Suspicious Behavior Detection

Boštjan Kaluža
Department of Intelligent Systems, Jozef Stefan Institute, Slovenia
Paul Scerri
Robotics Institute, Carnegie Mellon University, USA
Gal Kaminka
MAVERICK Group, Bar Ilan University, Israel
Milind Tambe
Teamcore Research Group, University of Southern California, USA

University of Southern California, September – December 2010

Introduction

- Definition:
  - Anomalous behavior – different from all others
  - Suspicious behavior – specific behavior
- Related work
  - Single camera view
  - Detection of anomalous trajectories
- Our goal
  - Assume all passengers can be suspicious
  - Monitor behavior all the time (e.g., Raytheon vision system)
  - Focus on actions instead on trajectory paths

Outline

- Modeling suspicious traces in simulator
  - Simulator
  - Suspicious behavior
- Detecting suspicious trajectories
  - State presentation
  - HMM problem setting, experiments
  - Future directions
- Detecting suspicious events
  - HMMs, CHMMs
  - LHMMs, utility-based plan recognition, penalty-based accumulation
- Future directions

Modeling Suspicious Traces in Simulator

Multi-agent Simulator (Tsai et al.)

- Different behaviors
  - Shopping
  - Wandering
  - Going to specific location
  - Families and children
- Authority figures
  - Patrolling
  - Evacuating people
- Cognitive mechanisms
  - Shopping over wandering
  - Surviving mechanism
  - Emotional contagion (spreading fear/calmness)
  - Incomplete knowledge
  - Social comparison (mimicking)

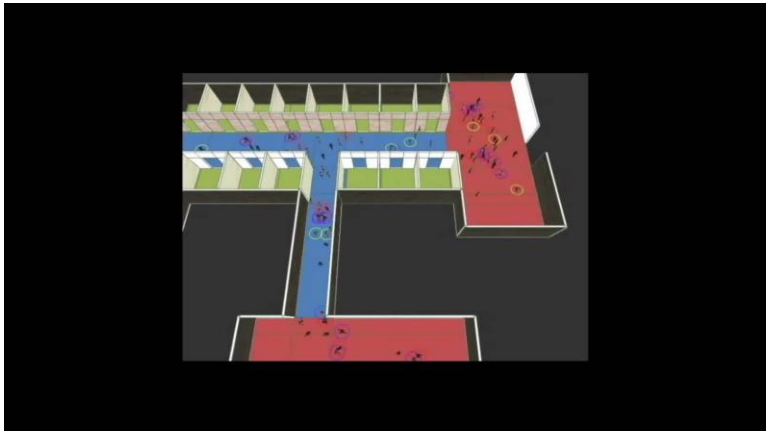

J. Tsai et al. (AAMAS 2011): 3D visualization in Massive.

Modeling Suspicious Traces

- Suspicious person
  - Goes from point A to point B
  - Avoids authorities
  - Models authority’s viewpoint

Modeling Suspicious Traces

- Suspicious person
  - Goes from point A to point B
  - Avoids authorities
  - Models authority’s viewpoint

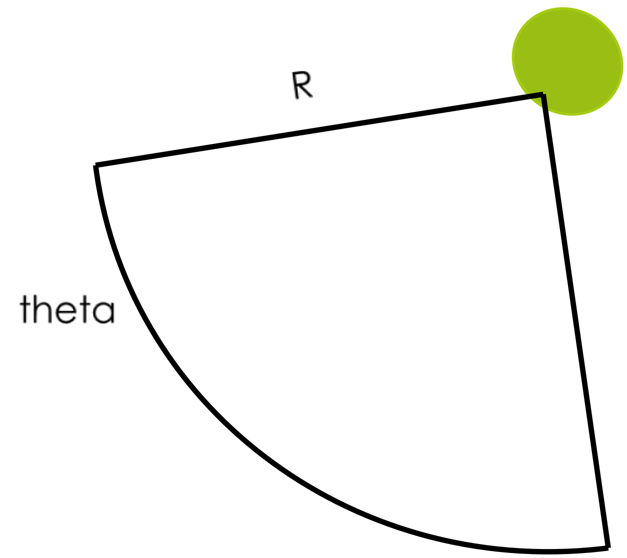

- Modeling authority's viewpoint
  - Authority: position and orientation
  - Viewpoint is bounded by angle *theta* and distance *R*

Modeling Suspicious Traces

- Suspicious person
  - Goes from point A to point B
  - Avoids authorities
  - Models authority's viewpoint

- Modeling authority's viewpoint
  - Authority: position and orientation
  - Viewpoint is bounded by angle *theta* and distance *R*
  - Ability to see is modeled with bivariate Gaussian distribution: *N(theta, R)*

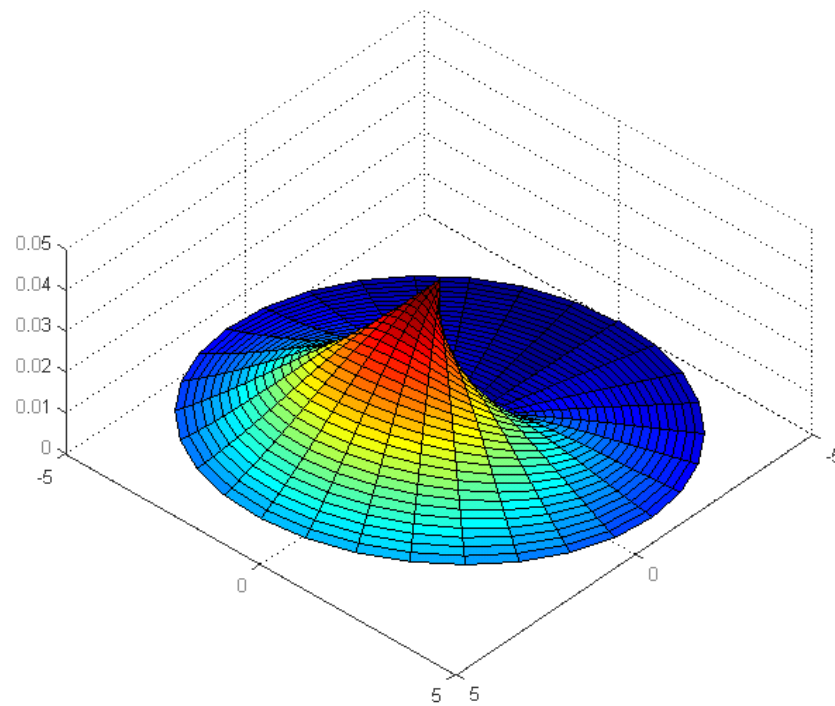

Authority's viewpoint (ability to see person)

Simulator

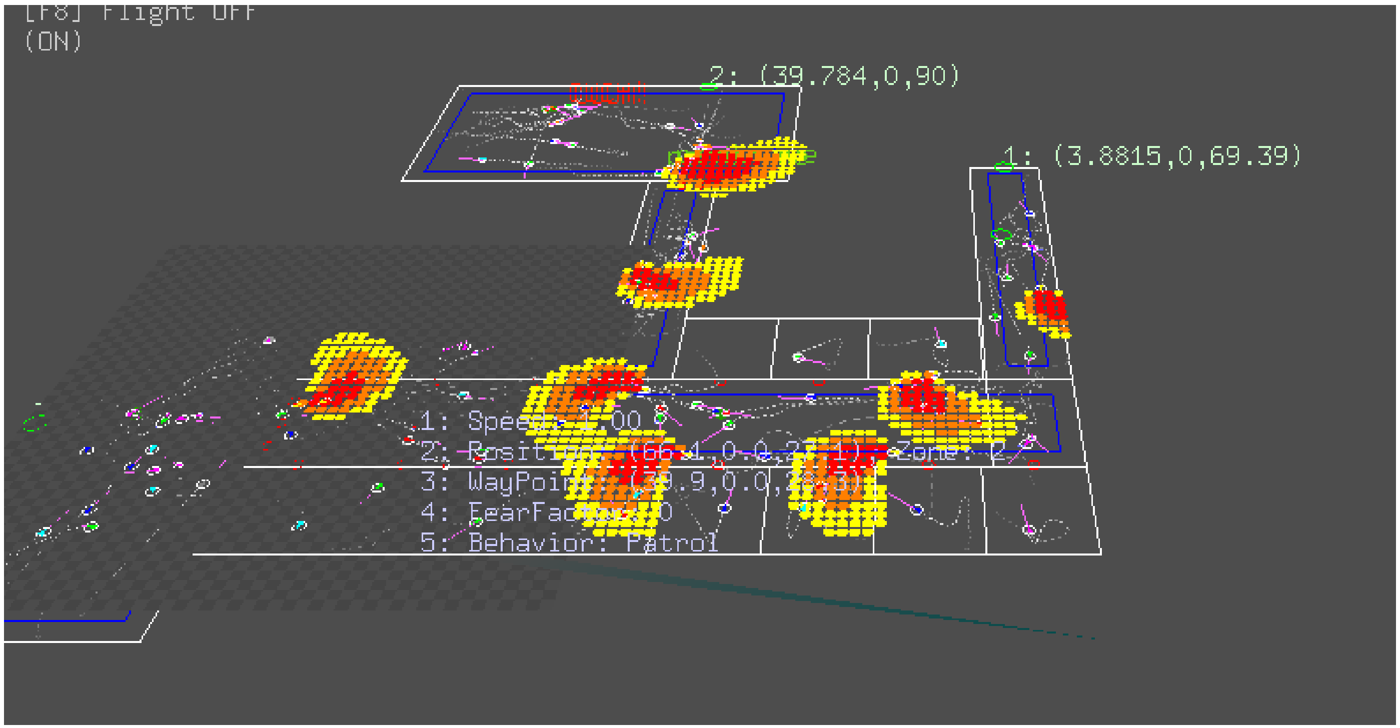

Authorities' viewpoints at the airport terminal

Modeling Suspicious Traces

- Behavior of suspicious agent
  - Suspicious person is at particular position
  - Compute probabilities for being seen by any authority figure in range *2R*

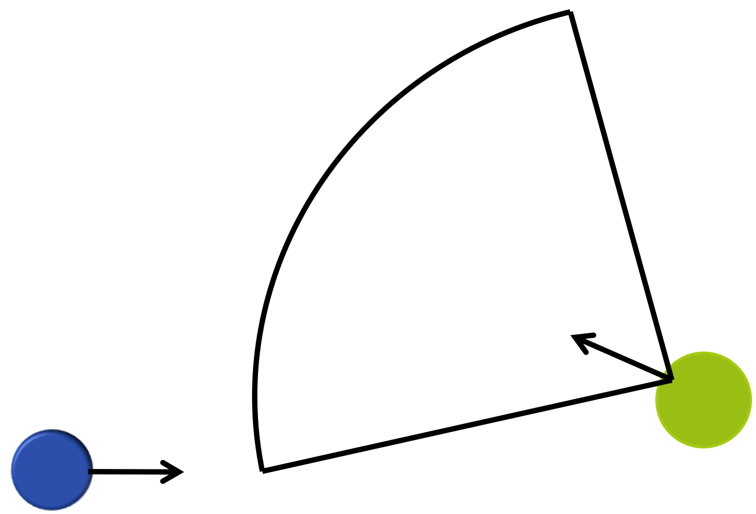

Modeling Suspicious Traces

- Behavior of suspicious agent
  - Suspicious person is at particular position
  - Compute probabilities for being seen by any authority figure in range *2R*
  - If probability exceeds a threshold value
    - Choose eight possible points around current position to avoid
    - Pick a point with the cheapest path from current position to the point (cost is probability for being noticed by authorities)

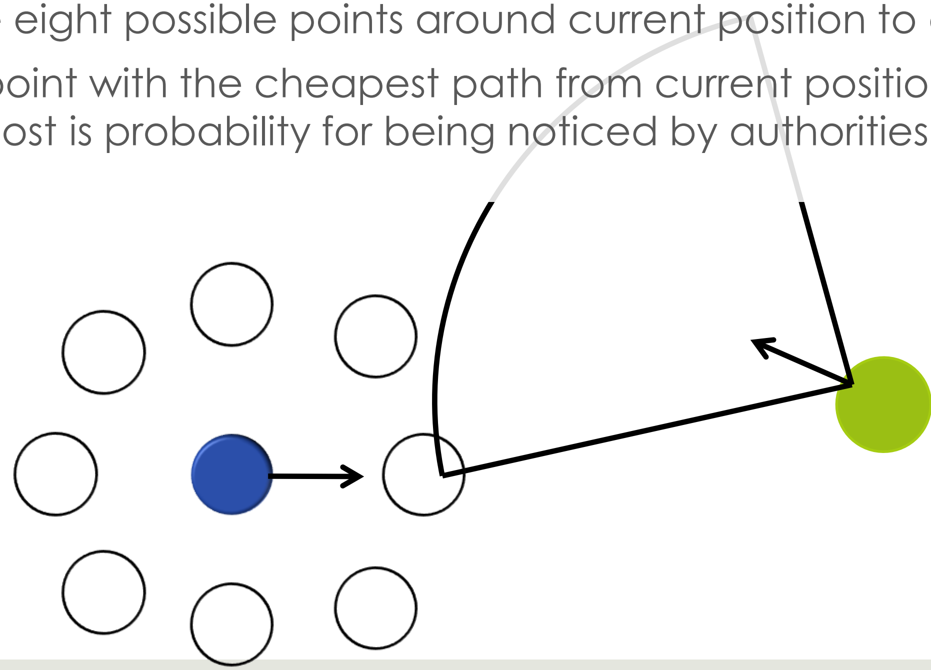

Modeling Suspicious Traces

- Behavior of suspicious agent
  - Suspicious person is at particular position
  - Compute probabilities for being seen by any authority figure in range *2R*
  - If probability exceeds a threshold value
    - Choose eight possible points around current position to avoid
    - Pick a point with the cheapest path from current position to the point (cost is probability for being noticed by authorities)

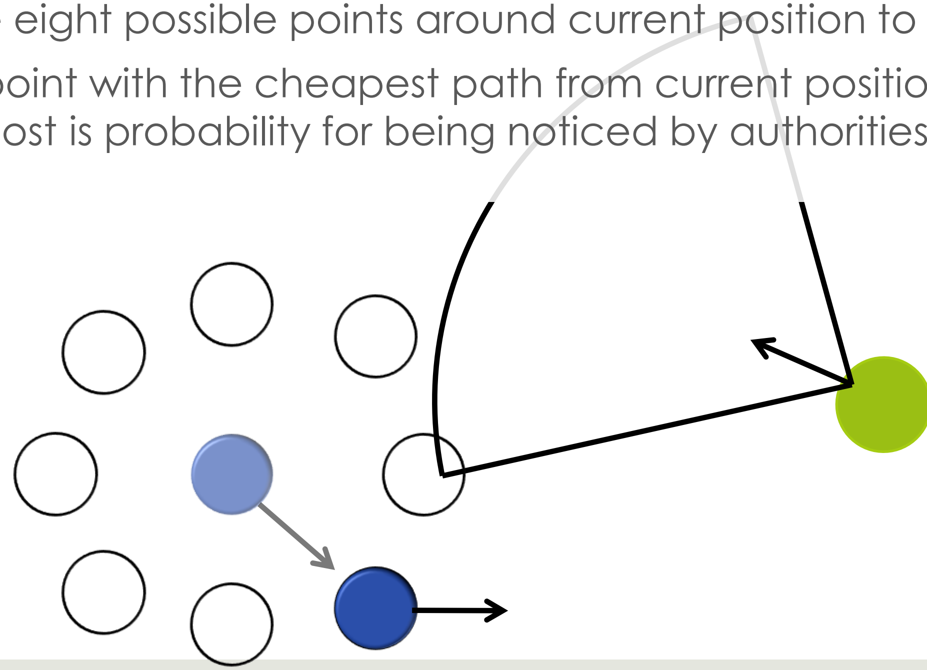

Modeling Suspicious Traces

- Behavior of suspicious agent
  - Suspicious person is at particular position
  - Compute probabilities for being seen by any authority figure in range *2R*
  - If probability exceeds a threshold value
    - Choose eight possible points around current position to avoid
    - Pick a point with the cheapest path from current position to the point (cost is probability for being noticed by authorities)

- Suspicious behavior in practice
  - Avoids in a half-circle
  - Makes u-turns
  - Hides in a nearby store checking the doors
  - Goes around the corner for a while

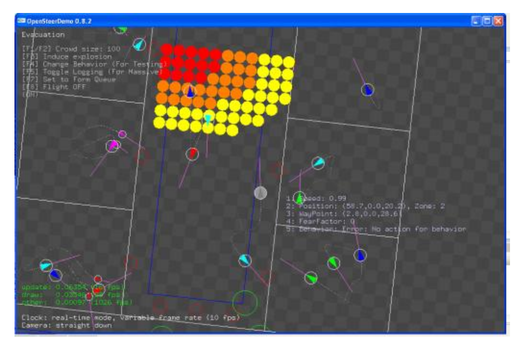

Modeling Suspicious Traces

- Simulator input
  - Airport map
  - Authority agents
  - Suspicious agent
  - Agents with normal behavior
    (Shopping, wandering, families)
- Result: 2D traces of all agents

Traces at the Airport

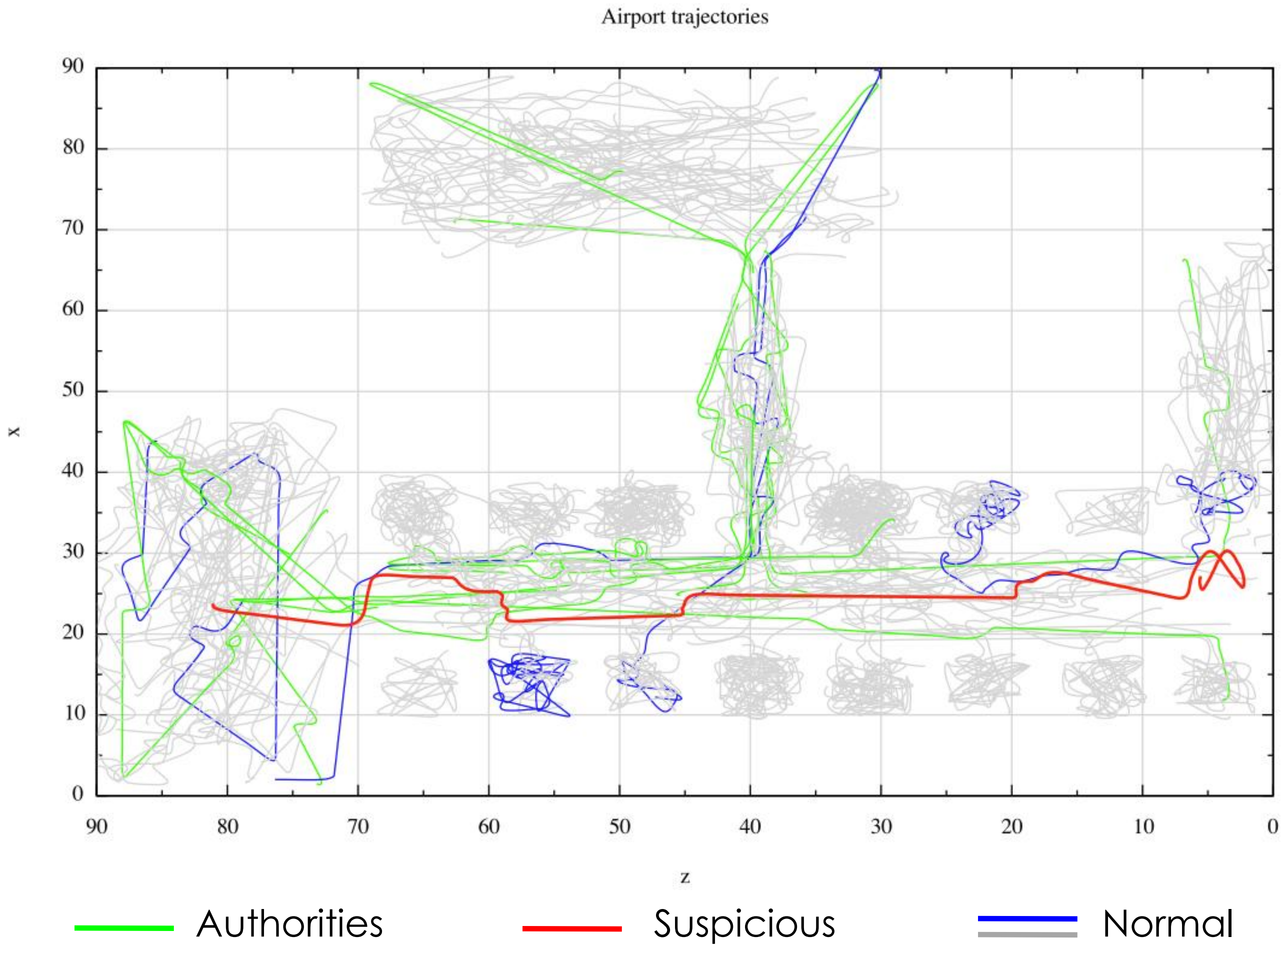

Detecting Suspicious Trajectories

State Presentation

- Trace presented as sequence
- Fixed presentation
  - Coordinates (x, y) on the mesh
  - Numbered fields
- Relative presentation
  - Actions – move North/West/East/South
  - Actions – move Forward/Back/Left/Right

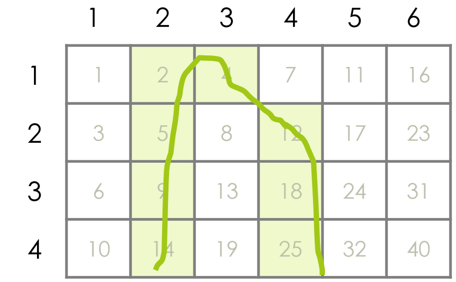

- Example

Coordinates	(4, 2)	(3, 2)	(2, 2)	(1, 2)	(1, 3)	(2, 4)	(3, 4)	(4, 4)
Fixed	14	9	5	2	4	12	18	25
North/West/East/South	N	N	N	E	SE	S	S	S
Forward/Back/Left/Right	F	F	F	L	LF	LF	F	F

Detecting Suspicious Traces

One layer, complete trajectories:

- Hidden Markov Models (Rabiner, 1989)
  - Single subject
  - Model sequence of actions
  - Fixed, relative state presentation

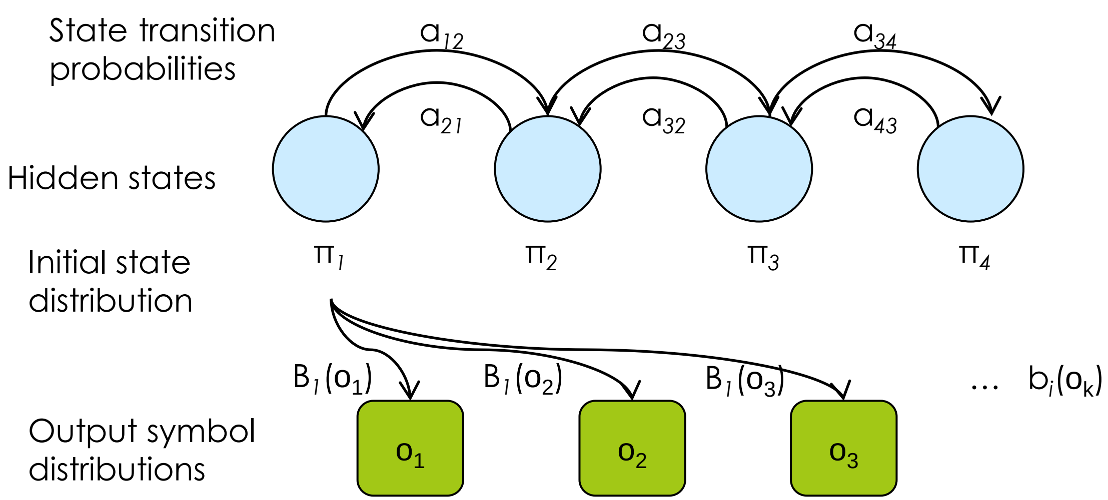

Detecting Suspicious Traces

One layer, complete trajectories:

- Hidden Markov Models (Rabiner, 1989)
  - Single subject
  - Model sequence of actions
  - Fixed, relative state presentation
- HMM as Dynamic Bayesian network (DBN)

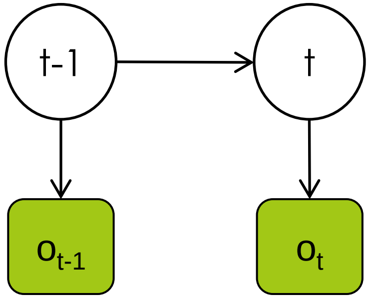

HMM Setting

- One HMM for
  - Normal behavior
  - Suspicious behavior
- Classification of sequence *s*
  - argmax(P(HMM$_i$(s))
- Algorithms
  - Training: Baum-Welch
  - Sequence probability: Forward-backward
  - Toolbox: BNT (Matlab, Kevin P. Murphy)

HMM Setting

- Number of observations
  - 9 for relative state presentations
  - 100 for fixed presentation
- Number of hidden states
  - One state for each observation
- Sequence length
  - Varied
- Initial parameters
  - Random
  - Possible improvement: clustering

Experiments

- Multi-agent simulator
  - 100 normal people
  - 10 authorities
  - 1 suspicious person
- Dataset
  - ~1000 traces
  - 99.1% normal
  - 0.9% suspicious
- Validation
  - 10-fold-cross

Results

- Detecting suspicious traces

	HMM Absolute position	HMM Relative position	HMM Relative position and orientation
Recall	62.63	66.23	40.86
Precision	7.40	10.42	11.54
F-measure	13.24	18.01	18.00

Detecting Suspicious Events

Detecting suspicious events

- Goal: detect and accumulate suspicious events
- Extract all pairs:
  - One person is authority
  - Distance < D_MIN
- Two datasets
  - Authority : Normal person
  - Authority : Suspicious person

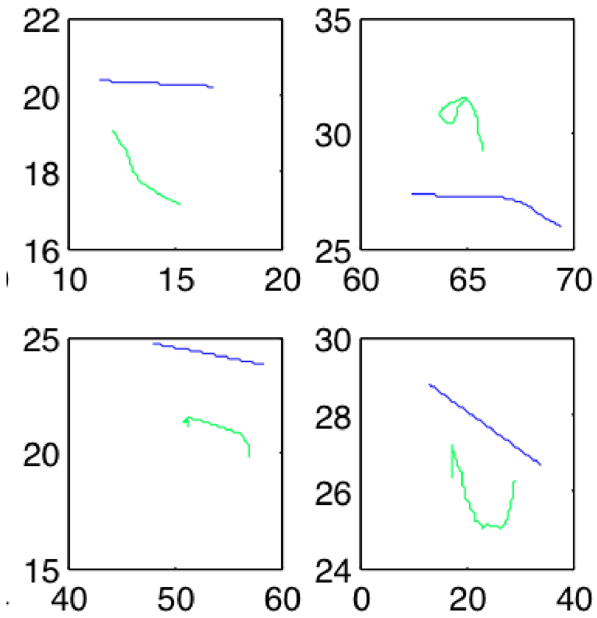

Detecting suspicious events

Two layers:

- Detection of suspicious events
  - Detect u-turns, avoidance, hiding
  - Coupled HMMs (others in progress)

- Accumulation of suspicious events
  - Layered HMMs (Oliver et al., 2004)
  - Utility-based plan recognition (Kaminka et al., 2007)
  - Penalty-based accumulation (Kaluza, Tambe)

Layer 1: Detection of Suspicious Events

- HMMs (Rabiner, 1989)
  - Single subject
  - Model sequence of actions

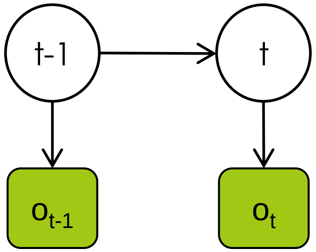

Layer 1: Detection of Suspicious Events

- HMMs (Rabiner, 1989)
  - Single subject
  - Model sequence of actions

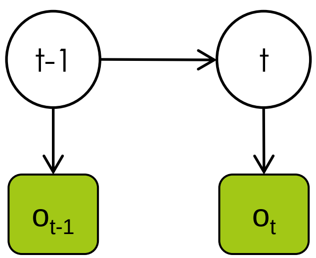

- Coupled HMMs (Brant et al., 1996)
  - Two or more subjects
  - Model sequence of actions and their interactions

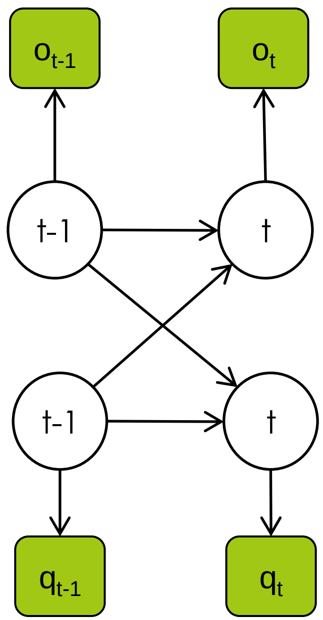

Layer 1: CHMM Input

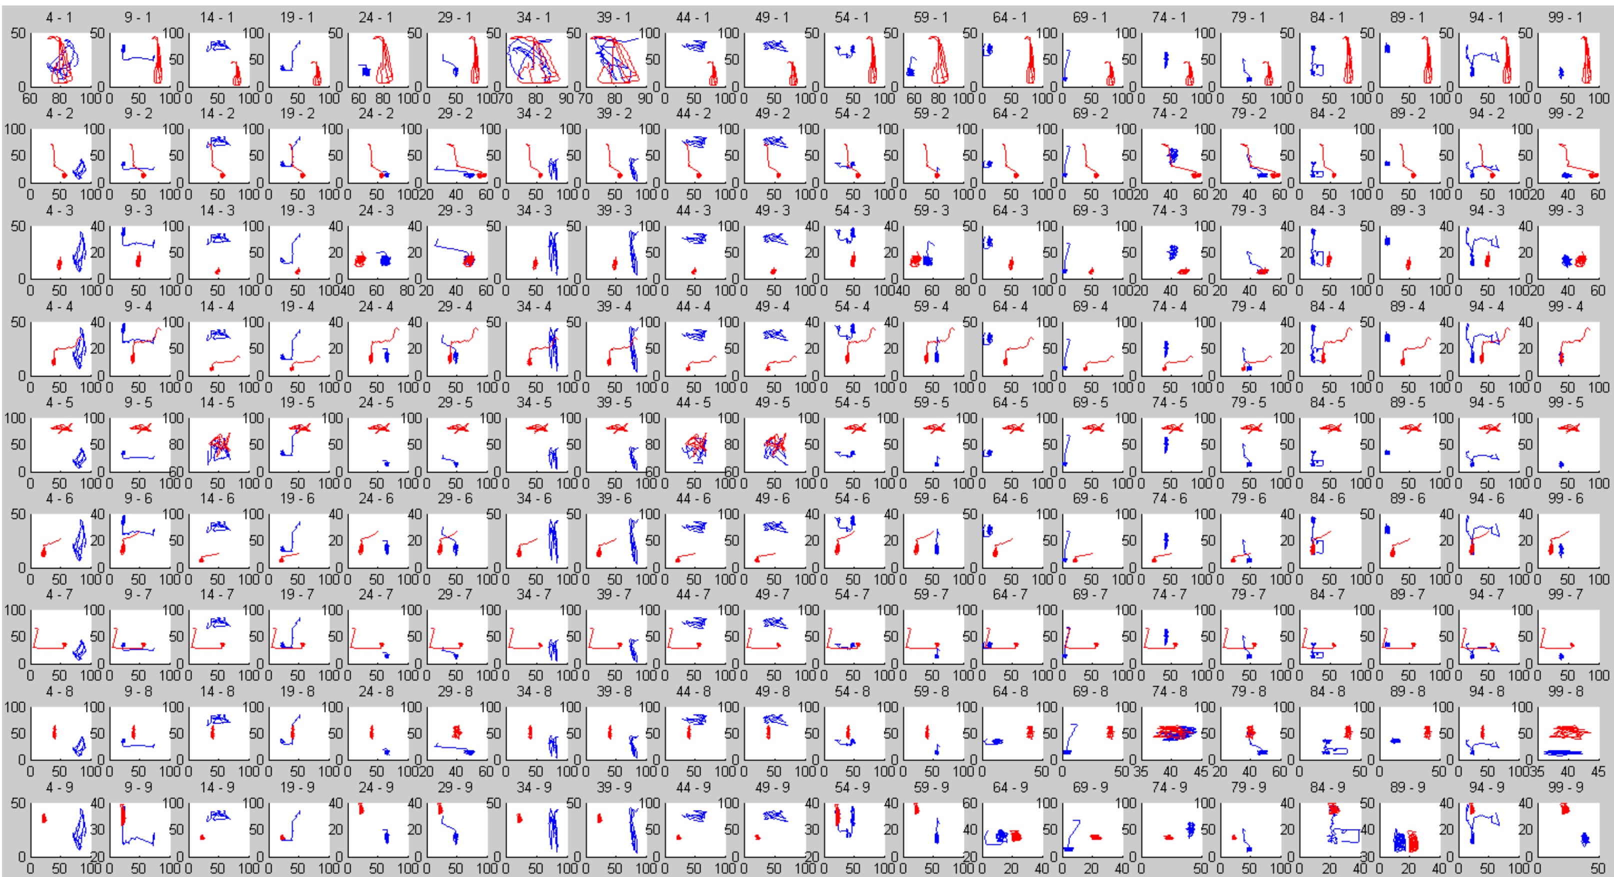

Layer 1: CHMM Input

- State presentation
  - Relative position (North, South, East, West)

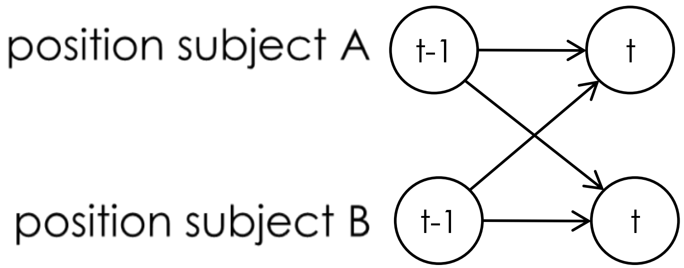

  - Relative speed and relative position

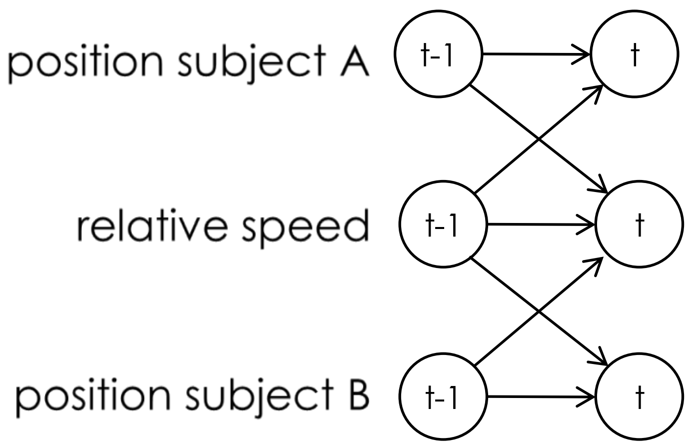

Layer 1: Experiments (Suspicious Events)

- Dataset
  - ~3000 interactions
  - 81.7% normal
  - 18.3% suspicious
- Results – suspicious events

	CHMM relative position	**CHMM relative position and speed**
Recall	38.75	61.38
Precision	28.44	29.94
F-measure	32.80	40.25

Layer 2: Baseline

- Results – suspicious passengers
  - If ∃ suspicious event → person is suspicious

	CHMM relative position	**CHMM relative position and speed**
Recall	90.00	100.00
Precision	18.37	20.00
F-measure	30.51	33.33

Layer 2: Baseline

- Results – suspicious passengers
  - If ∃ suspicious event → person is suspicious

	CHMM relative position	**CHMM relative position and speed**	**HMM Relative position**
Recall	90.00	100.00	66.23
Precision	18.37	20.00	10.42
F-measure	30.51	33.33	18.01

Layer 2: Layered HMMs (Oliver et al., 2004)

- Consists of N levels of HMMs
- Models hierarchical relations
- Equivalent to single HMM, but
  - Less likely to over-fit
  - Less training data
  - Can be (re)trained separately

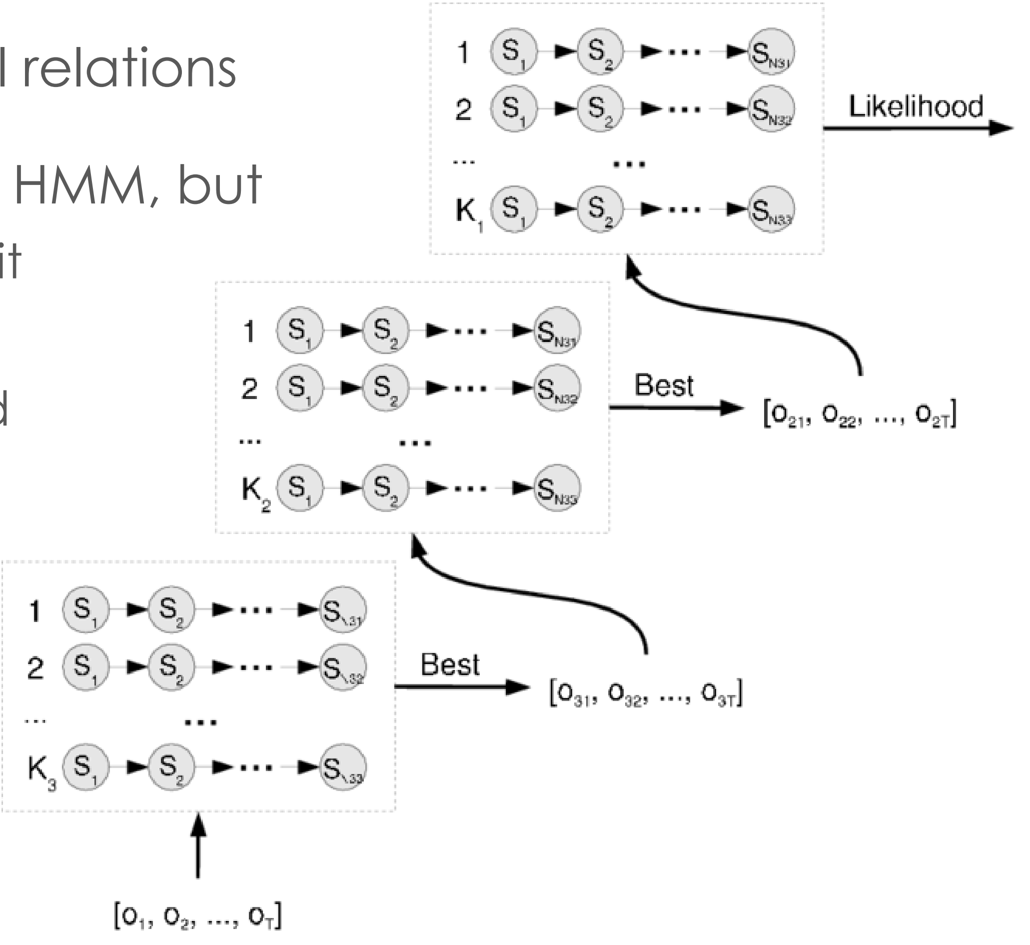

Layer 2: Layered HMMs

- Two layers
- Layer 1: Detection of suspicious events
  - Two CHMM: normal, suspicious events
- Layer 2: Detection of suspicious behavior
  - Two HMMs: normal, suspicious behavior
  - Select argmax($HMM_i(s)$)

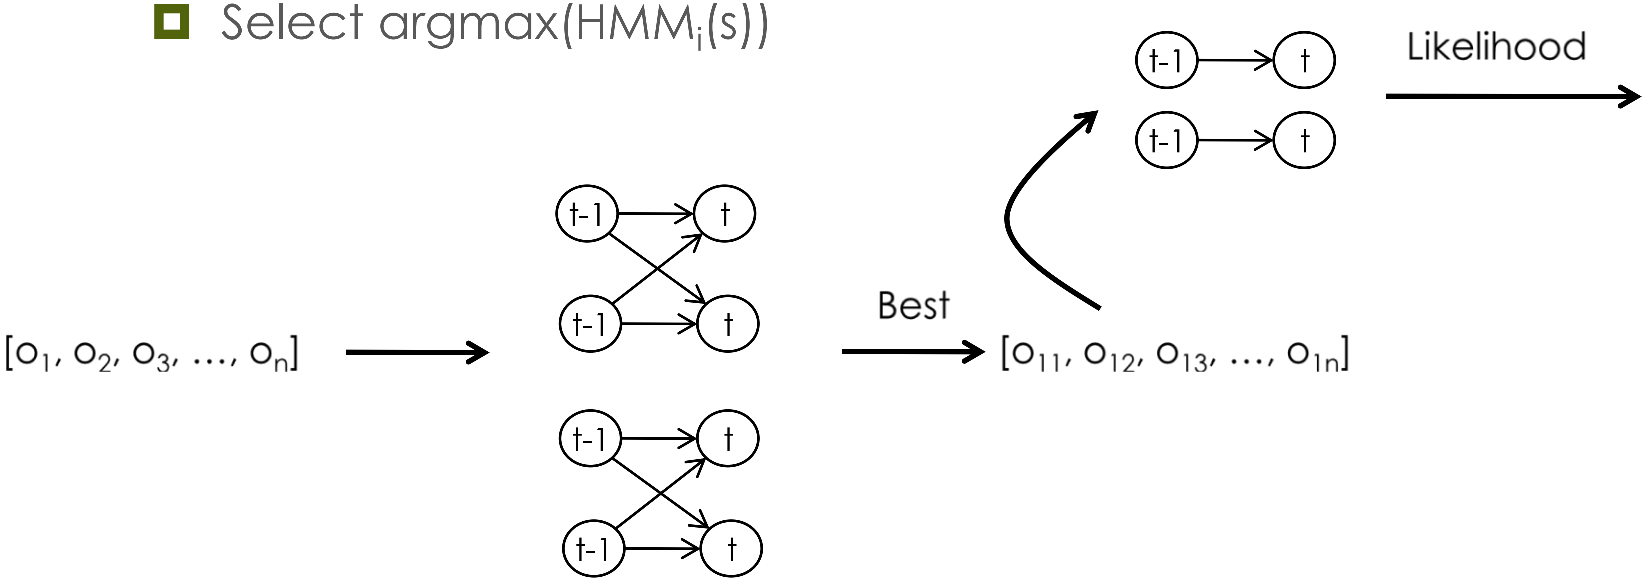

Layer 2: Utility-based plan recognition

(Kaminka et al., 2007)

- Idea
  - Low-likelihood events may be overlooked
  - UPR incorporates observer biases
  - Choose hypothesis by expected utility to the observer.
- Probability normalization

$$S'(e) = \frac{S(e)}{S(e) + N(e)}$$

$$N'(e) = \frac{N(e)}{S(e) + N(e)}$$

- Multiply
  - Posterior probability of an event
  - Utility to the observer

$$UPR(s) = \sum_{e \in s} S'(e) * u_S + N'(e) * u_N$$

Layer 2: Utility-based plan recognition

- Example CHMM + UPR
  - u_N=-1, u_S=10

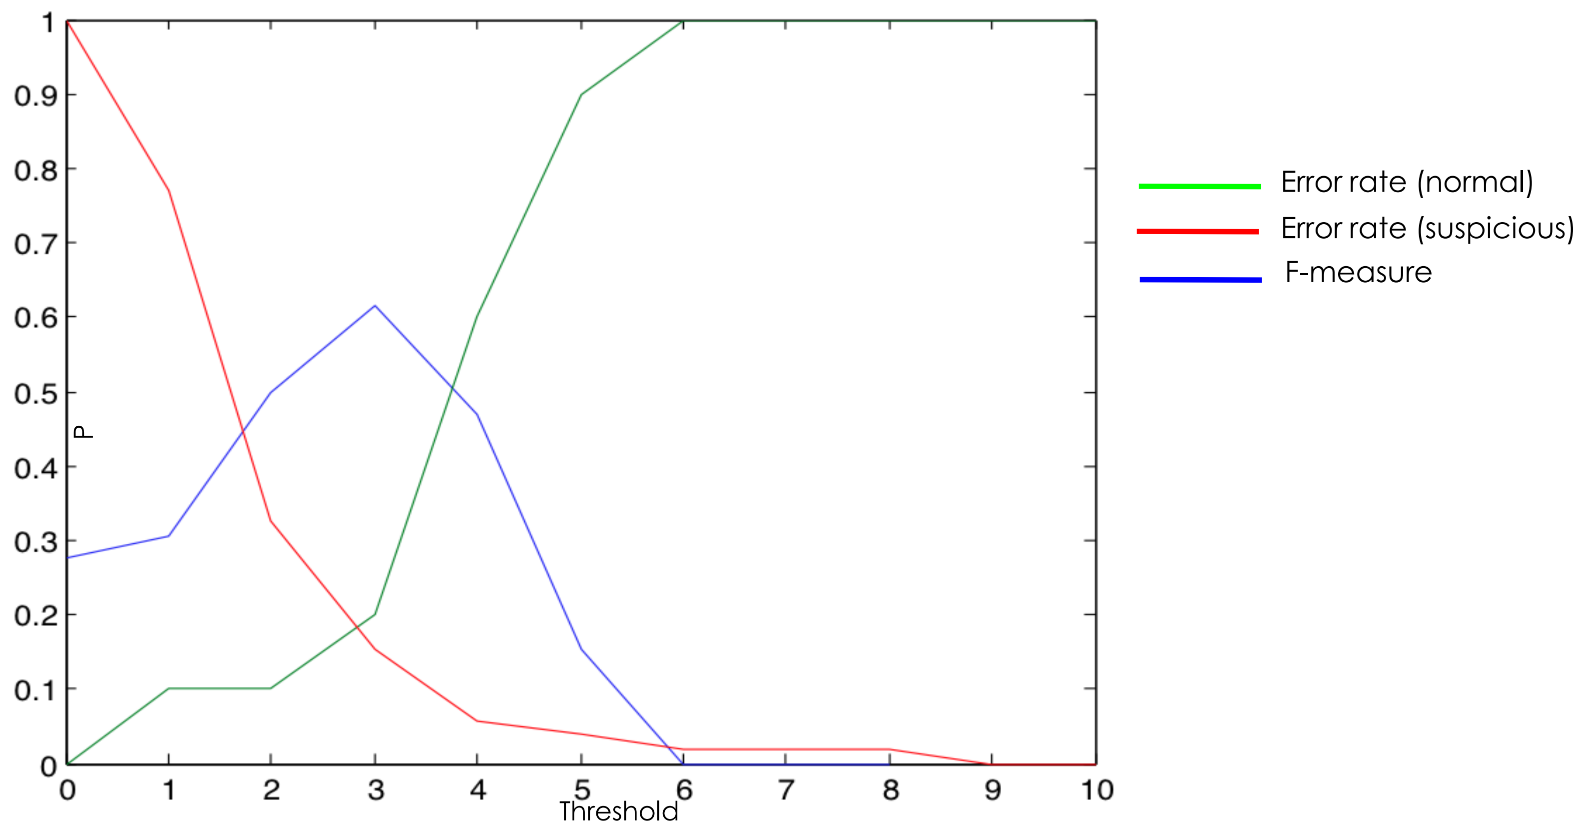

Layer 2: Penalty-based accumulation

- Idea
  - Suspicious events increase overall suspiciousness
  - Normal events decrease overall suspiciousness

- Probability normalization

$$r = pdf_{N(\sigma,\mu)}\left(\frac{S(e)}{N(e)}\right), \quad r \in [0,1]$$

Layer 2: Penalty-based accumulation

- Idea
  - Suspicious events increase overall suspiciousness
  - Normal events decrease overall suspiciousness
- Probability normalization

$$r = pdf_{N(\sigma,\mu)}(\frac{S(e)}{N(e)}), \quad r \in [0,1]$$

$$c_S(s_{1,t}) = \#\, of\ suspicious\ events \in s_{1,t}$$

$$c_N(s_{1,t}) = \#\, of\ normal\ events \in s_{1,t}\ after\ last\ suspicious$$

- Penalty function

$$PACC_{t+1} = PACC_t + [c_S(s_{1,t})^{A(r-0.5)} - 1]$$

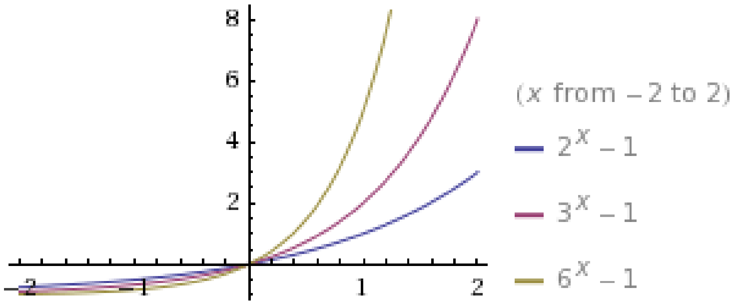

Layer 2: Penalty-based accumulation

- Idea
  - Suspicious events increase overall suspiciousness
  - Normal events decrease overall suspiciousness
- Probability normalization

$$r = pdf_{N(\sigma,\mu)}(\frac{S(e)}{N(e)}), \quad r \in [0,1]$$

$$c_S(s_{1,t}) = \#\, of\ suspicious\ events \in s_{1,t}$$

$$c_N(s_{1,t}) = \#\, of\ normal\ events \in s_{1,t}\ after\ last\ suspicious$$

- Penalty function

$$PACC_{t+1} = PACC_t + [c_S(s_{1,t})^{A(r-0.5)} - 1]$$

- Time decay

$$PACC_{t+1} = PACC_t \cdot (e^{-\frac{c_N(s_{1,t})}{B + c_S(s_{1,t})}})$$

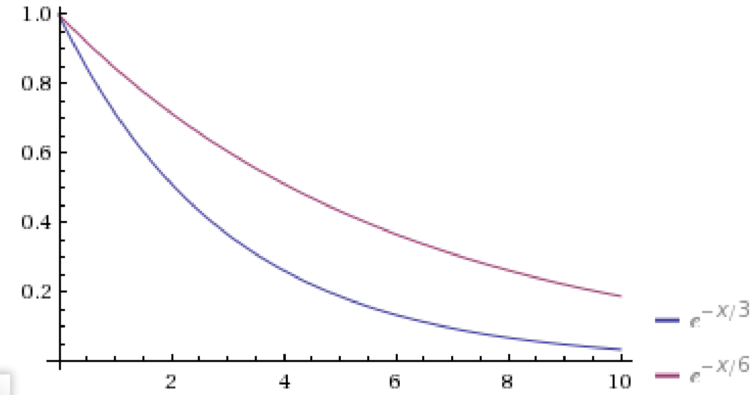

Layer 2: Penalty-based accumulation

- Example: detected events
  [s, n, n, n, n, n, s, n, n, n, n, n, s, n, n, n, n, n]

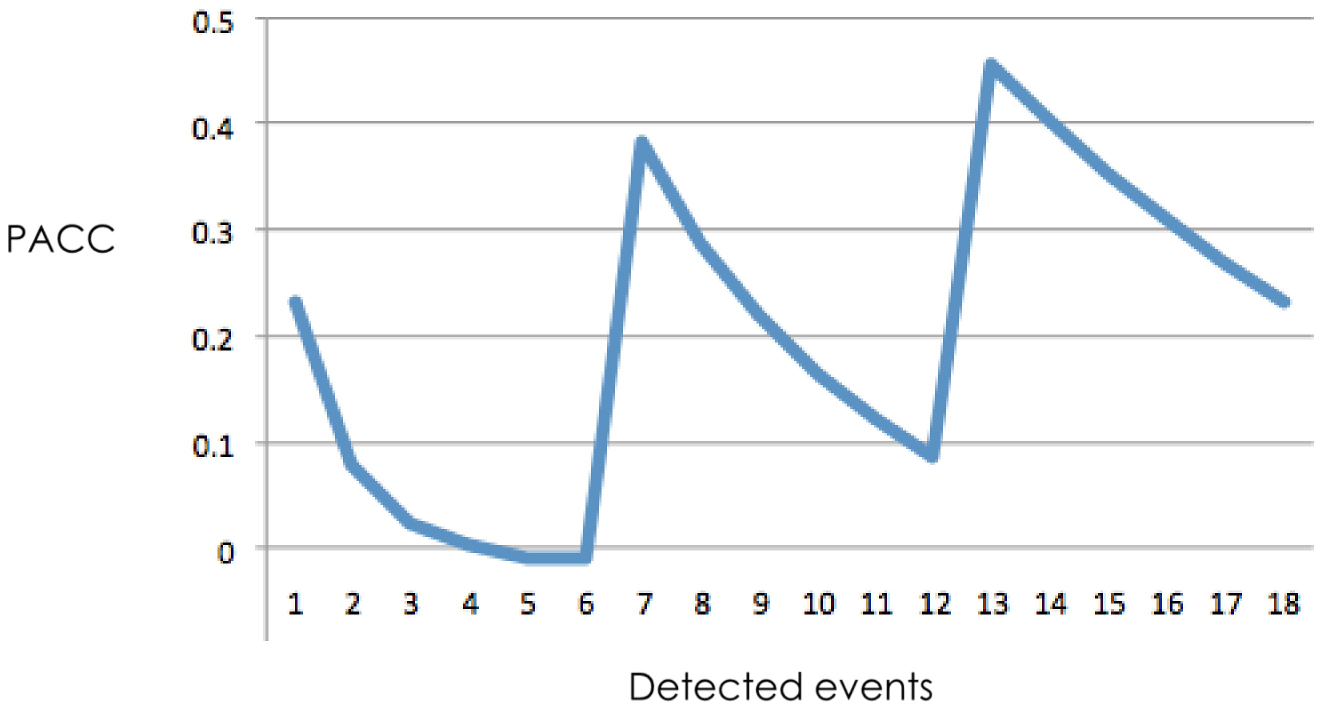

Results

Level 1	Level 2	Relative position	Relative position and speed
CHMM	If ∃suspicious event	R: 90.00 P: 18.37 F: 30.51	R: 100.0 P: 20.00 F: 33.33
	LHMM	R: 70.00 P: 31.82 F: 43.75	R: 50.00 P: 25.00 F: 33.33
	UPR	R: 70.00 P: 70.00 F: 70.00	R: 80.00 P: 50.00 F: 61.54
	PACC	R: 80.00 P: 50.00 F: 61.54	R: 90.00 P: 50.00 F: 64.29

Results

Level 1	Level 2	Relative position	Relative position and speed
CHMM	If ∃suspicious event	R: 90.00 P: 18.37 F: 30.51	R: 100.0 P: 20.00 F: 33.33
	LHMM	R: 70.00 P: 31.82 F: 43.75	R: 50.00 P: 25.00 F: 33.33
	UPR	R: 70.00 P: 70.00 F: 70.00	R: 80.00 P: 50.00 F: 61.54
	PACC	R: 80.00 P: 50.00 F: 61.54	R: 90.00 P: 50.00 F: 64.29
HMM		R:66.23 P: 10.42 F: 18.01	--

False positives

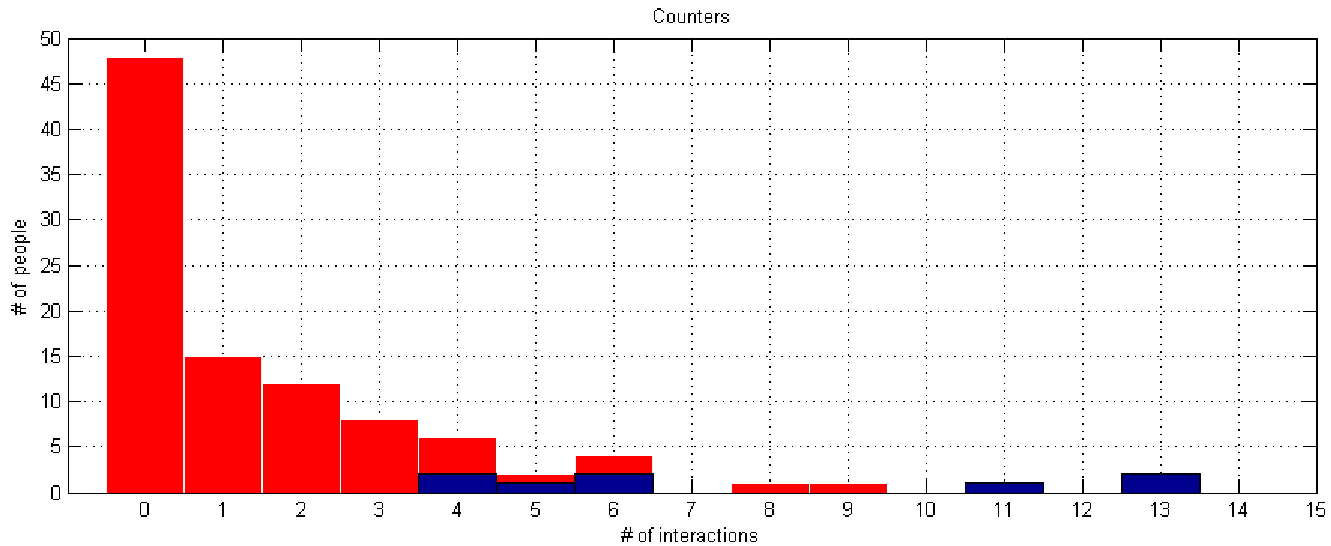

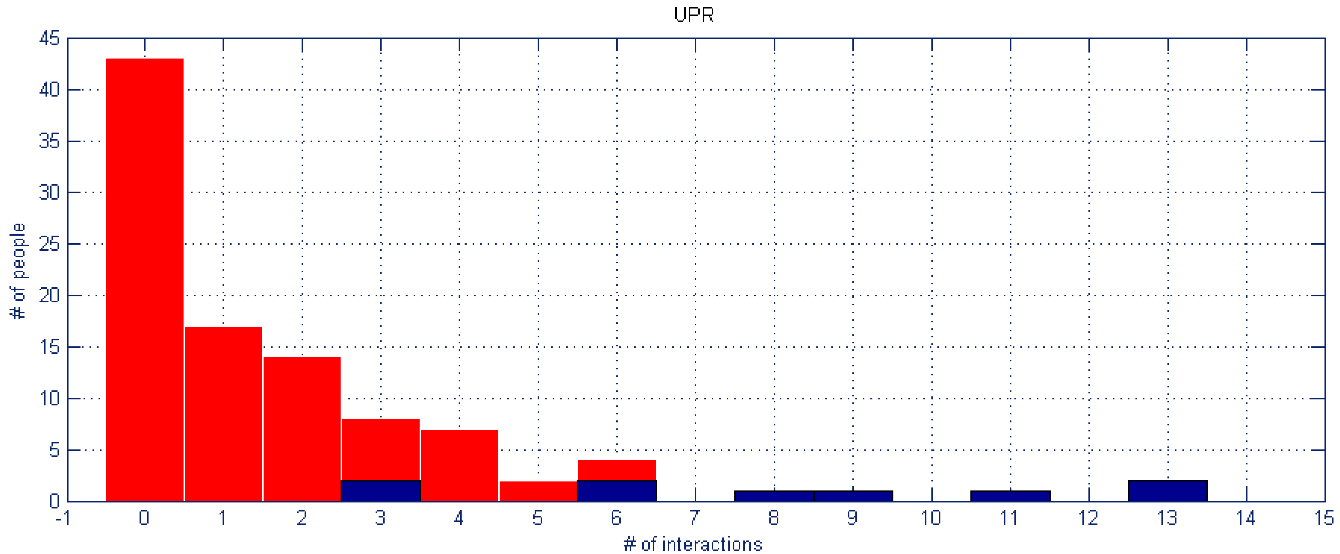

Conclusions

- Simulated suspicious behavior is convincing
  - Incorporate additional constraints (occlusion, uncertainty, real-world noise, etc.)
- HMM approach performs poor
  - Recall: ~66%, precision: ~10%
- Two-layered approach
  - CHMM + accumulation layer
  - Recall: >80%, precision: >50%
  - UPR and PACC outperform LHMM
  - Potential to exploit UPR and PACC

Future plans

- Experiments with real-world data
  - Near-realistic noise
  - Real airport data (LAX?)
- Recognize and connect multiple suspicious events
  - Improve detection of suspicious events
  - Recognize more events
- Actively trigger events to confirm hypothesis
  - Send authority to interact with passenger