Thank you for participating our second session of the 2023 Learning for Life Series.

Founded in 1975 by Shirley Welsh Ryan, the series aims to educate participants on important technological advances and current events. This year's program is focused on Artificial Intelligence.

Artificial Intelligence and Its Many Implications

Artificial intelligence (AI) has taken our society by storm. With the grand promise of reproducing human intelligence within machines, the implications of this rapidly advancing technology are broad. In this lecture series, we'll start with a broad overview of AI, its capabilities and limitations, and then look at its applications in healthcare and education. We'll end with a discussion of the importance of creativity to understand why artificial intelligence may always remain "artificial."

October 17th's program looked at Artificial Intelligence in Healthcare.

Summary from Tuesday, October 17, 2023's session

***Dr. Eric G. Neilson, MD**, is the Vice President for Medical Affairs and the Lewis Landsberg Dean at Northwestern University Feinberg School of Medicine. After 23 years at the University of Pennsylvania as the C. Mahlon Kline Professor of Medicine, chief of the Renal-Electrolyte and Hypertension Division, and director of the Penn Center for Molecular Studies of Kidney Diseases, he moved to the Vanderbilt University School of Medicine to serve for 13 years as the Hugh Jackson Morgan Professor and Chairman of the Department of Medicine, before becoming the Dean of Feinberg at Northwestern.*

Neilson is an active teacher of clinical medicine and has trained numerous graduate students and post-doctoral fellows in his laboratory. He has a special interest in the training of physician-scientists, has edited a book entitled What's Past is Prologue—The Personal Stories Of Women In Science At The Vanderbilt University School of Medicine to help mentor women interested in biomedical research, and helped initiate the Vanderbilt Prize for women scientists. Twenty-four of his former laboratory students and fellows have become professors of medicine, twenty-seven of his former fellows and faculty are now department chairs, eighteen are associate/assistant deans, vice-chancellors, or provosts, and eleven have been elected to the National Academy of Medicine.

Dr. Neilson is the recipient of numerous awards including being named a 2016 fellow of the American Academy of Arts and Sciences. In 2021, he was awarded a Mastership from the American College of Physicians.

Dean Eric Neilson began with an overview of the potential of AI in healthcare - - how the technology will break down barriers to access and dramatically improve quality, consistency, and efficiency. In fact, AI has been incorporated into medicine for quite a long time. Almost any piece of technology that we encounter, whether it's a blood pressure machine that automatically takes your blood pressure or

complex imaging studies like an MRI or a CAT scan - - all have tremendous amounts of AI wrapped into them, not only to operate the equipment, but to analyze the data AI has made telemedicine more effective and will continue to make our clinical work easier and less administratively burdensome. As an example, at Northwestern Medicine, the Mansueto Institute for Innovation has been hosting technology where visits can be recorded with later mining and analysis serving as a rich database. As patients, we should all expect AI to continue to bring powerful capabilities our way.

Dr. Abel Kho *received his MD from the Medical College of Wisconsin and completed a residency and Chief residency in Internal Medicine at the University of Wisconsin, Madison. He completed an NLM/NIH fellowship in Medical Informatics at the Regenstrief Institute before joining the faculty at Northwestern University. Abel is Professor of Medicine and Preventive Medicine in the Feinberg School of Medicine at Northwestern University and Founding Director of the Center for Health Information Partnerships and the Institute for Augmented Intelligence in Medicine He has served as PI for several regional or national projects including the ONC funded Chicago Health IT Regional Extension Center, the PCORI funded Chicago Area Patient Centered Outcomes Research Network, and the AHRQ funded Health Hearts in the Heartland consortium within the EvidenceNOW initiative. His research focuses on developing regional Electronic Health Record (EHR) enabled data sharing platforms for a range of health applications including high throughput phenotyping, cohort discovery, estimating population level disease burden, and quality improvement*

Dr. Abel Kho opened by reminding us that, despite the advances, we are still in the early days AI's impact on healthcare. To provide context, Dr. Kho took us back to the time when IBM's Watson was beating the best players on the game show Jeopardy. Watson next tackled healthcare and it quickly became obvious that healthcare is not the same as answering questions on Jeopardy. After many years and many efforts, IBM recently sold Watson Health to a private equity company. Healthcare is not an easy space – it is highly regulated and complex. Historically large technology companies have failed at healthcare. Amazon folded its Haven initiative a few years ago and Google Health has met many difficulties.

However, now appears to be a different time with today's methods for AI as the big difference. Eight to ten years ago, we were working with natural language processing models that were laborious. The methods analyzed text by determining the relationship of words to sentences and paragraphs, essentially mimicking how we might think. The recent realization is that we don't need these multiple steps – words and parts of words are instead converted into numeric representations to search for patterns and relations. As we have heard, the technology is not perfect, so there is still human intervention. All the major vendors use humans to refine the process, ranking and interacting with the system. We've seen models improving, from GPT2 to 3 to now 4. There are a number of interesting publications illustrating the effects of human guided trainings. And the interest in the new tools has been incredible. Netflix took almost 4 years to reach a million users while ChatGPT reached that number in only 5 days.

As physicians, like any profession, we worry if we will be replaced. Most people assume people prefer interacting with people rather than computers. An interesting article recently summarized work in this area. In studies of an online forum where both physicians and ChatGPT responded to patient questions, the patients rated ChatGPT as being 10x more empathetic as the physicians! Perhaps this is not unexpected as a computer will have unlimited time as compared to physicians.

Another interesting study was done at NYU where they worked with the data of the entire health system over the last decade – both clean and "unclean" data – notes, free text, etc. And with no training, using an open-source version of an AI tool, they were able to accurately predict important clinical outcomes, such as hospital mortality, 30-day readmissions, and other metrics.

We here at Northwestern launched the Institute for Augmented Intelligence in Medicine back in 2020. The institute has six centers and is adding a seventh. Some of our key partners are the McCormick School of Engineering, The Bluhm Cardiovascular Center, and Northwestern Medicine's Mansueto Innovation Institute. One of the centers is focused on Medical Education and Data Science enabling us to be one of the first medical schools to integrate digital health and data science into all four years of medical school. Ethics obviously is hugely important, so one of our first centers to be established was focused on bioethics and medical humanities.

Dr. Kho ended with highlighting examples throughout Northwestern. We have cardiologist Dr. Sanjiv Shah using AI to define subtypes of disease and to diagnose. Dr. Amy Paller has been using AI to look at children with very rare skin genetic conditions to provide early diagnosis. And we are using AI to look at different ways to approach diagnoses that are far more relevant for treatment. AI will affect medical treatment at every level.

Professor Maia Jacobs is the Wissner-Slivka Junior Assistant Professor of Computer Science with a joint appointment in preventative medicine at Northwestern University's Feinberg School of Medicine. She received her PhD in human centered computing from Georgia Institute of Technology in 2017 and was a postdoctoral fellow at Harvard University's Center for Research on Computation and Society. Jacobs' research focuses on the design, development, and assessment of novel technologies to support chronic disease management, such as adaptive health information systems and machine learning-based decision support tools. Her research was recognized in the 2016 report to the President of the United States from the President's Cancer Panel, which focused on improving cancer-related outcomes. She also received the iSchools Doctoral Dissertation Award and the Georgia Institute of Technology College of Computing Dissertation Award.

Professor Maia Jacobs focused on her work at the interface of humans and computers, studying how to design technology based on machine learning and AI to be usable and useful in real world contexts. We have seen how tools can help in areas such as automated diabetic retinopathy diagnoses, early sepsis, detection, improved tumor detection, faster speeds of diagnosis, and better imaging results. However, we are still lagging in the translation of technology advances into adoption in healthcare.

The deciding but sometimes overlooked factor for technology adoption is how tools are designed and then implemented. Projects within the Jacobs lab start by first deeply understanding healthcare socio-technical systems.

- What are these existing processes?
- What are the goals of the end users?
- Where is their opportunity for support?

The next step is to translate into an iterative user-centered design that involves nurses, doctors, patients, and staff to make the tools robust enough to work within complex healthcare settings.

Professor Jacobs highlighted two projects – a partnership with clinicians to improve care for people with major depressive disorders and then a new project to help pregnant people reduce stress with predictive models. Focusing on the study for major depressive disorders, Professor Jacobs worked with the Massachusetts General Hospital to study how to optimize antidepressant choices for individual patients. Right now, two thirds of patients do not see improvement in their symptoms after their first antidepressant trial and a third still do not see improvement in their symptoms even after four trials. Working with a team of machine learning and AI scientists, they developed a list of the thirty medications with the best likelihood of success for individuals. And then, within a randomized clinical trial, they studied how clinicians responded to and used the predictions. The most striking result was

how clinicians could be influenced when presented with “incorrect” predictions – highlighting the risk of over-trusting imperfect tools. Other important results were the effect of the presentation of the information – visual versus text and whether patients were included in the decision-making process.

Her newer project is focused on prenatal stress, defined as the feeling of being overwhelmed and unable to cope with very serious consequences for both the mother's health as well as for the infants. Prenatal stress is associated with and can predict prenatal depression and can also lead to preterm birth and low birth weight. With Professor Nabil Alshurafa, a team has developed AI models based on participant characteristics, demographic, self-reported data, and data from ECG sensors to predict next day physiological and perceived stress. With this information, we can provide evidence-based interventions to reduce stress.

All this begs the question of how to move this work into prevention, stopping the high stress episodes from starting. To build a real-world tool which co-designs AI with the general population who may have different levels of knowledge and background on machine learning and AI, Professor Jacob’s lab has built a toolkit to engage patients to communicate and collaborate to create this type of technology. The toolkit will soon be open-sourced, available to any researcher or industry practitioner who wants to engage the public in design. From studies with the toolkit, the most important factors were

- Adaptability - technology should continue to “understand” at a personal level. Participants want ways to help the technology better understand me as I use it over time.
- Feedback – technology should incorporate patients’ reactions
- Improvements - participants wanted a way to tell the technology to do better with the specifics of how and why

Interestingly, the work illustrated that participants had a very strong willingness to take on the extent of time and burden needed to provide feedback. data. This all loops back to what Dr. Kho highlighted – the criticality of people to making this technology work.