Advancing causal inference with machine learning and real-world data:

An application of targeted machine learning and super learners on hospital-acquired pressure injuries from MIMIC IV

MSR102

View all of Parexel's posters at ISPOR Europe

Gregg M[1], Wilson A[1,2], Streja E[1], Vanderpuye-Orgle J[1], Roessner M[1], Alderden J[3]

[1] Parexel International, Waltham, MA, USA [2] University of Utah, Salt Lake City, UT, USA [3] Boise State University, Boise, ID, USA

Background

Traditional causal methods typically rely on parametric statistical models that impose restrictive assumptions about underlying data structure. Recent advancements in targeted machine learning (ML) and super learning enable the identification of causal estimates in real-world data, irrespective of data complexity or structure, offering a flexible and comprehensive approach to causal inference. In our study, we leverage the Medical Information Mart for Intensive Care (MIMIC) IV database and ML models to ascertain the causes of hospital-acquired pressure injuries (HAPrI) and develop a risk prediction algorithm.

Methods

The MIMIC IV, version 2.2 dataset (1) is a deidentified electronic health records dataset from Beth Israel Deaconess Medical Center in Boston USA. The dataset captures admissions from 2008-2019 for 299,712 patients.

Eligible cases were defined as patients admitted to the intensive care unit (ICU).

We used cost-sensitive ensemble super learning to predict HAPrI in the ICU. Cost-sensitive super learners, assign different costs to different types of prediction errors, and take into account the clinical weight of false positives and false negatives and balance accordingly.

We then used a clinically-informed debiasing methods, directed acyclic graphs (DAG) and targeted maximum likelihood estimation (TMLE)(2), to estimate the potential causal effect of albumin (identified by super-learner) on HAPrI.

Results

- Of 28,395 eligible cases, 1,395 developed a pressure injury (4.9%).
- The DAG illustrates results of the ensemble super-learner, identifying albumin as a causal variable in pressure injury development. (Figure 1)
- The ensemble super learner had a cross-validated AUC of 0.8, with 45.6% sensitivity and 88.8% specificity.
- The crude odds ratio of low (below 3.0) albumin on pressure injury was significant: OR = 2.86, p<0.0001. (Figure 2)
- The TMLE-adjusted estimate was significant but attenuated: OR = 2.22, p<0.0001. (Figure 3)

With causal structural modeling, we are able to test hypotheses about pressure injury formation and both inform better interventions, but also elucidate mechanisms and better classify outcomes as preventable or intrinsic.

Figure 1. DAG and Causal Assumptions

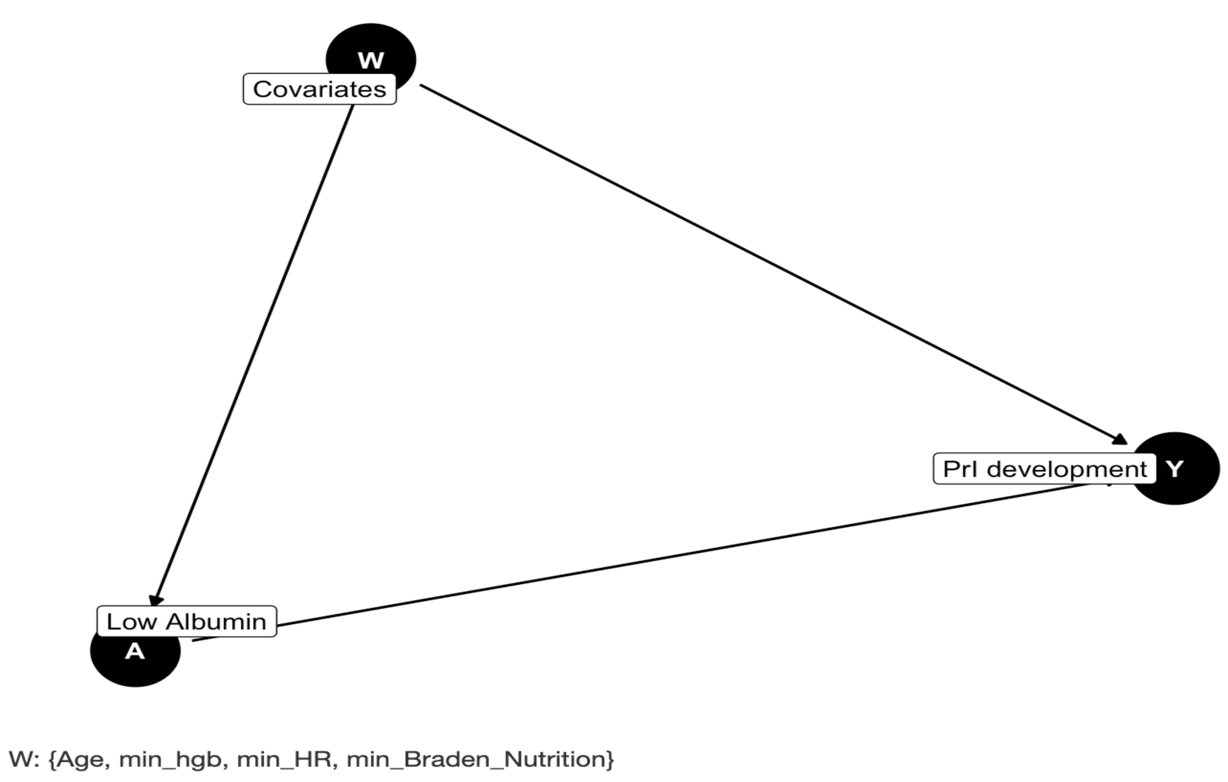

Figure 2. Crude Effects

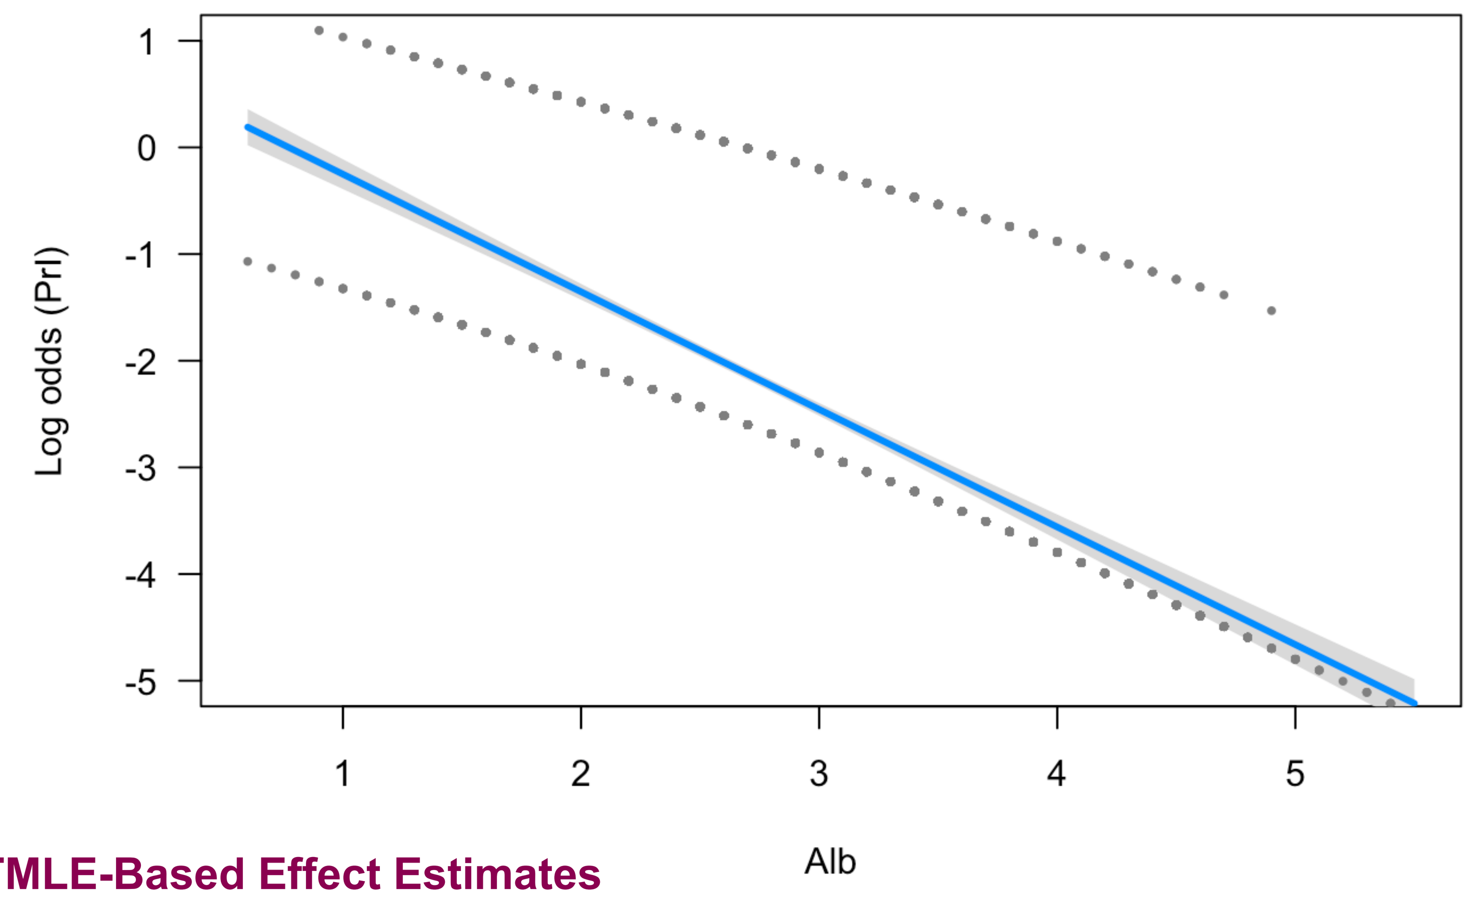

Figure 3. TMLE-Based Effect Estimates

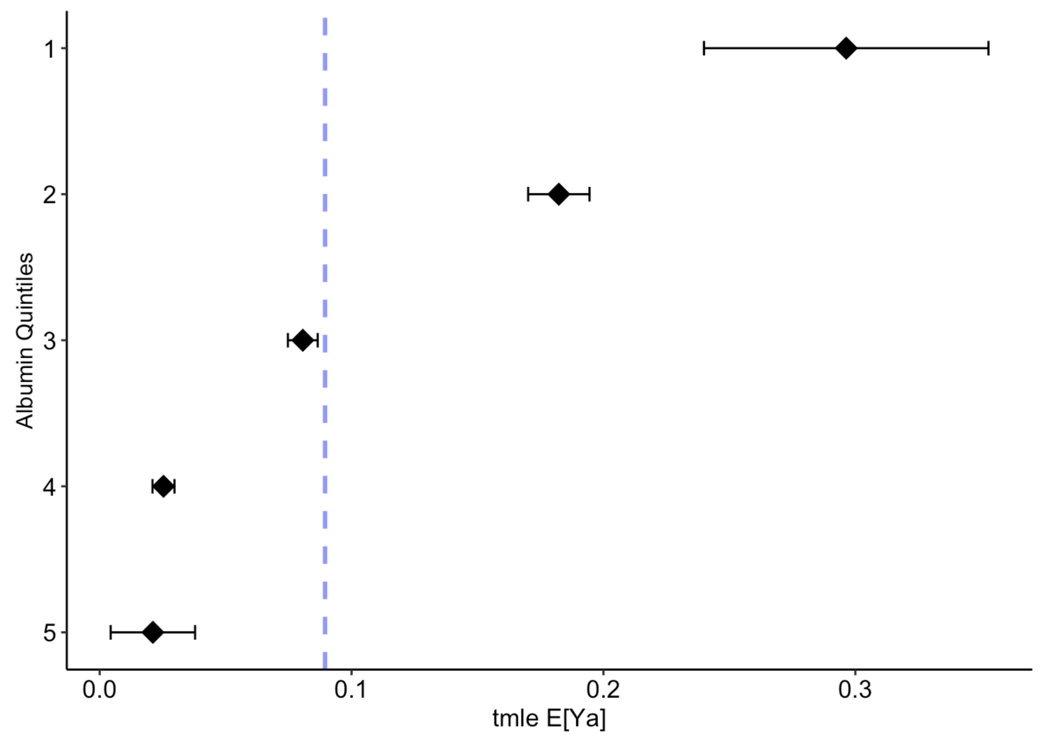

Figure 4. Causal Hypothesis and Revised Perspective

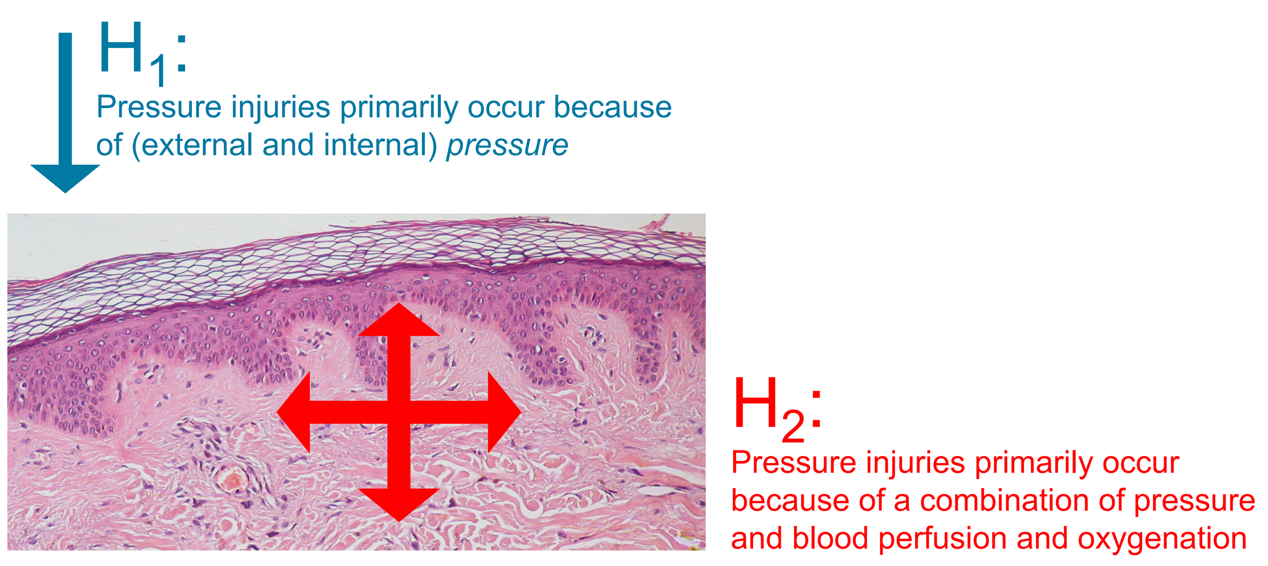

Conclusions

- Previous models predicting pressure injuries often favor overall accuracy and have a clinically uninformative sensitivity, whereas traditional Braden scales classify almost all patients as high-risk.
- Our initial models use (ensemble and discrete) cost-sensitive super learning to support risk prediction for HAPrI.
- We then developed structural causal models and leveraged targeted maximum likelihood estimation to evaluate the potential interventional effects of low albumin levels on the development of pressure injuries, while accounting for plausible confounding pathways.
- This ML approach allows for improved prediction and prevention of pressure injuries in the ICU, which can help to significantly improve patient outcomes and reduce healthcare costs.
- The true power of ML emerges when it's combined with causal modeling, offering both predictive prowess and deep insights into data relationships. The next frontier in clinical prediction seamlessly blends accuracy, clinical acumen, and the enriched utilization of real-world data through machine learning.

REFERENCES

[1] Johnson, A., Bulgarelli, L., Pollard, T., Horng, S., Celi, L. A., & Mark, R. (2023). MIMIC-IV (version 2.2). PhysioNet. https://doi.org/10.13026/6mm1-ek67.

[2] Gruber, S., & Laan, M. van der. (2012). tmle: An R Package for Targeted Maximum Likelihood Estimation. Journal of Statistical Software, 51, 1–35. https://doi.org/10.18637/jss.v051.i13

[3] Alderden, J., Pepper, G. A., Wilson, A., Whitney, J. D., Richardson, S., Butcher, R., Jo, Y., & Cummins, M. R. (2018). Predicting Pressure Injury in Critical Care Patients: A Machine-Learning Model. American Journal of Critical Care, 27(6), 461–468. https://doi.org/10.4037/ajcc2018525

www.parexel.com