

Continual Learning Strategies for Long-Term Autonomous Driving Model Evolution

Primitivo Esteban

Department of Computer Science, Binghamton University, Binghamton, NY, USA.
esteban323@binghamton.edu

Aleks Kolonko

Department of Electrical Engineering and Computer Science, University of Missouri,
Columbia, MO, USA.
kolonko1995@missouri.edu

Joacim Axelsson

School of Information Technology, University of Cincinnati, Cincinnati, OH, USA.
joacim.axelsson@uc.edu

Abstract

Long-term autonomous driving depends not only on the initial performance of perception, prediction, planning, and control models but also on their capacity to evolve safely after deployment. Continual learning offers a conceptual and infrastructural foundation for such evolution, allowing fleets to absorb new sensor characteristics, road geometries, traffic norms, weather patterns, and regulatory constraints without discarding prior competence. This paper examines continual learning strategies for long-term autonomous driving model evolution from a system-level perspective. It analyzes the structural trade-offs among stability, plasticity, latency, memory, compute, and safety assurance. It also considers the data and operational feedback loops that make continual adaptation possible, including fleet-scale collection, edge filtering, cloud training, validation, and staged rollout. Particular attention is given to catastrophic forgetting, distribution shift, robustness, fairness, governance, sustainability, and policy implications. The paper argues that continual learning in autonomous driving is not merely an algorithmic problem but a socio-technical infrastructure problem requiring coordinated choices in architecture, data governance, verification, and deployment. Successful long-term evolution will depend on hybrid modular designs, bounded end-to-end adaptation, auditable model lineage, human oversight, and regulatory frameworks that can accommodate incremental change without eroding public trust. The conclusion outlines a forward-looking agenda for resilient, equitable, and sustainable autonomous driving model evolution.

Keywords

continual learning, autonomous driving, model evolution, catastrophic forgetting, distribution shift, system architecture, safety governance, edge-cloud infrastructure, fairness, sustainability.

1. Introduction

Autonomous driving has moved from isolated demonstrations to sustained commercial experimentation, yet the long-term viability of the technology depends on how vehicle intelligence changes after initial deployment. Early autonomous systems such as Stanley and

subsequent urban driving platforms demonstrated that perception, localization, planning, and control could be integrated into functioning vehicles [1], [2]. Modern deep learning has expanded the perceptual and predictive capabilities of these systems, but it has also introduced dependence on large datasets, fragile generalization, and opaque failure modes [3]. In parallel, onboard computing has become a critical architectural constraint, shaping what can be updated, when, and under what assurance conditions [4]. The result is a field in which model evolution is no longer a secondary maintenance activity but a central operational capability.

Continual learning addresses the problem of learning from non-stationary data streams without catastrophic interference. Reviews of continual learning describe a tension between stability, which preserves prior knowledge, and plasticity, which enables adaptation to new situations [5], [20]. In autonomous driving, this tension has immediate safety consequences. A model that adapts too aggressively may forget rare but critical scenarios, such as unprotected left turns, emergency vehicle interactions, or construction-zone navigation. A model that adapts too slowly may fail to recognize new vehicle types, altered lane markings, seasonal road appearance, or revised traffic rules. The design challenge is therefore not to maximize adaptation in the abstract but to govern adaptation across a complex operational domain.

System-level continual learning for autonomous driving must also confront the fact that vehicles are safety-critical, geographically distributed, and subject to regulatory oversight. Unlike recommendation systems or language models, driving models can cause physical harm, and their failures are visible to the public. The infrastructure required for continual improvement therefore includes data pipelines, simulation environments, validation suites, over-the-air update mechanisms, rollback capabilities, and incident investigation processes. These elements connect technical choices to governance, liability, and public trust. As a result, the evolution of autonomous driving models should be understood as a problem of managed change rather than purely stochastic optimization.

This paper develops a long-form analysis of continual learning strategies for long-term autonomous driving model evolution. It argues that sustainable evolution depends on hybrid architectures, disciplined data governance, robust validation, and policy frameworks that treat model updates as regulated events. The discussion proceeds from architectural foundations to data infrastructure, robustness, governance, deployment, and policy. The central claim is that continual learning in autonomous driving is not merely a research technique but an infrastructural commitment that shapes how safety, fairness, and sustainability are realized over time.

2. Architectural Foundations and Continual Learning Trade-offs

The architecture of an autonomous driving stack determines what kinds of continual learning are practical. Modular systems separate perception, prediction, planning, and control, allowing localized updates to specific components. End-to-end systems map sensor inputs directly to driving actions or trajectories and may learn representations that are difficult to decompose [16], [17]. Between these extremes are hybrid designs that retain modular safety interfaces while using learned components for perception and behavior prediction. Each design carries different implications for continual learning. Modularity supports targeted updates, clearer validation, and interpretable rollback, but it can limit cross-component adaptation and create interface brittleness. End-to-end learning supports holistic adaptation

and potentially richer representations, but it complicates verification, attribution, and selective retention.

A central trade-off in continual learning is the management of catastrophic forgetting. Regularization-based methods such as elastic weight consolidation estimate the importance of parameters and penalize changes that would disrupt prior tasks [6]. Functional regularization approaches such as learning without forgetting preserve earlier outputs while learning new tasks [7]. Replay-based methods store or generate exemplars and interleave them with new data, as in incremental classifier and representation learning [8]. Parameter isolation and importance-based methods attempt to allocate capacity so that new learning does not overwrite critical older knowledge [9]. These strategies are often evaluated on classification benchmarks, but autonomous driving introduces continuous regression, temporal reasoning, multimodal sensing, and closed-loop control. The evaluation criteria must therefore include not only retention of old classes but also stability of trajectory quality, calibration, reaction time, and comfort.

Another architectural tension concerns where learning occurs. Onboard learning can adapt to local conditions but is constrained by compute, power, thermal limits, and safety certification. Cloud learning can exploit large datasets and powerful accelerators but depends on connectivity and delayed deployment. Fleet-scale learning can aggregate diverse experiences but raises privacy, ownership, and data-selection questions. A practical approach is to separate rapid local adaptation from slower fleet-level retraining. Local systems might adjust calibration or low-risk perception parameters, while cloud systems retrain larger models under controlled conditions. This separation mirrors broader machine learning production practice, where training and serving environments differ and where hidden technical debt accumulates when data dependencies, configuration, and feedback loops are poorly managed [11]. In autonomous driving, such debt can become a safety hazard if model lineage, data provenance, or validation assumptions are not maintained.

The choice of representation and memory also matters. Continual learning systems must decide what to store, what to summarize, and what to discard. Storing raw sensor data at fleet scale is expensive and privacy-sensitive. Storing embeddings, prototypes, or synthetic replay samples can reduce cost but may lose rare details. Time-series encodings and tiled representations have been used to make temporal patterns more amenable to visual or convolutional processing [14], but autonomous driving requires preserving causal order, sensor synchronization, and contextual metadata. A memory system that discards the wrong examples can induce bias toward common conditions and weaken performance in rare but safety-critical scenarios. Therefore, memory governance is not only a computational issue but also an epistemic one: it determines which experiences the fleet remembers and which it forgets.

3. Data, Infrastructure, and Operational Feedback Loops

Continual learning depends on data infrastructure that can turn vehicle operations into reliable learning signals. Autonomous fleets generate enormous volumes of camera, radar, lidar, inertial, and vehicle-state data. Most of this data is redundant, and uploading all of it is economically and environmentally unsustainable. Edge filtering, event detection, uncertainty estimation, and targeted sampling are therefore essential. The infrastructure must identify near-misses, disengagements, rare objects, sensor degradation, and atypical traffic behavior without introducing sampling biases that systematically exclude underrepresented regions or

populations. Data selection is thus an ethical and operational act, not merely a bandwidth optimization.

Operational feedback loops also create risks of feedback-induced distortion. If models influence what data is collected, and data determines future models, then deployment choices can narrow the fleet’s effective experience. For example, a vehicle that avoids certain neighborhoods may collect less data there, leading to weaker performance and further avoidance. A model that is overconfident in good weather may under-sample degraded visual conditions. To counter this, continual learning infrastructures need deliberate exploration, stratified data collection, and independent audit datasets that are not derived solely from current model behavior. The ML test score rubric and related production-readiness frameworks emphasize monitoring, data validation, and reproducibility [13]. In autonomous driving, these practices must extend to geographic coverage, demographic representation, and scenario diversity.

Cloud and edge infrastructure must also support versioning, reproducibility, and rollback. A model update is not simply a new weight file; it is a configuration of data, preprocessing, calibration, software dependencies, simulation parameters, and safety rules. If any of these elements change without traceability, incident investigation becomes difficult. A continual learning pipeline should maintain model lineage that records which data contributed to which update, what validation was performed, and under what conditions the update was deployed. This requirement connects to broader concerns in artificial intelligence safety, where specification, robustness, and unintended behavior must be addressed before and after deployment [12]. The infrastructure should treat every update as a controlled experiment with predefined success criteria and stop conditions.

Sustainability adds another dimension. Training and retraining large models consume energy, and frequent fleet-wide updates can multiply that cost [10]. A sustainable continual learning strategy may favor selective retraining, modular updates, distillation, quantization, and reuse of frozen representations. It may also schedule updates to avoid unnecessary computation when marginal performance gains are small. The environmental cost must be weighed against safety benefits, but the trade-off should be explicit. A system that continuously retrains everything is unlikely to be sustainable at scale. A system that never updates is unlikely to remain safe. The design target is therefore adaptive maintenance: enough change to address meaningful distribution shifts, but not so much that validation, energy, and operational complexity become unmanageable.

4. Robustness, Safety, and Validation under Distribution Shift

Autonomous driving models face distribution shift across seasons, weather, lighting, road surfaces, sensor aging, construction, and cultural driving norms. Continual learning is attractive because it offers a path to update models as these conditions change. However, adaptation can also introduce new vulnerabilities. An update that improves performance on a newly collected dataset may degrade performance on older but still relevant conditions. This is the stability-plasticity dilemma in operational form [5], [20]. Robustness validation must therefore evaluate both forward transfer, meaning improvement on new conditions, and backward transfer, meaning retention of prior competence. It must also test for negative transfer, where new learning harms unrelated capabilities.

Safety validation for continual updates requires more than aggregate accuracy. It requires scenario-based testing, simulation, closed-course evaluation, shadow mode deployment, and

staged rollout. Rare events, edge cases, and long-tail interactions must be represented, even if they are statistically uncommon. Reinforcement learning has demonstrated the ability to learn complex control policies in interactive environments [19], but its use in safety-critical driving requires constraints, shielding, and conservative exploration. End-to-end driving models, whether trained through imitation or conditional imitation, can learn useful mappings from perception to action [16], [17], [18], but they may lack explicit guarantees. Validation should therefore combine empirical testing with runtime monitors, fallback controllers, and uncertainty-aware arbitration.

Recent work on post-training in end-to-end autonomous driving highlights that pretrained models often require additional adaptation to align with downstream driving objectives, safety constraints, and deployment conditions [15]. This observation is important for continual learning because it suggests that model evolution is not a single training phase but a sequence of alignment, refinement, and evaluation stages. Post-training may improve behavioral fidelity, but it also introduces the risk of overfitting to narrow evaluation distributions or of weakening general capabilities. In a continual learning system, post-training must be embedded within a governance loop that checks whether adaptation preserves core safety invariants. The goal is not to reject post-training but to make it auditable, bounded, and reversible.

Validation must also address temporal stability. A model that performs well immediately after an update may drift as sensor calibration changes, as software dependencies evolve, or as the operational design domain expands. Continuous monitoring should track performance indicators such as disengagement rate, near-miss frequency, intervention severity, comfort metrics, and uncertainty calibration. When anomalies appear, the system should be able to isolate the cause, compare against previous versions, and roll back if necessary. This capability turns continual learning from an open-ended optimization process into a managed safety lifecycle. Without it, adaptation becomes a source of uncontrolled risk.

5. Governance, Sustainability, and Fairness in Model Evolution

Governance is inseparable from continual learning because model updates redistribute risk. A change that improves average safety may still harm a specific region, weather condition, vehicle type, or pedestrian group. Fairness in autonomous driving is often discussed in terms of equitable performance across neighborhoods, age groups, mobility patterns, and vulnerable road users. Continual learning can either reduce or amplify disparities. If data collection is skewed toward affluent or well-mapped areas, updates may improve those areas while neglecting others. If validation uses the same skewed data, disparities may remain invisible. Governance therefore requires disaggregated evaluation, independent audits, and mechanisms for affected communities to report concerns.

Policy frameworks for autonomous vehicles have traditionally focused on permits, testing, and incident reporting. Continual learning introduces the challenge of regulating change itself. Regulators may need to approve not only initial deployment but also categories of updates, thresholds for revalidation, and conditions under which an update requires new certification. A purely prescriptive approach may be too slow for rapidly changing software, while a purely performance-based approach may lack sufficient oversight. A hybrid model could combine pre-approved update envelopes, continuous monitoring, third-party audits, and mandatory reporting of safety-critical changes. This approach would treat model evolution as a lifecycle governed by evidence rather than a one-time approval.

Sustainability and governance also intersect through compute and data infrastructure. Large-scale training has environmental costs, and frequent updates can increase energy consumption across data centers and vehicle fleets [10]. A responsible continual learning strategy should consider carbon intensity, hardware efficiency, and the marginal safety value of additional training. It should also avoid planned obsolescence in vehicle hardware; if models grow too quickly for onboard accelerators, fleets may require premature replacement. Modular updates, model compression, and shared representations can extend hardware lifetimes and reduce electronic waste. Sustainability is therefore not an external constraint but an architectural design objective.

Finally, governance must address transparency and accountability. Users, regulators, and incident investigators need to understand what changed, why it changed, and what evidence supported the change. This does not require disclosing proprietary model weights, but it does require documenting data sources, validation coverage, known limitations, and update rationale. Such documentation supports public trust and enables independent scrutiny. In a continual learning regime, transparency is not a static report but a dynamic record. The governance system must be able to reconstruct the state of the model at any point in time and explain the operational decisions that followed from it.

6. Deployment Architectures and Cross-Domain Lessons

Deployment architecture determines how continual learning reaches the vehicle. A common pattern is cloud training followed by edge deployment, with simulation and staged rollout in between. More advanced patterns include federated learning, where vehicles contribute updates without sharing raw data, and hybrid architectures that combine centralized training with local calibration. Each pattern has trade-offs in privacy, bandwidth, latency, and control. Cross-domain lessons from cloud machine learning suggest that production systems accumulate technical debt when data dependencies, configuration, and feedback loops are not carefully managed [11]. Autonomous driving adds physical safety and regulatory constraints, making that debt more consequential.

Other domains offer instructive parallels. Industrial control systems use conservative update policies because unanticipated interactions can cause physical damage. Aerospace systems use rigorous verification and configuration management. Healthcare systems face fairness, privacy, and liability concerns that resemble those in autonomous mobility. In each case, continual learning must be embedded in an assurance case, not treated as an isolated algorithm. The automotive sector can borrow practices such as staged deployment, canary releases, rollback, red teaming, and independent validation. It must also adapt them to the unique temporal and spatial dynamics of driving.

From an architectural perspective, hybrid systems may offer the best near-term path. Learned perception and prediction modules can be updated more frequently, while planning and control retain formal or rule-based safety layers. End-to-end components can be trained and post-trained within bounded operational design domains [15], but they should be monitored by independent safety envelopes. This layered approach allows continual improvement without exposing the entire stack to unverified change. It also supports modular certification, where updates to one component trigger targeted revalidation rather than full-system recertification.

Cross-domain comparisons also highlight the importance of human oversight. In aviation, automation is paired with trained pilots and clear escalation procedures. In autonomous

driving, remote assistance, fleet operations, and incident response teams play analogous roles. Continual learning should be designed to support these human actors by providing explanations, uncertainty estimates, and actionable diagnostics. A model that evolves without interpretable signals may improve on benchmarks while making human oversight harder. The deployment architecture must therefore treat human-machine teaming as a first-class requirement.

7. Policy Implications and Forward-Looking Perspectives

Policy for continual learning in autonomous driving must balance innovation with precaution. Overly rigid rules could freeze models at a level of performance that is less safe than a well-governed update process. Overly permissive rules could allow rapid changes that outpace validation and public understanding. A workable framework might define risk tiers, where low-risk updates such as calibration improvements follow streamlined review, while high-risk changes to perception or planning require stronger evidence. It might also require continuous monitoring and mandatory reporting of near-misses and disengagements. Such a framework would recognize that safety is a dynamic property, not a static certification.

International coordination will be important because vehicles and data cross borders. Different regions have different privacy laws, road rules, and liability regimes. A continual learning system that adapts to one jurisdiction may behave differently in another. Policy should therefore support interoperability in safety documentation, incident reporting, and update auditing. Standards bodies can help by defining common formats for model lineage, validation evidence, and operational design domain declarations. These standards would reduce compliance costs while improving transparency.

A forward-looking research agenda should address several open problems. One is the development of continual learning methods that provide formal or statistical guarantees for safety-critical regression and control. Another is the design of memory systems that preserve rare events without violating privacy or inflating costs. A third is the creation of validation environments that can simulate long-tail scenarios and adversarial conditions at scale. A fourth is the integration of fairness metrics into continual learning objectives and governance. A fifth is the study of human factors in fleet oversight, including how operators understand and trust evolving models. These problems are interdependent and require collaboration among computer scientists, engineers, social scientists, policymakers, and communities.

The long-term vision is not a model that learns endlessly without constraint. It is a socio-technical system that learns deliberately, documents its changes, validates its adaptations, and remains accountable to the public. Continual learning strategies should therefore be evaluated by their contribution to safe, fair, and sustainable mobility, not only by benchmark improvement. The measure of success is whether the fleet becomes more capable over time without becoming less trustworthy.

8. Conclusion

Continual learning for long-term autonomous driving model evolution is a system-level challenge that spans architecture, data infrastructure, validation, governance, and policy. The stability-plasticity dilemma is not merely an algorithmic trade-off; it determines whether vehicles can adapt to new conditions while preserving rare but critical competencies. Modular and hybrid architectures, bounded end-to-end adaptation, disciplined memory management, and staged deployment offer practical pathways. Fleet-scale data collection must be paired with stratified sampling, privacy protection, and independent audit datasets to avoid feedback-

induced bias. Validation must assess forward and backward transfer, robustness, calibration, and temporal stability, not just average accuracy.

Governance must treat updates as regulated events with traceable lineage, disaggregated evaluation, and rollback capability. Sustainability requires attention to compute, energy, and hardware lifetime, while fairness requires equitable performance across regions and populations. Policy frameworks should enable responsible evolution through risk-tiered review, continuous monitoring, and international standards. The central conclusion is that continual learning in autonomous driving is not a single technique but an infrastructural commitment. It succeeds only when technical adaptation is embedded in a broader system of accountability, human oversight, and public trust. Future research should therefore pursue methods that are not only adaptive but also auditable, equitable, and sustainable over the full lifecycle of autonomous mobility.

References

1. Thrun, S., Montemerlo, M., Dahlkamp, H., Stavens, D., Aron, A., Diebel, J., ... & Mahoney, P. (2006). Stanley: The robot that won the DARPA Grand Challenge. *Journal of Field Robotics*, 23(9), 661–692. <https://doi.org/10.1002/rob.20147>
2. Levinson, J., Askeland, J., Becker, J., Dolson, J., Held, D., Kammel, S., ... & Thrun, S. (2011). Towards fully autonomous driving: Systems and algorithms. *IEEE Intelligent Vehicles Symposium*, 163–168. <https://doi.org/10.1109/IVS.2011.5940562>
3. Grigorescu, S., Trasnea, B., Cocias, T., & Macesanu, G. (2020). A survey of deep learning techniques for autonomous driving. *Journal of Field Robotics*, 37(3), 362–386. <https://doi.org/10.1002/rob.21918>
4. Liu, S., Tang, J., Zhang, Z., & Gaudiot, J. L. (2020). Computer architectures for autonomous driving. *Computer*, 53(10), 54–65. <https://doi.org/10.1109/MC.2020.3005618>
5. Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., & Wermter, S. (2019). Continual lifelong learning with neural networks: A review. *Neural Networks*, 113, 54–71. <https://doi.org/10.1016/j.neunet.2019.02.006>
6. Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., ... & Hadsell, R. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13), 3521–3526. <https://doi.org/10.1073/pnas.1611835114>
7. Li, Z., & Hoiem, D. (2018). Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12), 2935–2947. <https://doi.org/10.1109/TPAMI.2017.2773081>
8. Rebuffi, S. A., Kolesnikov, A., Sperl, G., & Lampert, C. H. (2017). iCaRL: Incremental classifier and representation learning. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2001–2010. <https://doi.org/10.1109/CVPR.2017.587>
9. Aljundi, R., Babiloni, F., Elhoseiny, M., Rohrbach, M., & Tuytelaars, T. (2018). Memory aware synapses: Learning what (not) to forget. *Proceedings of the European Conference on Computer Vision*, 139–154. https://doi.org/10.1007/978-3-030-01219-9_9

10. Rolnick, D., Donti, P. L., Kaack, L. H., Kochanski, K., Lacoste, A., Sankaran, K., ... & Bengio, Y. (2022). Tackling climate change with machine learning. *ACM Computing Surveys*, 55(2), 1–96. <https://doi.org/10.1145/3485128>
11. Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., ... & Dennison, D. (2015). Hidden technical debt in machine learning systems. *Advances in Neural Information Processing Systems*, 28, 2503–2511.
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. Breck, E., Cai, S., Nielsen, E., Salib, M., & Sculley, D. (2017). The ML test score: A rubric for ML production readiness and technical debt reduction. *IEEE International Conference on Big Data*, 1123–1132. <https://doi.org/10.1109/BigData.2017.8258038>
14. Wang, Z., & Oates, T. (2015). Encoding time series as images for visual inspection and classification using tiled convolutional neural networks. *Workshops at the Twenty-Ninth AAAI Conference on Artificial Intelligence*.
15. 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*.
16. Bojarski, M., Del Testa, D., Dworakowski, D., Firner, B., Flepp, B., Goyal, P., ... & Zieba, K. (2016). End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*.
17. Codevilla, F., Müller, M., López, A., Koltun, V., & Dosovitskiy, A. (2018). End-to-end driving via conditional imitation learning. *IEEE International Conference on Robotics and Automation*, 4693–4700. <https://doi.org/10.1109/ICRA.2018.8460487>
18. Chen, C., Seff, A., Kornhauser, A., & Xiao, J. (2015). DeepDriving: Learning affordance for direct perception in autonomous driving. *Proceedings of the IEEE International Conference on Computer Vision*, 2722–2730. <https://doi.org/10.1109/ICCV.2015.312>
19. 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>
20. De Lange, M., Aljundi, R., Masana, M., Parisot, S., Jia, X., Leonardis, A., Slabaugh, G., & Tuytelaars, T. (2022). A continual learning survey: Defying forgetting in classification tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(7), 3366–3385. <https://doi.org/10.1109/TPAMI.2021.3057446>