Machine Learning-Based Traffic Prediction for Frequency-Referenced Optical Multiplexing

Kalpit R. Bahri
Department of Computer Science, University of Houston, Houston, TX, USA.
kalpitwork@uh.edu

Thijmen Giselmeyer
Department of Computer Science, University of North Texas, Denton, TX, USA.
hellothijmen@unt.edu

Abstract

Frequency-referenced optical multiplexing introduces a distinctive control problem for high-capacity networks because it couples spectral assignment, latency guarantees, and synchronization discipline within a shared physical infrastructure. Machine learning-based traffic prediction has therefore emerged as a critical system-level capability rather than a narrow forecasting exercise. This paper examines how predictive models can support frequency-referenced multiplexing by anticipating demand shifts, reducing provisioning latency, improving spectral utilization, and enabling proactive resource governance. It argues that the value of prediction depends less on raw accuracy than on integration with optical layer constraints, control plane timing, fault management, and policy safeguards. The paper reviews relevant machine learning architectures, including sequence models, attention mechanisms, graph-based learning, and reinforcement learning, while emphasizing that each introduces operational trade-offs in observability, training data quality, interpretability, and deployment cost. It then analyzes architectural choices across centralized, distributed, and hybrid control planes, showing that prediction must be calibrated to the timescales of wavelength assignment, amplifier dynamics, and service-level agreements. Robustness and fairness are treated as first-order design requirements because prediction errors can cascade into unnecessary overprovisioning, service degradation, or inequitable access across tenants and regions. Sustainability concerns are also examined, particularly the energy implications of overprovisioning and the embodied costs of additional optical infrastructure. Through cross-domain comparisons with cloud scheduling, energy forecasting, and transport logistics, the paper identifies governance mechanisms needed to make predictive multiplexing trustworthy and accountable. The conclusion proposes that future systems should adopt uncertainty-aware, auditable, and federated prediction pipelines that preserve operator control while enabling adaptive spectral efficiency.

Keywords

optical multiplexing, traffic prediction, machine learning, frequency references, network governance, resource allocation, robustness, sustainability.

1. Introduction

Optical transport networks are under sustained pressure to support cloud services, artificial intelligence workloads, real-time industrial control, and densely distributed sensing applications [1], [2]. Conventional multiplexing strategies have achieved remarkable capacity by expanding wavelength density, coherent modulation, and flexible grid allocation, yet they

often treat traffic demand as an external input to be accommodated after the fact [7], [8]. Frequency-referenced optical multiplexing changes this assumption by anchoring transmission channels to a stable frequency reference, which can improve spectral discipline and support guaranteed bandwidth with low latency [13]. The resulting system is not merely a physical-layer innovation. It is a socio-technical infrastructure in which prediction, scheduling, synchronization, and policy enforcement become intertwined.

Traffic prediction for such networks must operate across multiple horizons. Short-term prediction at sub-second or millisecond scales can inform rapid wavelength reassignment, burst absorption, and protection switching. Medium-term prediction at minutes to hours can guide capacity activation, spectrum defragmentation, and service migration. Long-term prediction at days to months can shape infrastructure investment, geographic expansion, and energy procurement. A model that performs well at one horizon may fail at another because the statistical properties of demand, the cost of error, and the available control actions differ substantially. This paper therefore treats machine learning-based traffic prediction as a control-plane function embedded in a layered architecture, not as an isolated data science task.

The central argument is that frequency-referenced multiplexing intensifies the need for prediction while also raising the cost of prediction failure. When channels are tightly referenced and latency guarantees are explicit, inaccurate forecasts can produce either underutilization or service violations. Overprovisioning may preserve performance but wastes spectrum and energy. Underprovisioning may violate guarantees and erode trust. The design challenge is to create predictive systems that are accurate enough, fast enough, interpretable enough, and governable enough for critical infrastructure. This requires attention to architecture, deployment, robustness, fairness, and sustainability, as well as to the institutional arrangements that determine who controls data and how algorithmic decisions are reviewed.

2. Frequency-Referenced Optical Multiplexing as a System Problem

Frequency-referenced multiplexing can be understood as a coordination mechanism that binds multiple optical carriers to a common frequency reference, thereby reducing uncertainty in channel spacing and enabling tighter spectral packing. The approach is attractive for guaranteed bandwidth and low latency because it provides a more deterministic relationship between allocated spectrum and delivered service [13]. However, determinism at the physical layer does not automatically produce determinism at the service layer. Traffic arrivals remain stochastic, failures remain unpredictable, and tenant behavior can shift rapidly in response to external events. Prediction is therefore needed to translate physical determinism into operational readiness.

The system problem is complicated by the fact that optical resources are not perfectly fungible. Wavelengths, subcarriers, amplifier bands, and routing paths have different noise, nonlinearity, and dispersion characteristics [8], [9]. A prediction that identifies rising demand in a particular region must be matched to feasible spectral slots, regeneration sites, and transponder capabilities. If the forecast is spatially coarse, it may activate capacity where it cannot be used. If it is temporally coarse, it may arrive after the relevant scheduling window has closed. Frequency-referenced multiplexing thus requires prediction models that are sensitive not only to volume but also to location, direction, service class, and latency tolerance.

Governance considerations arise early because the reference itself becomes a shared dependency. If many channels rely on the same frequency discipline, a fault or

misconfiguration can affect multiple tenants simultaneously. Prediction systems may be asked to identify precursors to such faults, but they may also create new risks if they recommend aggressive reconfiguration based on uncertain data. The control plane must therefore distinguish between predictive signals that are advisory and those that are authorized to trigger changes. This distinction is architectural, not merely procedural, because it determines whether predictions are handled in a centralized orchestrator, an edge controller, or a distributed agent.

3. Traffic Prediction Requirements in Optical Control Planes

The control plane must reconcile prediction with provisioning. In a frequency-referenced environment, provisioning involves assigning spectrum, configuring transceivers, verifying synchronization, and updating service guarantees. Each step has latency and failure modes. A forecast that is accurate but arrives too late is operationally useless. A forecast that is timely but systematically biased can cause repeated overcorrection. The control plane therefore needs prediction outputs that include uncertainty estimates, confidence intervals, and expected costs of error. These outputs allow orchestration logic to choose between conservative and aggressive actions depending on service criticality.

Data quality is a foundational requirement. Optical networks generate telemetry from transponders, amplifiers, routers, and management systems, but these data are often siloed, inconsistently timestamped, and sampled at different rates [10], [11]. Prediction models trained on incomplete or biased data can reproduce historical allocation patterns rather than actual demand. For example, if a region has historically received less capacity, its traffic may appear low simply because users have adapted to scarcity. A naive model may then continue underprovisioning that region. Fairness-aware data governance is therefore necessary to distinguish latent demand from observed demand, especially when public or regulated services are involved.

The control plane also requires model lifecycle management. Traffic patterns drift due to new applications, seasonal events, pricing changes, and infrastructure upgrades. A model that performs well during training may degrade silently. Operators need monitoring for distribution shift, fallback procedures when predictions are unreliable, and mechanisms for human override. In frequency-referenced multiplexing, where guarantees are explicit, fallback cannot mean simply ignoring prediction. It must mean reverting to a safe allocation policy that may be less efficient but more predictable. This safe mode should be tested and auditable, much like protection switching in conventional transport networks [16].

4. Machine Learning Architectures for Predictive Multiplexing

Sequence models remain a natural starting point for traffic prediction. Long short-term memory networks can capture temporal dependencies and have been widely used in network forecasting because they handle variable-length histories and nonlinear patterns [4]. They are relatively well understood and can be deployed in centralized or distributed settings. However, they may struggle with very long-range dependencies and with sudden regime changes caused by failures or large-scale events. Their performance depends heavily on feature engineering and on the availability of high-resolution telemetry.

Attention-based models, particularly transformer architectures, offer greater flexibility in weighting past observations and combining heterogeneous inputs [3]. They can incorporate topology, service class, time of day, and fault indicators into a unified representation. For frequency-referenced multiplexing, attention mechanisms can help the model focus on the

most relevant prior intervals, such as previous bursts on the same route or correlated demand across tenants. The trade-off is computational cost and data appetite. Transformers may be difficult to run at the edge or in real time without specialized hardware, and their internal representations are often less interpretable than simpler statistical models.

Graph-based learning is valuable because optical networks are spatially and topologically structured. Graph convolutional networks can propagate information across nodes and links, allowing predictions to account for correlated demand between data centers, aggregation points, and access networks [5]. This is particularly relevant when a frequency reference is shared across a metro or regional domain. A graph model can learn that congestion on one route often precedes demand shifts on another. Yet graph models require accurate topology and may be sensitive to missing links or stale inventory data. In multi-operator environments, topology sharing may conflict with commercial confidentiality, limiting the practical scope of graph-based prediction.

Reinforcement learning offers a different paradigm by learning control policies rather than point forecasts [17]. Instead of predicting demand and then optimizing allocation, an agent can learn to assign spectrum, adjust references, and route traffic based on rewards such as latency compliance and energy efficiency. This approach can handle sequential decision-making and delayed consequences. However, it raises significant safety and governance challenges. Exploration in a live optical network is unacceptable if it risks service violations. Simulation-based training can help, but simulators may not capture all physical and operational complexities. Reinforcement learning is therefore best viewed as a longer-term complement to supervised prediction rather than an immediate replacement for rule-based safeguards.

Hybrid architectures are likely to dominate early deployments. A supervised model can produce demand forecasts, while a rule-based or optimization layer enforces hard constraints. A lightweight anomaly detector can flag unusual patterns, and a reinforcement learning agent can suggest long-horizon policy adjustments under human review. This layered approach preserves predictability while allowing incremental automation. It also aligns with the broader trend in optical network automation, where machine learning is used for quality-of-transmission estimation, fault localization, and resource optimization [10], [11], [12], [14], [15].

5. Architectural and Deployment Trade-offs

Centralized prediction offers advantages in data availability and model complexity. A central orchestrator can aggregate telemetry from multiple domains, train large models, and coordinate spectrum allocation across regions. This is attractive for frequency-referenced multiplexing because the reference discipline may span many nodes. However, centralization introduces latency, single points of failure, and privacy concerns. If the central system is unavailable, local controllers may lack predictive guidance. It also concentrates control over infrastructure, which can be problematic when multiple operators or tenants share the same physical plant.

Distributed prediction moves models closer to the data source. Edge controllers can make rapid decisions about local bursts, protection switching, and short-term spectrum adjustments. This reduces latency and improves resilience. Yet distributed models may lack global context, leading to conflicting actions across domains. They also require coordination to avoid oscillation, where one controller’s response to predicted congestion creates congestion

elsewhere. Federated learning can mitigate some privacy issues by training shared models without centralizing raw data, but it introduces complexity in model aggregation, trust, and version control.

Hybrid architectures attempt to balance these trade-offs. A central system can provide medium- and long-term forecasts and policy constraints, while edge agents handle short-term adaptation. The central system can also maintain a digital twin of the network, enabling simulation of proposed actions before deployment. In frequency-referenced multiplexing, a digital twin can model reference stability, amplifier behavior, and latency guarantees, allowing operators to test prediction-driven policies under stress. The cost is integration effort and the need for accurate synchronization between the twin and the live network.

Deployment also depends on the operational maturity of the organization. Some operators have extensive telemetry and automated provisioning, while others rely on manual processes. A prediction system that assumes real-time control may be unusable in a manual environment. Conversely, a conservative system may fail to exploit the capabilities of a highly automated network. Deployment should therefore be staged, starting with decision support, then closed-loop assistance, and finally limited autonomy in well-bounded domains. Each stage should include evaluation criteria that go beyond accuracy, such as reduction in service violations, improvement in spectral utilization, and operator trust.

6. Robustness, Fairness, and Governance

Robustness is a first-order requirement because optical infrastructure is critical. Prediction models must handle missing data, corrupted telemetry, adversarial manipulation, and rare events. They must also degrade gracefully. A model that produces confident but wrong forecasts during a fiber cut or synchronization fault can worsen the incident. Uncertainty quantification, anomaly detection, and out-of-distribution monitoring are therefore essential. Operators should define thresholds beyond which predictions are ignored and safe policies are invoked. These thresholds should be transparent to human operators and consistent with service-level agreements.

Fairness is less obvious in optical networking than in domains such as credit or hiring, but it is still relevant. Spectrum and latency are scarce resources. If prediction models allocate capacity based on historical usage, they may systematically disadvantage new entrants, rural regions, public institutions, or tenants with bursty but legitimate workloads. Fairness can be framed in terms of equal access, proportional service, or priority for essential services. The choice is political and regulatory, not purely technical. Prediction systems should therefore expose allocation outcomes to audit and allow policy makers to adjust constraints. Without such safeguards, predictive multiplexing could entrench existing inequalities under the guise of efficiency.

Governance must address accountability. When an automated prediction leads to a service violation, responsibility should be traceable. This requires logging of model inputs, versions, decisions, and overrides. It also requires clear roles for operators, vendors, and regulators. Principles such as transparency, justice, beneficence, non-maleficence, and explicability provide a useful starting point for algorithmic governance [18]. The ethics of algorithms literature further emphasizes the importance of contestability and the distribution of power between those who design systems and those affected by them [19]. In the optical context, this means that tenants and public authorities may need mechanisms to challenge allocation decisions or request priority during emergencies.

7. Sustainability and Infrastructure Policy

Sustainability is increasingly central to network planning. Optical networks are more energy efficient than many alternatives, but they are not free of environmental impact. Overprovisioning spectrum and cooling capacity wastes energy, while additional transponders and amplifiers have embodied carbon costs. Prediction can reduce waste by matching capacity to demand, but only if it is accurate and if the control plane can act on forecasts. Poor prediction can just as easily increase energy use by triggering unnecessary activation or by causing repeated reconfiguration. Sustainability metrics should therefore be integrated into model objectives alongside latency and utilization.

The infrastructure policy dimension is equally important. Frequency-referenced multiplexing may require new synchronization infrastructure, updated transceivers, and coordinated spectrum management. These investments have long payback periods and may be unevenly distributed across regions. Policy makers must consider universal service obligations, competition, and resilience. If predictive systems are proprietary, smaller operators may lack access to the models or data needed to compete. Open interfaces and standardized telemetry could reduce barriers, but they must be balanced against security and commercial confidentiality. The concentration of data and model expertise in a few large vendors is a governance risk that deserves attention [20].

Sustainability also intersects with repair and maintenance. Predictive maintenance can extend the life of optical equipment, reducing waste. Traffic prediction can inform when to upgrade versus when to optimize existing assets. However, prediction-driven optimization may encourage running equipment closer to its limits, potentially shortening lifespan. A holistic sustainability assessment should account for operational energy, embodied carbon, electronic waste, and service quality. This requires life-cycle analysis and transparent reporting, not just short-term efficiency gains.

8. Case Illustrations and Cross-Domain Comparisons

Consider a metro network serving data centers, hospitals, schools, and industrial campuses. Frequency-referenced multiplexing can provide guaranteed latency for telemedicine and industrial control while flexibly serving best-effort cloud traffic. A prediction system might forecast a surge in hospital imaging transfers after a public health event, a school video traffic peak at the start of classes, and a factory control cycle that is highly periodic. A single global model may miss these distinct patterns. A hybrid model that combines global trends with local context is more appropriate. The control plane can reserve protected spectrum for critical services and use prediction to schedule flexible traffic around them.

A second illustration involves a national backbone connecting regions with different economic profiles. Historical data may show low traffic in a rural region because capacity has been limited. A naive model might predict continued low demand. A fairness-aware model could incorporate policy targets, population trends, and applications such as remote education to estimate latent demand. This would require data sharing between network operators and public agencies, raising privacy and governance questions. The value of prediction here is not just efficiency but equitable infrastructure development.

Cross-domain comparisons are instructive. Cloud schedulers use prediction to place workloads and scale resources, but they operate in controlled data centers with relatively homogeneous hardware. Optical networks face greater physical heterogeneity and longer provisioning times. Energy forecasting uses weather and demand data to balance grids, and it

has developed sophisticated uncertainty quantification because errors are costly. Optical prediction can borrow these methods, especially for long-horizon planning. Transport logistics uses graph-based optimization and real-time routing; it offers lessons for multi-domain coordination but also shows how local optimization can create systemic congestion. These comparisons suggest that optical traffic prediction should be treated as a socio-technical control problem, not merely a time-series benchmark.

9. Future Directions

Future research should prioritize uncertainty-aware prediction that produces actionable risk assessments rather than single-point forecasts. Models should be evaluated on decision quality, not just mean absolute error. This requires simulation environments that capture physical layer constraints, control plane latency, and failure modes. Digital twins will likely play a central role, but they must be validated against real network behavior and updated as infrastructure changes.

Federated and privacy-preserving learning is another promising direction. Operators could train shared models without exposing raw traffic data, enabling broader generalization while respecting commercial boundaries. However, federated systems require robust authentication, secure aggregation, and incentives for participation. Standardization bodies could define interfaces for telemetry and model updates, reducing integration costs and enabling multi-vendor interoperability.

Finally, governance must evolve alongside technical capability. Regulators may need to require impact assessments for automated resource allocation, especially when essential services are affected. Operators should publish high-level performance and fairness metrics without revealing sensitive commercial information. Researchers should study how human operators interact with predictive recommendations, because trust and overreliance are both risks. The goal is not full autonomy but accountable augmentation, where machine learning improves human and institutional decision-making without obscuring responsibility.

10. Conclusion

Machine learning-based traffic prediction for frequency-referenced optical multiplexing is a system-level challenge that sits at the intersection of photonics, control theory, data science, and public policy. Frequency-referenced multiplexing offers stronger guarantees and tighter spectral coordination, but it also makes prediction errors more consequential. Accurate forecasts can improve spectral utilization, reduce latency violations, and support sustainable capacity management. Inaccurate or poorly governed forecasts can waste resources, entrench inequities, and undermine trust in critical infrastructure.

This paper has argued that prediction should be embedded in a layered architecture with clear boundaries between advisory signals and authorized actions. Sequence models, attention mechanisms, graph learning, and reinforcement learning each offer distinct capabilities, but none eliminates the need for uncertainty quantification, fallback policies, and human oversight. Centralized, distributed, and hybrid deployments present different trade-offs in latency, resilience, privacy, and control. Robustness and fairness must be treated as design requirements, not afterthoughts. Sustainability and infrastructure policy further shape what counts as a good prediction, because efficiency gains are meaningful only when they reduce environmental and social costs.

The path forward requires interdisciplinary collaboration. Optical engineers must specify constraints and failure modes. Machine learning researchers must develop models that are robust and interpretable. Operators must stage deployment and maintain safe modes. Policy makers must ensure accountability, competition, and equitable access. If these conditions are met, predictive frequency-referenced multiplexing can become a foundation for more adaptive, resilient, and inclusive optical infrastructure.

References

1. Agrawal, G. P. (2020). Fiber-optic communication systems (5th ed.). Wiley.

2. Cisco. (2020). Cisco annual internet report (2018–2023) white paper. Cisco Systems.

3. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30, 5998–6008.

4. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735–1780.

5. Kipf, T. N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks. In International Conference on Learning Representations.

6. 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.

7. 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.

8. Poggiolini, P., Bosco, G., Carena, A., Curri, V., Jiang, Y., & Forghieri, F. (2014). The GN-model of fiber non-linear propagation and its applications. Journal of Lightwave Technology, 32(4), 694–721.

9. 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 networks. IEEE Communications Surveys & Tutorials, 21(2), 1383–1408.

10. Zibar, D., Piels, M., Jones, R., & Schaeffer, D. (2016). Machine learning techniques in optical communication. Journal of Lightwave Technology, 34(6), 1442–1452.

11. 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.

12. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521(7553), 436–444.

13. Zhou, Z., Wei, J., Luo, Y., Clark, K. A., Sillekens, E., Deakin, C., ... & Liu, Z. (2023). Communications with guaranteed bandwidth and low latency using frequency-referenced multiplexing. Nature Electronics, 6(9), 694-702.

14. 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.

15. Barletta, L., Giusti, A., Rottondi, C., & Tornatore, M. (2017). QoT estimation for unestablished lightpaths using machine learning. In Optical Fiber Communication Conference (pp. Th1J-1). Optical Society of America.

16. Alvizu, R., Maier, G., Pattavina, A., & Tornatore, M. (2017). Comprehensive survey on T-SDN: Software-defined networking for transport networks. IEEE Communications Surveys & Tutorials, 19(4), 2232–2283.

17. Sutton, R. S., & Barto, A. G. (2018). Reinforcement learning: An introduction (2nd ed.). MIT Press.

18. Floridi, L., & Cowls, J. (2019). A unified framework of five principles for AI in society. Harvard Data Science Review, 1(1).

19. Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. Big Data & Society, 3(2).

20. Crawford, K. (2021). Atlas of AI: Power, politics, and the planetary costs of artificial intelligence. Yale University Press.