

AI-Assisted Frequency Resource Management in Multi-Tenant Optical Networks

Najwa Pradipta

Department of Computer Science, George Mason University, Fairfax, VA, USA.
najwa.pradipta405@gmu.edu

Abstract

Multi-tenant optical networks are becoming a foundational infrastructure for cloud services, enterprise connectivity, research collaborations, and emerging low-latency applications. In such environments, frequency resources are no longer managed solely as an engineering constraint but as a governed, contested, and economically significant asset. This paper examines artificial intelligence assisted frequency resource management in multi-tenant optical networks from a system-level and socio-technical perspective. It argues that effective management requires more than improved spectrum assignment algorithms; it demands architectural choices, operational policies, fairness safeguards, accountability mechanisms, and sustainability considerations that operate together across optical, computational, and institutional layers. The paper first reviews elastic optical networking and multi-tenancy as structural conditions that create both flexibility and coordination problems. It then analyzes frequency resource management as a control problem shaped by tenant isolation, service-level expectations, and infrastructure economics. Subsequent sections explore machine learning for traffic prediction, routing and spectrum assignment, quality-of-transmission estimation, defragmentation, and failure management. The discussion emphasizes trade-offs among efficiency, robustness, fairness, and explainability. Governance is treated as a first-order design concern, including auditability, contestability, and equitable allocation under scarcity. Deployment and sustainability are examined through energy use, lifecycle impacts, and the risk that AI optimization may shift rather than eliminate resource burdens. The paper concludes that AI-assisted frequency management should be embedded in accountable institutional frameworks rather than treated as a purely technical substitute for policy judgment.

Keywords

multi-tenant optical networks; frequency resource management; artificial intelligence; spectrum allocation; network governance; tenant fairness; sustainability; optical infrastructure.

1. Introduction

The rapid expansion of cloud computing, distributed artificial intelligence, high-definition media, industrial automation, and scientific data exchange has placed sustained pressure on optical transport infrastructure. Optical networks are no longer merely backhaul for relatively stable traffic patterns; they are dynamic platforms that must support heterogeneous services with different bandwidth, latency, reliability, and privacy requirements. Elastic optical networking emerged as a response to this heterogeneity by replacing fixed wavelength grids with flexible spectrum allocation and bandwidth-variable transmission [3]. This shift increased spectral efficiency and architectural adaptability, but it also introduced complex coordination problems because spectrum, transponder capacity, regeneration resources, and routing paths must be allocated jointly rather than independently [4]. Early access and passive

optical architectures already demonstrated that access and metro domains would need scalable, service-aware resource models [1]. As networks become multi-tenant, these problems intensify because multiple independent organizations share the same physical substrate while expecting isolation, predictable performance, and contractual guarantees.

Frequency resource management in multi-tenant optical networks is therefore not only an optimization problem. It is a governance problem, an infrastructure design problem, and a socio-technical control problem. Spectrum is finite, spatially constrained, and subject to interference-like effects when channels are placed poorly or when fragmentation accumulates. A tenant may demand dedicated capacity for safety-critical operations, while another may tolerate variable performance in exchange for lower cost. The operator must balance utilization, revenue, resilience, and fairness under uncertainty. Artificial intelligence offers powerful tools for prediction, classification, anomaly detection, and sequential decision-making in such settings [9]. Machine learning has been applied to traffic forecasting, routing and spectrum assignment, quality-of-transmission estimation, fault localization, and resource defragmentation [11]. Deep learning has also been proposed for intelligent traffic control and network management [10]. However, the presence of AI does not dissolve structural tensions. It may even amplify them if models are opaque, biased toward high-value tenants, or brittle under rare failures.

This paper positions AI-assisted frequency resource management within a broader system architecture that includes optical hardware, control planes, software-defined orchestration, service-level agreements, regulatory constraints, and sustainability commitments. It argues that the most credible deployments will be those that treat AI as one component in a layered governance regime rather than as an autonomous substitute for policy. Recent work on frequency-referenced multiplexing, for example, illustrates how physical-layer innovations can provide guaranteed bandwidth and low latency, but such guarantees still require allocation rules, monitoring, and enforcement mechanisms to become tenant-facing assurances [6]. The discussion proceeds from architectural foundations to allocation methods, governance, deployment, resilience, and policy implications. The central claim is that multi-tenant optical frequency management is a long-term institutional and infrastructural challenge in which technical efficiency, fairness, and accountability must be co-designed.

2. Architectural Foundations of Multi-Tenant Optical Networks

Multi-tenant optical networks rest on a layered architecture that separates physical transmission, spectrum routing, control, and service orchestration. In elastic optical networks, the frequency grid is divided into finer slices than traditional wavelength division multiplexing, allowing channels of variable width and adaptable modulation formats [3]. Bandwidth-variable transponders and flexible-grid reconfigurable optical add-drop multiplexers enable the optical layer to allocate spectrum more closely to demand. The architecture is often described as spectrum-sliced or sliceable, meaning that a single physical infrastructure can support multiple logical networks with different topologies, capacities, and protection levels [4]. This virtualization is attractive because it allows infrastructure providers to serve cloud operators, enterprises, mobile backhaul providers, and research networks without building separate physical systems. Yet virtualization also creates new dependencies. If logical networks share amplifiers, fibers, and switching elements, then a fault or misconfiguration in one tenant's slice can affect others unless isolation is carefully designed.

The European IDEALIST vision articulated a comprehensive view of elastic optical networks as programmable, adaptive, and scalable infrastructures capable of supporting heterogeneous

traffic [5]. That vision anticipated many of the coordination problems now seen in multi-tenant environments. Spectrum assignment cannot be separated from routing because the chosen path determines the available frequency slots and the accumulation of physical impairments. Routing and spectrum assignment algorithms therefore seek to minimize blocking, reduce fragmentation, and respect continuity and contiguity constraints [7]. More sophisticated allocation models consider modulation format, reach, regeneration, and the economic value of each request [8]. In a multi-tenant setting, these decisions also involve tenant priority, contractual penalties, security boundaries, and fairness criteria. A purely throughput-maximizing algorithm may repeatedly favor large tenants whose traffic patterns are easy to predict, while smaller or intermittent tenants suffer higher blocking. Over time, such bias can become a competitive distortion rather than a neutral engineering outcome.

Frequency-referenced multiplexing represents an important physical-layer development because it addresses the challenge of providing guaranteed bandwidth and low latency in a shared optical medium [6]. By anchoring transmission to stable frequency references, such approaches can reduce uncertainty in channel placement and improve determinism. For multi-tenant networks, this is valuable because tenants often require predictable performance rather than merely high average throughput. However, deterministic physical-layer mechanisms must be paired with deterministic allocation and admission control. If a tenant is granted a frequency reference but not guaranteed access to the corresponding spectrum during congestion, the physical guarantee is incomplete. Architectural foundations therefore include not only transponders and switches but also the policy engines, telemetry systems, and dispute-resolution processes that determine who gets which resources and under what conditions. The network becomes a computational and institutional artifact as much as a photonic one.

3. Frequency Resource Management as a Socio-Technical Control Problem

Frequency resource management is often framed as an optimization problem with constraints such as spectrum continuity, contiguity, capacity demand, and physical reach. That framing is necessary but insufficient. In multi-tenant networks, spectrum is a common-pool resource with rivalrous use and exclusion costs. One tenant's occupancy prevents another from using the same frequency slot on the same fiber, at least within the relevant interference domain. The operator must decide how much spectrum to reserve, how much to share dynamically, how to price access, and how to respond when demand exceeds supply. These decisions are not purely technical because they distribute economic opportunity, risk, and service quality across organizations [2]. A hospital network, a financial exchange, a video streaming service, and a research collaboration may all value latency and reliability differently. Treating them as interchangeable traffic flows obscures the social and contractual meanings of the services they support.

AI-assisted control adds another layer of complexity. Machine learning models can improve prediction and adaptation, but they also introduce dependencies on data quality, feature stability, and feedback loops. A model trained on historical traffic may perform well during normal conditions and poorly during emergencies, exactly when resource management matters most. Reinforcement learning can discover non-obvious allocation policies, but its reward function encodes value judgments about which outcomes matter. If the reward emphasizes revenue or average utilization, the resulting policy may neglect worst-case tenant experience. If it emphasizes strict fairness, it may sacrifice efficiency and increase energy use per bit. These trade-offs are not resolved by algorithmic sophistication; they require explicit

governance choices about priorities, thresholds, and acceptable forms of discrimination [15]. The control problem is therefore socio-technical because the controller's objectives are shaped by institutions and because the consequences of control are distributed across tenants and users.

Accountability mechanisms become essential when AI systems make or recommend allocation decisions. Tenants need to know whether their service degradation is caused by physical faults, market-based congestion, model error, or deliberate prioritization. Without such knowledge, trust in the multi-tenant platform erodes. Information accountability frameworks suggest that systems should maintain records sufficient to explain past decisions and support contestation [17]. In optical frequency management, this could mean logging spectrum assignments, model versions, input telemetry, confidence scores, and human overrides. Such records must be designed with privacy and commercial confidentiality in mind, because tenants may not want their traffic patterns exposed to competitors. The governance challenge is to provide enough transparency for fair operation while protecting legitimate secrets. That balance cannot be achieved by technical design alone; it requires policy, contracts, and independent oversight.

4. AI-Assisted Allocation, Prediction, and Adaptation

AI can contribute to frequency resource management across several functional areas. Traffic prediction is one of the most mature. Machine learning models can forecast demand at different time scales, allowing the operator to pre-allocate spectrum, adjust modulation, or reconfigure lightpaths before congestion occurs [9]. Classification models can estimate whether a proposed lightpath will meet a required quality of transmission, reducing the need for conservative margins [11]. Anomaly detection models can identify degradation before it becomes a hard failure, enabling proactive maintenance or rerouting [14]. Reinforcement learning can be used for sequential decision-making in routing and spectrum assignment, where the effect of current choices on future fragmentation is difficult to capture with static rules [10]. These capabilities are attractive because elastic optical networks generate large volumes of telemetry and because manual management cannot keep pace with dynamic multi-tenant demand.

However, AI-assisted allocation must confront the combinatorial nature of routing and spectrum assignment. The search space is large, constraints are coupled, and the optimal decision for one request may create blocking for future requests. Heuristic and exact methods have long been used for these problems [7], and AI does not remove the need for sound optimization foundations. Instead, AI can complement optimization by learning good initial solutions, predicting which constraints are likely to bind, or adapting parameters in response to changing conditions [12]. In some cases, learned models can approximate expensive physical-layer simulations, enabling faster what-if analysis [13]. This is particularly relevant in multi-tenant settings where the operator must evaluate many candidate allocations under different service-level agreements. Yet approximation introduces risk. If the model underestimates impairment or overestimates available margin, the network may accept a lightpath that later violates tenant expectations. Calibration, uncertainty quantification, and fallback mechanisms are therefore essential.

Adaptation is equally important. Multi-tenant traffic is non-stationary because tenants change applications, scale workloads, and enter or leave the infrastructure. Spectrum fragmentation accumulates as connections are established and torn down, reducing the effective capacity of the fiber even when total occupancy is moderate. Defragmentation techniques can reorganize

existing lightpaths to create larger contiguous blocks, but reconfiguration may cause temporary disruption and may violate service-level agreements if not carefully scheduled. AI can help identify when defragmentation is worth the risk, which tenants can tolerate reconfiguration, and how to sequence moves to minimize impact [11]. Frequency-referenced multiplexing may reduce some sources of uncertainty, but it does not eliminate the need for lifecycle management of spectrum [6]. The operator must decide whether to prioritize short-term tenant stability or long-term spectral efficiency. That decision is economic and political as much as technical.

A further challenge is the distribution of intelligence. Centralized AI controllers can optimize globally but may suffer from scalability limits, single points of failure, and governance concentration. Distributed or hierarchical controllers can respond faster and respect tenant boundaries, but they may produce suboptimal global outcomes and make accountability harder. Federated learning offers a middle path by allowing models to learn from tenant data without centralizing raw telemetry, but it introduces questions about model ownership, poisoning attacks, and equitable participation. In multi-tenant optical networks, the choice between centralized, distributed, and hybrid intelligence is architectural. It determines who can observe what, who can change allocations, and who bears the cost of error. The most robust designs are likely to combine global coordination for strategic planning with local autonomy for fast response, supported by clear escalation paths and human oversight.

5. Governance, Fairness, and Tenant Isolation

Governance is central to multi-tenant frequency resource management because allocation decisions distribute scarce capacity among organizations with unequal bargaining power. Fairness can be defined in multiple ways. Equal treatment gives each tenant the same access rules regardless of size or willingness to pay. Proportional fairness allocates resources in relation to demand or contribution. Priority-based fairness favors tenants with critical services or contractual guarantees. Each definition implies a different architecture. Strict isolation using dedicated spectrum simplifies guarantees but reduces statistical multiplexing gains. Dynamic sharing improves utilization but requires mechanisms to prevent one tenant from consuming capacity needed by others. A governance framework must specify which fairness principle applies, how it is measured, and what remedies exist when violations occur [15]. Without such specification, AI systems will optimize whatever proxy is convenient, often defaulting to revenue or average latency.

Tenant isolation is both a technical and a contractual requirement. At the optical layer, isolation can be achieved through dedicated fibers, dedicated spectrum blocks, or logical slices with reserved resources. At the control layer, isolation involves role-based access, policy separation, and telemetry boundaries. At the service layer, isolation means that one tenant's congestion or security incident does not degrade another's experience. These layers interact. A logically isolated slice may still share amplifiers and therefore be vulnerable to physical-layer effects. A dedicated spectrum block may still be affected by a faulty transponder or a misconfigured control plane. Multi-tenant design must therefore consider failure domains, not just nominal resource assignment. Governance should require regular isolation testing, transparent incident reporting, and clear liability when cross-tenant harm occurs.

AI introduces specific governance risks. Models can encode historical biases, such as favoring tenants that were previously prioritized or that generate cleaner data. They can also create accountability gaps when decisions emerge from complex ensembles rather than explicit rules.

Regulatory and ethical scholarship has emphasized the need for transparency, explainability, and contestability in algorithmic systems [16]. In optical networking, full explainability may be impractical for every micro-decision, but operators can provide meaningful explanations at the policy level. For example, a tenant might be told that a request was denied because the service class required protected spectrum and the available path had insufficient margin, rather than being given a raw model score. Such explanations are more useful if they connect to contractual categories and appeal processes. Governance should also address model updates, because an allocation policy may change without any physical change in the network. Tenants need notice, versioning, and the ability to audit changes that affect their service.

Cross-domain comparisons are instructive. Cloud computing platforms have developed sophisticated isolation, quota, and fair-use policies, though they continue to face criticism for opaque prioritization and lock-in. Wireless spectrum sharing has generated detailed regulatory frameworks for licensed and unlicensed access, dynamic spectrum access, and interference protection. Multi-tenant optical networks can learn from both. They need cloud-like automation and elasticity, but they also need spectrum-like rules for exclusive use, coordination, and enforcement. The most promising direction is a layered governance model in which physical-layer guarantees, control-plane policies, and contractual service levels are aligned. AI can support this model by providing prediction and adaptation, but it should not be the sole source of authority. Human institutions must set objectives, adjudicate conflicts, and update rules as technology and society change.

6. Deployment, Operations, and Sustainability

Deployment of AI-assisted frequency management occurs within existing operational and economic constraints. Optical networks are long-lived infrastructures with high capital costs and slow upgrade cycles. Operators cannot replace all hardware to support every new AI feature. They must integrate AI controllers with legacy network management systems, vendor-specific interfaces, and multi-vendor environments. This integration burden often determines whether an advanced algorithm is useful in practice. Data quality is another constraint. AI models require telemetry that may be incomplete, delayed, or inconsistent across domains. In multi-tenant settings, telemetry may also be commercially sensitive, limiting data sharing. Successful deployment therefore depends on standardized data models, open interfaces, and governance agreements that allow learning without violating tenant confidentiality [9]. The organizational dimension is equally important: network operations teams must trust and understand AI recommendations, and they need clear procedures for overriding them during incidents.

Sustainability adds a further dimension. Optical transport is often more energy-efficient per bit than electronic switching, but total traffic growth can still increase energy consumption. Data centers, cooling, amplifiers, transponders, and control infrastructure all contribute to the lifecycle footprint. AI itself consumes energy for training and inference, and large-scale optimization may shift energy use from transmission to computation [18]. A narrow focus on spectral efficiency can miss these broader effects. For example, aggressive defragmentation may reduce blocking but require more frequent reconfiguration and signaling, increasing control-plane overhead. Predictive routing may improve utilization but keep idle capacity powered for anticipated demand. Sustainable multi-tenant management should evaluate energy per delivered service, not only energy per bit, and should consider embodied carbon in hardware. AI can help by forecasting renewable energy availability, shifting non-urgent

traffic, and optimizing sleep modes, but these strategies must be reconciled with tenant service guarantees.

Deployment also raises questions of market structure. If a small number of infrastructure providers control AI-driven allocation, they may gain disproportionate influence over which services receive priority. Regulators may need to require nondiscrimination, transparency, and interoperability. At the same time, overly prescriptive rules can stifle innovation and increase costs. The policy challenge is to define outcome-based requirements, such as fair access, reliability, and security, while leaving implementation flexible. Multi-tenant optical networks can become critical public infrastructure, comparable to electricity and telecommunications. Their governance should reflect that status without freezing the technical evolution that makes them valuable. AI-assisted management should be deployed with staged trials, independent evaluation, and mechanisms for correcting harmful allocations. Sustainability and fairness are not peripheral concerns; they are criteria for whether the system deserves long-term social license.

7. Robustness, Resiliency, and Failure Management

Robustness in multi-tenant optical networks requires resilience against failures, attacks, and extreme demand. Optical fibers can be cut, amplifiers can fail, transponders can drift, and control planes can be misconfigured. In a shared infrastructure, the consequences of failure can propagate across tenants, making resilience a collective as well as individual concern. AI can improve resilience by predicting failures, detecting anomalies, and recommending restoration paths [14]. Machine learning models can analyze telemetry from optical performance monitoring to distinguish gradual degradation from sudden faults [11]. They can also classify root causes, reducing mean time to repair. However, resilience cannot rely solely on prediction. Networks need redundancy, diversity, and graceful degradation. AI should support these mechanisms, not replace them. If a model fails during a major incident, operators must still have deterministic fallback procedures.

Frequency resource management under failure is especially difficult because restoration may require reallocating spectrum across multiple tenants. A protected service may need reserved capacity, while a best-effort tenant may be temporarily throttled. These decisions involve contractual obligations and may need regulatory oversight if critical services are affected. AI systems should be designed to operate under uncertainty, with conservative behavior when confidence is low. They should also be resistant to adversarial manipulation. An attacker might inject false telemetry to mislead traffic prediction, cause unnecessary rerouting, or induce fragmentation. Security for AI-assisted optical networks therefore includes data integrity, model robustness, access control, and audit trails. The governance frameworks discussed earlier apply here as well: incident logs, human escalation, and independent review are necessary for accountability.

Resilience also has a temporal dimension. Short outages may be tolerable for some tenants but catastrophic for others. Multi-tenant architectures should allow differentiated protection classes, but these classes must be priced and governed fairly. AI can help match protection levels to actual need, avoiding over-provisioning for tenants that do not require it. Yet the operator must be careful not to create hidden tiers that undermine public interest obligations. For example, emergency services, hospitals, and public safety networks may require guaranteed priority regardless of market power. Embedding such priorities into AI reward functions is a policy decision, not a technical one. Robustness therefore connects directly to

fairness and governance. A network that is efficient in normal times but unfair in crises is not resilient in the sense that matters for society.

8. Policy Implications and Cross-Domain Lessons

The policy implications of AI-assisted frequency management extend beyond network operations. Spectrum and optical infrastructure are often regulated as essential facilities, and multi-tenancy can blur the boundary between infrastructure provider and service provider. If the same entity controls the physical layer, the AI controller, and the service platform, it may have incentives and capabilities to favor its own offerings. Regulators may need to require functional separation, transparent interfaces, or nondiscriminatory access. They may also need to mandate reporting on allocation outcomes, such as blocking rates by tenant class, latency distribution, and restoration performance. Such reporting can be designed to protect commercial confidentiality while enabling oversight. The goal is not to micromanage algorithms but to ensure that market power and technical opacity do not undermine fair competition or public interest.

Cross-domain lessons from cloud computing, electricity markets, and wireless spectrum suggest that governance must be adaptive. Static rules cannot keep pace with AI capabilities or changing traffic patterns. Regulatory sandboxes, periodic audits, and multi-stakeholder advisory bodies can help. Information accountability provides a useful principle: systems should be able to explain decisions after the fact and allow affected parties to contest them [17]. Ethical analysis of algorithms warns against reducing complex social values to single metrics [15]. Governance of AI more broadly requires legal, technical, and ethical coordination [16]. For optical networks, this means that performance metrics should be complemented by procedural safeguards. Who decides the reward function? How are tenants represented? What happens when the model harms a tenant? These questions are not peripheral to engineering; they determine whether the system can be trusted at scale.

Policy should also address sustainability. If AI optimization reduces energy use in one domain but increases it in another, regulators may need lifecycle accounting and disclosure. They may encourage shared infrastructure to avoid duplication, while guarding against concentration. They may support research into energy-aware allocation and low-carbon operations [18]. In addition, policy should promote interoperability so that tenants can move between providers without prohibitive switching costs. Multi-tenant optical networks can support competition and innovation, but only if access is fair and transparent. AI-assisted management should be evaluated not only by spectral efficiency but by its contribution to reliable, affordable, and sustainable connectivity. That evaluation requires collaboration among engineers, economists, legal scholars, and policymakers. The technical community cannot settle these questions alone, but it can design systems that make governance possible.

9. Conclusion

AI-assisted frequency resource management in multi-tenant optical networks is a system-level challenge that joins photonics, machine learning, operations research, economics, and governance. Elastic optical networking provides the flexibility needed to serve heterogeneous tenants, but it also creates coordination problems that cannot be solved by spectrum assignment alone. AI can improve prediction, allocation, adaptation, and failure management, yet its use introduces questions of bias, opacity, accountability, and robustness. The most credible architectures will combine physical-layer innovations, such as frequency-referenced multiplexing, with policy engines that enforce fairness and isolation. They will use centralized

and distributed intelligence selectively, maintain human oversight, and provide audit trails that allow decisions to be explained and contested. Deployment must account for legacy systems, data quality, market structure, and sustainability. Resilience must include graceful degradation and security against adversarial manipulation. Policy must ensure that essential infrastructure remains open, nondiscriminatory, and aligned with public interest. Ultimately, frequency resource management is not merely a matter of optimizing slots and routes. It is a matter of deciding how shared optical infrastructure will serve a pluralistic digital society. AI can assist that decision, but it cannot replace the institutional choices that give the network its purpose.

References

1. Kani, J., Bourgart, F., Cui, A., Rafel, A., Campbell, M., Davey, R., & Rodrigues, S. (2009). Next-generation PON-part I: Technology roadmap and general requirements. *IEEE Communications Magazine*, 47(11), 43-49. <https://doi.org/10.1109/MCOM.2009.5307465>
2. Talebi, S., Alam, M. J., Katib, I., Khamis, M., Salama, R., & Rouskas, G. N. (2014). Spectrum management techniques for elastic optical networks: A survey. *Optical Switching and Networking*, 13, 34-48. <https://doi.org/10.1016/j.osn.2014.02.003>
3. Jinno, M., Takara, H., Kozicki, B., Tsukishima, Y., Sone, Y., & Matsuoka, S. (2009). Spectrum-efficient and scalable elastic optical path network: Architecture, benefits, and enabling technologies. *IEEE Communications Magazine*, 47(11), 66-73. <https://doi.org/10.1109/MCOM.2009.5307468>
4. Gerstel, O., Jinno, M., Lord, A., & Yoo, S. J. B. (2012). Elastic optical networking: A new dawn for the optical layer? *IEEE Communications Magazine*, 50(2), s12-s20. <https://doi.org/10.1109/MCOM.2012.6146481>
5. Napoli, A., Bohn, M., Rafique, D., Stavdas, A., Sambo, N., Poti, L., Nölle, M., Fischer, J. K., Riccardi, E., Pagano, A., Di Giglio, A., Moreno, F., Temprana, E., & Tomkos, I. (2015). Next generation elastic optical networks: The vision of the European research project IDEALIST. *IEEE Communications Magazine*, 53(2), 152-162. <https://doi.org/10.1109/MCOM.2015.7045404>
6. Zhou, Z., Wei, J., Luo, Y., Clark, K. A., Sillescu, E., Deakin, C., ... & Liu, Z. (2023). Communications with guaranteed bandwidth and low latency using frequency-referenced multiplexing. *Nature Electronics*, 6(9), 694-702.
7. Klinkowski, M., & Walkowiak, K. (2011). Routing and spectrum assignment in spectrum sliced elastic optical path network. *IEEE Communications Letters*, 15(8), 884-886. <https://doi.org/10.1109/LCOMM.2011.060811.110281>
8. Christodoulopoulos, K., Tomkos, I., & Varvarigos, E. A. (2011). Elastic bandwidth allocation in flexible OFDM-based optical networks. *Journal of Lightwave Technology*, 29(9), 1354-1366. <https://doi.org/10.1109/JLT.2011.2125777>
9. Boutaba, R., Salahuddin, M. A., Limam, N., Ayoubi, S., Shahriar, N., Estrada-Solano, F., & Caicedo, O. M. (2018). A comprehensive survey on machine learning for networking: Evolution, applications and research opportunities. *Journal of Internet Services and Applications*, 9(1), 1-99. <https://doi.org/10.1186/s13174-018-0087-2>

10. Fadlullah, Z. M., Tang, F., Mao, B., Kato, N., Akashi, O., Inoue, T., & Mizutani, K. (2017). State-of-the-art deep learning: Evolving machine intelligence toward tomorrow's intelligent network traffic control systems. *IEEE Communications Surveys & Tutorials*, 19(4), 2432-2455. <https://doi.org/10.1109/COMST.2017.2707140>
11. Musumeci, F., Rottondi, C., Nag, A., Macaluso, I., Zibar, D., Ruffini, M., & Tornatore, M. (2019). An overview on application of machine learning techniques in optical communication. *Journal of Lightwave Technology*, 37(3), 710-723. <https://doi.org/10.1109/JLT.2018.2889519>
12. Mata, J., de Miguel, I., Durán, R. J., Merayo, N., Singh, S. K., Jukan, A., & Chamania, M. (2018). Artificial intelligence (AI) methods in optical networks: A comprehensive survey. *Optical Switching and Networking*, 28, 43-57. <https://doi.org/10.1016/j.osn.2018.01.002>
13. Zibar, D., Piels, M., Jones, R., & Schaeffer, G. (2016). Machine learning concepts for fiber-optic communications. *Journal of Lightwave Technology*, 34(17), 4110-4117. <https://doi.org/10.1109/JLT.2016.2569106>
14. Vela, A. P., Ruiz, M., Fresi, F., Sambo, N., Cugini, F., Meloni, G., Poti, L., Velasco, L., & Castoldi, P. (2018). BER degradation detection and failure identification in elastic optical networks. *Journal of Lightwave Technology*, 36(20), 4595-4604. <https://doi.org/10.1109/JLT.2018.2859978>
15. Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2), 1-21. <https://doi.org/10.1177/2053951716679679>
16. Cath, C. (2018). Governing artificial intelligence: Ethical, legal and technical opportunities and challenges. *Philosophical Transactions of the Royal Society A*, 376(2133), 20180080. <https://doi.org/10.1098/rsta.2018.0080>
17. Weitzner, D. J., Abelson, H., Berners-Lee, T., Feigenbaum, J., Hendler, J., & Sussman, G. J. (2008). Information accountability. *Communications of the ACM*, 51(6), 82-87. <https://doi.org/10.1145/1349026.1349043>
18. Baliga, J., Ayre, R. W. A., Hinton, K., & Tucker, R. S. (2011). Green cloud computing: Balancing energy in processing, storage, and transport. *Proceedings of the IEEE*, 99(1), 149-167. <https://doi.org/10.1109/JPROC.2010.2060451>