

Knowledge-Defined Networking

Learning how to route

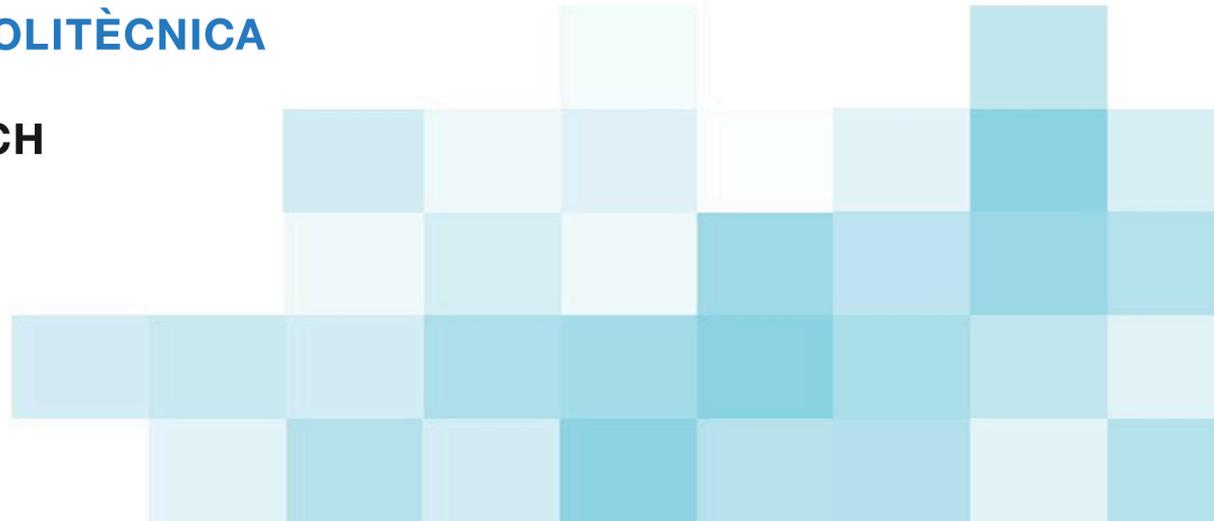
Albert Cabellos (UPC/BarcelonaTech, Spain)
albert.cabellos@gmail.com



UNIVERSITAT POLITÈCNICA
DE CATALUNYA
BARCELONATECH

EUCNC

Athens, June 2016



Thanks to:

- Prof. Jean Walrand
- Fabio Maino, John Evans, Chris Cassar, Hugo Latapie,
- Shyam Parekh
- David Meyer
- Sharon Barkai
- Mike J Hibbett, Giovanni Estrada
- Albert Mestres, Josep Carner, Alberto Rodriguez, Eduard Alarcón, Pere Barlet
- Victor Muntés, Marc Sole



Contextualization

Applying Machine Learning to Networks

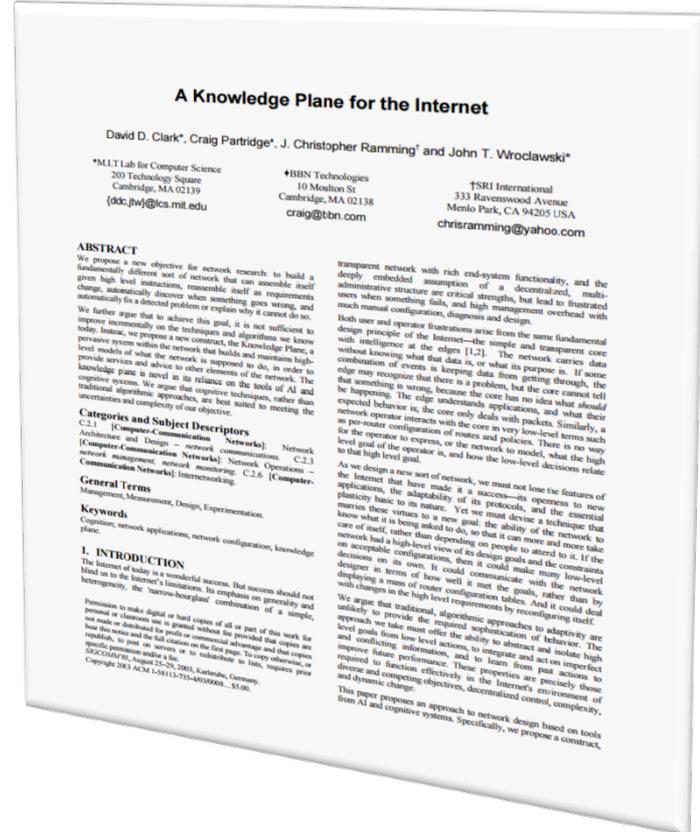


D. Clark (MIT) “A Knowledge Plane for the Internet”, 2003

“we propose a new construct, the Knowledge Plane, a pervasive system within the network that builds and maintains high-level models of what the network is supposed to do”

“The knowledge plane is novel in its reliance on the tools of AI and cognitive systems.”

Clark, David D., et al. "A knowledge plane for the internet." *Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*. ACM, 2003.



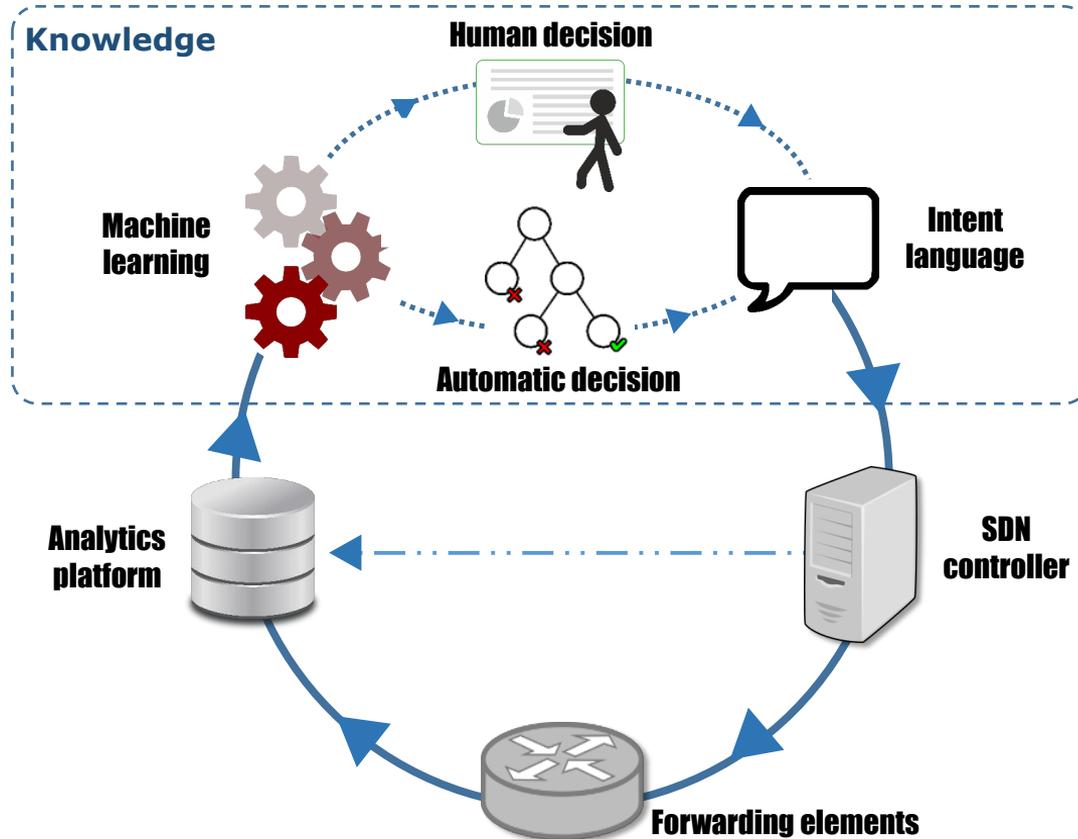
Why we are not there?

- Traditionally networks have been **distributed** systems
 - Partial view and control
- Beyond programmability, SDN provides **centralization**:
 - **Full control** over the network
- Data-Plane nodes are now equipped with computing and storage capabilities
 - Network telemetry and analytics
 - **Rich view** of the network

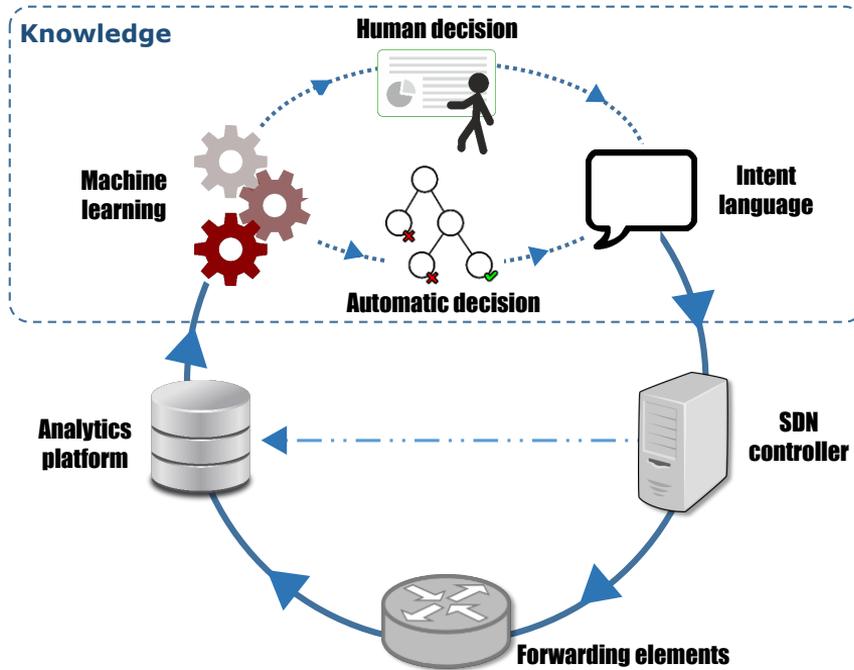
Knowledge-Defined Networking

- Apply ML techniques to Networking:
 - Control (fast dynamics)
 - E.g, routing, resource allocation (NFV/SFC), PCE, optimization, congestion detection
 - Management (slow dynamics)
 - E.g., network planning, resource management, load estimation
 - Recommendation mechanisms
- Towards **self-driving networks**
- **Knowledge-Defined Networking** paradigm

Knowledge-Defined Networking Paradigm



Benefits of KDN



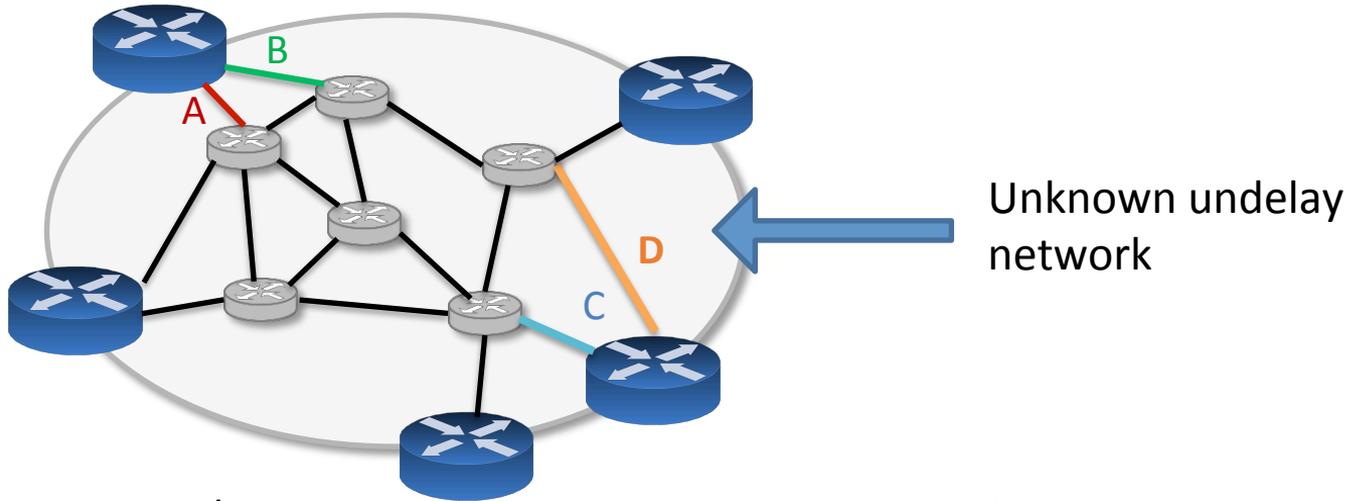
- Recommendation
- Optimization
 - Hidden Information
 - Complex systems
- Estimation
 - Performance/Cost
- Validation
 - Performance/Cost
- Knowledge discovery

Motivatation

Can we learn how to route?

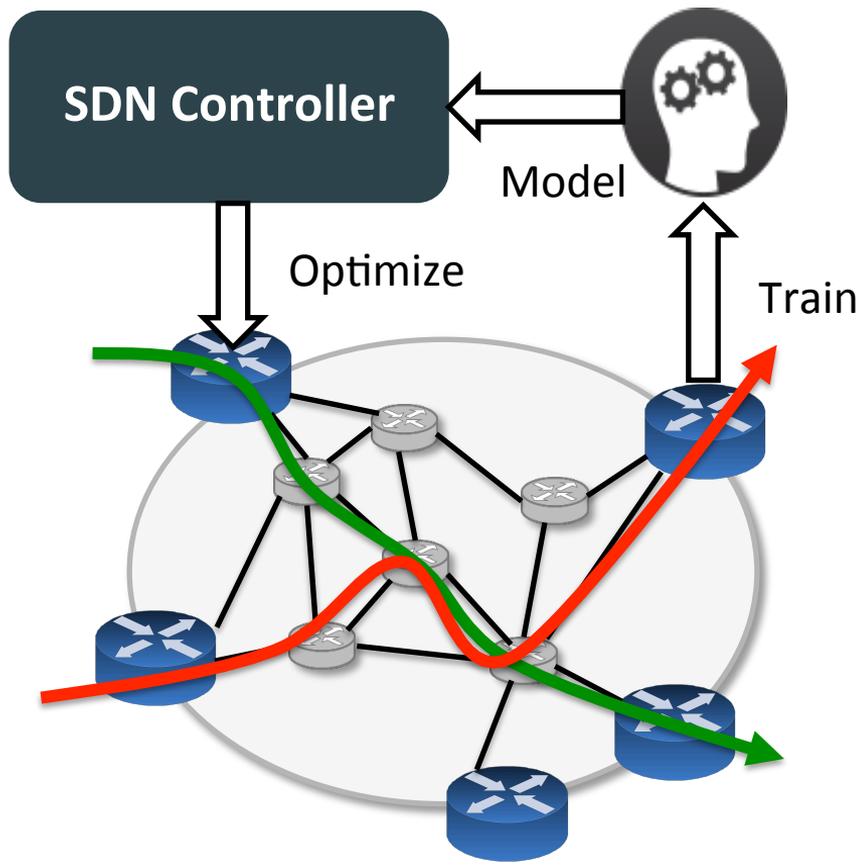


Can we learn how to route?



- Which egress/ingress links should overlay routers use? E.g. **A** or **B** and **C** or **D**?
 - Underlay is assumed that has an arbitrary constant routing
 - Underlay is assumed as hidden and out-of-control
 - Overlay protocol is assumed to be able to choose egress and ingress links, we refer to this as routing policy
- **Goal: Achieve overall minimum latency**

Can we learn how to route?



- Train
 - Ingress/Egress policy
 - Traffic (source, destination, bandwidth)
 - Resulting performance: delay
- Generate a model
 - $f(\text{ingress/egress policy, traffic}) = \text{delay}$
- Optimize
 - Pick, for a given traffic matrix and for each blue node, an ingress/egress link configuration that minimized the delay

Experimental Setup



Is it feasible to learn how to route?

Methodology

- Understand the accuracy of ML-based regressors under various network parameters
- Train a set of ML-based estimators (NN, SVM, etc)
 - $f(\text{ingress/egress policy, traffic}) = \text{delay}$
 - Try to find the optimal performance of the regressors (search over meta-parameters)
 - Datasets: 10.000 samples
 - Cross-validation (60% training, 40% evaluation)
- Evaluate its accuracy when varying different network parameters
 - Size, active stations, routing, etc

Training Set: Packet-Level Simulator: Omnet++

Parameter	Variation
Topology	Star, Ring and Scale-free
Traffic distribution	Poisson, Binomial, Uniform and Deterministic
Size of the network	3-15
Active Stations	3-15
Underlay routing policy	10 (random variations of traffic sent through each path)
Link Saturation	4 levels, level 3 means that at least 1 link is saturated

Regressors

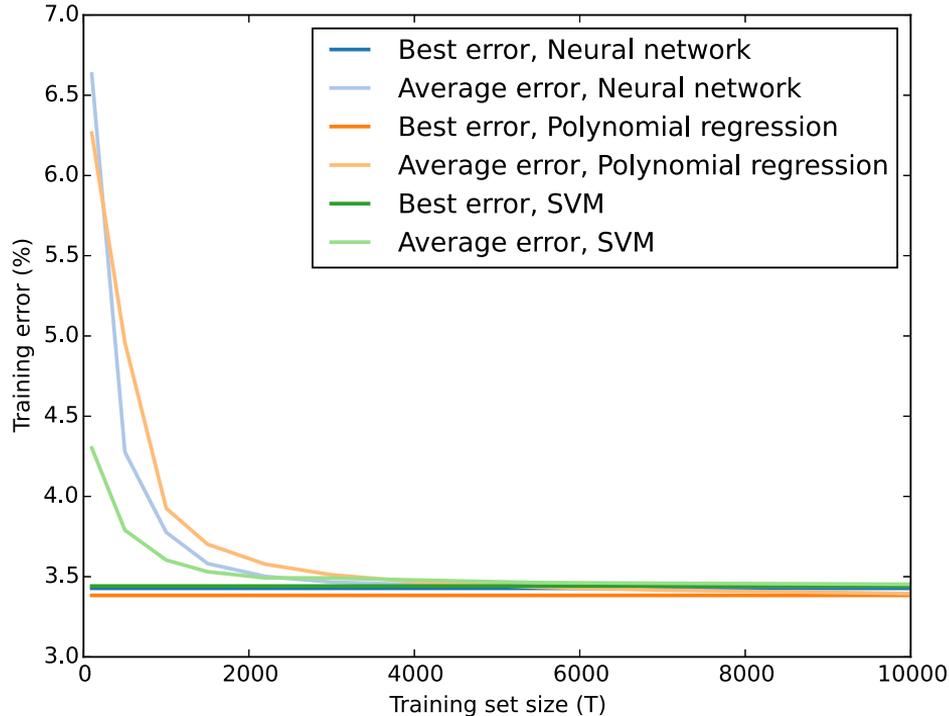
- Single-layer Neural Network
 - We iterate over sizes: 3-200
 - Activation functions: sigmoid, rectified linear unit, hypervolic tangent
- Polynomial regression
 - Linear search of the degree: 1-20
- Support Vector Machine
 - C parameter randomly chosen between 10^{-6} and 100
 - Kernels: Polynomial, Radial Basis Function and Logistic

Experimental Results



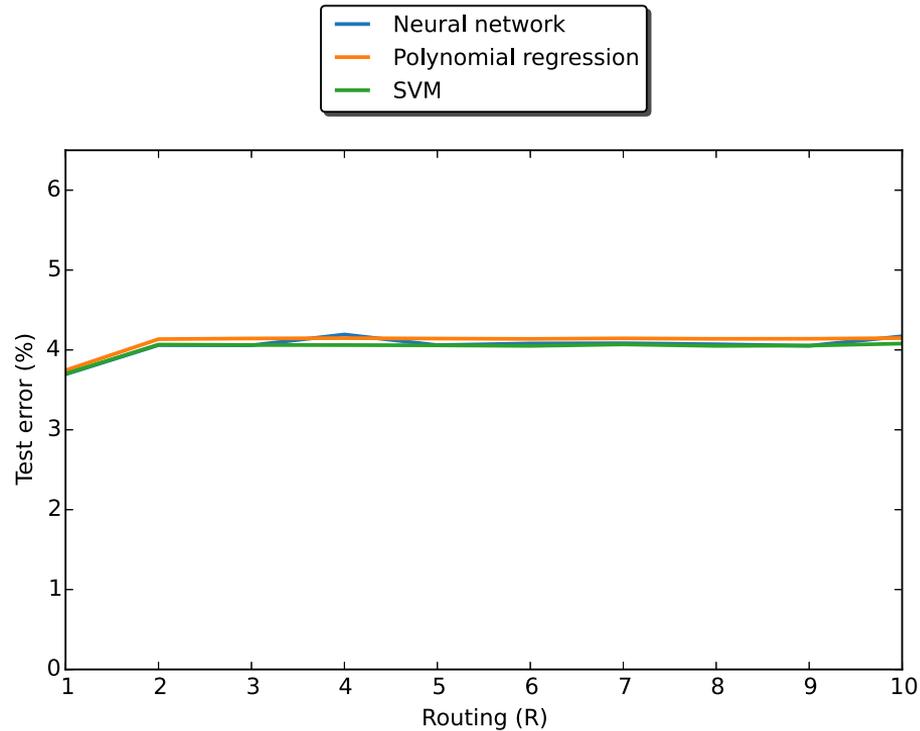
MSE vs. Training set size

(scale-free, poisson traffic, 9 active stations)



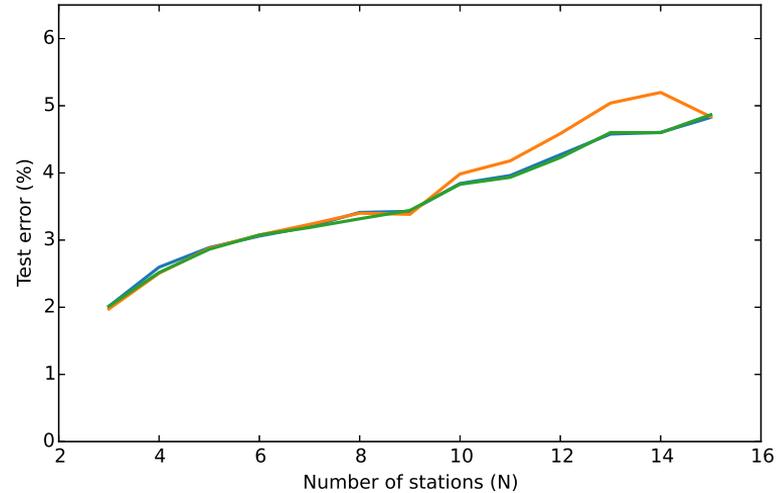
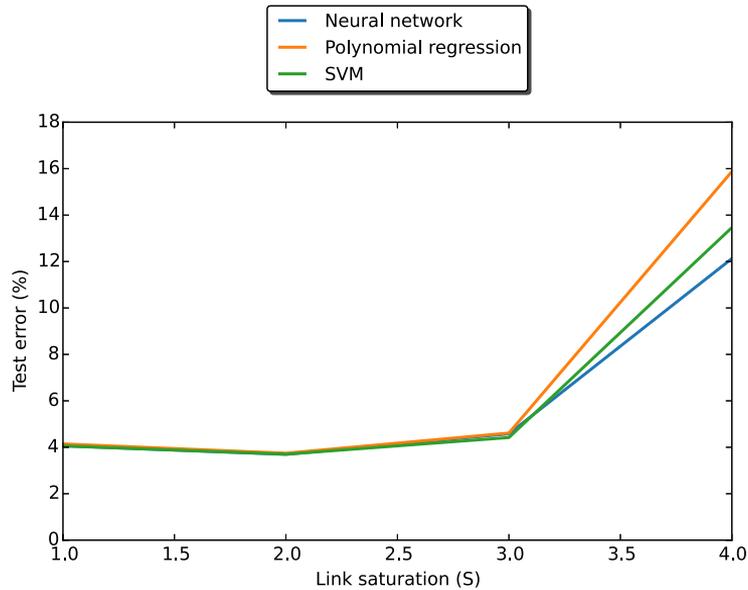
MSE vs. Routing policy

(scale-free, poisson traffic, 9 active stations)



MSE vs. Load

(scale-free, poisson traffic, 9 active stations)



Conclusions & Future Work



Conclusions & Future Work

- Results suggest that learning how to route is feasible
 - Low error for all three estimators
 - All three estimators converge to the (almost) same error
 - Polynomial regressor (order 2) is way faster to train.
- Increased load in the network leads to larger estimator error
 - This may be due to the higher randomness in the delays
- **This represents a new breed of network modeling algorithms**
- Future work
 - Test with larger networks
 - How can we represent the topology?

Thanks!!!

- More information about KDN:

- Albert Mestres, Alberto Rodriguez-Natal, Josep Carner, Pere Barlet-Ros, Eduard Alarcón, Marc Solé, Victor Muntés-Mulero, David Meyer, Sharon Barkai, Mike J Hibbett, Giovanni Estrada, Florin Coras, Vina Ermagan, Hugo Latapie, Chris Cassar, John Evans, Fabio Maino, Jean Walrand and Albert Cabellos “**Knowledge-Defined Networking**” in Arxiv.org (<http://arxiv.org/pdf/1606.06222.pdf>)

- Contribute to the NML WG at IRTF

- <https://datatracker.ietf.org/rg/nmlrg/charter/>



- Have a dataset? Want to start training your neural-network?

- Public data-sets available at: <http://knowledgedefinednetworking.org>