

The Good Bot: Artificial Intelligence, Health Care, and the Law — Navigating FDA's 2025 AI Guidance: Risk-Based Framework, Public Comments, and Generative Models

Host: Brett Mason

Guest: Kyle Dolinsky

Recorded: August 22, 2025

Aired: December 12, 2025

Brett Mason:

Welcome to *The Good Bot*, a podcast focusing on the intersection of artificial intelligence, healthcare, and the law. I'm Brett Mason, your host. As a trial lawyer and partner at Troutman Pepper Locke, my primary focus is on litigating and trying cases for life sciences and healthcare companies. However, as a self-proclaimed tech enthusiast, I'm also deeply fascinated by the role of technology in advancing the healthcare industry. Our mission with this podcast is to equip you with a comprehensive understanding of artificial intelligence technology, its current and potential future applications in healthcare, and the legal implications of integrating this technology into the health care sector.

I'm excited to welcome back to the podcast today, one of my colleagues and fellow partners, Kyle Dolinsky. Thanks for joining us, Kyle.

Kyle Dolinsky:

Thanks for having me.

Brett Mason:

Kyle, your specialty really focuses on FDA regulation. It's one of your specialties, I should say. We know that the FDA has been doing a lot around AI, especially when it comes to use of AI in the development of drugs and biologics. Wanted to have you go ahead and chat with us. What are the updates? Have there been any new guidances from the FDA in 2025 on that front that folks who are involved with drug and biologic development should be aware of and should be taking into consideration?

Kyle Dolinsky:

Thanks, Brett. Today, I really want to talk about some draft guidance that the FDA issued back in January this year. A bit of a while ago now. But we've got some more information that's come out, and some things to consider based on what's happened at FDA in the last couple of months. That draft guidance was on considerations for the use of artificial intelligence to support regulatory decision-making for drug and biological products. Like I said, it came out in January of this year. That guidance notes that the use of AI to produce data and information about the safety, the effectiveness, quality of drug and biologic production, it's really increased exponentially since 2016, which is the last time the FDA really took on this piece of it in

guidance. Because of that, FDA wanted to update and provide some guidance on how to incorporate AI into these processes.

The important thing is that this covers, really it parallels what the FDA regulates. It covers pre-clinical development, clinical development, post-marketing and manufacturing phases of the development cycle. The important things are that it does not cover drug discovery and internal workflow-type efficiencies.

Brett Mason:

Let me just pause there for a second for our listeners who may not be as familiar with how the FDA issues these guidances. They issued the draft guidance in January of this year. Is there a timeline for the final guidance to come out and what is the process to get to that final guidance?

Kyle Dolinsky:

FDA issues draft guidances and generally, excepting emergency situations, they will open it up for public comment. Then, FDA usually receives a good amount of comments and they'll incorporate that feedback, make some revisions. Generally, they tend to be not too extensive, but sometimes we've seen some fairly extensive revisions from the FDA. Then FDA will issue the final guidance. That timeline really varies. I mean, there have been draft guidances that have been on the books for years before the final comes out. Sometimes these guidances only stay in draft form. While they're not technically a finalized guidance, they do still provide guidance for the industry. In the absence of other information out there, we look to the guidance, even if it's a draft guidance, to give a sense of what FDA is thinking.

In this case, there was a comment period that ended back in April, but FDA has extended that period. There's no real deadline right now, but they've said that they're accepting late comments, so we've seen those continue to fill in and we can talk about those a little bit.

Brett Mason:

Yeah. Before we talk about the comments, let's just talk about the draft guidance itself. My understanding is that FDA is recommending a risk-based framework for companies to consider the use of AI that has seven different steps. Can you tell us about those steps and give us an explanation of the FDA's framework there?

Kyle Dolinsky:

The whole reason the FDA is doing this is they want to make sure that there are well-defined AI models that are credible and turning out reliable data and are being used in appropriate ways to help in the drug development activities that companies are engaging in. The seven steps are the following. The first one is to define the question of interest that AI will address and seems pretty basic, but you need to know the question that you want to answer.

The next one is defining the context of use. This is an important one. This is the role and the scope of the model. What is actually going to be modeled? How are the model outputs going to

be used? Importantly, is the model output going to be used in conjunction with other information to answer the question? Or is it going to be the bottom line? That's really going to inform the level of risk.

The next step is going to be assessing the risk associated with the model. You think of this in terms of two axes. On the one hand, you have the level of model influence. On the other, you have the consequences of the decision. Where the model has a lot of influence in the decision-making, it's going to be higher risk and also, where the consequences of that decision are specifically important, it's going to have a higher risk level. Depending on the risk level, the more rigorous you need to be in terms of ensuring the credibility of the model.

After that, you need to develop a plan to establish the credibility of the model within the context of use. This is the heart of the guidance, goes into lots of sub steps here, but it's really the important things that you would do when you're developing a model. It's the model design, talking about the data that you're using to develop a model and train it, model training, model evaluation. Then after that, you need to execute the plan and document its results.

Finally, at the end, you are determining the adequacy of the AI model for the context of use. If you determine that it's not adequate for that context of use, you have a few options. You can drop the AI model or make substantial revisions to it. You can downgrade the influence that model is having on your decision making. You can increase the rigor. You can establish certain risk mitigation controls. FDA is not saying that if this doesn't turn out exactly like you think, you need to abandon it, but you do need to recalibrate. In some cases, you need to scrap the model, or make serious changes to it.

Brett Mason:

Kyle, let me jump in there, now that we've talked about a couple of the steps that the FDA set forward in that kind of risk framework. One of the things that I found really interesting was that through those first couple of steps, the FDA is inviting a lot of interaction from companies to get feedback from the FDA about the AI model risk and about the context of use. Can you talk a little bit about that early engagement that FDA is recommending to ensure the appropriate credibility assessment activities are going on?

Kyle Dolinsky:

Yeah. Part of the guidance here is that you should engage early and often with the FDA to really make sure that everybody's on the same page in terms of defining the context of use and defining the utility of these models and understanding what they're going to be used for. Think about your submissions, especially an NDA, or a BLA if it's biologic. You need to provide everything and anything for the FDA. If AI is informing your decision making, you need to include those details there. Like other issues that are going to be major pressure points with the FDA, you want to have communication and have a sense going in that FDA is on the same page. Otherwise, you're investing a lot of money, not just into your AI model, but also into the rest of the development process that might turn out not in your favor.

Brett Mason:

What are the final couple of steps that the FDA sets forth in the guidance?

Kyle Dolinsky:

Like I said, once you've developed that plan, the rest flow out quickly. But it's executing the plan, it's documenting the result of that plan, and again, determining the adequacy is that final step, which can either be "this works", "it doesn't work", or somewhere in between where you're making changes.

Brett Mason:

Let's talk a little bit about the comments that have come in. Have you taken a look at those?

Kyle Dolinsky:

Yeah. Like I mentioned, the comment period officially ended back in April, but the FDA is allowing late comments, and they'd still come in. As of yesterday, at least, there were 113 on the docket. They've come from everybody. I mean, they've come from individuals, they've come from startups, they've come from midsize pharma companies, major large-scale pharma companies, industry groups, like pharma, and patient and provider groups, too. They run the gamut. There are comments on specific pieces of language. There are overarching comments. But there's three general themes that I wanted to mention, because they were recurring ideas that were coming through in the comments.

One which isn't really a surprise and something that we tend to see all the time in comments in the FDA guidances is a request for more specific examples and scenarios. The draft guidance gives some examples, but obviously, the more examples that industry has, the more it can rely on and be confident that it is going to be complying with FDA's regulations and with FDA's enforcement policies. There's usually a balance there at some point. You can have too many examples, to the point where the guidance doesn't become useful as broad advice. We might expect FDA to add some more examples and hopefully, geared towards the next bucket of comments I'm going to talk about. But I wouldn't expect to see just pages and pages of additional examples.

Brett Mason:

Let me ask you about that, this desire for examples. Do we think that in this area of AI and how AI can be used, should be used, and how FDA is viewing it, might be difficult for them to give examples, because there's so many new things and things changing so rapidly?

Kyle Dolinsky:

Yeah, I think that's probably true. At a certain point, the examples stop becoming useful, because to be a useful example, you want to be specific. But if you're being specific, you're excluding a universe of other possibilities. The other issue, too, is that I think and it's something

that we see all the time, law lags technology, and here technology is moving extremely quickly. It's something that we saw in the guidance and resulting in the comments, which lends itself to this next bucket of comments that I wanted to mention, which is a lot of the commenters noted the lack of commentary on generative AI and foundation models.

The guidance talks about AI very generally and it talks about machine learning very generally, but it doesn't provide specific guidance on generative AI and foundation models, which are the heart of a lot of what's going on with AI these days. Companies want examples and guidance that are tailored to that, and thinking we don't know the exact timeline of drafting this draft guidance, but you have to imagine that it took place over a significant amount of time. In that time, things have already changed and they've already changed in the several months, since the draft guidance was issued.

Brett Mason:

What are some of the other comments that are being submitted by stakeholders?

Kyle Dolinsky:

The last bucket, and I think it's a really important one that the guidance really doesn't take into account, but practically speaking, is going to be a huge deal is there's lack of clarity about third-party supplied models. If you think about the real giants in the industry, it's likely that they have proprietary internal models that they're using in many cases. But certainly, not even in all cases. If you think about the startups and the midsize pharma companies, a lot of times they are bringing in a product from a third party, and a lot of the development and feedback process that FDA is expecting in its draft guidance is difficult if not impossible when you're dealing with somebody else's product. You might not be training it, or you might not be training it at the level that FDA is expecting, and you might not have access to make the kinds of tweaks that FDA is hoping that you can make. That's a piece of it.

Then another piece of it is in your communications with the FDA a lot of times, obviously, these models have a lot of confidential proprietary trade secret type information in them. There's going to be challenges in terms of being clear and open with the FDA and having to navigate those kinds of confidentiality and IP issues.

Brett Mason:

That makes sense, because at the end of the day, the FDA is regulating really the company that's putting out the drug, or biologic, it's not regulating the third-party vendor that's providing the AI software, or the AI program that they want to use. That's definitely between a rock and a hard place for companies who want to take advantage of that technology, but at the end of the day, have to answer questions on technology they may not have the answers to.

Kyle Dolinsky:

Right. Like you mentioned, the sponsor, or the applicant here has the regulatory responsibility to the FDA. They're taking on responsibility for this model that isn't theirs and that they might have had very little stake in developing.

Brett Mason:

We can see so many efficiencies that could come from drug and biological development, including these types of technologies and think about how much faster could be to create a new drug application using generative AI that can review the data, draft, and then just be overseen, or reviewed and revised by humans. It seems like there's still a lot of questions that have to be answered on one, how that's going to work from a practical standpoint, and then two, how's the FDA going to oversee that and ensure that the use of that technology is providing safe and effective drugs and biologics?

Kyle Dolinsky:

Yeah, I think that's true. Those challenges are the challenges that we're seeing with the use of AI in all industries across the board. It's just having to balance the incredible utility of this and understanding that, especially with new models, they're not perfect and we need to be aware and vigilant to make sure that the information that's being produced is accurate, and especially in this context is leading to the development of safe and effective products.

Brett Mason:

Kyle, one last question before we finish up. Clearly, the guidance was issued back in January of 2025 prior to a change in administration. Given the different perspectives on the administrations around artificial intelligence, do we expect that to result in further changes to the FDA guidance?

Kyle Dolinsky:

It's possible that we'll see some changes when the guidance is finalized to put this administration's stamp on it. But given the new administration's focus on AI and using AI going forward, I wouldn't expect, for example, the draft guidance be scrapped, or there to be wholesale changes that are based on the change in administration, but it's something to keep an eye on in terms of minor changes that we might see in their approach.

Brett Mason:

Kyle, the way that I read the guidance and I'd be interested in your thoughts as well is not that the FDA is trying to limit, or prohibit companies from using artificial intelligence for these types of activities, but rather, trying to provide a framework within which the companies can use AI. Given the current administration's support of use of AI, I wouldn't think that they would be too troubled by this guidance. Again, it's not limiting, or prohibiting the use of AI. What are your thoughts on that?

Kyle Dolinsky:

Yeah. I think the draft guidance makes clear that not only do they support the use of AI in drug development going forward, but that it's going to be inevitable. Really, this is about making sure everybody's on the same page and making sure that AI models that are being used in drug development are reliable and credible and that the information that's coming out of them is going to be good and useful. I wouldn't expect that to change from one administration to the next.

Brett Mason:

Well, Kyle, thanks so much for being on and chatting with us about this draft guidance and the ongoing comments that the FDA is receiving about the guidance. We'll be interested to hear more from you once the guidance is finalized, and appreciate all of your expertise in this area.

Kyle Dolinsky:

Thanks, Brett. Happy to be on.

Brett Mason:

Thanks to our listeners. Please, don't hesitate to reach out to me at brett.mason@troutman.com if you have any questions, or comments, or topic suggestions. You can also subscribe and listen to other Troutman Pepper Locke podcasts wherever you listen to podcasts, including Apple, Google and Spotify. Until next time.

Copyright, Troutman Pepper Locke LLP. These recorded materials are designed for educational purposes only. This podcast is not legal advice and does not create an attorney-client relationship. The views and opinions expressed in this podcast are solely those of the individual participants. Troutman does not make any representations or warranties, express or implied, regarding the contents of this podcast. Information on previous case results does not guarantee a similar future result. Users of this podcast may save and use the podcast only for personal or other non-commercial, educational purposes. No other use, including, without limitation, reproduction, retransmission or editing of this podcast may be made without the prior written permission of Troutman Pepper Locke. If you have any questions, please contact us at troutman.com.