Directed Exploration of Complex Systems

AISR/CIDU Joint Meeting
NASA Ames Research Center
2009/10/16

Michael C. Burl*,
Esther Wang^
Brian Enke+
William J. Merline+

*Jet Propulsion Laboratory
^California Institute of Technology
+Southwest Research Institute

Overview

Motivation: Physics-based simulations are widely used to model complex systems

+ high-fidelity representation of actual system behavior

-- cumbersome and computationally expensive

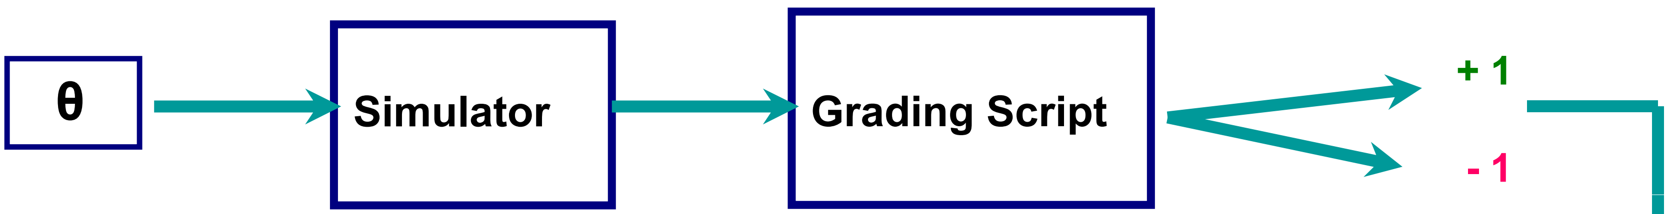

Goal: Build a simplified predictive model of system behavior and use this model to understand which input parameters produce a desired output (+1) from the simulator.

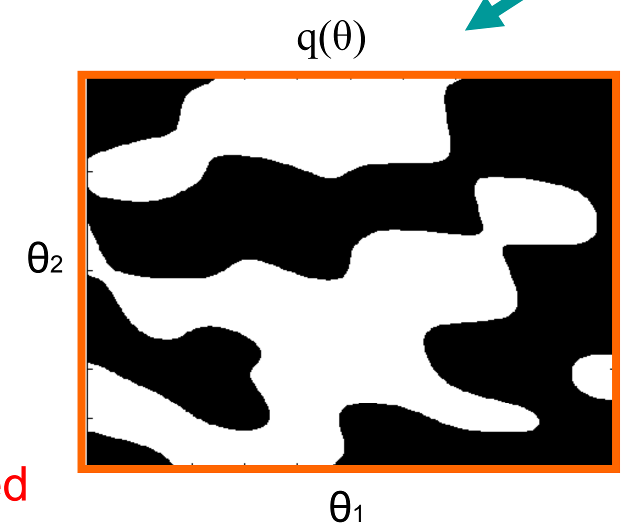

❖ Really want a function:

$$\hat{q} : \Theta \rightarrow \{-1, +1\}$$

such that $\hat{q}$ agrees with **q** over most of domain

Assumptions: the simulator is deterministic (and non-chaotic) and the learner is only informed of a binary-valued outcome.

Example: Asteroid Collisions

Currently working with asteroid collision simulator:

SPH + N-body gravity code

Various science questions can be studied with this simulator:

Under what conditions is a collision "catastrophic"?

How are asteroid satellites (e.g., Ida-Dactyl) generated?

How were particular asteroid families (e.g., Karin) formed?

But … Input space is 5-dimensional

Tgt body size (km)
Impactor velocity (km/s)
Impactor angle (degrees)
Impactor size (km)
Target composition (rubble/solid)

And … Each trial takes 1 CPU-day!

Grid sampling scales ~k^d where k is number of steps along each dimension

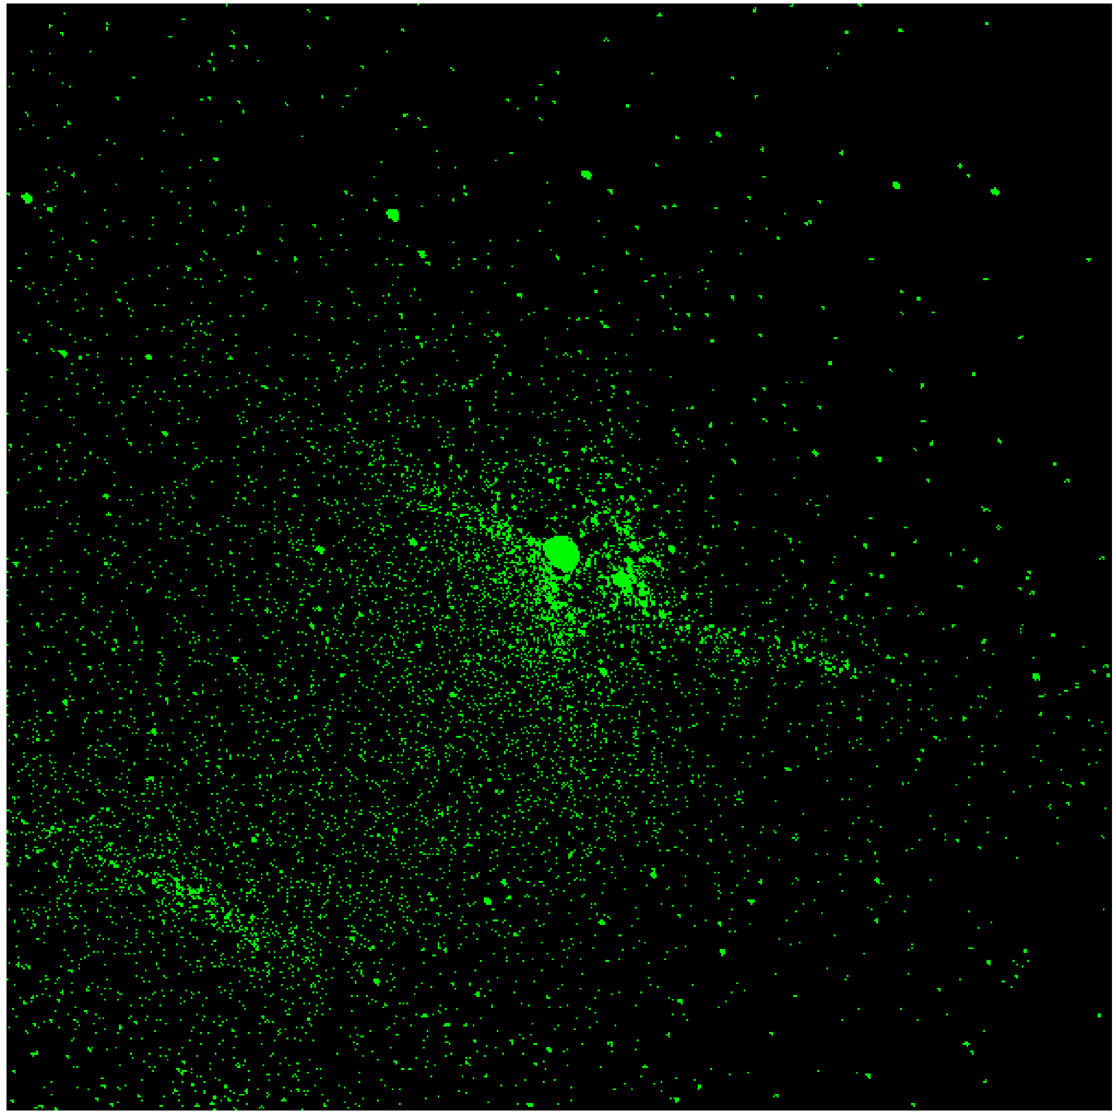

Movie courtesy B. Enke, D. Durda, SwRI.

Directed Exploration Approach

Current Knowledge

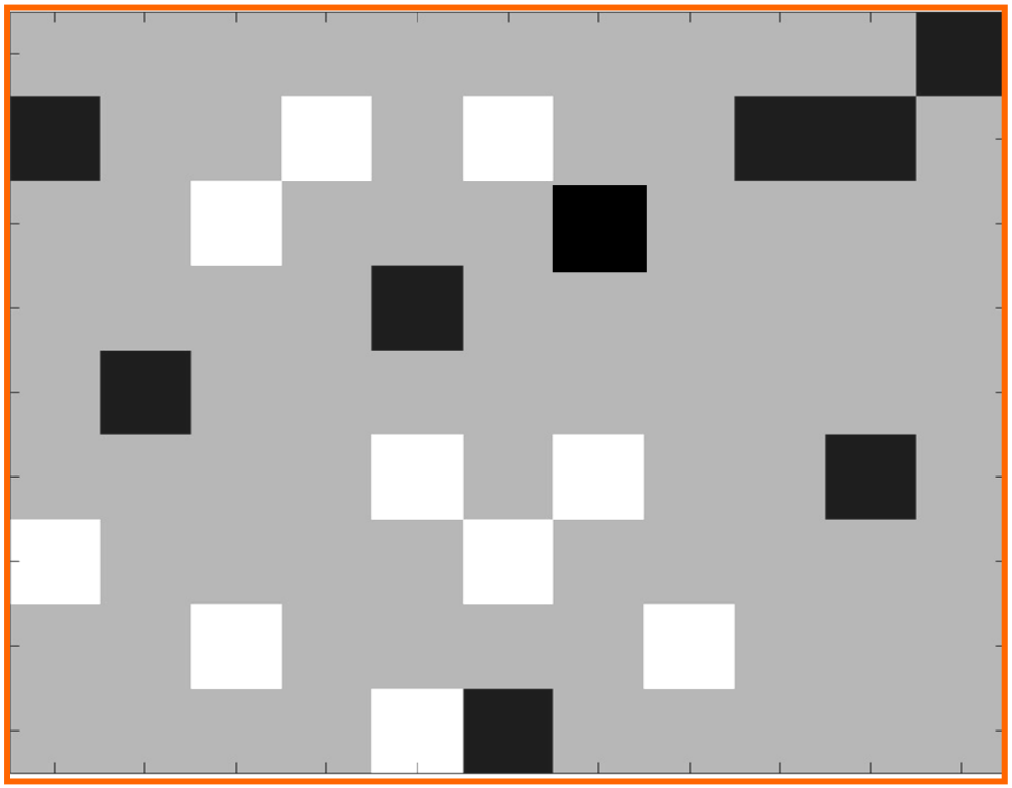

Grey: **Points with unknown labels**
White: **Points with known labels + 1**
Black: **Points with known labels - 1**

1. **Learn a model from current knowledge.**

2. **Decide which unlabeled point would be most valuable to label.**

3. **Get the label for the selected point from the oracle (by running the simulator/ grading script) and add that point and its label into "current knowledge".**

4. **Repeat until we have a good understanding of the oracle.**

Two Main Steps

1. **Supervised Learning:** Learn a model from current knowledge.
   - Support Vector Machines (SVM)
   - Kernel Density Estimation (KDE)
   - Gaussian Process Classifier (GP)

2. **Active Learning:** Decide which unlabeled point would be most valuable to label next.
   - **Passive** – randomly select an unlabeled point.
   - **Most-Confused Point (MCP)** – select the unlabeled point whose label is most uncertain.
   - **Most-Informative Point (MIP)** – select the unlabeled point such that the expected information gain about the entire set of unlabeled points is maximized. [Holub et al, 2008]
   - **Meta-strategy** – treat the base active learner’s valuation as expressing a preference over points; choose randomly while honoring the strength of the preference. (Similar to the epsilon-greedy approach used in RL to balance exploration and exploitation.) [Burl et al, 2006]

Active Learning

- **Determine which new point(s) in input space, if labeled, would be most valuable.**

- **Imagine discretized input space:**
  - **$L^{(n)}$ = set of labeled instances at end of trial n (red or green)**
  - **$U^{(n)}$ = set of unlabeled instances at end of trial n (black)**

- **Ideally, choose $\theta^{(n+1)}$ from $U^{(n)}$ such that the classifier learned from**

$$L^{(n+1)} = L^{(n)} + (\theta^{(n+1)}, q(\theta^{(n+1)}))$$

will bring $\hat{q}$ closer to q.

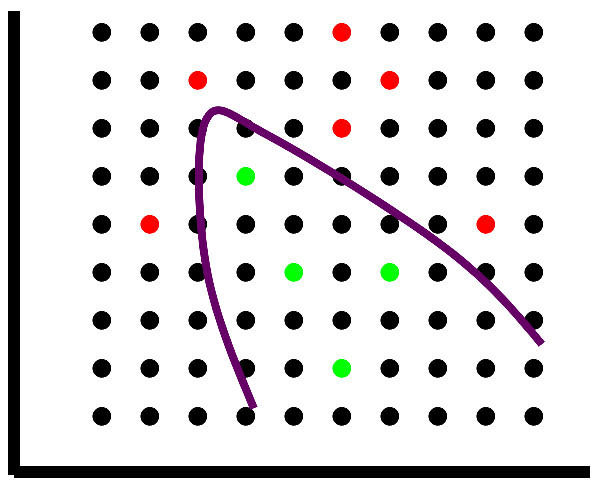

MCP: Choose points near current decision boundary

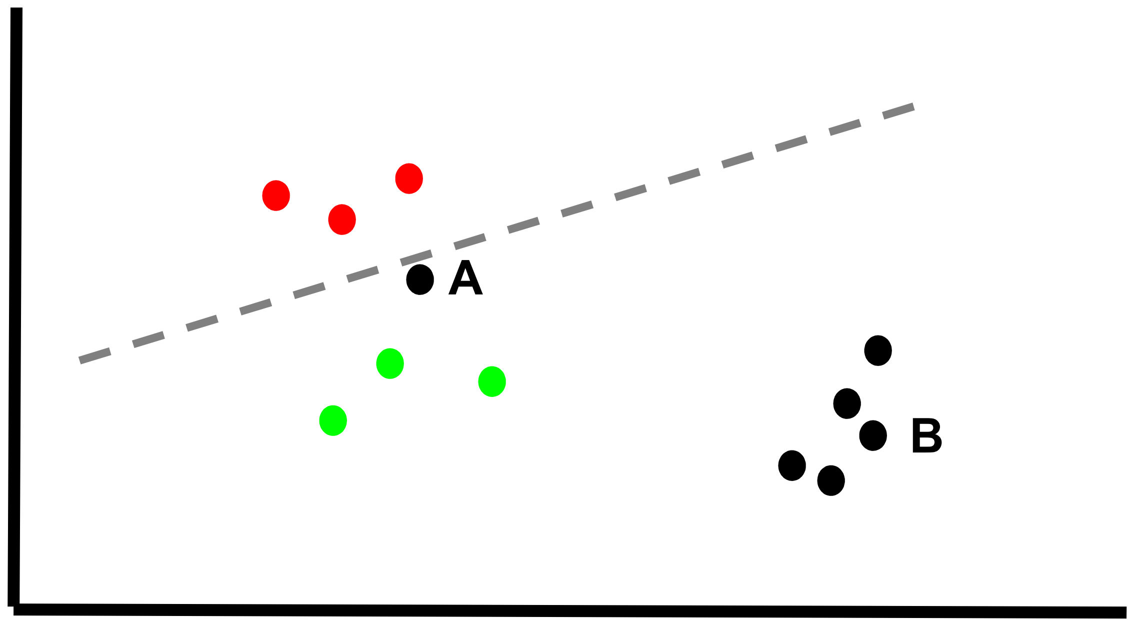

Figure from [Holub et al, 2008].

•MCP would choose Point A, but knowing the label of this point would not reveal much about the labels of the other unlabeled points.

•On the other hand, knowing the label of Point B would probably reveal the labels of the other nearby unlabeled points.

MIP: Most Informative Point

Current Knowledge

Grey: Points with unknown labels
White: Points with known +1 labels
Black: Points with known -1 labels

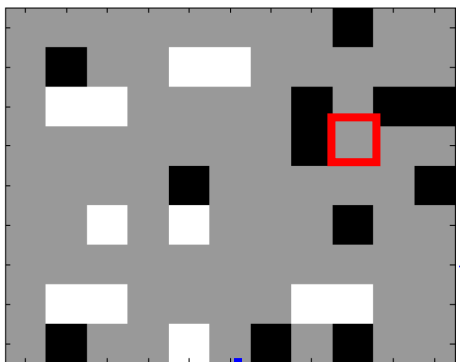

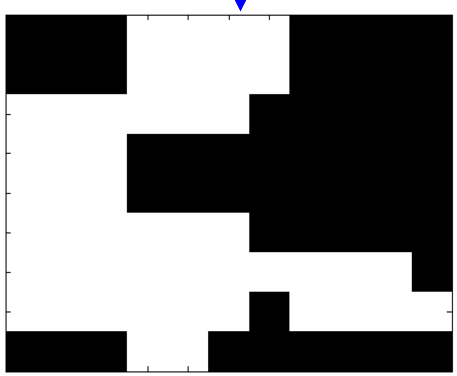

Prediction given Current Knowledge

Probability of labels being +1 (assuming y_h = +1)

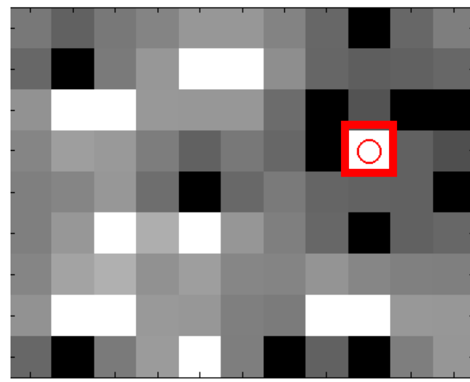

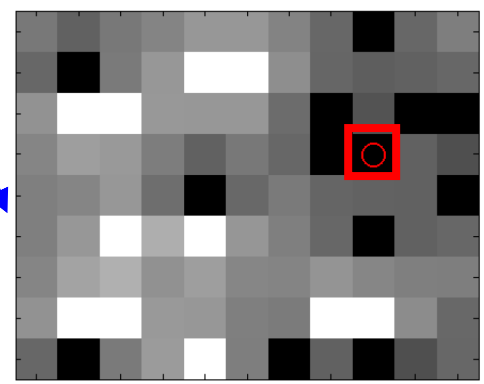

Probability of the labels being +1 (assuming y_h = -1)

Expected Information Gain

$$I_{gain} = H_0 - (H_+p_+ + H_-p_-)$$

$$H(Y) = -\Sigma\, p(y_i) \log p(y_i)$$

H(Y): Entropy/measure of uncertainty
$p(y_i)$: Probabilities of label i being +1

Synthetic Oracles

- Difficult to evaluate performance with real simulator since so slow and the true q() function is unknown => Use synthetic oracles initially.
- Added benefit: other researchers can try to replicate or improve upon results.

1

2

3

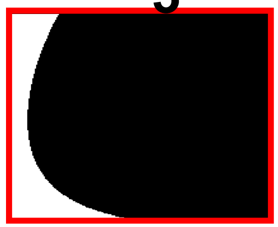

4

5

6

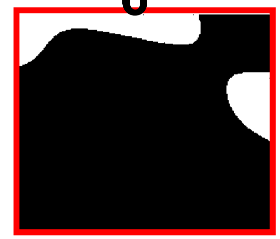

7

8

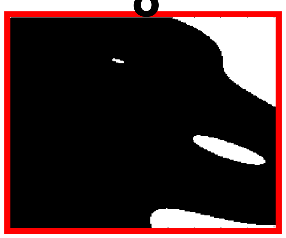

9

13

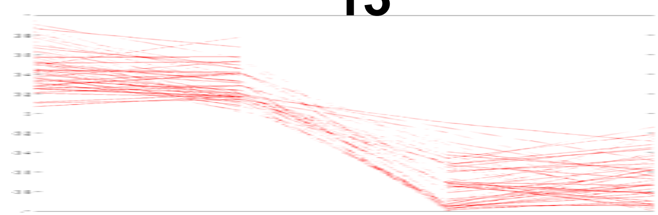

10

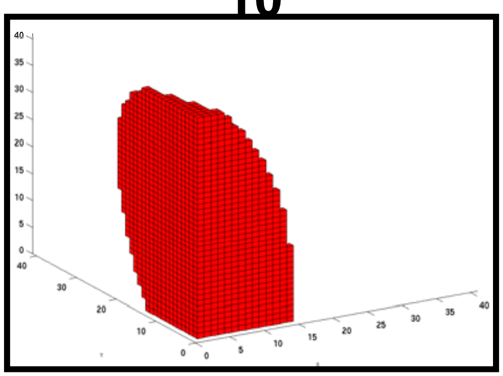

11

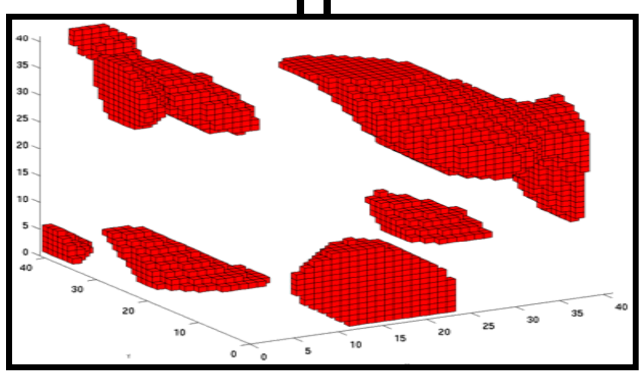

12

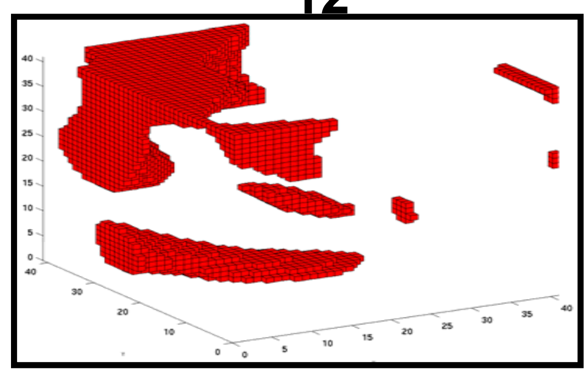

Results

Run each algorithm-oracle combination for 400 rounds of active learning.

Repeat seven times to get handle on variance.

Use classifier learned after r rounds to make a probabilistic prediction at each point in a fine discretization of the domain.

Threshold probabilities and compare against ground truth to generate ROC curves.

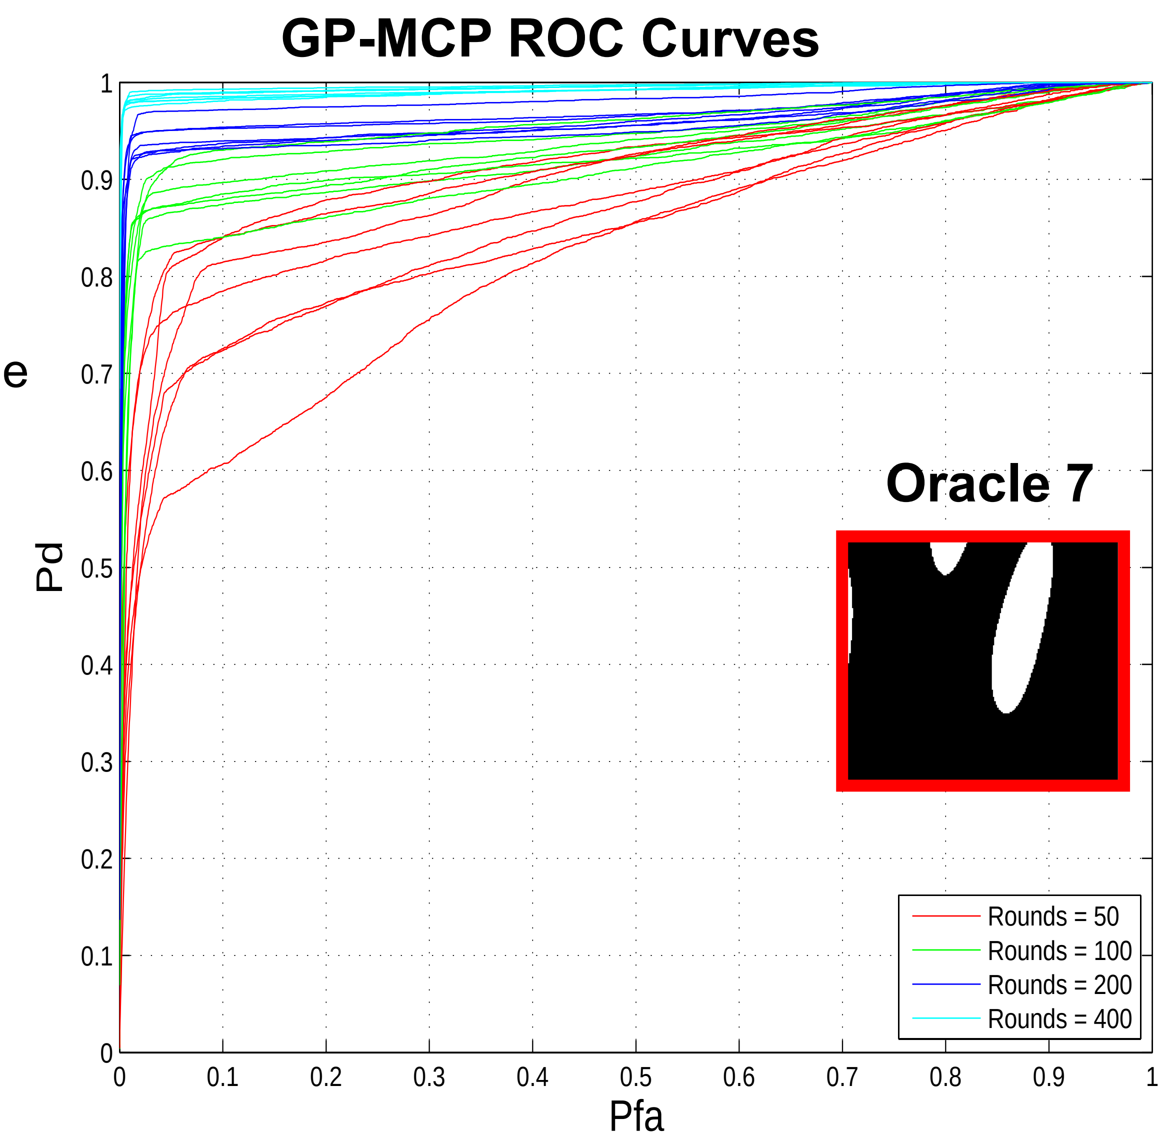

Pd @ Pfa = 0.05 on Oracle 9

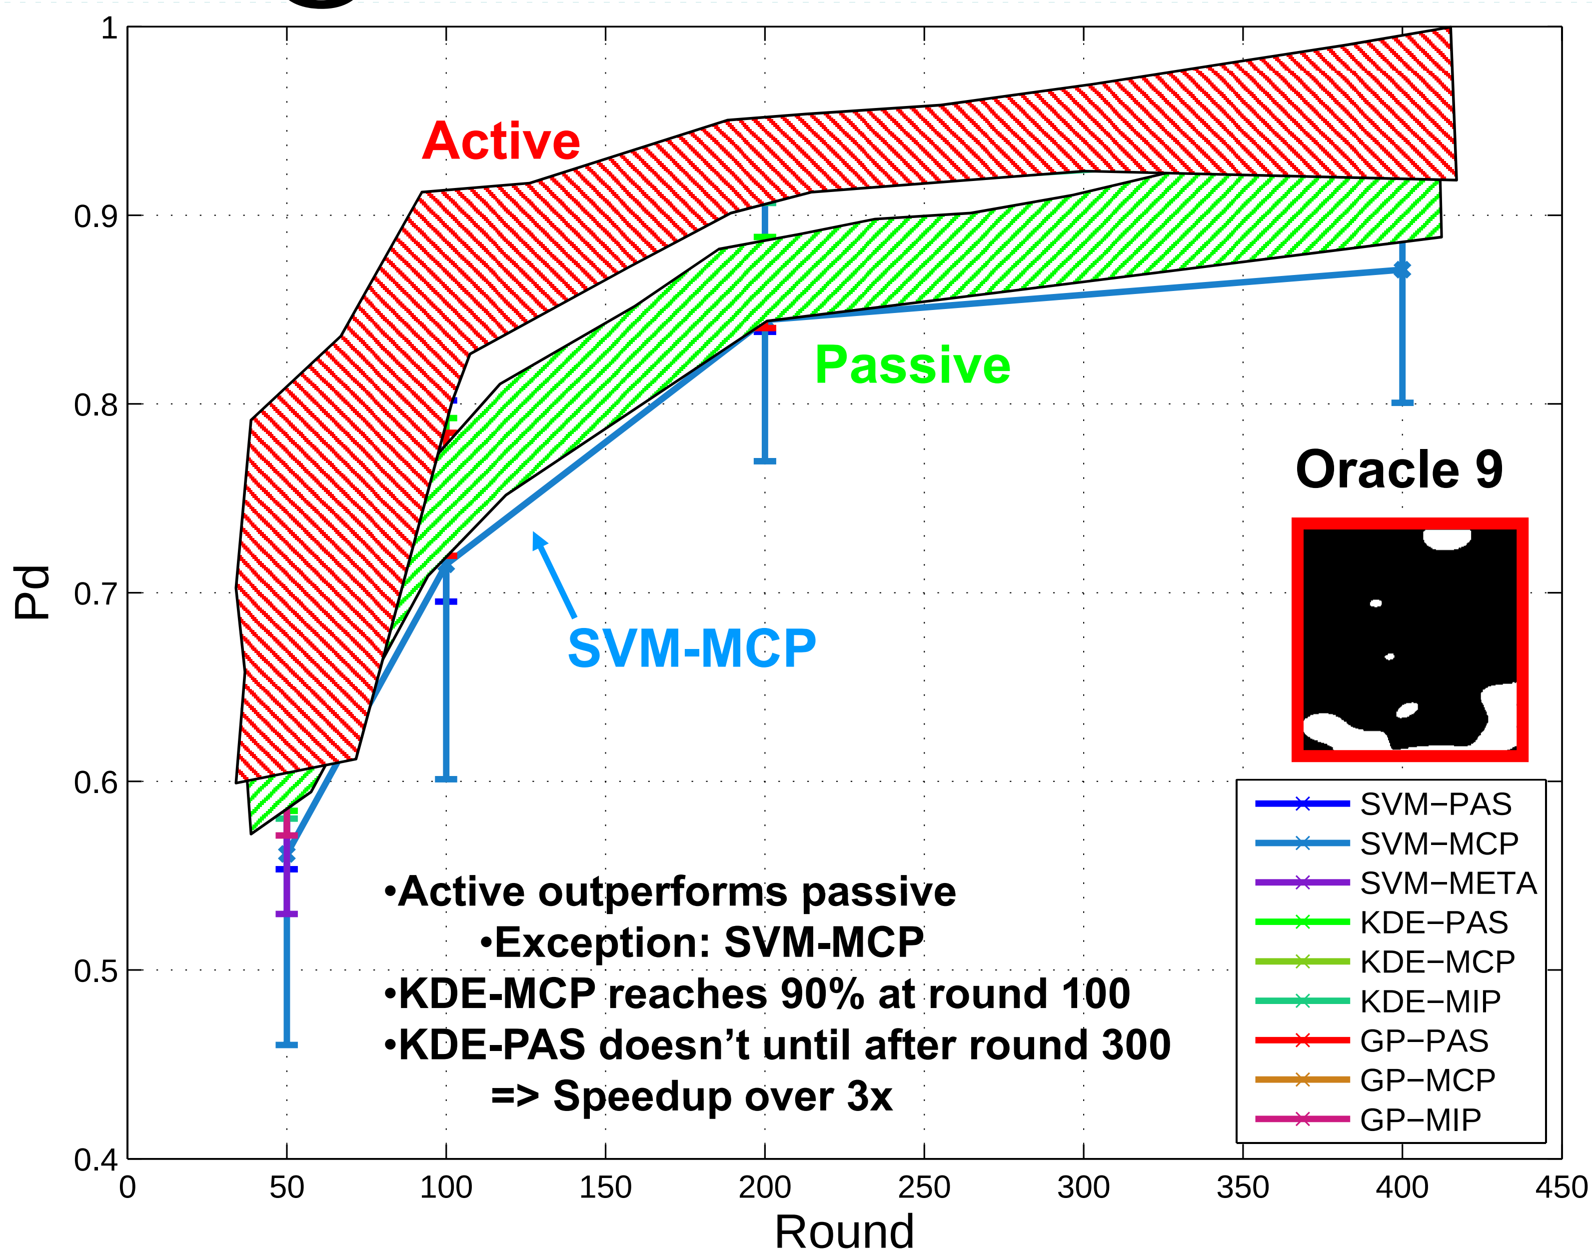

Initial Grid Evaluation (5D: 3 x 5 x 4 x 4 x 6)

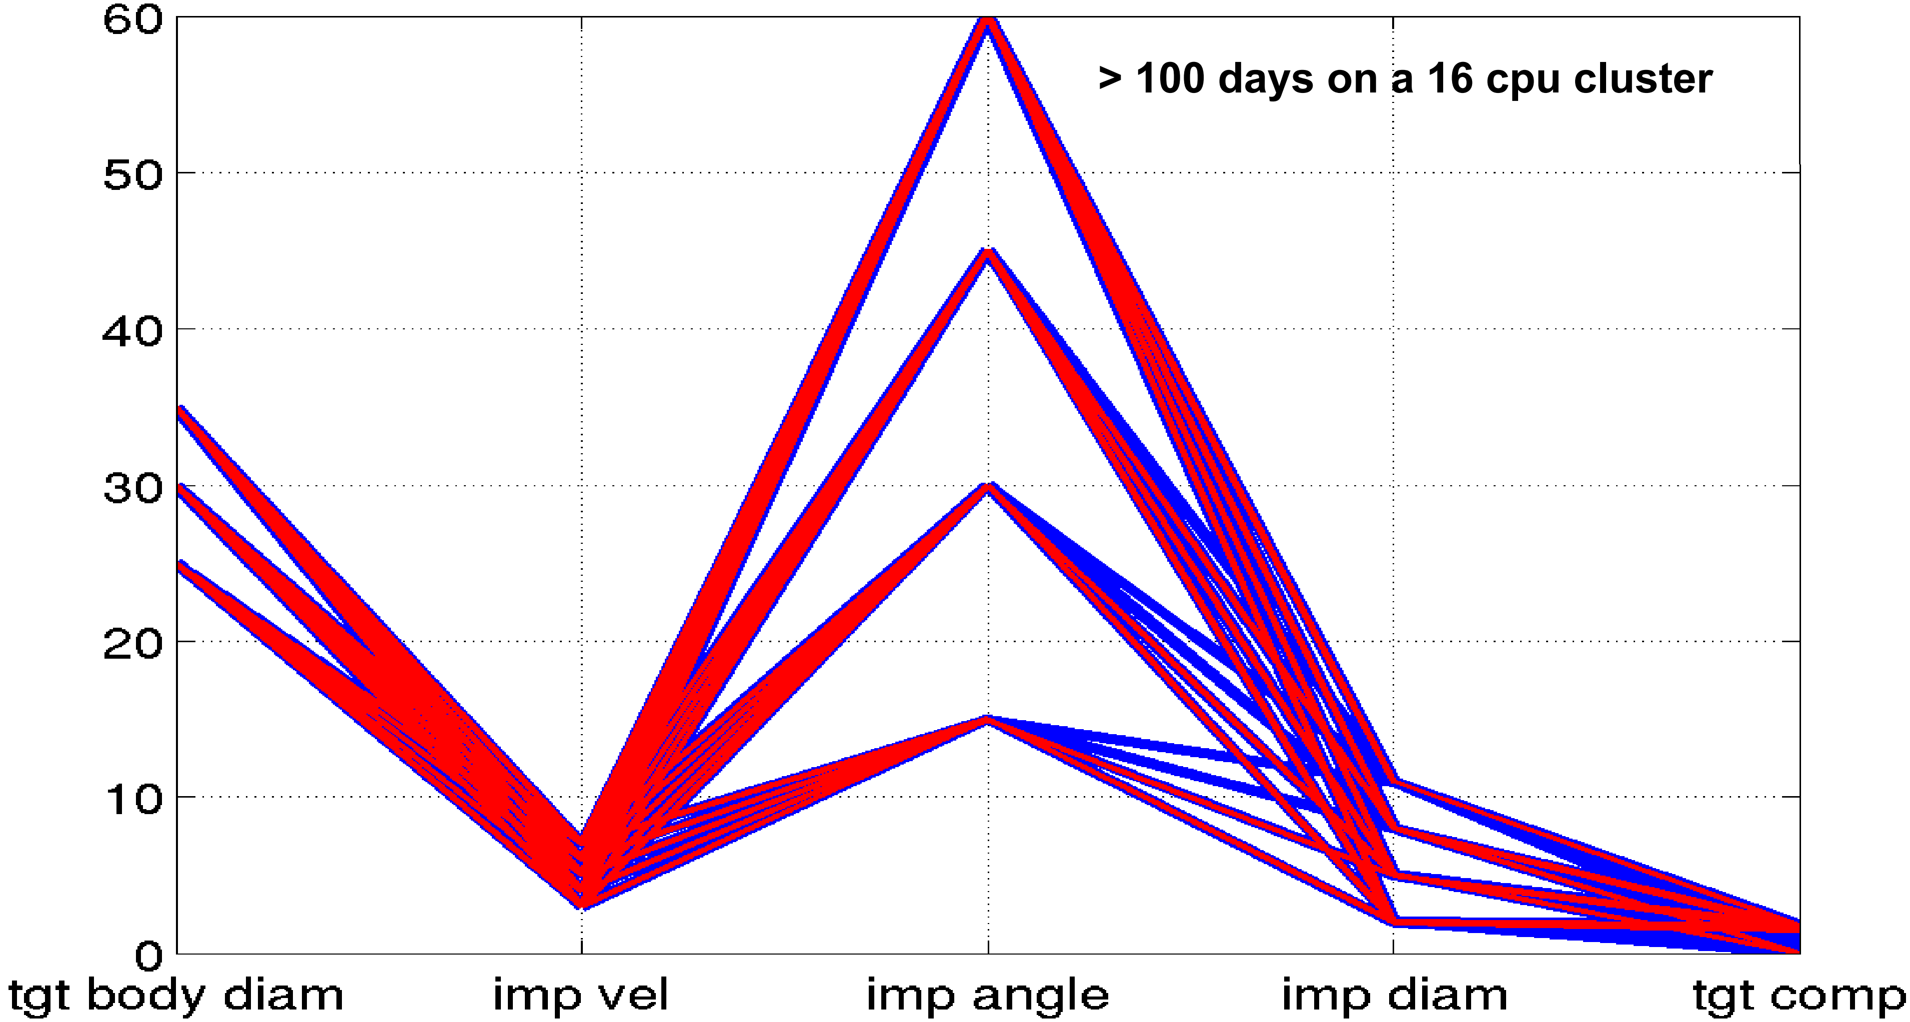

Conclusion

- Directed exploration shows promise for efficiently understanding the behavior of complex systems modeled by numerical simulations.
  - Use simulator as an oracle to sequentially generate labeled training data.
  - Learn predictive model from currently available training data.
  - Use active learning to choose which simulation trials to run next.
  - Get new labeled example(s) and repeat.

- Performance was systematically evaluated over synthetic oracles.

- Active learning yielded significant improvement over passive learning.
  - In some cases, 3x reduction in number of trials to reach 90% PD at PFA = 0.05

- Exploiting the time-savings
  - Same final result can be obtained with fewer simulation trials.
  - More simulation trials can be conducted in a given amount of time.
  - Higher-fidelity (e.g., finer spatio-temporal resolution) simulation trials can be used.
  - Concentrate trials at the region between interesting and non-interesting regions.

- Initial Set of Grid Runs and Refinement Completed

Publications and Reports

- **A. Holub, M.C. Burl, P. Perona, "Entropy-based Active Learning for Object Recognition", CVPR Workshop on Online Learning for Classification, (Jun 2008) – *Best Paper Award***

- **E. Wang, M.C. Burl, "Learning Simplified Predictive Models of Complex Dynamical Systems", NASA Conference on Intelligent Data Understanding (CIDU), (Sep 2008)**

- M.C. Burl, E. Wang, "The Expected Information Gain from a Single Training Point", Tech Memo, B-0130, (Oct 07, 2008, Revised Dec 12, 2008)

- M.C. Burl and E. Wang, "Surrogate Oracles for Active Learning", Tech Memo, B-0132, (Feb 2009)

- **M.C. Burl, E. Wang, "Directed Exploration of Complex Dynamical Systems", Int. Conf. on Machine Learning, (ICML), (Jun 2009)**

Future Work

- Experiments with Simulator in Loop
- Hill-climbing for point selection
- Kernel contraction
  - Shrink kernel bandwidth parameter as more data is acquired
- Maximize throughput on computer cluster
  - choose new point while other runs in-progress
- Scalability of MIP

Acknowledgements

D. Nesvornoy, D. Durda
A. Holub, P. Perona, D. Decoste, D. Mazzoni, L. Scharenbroich
Erika C. Vote SURF Fellowship, The Housner Fund
NASA Applied Information Systems Research Program, J. Bedekamp