

Adaptive Modulation and Wavelength Assignment for High-Capacity Optical Networks Using Deep Reinforcement Learning

Adya J. Tandon

Department of Computer Science, Binghamton University, Binghamton, NY, USA.
adya262@binghamton.edu

Abstract

High-capacity optical networks are confronting a structural inflection point in which static provisioning, conservative margins, and human-designed control loops are increasingly misaligned with dynamic traffic, heterogeneous service classes, and stringent sustainability constraints. Adaptive modulation and wavelength assignment have emerged as central mechanisms for matching transmission formats and spectral resources to time-varying channel conditions, yet their operational value depends on whether decisions can be made rapidly, robustly, and equitably across multi-vendor, multi-domain infrastructures. This paper examines deep reinforcement learning as a system-level control paradigm for adaptive modulation and wavelength assignment in high-capacity optical networks. Rather than treating learning as a narrow optimization patch, the analysis frames it as an architectural layer that interacts with physical-layer uncertainty, network governance, deployment economics, energy intensity, and fairness among tenants and regions. The discussion emphasizes structural trade-offs among spectral efficiency, latency, resilience, operational complexity, and policy accountability. It also considers how simulation-to-reality gaps, reward misspecification, and fragmented telemetry can undermine learning agents, and why observability, safety envelopes, and human oversight remain essential. Through conceptual analysis, case illustrations, and cross-domain comparisons with data-center interconnects and cloud resource management, the paper argues that deep reinforcement learning can improve wavelength utilization and modulation adaptivity only when embedded within transparent, sustainable, and equitable network governance. The conclusion outlines a research agenda for trustworthy learning-based optical control.

Keywords

optical networks, adaptive modulation, wavelength assignment, deep reinforcement learning, resource allocation, network governance, sustainability, robustness.

1. Introduction

High-capacity optical networks have become the computational and communicative substrate of contemporary digital society, carrying traffic that ranges from latency-sensitive financial transactions and industrial control to high-volume data replication and distributed artificial intelligence training. The classical design assumption that traffic matrices change slowly and that worst-case margins can be engineered once and then forgotten has become increasingly untenable. Coherent transmission, flexible grid architectures, and software-defined control have expanded the decision space, but they have also introduced combinatorial allocation problems that resist static planning. Capacity limits, nonlinear propagation, and the interplay between modulation format and reach remain fundamental constraints [1], while the evolution

of fiber-optic transmission and networking continues to shift the boundary between physical-layer capability and network-level intelligence [2]. In this environment, adaptive modulation and wavelength assignment are no longer peripheral tuning functions; they are constitutive elements of how capacity, latency, and reliability are produced and distributed.

The need for adaptivity is further reinforced by the diversity of service requirements. A cloud provider may prioritize aggregate throughput, a financial network may prioritize deterministic latency, and a public research infrastructure may prioritize equitable access across institutions. The flexible optical networking paradigm promised to reconcile such diversity through tunable spectrum allocation and reconfigurable transceivers [6], but flexible hardware alone does not resolve the control problem. Routing, modulation, and spectrum assignment must be coordinated across layers, domains, and time horizons, often under incomplete information. Classical optimization methods can produce high-quality plans in offline settings, yet they struggle when traffic, impairments, and failures evolve faster than the planning cycle. This gap has motivated interest in machine learning, and particularly in deep reinforcement learning, as a method for learning control policies from interaction rather than relying solely on explicit models [7], [8].

Deep reinforcement learning is attractive because it can approximate complex policies in high-dimensional state spaces and can optimize long-term objectives rather than myopic rewards. In optical networks, an agent could learn to select modulation formats, assign wavelengths, reserve spectral slots, and reroute traffic in response to telemetry from transceivers, amplifiers, and monitors. Such a policy could, in principle, reduce margins, improve spectral efficiency, and respond to failures more quickly than rule-based systems. However, the translation from simulated optical environments to production networks is not merely an engineering detail. It raises questions about interpretability, safety, accountability, and the distribution of benefits. If learning agents optimize aggregate throughput while silently degrading service for low-volume or remote users, they may reproduce or amplify existing inequities. If they consume substantial computational resources during training, they may undercut the sustainability gains they promise. If they remain opaque, they complicate regulatory oversight and multi-operator coordination.

This paper therefore approaches deep reinforcement learning for adaptive modulation and wavelength assignment as a socio-technical systems problem. It does not present a new algorithm or a mathematical derivation. Instead, it develops a structured analysis of architecture, trade-offs, governance, deployment, sustainability, robustness, and fairness. The central argument is that learning-based optical control can be valuable, but only when it is treated as an accountable infrastructure layer rather than as a black-box optimizer. The next section outlines the system architecture and problem context. Subsequent sections examine deep reinforcement learning as a control layer, structural trade-offs, robustness and fairness, deployment and sustainability, case illustrations, and open challenges.

2. Architecture and System-Level Problem

A high-capacity optical network can be understood as a layered socio-technical system in which physical transmission, optical switching, electronic grooming, control planes, and business agreements interact. Adaptive modulation determines how many bits per symbol can be transmitted over a given lightpath, subject to reach, signal-to-noise ratio, nonlinear interference, and aging infrastructure. Wavelength assignment determines which spectral resources are allocated to which demands, and it must respect continuity, fragmentation, regeneration, and protection constraints. Together, these decisions influence the effective

capacity of the network, the latency experienced by packets, the energy consumed by transceivers and amplifiers, and the resilience of services under failure. The problem is not simply one of packing more wavelengths into fiber. It is a problem of coordinating physical-layer feasibility with service-level expectations and operational policies.

The control architecture for such coordination may be centralized, distributed, or hybrid. A centralized controller can collect telemetry from multiple domains and optimize globally, but it introduces single points of failure, scalability limits, and governance concerns related to data ownership. A distributed approach can improve resilience and local responsiveness, yet it may produce suboptimal global outcomes and complicate coordination across administrative boundaries. Hybrid architectures attempt to combine global policy with local autonomy, but they require well-defined interfaces, trust mechanisms, and fallback rules. Deep reinforcement learning can operate at any of these levels, but its suitability depends on the timescale of decisions. Fast modulation adaptation may need to occur at the transceiver or line-card level, while wavelength assignment across domains may require coordination over longer intervals. The architecture must therefore specify what is learned, where learning occurs, and how learned actions are validated before they affect live traffic [10].

A further structural issue is observability. Reinforcement learning agents depend on state representations, and optical networks provide noisy, partial, and delayed measurements. Transceiver metrics may be unavailable for unestablished lightpaths, and quality-of-transmission prediction may require models that generalize beyond historical configurations [18]. Telemetry may be aggregated differently across vendors, domains, and generations of equipment. If the agent sees only a simplified state, it may learn policies that exploit simulation artifacts rather than genuine physical relationships. If it sees too much state, it may overfit to transient conditions and become brittle. The design of the observation space is therefore a governance decision as much as a technical one, because it determines which phenomena are visible to the optimizer and which remain externalized. The network architecture must include mechanisms for uncertainty estimation, anomaly detection, and human review, so that learning does not become a substitute for accountability.

3. Deep Reinforcement Learning as a Control Layer

Deep reinforcement learning combines reinforcement learning principles with deep neural networks that approximate policies or value functions. The agent observes a state, selects an action, receives a reward, and updates its parameters to improve long-term performance [8]. In optical networks, the state might include traffic demand, lightpath quality, spectral occupancy, queue lengths, failure indicators, and energy prices. Actions might include modulation format selection, wavelength assignment, spectrum defragmentation, rerouting, or admission control. Rewards might reflect throughput, latency, blocking probability, energy consumption, or a weighted combination of these objectives. The appeal of this framework is that it can learn sequential policies without requiring a complete analytical model of the network. It can adapt to changing conditions and can discover non-obvious strategies that human operators might not design explicitly [9].

Several studies have demonstrated the potential of deep reinforcement learning for optical resource allocation. Routing, modulation, and spectrum assignment have been formulated as sequential decision problems in which an agent learns to allocate resources in elastic optical networks [11]. Routing in optical transport networks has also been addressed with deep reinforcement learning, showing that learned policies can outperform heuristic baselines under certain traffic patterns [12]. More recent work has explored graph neural networks to

improve generalization across topologies and demand distributions, indicating that representation learning can reduce the need for retraining when network conditions change [13]. These advances suggest that learning-based control is technically feasible, but they also highlight a tension between specialization and generalization. A policy trained on one topology may fail on another, and a policy optimized for one traffic mix may behave poorly when demand shifts.

The control layer must also address reward design. A reward function is not a neutral technical choice; it encodes priorities. If throughput dominates, the agent may accept higher latency or lower fairness. If latency dominates, the agent may underutilize spectrum. If energy is included, the agent may trade off performance against carbon intensity, but only if energy data are accurate and timely. Multi-objective reinforcement learning can represent trade-offs more explicitly, but it does not eliminate the need for value judgments. Network operators, regulators, and users may disagree about the appropriate balance. A trustworthy system must therefore expose reward weights, constraints, and performance envelopes to human oversight. It must also provide fallback policies when the learned agent encounters out-of-distribution conditions. The control layer is thus not merely a neural network; it is a governance arrangement that includes objective setting, monitoring, auditing, and intervention.

4. Structural Trade-offs in Adaptive Modulation and Wavelength Assignment

Adaptive modulation offers a clear trade-off between spectral efficiency and reach. Higher-order modulation formats transmit more bits per symbol but are more susceptible to noise and nonlinearity, reducing the distance over which they can be used without regeneration. Lower-order formats are more robust but consume more spectrum for a given data rate. In a high-capacity network, the optimal choice depends on the specific lightpath, the condition of the fiber, the presence of amplifiers, and the interference from other channels. Wavelength assignment adds another layer of complexity because it determines which channels interact and how spectrum is fragmented. A modulation decision that is optimal for one wavelength may be suboptimal for another due to cross-phase modulation, four-wave mixing, or amplifier gain ripple. The system-level problem is therefore not separable; modulation and wavelength assignment must be considered together [5], [6].

Deep reinforcement learning can navigate this coupled space by learning policies that account for both decisions. However, the learned policy may exploit simulator inaccuracies. If the simulation underestimates nonlinear interference, the agent may choose aggressive modulation formats that fail in deployment. If the simulation overestimates component reliability, the agent may underprovision protection. The sim-to-real gap is a central challenge. Domain randomization, conservative safety constraints, and digital twins can mitigate the gap, but they cannot eliminate it. The network must include runtime verification that pauses or overrides learned actions when measured performance deviates from predictions. This is especially important in high-capacity systems where a single misconfiguration can affect many services. The trade-off between exploration and exploitation also has operational consequences: an agent that explores too aggressively may cause outages, while an agent that exploits too narrowly may fail to adapt.

Another trade-off concerns centralized versus local optimization. Centralized learning can coordinate wavelength assignment across the entire network, reducing fragmentation and improving global utilization. But centralization requires extensive telemetry and low-latency control channels, and it may create bottlenecks. Local learning can respond quickly to transceiver conditions, but it may not see the global consequences of its actions. Hybrid

approaches that combine global policy with local safety filters are promising, but they require clear authority boundaries. Who is responsible when a learned action violates a service-level agreement? How are conflicts resolved between domains? These questions are not purely technical. They concern liability, trust, and the distribution of operational control. A network that delegates too much authority to opaque agents may become difficult to govern, while a network that constrains agents too tightly may forgo the benefits of adaptivity.

5. Robustness, Fairness, and Governance

Robustness in learning-based optical control has multiple dimensions. The agent must be robust to noisy telemetry, partial failures, adversarial manipulation, and distributional shift. It must also be robust to its own errors. A learned policy that performs well on average may produce rare catastrophic decisions. In optical networks, rare events matter because a single wavelength misassignment or modulation overreach can cause widespread service degradation. Robustness therefore requires redundancy, conservative fallbacks, and continuous monitoring. It also requires interpretability. Operators need to understand why the agent chose a particular action, especially when the action is counterintuitive. Explainable artificial intelligence methods can provide post hoc explanations, but explanations are not a substitute for designed transparency. The control system should log state, action, reward, and outcome in a way that supports auditing and root-cause analysis [14], [16].

Fairness is often overlooked in optical network optimization because the primary metrics are aggregate throughput and blocking probability. Yet fairness matters in several ways. First, tenants with different service-level agreements may receive different levels of protection. If the learning agent prioritizes high-revenue traffic, it may degrade best-effort or public-interest traffic. Second, geographic fairness matters. Rural or underserved regions may have fewer parallel paths and less redundancy, making them more vulnerable to optimization that concentrates resources on dense routes. Third, temporal fairness matters. A policy that smooths performance over time may still create short outages that harm latency-sensitive applications. Fairness constraints can be incorporated into reinforcement learning through constrained optimization, reward shaping, or separate auditing. But the choice of fairness metric is normative. It should be made through inclusive governance processes rather than left to engineers alone.

Governance also concerns data and model ownership. Deep reinforcement learning for optical networks depends on telemetry that may span multiple operators, vendors, and jurisdictions. Who owns the data? Who can train on it? Who can audit the trained model? If models are proprietary, regulators may lack the ability to inspect their behavior. If data are siloed, learning may be biased toward the domains with the most data. Multi-party computation, federated learning, and standardized interfaces could help, but they introduce complexity and potential vulnerabilities. A governance framework should define roles for operators, vendors, regulators, and users. It should include requirements for documentation, impact assessment, incident reporting, and redress. Without such a framework, the adoption of deep reinforcement learning may exacerbate existing asymmetries of information and power.

6. Deployment, Infrastructure, and Sustainability

Deployment of learning-based adaptive modulation and wavelength assignment will occur incrementally rather than as a wholesale replacement of existing control systems. Operators must integrate new agents with legacy equipment, operations support systems, and business support systems. This integration is often more difficult than training the model. Data

pipelines must be reliable, control interfaces must be secure, and fallback procedures must be tested. The deployment strategy should begin with bounded use cases where the action space is limited and the consequences of error are manageable. For example, an agent might recommend modulation format changes for review by a human operator, or it might adjust wavelength assignments within a single domain before expanding to multi-domain coordination. Such staged deployment allows operators to build trust and gather evidence. It also creates opportunities for regulatory sandboxes where policies can be evaluated without risking critical infrastructure.

Sustainability is an increasingly important consideration. Optical networks already consume significant energy through transceivers, amplifiers, cooling, and switching. Adaptive modulation can reduce energy per bit by selecting efficient formats when channel conditions allow, but it can also increase energy consumption if it requires more frequent reconfiguration or if it encourages higher-order formats that demand more powerful digital signal processing. Wavelength assignment can reduce energy by minimizing regeneration and routing traffic through the most efficient paths, but it can also concentrate traffic in ways that increase cooling demands. Deep reinforcement learning training itself consumes computational resources, and the carbon intensity of that training depends on data-center location and energy mix. A holistic sustainability assessment must account for embodied emissions, operational emissions, and the rebound effect, where efficiency gains lead to higher overall traffic and energy use [2]. Recent work on low-latency wavelength-switched intra-data center interconnects illustrates how physical-layer choices such as hollow-core fiber can alter latency and energy trade-offs, reminding us that learning-based control cannot compensate for unsustainable infrastructure decisions [15].

Infrastructure policy also shapes deployment. Governments may incentivize broadband expansion, fund research testbeds, or mandate open interfaces. These policies affect the data available for learning and the markets in which learned control can be deployed. If open interfaces are required, third-party developers can create competing agents, but interoperability and security must be ensured. If infrastructure is concentrated in a few providers, learning-based optimization may reinforce market power unless regulated. Sustainability policy may require energy reporting or carbon caps, which can be incorporated into reward functions. However, policy must be carefully designed to avoid perverse incentives. For example, a narrow focus on operational energy might encourage shifting emissions to manufacturing or to less regulated regions. A socio-technical approach to deployment must therefore consider the full life cycle and the distribution of costs and benefits [17], [19].

7. Case Illustrations and Cross-Domain Comparisons

One illustrative case is a national research and education network that connects universities, laboratories, and supercomputing centers. Such networks often have diverse traffic patterns, including large data transfers, interactive remote instrumentation, and real-time collaboration. A deep reinforcement learning agent could adapt modulation and wavelength assignment to prioritize latency-sensitive flows while batching bulk transfers into off-peak windows. However, the network also has a public-interest mission, so fairness across institutions matters. A policy that optimizes aggregate throughput might starve smaller institutions with less traffic. A better design would include fairness constraints and transparent reporting. The agent might also need to coordinate with domain-specific controllers, because each campus

may have its own policies. This case shows that technical optimization must be embedded in institutional governance.

A second case is a cloud provider operating intra-data center and inter-data center optical interconnects. Here, traffic is highly dynamic, and latency directly affects application performance. Adaptive modulation can help manage reach and capacity, while wavelength assignment can reduce contention. Deep reinforcement learning has been explored for routing and resource allocation in such environments [12], [13]. Yet the provider must also consider tenant isolation, security, and energy costs. A learned policy that reroutes traffic to save energy might violate tenant expectations if it increases latency or jitter. The provider may therefore use multi-objective rewards and strict service-level constraints. The hollow-core fiber example mentioned earlier highlights that physical-layer latency improvements can change the feasible region for learning-based control [15]. If the underlying fiber has lower latency, the agent may have more flexibility in modulation and wavelength choices. This illustrates the co-evolution of infrastructure and control intelligence.

Cross-domain comparisons can illuminate these issues. In cloud computing, reinforcement learning has been used for resource management, autoscaling, and energy-aware scheduling. The challenges are similar: high-dimensional state spaces, delayed rewards, and the need for safety. Cloud systems have developed sophisticated observability, rollback, and canary deployment practices that optical networks could adapt. In power grids, reinforcement learning is being explored for demand response and voltage control, where safety and regulatory constraints are paramount. Power grids have strong traditions of contingency analysis and operator training, which are relevant to optical network control. In transportation, reinforcement learning for traffic signal control has faced challenges in fairness and public accountability, because optimizing average travel time can disadvantage certain neighborhoods. These comparisons suggest that optical networks should not treat deep reinforcement learning as a purely technical upgrade. They should learn from domains that have already grappled with the social implications of automated infrastructure control [14], [16].

8. Open Challenges and Research Agenda

Several open challenges remain. First, there is a need for standardized simulation environments that accurately represent optical physics, hardware behavior, and operational constraints. Without such environments, results are difficult to compare, and policies may not transfer. Second, there is a need for methods that can guarantee safety under uncertainty. Probabilistic guarantees, shielding, and runtime enforcement are promising, but they must be integrated with optical network control loops. Third, there is a need for multi-agent reinforcement learning that can coordinate across domains without exposing sensitive data. Federated learning and secure multi-party computation may help, but their overhead and security properties must be evaluated in realistic settings. Fourth, there is a need for interpretable models that operators can trust. Post hoc explanations are useful, but they may not capture the causal structure of optical impairments. Causal inference and digital twins could provide more reliable insights [3], [4].

A fifth challenge is the evaluation of fairness and sustainability. Current benchmarks often focus on throughput and blocking probability. Future benchmarks should include energy consumption, carbon intensity, latency distribution, and fairness metrics across tenants and regions. They should also include adversarial scenarios and rare failure events. Sixth, there is a need for governance frameworks that define accountability for learned actions. Who is

liable when an agent makes a decision that causes an outage? How are redress mechanisms designed? How can regulators audit proprietary models? These questions require interdisciplinary collaboration among engineers, legal scholars, economists, and policymakers. Seventh, there is a need for workforce development. Network operators will need new skills in machine learning, data engineering, and responsible artificial intelligence. Educational programs should integrate these skills with traditional optical networking knowledge [10], [18].

Finally, there is a need for long-term field trials. Laboratory results are valuable, but they cannot capture the full complexity of production networks. Field trials should be designed with clear metrics, independent evaluation, and mechanisms for sharing lessons learned. They should also include diverse stakeholders, including users who may be affected by automated decisions. Such trials can help build confidence and identify unintended consequences. They can also inform standards development and regulatory policy. The research agenda should therefore be both technically ambitious and socially grounded. The goal is not merely to build agents that can optimize wavelength assignment, but to build infrastructures that are high-capacity, resilient, sustainable, and fair.

9. Conclusion

Adaptive modulation and wavelength assignment are essential for high-capacity optical networks, and deep reinforcement learning offers a powerful approach for controlling them in dynamic environments. The technology can learn policies that coordinate modulation, spectrum, routing, and protection in ways that static rules cannot. However, its value depends on how it is embedded in the broader system. The analysis in this paper has emphasized that deep reinforcement learning is not a neutral optimization tool. It encodes priorities, depends on observability, and interacts with governance, deployment, sustainability, and fairness. Structural trade-offs among spectral efficiency, latency, resilience, and energy cannot be resolved by algorithms alone. They require human judgment, transparent objectives, and inclusive decision-making.

The paper has also highlighted the importance of robustness and accountability. Learning agents must operate under uncertainty, with noisy data and out-of-distribution conditions. They must be monitored, audited, and overridden when necessary. The deployment of such agents should be staged, with bounded action spaces and human oversight. Sustainability must be assessed holistically, including training costs and infrastructure life cycles. Fairness must be defined and measured, not assumed. Governance must clarify data ownership, model auditing, liability, and redress. Without these elements, learning-based optical control may improve some metrics while creating new risks and inequities.

A forward-looking research agenda should prioritize standardized simulation, safe reinforcement learning, multi-agent coordination, interpretability, fairness benchmarks, sustainability metrics, and field trials. It should also foster interdisciplinary collaboration and workforce development. If these challenges are addressed, deep reinforcement learning can help optical networks become more adaptive, efficient, and resilient. If they are ignored, the result may be brittle automation that concentrates power and externalizes costs. The choice is not between learning and no learning, but between learning as a black-box optimizer and learning as an accountable infrastructure capability. The latter is the more difficult and more necessary path.

References

1. Essiambre, R. J., Kramer, G., Winzer, P. J., Foschini, G. J., & Goebel, B. (2010). Capacity limits of optical fiber networks. *Journal of Lightwave Technology*, 28(4), 662–701.
2. Winzer, P. J., Neilson, D. T., & Chraplyvy, A. R. (2018). Fiber-optic transmission and networking: The previous 20 and the next 20 years. *Optics Express*, 26(18), 24190–24239.
3. Agrawal, G. P. (2021). *Fiber-optic communication systems* (5th ed.). Wiley.
4. Poggiolini, P. (2012). The GN model of non-linear propagation in uncompensated coherent optical systems. *Journal of Lightwave Technology*, 30(24), 3857–3879.
5. Sambo, N., Cugini, F., Bottari, G., Iovanna, P., & Castoldi, P. (2012). Routing, code, and spectrum assignment in flexi-grid optical networks. *Journal of Lightwave Technology*, 30(4), 572–579.
6. Tomkos, I., Azodolmolky, S., Sole-Pareta, J., Careglio, D., & Palkopoulou, E. (2014). A tutorial on the flexible optical networking paradigm: State of the art, trends, and research challenges. *Proceedings of the IEEE*, 102(9), 1317–1337.
7. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529–533.
8. Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd ed.). MIT Press.
9. Arulkumaran, K., Deisenroth, M. P., Brundage, M., & Bharath, A. A. (2017). Deep reinforcement learning: A brief survey. *IEEE Signal Processing Magazine*, 34(6), 26–38.
10. Zhou, D., & Yi, X. (2026). DISTRIBUTED AI-ASSISTED RISK ASSESSMENT FOR ZERO-TRUST MICROSERVICES: A TAXONOMY AND CONCEPTUAL FRAMEWORK. *AI and Data Science Journal*, 3(2), 1-13.
11. Chen, X., Guo, J., Zhu, Z., Proietti, R., Castro, A., & Yoo, S. J. B. (2019). Deep-RMSA: A deep reinforcement learning framework for routing, modulation and spectrum assignment in elastic optical networks. *Journal of Lightwave Technology*, 37(16), 4065–4073.
12. Suárez-Varela, J., Mestres, A., Yu, J., Kuang, L., Feng, H., Barlet-Ros, P., & Cabellos-Aparicio, A. (2019). Routing in optical transport networks with deep reinforcement learning. *Journal of Optical Communications and Networking*, 11(11), 547–558.
13. Almasan, P., Suárez-Varela, J., Badia-Sampera, A., Rusek, K., Alcaraz, J., Argente, E., Cabellos-Aparicio, A., & Barlet-Ros, P. (2022). Deep reinforcement learning meets graph neural networks: Exploring a routing optimization use case. *Computer Communications*, 191, 407–419.
14. 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), 16.

15. Zhou, Zichuan, et al. "Low-latency wavelength-switched clock-synchronized intra-data center interconnects enabled by hollow core fiber." *Optics Express* 31.15 (2023): 24739-24748.
16. Mata, J., de Miguel, I., Durán, R. J., Merayo, N., Singh, S. K., Jukan, A., & Chamanian, M. (2018). Artificial intelligence (AI) methods in optical networks: A comprehensive survey. *Optical Switching and Networking*, 28, 43–57.
17. Musumeci, F., Rottondi, C., Nag, A., Macaluso, I., Tornatore, M., & Pattavina, A. (2018). A survey on application of machine learning techniques in optical networks. *IEEE Communications Surveys & Tutorials*, 20(3), 2134–2160.
18. Rottondi, C., Barletta, L., Giusti, A., & Tornatore, M. (2018). Machine-learning method for quality of transmission prediction of unestablished lightpaths. *Journal of Optical Communications and Networking*, 10(2), A286–A297.
19. Khan, F. N., Fan, Q., Lu, C., & Lau, A. P. T. (2019). An optical communication's perspective on machine learning and its applications. *Journal of Lightwave Technology*, 37(2), 493–516.
20. Velasco, L., Castro, A., Ruiz, M., & Junyent, G. (2014). Solving routing and spectrum allocation related optimization problems: From off-line to in-operation flexgrid network planning. *Journal of Lightwave Technology*, 32(16), 2900–2911.