Austrian Journal of Technical
and Natural Sciences, 2026 No 7–8
ISSN 2310-5607

PREMIER Publishing
ppublishing.org

Section 5. Technical sciences in general

DOI:10.29013/AJT-26-7.8-86-91

OPTIMIZING EDGE TOPOLOGIES: A REVIEW OF PREDICTIVE NETWORK ROUTING ALGORITHMS

Kirill Andreevich Melnikov [1]

[1] Head of the Center for Applied Artificial Intelligence,
Synergy of Digital Solutions JSC, St. Petersburg

Cite: *Melnikov K.A. (2026). Optimizing Edge Topologies: A Review of Predictive Network Routing Algorithms. Austrian Journal of Technical and Natural Sciences 2026, No 7–8.* https://doi.org/10.29013/AJT-26-7.8-86-91

Abstract

The literature on machine-learning routing rests on an assumption that is rarely stated and almost never tested: network topology is treated as a fixed input over which a routing model operates. At the edge, where reconfiguration of links and nodes is the ordinary operating condition, this assumption does not hold. Packet-level evaluation of reinforcement-learning routing policies indicates that models trained under simplified network dynamics degrade once deployed against realistic traffic, and the degradation traces to the training regime. Drawing on fault-tolerance research from multi-agent coordination architectures outside the networking literature, this review separates topology-level disruptions, which require renegotiation across subgraphs and carry lasting downstream consequences, from resource-level disruptions, which recover within a fixed and comparatively short interval. The same body of research documents a general tendency of learned policies toward overconfidence under distributional shift, of which the routing degradation examined here is a specific instance. The review concludes that the outstanding problem is one of remedy transfer: whether a predicate-based safety architecture, proposed for this overconfidence problem in bounded robotic settings, can be constructed for network topologies whose state space is open-ended and subject to continuous reconfiguration.

Keywords: *predictive network routing; edge network topology; reinforcement learning routing; topology reconfiguration; distributional shift; fault-tolerant coordination*

Introduction

Reinforcement-learning routing policies trained on fluid network models degrade once packet-level dynamics and TCP traffic are introduced, and this degradation appears at exactly the timescale on which edge topologies now operate: single-digit milliseconds (Boltres et al., 2024). Policies that outperform static baselines under simulation fall below those same baselines once evaluated in a packet-level environment under realistic traffic mixes (Boltres et al., 2024). The competence such policies display in simulation is an artifact of the environment in which they were trained, a property that does not extend to the traffic they were built to route. A related, though structurally distinct, constraint has affected Q-routing since its introduction in 1994: limitations in router computing power and network-layer design have prevented deployment at any scale that would test its claims against real traffic (Yang et al., 2022).

These two limitations, one located in the training signal and one in the deployment substrate, converge on a single conclusion: the gap between predictive routing as studied and predictive routing as deployed shows no sign of closing on its own terms. A structurally comparable gap has already been documented in a different setting. Within distributed multi-agent coordination architectures, a partition in the communication topology forces the system out of its nominal coordination regime once packet loss exceeds approximately ten percent or latency exceeds forty milliseconds; beyond this threshold, coordination proceeds only within reachable subgraphs (Mikhailiuk, 2026). What triggers this regime change is the topology itself, precisely the variable that the routing literature under review treats as fixed.

Existing surveys of AI-driven routing tend to attribute the deployment gap to insufficient data or hardware capacity (Xiao et al., 2021; Almasan et al., 2021). This attribution does not survive scrutiny of where the degradation actually occurs. Graph neural network approaches to routing organize their contributions around network modeling, routing optimization, and dynamic reinforcement-learning tasks (Jiang et al., 2024), and the benchmarks underlying all three categories are constructed on topologies held fixed across training and evaluation. Surveys of AI-enabled routing spanning SDN, 6G, and satellite domains observe that the field itself is fragmented along these subdomains, with each treating the surrounding topology as a fixed backdrop (Aktas et al., 2025).

Topology churn is the ordinary operating condition at the edge. Surveyed approaches evaluate routing decisions against traffic load, congestion, and latency; none evaluates them against topology instability itself.

Literature Review

Across a first cluster of sources, network topology is treated as a fixed input to the learning process: something the model observes and routes over, without requiring continuous renegotiation on the algorithm's part. Graph neural network routing methods belong squarely to this cluster, since their taxonomy of supervised network modeling, supervised routing optimization, and reinforcement-learning-based dynamic routing all presuppose a topology sufficiently stable to support a learned representation (Jiang et al., 2024). RouteNet, among the founding architectures in this line of work, learns to predict per-flow performance metrics from a graph representation of the network built once and held fixed for the duration of training (Rusek et al., 2020). Concrete deployments of this pattern extend the same premise: routing optimization built on a graph neural network layered over deep reinforcement learning, validated within a single knowledge-defined-networking topology (He et al., 2023), and a deep-deterministic-policy-gradient routing scheme for software-defined networks evaluated on a fixed switch topology throughout (Kim et al., 2022). Work situating GNNs within communication networks acknowledges generalization across topologies as an open problem, which at least names the difficulty even where it does not resolve it (Suárez-Varela et al., 2022). Comparative surveys of machine-learning routing schemes report a parallel pattern: the training and deployment frameworks proposed for intelligent routing are validated against fixed simulated topologies, leaving cross-topology transfer for subsequent work (Yang et al., 2022; Li et al., 2022).

One study within this cluster addresses topology change directly: a multi-agent deep reinforcement learning scheme combined with graph neural networks is evaluated explicitly against link and node failures, with routing performance reported to degrade gracefully under such changes instead of collapsing outright (Bhavanasi et al., 2023). This partially addresses the gap identified here. The evaluation still holds the space of possible topologies fixed in advance; topology reconfiguration is not treated as an open-ended variable of the kind edge deployments present.

A comparable blind spot appears, under a different name, in a different coordination substrate. Mikhailiuk (2026) assigns each of several coordination paradigms, centralized, decentralized, consensus-based, and predictive, a distinct characteristic degradation pattern. For predictive coordination, the paradigm architecturally closest to a learned routing policy insofar as both operate by exchanging anticipated future states, the documented pattern is model drift. The same account holds that the organization of coordination is bound to the topology through which information propagates, treating topology as a variable coupled to the degradation it produces. The GNN routing literature reviewed above does not test its fixed-topology assumption against this coupling. Removing the Mikhailiuk citation from this passage would leave the review without its only account linking topology and degradation pattern in structural terms.

A second cluster treats topology as background architecture: the fixed structure inside which some other quantity is optimized. Surveys of AI-based edge computing classify edge devices into front-end, near-end, and far-end tiers and compare response time and storage capacity within that fixed classification (Pooyandeh & Sohn, 2021). Cost-optimization surveys of edge computing extend the same logic, evaluating resource-scheduling, task-offloading, and service-provisioning decisions against a topology held constant across the evaluation window (Cao et al., 2024). Federated-learning research supplies the exception: topology-aware federated learning explicitly models star, tree, decentralized, and hybrid topologies as design variables bearing on communication cost and convergence (Wu et al., 2024).

Mikhailiuk's regime model, applied to this second cluster, marks the same distinction in numerical terms. Different admissible bounds are assigned within that framework to different disruption sources: a communication-topology partition is bounded by a subgraph convergence requirement of 200 milliseconds, whereas a node loss, the resource-level counterpart of a cost-optimization decision, is bounded by a coordination-error threshold of 0.05 together with a five-second recovery window (Mikhailiuk, 2026). These two sources of disruption are not interchangeable within that framework, yet most edge-optimization surveys fold them together under a single heading of "optimization." Cost-optimization decisions of the kind reviewed by Cao et al. (2024), such as offload placement or service scheduling within a fixed topology, correspond to the resource-level disruption and its shorter, less costly recovery path; topology reconfiguration corresponds instead to the disruption requiring subgraph-wide renegotiation. Collapsing both into a single, undifferentiated efficiency problem discards a distinction that carries an explicit numerical basis in Mikhailiuk's model.

A third cluster of sources, addressing objectives instead of topology, treats link utilization and network availability under an outage as a single scalar target subject to smooth trade-off. Traffic-engineering research applying Value at Risk to WAN routing rejects this treatment directly: operators pursuing a pure utilization objective sustain markedly less throughput at a given availability level than operators using a risk-explicit formulation, by a factor of approximately two across the reported benchmarks (Bogle et al., 2019). A subsequent learning-accelerated WAN traffic-engineering system reports substantial speedups in solving the underlying optimization problem but inherits the same single-objective framing, evaluating solution quality chiefly against utilization (Xu et al., 2023). Surveys of AI-enabled routing that list delay, loss rate, and signal-to-noise ratio as parallel objectives (Aktas et al., 2025) do not distinguish those objectives that behave like utilization, tradeable one against another, from those that behave like availability, which call for a risk-bounded guarantee in place of a weighted average.

A utilization trade-off carries limited consequence, since the next traffic-engineering cycle can rebalance load without lasting effect. An outage on a routed link carries a different order of consequence: it propagates through every flow on the affected path before any correction can take effect. Bogle et al. (2019) treat it, for this reason, as belonging to a distinct risk category, separate from the optimization curve along which utilization is ordinarily assessed.

Discussion

Two unresolved gaps emerge from the literature review. Predictive routing research treats topology overwhelmingly as fixed input, even where the edge renders topology the actual decision variable. Accounts of why learned policies degrade silently under distributional shift remain untested against the specific pattern documented in the packet-level literature. A diagnosis of policy overconfidence carries practical value where the systems it describes are fitted with a mechanism reporting their approach to the boundary of competence; the routing literature reviewed here supplies no such mechanism.

Yang et al. (2022) frame the persistent deployment gap for intelligent routing chiefly as a hardware-scaling problem, pending faster routers and larger state-action spaces. The packet-level evidence points instead toward an architectural explanation. FieldLines and M-Slim, the two algorithms proposed specifically to correct the fluid-model mismatch, require no additional inference-time compute relative to the reinforcement-learning baselines they outperform; what they require instead is a different training signal, derived from packet-level simulation instead of fluid simulation (Boltres et al., 2024). A faster router executing a policy trained on the wrong abstraction of network dynamics does not thereby converge on a correct one.

Mikhailiuk's account of neural-policy behavior under distributional shift performs analytical work here that extends beyond naming the pattern. Developed in the setting of safety-critical multi-agent control, the claim holds that neural policies remain structurally uncontrolled outside their training distribution and are frequently overconfident in precisely the states where they prove least reliable (Mikhailiuk, 2026). The routing degradation reported by Boltres et al. (2024) is a specific instance of this general property, observed here in a routing policy instead of the robotic control policies for which the claim was originally developed.

Mikhailiuk's own architectural response to that overconfidence problem, however, does not transfer to routing as readily as the underlying diagnosis. The proposed remedy takes the form of a hybrid neural-symbolic architecture, in which a formally verified symbolic filter, whose safety guarantee is independent of the neural component's confidence, intercepts unsafe actions before execution and issues a tier-transition signal to a degradation controller whenever no admissible action survives filtering (Mikhailiuk, 2026). The guarantee this filter offers depends on predicate coverage: a library of formally specified conditions sufficiently broad to describe the states the system will actually encounter. Predicate coverage of this kind becomes intractable to construct once the underlying state space is high-dimensional and continuous, a limitation Mikhailiuk states explicitly, not one left to inference. A routing topology at the edge, reconfiguring across an open-ended and continually changing set of links and nodes, presents precisely this kind of state space.

Conclusion

The routing degradation documented at packet-level, millisecond-scale operation is a training-regime problem. No routing system reviewed here carries an observable trigger for recognizing when it operates outside the regime under which it was validated. Reading this literature together makes that visible in a way no single source in it does on its own.

Three assumptions traced through the literature review, that topology constitutes fixed input, that topology-level and resource-level decisions carry equivalent automation risk, and that utilization and availability trade off along a single curve, converge on a common blind spot: each treats a structural property of the system as background.

The practical consequence of this blind spot is specific to where predictive routing is next deployed. A misrouted flow within a data-center fabric constitutes, in the terms set out by Bogle et al. (2019), a utilization problem that the next traffic-engineering cycle absorbs. A misrouted flow across an edge topology that has just reconfigured, carrying safety-critical telemetry or a control-plane update bearing on the reconfiguration itself, presents a different order of problem: the topology against which correction would need to be made may already have changed again before the next cycle runs. Silent degradation of the kind Boltres et al. (2024) report proves tolerable in the first setting and considerably less so in the second, a distinction the routing literature under review does not yet draw when reporting a model's performance.

The open question that remains is a narrow one: whether a predicate library, of the kind Mikhailiuk's hybrid architecture requires for its symbolic safety filter, can be constructed for a routing topology at edge scale, or whether the state space presented by arbitrary, continuously reconfiguring network topologies is precisely the high-dimensional, open-ended environment that the source identifies as resistant to such treatment. Either outcome leaves intact the diagnosis that routing degradation constitutes a distributional-shift problem; a negative answer would indicate only that the remedy proposed for this problem in an adjacent domain does not carry over to routing without further development on the predicate side.

Mikhailiuk (2026) specifies what a formally verified safety filter requires in order to function: a predicate library covering the operationally relevant state distinctions the system will encounter. Whether such a library can be constructed for network topologies at the scale and reconfiguration rate documented in this review remains an open question, and it is the question the next stage of this line of inquiry must answer before a hybrid neural-symbolic remedy can be extended to routing systems, beyond the diagnosis of the underlying problem alone.

References

Aktas, F., Shayea, I., Ergen, M., Saoud, B., Yahya, A. E., & Aldasheva, L. (2025). AI-enabled routing in next generation networks: A survey. Alexandria Engineering Journal, – 120. – P. 449–474. URL: https://doi.org/10.1016/j.aej.2025.01.095

Almasan, P., Suárez-Varela, J., Wu, B., Xiao, S., Barlet-Ros, P., & Cabellos-Aparicio, A. (2021). Towards real-time routing optimization with deep reinforcement learning: Open challenges. In 2021 IEEE 22nd International Conference on High Performance Switching and Routing (HPSR) (pp. 1–6). IEEE. URL: https://doi.org/10.1109/HPSR52026.2021.9481864

Bhavanasi, S. S., Pappone, L., & Esposito, F. (2023). Dealing with changes: Resilient routing via graph neural networks and multi-agent deep reinforcement learning. IEEE Transactions on Network and Service Management, – 20(3). – P. 2283–2294. URL: https://doi.org/10.1109/TNSM.2023.3287936

Bogle, J., Bhatia, N., Ghobadi, M., Menache, I., Bjørner, N., Valadarsky, A., & Schapira, M. (2019). TeaVaR: Striking the right utilization-availability balance in WAN traffic engineering. In Proceedings of the ACM Special Interest Group on Data Communication (SIGCOMM '19) (P. 29–43). Association for Computing Machinery. URL: https://doi.org/10.1145/3341302.3342069

Boltres, A., Freymuth, N., Jahnke, P., Karl, H., & Neumann, G. (2024). Learning sub-second routing optimization in computer networks requires packet-level dynamics. Transactions on Machine Learning Research.

Cao, L., Huo, T., Li, S., Zhang, X., Chen, Y., Lin, G., Wu, F., Ling, Y., Zhou, Y., & Xie, Q. (2024). Cost optimization in edge computing: A survey. Artificial Intelligence Review, – 57(11). URL: https://doi.org/10.1007/s10462–024–10947–4

He, Q., Wang, Y., Wang, X., Xu, W., Li, F., Yang, K., & Ma, L. (2023). Routing optimization with deep reinforcement learning in knowledge defined networking. IEEE Transactions on Mobile Computing, – 23(2). – P. 1444–1455. URL: https://doi.org/10.1109/TMC.2023.3235446

Jiang, W., Han, H., Zhang, Y., Wang, J., He, M., Gu, W., Mu, J., & Cheng, X. (2024). Graph neural networks for routing optimization: Challenges and opportunities. Sustainability, – 16(21). – 9239 p. URL: https://doi.org/10.3390/su16219239

Kim, G., Kim, Y., & Lim, H. (2022). Deep reinforcement learning-based routing on software-defined networks. IEEE Access, – 10. – P. 18121–18133. URL: https://doi.org/10.1109/ACCESS.2022.3151081

Li, T., Zhu, K., Luong, N. C., Niyato, D., Wu, Q., Zhang, Y., & Chen, B. (2022). Applications of multi-agent reinforcement learning in future internet: A comprehensive survey. IEEE Communications Surveys & Tutorials, – 24(2). – P. 1240–1279.

Mikhailiuk, M. (2026). Predictive technologies and autonomous decision making. LAP LAMBERT Academic Publishing.

Pooyandeh, M., & Sohn, I. (2021). Edge network optimization based on AI techniques: A survey. Electronics, – 10(22). – 2830 p. URL: https://doi.org/10.3390/electronics10222830

Rusek, K., Suárez-Varela, J., Almasan, P., Barlet-Ros, P., & Cabellos-Aparicio, A. (2020). RouteNet: Leveraging graph neural networks for network modeling and optimization in SDN. IEEE Journal on Selected Areas in Communications, – 38(10). – P. 2260–2270.

Suárez-Varela, J., Almasan, P., Ferriol-Galmés, M., Rusek, K., Geyer, F., Cheng, X., Shi, X., Xiao, S., Scarselli, F., Cabellos-Aparicio, A., & Barlet-Ros, P. (2022). Graph neural networks for communication networks: Context, use cases and opportunities. IEEE Network, – 37(3). – P. 146–153. URL: https://doi.org/10.1109/MNET.001.2200269

Wu, J., Dong, F., Leung, H., Zhu, Z., Zhou, J., & Drew, S. (2024). Topology-aware federated learning in edge computing: A comprehensive survey. ACM Computing Surveys, – 56(10). – P. 1–41. URL: https://doi.org/10.1145/3659205

Xiao, Y., Liu, J., Wu, J., & Ansari, N. (2021). Leveraging deep reinforcement learning for traffic engineering: A survey. IEEE Communications Surveys & Tutorials, – 23(4). – P. 2064–2097. URL: https://doi.org/10.1109/COMST.2021.3102580

Xu, Z., Yan, F. Y., Singh, R., Chiu, J. T., Rush, A. M., & Yu, M. (2023). Teal: Learning-accelerated optimization of WAN traffic engineering. In Proceedings of the ACM SIGCOMM 2023 Conference (P. 378–393). Association for Computing Machinery.

Yang, S., Tan, C., Madsen, D. Ø., Xiang, H., Li, Y., Khan, I., & Choi, B. J. (2022). Comparative analysis of routing schemes based on machine learning. Mobile Information Systems, 2022, Article 4560072. URL: https://doi.org/10.1155/2022/4560072

submitted 04.08.2026;
accepted for publication 18.08.2026;
published 31.08.2026

Contact: kirill.andreev.melnikov@gmail.com