Toward a Two-Tier Clinical Warning System for Hospitalized Patients

Gregory Hackmann[1], Minmin Chen[1], Octav Chipara[1*], Chenyang Lu[1], Yixin Chen[1], Marin Kollef[2], Thomas C. Bailey[2]

[1] Department of Computer Science and Engineering, Washington University

[2] Department of Medicine, Washington University School of Medicine

[*] Now with Department of Computer Science, University of Iowa

Vision: Two-Tiered Clinical Early Warning

- Tier 1: identify at-risk patients from existing medical record data

+

- Tier 2: issue real-time warnings from existing medical records and real-time vital sign monitoring

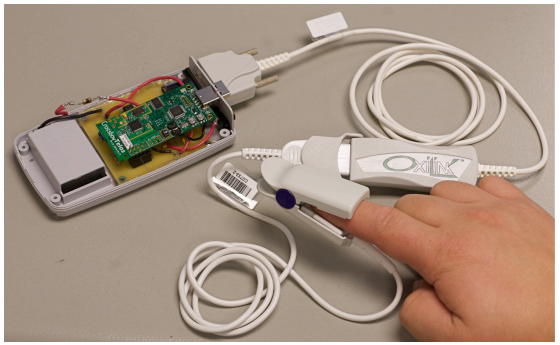

Early Warning System

- Electronic medical record systems aggregate a wealth of data about a patient's condition
- Challenge: how to determine the importance of these data?

- Feasibility study using approach based on *logistic regression*

Algorithm Overview

- Logistic regression assigns a weight to each kind of input in predicting an outcome
- Standard logistic regression does not:
  - Operate on time-series data
  - Handle missing data

Algorithm Overview

- Logistic regression assigns a weight to each kind of input in predicting an outcome
- Standard logistic regression does not:
  - Operate on time-series data
  - Handle missing data

- Split window of data into *n* equally-sized "buckets"
- Calculate min/mean/max of each bucket
- Find separate weights for all 3*n* values

Algorithm Overview

- Logistic regression assigns a weight to each kind of input in predicting an outcome
- Standard logistic regression does not:
  - Operate on time-series data
  - Handle missing data

- Try to fill in empty buckets by "carrying over" most recent value
- If patient had *no* data for a variable, use mean over entire historical dataset as a fallback

Evaluation: Retrospective Analysis

- Dataset of 28,927 hospital visits from 19,116 patients
- 36 categories of data + outcome (ICU transfer)

- Snapshot of 24 hours' data for each patient, divided into 6 buckets
- Use first half of dataset to train logistic model
- Use second half of dataset to test model against known outcome

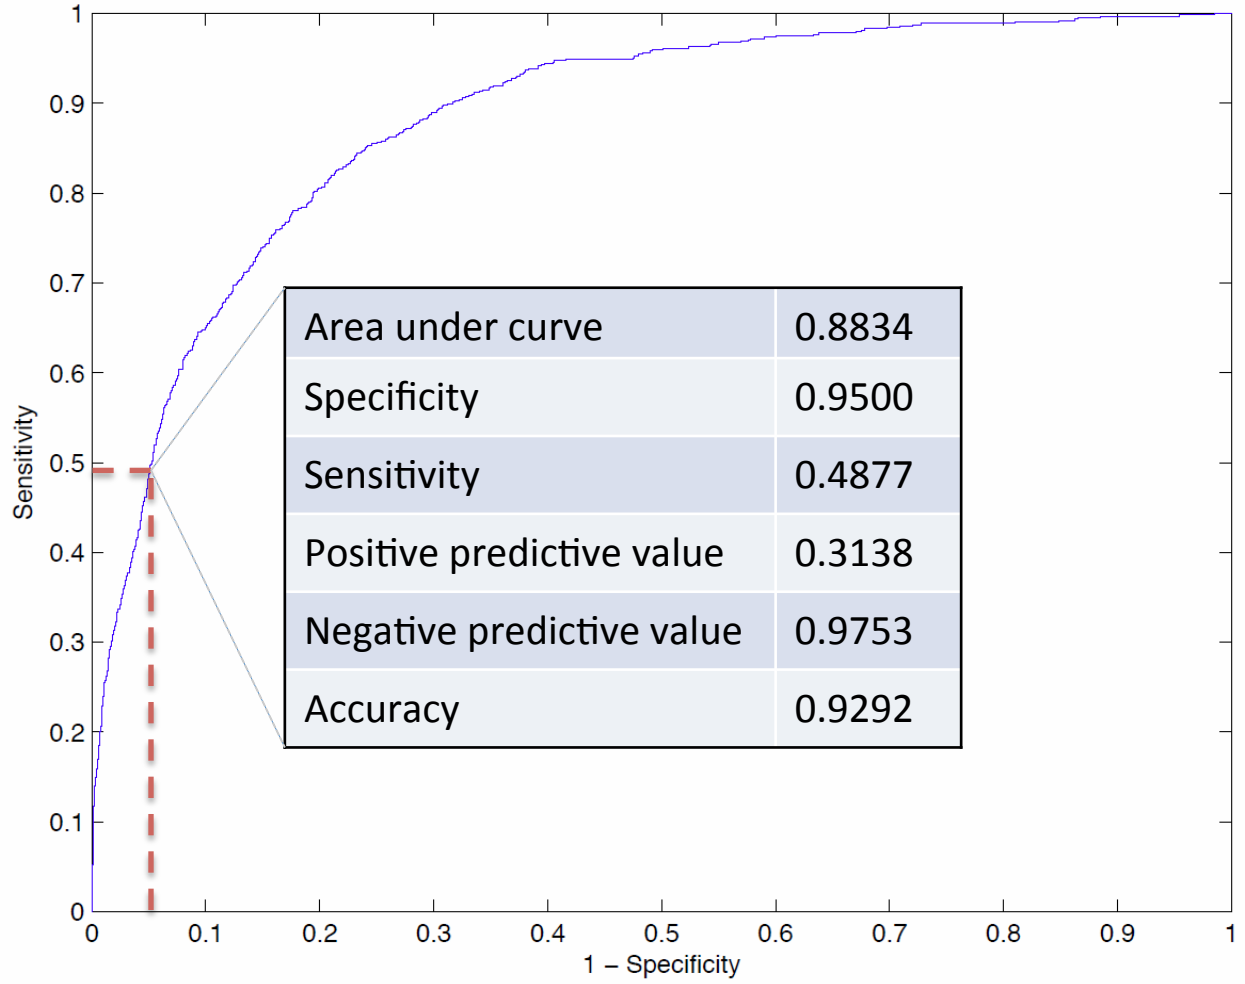

Area under curve	0.8834
Specificity	0.9500
Sensitivity	0.4877
Positive predictive value	0.3138
Negative predictive value	0.9753
Accuracy	0.9292

Retrospective Analysis Predictive Performance

Evaluation: Real-Time Simulation

- Retrospective study looks at one 24-hour window per patient
- Real-time detection system would produce a series of scores over entire hospital stay

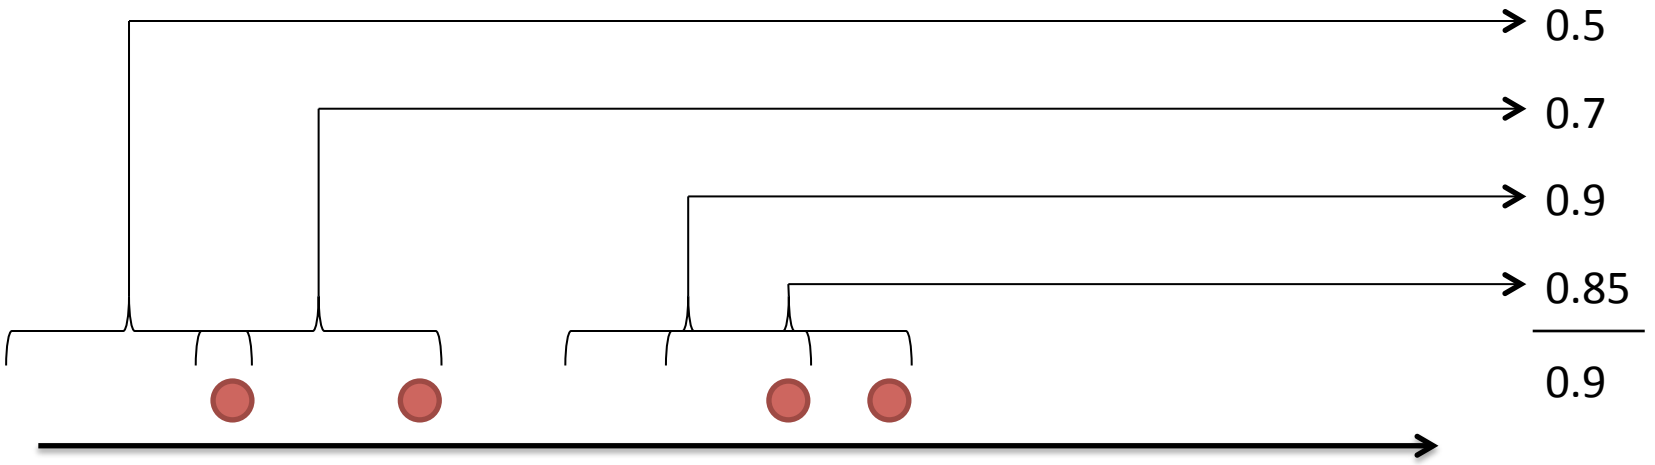

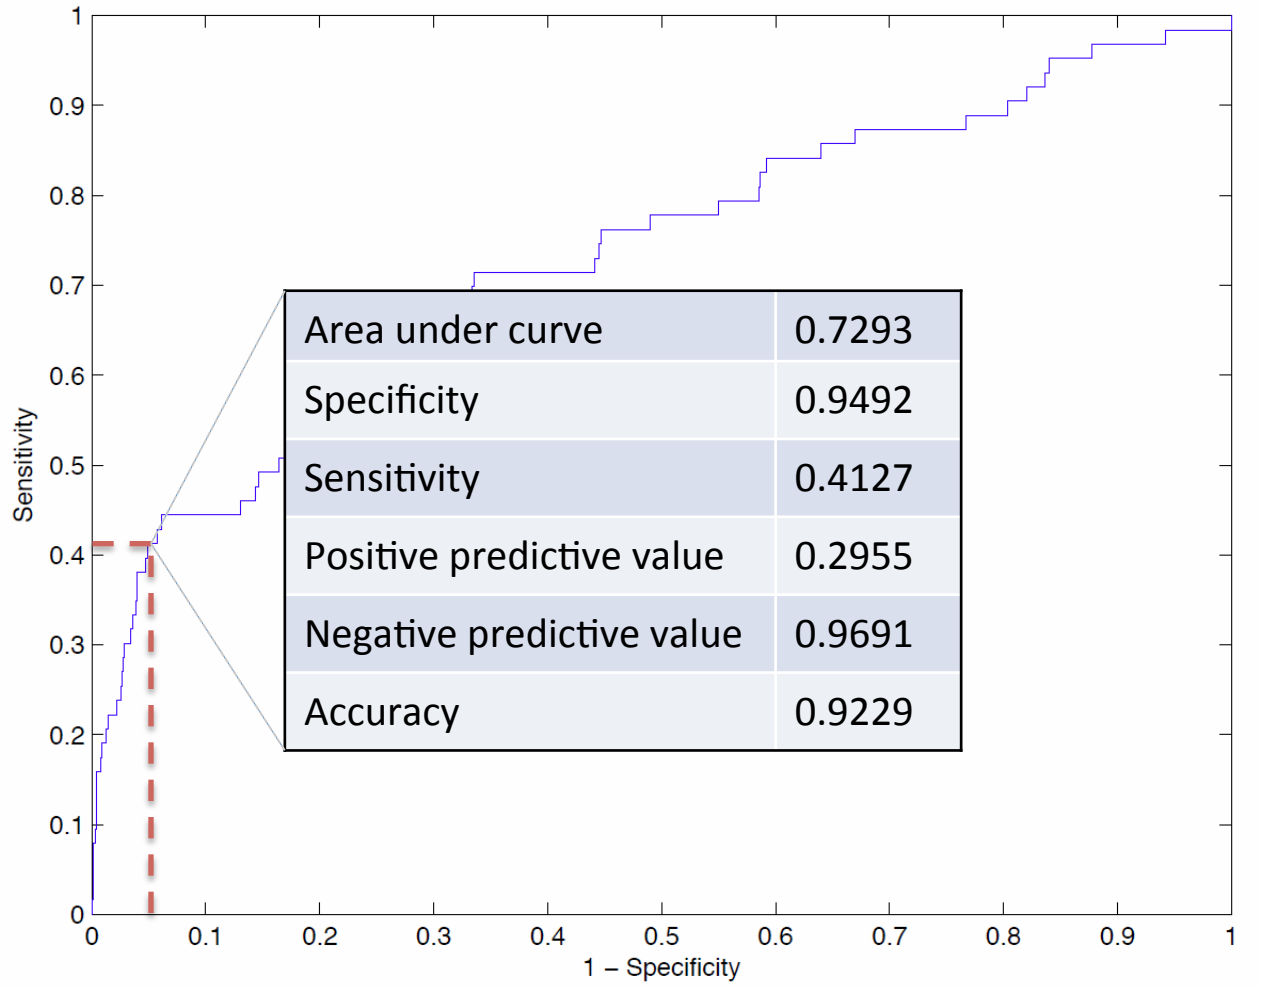

Area under curve	0.7293
Specificity	0.9492
Sensitivity	0.4127
Positive predictive value	0.2955
Negative predictive value	0.9691
Accuracy	0.9229

Real-Time Simulation Predictive Performance

Conclusion

- Tier 1: early warning feasibility study
- Sensitivity of 41.3% at a manageable alarm rate

- Tier 2: real-time vital sign monitoring
- Under clinical trial in four units at Barnes-Jewish Hospital

- Ongoing work: enhancements to Tier 1 performance, development of Tier 2 real-time detection algorithm