Adaptive embedded systems
Two applications to society

Devika Subramanian
Department of Computer Science
Rice University
http://www.cs.rice.edu/~devika

RICE

Socially relevant computing

- It is computing ……
  - for a cause, for a purpose.
    - Can we evacuate Houston in 72 hours?
    - Can we predict the efficacy of a cancer drug for patients by using their genomic and proteomic profiles?
  - that meets a need in some context.
    - How can computation help me organize my music, my thoughts?
- *Views computer science as the primary technological enabler for solving real-world problems; creating needs and meeting them!*

Embedded Systems

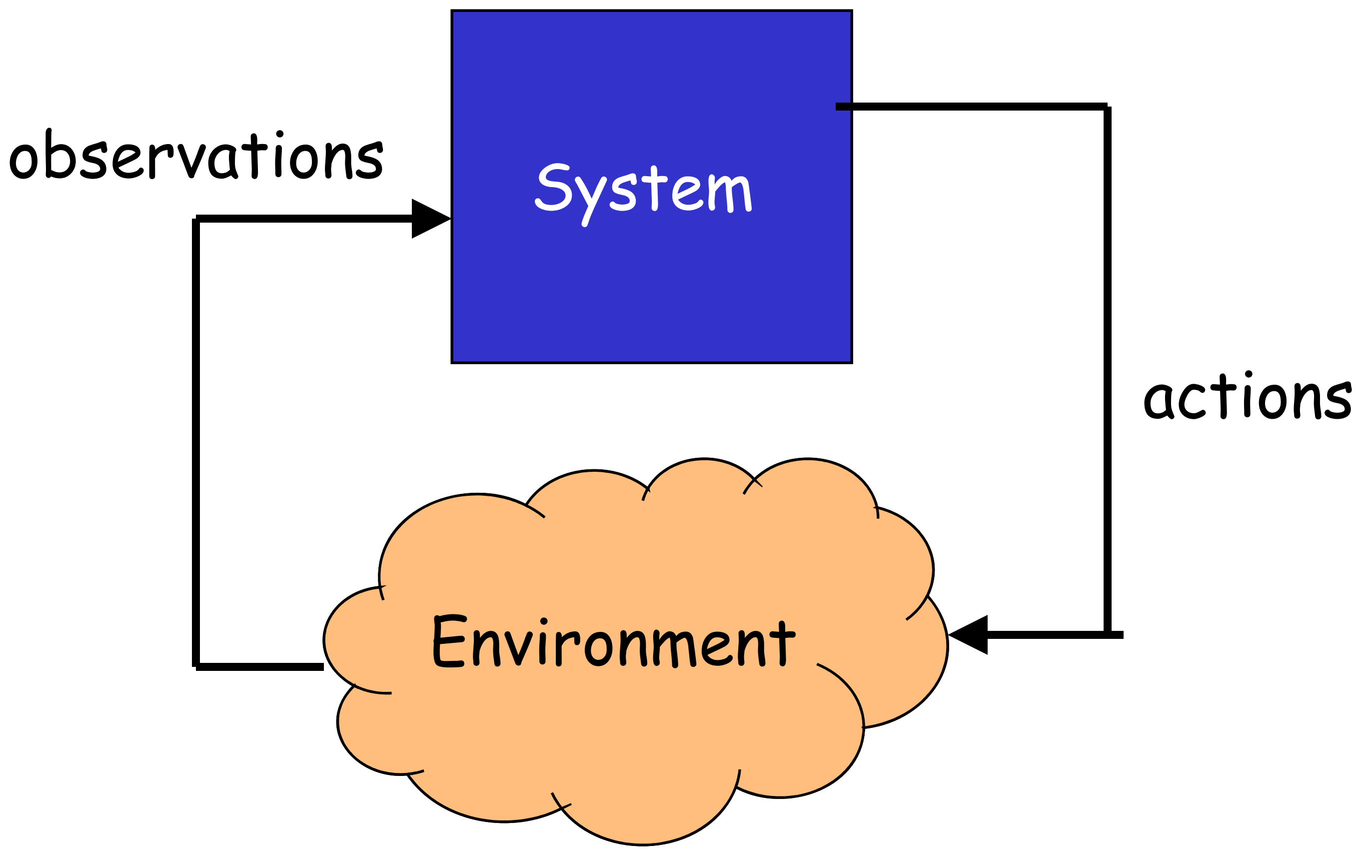

Observation driven, task-specific decision making

Machine Learning

Embedded Adaptive Systems

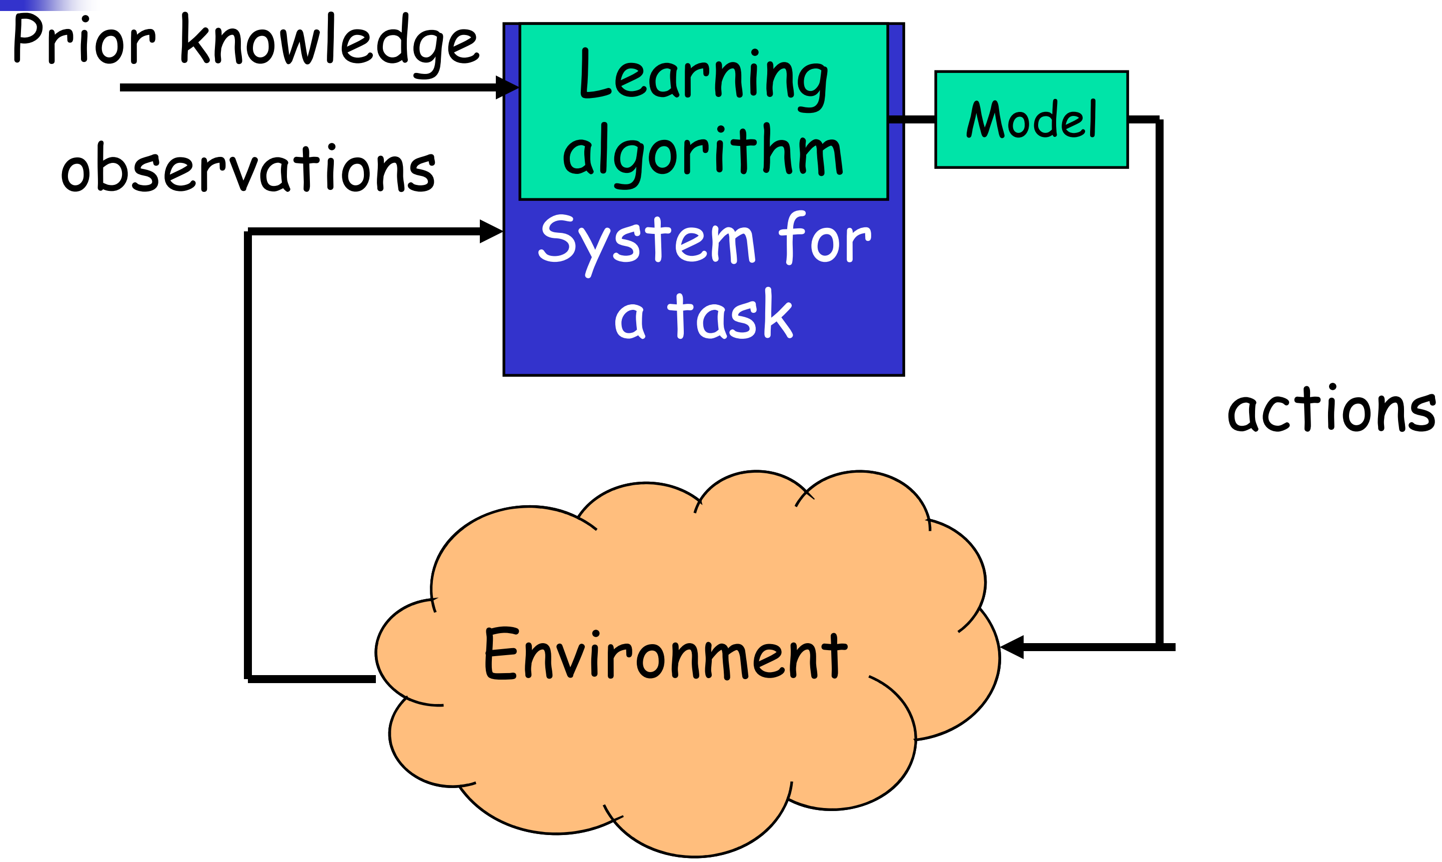

Calculate decisions on the basis of learned models of systems

Why embed learning?

- We cannot calculate and implement an action-choice/decision-making strategy for the system at design time.
  - System dynamics are unknown/partially known.
  - System dynamics change with time.
  - A one-size-fits-all solution is not appropriate – customization is needed.

Research questions in adaptive embedded system design

- Representation: What aspects of the task, environment and system dynamics do we need to observe and model for decision making?
- Learning: How can we build and maintain embedded models in changing environments?
- Decision making/acting: How can we use these models effectively to make decisions with scarce resources in changing environments?

Approach

- Design and validate algorithms for large-scale real world, socially relevant problems.
- Publish in the application community journals; get community to routinely use the methods.
- Abstract task-level analysis and present methods to the AI community.

Roadmap of talk

- Two case-studies
  - Unknown system, changing dynamics
    - Tracking human learning on a complex visual-motor task.
    - Predicting the evolution of international conflict.

Submarine School 101
The NRL Navigation Task

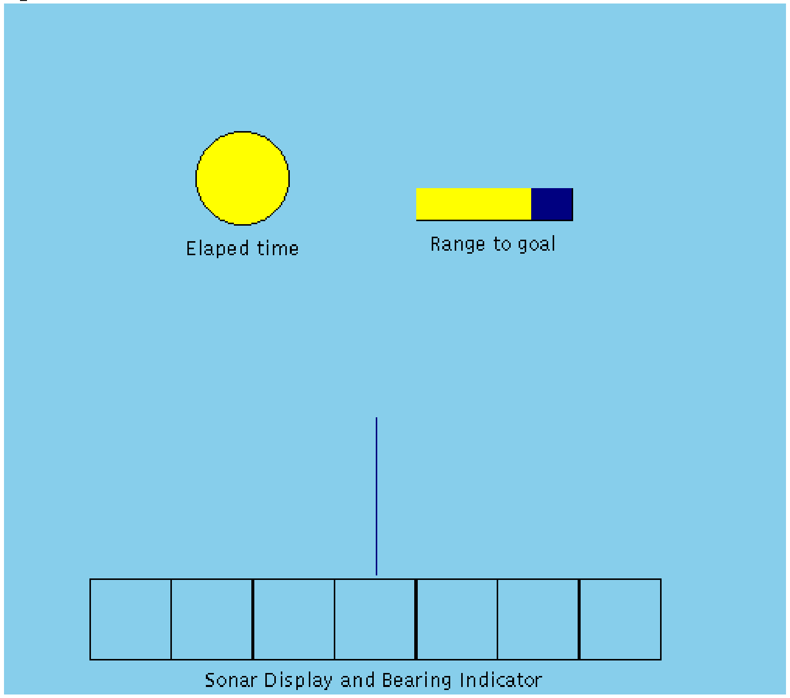

- Pilot a submarine to a goal through a minefield in a limited time period
- Distance to mines revealed via seven discrete sonars
- Time remaining, as-the-crow-flies distance to goal, and bearing to goal is given
- Actions communicated via a joystick interface

RICE

50% of class weeded out by this game!

The NRL Navigation Task

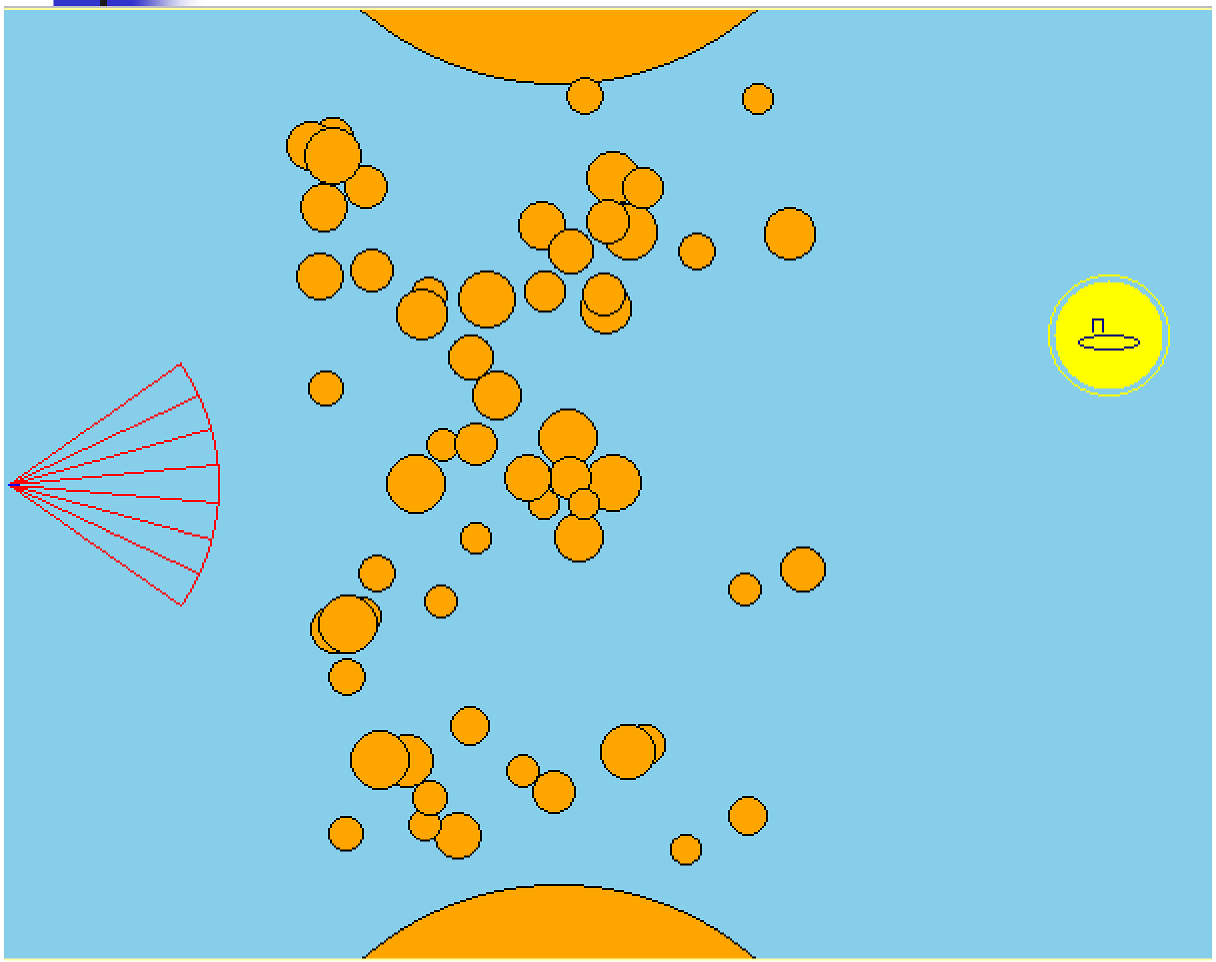

Mine configuration changes with every game.

Game has a strategic and a visual-motor component!

Learning curves

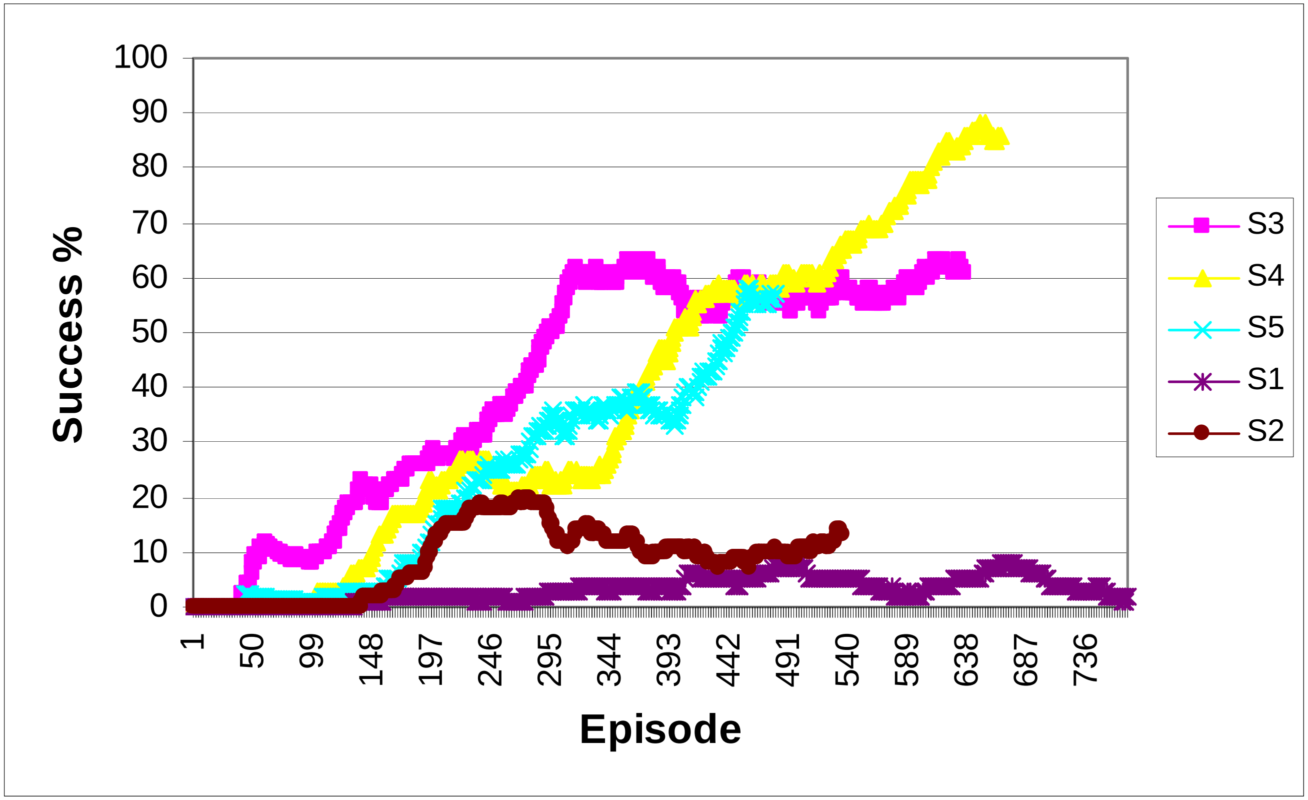

Successful learners look similar: plateaus between improvements

Unsuccessful learners are DOA!

Navy takes 5 days to tell if a person succeeds/fails.

Task Questions

- Is the game hard? What is the source of complexity?
- Why does human performance plateau out at 80%? Is that a feature of the human learning system or the game? Can machine learners achieve higher levels of competence?
- Can we understand why humans learn/fail to learn the task? Can we detect inability to learn early enough to intervene?
- How can we actively shape human learning on this task?

Mathematical characteristics of the NRL task

- A partially observable Markov decision process which can be made fully observable by augmentation of state with previous action.
- State space of size 10^{14}, at each step a choice of 153 actions (17 turns and 9 speeds).
- Feedback at the end of up to 200 steps.
- Challenging for both humans and machines.

Reinforcement learning

"a way of programming agents by reward and punishment without needing to specify *how* the task is to be achieved"

[Kaelbling, Littman, & Moore, 96]

Reinforcement learning

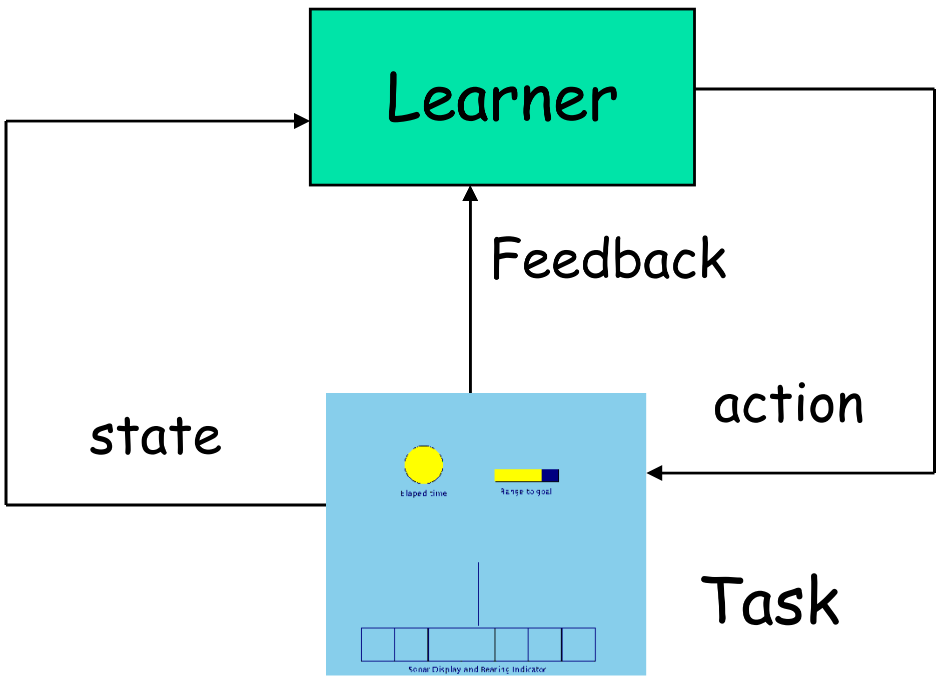

1. Observe state, s_t
2. Decide on an action, a_t
3. Perform action
4. Observe new state, s_{t+1}
5. Observe reward, r_{t+1}
6. Learn from experience
7. Repeat

$$\pi : S \rightarrow A$$

Reinforcement Learning, Barto and Sutton, MIT Press, 1998.

Reinforcement learning/NRL task

- Representational hurdles
  - State and action spaces have to be manageably small.
  - Good intermediate feedback in the form of a non-deceptive progress function needed.
- Algorithmic hurdles
  - Appropriate credit assignment policy needed to handle the two types of failures (timeouts and explosions are different).
  - Learning is too slow to converge (because there are up to 200 steps in a single training episode).

State space design

- Binary distinction on sonar: is it > 50?
- Six equivalence classes on bearing: 12, {1,2}, {3,4}, {5,6,7},{8,9}, {10,11}
- State space size = $2^7 * 6 = 768$.
- Discretization of actions
  - speed: 0, 20 and 40.
  - turn: -32, -16, -8, 0, 8, 16, 32.

Automated discovery of abstract state spaces for reinforcement learning, Griffin and Subramanian, 2001.

The dense reward function

r(s,a,s') = 0 if s' is a state where player hits mine.
= 1 if s' is a goal state
= 0.5 if s' is a timeout state

Feedback at the end

= 0.75 if s is an all-blocked state and s' is a not-all-blocked state
= 0.5 + Diff in sum of sonars/1000 if s' is an all-blocked state
= 0.5 + Diff in range/1000 + abs(bearing - 6)/40 otherwise

Useful feedback during play

Credit assignment policy

- Penalize the last action alone in a sequence which ends in an explosion.
- Penalize all actions in sequence which ends in a timeout.

Simplification of value estimation

- Estimate the average local reward for each action in each state.

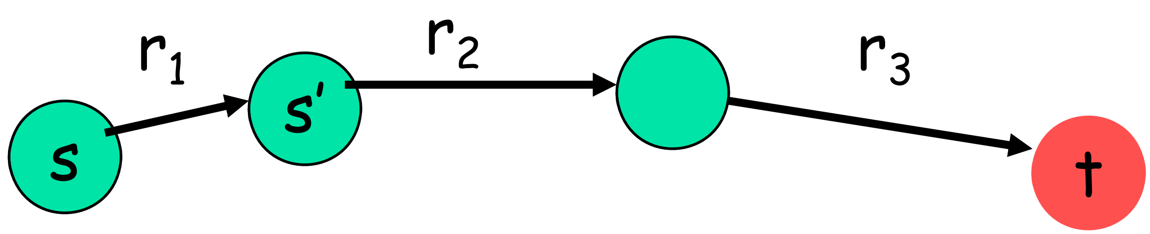

Instead of learning Q

$$Q(s,a) = \alpha[r + \max_{a'} Q(s',a')] + (1-\alpha)Q(s,a)$$

We maintain an approximation

$$Q^{'}(s,a) = (\text{running avg of rewards at } s \text{ for } a) * \text{pct of wins from } s$$

Q(s,a) = is the sum of rewards from s to terminal state.

Open question: When does this approx work?

Results of learning complete policy

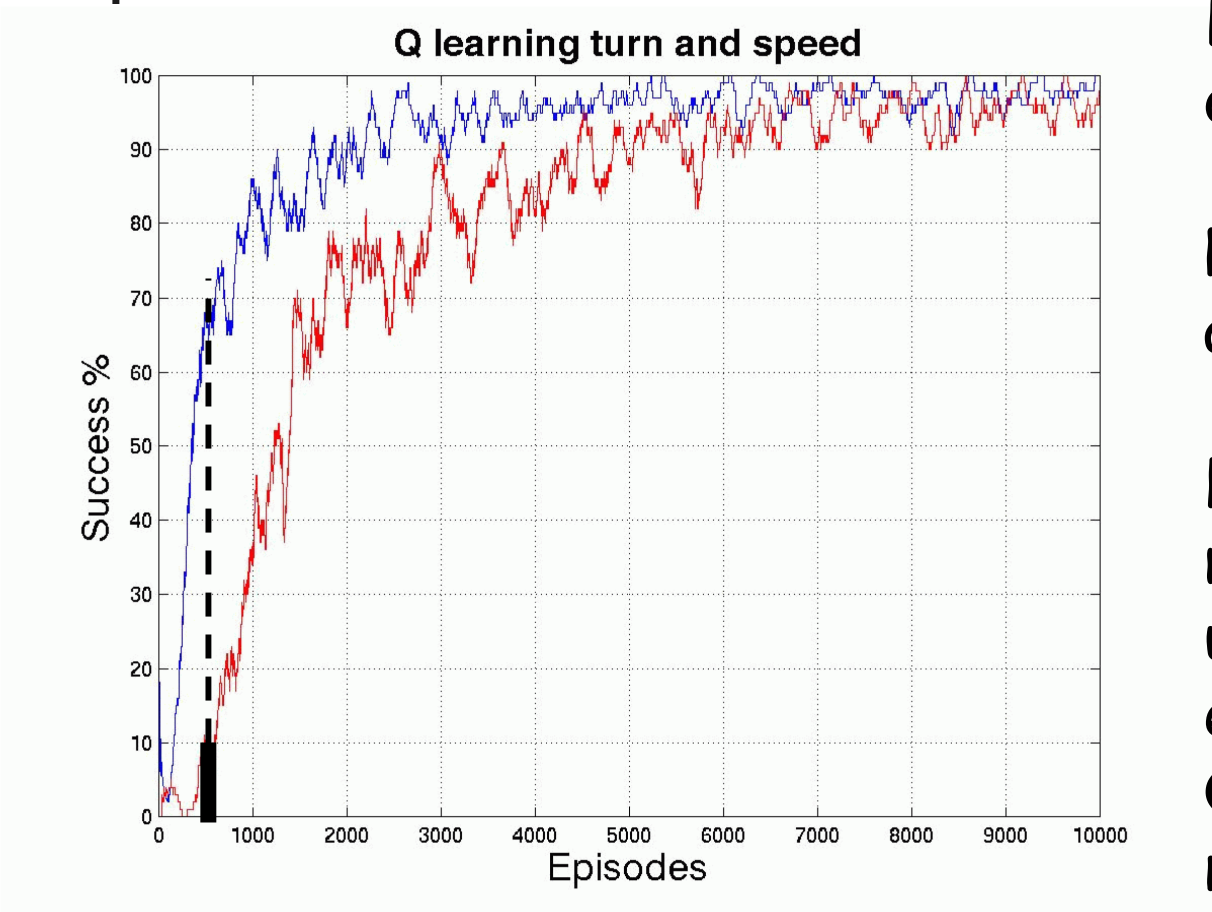

Blue: learn turns only

Red: learn turn and speed

Humans make more effective use of training examples. But Q-learner gets to near 100% success.

Griffin and Subramanian, 2000

Full Q learner/1500 episodes

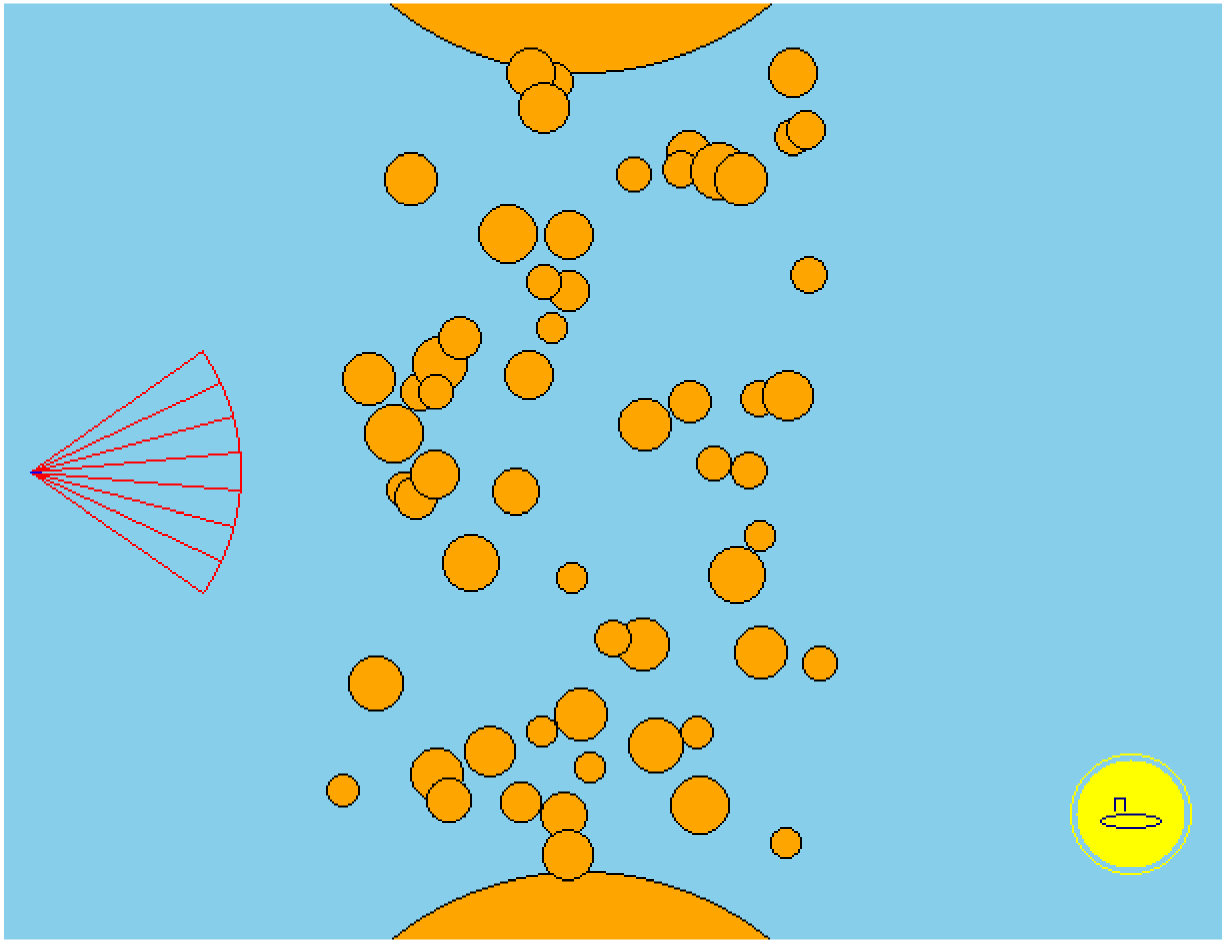

Full Q learner/10000 episodes

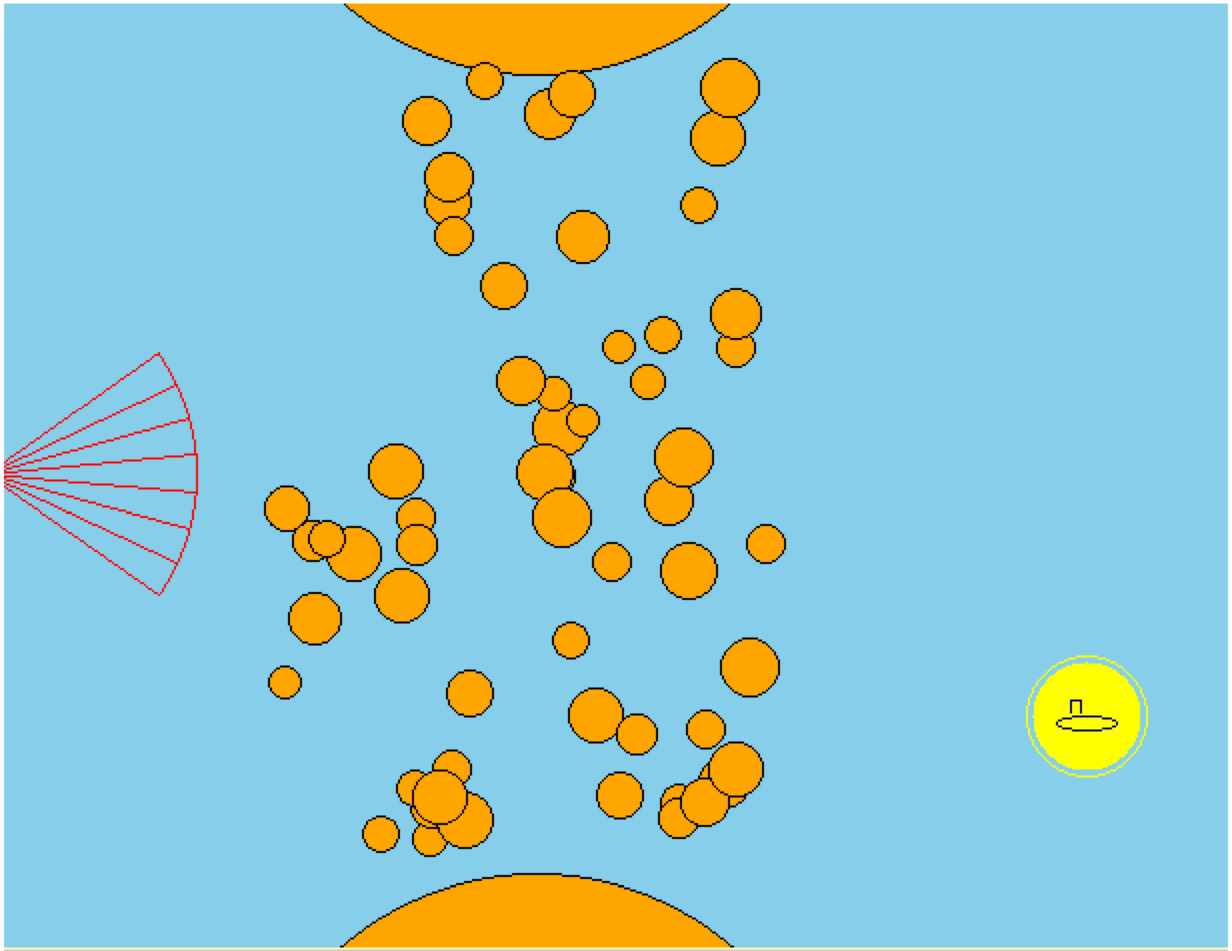

Full Q learner/failure after 10K

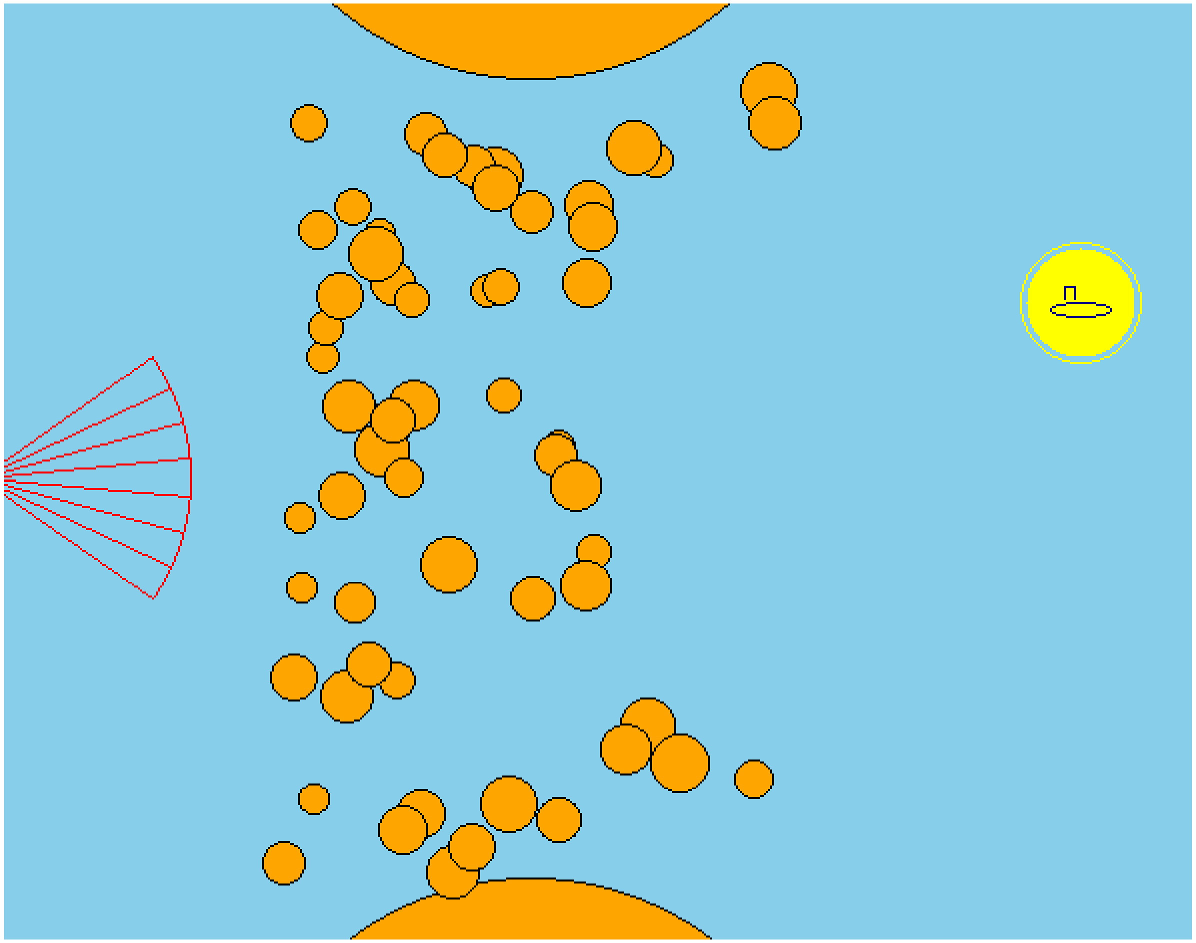

Why learning takes so long

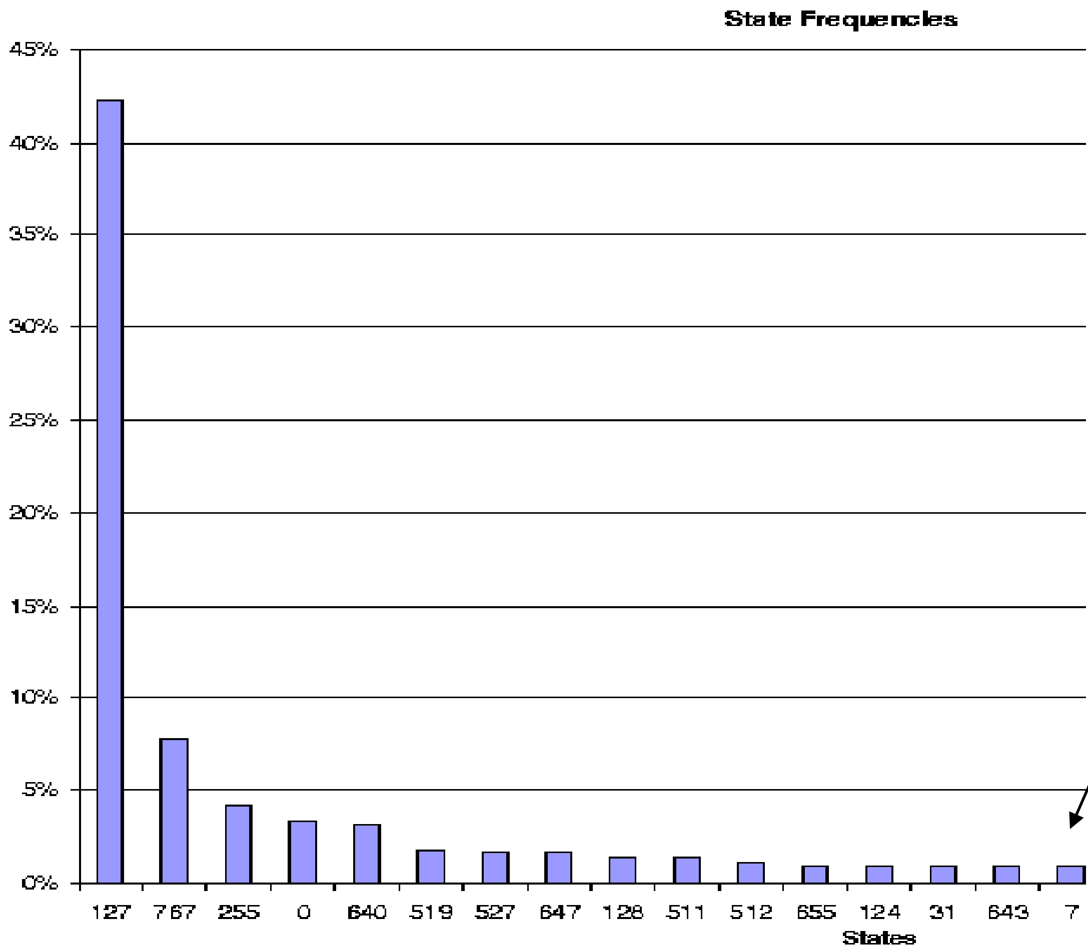

States where 3 or fewer of the 153 action choices are correct!

Griffin and Subramanian, 2000

Lessons from machine learning

- Task level
  - Task is hard because states in which action choice is critical occur less than 5% of the time.
  - Staged learning makes task significantly easier
  - A locally non-deceptive reward function speeds up learning.
- Reinforcement learning
  - Long sequence of moves makes credit assignment hard; a new cheap approximation to global value function makes learning possible for such problems.
  - Algorithm for automatic discretization of large, irregular state spaces.

Griffin and Subramanian, 2000, 2001

Task Questions

- Is the game hard? Is it hard for machines? What is the source of complexity?
- Why does human performance plateau out at 80%? Is that a feature of the human learning system or the game? Can machine learners achieve higher levels of competence?
- Can we understand why humans learn/fail to learn the task? Can we detect inability to learn early enough to intervene?
- How can we actively shape human learning on this task?

Tracking human learning

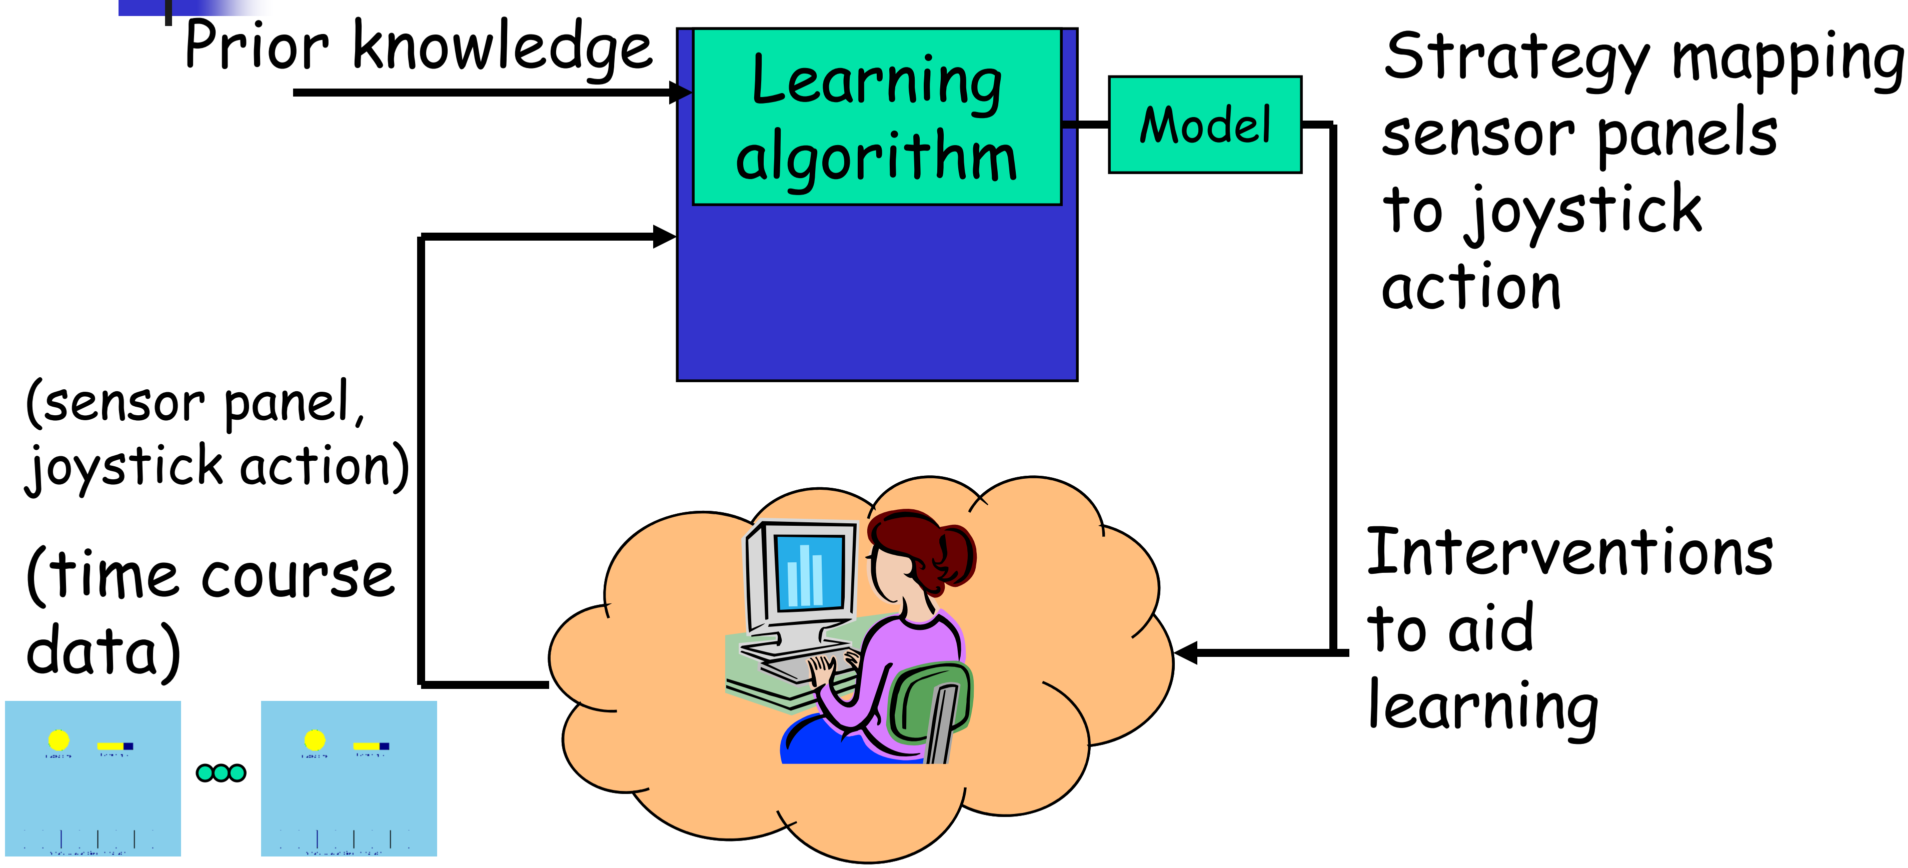

Extract strategy and study its evolution over time

Challenges

- High-dimensionality of visual data (11 dimensions spanning a space of size 10^{14})
- Large volumes of data
- Noise in data
- Non-stationarity: policies change over time

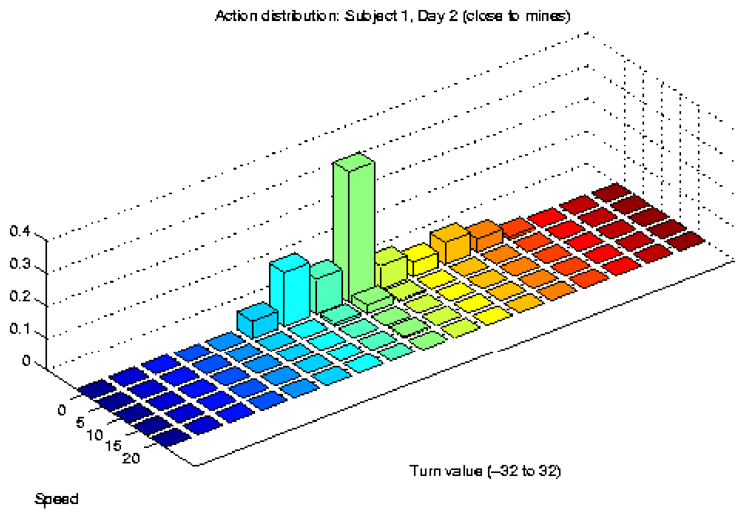

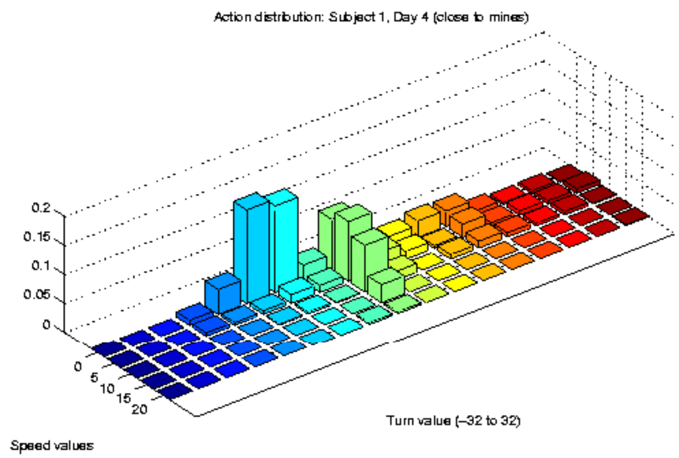

Embedded learner design

- Representation
  - Use raw visual-motor data stream to induce policies/strategies.
- Learning
  - Direct models: lookup table mapping sensors at time t and action at t-1 to distribution of actions at time t. (1st order Markov model)
- Decision-making
  - Compute "derivative" of the policies over time, and use it (1) to classify learner and select interventions, (2) to build behaviorally equivalent models of subjects

Strategy: mapping from sensors to action distributions

w

Window of
w games

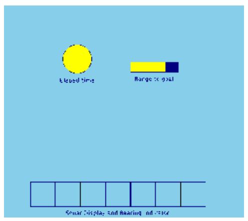

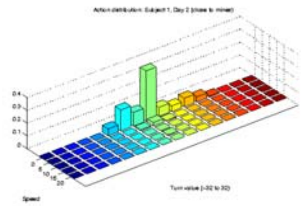

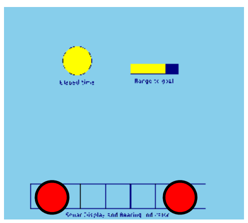

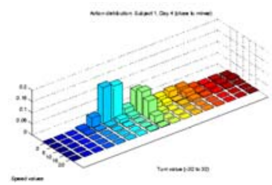

•

•

•

Surely, this can't work!

- There are 10^{14} sensor configurations possible in the NRL Navigation task.
- However, there are between 10^3 to 10^4 of those configurations actually observed by humans in a training run of 600 episodes.
- Exploit sparsity in sensor configuration space to build a direct model of the subject.

How do strategies evolve over time?

- Distance function between strategies: KL-divergence

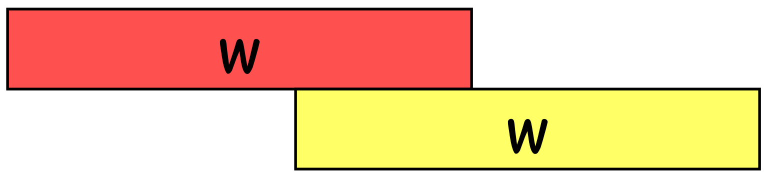

Overlap = s

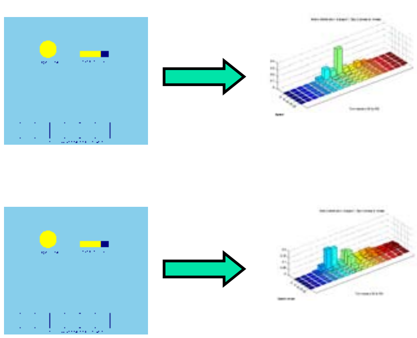

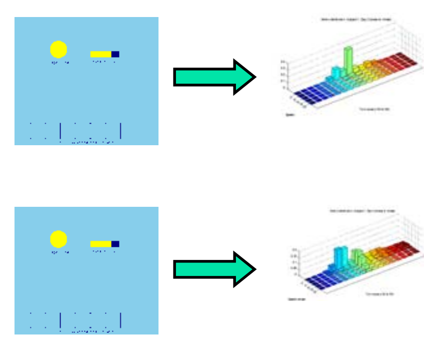

$$\Delta_P(\Pi(i, i+w), \Pi(i+w-s, i+2w-s))$$

Results: model derivative

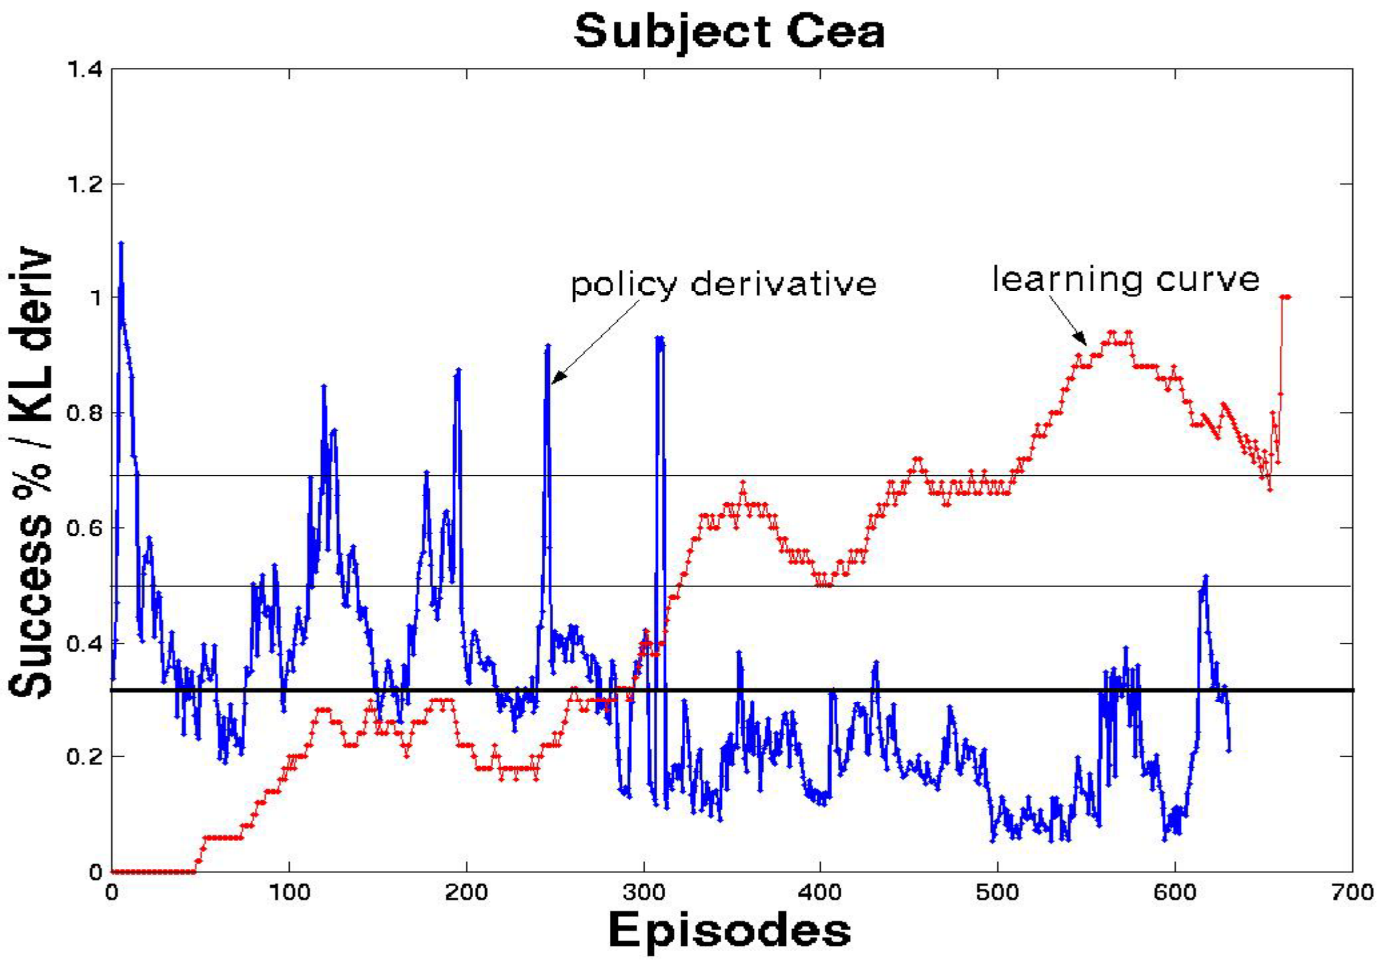

Siruguri and Subramanian, 2002

Before shift (episode 300)

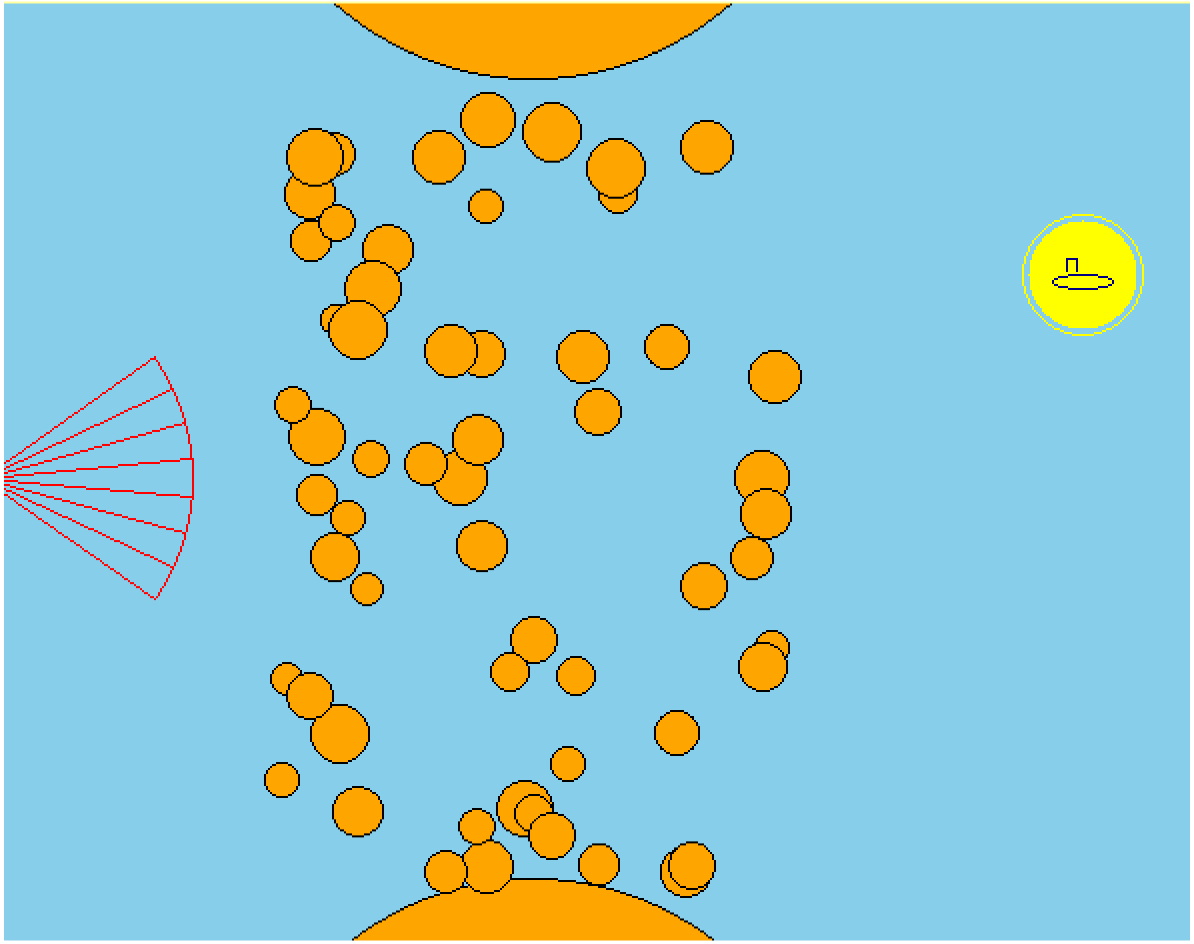

After shift (episode 320)

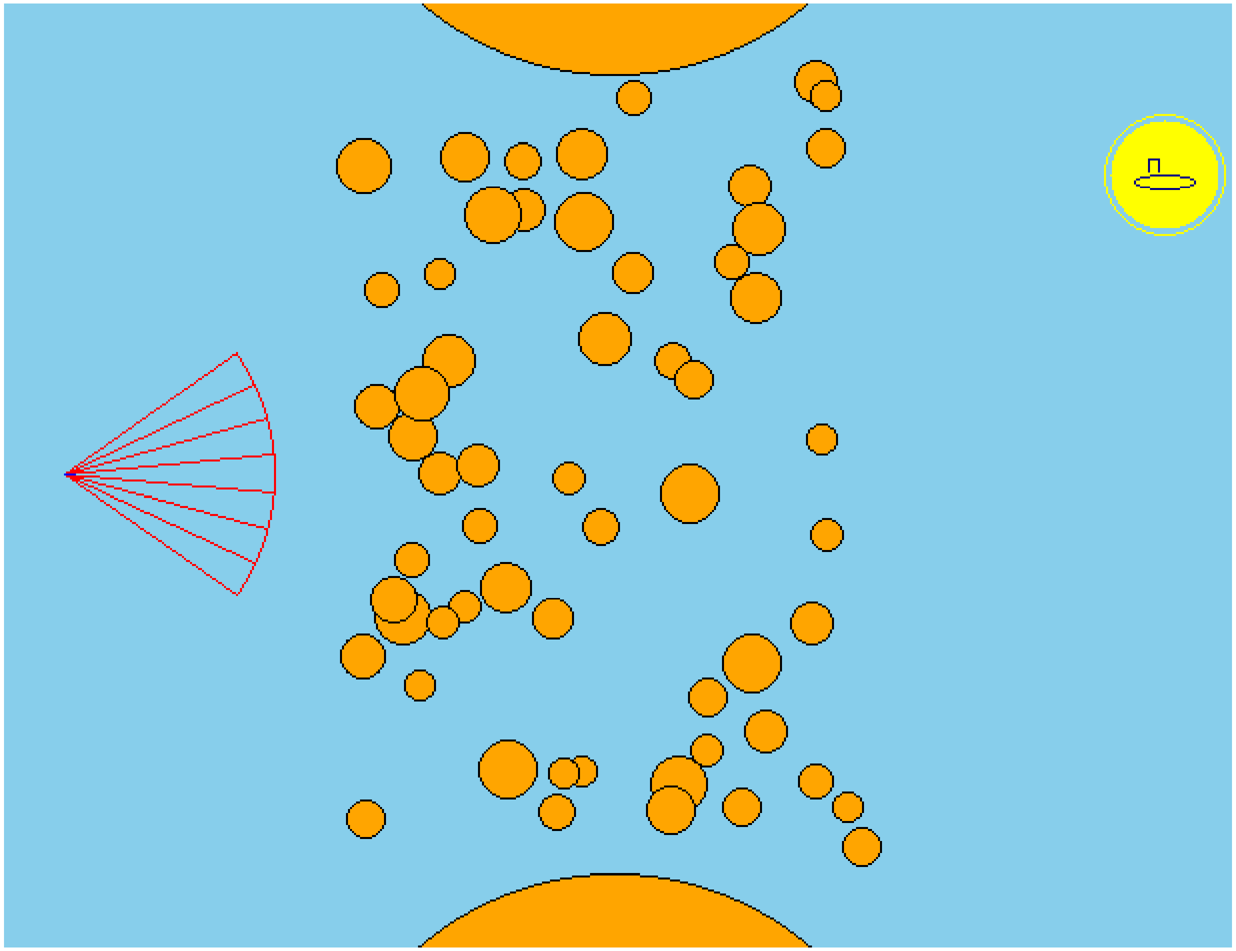

Model derivative for Hei

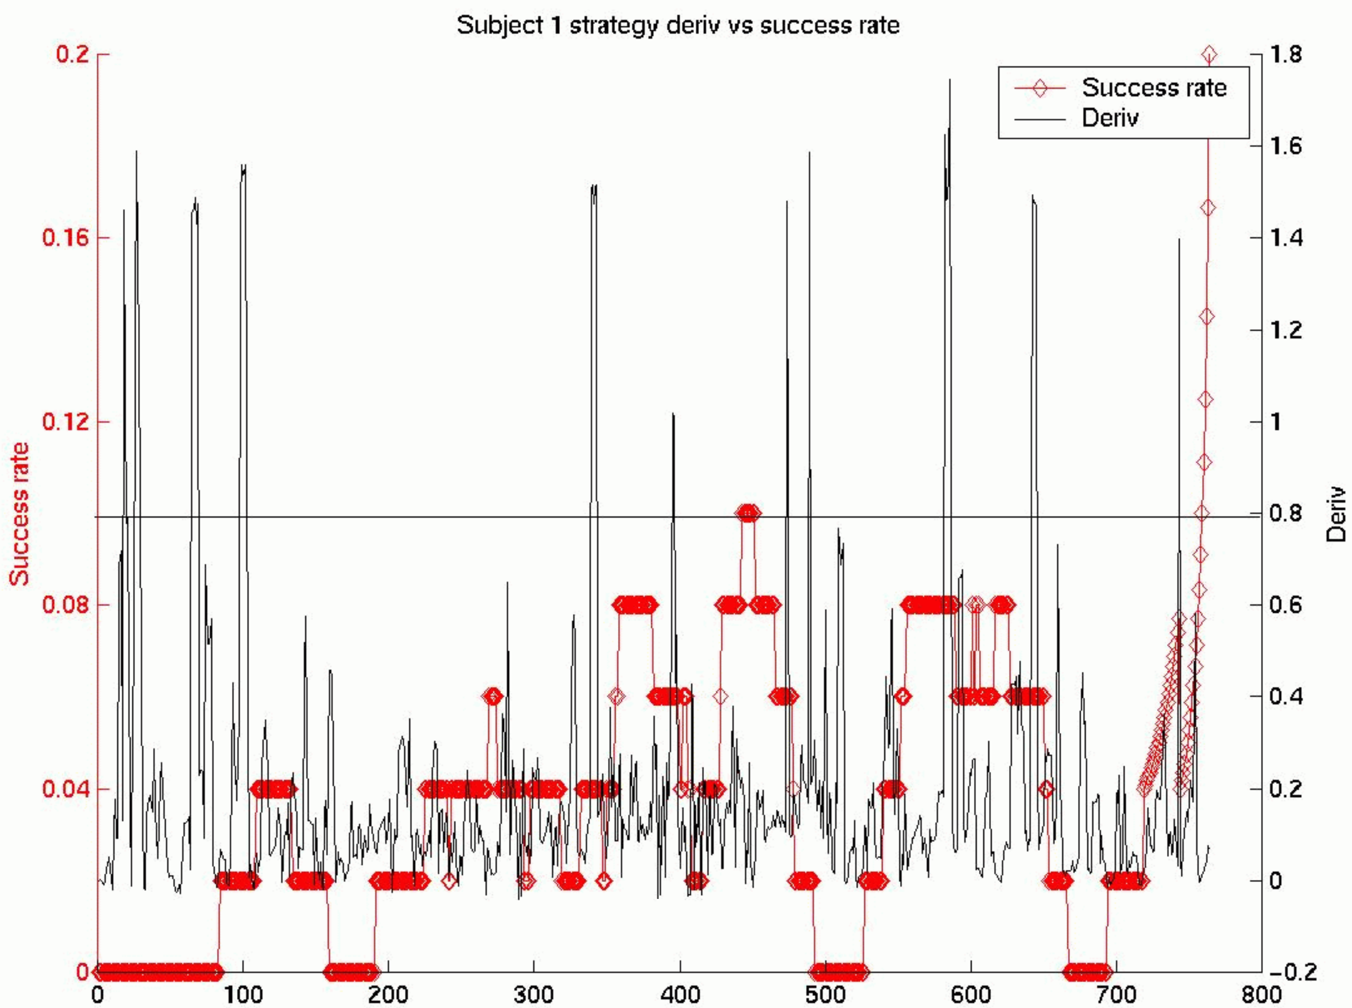

Siruguri and Subramanian, 2002

How humans learn

- Subjects have relatively static periods of action policy choice punctuated by radical shifts.
- Successful learners have conceptual shifts during the first part of training; unsuccessful ones keep trying till the end of the protocol!

Behaviorally equivalent models

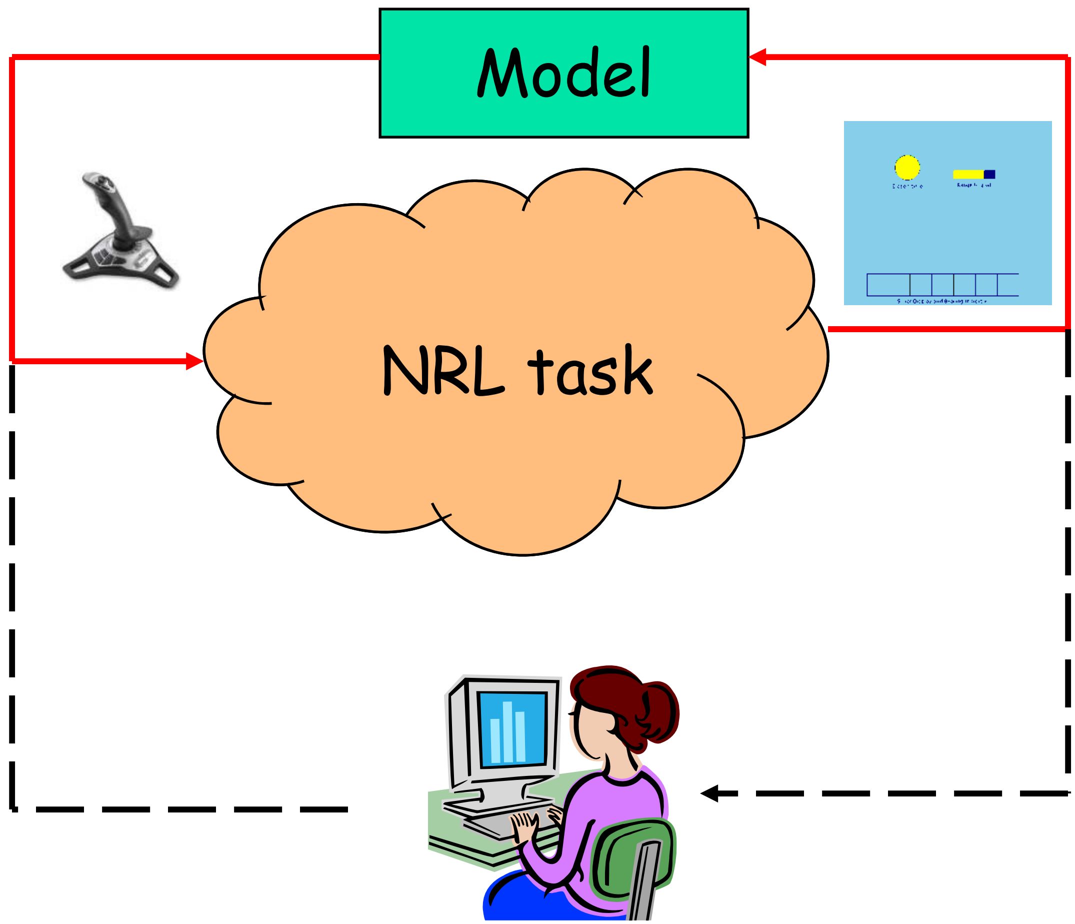

Generating behaviorally equivalent models

- To compute action a associated with current sensor configuration s in a given segment,
  - take 100 neighbors of s in lookup table.
  - perform locally weighted regression (LWR) on these 100 (s,a) pairs.

Subject Cea: Day 5: 1

Subject

Model

Subject Cea: Day 5: 2

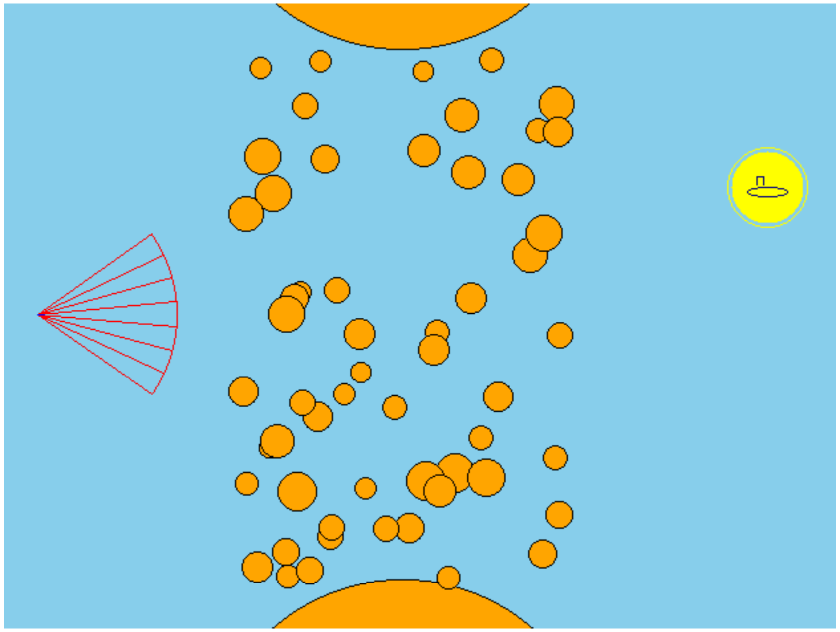

Subject

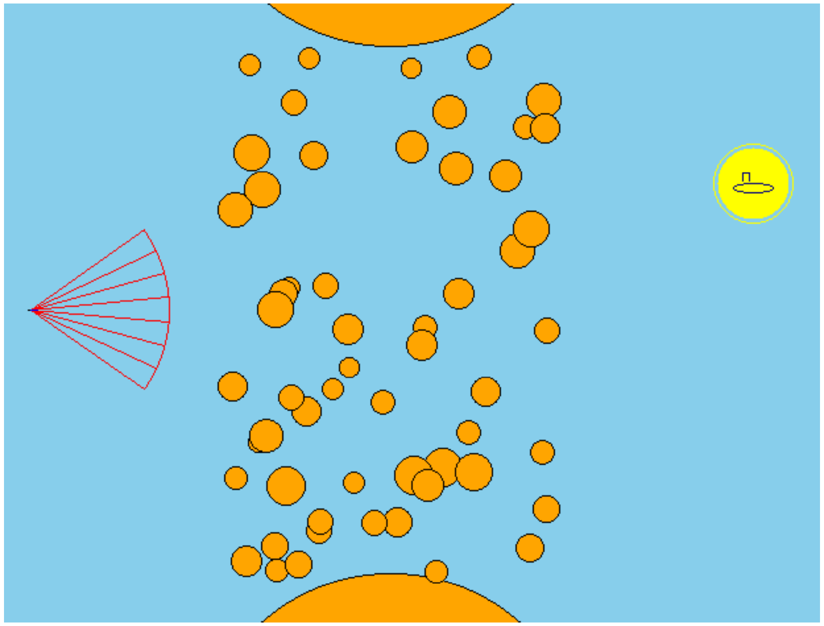

Model

Subject Cea: day 5: 3

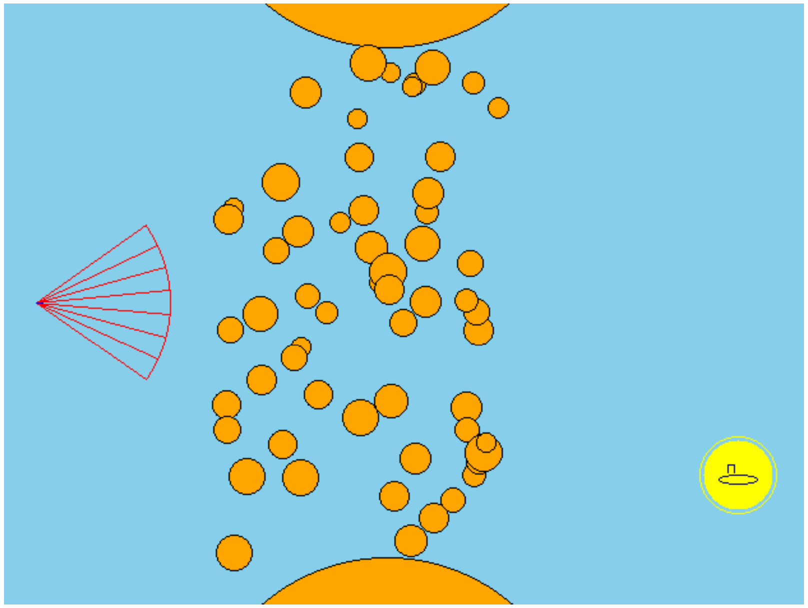

Subject

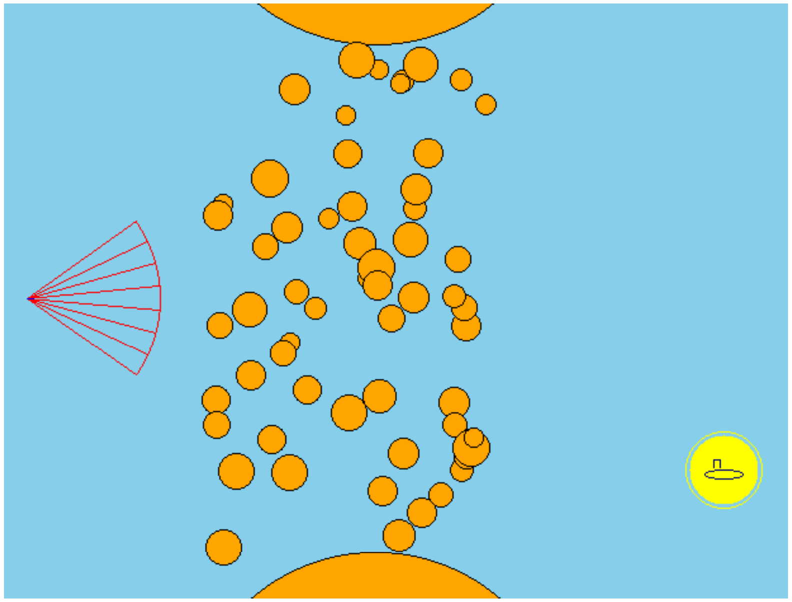

Model

Subject Cea: Day 5: 4

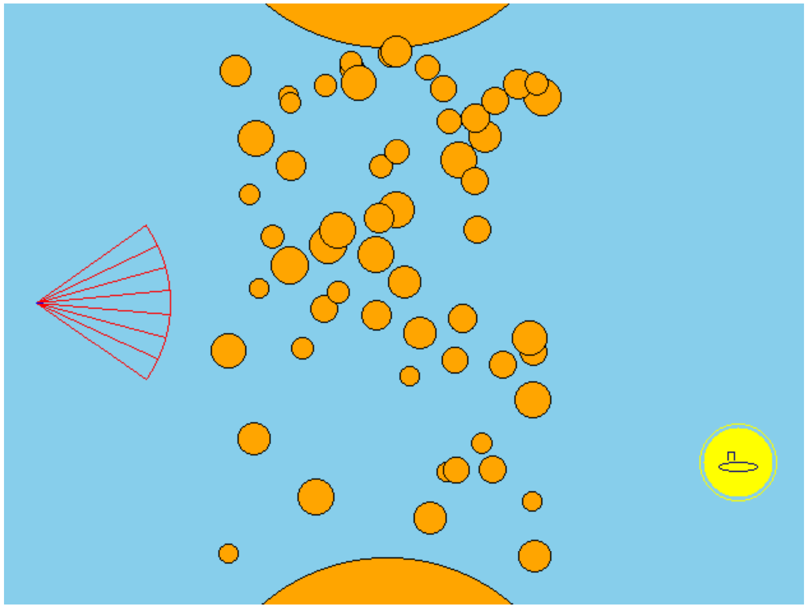

Subject

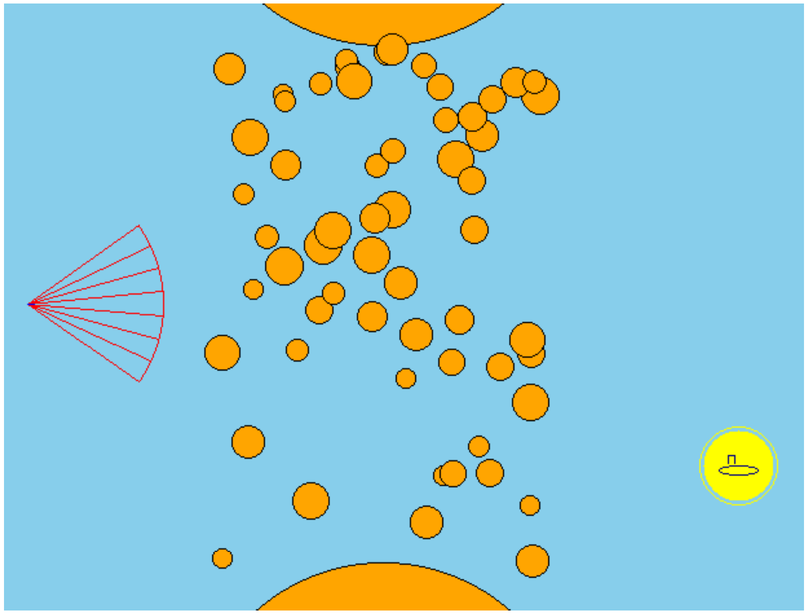

Model

Subject Cea: Day 5: 5

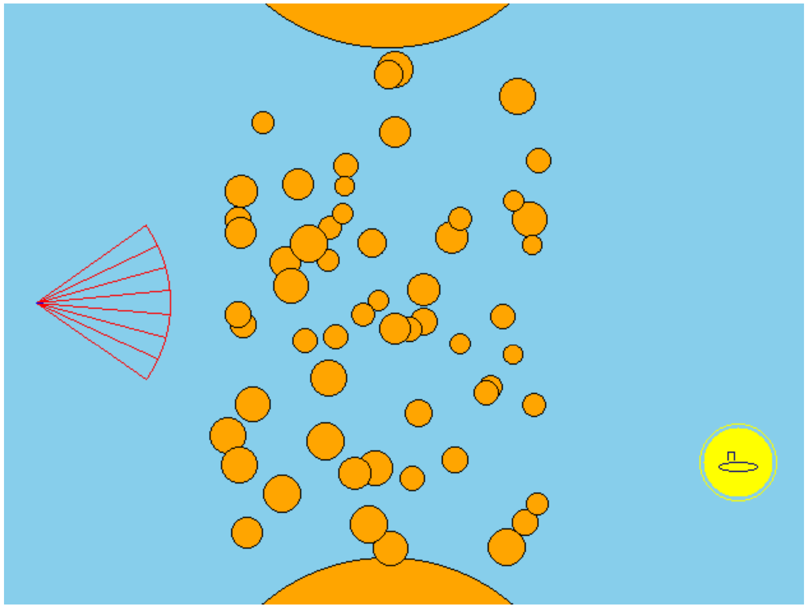

Subject

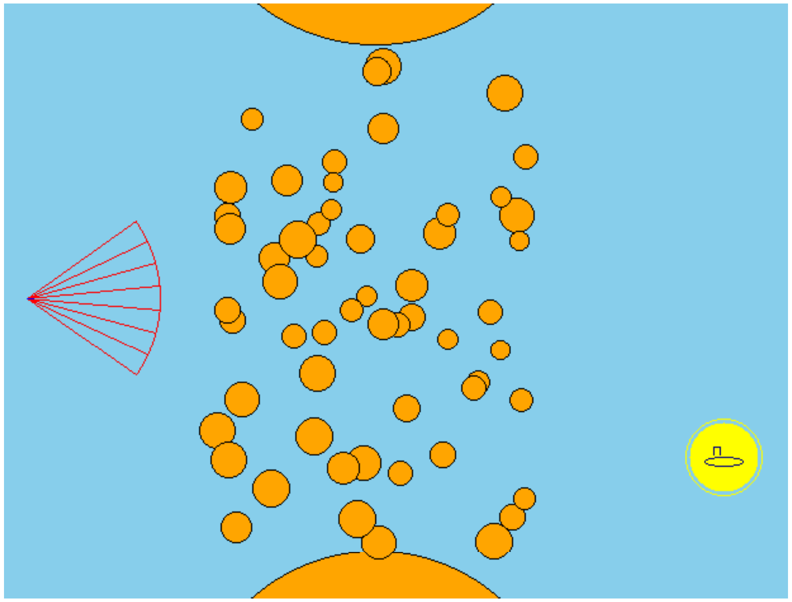

Model

Subject Cea: Day 5: 6

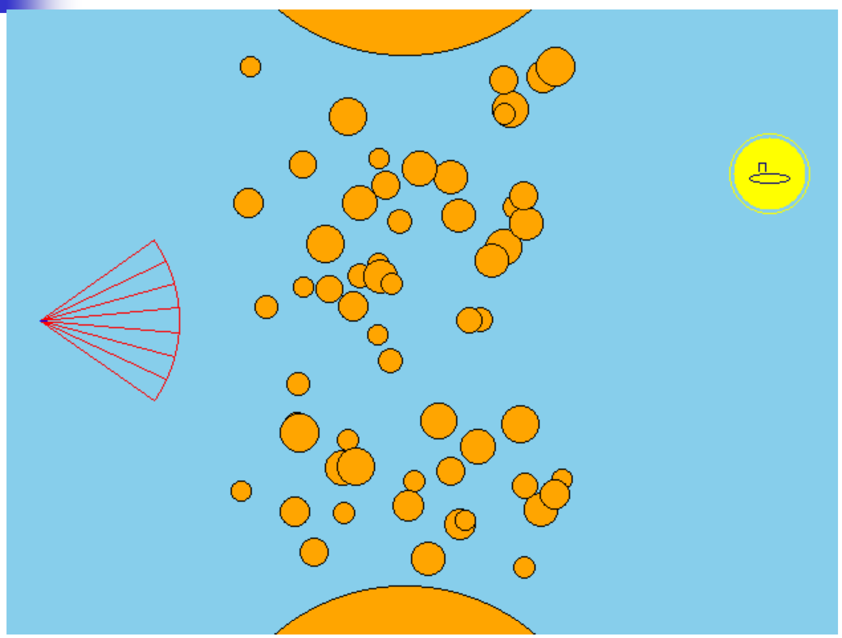

Subject

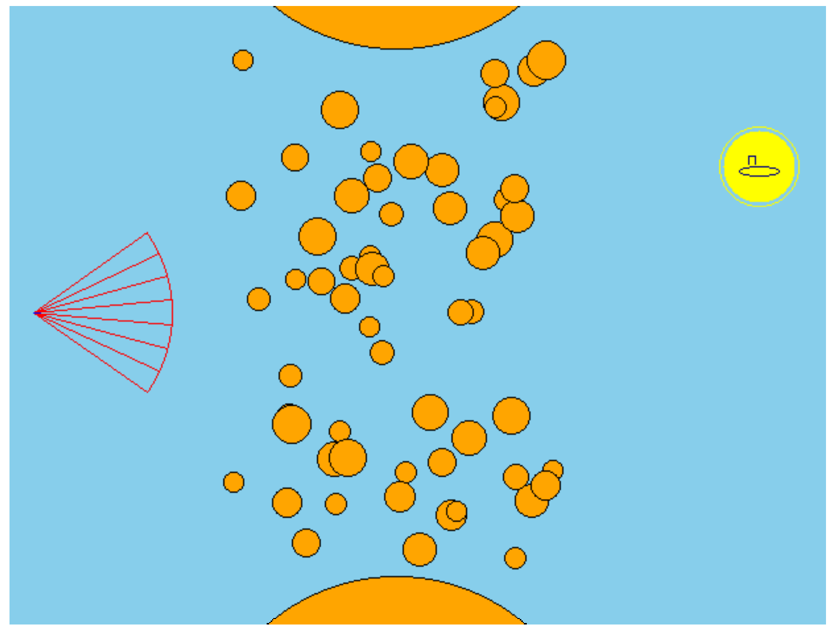

Model

Subject Cea: Day 5: 7

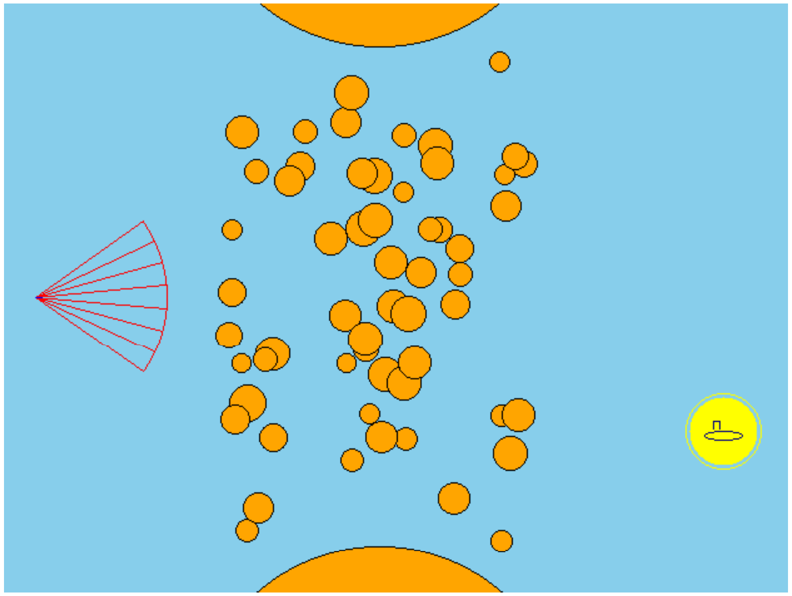

Subject

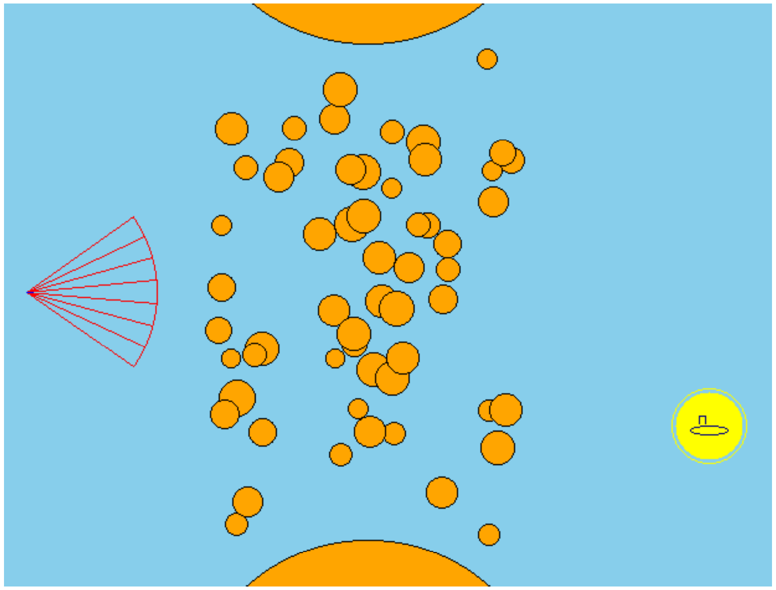

Model

Subject Cea: Day 5: 8

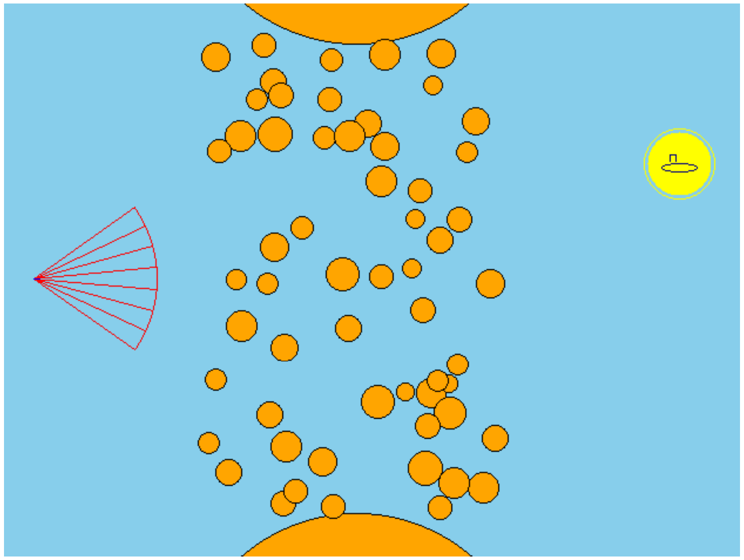

Subject

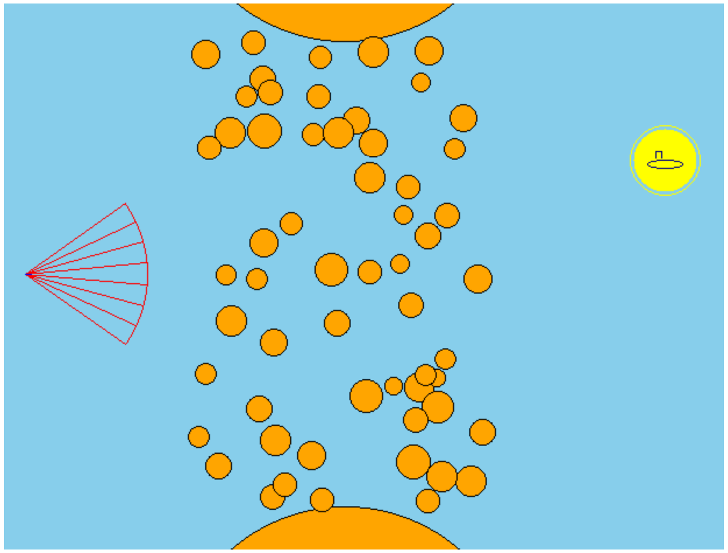

Model

Subject Cea: Day 5: 9

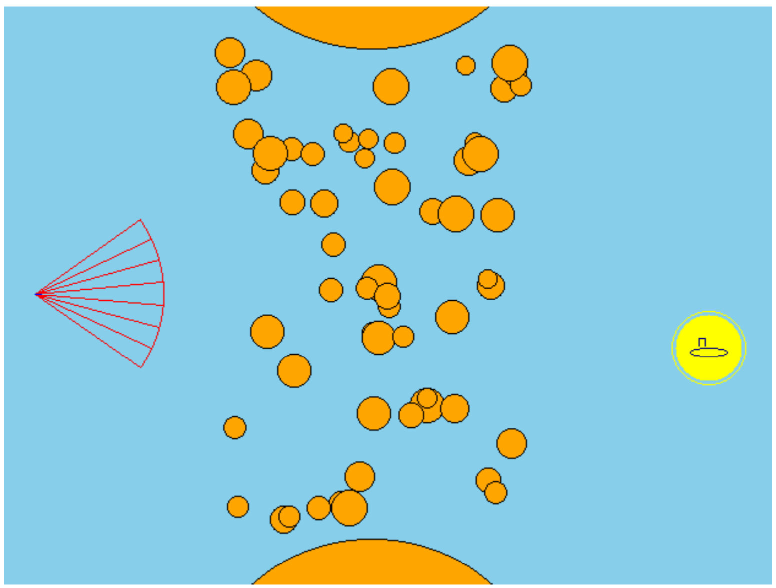

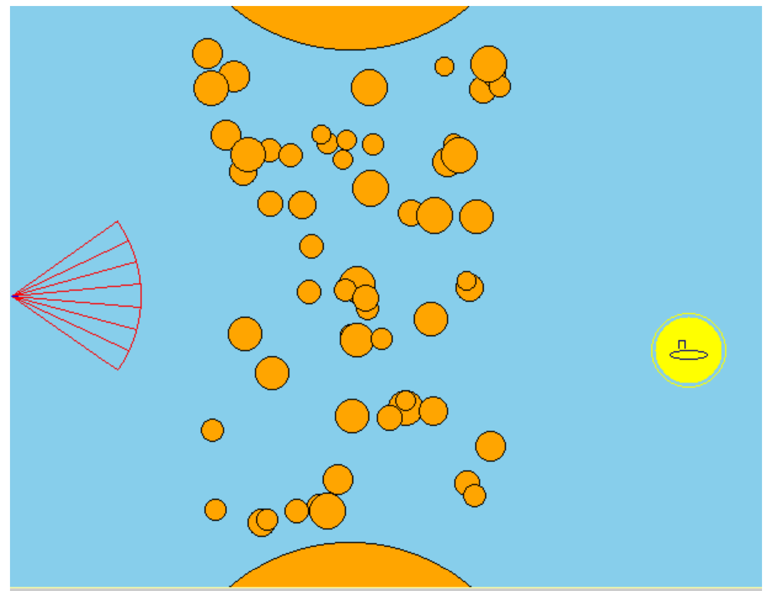

Subject

Model

Comparison with global methods

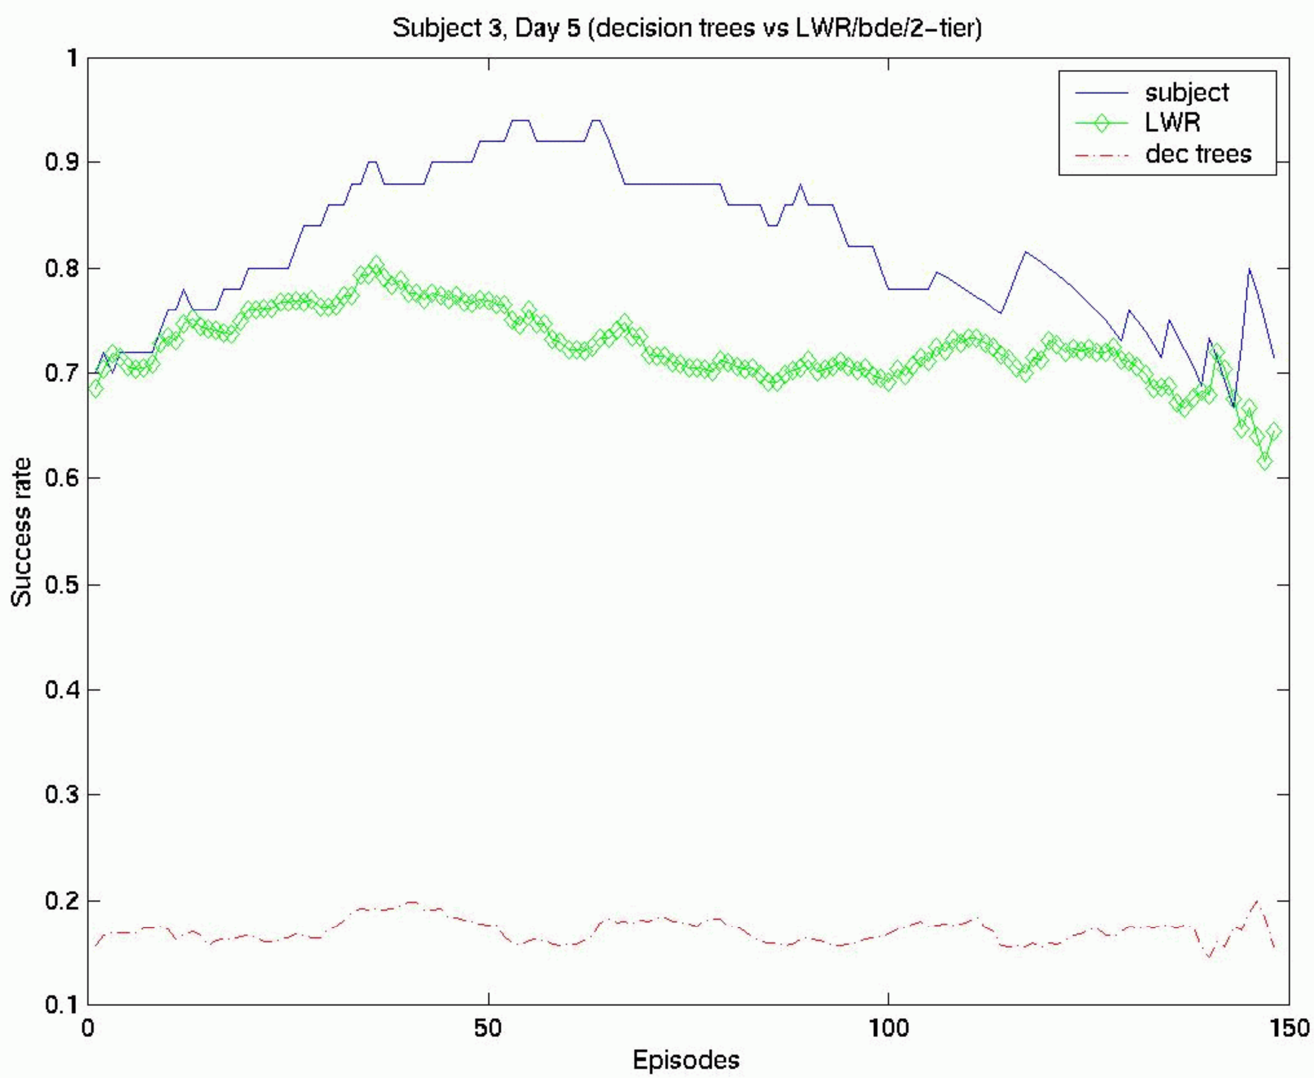

Siruguri and Subramanian, 2002

Summary

- We can model subjects on the NRL task in real-time, achieving excellent fits to their learning curves, using the available visual-motor data stream.
- One of the first in cognitive science to directly use objective visual-motor performance data to derive evolution of strategy on a complex task.

Where's the science?

Science Magazine

Lessons

- Learn simple models from objective, low-level data!
- Non-stationarity is commonplace, need to design algorithms robust with respect to it.
- Fast new algorithms for detecting change-points and building predictive stochastic models for massive, noisy, non-stationary, vector time series data.

Neural correlates

- Are there neural correlates to strategy shifts observed in the visual-motor data?

Task Questions

- Can we adapt training protocols in the NRL task by identifying whether subjects are struggling with strategy formulation or visual-motor control or both?
- Can we use analysis of EEG data gathered during learning as well as visual-motor performance data to correlate 'brain events' with 'visual-motor performance events'? Can this correlation separate subjects with different learning difficulties?

The (new) NRL Navigation Task

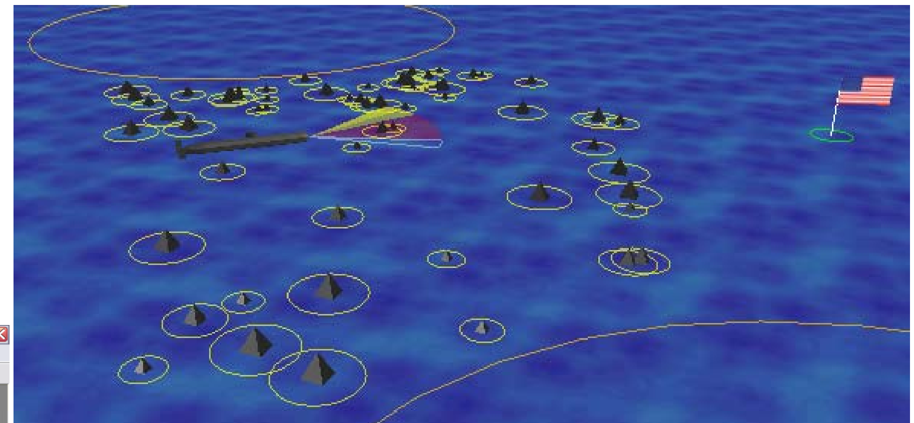

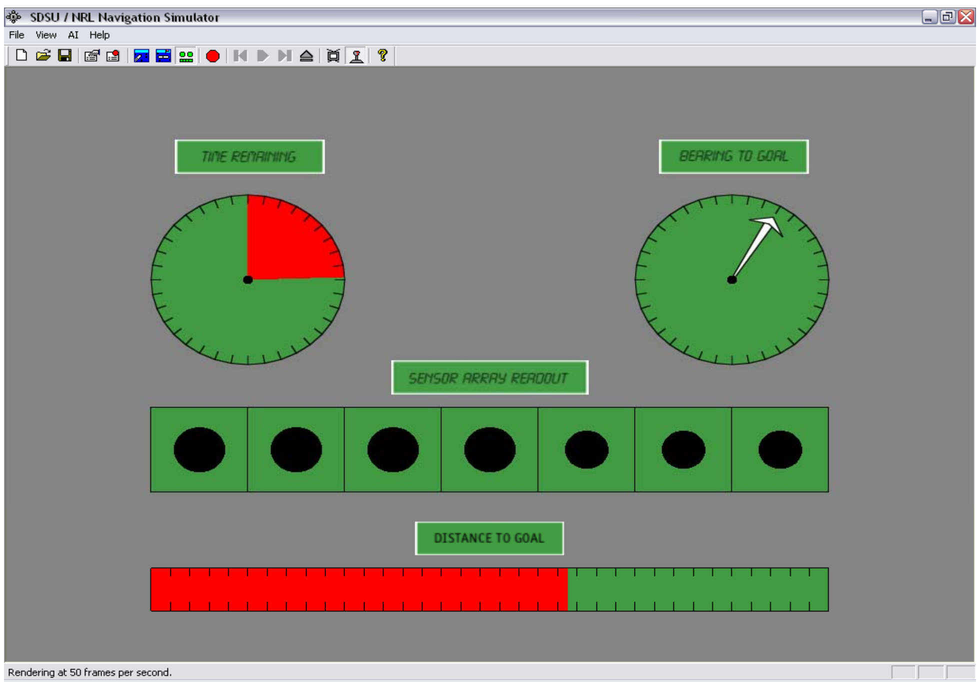

Gathering performance data

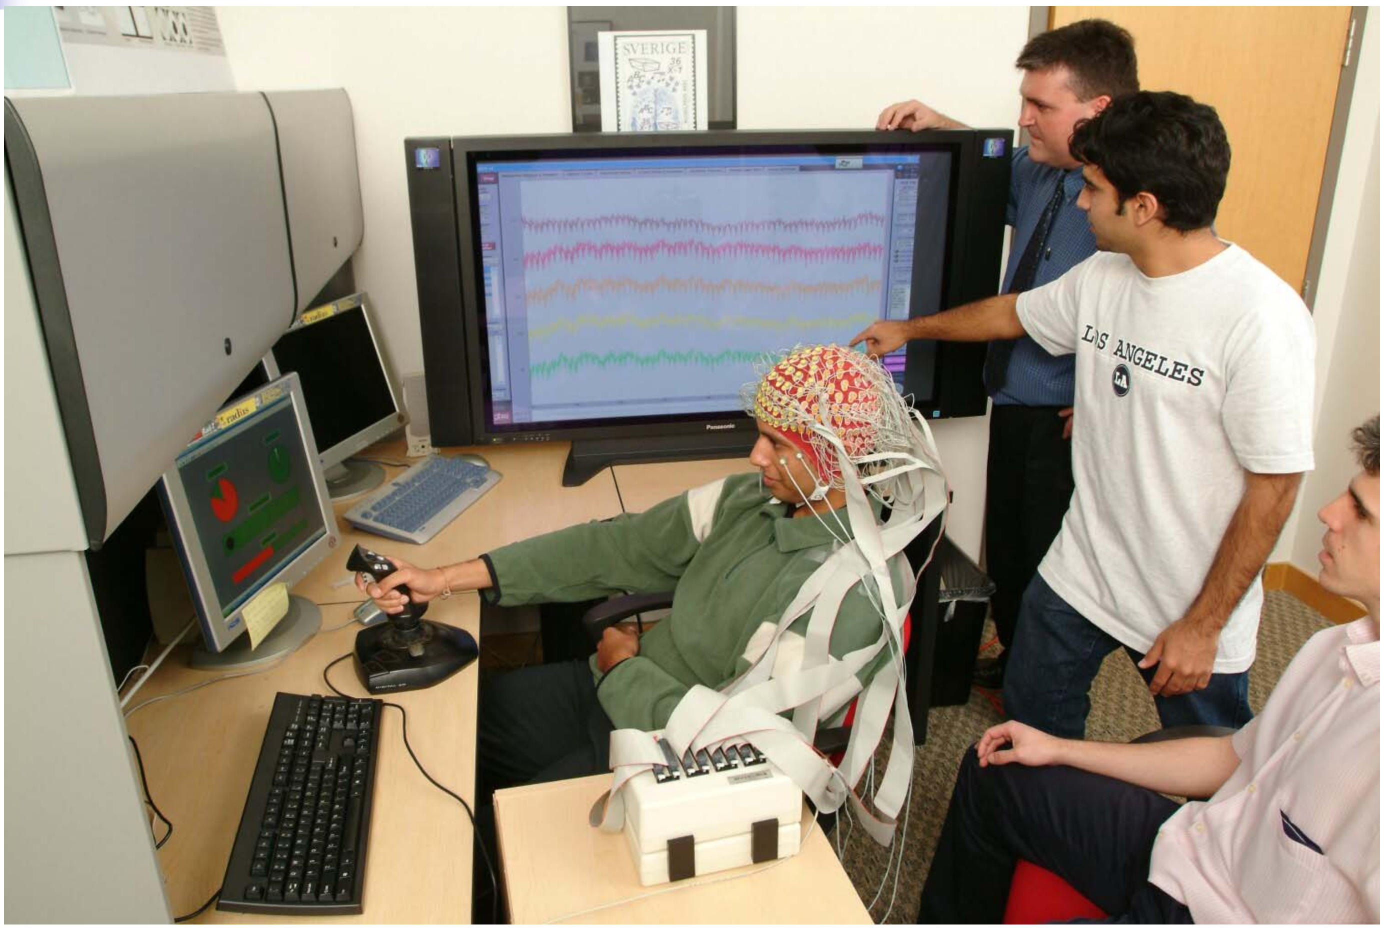

256 channel EEG recording

Fusing EEG and visualmotor data

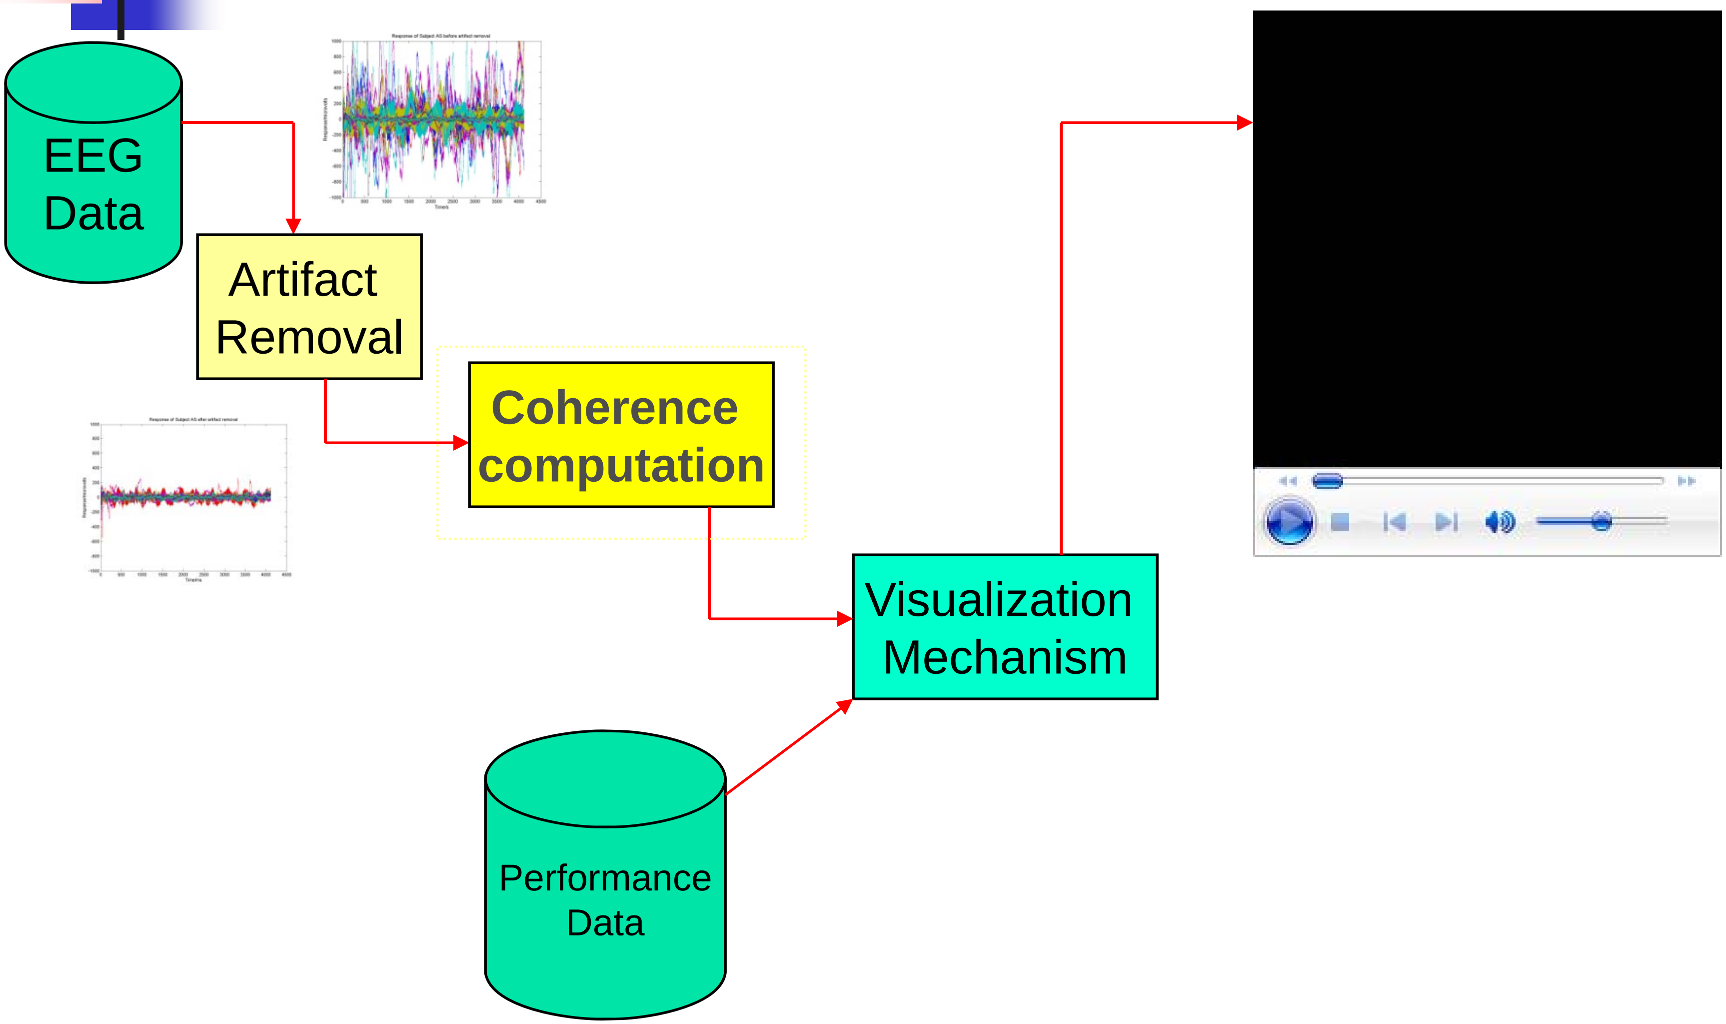

Measuring functional connectivity in the brain

- *Coherence* provides the means to measure synchronous activity between two brain areas
- A function that calculates the normalized cross-power spectrum, a measure of similarity of signal in the frequency domain

$$C_{xy}(f) = \frac{|S_{xy}(f)|^2}{[S_{xx}(f)S_{yy}(f)]}$$

Topological coherence map

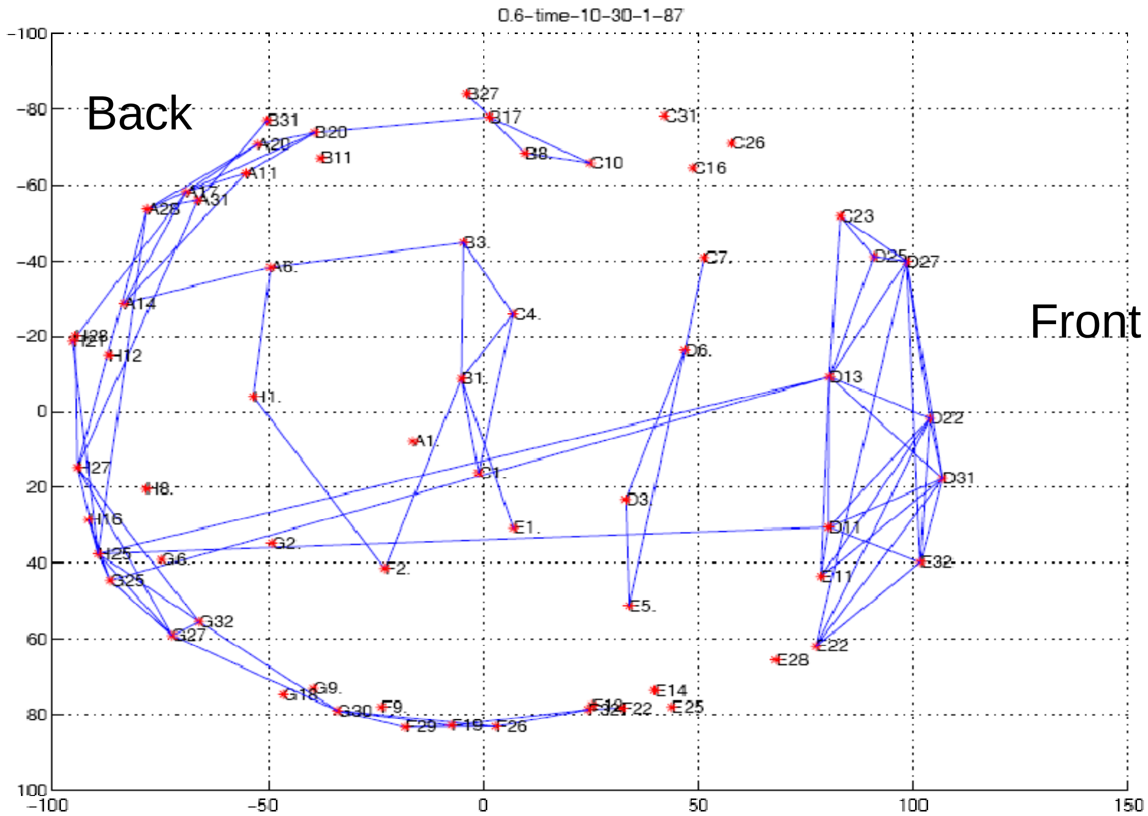

Frequency bands

- Coherence map of connections in each band
  - Δ (0-5 Hz)
  - θ (5-9 Hz)
  - α (9-14 Hz)
  - β (14-30 Hz)
  - γ (40-52 Hz)

Subject moh progression chart

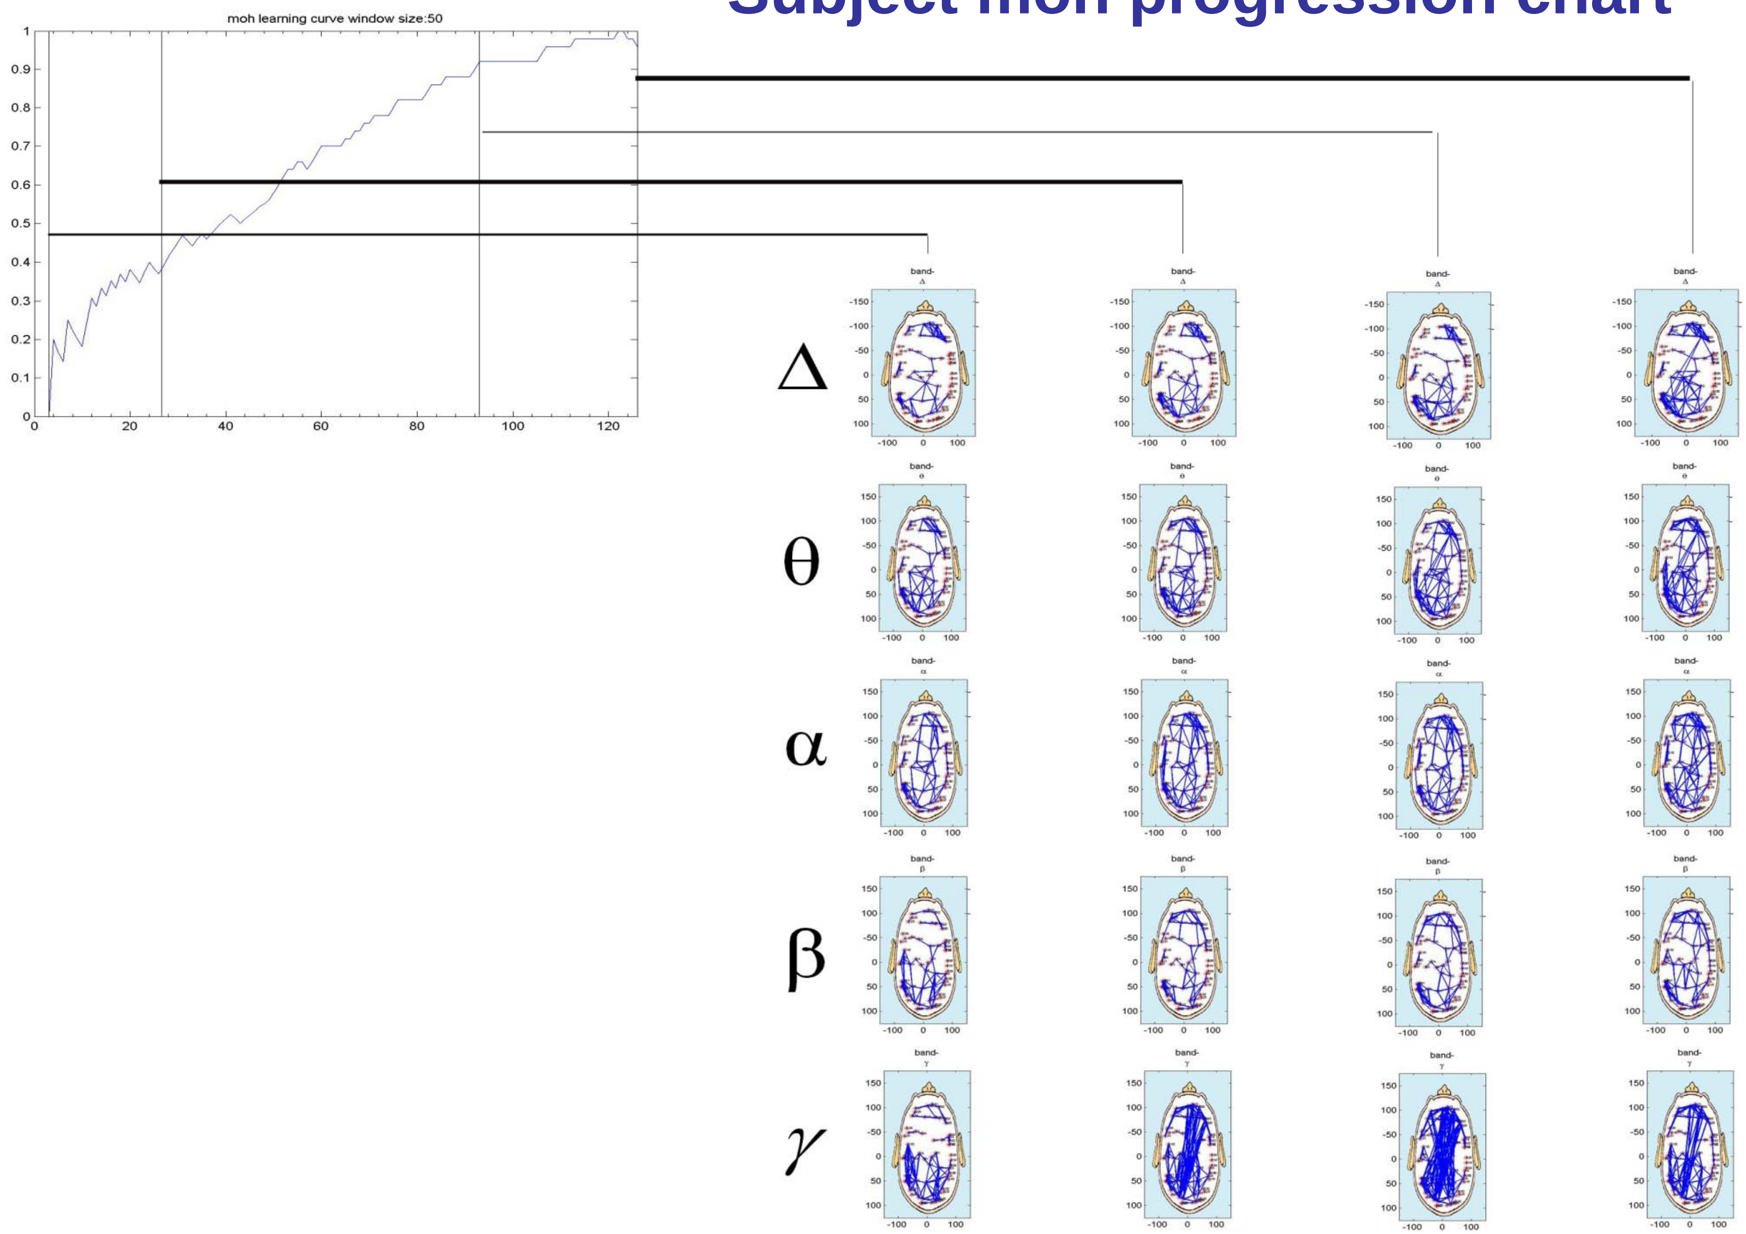

Results (subject moh)

Subject bil progression chart

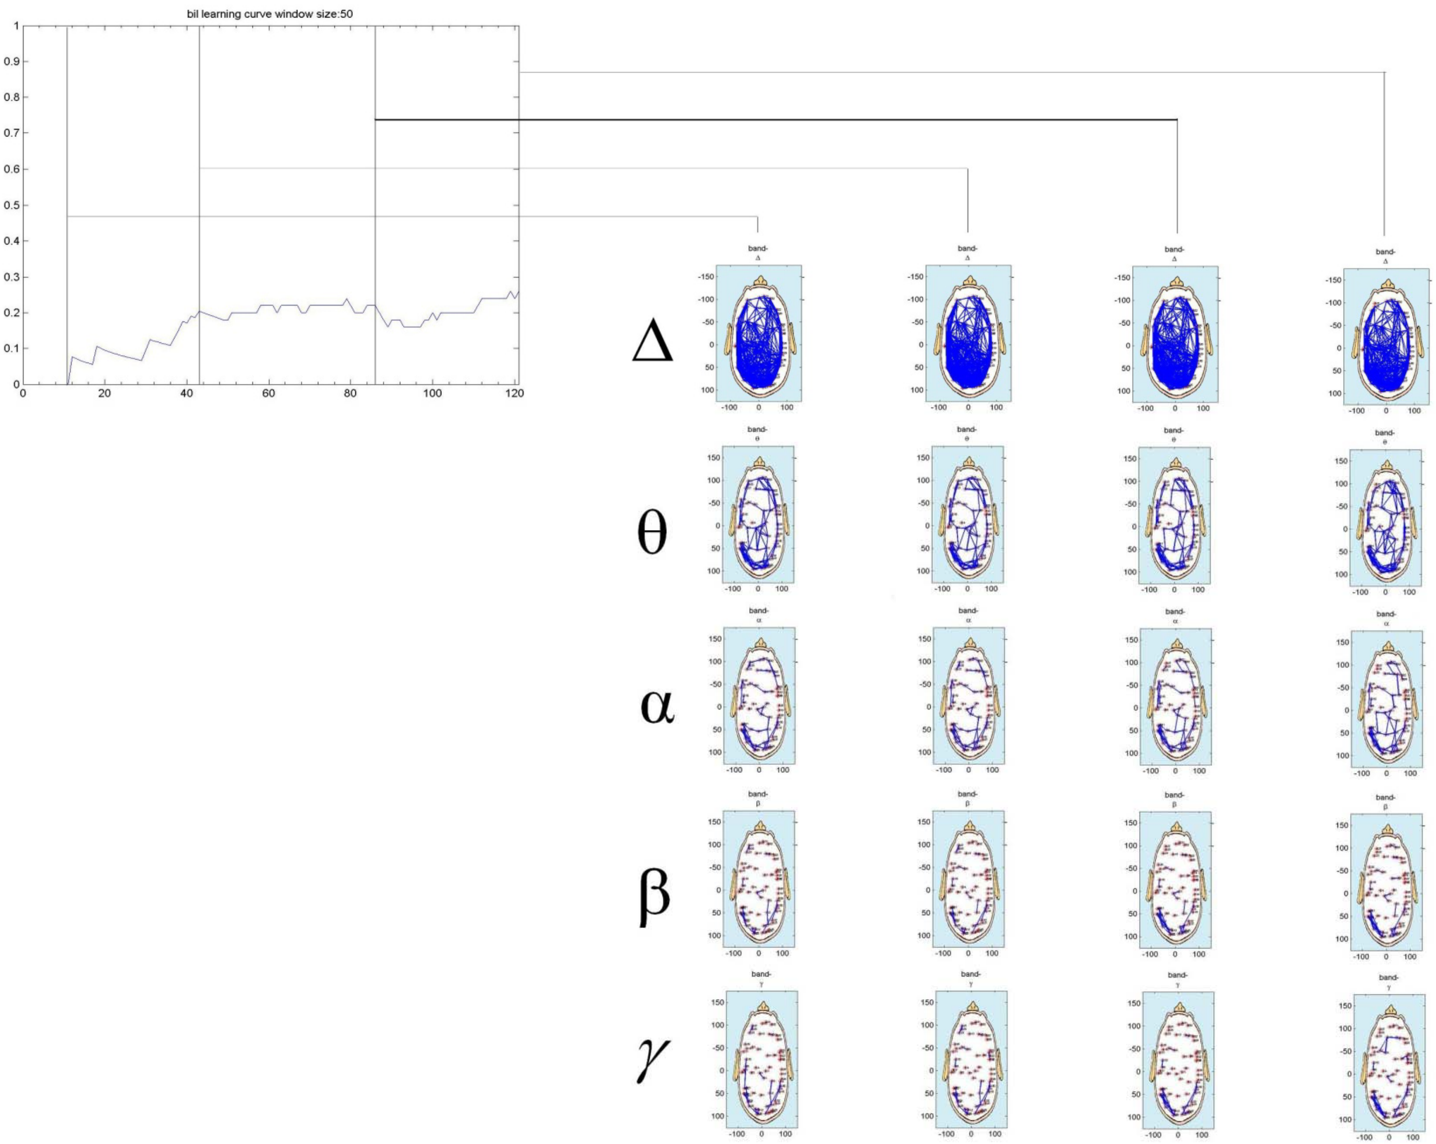

Results (subject bil)

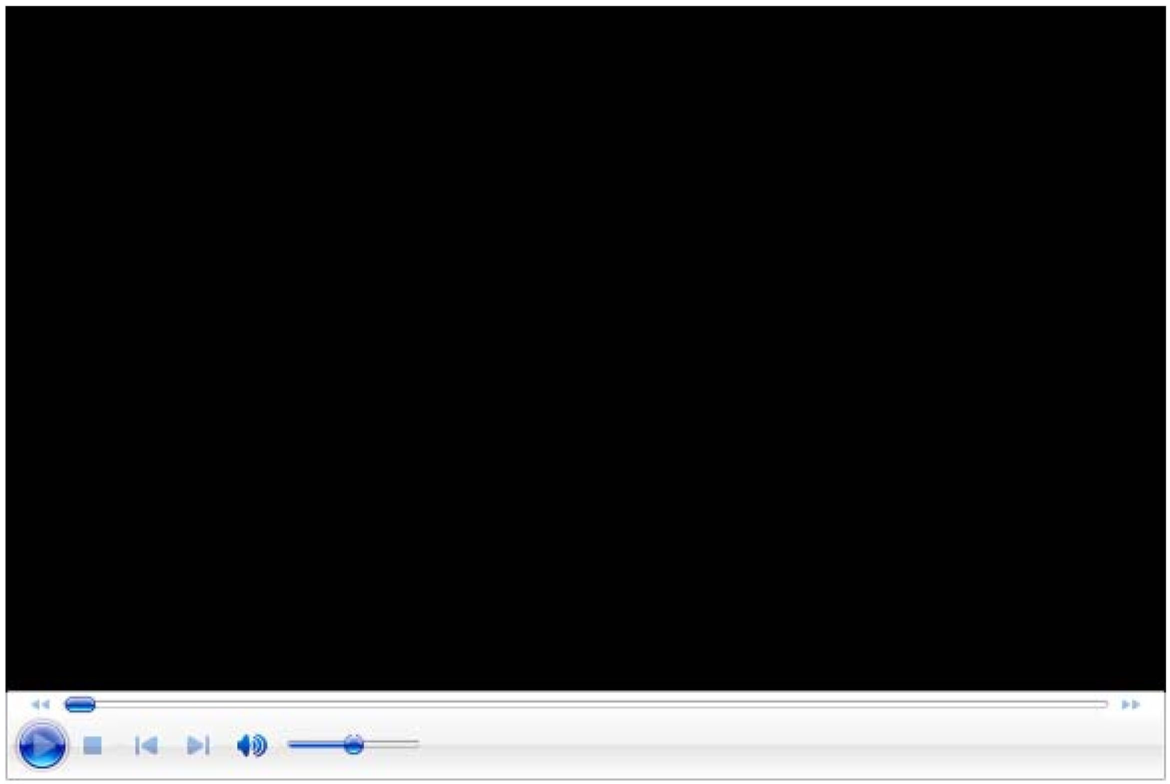

Baluch, Zouridakis, Stevenson and Subramanian, 2005, 2006

Subject G

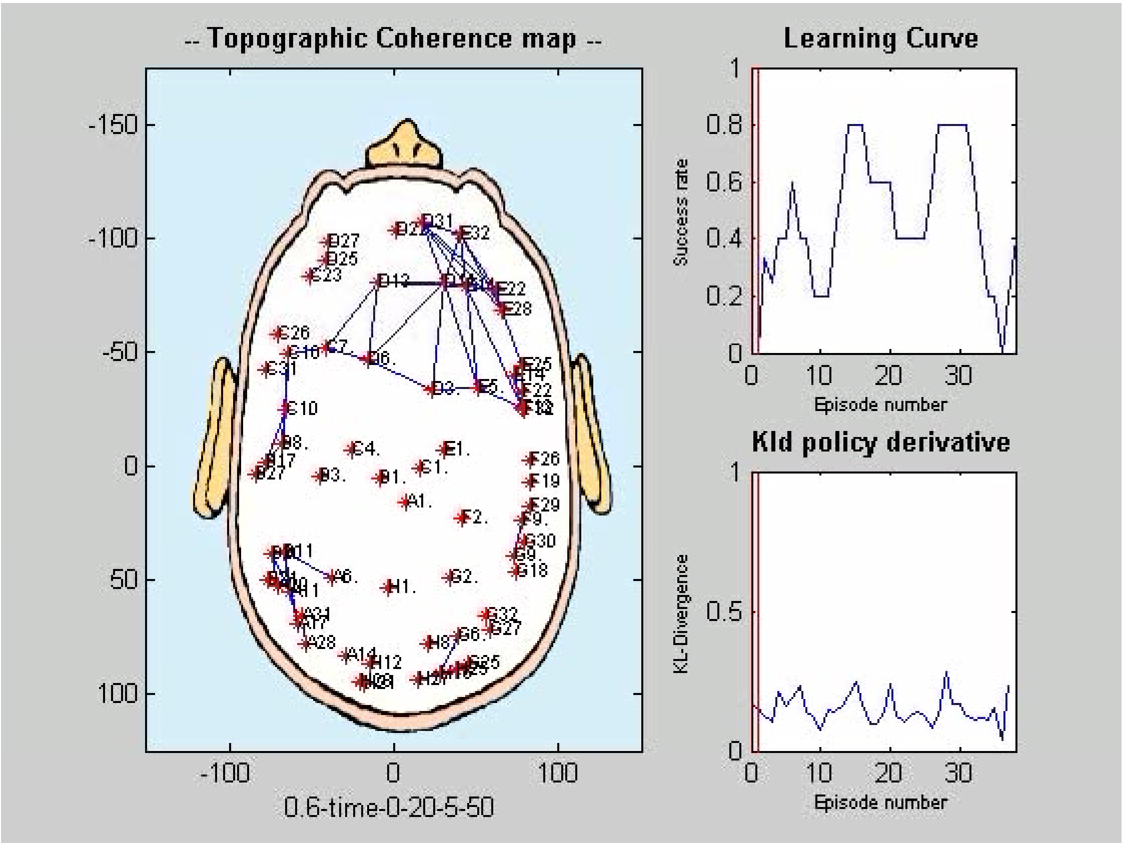

Subject is a near-expert performer

Subject is in skill refinement phase

Baluch, Zouridakis, Stevenson and Subramanian, 2005, 2006

Subject V

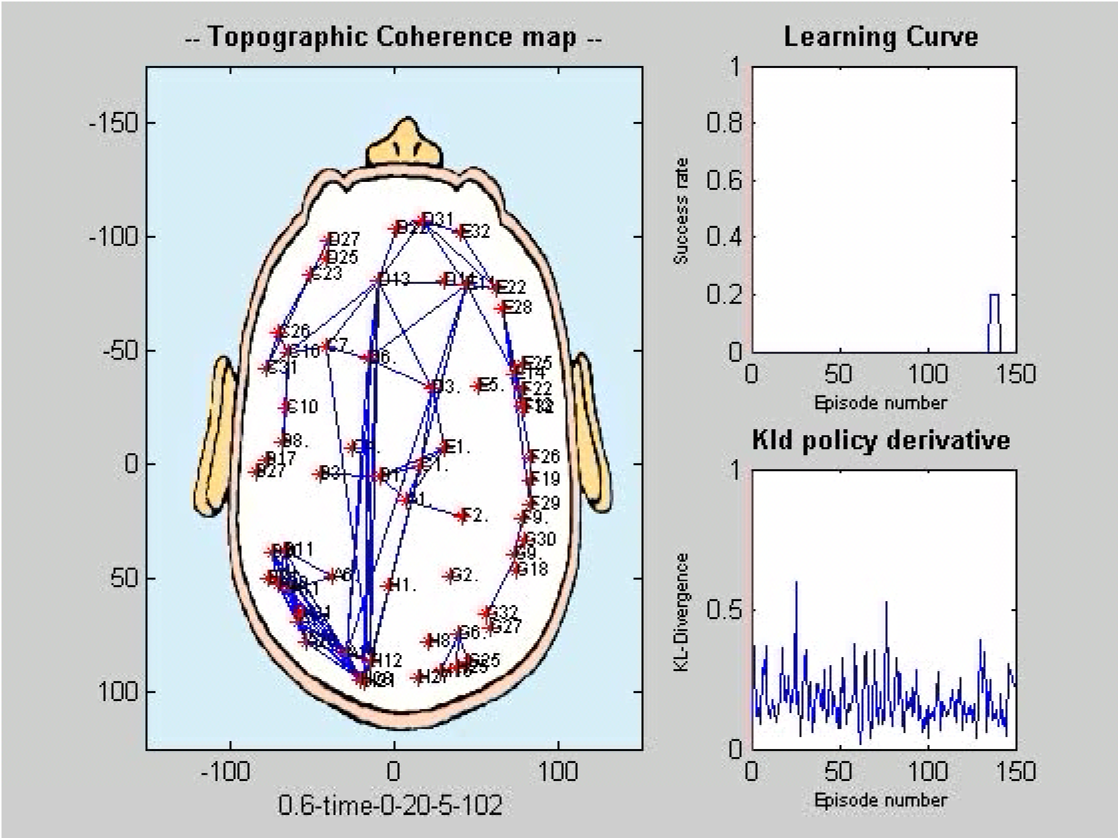

Subject never learned a good strategy

It wasn't for lack of trying..

Baluch, Zouridakis, Stevenson and Subramanian, 2005, 2006

Subject B

Results

- There are distinct EEG coherence map signatures associated with different learning difficulties
  - Lack of strategy
  - Shifting between too many strategies
- Subjects in our study who showed a move from a low level of performance to a high level of performance *show front to back synchrony in the gamma range or long range gamma synchrony (LRGS).* [Baluch,Zouridakis,Stevenson,Subramanian 2007]
- We are conducting experiments on more subjects to confirm these findings. (14 subjects so far, and more are being collected right now.)

What else is this good for?

- Using EEG readouts to analyze the effectiveness of video games for relieving pre-operative stress in children (A. Patel, UMDNJ).
- Using EEG to read emotional state of players in immersive video games (M. Zyda, USC).
- Analyzing human performance on any visualmotor task with significant strategic component.

Details, please……

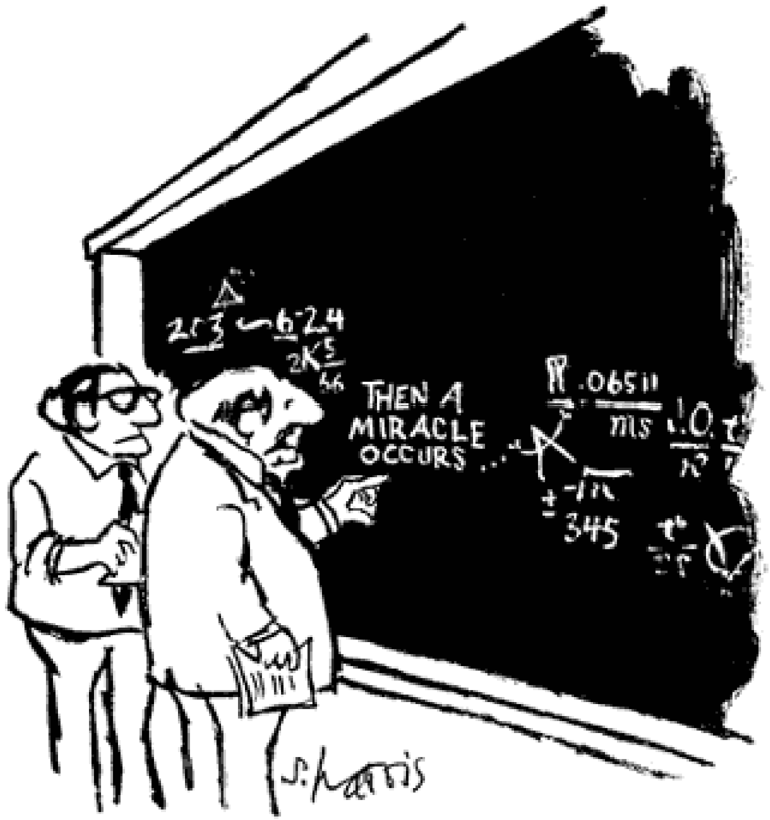

"I think you should be more explicit here in step two."

Publications

- Long-Range Gamma-Band Synchronization During Learning of a Complex Visuomotor Task, George Zouridakis, Farhan Baluch, Javier Diaz, Devika Subramanian, Ian Stevenson, IEEE EMBS 2007.
- Human Learning and the Neural Correlates of Strategy Formulation, F. Baluch, D. Subramanian and G. Zouridakis, 23rd Annual Conference on Biomedical Engineering Research, 2006.
- Understanding Human Learning on Complex Tasks by Functional Brain Imaging, D. Subramanian, R. Bandyopadhyay and G. Zouridakis, 20th Annual Conference on Biomedical Engineering Research, 2003.
- Tracking the evolution of learning on a visualmotor task Devika Subramanian and Sameer Siruguri, Technical report TR02-401, Department of Computer Science, Rice University, August 2002.
- Tracking the evolution of learning on a visualmotor task Sameer Siruguri, Master's thesis under the supervision of Devika Subramanian, May 2001.
- State Space Discretization and Reinforcement Learning, S. Griffin and D. Subramanian, Technical report, Department of Computer Science, Rice University, June 2000.
- Inducing hybrid models of learning from visualmotor data , *Proceedings of the 22nd Annual Conference of the Cognitive Science Society*, Philadelphia, PA, 2000.
- Modeling individual differences on the NRL Navigation task, *Proceedings of the 20th Annual Conference of the Cognitive Science Society*, Madison, WI, 1998 (with D. Gordon).
- *A cognitive model of learning to navigate*, *Proceedings of the 19th Annual Conference of the Cognitive Science Society*, Stanford, CA, 1997 (with D. Gordon).
- Cognitive modeling of action selection learning, *Proceedings of the 18th Annual Conference of the Cognitive Science Society*, San Diego, 1996 (with D. Gordon)

Roadmap of talk

- Four case-studies
  - Unknown system, changing dynamics
    - Tracking human learning on a complex visual-motor task.
    - Predicting the evolution of international conflict.

Adaptive Systems

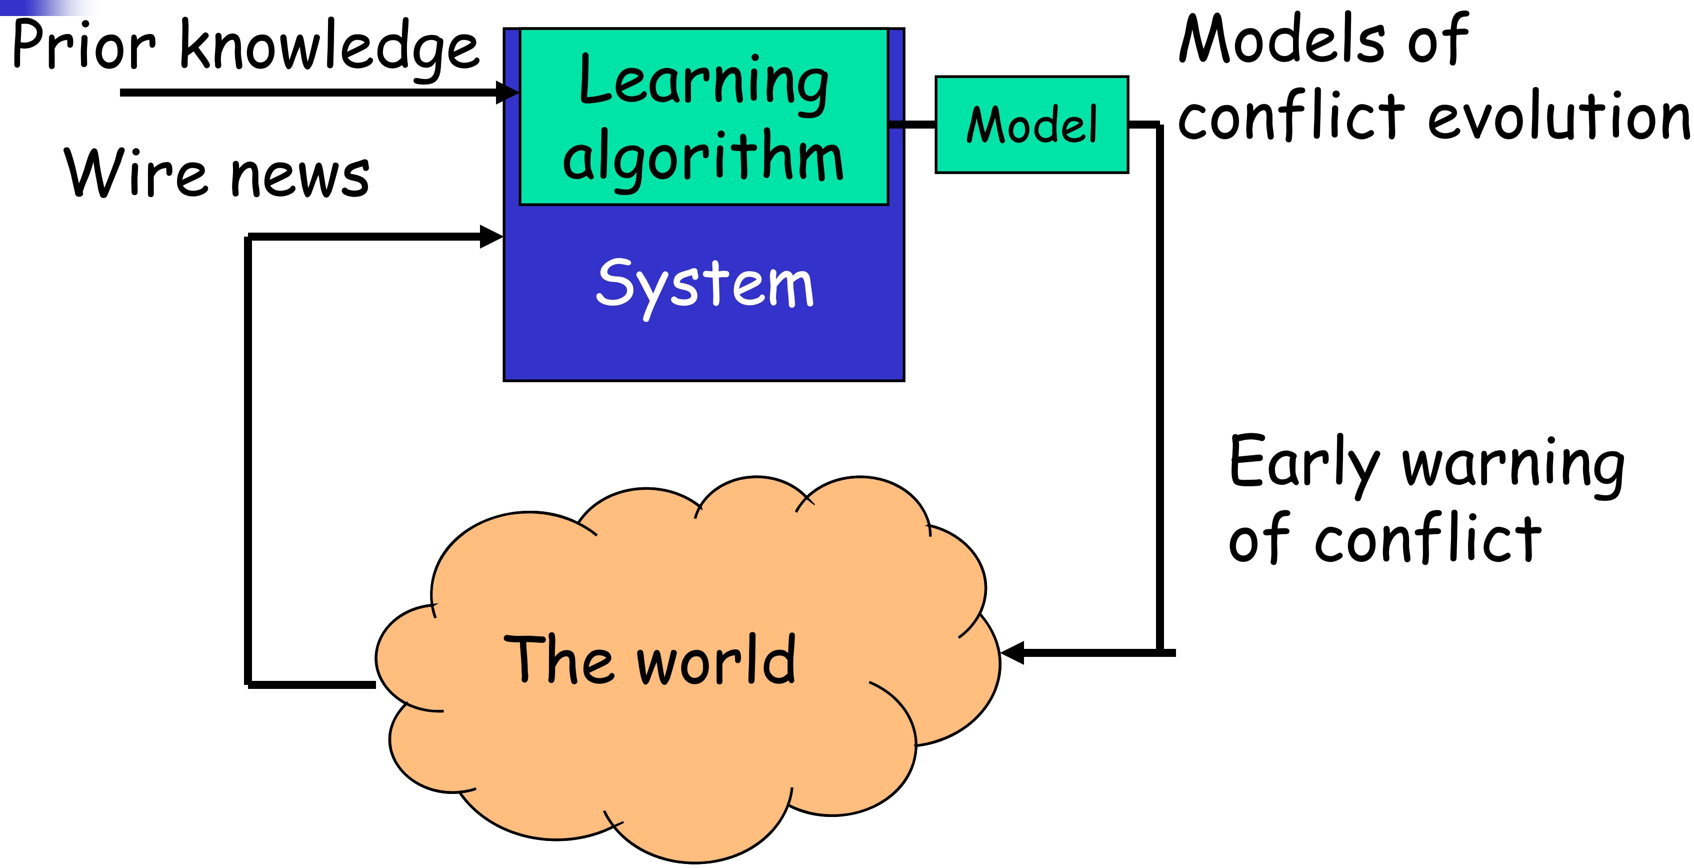

Task Question

- Is it possible to monitor news media from regions all over the world over extended periods of time, extracting low-level events from them, and piece them together to automatically track and predict conflict in all the regions of the world?

The Ares project

http://ares.cs.rice.edu

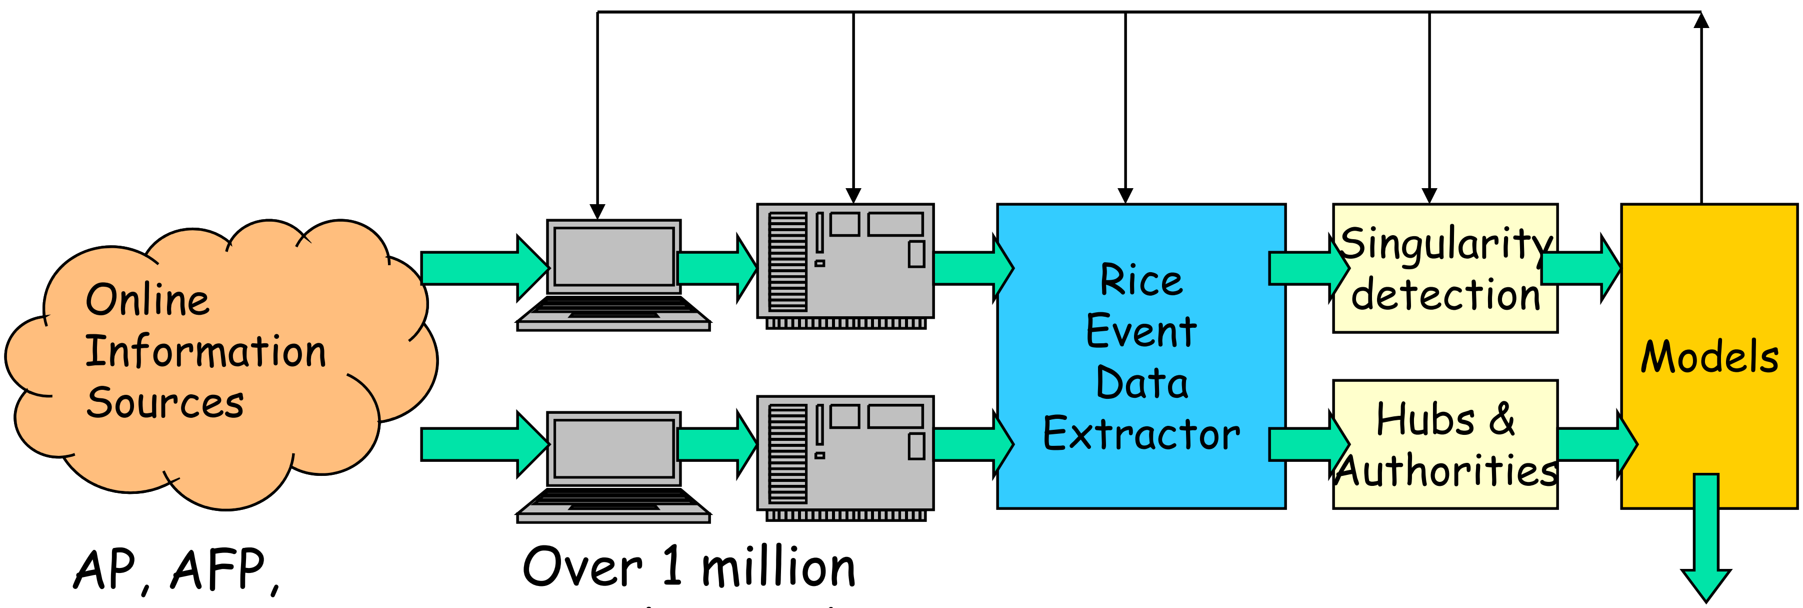

AP, AFP, BBC, Reuters, …

Over 1 million articles on the Middle East from 1979 to 2005 (filtered automatically)

Analysis of wire stories

"President Bill Clinton said on Monday the United States sought no confrontation with Iraqi President Saddam Hussein but declined to say whether that meant he would forego immediate air strikes on Iraq."

Relevance filter

Date	Actor	Target	Weis Code	Wies event	Goldstein scale
790415	ARB	ISR	223	(MIL ENGAGEMENT)	-10
790415	EGY	AFD	194	(HALT NEGOTIATION)	-3.8
790415	PALPL	ISR	223	(MIL ENGAGEMENT)	-10
790415	UNK	ISR	223	(MIL ENGAGEMENT)	-10
790415	ISR	EGY	31	(MEET)	1
790415	EGY	ISR	31	(MEET)	1
790415	ISRMIL	PAL	223	(MIL ENGAGEMENT)	-10
790415	PALPL	JOR	223	(MIL ENGAGEMENT)	-10
790415	EGY	AFD	193	(CUT AID)	-5.6
790415	IRQ	EGY	31	(MEET)	1
790415	EGY	IRQ	31	(MEET)	1
790415	ARB	CHR	223	(MIL ENGAGEMENT)	-10
790415	JOR	AUS	32	(VISIT)	1.9
790415	UGA	CHR	32	(VISIT)	1.9
790415	ISRGOV	ISRSET	54	(ASSURE)	2.8

Singularity detection on aggregated events data

Hubs and authorities analysis of events data

Embedded learner design

- Representation
  - Identify relevant stories, extract event data from them, build time series models and graph-theoretic models.
- Learning
  - Identifying regime shifts in events data, tracking evolution of militarized interstate disputes (MIDs) by hubs/authorities analysis of events data
- Decision-making
  - Issuing early warnings of outbreak of MIDs

Identifying relevant stories

- Only about 20% of stories contain events that are to be extracted.
  - The rest are interpretations, (e.g., op-eds), or are events not about conflict (e.g., sports)
- We have trained Naïve Bayes (precision 86% and recall 81%), SVM classifiers (precision 92% and recall 89%) & Okapi classifiers (precision 93% and recall 87%) using a labeled set of 180,000 stories from Reuters.
- Surprisingly difficult problem!
  - Lack of large labeled data sets;
  - Poor transfer to other sources (AP/BBC)
  - The category of "event containing stories" is not well-separated from others, and changes with time

Lee, Tran, Singer, Subramanian, 2006

Okapi classifier

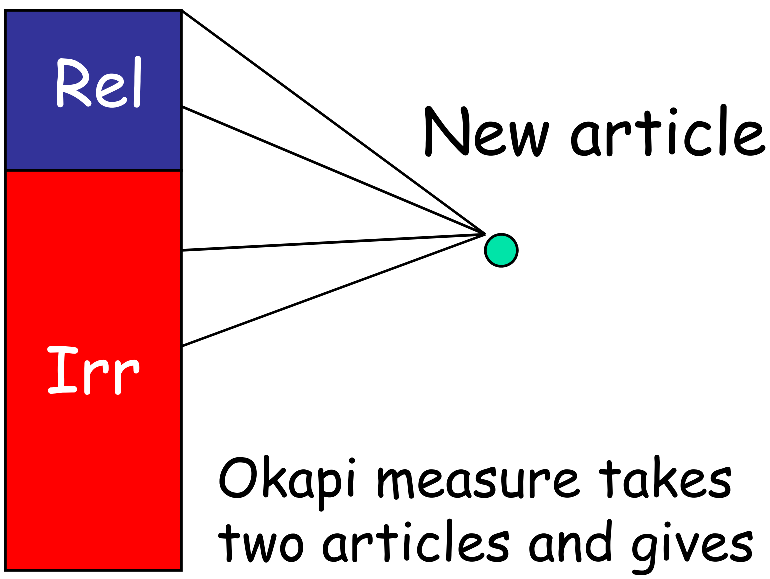

Okapi measure takes two articles and gives the similarity between them.

- Reuters data set: relevant categories are GVIO, GDIP, G13; irrelevant categories: 1POL, 2ECO, 3SPO, ECAT, G12, G131, GDEF, GPOL

Decision rule: sum of top N Okapi scores in Rel set > sum of top N Okapi scores in Irr set then classify as rel; else irr

Event extraction

"President Bill Clinton said on Monday the United States sought no confrontation with Iraqi President Saddam Hussein but declined to say whether that meant he would forego immediate air strikes on Iraq."

Fragment	Event data		
President Bill Clinton said on Monday the United States sought no confrontation with Iraqi President Saddam Hussein	USA	Comment	USA
the United States sought no confrontation with Iraqi President Saddam Hussein	USA	Deny	IRQ
President Bill Clinton declined to say whether that meant President Bill Clinton would forego immediate air strikes on Iraq	USA	Comment	USA
President Bill Clinton would forego immediate air strikes on Iraq	not part of event phrase (did not actually happen)		

Parse sentence

1. *President Bill Clinton said on Monday the United States sought no confrontation with Iraqi President Saddam Hussein but declined to say whether that meant he would forego immediate air strikes on Iraq.*

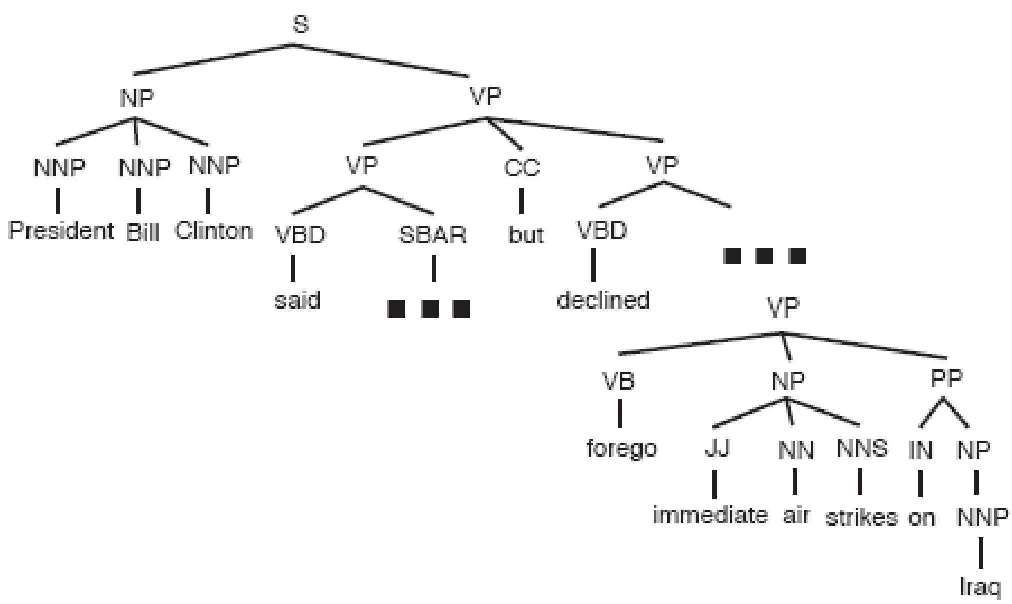

Klein and Manning parser

Pronoun de-referencing

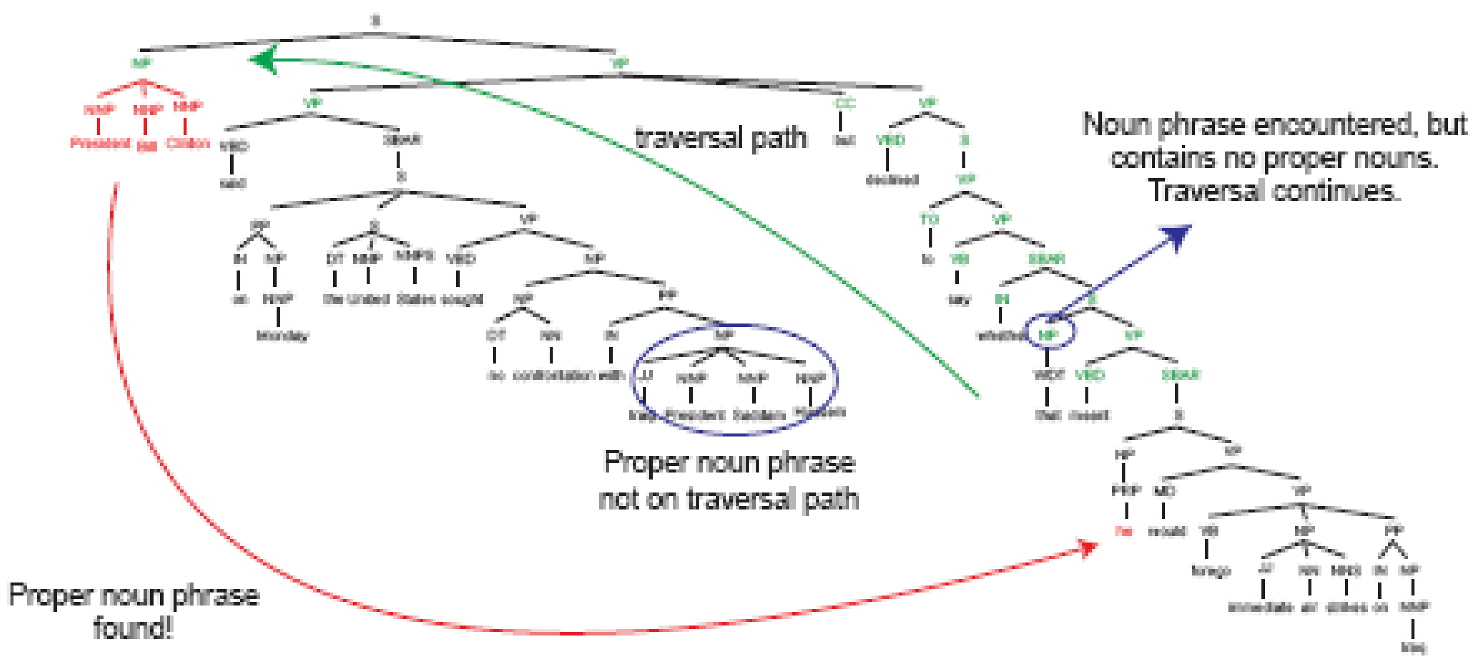

Sentence fragmentation

"President Bill Clinton said on Monday the United States sought no confrontation with Iraqi President Saddam Hussein but declined to say whether that meant he would forego immediate air strikes on Iraq."

Correlative conjunctions

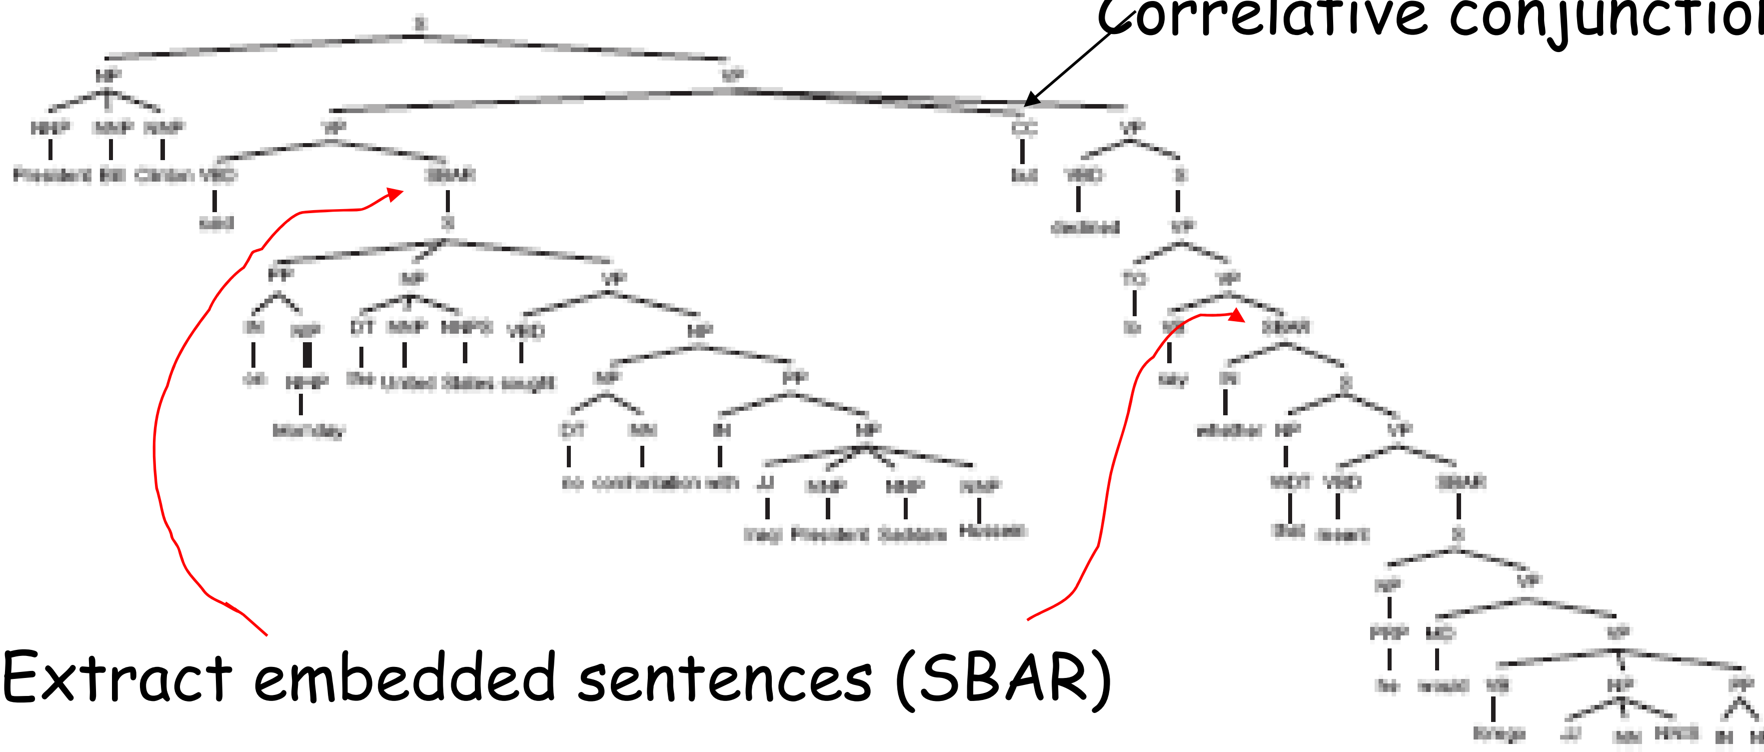

Extract embedded sentences (SBAR)

Conditional random fields

We extract who (actor) did what (event) to whom (target)

Actor & target labels	Event category labels
-countries, capitals, nationalities	-specific event keywords
-words in proper noun phrases	-words in main verb phrase
-actors occur before main verb	-specific parts of speech
-actors at higher levels of tree	-not modified by negative words
-targets occur after main verb	-part of event phrase
-targets at lower levels of tree	

Not exactly the same as NER

Results

TABARI is state of the art coder in political science

Table 1: Results for 22 (Force) category

Coder	Accuracy	Recall	Precision
TABARI	22%	7%	50%
TABARI with frag	20%	8%	83%
CRF	72%	70%	91%

Table 2: Results for 02 (Comment) category

Coder	Accuracy	Recall	Precision
TABARI	81%	31%	67%
TABARI with frag	88%	54%	93%
CRF	89%	96%	68%

200 Reuters sentences; hand-labeled with actor, target, and event codes (22 and 02).

Stepinksi, Stoll, Subramanian 2006

Events data

Date	Actor	Target	Weis Code	Wies event	Goldstein scale
790415	ARB	ISR	223	(MIL ENGAGEMENT)	-10
790415	EGY	AFD	194	(HALT NEGOTIATION)	-3.8
790415	PALPL	ISR	223	(MIL ENGAGEMENT)	-10
790415	UNK	ISR	223	(MIL ENGAGEMENT)	-10
790415	ISR	EGY	31	(MEET)	1
790415	EGY	ISR	31	(MEET)	1
790415	ISRMIL	PAL	223	(MIL ENGAGEMENT)	-10
790415	PALPL	JOR	223	(MIL ENGAGEMENT)	-10
790415	EGY	AFD	193	(CUT AID)	-5.6
790415	IRQ	EGY	31	(MEET)	1
790415	EGY	IRQ	31	(MEET)	1
790415	ARB	CHR	223	(MIL ENGAGEMENT)	-10
790415	JOR	AUS	32	(VISIT)	1.9
790415	UGA	CHR	32	(VISIT)	1.9
790415	ISRGOV	ISRSET	54	(ASSURE)	2.8

177,336 events from April 1979 to October 2003 in Levant data set (KEDS).

What can be predicted?

"IF THESE NUMBERS ARE CORRECT, THEN EVERYTHING IS GOING TO HAPPEN AT ONCE TOMORROW MORNING AT 10:35."

Singularity detection

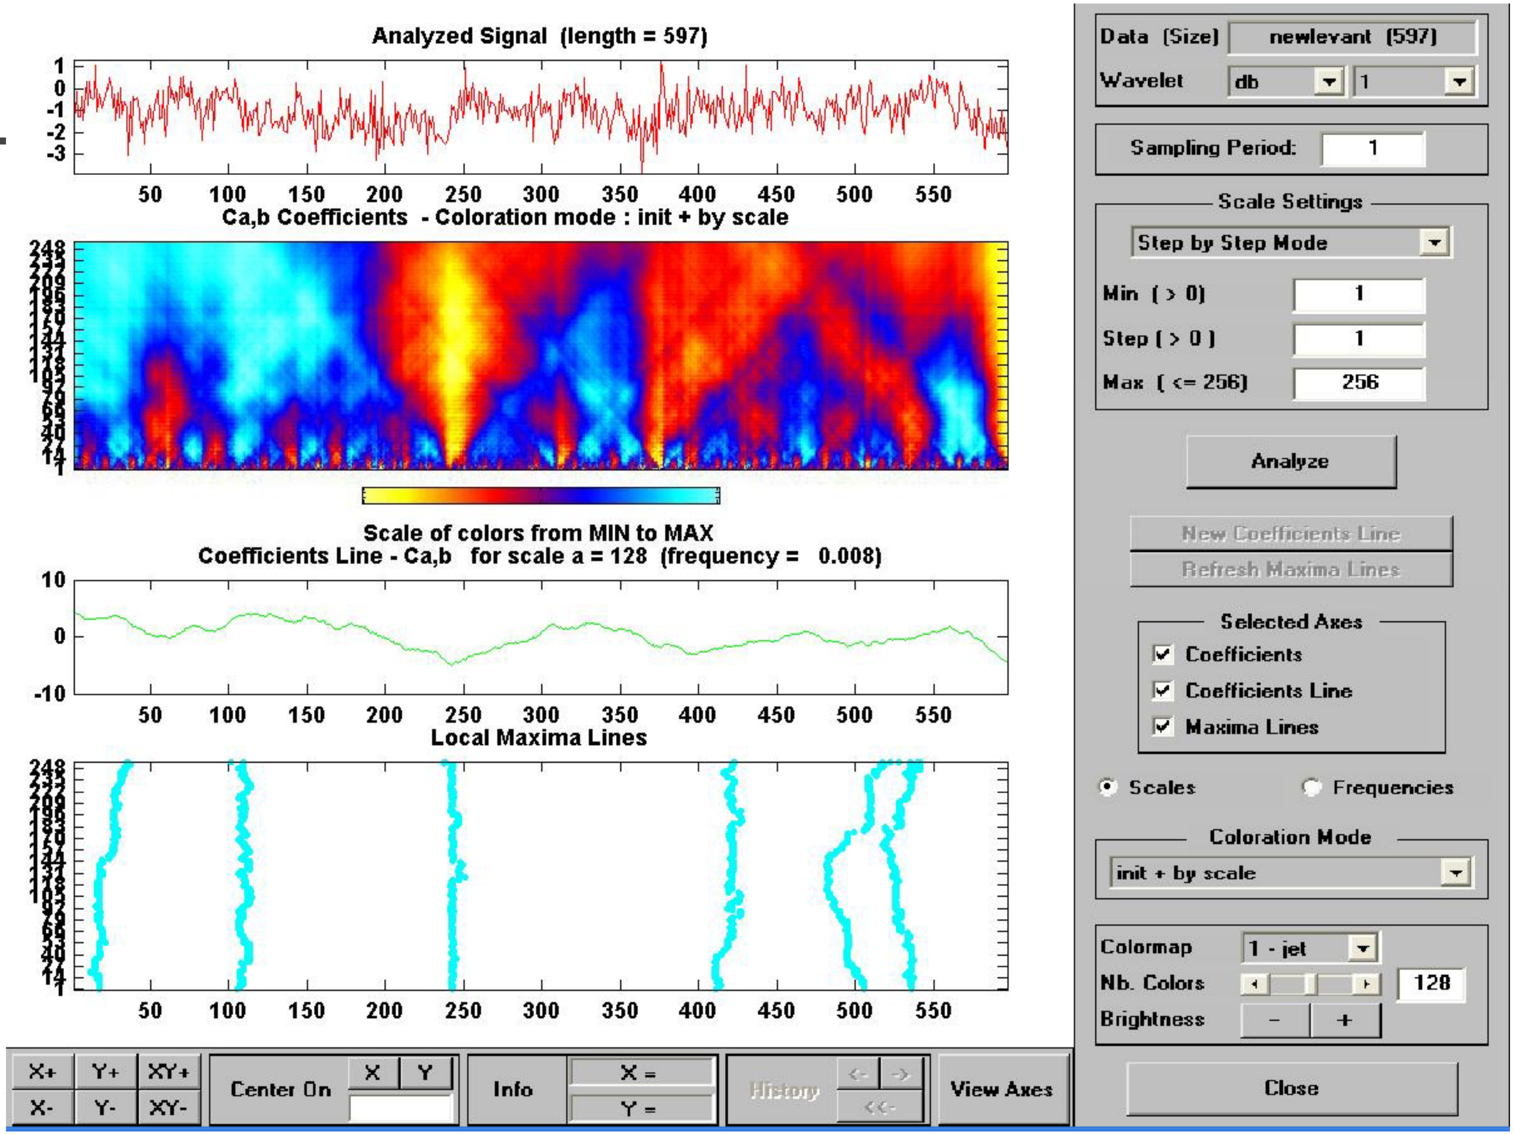

Stoll and Subramanian, 2004, 2006

Singularities = MID start/end

biweek	Date range	event
17-35	11/79 to 8/80	Start of Iran/Iraq war
105-111	4/83 to 7/83	Beirut suicide attack, end of Iran/Iraq war
244	1/91 to 2/91	Desert Storm
413-425	1/95 to 7/95	Rabin assassination/start of Intifada
483-518	10/97 to 2/99	US/Iraq confrontation via Richard Butler/arms inspectors
522-539	4/99 to 11/99	Second intifada Israel/Palestine

Interaction graphs

- Model interactions between countries in a directed graph.

Date	Actor	Target	Weis Code	Wies event	Goldstein scale
790415	ARB	ISR	223	(MIL ENGAGEMENT)	-10
790415	EGY	AFD	194	(HALT NEGOTIATION)	-3.8
790415	PALPL	ISR	223	(MIL ENGAGEMENT)	-10
790415	UNK	ISR	223	(MIL ENGAGEMENT)	-10
790415	ISR	EGY	31	(MEET)	1

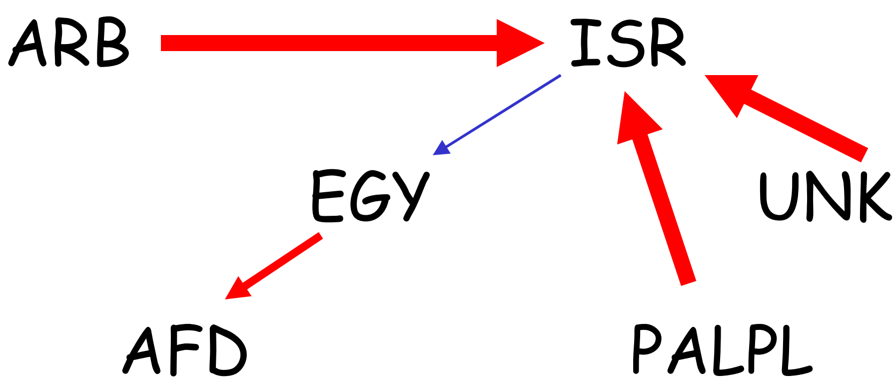

Hubs and authorities for events data

- A hub node is an important initiator of events.
- An authority node is an important target of events.
- Hypothesis:
  - Identifying hubs and authorities over a particular temporal chunk of events data tells us who the key actors and targets are.
  - Changes in the number and size of connected components in the interaction graph signal potential outbreak of conflict.

Hubs/Authorities picture of Iran Iraq war

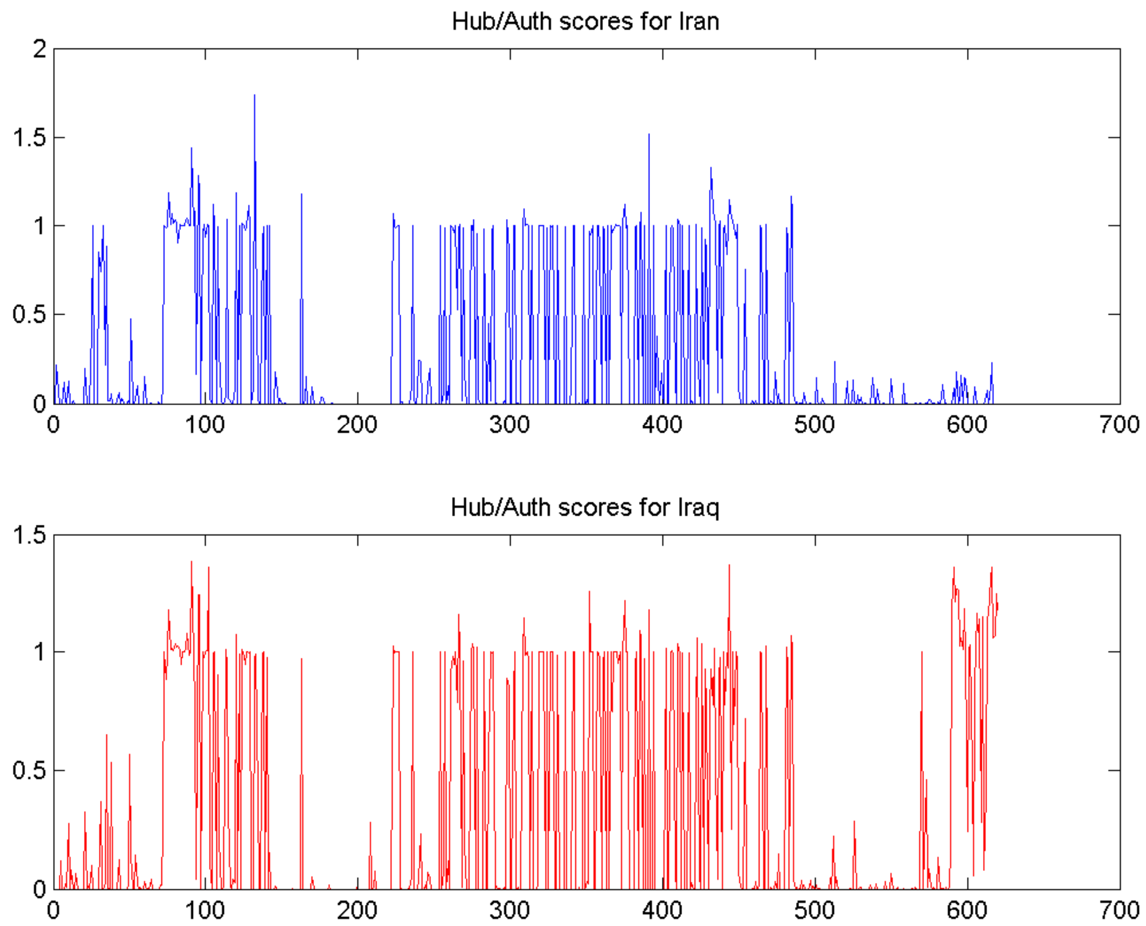

2 weeks prior to Desert Storm

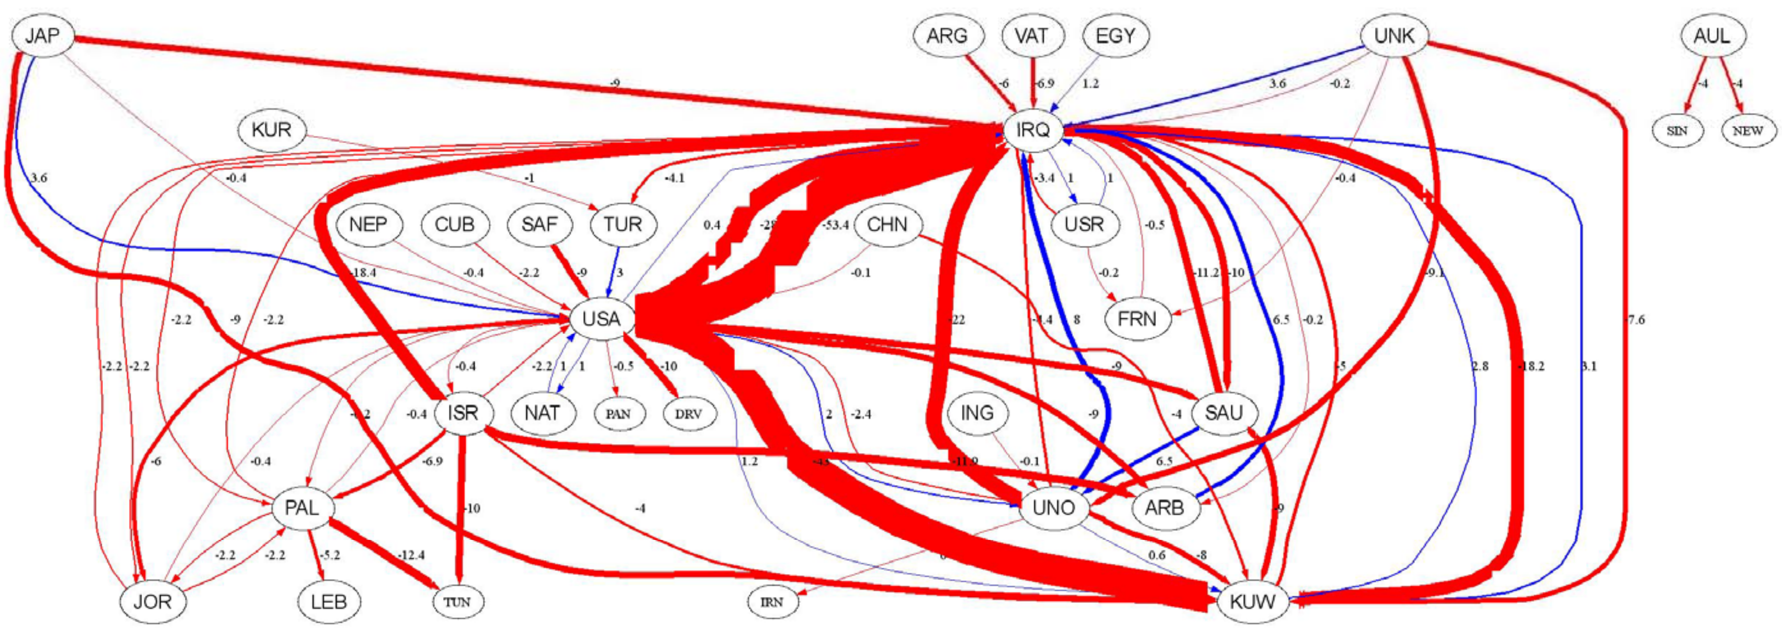

Validation using MID data

- Number of bi-weeks with MIDS in Levant data: 41 out of 589.
- Result 1: Hubs and Authorities correctly identify actors and targets in impending conflict.
- Result 2: Simple regression model on change in hubs and authorities scores, change in number of connected components, change in size of largest component 4 weeks before MID, predicts MID onset.
- Problem: false alarm rate of 16% can be reduced by adding political knowledge of conflict.

Stoll and Subramanian, 2006

"HERE'S THE GROUND RULE: DON'T TELL ME WHAT I SHOULD HAVE DONE."

Current work

- Extracting economic events along with political events to improve accuracy of prediction of both economic and political events.

Publications

- An OKAPI-based approach for article filtering, Lee, Than, Stoll, Subramanian, 2006 Rice University Technical Report.
- Hubs, authorities and networks: predicting conflict using events data, R. Stoll and D. Subramanian, International Studies Association, 2006 (invited paper).
- Events, patterns and analysis, D. Subramanian and R. Stoll, in Programming for Peace: Computer-aided methods for international conflict resolution and prevention, 2006, Springer Verlag, R. Trappl (ed).
- Four Way Street? Saudi Arabia's Behavior among the superpowers, 1966-1999, R. Stoll and D. Subramanian, James A Baker III Institute for Public Policy Series, 2004.
- Events, patterns and analysis: forecasting conflict in the 21st century, R. Stoll and D. Subramanian, Proceedings of the National Conference on Digital Government Research, 2004.
- Forecasting international conflict in the 21st century, D. Subramanian and R. Stoll, in Proc. of the Symposium on Computer-aided methods for international conflict resolution, 2002.

The research team

Embedded Adaptive Systems

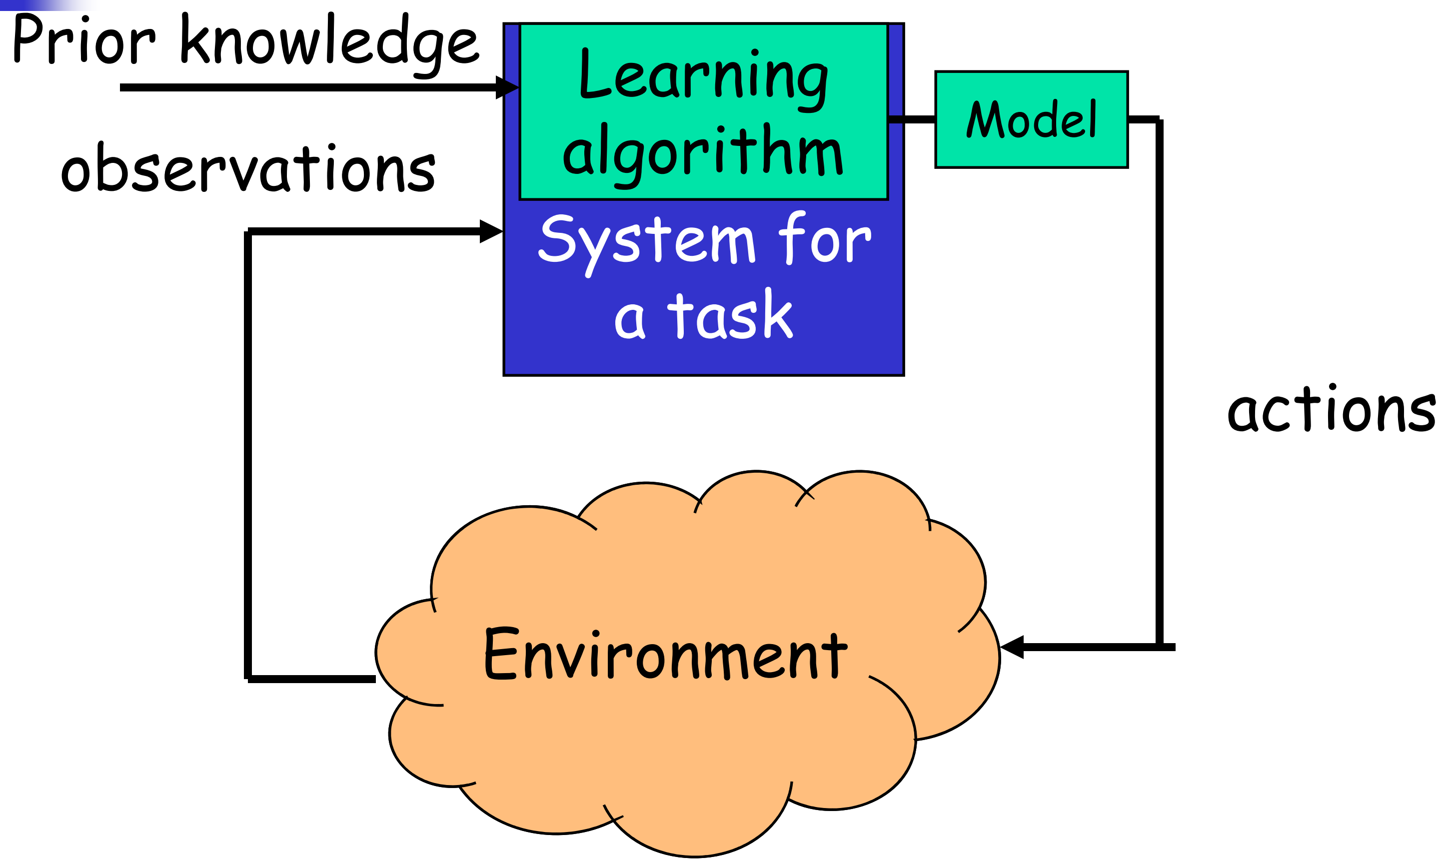

Calculate decisions on the basis of learned models of systems

The fine structure of adaptive embedded systems

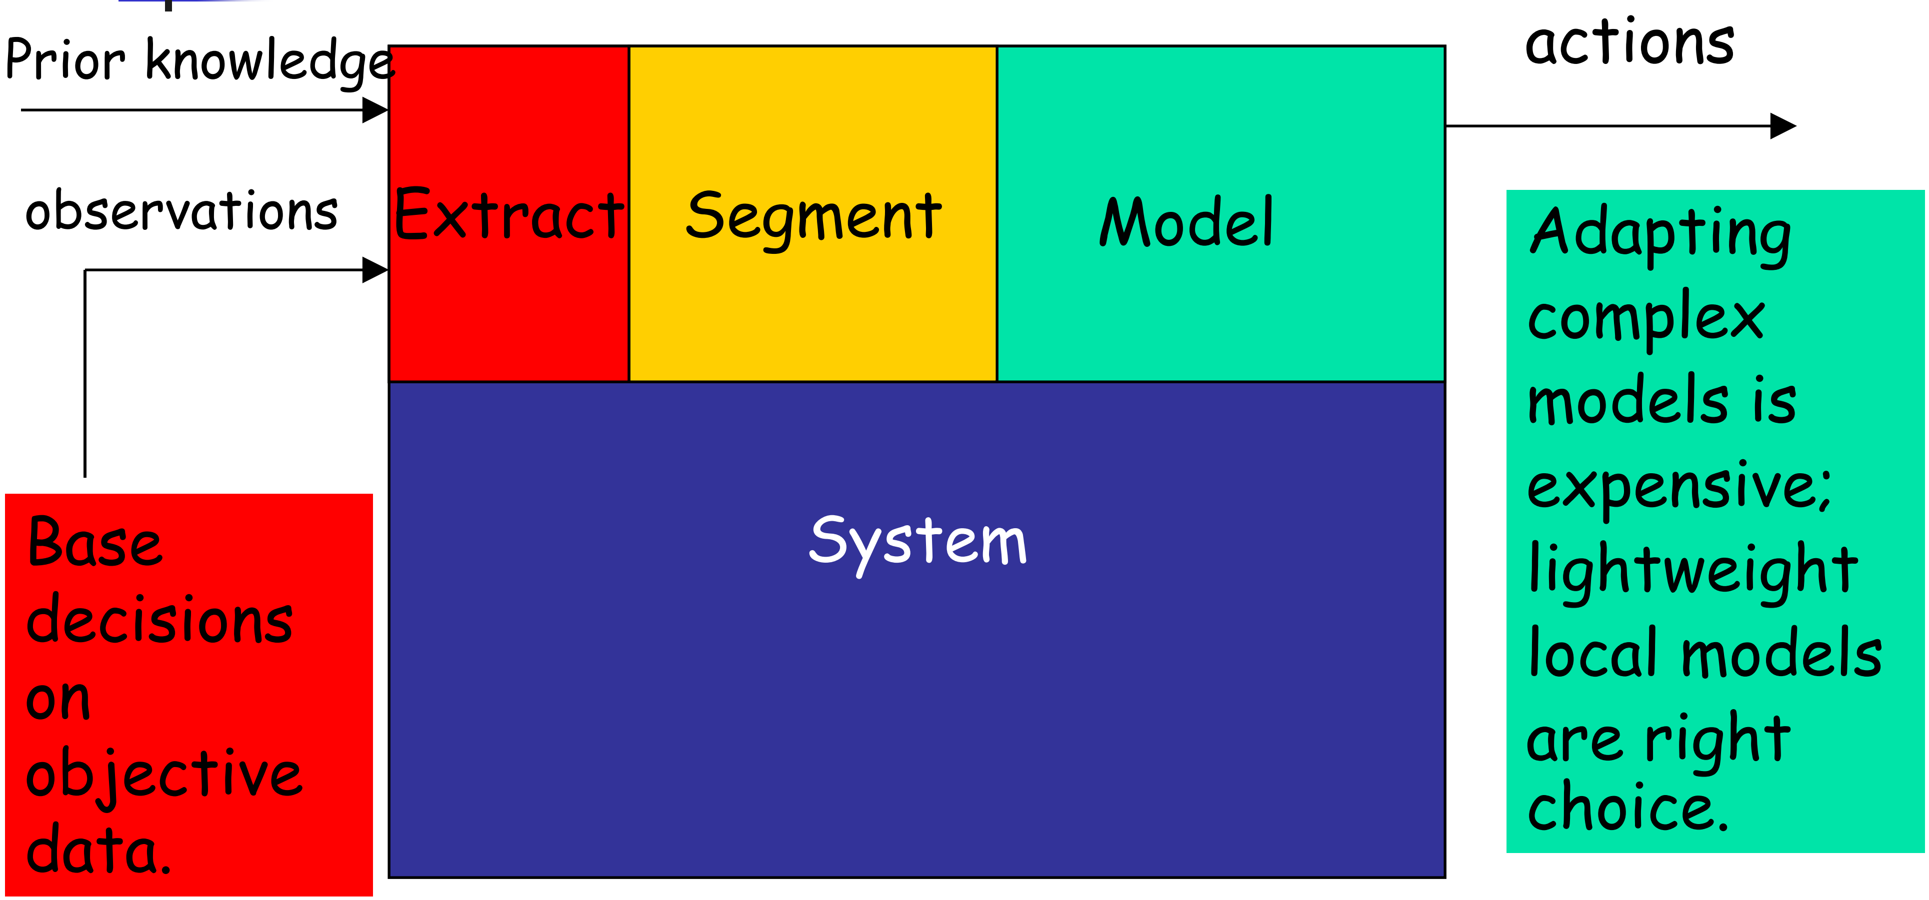

Non-stationarity is pervasive. Robust algorithms for detection

The vision

- “System identification” for large, non-stationary (distributed) systems.
- Off-the-shelf components for putting together feedback controllers with performance guarantees for such systems.

Collaborators

- Tracking human learning
  - Diana Gordon, ONR/University of Wyoming and Sandra Marshall, San Diego State University, George Zouridakis, University of Houston
- Tracking conflict
  - Richard Stoll, Rice University

Students

- Human learning
  - Richard Thrapp, National Instruments
  - Peggy Fidelman, PhD in CS/UT Austin
  - Igor Karpov, PhD in CS/UT Austin
  - Paul Ramirez
  - Gwen Thomas, Green Hills
  - Tony Berning
  - Gunes Ercal (CRA mentee)
  - Deborah Watt (CRA mentee)
  - Scott Griffin, Rational
  - Scott Ruthfield, Microsoft
  - Chris Gouge, Microsoft
  - Stephanie Weirich (Asst. Prof. at UPenn)
  - Sameer Siruguri, MS
  - Lisa Chang, MS, IBM
  - Nuvan Rathnayake, Rice junior
  - Ian Stevenson, PhD neuroscience, Northwestern
  - Farhan Baluch, University of Houston, MS 2006

Students

- Conflict
  - Michael Friedman, Rice sophomore
  - Adam Larson, Rice senior
  - Adam Stepinski, Rice sophomore
  - Clement Pang, Rice junior
  - Benedict Lee, MS 2007
  - Derek Singer, Rice junior

Sponsors

- Conflict analysis: NSF ITR 0219673
- Human learning: ONR N00014-96-1-0538