

Changing Higher Ed Podcast 145 with Host Dr. Drumm McNaughton and Guest Michael Feldstein

How Machine Learning and AI Can Benefit Higher Ed

AI, ML, ChatGPT, and Higher Ed

March 7, 2023

Welcome to Changing Higher Ed, a podcast dedicated to helping higher education leaders improve their institutions, with your host, Dr. Drumm McNaughton, CEO of the Change Leader, a consultancy that helps higher ed leaders holistically transform their institutions. Learn more at changinghighered.com. And now, here's your host, Drumm McNaughton.

Drumm McNaughton 00:31

Thank you, David. Our guest today is Michael Feldstein, chief accountability officer at e-Literate. Michael has a strong background in both k-12 and post-secondary education, including being a partner at MindWires Consulting, senior program manager of MindTap at Cengage Learning, and principal product strategy manager for Academic Enterprise Solutions. Michael's latest passion is higher ed, especially technology in higher ed. He joins us today to talk about artificial intelligence and machine learning and how they can help and hurt higher ed institutions.

Michael, welcome to the program.

Michael Feldstein 01:13

Thank you, Drumm. It's a pleasure to be here today.

Drumm McNaughton 01:15

I'm looking forward to our conversation. First, we will go into that big, bad, ugly topic called artificial intelligence, which drives higher education institutions. It's crazy, at least from what you read in publications, because they're concerned about many things. We will get to the truth today by exploring what's important and what folks need to be aware of. But before we get there, tell us a little bit about yourself. You're an expert in many areas around AI and machine learning. What's your background? How did you get here?

Michael Feldstein 01:52

Well, I'm not an expert in any area. I'm a jack of all trades and master of none. I am a teacher from a family of teachers. Even though I haven't been in the classroom in decades, I always approach everything from an educator's perspective. I have a love of technology and a particular passion for thinking through the process of teaching, which I grew up talking about around our kitchen table and thinking through how new technologies can help the mission and the education process. So, my interests and knowledge of different technologies, including artificial intelligence, have constantly evolved as new opportunities arise.

But I call myself a passionate amateur in artificial intelligence and more of a professional educator who can translate between the two groups.

Drumm McNaughton 02:57

That was well-spoken. Thank you. We got to know each other through [Phil Hill, an OPM expert](#). My perspective is that if Phil recommended you, you're the man.

Michael Feldstein 03:10

Phil and I have blogged together for eight or nine years and consulted for seven. He's an old friend and a brilliant guy.

Drumm McNaughton 03:22

Let's get into the topic today. Artificial Intelligence or AI. What is AI? I mean, we've read so much in the news lately. What is it? Explain it to us in layman's terms for those who look at a computer and go, "Oh, I see what it's doing, but I have no idea how it works."

Michael Feldstein 03:39

Right. Without getting into the technical details, we're interested in two families of algorithms. One is called [machine learning or ML, which is a cousin of artificial intelligence, or AI](#). We won't make too much of a distinction between the two today because that wouldn't serve the audience and can send us down some rabbit holes.

But the essential commonality between the two is that there are algorithms that we design on how they can evolve themselves without specific instruction. For example, back in the days when your spam filter wasn't run by Google or your email service, you had a plugin that would learn how to get better at recognizing emails that you would consider spam or not spam, even though you didn't tell it explicitly what about the email might allow it to recognize that the email was spam. You just said this is spam, or it isn't.

Drumm McNaughton 04:48

So, what you're talking about is when I'll get an email about how I have \$250,000 approved for financing, etc., I'll mark this as "junk mail," and it goes in there. Is that what you're talking about?

Michael Feldstein 05:04

Exactly. And it gets it wrong sometimes, but it gets better over time without us telling it, "Well, I noticed that this came from a weird email address, even though it says 'Bank of America' on it." The spam filter knows that you've scored this as spam, even though it looks like it's not spam to the spam filter. So, it will go through all the different parameters, figure out which aspects make this email look suspicious, and adjust it so that the next time a similarly suspicious-looking email comes in, it'll be more likely to flag it as spam.

Drumm McNaughton 05:54

So, that's machine learning, AI, or both?

Michael Feldstein 05:59

It's both. Again, I'm not a technical expert on AI. But I know well enough to tell you that that is a feature common between the two. AI often adds layers of complexity in terms of how it teaches itself and how the software is written to evolve, which machine learning does not have. An example of this we'll get into later when we talk about ChatGPT is "type ahead" on your phone. You might start typing a word, and it will guess the word you want to complete. Or, your word processor might guess the phrase or sentence you're trying to complete. That is all machine learning and artificial intelligence, too. It can get quite complex and interesting when you bridge the gap between machine learning and artificial intelligence.

Drumm McNaughton 06:57

That makes perfect sense because the machine or the intelligence has to learn based on what it's observed without being given specific directions. When I bought my last car, the dealer said, "Don't hit the brakes hard for the first 500 miles." I was curious why, and he said, "The car is going to learn how you generally brake and will adjust over time." If you hit the brakes hard when it's not an emergency, that car will start thinking that that is your regular braking pattern and adjust to that. That's AI, isn't it?

Michael Feldstein 07:45

Yes, and what comes out in that example are two critical factors in understanding some limitations of how good machine learning and AI can be. One is data. So, if you started driving that car and decided that the first thing you're going to do is teach your 16-year-old child how to drive by putting them in the driver's seat, there will be a lot of jerking and slamming on the brakes while the car is going to work from this data. So, the algorithm for the braking is going to adjust only based on the data that you give it.

The other aspect is that the person who designed that algorithm had to think about different aspects to tell the car to pay attention to, for example, if you are going uphill or downhill. Or if it's tied into your GPS, does it know whether you're at an intersection? Those factors might influence how it learns about your braking patterns, but only if the programmer tells it to pay attention to those details in advance. Like with a spam filter, it might notice that most spam comes on Tuesdays and Wednesdays, but only if the spam filter developer tells it to look at the day of the week.

Drumm McNaughton 09:20

That helps me put into context how complex this is. It depends not only on the programmer but on the information it will gather based on what the programmer tells it to look at.

Michael Feldstein 09:40

Yes. These are incredibly powerful tools. We can do all kinds of crazy new things that we couldn't do with them before. Just like tools we've been creating since the beginning of time, like the wheel, the inclined plane, and the pulley, they have very specific uses, and the better that we understand how they work, the better we'll be able to know what they're good for and what they're not so good for.

Drumm McNaughton 10:14

This brings us to our next topic: AI as a tool. I love Shakespeare's quote, "Nothing is good or bad. It's thinking that makes it so." So, is AI good? Is AI bad? It depends on how you think of it. But, in essence, it's a tool. So, what are some of the problems that could be useful in solving from an education perspective?

Michael Feldstein 10:42

Let's talk about a couple of teaching and learning problems. I'll give just a few quick examples and then broaden to how the university runs and some other essential aspects of teaching and learning.

For example, we know that Google has an algorithm that can take a poorly worded search phrase and interpret it by saying, "Well, I think what you mean is this, so I'm going to give you good results." It's very good at understanding the meaning of, at least, short phrases. We also know software applications detect plagiarism and say, "Well, this looks a lot like this other source, so it might be copied from that other source, right?" Combining those two results in a tool for evaluating whether students paraphrase well. Are they going too far and copying the text? Or are they interpreting the text in a way that differentiates it enough from its original meaning? Are they getting that sweet spot in the middle?

Second, the educational application has as much to do with us educators as it does with students. Sometimes when teaching students, we don't reach them, and they don't learn. Some classes might have a pattern of how those students don't learn. It can be challenging for a human educator to catch those systematic errors that the students are making. But suppose you have well-written questions with a well-written learning objective and software that analyzes multiple students' patterns of answers. In that case, not only can we learn whether a student is progressing towards competency or where they might get stuck, but we can also begin to determine if there's a problem with a particular part of the course design and where students get stuck on a specific topic. That way, we can do a better job of teaching the course in the future.

If we step back more broadly, we can use these technologies to, for example, identify earlier those students who might be in danger of dropping out because they've missed a week or class or two. By itself, this might or might not be a red flag, but it would be different if we knew that they work two jobs and commute to campus. If we combine this data, which the average person might miss, we can begin to discover the patterns of what and how students at a particular institution are likely to drop out.

Likewise, we can use that information to respond to students, which is a huge thing right now via [chatbots](#). At universities, chatbots are helping us see students who need help in a way that might prevent them from registering, getting on campus, passing their classes, or getting the financial aid they need. Students often ask the chatbot a poorly-worded question, like how we might phrase a web search term. The software can do an excellent job of helping these students find the correct answers. Then for questions that need a human to answer, these chatbots direct the students to humans who can help them. We're ensuring our student support professionals respond to students who need them instead of those who need a quick answer.

Drumm McNaughton 14:42

You brought up three excellent examples of how AI can be used. One is assessing student learning outcomes, including redesigning courses. Two is identifying students who are slipping behind or are at risk of slipping behind by recognizing patterns. And three, reaching out to students via chatbots. Personally, the ones I've seen are pretty useless, but I'm sure they'll improve. AI will teach itself. But being able to reach out to students before or when they start needing help is crucial. So that's those are three excellent examples.

Michael Feldstein 15:34

Let me add a fourth one since you mentioned whether chatbots are helpful. We often use sophisticated statistical analysis when trying to understand how well those chatbots work. For example, is this effective if we're doing a research project? That statistical analysis may be close to the algorithm you need to inform the chatbot to learn how to improve. You'll begin to measure specific parameters, like how much it matters that students receive quick responses or the types of responses that chatbots can more successfully help students achieve their goals or make more effective actions than others. Once we gain those insights through the research, we can start to automate some of that insight to refine itself.

Drumm McNaughton 16:39

Your explanations are just lighting sparks inside my brain. What about maintenance and buildings? Plant equipment? You can easily start to see these kinds of things. I hate to say it, but I'm going to. The possibilities are endless, or nearly endless. You do need a human component to make certain decisions. [This could even work with strategic planning.](#)

Michael Feldstein

Yes.

Drumm McNaughton

It requires good data that is organized correctly, but AI can find patterns that an average person would miss, e.g., [enrollment patterns](#).

Michael Feldstein 17:24

You bring up a critical point: the best use of machine learning and artificial intelligence is often not to replace the human but to augment the human. The term is: [human in the loop](#).

Drumm McNaughton 17:39

Okay. That's a new one for me. It's just amazing. We've seen so much negative hype. I haven't seen anybody talking about the positive aspects of this. I'm specifically speaking to ChatGPT. They're saying, "Oh my God! Students are going to be able to plagiarize all their papers!" Why is it getting so much attention?

Michael Feldstein 18:07

One of the reasons is that ChatGPT is a toy. I mean this in a neutral sense. I don't mean to denigrate. By toy, I mean that it invites play. It's easy for anyone who's not a technologist to use it, learn how to get better results, and find out what it's good at and not good at.

ChatGPT has improved tremendously, partly because it's getting all this data from constant usage. People can ask it anything, so they do ask it anything. So, the creators are learning what people will ask it, what it's good at, and what it's not. One example that might surprise you is that ChatGPT is terrible at math.

Drumm McNaughton

Really!

Michael Feldstein

This is a program that has billions of parameters. It's one of the most complex, expensive, and carbon-intensive software programs on the planet. But what it really is—at its heart—is a very, very, very sophisticated version of your “type ahead” program that you can use on your word processor. I'm oversimplifying here, but the basic idea is it can guess that the next letter you type might be a “C” if the letter you just typed is an “S.” But it won't be a “C” if you typed an “X,” right? It can guess that if the word you typed was “school,” the next word you might type is “bus.” But it probably won't be “chicken.”

Drumm McNaughton 20:00

Unless we're talking mascots.

Michael Feldstein 20:03

We might be talking mascots. It's not impossible, right?

You can think about this in terms of letters, words, phrases, sentences, or entire paragraphs. Remember that ChatGPT was trained on a big slice of the internet and many books. So, it can look across vast swaths of texts. If you're talking about cellular mitosis, for example, it's read a lot about that topic. It can anticipate all kinds of combinations of words, sentences, and ideas that might come together on that topic. So, it can give you responses that are spooky. They sound like ChatGPT understands the question that you're asking and is giving you an intelligent, insightful answer even though what it's done is boiled down a lot of information from the internet and rearrange it in ways that are so complex that we can't even follow precisely what it's doing.

Math, on the other hand, is a straightforward and very deterministic process. So it is surprisingly difficult to figure out how to take a function as simple as a basic calculator, not even a scientific calculator, and attach it to something that works like ChatGPT. First, there was a news story about how ChatGPT couldn't do math. Then the people who created it said, “Well, we fixed that.” And just last week, I asked ChatGPT, “Are you good at math?” It said, “I'm pretty good. There are certain things I can't do. But I certainly can do basic arithmetic.” So, I asked it a basic arithmetic question using a couple of relatively long numbers. It was an additional problem, and it got the wrong answer.

If I asked it today, it would give the correct answer because the programmers are constantly working on it, and the software is always learning. But it just shows you that if we understand that ChatGPT is a very sophisticated tool, rather than magic, we can begin to understand that, for example, it makes sense that it wouldn't be good at math.

Drumm McNaughton 20:19

That makes perfect sense to me. So those are some of the cons. Something else we talked about was that it's not predictable.

Michael Feldstein 20:29

Yeah. So, it was trained on the internet....

Drumm McNaughton

Which we all know is truthful, right?

Michael Feldstein 20:29

Right. People say all kinds of things on the internet, and ChatGPT has been known for giving very confident wrong answers. Right? That's bad enough when trying to find the correct answer. So, when people ask whether ChatGPT will replace search, it's important to understand that search is designed as a different tool. Making these two types of tools work together is challenging.

When we start thinking about students using ChatGPT to write essays for them, one of the many problems educators worry about is that ChatGPT will write an essay that is full of falsehoods or maybe even hateful or biased language because it doesn't know any better unless we teach it otherwise.

Drumm McNaughton 23:36

It doesn't have a moral compass. Is that what you're telling me?

Michael Feldstein 23:36

It doesn't. The folks who made it are working very hard to give it one, but because it's so complex a program, it's not entirely predictable.

Drumm McNaughton 23:52

That brings to mind the movie War Games with Matthew Broderick: "Would you like to play a game called thermonuclear war?" It's hard to set rules, etc. But there are good things, and we've discussed some for higher ed. We discussed its implications for teaching, assessment, student services, and even running plants. There are many things there. Now, we're starting to see some of these platforms. Google has its Bard. Bing has been in the news lately about how it named itself and some other crazier things. What are some things you anticipate will happen with AI over the next two years?

Michael Feldstein 24:50

Well, it's tough to predict. As we've been discussing, sometimes these technologies can seem miraculous and really are miraculous. I was a philosophy major in college, and, just as an exercise, I had an interesting conversation with ChatGPT about the [philosopher David Hume and his understanding of what knowledge is](#). I actually learned something from that conversation. So, it can surprise us in positive ways and negative ways.

I like to talk about AI development in three stages. First, there's the miracle, the grind, and the wall. The miracle is when you have an experience with a new technology, like ChatGPT. I'm sure we'll have it with Bard from Google and GPT-4, which will be the next version of the model that underlies ChatGPT.

These chatbot-like models we have been talking about are called large language models, and they are AI programs that work very differently from them. So each one will come with a miracle, where we'll say, "Oh, that's going to change everything." Then we'll realize, "Well, there are these little problems. It just can't get math right. Well, that should be easy to fix."

That's when we go from miracle to grind because we discover it's not easy to fix. And by the way, it's not easy to do what developers call QA or quality assurance. So how do you test to make sure you're going to get reliable, correct answers? And this is really important when dealing with students' lives or their learning.

Drumm McNaughton

Or airplane software.

Michael Feldstein 24:38

Or airplane software or automated driving. There are lots of different applications. But if you get a bad search result, that's not nearly as bad as if you get a bad plane landing vector, right? So, you go through the grind of figuring out how to solve the problem.

Then, sometimes you'll hit a wall unpredictably. For example, we have these grand aspirations for intelligent tutors and personal learning software, but it may turn out that the data we have from the learning systems is not what the software requires to understand student learning. So, the whole system this AI algorithm plugs into is designed poorly for that kind of software learning application and needs to be rethought from scratch. That's the wall.

Drumm McNaughton 28:00

Yeah. My brain just hit the wall because my jaw hit the desk. From a student services perspective, we must rethink how we've done student services and what's important. We've seen some institutions do this. For example, Amarillo College's President, [Dr. Russell Lowery-Hart](#), who's just an incredible visionary, [improved persistence and graduation rates](#) because he knows who his student population is. He knows what their limitations are and brought those services in that were needed. That's rethinking how student services are done at a particular college. It's not just one size fits all, either. This works for his demographic. But another demographic at schools like Harvard or Yale would be different. So, each institution is going to have to figure out what the parameters are. That's the necessary critical thinking aspect that comes from humans. Then, bring it into ChatGPT or another AI platform, and present the data about what's needed to analyze trends, etc. Does that make sense?

Michael Feldstein 29:17

It does. There's a third piece in the middle that you must be thinking about, too. If I'm a human counselor, I'm interacting with students and taking notes so that the next time the student calls, we have a record and can respond to that student. What am I capturing? Is it in my notes about my interaction with the student? What am I not capturing? Am I capturing it in a machine-readable way so that the AI, machine learning, or data-mining algorithm can make sense of it and do something with it? There are a lot of human contexts that we all understand intuitively but that the software doesn't. So, we have to make sure that we're explicitly encoding for the machine to learn and for the human to learn when recording our interactions with the students.

Drumm McNaughton 30:23

That brings up an interesting point about how certain programs convert an audio transcript into text. I'm guessing this is what we call a WAG, a wild, you-know-what guess. If you can record these conversations and then upload them in text, can you give the computer this kind of context? Does that make sense?

Michael Feldstein 30:58

You and I agreed to try something like this for this podcast. You will give me a clean transcript; I will write a prompt for ChatGPT, asking it to summarize; then we will share the transcripts, the prompt to ChatGPT, and

the output with your audience, and we'll see how those things turn out. We'll see how good ChatGPT is with the input of the text and some prompting from me to produce a viable summary.

We can imagine taking that further. I write a web blog called the e-Literate and have very idiosyncratic writing and thinking styles. I've written a lot. I would love to point one of these algorithms at the thousands of posts I've written and see if it can capture the sentence structures and some of my thought patterns and expressions. I'm not looking for it to replace me. I would like to see if we can have a dialogue where I can find a like mind.

Drumm McNaughton 32:17

Well, you have a like mind in me. I don't know if you will have it in the AI. But I think that's going to be a wonderful experiment.

Michael, what are three takeaways for university presidents and boards that they need to consider with this topic?

Michael Feldstein 32:31

First, don't be overly afraid or excited. Focus on understanding. The second is always to think about the human in the loop. A university is a collection of the ultimate knowledge workers. You don't want to waste those big brains by having them conduct many menial tasks that software can perform. For the third, I want to mention the following expression: there's a solution when searching for a problem. We must strike a balance between what that expression implies and the other way around. Sometimes, we can suffer from Stockholm Syndrome regarding the technology and the processes we live with. Sometimes, new technologies can suggest new approaches. If we understand what's important, we can stay focused.

If you're a board member, president, or senior leader, stay focused on the metrics, goals, and values that matter. But remain flexible and open-minded about the new ways to achieve them more effectively and cost-efficiently.

Drumm McNaughton 33:45

Excellent. Thank you. This is going to change the way we work. We have to recognize that it has strengths and weaknesses. But just like the wheel, we won't have to drag things. We're going to be able to roll things. The question is, is it going to roll smoothly? Or is it going to roll over us?

Michael Feldstein 34:08

We still have to drive to work, but we still have to go to work to do work.

Drumm McNaughton 34:14

Exactly. So, what's next for you, Michael, besides our next experiment?

Michael Feldstein 34:20

That's really what life is all about, isn't it? The next experiment.

Drumm McNaughton 34:25

It certainly is. So, Michael, thank you for being on the show. This has been a fabulous conversation for me. I look forward to our experiment. This is going to be fun.

Michael Feldstein 34:38

I agree. I'm curious, too.

Drumm McNaughton 34:40

Take care, my friend. Thank you.

Michael Feldstein 34:41

Thank you, Drumm. I really enjoyed it.

Drumm McNaughton 34:44

Thanks for listening today. I want to thank our guest, Michael Feldstein, for sharing his thoughts on the benefits and drawbacks of artificial intelligence and machine learning and how they can affect and help higher education institutions.

Our next guest is Elissa Sangster, CEO of the Forté Foundation. Elissa will join us to discuss the support women and minorities need to ascend the leadership ladder in higher ed. Elissa brings with her experience in both higher ed and corporate America. This talk will be enlightening for higher ed leaders who are striving to improve their diversity and leadership positions.

35:28

Changing Higher Ed is a production of The Change Leader, a consultancy committed to transforming higher ed institutions. Find more information about this topic and show notes on this episode at changinghighered.com. If you've enjoyed this podcast, please subscribe to the show. We would also value your honest rating and review. Email any questions, comments, or recommendations for topics or guests to podcast@changinghighered.com. Changing Higher Ed is produced and hosted by Dr. Drumm McNaughton. Post-production is by David L. White.