Activity Recognition with a Smart Walker

Friday, October 1st, 2010

CS886: AI for Healthcare Applications

Pascal Poupart

Associate Professor

Cheriton School of Computer Science

University of Waterloo

Smart Walker Project

- Smart walker: rollating walker instrumented with sensors
- Problem: user activity recognition
- Contributions:
  - HMM and CRF modeling
  - Comparative analysis of supervised and unsupervised algorithms
  - Experimentation with regular walker users

Outline

- Motivation
- Experimental data
- Probabilistic models
  - Hidden Markov Model
    - Maximum likelihood
    - Bayesian learning
  - Conditional Random Field
    - Conditional maximum likelihood
- Results
- Conclusion & Future Work

Disability Statistics

- US National Business Service Alliance (July 2006)

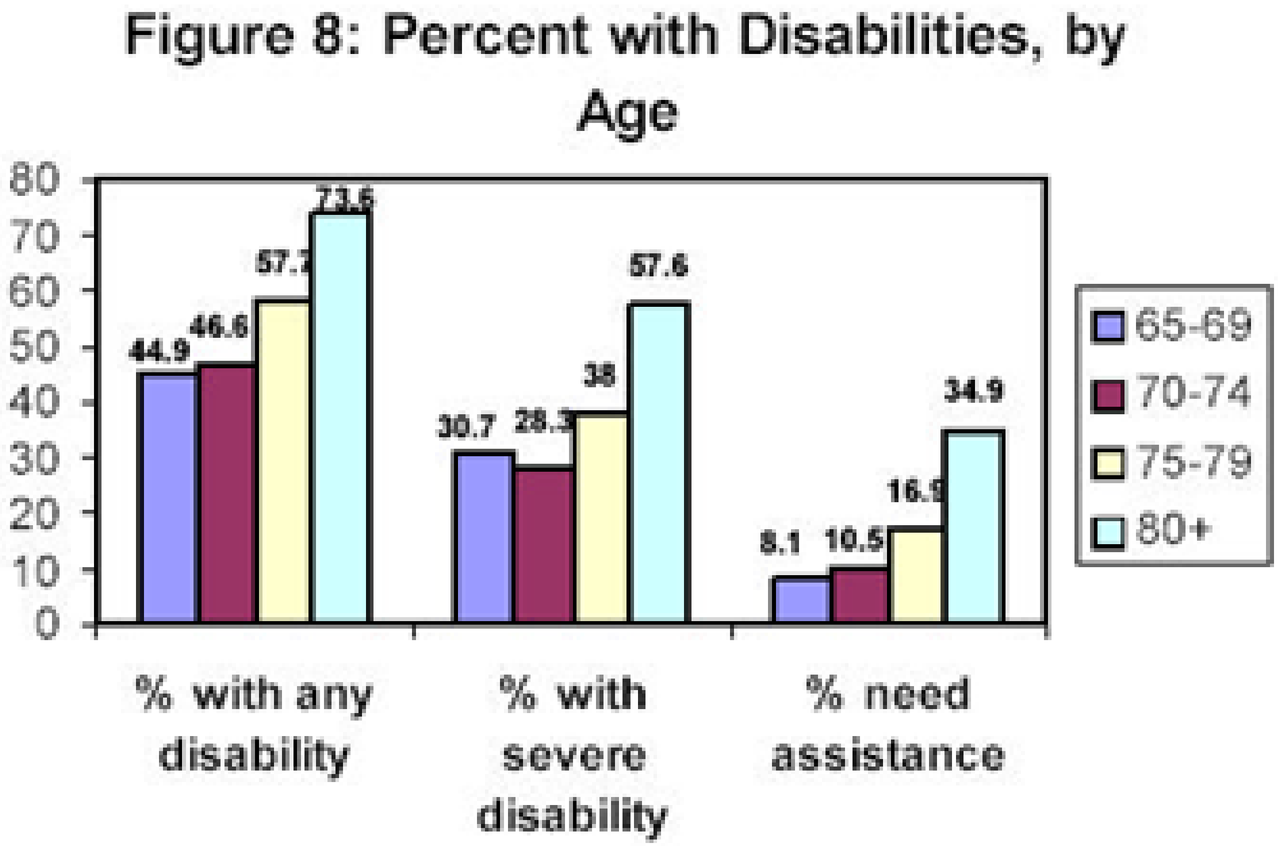

Disability Statistics

- US National Business Service Alliance (July 2006)

Types of disabilities

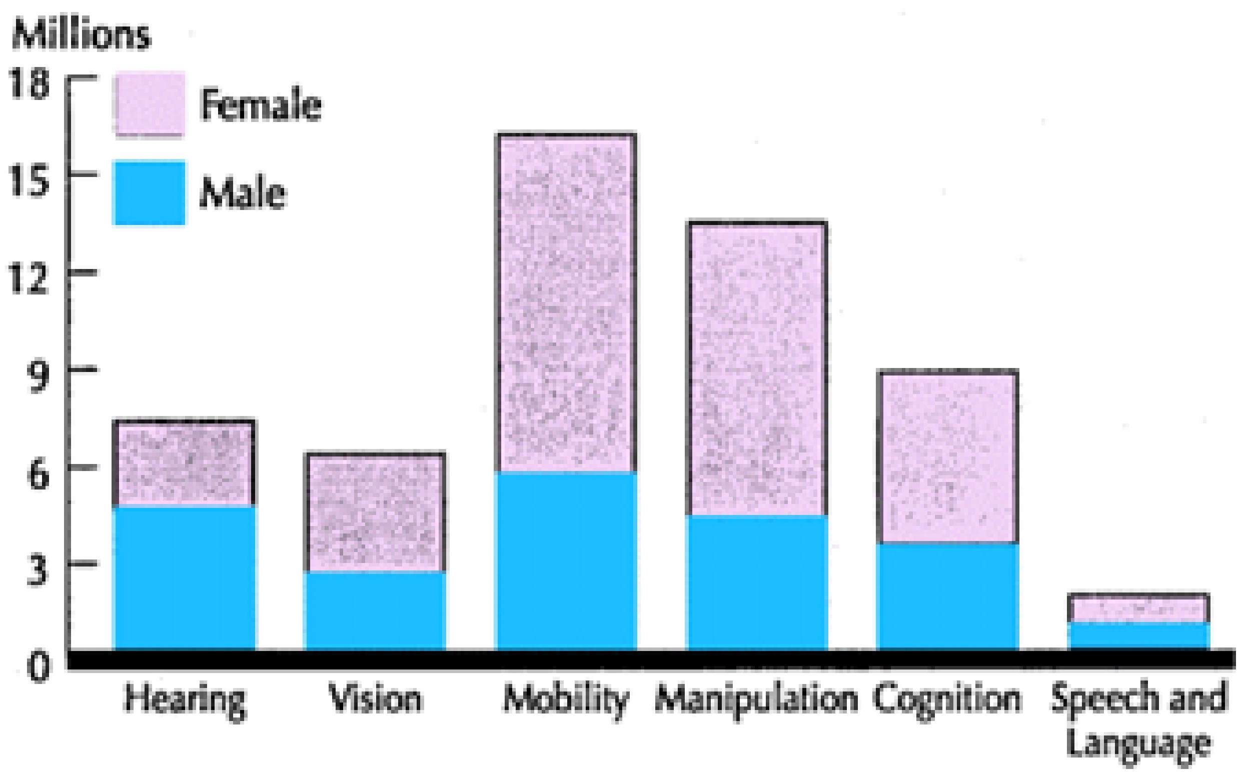

Mobility aids


Encourage walking
Increase safety


Discourage walking

Smart Walker

+

Force sensors
Load sensors
Video cameras
Microphone
Speech synthesizer
Servo-brakes
etc.

walker

devices

Smart walker

assist

caregivers

users

Current Prototype

- Designed by James Tung, Bill McIlroy and Kamran Habib (U of Waterloo and Toronto Rehab Institute)
- Overall objective: enhance mobility of elderlies
  - Monitoring: measure stability, detect falls
  - Assistance: braking, steering
  - Clinical research: collect usage data in non-laboratory setting
- New paradigm: ubiquitous computing in healthcare

Clinical Assessment

- Kinesiologists at Toronto Rehab and the Village of Winston Park assess walking abilities of patients/residents on a regular basis
  - Traditionally: timed-up-and-go test and balance control assessment
  - Recently: walking course with instrumented walker
- Quantitative data used to
  - Prescribe mobility aid
  - Recommend exercise schedule
  - Monitor improvement/deterioration of walking abilities

Activity Recognition

- State of the art: kinesiologists hand label sensor data by looking at video feeds
  - Time consuming and error prone!

Backward view

Forward view

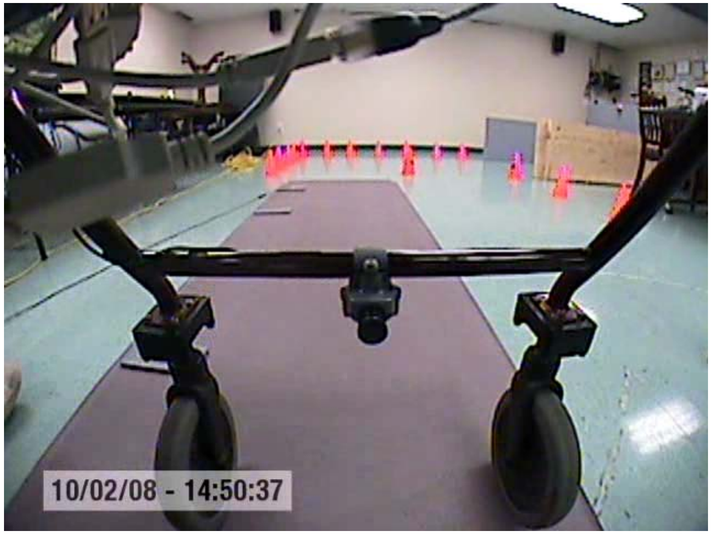

Activity Recognition

- Our contribution: automated techniques for activity recognition based on the non-video sensors

- Related work:
  - Alwan et al. (2004): infer gait characteristics the load applied to the walker handles
  - Hirata et al. (2006): infer whether the user is walking, sitting or in situation of emergency by measuring distance between walker and user

Outline

- Motivation
- Experimental data
- Probabilistic models
  - Hidden Markov Model
    - Maximum likelihood
    - Bayesian learning
  - Conditional Random Field
    - Conditional maximum likelihood
- Results
- Conclusion & Future Work

Raw Sensor Data

- 8 channels:
  - Forward acceleration
  - Lateral acceleration
  - Vertical acceleration
  - Load on left rear wheel
  - Load on right rear wheel
  - Load on left front wheel
  - Load on right front wheel
  - Wheel rotation counts (speed)

- Data recorded at 50 Hz and digitized (16 bits)

Experiment 1

- 17 healthy young subjects (19-53 years old)
- Each person performed the course twice

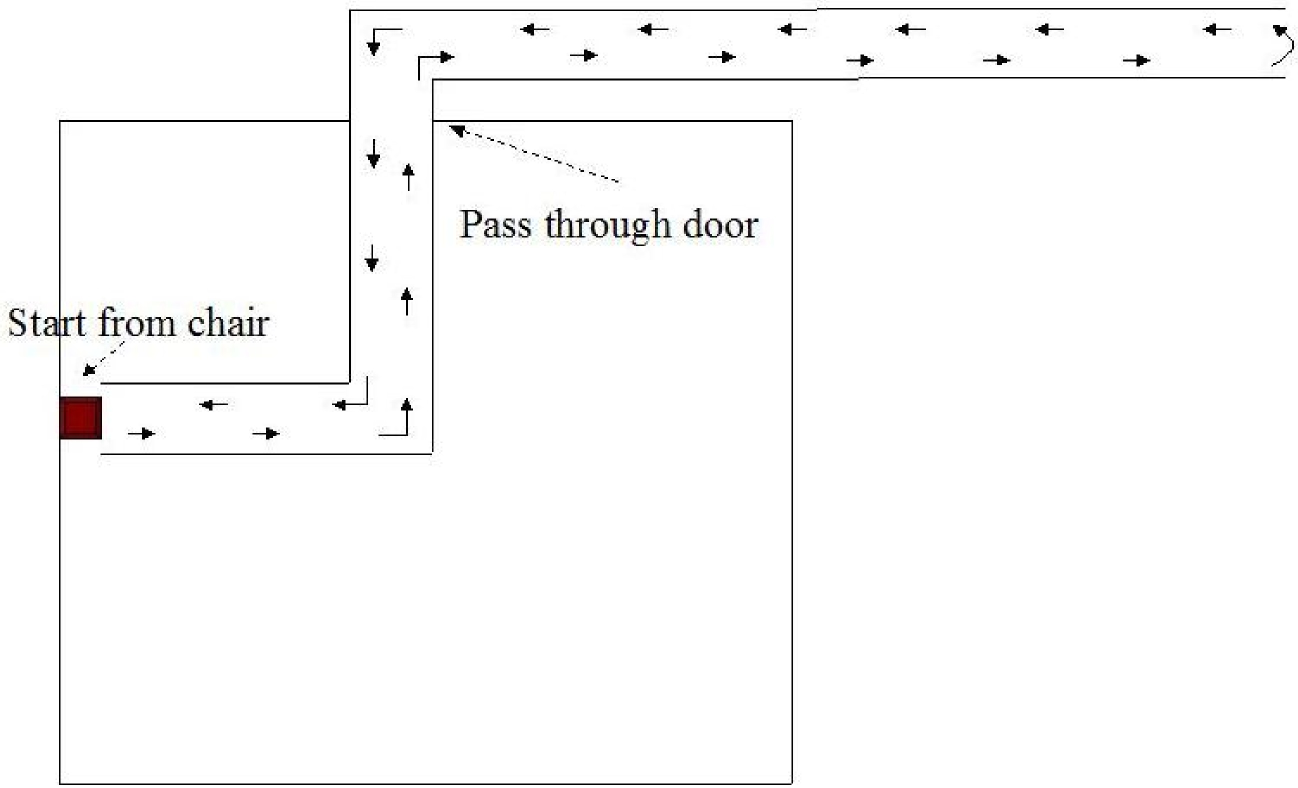

Activities in Experiment 1

- Activities
  - Not touching the walker (NTW)
  - Transfers (sit $\leftrightarrow$ stand) (TR)
  - Standing (ST)
  - Walking Forward (WF)
  - Turning Left (TL)
  - Turning Right (TR)
  - Walking Backwards (WB)

Experiment 2

- 8 walker users at Winston Park (84-97 years old)
- 12 older adults (80-89 years old) in the Kitchener-Waterloo area who do not use walkers

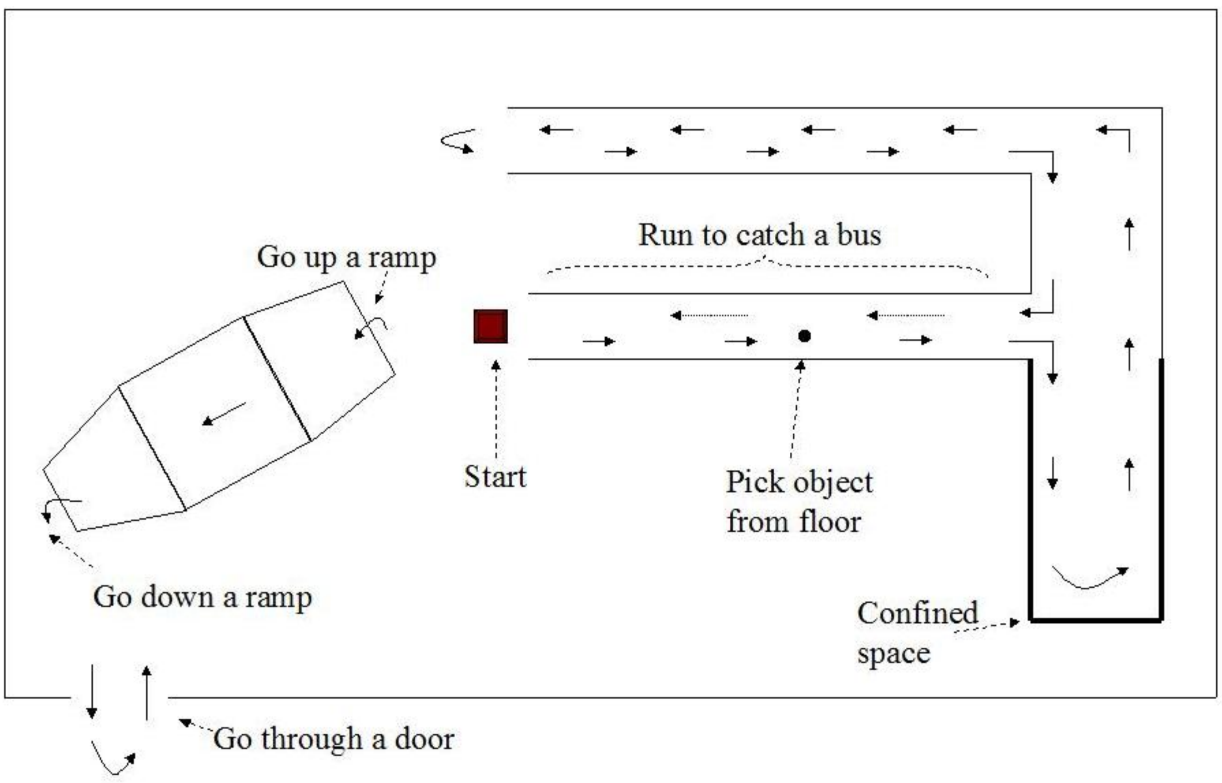

Activities in Experiment 2

- Activities
  - Not Touching the Walker (NTW)
  - Standing (ST)
  - Walking Forward (WF)
  - Turning Left (TL)
  - Turning Right (TR)
  - Walking Backwards (WB)
  - Sitting on the Walker (SW)
  - Reaching Tasks (RT)
  - Going Up Ramp/Curb (GUR/GUC)
  - Going Down Ramp/Curb (GDR/GDC)

Hypotheses I

- Activities expected to be easily distinguishable

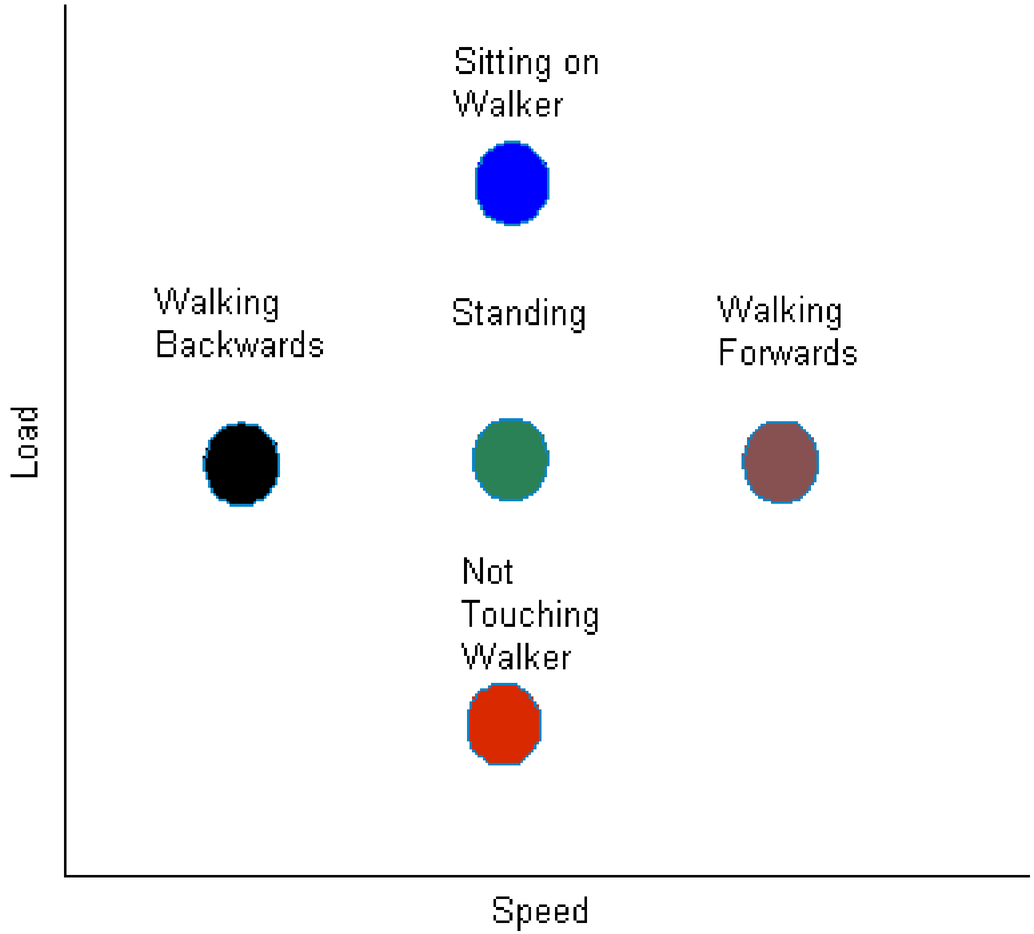

Hypotheses II

- Less obvious distinctions due to a variety of behaviours
  - Turns:
    - load fluctuations on side of turn
    - reduced speed compared to walking
    - Mild lateral force
  - Vertical transitions:
    - Fluctuations in vertical acceleration
    - Load fluctuations
  - Transfers:
    - Fluctuations in rear load
    - Mild movements
  - Reaching tasks
    - Unclear effect on sensors

Outline

- Motivation
- Experimental data
- Probabilistic models
  - Hidden Markov Model
    - Maximum likelihood
    - Bayesian learning
  - Conditional Random Field
    - Conditional maximum likelihood
- Results
- Conclusion & Future Work

Probabilistic Models

- Hidden Markov Model (HMM)
  - Supervised
    - Maximum likelihood (ML)
  - Unsupervised
    - Expectation maximization (EM)
    - Bayesian Learning
- Conditional Random Field (CRF)
  - Supervised
    - Maximum conditional likelihood

Hidden Markov Model (HMM)

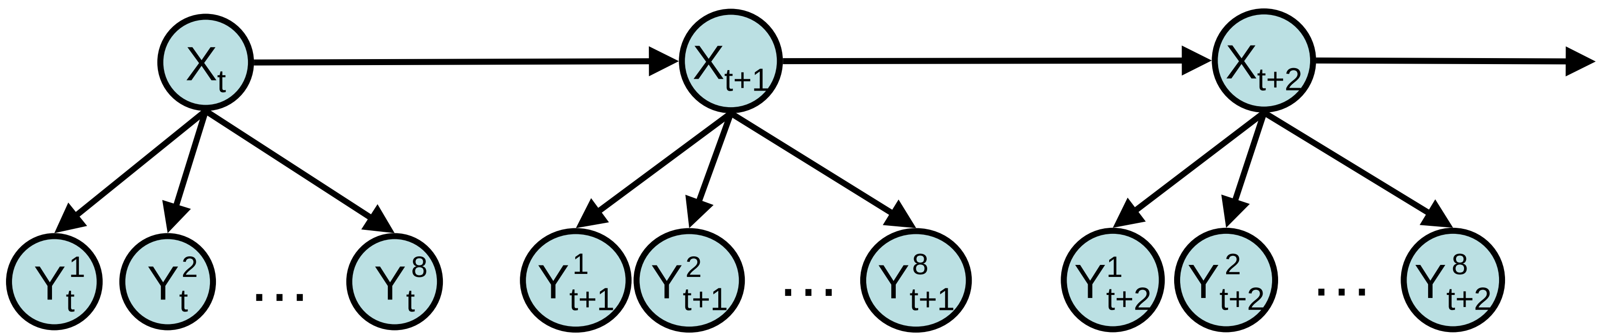

- Parameters
  - Initial state distribution:
  - Transition probabilities:
  - Emission probabilities:

Supervised Maximum Likelihood

- Supervised learning:
  - Relative frequency counts
  - Alternatively, to avoid bias due to predefined course

Unsupervised Maximum Likelihood

- Expectation maximization (EM)
  - Expectations
  - Improvement
- May overfit or get stuck in local optima!

Bayesian Learning

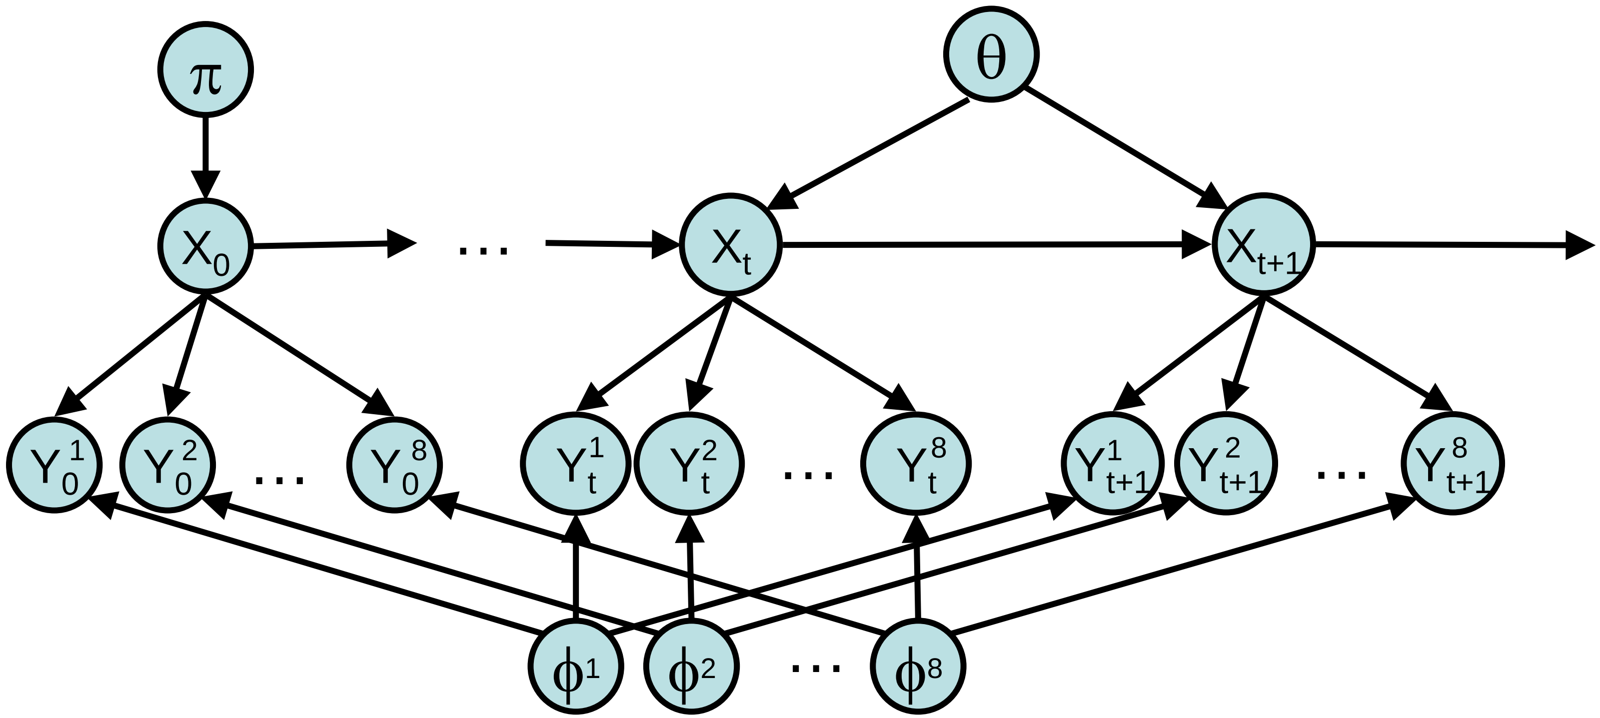

- Parameters treated as random variables
- More robust to overfitting and local optima

Bayesian Learning

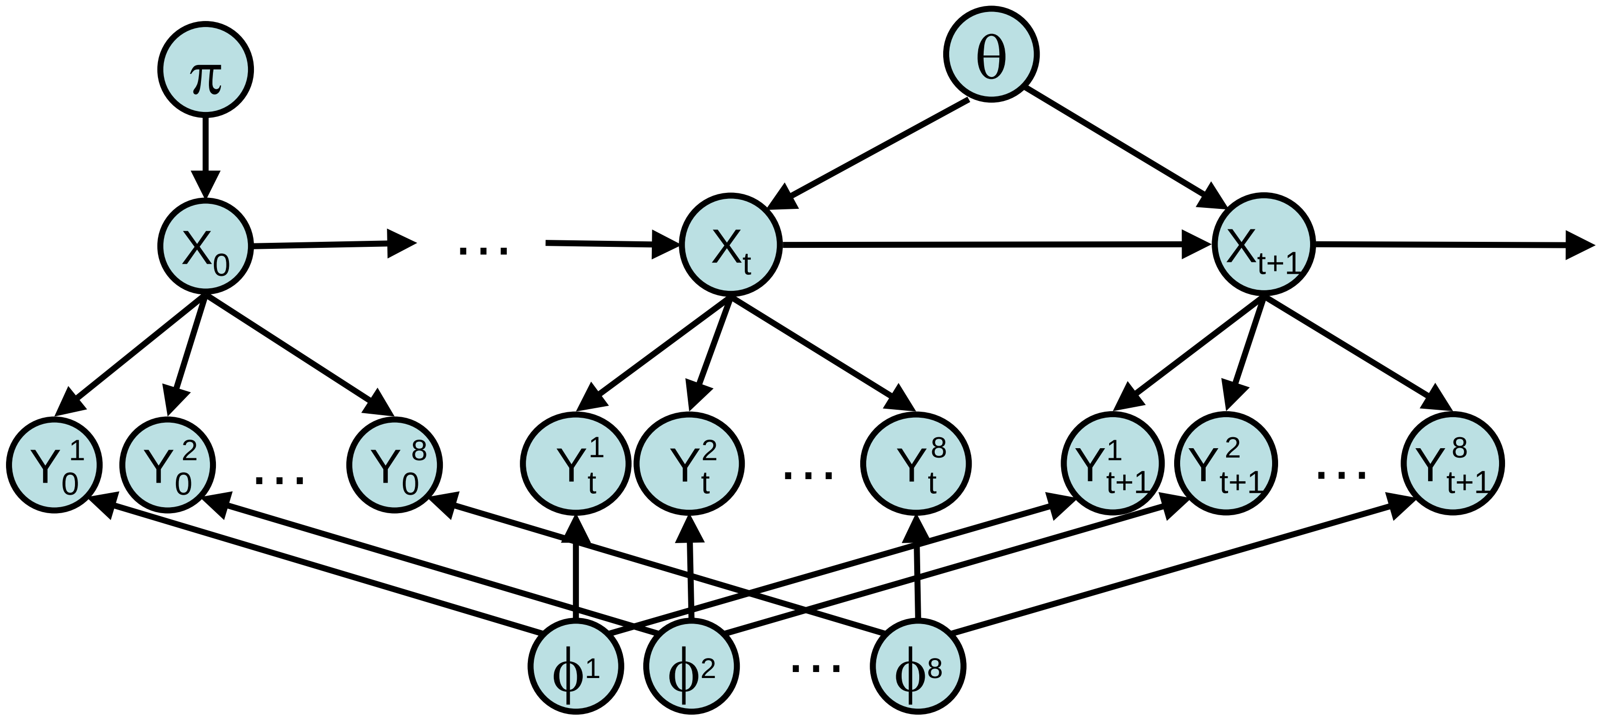

Conditional distributions

Priors

Dirichlet Distributions

- Dirichlets are monomials over discrete random variables

- Properties of Dirichlets
  - Easy to integrate
  - Conjugate prior for discrete likelihood distributions

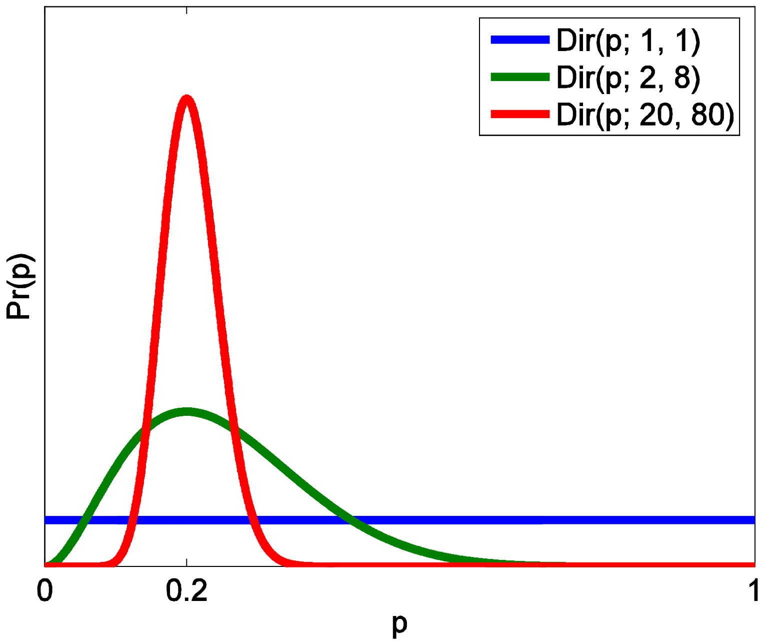

Bayesian Learning

- Parameter learning: compute posterior

- Intractable: exponential mixture of Dirichlets

- Approximate solution: Gibbs sampling

Gibbs Sampling

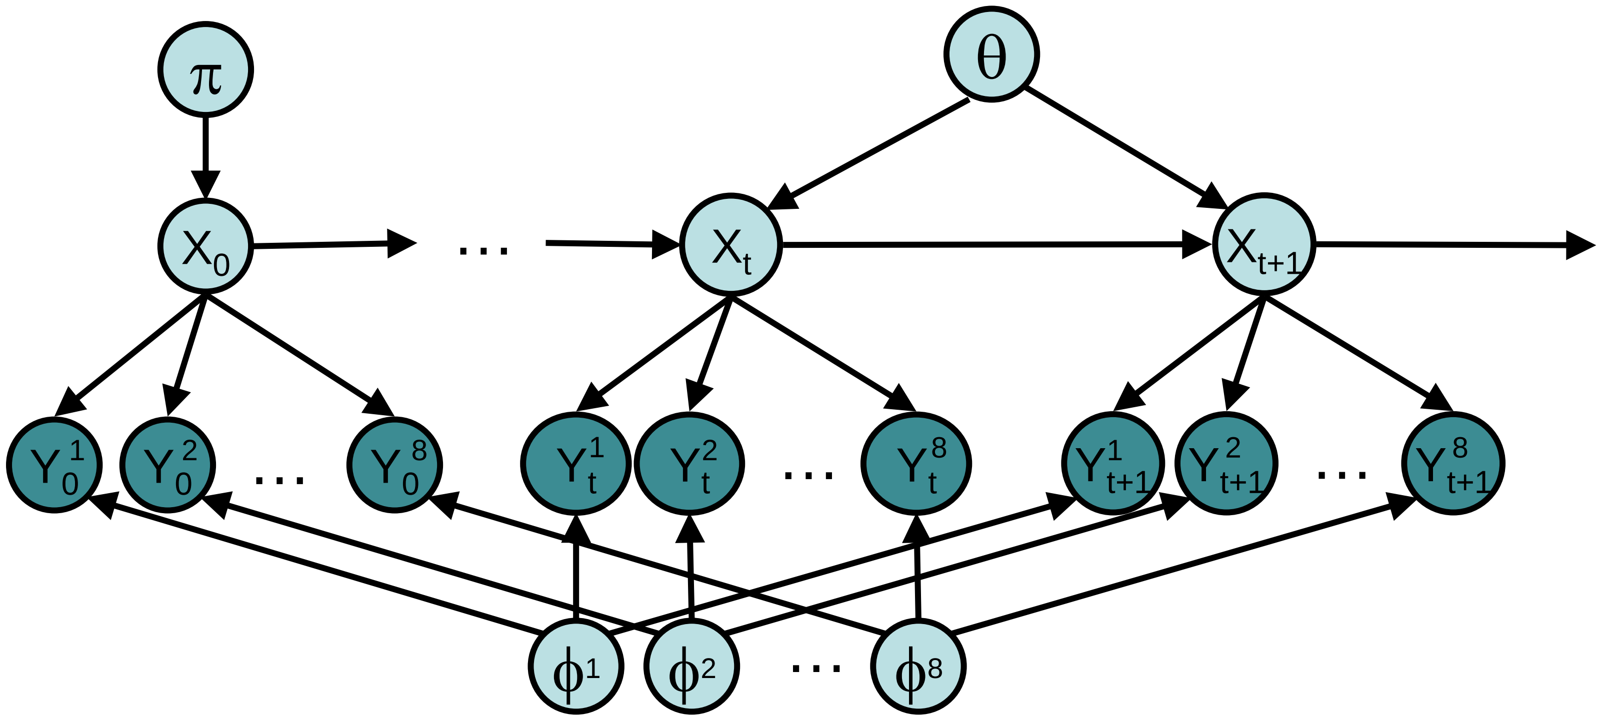

- Markov chain that converges to the posterior
- Repeatedly sample

Collapsed Gibbs Sampling

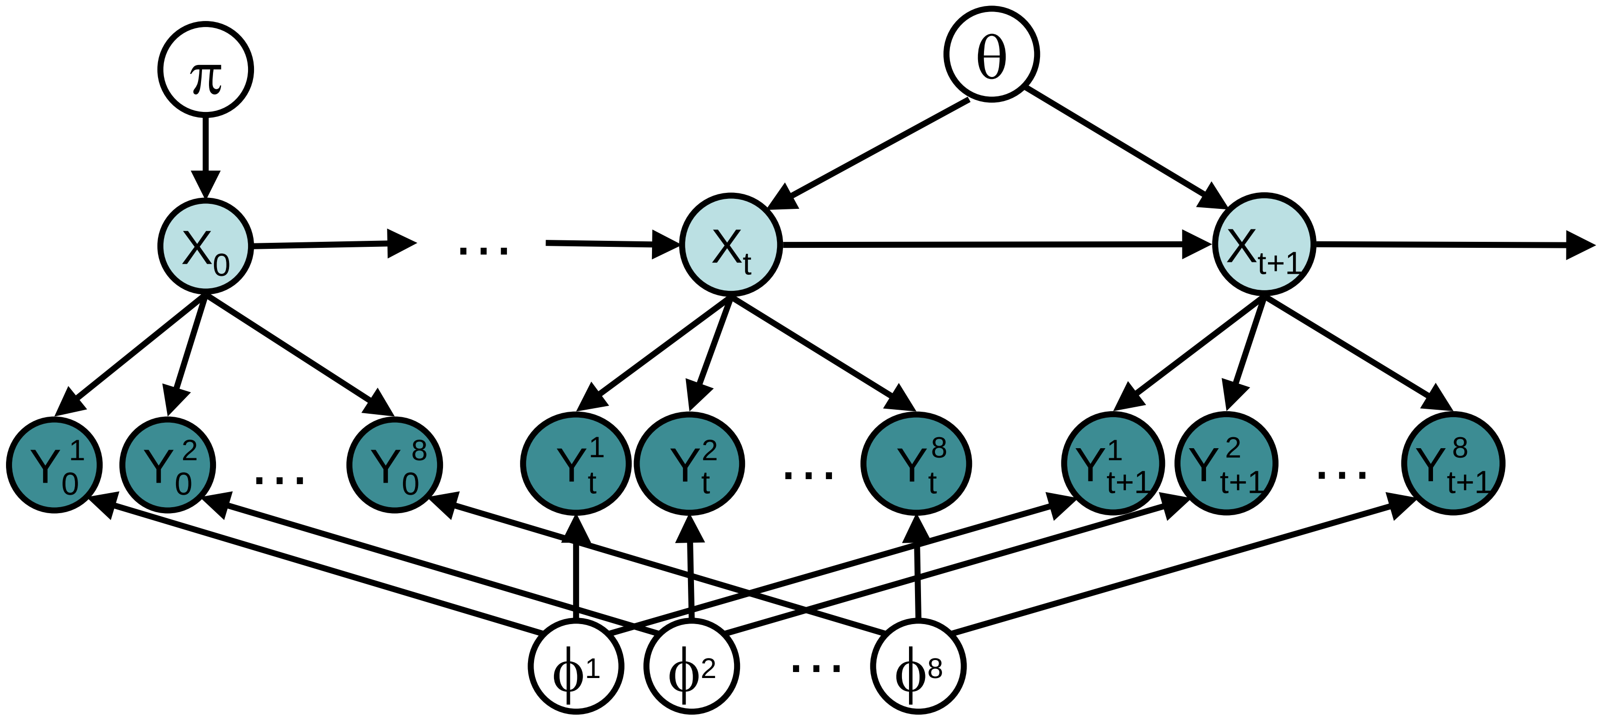

- Faster convergence
- Repeatedly sample while integrating out

Prediction

- Maximum aposteriori filtering

- Recursive belief monitoring

- Online prediction in real-time

Conditional Random Field (CRF)

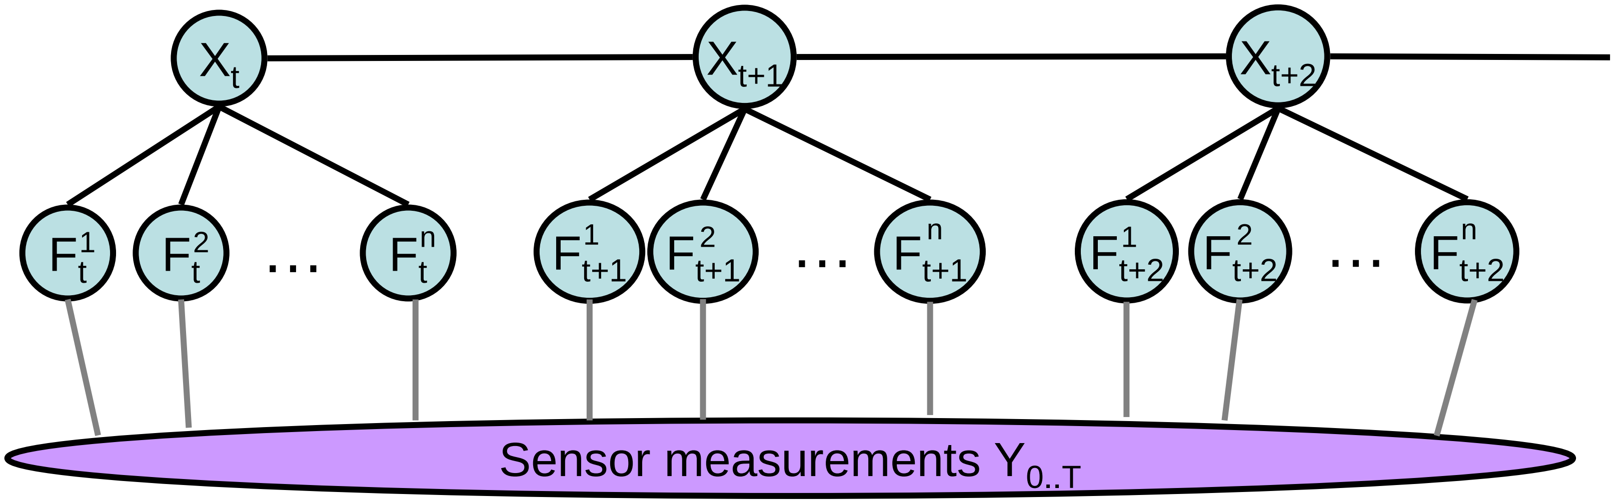

- Discriminative training: maximum conditional likelihood
- Does not assume conditional independence of sensor measurements
- Use features instead of raw measurements

Conditional Random Field (CRF)

- Probabilistic model

- Transition features: capture persistence

Conditional Random Field (CRF)

- Sensor features: thresholded average and std
  - Let $g(y_{t\text{-}w..t})$ be the average or std of
    - speed, total load, center of pressure (media-lateral, anterior-posterior), acceleration (forward, lateral, vertical)
    - For windows of 1, 5, 25 and 50 measurements
  - Features take the form
  - Thresholds
    - Initially: handcrafted based on data inspection
    - Later: learned by linear regression and logistic regression

Conditional Random Field (CRF)

- Training: maximize conditional log-likelihood
  - Concave optimization
  - Conjugate gradient
- Prediction:
  - Online and real-time

Outline

- Motivation
- Experimental data
- Probabilistic models
  - Hidden Markov Model
    - Maximum likelihood
    - Bayesian learning
  - Conditional Random Field
    - Conditional maximum likelihood
- Results
- Conclusion & Future Work

Demo

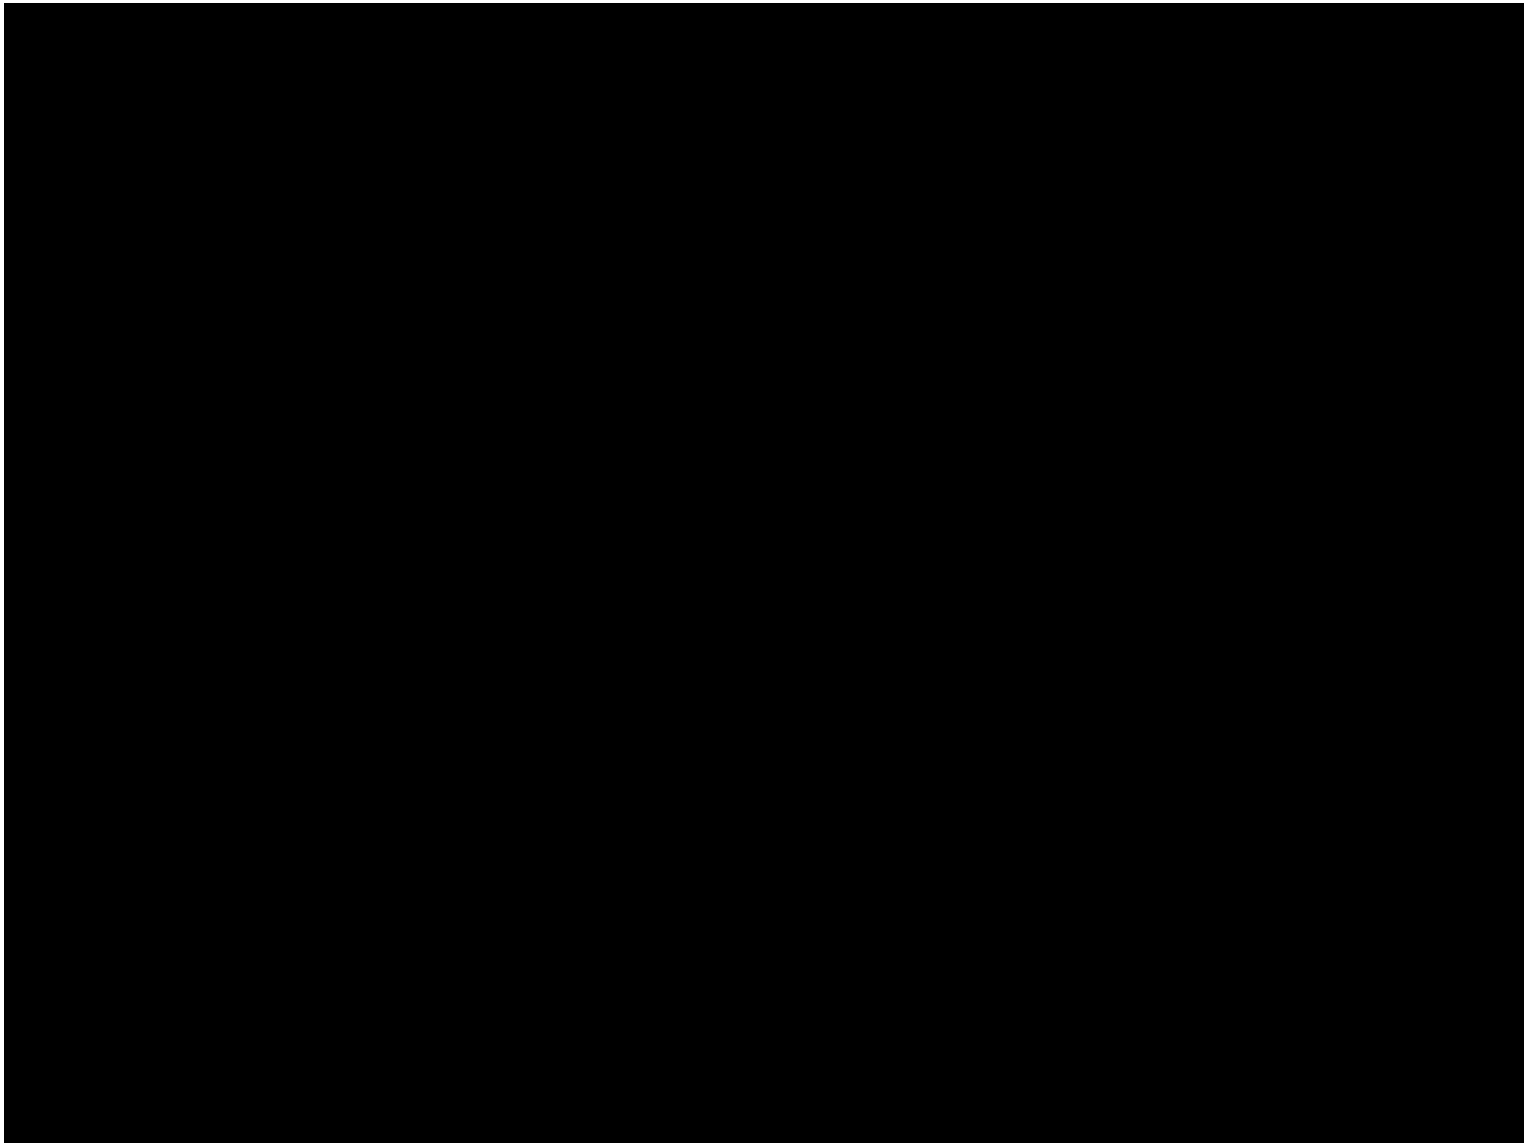

Evaluation Methodology

- Leave-one-out cross validation
- Performance measures:
  - Activity prediction
  - Activity changes:
- A prediction is counted as correct if the corresponding activity (change) occurs within a window of $\pm w$
  - $w \in [0, 50]$

Results

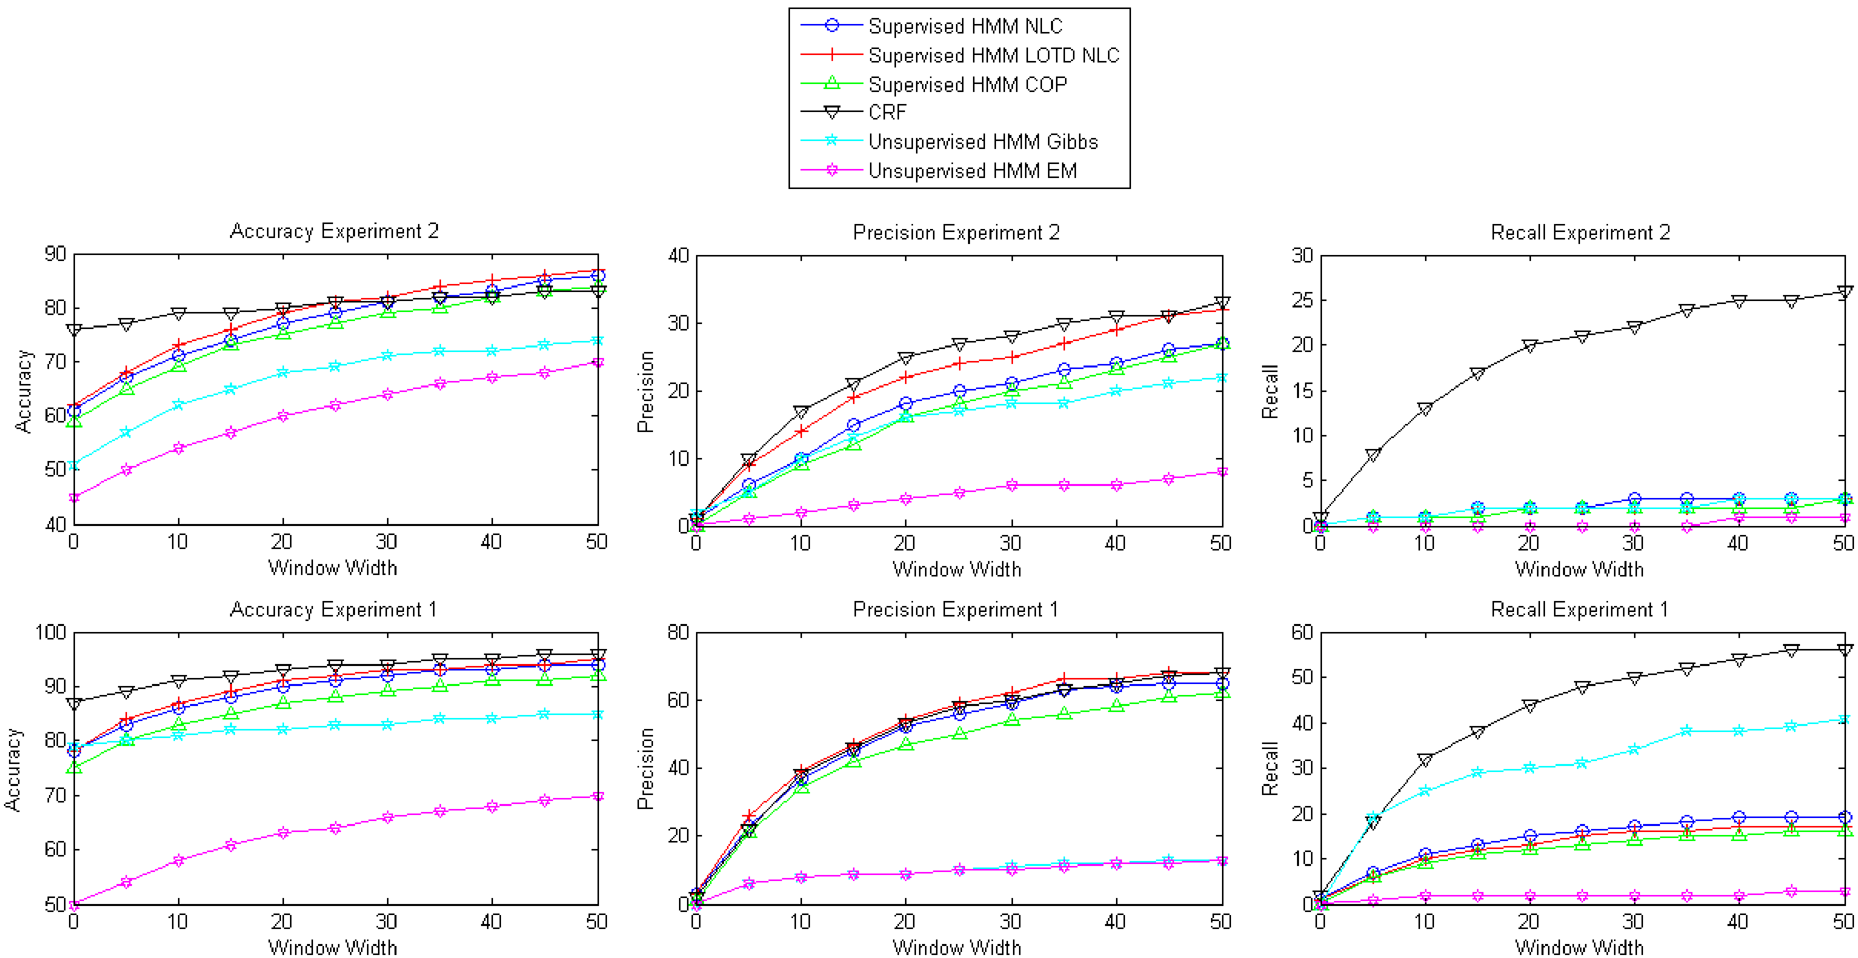

Supervised HMM NLC
Supervised HMM LOTD NLC
Supervised HMM COP
CRF
Unsupervised HMM Gibbs
Unsupervised HMM EM
Accuracy Experiment 2
Precision Experiment 2
Recall Experiment 2
Accuracy Experiment 1
Precision Experiment 1
Recall Experiment 1
Accuracy
Precision
Recall
Window Width

Results

- Accuracy in %, window size: 25

Model	NTW	ST	WF	TL	TR	WB	TRS	overall
HMM (ML)	65	88	95	96	92	95	85	91
CRF	96	82	98	89	80	94	70	93
HMM (EM)	100	14	48	94	90	97	0	61
HMM (Gibbs)	100	72	95	66	72	44	17	83

Model	NTW	ST	WF	TL	TR	WB	RT	SW	GUR	GDR	GUC	GDC	total
HMM (ML)	50	71	73	81	73	21	52	99	86	86	94	85	81
CRF	78	94	95	51	20	0	5	99	31	23	67	13	81
HMM (EM)	73	33	54	63	63	7	30	99	40	23	15	91	62
HMM (Gibbs)	100	42	77	69	81	0	62	91	57	52	46	2	69

Discussion

- Experiment 1: better results than Experiment 2
  - Participants perform the course twice
  - Less complex activities

- CRF: better results than HMM
  - Discriminative approach well suited
  - Sensor measurements violate conditional independence

- Latent states obtained by unsupervised techniques:
  - Manually match each latent state with the activity it is most frequently predicted as
  - Sub-dividing or merging activities may improve the results

Outline

- Motivation
- Experimental data
- Probabilistic models
  - Hidden Markov Model
    - Maximum likelihood
    - Bayesian learning
  - Conditional Random Field
    - Conditional maximum likelihood
- Results
- Conclusion & Future Work

Conclusion

- Contributions
  - Introduced the problem of recognizing the activity of walker users from streams of low-level sensor measurements
  - Presented supervised and unsupervised algorithms based on HMMs and CRFs
  - Showed results with real data from regular walker users and people who do not use walkers
  - Proposed algorithms can be regarded as “baselines” that yield good results that we plan to further improve.

Future Work

- Improve accuracy
  - Extract features from video
  - Develop more sophisticated models (e.g., hierarchical, time)
  - Combine with limb tracking
- Evaluation
  - Long-term recordings
  - Undefined walking course
- Use activity models to
  - Assist braking and navigation safely
  - Evaluate risk of falling
  - Evaluate activity patterns and level of exercise

Vision: Robotic Nurse

- Navigation assistance
  - Obstacle avoidance
  - Route to some destination
- Personal trainer
  - Monitor and coach exercises
- Reminder system
  - Medication, appointments, location
- Alert system
  - Notify caregiver when in need or fallen
- Conversational companion
  - Speech recognizer/generator

Pascal's Health Related Projects

- Behaviour recognition with a smart walker
  - Farheen Omar (PhD), Mathieu Sinn (Postdoc)
- Limb tracking with a smart walker
  - Richard Hu (MMath)
- Automated feature generation in temporal models
  - Adam Hartfiel (PhD), Mathieu Sinn (Postdoc)
- Voice, Location and Activity Monitoring for Alzheimer
  - James Tung (Postdoc)

Thank You

Questions?