DIHI

VOL 23

IMPACT

Duke Institute for Health Innovation

About DIHI

The Duke Institute for Health Innovation (DIHI) catalyzes transformative innovations in health and healthcare through implementation of high-impact innovations, leadership development, and cultivation of a community of entrepreneurship. We bring innovative solutions to the most pressing challenges in health and healthcare through multidisciplinary teamwork across Duke University and Duke Health and by fostering collaborations with national and international thought leaders.

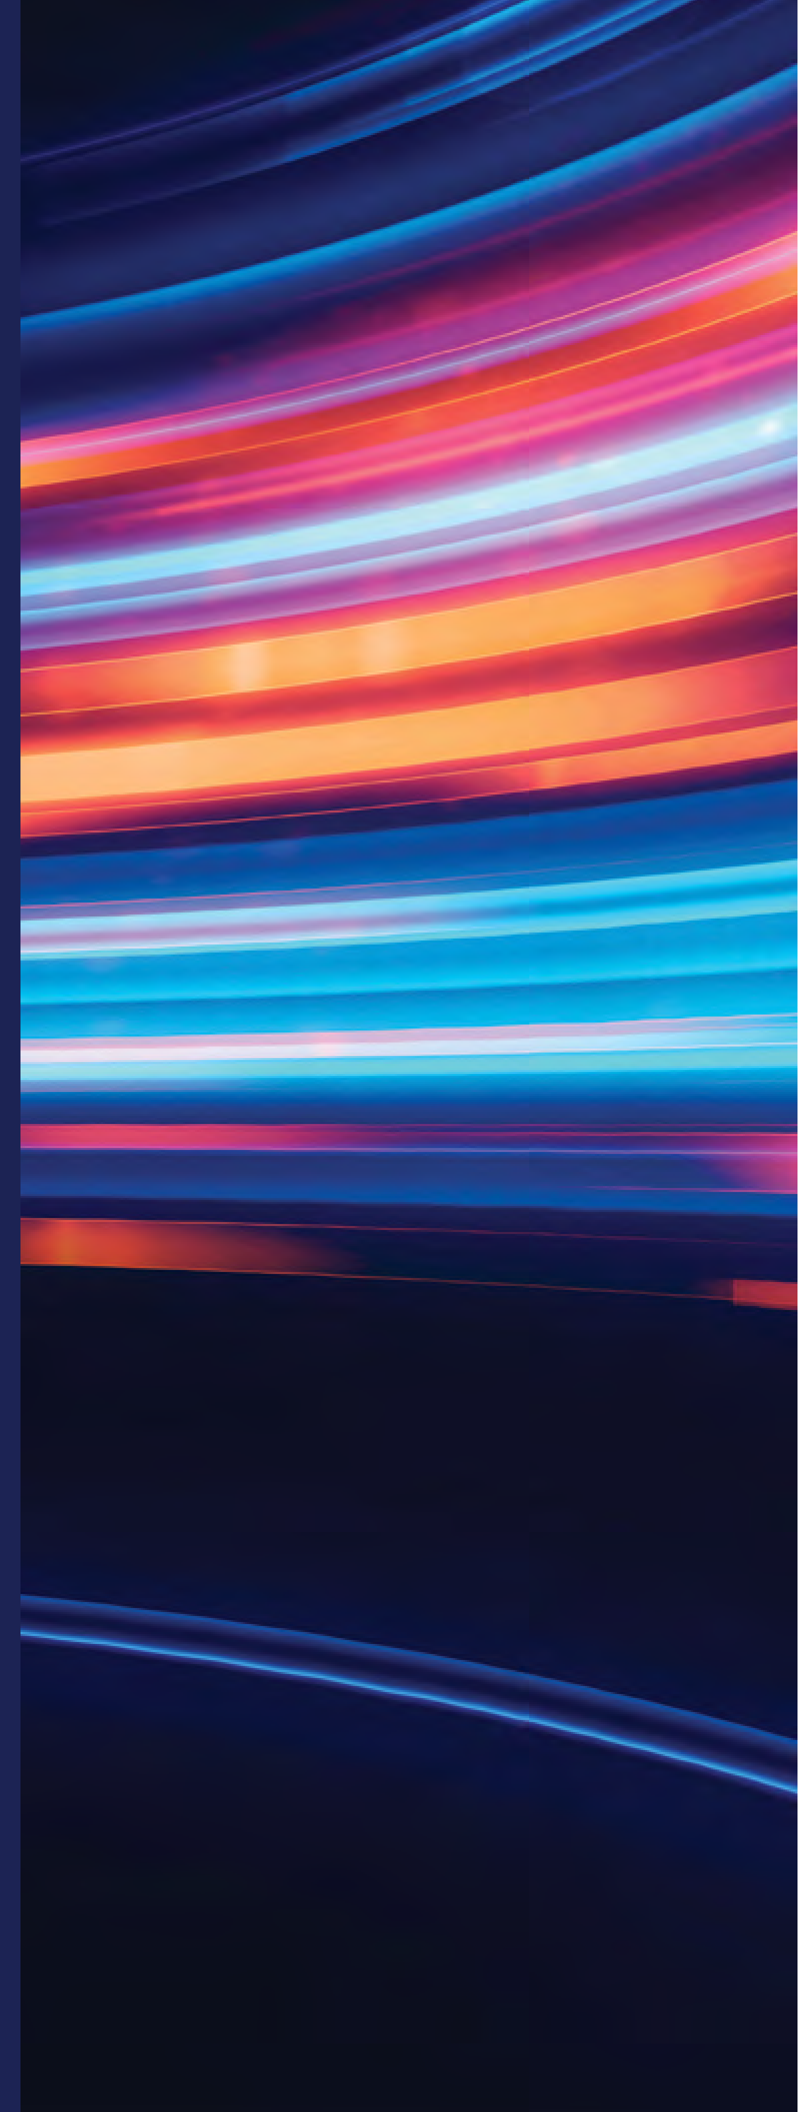

Contents

Letter from the Directors

Craig Albanese, MD, MBA

Executive Director for DIHI
Executive Vice President
and Chief Operating Officer, DUHS

Suresh Balu, MBA, MS

Director for DIHI
Associate Dean,
Innovation and Partnership,
School of Medicine

As we continue to navigate through COVID-19, new models for business and sustainability and a renewed emphasis on a culture of inclusion and belonging, our innovation efforts this past year have kept a laser focus on two important groups of stakeholders – our people and our patients. The principal themes this year were, therefore, areas that have direct impact on improving outcomes for our patients and the health and well-being of our many providers and staff who make up the clinical enterprise.

- Advance health equity
- Improve patient and community engagement and experience
- Accelerate population health strategies and solutions
- Improve value of care through novel strategies
- Grow digital care pathways and remote monitoring solutions
- Enhance care team experience & well-being; reduce workload

This report highlights some of the projects that align with the above themes that were selected through a competitive process for developing further, implementing in our clinical workflows, and in many cases, scaling across the clinical enterprise.

In a year filled with challenges, our exceptional faculty, staff, and learners found powerful opportunities to collaborate, innovate, translate data science to health benefit and accelerate solution deployment for real-world issues facing our patients and care teams. In the weeks and months when we could have hunkered down and paused work on untested processes and solutions, we instead chose to push ahead and expand the scope of what was possible. Not only did we adapt to the constraints placed by the pandemic and other circumstances, but Duke Health also led the way in fostering an ecosystem that made ideation and implementation of innovative solutions more mainstream. Earlier this year, we hosted the annual Machine Learning for Healthcare Conference, where we had an opportunity to showcase our implementation success stories on the one hand, but also share with our peer institutions and learn from them about best practices for implementing innovative machine learning models for greatest impact.

In the new academic year, we look forward to continuing the momentum, and leveraging our incredibly talented workforce – faculty, staff, and learners – to drive us to even higher levels of excellence and impact and build a healthier tomorrow for all.

Craig Albanese, MD, MBA

Suresh Balu, MBA, MS

DIHI Perspectives

DIHI Experiences

Diffusion and Scaling

Innovation Portfolio Impact

Health Guard for Advanced Care Planning

Prior to the development and integration of the DIHI mortality model, only 3% of patients had a goals-of-care conversation in the six months leading up to their death. Post implementation, in the most recent data, this figure has risen to 55% and is increasing.

The Gynecologic Oncology Surgery Postoperative Opioid Use Predictive (GO-POP) Calculator

The GO-POP calculator led to a 33% decrease (15 to 10) in the median number of opioids prescribed following surgery and an increase in application use from 33% to 64%. There was no increase in post-operative uncontrolled pain interventions or patient-reported pain.

Development of a Machine Learning Model for Early Detection of Pediatric Sepsis

We are implementing the real-time prediction model and will be evaluating its performance using metrics such as sepsis-associated mortality rate, sepsis-recognition time, hospital length of stay and equity of care.

Improving Equity and Efficiency in Access to Organ Transplantation

Through our data quality assurance process, we reduced data missingness and inaccuracies across eleven access-related variables from a range of 2.57-69.82% missingness to 0.0%-10.92% missingness. More accurate measurement of inequities improves understanding of personal and structural barriers encountered throughout the transplant selection process. This also enables design and effectiveness measurement of equity-improvement programs.

Improving Peripheral Artery Disease (PAD) Care at a Population Level

We implemented a PAD risk predictive model and created a population health rounding program for Duke patients. Six months into the project, over 200 patients have been identified as having PAD, and over 100 have received an intervention to mitigate progression of their disease.

Predicting Hospital Admissions and ER Visits from Immune-related Adverse Events

We have deployed and are preparing to pilot a tool to identify immunotherapy patients who are at high risk of an admission or visit to the Emergency Department. Our goal is to reduce unplanned visits to the ED and hospital and improve outcomes through early multi-disciplinary interventions.

Caremap: A Digital Personal Health Record for Complex Care Coordination

The Caremap app supports complex care pediatric patients and providers to securely organize and coordinate care longitudinally. Caremap has been developed and deployed. Patient enrollment has begun. We are studying efficacy of the app.

A Geriatric-Specific Morbidity and Mortality Risk Stratification Tool

We optimized post-op mortality and morbidity models and a user-interface to provide risk event rates and percentiles within five seconds at any time a surgeon or anesthesiologist needs. Specialized geriatric models developed are as performant and reliable as models trained for the full adult population.

Organization and Clean Up of the Electronic Health Record (EHR) Problem List

We improved the organization of the problem list in the EHR and will soon implement automated cleanup of over one million duplicate and lapsed diagnoses across 99,000 Duke patients.

A Predictive Model for Lung Transplant Post-Operative Complications and Survival

The model will lighten the workload of the Lung Transplant Clinic by helping providers prioritize transplants among waitlisted patients and prepare the patient for transplant. In the long-term, we expect to increase the transplant 1-year survival rate.

Algorithm Development for Duke Emergency Pre-hospital Capacity Management

We are working to implement and assess a solution to predict and prevent overcrowding in Duke Emergency Departments. We are aiming to improve time to intervention, ED staff resilience, ED wait times, and rates of patients left-without-being-seen.

GO-POP: Integration of a Personalized Post-operative Opioid Use Calculator in Gynecologic Oncology Surgery

Problem

Prescription of opioids following surgery is central to pain management. However, misuse of opioids is an immense problem in the United States, and, unfortunately, during the COVID-19 pandemic this problem only grew.[1,2] The over-prescription of pain medication contributes to the epidemic of abuse and many physicians, including gynecologists, often over-prescribe postoperatively.[3,4] Despite this, there is little evidence-based guidance on post-operative opioid prescription that allows for minimizing excess post-operative prescribing while precisely prescribing the right amount to meet each patient's unique needs.

Physicians and trainees need an evidence-based tool that can help to provide patient-personalized guidance on appropriate post-operative opioid prescriptions. In two prior prospective studies enrolling 382 subjects, we had recently developed and validated a predictive nomogram post-operative opioid prescribing for women undergoing surgery in the Gynecologic Oncology division for either benign or malignant indications. This application was incorporated into a Shiny app (an open-sourced R package allowing web applications to be built using the R, the statistical programming language). However, the use of this app required manual data entry during every use and the studies did not incorporate methods of applying this innovation into clinical practice.

01 Team

Nicole C. Zanolli
Will Knechtle, MPH, MBA
Marshall Nichols, MS
Laura J. Havrilesky, MD
Brittany A. Davidson, MD

02 Project in Brief

PROBLEM:
Over-prescription of pain medication contributes to the epidemic of opioid abuse, including postoperative prescriptions. Few evidence-based tools provide easy-to-use guidance on appropriate post-operative opioid prescription.

SOLUTION:
We built an application to integrate a Gynecologic Oncology Postoperative Opioid Use Predictive (GO-POP) calculator into clinical practice. A quality improvement pilot kicked off including phone calls to patients and real-time assessment of opioid prescription and application use.

IMPACT:
Use of GO-POP with the Gynecologic Oncology division led to 33% decrease in the median number of opioids prescribed following surgery, without an increase in post-operative uncontrolled pain interventions or patient-reported pain.

Therefore, the problem we are trying to solve is the need for an easy-to-use application that would decrease the number of opioid pills prescribed and minimize the number of superfluous pills in the community without compromising pain control, post-operative pain complications, or patient satisfaction.

Solution

We developed a Tableau mobile application that uses the validated Gynecologic Oncology Postoperative Opioid Use Predictive (GO-POP) algorithm to generate an estimated number of opioid pain pills a patient will need following surgery.[5] The application automatically identified current post-operative Gynecologic Oncology division patients. It immediately retrieved the GO-POP algorithm's required information from the electronic health record (EHR). Automated nomogram inputs included operative time, pregabalin administration, patient age, patient education attainment and smoking history, as well as their pre-operative anxiety regarding surgery and anticipated need for post-operative pain medication. The application presented these inputs on the left, while allowing the care provider using it to update them on the right if they saw fit. Furthermore, this application's integration with the EHR allowed the algorithm to calculate and auto-populate values into a phone-sized Tableau dashboard, easing use for the prescribing provider. The application showed the user a single predicted number of pills that a specific patient was predicted to require, as well as the percentage of patients with similar demographics who needed more than 5, 10 and 15 pills (Figure 1).

Impact

During the five-month study period, GO-POP algorithm use increased from 33% in November to 64% in March ($p = .087$, Figure 2). Compared to a historical cohort, the implementation of GO-POP resulted in a 33% decrease in the median number of post-operative opioids prescribed from 15 to 10 ($p < .001$). Thankfully, overall, GO-POP use did not change the number of refill prescriptions, provider emergency or urgent care visits, or readmissions for pain during the follow-up time period (Table 1). Use of GO-POP also did not change how patients rated their average pain over the prior week (scale of 1-10) at the time of the first follow-up (median 4 vs 4, $p = .834$). Given this, we believe the development and implementation of the GO-POP application was successful in providing an evidence-based tool to providers that allowed them to minimize their opioid prescribing without sacrificing quality of pain management. Integration of GO-POP allowed for evidence-based opioid prescribing that maximized pain control while reducing the number of excess opioids entering our community.

Figure 1. GO-POP mobile application interface

Next Steps

Though launched within the Gynecology Oncology division, patients with both benign and malignant indications for surgery were included in the original cohort used to build the GO-POP algorithm. Given this, we believe that the GO-POP application would benefit many more patients undergoing surgery in the Gynecology Department and further reduce excess opioids entering our community. We aim to expand the use of the application within the department.

To utilize GO-POP appropriately, pre-operative questions regarding the patient's pre-operative anxiety and anticipated need for post-operative pain medication will be integrated into pre-operative appointments, and providers will be educated on app updates.

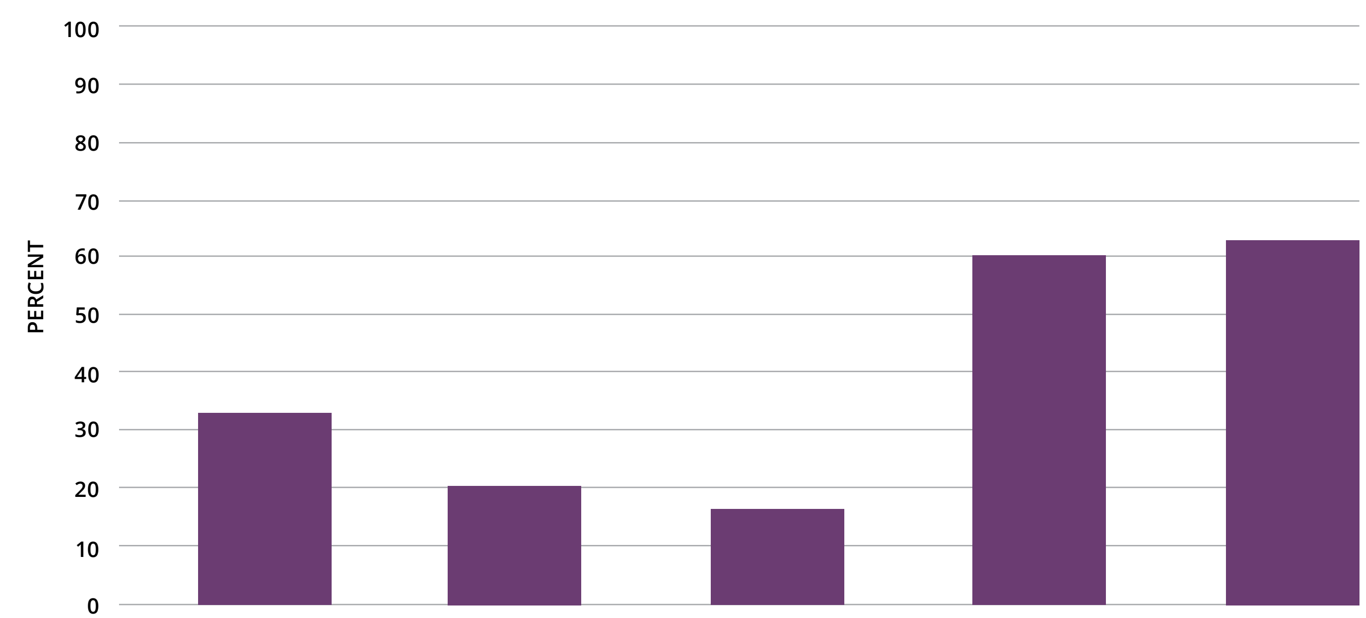

Figure 2: Rate of GO-POP use in post-operative opioid prescribing by month

PAIN INTERVENTION	GO-POP USE (n=120)	GO-POP USE (n=173)	p-VALUE
Refill prescription, n (%)	6 (5.00)	16 (9.25)	.175
Provider visit for pain, n (%)	1 (0.83)	1 (0.58)	.794
ED or urgent care visit for pain, n (%)	0 (0.00)	4 (2.31)	NA
Readmissions for pain, n (%)	0 (0.00)	3 (1.73)	NA

Table 1. Comparison of GO-POP use on post-operative uncontrolled pain interventions

Academic output

POSTER PRESENTATIONS:

01

Integration of a Post-operative Opioid Calculator into an Academic Gynecologic Surgery Practice

Zanolli N, Knechtle W, Havrilesky L, Davidson B
North Carolina Obstetrical and Gynecologic Society (NCOGS) 2022 Annual Meeting. Kiawah, SC, 2022.

02

Integration of a Post-operativeOpioid Calculator into an Academic Gynecologic Surgery Practice

Zanolli N, Knechtle W, Sendak M, Havrilesky L, Davidson B. Machine Learning Conference for Healthcare. Durham, NC, 2022.

References

1. Drug Overdose Deaths in the U.S. Top 100,000 Annually. https://www.cdc.gov/nchs/pressroom/nchs_press_releases/2021/20211117.htm. Accessed May 23, 2022.
2. Products - Vital Statistics Rapid Release - Provisional Drug Overdose Data. https://www.cdc.gov/nchs/nvss/vsrr/drug-overdose-data.htm. Accessed May 23, 2022.
3. As-Sanie S, Till SR, Mowers EL, et al. Opioid prescribing patterns, patient use, and postoperative pain after hysterectomy for benign indications. Obstet Gynecol. 2017;130(6):1261-1268. doi:10.1097/AOG.0000000000002344
4. Lamvu G, Feranec J, Blanton E, Perioperative pain management: an update for obstetrician-gynecologists. Am J Obstet Gynecol. 2018;218(2):193-199. doi:10.1016/j.ajog.2017.06.021
5. Davidson B, Jelovsek JE, Rodriguez I, et al. Development and validation of the Gynecologic Oncology Predictor Of Postoperative opioid use (GO-POP) model. Gynecol Oncol. 2021;162:S55. doi:10.1016/S0090-8258(21)00748-4

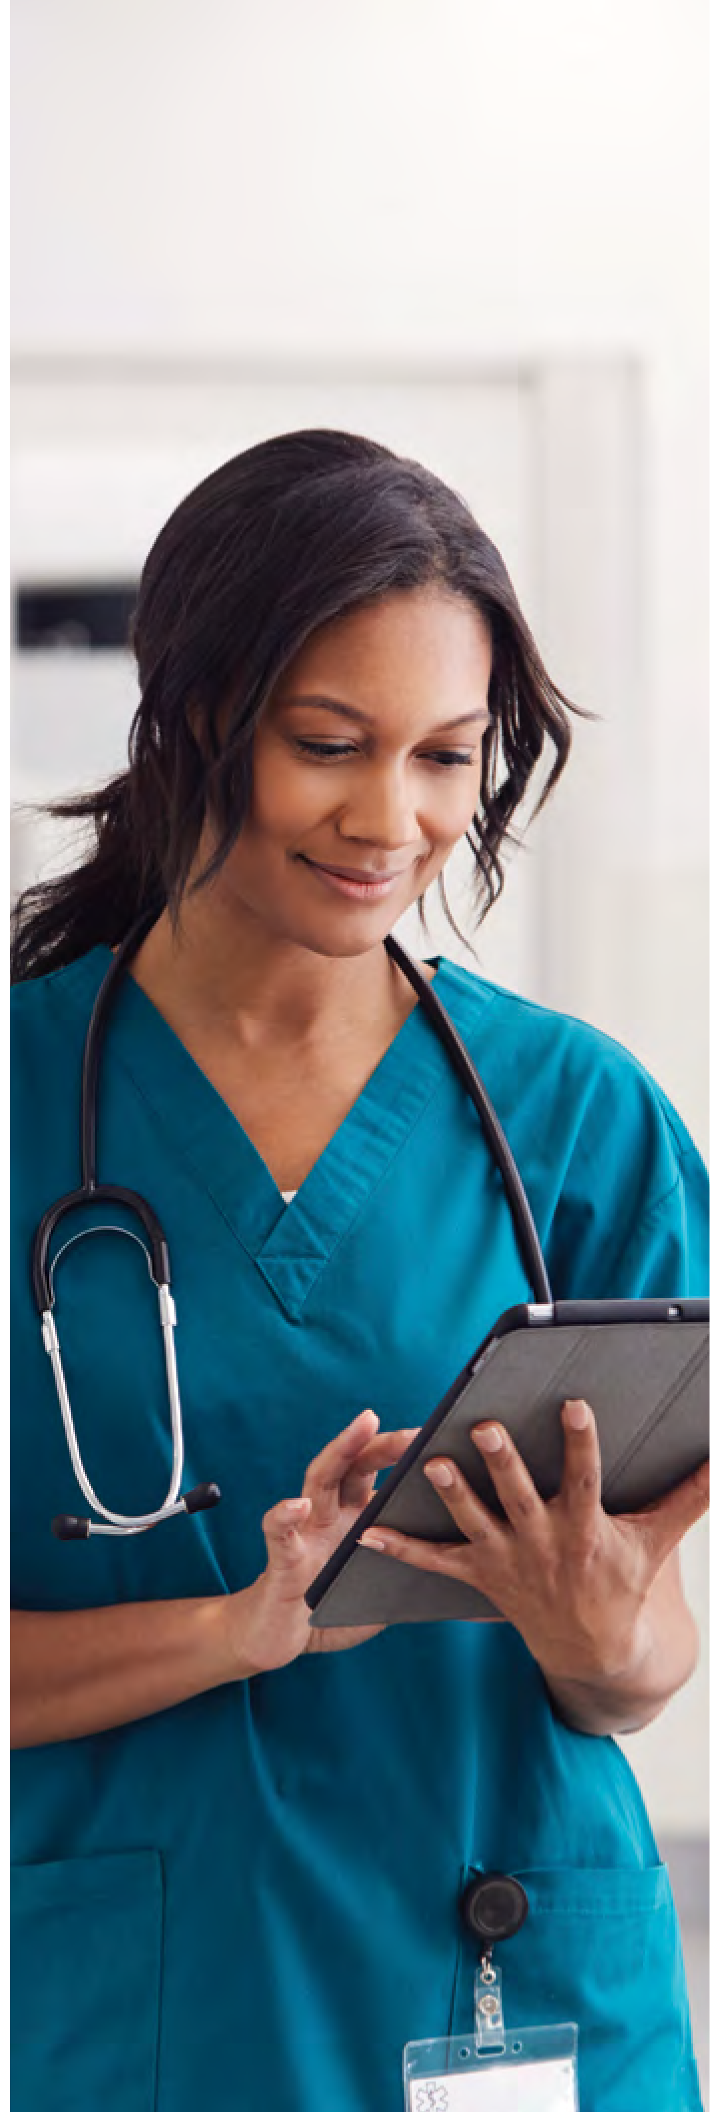

Development of a Machine Learning Model for Early Detection of Pediatric Sepsis

Problem

Sepsis is a dysregulated response to an infection and is a leading cause of morbidity and mortality among hospitalized children. From 2004 to 2012, across 43 hospitals, the prevalence of severe pediatric sepsis was 7.7%, with an associated mortality rate of 14.4%.[1] At Duke, sepsis-associated mortality is 8%, which is 30-50% higher compared to similar health systems.[2] Timely detection and treatment of sepsis is crucial to reducing the sepsis-associated mortality rate. However, the low specificity of associated abnormal vital signs makes detecting sepsis challenging among children. Sepsis indicators for adult patients, such as abnormal temperature, low blood pressure, and elevated heart rate, are less specific to sepsis in pediatric patients. Only a limited number of machine learning models exist to help recognize sepsis in children, and their performance has not been evaluated in real-time.[3, 4]

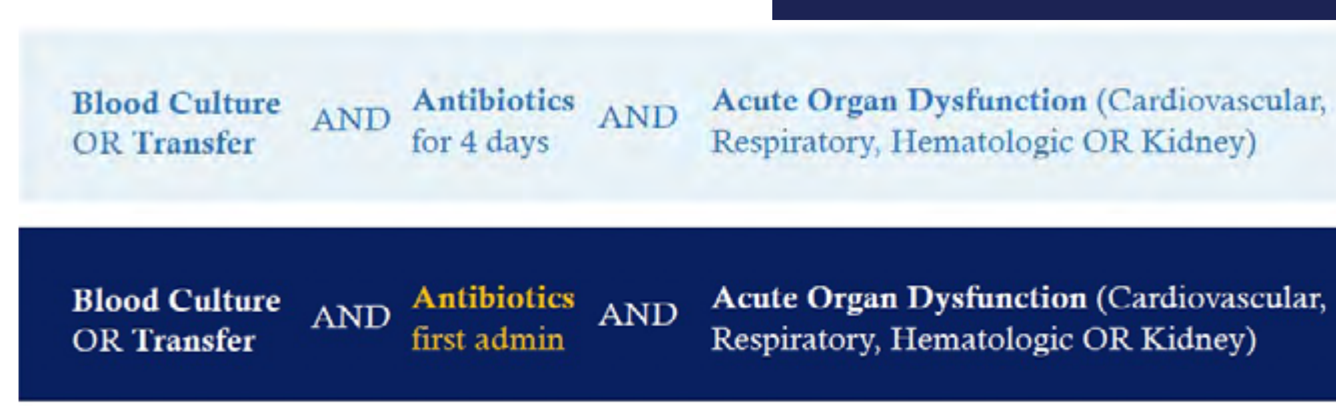

Figure 1: Comparison between the original Weiss definition (top) and the real-time Weiss definition (bottom).

01

Team

Linda Tang | Will Ratliff, MBA
Mark Sendak, MD, MPP | Michael Gao, MS
Jiayu Yao, PhD | Tim Ochoa
Marshall Nichols, MS | Mike Revoir
Faraz Yashar | Suresh Balu, MBA, MS
Neel Subramanian, MD | Crystal Crider, RN
Tammy Uhl, RN | Jordan Pung, MD
Liset Denis, RN | Emily Sterrett, MD, MS

02

Project in Brief

Sepsis in hospitalized pediatric patients has a high rate of poor outcomes but is challenging to detect early. To assist physicians in intervening earlier on patients who are at high risk of developing sepsis, our team developed a machine learning model that alerts clinicians of the onset of pediatric sepsis within the subsequent six hours. Currently, we are implementing the model in real-time and evaluating its performance using metrics such as sepsis-associated mortality rate, sepsis recognition time, and hospital length of stay.

Solution

To address this gap, we developed a pediatric version of the adult sepsis machine learning model, which accounts for the differences in vital sign and lab result interpretation, among other model inputs, between pediatric and adult patients. We constructed a predictive model that assesses patients' risk of meeting the sepsis phenotype in real-time to assist clinicians in providing timely intervention for patients.

We used Emergency Department (ED) and inpatient data collected from 17,491 hospitalizations for 10,492 unique patients between 30 days and 18 years old at Duke University Hospital (DUH) over the time period of November 2015 through December 2020. Pediatric sepsis was defined by the Weiss definition, which included the co-occurrence of infection and an indication of acute organ dysfunction (Figure 1).[5] However, since this phenotype requires four days of antibiotics, it did not support real-time identification of sepsis onset. So, we created a real-time version of the Weiss definition by reducing the four days antibiotics requirement to the first administration of antibiotics (Figure 1). We then adjudicated this real-time Weiss definition on Duke pediatric patients. Of the thirty three patients we adjudicated who met the real-time Weiss definition, 73% met the clinical presentation of sepsis. Figure 2 is a dashboard we utilized for validation that shows patients in DUH who met the real-time Weiss definition.

We then constructed a model to predict the occurrence of pediatric sepsis onset within the subsequent six hours. For model construction, we used patients who met the original Weiss definition as the positive samples. For these patients, we used the timestamp associated with the real-time Weiss definition as the time of onset of sepsis. We used both static features, such as demographics and comorbidities, as well as dynamic features such as vitals, labs, and medication administrations to train the model. We used 70% of the data for training, and of the 30% remaining and used 15% each for validation and testing. We evaluated the model performance using AUC-ROC and average precision score and chose the best model with the highest validation scores.

Initially, the model was trained using LightGBM, a gradient-boosting decision tree framework. However, the LightGBM framework examined each hour of a patient's encounter independently. This meant it ignored the temporal dependencies. Moreover, due to the rare occurrence of sepsis, treating each hour of a patient's encounter as a separate observation caused the dataset to be heavily imbalanced (positive outcomes: 0.4%). This imbalance limited the performance of the Light GBM model.

	SENSITIVITY	SPECIFITY	POSITIVE PREDICTIVE VALUE (PPV)
LSTM (no snooze)	0.54	0.99	0.29
LSTM (with 3-hour snooze)	0.54	0.99	0.33

Table 1: The performance of the LSTM model on the testing set at a threshold of 0.789

Hence, we explored a time-series-based modeling approach to incorporate the temporal trend between observations. We constructed a Long Short-Term Memory network (recurrent neural network) model. Given that most pediatric sepsis is community–acquired,[6] and to limit the imbalance of positive-to-negative outcome hours, we only included one week of data up to the time of sepsis occurrence and the first week of data for non-septic patients while training the model. We also applied a snoozing window of three hours to reduce alarm fatigue compared to an hourly alert while mitigating the likelihood of missing a timely detection as compared to a longer snooze window.

Outcomes

Our LSTM model (with a three-hour snooze incorporated) achieved a performance of AUC 0.919 and AUPRC 0.307. At a threshold of 0.789, the LSTM model achieves a sensitivity of 0.54, specificity of 0.99, and PPV of 0.33 on the testing set.

Next Steps

We are piloting the pediatric sepsis solution in the DUH ED to evaluate its performance. The clinical workflow utilizes push notifications and snoozing logic to alert the ED attending when a patient is at high risk of meeting or has met the real-time Weiss definition. At the conclusion of the pilot period, we will evaluate the impact of the solution on sepsis-recognition time, time to antibiotics, hospital length of stay, intensive care unit (ICU) requirement and ICU length of stay, and in-hospital mortality.

Academic output

Our work contributed to an abstract, poster, and spotlight presentation at the Machine Learning for Healthcare Conference, which was held in August 2022 in Durham, NC.

Identifying Sepsis in real-time for Duke University Hospital Pediatric Patients [Poster presented].

Tang, L, Ratliff, W, Sendak, M, Gao, M, Nichols, M, Revoir, M, Yashar, F, Yao, J, Balu, S, Subramanian, N, Uhl, T, Denis, L, Sterrett, E
Machine Learning for Healthcare (MLHC) Conference. August 2022. Durham, NC, USA.

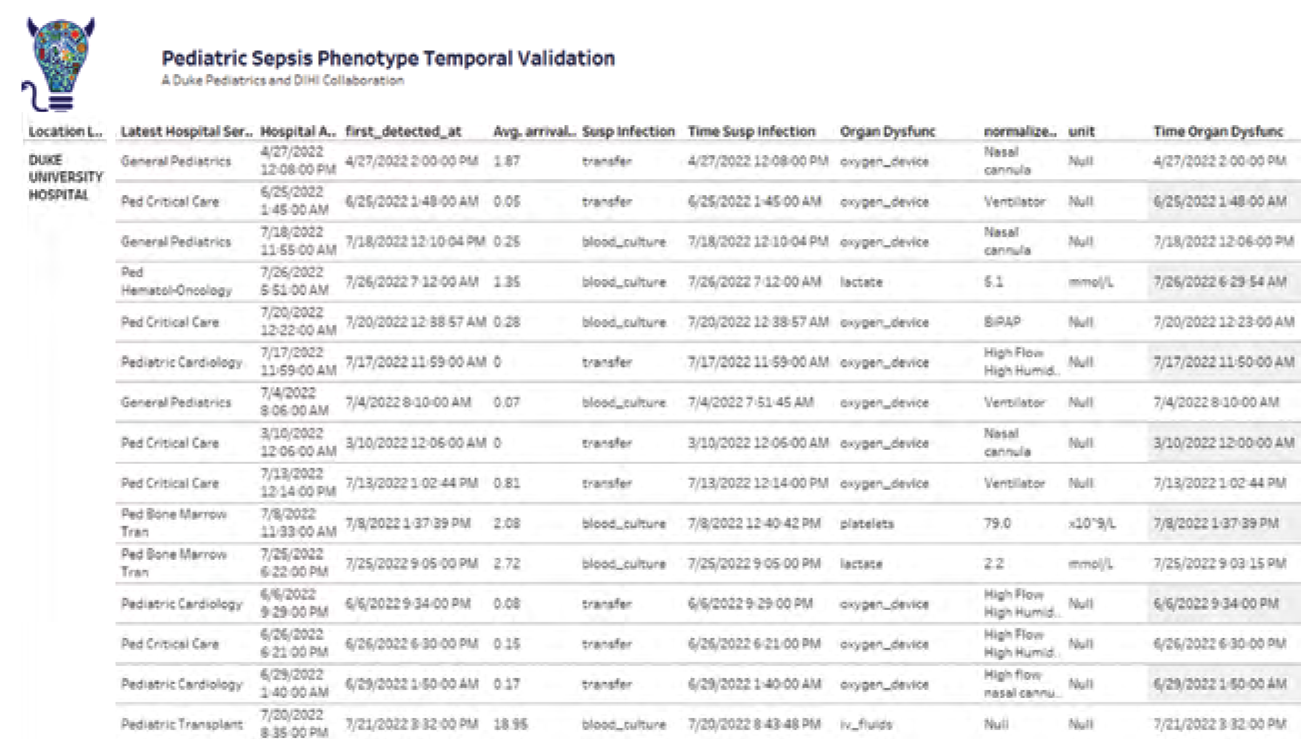

Pediatric Sepsis Phenotype Temporal Validation
A Duke Pediatrics and DIHI Collaboration

Location L..	Latest Hospital Ser..	Hospital A..	first_detected_at	Avg. arrival..	Susp Infection	Time Susp Infection	Organ Dysfunc	normalize..	unit	Time Organ Dysfunc
DUKE UNIVERSITY HOSPITAL	General Pediatrics	4/27/2022 12:08:00 PM	4/27/2022 2:00:00 PM	1.87	transfer	4/27/2022 12:08:00 PM	oxygen_device	Nasal cannula	Null	4/27/2022 2:00:00 PM
	Ped Critical Care	6/25/2022 1:45:00 AM	6/25/2022 1:48:00 AM	0.05	transfer	6/25/2022 1:45:00 AM	oxygen_device	Ventilator	Null	6/25/2022 1:48:00 AM
	General Pediatrics	7/18/2022 11:55:00 AM	7/18/2022 12:10:04 PM	0.25	blood_culture	7/18/2022 12:10:04 PM	oxygen_device	Nasal cannula	Null	7/18/2022 12:06:00 PM
	Ped Hematol-Oncology	7/26/2022 5:51:00 AM	7/26/2022 7:12:00 AM	1.35	blood_culture	7/26/2022 7:12:00 AM	lactate	5.1	mmol/L	7/26/2022 6:29:54 AM
	Ped Critical Care	7/20/2022 12:22:00 AM	7/20/2022 12:38:57 AM	0.28	blood_culture	7/20/2022 12:38:57 AM	oxygen_device	BiPAP	Null	7/20/2022 12:23:00 AM
	Pediatric Cardiology	7/17/2022 11:59:00 AM	7/17/2022 11:59:00 AM	0	transfer	7/17/2022 11:59:00 AM	oxygen_device	High Flow High Humid..	Null	7/17/2022 11:50:00 AM
	General Pediatrics	7/4/2022 8:06:00 AM	7/4/2022 8:10:00 AM	0.07	blood_culture	7/4/2022 7:51:45 AM	oxygen_device	Ventilator	Null	7/4/2022 8:10:00 AM
	Ped Critical Care	3/10/2022 12:06:00 AM	3/10/2022 12:06:00 AM	0	transfer	3/10/2022 12:06:00 AM	oxygen_device	Nasal cannula	Null	3/10/2022 12:00:00 AM
	Ped Critical Care	7/13/2022 12:14:00 PM	7/13/2022 1:02:44 PM	0.81	transfer	7/13/2022 12:14:00 PM	oxygen_device	Ventilator	Null	7/13/2022 1:02:44 PM
	Ped Bone Marrow Tran	7/8/2022 11:33:00 AM	7/8/2022 1:37:39 PM	2.08	blood_culture	7/8/2022 12:40:42 PM	platelets	79.0	x10^9/L	7/8/2022 1:37:39 PM
	Ped Bone Marrow Tran	7/25/2022 6:22:00 PM	7/25/2022 9:05:00 PM	2.72	blood_culture	7/25/2022 9:05:00 PM	lactate	2.2	mmol/L	7/25/2022 9:03:15 PM
	Pediatric Cardiology	6/6/2022 9:29:00 PM	6/6/2022 9:34:00 PM	0.08	transfer	6/6/2022 9:29:00 PM	oxygen_device	High Flow High Humid..	Null	6/6/2022 9:34:00 PM
	Ped Critical Care	6/26/2022 6:21:00 PM	6/26/2022 6:30:00 PM	0.15	transfer	6/26/2022 6:21:00 PM	oxygen_device	High Flow High Humid..	Null	6/26/2022 6:30:00 PM
	Pediatric Cardiology	6/29/2022 1:40:00 AM	6/29/2022 1:50:00 AM	0.17	transfer	6/29/2022 1:40:00 AM	oxygen_device	High flow nasal cannu..	Null	6/29/2022 1:50:00 AM
	Pediatric Transplant	7/20/2022 8:35:00 PM	7/21/2022 3:32:00 PM	18.95	blood_culture	7/20/2022 8:43:48 PM	iv_fluids	Null	Null	7/21/2022 3:32:00 PM

Figure 2: Dashboard for the real-time Weiss phenotype

References

1. Ruth, A, McCracken, C, E, Fortenberry, J, D, Hall, M, Simon, H,K, & Hebbar, K , B, (2014). Pediatric severe sepsis: current trends and outcomes from the Pediatric Health Information Systems database. Pediatric critical care medicine: a journal of the Society of Critical Care Medicine and the World Federation of Pediatric Intensive and Critical Care Societies, 15(9), 828–838. https://doi.org/10.1097/PCC.0000000000000254

2. Ames SG, Davis BS, Angus DC, Carcillo JA, Kahn JM. Hospital Variation in Risk-Adjusted Pediatric Sepsis Mortality. Pediatr Crit Care Med. 2018 May;19(5):390-396. DOI: 10.1097/PCC.0000000000001502. PMID: 29461429; PMCID: PMC5935525. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5935525/

3. Lamping, F, Jack, T, Rübsamen, N, Sasse, M, Beerbaum, P, Mikolajczyk, R, T, Boehne, M, & Karch, A (2018). Development and validation of a diagnostic model for early differentiation of sepsis and non-infectious SIRS in critically ill children - a data-driven approach using machine-learning algorithms. BMC Pediatrics, 18(1), 112. https://doi.org/10.1186/s12887-018-1082-2

4. Kamaleswaran, R, Akbilgic, O, Hallman, M,A, West, A, N, Davis, R, L, & Shah, S, H, (2018). Applying Artificial Intelligence to Identify Physiomarkers Predicting Severe Sepsis in the PICU. Pediatric critical care medicine: a journal of the Society of Critical Care Medicine and the World Federation of Pediatric Intensive and Critical Care Societies, 19(10), e495–e503. https://doi.org/10.1097/PCC.0000000000001666

5. Weiss, S, L, Peters, M, J, Alhazzani, W, Agus, M, Flori, H, R, Inwald, D, P, Nadel, S, Schlapbach, L, J, Tasker, R, C, Argent, A, C, Brierley, J, Carcillo, J, Carrol, E, D, Carroll, C, L, Cheifetz, I, M, Choong, K, Cies, J, J, Cruz, A,T, De Luca, D, Deep, A, ... Tissieres, P (2020). Surviving sepsis campaign international guidelines for the management of septic shock and sepsis-associated organ dysfunction in children. Intensive care medicine, 46(Suppl 1), 10–67. https://doi.org/10.1007/s00134-019-05878-6

6. Page DB, Donnelly JP, Wang HE. Community-, Healthcare-, and Hospital-Acquired Severe Sepsis Hospitalizations in the University HealthSystem Consortium. Crit Care Med. 2015 Sep;43(9):1945-51. doi: 10.1097/CCM.0000000000001164. PMID: 26110490; PMCID: PMC4537676. https://pubmed.ncbi.nlm.nih.gov/26110490/

DIHI STAFF PERSPECTIVE

Michael Gao

Federated Learning in Healthcare

As artificial intelligence (AI) and machine learning models become increasingly prevalent in healthcare, several challenges regarding the training and implementation of these models have surfaced. Foremost is the failure to generalize to different patient populations. Machine learning (ML) research communities have well established that problems with generalization can often be alleviated by using more training data that represents the population of interest. However, obtaining representative training data frequently seems like an insurmountable barrier, exacerbated by the inherently private nature of patient data.

> “The only thing that has to be shared between sites is the model itself!

In the Duke Institute for Health Innovation's (DIHI's) experience working with other institutions, data-sharing agreements are often a non-starter. Even if an agreement is reached, the logistics of deidentification and/or secure computation of this data is exceedingly complicated. The ML community has started to turn to other approaches to progress ML in medicine via the acquisition of representative training data. One of these other approaches is federated learning.

At its core, federated learning offers an appealing promise: train machine learning and AI models with data from other institutions without ever actually transferring data between sites. At first glance, this claim seems difficult to believe – how can a model learn from data without ever seeing it? To understand this, we will briefly discuss the idea at a high level and then provide references for further investigation.

Modern neural networks are trained in an iterative fashion to learn the relationship between input data and a target outcome. In a way similar to learning facts from flash cards, neural networks are provided with prompts and the appropriate answer. Then, the neural network is tweaked in order to move the networks' original answer to the correct answer. As the models see more and more data, this tweaking – also known as weight updating – moves the model's prompts closer and closer to the correct answers. How does federated learning leverage this learning mode? The key insight is that this weight updating can take place separately at each site where data is held. That is, we can begin by training the network at one site and then send the current model over to another site and train independently there. In this way, the only thing that has to be shared between sites is the model itself! (The model in its most basic sense, is an algorithm with features, weights, and parameters – it is independent of protected health information.) It is immediately clear the potential that this has in healthcare, where agreements to share a model are sure to be more palatable than sharing agreements.

A natural follow-up question might be, "has this been tried before?" In 2021, NVIDIA, in partnership with twenty institutions, implemented the largest pragmatic application of federated learning in healthcare to date. In this work, published in Nature Medicine, the team created a model that predicts the future oxygen requirements of COVID-19 patients using vital signs, laboratory, and x-ray data. This model was trained in a truly federated fashion because data stayed at each individual site. In other words, the model was shared across sites, yet the model weighting took place separately at each site.

That project was remarkable, both in scale and in implementation, and hopefully will lead to the adoption of similar technologies for other disease processes.

Despite the incredible success of this project, it should not be a given that future projects will experience similar success. In this case, COVID-19 presented a clear and present danger that motivated sites to collaborate in order to disseminate tools that could help combat the pandemic. It is less likely that a similar motivation would be present for conditions that are endemic and that health systems face on a day-to-day basis. In addition, there is still manual work associated with modern federated learning methods. That is, even though models can be trained at different sites, there is a tacit assumption that the data format is the same at every client site. This is an assumption that cannot be made unless there is explicit instruction on how to format data. Explicit data formatting instruction between sites is an indisputable need: anyone who has tried to evaluate models at different sites will be intimately familiar with the challenges of data harmonization. One last challenge is establishing institutional trust. Although federated learning techniques exist to ensure that data privacy is preserved even under the inspection of how weights are changed, institutions will need to be comfortable with these techniques and the associated methodology.

In light of these barriers, it is important to note that there are countermeasures. These barriers should not prevent institutions from attempting to undertake federated learning projects. While federated learning is itself a countermeasure to data sharing disagreements or data security fears, other countermeasures include:

1. Working on problems that are of enough importance that institutions are motivated to create general solutions,
2. Code audits to make sure that federated processes preserve privacy (i.e., protocols such as secure multiparty computation, homomorphic encryption, or differential privacy are correctly implemented), and
3. Relying on existing data standards such as FHIR, PCORNET, or OMOP to streamline data harmonization requirements.

Federated learning provides an exciting look into the future of machine learning and AI in healthcare, where a vast amount of digital information exists in distinct silos across the United States and the world. With federated learning, digital information can be leveraged to better human health everywhere, not just in places with rich data sources.

DIHI STUDENT/NEW HIRE EXPERIENCE

Linda Tang

I started working at the Duke Institute for Health Innovation (DIHI) during the summer of my junior year at Duke. I joined the project "Development of a Machine Learning Model for Early Detection of Pediatric Sepsis." First, I learned about the current challenges in identifying pediatric sepsis early. Next, I constructed a digital version of the Weiss phenotype.[1] Then, I continuously improved this phenotype with feedback from physicians until it aligned with the clinical presentation of sepsis. Given that pediatric sepsis is a rare outcome among pediatric patients, I explored several modeling approaches under the guidance of DIHI computer scientists. Currently, we are working with clinicians to integrate the model into clinical workflow.

Despite having taken classes in both biology and statistical science, DIHI provided the first experience where I got to integrate these disciplines within one project. My participation on the DIHI team fostered my desire to continue improving patient outcomes and healthcare deliveries through emerging technologies throughout my future career. Moreover, my experience at DIHI helped me understand how large healthcare systems operate and understand decision-making at a systems level. One of these lessons was just how much solving healthcare challenges requires close collaboration across disciplines.

> My experience at DIHI has strengthened my ability to engage with diverse stakeholders.

I am very excited to join DIHI as a full-time team member in 2022. I look forward to being involved in additional ongoing projects and continuing my learning experience.

1. Weiss SL et al. Surviving Sepsis Campaign International Guidelines for the Management of Septic Shock and Sepsis-Associated Organ Dysfunction in Children. Pediatr Crit Care Med. 2020 Feb;21(2):e52-e106. doi: 10.1097/PCC.0000000000002198. PMID: 32032273.

DIHI SCHOLAR EXPERIENCE

Norine Chan

My experience as a Duke Institute for Health Innovation (DIHI) Scholar was an immersive and engaging time that equipped me with foundational knowledge and skills in innovative health technology. As an aspiring surgeon and health equity researcher, I had the opportunity to curate and analyze abdominal transplant patient data to understand inequities along the transplant selection process continuum for my DIHI project. These quantitative tasks allowed me to develop skills in a programming language I had never learned before (Python) and encouraged me to think critically about healthcare system data collection, storage, extraction, and infrastructure.

DIHI also provided the opportunity to learn qualitative skills through stakeholder interviews and presentations to health system leadership. I became familiar with thematic coding analysis and utilizing powerful change management and entrepreneurial skills to tell a convincing narrative. These project-focused skills were augmented through other elements of the DIHI curriculum, including journal clubs and fireside chats examining innovation, leadership, and the future of healthcare policy and technology. The language and proficiency with which I can engage in these conversations surprised me from week to week. I will continue to build on this knowledge as I look ahead to a multidisciplinary career.

I am leaving DIHI a much-changed individual than when I began. I found a home base in our Morris Street office where ideas were always brewing. I never failed to find a friendly colleague to answer my questions about coding, professional identity, or chat about life. The Scholar experience often demanded that we find new boundaries for our capabilities—and I know I have gained confidence, professional skills, and a more creative and flexible outlook on my future career because of it. I am certain my future pursuits as a surgeon and researcher will draw deeply from the soft and hard skills that I have cultivated from working with DIHI.

Improving Equity and Efficiency in Access to Organ Transplantation

Problem

Although nearly 40,000 organ transplants were performed in 2019, significant disparities remain in the ability of patients from historically marginalized groups to undergo transplantation successfully. This is in part due to the complexity of the transplant selection process. Following referral to a transplant center, patients must navigate a multi step conditional pathway that involves screening to determine a patient's suitability for evaluation, in-person evaluation via multiple multidisciplinary outpatient visits (e.g., social work, case management, surgery, medicine, psychiatry), diagnostic testing and transplant education, and review of the case by a multidisciplinary committee of transplant clinicians who approve or deny listing for transplant. Patients from historically marginalized groups face unique challenges in completing the process and, as a result, have disproportionately lower rates of accessing the transplant waiting list.

Duke Transplant Center (DTC) rates of listing for transplant among patients referred are estimated at 25%. Reliable quantification of elimination rates at each step of the process is required to fully characterize inequities in access to the transplant waitlist. However, no reliable method exists to monitor patients as they progress through the transplant selection process.

The primary barriers to understanding and advancing equity in transplantation are twofold: limited Social Determinants of Health (SDOH) data is collected on patients referred for transplant, and there is no monitoring of patients during the transplant selection process. We sought to create a custom Electronic Health Record (EHR) query to allow a more accurate assessment of disparities within the selection process, and assess the extent of SDOH data collection among patients referred for transplant. We additionally investigated the DTC culture and capacity to change to identify barriers to the implementation of programs designed to improve equity.

Solution

To assess health system data infrastructure, we extracted Epic (Epic Systems Corporation, Verona, WI) EHR data on adult (≥18 years old) patient referrals to the Duke Transplant Center (DTC) for kidney or liver transplant from January 1, 2017, to December 31, 2020 (N=7,259). Preliminary extraction and analysis of this referral cohort performed prior to our study exhibited inaccuracies in transplant selection process dates at a rate of >20%. We used Structured Query Language (SQL, https://www.mysql.com/) queries to obtain patient demographic data, transplant selection process dates, and transplant evaluation notes from the Epic data warehouse known as Clarity. Focusing on initial data querying of selection process dates, we developed five phase definitions for reliable data extraction for our transplant referral cohort (Table 1).

We performed three rounds of quality review using the Python programming language (Python Software Foundation, https://www.python.org/) to evaluate for completeness, conformance, and plausibility of the demographic and selection process data in our referral cohort. For variables with >5% data missingness, we made code modifications and explored alternative Clarity data sources to ensure thorough and comprehensive inclusion of all available EHR data.

We next examined the percentage of inclusion of the 28 PhenX toolkit variables in Epic data collection forms and assessed the missingness of these data for our referral cohort on initial review. Patient-level data elements and most current collection forms were reviewed. Each source's point of contact was contacted to verify initial findings.

Finally, we performed a qualitative organizational assessment of the Duke Transplant Center to assess the current culture, capacity, and readiness of the organization to accept equity-focused interventions. We performed fourteen preliminary stakeholder interviews with abdominal transplant coordinators, surgeons, transplant nephrologists, transplant hepatologists, a pharmacist, referral/intake specialists, and a financial coordinator regarding their perceptions of transplant equity, barriers to patient success, and challenges facing the DTC.

01

Team

Norine Chan
Will Knechtle, MBA, MPH
Mark Sendak, MD, MPP
Hamed Zaribafzadeh, MBE
Lisa McElroy, MD, MS
Debra L Sudan, MD
Stuart J Knechtle, MD

02

Project in Brief

PROBLEM:
A disproportionate number of disenfranchised patients are eliminated during the transplant selection process (referral, screening, evaluation, committee deliberation, and decision). Major barriers in advancing equity in transplantation are (a) little research on the process and criteria used by transplant centers and (b) lack of data.

SOLUTION:
We will create a custom EHR-based pathway that consists of smart groups for each phase of the transplant selection process. The smart groups will be integrated into the transplant data pipeline to allow automated tracking of patients through each phase in the transplant selection process and monitoring of patients eliminated for inequities.

IMPACT:
More accurate measurement of these inequities will improve our understanding of personal and structural barriers encountered at each step in the transplant selection process and allow targeted design and accurate measurement of the effectiveness of programs designed to improve equity.

Impact

The overall results of our study support our core hypothesis that a combination of variable data infrastructure, SDOH documentation, and provider perspectives negatively impacts the ability of marginalized patients to successfully complete the transplant selection process. Our study highlights robust opportunities to address inequities in access to solid organ transplantation via (1) improved SDOH data collection infrastructure, (2) continued data monitoring and inequity identification, and (3) implementation of equity-focused education and quality metrics into the transplant center structure.

Analysis of our kidney and liver transplant referral cohort revealed overall decreased odds of listing for transplant and higher odds of elimination at both referral and evaluation phases in the selection process for patients in marginalized groups compared to their privileged counterparts.

A total of 18 variables (64.3%) were included as discrete data collection fields within SDOH forms in Epic, including variables that were poorly represented in the national data source review (access to health services, gender identity, sexual orientation, food insecurity, spirituality, and wealth). Of these eighteen variables, seven variables exhibited 100% missingness for the transplant referral cohort on initial review and after performing a quality review. The other eleven variables ranged in missingness within the cohort from 2.57-69.82% on the initial review and from 0.00-10.92% after data validation.

Our stakeholder interview thematic analysis found four major themes regarding the organizational assessment of the DTC:

1. Disconnect from community
2. Lack of tools to meet patient needs
3. Lack of ownership/accountability
4. Clinician/staff knowledge.

Next Steps

HEALTH SYSTEM DATA INFRASTRUCTURE:

The data query method developed by DIHI will be operationalized for both research and quality improvement.

Quality improvement:

The DTC data team will use the data query to begin tracking patients who are referred but not listed for transplant, with a review of these patients integrated into DTC Quality Assurance Project Improvement (QAPI) process.

PHASE	NAME	DEFINITION
I	Referral/Screening	Date corresponding to the receipt of referral by the Duke Transplant Center and initiation of EHR documentation by transplant coordinator
II	Evaluation	Date corresponding to the first visit either to a transplant specialist (e.g., surgery, nephrology, cardiology) or to obtain diagnostic testing (e.g., computed tomography scan, echocardiogram) for evaluation of transplant candidacy
III	Committee Review/ Decision	Date, of the committee review where a decision regarding eligibility for transplant (approved, declined, needs re-representation) was made.
IV	Waitlist	Date, the transplant candidate, was documented by the transplant nurse coordinator as being added to the United Network of Organ Sharing (UNOS) waitlist.
V	Transplant	Date the transplant surgery was performed.

Research:

A study funded by the American Surgical Association will begin July 1st that integrates the data query into two additional centers (Houston Methodist and University of Michigan) for external validation.

SOCIAL DETERMINANTS OF HEALTH DATA COLLECTION:

The DTC will form a new workgroup to focus on SDOH data collection with representation from transplant social workers, care coordinators, surgeons, nephrologists, hepatologists, and pharmacists. The team will focus on ensuring the completeness of SDOH data collection throughout the transplant selection process.

CULTURE AND CAPACITY FOR CHANGE:

Building a foundational knowledge base regarding equity and creating a culture of inclusion is a critical starting point for advancing equity within the DTC. The DTC will follow the model established by Population Health Sciences and institute a Diversity, Equity, Inclusion (DEI) initiative over the next six months that includes: center-wide climate assessment, online and in–person educational sessions, and implicit association testing of clinical faculty with follow up group discussion sessions.

CLINICAL OPERATIONS AND CARE IMPROVEMENT:

A variety of ongoing equity-focused efforts are already being pursued by DTC clinicians and staff, ranging in topics from pharmacoequity to food insecurity, financial strain, and access to technology; Later this year, the DTC leadership will meet with the project heads to establish objectives and determine what support is required to ensure successful completion. In 2023, regularly scheduled meetings with these project leads and transplant center leadership will begin to monitor progress.

Academic output

01

The American Surgical Association Foundation Award (PI: McElroy) was awarded in November 2021 to continue this work. The award period begins on July 1, 2022. The project aims are to:

1. Implement a data architecture to track patients along the continuum of transplant care at three centers.
2. Quantify disparities in access to the transplant waitlist based on manually extracted enhanced SDOH data.
3. Develop a clinical decision-making support tool to inform multidimensional risk assessment by transplant selection committees.

02

"Social Determinants of Health Data Capture Within National and Health System Data Sources" is accepted for an oral presentation in the Scientific Forum at Clinical Congress 2022, taking place October 16-20 in San Diego, CA.

03

PUBLICATION:

Social Determinants of Health Data in Organ Transplantation: National Data Sources and Future Directions.

Chan N, Moya Mendez M, Henson J, Zaribafzadeh H, Sendak M, Bhavsar N, Balu S, Kirk A, McElroy LM. American Journal of Transplantation. Am J Transplant. 2022 May 18

Improving Peripheral Artery Disease Care at a Population Level

Problem

Peripheral artery disease (PAD) remains underdiagnosed and undertreated. PAD patients suffer an increased risk of heart attack, stroke, amputation, and death, with minority and lower socioeconomic status patients more affected. Untreated Duke patients with PAD have a 49% hospitalization rate, 8% amputation rate, and 15% mortality rate, which is at the poorer end of outcomes nationally.[1,2] These poor clinical outcomes are associated with increased costs to patients and our health system. At Duke, there is an opportunity to expand care for this disease through population health monitoring of PAD.

Solution

Our team is improving the care of Duke patients with PAD through proactive identification and intervention to address gaps in their care. We designed and implemented a PAD virtual rounding program that includes a PAD predictive model and generates intervention recommendations to patients' primary care physicians (PCPs) ahead of scheduled visits. Our goal is to reduce downstream costs and improve equity and outcomes of PAD care within our health system.

We implemented a previously developed and validated logistic regression machine learning model that successfully identified PAD patients based on diagnosis codes and another encounter history within the Electronic Health Record (EHR).[3] Our model runs on all adult patients with a Duke University Health System (DUHS) clinical encounter from the start of the EHR that includes at least one of 108 PAD-related diagnosis codes. A risk score is generated for each patient in the model cohort, indicating the likelihood that the patient has PAD. For our pilot, we focus on patients with primary care provider (PCP) appointments in the upcoming week to provide an actionable recommendation to the PCP at a time when they are thinking about the patient.

01

Team

Rebecca Shen
E. Hope Weissler, MD, MHS
Will Ratliff, MBA
Marshall Nichols, MS
Bradley Hintze, PhD
Michael Gao, MS
Pamela Cohen
Holly Alvarado, PharmD
Dennis Narcisse, MD
Mary Schilder, RN
Tara Kinard, MSN, MBA
Benjamin Smith, PharmD
Daniel Costello, MPA
Steven Lippmann, PhD
Mark Sendak, MD, MPP
Schuyler Jones, MD

02

Project in Brief

Peripheral artery disease (PAD) remains underdiagnosed and undertreated. This puts patients, especially ones with lower socioeconomic status, at risk for cardiovascular and neurovascular complications. We utilize a machine learning model to identify likely PAD patients alongside a population-level rounding process implemented in collaboration with the Duke Population Health Management Office (PHMO). We aim to equitably reduce PAD complications and their downstream costs.

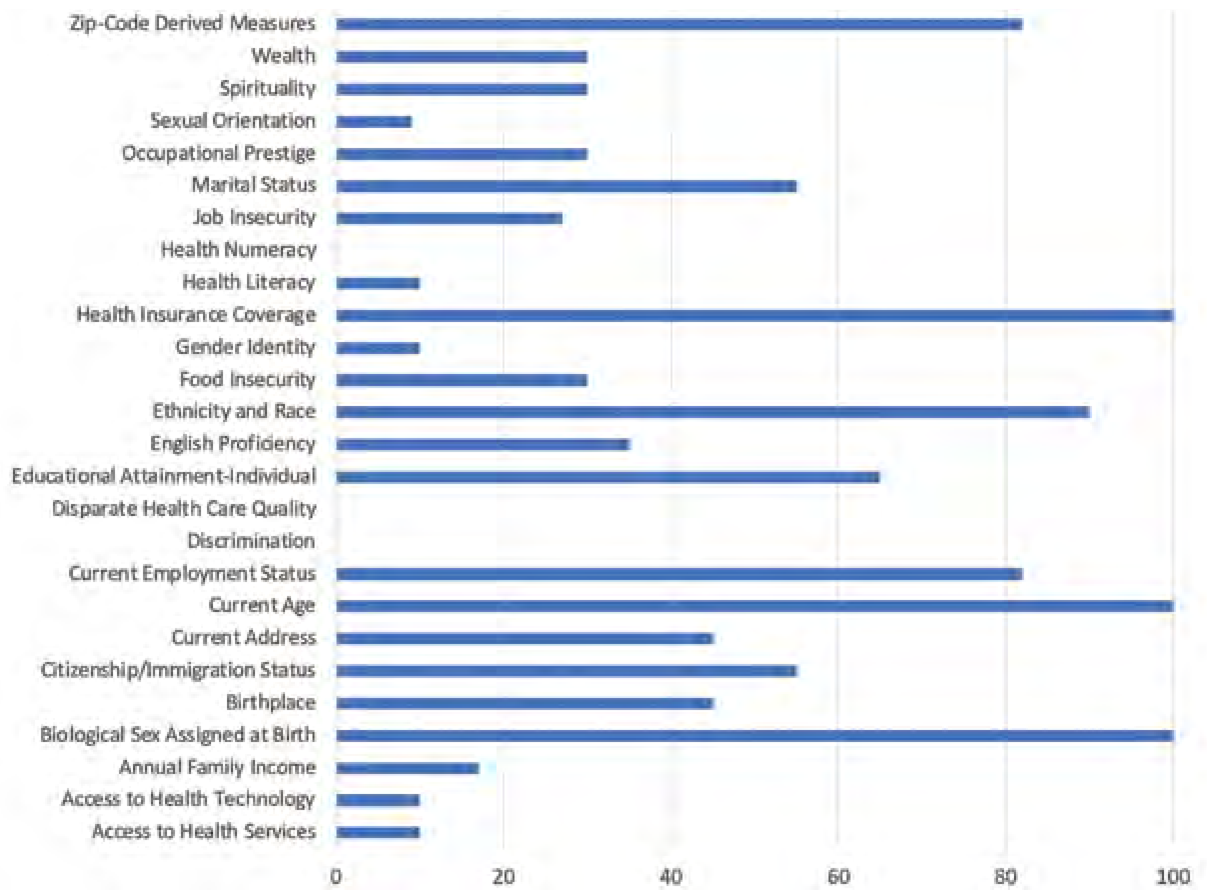

Figure 1. National and health system data source inclusion rate (%) of the social determination of health variables

[1]Zip code-derived variables consist of air quality index, concentrated poverty, and educational attainment—community. Inclusion of the zip code was counted as three variables toward total. Norine Chan and Mary Moya-Mendez led this data aggregation.

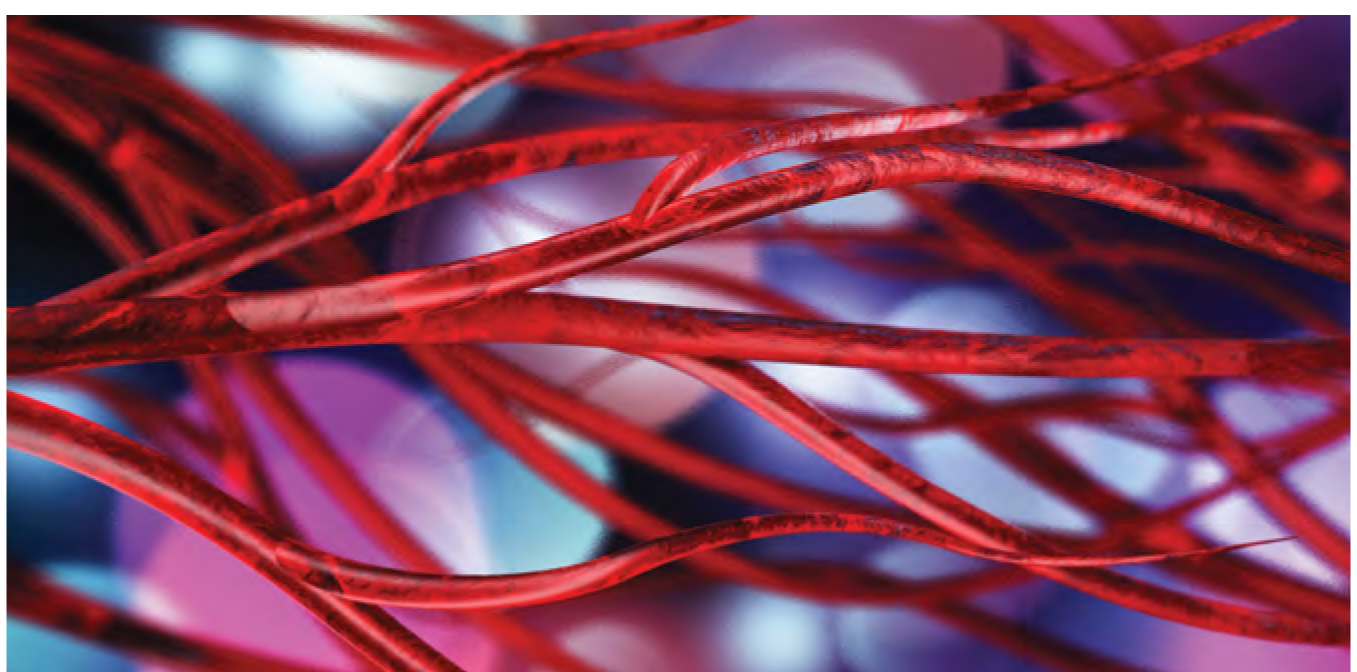

Patients above the threshold are put into a rounding cohort, then assessed by a PAD specialist for inclusion in the virtual rounds discussion. The specialist verifies the patient's PAD disease status and whether the patient would benefit from an intervention, most commonly lipid-lowering medication adjustment (statins and PCSK9 inhibitors), and smoking cessation clinic referrals. Final recommendations are discussed weekly on a multidisciplinary team rounds discussion that includes the PAD specialist, a pharmacist, and a Duke population health lead. Recommended changes to care are communicated to the patient's PCP, who can then tailor the patient's care plan at their upcoming appointment. The model runs on a weekly interval to identify new patients and update scores for existing patients. Figure 1 illustrates the workflow of these ongoing rounds.

Impact

Starting in January 2022, we successfully implemented an ongoing population health rounding intervention for PAD patients that identifies PAD patients within the DUHS. Patients with a model-generated risk score above the threshold enter the rounding cohort, where they are discussed during weekly PAD population-level rounds. In the initial six months of rounds, 237 patients entered active rounds discussions, 56 were recommended for medication adjustments, and 45 for smoking cessation referrals. Other interventions included ankle-brachial index assessment and care management referral.

Next steps

We will evaluate the impact of our pilot on clinical and cost outcomes for our patients, as compared with Duke patient population pre-implementation. We will continue to monitor model accuracy and will assess an alternative natural language processing model, which was previously shown to identify PAD patients with higher accuracy. We will maintain the PAD rounding process, seeking feedback for optimization from key stakeholders, especially from our PCPs. We aim to expand similar virtual rounds-based workflows for PAD at other institutions.

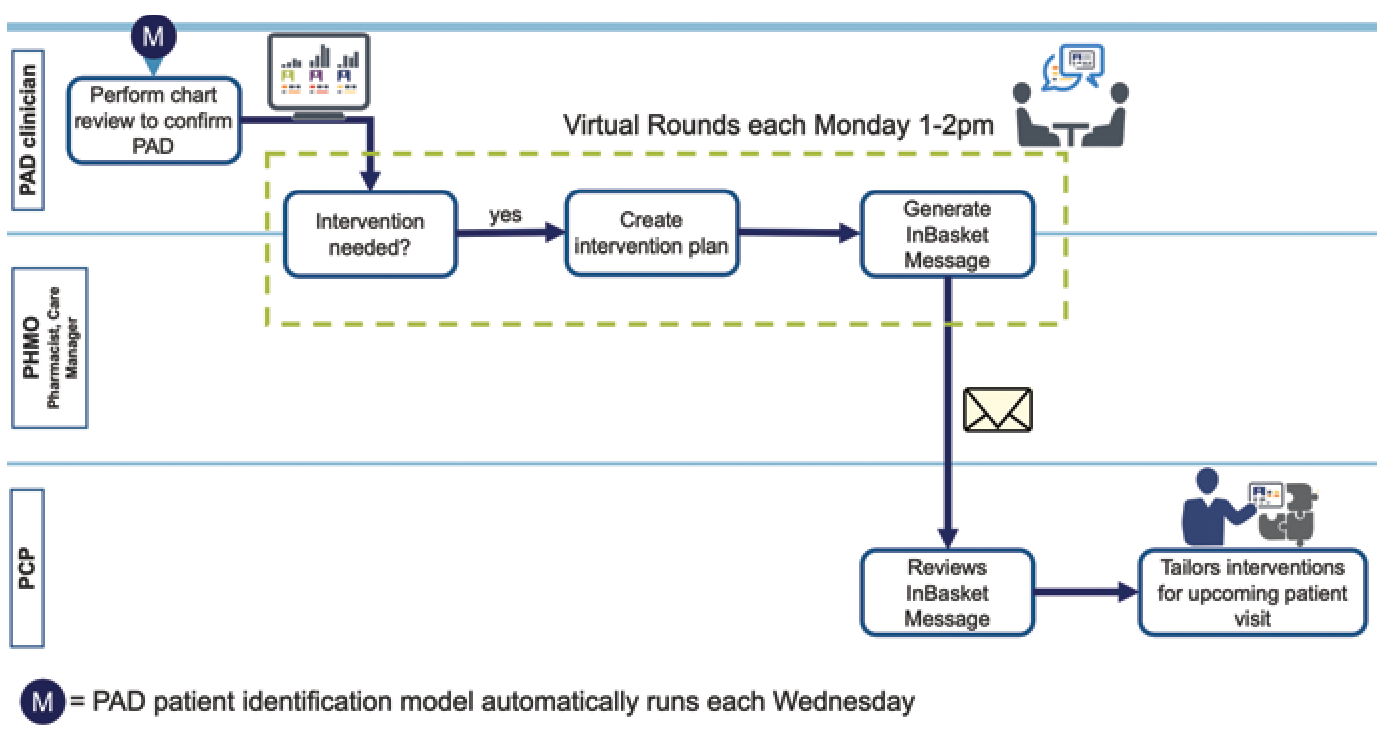

Figure 1. DUHS Virtual Rounds Workflow: PAD clinician + PHMO +PC

Academic output

PUBLICATION:

Improving Equity and Value of Peripheral Artery Disease Care at a Population Level. Machine Learning for Healthcare 2022 – Clinical Abstract Software and Demo Track. August 5-6, 2022.

Rebecca Shen, E. Hope Weissler, William Ratliff, Marshall Nichols, Bradley Hintze, Michael Gao, Pamela Cohen, Holly Alvarado, Dennis Narcisse, Mary Schilder, Tara Kinard, Benjamin Smith, Daniel Costello, Steven Lippmann, Mark Sendak, Suresh Balu, Schuyler Jones.

References

1. Kalbaugh CA, Loehr L, Wruck L, et al. Frequency of care and mortality following an incident diagnosis of peripheral artery disease in the inpatient or outpatient setting: the ARIC (Atherosclerosis Risk in Communities) study. J Am Heart Assoc 2018;7(8):e007332. https://doi.org/10.1097/PCC.0000000000000254
2. Weissler EH, Ford CB, Narcisse DI, Lippmann SJ, Smerek MM, Greiner MA, Hardy NC, O'Brien B, Sullivan RC, Brock AJ, Long C, Curtis LH, Patel MR, Jones WS. Clinician Specialty, Access to Care, and Outcomes Among Patients with Peripheral Artery Disease. Am J Med. 2022 Feb;135(2):219-227. doi: 10.1016/j.amjmed.2021.08.025. Epub 2021 Oct 7. PMID: 34627781; PMCID: PMC8840959. https://pubmed.ncbi.nlm.nih.gov/34627781/
3. Weissler EH, Lippmann SJ, Smerek MM, Ward RA, Kansal A, Brock A, Sullivan RC, Long C, Patel MR, Greiner MA, Hardy NC, Curtis LH, Jones WS. Model-Based Algorithms for Detecting Peripheral Artery Disease Using Administrative Data From an Electronic Health Record Data System: Algorithm Development Study. JMIR Med Inform. 2020 Aug 19;8(8):e18542. DOI: 10.2196/18542. PMID: 32663152; PMCID: PMC7468640. https://pubmed.ncbi.nlm.nih.gov/32663152/

DIHI SCHOLAR EXPERIENCE

Rebecca Shen

Spending a year as a Duke Institute for Health Innovation (DIHI) scholar is a truly unique medical student research year experience. DIHI is a leader in healthcare innovation at Duke and nationally. I was grateful for the opportunity to learn from team members from a variety of backgrounds, including different areas of data science, population health, economics, management, and policy expertise. This past year, I supported two Request-for-Applications (RFA) source projects. The first was "Improving Peripheral Artery Disease Care at a Population Level." DIHI helped implement a machine learning model to identify patients with peripheral artery disease and built the framework used by our population health collaborators to intervene in gaps in care for these patients. It was gratifying to see the actual patients from across the Duke University Health System benefit directly from our innovation. The second "Algorithm Development for Duke Emergency Pre-hospital Capacity Management", was essentially an Emergency Department Rerouting project, and I helped DIHI build a tool that will provide Durham Emergency Medical Services with real-time decision support for patient transport.

Outside of my projects' applications of data science, my year at DIHI introduced me to the discourse around effective leadership, the business of healthcare, policy efforts, and even the ethical considerations of machine learning algorithms. DIHI often challenged me to step outside my comfort zone, created an environment for growth, and gave me valuable insight into healthcare systems strategy and operations. I am excited to apply the knowledge and skills I gained at DIHI to continue advancing patient care throughout my career.

DIHI STAFF PERSPECTIVE

Will Ratliff

Discovering the Economics of DIHI Projects

"The cost of healthcare" is a complex, often charged topic that has distinct meanings depending on the context. At a national level, the cost of healthcare is a ballooning issue that negatively impacts millions of Americans. The problem is especially poignant when looking at the trend of healthcare expenditures per capita versus life expectancy: Across world countries, the United States is a clear outlier on expenditures ($9K, compared to $3K-$5K for most developed nations), yet lags behind most developed nations by 3+ years in life expectancy. Research on this topic points to the price differences and types of care being provided in the United States, with an emphasis on the use of higher-cost advanced technologies and lower allocation of resources to support preventive care. While it is true that Duke utilizes a wide variety of advanced technologies and interventions to care for highly complex patients, it is also true that Duke invests in high-value, proactive care to mitigate the risk of needing high-cost interventions downstream.

> "At the Duke Institute for Health Innovation (DIHI), we have the privilege of supporting many of leadership's investments in the proactive/preventive care opportunities identified by our front-line clinical experts.

At the Duke Institute for Health Innovation (DIHI), we have the privilege of supporting many of the leadership's investments in the proactive/preventive care opportunities identified by our front-line clinical experts. As part of this privilege, we are progressing the economic analyses of solutions we build to try to answer the question, "what is the net value (benefit-cost) of this solution, in dollars?". I use the word "benefit" as a term to encompass positive financial impact, namely the creation of additional revenue or the avoidance of costs (i.e., savings) to the health system. Below are two case studies that exemplify our ongoing work to illuminate the economic impact of our DIHI projects on Duke Health.

Case Study 1:

The DIHI RFA project "Development of a Machine Learning Model for Early Detection of Pediatric Sepsis", described in this impact report (see page 11), aims to improve clinical outcome metrics related to pediatric sepsis by reducing the time to recognition and treatment. By identifying and treating sepsis earlier, we hope to reduce adverse clinical outcomes, including sepsis-associated mortality, intensive care unit (ICU) requirement, ICU length of stay, and hospital length of stay. These metrics not only have a clear and positive clinical implication but also have a positive financial impact. The positive financial impact comes in the form of cost avoidance for Duke University Health System (DUHS). Once we have completed a pilot period of pediatric sepsis solution implementation, we will compare pre- vs. post-implementation cohorts to measure the impact on these metrics. We will work with our DUHS Finance colleagues to translate any measured clinical improvements to a Duke pediatric patient-specific dollar amount. For instance, for hospital length of stay, we can obtain an estimate of costs avoided for one day's worth of care provided to a patient at the end of their hospitalization when they are ready for discharge. We will then annualize the amount of these impacts to arrive at a cost avoidance amount per year through the use of the solution. This would be the per year benefit dollar amount achieved by the project for Duke Health.

Meanwhile, we will also calculate the cost of using the solution. We can think of the solution's costs in two parts: the technical overhead to support it and the effort (time spent) by a front-line clinician to use it. We have a general estimation for the technical overhead from our work with DHTS and other teams to stand up and maintain the infrastructure. To calculate the clinical effort spent, we analyze the retrospective performance of the model and the prevalence of the real-time sepsis phenotype. Specifically, we count the total high sepsis-risk patients identified by the machine learning model (i.e., the hourly true-positive and false-positive alarms) over the model evaluation time period, add the patients who met the real-time sepsis phenotype but were not previously deemed as "high risk for sepsis" by the model (i.e., the false negatives) and divide that sum by the count of twelve-hour shifts in the evaluation time period. We, therefore, arrive at the count of total alarms per twelve-hour shift. Then, if we apply a workload requirement estimate per alarm, such as a chart review (five minutes) and bedside assessment (fifteen minutes), we get an estimate of the time required to use the pediatric sepsis solution. From a resource standpoint, we can apply a dollar amount for the time spent by a nurse, physician, fellow, etc. (i.e., the direct and indirect responders to the alarms) to arrive at an estimate of cost in dollars to use the tool. Depending on the optimal workflow for the patients and clinicians, we can apply certain workflow-driven tactics to reduce the per-shift workload estimate. For this pediatric sepsis project, we apply an alarm snoozing mechanism that reduces the per-shift workload without sacrificing attention to the patient. This reduced the dollar estimate for the clinical use of the solution by 40%.

Case Study 2:

The DIHI RFA project, "Improving Peripheral Artery Disease Care at a Population Level," described in this impact report (see page 23), highlights a solution to support preventive care for peripheral artery disease (PAD) at a population health level. The solution combines a predictive model for PAD in Duke patients with a virtual rounds workflow to produce tailored intervention suggestions to PAD patients' primary care physicians (PCPs) just ahead of an upcoming clinic visit. The PCP then decides whether to utilize these suggestions as part of their care. Should the PCP decide to move forward with the suggestion, the intervention "benefit" is both revenue-generating (e.g., smoking cessation referral) and, hopefully, cost-avoiding (e.g., long-term statin use prevents the need for limb amputation). To understand the cost avoidance portion of the "benefit", we plan to study the patients who received an intervention over a longer time period. In terms of the cost to use the solution, we will similarly assess the technical overhead to support it and the effort by the clinical team to perform the workflow. Specifically, this would entail our clinician's time spent to review high-risk patients, the team members' time during virtual rounds, and an estimate of the PCP's time to review the suggestion. We will work with Duke's Population Health Management Office (PHMO) and Duke Finance to achieve the components of this net value equation.

While it is true that clinical outcomes for our patients should be the primary driver for patient care decision-making at Duke, we cannot separate the cost of care from this equation. In a world of limited resources, we should strive to understand the net value of our patient care decisions such that we can optimize for both clinical and financial outcomes. We can and should reinvest savings to support health and healthcare strategies that preserve and prolong life. By doing this, we will progress the economics of healthcare in the U.S. and hopefully spend less to live longer.

Predicting Hospital Admissions and Emergency Room Visits from Immune-Related Adverse Events

Problem

Immune checkpoint inhibitors (ICIs) have become a mainstay of cancer therapy since the approval of Ipilimumab in 2011. ICIs have the potential to dramatically improve outcomes even for patients with advanced disease.[1] Despite these innovative therapies' dramatic potential, ICIs come with side effects that can be extensive. Immune-related adverse events (irAEs) have been reported in as many as 12-79% of patients taking ICIs and contribute to increased healthcare utilization by cancer patients.[2, 3] Analysis of irAEs is further complicated by a lack of unified nomenclature or diagnostic codes to categorize these events.

Solution

Our team created a novel machine learning model to identify patients at risk of two irAEs: hospital and emergency department (ED) visits within the next three weeks of the patient's latest infusion. Oncology physicians and advanced practice providers will use the model to review patients for intervention in the outpatient setting before they require admission or emergency care. Allowing providers to proactively intervene on patients at risk of an adverse event will increase the utility and minimize the risk of this important and rapidly growing field of cancer therapy.

Our model's training cohort consisted of all adult patients who received CTLA4 inhibitors, PD-1/PD-L1 inhibitors (cytotoxic T-lymphocyte associated protein 4, Programed Death-1, Programed Death ligand, respectively), or both at Duke University Hospitals between October 2015-July 2021. We predicted the risk of any ED visit or hospital admission within the next three weeks from time of prediction for the first six months a patient receives an ICI or until death or ICI cessation. The model and user interface were designed to update daily for patients receiving infusions at Duke during the prior two weeks. Inputs for the model consisted of over 150 features, including lab values, vitals, medication types and duration, demographics, comorbidities, visit reasons, and encounter-level data. Prior to training, we assessed the quality of our data for conformance, completeness, and plausibility. The model was trained using a variety of standard machine-learning techniques, such as logistic regression, random forests, and gradient-boosting models.

Once trained, tested, and validated, our model will be implemented in our outpatient genitourinary (GU) and thoracic oncology clinics. Clinicians will receive automated alerts when their patient is a newly high risk, gather additional information from the patient using a preexisting triage line, and intervene if necessary by providing patients with education about symptoms and when to visit the emergency room. If additional information gathered or known about the patients indicates that more acute care is needed, the clinicians may schedule the patient for a visit in our same-day Duke Cancer Center Acute Care Clinic.

Retrospective analysis comports with previously published studies. In our data, 35.5% of our 5,298 unique patients had a hospital admission or ED visit in the first six months after ICI initiation with a median time of 57 days from the time of ICI initiation (Figure 1). Outcomes were similar across medication types (36.7, 35.3, and 36.1%, respectively). Consistent with the known difficulty capturing irAEs, only 14.4% of our outcomes (5.0% of the total cohort) could be clearly attributable to irAEs as measured by clinician-identified ICD-10 codes. In order to establish the ideal timeline for intervention, we evaluated the elapsed time between an adverse outcome and a patient's most recent scheduled visit prior to their outcomes.

In our cohort, 57.9% of patients with outcomes received immunotherapy less than two weeks prior with a median last-seen interval of 11 days. 12.2% of patients with hospital encounters were seen that same day for ICI administration.

The model yielded an Area Under the Receiver Operating Characteristic (AUC) of 0.76 (Figure 2).

- At a high-risk threshold of 0.3, 59.3% specificity, 4.4% sensitivity
- At a medium-risk threshold of 0.1, 27.1% specificity, 48.9% sensitivity

OUTCOME TIMECOURSE

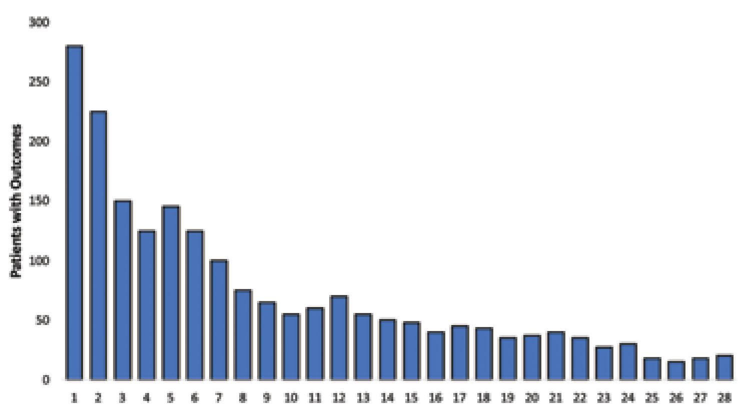

Figure 1. Weeks from first ICI administration to first ED visit or Hospital Admission.

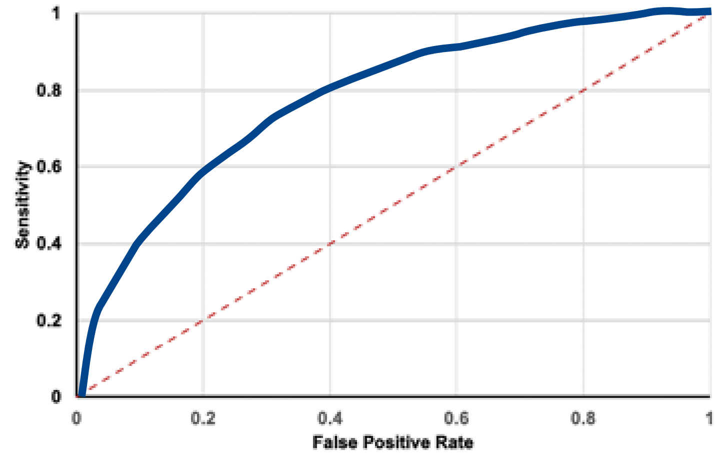

Figure 2. ICI irAE model AUROC

01

Team

Kira Niederhoffer
Will Knechtle, MBA, MPH
Will Ratliff, MBA
Mark Sendak, MD, MPP
Michael Gao, MS
Marshall Nichols, MS
Bradley Hintze, MS
Mike Revoir
Yousuf Zafar, MD, MHS
Carrie Diamond
Jeffrey Clarke, MD
Hope Uronis, MD, MHS
Suresh Balu, MBA, MS
Afreen Shariff, MD, MBBS

02

Project in Brief

Immune checkpoint inhibitors (ICI) are a promising cancer treatment that improve immune system targeting of tumor cells. This comes with the risk of immune-related adverse events (irAE). Our solution was to build a machine-learning-based Adverse Event-Early Warning System using electronic health records and develop a plan for an Adverse Event-Early Intervention Program (AI-EIP). We aimed to reduce the rate of Emergency Department and Hospital admissions.

The most relevant feature for model prediction overall was the length of time on immunotherapy (Figure 3). The most important features of each category include:

LABS:
Albumin, white blood cells, thyroid stimulating hormones, platelets, sodium, platelets, alkaline phosphatase

VITALS:
Pulse, temperature, pulse oximetry

ENCOUNTER LEVEL:
Prior outpatient appointments at Duke in six months prior to ICI initiation

COMORBIDITIES:
Chronic Obstructive Pulmonary Disease (COPD)/ bronchiectasis, hypertension (high blood pressure), anxiety disorders

Impact

The primary outcomes we will measure are emergency room visits, hospital admissions, and their corresponding rates among the ICI cohort. We also will follow cancer progression, patient satisfaction, clinician satisfaction, and ease of score utility.

DIHI EXPERIENCE

Willie Boag

On July 25, 2022, I began work as a Research Scientist with the Duke Institute for Health Innovation (DIHI), specializing in machine learning (ML). I spent the prior six years of graduate school doing research in ML for healthcare. I first encountered DIHI two years ago during a project where I interviewed researchers from over a dozen healthcare organizations about scoping projects and working with clinical collaborators. DIHI stood out among them.

During my last four years' work in policy (especially healthcare policy), I saw how so many barriers to progress are caused by bureaucracies and inertia. Oftentimes, the required technology existed, but there were challenges in the socio-technical system (e.g., poorly designed Electronic Health Records (EHRs) that led to errors and caused burnout. What impressed me so much about DIHI was their focus on innovation and their Request for Applications (RFA) management process as a vehicle for improvement. RFA projects get both top-down and bottom-up support: projects are sought according to themes and strategy from executive leadership, identified by improvable opportunities sourced directly from the frontline staff, and approved with leadership support. I've seen so many instances of healthcare health information technology tools fail to help anyone because of non-technical reasons (e.g., lack of leadership champion, lack of buy-in: willingness to actively participate in and support the tool). Consequently, when I heard about DIHI's work to remedy this and the projects they've done, I was very interested in working there. Since defending my dissertation in the Spring of 2022, I have been working at DIHI on some of the areas of projects that need the most support. I have had the privilege of working work side-by-side with doctors to understand the data-generating process for our projects' data elements so that we can use meaningful model features and avoid the garbage-in, garbage-out scenarios published in some scientific journals. I have been able to work with members from the Artificial Intelligence/ Machine Learning Consortium to Advance Health Equity and Researcher Diversity (AIM-AHEAD) consortium to promote capacity-building for under-served communities. One of our goals is for these under-served community members and institutions to be able to solve problems without needing to depend on others. Most excitingly, I've been able to use public policy evaluation tools to estimate how many people one of our deployed tools has helped.

Good things don't just happen on their own, which is why I'm so excited to have joined a team with a nationally-visible track record of getting it done!

Next Steps

We have created a model and pilot governance team with one or two representatives from each oncologic specialty providing ICI care. At the time of publication, this team is testing whether the model is presenting the correct ICI cohort of the previous two weeks and reviewing model alerts at various thresholds. Once the clinical governance team is satisfied with real-time model validity and the utility of alerts, we will present the model and model-response plan for committee review. The validation and review will inform which specialty to pilot the model with first. The specialty representatives on the governance team and DIHI project managers assume responsibility for developing materials and presentations that educate a piloting team. The governance team will use a model monitoring dashboard and frequent meetings to monitor pilot progress.

Academic output

POSTER PRESENTATION DURHAM, NC, 2022:

A Machine Learning Model to Predict Hospital Admissions and Emergency Department Use in Patients' Immune Checkpoint Inhibitors. Machine Learning Healthcare Conference (MLHC) 2022 poster.
Niederhoffer K, Knechtle W, Uronis H, Shariff A, et al. (2022)

A Machine Learning Model to Predict Hospital Admissions and Emergency Department Use by Patients Receiving Immune Checkpoint Inhibitors. Duke University School of Medicine: Student Thesis
Niederhoffer K, et al. (2022)

irAE MODEL FEATURE RELATIVE IMPORTANCE VALUE

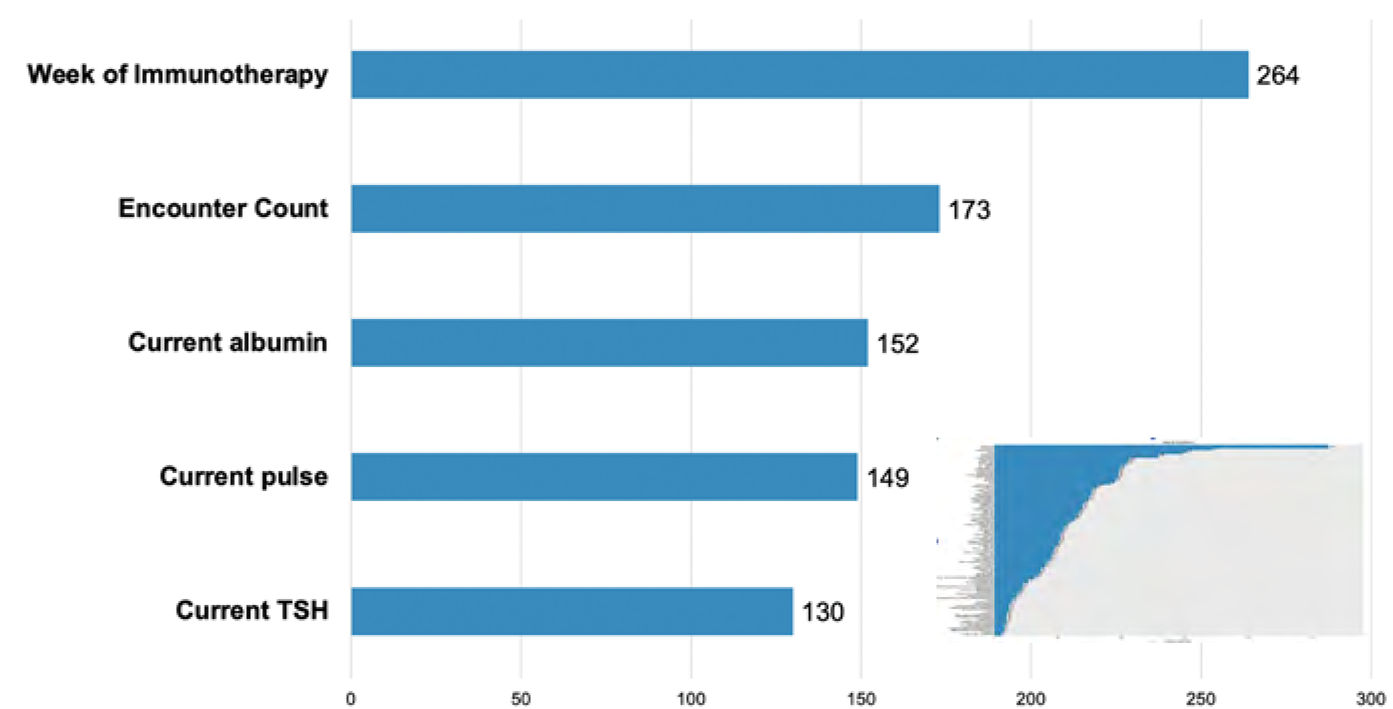

Figure 3. ICI irAE model measures of feature importance.

DIHI SCHOLAR EXPERIENCE

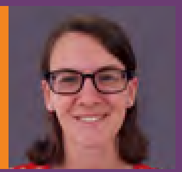

Kira Niederhoffer

The Duke Institute for Health Innovation (DIHI) gave me a unique opportunity to engage with leaders in healthcare innovation and artificial intelligence, develop my own data science skills, and interface with clinical leaders stepping outside traditional roles. I am deeply grateful for the opportunity it provided. My work over the year focused primarily on the project, "Predicting hospital admissions and emergency room visits from immune-related adverse events." This developed a machine learning (ML) model for cancer patients treated with immunotherapy. I had heard the buzzwords of "ML methods" and "artificial intelligence" before my time at DIHI, yet I had a limited ability to evaluate these techniques. I certainly had no firsthand experience with the process of creating them. Over the course of the year, I developed a concrete skill set in data science basics and ML project assessment. Most importantly, I perceived and valued the power this approach can bring to clinical questions. In my prior research in a translational oncology lab, we were lucky to enroll a few hundred patients in a study.

At DIHI, we were able to evaluate data from thousands of patients over the course of a day with a few lines of code. In addition to the hard skills of data analysis, interfacing with the various stakeholders on our project gave me a thorough appreciation for the importance of consensus building and diverse perspectives. We scholars learned to navigate a diverse group of stakeholders with slightly different priorities, and that an implementation pathway fitting clinical workflow is just as critical to model success as any performance statistic. Beyond the project itself, I cherished learning to balance the fundamentals of meaningful change: creativity and practicality, soft and hard skills, attention to detail, and big-picture thinking. I am deeply grateful for my experiences from the year and will strive to take the values back with me into clinical practice.

DIHI FELLOW EXPERIENCE

Kaivalya (Kai) Deshpande

The Duke Institute for Health Innovation (DIHI) provided a career segment filled with tremendous growth - catapulting me to heights that I otherwise would have never reached. Despite being equipped with a math background while completing training in medicine, my career plan fell short without DIHI. My desires were not yet met because this background did not bridge the full spectrum between math and medicine without applied research and machine learning (ML) modeling. Spanning the spectrum would not have been possible without the right environment: DIHI. The atmosphere at DIHI allowed me to extend myself as I learned coding in the Python programming language, familiarized myself with the different Python libraries (a set of useful functions that eliminate the need to write code from scratch), and was immersed in a fluid creative process. The learned coding skills were directly used to capture data inputs; process, clean, and store data; and prepare interacting designs and capabilities to put the model into practice.

Additionally, working on different projects illuminated the importance of combining a team of data scientists and medical professionals to create the optimal model and minimize chances of non-adoption or abandonment. I have learned many such lessons not only through direct experiences but also from the literature provided in teaching sessions and journal clubs of the DIHI curriculum. I have also had the opportunity to participate in conferences, lead group discussions at international conferences, and serve as a clinical reviewer for paper submissions (a great honor!). These lessons and opportunities were all a result of the sponsorship I received from DIHI; in particular, Dr. Mark Sendak. I am grateful to have experienced the pervasive sense of unity within the DIHI team, and I cannot think of a better set of helpful, hardworking, and fun team members.

Caremap: A Digital Personal Health Record for Complex Care Coordination

Problem

Children and youth with special healthcare needs (CYSHCN) and adults with multiple chronic conditions (MCC) represent the largest populations of high-need, high-cost patients.[2] CYSHCN and adults with MCC are common – they represent approximately 20% of all children and over 40% of all adults in the United States, respectively[3,4] – and despite their age differences, both groups share the common experience of frequently receiving fragmented, uncoordinated care.

Care coordination interventions for CYSHCN[5] and adults with MCC[6] are associated with better health outcomes; however, many receive inadequate care coordination.[7] For example, a subset of CYSHCN have the highest health needs and costs[8] because they often have multiple chronic conditions, intensive home care needs (e.g., home health), functional limitations that impact daily living (e.g. reliance on a feeding tube), and require long-term multi-specialty follow-up.[9] As a result, their "medical neighborhood"[10] can involve a dizzying team of providers and health and social agencies (Figure 1). Parents/caregivers spend substantial time and effort – often >20 hours per week[11] – coordinating. Importantly, this is very stressful for families of CYSHCN as they are primarily responsible for care coordination on their own. The overall care experience is poor, leaving many families isolated and unsupported. One parent of a CYSHCN cared for at Duke shared that: "We, the parents, were the keepers of [his] lifelong medical chart... we were lay people, with a very complicated child, who was only growing more complicated." This same stressful patient/family experience of fragmented, uncoordinated care affects adults with MCC and their

01

Team

David Ming, MD
Will Ratliff, MBA
Nitin Gujral, MBA
Sarah Gonzales
Willis Wong
Heather King, PhD
Nirmish Shah, MD
Richard Antonelli, MD

02

Project in Brief

Children and youth with special healthcare needs (CYSHCN) and adults with multiplate chronic conditions (MCC) represent the largest populations of high-need, high-cost patients and share the common experience of frequently receiving fragmented, uncoordinated care. There is a critical need for better care coordination solutions for parents of CYSHCN and adults with MCC. We refined a novel third party mobile application – Caremap – that integrates with electronic health records (EHR) to facilitate better coordinated care for this population. The app allows patients, families, and providers to securely organize and coordinate care longitudinally.

loved ones. In order to reduce care fragmentation, improve the patient/family experience, and improve health outcomes, there is a critical need for better care coordination solutions for parents of CYSHCN and adults with MCC.

To manage a CYSHCN or adult with MCC's lifelong medical record, patients and/or their families maintain extensive health records (e.g. history, procedures, medications) over time and across different health systems and providers in the forms of longitudinal care plans or paper binders.[12,13] However, such manually curated, often paper-based recordkeeping is time/labor-intensive and not synchronized with the electronic health record (EHR).

Solution

A personal health record (PHR) – a digital application through which parents can securely access, manage, and share their child's health information[14] – would be transformative because a digital PHR: (1) accesses and organizes information across multiple health information systems (e.g., different EHRs); (2) includes parent/patient-reported information (e.g., care goals, symptom burden), and (3) facilitates transparent, bi-directional digital communication of PHR contents with providers.

A team of clinicians, developers, and family partners at Duke Health and Boston Children's Hospital previously built a Fast Healthcare Interoperability Resources (FHIR)-enabled digital PHR mobile app called Caremap.[18] Caremap is an app to coordinate care for CYSHCN. However, FHIR-enabled EHR integration of the app as a care coordination solution in a real-world setting for CYSCHCN and adults with MCC is yet to be implemented.

The Caremap mobile app leverages the Apple iOS© platform and can be used on Apple© personal devices (tablet, phone). The Caremap mobile app has three core features (Figures 2 and 3). First is FHIR-enabled access to structured data from the child's EHR chart (e.g., medication list, allergies, problem list) and visualization of data within the app, thereby allowing patients/families to view the same health information seen by their Duke provider in the EHR. Second is a tracking feature so that patients/families can track their progress towards patient-centered care goals (e.g., fewer missed school days due to illness, and adherence to routine medications).Third is a clinician dashboard that allows providers to visualize a summary of patient/family-reported trends, progress, and insights within the EHR. These core features were co-designed with parents/caregivers and providers to tailor the app to best meet their needs. Previously conducted usability testing demonstrated proof-of-concept and strong support from families for Caremap as a promising digital health solution for care coordination.

> "We have completed significant activities in three core areas during the project period: embedment of participant recruitment within routine clinical workflows, information security, and technical development.

We are live with the use of Caremap mobile app and have begun enrolling participants as of September 2022. To achieve our go-live with Caremap, we have completed significant activities in three core areas during the project period: embedment of participant recruitment within routine clinical workflows, information security, and technical development. This included exciting integration of the clinician dashboard with the mobile app and Epic© EHR via a partnership with a third-party platform (Xealth©). Furthermore, we completed the Duke School of

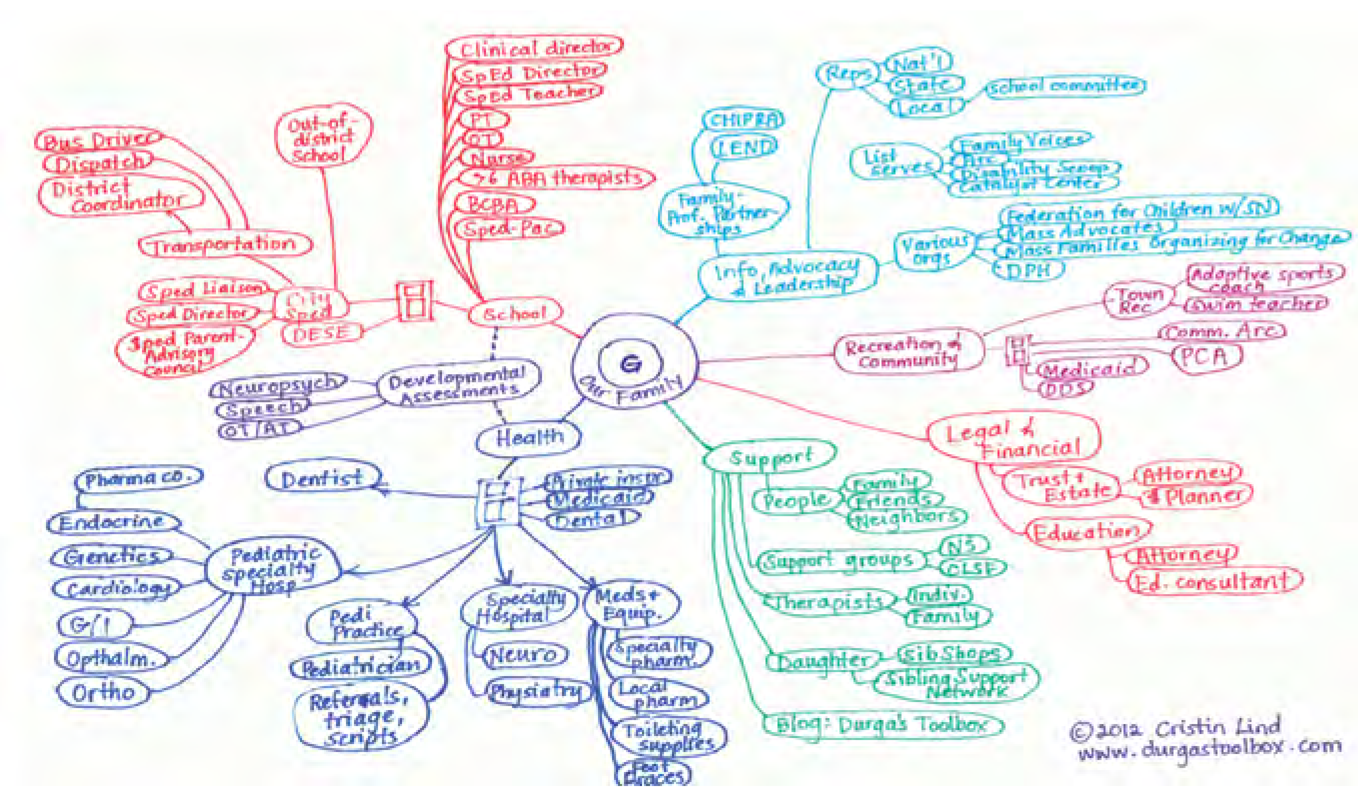

Figure 1. Care map of services and providers in one child's medical neighborhood

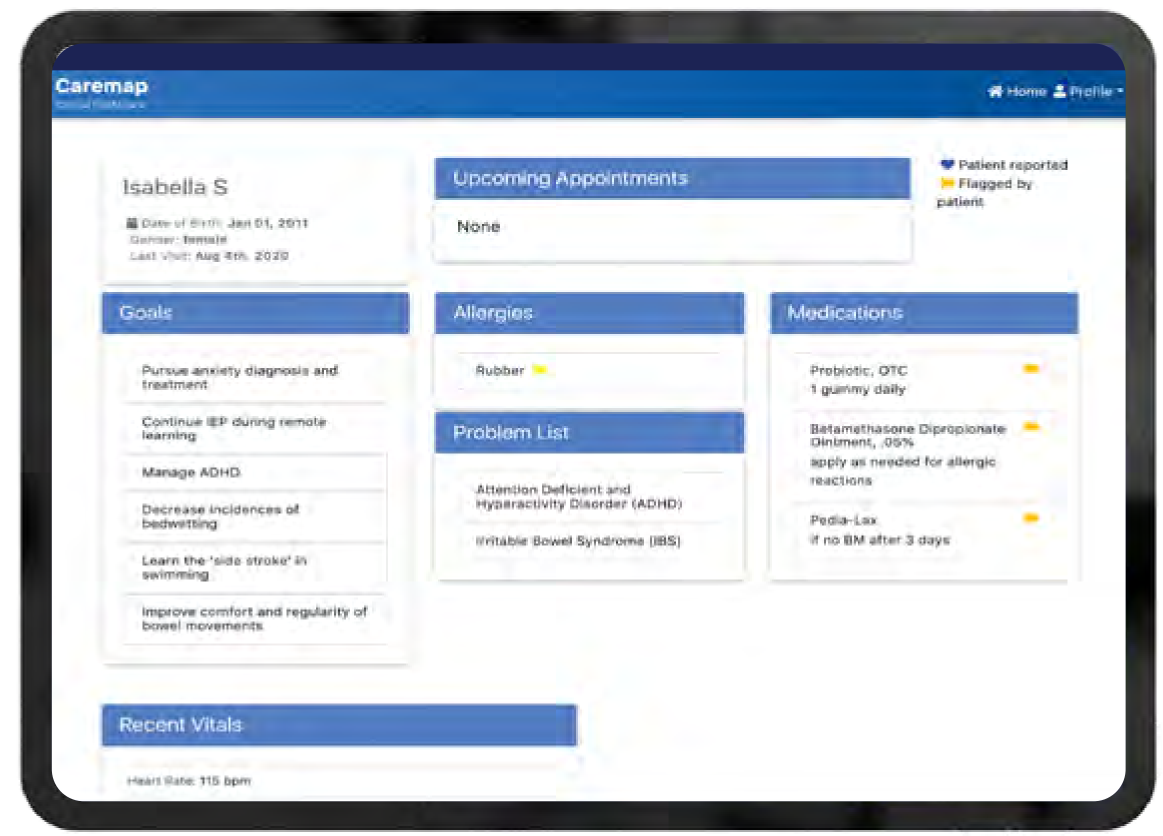

Figure 2. Caremap Clinician's Dashboard Prototype Interface

Prototype clinician's dashboard built into Epic© EHR showing overview of the patient's pertinent health information, including active care goals, allergies, problems list and medications. The dashboard also shows items flagged by the Caremap mobile app user (yellow flags) for review. Lastly, the dashboard also summarized family-reported outcomes for the patient.

Medicine and Duke Information Security Office (ISO) regulatory processes to ensure the security of data systems involved in the mobile app. This included meeting stringent review standards and comprehensive third-party penetration testing. Finally, we developed comprehensive study recruitment procedures and infrastructure using a novel digital prescribing process embedded within routine workflows. The digital prescription ('e-prescription') is a platform that leverages the EHR online patient portal (Epic MyChart©) to deliver all study-related materials – including onboarding materials (newly created video and written content), app download link, and informed consent – electronically directly to patients/families. Partnership with clinician champions in four Duke specialty care sites – two pediatric (pediatric pulmonology; neurodevelopmental pediatrics) and two adults (geriatrics; pulmonary transplant) – allowed the creation of workflows that will seamlessly integrate study recruitment and patient/family use of the app within clinical care at participating sites.

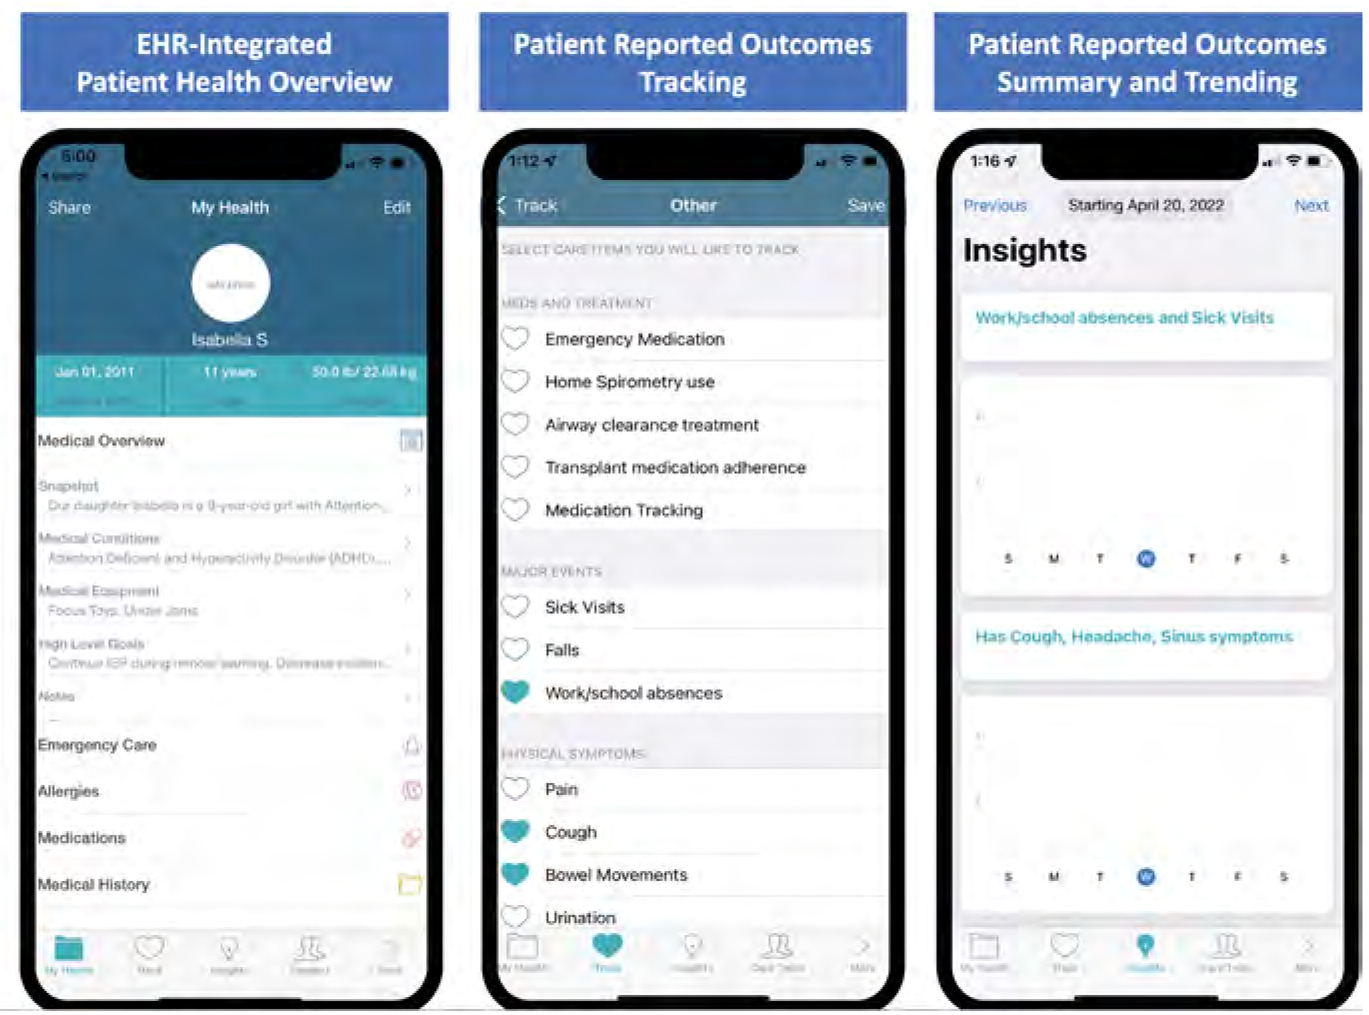

Figure 3. Caremap mobile app screenshots.
From left to right: 'My Health' tab showing personal health records features; 'Tracking' tab showing sample of trackable family reported health outcomes for users; 'Insights' tab showing graphical summary design of user-tracked outcomes; 'Care Team' tab displaying contact cards for relevant members of the care ecosystem.

Impact

Our team's experiences generated critical lessons learned that will be foundational information for other Duke innovators seeking to efficiently navigate the digital security review process with future EHR-integrated mobile apps.

We aim to enroll forty participants (ten participants per site) and gather prospective quantitative and user-reported data that will determine feasibility and impact of the mobile app on care coordination for CYSHCN and adults with MCC in specialty care settings.

Next Steps

We plan to collect qualitative data via semi-structured participant interviews in order to identify barriers and facilitators to implementation. We plan to complete the steps in the invention disclosure application process and have collaborated with internal and external partners to integrate the Caremap mobile app platform into interventions submitted for external grants.

Academic output

We presented this study protocol as a peer-reviewed research abstract poster at the national Pediatric Academic Societies' annual meeting in spring 2022. We have two peer-reviewed manuscripts in progress that were submitted for publication in fall 2022.

THANKS: These technical outcomes were made possible through extensive collaboration with multiple partners and groups internal (e.g., Digital Strategy Office; Duke Health Technology Services; Departments of Pediatrics, Medicine, and Population Sciences; Office of Research Administration) and external to Duke (e.g., Boston Children's Hospital).

References

1. Hospital BCs. Organizing Care and Relationships for Families: Care Map. http://www.childrenshospital.org/integrated-care-program/care-mapping. Accessed July 17, 2020.
2. Blumenthal D, Chernof B, Fulmer T, Lumpkin J, Selberg J. Caring for High-Need, High-Cost Patients - An Urgent Priority. The New England journal of medicine. 2016;375(10):909-911.
3. Musumeci MC, Priya Medicaid's Role for Children with Special Health Care Needs: A Look at Eligibility, Services, and Spending. San Francisco, CA June 12, 2019 2019.
4. Buttorff C RT, Bauman M. Multiple Chronic Conditions in the United States. Santa Monica, CA: RAND Corporation;2017.
5. Council on Children with D, Medical Home Implementation Project Advisory C. Patient- and family-centered care coordination: a framework for integrating care for children and youth across multiple systems. Pediatrics. 2014;133(5):e1451-1460.
6. Vasan A, Morgan JW, Mitra N, et al. Effects of a standardized community health worker intervention on hospitalization among disadvantaged patients with multiple chronic conditions: A pooled analysis of three clinical trials. Health services research. 2020;55 Suppl 2:894-901.
7. Cordeiro A, Davis RK, Antonelli R, et al. Care Coordination for Children and Youth With Special Health Care Needs: National Survey Results. Clinical pediatrics. 2018;57(12):1398-1408.
8. Cohen E, Berry JG, Sanders L, Schor EL, Wise PH. Status Complexicus? The Emergence of Pediatric Complex Care. Pediatrics. 2018;141(Suppl 3):S202-S211.
9. Cohen E, Kuo DZ, Agrawal R, et al. Children with medical complexity: an emerging population for clinical and research initiatives. Pediatrics. 2011;127(3):529-538.
10. Greenberg JO, Barnett ML, Spinks MA, Dudley JC, Frolkis JP. The "medical neighborhood": integrating primary and specialty care for ambulatory patients. JAMA internal medicine. 2014;174(3):454-457.
11. Kuo DZ, Cohen E, Agrawal R, Berry JG, Casey PH. A national profile of caregiver challenges among more medically complex children with special health care needs. Archives of pediatrics & adolescent medicine. 2011;165(11):1020-1026.
12. Klitzner TS, Rabbitt LA, Chang RR. Benefits of care coordination for children with complex disease: a pilot medical home project in a resident teaching clinic. The Journal of pediatrics. 2010;156(6):1006-1010.
13. McAllister JW, Keehn RM, Rodgers R, Lock TM. Care Coordination Using a Shared Plan of Care Approach: From Model to Practice. J Pediatr Nurs. 2018;43:88-96.
14. Saripalle R, Runyan C, Russell M. Using HL7 FHIR to achieve interoperability in patient health record. J Biomed Inform. 2019;94:103188.
15. Braunstein ML. Healthcare in the Age of Interoperability: Part 3. IEEE Pulse. 2019;10(1):26-29.
16. Mandel JC, Kreda DA, Mandl KD, Kohane IS, Ramoni RB. SMART on FHIR: a standards-based, interoperable apps platform for electronic health records. Journal of the American Medical Informatics Association : JAMIA. 2016;23(5):899-908.
17. Bloomfield RA, Jr., Polo-Wood F, Mandel JC, Mandl KD. Opening the Duke electronic health record to apps: Implementing SMART on FHIR. Int J Med Inform. 2017;99:1-10.
18. Caremap. http://caremap.health/index.html. Published 2016. Accessed.

DIHI STAFF PERSPECTIVE

Mark Sendak

So You Want to Transform Health and Healthcare? Get Out of the Valley

In April 2021, our team was invited to present at the annual Artificial Intelligence in Medical Imaging conference at Stanford. The instructions asked speakers to present an idea that is new or surprising, or challenges a belief the audience already has. Having grown up in NorCal (native speak for Northern California), I was familiar with the ideas and assumptions the audience at Stanford would have about health innovation. One that I held onto until recently goes something like this: the transformation of healthcare into Artificial Intelligence/Machine Learning (AI/ML)–enabled services will be centered in Silicon Valley and will be driven by technology firms.

> “Healthcare is local and the geographic disparities and differences in America across measures of cost, quality, and outcomes are striking.

That strongly held belief is wrong. The transformation of healthcare into AI/ML-enabled services will be centered in geographies with significant socioeconomic diversity, chronic disease burden, and consolidated markets. In other words, not Silicon Valley, which features extreme wealth, a young population, and fierce competition among legacy and emerging players. In addition, the transformation will be driven by forward-thinking healthcare delivery organizations that gradually embed AI/ML products into professional workflows.

Geographic Advantage

Healthcare is local, and the geographic disparities and differences in America across measures of cost, quality, and outcomes are striking. Poverty, inequality, diabetes, obesity, and many other chronic conditions that impact health outcomes disproportionately affect people in the American South and Southeast. Political polarization is also rampant in the South and Southeast, which shapes how individuals interact with healthcare providers to change behavior and manage chronic conditions. As the famous entrepreneur Steve Blank says, when you're trying to build a great product or service, you need to "get out of the building."[1] You need to talk to people, seeking input from potential users. In healthcare, that means going to where the problems are and deeply embedding yourself in the social and political environments that propagate the problems.

In North Carolina, not only are the healthcare problems massive, but the levers to drive change are also particularly strong. There is strong insurance concentration, with Blue Cross Blue Shield capturing 71%[2] of the large group insurance market and 97% of the individual market.[3] Similarly, there is strong provider market concentration, with Duke Health operating every inpatient bed in Durham County (estimated population 325,000)[4] along with the largest primary care network in the greater Triangle area (estimated population over two million).[5] This means that through risk-based contracts, Duke Health is able to successfully execute population health management programs with a small number of key partners. This also means that longitudinal data across care delivery settings is abundant for a large population of patients. Taken together, North Carolina, and Durham specifically, is fertile ground for healthcare transformation.

Professional Service Advantage

In 2017, Ziad Obermeyer and Thomas Lee wrote in the New England Journal of Medicine: "There is little doubt that algorithms will transform the thinking underlying medicine. The only question is whether this transformation will be driven by forces from within or outside the field. If medicine wishes to stay in control of its own future, physicians will not only have to embrace algorithms, they will also have to excel at developing and evaluating them, bringing machine-learning methods into the medical domain."[6]

At the Duke Institute for Health Innovation (DIHI), front-line clinicians play a critical role in all phases of a project, ranging from problem identification to peer training and education during roll-out of the innovation. We are also proactively training a new generation of physicians who deeply understand the AI/ML model development, evaluation, and integration lifecycle.[7] By driving digital transformation alongside clinical partners, our team has enabled Duke Health to more rapidly integrate a portfolio of AI/ML systems into the Electronic Health Record than any other health system in America.[8]

Even so, the competitive advantage of healthcare delivery firms goes beyond clinician engagement. For enterprises to effectively integrate AI/ML into operations, significant effort is required to reconfigure the organization.[9] In our own setting, this has required various forms of incentive alignment among stakeholders and repair work to ensure that disruptions to workflows are balanced by immediate gains and benefits for end-users.[10, 11] This can mean creating new reimbursement mechanisms, performance measures, or reporting structures. These types of organizational changes require strong buy-in from senior leadership and front-line workers, as well as, a deep sense of trust with the technology supplier. By operating within a health service organization, we have an unfair advantage to rapidly and repeatedly drive organizational change to integrate AI/ML.

Thankfully, the crowd at Stanford was kind and didn't throw me out. After the talk, attendees from Big Tech companies asked how they could better partner with teams like ours at Duke Health. The answer was simple. Solve a problem we have at Duke Health better than we can solve it ourselves. The first step towards achieving that would be to get out of the Valley.

References

1. https://www.youtube.com/watch?v=a-J_SwmMJyo
2. Market Share and Enrollment of Largest Three Insurers – Large Group Market. Kaiser Family Foundation. https://www.kff.org/other/state-indicator/market-share-and-enrollment-of-largest-three-insurers-large-group-market/?currentTimeframe=0&sortModel=%7B%22colId%22:%22Location%22,%22sort%22:%22asc%22%7D
3. Market Share and Enrollment of Largest Three Insurers – Individual Market. Kaiser Family Foundation. https://www.kff.org/private-insurance/state-indicator/market-share-and-enrollment-of-largest-three-insurers-individual-market/?currentTimeframe=0&sortModel=%7B%22colId%22:%22Location%22,%22sort%22:%22asc%22%7D
4. https://www.census.gov/quickfacts/durhamcountynorthcarolina
5. https://corporate.dukehealth.org/clinical-care/health-system-overview
6. Obermeyer Z, Lee TH. (2017) Lost in thought – the limits of the human mind and the future of medicine. N Engl J Med; 377:1209-1211. DOI: 10.1056/NEJMp1705348 https://www.nejm.org/doi/10.1056/NEJMp1705348
7. Sendak MP, Gao M, Ratliff W, Whalen K, Nichols M, Futoma J, Balu S. (2021). Preliminary results of a clinical research and innovation scholarship to prepare medical students to lead innovations in health care. Healthc (Amst). 9(3):100555. doi: 10.1016/j.hjdsi.2021.100555. Epub 2021 May 3.
8. Sendak MP, D'Arcy J, Kashyap S, Gao M, Nichols M, Corey, Ratliff W, Balu S. (2020). A path for translation of machine learning products into healthcare delivery. EMJ Innov. DOI/10.33589/emjiinnov/19-00172.
9. Elish MC, Mateescu A. (2019). AI in context: the labor of integrating new technologies. Data & Society. https://datasociety.net/library/ai-in-context/
10. Elish MC, Watkins EA. (2020) Repairing innovation: A study of integrating AI in clinical care. Data & Society https://datasociety.net/library/repairing-innovation/
11. Kellogg KC, Sendak M, Balu S. (2022). AI on the front lines. MIT Sloan Management Review. https://sloanreview.mit.edu/article/ai-on-the-front-lines/

A Geriatric-Specific Morbidity and Mortality Risk Stratification Tool

Problem

Adults over the age of sixty-five account for nearly 40% of inpatient surgical procedures in the US, and their hospital stays often require heavy resource utilization.[1,2] They are a highly vulnerable patient population at increased risk of perioperative morbidity and mortality. As the mean age of the United States population rises, the ability of a surgeon to make decisions by accurately incorporating goals of care, the utility of surgery, and the risk of harm is becoming an increasingly critical skill. To effectively identify which patients are at high risk for surgery, physicians require objective measures of the clinical factors that confer this elevated risk status. At project inception, there were no standardized geriatric-specific perioperative risk stratification tools that incorporated pertinent variables such as function, cognition, nutrition, polypharmacy, and code status.

Solution

We trained separate models to predict surgical outcomes at the time the surgical team decided to schedule a surgery and the day before surgery. Separate models were needed as the surgery and scheduling decisions are made well before the preoperative interventions. The model fitting our use case informs the prediction of surgery at the time of scheduling. The other informs the prediction of surgical outcomes on the day before surgery. We studied how the scheduling-to-surgery time interval and interventions impact model performance and decided to present both models' results to perioperative teams at their respective dates. Furthermore, we redesigned the user interface to display model results for any combination of medical record number (MRN) and current procedure terminology (CPT, the codes by which surgeries are identified for billing) within seconds (Figure 1). This allowed the surgical team to explore surgical opportunities for any patient at any time.

Impact

We learned that the optimal utility of surgical predictive models required accurate model results well before the date of surgery and even before the surgical case was scheduled. We realized that the original model required substantial adaption to meet this need. Statistics from the thirty-day all-surgery case creation and admission models designed to this end are presented in Figure 2. The AUC for the 30-day mortality model trained with data before case creation and evaluated on patients 65 years old or older was 0.78. Furthermore, we learned that structured data supporting Mini-Cognitive Assessment scores, the Duke Activity Status Index (DASI) score, and Best Evaluation Systems Test (MiniBEST) scores was not consistent through the model training and validation cohort from 2015-2019, principally because the scoring and documentation processes were implemented at Duke more consistently in the years (2019-2022). In the spirit of continuous improvement, we recommend strident maintenance of score documentation, exploring enhanced model versions with 2018-2022 data, and consideration of processing language from surgical progress notes. Starting September 2022, steps were also be taken to increase data availability through a remote patient monitoring study.

Next Steps

We will present and silently test the mobile and desktop applications with surgeons and clinicians who were high users of previous models. Without influencing clinical decision-making, the surgeons will provide feedback about model performance on recent retrospective patients, application design, and ease of use.

Subsequently, the user applicability of the models will be studied and tested in the surgical outpatient and preoperative anesthesia surgical screening clinics. This continued study is supported by Duke Health leadership and the DIHI-funded project, Improving Perioperative Care Coordination via Enhanced Decision Support Tools, led by Dr. Jeanna Blitz.

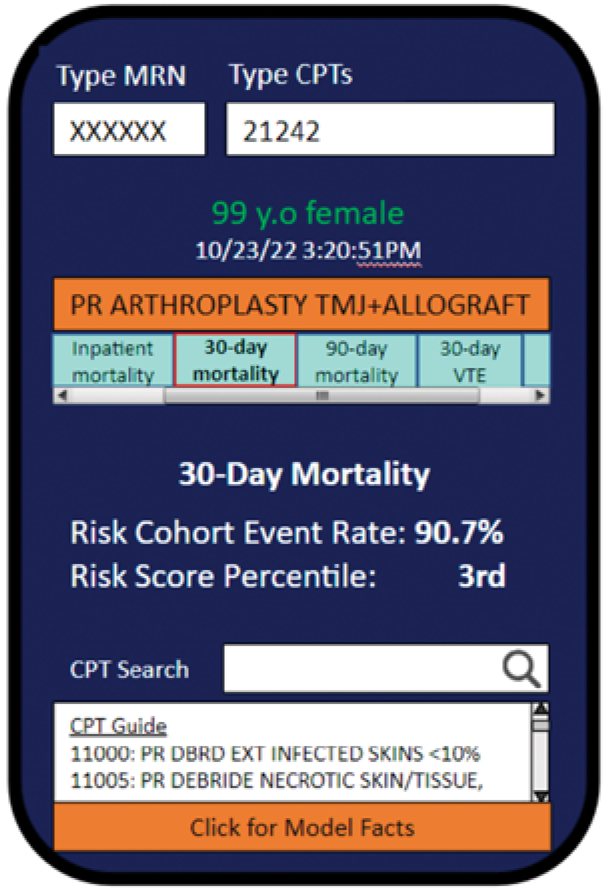

Figure 1. User interface for surgical prediction models by DIHI

01 Team

Hayley Premo | Harvey Shi
Will Knechtle MBA, MPH | Faraz Yashar
Sai Harshith Rachakonda | Marshall Nichols, MS
Michael Gao, MS | Mark Sendak, MD, MPP
Willie Boag, PhD | Bradley Hintze, PhD
Suresh Balu, MS, MBA | Jeanna Blitz, MD/FASA
Shelley McDonald, DO/PhD | Hadiza S Kazaure, MD
Sandhya Lagoo-Deenadayalan, MD/PhD

02 Project in Brief

PROBLEM:
Flow of patient care between outpatient visits, PASS/PAT visits, POSH visits, and surgery is suboptimal. Adults ≥65 comprise approximately 35-40% of annual surgical volume and have increased risk of perioperative mortality, morbidity, and resource consumption.

SOLUTION:
Enhance surgical outcome predictions by tailoring them to older surgical patients, incorporating variables relating to function, cognition, polypharmacy, and code status. This tool will aid triage with risk-tiered workflows, leading to efficient perioperative care.

IMPACT:
Learned ways to optimize the model and user interface to immediately provide the right information at the most useful times for the preoperative team to make decisions. Developed, trained, and tested multiple predictive models for different points in the care continuum.

TIME OF OUTPUT	AUC	PRC
CASE CREATION	0.86 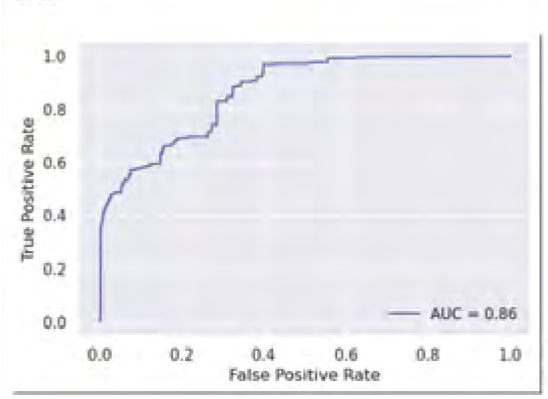	0.27 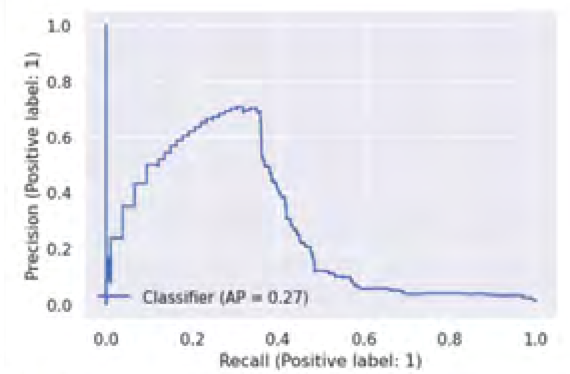
ADMISSION	0.80 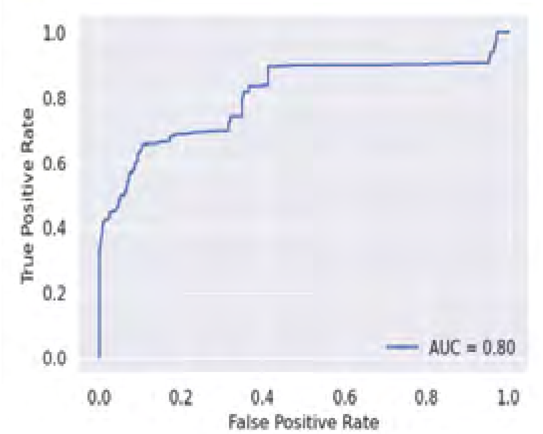	0.40 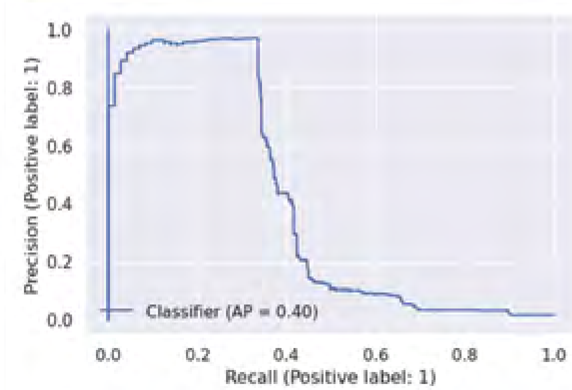

Figure 2. AUROC and AUPRC Curves for 30-day mortality models for all surgical cases. (218,016 encounters; 161,414 patients) designed for two different dates of the surgical perioperative process.

Academic output

POSTER PRESENTATION DURHAM, NC, 2022:

A Geriatric-Specific Morbidity and Mortality Perioperative Risk Stratification Tool. Machine Learning for Healthcare (MLHC) 2022.

Premo H, Shi H, Knechtle W, Kazaure H, et al. (2022).

References

1. McDermott KW (IBM Watson Health), Freeman WJ (AHRQ), ElixhauserA (AHRQ). Overview of Operating Room Procedures During Inpatient Stays in U.S. Hospitals, 2014. HCUP Statistical Brief #233. December 2017. Agency for Healthcare Research and Quality, Rockville, MD.
2. Dall T.M., Gallo P.D., Chakrabarti R., et. al. An aging population and growing disease burden will require a large and specialized health care workforce by 2025.Health Aff 2013; 32: pp. 2013-2020.
3. Corey KM, Kashyap S, Lorenzi E, et al. Development and validation of machine learning models to identify high-risk surgical patients using automatically curated electronic health record data (Pythia): A retrospective, single-site study. PLoS Med. 2018;15(11):e1002701. Published 2018 Nov 27. doi:10.1371/journal.pmed.100270

DIHI STAFF PERSPECTIVE

Marshall Nichols

The Unknown Unknowns of Data Used in AI/ML

The Electronic Health Record (EHR) is a shockingly messy place. I don't believe that statement is a surprise to anyone who has worked with or used EHRs since their beginning in the 1960s, growth in the 1980s, or promotion of their "meaningful use" by the Health Information Technology for Economic and Clinical Health (HITECH) act in 2009. At its core, the EHR is designed to serve three principal purposes:

1. Detailed chronological documentation of exactly what happened to cover any legal audit trail.
2. Accurately account for activities to enable appropriate billing to support the business of health care.
3. Capture and provide access to patient information so that a trained care practitioner can improve the health and lives of patients.

The EHR attempts to capture relevant patient information consistently enough to service these three goals. The third purpose is benevolent and the best-marketed.

As a data engineer who regularly collaborates with care practitioners, I have received unfiltered opinions about, and have experienced, the grime of EHR design. The primary focus on auditing and accounting over design for practitioners' use is obvious. The EHR front-end and back-end interfaces are far from intuitive. Seasoned and trained professionals regularly have challenges knowing where to go to find or document the discrete information they obtain during real-world interactions with their patients. Knowing this, it becomes easier to understand why practitioners highly value documentation for understanding the continuity of their patients' experience. Practitioners value unstructured notes and their patients so much that they sacrifice their own health and burn themselves out by extending clinical note documentation sessions well into their personal evening hours. If improving the lives of patients and enabling their practitioners were a central priority at the inception of EHRs, it would have been designed differently!

> "The EHR attempts to capture relevant patient information consistently enough to service these three goals.

There is still plenty of useful information in the EHR, but shadowing physicians reveals the subtle attitude differences impacted by the structured billing process vs. unstructured documentation. The billable process focuses on the status of care processes: what has been ordered (like a creatinine test), what tests were resulted, if a medication has been administered, if a patient was transferred to a different floor, etc. This information is readily available and helps audit, bill, and manage the patient experience. Nevertheless, if a care practitioner or healthcare analyst wants to know what's going on with the patient, notes are where they go today. Why? You may be able to find a structured value of a test result (e.g. A1C=7.9 on 10/05/2022 at 11:15 AM EST) but notes may tell you why. Let's examine the anonymous anecdote in Table 1 to understand the difference.

IMPACT OF EHR FOCUSED ON BILLING PROCESS	IMPACT OF NOTES
"I ordered a blood glucose test that had a result that was higher than expected. To cover our bases, we should order an A1C and a few other things that are justified by the elevated glucose lab test result. I'm also placing an order for a follow-up with a Nutritionist."	"My patient has high blood sugar and has difficulty accessing healthy food. I'd like to recommend that they follow up with a nutritionist, but it's unlikely that they'd be able to consistently make the appointments. Consider health system pathways to assist patient with access prior to nutritionist recommendation. Should be addressed with larger care team."
• Emphasis is placed indirectly on 'bill-ability' • Justifies further testing based on results • Things are being done to/for the patient without a process that is aware of overarching patient context	• Relays physician understanding of the patient situation. • Indirectly acknowledges the discrete-data-agnostic issues that must be addressed before treatment / follow-up can be effective for the patient.

Table 1. How styles of documentation affect content and care practitioner attitude

The current EHR design not only makes life unnecessarily difficult for care practitioners, it also creates difficulties for data engineers, who sift through these data to glean some beneficial insight or summary. The EHR is most interested in capturing what happened: a transaction. The fact that these transactions (events, orders, medication administrations, transfers, etc.) occurred is captured consistently. However, the discrete data that is captured as part of these transactions varies wildly. When we start digging deeper into these discrete data, the analytics processes consuming these raw EHR data come to a grinding halt.

Here is an example which highlights the situation: A creatinine lab test was ordered, but which one? At the time of this article's publication, there were ninety nine different laboratory tests with the word 'CREATININE' in their active Duke Health labels. Some are ratios, some are scaled to body weight, some are whole blood measurements, some are urine.

We start seeing names like these:

CREATININE, WHOLE BLOOD (BKR)
CREATININE WHOLE BLOOD (BKR)
CREATININE (BKR)
CREATININE LABCORP

Are these all creatinine measurements? Probably. Does any individual know enough about all 99 different tests to be able to make sense of them? Is there a system to capture the institutional knowledge from this person? Can this proliferation of names be arrested? Unclear.

Next, we start asking questions to assess the usability of the data. How consistent should we expect the results of these lab tests to be? We should check to see what units the test results come back in!

CREATININE mg/dL
CREATININE (BKR) mg/dL
CREATININE NULL
CREATININE (BKR) NULL

So far so good: what we find is fairly consistent around some version of milligram per deciliter with some identically named elements missing units! But, we also observe inconsistent measures like:

BUN/CREATININE RATIO EXTERNAL MG/DL
BUN/CREATININE RATIO EXTERNAL RATIO
BUN/CREATININE RATIO EXTERNAL RATIO

Ratios that come back with real units adds significant complexity and liabilities:

CR (CREATININE EXTERNAL) MG/DL
CR (CREATININE EXTERNAL) G/DL
CR (CREATININE EXTERNAL) /DL
CR (CREATININE EXTERNAL) K/UL
CR (CREATININE EXTERNAL) MG/G
CR (CREATININE EXTERNAL) MG/D

In the past seven years since July 2014, our EHR has collected ~630 million results associated with orders, with the majority of these (~430 million) being from laboratory tests. These ~630 million results are composed of ~25,500 unique result names and unit combinations. Within this set there are ~13,300 unique result names and unit combinations associated with laboratory tests.

Unification of these results into 'common', analyzable units or an even broader 'common understanding' of what a lab test result is, without prior clinical knowledge, is extremely challenging. There are worldwide efforts ongoing to standardize this information into a consumable form. LOINC codes are one example of this effort. These are highly detailed codifications of hierarchies of laboratory tests and their interpretations. The Duke Clinical Lab management team has done a great job applying these LOINC codes to appropriate laboratory test results produced internally. Of the ~430 million lab test results, ~300 million have an associated LOINC code we can use to assist with aggregation. Unfortunately, this leaves ~130 million records, typically lab test results coming back from external vendors, or tests tied to particular lab instrumentation

DIHI SCHOLAR EXPERIENCE

Gaurav Sirdeshmukh

I am a senior at Duke majoring in Statistics with minors in Chemistry and Mathematics and a concentration in Data Science. In the spring of 2022, I joined the Duke Institute for Health Innovation (DIHI) as a data science intern and worked with Dr. Mark Sendak to develop the Machine Learning Data Quality Assurance (ML-DQA) framework. The ML-DQA framework is designed to create and assign comprehensive data quality checks for ML projects. An especially rewarding aspect of the project is the potential impact of the framework: it was applied to five different ML applications in multiple healthcare settings and will be an integral part of DIHI's data quality assurance process going forward. In August of 2022, I was extremely excited that our research paper describing the ML-DQA framework was published and presented at the Machine Learning for Healthcare (ML4HC) conference. I enjoyed attending the conference, sharing our research, and learning about the tremendous work that is being done by the ML community.

DIHI's diverse team of data scientists, scholars, and undergraduate interns are incredibly good at what they do and fun to work with. The standard and quality of work are high and the DIHI team consistently meets these expectations. I think this positive culture stems from the collaborative nature of the group and its prioritization of learning and innovating. DIHI is a unique environment, and I look forward to spending my senior year at DIHI. I'm glad to be a part of the team!

like “REMISOL BUN/CREAT RATIO (BKR)” that have no LOINC code and often have missing, confusing, or explicitly incorrect unit representations (see CR (CREATININE EXTERNAL) above).

> Since July 2014 our EHR has collected ~630 million results associated with orders, the majority of these (~430 million) are from laboratory tests.

One particularly jarring issue with vendor-returned results that complicate downstream analysis and should be addressed promptly — cell or unit counts that return with units like `x10E3`. Another kudos to the Duke Clinical Lab management team, internal measurements of these counts are returned with correct units like `x10^3`. The vendor’s unit `x10E3` in long form is `10 * 10^3` or `1 * 10^4`. The scale of the internal numeric results represented by `x10^3` and the vendor results represented by `x10E3`, are the same. Meaning it’s possible to have a White Blood Cell Count of 9.1 x10^9 cells/L from a Clinical lab measurement and 9.1 x10E3 cells/uL from a vendor measurement. Once converted to the preferred `x10^9/L` unit, the vendor measurement becomes 91 x10^9 cells/L. A dangerously incorrect result!

The amount of prior knowledge of the content and structure of an EHR, the nuances of the particular implementation, the legal barriers between Epic, and the ability to start even gaining this knowledge are disappointing at best. The FHIR initiative may help with some of this as it provides a universal, common data interchange format. It would be great for abstracting us away from the EPIC-ness of our EHRs; but even FHIR won’t address these core, critical inconsistencies with the data we collect every day. After an immense amount of FHIR migration effort, our data model will be common, and we will understand its structure, but it won’t help analysis.

This is where the EHR’s design preference for ‘knowing a thing happened’ over ‘curation of exactly what happened’ begins to show. As expected, the most frequently used lab tests tend to be more consistent in their returned-results than the ones not used often. Notably, data and meta data for tests managed by the Duke Clinical Lab team tend to be much more consistent and well-maintained than the tests farmed out to vendors. However, to the data scientist or engineer trying to use all available creatinine measurements to build a model to help predict patient kidney function trajectories, this is a minefield. Unfortunately, this hazard is more the norm than the exception.

We at DIHI have long worked with our own internal ‘Data Pipeline’ to provide tooling that attempts to address some of these content concerns. However, we run into issues dealing with the compute performance required to deliver millions of unified, converted, and validated results in real-time. We’ve been able to deliver this performance across smaller subsets of data feeding our models, but have still struggled with universal conversion for the entire EHR. Luckily, this tooling is under constant development and we have some upgrades landing in 2022-2023 that should improve EHR-wide data curation performance and help transparently address the inconsistencies we find in our source data. These changes are all to support the growing business need for real-time model implementations and quick-access filtering of highly complex datasets. More to come, stay tuned!

Organization and Clean Up of the Electronic Health Record Problem List

01

Team

Eugenia R McPeek Hinz, MS
Lisa Nadler, MD
Will Ratliff, MBA
Matt Gardner
Suresh Balu MS, MBA
Momen Wahidi, MD

02

Project in Brief

The Problem List in the Electronic Health Record (EHR) is a tool used to communicate a patient’s current, active and notable medical diagnoses as well as symptoms and contributing health conditions. Unfortunately most PLs are incomplete at best to inaccurate overloaded with inconsistent or inaccurate diagnoses that are relics from the past and do not reflect the patient’s current health status. In this project we took steps to improve the organization of the PL and soon will implement automated cleanup of select clinical entities for the Duke Health problem list across more than two million patients in our EHR.

Problem

The Electronic Health Record (EHR) problem list (PL) supports communication of the patient’s problems across a wide range of clinical environments and patient caregivers. An accurate problem list serves as a foundation for clinical care, population health management and multiple secondary processes including research and severity of illness risk scoring. Maintaining an up-to-date problem list is essential to patient-centered care but is often a secondary process compared to the needs of direct patient care.

The PL suffers from an accumulation of diagnoses that can be either outdated, no longer accurate, duplicative, similar or conflicting.[1] A recent study on problem list completeness found wide variations of for diagnoses that have been used as visit diagnoses for inclusion in problems such as diabetes or asthma.[2] Another large cohort study found 22% incidence of the common diagnosis hypertension to not be included in the problem list and multiple duplications of diagnoses like asthma, Crohn’s disease and diabetes.[3]

In addition to the problem of volume and accuracy of the PL content, PLs are also not organized in a format that supports easy identification of related or outdated diagnoses, which also hinders cleanup. In our investigations, we found anywhere from 20% to 50% of diagnoses on a patient’s PL are not current and up to date for the patient. Because EHR data for the PL is at the individual level, broader population-level reviews of the inaccuracies have been not well identified until this project was undertaken.

This summary seeks to describe the work accomplished to improve the organization and cleanup of the Duke Health problem list in our EHR.

Solution

We developed 21 system/condition-based groupers using SNOMED-CT hierarchal concepts refined with Boolean logic (Figure 1). System groupers included traditional medical specialty categories and clinically relevant care coordination and procedure-based groupings. Specialty organization of diagnoses are an improvement over foundation default PL organization in Epic that organizes problem lists primarily in an alphabetized format. While there are other options for PL display such as priority organization, these require the user to apply a prioritization level one by one and are not used by a majority of our clinical users.

SNOMED-CT codes are translated in a one-to-many format for ICD-10 codes, allowing for more complete ICD10 groupers out of the potential more than 100,000 ICD10 codes available. For example, the Neurology specialty grouper with 167 SNOMED-CT concepts mapped to 9,243 ICD10 codes.

Outcome

The 21 specialty groupers were iteratively built by the primary study lead in the summer/fall of 2021 and implemented to be available for all Duke users in Epic in Dec 2021 and updated again in July 2022. All new users and most ancillary staff are now defaulted to new System Level organization of the PL for Duke Health users. Feedback has been very positive from all users but especially for staff who do chart reviews. To better understand the potential opportunity for automated cleanup of the PL, we analyzed a representative sample of PL diagnoses across 79 patients identified for specific focus on Oncology, Cardiology, Neurology and Orthopedics diagnoses. This format supported analysis of PL content and disorganization at more of an aggregate level.

Across the 79 patient cohort we identified 2,835 diagnoses ranging from 112 to 5 diagnoses with a mean of 36 diagnoses per patient. We found 1,508 (53.2%) singular concept diagnoses (i.e., diagnoses without duplicate, conflicting or similar diagnoses). In comparison, we found 1,327 (46.8%) diagnoses across 634 concepts with duplicate, related, conflicting, similar, lapsed or no longer active diagnoses. Notably this review was of the PL diagnoses alone and did not include chart review. The goal of this evaluation was to identify diagnoses that could be potentially automatically resolved or placed on the past medical history section.

In characterizing the state of the PL diagnoses with potential opportunities for cleanup, we found 7.7% of diagnoses were lapsed and potential targets for automatic conversion to history of diagnoses. The majority of these being for acute myocardial infarctions or acute stroke diagnoses where the date of onset of the problem was clearly more than 30 days old. We additionally found 124 exact duplicate diagnoses across 55 concepts identifying another significant opportunity for automated clean-up. Overall, 12.1% of the diagnoses from our evaluation subset had the potential opportunity for an automated clean-up.

Impact

We continue to work to implement an automated solution to clean-up the problem lists for all Duke Health patients in the Chronicles database in late 2022 to early 2023. This solution will continue to be applied to our problem lists in an ongoing format as a means to keep the PLs more accurate for our patients. We anticipate that this cleanup will have long-lasting effects across direct patient care and secondary processes that rely on problem lists diagnoses as a means to under the severity of illness of our patients.

Next Steps

We continue to work to implement an automated solution to clean-up the problem lists for all Duke Health patients in the Chronicles database in late 2022 to early 2023. This solution will continue to be applied to our problem lists in an ongoing format as a means to keep the PLs more accurate for our patients. We anticipate that this cleanup will have long-lasting effects across direct patient care and secondary processes that rely on problem lists diagnoses as a means to under the severity of illness of our patients.

References

1. Kreuzthaler M, Pfeifer B, Ramos JA, Kramer D, Grogger V, Bredenfeldt S, Pedevilla M, Krisper P, Schulz S. EHR problem list clustering for improved topic-space navigation. BMC Medical Informatics and Decision Making. 2019 Apr 1;19(3):72.
2. Wright A, McCoy AB, Hickman TT, Hilaire DS, Borbolla D, Bowes III WA, Dixon WG, Dorr DA, Krall M, Malholtra S, Bates DW. Problem list completeness in electronic health records: a multi-site study and assessment of success factors. International journal of medical informatics. 2015 Oct 1;84(10):784-90.
3. Wang EC, Wright A. Characterizing outpatient problem list completeness and duplications in the electronic health record. Journal of the American Medical Informatics Association. 2020 Aug;27(8):1190-7

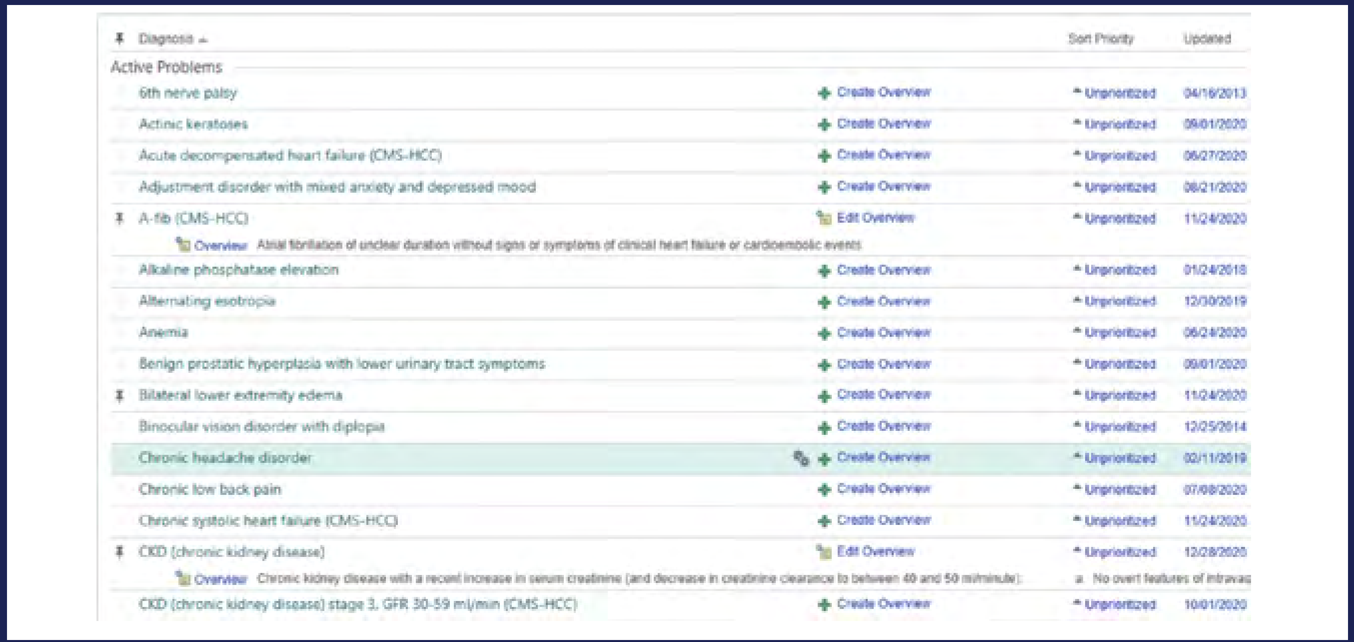

Figure 1: Histogram of count of patients in the ED in a given hour, for the 26,328 total observed hours

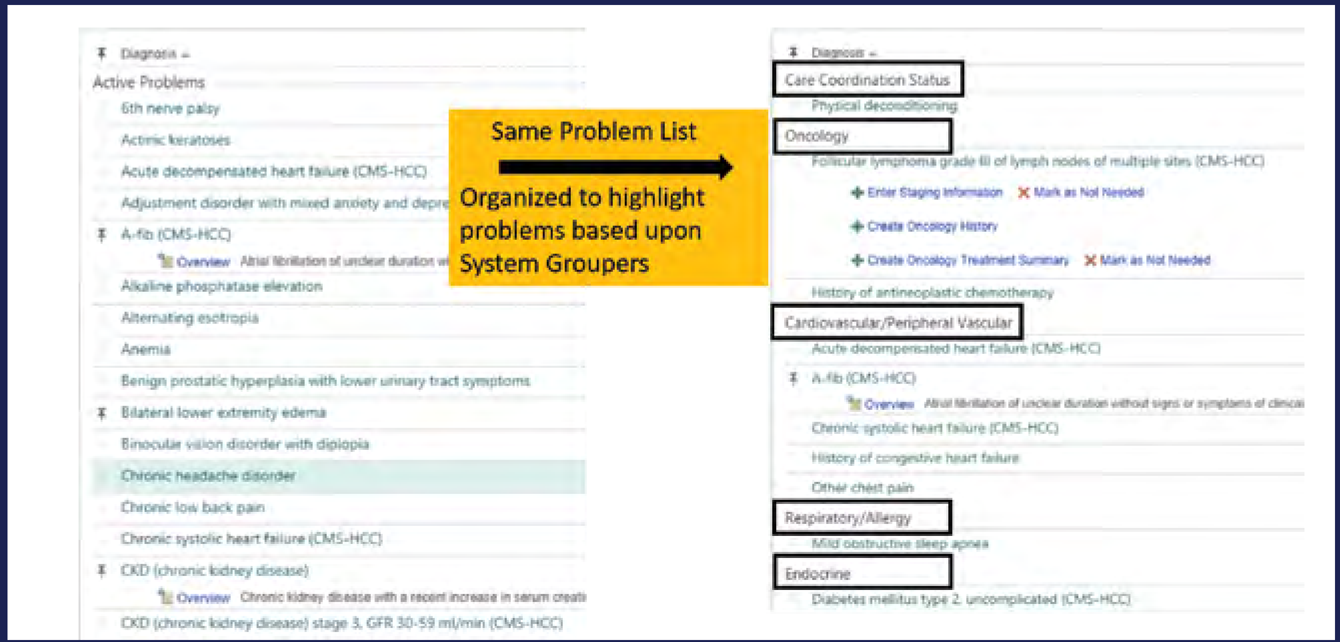

Figure 2: Problem List appearance before (left side) versus after (right side) grouping algorithm was applied. The grouping algorithm applied Boolean logic with SNOMED-CT codes, which are mapped to individual ICD codes, to organize the Problem List content into specialty categories.

DIHI SCHOLAR EXPERIENCE

Hayley Premo

The Duke Institute for Health Innovation (DIHI) provided the opportunity to explore the intersection of technology and healthcare. I had a growing interest in this after spending a year in the Duke Pathology Department amid conversations about the role of automation in making diagnoses. Looking at this future through the lens of surgery and the insights of DIHI was the ideal vantage point for my third year of medical school. I was able to support the project, "A Geriatric-Specific Mortality Risk Stratification Tool for Older Surgical Adults." This project included an amazing team of clinicians and data scientists who were creating a risk assessment tool designed to predict the risk of 30-day mortality in the perioperative period for geriatric patients. The tool incorporated a variety of data elements regarding comorbidities, vitals, lab results, flowsheet values, and medications to estimate patient-specific mortality risks not only on the day of surgery but also at the time of scheduling it. It will hopefully be used clinically to identify high-risk patients in whom surgery may need to be delayed or avoided entirely.

Aside from opportunities to contribute to machine learning (ML) model development, one of the most valuable aspects of DIHI was exposure to current leaders and research in the field of ML and data science. Being exposed to the myriad of ongoing advances in healthcare has been a privilege and has undoubtedly altered how I now view healthcare delivery. From streamlining workflows to optimizing cost to addressing inequities, there are people working on solutions that I barely realized existed prior to starting at DIHI. I am glad to be able to take my experiences from DIHI and apply new skills and knowledge in my career going forward.

DIHI SCHOLAR EXPERIENCE

Tim Ochoa

The Duke Institute for Health Innovation (DIHI) gave me the opportunity to work outside of traditional health care research and focus on innovative machine learning (ML) projects. My primary project during my third year of medical school was "Building a Predictive Model for Post-Operative Complications and Survival after Lung Transplantation". As part of the DIHI team, clinical leaders aimed to develop ML mortality and complications risk models for Duke University Health System patients waitlisted for lung transplants. Duke Transplant Center clinicians and administrators would be able to view these models during their committee review, identify lung transplant candidates at greatest risk for complications or mortality, and react accordingly. This project allowed me to explore the intersection between technology and health care. I developed skills I would likely not otherwise have been exposed to within traditional medical education.

Many of the skills I have been able to develop revolved around learning a programming language, Python, and how it is applied to machine learning model development and implementation. Programming had always been an interest of mine. DIHI's mentorship and teaching allowed me to develop and improve these programming skills within a healthcare setting. Before DIHI's scholarship experience, I lacked computer science education and had difficulty feeling like I could break into the tech world. Python had quite literally been a foreign language to me before this scholarship. Thanks to my time at DIHI, I eventually became comfortable working with complex Python code and feel like I belong among others in the health tech industry! My experiences at DIHI have been so uniquely transformative for me that I will be taking a second research year with DIHI to further explore innovative, tech-focused, health care study. I am looking forward to continuing to learn how to push boundaries within the health system and healthcare delivery.

Building a Predictive Model for Post-operative Complications and Survival after Lung Transplantation

Problem

Lung transplant program performance is graded on one-year survival yet current methods of predicting post-transplant mortality are poor. A lung allocation score (LAS) exists at a national level, using overall data reported by all US lung transplant centers, and is determined by the United Network of Organ Sharing (UNOS). This tool is used to help with organ allocation for lung transplant patients, but transplant surgeons and pulmonologists report that it does not provide a satisfying prediction of post-transplant mortality and does not address the risk of early post-operative complications. Notably, the congregated data and reports provided by UNOS are released infrequently (every six months); therefore, the most recent data could be from six months ago. Furthermore, the LAS calculation uses data between that most recent date and the prior eighteen months, such that the oldest data is two years old.

The Duke Lung Transplant Program is one of the largest lung transplant programs in the world with a diverse patient population with more than 2,250 lung transplants since its inception in 1992. The Duke Lung Transplant Program's higher risk (often declined at other centers) patient population may not be adequately represented with the current LAS model. Currently, the program hosts a weekly committee review in which the extensive evaluation testing is done with inputs from social workers, psychologists, financial coordinators,

01

Team

Timothy N. Ochoa
Will Knechtle, MPH, MBA
Kai Deshpande, MD
Jacob Klapper, MD
Brandi Bottiger, MD
Matthew Hartwig, MD
Laurie Snyder, MD
Linda Tang
Michael Gao, MS

02

Project in Brief

Lung transplant program performance is graded on 1-year survival, yet methods for predicting post-transplant mortality or complications are poor. Duke lung transplant referrals and recipients are unique. Our solution is to predict patient postoperative survival and complications based on patients' pre-lung-transplant EHR data. We will present waitlisted patients' likelihood of poor outcomes in the context of a committee review meeting. The model will lighten the workload of the Lung Transplant Clinic by helping providers prioritize transplants among waitlisted patients and prepare the patient for transplant. In the long-term, we aim to increase transplant 1-year survival.

and pulmonary rehabilitation specialists. Based on this review, the program subjectively deems some patients “high risk.” Early objective determination of high risk would lighten the effort required for the clinical review and, more importantly, help the program prepare patients for the transplant operation. Improving patient preparation and care execution during the transplant operation and hospital admission can dramatically affect the risk of 1-year mortality.

Solution

In an effort to address these shortcomings, we built a more comprehensive predictive model for lung transplant patients. In clinical practice, the goal of this model was to be able to provide real-time risk scores of patient mortality and post-operative outcomes to the clinical team during their pre-transplant committee review of the patient. The model was designed to: (1) predict post-operative “textbook”[2] outcomes in the pre-transplant setting and (2) predict 90-day and 1-year mortality while incorporating data from lung pre-transplant clinic visits and earlier patient history. Textbook Outcomes were defined as freedom from intensive care unit (ICU) or hospital readmission within 30 days of surgery, organ rejection within 45 days of surgery, Grade-3 primary graft dysfunction within 72 hours, tracheostomy within seven days, inpatient extracorporeal membrane oxygenation (ECMO), inpatient dialysis, and reintubation/extubating more than 48 hours after surgery. This model will help inform and guide patient selection criteria and patient clinical care by stratifying high-risk patients.

The model was selected based on technical feasibility as well the potential impact of the model in clinical workflows. Features were built using patients’ inpatient and outpatient data including, but not limited to, demographic data (age, sex, race/ethnicity, comorbidities), laboratory/imaging data, vitals, and medication administrations during the encounters. LASSO and Random Forest models were tested.

The models predicting textbook outcomes after the lung transplant clinic and on the day of admission had AUCs of 0.69 and 0.70, respectively, and their precision was 0.53 and 0.45. A model predicting 90-day mortality at the time of admission had an AUC of 0.78 and an average precision of 0.16 (Figure 3).

Impact

The project outcome desired is to increase the percent of patients having textbook outcomes (no perioperative complications) from 25% to 30%. Primary textbook outcomes to aim for are to reduce post-operative hospital length of stay, time in the ICU, and tracheostomies. Improving textbook outcomes is expected to reduce hospitalization cost by both reducing the time spent in the hospital and procedures needed. A longer-term outcome and ultimate goal is to reduce the rate of 1-year mortality. Mortality and textbook outcome prediction at the time of committee review will accomplish this by guiding resource allocation decisions, preoperative interventions, surgical scheduling, and awareness of potential complications.

Next Steps

First, we will set and discuss thresholds at which the models best differentiate useful information about risk. Second, we will produce and validate model results for patients in real-time: each Monday on the last two weeks of lung transplant clinic visit patients, and upon admission for the transplant. Third, we will begin to incorporate the use of this model into lung transplant committee review meeting agendas. One person will assume responsibility for stewarding presentation of model outcomes, especially if early versions are tested for use in a digital interface outside Epic (like Tableau on the ACE-DIHI server). Committee review stakeholders will be educated about the model with documents and presentations from project clinical leads. They will have the opportunity and weeks to fully review and validate the model according to individual perceptions before they might use it in the context of their other clinical knowledge.

The next version of the prediction models should incorporate patient donor information and UNOS’ compiled data on Duke patients. We expect this data to strengthen prediction as well as increase its applicability to other lung transplant centers.

Academic output

Drs. Snyder and Hartwig recently received a NHLBI U01 grant to expand on this model to other transplant centers and incorporate biological data in the prediction.

Development of a Machine Learning Model for Prediction of Mortality in Lung Transplant Patients. Poster presentation at Machine Learning for Healthcare (MLHC). August 2022, Durham, NC. Ochoa T, Knechtle W, Sendak M.

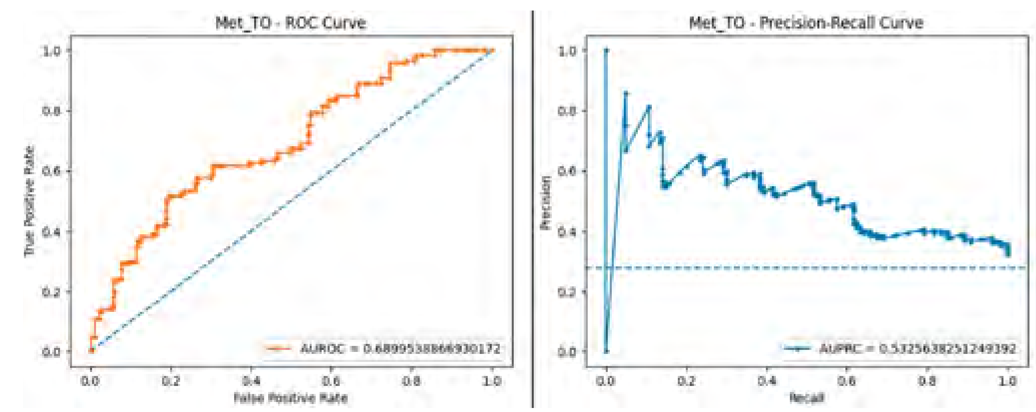

Figure 1.
Predicting textbook outcomes among waitlisted lung transplant clinic patients. (N=1524, AUC: 0.69, PRC: 0.53)

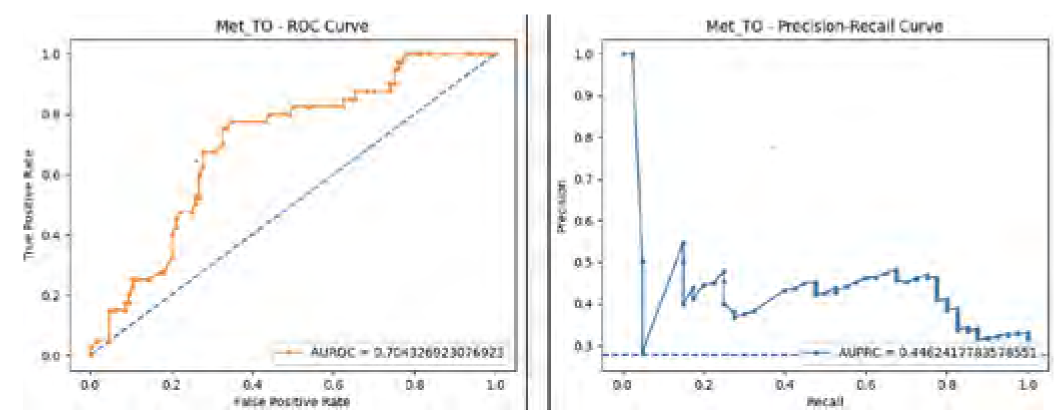

Figure 2.
Predicting textbook outcomes among lung transplant admissions (N=714, AUC: 0.70, PRC: 0.45)

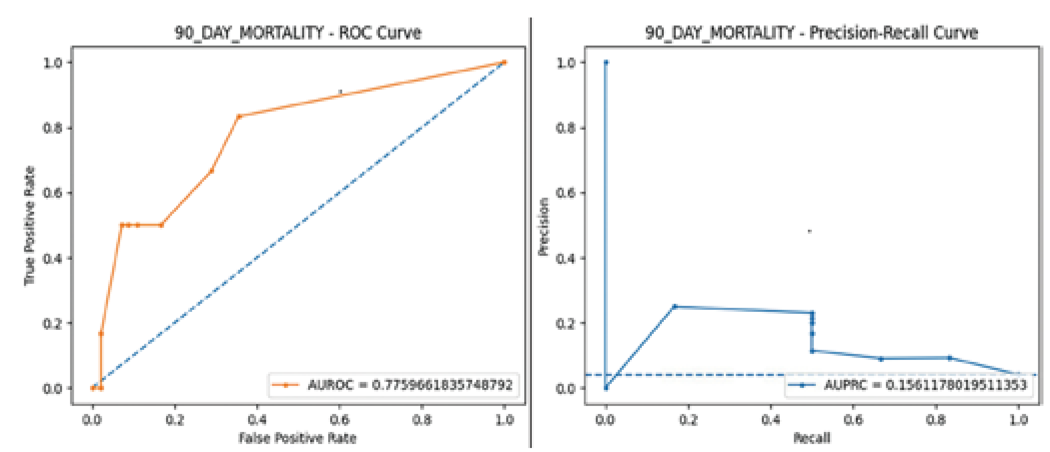

Figure 3.
Lung Transplant 90-day mortality at the time of admission. (N=714. AUC: 0.78, PRC: 0.16)

References

1. Health Resources and Services Administration. (2022, January 6). Scientific Registry of Transplant Recipients: Program Specific Report. Scientific Registry of Transplant Recipients. Retrieved April 10, 2022, from https://www.srtr.org/PDFs/012022_release/pdfPSR/NCDUTX1LU202111PNEW.pdf
2. Halpern SE, Moris D, Gloria JN, Shaw BI, Haney JC, Klapper JA, Barbas AS, Hartwig MG. (2021). Textbook Outcome: Definition and Analysis of a Novel Quality Measure in Lung Transplantation. Ann Surg. DOI: 10.1097/SLA.0000000000004916. Online ahead of print.

DIHI STAFF PERSPECTIVE

Will Knechtle

DIHI: The Case for Catalyzing Innovation to Stabilize and Grow Operations

To stabilize health system operations, leaders must think innovatively early and often. This urgency may sound counterintuitive because prevailing business logic dictates that executives should set an organization's direction and stabilize operations before they innovate. After all, in order to advance there has to be a foundation to build upon, right? To make a case for catalyzing innovation in parallel with stabilizing operations, we must start with agreement on three principles.

First, we must understand that innovation is not only the creation of radically advanced cutting-edge processes for the digital future, but also a method for thinking about and managing fundamental processes of today. Second, we should agree that a foundation already exists: Healthcare is not a new industry. The 'true north' has long been to provide high-quality healthcare safely, promptly, and equitably at the lowest reasonable cost to patients. Third, we must each see a need for foundations and mindsets to change. The foundation of United States (U.S.) healthcare has long needed reform. Regardless of whether a health system (defined here as hospitals and physicians under shared organizational management) is starting or redesigning itself, it sits within a long-existing health industry in need of problem-solving.

As Einstein stated, "We cannot solve our problems with the same level of thinking that created them."[1] Therefore, a health system community needs to think at a different level in order to convert itself from instability to stability. Health systems need to efficiently improve their ability to innovate: to practically implement ideas that result in new goods or services or an improvement in current goods and services.[2]

The Duke Institute for Health Innovation (DIHI, pronounced "Dee-High") provides a model for catalyzing innovative thoughts and services while adhering to lean management principles that lead to operational stability. This model will help health systems renew themselves and stabilize operations. This model will also set a new foundation allowing a health system to pivot from a strategy of renewal to one of growth.

> "Health systems need to efficiently improve their ability to innovate: to practically implement ideas that result in new goods or services or an improvement in current goods and services.[2]

DIHI's use of the word "catalyze" in its mission statement (inside cover page of this report) is intentional. A catalyst increases the rate of change without itself being consumed or permanently changed. When energy is required to move from a state of instability to stability, the catalyst reduces the energy required to activate the change. DIHI is a catalyst first because it is not an enforcer of change but aligns itself with enterprise strategy and transformations already underway in our health system. Second, DIHI's approach is to understand current and preferred user workflow, the core problems (needs) and desired outcomes; then collaboratively develop concepts and rapid prototypes, and finally, iterate through countermeasures. DIHI's thoughtfully engineered innovation management is closely aligned with lean management principles and lean start-up culture.

The first way DIHI practices lean innovation management in alignment with DUHS mission and values. This is exemplified by DIHI's Request for Applications (RFA) process for sourcing and selecting innovations to catalyze for the next year. DIHI leadership meets regularly with Duke Health executive leadership to understand our enterprise's strategic priorities. Every autumn, Duke faculty, staff, students, and trainees are invited to engage in the DIHI RFA by submitting novel project ideas that align with those priorities. Duke Health leaders who are familiar with DIHI and who feel the actual scale and urgency of the problem evaluate the applications on the following seven criteria: alignment with enterprise strategic priorities, potential impact, feasibility of implementation, quality of the team, clarity of the plan and metrics, funding and resource needs, and whether they would personally invest in the innovation project. Every innovation that DIHI spawns starts with top-down strategic alignment and care-unit-up testaments about day-to-day problems and valued solutions faced in work and/or by patients.

Second, DIHI practices lean innovation management by going to the actual place(s) of the problem and solution (i.e., "going to the Gemba"). DIHI's RFA process gets to the spirit of the problem by sourcing innovation from the people and places that directly touch and experience the problem. Yet this is not enough. Once the innovation project is selected, DIHI team members visit the physical locations,* shadow project members, and survey or interview stakeholders who face the problem and will contribute to the solution. A DIHI product manager builds a team of these stakeholders who govern the project and the implementation. DIHI staff continue to meet with the governance team and solicit feedback to iteratively plan, study, check, and act upon analyses and changes made. This innovation design is the best way to see how the innovation touches the patient and whether stakeholders are willing to put the innovation into practice.

Third, when anything is knowingly amiss, DIHI freely communicates an alert and stops the process (pulls the "Andon Cord": a metaphorical rope). DIHI has set up technology and work culture to make this possible. DIHI has created a data pipeline as well as software tools and dashboards to monitor data and meta-data drifts. DIHI closely monitors the pipeline and projects for drifts or errors and alerts the right personnel to ensure safety and quality of the projects consume the information. For example, a predictive mortality model input dashboard exists behind the scenes to track hourly trends of data inputs to the model so that, if any unusual variance occurs, the flow can be identified and fixed.

Furthermore, DIHI has also developed a thorough data quality assurance framework for machine learning in healthcare: A total of 2,362 quality checks and twenty three quality reports were generated across five projects.[3] This process prevents a model from moving towards the next stage of development or use until the quality-related problems are resolved. Not only are quality checks rigorous, but DIHI's culture allows a process to completely stop for root cause problem-solving and quick countermeasures. This full stop causes delays which, admittedly, distress

* COVID-19 prevented going to the Gemba as much as desirable, which negatively impacted innovation implementation.

clinical and operational project leaders. Nevertheless, DIHI managers prioritize the quality of patient care, the mission of executive leadership, and the needs of data scientists over project turnover, publications, or profit. Sounding the alarm is not as fun as visiting the sites of work or putting a solution into practice, but DIHI does it unreservedly.

Finally, DIHI operates like and maintains the work culture of a lean start-up.[4] The innovation RFA process and project are designed to align problems, solutions, and impact. Collaboration on the front lines forces intensive project scoping, exploratory data analysis, and rapid feedback cycles. DIHI often uses Tableau, for example, because it allows rapid creation of and feedback on minimum viable products–an app prototype can be made and tested within weeks. Some might fault DIHI for not running many large clinical trials. Nevertheless, DIHI is being lean when it focuses first on proving value is present and accepted on the frontlines. In the case of Sepsis Watch, DIHI delivered value without starting a clinical trial, but now that it intends to scale that value, it has initiated a clinical trial.[5]

> Not only are quality checks rigorous, but DIHI's culture allows a process to completely stop for root cause problem solving and quick countermeasures.

When monitoring and controlling scaled projects requires six-sigma blackbelt rigor, we maximize partnerships with groups like Performance Services Care Redesign and NSQIP nurses. Even if a project does not scale, we value the lessons learned from the project as much as its clinical or vocational impact. Furthermore, project managers recognize that management tools and templates such as RAILS, Trello, Agile, and Sprints can create a false sense of task completion and be a distraction from getting real work done: we self-evaluate weekly and adapt tools to directly add value through innovation lessons or project execution. Most importantly, a culture of meaningful work is created by three factors: (1) Alignment with patients and the front lines, leaders, and executives; (2) High-valuation of learning, rapid feedback, and quick impact; and (3) Leadership support to stop the line if the first factors deteriorate.

Innovation in healthcare is inherently focused on creating value by advancing systems and technology. Operational innovation begins with ensuring quality standards and technological engineering have been faithfully developed and executed. A lean start-up attitude starts and ends with incorporating user feedback, aligning with leaders, and ensuring a positive experience of safe and optimal utility. For these reasons, innovation catalysts and lean management engineers are symbiotic interdependent partners advancing operational stability and growth.

References

1. Albanese, C,T, Aaby D, Platchek T. (2014). Advanced Lean In Healthcare. Create Space Independent Publishing Platform, North Charleston, SC.
2. Schumpeter, Joseph A, 1883–1950 (1983). The theory of economic development : an inquiry into profits, capital, credit, interest, and the business cycle. Opie, Redvers,, Elliott, John E. New Brunswick, New Jersey. ISBN 0-87855-698-2. OCLC 8493721. (https://en.wikipedia.org/wiki/Innovation)
3. Sendak, M, Sirdeshmukh, G, Ochoa, T, Premo, H, Tang, L, Niederhoffer, K, et al. Development and Validation of ML-DQA -- a Machine Learning Data Quality Assurance Framework for Healthcare. Proceedings of Machine Learning Research. August 2022.
4. Ries, E. (2011). The lean startup: How today's entrepreneurs use continuous innovation to create radically successful businesses. Crown Business.
5. Implementation and Evaluations of Sepsis Watch. https://clinicaltrials.gov/ct2/show/NCT03655626

* COVID-19 prevented going to the Gemba as much as desirable, which negatively impacte innovation implementation.

DIHI STAFF PERSPECTIVE

Linda Tang

Developing a Machine Learning Data Quality Assurance (DQA) Framework

Real-world data is messy. This is especially true of electronic health records (EHR) data. A single data element such as creatinine could be represented in multiple ways and located in multiple tables in the dataset. EHR data also is likely to be incomplete and inconsistent across measures and time. It is prone to errors in the data generating process. Regardless of error, EHR data is dynamic because best practices change over time. For instance, a data shift could result from a lab value being measured in a new way or calculated using different formulas. Maintaining a well-curated EHR dataset is challenging because it requires a high level of domain-specific knowledge and is too large for manual inspections of all data points. Poor data quality has a cascading effect and hinders the validity of machine learning models constructed using these datasets.

To address this gap, our team started developing a framework to systematically assess and improve the quality of EHR data for machine learning in Fall 2021. The workflow of the data quality assurance (DQA) framework is comprised of three phases (preprocessing, data quality checks and adjudications) as shown in Figure 1.

During the processing phase, different representations of the same data element are combined into "groupers". In addition, we apply rule-based transformations. These transformations include standardizing measurements to the same unit (e.g. convert mg/mL to mg/dL), remove outliers (e.g. pulse oximetry values above 100%), and parsing text into numeric values (e.g. converting 120/80 mmHg to systolic and diastolic blood pressure).

Next, the data quality is assessed from three perspectives: completeness, conformance and plausibility, a framework adapted from the Patient Centered Outcomes Research Institute (PCORI). Completeness checks inspect patterns of missingness in the data. Conformance checks assess whether the variable type, value, and range match prespecified expectations. For instance, a conformance check could assess whether all Glasgow Coma Scale (GCS) values are between 3 and 15. Lastly, we assess whether a data element is plausible (i.e. credible within a clinical

	PEDIATRIC SEPSIS	LUNG TRANSPLANT	JEFFERSON HEALTH SEPSIS	irAE	MEWS
TOTAL NUMBER OF DATA MEMBERS	181	144	89	80	322
TOTAL NUMBER OF DQA CHECKS	389	432	267	1,034	877
TOTAL NUMBER OF GROUPERS TRANSFORMED	11	9	13	18	10
TOTAL NUMBER OF GROUPERS EXCLUDED	9	22	9	12	4

Table 1. The impact of implementing DQA of 5 projects

context). One sign of plausibility is stability over time; hence, we plot the distribution of data elements across months and years to assess trends over time. While completeness and conformance can be designed and interpreted almost entirely with data engineering and analytic expertise, plausibility checks require much more clinical expertise. Plausibility of values is dependent on the patient population, care environment, and time of care.

In the third and final phase, adjudication, the clinical lead of the project reviews the data quality reports and documents whether a data element is sufficiently complete, conformant and plausible. If a data element does not meet these requirements, the clinical lead communicates the transformations necessary to improve data quality. They may suggest eliminating the data element from modeling all together. With the feedback from clinicians, data scientists apply these transformations and refine the cleaning code.

In Spring 2022, the newly developed DQA framework was used to support four internal Duke Institute for Health Innovation (DIHI) machine learning projects and one project at Jefferson Health. On average, applying the DQA framework resulted in transforming 12.2 groupers and removing 11.2 groupers from modeling per project (Table 1). For example, the DQA process helped refine the bounds used to clean the data (e.g. respiratory rates greater than 100) and helped us catch sudden changes in the data elements (e.g. the mean of eGFR increased from one month to another. Consequently, this grouper was removed from modeling). This outcome suggests that applying the DQA process results in a significant improvement of data quality, safeguards the validity of the models we develop, and ameliorates the results from subsequent analysis. Moreover, this work draws attention to an often overlooked and taken-for-granted process in building machine learning models: multidisciplinary collaboration and iteration. Implementing the DQA framework involved input from clinicians, data scientists, data engineers and project managers. The DQA process is iterative and often required many rounds of discussions to identify the best approach to process a data element. This highlights the amount of effort and the level of collaboration necessary to assure data quality.

There are two directions by which our team hopes to improve the DQA framework in the future. First, we hope to create a standardized DQA programming template that is adaptable to different projects and easy to users to navigate. Second, we hope to involve more key stakeholders, such as nurses, in the DQA process to better understand data generation.

Reference

1. Sendak, M, Sirdeshmukh, G, Ochoa, T, Premo, H, Tang, L, Niederhoffer, K, et al. Development and Validation of ML-DQA -a Machine Learning Data Quality Assurance Framework for Healthcare. Proceedings of Machine Learning Research. August 2022.

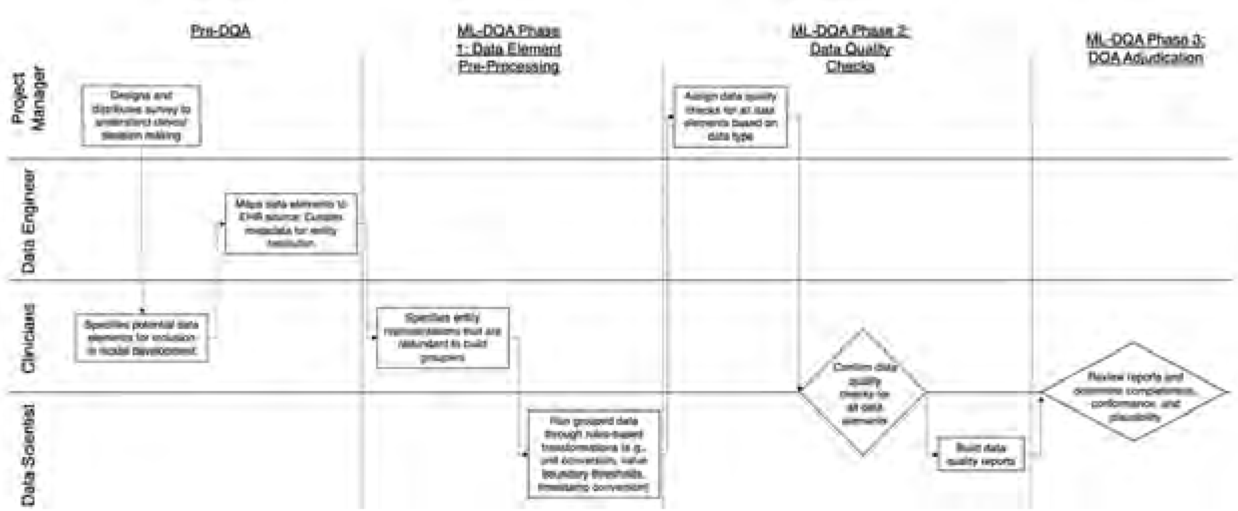

Figure 1. The DQA workflow diagram

Algorithm Development for Duke Emergency Pre-hospital Capacity Management

Problem

Emergency Department (ED) overcrowding is a growing national problem affecting the quality of and access to healthcare. Overcrowding can cause the ED to divert incoming emergency medical services (EMS), a decision which is often reactive and subjective. The decision is often made without pre-determined guidelines or sufficient coordination with nearby EDs. Diversion status compromises patient safety and quality of care (e.g.,long wait times, high left-without-being-seen rates) and diminishes provider well-being and performance. It also increases downstream transfers to other hospital EDs, further detracting from efficiency and patient care.

Solution

We proposed to provide our health system and EMS provider with pre-hospital traffic control using real-time ED census data. We believe this will be an effective approach to ensuring incoming patients are transported to the most appropriate destination to receive timely care. Diversion will be reduced or, at minimum, their effects will be mitigated via the optimization of Duke Health ED capacity.

In 2022, we began implementing an algorithmic tool for capacity management with the aim of improving patient care and coordination between the Duke University Hospital (DUH) and Duke Regional Hospital (DRH) EDs. The algorithm was based on hourly census data from the DUH and DRH EDs for a three-year period between 7/1/2018 and 6/30/2021 (26,328 hours in total). We gathered and analyzed data using conventional methods for measuring ED crowding:[1] number of ED patients,

01

Team

Rebecca Shen
Will Ratliff, MBA
Mark Sendak, MD, MPP
Neel Kapadia, MD
Anjni Joiner, DO, MPH
Will Knechtle, MBA, MPH
Andrew Godfrey, MD
Joshua Boyd, MD
Zachary Cockerham, RN
Suresh Balu, MS, MBA
Jason Theiling, MD
Brian Burrows, MD

02

Project in Brief

The Duke Emergency Departments do not currently have an objective approach to support diversion. This leads to Emergency Department (ED) overcrowding, which results in poor process and clinical outcomes. DIHI and the Duke EDs have created an ED capacity monitoring solution, and are coordinating with Durham County Emergency Medical Services (EMS) to implement it into their ED routing approach. In the autumn of 2022, we began implementing and assessing the solution for measurable, high-impact improvement for Duke EDs: time to intervention, ED staff resilience, ED wait times, and rates of patients left-without-being-seen.

waiting room status, ventilator use, patient acuity levels (ESI), number of psychiatry boarders, and EMS transports. Figure 1 shows the distribution of patients in the ED by hour over our three-year period. We created new measures such as cumulative time elapsed: the elapsed time spent in the ED summed over all patients. Percent raw capacity was defined as the number of patients out of the max count ever observed for that ED over the three-year historical period. We defined 'imminent ED workload,' our primary outcome label, to be larger when the number of total patients was high and the median of those patients' time since arrival was low compared to historical trends (i.e., more work needed to triage newly arrived patients) in Figure 2.

"We found that, retrospectively, the maximum ratio of high and critical workload labels occurred during the eighteen hours prior to the start of historical diversion events.

Data granularity was set to the hourly level (26,328 hours summarized per ED over the three-year period), allowing the inclusion of metrics like "arrived or triaged in the past hour" and "time of day comparisons".

We validated our imminent ED workload outcome label using historical diversion data as well as components of the National Emergency Department Overcrowding Study (NEDOCS) scoring system.[1]

Our algorithm modifies raw capacity based on the number of severe acuity cases (ESI level 1, 3% capacity increase per case), the number of ventilators in use (10% capacity increase per ventilator), and the percentage of triaged patients (10% capacity increase if less than half of patients have been triaged). We set high workload to be when modified capacity was greater than 55% and medium workload to be when modified capacity was between 33-55%. Critical workload occurred when modified capacity exceeded 55% and the cumulative time elapsed was less than the observed median for that hour of the day between 7/1/2018 and 6/30/2021. Running our algorithm retrospectively over the three-year timeframe, DUH had 1,451 critical hours, 6,222 high workload hours, and 12,418 medium workload hours. We found that, retrospectively, the maximum ratio of high and critical workload labels occurred during the eighteen hours prior to the start of historical diversion events. Future EMS validation of Durham County EMS decisions in the field will be based on real-time ED data inputs.

We developed an algorithm that identifies high and critical ED workload status, which will be incorporated into an AI tool that helps guide EMS transport decisions and hopefully reduces ambulance diversions. The relative ease of implementation was key for timely integration into existing systems and allowed our algorithm to be tailored to the local environment Duke's Durham EDs operate in. Our algorithm was also highly adaptable to ED-specific protocol changes and patient volume/acuity, both due to and independent of the COVID-19 pandemic.

Next Steps

Improved management of EMS transport destinations is a potentially high-value approach to avert downstream overcrowding crises and elevate patient care and safety. We are working with DUH, DRH, and Durham EMS to pilot the solution in early 2023. Then, post-pilot period, we will assess impact of the solution on time to intervention, ED wait times, rates of patients left without seen, and ED staff resilience. Future work will include incorporation of capacity status into DCEMS protocols and inclusion of patient conditions and patient preference into ED destination decision. We also plan to extend this work to support other health system EDs and their patients.

Academic output

Algorithm Development for Duke Emergency Pre-hospital Capacity Management. Machine Learning for Healthcare 2022 – Clinical Abstract, Software, and Demo Track. August 5-6, 2022.

Rebecca Shen; Will Ratliff, MBA; Mark Sendak, MD, MPP; Neel Kapadia, MD; Anjni Joiner, DO, MPH; Will Knechtle, MBA, MPH; Andrew Godfrey, MD; Joshua Boyd, MD; Zachary Cockerham, RN; Suresh Balu, MS, MBA; Jason Theiling, MD; Brian Burrows, MD.

Reference

1. Weiss SJ, et al. Estimating the degree of emergency department overcrowding in academic medical centers: results of the National ED Overcrowding Study (NEDOCS). Acad Emerg Med. 2004 Jan;11(1):38-50.

DUH CAPACITY, THREE-YEAR RETROSPECTIVE

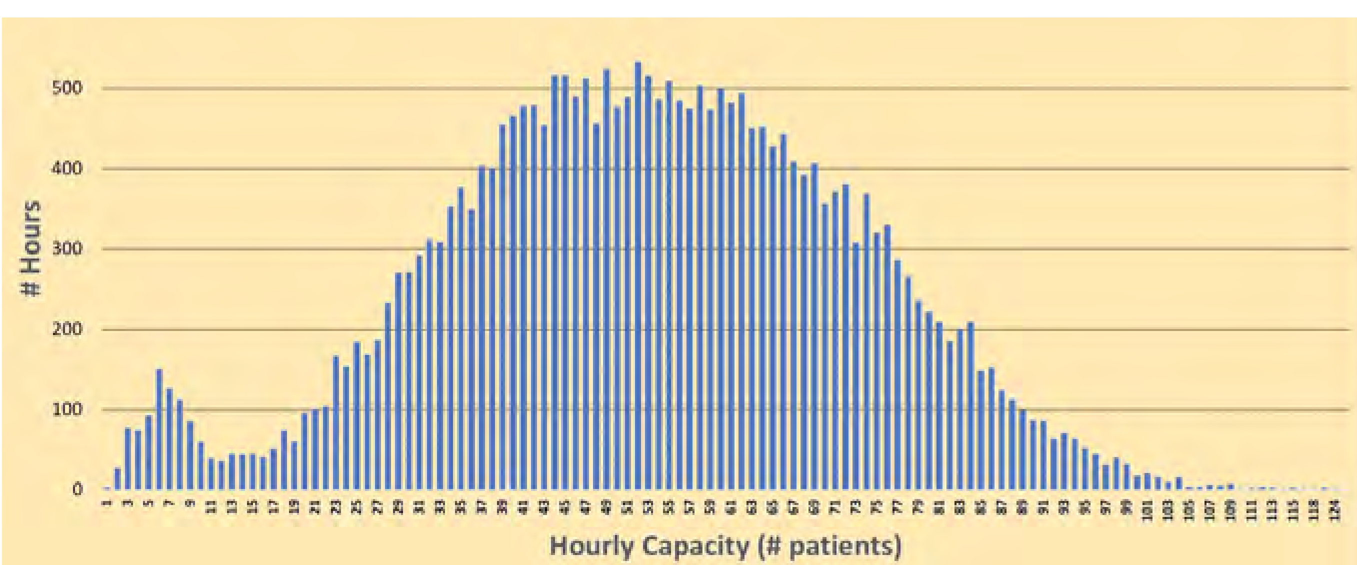

Figure 1: Histogram of count of patients in the ED in a given hour, for the 26,328 total observed hours

MEDIAN ELAPSED TIME IN DUH ED PER PT

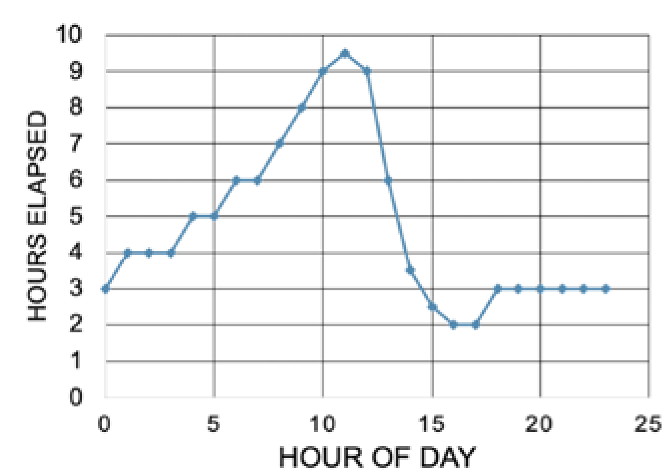

Figure 2: Median elapsed time spent in ED per patient, from ED arrival time until a given hour, for the 26,328 total observed hours

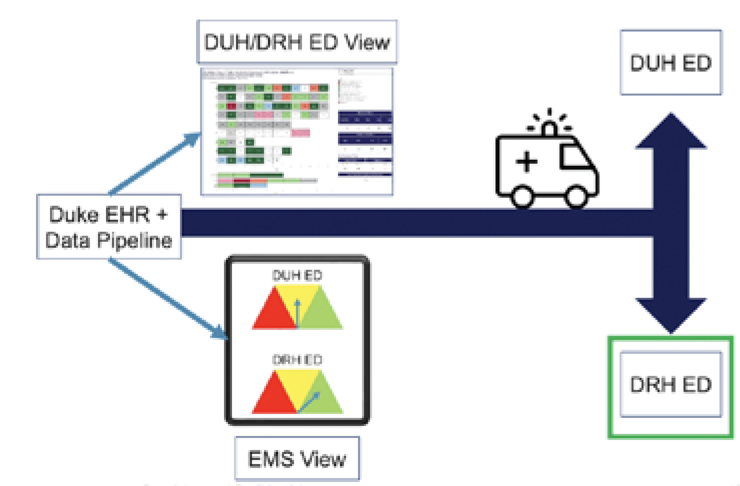

Figure 3: Workflow visualization of real-time Duke ED + Durham EMS Support Solution

Will Knechtle

Hospital at Home

As one of DUHS's strategic innovations in 2020, Duke Institute for Health Innovation (DIHI) led the design, development and piloting of a hospital-at-home care delivery model at Duke Raleigh Hospital with Drs. Vidhya Aroumougame and Mike Spiritos as principal innovators. This culminated in the safe and successful treatment of a Duke Raleigh patient in their home in April 2021. The torch of Hospital at Home's diffusion is being carried on in late 2022 by Duke leaders representing DukeHomeCare and Hospice, Internal Medicine, Performance Services, Finance, Case Management, Duke Health Technology Services (DHTS), and the Duke Institute for Health Innovation (DIHI).

In DIHI Impact Issue 22, we shared details and lessons learned from Duke Raleigh Hospital's 2020-2021 Hospital at Home (H@H) experience. We also released a 77-page H@H "playbook" providing a script for enrolling and caring for patients, workflow diagrams, organization charts, role-based checklists, budgeting information, and other details about execution logistics. In recognition of the success of pioneering an Acute Hospital Care at Home program in North Carolina, Dr. Vidhya Aroumougame, Dr. Mike Spiritos, and Will Knechtle were invited to accept the North Carolina Healthcare Association's Highsmith Award for Innovation on July 22, 2022.

For the diffusion and scaling of H@H, DIHI will continue to provide expertise in development and integration of algorithms, heuristics, decision support tools, and data science to identify candidate patients for H@H, support care for enrolled patients, and support market planning efforts.

In particular, DIHI has developed an infrastructure and tools to quickly search and analyze seven million records in a matter of seconds or identify patient changes in near real-time (e.g., the latest five minutes from this second). An example of this capability developed can be experienced by using DIHI's real-time Bed Watch analytics tool. This tool was put into use by DIHI when implementing machine learning models predicting admission risk for patients in Duke Emergency Departments as well as when monitoring COVID-19 across Duke hospitals (DIHI Impact Issue 21). Today, every hospital bed is monitored, visualized and color-coded by patient data (demographics, admission status, age bracket, length of stay, infection status, sepsis risk, etc.)

Starting in November 2022, DIHI Bed Watch is being used to help identify currently hospitalized patients who may meet Hospital at Home criteria, refine the criteria, and communicate patient availability and location quickly to Hospital at Home physicians and operational leaders. An example of this use is visible in Figure 1. The Duke Hospital at Home team plans to use and further develop this tool to expedite Hospital at Home patient enrollment and management.

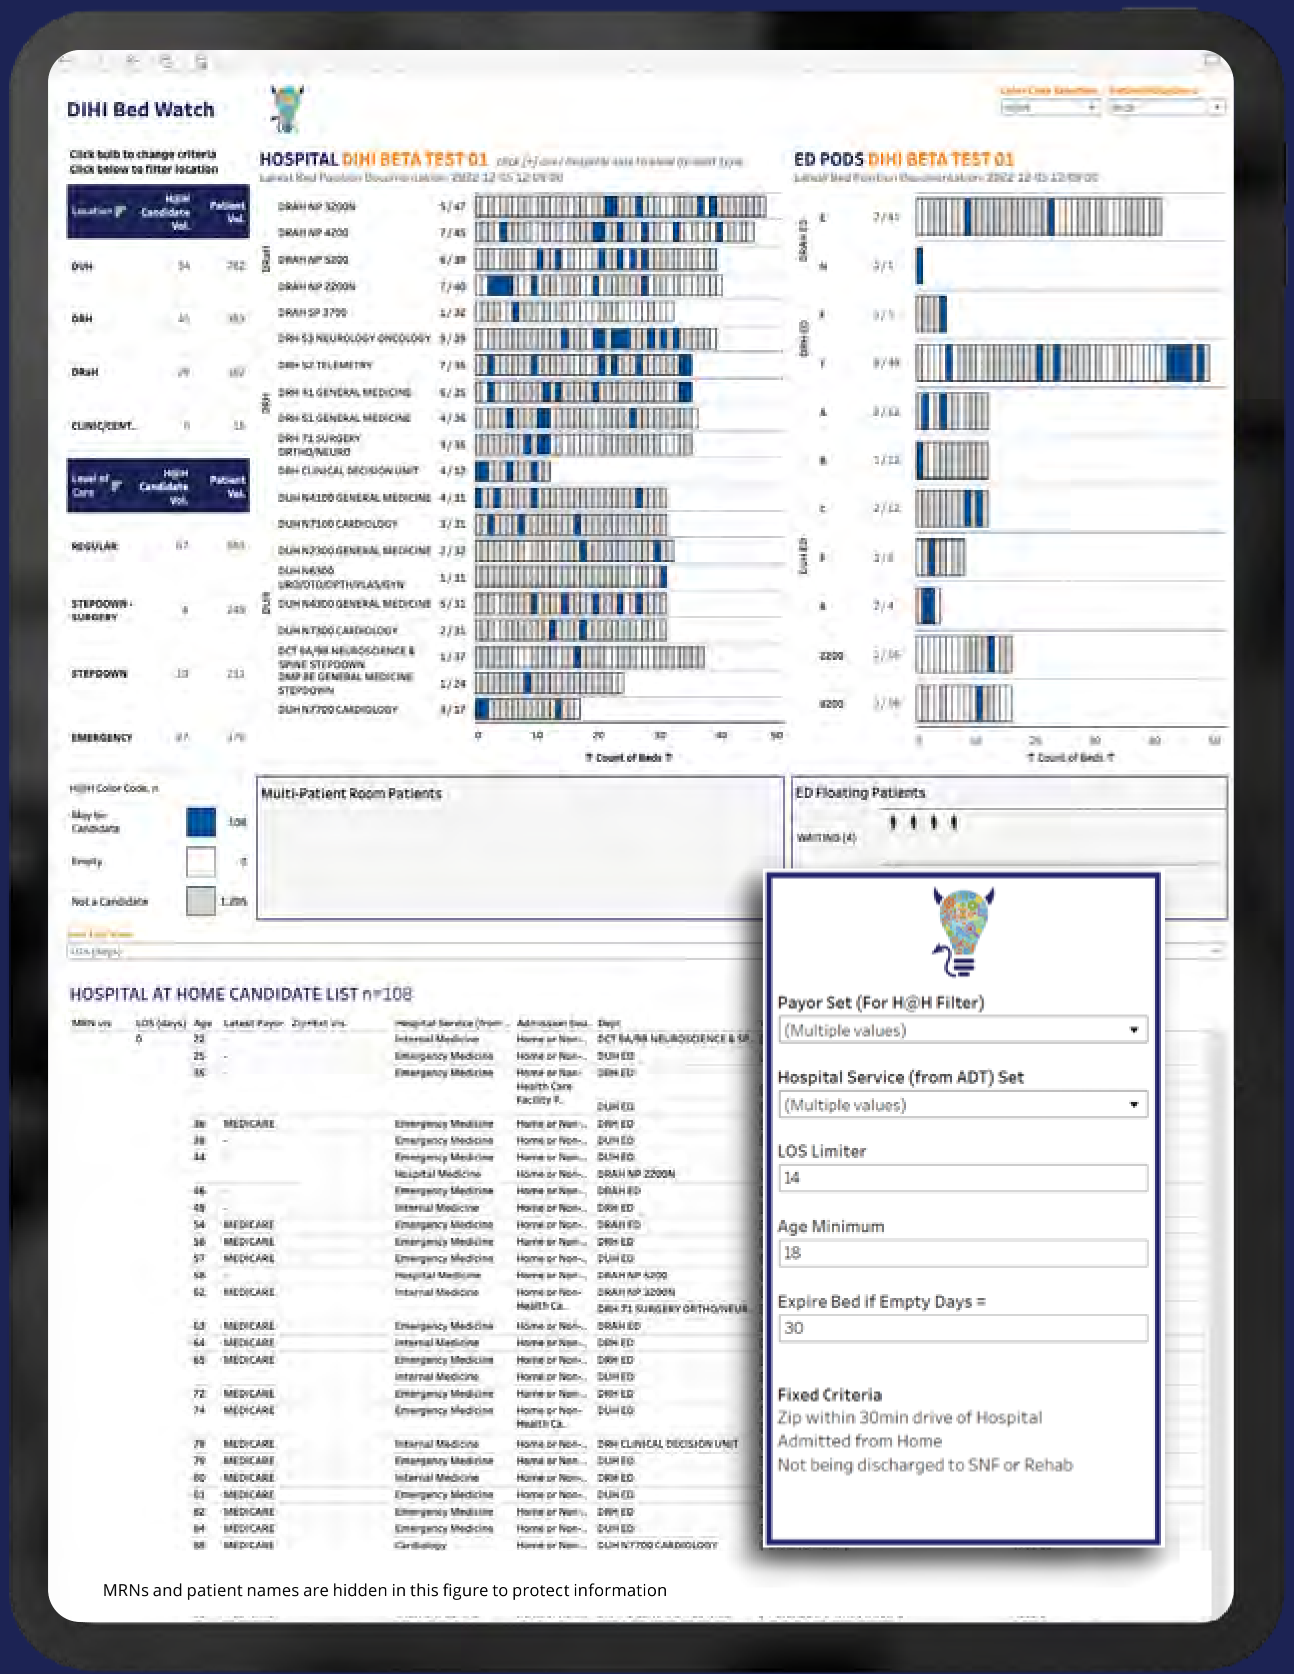

MRNs and patient names are hidden in this figure to protect information

Will Ratliff

Deterioration Program (Sepsis Watch, Cardiac Decompensation, Adult Deterioration)

Duke Health advances science and extends patient lives via advanced interventions, including first-ever surgeries, such as the partial heart transplant in September 2022. Harmoniously, at the Duke Institute for Health Innovation (DIHI) we aim to keep pace with ongoing improvements to identify patients experiencing imminent deterioration in the hospital. Our goal is to support Duke front-line clinicians in identifying deterioration earlier and providing actionable context to inform subsequent decision-making. Through our Sepsis Watch 2.0, Adult Deterioration, and Cardiac Decompensation projects, we focus on advancing methods for early identification of deterioration, as well as methods to scale these learnings throughout applicable settings in our health system. A primary example of this is our Sepsis Watch 2.0 product, which extends the prediction of sepsis to all hospitalized patients, not just those in the Emergency Department (ED), and predicts subsequent sepsis events after a cool-down period once sepsis is met. We are currently validating this model and optimizing the Sepsis Watch interface to handle the increased volume of patients. Meanwhile, we are diffusing the current Sepsis Watch solution to areas of the health system where it could be useful.

In March 2022, we went live with a Sepsis Watch notification system for all sepsis patients hospitalized at Duke Raleigh. In this workflow, the Duke Raleigh Critical Care Rounding (DRAH CCR) Nurse receives a page within ten minutes of any hospitalized patient meeting the Sepsis Watch definition (excluding surgical departments), and promptly rounds on that patient to discuss sepsis confirmation and/or bundle treatment with the bedside care team. Once a patient meets sepsis, notifications for that patient are snoozed for thirty hours before the patient is re-assessed for sepsis. This equates to 2.2-page notifications per day and is meant to extend our Sepsis Watch 1.0 program with Duke Raleigh beyond just patients in the ED. As we implement our Sepsis Watch 2.0 model this fall, we are adding logic to notify the DRAH CCR Nurse for patients at high risk of becoming septic, with a distinct six-hour snooze between notifications. At Duke University Hospital (DUH), we will similarly use the push notification methodology to alert the DUH Patient Response Program (PRP) nurses and physicians of any hospitalized patient who is septic or at high risk of becoming septic. The functionality will complement the Sepsis Watch 2.0 Web Application and minimize the burden of checking the Web App.

PATIENT GROUP	IP MORTALITY	ICU TRANSFER
Adult DUH patients 10/15-08/18	3.0%	4.1%
Hypotension	11.2%	8.6%
End organ dysfunction	6.4%	5.9%
Hypoperfusion	21.0%	12.3%
Vasoactive meds	19.2%	12.7%
Resp decline	9.3%	7.5%
Resp intervention	15.5%	12.0%
Met no phenotypes	0.6%	2.2%

Table 1

Deterioration outside of the realm of sepsis is a concurrent priority at DIHI. Our adult deterioration model, which predicts an intermediate/step-down patient's risk of needing an Intensive Care Unit (ICU) transfer within the next twelve hours, identifies all-cause deterioration early. We are integrating this model into existing workflows with the PRP team to support their rounding and interventions as needed. We expect to prevent the need to escalate care via ICU transfer. While this model is trained to give good direction on which patients need attention, it does not provide additional context as to why the patient is deteriorating. For this, we will leverage our Cardiac Decompensation phenotypes from a prior RFA project. In June 2021, we implemented a Tableau Dashboard displaying DUH Cardiology Unit patients who met one or more of the Cardiac Decompensation phenotypes (hypotension, end-organ dysfunction, vasopressor administration, respiratory decline, and respiratory intervention) within the past twelve hours. Through retrospective analysis and prospective adjudication, we've observed that these phenotypes are correlated with higher rates of ICU transfer and mortality and also that they are relevant for clinical deterioration (i.e., beyond just cardiovascular-driven deterioration) (Table 1). Moreover, they synthesize clinically meaningful data points related to the context of a patient's deterioration.

Our vision is to infuse the Sepsis Watch product, the adult deterioration prediction model, and these Cardiac Decompensation phenotypes to yield a comprehensive deterioration program that is more than the sum of its parts. If we can help identify patients earlier in their deterioration and reduce the real-time cognitive load on Duke clinicians to determine the cause and appropriate interventions, we hope our deterioration program will save lives and reduce the burden on our front-line colleagues.

DIFFUSION AND SCALING

Will Knechtle

Maternal Early Warning System (MEWS)

Lessons continue to be learned from the project "Development of a Maternal Early Warning System (MEWS) Using Machine Learning (ML)" summarized in DIHI's 22nd Impact Issue. The original goal of MEWS was to first identify obstetric patients at the highest risk for severe maternal morbidity (SMM) and mortality and then to effectively integrate alerts into workflows to improve patient outcomes. The first iteration of machine learning models was retrospectively evaluated in 2021 and the next step was to view the MEWS outputs as a project team and verify that the early warnings met clinical practice and expectations.

DIHI executed a collaborative MEWS output review with Maternal Fetal Medicine physicians and specialized anesthesiologists. Two more user-interface purposes emerged that incoming fellows and physicians believed would strongly increase buy-in. According to these use cases and care practitioner review, two methods for identifying severe maternal morbidity diverged, and two of the eight outcomes required updated definitions. While the original use case was to encourage extra checks on patients an ML model reliably encouraged, we learned it would be as useful to highlight patients who had already met official outcomes definitions (CDC diagnoses) so that unit leaders could remain aware of those patients' conditions, locations, and progression. Furthermore, it would be helpful to alert the team to patients experiencing extreme vital signs, regardless of risk calculations or diagnosis. Finally, practitioners believed we could improve the match between the real-time clinical signal of a phenotype and the actual diagnosis.

To predict whether or not an event will occur in the next four hours, the model needed to be trained with data describing the precise minute the outcome occurred. This means that diagnosis codes, the official method for tracking SMM over time, were insufficient – they could inform about whether a condition was met during a visit but not the hour of the visit. Consequently, the 2020-2021 MEWS project team used literature review and clinical experience to define real-time clinical phenotypes for each SMM outcome modeled. The MEWS ML is therefore predicting a timely signal that a patient is starting to meet criteria often used to diagnose a condition, but is far from providing a diagnosis itself. For example, the eclampsia clinical phenotype occurrence was ~5% whereas diagnosis was ~0.3%. This difference is helpful for timely prediction and for a high fraction of relevant instances of the outcome to be achieved (recall, sensitivity). However, it also lowers precision.

During a simulation using retrospective data, too many alerts were sent out for eclampsia that providers found were irrelevant to whether a patient would be diagnosed with it during their hospital stay. Accordingly, to prepare for MEWS diffusion, our team is focusing on improved real-time definitions of both eclampsia and hemorrhage and including (separate from the ML) monitoring of morbidity and extreme vitals within the user interface (Figure 1).

OVERVIEW INTERFACE

PATIENT DETAIL INTERFACE

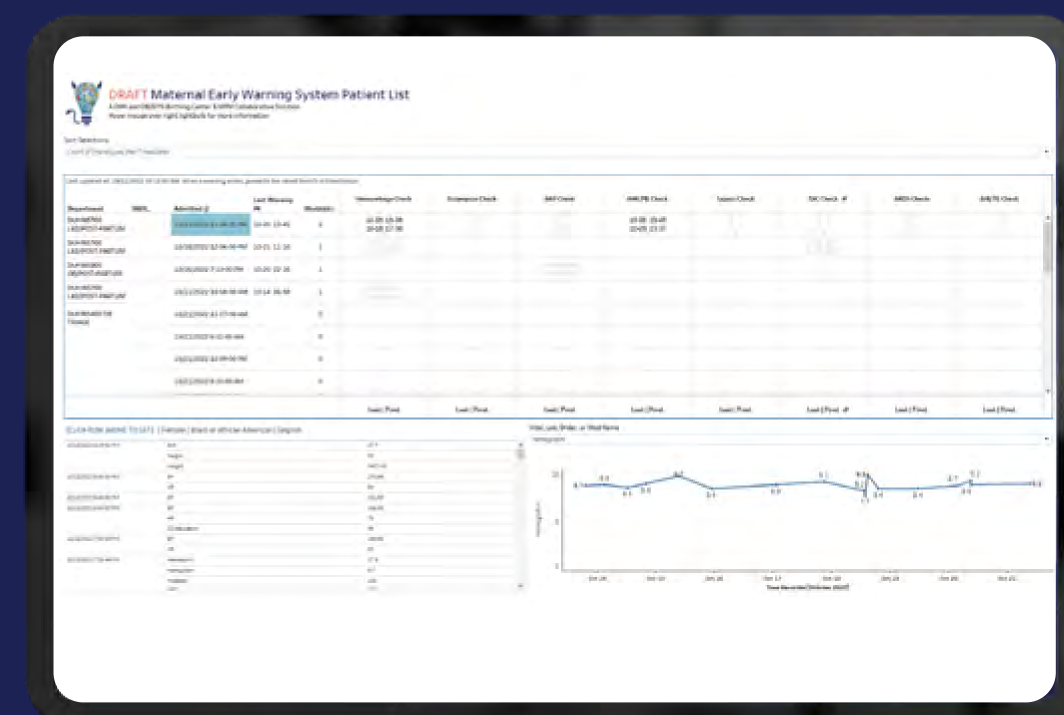

USER INTERFACE:
Tableau dashboard, refreshed with updated data every fifteen minutes and model updating every hour. Data displayed: most recent time the patient met the phenotype is in orange. If phenotype has not been met, then the risk of meeting the outcome within four hours is displayed in yellow. Additional info shows underlying phenotype data.

DIFFUSION AND SCALING

Michael Gao

Health Guard for Advanced Care Planning

DIHI's work in identifying patients at high-risk for mortality using machine learning has continued to evolve as the health system continues to focus on increasing rates of advance care planning documentation. Specifically, as the project sees increased adoption, we have made improvements in the delivery as well as the measurement of the model's impact. Two of the limitations of our original model, which flagged high risk patients of mortality for physicians to review, were a manual process for reviewing whether patients were correctly flagged and a lack of analysis on the true impact the model was having. We have addressed these two main limitations in our ongoing work.

First, our Analytic and Logic Driven Intimation System (ALDIS, a notification system) has relieved much of the burden of reviewing patients and passing information along to the right end users. The ALDIS notification system interacts with all of our machine learning models and forwards notifications to the appropriate user at the appropriate time, using the appropriate method. This system has allowed us to send critical information to the appropriate clinical team member throughs secure text messages, Microsoft Teams messages, pages, and other methods. The mortality model has been integrated into this system, and the model results are now sent to the appropriate team member via automated email. This allows us not only to remove the burdensome review process, but also to track metadata about how many notifications were sent and to whom.

Not only has ALDIS addressed the delivery of the model, but facilitated the measurement of model results. With more granular tracking and delivery functionality, we were able to kick off a clinical trial, where members of the DUH general medicine teams were randomized into either receiving the notification or not, to estimate the true effect that our notification system has on subsequent advanced care planning documentation. ALDIS handles the randomization and tracking of control and treatment groups and ensures that the integrity of the trial is kept intact.

Though we are awaiting full results from the trial, it is difficult to ignore the impact that the notification system, when coupled with other hospital programs, has had on the rates of advanced care planning. In a recent report, it was stated that in the months prior to the development and integration of the DIHI mortality model, only 3% of patients had a goals-of-care conversation in the six months leading up to their death. That number now stands at 55% in the most recent data. The impact that this model has had, both in identifying patients and catalyzing health system initiatives, is something that epitomizes the mission of DIHI as a whole; and we hope to continue to support this model and others in realizing that mission.

DIHI STAFF PERSPECTIVE

Mark Sendak

Building Bridges and Islands to Scale DIHI's Impact

"

If you can prove what is possible, the world will change… That is one of the cardinal mistakes of Silicon Valley. That's not how businesses or institutions or the world changes…

Islands are incredibly important. It is true that we need exemplars to understand what's possible. Islands create demand for change, but they are not the supply for change…

A bridge is something that starts where your customers are today, but it goes to the island. Bridges create the supply for change.

— Adam Pisoni from Abl Schools, November 7, 2018

Over the last ten years, our team at the Duke Institute for Health Innovation (DIHI) has refined an approach to successfully execute a yearly portfolio of pilot projects. We have completed nearly a hundred pilot projects. And while the number of projects we work on each year has remained relatively stable, the cumulative impact of successful projects has grown tremendously. Our team is an exemplar of how innovation can be done within a health system. But until recently, we were unable to supply settings beyond Duke with the ingredients needed for change.

In 2021, three different drivers for growth created a watershed moment. Below we describe the three drivers for growth and acknowledge the collaborators and funders who are helping scale DIHI's impact beyond Duke.

First, a number of DIHI innovators secured external research grants:

Timothy Dunn was awarded a five-year National Institutes of Health Research Project Grant (NIH R01) titled "Building and Implementing a Traumatic Brain Injury (TBI) Prognostic Model Featuring Real-Time Analysis of Brain CT Images." This project will build off prior work that developed machine learning algorithms to predict TBI prognosis. As part of this project, we will evaluate algorithms built at Duke on data from external health systems.

Lance Okeke and Meredith Edwards were awarded a five-year NIH R01 titled "Leveraging Local Health System Electronic Health Record Data to Enhance Pre-exposure Prophylaxis (PrEP) Access in Southeastern Louisiana: A Community-Informed

Approach." This project builds off prior work to curate local Electronic Health Record (HER) wi data to develop and validate a machine learning algorithm to predict incident HIV. As part of this project, we will externally validate the algorithm at two health systems in Louisiana and then there will be a prospective clinical trial to evaluate the impact of the program on PrEP update.

Manesh Patel and Cara O'Brien were awarded an eighteen-month fast-track NIH Small Business Technology Transfer (STTR) titled "Solving Sepsis: Early Identification and Prompt Management Using Machine Learning." This project builds off prior work to develop, validate, and integrate Sepsis Watch into clinical care across Duke Health emergency departments. As part of this project, we will externally validate the Sepsis Watch algorithm at two health systems and will prospectively validate the algorithm to assess model robustness over time.

Second, DIHI was invited to participate in the Data & Research Core of AIM-AHEAD

The NIH launched AIM-AHEAD (Artificial Intelligence /Machine Learning Consortium to Advance Health Equity and Researcher Diversity) and OCHIN invited DIHI to participate in the Data & Research Core (OCHIN: Oregon Community Health Network, a nonprofit health care innovation center designed to provide knowledge solutions that promote quality, affordable health care for all). OCHIN recognized our team's expertise in the curation of real-world data to develop and validate AI/ML algorithms and our data pipeline infrastructure to integrate AI/ML algorithms into the EHR. As part of this project, we will integrate the DIHI data pipeline into OCHIN's EHR to demonstrate the feasibility of prospectively validating AI/ML algorithms.

Third, DIHI received several grants from foundations to build capacity for AI/ML software adoption in diverse settings

The Gordon and Betty Moore Foundation awarded DIHI a grant titled "Health AI Partnership: An Innovation and Learning Network to Facilitate Safe, Effective, and Responsible Diffusion of AI Software Applied to Health Care Delivery Settings." This project brings together regulatory, engineering, social science, legal, and technology leaders from Duke Health, Mayo Clinic, DLA Piper, and UC Berkeley to promote the ethical adoption of AI software used in clinical care. The project builds off prior work helping the American Medical Association advance policies related to AI software development and reimbursement. As part of this project, we will build an online resource to help clinical and operational teams through the key decision points throughout the AI software adoption life cycle. We will also host quarterly workshops that bring together teams from across sites and expert discussants to develop actionable guidance that health systems can put into practice to tackle the most complex challenges.

The Patrick J. McGovern Foundation awarded DIHI a grant titled "Analytical Tools and Documentation Frameworks for Health AI Software Audits, Evaluations, and Monitoring." This project features a close collaboration between Duke Health and Aga Khan University (AKU) in Karachi, Pakistan to advance how health systems across settings assess AI software. The project builds off prior work developing Model Facts labels and educational materials to help clinical stakeholders effectively use AI software. As part of this project, we will develop documentation templates for Sepsis Watch and inpatient mortality algorithms at Duke Health, as well as two algorithms at AKU to disseminate publicly.

Our team at DIHI is very fortunate to be able to launch these new initiatives to scale technologies and build capabilities beyond Duke. We are especially excited to be able to diffuse AI software beyond high-resource settings like Duke Health to help close the digital divide that prevents many low-resource settings from being able to effectively use AI. We are grateful to our funders and collaborators for engaging us in challenging work that furthers our mission to catalyze transformative innovations in health and health care.

DCRI-DIHI Campfire

END OF YEAR PROJECT REPORT

Catalyzing Innovations for Clinical and Translational Research

Matthew Wilson, Nikki Harding and Natalie Sayewich

The Duke Clinical Research Institute (DCRI) and Duke Institute for Health Innovation (DIHI) continue to inspire and foster the development of novel ideas for improving clinical research through the Innovation Campfire program in 2022. The second annual event received several team submissions aligned with strategic priorities for this year: (a) Lead in engaging people, (b) Lead in hybrid/virtual trials. Related topics included:

- Increase engagement of underrepresented populations in research
- Attraction and attrition
- Digitization of studies
- RWD and RWE generation

Through due diligence and evaluation through a pitch presentation to all DCRI faculty and staff, four high-potential teams and projects were chosen to receive funding totaling over $500,000 to develop and implement their Innovation Campfire II ideas:

- Home is Where the Heart Gets Fit: An App-enabled Remote Fitness and Rehab Program for Children with Heart Disease. Reid Chamberlain, Kevin Hill, Greg Fleming, Kathleen Wood, Ilana Osten, Jennifer Martin, Stefany Olague, Anthony Cunningham
- We're Going to Talk about BRUNO: Building Relationships that are Unified & Nicely Organized. Jenny Cook, Emily O'Brien, Renee Leverty, Rashad Rahman, Vincent Miller
- Assessing Effectiveness of Patient-Directed Data Capture. Keith Marsolo, Lauren Cohen, Darcy Louzao, Gretchen Sanders, Lisa Wruck, Adam Devore, Manesh Patel
- The Pediatric Cardiac Screening Data Warehouse. Salim Idriss, Valarie Morrow, Vincent Miller, Samia Baluch

Tyrus Rorick, Head of Operations for DCRI said, "The winning 2.0 projects are aligned with key DCRI and industry strategies in operationalizing innovative research. From new apps to use of data to relationship building, the work these teams have committed to deliver will continue to further DCRI's mission."

As project milestones are reached, further updates will be shared periodically. The impact of inaugural campfire projects is described later in this section.

01 DCRI Repository for Research

Sybil Wilson RN, Ben Goldstein, PhD,
Jack Shostak, Gary Dunn, Kevin Anstrom PhD

SUMMARY

There was no way to use the vast amount of gained knowledge to answer questions or problem-solve. We envisioned a large repository where data could be loaded in a usable format to answer questions and be available for research. Our use case for the project loss is to follow-up (LTFU). This would result in a common repository for data that would start small and continue to build across historical and current data.

REPORT

Collectively, we had already spent our time and effort to gain data. At that time, access and use to improve our knowledge base or abilities was limited. There was no common repository for historical data nor current data and what we had was not set up for data mining or AI. We needed to explore. We were able to use two trials and show that having data set up across large trials eases additional use-case scenario runs. While one run of the data did not yield what we were looking for, we were able to see the advantage of data set-up/scaffolding to be able to answer many questions in the future. Our project work demonstrated that we would be able to build an internal scaffold so that analytics can be run ad hoc to further research. We were also able to glean the potential for additional research into the LTFU vertical. Discussions have begun for next steps for LTFU and we are in the planning stage. Elements of this discussion have been presented to Ty Rorick, Duke Clinical Research Institute (DCRI) Head of Operations.

AUC SCORES FOR LOGISTIC LASSO REGRESSION MODELS

WHAT WE PROPOSED

DCRI Repository for Research
Vision of a large repository where data can be loaded in a usable format to answer questions and be available for research.

The solution was the creation of a DCRI Repository for Research with scaffolding (formatting) to be able to access and store data that can enhance and/or create solutions for researchers. We would use the LTFU case scenario.

Resulting Project Campfire
Take 2 projects (Tecos/Excel) and run a retrospective proof-of-concept study for the ability to train machine learning models on relevant trial data as a way to provide, test, and measure the accuracy of predicting outcomes.

What is out there...	Offerings/Hosting/Tools to help build	Internal to DCRI	Current and ongoing efforts in data platforms with TDS and Digital Solutions allowed for the project to be re-scoped to the cross study data evaluation of LTFU and a model with which to be able to use the data
Converge Health, Deloitte	Data Minor	DCRI parallel review during proof-of-concept study	
Zaloni	Data Lake in a Box	DCRI parallel review during proof-of-concept study	
Snowflake	Lake (collection of data assets, less formal) vs. Warehouse (a defined construct)	DCRI parallel review during proof-of-concept study	
DCRI internal build	Pull studies into format for research access	Proposal for Campfire innovation	

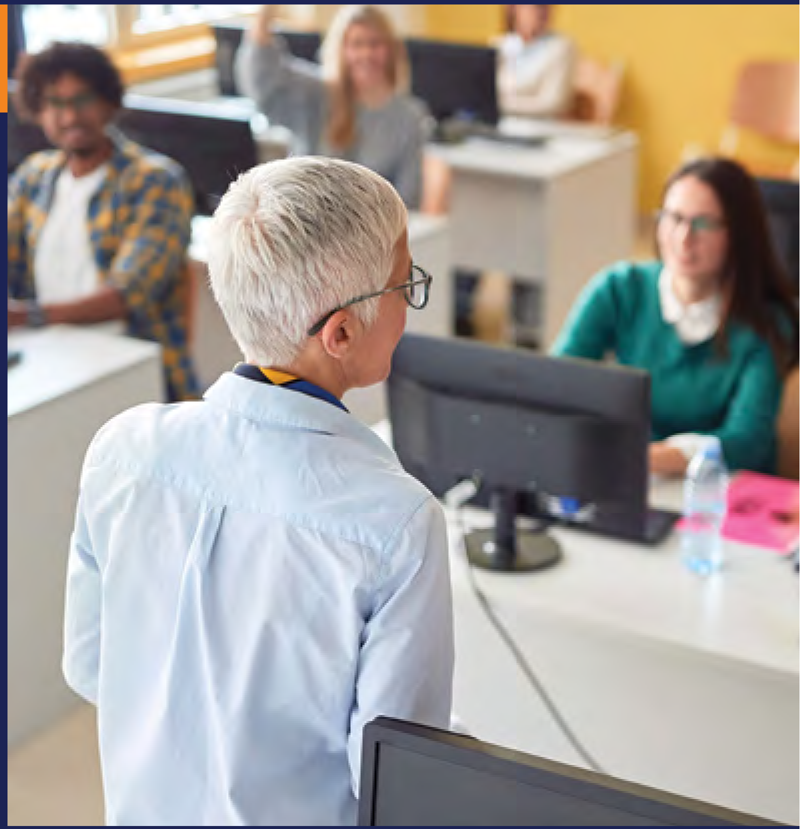

02 Expanding Delivery Science at the DCRI

Hayden Bosworth, Ph.D., Neha Pagitipati, MD, MPH, Lesley Curtis, Ph.D., Megan Oakes, MS, Monica Leyva, MHA

SUMMARY

The DCRI had no defined implementation science offering which would allow for a sharing of knowledge and experience. The aim of the team was to add implementation science educational tools and resources to improve internal and external stakeholder understanding of implementation science and its role in current and future research projects. Earlier adoption of implementation methods in research offerings can potentially translate into earlier adoption, reduce the cost of delays, and help identify proven strategies for success.

REPORT

The Duke Clinical Research Institute (DCRI) had no defined implementation science offering which would allow for a sharing of knowledge and experience. Implementation of evidence-based therapies and strategies in clinical practice are often delayed and inefficient. This delay costs health systems billions of dollars.

Implementation of science education, tools and resources will increase the value of what DCRI has to offer to internal and external stakeholders. We proposed to add implementation science educational tools and resources to improve internal and external stakeholder understanding of implementation science and its role in current and future research projects. The initial investment brings together population health faculty, DCRI operations and clinical leadership to fill the current void in tools and resources. The personnel identified will assist with protocol and business model development, as well as implement communication and marketing strategies. By leveraging talent and resources, the DCRI will be able to add implementation offerings to current research projects and proposals.

Swifter adoption of implementation methods in research offerings can potentially translate into earlier adoption, reduce the cost of delays, and help identify proven strategies for success.

PROJECTS DELIVERABLES INCLUDE:

Workforce Education Development
Seminars, short courses, and quick reference cards
- Quick Reference Cards (QRCs) attached
- Short course learning modules (external vendor)

Toolkit
Additional resources for internal and external stakeholders
- Lecture series
- Learning labs for external stakeholders

Consulting/team
- Team to assist in protocol development and bring in new business (at time of publication, working with several project leaders to incorporate implementation science in their projects and proposals)

Initial output includes a signed contract for a HF stakeholder project with Boehringer Ingelheim. The project includes qualitative interviews of stakeholders in the HF arena. The information gained from the qualitative interviews and analysis will be used to identify gaps in existing implementation strategies (approaches that will facilitate the uptake of evidence-based practice or therapy), which will lead to future programs that develop and test these interventions.

Project contract total was $135, 019.00. The project will include an academic output of a Landscape analysis and publications to follow. At the time of publication, discussions were already underway with the funder for a "Phase 2" project, which would immediately follow the final project deliverable.

GOALS FOR 2023 AND BEYOND INCLUDE:

- Establishing DCRI as a leader in the implementation/delivery science research and training
- Increasing the value proposition and meet growing demand
- Continuing education for internal and external stakeholders to build name recognition and support marketing
- Sustaining internal teams to meet demand and disseminate new knowledge and methods and help others pitch, plan and fund implementation science studies
- Sustaining an expert team to work on grant and funding proposals
- Formalizing a Center of Excellence branching into all therapeutic areas that focuses on building implementation science program and bringing in new business
- Communicating a sustainable model to solicit, leverage, and build implementation science resources
- Hybridizing a model of partnership between DCRI and Population Health that provides a mentor/trainee transition process. This would create a diverse pipeline of researchers specializing in implementation work. More work could transition to junior staff. (Initial Leadership: Hayden and Neha)

GETTING STARTED WITH INNOVATIONS IN SERVICE DELIVERY

LECTURE / ACTIVITY
- Key ingredients of service delivery projects (L)
- Defining the problem and intervention strategy (GW)

DEFINING ACTIVE INGREDIENTS AND IDENTIFYING THE RIGHT FRAMEWORK
- Example Study: Identifying a framework for a multilevel intervention (L)

DESIGN CONSIDERATIONS
- Design considerations for service delivery innovations (L)
- Identyfying design components (GW)
- Time plan, project layout (GW)
- Expert Debriefing (FF)

MEASURING IMPLEMENTATION
- Measuring Implementation Outcomes (L)
- Defining Implementation Measurements (GW)
- Successful Service Delivery Proposal (L)

L=LECTURE, GW= GROUP WORK, FF=FORMATIVE FEEDBACK

03 Project Research Voice

Manesh Patel, MD, Neha Pagidapati, MD, Nishant Shah, MD, Sharon Califf, Linda Lillis, Malaika Bhayana

SUMMARY

To improve clinical research enrollment and engagement we will test a video platform to capture the perspectives of any person involved in research and provide them with a community for learning and support of all kinds.

REPORT

Clinical research suffers from poor enrollment and engagement. The following is essential to clinical research engagement: Trust in the research process, inclusion and diversity, understanding the research study, accessibility to asking questions of your local healthcare team and getting to the patient. There is no easy platform for local clinicians and patients to generate media that captures their experience – their VOICE.

Our solution to improve enrollment and site engagement is "Test Kitchen." This is a video platform that can serve and capture the perspective of people involved in research: patients, coordinators, physicians, statisticians and others. It will support enrollment, patient education, broader research community participation. This "Test Kitchen" may be thought of as a "DCRI Tik Tok for Research Studies" and include commentary, stories, songs, dances, and poems. It may include 30-60 second perspectives on diseases, therapies people take, ongoing studies, or starting studies. Videos may be shared on the YouTube platform for application in other studies as applicable.

To monitor outcomes and study the impact, we will track enrollment, site operations, and site engagement in correlation with video viewing metrics. We aim for enrollment improvement and increased site engagement as video participation and viewing metrics increase. Positive outcomes would promote a cost-effective and sustainable tool to engage and promote research objectives. With success, we could consider a pricing structure, propose additional language adaptions for future projects, and publish our findings.

04 Direct to Patient Digital Health Platform

Kevin Hill, MD, Christoph Hornik, MD, PhD, MPH, Sarah Tallent, NP, DNP, Annie Schmelzer, NP, Lillian Kang, MD, Reid Chamberlain, MD, Stefany Olague, MPH, Anthony Cunningham, RRT

SUMMARY

In pediatric patients with congenital heart disease (CHD), research, clinical care, and patient engagement are complex, competing, and affect outcomes. Our solution is to develop and deploy the Kid's Heart mobile application (Kid's Heart app) for families who have children with single ventricle heart defects. A pilot study will inform further development of the application to enable a roll-out to other conditions and to include additional modules such as cardiac rehabilitation tools. The application continues to provide the team with further expertise in delivering research and care directly to participants and their families. Lastly, the application increases the availability of care to communities that are otherwise distant from adequate healthcare resources.

REPORT

Congenital heart disease (CHD) is rare and has a heterogeneous patient cohort, one that is difficult to enroll into research. Most CHD research is retrospective and based on siloed registries. Staff from the Duke Clinical Research Institute and the Duke Pediatric and Congenital Heart Center teamed up to develop and launch the Kid's Heart app for families who have children with single ventricle heart defects. Powered by Pattern Health, the application provides a mechanism to easily report at-home monitoring results, including vital signs and feeding. Care content is also accessible to the users and includes modules such as cardiac medications, infant CPR, and planning for a stay at the hospital. Red flag monitoring is built into the application and is visible to the care team through the monitoring portal. Also available to users (only) is a support group chat function. Additionally, the future addition of a Global Unique Identifier (GUID) to connect registries transforms the potential for both retrospective and prospective research.

At the time of publication, this pilot had enrolled 24 participants. We continue to work on the integration of the GUID. The care team has gained insight into the use of patient-reported outcomes in continuing care. The team is actively working on a conference abstract and manuscript using the data obtained. The next steps include the development of additional modules for use in conjunction with each other or as standalone mobile applications, including, but not limited to:
(1) a pediatric remote cardiac rehabilitation and exercise program, (2) a virtual neurodevelopment assessment for pediatric patients with CHD, and (3) the delivery of application content in Spanish to better support our community.

With FDA approval, we have utilized funding from the Global Pediatric Clinical Trials Network award to support our efforts in developing and deploying the platform. Many of the modules developed to support the project can be replicated as-is by research teams running projects with similar objectives.

05 DUKE Empower:

A Platform to Engage, Inform, Probe, Promote Wellness, and Foster Community Among Patients with Rare Diseases.

Rosalia Blanco, MBA, Vincent Miller, MMCi, Reed Johnson, PhD, Scott Palmer, MD, MHS, Shelby Reed, PhD, Aparna Swaminathan, MD, MHS

SUMMARY

Thirty million people in the US have a rare disease. Patients with rare diseases have difficulties identifying and accessing specialized care, resulting in delays and inaccuracies in treatments as well as gaps in patient education. We develop a platform to Engage, inform, Probe, promote Wellness, and foster community among patients with Rare diseases. This platform will create a community of interstitial lung disease (ILD) patients at Duke; educate them about clinical research opportunities at Duke and about ILD diseases; and enable them to track disease symptoms. This ultimately aims to improve ILD patient health, eating, and exercise habits.

REPORT

Patients with interstitial lung disease (ILD) have difficulty identifying and accessing specialized care. This results in delays and inaccuracies in treatments as well as gaps in patient education. While Duke has a large ILD clinic, many patients are not aware of ongoing research studies and, consequently, are not enrolled. Furthermore, many patients live far away and do not come to Duke regularly, but would be willing to come far to Duke to participate in a research study. Clinicians would also be interested in remotely gathering more data on patients, including symptoms and steps. Furthermore, exercise is an important component.

To help remediate these problems, we will build an App to Engage patients in research, inform, Probe, promote Wellness, and foster community among patients with rare diseases. At the time of this publication, the project is in progress. Even so, we have received the following important feedback from patients using the app:

- Patients are expressing interest in participating in clinical research studies
- Patients are providing feedback about features and information that they would like to see included in the App
- They are completing surveys and questionnaires that can be used for research purposes

This App provides a framework to expand to other rare diseases that face similar challenges to patients with ILD. It would particularly benefit conditions where patients would benefit from behavioral modification or tracking aspects such as nutrition and exercise. Sponsors have expressed interest in using the App for other studies, including to facilitate direct to patient recruitment and patient engagement in expansion of an ongoing ILD registry. We aim to present and publish our results.

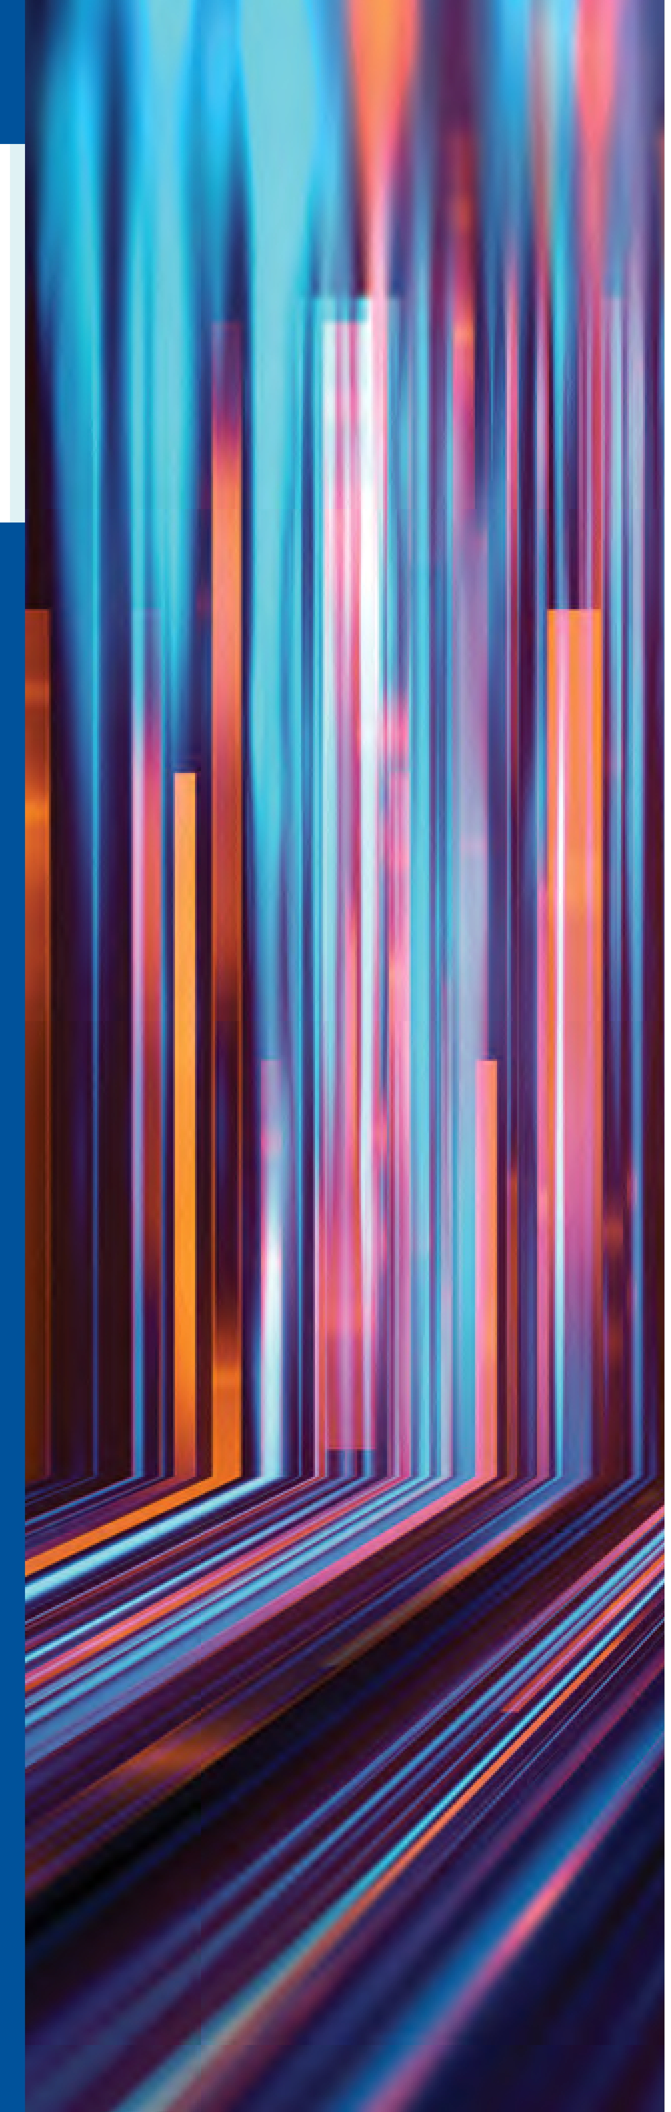

Presentations

1. Chan, N. W. et al. "Social Determinants of Health Data Capture Within National and Health System Data Sources", is accepted for an oral presentation in the Scientific Forum at Clinical Congress 2022, taking place October 16-20 in San Diego, CA.
2. Niederhoffer K, Knechtle W, Uronis H, Shariff A, et al. (2022). A Machine Learning Model to Predict Hospital Admissions and Emergency Department Use in Patients' Immune Checkpoint Inhibitors. Machine Learning Healthcare Conference (MLHC) 2022 poster. Durham, NC, 2022
3. Ochoa T, Knechtle W, Sendak M. Development of a Machine Learning Model for Prediction of Mortality in Lung Transplant Patients. Poster presentation at Machine Learning for Healthcare (MLHC), August 2022, Durham, NC.
4. Premo H, Shi H, Knechtle W, Kazaure H, et al. (2022). A Geriatric-Specific Morbidity and Mortality Perioperative Risk Stratification Tool. Machine Learning for Healthcare (MLHC) 2022. Poster presentation. Durham, NC, 2022.
5. Shen R; Ratliff W; Sendak M; Kapadia N; Burrows B, et al. Algorithm Development for Duke Emergency Pre-hospital Capacity Management. Machine Learning for Healthcare 2022 – Clinical Abstract, Software, and Demo Track. August 5-6, 2022.
6. Shen R, Weissle EH, Ratliff W, Nichols M, Hintze B, Gao M, Sendak M, Balu S, Jones S, et al. Improving Equity and Value of Peripheral Artery Disease Care at a Population Level. Machine Learning for Healthcare 2022 – Clinical Abstract, Software, and Demo Track. August 5-6, 2022.
7. Tang, L, Ratliff, W, Sendak, M, Gao, M., Nichols, M, Revoir, M, Yashar, F, Yao, J, Balu, S, Subramanian, N, Uhl, T, Denis, L, Sterrett, E. Identifying Sepsis in real-time for Duke University Hospital Pediatric Patients [Poster presented]. Machine Learning for Healthcare (MLHC) Conference. August 2022. Durham, NC, USA.
8. Zanolli N, Knechtle W, Havrilesky L, Davidson B. Integration of a Post-operative Opioid Calculator into an Academic Gynecologic Surgery Practice. North Carolina Obstetrical and Gynecologic Society (NCOGS) 2022 Annual Meeting. Kiawah, SC, 2022.
9. Zanolli N, Knechtle W, Sendak M, Havrilesky L, Davidson B, Integration of a Post-operative Opioid Calculator into an Academic Gynecologic Surgery Practice. Machine Learning Conference for Healthcare. Durham, NC, 2022.

Publications

1. Bellantoni, Julia B. MD, et al. "Implementation of a Telehealth Videoconference to Improve Hospital-to-Skilled Nursing Care Transitions: Preliminary Data." Journal of the American Geriatrics Society, Mar. 2022, pp. 1–10, https://doi.org/10.1111/jgs.17751.
2. Corey, Kristin M, et al. "Exposure and Outcome in Practice: A Retrospective Cohort Study between Fibrinolytic Suppression and Hypercoagulability, the Severity of Hypoxemia, and Mortality in COVID-19 Patients." Anesthesiology, Apr. 2022, p. 10.1097/ALN.0000000000004239, https://doi.org/10.1097/ALN.0000000000004239.
3. Dohlman, Anders B, et al. "The Cancer Microbiome Atlas: A Pan-Cancer Comparative Analysis to Distinguish Tissue-Resident Microbiota from Contaminants." Cell Host & Microbe, vol. 29, no. 2, 2021, pp. 281-298.e5, https://doi.org/10.1016/j.chom.2020.12.001.
4. Fenn, Alexander, et al. "Development and Validation of Machine Learning Models to Predict Admission From Emergency Department to Inpatient and Intensive Care Units." Annals of Emergency Medicine, vol. 78, no. 2, 2021, pp. 290–302, https://doi.org/10.1016/j.annemergmed.2021.02.029.
5. Honeycutt, Christopher Cole, et al. "Assessment of Spanish Translation of Websites at Top-Ranked US Hospitals." JAMA Network Open, vol. 4, no. 2, Feb. 2021, p. e2037196, https://doi.org/10.1001/jamanetworkopen.2020.37196.
6. Kansal, Aman, Cynthia L. Green, et al. "Electronic Health Record Integration of Predictive Analytics to Select High-Risk Stable Patients With Non–ST-Segment–Elevation Myocardial Infarction for Intensive Care Unit Admission." Circulation: Cardiovascular Quality and Outcomes, vol. 14, no. 4, 2021, https://doi.org/10.1161/CIRCOUTCOMES.120.007602.
7. Kansal, Aman, Michael Gao, et al. "Impact of Diagnosis Code Grouping Method on Clinical Prediction Model Performance: A Multi-Site Retrospective Observational Study." International Journal of Medical Informatics, vol. 151, 2021, p. 104466, https://doi.org/10.1016/j.ijmedinf.2021.104466.
8. McClain, Micah T, et al. "A Blood-Based Host Gene Expression Assay for Early Detection of Respiratory Viral Infection: An Index-Cluster Prospective Cohort Study." The Lancet Infectious Diseases, vol. 21, no. 3, 2021, pp. 396–404, https://doi.org/10.1016/S1473-3099(20)30486-2.
9. Sendak, Mark, et al. "AI Transparency: Why We Should Label Algorithms Like Food Products." Data Science Institute; American College of Radiology, 20 Oct. 2021, https://www.acrdsi.org/DSIBlog/2021/10/20/AI-Transparency-Why-We-Should-Label-Algorithms-Like-Food-Products.
10. Sendak, Mark P, et al. "Looking for Clinician Involvement under the Wrong Lamp Post: The Need for Collaboration Measures." Journal of the American Medical Informatics Association, Sept. 2021, p. ocab129, https://doi.org/10.1093/jamia/ocab129.
11. Sendak, Mark P, et al. "Preliminary Results of a Clinical Research and Innovation Scholarship to Prepare Medical Students to Lead Innovations in Health Care." Healthcare, vol. 9, no. 3, 2021, p. 100555, https://doi.org/10.1016/j.hjdsi.2021.100555.
12. Xiu, Zidi, et al. "Variational Disentanglement for Rare Event Modeling." Proceedings of the... AAAI Conference on Artificial Intelligence. AAAI Conference on Artificial Intelligence, vol. 35, no. 12, May 2021, pp. 10469–77. https://ojs.aaai.org/index.php/AAAI/article/view/17253/17060
13. Yang, Zhou, et al. "Advancing Primary Care with Artificial Intelligence and Machine Learning." Healthcare, vol. 10, no. 1, 2022, p. 100594, https://doi.org/10.1016/j.hjdsi.2021.100594.
14. Zanolli, Nicole, et al. Integration of a Post-Operative Opioid Calculator into an Academic Gynecologic Surgery Practice. NC OBGYN-ACOG Annual meeting, Kiawah Island Golf Resort, Kiawah Island, SC.
15. Zullig, Leah L, et al. "Low-Touch, Team-Based Care for Co-Morbidity Management in Cancer Patients: The ONE TEAM Randomized Controlled Trial." BMC Family Practice, vol. 22, no. 1, Nov. 2021, p. 234, https://doi.org/10.1186/s12875-021-01569-8.

DIHI Innovation Staff


Craig Albanese, MD, MBA
Executive Director for DIHI
Executive Vice President and Chief Operating Officer, DUHS


Suresh Balu, MBA, MS
Director for DIHI
Associate Dean, Innovation and Partnership, School of Medicine


Willie Boag, PhD
Data Scientist


Jamie Daniel, BS
Solutions Architect


Michael Gao, MS
Data Scientist


Matt Gardner, BS
Solutions Architect


Alifia Hasan, B.Pharm, MBA
Innovation Program Manager


Bradley Hintze, PhD
Data Engineer

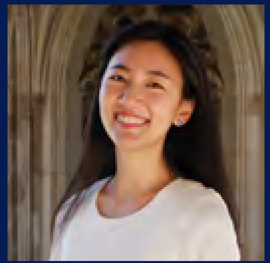
Jee Young Kim, PhD
Data Scientist


Will Knechtle, MBA, MPH
Innovation Program Manager


Ebony Nash, CEAP
Program Coordinator


Marshall Nichols, MS
Data Engineer


Will Ratliff, MBA
Innovation Program Manager


Mike Revoir, BS
Solutions Architect


Mark Sendak, MD, MPP
Clinical Data Scientist


Linda Tang, BS
Data Scientist

DIHI Innovation Scholars

Kaivalya (Kai) Deshpande
Pulmonary and Critical Care Fellow

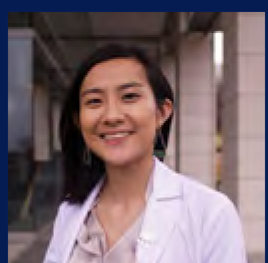

Norine Chan
DIHI Scholar

Kira Niederhoffer
DIHI Scholar

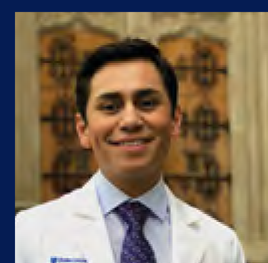

Tim Ochoa
DIHI Scholar

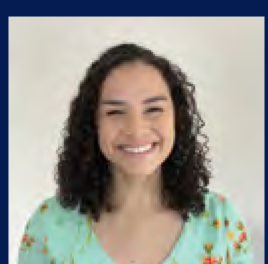

Hayley Premo
DIHI Scholar

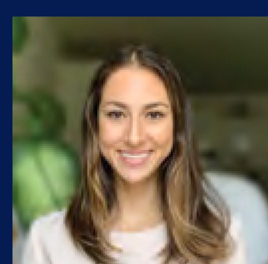

Connie Scoggins
DIHI Incoming Scholar

Rebecca Shen
DIHI Scholar

DIHI Data Science Scholars

Claire Carroll
Cole Fitzgerald
Freya Gulamali
Jeff Hogg
Joanne Kim
Megan Richards
Harvey Shi
Gaurav Siredeshmukh
Sabrina Wong
Jiayu Yao
Jennifer Young

DIHI Collaborators

Yvonne Acker, RN, BSN
Terrence Allen, MBBS
John Anderson, MD
Jonathan Bae, MD
Matthew Barber, MD, MHS
Melissa Bauer, DO
Armando Bedoya, MD, MMCi
Andrew Berchuck, MD
Amanda Bisset, MD
Jeanna Blitz, MD
John Bonnewell, MD, MSc
Brandi Bottiger, MD
Ebony Boulware, MD, MPH
Kaitlin Boyle
Tres Brown III
Dan Buckland, MD, PhD
Brian Burrows, MD
Blake Cameron, MD, MBI
Krista Camuglia
David Casarett, MD, MA
Anisha Chandiramani, MD
David Chang Villacreses, MSc
Devon Check, PhD
Seuphy Chen, MD
Alex Cho, MD, MBA
Sammy Chouffani El Fassi
Richard Chung, MD
Meredith Clement, MD
Kathleen Cooney, MD
Kristin Corey, MD
Daniel Costello, MPA
Tom Daigle
Brittany Davidson, MD
Liset Denis, BSN, RN
Timothy Dunn, PhD
Jill Engel, DNP
Robin Famiglietti, PhD, MBA
Jeffrey Ferranti, MD, MS
Christine Fowler
Mary Ann Fuchs, DNP, RN
William F Fulkerson, MD
Anthony Fuller, MD, MScGH
Katie Galbraith, MBA
Ziad Gellad, MD, MPH
Charles Gerardo, MD
Jennifer Gilner, MD, PhD
Zac Ginsberg, MD, MPP
Ben Goldstein, PhD
Lawrence Greenblatt, MD
Barbara Griffith, MD
Padma Gulur, MD
Ashraf Habib, MBBCh, MSc
Michael Haglund, MD, PhD, Med
Brian Halstater, MD
Matthew Hartwig, MD
Laura Havrilesky, MD, MHSc
Ricardo Henao, PhD
Adrian Hernandez, MD, MHS
Brenna Hughes, MD, MSc
Kimberly Jackson, MD
J. Eric Jelovsek, MD, MMEd, MSDS
Anjni Joiner, DO, MPH
Schuyler Jones, MD
Ibukun C. Kalu, MD
Aparna Kamath, MD
Neel Kapadia, MD
Jason Katz, MD, MHS
Hadiza Kazaure, MD
Joe Kelly
Michael Kent, MD
Kelly Kester, MSN
Meenal K. Kheterpal, MD
Tara Kinard, RN, MSN, MBA
Allan Kirk, MD, PhD
Jacob Klapper, MD
Mary Klotman, MD
Stuart Knechtle, MD
Charley Kneifel, PhD

DIHI Collaborators

Karan Kumar, MD, MS
Maribeth Kuntz, PA-C
Andrii Kuraksa, MBA
Andrii Kuraksa, MBA
Walter Kwiatek
Sandhya Lagoo-Deenadayalan, MD, PhD
Kristen Lakis, MSW
Shamayla Lando
Jeffrey Langdon, MHA
Jennifer Li, MD
Kelly Lindblom, PhD
Cooper Linton, MSHA
Michael Lipkin, MD
Keith Marsolo, PhD
Mary Martin, MPA
Leslie Mason, PhD, MSN
Joseph Mathew, MD, MBA, MHSc
Shelley McDonald, DO, PhD
Lisa McElroy, MD, MS
Eugenia McPeek Hinz, MD, MS
Rob Mentz, MD
David Ming, MD
Andrew Muir, MD, MHS
Lisa Nadler, MD
Susanna Naggie, MD, MHS
Dennis Narcisse, MD
Kristin Newby, MD, MHS
Paul R. Newman, MHA
Nick Nguyen, MHA
Michelle Nixon, PhD
Cara O'Brien, MD
Rob Odom, MA
Lance Okeke, MD
Vicky Orto, DNP, RN, NEA-BC
Karen Osborne, RN, BSN
Marcia Owen
Thomas Owens, MD
Neha Pagidipati, MD, MPH
Theodore Pappas, MD
Kishan Parikh, MD
Christine Park
Amanda Parrish, PhD
Akash Patel
Chetan Patel, MD
Manesh Patel, MD
Yuval Patel, MD, MHS
Michael Pencina, PhD
Donna Phinney, RN, MSN
Christopher Polage, MD, MAS
Noppon Pooh Setji, MD
Eric Poon, MD, MPH
John Ragsdale, MD
Vincent Ramos, PhD, MPH, RN
Priscilla Ramseur, DNP, RN, CNOR, NEA-BC
Robin Rasor, MS
Robin Rasor, MS
Ann Reed, MD
Henry Rice, MD
Matthew Roman, MHA, MMCi
Tyrus Rorick
Adia Ross, MD
Matt Rougeux, MHA
Moira Rynn, MD
John Sampson, MD, PhD, MBA, MHSc
Devdutta Sangvai, MD, MBA
Natalie Sayewich, MA
Mary Schilder, RN
Teri-Lynne Sennett
Kevin Shah, MD
Svati Shah, MD
Colleen Shannon, JD
Richard Shannon, MD
Afreen Shariff, MD
Shilpa Shelton, MHA, FACHE
John Shepherd
Harvey Shi
Kristen Shirley, MD
Gail Shulby, RN, MA, CPPS
Lauren Siewny, MD
Christina Silcox, PhD
Justin Silverman, MD, PhD
Stephen Skelton, MA
Ryan Smith, MPP
Becky Smith, MD
Stuart Smith
Laurie Snyder, MD, MHS
Denise Snyder, MS, RD
Michael Spiritos, MD
Emily Sterrett, MD
Allison Stewart
Neel Subramanian, MD
Debra Sudan, MD
Dustin Tart, RN
James Tcheng, MD
B. Jason Theiling, MD
Tammy Uhl, BSN, RN
Carey Unger, MHA
Sreekanth Vemulapalli, MD
Anthony Viera, MD, MPH
Momen Wahidi, MD, MBA
Andrew Wang, MD
Wendy Webster, MA, MBA
Zach Wegermann, MD
E. Hope Weissler, MD
Julius Wilder, MD, PhD
Matt Wilson, RN
Samantha C. Wong, BS
Mary Cooter Wright, MS
Yousuf Zafar, MD, MHS

dihi.org

@DukeInnovate