Machine Learning-Assisted Performance Prediction for Hollow-Core Fiber Interconnects Under Dynamic Traffic Loads

Chakradev Acharya
Department of Computer Science, University of New Hampshire, Durham, NH, USA.
acharyachakradev@unh.edu

Berenice Casarin
Department of Computer Science, University of Central Florida, Orlando, FL, USA.
berenice719@ucf.edu

Adya Khurana
Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA.
adyak@buffalo.edu

Abstract

Hollow-core fiber interconnects are emerging as a transformative substrate for low-latency, high-capacity data center and high-performance computing fabrics, yet their performance under dynamic traffic loads remains difficult to predict because optical, thermal, mechanical, and scheduling dynamics interact across multiple timescales. This paper examines machine learning-assisted performance prediction for hollow-core fiber interconnects from a system-level perspective. It argues that accurate prediction cannot be reduced to link-level regression or isolated optical performance monitoring; rather, it requires an integrated architecture that couples physical-layer telemetry, switch and routing state, workload traces, energy telemetry, and governance metadata. The paper analyzes structural trade-offs among latency, capacity, energy efficiency, reliability, and fairness when predictive models are embedded in network control loops. It discusses data acquisition and feature engineering challenges, including sparse failure events, distribution drift, adversarial manipulation, and the need for explainable decisions in critical infrastructure. It also considers deployment constraints, sustainability implications, supply-chain dependencies, and policy questions regarding transparency, accountability, and equitable access. Through cross-domain comparisons with cloud networking, transport software-defined networking, cyber-physical systems, and high-performance computing, the paper shows that hollow-core fiber interconnects are not merely faster pipes but socio-technical infrastructures whose benefits depend on institutional coordination and robust model governance. The conclusion outlines a research agenda for validated digital twins, federated and energy-aware learning, uncertainty quantification, and standards that make machine learning-assisted prediction auditable and operationally trustworthy.

Keywords

hollow-core fiber; machine learning; optical interconnects; dynamic traffic; performance prediction; data center networks; sustainability; robustness; governance.

1. Introduction

The growth of cloud computing, large-scale artificial intelligence training, and latency-sensitive interactive services has transformed the data center from a collection of servers into a tightly coupled computational organism whose performance depends on the interconnect fabric as much as on processors and memory [1]. As warehouse-scale machines scale outward, electrical interconnects encounter fundamental limits in reach, bandwidth density, and energy per bit, making optical technologies central to future system design [6]. At the same time, computing architectures are entering an era of domain-specific acceleration in which gains no longer come automatically from transistor scaling, further increasing the pressure on communication pathways [14]. In this environment, hollow-core fiber has attracted attention because it guides light in air rather than in solid glass, potentially reducing propagation latency and nonlinear impairment while offering a path toward higher capacity and lower latency than conventional silica fibers [4]. The operational value of such fiber, however, depends on whether network operators can predict how it behaves when traffic loads fluctuate rapidly across tenants, services, and artificial intelligence workloads.

Dynamic traffic loads are not merely a nuisance for performance modeling; they are a defining condition of modern infrastructure. East-west traffic between accelerators, storage, and inference servers can change within milliseconds, while scheduled batch jobs, live migrations, and failure recovery can create abrupt shifts in demand that propagate through the fabric. Machine learning has therefore become an attractive tool for network management because it can learn nonlinear relationships from telemetry, anticipate congestion, classify impairments, and support resource allocation in software-defined and transport networks [8], [10]. Yet the application of machine learning to hollow-core fiber interconnects raises distinctive questions. The fiber is not simply a faster replacement for conventional fiber; its propagation characteristics, connectorization, bending sensitivity, and manufacturing variability create new dependencies between physical deployment and logical performance. Prediction models that ignore those dependencies may appear accurate in controlled experiments but fail under real traffic, maintenance, and multi-tenant conditions.

This paper develops a system-level analysis of machine learning-assisted performance prediction for hollow-core fiber interconnects under dynamic traffic loads. It treats prediction as a socio-technical capability that spans sensing, data pipelines, model training, control integration, and governance. The discussion is organized around architecture, trade-offs, robustness, fairness, sustainability, deployment, and policy. The central argument is that the value of prediction is not limited to forecasting a single metric such as latency or bit error rate. Instead, prediction must help operators decide when to route, when to reserve, when to degrade gracefully, and when to intervene in ways that are defensible to users, regulators, and the public. This requires models that are not only accurate but also interpretable, energy-aware, and institutionally accountable.

2. Hollow-Core Fiber Interconnect Architecture and Dynamic Traffic Conditions

Hollow-core fiber changes the physical assumptions of optical networking by confining light in a central air or vacuum region surrounded by a microstructured cladding. In photonic bandgap and anti-resonant designs, the guiding mechanism reduces the overlap between the optical mode and the glass, which can lower nonlinear effects and propagation delay relative to conventional solid-core fiber [4]. For data center interconnects, these properties are attractive because they address both capacity and latency simultaneously. Capacity improvements in conventional fiber have historically relied on space-division multiplexing and advanced modulation, but those approaches introduce complexity in amplifiers, switches,

and transceivers [3]. Hollow-core fiber offers a complementary path by improving the medium itself, especially for short-reach and intra-data center distances where latency, energy per bit, and synchronization are critical.

The architecture of a hollow-core interconnect is not confined to the fiber strand. It includes transceivers, wavelength-selective switches, arrayed waveguides, clock synchronization subsystems, connectors, splicing, and thermal management. Recent work on wavelength-switched and clock-synchronized intra-data center interconnects illustrates that hollow-core fiber can be integrated with switching and synchronization mechanisms to reduce latency while maintaining reconfigurability [18]. This matters because dynamic traffic loads are often handled through wavelength reallocation or path switching rather than through changes in the physical medium. If the fiber introduces lower propagation delay but the switching and synchronization layers add jitter or reconfiguration delay, then system-level performance may not improve as much as physical-layer measurements suggest. Prediction models must therefore operate across layers, learning how optical path changes, queue build-up, and clock recovery interact under bursty demand.

Dynamic traffic in such an environment has several distinctive features. It is heterogeneous, because storage replication, distributed training, remote procedure calls, and video delivery have different latency, bandwidth, and loss tolerances. It is bursty, because artificial intelligence training can synchronize thousands of accelerators in collective operations that generate intense, short-lived traffic patterns. It is interdependent, because congestion in one rack can trigger retries, load balancing, and migrations that amplify rather than dampen demand. It is also partially predictable, because schedulers often know job placement and service-level objectives before traffic arrives. Machine learning-assisted prediction can exploit this partial predictability, but only if the feature space includes workload metadata, switch queue state, transceiver telemetry, and optical performance monitoring data. A purely optical model cannot anticipate a traffic burst caused by a synchronization barrier, and a purely workload-driven model cannot anticipate a fiber bend or connector degradation.

3. Machine Learning for Performance Prediction in Optical Interconnects

Machine learning has been applied to optical communication for impairment classification, quality-of-transmission estimation, nonlinear equalization, failure detection, and resource allocation [8], [9]. In conventional optical networks, models often learn mappings from signal features to bit error rate, optical signal-to-noise ratio, or alarm states. For hollow-core fiber interconnects under dynamic loads, the prediction target is broader. Operators may need short-horizon forecasts of latency, throughput, queue depth, and energy consumption, as well as longer-horizon forecasts of degradation and maintenance risk. A model that predicts average latency over minutes may be useless for a control loop that must reconfigure wavelengths within milliseconds. Conversely, a model optimized only for microsecond decisions may miss thermal and mechanical trends that determine whether a link will remain reliable over days.

Supervised learning is useful when historical telemetry includes labeled outcomes such as congestion events, link failures, or service-level violations. Unsupervised learning can identify novel operating regimes or anomalous traffic patterns that were not present in training data. Reinforcement learning can optimize routing and wavelength assignment by interacting with a simulated or live environment, but its deployment in critical infrastructure requires safeguards because exploration can cause real service disruption [10]. Deep learning models can capture complex temporal dependencies in telemetry, yet their opacity can

complicate root-cause analysis when predictions fail. In optical networks, hybrid approaches that combine physics-informed features with data-driven models often provide a more credible balance between accuracy and interpretability than purely black-box methods [9].

Software-defined networking and transport software-defined networking provide the control plane through which predictions become actions [12], [13]. Without a programmable control plane, prediction remains advisory. With such a control plane, prediction can inform traffic engineering, wavelength switching, protection switching, and energy-aware routing. This integration creates a feedback loop: the network generates data, the model predicts performance, the controller acts, and the action changes future data. That loop can improve efficiency, but it can also produce instability if the model overreacts to noise or if multiple controllers optimize conflicting objectives. For hollow-core fiber interconnects, where low latency is a primary value proposition, even small errors in prediction can cause unnecessary path changes that increase jitter and undermine the very benefit the fiber was deployed to provide.

4. Data, Telemetry, and Feature Engineering for Predictive Models

The quality of prediction depends on the observability of the system. Hollow-core fiber interconnects can expose physical-layer telemetry from transceivers, optical channel monitors, temperature sensors, and power meters; logical telemetry from switches, routers, and network interface cards; and workload telemetry from schedulers, orchestration platforms, and application-level tracing. Effective feature engineering must align these streams in time and space, because a latency spike may originate from a traffic burst, a thermal transient, a switch contention event, or a synchronization error. The challenge is not merely data volume but semantic coherence. A feature that is meaningful in one deployment, such as a particular wavelength assignment, may be unavailable or differently named in another vendor environment.

Label scarcity is a persistent problem. Failures in well-maintained optical systems are rare, so supervised models for failure prediction must cope with severe class imbalance. Simulated faults can augment data, but simulations may not capture the messy interactions of real connectors, dust, vibration, and maintenance errors. Transfer learning and domain adaptation can help models trained in one data center adapt to another, but they require careful validation because hollow-core fiber deployment practices are still maturing. Distribution drift is also significant: traffic patterns change with new applications, fiber ages, and software updates alter scheduling behavior. A model that was accurate during initial deployment may silently degrade unless monitoring detects drift and triggers retraining [10]. Governance metadata, including model version, training data provenance, and validation scope, is therefore part of the technical data pipeline, not an administrative afterthought.

Privacy and confidentiality complicate data sharing across tenants and operators. Traffic traces can reveal business relationships, user behavior, and security-sensitive timing patterns. Federated learning and secure aggregation offer partial remedies, but they introduce trade-offs in model accuracy and operational complexity. In multi-tenant data centers, fairness requires that prediction and control do not systematically favor large tenants whose workloads are better instrumented or more predictable. A small tenant with bursty but legitimate traffic may be penalized if the model treats unpredictability as a risk to be suppressed. Cyber-physical systems research has long emphasized that sensing, computation, and actuation must be co-designed with safety and real-time constraints [20]. The same principle applies here:

prediction models must be evaluated not only by statistical error but by their effect on service equity, stability, and human operators' ability to understand and intervene.

5. Trade-offs in Latency, Capacity, Energy, and Reliability

Hollow-core fiber promises lower latency and lower nonlinearity, but those benefits interact with capacity, energy, and reliability in ways that are not always additive. Lower propagation delay can reduce the round-trip time available for congestion control, which may improve application performance or may require retuning of transport protocols. Lower nonlinearity can support higher launch power or advanced modulation, potentially increasing capacity, but the supporting transceivers and switches may consume more energy. Reliability is also multidimensional: a fiber with excellent optical properties may still be vulnerable to connector contamination, bending loss, or manufacturing defects that are not fully represented in laboratory measurements. Prediction models must therefore incorporate uncertainty and not present a single deterministic forecast as if it were certain.

Energy efficiency is a central system constraint. Data centers already consume substantial electricity, and although global estimates require careful recalibration, the direction of growth in artificial intelligence workloads makes energy-aware design imperative [7]. Machine learning models themselves consume energy during training and inference, and a prediction system that saves a small amount of network energy while consuming a large amount of compute energy is not sustainable [15]. The trade-off is particularly sharp in hollow-core fiber interconnects because the fiber is often justified by latency and capacity, not directly by energy. If predictive control reduces overprovisioning by allowing tighter resource allocation, it may improve energy efficiency. If it causes frequent reconfiguration or conservative reservations, it may increase both energy use and operational complexity.

Reliability trade-offs appear in the choice between proactive and reactive control. Proactive prediction can reroute traffic before congestion occurs, but false positives can waste capacity and increase path length. Reactive control waits for evidence, which is simpler but may allow service-level violations. In hollow-core fiber systems, low latency may make reactive control more feasible for some events, but physical-layer degradation may evolve faster than logical telemetry can reveal. The required balance depends on the cost of failure, the predictability of the workload, and the governance framework that defines acceptable risk. A financial trading fabric may tolerate aggressive proactive control to avoid microseconds of delay, while a public-sector cloud may prioritize stability and equitable service. Prediction models must therefore be parameterized by policy, not only by performance data.

6. Robustness, Fairness, and Governance in Machine Learning-Assisted Networks

Robustness in machine learning-assisted optical networks includes resilience to noisy telemetry, missing data, adversarial manipulation, and model drift. An attacker who can influence traffic patterns or telemetry could induce the controller to make harmful routing decisions. Even without malice, a misconfigured scheduler or a faulty sensor can create input distributions that lie outside the model's training domain. Robustness engineering requires redundancy, uncertainty quantification, fallback rules, and human oversight. In critical infrastructure, a model that cannot explain why it recommends a reconfiguration may be unacceptable, because operators need to distinguish a genuine performance threat from a spurious correlation [17]. Explainability is not a cosmetic feature; it is a condition for trust and accountability.

Fairness in this context has multiple meanings. At the tenant level, predictive control should not systematically degrade service for users whose traffic is harder to forecast. At the regional level, deployment of advanced hollow-core interconnects may concentrate benefits in wealthy data center hubs unless policy intervenes. At the labor level, automation may change the skills required of network operators, shifting authority from local technicians to centralized model maintainers. Governance must address who has access to telemetry, who can retrain models, who is accountable when a prediction causes an outage, and how affected parties can contest automated decisions. Principles such as transparency, justice, and beneficence are necessary but insufficient without institutional mechanisms for enforcement [19]. The history of large-scale machine learning also warns that scaling up models without considering environmental and social costs can externalize harm onto marginalized communities [16].

Governance must therefore be embedded in the architecture. Model registries, audit logs, versioned control policies, and rollback mechanisms are technical instruments of accountability. They allow regulators and internal auditors to reconstruct why a link was reconfigured or why a tenant’s traffic was deprioritized. In transport software-defined networking, centralized controllers already raise questions of trust and jurisdiction because a single logical entity may govern many physical domains [13]. Adding machine learning increases the opacity and speed of decisions, making governance more difficult. A hollow-core fiber interconnect that spans multiple data centers, operators, or countries may be subject to conflicting rules about data localization, energy reporting, and critical infrastructure protection. Prediction systems must be designed to operate under those constraints, not as if they exist in a policy vacuum.

7. Deployment, Sustainability, and Policy Implications

Deployment of hollow-core fiber interconnects will be incremental rather than instantaneous. Existing data centers contain large investments in conventional fiber, connectors, transceivers, and operational knowledge. Hollow-core fiber may first appear in latency-critical paths, such as links between accelerator clusters or between trading systems, before expanding to broader fabrics. This incremental path creates heterogeneity: prediction models must handle mixed fiber types, mixed vendor equipment, and mixed traffic classes. The transition also affects supply chains, because hollow-core fiber manufacturing, splicing, and connectorization require specialized skills and equipment. Policy makers concerned with digital infrastructure may need to consider whether incentives for advanced fiber deployment are aligned with public interest, including rural connectivity, research networks, and energy grid constraints.

Sustainability is both a local and a global issue. More efficient interconnects can reduce energy per bit, but they can also enable more intensive computation and data movement, a rebound effect that may offset savings. Data center energy estimates are contested, and the rapid growth of artificial intelligence training has made forecasting difficult [7]. Machine learning-assisted prediction can contribute to sustainability if it reduces idle capacity, improves cooling coordination, or extends equipment lifetime. It can also worsen sustainability if it requires continuous high-resolution telemetry, frequent model retraining, and redundant control planes. The environmental cost of training large models is well documented [15], and even smaller network models have a cumulative footprint when deployed across thousands of switches. Sustainability assessment must therefore include the full lifecycle of fiber, transceivers, servers, and models.

Policy implications extend to standards, competition, and security. Standards bodies must define interfaces for telemetry, model interchange, and safety constraints so that operators are

not locked into a single vendor's prediction stack. Competition authorities may need to examine whether control over predictive analytics becomes a barrier to entry in cloud and interconnection markets. Security agencies must consider whether hollow-core fiber interconnects, despite physical advantages, introduce new attack surfaces through software-defined control and telemetry channels. Because machine learning models can encode biases and vulnerabilities, procurement rules may require documentation of training data, validation methods, and failure modes. The governance of cyber-physical systems offers a precedent: safety, security, and real-time performance must be addressed together rather than sequentially [20].

8. Case Illustrations and Cross-Domain Comparisons

A useful illustration is a cloud data center that operates a mixture of conventional and hollow-core links between racks of accelerators. During a distributed training job, collective communication creates synchronized bursts that can saturate a path for milliseconds. A prediction model that knows the job schedule, the collective algorithm, and the current queue depth can anticipate the burst and reserve wavelength capacity in advance. Without workload metadata, the same model may see the burst as an anomaly and react too late. The case shows that hollow-core fiber's low latency is most valuable when paired with predictive scheduling that understands application semantics. It also shows that fairness matters: if only premium tenants receive workload-aware prediction, then the interconnect becomes a mechanism for tiered service quality rather than a neutral substrate.

In transport networks, software-defined control has already demonstrated the value of centralized optimization and programmable traffic engineering [12], [13]. However, transport networks typically operate at longer timescales and with more stable traffic than intra-data center fabrics. A model that works for inter-city wavelength routing may be too slow for rack-scale interconnect decisions. Conversely, a model trained only on data center microbursts may lack the physical-layer generalization needed for long-haul hollow-core fiber, where temperature, pressure, and geospatial conditions vary. Cross-domain comparison suggests that prediction systems should be modular: fast local models for microsecond control, regional models for path optimization, and slow analytical models for maintenance and capacity planning.

High-performance computing offers another comparison. Supercomputing interconnects have long used topology-aware routing and application-aware scheduling to manage congestion. Cyber-physical systems emphasize safety envelopes and deterministic response [20]. Hollow-core fiber interconnects can borrow from both traditions. From high-performance computing, they can adopt topology-aware prediction and collective communication modeling. From cyber-physical systems, they can adopt safety cases, runtime verification, and fallback controllers. The synthesis is a layered architecture in which machine learning improves efficiency within bounds set by formal or rule-based safety mechanisms. This approach acknowledges that prediction is powerful but not infallible, and that the consequences of error in critical infrastructure can be severe.

9. Future Directions and Research Challenges

Several research directions follow from this analysis. First, validated digital twins of hollow-core fiber interconnects are needed to train and test prediction models without disrupting live service. A digital twin must represent optical propagation, thermal dynamics, switching behavior, and workload demand at multiple timescales. It must also be calibrated against real

telemetry and updated as the physical system ages. Second, federated and privacy-preserving learning can allow multiple operators to benefit from shared models without exchanging sensitive traffic traces. However, federation introduces challenges in heterogeneity, incentive design, and adversarial participants. Third, uncertainty quantification must become a first-class output of prediction. Operators need to know not only the expected latency but also the confidence interval and the conditions under which the model should defer to a fallback rule.

Fourth, energy-aware machine learning for network control must account for the energy cost of sensing, communication, and computation. A model that predicts congestion accurately but consumes more energy than the congestion would have wasted is not useful. Fifth, explainable artificial intelligence should be co-designed with network operations tools so that explanations are actionable rather than abstract. An operator needs to know which switch, wavelength, or workload feature drove a recommendation, and whether the recommendation conflicts with service-level agreements. Sixth, standards and regulatory frameworks must define minimum requirements for logging, auditing, and redress. The governance of artificial intelligence in critical infrastructure cannot rely solely on voluntary principles [17], [19]. It requires enforceable rules for transparency, security, and liability.

Finally, the research community should avoid treating hollow-core fiber interconnects as a purely technical optimization problem. The benefits of low latency and high capacity will be mediated by business models, energy markets, labor practices, and geopolitical supply chains. A holistic research agenda should include interdisciplinary work with economists, legal scholars, environmental scientists, and social scientists. The goal is not to slow innovation but to ensure that performance prediction serves resilient, equitable, and sustainable infrastructure. Machine learning can help operators navigate the complexity of dynamic traffic, but only if its limitations are understood and its authority is bounded by democratic and institutional accountability.

10. Conclusion

Machine learning-assisted performance prediction for hollow-core fiber interconnects under dynamic traffic loads is a systems problem that spans physics, computation, control, and governance. Hollow-core fiber offers compelling advantages in latency and nonlinearity, but those advantages are realized only when integrated with switching, synchronization, telemetry, and workload-aware control. Prediction models can improve resource allocation, energy efficiency, and reliability, yet they also introduce risks of instability, bias, opacity, and unsustainable computational overhead. This paper has argued for a layered architecture in which machine learning operates within safety bounds, produces uncertainty estimates, and remains auditable by human operators and external regulators. It has also emphasized that fairness and sustainability are not secondary concerns but design constraints that shape which predictions are acceptable and which actions are legitimate. Future work should focus on validated digital twins, federated and energy-aware learning, explainable control, and standards that make automated decisions contestable. By treating hollow-core fiber interconnects as socio-technical infrastructures rather than mere high-speed links, researchers and operators can pursue performance gains without sacrificing resilience, equity, or public trust.

References

1. Kachris, C., & Tomkos, I. (2012). A survey on optical interconnects for data centers. IEEE Communications Surveys & Tutorials, 14(4), 1021–1036. https://doi.org/10.1109/SURV.2011.122111.00069

2. Bogaerts, W., De Heyn, P., Van Vaerenbergh, T., De Vos, K., Kumar Selvaraja, S., Claes, T., Dumon, P., Bienstman, P., Van Thourhout, D., & Baets, R. (2012). Silicon microring resonators. Laser & Photonics Reviews, 6(1), 47–73. https://doi.org/10.1002/lpor.201100017

3. Richardson, D. J., Fini, J. M., & Nelson, L. E. (2013). Space-division multiplexing in optical fibres. Nature Photonics, 7(5), 354–362. https://doi.org/10.1038/nphoton.2013.94

4. Poletti, F., Wheeler, N. V., Petrovich, M. N., Baddela, N., Numkam Fokoua, E., Hayes, J. R., Gray, D. R., Li, Z., & Richardson, D. J. (2013). Towards high-capacity fibre-optic communications at the speed of light in vacuum. Nature Photonics, 7(4), 279–284. https://doi.org/10.1038/nphoton.2013.45

5. 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. https://doi.org/10.1364/OE.26.024190

6. Barroso, L. A., Clidaras, J., & Hölzle, U. (2013). The datacenter as a computer: An introduction to the design of warehouse-scale machines (2nd ed.). Morgan & Claypool. https://doi.org/10.2200/S00516ED2V01Y201306CAC024

7. Masanet, E., Shehabi, A., Lei, N., Smith, S., & Koomey, J. (2020). Recalibrating global data center energy-use estimates. Science, 367(6481), 984–986. https://doi.org/10.1126/science.aba3758

8. 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(14), 3516–3528. https://doi.org/10.1109/JLT.2019.2900035

9. Zibar, D., Piels, M., Jones, R., & Schäffer, C. G. (2016). Machine learning techniques in optical communication. Journal of Lightwave Technology, 34(6), 1442–1452. https://doi.org/10.1109/JLT.2015.2501625

10. 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. https://doi.org/10.1186/s13174-018-0087-2

11. 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

12. Kreutz, D., Ramos, F. M. V., Veríssimo, P. E., Rothenberg, C. E., Azodolmolky, S., & Uhlig, S. (2015). Software-defined networking: A comprehensive survey. Proceedings of the IEEE, 103(1), 14–76. https://doi.org/10.1109/JPROC.2014.2371999

13. Alvizu, R., Maier, G., Kaczmarek, S., & Duran, A. (2017). Comprehensive survey on T-SDN: Software-defined networking for transport networks. IEEE Communications Surveys & Tutorials, 19(4), 2232–2283. https://doi.org/10.1109/COMST.2017.2715220

14. Shalf, J. (2020). The future of computing beyond Moore's Law. Philosophical Transactions of the Royal Society A, 378(2166), 20190061. https://doi.org/10.1098/rsta.2019.0061

15. Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, 3645–3650. https://doi.org/10.18653/v1/P19-1355

16. Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, 610–623. https://doi.org/10.1145/3442188.3445922

17. Mittelstadt, B. (2019). Principles alone cannot guarantee ethical AI. Nature Machine Intelligence, 1(11), 501–507. https://doi.org/10.1038/s42256-019-0114-4

18. 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.

19. 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

20. Rajkumar, R., Lee, I., Sha, L., & Stankovic, J. (2010). Cyber-physical systems: The next computing revolution. Proceedings of the 47th Design Automation Conference, 731–736. https://doi.org/10.1145/1837274.1837461