2018 Impact Report

Duke Institute for Health Innovation

Implementation of Deep Learning Technologies for Sepsis Management
Page 26

Bringing innovative solutions to pressing challenges in health and healthcare

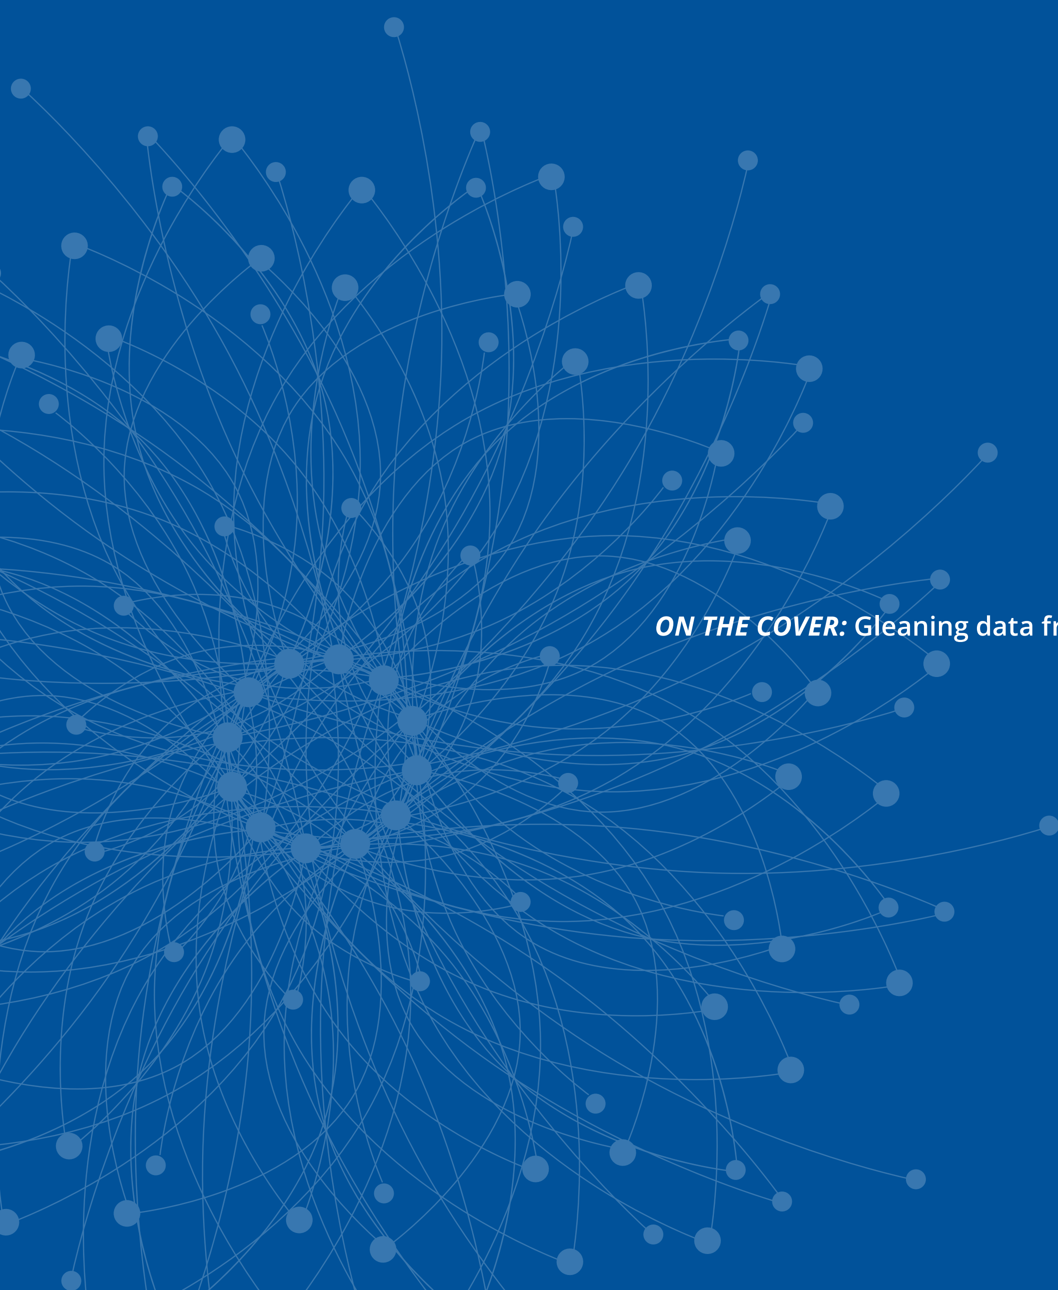

ON THE COVER: Gleaning data from bacterial growth on agar

Duke Institute for Health Innovation

2018 Impact Report

About DIHI

The Duke Institute for Health Innovation (DIHI) promotes innovation in health and health care through high-impact innovation pilots, leadership development, and cultivation of a community of entrepreneurship.

DIHI brings innovative solutions to the most pressing challenges in health and health care by catalyzing multidisciplinary teamwork across Duke University and Duke Medicine and by fostering collaborations with national and international thought leaders.

Letter from the Directors

FOR DUKE INNOVATORS AND SCHOLARS, 2018 MARKED A YEAR WHERE PATIENTS WERE SQUARELY IN THE CENTER OF OUR INNOVATION EFFORTS.

We are proud to launch the inaugural edition of The DIHI Innovation Brief. This publication will highlight the work of the Duke Institute for Health Innovation (DIHI), the people who are engaged in innovations in health and healthcare at Duke, as well as forward-thinking pieces about the future of healthcare innovation. We hope that this publication illustrates and informs, while inspiring faculty, staff, students and trainees to explore and to push the boundaries of what might be achieved through leveraging innovation, teamwork and a multidisciplinary approach to addressing the future of health and healthcare.

Over the past year, we have expanded our capabilities within data science, focused on the implementation of machine learning models, and continued to engage 3rd year medical students and trainees in research and innovation activities. Our data science team has undertaken the work of automating the ingestion of raw Medicare claims by cleaning and normalizing ACO claims data. This is a first for Duke and we believe that through this new capability, we can leverage not only clinical data to make better predictions but also claims data. We have used this capability to predict the need for palliative care consultations as well as to predict first admission to the hospital.

Over the past decades, much has been written about machine learning and the use of data and technology to enhance medical decision support. This certainly presents a tremendous opportunity, but the last mile to show the promise of machine learning in healthcare is the implementation, evaluation and scaling of theoretical models to real-world prognoses, diagnoses and clinical decision-making. Over the past year, we have developed a robust platform and the infrastructure needed to support the implementation of machine learning into clinical care.

For Duke innovators and scholars, 2018 marked a year where patients were squarely in the center of our innovation efforts. In the new academic year, we foresee innovation will continue to be a major driver of quality, productivity and outcomes for our patients and families at Duke and eventually for patients and populations everywhere.

Sincerely,

William Fulkerson, MD, MBA
Executive Director
Duke Institute for Health Innovation
Executive Vice President
Duke University Health System

Suresh Balu, MBA
Program Director
Duke Institute for Health Innovation
Associate Dean, Innovation and Partnership
Duke University School of Medicine

2017 DIHI RFA Project Summaries

Real-time Cost Transparency
in a Value Based Environment

Project Team
David E Attarian, MD
Ben Alman, MD
Chad Mather, MD
Shilpa Shelton, MHA

DIHI Team
Will ElLaissi MBA, MHA
Suresh Balu, MBA

PROBLEM
As we transition from a Fee-for-Service to a Value-Based Environment, Cost, the denominator, is often not accessible to the provider or the patient.

$$\text{Value} = \frac{\text{Quality} + \text{Customer Satisfaction}}{\text{Cost}}$$

SOLUTION
We developed comparative historical performance reports to send to total joint surgeons via Epic InBasket when an elective case is posted. Additionally, we can now share patient "out of pocket" cost estimates with providers for each posted elective case.

IMPACT
Surgeons can now access comparative historical performance reports and "out of pocket" costs per patient, closing the information gap and allowing for better shared decision-making. Further analyses will be needed to measure historical cost trends for surgeons.

Further iterations of this work will move from quarterly, static historical cost reports to monthly reports to see and catch trends faster. The Patient Revenue Management Office (PRMO) is investigating labor resources needed to ensure accuracy of patient out of pocket cost estimates, which will be sent to the physician prior to the scheduled case, allowing for shared decision-making with the patient.

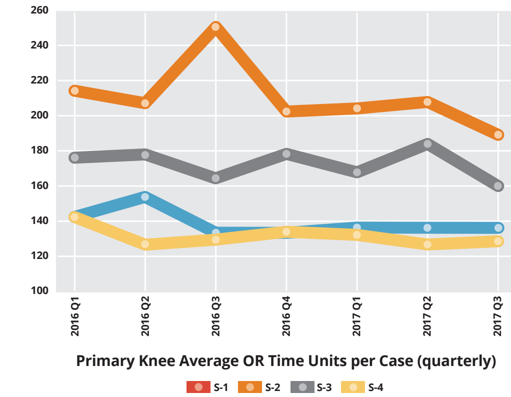

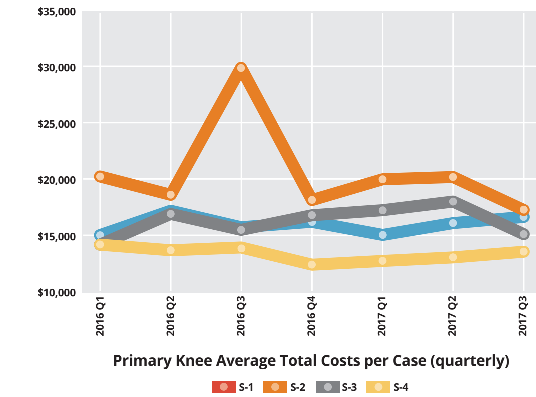

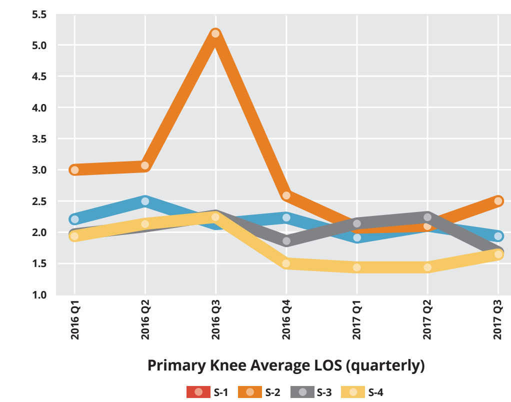

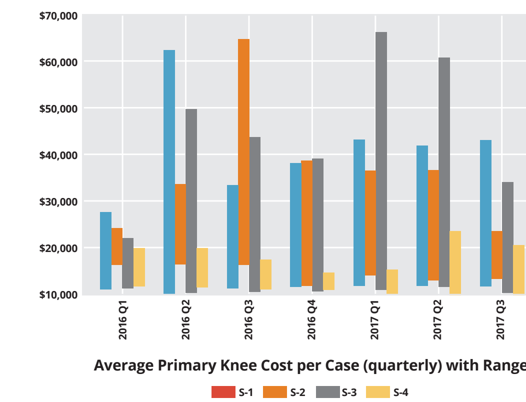

(Clockwise from top left) Quarterly trending comparison of four surgeons for average time per case, average cost per case, average cost range per case, and average length of stay (LOS)

"This is an issue that required hyper focus regarding historical costs of our surgeons as well as what expected out of pocket costs are for our patients to aid in shared decision-making... I look forward to collaborating with DIHI in the future to expand our capabilities."

– David Attarian, MD

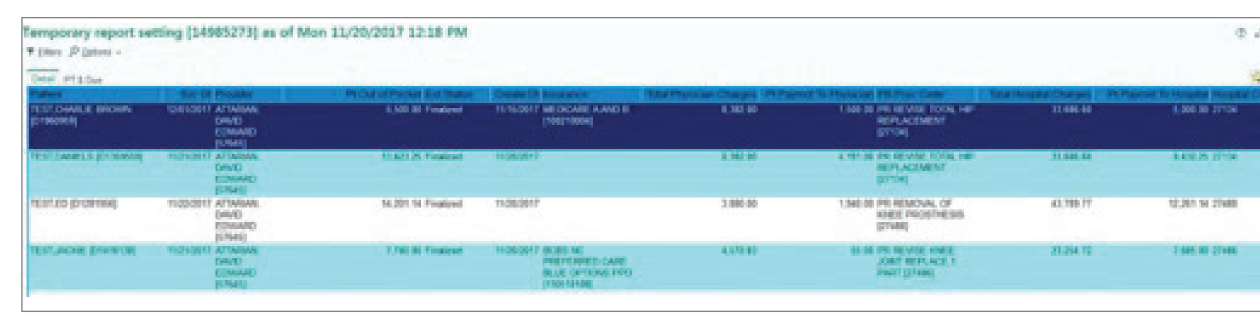

Charting functionality in Maestro Care showing detailed encounter cost information, including out of pocket costs to the patient and provider and hospital charges

Pallialytics

Project Team

Jonathan Fischer, MD
Eugenie Komives, MD
Cynthia Rudin, PhD
Stephanie Brinson
Leslie Calihman Alabi

DIHI Team

Michael Gao
Will ElLaissi MBA, MHA
Xiaozhou Wang
Daniel Costello, MPA

PROBLEM

Duke Connected Care/PHMO established a multidisciplinary palliative care virtual rounds initiative in Fall 2016. The means to identify patients with palliative care needs was highly inefficient, relying heavily on several staff to validate patient fit via chart audit. The time spent to manually identify patients came at the expense of the care management activities that would benefit patients and the healthcare system.

SOLUTION

Duke Connected Care (DCC) partnered with DIHI to design a novel machine-learning predictive model that combined claims and clinical data to more effectively target DCC's existing palliative care interventions among the MSSP population. In parallel, we developed a regression model to support and validate the novel model's predictions. Using the dual-model approach, we seek to predict up to four outcomes per patient over a 12-month time horizon: mortality, hospitalization, high Medicare costs, increasing rate of costs.

IMPACT

We produced a viable predictive model to positively affect patient care workflows. The most recent iteration of the regression model has an AUC of approximately 0.81, meaning that the model will correctly rank two random patients 81% of the time such that the first patient in the model's output is more likely to benefit from palliative care than is the second patient. Moreover, Dr. Fischer's appraisal of the regression model's most recent output indicates a 45% success rate, where success means that the model accurately identified a patient whom Dr. Fischer determines as a "good candidate" for palliative care based on chart review. In contrast, the original methods of patient identification via non-predictive algorithmic reports demonstrated a success rate of less than 20%.

Dr. Fischer has referred patients identified through Pallialytics to the DukeWELL complex care management program and has advised primary care managers about palliative care needs for additional patients already engaged in DukeWELL.

As we continue to iterate the regression model, we anticipate its success rate to increase. Additionally, we plan to continue to evaluate the effect of Pallialytics on anticipated long-term outcomes: increased volume of patients engaged in palliative care, increased referrals to palliative care clinic, improved patient/caregiver experience, increased appropriate use of hospice, and reduced unplanned admissions.

AS WE CONTINUE TO ITERATE THE REGRESSION MODEL, WE ANTICIPATE ITS SUCCESS RATE TO INCREASE.

DIHI Innovation Scholars

The Duke Institute for Health Innovation (DIHI) Clinical Research and Innovation Scholarship was founded in 2015 to support Duke medical students to help lead high-impact, high-visibility innovation projects. This scholarship is the first third year scholarship dedicated to supporting students to dedicate a single year driving an innovation project with senior faculty, operational leaders, and technology developers.

"Being surrounded by people and an environment that not only encouraged self-learning and innovative thinking, but required it, I dove in head first."

Kristin Corey

I initially chose to apply for the DIHI medical student scholarship because I had an interest in biomedical informatics and had heard about DIHI's inspiring work with EHR data. I had learned that their projects were aimed to better inform and even drive clinical decision making, improving patient and health system outcomes using patient data. So, I applied to work on a surgical data science project building a machine learning model to predict post-operative patient outcomes. Truthfully, I had no idea what "data science" and "machine learning" actually were at the time. But, being surrounded by people and an environment that not only encouraged self-learning and innovative thinking, but required it, I dove in head first, mimicking everyone around me.

Success in academic medicine is defined by certain traditional concrete metrics. Conferences attended, papers published, and influential networks joined form the mainstay measurement of academic achievement today. Because Duke University School of Medicine allows its students to experience a research year with the intellectual freedom to choose one's research topic with no regular grading system in place, this year is especially vulnerable to these classic methodologies of measuring success. Objectively reflecting back on my third year of research with DIHI, I can say that I met all three of these metrics. I've had the privilege of presenting my team's work to leadership and at national conferences while having the opportunity to formally write up our work. In the traditional medicine sense, my research year checks the intentional boxes that it was supposed to. But what I really gained from my DIHI experience was the new mindset that anything is actually possible with the right people around you. Walking away with this knowledge far surpasses those traditional accomplishments by enhancing who I am as a researcher as well as a person.

NSTEMI

Rational coronary care unit triage for stable patients with NSTEMI: Evaluating the safety and costs of a risk score-based triaging system

Project Team
Alexander Fanaroff, MD
Manesh Patel, MD
Tracy Wang, MD, MHS, MSc
Schuyler Jones, MD

DIHI Team
Aman Kansal
Krista Whalen
Mark Sendak

PROBLEM

Nationwide, intensive care unit (ICU) utilization for initially stable patients with non-ST segment elevation myocardial infarction (NSTEMI) is highly variable and does not appear to be related to patient risk. Patients admitted to the ICU have similar risk of in-hospital mortality compared with those who are not admitted to the ICU. Inappropriate admission of low-risk patients to the ICU increases costs and occupies limited ICU beds without associated benefit to the patient. Inappropriate admission of high-risk patients to a non-ICU setting may lead to ICU transfers and worse patient outcomes. Risk-based ICU utilization—that is, admitting those at highest risk of developing complications requiring ICU care to the ICU and admitting those at lower risk to a non-ICU setting—has the potential to better align resource use with patient needs.

SOLUTION

In partnership with Duke's Division of Cardiology leaders, we determined an appropriate threshold ACTION ICU score for ICU admission. For initially stable (no cardiac arrest or shock on first presentation) NSTEMI patients with a score ≤ 5, we would recommend non-ICU admission, and for patients with a score ≥ 6, we would recommend ICU admission. Working with ED physicians, we created a modified best practice advisory (BPA) that triggers each time a patient with a serum troponin level above the upper limit of normal is seen in the emergency department.

IMPACT

Broadly, risk-based ICU utilization did not appear to affect most of our outcomes in the limited time it was evaluated: the proportion of patients admitted to the ICU from the ED remained constant, there was no change in the proportion of patients transferred from a non-ICU bed to the ICU after admission, and there was no change in ICU length of stay, hospital length of stay (Figure 5), or mortality. We were surprised to discover that ICU utilization for NSTEMI patients was substantially lower than we had anticipated, with the majority of patients not being admitted to the ICU before roll-out of the BPA. We did find that, after roll-out of the BPA, the groups admitted to the ICU and not admitted to the ICU became more homogeneous in terms of risk.

This paradigm is useful in a number of areas. First, we are currently working with the transfer center to roll out a program where the ACTION ICU score is used to risk-stratify patients presenting with NSTEMI to outside hospitals, in order to determine priority for transfer. Especially at times of limited bed availability, use of the risk score may enable us to transfer high-risk patients preferentially compared with low-risk patients. Second, we would like to work with Duke Regional and Duke Raleigh Hospitals to roll out a similar BPA to calculate the ACTION ICU risk score and determine which patients with initially stable NSTEMI should be admitted to the ICU. Third, risk-based ICU admission for hemodynamically stable NSTEMI patients using an EHR-integrated risk calculator could potentially benefit a number of hospitals in the Duke-Lifepoint network and beyond.

ACTION ICU

NSTEMI

BPA triggers, physician opts to:

a) Hide BPA (patient not yet seen)

b) Cancel BPA (patient not stable, did not have NSTEMI)

c) Calculate ACTION ICU score (Figure 1)

When patients select the option to calculate the risk score, the risk score calculator pops up on the right side of the screen (Figure 2).

Five calculator elements of the risk score auto-populate, and the other four are entered by the physician entering "yes" or "no". Once these questions are answered, ACTION ICU calculates the score, displays it to the physicians, and recommends where to admit the patient (ICU vs. non-ICU) (Figure 3).

WE WERE SURPRISED TO DISCOVER THAT ICU UTILIZATION FOR NSTEMI PATIENTS WAS SUBSTANTIALLY LOWER THAN WE HAD ANTICIPATED.

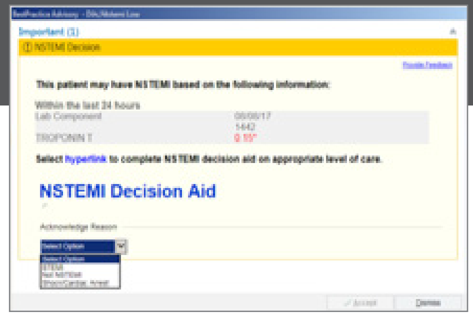

Figure 1: BPA trigger

The BPA triggers for all patients seen in the emergency department with a positive troponin. ED providers select the hyperlink if the patient is hemodynamically stable and has an NSTEMI to calculate the ACTION ICU score. If the patient is not hemodynamically stable, has a STEMI (rather than NSTEMI), or has shock/cardiac arrest, providers select this reason and cancel the BPA.

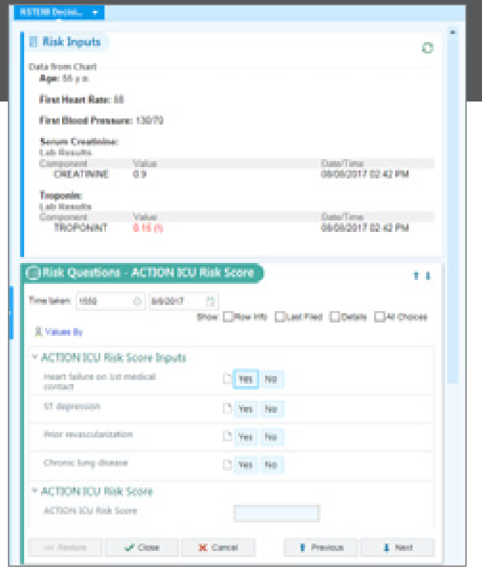

Figure 2: Risk calculator

The risk score calculator appears when providers select the option to calculate the ACTION ICU score. Age, heart rate, blood pressure, creatinine, and troponin auto-populate in the calculator, and operators select whether the patient has heart failure on exam, ST segment depression on ECG, prior revascularization, or chronic lung disease.

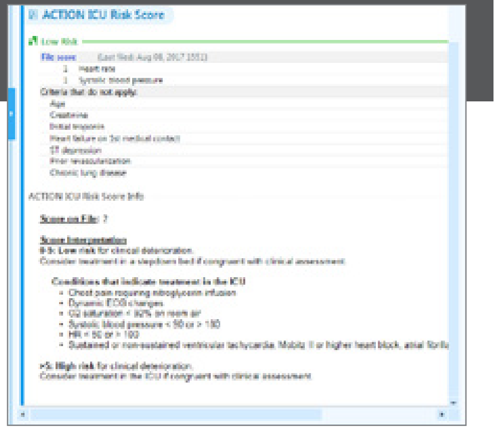

Figure 3: Display after risk calculation

The risk score calculator displays the score along with recommendations for location of admission based on the score.

NSTEMI, *continued*

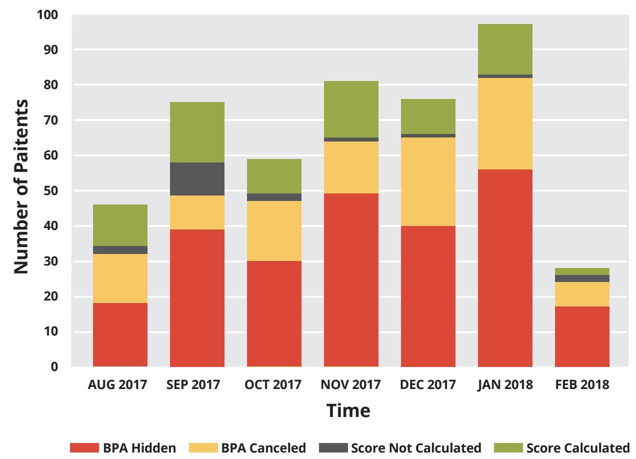

Figure 4: Actions when BPA was triggered

The BPA was triggered on 462 individual patient encounters. It was used properly (i.e., a score was calculated or a reason was given that a score was not calculated) 46% of the time. Number of triggers is lower in August 2017 and February 2018 because data collection began on August 14, 2017 and ended February 14, 2018.

RISK-BASED ICU ADMISSION FOR HEMODYNAMICALLY STABLE NSTEMI PATIENTS USING AN EHR-INTEGRATED RISK CALCULATOR COULD POTENTIALLY BENEFIT A NUMBER OF HOSPITALS IN THE DUKE-LIFEPOINT NETWORK AND BEYOND.

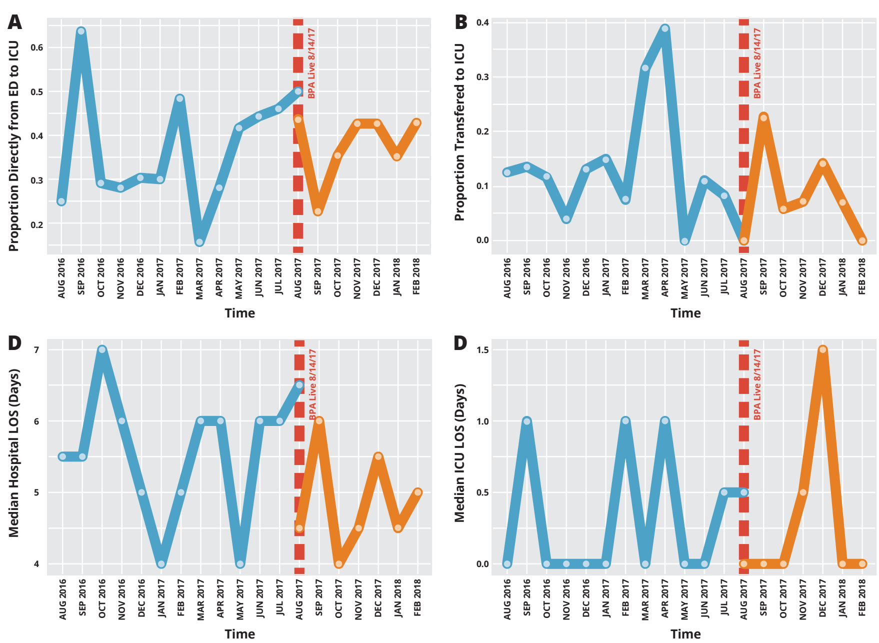

Figure 5: System-level outcomes before and after BPA roll-out

Panel A shows proportion of patients admitted from the ED to the ICU by month; panel B shows proportion of patients not originally admitted to the ICU that were transferred there by month; panel C shows median hospital length of stay by month; panel D shows median ICU length of stay by month. The BPA went live in August 2017, and outcomes are shown before and after the BPA went live. Blue lines represent the time period before BPA implementation; orange lines represent the time period after BPA implementation.

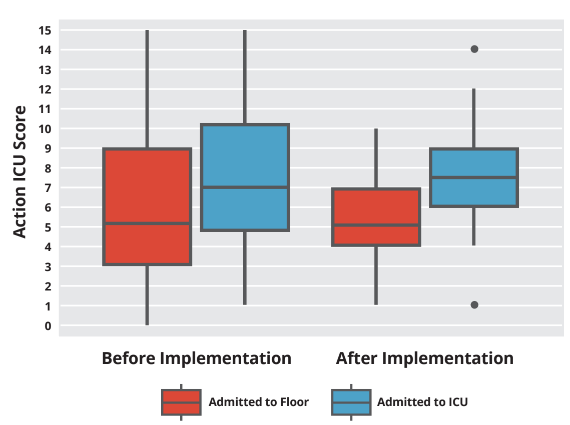

Figure 6: ACTION ICU score distribution of patients that were admitted to the ICU and not admitted to the ICU before and after BPA roll-out

Box and whisker plots show IQR and median ACTION ICU score for patients admitted to the ICU (blue) and patients not admitted to the ICU (red) before and after implementation of the BPA. Though the median scores of patients admitted versus not admitted to the ICU did not change substantially before and after implementation, there was less heterogeneity after implementation of the BPA.

"ACP @ DOC": A Pop Health Approach to Advance Care Planning in Primary Care

Project Team

Azalea Kim MD, MBA, MPA

Jared Lowe, MD

Larry Greenblatt, MD

Lynn Bowlby, MD

DIHI Team

Krista Whalen

PROBLEM

Primary care practices lack a pathway to meet the advance care planning needs of their patient populations. We defined advance care planning as the process of supporting our patients in the understanding and sharing their personal values, life goals, and preferences regarding future medical care..

SOLUTION

Our advance care planning (ACP) intervention had three main objectives:

1. Apply a predictive model to identify patients who might most benefit from ACP;
2. Complete ACP-dedicated appointments with non-physician Patient Navigators;
3. Use EMR tools to standardize ACP documentation and make data accessible across care settings

We piloted a population-health based pathway to provide ACP at the Duke Outpatient Clinic (DOC), and developed a model to identify and risk-stratify patients who would benefit from an ACP intervention.

IMPACT

Our data science model identified 480 patients appropriate for an ACP appointment. From July 1 to November 10, 2017, we completed outreach to 245 patients, scheduled 129 patients for ACP appointments, and completed 114 ACP appointments. 112 patients had completed ACP visit notes, and 103 patients had completed HCPOA forms (90%). Our social work-trained patient navigators provided bandwidth to engage patients through a new type of clinic encounter that can be reimbursed successfully through designated Medicare CPT codes for ACP.

Although we had some heavy lifting to do in workflow development and patient engagement, through our implementation we were able to focus our limited resources on higher-risk patients thanks to the data model, and better align care providers across settings with the help of the EHR tools. We found success in a care planning script that leveraged the relationship and trust between patient and PCP, and used language appropriate for patients with low health literacy.

ACADEMIC OUTPUT

We were honored to present our work at the following conferences and programs:

The Society for General Internal Medicine, Oral Plenary Presentation, April 2018

The North Carolina American College of Physicians, Poster Presentation, March 2018

CNCC Palliative Care Group, March 2018

Geriatric Workforce Enforcement Program (GWEP), September 2017

Improving Goals of Care Conversations

Project Team

Kathryn Pollak, PhD
David Casarett, MD
Brian Griffith, MD, MMCi

DIHI Team

Will ElLaissi, MBA, MHA

PROBLEM

Many patients who are near the end of life and facing difficult choices have not had a conversation with a health care provider to define their goals of care. Aggressive care might not match the patient's goals and values and can also have unnecessary costs for families and the health system. Even when goals of care are elicited, varied documentation practices may make it difficult for other providers to locate in the EHR, resulting in this valuable information not being known to future providers and not guiding the patient care.

SOLUTION

We developed and implemented a project to identify patients who may benefit from a goals of care discussion, and to help hospitalists better facilitate that conversation. To inform **when** to have the conversation, we designed an EHR alert based on triggering criteria determined with help from the hospital medicine program. To improve **how** hospitalists facilitate the conversation, the coach (Dr. Pollak) met with each hospitalist to educate on use of the "SUPER" script for goals of care conversations. The hospitalists then audio recorded goals of care conversations for identified patients on an encrypted iPod, which uploaded to HIPAA compliant cloud storage. After reviewing and coding on the audio transcriptions, the coach held feedback sessions with each hospitalist. Additionally, the hospitalists received education on documenting goals of care conversations in The Advance Care Planning Module in Maestro Care.

IMPACT

Hospitalists involved rated the intervention highly. 80% rated the intervention as "very helpful," that they had "made changes in their clinical practice," that the coaching would "have an impact on how effectively they communicate with patients," and that they would "definitely recommend to a colleague." Early analysis of the patients for whom an EHR alert was triggered provides some information about our study: approximately 12% of patients with an EHR alert died over the course of the study, suggesting the triggering conditions did correctly identify an ill cohort of patients at risk of near term mortality. 12% of those patients received palliative care consultation, with 9% being discharged to hospice care. Nearly half of the discharge summaries for these patients note that goals of care were addressed by the treatment team, but ultimately that conversation was documented in the Advance Care Planning note only a minority (n=15) of times. This suggests that the workflow of documenting these conversations in a distinct area of the chart rather than in conventional areas (progress notes, discharge summaries) was uncommon and that the training/ education and rationale provided for doing so was insufficient to change provider behavior.

SUPER:
Setup
Understand
Prognosis/Priorities
Emotion
Recommend/Review

HOSPITALIST FEEDBACK:

"The personalized assessment and review of my encounters with patients was most helpful. I learned what I was doing well and was given insight as to why the various techniques were effective."

"I found the goals of care discussion template as well as the personal feedback on my discussions very helpful."

PROMISE: Perioperative Risk Optimization with Machine Learning for an Improved Surgical Experience

Project Team

Sandhya Lagoo-Deenadayalan, MD, PhD
Mitch Heflin, MD
Shelley McDonald, DO, PhD
Madhav Swaminathan, MD, MBBS
Elizabeth Lorenzi
Annemarie Thompson, MD
Katherine Heller, PhD

DIHI Team

Kristin Corey
Mark Sendak, MD, MPP
Sehj Kashyap
Krista Whalen

PROBLEM

Two important problems exist in the current high volume perioperative clinical service:

1. Lack of a systemwide, streamlined automated process for rapid preoperative patient risk stratification and management by appropriate perioperative teams, resulting in inefficiencies like case cancellations/postponement.
2. Lack of timely identification and modification of surgical risk results in even more serious consequences for the patient and the health system: surgical/postsurgical complications, longer hospital stays, readmissions and general dissatisfaction among patients and family members.

SOLUTION

Using available surgical patient data, we developed a model for preoperative risk stratification, specifically the identification of key risk factors that predict complications and mortality after surgery. Using variables and risk strata identified in the model, we plan to prospectively pilot a risk stratification process at the point of surgical referral. This would provide clearer estimates of patients in need of optimization.

IMPACT

We developed an easy-to-use calculator that allows providers in surgery, anesthesia, medical specialties or primary care to evaluate the risk of complications for their patients. We plan to validate our analytical models using retrospective data from patients previously seen in the Perioperative Optimization of Senior Health (POSH), Preoperative Enhancement Team (POET) and Preoperative Anesthesia Testing (PAT). We will refine the predictive capacity of the tool by including variables collected in POSH clinic to the risk assessment model (e.g., function, falls, cognitive performance, gait speed, nutritional status, social and financial vulnerability) and evaluate the impact on outcomes.

ACADEMIC OUTPUT

PYTHIA: Automated Surgical Outcomes Data Pipeline and Prediction Engine.

Kristin Corey, Sehj Kashyap, Elizabeth Lorenzi, Krista Whalen, Mark Sendak, M.D., Mitchell Heflin, M.D., Shelley McDonald, D.O., Katherine Heller, Ph.D., Madhav Swaminathan, M.D., Sandhya Lagoo-Deenadayalan, M.D. Ph.D; Presented at the Machine Learning in Healthcare Conference (MLHC) August 2018

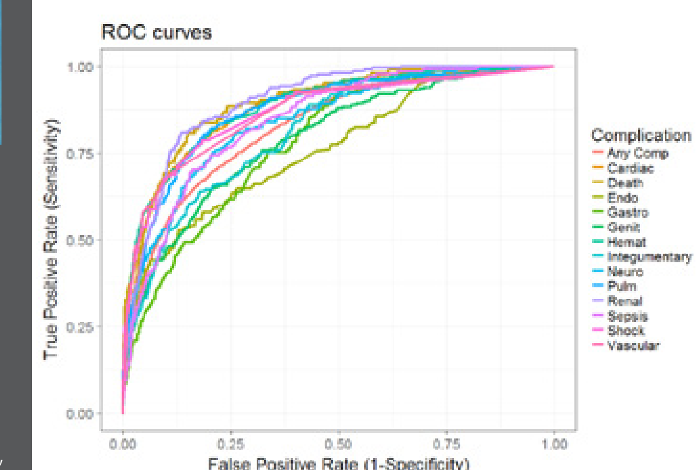

All models perform strongly with C-statistics (calculated on a held-out test set of 10,000 encounters) between 0.78-0.90 for the full model and 0.68-0.87 for the geriatric patient model (POSH).

DATA: Dashboard for Aggressive Treatment Analysis

Project Team

David Casarett, MD
Nrupen Bhavsar, PhD
Benjamin Goldstein, PhD
Matthew Phelan, MS
Yousuf Zafar, MD

DIHI Team

Krista Whalen
Will ElLaissi, MBA, MHA

PROBLEM

We were seeking to understand the causes of over-aggressive care at the end of life. We identified that most aggressive care was due to over admission of patients (as opposed to ED visits or chemotherapy treatments).

SOLUTION

This project involved an analysis of who received over-aggressive at the end of life. We took a mixed methods approach. First we interviewed 6 pancreatic oncology doctors about their definitions, perception of and reasons for over-aggressive end of life care. Next, we extracted EHR data on patient treatment during the last year of life. Initially, we considered this among only pancreatic cancer patients. Since the sample size was small we decided to expand this to breast, prostate, and lung cancer patients as well. One proposed solution is the creating of an acute care oncology clinic.

IMPACT

We identified characteristics of aggressive care at the end of life, at Duke, among cancer patients. We are considering expanding this analysis to other disease areas, such as Heart Failure and End Stage Renal Disease. We also plan to incorporate claims information to get a better understanding of receipt of care outside of DUHS. We are also engaged in a collaboration with Performance Services to replicate this project across DUHS (DATE: Dashboard of Aggressive Treatment at the End of life). That project uses the same methodology and leverages Death Masterfile data and North Carolina death certificates to create a 95% complete record of the end-of-life care of patients served by Duke Health. Key metrics in that dashboard will include use of inpatient care and ED visits in the last month of life, and 30-day mortality. That dashboard will be included in reports and action plans of all mortality stakeholder groups (oncology, cardiovascular, med-surg, neurosciences) and will be shared with department chairs and division chiefs.

WE IDENTIFIED CHARACTERISTICS OF AGGRESSIVE CARE AT THE END OF LIFE, AT DUKE, AMONG CANCER PATIENTS.

2017 DIHI RFA Project Summaries

ED LOS: Predictive analytics to reduce emergency department length of stay for youth with behavioral health disorders

Project Team

Gary Maslow, MD
Katherine Hobbs Knutson, MD
Nicole Heilbron, PhD
Nathan Copeland, MD
Kristen Stefureac, MSW
Kyla Machell, PhD
Spenser Radtke, LCSWA, MSW
Ben Goldstein, PhD
Clay Bordley, MD
Fred Johnson, MBA
Eva Mihm, MSW
Teka Dempson
Heather Copley, MSW
John Orr

Alliance Behavioral Health Team

Kate Peterson, MS
Katherine Benzaquen, MSW
Towanda Witherspoon, MA
Jennifer Meade, MPA

DIHI Team

Will ElLaissi, MBA, MHA
Krista Whalen

PROBLEM

Preliminary data suggested that emergency department (ED) length of stay (LOS) is longest for youth waiting for transfer to inpatient psychiatric units. Therefore, we developed a predictive model to identify individuals at risk for admission to an ED or inpatient psychiatric unit using Medicaid physical and behavioral healthcare claims from Alliance Behavioral Healthcare (ABH).

SOLUTION

We developed a predictive model using Medicaid health insurance claims to identify individuals at risk for admission to EDs or inpatient psychiatric units. For youth identified by the model, we intervened with care management to improve engagement in community-based behavioral health services, thus potentially reducing use of ED and inpatient psychiatric services.

IMPACT

In this project, 65% of cases that were reviewed by the DUHS-ABH multidisciplinary care management team were successfully engaged with community-based behavioral health services by the end of the project period. Specifically, over three months, the multidisciplinary team met for one hour weekly and reviewed clinical and social data for 31 at-risk youth identified by the predictive model. Following these multidisciplinary case reviews and outreach to families, approximately 65% of these at-risk youth were engaged in community-based behavioral health care (defined as having attended at least two visits with a single treatment provider). In some cases, there is continued outreach to families to support engagement in community-based services, thus potentially improving treatment engagement for this cohort. At the multidisciplinary team meetings, care management objectives were identified for each at-risk youth in review—objectives related to treatment referral, addressing SDOH needs, or other issues. In approximately 75% of cases, the primary care management objective was completed. The predictive model is run monthly at Alliance Behavioral Healthcare, and it identifies youth and adults who are at risk for using Emergency Department and inpatient psychiatric services. There are existing multidisciplinary rounding groups for adults at Duke, including the Familiar Faces program. Thus, by combining these existing efforts, we may further scale and disseminate this work.

Bridge-building between the Social and Health Sciences

Jessica Sperling, PhD

DR. SPERLING WORKS WITH DIHI TO DESIGN EVALUATION METHODOLOGIES THAT ARE INTEGRAL TO THE SUCCESS AND SCALING OF OUR INNOVATION PROJECTS.

Jessica Sperling, PhD, leads the Evaluation & Engagement area at the Social Science Research Institute (SSRI). Dr. Sperling works with DIHI to design evaluation methodologies that are integral to the success and scaling of our innovation projects. Through systematic and empirical investigation, Dr. Sperling designs evaluation and measurement processes that gauge implementation and/or progression towards intended outcomes, and uses this information to inform project development. Dr. Sperling has partnered with DIHI to provide measurement and evaluation expertise to several of DIHI's pilot innovation projects.

In addition, SSRI and DIHI recently collaborated to offer the new undergraduate Social Science Research Lab program—an SSRI-based curricular program that provides students applied experience in social science research methods. With DIHI's Will ElLaissi, Dr. Sperling developed and led the course, "Evaluating Health Innovation," which afforded students the opportunity to develop, pitch, and implement evaluation study designs with select DIHI innovation project partners. The partnership with SSRI and DIHI has provided student with diverse, unique educational opportunities to engage in real and relevant innovation projects, progressing efforts at Duke to meaningfully align the social and health sciences.

Anthony Lin

During my scholarship year, I was fortunate to get involved in a number of different projects. I worked with Dr. Cara O'Brien from the Division of General Internal Medicine to develop and deploy a machine learning model for early detection of sepsis. I led an analysis to better understand the clinical and operational implications of using different sepsis phenotypes in our health system sepsis redesign work. Our findings identified a specific patient population that could stand to benefit from earlier sepsis intervention and our model was designed to target that phenotype. I also helped design and plan for implementation of this sepsis early warning system at Duke University Hospital. I worked with key stakeholders in Duke Health leadership, the Rapid Response Team, and the Emergency Department to integrate this technology into clinical workflow and identify the actions that needed to be taken when a patient is identified as high risk.

Alongside deployment of the sepsis early warning system, I also helped develop and evaluate a new reinforcement learning algorithm to make personalized treatment recommendations for patients with sepsis. Our findings suggest that prescriptive actions recommended by our model may have improved care and open up a discussion about the role that reinforcement learning may one day be able to play in healthcare. Our study highlights the need for further collaborations, both technical and clinical, to thoughtfully incorporate new prescriptive analytic models into clinical practice.

Lastly, I helped develop novel data infrastructure for health data that promises to improve data quality, access, and timeliness for health systems and investigators seeking to derive more meaningful insights from clinical care data. Our team is now working with Duke Health leadership to leverage this resource within our health system to power research, operations, learning health, and medical education.

My time at DIHI has shown me the potential for clinical informatics and data analytics to help improve our diagnostic and prognostic power. The experience has given me an appreciation for not only the challenging process of acquiring health data, but also the difficulty in thoughtfully implementing data-driven insights to improve patient care. I have loved working with such a dynamic, effective, and caring team, and aspire to one day build my own teams with people of such caliber and character. DIHI's vision for healthcare, focus on rapid innovation, and dogged commitment to "doing what needs to be done" make it a truly unique organization within Duke and inspire students like me to continue playing our part to push open the bounds on healthcare.

"My time at DIHI has given me an appreciation for not only the challenging process of acquiring health data, but also the difficulty in thoughtfully implementing data-driven insights to improve patient care."

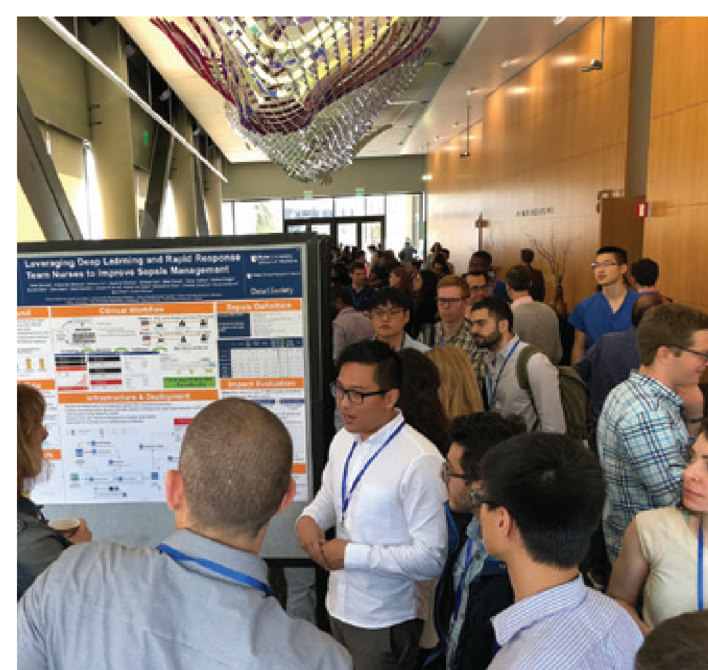

Machine Learning in Healthcare

A Critical Appraisal and Opportunities

Mark Sendak, MD, Michael Gao, Marshall Nichols,
Anthony Lin, Suresh Balu, MBA

Despite excitement surrounding machine learning in healthcare, health systems that fully integrate machine learning models into clinical care operations are the exception rather than the rule. Bringing machine learning models from the blackboard to care at the bedside requires intense transdisciplinary collaboration, alignment of goals, and capabilities that are hard to find in healthcare today. At the Duke Institute for Health Innovation (DIHI), we are in our fourth year developing, piloting, and implementing machine learning technologies in clinical care. To benefit from the full potential of machine learning in healthcare, we must step back from the trenches to systematically acknowledge breakthroughs in technology and adoption, address barriers to progress, and critically reflect on the strategic priorities necessary to bring healthcare into a new digital age.

AT THE DUKE INSTITUTE FOR HEALTH INNOVATION (DIHI), WE ARE IN OUR FOURTH YEAR DEVELOPING, PILOTING, AND IMPLEMENTING MACHINE LEARNING TECHNOLOGIES IN CLINICAL CARE.

The Good

In the last year, machine learning methods were prominently featured in mainstream medical literature. JAMA alone has presented three deep learning models that classified images of retinopathy and breast cancer metastases at a level equal to or better than clinical experts.[1-3] Using tens of millions of data points from our electronic health record (EHR), our transdisciplinary team developed a deep learning model to predict onset of sepsis.[4] Well-developed models demonstrate diagnostic acumen that surpasses human capabilities and do so at scale.

Although small in number, there are emerging uses of machine learning in healthcare operations. For example, Epic's cognitive computing platform will support machine learning models starting in 2018; the Food and Drug Administration has approved software to assist with medical imaging segmentation; and patient deterioration models are commonly built into EHRs. The adoption of such models and technologies serves as a foundation for machine learning to diffuse across institutions into clinical care operations.

The Bad

Historically, statistical models in healthcare found patterns in data that enhanced clinical reasoning. This expectation is often applied to machine learning models, but machine learning and clinical reasoning are not always coupled. Clinical reasoning is often cultivated across institutions, while machine learning models are often developed using data from a single institution and have limited generalizability. For example, a *Clostridium difficile* model tested at two AMCs revealed variables that were top risk factors in one setting and protective in the other.[5] Clinical care processes that generate and capture data vary widely across institutions and local biases are baked into machine learning models. However, even if a model cannot enhance clinical reasoning, it can still augment workflow-specific decisions at a local level.

If health system leaders want to test a newly validated machine learning model in their local environment, they must prepare for significant investment in personnel and technology. Culling through raw healthcare data to construct model features is expensive and time-intensive. At our institution, the cost of developing, validating, and integrating a single analytics tool to identify patients at high-risk of dialysis was $220,000.[6] At a national level, the cost for physician practices to abstract, normalize, and report on quality measures captured in the EHR is $15.4 billion.[7] Resource requirements prevent even the most generalizable model from efficiently scaling across institutions. When Kaiser Permanente, the nation's largest integrated health system, launched their EHR in 2003, they mandated interoperability and a common data model across regions. This enabled rapid diffusion of technologies across Kaiser's regions, but not every health system prioritized data standardization over customization.

Almost all research at the intersection of machine learning and healthcare is performed on remotely collected, stale data without appropriate domain expertise. During 2015 – 2017, the *Journal of Machine Learning Research* had three issues dedicated to healthcare, including a special feature and two proceedings for the "Machine Learning in Healthcare Conference". Of 40 publications, 23 (57.5%) had a clinical collaborator, 10 (25%) used non-Medical Information Mart for Intensive Care (MIMIC) EHR data, and only seven (17.5%) had both a clinical collaborator and used locally collected, non-MIMIC EHR data. Three of the seven papers were projects our group worked on and all seven were from AMCs. Without engaging partners across domains to solve relevant problems, machine learning will continue to struggle with adoption by both clinicians and health information technology leaders.

The Ugly

Personalized medicine will require mass customization of models that are trained and re-calibrated at the hospital- and cohort-level. Modern machine learning techniques focus on generalization beyond a training dataset, not on generalization to different sites. Transfer learning methods require further development to help address this problem and in the meantime generalization must be achieved through localization. This will require either the skills to recreate datasets and retrain models at every site or a willingness to leverage capabilities from outside institutions.

Methods for evaluating and monitoring models to ensure continued accuracy and performance are in their infancy. The underlying data structure of EHRs is highly dynamic and can result in errors when models are evaluated. Machine learning models and infrastructure need to account for these changes so that their results are robust to the underlying conditions. In addition, although machine learning has developed methods for model validation such as training-test-validation splits and k-fold cross-validation, further validation of a model post-implementation requires new techniques. Consider the case of a machine learning model used to predict the onset of sepsis. If action taken as a result of the model prevents infection, the counterfactual to this event is not observed and there is no clear way to classify the event as a false positive or a successful intervention. Compounding this issue, if the model is retrained at a later time using data from the post-imple- ▶

mentation period, the results can be biased in ways that are difficult to ascertain. New methods and technology infrastructure must be developed to address these complex issues.

The Opportunity

Institutions that are interested in embedding machine learning into clinical care operations must coalesce a workforce with new competencies, harness transdisciplinary resources, and invest in platforms to support machine learning. In 2015, Thomas Davenport and Julia Kirby characterized five ways knowledge workers can respond to automation.[8] Healthcare providers can *step up* (consider the big picture of the industry), *step aside* (develop strengths that aren't codifiable cognition), *step in* (modify and monitor software), *step narrowly* (specialize in something for which no computer program has yet been developed), or *step forward* (build the next generation of technology). At worst, financial and cultural pressures drive clinicians to *step narrowly* to specialize and hide from technology. At best, informatics and statistics training drive clinicians to *step in* to modify and monitor software. If clinicians are to *step forward,* health system leaders must invest in programs that empower the clinical workforce to develop next-generation technologies. This requires health system leaders to shift from viewing development as an expense to viewing development as an investment in future growth.

AMCs that align and value transdisciplinary collaboration and training are ideally suited to embed machine learning in clinical care. To cultivate this type of ecosystem at our institution, we embed statistics and computer science students with medical students on teams led by clinical and quantitative science experts. In two years, we trained 40 students studying statistics and computer science as well as eight medical student research scholars. Through this process, we have coupled career development with successful pilot implementations to improve clinical care.

The time has come to refocus our attention and energy from siloed applications of machine learning in healthcare to the underlying platforms required to efficiently scale machine learning across healthcare. In our local setting and as an industry, we have witnessed a critical mass of successful projects and collaborations. Now is the opportunity to reflect on learnings and expose and break down barriers to build platforms that support many machine learning applications. We must redefine success from optimizing performance metrics of a single model to optimizing scalable growth in the number of high-impact collaborations between clinical researchers and machine learning experts.

> IN OUR LOCAL SETTING AND AS AN INDUSTRY, WE HAVE WITNESSED A CRITICAL MASS OF SUCCESSFUL PROJECTS AND COLLABORATIONS.

[1] Ting DSW, Cheung CY-L, Lim G, et al. Development and Validation of a Deep Learning System for Diabetic Retinopathy and Related Eye Diseases Using Retinal Images From Multiethnic Populations With Diabetes. *JAMA.* 2017;318(22):2211-2213. doi:10.1001/jama.2017.18152.

[2] Ehteshami Bejnordi B, Veta M, Johannes van Diest P, et al. Diagnostic Assessment of Deep Learning Algorithms for Detection of Lymph Node Metastases in Women With Breast Cancer. *JAMA.* 2017;318(22):2199–12. doi:10.1001/jama.2017.14585.

[3] Gulshan V, Peng L, Coram M, et al. Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs. *JAMA.* 2016;316(22):2402-2410. doi:10.1001/jama.2016.17216.

[4] Futoma J, Hariharan S, Heller K, et al. An Improved Multi-Output Gaussian Process RNN with Real-Time Validation for Early Sepsis Detection. *Proceedings of Machine Learning for Healthcare.* August 2017:1-12.

[5] Makar M, Oh J, Fusco C, et al. A Data-Driven Approach to Predict Daily Risk of Clostridium Difficile Infection at Two Large Academic Health Centers. *Infectious Diseases Society of America.* October 2017.

[6] Sendak MP, Balu S, Schulman KA. *Barriers to Achieving Economies of Scale in Analysis of EHR Data: a Cautionary Tale.* Vol 8. 2017:826-831. doi:10.4338.

[7] Casalino LP, Gans D, Weber R, et al. US Physician Practices Spend More Than $15.4 Billion Annually To Report Quality Measures. *Health Aff (Millwood).* 2016;35(3):401-406. doi:10.1377/hlthaff.2015.1258.

[8] Davenport T, Kirby J. Beyond Automation: Strategies for Remaining Gainfully Employed in an Era of Very Smart Machines. *Harvard business review.* May 2015:58-65.

Clustering Data

Aman Kansal and Sarah Scharber

Making a Case for "Why Cluster?"

Patients' analyte values can be useful predictors, or features, in machine learning and statistical modeling. Automated models, while powerful, are sensitive to input and input not properly labeled can result in observations and predictions that are misinformed. Our group first encountered this problem when working on the Chronic Kidney Disease Project. When trying to follow creatinine, for example, we found over 15 distinct test names corresponding to creatinine, with no way of grouping names together into more meaningful categories. We again encountered this problem in our Sepsis Watch Project. From July to September 2014, the component name "report" corresponded to the order description "culture, blood". Unfortunately, physicians had difficulty distinguishing "report" in a patient's chart and often reordered the blood culture, resulting in a spike in orders for those three months. After September, the component name changed to "culture blood (bkr)", which was more easily found. Not having those two component names clustered under a standardized "common name" resulted in danger to patient safety and also incomplete input to our model, leading to incorrect predictions.

The Dynamic Data Wrinkle

When trying to come up with a solution, we realized even if analytes are standardized at a single point, there are no guarantees they will not change. So, could we develop a way to automate the clustering process? In order to solve this problem, we used analyte data derived from October 2014 until October 2017. We included pertinent raw fields, such as "component name", "reference unit", and "value", in order to develop an algorithm to group similar analytes under the standard common names. At first, we tried using supervised and unsupervised algorithms. However, there were not enough informative features to generate a reliable supervised learning model, and unsupervised learning models did not reach a threshold for guaranteed accuracy required for clinical decisions. We then explored the Bhattacharya distance, a method to compare probability distributions. We compared analyte distributions to a gold standard analyte for each common name and ranked analytes from least to most distance to develop a recommendation engine to be used in conjunction with physician curation. We tested our algorithm with some success on 12 common names in particular: bicarbonate, creatinine, glucose, hematocrit, lactate, magnesium, potassium, PCO2, platelets, PO2, sodium, troponin, and white blood cell.

A Path Forward

Clustering helps make sense of, and trend, valuable predictors for machine learning models. We have proposed a way to retroactively group analytes into more meaningful data. Furthermore, we propose a prototype workflow for prospectively grouping analytes: 1) separate clustering into two levels of grouping – the highest level corresponding to the common name and the more granular level corresponding to categories within the common name (e.g., outpatient vs. inpatient vs. OR, percentage vs. raw); 2) monitor new analyte names; 3) assign new analytes to appropriate existing or new common names. Establishing a strong foundation to better understand and group analytes is crucial towards future success.

DELPHI

Duke Environment for Learning and Promoting Health Innovation

Anthony Lin

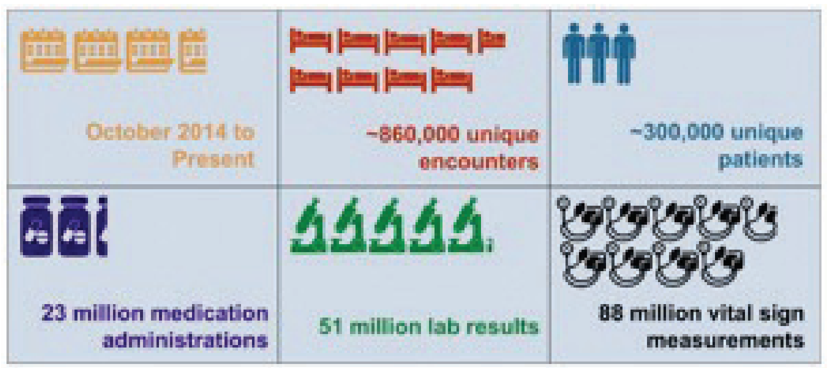

The Opportunity in Health Data

Since its implementation in 2013, Epic has enabled the creation of a comprehensive health record for all patient encounters at Duke Health to aid clinicians in health-care delivery. However, given the 125,000 different data fields used to store information and the inherent complexity of Epic's data backend, the health system's ability to derive meaningful in-sights from clinical care data has become bottlenecked. Our ability to use data-driven insights to in-form academic, operational, and learning health system strategy becomes limited by the technical barrier to access, clean, and vali-date health system data.

Our Experience

The Duke Institute for Heath Innovation (DIHI) has curated large datasets for more than a dozen data science projects over the past five years. These datasets have enabled predictive modelling of chronic kidney disease, first hospi-tal admissions, and sepsis, as well as informed quality improvement of new care delivery models and innovation pilots. Over the course of gathering and validating these datasets with clinical experts, DIHI has developed a suite of exten-sible and reproducible tools to rapidly curate datasets.

A Culmination of Effort

DELPHI (Duke Environment for Learning and Promoting Health Innovation) is the culmination of years of work in understand-ing how to leverage Duke clini-cal operations data to support investigators and health system leaders hoping to drive change in healthcare delivery. This data asset facilitates the rapid explo-ration and analysis of validated clinical data in a large, diverse, and comprehensive inpatient population at Duke University Hospital. DELPHI extracts clinical care data from our EHR relational reporting database and cleans, normalizes, and standardizes the data elements. We've worked with clinical domain experts to validate the data elements and populate them with meaningful metadata to enable grouping of high-level features and facilitate custom disease phenotyping. To-date, DELPHI contains mil-lions of data points ranging from encounter characteristics, patient demographics, and transfer times to laboratory results, vital signs, and medication administrations.

DELPHI's breadth of curated features, speed of access, and complete transparency of data curation processes enable our healthcare community to leverage data-driven insights in a matter previously unrealized. It great-ly reduces the time to procure meaningfully curated health data and empowers investigators and health system leaders with a clear-er understanding of the clinical care they seek to improve.

Pythia

Open Science Data Platform for Surgical Innovation

Kristin Corey

Insights from Aggregated Surgical Data

Pythia sprouted from a data science project on surgical out-comes. Specifically, the project began as a post-operative patient outcome predictor for high-risk geriatric patients. However, given the creative freedom to explore endless possibilities with coding and access to data, my colleague, Sehj Kashyap, and I found our-selves saying "let's see if we can go bigger". And so we went bigger. With mentored guidance from Mark Sendak, MD and Suresh Balu, in six months we were able to complete an initial version of a surgical data mart, housing over 145,000 invasive surgeries on over 90,000 patients since 2013. The features of our data repository include curated clinical features, such as comorbidities, compli-cations, CPT codes, outpatient medications, demographics, as well as inpatient encounter infor-mation, such as length of stay and mortality. Given the data science medium and the concurrent work of our colleagues at DIHI (see "DELPHI"), we decided to name this resource "Pythia". Working with Liz Lorenzi, a PhD student in the Department of Statistical Sciences, we created a machine learning model that predicted risk of post-operative complications not only for geriatric patients, but for all of Duke's surgical patients. Our team submitted the work for publication and were honored to present at the annual Machine Learning in Healthcare Confer-ence at Stanford University in August.

Scaling Up

After just one year "let's see if we can go bigger" was said yet again and over the summer we were able to re-write Pythia's source code. Our data repository now incorporates all procedures at DUHS (medical and surgical) from all three hospitals. There are almost a million procedures, making up over 614,000 surgeries on more than 311,000 patients since May 2013. With the primary aim to create an open science research platform to enhance our DukeHealth community research and innovation, our team's work has been significantly recognized. We are now fully supported by the Department of Surgery, with Dr. Alan Kirk, the department's chairman, as our PI. With a whole new year ahead of us, who knows where we'll be in twelve months.

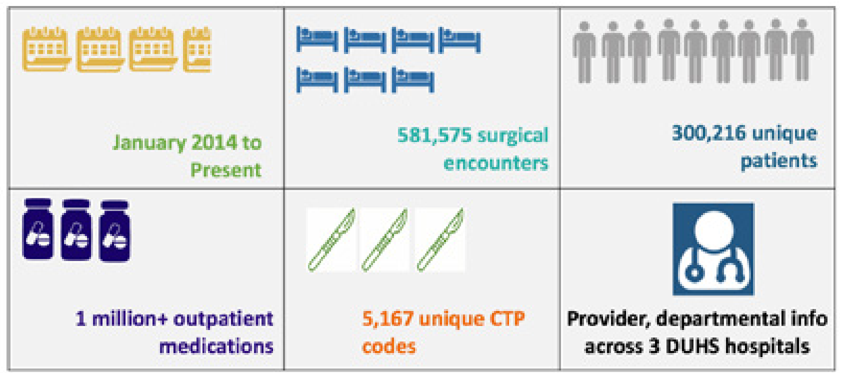

Implementation of

Deep Learning Technologies for Sepsis Management

COVER STORY

Anthony Lin

The Opportunity

The widespread adoption of electronic health records (EHRs) has enabled health systems to harness insights derived from clinical data like never before. With new availability of health data, clinical decision support tools based on predictive analytics have become popular means by which to improve diagnostic accuracy.

At the Duke Institute for Health Innovation, we worked with Dr. Cara O'Brien and her team to understand how we could leverage our EHR to identify patients at risk of sepsis before they might present clinically. Within a year and a half, we worked with our local technology solutions team and statistical partners to pull retrospective health data and develop a sepsis prediction model that outperforms traditional clinical risk scores and other standard machine learning techniques. Our analysis showed that our early warning system could predict sepsis a median time of 5 hours before clinical presentation and, given the high morbidity and mortality of sepsis, had the potential to save 8 lives a month.

However, the question still remained: how do you develop an analytics-driven clinical workflow to bring the full potential of such a technology to realization? Our approach was three-pronged: 1) evaluate our prediction model for its ability to provide clinicians with prospective and actionable insights, 2) engage key stakeholders in our effector arm to ensure seamless integration of this technology with clinical practice, and 3) build capacity and infrastructure across Duke Health to promote sustainable growth in the field of predictive health analytics.

Deriving Actionable Insights

Data-driven insights are useless unless they provide relevant and actionable information to clinicians. A prediction engine that warns of deterioration 72 hours in advance for a patient who looks completely healthy creates confusion in care delivery and raises a medical and ethical conundrum. To avoid such a predicament, we trained our model to predict sepsis within a more actionable time window of 12 hours to drive clinical decisions for all patients captured by our early warning system.

End-to-End Integration with Clinical Practice

After refinement of our model, we needed to ensure that we constructed a proper workflow to rapidly triage, reassess, and treat sepsis for all patients in our pilot site. We estimated caseload across different sites in the hospital and ultimately decided to partner with the Emergency Department (ED) as our initial site for deployment. We trained a centralized team of Rapid Response Team (RRT) nurses, with previous experience in managing clinical ▶

HOW DO YOU DEVELOP AN ANALYTICS-DRIVEN CLINICAL WORKFLOW TO BRING THE FULL POTENTIAL OF SUCH A TECHNOLOGY TO REALIZATION? OUR APPROACH WAS THREE-PRONGED.

"I have now had the pleasure of working with DIHI for over two years on our sepsis pilot. We have made significant strides because DIHI provided the necessary support to bring our idea from concept to reality. DIHI brings together front-line clinicians, data scientists, computer scientists, and project managers to design high impact innovative projects. DIHI's multidisciplinary approach brings everyone's individual expertise to the table. When a problem or situation arises, there is always someone who has a solution. By creating an atmosphere for collaboration, creativity, and productivity for innovators from all disciplines, DIHI has brought us together to improve sepsis care in Duke Health."

– Armando Bedoya, MD, MMCi
Fellow, Pulmonary and Critical Care

deterioration, to evaluate model output on a custom-developed risk stratification tool. Lastly, we worked with key stakeholders in Duke Health leadership, RRT, and the ED to develop the touchpoints between clinician and machine and establish a protocol for elevating clinical concern for at-risk patients.

Building Capacity for the Future

Our sepsis program will be one of Duke Health's first clinical workflows to leverage machine learning to prospectively identify and treat disease. To facilitate an environment in which superior analytic techniques are regularly incorporated to inform clinical practice, we as a health system must develop capacity in infrastructure and people. Our team has been working with clinical and technical leaders to develop the governance and monitoring of such analytics-driven clinical workflows. Only through a better understanding of the strengths and limitations of such analytic approaches, thoughtful integration into clinical practice, and a clear strategy for monitoring and maintaining these systems can we develop a novel health system framework to leverage these technologies to improve our patient care.

OUR SEPSIS PROGRAM WILL BE ONE OF DUKE HEALTH'S FIRST CLINICAL WORKFLOWS TO LEVERAGE MACHINE LEARNING TO PROSPECTIVELY IDENTIFY AND TREAT DISEASE.

The sepsis web dashboard (pilot phase) features trended modular output along with key vitals, labs, and status of relevant treatment bundle components.

Sarah Scharber

My name is Sarah Scharber and I worked as a DIHI scholar for part of my third-year research experience. During my time at DIHI, I was introduced to a whole new field: data science. There was a very steep learning curve but my skills in programming and statistics grew exponentially over the first few months. Throughout the year I learned a great deal about applying these skills to innovation and the improvement of healthcare delivery. I worked on a laboratory test appropriateness project to help ensure that Duke is providing the highest quality care in the most cost-effective way. Hopefully I can continue to utilize all I have learned and make meaningful change as I progress in my future career as a physician.

"I worked on a laboratory test appropriateness project to help ensure that Duke is providing the highest quality care in the most cost-effective way."

Aman Kansal

This year at DIHI gave me the rare opportunity to step completely outside my comfort zone and firmly into the intersection of medicine and data science. I have learned technical skills including R, SQL, and Python and have developed a strong understanding in data reading, manipulation, visualization, and machine learning model creation. I have worked with experts in multiple domains tackling projects such as incorporating diagnosis codes in prediction models, developing an algorithm to cluster similar labs, and evaluating the use and safety of a new triage system. Most importantly, through the incredible mentorship and vision pervasive at DIHI, I have been able to embrace our core tenets of curiosity, bold humility, adaptability, resilience, and Ubuntu—the spirit of "I am because we are." I am so grateful to be a part of this team and am excited for what the future holds!

"This year at DIHI gave me the rare opportunity to step completely outside my comfort zone and firmly into the intersection of medicine and data science."

Digital Health

DIHI's mobile app Highlights

Mike Revoir and Jamie Daniel

DIHI's Mobile Apps

As health systems continue to innovate, opportunities for technology-enabled health-care access and insight for consumers and providers will continue to expand. Now more than ever, digital functionality has significant implications for improving quality outcomes by providing infrastructure for patient-centered, patient-driven care.

One Thing Straight

One Thing Straight was developed as a posture coach for individuals with Parkinson's Disease in partnership with the LiveWell Rehabilitation Engineering Research Center at Duke University. The user's posture is determined via a wearable sensor which communicates with the user's mobile device using Bluetooth. When bad posture is determined, a discrete notification is sent to the user indicating for them to correct their posture.

One Thing Straight won 1st place at the 2018 LiveWell Student App Challenge.

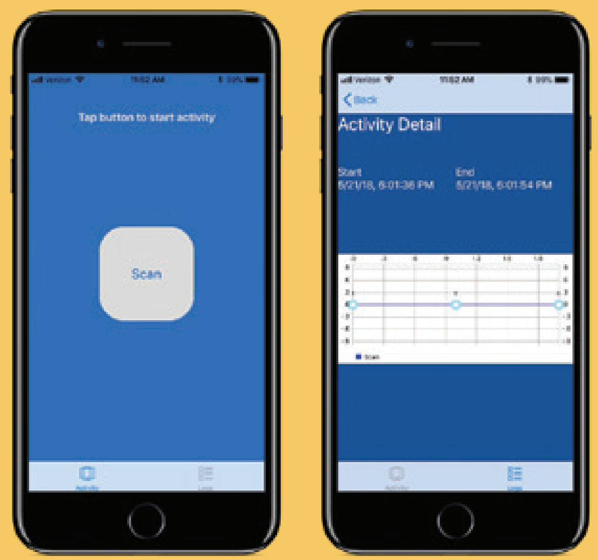

Ambulatory QR mobile app

The idea of a mobile application to assist patients in walking after a surgical procedure was brought to DIHI by 2nd year med students Dylan Eiger and Joshua Helmkamp. We created the app to scan QR codes that would be placed on the hospital hallways that would allow a patient to scan the code thereby measuring the walk waypoints. The main goals of the app are:

- To help physicians quantify the amount of walking a patient has performed
- To identify the speed and distance that the patient has walked
- Compare the previous walks to the current walk

All of this data helps to motivate and empower the patient while giving the care team insight on their recovery while still in an inpatient setting.

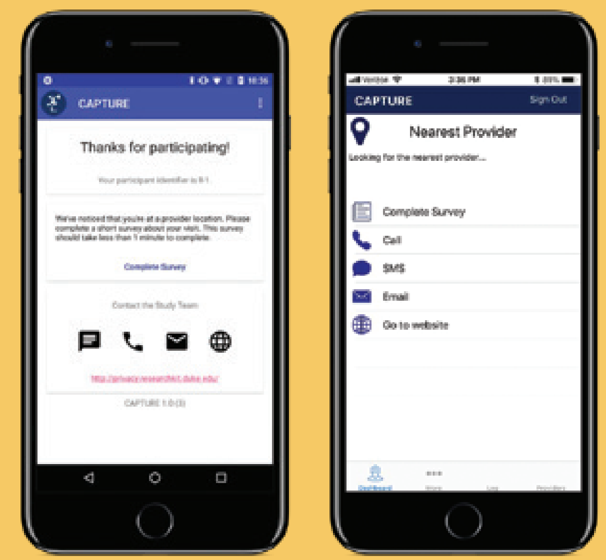

CAPTURE Mobile App

CAPTURE was developed in partnership with the Duke Clinical Research Institute to identify when patients have been hospitalized utilizing readily available geo-mapping technology on mobile devices. CAPTURE will be used to capture clinical events, typically hospitalizations, which may not be captured through traditional trial mechanisms. By establishing geofences around nearby medical facilities, when the patient has been at the hospital for a predetermined time, a notification is sent to the study coordinator who will then be able to follow-up with the patient / caregiver to identify whether a hospitalization has occurred. As patients travel, the list of nearby, geofenced medical facilities is updated.

DukeCare

DukeCare was developed for the Duke Colon and Rectal Surgery Clinic to provide colorectal ERAS patients with educational materials and tools related to their surgery. Instructional videos tailored to their procedure provided patients with information about their upcoming surgery, what to expect afterwards and other related topics. Patients were also able to keep an Activities of Daily Living diary, visualize activity level insights and contact the clinic.

The Importance of

Monitoring Data

Michael Gao

Data Can Help Us Solve Healthcare Challenges...

As a core member of DIHI's Data Science team, I have been both fortunate and unfortunate enough to have worked with electronic health record (EHR) data. Fortunate in that the potential for EHR data to revolutionize the way that healthcare is delivered and optimized is truly exciting. Unfortunately, in their current state, EHRs are messy and take an enormous effort to extract value. But, the rewards for these efforts have the potential to meaningfully impact today's patient care challenges. Specifically, our team has coupled EHR data with machine learning models to try to predict readmissions, recommend palliative care consults, and detect sepsis in the hospital before it occurs. And this is just the tip of the iceberg; we are working to expand our capabilities each and every year.

At their core, each of these machine learning models works by finding patterns in the data. The machine learning algorithms we employ are able to sift through a vast quantity of Duke patient data to find signs that may help predict future events. Although these methods are immensely powerful, they are inherently limited by the quality of the data. The saying "garbage in, garbage out" holds true in this setting. If you develop models with messy data, you'll get messy results. This means that the integrity of our EHR data and the way that we structure it directly affects the impact that our sophisticated technological approaches have on informing clinical care.

...But Only If We're Willing to Work With It

In addition to the problems of missingness (the manner in which data are missing from a population sample), lack of structure, convoluted relationships, and lack of standard entry mechanisms, the EHR is a dynamic and living system. Every week, data gets entered into our EHR that does not conform to anything that we have seen previously. Perhaps a new medication has just been put into practice, or maybe a new lab test has been ordered for the first time. However, due to the unstructured nature of the data, even a new dosage for an existing medication or a new name for the same laboratory test can look different once it gets entered into the record. How, then, are we supposed to develop algorithms that are robust to these changes? The answer, we believe, lies in monitoring the data.

To use a concrete example, let's say that one data element that is in use in one of our models is a blood culture. The first step in creating the machine learning model is to combine all of the different ways that blood cultures are represented in the system so that the next time the model sees data on a blood culture, the model knows how to identify it. However, what if the name of a blood culture changes in the system? In fact, we can tell through retrospective analysis that this exact scenario has happened, and if it were to happen in our sepsis model, it may drastically affect its performance.

Learning and Progression

Here at DIHI, we believe that in order to bring cutting-edge technology research into healthcare, we have to borrow best practices from other industries. The problem of monitoring data is not something new—many other industries have been tackling this for years. One of these industries which we have used as inspiration is quantitative trading. At its core, quantitative trading is all about trying to predict changes in time series data, such as a stock price of a company over time. If you picture the quantity of laboratory tests and medications ordered over all of Duke Health on a daily, weekly, and monthly basis, you might see that the resulting graphs look exactly like stock prices. There are some days where we might order more of one medication over another, and intra-weekly trends are easily noticeable. In addition, we might see trends that operate over larger time scales; maybe the institution makes a decision to consolidate laboratory providers and the names of laboratory values provided by that provider spikes suddenly. Using quantitative trading, signal processing, and other algorithms, our goal is to detect these changes as they occur and to take appropriate action. If we can catch when these anomalies occur, and get to the root of the problem, we can make sure that our models are robust to inherent dynamics of our EHR.

As with everything we do, we don't want to just stop there. Our data science team is constantly thinking of new ways to make sure that we are being responsible with new technologies that are making their way into the care delivery setting. Monitoring trends in data is just one of the many methods we employ to ensure the success of our pilot programs, many of which involve making predictions that can ultimately save the lives of our patients. The mission for the Duke Institute for Health Innovation has always been to catalyze change within Duke Health. I believe that the problems that our team have tackled and continue to tackle will ultimately lay the groundwork for a health system that is able to leverage the promise of machine learning, artificial intelligence, and all manners of cutting edge technology to deliver better quality and more efficient care to our patients.

```python
import numpy as np
import pandas as pd
from datetime import datetime
import tensorflow as tf
from tensorflow.python.framework import ops
from time import time
from datetime import timedelta

# Data cleaning functions
def process_encounter_baselinecovs(enc_dat):
    """
    Takes in db entry from the encounter table,
    returns a vector of baseline covariates with
    the demographics and status variables filled in:
        Age, Wt, Transfer, Emerg, Urgent, Gender, Race, Prior Sepsis
    Afterwards, will modify this covariate vector with any comorbiditie
    """
    start_time = enc_dat.iloc[0]['start_time'] #Use as t=0 for later
    cov = np.zeros(n_baseline_covs)
    dob = datetime.combine(enc_dat['birth_date'][0],datetime.min.time(
    if dob.year>2000:
        dob = dob.replace(year=dob.year-100)
    age = (start_time-dob).total_seconds()/60/60/24/365.24
    cov[COV_DICT_MODEL['Age']] = (np.log(age)-LOGAGE_MEAN)/LOGAGE_STD
    cov[COV_DICT_MODEL['Is_Male']] = 1-enc_dat['sex'][0]
    cov[COV_DICT_MODEL['Is_Not_White']] = float(enc_dat['race']!='CAUCA
    sep = enc_dat.sepsis_tally[0]
    if sep is None: sep = 0
    if sep==1:
        cov[COV_DICT_MODEL['Prior_Sepsis_1']] = 1
    if sep>=2:
        cov[COV_DICT_MODEL['Prior_Sepsis_2+']] = 1
    wt = enc_dat['weight'][0]
    if wt is None: wt = np.nan
    if np.isnan(wt):
        wt = 149.495-1.936*cov[COV_DICT_MODEL['Is_Not_White']]-0.0627*
    if wt < 60:
        wt = wt*2.206
    cov[COV_DICT_MODEL['Weight']] = (np.log(wt)-LOGWT_MEAN)/LOGWT_STD
```

DIHI Innovation Scholar

Sehj Kashyap

I joined DIHI in search of real-world experience handling electronic medical record healthcare datasets and improving my data science and programming skills.

Kristin Corey, my fellow medical student classmate, and I were the two lead data scientists and innovation scholars on the PROMISE project. We created the data repository called Pythia of over 150,000 surgical patients and their post-surgical outcomes. Then, we worked with a PhD statistician to create predictive models for post-surgical complications. Finally, I created a web applet that could be used as a calculator to predict post-surgical complications from 19 patient pre-surgical variables. My coding in SQL, R and python, analysis and documenting practices improved significantly as a result of working on the project and within our interdisciplinary team.

Ultimately, we presented Pythia as a surgical data resource to surgical department leadership, including Dr. Allan Kirk, under whom we have been working on applications of Pythia for outcomes and quality improvement projects. These experiences have completed a loop between ideation, research and implementation and have taught me how this cycle works in a learning healthcare system.

Throughout my fellowship, I was living in India. Despite this, DIHI allowed me to work on projects remotely and part-time. The flexibility was only one way in which I felt empowered by DIHI—Mark, Suresh, and others continually shared resources, advice and mentorship that helped me learn quickly. I would recommend the fellowship to anyone who wants an intimate understanding of how to scale innovation in a learning health system or strengthen their technical skills.

"I would recommend the fellowship to anyone who wants an intimate understanding of how to scale innovation in a learning health system or strengthen their technical skills."

DIHI Team

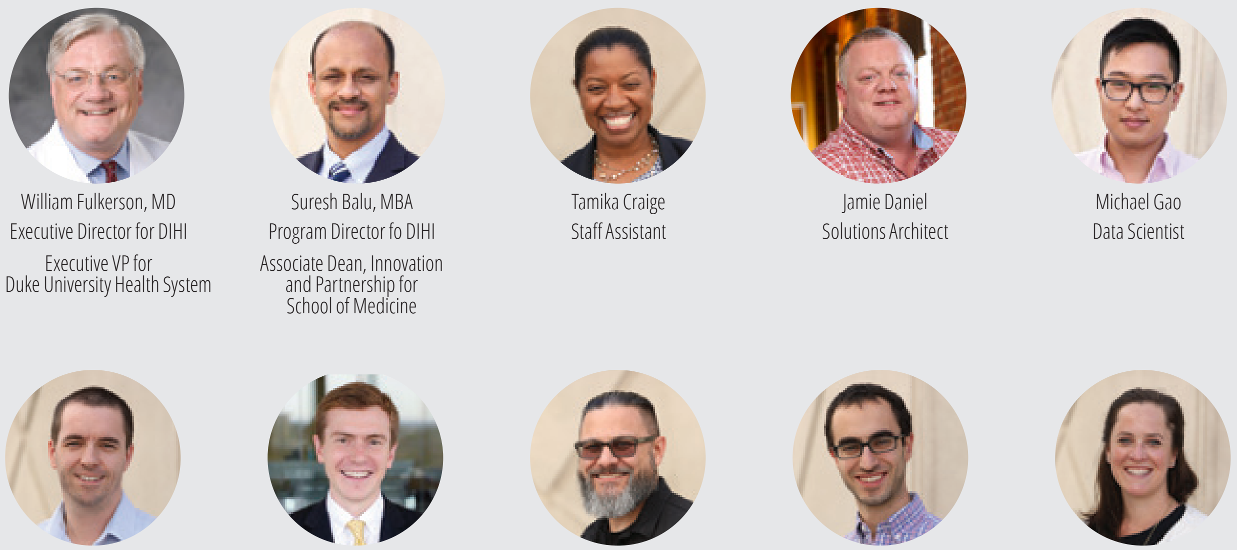

William Fulkerson, MD
Executive Director for DIHI
Executive VP for Duke University Health System

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

Tamika Craige
Staff Assistant

Jamie Daniel
Solutions Architect

Michael Gao
Data Scientist

Marshall Nichols
Systems & Analytics Manager

Will Ratliff, MBA
Innovation Program Manager

Mike Revoir
Solutions Architect

Mark Sendak, MD
Clinical Data Scientist

Krista Whalen
Innovation Program Manager

DIHI Innovation Scholars

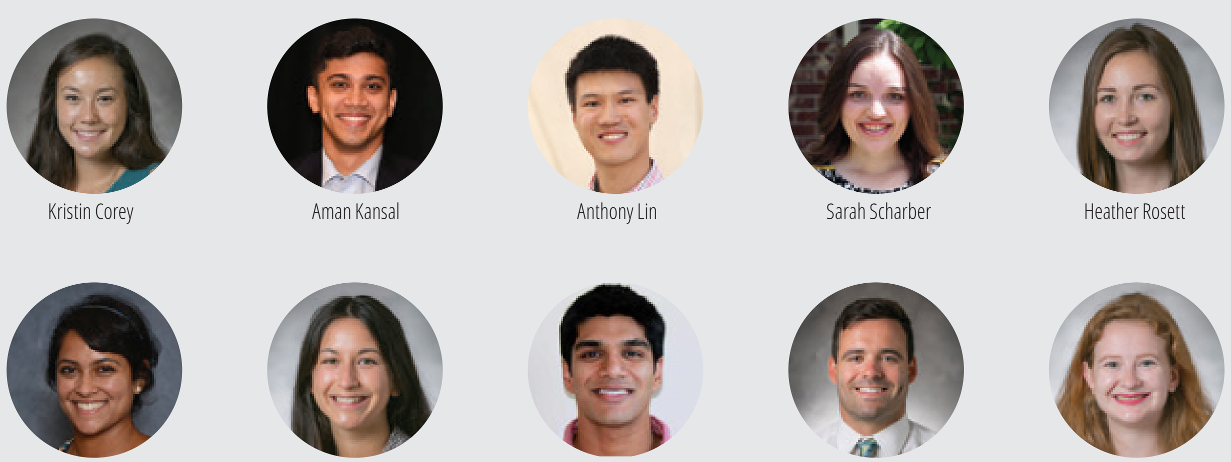

Kristin Corey

Aman Kansal

Anthony Lin

Sarah Scharber

Heather Rosett

Karishma Sriram

Christelle Tan

Sehj Kashyap

Joshua Helmkamp

Morgan Simons

DIHI Collaborators

Benjamin Alman, MD
John Anderson, MD
David Attarian, MD
Armando Bedoya, MD
Clay Bordley, MD
Hayden Bosworth, PhD
Lynn Bowlby, MD
Stephanie Brinson, FNP
Marion Broome, MD
Rob Califf, MD
Leslie Calihman Alabi
David Casarett, MD
Alex Cho, MD
Nathan Copeland, MD
Daniel Costello, MPA
Alexander Faranoff, MD
Jeff Ferranti, MD
Jonathan Fischer, MD
Joe Futoma, PhD
Benjamin Goldstein, PhD
Larry Greenblatt, MD
Brian Griffith, MD
Mitch Heflin, MD
Nicole Heilbron, PhD
Katherine Heller, PhD
Adrian Hernandez, MD
Katherine Hobbs Knutson, MD
Fred Johnson, MBA
Schuyler Jones, MD
Azalea Kim, MD
Allan Kirk, MD
Mary Klotman, MD
Eugenie Komives, MD
Sandhya Lagoo-Deenadayalan, MD, PhD
Elizabeth Lorenzi
Jared Lowe, MD
Kyla Machell, PhD
Gary Maslow, MD
Chad Mather, MD
Shelley McDonald, PhD
Eva Mihm, MSW
Paul Newman, MHA
Cara O'Brien, MD
Thomas Owens, MD
Theodore Pappas, MD
Manesh Patel, MD
Matthew Phelan, MS
Kathryn Pollak, PhD
Eric Poon, MD
Spenser Radtke, LCSWA
Dev Sangvai, MD
Kevin Shah, MD
Shilpa Shelton, MHA
Jessica Sperling, PhD
Kristen Stefureac, MSW
Madhav Swaminathan, MD
Annemarie Thompson, MD
Tracy Wang, MD
Yousuf Zafar, MD

Duke Institute
for Health Innovation

DukeHealth

@DukeInnovate

dihi.org