Predicting intervention onset in the ICU with switching state space models

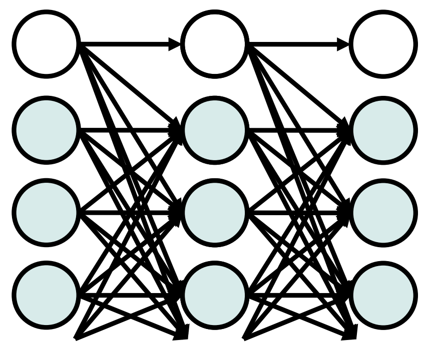

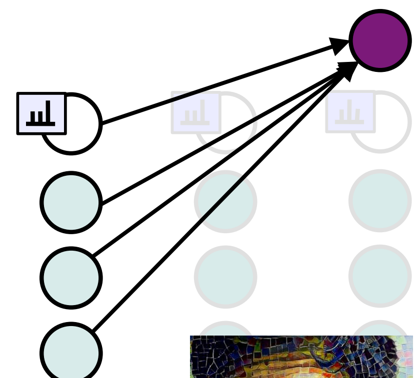

Marzyeh Ghassemi, Mike Wu,
Michael C. Hughes, Peter Szolovits,
and Finale Doshi-Velez

MIT, Yale University, and **Harvard University**

Problem: When will ICU patient need ***intervention***?

e.g.

mechanical ventilation

vasopressor (blood pressure drug)

or fluid transfusion

Early prediction helps:

prepare patient

plan staffing

try less aggressive options early

Possible Approaches

What to predict?

– lots of work on general risk scores

  - mortality, SAPS, APACHE

– less work on **actionable interventions**

How to represent patient state?

hand-engineered features

continuous-state temporal models

discrete switching-state temporal models

Lehman et al. 2015
Caballero Barajas et al. 2014

Contribution

We show that an

unsupervised auto-regressive Markov model

trained on

large cohort of 36,000 patients

can improve predictions for

5 interventions several hours ahead

mechanical ventilation	*red blood cell transfusion*
vasopressor	*plasma transfusion*
	platelet transfusion

Cohort from MIMIC-III dataset

mimic.physionet.org

(Johnson et al. Sci. Data 2016)

36,050 patients

recorded at Beth-Israel Deaconess in Boston between 2001-2012

kept all adults with record within 6-360 hours

Intervention	Training Num Positive	Training Num Control	Heldout Num Positive	Heldout Num Control
Vasopressor	6987	21865	1737	5461
Red blood cell transfusion	19171	9681	4776	2422
Fresh frozen plasma transfusion	2759	26093	620	6578
Platelet transfusion	27818	1034	6944	254
Mechanical Ventilation	13710	15142	3393	3805

Observed data

7 nurse-validated vital signs (hourly)

heart rate, blood pressure, temp., SpO2, ...

11 lab measurements (much less than hourly)

hematocrit, lactate, ...

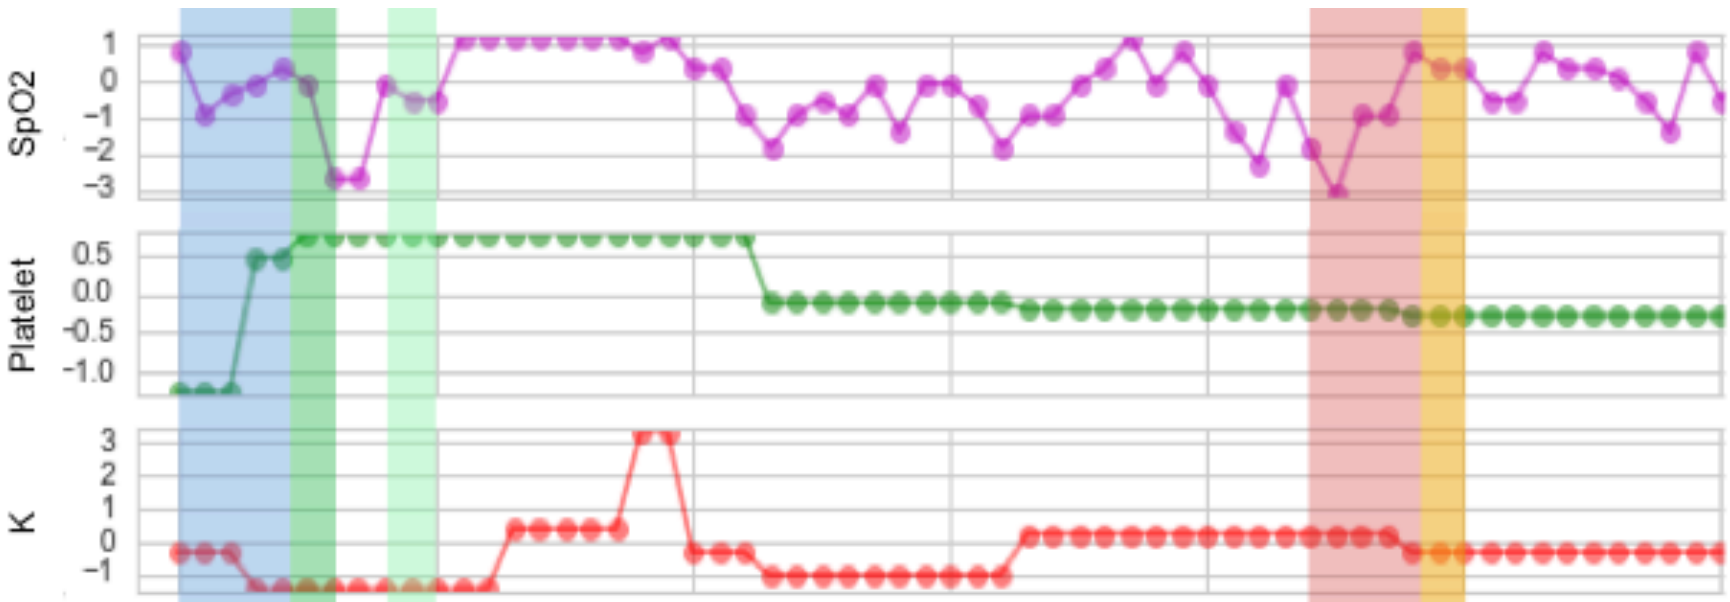

each channel standardized to mean=0, var=1 with carry-and-hold for missing data

Switching Autoregressive Model

Latent State

one of K possible values

Observed Vitals

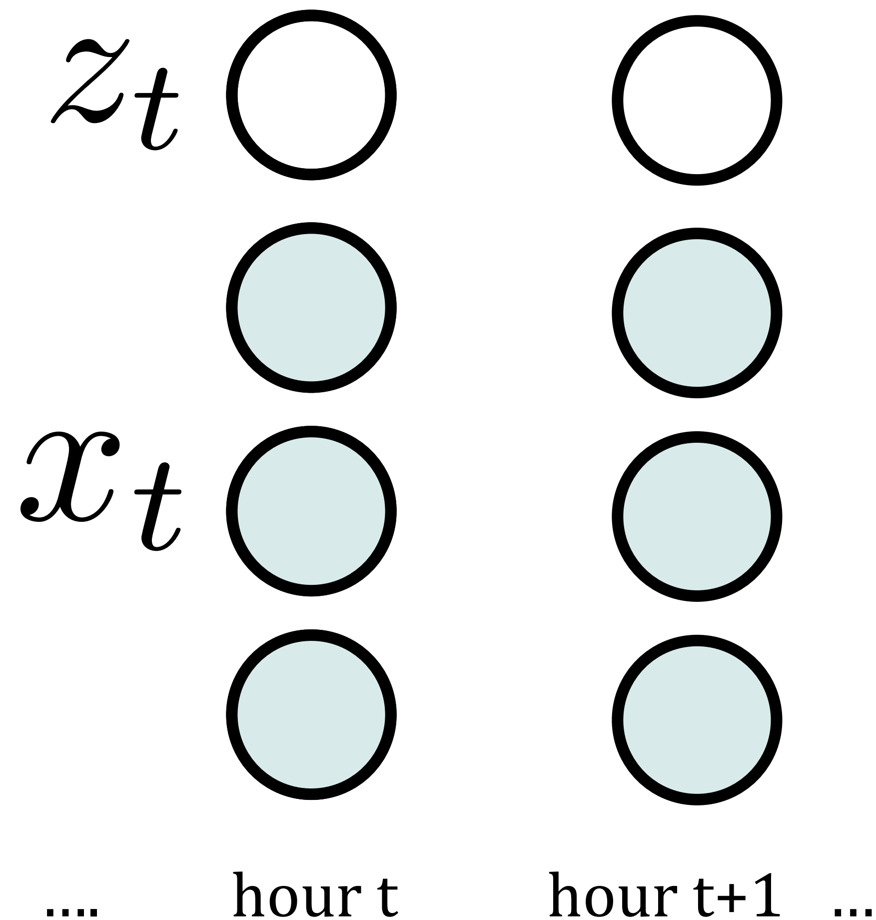

Switching Autoregressive Model

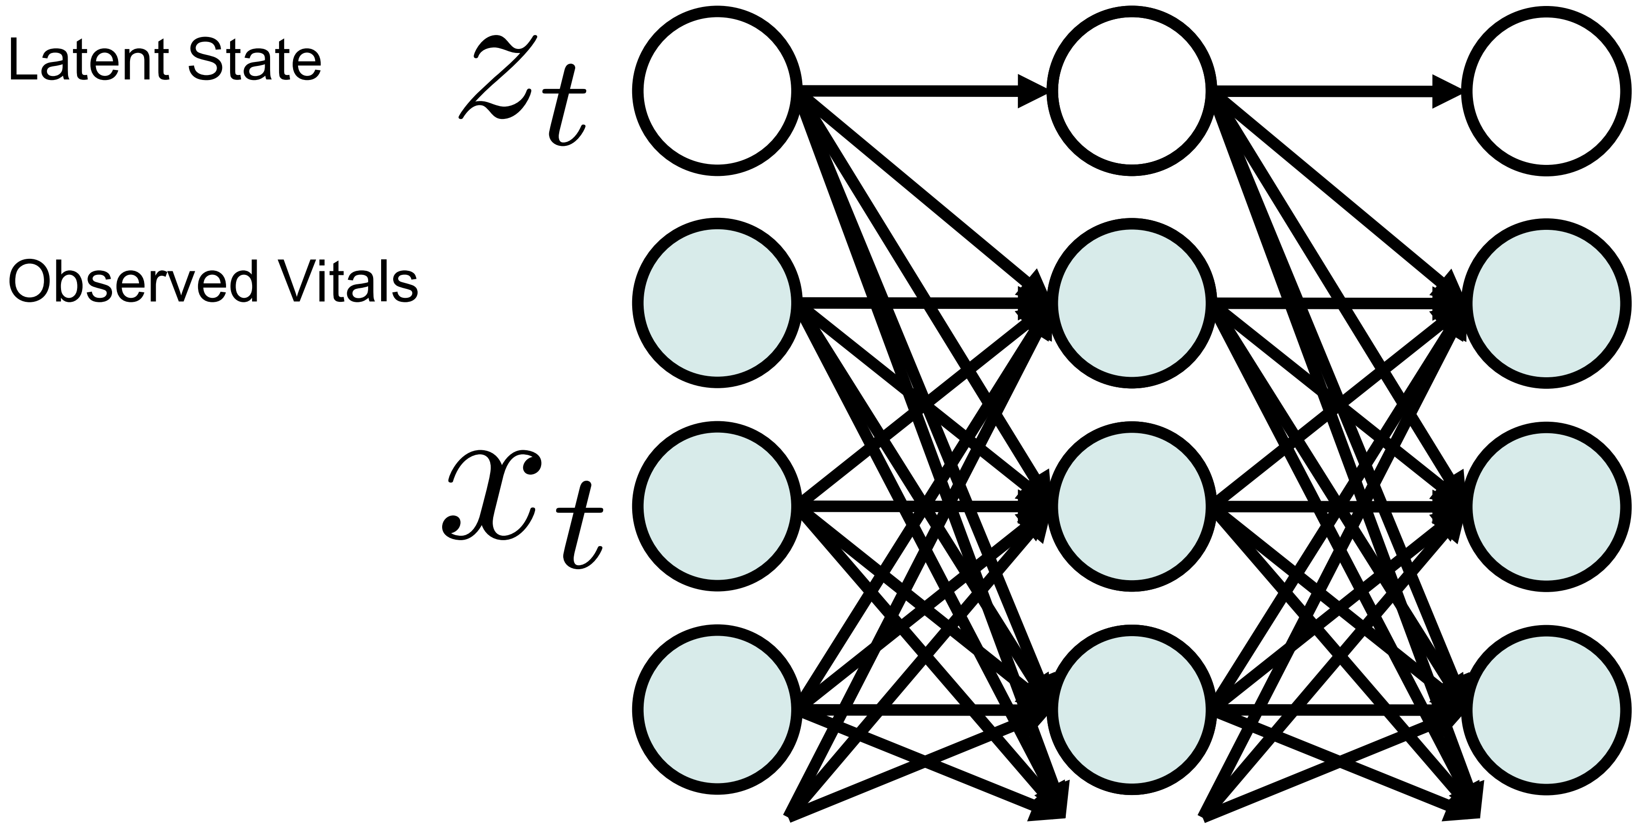

$$x_t | z_t = k \sim \mathcal{N}(A_k x_{t-1} + \mu_k, \Sigma_k)$$

autoregressive Gaussian allows modeling trajectories/trends in vitals

Training Phase

Learn model parameters from many patients

variational EM algorithm

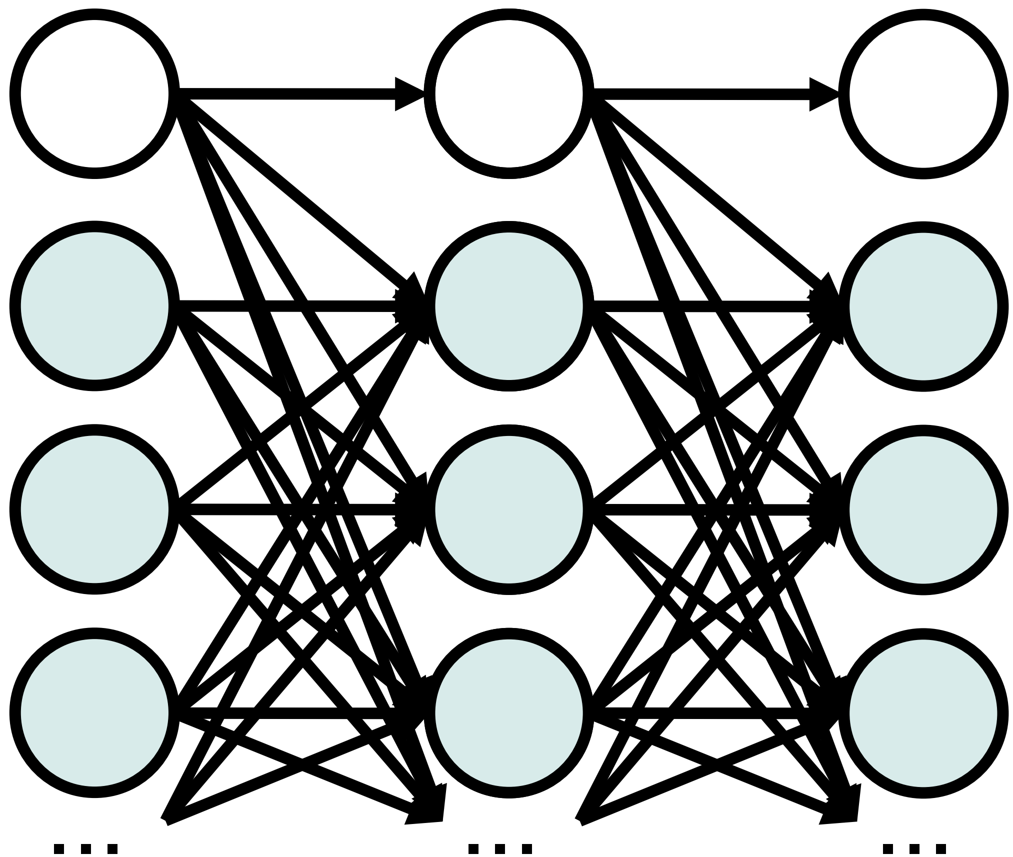

Prediction Step 1: Belief features

Infer distribution over hidden states at each timestep

HMM dynamic programming (forward alg.)

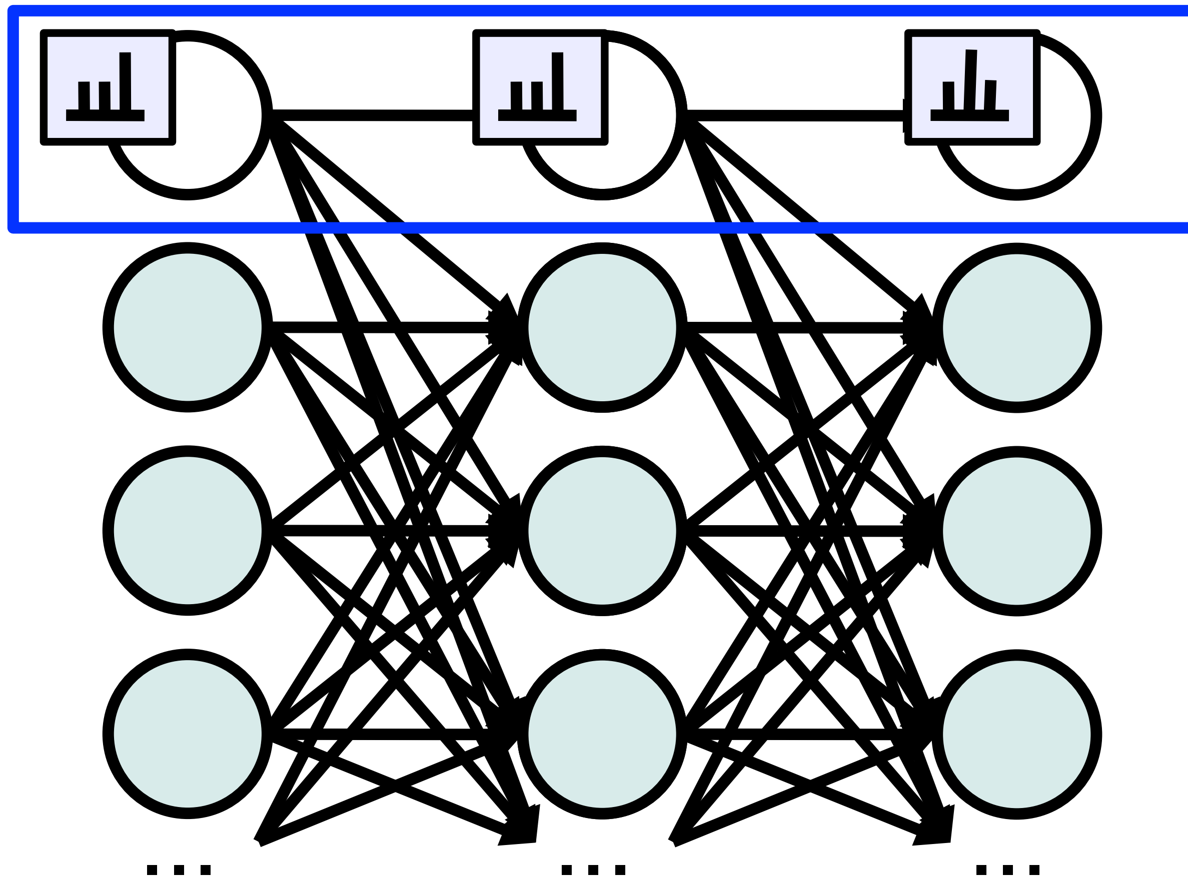

Step 2: Classify given features

Binary Intervention
(did ventilate at hour *t*)

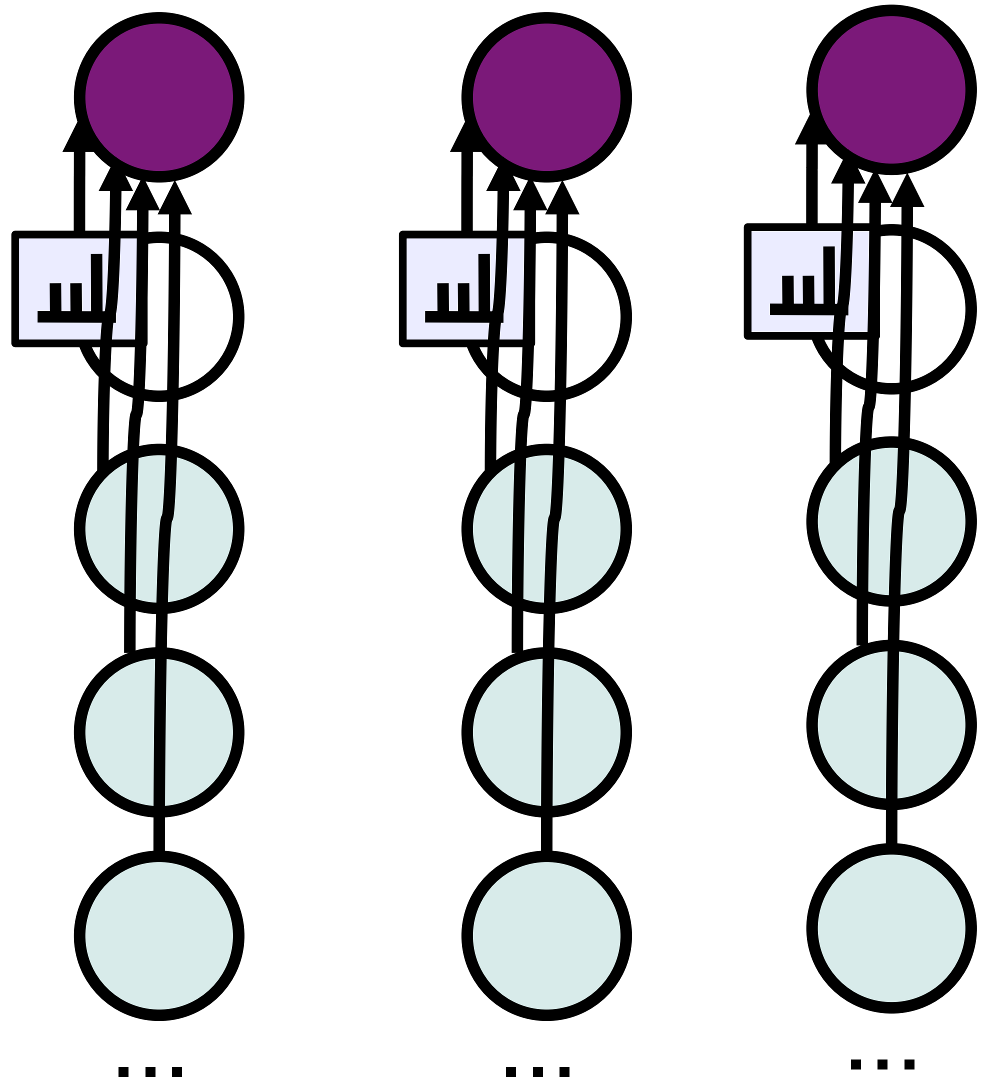

Logistic regression
(with label-balanced cost function)

Task: predict onset **in advance**

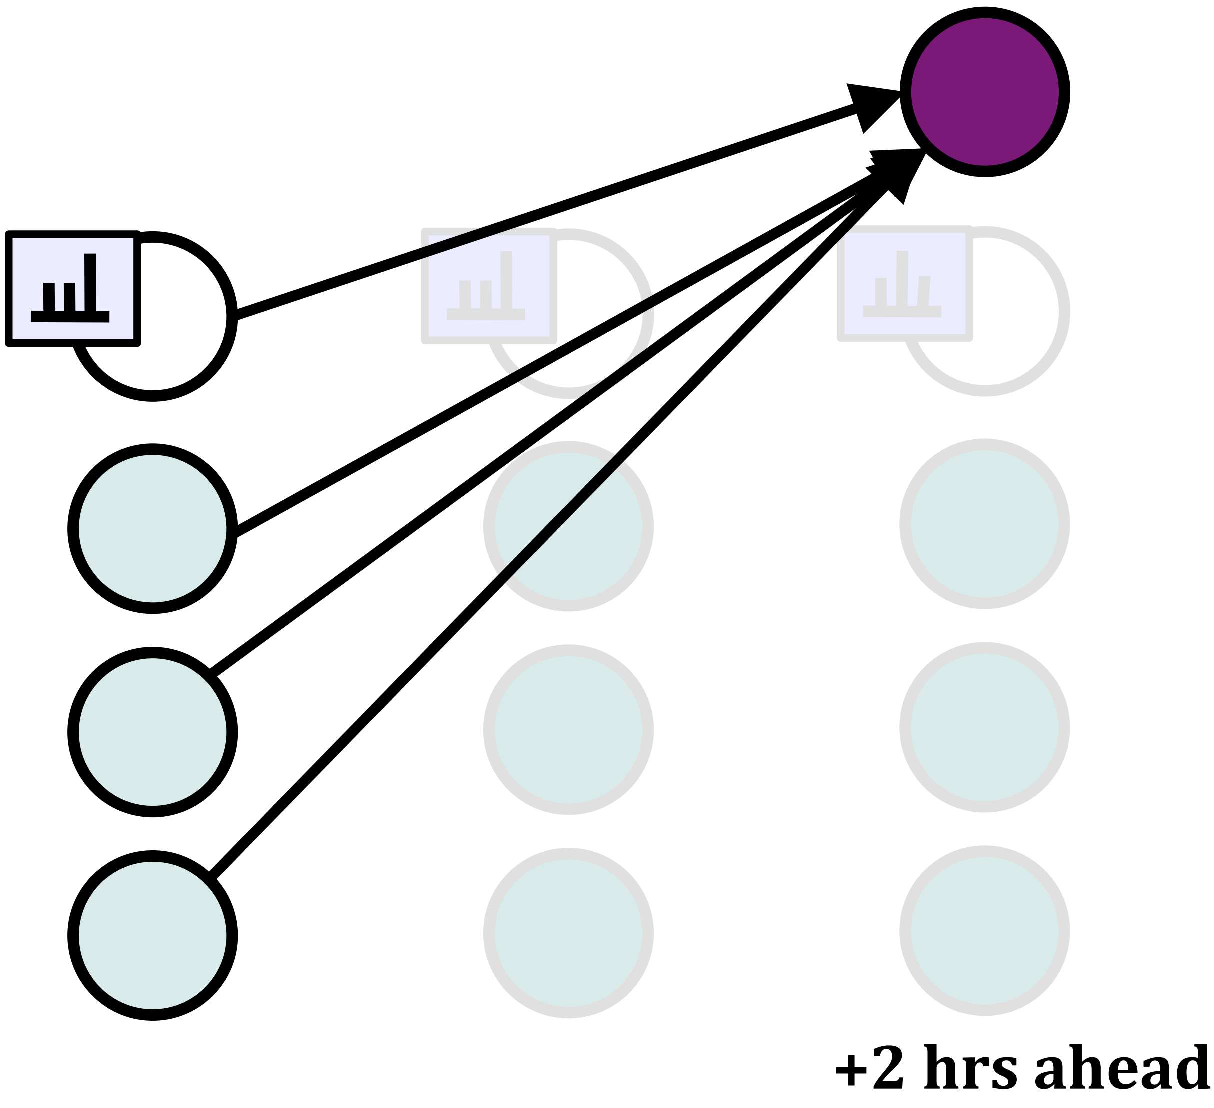

Vasopressor prediction : 1 hr ahead

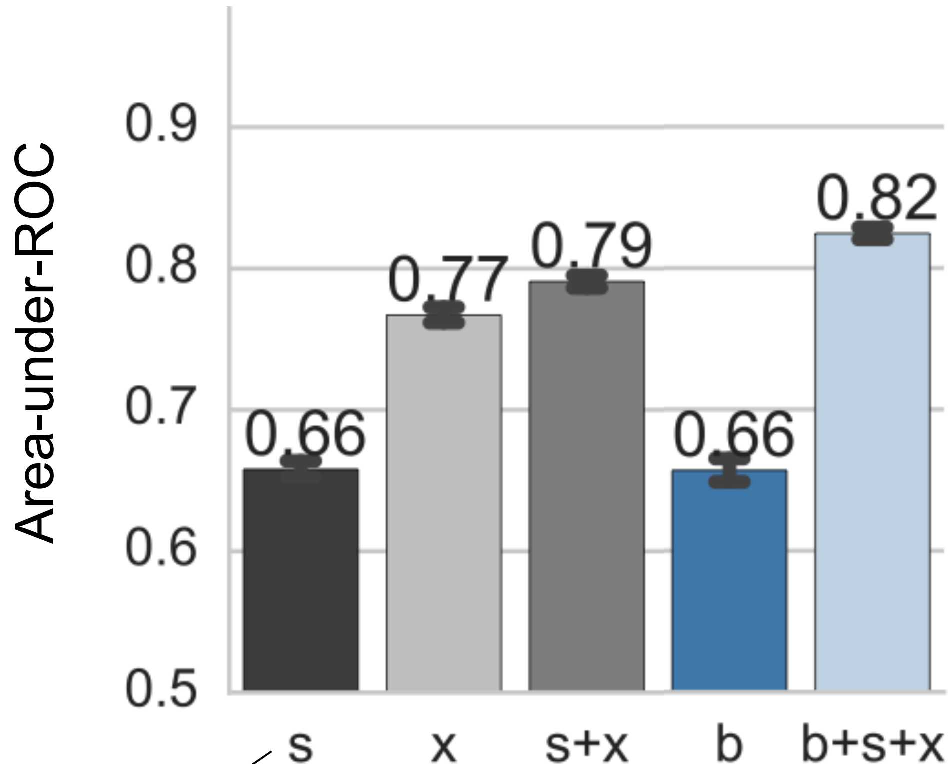

static demographics

dynamic patient vitals at time t

SSAM belief vector at time t using 10 states

Vasopressor prediction : 4 hr ahead

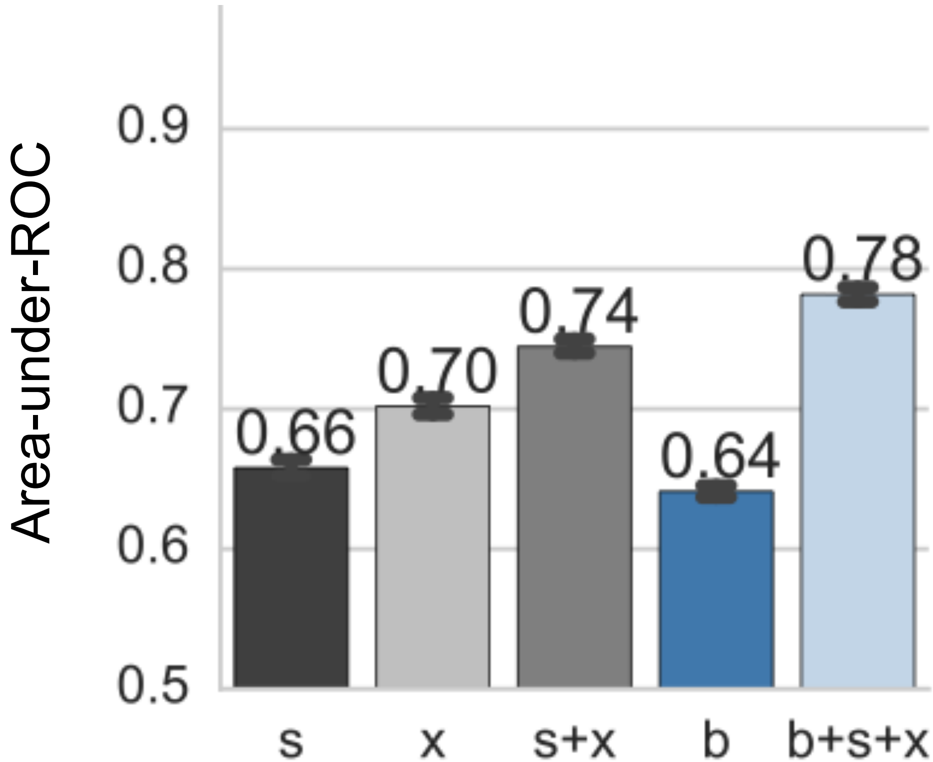

Ventilator : 4 hr ahead

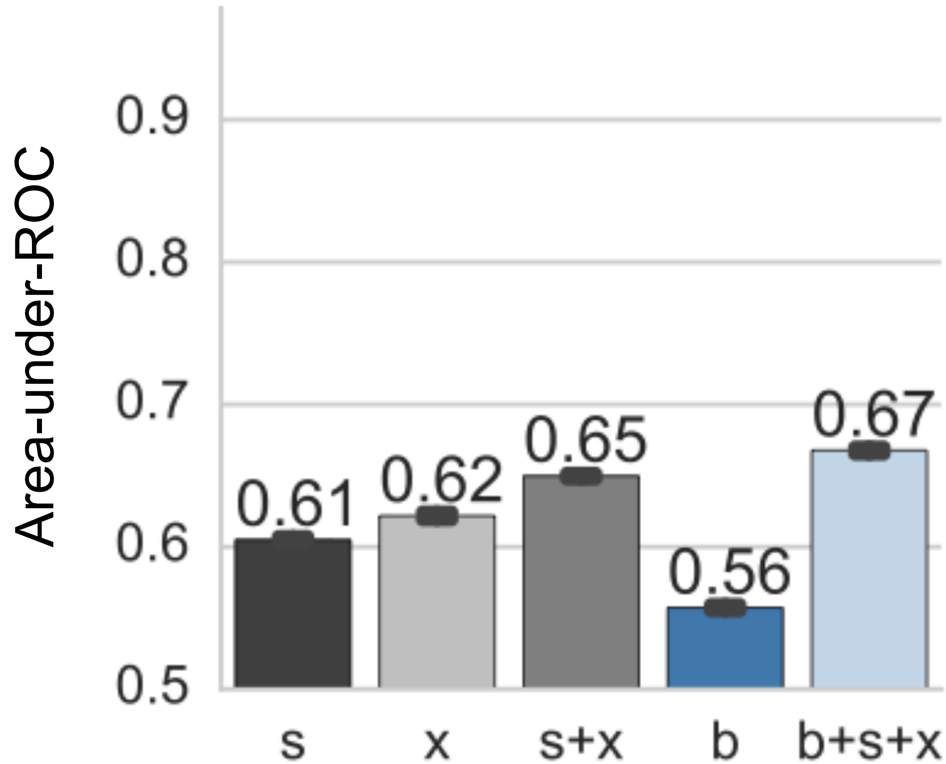

Fresh Frozen Plasma : 4 hr ahead

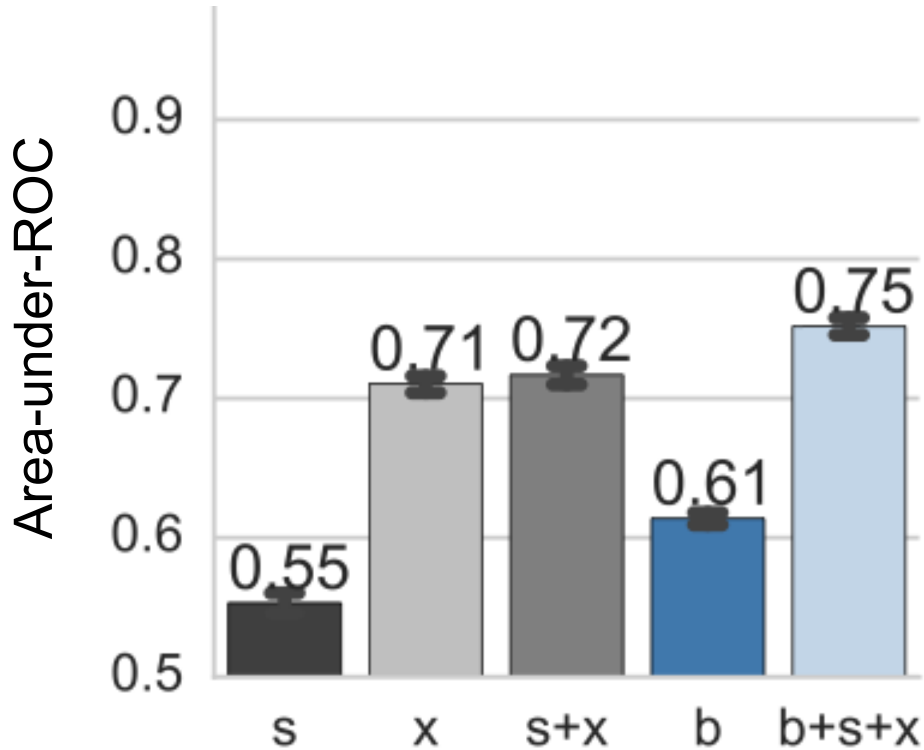

Interpreting Latent States

Inspect classifier weights across all 10 states

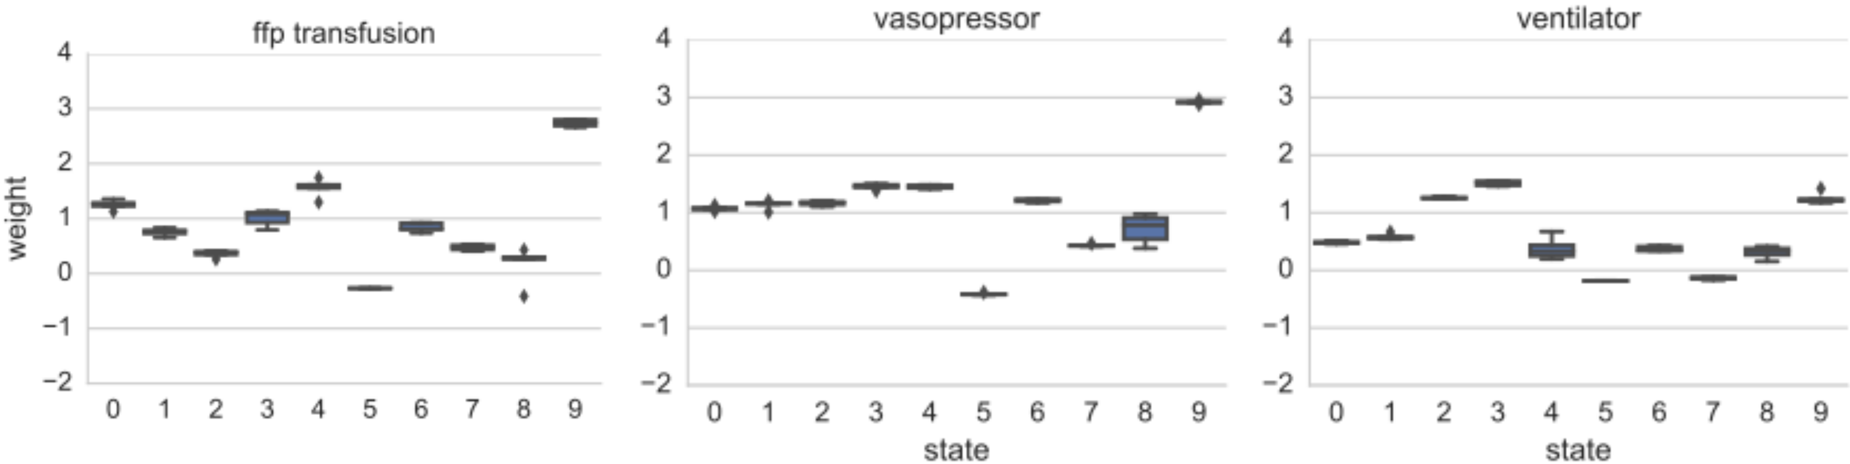

Inspect data associated with belief state 9

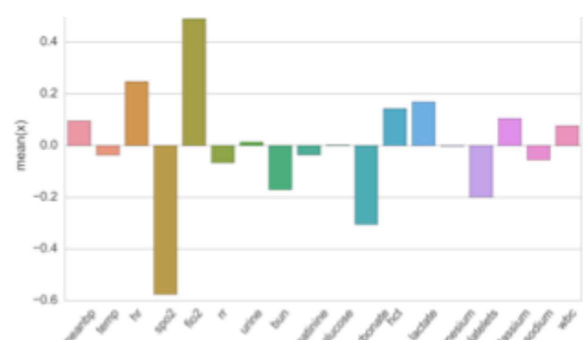

increased lactate,
lowered SpO2 and bicarbonate

Conclusion: state 9 seems to capture general physiological decline

Future Directions

Can we optimize generative models for particular downstream tasks without losing (too much) generalization?

Compare to alternative representation learning

- auto-encoders
- RNNs, LSTMs, etc

Move towards **reinforcement learning** approach

Predicting intervention onset in the ICU with switching state space models

Marzyeh Ghassemi, Mike Wu, **Michael C. Hughes**
Peter Szolovits, and Finale Doshi-Velez

Summary

unsupervised auto-regressive Markov model
large cohort of 36,000 patients
improves prediction on 5 interventions **several hours ahead**

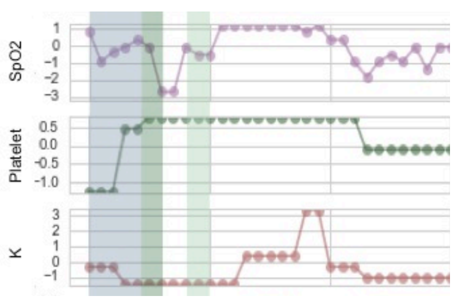

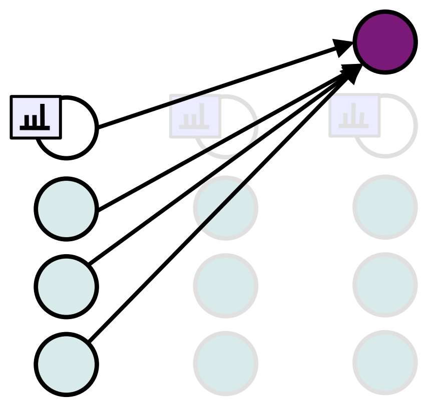

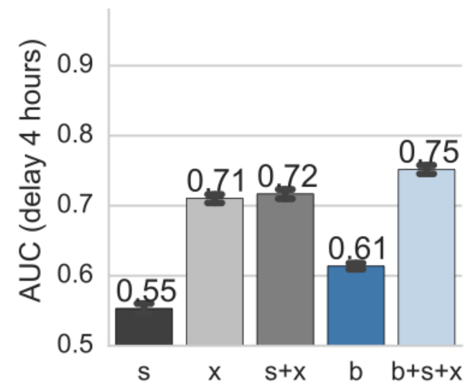

Acknowledgments

MCH supported by Oracle Labs
MG funded by the Intel Science and Technology Center for Big Data
and a National Library of Medicine Biomedical Informatics Research Training grant