able to take it through the lens of saying, "Okay, here's the technical aspects of this problem. I know I can get the solution this far." But then having the healthcare hat on of working within the system, knowing how complex it is, and knowing how patients flow through the system, and then being able to then map that back towards the technical problem is the key.

I think with healthcare, there's no shortage of technical solutions. The shortage is in their ability to implement and integrate into the workflow effectively. So it is really the combination of understanding the workflow, understanding “will I use this?” as a physician, and then understanding how to solve that as a problem and meet in the middle. And I think physicians are brilliant at not using stuff. If it makes their lives any harder at all, they'll find a way around using something. And adoption of software in healthcare and in health systems has been very hard historically, whereas if it saves their lives, or if it saves their working life some time, they will adopt it.

And we're seeing this with OpenEvidence, used by 80 percent of all physicians in the US because it just saves them time to access that knowledge. You have the ambient scribes, which are adopted widely because it saves them time with their documentation. So if you're able to understand what they need and the workflow that they work in and how to integrate into that, as well as technically how to build that, I think that's a very powerful position to be in.

And I think we've seen that with these other companies. Shiv Rao, who's the founder of Abridge, which is one of the biggest ambient scribes now in the US, is a cardiology doctor by background. You have Heidi, which is one that's very big in Europe; the founder is a surgeon by background. So you have these doctors coming through, finding it easier and easier to build software, but then combining that with the understanding of, “would I use this?” So you're starting to see this pattern repeat over and over again.

Dr. McDonough:
You bring up a couple of really interesting points. First of all, I can imagine a lot of our viewers and listeners are physicians, obviously, and they're probably agreeing with you, saying, “Yes, we've been wanting somebody who thinks like a doctor to put these things through." You really did touch on some key points—the fact that we want to save time and that we're tech-averse if it's going to cause problems, and you try to find your own way to do things.

I do want to ask you—you talked about institutional knowledge, and it's so true. I knew numbers. I knew how to dial the emergency room or security or this or that, but basically, I never knew how to really put things into play, and I also didn't think about it. We sent somebody

© 2026 ReachMD

from the emergency department to the ICU. Until I got involved with the informatics side and the computer side, I didn't know all the work that was required to get that person just to the ICU. The monitor travels with them. Nursing has to know. You have to have somebody cleaning the ICU to prepare for it. They have to set up the bed. There's so much that goes into it. I can't imagine how you built something that takes that institutional knowledge. How tough was that challenge?

Dr. Kelly:
Yeah, it's an extremely difficult challenge because I would say only 20 percent of that is a technical challenge. The other 80 percent is a change management problem. It's a key stakeholder engagement problem. It's finding where this information lives and getting it onto the platform. Sometimes, integrating with the systems can be technically challenging, and that makes it slightly more difficult. But yeah, it's finding where all of this knowledge lives.

And I suppose if you look at the external world's medical knowledge, it might seem that there's so much more knowledge there in the world, and that should be harder, but that information is generally quite well organized now. There's probably about 50 to 100 journals where all of that information lies. There's probably about 200 to 300 national guideline bodies that all of that information lies. Most of it's in textual or pictorial format, very easy to take down and then organize and fine-tune a model on that. Whereas all of the stuff that exists within your own health system is scattered across everywhere and takes a lot of key stakeholders to manage it. So it's very complex.

And the other complexity lies within the combination of clinical and operational questions that get asked. So we now have about 400,000 clinicians using the platform, and they're using it millions of times per month, and we can see exactly what they're asking in real time and when in natural language. And we generally categorize that into, is this an operational question or is this a clinical question? Clinical question: I have a child with meningitis. What cephalosporin do we use? What antibiotic do we use? And that's very much a clinical question specific to that healthcare system. And it might contain dosing. There might be a clinical risk associated with that answer that you give.

But the operational question, as you said: how do I transfer a patient to the ICU? Or, for example, I chatted with my wife this weekend, and she's an oncologist; very often, when she's on call, she'll get someone who may need radiotherapy to their spine or radiotherapy for a lung tumor or something like that in an emergency setting. And there's a clinical decision of, does this patient meet the criteria for needing emergency radiotherapy? But then 90 percent of the rest of that decision is, how do I activate the radiographers and the radiotherapists to start the lab up so I can get them transferred? This patient is in a different district general hospital or in a different site. How do I organize a transfer across here? Do we have a bed? Do we not have a bed? Et cetera. And she's having to liaise with all the people in the team to try to flow that decision through. So only part of that decision's actually a clinical decision.

A lot of the other factors into that decision are operational, and so many of the decisions that we make in healthcare as physicians are a blended combination of operational and clinical decisions. As a primary care physician, you'll see someone in your clinic, and more often than not, the decisions that you make for that patient are informed by their social circumstances, their family in the room, et cetera, and it's not just purely a clinical decision. So having access to all of that internal contextual information is really important to help inform, should you actually do this for them, or shouldn't you? How do you do it? And that can be difficult to get that. And usually, how we've got that is you have to work within the system for 5 to 10 years to realize how it all works, and then you've finally got it, and then eventually then you leave the system, and that knowledge leaves with you. So that's really what we're trying to tackle.

Dr. McDonough:
So you gave some good examples, but for a clinician who's never used Eolas, walk us through a question. You can pick one from start to finish, like what's being searched, what's retrieved, what's generated. Maybe an example like that.

Dr. Kelly:
Yeah, so a very good example where we normally start in a health system is in infectious diseases. It's the perfect example because usually there's local antibiograms. So based on the resistance pattern of the bacteria in that hospital or that region, the antimicrobial pharmacist and the infectious disease doctors will decide, "Actually, you know what? We have a high resistance to this antibiotic, so therefore we should use this one," et cetera. So external knowledge doesn't fully cut it, or else doctors end up using very broad-spectrum antibiotics.

So how it exactly works is that team would upload all of the content, so the clinical policies, procedures, and guidelines around all of their antibiograms and all of their procedures for how to treat sepsis, et cetera. And then the AI would chunk all of that information up, so it doesn't matter what structure it's in. If it's charts, if it's formatted badly, et cetera, the AI is able to understand all of that because we use a technique that uses computer vision to do that.

And then it means that the doctor can then just take the phone out or go onto the EHR and just ask a question: "Hey, I have a three-year-old. They've got meningitis. They've got a penicillin allergy. What antibiotic should I give?" And then it's able to then query the data

© 2026 ReachMD

that you've uploaded and then just give the answer straight away. So that's a high-level overview of how it works.

Dr. McDonough:
So let's say I get an answer. What can I inspect or look at to verify it? Obviously, I want to trust AI, and I want to trust the information, but what do you normally do? And can well-referenced answers go wrong? Have you seen that?

Dr. Kelly:
Yeah, it's a great question. 100 percent of the time, the answer will be referenced to the source that it's come from. So it'll give you the answer, and then it'll say, "This is where I got that answer from," and you click the hyperlink, and it brings you to the point of that guideline so you can independently verify it.

AI, by its very nature, isn't deterministic; it's probabilistic. So for maybe non-technical people listening, software in general can be deterministic—so if this happens, do this thing. If this happens, then do this thing. So every single time, the same input will result in the same output, and that's, up until the last couple of years, how our relationship with software was. We expect the computer to do this. It's ones and zeros. If this happens, do this.

That has now drastically shifted to probabilistic, and ultimately, that is a statistical model under the hood. These large language models, as powerful as they are, they're just very powerful word prediction machines. So therefore, because they're probabilistic, they can never be 100 percent accurate. You put an input in, 99 times the output will be the same or similar, but one time out of that 100 there, there might be an error and might be a mistake. And that's really important because we need to build systems around how that works and how we verify it, so that's why we always do the independent verification.

Our aim is to keep that as low as possible and to study that and to monitor that. So we had independent physicians review answers, and we had them ask, "Was this different than the answer that was given? And if it was different in any shape or form, was it clinically safe?" So 100 percent of the time, they deemed that the answers were clinically safe. We then scaled that up to a large language model as a judge system that we do continuously. And we did 6,000 questions, and it found that 9 out of the 6,000 actually had an error in it where it extracted the information wrong. 99.85 percent of the time, it was accurate and safe, but always, there is that aspect it can get the answer wrong. That's why there's a human in the loop, and that's why the references are always easily accessible and referenced. So it's really important that physicians using these types of technologies understand that.

The final point I'll make on this is, this isn't new in healthcare for us. We're used to having a tool that's not 100 percent sensitive or specific and then making a decision based on that tool. And the example that I always use, as an emergency physician, is a troponin blood test. So if someone has chest pain and you do a troponin test three hours after their chest pain, it's probably about 98 percent sensitive roughly of picking up a heart attack. And therefore, if you decide to send everyone home, you're going to miss 2 percent of heart attacks. Does that mean that system is bad? No, it's part of a wider system. So then you take a history, you look at the EKG, you try to understand that more deeply, and then you repeat the troponin blood test, and then you're able then to rule in or out more confidently.

So everything we do in healthcare, all the tests and all the tools we use, none of them are 100 percent, and we combine them all to then use our clinical judgment to make a decision. These AI tools are exactly the same; they're never going to be 100 percent, and we need to then take that into consideration and make a decision based on that or build a safety system around them.

Dr. McDonough:
I think that's a wonderful point you're making. One other question that I have in the back of my head is, let's say a hospital's protocol's outdated or it conflicts with newer evidence. It doesn't fit the patient in front of you or that kind of a thing. What do we do then? Are you able to see if they change things? What should the clinician do next?

Dr. Kelly:
That's a great question. So broadly speaking, we solve problems for two people. One is the people who manage and maintain the guidelines, and the other is for the people who want to try to access those guidelines to make a clinical decision quickly. And for the people maintaining it, the problem that you're ultimately describing is, we have this hospital guideline; is it in keeping with national standards, or is it out of date? So our platform actually is able to scan national guidelines, compare that against your own, and then highlight the difference between them. And then you can then decide, "No, that's a deliberate difference. We're happy to still use different medicines because we don't have that medicine on formulary." Or it highlights and says, "Oh, wait a minute. Yeah, we probably should be changing the dose or changing the medicine or changing our approach based on the latest evidence." So we do that for the people who manage and maintain the guidelines. Then that goes live, and then the physician can confidently get that latest version.

But it's a massive problem when a lot of the time, we see that these guidelines are actually out of date, and sometimes they're 10 years

© 2026 ReachMD

old and they've never been updated, and medical knowledge changes. It doubles every 73 days. So these things should be changed more frequently, and AI's perfectly placed to help health systems do that.

Dr. McDonough:
I know you refer to the golden triangle of clinical AI: patient conversation, the chart, and clinical knowledge. Tell me a little bit about that. It's an interesting concept.

Dr. Kelly:
I think that was really formed around the idea of, okay, where is all this going? AI is moving so fast. The whole world's going a bit mad. It's hard to predict where things are going. And very quickly, the pattern we recognized after doing a large amount of study and having a lot of conversations with health system leaders was there are three main buckets of technology that are now being adopted using this generative AI that focus on the patient-doctor interaction. And it essentially has the ability now to fully digitize everything that happens downstream and during that interaction.

And what do I mean by that? So if we zoom out and think, okay, what does a physician do when they're talking to a patient? You generally do three things. One is you listen to the patient and do a history and examination. “We know why you're here, what are your signs and symptoms?” That’s examining the patient. But you can't take that out of context. You have to know then everything that's happened to that patient previously. So what's all their past history? What are all the previous blood tests and labs and scans that they've ever had? And that's taken into context too. So that forms two points of the triangle. Then, as a physician, what you do is you take that information, and you query a knowledge base. And that knowledge base, historically, was the stuff you learned in medical school. And then it became the stuff you learned in medical school and your colleagues' opinions and maybe looking up the latest evidence. But ultimately, you were querying a knowledge base to try to make a decision, and then you'd make a decision, and everything from that point of decision onward informed what's going to happen to that patient next.

And ultimately, those three points of the triangle, which is everything that's happened to the patient to this point, i.e., the chart summarization, what's happening in the room, the conversation synthesis, and the history and examination, and then all of the world's medical knowledge internally in your own hospital as well as externally, all of that now is becoming streamlined by AI. So you're having all of these ambient scribe technologies able to fully record the patient history. You're having all of these pre-charting technologies, mostly led by Epic and Cerner and the EHRs, able to summarize everything that's happened to that patient to date, and then in real time, you're able to get access to the world's medical knowledge at your fingertips using AI.

So all three of those things then can then decide what's the best next decision for that patient. And that is now at the point where it's probably more accurate to be done by an AI than it is for a doctor to think, "Oh, I need to think about what do I know or what do I not know?" Medical knowledge doubles every 73 days. That's more information than any doctor can keep up to date in a lifetime. So that's the thing that we should be leaving to AI. We should be turning and facing the patient more and being more empathetic, delivering information in an empathetic way. It's resulted in an enormous opportunity. There's never going to be enough capacity in any of the systems that we have, and this unlocks that capacity. It helps double or triple the healthcare capacity that we're able to deliver.

Dr. McDonough:
You can look at things in a unique way that many of us can't. When you moved to the US hospitals, what surprised you about how an ordinary clinical day is organized? Will you give an example where workflow might have been different if it in fact was different?

Dr. Kelly:
I think the way health systems are structured is very different. So if we look at, for example, the NHS, the clinicians, so the doctors, nurses, and pharmacists, are all employees under that one health system or that trust and are directly employed by that health system, all working towards that goal. And that's an easier concept to think about, because it just is nicely wrapped, essentially. Everyone's employed. The boss is the CEO of the trust ultimately, and then everyone below them is below them in the org chart, and that's how that works.

The US is extremely different from that. So you'll have patients flowing in, coming into the emergency department. The emergency department is owned by the hospital, but it's run by a different company, potentially. So you have a different company staffing that emergency department, and those emergency physicians will work for TeamHealth or some other big company that supplies those doctors. The patient might be very sick and might have to go to the critical care unit. Sometimes that critical care unit is run by the health system, or it's maybe outsourced again. So then you have these different structures and lots of different dynamics. How medicines are bought and how they're priced in the US, and then how they're used and reimbursed throughout that system, is drastically different as well at every single point in time. So there's lots of different competing incentives throughout the health system as people are moving.

So I suppose the point I'm making is it's enormously more complex in the US than it is in the NHS and in other systems. Healthcare is

© 2026 ReachMD

extremely complex everywhere, but because of all those extra points, it becomes much more complex to get through the system. It also results in a lot more opportunity, because there's a lot more problems that need solving and a lot more people who are willing to look at solutions to help solve their problems because of that complexity as well. And I suppose they're the two major differences from a systems and an operationalist point of view.

Dr. McDonough:
What should US hospitals learn from the NHS? From the standpoint of care for the patient and how it impacts them at the bedside, what could they learn? And vice versa; obviously, there's two sides to this.

Dr. Kelly:
Yeah. I think there's two sides. I think the move to value-based care is the thing that it is actively learning as a system as a whole across the US. So it's a "show me the incentive and I'll show you the outcome" type of thing. And a fee-for-service model obviously creates a lot of a disincentive to actually move the needle on care. If you come into the clinic, I could give you every test under the sun, but it's not going to move the needle for you from a quality-of-care point of view. There's a reason why out of the top 15 countries by GDP, the US spends the most in healthcare but has relatively the worst outcomes. And that's now starting to change, whereas in other systems, it's very much incentivized towards, "Was the quality of the care improved?" And that's what we should invest in. So I think that's probably the biggest thing it should learn, at a very high level, and it is learning that lots of systems are starting to move towards value-based care as opposed to fee for service.

The flip side of that is I think the US is much better at asking the very sensible question of, "Does this help us make more money, or does this help us save money?" Because if it does help us do that, we have the theoretical ability to treat more patients and cure more disease and help more people. The NHS and other health systems are very poor at doing that, actually. They're very bad at analyzing, "Does this actually help us save more money and therefore treat more patients?" And we're starting to build that muscle a bit better, but the costs are just bloating, and we haven't really learned that lesson fully. There's a lot of things that make sense to do, and then we do them, but ultimately, they actually don't affect the bottom line. It doesn't save more money, and if we're not saving more money and making healthcare economics deflationary, we'll never have the ability to actually treat more patients and expand the capacity of the system. So I think that's one thing that the UK should be learning from the US.

Dr. McDonough:
So a little bit more on hospital operations—there were two company-published examples. The Sheffield Teaching Hospitals timed 11 staff finding five guidelines reported. They saved 60 seconds per guideline, but patient outcomes weren't measured. A Brigham and Women's infectious disease and pharmacy case study described mobile access on rounds, and prescribing changes had not been measured. But they're very important reports, and I wanted to maybe look at the Sheffield report that I alluded to, and it shows faster access to guidelines. But how do you test whether that leads to better decisions and better adherence? Did it really make a difference from your perspective?

Dr. Kelly:
Yeah, it's a great question. "How do you measure clinical impact of an intervention" essentially is ultimately the question. Better, cheaper, faster is ultimately the framework that everything comes down to. So in healthcare, better is, "Did this improve the quality and safety of the care? Was it cheaper? Was it deflationary? Did it reduce the cost of care or help us make more money from that care? Or did it improve the efficiency of care?" And that's ultimately the framework through this.

So from Sheffield's point of view, they showed that it was nine times faster to access the information than using another system. It improves the efficiency, so it means that doctors are spending less time on a certain task and able to spend more time decreasing cognitive load. Importantly, as well, that study showed that it increased the usage of the guidelines, and if you increase and adhere to guidelines of evidence-based practice, medical error decreases, and it actually directly impacts practice.

Another study that you hadn't mentioned there was a study that was done with Imperial, and they studied the AI platform on prescribing behaviors versus just using something that was on their intranet, like a PDF, and they found that every other time that a physician used the PDF to make a prescribing decision, there was actually a mistake made. So whether length of medicine or type of medicine, whatever it was, but every other time, there was a mistake made. Using the AI, there were zero mistakes made. So it reduced that error from about 50 percent down to zero, so that has a direct clinical outcome for patients: decreased medical error and decreased litigation for the health system as well. So that's our main focus.

Carilion Clinic in the US as well, in Virginia, they cover about a million patients. They did a study showing that it increased the utilization of the guidelines by about five times, which decreased broad-spectrum antibiotic use and also decreased antibiotic use in general for viral upper respiratory tract infection in kids, so decreased antibiotic use there as well. So you're showing those direct clinical outcomes

© 2026 ReachMD

by increasing the utilization and adherence to the protocols and policies.

Dr. McDonough:
I want to shift a little bit to who owns the knowledge and who has the responsibility. So let's say an approved guideline changes tomorrow; who updates it? How do users know? And what prevents the old information from shaping an answer?

Dr. Kelly:
Yeah, so the system itself is full version control. Once you change it, it automatically shows the new version, and then you can roll back to old versions. You have everything there from a posterity point of view to know what's been changed and when.

Who owns it generally depends on the team. So usually it's decentralized, as in the infectious disease team will own those guidelines, or the cardiology team will own those guidelines, and then they update it, and they go through a governance process of who owns it. Once that's updated and changed, you have the option to send push notifications to everyone on the platform based on user groups to say, "Hey, our sepsis policy is updated. Here's the changes." So that then gets that information disseminated.

But most importantly, then, when the physician asks the question, they always get shown the latest version. So they'll not get shown outdated versions. Most people go to their emails. "Oh, this was emailed somewhere. Is this the final version, or is this the final version?" Or go to a chat app or whatever it may be. So it just solves all of that mess.

Dr. McDonough:
You're obviously responsive to what the providers, physicians, and others are saying. So where have users pushed back or found errors, and has it had an impact on changing the product or workflow that came from it?

Dr. Kelly:
Yeah, we have very deep and granular analytics around how users are using it. We have an enormous amount of quantitative data that we look at on a daily basis and make product decisions based on that. But we also have an enormous amount of qualitative data where we very regularly jump on calls with users and physicians and ask their opinions, and they're thankfully very vocal. They're able to give feedback unabashedly, which is really useful for us when building a product. And it's everything and anything from, "Hey, this answer is slightly too long and too verbose," to, "This answer is too short and doesn't have enough information in it." And then we react based on that. And we also react to the admins who manage and maintain the content of, "Hey, if we had this feature, it would help us solve these problems." So we're constantly listening and iterating on the product based on that feedback.

Dr. McDonough:
I'll probably have to say I'm guilty of this, but I know whether I was using Cerner or Epic, obviously, things are built in to standardize, but we all learn workarounds. We find ways to make it easier. It doesn't necessarily help the analytics, but we're like, "Hey, I've got to get through the day." So making local practice easier to follow is important. But also, if somebody like me is working around the system, how do you avoid making poor local habits easier to reproduce at other locations as well? I can see that happening.

Dr. Kelly:
I think ultimately, the lens through which we see this is back to that idea of solving the problem top-down for someone at the top of the hospital who manages these policies and wants to change practice, and then bottom-up, i.e., solving the problem for the physician. And the other way to look at that is carrot or stick. And usually the health system will say, "Hey, we have a new policy. You have to use this." And then they find it very difficult to try to implement that because physicians will find workarounds.

What we've seen with the ambient scribe technology is the carrot always wins, actually, when it comes to adoption. So as long as you can align with something that's actually a problem for the physician and solve that problem, they will continue to use it. They're avoiding it deliberately, and they're avoiding it because they've got busy lives. They have more than enough things to adhere to, and if they're trying to avoid it deliberately, it's not going to actually impact change and move the needle for that health system. So it has to be very well aligned to the problems that they actually experience or else they won't use it.

You then have to make sure that it's also aligned to the top of the health system. You have to make sure that it's actually solving a problem for that health system as well, and ambient scribe is that perfect example of that combination of both. So if the physicians don't want to use the tool that we've implemented, we're not solving a problem bottom up for those clinicians. And as I said, seven of the members of our team are physicians. We don't want to build tools that physicians don't want to use. We want to make sure that this is actually solving a problem that they want to solve, and that's probably the core of our obsession.

Dr. McDonough:
That's such an advance from when they first came out with a lot of these electronic records. At least in the United States, it was built with an engineering background, but they weren't necessarily always thinking about physicians. It was also sold to the C-suite, and then

© 2026 ReachMD

physicians were like, "What the heck am I doing and why?" It seems like if you give people a reason to use it, they'll take that extra step if I can get data reproduced and it moves the needle for my patients.

So let's bring us to the present time. What should physicians pay attention to now? What should they be thinking of hearing this interview? Where is this going to impact their lives, or where could it impact their lives if their institution adopts it or has adopted it? Tell me about that.

Dr. Kelly:
So their ability to more confidently make a decision faster, that they know it's the safest decision for that patient, and that they get access to that information faster is ultimately the problem that it solves. I never got out of my head one of my seniors who had a patient in the emergency department, a woman with very bad hyperemesis, and she needed to be admitted, and there was a couple of nuances of where and how this patient needed to be admitted. He knew there was a policy, he knew there was a flow for this referral and who to refer to and when, and he knew it was on the intranet somewhere, and he actually shook the screen of the computer and said, "Where is this? I know this is in here." And he shouted at the screen, and I think that those 60 seconds where you're so busy, you have so much to do, and you need access and you know it's somewhere, if we're able to get those 60 seconds down to four seconds of a question, "Hey, I've got this patient," and it gets you the answer, and you can go back to your day and actually delivering care, that's the core of what this solves. And at scale across millions of users, that's actually quite powerful, and it saves an enormous amount of time for health systems.

Dr. McDonough:
What is the one thing you wish every physician understood about clinical AI and the knowledge inside their own hospital? Because obviously all these things are in your brain, but what is that one thing that you would say?

Dr. Kelly:
So the word of the day, or the word of the year, I suppose, is "agents." And I would say I think we're going to look back, probably in five or 10 years, or maybe sooner, and go, "Oh, remember those cute chat apps? They were really nice, weren't they? And you got to chat with it?" And it's just going to look so silly. That was the start of what this actually is.

What this transformer-based technology enables for healthcare is truly autonomous agentic workflows. And I think trying to deeply understand how that benefits patients and how that benefits health systems as a whole is really important because there will be pushback against it. It's already quite polarizing as a topic, AI in general, for multiple different reasons. But I think I see it through the lens of this technology is the worst it is ever going to be, and it's just going to get better and better, and if you stack it on top of each other, we have the ability to completely change how health systems work. We have the ability to launch 10,000 agents into health systems now to do referrals, to do triage, and to do lots of things that doctors shouldn't be doing when they should be doing other things. And being open to that and pushing that forward as a future only benefits patients. So that's probably the lens through which I see it.

Dr. McDonough:
Couple more things. I alluded to workarounds before. What is one workaround you accepted as a doctor that now makes the software engineer in you want to fix immediately?

Dr. Kelly:
Other than the problem that we solve at Eolas, I would say the manual referral process across health systems in general. I think ultimately, a patient's journey is just the transferring of one point of time and type of care that they're getting to another physician or another service line, and those bits are extremely fragmented and broken. So whether it's the patient phoning the primary care physician in the first instance, and they haven't seen anybody, that's extremely broken. Then the primary care physician referring to a secondary care physician, whether it's a urologist or a breast surgeon or whatever it may be, that's extremely broken in those handoffs. Or whether it's an emergency doctor seeing a patient and trying to refer them to critical care or refer them to cardiology, that's all extremely broken, and ultimately, it's a data transfer problem and a decision-making problem, and AI can solve an enormous amount of that.

I think empowering the patients is a big thing that we've done very poorly over the last two or three decades, and I think that's very much shifting. Empowering patients is, I think, going to be more and more. People are going to take that into their own hands. The ability to track your Amazon package down to the mile of where exactly it is now, but you're not able to track where my journey is through healthcare and the most important part of my life—it's just insane. I think patients are going to take back the autonomy over their own healthcare a lot more over the next 10 years compared to the last 10 years.

Dr. McDonough:
Dr. Declan Kelly, I can't thank you enough for joining us on *The Convergence.* I have to tell you, your clinical knowledge and your experience from actively seeing patients really comes through. So many of the things you're saying, I'm like, "Oh my gosh, that's so

© 2026 ReachMD

true." And I'm sure a lot of our audience is saying the same thing, like, "Wow, you really have targeted." So it's great to see that Eolas is looking at these things, and you also have a really good sense of what could be fixed and where we can make gains. So thank you so much for joining me on *The Convergence.* I appreciate it.

Dr. Kelly:
Thanks very much for having me. It was great chatting. Thank you.

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 healthcare, visit *The Convergence* on ReachMD.com, where you can Be Part of the Knowledge. Thanks for listening, and we'll see you next time.

© 2026 ReachMD