Machine Learning for Real-Time Transient Detection

Kiri L. Wagstaff, David R. Thompson, and Walid A. Majid
Jet Propulsion Laboratory, California Institute of Technology
Contact: firstname.lastname@jpl.nasa.gov

March 24, 2010

 US Government support acknowledged.

Machine Learning & Fast Transients

- Machine Learning: robust analysis inferred from data
  - Automatically choose tests and detection thresholds
- Detect short-duration radio events: major SKA goal
  - Pulsars, RRATs, merging neutron stars, evaporating black holes, gravitational wave events, etc.
  - Fast "trigger" decision to save buffer for offline analysis
    - atically lume for next stage of analysis

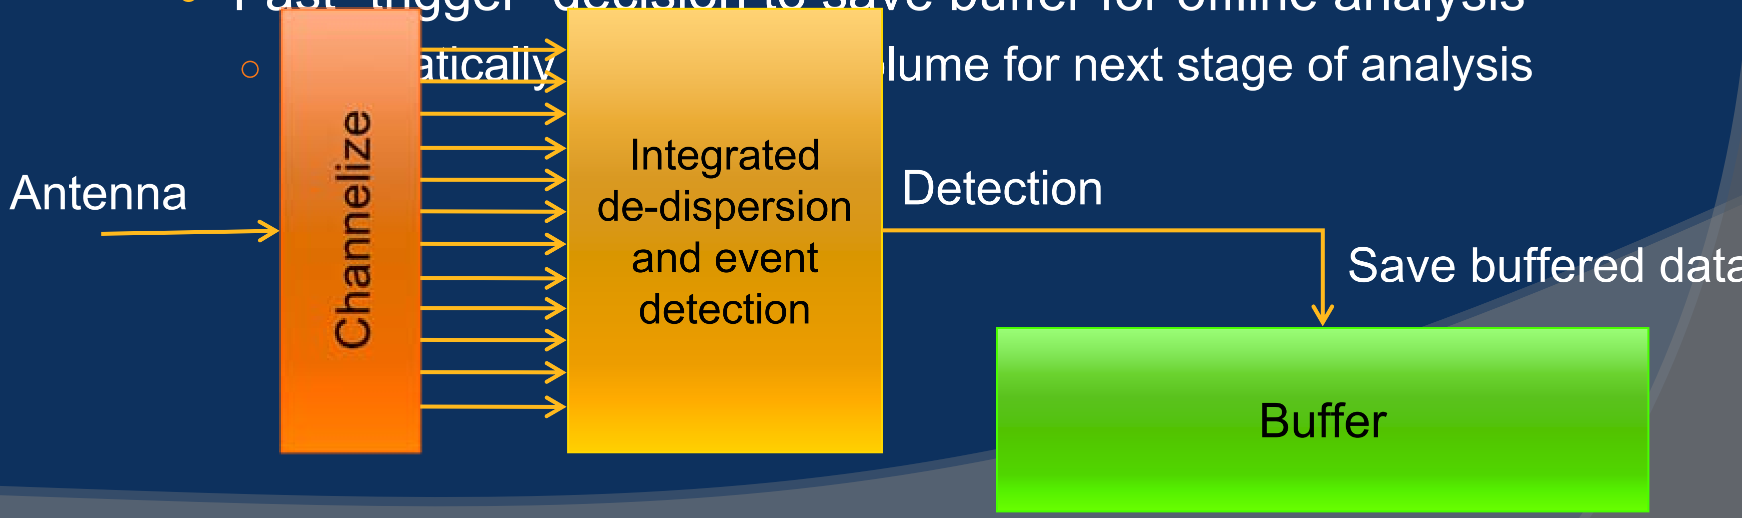

Signal Processing

- Dispersion
  - Delays caused by intervening free electrons; stronger for lower frequencies
- Incoherent de-dispersion (assuming given DM)
  - Sum pulse power across channels according to expected dispersion
- Event detection
  - Is summed power > θ?

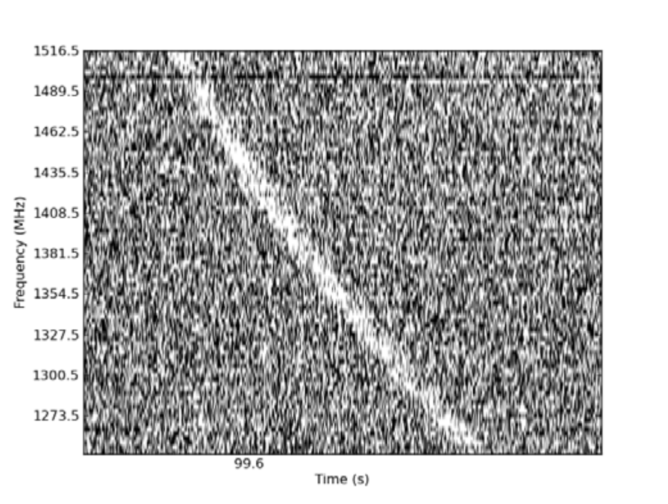

Pulsar J0742-2822, DM 73.78
Parkes Beam 5 [Edwards et al., 2001]

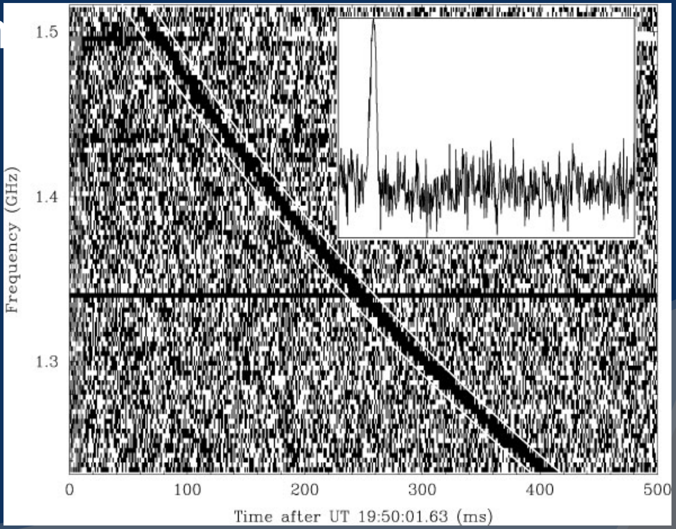

4.6-ms pulse at DM 375 (August 21, 2001)
[Lorimer, 2007]

De-dispersion & Event Detection

- Standard: full DM search, then event detection
  - Constant or linearly increasing spacing
  - All signals “cost” the same
- Machine Learning: learn which DMs to test for a given signal
  - Coarse-to-fine DM search
  - Flexible: adapts to data properties

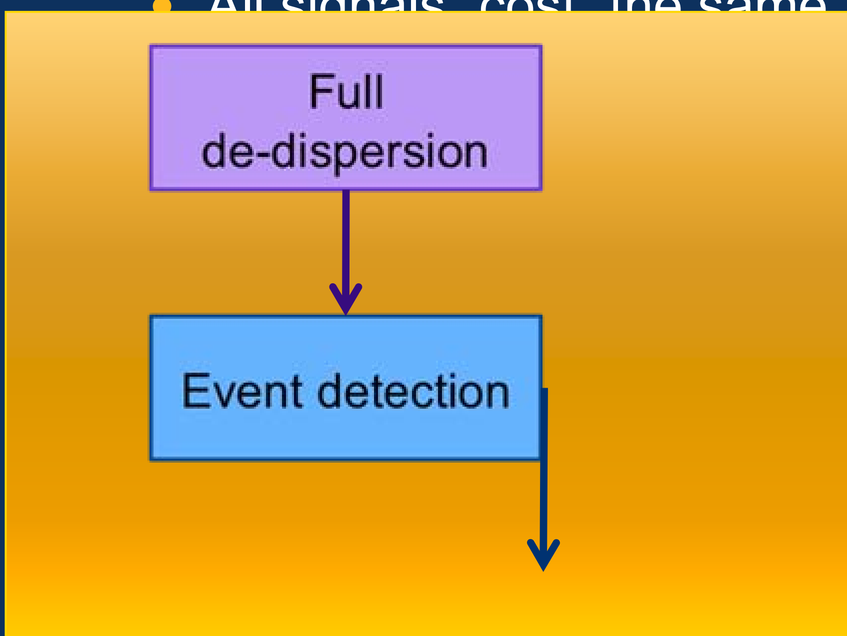

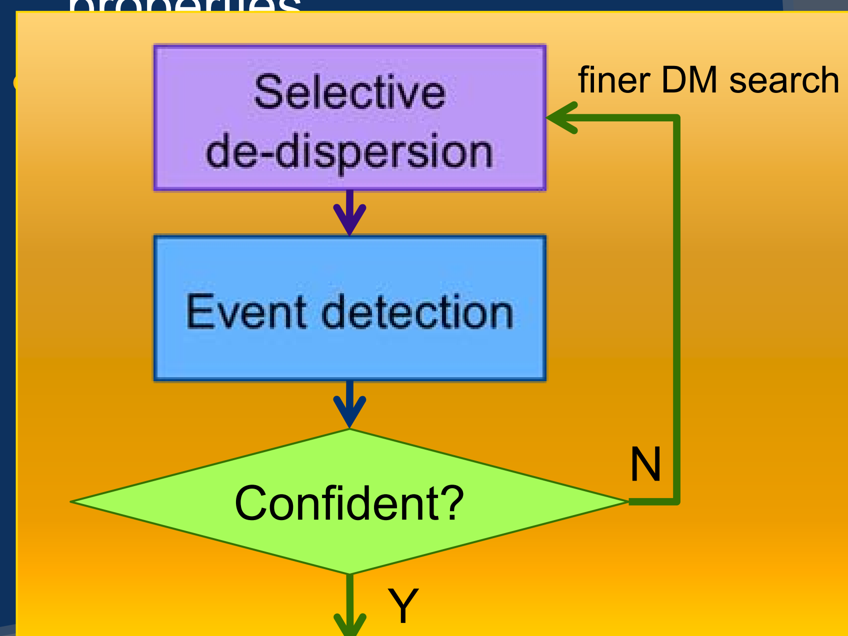

Decision Tree

- Input:
  - Observations, labeled as pulse or non-pulse
  - Cost of missed detection, false detection
  - Capacity: Max. number of parallel de-dispersions conducted (applies to each tier of the tree); constrain HW cost & power

Choice of DMs, thresholds θ tuned to data

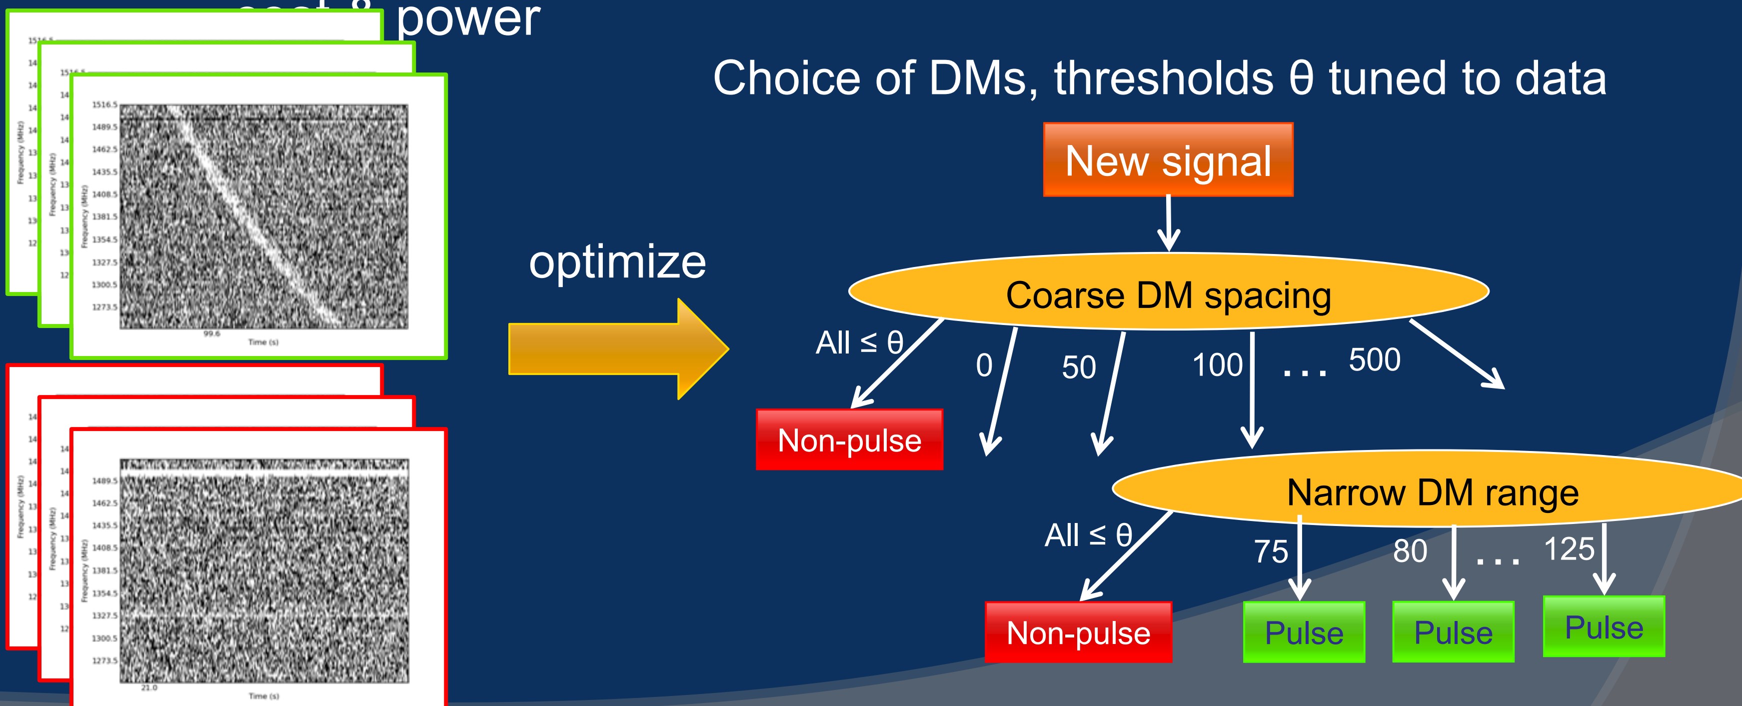

Adaptive Priors

- What if training data doesn't match new data?

Radio Sky from Green Bank, WV

NRAO/AUI/NSF

Galactic Center

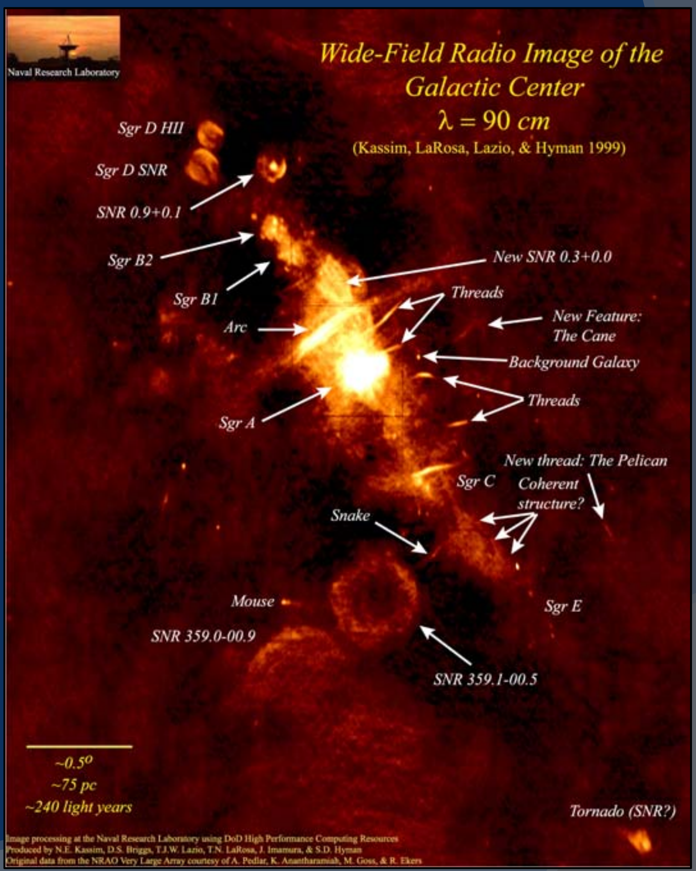

N. E. Kassim, D. S. Briggs, T. J. W. Lazio, T. N. LaRosa, J. Imamura (NRL/RSD)

Adaptive Priors

- What if training data doesn't match new data?
  - Modify old training data according to new target
    - Increase weight of similar signals
    - Decrease weight of dissimilar signals

Probability of pulse
with given DM
depends on galactic longitude:

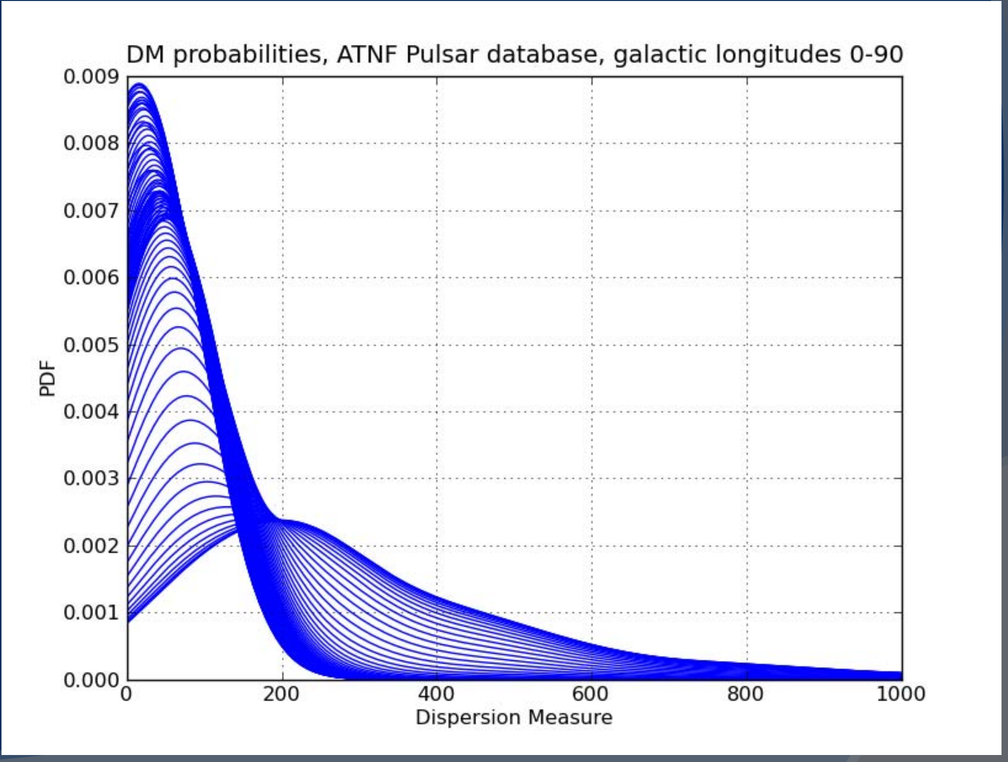

Could also model:
- Extragalactic sources
- Negative DMs
- Etc.

Simulation

- Assumptions
  - P(pulse) = $1\text{x}10^{-6}$; P(non-pulse) = 1 – $1\text{x}10^{-6}$
  - Hardware de-dispersion in parallel; vary capacity of system
- Data
  - Parkes Multibeam [Edwards et al., 2001]
    - 13 beam positions, 1.4 GHz (288 MHz bandwidth), 125 μs sample time, 96 channels
  - Real "noise" data + synthesized pulses
    - Uniform prior over amplitudes, DMs from 0-500 cm^{-3}pc, SNRs (5-10), pulse width (0.5 ms)
    - 11,668 training signals (uniform)
    - 5,074 test signals (galactic lat, lon = 90, 90)

Simulation Results

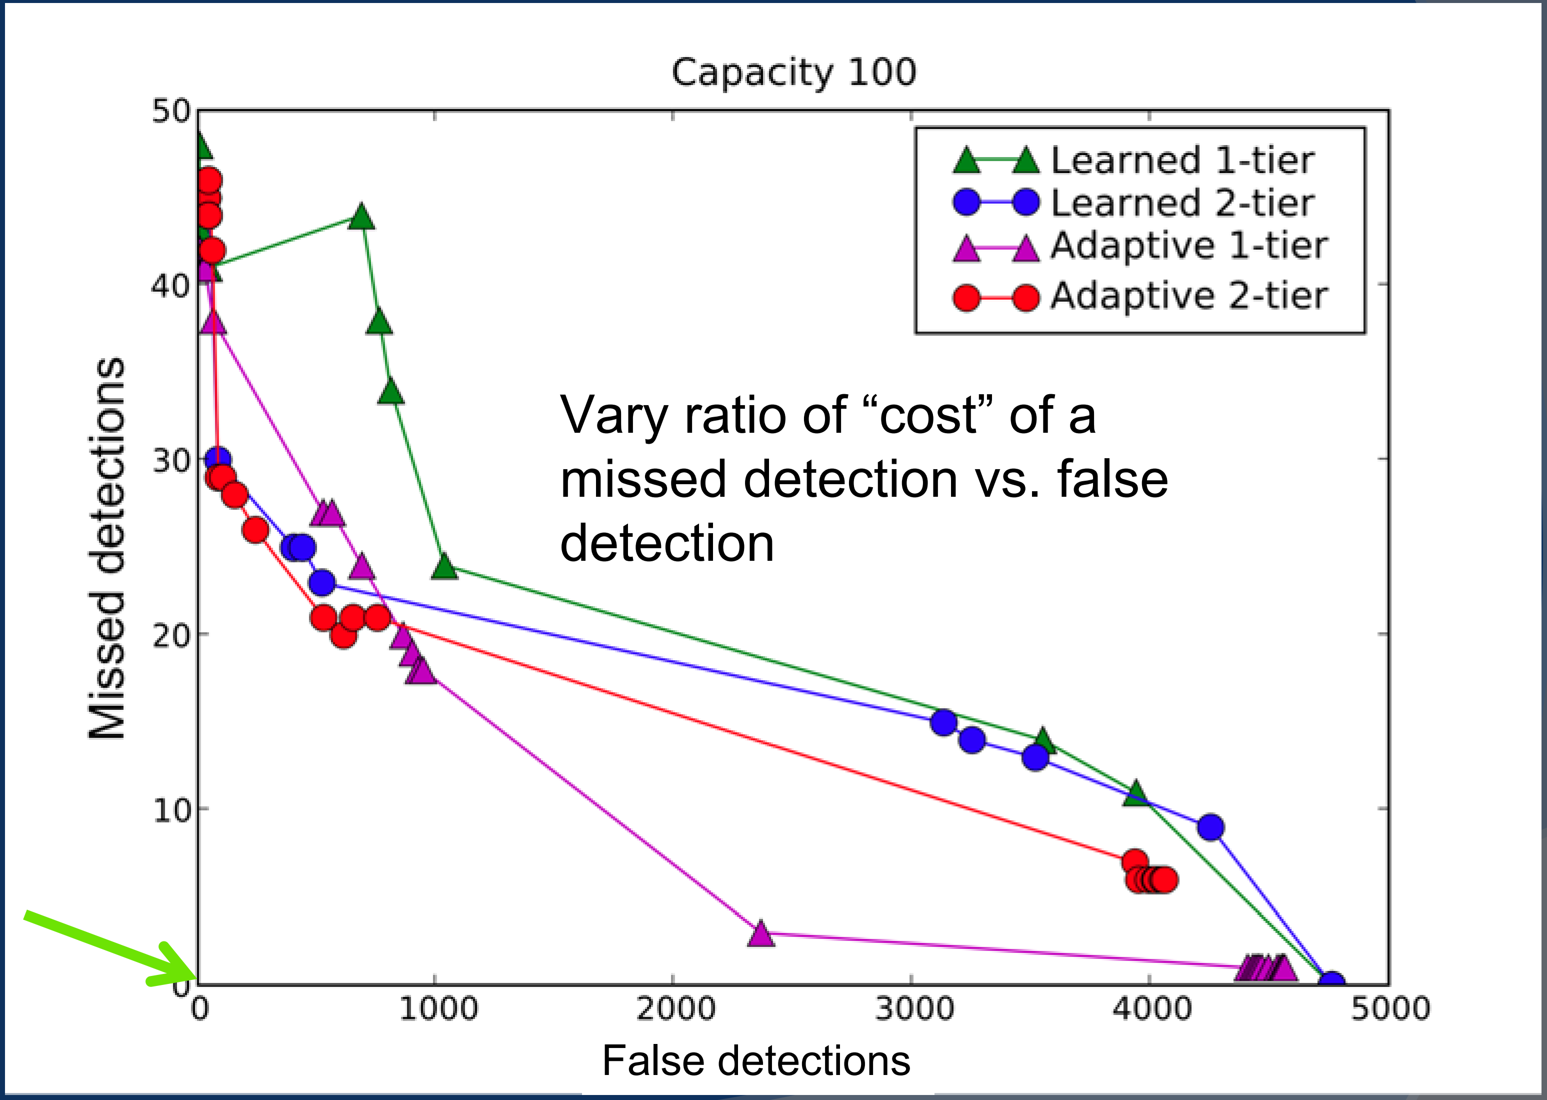

Perfect (0 errors):

Simulation Results

- Machine learning, 2-tier trees cut hardware needs (and power consumption) by ~20%
  - Also reduce false detections
  - Adaptive trees further reduce false detections

To achieve <= 20 missed (of 313):

Model type	Capacity required	False detections
Standard search	**100 (125)**	2000 (1000)
Learned 2-tier	**100**	1500
Adaptive 1-tier	**100**	900
Adaptive 2-tier	**100**	700

- Note: current results are data-limited

What was Learned?

- Top threshold set very low
- Threshold for DM ~0 sub-tests set high
  - Automatic RFI excision

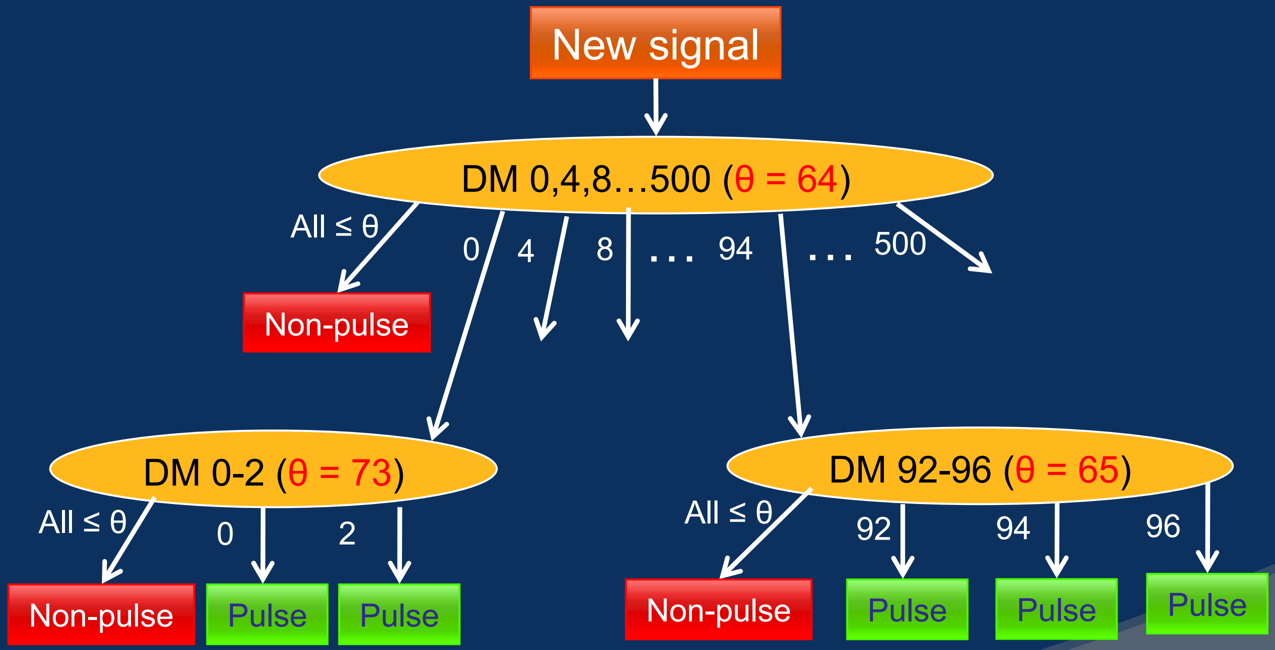

Summary

- Machine learning for fast transient detection
  - 2-tier approach achieves same performance as standard search with 20% fewer DM units, fewer false positives
  - Adaptive trees enable extrapolation to new pointings
  - Some automatic RFI excision
  - Extensible: Can re-train tree as new signal types are discovered
- Next steps:
  - Incorporate information from multiple beams (RFI excision)
  - Other tests in decision tree nodes
    - E.g., event detection in time-frequency images
- Thank you: Sarah Burke-Spolaor, Dayton Jones, Bob Preston, Larry D'Addario, Robert Navarro

This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.