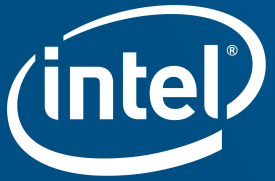

INTEL® HPC DEVELOPER CONFERENCE

FUEL YOUR INSIGHT

USING MACHINE LEARNING TO AVOID THE UNWANTED

Justin Gottschlich, Senior Staff Research Scientist

Intel Corporation

November 2016

Outline

- ***Background***
  - Anomalies
  - Anomaly Detection and Management
    - Challenges
    - Impact and applications
- Research at Intel Labs
  - Inverted time-series DNN
- Future Directions

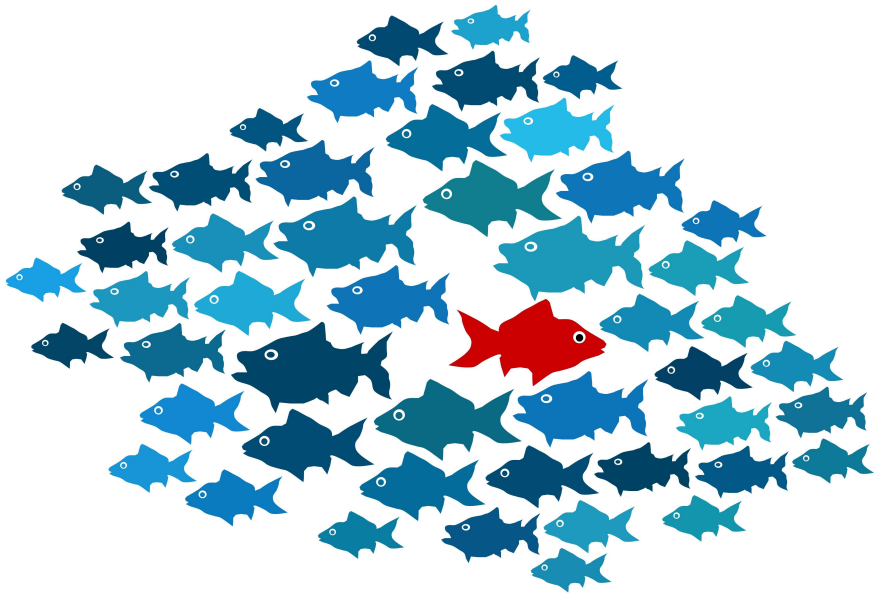

Anomalies

Anomaly - something that deviates from what is standard, normal, or expected.

- Examples of anomalies
  - August 2013: Amazon website 49min outage ($2M, $70K min)
  - January / May 2016: Tesla Autopilot fatalities
  - September 2016: SpaceX rocket explosion
- Impact of anomalies
  - Monetary losses
  - Property damage
  - Loss of life

"...from first signs of an anomaly to loss of data is about ***93 milliseconds*** or less than 1/10th of a second."

http://www.spacex.com/news/2016/09/01/anomaly-updates

Anomalies

- Time
  - Discrete (e.g., power outage)
  - Continuous (e.g., wear on tire tread)
- Impact
  - Efficiency / performance
  - Correctness
- Types
  - Harmful
  - Benign (e.g., robot experiencing novelty, road closure away from route, etc.)

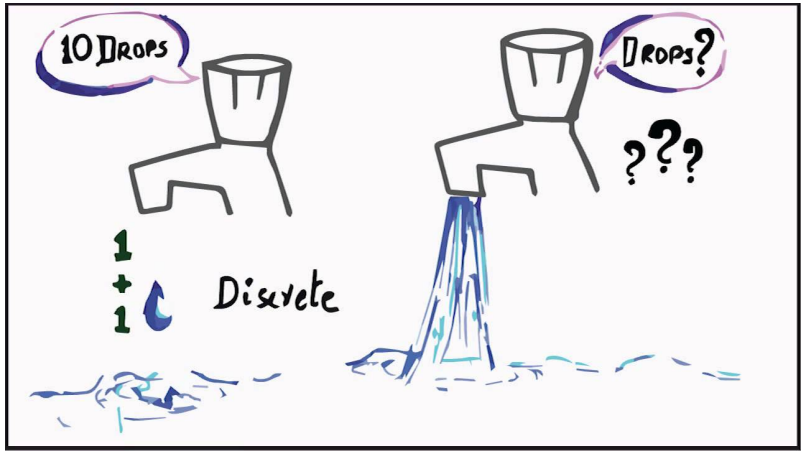

Anomaly Systems

- *Anomaly detection*
  - Identify / predict efficiently & accurately
  - ***Neural networks leading in this space***
- *Anomaly management*
  - Minimize negative side-effects
- Example: car braking assistance
  - Identify potential collision (detection)
  - Apply brakes to avoid it (management)

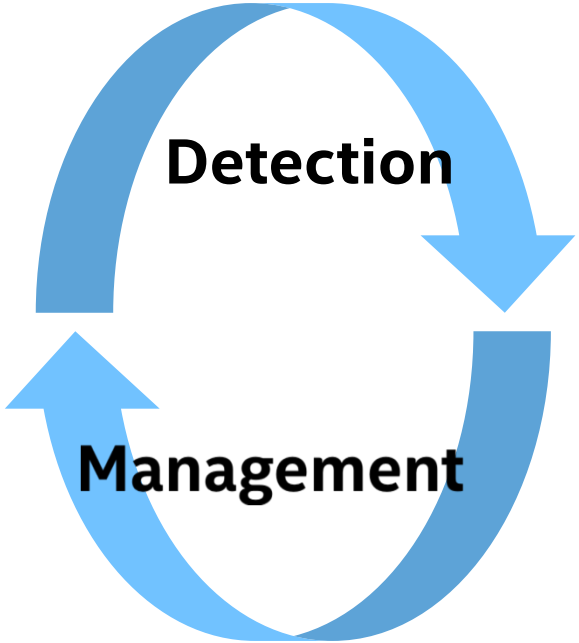

Anomaly Nomenclature

	Positive Prediction	Negative Prediction
Actual Positive	***True Positive*** +	False Negative -
Actual Negative	False Positive +	***True Negative*** -

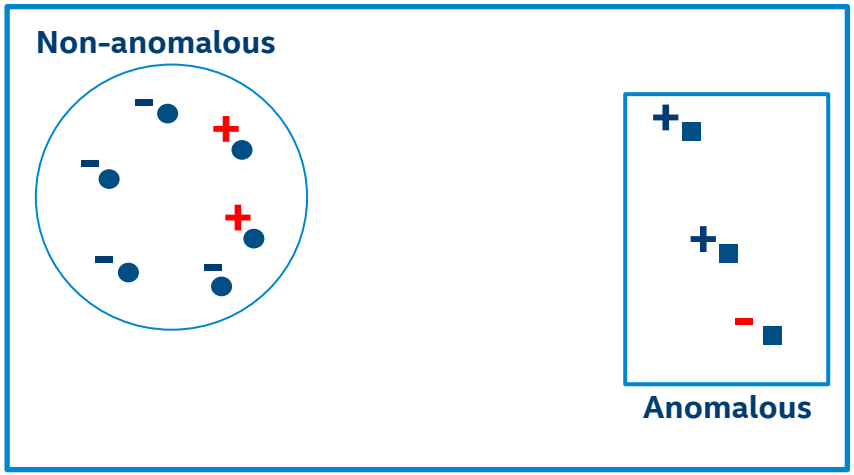

4 true negatives, 2 false positives
2 true positives, 1 false negative

Correct predictions = true positives and true negatives
Incorrect predictions = false positives and false negatives

Anomaly Detection

- Ranking an anomaly detection system
  - **Precision** = true positives / (true positives + false positives)
  - **Recall** = true positives / all actual positives

$$F1 = 2 * \frac{precision * recall}{precision + recall}$$

- Detect new anomalies
  - Rare for all anomalies to be known
    - E.g., software security exploits

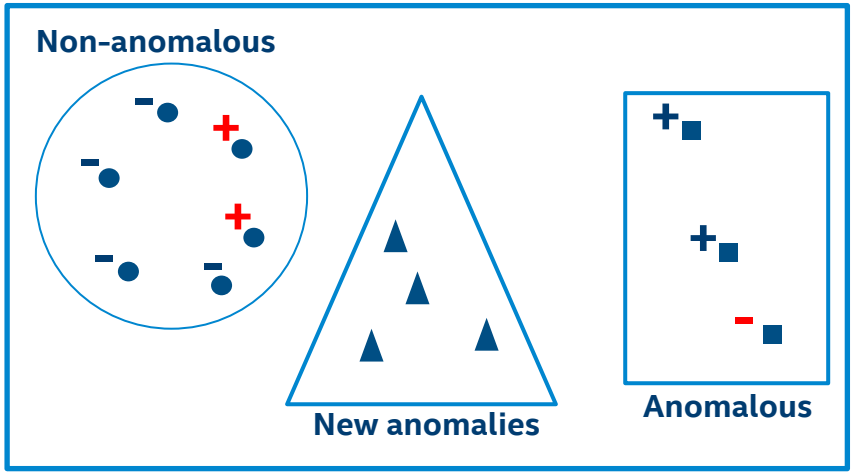

4 true negatives, 2 false positives
2 true positives, 1 false negative

Anomaly Detection

- Identification is (sometimes) insufficient
  - Need predictive time-band
    - Adds new inference component: time-to-anomaly (TTA)

- Detection constraints are domain-specific
  - Detect cancer
    - Prefer false positive to false negative
  - Detect spam/junk mail
    - Prefer false negative to false positive

Anomaly Management

- Mitigation
  - Predict hurricane then evacuate residents
- Avoidance
  - Identify anomaly and act to prevent it
  - E.g., self-driving car to identify / prevent accidents
- Notions of self-preservation
  - ***Ideal systems identify and manage their anomalies***
  - Ethical murkiness; AI ethics board / governance

Challenges

- *Problem: Most anomalies occur over time*
  - Solution: Recurrent neural nets, such as LSTMs

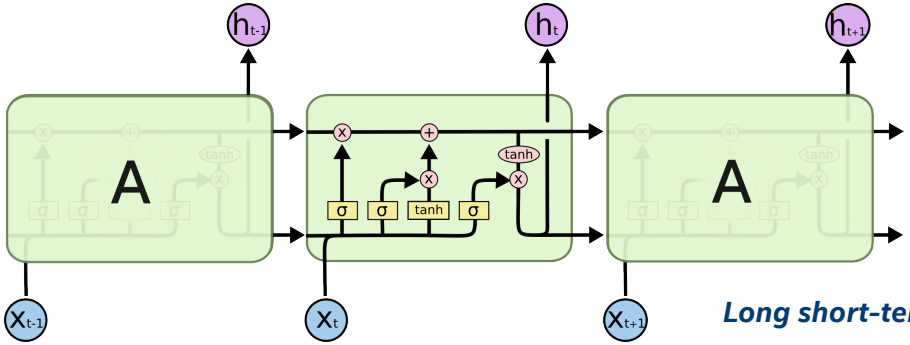

Long short-term memory network cell

(Source: "Understanding LSTM Networks" by Chris Olah, http://colah.github.io/posts/2015-08-Understanding-LSTMs/)

- *Problem: Anomalous data is rare*
  - Solution: Generative models
- *Problem: New anomalies at run-time*
  - Solution: Unsupervised learning

How Does It Impact You?

- Do you …
  - Write software?
  - Build robots?
  - Create algorithms?
  - Run experiments / benchmarks?
  - Optimize code?
- Example
  - ACT, ISCA 2016; identified multithreaded bugs
    - “Production-Run Software Failure Diagnosis via Adaptive Communication Tracking”
  - Collaboration Intel + ACT team: *performance anomalies*

Technical Applications

- Areas of Anomaly Detection
  - Medical / health anomalies
  - Intrusion detection for servers / data centers
  - Malware detection
  - Spacecraft anomalies
  - Autonomous vehicles
  - Software correctness and performance
  - Robotics and self-repair
  - Ambient computing
  - And so on …

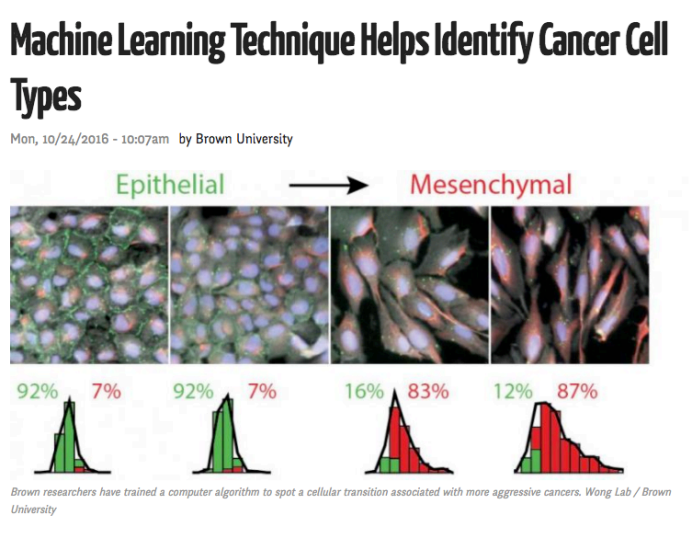

"Another intriguing result from the study was … about 10 percent of cells … defy categorization.

It's possible … this indicates an intermediate cell type that is somewhere between epithelial and mesenchymal."

(Source: Scientific Computing)

Outline

- Background
  - Anomalies
  - Anomaly Detection and Management
    - Challenges
    - Impact and applications
- ***Research at Intel Labs***
  - ***Inverted time-series DNN***
- Future Directions

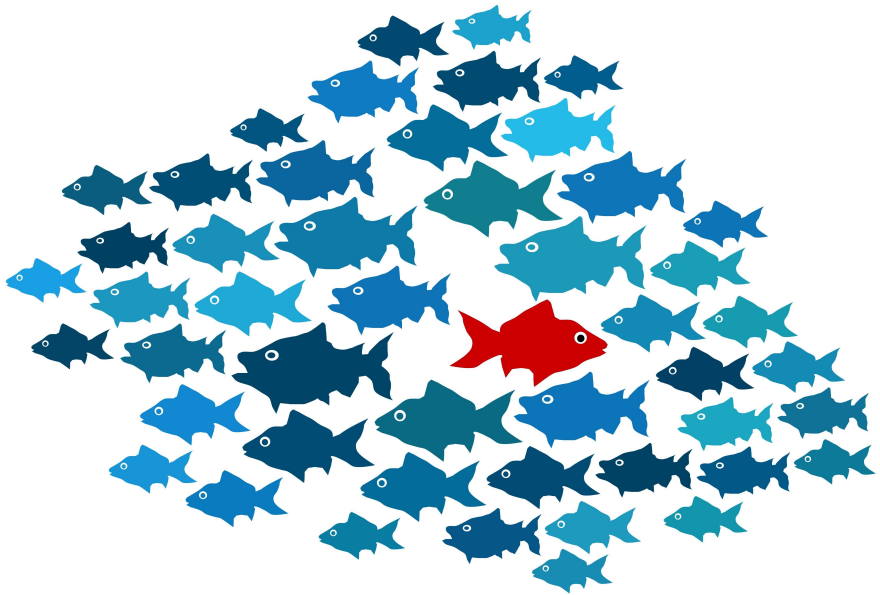

Research at Intel Labs: **High-Level Overview**

- Primary
  - Autonomous vehicles (specifically self-driving cars)
  - Data centers (general and HPC)
  - Anomalous dataset generation and modeling
- Secondary
  - Self-learning systems
  - Ambient computing

Research at Intel Labs: **Self-Driving Cars**

- Detection
  - SDC fleet intelligence to detect outliers
  - Anomalous human cues
- Management
  - Blind guidance system to assist malfunctioning SDCs
- Detection and management
  - Third person perspective to identify and prevent possible collisions

Research at Intel Labs: **Anomaly Detection**

- ***Novel DNN: inverted time-series***
  - Zero positive learning
  - Real-time learning
  - Unsupervised learning
- Dataset generation and modeling

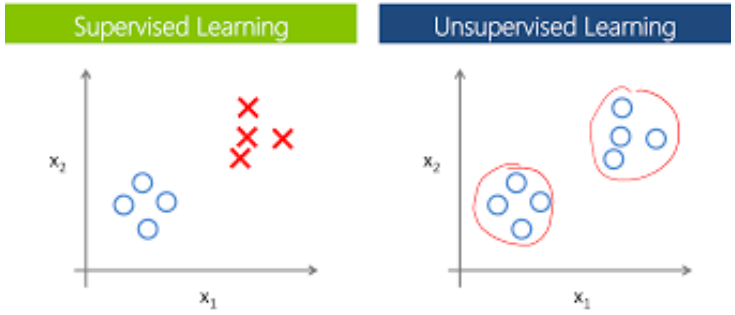

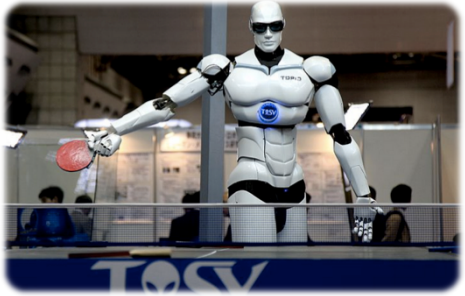

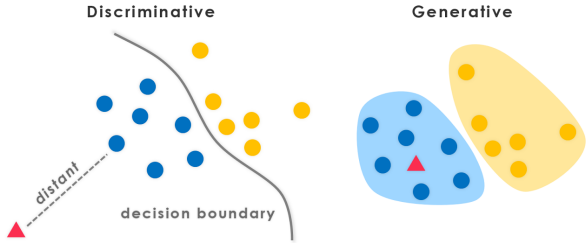

Inverted Time-Series DNN: **Data Adapters**

High-level Intuition:

Flexible data processing from various inputs.

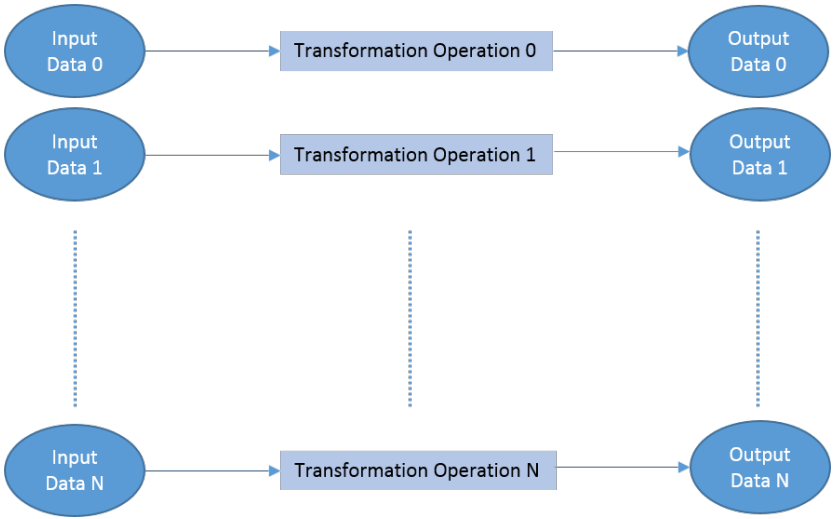

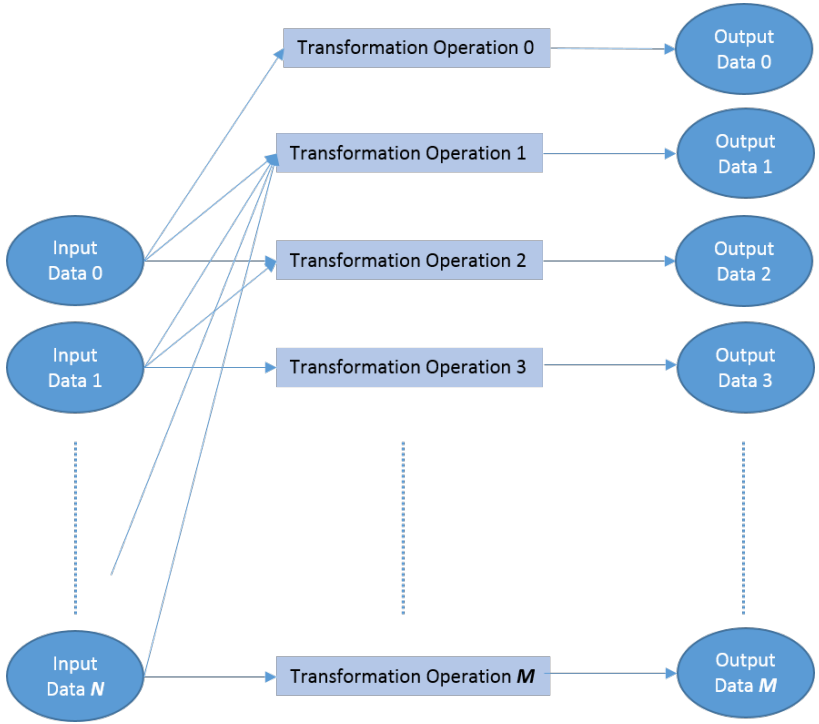

Inverted Time-Series DNN: **Data Aggregator**

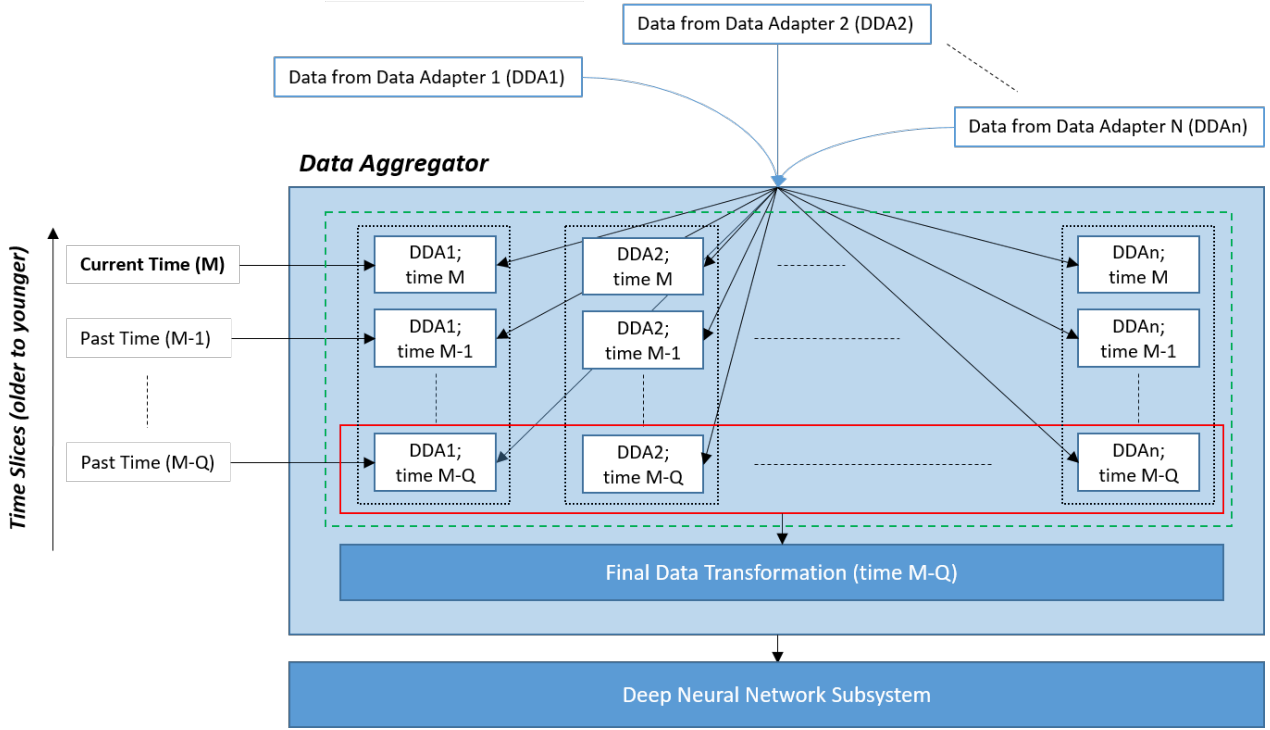

High-level Intuition:

Data buffering and ordering system to ensure proper time-ordered data.

Inverted Time-Series DNN: **Cooperative Learning**

High-level Intuition:

Uses two DNNs that consider data differently for a broader understanding of anomalies.

They also learn from each other, unsupervised.

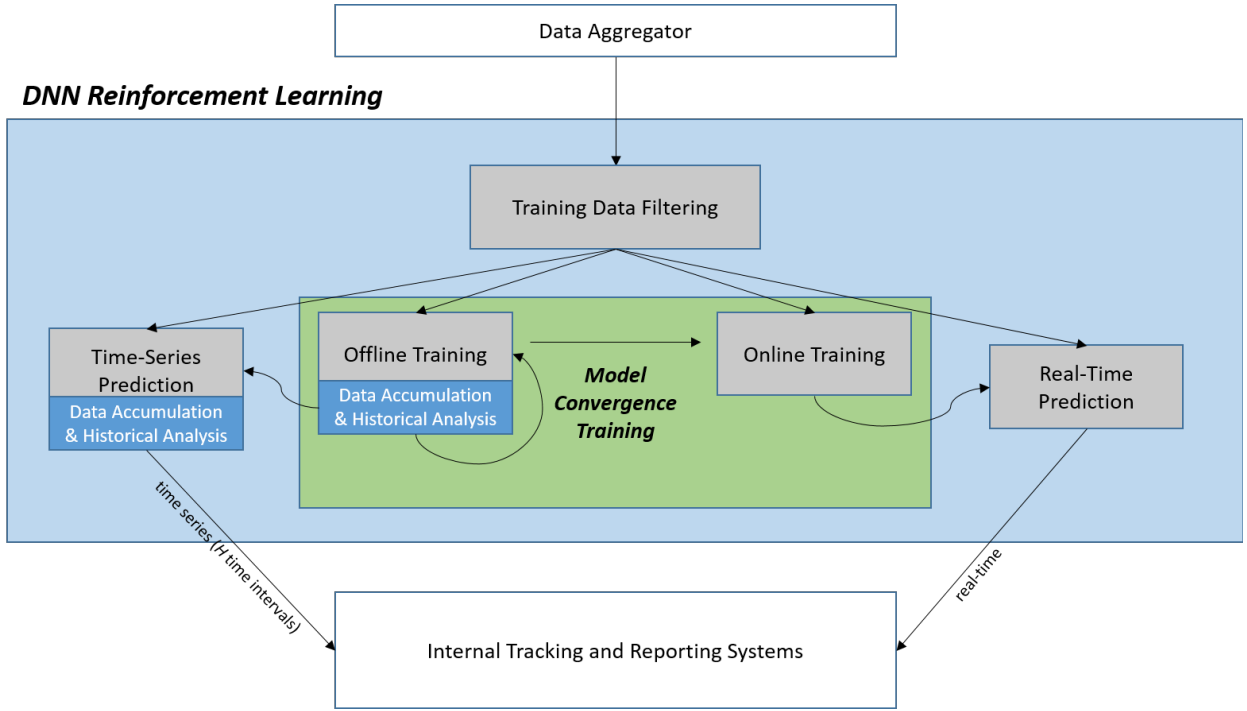

Inverted Time-Series DNN: **Novelties**

- **Inverted time-series predictions**
  - Offline DNN anomaly detection:
    - $Prediction_R = Time_R \ldots Time_{R+N}$

- **Unsupervised, reinforcement training**
  - If offline / online predictions diverge
    - Online DNN is negatively reinforced
  - If offline / online predictions converge
    - Potentially positively reinforce offline DNN (possible overfit)

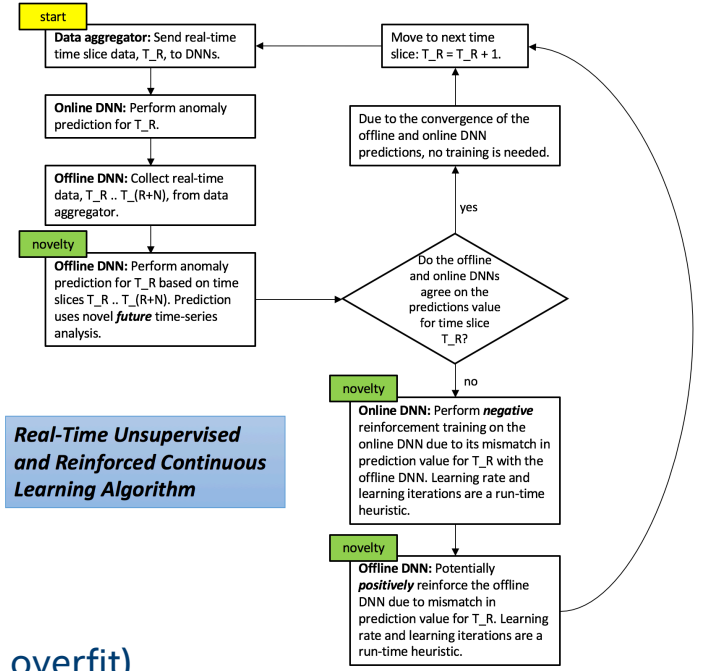

Real-Time Unsupervised and Reinforced Continuous Learning Algorithm

Future Directions

- Anomaly detection
  - Real-time, unsupervised, zero positive learning
    - Identification and prediction of all new anomalies
  - Low power edge device (ambient computing)
  - Big *time-series* data
    - Intel, Brown, MIT
  - Anomalous datasets generation and modeling
    - Intel, Stanford

Future Directions

- Anomaly management
  - HPC & distributing computing opportunities
  - Fleet (swarm) game theory
    - Function as unit to minimize negative impact
  - Cooperative, unsupervised learning
    - Identify novelty
    - Machine-to-machine learning
    - Knowledge sharing

- ***Goal: fully autonomous anomaly detection / management across domains***

INTEL® HPC DEVELOPER CONFERENCE

FUEL YOUR INSIGHT

THANK YOU FOR YOUR TIME

Justin Gottschlich

justin.gottschlich@intel.com

www.intel.com/hpcdevcon

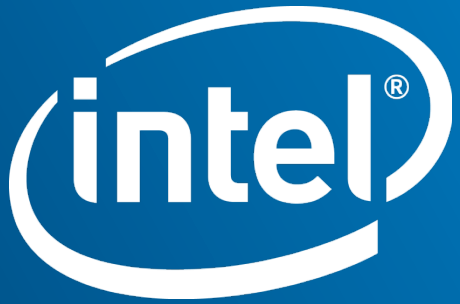
intel®


intel®
experience
what's inside™