

Internet Research Task Force
Internet-Draft
Intended status: Informational
Expires: 19 September 2025

J. François
University of Luxembourg and Inria
A. Clemm
Independent
D. Papadimitriou
3NLab Belgium Research Center
S. Fernandes
Central Bank of Canada
S. Schneider
Digital Railway (DSD) at Deutsche Bahn
18 March 2025

Research Challenges in Coupling Artificial Intelligence and Network
Management
draft-irtf-nmrg-ai-challenges-05

Abstract

This document is intended to introduce the challenges to overcome when Network Management (NM) problems may require coupling with Artificial Intelligence (AI) solutions. On the one hand, there are many difficult problems in NM that to this date have no good solutions, or where any solutions come with significant limitations and constraints. Artificial Intelligence may help produce novel solutions to those problems. On the other hand, for several reasons (computational costs of AI solutions, privacy of data), distribution of AI tasks became primordial. It is thus also expected that networks are operated efficiently to support those tasks.

To identify the right set of challenges, the document defines a method based on the evolution and nature of NM problems. This will be done in parallel with advances and the nature of existing solutions in AI in order to highlight where AI and NM have been already coupled together or could benefit from a higher integration. So, the method aims at evaluating the gap between NM problems and AI solutions. Challenges are derived accordingly, assuming solving these challenges will help to reduce the gap between NM and AI.

This document is a product of the Network Management Research Group (NMRG) of the Internet Research Task Force (IRTF). This document reflects the consensus of the research group. It is not a candidate for any level of Internet Standard and is published for informational purposes.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 19 September 2025.

Copyright Notice

Copyright (c) 2025 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Table of Contents

1. Introduction	3
2. Acronyms	5
3. Difficult problems in network management	6
4. High-level challenges in adopting AI in NM	9
5. AI techniques and network management	11
5.1. Problem type and mapping	11
5.1.1. Sub-challenge: Suitable Approach for Given Input . .	12
5.1.2. Sub-challenge: Suitable Approach for Desired Output	13
5.1.3. Sub-challenge: Tailoring the AI Approach to the Given Problem	14
5.2. Performance of produced models	15
5.3. Lightweight AI	17
5.4. Distributed AI	19
5.4.1. Network management for efficient distributed AI . . .	19
5.4.2. Distributed AI for network management	20
5.5. AI for planning of actions	21
6. Network data as input for ML algorithms	23

6.1.	Data for AI-based NM solutions	23
6.2.	Data collection	25
6.3.	Usable data	26
7.	Acceptability of AI	28
7.1.	Explainability of Network-AI products	28
7.2.	AI-based products and algorithms in production systems	29
7.3.	AI with humans in the loop	31
8.	Security considerations	32
8.1.	AI-based security solutions	32
8.2.	Security of AI	33
8.3.	relevance of AI-based outputs	34
9.	IANA Considerations	34
10.	References	34
10.1.	Normative References	34
10.2.	Informative References	35
	Acknowledgments	47
	Authors' Addresses	47

1. Introduction

The functional scope of Network Management (NM) is very large, ranging from monitoring to accounting, from network provisioning to service diagnostics, from usage accounting to security. The taxonomy defined in [Hoo18] extends the traditional Fault, Configuration, Accounting, Performance, Security (FCAPS) domains by considering additional functional areas but above all by promoting additional views. For instance, network management approaches can be classified according to the technologies, methods or paradigms they will rely on. Methods include common approaches as for example mathematical optimization or queuing theory but also techniques which have been widely applied in last decades like game theory, data analysis, data mining and machine learning. In management paradigms, autonomic and cognitive management are listed. As highlighted by this taxonomy, the definition of automated and more intelligent techniques has been promoted to support efficient network management operations. Research in NM and more generally in networking has been very active in the area of applied ML [Bou18] and allows progresses in different domains such as network security [Ke23], Software-Defined Networks (SDN) [Ham21], vehicular network [Tan21]. This large adoption has led to a strong interlinkage between networks and AI, laying the foundation of future networks such as 6G or B5G networks expected to be AI-native [Sha25].

However, for maintaining network operational in pre-defined safety bounds, NM still heavily relies on established procedures. Even after several cycles of adding automation, these procedures are still mostly fixed and set offline in the sense that the exact control loops and all possible scenarios are defined in advance. They are so

mostly deterministic by nature or at least with sufficient safety margin. Obviously, there have been a lot of propositions to make network smarter or intelligent with the use of Machine Learning (ML) but without large adoption for running real networks because it changes the paradigms towards stochastic methods.

ML includes regression analysis, statistical learning (SVM and variants), deep learning (ANN and variants), Reinforcement Learning (RL), Large Language Models (LLMs), etc. It is a sub-area of Artificial Intelligence (AI) that concentrates the focus nowadays but AI encompasses other areas including knowledge representation, inductive logic programming, inference rule engine or by extension the techniques that allow to observe and perform actions on a system.

It is thus legitimate to question if ML or AI in general could be helpful for NM in regard to practical deployment. This question is actually tight with the problems the NM aims to address. Independently of NM, ML-based solutions were introduced to solve one type of problems in an approximate way which are very complex in nature, i.e. when finding an optimal solution is not possible (in polynomial time). This is the case for NP-hard problems. In those cases, solutions typically rely on heuristics that may not yield optimal results, or algorithms that run into issues with scalability and the ability to produce timely results due to the exponential search space. In NM, those problems exist: allocation of resources in case of Service Function Chaining (SFC) or network slicing among others are recent examples which have gained interest in our community with SDN. Many propositions consist of modelling the optimization problem as a MILP (Mixed-Integer Linear Programming) and solve it by means of heuristics to reach a satisfactory trade-off between solution quality (gap to the optimal solution) - computation time and model size/dimensionality. Hence, ML is recognized to be well adapted to progress on this type of problem [Kaf19].

However, all computational problems of NM are not NP-hard. Due to real-time constraints, some involve very short control loops that require both rapid decisions and the ability to rapidly adapt to new situations and different contexts. So, even in that case, time is critical, and approximate solutions are usually more acceptable. Again, it is where AI can be beneficial. Actually, expert systems are AI systems [Ste92] but this kind of rule-based systems are not designed to scale with the volume and heterogeneity of data we can collect in a network today. In contrast, ML is more efficient to automatically learn abstract representations of the rules, which can be eventually updated.

On one hand, another type of common problem in NM is classification. For instance, classifying network flows is helpful for security purposes to detect attack flows, to differentiate QoS among the different flows (e.g. real-time streams which need to be prioritized), etc. On the other hand, ML-based classification algorithms have been widely used with satisfactory results when properly applied leading to their application in commercial products. There are many algorithms including decision trees, Support Vector Machine (SVM) or (deep) Neural Networks (NNs) which have been to be proven efficient in many areas and notably for image and natural language processing.

Finally, many problems also still rely on humans in the loop: from support issues such as dealing with trouble tickets to planning activities for the roll-out of new services. This creates operational bottlenecks and is often expensive and error prone. This kind of tasks could be either automated or guided by an AI system to avoid individual human bias. It is worth noting that ML relying on training data generated by human can also indirectly suffer from a collective human bias. Indeed, the balance between human resources and the complexity of problems to deal with is very imbalanced and this will continue to increase due to the evolving size of networks, heterogeneity of devices, services, etc. Hence, human-based procedures tend to be simple in comparison to the problems to solve or time-consuming. Notable examples are in security where the network operator should defend against potential unknown threat.

Actually, all aforementioned problems are exacerbated by the situation of more complex networks to operate on many dimensions (users, devices, services, connections, etc.). Therefore, AI is expected to enable or simplify the solving of those problems in real networks in the near future [czb20] [Yan20] because those would require reaching unprecedented levels of performance in terms of throughput, latency, mobility, security, etc.

2. Acronyms

- * AI: Artificial Intelligence
- * CNN: Convolutional Neural Network
- * FL: Federated Learning
- * GAN: Generative Adversarial Network
- * GNN: Graph Neural Network
- * IBN: Intent-Based Networking

- * LLM: Large Language Model
- * LSTM: Long Short-Term Memory
- * MAE: Mean Absolute Error
- * ML: Machine Learning
- * MILP: Mixed-Integer Linear Programming
- * MLP: Multilayer Perceptron
- * MSE: Mean Squared Error
- * NM: Network Management
- * RL: Reinforcement Learning
- * SDN: Software-Defined Networks
- * SFC: Service Function Chaining
- * SVM: Support-Vector Machine
- * VNF: Virtual Network Function

3. Difficult problems in network management

As mentioned in introduction, problems to be tackled in NM tend to be complex and exhibit characteristics that make them good candidates for AI-based solutions:

- * C1: A very large solution space, combinatorically exploding with the size of the problem domain. This makes it impractical to explore and test every solution (again NP-hard problems).
- * C2: Uncertainty and unpredictability along multiple dimensions, including the context in which the solution is applied, behaviour of users and traffic, lack of visibility into network state, etc. In addition, many networks do not exist in isolation but are subjected to myriads of interdependencies, some outside their control. Accordingly, there are many external parameters that affect the efficiency of the solution to a problem and that cannot be known in advance: user activity, interconnected networks, etc.

- * C3: The need to provide answers (i.e. compute solutions, deliver verdicts, make decisions) in constrained or deterministic time. In many cases, the context changes dynamically and decisions need to be made quickly to be useful.
- * C4: Data-dependent solutions. To solve a problem accurately, it can be necessary to rely on large volumes of data and to deal with issues that range from data heterogeneity to incomplete data to general challenges of dealing with high data velocity.
- * C5: Need to be integrated with existing automatic and human processes.
- * C6: Solutions must be cost-effective as resources (bandwidth, CPU, human, etc.) can be limited, notably when part of processing is distributed at the network edge or within the network.

Many decision/optimization problems are affected by multiple criteria. Below is a non-exhaustive list of complex NM problems for which AI and/or non-AI-based approaches have been proposed:

- * Computation of optimal paths: Packet forwarding is not always based on traditional routing protocols with least cost routing, but on computation of paths that are optimized for certain criteria - for example, to meet certain level objectives, to result in greater resilience to balance utilization, to optimize energy usage, etc. Many of those solutions can be found in SDN, where a controller or path computation element computes paths that are subsequently provisioned across the network. However, such solutions generally do not scale to millions of paths (C1) and cannot be recomputed in sub-second time scales (C3) to take into account dynamically changing network conditions (C2). To compute those paths, operations research techniques have been extensively used in literature along with AI methods [Lop20]. Mobility and dynamicity are two conditions that make the problem of computing optimal paths even harder. Hence, adaptive routing based on RL has been proposed as it allows an agent to promptly react to topological changes [Sin22]. As such, this problem can be considered as close to big data problems with some of the different Vs: volume, velocity, variety, value 変ヲ
- * Classification of network traffic: Without loss of generality a common objective of network monitoring for operators is to know the type of traffic going through their networks (web, streaming, gaming, VoIP). Such a task analyses data (C4) which can vary over time (C2) except in very particular scenarios like industrial isolated networks. However, the output of the classification technique is time-constrained only in specific cases where fast

decisions must be made, for example to reroute traffic. Simple identification based on IANA-assigned TCP/UDP ports numbers were sufficient in the past. However, with applications using dynamic port numbers, signature techniques can be used to match packet payload [Sen04]. To handle applications now encapsulated in encrypted web or VPN traffic, various ML-based approaches have been thus adopted [Bri19][Naj24][Wan24][Ake24] including LLMs [Gin24].

- * **Network diagnostics:** Disruptions of networking services can have many causes and thus can rely on analysing many sources of data (C4). Identifying the root causes is of high importance, so that repair actions can address them versus just working around the symptoms. Such repair actions may involve human actions (C5). Further complicating the matter are scenarios in which disruptions are not complete but involve only a degradation of service level, and where disruptions are intermittent, not reproducible, and hard to predict. Artificial intelligence techniques can offer promising solutions. Especially, anticipation of faults is of paramount of importance and will lead to the development of predictive maintenance in future networks[Mut24].
- * **Network observability:** having deeper insights of network status can rely on monitoring techniques to gather data from various sources. A major issue is to aggregate all these data in a valuable format [Zha21]. When it is not directly used to automate some actions, the aggregation of the data needs to be presented in an interpretable manner to human operators. In this area, visualisation techniques are helpful and also rely on AI techniques to provide the best outputs by reducing the number of dimensions (C4) and adapting the visualisation of data for human-handled processes (C5) [Ami24].
- * **Intent-Based Networking (IBN):** Roughly speaking, IBN refers to the ability to manage networks by articulating desired outcomes without the need to specify a course of [RFC9315]. The ability to determine such courses of actions, in particular with multiple interdependencies, conflicting goals, large scale, and highly complex and dynamic environments is a huge and largely unsolved challenge (C1, C2, C3). As an illustration, a major problem with intent is to interpret them correctly knowing that different intent formats have been proposed including natural language. Without good interpretation of the intent, i.e. the expected outcomes, the derived actions will not be adequate. In case the intent is correctly interpreted, a major problem is to find concrete solutions to realize the intents which implicitly needs to optimize the actions to be taken. Artificial Intelligence techniques can be of help here in multiple ways, from accurately

classifying dynamic context to determine matching actions to reframing the expression of intent as a game that can be played (and won) using artificially intelligent techniques. As an example, LossLeap even goes further by trying to predict the gap between a targeted objective and the predicted impact of the intent realization [Col22].

- * VNF (Virtual Network Function) placement and SFC (Service Function Chain) design: VNFs need to be placed on physical resources and SFCs designed in an optimized manner to minimize the use of networking resources and energy (C1,C6). As it is known to be a NP hard problem, many heuristic- or machine learning based approaches have been proposed. The VNF paradigm actually emerges alongside 5G networking and orchestration methods [Att23].
- * Smart admission control to avoid congestion and oversubscription of network resources: Admission control needs to be set up to ensure service levels are optimized in a manner that is fair and aligned with application needs, congestion avoided, or its effects mitigated (C6). This field of research has notably been extended to the context of network slicing [Vin21][Sul23].

4. High-level challenges in adopting AI in NM

As shown in the previous section, AI techniques are good candidates for the difficult NM problems. There have been many propositions but still most of them remain at the level of prototypes or have been only evaluated with simulation and/or emulation. It is thus questionable why our community investigates much research in this direction but has not widely adopted those solutions to operate real networks. There are different obstacles.

First, AI advances have been historically driven by the image/video, natural language and signal processing communities as well as robotics for many decades. As a result, the most impressive applications are in this area including recently the generalization of home assistants, chatbots or the large progress in autonomous vehicles. However, the network experts have been focused on building the Internet, in particular designing protocols to make the world interconnected and with always better performance and services. This trend continues today with the 5G networks in deployment and beyond 5G under definition. Hence, AI was not the primary focus even if increased network automation calls for AI and ML solutions. However, AI is now considered as a core enabler for the future 6G networks which are sometimes qualified as AI-native networks [Sha25].

While we can see major contributions in AI-based solutions for networking over more than two decades, only a fraction of the community was concerned by AI at that time. Progress as a whole, from a community perspective, was so limited and compensated by relying on the development of AI in the communities as mentioned earlier. Even if our problems share some commonalities, for example on the volume of data to analyse, there are differences: data types are completely different, networks are by nature heavily distributed, etc. If problems are different, they should require distinct solutions or at least in-depth adaptation. In a nutshell, network-tailored AI was overlooked and leads to a first set of challenges described in Section 5.

Second, many AI techniques require data representative enough. For example, (deep) learning techniques mostly rely on having vectors of (real) numbers as input which fits some metrics (packet/byte counts, latency, delays, etc) but needs some adjustment for categorical (IP addresses, port numbers, etc) or topological features. Conversions are usually applied using common techniques like one-hot encoding or by coarse-grained representations [Scoll]. However, more advanced techniques can be proposed to embed representation of network entities rather than pure encoding as illustrated in [Rin17][Evr19][Sol20].

Besides, AI techniques that involve analysis of networking data can also lead to the extraction of sensitive and personally identifiable information, raising potential privacy concerns and concerns regarding the potential for abuse. Actually, this is a common and known problem that applies to many application domains [Liu22]. For example, AI techniques used to analyse encrypted network traffic with the legitimate goal to protect the network from intrusions and illegitimate attack traffic could be used to infer information about network usage and interactions of network users [Hoa21]. Intelligent data analysis and the need to maintain privacy are in many ways contradictory in nature, resulting in an arms race. Similarly, training ML solutions on real network data is often preferable over using less-realistic synthetic data sets [Liu22b]. However, network data may contain private or sensitive data, the sharing of which may be problematic from a privacy standpoint and even result in legal exposure. The challenge concerns thus how to allow AI techniques to perform legitimate network management functions and provide network owners with operational insights into what is going on in their networks, while prohibiting their potential for abuse for other (illegitimate) purposes. Challenges related to network data as input to ML algorithms is detailed in Section 6.

Finally, networks are already operated thanks to (semi-)automated procedures involving many resources which are synchronized with management or orchestration tools. Adding AI supposes so a seamless integration within pre-existing processes. Although the goal of these procedures might be solely to provide relevant information to operators through alerts or dashboards in case of monitoring applications, this can be defined to trigger actions on the different resources, which can be local or remote. The use of AI or any other approaches to derive NM actions adds further constraint on them, especially regarding time constraints [Liu21] and synchronization to maintain a coherence over a distributed system.

A related challenge concerns the fact that to be deployed, a solution needs to provide a technical solution but also be acceptable to users - in this case, network administrators and operators. With automated solutions concerns that users want to feel 寔佛n control寔and able to understand what is going on, even more so if ultimately those users are held accountable for whether or not the network is running smoothly. To mitigate those concerns, aspects such as the ability to explain actions that are taken - or about to be taken - by AI systems become important [Sen24].

Beyond reasons of making users more comfortable, there are potentially also legal or regulatory ramifications to ensure that actions taken are properly understood. For example, agencies such as the FCC may impose fines on network operators when services such as E911 experience outages. In investigating causes for such outages, the underlying behaviour of the systems has to be properly understood, and even more so the reasons for actions that fall under the realm of network operations. All these aspects about integration and acceptability of the integration of AI in NM processes is detailed in Section 7.

5. AI techniques and network management

5.1. Problem type and mapping

An increasing number of different AI techniques have been proposed and applied successfully to a growing variety of different problems in different domains, including network management [Mus18], [Xie18]. Some of the most recently proposed AI approaches are clearly advancements of older approaches, which they supersede. Many other AI approaches are not predecessors or successors but simply complementary because they are useful for different problems or optimize different metrics. In fact, different AI approaches are useful for different kinds of problem inputs (e.g., tabular data vs. text vs. images vs. time series) and also for different kinds of desired outputs (e.g., a predicted value, a classification, or an

action). Similarly, there may be trade-offs between multiple approaches that take the same kind of inputs and desired outputs (e.g., in terms of desired objective, computation complexity, constraints).

Overall, it is a key challenge of using AI technique for network management to properly understand and map which kind of problems with which inputs, outputs, and objectives are best solved with which kind of AI (or non-AI) approaches. Given the wealth of existing and newly released AI approaches, this is far from a trivial task.

5.1.1. Sub-challenge: Suitable Approach for Given Input

Different problems in network management come with widely different problem parameters. For example, security-related problems may have large amounts of textual or encrypted data as input, whereas forecasting problems have historical time series data as input. They also vary in the amount of available data.

Both the type and amount of data influences the selection of an appropriate AI technique. On one hand, in scenarios with small dimensional data, classical machine learning techniques (e.g., SVM, tree-based approaches, etc.) are often sufficient and even superior to NNs [Grel9]. On the other hand, NNs have the advantage of learning complex models from large amounts of data without requiring feature engineering. Here, different neural network architectures are useful for different kinds of problems. The traditional and simplest architecture are (fully connected) Multi-Layer Perceptrons (MLPs), which are useful for structured, tabular data. For images, videos, or other high-dimensional data with correlation between “close” features, convolutional neural networks (CNNs) are useful. Recurrent neural networks (RNNs), especially the Long Short-Term Memory (LSTM) architecture, and attention-based neural networks (transformers) are great for sequential data like time series or text. This evolution leads to the era of LLMs also impacting research in networking [Hua25][Wu24].

It is worth noting that Graph Neural Networks (GNNs) can incorporate and consider the graph-structured input, which is very useful in network management [Jie22], e.g., to represent the network topology.

The aforementioned rough guidelines can help identify a suitable AI approach or NN architecture. Still, best results are often achieved with a sophisticated combination of different approaches. For example, multiple elements can be combined into one architecture [Ham23], e.g., with both CNNs and LSTMs, and multiple separate AI approaches can be used as an ensemble to combine their strengths [Das23]. Here, simplifying the mapping from problem type and input

to suitable AI approaches and architectures is clearly an open challenge. Future work should address this challenge by providing both clearer guidelines and striving for more general AI approaches that can easily be applied to a large variety of different problem inputs.

5.1.2. Sub-challenge: Suitable Approach for Desired Output

Similar to the challenge of identifying suitable AI approaches for a given problem input, the desired output for a given problem also affects which AI approach should be chosen. Here, the format of the desired output (single value, class, action, etc.), the frequency of these outputs and their meaning should be considered.

Again, there are rough guidelines for identifying a group of suitable AI approaches. For example, if a single numerical value is required (e.g., the amount of resources to allocate to a service instance), then a supervised regression approaches should be considered as a first candidate option as it is lightweight for this type of task. In case of classification (e.g., of malware or another security issue [Abd10]) instead of predicting a value is desired, supervised methods can be used if labelled training data is available. There are also cases where a single class of training data is available, as for example in the context of anomaly detection where the model is fitted to normal data. In that case, one-class supervised techniques can be considered as a good candidate. Alternatively, unsupervised machine learning can help to cluster given data into separate groups, which can be useful to analyse networking data, e.g., for better understanding different types of traffic or user segments. Furthermore, the quality of the data [Liu22b] directly impacts on the robustness of a ML model with the risk of biased models due to overfitting. As highlighted with these few examples, finding a suitable approach to a problem depends on many factors including the type of problem to handle but also other contextual elements such as the availability and the quality of data. To help in building AI-based solutions, pipeline generators have merged with automated capabilities, paving the way to the field of AutoML [Urb23].

In addition to these classical supervised and unsupervised methods, Reinforcement Learning (RL) approaches allows active, sequential decisions rather than simple predictions or classification. This is often useful in network management, e.g., to actively control service scaling and placement [Sah23]. RL agents autonomously select suitable actions in a given environment and are especially useful for self-learning network management. In addition to model-free RL, model-based planning approaches (e.g., Monte Carlo Tree Search) also allow choosing suitable actions but require full knowledge of the environment dynamics. In contrast, model-free RL is ideal for scenarios with unknown environment dynamics, which is often the case in network management.

Like the previous sub-challenge, these are just rough guidelines that can help to select a suitable group of AI approaches. Identifying the most suitable approach within the group, e.g., the best out of the many existing reinforcement learning approaches, is still challenging. And, as before, different approaches could be combined to enable even more effective network management (e.g., heuristics + RL, LSTMs + RL, etc.). Here, further research can simplify the mapping from desired problem output to choosing or designing a suitable AI approach.

5.1.3. Sub-challenge: Tailoring the AI Approach to the Given Problem

After addressing the two aforementioned sub-challenges, one may have selected a useful kind of AI approach for the given input and output of a network management problem. For example, one may select regression and supervised learning to forecast upcoming network traffic or select reinforcement learning to continuously control network and service coordination (scaling, placement, etc.). However, even within each of these fields (regression, reinforcement learning, etc.), there are many possible algorithms and hyper-parameters to consider. Selecting a suitable algorithm and parametrizing it with the right hyper-parameters is crucial to tailor the AI approach to the given network management problem.

For example, there are many different regression techniques (classical linear, polynomial regression, lasso/ridge regression, support vector regression, regression trees, neural networks, etc.), each with different benefits and drawbacks and each with its own set of hyper-parameters. Choosing a suitable technique depends on the amount and structure of the input data as well as on the desired output. It also depends on the available amount of compute resources and compute time until a prediction is required. If resources and time are not a limiting factor, many hyper-parameters can be tuned automatically.

This sub-challenge holds for all fields of AI: supervised learning (regression and classification), self-supervised learning, unsupervised learning, and reinforcement learning, each are broad and rapidly growing fields. Selecting suitable algorithms and hyper-parameters to tailor AI approaches to the network management problem is both an opportunity and a challenge. Here, future work should further explore these trade-offs and provide clearer guidelines on how to navigate these trade-offs for different network management tasks. As already mentioned, the AutoML field of research provides solutions to better customize ML algorithms and pipeline. However, such kind of optimization should be optimized according to domain-specific metrics rather than pure-AI metrics only. For instance, the integration of network-specific knowledge can be done through human feedback [Arz21].

5.2. Performance of produced models

From a general point of view, any AI technique will produce results with a certain level of quality. This leads to two inherent questions: (1) what is the definition of the performance in a context of a NM application? (2) How to measure it? (3) How to ensure the quality of produced results by AI is aligned with NM objectives? (4) How to maintain or improve the quality of produced results?

Many metrics have been already defined to evaluate the performance of an AI-based technique according to its NM-level objectives. For example, QoS metrics (throughput, latency) can serve to measure the performance of a routing algorithm along with the computational complexity (memory consumption, size of routing tables). The question is to model and measure these two antagonist types of metrics. Number of true/false positives/negatives are the most basic metrics for network attack detection functions. Although the first two questions are thus already answered even if improvement can be done, question (3) refers to the integration of metrics into AI algorithms. Its objective is to obtain the best results which need to be quantified with these metrics. Depending on the type of algorithm, these metrics are either evaluated in an online manner with a feedback loop (for example with reinforcement learning) or in batch to optimize a model based on a particular context (for example described by a dataset for machine learning).

The problem is two-fold. First, the performance can be measured through multiple metrics of different types (numerical or ordinal for example) and some can be constrained by fixed boundaries (like a maximum latency), making their joint use challenging when creating an AI model to resolve a NM problem. Second, the scale of the metrics differs from each other in terms of importance or impact and can eventually varies on their domains. It can be hard to precisely

assess what a good or bad value is (as it might depend on multiple other ones) and it is even more difficult to integrate this in an AI technique, especially for learning algorithms to adjust their models based on performance. Indeed, many learning algorithms run through multiple iterations and rely on internal metrics, Mean Absolute Error (MAE) or Mean Squared Error (MSE) for neural network, gini index or entropy for decision trees, distance to an hyperplane for SVMs, etc) which are not strongly correlated to the final operational metrics of the NM application. AI-internal metrics such as the loss do not match well the metrics related to the final NM objectives, thus the significance and impact of the AI errors cannot be easily translated into the NM domain.

For instance, a decision tree algorithm for classification purposes aims at being able to create branches with a maximum of data from the same classes and so avoid mixing classes. It is done thanks to a criterion like the entropy index but this kind of index does not assume any difference between mixing class A and B or A and C. Assuming now that from an operational point of view, if A and B are mixed in the predictions is not critical, the algorithm should have preferred to mix A and B rather than A and C even if in the first case it will produce more errors. Therefore, the internal functioning of the AI algorithms should be refined, here by defining a particular criterion to replace the entropy as a quality measure when separating two branches. It assumes that the final NM objectives are integrated at this stage.

Another concrete example is traffic predictors which aim at forecasting traffic demands. They only produce an output that is not necessarily simple to be interpreted and used by, e.g., capacity allocation strategies/policies. A traditional traffic prediction that tries to minimize /(perfectly symmetric) MAE/MSE treats positive and negative errors in identical ways, hence is agnostic of the diverse meaning (and costs) of under- and over-provisioning. And, such a prediction does not provide any information on, e.g., how to dimension resources/capacity to accommodate the future demand avoiding all under-provisioning (which entails service disruption) while minimizing overprovisioning (i.e., wasting resources). In other words, it forces the operator to guess the overprovisioning by taking (non-informed) safety margins. A more sensible approach here is instead forecasting directly the needed capacity, rather than the traffic [Beg19].

While the one above is just an example, the high-level challenge is devising forecasting models that minimize the correct objective/loss function for the specific NM task at hand (instead of generic MAE/MSE). In this way, the prediction phase becomes an integral part of the NM, and not just a (limited and hard-to-use) input to it. In ML terms, this maps to solving the loss-metric mismatch in the context of anticipatory NM [Hua19].

Another issue for statistical learning (from examples/observations) is mainly about extracting an estimator from a finite set of input-output samples drawn from an unknown probability distribution that should be descriptive enough for unseen/new input data. In this context online monitoring and error control of the quality/properties of these point estimators (bias, variance, mean squared error, etc.) is critical for dynamic/uncertain network environments. Similar reasoning/challenge applies for interval estimates, i.e., confidence intervals (frequentist) and credible intervals (Bayesian).

Finally, question (4) refers to the ability of an AI solution to remain efficient and to eventually improve over time. This requires dynamic methods capable to adapt to a changing environment. As already highlighted, the models can be dynamically adjusted based on the errors they produced. In the context of ML, the models can be also updated based on new data, either through a complete re-learning phase, fine-tuning or transfer-learning. This assumes to collect and ingest continuously new data. However, as highlighted in [Sha21b], this type of ML, qualified as online or incremental, raises several challenges when applied to traffic analysis. For example, there is a set of related challenges related to select or discard some data over a time horizon and to label data real-time. Other challenges are more generic to this ML research area such as class imbalance or concept drift.

5.3. Lightweight AI

Network management and operations often need to be performed under strict time constraints, i.e. at line rate, in particular in the context of autonomic or self-driven networks. Locating NM functions as close as possible where forwarding is achieved is thus an interesting option to avoid additional delays when these operations are performed remotely, for example in a centralized controller. Besides, forwarding devices may offer available resources to supplement or replace edge resources. In case of AI coupled with network management, AI tasks can be offloaded in network devices, or more generally embedded within the network. Obviously, time-critical tasks are the best candidates to be offloaded within the network. Costly learning tasks should be processed in high-end servers but created models can be deployed, configured, modified and tuned in

switches.

Recent advances in network programmability ease the programming of specific tasks at data-plane level. P4 [Bos14] is widely used today for many tasks including firewalling [Dat18] or bandwidth management [Che19]. P4 is prone to be agnostic to a specific hardware. It is based on the RMT (Reconfigurable Match Table) architectural model [Bos13] that is generally accepted to be generic enough to represent limited but essential switch architecture components and functionalities. The RMT model allows reconfiguring match-action tables where actions can be usual ones (rewrite some headers, forward, drop...). Actions are thus applied on the packets when they are forwarded. Actions can also be more complex programs with some safeguards: no loop, resistivity, etc. The impact on the program development is huge. For example, real number operations are not available by default while they are widely used in many AI algorithms.

In a nutshell, the first challenge to overcome of embedding AI in a network is the capacity of the hardware to support AI operations (architectural limitation). Considering software equipment such as a virtual switch simplifies the problem but does not totally resolve it as, even in that case, strong line-rate requirement limits the type of programs to be executed. For example, BPF (Berkeley Packet Filter) [McC93] programs provide a higher control on packet processing in OVS [Cha18] but still have some limitations, as the execution time of these programs are bounded by nature to ensure their termination, an essential requirement assuming the run-to-completion model which permits high throughput.

The second challenge (resource limitation) of network-embedded AI in the network is to allocate enough resources for AI tasks with a limited impact on other tasks of network devices such as forwarding, monitoring, filtering... Approximation and/or optimization of AI tasks are potential directions to help in this area. For instance, many network monitoring proposals rely on sketches and with a well-tuned implementations for data-plane [Liu16][Yan18]. However, no general optimized AI-programmable abstraction exists to fit all cases and proposals are mostly use-case centric. There have been many propositions to develop specific P4 programs for many NM tasks, including involving ML. For each, this requires a specific adaptation [Hau23] with a few attempts to propose generic programs to be reusable or composed as a kind of libraries such as [Zha24] leveraging quantization, [Jos21] relying on pre-computed lookup tables of real-value functions or [Swa23] proposing function templates for common operations. Besides, distributed processing is a common technique to distribute the load of a single task between multiple entities. AI task decomposition between network elements, edge servers or controllers has been also proposed [Gup18].

5.4. Distributed AI

Distributed AI assumes different related tasks and components to be distributed across computational and possibly heterogeneous resources. For example, with advances in transfer and Federated Learning (FL), models can be learned, partially shared and combined or data can be also shared to either improve a local or global model. By nature, a network and a networked system is distributed and is thus well adapted to any distributed application. This is exacerbated with the deployment of fog infrastructure mixing network and computational resources. Hence, network management can directly benefit to the distributed network structure to solve its own particular problems, but any other type of AI-based distributed applications also assumes communication technologies to enable interactions between the different entities. This leads to the two sub-challenges described hereafter.

5.4.1. Network management for efficient distributed AI

Distributed AI relies on exchanging information between different entities and comes with various requirements in terms of volume, frequency, security, etc. This can be mapped to network requirements such as latency, bandwidth or confidentiality. Therefore, the network needs to provide adequate resources to support the proper execution of the AI distributed application. While this is true for any distributed application, the nature of the problem that is intended to be solved by an AI application and how this would be solved can be considered [Lin21]. For example, in the context of

optimizing FL [Che22], local models can be shared to create a global model. In case of failure of network links or in case of too high latency, some local models might not be appropriately integrated into the global model with a possible impact on AI performance. Depending on the nature of the latter, it might be better to guarantee high performance communications with a few numbers of nodes or to ensure connectivity between all of them even with lower network performance. Coupling is thus necessary between the network management plane and the distributed AI applications which leads to a set of questions to be addressed about interfaces, data and information models or protocols. While the network can be adapted or eventually adapt itself to the AI distributed applications, AI applications could also adapt themselves to the underlying network conditions [Raj24]. It paves the way to research on methods to support AI-aware NM or network-aware AI applications or a mix of both.

5.4.2. Distributed AI for network management

For network management applications relying on distributed AI, challenges from Section 5.4.1 are still valid. Furthermore, network management problems also consider network-specific elements like traffic to be analysed or configuration to be set on distributed network equipment. Co-locating AI processing and these elements (fully or partially) may help to increase performance. For example, pre-calculation on traffic data can be offloaded on network routers before being further processed in high-end servers in a data-center. Besides, as data is forwarded through multiple routers, decomposition of AI processes along the forward path is possible [Jos22]. In general, distributed AI-based network management decisions could be made at different nodes in the network based on locally available information [Sch21]. Hence, deployment of AI-based solutions for network management can also consider various network attributes like network topology, routing policies or network device capability. In that case, management of computational and network resources is even more coupled than in Section 5.4.1 since the network is both part of the AI pipeline resources and the managed object through AI.

A primary application for distributed AI is for management problems that have a local scope. One example concerns problems that can be addressed at the edge, involving tasks and control loops that monitor and apply local optimizations to the edge in isolation from activities conducted by other instances across the network. However, distributed AI can involve techniques in which multiple entities collaborate to solve a global problem. Such solutions lend themselves to problems in which centralized solutions are faced with certain foundational challenges such as security, privacy, and trust: The need to maintain complete state in a centralized solution may not be practical in some cases due to concerns such as privacy and trust

among multiple subdomains which may not want to share all of their data even if they would be willing to collaborate on a problem). Other foundational challenges concern issue related to timeliness, in which distributed solutions may have inherent advantages over centralized solutions as they avoid issues related to delays caused by the need to communicate updates globally and across long distances.

5.5. AI for planning of actions

Many tasks in network management revolve around the planning of actions with the purpose of optimizing a network and facilitating the delivery of communication services. For example, communication paths need to be planned and set up in ways that minimize wasted network resources (to optimize cost) while facilitating high network utilization (avoiding bottlenecks and the formation of congestion hotspots) and ensuring resiliency (by making sure that backup paths are not congruent with primary paths). Other examples were mentioned in Section 3.

The promise of central control is that decisions can be optimized when made with complete knowledge of relevant context, as opposed to distributed control that needs to rely on local decisions being made with incomplete knowledge while incurring higher overhead to replicate relevant state across multiple systems. However, as the scale of networks and interconnected systems continues to grow, so does the size of the planning task. Many problems are NP-hard. As a result, solutions typically need to rely on heuristics and algorithms that often result in suboptimal outcomes and that are challenging to deploy in a scalable manner [Ahm21].

The emergence of Intent-Based Networking (IBN) emphasizes the need for automated planning even further. IBN should allow users (network operators, not end users of communication services) to articulate desired outcomes without the need to specify how to achieve those outcomes. An Intent-Based System is responsible for translating the intent into courses of action that achieve the desired outcomes and that continue to maintain the outcomes over time [RFC9315]. How the necessary courses of action are derived and what planning needs to take place is left open but where the real challenge lies. Solutions that rely on clever algorithms devised by human developers face the same challenges as any other network management tasks.

These properties (problems with a clearly defined need, whose solution is faced with exploding search spaces and that today rely on algorithms and heuristics that in many cases result only suboptimal outcomes and significant limitations in scale) make automated planning of actions an ideal candidate for the application of AI-based solutions [Abd24].

A much-publicized leap in AI has been the development of AlphaGo [Sil16]. Instead of using AI to merely solve classification problems, AlphaGo has been successful in automatically deriving winning strategy for board games, specifically the game of Go which features a prohibitively large search space that was long thought to put the ability to play Go at a world class level beyond the reach of problems that AI could solve. Among the remarkable aspects of Alpha Go is that it is able to identify winning strategies completely on its own, without needing those strategies to be taught or learned by observations assuming the system is aware of rules.

The challenge for AI in network management is hence, where is the equivalent of an Alpha Go that can be applied to network management (and networking) problems? Specifically, better solutions are needed for solutions that automatically derive plans and courses of actions for network optimization and similar NP-hard problems, such as provided today with only limited effectiveness by controllers and management applications.

Although AI-based solutions for the automated planning of actions, including the automated identification of courses of action, have to this point not been explored as much as classification [Sal20]. problems, the quest for autonomous networks in the last decade and the advent of 5G and B5G have led to a quick increase in proposed solutions, as for example within the context of Zero Touch Management [Cor22].

Also, the evaluation of AI algorithms to derive courses of actions is complex. Contrary to game playing, solutions need to be applied in the real world, where actions have real effects and consequences. Different orientations can be envisioned. First, incremental application of AI decisions with small steps can allow us to carefully observe and detect unexpected effects. This can be complemented with roll-back techniques. Second, verification techniques can be leveraged to verify decisions made by AI are maintained within safety bounds [Xin24]. Third, sandbox environments can be used but they should be representative enough of the real world. After progress in simulation and emulation, recent research advances lead to the definition of digital twins which implies a tight coupling between a real system and its digital twin to ensure a parallel but synchronized execution [Wu21]. Alternatively, transfer

learning techniques in another promising area to be able to capitalize on ML models applicable on a real word system, for instance to learn an intrusion detector which can be instantiated in multiple environment [Sha24][Ans22]. Generally, it is also an open problem to make the use of AI more acceptable as highlighted in the dedicated section.

6. Network data as input for ML algorithms

Many applications of AI take data as input. The quality of the outputs of ML-based techniques are highly dependent on the quality and quantity of data used for learning but also during the inference. For example, as modern network infrastructures move towards higher speed and scale, they aim to support increasingly more demanding services with strict performance guarantees. These often require resource reconfigurations at run time, in response to emerging network events, so that they can ensure reliable delivery at the expected performance level. Timely observation and detection of events is also of paramount importance for security purposes and can allow faster execution of remedy actions thus leading to reduced service downtime.

Thus, the challenge of data management is multifaceted as detailed in next subsections.

6.1. Data for AI-based NM solutions

Assuming a network management application, the first problem to address is to define the data to be collected which will be appropriate to obtain accurate results. This data selection can require defining problem-specific data or features (feature engineering). Furthermore, machine learning algorithms only work as desired when data to be analysed respects given properties. Many methods rely on vector-based distances which so supposes that the data encoded into the vector respects the underlying distance semantic. Taking the first n bytes of a packet as vectors and computing distances accordingly is possible but does not embed the semantic of the information carried out in the headers. A solution to circumvent the problem of network traffic representation is to transform traffic data into a specific format that can be easily handled by NN architectures, for instance by creating images analysed by convolutional neural networks [Sha21]. Data can also be easily represented as graphs because topologies or communication graphs would require less adaptation to be given as inputs to GNNs [Jie22]. As an example, graph-based representations are considered as practically efficient due to their intrinsic structural similarities with network data as illustrated in [Bar23] for attack detection.

Since data to handle can also be in a schema-free or eventually text-based format, one example could be the automated annotation of management intents provided in an unstructured textual format (policies descriptions, specifications,) to extract from them management entities and operations. For that purpose, suitable annotation models need to be built using existing NER (Named Entity Recognition) techniques usually applied for NLP. However, this shall be carefully crafted or specialized for network management (intent) language which indirectly bounces back to the challenges of AI techniques for NM specified earlier. Today, with the progress on LLMs, different propositions have emerged as for example in [Mek24][Fua24][Dze23].

Secondly, similar to the problem of mapping AI algorithms with NM problems in section Section 5.1, data to be collected also depends on the NM problem to be solved. The mapping between the data sources and the problem is not straightforward as all dependencies or correlations are not known in advance and some might be discovered by the AI algorithms themselves [For23]. In addition, the types of data to collect can vary over time to maintain the performance of an AI-based application or to adapt it to a new context when learned models are updated dynamically. The problems of collecting relevant data and updating models should thus be handled conjointly.

Thirdly, the behaviour of any network is not just derived from the events that can be directly observed, such as network traffic overload, but also from events occurring outside the environment of the network. The information provided by the detectors of such kinds of events, e.g. a natural incident (earthquake, storm), can be used to determine the adaptation of the network to avoid potential problems derived from such events. Those can be provided by big data sources as well as sensors of many kinds. The challenge related to this task is to process large amounts of data and associate it with the effects that those events have on the network. It is hard to determine the static and dynamic relation between the data provided by external sources and the specific implications it may have in networks. For instance, the effect of a “flash crowd” detected in an external source may require adapting the network service configuration to support is related workload. The objective is to complement a control-loop, as shown in [Mar18], by including the specific AI engines into the decision components as well as the processes that close the loop, so the AI engine can receive feedback from the network in order to improve its own behaviour.

6.2. Data collection

Once defined, the second problem to address is the collection of data. Monitoring frameworks have been developed for many years such as IPFIX [RFC7011] and more recently with SDN-based monitoring solutions [Yu14][Ngu20]. However, going towards more AI for actions in network management supposes also to retrieve more than traffic related information. Actually, configuration information such as topologies, routing tables or security policies have been proven to be relevant in specific scenarios. As a result, many different technologies can be used to retrieve meaningful data. To support improved QoE, monitoring of the application layer is helpful but far from being easy with the heterogeneity of end-user applications and the wide use of encrypted channels. Monitoring techniques need to be reinvented through the definition of new techniques to extract knowledge from raw measurement [Bri19] or by involving end-users with crowd-sourcing [Hir15] and distributed monitoring. Also, the data-mesh concept [Mac22] proposes to classify data into three categories: source-aligned, aggregate and consumer-aligned. Source-aligned data are those related to the same operational domain, and it is important to correlate or aggregate them with higher planes: management-, control- and forwarding plane. An issue is the difference, not only in the nature of data, but in their volumes and their variety. Some may change rapidly over time (for example network traffic) while other may be quite stable (device state).

The collecting process requirements depend on the kind of processing. We can distinguish two major classes: batch/offline vs real-time/online processing. In particular, real-time monitoring tools are key in enabling dynamic resource management functions to operate on short reconfiguration cycles. However, maintaining an accurate view of the network state requires a vast amount of information to be collected and processed. While efficient mechanisms that extract raw measurement data at line rate have been recently developed, the processing of collected data is still a costly operation. This involves potentially sampling, evaluating and aggregating a vast amount of state information as a response to a diverse set of monitoring queries, before generating accurate reports. One difficult problem resides also in the availability of data as real-time data from different sources to be aggregated may not arrive at the same time requiring so some buffering techniques. Machine learning methods, e.g. based on regression, can be used to intelligently filter the raw measurements and thus reduce the volume of data to process. For example, in [Tan20] the authors proposed an approach in which the classifiers derived for this purpose (according to measurements on traffic properties) can achieve a threefold improvement in the query processing capability. A residual question is the storage of raw measurements. In fact, predicting the lifetime

of data is challenging because their analysis may not be planned and triggered by a particular event (for example, an anomaly or attack). As a result, the provisioning of storage capacity can be hard.

In parallel to the continuously increasing dynamicity of networks and complexity of traffic, there is a trend towards more user traffic processing customization [RFC8986][Li19]. As a result, fine grained information about network element states is expected and new propositions have emerged to collect on-path data or in-band network telemetry information [Tan20b]. These new approaches have been designed by introducing much flexibility and customization and could be helpful to be used in conjunction with AI applications. However, the seamless coupling of telemetry processes with packet forwarding requires careful definition of solutions to limit the overhead and the impact of the throughput while providing the necessary level of details. This shares commonalities with the lightweight AI challenge.

6.3. Usable data

Although all agree on the necessity to have more shared datasets, it is quite uncommon in practice. Data contains private or sensitive information and may not be shared because of the criticality of data (which can be used by ill-intentioned adversaries) or due to laws or regulations, even within the same company. To solve this issue, anonymization techniques [Dij19] can be enhanced to optimize the trade-off between valuable data and sensitive information (potential) leakage or reconstruction. Whatever the final user of data, regulations and laws impose rules on data management with potentially costly impact if they are not respected voluntarily or not. Defining a new monitoring framework should always consider security and privacy aspects, for example to let any user/customer or access/remove its own data with General Data Protection Regulation (GDPR) in EU. The challenge resides here in the capacity of qualifying what is critical or private information and the capacity for an adversary to reconstruct it from other sources of data. Hence AI/ML based solutions will require more data but also more administrative, legal and ethical procedures. Those can last long and so slow down the deployment of a new solution. In addition, this requires interaction with experts from different domains (e.g. AI engineer and a lawyer) for example to ensure by-design privacy in traffic analysis techniques [Wan22]. The integration of these non-technical constraints should be considered when defining new data to be collected or a new technique to collect data. However, knowing the final use of data is most of the time necessary for ethical and legal assessment which assumes that those considerations should be integrated from the early design of new AI-based solutions.

For supervised or semi-supervised training, having a labelled dataset is a prerequisite. It constitutes a major challenge as well. While probes exist to collect real network data, those data are typically unlabelled. This limits application of ML to unsupervised learning tasks (learning from data). Because manual labelling is a tedious task, one option is to leverage AI to guide humans. This may also support a better generalization of a learned model. Indeed, an underlying challenge is the genericity or coverage of the datasets. Labels encode values of an objective function, the challenge posed by the design of such tools is tremendous since for involving a M:N relationship: 1 data type may be associated to M objective function values and N data types may be associated to 1 objective function. As a result, most datasets used for research encodes a single label for a particular application like attack label for datasets to be used in the context of intrusion detection or application type for network traffic used for classification where the value of a single dataset could be capitalized in several applications. More generally, in the context of intrusion detection, using raw data rather than pre-processed data as it is common in open dataset has been demonstrated to be inefficient [Dub24].

Again, researchers need empirical (or at least realistic) datasets to validate their solutions. Unfortunately, as highlighted above, having such data from real deployments for various reasons (business secrets, privacy concerns, concerns that vulnerabilities are revealed by accident, raw unlabelled data, etc.) is tough. Even if such a dataset is available it might not be enough to convincingly validate a new algorithm. Instead of falling back to artificial testbed experiments or simulation, it would be useful to have the capability to generate datasets with characteristics that are not 100% identical but like the characteristics of one or more real datasets. Such synthetic networks can be used to validate new management algorithms, intrusion detection systems, etc. The usage of AI, Generative Adversarial Networks (GANs), in this area [Hui22] is not yet widespread and there are still many concerns that deter researchers, e.g. the fear of leaking sensitive information from the original dataset into the synthetic dataset. Furthermore, a major underlying challenge is to generate data realistic enough. This requires formalizing and integrating network-specific constraints, such as protocol specification, when generating data [Jia24].

7. Acceptability of AI

Networks are critical infrastructures. On one hand, they should be operated without interruption and must be interoperable. Networks, except in a lab, are not isolated which slow down innovation in general. For example, changing Internet routing protocols need to be accepted by multiples entities such as operators of interconnected networks. The same applies for protocols. Even if there have been several versions of major protocols in use like TCP or DNS, there are still some security issues which cannot be patched with 100% guarantee. On the other hand, results provided by AI solutions are uncertain by nature. The same technique applied in different environments can produce different results. AI techniques need some effort (time and human) to be properly configured or to be stabilized. For instance, RL needs several iterations before being able to produce acceptable results. These properties of AI techniques are thus a bit antagonist with the criticality of network infrastructures. With that in mind, acceptability of AI by network operators is clearly an obstacle for its larger adoption.

7.1. Explainability of Network-AI products

A common issue across many ML applications is their lack to provide human-understandable reasoning processes. This means that, after training, the knowledge acquired by ML models is unintelligible to humans. As a result, offering hard guarantees on performance is a very challenging issue. In addition, complex ML models like neural networks -that often have more than hundreds of thousands of parameters- are very hard to debug or troubleshoot in case of failure.

While this is a common issue for all applications of AI, many areas work well with uncertainty and the black-box behaviour of AI-based solutions. For instance, users accept an inherent error in recommender systems or computer vision solutions.

The networking field has already produced a set of well-established network management algorithms and methods, with clear performance guarantees and troubleshooting mechanisms [Rex06][Kr14]. As such, improving debugging, troubleshooting and guarantees on AI-based solutions for networking is a must. AI researchers and practitioners are devoting large research efforts to improve this aspect of ML models, which is commonly known as explainability [XAI].

This set of techniques provides insights and, in some cases, guarantees on the performance and behaviour of ML-based solutions. Understanding such techniques, researching and applying them to network AI is critical for the success of the field.

There exist several ML-based methods that are human-understandable, although not widely used today. For instance, [Mar20] shows a method for building anticipation models (prediction) that provide explanations while determining some actions for tuning some parameters of the network. There are other challenges that should be addressed, such as providing explanations for other ML methods that are quite extended. For instance, xNN/SVM models can be accompanied by Digital Twins of the network that are reversely explored to explain some output from the ML model (e.g., xNN/SVM). In this context, there already exist several methods [Zil20][Puj21] that produce human-readable interpretations of trained NN models, by analysing their neural activations on different inputs. It is worth noting that Digital Twins are not considered per se an AI approach; they merely serve to provide a digital representation of a network that can serve as its proxy and offer a layer of indirection between management applications and actual network resources. However, it is conceivable that AI-based management applications can be combined and operate in conjunction with Digital Twin technology, for example to use a Digital Twins as an experimentation sandbox or staging ground for AI-driven applications.)

7.2. AI-based products and algorithms in production systems

AI-based network management and optimization algorithms are first trained, then the resulting model is used to produce relevant inferences in operation, either in management or optimization scenarios. A relevant question for the success of AI-based solutions is: where does this training occur?

Traditionally, AI-based models have been trained in the same scenario where they operate [Vall17][Xu18], this is the customer network. However, this presents critical drawbacks. First, training an AI model for management and operation typically requires generating network configurations and scenarios that can break the network. This is because training requires seeing a broad spectrum of scenarios. Thus, training in production networks is very challenging. Second, customer networks may not be equipped with the monitoring infrastructure required to collect the data used in the training process (e.g., performance metrics). Performing learning directly into a production network is possible assuming imperfect models and the need of several step of refinement before it gets stable. For non-critical management task, such assumption can hold, and additional safe-guarding mechanisms should be considered in order to keep outputs of ML algorithms (such as decisions) within acceptable boundaries.

A more sensible approach is to train the AI-based product in a lab, for instance in the vendor's premises. In the lab, AI models can be trained in a controlled testbed, with any configuration, even ones that break the network. However, the main challenge here arises from the fundamental differences between the lab's network and the customer networks. For instance, the topology of the lab's network might be smaller, etc. As a result, there is a need for the learned models to generalize. In this context, generalization means that models should be able to operate in other scenarios not seen during training, with different topologies, routing configurations, scheduling policies, etc. This well-known problem is due to the dependency between training and testing data when benchmarking models.

In order to address this generalization problem, multiple complementary approaches are possible:

One approach is training on diverse data that represents large parts of the expected problem space. For example, training with various traffic patterns will help improving the generalization of an AI-based traffic analyser. Another approach is to leverage AI designs or architectures that facilitate generalization. One example is GNNs [gnn1][gnn2]. They are a type of neural network able to operate and generalize over graphs. Indeed, networks are fundamentally represented as graphs: topology, routing, etc. With GNNs, vendors can train the AI model in a lab with a certain topology and then directly use the resulting model in different customer networks, even with different network topologies. Finally, another approach is Transfer Learning [tl1]. With this technique, the knowledge gained in the lab's training is used to operate in the customer network. Transfer Learning still requires that some data from the customer is used to re-train and fine-tune the model (e.g., accurate performance measurements). This means that, for each customer network, re-training is required. This may be problematic since it requires added cost and access to customer data.

In addition to the challenge of generalizing from training to production environment, there are also challenges in terms of interoperability between different AI approaches and different deployment environments. As mentioned above, AI approaches may be deployed in diverse environments, e.g., for training and production, but also for local development, for testing, and for validation or in different part of the production systems. These environments may differ in available compute resources, network topology, operating systems, cloud providers, etc. (single node machine, single cluster, many distributed clusters, etc.). Deploying the same AI solutions in these different environments can lead to various challenges in terms of interoperability. Common AI frameworks support scaling across networks of different size. Yet, many frameworks are often combined,

e.g., for data collection, processing, predictions, validation, etc. Again, ensuring interoperability between these frameworks can be tedious.

This shares similarities with problems described in Section 5.4 and particularly emphasizes the need for network environments to provide interfaces and descriptions suitable for AI solutions to be properly instantiated and configured.

One approach to address these interoperability challenges is through meta-frameworks that interface with most available AI frameworks. These meta-frameworks provide a higher level of abstraction and often allow seamless deployment across different environments (e.g., on-premises or at different cloud providers) [Mor18].

7.3. AI with humans in the loop

Depending on the network management task, AI can automate and replace manual human control, or it can complement human experts and keep them in the loop. Keeping humans in the loop will be an important step of building trust in AI approaches and help ensure the desired outcomes. There are various ways of keeping humans in the loop in the different fields of AI [Wu22], which could be useful for different aspects of network management.

In classification tasks (e.g., detecting security breaches, malware or detecting anomalies), trained AI models provide a confidence score in addition to the predicted class. If the confidence is high, the prediction is used directly. If the confidence is too low, a human expert may jump in and make the decision - thereby also providing valuable training data to improve the AI model. Such approaches are already being used in industry, e.g., to automatically label datasets (AWS SageMaker). Similar approaches could also be used for other supervised learning tasks, e.g., regression. Still, it is an open challenge to keep humans in the loop in all phases of the learning process.

When using RL, e.g. to control service scaling and placement or route traffic flows, the agents typically interact with the environment (i.e., the simulated or real network) completely autonomously without human feedback. However, there is a growing number of approaches to put human experts back into the loop. One approach is offline reinforcement learning, where the training data does not come from the reinforcement learning agent's own exploration but from pre-recorded traces of human experts (e.g., placement decisions that were made by humans before). Another approach is to reward the RL agent based on human feedback rather than a pre-defined reward function [Lee21]. Again, while there are first promising approaches, more

work is required in this area. Overall, it is an open challenge to both leverage the benefits of AI but keep human experts in the loop where it is useful.

8. Security considerations

This document introduces the challenges of coupling AI and NM. Since the aim of this document is not to address a particular NM problem by defining a solution and because many possible ones can be developed further, it is not possible at this stage to define security concerns specific to a solution. However, examples of applications mentioned and cited in the different sections may face their own security concerns. In this section, our objective is to highlight high-level security considerations to be considered when coupling AI and NM. Those concerns serve as the common basis to be refined according when a particular NM application is developed.

8.1. AI-based security solutions

The first security consideration refers to the use of AI for NM problem related to security of the managed networked systems. There are multiple scenarios where AI can be leveraged: to perform traffic filtering, to detect anomalies or to decide on target moving defence strategies. In these cases, the performance of the AI algorithms impacts on the security performance (e.g. detection or mitigation effectiveness) like any other non-AI system. However, AI methods generally tend to obfuscate how predictions are made and decisions taken. Explainability of AI is thus highly important and is addressed globally in section Section 7.1.

Assuming a ML trained model, there is always an uncertainty regarding reachable performance on the wild once the solution is deployed as it can suffer from a poor generalization due to different reasons. There are two major problems which are well known in the ML field: overfitting or under-fitting. In the first case, the learned model is too specific to the training data while in the second case the model does not infer any valuable knowledge from data. To avoid these issues, hyper-parameter fine-tuning is necessary. For example, the number of iterations is an essential hyper-parameter to be adjusted to learn a neural network model. If it is too low, the learning does not converge to a representative model of the training data (underfitting). With a high value, there is a risk that the model is too close to the training data (overfitting). In general, finding the right hyper-parameters is helpful to find a good ML algorithm configuration. There are different techniques ranging from grid-search to Bayesian optimization falling into the AI area of Hyper-Parameter Optimization (HPO) [Bis23]. This consideration goes beyond the sole problem of hyper-parameters settings but as a full

analysis pipeline also assumes the ML algorithm to be selected or the data pre-processing to be configured. This reflects challenges from the AI research area covered by AutoML technique. In this document, this is also referred in section Section 5.1.3 when considered in the context of NM. As highlighted in the aforementioned section, some expertise or area-specific knowledge can help guiding automated configuration processes.

Besides, machine learning assumes to have representative training data. The quality of dataset for learning is a vast problem. Additionally, the representation of data needs to be addressed carefully to be properly analysed by AI models, for example with pre-processing techniques to normalize data, balance classes or encode categorical features depending on the type of algorithms. Actually, section Section 6 of this document fully addresses the concerns related to data in regard to NM problems.

8.2. Security of AI

Although ensuring a good performance of AI algorithms is already challenging, assuming an attacker aiming at compromising it emphasizes the problem. Adversarial AI and notably adversarial ML have attracted a lot of attention over the last year. Adversarial AI and ML relates to both attack and defences. While this is out of scope of the document, evaluating threats against an ML system before deploying it is an important aspect. This supposes to assess what types of information the attacker can access (training data, trained model, algorithm configuration...) and the performable malicious actions (inject false training data, test the system, poison a model, etc.) to evaluate the magnitude of the impact of possible attacks.

For illustration purposes, we refer hereafter to some examples. In the case of an intrusion detector, an attacker may try to poison training data by providing adversarial samples to ensure that the detector will miss-classify the future attacks [Jmi22]. In a white-box approach where the model is known from the attacker, the attack can be carefully crafted to avoid being properly labelled. For instance, packet sizes and timings can be easily modified to bypass ML-based traffic classification system [Nas21]. In a black-box model, the attacker ignores the functioning or training data of the ML systems but can try to infer some information. For example, the attacker can try to reconstruct sensitive information which have been used for training. This type of attack qualified as model inversion [Fre15] raises concern regarding privacy.

All these threats are exacerbated in the context of solutions relying on distributed AI involving multiple entities that are not necessarily controlled by the same authority. Once the threats are

assessed, solutions need to be developed and deployed which can be either proactive by providing some guarantees regarding the involved entities using authentication or trust mechanisms but also reactive by validating data processing through voting mechanisms or knowledge proofs. Other solutions include defensive techniques to rate limit or filter queries to a particular deployed model. All these examples are for illustration purposes and are not exhaustive.

8.3. Relevance of AI-based outputs

Security breaches can be created by an AI-driven application. Generally, any system that will be used to guide or advice on actions to be perform on network raise the same issue. For example, if an AI algorithm decides to change the filtering tables in a network it may compromise access control policies. Irrelevant results could be also produced. In the area of QoS, an AI system could allocate a bandwidth to a flow higher to the real link capacity. As shown from these two examples, an AI can produce decisions or values which are out of bounds of normal operations. To avoid such issues, safeguards can be added to discard or correct irrelevant outputs. Detecting such type of outputs can be also challenging in complex and distributed systems such as a network. Formal verification methods or testing techniques are helpful in that context.

9. IANA Considerations

This document has no IANA actions.

10. References

10.1. Normative References

- [RFC7011] Claise, B., Ed., Trammell, B., Ed., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, DOI 10.17487/RFC7011, September 2013, <<https://www.rfc-editor.org/info/rfc7011>>.
- [RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.
- [RFC9315] Clemm, A., Ciavaglia, L., Granville, L. Z., and J. Tantsura, "Intent-Based Networking - Concepts and Definitions", RFC 9315, DOI 10.17487/RFC9315, October 2022, <<https://www.rfc-editor.org/info/rfc9315>>.

10.2. Informative References

- [Abd10] Jalil, K. A., Kamarudin, M. H., and M. N. Masrek, "A Diagnosis Expert System for Network Traffic Management", 2010. IEEE international conference on networking and information technology
- [Abd24] Bouroudi, A., Outtagarts, A., and Y. Hadjadj-Aoul, "A dynamic AI-based algorithm selection for Virtual Network Embedding", 2024. Springer Annals of Telecommunication
- [Ahm21] Ahmad, S. and A. H. Mir, "Scalability, Consistency, Reliability and Security in SDN Controllers: A Survey of Diverse SDN Controllers", 2021. Springer Journal of Network and Systems Management, 29(2)
- [Ake24] Akem, A. T. J., Fraysse, G., and M. Fiore, "Encrypted Traffic Classification at Line Rate in Programmable Switches with Machine Learning", 2024. IEEE/IFIP Network Operations and Management Symposium (NOMS)
- [Ami24] Amin, R. M., Hammer, P., and A. Butz, "Using Machine Learning to Improve Interactive Visualizations for Large Collected Traffic Detector Data", 2024. Proceedings of the 29th International Conference on Intelligent User Interfaces (ITUI'24), ACM
- [Ans22] Anser, O., Franois, J., and I. Chrisment, "Auto-tuning of Hyper-parameters for Detecting Network Intrusions via Meta-learning", 2022. IEEE/IFIP Network Operations and Management Symposium (NOMS) - AnNet workshop
- [Arz21] Arzani, B., Hsieh, K., and H. Chen, "Interpretable Feedback for AutoML and a Proposal for Domain-customized AutoML for Networking", 2021. ACM Workshop on Hot Topics in Networks (HotNets)
- [Att23] Attaoui, W., Sabir, E., Elbiaze, H., and M. Guizani, "VNF and CNF Placement in 5G: Recent Advances and Future Trends", 2023. IEEE Transactions on Network and Service Management, 20(4), pp. 4698-4733
- [Bar23] Barsellotti, L., Marinis, L. D., Cugini, F., and F. Paolucci, "FTG-Net: Hierarchical Flow-to-Traffic Graph Neural Network for DDoS Attack Detection", 2023. IEEE International Conference on High Performance Switching and Routing (HPSR)

- [Beg19] Bega, D., Gramaglia, M., Fiore, M., Banchs, A., and X. Costa-Perez, "DeepCog: Cognitive Network Management in Sliced 5G Networks with Deep Learning", 2019. IEEE INFOCOM
- [Bis23] Bischl, B., Binder, M., Lang, M., Pielok, T., Richter, J., Coors, S., Thomas, J., Ullmann, T., Becker, M., Boulesteix, A.-L., Deng, D., and M. Lindauer, "Hyperparameter optimization: Foundations, algorithms, best practices, and open challenges", 2023. Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery, 13(2)
- [Bos13] Bosshart, P., Gibb, G., Kim, H.-S., Varghese, G., McKeown, N., Izzard, M., Mujica, F., and M. Horowitz, "Forwarding metamorphosis: Fast programmable match-action processing in hardware for SDN", 2013. ACM SIGCOMM
- [Bos14] Bosshart, P., Daly, D., Gibb, G., Izzard, M., McKeown, N., Rexford, J., Schlesinger, C., Talayco, D., Vahdat, A., Varghese, G., and D. Walker, "P4: programming protocol-independent packet processors", 2014. SIGCOMM Comput. Commun. Rev. 44
- [Bou18] Boutaba, R., Salahuddin, M. A., Limam, N., Ayoubi, S., Shahriar, N., Estrada-Solano, F., and O. M. Caicedo, "A comprehensive survey on machine learning for networking: evolution, applications and research opportunities", 2018. Journal of Internet Services and Applications 9, 16
- [Bri19] Brissaud, P.-O., Franois, J., Chrisment, I., Cholez, T., and O. Bettan, "Transparent and Service-Agnostic Monitoring of Encrypted Web Traffic", 2019. IEEE Transactions on Network and Service Management, 16 (3)
- [Chal8] Chaignon, P., Lazri, K., Franois, J., Delmas, T., and O. Festor, "Oko: Extending Open vSwitch with Stateful Filters", 2018. ACM Symposium on SDN Research (SOSR)
- [Che19] Chen, Y., Yen, L., Wang, W., Chuang, C., Liu, Y., and C. Tseng, "P4-Enabled Bandwidth Management", 2019. Asia-Pacific Network Operations and Management Symposium (APNOMS)
- [Che22] Chen, H., Huang, S., Zhang, D., Xiao, M., Skoglund, M., and H. V. Poor, "Federated Learning Over Wireless IoT Networks With Optimized Communication and Resources", 2022. IEEE Internet of Things Journal, 9(17)

- [Col22] Collet, A., Banchs, A., and M. Fiore, "Using Machine Learning to Improve Interactive Visualizations for Large Collected Traffic Detector Data", 2022. IEEE Conference on Computer Communications (Infocom)
- [Cor22] Coronado, E., Behravesh, R., Subramanya, T., Fernndez-Fernndez, A., Siddiqui, M. S., and X. Costa-Prez, "Zero Touch Management: A Survey of Network Automation Solutions for 5G and 6G Networks", 2022. IEEE Communications Surveys & Tutorials, 24(4)
- [czb20] Clemm, A., Zhani, M. F., and R. Boutaba, "Network Management 2030: Operations and Control of Network 2030 Services", 2020. Springer Journal of Network and Systems Management (JNSM)
- [Das23] Dasari, A.K., Biswas, S. K., Thounaojam, D. M., Devi, D., and B. Purkayastha, "Ensemble Learning Techniques and Their Applications: An Overview", 2023. Advances in Cognitive Science and Communications
- [Dat18] Datta, R., Choi, S., Chowdhary, A., and Y. Park,, "P4Guard: Designing P4 Based Firewall", 2018. IEEE Military Communications Conference (MILCOM)
- [Dij19] Dijkhuizen, N. V., Ham, J. V. D., and X. Li, "A Survey of Network Traffic Anonymisation Techniques and Implementations", 2014. ACM Comput. Surv. 51, 3, Article 52
- [Dub24] Dube, R., "Faulty use of the CIC-IDS 2017 dataset in information security research", 2024. Journal of Computer Virology and Hacking Techniques
- [Dze23] Dzevaroska, K., Lin, J., Tizghadam, A., and A. Leon-Garcia,, "LLM-Based Policy Generation for Intent-Based Management of Applications", 2023. International Conference on Network and Service Management (CNSM)
- [Evr19] Evrard, L., Franois, J., Colin, J.-N., and F. Beck, "port2dist: Semantic Port Distances for Network Analytics", 2019. IFIP/IEEE Symposium on Integrated Network and Service Management (IM)
- [For23] Foroughi, P., Brockners, F., and J.-L. Rougier, "ADT: AI-Driven network Telemetry processing on routers", 2023. Elsevier Computer Networks (220)

- [Fre15] Fredrikson, M., Jha, M. S., and T. Ristenpart, "Model Inversion Attacks that Exploit Confidence Information and Basic Countermeasures", 2015. ACM SIGSAC Conference on Computer and Communications Security
- [Fua24] Fuad, A., Ahmed, A. H., Riegler, M. A., and T. ii, "An Intent-based Networks Framework based on Large Language Models", 2024. IEEE International Conference on Network Softwarization (NetSoft)
- [Gin24] Ginige, Y., Dahanayaka, T., and S. Seneviratne, "TrafficGPT: An LLM Approach for Open-Set Encrypted Traffic Classification", 2024. Asian Internet Engineering Conference
- [gnn1] Battaglia, P. W. and E. al, "Relational inductive biases, deep learning, and graph networks", 2018. arXiv preprint arXiv:1806.01261
- [gnn2] Rusek, K., Surez-Varela, J., Mestres, A., Barlet-Ros, P., and A. Cabellos-Aparicio, "Unveiling the potential of Graph Neural Networks for network modeling and optimization in SDN", 2019. ACM Symposium on SDN Research
- [Gre19] Greeshma, K. V. and K. Sreekumar, "Fashion-MNIST classification based on HOG feature descriptor using SVM", 2019. Int. J. Innov. Technol. Explor. Eng. 8, 5 (2019)
- [Gup18] Gupta, A., Harrison, R., Canini, M., Feamster, N., Rexford, J., and W. Willinger, "Sonata: query-driven streaming network telemetry", 2018. ACM SIGCOMM Conference
- [Ham21] Hamdan, M., Hassan, E., Abdelaziz, A., Mohammed, B., Khan, S., Vasilakos, A. V., and M.N. Marsono, "Adversarial Machine Learning for Network Intrusion Detection Systems: A Comprehensive Survey", 2021. Journal of Network and Computer Applications, Elsevier
- [Ham23] Hammi, S., Hammami, S. M., and L.H. Belguith, "Advancing aspect-based sentiment analysis with a novel architecture combining deep learning models CNN and bi-RNN with the machine learning model SVM", 2023. Soc. Netw. Anal. Min., 13(117)

- [Hau23] Hauser, F., Hberle, M., Merling, D., Lindner, S., Gurevich, V., Zeiger, F., Frank, R., and M. Menth, "A survey on data plane programming with P4: Fundamentals, advances, and applied research", 2023. Elsevier Journal of Network and Computer Applications (212)
- [Hir15] Hirth, M., Hossfeld, T., Mellia, M., Schwartz, C., and F. Lehrieder, "Crowdsourced network measurements: Benefits and best practices", 2015. Computer Networks. 90
- [Hoa21] Hoang, N. P., Niaki, A. A., Gill, P., and M. Polychronakis, "Domain name encryption is not enough: Privacy leakage via IP-based website fingerprinting", 2021. Proceedings on Privacy Enhancing Technologies
- [Hoo18] Hoof, J. V. D., Claeys, M., Bouten, N., Wauters, T., Schnwlder, J., Stiller, A. P. B., Charalambides, M., Badonnel, R., Serrat, J., Santos, C. R. P. D., and F. D. Turck, "Updated Taxonomy for the Network and Service Management Research Field", 2018. Journal of Network System Management (JNSM) 26, 790808
- [Hua19] Huang, C., Zhai, S., Talbott, W., Bautista, M. A., Sun, S.-Y., Guestrin, C., and J. Susskind, "Addressing the Loss-Metric Mismatch with Adaptive Loss Alignment", 2020. ICRL
- [Hua25] Huang, Y., Du, H., Zhang, X., Niyato, D., Kang, J., and Z. Xiong, "Large Language Models for Networking: Applications, Enabling Techniques, and Challenges", 2025. IEEE Network, 39(1)
- [Hui22] Hui, S., Wang, H., Wang, Z., Yang, X., Liu, Z., Jin, D., and Y. Li, "Knowledge Enhanced GAN for IoT Traffic Generation", 2022. ACM Web Conference 2022 (WWW)
- [Jia24] Jiang, X., Liu, S., Gember-Jacobson, A., Bhagoji, A. N., Schmitt, P., Bronzino, F., and N. Feamster., "NetDiffusion: Network Data Augmentation Through Protocol-Constrained Traffic Generation", 2024. ACM on Measurement and Analysis of Computing Systems
- [Jie22] Jieng, W., "Graph-based deep learning for communication networks: A survey", 2022. Computer Communications, Elsevier, vol. 185

- [Jmi22] Jmila, H. and M. I. Khedher, "Adversarial machine learning for network intrusion detection: A comparative study", 2022. Computer Networks
- [Jos21] Jose, M., Lazri, K., Franois, J., and O. Festor, "InREC: In-network REal Number Computation", 2021. IFIP/IEEE International Symposium on Integrated Network Management (IM)
- [Jos22] Jose, M., Lazri, K., Franois, J., and O. Festor, "NetREC Network-wide in-network REal-value Computation.", 2022. IEEE International Conference on Network Softwarization (NetSoft)
- [Kaf19] Kafle, V. P., Martinez-Julia, P., and T. Miyazawa, "Automation of 5G Network Slice Control Functions with Machine Learning", 2019. IEEE Communications Standards Magazine, vol. 3, no. 3, pp. 54-62
- [Ke23] Ke, H., Dongseong, K. D., and M. R. Asghar, "Adversarial Machine Learning for Network Intrusion Detection Systems: A Comprehensive Survey", 2023. IEEE Communications Surveys & Tutorials, 25(1), 538-566.
- [Kr14] Kreutz, D., Ramos, F. M., Verissimo, P. E., Rothenberg, C. E., Azodolmolky, S., and S. Uhlig, "Software-defined networking: A comprehensive survey", 2015. Proceedings of the IEEE, vol. 103, no. 1, pp. 14-76
- [Lee21] Lee, K., Smith, L., and P. Abbeel, "Feedback-efficient interactive reinforcement learning via relabeling experience and unsupervised pre-training", 2021. arXiv preprint arXiv:2106.05091
- [Li19] Li, R., Makhijani, K., Yousefi, H., Westphal, C., Dong, L., Wauters, T., and F. D. Turck., "A Framework for Qualitative Communications Using Big Packet Protocol", 2019. ACM SIGCOMM Workshop on Networking for Emerging Applications and Technologies (NEAT)
- [Lin21] Lin, Z., Bi, S., and Y. -J. A. Zhang, "Optimizing AI Service Placement and Resource Allocation in Mobile Edge Intelligence Systems", 2021. IEEE Transactions on Wireless Communications, 20(11)

- [Liu16] Liu, Z., Manousis, A., Vorsanger, G., Sekar, V., and V. Braverman, "One Sketch to Rule Them All: Rethinking Network Flow Monitoring with UnivMon", 2016. ACM SIGCOMM Conference
- [Liu21] Liu, Q., Choi, N., and T. Han, "Constraint-Aware Deep Reinforcement Learning for End-to-End Resource Orchestration in Mobile Networks", 2021. IEEE International Conference on Network Protocols (ICNP)
- [Liu22] Liu, B., Ding, M., Shaham, S., Rahayu, W., Farokhi, F., and Z. Lin, "When Machine Learning Meets Privacy: A Survey and Outlook", 2022. ACM Comput. Surv., 54(2)
- [Liu22b] Liu, L., Engelen, G., Lynar, T., Essam, D., and W. Joosen, "Error Prevalence in NIDS datasets: A Case Study on CIC-IDS-2017 and CSE-CIC-IDS-2018", 2022. IEEE Conference on Communications and Network Security (CNS)
- [Lop20] Lopez, J., Labonne, M., Poletti, C., and D. Belabed, "Priority Flow Admission and Routing in SDN: Exact and Heuristic Approaches", 2020. IEEE International Symposium on Network Computing and Applications (NCA)
- [Mac22] Machado, I. A., Costa, C., and M. Y. Santos, "Data Mesh: Concepts and Principles of a Paradigm Shift in Data Architectures", 2022. Procedia Computer Science, Volume 196
- [Mar18] Martinez-Julia, P., Kafle, V. P., and H. Harai, "Exploiting External Events for Resource Adaptation in Virtual Computer and Network Systems", 2018. IEEE Transactions on Network and Service Management, Vol. 15, N. 2,
- [Mar20] Martinez-Julia, P., Kafle, V. P., and H. Asaeda, "Explained Intelligent Management Decisions in Virtual Networks and Network Slices", 2020. Conference on Innovation in Clouds, Internet and Networks and Workshops (ICIN)
- [Mcc93] McCanne, S. and V. Jacobson, "The BSD packet filter: a new architecture for user-level packet capture", 1993. USENIX Winter Conference
- [Mek24] Mekrache, A., Ksentini, A., and C. Verikoukis, "Intent-Based Management of Next-Generation Networks: an LLM-Centric Approach", 2024. IEEE Network, 38(5)

- [Mor18] Moritz, P., Nishihara, R., Wang, S., Tumanov, A., Liaw, R., Liang, E., Elibol, M., Yang, Z., Paul, W., Jordan, M., and I. Stoica, "Ray: A Distributed Framework for Emerging AI Applications", 2018. USENIX Symposium on Operating Systems Design and Implementation (OSDI)
- [Mus18] Musumeci, F., Rottondi, C., Nag, A., Macaluso, I., Zibar, D., Ruffini, M., and M. Tornatore, "An overview on application of machine learning techniques in optical networks", 2018. IEEE Communications Surveys & Tutorials, 21(2), 1383-1408.
- [Mut24] Muteba, F., Dahanayaka, T., and S. Seneviratne, "Digital twin (DT)-based predictive maintenance of a 6G communication network", 2024. Procedia Computer Science, Elsevier, vol. 238
- [Naj24] Najm, I. A., Saeed, A. H., Ahmad, B.A., Ahmed, S. R., Sekhar, R., Shah, P., and B.S. Veena, "Enhanced Network Traffic Classification with Machine Learning Algorithms", 2024. Proceedings of the Cognitive Models and Artificial Intelligence Conference
- [Nas21] Nasr, M., Bahramali, A., and A. Houmansadr, "Defeating DNN-Based Traffic Analysis Systems in Real-Time With Blind Adversarial Perturbations", 2021. USENIX Security Symposium
- [Ngu20] Nguyen, T. G., Phan, T. V., Hoang, D. T., Nguyen, T. N., and C. So-In, "Efficient SDN-based traffic monitoring in IoT networks with double deep Q-network", 2020. International conference on computational data and social networks, Springer
- [Puj21] Pujol-Perich, D., Surez-Varela, J., Xiao, S., Wu, B., Cabello, A., and P. Barlet-Ros, "NetXplain: Real-time explainability of Graph Neural Networks applied to Computer Networks", 2021. MLSys workshop on Graph Neural Networks and Systems (GNNSys)
- [Raj24] Rajasekaran, S., Ghobadi, M., and A. Akella, "CASSINI: Network-Aware Job Scheduling in Machine Learning Clusters", 2024. USENIX Symposium on Networked Systems Design and Implementation (NSDI)
- [Rex06] Rexford, J., "Route optimization in IP networks", 2006. Handbook of Optimization in Telecommunications (pp. 679-700), Springer

- [Rin17] Ring, M., Dallmann, A., Landes, D., and A. Hotho, "IP2Vec: Learning Similarities Between IP Addresses", 2017. IEEE International Conference on Data Mining Workshops (ICDMW)
- [Sah23] Saha, N., Zangoeei, M., Golkarifard, M., and R. Boutaba, "Deep Reinforcement Learning Approaches to Network Slice Scaling and Placement: A Survey", 2023. IEEE Communications Magazine, 61(2)
- [Sal20] Salman, O., Elhajj, I. H., Kayssi, A., and A. Chehab, "A review on machine learningbased approaches for Internet traffic classification", 2020. Springer Annals of Telecommunication
- [Sch21] Schneider, S., Qarawlus, H., and H. Karl, "Distributed Online Service Coordination Using Deep Reinforcement Learning", 2021. IEEE International Conference on Distributed Computing Systems (ICDCS)
- [Scol1] Coull, S. E., Monrose, F., and M. Bailey, "On Measuring the Similarity of Network Hosts: Pitfalls, New Metrics, and Empirical Analyses", 2011. NDSS
- [Sen04] Sen, S., Spatscheck, O., and D. Wang, "Accurate, scalable in-network identification of p2p traffic using application signatures", 2004. ACM International conference on World Wide Web (WWW)
- [Sen24] Senevirathna, T., La, V. H., Marchal, S., Siniarski, B., Liyanage, M., and S. Wang, "A Survey on XAI for 5G and Beyond Security: Technical Aspects, Challenges and Research Directions", 2024. IEEE Communications Surveys & Tutorials
- [Sha21] Shapira, T. and Y. Shavitt, "FlowPic: A Generic Representation for Encrypted Traffic Classification and Applications Identification", 2023. IEEE Transactions on Network and Service Management, 18(2), pp. 1218-1232
- [Sha21b] Shahraki, A., Abbasi, M., Taherkordi, A., and A. D. Jurcut,, "A comparative study on online machine learning techniques for network traffic streams analysis", 2022. Elsevier Computer Networks (207)
- [Sha24] Latif, S., Boulila, W., Koubaa, A., Zou, Z., and J. Ahmad, "DTL-IDS: An optimized Intrusion Detection Framework using Deep Transfer Learning and Genetic Algorithm", 2024. Elsevier Journal of Network and Computer Applications

- [Sha25] Sharma, G. and B. Priya, "Chapter 26 - Future now: unleashing the potentials of machine learning algorithms in 6G", 2025. *Advances in Computational Methods and Modeling for Science and Engineering*, Morgan Kaufmann, 331-339
- [Sil16] Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Driessche, G. V. D., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Silver, D., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., and D. Hassabis, "Mastering the game of Go with deep neural networks and tree search", 2016. *Nature*, vol. 529 (2016), pp. 484-503
- [Sin22] Singh, A., Prakash, S., and S. Singh, "Optimization of reinforcement routing for wireless mesh network using machine learning and high-performance computing", 2022. *Concurrency Computat Pract Exper.*, 35(15)
- [Sol20] Soliman, H. M., Salmon, G., Sovilij, D., and M. Rao, "A Graph Neural Network Approach for Scalable and Dynamic IP Similarity in Enterprise Networks", 2020. *IEEE International Conference on Cloud Networking (CloudNet)*
- [Ste92] Stern, D. and P. Chemouil, "A Diagnosis Expert System for Network Traffic Management", 1992. *Networks*, Kobe, Japan
- [Sul23] Sulaiman, M., Moayyedi, A., Ahmadi, M., Salahuddin, M. A., Boutaba, R., and A. Saleh, "Coordinated Slicing and Admission Control Using Multi-Agent Deep Reinforcement Learning", 2023. *IEEE Transactions on Network and Service Management*, 20(2), pp. 1110-1124
- [Swa23] Swamy, T., Zulfiqar, A., Nardi, L., Shahbaz, M., and K. Olukotun, "Homunculus: Auto-Generating Efficient Data-Plane ML Pipelines for Datacenter Networks", 2023. *ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*
- [Tan20] Tangari, G., Charalambides, M., Pavlou, G., Grazian, C., and D. Tuncer, "Classification-assisted Query Processing for Network Telemetry", 2020. *Network Traffic Measurement and Analysis Conference (TMA)*
- [Tan20b] Lizhuang, T., Wei, S., Zhenyi, Z., Jingying, M., Xiaoxi, L., and L. Na, "In-band Network Telemetry: A Survey", 2020. *Computer Networks*. 186. 10.1016

- [Tan21] Tang, F., Mao, B., Kato, N., and G. Gu, "Comprehensive Survey on Machine Learning in Vehicular Network: Technology, Applications and Challenges", 2023. IEEE Communications Surveys & Tutorials, 23(3), 2027-2057.
- [tl11] Torrey, L. and J. Shavlik, "Transfer learning", 2010. Handbook of research on machine learning applications and trends: algorithms, methods, and techniques
- [Urb23] Urbanowicz, R., Zhang, R., Cui, Y., and P. Suri, "STREAMLINE: A Simple, Transparent, End-To-End Automated Machine Learning Pipeline Facilitating Data Analysis and Algorithm Comparison", 2023. Genetic Programming Theory and Practice XIX. Genetic and Evolutionary Computation. Springer
- [Val17] A., V., M., S., D., S., and T. A., "Learning to route", 2017. ACM HotNets
- [Vin21] Vincenzi, M., Lopez-Aguilera, E., and E. Garcia-Villegas, "Timely admission control for network slicing in 5G with machine learning", 2021. IEEE Access
- [Wan22] Wang, J., Han, H., Li, H., He, S., Sharma, P. K., and L. Chen, "Multiple Strategies Differential Privacy on Sparse Tensor Factorization for Network Traffic Analysis in 5G", 2022. IEEE Transactions on Industrial Informatics, 18(3)
- [Wan24] Wang, X., Yuan, Q., Wang, Y., Gou, G., Gu, C., Yu, G., and G. Xiong, "Combine intra- and inter-flow: A multimodal encrypted traffic classification model driven by diverse features", 2024. Computer Networks, Elsevier, vol. 245
- [Wu21] Wu, Y., Zhang, K., and Y. Zhang, "Digital Twin Networks: A Survey", 2021. IEEE Internet of Things Journal, 8(18)
- [Wu22] Wu, X., Xiao, L., Sun, Y., Zhang, J., Ma, T., and L. He, "A survey of human-in-the-loop for machine learning", 2022. Elsevier Future Generation Computer Systems
- [Wu24] Wu, D., Wang, X., Qiao, Y., Wang, Z., Jiang, J., Cui, S., and F. Wang, "NetLLM: Adapting Large Language Models for Networking", 2024. ACM SIGCOMM 2024 Conference
- [XAI] Samek, W., Wiegand, T., and K.-R. Mller, "Explainable artificial intelligence: Understanding, visualizing and interpreting deep learning models", 2017. arXiv preprint arXiv:1708.08296

- [Xie18] Xie, J., Yu, F. R., Huang, T., Xie, R., Liu, J., Wang, C., and Y. Liu, "A survey of machine learning techniques applied to software defined networking (SDN): Research issues and challenges", 2018. IEEE Communications Surveys & Tutorials
- [Xin24] Xing, J., Hsu, K.-F., Xia, Y., Cai, Y., Li, Y., Zhang, Y., and A. Chen, "Occam: A Programming System for Reliable Network Management", 2024. ACM European Conference on Computer Systems (Eurosys)
- [Xu18] Z., X., J., T., J., M., W., Z., Y., W., H., L. C., and Y. D., "Experience-driven networking: A deep reinforcement learning based approach", 2018. IEEE INFOCOM
- [Yan18] Yang, T., Jiang, J., Liu, P., Huang, Q., Gong, J., Zhou, Y., Miao, R., Li, X., and S. Uhlig, "Elastic sketch: adaptive and fast network-wide measurements", 2018. ACM SIGCOMM Conference
- [Yan20] Yang, H., Alphones, A., Xiong, Z., Niyato, D., Zhao, J., and K. Wu,, "Artificial-Intelligence-Enabled Intelligent 6G Networks", 2020. IEEE Network, vol. 34, no. 6, pp. 272-280
- [Yu14] Yu, Y., Qian, C., and X. Li, "Distributed and collaborative traffic monitoring in software defined networks", 2014. ACM Hot topics in software defined networking
- [Zha21] Zhao, J., Jing, X., Yan, Z., and W. Pedrycz, "Network traffic classification for data fusion: A survey", 2021. Information Fusion, Elsevier, vol. 72
- [Zha24] Zhang, K., Samaan, N., and A. Karmouch, "A Machine Learning-Based Toolbox for P4 Programmable Data-Planes", 2024. IEEE Transactions on Network and Service Management, 21(4), pp. 4450-4465
- [Zil20] Meng, Z., Wang, M., Bai, J., Xu, M., Mao, H., and H. Hu, "Interpreting Deep Learning-Based Networking Systems", 2020. ACM SIGCOMM

Acknowledgments

This document is the result of a collective work. Authors of this document are the main contributors and the editors, but contributions have been also received from the following people we acknowledge Laurent Ciavaglia, Felipe Alencar Lopes, Abdelkader Lahamdi, Albert Cabellos, Jos Surez-Varela, Marinos Charalambides, Ramin Sadre, Pedro Martinez-Julia and Flavio Esposito

This document was also partially supported by project AI@EDGE, funded from the European Union's Horizon 2020 H2020-ICT-52 call for projects, under grant agreement no. 101015922. This research is funded in part, by the Luxembourg National Research Fund (FNR), grant reference C23/IS/18088425/COCTEL. The views expressed in this document do not necessarily reflect those of the Bank of Canada's Governing Council.

Authors' Addresses

Jrme Franois
University of Luxembourg and Inria
6 Rue Richard Coudenhove-Kalergi
L- Luxembourg
Luxembourg
Email: jerome.francois@uni.lu

Alexander Clemm
Independent
United States of America
Email: ludwig@clemm.org

Dimitri Papadimitriou
3NLab Belgium Reseach Center
Leuven
Belgium
Email: papadimitriou.dimitri.be@gmail.com

Stenio Fernandes
Central Bank of Canada
Canada
Email: stenio.fernandes@ieee.org

Stefan Schneider
Digital Railway (DSD) at Deutsche Bahn
Germany
Email: stefanbschneider@outlook.com