Applications of Machine Learning and Intelligent Algorithms for SDN and NFV

May 16, 2017

Daniel King, University of Lancaster

Intelligent Networking
Next generation network challenges

- Convergence means the integration of these services over a common infrastructure and provision through a single point of attachment
  - Combining compute and connection resources for end-to-end services
- The customer expects:
  - Rapid delivery of new services
  - Greater bandwidth and scalability
  - Higher QoE and More sophisticated SLAs
- The provider needs to:
  - Drive up income from deployed resources
  - Find a way to deliver QoE and meet SLAs
  - Deploy agile management and controller platforms
  - Whilst reducing operational costs

Intelligent Networking
Solving current network problems

- Throwing resources at the problem?
  - A guaranteed fat pipe or over provisioning of compute resources is a good way to deliver quality
- But. Exceeding bandwidth and resource requirements will be expensive and impractical in the long-term
  - Inevitably, even in a lightly used network, some links reach critical utilisation
  - It can be hard to predict which links or compute resources will be affected in failure scenarios
  - New customers can cause unforeseen congestion points
- Better network planning and appropriate reoptimisation of services
  - Requires complex path computation capabilities
  - Model the entire network (multi-layer modelling)
  - Consider all current services and compute in parallel not sequentially
  - Respond to network events and deliver services in real-time
  - Requires online path computation capabilities

Intelligent Networking
Solving current network problems

- Support of complex transport services:
  - P2P and P2MP based service types
  - High levels of QoS demand multiple constraints
  - Minimal cost, minimal delay, high bandwidth
  - Constraints may be conflicting
- Multiple resources to support one service
  - Compute, storage, function and connectivity
- Resource continuity issues
  - Multi-layer networking
  - Non-Linear effects and Wavelength continuity
- Path diversity or congruence:
  - Mesh protection resource sharing
  - m:n protection
- Concurrent network-wide optimisation and frequent reoptimisation

Intelligent Networking
Network Operations

- Questions Operators asked themselves…
  - Where is my traffic flowing today?
  - Where do I place new resources, such as links, switches and functions?
  - What resource capacities do I require?
  - How do I design my network to minimise or negate the impact of resource failures?
  - What configuration metrics do I place on the network equipment that will influence traffic flows and quality of service?
  - Where is the most cost-effective place to add new resources to accommodate anticipated traffic growth?
  - What is the most effect mechanism for carrying new types of services?
  - Which protection mechanism is most effective for network topology and service types I currently have?
  - What if…?

Intelligent Networking
Existing computation techniques

- Single-service computations
  - CSPF is perfectly functional
    - Optimal paths for single LSPs with multiple constraints
  - Modified CSPF can compute multiple paths
    - Good for solving k-disjoint paths
  - Conventionally used to satisfy real-time requirements
- Linear programming can optimise a whole network
  - Can take long periods to develop
  - Not flexible to changing demands, new topologies, new constraints, or new service types
- But can it do it fast enough?
  - More constraints mean slower computation times
  - More paths mean more complex computation
  - Increasingly we see competing constraints
  - Larger networks are phenomenally complicated

Intelligent Networking
Comparing techniques

- Conventional algorithms are deterministic
  - Same solution every time
  - Normally tuned to the specific technologies and services
  - Not good at handling multiple service types
  - Generally slow when handling large networks with many elements
- Non-heuristic processes assess a wider variety of data to derive solutions
  - May produce a different, but correct solution each time
  - Is able to handle a variety of topologies
  - Would be capable of managing different service types

Intelligent Networking
What is Machine Learning?

- The complexity in traditional computer programming is in the code (programs that people write).
- In machine learning, learning algorithms are in principle simple and the complexity (structure) is in the data.
  - The way we automatically learn that structure is the heart of machine learning.
- A trained learning algorithm (e.g., neural network, boosting, decision tree, SVM, …) is highly flexible, capable of solving complex problems

Traditional Programming

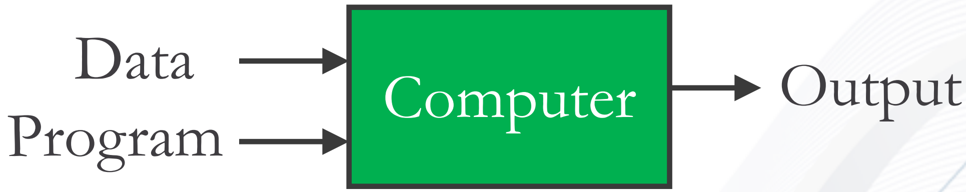

Machine Learning

Intelligent Networking
When to use Machine Learning?

- When patterns exists in our data
  - Even if we don't know what they are
    - Or perhaps especially when we don't know what they are
    - Or if they are just noise
- We can not pin down the functional relationships mathematically
  - Else we would just code up the algorithm
  - Neural networks as function approximators
    - Need this for scale
- When we have lots of (unlabeled) data
  - Labeled training sets harder to come by
  - Data is of high-dimension
    - High dimension "features"
    - For example, sensor data
  - Want to "discover" lower-dimension representations
    - Dimension reduction

Intelligent Networking

Are we in the next Hype Cycle?

Wait..

Are we are here (again)?

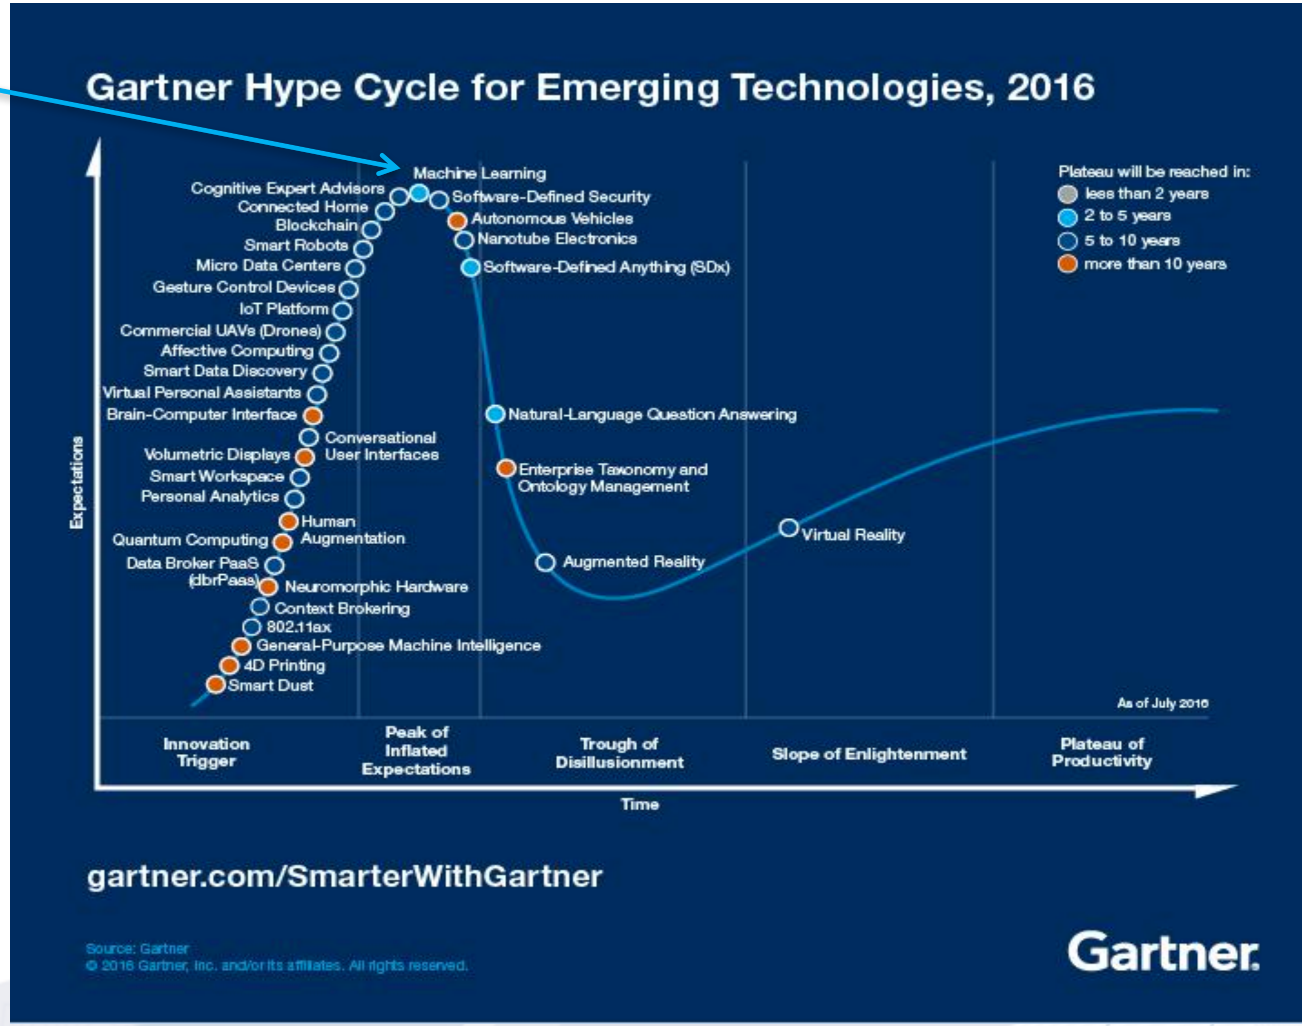

Intelligent Networking
Next Steps for applying ML to Networking?

- Lots of open source ML frameworks, Cloud APIs, Examples, …
  - Google Tensorflow (tensorflow.org)
  - Facebook Torch (torch.ch)
  - Microsoft CNTK (github.com/Microsoft/CNTK)
  - Amazon MXNET (aws.amazon.com/mxnet)
  - Keras (keras.io/)
  - Theano (github.com/Theano)
- Lots of Network Data, but…
  - Standardized and labelled datasets are scarce
  - Most network data sources (e.g., NETFLOW) not designed for ML
  - Arbitrary data might be over-valued
- Skills gap persists , but… it still requires skill/experience to
  - Build DNN architectures/models
  - Finding ”good” settings for hyper-parameters

Intelligent Networking
Join IETF IDNet (Intelligence-Defined Network)

- The IDNet (Intelligence-Defined Network) is aiming to apply Machine Learning mechanisms network environment and react to dynamic situations.
  – https://www.ietf.org/mailman/listinfo/idnet
- Topics include, but not limited:
  – Better management & control of network technologies
  – Architecture or reference model for the Intelligence-Defined Network
  – Integrate IDN with various network infrastructure architectures and IETF standards
  – Data requirements for AI network controlling, including new measurement technologies
  – Network data selection, data structure & protocol for data transmission