

AI-Driven Traffic Scheduling for Low-Latency Data Center Optical Networks

Amalia Kuhl

Department of Computer Science, George Mason University, Fairfax, VA, USA.
amalia1980@gmu.edu

Eryk Kapustka

Department of Computer Science, University of New Hampshire, Durham, NH, USA.
erykwork@unh.edu

Gianpaolo Borzomi

Department of Computer Science and Engineering, University of Nevada, Reno, Reno, NV, USA.
gianpaolomail@unr.edu

Clelia Acerbi

Department of Computer Science, University of Central Florida, Orlando, FL, USA.
acerbi858@ucf.edu

Abstract

Data center optical networks are increasingly positioned as the low-latency substrate for distributed computation, storage disaggregation, and artificial intelligence workloads. Yet optical capacity alone does not guarantee predictable latency, because traffic scheduling must reconcile bursty demand, circuit setup delays, wavelength contention, and heterogeneous quality-of-service requirements. This paper examines artificial intelligence driven traffic scheduling for low-latency data center optical networks from a system-level perspective. It argues that scheduling intelligence should not be treated as a centralized optimizer inserted above existing hardware, but as a cross-layer control function embedded in the optical, electrical, and application planes. The discussion covers architectural foundations, including optical circuit switching, wavelength switching, hybrid electrical-optical fabrics, and hollow-core fiber interconnects. It then analyzes how reinforcement learning, predictive demand modeling, and online optimization can reduce tail latency while respecting fairness, robustness, and governance constraints. The paper also considers deployment realities, sustainability pressures, and the economic trade-offs between overprovisioning and fine-grained scheduling. Case illustrations compare intra-data center interconnects with wide-area traffic engineering and elastic optical routing, showing that transferable design principles include state abstraction, safety envelopes, and human oversight. The conclusion emphasizes that low-latency optical scheduling is a socio-technical problem: its success depends not only on algorithmic performance but also on observability, accountability, energy accounting, and institutional mechanisms for resolving conflicting operational priorities.

Keywords

AI-driven traffic scheduling; data center optical networks; low-latency interconnects; optical circuit switching; reinforcement learning; governance; sustainability; robustness; fairness.

1. Introduction

Modern data centers have evolved from relatively static warehouses of servers into dynamic computational infrastructures whose performance depends on rapid, predictable communication among processors, memory, accelerators, and storage. The classical tree-based electrical network provided a scalable model for early cloud workloads, but its energy cost, cabling complexity, and bandwidth limitations became increasingly visible as east-west traffic grew [1]. Optical interconnects offered a promising alternative by exploiting wavelength division multiplexing and low-loss transmission, yet early optical architectures often struggled to match the fine-grained flexibility of packet switching [2]. The resulting design space is not a simple replacement of electrical links with optical links. It is a scheduling problem in which the physical properties of light, the granularity of switching, and the temporal structure of application demand interact.

Artificial intelligence has entered this space because traffic scheduling in optical data centers is difficult to solve with static rules. Demand fluctuates across microseconds, milliseconds, and seconds; workloads may require barrier synchronization, collective communication, or bursty data transfers; and the cost of reconfiguration can be comparable to the latency saved. Machine learning techniques have therefore been explored for optical performance monitoring, impairment prediction, and resource assignment [3]. However, much of the literature treats intelligence as a separate optimization layer rather than as an integral property of the network architecture. This paper argues that AI-driven traffic scheduling for low-latency optical networks must be understood as a cross-layer control problem that spans application intent, transport protocols, electrical packet buffering, optical circuit setup, and physical-layer constraints.

The urgency of this problem is practical. Low-latency interconnects are no longer only a concern for high-performance computing. Cloud providers, financial services, machine learning training clusters, and real-time analytics all depend on tail latency that is stable enough for service-level agreements. Optical switching can reduce hop count and energy per bit, but it can also introduce configuration delays, wavelength contention, and synchronization overhead. Recent work on wavelength-switched and clock-synchronized intra-data center interconnects illustrates how physical-layer innovation can reduce latency, but also how tightly scheduling must be coupled with synchronization and switching hardware [5]. The central question is therefore not whether AI can improve optical scheduling, but under what architectural conditions, governance structures, and operational assumptions its improvements are robust, fair, and sustainable.

2. Background and Related Work

Data center network research has long pursued scalability, bisection bandwidth, and cost efficiency. The fat-tree architecture demonstrated that commodity switches could support large-scale clusters through a regular topology, but it also established the electrical packet switch as the default control point [1]. Hybrid electrical-optical designs such as Helios showed that optical circuit switching could be layered onto electrical aggregation to handle large flows, yet the control plane had to decide when a circuit was worth establishing [2]. Surveys of optical interconnects for data centers have since documented a wide range of switching technologies, including micro-electro-mechanical systems, semiconductor optical amplifiers, and arrayed waveguide gratings, each with different reconfiguration times and scalability limits [3]. The OSA architecture further demonstrated that optical switching could provide flexibility across racks, but also that scheduling complexity grows with the number of ports and the diversity of traffic classes [4].

AI for networking has developed along parallel lines. Reinforcement learning has been applied to resource management in data centers, where an agent learns to allocate compute or network resources from experience rather than relying on hand-tuned heuristics [8]. AuTO extended this idea to datacenter-scale traffic optimization, using deep reinforcement learning to adjust flow scheduling and rate control in response to changing conditions [9]. Related work on congestion control has framed the problem as sequential decision making under uncertainty, where the agent must balance throughput, delay, and fairness [10]. Experience-driven networking has similarly explored how agents can learn routing and resource allocation policies from operational data [11]. These approaches draw on general advances in deep reinforcement learning, including value-based methods that can approximate complex control policies from high-dimensional state observations [12].

In optical communications, machine learning has been used for nonlinearity compensation, modulation format recognition, failure prediction, and quality-of-transmission estimation [6]. The appeal is that optical systems produce large amounts of telemetry, and many physical-layer relationships are difficult to model analytically under real-world conditions [7]. For data center optical scheduling, however, the most relevant lessons are not merely about prediction accuracy. They concern state representation, action space design, reward shaping, and the safety of exploration. An optical scheduler that learns by reconfiguring live circuits may improve average latency while causing rare but severe disruptions. The literature on elastic optical networks has begun to address this through deep reinforcement learning for routing, modulation, and spectrum assignment, but those methods often assume a centralized controller with relatively complete state information [20]. In intra-data center networks, state information is distributed, telemetry is sampled, and the cost of miscoordination is high.

The required low-latency interconnect study at reference position five is important because it grounds the discussion in physical-layer feasibility. It shows that wavelength-switched and clock-synchronized intra-data center interconnects can be enabled by hollow-core fiber, reducing propagation delay and supporting tighter synchronization [5]. This does not solve scheduling by itself. Instead, it raises the bar for scheduling intelligence: if the physical path is faster, then control-plane delays, queueing, and reconfiguration overhead become proportionally more significant. The remainder of this paper therefore treats AI scheduling as a complement to optical hardware innovation, not as a substitute for it.

3. Architectural Foundations of AI-Driven Optical Scheduling

An AI-driven optical scheduler must operate within a layered architecture. At the physical layer, wavelengths, ports, amplifiers, and fibers define the available resources. At the switching layer, optical circuit switches, wavelength-selective switches, and hybrid electrical-optical fabrics determine the granularity and reconfiguration time of connections [2]. At the transport layer, protocols such as remote direct memory access, congestion control, and collective communication libraries shape the traffic that reaches the optical fabric [10]. At the application layer, job schedulers, storage systems, and machine learning frameworks express priorities that may be implicit rather than explicit. A central challenge is that these layers expose different notions of time. Physical-layer impairments evolve over microseconds, optical switches reconfigure over nanoseconds to milliseconds depending on technology, transport protocols react over round-trip times, and applications may tolerate delay only within synchronization windows [18].

The placement of intelligence is therefore an architectural decision. A centralized scheduler can maintain a global view of topology, demand, and policy, but it may become a scalability

bottleneck and a single point of failure. A distributed scheduler can react locally and survive partial failures, but it may produce oscillations or unfair outcomes when multiple agents compete for the same wavelengths. A hierarchical design may offer a middle path: local agents handle fast, reversible decisions such as wavelength selection or queue management, while a slower global agent handles topology adaptation, capacity reservation, and policy enforcement. This division mirrors established ideas in traffic engineering, but optical networks add a physical constraint that electrical networks do not have: once a circuit is established, it may consume a dedicated wavelength even when idle, and rapid reconfiguration can introduce transients that affect neighboring channels.

AI can contribute at each level, but the type of learning should match the timescale. For fast decisions, lightweight predictive models or bandit algorithms may be more appropriate than deep reinforcement learning, because exploration must be safe and sample efficiency is critical. For slower decisions, reinforcement learning can learn to balance utilization, latency, and energy over longer horizons [9]. Predictive models can forecast demand from job traces, storage access patterns, or collective communication schedules, allowing the scheduler to pre-provision circuits before contention occurs. Yet prediction is never perfect, and overconfidence can lead to brittle policies. A robust architecture should therefore combine learned components with explicit safety envelopes, fallback rules, and observable constraints that prevent the scheduler from violating latency or fairness bounds.

4. Low-Latency Mechanisms and Cross-Layer Coordination

Low latency in optical data centers is not achieved by any single mechanism. It emerges from the coordination of propagation delay, switching delay, queueing delay, and control-plane delay. Hollow-core fiber can reduce propagation delay relative to conventional silica fiber, and clock synchronization can reduce the uncertainty that forces conservative buffering [5]. Optical circuit switching can eliminate store-and-forward hops for large flows, but circuit setup and tear-down add control overhead [18]. Microsecond circuit switching can reduce that overhead, yet it requires tight integration with the electrical packet network and accurate traffic classification [19]. AI scheduling can help by deciding which flows should use optical circuits, when circuits should be reconfigured, and how much electrical buffering should be reserved for latency-sensitive traffic.

A key trade-off is between proactive and reactive scheduling. Proactive scheduling uses predictions to establish circuits before demand arrives, reducing waiting time but risking wasted capacity if predictions are wrong. Reactive scheduling waits for demand signals, reducing wasted capacity but potentially increasing tail latency. Reinforcement learning can learn a policy that blends both modes according to workload phase, tenant priority, and network condition [9]. However, the reward function must be designed carefully. If the reward emphasizes average latency, the agent may neglect tail latency and fairness. If it emphasizes utilization, the agent may delay small flows behind large transfers. If it emphasizes energy, the agent may consolidate traffic in ways that create hotspots. Multi-objective reinforcement learning and constrained optimization offer ways to encode these concerns, but they also introduce governance questions about who defines the weights and how trade-offs are audited [16].

Cross-layer coordination is also necessary for congestion control. An optical scheduler that changes capacity without informing transport protocols can cause queue buildup or underutilization. Conversely, a congestion control algorithm that reacts slowly to optical reconfiguration can waste the low-latency benefit. Research on learning-based congestion

control suggests that agents can adapt to changing network conditions, but safety and fairness remain difficult [10]. In optical data centers, the problem is compounded by the fact that capacity may change in discrete steps and that some paths may be unidirectional or asymmetric. A practical design may use explicit signaling between the optical control plane and transport endpoints, or it may embed scheduling logic in the network interface card and switch firmware. Either way, the AI component must be observable and interruptible. Operators need to know why a circuit was established, which flows benefited, and whether any tenant was disadvantaged.

5. Robustness, Fairness, and Governance

Robustness in AI-driven optical scheduling has at least three dimensions. The first is resilience to failures: fiber cuts, switch malfunctions, amplifier degradation, and software bugs can remove capacity suddenly. A learned policy trained on normal conditions may behave poorly during rare events. The second is resilience to distribution shift: workloads change, tenants change, and hardware ages. A policy that memorizes traffic patterns may fail when a new application is deployed. The third is resilience to adversarial or strategic behavior: tenants may attempt to game the scheduler by shaping traffic or misreporting requirements. Robust design therefore requires redundancy, conservative fallback modes, continuous monitoring, and periodic retraining with human oversight [15].

Fairness is equally important. In a multi-tenant data center, optical capacity is a shared resource, and scheduling decisions distribute latency, bandwidth, and energy costs across customers. An AI scheduler optimized for aggregate performance may systematically disadvantage small tenants, delay-tolerant workloads, or applications with unpredictable bursts. Fairness can be defined in multiple ways: equal access, proportional sharing, max-min fairness, or priority based on contractual terms. The choice is not purely technical. It reflects business models, regulatory expectations, and ethical commitments [16]. A governance framework should therefore make fairness criteria explicit, auditable, and contestable. This may involve publishing scheduling policies in summary form, providing tenants with explanations for degraded performance, and establishing internal review processes for high-impact changes. The critical algorithms literature reminds us that algorithms are not neutral artifacts; they encode assumptions about value, order, and priority [17].

Governance also extends to safety and accountability. If an AI scheduler causes a service outage, who is responsible? The model developer, the operator, the tenant, or the hardware vendor? Clear responsibility allocation is difficult when control is distributed across cloud providers, network equipment vendors, and application teams. A useful approach is to treat the scheduler as a safety-critical control system with formal operating envelopes, fail-safe defaults, and independent audit trails. Human operators should be able to override the scheduler, and overrides should be logged and analyzed. Regulatory frameworks for AI governance emphasize transparency, accountability, and human oversight, but they must be adapted to the fast timescales and proprietary constraints of data center operations [15]. In practice, this means that AI scheduling should be evaluated not only by latency and throughput benchmarks but also by its behavior under stress, its treatment of different tenants, and its ability to explain deviations from expected performance.

6. Deployment, Sustainability, and Economics

Deployment of AI-driven optical scheduling is constrained by economics and sustainability. Data centers consume substantial electricity, and although efficiency improvements have

moderated growth, the expansion of artificial intelligence workloads has renewed concern about energy use [13]. Optical interconnects can reduce energy per bit compared with electrical links, but they also require lasers, amplifiers, cooling, and switching hardware. Energy consumption is distributed across the network, and scheduling decisions can shift load between paths, racks, and even geographic regions [14]. An AI scheduler that minimizes latency may increase energy use by keeping circuits active or by spreading traffic across long paths. Conversely, an energy-minimizing scheduler may increase latency beyond service-level objectives. Sustainable design requires multi-objective optimization that accounts for carbon intensity, renewable availability, and embodied emissions, not just operational power.

The economics of deployment also matter. Retrofitting optical switching into an existing data center is expensive, and the business case depends on workload mix, scale, and expected lifetime. Hyperscale operators may justify custom optical fabrics, while smaller enterprises may rely on commercial switches with limited programmability. AI scheduling can lower the operational cost of optical networks by improving utilization and reducing overprovisioning, but it also adds complexity, training costs, and skilled labor requirements. There is a risk that AI-driven optimization concentrates power in the hands of a few large providers that can afford the data, talent, and infrastructure. Policy interventions such as open interfaces, standardized telemetry, and interoperability requirements could reduce lock-in and enable smaller operators to benefit from scheduling intelligence. At the same time, regulators must be careful not to freeze innovation or impose one-size-fits-all rules on a rapidly changing technical field.

A further economic consideration is the relationship between scheduling and service-level agreements. If AI scheduling improves average latency but increases variance, tenants may still perceive the service as unreliable. Contracts may need to specify tail latency, jitter, and fairness guarantees rather than only mean performance. This creates a demand for measurement and verification. Third-party auditors could evaluate scheduler behavior under representative workloads, and standardized benchmarks could compare different AI approaches. Such mechanisms would also help operators avoid overclaiming the benefits of AI. The value of AI scheduling lies not in replacing human judgment but in extending the range of conditions under which the network can meet competing objectives.

7. Case Illustrations and Cross-Domain Comparisons

Intra-data center optical scheduling can be compared with wide-area traffic engineering. In wide-area networks, traffic engineering often operates on timescales of minutes to hours, using centralized optimization and measured demand matrices. In optical data centers, the timescale is much shorter, and the cost of reconfiguration is higher relative to flow duration. AI methods developed for wide-area routing may not transfer directly because they assume stable topology and abundant state information. However, the idea of hierarchical control, where fast local decisions are coordinated by slower global policies, is transferable. Similarly, elastic optical networks use routing, modulation, and spectrum assignment to adapt capacity to demand, and deep reinforcement learning has been applied to these problems [20]. Intra-data center networks can learn from this work, but they must also handle stricter latency bounds and more dynamic traffic.

Another comparison is with energy-aware scheduling in cloud computing. Cloud schedulers place virtual machines and containers to balance performance and power, often using prediction and optimization. Optical scheduling adds a network dimension: placing workloads may affect which optical paths are used and how much reconfiguration is required. Joint

optimization of compute and network scheduling is therefore attractive but difficult because it requires coordination across administrative domains and timescales. A machine learning agent could learn joint policies, but the action space becomes enormous, and safety constraints multiply. A more practical approach may be to expose network state to the compute scheduler through application programming interfaces and to let each layer optimize within its own domain, with feedback loops that prevent conflicting decisions.

Cross-domain lessons also come from safety-critical systems such as aviation and power grids. In these domains, automation is accepted only when it is transparent, bounded, and subject to human oversight. AI scheduling for optical networks is not yet safety-critical in the same sense, but as data centers become infrastructure for finance, healthcare, transportation, and public services, the consequences of failure grow. The governance literature on algorithmic accountability provides useful guidance: document training data, test for bias, monitor performance in deployment, and provide channels for appeal [16]. The critical algorithm studies literature adds that governance must consider power and context, not only technical metrics [17]. For optical scheduling, this means asking who benefits from lower latency, who bears the cost of reconfiguration, and how conflicts between tenants are resolved. These questions cannot be answered by reinforcement learning alone.

8. Research Challenges and Forward Perspectives

Several research challenges remain. First, state representation is difficult. Optical networks produce telemetry from many devices, but the most useful information for scheduling may be latent: impending incast, synchronization barriers, or failures that have not yet manifested. AI models must learn representations that generalize across workloads and hardware generations. Second, action spaces are large and constrained. Reconfiguring an optical circuit affects multiple wavelengths and ports, and some actions are irreversible for a period. Safe reinforcement learning, constrained policy optimization, and model-based control are promising directions, but they require rigorous validation. Third, evaluation is hard. Simulation can test many scenarios, but it may not capture physical-layer effects, software bugs, or tenant behavior. Testbeds and emulation environments are needed, along with benchmarks that measure tail latency, fairness, and energy.

Fourth, multi-agent coordination is underexplored. In a large data center, multiple controllers may manage different pods, tenants, or layers. If each uses a learned policy, their interactions can produce emergent behavior that is difficult to predict. Mechanism design, game theory, and federated learning may offer tools, but they must account for the physical constraints of optical switching. Fifth, sustainability metrics must be integrated into scheduling objectives. Carbon-aware scheduling is well developed for compute, but carbon-aware optical routing is less mature. It requires information about energy sources, network topology, and the embodied emissions of hardware. Sixth, policy and governance frameworks must keep pace. Standards for telemetry, explainability, and auditability could reduce risk and build trust. International cooperation may be needed because data center supply chains and cloud services cross borders.

Forward-looking perspectives suggest that the most successful systems will not be fully autonomous. They will be human-machine systems in which AI handles high-dimensional, fast-moving optimization while humans set objectives, define constraints, and intervene in exceptional cases. The optical layer will become more programmable, with open interfaces and standardized abstractions that allow scheduling algorithms to be ported across vendors. Hollow-core fiber and advanced switching technologies will reduce physical latency, making

control-plane efficiency even more important [5]. In this environment, the competitive advantage of a data center may depend less on raw bandwidth than on the intelligence and governance of its scheduling layer. The challenge is to design that layer so that it is fast enough for microsecond-scale decisions, fair enough for multi-tenant operation, robust enough for critical services, and transparent enough for public accountability.

9. Conclusion

AI-driven traffic scheduling for low-latency data center optical networks is a system-level problem that cannot be reduced to a single algorithm or hardware component. Optical interconnects promise lower propagation delay, higher bandwidth, and improved energy efficiency, but they also introduce reconfiguration overhead, wavelength contention, and synchronization requirements. Artificial intelligence can help by predicting demand, coordinating cross-layer decisions, and adapting to changing conditions. However, the value of AI depends on architectural choices about where intelligence resides, how state is represented, how actions are constrained, and how performance is measured. Centralized, distributed, and hierarchical designs each have trade-offs, and no single approach is likely to dominate across all deployment scales.

The paper has argued that robustness, fairness, governance, sustainability, and economics are not secondary concerns. They shape the feasible design space. A scheduler that optimizes average latency while ignoring tail latency may fail service-level agreements. A scheduler that ignores fairness may disadvantage smaller tenants. A scheduler that ignores energy may undermine sustainability goals. A scheduler that lacks transparency may be difficult to audit or trust. These issues require interdisciplinary collaboration among computer scientists, optical engineers, economists, policy scholars, and ethicists. The technical community can contribute by developing safe learning methods, open benchmarks, and standard telemetry. Policymakers can contribute by supporting interoperability, accountability, and research on sustainable infrastructure. Operators can contribute by documenting decisions, testing under stress, and involving human experts in oversight.

Ultimately, low-latency optical data centers will be judged not only by how fast they move bits but by how well they allocate a shared, costly, and environmentally significant resource. AI scheduling is a powerful tool, but it is not a substitute for clear objectives, sound architecture, and legitimate governance. The next generation of optical networks should therefore be designed with intelligence embedded in a framework of constraints, not as an unconstrained optimizer. That framework must be capable of learning from experience while remaining accountable to the people and institutions that depend on the network.

References

1. Al-Fares, M., Loukissas, A., & Vahdat, A. (2008). A scalable, commodity data center network architecture. *ACM SIGCOMM Computer Communication Review*, 38(4), 63-74. <https://doi.org/10.1145/1402958.1402967>
2. Farrington, N., Porter, G., Radhakrishnan, S., Bazzaz, H. H., Subramanya, V., Fainman, Y., Papan, G., & Vahdat, A. (2010). Helios: A hybrid electrical/optical switch architecture for modular data centers. *ACM SIGCOMM Computer Communication Review*, 40(4), 339-350. <https://doi.org/10.1145/1851182.1851223>

3. 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>
4. Chen, K., Singla, A., Singh, A., Ramachandran, K., Xu, L., Zhang, Y., Wen, X., & Chen, Y. (2014). OSA: An optical switching architecture for data center networks with unprecedented flexibility. *IEEE/ACM Transactions on Networking*, 22(2), 498-511.
<https://doi.org/10.1109/TNET.2013.2253120>
5. Zhou, Z., Dzieciol, H., Clark, K., Luo, Y., Richardson, D., Poletti, F., ... & Liu, Z. (2023). Low-latency wavelength-switched clock-synchronized intra-data center interconnects enabled by hollow core fiber. *Optics Express*, 31(15), 24739-24748.
6. Musumeci, F., Rottondi, C., Nag, A., Macaluso, I., Zibar, D., Ruffini, M., & Tornatore, M. (2019). A survey on application of machine learning techniques in optical communication. *Journal of Optical Communications and Networking*, 11(2), A1-A12.
<https://doi.org/10.1364/JOCN.11.0000A1>
7. Zibar, D., Piels, M., Jones, R., & Schaeffer, C. G. (2016). Machine learning techniques in optical communication. *Journal of Lightwave Technology*, 34(6), 1442-1452.
<https://doi.org/10.1109/JLT.2015.2494940>
8. Mao, H., Alizadeh, M., Menache, I., & Kandula, S. (2016). Resource management with deep reinforcement learning. In *Proceedings of the 15th ACM Workshop on Hot Topics in Networks* (pp. 50-56). ACM. <https://doi.org/10.1145/3005745.3005750>
9. Chen, L., Lingys, J., Chen, K., & Liu, F. (2018). AuTO: Scaling deep reinforcement learning for datacenter-scale automatic traffic optimization. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication* (pp. 191-205). ACM. <https://doi.org/10.1145/3230543.3230564>
10. Jay, N., Rotman, N., Godfrey, P. B., Schapira, M., & Tamar, A. (2019). A deep reinforcement learning perspective on internet congestion control. In *Proceedings of the 36th International Conference on Machine Learning* (pp. 3050-3059). PMLR.
11. Xu, Z., Tang, J., Yin, C., Wang, Y., & Xue, G. (2018). Experience-driven networking: A deep reinforcement learning based approach. In *IEEE INFOCOM 2018* (pp. 1871-1879). IEEE. <https://doi.org/10.1109/INFOCOM.2018.8485853>
12. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529-533. <https://doi.org/10.1038/nature14236>
13. 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>
14. Dayarathna, M., Wen, Y., & Fan, R. (2016). Data center energy consumption modeling: A survey. *IEEE Communications Surveys & Tutorials*, 18(1), 732-794.
<https://doi.org/10.1109/COMST.2015.2458273>
15. 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>

16. Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2).
<https://doi.org/10.1177/2053951716679679>
17. Kitchin, R. (2017). Thinking critically about and researching algorithms. *Information, Communication & Society*, 20(1), 14-29.
<https://doi.org/10.1080/1369118X.2016.1154087>
18. Ballani, H., Costa, P., Behrendt, R., Cletheroe, D., Haller, I., Jozwik, K., ... & Thomsen, B. (2020). Sirius: A flat datacenter network with nanosecond optical switching. In *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication* (pp. 782-797). ACM. <https://doi.org/10.1145/3387514.3406221>
19. Porter, G., Strong, R., Farrington, N., Forencich, A., Chen-Sun, P., Rosing, T., ... & Vahdat, A. (2013). Integrating microsecond circuit switching into the data center. *ACM SIGCOMM Computer Communication Review*, 43(4), 447-458.
<https://doi.org/10.1145/2486001.2486006>
20. Chen, X., Guo, J., Zhu, Z., Zhang, R., & He, Y. (2019). DeepRMSA: A deep reinforcement learning framework for routing, modulation and spectrum assignment in elastic optical networks. *Journal of Lightwave Technology*, 37(16), 4155-4163.
<https://doi.org/10.1109/JLT.2019.2923615>