

PHYSICS-INFORMED REINFORCEMENT LEARNING FOR AUTONOMOUS VEHICLES: BRIDGING THE SIMULATION- REALITY GAP

Singh Atul Sudhakar

Pacific Academy of Higher Education and Research University, Udaipur, India

Faculty of Computer Science

Abstract- This review analyses a new theory using physics-informed reinforcement learning (PI-RL) to tackle ongoing issues when moving simulations to real-world use in self-driving cars. Instead of relying solely on data-driven methods, it combines Physics-Informed Neural Networks (PINNs) with policy gradient techniques, building core physical rules - such as Newton's laws, friction behavior, motion limits, and energy preservation - directly into the learning process^[1]. Unlike standard domain randomization strategies, this approach uses known physics as fixed structure rather than letting such patterns emerge indirectly^[2]. It introduces several elements: an overview of common breakdown points, equations for physics-guided policy updates, layered safety systems, and structured testing designs. Still, no experiments were run - the study stays purely conceptual - and demands high computing power, slowing training by 30–50%. There is also tension between accurate modeling and practical optimization^[3]. Despite these barriers, clear limitation disclosures and thorough method planning offer solid starting points for later

hands-on research; however, actual implementation remains far off. This assessment explores underlying concepts, highlights missing components, compares suggested tools to existing top-tier solutions, and suggests steps toward experimental confirmation across various driving conditions.

Keywords- Physics-Informed Neural Networks, Reinforcement Learning, Autonomous Vehicles, Sim-To-Real Transfer, Control Barrier Functions, Vehicle Dynamics, Safe Learning, Policy Optimisation, Sensor Fusion, Domain Randomisation

I. INTRODUCTION

Contextualising the Simulation-Reality Challenge

The gap between simulated and real-world performance is a major hurdle in today's self-driving car research, often leading to sudden system breakdowns when models trained in digital settings face actual driving conditions ^[4]. Instead of blaming poor graphics or visuals, the study highlights hidden flaws in how physics are modeled - differences that pop up unexpectedly, shifting attention away from

testing every possible situation toward building core physical rules into learning systems. A tragic crash involving an Uber test vehicle in 2018 illustrates this well: despite heavy simulation use, safety wasn't ensured because rare events revealed weak approximations in the training setup [^5].

Common fixes like randomizing environments, creating hard edge cases, or scaling task difficulty help slightly - but they react after problems arise instead of preventing them by embedding realistic dynamics upfront. Here lies the key idea: stop chasing hyper-realistic simulations; start shaping algorithms that bake in known laws of motion directly, much like CNNs build in spatial consistency via their structure[^6]. This shift matters now more than ever, as large-scale AI models show strong perception skills but still fall short in safe deployment due to shaky grounding in physical reality[^7]

The combination of PINNs - first made for solving differential equations - with policy gradient methods is new, yet untested in practice. Although the study admits its limits clearly: there's no real-world testing, no broad review of existing work, nor promises of ready-to-use systems. This openness sets it apart from exaggerated claims often seen in self-driving vehicle studies, offering balanced views on a concept needing thorough experimentation.

II. LITERATURE REVIEW: POSITIONING PHYSICS- INFORMED APPROACHES WITHIN CONTEMPORARY RESEARCH

Reinforcement Learning for Autonomous Driving: Evolution and Persistent Challenges

Modern self-driving car controls have grown more complex in their algorithms, shifting from step-by-step decisions to smooth systems that turn sensor data straight into actions [^8]. The paper places this shift correctly within broader trends, citing wide-ranging studies by Kiran et al. (2022) and Chen et al. (2023), while pointing out ongoing issues - like needing fewer training examples, ensuring safety, or moving models from simulation to real-world use - that aren't solved just with better code[^9]. Instead of "and", these problems often interact; for instance, low sample efficiency worsens deployment delays because testing takes longer. Soft Actor-Critic (SAC), a type of policy gradient method using entropy tuning, performs better than PPO in learning quickly from limited experience. Still, findings by Yang and Zhao (2023) show even strong methods fail when real vehicle dynamics don't match simulated ones - highlighting limits of algorithm choice alone.

Model-based reinforcement learning methods - like those in Chua et al. (2018), using probabilistic dynamics models - gain much higher sample efficiency by predicting future

states and refining action sequences. Still, such trained models often miss key real-world physics when facing rare events; that is exactly when correct modeling becomes vital for safe operation^[10]. Because of this shortcoming, the discussed study shifts focus: instead of extracting physical rules from observations, it builds established laws directly into the optimization framework

While integrated learning methods cut out separate modules - boosting speed and efficiency - recent work by Hu et al. (2023) shows gains in planning tasks, whereas Wu et al. (2024) highlight improvements using GenAD for forecasting traffic scenes^[11]. However, Shao et al. (2023) find such combined designs can fail under rare conditions, producing unrealistic actions that rule-based or physics-checked systems would quickly filter out^[12]. Because of this conflict - streamlined design versus real-world feasibility - hybrid techniques emerge, blending data-driven training with built-in physical rules without breaking gradient flow

Physics-Informed Neural Networks: Theoretical Foundations and Domain Applications

Physics-informed neural networks changed how science uses computers - by building key equations right into the training process. As a result, they need less data while getting better at predicting outside their original scope^[1].

Instead of just fitting patterns, these models balance errors in both data and equation rules. Raissi et al. (2019) showed such networks can tackle complex math problems, including those where solutions aren't fully known upfront. Their method mixed measurement mismatches with deviations from physical laws. The paper under review takes this idea but shifts focus: not to solve equations, yet to guide decision-making policies. That shift matters - it moves the goal from simulating systems to shaping actions within real-world limits.

Cuomo et al. (2022) examined structural issues in physics-guided models, revealing distinct gradient problems - such as sharp differential equations that disrupt training or mixed-scale effects creating clashing gradient inputs^[13]. Instead of fixed weights, Wang et al. (2021) introduced dynamic adjustments to align data-based and physics-based losses; this approach could support the friction, motion, power, and movement rules outlined in the framework under review^[3]. However, the paper overlooks these practical hurdles during its method description, leaving a key shortfall needing real-world testing.

Domain-specific PINNs show strong results in scientific computing. Instead of relying only on data, Sharma et al. (2023) found that embedding Navier-Stokes equations as soft constraints improves flow predictions - even with little training input^[14]. In their review, Cai et al. (2021) noted better accuracy and

physics alignment versus pure machine learning models; some cases saw a tenfold gain in data efficiency for turbulence modeling^[15]. Meanwhile, Yuan et al. (2024) introduced loss attention and adaptive weighting - methods later used here for multi-condition optimization in vehicle control - but these have not yet been tested in decision-making systems^[16]

Recent advances in PINN training methodologies address fundamental optimisation challenges. Wang et al.'s (2023) expert training guides emphasise causal training sequences respecting temporal dependencies, careful initialisation strategies preventing gradient vanishing, and adaptive learning rate scheduling accommodating multi-scale physics^[17]. McClenny and Braganeto's (2023) self-adaptive PINNs automatically adjust sampling strategies to focus computational resources on regions where physics violations are most severe, achieving 40% faster convergence compared to uniform sampling in benchmark problems^[18]. These methodological advances directly inform the reviewed paper's proposed physics-regularised RL training pipeline, though specific adaptation strategies for policy optimisation contexts require investigation.

Safe Reinforcement Learning: Control Barrier Functions and Integrated Approaches

Safety in reinforcement learning has shifted from checking constraints after training to built-in methods that maintain safety throughout operation. Instead of adding rules later, current approaches ensure systems stay within safe bounds at all times. Ames et al. (2019) laid solid groundwork for Control Barrier Functions - tools that use mathematical conditions similar to Lyapunov stability to keep behavior safe. These functions act like virtual fences guiding system dynamics away from danger. Cheng et al. (2019) applied such barriers directly in learning tasks involving continuous actions, showing how neural network controllers can obey safety limits reliably. Their method trains agents while guaranteeing real-time adherence to predefined boundaries - an alternative to physics-based models being discussed here - but it depends heavily on precise model estimates and clear definitions of what counts as "safe"^[20].

Zhao et al. (2023) merge control Lyapunov functions with barrier functions inside actor-critic setups - ensuring both stability and safety, while improving performance through joint optimization^[21]. Still, their method depends on real-time quadratic programming every step, leading to delays of 15–30ms per iteration. In contrast, the technique in the reviewed study evaluates physical constraints faster, taking only 2–5ms. Such a gap implies that embedding physics directly into learning could yield lighter, more efficient safety layers.

Yet, solid comparisons between CBF-driven safe reinforcement learning and physics-guided variants are still missing, making firm judgments premature.

Emam et al. (2022) use strong control barrier functions to handle model errors by assuming limited disturbances along with cautious safety limits; the analysed study adopts this idea via statistical friction estimates and dynamic lambda adjustment^[22]. Combining uncertainty analysis with physical laws strengthens the method's foundation - however, applying these uncertainties across detailed vehicle models demands precise computational methods, which the evaluated research does not thoroughly cover. In contrast, Taylor et al. (2020) show that barrier parameters can come from real-world data without losing safety assurances, opening paths toward mixed strategies where learned barriers work together with rule-based controls^[23].

Identified Research Gap and Theoretical Positioning

The literature review shows a clear shortfall: although there's broad work in reinforcement learning for self-driving systems, physics-based models, and safety-aware training, none combine these into one cohesive approach - specifically, physics-guided RL tailored for vehicle control. Instead of merging them, current techniques handle physical rules after

policy generation, using them only as checks; meanwhile, physics-informed networks mostly solve equations, not policies within dynamic environments. The study tackles this issue by introducing both conceptual and structural designs where physical laws are built directly into the learning process - as smooth, optimisable terms in the loss function, not standalone filters

This setup brings some conceptual benefits: first, physical rules shape exploration from the start instead of just rejecting actions later. Because gradients pass through real-world laws, systems can learn realistic behaviors more consistently. Clear constraint definitions also improve transparency, making outcomes easier to evaluate by oversight bodies. Additionally, built-in knowledge structures cut the need for testing endless situations. Still, without experimental proof, it's unclear if these ideas actually boost real-world results. Combining them is a fresh approach that deserves testing - provided issues like computational load, accuracy limits, and tuning difficulties are fully addressed, as detailed in the paper's evaluation part

III. THEORETICAL FRAMEWORK ANALYSIS: MATHEMATICAL RIGOUR AND PRACTICAL CONSIDERATIONS

Physics-Regularised Objective Function:
Conceptual Foundation

The study's main idea adjusts the usual reinforcement learning goal by adding penalties for unrealistic movements based on physics. Instead of treating physical rules separately, they're built directly into the training process using adjusted weights for dynamics, friction, energy, and motion limits. This setup allows errors to propagate back through physical equations during updates. As a result, the model learns smoother behaviors that follow real-world rules without losing effectiveness on assigned tasks.

The augmented objective takes the form:

$$\begin{aligned} \mathbf{L}_{\text{total}} = & \mathbf{L}_{\text{RL}} + \lambda_{\text{dynamics}} \cdot \\ & \|\mathbf{F}_{\text{dynamics}}\|^2 + \lambda_{\text{friction}} \cdot \|\mathbf{F}_{\text{friction}}\|^2 + \\ & \lambda_{\text{energy}} \cdot \|\mathbf{F}_{\text{energy}}\|^2 + \lambda_{\text{kinematic}} \cdot \\ & \|\mathbf{F}_{\text{kinematic}}\|^2 \end{aligned}$$

where \mathbf{L}_{RL} represents standard policy gradient loss (PPO, SAC, or TD3), \mathbf{F} terms represent physics constraint residuals, and λ coefficients control relative importance. This approach uses physics limits as flexible guidelines instead of strict conditions, recognising actual vehicle behaviour is more complex than models can fully capture. The λ parameter adjusts balance - larger settings push closer adherence if predictions are confident; smaller ones allow looser alignment under high uncertainty.

Still, some math issues need close look. Because friction limits use subdifferentials for $\max(0, \cdot)$, trouble can arise when switching

between bound and free states - this might disrupt solver performance. Even though the authors add probability adjustments to handle such shifts, they don't fully explore how these affect optimization behavior. As seen in Wang et al. (2021), PINNs face stiff gradients; likewise, combining multiple constraints here could create comparable problems - especially if timeframes differ sharply, like quick tire responses (ms) against gradual heat changes (sec-min).

The chain rule's use in gradient backpropagation depends on smoothness across the entire computation path. When considering the friction constraint $\mathbf{F}_{\text{friction}}$ in relation to policy variables θ :

$$\begin{aligned} \partial \mathbf{F}_{\text{friction}} / \partial \theta = & \partial \mathbf{F}_{\text{friction}} / \partial \mathbf{a}_{\text{lateral}} \cdot \\ & \partial \mathbf{a}_{\text{lateral}} / \partial \mathbf{a} \cdot \partial \mathbf{a} / \partial \theta \\ & + \partial \mathbf{F}_{\text{friction}} / \partial \mathbf{a}_{\text{longitudinal}} \cdot \\ & \partial \mathbf{a}_{\text{longitudinal}} / \partial \mathbf{a} \cdot \partial \mathbf{a} / \partial \theta \end{aligned}$$

Every term needs a clear definition alongside numerical reliability. Where $\max(0, \cdot)$ is involved, its subdifferential creates jumps; autodiff systems manage these via chosen subgradients, which might lead to shaky optimization close to limits. Using smooth substitutes - like softplus or ELUs - may reduce instability, however such adjustments weaken exact constraint adherence, a compromise the analysed study doesn't directly examine

Detailed Physics Constraint Formulations: Scope and Limitations

The four types of constraints - related to motion (Newton's second law), grip (friction circle), energy balance (conservation principle), and steering design (Ackermann geometry) - together capture core aspects of vehicle behavior. Since Newton's second law links forces to changes in speed, its mismatch signal reflects how well real-world data aligns with basic physics.

$$F_{\text{dynamics}} = |m \cdot a_{\text{longitudinal}} - (F_{\text{traction}} - F_{\text{drag}} - F_{\text{rolling}})|^2 + |m \cdot a_{\text{lateral}} - F_{\text{cornering}}|^2 + |I_z \cdot \alpha_{\text{yaw}} - \tau_{\text{steering}}|^2$$

This model treats the vehicle as a rigid body with fixed mass and rotational inertia I_z , ignoring how weight shifts between axles when braking while turning. Under hard braking, more load moves forward - raising front wheel grip by 20–40% - a key effect absent in this basic approach.

The friction circle idea clearly shows how tire forces are limited by this equation:

$$F_{\text{lateral}}^2 + F_{\text{longitudinal}}^2 \leq (\mu \cdot N)^2$$

where μ stands for the grip level between tyre and road, while N is the vertical force. Instead of capturing quick changes - like delay in side-slip response or lag during cornering - it uses a steady-state simplification, which fails under sudden manoeuvres when safety is critical.

Models like Magic Formula, or other nonlinear versions, offer better accuracy; however, they require more computing power and rely on multiple hard-to-measure values such as maximum friction, curve shapes, and how load affects traction

Energy limits stop impossible speed-ups that go beyond available power

$$F_{\text{energy}} = |dE_{\text{total}}/dt - (P_{\text{engine}} - P_{\text{drag}} - P_{\text{rolling}} - P_{\text{brake}})|^2$$

where E_{total} includes kinetic, potential, and lost energy parts. Such models need precise engine power data - yet these depend on RPM, throttle setting, and intake pressure in highly nonlinear ways. Aerodynamic drag values (C_d) often lack accuracy due to real-world variability influenced by temperature shifts. Rolling resistance inputs show similar unpredictability under changing environmental conditions. Detailed performance maps are seldom accessible at required resolution levels for digital simulation environments.

The motion rule that maintains proper Ackermann steering behaviour:

$$F_{\text{kinematic}} = |v_{\text{lateral}} - v_{\text{longitudinal}} \cdot \tan(\delta)|^2$$

stands as the simplest geometric form, based on the rule that the turning center must align with the rear axle line. Yet this basic bike setup overlooks how suspensions move - like shifts in wheel tilt or direction - and ignores flex in

parts such as bushings or the steering shaft; it also leaves out tire slippage angles, which strongly affect real-world performance at driving limits. Tools using multi-part physics models can include these factors, though they require many inputs and demand heavy computing power

Model Uncertainty and Approximation Error: Honest Assessment

The study handles model uncertainty better than most overly positive tech papers. It openly states that actual vehicle behaviour is more complex than any simple mathematical model can capture - showing clarity about limits. Instead of treating physical laws as strict boundaries, it suggests using them flexibly, which reflects a realistic approach. By including uncertain friction values in its probability framework, the method offers a solid way to measure uncertainty

$$F_{\text{friction}} = \max(0, (a_{\text{lateral}}^2 + a_{\text{longitudinal}}^2) - (\hat{\mu} \cdot g \cdot \cos(\theta))^2 - k \cdot \sigma_{\mu})$$

Here, $\hat{\mu}$ stands for the estimated friction coefficient, σ_{μ} indicates the standard deviation of uncertainty, while k sets the degree of caution - often set to 3 for three-sigma limits. This approach turns strict physical rules into probabilistic guidelines that suggest direction instead of enforcing rigid limits. Yet, how to estimate σ_{μ} accurately across changing environmental conditions is still not clearly defined.

The debate over computation speed versus model accuracy reveals a real challenge: basic models - like fixed friction or bike-style movement - allow faster calculations yet overlook key behaviors; advanced setups - including precise tyre effects, flexible suspensions, and air resistance interactions - better reflect reality but bring calculation hiccups and unclear parameter values. This study suggests a practical approach - "begin with stripped-down versions and include extra details only when actual tests show mismatches between simulated and real outcomes" - recognising the conflict without giving strict rules for engineers."

The management of gradient issues via adaptive weight methods draws on McClenny and Braga-Neto's (2023) approach to self-adjusting PINNs, yet practical steps are still unclear. Because vehicle systems operate at different scales - tire deformation in microseconds, body motion in milliseconds, heat changes over seconds - gradient conflicts may arise, which standard auto-differentiation tools might not resolve effectively. Instead of fixed weights, using loss-guided strategies that adjust constraint importance according to real-time gradient size could help; these would raise λ when constraints aren't met enough, while lowering it where they're already satisfied.

IV. SYSTEM ARCHITECTURE AND IMPLEMENTATION: BRIDGING THEORY TO PRACTICE

Multi-Layered Safety Architecture: Defence-in-Depth Principles

The suggested three-level safety setup uses a physics-guided policy as main control, Model Predictive Control for real-time checks, while fast emergency responses act last - this mirrors layered protection suited to high-risk applications. The structure accepts that trained behaviours - even when grounded in physical laws - can't ensure full safety on their own, so separate oversight layers are added, each working across distinct time frames and decision hierarchies.

Layer 1 (Physics-Guided Control): Runs at 10–20 Hz, using just 5–10 ms per cycle. Built via a physics-aware reinforcement learning method, it manages routine driving when enough processing power is available for detailed physical modeling. Instead of raw data, it uses vehicle speed, acceleration, surrounding cars' locations, and road shape as input - then generates smooth throttle, braking, and steering commands. Constraints from real-world physics are directly built into its design.

Layer 2 (MPC Verification): Runs at 5–10 Hz, using around 10–20 ms per cycle. Instead of fixed paths, it computes optimal trajectories over short time windows - typically 1 to 3 seconds - forecasting how states evolve given possible inputs. Whenever these forecasts

suggest a breach of safety rules - like collisions, exceeding grip limits, or losing balance - it adjusts or replaces commands from Layer 1. Rather than just tracking goals, the MPC setup handles restrictions tied to vehicle dynamics while reducing errors relative to desired motion.

Layer 3 (Handles emergency reactions): works at 50–100+ Hz with less than 1ms processing time. Uses fixed rules to respond to sudden threats - like possible crashes found through time-to-contact math, loss of balance detected by tracking spin speed, or broken sensors spotted via anomaly checks. When triggered, it skips any learned models and runs preset actions such as full brake engagement, slide correction routines, or safe shutdown sequences. Thresholds for activation are tuned carefully - to avoid unnecessary triggers while still reacting fast enough in real crises. Timing allocations per component - around 5–10ms for policy decisions, 10–20ms for model-based planning, under 1ms for reflexes - are feasible on car-ready hardware; however, total delay near 70 μ s may push limits where smoother operation needs faster cycles (above 20Hz ideal). Choosing to embed physical laws directly into the main controller - not just using them during training - as opposed to adding safety layers after learning - is different from standard RL safety techniques. This setup helps keep exploration within realistic

boundaries early on, potentially cutting risky trial behaviors before they occur.

Yet a three-level setup can make systems harder to certify. Instead, regulators need to check the trained decision logic along with how levels interact. When shifting from one level to another, communication rules matter too. Problems may arise if layers conflict - say Level 1 suggests a move, Level 2 blocks it, yet its own choice turns out worse. These breakdowns in hierarchy demand structured review. Methods like fault trees or FMEA help assess risks. Also, broad simulations under varied malfunction cases are essential

The suggested training plan uses high-quality simulations first, combining solid physics rules with varied environments that still obey physical laws - then adds minimal real-world adjustments using a flexible lambda schedule. This step-by-step method follows current ideas in transfer learning - but also includes unique physics-aware elements missing in typical domain-adaptation techniques.

Stage 1 (High-Fidelity Simulation Pre-Training): Training lasts 1–2 million steps - around 5.5 to 11 hours real time, accounting for 70% efficiency. The setup runs on IPG CarMaker, rFactor Pro, or CARLA, each extended with advanced physics modules. Settings prioritise physical consistency, using regularisation weights: ($\lambda_{\text{dynamics}}$ at 1.0, $\lambda_{\text{friction}}$ set to 1.0, λ_{energy} at 0.5, and

$\lambda_{\text{kinematic}}$ fixed at 0.8.) Scenarios cover highway lane changes under differing traffic loads and relative speeds; urban intersections involving shielded or unshielded turns plus crossing pedestrians; crisis actions like dodging obstacles or hard stops; also difficult driving states such as rain-slick roads, ice patches, side gusts, and fog.

Task success means finishing at least 95% of tasks, keeping physics errors under 1%, while maintaining steady learning progress. A small number of failures are expected - like collisions during impossible starts - so 95% is considered full performance. Staying below 1% rule breaks shows the model mostly follows physical laws. Instead of strict perfection, this allows minor exceptions due to extreme cases. Learning trends are checked for flat lines, sudden shifts, or repeating swings. These signs help decide changes in settings like step size, penalty values, or network scale. Patterns over time guide tuning more than single results do.

Stage 2 (Domain Randomisation with Physics Constraints): Training lasts 500,000 to 1 million steps - roughly 3.3–6.7 hours real-time. The setup keeps physics active using moderate λ settings: $\lambda_{\text{dynamics}}$ at 0.8, $\lambda_{\text{friction}}$ also 0.8. Instead of altering core mechanics, it changes visuals like car types, surface textures, and material properties. Weather varies too, including rain strength, fog levels, and time of day. Traffic gets adjusted in

terms of volume, driver behavior, and vehicle mix. Lighting shifts cover dawn or dusk scenarios, shadow angles, and glare effects. Importantly, keeping basic physical laws fixed sets this apart from typical domain randomization methods - which often tweak mass, inertia, or friction beyond realistic bounds.

This approach uses controlled randomness to improve image-based stability while avoiding unrealistic motion effects. Still, it needs simulation tools that allow distinct adjustments for visuals and physical rules - a feature missing in some systems. For example, CARLA lets users set graphics and physics separately, whereas programs such as CarMaker might need tailored changes. Success is measured similarly to Phase 1: consistent results under different visual settings, stable physical behavior, plus confirmation that outcomes remain close to original performance levels.

Stage 3 (Real-World Fine-Tuning): Driving period depends on how fast the system adapts and safety needs - usually between 1,000 and 10,000 km with human oversight. The setup uses a changing λ schedule: keeps early settings at first, then slowly lowers them (say, cutting by 10% per 1,000 km) so the model adjusts to real-world differences from simulation. A person stays ready to take control during tuning, while full data tracking

records states, actions, broken constraints, and any manual overrides.

The adaptive λ method tries to keep safety levels from simulation while allowing updates for real use. Yet it may weaken adherence to physical laws if field changes demand breaking simplified assumptions - a conflict needing close oversight. Other options involve keeping physics rules fixed but adjusting just certain network parts - later stages adapt to surroundings, earlier ones retain rule-based actions - or using residual techniques where a core physics-guided policy gets corrections shaped by recurring prediction mistakes.

The proposed 1,000-10,000 kilometre fine-tuning budget appears optimistic given contemporary autonomous vehicle testing practices. Waymo reports accumulating 20+ million miles before deployment consideration; Cruise exceeded 1 million miles supervised testing. The reviewed paper's curriculum may underestimate sim-to-real gap closure difficulty, though the focus on physics-informed training specifically aims to reduce required real-world mileage compared to purely data-driven approaches. Empirical validation must assess whether physics-regularised training delivers hypothesised data efficiency improvements

Computational Architecture: Hardware needs include a fast multi-core processor - like Intel i9 or AMD Ryzen 9 - paired with an RTX

4000-level graphics card, at least 32GB of memory, plus a solid-state drive larger than 500GB; these specs demand high processing power far beyond standard car computers. While automotive platforms like NVIDIA Drive AGX Orin use eight ARM cores, built-in graphics, and shared memory between 32GB and 64GB, they remain more limited in raw performance. Still, progress in on-board computing now allows vehicles to handle similar hardware within ruggedized units designed for extreme conditions. Such systems must endure cold and heat ranging from -40°C up to 85°C, resist physical shaking tested under MIL-STD-810G standards, along with maintaining function amid electrical noise defined by ISO 11452 guidelines.

The suggested improvements tackle deployment computing limits in various ways - by adjusting resource use, streamlining processes, while enhancing efficiency; each approach targets a specific bottleneck without increasing complexity

Once trained using exact physics simulations, constraints are transferred into a compact neural model - enabling quicker predictions with minimal loss in precision. Instead of direct computation, a simple network (3–5 layers, each with 128–256 nodes) learns from simulation data pairs to deliver results under 1 ms on GPU hardware while maintaining over 95% accuracy. Still, estimation inaccuracies may arise, especially in rare situations where

precise physical behavior is critical; therefore, thorough testing is needed because reliability cannot be taken for granted.

Selective Constraint Evaluation: Use complete physics analysis just in urgent cases - like sudden stops or sharp turns; rely on simpler models otherwise. This cuts typical computational load to 10–15%, yet keeps precision when it matters most. But spotting those key moments accurately depends on rules or trained predictors, which might misclassify events due to ambiguous signals.

Hardware acceleration uses dedicated devices to run physics calculations faster - GPUs handle many constraints at once, especially useful when testing various action-state combos in learning phases. Instead of general chips, FPGAs give predictable response times, which helps systems that must react quickly and reliably, while using lower-precision math (like 16-bit integers) to save energy. For mass-produced vehicles, custom silicon such as ASICs can boost efficiency even further. On the downside, building tailored hardware demands heavy upfront effort: half a million to two million dollars spent over six to eighteen months - a cost only worth it if millions of units will ship.

Computational Architecture: Hardware needs include an 8+ core processor - like Intel i9 or AMD Ryzen 9 - paired with a high-end GPU such as the RTX 4000 series, alongside at

least 32GB of RAM and a 500GB SSD or better. These specs demand significant computing power, far beyond what's common in standard vehicle systems; for instance, NVIDIA Drive AGX Orin uses an 8-core ARM chip, built-in graphics, and shares between 32–64GB memory. Still, progress in edge devices now allows similar performance within ruggedized car-ready units. Such upgrades operate reliably under harsh conditions: extreme temps from -40°C up to 85°C, strong vibrations tested per MIL-STD-810G, along with noise resistance aligned to ISO 11452 standards.

The suggested improvements tackle computing limits during setup in various ways:

Neural Physics Approximation: Once trained using physics-based models, constraints are compressed into neural networks that infer quicker but with minor precision loss. Instead of complex calculations, a compact feedforward setup - just 3 to 5 layers, each holding 128–256 nodes - learns from input-output data generated by traditional simulations. This allows under-one-millisecond prediction times on GPUs while matching physical behavior more than 95% accurately. Still, estimation mistakes may occur; particularly in rare situations where exact dynamics are critical, such inaccuracies could weaken safety buffers - a compromise needing thorough testing instead of blind trust.

Selective Constraint Evaluation: Use complete physics analysis just in tough moments - like quick speed changes, edge-case moves, or emergencies; swap to simpler models when driving normally. That cuts overall load to 10–15%, but keeps accuracy high when safety matters most. Still, spotting those key moments well depends on rules or trained systems, which might misjudge and cause errors.

Hardware Acceleration: Run physics calculations on dedicated chips - use GPUs to handle many constraints at once, FPGAs for consistent timing in control systems, ASICs when building car tech. Today's GPUs process large groups of state-action sets quickly by spreading work across cores. Instead of general processors, FPGAs give predictable response times, which helps in time-sensitive applications while allowing tailored number formats like 16-bit integers that save energy. On the downside, designing special-purpose hardware demands heavy upfront effort - one to one-and-a-half years plus half a million to two million dollars - so it only makes sense if millions of units will ship.

Hierarchical Processing: Split quick reactive actions - above 100 Hz, few physical limits, aimed at stability and instant obstacle response - from slow thinking tasks around 10 Hz that check full dynamics, shape paths, and predict movements seconds ahead. Such a setup matches current autonomous vehicle designs

using rapid internal cycles for throttle or steering alongside delayed external stages handling route choices and driving behavior. Physical rules matter most during planning; reactive systems follow pre-approved motions instead of recalculating feasibility every step

V. CRITICAL ANALYSIS AND LIMITATIONS: INTELLECTUAL HONESTY ASSESSMENT

Computational Overhead: Empirical Quantification and Mitigation Analysis

The study openly admits a 30–50% rise in computation cost - this level of clarity stands out when contrasted with overly positive tech summaries. Early tests show around one-third drop in training speed: 1,500 versus 1,000 environment steps per second; at the same time, response delays grow by 2 to 5 milliseconds. These figures offer clear insight into performance under real conditions. Such data-based support sets this research apart from conceptual models that skip practical computing limits.

Examining where overhead comes from shows possible improvements in how things are set up.

Constraint Evaluation: Computing physics errors - like friction, motion, energy, or movement - needs extra runs through the dynamic formulas. Instead of forces, accelerations are assessed for every restriction; meanwhile, checks for friction limits occur

alongside energy gradient calculations, taking 0.5–1.5ms per condition and action set. When applying four restrictions, delays add up - that's how we reach 2–6ms, aligning with measured processing lag.

Gradient Computation: Backpropagation using physical rules expands computation graphs compared to basic policy gradients. While automatic differentiation keeps intermediate values for derivatives, this raises memory use by 50–100%, along with a 20–30% rise in processing time. Training speed suffers more than inference delay; hence, training drops by 33%, whereas inference slows just 2–5ms

Constraint Weighting and Aggregation: Computing weighted sum of constraint violations and combining with standard RL loss introduces minimal overhead (<0.1ms) given modern GPU parallelisation capabilities. However, adaptive weighting schemes that adjust λ values based on constraint violation histories could introduce iteration-dependent branching, degrading GPU efficiency.

However, the overhead assessment focuses on training throughput and inference latency without addressing memory footprint increases from maintaining physics constraint computation graphs during backpropagation. Physics-informed networks typically require storing intermediate activations for force, acceleration, and energy calculations,

potentially doubling memory requirements compared to standard RL (from ~8GB to ~16GB for typicalPPO implementations). For memory-constrained embedded systems, this limitation may prove more restrictive than computational latency, yet receives minimal attention in the reviewed work.

The proposed mitigation strategies exhibit varying practical viability. Neural physics approximation achieves <1ms inference through GPU-accelerated neural evaluation but introduces approximation errors potentially degrading safety margins, a trade-off requiring rigorous validation rather than assumed acceptability. Selective evaluation reducing average overhead to 10-15% demonstrates efficiency gains but demands reliable criticality heuristics potentially introducing classification errors. The discussion of training time implications (1.5×wall-clock duration) understates practical impact for large-scale development programmes where extended training directly translates to increased computational infrastructure costs, development timelines, and carbon emissions.

Model Fidelity Versus Tractability: Fundamental Tension the conflict between complex physics models and practical optimization is likely the core issue here. It appears in various ways - each demanding structured solutions instead of random fixes.

Dimensionality Curse: Simple car models - like bicycle-style ones with 4 states and 2 inputs - allow smaller-scale optimization that converges reliably. Models of medium detail - such as single-track systems with weight shift (8 states, 3 inputs) - include more effects while still being manageable. Highly detailed versions - with multi-part bodies, suspension, over 20 states, 10+ controls, like separate wheel forces - mirror actual vehicle motion closely; however, they greatly increase variables, leading to scalability issues where required data rises fast as dimensions grow

Gradient Pathology: Stiff equations from vehicle dynamics - with fast tire effects (~1 ms), medium-body motions (~100 ms), and slow heat changes (~10 s) - lead to unstable Hessians in optimization. When condition numbers go beyond 10^6 , gradient steps jump across physical limits, either crawling toward convergence or failing completely. Methods like adjusted learning speeds, limiting gradient size, or using curvature info (e.g., L-BFGS, natural gradients) help somewhat; however, they add settings that need manual setup.

Parameter Uncertainty: Complex models need many inputs - like tire stiffness, damping values, or air resistance - that are often uncertain. A high-detail tire description might use 10 to 15 variables per wheel, covering grip traits, response delays, and load effects; each carries around 10–20% error margin. Handling such variability under physical laws demands

statistical methods or multiple simulations, which sharply increase computing effort.

The study suggests beginning with basic models, then gradually increasing complexity if differences between simulation and reality appear during testing - yet offers little concrete advice. One method worth exploring is sensitivity analysis, which adjusts physical parameters slightly while observing effects on policy outcomes. Another option involves using information criteria like AIC or BIC to choose models that balance accuracy with simplicity. Additionally, cross-validation across different simulator levels can help: train on simpler physics, test on more detailed versions, identifying setups that maintain performance across settings.

Unaddressed Challenges: Honest Scope Acknowledgement

The study clearly lists unresolved issues - like misperceptions, hostile settings, advanced planning, or team-based control - which shows a realistic view of the model's boundaries. By setting these limits, it avoids broad claims while highlighting particular advances. Still, some listed problems need more analysis, especially how they connect with physics-guided training.

Perception Failures: The idea that flawed inputs lead to flawed outputs highlights how physical laws affect decisions after perception. While physics-based rules rely on estimated

states, wrong estimates cause bad choices - even when those rules are followed. Yet, physical limits might help check perception accuracy: a reported speed conflicting with prior acceleration could signal faulty sensors or tracking mistakes. This two-way link - where physics shapes state use while also testing perceptual reliability - needs proper design attention instead of brief references

Take a real example: a car's location (x,y) and speed (v_x , v_y) come from LiDAR data. Instead of assuming it's correct, we check if that speed matches recent movement using basic motion rules. When there's a big mismatch - say, 5 m/s claimed but only 2 meters moved in one second, meaning about 2 m/s - the system may have misread surroundings. Such gaps can prompt tighter blending of sensor inputs or warnings to human operators. Yet, adding this kind of reality check needs structural changes missing in current design proposals.

Adversarial Scenarios: The discovered weaknesses - oil patches lowering grip below normal levels, objects with odd traits like unusually high or low weight or strange flexibility, fake sensor inputs leading to incorrect system readings - are real risks in settings where attackers are present. Methods from robust machine learning such as training with adversarial examples, verified protections, and checking input data help somewhat; however, they have clear limits:

protection from unseen attacks is still not feasible, so layered security using several different strategies becomes necessary.

Physics-informed learning can have new weaknesses: when attackers know the built-in physical rules, they might design situations where real-world conditions differ from those rules. Because of this mismatch, problems can arise - like assuming friction values stay between 0.2 and 0.9, while actual surfaces slip below 0.1 using ice or special coatings. Since these cases break model expectations, control strategies relying on physics may fail more than simpler methods ignoring friction entirely. To find such flaws, stress tests by adversarial teams are needed; yet current studies don't mention them.

High-Level Planning: Splitting physics-based low-level actions from broader decisions - like path choice or lane use - matches how systems are often built. Still, the link between these levels needs clear rules; plans should fit what controllers can actually do physically, whereas controls must align with planner expectations. Methods that optimise both planning and control together, while keeping physical limits in mind across layers, show potential yet remain underexplored.

For example, top-tier planners might choose bold lane changes if they assume strong sideways acceleration. When physical limits restrict actual performance - like low grip,

heavy loads, or soft suspension - the plan fails during execution. Fixing mismatches like these means either planning more cautiously based on worst-case scenarios - which slows things down - or using repeated cycles between planner and controller to adjust plans, which demands high computing power and complicates live operation. The examined structure lacks clear design principles for aligning different decision layers.

Multi-Agent Coordination: Modern self-driving systems often use groups of vehicles working together - like forming platoons, merging smoothly, handling intersections, or improving traffic flow. Applying physics-based learning to these multi-agent setups brings key issues: first, more agents mean higher computation costs since each checks physical rules; second, teamwork strategies need to honor personal vehicle limits while reaching shared goals; third, tight communication channels restrict how much data can be exchanged; fourth, when visibility is low, agents must infer others' physical capabilities using sparse data.

An interesting idea needs testing: when agents use shared physical knowledge, they might coordinate better by acting in ways that naturally suggest intent. Because one car's brake pattern reveals its grip limits, others can anticipate moves without direct signals. Such physics-based interaction may cut reliance on heavy data exchange while making actions

easier to foresee - unlike opaque AI strategies that act unpredictably. Still, this concept lacks proof and demands thorough tests across multiple interacting systems.

VI. EXPERIMENTAL DESIGN AND VALIDATION PROTOCOLS: METHODOLOGICAL CONTRIBUTIONS

Benchmark Scenario Specifications: Concrete Testability

The suggested test procedures offer a clear method advance, giving specific details that allow consistent experiments. While the first scenario checks emergency stop performance across settings, the second evaluates how well merging works in varied highway conditions. Although the third case studies which constraints matter most, the fourth measures time and resource costs during operation. Because each setup targets a different aspect - adaptability, broad applicability, module impact, or speed - it supports thorough system validation without redundancy.

Emergency Braking Benchmark: This case tackles the main sim-to-real challenge via structured testing. During training, Agent A uses standard PPO combined with domain variation, while Agent B applies physics-constrained PPO - both learn in CARLA over half a million steps using matching setups: policy nets have three layers of 256 nodes, value nets use three layers of 128, learning rate is set at $3e-4$ with cosine decay, and batches

contain 2,048 samples per update. For Agent A, domain shifts include changes in lighting, surface looks, weather conditions along with small physical tweaks like mass adjustments up to $\pm 10\%$ and friction altered by $\pm 15\%$. In contrast, Agent B keeps physics stable during training, applying penalty weights $\lambda_{\text{friction}}=1.0$ and $\lambda_{\text{dynamics}}=0.8$ to limit unrealistic dynamics

During testing, two agents operate within IPG CarMaker or rFactor Pro - simulators offering stronger physics accuracy than CARLA, especially in tire behavior, suspension response, or brake representation. One thousand test runs include structured variations: starting velocities (30, 50, 70, 90, 110, 120 km/h), pavement types (dry asphalt $\mu=0.9$, wet asphalt $\mu=0.6$, ice-covered $\mu=0.2$), incline angles (-5° , 0° , $+5^\circ$ influencing load shift along with braking capacity), and car weights (1,200 kg small model, 1,500 kg midsize, 2,000 kg SUV, 2,500 kg heavy-duty truck)

Evaluation measures show how well knowledge transfers across tasks. When assessing Stop Distance Error, researchers use $|d_{\text{actual}} - d_{\text{optimal}}| / d_{\text{optimal}}$; here, d_{optimal} is calculated as $v^2 / (2\mu g \cdot \cos(\theta))$ plus $v \cdot t_{\text{reaction}}$ - this reflects the shortest physically possible stop. The score reveals whether policies grasp key elements like surface grip or brake response time. Values under 20% imply solid adaptation. In contrast,

results above 40% point to weak comprehension of physical rules. Stability breaches show how often wheels slip past grip limits - meaning they lock - or when sideways movement crosses road lines, suggesting poor steering response; also included is excessive rotational speed that hints at spinning danger. Comfort levels are judged by peak jerk values $|da/dt|_{\max}$ reflecting motion fluidity, where results under 3 m/s³ feel smooth while those above 5 m/s³ suggest sudden shifts.

Success criteria establish quantitative bars: Physics-regularised Agent B must achieve <20% stopping distance error with <5% stability violations, whilst standard Agent A expected to exhibit >40% distance error with >15% stability violations. Statistical analysis employs two-sample t-tests comparing means across 1,000 episodes, with Bonferroni correction for multiple comparisons (three primary metrics: $\alpha_{\text{corrected}} = 0.05/3 \approx 0.017$) controlling family-wise error rate. Effect sizes (Cohen's d) supplement significance tests, characterising practical importance beyond statistical significance.

Limitations and Refinements: The benchmark assumes CarMaker/rFactor Pro provides ground-truth physics, yet they too employ approximations. Ideally, validation would include limited real-vehicle testing (50-100 trials on proving grounds), though the paper's theoretical scope reasonably defers this. The 500,000-step training budget may

prove insufficient for convergence given physics constraint overhead; preliminary results may reveal need for 1-2million steps. The binary Agent A vs Agent B comparison could be strengthened by including additional baselines: Agent C (safe RL using CBFs), Agent D (model-based RL with learned dynamics), Agent E (hybrid physics-informed + CBF).

Highway Merge Generalisation: Cross-Vehicle Transfer

The highway merge test checks if physics-guided strategies use built-in mechanics insights to adapt across different car motions without new training. This tackles a key real-world issue: does grasping core physical principles - like grip, mass resistance, turning layout - allow instant adaptation when car specs shift, or do learned strategies become tied to specific vehicles, needing fresh tuning?

Experimental Design: Training uses a common sedan (weight: 1,500 kg, wheel distance: 2.7 m, width: 1.5 m, height: 1.4 m, air resistance factor: 0.32). In highway merge cases, the main car starts on a ramp at 60 km/h while joining traffic moving between 100–120 km/h, with speeds spread normally and variation of ± 5 km/h. Vehicle spacing changes per run - low (0.1 cars/meter, around 3 seconds apart), medium (0.05 cars/m, about 1.5 s gaps), heavy (0.03 cars/m, near 1-second intervals demanding bold entry maneuvers). Learning

proceeds until performance stabilizes, roughly after 800,000 steps observed through flat trends in improvement graphs.

During testing, the trained policy runs - no adjustments made - on three vehicles with clearly different physical traits. First, a compact car: 1,200 kg (down 20%), shorter frame at 2.5 m (−7%), narrower width (1.4 m, −7%), lower height (1.3 m, −7%), drag coefficient of 0.30 (−6%) - this setup allows quicker turns due to light weight yet struggles in balance because of low mass and slim stance. Next, an SUV: heavier at 2,200 kg (+47%), longer by 3.0 m (+11%), wider track (1.7 m, +13%), taller build (1.8 m, +29%), air resistance rises to 0.38 (+19%); its high body raises roll danger while added bulk slows pickup and stop - but grip improves under load. Lastly, a pickup truck hits 2,500 kg (+67%), stretches over 3.5 m wheelbase (+30%), spans 1.8 m wide (+20%), stands tall at 1.9 m (+36%), faces more wind with Cd 0.42 (+31%); extended length limits cornering ease whereas unladen rear heaviness shifts control behavior compared to front-heavy sedans.

Each vehicle type undergoes 500 merge attempts across traffic density variations (167 episodes per density level: sparse, moderate, dense). Evaluation metrics include Success Rate measuring percentage of collision-free merges completing within 150m without forcing highway traffic to perform emergency

maneuvers (decelerations $>4 \text{ m/s}^2$ or lateral accelerations $>3 \text{ m/s}^2$). Comfort assesses average lateral and longitudinal jerk during merge maneuver, calculated as RMS jerk: $\sqrt{\text{mean}[(da_{\text{lat}}/dt)^2 + (da_{\text{long}}/dt)^2]}$ with target $<2 \text{ m/s}^3$ for comfortable merging. Physics Violations measure percentage of control cycles where actions would exceed vehicle-specific friction limits (computed using actual vehicle mass, track width, centre-of-gravity height), kinematic constraints (vehicle-specific wheelbase in Ackermann geometry), or stability thresholds (rollover risk based on actual height/track ratio). Efficiency captures time to complete merge maneuver from on-ramp entry to stable highway lane position. Success criteria specify physics-regularised policy should achieve $>80\%$ success rate across all vehicle types (compact, SUV, pickup) with $<5\%$ physics violations per vehicle type. Standard policy trained only on sedan expected to achieve $<60\%$ success on alternative vehicles with $>15\%$ physics violations as learned behaviors fail to generalise to different dynamics. Statistical analysis employs chi-square tests comparing success rate distributions across vehicle types and policies, with ANOVA testing factors [policy type: standard vs physics-informed] \times [vehicle type: compact/SUV/pickup] \times [traffic density: sparse/moderate/dense] for main effects and interactions.

Constraint Importance Ablation: **Component Contribution Analysis**

The ablation study examines each physics constraint separately, tackling a core issue: which laws boost real-world performance most? By doing so, it distinguishes key rules from optional ones - guiding leaner designs focused on faster computation where resources are limited.

Experimental design tests six versions of the policy, each using different active physics rules; otherwise, they share the same setup - network structure stays fixed at three layers with 256 neurons each. The learning method is consistent across all: PPO algorithm, unchanged settings. Training conditions are also matched precisely, including the range of environments used. The variants are: (1) Baseline with standard PPO and no physics constraints, (2) Friction-Only with $\lambda_{\text{friction}}=1.0$, (3) Dynamics-Only with $\lambda_{\text{dynamics}}=1.0$, (4) Energy-Only with $\lambda_{\text{energy}}=0.5$, (5) Kinematic-Only with $\lambda_{\text{kinematic}}=0.8$, and (6) Full-Physics with all λ values at 0.8.

The Training Phase uses the CARLA simulator with fixed test situations for every version: emergency stops (200K steps), fast turns (200K steps), avoidance actions (200K steps), changing lanes (200K steps), adding up to 800K steps per version \times 6 versions = 4.8 million total training steps (~32 GPU-hours

including 30% extra load). In contrast, the Testing Phase checks all six versions inside a detailed simulator (CarMaker) using identical situation types, tracking how well skills transfer (success rates, frequency of rule breaks, safety scores tied to each case), learning speed (how quickly results improve during training, time to reach 90% max performance, peak outcome reached), along with computing demands (steps processed per second while training, response delay in milliseconds during use, memory used in GB per model allowing direct comparison of trade-offs).

Analysis methods use several comparisons - like Tukey HSD or Holm-Bonferroni - to spot which versions beat the baseline without raising false positive rates across many tests. Instead of simple addition, interaction checks ask if limits help more in certain tasks - for example, friction matters most when turning, while dynamic bounds affect stopping stronger. Rather than isolated gains, synergy evaluation assesses whether full-model results go beyond combined single improvements, suggesting constraints work better together

Expected outcomes suggest friction and dynamic limits bring main advantages, whereas energy or movement rules add smaller gains. Still, real-world data might show unexpected patterns: keeping energy constant, though it seems secondary, could block impossible chain reactions when immediate

restrictions fall short. Movement boundaries, even if rooted in shape and space logic, might become irrelevant if dynamic conditions already ensure realistic paths. This setup risks confusion due to task-specific influences - since rule impact probably shifts across situations, deeper insight demands separating results by task and examining how tasks interact with each constraint type

VII. DISCUSSION, FUTURE DIRECTIONS, AND CRITICAL RECOMMENDATIONS

Integration with Contemporary AV Research: Synergies and Tensions

The paper presents physics-informed learning as a supplement - rather than a substitute - for current AV tech like foundation models, perception tools, high-quality simulations, and traditional control techniques. By framing them together, it recognizes that autonomous driving's challenges call for varied methods instead of one-size-fits-all fixes. Still, the talk on merging these stays mostly abstract, without clear design details for real-world system integration.

Using foundation models shows potential when combining new vision-language systems with self-driving specific models that understand scenes well and predict paths effectively - yet they often ignore physical laws. One option is a layered setup: higher-level decisions like route choices or lane changes follow traffic rules and road

conditions using foundation models, while lower-level actions such as steering or braking rely on reinforcement learning guided by vehicle physics. These parts connect when the high-level model suggests movement goals, which the low-level controller carries out under real-world limits. To check safety, one approach uses the foundation model to suggest moves, then applies physics-based checks - if something isn't doable, it sends back limitations and asks for revised plans. Mixed strategies might merge abstract insights from large models (like interpreting complex environments) with physically valid responses that obey built-in restrictions, demanding new designs allowing training signals to pass across both meaning-focused and physics-aware modules.

Perception-Physics Coupling implies current systems that turn sensor data into top-down views might work better with two-way interaction. In one path, perception results move forward to guide control laws shaped by physical rules, while uncertainty estimates widen safety buffers. The reverse link introduces a new function: physical limits check whether perceived motion makes sense - like spotting mismatches between claimed acceleration and actual movement, or speed shifts not aligned with measured tire grip. Such inconsistencies can reveal faulty sensors or wrong object tracking. To use this idea effectively, system design must clarify which

physics rules are used for checking perception compared to guiding actions, define how detection of rule-breaking prompts updates in interpretation or weighting of sensor inputs, allocate enough processing power for instant checks, and decide responses when repeated violations suggest either broken hardware or rare yet valid events.

Classical Control Integration recognizes that Model Predictive Control, PID controllers, and trajectory tracking methods provide clear stability assurances along with established certification routes. Instead of treating physics-informed RL as main control backed by MPC checks, different setups deserve attention. In an MPC-First Setup, MPC acts as core controller - using trained models grounded in physical laws - and RL tunes these dynamic models to boost MPC results instead of forming direct actions, making approval easier since MPC delivers proven safety bounds while learning sharpens predictions. A Hybrid Switch Design might assign traditional controllers to routine tasks like lane centering or speed regulation, triggering learned strategies only when situations grow difficult - such as heavy traffic or unclear road layouts - with physics rules defining safe transition zones and seamless shifts between systems. An Error-Correction Approach may rely on classic controllers for basic operation, applying learned adjustments to fix recurring inaccuracies or adapt locally, under the

condition that both output types together obey fundamental physical limits, preserving full-system integrity.

Promising Research Directions: Concrete Future Work

Multidisciplinary work on learning systems guided by physical laws now includes setups with multiple interacting agents - this shift brings both difficulties and potential that need focused study. A central issue is scaling: does added computation grow steadily as more agents join, or rise faster when they must jointly satisfy physical rules? Instead of assuming direct links between parts, researchers look at how common knowledge of real-world limits like friction may let agents anticipate each other's moves silently, cutting down data exchange needs while making actions easier to foresee. Studies into spontaneous patterns examine if groups using physics-based reasoning form distinct cooperation styles versus those ignoring such principles; early signs suggest these informed teams act bolder - for instance, driving closer together or merging quicker because outcomes are less uncertain. Another focus lies in combining smart vehicles aware of dynamics with regular ones lacking this insight, reflecting actual roads during gradual automation rollout.

Adaptive Constraint Weighting Using Meta-Learning tackles today's models that need

hand-picked λ settings - a slow method restricting usability. Instead of fixed choices, gradient-driven meta-learning may turn λ into adjustable factors refined through nested optimization: one level updates policies using present λ s; another modifies those λ s to boost test outcomes or generalization measures - yet this brings higher computation needs, possible fluctuations, and reliance on sensitive tuning parameters. Alternatively, reinforcement-style meta-learning can model λ selection as a control task: a top-level strategy adapts penalties depending on ongoing behavior, such as violation levels, improvement rates, and update patterns; changes in λ act as outputs, while gains reflect real-world effectiveness, fast convergence, and cross-task results.

Models aware of uncertainty improve existing methods by adding ways to measure doubt. Instead of fixed values, Bayesian PINNs use probability ranges for physical traits - like friction or weight - and update these ranges when new data arrives; this helps set safer limits based on how unsure the model is. Rather than one model, ensembles use several versions - with varied setups or training bits - to get both an average result and a sense of disagreement; if outputs differ widely, systems react with caution. Another method, conformal prediction, creates bounds that hold error rates low (say, under 5%) without assuming specific data patterns, giving reliable safety intervals. To gather better data, active learning picks

actions that reduce unknowns in physics estimates while still achieving goals, mixing goal-driven choices with probing steps that cut down overall model doubt over time.

VIII. CONCLUSION: ASSESSING THEORETICAL FOUNDATIONS AND PATH FORWARD

Summary of Contributions and Limitations

This broad review explores a theory using physics-guided reinforcement learning to reduce the lasting disconnect between simulated training and real-world use in self-driving vehicles. The study under discussion offers key advances: it builds a precise math structure that weaves Newton's laws, grip limits, energy rules, and motion boundaries into RL goals via smooth loss terms - instead of separate modules; it merges Physics-Informed Neural Networks, first made for science simulations, with policy gradient techniques to improve step-by-step decisions; detailed breakdowns show how gradients pass through physical equations, helping networks learn actions that obey real-world mechanics; a layered safety design links trained behaviors, model-based checks, and fast override responses within a tiered system that weighs independence against protection; its training unfolds in three phases - starting in rich sim environments, applying domain variation, then adjusting on actual data while keeping physics rules active at every stage; efficiency methods are introduced to ease computation load for

practical onboard use; clear test setups include measurable benchmarks allowing repeatable testing across four standard tasks; the authors openly discuss drawbacks like processing cost, trade-offs in accuracy, and unresolved issues; finally, they map related work carefully, highlighting missing areas and placing their findings among current efforts in driverless systems, physics-aware machine learning, and secure reinforcement learning.

The framework stands out from standard upbeat tech papers by openly admitting flaws - no real-world testing was done, computing demands are high (training drops 30–50%), core accuracy trade-offs remain unresolved. It tackles basic vehicle control but skips advanced planning, sensor fusion, and group coordination. Verification is tough due to physical limits; certifying RL systems stays a major hurdle, with approval processes expected to take 5–10 years after the tech matures.

Final Assessment: Promise and Required Validation

The study offers a clear theory tackling real self-driving car issues by using solid physics principles. Its math models show skill, the design reflects practical needs, tests suggest measurable ways to assess performance - yet key questions remain unanswered. For instance, whether adding physical rules into reinforcement learning goals can reliably

reduce mismatches between simulation and reality while improving safety and clarity still demands thorough testing. Despite careful structure and honest discussion of limits, this core issue awaits deeper experimental validation.

The suggested test setups offer clear methods to address the research query using repeatable evaluations. While one evaluates if incorporating physical rules enhances performance when moving from basic to detailed simulations, another examines whether model-based control strategies use structural insights for instant adaptation across different vehicles. A breakdown experiment highlights which physical constraints have the greatest impact on transfer success. At the same time, an efficiency review outlines real-world usability differences among various design choices.

If later studies back up theory, using physics in reinforcement learning might help create safer and more trustworthy self-driving cars. Advantages could involve less need for physical tests by improving simulated training transfer - greater safety thanks to rules based on physical laws - easier approval routes through transparent system limits - or better performance across different machines and situations. By directly building in physical principles, the method may simplify validation for agencies who prefer defined boundaries instead of guessing how black-box networks

operate - possibly speeding adoption without lowering security benchmarks.

If test outcomes fall short, the framework's math models, design choices, testing methods, and clear discussion of weaknesses still help later studies by showing which paths don't work - along with reasons. When negative findings reveal core barriers - like computation demands too high for live use, conflicts in model accuracy blocking consistent results across settings, or physical rules limiting strategy flexibility so much that performance drops unacceptably - the field can shift focus to better alternatives, advancing faster via useful dead ends. Knowing why physics-based strategies break down offers key insights, narrowing future options and avoiding wasted effort on deeply problematic routes.

The most probable situation shows varied results - physics-guided learning works better in certain cases, such as known vehicle dynamics or situations limited by physical rules; however, it falls short elsewhere. For example, in fast-changing, messy settings where simple models fail to reflect real-world complexity, performance drops. It also struggles when speed matters and computation slows response times, especially if purely data-driven methods outperform due to hidden pattern detection. Testing across conditions would map where gains outweigh drawbacks, guiding smarter use based on actual needs. Such a balanced view reflects typical progress

with advanced tech: no single solution dominates every area.

The journey from concept to working self-driving machines takes many years of careful study involving:(1) early simulations tested core ideas via suggested benchmark methods, (2) repeated improvements tackled detected errors while fine-tuning settings, (3) linking with related tools like smart sensing units or base AI models followed, (4) real-life testing expanded gradually - from closed tracks to open streets, (5) dialogue with officials supplied proof for approval claims, whereas (6) trust grew among users by openly sharing what the system can and cannot do. Every phase brought hurdles needing steady work from mixed groups skilled in areas such as trial-based learning, motion physics, regulation rules, code design, plus feedback controls.

Recommendation for Research Community: Prioritise hands-on testing using new benchmark setups - begin with simulated environments, then move toward small-scale field trials. Evaluate performance carefully against current techniques like domain randomisation or safety-aware reinforcement learning, not just basic RL; include comparisons of computing needs and setup difficulty across options. Explore ways to reduce computation load by running ablation tests that measure accuracy losses versus gains in speed. Create clear criteria for choosing

models, helping users match physics detail to task demands and hardware limits. Build stronger ties between research groups and companies to gain access to advanced tools such as realistic simulators, sensor-equipped cars, test sites, and private data collections, while supporting faster exchange of insights.

By examining various situations - different weather, car models, and driving environments - the self-driving car field can determine if physics-based learning works in practice or faces unsolvable challenges. This review offers clear theory plus step-by-step testing methods, helping researchers test these ideas thoroughly instead of drawing quick judgments from narrow results

The gap between simulation and reality won't vanish with just one method. Moving forward means openly acknowledging limits, testing options thoroughly under real-world settings - while weighing their pros and cons. Combining different methods can help - each contributing what it does best, provided researchers stay aware of how much they actually know so far. Learning that uses physics rules might play a useful role in solving this broader challenge - a role that's still uncertain, shaped by solid theory and test designs reviewed here. If this strategy changes, self-driving car progress - or fades into obscurity - will become clear only after data is gathered, studied, and confirmed across

multiple teams working in varied environments.

The scientific world must make a key decision: Test ideas carefully using clear methods while checking results against current approaches - rather than ignoring theory without quick proof. One way demands effort but could lead to real progress via structured study; the other might skip useful paths because they haven't been tested enough. This analysis supports careful testing, noting that answers will take years of shared research before we fully understand how physics-guided learning affects self-driving systems

References

- [1] Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378, 686-707.
- [2] Kiran, B. R., Sobh, I., Talpaert, V., Mannion, P., Al Sallab, A. A., Yogamani, S., & Pérez, P. (2022). Deep reinforcement learning for autonomous driving: A survey. *IEEE Transactions on Intelligent Transportation Systems*, 23(6), 4909-4926.
- [3] Wu, P., Jia, X., Chen, L., Yan, J., Li, H., & Qiao, Y. (2024). GenAD: Generative end-to-end autonomous driving. *arXiv preprint arXiv:2402.11502*.

- [4] Hua, J., Zeng, L., Li, G., & Ju, Z. (2024). Learning for vehicle-to-vehicle cooperative perception under lossy communication. *IEEE Transactions on Intelligent Transportation Systems*, 25(4), 3560-3572.
- [5] Vemula, D., & Patel, D. (2024). Sim-to-real transfer for autonomous vehicle control using soft actor-critic. *International Journal of Intelligent Transportation Systems Research*, 22(1), 145-159.
- [6] Chen, L., Wu, P., Chitta, K., Jaeger, B., Geiger, A., & Li, H. (2023). End-to-end autonomous driving: Challenges and frontiers. *arXiv preprint arXiv:2306.16927*.
- [7] Yang, S., & Zhao, W. (2023). Autonomous vehicle control based on DQN and PPO deep reinforcement learning algorithm. *Applied Sciences*, 13(11), 6692.
- [8] Shao, S., Pei, H., Chen, L., Liu, J., & Niu, W. (2023). Reinforcement learning-based decision-making for autonomous vehicles: Recent advances and prospects. *World Electric Vehicle Journal*, 14(9), 263.
- [9] Sharma, P., Evans, W. T., Tindall, J., Negi, P. S., & Kimber, M. L. (2023). A review of physics-informed machine learning in fluid mechanics. *Energies*, 16(5), 2343.
- [10] Wang, S., Sankaran, S., & Perdikaris, P. (2023). Respecting causality for training physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 421, 116813.
- [11] McClenny, L. D., & Braga-Neto, U. M. (2023). Self-adaptive physics-informed neural networks. *Journal of Computational Physics*, 474, 111722.
- [12] Zhao, W., He, J., Qian, Y., Huang, Y., & Liu, T. (2023). Integrated decision and control: Toward interpretable and efficient driving intelligence. *IEEE Transactions on Cybernetics*, 53(10), 6232-6245.
- [13] Kiran, B. R., Sobh, I., Talpaert, V., Mannion, P., Al Sallab, A. A., Yogamani, S., & Pérez, P. (2022). Deep reinforcement learning for autonomous driving: A survey. *IEEE Transactions on Intelligent Transportation Systems*, 23(6), 4909-4926.
- [14] Cuomo, S., Di Cola, V. S., Giampaolo, F., Rozza, G., Raissi, M., & Piccialli, F. (2022). Scientific machine learning through physics-informed neural networks: Where we are and what's next. *Journal of Scientific Computing*, 92(3), 88.
- [15] Emam, Y., Glotfelter, P., Bobadilla, L., & Notomista, G. (2022). Adaptive model-free safety-critical control with application to collision-free multi-robot coordination. *IEEE Robotics and Automation Letters*, 7(2), 2699-2706.
- [16] Wang, S., Teng, Y., & Perdikaris, P. (2021). Understanding and mitigating gradient flow pathologies in physics-informed neural networks. *SIAM Journal on Scientific Computing*, 43(5), A3055-A3081.
- [17] Cai, S., Mao, Z., Wang, Z., Yin, M., & Karniadakis, G. E. (2021). Physics-informed neural networks (PINNs) for

- fluid mechanics: A review. *Acta Mechanica Sinica*, 37, 1727-1738.
- [18] Taylor, A., Singletary, A., Yue, Y., & Ames, A. (2020). Learning for safety-critical control with control barrier functions. In *Proceedings of the 2nd Conference on Learning for Dynamics and Control* (pp. 708-717). PMLR. [↩](#)
- [19] Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378, 686-707.
- [20] National Transportation Safety Board. (2019). *Collision between vehicle controlled by developmental automated driving system and pedestrian* (Highway Accident Report NTSB/HAR-19/03). Author.
- [21] Ames, A. D., Coogan, S., Egerstedt, M., Notomista, G., Sreenath, K., & Tabuada, P. (2019). Control barrier functions: Theory and applications. In *2019 18th European Control Conference (ECC)* (pp. 3420-3431). IEEE.
- [22] Cheng, R., Orosz, G., Murray, R. M., & Burdick, J. W. (2019). End-to-end safe reinforcement learning through barrier functions for safety-critical continuous control tasks. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 33, No. 01, pp. 3387-3395).
- [23] Haarnoja, T., Zhou, A., Abbeel, P., & Levine, S. (2018). Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *Proceedings of the 35th International Conference on Machine Learning* (pp. 1861-1870). PMLR.
- [24] Chua, K., Calandra, R., McAllister, R., & Levine, S. (2018). Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Advances in Neural Information Processing Systems* (Vol. 31). Curran Associates, Inc.
- [25] Dosovitskiy, A., Ros, G., Codevilla, F., Lopez, A., & Koltun, V. (2017). CARLA: An open urban driving simulator. In *Proceedings of the 1st Annual Conference on Robot Learning* (pp. 1-16). PMLR.
- [26] 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.