White Paper

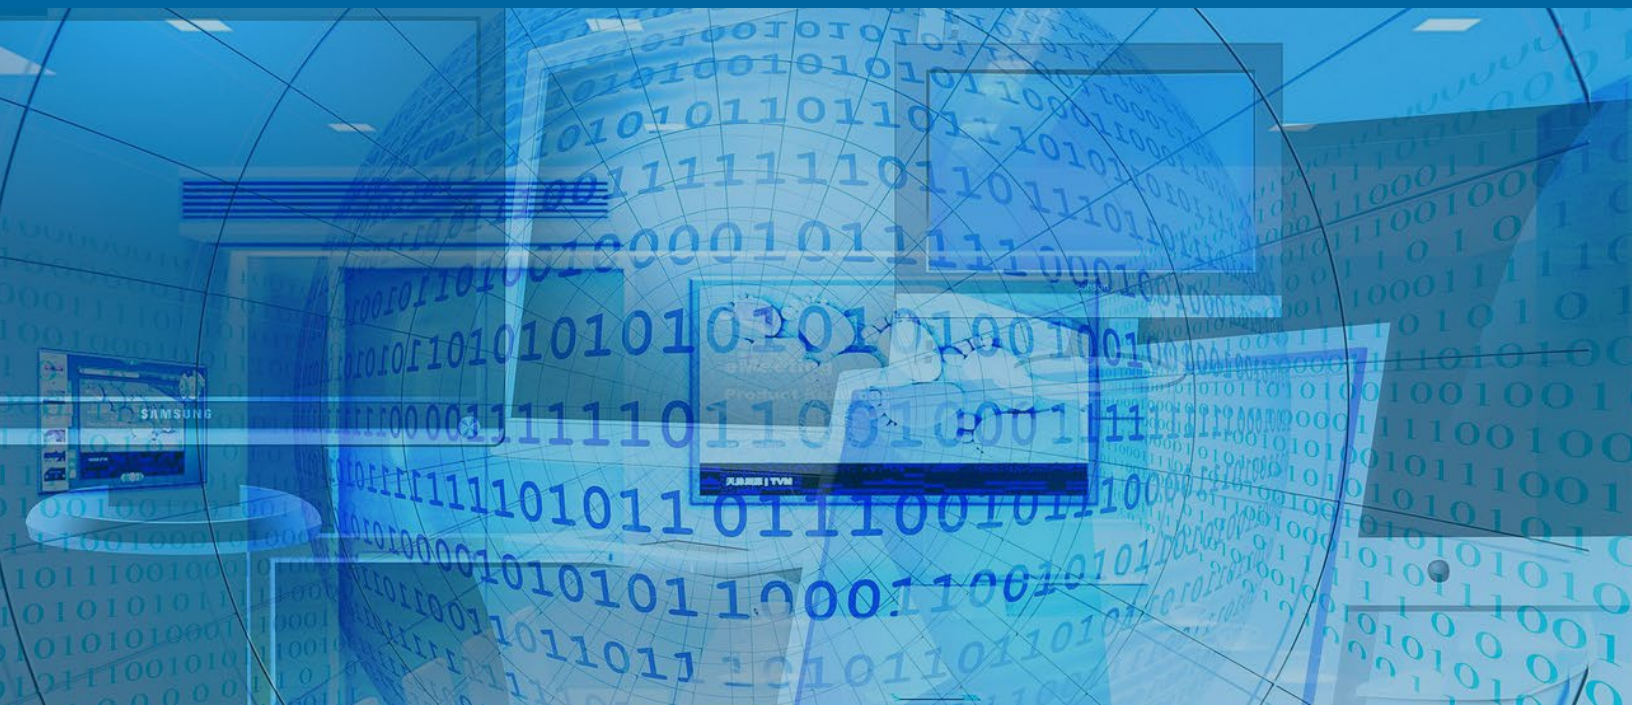

INTERNET POLICY IN THE AGE OF MACHINE LEARNING

November 2018

Nishanth Sastry and Mischa Dohler, King's College London, UK

Disclaimer

Table of Contents

ABSTRACT

Communication networks are traditionally underpinned by standards and policies which establish operating rules and ensure interoperability. Such governance allows for a rich stakeholder system of Content Providers, Internet Service Providers, Mobile Operators, and consumers to grow and drive much of the world's economy today. However, with the increasing complexity of the fixed and wireless internet, machine learning and artificial intelligence (AI) have become an integral part of the day-to-day operation of these communication networks. What was a network run by a few, largely myopic, algorithms and protocols is slowly morphing into an infrastructure governed by a rapidly increasing set of highly intelligent and often unaccountable algorithms.

In this white paper we examine the growing use of machine learning and AI in the internet. We begin by examining how machine learning is being incorporated into the working principles of the internet, from a researcher's perspective. We then look at the possible future role for AI and machine learning in ensuring the continued functioning of the internet, and elaborate the impact of this development on current and future internet policies. We conclude by discussing a few case studies.

CONSENSUS AND THE INTERNET

The internet[1], which started as an academic research project, has now grown to underpin a large portion of today's economy, connecting society as never before, with applications like social networks and instant messaging services that link friends across continents, and e-governance portals that provide basic needs of the citizens. As the internet has become a basic necessity of modern life, and a basic human right according to the U.N. Human Rights Council[2], it has become imperative for regulatory and policy bodies to consider whether and how to shape this technology that has woven itself into the fabric need of everyday life. To function smoothly, the internet needs agreement amongst different stakeholders, and a policy roadmap on two levels:

The first level of consensus is required on the governance and husbandry of shared or common resources. This involves questions such as: Should the internet remain neutral to all the applications using it, or should prioritized delivery services be allowed, or a hybrid thereof? How should competition among different operators (mobile as well as fixed-line broadband) be regulated? Or, how should scarce spectrum be shared or dispensed to mobile network operators? These topics are typically under the purview of *national* regulatory bodies such as the Federal Communications Commission (FCC) in the U.S. or the Office of Communications (Ofcom) in the U.K. Being national bodies, these regulators often derive their power from their legislatures, and their agenda relates to national infrastructure and practices. This means that different regulators can sometimes be at odds with each other. For instance, Brazil, Chile and The Netherlands have strong network neutrality laws, whereas the U.S. is currently changing her stance on network neutrality. Since internet services can span international borders (e.g., a website hosted in the U.S. can be accessed by a user in Brazil), such national differences in policy and regulation can lead to the conflicts or sub-optimal performance of the overall system.

A second level of agreement is needed on a technical ground – standards are essential to ensure that all entities can "speak" to each other and maintain the global nature of the internet. For instance, all applications today operate on top of the Internet Protocol (IP). Agreement over this is what enables the global traffic flows across the internet. Achieving a new consensus, moving from IP version 4 to version

---

[1] In this white paper, we use the word internet in its colloquial sense, to encompass the hardware infrastructure such as routers and other specialist devices, and the suite of protocols such as TCP and IP that run on top of it, which together work to provide access to a vast amount of information, including, but not limited to the World Wide Web (WWW), as well as various applications and so-called 'Apps' in the mobile world.

[2] https://www.article19.org/data/files/Internet_Statement_Adopted.pdf

6 has taken over 20 years, and is not fully complete yet. This fine level of technical synchronization is achieved through the operation of various standards bodies, such as the IEEE, Internet Engineering Task Force (IETF), International Telecommunications Union (ITU), the 3G Partnership Programme (3GPP), and ETSI (European Telecommunications Standards Institute). These bodies typically do not have a regional or national agenda (ETSI is an exception, but despite the European name, associate members from several non-European countries actively participate in the discussions). However, in many cases, there is a well-understood separation of duties: the IETF standardizes various protocols related to the internet, especially those relating to fixed-line internet; the IEEE focuses on broadband wireless access and the 3GPP focuses on cellular networks; and ETSI has taken a lead on technologies such as Network Functions Virtualisation (NFV) and Multi-access Edge Computing (MEC).

Consensus at these two levels is required for all aspects of the internet. In this white paper, we examine one specific and new aspect – the growing use of machine learning and AI in the internet. We start off by examining how machine learning is being incorporated into the working principles of the internet, from a researcher's perspective. We then look at the possible future role for AI and machine learning in ensuring the continued functioning of the internet. We then elaborate the impact of this development on the current internet policies and further developments in the future, and finally conclude this white paper by discussing a few case studies.

EVOLVING COMPLEXITY IN TELECOMMUNICATIONS AND INTERNET PROTOCOLS

The internet was built on a philosophy of simplicity (Clark, 1988). Indeed, the design of the internet was a reaction to the highly complex protocols used in telecommunications networks to connect phones (land lines) together. The telecommunications networks operated on a paradigm of setting up an end-to-end connection between the caller and the called, guaranteeing resources on each intermediate entity in order to ensure that the calls went through without a hitch. Instead, the internet adopted a "best effort" model, where each router or intermediate node in the internet would do its best to forward data packets along, with no guarantees whatsoever on whether delivery would happen. The onus is on the end hosts which are communicating with each other to add reliability through mechanisms such as retransmitting after waiting for a reasonable period of time.

The simplicity and standardized interface enabled by this so-called end-to-end principle (Saltzer, 1984) set the base for the growth of the *internet*, as a *network of networks* operated by mutually independent domains with potentially different

underlying technologies, that enter into a business relationship with each other to provide mutual interconnection to each other's networks, and beyond. Because only the end hosts had to agree about complex issues, it enabled an immense amount of innovation and heterogeneity.

However, in the recent years, there has been a tremendous increase in complexity. Let us illustrate this with video streaming as an example. With the growth of applications such as YouTube and Netflix, video has come to dominate the internet. Some estimates suggest that 80-90% of all traffic in the near future will be video[3]. Video requires to be played out at a minimum frame rate (e.g., 30 frames per second) with a resolution suitable for the display device. At a network level, this translates to support for good amount of bandwidth and a low variation or "jitter" in the bandwidth over time. This is difficult to guarantee over the Best Effort Internet. Any number of issues can arise, ranging from increased queueing delay due to a sudden burst of packets from other flows which share the same bottleneck link, to dropped packets due to router faults or overfull queues, to packets which are lost due to mistakes in the routing tables that direct packets from source to destination.

The device that is playing the video can insure against some amount of bandwidth variation by maintaining a small buffer of packets before playing them out, and using this buffer when there are temporary fluctuations in bandwidth. However, this is by no means sufficient. For example, a short run of missed packets may empty the entire buffer, causing stalls. Scaling video streaming globally has been achieved by abandoning the classical "end-to-end" principle, and introducing content delivery servers in the middle, which are distributed closer to the users. A large amount of network telemetry is used to identify the "best" servers for each user, and the user is directed to a server that is closer to them, and provides good performance. The bandwidth at which the videos are streamed from these servers is also adapted based on current network conditions. Content Delivery Networks have been extremely successful, but the problem is by no means solved. Exciting new research is emerging that uses deep neural networks to help decide the best rate (Mao, 2017) and best servers for each client.

This rise in complexity is reflected in other areas as well. For instance, middle boxes – specialized nodes such as Network Address Translators, Load Balancers, Firewalls and Intrusion Detection Systems – have seen a sharp rise in growth, to the point that they are on par with routers and switches in terms of numbers[4]. Middleboxes such as load balancers change the last hops near servers; network address translators can make multiple clients appear as the same host; firewall rules can

---

[3] https://www.cisco.com/c/en/us/solutions/service-provider/visual-networking-index-vni/vni-infographic.html

[4] https://www.cs.cornell.edu/courses/cs5413/2014fa/lectures/26-datacenter-middleboxes.pdf.

grant access on some ports for some protocols while at the same time preventing reachability over another port for another protocol. Thus, the increased use of middleboxes has made it difficult to reason about how paths are chosen through the internet. If a certain destination is not reachable, it becomes very difficult to debug and pinpoint the cause of the issue.

Similarly, older and more established protocols, which started off as simple, expedient stop-gap fixes, have outgrown their initial designs and are causing increasing difficulties for operators. The Border Gateway Protocol (BGP), the inter-domain routing protocol that is used to discover and advertise about reachability beyond immediate neighbors, is famously known as the "two napkin" protocol because its inventors came up with the initial design on two paper napkins during a lunch meeting. BGP allows autonomous systems (independent networks) to advertise routes to neighboring autonomous systems without giving away the internal network topology or the business deals and arrangements. Trusting the neighbors' advertisements, which is the basis of BGP operation, can cause errors in routing tables that can propagate far. In 2008, many networks around the world started believing that YouTube had moved to Pakistan, because of a BGP advertisement[5]. Nevertheless, this mistake was corrected quickly by YouTube by advertising a more authentic shorter path to itself, this exposes the vulnerability of the current internet to misconfigurations and attacks. Similarly, the Domain Name System which translates from human-friendly names such as www.ieee.org to an equivalent IP address relies on a hierarchical name resolution infrastructure which, if hacked, can lead to attacks such as redirecting www.bank.com to the attacker's server, which can be made to look like the real bank website. Designed merely as an aid for human memory, DNS servers now represent a highly critical link in the internet infrastructure that must be secured with great diligence.

Another sharp rise in complexity is due to an exponentially increasing number of mobile end-points in the internet. Ever since the IEEE and 3GPP standards have enabled natively embedding IP into the mobile end points, such as laptops, mobile phones and Internet of Things (IoT) devices, the internet has not only grown in the number of end-points but also in terms of characteristics including mobility, unreliability due to the wireless links (and not due to the congested routers), heterogeneity of devices, among others. While both fixed and mobile networks were largely running in parallel over past decades, we will shortly witness strong convergence in the 5G era which in turn constitutes a significant increase in network management complexity. The methodologies to manage such networks are being standardized, such as ETSI's Management and Orchestration (MANO) standards; however, the algorithmic frameworks to actually do the orchestration are scarce and largely based on machine learning.

[5] https://www.wired.com/2008/02/pakistans-accid/

THE ROLE OF MACHINE LEARNING & AI IN THE FUTURE INTERNET

Examples such as the above have given rise to the idea of using predictive and intelligent oversight approaches to deal with the rise in complexity. This trend has been fostered by the growth of available data, as well as increased availability of telemetry and network analytics tools to process this data. However, beyond basic analytics such as performance dashboards, predictive approaches typically involve machine learning in one form or another.

According to Tom Mitchell, a pioneer in the field of statistical machine learning, "A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P if its performance at tasks in T, as measured by P, improves with experience E." (Mitchell, 1997). This definition applies equally to data-driven methods used to improve performance on the internet. Note that the term *Artificial Intelligence (AI)*, which is sometimes used synonymously with machine learning in literature, actually encompasses the idea of machines learning and extends beyond it to several other approaches that involve some form of "intelligence" other than learning (e.g., the use of logic and reasoning). For the purposes of this white paper, it is not essential to define or enumerate all forms of AI, but it is important to be aware that there are approaches different from machine learning which have proven to be useful in various application contexts, including in networking and telecommunications.

While there are any number of machine learning algorithms which can be applied to a given problem domain such as networking and communications, a traditional taxonomy is to consider whether the items or instances being learned are labelled or not. In a case where labelled instances are available, the learning mechanisms or algorithms are termed as supervised learning algorithms. A learning agent can take advantage of a "teacher" or "oracle" which has labelled instances in a training set (e.g., as positive or negative for a particular condition), and learn to give these labels on a test set. In some cases, labels are available only for a small sub-sample, and learning algorithms that function in this condition are called semi-supervised learning algorithms. In contrast, when there is no label available, the agent is forced to use unsupervised learning algorithms that learns distinct classes or patterns in high dimensional data. More often seen in networking related applications is the use of reinforcement learning where an agent is given a reward for successful performance, and a punishment (or regret) for unsuccessful performance, and optimizes its operation to maximize the reward or minimize regret.

Examples of the success of supervised and semi-supervised learning include tools such as Spam Assassin[6], which uses a range of Bayesian filtering techniques. Spam Assassin and similar tools are trained based on emails which are labelled as spam or not spam, and learn the prior probability distribution of words which are highly frequent in spam emails but not in regular emails. Note that since the training of spam and not spam can be given by individual users, Spam Assassin can adapt itself to the vocabulary of emails received by that user.

Another nice illustration of the ability to learn from experience through reinforcement is found in the recently proposed TCP Ex Machina approach to designing protocols for congestion control (Weinstein, 2013). One consequence of the end-to-end principle previously alluded to has been a long history of research into congestion control protocols. The goal of congestion control is to understand when the end hosts are congesting the network and adapt their transmission rates accordingly. Different congestion control protocols have been designed to meet various network conditions, and different application requirements. In most cases, this translates to two kinds of rules: one for increasing the bandwidth available to the communicating hosts, and one for decreasing the bandwidth. Bandwidth used by a flow is increased (resp. decreased) when free capacity (resp. congestion) is detected, either by an explicit directive from the network or by heuristics that are employed by the communicating hosts. These rules are typically hand crafted, to meet application- or network-specific requirements. Instead, TCP Ex Machina proposes to learn from prior experience. By playing out a model of the congestion control state in heterogeneous network conditions for a "design range" of operation, it proposes to learn a congestion control algorithm for a specific state. The evaluation shows that this *Ex Machina* or "of the machine" congestion control scheme performs better than carefully tuned handcrafted rules.

It should be noted that the above examples are only chosen as illustration. The range of predictive approaches which can be used (and have been successfully applied) with the internet is broader than just the traditional approaches of machine learning. There have been proposals based on control theoretic approaches that predict variables important for the problem and then solve an optimization problem based on this prediction. This method, known as model predictive control, has been applied to solve the problem of dynamically and adaptively choosing the best bitrate for video streaming (Yin, 2015). Other approaches have involved formally specifying beliefs about the properties desired of a network and then reasoning about whether these beliefs hold, in the face of varying network configurations (Lopes, 2015), and declarative networking (Loo, 2009), which espouses writing down a laundry list of expected properties and then generating code or protocols that satisfy such properties.

---

[6] https://en.wikipedia.org/wiki/SpamAssassin

One of the broadest applications of Artificial Intelligence into networks is seen in Clark *et al.*'s bold vision of the Knowledge Plane (Clark, 2003). It is one of the first papers to recognize and articulate the need to incorporate cognitive approaches to tame the growing complexity of network protocols. The Knowledge Plane calls for a network that can reason about itself, and aims to develop a network architecture that is sufficiently self-aware to be able to identify *why* a problem occurs when one occurs. For instance, if a user is not able to connect to www.ieee.org, the network should be able to tell whether the problem is in the first hop (e.g., the user is not authenticated to a captive portal, or the broadband provider has cut off network connectivity because a bill was not paid), or whether the problem is within the network (e.g., a routing misconfiguration is causing packets to go in a loop). Going beyond this self-diagnostic property, Clark *et al.* suggest a self-managing/self-healing network, one that can *fix* the problem when it occurs. Notice that the network may not be able to solve the problem by itself *all* the time – e.g., if a user does not have connectivity because of an unpaid bill to the broadband provider, this necessarily requires the user's intervention.

Systems that are able to take care of themselves in this fashion are sometimes dubbed as autonomous or cognitive systems. A holy grail for an autonomous or cognitive system is to have various self-* properties, such as being self-configuring (choosing appropriate values for different system parameters), self-healing (fixing errors by itself), self-optimizing (adjusting system parameters automatically to achieve best performance), and self-protecting (guard against viruses and other attacks without external intervention). The Knowledge Plane and related proposals are an attempt to bring this world view into the domain of networking and communications. As with many proposals for autonomic systems, the knowledge plane has been a motivation for additional follow-on research, but has not yet been fully realized as a concrete implementation, despite a decade and a half of work towards it.

The Knowledge Plane proposal recognizes the complexity involved in the task it has set itself. Most of the functionality in networks is organized into a layered stack, with each layer hiding the complexity from layers above it; and providing a set of functions which can be relied upon by the higher layers. Thus, for example, the network layer may provide the functionality of packetizing information and routing these packets to a network layer address (IP address in the internet). The transport layer is a higher layer that relies on this packet forwarding function of the network layer and adds additional functionality such as reliable data delivery or ensuring that the packets that arrive at the destination are ordered in the same way as at the sender. In this paradigm, a key question is deciding in which layer of the network should its cognitive capabilities reside. Clark and his co-authors recognize that the cognitive functions of a network will span *all* the layers, and therefore envision a Knowledge Plane, that permeates all the layers and becomes a plane of functionality similar to the data plane and the control plane. In doing so, the knowledge plane gives a first-class status to the data about the network and the metadata about the

data being transported, and thereby enables reasoning about the network and its functioning.

It should be noted that cognitive properties have been considered and independently researched upon in the context of radios for wireless communications for slightly longer than the first articulation of the Knowledge Plane (Mitola, 1999). Indeed, an early vision was to increase the spectrum efficiency by allowing for secondary users to use the same spectrum as primary users; this, in turn, required advanced sensing and then autonomous decision-making approaches due to the distributed nature of the architecture. The thus evolving Cognitive (i.e. learning) Radios were soon extended into Docitive (i.e. teaching) Radios[7] with much better performance in terms of speed and precision of convergence.

With the research field of Cognitive and Docitive Radios gaining in momentum, the industry picked up on the potential and, for example, 3GPP introduced first self-organizing networking (SON) mechanisms into its feature portfolio. Notably, Release 8 enabled base station self-configuration (e.g. automatic neighbor relation, automatic physical cell ID (PCI) assignment, automatic inventory and automatic software download); Release 9 was about network optimization procedures (e.g. load balancing optimization and inter-cell interference coordination (ICIC)); Release 10 was about overlaid networks; and Releases 11 and 12 about SON in highly heterogeneous deployment settings. While the standards framework enabled SON and the implementation of machine learning / AI, no viable industry solution was available at that time to handle the very high number of conflicting requirements, along with the very high degree of freedom in the network. Some light SON features are active in today's networks, but it has not lived up to its expectations.

Ever since these earlier standards releases in 3GPP, the complexity in the wireless edge has further increased: a solid response to the smart phone adoption as well as a steady increase in the number of IoT devices connected means that humans are not able to optimize lest run such networks at scale. Machine learning and AI seem to be the only possible solution to ensure a reliable operation of the wireless internet.

---

[7] L. Giupponi, A. Galindo-Serrano, P. Blasco and Mischa Dohler, "Docitive Networks – An Emerging Paradigm for Dynamic Spectrum Management," IEEE Wireless Communications Magazine, Special Issue on Dynamic Spectrum Management, Vol 17, Issue 4, pp. 47-54, 2010.

IMPACT ON INTERNET POLICY

From the above it has become clear that artificial intelligence (AI), and machine learning (ML), is crucial for the well-being and proper functioning of our fixed and mobile internet. However, it is imperative that we examine the impact of using machines to run our networks, some of which have become highly mission-critical infrastructure.

Given the growing worries about AI taking control of many aspects of human life (prominent examples include Elon Musk's warnings[8], Stephen Hawking's ambivalence[9], etc.), it is timely to ask if specific policies on this topic are needed to ensure technology is put to its best use, and provide society with a means to manage the complexity of providing universal interconnection. To this end, we first inspect the impact of AI on various prominent internet policies; and then argue for the need of a new policy framework on AI.

In the context of the internet, three issues have to be settled in relation to the use of AI and machine learning: Accountability, privacy and universal access.

Policy on Accountability:

Accountability is important as it settles disputes. However, due to the convolutional and "black-box" nature of many algorithms underpinning AI, accountability is not only poorly defined but also difficult to achieve – if a failure occurs with "black box" AI algorithms, it can be difficult to pinpoint why the failure occurred, and ensure accountability. The algorithms may perform as desired in 99.99999% of all cases; who is accountable and how do we find the engineering problem in the system in the rare case of failures? Traditional AI approaches are clearly not the answer whereas the pioneering work on "Explainable AI" might yield the future of artificial intelligence in more industrial applications. Insisting on Explainable AI may mean giving up a number of potentially powerful AI tools and approaches, but this may be unavoidable to ensure accountability. We note that Explainable AI is far from a solved problem, and is the subject of a current DARPA challenge[10].

---

[8] https://www.vanityfair.com/news/2017/03/elon-musk-billion-dollar-crusade-to-stop-ai-space-x

[9] Stephen Hawking has previously said AI could destroy humanity (cf: http://www.bbc.co.uk/news/technology-30290540) but more recently expressed an ambivalent attitude that AI could either be the best or worst thing for the human race (cf: https://www.theguardian.com/science/2016/oct/19/stephen-hawking-ai-best-or-worst-thing-for-humanity-cambridge) .

[10] https://www.darpa.mil/program/explainable-artificial-intelligence

Detecting faults or failures in networks is an example of how AI and Machine Learning can be used to enable accountability. In a complex ecosystem such as the internet, there are a number of players that rely on each other. Many of these relationships are governed by strong Service Level Agreements (SLAs). However, given the scale at which the internet operates, it may sometimes not be possible to identify whether there has been a fault, or if there is has been a failure, who is to be held accountable. AI-based approaches, building on the previously mentioned Knowledge Plane concept, can potentially be a solution. Going forward, it may even be possible to build predictive models of potential failures, and then using these to fix these errors (e.g., congestion in the connection between two routers in the Mediterranean Sea may mean a loss of connectivity for a country in Sub Saharan Africa that links to Europe and the rest of the world through that router. If a pattern of instability is detected using AI approaches *before* the link fails, that information can be used to bring up a backup option. If that is not possible, pattern recognition can be used to identify the *mode* of failure when it happens, and attach a degree of belief – a probability estimate – that the fault was due to a cable in the Mediterranean. Additional forensics can then be used to determine whether this was really the cause of the fault, which SLA was violated, and what the penalty would be).

Policy on Privacy:

AI can be a strong support in monitoring and enforcing privacy. However, due to the convolutional and "black-box" nature of many algorithms underpinning AI, it is almost certain that personal information is natively embedded into a large-scale infrastructure; if the system gets compromised or goes rogue, there is no guarantee that private information will not be released or misused.

An important point to bear in mind is that most AI approaches (especially those based on machine learning) require lots of data to be collected, and this is typically mentioned in Terms of Service that end users sign up to. Thus, the growing use of AI has led to an increase in ("authorized") collection, use and distribution of many aspects of private information. Private companies have sometimes tended to collect more data than required for current services in anticipation of enabling potential analytics in the future. This may, in some cases, and some countries, be seen as "misuse" of data, and Policy bodies need to look at this more carefully to ensure and prohibit intended or inadvertent misuse of personal data.

Upcoming legislation such as the General Data Protection Regulation (GDPR) in the EU has already started to look at this, and will provide fundamental protections[11] to individuals regarding their personal data. In the context of the internet, this includes information such as the IP address of a device belonging to an end user. In the U.S.,

[11] https://ico.org.uk/media/for-organisations/data-protection-reform/overview-of-the-gdpr-1-13.pdf

IP address by itself is not considered Personally Identifiable Information (PII), but it is considered as linked PII (McCallister, 2010; Section 3.3.2), i.e., a user can be traced by correlating IP address with other records, for example a log that contains domain login information (which IP address a username logged in from).

Going beyond privacy of personal information, other sensitive information, such as the "secret sauce" of different companies, which needs to be kept private. BGP arose in order to allow different autonomous systems (ASes) to keep their internal topology information private while at the same time co-operating to advertise routes that ensure global reachability on the internet. With AI and pattern recognition comes the possibility that such internal topologies can be reverse engineered. Such practice, if allowed, may undermine the careful web of trust that keeps the internet together, and therefore may require policy intervention.

Policy on Equal/Universal Access:

AI can potentially play a positive role in network access since, if well designed, the holistic non-myopic nature of AI (as compared with the relatively myopic nature of traditional networking protocols and heuristics) ensures that universal access can be monitored and enforced at scale for every single individual concerned. However, the potential for building global AI-based monitoring mechanisms, even if it be for well-intentioned reasons such as ensuring accountability, also showcases the power imbalances inherent in this complex ecosystem of stakeholders.

In many countries, the internet currently functions on an implicit assumption of equal access, i.e., no content provider is discriminated against by individual ISPs, or treated favorably in comparison with other content providers. In an "AI-enhanced" internet, it is imperative that policy makers clarify how AI may affect such equal access policies. For instance, technically, nothing prevents ISPs from replacing the ads on the content providers pages with ads injected on the fly. There may be valid and good reasons for providing differential treatment to different players or even for injecting content such as ads into an existing stream of bytes (e.g., when the network is congested, an ISP may be able to deliver the message if a high-quality image is replaced with a lower quality one). However, whether this should be allowed or not, and if allowed, under what circumstances, is an important question that requires vigorous policy debates.

Geographical and Other Contexts:

Policies are situated within the context of different countries that are making these policies. For instance, in Europe, the U.S., and several other countries, it is considered illegal for Internet Service Providers to log and track the browsing activities of end users. However, in other countries, doing so is not explicitly forbidden. One of the major problems for internet advertisers trying to build behavioral profiles of their users is that they lack visibility into the full universe of websites that the users visit. The ISP of the user is able to see all traffic of the user, and is therefore able to build a much more accurate behavioral profile than any single advertiser or website. This would be seen in some countries as a massive

violation of privacy; in other jurisdictions, this may potentially be seen as an aid to the user in that the ads delivered by learning a more complete behavioral profile may allow for better targeting. Some countries actually require retention of browsing history for purposes of law enforcement. This has been a significant concern for civil libertarians.

Importantly, policy relies on three pillars i) *observability* (i.e. are we able to detect/monitor the policy construct); ii) *implementability* (i.e. are we able to implement mechanisms which execute a given policy); and iii) *enforceability* (i.e. are we able to enforce policies and disobedience thereof). In the context of AI, we struggle with all three but above all with "observability" which relies on a "platform of observation" by humans. In the absence of such an objective point of observation, one would not even know when/if AI isn't going beyond its designed purpose. **Therefore, large efforts ought to be directed towards policy frameworks which enable and enact common platforms of observation which allow us (humanity, or society) to gauge if AI is acting beyond a designed purpose, and whether it is serving its intended purpose, which is the provision of fit-for-purpose, and acceptable connectivity solutions to humanity.**

CONCLUDING REMARKS

Thus far, the internet has been built, managed and extended by engineers. One of the founding beliefs of the Internet Engineering Task Force has been the following dictum from the aforementioned David Clark: "We reject kings, presidents and voting. We believe in rough consensus and running code"[12].

Apart from the running code, there was a separate code – a code of ethics, which ensured that the internet was available to all as a common good. Indeed, Tim Berners-Lee, the father of the World Wide Web, refused to patent anything, instead preferring to place all Intellectual Property in the public domain, to be used by commercial as well as non-commercial entities. Much of the internet is run on open source code such as Linux, BSD and Apache. There exists a strong culture of "giving back" to the community, with virtually all major big commercial players including Google, Facebook, Twitter, Apple, Microsoft and IBM making large amounts of their internal code available as open source, for others to build on.

Although this outlook of mutual respect based on technical competence and forward thinking cooperation amongst competing entities has served us well until now, setting up an incredible pace of innovation, and creating a massive infrastructure that has completely transformed every aspect of our lives within a generation, we

[12] https://www.ietf.org/tao.html

have come to a point where we need to stop and think more carefully about the societal aspects of future improvements to the internet. Running code can be good or evil. The nature of the internet as a shared medium means that one entity's innovation may end up yielding profits to itself but inadvertently harm another entity or entities. Although this danger of a "tragedy of the commons" existed even in the early years of the internet, the code which was running the internet was code handcrafted by humans. Code, which, despite its complexity, could be understood by other humans. Admittedly these humans were a selected elite – a highly trained and highly competent band of engineers, but the code was under human control nevertheless.

As we move from this regime of mutual human and organizational-level co-operation towards a regime where AI controls part of the internet, we need to ensure that societal needs such as privacy, accountability and fairness are met. In this paper, we presented case studies and potential scenarios where such societal needs may be affected by the growing use of Artificial Intelligence. Such scenarios deserve attention from a Policy angle, to ensure the utility of the internet and to avoid a possible near-term future where privacy can be destroyed as a result of excessive and unregulated data collection. A future where governments and businesses need to be accountable not only to each other but also to the citizens and end consumers who rely on them for services through the internet.

We also painted a picture where AI can be useful, not only in enhancing the performance of the underlying internet infrastructure, but also in ensuring its safety, and in creating a new economy which is not possible without a machine-controlled infrastructure that improves upon itself through learning. Machine learning can help provide the basic technical underpinning to help navigate the complicated and complex ecosystem of today's internet but appropriate regulation is required to make sure that this happens in the way that society wants it to: The vast amounts of data available today can help address information imbalances if we build monitoring tools that use collaborative big-data methods to ensure that a) global monitoring mechanisms are not being used without explicit user consent and b) if and when the global monitoring happens, it is tailored to the wishes of the users whose data is being mined.

BIBLIOGRAPHY

Mitchell, Tom. (1997). Machine Learning. McGraw Hill. p. 2. ISBN 0-07-042807-7.

Clark, David. "The design philosophy of the DARPA Internet protocols." ACM SIGCOMM Computer Communication Review18, no. 4 (1988): 106-114.

Clark, David D., Craig Partridge, J. Christopher Ramming, and John T. Wroclawski. "A knowledge plane for the internet." In Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications, pp. 3-10. ACM, 2003

Yin, Hao, Yong Jiang, Chuang Lin, Yan Luo, and Yunjie Liu. "Big data: transforming the design philosophy of future internet." IEEE network 28, no. 4 (2014): 14-19.

Huang, Haojun, Hao Yin, Geyong Min, Hongbo Jiang, Junbao Zhang, and Yulei Wu. "Data-Driven Information Plane in Software-Defined Networking." IEEE Communications Magazine (2017).

Mao, Hongzi, Ravi Netravali, and Mohammad Alizadeh. "Neural Adaptive Video Streaming with Pensieve." In Proceedings of the Conference of the ACM Special Interest Group on Data Communication, pp. 197-210. ACM, 2017.

Loo, Boon Thau, Tyson Condie, Minos Garofalakis, David E. Gay, Joseph M. Hellerstein, Petros Maniatis, Raghu Ramakrishnan, Timothy Roscoe, and Ion Stoica. "Declarative networking." Communications of the ACM 52, no. 11 (2009): 87-95.

Lopes, Nuno P, Nikolaj Bjørner, Patrice Godefroid, Karthick Jayaraman, George Varghese. 2015. “Checking Beliefs in Dynamic Networks.” The 12th USENIX Symposium on Networked Systems Design and Implementation, 499–512.

McCallister, Erika, Timothy Grance, and Karen A. Scarfone. Guide to protecting the confidentiality of personally identifiable information (PII). No. Special Publication (NIST SP)-800-122. 2010.

Mitola, Joseph, and Gerald Q. Maguire. "Cognitive radio: making software radios more personal." IEEE personal communications 6, no. 4 (1999): 13-18.

Saltzer, Jerome H., David P. Reed, and David D. Clark. "End-to-end arguments in system design." ACM Transactions on Computer Systems (TOCS) 2, no. 4 (1984): 277-288.

Yin, Xiaoqi, Abhishek Jindal, Vyas Sekar, and Bruno Sinopoli. "A control-theoretic approach for dynamic adaptive video streaming over http." ACM SIGCOMM Computer Communication Review 45, no. 4 (2015): 325-338.

Winstein, Keith, and Hari Balakrishnan. "Tcp ex machina: Computer-generated congestion control." In ACM SIGCOMM Computer Communication Review, vol. 43, no. 4, pp. 123-134. ACM, 2013.