# Transforming Healthcare with AI-Hype-Hope-Evidence

July 18, 2025, 12:29PM

1h 2m 56s

**Kamat, Deepak M** started transcription

**Kamat, Deepak M** 0:37

Good morning. It's 7:30 in San Antonio and so therefore it's time to start your grand round. This is the first grand round of our new academic year 202526. As always, the CME code is in the chat box and we'll keep repeating it.

Every few minutes so you will have it. Please complete the evaluation at the end of the grand round so we can provide feedback to the speaker. It's my great honour and privilege to introduce this morning's grand round speaker.

Doctor Dinisha Kuditipuri, who is a Robert F McDormand Endowed Chair in in Engineering, Professor in Electrical and Computer Engineering, Computer Sciences and Founding Director of the Matrix AI Consortium at the University of Texas San Antonio.

She received her B.Tech from University in India, Miss. from Wright State University in Dayton, Oh, and PhD from the University of Texas, San Antonio. She's nationally recognized leader in neuro inspired AI designing systems that mimic how humans solve.

Problems while using minimum energy. For the past decades, she has institutionalized AI research at 2 universities, initiated the country's first MDMS program in AI, and has National Centers of Excellence focused on human well-being and energy efficient technologies.

She leads large multidisciplinary AI teams, including hundreds of scientists with transparency and adaptability, fostering a culture with innovation and collaboration drive. The title for her today's presentation is on the screen Transforming Healthcare with AI Hype, Hope and.

Evidence. Dr. Kutitbiri, thank you very much for accepting my invitation and the floor is yours.

**Dhireesha Kudithipudi** 2:32

Thank you Dr. Kamath for this invitation and good morning everyone. This is my first talk for a grand rounds and what I thought would be good was to give a quick or a

brief overview of what we do at the Matrix Consortium.
since we are going to be one institution very soon and maybe this resource might be of advantage to you before I talk about the technology itself.
I have no actual or potential conflict of interest in relation to this program or presentation any of the content that I'm showing.
So the mission of Matrix AI Consortium is really to think about inventing, deploying and advancing AI at scale.
With the primary goal of building technologies that advance or promote human well-being, and quite often when we say human well-being, health is of prime importance, but in general it could be any techno, any aspect of human well-being. That could be boosted or accelerated with these technologies and in order to accomplish this, we really take a transdisciplinary approach. So in the consortium we have AI researchers.
ML Engineers, Computer scientists.
Physicists, but most importantly, we have domain experts, whether it's the clinicians or whether it's the geoscience researchers, neuroscience researchers, helping them solve the complex problems using AI.
And it's beyond a single institution. We have, I guess this is 3 organizations now pretty soon, but we have UTUTSA, UT Health, Southwest Research Institute and Texas Biomedical Research Institute as partners.
And we have a total of 87 scientists in the consortium, out of which 40 or 45 of them are core members of the consortium.
And we have identified 4 thrusts within matrix that we focus on. The first one is how do we augment human capabilities and this is really our community facing thrust and.

**Akbar, Asra** 5:09
Yeah.

DK

**Dhireesha Kudithipudi** 5:10
The second one is more thinking about helping researchers who are not necessarily working in the machine learning domain in identifying the right technologies for deployment in different domains. So this is really a translational thrust and we have a large group of scientists in this thrust.
And the third one is neuro inspired AI, which is more our basic or foundational

research thrust. We are really building the next generation of AI technologies using inspiration from biological processes.
Or a lot of it is looking at how humans process information, how we make predictions, how we learn in different environments, and how we do that in the most efficient way.
But most importantly, when we are deploying any of these technologies, we have to think about whether the solutions we are building are trustworthy, whether the solutions are robust, fair, unbiased, whether the the models are interpretable.
So we have a thrust on trustworthy A I and ethics is a cross cutting theme across all of these. And today the consortium has centers of excellence in each of these thrusts supported by different agencies.
So we have about $45 million in funding portfolio with 30 plus research labs and as I said, 5 centers and students engaged from multiple grad programs and about 500 students from different research labs who are part of the consortium.
And we have a I compute infrastructure. I think this is a modest infrastructure. If you think of the resources that some of the computational researchers would have, we have DGX servers, Lambda servers and we have a very niche infrastructure that we recently got funding from.
National Science Foundation. We are the national hub for a unique AI infrastructure that's called neuromorphic computing and we are going to get the hardware later this summer.
They just shipped it from Germany and we're going to get a partnership with a few other industries in order to get this kind of infrastructure.
OK. So that's kind of a brief overview of Matrix and I'm happy to have conversations offline about any interest about what Matrix is doing and how to engage and so forth.
So for the rest of the talk today, I'm going to focus on kind of where a I is today in a very specific aspect of how algorithms are doing today and what are some of the more recent advancements and some opportunities specifically in paediatric care.
Yeah.
So if you look at the trends that are out there, I think this is publication or data points from Harvard School for Public Health. It shows that there is a.
Projected reduction in treatment costs. When using AI for diagnosis, you could get like a 50 percent reduction in treatment costs. That is a significant reduction, high cost efficiency, so.

Why they're not consider this in the loop in the for diagnosis. And similarly there is a projected improvement in health outcomes up to 40%. I'm sure there are a lot of caveats to these projections.
But in general, there seems to be a huge opportunity in achieving some savings by using these tools and there is an estimated market.
Um.
Rise that you would see for healthcare in and that is a significant investment that is coming in. So it's really, I would say these are not ordinary trends and I think this is only going to increase in the coming years.
So it's very important to think about AI not as a passive tool, but how do we actually actively engage in the diagnosis or in the prediction or in?
In your workflow, how to actively use it. So when I say passively using it as a tool, I I am sure some of you are using ChatGPT or some form of large language model.
Of large language model today and maybe querying it in some forms, but is it part of your workflow where it can help you in more active ways? I'm.
I think most clinicians are not doing that at this stage and I'm sure for several reasons beyond the technical flows that are not designed to support it, it's also.
Infrastructure support and so forth, but it's very important to think in this day and age to actively use this in the design flow. So at Matrix we kind of look at different projects that are focused on.
AI centers of excellence that are related to health. I'll give example of one of it in the latest slides, which is the chatbot we are building for the trauma research supported by the state where we are collecting data, trauma data for the entire state of Texas.
And there are a few other projects where we are trying to build tools for physicians and clinicians in adopting AI in their workflows through AIM Ahead and other projects.
But I just want to give you a context in terms of the hype that is there about AI. So several of you might be familiar with macro burst videos if you look at this.
Video in here you see that this animal is able to really adapt to the changes in its environment, come up with creative solutions to get to the destination.
And it is continuously improving on how it is doing this. This is a complex environment, whereas if you think of modern day AI, it is not able to adapt in this form this kind of flexibility and real time.
Adaptation uh for different tasks is not there so.
If you think of current AI models, whether they're AI agents or not, they're really

incapable of such continual learning or adapting to unforeseen challenges in real time. And a lot of times there are huge data sets that are needed to solve these problems.
So we did this project with DARPA. This is almost like five years or four years old now, where we were trying to solve this problem of how do you build AI models that learn continually in different environments.
So um.
I hope you're able to see this video. So we built this new types of algorithms that are able to adapt and this is deployed on this drone in a photo realistic simulation where it is able to forage through this dynamic environment and try to reach.
To the, you know, people who need rescue operation. So in search and rescue operations you can perhaps consider these kinds of algorithms. If you look at the current ones, this kind of adaptability is not there.
And of course this has to be done in in a matter of milliseconds. So which means latency is a key aspect as well, not just adaptability, but how quickly it is able to do it.
So these were some early attempts in addressing this problem and since then we have advanced the algorithms quite a bit in our work and there is a lot going on in the community in order to think about how do we design a I that learns continually similar to humans a life.
Long learner.
2.
So for the next part of this talk, I'm going to show some examples in settings that might be familiar to you in how AI is.
Being used or can be used and potentially used.
So very early on, I think one project that we worked on was with the city of San Antonio, where during the peak of COVID, they wanted support on trying to see if AI could help them in.
In understanding the patterns of, you know.
The mobility pattern, sorry, in the city and based on that, if we can predict how the COVID spread is going to be impacted within the city and these were some things we have done.
On almost a weekly basis early on and then it almost became a daily report that we gave to the City of San Antonio. These models were built by graduate students in matrix to make those predictions in like you know, if you increase the mobility by this factor.

What would be the effect if you increase by 50 percent, 75% and so forth. So this was I think helpful for the city at that point of time to get have a real time dashboard and have a model that.
Make these predictions.
So if there are situations where real time decision is needed with different types of data in your workflows, I think this is something Matrix can support it.
The other thing we have done more in advanced technology is like you know quite often, well these especially in during COVID-19, you have these different strains that have evolved, right? Like you have the alpha strain, the beta strain, the gamma and so forth.
What happens is if you design an algorithm, like when we designed it, the algorithm was looking at certain strains, but a new strain might have.
Other features that need to be embedded in the algorithm. So we have started looking at the continual learning algorithms that could help in predicting these kind of changes without forgetting what it has learned in the past.
Because a lot of algorithms, they're very good on solving the current task, but they forget what they have learned in the past and this is called catastrophic forgetting or catastrophic interference problem in machine learning community. So we're looking at algorithms that would solve this and this is a.
A great case study for it in how you can build dynamic architectures so that when the strains or when these viruses are evolving, you could use more modern architectures to solve that problem without losing the information you have.
Learned in the past, similar to how humans do it.
And if we look at other ways, if you look at this cartoon, I think.
There are different sources of information from which you can.
Learn and try to get a full picture of a patient's health. So here you see the diagnosis codes from the doctors, or you could have textual notes that come in from the.
Clinical data or there is imaging data of different modalities. There is cell imaging data and there is also these nodes that are written or we would call this unstructured data.
And you can think of other, you know, EGCG, you can have claims data, you can have genomics data for this patient. So these are what we call multimodal data.
As in like you know the data is in different formats, it's coming from different sources and you use this information and fuse it with multiple algorithms to understand or diagnose the patient's health.

And for example, what I'm going to show in the next slides or what you might be using something called ChatGPT is using one modality or two modalities of information. I guess it's using three right now. It's using text, voice and image data. So.
If you have this multimodal input, the algorithms become lot more challenging in order to get the right output. It it's the algorithms don't get challenging, it gets challenging to design the algorithms.
But um.
It's also important to understand where the data is coming from and if the data is being fused in the right format to make the right decision. So the sources of all these data formats becomes very important.
So we kind of did this a single modality one in the trauma project that we are working with the state, which is the TRC for funding for this. So Doctor Kutub from Matrix is leading this project.
So you can see that. Let me play this video. I hope you're able to see it. So a trauma assistant was built where it gives some more nuanced information of what action should be done in an emergency situation.
And it runs in like, you know, a matter of milliseconds. So here is an information. A 28 year old male falls from a height of 10 feet and presents with numbness and weakness in both legs. And he has normal sensation in his upper limbs, but is unable to move his legs and has absent patellar reflexes.
So what action should you take and does this patient require a level 1 trauma center? So, so when we run that.
Hope the videos.
So it is giving guidance on the actions that could be taken like you know, immediate actions, whether there's immobilization, assessment, vital signs, monitoring, establishing IV access and then referral through the trauma centers.
And the level of care that is needed. So here in summary, ensured rapid transport transport to a level 1 trauma center is initiated while monitoring and stabilizing the patient as per guidelines. But then you can think and imagine further where it.
Use like dynamic information on what what is the trauma center that could be.
I mean though the I think ground level crew would know what the nearest trauma center. I think you could have access to the usage, the wait times, everything dynamically in a visual so that you you and the traffic that you would have all this kind.

Of information in a single flow where you could be guided to go to the right trauma center and get the care immediately. So this kind of application is possible where you have here a singular modality of information. This is not multi.
Model, but multiple other inputs can be given to this assistant and that's what this was a pilot study and we are advancing this with now collecting information of all the data from the state of Texas and Amina is looking at advancing.
This chatbot to do a lot of these additional features that would support.
Care in these situations.
The other tool I thought you should know that got released literally last week or.
Yeah, 10 days back. It is called Medcamel. It is again a language model, but it's particularly trained with medical data to support.
Diagnosis or a lot of pre analysis and care. So I wanted to show some of the demos for the because I think if you want to explore or just, you know, even play with.
Uh.
And see what it is capable of. This would be a good tool to explore. Sorry.
So, so if you have a use case, you can define the use case. Let's say you want to generate a text report and you want to get an evaluation, but you want to get performance at a certain level like you know.
Do you want? You can set the performance that you would want, whether you want a high performance or a low performance, and I'll tell you in a second why that would change the outcomes in some cases, but so.
You could look at the data you you have the input in here with multiple images and then you're looking at certain cases and trying to see what based on your setting of what the performance is.
You get a certain response.
So the reason you're tuning that performance is that you're able to process it faster on a slighter model, and this is particularly important when you're thinking about.
Embedding these kinds of tools in your workflows because they're they consume lot of energy and which means a lot of money. It costs a lot because of the computer resources they need. So you would want to reduce that burden and.
One way is to kind of use lighter models and if let's say 99% accuracy is not needed in a certain situation, maybe 90% gives you the confidence to go in a certain route. I think you could change those or.
Adjust those performance metrics in this model. So it made a big splash in the last one week. It's a 4B model is what they call it, so I would encourage you to explore it.

And they kind of show it how like you know they they demonstrate it like you know how it looks like for an intake optimization with this tool. So there there is a built-in agent that is going to ask questions to gather.
Uh information about the patient. So it kind of generates an intelligent previsit report that would help you as a Uh provider or a caretaker Uh giver.
And I'll give you a detailed evaluation of the previous report.
So this is just a simulation of how it runs with.
What it would ask the end user, So what they're visiting for and based on that, I think what they feed in and the questions that or what the queries that gives to the patient, it kind of gives you a generic summary.
And you can see that some details that are not mentioned and all of this additional information that would help you to.
Diagnose or help you to provide the right care for this patient.
So obviously I'm sure you have lots of questions on, you know, whether.
The output that this tool is providing is reliable, but I think this is where collaboration with humans experts is very critical.
Because based on the data, the generalized databases that it is looking from, the outputs that could it could you sometimes may have mismatches, may have incorrect information and that is where the expert is able to quickly say well that.
For example, those triggers do not fit with this portfolio. How and you would call it out and the model is able to adapt. So I would say use this can be used as a secondary or a surrogate.
Still, because of some of these concerns on how the data is being processed, the source of the data, but I think compared to several other models that are out there, this model seem to be doing really well with medical data.
Some other examples where you can think of where a I could support is let's say a lot of models that are built today or a lot of analysis that is done is done on adult medical data.
Not as much. Let's say pediatric data sets are not there for a lot of diseases. So what? I mean, of course it would help to evaluate on pediatric data sets.
And get the right diagnosis. But there are also AI models that could transfer the knowledge that is there from the adult data sets and try to make diagnosis with pediatric data.
So.
This is again something that is still in the works and a lot of models. For example,

what we have seen just now with Metgamma, one of the models, the large language models, there is a lot of personalization.
But there is also opportunity to translate complex information with a language that patients can easily process or understand. So the tone the.
The terms, the jag and all could be reduced and when you're communicating with kids and also there are tools including ones we are building on that provide real time monitoring.
And decision making and also there are researchers and matrix are looking at how do we ensure that what we are building is secure and information is.
Sensitive information is protected. So for example, in the case of trying to predict the variations, you could think of AI that is processing vast amounts of data basically.
And it's pinpointing certain signs of a disease that might otherwise go unnoticed. So this will help in early intervention and all of that. But well, what would you do if you didn't have the data for the specific disease you're?
Looking at or if it is really scarce and difficult to collect data for that specific situation, so then it's really challenging for the model to make the prediction. So there are newer techniques like transfer learning and the continual learning models. One of the feature is transfer learning.
That could leverage existing data science and adapt this knowledge to the unique situation or characteristics that you were looking at and try to give a better diagnosis or pre-diagnosis or so forth.
And similarly the large language models where I said you could have an opportunity to think about how to generate information that is more digestible to the end audience in a say in the in this case.
You could have the model, let's say the. If you look at the example on the right here, you see that there is a specific information about.
This condition that was provided by the by the doctor and you could translate that in a different way to an adult and a different way to the kid and I think you do that on real time on a daily basis and.
You could have the models also do this not in real time, perhaps before during the entry free intake time or post post visit care in in different situations you could think of.
Using this type of individualized and age appropriate or customized responses.
And you could. The other example I was saying was where you could do the real-time monitoring and prediction. There might be situations. I think there are smart

health homes is like a classic example where we can think of using these kinds of sensors.
That monitor vital signs and use that information to make a certain prediction if there is an anomaly in the signal and.
You can think of using this in different contexts, not just within a smart health home on a regular basis. Like you know you have all of us have some kind of smart device now monitoring our critical signals, but then fusing them in trying to make the predictions. I think you need really.
A light device that is able to do in again a matter of less than I think 10 milliseconds. So the neuromorphic computing technologies we are building help in this kinds of analysis type and prediction.
And the last one I want to highlight is about privacy. So a lot of times these datasets, whether those are proprietary datasets that are designed.
In your clinical setting or design or or something that you're using out of a general database, you might want to secure this data and I know there are certain standards that you follow with the HIPAA compliance and all of that.
But in general, instead of having a single cloud infrastructure or a server where you store all of this data, you would want to think about something called a Federated model.
In in how this data is being saved and how you learn in from this data sets that are stored in this cloud so.
What a federated learning infrastructure does is there is distribution of information across multiple clouds, mini clouds, fog. There are a lot of these terms that are used in the compute community, so you could have that distributed and you're just not learning from a single.
A single data source, but then you're learning from these different data sources and making the decision whether it is diagnosis or treatment or so forth. And this is something that is commonly used nowadays when you're trying to chat or messaging or.
In any of those tools, federated learning is naturally embedded in your smartphones or devices, and these could be extended to more of the medical data to ensure that the data is more secure so the end user only can see data that is relevant to them.
Not have the visibility of the generic or other data, so it is more secure in this using this federated learning approaches.
Um.

I'm sorry, I think this is again an extension of another example of Met Gamma that I wanted to show in diagnostic support. So you can see the prompts that are given here. Again, if you're an expert radiologist.
Please describe the findings for the about chest X-ray and you see the radiologist impression on the left and and you could also see on the prompt on the right what what do you think is the most likely cause in this image?
And what would be the next steps that you would suggest? And this is like you know more a quick response to see like well consulted dermatologist for this additional steps. So you can think of using this in different ways in your workflows.
It could, you know, integrate multiple images again here, both 2D and 3D scans. It can generate comprehensive reports and there is visual question answering is what you see on the right is called the VQA technology.
Where you are giving that image and asking it to analyze it and and you can query about the image and you can also think about using it in clinical decision support. Again, this is a caveat with it comes with a big caveat, right? You want to really think about.
The data sources and inbuilt models and the level of accuracy this model is given. But since it has just been recently released, I would look at some of the example scenarios and try to work with it and see how it does without giving any proprietary. Data.
Um. And similarly, I think I've already shown this, so I'm gonna skip that so.
As you're thinking about adopting or integrating more actively any of these models in your workflows, there are multiple things to think about. Well, the model does great in terms of performance.
But does it work in all the situations? Is the model fair and why would this happen? Why would the model be not fair? The reason is.
It might be strained only on a subset of the data. It might be trained only on in certain geographical regions. Maybe does not translate well across the country in different scenarios, different geographical regions, so.
It's very important to be aware what was the training data for the model and what are its limitations before using a model so.
So that it doesn't perpetuate incorrect information or you know, it might not amplify unfair predictions in some scenarios, similarly with.
With the output that is generated from this tool, if the data set it is relying on is not a robust one, it can generate incorrect information and this is where human AI teaming

and collaboration is very important to correct the outputs of this model.
As an expert, you're able to do that, so it's important to not fully rely on the outputs that are viewed from the models.
And and I think this is probably more of concern in human rich environments and this is a.
Or in mission critical environments where if you're using AI, you have to think about the accountability and I'm sure you are lot more.
Cognizant and trying to embed this in your workflows more actively as to how to reduce any kind of compliance issues. So when you're using these types of models, what if the model has predicted?
Something wrong and you have relied on that information and has given the treatment based on that. So how do you how do you ensure that?
This this kind of impact is reduced, so it's very important again to tie to the earlier things, whether the model is fair, whether it is giving the right information, whether it is trained on the right data sets. Is the model matured enough to make accurate predictions?
So these kinds of information, this kind of information is helpful whether to in deciding whether to use that model because of these kinds of issues on accountability and so forth. And the other point to think about is whether the model is.
Even trained on open data sets or are there going to be some IP ownership issues on where the model is getting its data from or where it was originally trained on? There are models which are originally trained on and then you just deploy them in your scenario. They don't need that additional.
Training, but there are other models that are continuously training and updating based on the information that you're given. So in either of these scenarios, it's important for you to think about where that data is.
Coming from. So I guess you see that underlying theme for all of these is to be knowledgeable about the data that the model is trained on and the source of the data and whether the reliability of the data.
It's your fault.
So if you're more interested in trying to understand some of these, we do have sessions in tinkering with an AI scientist or we do boot camps. Sorry for MD students in MDMS and AI program. We were doing boot camps earlier, so if you're interested,

we would keep you in the loop when the next training.
Session happens. So thank you so much for your time and I'm open to questions.

**Kamat, Deepak M** 44:21

Thank you, Doctor Kuditipudi, for that wonderful presentation on transforming healthcare with AI. Let's see if anybody has any questions. So let me start with me. So most of us are sort of scared of AI. We have not used it. Where do we start as a novice in AI in medical field is the.
Gamma is a good, good place to start.

**Dhireesha Kudithipudi** 44:44

It's a good place to play. What I would say is I think no matter which profession we are in, everybody should be an A I generalist in this day and age. So what does that mean is to be aware of some of the tools.
Um. And actually try to um.
It takes very, uh, minimal effort in, you know, um, playing with these tools because the user interfaces have improved significantly.
Definitely. I think large language models is something I think that is, if not today, pretty soon going to be embedded in your workflows and thinking. So it's a great tool for you to play with and Metgamer is yes, definitely something I would play with. And and I think you can immediately see, Oh well, it does really well in the situation, but it does terrible here. But you're giving the guidance to have better models that will help you. But also I think in general trying to understand the basics of how these models work.
You don't have to understand the full intricacy of the exact models, but I think as I said, being an AI generalist, trying to understand well what does the pipeline look like as in.
How the data is being taken into the model, as in what kind of pre-processing is done, how it is fed into the model, and then once it is fed into the model, what kind of outputs the model is giving and how reliable are they?
So I think from the AI developers itself, there are things that we should be doing in reporting these statistics on the the situations in which the model should be used and should not should not be used, but I think from.

**Kamat, Deepak M** 46:28
OK.

**Dhireesha Kudithipudi** 46:46
To answer your question, I think it's really important for everyone to be an AI generalist and this is a good tool to start with. I was actually going to compile a list of tools that would help you and I can share this after the talk and that might if anybody is interested in playing with.
Series of tools.

**Kamat, Deepak M** 47:06
Yeah. So there's a question in the comment box by Michael Lipsitz. Very interesting. Thank you. Will AI eventually completely replace the physician?

**Dhireesha Kudithipudi** 47:17
I think um.
This is where the hype is, I would say, right? I mean, yes, AI is able to do some of the tasks that are very streamlined or you know, where the data is streamlined, where the data is structured, where the environment is structured. It doesn't do very well in unstructured environment.
Still, it doesn't learn continually or adapt as well as humans do in different situations. We are far from that. There are a lot of claims that you know today we have an AGI is going to come, which is artificial general intelligence.
Or you might hear super intelligence. You hear all these terms that are thrown out there that it's going to happen in the next couple of years. And a lot of researchers who are in the field do not think it's there because of.
Several underlying fundamental problems that need to be solved with these models. And again, I think we humans, we are 3D structures, right? So we move in environments and we do a lot of things beyond just processing this.
Textual or information from the images. So if you think of robotics or other.
Tools that are going to support the clinicians or physicians, they are also, they also need quite a bit of evolution to even think of working in these unstructured environments. They work very well in factory settings.
Where everything is streamlined so.

I I personally do not think that they are going to replace the physician, but there is going to be a lot of teaming up that's going to happen. So This is why I would encourage everyone to be an AI generalist.

**Kamat, Deepak M** 49:26
Thank you. There is a Doctor Hedge and Beller. Do you have a question? I I I see you have raised your hand.

**Hatzenbiler, Zoe** 49:30
Yeah, sorry, Dr. Kamat. This is Doctor Asanasan. I'm here with my residents on the inpatient service. Dr. Hedberg was actually the one that logged in for us. So again, my name is Doctor Chachwanasan. I'm a pediatric hematologist, oncologist and I also work for the physicians practice.

**Kamat, Deepak M** 49:37
Oh, OK.

**Hatzenbiler, Zoe** 49:50
As an administrator, you mentioned the issues with regard to HIPAA and utilizing the federated model. There are currently AI platforms in use for physicians kind of collecting.
Data with regard to the interview process. I think I can mention a bridge here. They have mentioned that they erase the data after a period of time as AI generalists. What should we be looking for? Really the data model of how to?

**Dhireesha Kudithipudi** 50:13
Mhm, mhm.

**Hatzenbiler, Zoe** 50:26
To make sure that we're using secure mechanisms? Or should we feel fairly comfortable when people say that they're erasing the information before anybody outside may be able to access it?

**Dhireesha Kudithipudi** 50:44
Um, very good question. So.

I mean when someone is saying they're erasing data, well, I think perhaps, I mean the way they have designed it, perhaps it is like a three-month time frame, six month time frame the data is erased. But what happens is it really depends on the underlying architecture it happened with.
Some of the several of you might be using messaging platforms like WhatsApp. I don't know if you're familiar with it. There are some of these messaging platforms where the user was thinking or or informed that.
Data is being erased, but in actuality, because of how architectures are built behind where the models models are deployed, the data could be residing in some.
Or some artifacts of this could be residing in some form in one of these architectural building blocks. So how do we know as an end user? Sometimes this is really only experts can see that and sometimes these things come up.
Out, you know, much later than you would like to know about. So what? I mean, I'm not trying to scare you here, but I think it's very important to ask questions about the sources of data and the architecture they're using.
And the confidence rate at which they think this data is going to be data is being erased and ask them to show you the flow, how data is being erased, I mean probe on it further to kind of see what is the process.
And what I would also ask is like have consult with an AI specialist, not just with the vendor you are working with, but consult with an AI specialist. I mean there are several.
That we could refer to from the Matrix Consortium for example, who may not be expert in the exact technology you're using, but they could help you ask the right questions also. So I would use that layer approach.
Trust, but I think be cautious.

**Kamat, Deepak M** 53:13
Thank you, Dirisha. Then there are comments by Doctor Jones, one HIPAA. Let's see.

**Jones, Woodson S** 53:19
I can go ahead and comment Dr. Kamad. I think you know I interesting. I presume you guys know about open evidence again out by Mayo and in collaboration with New England Journal of Medicine and now actually JAMA Network as a tool. It's free and available to clinicians. You have to have an MPI.
Number to use and it's HIPAA compliant now as well. And I actually just ran your

question through on the trauma patient that you had and he actually gave a rather remarkable answer, very thorough with the references that you can then tap on. So open evidence is actually.

**Dhireesha Kudithipudi** 53:58
Mhm.

**Jones, Woodson S** 53:59
Actually used. Now again, you talk about a place to start. I warn the residents routinely when they use these tools. Prompt engineering is still very, very important, and particularly in Pediatrics if you miss a key element and age is one of them.
To put into these tools, you will get very convincing wrong answers. And so that's still one of the biggest concerns I have as we begin to introduce this to our residents is if you don't have that knowledge and that background to see.

**Dhireesha Kudithipudi** 54:18
Mhm.

**Jones, Woodson S** 54:29
Or to when you get that wrong answer to at least turn your head and say, man, that doesn't sound quite right, you're going to miss it because it again, it's so confidently wrong that you will make errors. And so prompt engineering I think is is such a key element to begin with when you train people how to use these.

**Dhireesha Kudithipudi** 54:39
Mhm.
Right.

**Jones, Woodson S** 54:48
These tools.

**Dhireesha Kudithipudi** 54:51
Right. There are bots that are being generated to help with the prompt engineering as well. Hopefully that those are made more available for different platforms pretty soon because I completely agree with you.

Missing keywords can give you completely different answers and or inaccurate information that shouldn't be relied on.
And the other thing about references, sorry to just add to that earlier tools like at least one year or six months back, it would give you references. So you would think that well this information is reliable. Lot of times those references are not existing publications.

**Jones, Woodson S** 55:17
Right.

**Kamat, Deepak M** 55:18
There's a question, but.

**Dhireesha Kudithipudi** 55:37
It would make up references with high confidence, so it's very important to check the references that are being pointed to.

**Kamat, Deepak M** 55:38
Yes.
I I agree. I ran into it as a editor. Actually, I always check the references because those are not right. Anyway, there's a question by Doctor Risby. How good is the news of AI Hospital in China?

**Dhireesha Kudithipudi** 55:48
Yeah.

**Kamat, Deepak M** 56:03
OK.

**Dhireesha Kudithipudi** 56:04
Uh.
How good is the news?

**Kamat, Deepak M** 56:11
Of AI Hospital in China.

**Dhireesha Kudithipudi** 56:13

I I don't know if I read this specific news article or information, but I know there is a lot more active adoption of tools at all levels.

Think.

In some countries, so.

I think there is perhaps some significant amount of training data that you would have these models being trained on and over time they would become stronger and more accurate, but.

Um.

Yeah, I I think there are several aspects within a hospital that you could use smart technologies and which I could see happening in US as well. But it wouldn't be a complete AI hospital as I said like I think the human element is still.

Quite significant in here.

**Kamat, Deepak M** 57:18

Thank you. Any other questions, comments for Dr. Kuridipudi?

**Jones, Woodson S** 57:23

I I have one other. Again, I think one of our biggest challenges as I've been meeting with Doctor Sankry and also our CIO here or Vice President for Information Technology here at UTL San Antonio is a lack of a ready access to a enterprise or private level LL.

You know, we've been trying to use collators as some people will label them like you.com and other ways to improve efficiency, but we're still limited even though we have a very powerful list of 23 plus LLM that we can use.

**Dhireesha Kudithipudi** 57:43

Mhm.

**Jones, Woodson S** 57:59

They're not enterprise, they're not private. And so it really limits what we can load up into our sources or what people call a rag or interpretation or what we can ask as far as questions and data we can share. I mean, do you see things at your level coming out of the UT system or in our partnership with UT?

**Dhireesha Kudithipudi** 58:08
Mhm.


**Jones, Woodson S** 58:19
UTSA which are merging with UTSA which will move us towards having enterprise or private level LLM's for doing some of our work cuz it is limiting our office in the office of GME with moving forward and I'm watching Stanford and other people I talk to all the time.

**Dhireesha Kudithipudi** 58:32
Mhm.
Yeah.


**Jones, Woodson S** 58:38
Tell them what we're doing, they take and run it, and they're already ahead of us within months because we just seem to be limited on the resources side of this equation with AI.

**Dhireesha Kudithipudi** 58:50
Yeah, I think there are two very good points to think about here regarding the private instance of LLM. I'm actually having a meeting with the CIO next week with your CIO. So we're trying to think about that structure and facilitating tools for.
Um.
For physicians and researchers in UT Health, so I'm sure we will. This is of prime importance and the VPRS office is already thinking about it. So there's just build this team to think about how do we help researchers.
The other part is think about the data that you have or the assets that you have within your groups. I personally believe in open sourcing, you know, whatever models we are building and everything.
But there are valuable assets that you have within your units, whether it is data that is collected from the patients or even in the intake forms, there's several things, right, that are really valuable assets.
So if you can think of what are the digital assets that you would have within your units and I would not give it away for free to any entity, I would really think about it.

Um, I know Stanford. Um uh.
I think probably you're talking about their Research Center on health that they are being very smart and pretty aggressive in collecting data from multiple States and trying to deploy their models. I know we could do better and we could move faster. But independent of which group is using these assets, I would think for you as a clinician in your or if you're a chair of a department, what are the assets that you have within your unit?
It may not always be a cost associated with you giving those assets, but there could be other negotiations that could happen. We're not, at least in academic settings from the engineering side, we are not generally trained to think that way, but I think it's very important to.
Add value, I mean get value out of those assets that you would have.

**Kamat, Deepak M** 1:01:21
Thank you. One last questions, a quick question. Can we get an AI tool to automate billing? So billing for the insurance companies or the for the patients, can this be automated?

DK

**Dhireesha Kudithipudi** 1:01:33
Uh.
I think there are multiple tools. I know I don't know how it fits within your workflow, but I think this is this aspect of the problem is a lot more. I wouldn't say easier, but I think it is being adopted in multiple.
Industries already, so there are tools to do it. OCR technology has matured so much, mature recognition technology, even if it is all like, you know, legacy paperwork that you have, it could all be screened and fed into models to automate the process, so.
I think there are tools to look at.

**Kamat, Deepak M** 1:02:14
OK. Thank you, Doctor Putitugri, for that wonderful presentation. Really enjoyed your presentation. Thank you all for attending this morning's grand round. I'm going to conclude this morning's grand round. Please fill the complete the evaluations which we'll send shortly.
And again, thank you and have a wonderful Friday and a wonderful weekend. Again, thank you, Doctor Kutitipedi for that wonderful presentation.

**Dhireesha Kudithipudi** 1:02:37
Thank you so much and I will share the list of tools.

**Kamat, Deepak M** 1:02:40
OK. And then I'll I'll forward it to the faculty. Thank you.

**Dhireesha Kudithipudi** 1:02:42
It.
Bye.

**Kamat, Deepak M** stopped transcription