Patrick Kiessling:

Hello, everyone. This is Patrick Kiessling, one of the medical students on the team behind ENT in a Nutshell. Don't forget to subscribe to this podcast on Apple Podcasts, Podbean, Stitcher, Google Podcasts, or Spotify so that you never miss an episode. Now, onto the episode.

Dr. Ashley Nassiri:

Hi. Welcome to ENT in a Nutshell. My name is Ashley Nassiri. Today, we are here with Dr. Naweed Chowdhury to discuss the basics of artificial intelligence in otolaryngology. Dr. Chowdhury, thank you so much for being here.

Dr. Naweed Chowdhury:

Hi. Nice to be here and really excited to chat a little bit about this topic.

Dr. Ashley Nassiri:

So before we begin, I'd like to introduce our guest speaker. Dr. Chowdhury is a rhinologist who has a keen interest in data analysis and statistical research. His current research focuses on predictive analytics in applications of artificial intelligence or AI within otolaryngology.

Today, we're discussing a topic that doesn't really follow our usual recipe for patient workup and treatment, but has wide implications in the future of medicine and data analytics. We will go over the technology known as artificial intelligence, it's medical and surgical applications, limitations of the technology, and future implications for practice. So let's jump right in. Dr. Chowdhury, let's start at the beginning. What technology does the term artificial intelligence specifically refer to?

Dr. Naweed Chowdhury:

So it's a pretty broad kind of basket of different algorithms and techniques, but in general, it's technology that allows computer systems to perform really narrow tasks at this point, and hopefully, at some point, very broad tasks. But these tasks are typically things that required normal human intelligence. So for example, looking at a picture and identifying what it is. Is that a cat? Is it a dog? Is it tumor and whatnot? Also, looking at things like object and word recognition and decision making rules, interpreting language, and one of the more fun ones is competitive gaming.

So you may have seen IBM's AI beat Ken Watson during Jeopardy and playing chess and beating grandmasters at chess and that kind of thing. So it's really exciting. And what's really caused this to take off more recently is that we've got immense amounts of data, and we can store it really cheaply, and we can actually process it faster than we've ever been able to do so. And so those things have really led to, I'd say maybe in the last decade, a real explosion of these algorithms and finding applications for them in different realms.

Dr. Ashley Nassiri:

So how would you say that AI is different than data mining or statistical methods that look at large databases to develop predictive models?

Dr. Naweed Chowdhury:

I think it all is essentially on a spectrum. And there's a lot of overlap and a lot of really gray areas. Some machine learning models are essentially common statistical methods like logistic regression, others are much more complicated, things like a convolutional neural network or an artificial neural network. And so I think depending on who you ask, you'll get a different answer to that question. So I think a statistician would have a much different answer than someone who's in the computer science realm as to what exactly machine learning and artificial intelligence is.

But my personal take on this is that these are generally methods that use flexible modeling tools that use iterative approaches to identify patterns in data. Essentially, you are giving the computer a set of data and then the answers are what you want it to learn from the data, and then kind of telling it, hey, figure out the rules to this, come up with a model that fits this data to the answers.

And then frequently, you then are looking to apply this to a new set of data and see, well, how well did this machine actually learn the rules? And oftentimes, there's a random component to this, which I think is a little bit different than things like a linear regression model where all the terms are fixed with different coefficients and they're all added together. The models can get much more complex with artificial intelligence.

And the trade off really is that you can have a model that is not easily explained, so it's kind of this black box, it kind of spits out an answer, and it might be right, but you don't actually know how it got that answer. We have to be a little bit careful with that compared to something like a regression model where you can say, okay, you just take the age and add this coefficient and that's the answer.

Dr. Ashley Nassiri:

So we already started to discuss this a little bit, but what does AI offer that we currently can't do with statistics and large data sets?

Dr. Naweed Chowdhury:

What AI is able to do is essentially identify complex patterns in data that humans may not necessarily recognize because of either the sheer amount of data or the amount of collinearity between data. And so as a result, it's I think really great for exploratory analysis and generating hypotheses on large data, in particular, looking at nonlinear relationships. And the example that I'll give is, think of say some data that's described by Y equals X squared. And so it's like a parabola basically.

And you try to model that with a linear model, i.e. a line, you're not really going to find that relationship, you're actually likely to find a zero relationship or a no relationship when really the variable is completely described by the input. While you can model a linear regression with an X squared term, it's really unusual to do that. I don't think I've seen many examples of that in ENT, and so we are often missing these nonlinear interaction terms when just fit standard models.

And so some of what machine learning methods can do is pick up on that without explicitly being told, hey, I think that there's this relationship going on. And so it really helps say, hey, maybe there is this non-linear term, and then we can specifically design an experiment to look at that in the future.

Another good use for AI is to look at really informationally and computationally intense tasks, specifically things like parsing out tumor volumes on a CT scan or reading pathology slides, things where the data is not on a spreadsheet or quantified in a number. And so we can automate this process potentially, and instead of using tumor staging, we could just feed the pathology images and the CT scan into the computer and it could automatically say, hey, there is X probability that this is going to be a malignancy versus a benign tumor. And this is what the likely outcome is. So instead of compressing data where we're using the full breadth of what's available for us.

Dr. Ashley Nassiri:

Outside of medicine, what are some examples of AI technology being used today?

Dr. Naweed Chowdhury:

So I think one of the fascinating things is how intertwined AI is already in our daily lives and how often we use it without even knowing. So something as simple as your spam filter on your email is a complex AI algorithm that has studied millions, if not billions of emails that have been labeled as spam and predicts and filters out new emails that might be spam sometimes accurately and sometimes not.

Also, when you're watching Netflix and it shows you these are the new movies based on your history that we think you might like, or when Amazon tells you to buy more stuff and you're like, okay, this is actually stuff that I wouldn't buy, and it's kind of scary how it can predict that, all of that is using some type of machine learning algorithm. And it's been feasible because these companies have so much data on you more than we even know. And so they're all using that to try to figure out how to sell us more stuff and get us to watch more stuff.

Dr. Ashley Nassiri:

So now that we've talked a little bit about what AI is, I think it's always important to take a look back on development in the history of topics. So can you tell us a little bit about how AI was initially developed?

Dr. Naweed Chowdhury:

Another one of the fascinating things about AI, I think is that the key concepts of it were really developed back in the 1950s. One of the first conceptual understandings of AI was by Alan Turing. And he came up with this idea of the Turing test, which he called The Imitation Game. And that's why the movie is named The Imitation Game about his life.

And essentially, he posed this question, how would you know that a machine was thinking or a machine would have human-level intelligence? And he said that it would be when you could talk to this machine and this discussion with this machine would be indistinguishable from a human. And so that's always been the standard and it's been evolved a little bit since then, but that planted the seed of artificial intelligence in consciousness.

And so afterwards, in the 1950s, a computer scientist named Marvin Minsky developed the first neural network. So I think we think about, this neural network is this cool new hot topic and it was actually done in the 1950s using a bunch of vacuum tubes and they created a 40-neuron artificial neural network at that time. And then soon afterwards, another scientist, Arthur Samuel, developed a program that could play checkers. And Arthur Samuel is actually credited with coming up with the term machine learning as well.

What's interesting is in the history of AI, there've been these booms and busts. And actually, right after that period, there was a kind of a bust and it's called the AI winter. And essentially, people really got disheartened and they said, we have all these cool ideas, but there wasn't enough funding to really maintain things. And more importantly, these algorithms were so computationally complex. They could say, if you could have a machine that could do X, Y, and Z, you could have artificial intelligence. But there was no machine at that time that could actually do any of that. And so people really got frustrated and a lot of people went into other fields.

But then in the late 1980s, you started to see the pace of computing takeoff and a computer scientist named Geoffrey Hinton came up with this article on backpropagation, which is essentially how a neural network could learn from data and the process by which a neural network could learn from data. And even today, most of the neural networks are using some variation of backpropagation to learn. And after that, there was essentially an explosion of AI research again, you had IBM's Deep Blue beat Garry Kasparov in chess in the mid 90s.

And then really, in the next decade, CPU and GPU hardware just begins to just get faster and faster and faster. There's been almost a democratization of this technology where today, you could buy a computer that essentially would be considered a super computer back in the 90s. And so people have really started to use these tools to apply them to their own problems, which I think is really exciting, and that's the era that we're in now.

Dr. Ashley Nassiri:

That's really interesting. Before we get into the medical applications of this technology, can you explain how AI technology actually works?

Dr. Naweed Chowdhury:

There's a couple of different specific algorithms that is generally considered AI, but from a high level view, essentially, these algorithms are rather than explicitly given rules to figure out how to interpret data. They're expected to learn the rules from the answers. And so one analogy would be instead of giving a medical student a textbook of pathology and say, "Okay, I want you to read this and study this," you essentially would give them a bunch of images of squamous cell pathology slides and just say, "I want you to learn what makes a squamous cell carcinoma look different from a normal tissue," and that's it.

You just give them the images and the answers and say, "Good luck." And so that's what AI does that's different from, I guess traditional programming. And then the next phase of that is then to check understanding. And so you would ask this AI or medical student to classify new images and see where they or how accurately they predict them and how what they have learned, I guess, from the data.

Dr. Ashley Nassiri:

So we've used the term neural network a few times already. Can you explain what that means exactly?

Dr. Naweed Chowdhury:

So a neural network is loosely based on the concept of a neural network in biology and how biological nervous systems operate. And essentially, you can think of a neuron biologically as something that takes in electrical signal, does some sort of processing to it, and then outputs it to other neurons. And so there's a similar electronic neuron, if you will, within a neural network. And often, there is several, if not millions of them, in really some of the bigger neural networks. Essentially, they're taking in an input, doing some kind of processing, and sending it to another node or another neuron that is then taking in that input, doing processing, and sending it onwards.

And it seems somewhat basic, but there are actually some mathematical proofs that say using an appropriately sized neural network, you can model any function in the world. And so that's where the power of the neural network really is. And so there is some term such as nodes, which are basically just different. I guess the analogy would be like cell bodies basically. There's often this term called a hidden layer, which essentially, you have your input, and then you have a bunch of processing that happens in these nodes. And sometimes, there is just one set of nodes that's doing that processing. So there is one hidden layer sometimes for things like convolutional neural networks, there is several hidden layers where processing is happening, and you can think of each layer as being a more and more abstract level of processing.

So you could say the analogy in a person would be first, the cochlea is detecting frequencies, and then you're sending these frequencies into the brain, and then the brain is actually taking these frequencies and turning them into perceptions of sound. And then those sounds are words. And there's another level of abstraction about that that's changing and interpreting sounds into words and then maybe words into emotions. So there's all of these different abstractions and different areas of the neural network are doing different processing calculations similar to how different areas of your brain are doing different processing of the same input signal.

The other thought is this concept of the deep neural network. And that's also a loosely defined term. There's not a point at which a neural network, I guess, becomes deep. It's just known as a neural network that has a multitude of hidden layers where there's a lot of processing that's going on underneath the hood. But the main thing for a deep neural network is that they're often very difficult to interpret. And so a really simple neural network, you can actually look at the different weights and you can say, okay, this is how this is working. And it's more like an equation, but for a deep neural network, there's just so much overlap and so much complexity that it's often a bit of a black box.

Dr. Ashley Nassiri:

Thank you. I think that was a great overview of a very complicated technology. I appreciate the hearing reference. Outside of neural networks, we hear a lot of buzz words around AI. Can you go into the different sub-fields of AI and what each involves?

Dr. Naweed Chowdhury:

Yeah. There is a lot of different areas that are being investigated in terms of AI. And each one is probably looking at different input signals, I suppose, would be one way to look at it, or different data. So there's natural language processing, which I use every day. I'm using Dragon and I'm dictating into it and it's spitting out my words usually pretty well. And that's essentially what that's trying to do is to process language and interpret language eventually.

So maybe it doesn't have this capability now, but perhaps I could say, "Order Augmentin," and it just knows that that's what I'm meaning and it does it automatically. And so that's where that technology is going. There is the concept of machine learning, which I think is broadly just any algorithm that learns from data. And so I think even an algorithm as simple as the mean of a variable is a form of machine learning if you do it on a computer, but obviously, it's generally used to refer to more complicated algorithms that are often non-linear and sometimes random.

And so within that, there is a couple of different subdivisions. You could think of supervised learning, which is what we've been talking about mostly here today, and I think that's probably one of the more common uses of machine learning. And what that means is when you're supervising learning, you're giving the algorithm, the answers, and expecting it to learn rules and relatIonships. So you would essentially give it a set of data, what you're looking for, and say, hey, figure this out.

In contrast to that, there is the concept of unsupervised learning where you are giving the algorithm just a set of data and saying, I'm not actually sure what the relationships are, and I'm trying to see if you can figure out how many, maybe different groups or different subdivisions there might be in this data. And so one of the real common uses for this, especially in the research that we do, is clustering.

So there's an interest in the study of chronic sinusitis of trying to determine what different inflammatory endotypes would be. And so cluster algorithms are essentially using biomarkers and trying to say, hey, when we look at this from a statistical perspective, it seems like there's maybe seven groups or 10 groups of inflammatory biomarkers that seem to correlate with each other.

The downside to unsupervised learnings, because we're asking a computer to do learning without knowing the answer, there's a lot of variability. And that's why in endotyping, there is about four or five different papers that have looked at this and they all have different answers. And it's all about your inputs and where you draw the line for your outputs. But it's still an active area of research.

I talked a little bit about deep learning, which is typically using very large neural networks to perform this learning process. There's a couple of interesting new things that have come out just in the past year. I think the most fascinating but also very scary is this idea of a GAN or a generative adversarial network. And essentially, they are two neural networks that are training each other.

This sounds like Skynet, but it's actually pretty impressive. The neural networks are actually trying to fool each other. And so there's some evidence that shows that this method of training and network is faster and more efficient and can do really well with certain problems, particularly generating images. And so you may come across some people who have posted synthetic images of faces that look very real, and those are usually done with GANs.

The Holy grail of all of this and what people are really working towards is this concept of artificial general intelligence. And so this would be something really akin to human intelligence, and it would be something that would be adaptable beyond just really narrow tasks. So something that could essentially have true human level intelligence. This is associated with a concept called the singularity.

And this is the idea that at some point, we could develop general intelligence that could be so smart that it could actually create smarter and smarter versions of itself and eventually surpass human intelligence by several orders of magnitude. And obviously, we've probably seen the implications of that in lots of sci-fi movies. But it is something that I think is increasingly possible and maybe even likely there's different theories on that, but let's just hope that we don't have Skynet and have to have the Arnold save us at some point

Dr. Ashley Nassiri:

That's just absolutely fascinating. So let's say if we had access to machine learning technology and we had a data set that we wanted to analyze and develop predictions to a specific question. More broadly, what are the steps that would go into doing something like this?

Dr. Naweed Chowdhury:

The fundamental thing about a lot of machine learning is it's all about the data. And I think sometimes people lose sight of that. And it's all about really getting high quality data upfront, which is the same problem that we have with standard research problems. So let's say we had this database of images of different say laryngeal tumors and we want to identify which ones are just benign nodules and which ones are squamous cell carcinomas. So ideally, you would have a database of images, and with machine learning more is better.

Some of the top machine learning and deep learning networks that have been published by Google and Facebook all use millions of images for them. So it would be something on that order of magnitude that you would want, especially for something like this. You would essentially separate this into a couple of sets. So one, you'd have a training data set where it would be a bunch of images, and typically, it's anywhere from 70% to 80% of the images that you would use to teach the algorithm how to differentiate between cancers and non-cancers. And you would ideally also want a fairly even distribution of both within your data set so it could learn the features of both of those.

And then you would have another set of data that you just keep in your back pocket, and it's something that you would only really use at the very end to test the performance of your model and test the predictive accuracy of your model. And so you can then label these images afterwards and then feed the images into your algorithm of choice. And there's probably a whole different lecture that we could do on that topic, but more commonly, for this sort of thing, we would probably use a deep neural network. It seems to have the best performance. But you can use any of a number of them.

Essentially, at the very beginning, your network doesn't know anything. And so you initialize the network with just your best guess, just almost random values. And then it would run the model once and it would say, hey, I missed 90% of these images. And then what it'll do is create an error function that determines, okay, this is how I could adjust the weights of my algorithm to improve prediction on the next iteration. And then it runs the whole algorithm again, and then hopefully, your [inaudible 00:23:54] is better.

And then again, it's this iterative process where it keeps adjusting how much weight it's giving to, say different pixels or different aspects of the image in a circular fashion until it converges on a final-ish algorithm. And then you would have this trained algorithm, and then you would test it on your dataset and hopefully it works really well. And then you publish it and you get to do AI podcasts on.

Dr. Ashley Nassiri:

So it sounds like a lot of the success of the model obviously depends on what kind of data you put into the algorithm and how it's developed. What are some of the other limitations of artificial intelligence?

Dr. Naweed Chowdhury:

So for most AI algorithms to be effective, you really need a lot of data. And this is almost an order of magnitude, I think, beyond what we typically would consider big. Even something like 1,000 patients worth is pretty small by machine learning standards. When Amazon and Netflix are training their algorithms, they're literally using millions of inputs from millions of people. AI really works well for that. But you have to be careful because you have to see how the data is being generated and what are the biases of the data itself? Just because the data is being generated in a particular fashion and maybe it's not being collected by someone, it doesn't mean that it's unbiased.

And one of the recent examples that came out of this was this healthcare prediction algorithm that was run by a big insurance company. And they basically wanted to automate the process of who should get what tests and care? And they used their "unbiased" data to build this algorithm. And what they actually found is that the algorithm was systematically recommending less care for African Americans, for example, even though the only real difference was their race, and these people were just as sick or had the same co-morbidities as other populations.

And it's because of the biases in our healthcare system that cause people or that generally deliver less care to African Americans. That underlying bias was just reflected in the algorithm. Fortunately, it was caught before it was deployed on a widespread level, but it just shows you that it doesn't undo any of the biases of the data that's being fed into the algorithm. We have to be, almost I think, more careful with these kinds of algorithms because if say you fit a regression model where you can maybe very easily look at the different parameters and you could say, hey, our beta for African Americans was really weird compared to everything else, something like a neural network, you don't have that number.

You don't have a statistical way to quantify that. And so you could just blindly deploy it and say, hey, it does a great job 90% of the time, but for those 10% of the people that it's not doing a great job from it, it may not be a random 10% and maybe systematically biased.

Dr. Ashley Nassiri:

Well, that leads right into our next question is that certainly, this technology has the capacity to do great things, but it also can have some negative outcomes. I think the example you listed is a tremendous example of something that we don't want to happen in the future. We do hear some of these negative opinions about AI and the potential dangers down the road. Aside from including some of the biases that are inherent in society today, what are some of the other concerns or how do we address these concerns?

Dr. Naweed Chowdhury:

So there's obviously a concern about computers working autonomously without human inputs or controls, especially if we do develop a algorithm that has general intelligence or something that's not tailored to a specific task. As we discussed in the last question, because of these biases, we need to actually be even more careful about what we tell computers and what we use them for, and to be critical of the conclusions that they're drawing. Sometimes, there's a really [inaudible 00:28:06] quality about artificial intelligence, the promise of offloading some task that is not considered to be very interesting or some [inaudible 00:28:16] thing that we can just have the AI algorithm do. But if we are making decisions on those outputs without being critical about them, we may be doing people a disservice, especially in healthcare for example. We should definitely keep having the humans be critical in service checkpoints for the conclusions of artificial intelligence.

There's also this concept of the black box. That's a danger of artificial intelligence where oftentimes, we don't necessarily know why the algorithm is making the decisions that it is. And we saw that in the last question. But even more so, there's issues with interpretability and you could imagine somebody publishing an algorithm and it says, okay, well, based on our data, there's no difference between people with polyps and no polyps. And there may be some underlying bias within that data set that caused them to lead to those conclusions.

And finally, I think from a political standpoint, there's a lot of economic considerations that are present. And I think you're hearing that around the ideas of universal basic income and what happens when robots and artificial intelligence takes jobs that are usually done by humans. And I think we're still ways out from that, but you could certainly imagine if artificial intelligence algorithms become very potent, even jobs in healthcare, perhaps radiology jobs or pathology jobs may be at risk, especially if you have an algorithm that can automatically process a lot of this data on its own.

With rising healthcare costs, you could easily see an insurance company saying, well we're only going to get a pathologist to read this if there's something really weird, or if it's some really rare tumor and our algorithm doesn't have good performance on that. So I think these are some of the things that are being discussed in society now, but I think will just become a bigger issue down the road.

Dr. Ashley Nassiri:

So it sounds like we need to teach the machines to incorporate morals and ethics into their algorithms as well. So taking all of that into account, the capacity for this tremendous technology, but knowing that it can have limitations that can have significant consequence, would you envision for the future any specific ongoing projects in otolaryngology that you think that are headed somewhere promising?

Dr. Naweed Chowdhury:

So I think there's a lot of interesting things that are going on that have already been published, and I think we're just at the very tip of what's possible. One of the things, I guess this is fairly broad, but I know a couple of different folks are working on the use of voice recognition, not specifically in otolaryngology, but in general, to potentially perform automated clinical documentation. So I think you have an algorithm listening to your conversation with the patient in the background, and it's just automatically generating your note.

So I think that would be something that would be really cool, certainly improving our quality of life as physicians. But then more specifically, I think we're doing some work here looking at vestibular schwannomas and trying to predict which tumors will grow, which tumors won't grow, trying to identify characteristics based on imaging. One of the things that I'm particularly interested in is using machine learning and artificial intelligence to understand disease insights. And in chronic rhinosinusitis, we're at this era in biology where we're getting such granular high-dimensional data on inflammatory processes that we have never really been able to get at this scale before.

Taking that data, taking clinical data inputs and combining them together to investigate and identify non-linear interactions between the two and really gain a granular understanding of why people get chronic sinusitis and why people do well, don't do well, et cetera. There's also a growing interest in image segmentation and automated imaging and identifying characteristics of features of images that may be useful from a clinical standpoint without having somebody explicitly create like a [inaudible 00:32:52] score or something like that. I think that's pretty exciting and interesting.

Ultimately, the holy grail, I think would be to have some type of algorithm that is combining all these different elements of say, pathology slides, imaging, maybe endoscopy images, inflammatory data, maybe even the patient's story, you throw in a little natural language processing, and it's able to really tell based on a discussion if this patient has clinical symptoms and signs of CRS or not.

And putting that all into one algorithm, that then could spit out a probability score of, okay, there's a 90% chance this person does have chronic rhinosinusitis. And if you do surgery on them, their chance of improving their SNOT-22 is 75% versus for medicine it's 25% and vice versa. So I think that's what the grand goals and schemes that we're thinking about are, but there's certainly a little bit of ways to get there at this point.

Dr. Ashley Nassiri:

That would be tremendous if we had help with decision making and taking some of the guesswork out of how to treat our patients. One question that comes to mind as we're talking about this is, do you see AI as a technology that would replace statistics, or do you think that this is something different than statistics in a way, and statistical analysis will still play a role in the future?

Dr. Naweed Chowdhury:

Statistics is always going to be there. I think AI is certainly not a replacement for statistics by any means. And with so much focus on AI, I think we often forget that classical statistical methods are very well validated. And often, for some of the questions that we really want answers to, particularly questions about causality and how things relate to each other, if a treatment will help or not, there's a reason why the randomized controlled trial is the gold standard. And it's just such an elegant tool in terms of addressing confounders and using random assignment to make sure that groups are equal.

And most AI algorithms are really based on observational data. And there's inherent limitations to observational data, especially with respect to causality. And so while you can adjust for them, and there is different methods that you can use to analyze that data in a more rigorous way, ultimately, I think that they will exist hand-in-hand. And I see them very much as complimentary tools rather than one surpassing the other. At least now. And maybe if you have this brilliant AI statistician that can rationalize and reason beyond our current understanding, at some point, that'll change. But for now, I think that they're both just tools.

Dr. Ashley Nassiri:

So for our listeners, do you have any suggestions on how to learn more about AI and how this technology can be leveraged to benefit otolaryngology research and patient care?

Dr. Naweed Chowdhury:

So I think because of the tools that we learn as physicians and the tools that AI requires, there's a relatively steep learning curve for an individual practitioner. A lot of the core concepts of AI and machine learning are in the realm of math, probability, statistics, and computer science, and you need a little bit of all of those topics to really understand AI on a very deep level. I thought I knew a lot when I was first starting to learn about AI, and it's one of those things where the more you learn about it, the more you understand what you don't know.

But I think that's also been a really big drive for me to just keep learning and expanding. And I'm personally just ... This is going to be super nerdy, but [inaudible 00:37:00] cracking open old textbooks from college and trying to relearn linear algebra, that kind of stuff, just because I think it's so important to really understand what's happening within an AI algorithm. So I think there's that aspect of it.

I think it really helps to have some kind of background in coding or some experience with coding. Most of the cutting edge techniques are using either iPhone or R or C++. And for those who don't know, those are all fairly commonly used programming languages. And so it's important to be able to implement some of these ideas in code, especially with your data.

Now, it doesn't mean that everybody has to go back and learn all of this, I do think it's helpful in order to determine where machine learning and AI can be most useful for us as researchers. But I also think a great option, especially for those of us in academics who have a lot of collaborators who are PhDs in these fields is to try to collaborate with those people locally in your departments and discuss what your vision is and what problems you're trying to answer and get their inputs and insights into that.

For those who do want to tackle it head on and try to learn things, we live in this amazing era where there's a lot of really great resources on YouTube and Coursera and edX and everything to really brush up on key machine learning and statistical concepts. And for those who to take it one step further, I think doing some postgraduate training in that, even doing a PhD in that is also an option for the really serious-minded.

Dr. Ashley Nassiri:

Thank you, Dr. Chowdhury. I think that sums up some of the practical advice that we have regarding this topic for surgeons. At this point, I'll go ahead with a summary of our discussion. Artificial intelligence leverages massive computing power to process information and learn patterns from a data set in an effort to develop an algorithm with predictive capacity for a specific question. While statistical analysis of a large data set can also identify models with predictive values, statistical modeling requires human input on variable selection, while AI can process data and identify patterns without human input.

AI is broadly divided into two subgroups, including natural language processing and machine learning. Machine learning is further divided into supervised, unsupervised, and deep learning. While the theoretical practice of AI holds endless promise, practical applications of the technology require thorough and accurate databases. Usually, they need to be larger than what we use on a day-to-day basis in otolaryngology. The quality of the data used to develop the learned algorithm is critical, and that leads us into the limitations which require high quality data for an accurate algorithm development and human input as a checkpoint. Dr. Chowdhury, thank you so much for being here today. It's been an absolute pleasure interviewing you. Is there anything else that you would like to add?

Dr. Naweed Chowdhury:

So, I think the take home point is that really just to remember that AI and machine learning are tools and they're not really truth machines, even though we want them to be. There's a place for them along with more traditional methods of research. Like I mentioned before, at the end of the day, I think a well-designed study will give you more insights into a specific question with less patients and less work than even the best ML and AI models.

While there is a lot of AI hype these days as there has been in the past, it's still our job as clinician scientists to keep focusing on good questions and good data and be critical of our findings, especially if they were generated by artificial intelligence. On a side note, I want to say congratulations, Ashley, on almost being done with residency. We're about one week away from graduation. We'll all really miss you here at Vanderbilt, but I know you're going to do amazing things moving forward.

Dr. Ashley Nassiri:

Thank, Dr, Chowdhury. That's very well put. Well, folks, that wraps up another episode of ENT in a Nutshell. Thanks for listening, and we'll see you next time.