



Proactively Managing Aviation-System Safety Risk

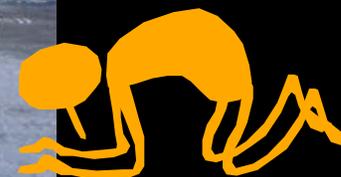
Recent advances in mining aviation safety data

ARMD Technical Seminar
NASA Headquarters
September 22, 2006

Dr. Irving C. Statler
Irving.C.Statler@nasa.gov

Dr. Ashok N. Srivastava
Ashok.N.Srivastava@nasa.gov

The Forensic (Historic) Approach to Accident Prevention



VS....



... a More Prognostic Approach



Leads to
Decisions

*before an Accident
or Incident Occurs*



The Voluntary Aviation Safety Information-sharing Process (VASIP)



Provides an acceptable process for U.S. airlines to comply with current information-sharing regulations.

- FAA created the Voluntary Safety Information Sharing (VSIS) Aviation Rulemaking Committee (ARC) to:
 - Serve as a forum for interaction between industry representatives and FAA.
 - Provide advice to FAA on safety-information sharing.
 - Prepare recommendations for rulemaking.
- NASA built and exercised distributed national archives.



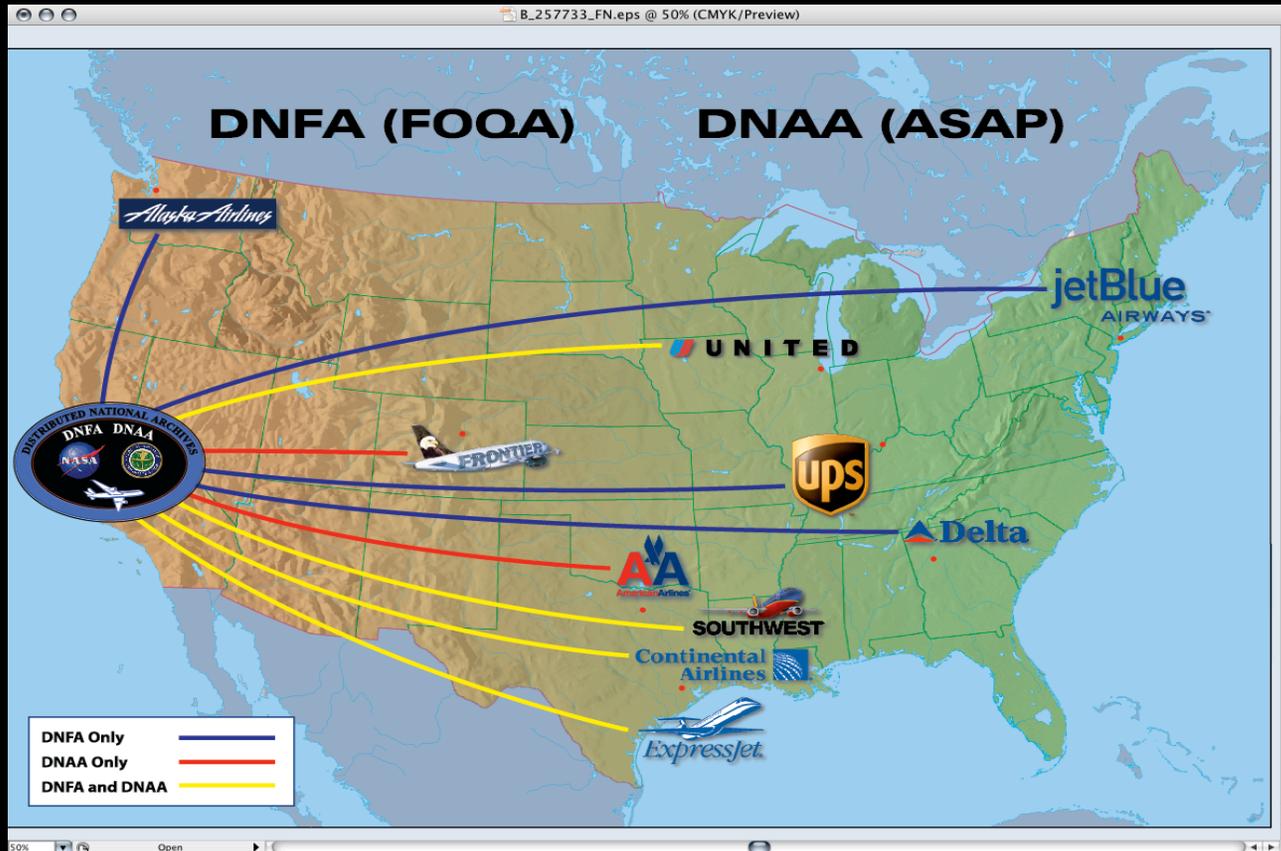
Status of National Archives as of end of FY'06



- DNFA - Distributed National Flight Operational Quality Assurance (FOQA) Archive is fully operational
 - Total of **700,000 flights**. Adding approx **100,000 flights per month**.
 - From **8** participating **airlines** (passenger & cargo). Data from **9 aircraft fleets** (22 subfleets)
- DNAA - Distributed National Aviation Safety Action Program (ASAP) Archive is fully operational
 - Database is still small; only **three** airlines contributing data **during most of '06**.
 - Total of almost **10,000 reports**. Adding approx **1,000 events per month**.
 - From **6** participating **airlines** (passenger only). Data from **4 aircraft fleets**
- Utilized to support a VSIS ARC Working Group in a study of Ground Proximity Warning System alerts nationally.



Distributed National Archives





New Technical Challenges



Enlarge

- Increase the *number of air carriers participating in the DNFA and the DNAA* by up to ten new Part 121 airlines within the next two years.

Expand

- Add *new sources of information* to augment insight into systemic issues in the air transportation system; e.g., add data sources from *maintenance, ATC, CAST, NTSB, the available published literature.*

Enhance

- *Uncover system-level vulnerabilities* in any of the **heterogeneous** (continuous digital, discrete digital, analog, and textual) aviation-safety *data sources*
- *Acquire and fuse information* extracted from very *distributed, multiple, heterogeneous* data sources.



Recent Advances in Mining Aviation Safety Data

Ashok N. Srivastava, Ph.D.

ashok@email.arc.nasa.gov

Leader: Intelligent Data
Understanding Group
NASA Ames Research Center



The Data Mining Team



Lead Investigators

Irv Statler, Ph.D.

Ashok N. Srivastava, Ph.D.

Team Members

Suratna Budalakoti

Pat Castle

Tom Ferryman, Ph.D.

Dave Iverson

Amy Mai

Rodney Martin, Ph.D.

Bryan Matthews

Dawn McIntosh

Matthew Otey, Ph.D.

Nikunj Oza, Ph.D.

Loren Rosenthal

Mark Schwabacher, Ph.D.

Eugene Turkov

Richard Watson

David Wolpert, Ph.D.

Team Members are NASA Employees, Contractors, and Students.



Key Programs



- Aeronautical Research Mission Directorate - ASIAS
- NASA Engineering and Safety Center
- Exploration Systems Mission Directorate - Exploration Technology Development Program, ISHM Project
- Shuttle Program - Wing Leading Edge Impact Detection
- Science Mission Directorate - AISRP

All schematic diagrams and pictures in this presentation are publicly available on the Internet.



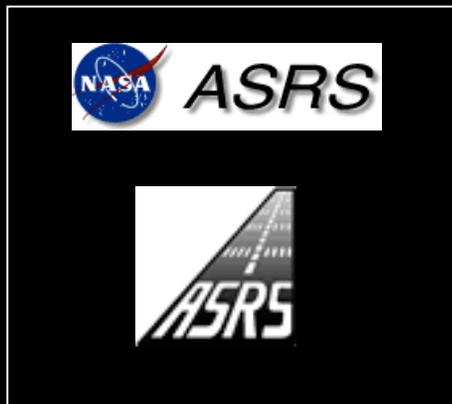
Outline of Talk



Categorizing and detecting anomalies described in safety documents

Detecting anomalies in cockpit switching sequences

Detecting Shuttle wing heating and engine anomalies





Aviation Safety Reporting System (ASRS) Reports



- ASRS collects, analyzes, and responds to voluntarily submitted aviation safety incident reports to improve aviation safety.
- Authors are pilots, crew, and others involved in aviation operations.
- Each airline has an Aviation Safety Action Program (ASAP).
- Approximately **135,000 ASRS reports** with **60 overlapping anomaly categories**.
- Reports can be 0.5-4 pages long.
- Reports accrue at about 3,000 per month.



ASRS Report Excerpt



JUST PRIOR TO TOUCHDOWN, LAX **TWR** TOLD US TO GO AROUND BECAUSE OF THE **ACFT** IN FRONT OF US. BOTH THE **COPLT** AND I, HOWEVER, UNDERSTOOD TWR TO SAY, '**CLRED** TO LAND, **ACFT** ON THE **RWY**.' SINCE THE **ACFT** IN FRONT OF US WAS **CLR** OF THE **RWY** AND WE BOTH **MISUNDERSTOOD TWR'S** RADIO CALL AND CONSIDERED IT AN ADVISORY, WE LANDED...



Classification Task



- Automatically map safety reports into Distributed National ASAP Archive (DNAA) anomaly categories.
- New reports entering the DNAA can then be automatically categorized by the classifier.
- Comparison among Natural Language Processing (NLP), statistical methods, and Mariana, which is based on advanced machine learning techniques.



Automatic Categorization of ASRS Reports

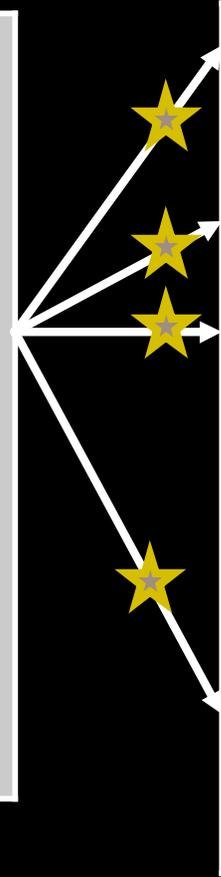


ASRS Report Excerpt

JUST PRIOR TO TOUCHDOWN, LAX TWR TOLD US TO GO AROUND BECAUSE OF THE ACFT IN FRONT OF US. BOTH THE COPLT AND I, HOWEVER, UNDERSTOOD TWR TO SAY, 'CLRED TO LAND, ACFT ON THE RWY.' SINCE THE ACFT IN FRONT OF US WAS CLR OF THE RWY AND WE BOTH MISUNDERSTOOD TWR'S RADIO CALL AND CONSIDERED IT AN ADVISORY, WE LANDED...

Sample of 60 ASRS Anomaly Categories

- Non Adherence to ATC Clearance
- Critical Equipment Problem
- Runway Incursion
- Landing without a Clearance
- Air Space Violation
- Altitude Deviation Overshoot
- Fumes
- Altitude Deviation Undershoot
- Ground Encounter, Less Severe
- ...





Machine Learning Approach



- Convert documents into a vector space representation “Bag of Words” matrix.
- Learn the mappings from documents to categories.
- Typical matrix:
 - 30,000 rows
 - 40,000 dimensions

Frequency
of term in
document

	Term 1	Term 2	Term 3	..
Document 1	0	1	0	4
Document 2	0	3	0	0
...	2	8	1	0



Some Classification Methods



Standard Methods

- Linear Discriminant Analysis
- Logistic Regression
- Neural Networks
- Decision Trees

- Methods may make strong assumptions and may not work well on high dimensional data.
- Some nonlinear methods converge to a local error minimum.

Kernel Methods

- Gaussian Process Regression
- Support Vector Machines

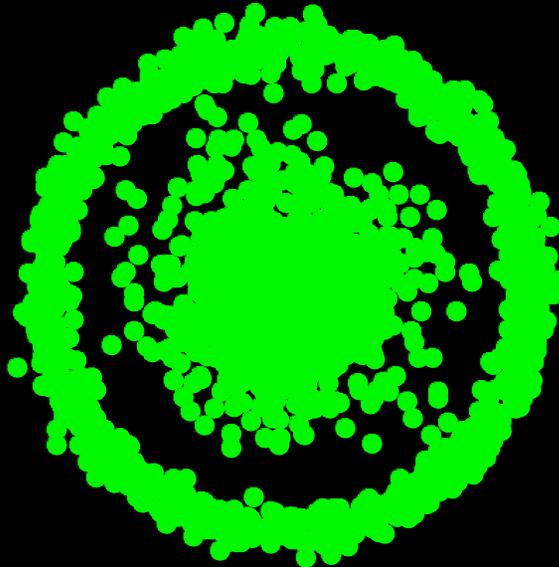
These methods can provide high accuracy on high dimensional data and converge to a global error minimum.



How Kernel Methods Work



Example
Problem:
separate the
inner cloud of
dots from the
outer ring.



Each dot represents a point in two dimensions.
Each region represents a different category.

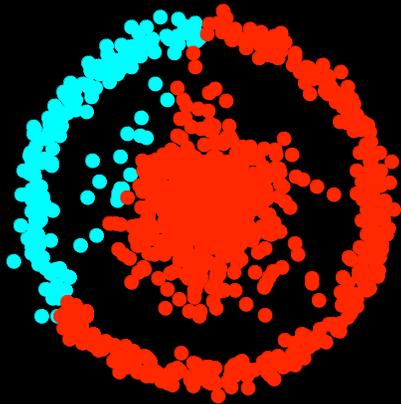
**In real text mining problems, each dot is a document
that lives in a $\sim 40,000$ dimensional space.**



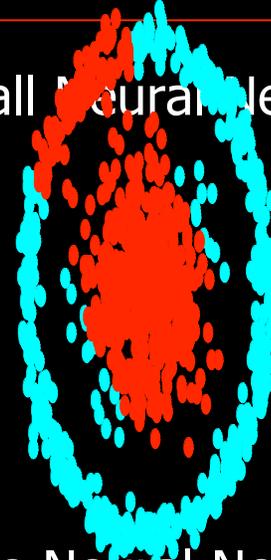
Classifications with Standard Methods



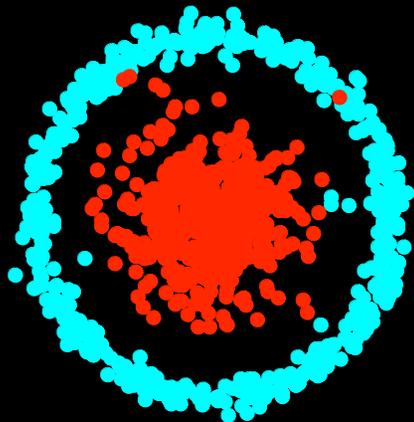
Logistic Regression



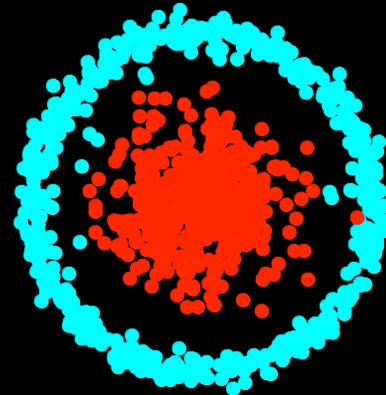
Small Neural Network



Decision Tree



Large Neural Network



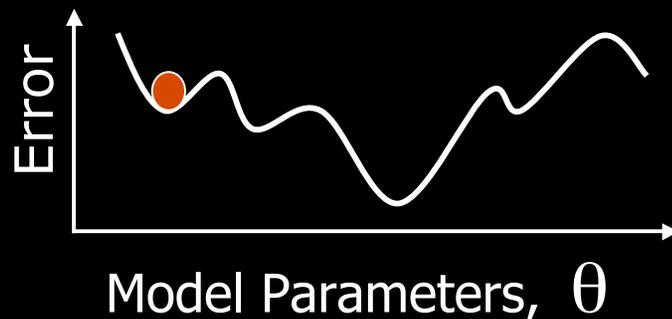


Analysis of Standard Solutions

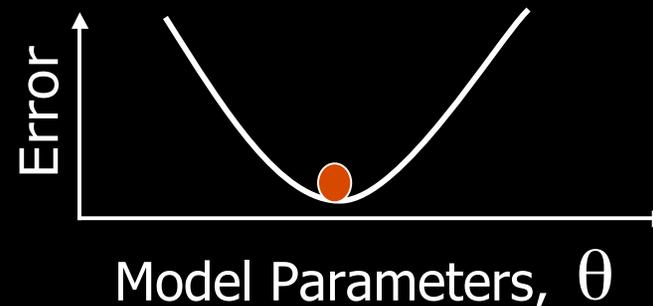


- Linear methods may be restrictive for complex tasks.
- Many nonlinear classifiers converge to local error minima and miss the global minimum.
- Kernel classifiers can find the best model given the data and hyperparameters and converge to the global minimum.

Standard Methods



Kernel Methods

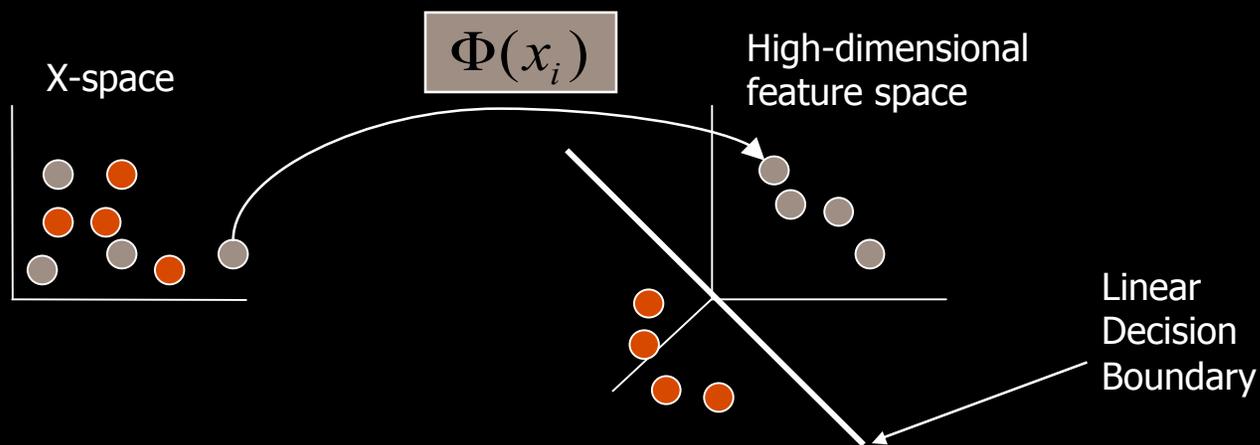




The Support Vector Machine



- Given a set of p -dimensional data $\{x_i\}_{i=1}^N$
- Use a possibly infinite dimensional operator $\Phi(x_i)$ to map the data into a **feature space**.
- Perform linear operations in the feature space.
- Map result back to the original space.
- Can do this operation without explicitly computing $\Phi(x_i)$

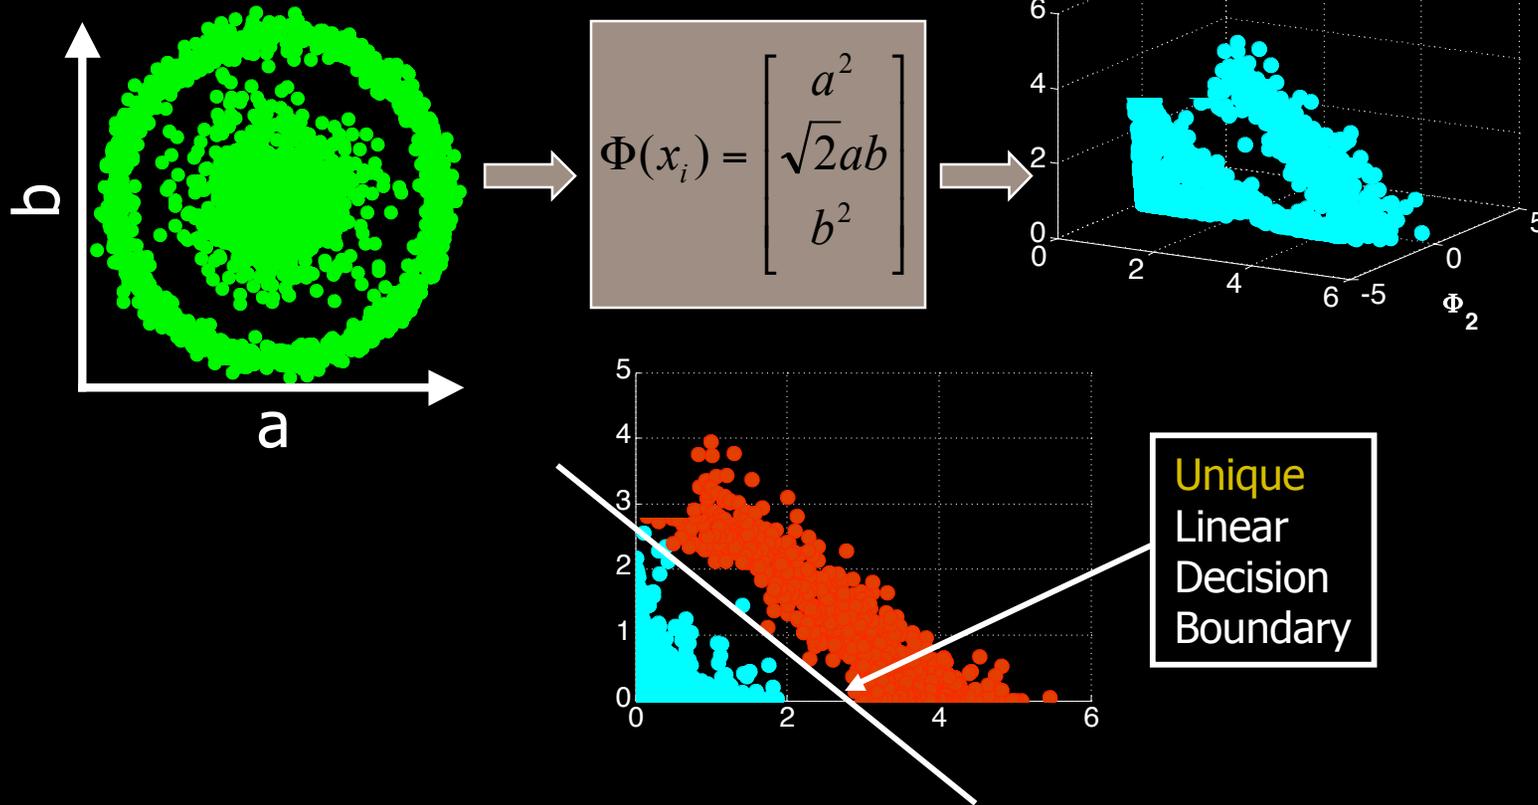




An Example Mapping



Using a kernel function $K(x_i, x_j) = \langle x_i, x_j \rangle^2$, **two**-dimensional data gets mapped to **three** dimensions.





Text Mining with SVMs

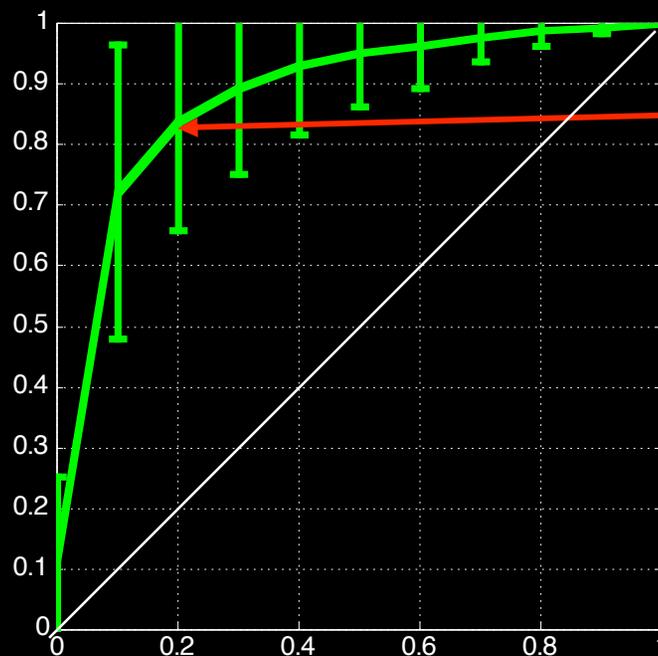


- We built 23 instances of a Support Vector Machine, each tuned to classify ASAP documents into DNAA anomaly categories with advanced noise reduction methods.
- We developed Mariana, an advanced Markov Chain Monte Carlo (MCMC) algorithm to find the best SVM hyperparameters.
- Kernel induces an **infinite dimensional** feature space.

$$K(x_i, x_j) = \Phi^T(x_i)\Phi(x_j) = \exp\left[-\frac{\|x_i - x_j\|^2}{2\sigma^2}\right]$$

Receiver Operator Characteristic (ROC) for Mariana on ASAP Text for 23 DNAA Categories

Study performed on test data to mimic human categorization.



Mean Accuracy

20% False Positive Rate
83% True Positive Rate

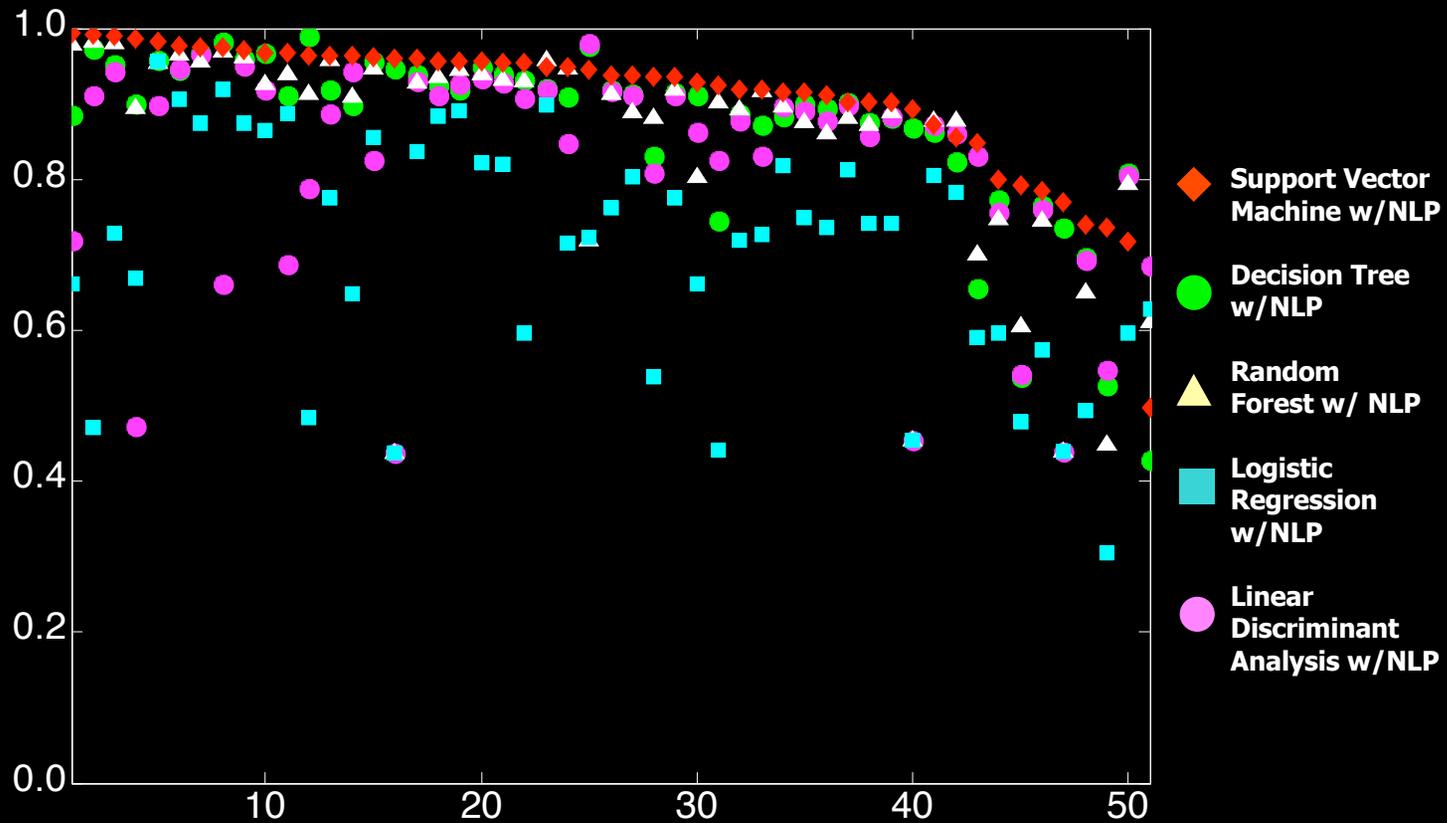
Median Accuracy

20% False Positive Rate
90% True Positive Rate

Area Under Curve

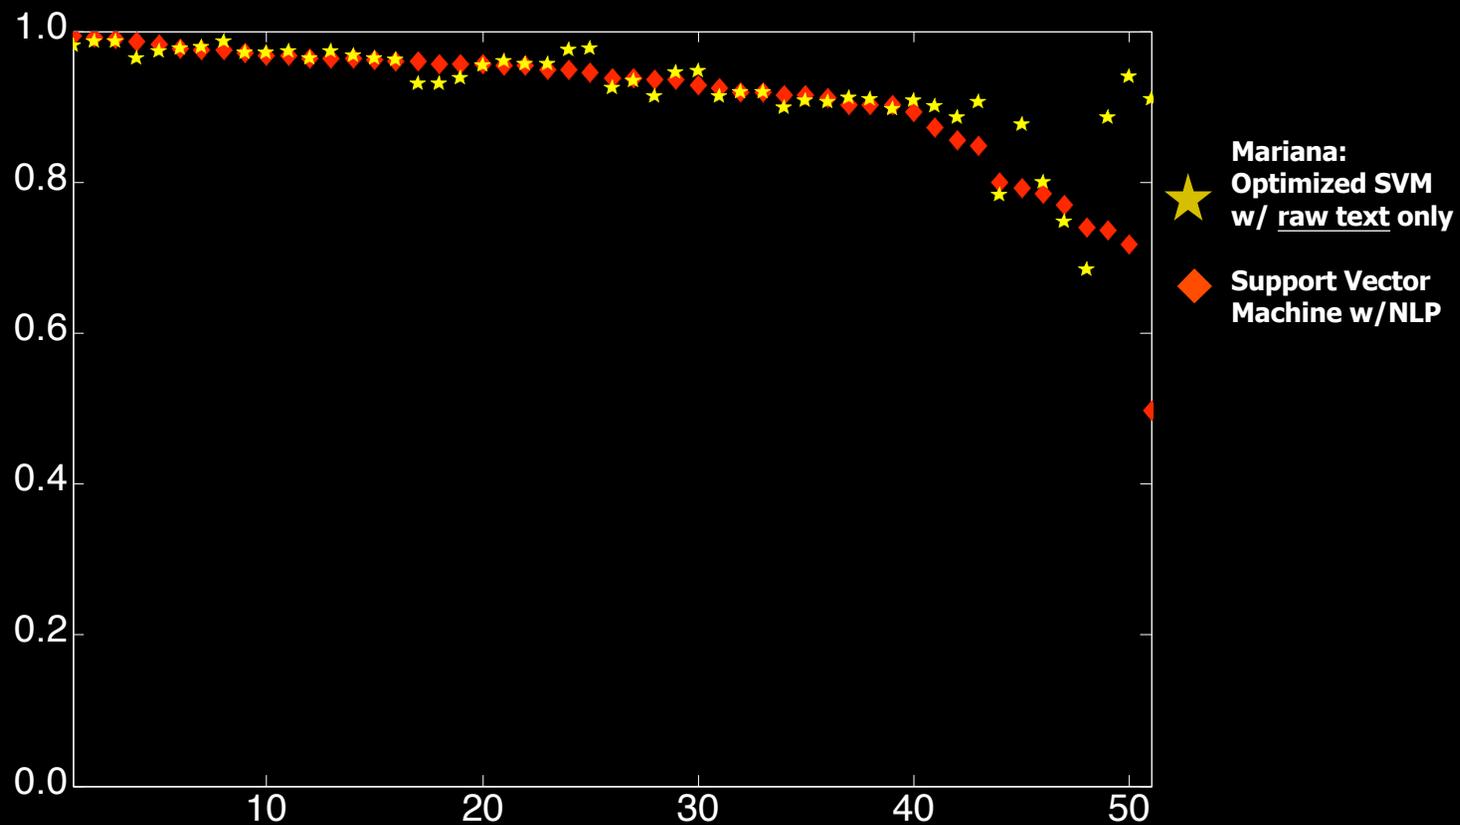
Area = 0.9 (max 1.0)
Overall measure of
classifier performance.

Comparison of un-optimized SVM and Standard Methods



SVM beats other methods 43 out of 51 times.

Comparison of Mariana with Raw Text and SVM with Raw Text + NLP





Our Innovations



- Mariana, an advanced method to search for the best hyperparameters of SVM using Markov Chain Monte Carlo approach.
- Mariana **performs as well as or better** than the SVM built using NLP techniques **without the overhead**.
- Our methods for term selection and noise reduction reduce false positive rates by as much as 30%.



Searching for Recurring Anomalies



*Enabling discovery of
anomalous trends in
complex aerospace
systems*

*Research sponsored by:
NASA Engineering and
Safety Center*

NASA Mishap and Anomaly Information Systems (MAIS) - Mozilla Firefox

File Edit View Go Bookmarks Tools Help

http://jerusington.aen.nasa.gov:8090/mais-web/

Customize Links Free Hotmail Like Music - Try AOL! Windows Media Windows

Engineering For Complex Systems

Home Taxonomic Search Trend Charting Report Browsing Mishap Classification

User Account Request

Netmark XDB Search Context Content DB NESC CARS Submit Links:

Text Mining

ECS Mishap and Anomaly Information System

Features

- Taxonomy Analysis
- What / Why Population Report Graph
- Mishap Reports
- NETMARK Search

© 2006 AEN In6Lab Group, NASA Ames Research Center

Transferring data from jerusington.aen.nasa.gov...



Searching for Recurring Anomalies



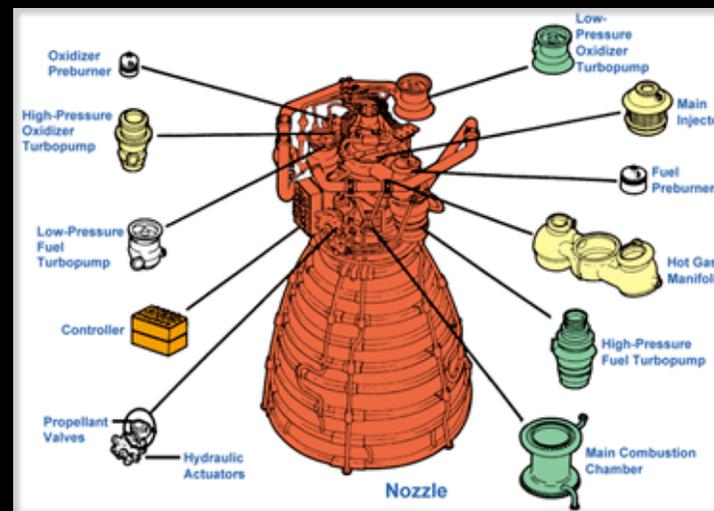
- These reports *do not* have an anomaly category associated with them.
- Potentially several hundred thousand reports.
- Some systems have been around for decades.
- Enables analysis of trends of anomalies (trending).
- Can't be addressed using standard clustering techniques.
- Our systems use content-based similarity as well as statistical similarity.



NESC Definition of Recurring Anomalies



- Recurrent failures described in text reports.
- Problems that cross traditional system boundaries.
- Problems that have been accepted by repeated waivers.
- Discrepant conditions repeatedly accepted by routine analysis.
- Events with unknown causes.





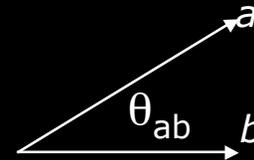
Detecting Recurring Anomalies



1. Calculate cosine similarity between all document vectors.

	Term 1	Term 2	Term 3	Term 4	...
Doc a	3	2	1	5	...
Doc b	0	1	4	1	...
...

$$\cos \theta_{ab} = \frac{\langle a, b \rangle}{\|a\| \cdot \|b\|}$$

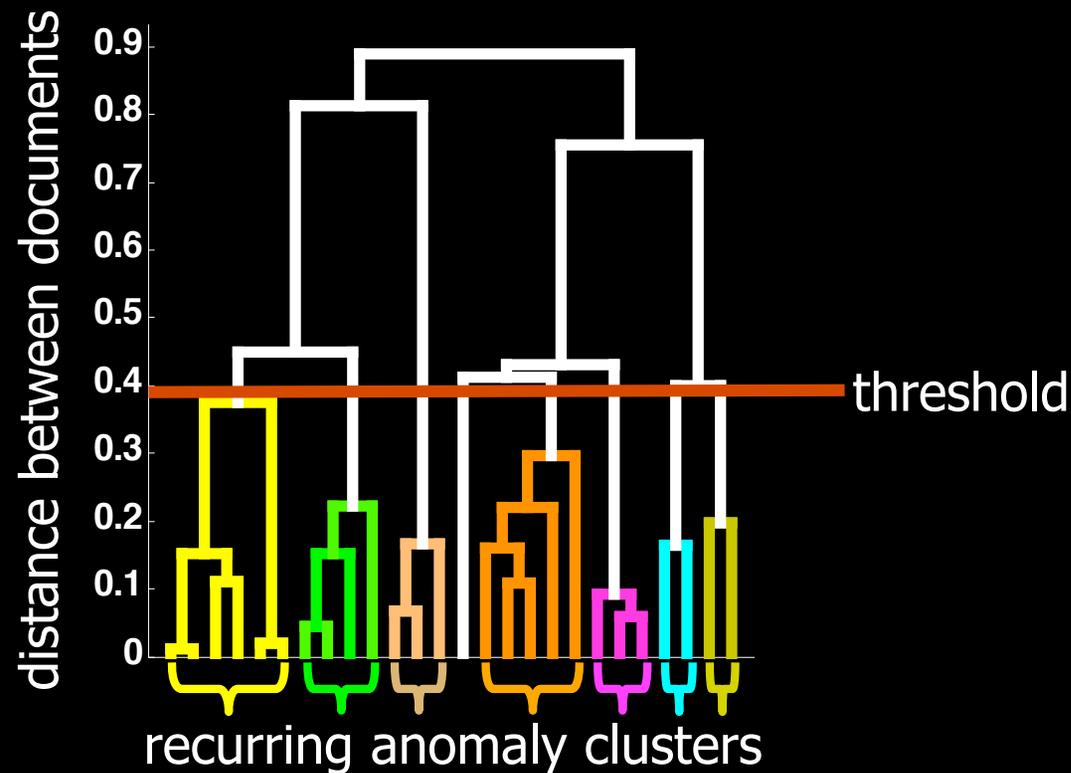




Detecting Recurring Anomalies



2. Apply agglomerative clustering.

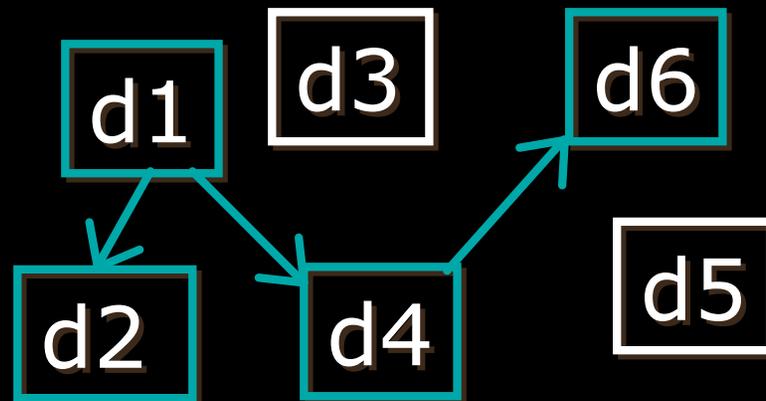




Detecting Recurring Anomalies



3. Identify referenced documents.



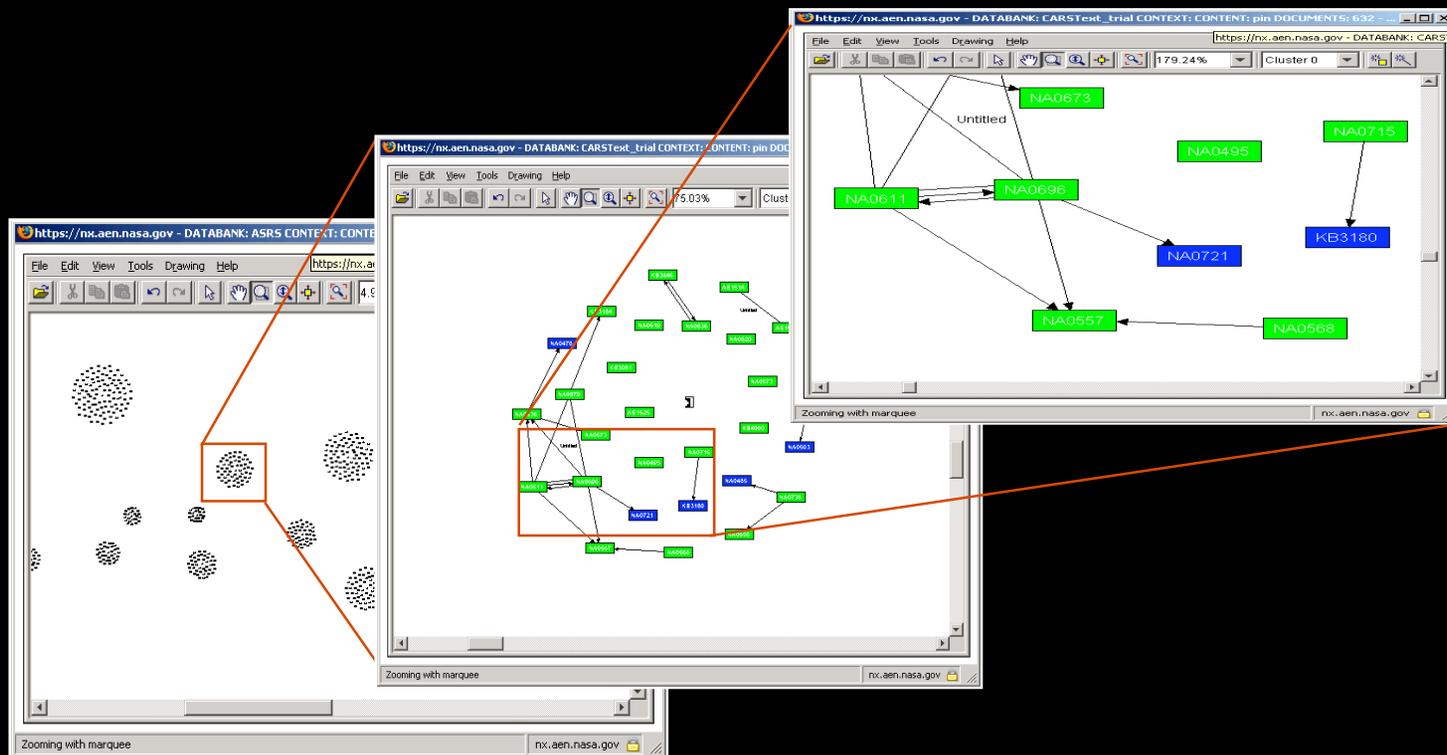
If d1 refers to d2 and d4, and d4 refers to d6, then d1, d2, d4, & d6 are considered a recurring anomaly.



Detecting Recurring Anomalies



4. Identify & visualize possible recurring anomalies.





Testing the Recurring Anomaly Detection System (ReADS)



- Subject experts reviewed a subset of the Shuttle Orbiter Corrective Action Records (CARs) specifically to identify recurring anomalies.
- We extracted 333 reports to test the performance of our system called REcurring Anomaly Detection System (ReADS).
- Of those 333 reports, the experts identified 20 recurring anomalies and ReADS identified 39 recurring anomalies.



Performance of ReADS



On a subset (333) of the Shuttle Orbiter Records:

58% of the records were eliminated as non-recurring anomalies (RAs) by ReADS.

12 exact matches between RAs discovered by experts and RAs discovered by ReADS.

6 previously unidentified RAs discovered by ReADS which were confirmed by experts.

1 record was identified by experts as being part of an RA and was missed by ReADS.

5% of the expert RAs were separated by ReADS into more than one RA.

8% of the ReADS RAs combined two expert RAs into a single RA.



Our Innovations



- Enable analysis of anomaly trends using a combination of content and statistical search methods.
- ReADS is a novel tool designed especially for identifying recurring anomalies across multiple databases.
- Development of robust platform to analyze and visualize recurring anomalies.



Detecting Anomalies in Cockpit Switch Sequences



*Enabling
discovery of
anomalous
switching events*

*Research
sponsored by:
NASA ARMD*





Background



- sequenceMiner analyzes large repositories of discrete sequences and identifies operationally significant anomalies.
- Learns the typically observed switching patterns directly from discrete data streams.
- This method outperforms others in terms of speed, comprehensibility, and stability, and does not require knowledge of Standard Operating Procedures.



Example Sequence Anomaly Detection Problem

Typically Observed Switching Patterns

A B C D A D D A G F Q ...

Example Observed Switching Sequence

A B G F Q C D A D D A ...

Problems: (1) Discover Typically Observed Switching patterns given thousands of flights.

(2) Discover outlying sequences.



Outline of Approach



- sequenceMiner discovers typically observed switching patterns using Multiple Sequence Alignment.
 - Normalized Longest Common Subsequence as a similarity measure
 - Optimized for speed. Analyzes 7400 flights in 6 minutes.
- sequenceMiner discovers:
 - Switches **absent** in an **expected** sequence position.
 - Switches **inserted** in an **unexpected** sequence position.
 - Switches that **are out of order** from what is expected.
- sequenceMiner describes why flights are called anomalous and provides a degree of anomalousness.



Multiple Sequence Alignment (MSA)



- Used in bioinformatics to compare DNA sequences of organisms descended from a common ancestor.
- Can identify mutation inside a sequence by comparing it to other sequences.
- In the context of flights, these mutations are the points where a flight deviated from the norm.



Incorporating Operational Information



- Weighting of Switches
 - Measures its importance to flight.
 - Used during clustering and anomaly detection.
 - Sequences are identified that have more highly weighted switches out of sequence, instead of simply the number of switches out of sequence.
- Ignore order of switches within a one-minute time interval.
 - This step reduces alarms by around 30%.

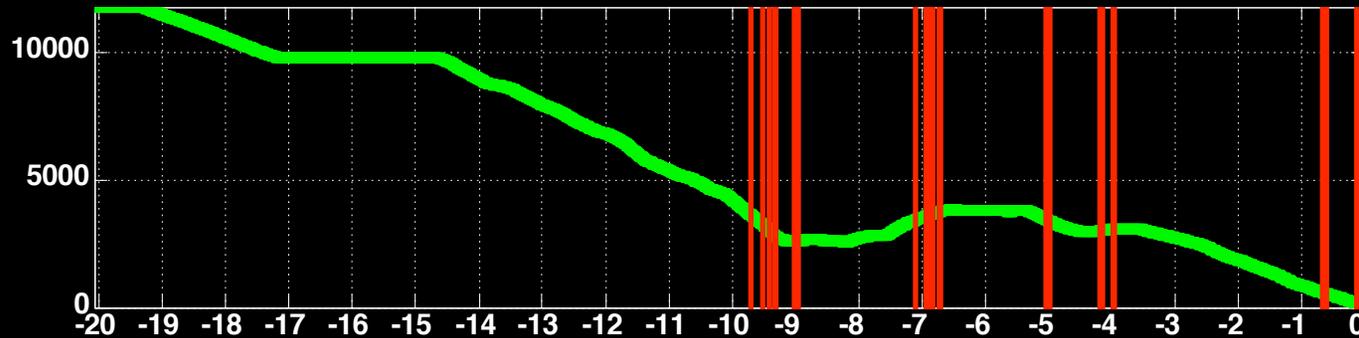


Data and Methodology



- Initial Dataset
 - 7400 flights from a single fleet and airline.
 - Recordings of 1038 primary and secondary binary switches.
 - 111 primary switches were selected from a subset of 2225 flights.
 - Landing phase to a specific destination airport.
- The 13 most anomalous flights identified by sequenceMiner were analyzed by a 747 pilot who was our expert.
 - 5 were judged to be **bad data**.
 - 3 were judged to be **normal**.
 - 5 were judged to be **operationally significant anomalies**.

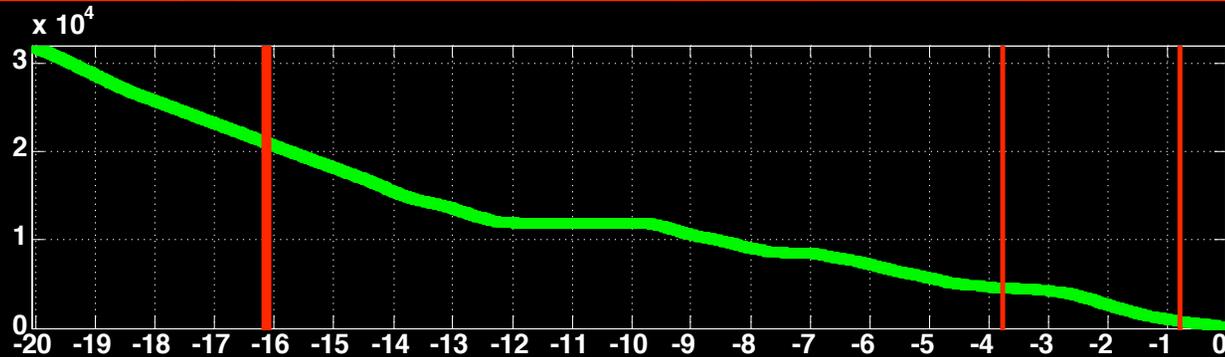
sequenceMiner Discovered Anomalous Presses of the Igniter Switch (Red Bars)



Expert: "Pilot switched igniter on and off at atypical times. Possible engine malfunction."



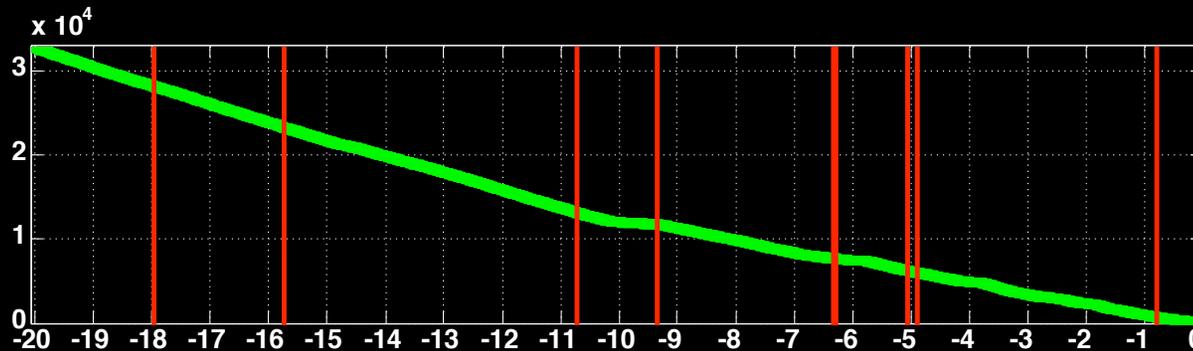
sequenceMiner Discovered Anomalous Engagement of the Autopilot (Red Bars)



Expert: "Auto-pilot used too many times.
Possible case of mode confusion."



sequenceMiner Discovered Anomalous Usage of Speed Brakes (Red Bars)



Expert: "Overuse of speed brakes. Possibly a high energy approach."



Our Innovations



- sequenceMiner is a fast and reliable system to learn typically observed switching patterns from large volumes of discrete data.
- This system outperforms other algorithms in terms of speed and reliability.
- Discovers operationally significant events such as mode-confusion and high-energy approaches.

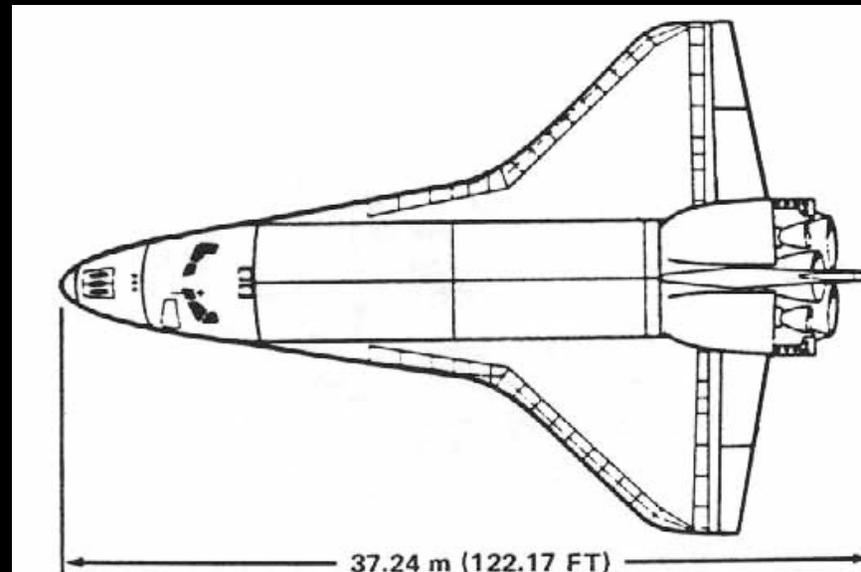


Detecting Anomalies in Shuttle Systems



Enabling discovery of anomalies in continuous data streams

*Research sponsored by:
NASA ESMD ETDP -
ISHM Program*

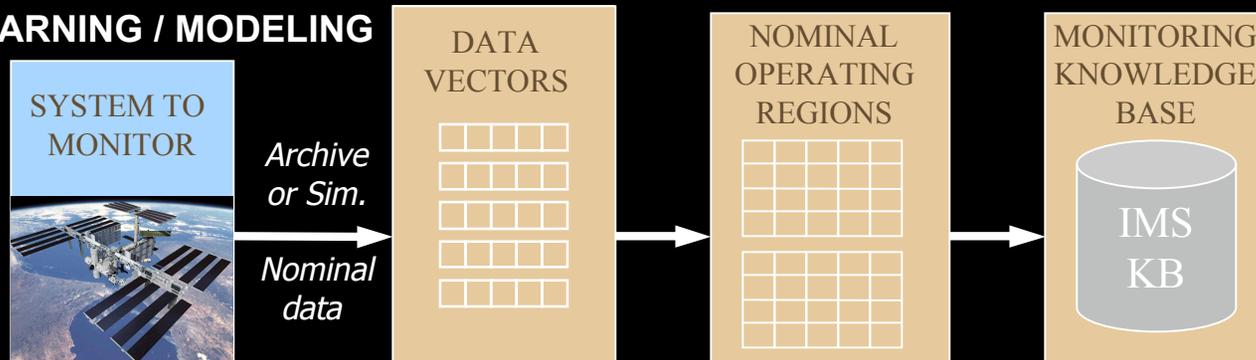




Inductive Monitoring System

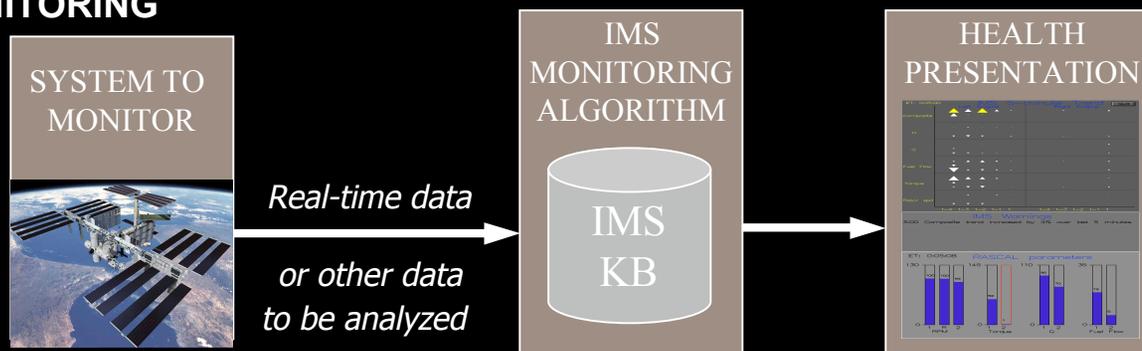


LEARNING / MODELING



IMS learns nominal system behavior from archived or simulated system data, automatically builds a "model" of nominal operations, and stores it in a knowledge base.

MONITORING



IMS real-time monitor & display informs users of degree of deviation from nominal performance. Trend analysis can detect conditions that may indicate incipient failure or required system maintenance.



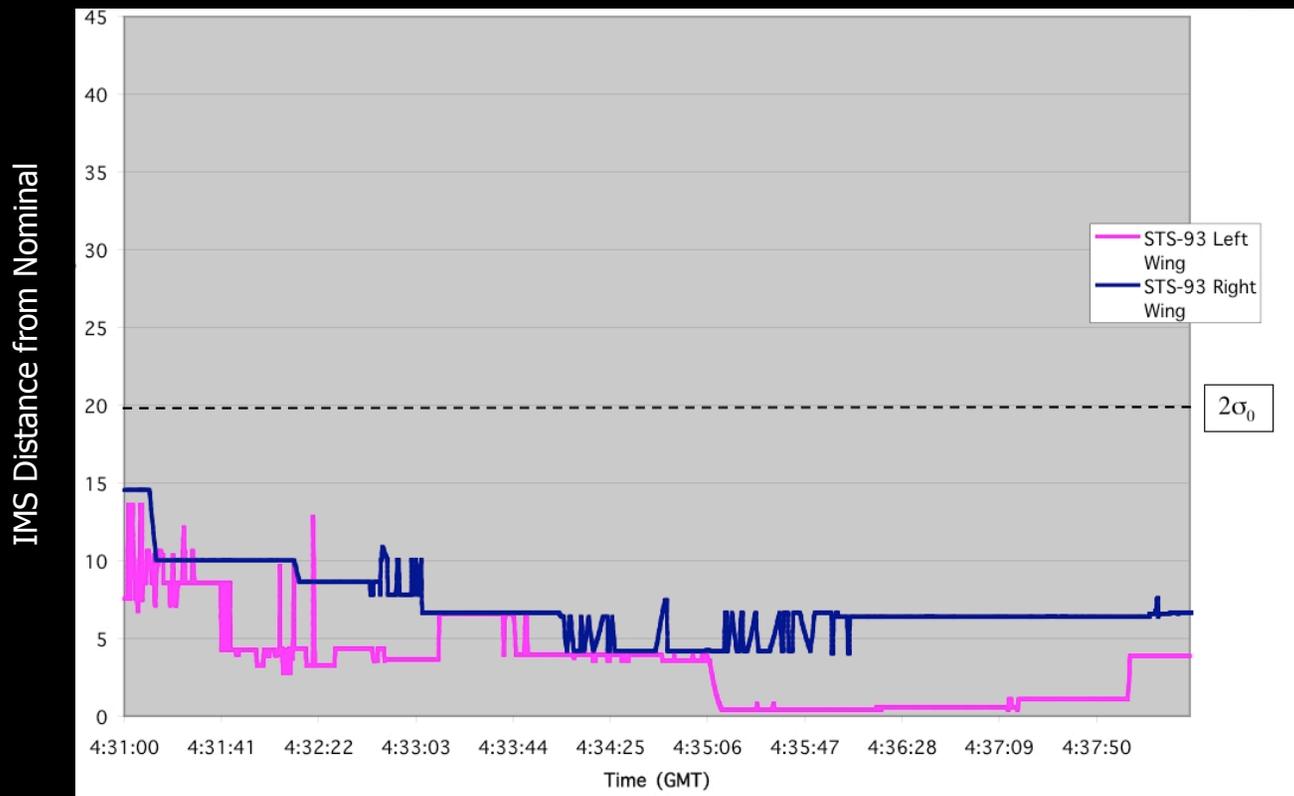
STS-107 Launch Analysis



- The IMS method can help identify subtle but meaningful changes in system behavior.
- A comparison of STS-107 ascent telemetry data to data from previous Columbia flights indicates that there may have been enough information to detect a wing-heating anomaly.

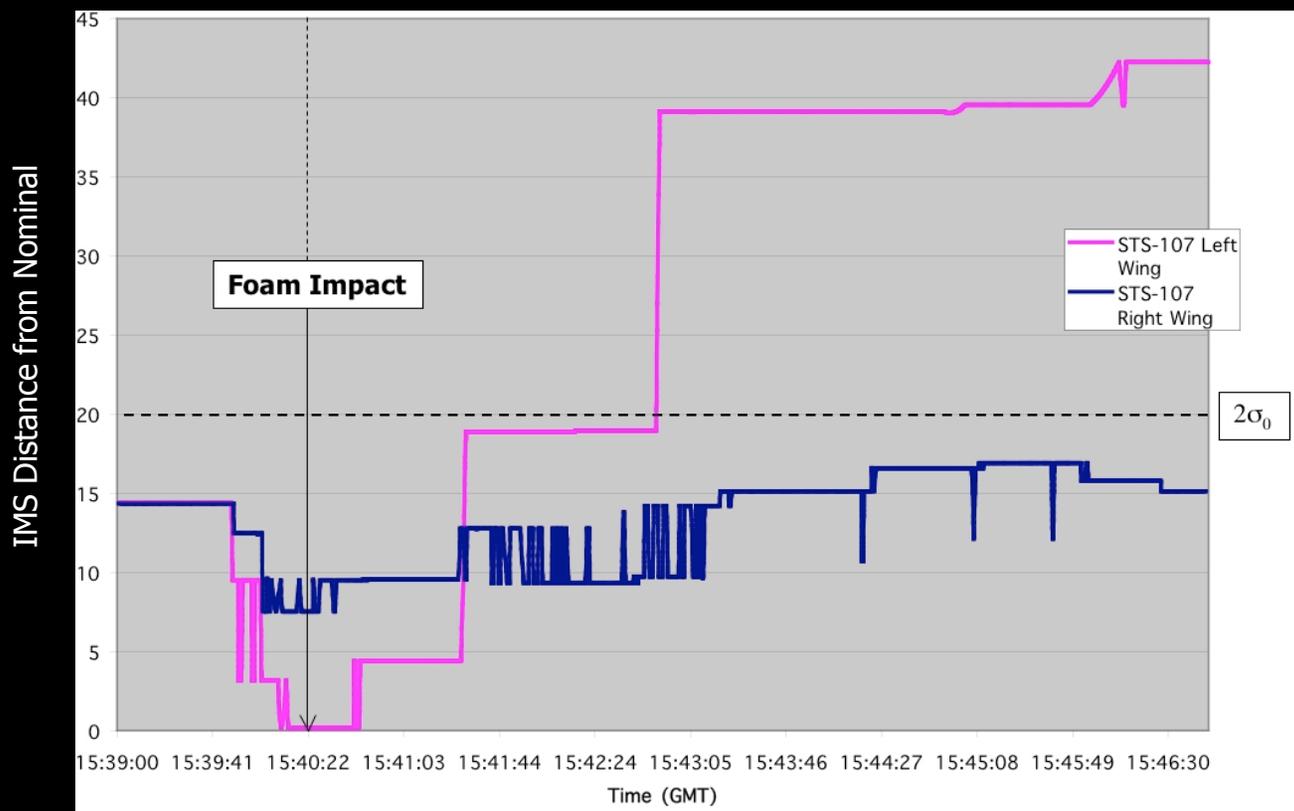


STS-93 Launch IMS Analysis





STS-107 Launch IMS Analysis





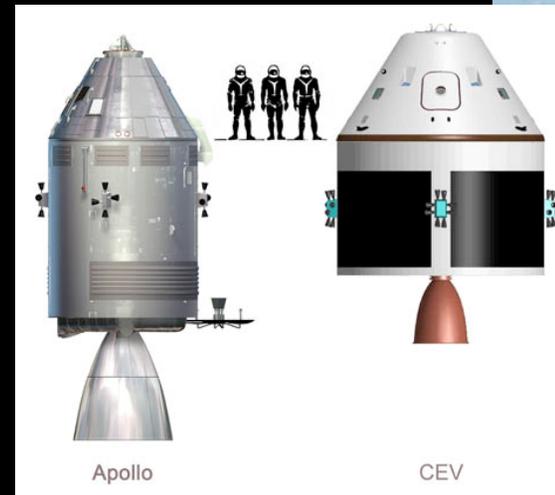
Anomaly Detection for Space Shuttle Main Engine



- Development of algorithms for failure detection and prognostics on SSME.
- Serves as a testbed for CEV/CLV.
- CEV/CLV – Methods implemented on SSME data will be improved, extended, and used for future platforms.



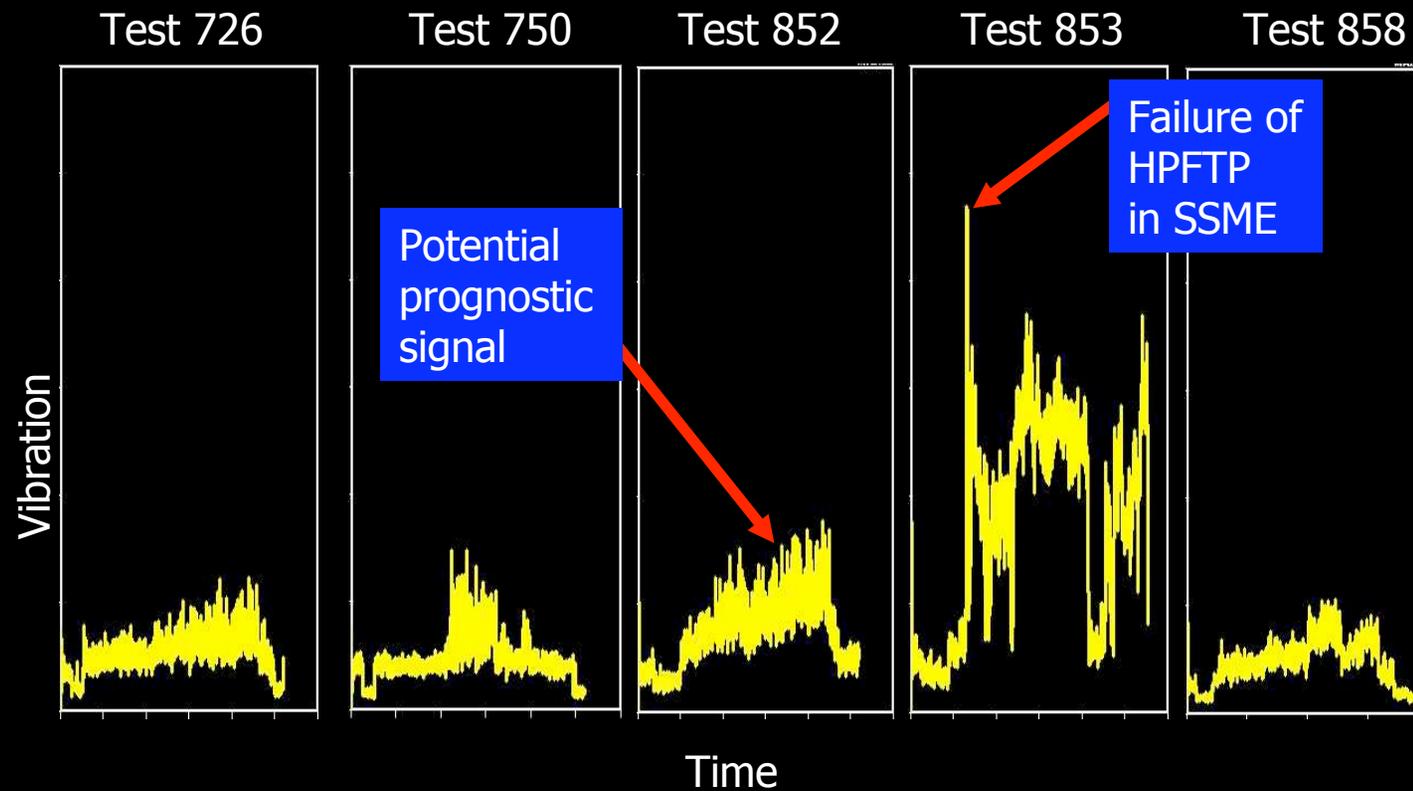
CLV



Apollo

CEV

Example: Potential Prognostic Signal for High Pressure Fuel Turbopump Failure on SSME



Using the Orca anomaly detection system, we discovered that the failure can be detected along with a potential prognostic signal.



Our Innovations



- The IMS system automatically learns a model for nominal behavior to detect system anomalies.
- Orca provides a flexible platform to detect anomalies in massive data sets.
- IMS is used to detect wing impacts in support of STS-121 and STS-115.
- IMS will be deployed on Console at Mission Operations Directorate, JSC.



Conclusions



- Demonstrated transparent mining of discrete, continuous, and textual information to uncover safety anomalies.
- Enabling automated analysis of the Distributed National ASAP and FOQA Archives.
- The methods we discuss provide a comprehensive capability to monitor, detect, and analyze system anomalies.



Future Directions



- Advanced methods to analyze heterogeneous data sets.
- Prognostic and diagnostic methods for aircraft and space systems.
- Potential new book on text mining (Srivastava and Sahami): A collaboration between NASA and Google.
- Text Mining Competition: Classification of ASRS reports, sponsored by NASA.
- Data Mining Consultancy: A team to provide high quality data mining support across the agency.



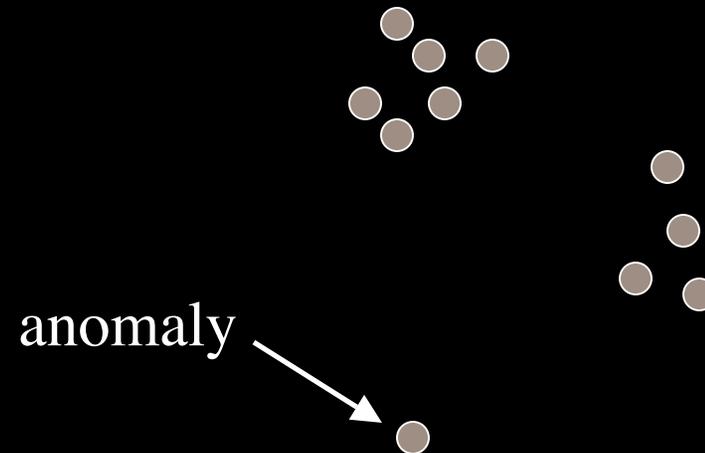
Appendix



Orca

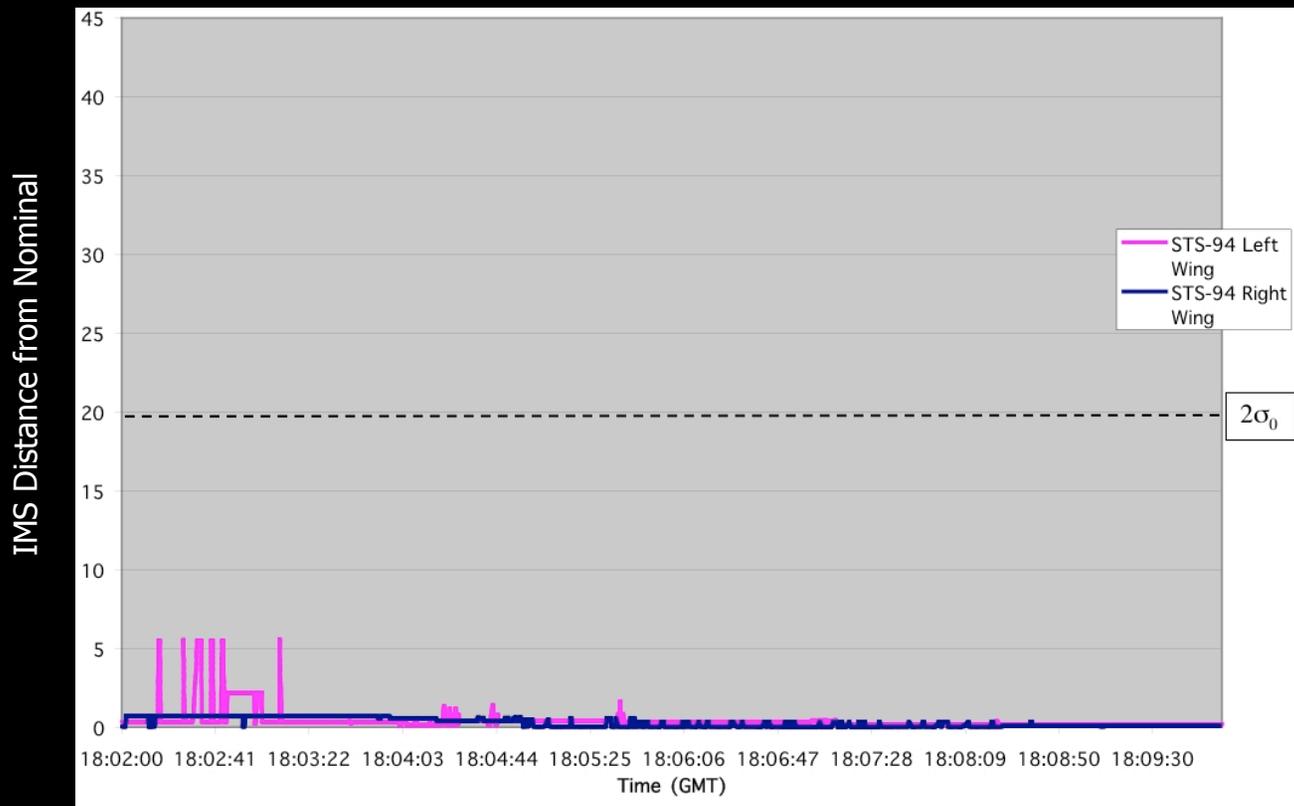


- An unsupervised anomaly detection algorithm
- Uses a nearest-neighbor approach, defining a point to be an anomaly if its nearest neighbors are far away from it
- Uses a novel pruning rule to run in nearly linear time

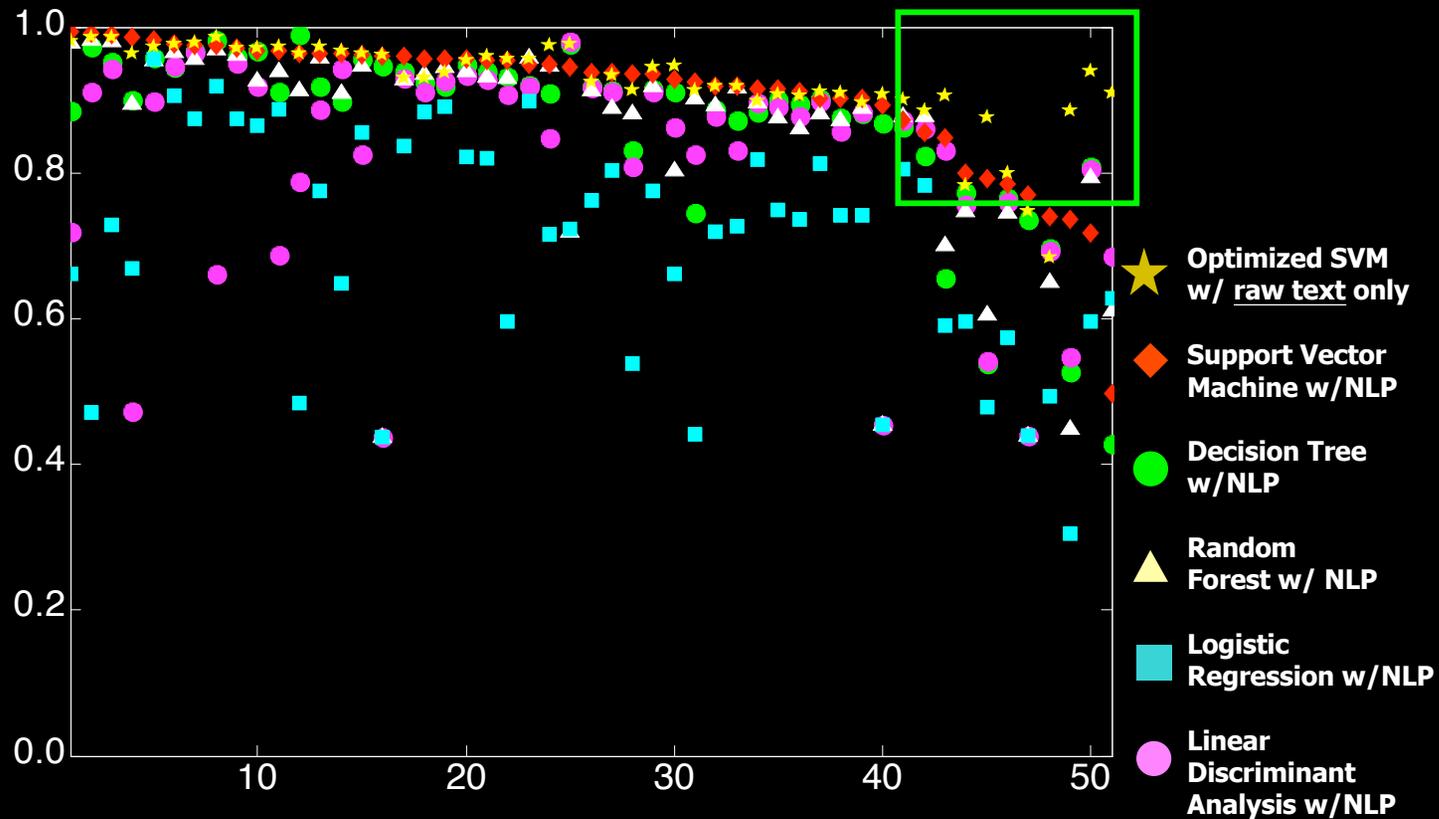




STS-94 Launch IMS Analysis



Optimized SVM vs. Methods using Natural Language Processing



The MCMC Optimized SVM can give substantially better results over other methods (green box), even when processing **raw text**.



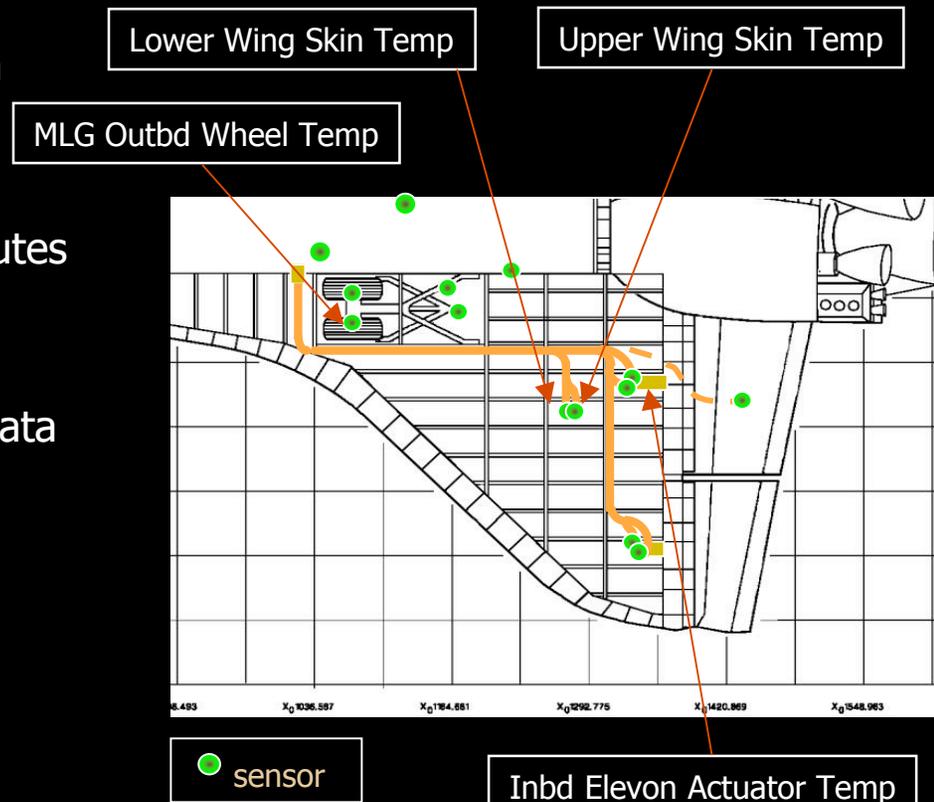
STS-107 Ascent - IMS Analysis



- Data vectors formed from 4 temperature sensors inside the wing
- Data covered first 8 minutes of each flight (Launch to Main Engine Cut Off)
- Trained on telemetered data from 10 previous Columbia flights

Normalization:

- Data expressed as value relative to a reference sensor



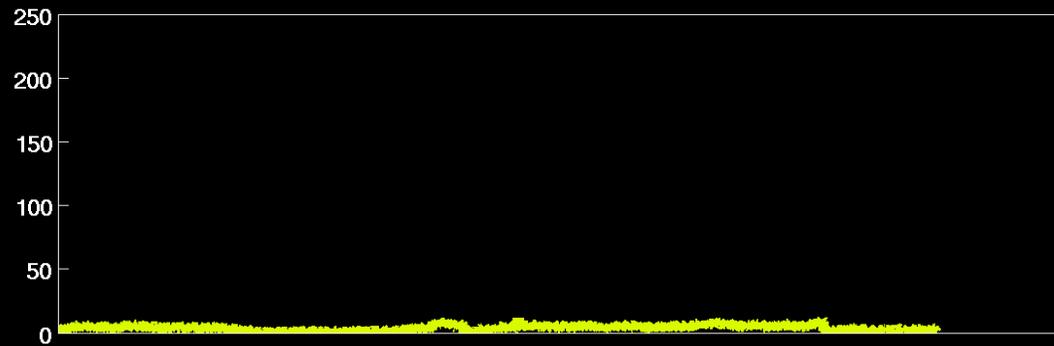


Example: IMS detection of High Pressure Fuel Turbopump Failure in SSME



IMS was trained on nominal data from test 851

Test 852



Test 853

