

Risk-Aware Trajectory Optimization for Complex Multi-Agent Transportation Systems

Ayesha Dowley

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

Abstract

Complex multi-agent transportation systems increasingly combine connected vehicles, automated fleets, drones, rail operations, logistics robots, and human road users within shared physical and digital infrastructures. Trajectory optimization in such settings cannot be reduced to collision avoidance or local efficiency, because risk is produced by interactions among heterogeneous agents, uncertain predictions, infrastructure constraints, regulatory choices, and social expectations. This paper develops a system-level account of risk-aware trajectory optimization for multi-agent transportation. It argues that risk should be treated as a structural property of socio-technical architectures rather than as a residual probability attached to individual maneuvers. The discussion examines layered coordination, centralized and decentralized planning, edge and cloud intelligence, vehicle-to-everything communication, and the trade-offs among latency, scalability, privacy, resilience, and fairness. It then analyzes optimization under uncertainty, including prediction error, distributional shift, rare events, and adversarial conditions, while considering how scenario generation and simulation can stress-test planners in high-density environments. The paper further considers governance, infrastructure investment, certification, liability, environmental sustainability, distributive justice, and public legitimacy. It identifies robustness and graceful degradation as central design requirements, especially where communication fails, sensors degrade, or cascading disruptions cross system boundaries. Case illustrations from urban intersections, highway platooning, ports, aviation, and rail clarify how governance regimes and safety cultures shape technical choices. The conclusion proposes an interdisciplinary research agenda that integrates risk metrics, institutional design, infrastructure planning, and participatory policy. The central claim is that risk-aware trajectory optimization must be evaluated not only by accident reduction but also by its effects on system resilience, access, environmental performance, and democratic accountability.

Keywords

risk-aware trajectory optimization; multi-agent transportation; socio-technical systems; governance; resilience; fairness; infrastructure; deployment.

1. Introduction

Transportation systems are becoming computationally intensive environments in which vehicles, infrastructure operators, logistics platforms, and public agencies continuously exchange information and make interdependent decisions. In this context, trajectory optimization refers to the planning and adjustment of movement paths for multiple agents over time, but its significance extends far beyond kinematic feasibility. It determines how road capacity is used, how risk is distributed, how energy is consumed, how public space is allocated, and how institutions justify automated decisions. The rise of connected and automated mobility has therefore made trajectory optimization a problem of system

architecture and governance, not merely a problem of control engineering. Surveys of motion planning for self-driving urban vehicles show that technical methods have matured across search, sampling, optimization, and learning-based approaches, yet the system-level challenges of coordination, verification, and accountability remain unresolved [1]. Reviews of highway motion planning similarly reveal that planning algorithms are often evaluated in isolated scenarios rather than in the dense, mixed, and institutionally constrained environments where deployment occurs [2].

The multi-agent character of transportation creates distinctive risks. Each agent's optimal action depends on predictions about others, and those predictions depend on social norms, vehicle capabilities, communication quality, and local rules. In partially observable settings, agents must plan under uncertainty about hidden intentions and future observations, a problem long recognized in decision-theoretic planning [3]. Computational methods for sequential decision making under uncertainty provide useful foundations, but they do not by themselves resolve the collective risk problems that arise when many agents optimize simultaneously [4]. Formal models of safe self-driving cars attempt to define responsibility-sensitive safety conditions, yet their assumptions about perception, communication, and regulatory enforcement must be examined carefully when scaled to heterogeneous fleets [5]. Risk-aware trajectory optimization must therefore combine algorithmic reasoning with an understanding of how failure propagates through physical, digital, and institutional layers.

This paper argues that risk in multi-agent transportation is not simply the probability of a collision. It is a structural property generated by architecture, incentives, infrastructure, and governance. A planner that minimizes expected harm for one vehicle may increase systemic risk by inducing unpredictable behavior, degrading throughput, or shifting danger onto pedestrians and cyclists. A policy that prioritizes smooth traffic flow may improve emissions but exclude users who cannot afford connected vehicles. A deployment that relies on centralized coordination may achieve efficient trajectories but become vulnerable to communication outages, cyberattacks, or single-point failures. The paper therefore proceeds from conceptual foundations to system architecture, optimization under uncertainty, governance, sustainability, fairness, resilience, cases, and a forward-looking research agenda. The aim is to articulate an interdisciplinary framework for evaluating risk-aware trajectory optimization as a socio-technical design problem.

2. Risk as a Structural Property of Multi-Agent Systems

Risk in transportation has traditionally been framed through accident statistics, exposure measures, and safety envelopes. In multi-agent systems, however, risk emerges from interaction patterns. Two vehicles may each behave safely in isolation while creating danger through mutual anticipation, competitive merging, or ambiguous signaling. The problem becomes acute when automated agents interact with human drivers, pedestrians, and cyclists whose behavior is shaped by culture, age, impairment, and local practice. Motion prediction and risk assessment surveys show that uncertainty in other agents' intentions is a central obstacle for intelligent vehicles [10]. A risk-aware planner must therefore represent not only physical uncertainty but also social and institutional uncertainty. It must ask how other agents interpret its behavior, how rules are enforced, and how liability is assigned when norms conflict.

A second structural dimension is cascading failure. In tightly coupled systems, small local disturbances can propagate through platoons, intersections, logistics hubs, and communication networks. Normal accident theory emphasizes that catastrophic failures often

arise from unexpected interactions rather than from a single broken component [8]. Resilience engineering adds that safety depends on the system's capacity to anticipate, monitor, respond, and learn [9]. For trajectory optimization, this means that risk metrics should capture systemic fragility, not only instantaneous collision probability. A planner that assumes perfect communication may be efficient under nominal conditions but brittle during outages. A planner that reserves excessive safety margins may be robust but wasteful, reducing capacity and increasing emissions. The design challenge is to allocate risk intelligently across agents, time, and infrastructure layers.

A third dimension is distributional risk. Different road users face different levels of exposure and have different capacities to avoid harm. Pedestrians, cyclists, wheelchair users, and low-income commuters may bear risks created by automated fleets optimized for vehicle throughput. Ethical debates about automated vehicle crashes have highlighted dilemmas involving unavoidable harm, but real-world fairness concerns extend to everyday routing, curb access, crossing time, and enforcement [16]. Ethical frameworks for autonomous cars often focus on discrete dilemmas, yet systemic justice requires examining how risk is routinely distributed through design and policy [17]. If risk-aware trajectory optimization is to be legitimate, it must make distributional choices visible and contestable.

3. Architectural Layers and Coordination Mechanisms

Multi-agent transportation systems can be organized through centralized, decentralized, and hierarchical architectures. Centralized coordination can compute globally efficient trajectories, especially in closed environments such as ports, warehouses, and controlled highway segments. However, centralization concentrates failure risk, requires high-bandwidth communication, and raises privacy and accountability concerns. Decentralized approaches allow agents to plan locally, improving scalability and resilience, but they may produce inefficient equilibria or oscillatory behavior. Hierarchical architectures combine strategic coordination at the infrastructure level with tactical planning at the vehicle or fleet level. The question is not which architecture is universally superior but how risk is allocated across layers. Decentralized self-adaptation research shows that local adaptation can be powerful but requires careful coordination to avoid unstable collective behavior [21].

Communication infrastructure is a decisive factor. Connected vehicle systems depend on reliable, low-latency links among vehicles, roadside units, pedestrians' devices, and traffic management centers [13]. Where communication is unavailable, planners must fall back to onboard perception and conservative rules. Where communication is available, they must guard against spoofing, latency, packet loss, and adversarial interference. Edge computing can reduce latency by processing data near the roadside, but it creates new attack surfaces and governance questions about data ownership. Cloud computing can support fleet-wide learning and scenario analysis, but it may be inappropriate for safety-critical control loops. A risk-aware architecture therefore needs graceful degradation, clear authority boundaries, and auditable decision logs.

Coordination at intersections and merging points illustrates these trade-offs. Surveys of connected and automated vehicle coordination show that intersection management can improve throughput and reduce stop-and-go traffic, but also that mixed traffic, pedestrian movements, and failure modes complicate deployment [11]. Traffic flow research indicates that connected and autonomous vehicles may stabilize flow or destabilize it depending on penetration rates and control assumptions [14]. System-level planning must therefore account for transition periods in which automated, connected, and human-driven agents coexist.

During these periods, risk-aware trajectory optimization must be robust to inconsistent capabilities, partial compliance, and divergent expectations. The architectural goal is not perfect coordination but resilient coordination under imperfect information and uneven participation.

4. Trajectory Optimization under Uncertainty and Risk Constraints

Trajectory optimization under uncertainty typically balances performance objectives such as travel time, energy use, and comfort against safety constraints. In multi-agent settings, uncertainty arises from sensor noise, prediction error, human intent, mechanical degradation, weather, and communication delay. Classical planning under partial observability provides a formal language for acting despite hidden state, but exact solutions are computationally intractable in dense environments [3]. Approximate dynamic programming and reinforcement learning offer scalable alternatives, yet they require careful treatment of safety during exploration and deployment [4]. Formal safety models can provide explicit responsibility conditions, but they depend on assumptions about perception and rule compliance that may not hold in mixed traffic [5]. Risk-aware optimization must therefore combine multiple methods rather than rely on a single formalism.

A central design choice is how to represent risk. Expected-value optimization can hide rare but catastrophic outcomes. Worst-case optimization can be overly conservative and reduce system capacity. Scenario-based optimization samples plausible futures, but its performance depends on the diversity and realism of scenarios. Distributionally robust methods account for uncertainty about the probability distribution itself, which is important when deploying in new cities or unusual conditions. Risk measures such as conditional value-at-risk can prioritize tail events, but they require value judgments about which tails matter and who bears them. These choices are not purely technical; they encode policy preferences about safety, efficiency, and equity.

Simulation and scenario generation are essential for testing risk-aware planners before deployment. High-density traffic and diverse agent-interaction scenarios are particularly important because rare interaction patterns can reveal failure modes that simple test tracks miss [15]. A trajectory generator for complex environments can help evaluate planners across heterogeneous agents, varying densities, and ambiguous social cues, but it cannot replace real-world validation. The gap between simulated and deployed behavior remains a major source of risk. Digital twins can partially bridge this gap by updating models with operational data, yet they also raise concerns about data quality, feedback loops, and accountability. The most credible approach combines simulation, closed-course testing, staged deployment, and continuous monitoring with clear criteria for rollback.

Risk-aware optimization also involves trade-offs between individual and collective rationality. A vehicle that aggressively claims right-of-way may reduce its own travel time while increasing systemic uncertainty. A fleet that coordinates to smooth traffic may improve fuel efficiency while disadvantaging non-connected users. A pedestrian detection system tuned for adult pedestrians may fail for children or wheelchair users. These examples show that optimization objectives must be specified through inclusive deliberation, not solely through engineering convenience. The system should be evaluated on how it performs across user groups, weather conditions, road types, and failure scenarios. Risk awareness thus becomes an ongoing socio-technical practice, not a one-time algorithmic setting.

5. Governance, Infrastructure, and Deployment

Deployment of risk-aware trajectory optimization depends on governance regimes that define standards, liability, data rights, and public accountability. Certification approaches developed for conventional vehicles assume a human driver and relatively stable operating conditions. Automated multi-agent systems challenge these assumptions because behavior emerges from software updates, communication networks, and fleet-level policies. Regulators must therefore develop adaptive oversight that can assess not only components but also interactions. Regulatory sandboxes, staged deployment corridors, and mandatory incident reporting can create learning environments while protecting public safety. However, governance must avoid capture by dominant firms and ensure that small municipalities and underserved communities are not excluded from decision making.

Infrastructure investment shapes what is technically feasible. Dedicated lanes, smart intersections, roadside sensing, high-definition mapping, and reliable power and communication are preconditions for advanced coordination. Yet infrastructure is unevenly distributed, and investment decisions can reinforce spatial inequality. A region with advanced digital infrastructure may attract automated logistics and mobility services, while rural and low-income areas remain dependent on aging roads and limited transit. Technological transitions research shows that systemic change involves alignment among niches, regimes, and landscapes, not merely the diffusion of devices [20]. Risk-aware trajectory optimization should therefore be planned within broader mobility transitions that include public transit, active travel, freight, and land use.

Deployment also raises questions about public legitimacy. If automated systems make risk-laden decisions that are difficult to explain, trust may erode. Transparency should not mean disclosing proprietary algorithms in full, but it should provide meaningful information about operating conditions, known limitations, and avenues for redress. Liability must be clear enough to incentivize safety without stifling innovation. Insurance models may need to shift from individual driver fault to shared responsibility among manufacturers, operators, infrastructure providers, and regulators. Governance should also address cybersecurity, because a compromised fleet or traffic management center could cause physical harm at scale. Risk-aware trajectory optimization is therefore inseparable from institutional design.

6. Sustainability, Fairness, and Policy Implications

Sustainability is often framed as an emissions problem, but multi-agent trajectory optimization affects sustainability through congestion, vehicle miles traveled, mode shift, and land use. Smoothing traffic flow can reduce energy consumption and emissions, especially for conventional and hybrid vehicles. However, optimized automated mobility may also induce additional travel demand, encourage sprawl, and compete with public transit. If shared automated vehicles replace walking, cycling, and bus trips, the net environmental effect may be negative. Policy must therefore integrate trajectory optimization with broader goals such as reducing vehicle dependence, prioritizing transit, and protecting public space. Risk-aware planning should include environmental risk as a first-class concern rather than treating it as an externality.

Fairness is equally central. Risk-aware systems can distribute safety benefits and burdens unevenly. A routing policy that avoids certain neighborhoods may reduce risk for passengers while concentrating traffic and pollution elsewhere. A curb management system that prioritizes automated deliveries may disadvantage small businesses or pedestrians. Justice theories remind us that fairness is not only about equal treatment but also about the basic structure of social cooperation and the recognition of difference [18]. A theory of justice that

focuses on the least advantaged can guide the evaluation of mobility policies, while attention to group difference reveals how universal design can mask exclusion [19]. Fairness in trajectory optimization requires disaggregated data, inclusive scenario design, and community participation in setting acceptable risk levels.

Policy implications include the need for risk disclosure, equity impact assessments, and open standards. Governments can require that automated mobility deployments report near misses, disengagements, and distributional impacts across neighborhoods. They can fund public infrastructure that supports safe interactions among all road users. They can establish data trusts or privacy-preserving mechanisms that allow learning without surveillance. They can also use procurement power to demand accessibility, interoperability, and environmental performance. The goal is not to slow innovation but to ensure that innovation is accountable to public values. Risk-aware trajectory optimization should be understood as a public infrastructure problem, not only a commercial product.

7. Robustness, Resilience, and Failure Modes

Robustness refers to the ability to maintain performance under expected disturbances, while resilience refers to the ability to absorb, adapt to, and recover from unexpected shocks. Multi-agent transportation systems face both. Sensor degradation, communication loss, adverse weather, and cyberattacks are foreseeable threats. Cascading failures, novel edge cases, and institutional breakdowns are less predictable. Systems thinking for safety emphasizes that accidents arise from interactions among components, organizations, and operators rather than from isolated component failures [7]. Trajectory optimization must therefore be designed with safety constraints that remain valid when perception, communication, or coordination degrades. A planner that cannot degrade gracefully may turn a temporary fault into a collision.

Resilience requires redundancy and diversity. Vehicles might rely on multiple sensing modalities, communication channels, and fallback rules. Infrastructure might support both centralized and decentralized coordination. Fleets might maintain heterogeneous planning strategies to avoid correlated failure. However, redundancy also adds cost, complexity, and maintenance burden. The design question is how much redundancy is justified for different risk levels and operational contexts. In aviation and rail, safety-critical systems use rigorous redundancy, certification, and incident learning. Road transportation has traditionally relied on distributed human judgment and informal norms. As automation increases, road systems may need to adopt stronger safety cultures without losing the flexibility that makes streets work in diverse contexts.

Artificial intelligence introduces additional failure modes. Concrete problems in AI safety include specification gaming, negative side effects, distributional shift, and reward hacking [6]. In transportation, a planner might optimize a proxy for safety while creating new hazards. It might overfit to simulation and fail in rare real-world conditions. It might learn discriminatory behavior from biased data. These risks are not purely algorithmic; they are amplified by deployment incentives, weak oversight, and opaque supply chains. Robustness therefore requires technical verification, operational monitoring, organizational learning, and regulatory enforcement. Resilience is not a property that can be certified once; it must be sustained through continuous adaptation.

8. Case Illustrations and Cross-Domain Comparisons

Urban intersections provide a compelling case. They concentrate conflicts among vehicles, pedestrians, cyclists, and micromobility users. Connected coordination can reduce delay and

emissions, but it must handle mixed traffic and vulnerable road users. A risk-aware system might prioritize pedestrian crossing time, allocate phases dynamically, and communicate intentions to human users. Yet if the system depends on smartphone connectivity, it may exclude people without devices or create privacy risks. The design must therefore combine infrastructure sensing, vehicle intelligence, and accessible interfaces. Cross-domain comparison with air traffic control shows that centralized coordination can be safe but expensive and procedurally rigid. Road systems are more open and heterogeneous, so they require different governance and fallback strategies.

Highway platooning illustrates trade-offs between efficiency and resilience. Closely spaced automated trucks can reduce fuel consumption and improve lane capacity, but they also create correlated risk. A cyberattack, communication outage, or mechanical failure in a platoon could affect multiple vehicles simultaneously. Platooning therefore requires secure communication, fail-safe braking, and clear responsibility among carriers and manufacturers. Ports and logistics hubs offer more controlled environments where centralized scheduling can work well. However, ports connect to public roads and rail networks, so risk can migrate across boundaries. Rail systems demonstrate the value of strict separation, formal signaling, and mature safety regulation, but they also show the cost of inflexibility. These cases suggest that risk-aware trajectory optimization should be tailored to the coupling, openness, and governance capacity of each domain.

9. Research Agenda and Forward-Looking Perspectives

A forward-looking research agenda should integrate technical, institutional, and social dimensions. First, risk metrics must become multidimensional. They should capture collision risk, near misses, cascading potential, environmental impact, accessibility, and distributional equity. Second, optimization methods must be evaluated in mixed, high-density, and adversarial scenarios rather than in narrow benchmarks. Scenario generation and digital twins can help, but they must be validated against real-world data and inclusive of vulnerable users [15]. Third, architectures must support graceful degradation across communication, sensing, and coordination failures. This requires research on fallback policies, authority transitions, and human oversight.

Fourth, governance must be designed alongside algorithms. Certification regimes should assess system-level behavior, not only component compliance. Liability rules should clarify responsibility when multiple agents and infrastructures contribute to harm. Data governance should protect privacy while enabling learning. Fifth, public participation should shape acceptable risk and priority setting. Communities should have a voice in where automated services operate, how data are used, and how benefits are distributed. Sixth, sustainability and fairness should be treated as constraints and objectives, not afterthoughts. A system that reduces emissions but excludes disabled users or concentrates pollution is not sustainable in a meaningful sense.

Finally, research should examine transition dynamics. The co-existence of human-driven, connected, and automated vehicles will persist for decades. Planners must be robust to uneven penetration rates, inconsistent capabilities, and changing regulations. They must also anticipate feedback effects: automation may change travel demand, land use, and labor markets. Technological transitions are evolutionary and contested, involving alignment across niches, regimes, and landscapes [20]. Risk-aware trajectory optimization should therefore be studied as part of a broader mobility transition, with attention to path dependence, lock-in,

and alternative futures. The most valuable innovations may be those that improve safety and access across modes, not only those that optimize automated vehicle motion.

10. Conclusion

Risk-aware trajectory optimization for complex multi-agent transportation systems cannot be reduced to better prediction or faster computation. It is a socio-technical design problem in which algorithms, infrastructure, governance, and public values interact. Risk emerges from interactions among heterogeneous agents, cascading failures, distributional burdens, and institutional choices. Effective systems must coordinate across centralized and decentralized architectures, communicate reliably, degrade gracefully, and remain accountable under uncertainty. They must optimize not only travel time and energy but also safety, resilience, fairness, and environmental performance. Deployment requires adaptive governance, inclusive infrastructure planning, and transparent mechanisms for learning from incidents. The central conclusion is that risk awareness should be embedded in the entire system lifecycle, from scenario design and architecture to certification, operation, and public deliberation. Only then can trajectory optimization contribute to transportation systems that are not merely efficient but also robust, just, and trustworthy.

References

1. Paden, B., Čáp, M., Yong, S. Z., Yershov, D., & Frazzoli, E. (2016). A survey of motion planning and control techniques for self-driving urban vehicles. *IEEE Transactions on Intelligent Vehicles*, 1(1), 33–55. <https://doi.org/10.1109/TIV.2016.2578706>
2. Claussmann, L., Revilloud, M., Gruyer, D., & Glaser, S. (2019). A review of motion planning for highway autonomous driving. *IEEE Transactions on Intelligent Transportation Systems*, 21(5), 1826–1848. <https://doi.org/10.1109/TITS.2019.2913998>
3. Kaelbling, L. P., Littman, M. L., & Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1–2), 99–134. [https://doi.org/10.1016/S0004-3702\(98\)00023-X](https://doi.org/10.1016/S0004-3702(98)00023-X)
4. Bertsekas, D. P., & Tsitsiklis, J. N. (1996). *Neuro-dynamic programming*. Athena Scientific.
5. Shalev-Shwartz, S., Shammah, S., & Shashua, A. (2017). On a formal model of safe and scalable self-driving cars. *arXiv preprint arXiv:1708.06374*.
6. Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., & Mané, D. (2016). Concrete problems in AI safety. *arXiv preprint arXiv:1606.06565*.
7. Leveson, N. G. (2011). *Engineering a safer world: Systems thinking applied to safety*. MIT Press.
8. Perrow, C. (1999). *Normal accidents: Living with high-risk technologies*. Princeton University Press.
9. Hollnagel, E., Woods, D. D., & Leveson, N. (Eds.). (2006). *Resilience engineering: Concepts and precepts*. Ashgate.
10. Lefèvre, S., Vasquez, D., & Laugier, C. (2014). A survey on motion prediction and risk assessment for intelligent vehicles. *ROBOMECH Journal*, 1(1), 1–14. <https://doi.org/10.1186/s40648-014-0001-z>

11. Rios-Torres, J., & Malikopoulos, A. A. (2017). A survey on the coordination of connected and automated vehicles at intersections and merging at highway on-ramps. *IEEE Transactions on Intelligent Transportation Systems*, 18(5), 1066–1077.
<https://doi.org/10.1109/TITS.2016.2600504>
12. Mahmassani, H. S. (2016). Autonomous vehicles and connected vehicle systems: Flow and operations considerations. *Transportation Science*, 50(4), 1140–1162.
<https://doi.org/10.1287/trsc.2016.0712>
13. Shladover, S. E. (2018). Connected and automated vehicle systems: Introduction and overview. *Journal of Intelligent Transportation Systems*, 22(3), 190–200.
<https://doi.org/10.1080/15472450.2017.1420494>
14. Talebpour, A., & Mahmassani, H. S. (2016). Influence of connected and autonomous vehicles on traffic flow stability and throughput. *Transportation Research Part C: Emerging Technologies*, 71, 143–163. <https://doi.org/10.1016/j.trc.2016.07.007>
15. Yang, R., Xu, Y., Chen, Y., Fu, Y., & Su, L. (2025). A Trajectory Generator for High-Density Traffic and Diverse Agent-Interaction Scenarios. *arXiv preprint arXiv:2510.02627*.
16. Goodall, N. J. (2014). Ethical decision making during automated vehicle crashes. *Transportation Research Record*, 2424(1), 58–65. <https://doi.org/10.3141/2424-07>
17. Lin, P. (2016). Why ethics matters for autonomous cars. In M. Maurer, J. C. Gerdes, B. Lenz, & H. Winner (Eds.), *Autonomous driving* (pp. 69–85). Springer.
https://doi.org/10.1007/978-3-662-48847-8_4
18. Rawls, J. (1971). *A theory of justice*. Harvard University Press.
19. Young, I. M. (1990). *Justice and the politics of difference*. Princeton University Press.
20. Geels, F. W. (2002). Technological transitions as evolutionary reconfiguration processes: A multi-level perspective and a case-study. *Research Policy*, 31(8–9), 1257–1274.
[https://doi.org/10.1016/S0048-7333\(02\)00062-8](https://doi.org/10.1016/S0048-7333(02)00062-8)
21. Weyns, D., Malek, S., & Andersson, J. (2010). On decentralized self-adaptation: Lessons from the trenches. In *Proceedings of the 2010 ICSE Workshop on Software Engineering for Adaptive and Self-Managing Systems* (pp. 84–93). Association for Computing Machinery. <https://doi.org/10.1145/1808984.1808994>