

BBN Technologies

Cognitive Learning And Decision Making for EW

Dr. Karen Zita Haigh, BBN
612-308-6726
khaigh@bbn.com

Date: *Aug 2015*

Distribution Statement A. Approved for Public Release, Distribution Unlimited.

BAE SYSTEMS

Raytheon
BBN Technologies

The Problem:

- Modern mobile communications networks operate in highly dynamic, potentially hostile environments
- Current approaches to EP and EA are usually limited to previously-seen RF environments

The Solution:

- Automatically learn to select actions that improve mission performance even in novel RF environments
 - **Characterize** the communications conditions
 - **Learn** the performance of the available responses
 - **Optimize** and implement the most effective strategy to improve mission performance

Learn how conditions affect mission success and optimize performance on-the-fly

PROBLEM FORMULATION

- Each node has a set of **observable parameters** that describe the signal environment
 - Often normalized, e.g., ranging from -1 to 1 (strongly “is not” to strongly “is”)
 - Local statistics
 - Shared (global) statistics if available

EXAMPLES

- Saturation
- Signal-to-noise ratio
- Error rates
- Gaussianness
- Repetitiveness
- Similarity to own communications signal
- Link and retransmission statistics
- Neighborhood size

Observables describe RF environment *behaviour*, not emitter names

- Each node has a set of **Controllable Parameters** that change radio behaviour
 - Each CP, c , has a known set of discrete values of size v_c
- **Strategy** is a combination of control parameters
 - Total of $\prod_{\forall c} v_c$ strategies
 - If all n CPs are binary on/off, then there are 2^n strategies, well beyond the ability of a human to manage.

EXAMPLES

- **Antenna:** e.g. beam forming, nulling
- **RF front end:** e.g. analog tunable filters, frequency-division multiplexing
- **PHY:** e.g. transmit power, notch filters, modulation scheme
- **MAC:** e.g. dynamic spectrum access, frame size, carrier sense threshold, reliability mode, unicast/broadcast, timers, contention window algorithm
- **Network:** e.g. neighbor discovery algorithm, thresholds, timers
- **Application:** e.g. compression (e.g., jpg 1 vs 10), method (e.g., audio vs video)

- Each node has a scalar **performance metric** that quantify how well the network satisfies requirements
- Operationally meaningful
 - Mission
 - Situational
 - Social (multi-user)
- Local estimates can be shared across the network to obtain measure of global performance

EXAMPLES

- Effectiveness:
 - Throughput
 - Latency
 - Bit-error-rate
 - EW BDA
- Cost:
 - Power
 - Overhead
 - Probability of detection

- Each node builds a model f that estimates how each candidate strategy s will perform in the current environment o_t

$$\forall s, \tilde{m}_s = f(o_t, s)$$

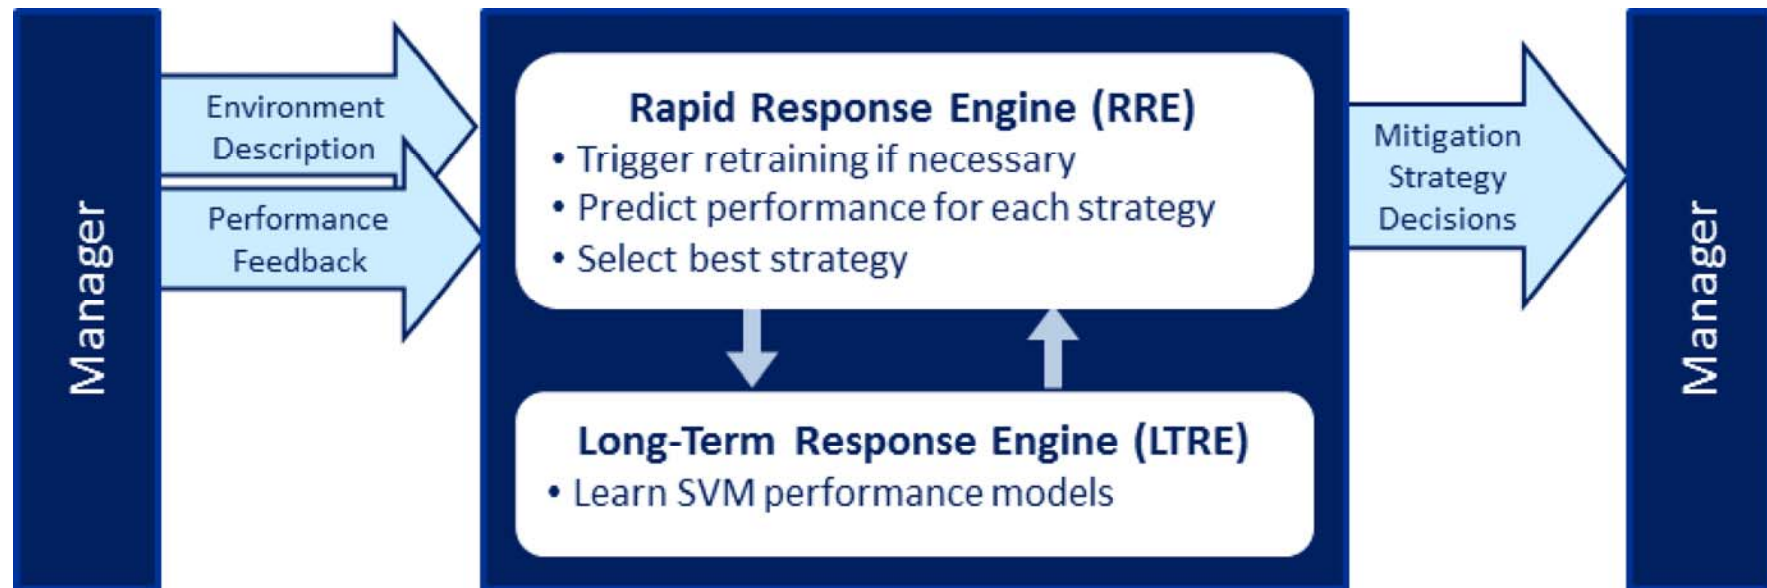
s	Strategy
m	Metric
f	Support Vector Model for metric
o_t	Observations at time t

- From training data, collected previously or during current mission
- Support Vector Regression Machines (Vapnik, 1995; Drucker et al, 1997)
- The model predicts performance for ALL possible strategies, whether or not they appeared in the training data

The Strategy Optimizer learns the performance of all controllables against all communications environments

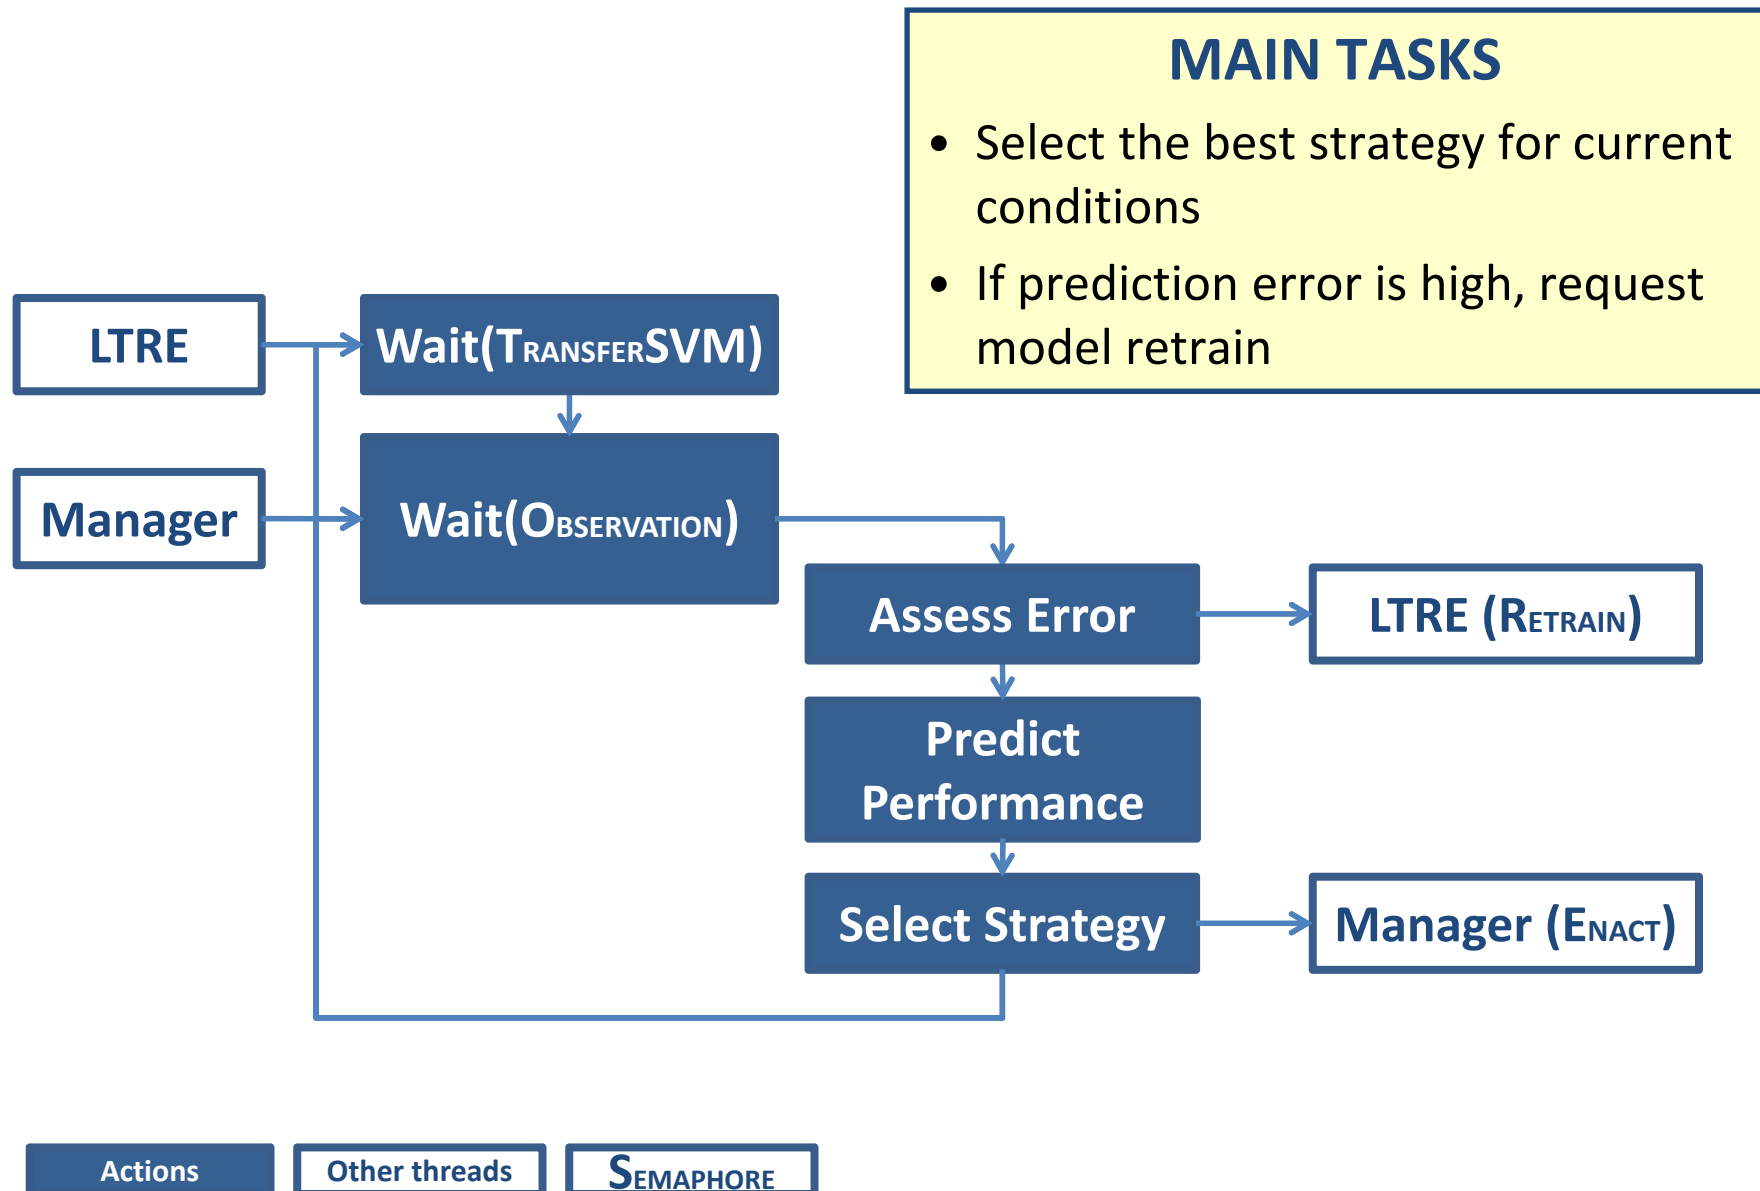
STRATEGY OPTIMIZER ARCHITECTURE

Parallel Learning & Decision Making

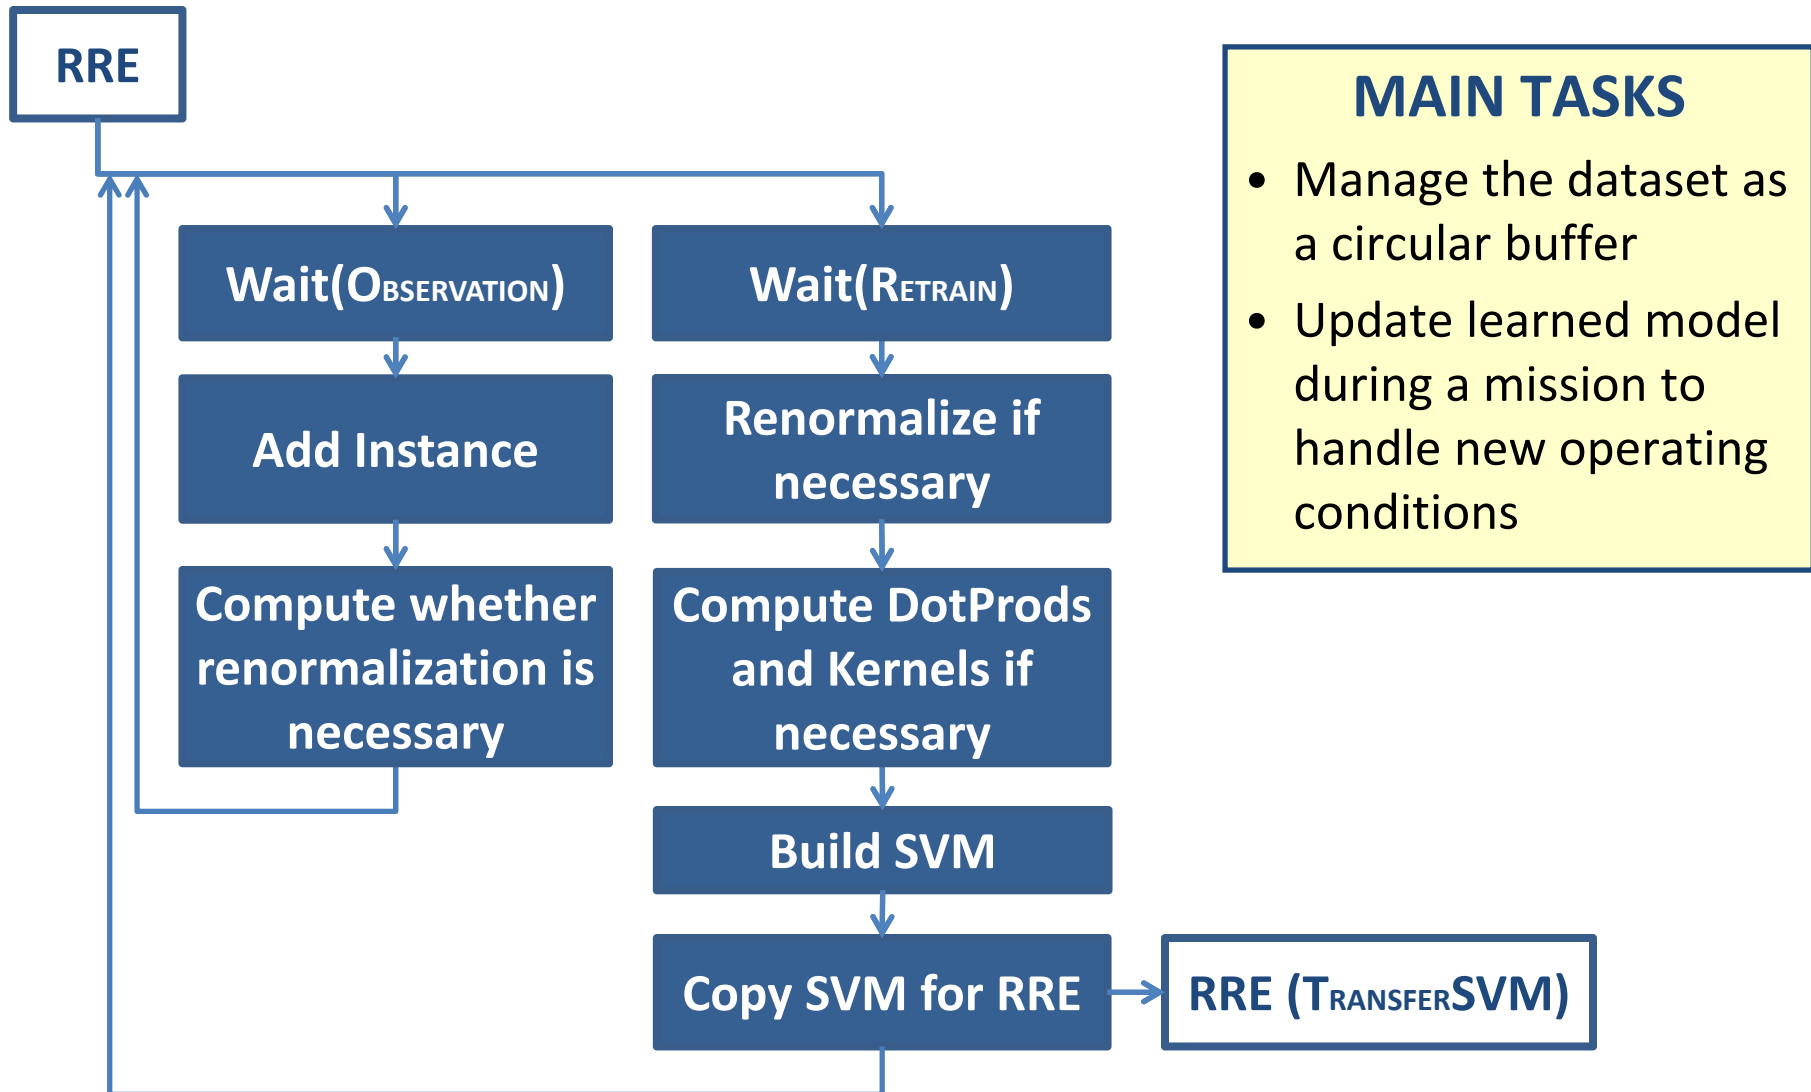


- **RRE:** Adaptively selects strategies in real-time to optimize performance metrics
- **LTRE:** A cognitive learning loop that builds models to describe new RF environments

Rapid Response Engine (RRE)



Long Term Response Engine (LTRE)



Actions

Other threads


SEMAPHORE

RESULTS

- Compare adaptive Strategy Optimizer to a static system
- Compare incremental learning system to adaptive system
- A detailed incremental learning example
- Aggregate incremental learning
- Parallel RRE decision making and LTRE incremental learning

Results: Adaptive vs Static System


Dynamic adaptive system performs better than static system



Compare (a) static system with one fixed strategy to (b) system that adaptively chooses strategy as RF conditions change

Results: Adaptive vs Static System


Dynamic adaptive system performs better than static system



Compare (a) static system with one fixed strategy to (b) system that adaptively chooses strategy as RF conditions change

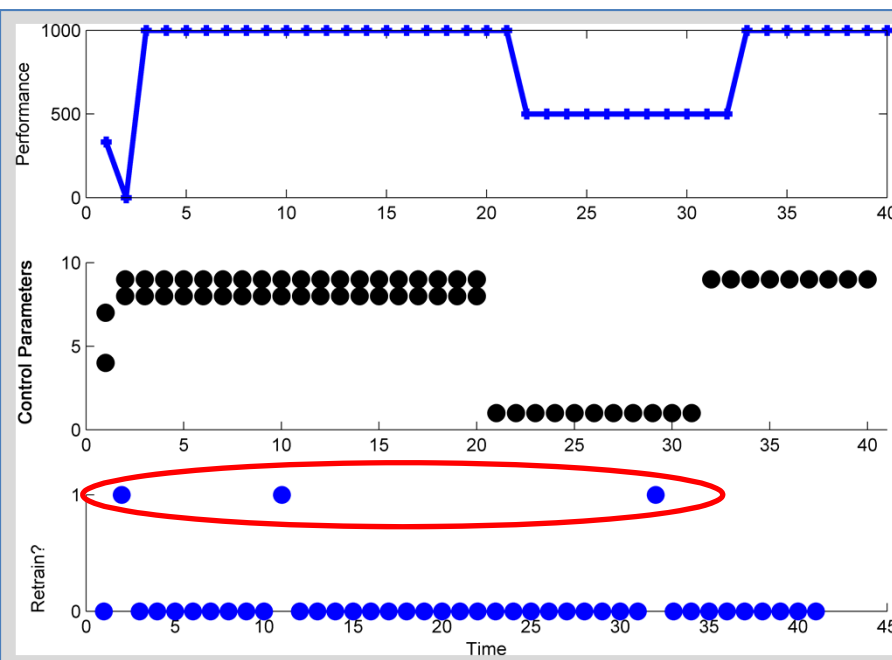
Results: Cognitive vs Adaptive System

Cognitive incremental learning performs better than dynamic adaptive system (even when both start with learned models)



Average performance = 275
(31% of optimal)

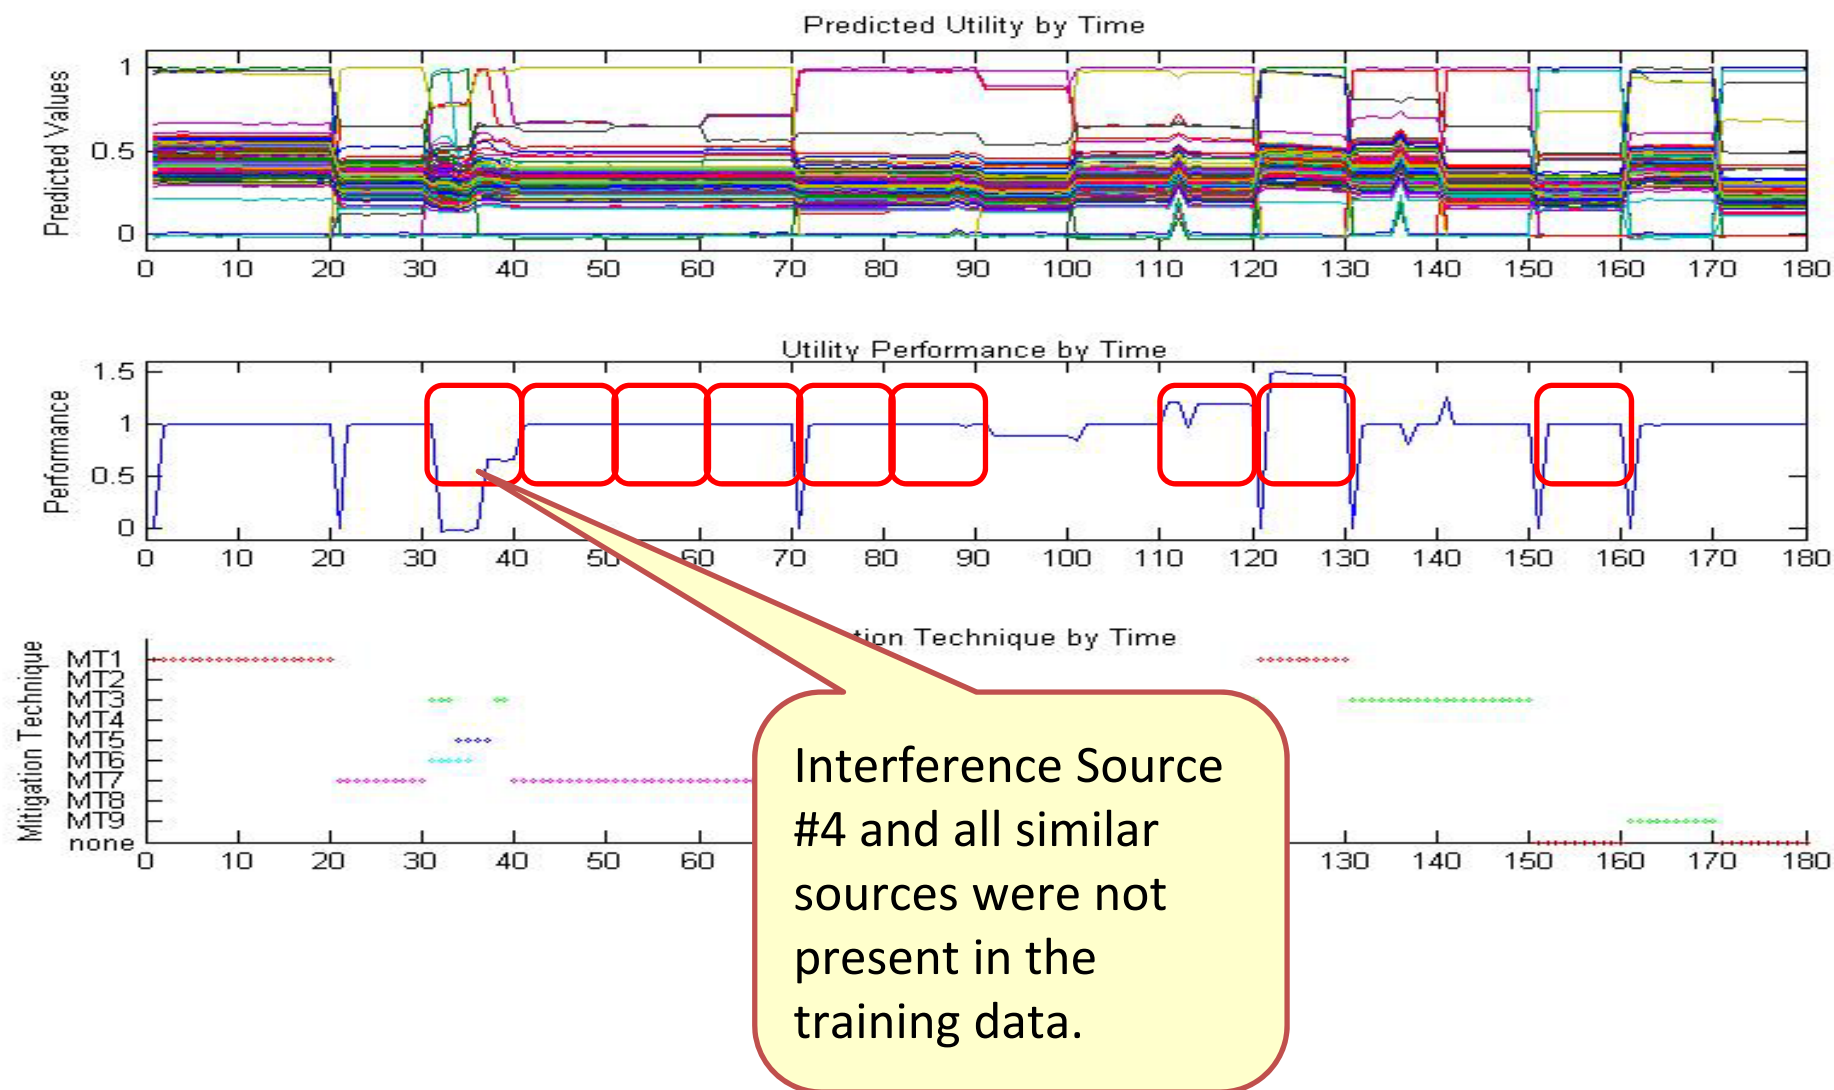
(a) Adaptive system that does not update models in-mission



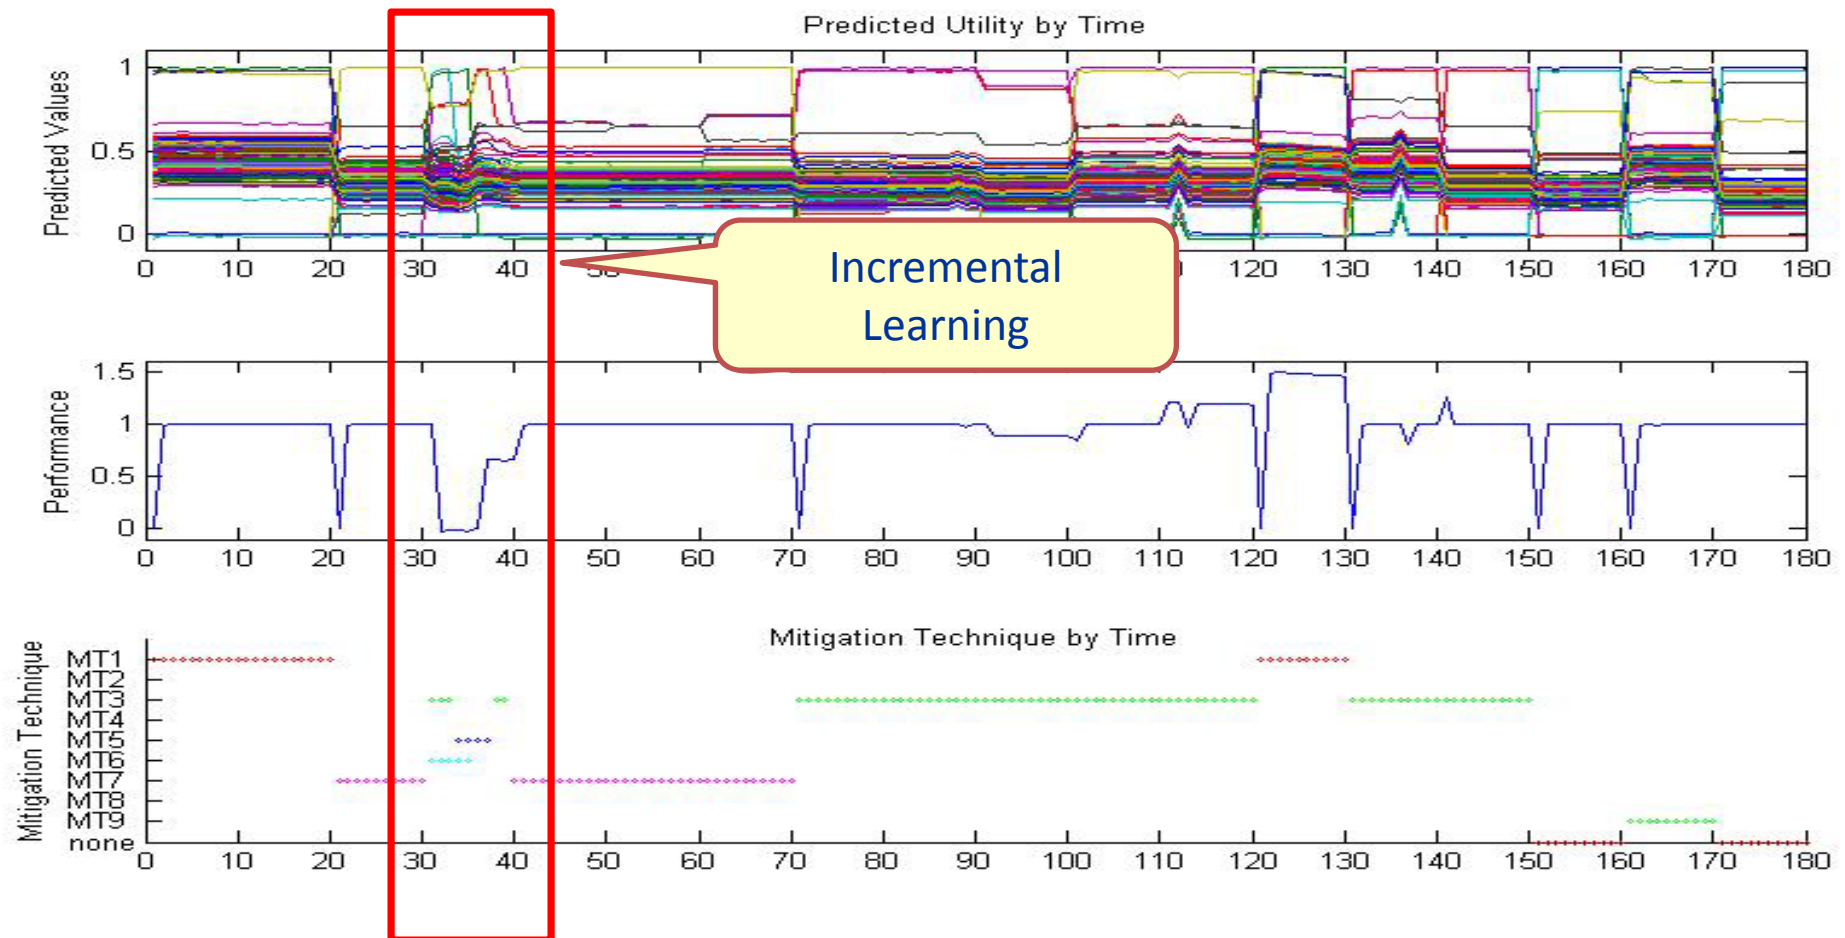
Average performance = 821
(94% of optimal)

(b) Cognitive adaptive system that incrementally learns models


Results: Detailed Incremental Learning (1)



Results: Detailed Incremental Learning (2)



Results: Detailed Incremental Learning (3)

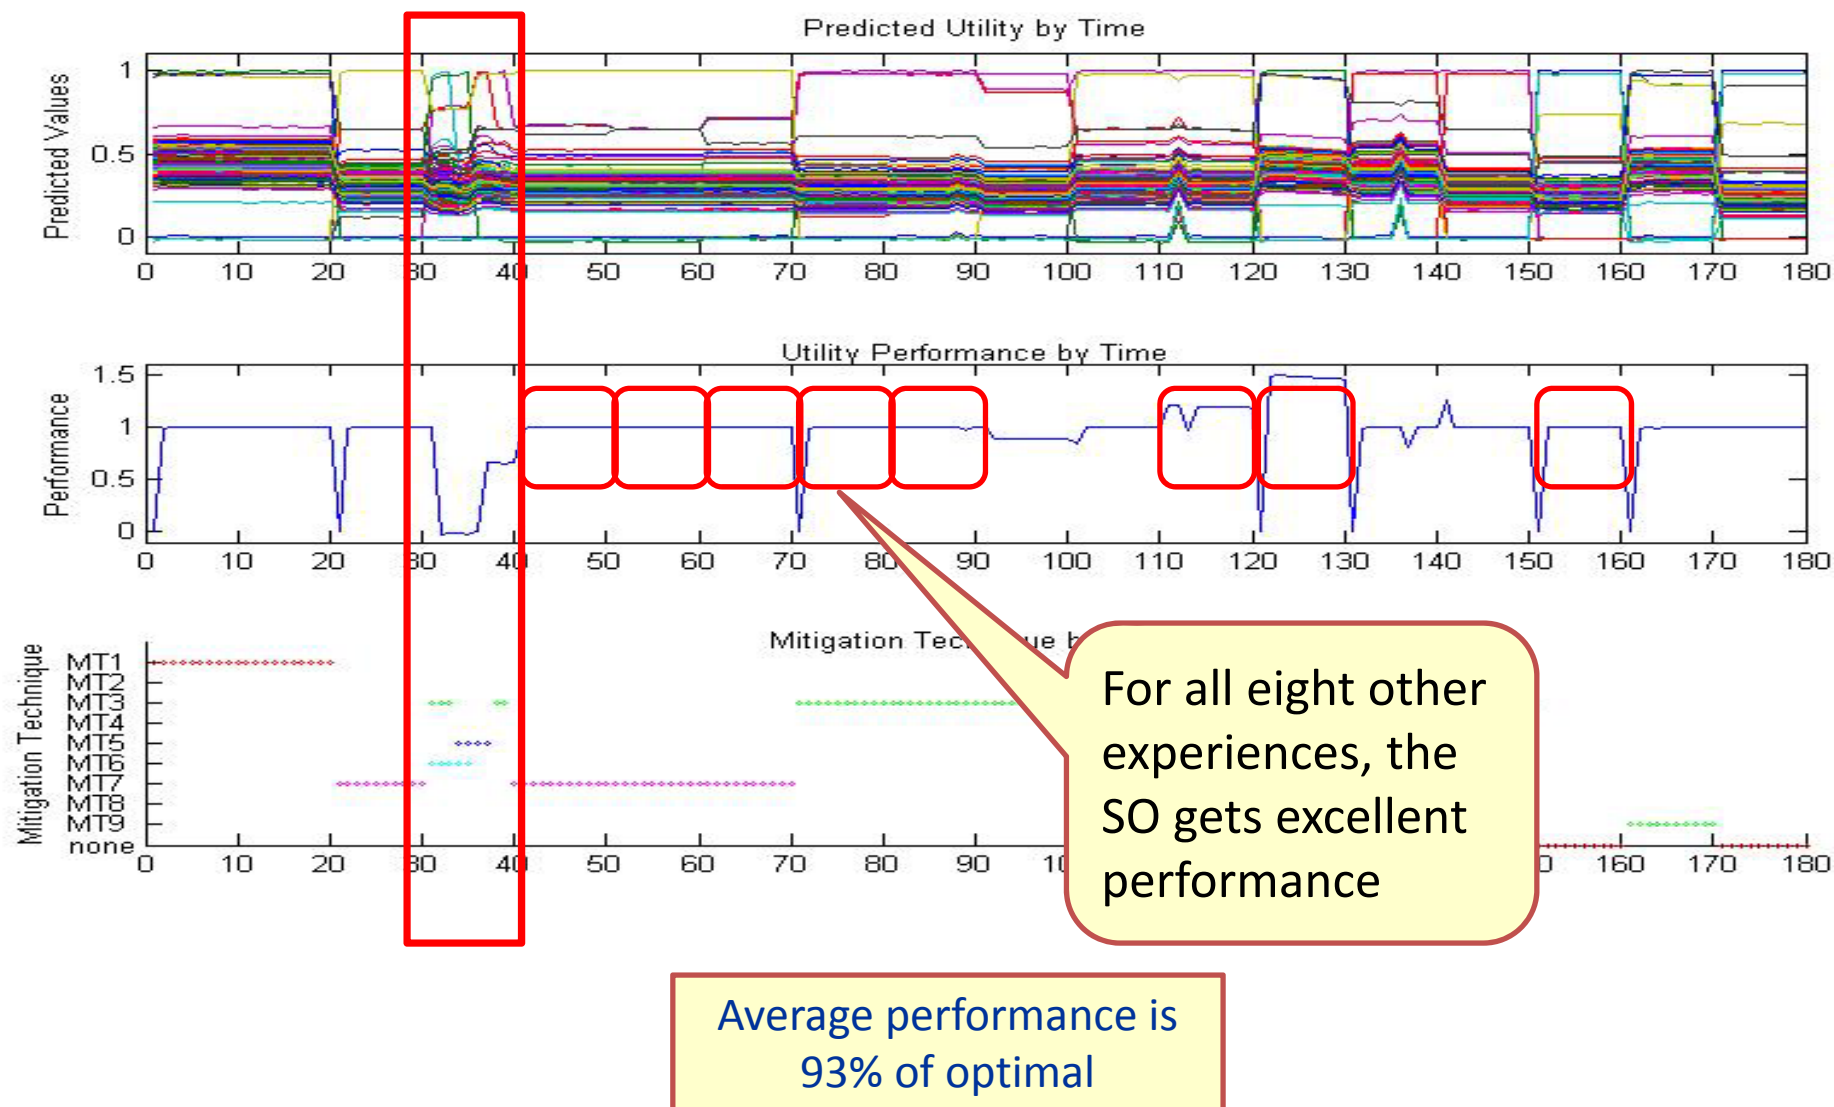


Try different strategies and learn from them until performance is sufficient for mission

Interference Source #4: Estimates of Throughput


Time	MT3	MT5	MT6+3	MT6+5	MT7	Observed
30	751.4	751.2	969.3	948.4	749.5	0.0
31	751.3	751.2	953.4	949.3	749.6	0.0
32	751.2	752.0	950.8	948.4	749.0	0.0
33	750.6	750.9	402.8	949.6	749.1	0.0
34	750.1	750.3	376.4	414.9	748.8	500.0
35	750.9	749.1	376.6	414.9	748.9	500.0
36	752.1	501.2	373.4	378.4	750.7	500.0
37	749.4	501.2	372.2	377.5	749.3	500.0
38	502.5	502.9	336.3	375.0	750.1	750.0
39	501.9	501.8	335.6	374.2	749.1	750.0

Results: Detailed Incremental Learning (4)



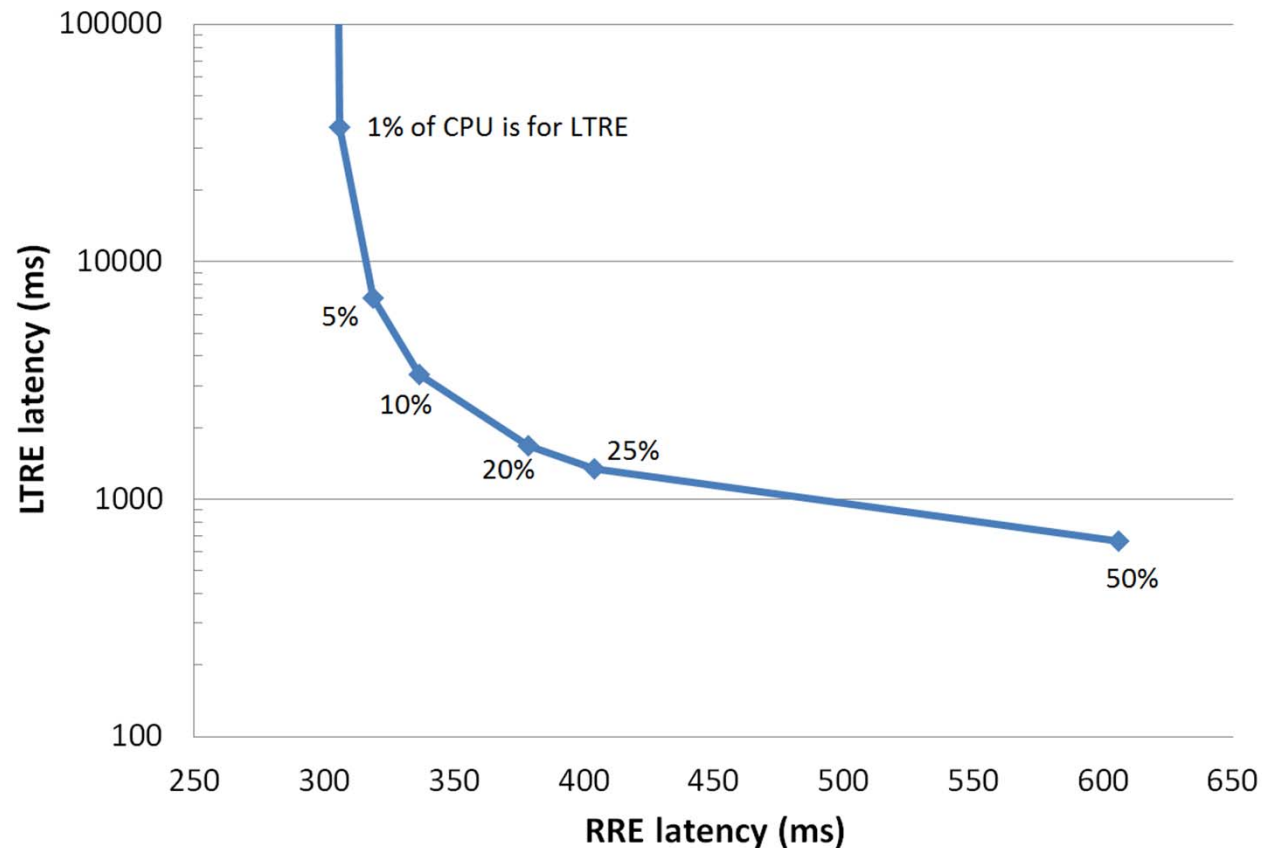
Results: Aggregate Incremental Learning

Cognitive incremental learning handles new communications conditions with only a small loss of optimality



Results: Sharing the processor

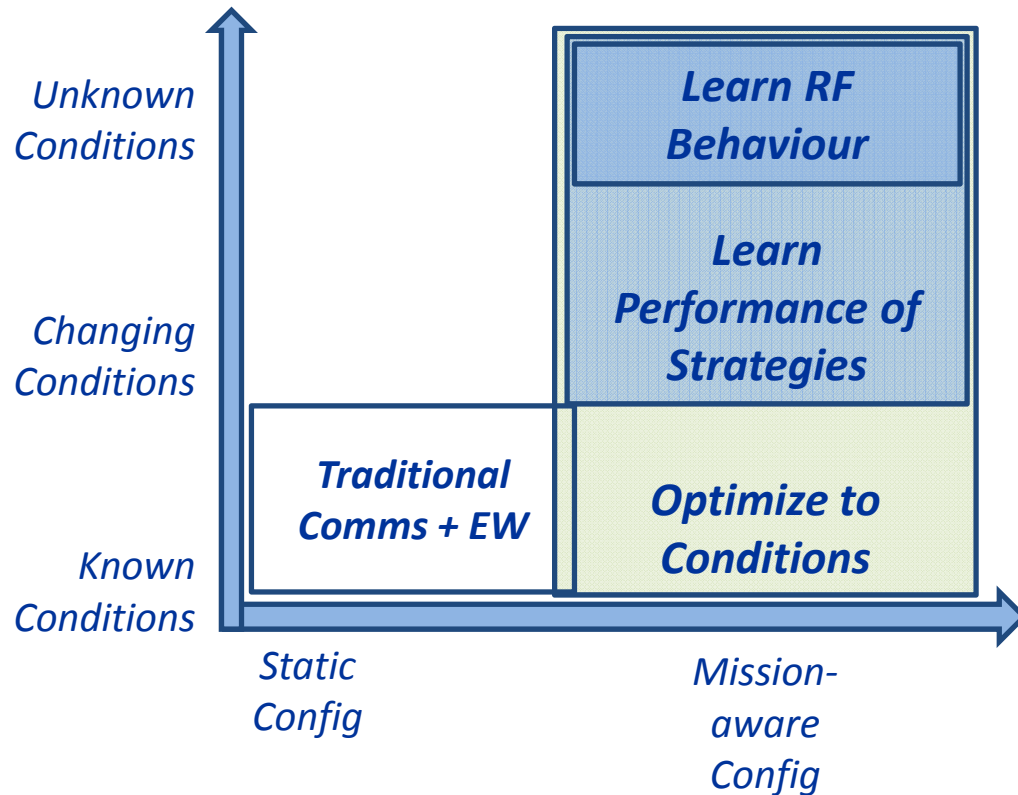
Trade RRE latency for LTRE latency, as a function of CPU sharing



	CPU	OS	Compiler
ARMv7	IBM ARMv7 rev 2 (v7I), 800MHz, 256 kB cache, 256MB RAM, vintage 2005	Linux version 2.6.38.8	g++ 4.3.3, 2009

SUMMARY

Strategy Optimizer Key Capabilities



- **Rapid adaptive decision making** selects actions in real-time to optimize mission performance
- **Incremental Learning** learns to optimize mission performance in complex, changing & unknown environments
- **Semantically Agnostic Architecture** supports easy deployment to new platforms and domains
 - does not depend on meaning of observables, controllables or performance metrics

**Rapid adaptive decision making + cognitive learning
for unknown environments**

Acknowledgements

Colleagues:

- Allan MacKay (BBN)
- Michael Cook (BBN)
- Li Lin (BBN)
- John Tranquilli (BAE)
- Amber Dolan (BAE)
- Dianne Egnor (JHU/APL)
- Vincent Sabio (formerly DARPA/STO)
- Bruce Fette (formerly DARPA/STO)

The views, opinions, and/or findings contained in this article/presentation are those of the author(s)/presenter(s) and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.