



HUMANITY  
CENTERED  
ROBOTICS  
INITIATIVE



BROWN

## Rich Caruana Microsoft Research

### TECH TALK: "Intelligible Machine Learning for HealthCare"



**Wednesday, October 26, 2016**  
**1-2pm**  
**CIT 477, Lubrano Conference Room**

**Abstract:** In machine learning often a tradeoff must be made between accuracy and intelligibility: the most accurate models usually are not very intelligible (e.g., neural nets, boosted trees, and random forests), and the most intelligible models usually are less accurate (e.g., linear or logistic regression). This tradeoff often limits the accuracy of models that can be applied in mission-critical applications such as healthcare where being able to understand, validate, edit, and ultimately trust a learned model is important. We have developed a learning method based on generalized additive models (GAMs) that is often as accurate as full complexity models, but remains as intelligible as linear/logistic regression models. In the talk I'll present two case studies where these high-performance generalized additive models (GA2Ms) are applied to healthcare problems yielding intelligible models with state-of-the-art accuracy. In the pneumonia risk prediction case study, the intelligible model uncovers surprising patterns in the data that previously prevented complex learned models from being deployed, but because it is intelligible and modular allows these patterns to easily be recognized and removed. In the 30-day hospital readmission case study, we show that the same methods scale to large datasets containing hundreds of thousands of patients and thousands of attributes while remaining intelligible and providing accuracy comparable to the best (unintelligible) machine learning methods.

Rich Caruana is a Senior Researcher at Microsoft Research. Before joining Microsoft, Rich was on the faculty in the Computer Science Department at Cornell University, at UCLA's Medical School, and at CMU's Center for Learning and Discovery. Rich's Ph.D. is from Carnegie Mellon University, where he worked with Tom Mitchell and Herb Simon. His thesis on Multi-Task Learning helped create interest in a new subfield of machine learning called Transfer Learning. Rich received an NSF CAREER Award in 2004 (for Meta Clustering), best paper awards in 2005 (with Alex Niculescu-Mizil), 2007 (with Daria Sorokina), and 2014 (with Todd Kulesza, Saleema Amershi, Danyel Fisher, and Denis Charles), co-chaired KDD in 2007 (with Xindong Wu), and serves as area chair for NIPS, ICML, and KDD. His current research focus is on learning for medical decision making, deep learning, and computational ecology.

Host: Amy Greenwald/HCRI