Artificial Intelligence and Novelty

Casey Robinson
caseyr@gmail.com

Division of Science and Mathematics
University of Minnesota, Morris

April 28, 2012

What is AI?

- "The theory and development of computer systems able to perform tasks that normally require human intelligence, such as visual perception, speech recognition, *decision-making,* and translation between languages" [Oxford English Dictionary]
- This definition includes both learning and static systems

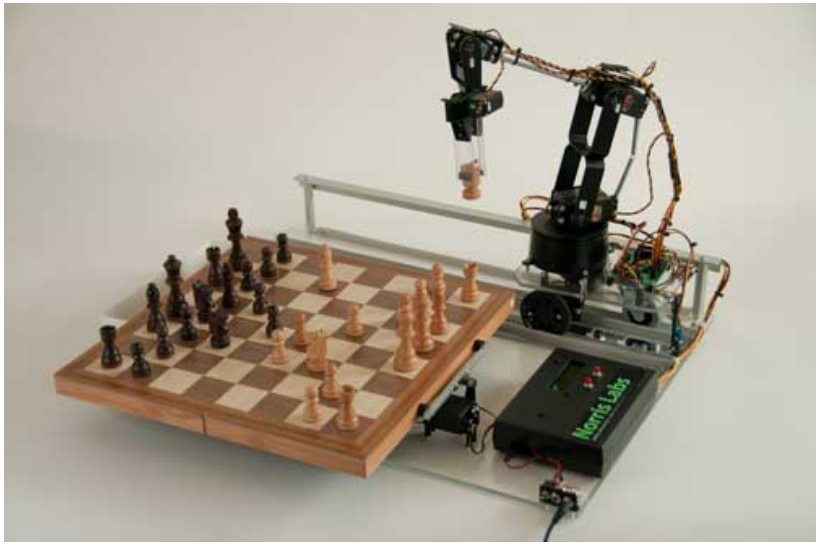

Photo credit Norris Labs, (norrislabs.com)

The Problem With Novelty

- Static systems often cannot cope with factors not considered during their design
- Learning systems require repeated exposure to data to encode an appropriate reaction
- Spurious or very infrequent events may have more importance than those seen more frequently, but are harder to learn about

Top photo credit: Adam Hart-Davis

Bottom: rickkirvan.net46.net

Overview

- Based upon a model of biological immune systems
- Detects behaviors not considered "normal"
- This can be useful in the context of network security, where detecting novel attacks would have great value

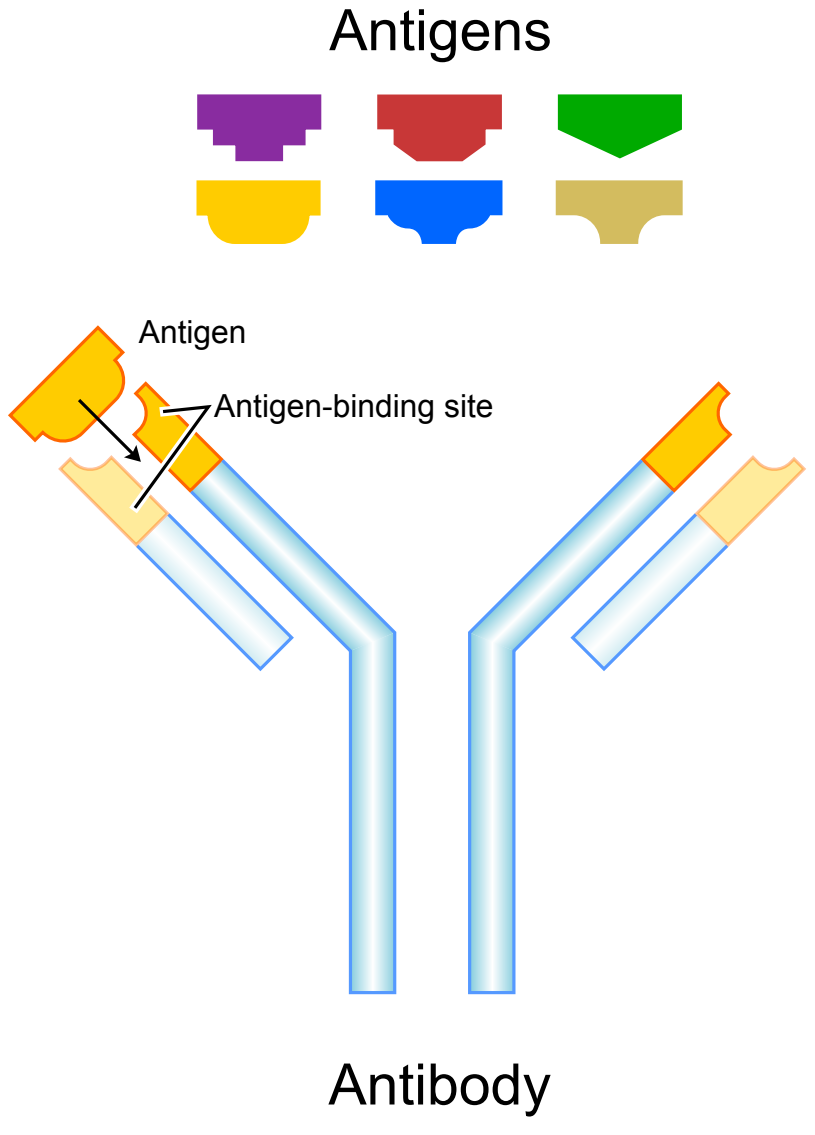

Detectors

- Detectors fit the role of antibodies in an immune system
- Detectors are bitstrings, compared to target data

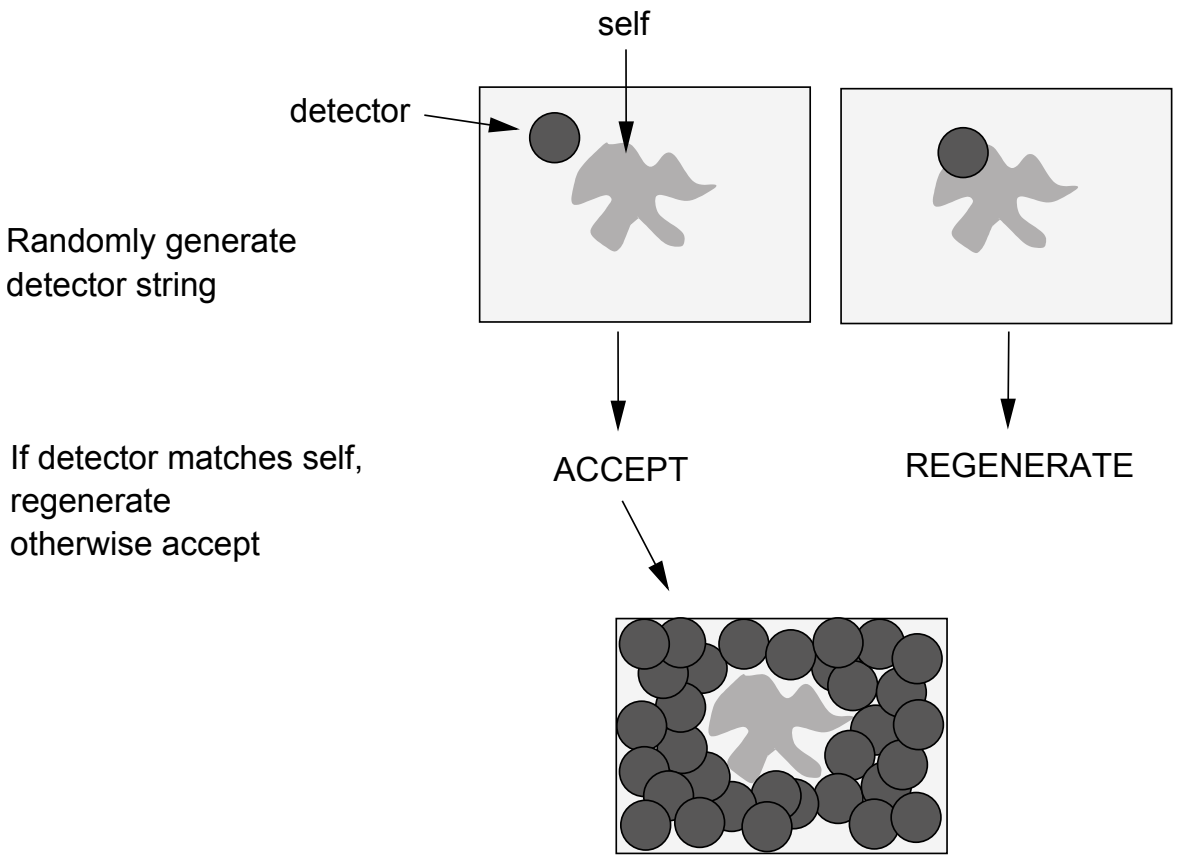

Figure: From [Hofmeyr-Forrest, 2000]

Coverage and Generalization

Adequate Fitting

- It is impossible to train an AIS on absolutely everything it should view as self
- An overzealous fit to the training data will yield an AIS which generates a lot of false positives
- A system which is too loosely fit, by contrast, will often fail to recognize non-self elements

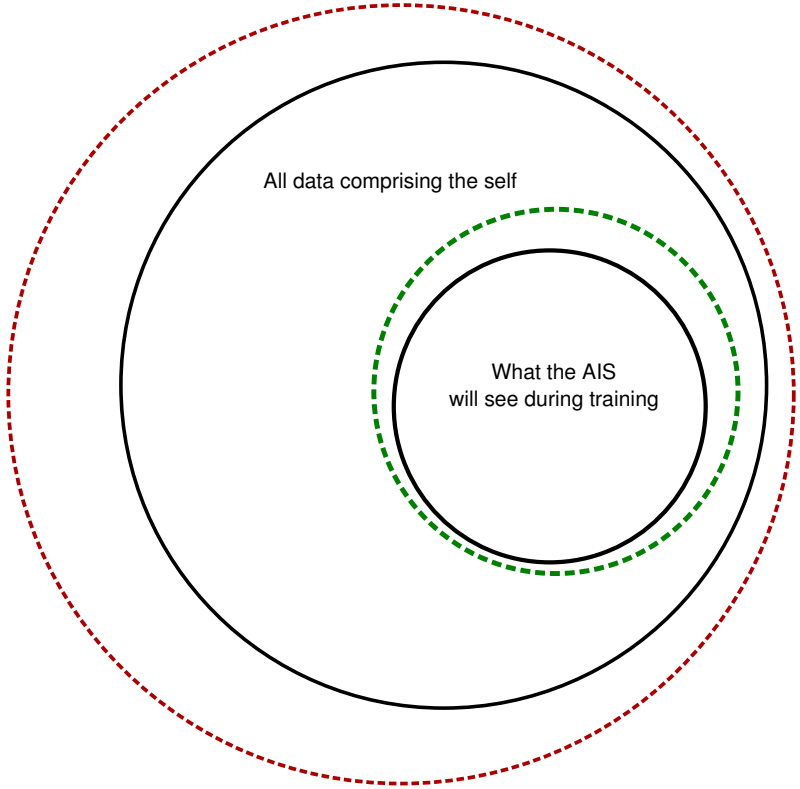

Figure: In the above diagram, the red dashed circle represents the result of underfitting, and the green circle represents the result of overfitting.

r-Contiguous Bits Matching

Improving Match Rates

- Chances of a complete match between a detector and a bitstring is very low for large detectors
- *r*-chunk matching scheme vastly increases the likelihood of a match

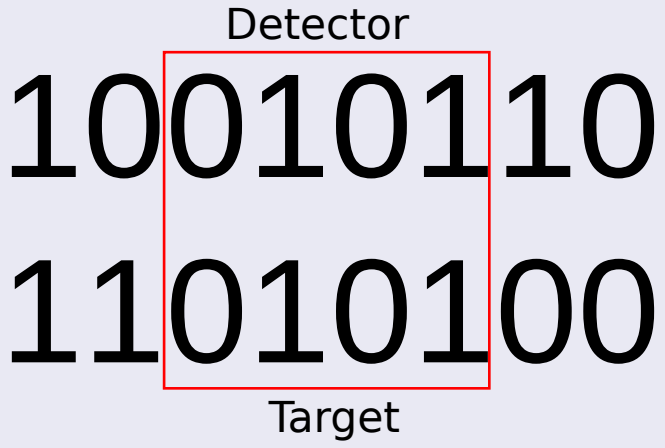

Figure: Two different bitstrings, which have a match for r=4

Unreachable Patterns

- Even using *r*-chunks detection, some patterns cannot be detected
- Removing the *contiguous* constraint will produce far too many matches

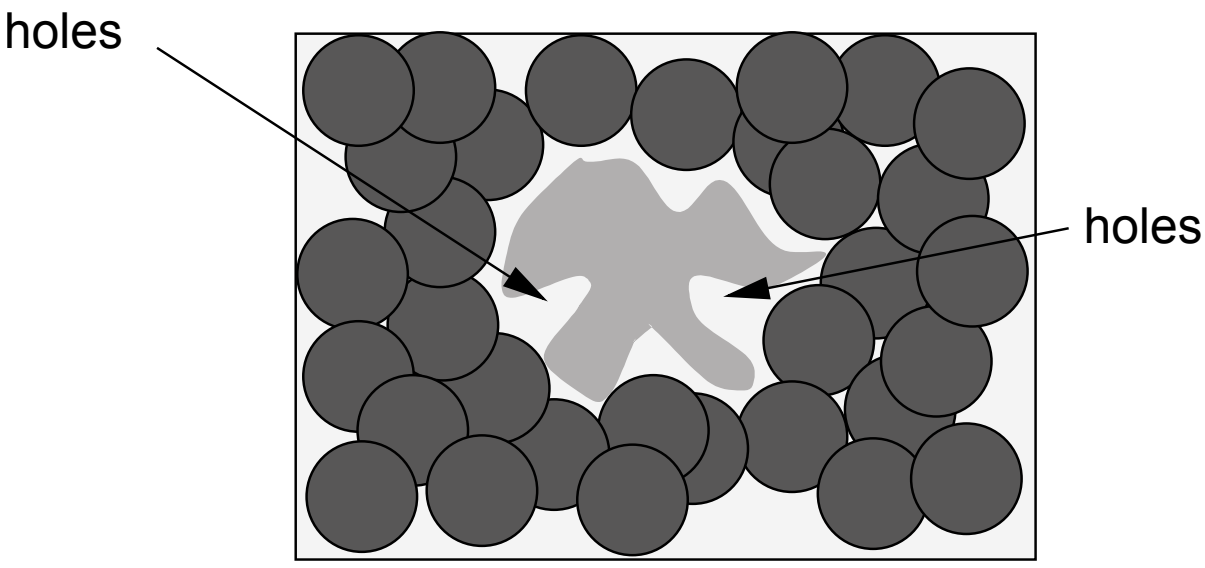

Figure: From [Hofmeyr-Forrest, 2000]

Permutation Masks

- Permutation masks reorder data in a predictable way, allowing detection of some such features

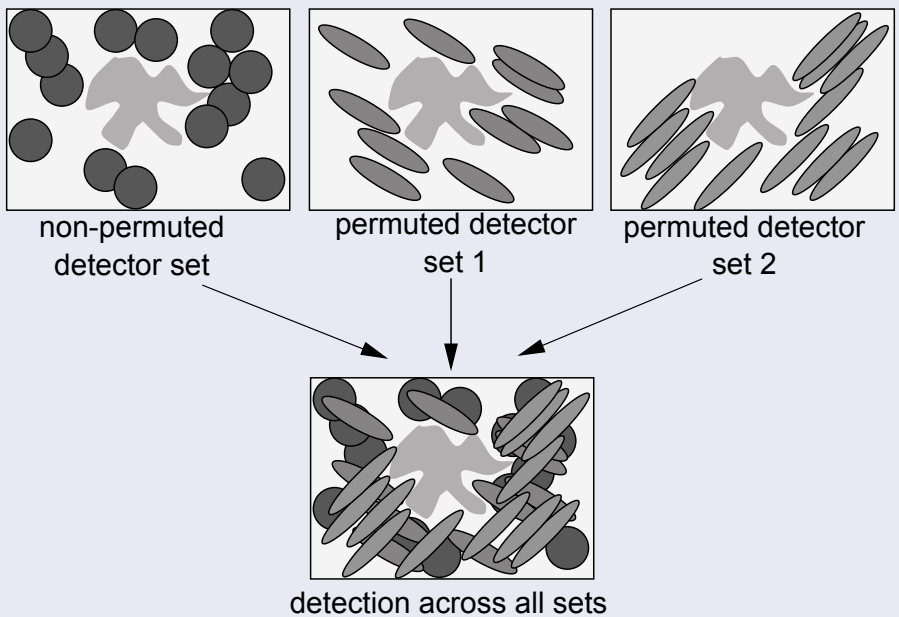

Figure: From [Hofmeyr-Forrest, 2000]

Permutation Mask Example

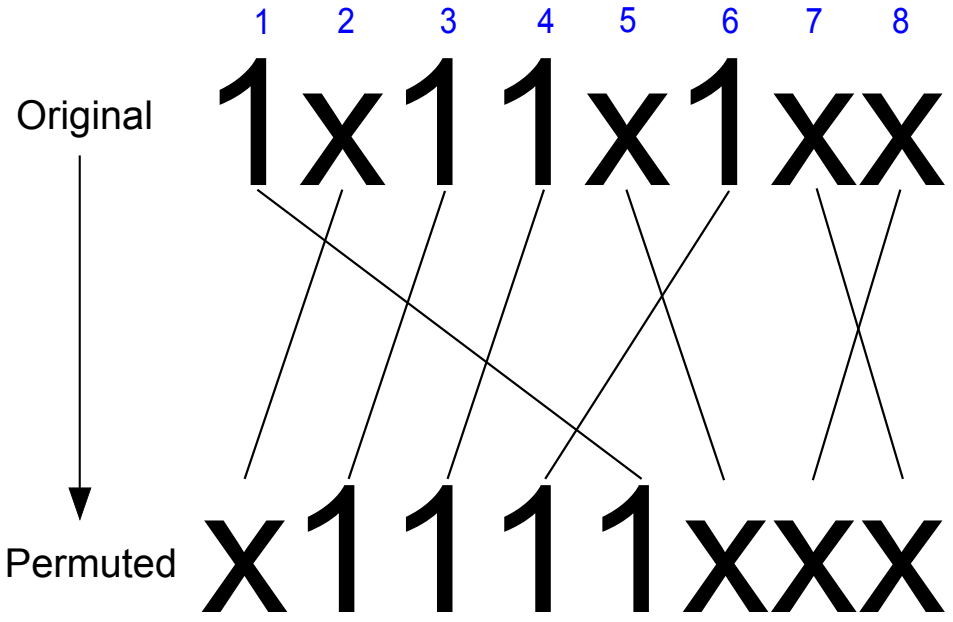

Figure: Application of a permutation mask 2-3-4-6-1-5-8-7

Overview

- Developed by Steven Hofmeyr and Stephanie Forrest
- Checks packets being transmitted over a network
- Does not use timing data
- Multiple detectors must match a packet in order for the packet to be considered non-self

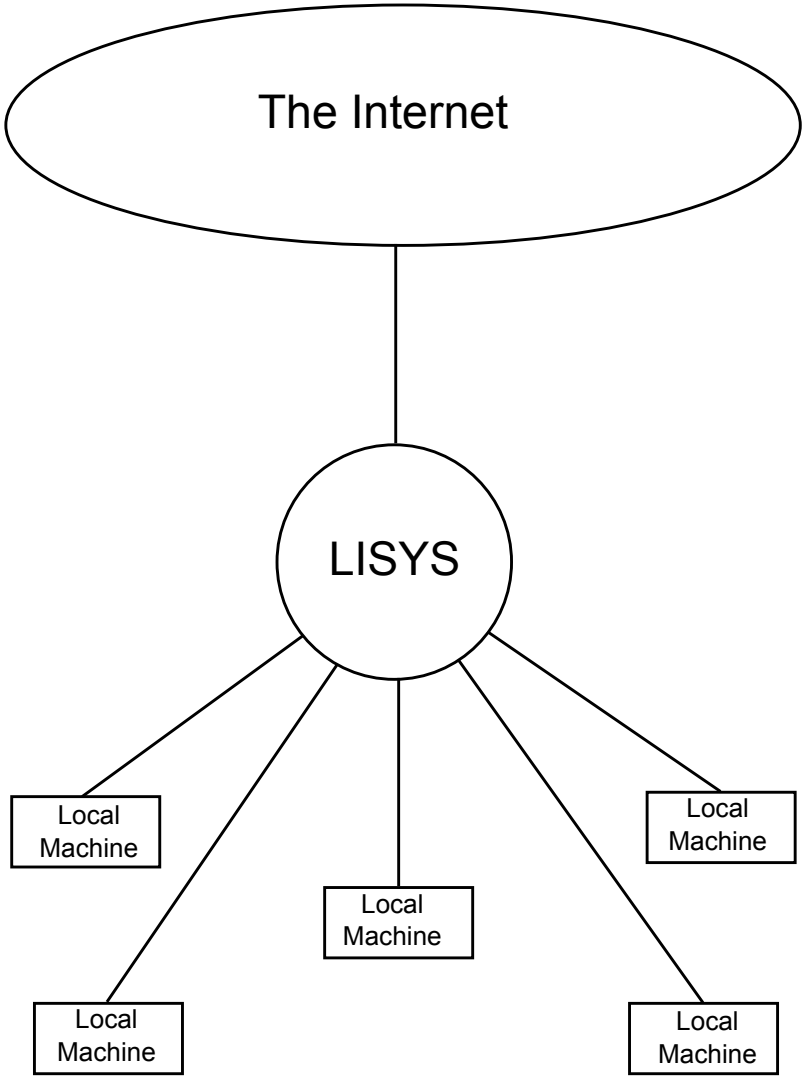

LISYS Packet Representation

- Checking every packet sent over a network would invoke a lot of overhead
- LISYS examines a compressed representation of SYN packets, sent when a connection is being established between two machines
- LISYS only decides whether a connection should be allowed

49-Bit Packet Representation

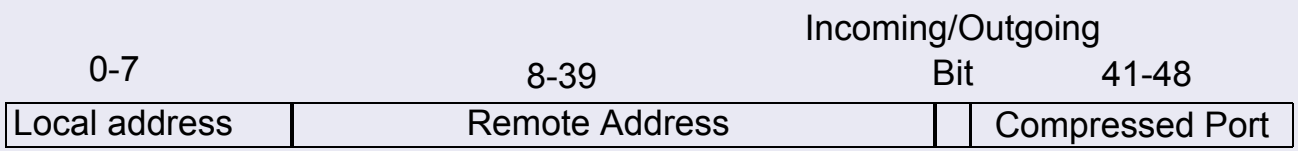

Experimental Setup

- Data set was 15,000 packets – roughly 8 days of normal traffic
- LISYS developed self-model by examining roughly half of this set
- Due to the fact that the same machines may connect many times, only 131 unique strings in training set
- Both normal detectors and detectors using permutation masks were generated
- After training, remaining half of data set with 400 packets of “attack” data shown

Results

Permuted Detectors

Matched almost all attack data, some had better false positive rates than others

Packet Representation

Local:Remote:IO:Port

Definitions

- **False positive** – Self-data which the system classified as non-self
- **True positive** – Attack data which the system classified as non-self.

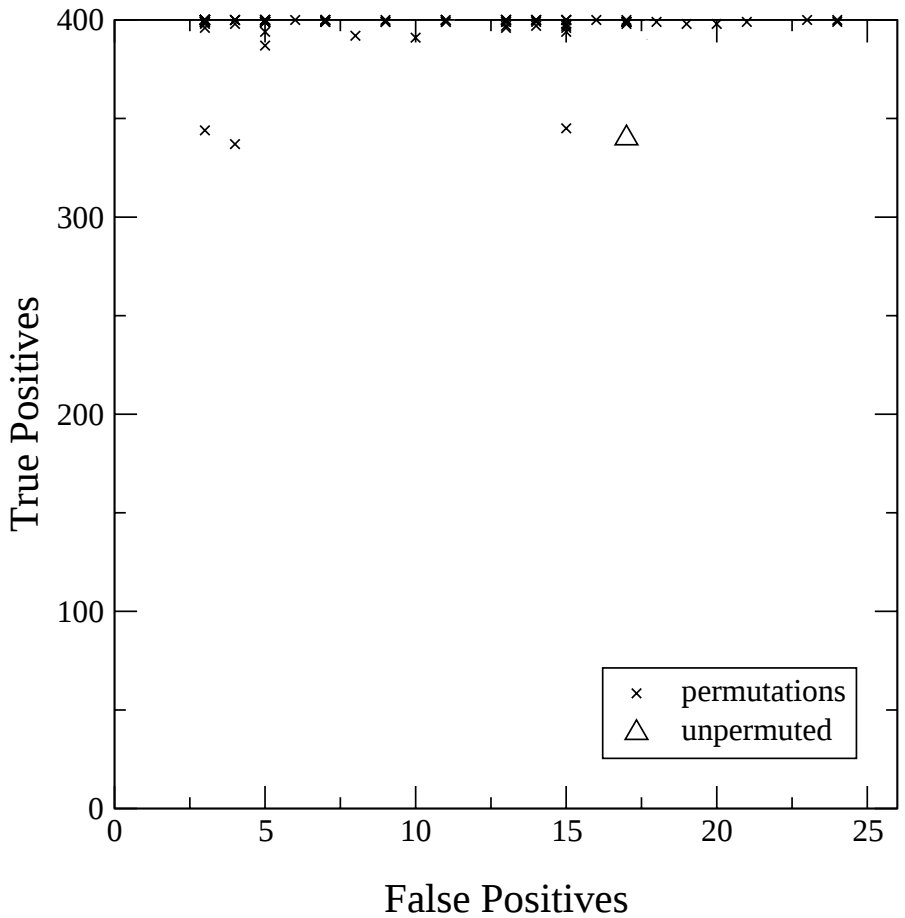

Figure: Based on [Balthrop *et al.*, 2002]

Results

Full-Length Detectors

Matched most attack data, but also a high proportion of **non-attack** data

r-Chunk Detectors

Matched even more attack data, tended to have fewer false positives (varying with value of *r*)

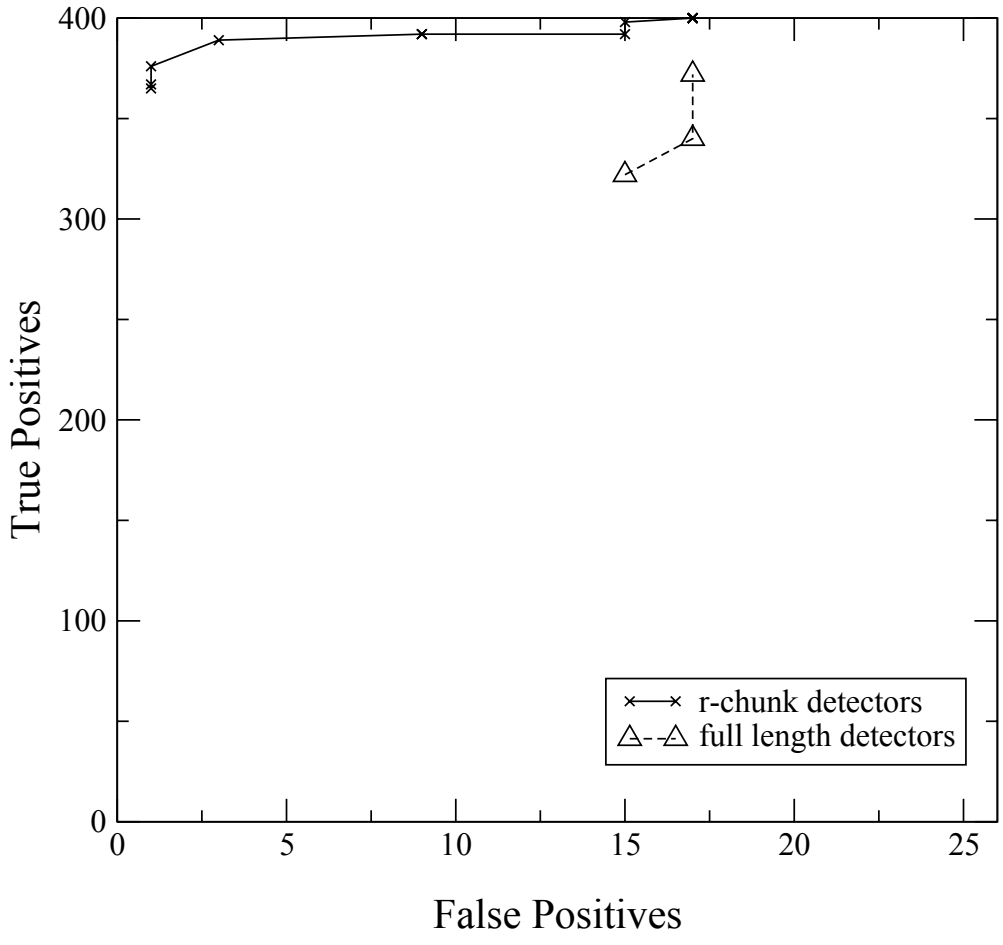

Figure: From [Balthrop *et al.*, 2002]

Analysis

Small *r* Results

- Good performance on small values of *r* was due to how dynamic / static IP addresses were assigned in the experiment
- Static IP addresses were (0..127).x.x.x, while dynamic IP addresses were (128..255).x.x.x

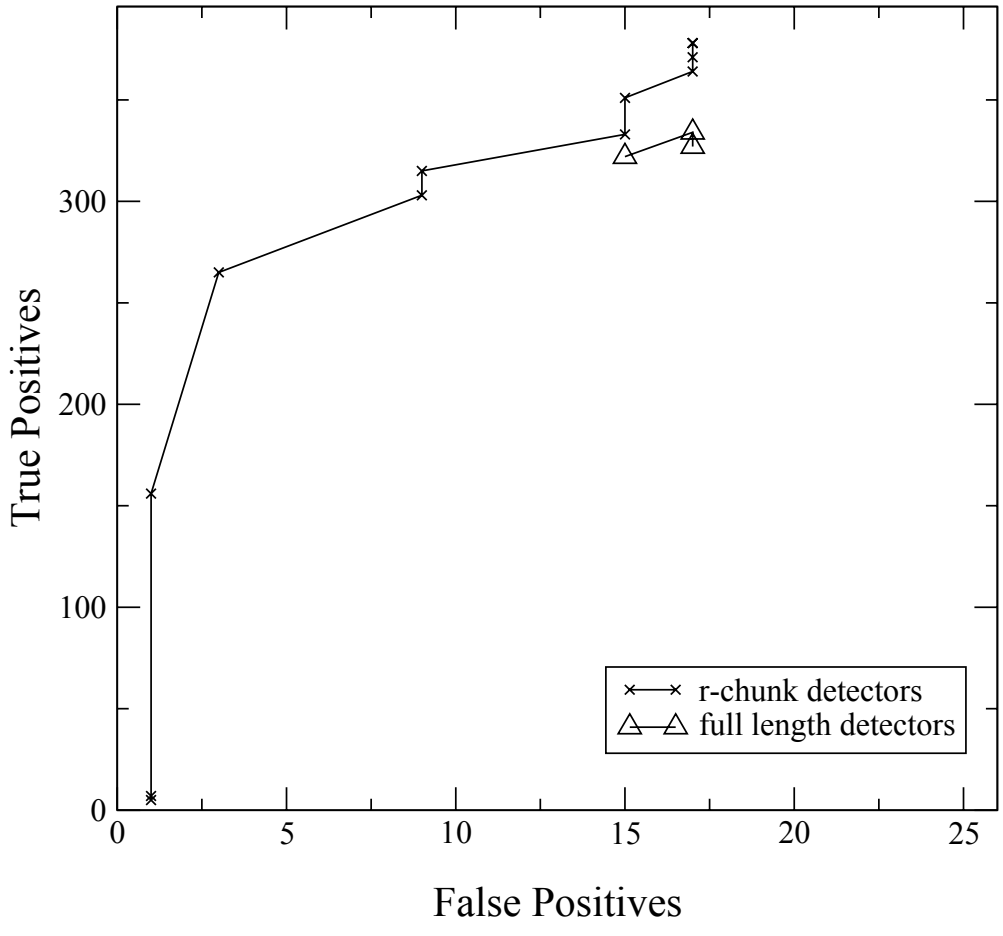

Figure: Results after removal of "magic bit." From [Balthrop, *et al.*, 2002].

General AIS Results

Strengths of AIS

- Given sufficient training on any self, AIS can recognize non-self rather efficiently
- Makes no assumptions about the state of non-self, so both familiar and novel non-self elements are detected

Weaknesses of AIS

- Cannot determine *why* something is non-self, only *that* it is non-self
- Ensuring an appropriate degree and type of generalization is very difficult

Other Applications of AIS

- Industrial settings
- Malware detection
- More?

A-Brain: A General Problem Solving Algorithm

- Extends the concept of evolutionary programming
- Develops solutions to problems, keeps a record of solutions for problems already seen
- Its ability to recognize solved problems is fragile – depends on accurate user input

A-Brain Operation

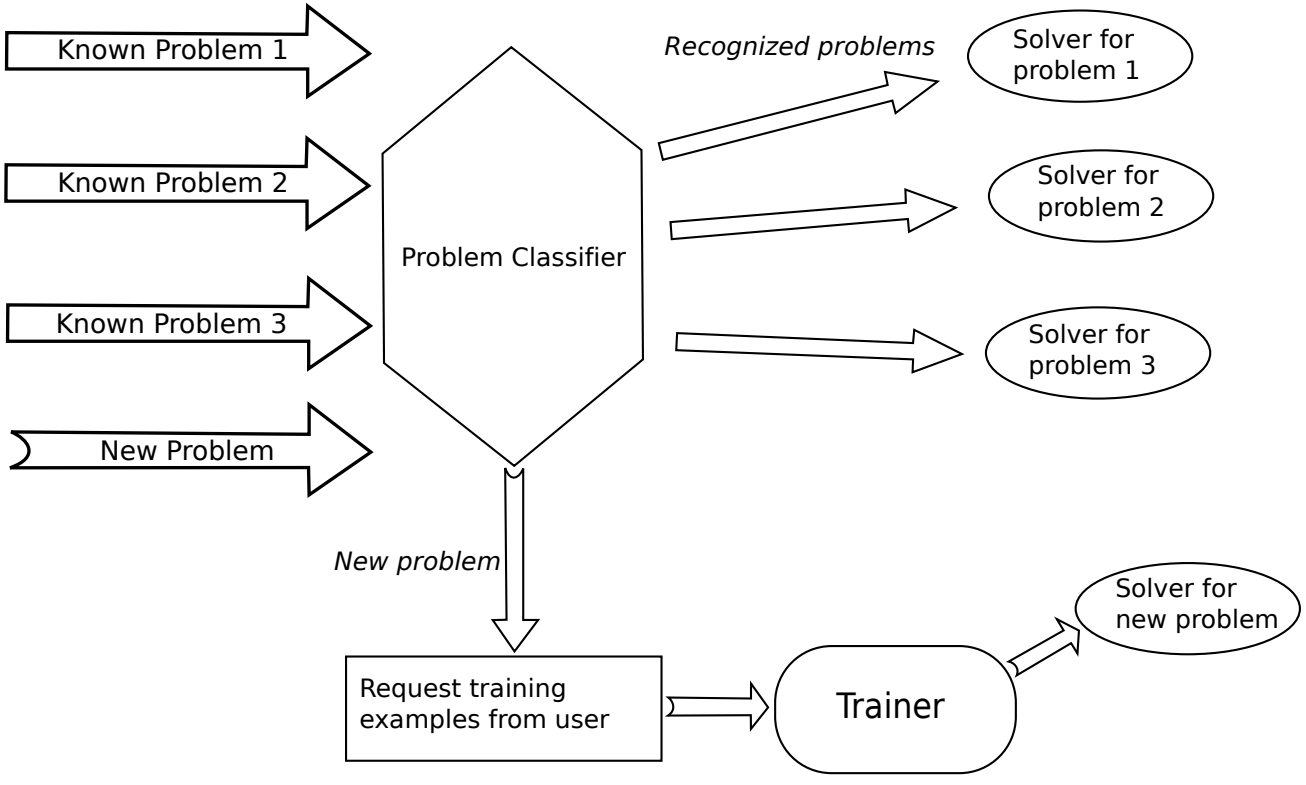

Figure: From [Oltean, 2007]

For further information on A-Brain, see the original text [Oltean 2007], or the paper associated with this presentation.

Conclusions

Two Sides of the Same Coin

- AIS good at recognizing novelty but cannot address it
- A-Brain capable of responding to novelty, but bad at recognizing it
- Both aspects are important, but neither one more so than the other

Importance of Novelty

An ability to appropriately handle novel data is key to creating a system which we would identify as “smart”

Thank You

- Comments?
- Questions?

References

References

1. Stephen Hofmeyr and Stephanie Forrest, *Architecture for an Artificial Immune System*, 2000.
2. Mihai Oltean, *A-Brain: A General System for Solving Data Analysis Problems*, 2007.
3. Balthrop et al, *Coverage and Generalization in an Artificial Immune System*, 2002.

Detectors, cont

Short Detectors

Detectors of length N can be used to examine strings of length $\geq N$.

Target	0	1	0	0	1	0	1
No match	1	0	0	1			
Match		1	0	0	1		
No match			1	0	0	1	
No match				1	0	0	1