Is Artificial Intelligence (AI) and Machine Learning (ML) the Solution to our Network Management Challenges?

Petros Mouchtaris
President, Perspecta Labs

IFIP/IEEE International Symposium on Integrated Network Management
April 10, 2019

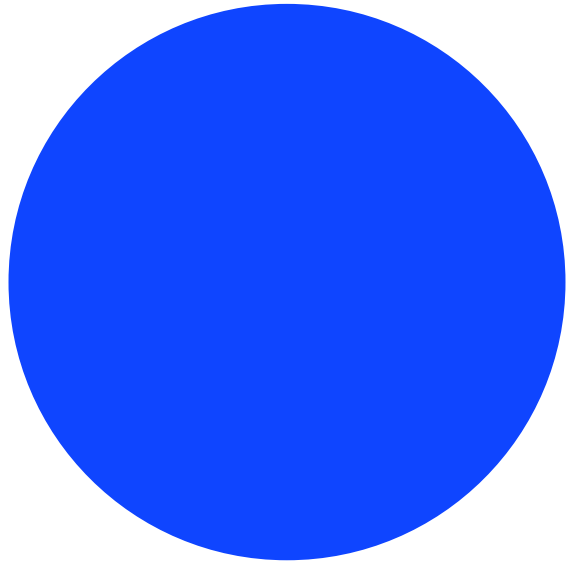

perspecta LABS

Overview

- Complexity is driving the need for automation in network management
- What are important characteristics of Artificial Intelligence (AI) and Machine Learning (ML)?
- Where has AI/ML been used successfully in the past?
- What are the limitations of AI/ML?
- How do we proceed?

Evolution of networks

- Several key aspects of networks are rapidly increasing:
  - Number of devices
  - Heterogeneity (different networks, different vendors, different devices)
  - Mobility
  - Speed of network and bandwidth
- At the same time, the number of human network management experts is not keeping pace
  - The scope and change of networks is difficult for humans to keep up with
- Security concerns are adding another dimension
  - IoT in particular

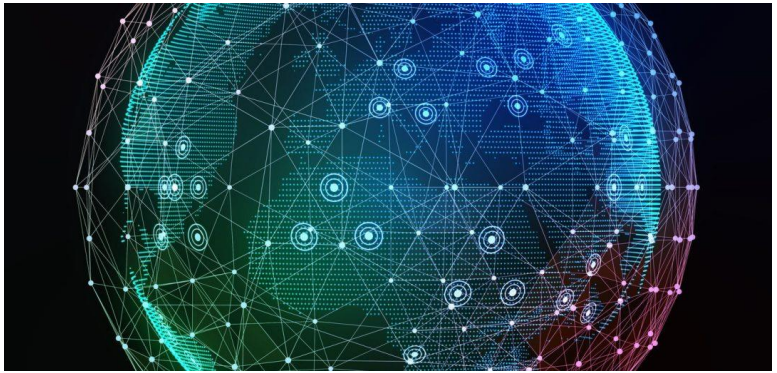

Evolution towards 5G

- Increased data rates
  - 10x to 100x improvement
- Low latency (1ms)
- Reduction in energy usage
- Increase in heterogeneity
- Large number of use cases
  - Diversity
- Enables Internet of Things (IoT)
  - Explosion in number of devices

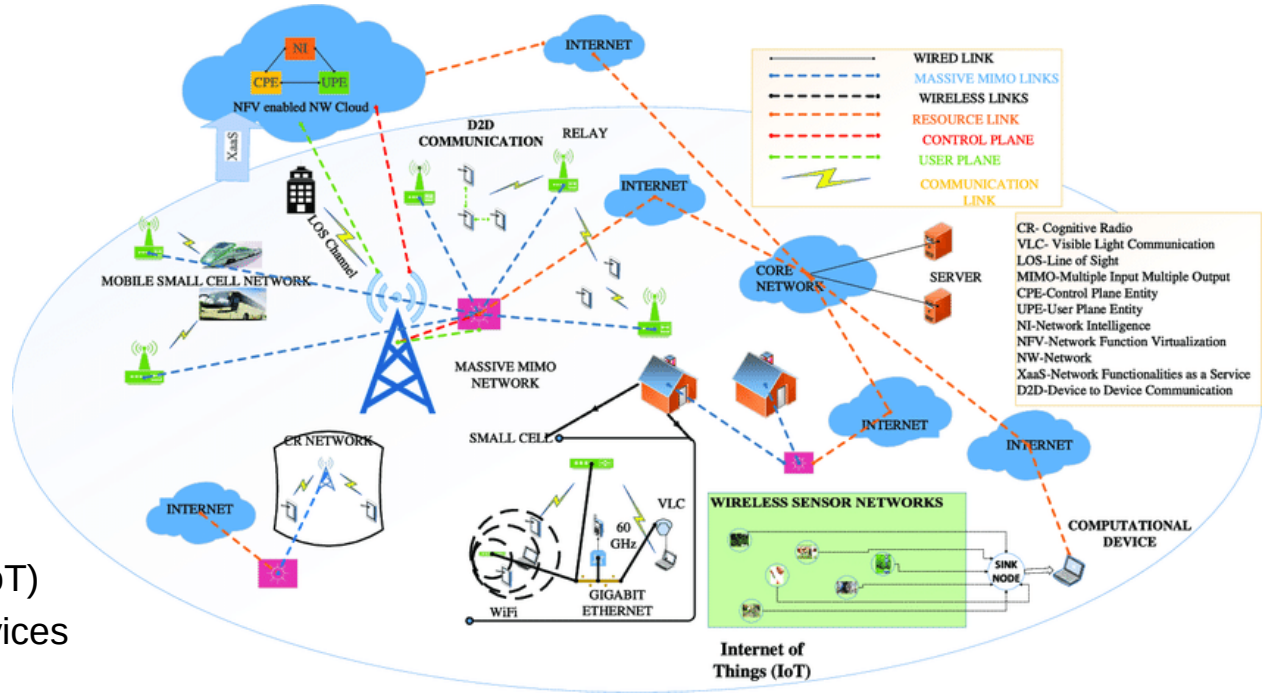

From A. Gupta, “A Survey of 5G Network: Architecture and Emerging Technologies”

Smart Cities and Internet of Things (IoT)

- 30B+ devices connected to the network by 2020
  - 10+ connected devices per person in US already
- Diversity of applications, vendors, protocols, service providers
- Who and how can anyone figure out what's wrong
- Introduces significant security and privacy concerns

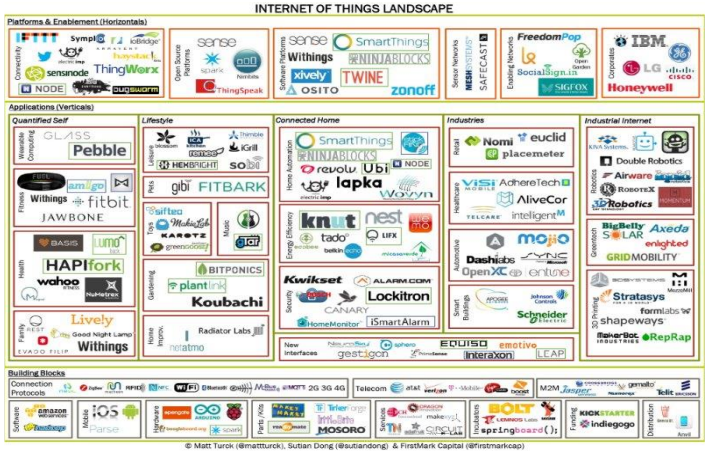

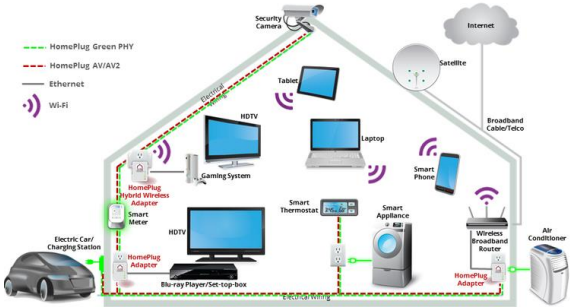

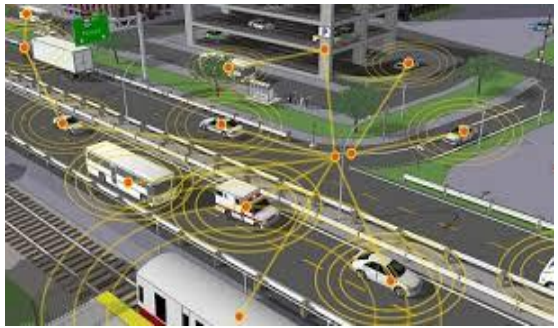

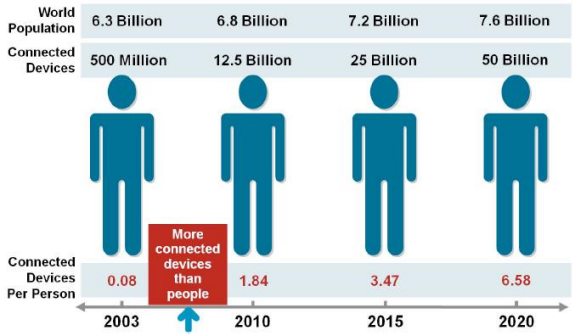

Source: CISCO IBSG

Challenge: Security and Privacy

- Low cost and small memory foot-print in IoT devices result in rudimentary security functionality
- IoT devices are difficult to patch due to scale and accessibility
- Many IoT builders use third party security solutions without deep understanding of the overall security architecture
- Internet connectivity and proliferation of attack software enable remote attacks
- Many IoT devices have wireless communication with no privacy
  - Wireless attack tools are becoming much more inexpensive

Mirai Botnet Map - 150,000 devices - 1Tbps DDoS traffic

perspecta LABS

Network Function Virtualization (NFV)

- Network devices (router, switch, middle-box, etc.) have been vertically integrated systems
  - Typically expensive with significant barriers to competition
  - Adding new functionality is a slow and expensive process (at the control of the vendor)
- Disaggregation has many advantages
  - Use commodity hardware (cheaper)
  - Mix and match best of breed
  - Scalability (adapt quickly to demand growth)
  - Opportunities for innovation (faster)
- Disaggregation introduces challenges
  - How components interact with each other
  - Identifying/addressing problems
  - Heterogeneity
  - Increased speed of change

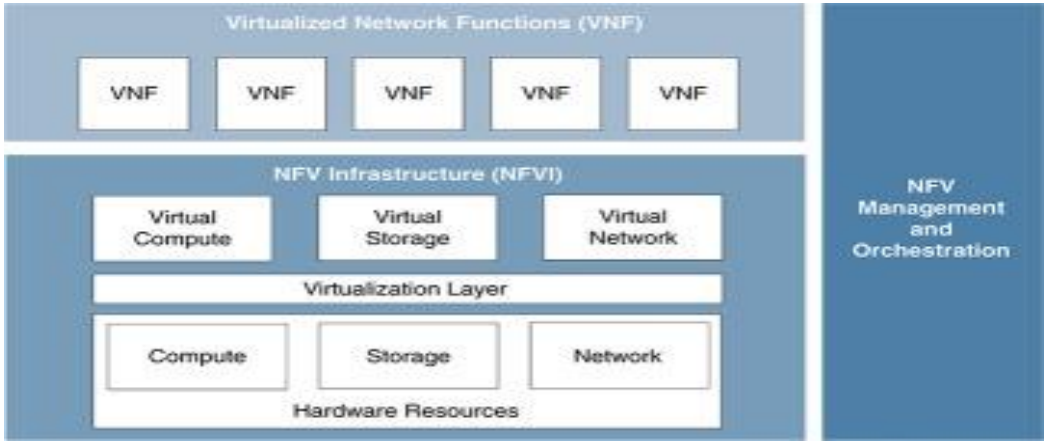

Software-Defined Networking (SDN)

- Software-Defined Networking (SDN)
  - open systems and standards-based components
  - Commodity L2/L3 Switch-Router hardware
- Key advantages:
  - Ability to centralize network provisioning (reduced cost?)
  - Improved security management
  - Reduced hardware costs (using commodity hardware)
- Increased complexity (especially with hybrid architectures)
- Another dimension of heterogeneity
- Increased speed of change

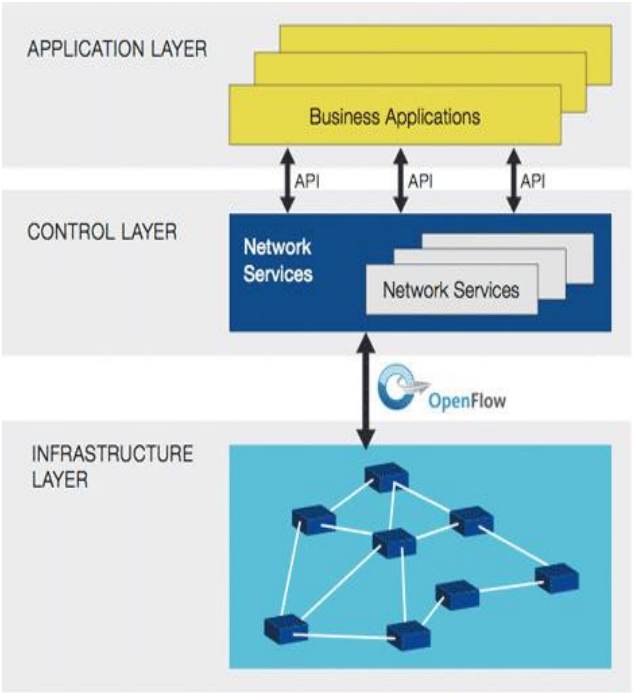

Is AI the solution to our Network Management Challenges?

- Network management requirements for future networks are becoming too difficult for humans:
  - Scalability
  - Heterogeneity
  - Speed of change
  - Amount of data
  - Speed of information
  - Security requirements

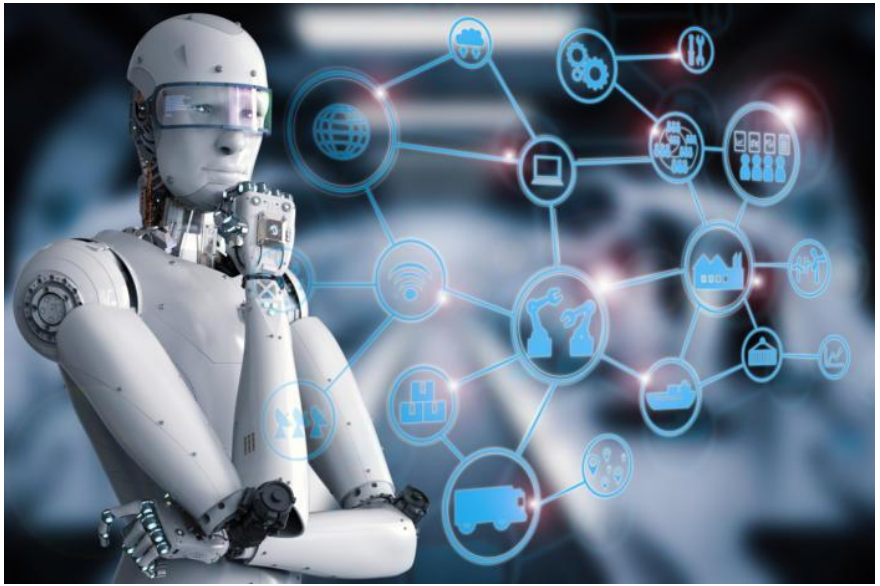

Source: The CyberSecurity Place

Basics of Artificial Intelligence

- Artificial Intelligence (AI) = A system's ability to correctly interpret external data, to learn from such data, and to use those learnings to achieve specific goals and tasks through flexible adaptation
  - Professor McCarthy coined the term "artificial intelligence" in 1955
- Machine Learning (ML) = The field of study that gives computers the ability to learn without being explicitly programmed".
  - Arthur Samuel (IBM) coined the term "Machine Learning" in In 1959,
- (Artificial) Neural Networks are a key part of machine learning. Neural Network = a computer system designed to work by classifying information in the same way a human brain does.
- ML is a subset of AI but key area of focus in today's research efforts

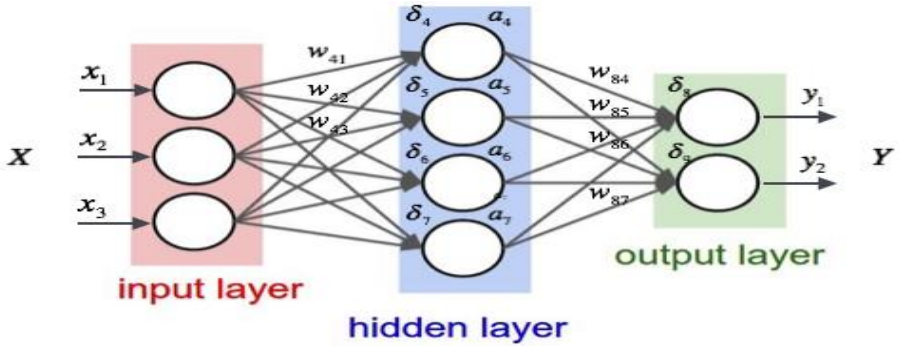

Machine Learning Basics

- Four broad categories of problems ML is being applied to: clustering, classification, regression, and rule extraction
- Typical learning approaches:
  - Supervised or semi supervised
  - Unsupervised (automated groupings in clusters)
  - Reinforcement Learning: agent interacts with the external world and learns by exploring the environment maximize cumulative reward
- Typical machine learning models:
  - Artificial neural networks. More recently Deep Learning number of layers)
  - Support Vector Machines (SVM)
  - Bayesian Networks (fault management)
  - Genetic algorithms

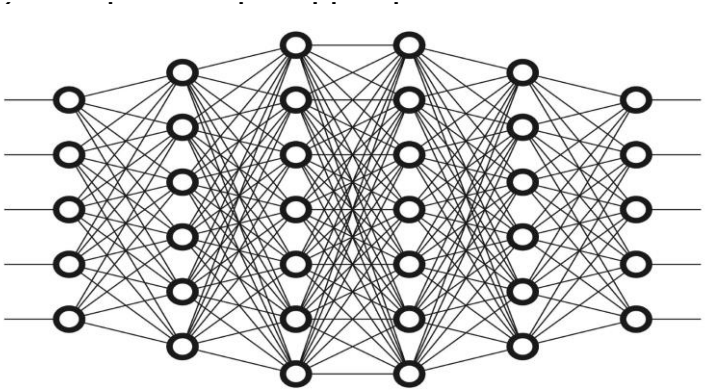

Enablers & Continuing Challenges in Machine Learning

- Explosion in availability of data (critical for ML)
- Significant improvements in ML techniques (e.g. Deep Learning)
- New computing platforms:
  - Cloud, Graphics Processing Units (GPUs), Tensor Processing Units (TPUs) provide accelerated
- However, AI/ML still requires deep expertise for designing, developing, and evolving systems
- Feature extraction
  - Tradeoffs between over-fitting for higher accuracy and lower computational overhead.
  - It is essential to select features that do not contradict underlying assumptions in the context of the problem.
- Dealing with a changing environment
- Creating a robust solution

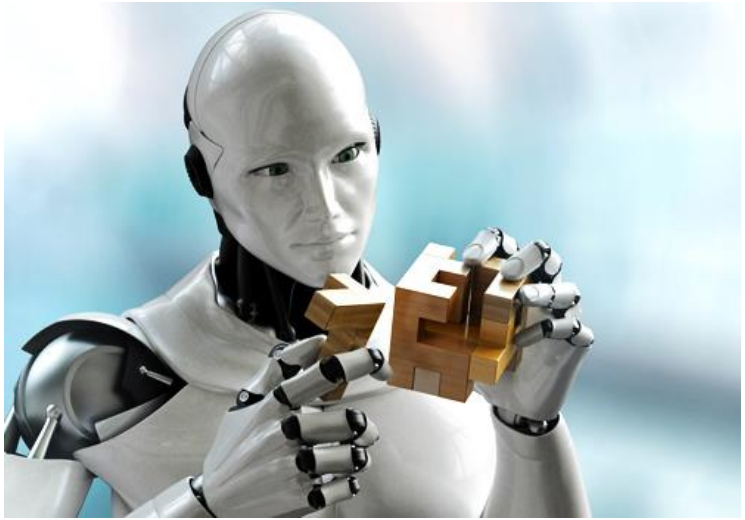

Successful Application of AI/ML

- Jeopardy, chess, various games
- Autonomous vehicles & drones
- Object detection, image recognition (e.g. Facebook)
- NLP, Google Translate, Siri , Alexa, and other assistants
- Recommender systems, online adds (Netflix, Amazon, Google)
- Automated response systems and chatbots
- Robotic Process Automation
- E-mail filters
- Google predictive searches
- Fraud detection
- Common characteristics:
  - Mostly in areas that mistakes are ok. Autonomous platforms an exception.
  - Need a lot of data for training

MIT's Cheetah Robot

What can machines do better than humans

- Performing repeatable and well structured tasks (no mistakes)
- Searching and analyzing large amounts of data (finding patterns)
- Speed of execution
- Remembering things
- Deal with stress, tiredness, emotions
- Making decisions with no biases
- Providing extremely accurate answers

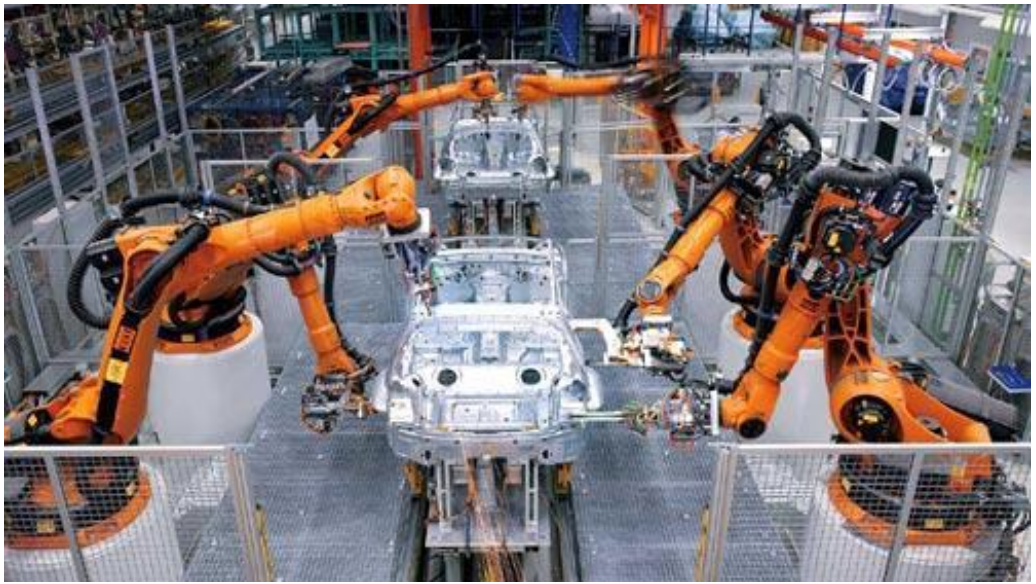

What can machines not do as well (at least as of today)

- Dealing with changing not well structured environments
- Dealing with unpredictability
- Understanding why
- Helping others or managing others
- Understanding hidden meanings
- Explain things
- Being creative
- Demonstrate empathy

Can AI/ML be used in Network Management?

- Large amounts of data available for learning
  - Traffic load, Performance data, Syslog, Trouble tickets, social media messages (e.g. Twitter), Numerical, text, …
- Significant progress in use of AI/ML already in certain areas *
  - Traffic analysis and prediction
  - Traffic classification
  - Resource management & admission control
  - Fault & performance management
  - Network security
- However networks are very diverse. What works in one may not work in another.
- Networks continue to evolve and change
- There are unpredictable events and surprises
- What choice do we have?

perspecta LABS

* Boutaba *et al. Journal of Internet Services and Applications* (2018)

Situational Awareness Of Large Systems Through Machine Learning

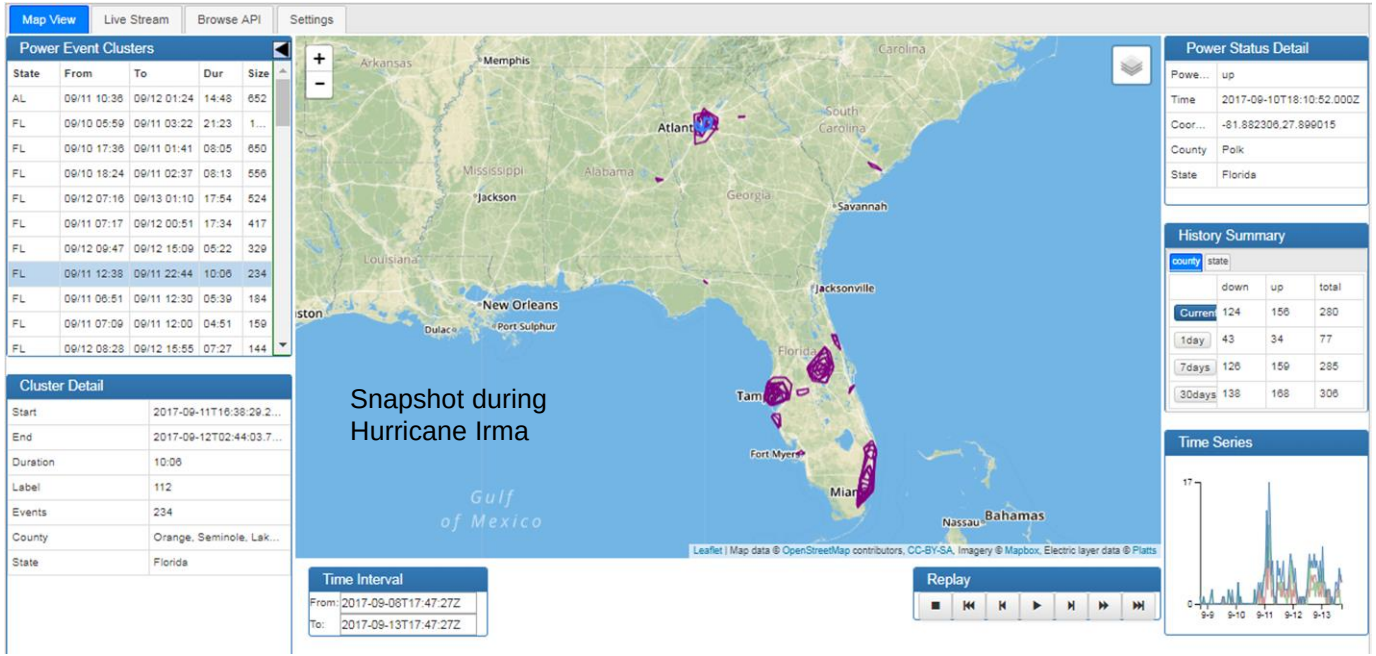

- Perspecta Labs MANTESSA data analytics extracts power grid situational awareness from non-grid sources: social media, Internet measurements, and satellite imagery
- Results show good correspondence to official outage reports and news

Machine Learning for Cyber Security In Action

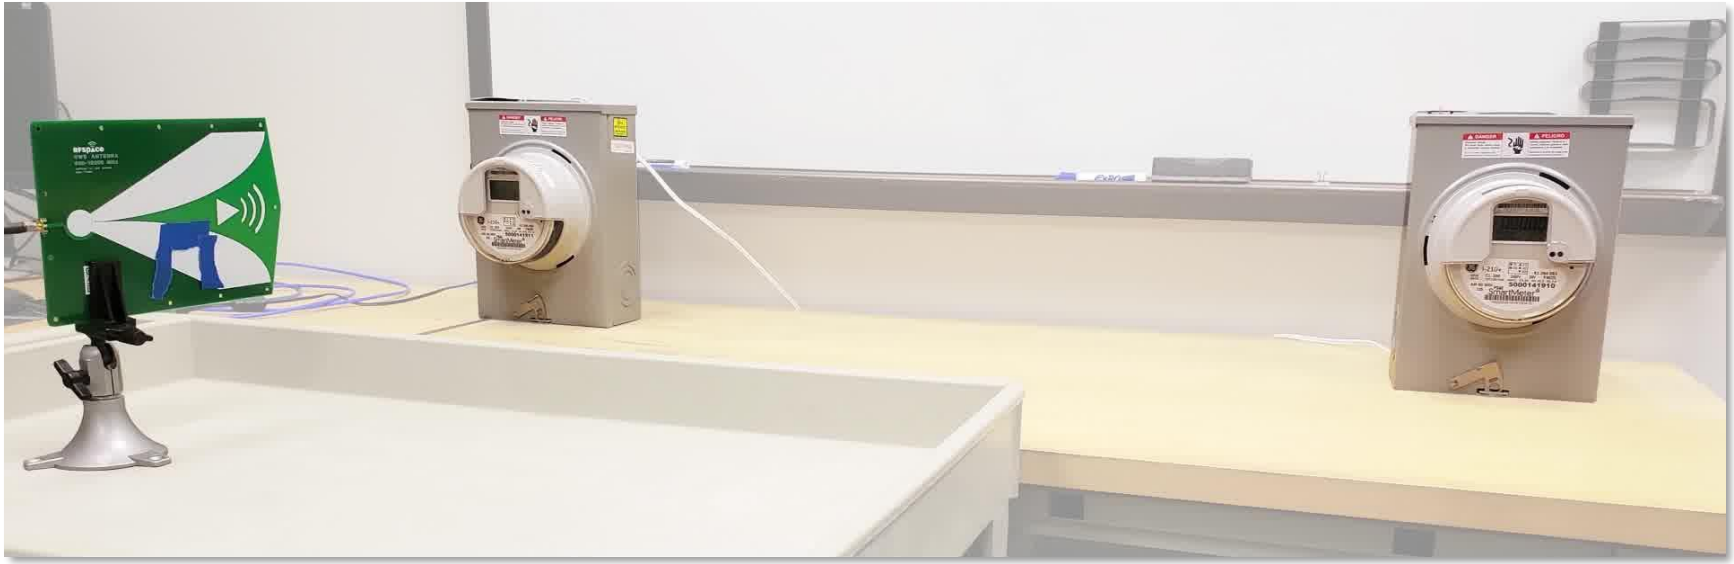

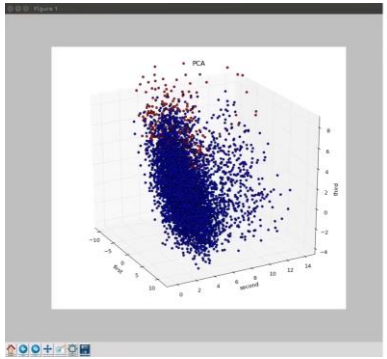

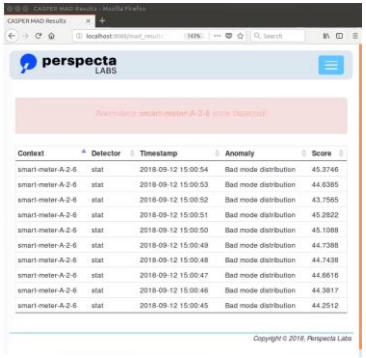

perspecta
LABS

Risks of Using AI/ML

- Accidents are still possible with AI/ML
  - The New York Times reported (2018) that Uber's self-driving technology worked on average for just 13 miles in Arizona before requiring human correction to avoid a crash. Waymo did do much better.
  - There have been 104 collision reports involving self-driving cars in CA since 2014 (Wired 10/2018)
  - There has been a small number of fatalities with self-driving cars
- There are also moral dilemmas to consider:
  - How to program decisions on what the vehicle should do when loss of life is unavoidable
  - There are concerns (and some indications) that AI decisions may still be biased

Adversarial Machine Learning

- Small perturbations (imperceptible to humans) in input data can result in misclassification by ML algorithms
  - Most examples in the space of image classification
    - Examples emerging in other domains, such as audio, text, cyber data

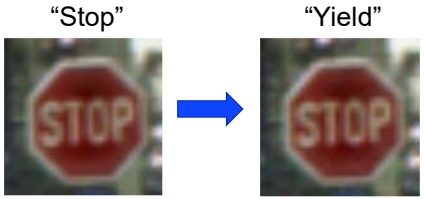

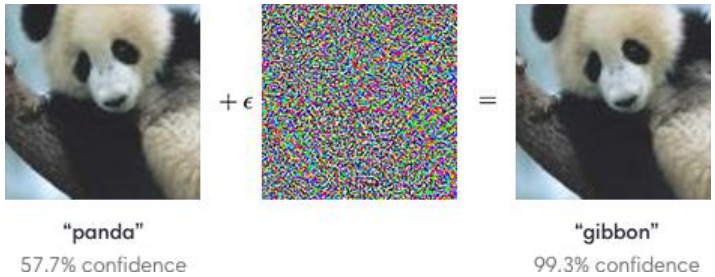

https://blog.openai.com/adversarial-example-research/

Adversarial Examples can hide in music

Carlini et al., Audio Adversarial Examples: Targeted Attacks on Speech-to-Text, DLS Workshop 2018

Yuan et al., CommanderSong: A Systematic Approach for Practical Adversarial Voice Recognition, USENIX Security 2018

Path forward

- Humans and machines need to collaborate with humans in charge (Tesla model)
- We should worry about:
  - Decision-making on key aspects of digital life being ceded to black boxes.
  - Humans need to continue understanding implications of AI discoveries and recommendations
  - Dependence on machine-driven networks may erode people's abilities to think for themselves
- Research needs to continue in key areas:
  - Human machine collaboration
  - Dealing with changing environments (Continuous learning)
  - Machine have to be able to explain decisions (Explainable AI)
  - Adversarial machine learning

Thank you

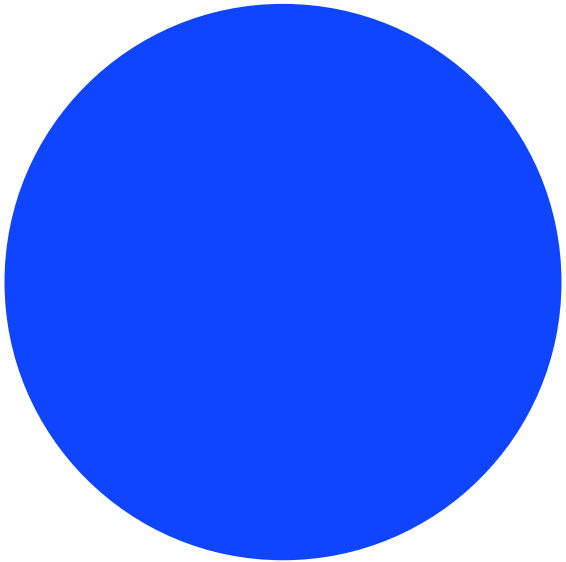

perspecta
LABS