Learning from Real-World Deployment

Suchi Saria, PhD
John C. Malone Associate Professor, Computer Science, Statistics, and Health Policy
Director, Machine Learning and Healthcare Lab at Hopkins
Research Director, Malone Center for Engineering in Healthcare

Founder, Bayesian Health

joint work w/ Adarsh Subbaswamy, Peter Schulam, Katie Henry, and Roy Adams

Support:

GORDON AND BETTY MOORE FOUNDATION

NSF

ALFRED P. SLOAN FOUNDATION 1934

NIH National Institutes of Health

DARPA

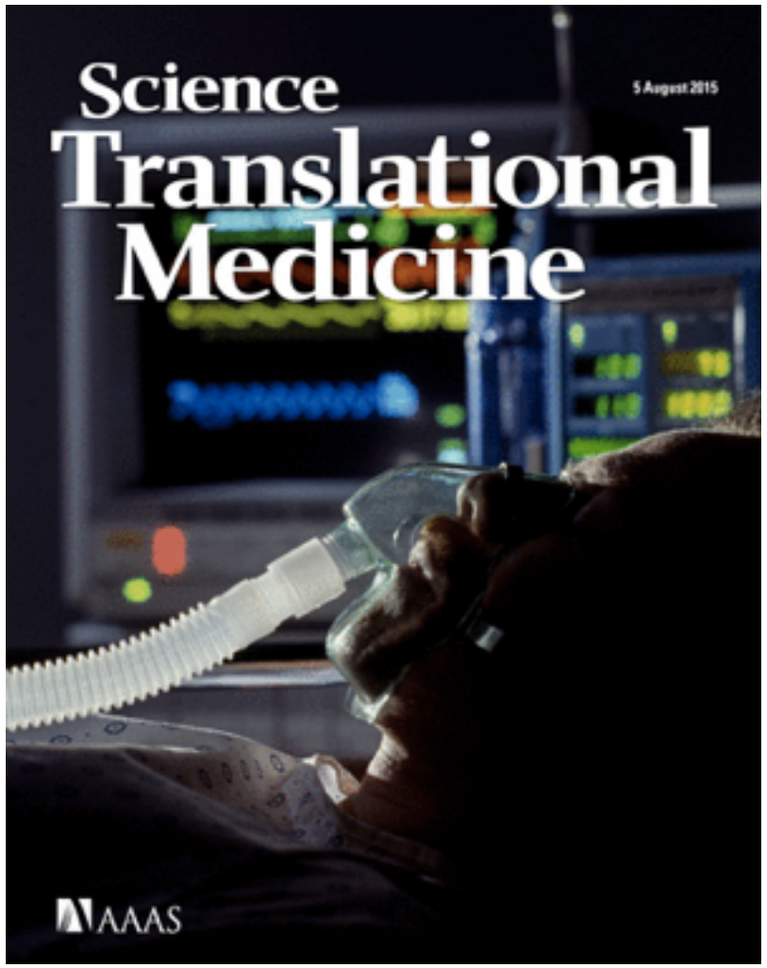

2015

Many innovations
for improving
quality of tools

Infrastructure
for deployment

2017

Real-world adoption
and behavior change

Sustainable Learning
core part of culture

Partnering with systems

Scalable Joint Models for Reliable Uncertainty-Aware Event Prediction

Hossein Soleimani, James Hensman, and Suchi Saria

Abstract—Missing data and noisy observations pose significant challenges for reliably predicting events from irregularly sampled multivariate time series (longitudinal) data. Imputation methods, which are typically used for completing the data prior to event prediction, lack a principled mechanism to ac
techniques can be used for jointly modeling tl
longitudinal observations. These approaches
signals with many observations. Our propose
joint model based upon sparse multiple-outpu
explain highly challenging structure including
predicting events using the distribution of the
delayed detection versus incorrect assessme
satisfy the derived confidence criteria. Experi
state-of-the-art techniques in event predictio

Index Terms—Uncertainty-aware prediction

Reliable Decision Support using Counterfactual Models

Peter Schulam
Department of Computer Science
Johns Hopkins University
Baltimore, MD 21211

Suchi Saria
Department of Computer Science
Johns Hopkins University
Baltimore, MD 21211

Can You Trust This Prediction? Auditing Pointwise Reliability After Learning

Peter Schulam
Department of Computer Science
Johns Hopkins University
pschulam@cs.jhu.edu

Suchi Saria
Department of Computer Science
Johns Hopkins University
ssaria@cs.jhu.edu

ı high stakes ap-
we need tools for
stem and evaluat-
ethods to improve
ire new learning
sian inference to
). An alternative

pass" of a predictive model and then train the model using gradient-based methods without much additional effort (e.g. Maclaurin et al. 2015; Abadi et al. 2016). As the barriers to building machine learning systems become lower, there is rising excitement around the idea of applying the technology in high-impact domains (e.g. Soleimani et al. 2018; Lipton et al. 2018).

Tools for quickly building machine learning systems, however, have generally outpaced the growth and adop-

Biostatistics (2020) **21**, 2, *pp.* 345–352
doi:10.1093/biostatistics/kxz041
Advance Access publication on November 19, 2019

From development to deployment: dataset shift, causality, and shift-stable models in health AI

ADARSH SUBBASWAMY

Department of Computer Science, Johns Hopkins University, 160 Malone Hall, 3400 N. Charles Street, Baltimore, MD, USA

SUCHI SARIA*

Department of Computer Science; Department of Applied Math & Statistics, and Department of Health Policy & Management, Johns Hopkins University, 160 Malone Hall, 3400 N. Charles Street, Baltimore, MD, USA

ssaria@cs.jhu.edu

Keywords: Causal inference; Dataset shift; Generalizability; Machine learning.

nature medicine

Explore our content Journal information

nature > nature medicine > perspectives > article

Perspective | Published: 19 August 2019

Do no harm: a roadmap for responsible machine learning for health care

Jenna Wiens ✉, Suchi Saria, Mark Sendak, Marzyeh Ghassemi, Vincent X. Liu, Finale Doshi-Velez, Kenneth Jung, Katherine Heller, David Kale, Mohammed Saeed, Pilar N. Ossorio, Sonoo Thadaney-Israni & Anna Goldenberg ✉

Original Clinical Report

Critical Care Explorations

OPEN

Comparison of Automated Sepsis Identification Methods and Electronic Health Record–based Sepsis Phenotyping: Improving Case Identification Accuracy by Accounting for Confounding Comorbid Conditions

Katharine E. Henry, MSE[1]; David N. Hager, MD, PhD[2]; Tiffany M. Osborn, MD, MPH[3]; Albert W. Wu, MD[4]; Suchi Saria, PhD[1,4,5]

[1]Department of Computer Science, Johns Hopkins University, Baltimore, MD.

[2]Division of Pulmonary and Critical Care, Department of Medicine, Johns Hopkins University, Baltimore, MD.

Objective: To develop and evaluate a novel strategy that automates the retrospective identification of sepsis using electronic health

Cite
Share
Favorites
Permissions

FOREWORD

Too Many Definitions of Sepsis: Can Machine Learning Leverage the Electronic Health Record to Increase Accuracy and Bring Consensus?

Saria, Suchi PhD; Henry, Katharine E. MSE Author Information

Critical Care Medicine: February 2020 - Volume 48 - Issue 2 - p 137-141
doi: 10.1097/CCM.0000000000004144

BUY EDITOR'S CHOICE
Metrics

Copyright © 2020 by the Society of Critical Care Medicine and Wolters Kluwer Health, Inc. All Rights Reserved.

Augmentation …

DR. AI

We should treat algorithms like prescription drugs

By Andy Coravos, Irene Chen, Ankit Gordhandas & Ariel Dora Stern · February 14, 2019

https://qz.com/1540594/treating-algorithms-like-prescription-drugs-could-reduce-ai-bias/

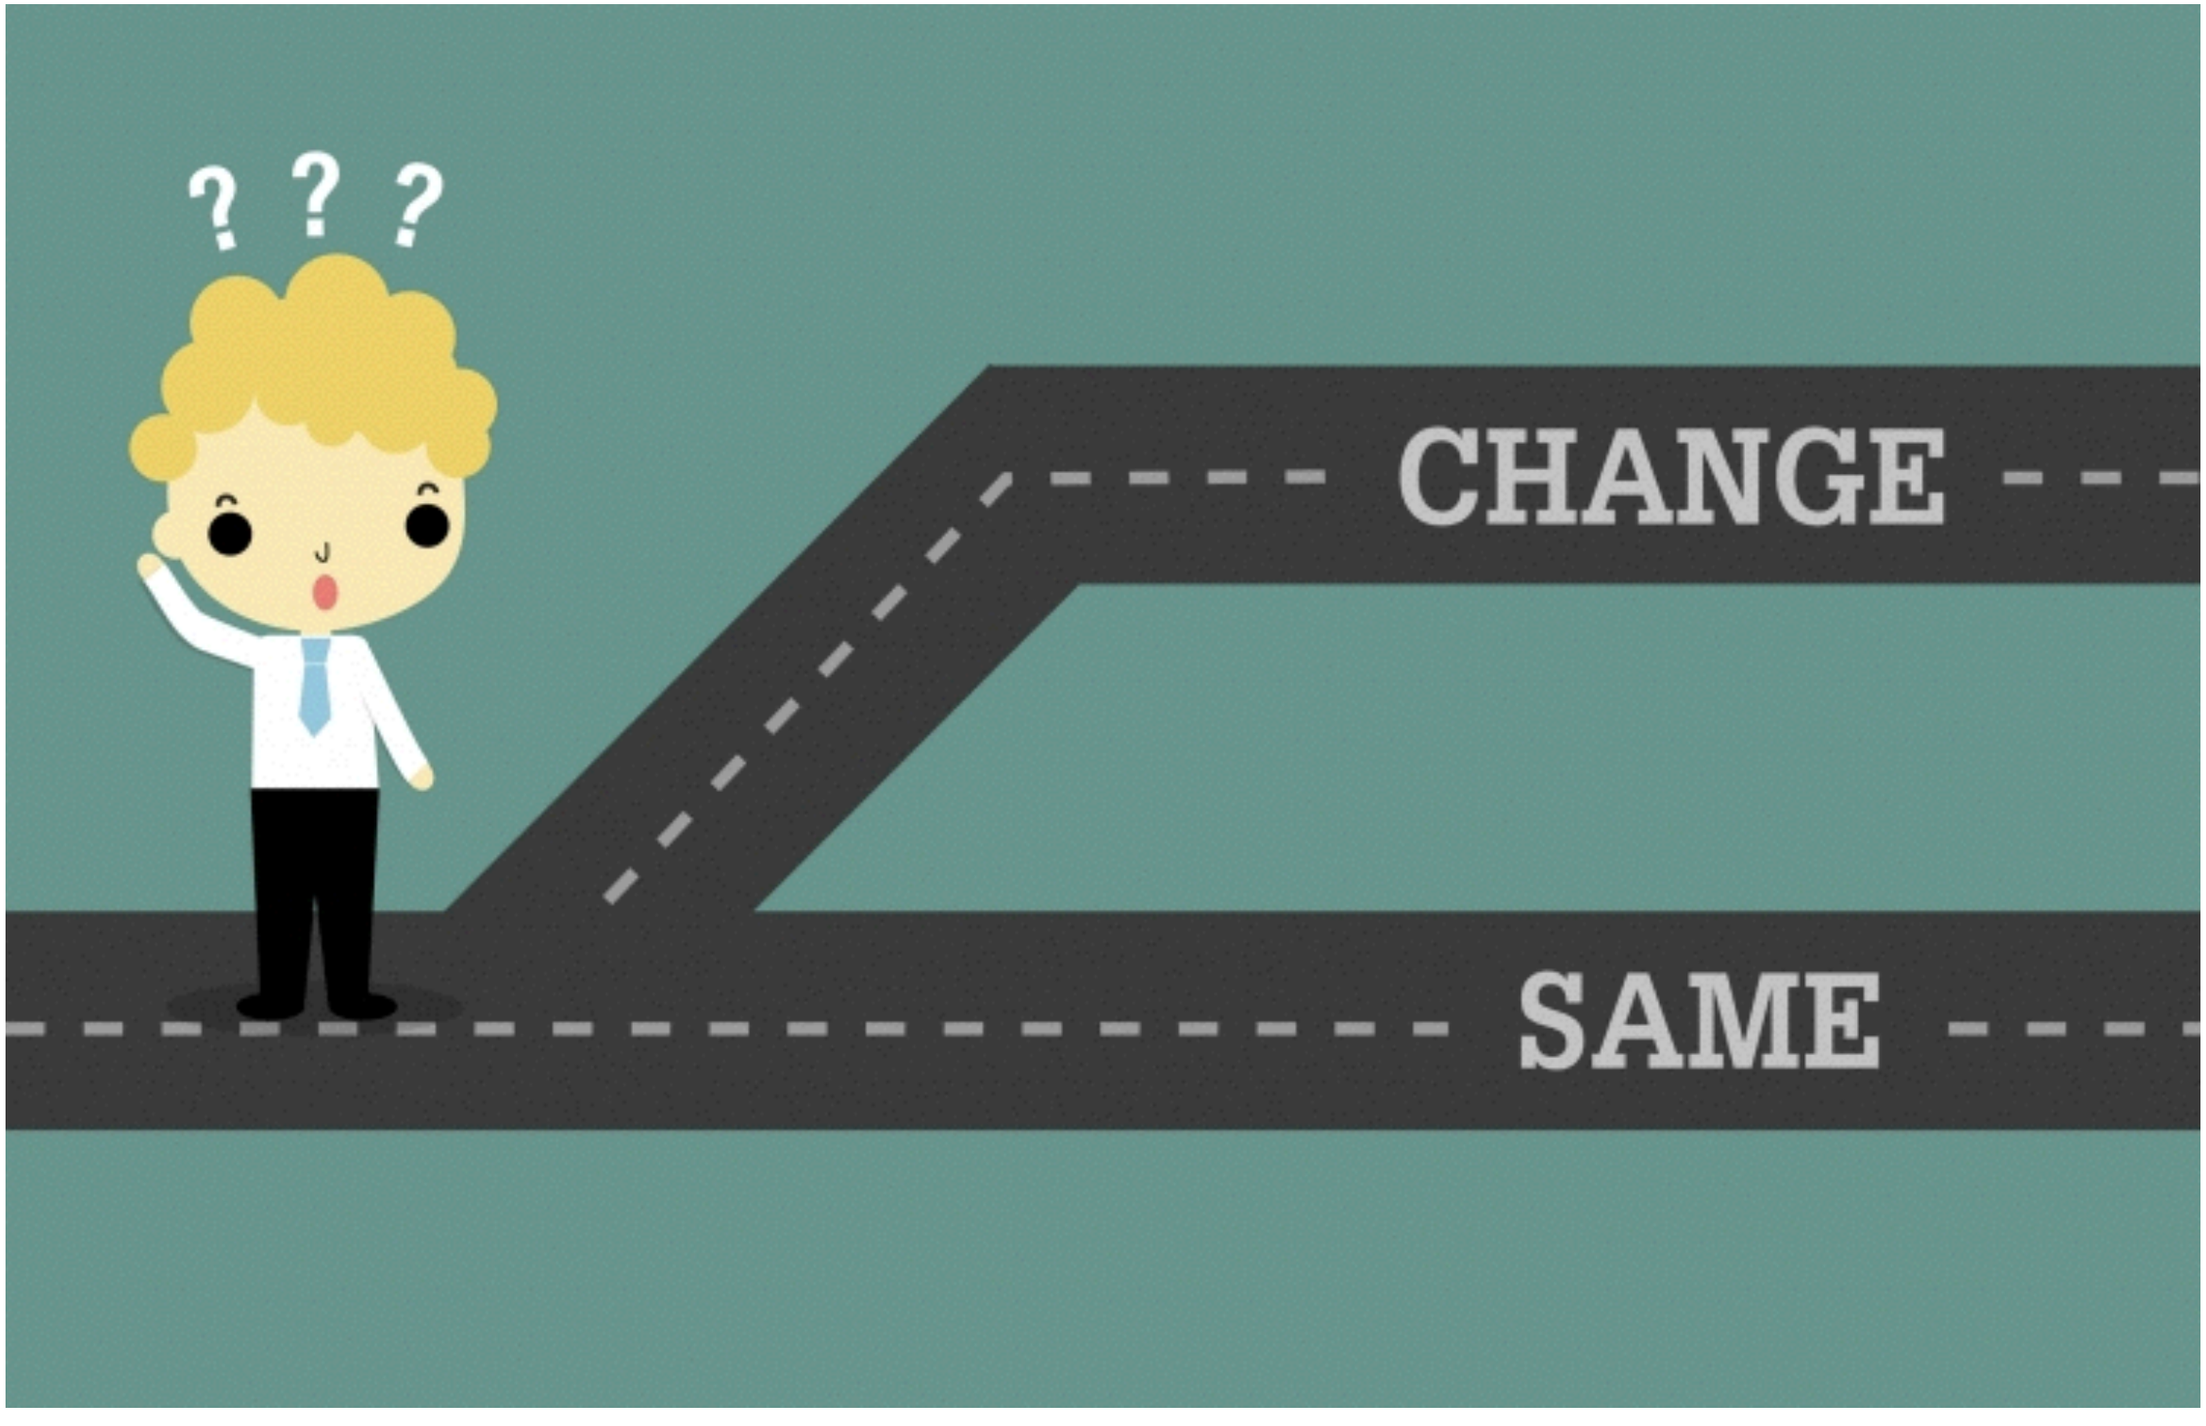
? ? ?
CHANGE
SAME

Thank you! Questions: ssaria@cs.jhu.edu I @suchisaria