

HALO LAB · WEBCAST

Reducing Clinician Burden with AI

A conversation about why AI in healthcare is a workflow problem before it is a technology one, and how to keep clinicians and patients at the center of it.

RECORDED 4 AUGUST 2026 • ⌚ 57 MIN • WEBCAST

HOST

Sergiy Sumnikov

General Manager, Healthtech · Halo Lab

GUEST

Kimberly Noel

Global Lead, AI Advocacy & Digital Health · Roche

KEY TAKEAWAYS

The session's guidance, distilled into steps a team can act on when bringing AI into clinical work.

01 Treat AI adoption as a workflow problem first.

Define the pain point and how you will measure it before you build. Most failures are human-workflow, not technology, and an unclear workflow only gets worse once automated.

02 Re-stratify tasks by risk before automating.

Automate low-risk work like refills, and reserve clinician attention for high-risk steps. Treating low and high risk the same drives alert fatigue and burnout.

03 Co-create with clinicians as partners, not end users.

Assign a clear human-in-the-loop owner for each step, involve clinicians (and ideally patients, ethicists, and legal) from design, and make iteration loops visible.

04 Set a real baseline and measure before and after.

If you cannot define the status quo, you cannot define improvement. Pair clinician and patient metrics such as documentation time, cognitive load, burnout, and encounter satisfaction.

05 Put checkpoints only where decisions are irreversible.

Mirror aviation-style checkpoints at high-impact steps, and avoid adding friction to the natural rhythm of clinical care everywhere else.

06 Guard against over-reliance and de-skilling.

Keep clinicians doing the first reasoning themselves, track acceptance and override rates, make overriding easy, and train on limitations and failure modes, not just capabilities.

07 Build trust on evidence, then monitor in production.

Validate algorithms with the rigor of any other treatment, then watch them live with surveillance, contingency plans, and protected clinician time.

IN THIS SESSION

Why reducing clinician burden with AI is a question of workflow, evidence, and trust before it is a question of technology.

00 Introduction

Kimberly's path from Deputy CMIO to Roche, Human in the Loop, and Nature's npj Digital Medicine.

01 Workflow, Not Technology

Why AI projects fail despite strong technology, and what real clinician co-creation looks like.

02 Human in the Loop & Cognitive Load

Why automation alone doesn't cut load, what success beyond compliance means, and measuring before and after.

03 Automation vs Control

What to automate versus keep under clinician control, readiness, and avoiding over-reliance.

04 Building Trust

Earning clinician trust through evidence, and monitoring an AI system after it ships.

About this transcript. Edited lightly for reading: filler words and false starts removed, and a few clear mishearings corrected. The speakers' meaning is unchanged, and nothing has been added. Pull quotes are drawn verbatim from the conversation.



Introduction

Kimberly's path from Deputy CMIO to Roche, the founding of Human in the Loop, and a seat on Nature's npj Digital Medicine.

Sergiy Sumnikov
HOST

We'll start punctually. Welcome, everyone. Today we have a webcast on reducing clinician burden with AI. My name is Sergiy Sumnikov, I'm from Halo Lab, a UX/UI design company focused on improving clinician adoption of AI products, and it's a pleasure to have Kimberly Noel with me. Hello, Kimberly, how are you?

Kimberly Noel
GUEST

Hi! I'm doing really well.

Sergiy Sumnikov
HOST

Great to have you. Before we dive in: you currently hold the role of Global Lead of AI Advocacy and Digital Health at Roche, and I know you have other activities as well. Could you tell us a bit more about your background and what you're focused on right now?

Kimberly Noel
GUEST

Absolutely, and thank you again for having me. I'm at an interesting crossroads. I'm very passionate about what's happening in AI, and notably in responsible AI. I recently founded a company called Human in the Loop to advocate for clinical validation studies, and I joined the editorial board of Nature's npj Digital Medicine. We're in a really interesting state right now, not only in pharma but in clinical practice, so I'd love to touch on all the different experiences I've had.

A little background: I've spent 13 years at the intersection of clinical medicine, digital health, and AI. In my past life I was a Deputy CMIO at Stony Brook Medicine, leading virtualization during the COVID pandemic. Those experiences, from academia to healthcare practice to my work at Roche, and now accelerating AI through this new venture and through peer review, all feel timely and relevant to our conversation today.

Sergiy Sumnikov
HOST

I really appreciate how diverse your background is, and how many perspectives you bring to what we'll discuss. Let's start general. What lessons from clinical AI adoption should pharma and healthcare organizations apply today?

01

Workflow, Not Technology

Why AI projects fail despite promising technology, and what real clinician co-creation looks like in practice.

Kimberly Noel

GUEST

The most important lessons are intrinsic to the business. In pharma it's about evidence. We would never dose a molecule without proper pharmacovigilance and post-market surveillance, and the same rigor and the same responsibility should apply to AI algorithms. So the lessons of pharma's past that are still relevant today are about clearly identifying evidence gaps and then addressing true unmet needs. There's also a lot to be said about understanding the technologies themselves. In an era of very rapid AI adoption there's a lot of noise, so being able to sit down and understand which insights are actually relevant to the practice, whether in pharma or beyond, is still the cornerstone of success.

Sergiy Sumnikov

HOST

I agree that this evidence-generation discipline is very strong at pharma companies, and I see pharma as very data-mature compared to health systems. But from my observation, pharma has less experience with frontline behavior change than health systems do. Would you agree?

Kimberly Noel

GUEST

Historically that was probably true, but I think the new frontier isn't the bulk of the data, it's the insights you can generate from it. We're seeing rapid adoption in clinical practice, notably with ambient scribing, AI diagnostic support, and clinical decision support, so there are learnings that flow both ways. Making sense of these vast amounts of data is the real unlock for what AI can do for patient outcomes. Where healthcare may actually be faster than pharma is in monitoring health outcomes directly. As we move toward patient-centered care, pharma is also adapting to new realities, perhaps more evidence burdens and new responsibilities toward patients, which aligns it more closely with the health systems.

Sergiy Sumnikov

HOST

Very interesting. And why do AI projects fail despite promising technology, from what you've seen?

Kimberly Noel

GUEST

Most people think AI adoption is a technology challenge when, in fact, most of the time it's a human workflow challenge. That's what inspired me to start Human in the Loop. It's about co-creation,

governance, defining the workflow clearly, and getting the right data to provide the right insight at the right time. There's a lot of pressure across industries for rapid adoption, this sense that we must do AI. But in healthcare, especially given the high-risk provisions that AI adoption brings, we need to think clearly about what is actually a true benefit.

A lot of early adoption has been toward automation, which isn't inappropriate. Refilling prescriptions and things like that should perhaps be automated. But we can learn from these early adoptions, notably ambient scribing, computer vision, and decision support, to build better workflows. One example: when you automate an activity, it should really be re-stratified first. Low-risk automation requires less cognitive burden, and treating low risk and high risk as the same is a big recipe for burnout, because you either flood people with alerts and create cognitive fatigue, or you miss things. What's missing at the center of this is respect for clinician attention. When you do have the clinician's attention, it should be spent at the right point in the workflow, at the right level of risk. So the big factor is that this is a workflow issue, not just a technology one, and a lot of the gains aren't about AI at all. They come from creating clarity around the workflow so people can see the true benefit of the design.

Most people think AI adoption is a technology challenge when, in fact, most of the time it's a human workflow challenge.

KIMBERLY NOEL

**Sergiy
Sumnikov**
HOST

That's exactly what I've seen: AI projects often have a research budget and an engineering budget, but not a workflow budget, and that's where adoption actually happens, where clinicians interact with the tool. It's also where we come in, and it ties to what you said about output arriving in the right place at the right moment, which is partly a UX/UI question. My advice from the vendor side is that a new technology shouldn't be evaluated only on the demo of the happy path, because clinicians will live in the edge cases too. That product mindset has to be there. You've already touched on the organizational side. What does successful clinician co-creation actually look like in practice?

Clinician co-creation starts with the design, and with truly understanding the pain point. If you haven't clearly defined the pain point and how you'll measure it before you start implementing AI, you can just exacerbate a broken workflow. So it's important from the very beginning to have clinicians engaged in identifying the right pain point. Once you have it, you develop a clear path for what gets automated and who owns responsibility for each step, and you define which human is in the loop at each step, because sometimes it's the ordering clinician, sometimes the CMIO, sometimes the chief medical officer, sometimes the vendor. It's ideally not just clinicians either, but patients, ethicists, and legal on board, so you can move through the process seamlessly.

Then, once you've defined how to measure improvement against a baseline, you make sure you have an iteration and reiteration plan, and that the loop is visible quickly. Even for something simple like automating refills, you want to think through who owns the initiation, who owns the review, what the refill rate is at baseline, and what it becomes after automation. It doesn't have to be complex or burdensome, but it has to be thought through. And developing governance that includes clinicians matters, because they're usually the first to understand why something isn't working. That's the shift: moving past the clinician as end user and treating the clinician as a key success partner in the journey toward modernization. We've seen great results. Look at the research Kaiser Permanente published on the savings they achieved with ambient scribing, or the rapid adoption of tools like OpenEvidence for clinical decision support. These are suites of tools. AI isn't one uniform technology; it's a suite of tools that can be implemented successfully with the right clinician partnership.

From our perspective, clinician co-creation is the foundation. Before we even write a specification for a UX/UI redesign, we shadow clinicians and do journey mapping. It's also important to put clickable prototypes in front of clinicians as fast as we can, because they won't comment on a requirements document, and with AI it's now easy to get to a clickable prototype quickly and change it. With Seattle Children's Hospital, for example, they gave us feedback this week and we had a new version of the prototype for them to react to next week. We even try to make the changelogs visible to clinicians, so they see the direct influence of their feedback: this was changed, this will be changed. They feel in control, which matters, because they're the users.

02

Human in the Loop & Cognitive Load

Why automation alone doesn't reduce load, what success beyond compliance looks like, and why you have to measure before and after.

Sergiy Sumnikov
HOST

We've touched on human in the loop and on cognitive load. Let's go deeper. Could you briefly outline the key human-in-the-loop principles, and how they relate to a clinician's cognitive load? My sense is that the relationship isn't black and white.

Kimberly Noel
GUEST

This is my favorite topic, clearly, having founded Human in the Loop. The core principles are essential for responsible AI: compliance and accountability, responsibility, iteration and reiteration, which is that feedback loop you described with Seattle Children's, plus clinical validation, safety, and risk mitigation, all with transparency and explainability. These principles alone won't reduce cognitive load or burnout; that's the balance. Load reduction has to come with design. It's not a given that AI implementation returns value on its own. Good workflows that capture clinician attention appropriately are what reduce cognitive load. Everyone knows about alert fatigue, and there's nothing more frustrating than being bombarded with alerts that aren't relevant to the clinical scenario you're in. The iteration and clarity these principles demand are what let workflows keep improving rather than going stale.

The key to real reduction of cognitive load, and hopefully of burnout, is turning data into insights that are actually relevant. That's where large language models hold a lot of promise, in amassing the right data to build a clinical narrative that makes sense. But we're not yet at the point where they replace human judgment. This isn't how we've historically integrated technology into healthcare. Think of EMR adoption: it happened to clinicians and created a crushing administrative burden. With AI we have a chance to do it differently, to rethink the dynamic and adopt these principles as a team of experts to improve patient care.

The last thing I'd say about human in the loop is that the human who matters most in the process is the patient. Not also the patient, the patient. So when we implement, we have to assess how it ultimately improves care for the patient. That's where humane implementation of AI, where clinicians help shape it and are accountable partners rather than just human proofreaders of automation, lets them orchestrate with the patient in mind and produces long-lasting results.

The human who matters most in the process is the patient.

KIMBERLY NOEL

**Sergiy
Sumnikov**
HOST

From the product side I'd add that every AI product has limitations, and those limitations shouldn't be buried in a PDF, they should be shown in the product, at the moment of use. That builds trust, and it supports the human-in-the-loop principles. Where there's uncertainty in an output, we try to give an evidence trail, so it isn't just a confidence level of 78% that a clinician can't interpret, but something visualized and clear, integrated into the product at the point of use. So, back to your point that automation doesn't automatically reduce cognitive load, especially when the clinician remains responsible for AI outputs. Could you elaborate?

Kimberly Noel
GUEST

Of course. There's an interesting study from an MIT lab on cognitive debt. They looked at the use of large language models like ChatGPT on brain activity, using EEGs, and they found that with more repeated use of the model there was a kind of cognitive deficit. People weren't thinking as hard. We can't afford that in medicine. You want clinicians at the top of their expertise, really using their cognitive abilities. So we have to think through what humane implementation means, and confront the clinician-as-proofreader challenge that is the new frontier. We need humans responsible at the key checkpoints where decisions are irreversible.

The responsibility a clinician takes on, to do no harm, is about accepting the risk of the art of medicine and making the best decisions with the information available. Clinicians are used to making evaluations under uncertainty, and adding AI is another dynamic that requires the same preparation, which means thinking through risk mitigation. As you reminded us with your product point, many of these models are probabilistic and Bayesian, so you never have 100% applicability to the exact situation in front of you. Clinicians have to think critically about a powerful tool that can accentuate diagnostic capability or surface information they might not have, while staying on guard to change a decision if the new information calls for it.

So again, I distinguish data from insights. Ideally, AI moves you toward insights that matter to the clinician, insights within which they can understand the uncertainty they're navigating, which is exactly what they're trained to do. Navigating uncertainty to arrive at a diagnosis and an optimal treatment is the commitment clinicians make to their patients. AI is one more tool in that toolkit, and it has to be adopted with the right workflows, risk mitigation, training, and surveillance.

Sergiy Sumnikov

HOST

From the product side, sometimes we put evidence first and the conclusion later, so clinicians don't just agree by default with something that pops up. On cognitive load, we had a specific project for the ICU at a children's hospital in the US, where they had built an AI tool internally on top of Epic. It was smart and genuinely helpful, but it displayed everything they measured at once, a huge amount of time-series data. Picture a doctor at 3am, with all the sounds of the ICU, trying to comprehend that one big monitor. We redesigned it using progressive disclosure, so you first see only what you need, and for the roughly 20% of cases that require it you can go deeper. That immediately improved adoption. The tool was built by data scientists, who think in their own way, but doctors often think in body systems, so we also used visualization: it was for a heart-center ICU, so we showed the heart schematically and indicated the parts where attention was needed. That completely changed adoption. ICU is number one for cognitive load and criticality, and it's directly tied to alarm fatigue, but the same ideas apply widely, even to administrative work. So, how would you define human-in-the-loop success beyond compliance?

Kimberly Noel

GUEST

Compliance is the floor. It's the license to operate. Beyond compliance, we should be measuring elevation, real improvement in healthcare: time to diagnosis, time to decision, and the metrics coming out of the ambient-scribing studies, like documentation minutes, cognitive-load scores, burnout, and retention. And the other part is bringing it back to the patient. We can see that Kaiser saved something like 16,000 hours, but it would be powerful to also design studies that capture patient-encounter satisfaction, so the narrative is richer, that adding a tool like ambient scribing improved the doctor-patient relationship or the time spent with the patient. Those are design principles that have to be governed and protected. Saving time means nothing if you don't protect the time, because then you don't get the meaningful, extended relationship in the encounter.

So getting the metrics right is one step toward success beyond compliance. The other is holding the duality of clinician and patient experience in parallel, pairing patient experience with clinician satisfaction. And from a product perspective, maybe you'll comment further, it's about that tight loop of improvement: how fast does clinician feedback become a product change? When that's slow, as it historically has been in healthcare, that's when you get lasting failure, a rise in distrust, and, importantly, compliance risk. That's the other reason clinicians being frontline matters. The whole clinical team knows when something isn't working, and being able to say early, this isn't working and here's why, so you can iterate quickly, is very valuable.

Sergiy Sumnikov

HOST

I want to stress what you said about measuring before and after, because in theory everyone knows it, but in practice we've seen many cases where it wasn't done. On cognitive-load projects we use the NASA Task Load Index before and after a redesign and measure the difference. It's subjective, but it's still an important metric, capturing mental demand, physical demand, temporal demand, performance, effort, and frustration. Sometimes after a redesign some dimensions stay the same and some improve, and sometimes some improve while others get worse, so you have to read it holistically. Please, go ahead.

Kimberly Noel

GUEST

It really matters who the stakeholder is in the decision, and how the insights will help them make future decisions for the practice. Sometimes, if the main decision maker is in finance, they may not need all of those dimensions to have enough data to fund a further iteration. So to be holistic, we capture the clinician and patient experience, but we also have to build trust among the decision makers at each level, asking what will actually move the needle. Then it isn't research for research's sake, it's a clear narrative about what we're trying to accomplish and what success looks like. And that baseline is not a given. A lot of people can't define the status quo, and if you can't define the status quo, you can't define improvement.

If you can't define the status quo, you can't define improvement.

KIMBERLY NOEL

Sergiy Sumnikov
HOST

I agree that clinical leaders belong on the decision committee for this. In the ICU case I described, the clinical leader was the head of the heart center and also the budget holder, and he cared a lot about clinician well-being, burnout, and performance. In the ICU, cognitive load is acute and directly tied to alarm fatigue, which is a big topic everywhere. That's why this angle was the main one for us. For a finance leader, though, the time saved with AI scribes might be the more relevant business case.

Kimberly Noel
GUEST

My point is more that we need to think about who is ultimately tasked with making the final decisions in clinical care. If you consider that early enough, you can be thoughtful about what data you collect. It shouldn't be an afterthought that the person who holds the budget gets engaged. A big part of being effective in AI implementation is capturing the insights that lead to improvement, and having consensus on what improvement is. And short of full consensus, because people have competing priorities, having at least enough data to substantiate change in the right direction, toward a better clinical experience.

03

Automation vs Control

What should be automated versus kept under clinician control, when a workflow is ready for AI, and how to avoid over-reliance.

Sergiy Sumnikov
HOST

Let's move to the next block, which we can roughly call automation versus control. How do you determine what should be automated versus what stays under clinician control?

Kimberly Noel
GUEST

I have a rubric I like: risk and relationship. Things that are low risk don't need a doctor to hold my hand. I don't need one to refill my prescription or get a routine form for my child's school; there you just want efficiency, and automation through AI works great. Automation also excels at scale, at assessing multiple sources of data and triaging them. Where we really need the relationship is through complexity and high risk. I know very few people, if any, who would prefer a cancer diagnosis from a bot rather than a human. That relational aspect of care is so important. The reality is that a big part of clinical practice sits in between, and that's the art of medicine.

There are different affinities toward that. There's a group studying AI affinity scores, recognizing that different people have different

emotional needs and different views on how technology should show up in their care. So we may be at a frontier of personalization that isn't only about someone's biology, but also about their behaviors and how they want to use the healthcare system. We have a lot to learn from consumer AI adoption there. I think this is especially important for chronic care, where people have to sustain behaviors and adherence over a long period. So personalization that enables automation and makes patients feel they're getting better care is a promising area.

Sergiy Sumnikov

HOST

At the last HLTH in Amsterdam there was a session on what health-tech product leaders can learn from consumer-app product leaders. One point I remember is that framing matters: it should read as a recommendation rather than telling someone they did something wrong. And different people expect different personalization and interaction styles. How do you know when a workflow is ready for AI, or whether it needs a redesign first?

Kimberly Noel

GUEST

The people on my team would laugh, because I'm always about the baseline. You have to have a baseline to understand the workflow, and that sounds simple, but gathering data on the current state of affairs is itself a task. Then, again, defining who's responsible and what the goal is. These sound like self-evident mantras, but they're make-or-break once you start including automation at scale. So start with the pain point you're really trying to solve, the metric that defines it, the data you have to collect and turn into insight, and for which stakeholder. Then stress-test it and see how you'd mitigate the risk if it breaks. Think about all of that in advance, and you're ready, because the implementation takes far less time than cleaning up after a failed one. It feels slow at first, but it's a way to go slow to go fast, and to create lasting improvement.

Sergiy Sumnikov

HOST

I agree. For us as a product partner, there's one marker for whether we should redesign a process before automating it: whether we can draw the current workflow as a service blueprint on the wall and get the people who actually do it to confirm that yes, this is how it works. If we can't, we have a process problem first. Sometimes a customer runs the same workflow differently across sites and can't articulate why, and then there's a lot of workflow redesign to do before any automation. So, how do you balance AI autonomy with human accountability without creating excessive friction?

Kimberly Noel

GUEST

You need checkpoints. We have a lot to learn from other industries, like aviation, in terms of checklists and checkpoints, and a lot of that has already been adopted in surgery. We have methods for identifying high-risk checkpoints in a workflow, and successful AI mirrors those checkpoints. You shouldn't add extra steps to what is already the natural rhythm of clinical care, because we train clinicians to follow that rhythm. There are moments, like gathering data and forming a diagnosis, when you want the clinician not to be interrupted.

But in the end it's humans who are accountable for any improper outcome for the patient, so the question is which humans. With clinical decision support, it should be clear that the clinician is using it as validation, not as a replacement for critical thinking. A cognitive-first approach, thinking through the diagnosis uninterrupted and then validating it with decision support, is an appropriate workflow that doesn't add friction. Each AI tool being implemented requires that kind of thought: what's the clinical rhythm, what information is introduced, for which insight, for which behavior. Take that approach and you reduce friction. But there's also appropriate friction. When an outcome is irreversible and high-impact, you do need a checkpoint, and that's the moment you want clinicians highly engaged, taking responsibility for the patient's outcome. Being explicitly clear about when that matters for us as humans interacting with AI.

*You shouldn't add extra steps to what is already
the natural rhythm of clinical care.*

KIMBERLY NOEL

**Sergiy
Sumnikov**

HOST

In our product work we often create friction exactly where it's needed, especially where irreversible consequences can occur. In consumer life friction shows up when you try to cancel a subscription, but here the goal is far more meaningful. And how do we prevent over-reliance on AI by clinicians?

Kimberly Noel

GUEST

This is an important topic for me, because I care about the future generations of clinicians entering this era of hybrid care with AI. My hope is that they can be at the top of their cognitive skills, pushing medicine to new frontiers, making diagnosis more equitable across

the world, and enabling new discoveries. For that to be true, humans have to feel valued in the system where they practice. The skills clinicians will need include digital literacy and understanding the Bayesian, probabilistic nature of these models and how it differs from the art of medicine, and how to incorporate new data into decision making. These skills have to be taught.

We also have to admit that the way we've taught medicine may be archaic. I was reading a discussion, I think in the New York Times, about the clinical note. I trained, like generations of doctors, by writing the clinical note, and writing it is what synthesizes our understanding of what's happening. It's an observation of the clinical space that you then synthesize into a summary. So with AI now able to write the note, we have to think through what that means for that process of cognitive thinking, and whether the act of writing is itself the cognitive reinforcement we need. I'd argue doctors should keep writing those first notes before they reach for AI. These are open questions that experts in education need to tackle, so the clinicians of the future keep the relevant skills to operate in this hybrid environment and don't lose them.

**Sergiy
Sumnikov**
HOST

In a narrower sense, training clinicians for a specific AI product isn't only about learning its capabilities, but also its limitations and possible failures, because when a vendor comes in it all sounds like a sales pitch. Providers and pharma companies should make sure there's enough training material on what can go wrong. On the product side we implement and measure many metrics, and one is the acceptance rate, or its inverse, the override rate. If your acceptance rate climbs toward 100% and clinicians never adjust the results, that flags a problem, and it should be on a dashboard for the governance committee. Sometimes the problem is on the product side, not the clinician side, because overriding an output has to be easy and clear. If it's hard to change what the AI produced, that creates unnecessary friction and can push people toward a default of just accepting. I try to monitor my own over-reliance too, in both my personal and professional life. We're almost out of time, but I'd love to cover the next block, which I really like: building trust. How do you build trust in AI among clinicians?

04

Building Trust

Earning clinician trust through evidence and validation, and monitoring an AI system after deployment.

Kimberly Noel

GUEST

Trust doesn't have to be slow, but it has to be earned. As we think about the clinicians of the future, we should be teaching clinical validation skills and how to evaluate AI in healthcare. That's one reason I joined the editorial board at Nature's npj Digital Medicine, to look at the rigor of evidence around AI algorithms. We need to be systematic about that, because evidence is the foundation of trust. It's natural for us to distrust as humans; you don't want clinicians blindly believing a technology will improve care. So it's the same principles of validation, peer review, and scientific rigor. And it can move fast, as we've seen with ambient scribing and clinical decision support. The flywheel should run in cycles: a healthy distrust, an evaluation of the data in front of you, meaningful adoption, then validating and evaluating the outcome, risk mitigation if it fails, and then fixing. Coming from pharma, where there's a lot of regulation, we should have the same vigilance over algorithms as over other treatments. Really thinking through what surveillance and clinical validation look like, and how to create humane monitors so frontline clinicians pay attention when it truly matters rather than just surveying automation, these are foundational to building trust.

Trust doesn't have to be slow, but it has to be earned.

KIMBERLY NOEL

Sergiy Sumnikov

HOST

I appreciate how much focus you've put on evidence from the very start of our conversation, so this clearly feels close to you. I agree that evidence comes first. In practice, though, we also see that the first 90 seconds of using a new app shape a lot of whether a clinician trusts it: are the interfaces appropriate, is it clear, is it even responsive enough. If it's messy or uncomfortable to use, that erodes trust even when the evidence is all there. So I'd highlight the interface as part of trust too. And a very last question, Kimberly: how should organizations perform continuous monitoring after an AI deployment?

Kimberly Noel

GUEST

As we said, you start with a clear baseline and the metrics that matter most to the decision makers, and that delta should be personalized to the exact technology being implemented. Understand that wins are real but sometimes modest, and that

iteration is part of the process. You should anticipate things breaking and have proper contingency plans; that's what creates strength for organizations. And the last piece is respecting clinicians and their attention. With AI there's a lot of fear and distrust, so engaging clinicians as partners and as experts in their domain is really important for aligning to the bigger goal, which is efficiency and improved patient outcomes. That's done through governance, protection of clinician time, and proper user-experience design. For systems that truly excel at AI, none of that is an afterthought.

End Closing

A short thank-you to close the webcast.

Sergiy Sumnikov

HOST

Thank you, Kimberly. I really enjoyed our conversation, and I have a feeling we could talk much longer, though the limits on our time probably make us more efficient and meaningful in what we do. Thank you very much for joining the webcast today. It was great to have you and to learn from your insights.

Kimberly Noel

GUEST

Thank you.

Sergiy Sumnikov

HOST

Thank you. Have a good day.

End of transcript. Recorded 4 August 2026 · Halo Lab · halo-lab.com