Center for AI Innovation in Healthcare

2nd Annual Holistic AI in Medicine (HAIM) Summit
November 14, 2024

Dinner To Go	**4:30p – 5:15p**
Safety Moment and Welcome • Barry Stein, MD, Vice President and Chief Clinical Innovation Officer, HHC	**5:15p – 5:20p**
Opening Remarks • Jeffrey Flaks, President and Chief Executive Officer, HHC	**5:20p – 5:30p**
HAIM Overview • Dimitris Bertsimas, PhD, Vice Provost (Open Learning), and Associate Dean of Master of Business Analytics, MIT	**5:30p – 5:35p**
Part 1: Project Presentations: MIT Researchers & HHC Leaders Moderated by Dimitris Bertsimas, PhD • Prediction of Left Ventricular Ejection Fraction from a 12-Lead Electrocardiogram • Prescribing SAVR vs. TAVR to Minimize Long-Term Mortality in Aortic Stenosis • Detection and Localization of Subdural Hematoma • Patient Screening for Early Detection of Aortic Stenosis • Predicting Psychiatric Admissions and Prognosis of Psychiatric Conditions • BrainSage: What Aging in the Brain Can Tell Us About Neurological Conditions • Rapid Identification of Psychosis Biotypes from the BSNIP Project Using AI/ML	**5:35p – 6:45p**
Break	**6:45p – 6:55p**
Part 2: Project Presentations: MIT Researchers & HHC Leaders Moderated by Dimitris Bertsimas, PhD • Predicting Readmission of Fragility Hip Fractures • Leveraging Artificial Intelligence in Utilization Management/ Appeals and Denials • Using Holistic AI to Diagnose Injury Grade for Liver Injuries • Predicting Medication-Related Readmission Risks • Eliminating Diagnostic Error in UTI Detection • Automated Data Extraction for the STS National Database	**6:55p – 7:55p**
MIT Leadership Panel Moderated by Barry Stein, MD • Cynthia Barnhart, PhD, Provost, MIT • Georgia Perakis, PhD, Dean (Interim) of MIT-Sloan School of Management • Dimitris Bertsimas, PhD, Vice Provost, MIT	**7:55p – 8:25p**
Looking to the Future • Dimitris Bertsimas, PhD	**8:25p – 8:40p**
Closing Remarks • Ajay Kumar, MD, Executive Vice President & Chief Clinical Officer, HHC	**8:40p – 8:45p**
Dessert Reception	**8:45p – 9:15p**

HAIM Website

Post Summit Survey

Background

Hartford HealthCare and MIT established an Innovation and Research collaboration partnership in 2017 focused on "Holistic AI in Medicine" (HAIM) to help shift the frontiers of AI in Healthcare. The differentiated capabilities we've assembled along the journey catalyzed the launch of HHC's "Center for AI Innovation in Healthcare", with the goal of unlocking the full potential of AI for our patients in a safe and trustworthy way.

Goals

1. Review HAIM initiatives, and potential transformative clinical impact, through the lenses of MIT researchers and HHC clinicians
2. Explore next horizon opportunities between HHC and MIT
3. Generate new ideas to continue building, and leveraging, the evolving capabilities of the "Center for AI Innovation in Healthcare"

Introduction to Projects & Participants

<table>
<tr><th colspan="2">Project #1: Prediction of Left Ventricular Ejection Fraction from a 12-Lead Electrocardiogram</th></tr>
<tr><td>Project Summary: Left ventricular ejection fraction (LVEF) serves as a critical indicator of cardiac function, traditionally assessed through echocardiography—a resource-intensive and time-consuming procedure. The initial phase of our study investigates a more accessible methodology for LVEF prediction utilizing solely electrocardiogram (EKG) data and basic patient health records. We employ a novel approach, transforming 12-lead EKG time series into tabular features, thereby enabling machine learning models to categorize a given EKG into normal, mildly, moderately, or severely reduced LVEF classifications for the corresponding patient. XGBoost is implemented as the primary model with strong performance on tabular features. Our findings highlight that EKGs may offer a viable alternative for LVEF estimation, opening up an exciting future where screening for symptoms that could be consistent with heart failure can be carried out already in EDs and PCP offices.</td><td>HHC Speaker: Steven Zweibel – Electrophysiology
MIT Speaker: Catherine Ning
Additional MIT Researchers: Yu Ma, Sara Pasquino, Cindy Wang</td></tr>
<tr><td colspan="2">MIT Speaker Bio: Catherine is a second-year Ph.D. student at the Operations Research Center, MIT. She graduated in June 2023 from the University of Oxford with a four-year integrated M.Eng. in Engineering Science. Born and raised in Luxembourg, Catherine has always been driven by curiosity, pursuing a broad range of interests and academic subjects, and she is fascinated by interdisciplinary fields. Recently, her research has focused on optimization and enhancing diagnostic and survival models in biomedicine and healthcare. She has led the HAIM project on the prediction of LVEF for the past year, which has inspired her to contribute to a more inclusive and universal healthcare system through data analysis, the multi-faceted capabilities of AI models, and shared knowledge.
  </td></tr>
</table>

Project #2: Prescribing SAVR vs. TAVR to Minimize Long-Term Mortality in Aortic Stenosis

Project Summary: SAVR (Surgical Aortic Valve Replacement) and TAVR (Transcatheter Aortic Valve Replacement) are essential procedures for treating severe aortic stenosis, a condition where the narrowing of the heart's aortic valve restricts blood flow, increasing the risk of heart failure and mortality. Choosing the most appropriate surgical option is crucial, as each procedure carries unique risks and benefits depending on patient characteristics such as age, overall health, and heart conditions. This project aims to prescribe the optimal surgery type for each patient to reduce 5-10 year mortality risk, leveraging a multimodal dataset that includes demographic details, clinical information, and medical notes from imaging scans. By integrating diverse patient data, we seek to build a model that can personalize treatment recommendations, ultimately improving long-term outcomes for patients.

HHC Speakers: Robert C. Hagberg, Jawad Haider
MIT Speaker: Yu Ma

MIT Speaker Bio: A final-year doctoral candidate, Yu has been collaborating with Hartford Hospital on early warning index and cardiovascular surgery. She has authored and co-authored research articles in Nature npj Digital Medicine, JAMA Surgery, Lancet Oncology, International Journal of Radiation Oncology and JCO Clinical Cancer Informatics. She was part of the original HAIM team that earned the Cognex prize at MIT's Machine Intelligence for Manufacturing and Operations Symposium, and was the graduate student representative on the MIT committee for the use of humans as experimental subjects.

Project #3: Detection and Localization of Subdural Hematoma

Project Summary: Subdural hematoma (SDH) is a common intracranial bleed on the brain's surface, with chronic SDH projected to become a leading neurosurgical concern in the U.S. by 2030 due to an aging population and increased use of antiplatelet and anticoagulant medications. Current diagnostic methods often lack timely, accurate detection, underscoring the need for advanced predictive tools. This study assesses machine learning (ML) approaches for predicting and localizing SDH using demographic, laboratory, medication, and imaging data. We employ deep learning architectures, including Convolutional Neural Networks (CNNs) for classification and U-Net models for segmentation. Our models demonstrate high accuracy in detecting and localizing SDH, and our AI tool provides clinicians with probabilistic assessments of CSDH presence, improving diagnostic accuracy, supporting timely interventions, and enhancing patient outcomes.

HHC Speaker: Tapan Mehta – Interventional Neuroradiology
MIT Speaker: Vassilina Stoumpou
Additional MIT Researchers: Rohan Kumar, Bernard Burman

MIT Speaker Bio: A third-year doctoral student, Vassilina grew up in Greece. She finished her five-year undergrad studies at the National Technical University of Athens in electrical and computer engineering, with a major in computer science. She worked at Tesla as an associate electrical design engineer before coming to MIT, where she is currently working on projects that focus on leveraging machine learning techniques in healthcare.

Project #4: Patient Screening for Early Detection of Aortic Stenosis

Project Summary: Aortic stenosis (AS) is a severe condition affecting 12% of older adults, typically diagnosed using echocardiograms (ECHOs). However, ECHOs are costly, highlighting the need for accessible, cost-effective screening methods. In this study, we develop a machine learning approach to diagnose AS using patient history, electrocardiograms (ECGs), and clinical notes, enabling earlier detection at the primary care level. Using data from over 130,000 patients, we trained XGBoost and TabNet models, evaluating the performance across a range of different data combinations. We achieved an F1 score of 0.81 when using tabular data, ECG data, and clinical notes in combination, which was an increase of +0.05 compared to tabular data alone. Our models identified key predictors aligned with known AS risk factors, such as age and hypertensive status, while also offering insights into ECG features linked to AS. As our models achieved a good balance of precision and recall, they provide a cost-effective screening alternative, with ECGs costing less than one-tenth of an ECHO. This approach could reduce unnecessary ECHO use, making AS screening more accessible. Our findings suggest that integrating diverse clinical data can enhance AS detection, offering a practical tool for early diagnosis in primary care settings.

HHC Speaker: Trevor Sutton – Cardiac Anesthesia,
MIT Speaker: Phevos Paschalidis
Additional MIT Researchers: Ikram Chairi, Muhammad Maaz, Dewang Kumar Agarwal

MIT Speaker Bio: Phevos is a senior at Harvard University studying computer science. He has been working with Hartford Hospital for more than a year to apply interpretable machine learning techniques to prescribe the optimal valve type to minimize post-operative pacemaker risk during Transcatheter Aortic Valve Replacement (TAVR) procedures. His other research interests include theoretical computer science and reinforcement learning.

Project #5: Predicting Psychiatric Admissions and Prognosis of Psychiatric Conditions

Project Summary: Using tabular EMR and clinical notes data, this study aims to predict key psychiatric outcomes, including 30-day and 1-year hospital admissions, as well as 1-year prognoses for schizophrenia, mood disorders, and anxiety disorders.

HHC Speaker: Manu Sharma – BHN
MIT Speaker: Jiayi Gu
Additional MIT Researchers: Emily Hahn, Nidhish Nerur

MIT Speaker Bio: Jiayi is a third-year doctoral at the Operations Research Center at MIT. She is passionate about healthcare and has worked in collaboration with Brigham and Women hospital in intimate partner violence using HAIM. Her research interests center around the synergy of machine learning and optimization in healthcare applications. She holds a bachelor's degree in Industrial Engineering and Global Health Studies from Northwestern University, graduating with a 3.99 GPA. She is particularly interested in applying holistic, data-driven approaches to address public health challenges.

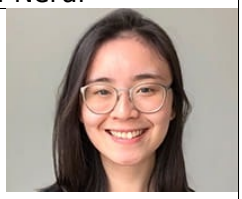

Project #6: BrainSage: What Aging in the Brain Can Tell Us About Neurological Conditions

Project Summary: Early detection of neurological disease and aging in the brain are challenging and crucial tasks in modern healthcare. Given the increasing availability of MRIs, especially T1-weighted (T1w) scans, there has been an effort to use brain scans to determine aging and for early detection and diagnosis of neurological disease. In this work, we develop deep learning models for brain age prediction using T1w MRI scans, leveraging a large data set of scans from the Olin Institute. Using these models, we investigate the relationship between aging in the brain and neurological conditions such as Autism, Alzheimer's, and Schizophrenia. We find some interesting and unintuitive relationships, in addition to some results that motivate future work and highlight the limitations of brain age as a proxy for brain health and function.

HHC Speaker: Michal Assaf – BHN
MIT Speaker: Matthew Hamilton Peroni
Additional MIT Researchers: Karl Zhu, Charles Dai, Joanna Kondylis

MIT Speaker Bio: Matthew is a third-year doctoral student. His research is focused on building interpretable, multi-modal, and robust machine learning systems for high-stakes decision making, especially in the healthcare domain. Prior to graduate school, he worked in industry as a machine learning engineer for a healthcare AI start-up. He received his bachelor's in mathematics and computer science from Cornell University.

Project #7: Rapid Identification of Psychosis Biotypes from the BSNIP Project Using AI/ML

Project Summary: Psychiatric diagnosis is based on symptoms but uses no biomarkers. Such diagnoses have good reliability, but zero validity, as their underlying causes are unknown, and unlike the majority of medical diagnoses there are no confirmatory or diagnostic tests for them. Biologically based diagnostic tests are foundational to developing rational treatments and identifying true disease entities, rather than crude syndromic classifiers. As part of the BSNIP study, scientists have developed unique biological "fingerprints" of the major psychotic syndromes, which have been clustered into 3 well-defined biotypes. The problem is that it takes approximately five hours to derive the data necessary to accurately classify psychosis patients into their various biotypes. Our findings for this project suggest that using MRI images, as well as readily available tabular data may give a straightforward tool for biotype designation that ideally could be employed in the medical office.

HHC Speaker: Godfrey Pearlson – BHN and Olin Institute
MIT Speaker: Kimberly M Villalobos Carballo
Additional MIT Researchers: Emily Hahn, Nidhish Nerur

MIT Speaker Bio: An MIT instructor and an incoming Assistant Professor at NYU, Kimberly's research integrates optimization and machine learning tools to enhance the performance of conventional algorithms. She is passionate about healthcare applications, and much of her research has been inspired by collaborations with hospitals (including work with Hartford Hospital on predictive machine learning models for length of stay reduction and identification of life-threatening patient events) that aim to improve the quality of their services and operations.

Project #8: Predicting Readmission of Fragility Hip fractures

Project Summary: Fragility hip fractures are defined as fractures due to a fall from at most standing height in individuals 50 years of age or greater. With over 10+ million incidences occurring globally from 1990 to 2019, and the number of incidences projected to nearly double by 2050, fragility hip fractures have great impact on the individual's physical and mental health; not only does the individual have an increased risk for additional fractures, but their mobility and independence is affected, resulting in sudden life changes and reduced quality of life. In particular, surgery is the main intervention, and readmission one of the main complications. We propose a machine learning model that leverages tabular data (demographics, labs, hospital visit information, past medical illnesses) and language data (discharge summary) to predict 30-day readmission risk. Our model has an AUC of 0.75. We note that readmission is a particularly difficult task, as there is no available data after the individual is discharged from the hospital; our model's performance is comparable to other readmission prediction models in the literature.

HHC Speaker: John Grady-Benson – BJI
MIT Speaker: Lisa Everest
Additional MIT Researcher: Catherine Ning

MIT Speaker Bio: Lisa is a third-year doctoral student and her research integrates optimization and machine learning with applications in healthcare and a special interest in prescription. She previously worked for three years as a quantitative analyst in Goldman Sachs' Special Situations Group in private asset management. She received bachelor's degrees in mathematics and management from MIT. In her free time, she enjoys snowboarding, figure skating, and baking.

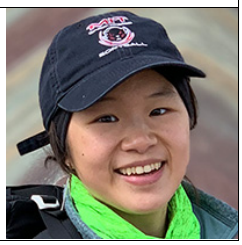

Project #9: Leveraging Artificial Intelligence in Utilization Management/Appeals and Denials

Project Summary: Utilization management is crucial for hospitals and healthcare systems to ensure proper reimbursement for services through accurate level-of-care reviews and denial management. However, complex payer policies, combined with the demands of patient-centered care, make it difficult for frontline providers to document effectively, often leading to unclear information, increased denials, and time-consuming manual reviews. To address these challenges, we are developing two multimodal AI tools aimed at supporting utilization management: a Clinical Decision Support Tool (DST) and an AI Summarizer. Our DST leverages multimodal machine learning to predict initial admission decisions in the emergency department (ED), using patient demographics, vitals, lab results, medications, and ED provider notes. The model achieves an area under the curve (AUC) of 87%, enabling practical decision support for admission assessments. For the AI Summarizer, we have developed a prototype that harnesses recent advances in large language models (LLMs) to generate patient summaries, aimed at assisting with insurance claims and appeals. These AI tools are anticipated to streamline utilization management, decrease denials and write-offs, and improve the efficiency of review processes. By providing timely and accurate support to healthcare providers, these tools will enhance documentation quality, ultimately benefiting both providers and patients.

HHC Speaker: Swathi Rachoor – ICP
MIT Speaker: Karl Zhu
Additional MIT Researchers: Yubing Cui, Cindy Wang

MIT Speaker Bio: Karl is a second-year graduate student specializing in operations research. Before coming to MIT, he was a graduate research assistant in the University of Auckland (New Zealand) Department of Engineering Science, where he built a strong foundation in applying statistical and optimization models. His work earned the Best Presentation of the Young Practitioner's Prize at the 2022 Operations Research Society of New Zealand conference. Karl holds a bachelor's degree in engineering science from the University of Auckland.

Project #10: Using Holistic AI to Diagnose Injury Grade of Liver Injuries

Project Summary: Liver injury is a critical concern in trauma care. As management depends on the severity of liver injury, cases are graded according to the American Association for the Surgery of Trauma (AAST) classification on a scale from I to V. However, recent literature has shown that CT reads by humans can sometimes misclassify liver injuries and machine learning frameworks have only been found effective in predicting high-grade injuries. We introduce a multimodal approach combining CT scans, as 3D images, unstructured clinical notes, vitals, and past medical history to predict the grading of liver injury for trauma patients. We apply our approach to 3,331 trauma patients at Hartford Healthcare to demonstrate the accuracy of our method. Our models achieve high averaged one-versus-rest AUC of 0.95 and demonstrate the value of multimodal AI in diagnosing medium-grade injuries.

HHC Speaker: Shea Gregg – Surgery
MIT Speaker: Carol Gao
Additional MIT Researchers: Ikram Chairi, Muhammad Maaz, Dewang Kumar Agarwal, Judith Brugman

MIT Speaker Bio: Carol is a first-year graduate student. She has collaborated with Hartford Healthcare on projects including liver injury prediction, infections prediction, and optimal surgery scheduling. She received her bachelor's degree in mathematics and economics from Smith College.

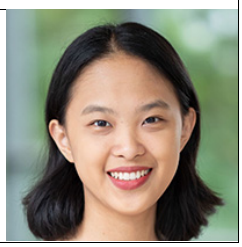

Project #11: Predicting Medication-Related Readmission Risks

Project Summary: When patients transition from an inpatient stay to home, they may face poor outcomes due to medication-related issues. A thorough review of each patient's medications by pharmacists could help reduce the risk of such complications before discharge. However, the limited pharmacist resources often make this approach impractical. In this presentation, we will introduce a multimodal machine learning model designed to prescreen inpatients prior to discharge, predicting their risk of readmission due to drug-related problems. This model enables the strategic allocation of pharmacist resources to patients who need them most.

HHC Speaker: Eric Arlia – Pharmacy
MIT Speaker: Yubing Cui
Additional MIT Researcher: Lisa Everest

MIT Speaker Bio: Yubing is a second-year doctoral student at MIT's Operations Research Center. She earned her bachelor's degree in Honors Mathematics and Computer Science from the University of Michigan, Ann Arbor. She has engaged in Machine Learning research since her undergraduate studies, covering theoretical foundations, algorithm development, and real-world applications, and has expanded into optimization and generative AI as new focus areas in her PhD work. She is passionate about applying these techniques in healthcare to make a positive impact in the real world.

Project #12: Eliminating Diagnostic Error in UTI Detection

Project Summary: Urinary Tract Infections (UTIs) account for approximately 3 million visits to the emergency department annually at a cost of approximately $3.5 billion. The perceived, ambiguous symptoms of the disease and the unreliability of clinical methods for UTI diagnosis, including urinalyses and urine cultures complicate care and facilitate the systematic overdiagnosis of UTI, which can lead to adverse drug events and delay targeted interventions. In this study, we propose a multimodal clinical support tool for pre-urinalysis UTI diagnosis by integrating tabular data extracted from the electronic health records of more than 30,000 patients and language data from clinician notes. Our model differs from existing work by making predictions before any tests have been ordered — thus addressing overprescription at its root cause — and by defining UTI diagnoses based on consultations from infectious disease experts rather than urine cultures which may contain false positives. Our model achieves a high AUROC of 0.843, surpassing comparative studies. Appropriate choice of thresholds, furthermore, could correspond to a 51% reduction in ordered urinalysis while accurately identifying 90% of UTI positive patients. We thus offer the first application of multimodal artificial intelligence to UTI prediction, with model performance surpassing comparable studies. We outline a process by which this clinical support tool can be integrated into existing hospital pipelines.

HHC Speakers: Ulysses Wu, Adam Steinberg
MIT Speaker: Phevos Paschalidis
Additional MIT Researchers: Carol Gao, Konstantina Rasvani

MIT Speaker Bio: Phevos is a senior at Harvard University studying computer science. He has been working with Hartford Hospital for more than a year to apply interpretable machine learning techniques to prescribe the optimal valve type to minimize post-operative pacemaker risk during Transcatheter Aortic Valve Replacement (TAVR) procedures. His other research interests include theoretical computer science and reinforcement learning.

Project #13: Automated Data Extraction for the STS National Database

Project Summary: We develop a pipeline that uses Large Language models to automatically extract hundreds of clinical variables for the STS National Database. Our pipeline is versatile and achieves very high accuracy standards, even outperforming human experts in some cases. Hence, with this work, we aim to standardize clinical data collection at scale, reducing the burden on healthcare professionals, while maintaining or improving data quality.

HHC Speakers: Bob Hagberg
MIT Leader: Dimitris Bertsimas
MIT Speaker: George Margaritis
Additional MIT Researcher: Periklis Petridis

MIT Speaker Bio: George is a fourth-year doctoral student, with research interests centered around AI for Healthcare and AI for Optimization. He finished his undergraduate degree in Electrical & Computer Engineering at the Technical University of Crete, achieving the highest grade in the school's 30-year history. George has also considerable experience in developing and scaling AI tools for real-world applications, which was gained through roles in both startups and major companies like Netflix.

MIT Leadership Attendees

Cynthia Barnhart, PhD, Provost, MIT

Cynthia Barnhart is MIT's Provost and the Abraham J. Siegel Professor of Management Science and Professor of Operations Research at MIT. She earned her SM in transportation in 1986 and her PhD in 1988 from MIT. Barnhart has a distinguished career at MIT, serving in many leadership roles, including most recently as MIT's Chancellor from 2014 to 2021. She also served as the Associate and Acting Dean for the School of Engineering. A member of the Institute's faculty since 1992 and an elected member of the National Academy of Engineering and of the American Academy of Arts and Sciences, Barnhart's teaching and research has focused on the areas of large-scale optimization and analytics, with applications in transportation and logistics systems. She has supervised the thesis research of scores of graduate and undergraduate students across a range of disciplines and has published widely in the flagship journals of her field.

Georgia Perakis, PhD, Dean (Interim) of MIT-Sloan School of Management

Georgia Perakis is the John C Head III Dean (Interim) of the MIT Sloan School of Management and a Professor of Operations Management, Operations Research & Statistics at the MIT Sloan School of Management. She is on leave from the roles of co-director of the Operations Research Center and Associate Dean for Social and Ethical Responsibility in Computing (SERC) in the Schwarzman College of Computing and MIT Sloan. Her widely published research has received many awards and focuses on analytics/AI, in particular, in the intersection of optimization and machine learning with applications in pricing, revenue management, supply chain, and healthcare among others. She received the PECASE Award from the Office of the President on Science and Technology. In 2016, she was elected as an INFORMS Fellow, and in 2021 as Distinguished MSOM Fellow.

Perakis has passion for supervising PhD, masters, and undergraduate students, having graduated 30 PhD and 59 masters students. She has received numerous awards for teaching including the Graduate Student Council Teaching Award (2002), the Samuel M. Seegal Award (2012), the Jamieson Prize for excellence in teaching (2014), and the Teacher of the Year Award (2017) at MIT Sloan. Perakis is currently the Editor in Chief of the *M&SOM* journal and has served as editor at a number of other publications. She holds a BS in mathematics from the University of Athens as well as an MS in applied mathematics and a PhD in applied mathematics from Brown University.

Dimitris Bertsimas, PhD, Vice Provost (Open Learning), and Associate Dean of Master of Business Analytics, MIT

Dimitris Bertsimas is the *Boeing Leaders for Global Operations Professor of Management*, a Professor of Operations Research, and the Associate Dean of Business Analytics at MIT. He was named Vice Provost for Open Learning in September 2024. He is a member of the US National Academy of Engineering, an INFORMS fellow, recipient of the John von Neumann theory prize, the Frederick W. Lanchester and the INFORMS president awards among many other research and teaching awards, supervisor of 98 completed and 24 current doctoral theses and co-founder of twelve AI companies, four of which have been sold. He has co-authored more than 300 scientific papers, 8 books. Bertsimas holds a BS in electrical engineering and computer science from the National Technical University of Athens, Greece, as well as an MS in operations research and a PhD in applied mathematics and operations research from MIT.