

AI and Machine Learning in Campus Networks

Robert Barton, P. Eng

Distinguished Systems Engineer, Cisco



BCNET[→]2019

AI Isn't Exactly What we Thought
it Was Going to be




AI is Changing the World

- In 2016 AlphaGo (Google Deepmind) defeated the world champion Go master, Lee Sedol
- AlphaGo used innovative moves not seen before



AI@Cisco



Reinvent the Network



Security is Foundational



Power a Multicloud World



Unlock the Power of Data



Create Meaningful Experiences



DNA



ETA



Network Early Warning



SD-WAN



AMP



Cloudlock



Cognitive Threat Analytics



NGFW



Stealthwatch



Talos



Umbrella



Hyperflex



Intersight



UCS



AppDynamics



Kinetic



Tetration



Accompany



CAM



Meraki



MindMeld



Talent Trends



Webex

How Cisco Approaches AI/ML



Consumption

Products use AI/ML to do things better




Enablement

Infrastructure Supporting AI/ML workloads

The Power of GPUs for Deep Learning


- Graphical Processing Units are specialized types of electronic circuitry designed to rapidly manipulate memory for graphics
- GPUs support parallel processing, accelerating their ability to execute algorithms that require parallel processes
- GPUs are at the heart of deep learning and neural networks



GPUS HAVE THOUSANDS OF
CORES TO PROCESS PARALLEL
WORKLOADS EFFICIENTLY



CPU
MULTIPLE CORES




GPU
THOUSANDS OF CORES

Comparing CPUs vs. GPUs

- CPUs are capable of almost any task
 - but at a price
- GPUs are highly-specialized processors used to solve complex math problems

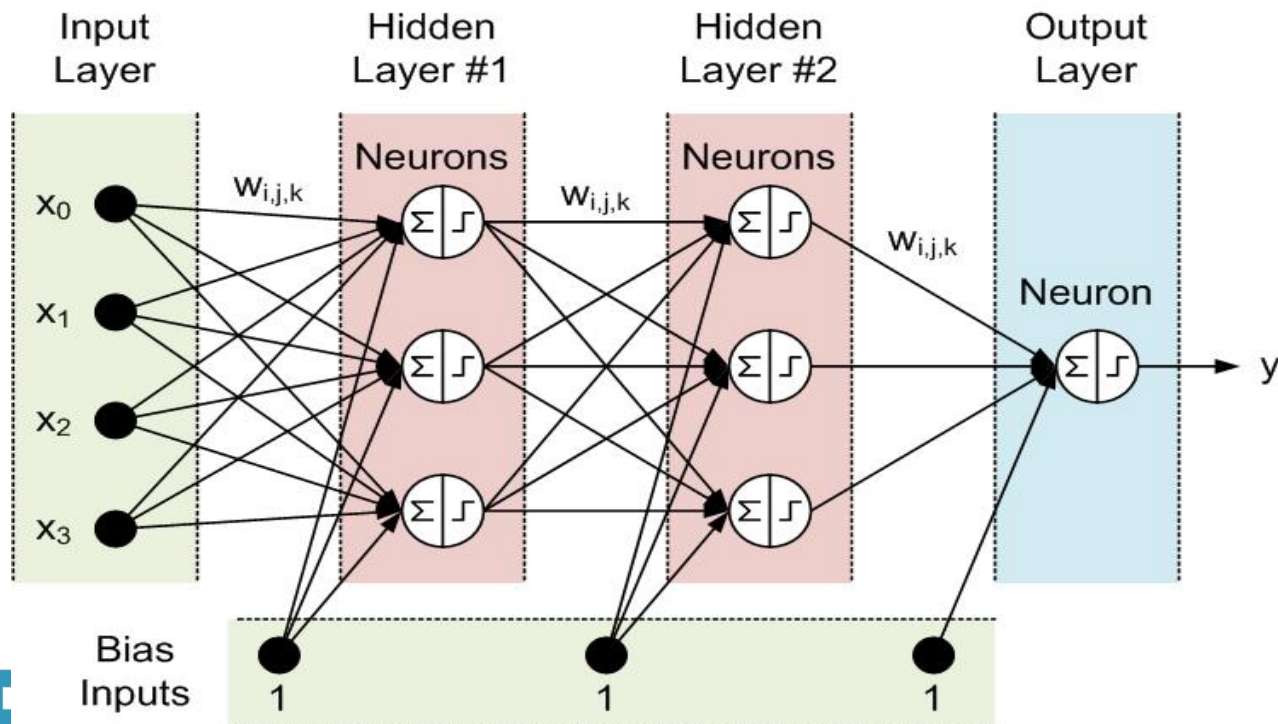
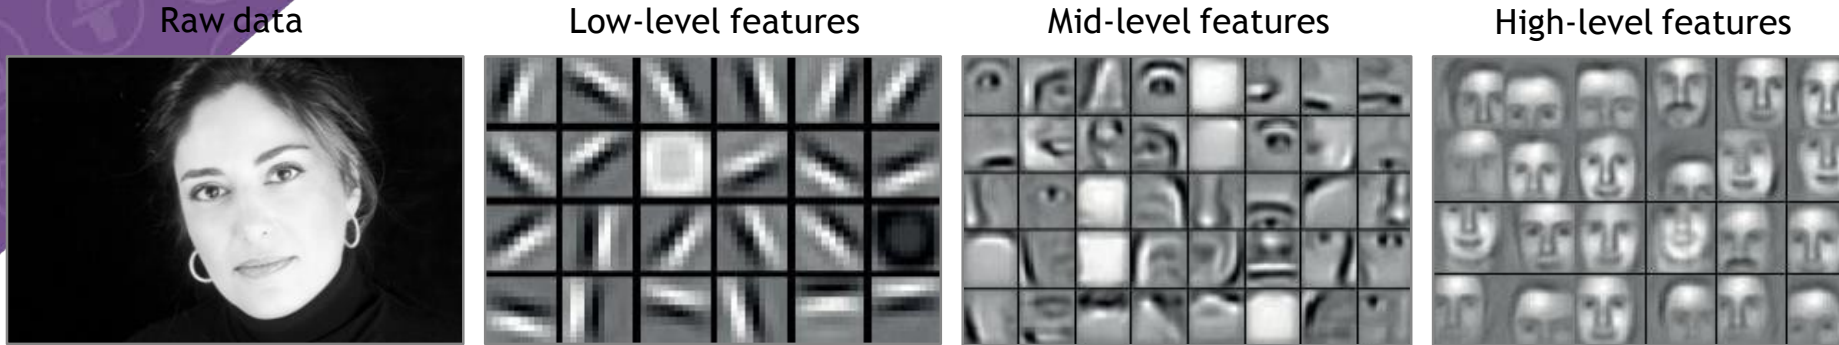


CPUs are like a swiss army knife



GPUs are like specialized surgical instruments

Example: Neural Networks & GPU




	Neural Networks	GPU
Inherently Parallel	✓	✓
Matrix Operations	✓	✓
Bandwidth	✓	✓

Example: Cisco UCS C480 ML



8 X V100 32GB: 1st GPU to break 100 teraflops
NVLink interconnect: >300GB/s bandwidth

RAID controller



Choice of 10/25 or
40/100G
Four PCIe slots
Two 10G Base-T
shared LOMs on I/O
module

Redundant fans



Storage
Up to 24 SAS/SATA SSD/HDDs
Up to 6 NVMe drives
M.2 SATA



CPUs
2 * Intel® Xeon® Scalable processors
(up to 28 cores per socket)
24 DDR4 DIMMs — up to 3 TB memory



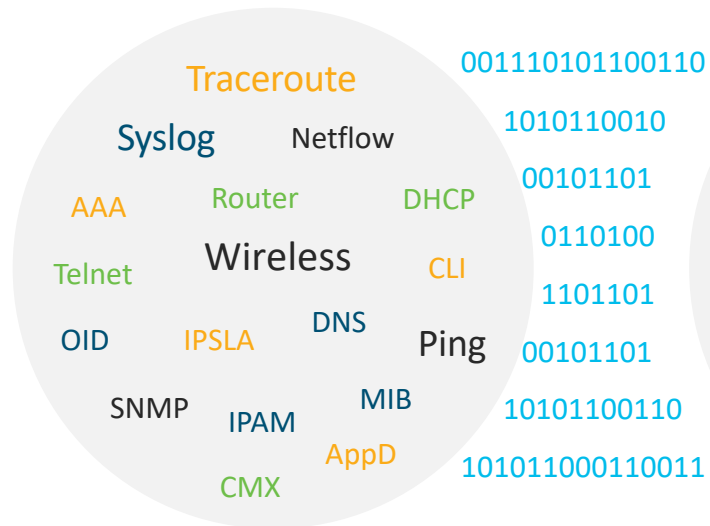
From Network Data to Business Insights

Network Telemetry
Contextual Data

Complex Event
Processing

Correlated Insights

Guided
Remediation



Everything In the Networks is a Sensor




Client | Applications | Wireless | Switching | Routing

Cisco DNA Analytics for Wireless, Wired networks and IoT

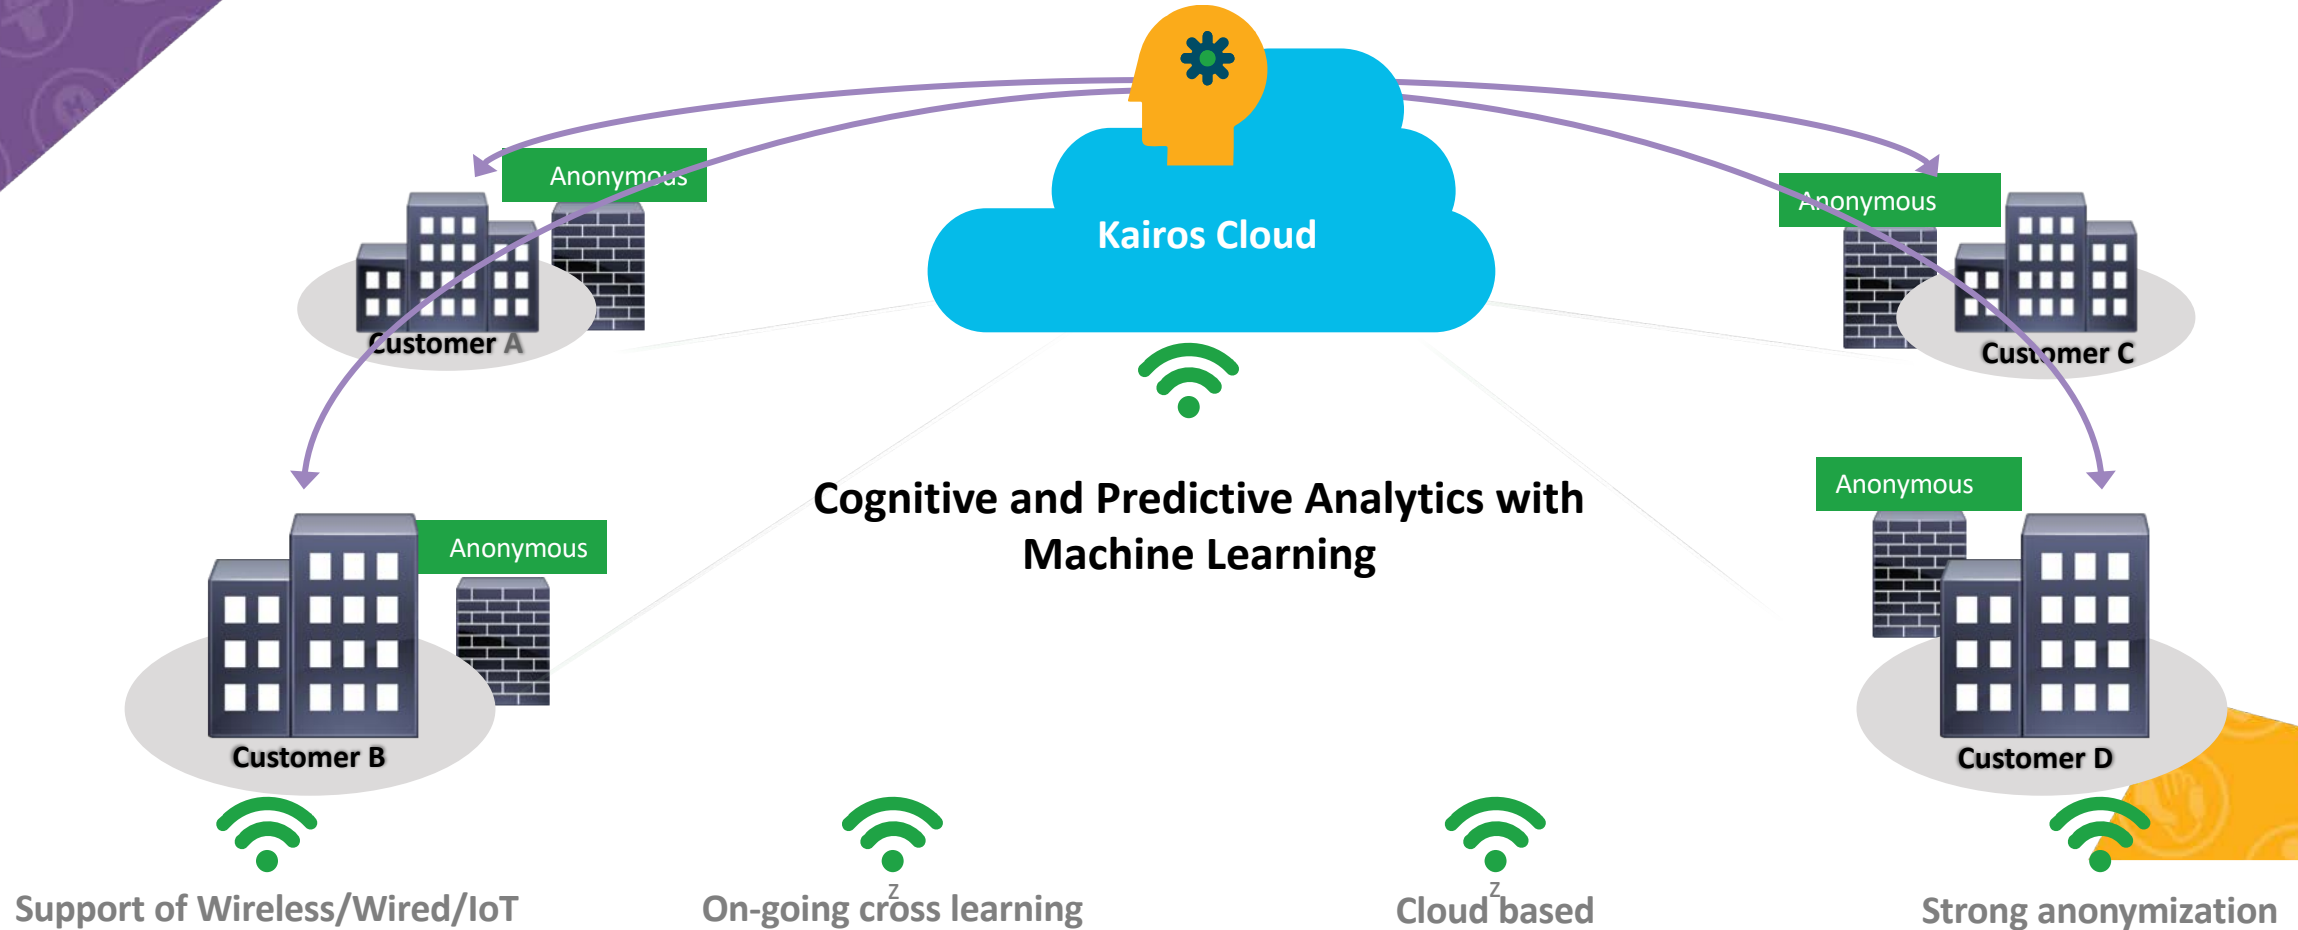


Cognitive Analytics

Anomaly detection across hundreds of
thousands of devices and thousands of
networks



Cisco DNA Analytics Architecture



Example: Radio Performance Issues

Category

Real-time Anomaly Detection

Context

University Customer

Findings

Throughput drops when interference increases as well as 100% of clients have low RSSI and SNR

Root Cause

Coverage issue

Actions

Review the RF design to provide better coverage in this area.



App Throughput – High Packet Retries

Description

APs in network are experiencing a drop in Media Applications throughput. These radios are in the 5GHz band.

Impact of Last Occurrence

Aug 28, 2018 9:30 pm to Aug 28, 2018 10:30 pm



Location:

1 Building



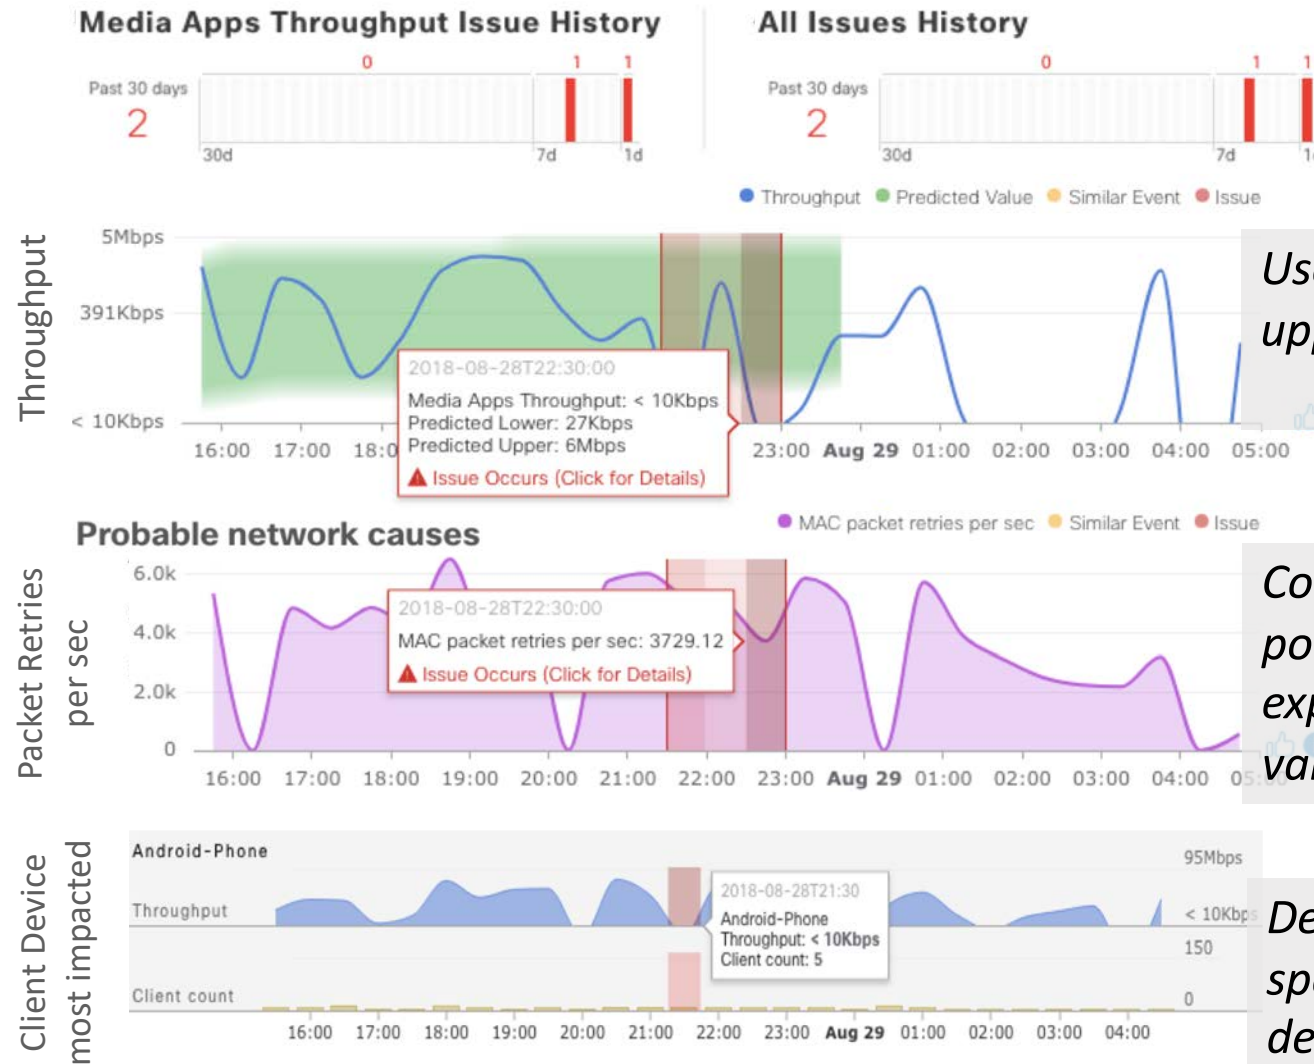
Clients

175 Wireless Clients

Additional Insight

Media Apps Throughput Issues Heatmap

Media Apps Throughput Peer Comparisons



Use regression to predict upper and lower band.

Correlate with other potential issues that are experiencing peaks / valleys in performance.

Determine if issue is specific to any specific devices.


AI-Powered Workspace and Collaboration



Webex Assistant




Face recognition



In-Room Analytics




Noise suppression



Speaker Tracking



Best view



Bots

Webex Endpoints Built on Powerful AI (Nvidia Jetson)



Webex Board 70



Webex Board 55




Room 70D



Room 70S



Room 55




Room Kit Plus



Room Kit

NVIDIA Jetson Platform - The same electronics engine powering self-driving cars




Audio distance: 1.90 Quality: 0.58 Far end: 0

F: 0.0% T: 86.4% U: 0.0% N: 0.0% S: 235

People count: 6

AI that detects you



AI that Recognizes you



John

Alicia

Addy

Rui

Andre

Dora

Basics of Signal Analysis



.wav file



Acoustical Energy



Electrical Energy




We can convert this into an image analysis problem

Classification: From Signals to Images


Voice

Wave Plot




Dog

Wave Plot




Siren

Wave Plot




Tapping


Wave Plot




Mel spectrogram




Mel spectrogram



Mel spectrogram



Mel spectrogram



Voices and “noise” have a distinct “image” that can be detected and filtered.

Encrypted Traffic Analytics

Malware Detection and Visibility without Decryption



Malware in Encrypted Traffic

Is the payload within the TLS session malicious?

- End to end confidentiality
- Channel integrity during inspection
- Adapts with encryption standards



Cryptographic Compliance

How much of my digital business uses strong encryption?

- Audit for TLS policy violations
- Passive detection of Ciphersuite vulnerabilities



Deeper Data + More Context = Better Security

We analyze massive amounts of data —
20 billion threats blocked daily



600 Billion
Email samples



16 Billion
Web requests



3.4 Billion
AMP queries

Our real-time datasets
are diverse, global, and
live



100 Billion
Internet
requests



80 Million
Daily active
users



12,000
Enterprise
customers



160
Countries
worldwide

Detecting Malware by Behavioral Analytics (ML/AI)

Google Search (good)




Initial Page Load

Autocomplete

Page Refresh

Bestafera (bad)



Self-Signed Certificate

Data Exfiltration

C2 Message

The Power of AI/ML in the Network



- **Anomaly detection**
 - Dynamic network performance at different times for different network conditions
 - Static thresholds (even if configurable) would likely raise many false positives or miss relevant events
- **Root cause analysis**
 - Automatic selection of relevant KPIs explaining an issue
 - Cross-correlation across multiple devices
- **Long-term trending**
 - Automatically identifying trends and behavior changes

The background of the slide features a dark blue area on the left, which is filled with a repeating pattern of circular icons. These icons include a smartphone with a gear, a person at a podium, a person with a stethoscope, a handshake, a location pin with an 'H', and a share symbol. A diagonal white line separates this blue area from a white area on the right.

Thank You!

BCNET[→]2019

In the bottom right corner, there are two overlapping geometric shapes, a yellow one on top of an orange one. Both shapes contain a pattern of the same circular icons seen in the blue background area.