

Simulation-to-Real Transfer Learning for End-to-End Autonomous Driving Models

Gertraut Haring

Department of Electrical Engineering and Computer Science, University of Missouri,
Columbia, MO, USA.

gertraut.haring516@missouri.edu

Irmingard Knappe

School of Information Technology, University of Cincinnati, Cincinnati, OH, USA.

irmingardk@uc.edu

Abstract

End-to-end autonomous driving models promise to unify perception, prediction, planning, and control within learned computational pipelines, but their deployment depends on whether capabilities acquired in simulation can be transferred reliably to real roads. This paper examines simulation-to-real transfer learning as a system-level problem rather than a narrow algorithmic challenge. It argues that transferability is shaped by architectural choices, data infrastructures, validation regimes, operational governance, and the socio-technical conditions of deployment. Simulation offers scalable exposure to rare and dangerous scenarios, yet it also introduces visual, dynamic, and semantic discrepancies that can induce brittle behavior when models encounter real-world distribution shifts. The paper analyzes domain randomization, adversarial adaptation, generative translation, and post-training strategies as complementary mechanisms that must be embedded within broader assurance architectures. It further discusses how scenario libraries, digital twins, fleet feedback, and model documentation can support continuous validation and accountability. Because autonomous driving systems operate in public space, questions of fairness, accessibility, liability, and sustainability cannot be treated as post-deployment concerns. The paper concludes that simulation-to-real transfer should be governed as an infrastructural capability, with shared benchmarks, traceable data provenance, safety cases, and policy frameworks that align technical performance with public interest.

Keywords

autonomous driving, simulation-to-real transfer, end-to-end learning, domain adaptation, safety assurance, governance, infrastructure.

1. Introduction

End-to-end autonomous driving models have emerged as a compelling alternative to modular stacks that separate perception, localization, prediction, planning, and control. By learning mappings from sensor inputs to driving actions or trajectories, such models can exploit large-scale data and reduce hand-engineered interfaces [1], [3]. This architecture is attractive because it can capture subtle correlations across perception and decision-making, and because it can be trained with imitation learning, reinforcement learning, or hybrid objectives. Nevertheless, the promise of end-to-end learning is constrained by a practical asymmetry: real-world driving data are expensive, unevenly distributed, and dangerous to collect at scale, whereas simulation can generate vast quantities of labeled and controllable experience [2],

[15]. Simulation-to-real transfer therefore becomes a central condition for the viability of end-to-end driving systems.

The transfer problem is often framed as a domain adaptation issue, but that framing is too narrow for deployment-scale autonomous driving. Real roads involve weather, illumination, sensor degradation, human unpredictability, construction, legal variation, and cultural driving norms that cannot be exhaustively enumerated in advance. Even a model that performs well in a simulator and on a held-out real dataset may fail under rare combinations of conditions that were absent from training and validation. This is not merely a machine learning problem; it is a systems problem involving data pipelines, compute infrastructure, testing facilities, regulatory oversight, and organizational accountability [12], [13]. Transfer learning must therefore be understood as one component within a larger assurance and governance architecture.

This paper advances three claims. First, simulation-to-real transfer for end-to-end driving should be designed as a layered capability that combines domain randomization, representation alignment, generative augmentation, and targeted real-world post-training. Second, the effectiveness of these techniques depends on the quality and diversity of the simulation and validation infrastructure, including scenario generation, sensor modeling, and digital twin maintenance. Third, because autonomous vehicles operate in public environments, transfer performance must be evaluated not only by average accuracy or collision rate but also by robustness, fairness, interpretability, environmental cost, and policy compliance [22], [23]. The analysis proceeds from the architecture of transfer to the infrastructure that sustains it, and then to the governance and sustainability questions that shape whether deployment is socially legitimate.

2. The Simulation-to-Real Gap as a System-Level Challenge

The simulation-to-real gap is commonly described in terms of visual appearance, dynamics, and semantics. Visual discrepancies arise from rendering limitations, sensor noise models, weather effects, and material properties that differ from real optics and reflectance. Dynamic discrepancies arise because simulated vehicles, pedestrians, and cyclists may not reproduce human reaction times, social norms, or multi-agent interactions. Semantic discrepancies arise when simulation simplifies road rules, signage, maps, or edge cases, producing labels and behaviors that do not correspond to real operational design domains. For end-to-end models, these discrepancies are especially consequential because the model learns directly from raw or minimally processed inputs and may encode simulator-specific shortcuts rather than generalizable driving concepts [2], [5].

A system-level view recognizes that the gap is not a single distance between two datasets. It is a set of mismatches distributed across sensors, actuators, environments, agents, and institutions. A vehicle trained in one simulator may transfer poorly to another simulator, and a model trained across multiple simulators may still fail on real roads if the simulators share common rendering assumptions. Similarly, a model fine-tuned on a small real dataset may overfit to geographic or demographic conditions. The relevant question is not whether simulation is realistic in general, but whether the simulation distribution covers the operational conditions that the deployed system will encounter, including rare and safety-critical events [14], [25].

Domain randomization has become a prominent response to this challenge. By varying textures, lighting, weather, sensor placement, and dynamic parameters during training,

randomization encourages the model to learn features that are invariant to incidental simulator characteristics [5], [6]. Dynamics randomization extends this logic to vehicle mass, friction, latency, and actuator response, which is important for control and planning. Yet randomization also introduces trade-offs. Overly broad randomization can make learning unstable or force the model to ignore informative cues, while narrow randomization may fail to cover real-world tail risks. The calibration of randomization ranges therefore becomes an empirical and governance question, not merely a hyperparameter choice.

Generative and adversarial methods offer another route. Image-to-image translation can make simulated scenes appear more realistic, and adversarial domain adaptation can align feature distributions between simulated and real domains [8], [9], [10]. These methods can improve visual transfer, but they may also hallucinate structures or remove safety-relevant details. A translated pedestrian may become visually plausible yet semantically misleading, and an adapted feature space may collapse distinctions that matter for rare-class detection. For this reason, generative transfer should be evaluated with downstream driving metrics and counterfactual tests, not only with visual realism or domain classifier accuracy [11]. A recent emphasis on post-training for end-to-end autonomous driving further suggests that broad pretraining in simulation must be followed by targeted alignment using real or high-fidelity data, especially for long-tail behaviors and decision boundaries [17].

3. Architectural Foundations and Transfer Mechanisms

The architecture of an end-to-end driving model determines how transferable its learned representations are likely to be. Some systems map camera images directly to steering and acceleration, while others predict intermediate representations such as bird's-eye-view features, occupancy grids, affordances, or trajectory heatmaps before producing control outputs [1], [4]. Intermediate representations can improve transfer by imposing spatial or physical structure, but they also introduce design assumptions that may not hold across domains. A model that relies on a particular sensor configuration, map format, or coordinate frame may transfer poorly even if its visual backbone generalizes well. Architectural modularity, therefore, is not the opposite of end-to-end learning; it can be a strategy for controlling where domain-specific and domain-invariant computation reside.

Transfer learning mechanisms can be grouped into input-level, feature-level, and parameter-level approaches. Input-level approaches modify simulated data through style transfer, sensor simulation, or augmentation. Feature-level approaches align latent representations using adversarial losses, maximum mean discrepancy, or contrastive objectives. Parameter-level approaches fine-tune pretrained models on real data, often with frozen early layers or low-rank adaptation. Each approach has different infrastructure and governance implications. Input-level methods require high-fidelity content generation and careful label preservation. Feature-level methods require access to representative real data and robust domain discriminators. Parameter-level methods require curated real-world datasets and mechanisms to prevent catastrophic forgetting. A practical system may combine all three, but the combination must be validated as an integrated pipeline rather than as isolated components [8], [10], [17].

The choice of learning objective also affects transfer. Imitation learning from expert demonstrations is common because it avoids reward specification and can leverage human driving data [1], [3]. However, imitation learning suffers from distribution shift: small errors can lead the vehicle into states not represented in the demonstration distribution. Reinforcement learning in simulation can explore beyond expert trajectories, but simulated

rewards may not capture real-world safety and comfort preferences. Hybrid approaches use simulation for pretraining and real-world fine-tuning for calibration, but they raise questions about how much real-world data is sufficient and how to prioritize collection. The answer depends on the operational design domain, the acceptable risk level, and the regulatory context, which means transfer strategy cannot be separated from deployment policy [14], [24].

Another architectural consideration is uncertainty estimation. End-to-end models can produce overconfident predictions in unfamiliar conditions, which is dangerous when the system must decide whether to continue, degrade, or request human intervention. Uncertainty-aware architectures, ensembles, and out-of-distribution detectors can support runtime monitoring, but they add computational cost and complexity. They also require calibration data that reflects real-world conditions, not only simulation. Transfer learning should therefore include not just performance transfer but also uncertainty transfer, ensuring that confidence estimates remain meaningful when the input distribution shifts. This is a system requirement because it affects fallback behavior, human-machine interaction, and liability allocation [12], [24].

4. Data, Scenario, and Validation Infrastructure

Simulation-to-real transfer depends on infrastructure that is often invisible in model-centric papers. Scenario libraries, map formats, sensor models, traffic agent behavior models, and compute clusters constitute the substrate of transfer learning. CARLA and similar simulators provide open environments for training and benchmarking, while real-world datasets such as KITTI, nuScenes, and Waymo Open Dataset provide anchors for validation [2], [15], [16], [18]. Yet no single dataset or simulator can cover the full operational design domain. A robust infrastructure must support composition, versioning, and provenance across multiple sources, allowing researchers and engineers to trace which scenarios, sensors, and labels contributed to a model's behavior.

Scenario generation is particularly important for safety-critical transfer. Real-world driving is dominated by nominal conditions, while collisions and near-collisions are rare. Simulation can generate adversarial, edge-case, and counterfactual scenarios that would be unethical or impractical to collect on public roads. However, generated scenarios must be plausible and relevant. A physically impossible cut-in or an absurd pedestrian trajectory may stress the model in ways that do not predict real-world failure. Scenario generation therefore requires collaboration among traffic engineers, human factors specialists, ethicists, and machine learning practitioners. The value of a scenario library lies not only in its size but in its coverage of causal mechanisms that produce accidents and disengagements [14], [24].

Validation infrastructure must also be continuous rather than episodic. End-to-end models are updated frequently, and each update can alter behavior in unanticipated ways. Regression testing across simulation, closed-course, and public-road evaluations can detect performance drift, but only if the test suite is representative and stable. Digital twins of cities or highways can provide high-fidelity environments for repeated testing, yet maintaining digital twins is costly and requires updating maps, construction zones, and traffic patterns. Organizations must decide how much fidelity is necessary for different assurance claims. A model intended for low-speed campus shuttles may require different simulation infrastructure than a highway pilot system. These decisions are infrastructural and economic, not purely technical [13], [21].

Data governance is another component. Simulation data may be synthetic and thus less privacy-sensitive, but real-world fleet data often contain personal information, including faces, license plates, and location traces. Transfer learning pipelines that mix simulated and real data

must therefore manage consent, retention, de-identification, and access control. Model cards and datasheets for datasets provide useful documentation practices, but they must be adapted to the dynamic and distributed nature of autonomous vehicle fleets [19], [20]. Without such governance, transfer learning can become a mechanism for opaque data accumulation rather than accountable improvement.

5. Robustness, Safety, and Operational Assurance

Robustness in end-to-end driving refers to the model's ability to maintain safe performance under distribution shift, sensor degradation, adversarial interference, and rare environmental conditions. Simulation-to-real transfer can improve robustness by exposing the model to diverse conditions, but it can also create false confidence if the simulation distribution is mistaken for the real distribution. A safety case for an end-to-end vehicle must therefore articulate the relationship between simulated evidence and real-world claims. It must explain which hazards were tested, which were not, and how residual risk will be monitored after deployment. This is a form of argumentation that cannot be reduced to a single benchmark score [14], [24].

Operational assurance requires runtime monitoring and fallback mechanisms. Even a well-transferred model may encounter conditions outside its training distribution, such as unusual road geometry, emergency vehicles, or extreme weather. The system should be able to detect degradation, reduce speed, change lanes, or hand over control to a human operator. These behaviors depend on the architecture of the end-to-end model and on the surrounding vehicle platform. If the model is a monolithic neural network, it may be difficult to isolate the cause of failure or to certify a safe fallback. Hybrid architectures that combine learned components with rule-based safety monitors can provide stronger guarantees, but they reintroduce interface complexity. Transfer learning must be evaluated within this broader safety architecture [12], [24].

Safety assurance also involves human factors. A vehicle that transfers well in simulation may behave in ways that are technically safe but socially confusing. For example, it may stop too conservatively, accelerate too abruptly, or fail to communicate intent to pedestrians and cyclists. These behaviors affect acceptance and can create secondary risks. Simulation can model human reactions, but only if the behavior models are calibrated to real populations. Transfer learning should therefore include social and communicative dimensions, not only perception and control. This requires interdisciplinary validation that includes cognitive science, urban planning, and law [25].

Finally, robustness must be assessed across the lifecycle. A model that works at deployment may degrade as roads change, sensors age, and software dependencies evolve. Continuous learning from fleet data can address this, but it also introduces feedback loops and potential bias. If updates are driven by frequently encountered conditions, rare but severe scenarios may remain underrepresented. A lifecycle assurance approach should combine simulation-based stress testing, real-world monitoring, and periodic revalidation. The required infrastructure is substantial, which is why simulation-to-real transfer is not only an algorithmic strategy but also an organizational commitment [13], [21].

6. Governance, Fairness, and Policy Implications

Autonomous driving systems operate in public space, and their benefits and risks are distributed unevenly. Simulation-to-real transfer can exacerbate fairness concerns if simulators and datasets underrepresent certain neighborhoods, road types, pedestrians, or

cultural driving practices. A model trained primarily on high-income, well-mapped urban areas may perform poorly in rural areas, informal settlements, or regions with different signage and enforcement patterns. Fairness in this context is not only about demographic parity in predictions; it is about equitable access to safe mobility and equitable exposure to risk. Transfer learning pipelines should therefore include geographic, demographic, and infrastructural diversity as explicit design goals [19], [20].

Policy frameworks are beginning to address these issues. The National Highway Traffic Safety Administration has issued guidance on automated vehicle safety, and the European Union's Artificial Intelligence Act proposes risk-based obligations for high-risk AI systems, including those used in transportation [22], [23]. These frameworks emphasize documentation, human oversight, cybersecurity, and post-market monitoring. For simulation-to-real transfer, regulators may need to assess whether simulated evidence is sufficient for a safety claim and how companies validate models across operational design domains. This requires new forms of technical evidence, such as simulation credibility assessments, scenario coverage reports, and model change logs. It also requires regulators to develop expertise in machine learning without losing sight of traditional vehicle safety principles [21], [24].

Liability is another governance challenge. If an end-to-end model fails after transfer, it may be difficult to determine whether the failure originated in simulation assumptions, training data, real-world fine-tuning, sensor limitations, or runtime decision-making. This complexity can obscure accountability and slow compensation for victims. Governance mechanisms such as audit trails, versioned datasets, and third-party testing can help, but they must be designed before incidents occur. Model cards and datasheets are useful starting points, but they are not sufficient for dynamic systems that learn from fleets. A stronger regime would require pre-deployment safety cases, continuous monitoring reports, and independent authority to suspend operations when evidence of unacceptable risk emerges [19], [20], [23].

Fairness also intersects with labor and economic policy. Autonomous driving may displace professional drivers while creating new roles in data annotation, simulation engineering, and remote assistance. Simulation-to-real transfer changes the geography of these roles, concentrating high-skill development in some regions while distributing low-wage annotation work elsewhere. A socially sustainable transition would require attention to workforce development, data labor conditions, and equitable distribution of benefits. These concerns are not peripheral to transfer learning; they shape the public legitimacy of the technology and therefore its long-term deployment [22], [25].

7. Sustainability and Deployment Economics

Sustainability in simulation-to-real transfer has both environmental and economic dimensions. Training large end-to-end models consumes significant compute, and high-fidelity simulation at scale requires energy-intensive rendering and data processing. If simulation is used to reduce real-world testing, it may lower emissions from prototype fleets, but it may also shift energy consumption to data centers. The net environmental effect depends on the energy mix, model efficiency, and the number of real-world miles avoided. Lifecycle assessment should therefore accompany transfer learning research, especially as models grow larger and simulation environments become more detailed [13], [25].

Economic sustainability depends on the cost of maintaining simulation infrastructure, curating real-world datasets, and validating updates. Building a high-fidelity simulator is expensive, but so is collecting and labeling rare real-world events. Organizations must decide whether to

invest in proprietary simulation, share industry consortia, or rely on open-source platforms. Each choice has implications for competition, innovation, and safety. Shared simulation benchmarks can reduce duplication and improve comparability, but they can also entrench particular assumptions. A pluralistic ecosystem with multiple simulators and datasets may be more robust, provided that models can be evaluated across them [2], [16], [18].

Deployment economics also affect transfer strategy. A company with a large fleet can collect real-world data continuously and fine-tune models at scale, while a smaller operator may depend more heavily on simulation. This asymmetry can concentrate market power and shape safety outcomes. Regulatory requirements that demand extensive real-world validation may favor incumbents, while requirements that accept simulated evidence may lower barriers but raise questions about credibility. Policy should aim for proportionate evidence standards that account for operational design domain, vehicle capability, and risk exposure. The goal is not to choose between simulation and reality but to design a transfer ecosystem that is both innovative and accountable [14], [21], [23].

8. Conclusion

Simulation-to-real transfer learning is a foundational capability for end-to-end autonomous driving, but it cannot be treated as a purely technical optimization problem. The transfer gap is distributed across sensors, dynamics, semantics, institutions, and human behavior. Closing it requires architectural choices that balance end-to-end learning with structured representations, transfer mechanisms that combine randomization, adaptation, generation, and post-training, and infrastructure that supports scenario generation, validation, and lifecycle monitoring. It also requires governance frameworks that address fairness, liability, privacy, environmental impact, and public accountability. A recent post-training perspective reinforces that simulation pretraining alone is insufficient; targeted real-world alignment and continuous evaluation are necessary for robust deployment [17].

The central argument of this paper is that simulation-to-real transfer should be governed as an infrastructural capability rather than a proprietary trick. Shared benchmarks, traceable data provenance, independent safety cases, and transparent model documentation can make transfer claims auditable. Continuous monitoring and revalidation can detect degradation after deployment. Fairness and sustainability must be integrated from the beginning, not added after market entry. If these conditions are met, simulation-to-real transfer can help autonomous driving systems learn from rare events without exposing the public to unnecessary risk. If they are neglected, transfer learning may produce systems that appear capable in simulation but fail unpredictably in the complex social environments where they must operate.

References

1. Codevilla, F., Müller, M., López, A., Koltun, V., & Dosovitskiy, A. (2018). End-to-end driving via conditional imitation learning. In 2018 IEEE International Conference on Robotics and Automation (ICRA) (pp. 4693–4700). IEEE.
<https://doi.org/10.1109/ICRA.2018.8460487>
2. Dosovitskiy, A., Ros, G., Codevilla, F., López, A., & Koltun, V. (2017). CARLA: An open urban driving simulator. In Proceedings of the 1st Annual Conference on Robot Learning (pp. 1–16). PMLR.

3. Bojarski, M., Del Testa, D., Dworakowski, D., Firner, B., Flepp, B., Goyal, P., Jackel, L. D., Monfort, M., Muller, U., Zhang, J., Zhang, X., Zhao, J., & Zieba, K. (2016). End to end learning for self-driving cars. arXiv preprint arXiv:1604.07316.
4. Kendall, A., Hawke, J., Janz, D., Mazur, P., Reda, D., Allen, J. M., Lam, V. D., Qiao, K., & Shah, A. (2019). Learning to drive in a day. In 2019 International Conference on Robotics and Automation (ICRA) (pp. 8248–8254). IEEE.
<https://doi.org/10.1109/ICRA.2019.8793742>
5. Tobin, J., Fong, R., Ray, A., Schneider, J., Zaremba, W., & Abbeel, P. (2017). Domain randomization for transferring deep neural networks from simulation to the real world. In 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (pp. 23–30). IEEE. <https://doi.org/10.1109/IROS.2017.8202133>
6. Peng, X. B., Andrychowicz, M., Zaremba, W., & Abbeel, P. (2018). Sim-to-real transfer of robotic control with dynamics randomization. In 2018 IEEE International Conference on Robotics and Automation (ICRA) (pp. 3803–3810). IEEE.
<https://doi.org/10.1109/ICRA.2018.8460528>
7. Sadeghi, F., & Levine, S. (2017). CAD2RL: Real single-image flight without a single real image. In Robotics: Science and Systems XIII.
<https://doi.org/10.15607/RSS.2017.XIII.034>
8. Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., & Lempitsky, V. (2016). Domain-adversarial training of neural networks. *Journal of Machine Learning Research*, 17(59), 1–35.
9. Tzeng, E., Hoffman, J., Saenko, K., & Darrell, T. (2017). Adversarial discriminative domain adaptation. In 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 7167–7176). IEEE. <https://doi.org/10.1109/CVPR.2017.316>
10. Zhu, J. Y., Park, T., Isola, P., & Efros, A. A. (2017). Unpaired image-to-image translation using cycle-consistent adversarial networks. In 2017 IEEE International Conference on Computer Vision (ICCV) (pp. 2223–2232). IEEE.
<https://doi.org/10.1109/ICCV.2017.244>
11. Wang, Z., She, Q., & Ward, T. E. (2021). Generative adversarial networks in computer vision: A survey and taxonomy. *ACM Computing Surveys*, 54(2), Article 37.
<https://doi.org/10.1145/3439723>
12. Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., & Mané, D. (2016). Concrete problems in AI safety. arXiv preprint arXiv:1606.06565.
13. Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., Chaudhary, V., Young, M., Crespo, J. F., & Dennison, D. (2015). Hidden technical debt in machine learning systems. In *Advances in Neural Information Processing Systems 28* (pp. 2503–2511). Curran Associates.
14. Shalev-Shwartz, S., Shammah, S., & Shashua, A. (2017). On a formal model of safe and scalable self-driving cars. arXiv preprint arXiv:1708.06374.
15. Geiger, A., Lenz, P., Stiller, C., & Urtasun, R. (2013). Vision meets robotics: The KITTI dataset. *The International Journal of Robotics Research*, 32(11), 1231–1237.
<https://doi.org/10.1177/0278364913491297>

16. Caesar, H., Bankiti, V., Lang, A. H., Vora, S., Liong, V. E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., & Beijbom, O. (2020). nuScenes: A multimodal dataset for autonomous driving. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 11621–11631). IEEE. <https://doi.org/10.1109/CVPR42600.2020.01164>
17. Yang, R., Wang, M., Chen, Y., Guo, T., Xu, Y., Cui, C., ... & Su, L. (2026). Post-Training in End-to-End Autonomous Driving. arXiv preprint arXiv:2607.08072.
18. Sun, P., Kretzschmar, H., Dotiwalla, X., Chouard, A., Patnaik, V., Tsui, P., Guo, J., Zhou, Y., Chai, Y., Caine, B., Vasudevan, V., Han, W., Ngiam, J., Zhao, H., Timofeev, A., Ettinger, S., Krivokon, M., Gao, A., Joshi, A., Zhang, Y., Shlens, J., Chen, Z., & Anguelov, D. (2020). Scalability in perception for autonomous driving: Waymo Open Dataset. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 2446–2454). IEEE. <https://doi.org/10.1109/CVPR42600.2020.00252>
19. Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model cards for model reporting. In Proceedings of the Conference on Fairness, Accountability, and Transparency (pp. 220–229). ACM. <https://doi.org/10.1145/3287560.3287596>
20. Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Daumé III, H., & Crawford, K. (2021). Datasheets for datasets. *Communications of the ACM*, 64(12), 86–92. <https://doi.org/10.1145/3458723>
21. Shladover, S. E. (2018). Connected and automated vehicle systems: Introduction and overview. *Journal of Intelligent Transportation Systems*, 22(3), 190–200. <https://doi.org/10.1080/15472450.2017.1421166>
22. National Highway Traffic Safety Administration. (2016). Federal automated vehicles policy: Accelerating the next revolution in roadway safety. U.S. Department of Transportation.
23. European Commission. (2021). Proposal for a regulation of the European Parliament and of the Council on laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). COM/2021/206 final.
24. Koopman, P., & Wagner, M. (2017). Autonomous vehicle safety: An interdisciplinary challenge. *IEEE Intelligent Transportation Systems Magazine*, 9(1), 90–96. <https://doi.org/10.1109/MITS.2016.2583491>
25. Yurtsever, E., Lambert, J., Carballo, A., & Takeda, K. (2020). A survey of autonomous driving: Common practices and emerging technologies. *IEEE Access*, 8, 58443–58469. <https://doi.org/10.1109/ACCESS.2020.2983149>