Efficient Multi-Instance Learning for Activity Recognition from Time Series Data Using an Auto-Regressive Hidden Markov Model

Xinze Guan, Raviv Raich, Weng-Keen Wong
School of EECS, Oregon State University
Email: {guan,raich,wongwe@eecs.oregonstate.edu}

Introduction

- Wearable sensors are everywhere



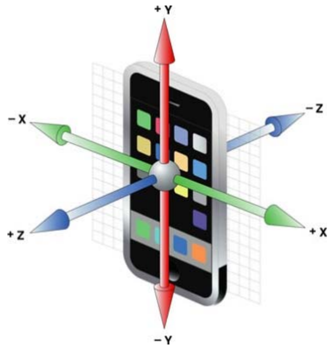

- Record human motion as a multivariate time series

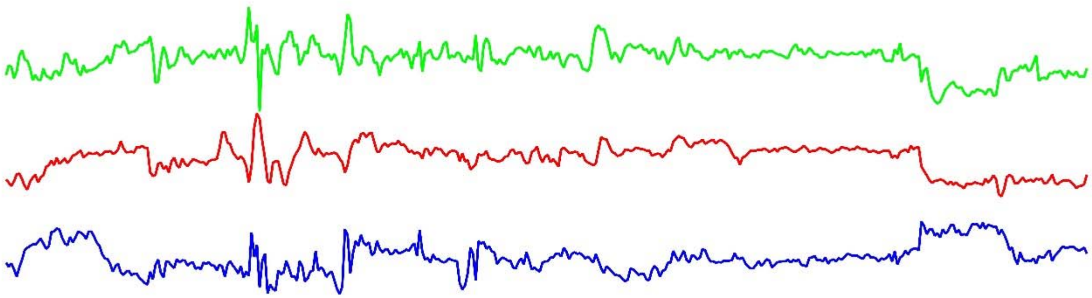

Introduction

- Goal: physical activity recognition

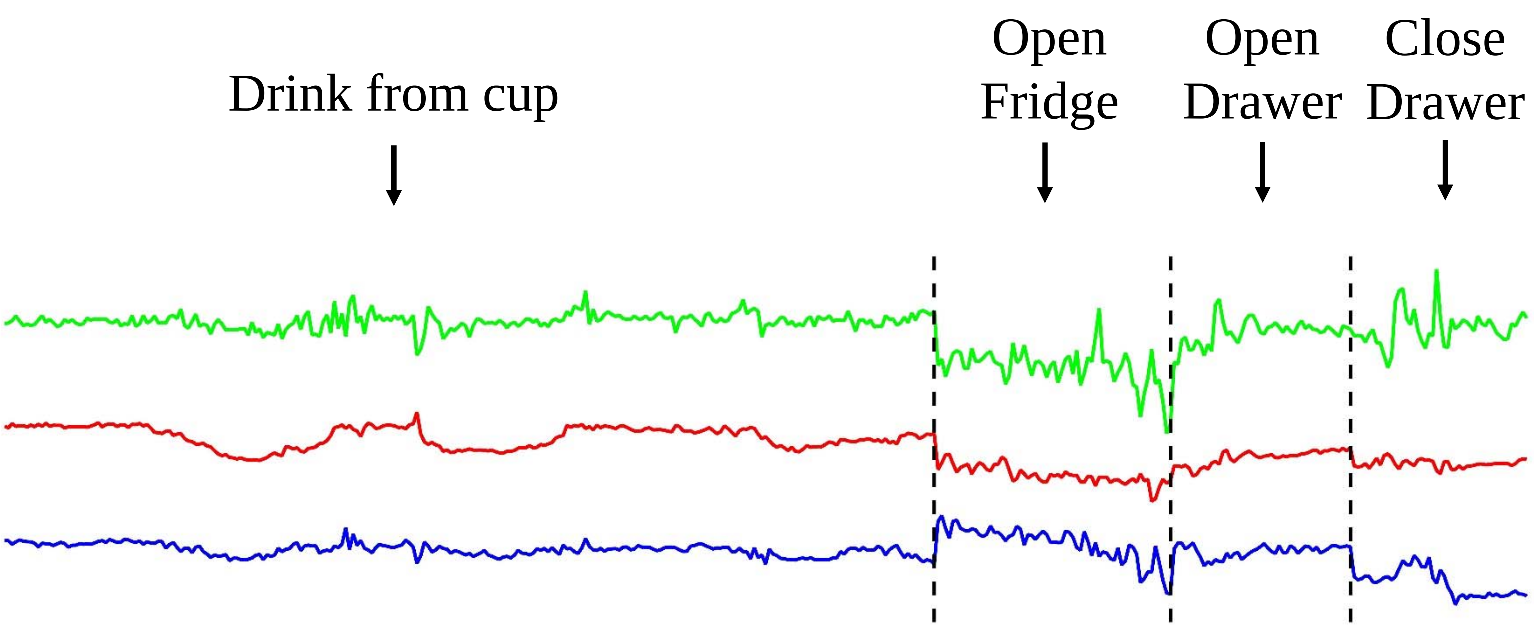

From the Opportunity dataset (Chavarriaga et al. 2013)

Introduction

Physical activity recognition important for:

Elder care

Assistance with cognitive disabilities

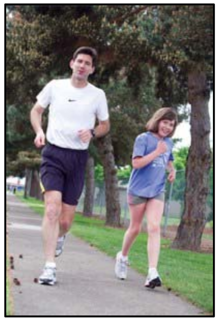

Health surveillance and research

Introduction

- Past work has typically applied standard supervised learning (eg. Bao and Intille 2004, Ravi et al. 2005, Zheng et al. 2013) or sequential approaches (Lester et al. 2005, van Kasteren et al. 2008, Wu et al. 2009)
- High annotation effort to label training data

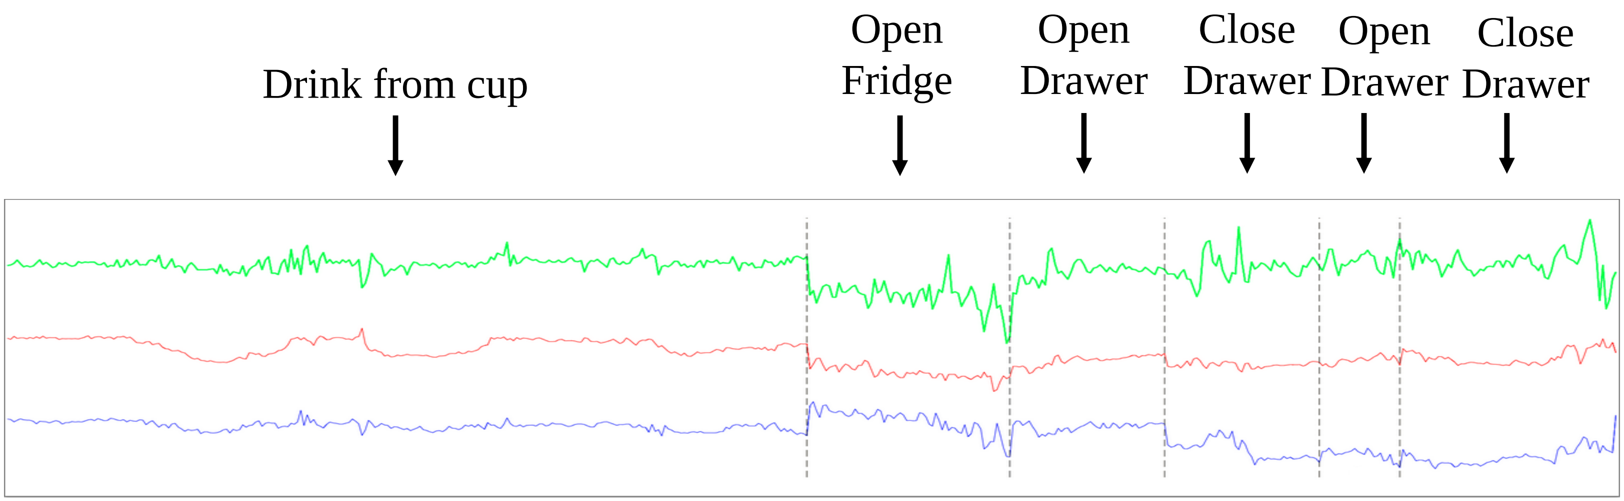

Introduction

- Stikic et al. (2011) proposed a weakly supervised approach based on multi-instance learning
- Trades off the ease of labeling with ambiguity in the labeling
- Our work builds on their approach

Methodology: MIL

Multi-instance Learning (Dietterich et al. 1997) :

- Data made up of bags of instances
- Bags can be labeled positive or negative

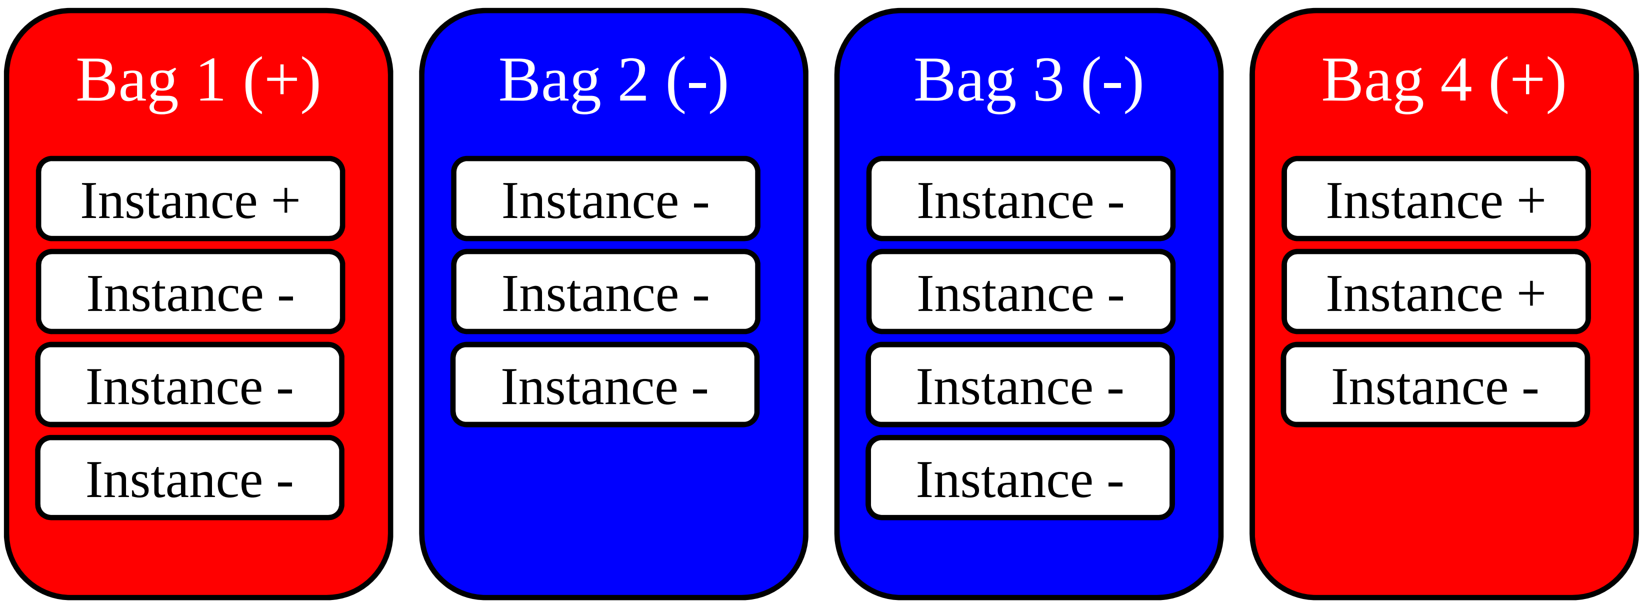

Methodology: MIL for Time Series

<u>Majority Labeling Scheme</u>:

Bag labeled + if the majority of the time ticks belong to the activity of interest (eg. "Drink from Cup")

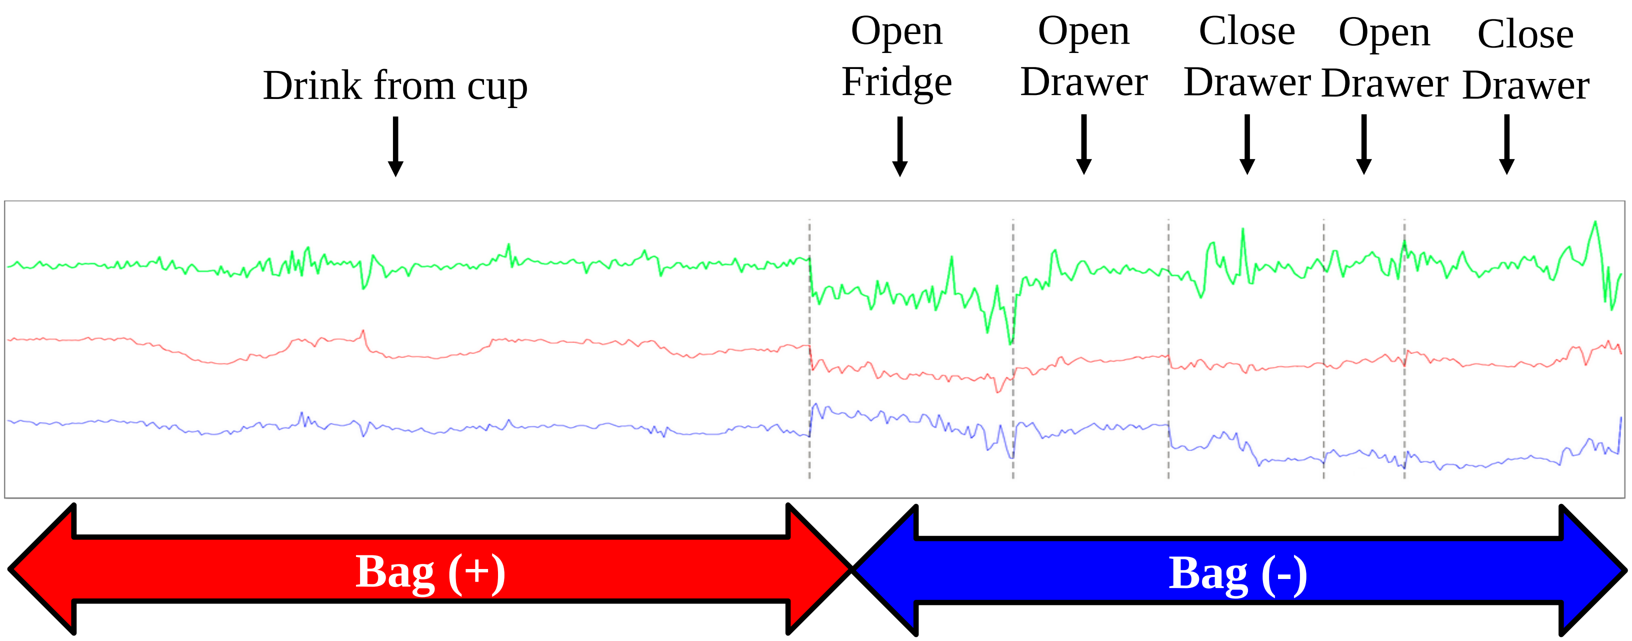

Related Work

Structured MIL

- Relationship between instances in a bag (Zhou et al. 2009, Warrell and Torr 2011)
- Relationship between instances in different bags (Deselaers and Ferrari 2010)
- Relationship between bags (Zhang et al. 2011)

Our work: models temporal dynamics between instances in a bag

Methodology: The Model

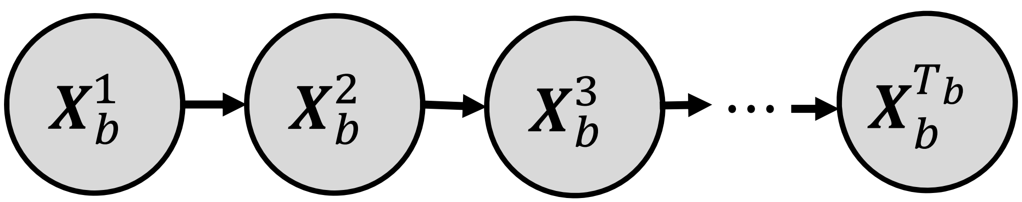

Methodology: The Model

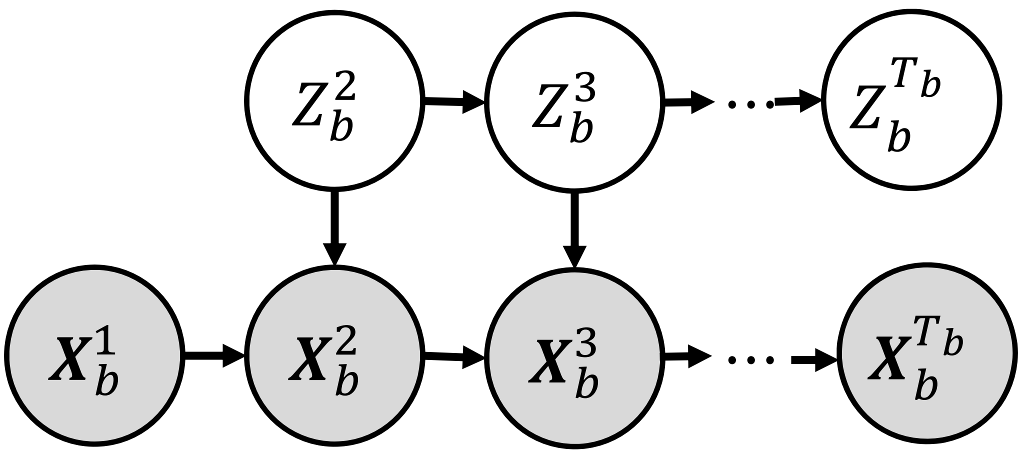

Methodology: The Model

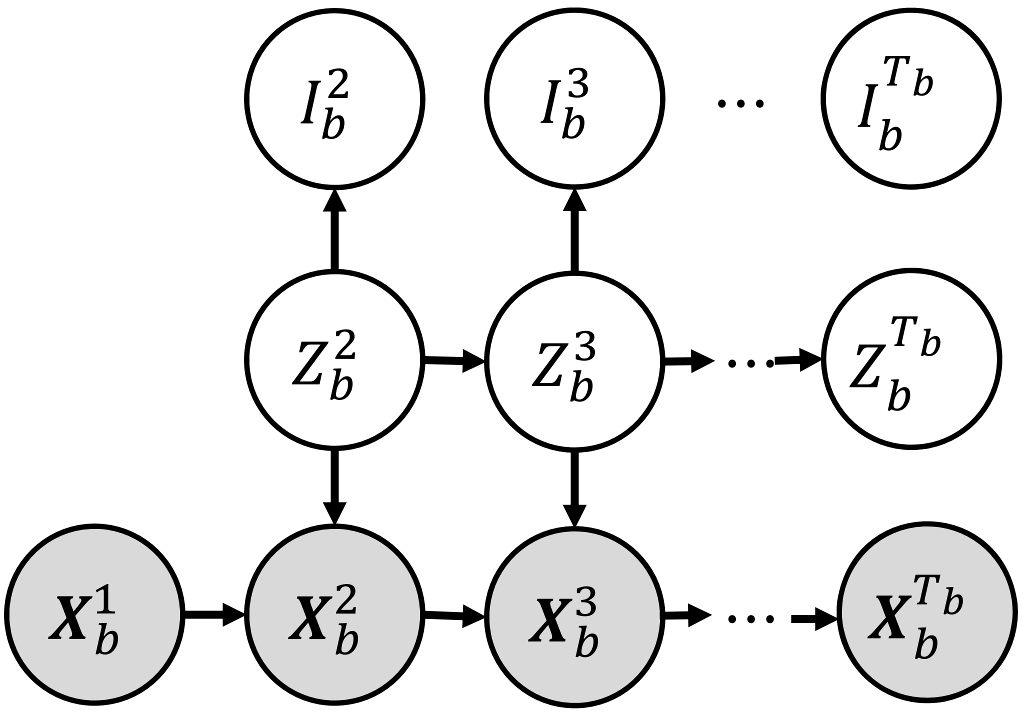

Methodology: The Model

$$Y_b = MIL\_Softmax(I_b^1, \ldots, I_b^{T_b})$$

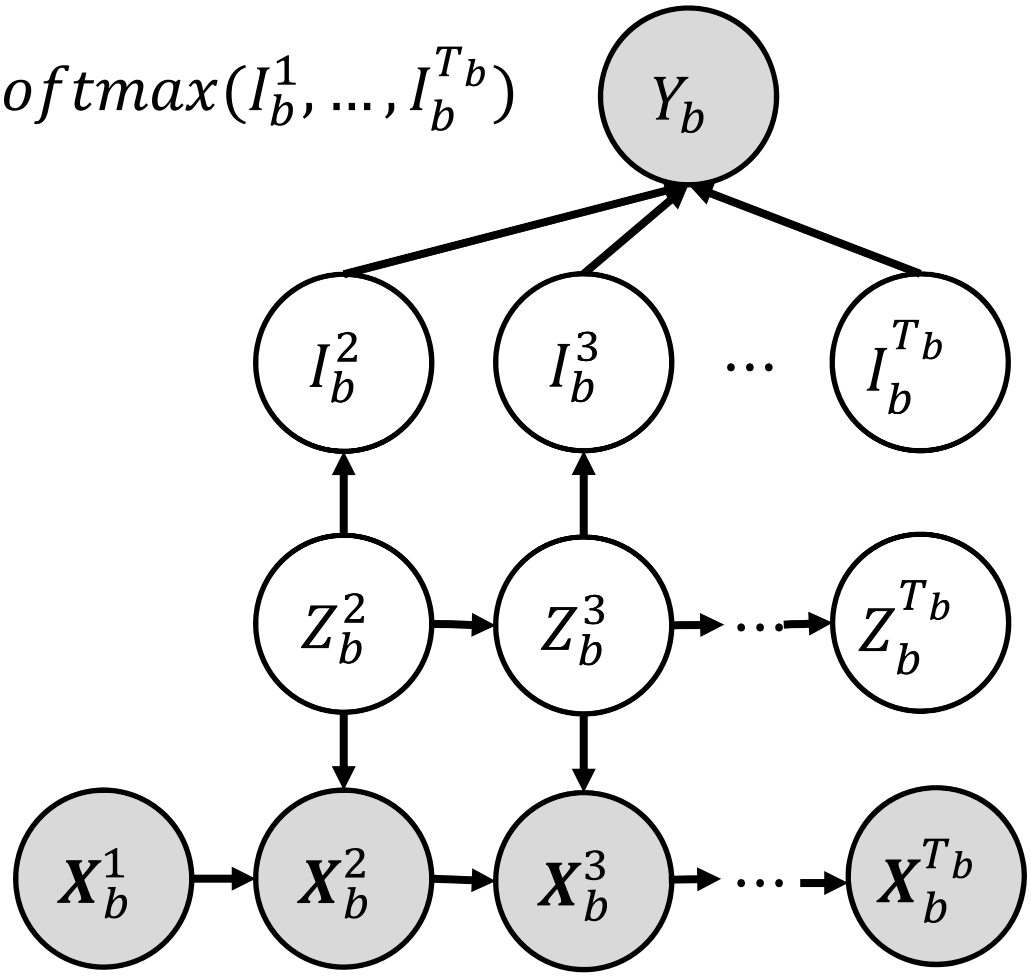

Methodology: The Model

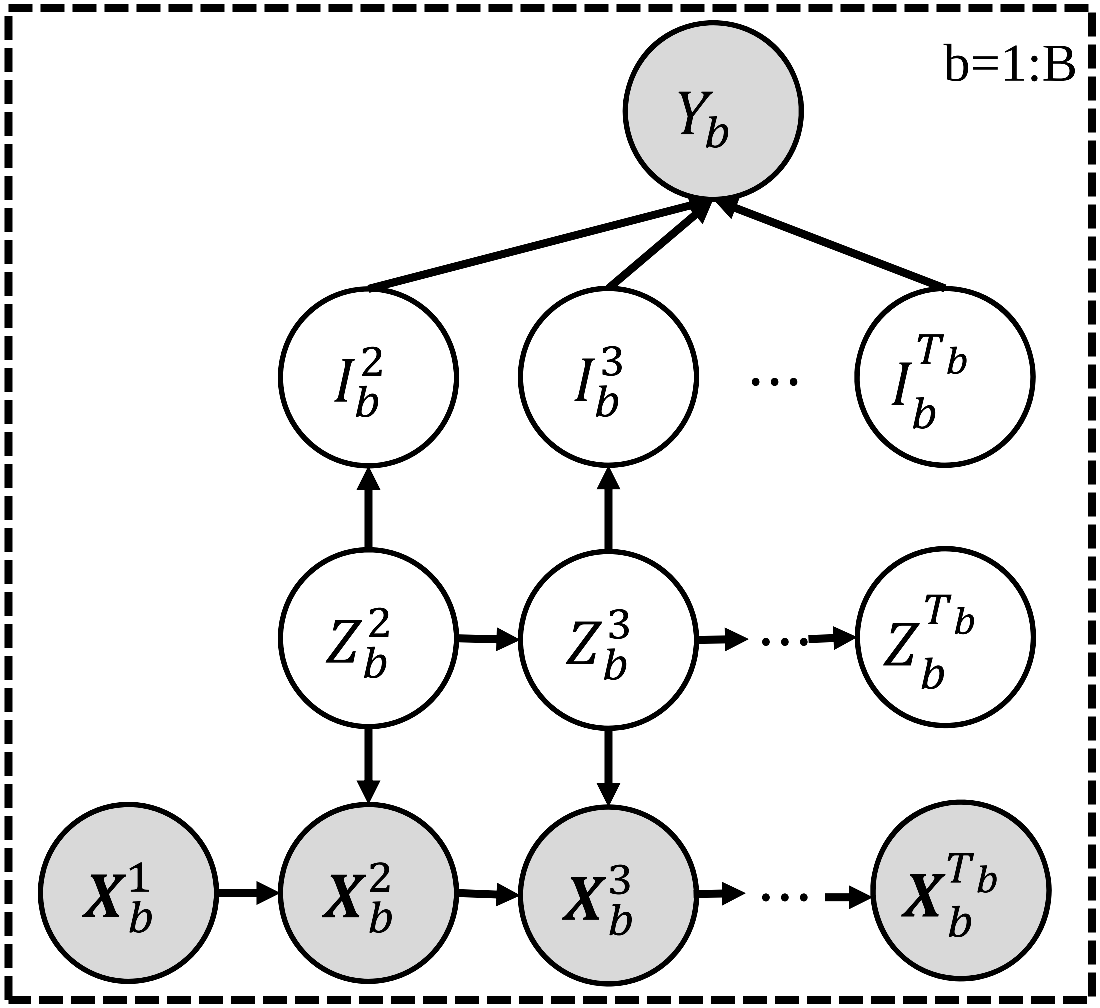

Methodology: The Model

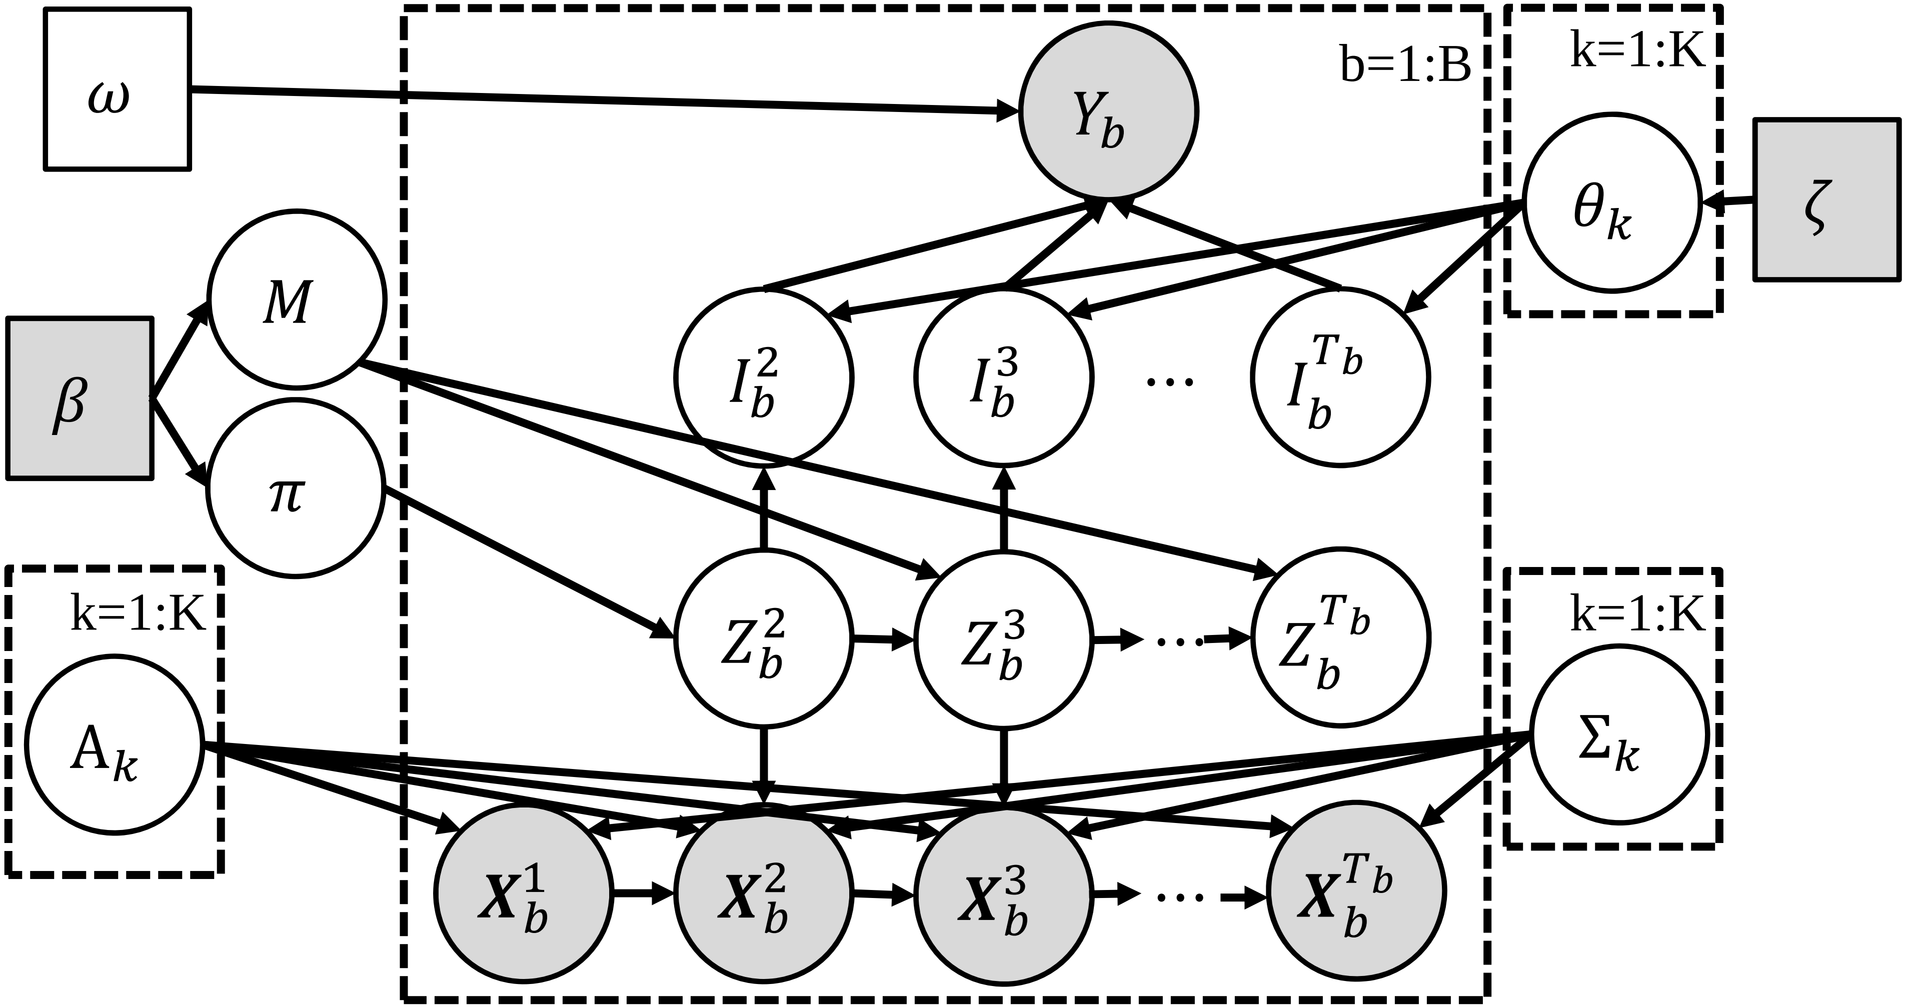

Methodology: Parameter Estimation

Expectation-Maximization:

1. M-step:
   - Straightforward
2. E-step:
   - Requires computation of
     $P\left(I_b^t, Z_b^t, I_b^{t-1}, Z_b^{t-1} | \boldsymbol{X}_b, Y_b\right)$
   - If done naively: $\mathrm{O}\left(2^{\mathrm{T_b}}\mathrm{K}^{\mathrm{T_b}}\right)$

Methodology: Efficient Message Passing

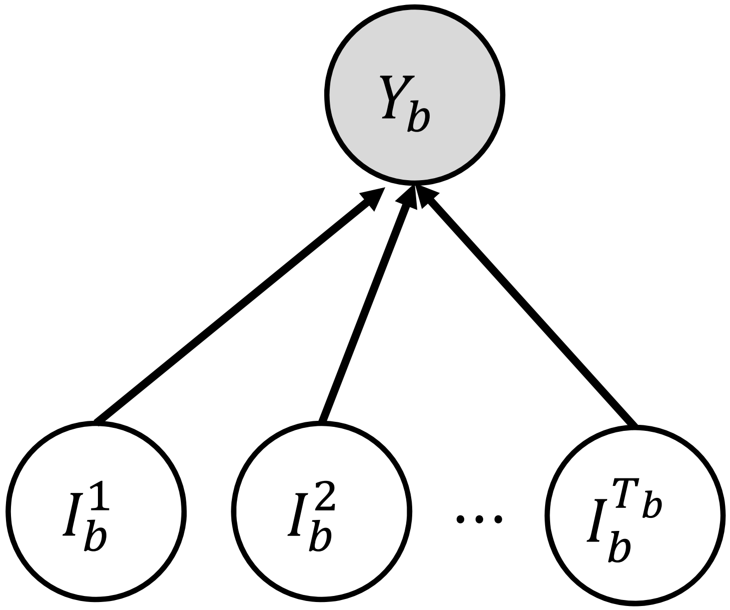

$$P\left(Y_b = 1 \middle| I_b^1, \ldots, I_b^{T_b}\right) = \frac{(\#\, positive\ instances) * \exp(\omega)}{(\#\, positive\ instances)\exp(\omega) + (\#\, negative\ instances)}$$

Methodology: Efficient Message Passing

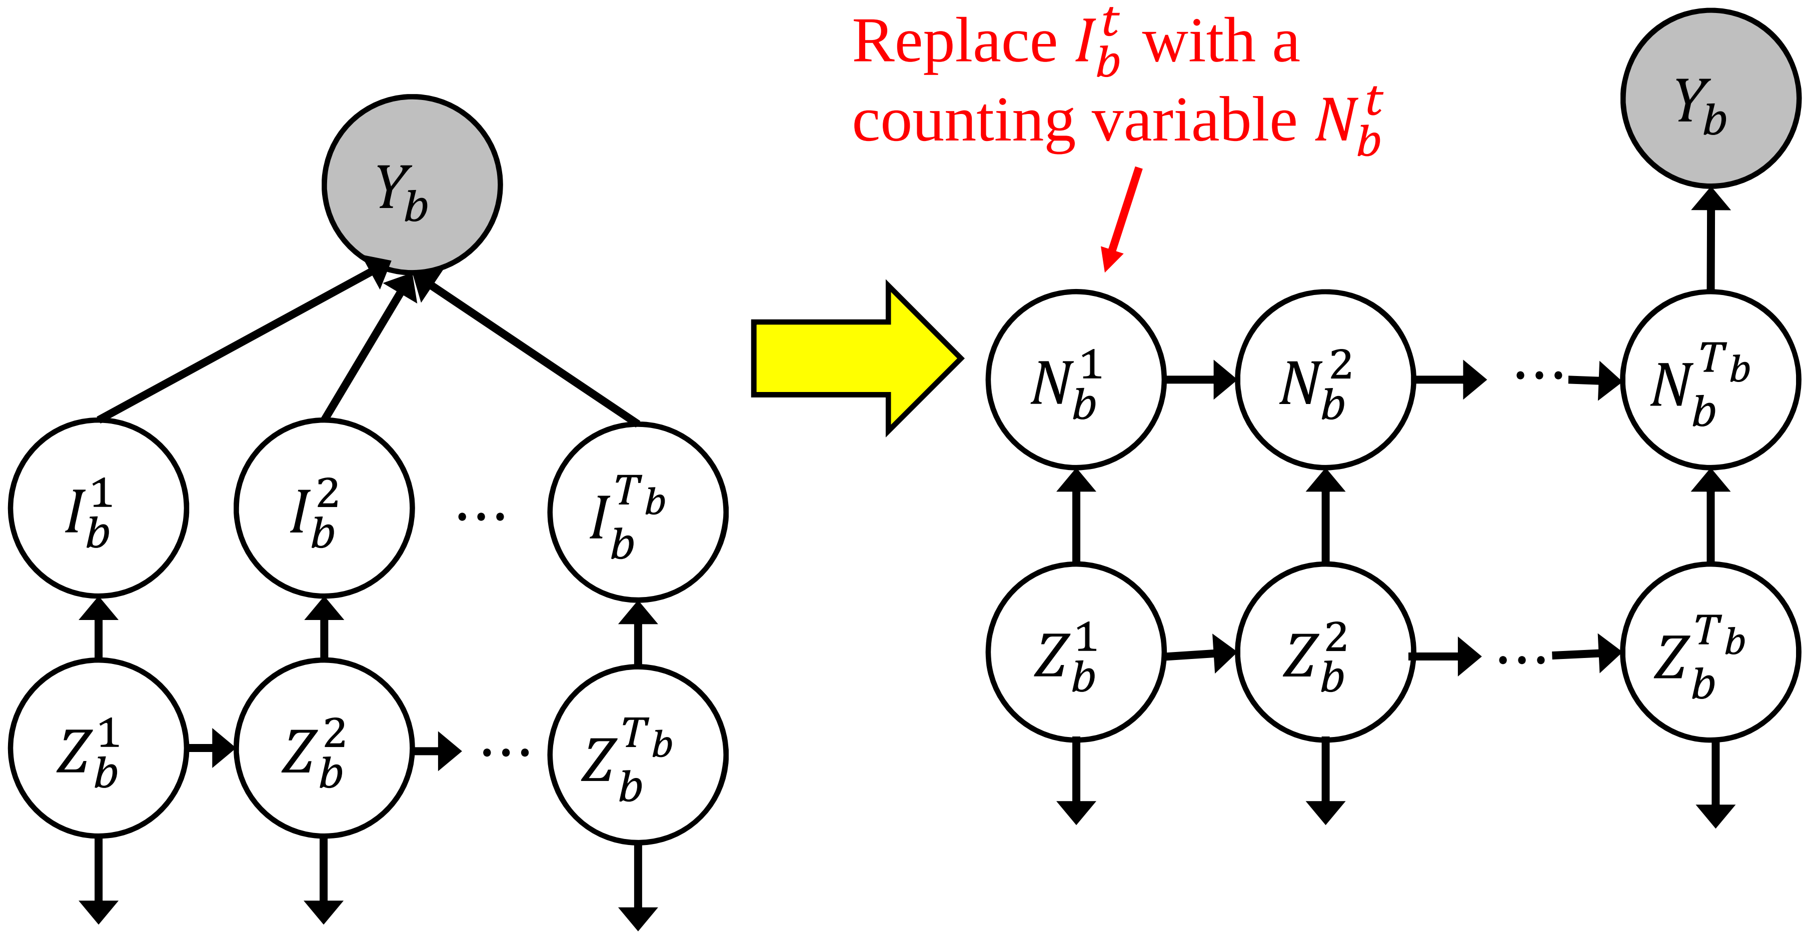

Methodology: Efficient Message Passing

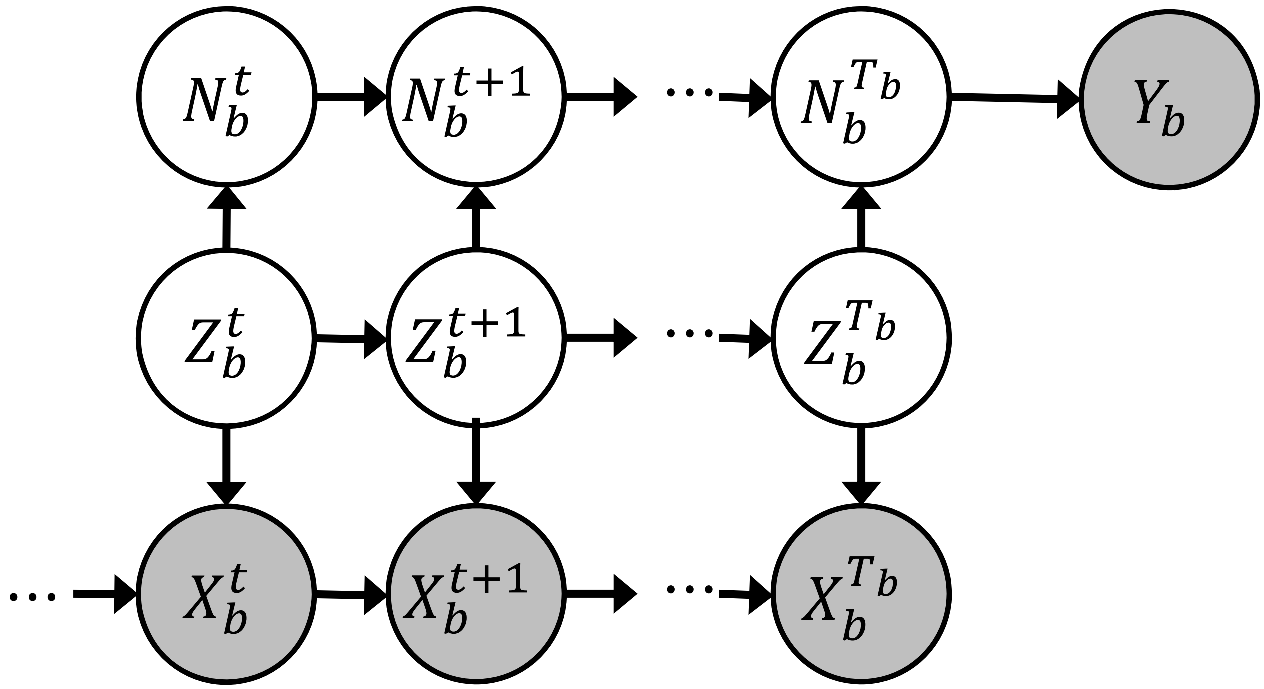

Methodology: Efficient Message Passing

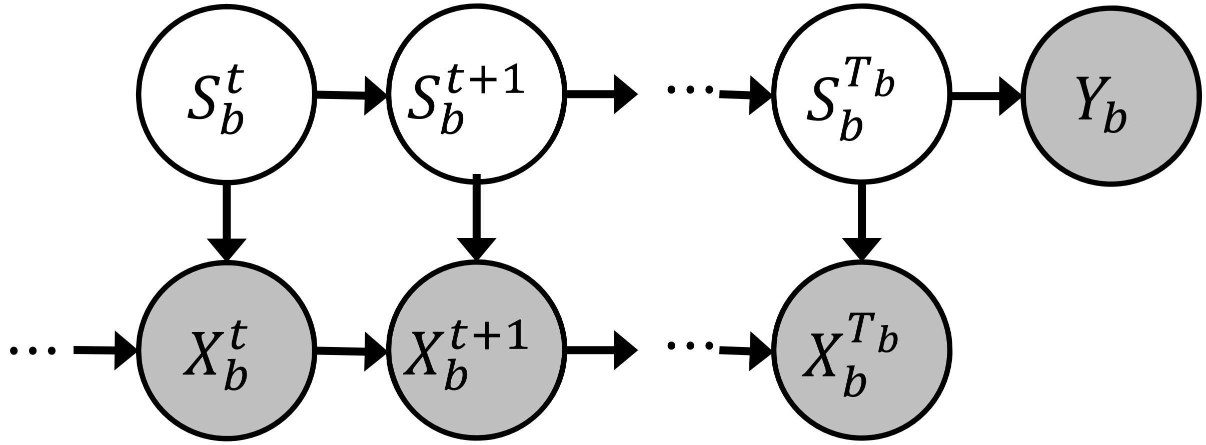

- Replace the N_b^t, Z_b^t nodes with a super-node $S_b^t = (N_b^t, Z_b^t)$
- Becomes an Auto-regressive Hidden Markov Model

Methodology: Efficient Message Passing

- Apply standard forward-backward message passing for ARHMM
- But can exploit a sparse transition matrix
- E-step computation is now $\mathrm{O}\left(K^2 {T_b}^2\right)$

Results: Algorithms

Using features from Stikic et al. (2011)

- miSVM (Andrew et al. 2003)
- DPMIL (Kandemir and Hamprecht 2014)
- miGraph (Zhou et al. 2009)

Using the raw time series:

- MARMIL (our NIPS workshop paper)
- ARHMM-MIL (ours)

Results: Experimental Setup

Datasets:

- Opportunity (Chavarriaga et al. 2013)
- Trainspotting1 (Berlin and Laerhoven 2012)
- Trainspotting2 (Berlin and Laerhoven 2012)

Results

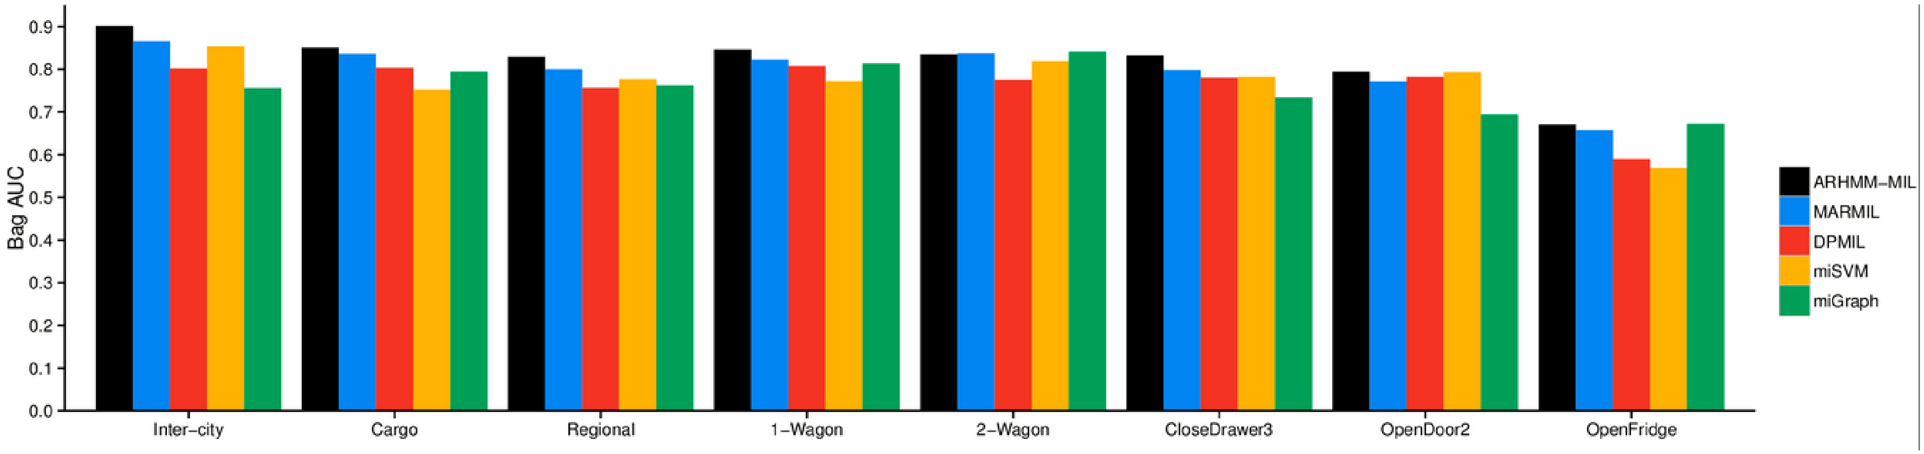

Bag AUC
0.0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
Inter–city
Cargo
Regional
1–Wagon
2–Wagon
CloseDrawer3
OpenDoor2
OpenFridge
ARHMM–MIL
MARMIL
DPMIL
miSVM
miGraph

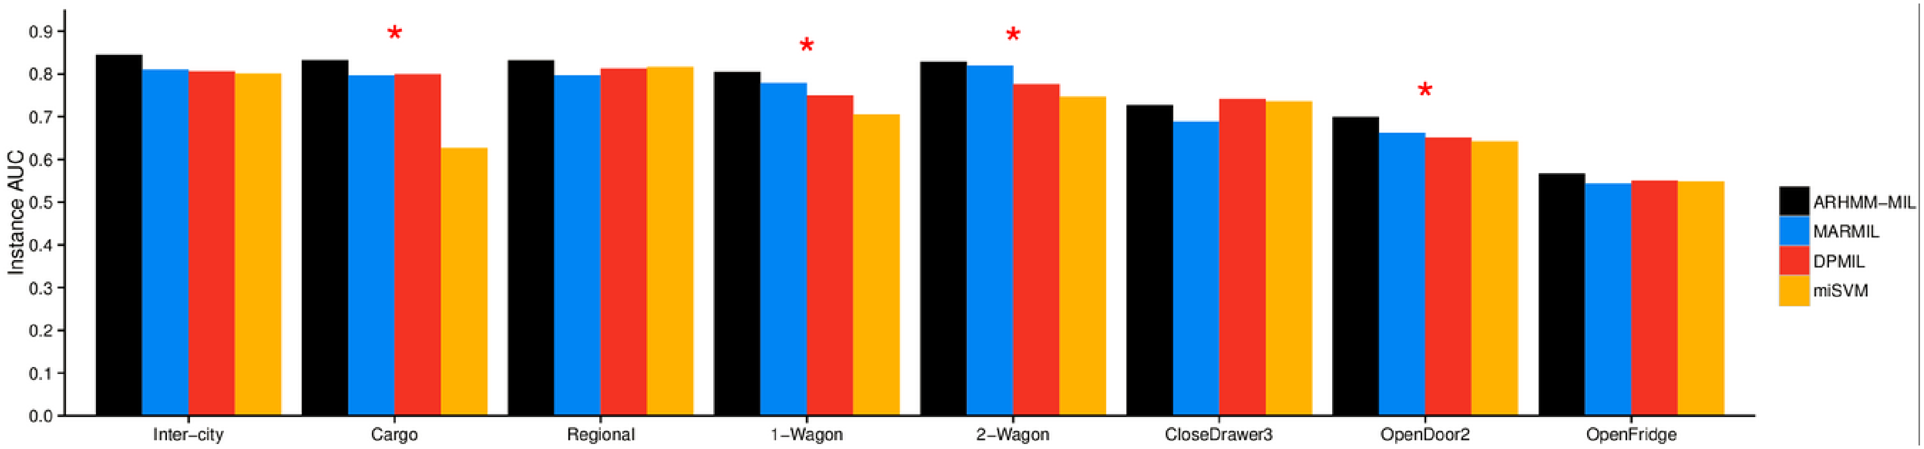

Instance AUC
0.0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
*
*
*
*
Inter–city
Cargo
Regional
1–Wagon
2–Wagon
CloseDrawer3
OpenDoor2
OpenFridge
ARHMM–MIL
MARMIL
DPMIL
miSVM

Conclusion

- ARHMM-MIL models temporal dynamics between instances in a bag
- Generative model that can:
  - Predict bag and instance labels
  - Allow deeper analysis of data by decomposing it into AR processes
  - Allow you to sample data from it

Future Work

Multi-Instance Multi-Label Approach

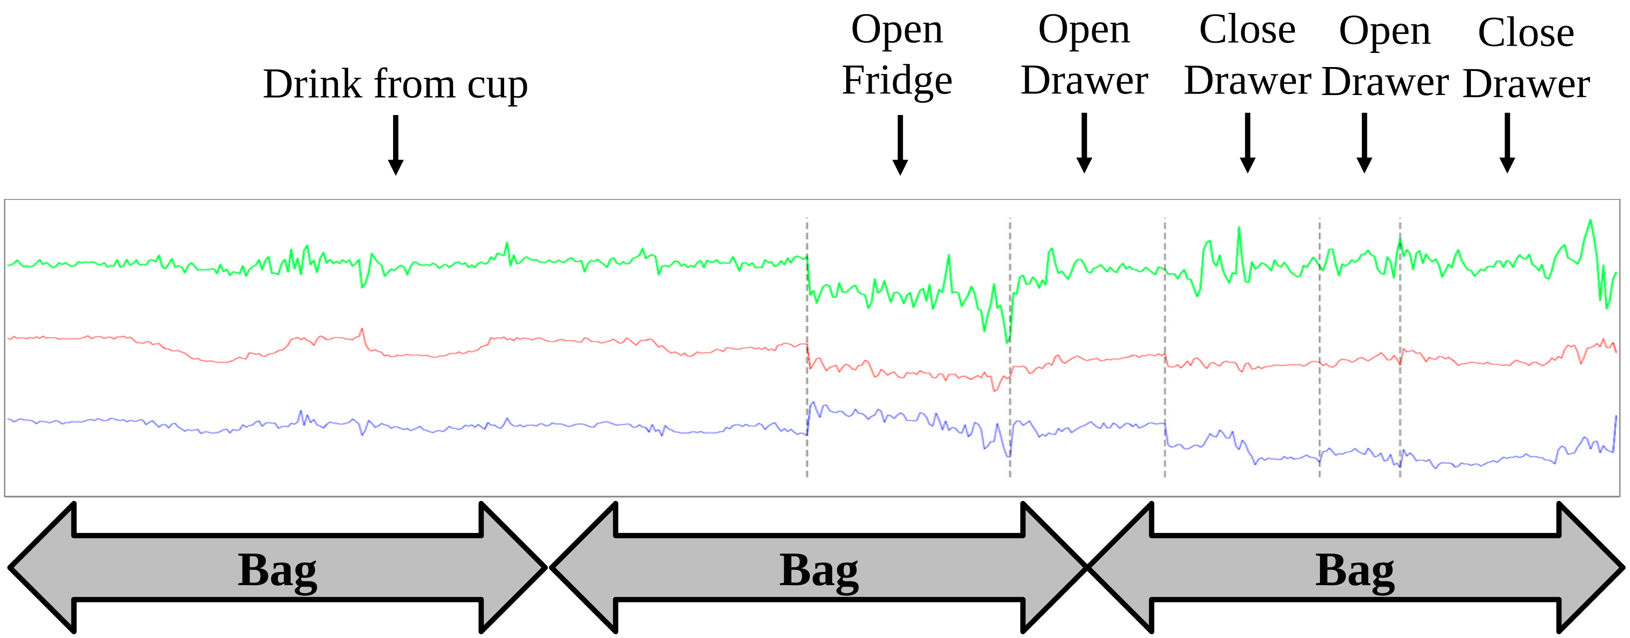

Drink from cup

Drink from cup, Open Fridge, Open Drawer

Open Drawer, Close Drawer

Thank you!

This work is partially supported by NSF grant CCF-1254218

Poster Session: Tues Morning

Questions?