Autonomy and Machine Learning: Good for Cybersecurity, and Maybe also Bad

Paul D. Nielsen

Director and CEO

Security concerns have grown from nuisance intrusion . . .

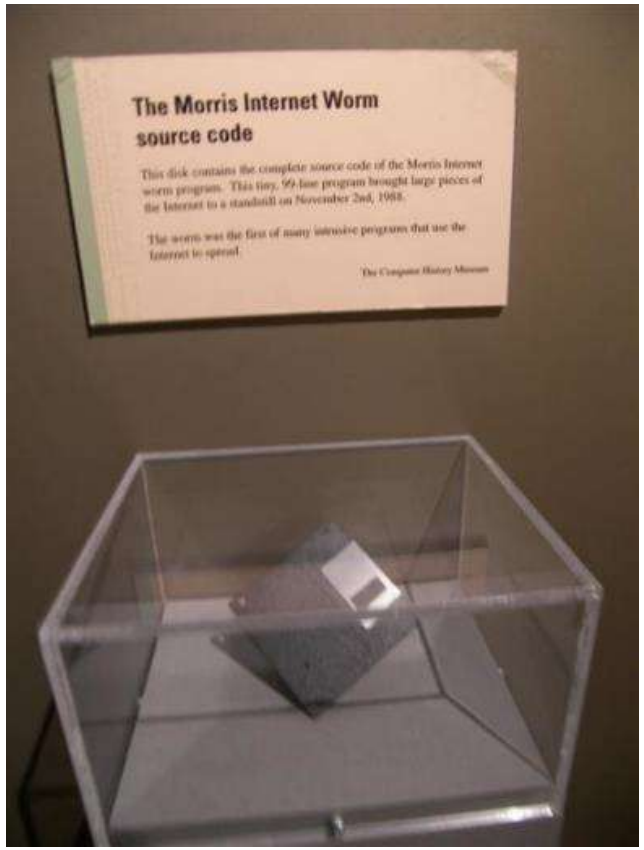

Meant to gauge the size of the fledgling Internet

Unintended consequence: denial-of-service attack
Prompted founding of CERT/CC

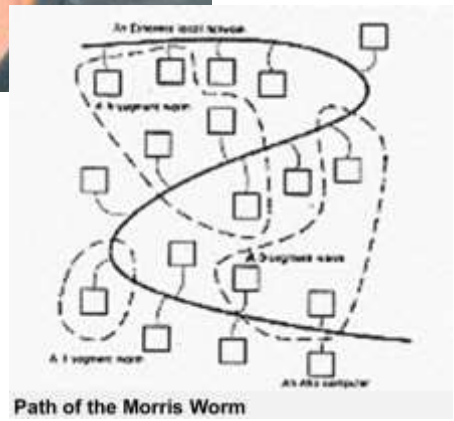
Path of the Morris Worm

The Morris Worm compromised about 10 percent of all systems connected to the ARPANET in 1988. It was named after its creator, Cornell graduate student Robert Tappan Morris.

. . . to intentional network attack, just a few years later

U.S. Air Force Research Laboratory, Rome Research Site (1994)

. . . to the explosion of cyber-attacks today

600,000 Facebook accounts hacked—per day

30,000 websites hacked—per day

110,000 attacks on U.S. Navy networks—per hour

18 victims of cybercrime—per second

$100 billion annual cost for global cybercrime

Evolution of the threat

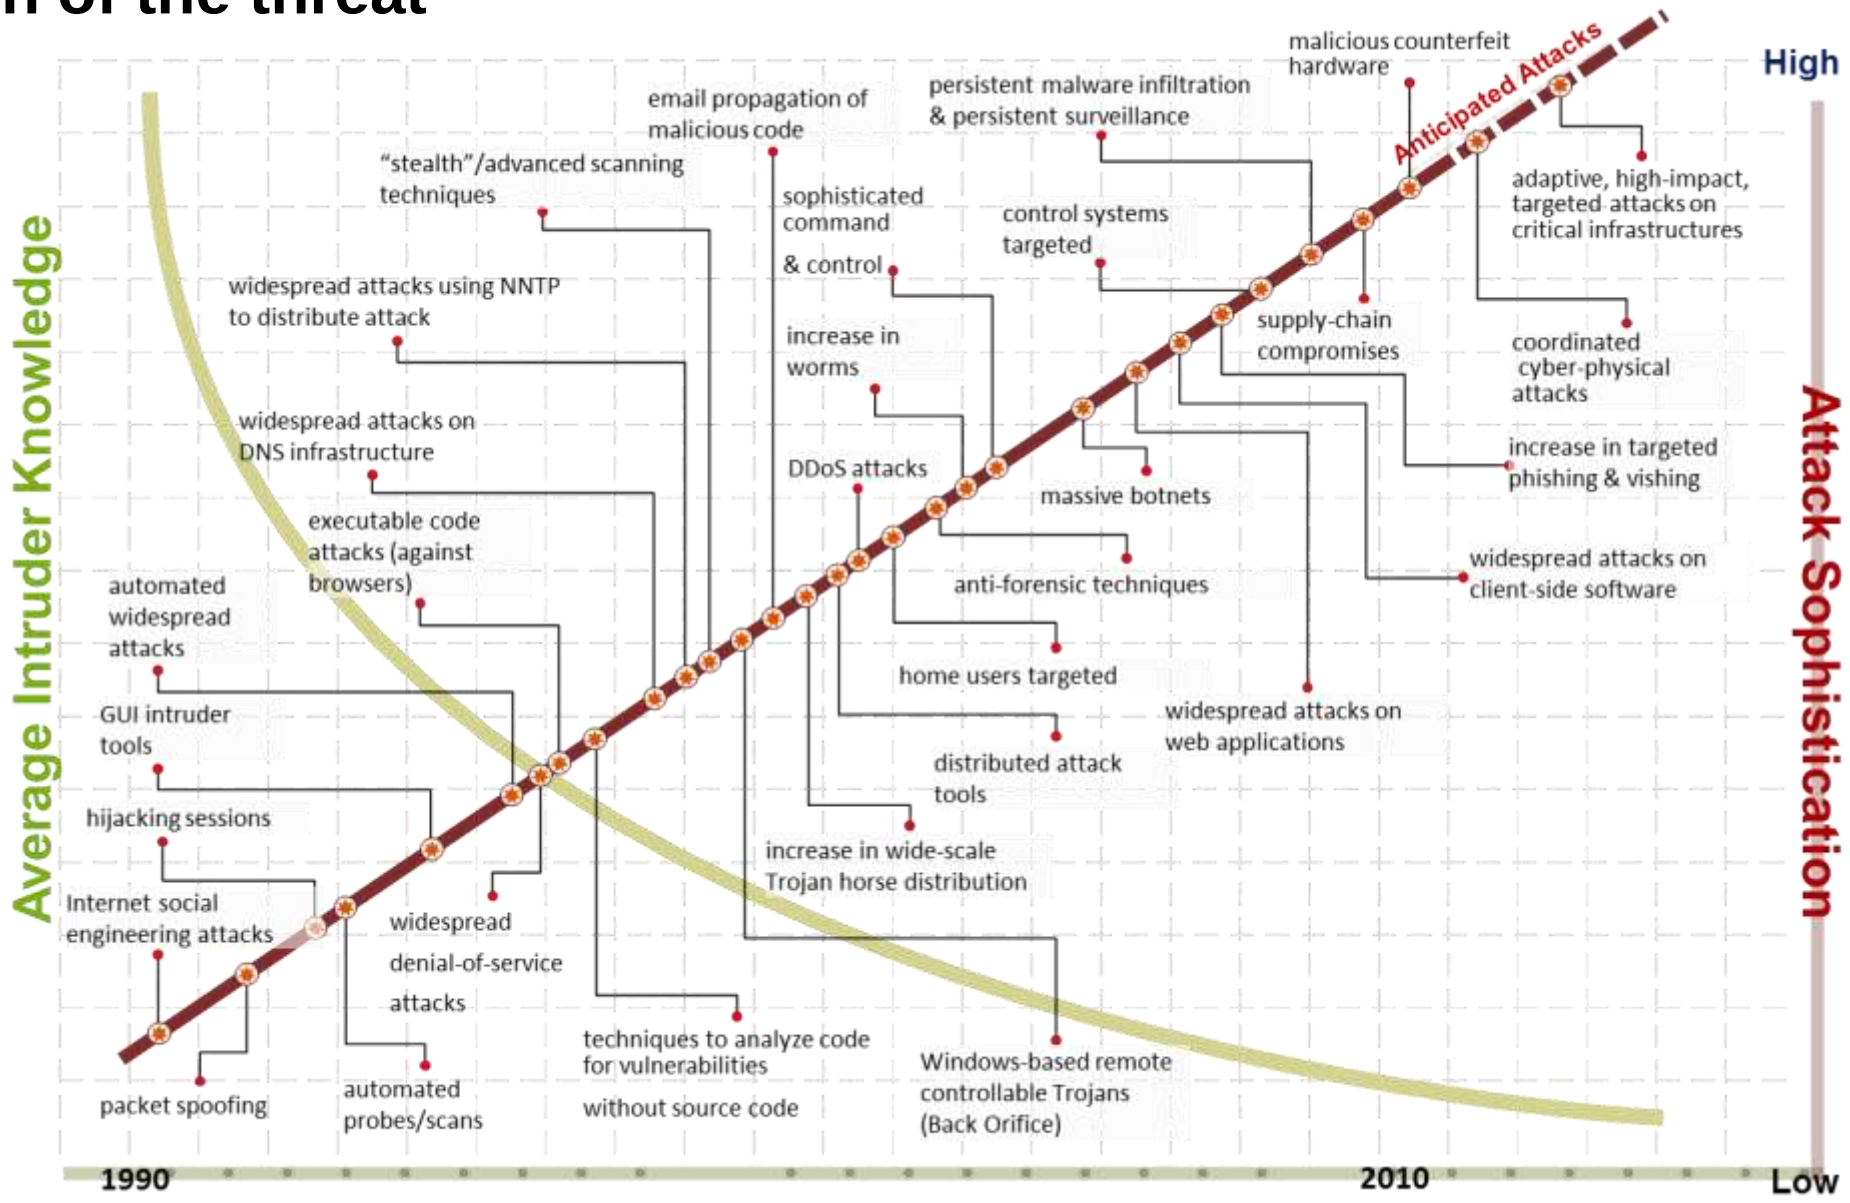

Meet the threat with new science

Autonomous systems can perform some tasks on their own based on their understanding of themselves, their situation, and their environment

Machine Learning technology offers a means to design autonomous systems, because it enables software to automatically improve with experience

Autonomous systems are changing traditional roles

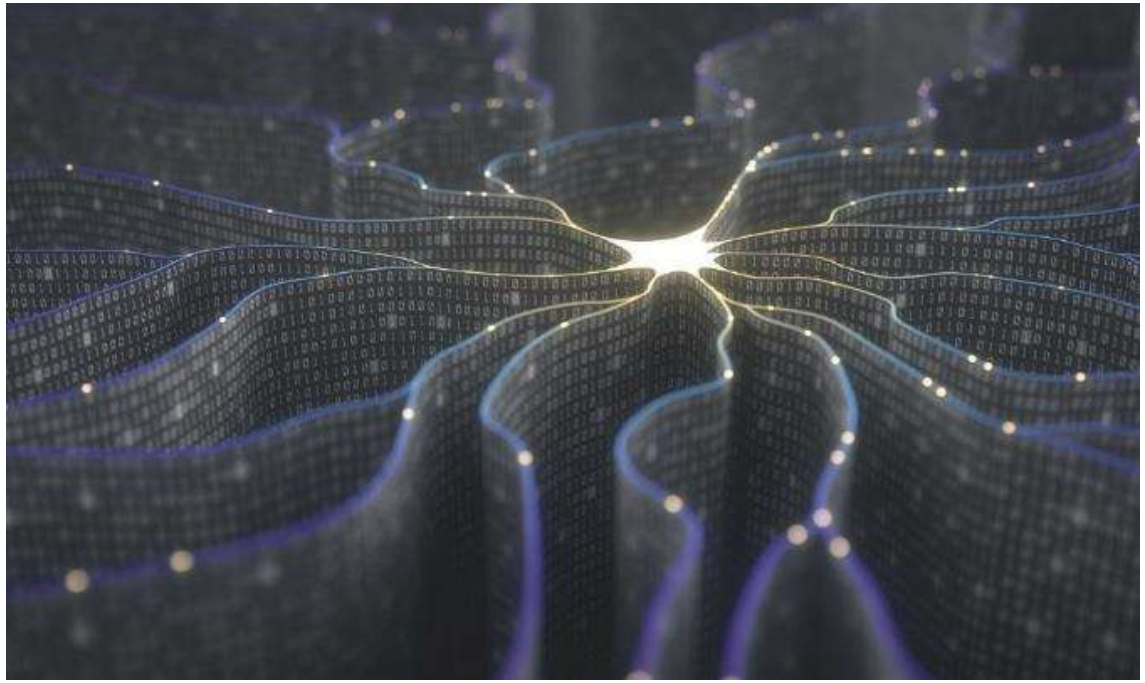

Traditional human-machine roles (human in-the-loop)

1. Machine reads “sensors” and offers alternatives to human.
2. Human chooses actions and monitors results.

But autonomous systems can also take actions

- Read and react, based on its interpretation of “mission” goals
- Human can expand system actions beyond original context (human on-the-loop)

Autonomous systems can improve productivity, operate continuously

Past

Present

Future

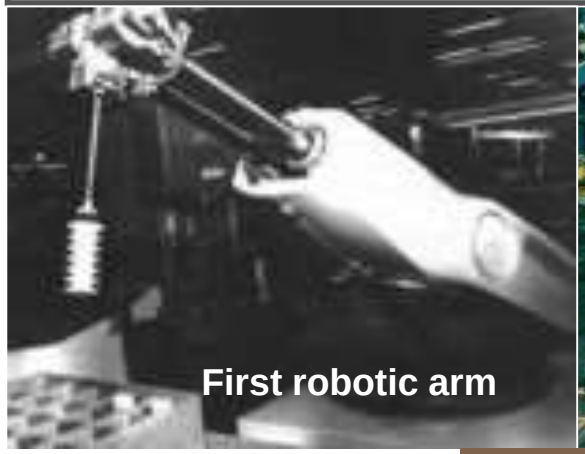

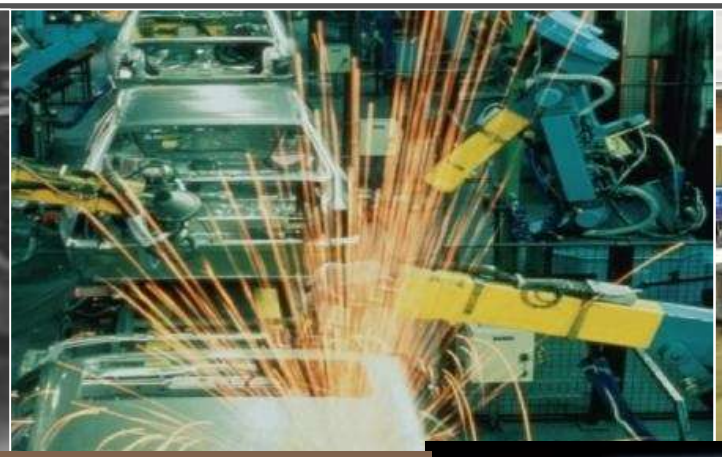

Baxter deep-learning robot

Mars Opportunity Rover

Soft robots: change shape and move via their own internal energy

Bio-inspired prototype "soft robot" material with greater dexterity and mobility (U. of Pitt.)

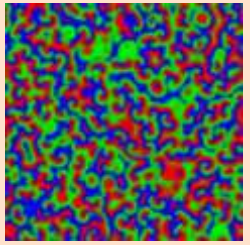

. . . process tremendous volumes of data

Past **Present** **Future**

Winner of 2005 DARPA grand challenge

2,800 microprocessors, with capacity to handle 80 trillion operations per second

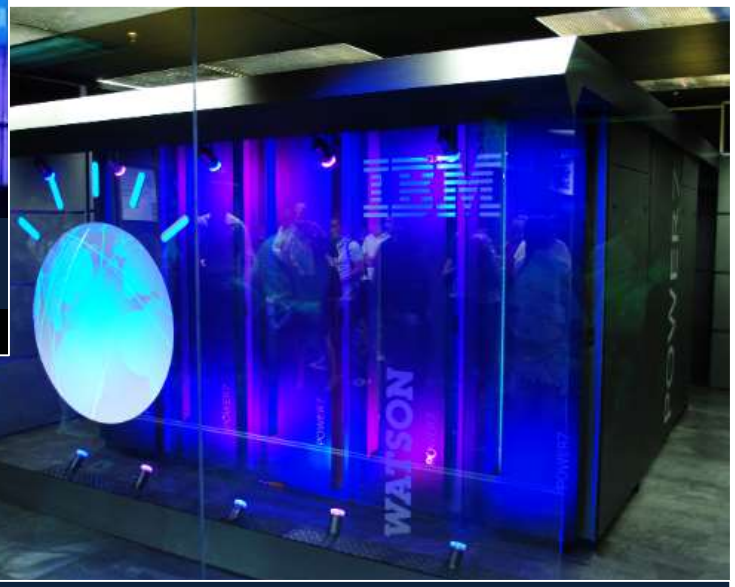

Memorial Sloan Kettering Cancer Center trains Watson to help oncologists make more nuanced treatment decisions more quickly

Diagnostic expert systems

Deep Blue

. . . work where we cannot safely go

Past **Present** **Future**

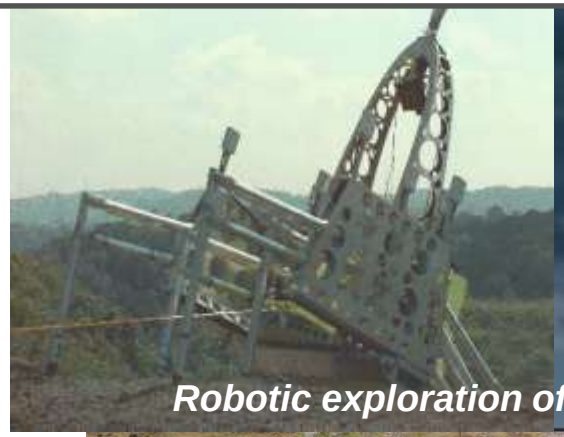

Robotic exploration of extreme terrains

World Trade Center, Iraq, Afghanistan

Search and rescue

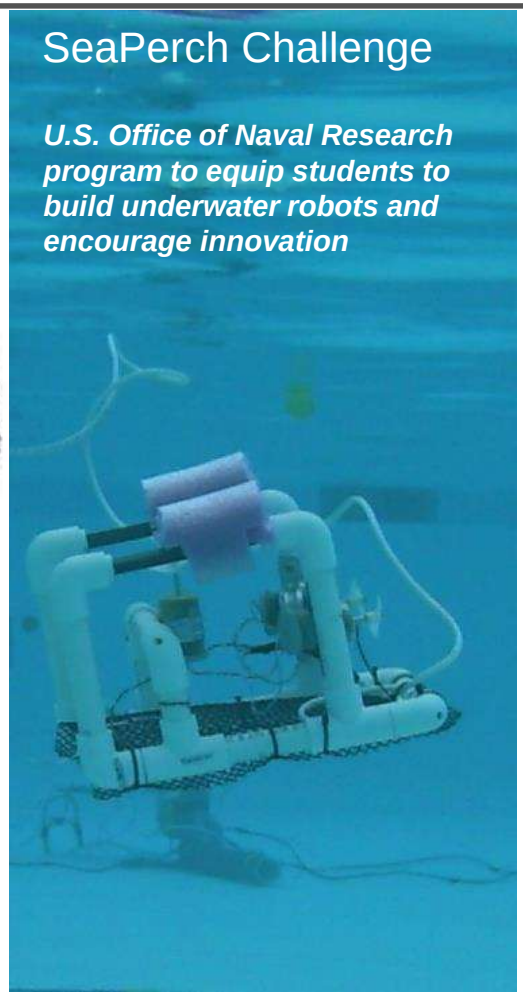

. . . and augment human decision-making

Machine-learned intelligence extends human effort in areas such as

- media analysis and media exploitation (e.g., automated video summarization)
- anomaly-(unlikely event) based detection
- emotion-sensing algorithms

Early research into autonomous system intelligence (through biometrics) has demonstrated

- extraction of heart rate from surveillance video in near-real-time
- recognition of facial micro-expressions, which unintentionally reveal emotions, from video

But face recognition algorithms can be tricked

The facial recognition software deployed at airports by the TSA has failed to catch any terrorists.

CMU privacy researcher Dr. Alessandro Acquisti noted that a baseball cap adorned with glowing LEDs will confuse face-finding algorithms.

Other researchers with OpenAI

- forced an algorithm to misidentify images, up to 97 percent of the time, in pictures taken with a smartphone
- determined that attacks would still be successful even with changes in lighting, perspective, camera optics, and digital processing

. . . and autonomous actions are vulnerable to cyber-attack

The systems are dominated by complex software and adaptive software architectures.

Even the best software exposes defects in operation:

- 1-5% of defects expose exploitable security vulnerabilities
- Criminals and other adversaries can create vulnerabilities, too

Also vulnerable to

- Mis-training
- Spoofing
- Hidden modes

Many attacks come about through steps made to look to the trained machine like acceptable requests

Still, Machine Learning (ML) can make software more secure

In software development, automatic identification and repair of software defects, where a pattern can be identified

In system operation

- Faster, more accurate identification and addressing of software vulnerabilities
- Reduction in time to learn about malware intrusions

ML to repair software code

Identifies and repairs errors automatically

Eliminates potential vulnerabilities at a greatly reduced cost

- Industry estimate: cost to repair an error at the integration testing phase is $100K

Repaired software ultimately reduces a system's vulnerability to attack

ML in vulnerability hunting: Mayhem

Fully autonomous system for finding and fixing computer exploitable security vulnerabilities

Creates, tests, and applies patches in real-time

- Won DARPA's first all-machine hacking tournament Grand Challenge in 2016
- Developed by ForAllSecure, a Carnegie Mellon University spin-off firm

ML to cut malware analysis time

Tools that

- reduce analysis from hours to seconds
- help analysts keep pace with adversary techniques

Useful for both trusted and malicious software code

Collaboration between CMU's Software Engineering Institute and Lawrence Livermore National Laboratory

ML and autonomous systems deliver benefits for cybersecurity

And, they issue new challenges

- Human-system trust
- Continuous runtime assurance
- Testing regimens
- Vulnerabilities/resilience
- Software maintenance and evolution

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

DM17-0890